Dynamic Optimization of Behavior

Emo Todorov, University of Washington

The control of rich dynamic behaviors involving underactuation and contact dynamics remains a mystery. The brain somehow does it, in a way that appears related to optimal control, however the algorithmic details are hard to infer from experimental data. While optimal control is a mature mathematical framework, available algorithms run into the curse of dimensionality. Trajectory optimization has emerged as the most scalable approach. The general idea dates back to the 1950's. Yet until recently, this form of optimization was limited to cleaning up pre-existing trajectories, or to optimizing the smooth segments between manually-specified contact events. And even this these limitations, it was too slow for online re-planning of complex movements.

In this talk I will describe our efforts to make trajectory optimization fully autonomous so that no montion capture or manual scripting is required, and speed it up so that it can be used for online re-planning. Our latest methods have been able to solve every problem we have thrown at them -- including walking, running, getting up, manipulating objects, performing cooperative actions. On a single computer, optimizing a movement without informative initialization takes minutes. Once a solution is found, re-optimizing it online (i.e. warm-starting from the previous iteration) takes tens of milliseconds. While there is still plenty of room for improvement, in particular by using more parallel computing, the existing methods can already solve complex control problems fully autonomously. Unfortunately there is no simple trick to make this happen; instead we have had to work hard on multiple fronts: developing a fast physics engine (MuJoCo), new models of contact dynamics that are realistic yet suitable for differentiation and optimization, and a long list of algorithmic refinements than enable our optimizers to navigate challenging cost landscapes.

Bio

Emo Todorov obtained his PhD in Cognitive Neuroscience from MIT in 1998. He was then a postdoc in Computational Motor Control at University College London, research scientist in Biomedical Engineering at USC, Assistant Professor in Cognitive Science at UCSD, and is now Associate Professor in Applied Mathematics and Computer Science & Engineering at UW. His research focuses on intelligent control in biology and engineering -- in particular using numerical optimization to understand how brains, and some day robots, can autonomously generate complex yet successful movements.