MARIE SKŁODOWSKA-CURIE POSTDOCTORAL FELLOWSHIPS 2025

EXPRESSION OF INTEREST FOR HOSTING MARIE CURIE FELLOWS

HOST INSTITUTION

NOVA Information Management School (NOVA IMS)

RESEARCH GROUP AND URL

https://magic.novaims.unl.pt/

SUPERVISOR (NAME AND E-MAIL)

Mauro Castelli mcastelli@novaims.unl.pt

SHORT CV OF THE SUPERVISOR

Mauro Castelli has a Ph.D. in Computer Science obtained at the Università di Milano Bicocca (Italy), a Master's in computer science, and Degree in Computer Science obtained at the Università di Milano Bicocca (Italy). He is currently a Full Professor at the Universidade Nova de Lisboa, director of the bachelor degree in Data Science, and member of the Scientific Council of NOVA Information Management School (NOVA IMS). He is also a researcher at the Information Management Research Center of this university. He participated as a principal investigator, coprincipal investigator, or work package leader in different research projects at national and international levels. He collaborated with several European universities on the application of machine learning methods for addressing complex real-world problems. It has international collaborations with researchers in more than twenty different countries and with universities recognized as leaders in the area of artificial intelligence. He is a member of ACM (Association for Computing Machinery). He is the author of more than 200 scientific publications and has presented about fifty seminars, conferences, and communications. He was awarded in 2013 and 2014, in the framework of the main European conference of Artificial Intelligence and Evolutionary Computing, for the quality and contribution of his scientific research. He is in the Stanford list of the the 2% of the world's most cited researchers. His research focuses on the following areas: Deep Learning, Large Language Models, Machine Learning, Evolutionary Computation. He was the supervisor of more than 140 master theses and nine Ph.D. theses in the field of Machine Learning and Artificial Intelligence.

5 SELECTED PUBLICATIONS

- Perezhohin, Y., Peres, F., & Castelli, M. (2024). Combining computational linguistics with sentence embedding to create a zero-shot NLIDB. Array, 24, 100368.
- Perezhohin, Y., Santos, T., Costa, V., Peres, F., & Castelli, M. (2024). Enhancing Automatic Speech Recognition: Effects of Semantic Audio Filtering on Models Performance. IEEE Access.
- Marchetti, F., Pietropolli, G., Verdù, F. J. C., Castelli, M., & Minisci, E. (2024). Automatic design of interpretable control laws through parametrized Genetic Programming with adjoint state method gradient evaluation. Applied Soft Computing, 111654.
- Santos, F. J., Gonçalves, I., & Castelli, M. (2023). Neuroevolution with box mutation: An adaptive and modular framework for evolving deep neural networks. Applied Soft Computing, 147, 110767.
- Philippi, D., Rothaus, K., & Castelli, M. (2023). A vision transformer architecture for the automated segmentation of retinal lesions in spectral domain optical coherence tomography images. Scientific Reports, 13(1), 517.

PROJECT TITLE AND SHORT DESCRIPTION

(1) **Unified Deep Learning and LLM Frameworks for Reasoning over Structured and Unstructured Data.** This project proposes a hybrid architecture combining deep neural networks and pretrained LLMs to jointly reason over tabular, textual, and visual data. The research will develop multimodal

embeddings that preserve structure-awareness while retaining language generalization capabilities. The fellow will focus on training regimes that enable cross-domain reasoning and few-shot generalization. This project addresses a key open issue: how to achieve integrated understanding across heterogeneous data types.

(2) **Large Language Models for Health: Trustworthy Clinical Decision Support.** This project investigates the deployment of LLMs in clinical contexts, focusing on generating and validating patient-specific recommendations from medical records and multimodal data. Key research questions include bias mitigation, factual accuracy, and the interpretability of generated content. The fellow will develop fine-tuning strategies incorporating domain knowledge and human feedback. Special attention will be given to explainable summarization and data privacy. The outcome will be a framework for certified LLM integration in digital health workflows.

(3) **Efficient Modular Architectures for Scalable Large Language Models.** This research will explore novel modular and sparsely-activated transformer architectures to address the computational bottlenecks of large-scale LLM training and inference. The fellow will design models that adaptively scale capacity depending on task complexity. Emphasis will be placed on energy-efficient training and inference. The project addresses the pressing need for resource-conscious LLM design, especially in decentralized or low-resource environments.

(4) **Causality-Aware Interpretable AI for Decision-Making in High-Stakes Domains.** This project aims to develop interpretable AI systems grounded in causal inference, enabling models to provide counterfactual explanations for their predictions. It will investigate how to disentangle spurious correlations from meaningful causal structures in domains such as finance and law. The fellow will explore hybrid models combining symbolic reasoning with deep learning. A key challenge is to ensure transparency without sacrificing accuracy or scalability.

(5) **Neuroevolutionary Architectures for Explainable Reinforcement Learning Agents.** The aim is to integrate neuroevolutionary algorithms with intrinsic explainability constraints to evolve deep policies that are transparent and verifiable. The fellow will explore genotype-to-phenotype mappings that favor compact, interpretable representations and employ novelty-driven search strategies to enhance generalization. Applications include robotics and autonomous systems. A core innovation is leveraging evolutionary diversity to discover naturally interpretable decision modules.

SCIENTIFIC AREA WHERE THE PROJECT FITS BEST*

ENG

***Scientific Area where the project fits best** – Please select/indicate the scientific area according to the panel evaluation areas: Chemistry (CHE) • Social Sciences and Humanities (SOC) • Economic Sciences (ECO) • Information Science and Engineering (ENG) • Environment and Geosciences (ENV) • Life Sciences (LIF) • Mathematics (MAT) • Physics (PHY)