Learning to Fly like a Bird

Russ Tedrake
MIT

In this talk, I'll describe some of our new work using machine learning to design nonlinear control systems for robotic birds. First, I will describe the problem of making a fixed-wing aircraft execute a post-stall maneuver in order to land on a perch. I'll describe our experimental system identification, which revealed a surprisingly compact lumped-parameter model for the problem, and our model-based approximate optimal feedback control algorithm which has produced some promising initial results. Then I'll describe our experiments with a large (2m wingspan) autonomous robotic bird, which flies with flapping wings, and try to convince you that these birds are solving the ultimate manipulation problem (manipulating the air). Finally, I'll describe results of using policy gradient reinforcement learning to optimizing the performance of a more careful flapping-wing experiment in a water tunnel. We have not, so far, identified any compact lumped-parameter models for this system, but I will show that our model-free optimal control algorithms can still optimize flapping efficiency (measured as a dimensionless cost of transport) in just a few minutes of trial-and-error with the real system. The main theme of this work is that high-fidelity dynamic models for these complicated fluid regimes may not be necessary for designing high-performance control systems (birds don't solve Navier Stokes!).