Training a Humanoid Walker: A Case Study for Physical Modelling, Optimization & Deep Reinforcement Learning for Controls System Design

Dr Francesco Ciriello

MathWorks

Monday 8 February 2020
2pm

Abstract:

In this session, Francesco presents on how to model the dynamics of a humanoid robot and train it to walk using two different techniques: first using a genetic algorithm and then using deep reinforcement learning. The case study uses physical modelling tools native to Simscape™ Multibody™, a simulation environment for 3D mechanical systems that is embedded within Simulink®. The example starts with a simple demonstration of how to import a plant model from a common robotic description. The plant model can be used to design the control system and test system-level performance. The robot interacts with a simple floor environment using new Simscape™ Multibody™ features for spatial contact force modelling. Francesco discusses how to avoid common pitfalls and bottleneck using contact proxies and the appropriate choice of contact parameters and solver settings.

The availability of a plant model in Simulink makes it easy to use global optimization algorithms for the design and tuning of a controller. In the example, the parameters of a central pattern generator are optimized to teach the robot to walk. The same plant model of the robot can also be used as an environment against which to train a reinforcement learning agent. Francesco introduces MATLAB® reinforcement learning tools and draws analogies between reinforcement learning and traditional control system design. Finally,

he demonstrates how to create a deep deterministic policy gradient (DDPG) reinforcement learning agent and train it to maximise policy action commands subject to a simple reward function that is designed to get the robot to walk in a straight line. This example is shipped with MATLAB® (R2020a and above) and can be fully accessed and run by attendees.