ISPR Final Lecture

Davide Bacciu

Dipartimento di Informatica
Università di Pisa

Intelligent Systems for Pattern Recognition (ISPR)

Lecture Outline

- Course wrap-up
  - A-posteriori view of the course
  - Final take home messages
- Overview of ML research @UNIPI
  - The CI&ML group
  - The Pervasive AI Laboratory
  - Research themes and projects
  - Opportunities
- ISPR Final exam
- Conclusions & Discussion

Course Posterior Analysis

$$P(ISPR|\boldsymbol{D}) = P(ISPR)P(\boldsymbol{D}|ISPR)$$

Hypothesis

Lectures

How is your view of ISPR after having taken the course?

Your expectations before the course

Your interpretation of the lectures based on your idea of the course

Deep Learning – Any Change in Perception?

What society thinks I do

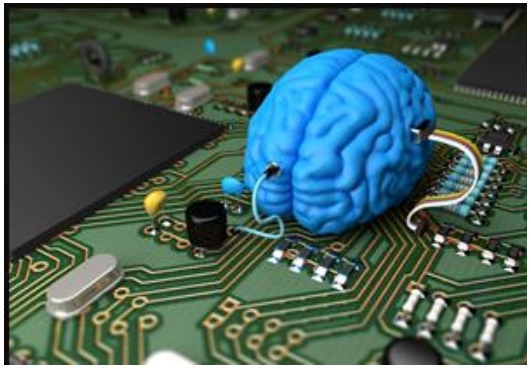

What my friends think I do

What other computer scientist think I do

What mathematicians think I do

What I think I do

```
from PyTorch import *
```

What I actually do

A Modern View on Pattern Recognition (I)

Dealing with complex data

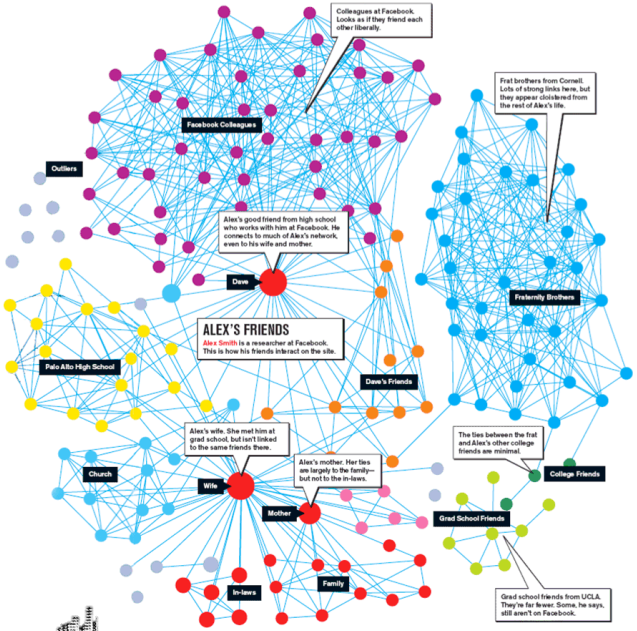

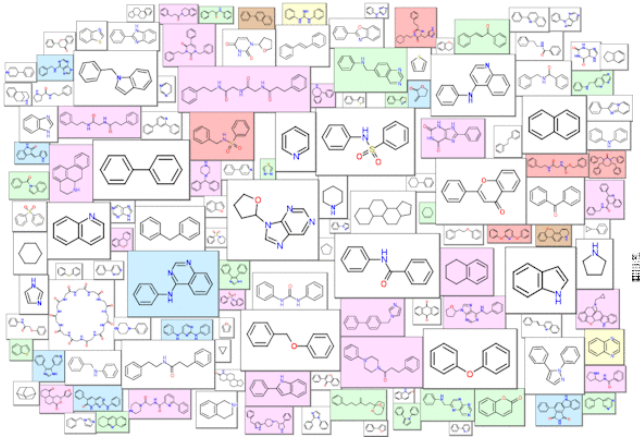

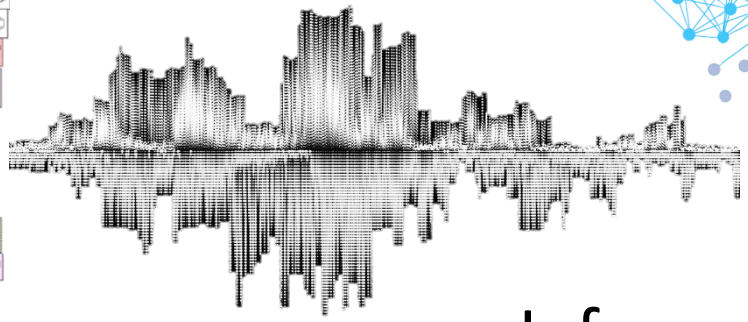

Information in context

Multimodal

Large scale

Raw and noisy

A Modern View on Pattern Recognition (II)

The goals are well past recognition

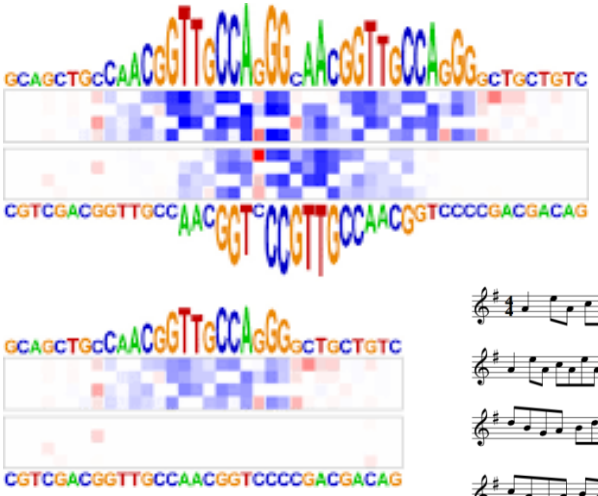

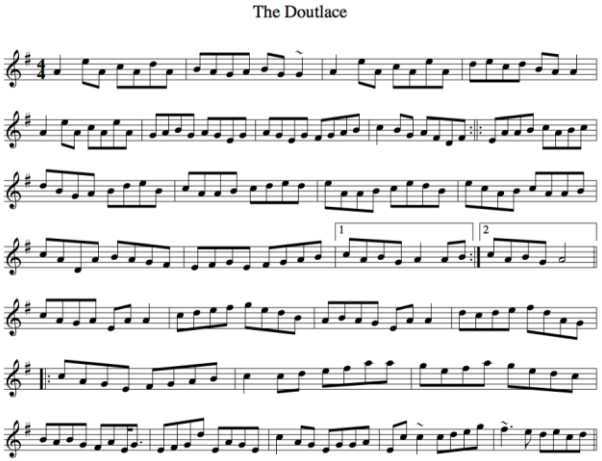

Creativity

Understanding,
reasoning and explaining

Generation

The Course in 1-Slide

- Old-school pattern recognition
  - Building blocks: convolution, filters, spectral analysis
- Generative models
  - Infer knowledge rather than just predict
  - Learn a generative process
  - Introduce prior 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

A Convergence of Paradigms

- Need the efficacy and efficiency of discriminative 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 models
  - Variational and generative DL

After Completing This Course, Hopefully…

- Know which ML models are best to start with for addressing a given PR problem
- Know what challenges your ML model will need to solve to realize a PR application
- Know a bag of tricks to modify a model to suit your needs
  - Message passing, variational approximations, sampling, latent representations, feature functions
  - Batch normalization, pretraining, end-to-end differentiability, distribution learning with NN, enhancing memory, attention

Machine Learning is the New Algorithmics

Can you derive EM for GMM?

When do you need to check you gradient?

Implement a Gaussian/edge filter

ML Research @ UNIPI

Computational Intelligence & Machine Learning Group (CIML)

https://ciml.di.unipi.it/

Faculty

Alessio Micheli (coordinator)

Davide Bacciu

Claudio Gallicchio

Vincenzo Lomonaco

Associates

Antonio Carta, Postdoc

Daniele Castellana, Postdoc

Andrea Cossu, Ph.D. Student

Valerio De Caro, Ph.D. Student

Giovanna Maria Dimitri, Postdoc

Daniele di Sarli, Ph.D. Student

Federico Errica, Ph.D. Student

Alessio Gravina, Ph.D. Student

Francesco Landolfi, Ph.D. Student

Danilo Numeroso, Ph.D. Student

Michele Resta, Ph.D. Student

Andrea Valenti, Ph.D. Student

Asma Sattar, Ph.D. Student

Domenico Tortorella, Ph.D. Student

Other Collaborators

Davide Serramazza, Research Associate

Dario Balboni, Data Science Ph.D. Student

Francesco Sansone, Ph.D. student (CNR)

Elisa Ferrari, Data Science, Postdoc

Giacomo Lanciano, Data Science Ph.D. Student

Pervasive AI Laboratory (PAILab)

A joint initiative by Dipartimento di Informatica @ UNIPI and Istituto Scienza e Tecnologia dell'Informazione @ CNR

pai.di.unipi.it pai-info@isti.cnr.it

Features

- Coordinators– Davide Bacciu (UNIPI) and Patrizio Dazzi (ISTI-CNR)
- ~40 members
- Coordinating 2 H2020 Projects and 1 KA; participation in 2 H2020 projects and 3 industrial projects
- 10M Euro secured grants

- Focus
  - AI as a ubiquitous component in ICT systems
  - Design communication and computing systems to support pervasive AI

PAILab @ UNIPI - H2020 TEACHING Laboratory

- **Focus**
  - Efficient machine learning for streaming data analysis
  - Learning from human state
- **Human-system interface**
  - High quality wearable sensors (EEG, GSR, Inertial, …)
  - EEG headset
  - Environmental microphones
  - Wide-screen display
  - SW for data collection & demoing
- **Edge computing facilities**
  - Jetson nano GPUs
  - Open-CL enabled FPGAs (Intel Arrya)
  - Gateway PC
  - SoC board for automotive-grade embedded applications
  - SW library for data stream processing (AI loads)
- **HPC facilities @ ITC**
  - 16x R74xd cluster (2x Intel XGold 6240R, 24 cores per socket, 768 GB RAM, Mellanox ConnectX-5)
  - 1x server (4x Intel XPlatinum 8260L, 768 GB RAM, 2 TB Optane, Mellanox ConnectX-5)

PAILab

PAILab @ CNR - AI@Edge Laboratory

- **Focus**
  - Edge and distributed computing for AI-based systems
  - Focus on autonomous vehicles, computer vision and 3D graphics, IoT/CPS, critical and industrial applications
  - Complementary to ICAR's AI datacenter
- **Resources**
  - Nvidia Jetson development kits
  - Intel Arria SoC FPGA
  - Integrated imaging & camera
  - 5G networking
  - Wearable and environmental sensors
  - Edge and dense servers
  - Rovers/UAV
  - 3D printing

PAILab

Research Overview

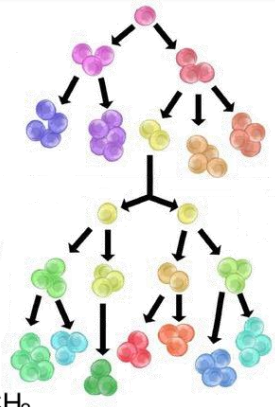

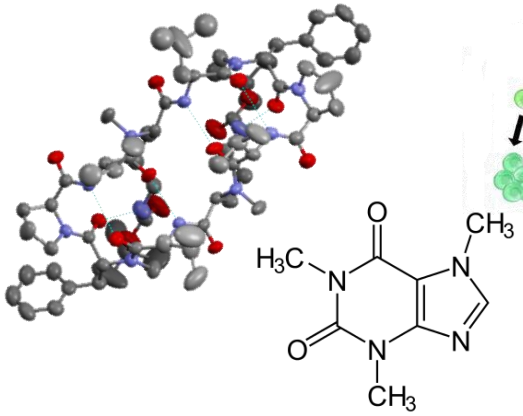

- Long standing on learning for structured data processing
- Deep learning and neural networks, probabilistic models, kernel methods

Funded by Italian, European and industrial projects

MONTE DEI PASCHI DI SIENA
BANCA DAL 1472

TEACHING

H&M

SCIENTIFIC INDEPENDENCE OF YOUNG RESEARCHERS

intel

biobeats
a HUMA company

rubicon

DELL

stm

doremi

Topics of current research (I)

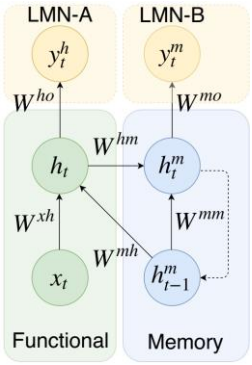

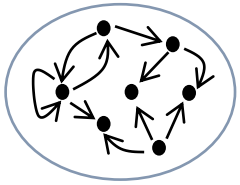

Recurrent neural networks

- (Deep) Reservoir computing and randomized networks
- Theoretical properties of neural memories
- New paradigms of dynamic memory

Learning with structured data

- Learning non-isomorph transductions
- Deep learning for graphs
- Graph generation

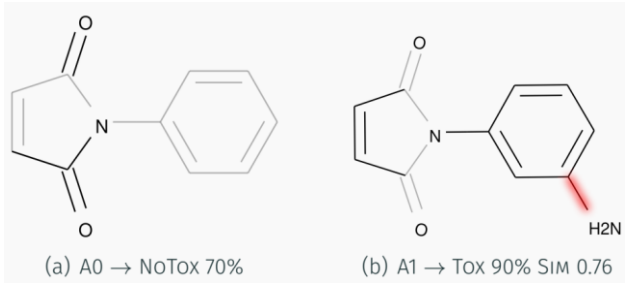

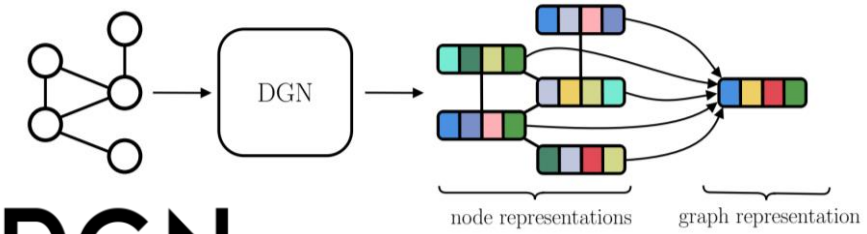

PYDGN

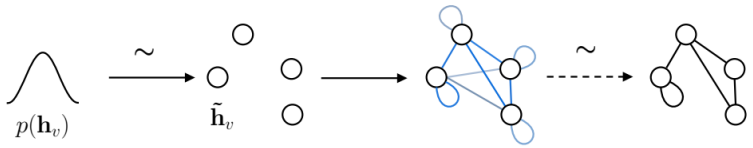

Topics of current research (II)

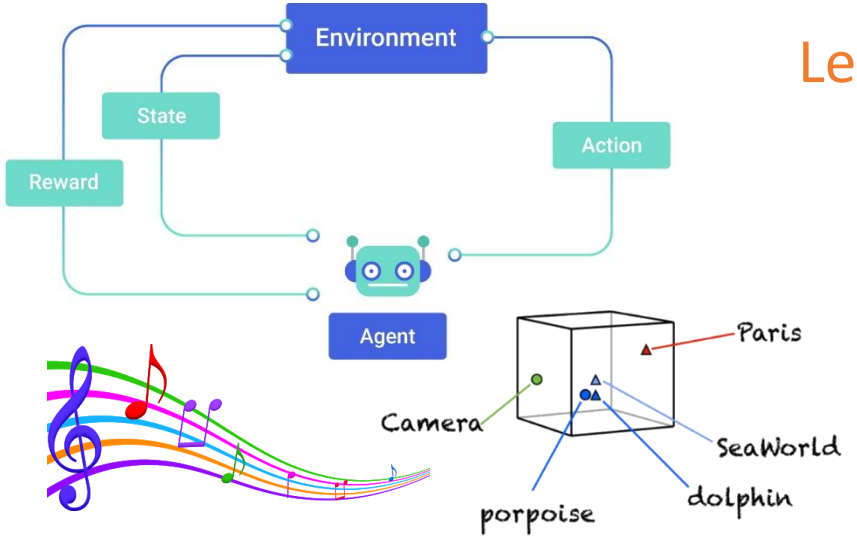

Learning Fundamentals

- Continual learning
- Federated learning
- Learning representations
- Reinforcement learning and learning under weak/self/noisy supervision
- Theoretical properties of deep NNs

Pervasive AI

- Distributed learning
- In-silico embedded intelligence
- AI on GPU/FPGA/Cluster computing
- ML as a service & ML-based application development
- Trustworthy & human-centered learning

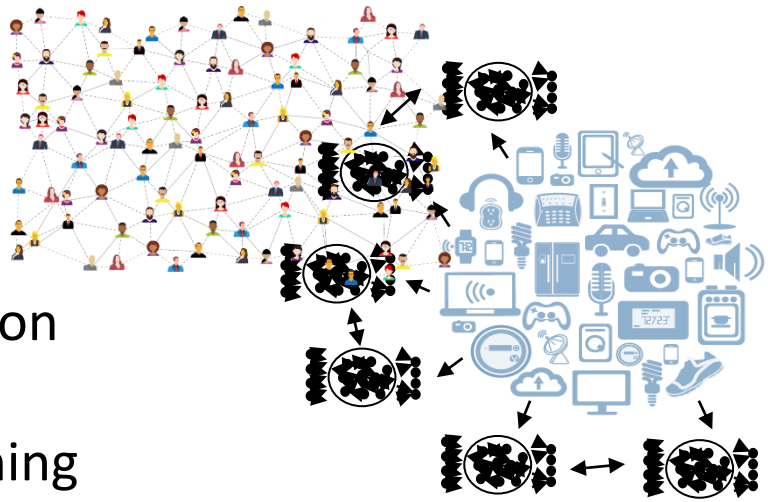

Topics of current research (III)

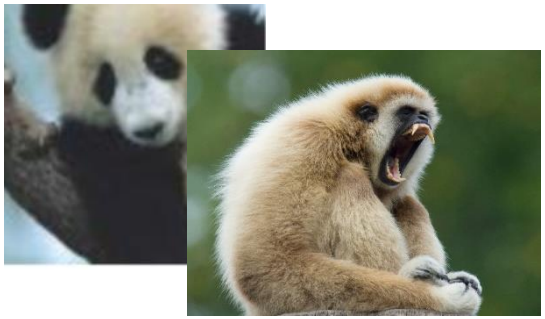

Trustworthy Neural Systems

- Protecting neural representations
- Robust, safe and secure deep learning
- Interpretability

Applications of learning systems

- Recommendation systems
- Machine vision & multimedia
- Bioinformatics
- Chemistry
- Robotics
- Autonomous Vehicles
- Physio-signal processing

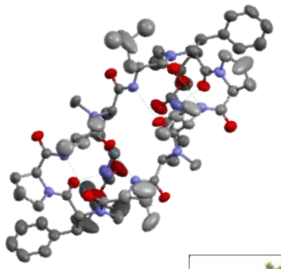

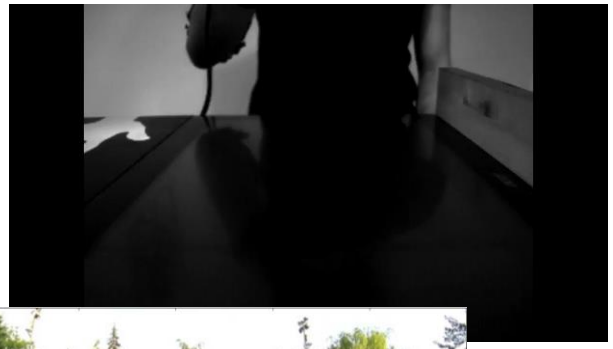

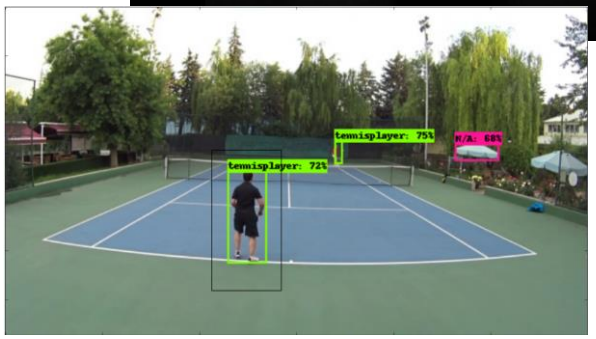

And many other things…

- Generative and unsupervised deep learning
- Computational creativity: music generation and style transfer; visual arts
- Medical imaging
- Integrating learning and reasoning
- …

Contact if interested in M.Sc. and Ph.D. projects

Active Projects & Initiatives

- EU H2020 TEACHING: A computing Toolkit for building Efficient Autonomous appliCations leveraging Humanistic INtelliGence (2020-2023)
- EU H2020 TAILOR: Trustworthy AI Integrating Learning, Optimization and Reasoning (2020-2023)
- CLAIRE: European AI Task force on COVID-19
- Intel-Dell Industrial competence centre on AI
- H&M Industrial collaboration on Deep Learning for Graphs
- Industrial research project: learning to predict critical events in sensor timeseries (2021)
- Industrial research project: Deep and Bayesian learning for perceived stress prediction (2019-2021)
- Regional project PH-REMIX: Remixing tools for multimedia data (2020-2022)
- Regional project PRE-MED2: Precision Medicine for Preventing Type 2 Diabetes: a Step Forward (2020-2023)

H2020 TEACHING (2020-2023)

A computing toolkit for building efficient autonomous applications leveraging humanistic intelligence

Short Facts
10 partners
5 countries
4M Euro

UNIPI role
Davide Bacciu (Project Coordinator)
Claudio Gallicchio (AI WP leader)

www.teaching-h2020.eu

TEACHING - Motivation & Vision

A human-centric perspective on autonomous CPSoS applications

TEACHING - Motivation & Vision

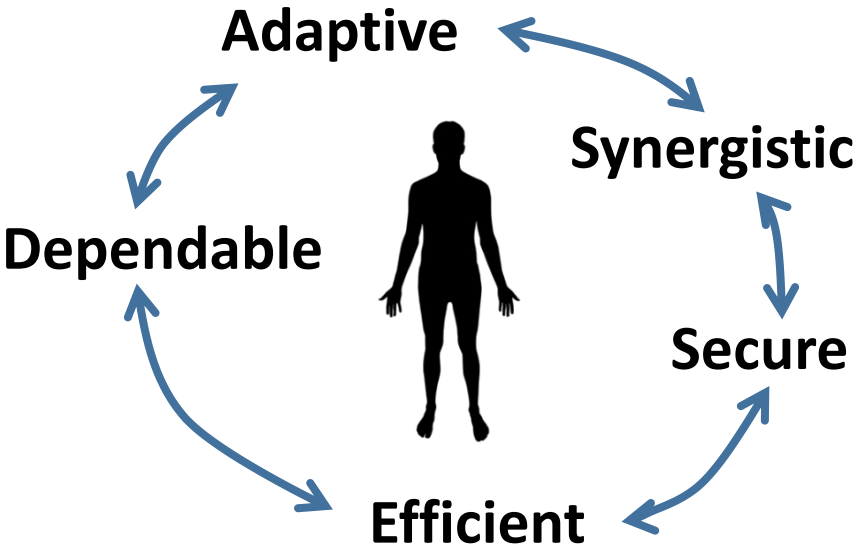

A human-centric perspective on autonomous CPSoS applications

Paradigmatic shift needing support at computing and system level

TEACHING - Conceptual Architecture & Approach

Autonomous CPSoS applications need intelligence

Human-centric AI as a distributed service leveraging PEC feedback

Mixed edge/cloud orchestration & software-level abstraction platform

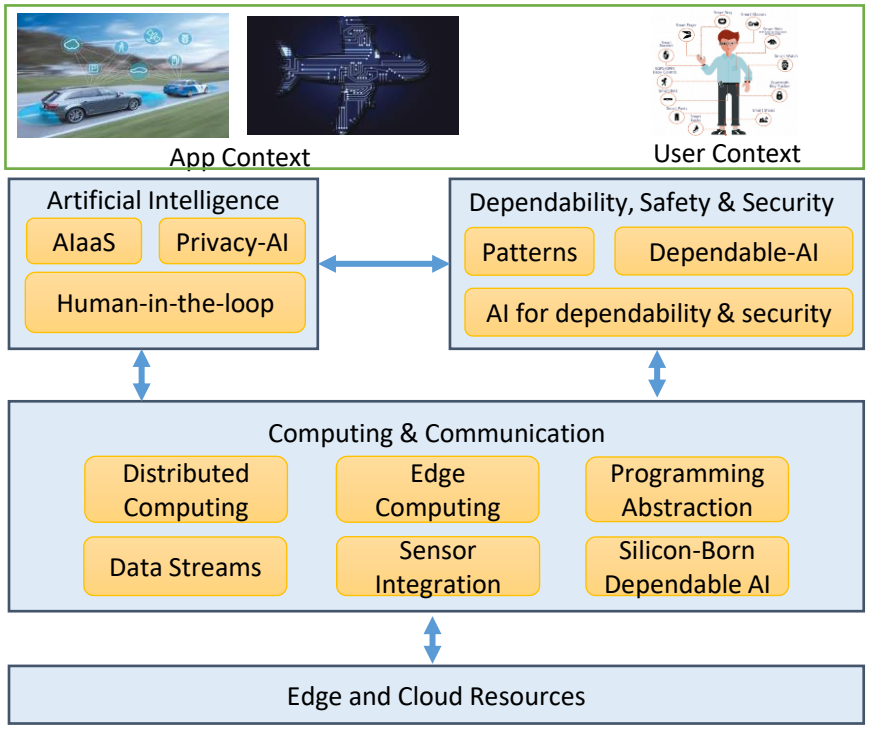

Human-state must be accounted for in autonomous system operation

Dependable system engineering for mission-critical CPSoS with AI components

Native support for AI & dependability (computing and communication)

H2020 TAILOR (2020-2023)

TAILOR

Trustworthy AI - Integrating Learning, Optimisation and Reasoning

ICT-48 Network of research excellence centres on foundations of Trustworthy AI

- 55 partners; >75 affiliated labs
- Research program to address grand challenges
- Connectivity fund for active dissemination to the larger AI community
- Network collaboration activities (exchanges, events, joint supervisions)

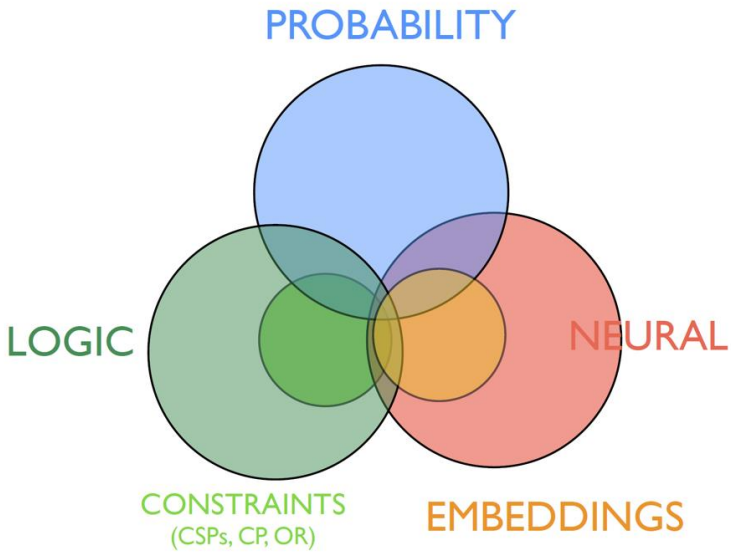

tailor-network.eu

CLAIRE AI

CLAIRE Confederation of Laboratories for Artificial Intelligence Research in Europe

claire-ai.org

- UNIPI role
  - Partner member
  - Member of CLAIRE network seed project H2020 TAILOR
  - CLAIRE national office for Italy (Attardi, Bacciu)

- CLAIRE COVID-19 Task Force
  - Coordination of Bioinformatics workgroup
  - Working on drug-repurposing, viral-hosts interaction analysis, genomics, clinical data analysis
  - Soon to be become a WG on AI for Health

covid19.claire-ai.org

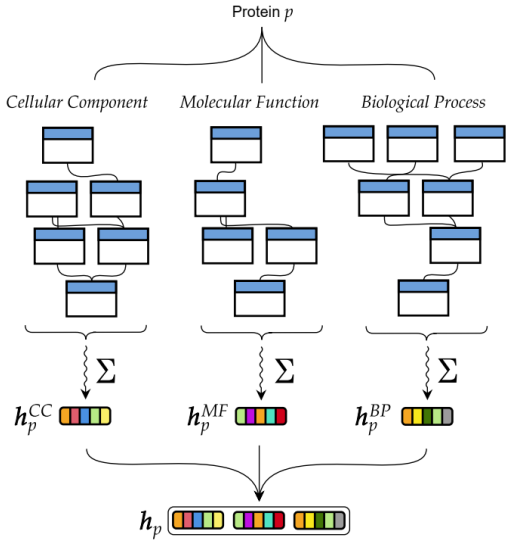

Midterm 4

- Out next week
- Format
  - Read 1 paper on a course topic
  - From a list of referenced papers
  - Prepare a 5 minutes presentation for the oral day with the following (rough) content
    - Introduction to the problem
    - Model description
    - Results
    - Comment on novelties, strong points and weaknesses
- Deliver the presentation by the Appello deadline

Final Projects (Alternative to MIDTERMS)

- List already on the course Moodle
- Survey
  - Read at least 3 relevant papers on a topic
  - Prepare a presentation: not a simple summary but find connections between the works and highlight open problems
- Software
  - Develop a software implementing a non-trivial learning model and/or an application relevant for the course
  - Prepare a presentation describing the software and its validation

Final Projects – What to Deliver?

- Two things need to be delivered (by the Appello deadline)
  - A written report on the project
  - A presentation on the project
- Presentation will last 15 minutes and will be given on the oral day
- Presentation tips
  - (Survey) Summarize the ideas, models and results
  - (Software) 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

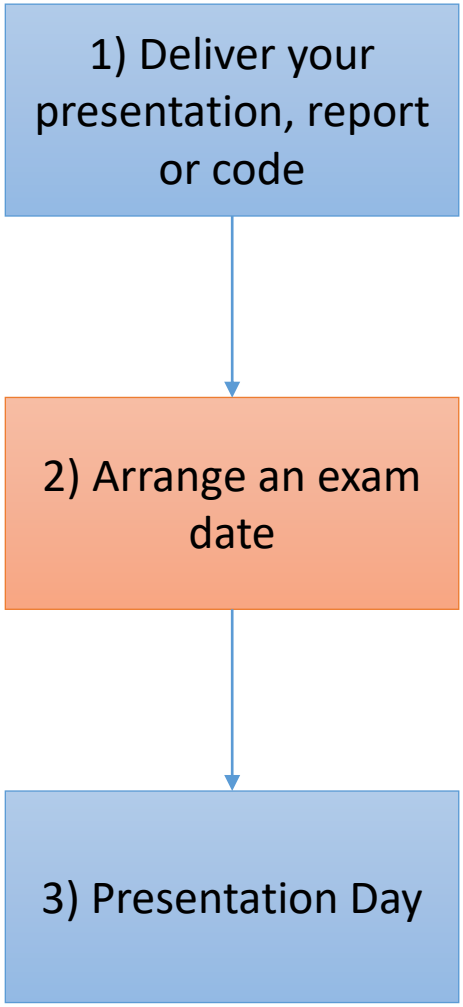

Typically >=5 days

Temporal slots for oral exams will be made available on the Esami platform

Presentation, Report & Code Delivery

Fixed and strict deadlines for handling the presentation, report and/or code to me

1. 11/06/2020 h. 18.00
2. 02/07/2020 h. 18.00
3. 22/07/2020 h. 18.00

Delivery through the ISPR moodle

- Midterm: submit presentation in PDF
- Project: submit presentation, report, code in a single archive file (no data!!!!)

On the Oral-Presentation Day

- At least until July 2021 exams will be held online on the TEAMS of the ISPR course
  - Check for specific calendar items on the Team
  - All students are welcome to attend
- I will ask questions after the presentation
  - On the content of the talk
  - On associated models
  - On other models, algorithms and applications discussed during the course lectures

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 September exams (included)
  - Yes, I will keep them even if you give the exam and fail it (not if you fail it because of plagiarism though)
- Other questions?