

Probabilistic Inference and Learning for Control

Marc Deisenroth (University of Washington)

January 4, 2011

We propose PILCO, a data-efficient and fully probabilistic model-based framework for autonomously learning transition dynamics and controllers in the absence of expert knowledge. In most autonomous learning scenarios either task-specific domain knowledge and/or many trials are required to learn a task. In practical applications, however, full knowledge about the underlying dynamics or thousands of trials might be unavailable/impractical. PILCO learns a probabilistic dynamics model from data only. By representing and incorporating model uncertainty into the decision making process, PILCO reduces model bias and fully automatically learns to solve fairly complicated control problems in only a few trials. Across multiple complicated control tasks, PILCO achieves an unprecedented degree of automation and an unprecedented speed of learning.