Synthesis of Complex Behaviors with Optimal Control

Emo Todorov, Yuval Tassa, Tom Erez, Igor Mordatch, Paul Kulchenko, Vikash Kumar
University of Washington

Summary

Optimal control has become the mainstream theory of systems-level information processing in the motor system. However the success stories are based on simple examples. Simple is of course a good place to start, but we already started decades ago. If current models indeed capture the essence of neural computation, what is our excuse for not making them scale? Having run out of excuses, we spent the past three years developing optimal control machinery that can handle some of the most complex motor behaviors that humans perform -- including walking, running, climbing, getting up, dexterous hand manipulation (with 20 degree-of-freedom hands), as well as cooperative actions. By "handle" we mean that our methods can automatically generate plausible-looking and physically-realistic movements, given only high-level cost functions encoding the desired outcome. For example, the task of spinning a pen between the fingers is encoded with a quadratic cost on the angular velocity of the pen. The actual finger movements and finger-object interactions emerge from optimizing this simple cost. While these methods are not foolproof, we have been able to make them work on every motor behavior we have focused on, with surprisingly little effort on our part and CPU times of a few minutes. This has led to three key insights: (i) contact dynamics are essential, and we had to develop new mathematical models of contact that are realistic yet amenable to numerical optimization; (ii) the best way to make progress is not to invent new learning algorithms (indeed everyone uses Newton's method eventually) but rather to invent problem formulations that are more tractable; (iii) learning global feedback controllers remains challenging, however optimizing individual movement trajectories is very efficient and can sometimes be done in real-time, giving rise to implicitly-defined feedback controllers with remarkable performance.

Additional Detail

This talk will summarize recent results we have published in robotics and computer graphics:

First-exit model predictive control of fast discontinuous dynamics: Applications to ball bouncing. *International Conference on Robotics and Automation* 2011
Trajectory optimization for domains with contacts using inverse dynamics. *International Conference on Intelligent Robots and Systems* 2012
Synthesis and stabilization of complex behaviors through online trajectory optimization. *International Conference on Intelligent Robots and Systems* 2012
Discovery of complex behaviors through contact-invariant optimization. *ACM SIGGRAPH* 2012
Contact-invariant optimization for hand manipulation. *Symposium on Computer Animation* 2012

Some images illustrating the range of results are shown below, but the best way to see what we are doing is watch a movie: http://homes.cs.washington.edu/~todorov/media/sfn.mp4

We will include new results obtained in the meantime, particularly in the context of the DARPA Robotics Challenge (see http://theroboticschallenge.org/) where we are one of the funded teams. We will also present applications to brain-machine interfaces (bottom plots), where we use eye tracking and speech recognition to let the user/patient specify objects of interest and desired outcomes, and leave the rest to automatic control, without drilling holes in anyone's head.

The proposed talk will have three segments:
1. Brief overview of optimal control modeling in sensorimotor control, emphasizing how simple the existing examples are and arguing that we need to do a lot better;
2. Intuitive introduction to our new methodology, covering methods/software for physics simulation, contact modeling, and numerical optimization (with minimal math);
3. Lots of movies, with live narrative clarifying where the challenges are and which aspects of our new methodology are needed to address them.