

AA274A: Principles of Robot Autonomy I

Course Notes

The last few chapters introduced some of the most widely used algorithms based on Bayes' filter for probabilistic robot localization and state estimation. However these fundamental algorithms still need further enhancements before application to many robot localization tasks, since in their standard form they don't incorporate a notion of a local *map*. For example, a particle filter could be applied in its original form to a problem of global localization based on GNSS measurements, but localizing based on range measurements requires knowledge about *what* object is being ranged, and *where* that object is with respect to the local environment (i.e. the map). In this chapter a more specific definition of mobile robot localization is considered, namely the problem of determining the pose of a robot relative to a *given map* of the environment.

17 Robot Localization

Localization with respect to a map can be interpreted as a problem of coordinate transformation. Maps are described in a global coordinate system, which is independent of a robot's pose. Localization can then be viewed as the process of establishing a correspondence between the map coordinate system and the robot's local coordinate system. Knowing this coordinate transformation then enables the robot to express the location of objects of interest within its own coordinate frame (a necessary prerequisite for robot autonomy).

In 2D problems, knowing the pose $\mathbf{x}_t = (x, y, \theta)^T$ of a robot is sufficient to establish this correspondence, and an ideal sensor would directly be able to measure this pose. However in practice no such sensor exists, and therefore *indirect* (often noisy) measurements \mathbf{z}_t of the pose are used. Since it is almost impossible to be able to reliably estimate \mathbf{x}_t from a single measurement \mathbf{z}_t , localization algorithm typically integrate additional data *over time* to build reliable localization estimates. For example, consider a robot located inside a building where many corridors look alike. In this case a single sensor measurement (e.g. a range scan) is usually insufficient to disambiguate the identity of the corridor from the others.

In this chapter it will be seen how this map-based localization problem can be cast in the Bayesian filtering framework, such that the algorithms from previous chapters can be leveraged.

17.1 A Taxonomy of Localization Problems

To understand the broad scope of challenges related to robot localization, it is useful to develop a brief taxonomy of localization problems. This categorization will divide localization problems along a number of important dimensions pertaining to the nature of the environment (e.g. static versus dynamic), the initial knowledge that a robot may possess, and how information about the environment is gathered (e.g. passive or active, with one robot or collaboratively with several robots).

17.1.1 Local vs. Global

Localization problems can be characterized by the type of knowledge that is available initially, which has a significant impact on what type of localization algorithm is most appropriate for the problem.

- *Position tracking* problems assume that the initial pose of the robot is known. In these types of problems only incremental updates are required (i.e. the localization error is generally always small), and therefore unimodal Gaussian filters (e.g. Kalman filters) can be efficiently applied.
- *Global localization* problems assume that the initial pose of the robot is unknown. In these scenarios the use of a unimodal parametric belief distribution cannot adequately capture the global uncertainty. Therefore it is more appropriate to use non-parametric, multi-hypothesis filters, such as the particle filter.
- The *kidnapped robot problem* is a variant of the global localization problem (i.e. unknown initial pose) where the robot can get “kidnapped” and “teleported” to some other location. This problem is more difficult than the global localization problem since the localization algorithm needs to have an awareness that sudden drastic to the robot’s pose are possible. While robots are typically not “kidnapped” in practice, the consideration of this type of problem is useful for ensuring the localization algorithm is *robust*, since the ability to recover from failure is essential for truly autonomous robots. Similar to the global localization problem, these problems are often best addressed using non-parametric, multi-hypothesis filters.

17.1.2 Static vs. Dynamic

Environmental changes are another important consideration in mobile robot localization, specifically whether they are static or dynamic.

- In *static* environments the robot is the only object that moves. Static environments are generally much easier to perform localization in.
- *Dynamic* environments possess objects other than the robot whose locations or configurations change over time. This problem is usually addressed by augmenting the state

vector to include the movement of dynamic entities, or by filtering the sensor data to remove the effects of environment dynamics.

17.1.3 Passive vs. Active

Information collected via measurements is crucial for robot localization. Therefore it is reasonable to consider localization problems where the robot can *explicitly* choose its actions to gather more (or more specific) information from the environment.

- *Passive localization* problems assume that the robot’s motion is unrelated to its localization process.
- *Active localization* problems consider the ability of the robot to choose its actions (at least partially) to improve its understanding of the environment. For example, a robot in the corner of a room might choose to reorient itself to face the rest of the room, so it can collect environmental information as it moves along the wall. Hybrid approaches are also possible, since it may be inefficient to use active localization all of the time.

17.1.4 Single Robot vs. Multi-Robot

It is of course also possible to consider problems where several robots all gather independent information and then share that information with each other.

- *Single-robot localization* problems are the most commonly studied and utilized approach, and are often simpler because all data is collected on a single platform.
- *Multi-robot localization* problems consider teams of robots that share information in such a way that one robot’s belief can be used to influence another robot’s belief if the relative location between robots is known.

17.2 Robot Localization via Bayesian Filtering

The parametric (e.g. EKF) and non-parametric (e.g. particle) filters from the previous chapters are all variations of the Bayes filter. In particular they rely on a Markov process assumption and the identification of probabilistic measurement models. In this section it is shown how map-based robot localization can be cast into this framework, such that the previously discussed algorithms can be applied.

Similar to the general filtering context from the previous chapters, at time t the state is denoted by \mathbf{x}_t , the control input is denoted by \mathbf{u}_t , and the measurements are denoted by \mathbf{z}_t . For example, a differential drive robot equipped with a laser range-finder (returning a set of range measurements r_i and bearings ϕ_i), the state, control, and measurements would be:

$$\mathbf{x}_t = \begin{bmatrix} x \\ y \\ \theta \end{bmatrix}, \quad \mathbf{u}_t = \begin{bmatrix} v \\ \omega \end{bmatrix}, \quad \mathbf{z}_t = \begin{bmatrix} r_1 \\ \phi_1 \\ \vdots \end{bmatrix}. \quad (1)$$

However, the critical new component is the concept of a *map* (denoted as \mathbf{m}), which is a list of objects in the environment along with their properties:

$$\mathbf{m} = \{m_1, m_2, \dots, m_N\}, \quad (2)$$

where m_i represents the properties of a specific object. Generally there are two types of maps that will be considered, location-based maps and feature-based maps, which typically have differences in both computational efficiency and expressiveness.

For location-based maps, the index i associated with object m_i corresponds to a specific *location* (i.e. m_i are volumetric objects). For example, objects m_i in a location-based map might represent cells in a cell decomposition or grid representation of a map (see Figure 1). One potential disadvantage of the cell-based maps is that their resolution is dependent on

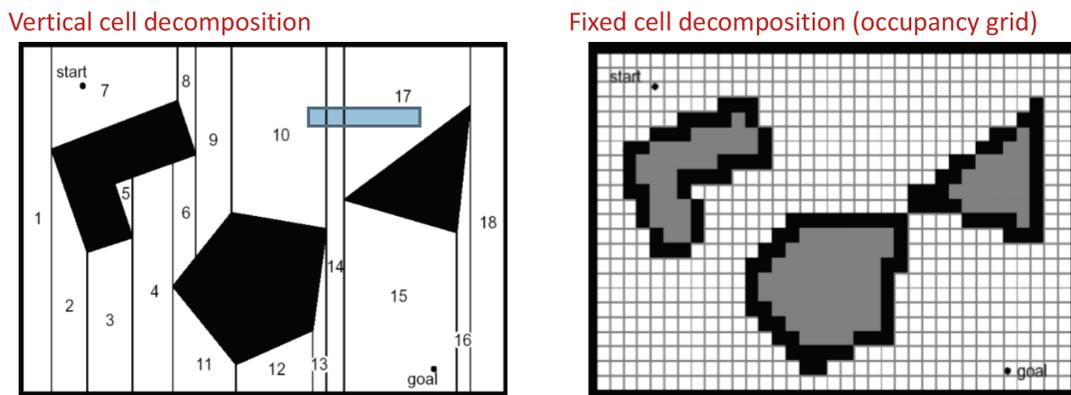


Figure 1: Two examples of location-based maps, both represent the map as a set of volumetric objects (i.e. cells in these cases).

the size of the cells, but their advantage is that they can explicitly encode information about presence (or absence) of objects in specific locations.

For feature-based maps, an index i is a feature index, and m_i contains information about the properties of that feature, including its Cartesian location. These types of maps can typically be thought of as a collection of landmarks. Figure 2 gives two examples of feature-based maps, one which is represented by a set of lines, and another which is represented by nodes and edges like a graph (i.e. a topological map). Feature-based maps can be more finely tuned to specific environments, for example the line-based map might make sense to use in highly structured environments such as buildings. While feature-based maps can be computationally efficient, their main disadvantage is that they typically do not capture spatial information about all potential obstacles.

17.2.1 State Transition Model

In the previous chapters on Bayesian filtering the probabilistic state transition model $p(\mathbf{x}_t | \mathbf{u}_t, \mathbf{x}_{t-1})$ describes the posterior distribution over the states that the robot could transition

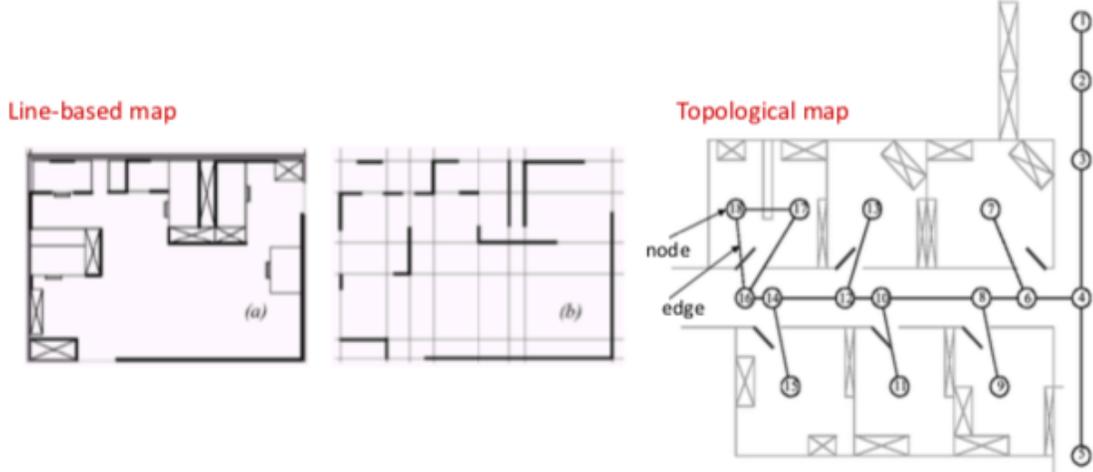


Figure 2: Two examples of feature-based maps.

to when executing control \mathbf{u}_t from \mathbf{x}_{t-1} . However in robot localization problems it might be important to take into account how the map \mathbf{m} could affect the state transition since in general:

$$p(\mathbf{x}_t | \mathbf{u}_t, \mathbf{x}_{t-1}) \neq p(\mathbf{x}_t | \mathbf{u}_t, \mathbf{x}_{t-1}, \mathbf{m}).$$

For example, $p(\mathbf{x}_t | \mathbf{u}_t, \mathbf{x}_{t-1})$ cannot account for the fact that a robot cannot move through walls since it doesn't know that walls exist!

However, a common approximation is to make the assumption that:

$$p(\mathbf{x}_t | \mathbf{u}_t, \mathbf{x}_{t-1}, \mathbf{m}) \approx \eta \frac{p(\mathbf{x}_t | \mathbf{u}_t, \mathbf{x}_{t-1})p(\mathbf{x}_t | \mathbf{m})}{p(\mathbf{x}_t)}, \quad (3)$$

where η is a normalization constant. This approximation can be derived from Bayes' rule by assuming that $p(\mathbf{m} | \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{u}_t) \approx p(\mathbf{m} | \mathbf{x}_t)$ (which is a tight approximation under high update rates). More specifically:

$$\begin{aligned} p(\mathbf{x}_t | \mathbf{u}_t, \mathbf{x}_{t-1}, \mathbf{m}) &= \frac{p(\mathbf{m} | \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{u}_t)p(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{u}_t)}{p(\mathbf{m} | \mathbf{x}_{t-1}, \mathbf{u}_t)}, \\ &= \eta' p(\mathbf{m} | \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{u}_t)p(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{u}_t), \\ &\approx \eta' p(\mathbf{m} | \mathbf{x}_t)p(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{u}_t), \\ &= \eta \frac{p(\mathbf{x}_t | \mathbf{u}_t, \mathbf{x}_{t-1})p(\mathbf{x}_t | \mathbf{m})}{p(\mathbf{x}_t)}, \end{aligned}$$

where η' and η are normalization constants (such that the total probability density integrates to one).

In this approximation the term $p(\mathbf{x}_t | \mathbf{m})$ is the state probability conditioned on the map which can be thought of as describing the “consistency” of state with respect to the map. The approximation (3) can therefore be viewed as making a probabilistic guess using the original state transition model (without map knowledge), and then using the consistency term $p(\mathbf{x}_t | \mathbf{m})$ to check the plausibility of the new state \mathbf{x}_t given the map.

17.2.2 Measurement Model

The probabilistic measurement model $p(\mathbf{z}_t | \mathbf{x}_t)$ from previous chapters also needs to be modified to take map information into account. This new measurement model can simply be expressed as $p(\mathbf{z}_t | \mathbf{x}_t, \mathbf{m})$ (i.e. measurement is also conditioned on the map). This is obviously important because the local measurements can have significant influence from the environment. For example a range measurement is dependent on what object is currently in the line of sight.

Additionally, since the suite of sensors on a robot may generate more than one measurement when queried, it is also common to make another measurement model assumption for simplicity. Suppose K measurements are taken at time t , such that:

$$\mathbf{z}_t = \begin{bmatrix} \mathbf{z}_t^1 \\ \vdots \\ \mathbf{z}_t^K \end{bmatrix}.$$

Then it can often be assumed that each of the K measurements are conditionally independent from each other (i.e. when conditioned on \mathbf{x}_t and \mathbf{m} the probability of measuring \mathbf{z}_t^k is independent from the other measurements). With this assumption the probabilistic measurement model can be expressed as:

$$p(\mathbf{z}_t | \mathbf{x}_t, \mathbf{m}) = \prod_{k=1}^K p(\mathbf{z}_t^k | \mathbf{x}_t, \mathbf{m}). \quad (4)$$

17.3 Markov Localization

With the probabilistic state transition and measurement models that include the map, the Bayes' filter can be directly modified as shown in Algorithm 1. As can be seen, this algorithm

```

Data:  $bel(\mathbf{x}_{t-1}), \mathbf{u}_t, \mathbf{z}_t, \mathbf{m}$ 
Result:  $bel(\mathbf{x}_t)$ 
foreach  $\mathbf{x}_t$  do
     $\bar{bel}(\mathbf{x}_t) = \int p(\mathbf{x}_t | \mathbf{u}_t, \mathbf{x}_{t-1}, \mathbf{m}) bel(\mathbf{x}_{t-1}) d\mathbf{x}_{t-1}$ 
     $bel(\mathbf{x}_t) = \eta p(\mathbf{z}_t | \mathbf{x}_t, \mathbf{m}) \bar{bel}(\mathbf{x}_t)$ 
end
return  $bel(\mathbf{x}_t)$ 

```

Algorithm 1: Markov Localization Algorithm

is conceptually identical to the Bayes' filter except for the inclusion of the model \mathbf{m} . This algorithm is referred to as the *Markov localization* algorithm, and the localization problem it is trying to solve is generally referred to as simply *Markov localization*¹.

¹Recall the use of the Markov property assumption in the derivation of the Bayes' filter.

The Markov localization algorithm can be used to address global localization, position tracking, and kidnapped robot problems, but generally some implementation details might be different. The choice for the initial (prior) belief distribution $bel(\mathbf{x}_0)$ is one such parameter that may be different depending on the type of localization problem.

Specifically, since the initial belief encodes any prior knowledge about the robot pose, the best choice of distribution depends on what (if any) knowledge is available. For example, in the position tracking problem it is assumed that an initial pose of the robot is known. Therefore choosing a (unimodal) Gaussian distribution $bel(\mathbf{x}_0) \sim \mathcal{N}(\bar{\mathbf{x}}_0, \Sigma_0)$ with a small covariance might be a good choice. Alternatively, for a global localization problem the initial pose is not known. In this case an appropriate choice for the initial belief would be a uniform distribution $bel(\mathbf{x}_0) = 1/|X|$ over all possible states \mathbf{x} .

Similarly to the original Bayes' filter from previous chapters, the Markov localization algorithm 1 is generally not possible to implement in a computationally tractable way. However, practical implementations can still be developed by again leveraging some sort of structure to the belief distribution $bel(\mathbf{x}_t)$ (e.g. through Gaussian or particle representations). Two commonly used implementations based on specific structured beliefs will now be discussed: extended Kalman filter localization and Monte Carlo localization.

17.4 Extended Kalman Filter (EKF) Localization

The extended Kalman filter (EKF) localization algorithm is essentially equivalent to the EKF algorithm presented in previous chapters, except that it also takes the map \mathbf{m} into account. In particular, it still makes a Gaussian belief assumption, $bel(\mathbf{x}_t) \sim \mathcal{N}(\boldsymbol{\mu}_t, \Sigma_t)$, to add structure to the filtering problem. As a brief review, the assumed state transition model is given by:

$$\mathbf{x}_t = g(\mathbf{u}_t, \mathbf{x}_{t-1}) + \boldsymbol{\epsilon}_t,$$

where $\boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_t)$ is Gaussian zero-mean noise. The Jacobian G_t is again defined by $G_t = \nabla_{\mathbf{x}} g(\mathbf{u}_t, \boldsymbol{\mu}_{t-1})$, where $\boldsymbol{\mu}_{t-1}$ is the expected value of the previous belief distribution $bel(\mathbf{x}_{t-1})$.

The main difference in EKF localization is the assumption that a feature-based map is available, consisting of point landmarks given by:

$$\mathbf{m} = \{m_1, m_2, \dots, m_N\}, \quad m_j = (m_{j,x}, m_{j,y}),$$

where N is the total number of landmarks, and each landmark m_j encapsulates the location $(m_{j,x}, m_{j,y})$ of the landmark in the global coordinate frame. Measurements \mathbf{z}_t associated with these point landmarks at a time t are denoted by:

$$\mathbf{z}_t = \{\mathbf{z}_t^1, \mathbf{z}_t^2, \dots\},$$

where \mathbf{z}_t^i is associated with a particular landmark and is assumed to be generated by the measurement model:

$$\mathbf{z}_t^i = h(\mathbf{x}_t, j, \mathbf{m}) + \boldsymbol{\delta}_t,$$

where $\delta_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_t)$ is Gaussian zero-mean noise and j is the index of the map feature $m_j \in \mathbf{m}$ that measurement i is associated with.

One fundamental problem that now needs to be addressed is the *data association problem*, which arises due to uncertainty in which measurements are associated with which landmark. To begin addressing this problem, the correspondences are modeled through a variable $c_t^i \in \{1, \dots, N+1\}$, which take on the values $c_t^i = j$ if measurement i corresponds to landmark j , and $c_t^i = N+1$ if measurement i has no corresponding landmark. Then, given a correspondence c_t^i of measurement i (associated with a specific landmark), the Jacobian H_t^i used in the EKF measurement correction step can be determined. Specifically, for the i -th measurement the Jacobian of the new measurement model can be computed by $H_t^{c_t^i} = \nabla_{\mathbf{x}} h(\bar{\boldsymbol{\mu}}_t, c_t^i, \mathbf{m})$, where $\bar{\boldsymbol{\mu}}_t$ is the predicted mean (that results from the EKF prediction step).

17.4.1 EKF Localization with Known Correspondences

In practice the correspondences between measurements \mathbf{z}_t^i and landmarks m_j are generally unknown. However, it is useful from a pedagogical standpoint to first consider the case where these correspondences $\mathbf{c}_t = [c_t^1, \dots]^T$ are assumed to be *known*.

In the EKF localization algorithm given in Algorithm 2, the main difference from the original EKF filter algorithm is that multiple measurements are processed at the same time. Crucially, this is accomplished in a computationally efficient way by exploiting the conditional independence assumption (4) for the measurements. In fact, by exploiting this assumption and some special properties of Gaussians, the multi-measurement update can be implemented by just looping over each measurement individually and applying the standard EKF correction.

```

Data:  $\boldsymbol{\mu}_{t-1}, \boldsymbol{\Sigma}_{t-1}, \mathbf{u}_t, \mathbf{z}_t, \mathbf{c}_t, \mathbf{m}$ 
Result:  $\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t$ 
 $\bar{\boldsymbol{\mu}}_t = g(\mathbf{u}_t, \boldsymbol{\mu}_{t-1})$ 
 $\bar{\boldsymbol{\Sigma}}_t = G_t \boldsymbol{\Sigma}_{t-1} G_t^T + \mathbf{R}_t$ 
foreach  $\mathbf{z}_t^i$  do
     $j = c_t^i$ 
     $S_t^i = H_t^j \bar{\boldsymbol{\Sigma}}_t [H_t^j]^T + \mathbf{Q}_t$ 
     $K_t^i = \bar{\boldsymbol{\Sigma}}_t [H_t^j]^T [S_t^i]^{-1}$ 
     $\bar{\boldsymbol{\mu}}_t = \bar{\boldsymbol{\mu}}_t + K_t^i (\mathbf{z}_t^i - h(\bar{\boldsymbol{\mu}}_t, j, \mathbf{m}))$ 
     $\bar{\boldsymbol{\Sigma}}_t = (I - K_t^i H_t^j) \bar{\boldsymbol{\Sigma}}_t$ 
end
 $\boldsymbol{\mu}_t = \bar{\boldsymbol{\mu}}_t$ 
 $\boldsymbol{\Sigma}_t = \bar{\boldsymbol{\Sigma}}_t$ 
return  $\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t$ 

```

Algorithm 2: Extended Kalman Filter Localization Algorithm

17.4.2 EKF Localization with Unknown Correspondences

For EKF localization with *unknown* correspondences, the correspondence variables must also be estimated! The simplest way to determine the correspondences online is to use maximum likelihood estimation, in which the most likely value of the correspondences \mathbf{c}_t is determined by maximizing the data likelihood:

$$\hat{\mathbf{c}}_t = \arg \max_{\mathbf{c}_t} p(\mathbf{z}_t \mid \mathbf{c}_{1:t}, \mathbf{m}, \mathbf{z}_{1:t-1}, \mathbf{u}_{1:t})$$

In other words, the set of correspondence variables is chosen to maximize the probability of getting the current measurement given the history of correspondence variables, the map, the history of measurements, and the history of controls. By marginalizing over the current pose \mathbf{x}_t this distribution can be written as:

$$\begin{aligned} p(\mathbf{z}_t \mid \mathbf{c}_{1:t}, \mathbf{m}, \mathbf{z}_{1:t-1}, \mathbf{u}_{1:t}) &= \int p(\mathbf{z}_t \mid \mathbf{c}_{1:t}, \mathbf{x}_t, \mathbf{m}, \mathbf{z}_{1:t-1}, \mathbf{u}_{1:t}) p(\mathbf{x}_t \mid \mathbf{c}_{1:t}, \mathbf{m}, \mathbf{z}_{1:t-1}, \mathbf{u}_{1:t}) d\mathbf{x}_t, \\ &= \int p(\mathbf{z}_t \mid \mathbf{c}_t, \mathbf{x}_t, \mathbf{m}) \overline{\text{bel}}(\mathbf{x}_t) d\mathbf{x}_t. \end{aligned}$$

Note that the term $p(\mathbf{z}_t \mid \mathbf{c}_{1:t}, \mathbf{x}_t, \mathbf{m})$ is essentially the assumed measurement model given *known* correspondences. Then, by again leveraging the conditional independence assumption for the measurements \mathbf{z}_t^i from (4), this can be written as:

$$p(\mathbf{z}_t \mid \mathbf{c}_{1:t}, \mathbf{m}, \mathbf{z}_{1:t-1}, \mathbf{u}_{1:t}) = \int \prod_i p(\mathbf{z}_t^i \mid c_t^i, \mathbf{x}_t, \mathbf{m}) \overline{\text{bel}}(\mathbf{x}_t) d\mathbf{x}_t.$$

Importantly, each decision variable c_t^i in the maximization of this quantity shows up in separate terms of the product! Therefore it is possible to maximize each parameter independently by solving the optimization problems:

$$\hat{c}_t^i = \arg \max_{c_t^i} \int p(\mathbf{z}_t^i \mid c_t^i, \mathbf{x}_t, \mathbf{m}) \overline{\text{bel}}(\mathbf{x}_t) d\mathbf{x}_t.$$

This problem can be solved quite efficiently since it is assumed that the measurement models and belief distributions are Gaussian². In particular, the probability distribution resulting from the integral is a Gaussian with mean and covariance:

$$\int p(\mathbf{z}_t^i \mid c_t^i, \mathbf{x}_t, \mathbf{m}) \overline{\text{bel}}(\mathbf{x}_t) d\mathbf{x}_t \sim \mathcal{N}(h(\bar{\boldsymbol{\mu}}_t, c_t^i, \mathbf{m}), H_t^{c_t^i} \bar{\boldsymbol{\Sigma}}_t [H_t^{c_t^i}]^T + \mathbf{Q}_t).$$

The maximum likelihood optimization problem can therefore be expressed as:

$$\hat{c}_t^i = \arg \max_{c_t^i} \mathcal{N}(\mathbf{z}_t^i \mid \hat{\mathbf{z}}_t^{c_t^i}, S_t^{c_t^i}),$$

²Similar to the previous chapters, in this case the product of terms inside the integral will be Gaussian since both terms are Gaussian.

where $\hat{\mathbf{z}}_t^j = h(\bar{\boldsymbol{\mu}}_t, j, \mathbf{m})$ and $S_t^j = H_t^j \bar{\boldsymbol{\Sigma}}_t [H_t^j]^T + \mathbf{Q}_t$. To solve this maximization problem, recall the definition of the Gaussian distribution:

$$\mathcal{N}(\mathbf{z}_t^i | \hat{\mathbf{z}}_t^j, S_t^j) = \eta \exp\left(-\frac{1}{2}(\mathbf{z}_t^i - \hat{\mathbf{z}}_t^j)^T [S_t^j]^{-1} (\mathbf{z}_t^i - \hat{\mathbf{z}}_t^j)\right),$$

where η is a normalization constant. Since the exponential function is monotonically increasing and since η is a positive constant, the maximum likelihood estimation problem can be equivalently expressed as:

$$\hat{c}_t^i = \arg \min_{c_t^i} d_t^{i,c_t^i}, \quad (5)$$

where

$$d_t^{ij} = (\mathbf{z}_t^i - \hat{\mathbf{z}}_t^j)^T [S_t^j]^{-1} (\mathbf{z}_t^i - \hat{\mathbf{z}}_t^j), \quad (6)$$

is referred to as the *Mahalanobis distance*.

The EKF localization algorithm with unknown correspondences is very similar to Algorithm 2, except with the addition of this maximum likelihood estimation step. This new algorithm is given in Algorithm 3.

```

Data:  $\boldsymbol{\mu}_{t-1}, \boldsymbol{\Sigma}_{t-1}, \mathbf{u}_t, \mathbf{z}_t, \mathbf{m}$ 
Result:  $\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t$ 
 $\bar{\boldsymbol{\mu}}_t = g(\mathbf{u}_t, \boldsymbol{\mu}_{t-1})$ 
 $\bar{\boldsymbol{\Sigma}}_t = G_t \boldsymbol{\Sigma}_{t-1} G_t^T + \mathbf{R}_t$ 
foreach  $\mathbf{z}_t^i$  do
    foreach landmark  $k$  in the map do
         $\hat{\mathbf{z}}_t^k = h(\bar{\boldsymbol{\mu}}_t, k, \mathbf{m})$ 
         $S_t^k = H_t^k \bar{\boldsymbol{\Sigma}}_t [H_t^k]^T + \mathbf{Q}_t$ 
    end
     $j = \arg \min_k (\mathbf{z}_t^i - \hat{\mathbf{z}}_t^k)^T [S_t^k]^{-1} (\mathbf{z}_t^i - \hat{\mathbf{z}}_t^k)$ 
     $K_t^i = \bar{\boldsymbol{\Sigma}}_t [H_t^j]^T [S_t^j]^{-1}$ 
     $\bar{\boldsymbol{\mu}}_t = \bar{\boldsymbol{\mu}}_t + K_t^i (\mathbf{z}_t^i - \hat{\mathbf{z}}_t^j)$ 
     $\bar{\boldsymbol{\Sigma}}_t = (I - K_t^i H_t^j) \bar{\boldsymbol{\Sigma}}_t$ 
end
 $\boldsymbol{\mu}_t = \bar{\boldsymbol{\mu}}_t$ 
 $\boldsymbol{\Sigma}_t = \bar{\boldsymbol{\Sigma}}_t$ 
return  $\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t$ 

```

Algorithm 3: EKF Localization Algorithm, Unknown Correspondences

One of the disadvantages of using the maximum likelihood estimation is that it can be brittle with respect to outliers and in cases where there are equally likely hypothesis for the correspondence. An alternative approach to estimating correspondences that is more robust to outliers is to use a *validation gate*. In this approach the Mahalanobis smallest distance d_t^{ij} must also pass a *thresholding* test:

$$(\mathbf{z}_t^i - \hat{\mathbf{z}}_t^j)^T [S_t^j]^{-1} (\mathbf{z}_t^i - \hat{\mathbf{z}}_t^j) \leq \gamma,$$

in order for a correspondence to be created.

Example 17.1 (Differential Drive Robot with Range and Bearing Measurements). Consider a differential drive robot with state $\mathbf{x} = [x, y, \theta]^T$, and suppose a sensor is available on the robot which measures the range r and bearing ϕ of landmarks $m_j \in \mathbf{m}$ relative to the robot’s local coordinate frame. Additionally, multiple measurements corresponding to different features can be collected at each time step:

$$\mathbf{z}_t = \{[r_t^1, \phi_t^1]^T, [r_t^2, \phi_t^2]^T, \dots\},$$

where each measurement \mathbf{z}_t^i contains the range r_t^i and bearing ϕ_t^i .

Assuming the correspondences are known, the measurement model for the range and bearing is:

$$h(\mathbf{x}_t, j, \mathbf{m}) = \begin{bmatrix} \sqrt{(m_{j,x} - x)^2 + (m_{j,y} - y)^2} \\ \text{atan2}(m_{j,y} - y, m_{j,x} - x) - \theta \end{bmatrix}. \quad (7)$$

The measurement Jacobian H_t^j corresponding to a measurement from landmark j is then given by:

$$H_t^j = \begin{bmatrix} -\frac{m_{j,x} - \bar{\mu}_{t,x}}{\sqrt{(m_{j,x} - \bar{\mu}_{t,x})^2 + (m_{j,y} - \bar{\mu}_{t,y})^2}} & -\frac{m_{j,y} - \bar{\mu}_{t,y}}{\sqrt{(m_{j,x} - \bar{\mu}_{t,x})^2 + (m_{j,y} - \bar{\mu}_{t,y})^2}} & 0 \\ \frac{m_{j,y} - \bar{\mu}_{t,y}}{(m_{j,x} - \bar{\mu}_{t,x})^2 + (m_{j,y} - \bar{\mu}_{t,y})^2} & -\frac{m_{j,x} - \bar{\mu}_{t,x}}{(m_{j,x} - \bar{\mu}_{t,x})^2 + (m_{j,y} - \bar{\mu}_{t,y})^2} & -1 \end{bmatrix}. \quad (8)$$

It is also common to assume that the covariance of the measurement noise is given by:

$$\mathbf{Q}_t = \begin{bmatrix} \sigma_r^2 & 0 \\ 0 & \sigma_\phi^2 \end{bmatrix},$$

where σ_r is the standard deviation of the range measurement noise and σ_ϕ is the standard deviation of the bearing measurement noise. This diagonal covariance matrix is typically used since these two measurements can be assumed to be uncorrelated.

17.5 Monte Carlo Localization (MCL)

Another approach to Markov localization is the Monte Carlo localization (MCL) algorithm. This algorithm leverages the non-parametric particle filter algorithm from the previous chapter, and is therefore much better suited to solving *global* localization problems (unlike EKF localization which only solves position tracking problems). MCL can also be used to solve the kidnapped robot problem through some small modifications, such as injecting new random particles at each step to ensure that a “particle collapse” problem does not occur.

As a brief review, the particle filter represents the belief $bel(\mathbf{x}_t)$ by a set of M particles:

$$\mathcal{X}_t := \{\mathbf{x}_t^{[1]}, \mathbf{x}_t^{[2]}, \dots, \mathbf{x}_t^{[M]}\},$$

where each particle $\mathbf{x}_t^{[m]}$ represents a hypothesis about the true state \mathbf{x}_t . At each step of the algorithm the state transition model is used to propagate forward the particles, and then the measurement model is used to resample particles based on the measurement likelihood. This algorithm is shown in Algorithm 4, and is nearly identical to the particle filter algorithm except that the map \mathbf{m} is used in the probabilistic state transition and measurement models.

Data: $\mathcal{X}_{t-1}, \mathbf{u}_t, \mathbf{z}_t, \mathbf{m}$
Result: \mathcal{X}_t
 $\bar{\mathcal{X}}_t = \mathcal{X}_t = \emptyset$
for $m = 1$ **to** M **do**
 | Sample $\bar{\mathbf{x}}_t^{[m]} \sim p(\mathbf{x}_t \mid \mathbf{u}_t, \mathbf{x}_{t-1}^{[m]}, \mathbf{m})$
 | $w_t^{[m]} = p(\mathbf{z}_t \mid \bar{\mathbf{x}}_t^{[m]}, \mathbf{m})$
 | $\bar{\mathcal{X}}_t = \bar{\mathcal{X}}_t \cup (\bar{\mathbf{x}}_t^{[m]}, w_t^{[m]})$
end
for $m = 1$ **to** M **do**
 | Draw i with probability $\propto w_t^{[i]}$
 | Add $\bar{\mathbf{x}}_t^{[i]}$ to \mathcal{X}_t
end
return \mathcal{X}_t

Algorithm 4: Monte Carlo Localization Algorithm

References

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