

LEARNING CONTROLLERS FOR HUMAN-ROBOT INTERACTION

By

Eric M. Meisner

An Abstract of a Thesis Submitted to the Graduate

Faculty of Rensselaer Polytechnic Institute

in Partial Fulfillment of the

Requirements for the Degree of

DOCTOR OF PHILOSOPHY

Major Subject: Computer Science

The original of the complete thesis is on file
in the Rensselaer Polytechnic Institute Library

Examining Committee:

Volkan Isler, Thesis Adviser

Jeff Trinkle, Thesis Adviser

Linnda Caporaël, Member

Sanmay Das, Member

Brian Scassellati, Member

Rensselaer Polytechnic Institute
Troy, New York

April 2009
(For Graduation May 2009)

In order for robots to assist and interact with humans, they must be socially intelligent. Social intelligence is the ability to communicate and understand meaning through social interaction. Artificial intelligence can be broadly described, as an effort to describe and simulate the property of human intelligence inside of a computational model. In most cases, this simulation happens in a vacuum. An agent, such as a robot, maintains a computational model which includes all information required to make decisions. This internal computational representation is what we consider its intellect. Information may enter this model through perception, and be expressed in the form of action. This separation of knowing and doing can be quite effective for representing certain types of intelligence. However it does not lend itself to simulating social cognition.

In order to communicate socially, an agent must be able to affect change in the mental representation of other agents, as well as the physical world. However, social interaction and the creation of meaning is inherently different than interactions with the physical world. Because there are no mathematical models which describe how actions and perceptions affect the mental representations of a human, we cannot hope to build an interactive agent by directly simulating this process. For this reason, when building artificial social intelligence, we need to pay attention to prevailing theories on how humans learn. Many popular theories from cognitive science, social psychology and language development suggest that action and perception are not subordinate to mental representations. Instead, mental representations are a result of action and perception that results from an agent's interaction with an environment and other agents. In particular, social learning theory says that the process which allows agents to understand one another happens from the ground up, starting with action and perception, and resulting in the shared mental representations, and understanding of how to affect change in the representations of others.

This thesis addresses the problem of building social intelligence into robotic systems using computational learning and adaptive control. We focus on the problem of how to use decision theoretic planning for learning to interact with humans from the bottom up. We first examine the use of affect recognition in designing

human-friendly control strategies. Next, we address the problem of defining subjective measures of interactivity by leveraging human expertise. Finally, we define and evaluate a method for participating in the process of socially situated cognition. We emphasize learning to predict and modulate observable responses of the human rather than attempting to directly infer their mental or emotional states. The effectiveness of this method is demonstrated experimentally using custom robotic systems.