Learning the Foundations for Humanoid Autonomy

Chad Jenkins
Brown University

Robotic humanoid systems are becoming a realization of science fiction into science fact. Despite popular assertions, a robot is defined as a machine (embodied in some physical form) that acts (or actuates its motors) based on information about its environment (provided by its sensors). Based on this definition, a robot is inherently autonomous and consists of three primary components: actuators, sensors, and a control policy. Advances in Engineering and Materials Science have culminated into humanoid systems that are approaching the actuation capabilities of human beings. In parallel, there is ongoing proliferation of increasingly smaller and cheaper sensing technologies providing robots with richer and more accurate information about their environment.

As sensing and actuation technologies advance with faster computing, one might believe crafting autonomous robot control policies would be a relatively achievable task. Unfortunately, we have found this not to be the case as autonomous control of humanoids has remained a difficult problem in robotics and related areas of research (e.g., computer animation, artificial intelligence, biomechanics). Even when restricted to limited domains, the manual crafting of a control policy can be a tedious and time-consuming endeavor that may not afford scaling to wider classes of behavior. The difficulty surrounding autonomous humanoids involves several challenges for low-level control (such as following desired motions, locomotion, maintaining balance, manipulating objects) and high-level decision making (i.e., directing low-level control to achieve objectives).

A viable approach to autonomous control is learning from human demonstration. Human beings have control policies that allow us to be functional for moving about in a dynamic world. These policies are internal and unconscious to humans and, thus, difficult to form into a computational mechanism. However, we are able to externally observe human performance resulting from these internal policies. Data collected from human demonstration have the potential to be reverse engineered into computational control policies for humanoid motion. In order to learn from demonstration, we must address three basic issues:

-Acquisition: how to acquire human motion performed in the real world into a digital form?

-Learning: how to uncover latent structures underlying collected motion data?

-Utilization: how to utilize uncovered structures for control and perception to further robot autonomy?

This talk will present recent research advancements in the acquisition, learning, and utilization of human demonstration for autonomous humanoid control. In terms of motion acquisition, we will describe traditional constrained approaches to human motion capture, untethered motion capture with inertial sensing, and markerless motion capture from multiple viewpoints. We will summarize recent methods in manifold learning and dimensionality reduction and the application of these methods to learning models of human motion. In particular, we describe our work on learning modular vocabularies of exemplar-based primitives expressing the nonlinear dynamics of human behavior. Lastly, applications of learned motion models to low-level humanoid control, dynamic interactions for physically animated humanoid characters, and perception of human activity will be presented.

Keywords:
Humanoid robot: a machine with a human-like physical structure that acts based on sensory information

Control policy: a function that maps that state of a robot's system into a control command for its actuators

Learning from demonstration: for robots, the process of automatically creating a control policy from an external demonstration

Motion capture: the recording of movement performed by physical entities for immediate or delayed analysis and playback

Manifold learning: statistical methods for uncovering the manifold underlying a dataset