Towards Efficient Reinforcement Learning by Non-Reversible Exploration
(Funded Master internship and PhD via selective funding)

Motivation: Animals explore their environments through embodied interactions, creating intrinsically correlated sequences of experiences. These correlations reflect the constraints imposed by their bodies, senses, and the continuous nature of space and time. In neuroscience, understanding how biological systems navigate and learn from such correlated trajectories is key to unraveling mechanisms of motor control and decision-making. Similarly, in robotics and machine learning, correlated experiences pose challenges for efficient exploration and learning, as traditional reinforcement learning frameworks often rely on independent and identically distributed data.

Goal: This interdisciplinary internship will focus on developing and analyzing non-reversible stochastic processes [1, 2] as a framework for efficient exploration and learning in high-dimensional action spaces [3, 4]. By investigating how these dynamics can decorrelate agent experiences and maximize robust action sampling, the project seeks to uncover novel computational principles for optimizing exploration and decision-making. This work will integrate numerical simulations of embodied systems with analytical approaches to study quantitative aspects of these processes, such as convergence rates and exploration efficiency. Later applications, if time permits, include understanding motor control and reinforcement learning in biological systems and designing algorithms for robotics that draw on these principles.

Environment: The intern will join an interdisciplinary research collaboration bridging computational neuroscience and mathematics. Supervised by Alexander Mathis (EPFL, Switzerland) and Manon Michel (UCA, France), the internship can be primarily based at either EPFL or UCA. This opportunity is open to Master's students, including first-year students (Master 1), who are passionate about exploring the intersection of neuroscience and mathematics. For Master 2 students, the internship will aim to lay the foundation for an interdisciplinary PhD project focused on developing and analyzing more data-efficient reinforcement learning algorithms for motor control through quantitative analytical and numerical approaches.

Profile: We are seeking a curious and motivated individual with a strong foundation in applied probability, statistical physics, neuroscience and/or machine learning. A passion for interdisciplinary research combining biology, physics, mathematics and computational methods is a plus. Experience in programming is highly desirable. For further details, please contact Alexander Mathis (alexander.mathis@epfl.ch) and Manon Michel (manon.michel@uca.fr).

References:

[1] M. H. Davis, *Piecewise-Deterministic Markov Processes: A General Class of Non-Diffusion Stochastic Models*, *JRSS B*, 46 (3): 353–388, 1984.

[2] A. Monemvassitis, A. Guillin and M. Michel. *PDMP characterisation of event-chain Monte Carlo algorithms for particle systems*, *Journal of Statistical Physics* 190, 66 (2023).

[3] A.S. Chiappa, A.M. Vargas, A.Z. Huang, and A. Mathis. *Latent Exploration for Reinforcement Learning*, *37th Conference on Neural Information Processing Systems (NeurIPS)*, 36 (2024)

[4] T.A. Berrueta, A. Pinosky, and T.D. Murphey. *Maximum diffusion reinforcement learning*, *Nat Mach Intell* 6, 504–514 (2024).