Estimation and control as dual trajectory optimization problems

Emo Todorov, University of Washington

***Abstract*—We introduce a computational framework for using trajectory optimization to solve both estimation and control problems online. The trajectory is defined over a sliding window centered at the present moment. Costs over future states correspond to task errors and energy, while costs over past states correspond to data likelihoods and priors. Physical consistency priors are evaluated by a general-purpose physics engine.**

Keywords: estimation; control; optimization

I. INTRODUCTION

Bayesian estimation and optimal control are closely related problems. The formal duality between them in the LQG case was known since Kalman, and was more recently extended to certain nonlinear systems [1, 2]. Such formulations however are global and suffer from the curse of dimensionality. In contrast, local trajectory-based methods for control have been able to generate remarkably complex behaviors, especially in contact-rich domains [3, 4]. This talk will summarize our latest results in terms of control, and then develop a related approach to estimation. The estimation approach is based on fixed-lag smoothing: the trajectory over some interval into the past is optimized with respect to the usual estimation costs. What is new here is the use of a general-purpose physics engine (MuJoCo) to define the prior. This results in dynamically-consistent estimation, incorporating not only the sensor data but also the laws of physics. The new prior is made possible by our convex contact model [5] which has an analytical inverse. Given a kinematic representation of the trajectory, we can recover both the control forces and contact forces, and evaluate all relevant cost terms. The computation scales linearly with the number of contacts and is faster than the forward dynamics simulation. This enables real-time estimation with 20-timestep trajectories for humanoid robots.

II. RESULTS

Results will be demonstrated on tracking a walking Darwin robot. The actual walking is done by a pre-existing ZMP-like controller. The focus here is on the estimation. We record sensor data from joint potentiometers, IMU, foot contact sensors, as well as motion capture markers for ground-truth comparison. The figure shows leave-one-out cross validation. The measured output of every sensor can be predicted given the trajectory estimated from the other sensors.

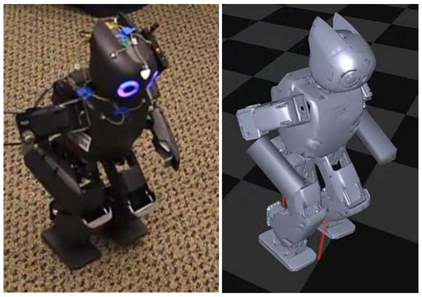

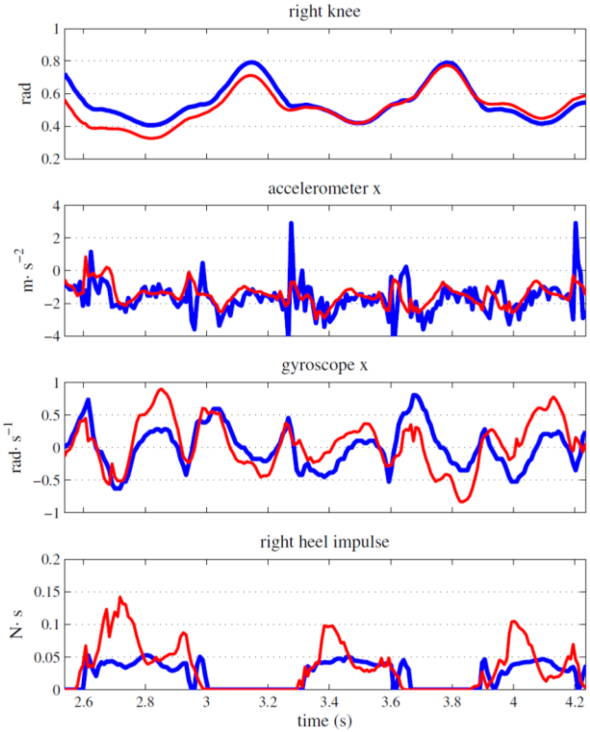

REFERENCES

[1] S. Mitter and N. Newton. A variational approach to nonlinear estimation. *SIAM J Control Opt*, 42, 2003.

[2] E. Todorov. General duality between optimal control and estimation. CDC 2009.

[3] I. Mordath, E. Todorov and Z. Popovic. Discovery of complex behaviors through contact-invariant optimization. SIGGRAPH 2012.

[4] Y. Tassa, T. Erez and E. Todorov. Synthesis and stabilization of complex behaviors through online trajectory optimization. IROS 2012.

[5] E. Todorov. Analytically-invertible dynamics with contacts and constraints: Theory and implementation in MuJoCo. ICRA 2014.