

AA274A: Principles of Robot Autonomy I

Course Notes

The previous chapter introduced a probabilistic framework that can be used to design algorithms for robot localization and state estimation. The chapter concluded with the introduction of the Bayes filter, which is a fundamental algorithm for maintaining and updating a belief distribution (a probability distribution over possible states). While the Bayes filter is generally intractable to implement in practice, it lays a mathematical foundation for the development of algorithms that can exploit structure or other approximations to generate tractable approaches. One such example is the discrete Bayes filter, which assumes that the number of possible states is finite such that the belief distribution can be represented by simply storing the probability of each state individually. This type of distribution is referred to as *non-parametric* since the belief distribution is not required to have a particular structure.

Alternatively, it is possible to develop tractable algorithms for probabilistic localization and state estimation by leveraging *parametric* belief distributions. Parametric distributions are distributions that are fully specified by a fixed number of parameters, for example Gaussian distributions are defined by the mean and covariance parameters. These *parametric filters* can generally be viewed as practical implementations of Bayes filter that exploit structure for efficiency, and include the Kalman filter family of algorithms.

15 Parametric Filters

Parametric filters are a family of algorithms for robot localization and state estimation that model the robot's belief with parametric distributions. Therefore, as the robot's state evolves and new measurement information arrives, updating the belief distribution is accomplished by simply updating the parameters that define the distribution. This can lead to practically implementable algorithms since the number of parameters is generally not too large. For example, a Gaussian distribution in one dimension only requires the specification of two parameters: the mean and standard deviation. Yet with these two parameters a probability distribution is defined over an infinite number of values! This is an example of how parametric filters exploit structure for efficiency.

15.1 Gaussian Distribution

The Gaussian distribution (also referred to as a normal distribution) is likely the most commonly used parametric distribution in many disciplines, including robotics. The probability density function for a one-dimensional (univariate) Gaussian distribution is given by:

$$p(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2}\frac{(x-\mu)^2}{\sigma^2}}, \quad (1)$$

where the parameters are the mean μ and standard deviation σ (the quantity σ^2 is referred to as the variance). A shorthand notation for saying that a random variable X is distributed according to a Gaussian (normal) distribution is $X \sim \mathcal{N}(\mu, \sigma^2)$. For the multi-dimensional case, the multivariate Gaussian distribution is defined by the probability density function:

$$p(\mathbf{x}) = \frac{1}{\sqrt{\det(2\pi\mathbf{\Sigma})}} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right), \quad (2)$$

where $\mathbf{x} \in \mathbb{R}^n$ and the parameters are the mean $\boldsymbol{\mu} \in \mathbb{R}^n$ and the covariance matrix $\mathbf{\Sigma} \in \mathbb{R}^{n \times n}$. Again, a shorthand to say a random variable X is distributed according to the multivariate Gaussian distribution is $X \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{\Sigma})$. These distributions are represented graphically for the univariate and bivariate case in Figure 1. These distributions exhibit several important

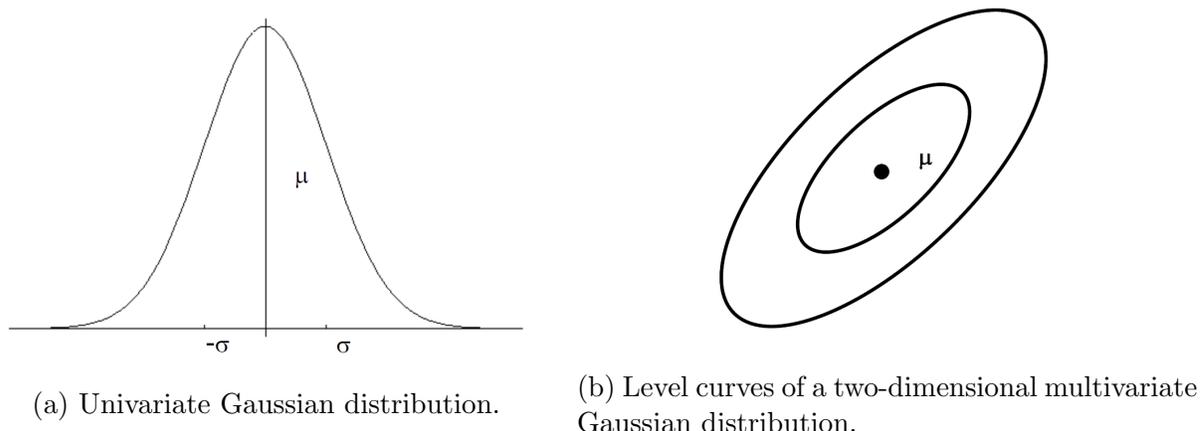


Figure 1: Univariate and multivariate Gaussian distributions.

properties which make them particularly attractive for algorithm development:

1. The affine transformation of a Gaussian random variable is also a Gaussian random variable. In particular, suppose the random variable X has a multivariate Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance $\mathbf{\Sigma}$. Then the random variable Y resulting from an affine transformation:

$$Y = AX + b,$$

also has a multivariate Gaussian distribution with expected value $A\boldsymbol{\mu} + b$ and covariance $A\mathbf{\Sigma}A^T$. In other words, if $X \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{\Sigma})$ then $Y \sim \mathcal{N}(A\boldsymbol{\mu} + b, A\mathbf{\Sigma}A^T)$.

2. The sum of two independent Gaussian random variables is also a Gaussian random variable. In particular, suppose X_1 and X_2 have multivariate Gaussian distributions with means $\boldsymbol{\mu}_1$ and $\boldsymbol{\mu}_2$ and covariances $\boldsymbol{\Sigma}_1$ and $\boldsymbol{\Sigma}_2$. Then the random variable Y given by the sum:

$$Y = X_1 + X_2,$$

also has a multivariate Gaussian distribution with expected value $\boldsymbol{\mu}_1 + \boldsymbol{\mu}_2$ and covariance $\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2$. In other words, if $X_1 \sim \mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$ and $X_2 \sim \mathcal{N}(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2)$ then $Y \sim \mathcal{N}(\boldsymbol{\mu}_1 + \boldsymbol{\mu}_2, \boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2)$.

3. The product of two Gaussian random variables is also a Gaussian random variable.

15.2 Kalman Filter

The Kalman filter is an extremely well known algorithm for state estimation that leverages the Gaussian distribution for efficiency. Unlike the discrete Bayes filter from the previous chapter, this filter can be applied to problems with *continuous* states. In particular, a multivariate Gaussian distribution is used to parameterize the belief distribution over possible states, in other words the state $\mathbf{x}_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$, and in long form this can be expressed as:

$$bel(\mathbf{x}_t) = p(\mathbf{x}_t) = \frac{1}{\sqrt{\det(2\pi\boldsymbol{\Sigma}_t)}} \exp\left(-\frac{1}{2}(\mathbf{x}_t - \boldsymbol{\mu}_t)^T \boldsymbol{\Sigma}_t^{-1}(\mathbf{x}_t - \boldsymbol{\mu}_t)\right).$$

15.2.1 Assumptions

To ensure that the belief *remains* Gaussian after the prediction and measurement update steps of the filtering algorithm, several additional assumptions are required. First, it is assumed that the initial belief $bel(\mathbf{x}_0)$ is Gaussian with $\mathbf{x}_0 \sim \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ and that the state transition model is linear and is given by:

$$\mathbf{x}_t = A_t \mathbf{x}_{t-1} + B_t \mathbf{u}_t + \boldsymbol{\epsilon}_t, \tag{3}$$

where \mathbf{x}_{t-1} is the previous state, \mathbf{u}_t is the most recent control input, and $\boldsymbol{\epsilon}_t$ is an independent *process noise* that is normally distributed with $\boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_t)$. Because of the properties of the Gaussian distribution presented earlier, this affine model preserves the Gaussian structure. In particular, the state transition model can be explicitly written as:

$$p(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{u}_t) = \frac{1}{\sqrt{\det(2\pi\mathbf{R}_t)}} \exp\left(-\frac{1}{2}(\mathbf{x}_t - A_t \mathbf{x}_{t-1} - B_t \mathbf{u}_t)^T \mathbf{R}_t^{-1}(\mathbf{x}_t - A_t \mathbf{x}_{t-1} - B_t \mathbf{u}_t)\right).$$

In other words, this can be expressed in shorthand as $\mathbf{x}_t \sim \mathcal{N}(A_t \mathbf{x}_{t-1} + B_t \mathbf{u}_t, \mathbf{R}_t)$.

Second, the measurement model is also assumed to be linear, which again preserves the Gaussian structure:

$$\mathbf{z}_t = C_t \mathbf{x}_t + \boldsymbol{\delta}_t, \tag{4}$$

where δ_t is an independent measurement noise that is normally distributed with $\mathcal{N}(\mathbf{0}, \mathbf{Q}_t)$. Again, this implies the measurement model can be expressed probabilistically as:

$$p(\mathbf{z}_t | \mathbf{x}_t) = \frac{1}{\sqrt{\det(2\pi\mathbf{Q}_t)}} \exp\left(-\frac{1}{2}(\mathbf{z}_t - C_t\mathbf{x}_t)^T \mathbf{Q}_t^{-1}(\mathbf{z}_t - C_t\mathbf{x}_t)\right),$$

and in shorthand as $\mathbf{z}_t \sim \mathcal{N}(C_t\mathbf{x}_t, \mathbf{Q}_t)$.

To summarize, if the belief is modeled as a Gaussian distribution and the state transition and measurement models are both linear with Gaussian noise, then the Bayes filter updates can be applied and the belief will always remain Gaussian (i.e. the prediction and measurement correction steps will not warp or alter the *structure* of the belief distribution)! This results in a very practically efficient algorithm since now only the parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ need to be updated by the algorithm.

15.2.2 Algorithm

The Kalman filter algorithm is a recursive Bayes filter whose prediction and measurement correction steps take on a special form due to the structure of the Gaussian belief distributions and the assumptions listed above. In particular, the Kalman filter algorithm is given in Algorithm 1.

Data: $\boldsymbol{\mu}_{t-1}, \boldsymbol{\Sigma}_{t-1}, \mathbf{u}_t, \mathbf{z}_t$
Result: $\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t$
 $\bar{\boldsymbol{\mu}}_t = A_t\boldsymbol{\mu}_{t-1} + B_t\mathbf{u}_t$
 $\bar{\boldsymbol{\Sigma}}_t = A_t\boldsymbol{\Sigma}_{t-1}A_t^T + \mathbf{R}_t$
 $\mathbf{K}_t = \bar{\boldsymbol{\Sigma}}_t C_t^T (C_t \bar{\boldsymbol{\Sigma}}_t C_t^T + \mathbf{Q}_t)^{-1}$
 $\boldsymbol{\mu}_t = \bar{\boldsymbol{\mu}}_t + \mathbf{K}_t(\mathbf{z}_t - C_t\bar{\boldsymbol{\mu}}_t)$
 $\boldsymbol{\Sigma}_t = (I - \mathbf{K}_t C_t)\bar{\boldsymbol{\Sigma}}_t$
return $\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t$

Algorithm 1: Kalman Filter Algorithm

In this algorithm, the first two steps define the predicted mean $\bar{\boldsymbol{\mu}}_t$ and covariance $\bar{\boldsymbol{\Sigma}}_t$, and the remaining steps perform the measurement correction. The matrix \mathbf{K}_t that is computed for the measurement correction is typically referred to as the *Kalman gain*. This prediction and correction process is also shown graphically in Figure 2.

15.2.3 Practical Considerations

Due to the exploitation of the Gaussian distribution, the Kalman filter is a computationally efficient algorithm that can handle continuous state values. However, the consideration of Gaussian beliefs also restricts the flexibility of the probabilistic model. In particular, the belief is forced to be unimodal which may limit performance. Additionally, the assumption about the linearity of the state transition and measurement models may not be very accurate for some robots and certain sensors, which can make the Kalman filter not perform well for some robotics applications.

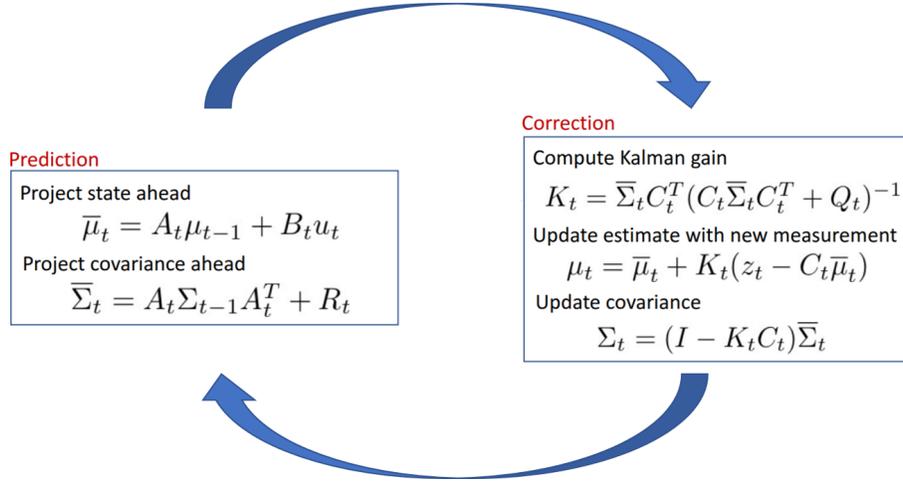


Figure 2: Kalman filter prediction and measurement correction procedure.

15.2.4 Derivation

While it is possible to derive the Kalman filter algorithm by evaluating the Bayes filter updates from the previous chapter (i.e. computing the integral of $p(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{u}_t)p(\mathbf{x}_{t-1}, \text{etc.})$), it is more intuitive to directly leverage the properties of Gaussians presented in the preceding section. First, from the prior belief distribution $bel(\mathbf{x}_{t-1}) \sim \mathcal{N}(\boldsymbol{\mu}_{t-1}, \boldsymbol{\Sigma}_{t-1})$ the predicted belief $\overline{bel}(\mathbf{x}_{t-1})$ can be computed by using the affine transformation property of Gaussian random variables and the sum of two independent Gaussian random variables property. Specifically, these properties can be applied to the assumed linear state transition model (3) to give:

$$\bar{\boldsymbol{\mu}}_t = A_t \boldsymbol{\mu}_{t-1} + B_t \bar{\mathbf{u}}_t + \mathbf{0},$$

where the $\mathbf{0}$ is the mean of the independent noise $\boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_t)$. The covariance properties can similarly be used to give:

$$\bar{\boldsymbol{\Sigma}}_t = A_t \boldsymbol{\Sigma}_{t-1} A_t^T + \mathbf{R}_t.$$

For the measurement update it is possible to again use the properties of Gaussians to simplify the derivation of the Kalman filter measurement correction step. In particular, that the product of two Gaussians is also Gaussian. In fact, the product of the two Gaussians:

$$bel(x_t) = p(\mathbf{z}_t | \mathbf{x}_t) \overline{bel}(\mathbf{x}_t) = \mathcal{N}(C_t \mathbf{x}_t, \mathbf{Q}_t) \mathcal{N}(\bar{\boldsymbol{\mu}}_t, \bar{\boldsymbol{\Sigma}}_t),$$

can be expressed as:

$$bel(x_t) = \eta \exp\left(-\frac{1}{2} J_t\right),$$

where η is a normalization constant and:

$$J_t = (\mathbf{z}_t - C_t \mathbf{x}_t)^T \mathbf{Q}_t^{-1} (\mathbf{z}_t - C_t \mathbf{x}_t) + (\mathbf{x}_t - \bar{\boldsymbol{\mu}}_t)^T \bar{\boldsymbol{\Sigma}}_t^{-1} (\mathbf{x}_t - \bar{\boldsymbol{\mu}}_t).$$

To determine the mean $\boldsymbol{\mu}_t$ and covariance $\boldsymbol{\Sigma}_t$ of this new Gaussian, one simple approach is just take the first and second derivative of J_t with respect to \mathbf{x}_t . The mean is found where the first derivative is zero, and the covariance is the (inverse) of the constant second derivative:

$$\begin{aligned} 0 &= -C_t^T \mathbf{Q}_t^{-1} (\mathbf{z}_t - C_t \boldsymbol{\mu}_t) + \bar{\boldsymbol{\Sigma}}_t^{-1} (\boldsymbol{\mu}_t - \bar{\boldsymbol{\mu}}_t), \\ \boldsymbol{\Sigma}_t^{-1} &= C_t^T \mathbf{Q}_t^{-1} C_t + \bar{\boldsymbol{\Sigma}}_t^{-1}. \end{aligned}$$

Through some algebraic manipulation the mean can be written in terms of the covariance $\boldsymbol{\Sigma}_t$:

$$\boldsymbol{\mu}_t = \bar{\boldsymbol{\mu}}_t + \boldsymbol{\Sigma}_t C_t^T \mathbf{Q}_t^{-1} (\mathbf{z}_t - C_t \bar{\boldsymbol{\mu}}_t),$$

and of course the covariance can be written as:

$$\boldsymbol{\Sigma}_t = (C_t^T \mathbf{Q}_t^{-1} C_t + \bar{\boldsymbol{\Sigma}}_t^{-1})^{-1}.$$

While it is technically possible to stop here, this is not quite the form of the Kalman filter equations. In particular a few more algebraic steps are needed, based on the matrix inversion lemma result:

$$(C_t^T \mathbf{Q}_t^{-1} C_t + \bar{\boldsymbol{\Sigma}}_t^{-1})^{-1} = \bar{\boldsymbol{\Sigma}}_t - \bar{\boldsymbol{\Sigma}}_t C_t^T (C_t \bar{\boldsymbol{\Sigma}}_t C_t^T + \mathbf{Q}_t)^{-1} C_t \bar{\boldsymbol{\Sigma}}_t.$$

By choosing to define the Kalman gain as $\mathbf{K}_t = \bar{\boldsymbol{\Sigma}}_t C_t^T (C_t \bar{\boldsymbol{\Sigma}}_t C_t^T + \mathbf{Q}_t)^{-1}$ it can be seen that the covariance can be expressed as:

$$\boldsymbol{\Sigma}_t = \bar{\boldsymbol{\Sigma}}_t - \mathbf{K}_t C_t \bar{\boldsymbol{\Sigma}}_t,$$

Through some additional algebra, the expression for the mean can also be expressed in terms of the Kalman gain and simplified to its final form:

$$\boldsymbol{\mu}_t = \bar{\boldsymbol{\mu}}_t + \mathbf{K}_t (\mathbf{z}_t - C_t \bar{\boldsymbol{\mu}}_t).$$

For more details on the algebraic steps see [TBF05].

15.3 Extended Kalman Filter (EKF)

The extended Kalman filter (EKF) is a modified version of the Kalman filter that revisits the assumption of linearity for the state transition and measurement models. This extension still exploits the Gaussian distribution to represent the belief in a computationally efficient parametric way, but by generalizing to more complex models the EKF can be applied to a

wider variety of robotics state estimation and localization problems. In particular, the EKF considers general nonlinear state transition and measurement models defined as:

$$\begin{aligned}\mathbf{x}_t &= g(\mathbf{u}_t, \mathbf{x}_{t-1}) + \boldsymbol{\epsilon}_t, \\ \mathbf{z}_t &= h(\mathbf{x}_t) + \boldsymbol{\delta}_t,\end{aligned}\tag{5}$$

where again $\boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_t)$ and $\boldsymbol{\delta}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_t)$ are normally distributed noise terms.

The EKF handles these nonlinear functions by taking advantage of their first order Taylor series expansions (which are linear functions, similar to those used in the Kalman filter). In particular the Taylor series expansion of the state transition model g is performed about the *most likely state* from the current belief distribution, which is the expected value $\boldsymbol{\mu}_{t-1}$:

$$g(\mathbf{u}_t, \mathbf{x}_{t-1}) \approx g(\mathbf{u}_t, \boldsymbol{\mu}_{t-1}) + G_t(\mathbf{x}_{t-1} - \boldsymbol{\mu}_{t-1}),$$

where $G_t = \nabla_{\mathbf{x}}g(\mathbf{u}_t, \boldsymbol{\mu}_{t-1})$ is the Jacobian of g evaluated at $\boldsymbol{\mu}_{t-1}$. From this linear approximation the state transition model can be expressed as:

$$p(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{u}_t) = \frac{1}{\sqrt{\det(2\pi\mathbf{R}_t)}} \exp\left(-\frac{1}{2}\Delta\mathbf{x}_t^T \mathbf{R}_t^{-1} \Delta\mathbf{x}_t\right),$$

where

$$\Delta\mathbf{x}_t = \mathbf{x}_t - g(\mathbf{u}_t, \boldsymbol{\mu}_{t-1}) - G_t(\mathbf{x}_{t-1} - \boldsymbol{\mu}_{t-1}).$$

From this result, the linear predictions that are used in the Kalman filter algorithm can be replaced by the nonlinear generalizations:

$$\begin{aligned}\bar{\boldsymbol{\mu}}_t &= g(\mathbf{u}_t, \boldsymbol{\mu}_{t-1}), \\ \bar{\boldsymbol{\Sigma}}_t &= G_t \boldsymbol{\Sigma}_{t-1} G_t^T + \mathbf{R}_t.\end{aligned}$$

As can be seen the prediction of the new mean is simply an evaluation of the nonlinear function g , and the updated covariance is very similar to the Kalman filter but leverages the Jacobian G_t .

A very similar procedure is used for the measurement corrections. The measurement model is also Taylor series expanded to yield (this time about the *predicted point* $\bar{\boldsymbol{\mu}}_t$):

$$h(\mathbf{x}_t) \approx h(\bar{\boldsymbol{\mu}}_t) + H_t(\mathbf{x}_t - \bar{\boldsymbol{\mu}}_t),$$

where $H_t = \nabla_{\mathbf{x}}h(\bar{\boldsymbol{\mu}}_t)$ is the Jacobian of h evaluated at $\bar{\boldsymbol{\mu}}_t$. The measurement model can then be expressed using this approximation as:

$$p(\mathbf{z}_t | \mathbf{x}_t) = \frac{1}{\sqrt{\det(2\pi\mathbf{Q}_t)}} \exp\left(-\frac{1}{2}\Delta\mathbf{z}_t^T \mathbf{Q}_t^{-1} \Delta\mathbf{z}_t\right),$$

where $\Delta\mathbf{z}_t = \mathbf{z}_t - h(\bar{\boldsymbol{\mu}}_t) - H_t(\mathbf{x}_t - \bar{\boldsymbol{\mu}}_t)$. From this result the measurement correction in the EKF can be shown to be similar to the Kalman filter, but where the Jacobians H_t are used:

$$\begin{aligned}\boldsymbol{\mu}_t &= \bar{\boldsymbol{\mu}}_t + \mathbf{K}_t(\mathbf{z}_t - h(\bar{\boldsymbol{\mu}}_t)), \\ \boldsymbol{\Sigma}_t &= (I - \mathbf{K}_t H_t) \bar{\boldsymbol{\Sigma}}_t,\end{aligned}$$

where the Kalman gain is $\mathbf{K}_t = \bar{\boldsymbol{\Sigma}}_t H_t^T (H_t \bar{\boldsymbol{\Sigma}}_t H_t^T + \mathbf{Q}_t)^{-1}$.

15.3.1 Algorithm

The extended Kalman filter algorithm is quite similar to the Kalman filter algorithm outlined in Algorithm 1. In particular the main differences are that the updates use the nonlinear functions and their Jacobians rather than assuming strictly linear models. The EKF algorithm is outlined in Algorithm 2.

```
Data:  $\boldsymbol{\mu}_{t-1}, \boldsymbol{\Sigma}_{t-1}, \mathbf{u}_t, \mathbf{z}_t$   
Result:  $\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t$   
 $\bar{\boldsymbol{\mu}}_t = g(\mathbf{u}_t, \boldsymbol{\mu}_{t-1})$   
 $\bar{\boldsymbol{\Sigma}}_t = G_t \boldsymbol{\Sigma}_{t-1} G_t^T + \mathbf{R}_t$   
 $\mathbf{K}_t = \bar{\boldsymbol{\Sigma}}_t H_t^T (H_t \bar{\boldsymbol{\Sigma}}_t H_t^T + \mathbf{Q}_t)^{-1}$   
 $\boldsymbol{\mu}_t = \bar{\boldsymbol{\mu}}_t + \mathbf{K}_t (\mathbf{z}_t - h(\bar{\boldsymbol{\mu}}_t))$   
 $\boldsymbol{\Sigma}_t = (I - \mathbf{K}_t H_t) \bar{\boldsymbol{\Sigma}}_t$   
return  $\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t$ 
```

Algorithm 2: Extended Kalman Filter Algorithm

15.3.2 Practical Considerations

The extended Kalman filter can provide more accurate results than the Kalman filter in many applications due to its ability to consider more general nonlinear models. However, the approximation of the nonlinear models by a Taylor series expansion can lead to the filter to diverge if the approximation is not accurate enough. Additionally, the EKF still suffers from the same unimodal modeling assumption as the Kalman filter since the beliefs are still represented by a single Gaussian distribution.

15.4 Unscented Kalman Filter

The unscented Kalman filter (UKF) is another variation of the Kalman filter (still uses Gaussian distribution to parameterize the belief). This filter is also similar to the EKF in that it can handle nonlinear state transition and measurement models. However, this filter improves upon the EKF by not relying on Taylor series expansions, which can cause filter divergence due to approximation errors. This is accomplished by representing the Gaussian through a set of *sigma-points* that are transformed through the nonlinear functions. Once each sigma-point has been updated, a new Gaussian distribution is computed to represent the updated belief.

References

- [TBF05] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic robotics*. MIT press, 2005.