Artificial Intelligence
A tour with view from Lyon

Patrice Abry, Pierre Borgnat, Aurélien Garivier, Rémi Gribonval, Mathurin Massias,
Nelly Pustelnik, Julian Tachella

IXXI, LPENSL, LIP, UMPA,
École Normale Supérieure de Lyon

June 2025

Roadmap

SCIENTIFIC STRENGTHS OF LYON – SAINT ÉTIENNE

2nd largest university site in France

Education

200 000 students, among which

- 15 400 engineering students
- 4 700 PhD students

Innovation

- 1st industrial region in France
- 2nd largest economic region in France
- 2nd region for start-up creation
- 451 AI-related companies (+50% since 2020)

Research

168 laboratories

- 53 laboratories on core AI and AI+X

Innovation indicators

15 000
SMEs

50
startups

70%
in industry & services

€ 3,2 billion
raised

34% face the challenge to establish relationships with the public sector

Scientific Strengths of Lyon – Saint Étienne

2nd largest university site in France

Education

200 000 students, among which

- 15 400 engineering students
- 4 700 PhD students

Innovation

- 1st industrial region in France
- 2nd largest economic region in France
- 2nd region for start-up creation
- 451 AI-related companies (+50% since 2020)

Research

168 laboratories

- 53 laboratories on core AI and AI+X

Innovation + Research in AI: keys indicators since 2019

427 projects	**479 M€** total funding	**117** ANR projects	**9** PIA4 PEPR
95 projects/year	**+106 M€** funding/year	**24** EU projects	**185** industrial / local /regional projects

AILYS

COLT 2025

Katrina Ligett
Hebrew Univ.

Matus Telgarsky
Courant Inst.

Francis Bach
INRIA

Q2

Computing Infrastructure

Goal: Smooth Access and Support for AILyS Participants

MESOLyS: Towards a mesocenter

ÉCOLE CENTRALE LYON

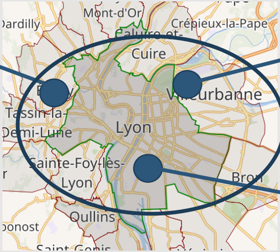

INSA INSTITUT NATIONAL DES SCIENCES APPLIQUÉES LYON

Lyon 1

ENS DE LYON

with engineering support

➡ 240 racks, 300 Million computing hours (2024)

➡ Among the largest mesocenters (i.e., regional data centers) in France

Jean-Zay for large needs

- GENCI IDRIS Jean Zay
- HPE SGI 8600 supercomputer

AILYS CONSORTIUM

Coordinator: ENS DE LYON

4 UNIVERSITIES 6 *GRANDES ÉCOLES*	2 MEDICAL INSTITUTIONS	4 NATIONAL RESEARCH INSTITUTES	4 ENTERPRISE PARTNERS
UCB Lyon 1	HCL Hospices Civils de Lyon	CNRS	ESKER
Université Lumière Lyon 2	Centre Léon Bérard	Inria	onepoint.
Université Lyon III Jean Moulin		Inserm	SOLVAY
Université Jean Monnet Saint-Étienne		IFP Energies nouvelles	VOLVO
INSA Institut National des Sciences Appliquées Lyon			
École Centrale Lyon			
Mines Saint-Étienne			
ENTPE			
ENS de Lyon			
CPE Lyon			

PULSALYS SATT Lyon St Etienne

Q1

AILYS CONSORTIUM

INDUSTRIAL PARTICIPANTS

Fields: health, tech, manufacturing, biotech, agriculture, trade, etc.

Algoé, Ansys, Authôt (You speak. We write.), BIOMÉRIEUX, BLUEWAY, COEXYA, HEAAN CRYPTO LAB

EDF, EXCELLERIA, FEALINX, FOXSTREAM (SMART VIDEO ANALYTICS), framatome, GAOMA Therapeutics, Gencovery

Glaaster., Limagrain (de la terre à la vie), Manitty, meersens, MISTRAL AI_, Optimistik, PHILIPS

PIXYL, PULSALYS (SATT LYON ST ETIENNE), SEB, SIEMENS, Valeo, VIDIUM (SYSTEMS BIOLOGY SOLUTIONS)

Q1

AILYS CONSORTIUM

INDUSTRIAL PARTICIPANTS

TERRITORIAL/POLITICAL STAKEHOLDERS

La Région Auvergne-Rhône-Alpes

MÉTROPOLE GRAND LYON

SÉM SAINT-ÉTIENNE la métropole

CCI LYON

axelera ENSEMBLE, CATALYSONS LES RÉUSSITES

MINALOGIC Auvergne-Rhône-Alpes

LYONBIOPOLE Auvergne - Rhône - Alpes

CARA EUROPEAN CLUSTER FOR MOBILITY SOLUTIONS

POLYMERIS

Tenerrdis Auvergne-Rhône-Alpes

Q1

AILYS CONSORTIUM

INDUSTRIAL PARTICIPANTS

TERRITORIAL/POLITICAL STAKEHOLDERS

PUBLIC/PRIVATE ENTITIES

BIOASTER MICROBIOLOGY TECHNOLOGY INSTITUTE

CRCL CANCER RESEARCH CENTER OF LYON

VetAgro Sup Campus Vétérinaire de Lyon

INRAE

SCIENCES PO LYON

CONSERVATOIRE NATIONAL SUPÉRIEUR MUSIQUE ET DANSE DE LYON

em lyon business school

Ecole supérieure d'art et design

IRIIG INTERNATIONAL RESEARCH INSTITUTE FOR INNOVATION & GROWTH

cnes CENTRE NATIONAL D'ÉTUDES SPATIALES

Q1

AILYS AXES AND CHALLENGES

	CHALLENGE 1 – FRUGAL AI	CHALLENGE 2 – RESPONSIBLE AI
AXIS 1 – CORE AI	frugal LLMs, frugality vs efficiency, information theoretic limits of sample complexity, etc.	explainable AI, fair AI, generative AI for privacy, federated learning, etc.
AXIS 2 – AI & HEALTH	neuroscience, few shot learning, ubiquitous AI, etc.	5P medicine to public health, hospitals' data warehouse, etc.
AXIS 3 – AI & HSS	ecological need for sobriety, frugality as a heuristic, etc.	ethics, transparency, creation, law, accountability, AI as discourse, etc.
AXIS 4 – AI & SCIE. ENG.	frugal infrastructure cooling, energy efficiency and rebound effect, etc.	smart cities, AI for education, etc.

AILYS RESEARCH STRENGTHS AND INTERNATIONAL POSITIONING

Q1

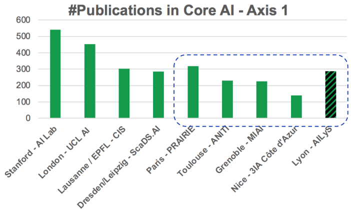

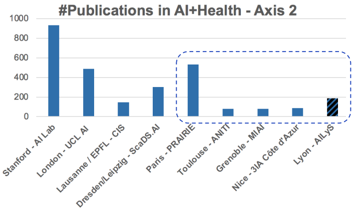

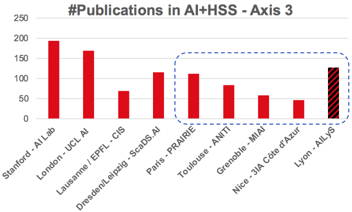

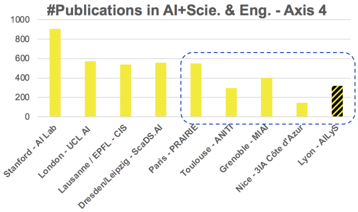

KEY INDICATORS ON AILYS
~550 RECENT PUBLICATIONS

82%

A* / Q1 journals

95%

A* / A / B conferences

Ambition

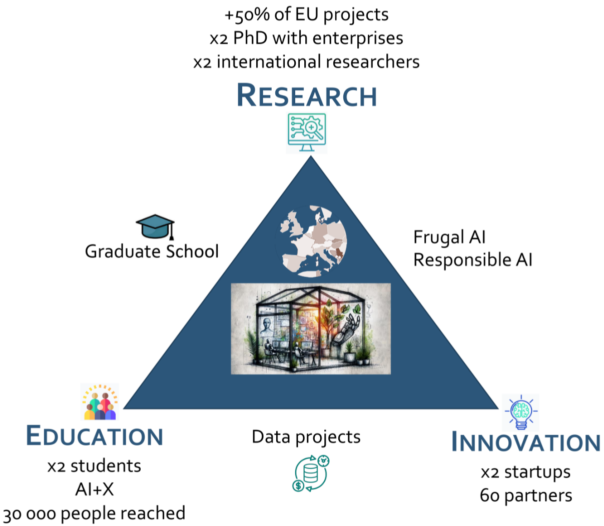

+50% of EU projects
x2 PhD with enterprises
x2 international researchers
RESEARCH
Graduate School
Frugal AI
Responsible AI
EDUCATION
x2 students
AI+X
30 000 people reached
Data projects
INNOVATION
x2 startups
60 partners

Example of Technical Approaches for Challenge 1 on frugal AI

Q8

- More frugal methods are required for the development of AI in the labs and in the industry
- Ecological necessity
- Better think the balance efficiency / sobriety
- More frugal computing infrastructures
- Sketching / dimensionality reduction
- Learning with few examples: informational limits, bio-inspired examples
- Frugality is an esthetic category with a heuristic value

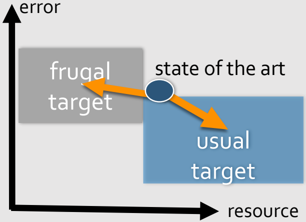

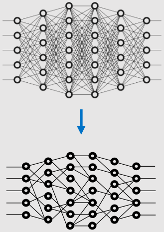

Ockham

AILYS APPROACH TO RESPONSIBLE AI

	CHALLENGE 2 – RESPONSIBLE AI
AXIS 1 – CORE AI	explainable AI, fair AI, generative AI for privacy, federated learning, etc.
AXIS 2 – AI & HEALTH	5P medicine to public health, hospitals' data warehouse, etc.
AXIS 3 – AI & HSS	ethics, transparency, creation, law, accountability, AI as discourse, etc.
AXIS 4 – AI & SCIE. ENG.	smart cities, AI for education, etc.

RESEARCH

- Interdisciplinarity (e.g., AI privacy + law, AI acceptability + psychology, bio-inspired AI + neurosciences, etc.)

INNOVATION

- EU AI Act

EDUCATION

- Responsible AI in education programs
- BSC Centrale /Ecole de Management
  "Data Science for Responsible Business"

DISSEMINATION TO THE LARGE PUBLIC

- AI appropriation and acceptability

AILYS CONSORTIUM IN AI PEPR PROJECTS

Q1

1. Project EMERGENCE: Near-physics emerging models for embedded AI

2. Project ADAPTING: Adaptive architectures for embedded artificial INtelliGence
led by Alberto Bosio (ECL)

3. Project HOLIGRAIL: HOLIistic approaches to GReener model Architectures for Inference and Learning
partner Florent de Dinechin (INSA Lyon)

4. Project REDEEM: Resilient, Decentralized and Privacy-Preserving Machine Learning

5. Project Project SAIF : Safe AI through Formal Methods

6. Project CAUSALI-T-AI: CAUSALIty Teams up with Artificial Intelligence

7. Project FOUNDRY: FOUNDations of Robustness and reliabilitY in machine learning
partner Aurélien Garivier (ENSL)

8. Project SHARP: Sharp Theoretical & Algorithmic Principles for frugal ML
led by Rémi Gribonval (INRIA-ENSL)

9. Project PDE-AI - Numerical analysis, optimal control and optimal transport for AI
partner Julie Digne (LIRIS)

AILYS CONSORTIUM IN OTHER PEPR PROJECTS

Q1

Digital Health:

1.1 M4DI: Méthodes et modèles pour l'intégration de données multimodales et multi-échelles

partner Sara Bouchenak (INSA Lyon)

1.2 AI4scMed: IA multi-échelle pour une médecine de précision en cellules uniques

led by Franck Picard (CNRS, ENSL)

1.3 REWIND: Médecine de précision avec données longitudinales

partner Delphine Maucort-Boulch (HCL, UCBL)

3.2 ChroniCardio: Climatologie de la cardiomyopathie non-ischémique chronique : prédiction à long terme avec des données et modèles multi-échelle

partner HCL

Other PEPRs:

- **ORIGINS** (from planets to life), WP led by Loïc Denis (UJM), partner Nelly Pustelnik (CNRS, ENSL)
- **TARANIS** (Cloud), led by Christian Perez (INRIA LIP)
- **CELCER-EHT** (durable, high-performance, low-cost EHT ceramic cells)
- **DIADEM** (AI for materials)
- **URBHealth** (sustainable city, urban surveillance), partner Lyon 2
- **WAIT4** (welfare: AI for agro-ecology), partner Céline Robardet (INSA Lyon)
- **acoustic sensing for gound and bulding motion in Rhône Vallée** (IRIMA)
- **Cybersecurity** WP led by Lilian Bossuet (UJM)

Roadmap

AXE 1 : CORE AI

CONTINUUM FROM THEORY TO ALGORITHMS


Alice Guionnet

- Understanding Deep Neural Networks generalization
- Optimal transport
- Optimization, generative methods
- Signal processing, graphs
- Data mining, machine learning
- Data management

STRUCTURING EFFECT ON EDUCATION : GRADUATE SCHOOL


Angela Bonifati

- Flagship master
- Masters IA+X
- Lifelong training (including PhD students in physics...)

DIFFERENTIAL PRIVACY

Private Quantiles Estimation in the Presence of Atoms
C. Lalanne, C. Gastaud, N. Grislain, A. Garivier, R. Gribonval
Information and Inference 12(3), Sep. 2023, p. 2197–2223

Private Statistical Estimation of Many Quantiles
C. Lalanne, A. Garivier, R. Gribonval
ICML Jul. 2023

On the Statistical Complexity of Estimation and Testing under Privacy Constraints
C. Lalanne, A. Garivier, R. Gribonval
TMLR Apr. 2023

About the Cost of Global Privacy in Density Estimation
C. Lalanne, A. Garivier, R. Gribonval
TMLR Aug. 2023

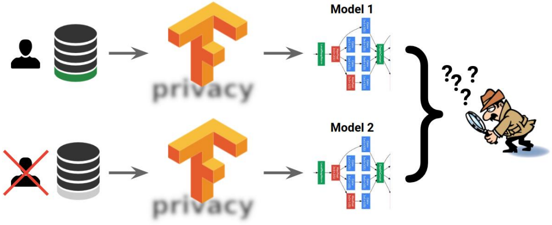

	ϵ-DP Equation (2.1)	**ρ-zCDP** Equation (2.2)
Lipschitz Equation (3.2)	Upper-bound: $O\left(\max\left\{n^{-2/3}, (n\epsilon)^{-1}\right\}\right)$ [BD14] & Proposition 1 Lower-bounds: -Point-wise: $\Omega\left(\max\left\{n^{-2/3}, (n\epsilon)^{-1}\right\}\right)$ Theorem 2 -Integrated: $\Omega\left(\max\left\{n^{-2/3}, (n\epsilon)^{-1}\right\}\right)$ [BD14] & Theorem 3	Upper-bound: $O\left(\max\left\{n^{-2/3}, (n\sqrt{\rho})^{-1}\right\}\right)$ Proposition 1 Lower-bounds: -Point-wise: $\Omega\left(\max\left\{n^{-2/3}, (n\sqrt{\rho})^{-1}\right\}\right)$ Theorem 2 -Integrated: $\Omega\left(\max\left\{n^{-2/3}, (n\sqrt{\rho})^{-4/3}\right\}\right)$ Theorem 3
Periodic Sobolev Smoothness β Equation (4.2)	Upper-bounds: -Pure DP: $O\left(\max\left\{n^{-\frac{2\beta}{2\beta+1}}, (n\epsilon)^{-\frac{2\beta}{\beta+3/2}}\right\}\right)$ Theorem 4 -Relaxed: $\tilde{O}\left(\max\left\{n^{-\frac{2\beta}{2\beta+1}}, (n\epsilon)^{-\frac{2\beta}{\beta+1}}\right\}\right)$ $\tilde{O}$ hides polylog factors. Proposition 6 Lower-bound: $\Omega\left(\max\left\{n^{-\frac{2\beta}{2\beta+1}}, (n\epsilon)^{-\frac{2\beta}{\beta+1}}\right\}\right)$ Theorem 5	Upper-bound: $O\left(\max\left\{n^{-\frac{2\beta}{2\beta+1}}, (n\epsilon)^{-\frac{2\beta}{\beta+1}}\right\}\right)$ Theorem 4 Lower-bound: $\Omega\left(\max\left\{n^{-\frac{2\beta}{2\beta+1}}, (n\epsilon)^{-\frac{2\beta}{\beta+1/2}}\right\}\right)$ Theorem 5

GENERATIVE AI IN AILYS

Q1

- **Deep-Fakes detection** think tank led by Minalogic (with MIAI)
- **Literature / NLP** (PEPR-IA Sharp, INRIA Ockham, ANR DIKé, AURA QABOT)
- **Medical Data** (ANR Orchid&ULYB, PEPR Santé-Chronicardio)
- **Synthetic network traffic generation** (ANR-NSF project MINT in collaboration with University of Chicago, Colgate University, and Hawaii University)

- **Seismology / Geology** (ERC Transcale ended ‘22)
- **Scene representation** (Edge AI Asterix, ERC Panoramix under review)
- **Human-Robot interaction** (ASLAN Pepper Mint)
- **Spatio-temporal processes**
- **Economic games**

AI for imaging inverse problems

Self-supervised learning for inverse problems: noisy/incomplete data, no ground truth
DeepInverse: open-source library (deepinv.github.io/) for solving imaging inverse problems with AI: +1000 monthly users, +40 international contributors, +15 imaging domains, hackathons, SOLL 2024 award.
Foundation models for computational imaging: AI models that can solve multiple imaging tasks, from medical imaging to remote sensing. Startup spin-off (J. Tachella).

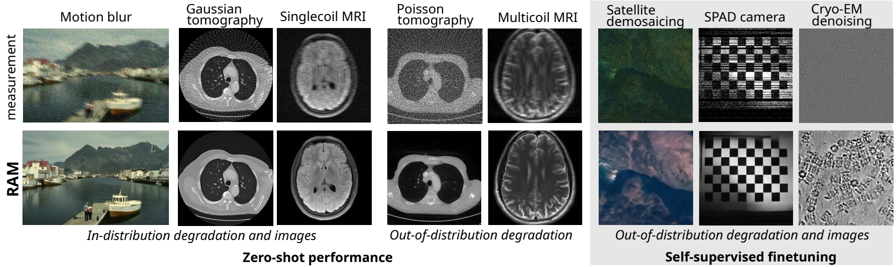

Roadmap

AXE 2 : HEALTH AND LIFE SCIENCES

INTEGRATION HCL/UCBL

LEADING HOURAA GCS

Delphine Maucort-Boulch

- Using AI to better understand complex biological systems
- 5P medicine
- Develop new AI methods for life sciences

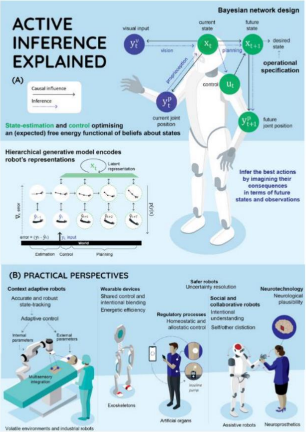

NEUROSCIENCE : MUTUAL INFLUENCES

Olivier Bernard

- 850 community, 10 ERC
- 2 CPJ, multidisciplinary
- bio-inspired neural networks human-like cognitive systems

SHAPE-MED@LYON: STRUCTURING ONE HEALTH APPROACH FOR PERSONNALIZED MEDICINE IN LYON

Q1

KEY FIGURES

- A 10-year project
- **28.1 M€** of funding (France 2030)
- 12 partners
- **5 000 researchers and clinicians**

OBJECTIVES

- Bringing together the **One Health** approach and **5P medicine**
- Build a research and training ecosystem with a **transdisciplinary** vision
- Reposition **the individual as the target and actor** in his or her environment

HCL and Philips Collaboration with Integration to Clinical Systems

Q1

Quantification of Lung Damage by COVID-19

Step 1: Automatic Segmentation of the Lungs 1

Step2: Automatic Segmentation of the Lesions 2

Step 3: Automatic Reporting 3

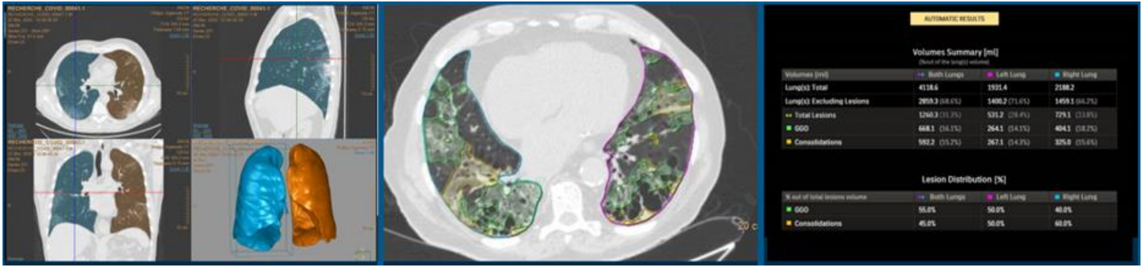

Roadmap

AXE 3 : AI AND HUMAN SCIENCES

AI FOR HUMAN SCIENCES

Gaëlle Marti

- Corpus exploration (museography, music, literature, technology watch), AI changes perceptions
- Need for sobriety (frugality)
- GenAI = pedagogical tool for creation (literature, music, live performance, visual arts, design)

HUMANITIES FOR A BETTER AI

Jean-Claude Dreher

- Neurosciences
- Social impact
- Ethics, Law, Privacy, Accountability
- Deconstructing discourses on AI
- Creation and authorality

Double bachelor **and Master** degrees: law&IT

⇨ CPJ 2024 in Lyon 2

Lyon Leading France-Wide PEPR Projects with HSS and AI&Data Science

Q1

Cultural and Creative Industries (ICCARE)

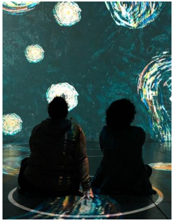

- A 6-year project
- **25 M€** of funding for all partners (France 2030)
- Art creation
- Interdisciplinarity (Art & Culture, AI & data science)

Sustainable City and Innovative Building (VDBI)

- A 8-year project
- **40 M€** of funding for all partners (France 2030)
- Modelling and urban simulation
- Cross-referencing, processing, and analysis of data

Hackathon: How does AI impact the practice of science?

Do Large Language Models –LLMs– (help) produce research ?

How ? At what Costs ?

How deeply is it changing requested research skills ?

How does it impact human-human interactions in research ?

main.pdf
PDF

could you produce a matlab toolbox that impleme,nts the estimation of the vector of selfsimilarity parameters from the bwavelet method proposed in the attached paper ?

Hackathon at a Glance

Challenge: complete a research task in 24hours

Using Large Language Models

5-groups, 30+ partcipants

Organized by IXXI

J.-P. Magué, R. Cazabet, N. Ferret, P. Abry ;

https://www.ixxi.fr/agenda/evenements/hackathon-grands-modele-de-langage

Yes, I can help you produce a MATLAB toolbox that implements the estimation of the vector of selfsimilarity parameters using the **bias-corrected multivariate wavelet eigenvalue method** described in the paper *"Multivariate selfsimilarity: Multiscale eigen-structures for selfsimilarity parameter estimation"*.

Key Tasks for MATLAB Toolbox Implementation

Based on the method in the paper, here's what the toolbox will need to include:

📦 **Toolbox Name:** `MultivariateSelfSimilarityEstimation`

1. Wavelet Decomposition Module

- Compute the Discrete Wavelet Transform (DWT) for multivariate signals.
- Support compactly supported wavelets like Daubechies2 with configurable vanishing moments.

2. Wavelet Spectrum Estimation

- Compute the wavelet spectrum matrix $S_Y(2^j)$ for each scale:

$$S_Y(2^j) = \frac{1}{n_j} \sum_{k=1}^{n_j} D_Y(2^j,k) D_Y(2^j,k)^\top$$

3. Eigenvalue Extraction and Bias-Correction

- Divide wavelet coefficients into fixed-size blocks across scales.
- Average log-eigenvalues across blocks to mitigate repulsion bias:

$$\log_2 \lambda_m(2^j) := \frac{1}{} \sum^{T_j} \log_2 \lambda^{(\tau)}(2^j)$$

Roadmap

AXE 4 : IA FOR NATURAL SCIENCES AND ENGINEERING

- Generative models
- Physics-informed AI
- NLP & technology watch
- High-level engineering schools with high activity in AI (partnerships health, EM-Lyon, etc.)

Marc Sebban

Modelling Lab Greenhouse

Loïs Denis

AIIYS

Extreme Heatwave Forecast with Deep Learning

Q1

Objective

- Forecast the occurrence of **extreme long-lasting heatwaves**
- **U**sing outputs of physics-driven models
- **Alternative to physics-driven models**

Constraints & Approach

- **Rare events**
- **Data imbalance**
- Transfer deep learning

Instantaneous map of the T_s and Z_g anomalies

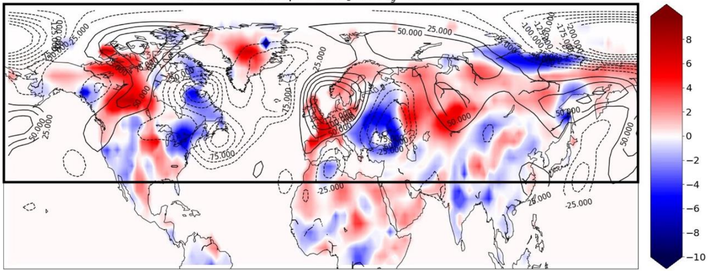

AILYS

UNDERSTANDING THE UNDERLYING PHYSICS OF LASER-MATTER INTERACTION

Processing the surface of matter with a (femto-second) laser allows one to create nanoscale patterns leading to specific properties such as a hydrophobic surface or avoiding the fixation of some bacterias. In this work, we were able to characterize and predict the formation of nanoscale patterns using a stochastic Swift-Hohenberg (SH) PDE model combined with a physics-guided machine learning strategy. Incorporating machine learning and physical knowledge allowed us to learn the complex relationship between irradiation conditions and patterns. This enables us to predict patterns in unexplored regions of the laser parameter space, and even potentially create new ones: in a sense, to teach matter how to move. Being able to do so drastically reduces the need for extensive experimental efforts.

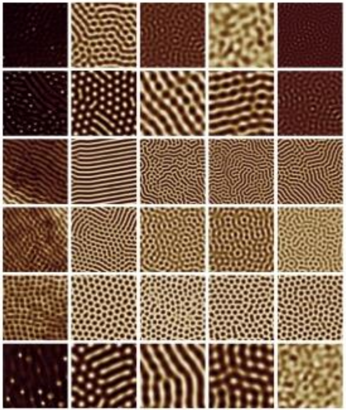

Joint work between AI specialists and physicists from the Hubert Curien laboratory, Saint-Etienne.

Reference:

E Brandao, A Nakhoul, S Duffner, R Emonet, F Garrelie, A Habrard, F. Jacquenet, F. Pigeon, Marc Sebban, J.-P. Colombier. Learning Complexity to Guide Light-Induced Self-Organized Nanopatterns Physical Review Letters 130 (22), 226201, 2023.

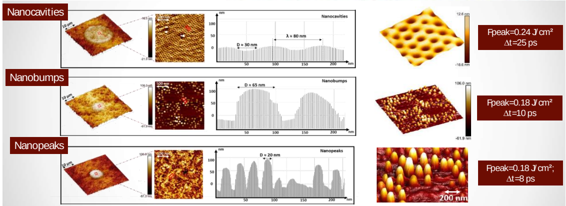

Roadmap

Context : network compression

Large network

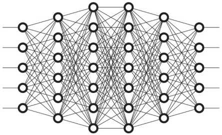

- includes **dense** layers
- **high-precision** float format

Compressed network

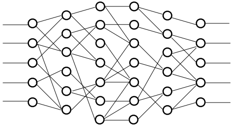

- **sparsely** connected
- **coarsely** quantized

Can we leverage usual know-how ?

From deep sparsity to butterflies

Classical sparsity"

- few but *irregular* connexions

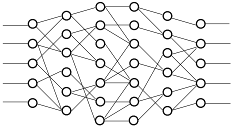

- **several challenges**
  - ➡ NP-hard / instable optimisation
  - ➡ *sparse format inefficient on GPUs*

Butterfly sparsity

- *structured* sparsity patterns
  - ex: DFT

$\mathbf{F} =$ 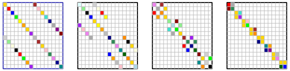

  - **scénario:** replace dense layers par (several, structured) sparse layers

see e.g. [T. Dao & al, Monarch: Expressive Structured Matrices for Efficient and Accurate Training, ICML, 2022]

- **recent advances**
  - flexible butterfly format
  - near-optimal approximation algorithm
  - efficient GPU implementation

Efficiency and guarantees

Behaviour of butterfly algorithm

Theoretical guarantees (for chainable architecture β)

➡ Existence of optimum of (1)

$$E(\mathbf{A}) = \inf \|\mathbf{A} - \mathbf{X}_1 \ldots \mathbf{X}_L\|_F \quad (1)$$

under Kronecker-sparse support constraints on factors

➡ Near-optimality of $(\hat{\mathbf{X}}_1, \ldots \hat{\mathbf{X}}_L) = \texttt{butterfly}(\mathbf{A}, \beta)$

$$\|\mathbf{A} - \hat{\mathbf{X}}_1 \ldots \hat{\mathbf{X}}_L\| \leq \sqrt{L}\ E(\mathbf{A})$$

Practical behaviour

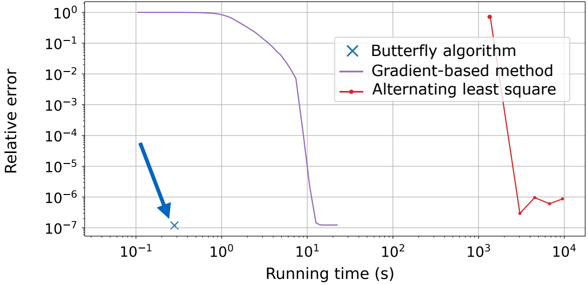

Quantized butterflies


Elisa Ricciotti

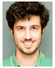
Théo Mary

Unquantized butterfly

- high-precision weights

$\mathbf{F} =$ 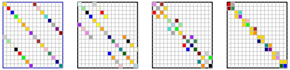

goal

Quantized butterfly

- *low-precision format (eg. float8)*

$\mathbf{F} =$ 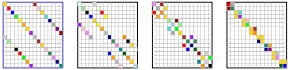

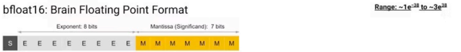

- **naive approach :** entry wise rounding
- **opportunity:** rescaling-invariance $\mathbf{XY} = (\lambda \mathbf{X})(\mathbf{Y}/\lambda)$

➡ **new algorithm**
- optimal for L=2 factors
- 30% less bits than naive rounding for same precision

Roadmap

Batch vs Sequential Learning

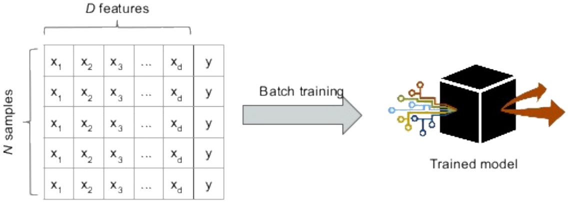

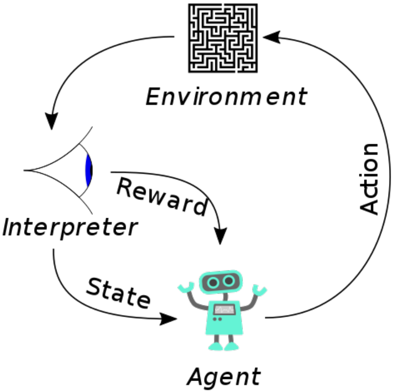

What is Reinforcement Learning?

- RL is a paradigm of machine learning where an agent learns by interacting with an environment.
- The agent takes actions, receives feedback in the form of rewards, and updates its policy.
- Goal: learn a strategy (policy) to maximize cumulative reward.

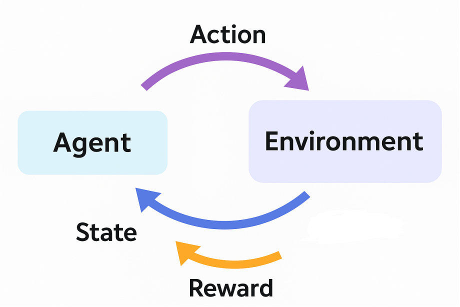

Applications of RL in Physics

- **Quantum control**: optimize pulse sequences for manipulating quantum states.
- **Plasma control**: real-time decisions in fusion reactors.
- **High-energy physics**: trigger systems and simulation optimization.

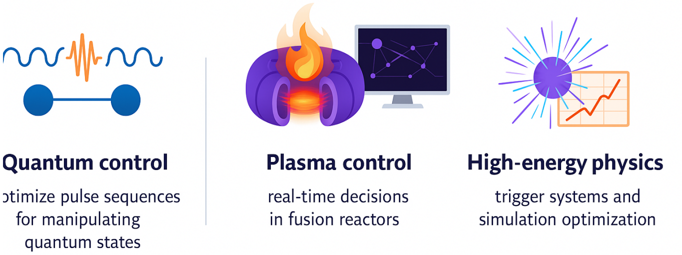

Applications of RL in Chemistry

- **Molecular generation**: design novel molecules with targeted properties.
- **Reaction optimization**: adaptive experimentation in synthesis pathways.
- **Automated drug discovery**: exploration of chemical space via reward-driven search.
- Integrates well with simulations and real-world experimentation.
- Promising for autonomous labs and closed-loop scientific pipelines.

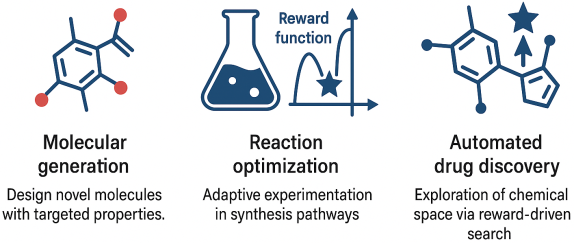

A-B Testing

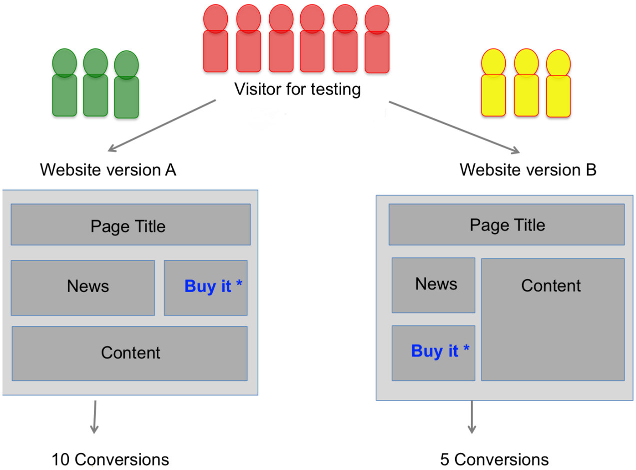

Visitor for testing
Website version A
Page Title
News
Buy it *
Content
10 Conversions
Website version B
Page Title
News
Content
Buy it *
5 Conversions

Monte Carlo Tree Search

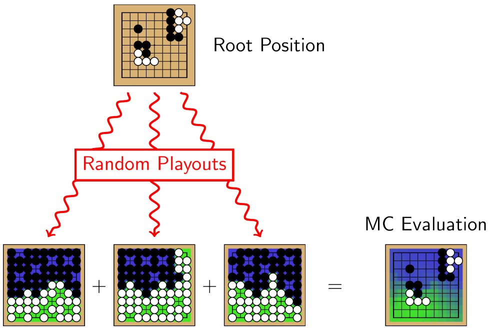

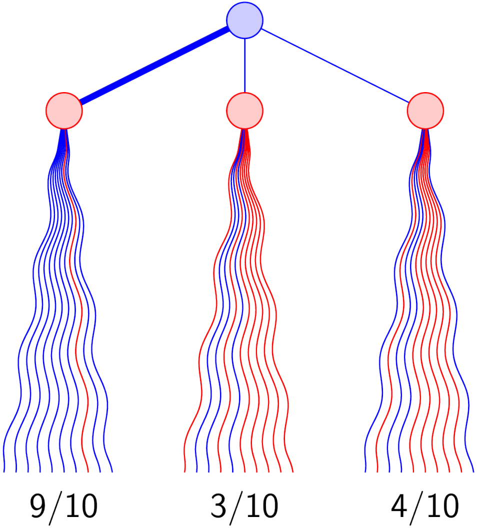

Optimal Discovery

Power system security assessment

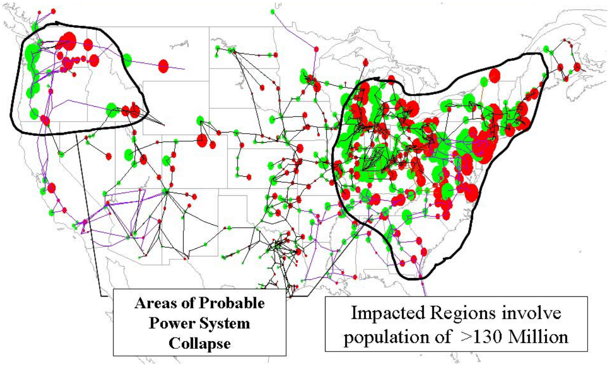

By Mark MacAlester, Federal Emergency Management Agency [Public domain], via Wikimedia Commons

How to identify quickly contingencies/scenarios that could lead to unacceptable operating conditions (dangerous contingencies) if no preventive actions were taken?

Black Box Optimization

Black-box interaction model:

- choose $X_1 = \phi_1(U_1)$, observe $Y_1 = F(X_1, \omega_1)$
- choose $X_2 = \phi_2(X_1, Y_1, U_2)$ observe $Y_2 = F(X_2, \omega_2)$,
- etc...

Target = optimize f: $f^* = \max_{\mathcal{X}} f$ (or min, or find level set, etc.)
Strategy: sampling rule $(\phi_t)_{t\geq 1}$, stopping time τ
PAC setting: for a risk δ and a tolerance ϵ, return $X_{\tau+1}$ such that

$$P_f\Big(f(X_{\tau+1}) < f^* - \epsilon\Big) \leq \delta \ .$$

The set of correct answers is $\mathcal{X}_\epsilon(f) = \Big\{x \in X : f(x) \geq f^* - \epsilon\Big\}$.
Goal: find a strategy minimizing $E_f[\tau]$ for all possible functions f
(no sub-optimal multiplicative constant!)

Vanilla Bandit Model

Which treatment is most efficient to cure a disease?

		Treatment			
Patient	1	1			
	2		0		
	3			1	
	4				1
	5			0	
	6	1			
	7		0		
	8				0
	9	0			
	10	0			
	11			1	
	12				1
	13				1
	14			1	
	15	?	?	?	?
Empirical efficiency		**50%**	**0%**	**75%**	**75%**

0: non-recovered patient
1: recovered patient

Thompson (1933)

Graph Bandits

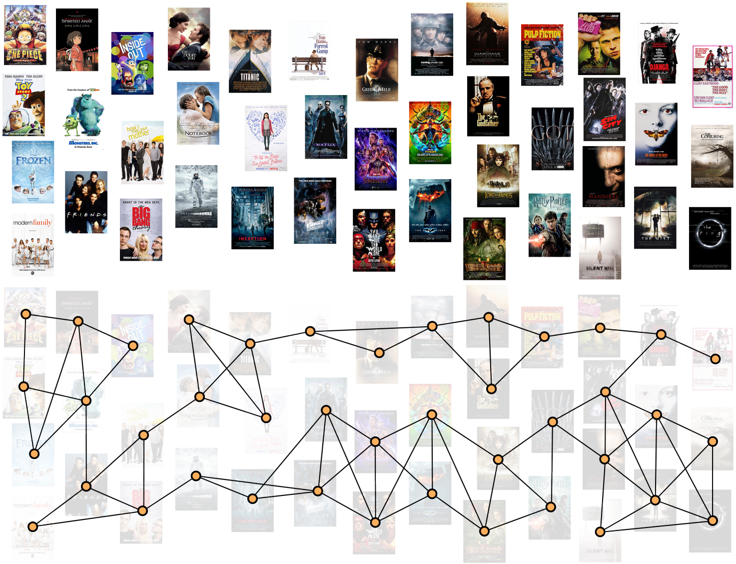

Graph Smoothness

See [*Spectral bandits for smooth graph functions*, by M. Valko, R. Munos, B. Kveton, T. Kocák]

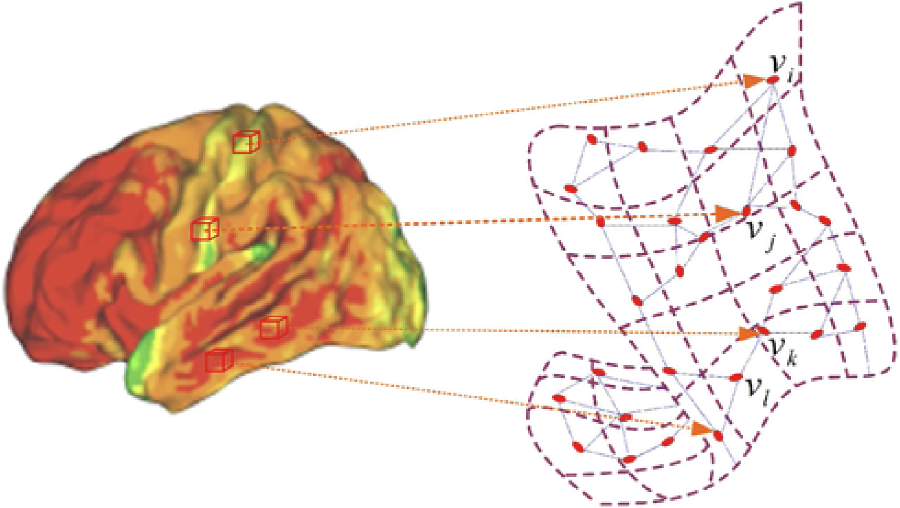

Black Box optimization

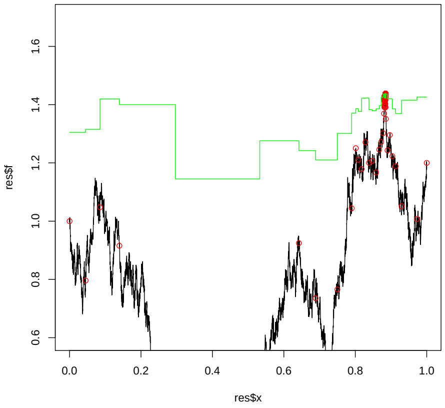

Stochastic Optimization

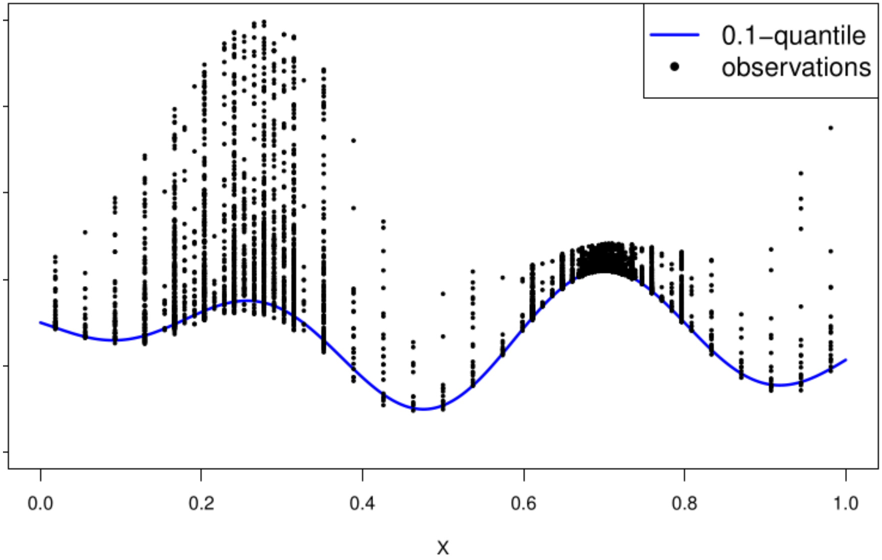

Roadmap

Distributional RL

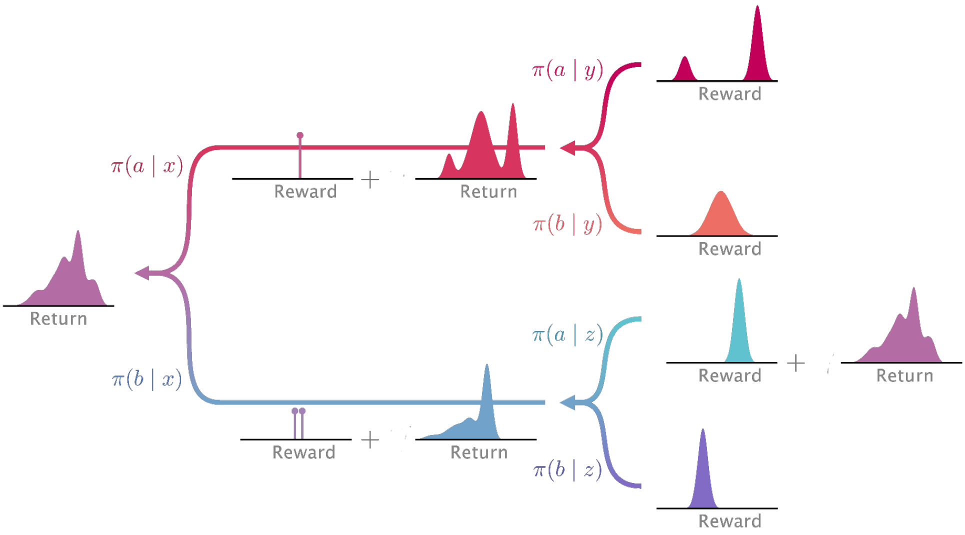

Risk-aware RL

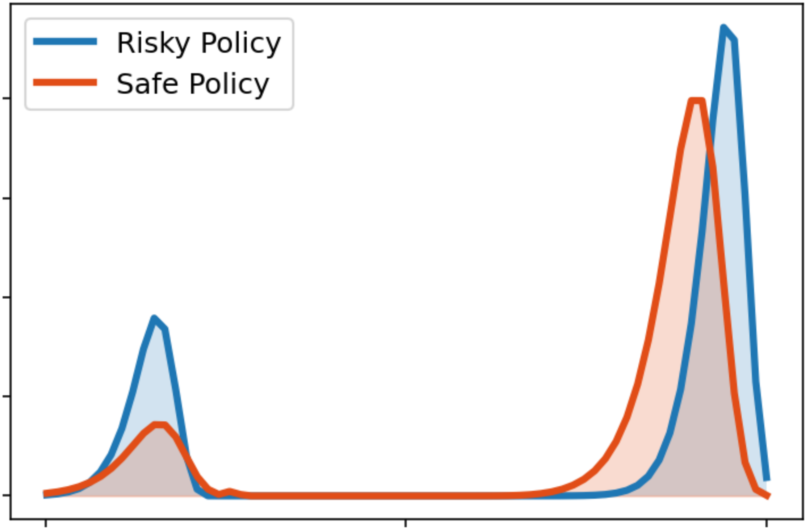

Example: Cliff

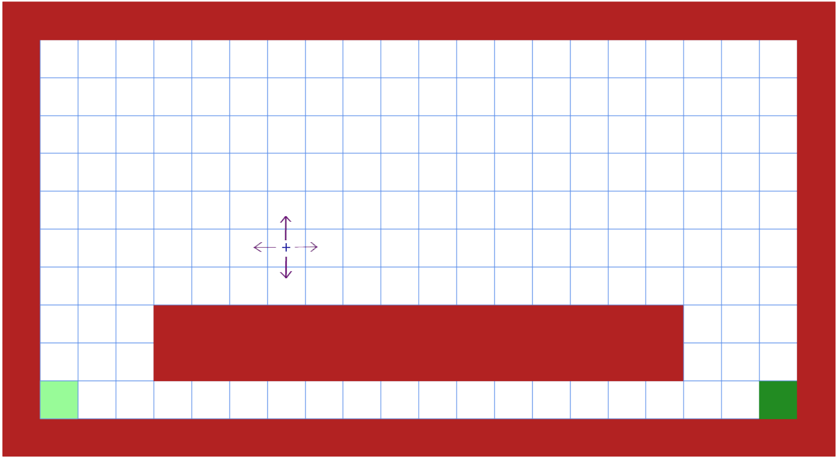

Example: Cliff

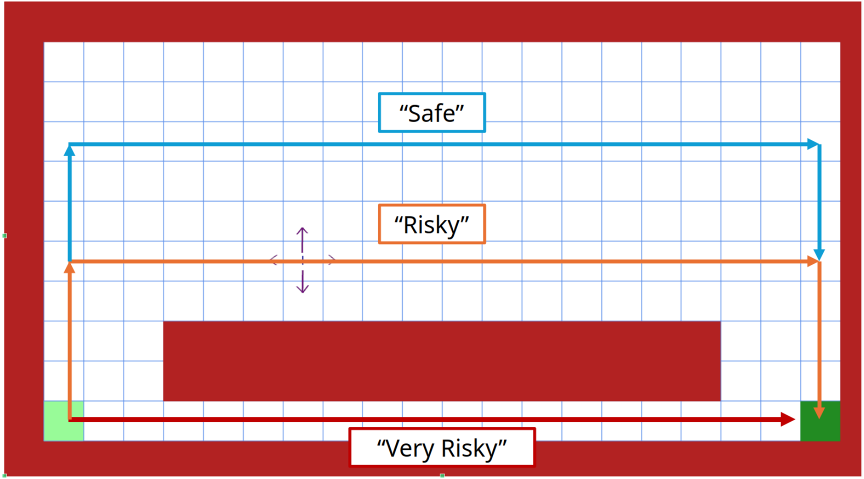

Example: Cliff environment

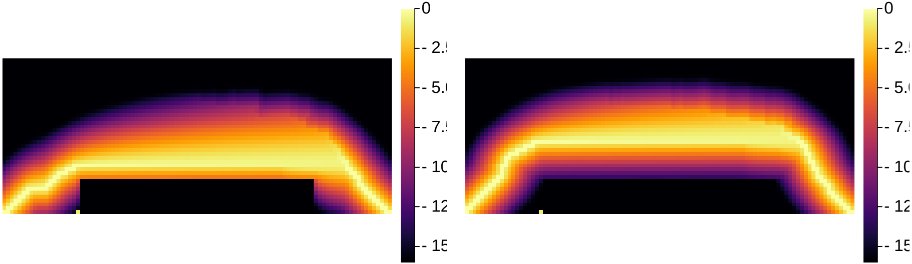

AILyS

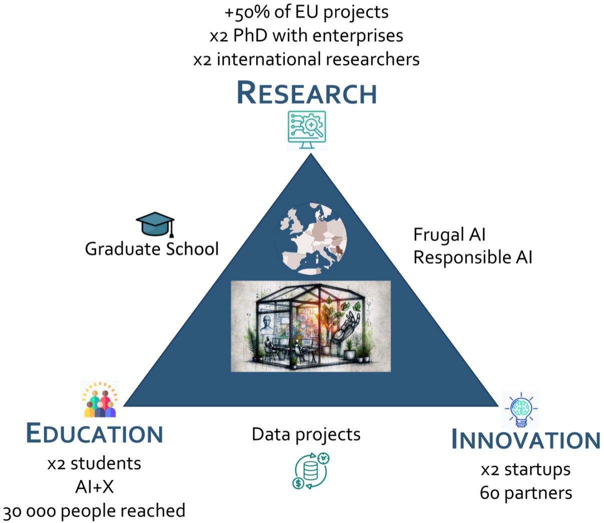
+50% of EU projects
x2 PhD with enterprises
x2 international researchers
RESEARCH
Graduate School
Frugal AI
Responsible AI
EDUCATION
x2 students
AI+X
30 000 people reached
Data projects
INNOVATION
x2 startups
60 partners