Available online at www.sciencedirect.com

ScienceDirect

ELSEVIER

Current Opinion in Neurobiology

Dynamic belief state representations

Daniel D Lee[1], Pedro A Ortega[1] and Alan A Stocker[1,2]

Perceptual and control systems are tasked with the challenge of accurately and efficiently estimating the dynamic states of objects in the environment. To properly account for uncertainty, it is necessary to maintain a dynamical belief state representation rather than a single state vector. In this review, canonical algorithms for computing and updating belief states in robotic applications are delineated, and connections to biological systems are highlighted. A navigation example is used to illustrate the importance of properly accounting for correlations between belief state components, and to motivate the need for further investigations in psychophysics and neurobiology.

Addresses
[1] Electrical and Systems Engineering Department, University of Pennsylvania, Philadelphia, PA 19104, United States
[2] Psychology Department, University of Pennsylvania, Philadelphia, PA 19104, United States

Corresponding author: Lee, Daniel D (ddlee@seas.upenn.edu)

Current Opinion in Neurobiology 2014, **25**:221–227

This review comes from a themed issue on **Theoretical and computational neuroscience**

Edited by **Adrienne Fairhall** and **Haim Sompolinsky**

For a complete overview see the Issue and the Editorial

Available online 15th March 2014

http://dx.doi.org/10.1016/j.conb.2014.01.018

Introduction

A key element of both perceptual and control algorithms is the need to estimate the dynamic state of a system. Consider the response of an animal as a potential predator approaches. Perceptually, it becomes important to accurately track the location of the predator as it nears, in order to decide when to flee. When the animal decides to run away, it becomes equally important to monitor the state of its own body posture, moving limbs, and muscle torques from noisy propioceptive feedback in order to maximize its running speed while maintaining balance and agility. Furthermore, a successful escape must keep track of its position relative to the predator and make accurate future predictions (Figure 1).

In this article, we first review some of the approaches used to model and track states in these situations. Then, we draw examples from engineering systems, in particular from robotics, and use these examples to motivate some key questions that arise with respect to potential dynamical belief state representations in neurobiology. Robots can be viewed as artificial model systems for understanding sensorimotor mechanisms, because their design and construction need to address many of the important challenges Nature had to face. Two crucial aspects which renders robotics appropriate as model systems are the embodiment [1] and the need for efficient, real-time processing of massive, high-dimensional sensorimotor data. Our objective is to communicate the insights we have gained with respect to dynamic belief state representations, complementing previous findings about Bayesian optimal decision making and sensorimotor integration in computational neuroscience [2–4].

State dynamics

Whether the state is the location of a predator or the angles and velocities of the leg and arm joints, there should be some predictive model of how the state changes over time. The state at time t can be written as a real valued vector $\vec{s}_t$. For example, in describing the position of an object, the state vector could contain the coordinates of the object in either rectilinear or polar coordinates. On the other hand, joint angles and their associated velocities would be described as a set of angles along with their time derivatives in the state vector.

Here we simplify our description by considering discrete time updates. In order to make accurate predictions, we would like to know how the state evolves from the previous time instant $t-1$. This can be described in terms of a motion model:

$$\vec{s}_t = f(\vec{s}_{t-1}, \vec{a}_{t-1}) \tag{1}$$

where the dynamics depend explicitly upon the previous state $\vec{s}_{t-1}$ and action $\vec{a}_{t-1}$.

A crucial issue is that the state is never directly observed. As assumed by regular hidden Markov models (HMMs) and partially-observable Markov decision processes (POMDPs), information about the underlying state is provided by observations in time, which may not fully specify the state:

$$\vec{o}_t = g(\vec{s}_t) \tag{2}$$

because the measurement function g may not be invertible. An example of such measurements includes monocular vision, where the reflected light from an object is projected upon a 2D retina array resulting in measurements with an unknown depth. In legged locomotion, information about the full body state is indirectly provided by vestibular and proprioceptive measurements,

Figure 1

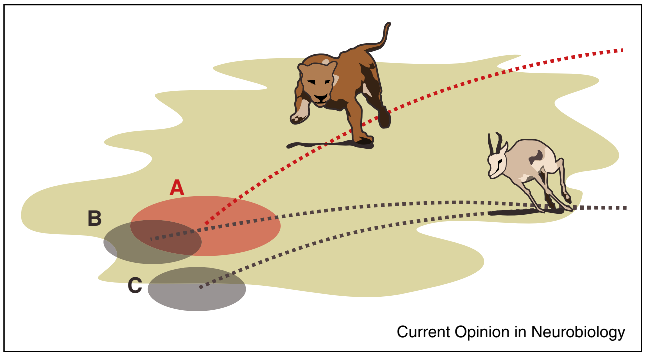

A successful escape from the predator must keep track of the positions and velocities and make accurate future predictions. Here, the gazelle corrects its original escape direction (from B to C) in order to decrease the risk of getting caught (at A).

including readings from IMUs[3] and encoders to measure body acceleration and orientation and rotations in joints respectively. The resulting dynamic position and orientation of the whole body need to be inferred by combining both idiothetic and alleothetic information coming from these indirect measurements.

Incorporating uncertainty

Unfortunately, there is uncertainty in both the motions as well as measurements. Thus, we enrich our previous model with a more complete description that incorporates noise terms into the dynamics and measurements:

$$\vec{s}_t = f(\vec{s}_{t-1}, \vec{a}_{t-1}) + \eta_t \quad (3)$$

$$\vec{o}_t = g(\vec{s}_t) + \varepsilon_t \quad (4)$$

where the noise terms η_t and ε_t are independent random variables.

Probabilistic representation

The noise terms can be viewed as random variables drawn from some underlying probability distribution. Thus, Eqs. (3) and (4) are more conveniently described in terms of the conditional distributions of the noise terms:

$$\vec{s}_t \sim p(\vec{s}_t | \vec{s}_{t-1}, \vec{a}_{t-1}) \quad (5)$$

$$\vec{o}_t \sim p(\vec{o}_t | \vec{s}_t) \quad (6)$$

For instance, a state evolution with Gaussian noise and no actions and measurements would result in Brownian motion of the state over time. Together, (5) and (6) specify a dynamic Bayesian graphical model as shown in Figure 2 [5].

[3] IMUs (Inertial Measurement Units) are electronic devices that measure the velocity, orientation and gravitational forces.

Figure 2

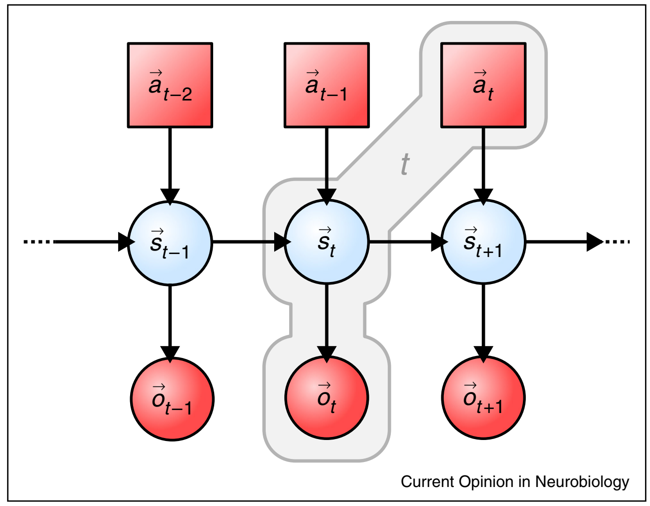

Dynamic Bayesian graphical model. This model characterizes the evolution of a hidden state $\vec{s}_t$ subject to the influence of an action $\vec{a}_t$. At each time step, the hidden state emits an observation $\vec{o}_t$. The grey area highlights the variables involved in time step t.

Belief states

According to the probabilistic view, the state $\vec{s}_t$ can be seen as being drawn from an underlying density $\pi(\vec{s}_t)$. This distribution is known as the belief state. Uncertainty in specifying the actual state is manifested in the entropy of the belief state. Consider the situation when the state is the pose of an object in two-dimensional space. The simplest specification of the pose state would consist of three variables, the two-dimensional translational position (x_t, y_t) along with the heading of the object θ_t. In this case, the belief state would be a distribution over these three components $\pi(x_t, y_t, \theta_t)$. Figure 3 shows an illustration of how a potential belief state may look at a particular time, and how it may evolve over time.

Figure 3

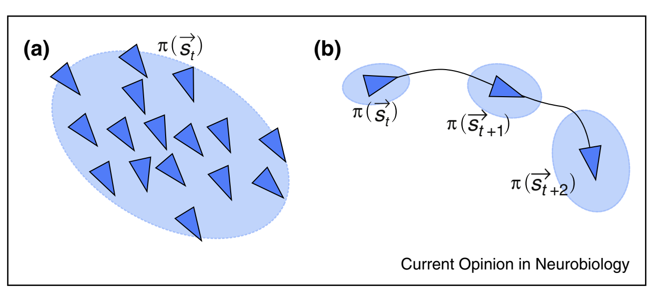

(a) Belief state representing possible poses, consisting of different locations and heading angles. **(b)** Propagation of belief state over time.