Optimality of active sensing

Emo Todorov and Alex Simpkins, UCSD

Active sensing via exploratory actions is an important yet little-understood aspect of motor behavior. It includes eye movements, exploratory finger movements, whisker and ear movements, vocalizations used for echo-location, command signals sent to muscle spindles via gamma motoneurons. Such actions do not normally have direct consequences in terms of rewards or punishments. Instead they enhance the flow and quality of sensory information, and thereby contribute indirectly, by improving the feedback control of "regular" actions. This contribution can be substantial – as for example in driving, where the wrong pattern of eye movements can have grave consequences.

Our goal is to develop a computational theory of active sensing in the broader context of sensorimotor integration. We consider active sensing to be a special case of stochastic optimal control combined with Bayesian inference. This general framework is becoming the theoretical framework of choice for studying perception and action. A growing body of evidence supports the view that sensory systems integrate all sources of information in a statistically-optimal manner, and send the resulting estimates to motor systems, which apply feedback control laws optimized to yield the best possible performance.

Despite the mathematical coherence of this framework, the sensory and motor aspects of sensorimotor integration are often studied separately, resulting in a gap which prevents the theory from reaching its full potential. This is partly because the information processing which the theory calls for is often beyond the reach of existing algorithms, forcing researchers to over-simplify their problems. In particular, existing models of motor control assume that motor commands are generated on the basis of point estimates, and ignore the uncertainty associated with those estimates. In special circumstances – involving linear dynamics, Gaussian noise and quadratic performance criteria – the optimal way to act is indeed independent of estimation uncertainty. But in general uncertainty is likely to matter. We need formal models which make such effects explicit. Active sensing is an ideal candidate in that regard, because reducing uncertainty is the only purpose of exploratory actions.

We have developed two formal models in which optimal control laws for active sensing can be efficiently approximated and their predictions compared to experimental data. The first model is a model of eye-hand coordination in the context of manual tasks involving multiple objects. Key to the model is the fall-off of visual acuity with distance away from the fixation point. We represent this with state-dependent sensory noise. Different patterns of eye movements cause different patterns of uncertainty in the estimates of object positions, and thereby affect the accuracy of the concurrent hand movements. The model yields concrete predictions which are in close agreement with data from eye-hand coordination experiments. In these experiments we manipulated the visual feedback available to the subject, in ways suggested by the model, and observed changes in eye-hand coordination as predicted by the model.

The second model addresses the tradeoff between exploration and exploitation. It applies to experiments where hand movements are mapped to screen cursor movements via a well-defined but unknown to the subject mapping. Even though the subject has a simple task – to track a moving target on the screen – this task imposes conflicting demands on the hand movement system: it requires both tracking (exploitation) as well as probing and learning the unknown mapping (exploration). We represent the online learning process in innovations form, which allows us to transform the partially-observed system into a fully-observed system whose augmented state incorporates the uncertainty about the hand-cursor mapping. The solution to the resulting stochastic optimal control problem is a control law which incorporates both exploration and exploitation, and achieves an optimal tradeoff between the two. Although this tradeoff has received a lot of attenion in Reinforcement Learning, it has previously been resolved heuristically.