SSM
Scuola Superiore Meridionale

MERC PhD Project Proposal 2024/2025

Title of the research project:

Learning, Planning and Control via Biologically Plausible Neural Computation

Keywords (up to five)

Autonomous agents, Biologically Plausible Neural Networks, Nonlinear Dynamics, Optimization, Decision-making

Supervisors (at least two from two different areas):

Supervisor 1: *Giovanni Russo (giovarusso@unisa.it)*
Dept. of Information and Electrical Engineering & Applied Mathematics (DIEM) at University of Salerno
https://www.sites.google.com/view/giovanni-russo
Reinforcement Learning, Nonlinear/Optimal Control, Complex Systems, Optimization

Supervisor 2: *Francesco Bullo (bullo@ucsb.edu)*
Center for Control, Dynamical Systems, and Computation, UC Santa Barbara, CA 93106 USA
https://fbullo.github.io
Complex Systems, Contraction Theory, Nonlinear Dynamics, Neural Networks

Project description (max 5000 characters)

The vision and the context of the project. We are still better than machines at solving complex tasks that demand learning, planning and adapting in challenging – nonlinear, nonstationary and stochastic – environments. Reinforcement Learning agents have made headlines: outpacing human drivers in Gran Turismo, mastering Atari games, controlling fusion plasmas, winning drone races against human champions. These are all remarkable achievements. And yet – these agents are fragile. Their skills often collapse even under slight environmental perturbations. A change in lighting, a shift in object color, or an unexpected disturbance can make the agents fail – and this can be catastrophic when the agent is interacting with a physical environment. In contrast, natural agents – with little or no training – can *navigate* complexity, cope with uncertainty, and adapt robustly. The common consensus among neuroscientists and cognitive scientists is that we owe this ability not only to learning and experience, but to something deeper: our biological neural networks. Yet, we do not know to design these network.

Given this analysis, the **overarching goal** of this project is to uncover the computational principles that make biological neural networks robust, adaptive, and efficient, and to apply these insights to the design of novel agents capable of resilient behaviors where they learn, plan and compute optimal policies despite complex, dynamic, and uncertain – potentially adversarial – environments. Unifying concepts from neuroscience, machine learning, optimization, dynamical systems, statistical physics and control theory, our technical approach is threefold:

(i) propose a general normative framework to translate learning, planning and control problems into dynamical systems that admit a neural circuit representation. State variables are neuronal activities and neurons are connected through a complex web of connections that evolve over time based on context, environment and the agent teleological behavior;

(ii) characterize key dynamical properties, such as positivity and convergence, of the resulting NN that justify biological plausibility. Salient features of the analysis involve devising dynamical mechanisms enabling the network synapses to autonomously re-wire and re-tune its weights to adapt to different operating conditions;

(iii) leverage the brand-new CoRE (Complexity and Risk Engineering) Lab at the Scuola Superiore Meridionale to bring these networks to life via (a) implementing the circuits on, e.g., electric boards, (b) integrating the networks onto robotic platforms for (multi-agent) robot navigation.

The results of this highly interdisciplinary project will be transversal to several communities, including neuroscience and cognitive science, systems and control engineering, and machine learning. The project also builds on a solid scientific foundation, including peer-reviewed contributions that have been recognized with a best paper award in the field (see references in bold for works of which the proposers are authors).

Project objectives and methodology. The detailed research program of this ambitious project will be shaped based on the **interests of students**. A preliminary list of concrete research objectives is given below.

O1 – *Normative Framework For Planning, Learning And Control Via Dynamical Systems***.** This objective tackles the development of the normative framework to systematically translate planning, learning and control problems into dynamical systems that have a neural interpretation. As a first step, given the motivation of the project, the starting point is distilling neuroscientists and cognitive scientists theories of planning, learning and control in animals into an optimal control formalism. The thousand brains theory, free energy principle and the Nobel-winning fast slow thinking are suitable theories to this aim: the aim to explain planning, learning and control lead to elegant control formulations. The second step is to obtain dynamical systems that address the corresponding optimization problems of interest. By exploiting tools from monotone operators and proximal gradients we will obtain dynamical systems with the desirable feature of having among the equilibria the optimal solution of the optimization problem. Finally, exploiting also the recent work by the proposers, neural interpretations for these dynamical systems will be obtained. **In summary**, O1 outputs biologically plausible computational models for planning, learning and control.

O2 – Resolution Engine. Given the framework from O1, we will investigate the salient properties that make a neural network biologically plausible: resilience, convergence to the equilibrium/optimal solution, robustness to disturbances, adaptability to context. To this aim, contraction theory – recently exploited by the proposers – can serve as a tool to guarantee, with a single condition all these properties. Work in this objective is twofold. First, the NNs of interest for the project will be studied by leveraging, e.g. contraction theory – should the networks not be contracting we will exploit other tools from e.g., Lyapunov theory. Second, to guarantee adaptability and resilience, we distil dynamic rules enabling the network to re-wire and tune its weights. This allows the network not only to e resilient against, e.g., malfunctioning neurons/synapses, but also to process heterogeneous streams of data based on the context. A particular emphasis will be given in investigating if known features of biological NNs, such as plasticity, emerge from this re-wiring mechanism. As a by-product, we also aim at investigating if our biologically plausible NNs would exhibit the ability, elusive for artificial NNs, to *forget* information that becomes irrelevant. In summary, O2 provides the dynamic engine that makes the

NN models from O1 truly able to solve complex learning, planning and control tasks in unknown and non-stationary environments.

O3 – Embodiment. With this objective, we aim at embodying the neural-synaptic dynamics from O1 and O2 into a concrete *object* that is able to process heterogenous streams of data and reliably return optimal plans and policies. Emphasis is given to effectively deploying the NNs onto hardware with limited power/computational capabilities, such as Raspberry PIs, Arduinos, electric boards, which will be available within the recently established CoRE Lab. We also aim at investigating if the architectures from O1 and O2 would resemble those of known connectomes (such as the C. Elegans connectome that has been fully mapped) from simple organisms. We expect that this analysis will help us to both map the functionalities arising from the NNs of this project into specific biological circuits and, if needed, to simply our NNs making it possible to embody them into real hardware. **In summary**, O3 brings the results from O1 and O2 to life using in-house equipment and thus providing hands-on experience, translating theory to real-world systems.

O4 – *Experiments*. Through the CoRE Lab, we will validate the results onto challenging and exciting applications. While the applications can be tailored towards the students' interests, also taking into account opportunity and time, the plan is to leverage the NNs obtained within this project to enable a network of multi-agent autonomous vehicles to navigate in noisy, challenging, city-like scenario. Vehicles will be equipped with our NNs and the experiments will verify if these networks allow the robots to learn, plan and navigate in an environment that, by its very definition, is nonlinear, non-stationary and stochastic. At the CoRE Lab, a choice of autonomous robots will be available (such as Duckiebots, Turtlebots). We expect experimentations to be a salient feature of this research that integrates with the methodological advances. The successful student, supported by supervisors, will establish an iterative process between theory and applications, where application aspects inform the development of the methodological tools. **In summary**, O4 brings the results of the project to an extremely timely – and relevant – application that is exemplary of real-world situations where agents need to interact with non-stationary, stochastic and nonlinear environments.

Workplan. The project will be developed in incremental tasks and periodic meetings will be scheduled with the supervisors. First, the student will start with becoming familiar with the existing literature in the areas related to the project. The output of this first step will be the definition of a preliminary methodology that has the potential to tackle the above objectives. Then, in the second phase, the student will develop the methodology and deploy/test the theory on small-scale problems. The final part of the project will see the student applying the methods and tools developed to a selection of real-world applications (see above) that are relevant to MERC.

See the list of references for further details on the different aspects of the project.

Relevance to the MERC PhD Program (max 2000 characters)

The project is, by its very definition, highly interdisciplinary. The ambition of the project is to disrupt the way NNs are designed and analysed. Our aims are clearly related to modern intelligent complex systems. The project also has a link with risk, not only through the application, but also through the fact that the embodied networks will need to safely operate in environments that are noisy.

See the list of references for further details on the different aspects of the project.

Key references

B. Lake, T. Ullman, J. Tenenbaum and S. Gershman, "Building machines that learn and think like people," Behavioral and brain sciences, vol. E253, pp. 1-72, 2017

C. Hoffmann, "Is AI intelligent? An assessment of artificial intelligence, 70 years after Turing," Technology in Society, vol. in press, pp. 1-38, 2022.

J. Bruineberg and E. Rietveld, “What's inside your head once you've figured out what you head's inside of,” Ecological pyschology, vol. 31, pp. 198-217, 2019.

D. Hall and Z. Altun, C. Elegans Atlas, Cold Spring Harbor Laboratory Press, 2008.

V. Centorrino, F. Bullo, G. Russo, "Modelling and Contractivity of Neural-Synaptic Networks with Hebbian Learning", Automatica, vol. 164, 111636, 2024

V. Centorrino, A. Gokhale, A. Davydov, G. Russo, F. Bullo, "Positive Competitive Networks for Sparse Reconstruction", Neural Computation, vol. 36, pp. 1163-1197, 2024

V. Centorrino, A. Gokhale, A. Davydov, G. Russo, F. Bullo, "Euclidean Contractivity of Neural Networks with Symmetric Weights", IEEE Control Systems Letters, vol. 7, pp. 1724-1729, 2023 (recipient of the 2024 IEEE Letters on Control Systems Outstanding Paper Award)

V. Centorrino, A. Gokhale, A. Davydov, G. Russo, F. Bullo, "Contracting Dynamics for Time-Varying Convex Optimization", IEEE Transactions on Automatic Control, 2023, Cond. Accepted.

E. Garrabé, H. Jesawada, C. Del Vecchio, G. Russo, "On Convex Data-Driven Inverse Optimal Control for Nonlinear Stochastic Non-stationary Systems", Automatica, vol. 173, 112015, 2025

F. Rossi, E. Garrabé, G. Russo, "Neo-FREE: Policy Composition Through Thousand Brains And Free Energy Optimization", 2025

A. Shafiei, H. Jesawada, K. Friston, G. Russo, "Robust Decision-Making Via Free Energy Minimization", 2025, Submitted to Nature Communications

A. Tversky, D. Kahneman “Judgment under Uncertainty: Heuristics and Biases: Biases in judgments reveal some heuristics of thinking under uncertainty,” Science, vol. 185, pp. 1124-1131, 1974.

J. Hawkins, A thousand brains: a new theory of intelligence, Basic Books, 2021.

D. Kahneman, Thinking, fast and slow, Penguin Books, 2011.

V.B. Mountcastle, “An organizing principle for cerebral function: the unit module and the distributed system,” The mindful brain, 1978.

V.B. Mountcastle, B.C. Motter, M.A. Steinmetz, A.K. Sestokas, “Common and differential effects of attentive fixation on the excitability of parietal and prestriate cortical visual neurons in the macaque monkey,” Journal of Neuroscience, vol. 7, pp. 2239-2255, 1987.

Joint supervision arrangements

Meetings will be scheduled on an as-needed basis, in order to ensure the effective development of the project. As a minimum, supervisor(s) will meet students at least once a week.

Location and length of the study period abroad (min 12 months)

The candidate will be able to spend a research period (or research periods) at Prof. Bullo's lab at the University of California, Santa Barbara. Moreover, while the expertise to carry-on the experiments is available at the CoRE Lab, if time and opportunity allows, the student will have the opportunity to spend research periods with a more application-oriented in other leading institutions, such as within Prof. Shorten's lab at Imperial College London (UK).

Any other useful information

The project is best suited for ***ambitious students*** *with a preference towards mathematical (proof-oriented) mindset, and with a background in nonlinear dynamical systems. For further details on the background students can contact their potential supervisors.*