Wednesday January 21st, 2026

Department of Civil and Environmental Engineering

Castigliano room

Toward Mechano Intelligence

9.30 – 9.45: Introduction

Alberto Corigliano

9.45 – 10.15: ***Sequential Planning and Spatial Navigation Through Active Inference***

Giovanni Pezzulo - Institute of Cognitive Sciences and Technologies – CNR Rome

Humans and other animals can form sequential plans, whether selecting a series of destinations during travel or assembling the steps needed to achieve a goal. In this talk, I will present studies that use the active inference framework to advance our understanding of how such sequential plans are formed and deployed in spatial navigation and other cognitive tasks that require flexible, multi-step behavior.

I will begin with simple simulations of sequential planning, illustrating how active inference accommodates both pragmatic drives such as maximizing expected utility and epistemic drives aimed at reducing uncertainty. These examples will also highlight how the underlying inferential mechanisms may correspond to neural computations. I will then turn to more complex scenarios that require integrating higher-level plans (for example, sequences of navigation goals and subgoals) with lower-level action policies (for example, step-by-step behaviors needed to reach each subgoal). To address these challenges, I will present active inference models grounded in hierarchical generative architectures that connect cognitive maps across multiple levels of abstraction, spanning both abstract "task space" and concrete "physical space," and offering insights into interactions between frontal and hippocampal systems.

By bringing together computational modeling and neurobiological evidence, this work sheds light on how the brain constructs and executes flexible, goal-directed plans.

Giovanni Pezzulo

Received a Ph.D. in Cognitive Psychology from the University of Rome ''La Sapienza''. He is Research Director at the ISTC-CNR, Rome. He uses a combination of theoretical, computational and empirical methods to study cognitive processing in humans and other animals, with a focus on goal-directed behaviour, decision-making, prediction, and planning. His research was funded by the ERC, the EU's Horizon 2020, the HFSP, and other agencies.

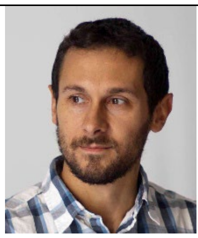

10.15 – 10.45: ***Emergent collective cognition in multi-agent systems: an integrated approach combining active inference and information theory***

Domenico Maisto - Institute of Cognitive Sciences and Technologies – CNR Rome

From animal groups to human societies and robotic swarms, collective behaviour is a hallmark of multi-agent systems. A key unresolved question is how such collectives develop functional capacities—such as joint agency or distributed knowledge—that exceed those of their constituent parts.

This presentation addresses this topic by illustrating how to model ensembles of agents that, governed by active inference principles, interact reciprocally, and how to detect and analyse emergent collective cognitive features using information-theoretic measures.

Two cases will be presented. In the first case, I will show a flocking system in which individual agents (birds) minimize free energy while engaging in local sensory coupling. By analysing the formation of Markov blankets, we can examine how the flock emerges as a higher-order agent with distinct internal, sensory, and active states. When subjected to external perturbations (e.g., a predator), the flock not only reacts more rapidly than individual agents but also encodes synergistic information about the threat—information that is not reducible to any single agent's knowledge.

The second model extends this framework to ensembles of self-propelled particles on a lattice that self-organise via a shared scalar field. Through leader-follower differentiation and field-based interactions, the system exhibits controlled aggregation and adaptive responses to external field traps, illustrating how local inference can generate global patterns without centralised control and steer group coordination.

Together, these studies provide a principled, information-theoretic account of how distributed intelligence arises from decentralized active inference. They offer a unified framework for studying the rise of collective cognition and multi-scale autonomy in biological and engineered systems.

Domenico Maisto
(M.Sc. in Physics, Ph.D. in Multiscale Modeling, Computational Simulations and Characterization in Material and Life Sciences) Has research interests in cybernetics, artificial intelligence, and machine learning, with a focus on computational neuroscience and cognitive science.
Currently, he is a senior researcher at the Institute of Cognitive Sciences and Technologies (ISTC) of the Italian National Research Council.

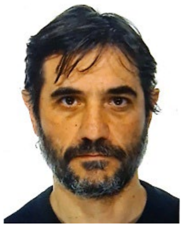

10.45 – 11.10: Coffee break

11.10 – 11.40: ***On the Structure of Compositional Intelligence: Hierarchical Representations and the Reuse of Neural Computation***

Francesco Donnarumma - Institute of Cognitive Sciences and Technologies – CNR Rome

A central set of open problems in neuroscience and artificial intelligence concerns how complex behavior is represented, learned, and flexibly recomposed across tasks and time scales in the brain. Despite major advances in machine learning and neural data analysis, current models still struggle with continual learning without *catastrophic forgetting*, compositional generalization, and the principled reuse of previously acquired knowledge.

This talk offers some hints and possible perspectives on how these challenges might be addressed from a distributed and representational viewpoint. It touches on questions concerning how high-level cognitive variables—such as goals, task structure, and abstract rules—may be abstracted and encoded in neural populations. Converging evidence suggests that such variables are expressed as low-dimensional neural manifolds embedded within high-dimensional activity spaces, and that some of the principles governing their formation, stability, and transformation across tasks and learning can be envisaged.

Open questions persist regarding how these latent structures are acquired, linked across contexts, and dynamically reconfigured during ongoing behavior, allowing for flexibility and reuse. The brain appears to repurpose the same neural resources across multiple functions, raising the question of how fixed neural substrates can support different computations depending on context. From this perspective, models of neural *programmability* — enabled by modulatory signals or auxiliary control variables rather than synaptic change — allow us to distinguish online inference from long-term learning, offering a useful conceptual lens.

Finally, the talk focuses on how some of these issues may be addressed within hierarchical active inference, proposing a compositional model in which each hierarchical level constitutes a complete generative model operating under the same variational free-energy objective.

The representations of latent variables at each level are recursively defined as compositions of lower-level sequences, with interactions mediated by probabilistic inference linking predictions and evidence across scales, thereby providing an elegant framework for a network performing abstraction and goal composition.

Francesco Donnarumma
Researcher at ISTC-CNR, in Rome. MSc in Physics (Cybernetics) at the University of Naples Federico II and Ph.D. in Computer and Information Science at the Department of Mathematics "R. Caccioppoli" of the University of Naples Federico II. Research activities focus on computational modelling of cognitive functions related to Social Interaction, Joint Action, Planning, and Decision Making in Spatial Navigation, Action Prediction, and Sensorimotor Communication. My models are mainly based on Dynamical Bayesian Networks and/or Dynamic Neural Networks. I also developed a "programmable" neural network architecture endowed with the concept of programmability that enables neural circuits with a fixed structure to exhibit multiple behaviors and "fast" switching among them.

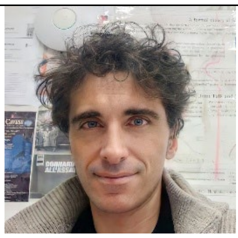

11.40 – 12.00: ***Recent Deep learning & reduced order modeling for simulation, discovery and control of large-scale systems***

Andrea Manzoni – MOX, Department of Mathematics – Politecnico di Milano

Among several recently proposed data-driven Reduced Order Models (ROMs), deep learning-based ROMs (DL-ROMs) and shallow recurrent decoder networks (SHRED) have proved to be successful strategies whenever aiming to construct non-intrusive, highly accurate surrogates for the real time simulation and control of large-scale dynamical systems. Key to both approaches is either the generation of a latent space, or the use of sensor measurements, in order to build the candidate solution, as a function of parameters and time, across a range of multiple scenarios.

In this talk I will provide a synthetic overview on DL-ROMs and SHRED, as well as on some recent extensions useful to deal with the control of distributed systems – such as those arising from fluid flows or mass transport – including (i) latent feedback control, (ii) low-data regimes, and (iii) multi-agent reinforcement learning.

Andrea Manzoni
Associate Professor of Numerical Analysis, he obtained a PhD in Mathematics from the Ecole Polytechnique Fédérale de Lausanne (EPFL) in 2012. Before joining PoliMI as a tenure-track Researcher in 2017, he was PostDoc Fellow at SISSA, Trieste and Researcher at EPFL (2014-2017). His research activity focuses on scientific machine learning and reduced order modeling for the simulation and control of large-scale systems across different Engineering fields. He is currently PI of a FIS (Italian Science Fund) project funded with a starting grant in 2023.

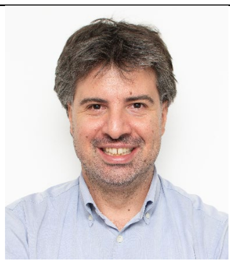

12.00 – 12.20: ***Active Inference for Perception, Action, and Learning in Mechanical Systems***

Matteo Torzoni - Department of Civil and Environmental Engineering – Politecnico di Milano

Active inference is a neuroscience-inspired Bayesian framework for decision-making under uncertainty. It relies on a generative model that interacts with a partially observable dynamical environment. I will present examples of active inference agents capable of learning and adapting in mechanical systems.

The first application addresses scenarios in which the generative model is initially unknown and must be learned online from experience. The resulting behavior combines goal-directed, information-seeking, and curiosity-driven components, enabling adaptation to novelty without catastrophic forgetting while completing a mechanical task.

The second application focuses on an active digital twin framework for structural health monitoring and predictive maintenance. The setup enables the joint optimization of structural health and maintenance costs while maintaining synchronization between the physical asset and its digital twin by actively acquiring new data when uncertainty about key system features increases.

Matteo Torzoni

I obtained a PhD in Structural, Seismic, and Geotechnical Engineering in 2024 from Politecnico di Milano, where I am currently a young researcher (RTDa) at the Department of Civil and Environmental Engineering.

My research focuses on data assimilation, prediction, and planning for digital twins of engineering systems, with an emphasis on endowing mechanical systems with learning and adaptive capabilities.

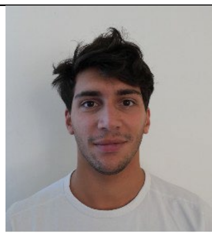

12.20 – 12.30: **Conclusion**

Alberto Corigliano

An initiative of the ERC AdG project IMMENSE

ERC – Immense