CS 369: Introduction to Robotics

Prof. Thao Nguyen

Spring 2026

HAVERFORD
COLLEGE

Admin

- Lab 3 grades posted on Moodle
- Lab 4 due tonight
- Lab 5 posted (due next Wednesday)

Outline for today

- Backpropagation
- Localization

Outline for today

- Backpropagation
- Localization

Backpropagation

https://cs231n.github.io/optimization-2/

Outline for today

- Backpropagation
- Localization

Localization

- Problem of determining the robot's pose relative to the environment map
- Local vs. global localization:
  - Position tracking: assumes that the initial robot pose is known
  - Global localization: initial robot pose is unknown
  - Kidnapped robot problem: during operation, the robot can get kidnapped and teleported to some other location
- Static vs. dynamic environments

Thrun, Burgard, and Fox,
Probabilistic Robotics

Example

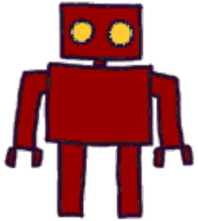

- Moving only in one dimension
- Known map of flower garden
- Simple flower detector
  - Beeps when you are in front of a flower
  - Gaussian distribution of a flower given a beep

Example

Initially, no idea where we are

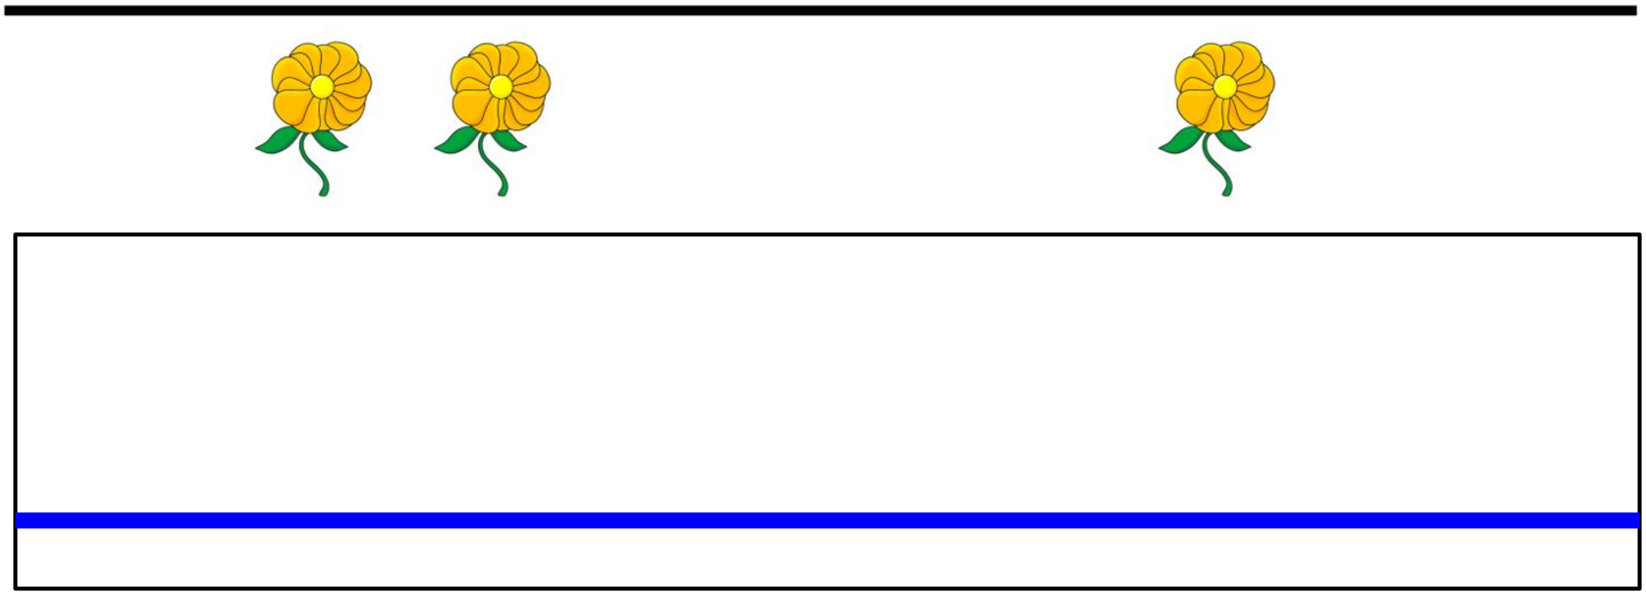

Example

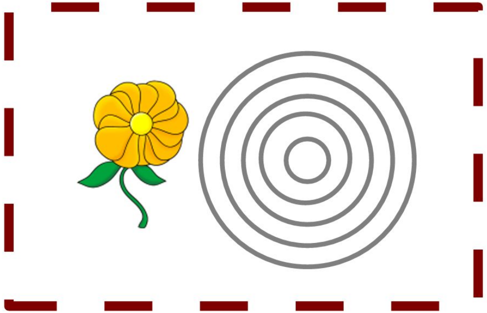

First observation -> update belief about location

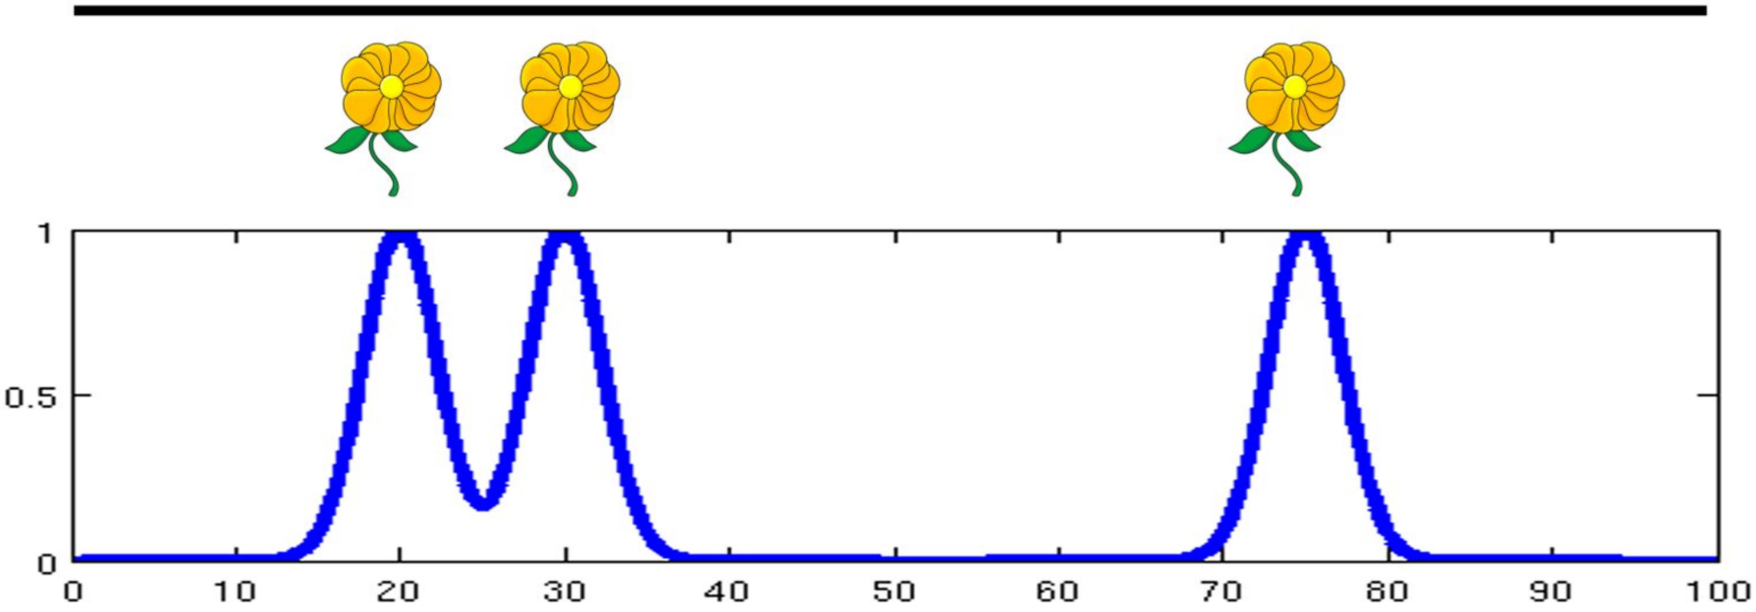

Example

Robot moves, motion update

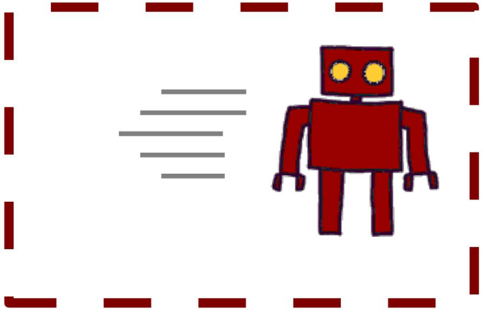

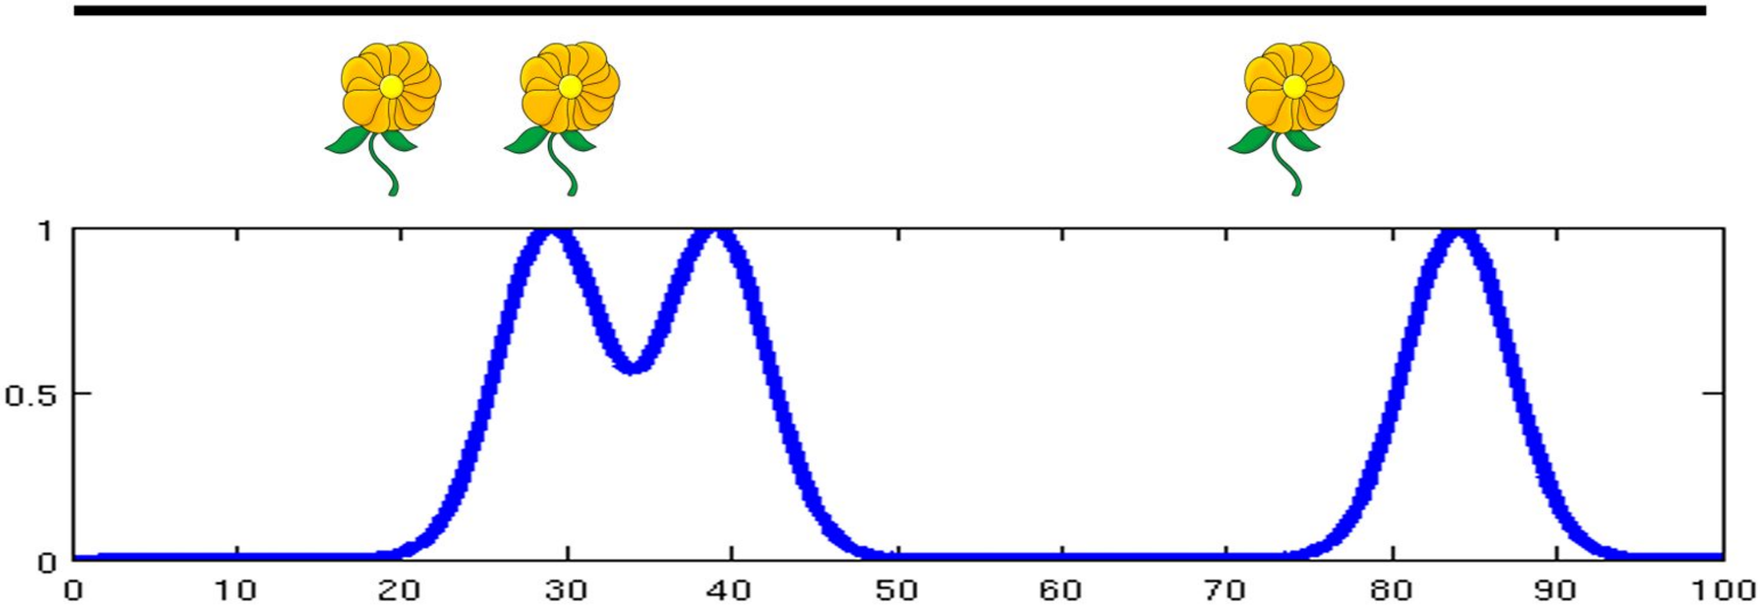

Example

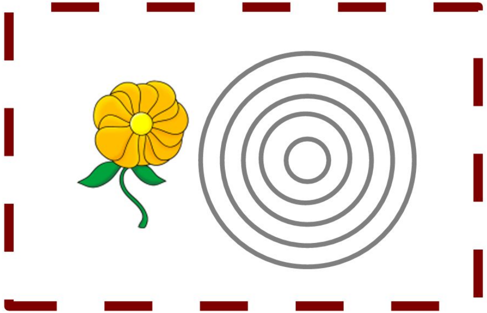

Observation update

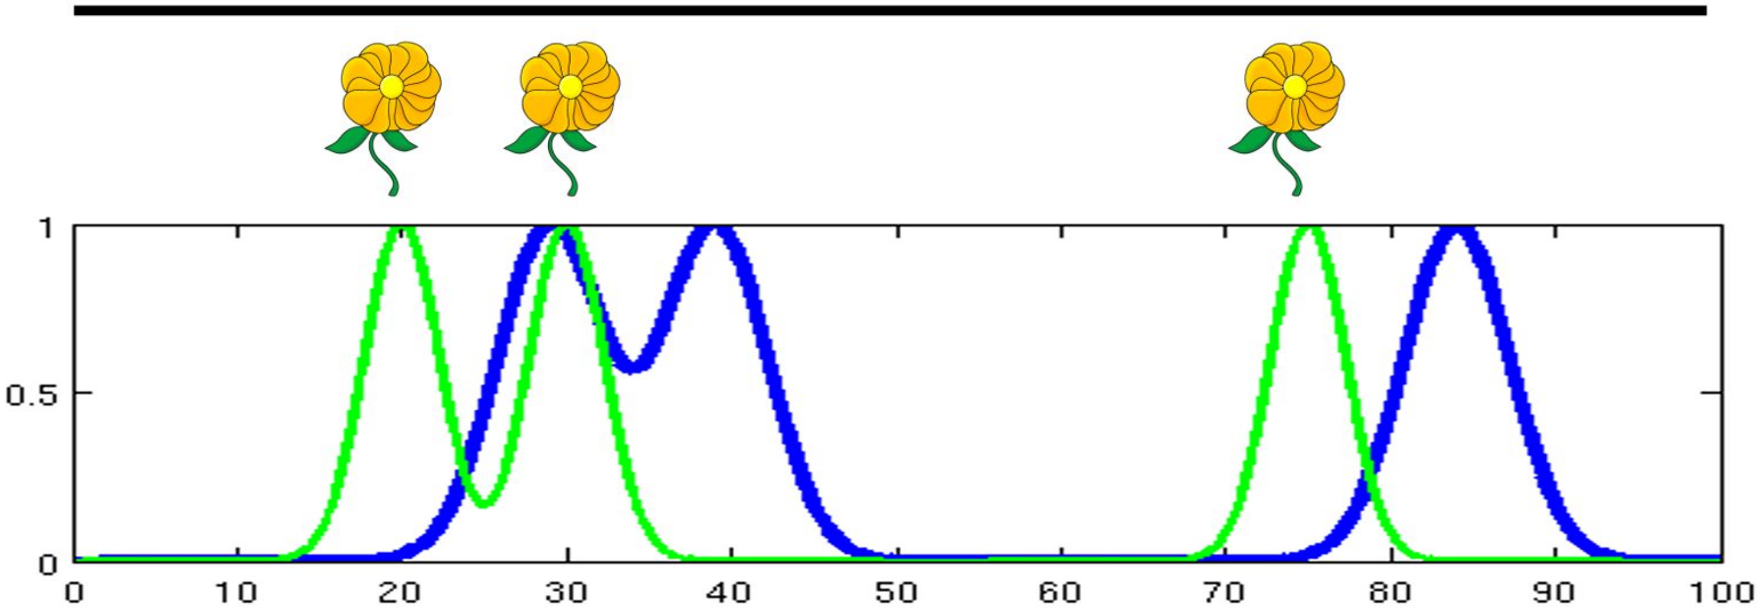

Example

Final belief

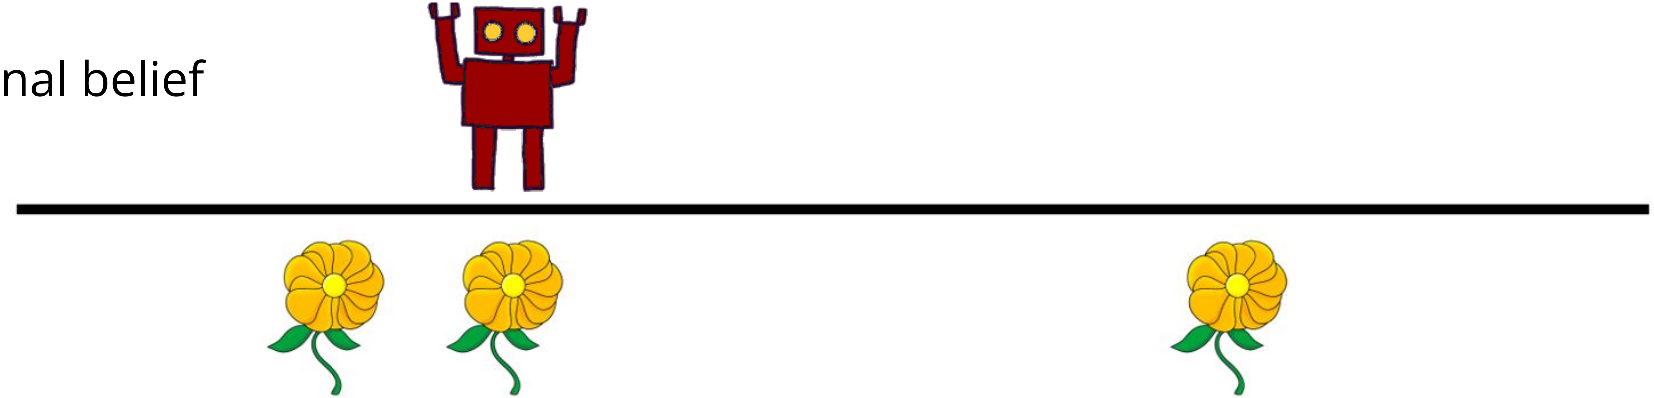

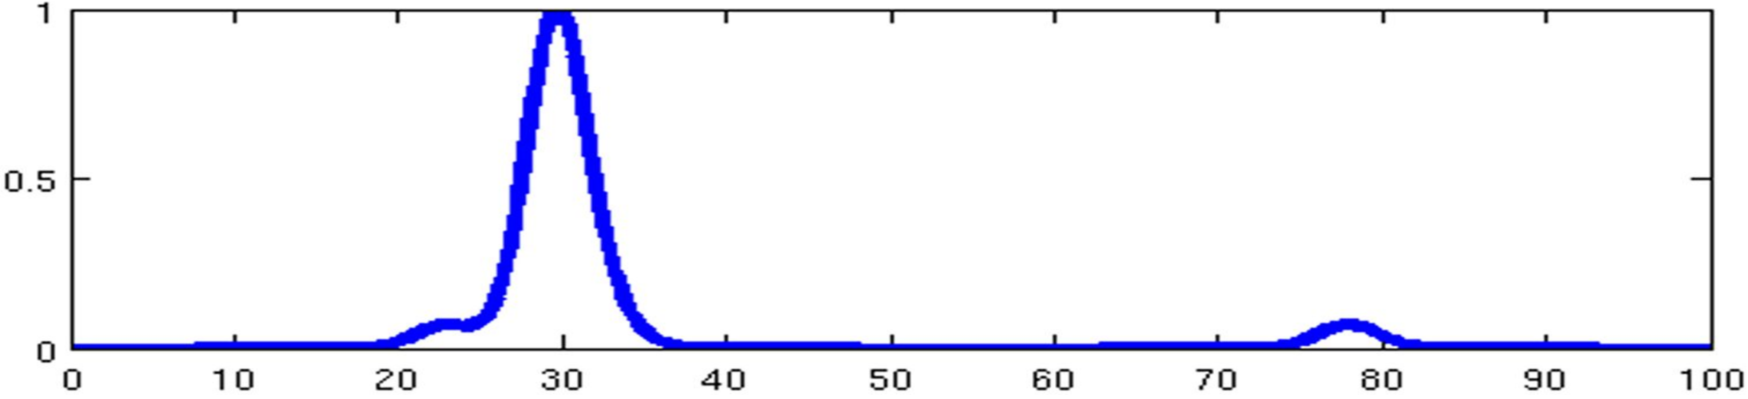

Bayes' rule

- $P(A,B) = P(A \mid B) P(B)$
- $P(A,B) = P(B \mid A) P(A)$

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

Bayes' rule with background knowledge

$$P(x \mid y, z) = \frac{P(y \mid x, z)\ P(x \mid z)}{P(y \mid z)}$$

Localization

$$\underbrace{Bel(x_t)}_{\text{belief}} = P(\underbrace{x_t}_{\text{state}} \mid u_1, \underbrace{z_1}_{\text{observation}} \ldots, \underbrace{u_t}_{\text{action}}, z_t)$$

$$= \frac{P(z_t \mid x_t, u_1, z_1, \ldots, u_t)\, P(x_t \mid u_1, z_1, \ldots, u_t)}{P(z_t \mid u_1, z_1, \ldots, u_t)}$$

$$= \eta\, P(z_t \mid x_t, u_1, z_1, \ldots, u_t)\, P(x_t \mid u_1, z_1, \ldots, u_t)$$

Modeling actions

P(x|u,x') for *u* = "close door":

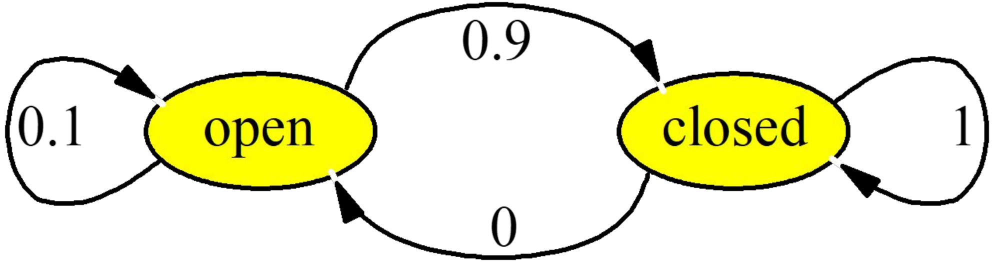

- Actions affect the state
- Actions are never carried out with absolute certainty
- Action model: **P(x|u,x')**
- Integrating the outcome of actions:

$$P(x \mid u) = \int P(x \mid u, x') P(x') dx'$$

Markov assumption

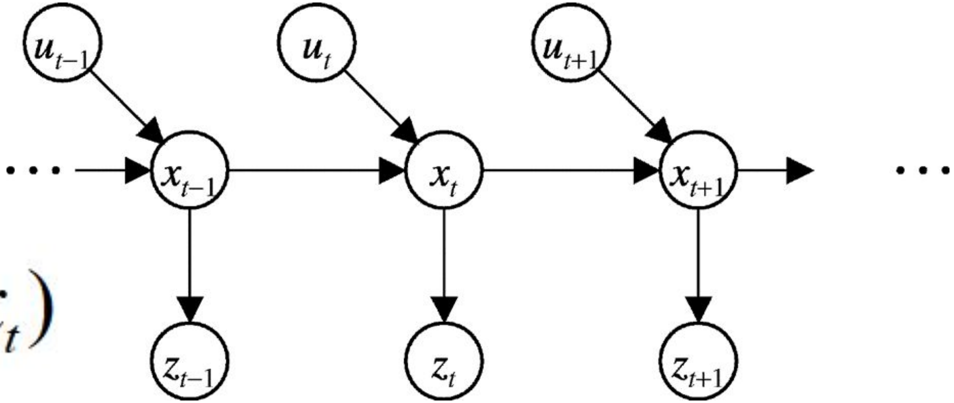

$$p(z_t \mid x_{0:t}, z_{1:t}, u_{1:t}) = p(z_t \mid x_t)$$

observation/sensor model

$$p(x_t \mid x_{1:t-1}, z_{1:t}, u_{1:t}) = p(x_t \mid x_{t-1}, u_t)$$

Underlying assumptions:

- Static environment
- Independent noise
- Perfect model, no approximation errors

Localization

$$Bel(x_t) = \eta \; P(z_t \mid x_t, u_1, z_1, \ldots, u_t) \; P(x_t \mid u_1, z_1, \ldots, u_t)$$

$$= \eta \; P(z_t \mid x_t) \int P(x_t \mid u_1, z_1, \ldots, u_t, x_{t-1})$$

$$P(x_{t-1} \mid u_1, z_1, \ldots, u_t) \; dx_{t-1}$$

$$= \eta P(z_t \mid x_t) \int P(x_t \mid u_t, x_{t-1}) \, P(x_{t-1} \mid u_1, z_1, \ldots, z_{t-1}) \; dx_{t-1}$$

$$= \eta \; P(z_t \mid x_t) \int P(x_t \mid u_t, x_{t-1}) \; Bel(x_{t-1}) \; dx_{t-1}$$

Bayes filter

Algorithm Bayes_filter($bel(x_{t-1}), u_t, z_t$**):**

  for all x_t do

    $\overline{bel}(x_t) = \int p(x_t \mid u_t, x_{t-1})\, bel(x_{t-1})\, dx$

    $bel(x_t) = \eta\, p(z_t \mid x_t)\, \overline{bel}(x_t)$

  endfor

  return $bel(x_t)$

Markov localization

Algorithm Markov_localization($bel(x_{t-1}), u_t, z_t, m$):

  for all x_t do

    $\overline{bel}(x_t) = \int p(x_t \mid u_t, x_{t-1}, m)\, bel(x_{t-1})\, dx$

    $bel(x_t) = \eta\, p(z_t \mid x_t, m)\, \overline{bel}(x_t)$

  endfor

  return $bel(x_t)$