

The Bayes Filter

Josiah Hanna

1 Notes on Notation

1. We will write $p(\cdot|y)$ to denote the conditional distribution over some random variable given that we have observed y .
2. We will write $x \sim p(\cdot|y)$ to denote that variable x has its outcome distributed according to the conditional distribution, $p(\cdot|y)$.

2 Robot Interaction Model

We will formalize a robot interacting with its environment as a discrete-time process. At time $t = 0$, the environment is in an initial state, x_0 . The robot chooses a control action, u_1 , and then the environment transitions to state $x_1 \sim p(\cdot|x_0, u_1)$. The robot then receives an observation, $z_1 \sim g(\cdot|x_1)$. The process then repeats with the robot choosing a new action. We will refer to the full sequence of states up to and including time t as $x_{0:t}$ and similarly for actions, $u_{1:t}$ and observations, $z_{1:t}$. Following common notation (e.g., in the textbook *Probabilistic Robotics*, the state index begins at $t = 0$ and control and observation indices begin at $t = 1$).

3 State Estimation

Robots only perceive observations from their sensors. In general, observations provide incomplete and potentially noisy information about the state of the world. However, robots need to infer the state of the world to make good decisions.

Why does the robot need to estimate state to make good decisions? A single observation is likely missing some critical information. For example, autonomous vehicles cannot observe pedestrians that are currently behind other vehicles even if their presence is known. An alternative approach to state estimation is to attempt to reason based on the full history of observations and controls. However, the number of possible observation/control histories is exponential in the time-step t and so will quickly become intractable to base decision-making on. The state strikes a balance between these extremes by summarizing the useful parts of the observation history into a compact representation containing all aspects that affect the future.

A final reason that state must be estimated is that it is constantly changing as the robot moves about its environment. The state may evolve if the robot takes action or not. Choosing not to act is another type of action.

Moreover, we want to be probabilistic and infer a belief distribution over states instead of simply guessing a single state as the current one. Doing so allows the robot to represent uncertainty about the state of the world.

4 Probabilistic State Estimation

A naïve approach to state estimation might simply estimate which state is most likely given the robot’s observations and actions up to time t . This approach is limited as it chooses a single state with certainty and thus does not allow the robot to represent uncertainty about the state of the world. Instead, we will estimate a *belief distribution*, which is a probability distribution over possible states in the world given everything that has been previously observed. We will write $\text{bel}(x_t)$ to denote the probability of state x_t under this belief distribution. Formally, $\text{bel}(x_t)$ is the posterior $\Pr(x_t|z_{1:t}, u_{1:t})$. This belief factors in all available information up to and including step t . We will find it useful to also refer to the belief, $\overline{\text{bel}}(x_t)$, which we define as $\Pr(x_t|z_{1:t-1}, u_{1:t})$ or the posterior distribution of x_t before the final observation has been received and factored in.

4.1 Naive Belief Computation

Bayes rule gives us a straightforward – but computationally complex – means to compute the robot’s belief at time t . Using Bayes Rule, we obtain:

$$\text{bel}(x_t) = \Pr(x_t|z_{1:t}, u_{1:t}) = \eta \sum_{x_{1:t-1}} \Pr(x_{1:t}|u_{1:t}) \Pr(z_{1:t}|x_{1:t}),$$

where $\sum_{x_{1:t-1}}$ is a summation over all possible state sequences from time 1 to time $t-1$ and η is a constant that ensures we have a valid probability distribution. With full knowledge of the state and observation probabilities, we can compute $\Pr(x_{1:t}|u_{1:t}) = p(x_1) \prod_{i=2}^t p(x_i|x_{i-1}, u_{i-1})$ and $\Pr(z_{1:t}|x_{1:t}) = \prod_{i=1}^t g(z_i|x_i)$.

Unfortunately, this approach to belief computation becomes intractable as the interaction sequence grows because the summation above will have an exponential number of terms. It is also memory-intensive as the robot would have to keep around the entire sequence of observations and past control actions in order to update its belief at each moment in time.

5 Recursive Belief Computation

The key idea of the Bayes filter is to update the robot’s belief recursively as each new control is taken and a new observation is received. The Bayes filter requires two steps to update the current belief to a belief that incorporates u_t and z_t .

1. **Prediction.** The robot computes a new belief, $\overline{\text{bel}}(x_t)$ by predicting the effect of u_t on $\text{bel}(x_{t-1})$. In general, the step *increases* the robot's uncertainty about its true state.
2. **Correction.** The robot uses the information in the new observation, z_t , to correct its prediction. This step amounts to applying Bayesian inference with $\overline{\text{bel}}(x_t)$ as the prior. In general, this step *decreases* the robot's uncertainty about its true state.

Algorithm 3 provides pseudocode for the complete belief update of the Bayes filter. When used sequentially on robot data, $u_{1:t}$ and $z_{1:t}$, the Bayes filter algorithm will compute the same posterior $\text{bel}(x_t)$ as the direct use of Bayes rule on all of the data at once. The proof of this fact is left as an exercise to the reader or office hours.

Algorithm 1 Bayes Filter

```

1: Input: Previous belief  $\text{bel}(x_{t-1})$ , control input  $u_t$ , measurement  $z_t$ 
2: Output: Updated belief  $\text{bel}(x_t)$ 
3: procedure BAYESFILTER( $\text{bel}(x_{t-1})$ ,  $u_t$ ,  $z_t$ )
4:   // Prediction step
5:   for all  $x_t$  do
6:      $\overline{\text{bel}}(x_t) \leftarrow \sum_{x_{t-1}} p(x_t | u_t, x_{t-1}) \cdot \text{bel}(x_{t-1})$ 
7:   end for
8:   // Correction step
9:   for all  $x_t$  do
10:     $\text{bel}(x_t) \leftarrow \eta \cdot g(z_t | x_t) \cdot \overline{\text{bel}}(x_t)$  //  $\eta$  is the normalizing constant.
11:  end for
12:  return  $\text{bel}(x_t)$ 
13: end procedure

```

6 Limitations of the Bayes Filter

There are several limitations of the Bayes filter and it is often more of theoretical interest than real application. Later this week and the following week, we will introduce the Kalman filter family and Particle filters which are significantly more applicable to real robots than Bayes filters.

In particular, the Bayes filter has the following limitations:

1. The computational complexity of exact summation. The prediction step often involves summing over the entire state space or integrating if the state space is continuous. Doing so is computationally infeasible for high-dimensional problems. The same limitation applies to computing the normalization constant, η , in the update step. The high computational complexity limits scalability and may make the filter too slow to use for real-time belief updates.

2. Markov assumption. The validity of the Bayes filter relies upon the Markov assumption, i.e., that $p(x_t|x_{t-1}, u_t) = p(x_t|x_{0:t-1}, u_{1:t})$. This assumption is unlikely to hold in practice.
3. Assumption of known models. The Bayes filter requires accurate knowledge of the robot’s state transition model, p , and sensor model, g . Errors or approximations in these models can lead to poor performance.
4. Scalability. The Bayes filter is difficult to implement with large or high dimensional state spaces. Approximations can be used but these may make it more likely the Markov assumption is violated.

7 Bayes Smoother

In some applications, we are not only interested in $\text{bel}(x_t)$ for the current time t but in $\text{bel}(x_t)$ for all t from $t = 0$ to some $T = T$. The Bayes filter can almost solve this problem as it produces a belief for each step $\text{bel}(x_t)$. However, it is not the optimal approach as it neglects to use information available in $z_{t:T}$ to refine the belief $\text{bel}(x_t)$. To see why observations after t can help infer x_t , consider listening to a person talking in a noisy room. You are unsure if they just said “The price will rise” or “The prize will rise” due to the similar sound of “price” and “prize.” As they continue talking, you hear “due to inflation” and so you infer that the more likely utterance is “The price will rise due to inflation.” In this case, you have used the most recent information to refine a belief about something that happened in the past.

The Bayes filter only considers information available up to time t when inferring $\text{bel}(x_t)$. To incorporate later observations, we will turn to the Bayes smoother. The Bayes smoother is a post-processing procedure that is ran over the sequence of beliefs computed by the Bayes filter. Pseudocode is provided in Algorithm 2.

8 Kalman Filter for Continuous State Spaces

We just introduced the Bayes Filter as a means to compute a robot’s posterior belief, $p(x_t|z_{1:t}, u_{1:t})$, online, using recursive updates in which computation does not increase with t . The introduction assumed a finite state-space since we had to sum over all possible states in the prediction step. For continuous state-spaces, at every time-step, the Bayes Filter alternates the two steps:

1. **Prediction:** The prior belief about the state is updated based on the system dynamics:

$$\overline{\text{bel}}(x_t) \leftarrow \int_{x_{t-1}} p(x_t | u_t, x_{t-1}) \cdot \text{bel}(x_{t-1}) dx_{t-1}.$$

2. **Update:** The prior belief is refined using sensor observations.

$$\text{bel}(x_t) \leftarrow \eta \cdot p(z_t | x_t) \cdot \overline{\text{bel}}(x_t).$$

Algorithm 2 Bayes Smoother

```
1: Input: Beliefs  $\text{bel}(x_t)$  for  $t = 1 \dots T$ , motion model  $p(x_{t+1} \mid x_t, u_{t+1})$ 
2: Output: Smoothed beliefs  $\text{bel}'(x_t)$  for  $t = 1 \dots T$ 
3: procedure BAYESMOOTH( $\text{bel}(x_t), u_t$ )
4:   // Initialization: Start with the final belief
5:    $\text{bel}'(x_T) \leftarrow \text{bel}(x_T)$ 
6:   // Backward smoothing step
7:   for  $t = T - 1$  to 1 do
8:     for all  $x_t$  do
9:        $\text{bel}'(x_t) \leftarrow \text{bel}(x_t) \cdot \sum_{x_{t+1}} \frac{p(x_{t+1} \mid x_t, u_{t+1}) \cdot \text{bel}'(x_{t+1})}{\text{bel}(x_{t+1})}$ 
10:    end for
11:  end for
12:  return  $\text{bel}'(x_t)$  for  $t = 1 \dots T$ 
13: end procedure
```

When the number of possible states is finite, the integral is replaced with a summation. A main limitation of Bayes Filter is that the updates in the prediction and update step are not tractable if the number of possible states is large or states are continuous. While the Bayes Filter is intractable for continuous state spaces in general, there are special cases where it can be applied even with continuous states. We will next introduce the Kalman Filter which will enable exact implementation of a Bayes Filter for the special case of a linear state transition and observation model with Gaussian uncertainty. Furthermore, the Kalman Filter assumes that the belief distribution is Gaussian, which allows it to be fully described by its mean vector, μ , and covariance matrix, Σ . These assumptions enable the use of linear algebra to efficiently compute the optimal state estimate through the prediction and update steps, making the Kalman Filter a computationally efficient realization of the Bayes Filter. The Kalman Filter is an optimal recursive algorithm used for estimating the state of a dynamic system from a series of noisy measurements. In addition to robotics, it is widely applied throughout control systems.

9 Kalman Filter System Model Assumptions

The basic Kalman Filter makes specific assumptions about the robot's environment. Specifically, we assume linear Gaussians for both state transitions and sensor observations. Formally, we have that:

$$x_t = Ax_{t-1} + Bu_t + w_t, \tag{1}$$

$$z_t = Hx_t + v_t, \tag{2}$$

where:

- x_t is the vector of state variables at time t ,

- u_t is the vector of control inputs at time t ,
- z_k is the vector of sensor observations at time t ,
- A is the state transition matrix,
- B is the control matrix,
- H is the observation matrix,
- $w_t \sim \mathcal{N}(0, Q)$ is noise in the robot's state transition,
- $v_t \sim \mathcal{N}(0, R)$ is noise in the robot's observation.

Note that we are assuming both transition noise and observation noise have Gaussian distributions with mean $\mathbf{0}$ and covariance matrices Q and R respectively. Ignoring the noise, we can see that both state transitions and observations are linear in x_t and u_t . Taken together, we are assuming that state transitions are produced by a linear transformation of the state and control at time t followed by adding Gaussian noise. A similar process is used to determine the observation at time t . This is the *linear-Gaussian* assumption.

Another way of writing the linear-Gaussian assumption is that $p(x_t|x_{t-1}, u_t) = \mathcal{N}(x_t; Ax_{t-1} + Bu_t, Q)$ and $p(z_t|x_t) = \mathcal{N}(z_t; Hx_t, R)$.

10 Kalman Filter

We are now ready to instantiate the Bayes Filter under the specific assumptions from the previous section. One final assumption we must make is that the robot's initial belief about its state is Gaussian: $\text{bel}(x_0) = \mathcal{N}(x; \mu_0, \Sigma_0)$. This assumption and the assumption of a linear-Gaussian environment means that the robot's belief will maintain a Gaussian form under Kalman Filter updates: $\text{bel}(x_t) = \mathcal{N}(x; \mu_t, \Sigma_t)$.

10.1 Prediction Step

First, we compute the posterior belief about the next state given the robot's control, u_t , but before incorporating the observation:

$$\bar{\mu}_t = A\mu_{t-1} + Bu_t \tag{3}$$

$$\bar{\Sigma}_t = A\Sigma_{t-1}A^\top + Q \tag{4}$$

where:

- $\bar{\mu}_t$ is the predicted mean of the belief distribution before the sensor observation is incorporated.
- $\bar{\Sigma}_t$ is the predicted covariance of the belief distribution.

10.2 Update Step

We then incorporate the sensor observation to produce the final posterior belief, $\mathbf{bel}(x_t)$.

$$K_t = \bar{\Sigma}_t H^T (H \bar{\Sigma}_t H^T + R)^{-1} \quad (5)$$

$$\mu_t = \bar{\mu}_t + K_t (z_t - H \bar{\mu}_t) \quad (6)$$

$$\Sigma_t = (I - K_t H) \bar{\Sigma}_t \quad (7)$$

where:

- K_t is known as the Kalman gain. The Kalman gain balances how much to trust the new observation instead of the prediction alone. If R is large then the measurement is very noisy and hence less trustworthy; consequently, K_t will be small.
- μ_t is the updated state estimate.
- Σ_t is the updated covariance.

Following these updates, we have that the robot's belief will be $\mathbf{bel}(x_t) = \mathcal{N}(x_t; \mu_t, \Sigma_t)$.

This prediction and update step are an exact implementation of a Bayes Filter under specific assumptions on the system model and initial belief. That is, the Kalman prediction and update steps are obtained from plugging the linear Gaussian system model into the Bayes Filter update steps (the actual proof is a bit more involved).

Algorithm 3 Kalman Filter

- 1: **Input:** Previous belief represented as a Gaussian with mean, μ_{t-1} , and covariance, Σ_{t-1} , control input u_t , measurement z_t
 - 2: **Output:** Updated belief represented as a Gaussian with mean, μ_t , and covariance Σ_t .
 - 3: **procedure** KALMANFILTER($\mu_{t-1}, \Sigma_{t-1}, u_t, z_t$)
 - 4: // Prediction step
 - 5: $\bar{\mu}_t = A\mu_{t-1} + Bu_t$
 - 6: $\bar{\Sigma}_t = A\Sigma_{t-1}A^\top + Q$
 - 7: // Correction step
 - 8: $K_t = \bar{\Sigma}_t H^T (H \bar{\Sigma}_t H^T + R)^{-1}$ // Compute Kalman gain
 - 9: $\mu_t = \bar{\mu}_t + K_t (z_t - H \bar{\mu}_t)$
 - 10: $\Sigma_t = (I - K_t H) \bar{\Sigma}_t$
 - 11: **return** μ_t, Σ_t
 - 12: **end procedure**
-

10.3 Strengths and Weaknesses

The Kalman Filter has the following strengths:

- The Kalman Filter inherits the optimality of a Bayes Filter assuming the system is actually linear with Gaussian noise.
- Like a Bayes Filter, the Kalman Filter is recursive which implies that computation does not grow with t . Furthermore, the linear algebra computations can be done in Polynomial time and are typically fast with modern linear algebra libraries (e.g., Eigen or Numpy).
- Like a Bayes Filter, the Kalman Filter provides uncertainty estimates through the covariance, Σ_t . If the determinant of Σ_t is large then the robot is uncertain about the true state. If the determinant of Σ_t is small then the robot is more certain that μ_t is its current state.

These advantages are offset by some important weaknesses:

- The Kalman Filter assumes Gaussian noise which limits its applicability to systems with non-Gaussian noise.
- If the system is highly nonlinear then the Kalman Filter will compute a sub-optimal belief.
- With a Kalman Filter, the belief is always Gaussian which means the robot cannot represent multi-modal uncertainty. For example, the robot cannot represent the belief that it is located in one of two locations.

A Review

The following appendices provide background on some mathematical concepts used in these notes. If you are familiar with these topics, you can skip over them.

A.1 Gaussian Distributions

The Gaussian (or multi-variate normal) distribution is one of the most common probability distributions for continuous random variables. Let $x \in \mathbf{R}^d$ be a d -dimensional vector. We say that x has a Gaussian distribution if:

$$p(x) = \frac{1}{\sqrt{2\pi \det \Sigma}} e^{-\frac{1}{2}(x-\mu)^\top \Sigma^{-1}(x-\mu)},$$

where $\mu \in \mathbf{R}^d$ is called the mean vector and $\Sigma \in \mathbf{R}^{d \times d}$ is the covariance matrix. The matrix Σ must be positive-definite.

A few additional notes on the Gaussian distribution:

1. When $d = 1$, the Gaussian distribution is the familiar bell curve.

2. Under a Gaussian distribution, $p(x)$ is maximized when $x = \mu$, i.e., the mean is the most likely outcome.
3. The covariance controls the spread of $p(x)$. When the covariance is large, then $p(x)$ is spread out and values of x that are farther from μ receive more probability density. When the covariance is small then most sampled values of x will be close to μ .

For notation, we will denote Gaussian distributions as $\mathcal{N}(\mu, \Sigma)$ and the density of x under a Gaussian with mean μ and covariance Σ as $\mathcal{N}(x; \mu, \Sigma)$.