Project Overview

The objective of this research is to enable operation of synthetic and cyborg insects in complicated environments, such as outdoors or in a collapsed building. As the mobile platforms and environment have significant uncertainty, learning and adaptation capabilities are critical. The approach consists of three main thrusts to enable the desired learning and adaptation: (i) Development of algorithms to efficiently learn optimal control policies and dynamics models through sharing the learning and adaptation between various instantiations of platforms and environments. (ii) Creation of control learning algorithms that can be run on low-cost, low-power mobile platforms. (iii) Development of algorithms for online improvement of policy performance in a minimal number of real-world trials.

Brief Summary of 2010-2011 Findings

Stimulation methods for the cyborg beetles:
(i) The wing-folding muscle: we investigated the beetle's muscle known as the wing-folding muscle. Anatomical study and high-frame rate camera experiments suggest it plays a significant role in turning. We confirmed this through free-flight experiments with beetles who had the muscle removed; and through active stimulation of the muscle. Further investigation is on-going.
(ii) *Conditioning:* in parallel to the direct muscle stimulation line of work, we are investigating whether the beetle could be controlled through conditioning, in particular operant conditioning. We built a virtual arena for our experiments. We have about everything in place to start performing the desired experiments.
Finding good control policies in risk-aware settings*:*
The expected return is a widely used objective in decision making under uncertainty and many algorithms, such as value iteration, have been proposed to optimize it. In risk-aware settings, however, the expected return is often not an appropriate objective to optimize. We have developed a new optimization objective for risk-aware planning and shown it has desirable theoretical properties. We have also drawn connections to previously proposed objectives for risk-aware planning: minmax, exponential utility, percentile and mean - variance. Our method applies to an extended class of Markov decision processes; it allows rewards to be stochastic as long as their distributions have moment generating functions. Additionally, we have developed efficient algorithms for optimizing our proposed objective.
Synthetic Insects:
To create a platform for robot learning experiments, we've extended the capabilities of our 35 gram synthetic crawling insect, the OctoRoACH. This robot features a differential drive-train controlled through manually-tuned PID controllers, as well as several on-board sensors (IMU, Back-EMF, cell-phone camera). Since the PID tuning is by no means optimal, our first learning objective is to improve its parameters, the motors' phasing, and the time-variant drive signals. A preliminary learning experiment made use of the extended capabilities to learn the policy that maximized the turning rate of the robot on carpet.