GDL final lecture

Generative and Deep Learning (GDL)

Davide Bacciu (davide.bacciu@unipi.it)

Università di Pisa

Lecture Outline

- Course wrap-up
  - A-posteriori view of the course
  - Final take home messages
- Overview of GDL research @UNIPI
- GLD Final exam
- Conclusions & Discussion

The Course

The Course in 1-Slide

- Probabilistic models
  - Infer knowledge rather than just predict
  - Introduce prior knowledge
  - Reason causally on models' knowledge
  - How to approximate distributions
  - How to sample distributions
- Deep Neural Network
  - Efficient and high predictive performance
  - Non-parametric and non-linear
  - Work on noisy, raw and heterogeneous data
  - Carefully design the inductive bias (or be ready to pay the toll)
- Deep Generative models
  - Learn a generative process
  - How to approximate complex distributions
  - How to sample complex distributions
  - Probabilities are close friends with vector fields

GDL: a Convergence of Neural-Generative Paradigms

- Need the efficacy and efficiency of neural models with the interpretability and generative ability of probabilistic-based models
- Modular approach
  - E.g. CRF on the top of CNN for semantic segmentation
  - Easily incorporate prior knowledge
- Inbreeding of paradigms
  - CRF as discriminative-generative-energy based models
  - Variational and generative DL
  - Statistics, optimization and differential equations
- Mutual support
  - Causal learning

After Completing This Course, Hopefully…

- Know which learning paradigms and models are best to start with for addressing a problem di modern ML
- Know what challenges your model will need to solve to realize an application
- Know a bag of tricks to modify a model to suit your needs
  - Priors, posteriors and likelihoods, message passing, variational approximations, sampling, latent representations, feature functions, diffusion processes, change of variables, flows and transport
  - Batch/activation normalization, pretraining, end-to-end differentiability, distribution learning with NN, dynamic memories, attention, adversarial learning, invertible neural layers, forward/backward information propagation, on the importance of gaussian noise
  - Randomized NNs, learning with graphs, causal Vs observational models

Machine Learning is the New Algorithmics

Can you derive EM for GMM?

When do you need to check you gradient?

GDL @ DI.UNIPI

My Research Overview

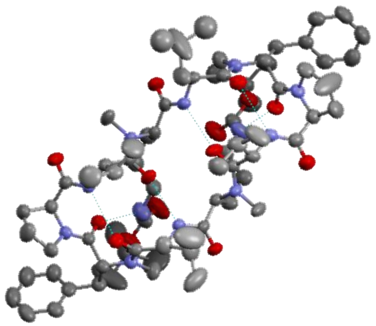

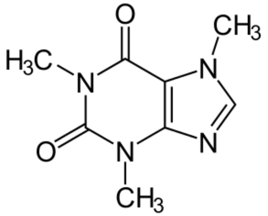

- Learning for graphs & structured data
- Deep learning and neural networks, generative models, RL, causality, AI for science
- Pervasive AI (distributed, embedded, morphological/neuromorphic, continual)

Funded by Italian, European and industrial projects

Snapshots of current research (Graphs)

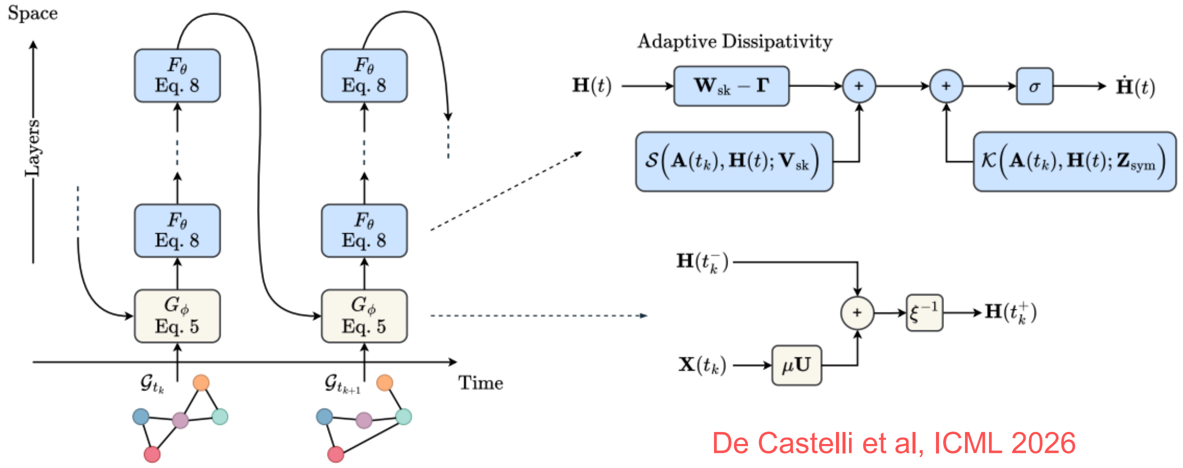

De Castelli et al, ICML 2026

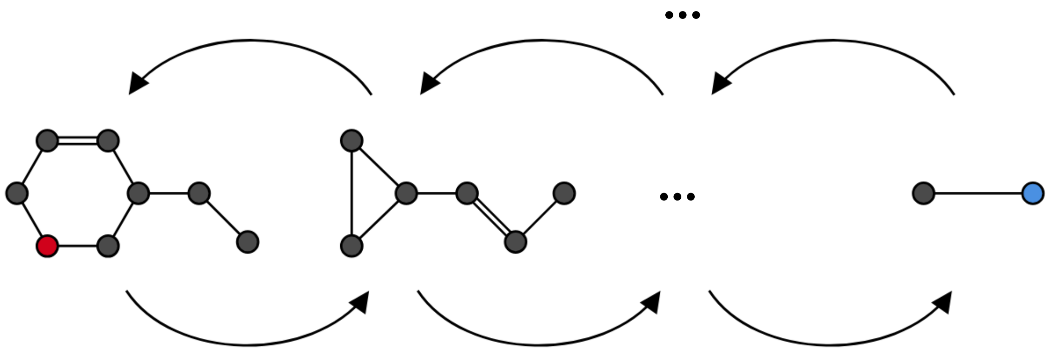

Ninniri et al, NeurIPS 2025

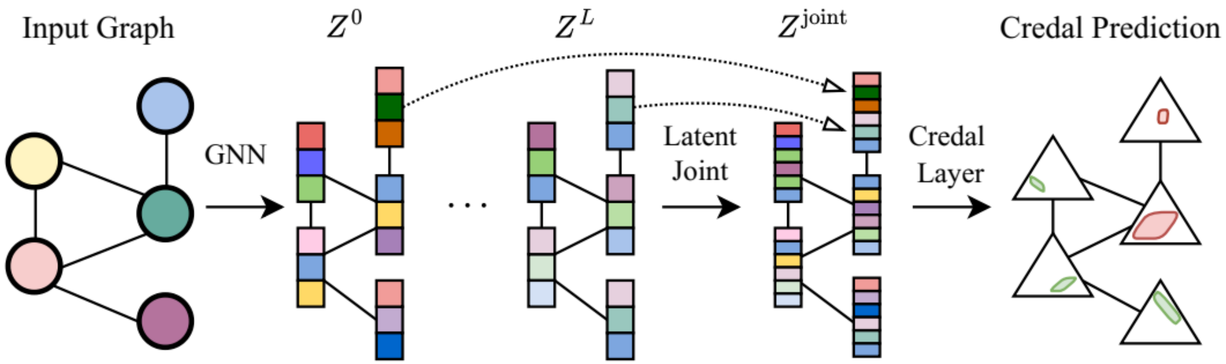

M. Tolloso et al, EurIPS 2025

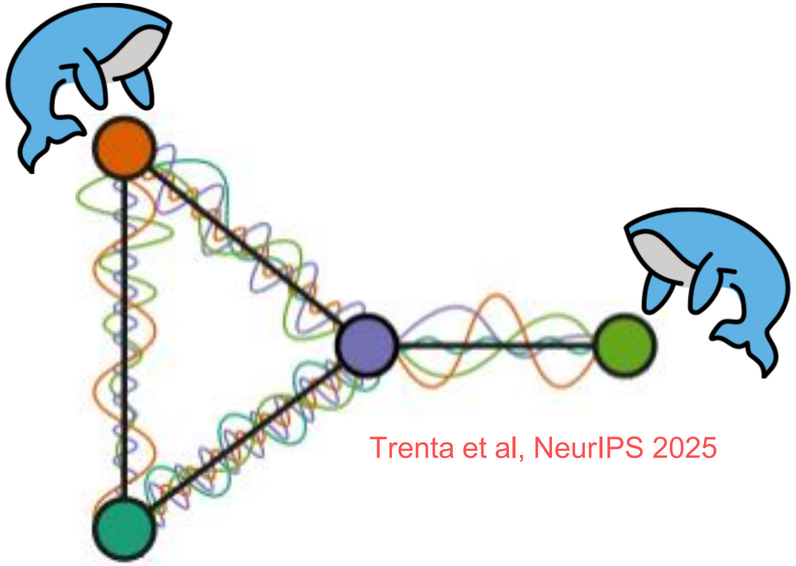

Trenta et al, NeurIPS 2025

Snapshots of current research (AI for science)

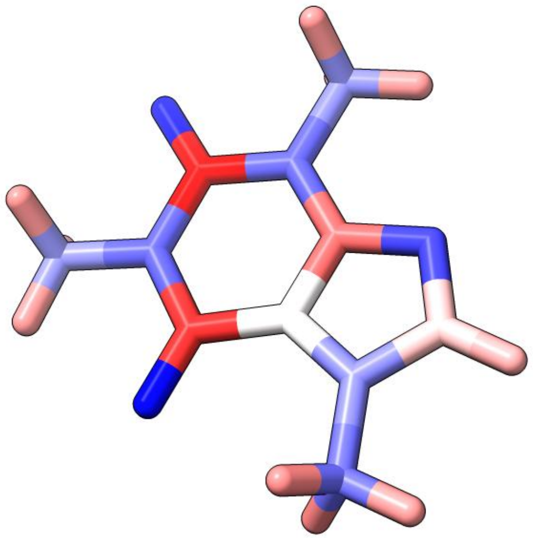

Miglior et al, ICLR 2026

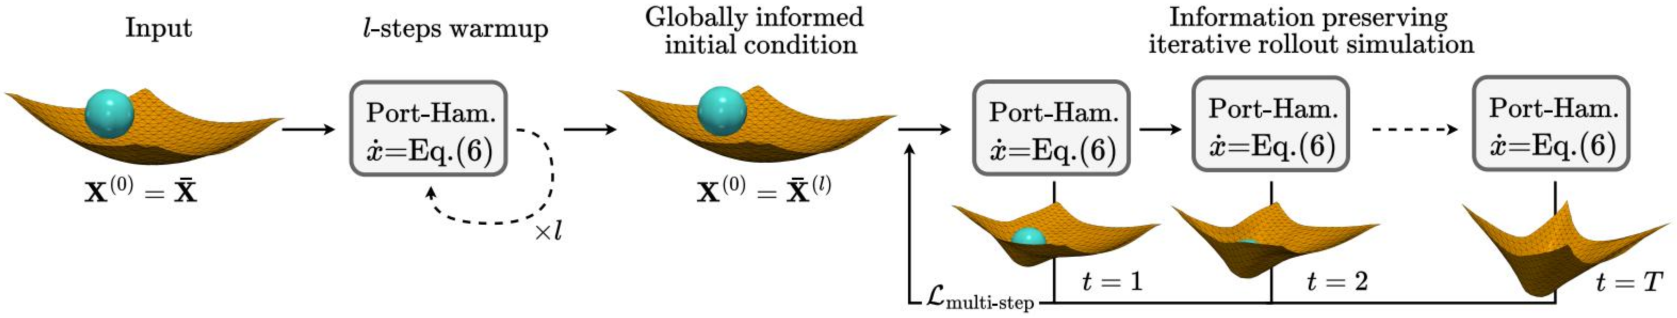

Hoang et al, ICLR 2026

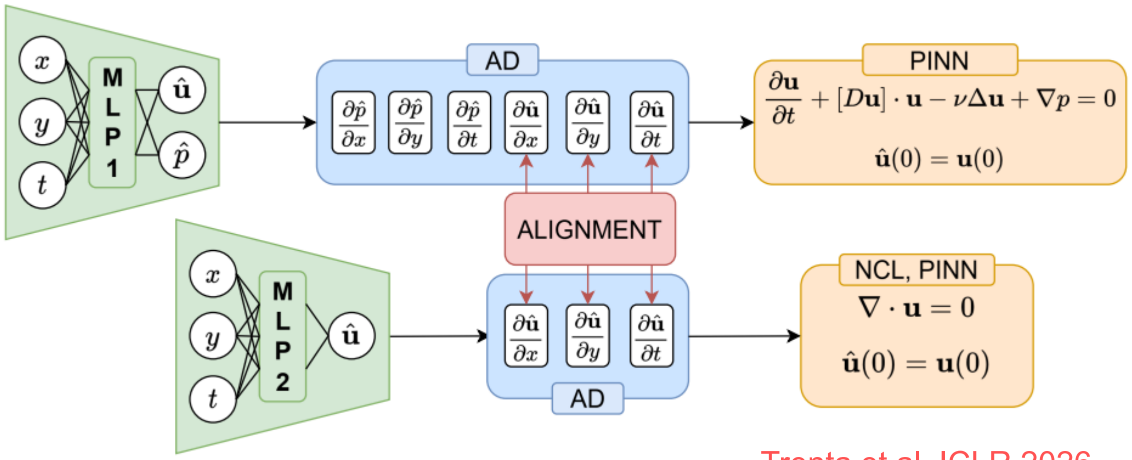

Trenta et al, ICLR 2026

Snapshots of current research (... and more)

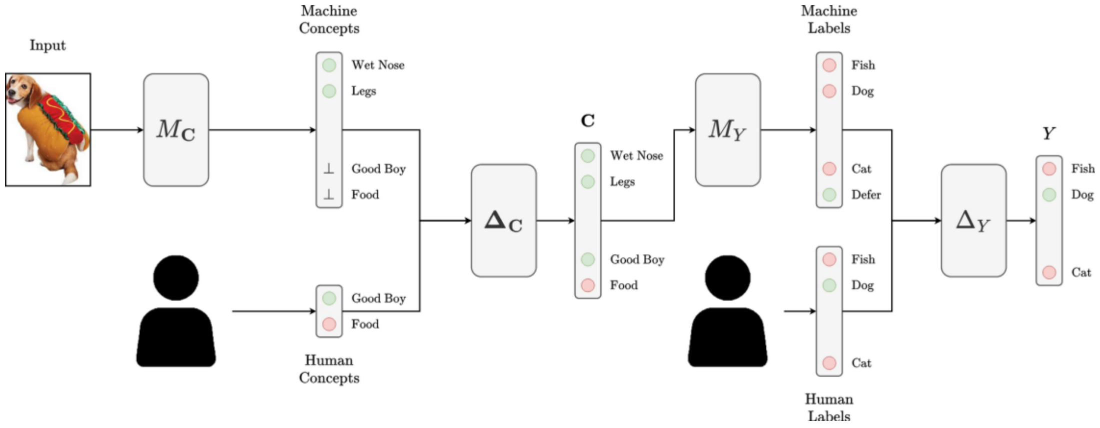

Pugnana et al, NeurIPS 2025

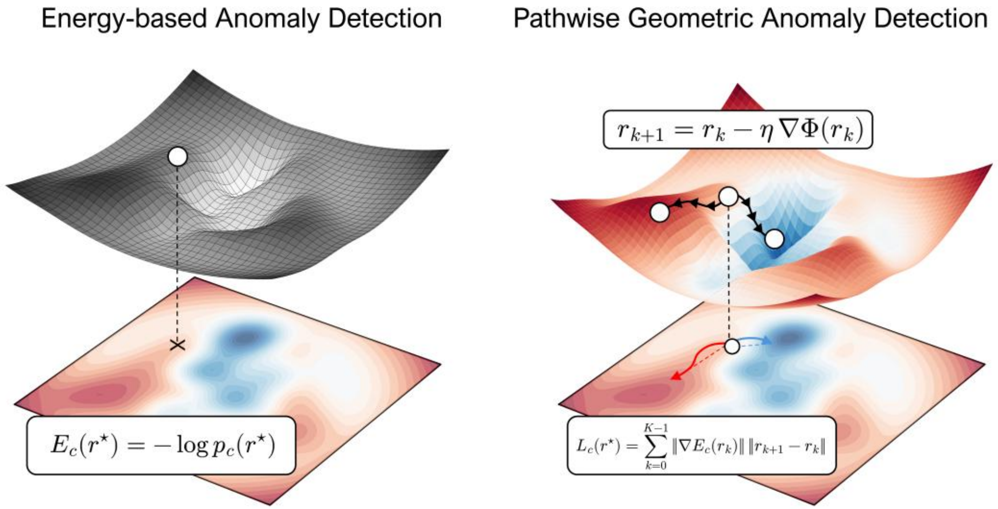

Recent Projects & Initiatives

- EU EIC Pathfinder SWIFT-BUILD (2026-2029)
- EU HorizonEU CoEvolution: Trustworthy Framework for Connected ML and (2024-2027)
- EU EIC Pathfinder EMERGE: Emergent awareness from minimal collectives (2022-2026)
- EU H2020 TEACHING: A computing Toolkit for building Efficient Autonomous appliCations leveraging Humanistic INtelliGence (2020-2023)
- FAIR AI research network: Continual learning transversal project & deep learning for graphs (2022-2025)
- EU H2020 TAILOR: Trustworthy AI Integrating Learning, Optimization and Reasoning (2020-2023)
- MIT-UNIPI international project on deep learning for optimization problems
- NEC-LABS industrial collaboration on learning inductive bias
- H&M Industrial collaboration on Deep Learning for Graphs
- Startups: Aptus AI, ContinualIST, QuantaBrain

TEACHING

TEACHING - Motivation & Vision

A human-centric perspective on autonomous CPSoS applications

TEACHING - Motivation & Vision

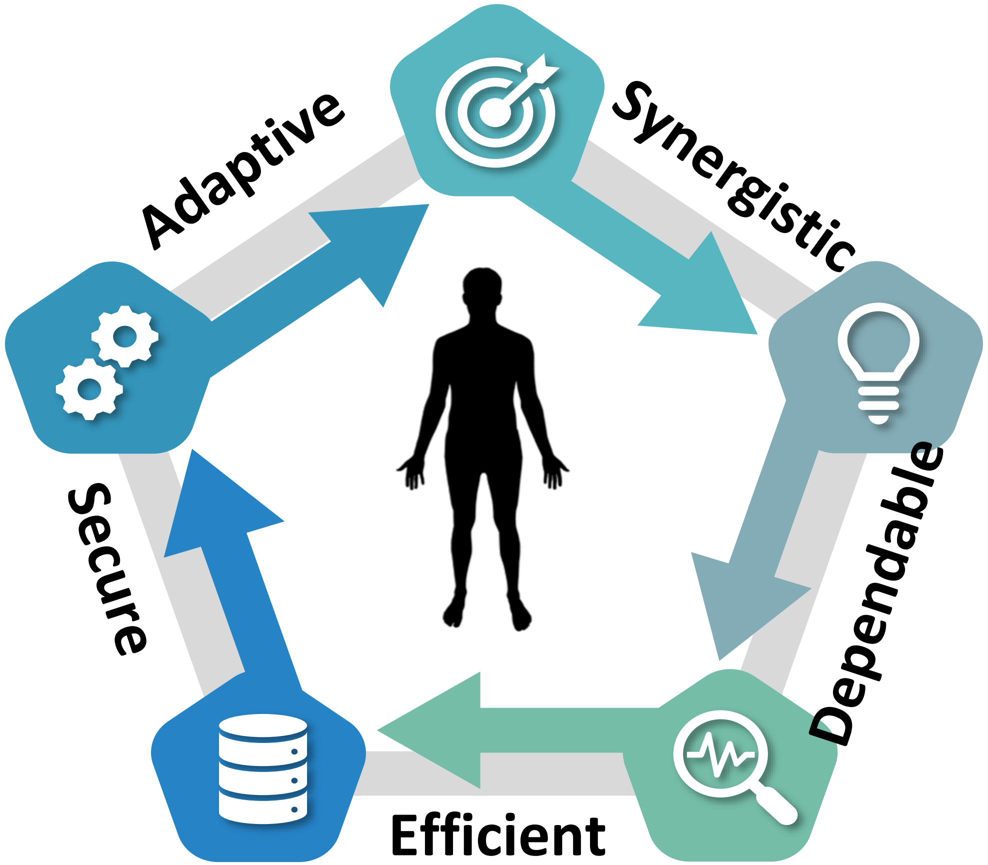

A human-centric perspective on autonomous CPSoS applications

Paradigmatic shift needing support at computing and system level

TEACHING – Autonomous Driving

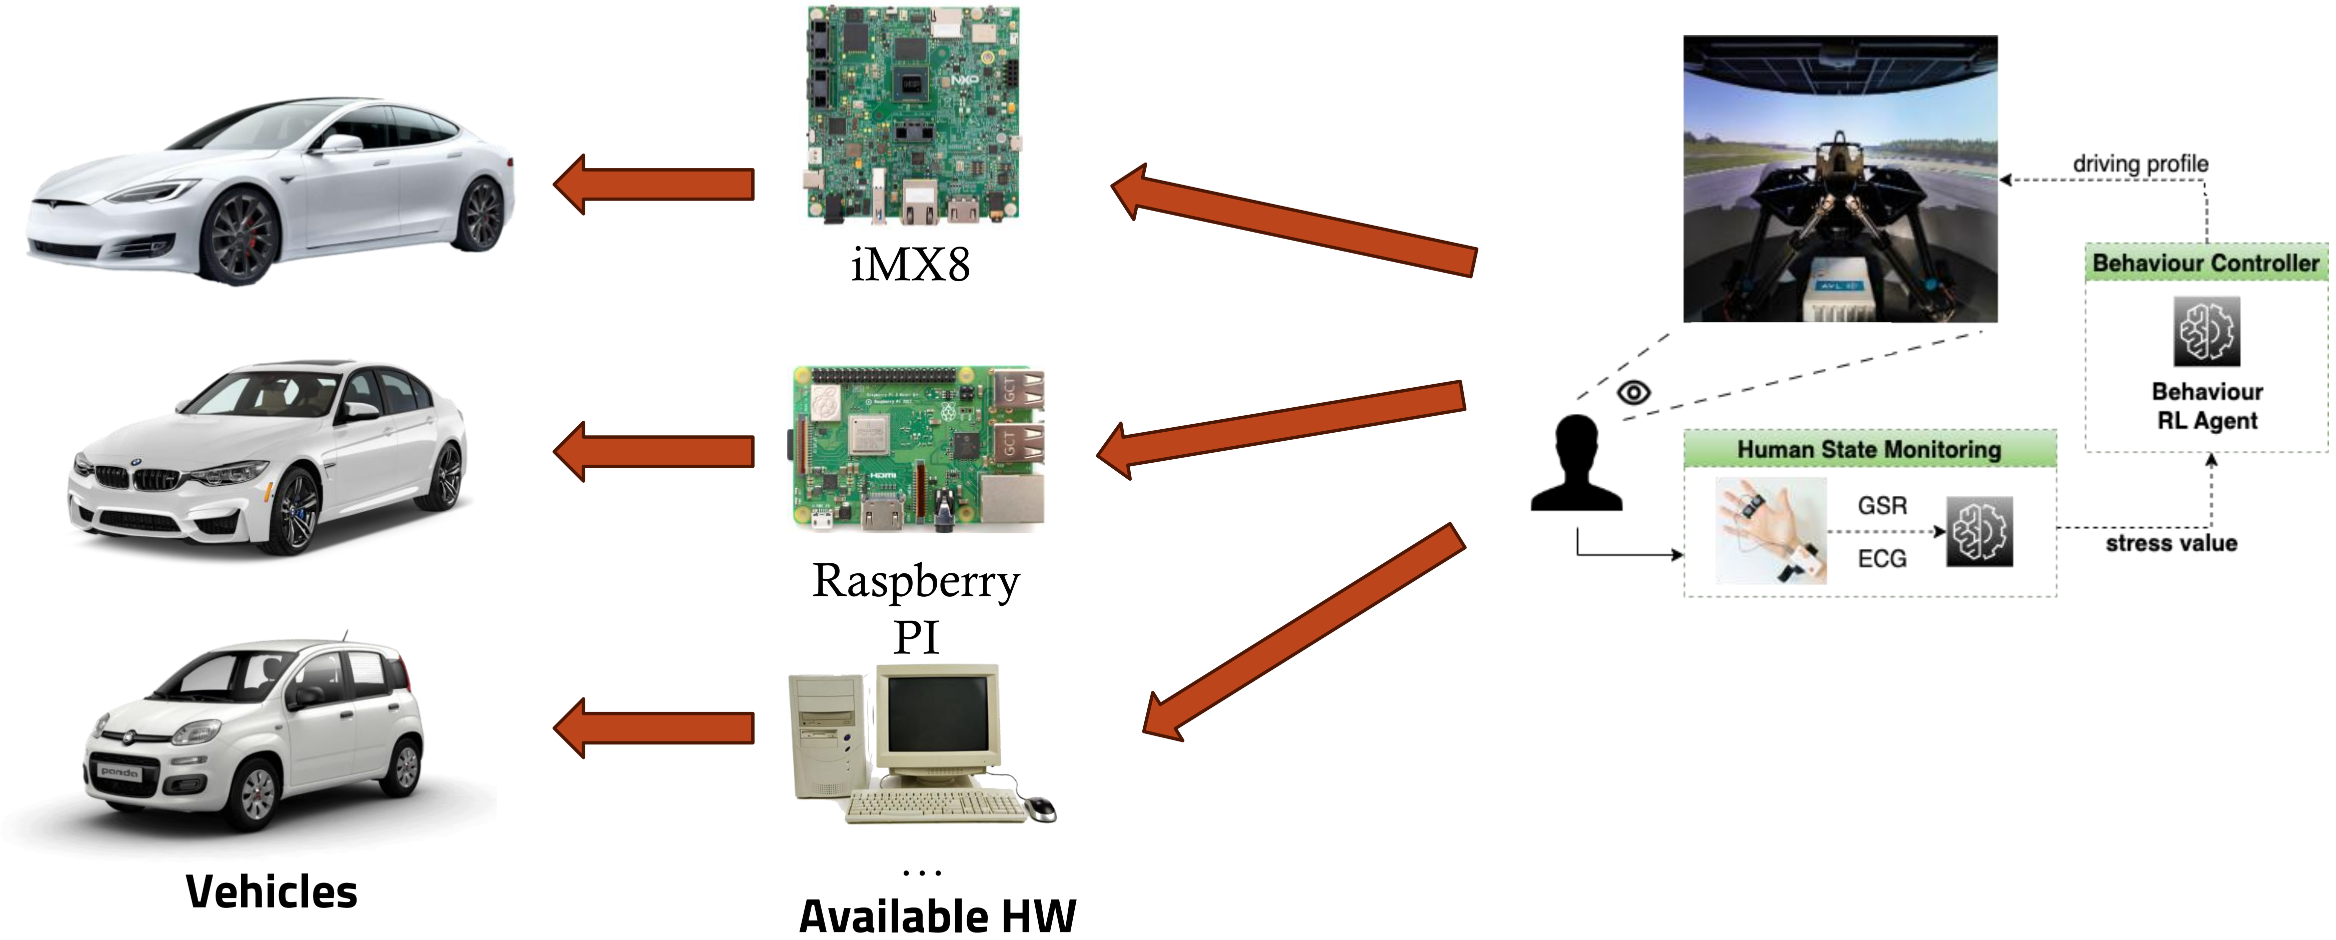

TEACHING Driving Platform @ PAILAB

Human state estimation

EMERGE – Awareness of Minimal Collectives

EMERGE – Aware Minimal Collectives

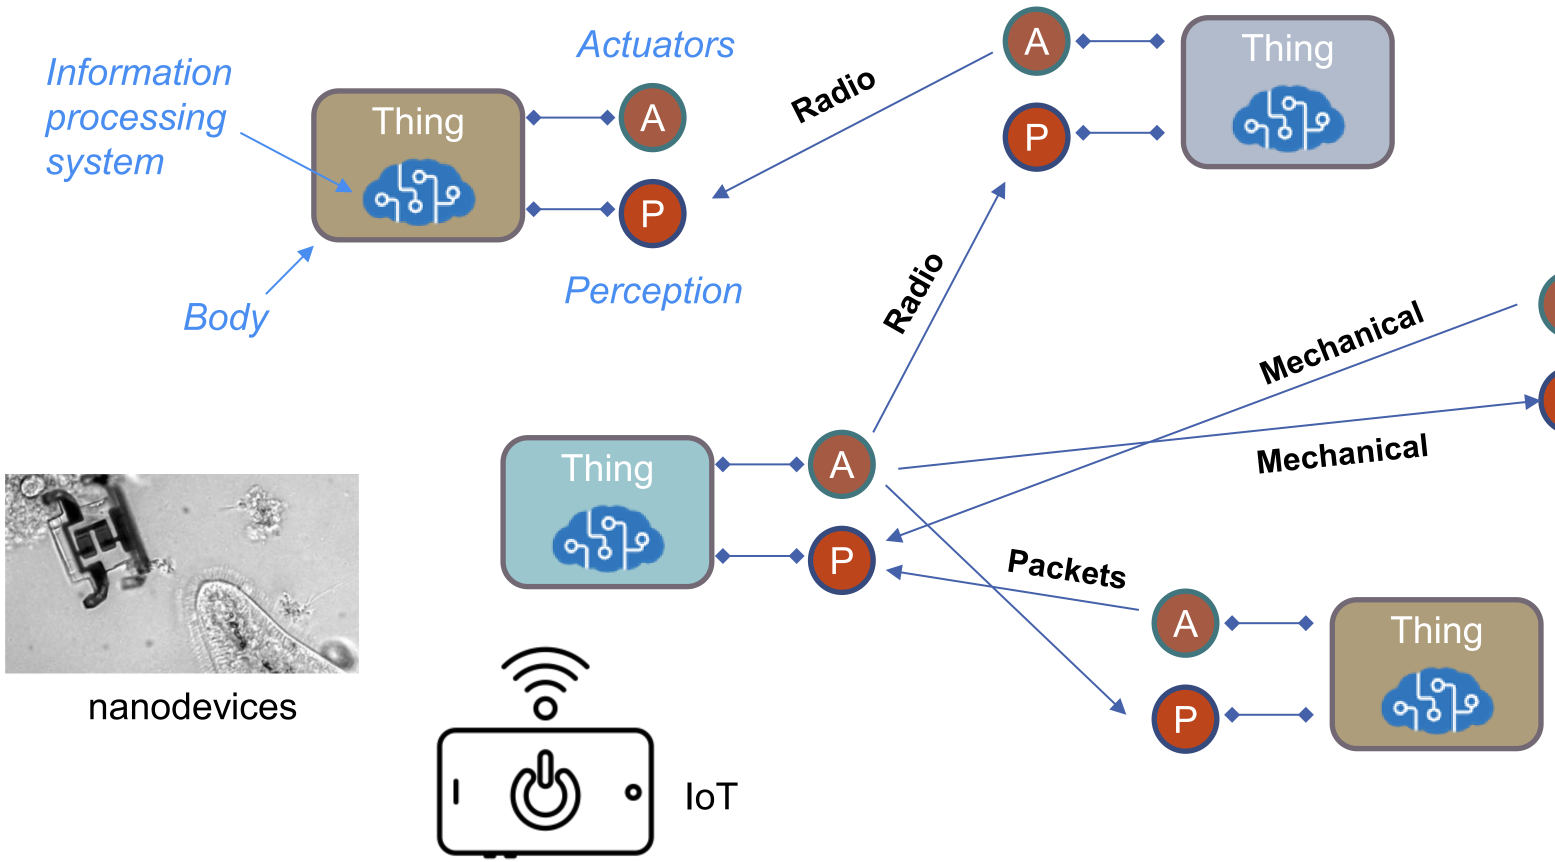

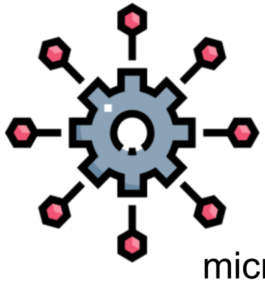

microservices

nanodevices

IoT

robots

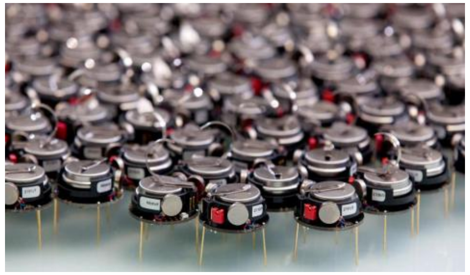

Novel Neural Computing Framework

Latent compressed representations for control

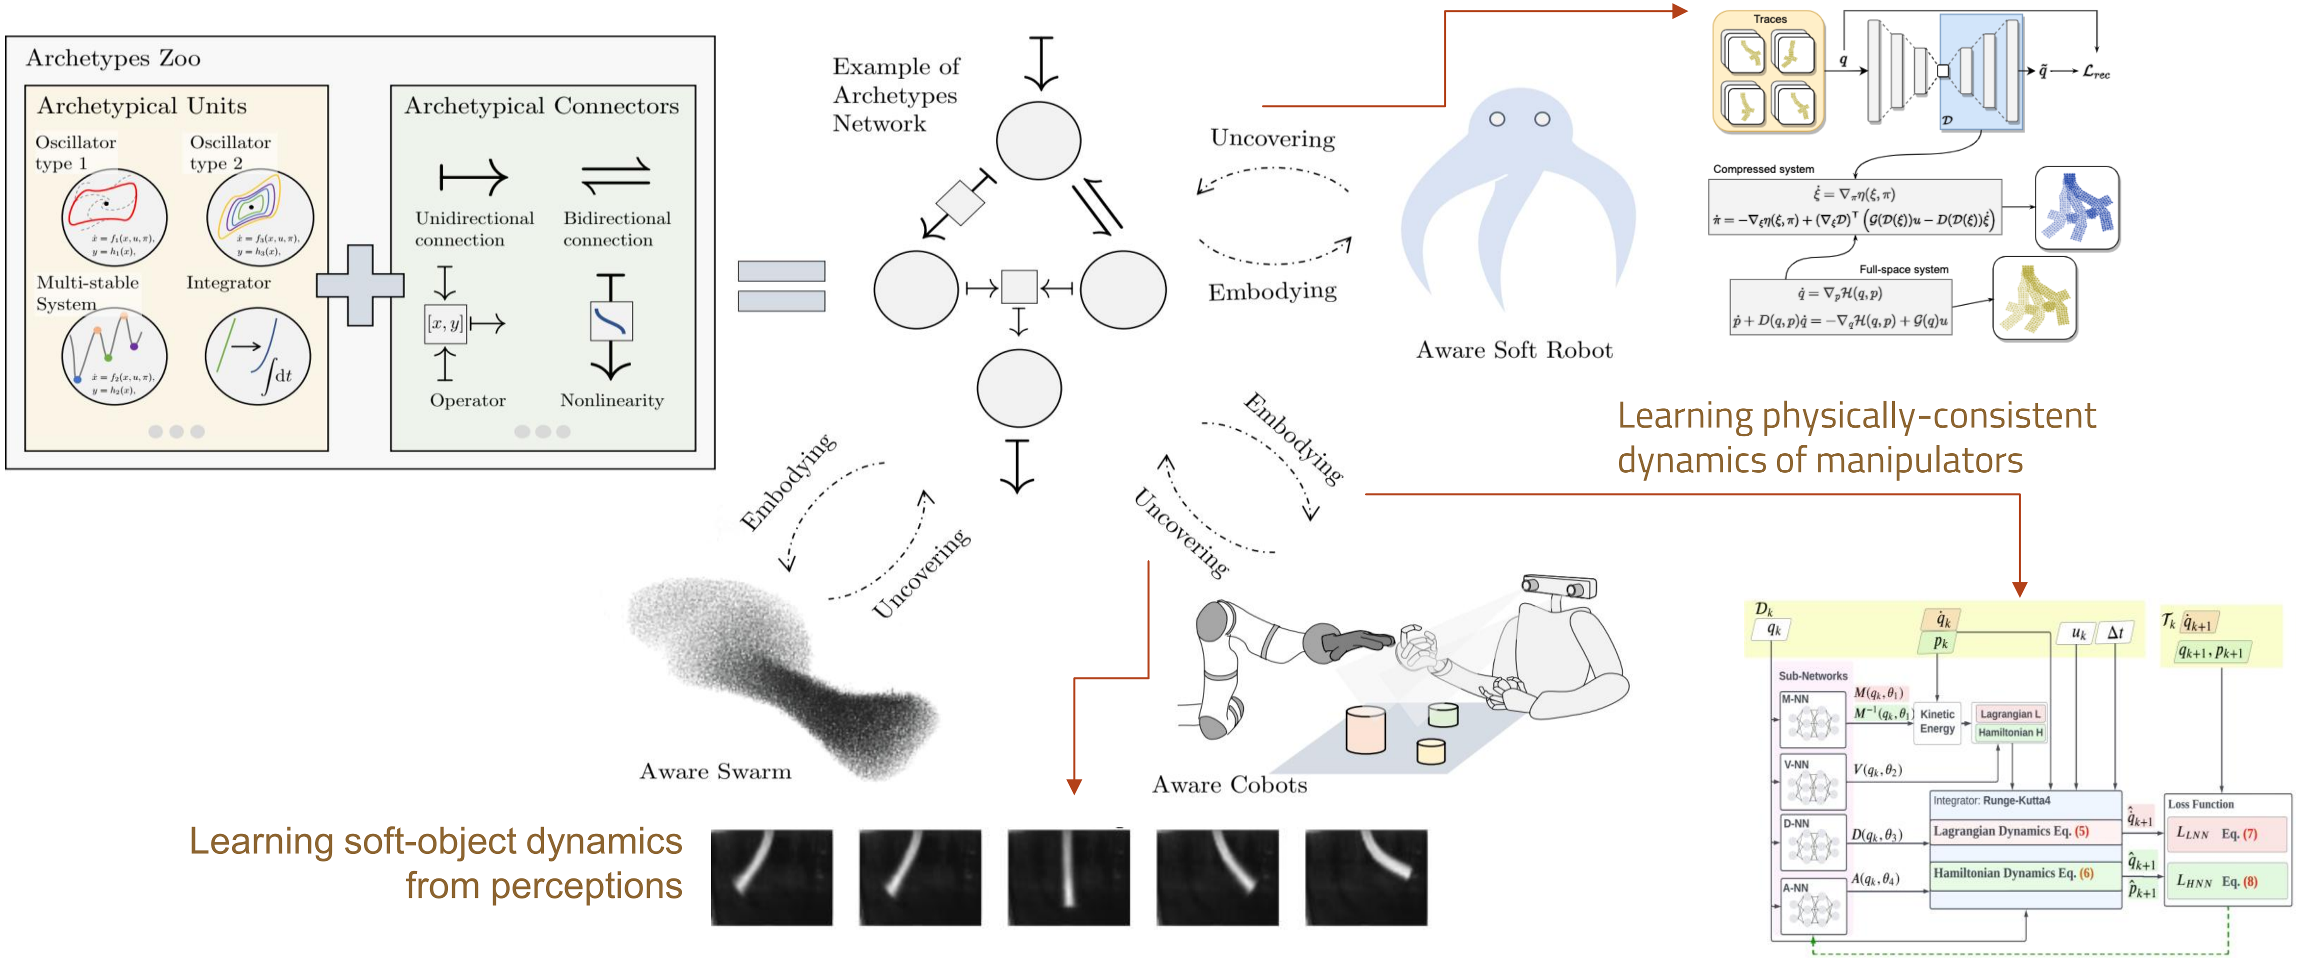

Learning physically-consistent dynamics of manipulators

Learning soft-object dynamics from perceptions

Neural Computing on Dynamical Systems

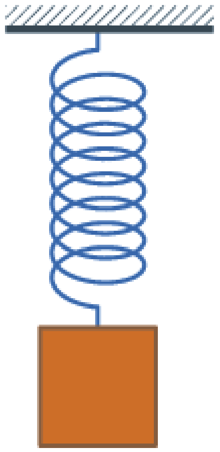

- Neural processing on multi-stable harmonic oscillators
- Physically implementable Neural Networks
- Neural Networks of Neural Networks

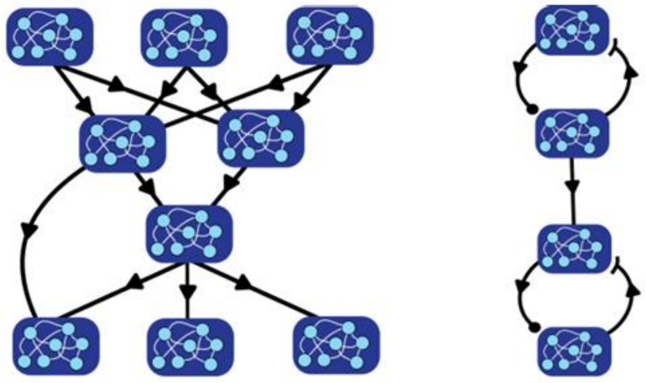

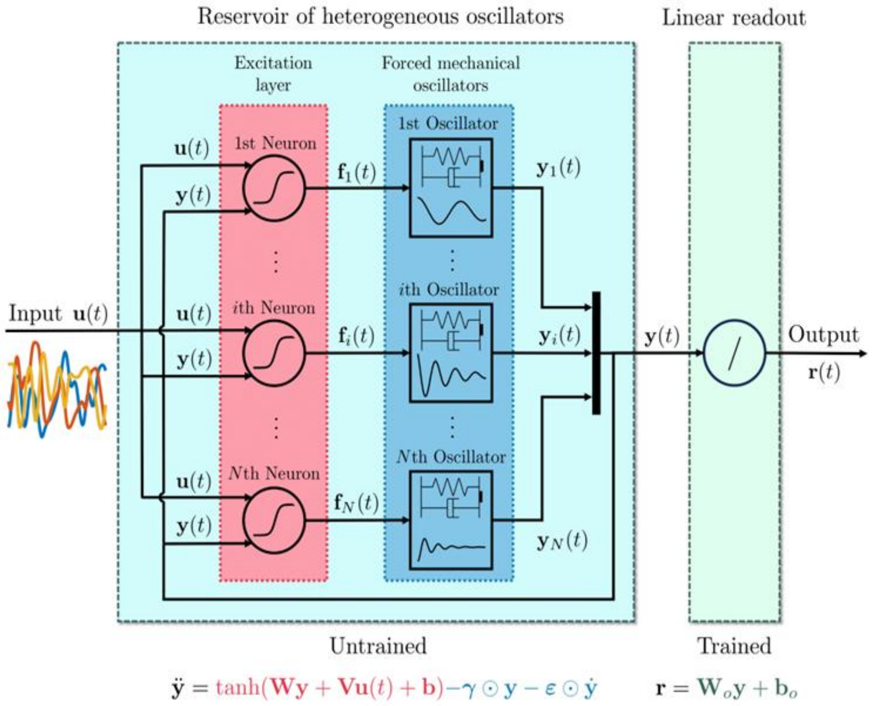

Neural computing with physical systems

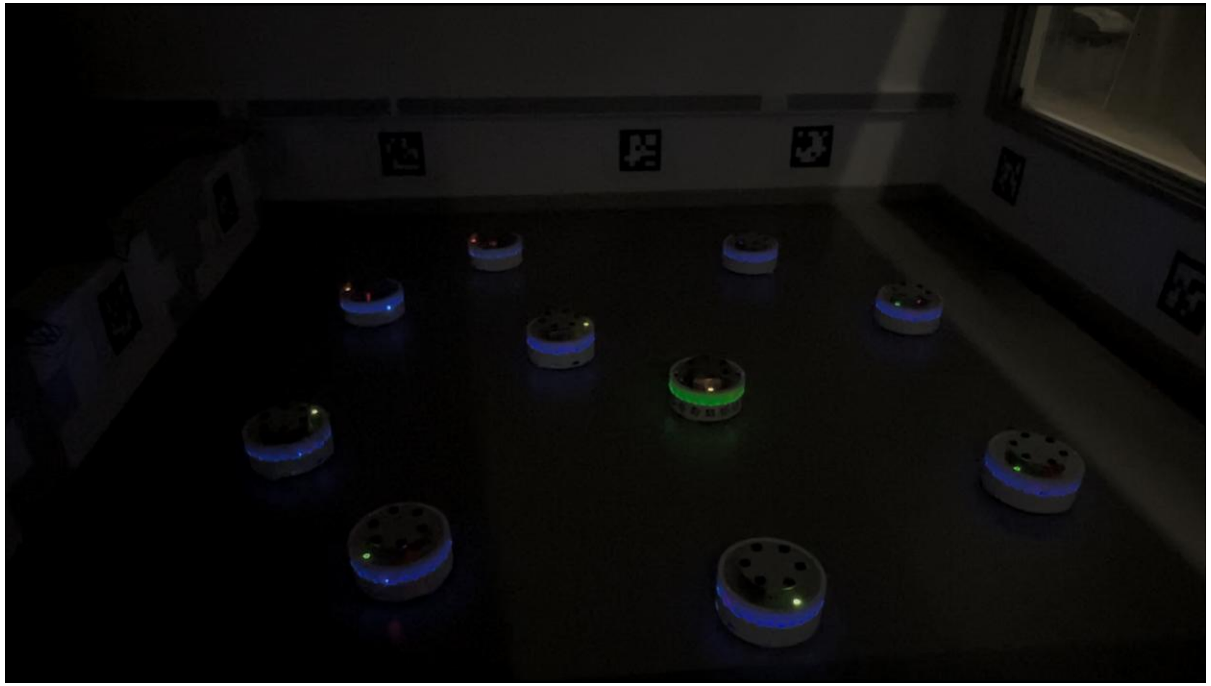

Robotic swarms

Mechanical neural networks

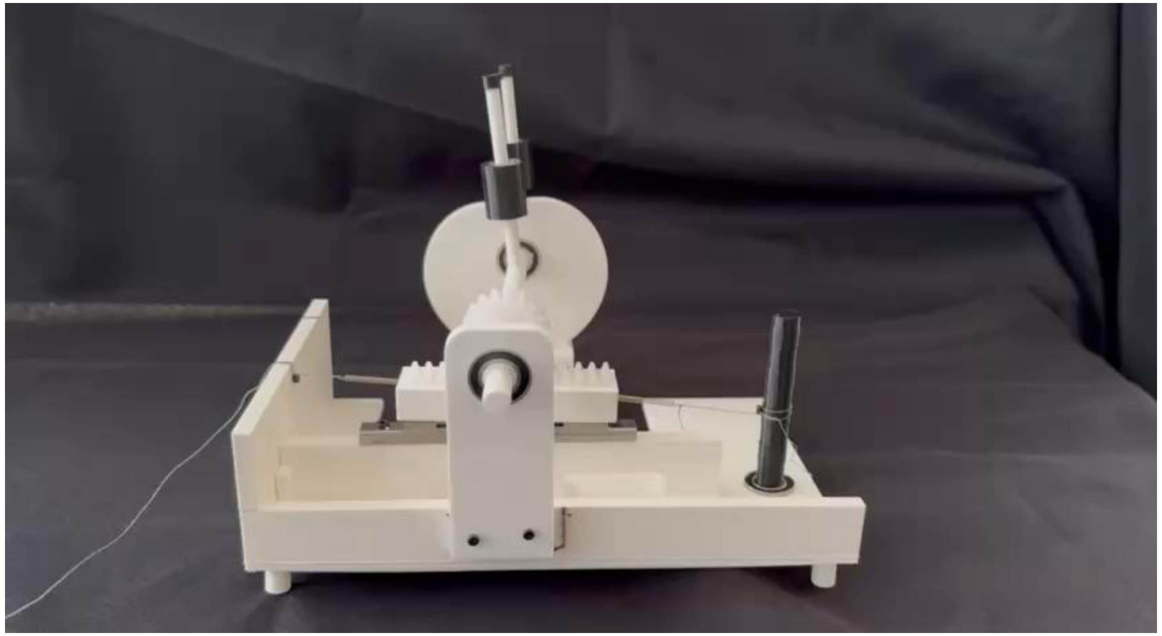

EMERGE

The **EMERGE** consortium:

Università di Pisa

TU Delft

University of Bristol

LMU Ludwig-Maximilians-Universität München

DA VINCI LABS

Funded by the European Union

Funded by the European Union under Grant Agreement 101070918

SWIFT-BUILD: SWarm-based Inverted Fabrication for Timber Buildings

Coming up soon!

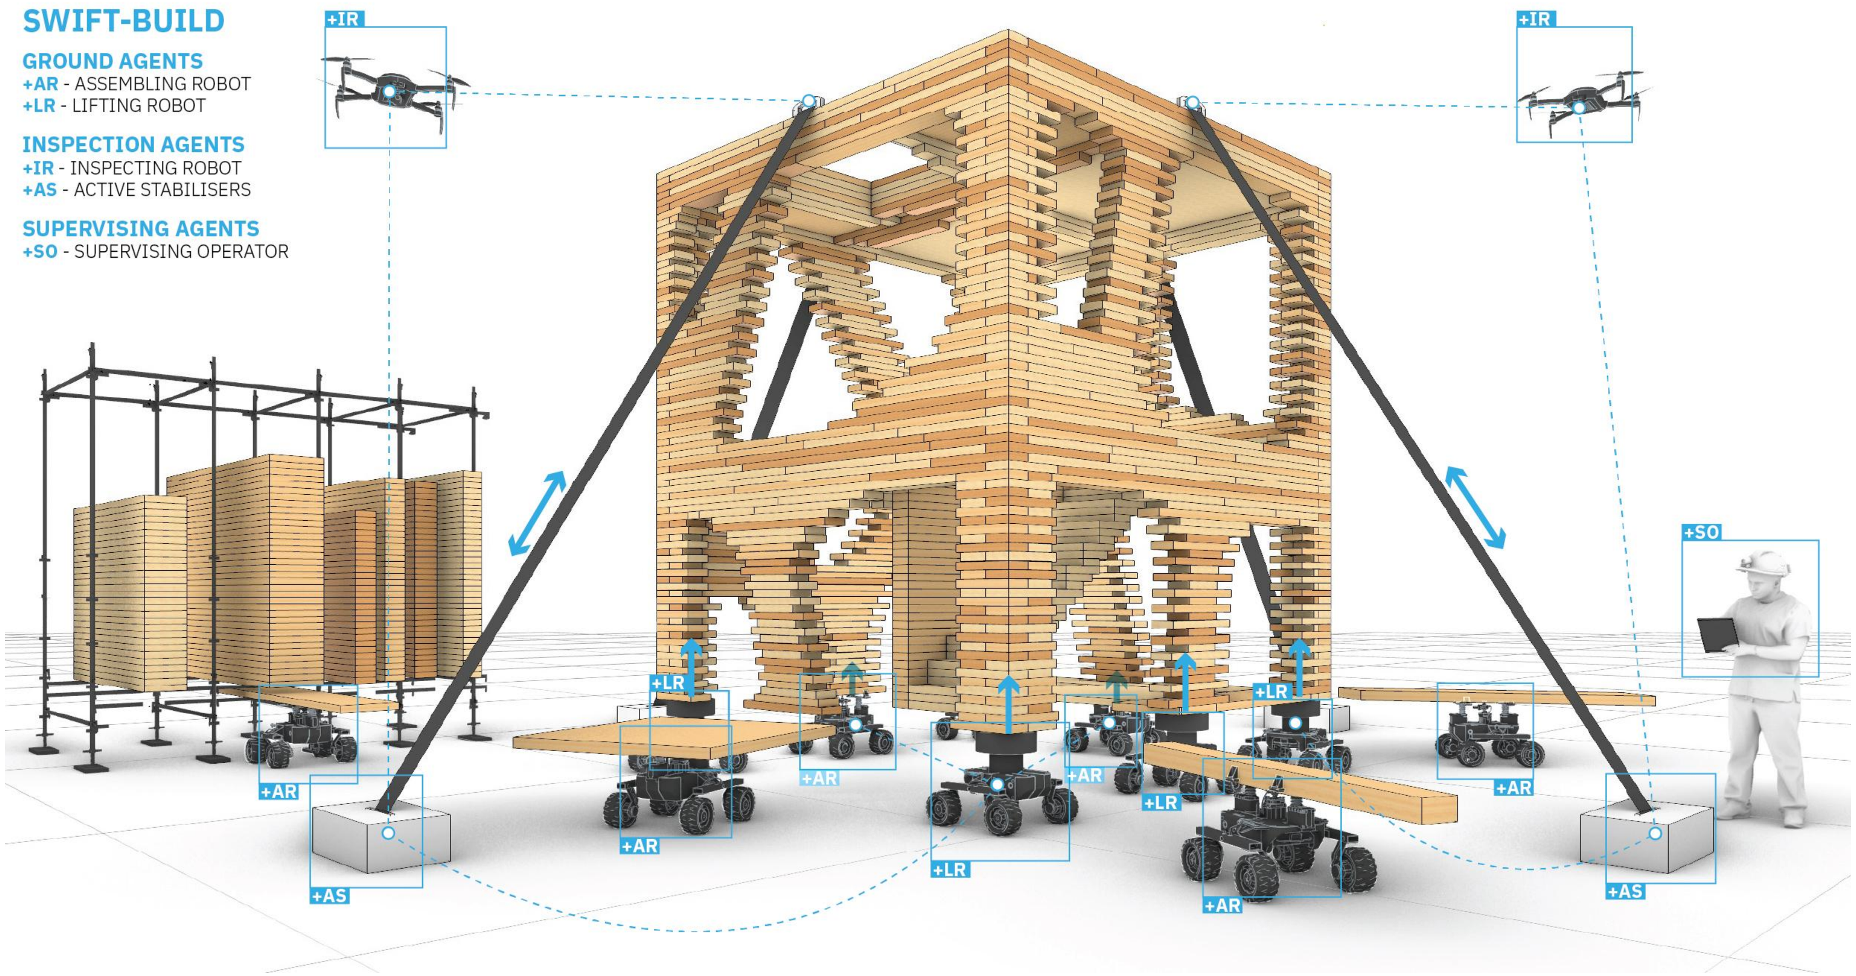

CoEvolution

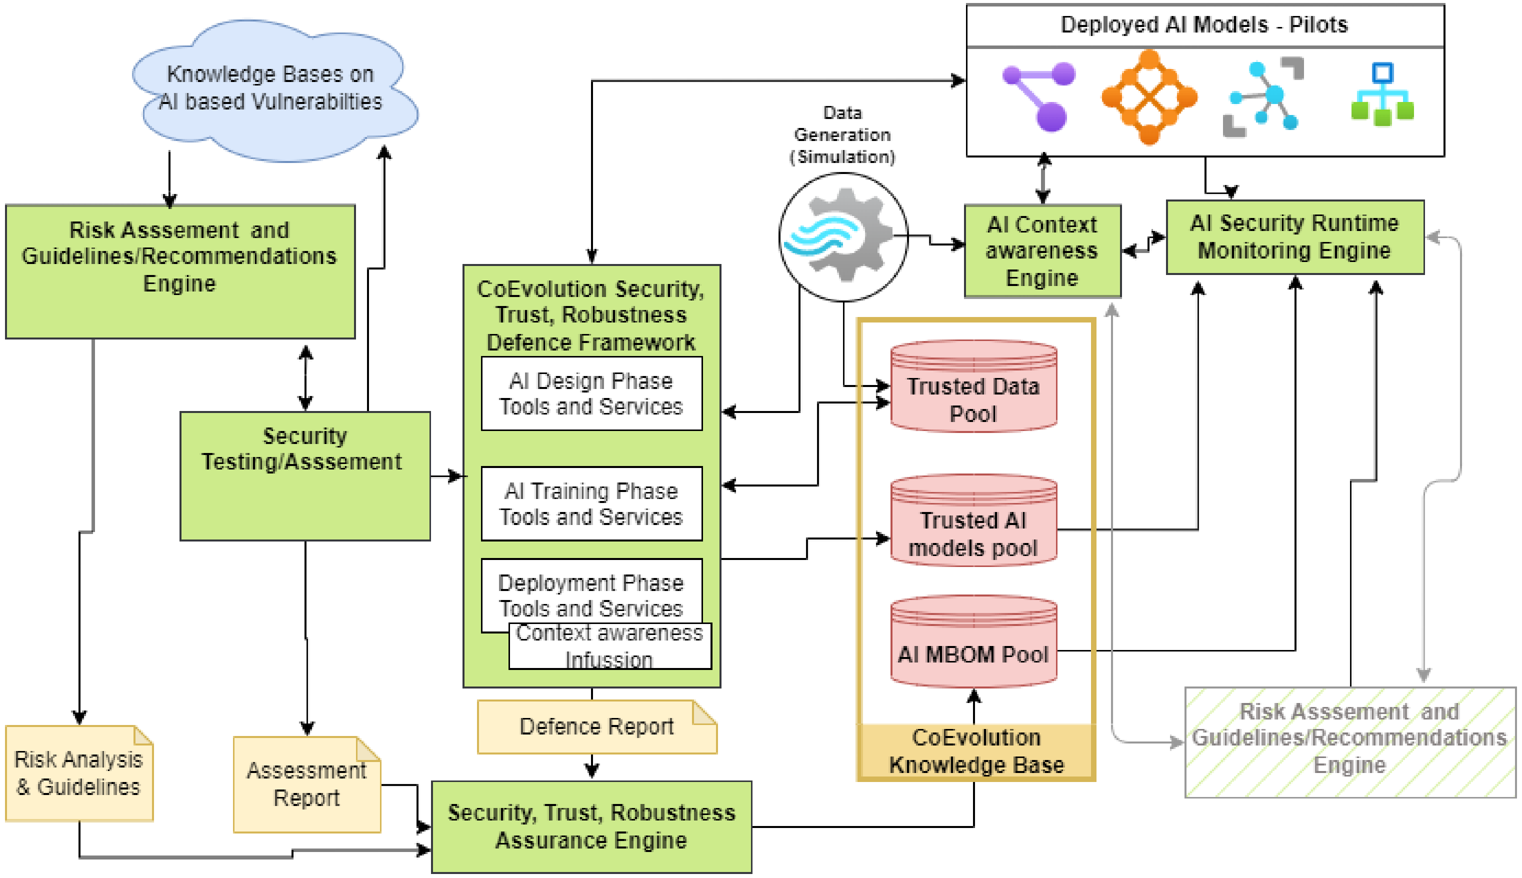

Post-deployment security and robustness through incremental and continual learning

The Exam

Midterm 4 – Poster Session Time

- Format
  - Read 1 paper on a course topic
  - From a list of provided papers
  - Prepare in group a poster using the structure and template here
    - Print it (in any copisteria, suggested format A0)
    - Hang it in room C1 (Fibonacci) on Thursday May 28th h. 11.00-14.00
    - Discuss it in front of me and Ph.D. students
- All group members are expected to contribute to the poster, to be present at the poster session and to be able to answer questions
- Groups are randomly generated: no changes, no exceptions

Final Projects (Alternative to MIDTERMS)

- Develop a tested and commented software implementing a non-trivial learning model and/or a DL-based application relevant for the course.
- Choose from a set of suggested topics or propose your own topic of interest

Final Projects – What to Deliver?

- Three things need to be delivered (by the Appello deadline)
  - A report describing the model and its validation (~10 pages, single column, standard margins)
  - The code for the project
  - A presentation on the project
- Presentation will last 15 minutes and will be given on the oral day
- Presentation tips
  - Describe the implemented model, the library and the experimental validation
  - My suggestion is to keep the number of slides around 15 (tops)

Final Exam Timeline

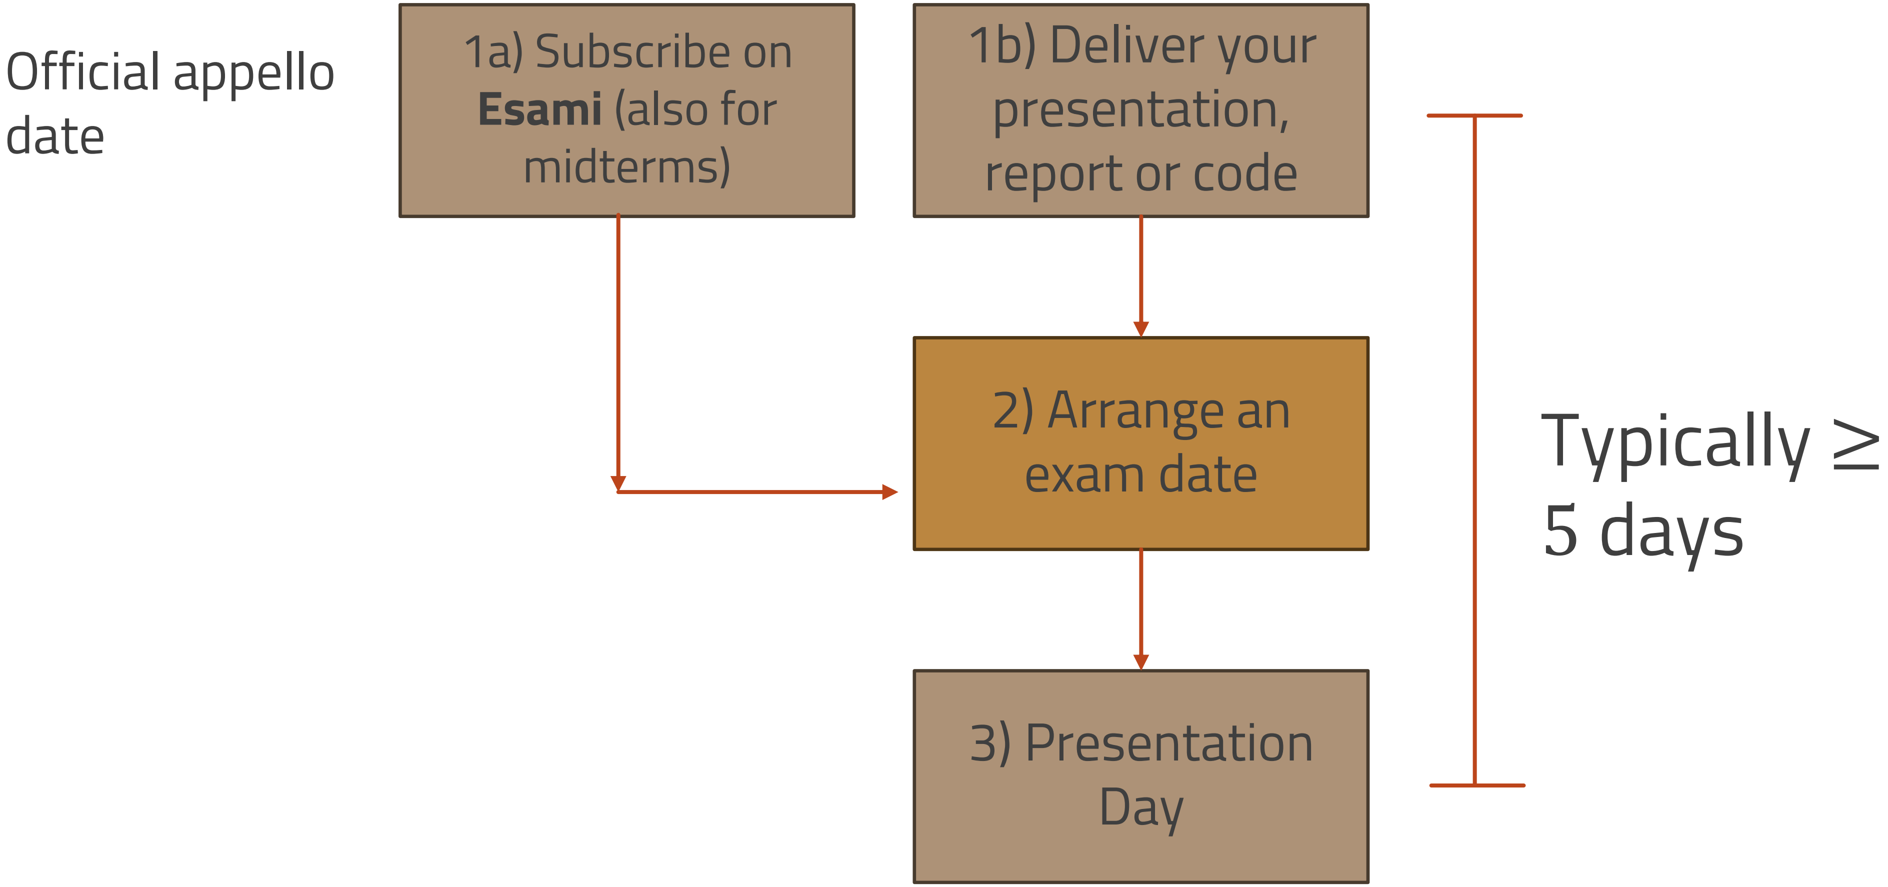

Presentation, Report & Code Delivery

Fixed and strict deadlines for submitting final project files to me

- 03/06/2026 h. 14.00
- 26/06/2026 h. 14.00
- 17/07/2026 h. 14.00

Delivery through the GDL moodle

Submit presentation, report or code in a single archive file (no data!!!!)

On the Oral-Presentation Day

- Oral exams will typically be held in my office at the Computer Science department
  - Information sent via mail to students subscribed to the appello (if with completed midterm/final projects)
  - All students are welcome to attend
- Non-midterm students will first deliver their presentation
- All students will be subject to an oral exam on models, algorithms and applications discussed during the course lectures
  - Deep learning for graph II: Advanced models lecture is not part of exam materials

Expectations about the oral exam (from Lecture 1)

Let's set our mutual expectations right and live happily ever after

- Yes, the oral exam is going to be about methods & models, also from a theoretical perspective
- Yes, I can ask you to derive statements and describe the models/methods formally
- Focus on reasoning not memorization: I will ask questions in ways that will show if you have learned things by hearth
- Structure of the oral discussion:
  - Expect roughly 3 questions of increasing level of complexity (higher complexity==higher grade)
  - If you cannot answer the current question in reasonable time (i.e. some 5-10 minutes) we conclude the oral and you get the grade related to the question(s) you have managed to answer
  - I don't do on-request random questioning until you get one you can answer
- **Minimum threshold to pass the course requires being able to formalize models seen as part of the course and being able to convincingly discuss their properties and use**

For ISPR students

- I allow you to take the exam either with
  - The “old” ISPR program and modality
  - The “new” GDL program and modality
- In either cases
  - The exam officially registered on your student profile will be ISPR
- When you subscribe to GDL on the Esami platform, remember to state in the notes that you are an ISPR student, and which program you would like to be assessed on (ISPR or GDL)
- Deadlines for material delivery and orals are the same as of GDL

FAQs

- What is the language for the report and the presentation?
  - Both need to be written in English, but the presentation can be given in either Italian or English
  - The oral exam can be in either Italian or English (your choice)
- How long do midterms last?
  - Until end of current solar year (September exams and all included)
  - Yes, I will keep them even if you give the exam and fail it (not if you fail it because of plagiarism/misbehaviour though)
- Other questions?

Wrap-up

Concluding

- I hope you have enjoyed the course
- I will be out of my office:
  - 03 June-5 June
  - 22-26 June
  - 06-10 July
  - 22-24 July

Enjoy the rest of you AI curriculum