

## Recommended reading

- Lesson 3 in <https://www.udacity.com/course/artificial-intelligence-for-robotics--cs373>
- Chapters 4.3 and 8.3 in the Probabilistic Robotics textbook

COMP417  
Introduction to Robotics and Intelligent Systems  
Particle Filters (~2 lectures)



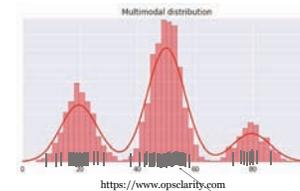
### KF vs EKF vs PF

	Kalman Filter	Extended Kalman Filter	Particle Filter
Dynamics model	Linear	Nonlinear	Nonlinear
Sensor model	Linear	Nonlinear	Nonlinear
Noise	Gaussian (Unimodal)	Gaussian (Unimodal)	Multimodal

One peak

Multiple peaks

### How can we represent multimodal distributions?



Higher density of particles means higher probability mass

$$\begin{aligned}
 \text{bel}(x_t) &= p(x_t | z_{0:t}, u_{0:t-1}) \\
 &= \sum_{m=1}^M \begin{cases} w_t^{[m]} / W & \text{if } x_t = x_t^{[m]} \\ 0 & \text{o.w.} \end{cases}
 \end{aligned}$$

$W$  = sum of all particles' weights

#### Idea #1: Histograms

Advantages: the higher the number of bars the better the approximation is

Disadvantages: exponential dependence on number of dimensions

Note: this approach is called the Histogram Filter. It is useful for low-dimensional systems.

#### Idea #2: Function Approximation

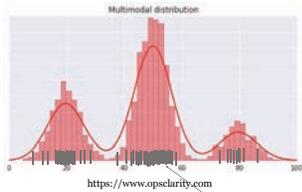
Unclear how to do Bayes' filter updates and predictions in this case, but a mixture of Gaussians is an option.

#### Idea #3: Weighted Particles $\{(x_t^{[1]}, w_t^{[1]}), \dots, (x_t^{[M]}, w_t^{[M]})\}$

Advantages: easy to predict/update by treating each particle as a separate hypothesis whose weight is updated.

Disadvantages: need enough particles to "cover" the distribution

# How can we represent multimodal distributions?



<https://www.opsclarity.com>

Higher density of particles means higher probability mass

## Idea #1: Histograms

Advantages: the higher the number of bars the better the approximation is

Disadvantages: exponential dependence on number of dimensions

Note: this approach is called the Histogram Filter. It is useful for low-dimensional systems but we will not study it in this class.

## Idea #2: Function Approximation

Unclear how to do Bayes' filter updates and predictions in this case.

## Idea #3: Weighted Particles $\{(x^{[1]}, w^{[1]}), \dots, (x^{[M]}, w^{[M]})\}$

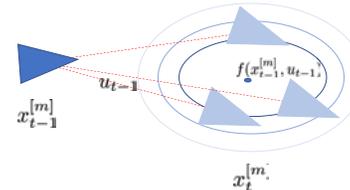
Advantages: easy to predict/update by treating each particle as a separate hypothesis whose weight is updated.

Disadvantages: need enough particles to "cover" the distribution

Want particles to be drawn from the belief at time t:

$$x_t^{[m]} \sim p(x_t | z_{0:t}, u_{0:t-1})$$

# Particle propagation/prediction



Simulate what is going to happen to the particle at the next time step by drawing a sample from the next state specified in the dynamics (a.k.a. one-step simulator)

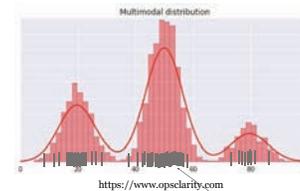
$$x_t^{[m]} \sim p(x_t | x_{t-1}^{[m]}, u_{t-1})$$

Usually

$$x_t^{[m]} = f(x_{t-1}^{[m]}, u_{t-1}) + w_{t-1}$$

$$w_{t-1} \sim \mathcal{N}(0, Q)$$

# How can we represent multimodal distributions?



<https://www.opsclarity.com>

Higher density of particles means higher probability mass

## Idea #1: Histograms

## Idea #3: Weighted Particles $\{(x^{[1]}, w^{[1]}), \dots, (x^{[M]}, w^{[M]})\}$

Advantages: easy to predict/update by treating each particle as a separate hypothesis whose weight is updated.

Disadvantages: need enough particles to "cover" the distribution

# Particle Update

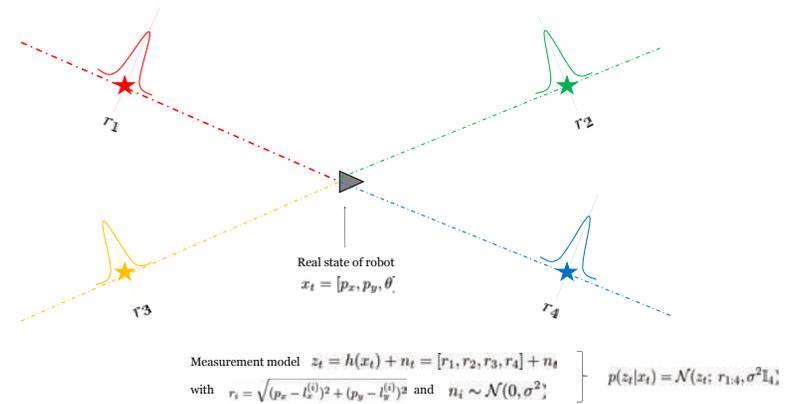
Want particles to be drawn from the belief at time t:

$$x_t^{[m]} \sim p(x_t | z_{0:t}, u_{0:t-1})$$

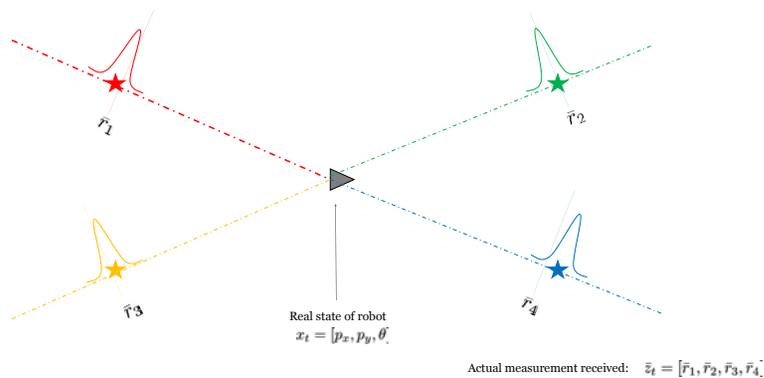
## Using particles to represent a distribution

- Each particle is a hypothesis
- Each particle moves forward using your dynamics model (motion model)
- Each particle is "confirmed" by comparing the actual observation to what the particle pose would predict
- These predictions are used to either:
  - change the particle density (unweighted particles)
  - change the particle weight

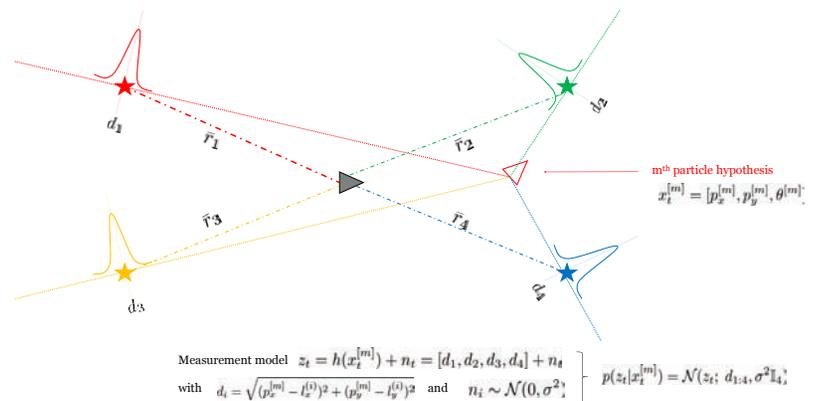
## How to update particle weights after an observation



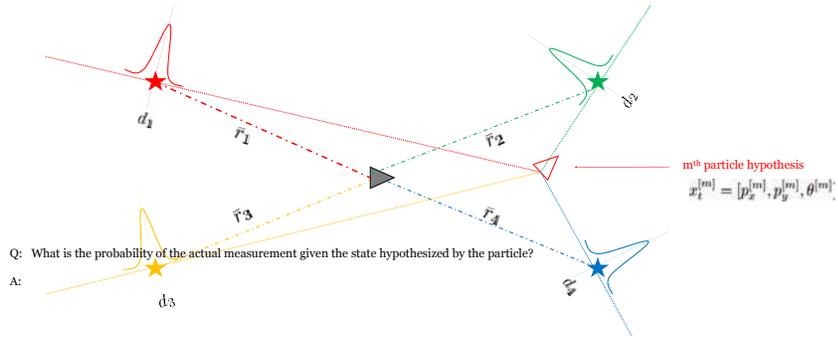
## How to update particle weights after an observation



## How to update particle weights after an observation

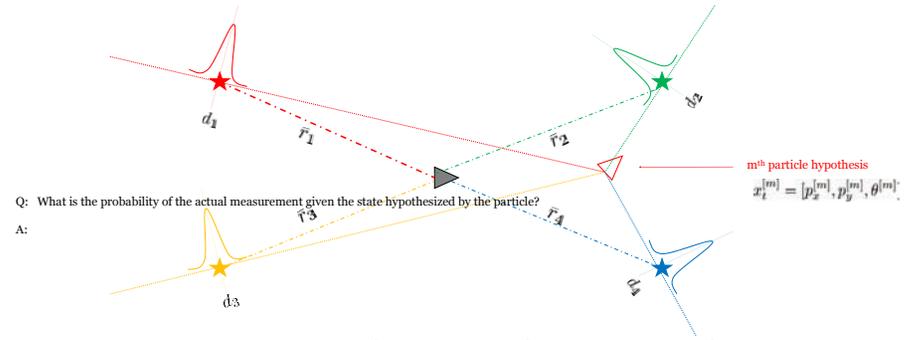


## How to update particle weights after an observation



$$p(\bar{z}_t | x_t^{[m]}) = \mathcal{N}(\bar{z}_t; d_{1:4}, \sigma^2 \mathbb{I}_4) = \eta \exp(-\|\bar{z}_t - d_{1:4}\|^2 / \sigma^2)$$

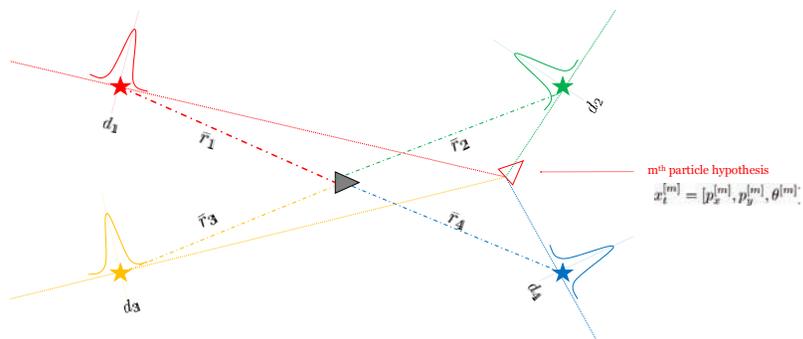
## How to update particle weights after an observation



Assuming range measurements are conditionally independent given state

$$p(\bar{z}_t | x_t^{[m]}) = \prod_{i=1}^4 p(\bar{r}_i | x_t^{[m]}) = \prod_{i=1}^4 \mathcal{N}(\bar{r}_i; d_i, \sigma^2) = \prod_{i=1}^4 \eta \exp(-(\bar{r}_i - d_i) / \sigma^2)$$

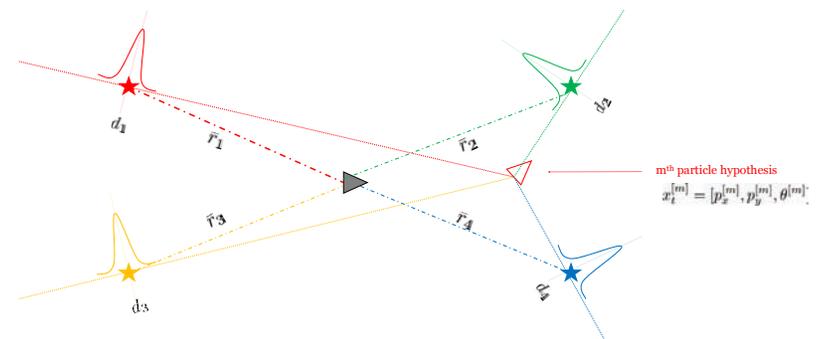
## How to update particle weights after an observation



$$p(\bar{z}_t | x_t^{[m]}) = \prod_{i=1}^4 p(\bar{r}_i | x_t^{[m]}) = \prod_{i=1}^4 \mathcal{N}(\bar{r}_i; d_i, \sigma^2)$$

In the figure above this probability would be low and this particle would be unlikely.

## How to update particle weights after an observation



Particle's (unnormalized) weight  $w_t^{[m]} \propto p(\bar{z}_t | x_t^{[m]})$

See Appendix 1 for why this choice was made for the weight

The distribution of the particles has not been updated yet. We only updated their weights. To update the distribution of particles we need to do resampling

The distribution of the particles has not been updated yet. We only updated their weights. To update the distribution of particles we need to do resampling

Sample particles with repetition/replacement, according to their updated weights.

## Resampling Particles

- Main goal: Get rid of unlikely particles (with too low weights) and focus on most likely particles (a.k.a. survival of the fittest).
- Main mechanism: Sample new set of particles from existing set, with replacement (repetition), so that same particle can be sampled more than once.  
Sample old particle  $i$  with probability  $\propto \text{weight}_i$
- Many possible ways to implement it. Here we present two algorithms.

## Four classes of resampling strategy

- **Multinomial**: imagine a strip of paper where each particle has a section, where the length is proportional to its weight. Randomly pick a location on the strip  $N$  times, and pick the particle associated with the section.
- **Residual**: object is to reduce the variance of the sample set. Do this by first allocating each particle their integer floor of the expected value, and leave the rest to multinomial resampling. E.g. a particle with an expected value of 2.5 will have 2 copies in the resampled set and another one with an expected value of 0.5.
- **Systematic**: take a ruler with regular spaced marks, such that  $N$  marks are the same length as your strip of paper. Randomly place the ruler next to your strip. Take the particles at the marks.
- **Stratified**: same as systematic, except that the marks on the ruler are not evenly placed, but added as  $N$  random processes sampling from the interval  $0..1/N$ .

## Resampling Particles: Algorithm #1

```

new_particles = []
sample u ~ Uniform[0,1]
idx = int( u * (N-1) )
beta = 0
max_w = max(weights)

for each of the N particles:
    sample v ~ Uniform[0,1]
    beta += v * 2 * max_w

    while beta > weights[idx]:
        beta -= weights[idx]
        idx = (idx + 1) % N

    p = particles[idx].copy()
    new_particles.append(p)
    
```

## Resampling Particles: Algorithm #2

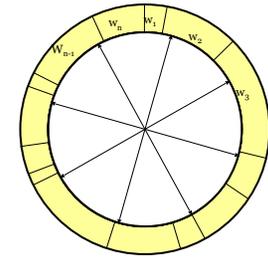
```

new_particles = []
sample r ~ Uniform[0, 1/N]
c = weights[0]
idx = 0

for n = 1..N:
    u = r + (n-1)/N

    while u > c:
        idx = idx + 1
        c = c + weights[idx]
        c computes cumulative distribution

    p = particles[idx].copy()
    new_particles.append(p)
    
```



- Stochastic universal sampling
- Systematic resampling
- Linear time complexity

## Resampling Particles: Example

- Suppose we only have 5 particles:

Particle index	Normalized weight
1	0.1
2	0.2
3	0.4
4	0.1
5	0.2

Q: What is the probability that after a round of resampling the highest probability particle (#3) is not sampled?

A:  $0.6^5 \approx 0.077$

i.e. there is nonzero probability that we will lose the highest-probability particle → it will happen eventually

## Resampling Particles: Example

- Suppose we only have 5 particles:

Particle index	Normalized weight
1	0.1
2	0.2
3	0.4
4	0.1
5	0.2

Q: What is the probability that after a round of resampling the highest-probability particle (#3) is not sampled?

A:  $0.6^5 \approx 0.077$

Q: What is the probability that after a round of resampling one of the lowest-probability particles (#1) is not sampled?

A:  $0.9^5 \approx 0.59$

## Resampling Particles: Consequences

- Weak particles very likely do not survive.
- ↓
- Variance among the set of particles **decreases**, due to mostly sampling strong particles (i.e. loss of particle diversity).
- ↓
- Loss of particle diversity implies **increased variance** of the approximation error between the particles and the true distribution.
- ↓
- Particle deprivation: there are no particles in the vicinity of the correct state

## How to address particle deprivation

- Idea #1: don't resample when only a few particles contribute
- Idea #2: inject random particles during resampling
- Idea #3: increase the number of particles (may be impractical depending on the computational complexity of the system)

## How to address particle deprivation

- Idea #1: don't resample when only a few particles contribute
  - Effective sample size:  $N_{\text{eff}} = \frac{1}{\sum_{i=1}^N w_i^2}$
  - When all particles have equal, normalized weights (1/N) then  $N_{\text{eff}} = N$
  - When a single particle carries the entire weight then and we have loss of particle diversity  $N_{\text{eff}} = 1/N$
  - Resample only when  $N_{\text{eff}} > N_{\text{thresh}}$
- Idea #2: inject random particles during resampling
- Idea #3: increase the number of particles (may be impractical depending on the computational complexity of the system)

## How to address particle deprivation

- Idea #1: don't resample when only a few particles contribute
- Idea #2: inject random particles during resampling
  - A small percentage of the particles' states should be set randomly
    - Pro: simple to code, reduces (but does not fix) particle deprivation
    - Con: incorrect posterior estimation even when there are infinitely many particles
- Idea #3: increase the number of particles (may be impractical depending on the computational complexity of the system)

# Particle Filter Algorithm

```

ParticleFilter( $\bar{z}_t, u_{t-1}$ )
     $\bar{S}_t = \{\}$     $\bar{W}_t = \{\}$ 
    for particle index  $m = 1 \dots M$ 
        sample  $x_t^{[m]} \sim p(x_t | x_{t-1}^{[m]}, u_{t-1})$ 
         $w_t^{[m]} = p(\bar{z}_t | x_t^{[m]})$ 
         $\bar{S}_t.append(x_t^{[m]})$ 
         $\bar{W}_t.append(w_t^{[m]})$ 

     $S_t = \{\}$ 
    for particle index  $m = 1 \dots M$ 
        sample particle i from  $\bar{S}_t$  with probability  $\propto w_t^{[i]}$ 
         $S_t.append(x_t^{[m]})$ 

    return  $S_t$ 

```

Actual observation and control received

# Particle Filter Algorithm

```

ParticleFilter( $\bar{z}_t, u_{t-1}$ )
     $\bar{S}_t = \{\}$     $\bar{W}_t = \{\}$ 
    for particle index  $m = 1 \dots M$ 
        sample  $x_t^{[m]} \sim p(x_t | x_{t-1}^{[m]}, u_{t-1})$ 
         $w_t^{[m]} = p(\bar{z}_t | x_t^{[m]})$ 
         $\bar{S}_t.append(x_t^{[m]})$ 
         $\bar{W}_t.append(w_t^{[m]})$ 

     $S_t = \{\}$ 
    for particle index  $m = 1 \dots M$ 
        sample particle i from  $\bar{S}_t$  with probability  $\propto w_t^{[i]}$ 
         $S_t.append(x_t^{[m]})$ 

    return  $S_t$ 

```

Particle propagation/prediction:  
noise needs to be added in order to make  
particles differentiate from each other.

If propagation is deterministic then particles  
are going to collapse to a single particle after a  
few resampling steps.

# Particle Filter Algorithm

```

ParticleFilter( $\bar{z}_t, u_{t-1}$ )
     $\bar{S}_t = \{\}$     $\bar{W}_t = \{\}$ 
    for particle index  $m = 1 \dots M$ 
        sample  $x_t^{[m]} \sim p(x_t | x_{t-1}^{[m]}, u_{t-1})$ 
         $w_t^{[m]} = p(\bar{z}_t | x_t^{[m]})$ 
         $\bar{S}_t.append(x_t^{[m]})$ 
         $\bar{W}_t.append(w_t^{[m]})$ 

     $S_t = \{\}$ 
    for particle index  $m = 1 \dots M$ 
        sample particle i from  $\bar{S}_t$  with probability  $\propto w_t^{[i]}$ 
         $S_t.append(x_t^{[m]})$ 

    return  $S_t$ 

```

Weight computation as measurement likelihood.  
For each particle we compute the probability of the  
actual observation given the state is at that particle.

# Particle Filter Algorithm

```

ParticleFilter( $\bar{z}_t, u_{t-1}$ )
     $\bar{S}_t = \{\}$     $\bar{W}_t = \{\}$ 
    for particle index  $m = 1 \dots M$ 
        sample  $x_t^{[m]} \sim p(x_t | x_{t-1}^{[m]}, u_{t-1})$ 
         $w_t^{[m]} = p(\bar{z}_t | x_t^{[m]})$ 
         $\bar{S}_t.append(x_t^{[m]})$ 
         $\bar{W}_t.append(w_t^{[m]})$ 

     $S_t = \{\}$ 
    for particle index  $m = 1 \dots M$ 
        sample particle i from  $\bar{S}_t$  with probability  $\propto w_t^{[i]}$ 
         $S_t.append(x_t^{[m]})$ 

    return  $S_t$ 

```

Resampling step

Note: particle deprivation heuristics are not  
shown here

# Particle Filter Algorithm

```

ParticleFilter( $\bar{z}_t, u_{t-1}$ )
 $\bar{S}_t = \{\}$   $\bar{W}_t = \{\}$ 
for particle index  $m = 1 \dots M$ 
  sample  $x_t^{[m]} \sim p(x_t | x_{t-1}^{[m]}, u_{t-1})$ 
   $w_t^{[m]} = p(\bar{z}_t | x_t^{[m]})$ 
   $\bar{S}_t.append(x_t^{[m]})$ 
   $\bar{W}_t.append(w_t^{[m]})$ 

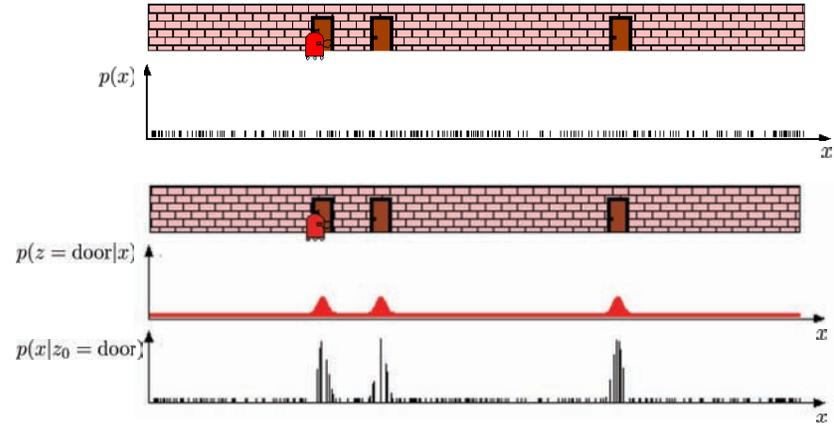
 $S_t = \{\}$ 
for particle index  $m = 1 \dots M$ 
  sample particle  $i$  from  $\bar{S}_t$  with probability  $\propto w_t^{[m]}$ 
   $S_t.append(x_t^{[m]})$ 

return  $S_t$ 
  
```

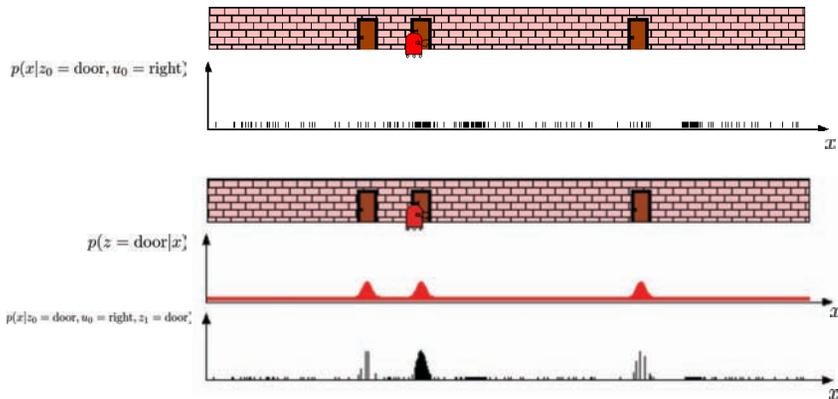
Note: here we work with a fixed number of particles but in many applications, such as localization, you could work with a reduced number of particles after the particles have converged to the true estimate.

Such implementations of particle filters are called adaptive. An example is the KLD-sampling adaptive particle filter, which is not going to be covered here.

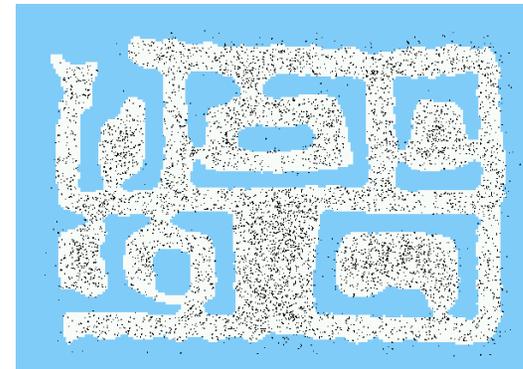
# Examples: 1D Localization



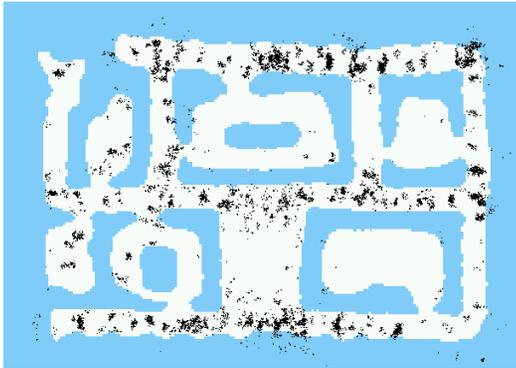
# Examples: 1D Localization



# Examples: Monte Carlo Localization



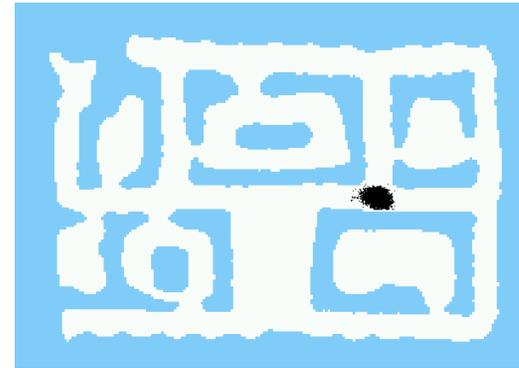
## Examples: Monte Carlo Localization



After incorporating 10 ultrasound scans

37

## Examples: Monte Carlo Localization



After incorporating 65 ultrasound scans

38

## Using Ceiling Maps for Localization

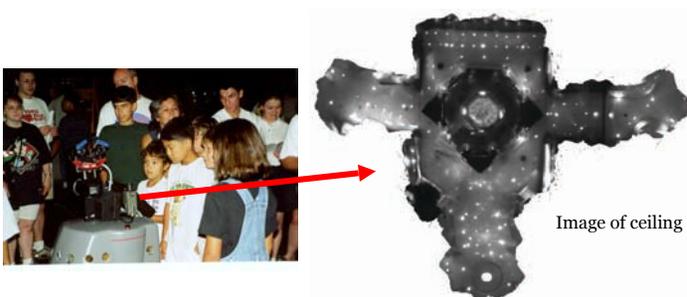
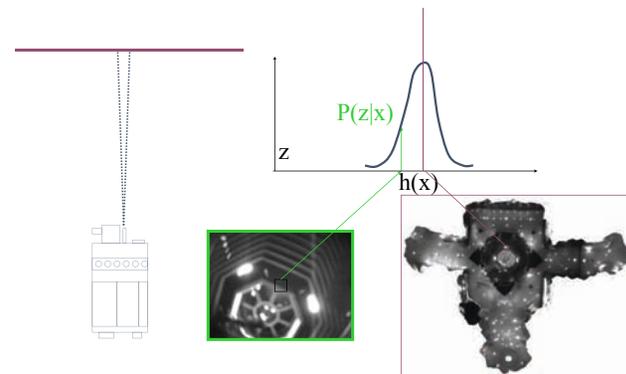


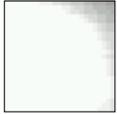
Image of ceiling

## Vision-based Localization



## Under a Light

Measurement  $z$ :



$P(z|x)$ :

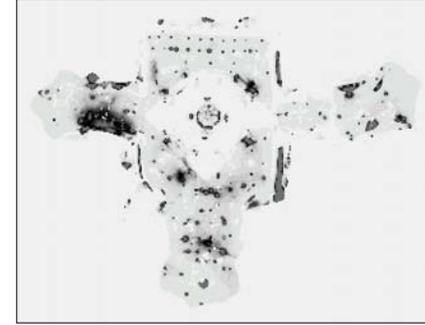


## Next to a Light

Measurement  $z$ :



$P(z|x)$ :

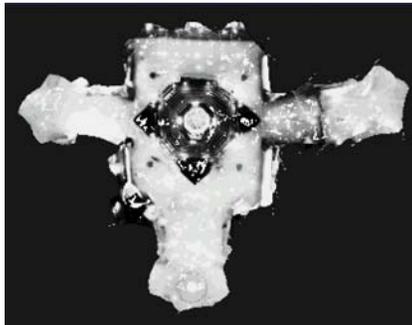


## Elsewhere

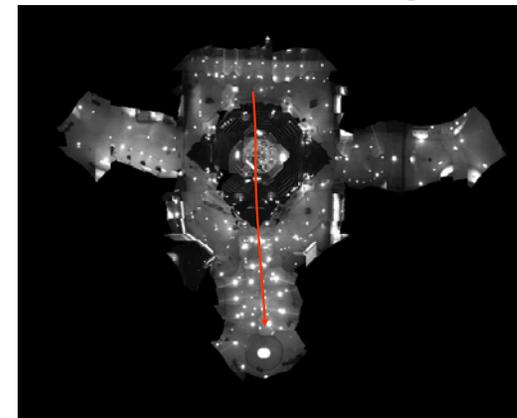
Measurement  $z$ :



$P(z|x)$ :

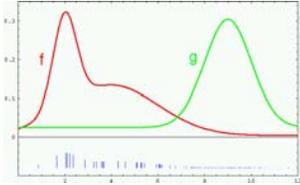


## Global Localization Using Vision



# Appendix 1

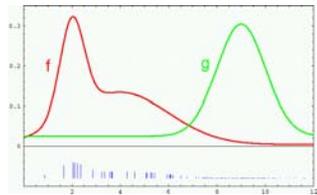
- Why did we choose  $w_t^{[m]} \propto p(z_t|x_t^{[m]})$ , as the importance weight for particle  $m$ ?
- Main trick: **importance sampling**, i.e. how to estimate properties/statistics of one distribution (f) given samples from another distribution (g)



For example, suppose we want to estimate the expected value of f given only samples from g.

# Appendix 1

- Why did we choose  $w_t^{[m]} \propto p(z_t|x_t^{[m]})$ , as the importance weight for particle  $m$ ?
- Main trick: **importance sampling**, i.e. how to estimate properties/statistics of one distribution (f) given samples from another distribution (g)



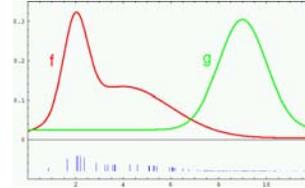
For example, suppose we want to estimate the expected value of f given only samples from g.

$$\begin{aligned} \mathbb{E}_{x \sim f(x)}[x] &= \int x f(x) dx \\ &= \int \frac{g(x)}{g(x)} x f(x) dx \\ &= \int \frac{x f(x)}{g(x)} g(x) dx \\ &= \mathbb{E}_{x \sim g(x)} \left[ x \frac{f(x)}{g(x)} \right] \\ &= \mathbb{E}_{x \sim g(x)} [x w(x)] \end{aligned}$$

Weights describe the mismatch between the two distributions, i.e. how to reweigh samples to obtain statistics of f from samples of g

# Appendix 1

- Why did we choose  $w_t^{[m]} \propto p(z_t|x_t^{[m]})$ , as the importance weight for particle  $m$ ?
- Main trick: **importance sampling**, i.e. how to estimate properties/statistics of one distribution (f) given samples from another distribution (g)

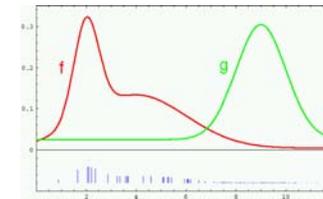


For example, suppose we want to estimate the expected value of f given only samples from g.

$$\begin{aligned} \mathbb{E}_{x \sim f(x)}[x] &= \int x f(x) dx \\ &= \int \frac{g(x)}{g(x)} x f(x) dx \\ &= \int \frac{x f(x)}{g(x)} g(x) dx \\ &= \mathbb{E}_{x \sim g(x)} \left[ x \frac{f(x)}{g(x)} \right] \\ &= \mathbb{E}_{x \sim g(x)} [x w(x)] \end{aligned}$$

# Appendix 1

- Why did we choose  $w_t^{[m]} \propto p(z_t|x_t^{[m]})$ , as the importance weight for particle  $m$ ?
- Main trick: **importance sampling**, i.e. how to estimate properties/statistics of one distribution (f) given samples from another distribution (g)



In the case of particle filters

$$f(x_t) = p(x_t|z_{0:t}, u_{0:t-1}) = \text{bel}(x_t), \quad g(x_t) = p(x_t|z_{0:t-1}, u_{0:t-1}) = \text{bel}(x_t)$$

Posterior belief after update

Belief after propagation, before update

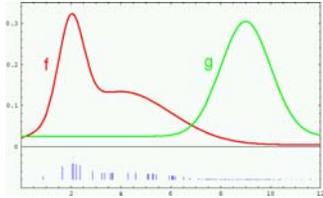
For example, suppose we want to estimate the expected value of f given only samples from g.

$$\begin{aligned} \mathbb{E}_{x \sim f(x)}[x] &= \int x f(x) dx \\ &= \int \frac{g(x)}{g(x)} x f(x) dx \\ &= \int \frac{x f(x)}{g(x)} g(x) dx \\ &= \mathbb{E}_{x \sim g(x)} \left[ x \frac{f(x)}{g(x)} \right] \\ &= \mathbb{E}_{x \sim g(x)} [x w(x)] \end{aligned}$$

Weights describe the mismatch between the two distributions, i.e. how to reweigh samples to obtain statistics of f from samples of g

# Appendix 1

- Why did we choose  $w_t^{[m]} \propto p(z_t|x_t^{[m]})$ , as the importance weight for particle m?
- Main trick: **importance sampling**, i.e. how to estimate properties/statistics of one distribution (f) given samples from another distribution (g)



$$\begin{aligned}
 w(x_t^{[m]}) &= \frac{f(x_t^{[m]})}{g(x_t^{[m]})} \\
 &\propto \frac{p(z_t|x_t^{[m]}) p(x_t^{[m]}|x_{t-1}^{[m]}, u_{t-1}) \text{bel}(x_{t-1}^{[m]})}{p(x_t^{[m]}|x_{t-1}^{[m]}, u_{t-1}) \text{bel}(x_{t-1}^{[m]})} \\
 &\propto p(z_t|x_t^{[m]})
 \end{aligned}$$

In the case of particle filters

$$f(x_t) = p(x_t|z_{0:t}, u_{0:t-1}) = \text{bel}(x_t), \quad g(x_t) = p(x_t|z_{0:t-1}, u_{0:t-1}) = \overline{\text{bel}}(x_t)$$

Posterior belief after update

Belief after propagation, before update