Boris Gutkin

Ecole Normale Superieure

Need Is All You Need: A Mathematical Theory of Biologically Rational Behavior Through Homeostatically Regulated Reinforcement Learning

Survival—the ultimate imperative of living systems—demands efficient regulation of internal milieu and robust defense against perturbations. In complex, dynamic environments, this requires organisms to learn predictively adaptive behavioral strategies: they must adjust to environmental challenges before those challenges prove catastrophic.

In this talk, I will survey a decade of research developing a mathematical theory of predictive homeostasis, or allostasis, that unifies homeostatic regulation with reinforcement learning. At the heart of this framework lies the classical concept of drives—linking motivation to interoception—and their minimization as a control principle. We propose a biological definition of primary rewards as action outcomes that fulfill physiological needs, translating mathematically into rewards and punishments as drive gradients conditioned on executed actions.

Our normative theory, Homeostatically Regulated Reinforcement Learning (HRRL), demonstrates how internal physiological states modulate the learning of motivated behaviors. We prove that reward-seeking is mathematically equivalent to maintaining physiological stability—establishing what we call the *physiological rationality* of behavior. This framework provides a formal basis for temporal discounting, showing that discount factors motivate animals to follow geodesic paths through physiological state space toward homeostatic setpoints.

The theory explains how animals learn to act predictively, preempting homeostatic challenges and responding to appetitive reinforcers in a state-dependent manner. Critically, agents transition from reactive homeostatic control to proactive allostatic strategies, deliberately deviating from setpoints to counter anticipated perturbations. Time permitting, I will discuss recent extensions modeling sodium appetite, state-dependent dopamine responses to rewards, and continuous-time, continuous-space deep HRRL architectures.