Available online at www.sciencedirect.com

ScienceDirect

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**



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:

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

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:

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

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**

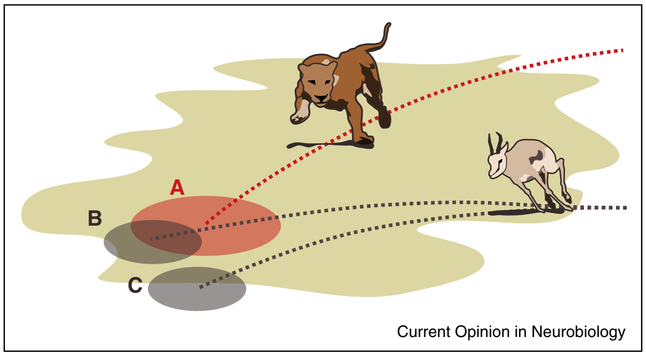


A successful escape from the predator must keep track of the positions and velocities and make accurate future predictions. Here, the gazelle corrects its original escape direction (from B to C) in order to decrease the risk of getting caught (at A).

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:

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

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

where the noise terms $\eta_t$ and $\varepsilon_t$ 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:

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

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

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**

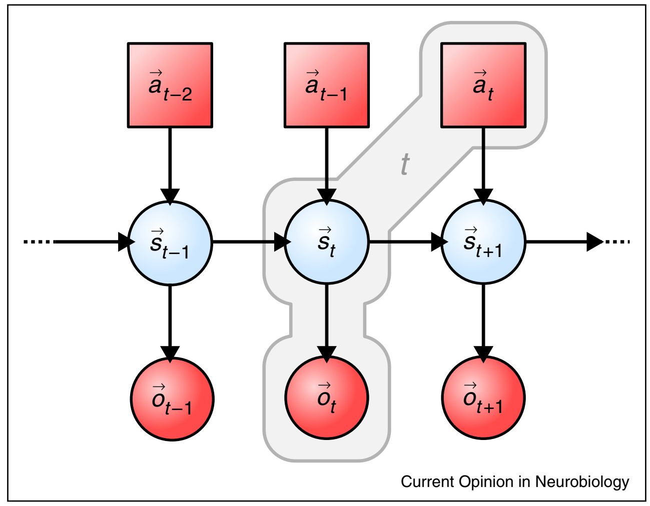


Dynamic Bayesian graphical model. This model characterizes the evolution of a hidden state $\vec{s}_t$ subject to the influence of an action $\vec{a}_t$. At each time step, the hidden state emits an observation $\vec{o}_t$. The grey area highlights the variables involved in time step $t$.

## 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 $\theta_t$. 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**

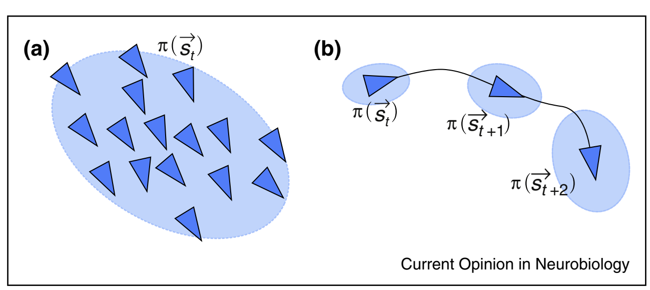


**(a)** Belief state representing possible poses, consisting of different locations and heading angles. **(b)** Propagation of belief state over time.

**Figure 4**

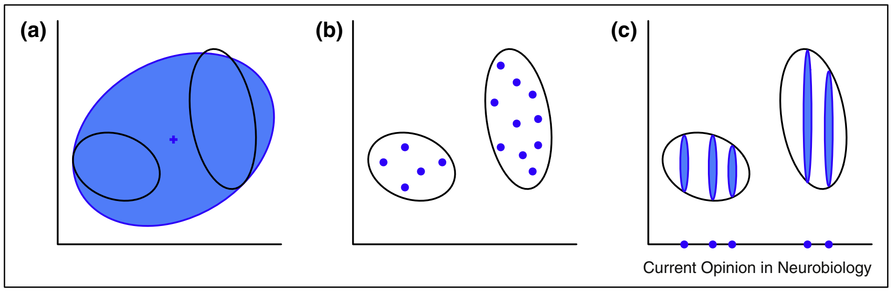


Comparison of belief representations: **(a)** Kalman filter; **(b)** particle filter; **(c)** Rao-Blackwellized particle filter.

## Estimating states versus belief states

What is the advantage of tracking the belief state rather than a single state vector, that is to say, a point estimate? If one simply estimates the state at each time, then any information about the underlying uncertainty in the estimate is lost. This can be very risky if the uncertainty has a multimodal evolution. For instance, say the predator hides behind one of two possible bushes. Keeping track of a point estimate would correspond to ignoring one of the possible locations — a very dangerous assumption for the prey if it is looking for a safe hideout. Proper accounting of uncertainty is needed if a future measurement can be used to resolve an ambiguity at a previous time. Thus, it is better to track the uncertain estimate in the form of the full distribution of the belief state.

## Probabilistic Bayes filter

Belief states are tracked using filters. A recursive filter calculates a belief state $\pi(\vec{s}_t)$ corresponding to the probability of the state $\vec{s}_t$ conditioned on all the previous actions $\vec{a}_{1:t-1}$ and observations $\vec{o}_{1:t-1}$, that is,

$$\pi(\vec{s}_t) = p(\vec{s}_t|\vec{a}_{1:t-1}, \vec{o}_{1:t-1}).$$

The next belief state is then calculated in two steps: given the belief of the state $\vec{s}_t$, we first update it using the latest observation and then propagate it using the dynamics. According to Bayes' rule, the posterior belief state $\pi(\vec{s}_t)$ after receiving the observation $\vec{o}_t$ is:

$$p(\vec{s}_t|\vec{a}_{1:t-1}, \vec{o}_{1:t}) = \frac{p(\vec{o}_t|\vec{s}_t)\pi(\vec{s}_t)}{\int d\vec{s}_t{}' \, p(\vec{o}_t|\vec{s}_t{}')\pi(\vec{s}_t{}')} \tag{7}$$

On the other hand, the predicted next state for a given action $\vec{a}_t$ is obtained by marginalizing over the current state:

$$\pi(\vec{s}_{t+1}) = \int d\vec{s}_t \, p(\vec{s}_{t+1}|\vec{s}_t, \vec{a}_t) p(\vec{s}_t|\vec{a}_{1:t-1}, \vec{o}_{1:t}). \tag{8}$$

Thus, by simply multiplying and convolving the previous time belief state with the dynamics and measurement likelihoods, the belief state can be recursively updated in time.

In general, the implementation of an exact Bayes filter is intractable. As in standard Bayesian inference, even its approximation can be very hard; for instance, the approximation of a Bayesian network is NP-hard [6]. Because of this, in robotics, there are essentially a limited number of approaches that work in practice, in the sense that they lead to efficient, real-time tracking schemes: Kalman filters [7,8], particle filters [9•,10] and Rao-Blackwellized particle filters [11–13].

## Kalman filters and particle filters

One canonical example of belief state filtering is the Kalman filter. In the Kalman filter, the belief state is described by a multivariate normal distribution:

$$\pi(\vec{s}_t) = N(\vec{\mu}_t, \Sigma_t) \tag{9}$$

with mean $\vec{\mu}_t$ and covariance $\Sigma_t$ (Figure 4a). When the motion and measurement models are both linear with Gaussian noise, the belief state will then remain a normal distribution over time. The Kalman filter equations then describe how the mean and covariance parameters change according to Eqs. (7) and (8).

However, with either nonlinear motion and measurement models, or with non-Gaussian noise, the belief state will no longer be a simple Gaussian density. In fact, it may become highly multimodal and non-ellipsoidal in shape. One method to handle such belief states is with a particle filter. In a particle filter, the belief state is represented using a set of weighted samples:

$$\pi(\vec{s}_t) = \sum_i w_t^{(i)} \delta(\vec{s}_t - \vec{s}_t^{(i)}), \tag{10}$$

where each sample $s_t^{(i)}$ ($i = 1, 2, \ldots, I$) follows a simulated trajectory and where the weights $w_t^{(i)}$ describe their probabilities (Figure 4b). Sequential Monte Carlo sampling techniques, such as *Sequential Importance Resampling* (SIR), are then used to generate new belief state samples over time in accordance with Eqs. (7) and (8) to track arbitrary belief states. For instance, the *bootstrap filter*, a variant of SIR, proceeds as follows:

1 *Dynamics:* Propagate the $I$ particles from $\vec{s}_{t-1}^{(i)}$ to $\vec{s}_t^{(i)}$ using the dynamical model:

$\vec{s}_t^{(i)} \sim p(s_t|\vec{s}_{t-1}^{(i)}, \vec{a}_{t-1})$.

2 *Update:* Update the weights using the observation $\vec{o}_t$:

$w_t^{(i)} \propto w_{t-1}^{(i)} p(\vec{o}_t|\vec{s}_t^{(i)})$.

3 *Resampling:* If the particles are degenerate, then resample them. For instance, this can be done by first checking whether the effective number of particles $N_{\text{eff}}$, estimated as[4]

$$N_{\text{eff}} = \frac{\left(\sum_i w_t^{(i)}\right)^2}{\sum_i \left(w_t^{(i)}\right)^2},$$

is smaller than a given threshold $N_0$ . If this is the case, then we draw $I$ new particles from the current particle set with probabilities proportional to their weights. Finally, we reset the weights of the new particles to a uniform distribution.

This procedure is repeated in each time step. The purpose of the resampling step is to avoid degeneracy, i.e. that the probability mass is concentrated on a few particles.

### Factorized representations

There are situations when neither Kalman filters nor particle filters can adequately represent the belief due to representational or efficiency concerns respectively. In such cases, there is the possibility of representing belief states as a hybrid between Kalman filters and particle filters, where the former capture the linear and the latter the non-linear aspects of the dynamics.

These hybrid representations are justified by a procedure known as *Rao-Blackwellization*. Sometimes the state vector can be partitioned into two separate subspaces of components $\vec{s}_t = \{\vec{s}_t^1, \vec{s}_t^2\}$, in which case the belief can be factorized as:

$$\pi(s_t) = \pi(s_t^2|s_t^1)\pi(s_t^1) \tag{11}$$

Given this partition, a belief state can now be maintained using different representations to track the various components. For example, a particle filter can be used for $\pi(s_t^1)$, while a Gaussian model can be used to represent the conditional beliefs $\pi(s_t^2|s_t^1)$, yielding an evolving mixture of Gaussians model [14] (Figure 4c).

[4] Here we estimate the effective number of particles using the Rényi entropy of order $\alpha = 2$.

### Performance in high dimensions

The recursive Bayesian filters presented above scale differently with higher dimensions. The Kalman filter extends smoothly as long as the dynamics is linear, but will be strongly biased when the dynamics are nonlinear. It is well-known that particle filters fail in very high dimensional systems due to the ‘‘curse of dimensionality’’ [15]. In particular, it has been shown that, unless the ensemble size is exponentially large in the variance in the observation log likelihood, the particle filter update suffers from a ‘‘weight collapse’’ in which a single particle claims all the probability mass for itself [16]. Rao-Blackwellized filters constitute a compromise between both approaches where exploiting knowledge about the structure of the correlations in the belief state yields a more efficient representation.

## Navigation example

We consider the simple example of a robot navigating in a two-dimensional world. We assume that there are discrete landmarks whose 2D coordinates are known which can be observed when they are in range. In this case, we consider a pose state consisting of three components, the two translational degrees of freedom, along with a heading direction.

Suppose that the robot is initially located very close to a landmark A, but the initial heading is unknown, e.g. due to some prior adverse environmental conditions. In this case, the pose belief states consists of well-localized positions, but completely unknown headings (Figure 5a). Now the robot takes a few steps forward. After moving, the robot turns out to have uncertainty over both the pose and the heading, and the corresponding belief state consists of locations in an annulus surrounding the initial landmark (Figure 5b). However, both heading and position will be highly correlated in this case.

After moving and an ensuing delay, the robot observes directly ahead a second landmark whose position is known. With this observation, the pose belief state will shrink onto the positions and heading angles that are consistent with the proper poses in the map (see Figure 5). This illustrates how keeping track of the full uncertainty of the belief state, including the joint dependencies among variables, can be crucial in order the

Figure 5

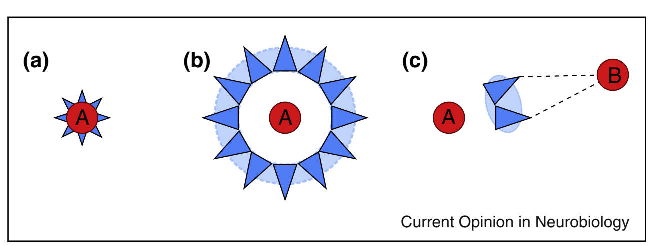


Evolution of pose belief state with uncertain initial heading.

resolve ambiguities. Some algorithms take this idea even a step further: *Simultaneous Localization and Mapping* (SLAM) algorithms assume that neither position nor landmarks are known, and they specialize in simultaneously updating both pose and map variables.

An example is given by [17••]. There, the authors presented a computational model based on spatially-responsive cells and functional engineering principles that display similar characteristics to rodent grid cells, and show how these cells may maintain and propagate multiple probabilistic estimates of pose, enabling the correct pose estimate to resolve even without uniquely identifying perceptual cues.

## Biological implications

The ability to make educated predictions about the future state of the world is essential for humans and other animals to successfully plan and perform actions. Full probabilistic representations and proper updating of belief states (Bayes filters) would be beneficial in the sense that it provides the optimal basis for making such predictions. Like robots, animals are limited in their representational and computational resources and thus it is unclear to what degree the animal brain is performing full Bayesian filtering and to what degree it pursuits computational shortcuts and approximations. This is an open question in the behavioral and computational neurosciences and in the following, we discuss behavioral evidence for belief state representations in biological brains, potential experiments to further test this hypothesis, and the implications for potential neural representations of such belief states.

### Behavioral evidence of probabilistic belief state representations

Recent studies in perceptual (e.g. [18,19]), and sensorimotor (e.g. [3]) integration have demonstrated that humans are able to properly combine uncertain sensory evidence with prior beliefs. While these experiments have demonstrated the ability to update prior beliefs with current sensory evidence, the involved tasks typically required subjects to only perform one inference step at a time that then was repeated over many trials. As a consequence, the data of these studies do not allow us to distinguish whether subjects retain a full probabilistic representation of belief states or whether they perform (Bayesian) state estimates because both lead to identical predictions in such one-step inference tasks. To address this question, we need to design experiments that require subjects to maintain continuously updated *dynamic* state representations, which will lead to different outcomes for representations that reflect belief states versus state estimates.

One suitable class of experiments addresses so-called "intuitive physics" [20,21], the ability of humans to have an intuitive understanding of the laws of physics. For example, this understanding enables a human observer to predict the trajectory of a moving object solely based on a partial observation of the object's initial state (position, motion). While humans make systematic errors in such predictions recent studies have argued that these prediction errors can be explained as the result of a rational Bayesian observer who applies Newton's laws to the noisy sensory observations of the initial object states. These models assume that observers mentally track the trajectory of the occluded object [22] or simulate it in order to predict future outcomes before they happen [23••]. Both, tracking and mental simulations are equivalent to the propagation of state estimates and do not require to maintain and update full belief states. An experiment that could potentially distinguish whether human observer maintain full belief states or not would need to probe situations for which predictions based on full belief states will substantially differ from predictions based on propagated state estimates. Figure 6 outlines such an experiment.

**Figure 6**

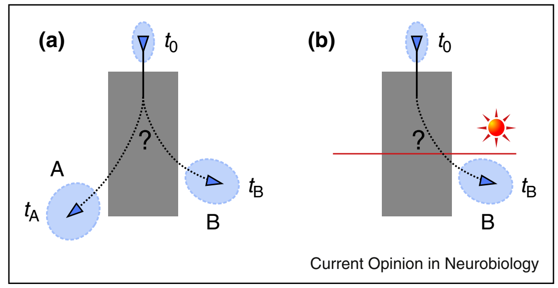


Psychophysical experiment to test belief state representations. **(a)** Based on the initial motion trajectory of a moving object at time $t_0$ an observer has to predict when and where the object reappears. Here, the object can undergo two very distinct trajectories A and B; when and where the object reappears depends on the initial state of the object (motion, position). Maintaining full belief states would allow the observer to provide educated predictions involving both potential trajectories. In particular, they would be able to update and refine their belief states when signaled halfway through which of the two potential trajectories the object is following **(b)**. In contrast, if the observer was mentally tracking the object's trajectory (*updating a state estimate*) they would perform very badly in predicting the object's future position if they were to track the wrong trajectory.

### Neural representations of belief states

There are several existing models of neural representations of static beliefs: [24–27]. However, dynamic belief states need to incorporate the continuous updating of beliefs in real time. How would these beliefs be represented and computed by neurons? Is it possible that different cortical circuits contain hypotheses about different potential states?

One clue may lie in the correlations between different components of the state being estimated. For instance, our navigation example showed the advantage of tracking the full belief state with correlations between positional and orientational components. For animal navigation, given that head directions and place locations appear to be represented in different brain areas (Enthorinal cortex [28] and Hippocampus [29], respectively), it would be interesting to see if there were any correlations in the neural activities between the two areas as would be predicted by a Rao-Blackwellized particle filter. Perhaps someday connectomics will be able to show whether there is a topographical organization in the connections between these two areas.

Another possibility is that only an approximation of the full belief state is maintained by neurons. A recent computational model considers the simultaneous estimation of eye position with noisy microsaccades and an image by independently factorizing the belief state into its marginal distributions [30••]. Such approximations may work in most cases, but they would fail to provide accurate estimates under certain conditions. Predictions of interesting failure cases can be provided given different approximation methods for dynamic belief states representations.

## Conclusions

The need for neural systems to accurately estimate and track dynamic states is clear. Due to noise and uncertainty, a solution necessarily involves tracking a dynamic belief state rather than a single state over time; as we have argued above, in some cases failing to do so might have fatal consequences. In robotics, there are some canonical algorithms for tracking dynamic belief states, including Kalman filters, particle filters, and state space decompositions in Rao-Blackwellized filters. We speculate that biological systems may track analogous representations in the brain. Exactly how these belief states are maintained and computed using neurons, however, is an important open problem. Future experiments and analysis may someday point to more concrete evidence of such representations implemented in neural circuits.

## References and recommended reading

Papers of particular interest, published within the period of review, have been highlighted as:

- • of special interest
- •• of outstanding interest


1. Pfeifer R, Scheier C: *Understanding Intelligence*. Boston, MA, USA: MIT Press; 1999, .

2. Todorov E: **Optimality principles in sensorimotor control**. *Nat Neurosci* 2004, **7**:907-915.

3. Koerding KP, Wolpert DM: **Bayesian integration in sensorimotor learning**. *Nature* 2004, **427**:244-248.

4. Shadmehr R, Smith MA, Krakauer JW: **Error correction, sensory prediction, and adaptation in motor control**. *Annu Rev Neurosci* 2010, **33**:89-108.

5. Kollar Daphne, Friedman Nir: *Probabilistic graphical models: principles and techniques*. London, England: The MIT Press; 2009.

6. Cooper GF: **The computational complexity of probabilistic inference using Bayesian belief networks**. *Artif Intell* 1990, **42**:393-405.

7. Kalman RE: **A new approach to linear filtering and prediction problems**. *Trans ASME-J Basic Eng* 1960, **82(Series D)**:35-45.

8. Durrant-Whyte H, Bailey T: **Simultaneous localization and mapping. Part I. Robotics & automation magazine**. *IEEE* 2006, **13**:99-110.

9. • Cappe O, Godsill SJ, Moulines E: **An overview of existing methods and recent advances in sequential Monte Carlo**. *Proc IEEE* 2007, **95**:899-924.

The authors present a thorough review of existing sequential Monte Carlo methods.

10. Fox D, Thrun S, Burgard W, Dellaert F: **Particle filters for mobile robot localization**. *Sequential Monte Carlo Methods in Practice*. New York, NY, USA: Springer; 2001, 401-428.

11. Chen R, Liu JS: **Mixture Kalman filters**. *J R Stat Soc B* 2000, **62**:493-508.

12. Doucet A, De Freitas N, Murphy K, Russell S: **Rao-Blackwellised particle filtering for dynamic Bayesian networks**. In *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers Inc.; 2000:176-183.

13. Montemerlo M, Thrun S, Koller D, Wegbreit B: **FastSLAM: a factored solution to the simultaneous localization and mapping problem**. *AAAI/IAAI*. 2002:593-598.

14. Vernaza P, Lee DD: **Rao-Blackwellized particle filtering for 6-DOF estimation of attitude and position via GPS and inertial sensors**. In *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006*; *IEEE: 2006:1571-1578*.

15. Bengtsson T, Bickel P, Li B: **Curse of dimensionality revisited: Collapse of the particle filter in very large scale systems**. In *Probability and Statistics: Essays in Honor of David A. Freedman, vol. 23*. Edited by Nolan D, Speed T. Institute of Mathematical Statistics; 2008:316-334.

16. Snyder C, Bengtsson T, Bickel P, Anderson J: **Obstacles to high-dimensional particle filtering**. *Mon Weather Rev* 2008, **136**:4629-4640.

17. •• Milford MJ, Wiles J, Wyeth GF: **Solving navigational uncertainty using grid cells on robots**. *PLoS Comput Biol* 2010, **6**:e1000995.

This work illustrates a computational model based on rodent spatially-responsive cells and functional engineering principles that display similar characteristics to rodent grid cells that resolve estimates of pose.

18. Stocker AA, Simoncelli EP: **Noise characteristics and prior expectations in human visual speed perception**. *Nat Neurosci* 2006 April:578-585.

19. Ernst MO, Banks MS: **Humans integrate visual and haptic information in a statistically optimal fashion**. *Nature* 2002, **415**:429ff.

20. Gilden DL, Proffitt DR: **Understanding collision dynamics**. *J Exp Psychol Hum Percept Perform* 1989, **15**:372-383.

21. Proffitt DR, Gilden DL: **Understanding natural dynamics**. *J Exp Psychol Hum Percept Perform* 1989 May, **15**:384-393.

22. Smith KA, Battaglia PW, Vul E: **Consistent physics underlying ballistic motion prediction**. In *Proceedings of the 35th Annual Meeting of the Cognitive Science Society*. 2013.

23. •• Battaglia PW, Hamrick JB, Tenenbaum JB: **Simulation as an engine of physical scene understanding.**. *Proc Natl Acad Sci USA* 2013.

The authors propose the interesting hypothesis that humans predict the outcomes of complex physical interactions of objects via repeated mental simulations. They suggest that this fast sampling of predicted state estimates represent a dedicated computational processing unit (''intuitive physics engine'') similar to the graphics processors in virtual reality systems.

24. Fischer BJ: **Bayesian estimates from heterogeneous population codes**. *Proc. IEEE Int'l. Joint Conf. on Neural Networks. IEEE*. 2010.

25. Fischer BJ, Pena JL: **Owl's behavior and neural representation predicted by Bayesian inference**. *Nat Neurosci* 2011, **14**:1061-1066.

26. Ganguli D, Simoncelli EP: **Implicit encoding of prior probabilities in optimal neural populations**. Edited by Lafferty J, Williams CKI, Shawe-Taylor J, Zemel RS, Culotta A.**Adv. Neural Information Processing Systems 23 (NIPS*10)**. Cambridge, MA: MIT Press; 2011:658-666.

27. Wei X-X, Stocker AA: **Bayesian inference with efficient neural population codes**. *Lecture Notes in Computer Science, Artificial Neural Networks and Machine Learning - ICANN 2012, vol. 7552*; *Lausanne, Switzerland, September 11–14: 2012:523-530*.

28. Taube JS, Muller RU, Ranck JB Jr: **Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis**. *J Neurosci* 1990, **10**:420-435.

29. O'Keefe J, Dostrovsky J: **The hippocampus as a spatial map. preliminary evidence from unit activity in the freely-moving rat**. *Brain Res* 1971, **34**:171-175.

30. •• Burak Y, Rokni U, Meister M, Sompolinsky H: **Bayesian model of dynamic image stabilization in the visual system**. *Proc Natl Acad Sci USA* 2010, **107**:19525-19530.

The authors present a Bayesian spiking model for dynamical image stabilization that tracks and estimates the unknown trajectory of the projected image on the retina.