Jan Peters

Inductive Biases for Robot Reinforcement Learning

The quest for intelligent robots capable of learning complex behaviors from limited data hinges critically on the design and integration of inductive biases—structured assumptions that guide learning and generalization. In this talk, Jan Peters explores the foundational role of inductive biases in robot learning, drawing from insights in control theory, neuroscience, and machine learning. He discusses how exploiting physical principles, modular control structures, symmetry, temporal abstraction, and domain-specific priors can drastically reduce sample complexity and improve robustness in robotic systems.

Through a series of concrete examples—including robot table tennis, tactile manipulation, quadruped locomotion, and dynamic motor skill learning on anthropomorphic arms—Peters illustrates how inductive biases enable efficient policy search, reinforcement learning, and imitation learning. These applications demonstrate how embedding prior knowledge about motor primitives, control hierarchies, or contact dynamics helps robots acquire versatile skills with minimal data. The talk concludes with a vision for future robot learning systems that integrate such structured biases with modern data-driven methods, enabling scalable, adaptive, and generalizable autonomy in real-world environments.