




Hochschule
Bonn-Rhein-Sieg
University of Applied Sciences



Monte Carlo Localisation

Localising Using a Particle Filter

Dr. Alex Mitrevski
Master of Autonomous Systems




- ▶ Reminder: Robot localisation essentials
- ▶ Monte Carlo localisation
- ▶ Probabilistic motion models

Reminder: Robot Localisation Essentials



Reminder: Localisation



- ▶ The process that determines a robot's pose in an environment is called **localisation**
- ▶ The localisation process is performed **with respect to a given map** — thus, **localisation requires a map to be given**
- ▶ Since a robot moves around, the localisation estimate needs to be continuously updated based on the **known motion commands** and **any relevant environment observations**

Reminder: Recursive State Estimation Using a Bayes Filter



- ▶ The process of continuously updating the belief based on performed motions and recorded measurements is called **recursive state estimation**; this is performed using a **Bayes filter**
- ▶ The Bayes filter continuously performs two steps:

Control update (prediction)

The belief is updated based on a performed motion \mathbf{u}_t :

$$\overline{\text{bel}}(\mathbf{x}_t) = \int p(\mathbf{x}_t | \mathbf{u}_t, \mathbf{x}_{t-1}) \text{bel}(\mathbf{x}_{t-1}) d\mathbf{x}_{t-1}$$

Measurement update (correction)

The belief is corrected based on a sensor measurement z_t :

$$\text{bel}(\mathbf{x}_t) = \eta p(z_t | \mathbf{x}_t) \overline{\text{bel}}(\mathbf{x}_t)$$

- ▶ The control update **increases the pose uncertainty**, while the measurement update **decreases it**
- ▶ Both update equations are based on the **Markov assumption**, according to which:
 - ▶ \mathbf{x}_t only depends on \mathbf{x}_{t-1} and \mathbf{u}_t
 - ▶ z_t only depends on \mathbf{x}_t


Recursive Bayes Filter Family

- ▶ The Bayes filter is not a single algorithm, but **a family of algorithms**, each of which is based on different assumptions
- ▶ We will particularly distinguish between the following variants of the recursive Bayes filter as applied to the robot localisation problem:
 - ▶ **Discrete Bayes filter**: Applicable when **the state is discretised** (discussed in MRC)
 - ▶ **Kalman filter**: Assumes that **the state is governed by a Gaussian distribution** (discussed last time)
 - ▶ **Particle filter**: Does not make an assumption about the underlying state distribution; **represents multiple hypotheses about the state by a set of particles** (discussed today)

Monte Carlo Localisation




- ▶ A particle filter is a **non-parametric representation of a Bayes filter**



- ▶ A particle filter is a **non-parametric representation of a Bayes filter**
- ▶ The filter represents a distribution by n particles

$$P = \{\mathbf{p}^1, \dots, \mathbf{p}^n\}$$

where **each particle is a weighted hypothesis** about the state, namely $\mathbf{p}^i = (\mathbf{x}^i, w^i)$



- ▶ A particle filter is a **non-parametric representation of a Bayes filter**


- ▶ The filter represents a distribution by n particles


$$P = \{\mathbf{p}^1, \dots, \mathbf{p}^n\}$$

where **each particle is a weighted hypothesis** about the state, namely $\mathbf{p}^i = (\mathbf{x}^i, w^i)$

- ▶ The objective is to include particles $\mathbf{p}^i, 1 \leq i \leq n$ in P with a **probability that is proportional to the estimated posterior distribution**:

$$\mathbf{x}_{t+1}^i \sim p(\mathbf{x}_{t+1} | \mathbf{x}_{0:t}, \mathbf{u}_{0:t}, \mathbf{z}_{1:t})$$





- ▶ A particle filter is a **non-parametric representation of a Bayes filter**

- ▶ The filter represents a distribution by n particles

$$P = \{\mathbf{p}^1, \dots, \mathbf{p}^n\}$$

where **each particle is a weighted hypothesis** about the state, namely $\mathbf{p}^i = (\mathbf{x}^i, w^i)$

- ▶ The objective is to include particles $\mathbf{p}^i, 1 \leq i \leq n$ in P with a **probability that is proportional to the estimated posterior distribution**:

$$\mathbf{x}_{t+1}^i \sim p(\mathbf{x}_{t+1} | \mathbf{x}_{0:t}, \mathbf{u}_{0:t}, \mathbf{z}_{1:t})$$

- ▶ A particle filter is a general instantiation of the Bayes filter and has been used in a variety of problems, for instance localisation, object tracking, or fault diagnosis

- ▶ In a particle filter, **both the motion and the measurement update are operations on the particle set P_t at time t :**

```
1: function PARTICLEFILTER( $P_t, \mathbf{u}_t, \mathbf{z}_t, \mathcal{M}$ )
2:    $\hat{P} \leftarrow \{\}$ 
3:   sampling_dist  $\leftarrow []$ 
4:   for  $i \leftarrow 1$  to  $n$  do
5:      $\mathbf{x}_t^i \leftarrow P^i.x$ 
6:      $\mathbf{x}_{t+1}^i \sim p(\mathbf{x}_{t+1} | \mathbf{x}_t^i, \mathbf{u}_t)$ 
7:      $w^i \leftarrow p(\mathbf{z}_{t+1} | \mathbf{x}_{t+1}^i, \mathcal{M})$ 
8:      $\hat{P} \leftarrow \hat{P} \cup (\mathbf{x}_{t+1}^i, w^i)$ 
9:     sampling_dist.append( $w^i$ )
10:  sampling_dist  $\leftarrow \eta \cdot \text{sampling\_dist}$ 
11:   $P_{t+1} \leftarrow \{\}$ 
12:  for  $i \leftarrow 1$  to  $n$  do
13:     $p^i \leftarrow \text{sample}(\hat{P}, \text{sampling\_dist})$ 
14:     $P_{t+1} \leftarrow P_{t+1} \cup p^i$ 
15:  return  $P_{t+1}$ 
```

```
1: function PARTICLEFILTER( $P_t, \mathbf{u}_t, \mathbf{z}_t, \mathcal{M}$ )
2:    $\hat{P} \leftarrow \{\}$ 
3:   sampling_dist  $\leftarrow []$ 
4:   for  $i \leftarrow 1$  to  $n$  do
5:      $\mathbf{x}_t^i \leftarrow P^i \cdot \mathbf{x}$ 
6:      $\mathbf{x}_{t+1}^i \sim p(\mathbf{x}_{t+1} | \mathbf{x}_t^i, \mathbf{u}_t)$ 
7:      $w^i \leftarrow p(\mathbf{z}_{t+1} | \mathbf{x}_{t+1}^i, \mathcal{M})$ 
8:      $\hat{P} \leftarrow \hat{P} \cup (\mathbf{x}_{t+1}^i, w^i)$ 
9:     sampling_dist.append( $w^i$ )
10:  sampling_dist  $\leftarrow \eta \cdot$  sampling_dist
11:   $P_{t+1} \leftarrow \{\}$ 
12:  for  $i \leftarrow 1$  to  $n$  do
13:     $p^i \leftarrow$  sample( $\hat{P}$ , sampling_dist )
14:     $P_{t+1} \leftarrow P_{t+1} \cup p^i$ 
15:  return  $P_{t+1}$ 
```

► In a particle, filter, **both the motion and the measurement update are operations on the particle set P_t at time t :**

► **Motion update:** The states of the particles are updated by sampling from a **motion model**

```
1: function PARTICLEFILTER( $P_t, \mathbf{u}_t, \mathbf{z}_t, \mathcal{M}$ )
2:    $\hat{P} \leftarrow \{\}$ 
3:   sampling_dist  $\leftarrow []$ 
4:   for  $i \leftarrow 1$  to  $n$  do
5:      $\mathbf{x}_t^i \leftarrow P^i \cdot \mathbf{x}$ 
6:      $\mathbf{x}_{t+1}^i \sim p(\mathbf{x}_{t+1} | \mathbf{x}_t^i, \mathbf{u}_t)$ 
7:      $w^i \leftarrow p(\mathbf{z}_{t+1} | \mathbf{x}_{t+1}^i, \mathcal{M})$ 
8:      $\hat{P} \leftarrow \hat{P} \cup (\mathbf{x}_{t+1}^i, w^i)$ 
9:     sampling_dist.append( $w^i$ )
10:  sampling_dist  $\leftarrow \eta \cdot$  sampling_dist
11:   $P_{t+1} \leftarrow \{\}$ 
12:  for  $i \leftarrow 1$  to  $n$  do
13:     $p^i \leftarrow$  sample( $\hat{P}$ , sampling_dist )
14:     $P_{t+1} \leftarrow P_{t+1} \cup p^i$ 
15:  return  $P_{t+1}$ 
```

► In a particle filter, **both the motion and the measurement update are operations on the particle set P_t at time t :**

- **Motion update:** The states of the particles are updated by sampling from a **motion model**
- **Measurement update:** The particles' weights are updated based on the **measurement model** with respect to a map \mathcal{M}

```
1: function PARTICLEFILTER( $P_t, \mathbf{u}_t, \mathbf{z}_t, \mathcal{M}$ )
2:    $\hat{P} \leftarrow \{\}$ 
3:   sampling_dist  $\leftarrow []$ 
4:   for  $i \leftarrow 1$  to  $n$  do
5:      $\mathbf{x}_t^i \leftarrow P^i \cdot \mathbf{x}$ 
6:      $\mathbf{x}_{t+1}^i \sim p(\mathbf{x}_{t+1} | \mathbf{x}_t^i, \mathbf{u}_t)$ 
7:      $w^i \leftarrow p(\mathbf{z}_{t+1} | \mathbf{x}_{t+1}^i, \mathcal{M})$ 
8:      $\hat{P} \leftarrow \hat{P} \cup (\mathbf{x}_{t+1}^i, w^i)$ 
9:     sampling_dist.append( $w^i$ )
10:  sampling_dist  $\leftarrow \eta \cdot$  sampling_dist
11:   $P_{t+1} \leftarrow \{\}$ 
12:  for  $i \leftarrow 1$  to  $n$  do
13:     $p^i \leftarrow \text{sample}(\hat{P}, \text{sampling\_dist})$ 
14:     $P_{t+1} \leftarrow P_{t+1} \cup p^i$ 
15:  return  $P_{t+1}$ 
```

- ▶ In a particle filter, **both the motion and the measurement update are operations on the particle set P_t at time t :**
 - ▶ **Motion update:** The states of the particles are updated by sampling from a **motion model**
 - ▶ **Measurement update:** The particles' weights are updated based on the **measurement model** with respect to a map \mathcal{M}
- ▶ After these two steps, the posterior distribution is updated — this is achieved by **resampling the particle set**

```
1: function PARTICLEFILTER( $P_t, \mathbf{u}_t, \mathbf{z}_t, \mathcal{M}$ )
2:    $\hat{P} \leftarrow \{\}$ 
3:   sampling_dist  $\leftarrow []$ 
4:   for  $i \leftarrow 1$  to  $n$  do
5:      $\mathbf{x}_t^i \leftarrow P^i \cdot \mathbf{x}$ 
6:      $\mathbf{x}_{t+1}^i \sim p(\mathbf{x}_{t+1} | \mathbf{x}_t^i, \mathbf{u}_t)$ 
7:      $w^i \leftarrow p(\mathbf{z}_{t+1} | \mathbf{x}_{t+1}^i, \mathcal{M})$ 
8:      $\hat{P} \leftarrow \hat{P} \cup (\mathbf{x}_{t+1}^i, w^i)$ 
9:     sampling_dist.append( $w^i$ )
10:  sampling_dist  $\leftarrow \eta \cdot$  sampling_dist
11:   $P_{t+1} \leftarrow \{\}$ 
12:  for  $i \leftarrow 1$  to  $n$  do
13:     $p^i \leftarrow \text{sample}(\hat{P}, \text{sampling\_dist})$ 
14:     $P_{t+1} \leftarrow P_{t+1} \cup p^i$ 
15:  return  $P_{t+1}$ 
```

- ▶ In a particle filter, **both the motion and the measurement update are operations on the particle set P_t at time t :**
 - ▶ **Motion update:** The states of the particles are updated by sampling from a **motion model**
 - ▶ **Measurement update:** The particles' weights are updated based on the **measurement model** with respect to a map \mathcal{M}
- ▶ After these two steps, the posterior distribution is updated — this is achieved by **resampling the particle set**
- ▶ Resampling is performed using a procedure called **importance sampling**, during which particles from P_t are **sampled with replacement proportional to their weights**; this results in a new particle set P_{t+1}

Importance Sampling

- ▶ Importance sampling is a technique based on which samples are generated from a **proposal distribution** $g(x)$ instead of a **target distribution** $f(x)$
 - ▶ g is defined so that, if $f(x) > 0$, then it should also be the case that $g(x) > 0$



Importance Sampling

- ▶ Importance sampling is a technique based on which samples are generated from a **proposal distribution** $g(x)$ instead of a **target distribution** $f(x)$
 - ▶ g is defined so that, if $f(x) > 0$, then it should also be the case that $g(x) > 0$
- ▶ Given a random variable X with a density $f(x)$, its expected value can be calculated as

$$\begin{aligned}
 E_f[X] &= \int f(x)x dx \\
 &= \int \frac{f(x)}{g(x)}g(x)x dx \\
 &= \int w(x)g(x)x dx = E_g[W(X)X]
 \end{aligned}$$

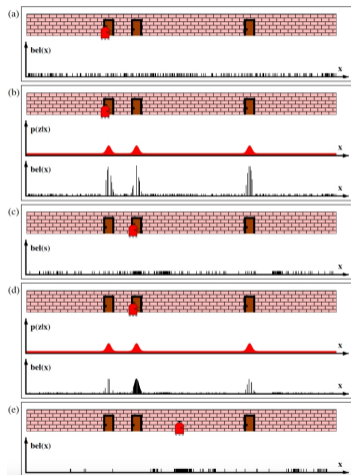
Importance Sampling

- ▶ Importance sampling is a technique based on which samples are generated from a **proposal distribution** $g(x)$ instead of a **target distribution** $f(x)$
 - ▶ g is defined so that, if $f(x) > 0$, then it should also be the case that $g(x) > 0$
- ▶ Given a random variable X with a density $f(x)$, its expected value can be calculated as

$$\begin{aligned}
 E_f[X] &= \int f(x)x dx \\
 &= \int \frac{f(x)}{g(x)}g(x)x dx \\
 &= \int w(x)g(x)x dx = E_g[W(X)X]
 \end{aligned}$$

- ▶ In a particle filter:

- ▶ the proposal distribution $g(\mathbf{x}) = p(\mathbf{x}_{t+1}|\mathbf{u}_{t+1}, \mathbf{x}_t) \underbrace{p(\mathbf{x}_{0:t}|\mathbf{u}_{0:t}, \mathbf{z}_{1:t})}_{\text{bel}(\mathbf{x}_t)}$
- ▶ the target distribution $f(\mathbf{x}) = p(\mathbf{z}_{t+1}|\mathbf{x}_{t+1})p(\mathbf{x}_{t+1}|\mathbf{u}_{t+1}, \mathbf{x}_t) \underbrace{p(\mathbf{x}_{0:t}|\mathbf{u}_{0:t}, \mathbf{z}_{1:t})}_{\text{bel}(\mathbf{x}_t)}$
- ▶ the weights $w^i \propto p(\mathbf{z}_{t+1}|\mathbf{x}_{t+1})$




- ▶ Localisation performed using particle filters is called **Monte Carlo localisation**
 - ▶ Monte Carlo methods represent a general category of techniques based on which solutions to problems are discovered using random sampling
- ▶ In the most general case, the initial state of the robot is fully unknown; this can be expressed by **sampling the initial particle set P_0 from a uniform distribution**

Particle Filter Illustration: Motion Update

During the motion update step, each particle is updated according to the motion model


```
1: function PARTICLEFILTER( $P_t, \mathbf{u}_t, \mathbf{z}_t, \mathcal{M}$ )
2:    $\hat{P} \leftarrow \{\}$ 
3:   sampling_dist  $\leftarrow []$ 
4:   for  $i \leftarrow 1$  to  $n$  do
5:      $\mathbf{x}_t^i \leftarrow P^i \cdot \mathbf{x}$ 
6:      $\mathbf{x}_{t+1}^i \sim p(\mathbf{x}_{t+1}^i | \mathbf{x}_t^i, \mathbf{u}_t)$ 
7:      $w^i \leftarrow p(\mathbf{z}_{t+1} | \mathbf{x}_{t+1}^i, \mathcal{M})$ 
8:      $\hat{P} \leftarrow \hat{P} \cup (\mathbf{x}_{t+1}^i, w^i)$ 
9:     sampling_dist.append( $w^i$ )
10:  sampling_dist  $\leftarrow \eta \cdot \text{sampling\_dist}$ 
11:   $P_{t+1} \leftarrow \{\}$ 
12:  for  $i \leftarrow 1$  to  $n$  do
13:     $p^i \leftarrow \text{sample}(\hat{P}, \text{sampling\_dist})$ 
14:     $P_{t+1} \leftarrow P_{t+1} \cup p^i$ 
15:  return  $P_{t+1}$ 
```



Particle Filter Illustration: Motion Update

During the motion update step, each particle is updated according to the motion model


```
1: function PARTICLEFILTER( $P_t, \mathbf{u}_t, \mathbf{z}_t, \mathcal{M}$ )
2:    $\hat{P} \leftarrow \{\}$ 
3:   sampling_dist  $\leftarrow []$ 
4:   for  $i \leftarrow 1$  to  $n$  do
5:      $\mathbf{x}_t^i \leftarrow P^i \cdot \mathbf{x}$ 
6:      $\mathbf{x}_{t+1}^i \sim p(\mathbf{x}_{t+1} | \mathbf{x}_t^i, \mathbf{u}_t)$ 
7:      $w^i \leftarrow p(\mathbf{z}_{t+1} | \mathbf{x}_{t+1}^i, \mathcal{M})$ 
8:      $\hat{P} \leftarrow \hat{P} \cup (\mathbf{x}_{t+1}^i, w^i)$ 
9:     sampling_dist.append( $w^i$ )
10:  sampling_dist  $\leftarrow \eta \cdot$  sampling_dist
11:   $P_{t+1} \leftarrow \{\}$ 
12:  for  $i \leftarrow 1$  to  $n$  do
13:     $p^i \leftarrow \text{sample}(\hat{P}, \text{sampling\_dist})$ 
14:     $P_{t+1} \leftarrow P_{t+1} \cup p^i$ 
15:  return  $P_{t+1}$ 
```



Particle Filter Illustration: Measurement Update

During the measurement update, each particle is reweighed based on a given observation — particles that result in likelier measurements are assigned higher weights


```
1: function PARTICLEFILTER( $P_t, \mathbf{u}_t, \mathbf{z}_t, \mathcal{M}$ )
2:    $\hat{P} \leftarrow \{\}$ 
3:   sampling_dist  $\leftarrow []$ 
4:   for  $i \leftarrow 1$  to  $n$  do
5:      $\mathbf{x}_t^i \leftarrow P_t^i \cdot \mathbf{x}$ 
6:      $\mathbf{x}_{t+1}^i \sim p(\mathbf{x}_{t+1} | \mathbf{x}_t^i, \mathbf{u}_t)$ 
7:      $w^i \leftarrow p(\mathbf{z}_{t+1} | \mathbf{x}_{t+1}^i, \mathcal{M})$ 
8:      $\hat{P} \leftarrow \hat{P} \cup (\mathbf{x}_{t+1}^i, w^i)$ 
9:     sampling_dist.append( $w^i$ )
10:  sampling_dist  $\leftarrow \eta \cdot$  sampling_dist
11:   $P_{t+1} \leftarrow \{\}$ 
12:  for  $i \leftarrow 1$  to  $n$  do
13:     $p^i \leftarrow \text{sample}(\hat{P}, \text{sampling\_dist})$ 
14:     $P_{t+1} \leftarrow P_{t+1} \cup p^i$ 
15:  return  $P_{t+1}$ 
```



Particle Filter Illustration: Particle Resampling

The particles are finally resampled (with replacement) proportional to their weights

```
1: function PARTICLEFILTER( $P_t, \mathbf{u}_t, \mathbf{z}_t, \mathcal{M}$ )
2:    $\hat{P} \leftarrow \{\}$ 
3:   sampling_dist  $\leftarrow []$ 
4:   for  $i \leftarrow 1$  to  $n$  do
5:      $\mathbf{x}_t^i \leftarrow P^i.\mathbf{x}$ 
6:      $\mathbf{x}_{t+1}^i \sim p(\mathbf{x}_{t+1}|\mathbf{x}_t^i, \mathbf{u}_t)$ 
7:      $w^i \leftarrow p(\mathbf{z}_{t+1}|\mathbf{x}_{t+1}^i, \mathcal{M})$ 
8:      $\hat{P} \leftarrow \hat{P} \cup (\mathbf{x}_{t+1}^i, w^i)$ 
9:     sampling_dist.append( $w^i$ )
10:  sampling_dist  $\leftarrow \eta \cdot \text{sampling\_dist}$ 
11:   $P_{t+1} \leftarrow \{\}$ 
12:  for  $i \leftarrow 1$  to  $n$  do
13:     $p^i \leftarrow \text{sample}(\hat{P}, \text{sampling\_dist})$ 
14:     $P_{t+1} \leftarrow P_{t+1} \cup p^i$ 
15:  return  $P_{t+1}$ 
```



- ▶ In localisation, we are typically **interested in a single estimate of a robot's state** — the states of all particles are not of direct interest




Figure 4.5 Different ways of extracting densities from particles. (a) Density and sample set approximation, (b) Gaussian approximation (mean and variance), (c) histogram approximation, (d) kernel density estimate. The choice of approximation strongly depends on the specific application and the computational resources.




Figure 4.5 Different ways of extracting densities from particles. (a) Density and sample set approximation, (b) Gaussian approximation (mean and variance), (c) histogram approximation, (d) kernel density estimate. The choice of approximation strongly depends on the specific application and the computational resources.

- ▶ In localisation, we are typically **interested in a single estimate of a robot's state** — the states of all particles are not of direct interest
- ▶ One way to achieve this is to **represent the state as the weighted mean of the particles**: $\tilde{x}_t = \sum_{i=1}^n w^i x_t^i$
 - ▶ This corresponds to **approximating the unknown distribution $f(x)$ by a Gaussian distribution**
 - ▶ The variance of this distribution provides **information about how certain we can be in the state approximation**
 - ▶ The variance information can be used as a heuristic to **reinitialise the filter**

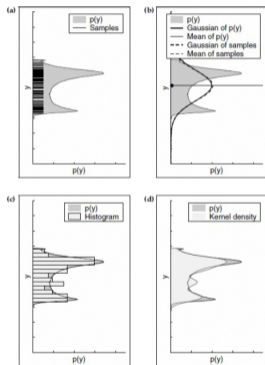


Figure 4.5 Different ways of extracting densities from particles. (a) Density and sample set approximation, (b) Gaussian approximation (mean and variance), (c) histogram approximation, (d) kernel density estimate. The choice of approximation strongly depends on the specific application and the computational resources.

- ▶ In localisation, we are typically **interested in a single estimate of a robot's state** — the states of all particles are not of direct interest
- ▶ One way to achieve this is to **represent the state as the weighted mean of the particles**: $\tilde{x}_t = \sum_{i=1}^n w^i x_t^i$
 - ▶ This corresponds to **approximating the unknown distribution $f(x)$ by a Gaussian distribution**
 - ▶ The variance of this distribution provides **information about how certain we can be in the state approximation**
 - ▶ The variance information can be used as a heuristic to **reinitialise the filter**
- ▶ The particles can also be used to approximate the density more accurately, for instance using a **histogram** or with **kernel density estimation**

Particle Deprivation and Kidnapped Robot Problem Avoidance




- ▶ Unlike the Kalman filter, the particle filter can, in principle, be used to solve the kidnapped robot problem and thus **recover from localisation failures**
 - ▶ This is due to the fact that the particle filter can model multimodal distributions, while the Kalman filter represents the underlying state distribution by a unimodal Gaussian distribution



- ▶ Unlike the Kalman filter, the particle filter can, in principle, be used to solve the kidnapped robot problem and thus **recover from localisation failures**
 - ▶ This is due to the fact that the particle filter can model multimodal distributions, while the Kalman filter represents the underlying state distribution by a unimodal Gaussian distribution
- ▶ The most common way to avoid localisation failures is to **include random particles in P_t during the resampling step**, typically sampled from a uniform distribution

- ▶ Unlike the Kalman filter, the particle filter can, in principle, be used to solve the kidnapped robot problem and thus **recover from localisation failures**
 - ▶ This is due to the fact that the particle filter can model multimodal distributions, while the Kalman filter represents the underlying state distribution by a unimodal Gaussian distribution
- ▶ The most common way to avoid localisation failures is to **include random particles in P_t during the resampling step**, typically sampled from a uniform distribution
- ▶ While random particles can be included continuously and with the same frequency, a better strategy for this is to **couple the inclusion of random particles to the evolution of the measurement likelihood**
 - ▶ This is an intuitive principle based on which more particles are included only if the particles' measurement likelihood reduces over time


Adaptive Monte Carlo Localisation (AMCL)



- ▶ The basic particle filter algorithm uses **a fixed number of samples n**
 - ▶ The number of particles required for accurately approximating the state is typically high (in the order of thousands or even tens of thousands) — although this depends on the map size



Adaptive Monte Carlo Localisation (AMCL)



- ▶ The basic particle filter algorithm uses **a fixed number of samples n**
 - ▶ The number of particles required for accurately approximating the state is typically high (in the order of thousands or even tens of thousands) — although this depends on the map size
- ▶ In practice, using a fixed number of particles is computationally suboptimal, **particularly once the particles converge to a state estimate**



Adaptive Monte Carlo Localisation (AMCL)



- ▶ The basic particle filter algorithm uses **a fixed number of samples n**
 - ▶ The number of particles required for accurately approximating the state is typically high (in the order of thousands or even tens of thousands) — although this depends on the map size
- ▶ In practice, using a fixed number of particles is computationally suboptimal, **particularly once the particles converge to a state estimate**
- ▶ A practically more useful approach is to **vary the number of particles depending on the distribution approximation quality** — this is the approach followed by commonly used localisation frameworks, such as AMCL in ROS



Adaptive Monte Carlo Localisation (AMCL)



- ▶ The basic particle filter algorithm uses **a fixed number of samples n**
 - ▶ The number of particles required for accurately approximating the state is typically high (in the order of thousands or even tens of thousands) — although this depends on the map size
- ▶ In practice, using a fixed number of particles is computationally suboptimal, **particularly once the particles converge to a state estimate**
- ▶ A practically more useful approach is to **vary the number of particles depending on the distribution approximation quality** — this is the approach followed by commonly used localisation frameworks, such as AMCL in ROS
- ▶ AMCL modifies the particle resampling step so that **particles are sampled until a desired distribution quality is reached**
 - ▶ The distribution quality is measured by **overlaying a discrete grid over the map** and **calculating an approximation bound based on the number of grid cells with non-zero particles**
 - ▶ The concrete sampling technique used by AMCL is called **KLD-sampling**



- ▶ The Monte Carlo localisation algorithm discussed above is actually only applicable to static environments — **dynamic obstacles will have a direct effect on the measurements, but this effect is not considered by the measurement model**

- ▶ The Monte Carlo localisation algorithm discussed above is actually only applicable to static environments — **dynamic obstacles will have a direct effect on the measurements, but this effect is not considered by the measurement model**
- ▶ Two main strategies can be used to handle dynamic obstacles during localisation:

- ▶ The Monte Carlo localisation algorithm discussed above is actually only applicable to static environments — **dynamic obstacles will have a direct effect on the measurements, but this effect is not considered by the measurement model**
- ▶ Two main strategies can be used to handle dynamic obstacles during localisation:
 - ▶ **Modelling dynamic obstacles in the measurement model:** This is the most general way of handling dynamic obstacles, but is typically inefficient — a variable state dimensionality needs to be considered

- ▶ The Monte Carlo localisation algorithm discussed above is actually only applicable to static environments — **dynamic obstacles will have a direct effect on the measurements, but this effect is not considered by the measurement model**
- ▶ Two main strategies can be used to handle dynamic obstacles during localisation:
 - ▶ **Modelling dynamic obstacles in the measurement model:** This is the most general way of handling dynamic obstacles, but is typically inefficient — a variable state dimensionality needs to be considered
 - ▶ **Outlier rejection:** Using this strategy, measurements from dynamic obstacles are rejected before the particle filter performs a measurement update

- ▶ The Monte Carlo localisation algorithm discussed above is actually only applicable to static environments — **dynamic obstacles will have a direct effect on the measurements, but this effect is not considered by the measurement model**
- ▶ Two main strategies can be used to handle dynamic obstacles during localisation:
 - ▶ **Modelling dynamic obstacles in the measurement model**: This is the most general way of handling dynamic obstacles, but is typically inefficient — a variable state dimensionality needs to be considered
 - ▶ **Outlier rejection**: Using this strategy, measurements from dynamic obstacles are rejected before the particle filter performs a measurement update
- ▶ Particularly outlier rejection requires a **model of dynamic obstacles**; such a model may be learned

Probabilistic Motion Models



- ▶ The particle filter performs a motion update of each particle based on a motion command \mathbf{u}_{t+1} ; this update is **not deterministic, but probabilistic** — by sampling from a motion model $p(\mathbf{x}_{t+1}|\mathbf{u}_{t+1}, \mathbf{x}_t)$ — so that it takes the motion uncertainty into account
- ▶ The used motion model depends on the **type of motion command that is sent to a robot**
- ▶ We will briefly look at two types of **motion models for planar motion**:
 - ▶ **Velocity motion model**: The motion command is specified by linear and angular velocities:
 $\mathbf{u} = (v, \omega)^T$
 - ▶ **Odometry motion model**: The motion command is given by relative odometry $\mathbf{u} = (\bar{\mathbf{x}}_t, \bar{\mathbf{x}}_{t+1})^T$, where $\bar{\mathbf{x}}_t = (\bar{x}, \bar{y}, \bar{\theta})^T$ and $\bar{\mathbf{x}}_{t+1} = (\bar{x}', \bar{y}', \bar{\theta}')$

Velocity Motion Model

- ▶ Given a velocity command $(v, \omega)^T$, a robot moves along a circle of radius $r = \left| \frac{v}{\omega} \right|$ with center $\left(x - \frac{v}{\omega} \sin \theta, y + \frac{v}{\omega} \cos \theta \right)^T$

- ▶ Based on the velocity motion model, a robot is **controlled by a velocity command under the influence of noise**:




Figure 5.3 The velocity motion model, for different noise parameter settings.

$$\begin{pmatrix} \hat{v} \\ \hat{\omega} \end{pmatrix} = \begin{pmatrix} v \\ \omega \end{pmatrix} + \begin{pmatrix} \epsilon_{\alpha_1} v^2 + \alpha_2 \omega^2 \\ \epsilon_{\alpha_3} v^2 + \alpha_4 \omega^2 \end{pmatrix}$$

- ▶ The robot is also **assumed to undergo a final orientation once the goal location is reached**; the final orientation is thus given by $\theta' = \theta + \hat{\omega} \Delta t + \hat{\gamma} \Delta t$ with $\hat{\gamma} = \epsilon_{\alpha_5} v^2 + \alpha_6 \omega^2$

- ▶ Under the velocity motion model, the pose update is given as

$$\begin{pmatrix} x' \\ y' \\ \theta' \end{pmatrix} = \begin{pmatrix} x \\ y \\ \theta \end{pmatrix} + \begin{pmatrix} -\frac{\hat{v}}{\hat{\omega}} \sin \theta + \frac{\hat{v}}{\hat{\omega}} \sin(\theta + \hat{\omega} \Delta t) \\ \frac{\hat{v}}{\hat{\omega}} \cos \theta - \frac{\hat{v}}{\hat{\omega}} \cos(\theta + \hat{\omega} \Delta t) \\ \hat{\omega} \Delta t + \hat{\gamma} \Delta t \end{pmatrix}$$

- ▶ In the particle filter, we are interested in sampling from the velocity motion model; this sampling can be done using the algorithm below

```
1: function SAMPLEVELOCITYMOTIONMODEL( $\mathbf{x}_t = (x, y, \theta)$ ,  $\mathbf{u}_t = (v, \omega)$ )
2:    $\hat{v} \leftarrow v + \text{sample}(\alpha_1 v^2 + \alpha_2 \omega^2)$ 
3:    $\hat{\omega} \leftarrow \omega + \text{sample}(\alpha_3 v^2 + \alpha_4 \omega^2)$ 
4:    $\hat{\gamma} \leftarrow \text{sample}(\alpha_5 v^2 + \alpha_6 \omega^2)$ 
5:    $x' \leftarrow x - \frac{\hat{v}}{\hat{\omega}} \sin \theta + \frac{\hat{v}}{\hat{\omega}} \sin(\theta + \hat{\omega} \Delta t)$ 
6:    $y' \leftarrow y + \frac{\hat{v}}{\hat{\omega}} \cos \theta - \frac{\hat{v}}{\hat{\omega}} \cos(\theta + \hat{\omega} \Delta t)$ 
7:    $\theta' \leftarrow \theta + \hat{\omega} \Delta t + \hat{\gamma} \Delta t$ 
8:   return  $(x', y', \theta')^T$ 
```




Figure 5.4 Sampling from the velocity motion model, using the same parameters as in Figure 5.3. Each diagram shows 500 samples.

Odometry Motion Model

- Given an initial pose x_t and a final pose x_{t+1} , the motion can be modelled by an **initial rotation** δ_{rot1} towards the goal, a **translational motion** δ_{trans} , and a **final rotation** δ_{rot2}




Figure 5.7 Odometry model: The robot motion in the time interval $(t-1, t)$ is approximated by a rotation δ_{rot1} , followed by a translation δ_{trans} and a second rotation δ_{rot2} . The turns and translations are noisy.

- Under the odometry motion model, **the components of the motion are all under the influence of noise:**

$$\begin{pmatrix} \hat{\delta}_{rot1} \\ \hat{\delta}_{rot2} \\ \hat{\delta}_{trans} \end{pmatrix} = \begin{pmatrix} \delta_{rot1} \\ \delta_{rot2} \\ \delta_{trans} \end{pmatrix} + \begin{pmatrix} \epsilon_{\alpha_1} \delta_{rot1}^2 + \alpha_2 \delta_{trans}^2 \\ \epsilon_{\alpha_1} \delta_{rot1}^2 + \alpha_2 \delta_{trans}^2 \\ \epsilon_{\alpha_3} \delta_{trans}^2 + \alpha_4 \delta_{rot1}^2 + \alpha_4 \delta_{rot2}^2 \end{pmatrix}$$




Figure 5.8 The odometry motion model, for different noise parameter settings.

- The pose update under the odometry motion model is thus given as

$$\begin{pmatrix} x' \\ y' \\ \theta' \end{pmatrix} = \begin{pmatrix} x \\ y \\ \theta \end{pmatrix} + \begin{pmatrix} \hat{\delta}_{trans} \cos(\theta + \hat{\delta}_{rot1}) \\ \hat{\delta}_{trans} \sin(\theta + \hat{\delta}_{rot1}) \\ \hat{\delta}_{rot1} + \hat{\delta}_{rot2} \end{pmatrix}$$

Odometry Motion Model Sampling

► The algorithm below can be used for sampling from the odometry motion model

- 1: **function** SAMPLEODOMETRYMOTIONMODEL($\mathbf{x}_t = (x, y, \theta)$, $\mathbf{u}_t = (\bar{x}_{t-1}, \bar{x}_t)$)
- 2: $\delta_{\text{rot}_1} \leftarrow \text{atan2}(\bar{y}' - \bar{y}, \bar{x}' - \bar{x}) - \theta$
- 3: $\delta_{\text{trans}} \leftarrow \sqrt{(\bar{x}' - \bar{x})^2 + (\bar{y}' - \bar{y})^2}$
- 4: $\delta_{\text{rot}_2} \leftarrow \bar{\theta}' - \bar{\theta} - \delta_{\text{rot}_1}$
- 5: $\hat{\delta}_{\text{rot}_1} \leftarrow \delta_{\text{rot}_1} - \text{sample}(\alpha_1 \delta_{\text{rot}_1}^2 + \alpha_2 \delta_{\text{trans}}^2)$
- 6: $\hat{\delta}_{\text{trans}} \leftarrow \delta_{\text{trans}} - \text{sample}(\alpha_3 \delta_{\text{trans}}^2 + \alpha_4 \delta_{\text{rot}_1}^2 + \alpha_4 \delta_{\text{rot}_2}^2)$
- 7: $\hat{\delta}_{\text{rot}_2} \leftarrow \delta_{\text{rot}_2} - \text{sample}(\alpha_1 \delta_{\text{rot}_2}^2 + \alpha_2 \delta_{\text{trans}}^2)$
- 8: $x' \leftarrow x + \hat{\delta}_{\text{trans}} \cos(\theta + \hat{\delta}_{\text{rot}_1})$
- 9: $y' \leftarrow y + \hat{\delta}_{\text{trans}} \sin(\theta + \hat{\delta}_{\text{rot}_1})$
- 10: $\theta' \leftarrow \theta + \hat{\delta}_{\text{rot}_1} + \hat{\delta}_{\text{rot}_2}$
- 11: **return** $(x', y', \theta')^T$




Figure 5.9 Sampling from the odometry motion model, using the same parameters as in Figure 5.8. Each diagram shows

Determining the Noise Parameters

- ▶ Both the velocity motion model and the odometry motion model represent **noise by a set of parameters $\alpha_{1:6}$ and $\alpha_{1:4}$, respectively**
- ▶ **The noise parameters are robot-specific** — the noise induced by a given motion varies between robots, as it depends on the concrete kinematic structure
- ▶ These parameters need to be **determined experimentally**, using a procedure as follows:
 1. **Collect a dataset of robot motions**
 2. **Calculate the probability $p(\mathbf{x}_{t+1}|\mathbf{u}_{t+1}, \mathbf{x}_t)$ under the chosen motion model** — algorithms for calculating the probability are given in the reference book
 3. **Perform parameter estimation**, for instance using maximum likelihood estimation (MLE)

Summary

- ▶ A particle filter approximates an unknown state distribution by a discrete set of particles
 - ▶ Each particle represents a weighted hypothesis about the state, where the weight is determined by the measurement model
- ▶ Particle filters update the approximation of the state distribution by performing importance sampling
 - ▶ At each iteration, particles are sampled with replacement; the probability that a particle is sampled is proportional to its weight
- ▶ Unlike a Kalman filter, a particle filter can, in principle, be used to recover from the kidnapped robot problem
 - ▶ For this, randomly distributed particles can be added during the particle resampling process
- ▶ Probabilistic motion models, such as the velocity motion model and the odometry motion model, can be used for recursive state estimation using particle filters