Learning Ego-motion Relations Via Sensorimotor Correlation

Matthew J. Marjanović

The Problem: To learn the causal relations for interactions of a complex embodied system with the world, by correlating motor commands with sensory input.

Motivation: Our group has been developing a robotic torso, called Cog, with of the intention of creating a test-bed on which to study theories of cognitive science and artificial intelligence[4]. The goal is to create a robot which is capable of interacting with the world — including both objects and people — in a human-like way, so that we may study human intelligence by trying to implement it.

Such interaction requires rich sensory and motor apparatus: our two-armed, two-eyed robot has over twenty actuated joints, and twice as many sensors, ranging from torque sensors on motors to the four cameras composing the eyes. To control and coordinate so many degrees of freedom, one could in theory measure the properties of each joint, limb, lens, and CCD, pull out a physics textbook, and work out the interacting kinematics by hand. However, this is a very brittle approach, requiring extensive analysis and simulation of a particular mechanism — an effort which, even if it could be completed for such a complex system, may be wasted if the mechanism is modified or less than perfectly calibrated. A far more robust approach is to allow the robot to interact with its environment and learn a predictive model of that interaction from the experience itself.

The goal of this project is to develop a relatively general system by which Cog can learn the causal relations between commands to its motors and input from its sensors, primarily vision and mechanical proprioception. This way, the robot can learn first hand how its own movement is reflected in perceptible activity in the external world. And, conversely, such a model will allow the robot to decide how to generate actions based on the intended effect.

Such causal relationships are the root of the sense of kinesthesia, as well as the beginnings of what could be considered a sense of self. By embedding knowledge of the effects of actions directly in their sensory results, one can avoid the classic "symbol grounding problem" of artificial intelligence.

Previous Work: Previous work in this area has involved learning maps between particular sensory and motor systems, via incremental or reinforcement learning, with the causal relationship implicit in the algorithm. Figure 1 shows the results of such a task: learning a map between an arm motor command and the resulting position of the hand in the visual field. It is known a priori that the visual position of the hand and the position as given by motor joint coordinates are related by a coordinate transform, which allows the problem to be solved a simple self-supervised learning scheme. This is a special case in which two sensorimotor systems are known to be completely and nearly instantaneously correlated.

The current task, however, is to explicitly discover unknown relationships between sensory and motor systems, with correlations which may be weak or contain temporal delays.

Approach: *Tabula rasa* learning of causal relationships between sensory and motor systems is mathematically similar to reinforcement learning of state-action mappings. Prospective techniques are variations on the U-Tree algorithm[2] and G-algorithm[1], extended to work on continuous state spaces.

The initial implementation will be an extension of the current reaching task on Cog[3]. Learning the correlation between arm motor commands and motion in the visual field will allow Cog to identify its arm in the presence of visual distractions, which currently present a problem for the reaching task.

Difficulty: *Tabula rasa* learning of causal relationships between sensory and motor systems is a general systems identification task. The difficulty arises in trying to efficiently analyze a large state space, i.e. finding the proper heuristics to compartmentalize the space into relevant, learnable chunks.

On Cog, part of this compartmentalization will be achieved through enforcement of a developmental progression, i.e. allowing motor systems to come on-line slowly, one-at-a-time, so that one can be well-explored before the next is available.

Impact: For Cog, this work will create a foundation for its interaction with the world. As new sensory and motor systems are developed for the robot, Cog will have a mechanism by which it can learn to use those systems automatically.

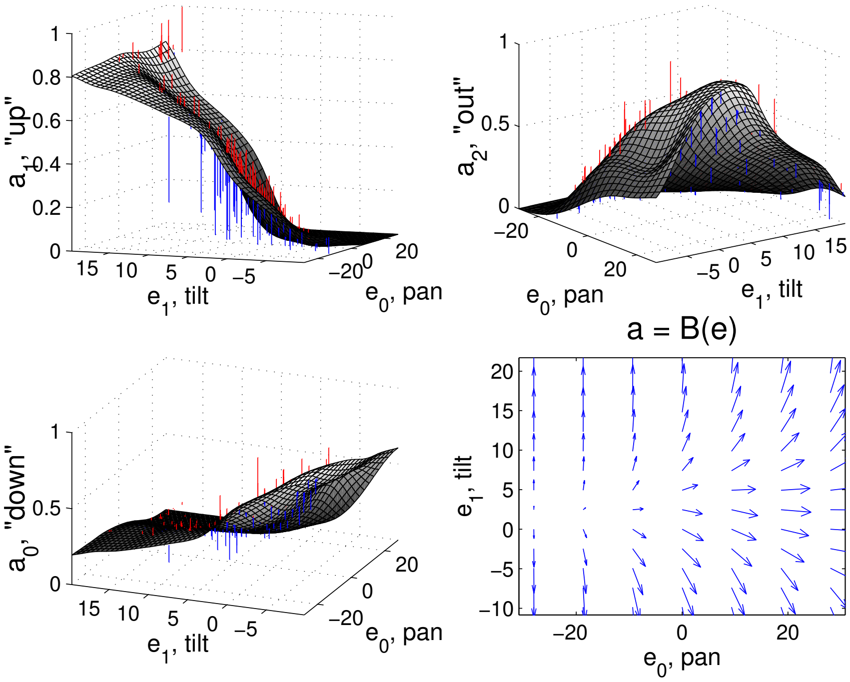

Figure 1: The relation between arm motor position, $\vec{a}$, and eye gaze position, $\vec{e}$, learned by reaching to random positions in the workspace and detecting the hand (arm endpoint) in the visual field.

For embodied robotic systems in general, this work will yield a mechanism which allows a system to automatically discover how sensory input and motor output are linked together. It will be able to learn a series of dynamic models particular to specific tasks, such as the differentiation of motion in the world from motion caused by camera movement, or the coordination of head and eye motors to track a moving object.

Aside from cognitive science test platforms like Cog, such a mechanism would be useful for any robotic system which has changing kinematic structure, or which is coupled to a changing kinematic structure — for example, a pair of adaptive prosthetic legs, which could learn to walk with a patient as much as the patient learns to walk with the legs.

Future Work: Hopefully, the architecture of this mechanism can be extended to learn the relations between higher-level activities, comprising sequences of motor and sensory states. One example is learning what tactile contact is caused by activation of an entire "reach to visual target" task.

Research Support: Parts of this research are funded by DARPA/ITO under contract number DABT 63-99-1-0012 and parts have been funded by ONR under contract number N00014-95-1-0600, "A Trainable Modular Vision System."

[1] David Chapman and Leslie Pack Kaebling. Learning from delayed reinforcement in a complex domain. Twelfth International Joint Conference on Artificial Intelligence, 1991.

[2] Andrew McCallum. Learning to use selective attention and short-term memory in sequential tasks. from *Animals to Animats 4*, ed. Pattie Maes and others, 1996, MIT Press, pp314-324.

[3] Matthew Marjanović, Brian Scassellati, Matthew Williamson. Self-Taught Visually-Guided Pointing for a Humanoid Robot. from *Animals to Animats 4*, ed. Pattie Maes and others, 1996, MIT Press.

[4] R. Brooks and L. A. Stein. Building Brains for Bodies. *Autonomous Robots*, 1994, vol 1:1, pp7-25.