Learning to Control an Octopus Arm with Gaussian Process Temporal Difference Methods

Yaakov Engel

Collaborators: Peter Szabo and Dmitry Volkinshtein

ALBERTA INGENUITY CENTRE FOR MACHINE LEARNING

UNIVERSITY OF ALBERTA

The Octopus Arm

- Can bend and twist at any point
- Can do this in any direction
- Can be elongated and shortened
- Can change cross section
- Can grab using any part of the arm

Virtually infinitely many DOF

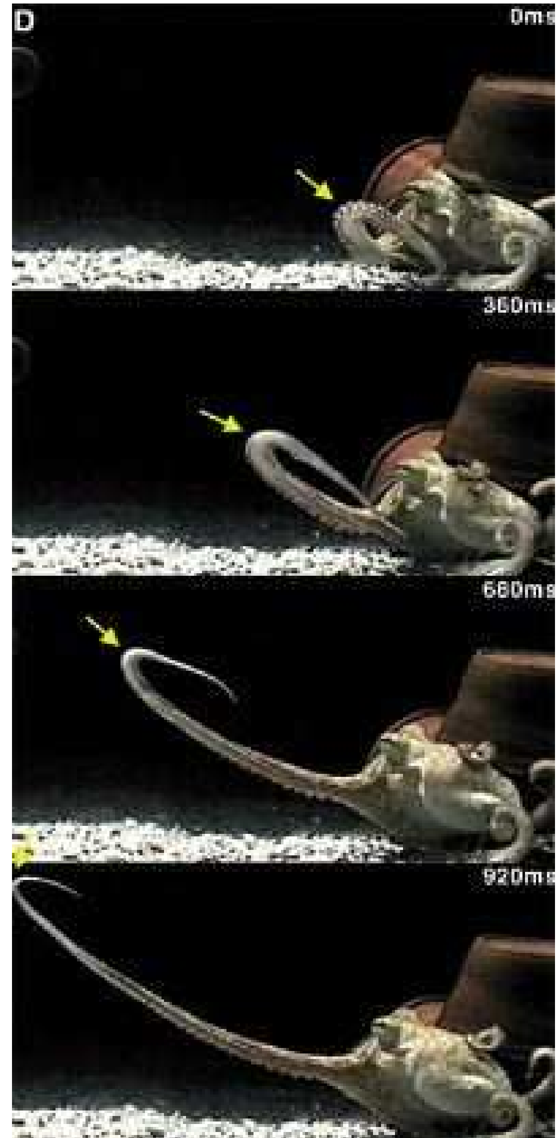

Octopus Arm Anatomy 101

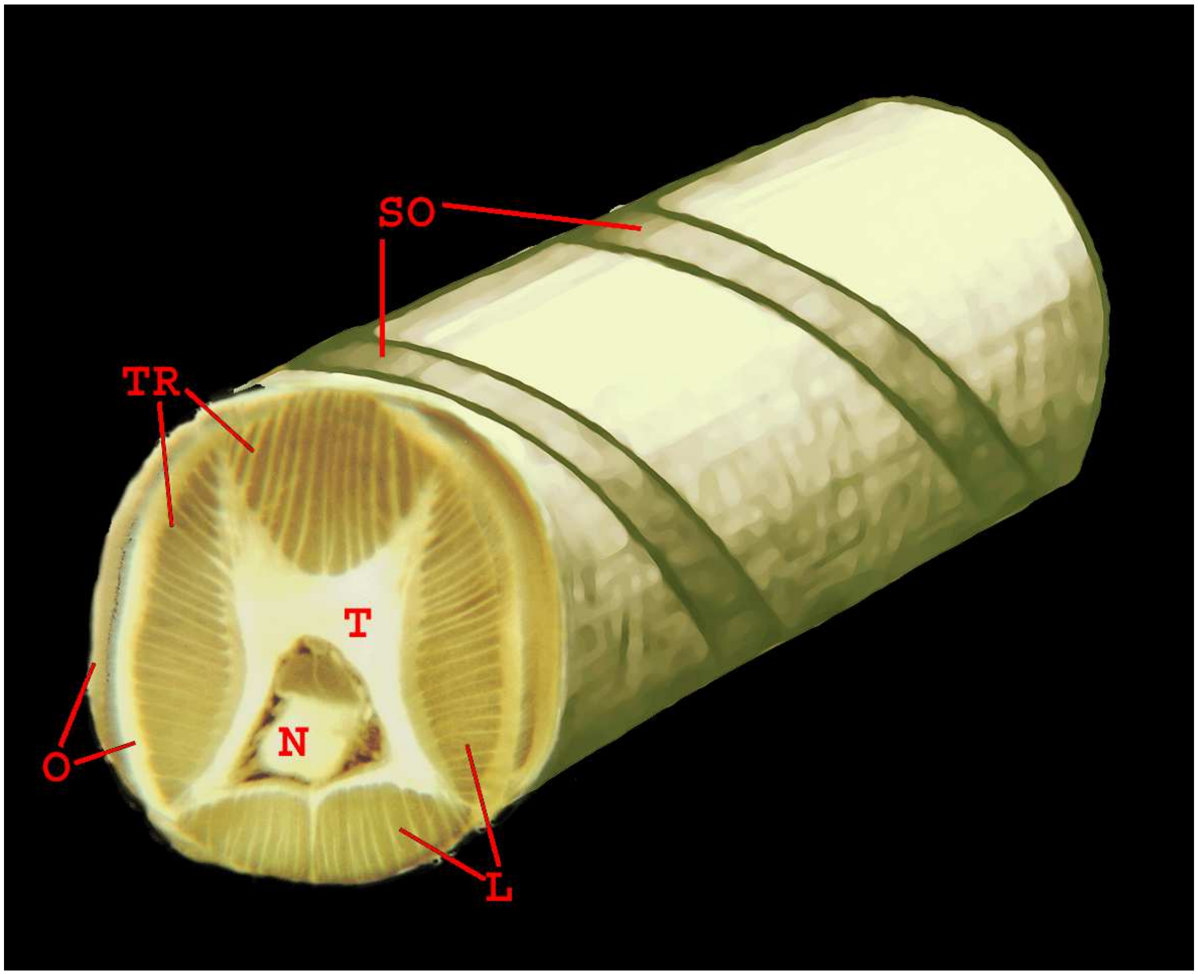

Our Arm Model

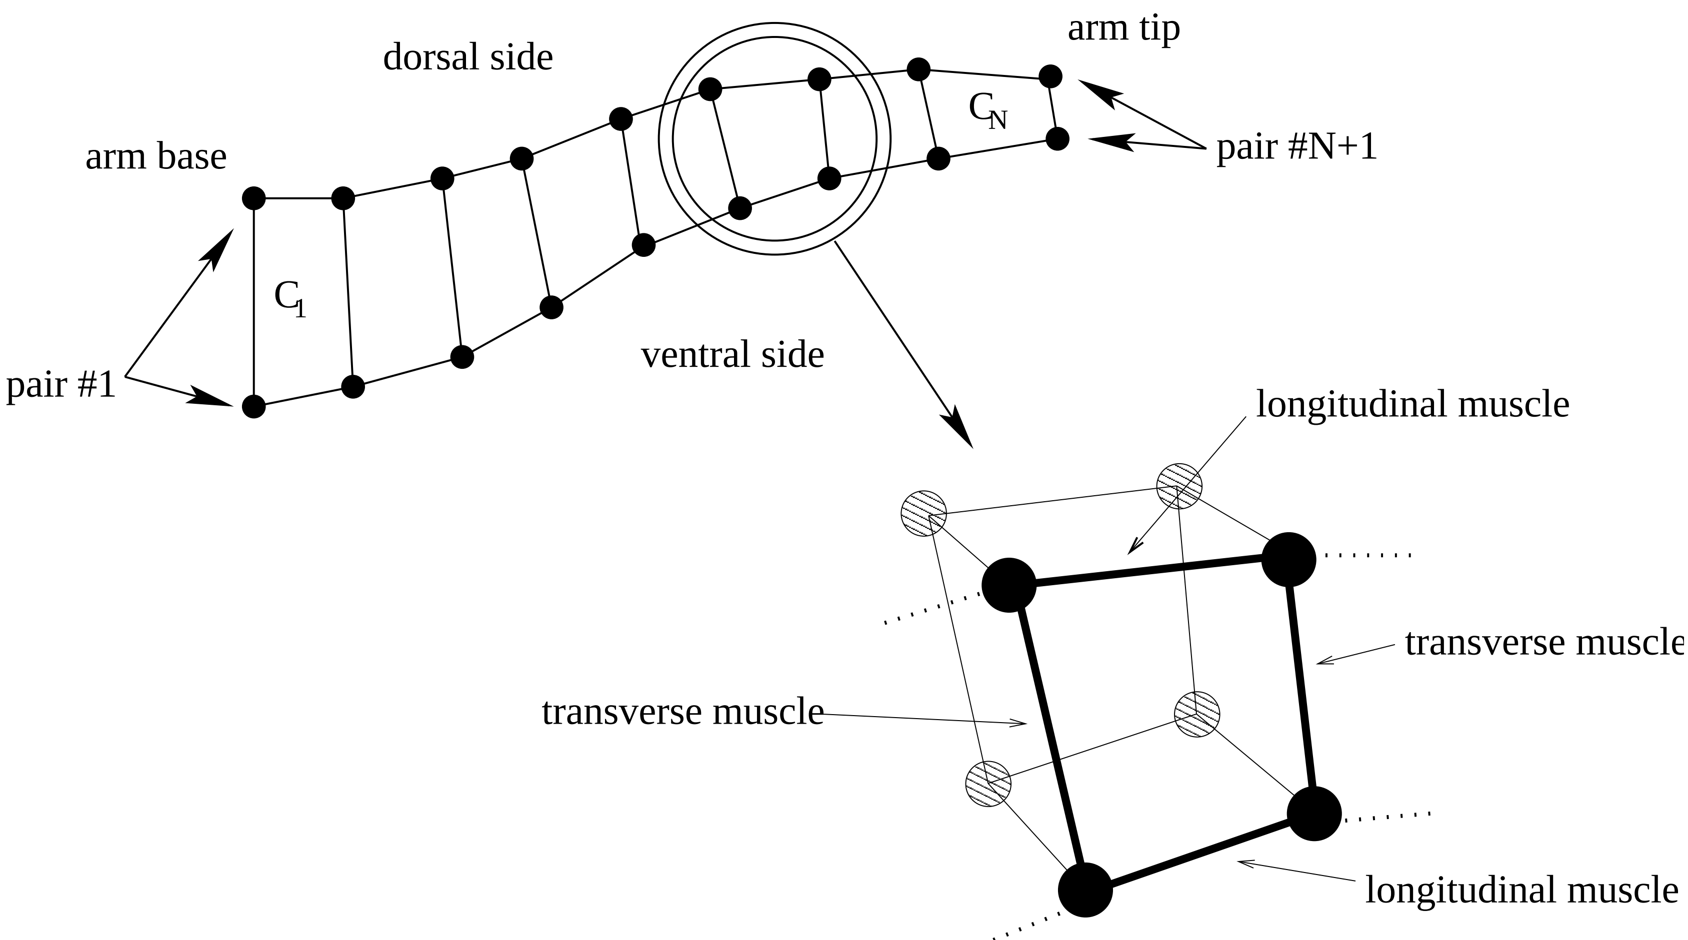

The Muscle Model

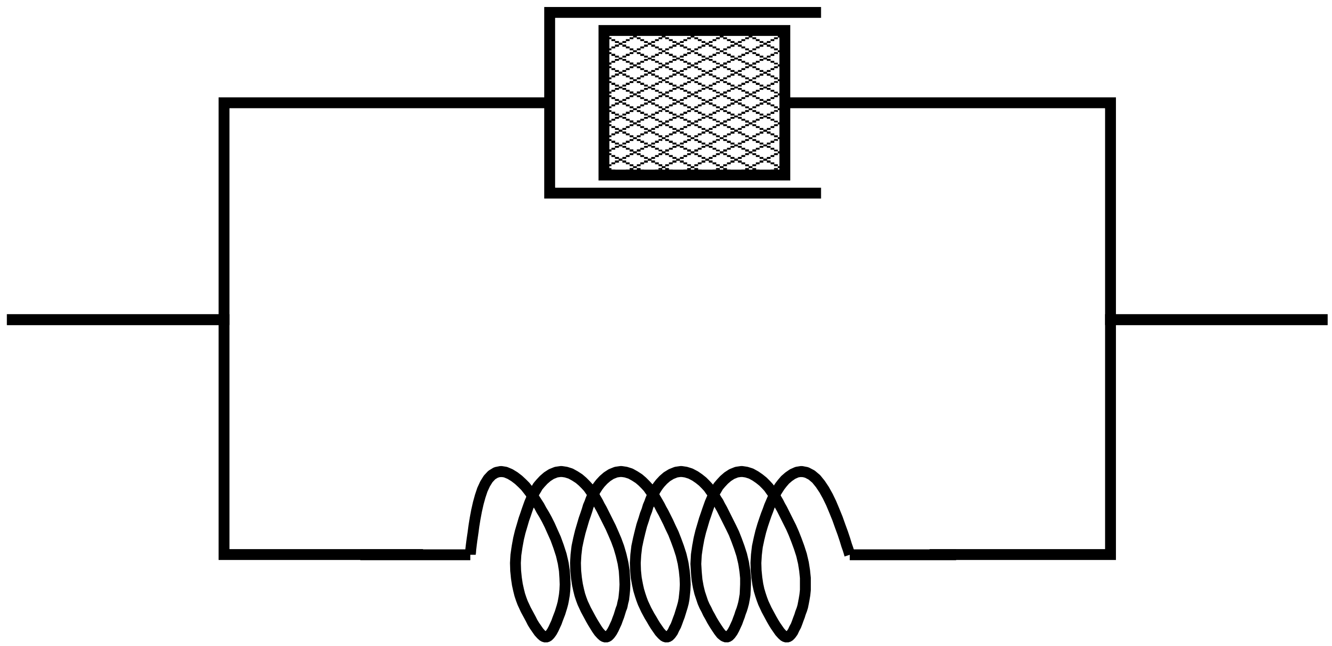

$$f(a) = (k_0 + a(k_{max} - k_0))\,(\ell - \ell_0) + \beta\frac{d\ell}{dt}$$

$$a \in [0, 1]$$

Bayes RL

Other Forces

- Gravity
- Buoyancy
- Water drag
- Internal pressures (maintain constant compartmental volume)

Bayes RL

Dimensionality

10 compartments $\Rightarrow$

22 point masses $\times$ $(x, y, \dot{x}, \dot{y})$

= 88 state variables

Bayes RL

The Control Problem

Starting from a random position, bring {any part, tip} of arm into contact with a goal region, **optimally**.

Optimality criteria:

Time, energy, obstacle avoidance

Constraint:

We only have access to sampled trajectories

Our approach:

Define problem as a MDP

Apply Reinforcement Learning algorithms

Bayes

RL

The Task

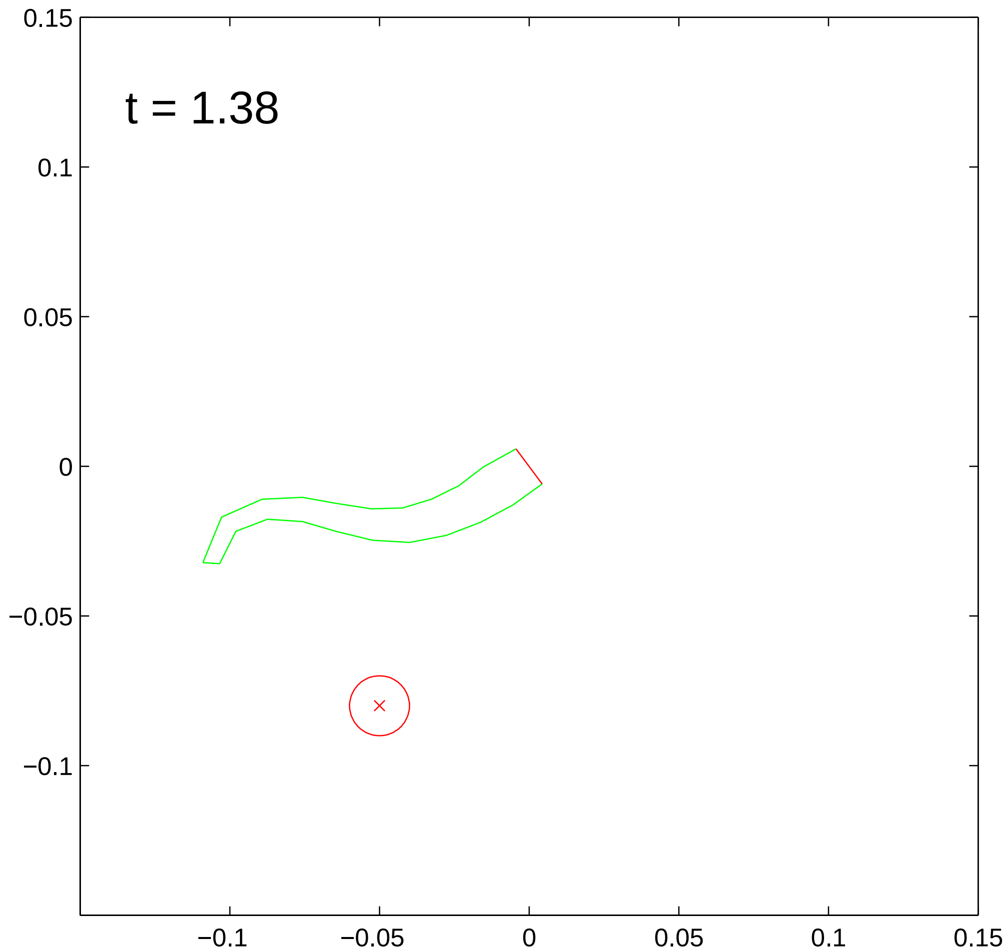

Actions

Each action specifies a set of fixed activations – one for each muscle in the arm.

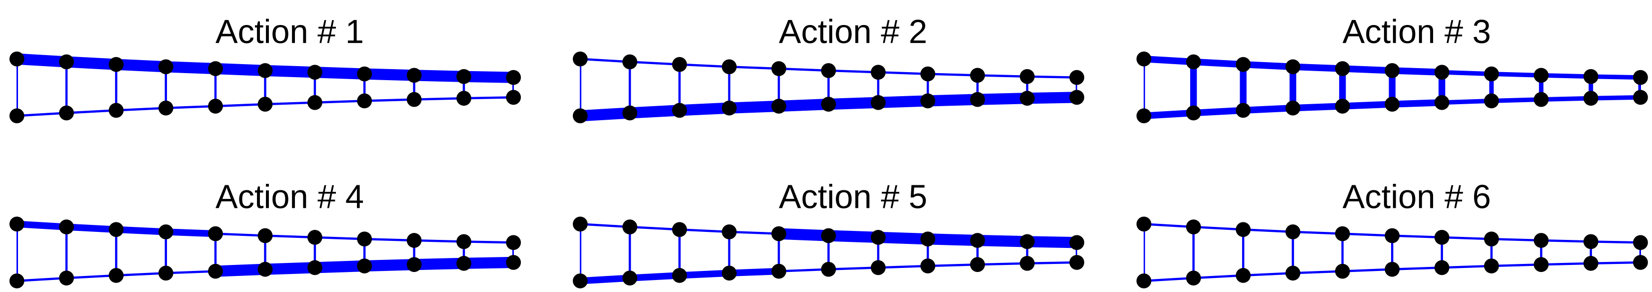

Base rotation adds duplicates of actions 1,2,4 and 5 with positive and negative torques applied to the base.

Rewards

Deterministic rewards:

+10 for a goal state,
Large negative value for obstacle hitting,
-1 otherwise.

Energy economy:

A constant multiple of the energy expended by the muscles in each action interval was deducted from the reward.

Bayes
RL

Now, to the Movies...

Fixed Base Task I

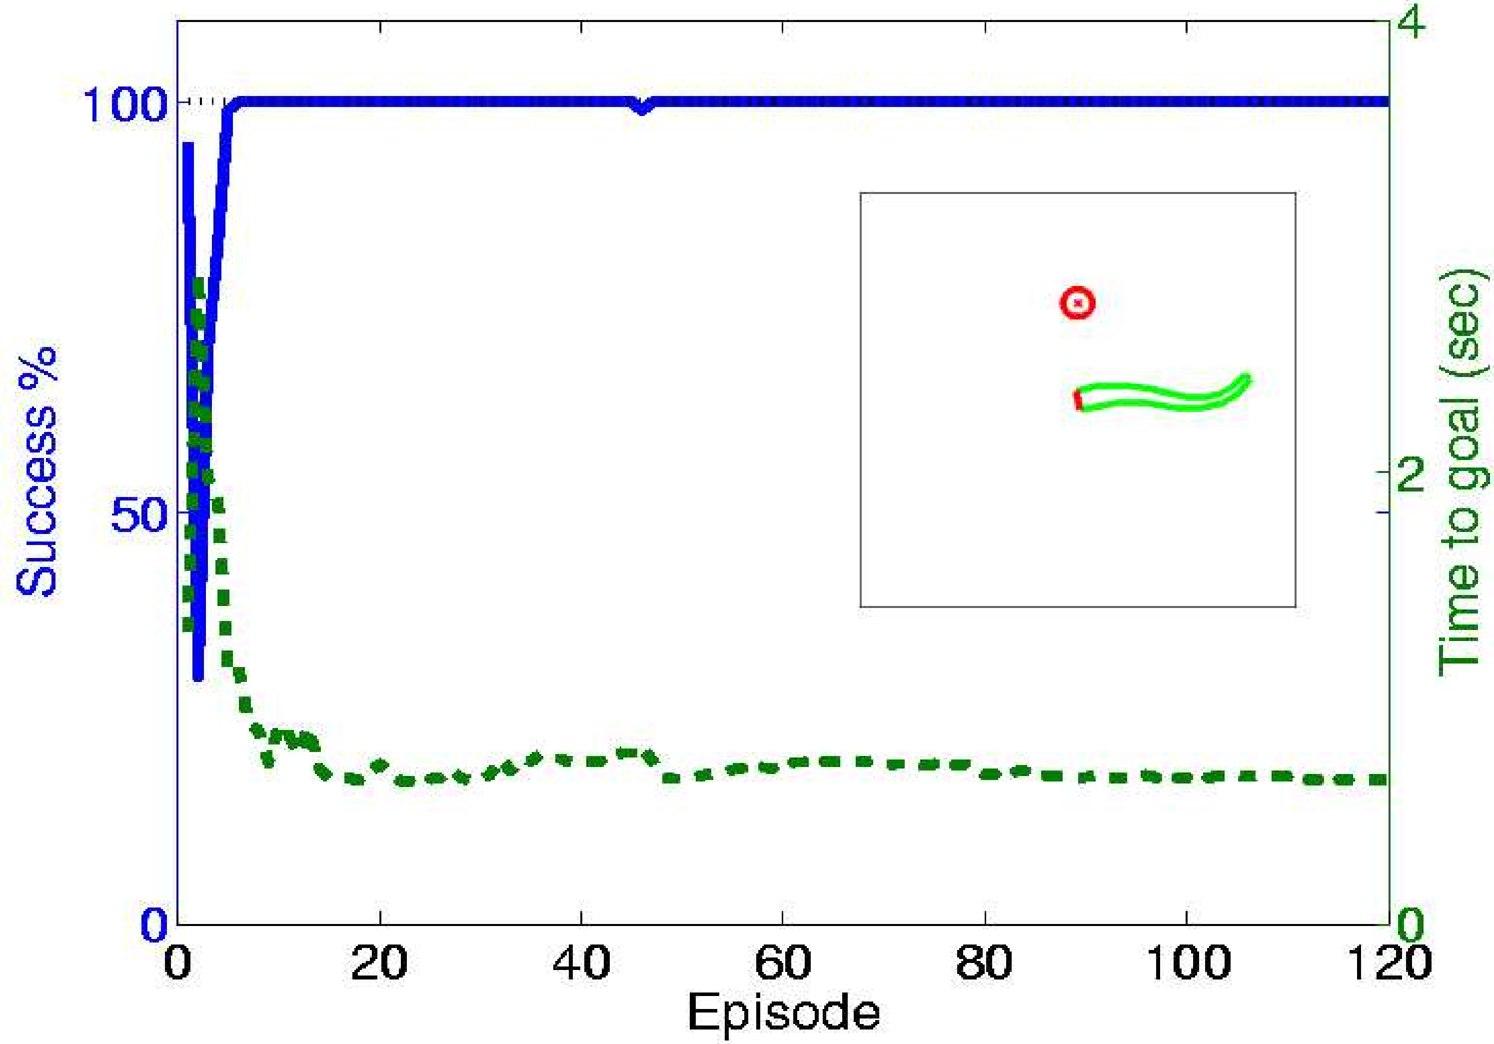

FIXED BASE TASK II

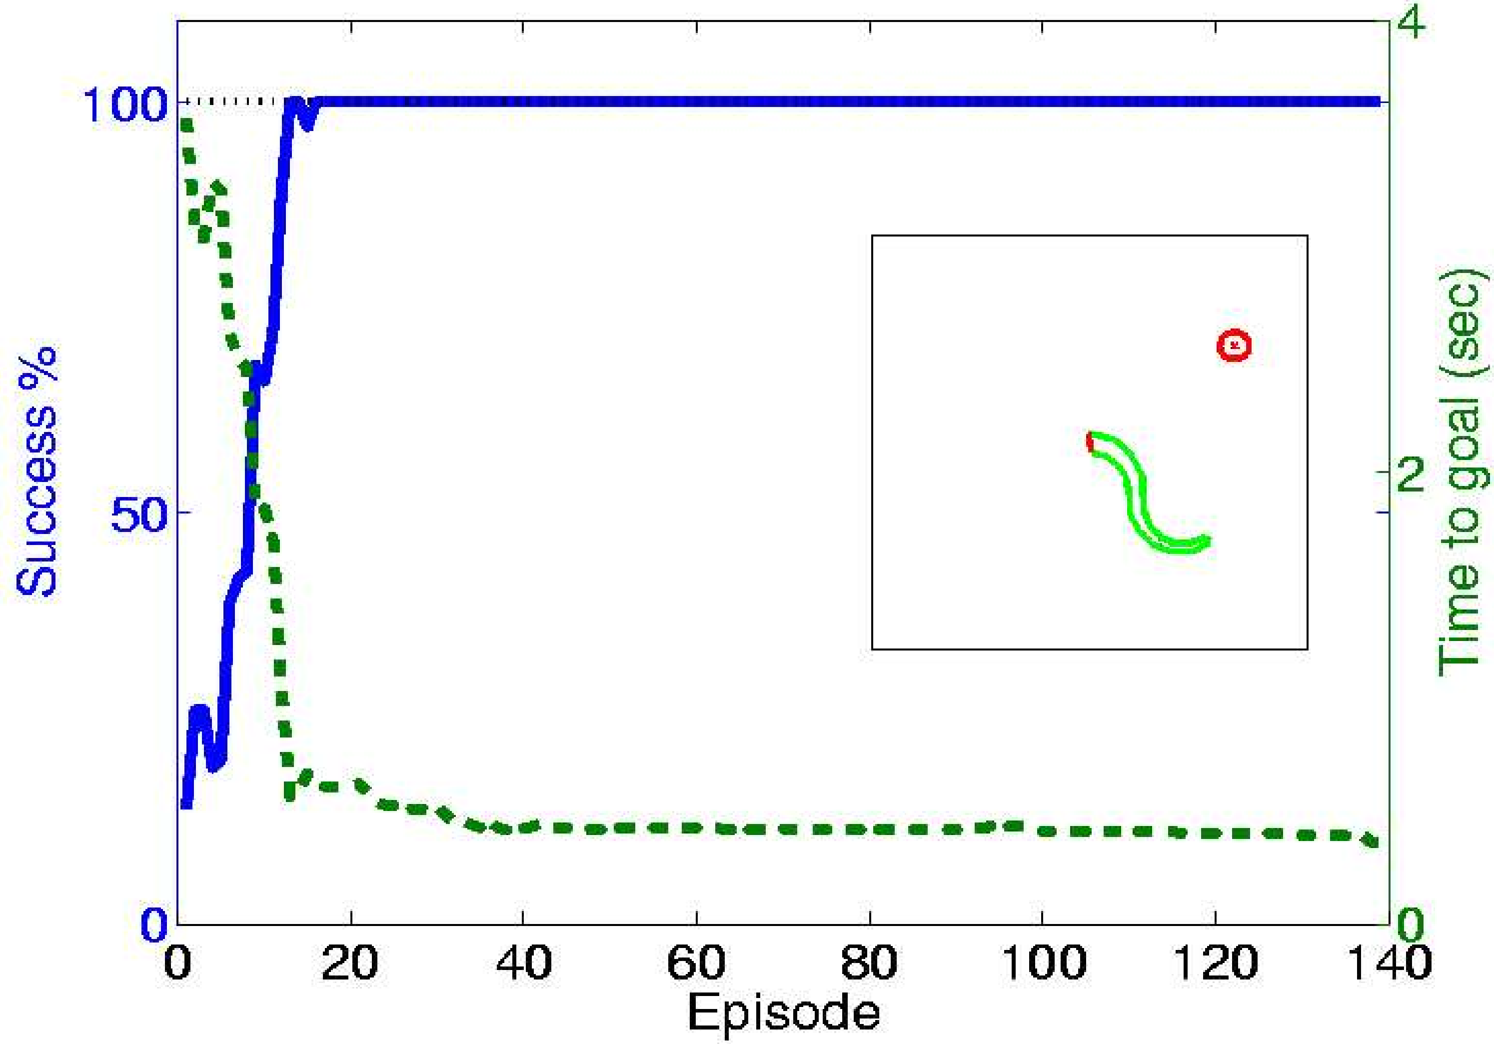

Rotating Base Task I

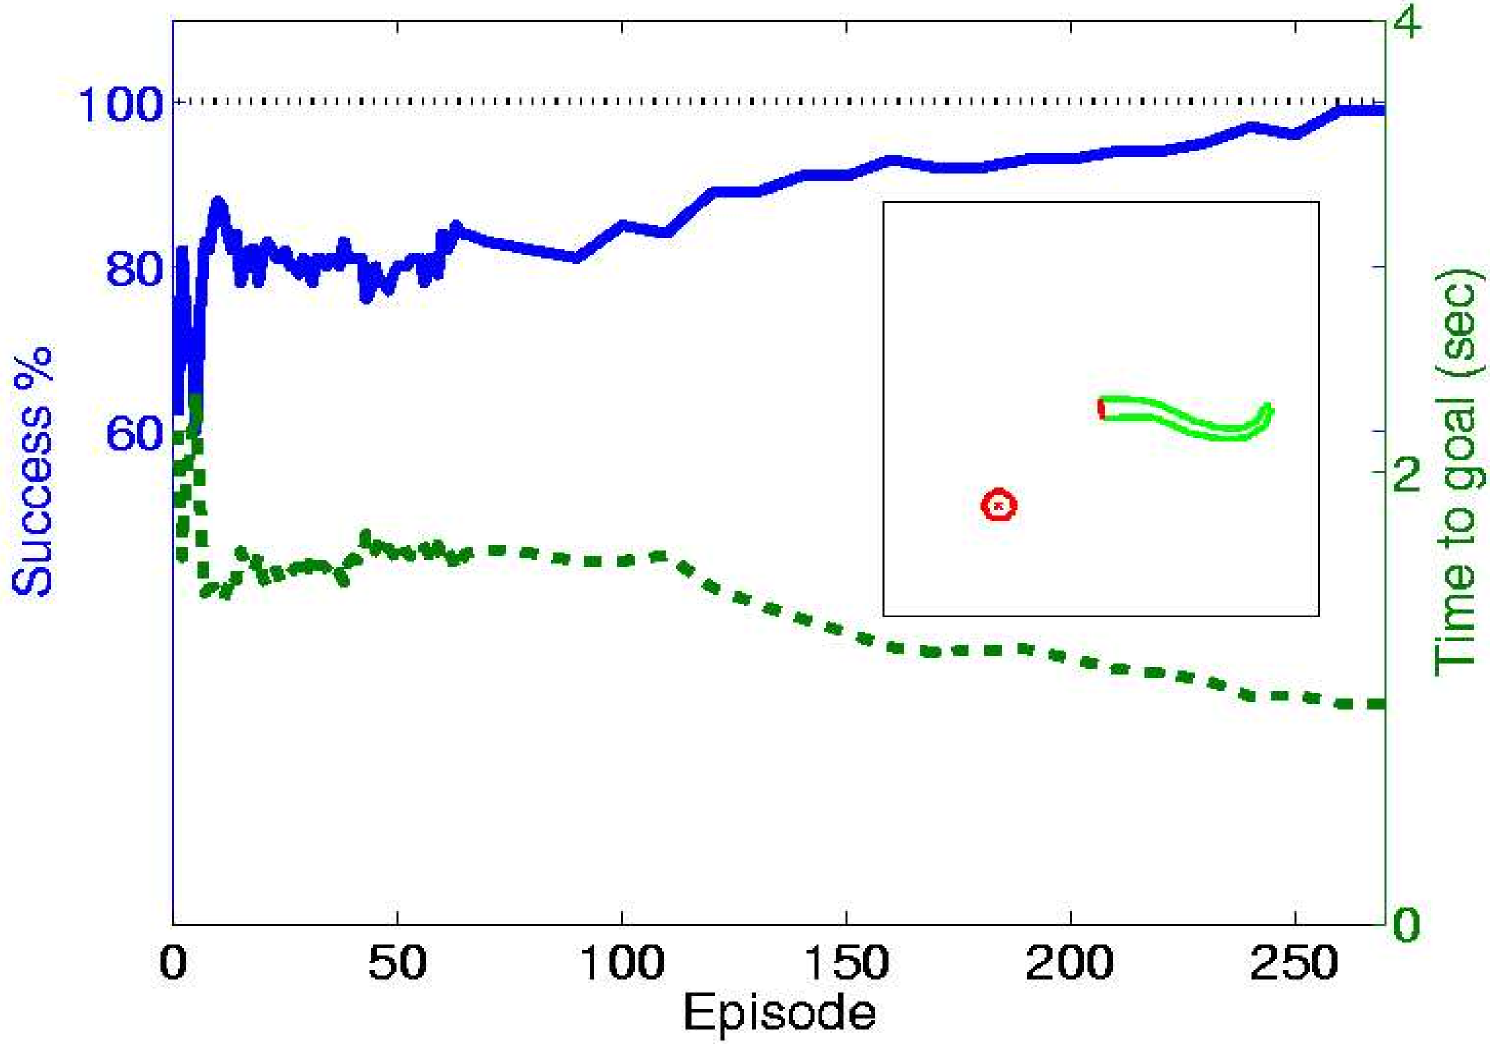

ROTATING BASE TASK II