COL864: Special Topics in AI

Semester II, 2020-21

Sate Estimation - I

Rohan Paul

Today's lecture

- Last Class
  - Planning Motions
- This Class
  - State Estimation
    - Recursive State Estimation
    - Bayes Filter
  - References
    - Probabilistic Robotics Ch 1 & 2
    - AIMA Ch 15 (till sec 15.3)

Acknowledgements

These slides are intended for teaching purposes only. Some material has been used/adapted from web sources and from slides by Nicholas Roy, Wolfram Burgard, Dieter Fox, Sebastian Thrun, Siddharth Srinivasa, Dan Klein, Pieter Abbeel and others.

Robot Environment Interaction

- Environment or world
  - Objects, robot, people, interactions
  - Environment possesses a true internal state
- Observations
  - The agent cannot directly access the true environment state.
  - Takes observations via its sensors which are error prone.
- Belief
  - Agent maintains a belief or an estimate with respect to the state of the environment derived from observations.
  - The belief is used for decision making
- Actions
  - Agent can influence the environment through its physical interactions (actuations, motions, language interaction etc.)
  - The effect of actions may be stochastic.
  - Taking actions affects the world state and the robot's internal belief with regard to this state.

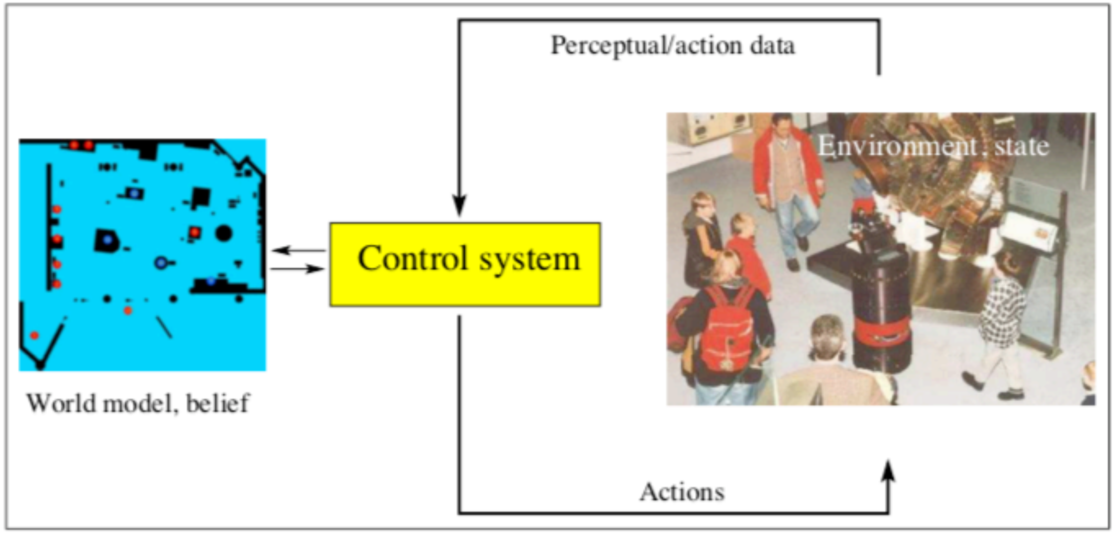

Robot Environment Interaction

- Sensing: Receive sensor data and estimate “state”
- Plan: Generation long-term plans (action sequences) based on the state & goal
- Act: Execute the actions to the robot

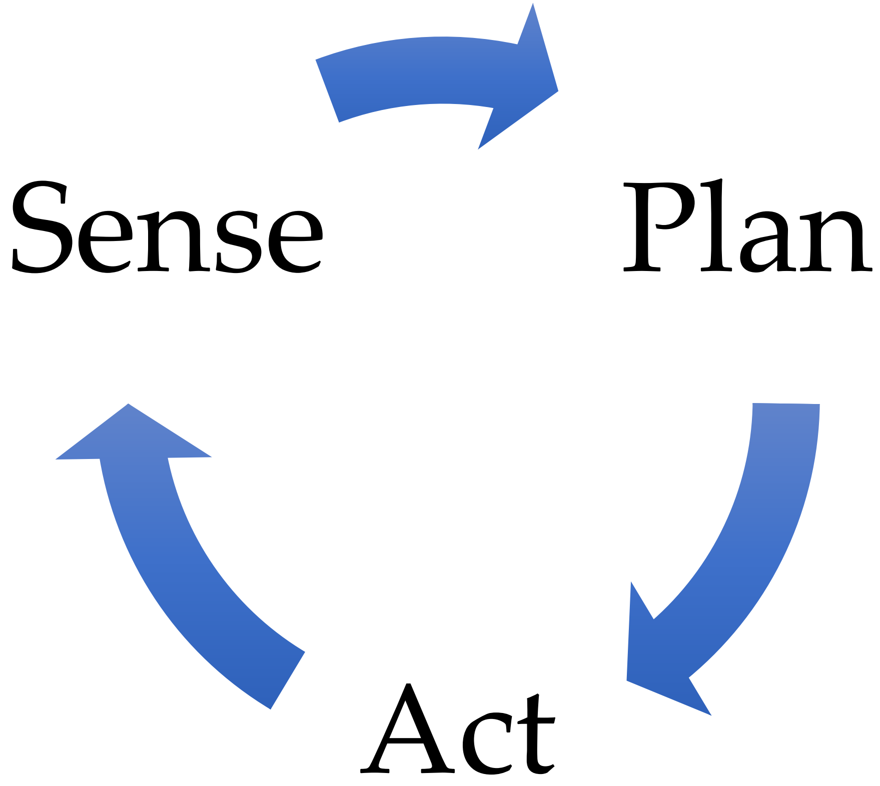

State Estimation

- Framework for estimating the state from sensor data.
- Estimating quantities that are *not directly observable. But* can be inferred if certain quantities are available to the agent.
- State estimation algorithms compute *belief distributions* over *possible states* of the world.
- Traditionally, probabilistic in nature. Can also be via function approximation.

State

- Environment is characterized by the state.
- *"A collection of all aspects of the agent and its environment that can impact the future"*
- A sufficient statistic of the past observations and interactions required for future tasks.

Figure courtesy Byron Boots

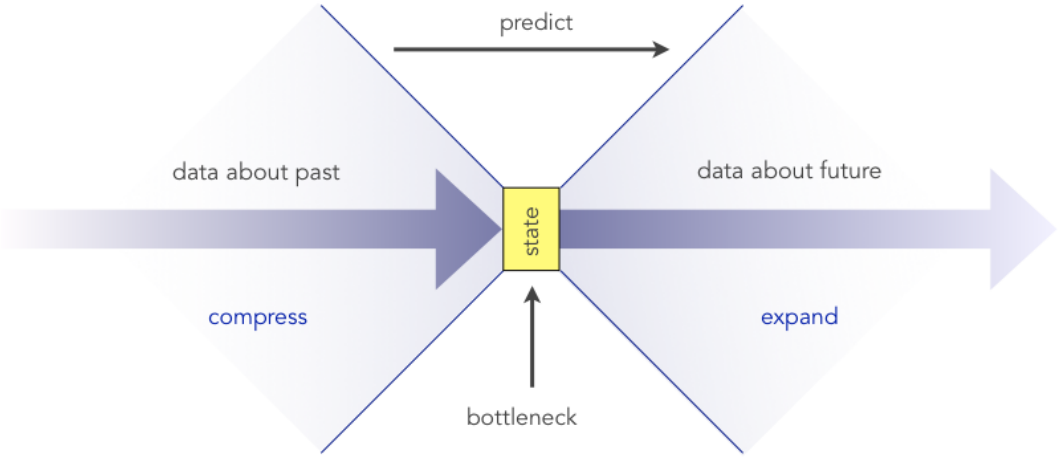

State: statistic of history sufficient to predict the future

Markovian assumption:

Future is independent of past given present

State

- What is typically part of the state, x?
  - Robot pose: position and orientation or kinematic state
  - Velocities: of the robot and other objects like people.
  - Location and features of surrounding objects in the environment.
  - Semantic states: is the door open or closed?
  - ……..
- What is put in the state is influenced by which task we seek to perform
  - Navigation
  - More complex example (e.g., delivery of hospital supplies)

Figure courtesy Byron Boots

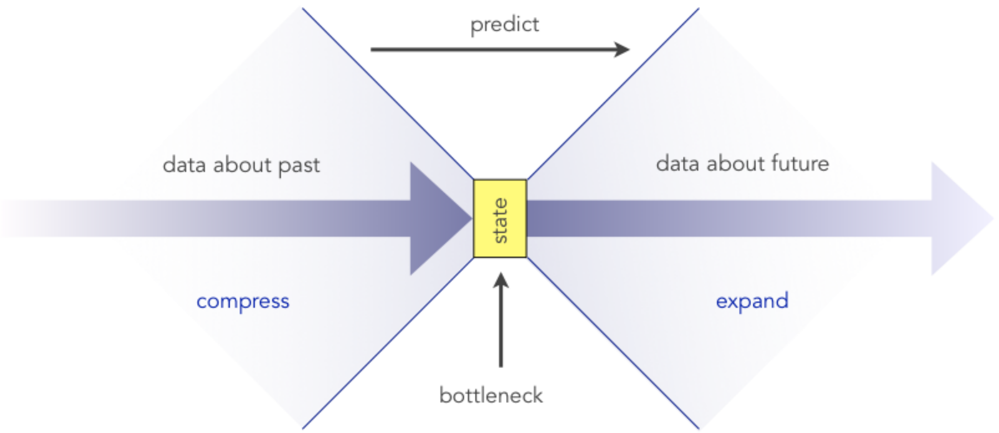

State: statistic of history sufficient to predict the future

Environment Interaction

- **Environment Sensor Measurements**
  - Camera, range, tactile, language query etc.
  - Denote measurement data as z_t
  - Noisy observations of the true state.
  - Measurements typically add information, *decrease uncertainty.*

- **Actions (or Controls)**
  - Physical interaction: robot motion, manipulation of objects, *NO_OP* etc.
  - Carry information about the change of state.
  - Source of control data: odometers or wheel encoders.
  - Denote control data as u_t
  - Actions are never carried out with absolute certainty.
  - In contrast to measurements, actions generally *increase uncertainty.*

Uncertainty

Explicitly represent uncertainty using probability theory.

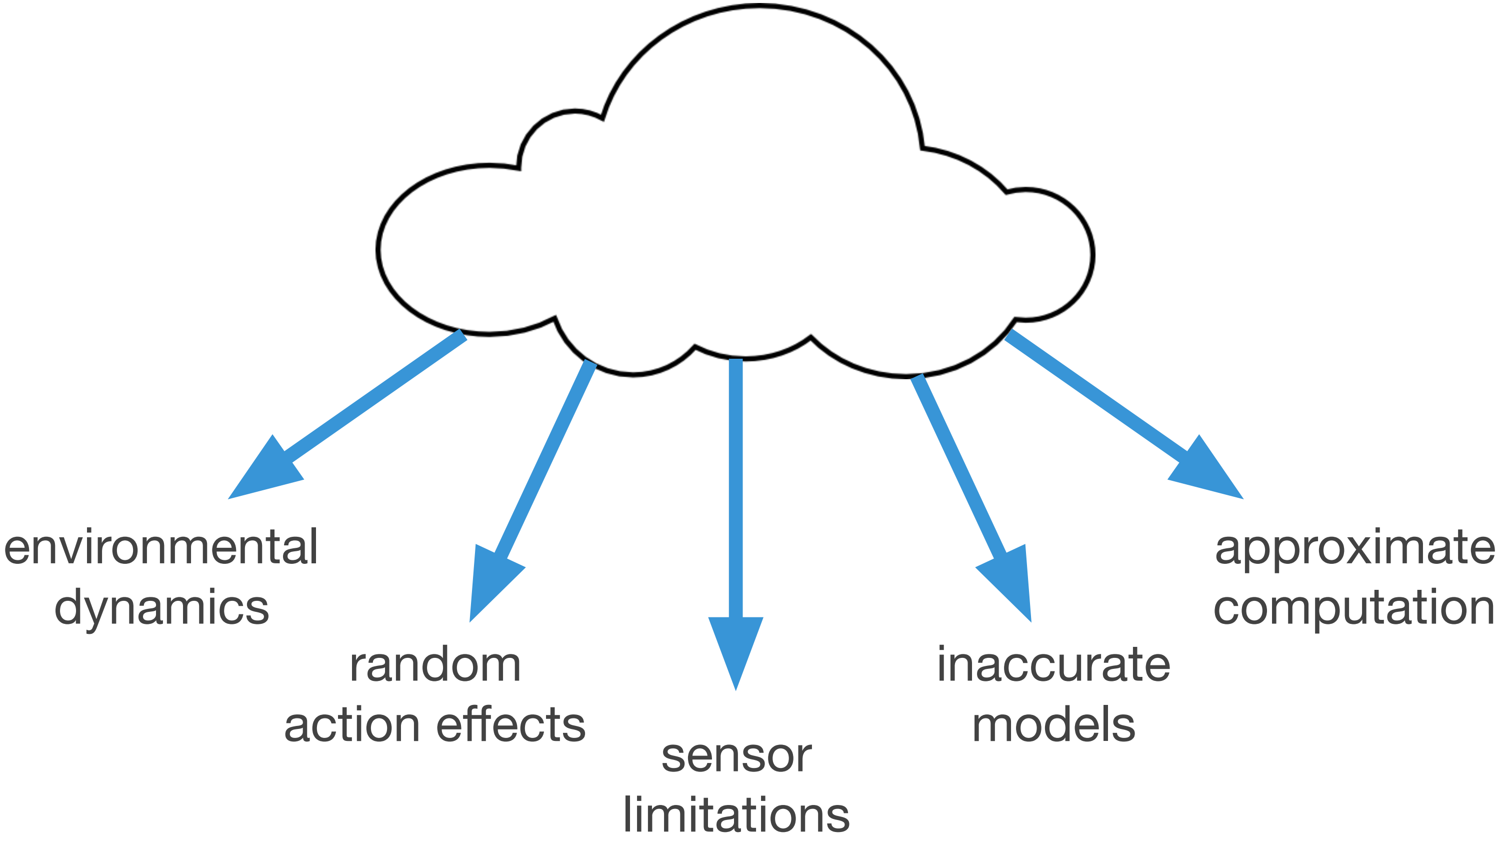

Probability

Independence

- X and Y are independent iff

  $P(x,y) = P(x)\, P(y)$

- $P(x \mid y)$ is the probability of x given y

  $P(x \mid y) = P(x,y) / P(y)$

  $P(x,y) = P(x \mid y)\, P(y)$

- If X and Y are independent then

  $P(x \mid y) = P(x)$

Marginalization

Discrete case	Continuous case
$\sum_x P(x) = 1$	$\int p(x)\, dx = 1$
$P(x) = \sum_y P(x,y)$	$p(x) = \int p(x,y)\, dy$
$P(x) = \sum_y P(x \mid y) P(y)$	$p(x) = \int p(x \mid y) p(y)\, dy$

Bayes Rule

$$P(x,y) = P(x \mid y)P(y) = P(y \mid x)P(x)$$

$$\Rightarrow$$

$$P(x \mid y) = \frac{P(y \mid x)\ P(x)}{P(y)} = \frac{\text{likelihood} \cdot \text{prior}}{\text{evidence}}$$

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

$$\eta = P(y)^{-1} = \frac{1}{\sum_x P(y \mid x)P(x)}$$

Conditioning

- Law of total probability:

$$P(x) = \int P(x, z) dz$$

$$P(x) = \int P(x \mid z) P(z) dz$$

$$P(x \mid y) = \int P(x \mid y, z) \, P(z \mid y) \, dz$$

Bayes Rule with Background Knowledge

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

Conditional Independence

- X and Y are conditionally independent given Z.

$$P(x, y \mid z) = P(x \mid z) P(y \mid z)$$

$$P(x \mid z) = P(x \mid z, y)$$

$$P(y \mid z) = P(y \mid z, x)$$

Example of State Estimation

- The robot wants to estimate the state of the door as closed or open
  - Has a noisy sensor that produces measurement, z
- Estimate: P(open|z)?
  - Likelihood that the true state of the door is open given that it was measured as open.

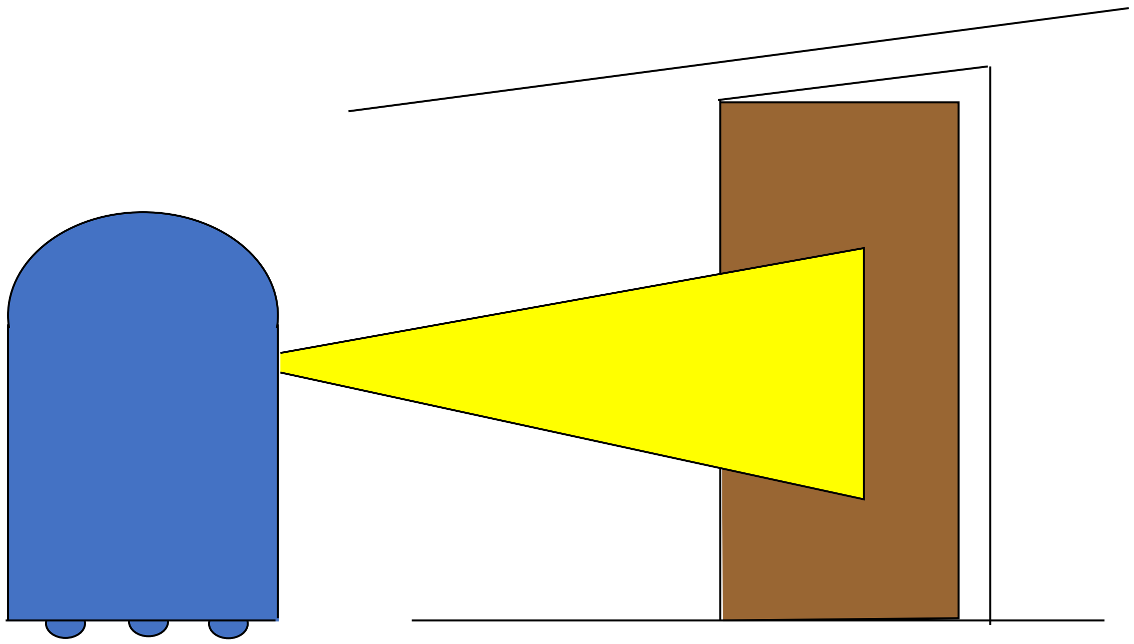

Causal vs. Diagnostic Reasoning

- P(open|z) is diagnostic reasoning
- P(z|open) is causal reasoning (can estimate by counting frequencies)
- Often causal knowledge is easier to obtain.
- Bayes rule enables the use of causal knowledge:

$$P(open \mid z) = \frac{P(z \mid open)P(open)}{P(z)}$$

Example

- Higher likelihood of observation z when the door **is** open compared to when the door is closed.
- The incorporation of the measurement *z* **raises** the probability that the door is open.

$$P(z \mid open) = 0.6 \qquad P(z \mid \neg open) = 0.3$$

$$P(open) = P(\neg open) = 0.5$$

$$P(open \mid z) = \frac{P(z \mid open)P(open)}{P(z \mid open)\,p(open) + P(z \mid \neg open)\,p(\neg open)}$$

$$P(open \mid z) = \frac{0.6 \cdot 0.5}{0.6 \cdot 0.5 + 0.3 \cdot 0.5} = \frac{2}{3} = 0.67$$

Combining Evidence

- Suppose the robot has another sensor that produces a second observation z_2
- How can we combine the measurement of the second sensor
- What is $P(open|z_1, z_2)$?
- In general, how to estimate $P(x| z_1...z_n)$

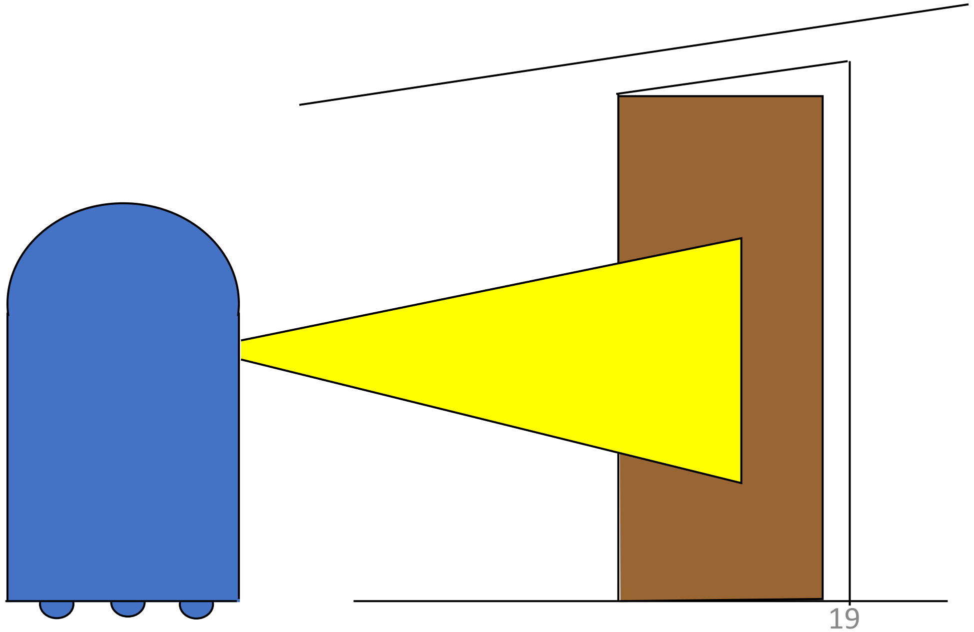

Recursive Bayesian Updating

$$P(x \mid z_1,\ldots,z_n) = \frac{P(z_n \mid x, z_1,\ldots,z_{n-1})\, P(x \mid z_1,\ldots,z_{n-1})}{P(z_n \mid z_1,\ldots,z_{n-1})}$$

Markov assumption: z_n is conditionally independent of $z_1,\ldots,z_{n-1}$ given x.

$$\begin{aligned} P(x \mid z_1,\ldots,z_n) &= \frac{P(z_n \mid x)\, P(x \mid z_1,\ldots,z_{n-1})}{P(z_n \mid z_1,\ldots,z_{n-1})} \\ &= \eta\, P(z_n \mid x)\, P(x \mid z_1,\ldots,z_{n-1}) \\ &= \eta_{1\ldots n} \prod_{i=1\ldots n} P(z_i \mid x)\, P(x) \end{aligned}$$

In our causal modeling view, the world state is causing all the observations.

Incorporating second sensor measurement

- Higher likelihood of observation z when the door is **not** open compared to when the door is open.
- The inclusion of the second measurement z_2 **lowers** the probability for the door to be open.

$$P(z_2 \mid open) = 0.5 \qquad P(z_2 \mid \neg open) = 0.6$$

$$P(open \mid z_1) = 2/3 \qquad P(\neg open \mid z_1) = 1/3$$

$$P(open \mid z_2, z_1) = \frac{P(z_2 \mid open)\, P(open \mid z_1)}{P(z_2 \mid open)\, P(open \mid z_1) + P(z_2 \mid \neg open)\, P(\neg open \mid z_1)}$$

$$= \frac{\frac{1}{2} \cdot \frac{2}{3}}{\frac{1}{2} \cdot \frac{2}{3} + \frac{3}{5} \cdot \frac{1}{3}} = \frac{5}{8} = 0.625$$

Modeling the Sensor

- Sensor model
  - Generative model of a sensor measurement given the true state.
  - A conditional distribution over observations given the true state.
  - Observations or measurements can be considered as the noisy projection of the state

$$p(z_t|x_t)$$

Modeling Actions

- Action or Motion model
  - Actions or controls change the state of the world.
  - Incorporate the outcome of an action u into the current "belief", we use the conditional distribution.
  - Specifies how does the state change by application of the action (from the state, x_{t-1} to the state, x_t by executing the action, u_t).

$$p(x_t | x_{t-1}, u_t)$$

Belief

- Belief
  - Expresses the agent's internal knowledge about the state of an aspect of the world.
  - Note: we do not know the true state.
  - The belief estimated from the sensor measurement data and the actions taken till now.

$$Bel(x_t) = p(x_t | z_{1:t}, u_{1:t})$$

Example: Closing the door

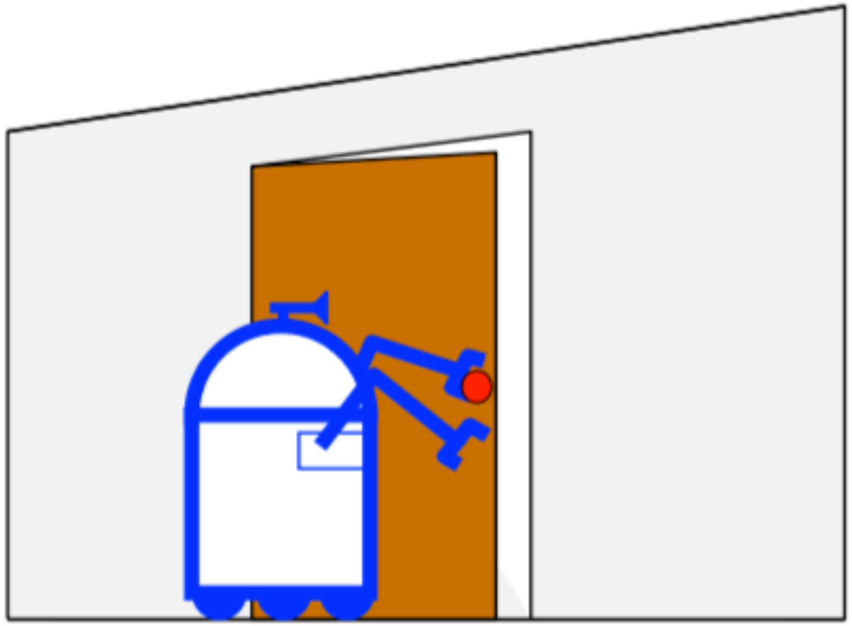

State Transitions

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

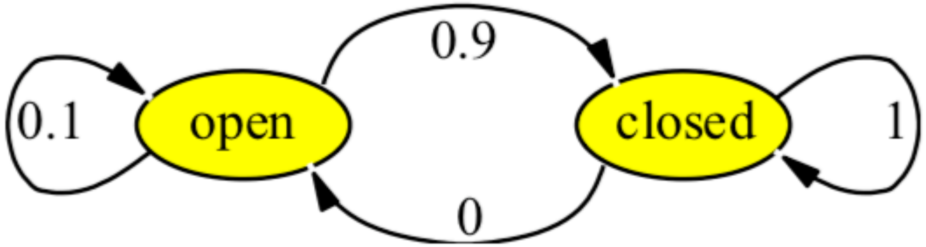

If the door is open, the action "close door" succeeds in 90% of all cases.

Example: The Resulting Belief

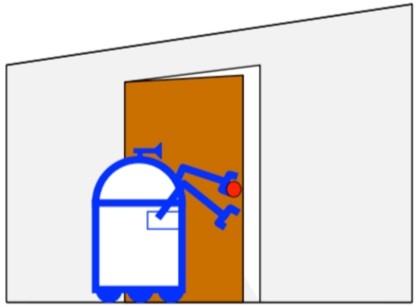

Marginalizing (integrating) out the outcome of actions

Continuous case:

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

Discrete case:

$$P(x \mid u) = \sum P(x \mid u, x')P(x')$$

$$\begin{aligned}
P(closed \mid u) &= \sum P(closed \mid u, x')P(x') \\
&= P(closed \mid u, open)P(open) \\
&\quad + P(closed \mid u, closed)P(closed) \\
&= \frac{9}{10} * \frac{5}{8} + \frac{1}{1} * \frac{3}{8} = \frac{15}{16}
\end{aligned}$$

$$\begin{aligned}
P(open \mid u) &= \sum P(open \mid u, x')P(x') \\
&= P(open \mid u, open)P(open) \\
&\quad + P(open \mid u, closed)P(closed) \\
&= \frac{1}{10} * \frac{5}{8} + \frac{0}{1} * \frac{3}{8} = \frac{1}{16} \\
&= 1 - P(closed \mid u)
\end{aligned}$$

Incorporating Measurements

- Bayes rule

$$P(x|z) = \frac{\boxed{P(z|x)}\,P(x)}{P(z)} = \frac{\text{likelihood} \cdot \text{prior}}{\text{evidence}}$$

Bayes Filter

- **Given:**
  - Stream of observations *z* and action data *u:* $d_t = \{u_1, z_2 \ldots, u_{t-1}, z_t\}$
  - Sensor model $p(z_t | x_t)$
  - Action model $p(x_t | x_{t-1}, u_t)$
  - Prior probability of the system state *P(x).*

- **What we want to estimate?**
  - The state at time t
  - A belief or the posterior over the states: $Bel(x_t) = P(x_t \mid u_1, z_2 \ldots, u_{t-1}, z_t)$

Generative Model

Assumptions

- Static world
- Independent noise
- Perfect model, no approximation errors
- Markov assumption (once you know the state the past actions and observations do not affect the future).

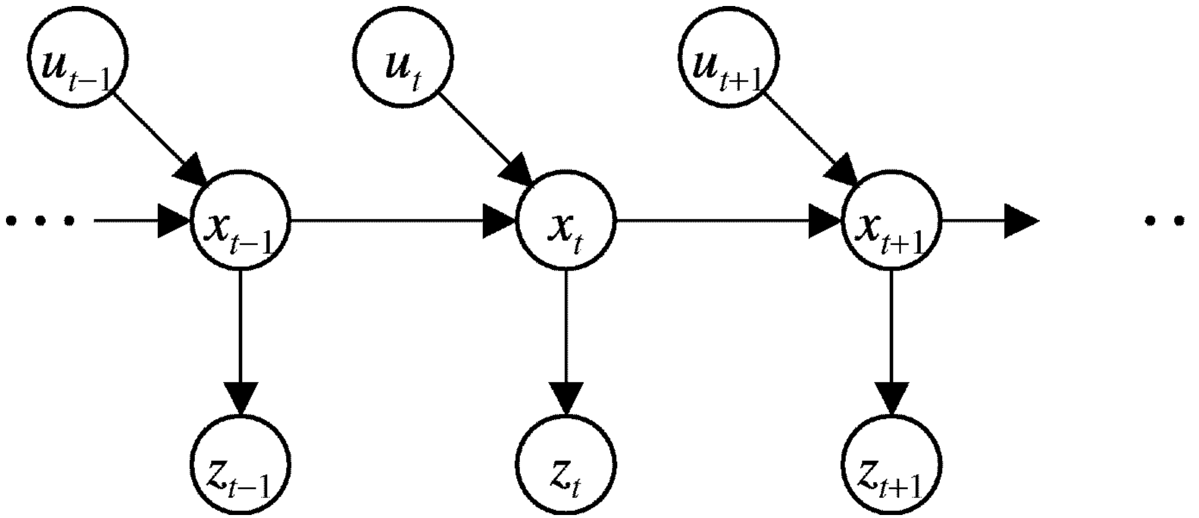

$$p(z_t \mid x_{0:t}, z_{1:t-1}, u_{1:t}) = p(z_t \mid x_t)$$

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

Bayes Filters

z = observation
u = action
x = state

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

Bayes

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

Markov

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

Total prob.

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

Markov

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

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

Bayes Filters Algorithm

Algorithm Bayes_filter (*Bel(x), d*):

n=0

If *d* is a perceptual data item *z* then

For all x do

$$Bel'(x) = P(z \mid x) Bel(x)$$

$$\eta = \eta + Bel'(x)$$

For all x do

$$Bel'(x) = \eta^{-1} Bel'(x)$$

Else if *d* is an action data item *u* then

For all x do

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

Return *Bel'(x)*

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

Bayes Filter: Takeaways

- Bayes filters are a probabilistic tool for estimating the state of with observations acquired over time.
- Bayes rule allows us to compute probabilities that are difficult to determine otherwise.
- Under the Markov assumption, recursive Bayesian updating can be used to efficiently combine evidence.

Hidden Markov Models

- The state of the world changes with time.
- Predict it with successive observations.
- Discrete states and observations.

$\mathbf{X}_t$ = set of unobservable state variables at time t
e.g., $BloodSugar_t$, $StomachContents_t$, etc.

$\mathbf{E}_t$ = set of observable evidence variables at time t
e.g., $MeasuredBloodSugar_t$, $PulseRate_t$, $FoodEaten_t$

This assumes **discrete time**; the step size depends on the problem

Notation: $\mathbf{X}_{a:b} = \mathbf{X}_a, \mathbf{X}_{a+1}, \ldots, \mathbf{X}_{b-1}, \mathbf{X}_b$

HMMs: Conditional Independences

- Future depends on past via the present
- Current observation independent of all else given current state
- Note: there is no explicit notion of controls or actions.

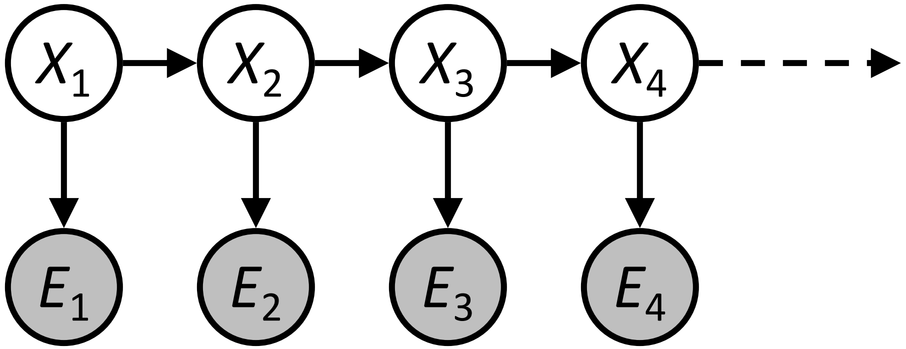

$$\mathbf{P}(\mathbf{X}_t \mid \mathbf{X}_{0:t-1}) = \mathbf{P}(\mathbf{X}_t \mid \mathbf{X}_{t-1})$$

$$\mathbf{P}(\mathbf{E}_t \mid \mathbf{X}_{0:t}, \mathbf{E}_{0:t-1}) = \mathbf{P}(\mathbf{E}_t \mid \mathbf{X}_t)$$

Example: Rain HMM

- Observations: a person carries an umbrella or not.
- States: rainy or not.
- Noisy transition and observation models.

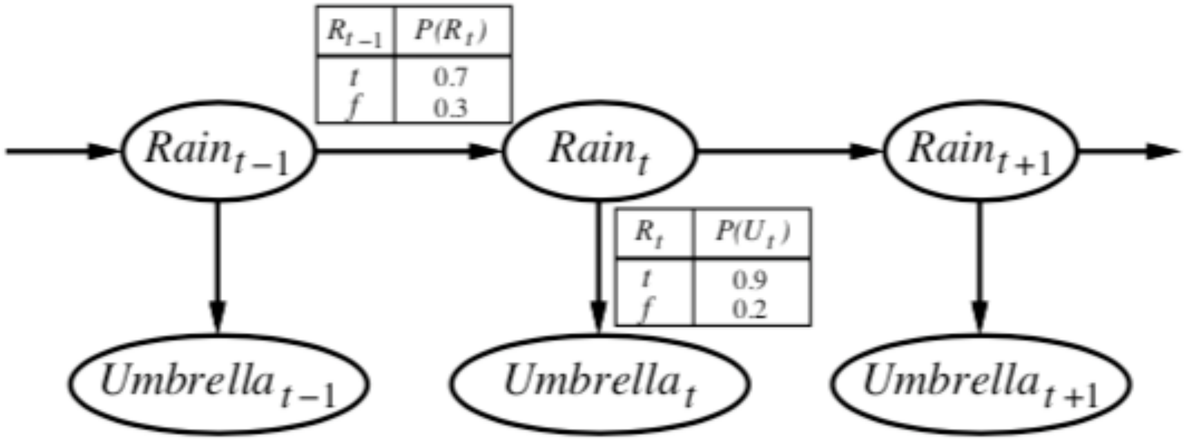

Inference Tasks

Filtering: $\mathbf{P}(\mathbf{X}_t|\mathbf{e}_{1:t})$

to compute the current belief state given all evidence

better name: state estimation

Prediction: $\mathbf{P}(\mathbf{X}_{t+k}|\mathbf{e}_{1:t})$ for $k > 0$

to compute a **future** belief state, given current evidence

(it's like filtering without all evidence)

Smoothing: $\mathbf{P}(\mathbf{X}_k|\mathbf{e}_{1:t})$ for $0 \leq k < t$

to compute a better estimate of past states

Most likely explanation: $\arg\max_{\mathbf{x}_{1:t}} P(\mathbf{x}_{1:t}|\mathbf{e}_{1:t})$

to compute the state sequence that is most likely, given the evidence

HMM Filtering

$$\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t+1}) = f(\mathbf{e}_{t+1}, \mathbf{P}(\mathbf{X}_t|\mathbf{e}_{1:t}))$$

$$\begin{aligned}\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t+1}) &= \mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t}, \mathbf{e}_{t+1}) \\ &= \alpha\mathbf{P}(\mathbf{e}_{t+1}|\mathbf{X}_{t+1}, \mathbf{e}_{1:t})\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t}) \\ &= \alpha\mathbf{P}(\mathbf{e}_{t+1}|\mathbf{X}_{t+1})\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t})\end{aligned}$$

$$\begin{aligned}\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t+1}) &= \alpha\mathbf{P}(\mathbf{e}_{t+1}|\mathbf{X}_{t+1})\Sigma_{\mathbf{x}_t}\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{x}_t, \mathbf{e}_{1:t})P(\mathbf{x}_t|\mathbf{e}_{1:t}) \\ &= \alpha\mathbf{P}(\mathbf{e}_{t+1}|\mathbf{X}_{t+1})\Sigma_{\mathbf{x}_t}\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{x}_t)P(\mathbf{x}_t|\mathbf{e}_{1:t})\end{aligned}$$

Inference: Estimate State Given Evidence

- We are given evidence at each time and want to know

$$B_t(X) = P(X_t|e_{1:t})$$

- Approach: start with $P(X_1)$ and derive B_t in terms of B_{t-1}
  - Equivalently, derive B_{t+1} in terms of B_t

- Two Steps:
  - Passage of time
  - Evidence incorporation

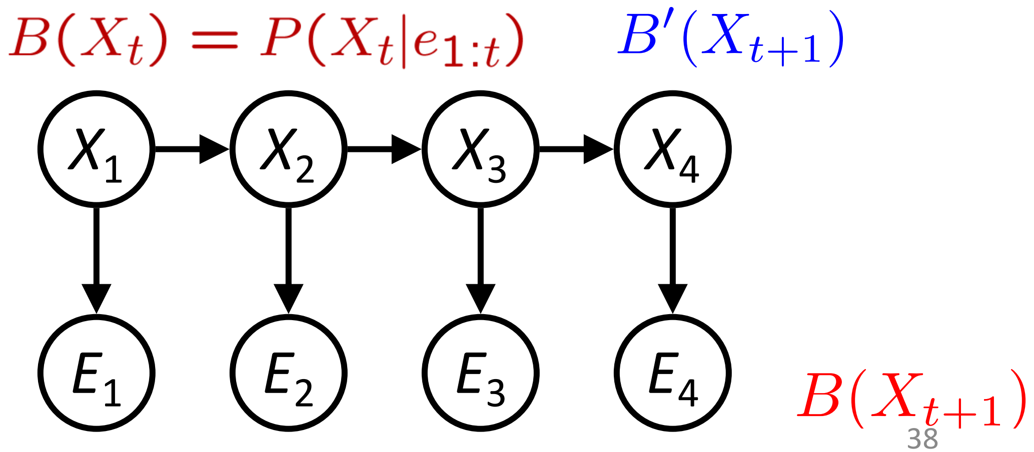

Passage of Time

Assume we have current belief P(X | evidence to date)

$$B(X_t) = P(X_t|e_{1:t})$$

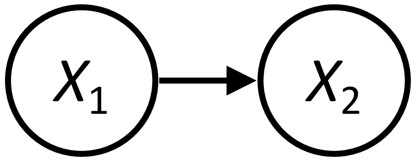

Then, after one time step:

$$P(X_{t+1}|e_{1:t}) = \sum_{x_t} P(X_{t+1}, x_t|e_{1:t})$$

$$= \sum_{x_t} P(X_{t+1}|x_t, e_{1:t})P(x_t|e_{1:t})$$

$$= \sum_{x_t} P(X_{t+1}|x_t)P(x_t|e_{1:t})$$

Basic idea: the beliefs get "pushed" through the transitions

Incorporating Observations

Assume we have current belief P(X | previous evidence):

$$B'(X_{t+1}) = P(X_{t+1}|e_{1:t})$$

Then, after evidence comes in:

$$P(X_{t+1}|e_{1:t+1}) = P(X_{t+1}, e_{t+1}|e_{1:t})/P(e_{t+1}|e_{1:t})$$

$$\propto_{X_{t+1}} P(X_{t+1}, e_{t+1}|e_{1:t})$$

$$= P(e_{t+1}|e_{1:t}, X_{t+1})P(X_{t+1}|e_{1:t})$$

$$= P(e_{t+1}|X_{t+1})P(X_{t+1}|e_{1:t})$$

View it as a "correction" of the belief using the observation

$$B(X_{t+1}) \propto_{X_{t+1}} P(e_{t+1}|X_{t+1})B'(X_{t+1})$$

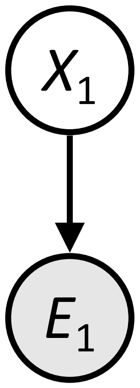

Example: Robot Localization

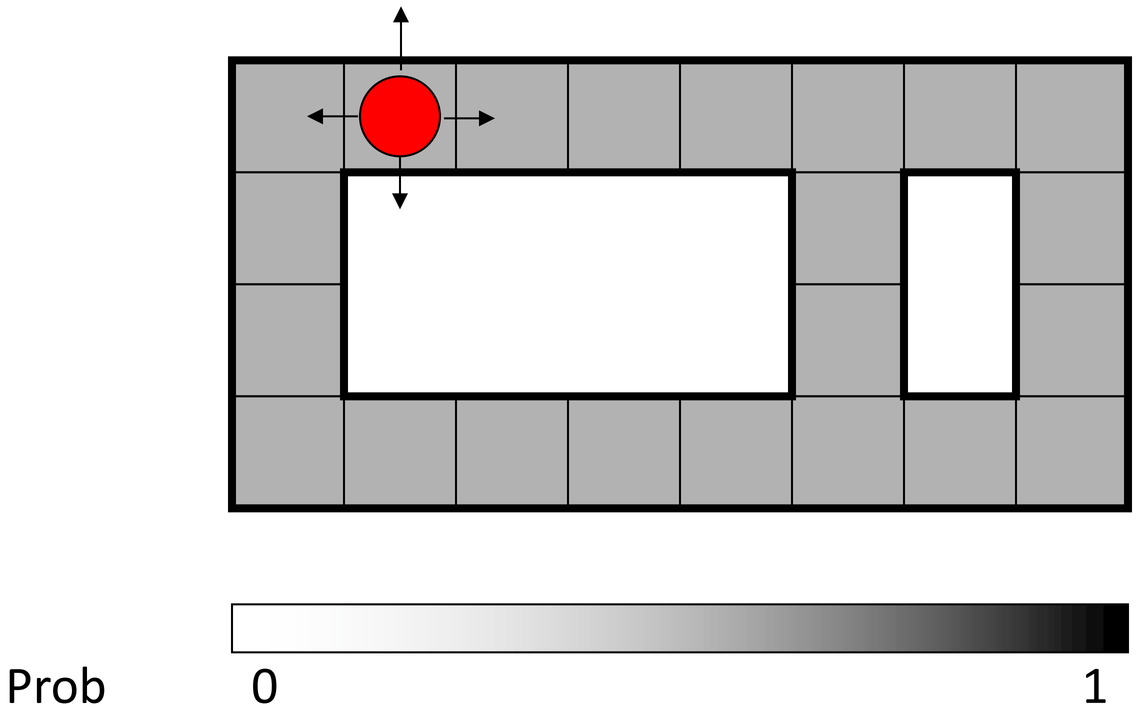

t=0

Robot can take actions N, S, E, W
Detects walls from its sensors

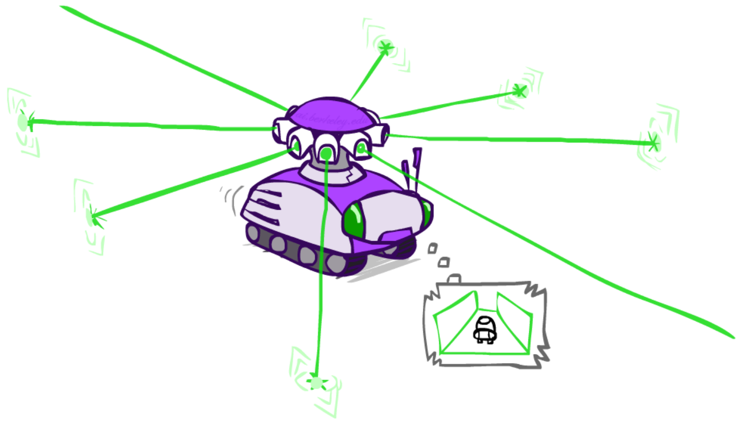

Sensor model: can read in which directions there is a wall, never more than 1 mistake

Motion model: may not execute action with small prob.

Example: Robot Localization

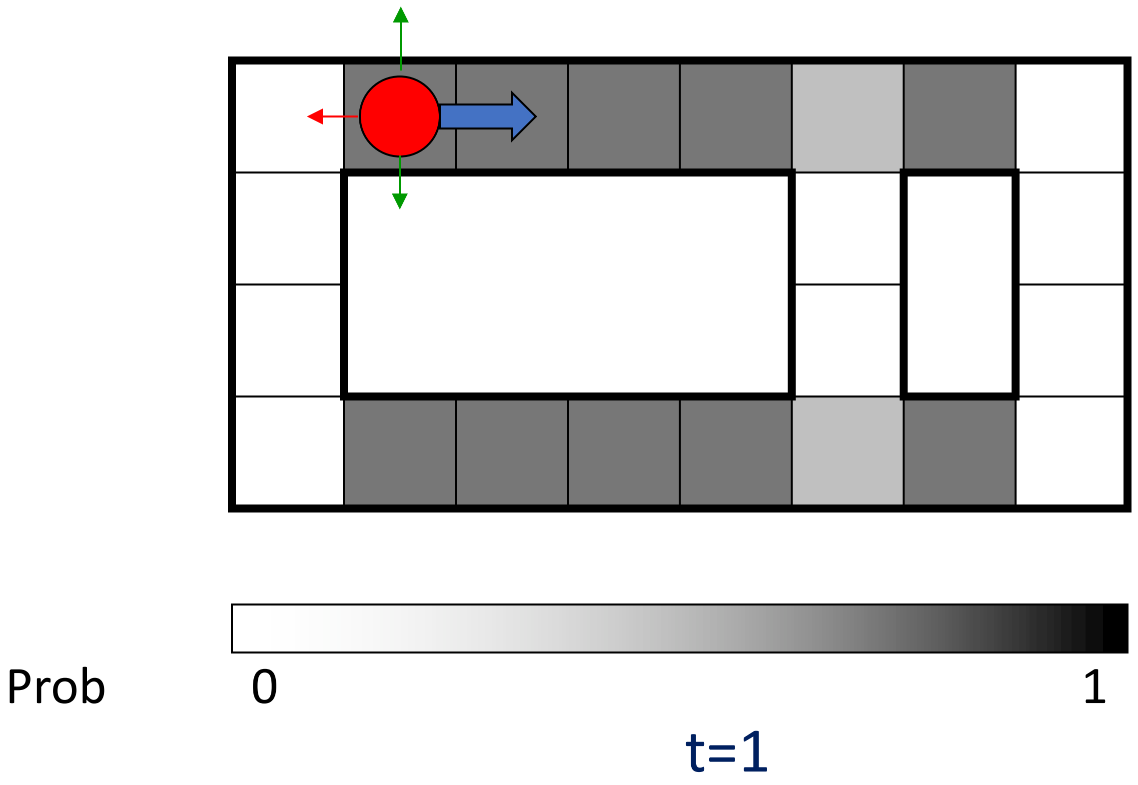

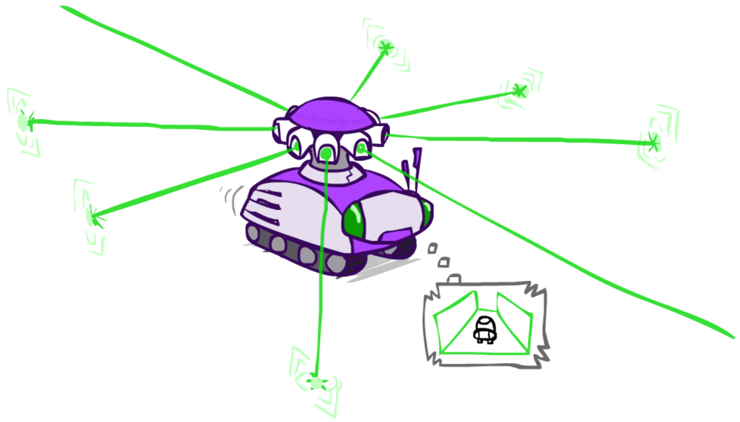

Lighter grey: was possible to get the reading, but less likely b/c required 1 mistake

Example: Robot Localization

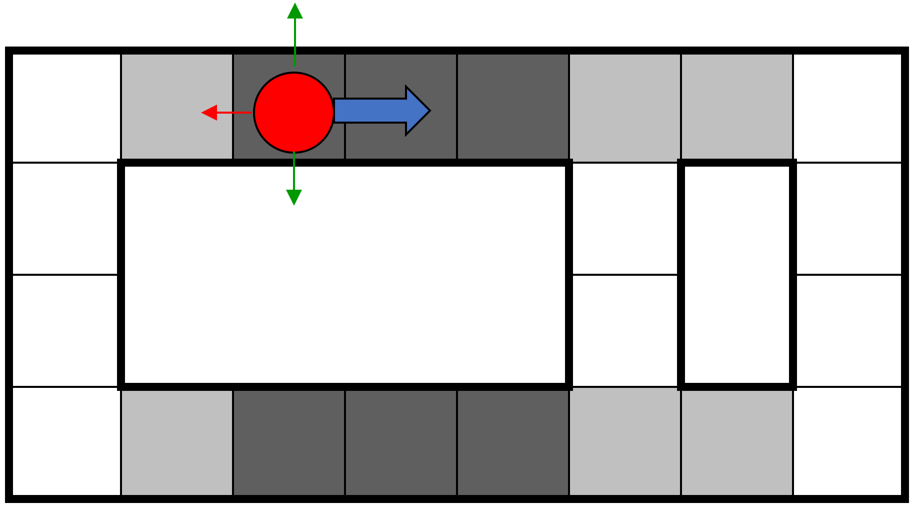

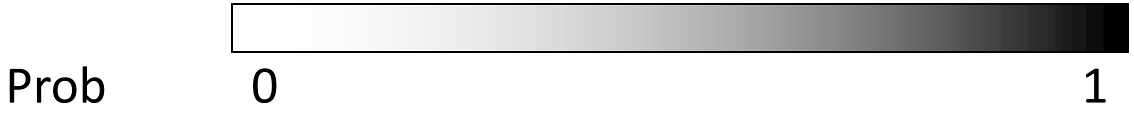

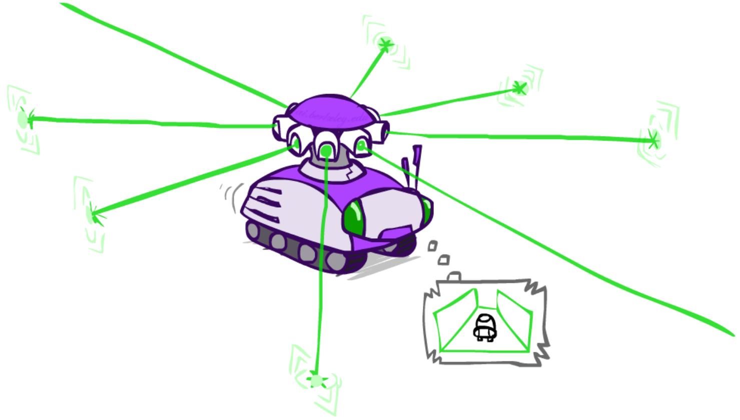

t=2

Example: Robot Localization

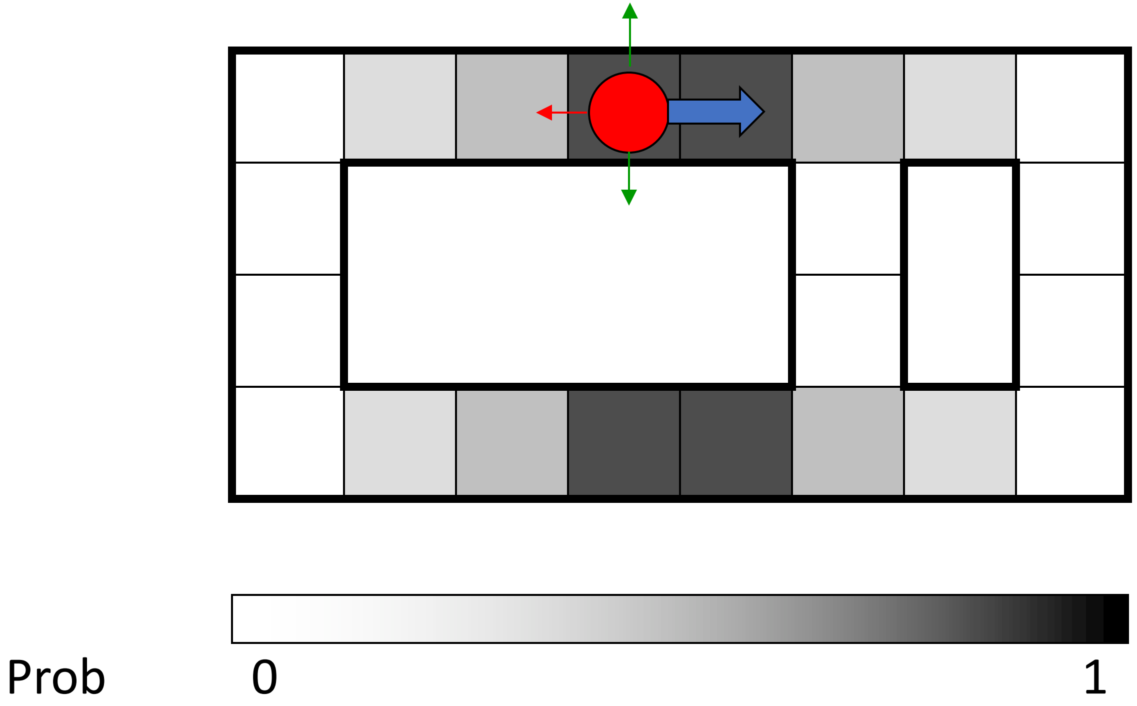

t=3

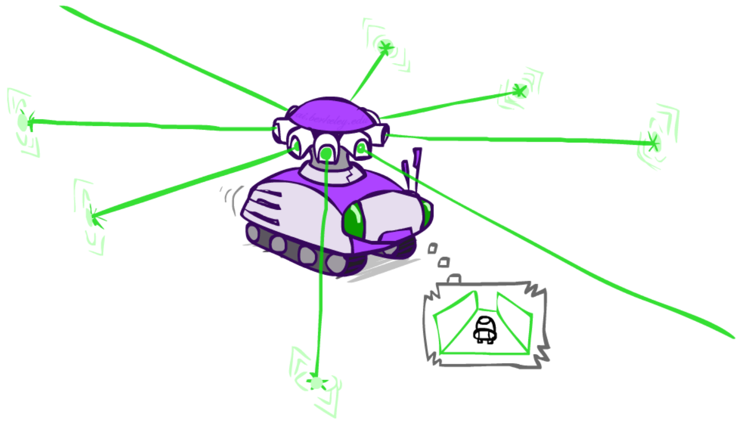

Example: Robot Localization

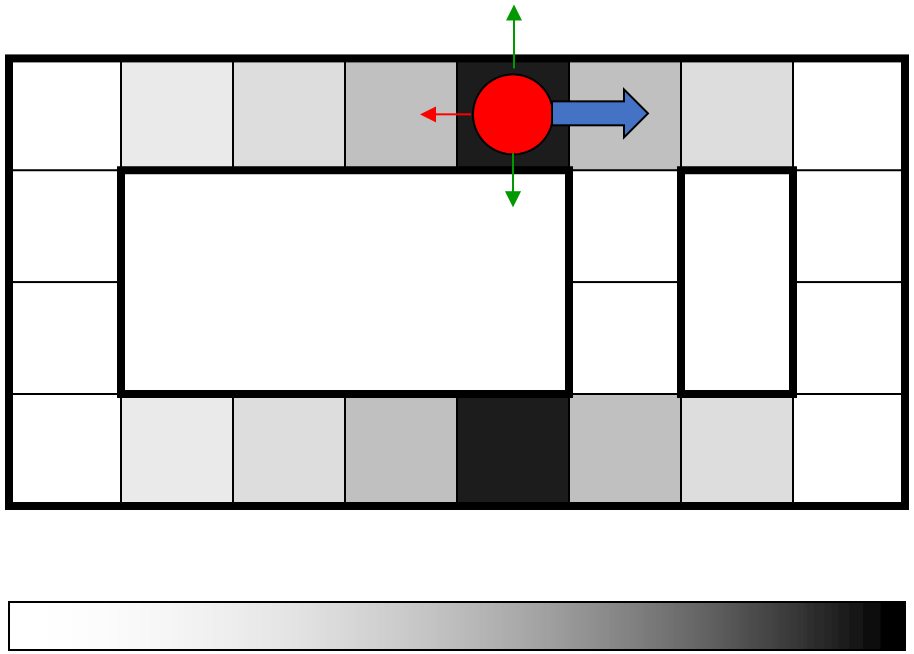

Prob 0 1

t=4

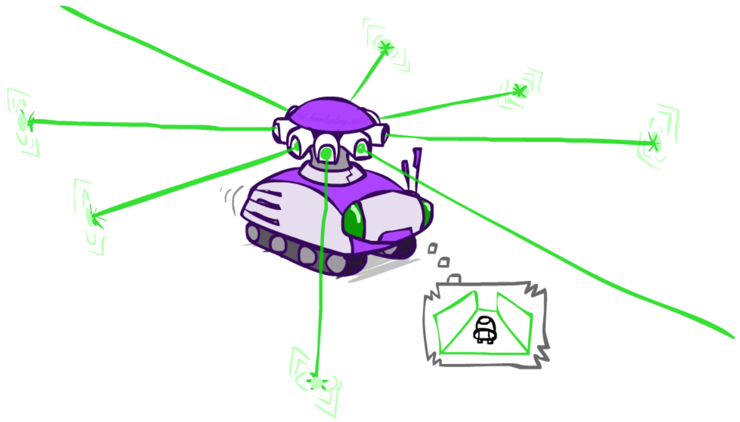

Example: Robot Localization

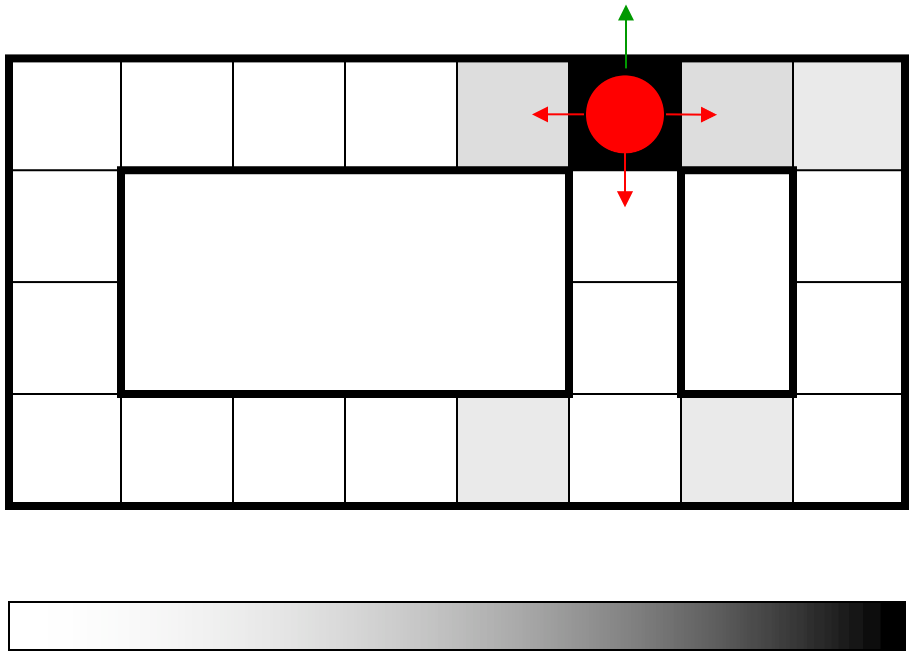

Prob 0 1

t=5

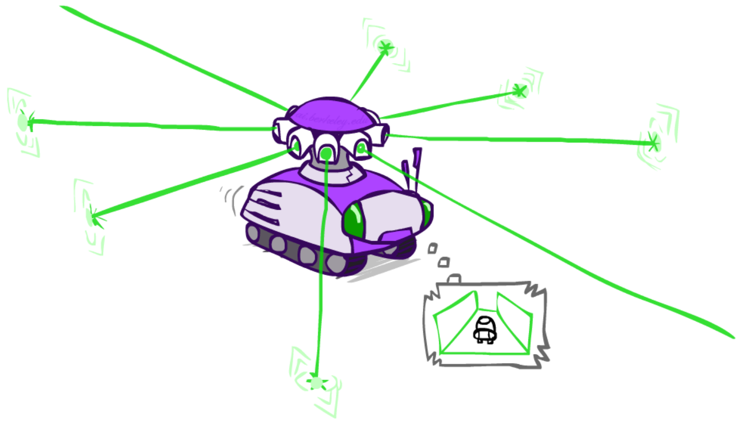

Particle Filtering

- Problem:
  - |X| may be too big to even store B(X)
  - E.g. X is continuous (though here we focus on the discrete case)
- Particle filtering
  - Track samples of X, not all values. Samples are called particles
  - Time per step is linear in the number of samples. Keep the list of particles in memory, not states
  - An approximation. Larger the number of particles, better the approximation.

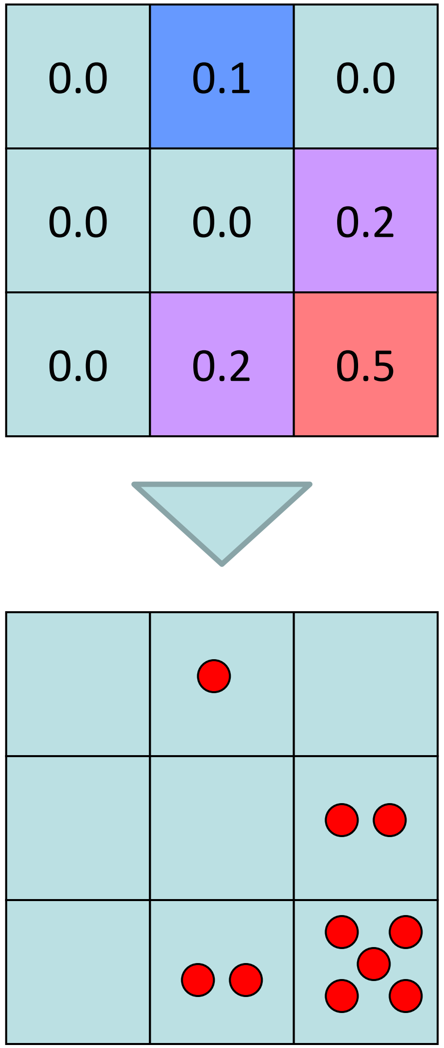

Particles

- Our representation of P(X) is now a list of N particles (samples)
  - Generally, N << |X|

- P(x) approximated by number of particles with value x
  - Several x can have P(x) = 0. Note that (3,3) has half the number of particles.

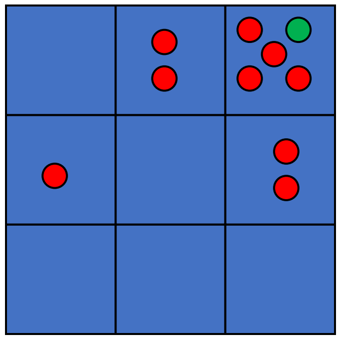

Particles:
(3,3)
(2,3)
(3,3)
(3,2)
(3,3)
(3,2)
(1,2)
(3,3)
(3,3)
(2,3)

Passage of Time

Each particle is moved by sampling its next position from the transition model

$$x' = \text{sample}(P(X'|x))$$

- Example
  - Most samples move clockwise, but some move in another direction or stay in place.
  - An outcome of the probabilistic transition model.

Particles:
(3,3)
(2,3)
(3,3)
(3,2)
(3,3)
(3,2)
(1,2)
(3,3)
(3,3)
(2,3)

Particles:
(3,2)
(2,3)
(3,2)
(3,1)
(3,3)
(3,2)
(1,3)
(2,3)
(3,2)
(2,2)

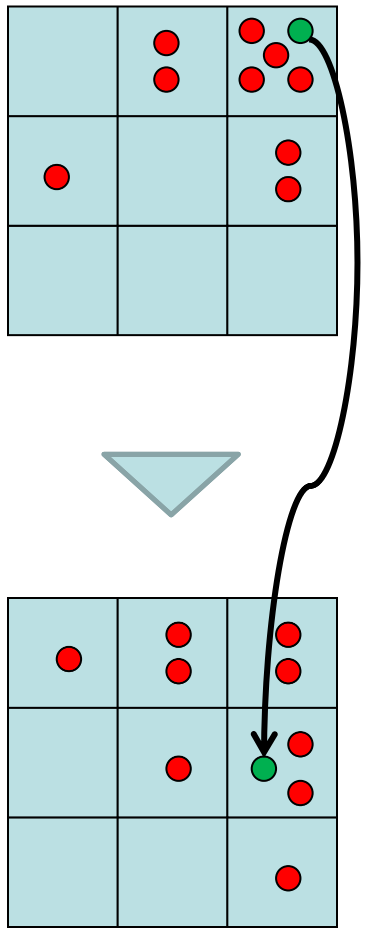

Incorporate Evidence

- Incorporating evidence adjusts or weighs the probabilities.
- Attach a weight to each sample.
- Weigh the samples based on the likelihood of the evidence.

$$w(x) = P(e|x)$$

$$B(X) \propto P(e|X)B'(X)$$

Particles:
(3,2)
(2,3)
(3,2)
(3,1)
(3,3)
(3,2)
(1,3)
(2,3)
(3,2)
(2,2)

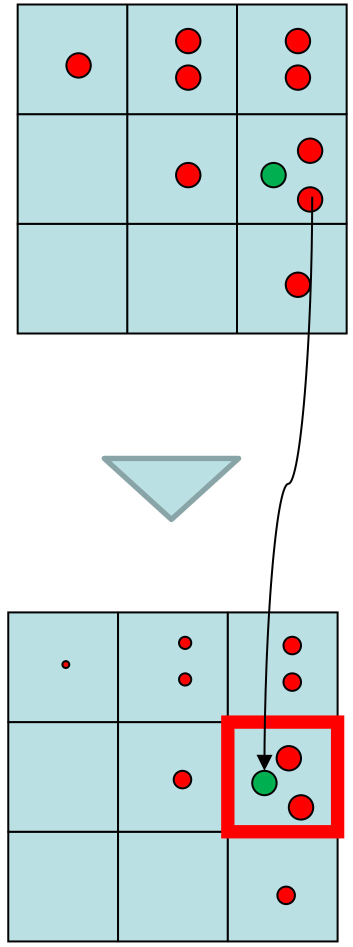

Particles:
(3,2) w=.9
(2,3) w=.2
(3,2) w=.9
(3,1) w=.4
(3,3) w=.4
(3,2) w=.9
(1,3) w=.1
(2,3) w=.2
(3,2) w=.9
(2,2) w=.4

Representation: Resample

- Rather than tracking weighted samples, we resample.
- We choose, N times, from our weighted sample distribution (i.e. draw **with** replacement)
- Now the update is complete for this time step, continue with the next one.

Particles:
(3,2) w=.9
(2,3) w=.2
(3,2) w=.9
(3,1) w=.4
(3,3) w=.4
(3,2) w=.9
(1,3) w=.1
(2,3) w=.2
(3,2) w=.9
(2,2) w=.4

(New) Particles:
(3,2)
(2,2)
(3,2)
(2,3)
(3,3)
(3,2)
(1,3)
(2,3)
(3,2)
(3,2)

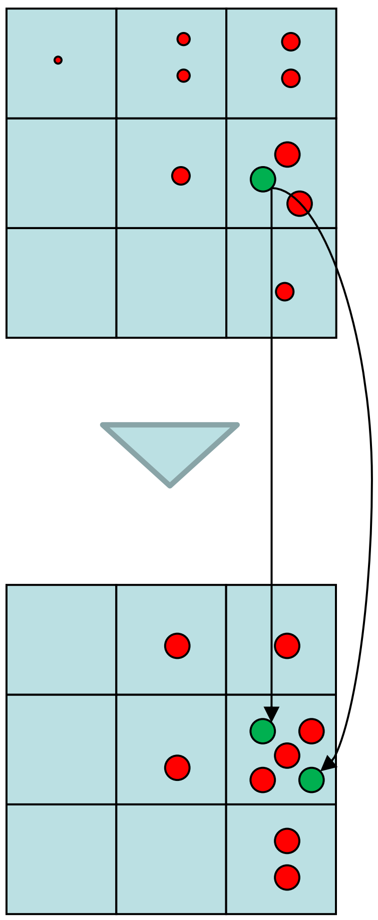

Key idea: maintain hypotheses (particles) in the region of probable states, discard others. Note that the sampling is with replacement.

Representation: Particles

Particles: track samples of states rather than an explicit distribution

Elapse | Weight | Resample

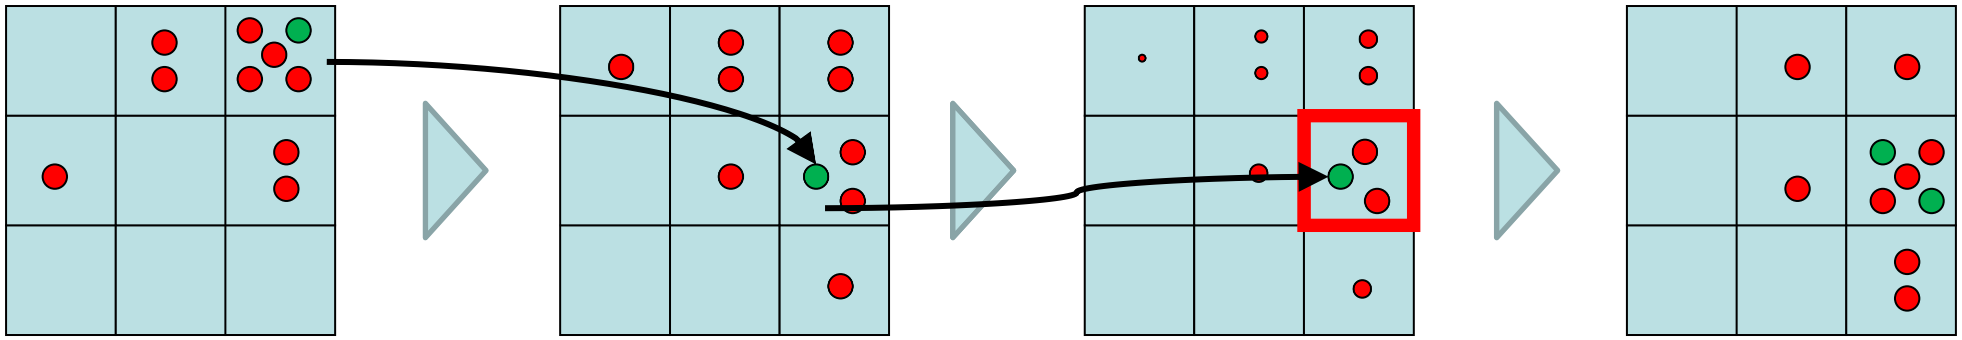

Particles:	Particles:	Particles:	(New) Particles:
(3,3)	(3,2)	(3,2) w=.9	(3,2)
(2,3)	(2,3)	(2,3) w=.2	(2,2)
(3,3)	(3,2)	(3,2) w=.9	(3,2)
(3,2)	(3,1)	(3,1) w=.4	(2,3)
(3,3)	(3,3)	(3,3) w=.4	(3,3)
(3,2)	(3,2)	(3,2) w=.9	(3,2)
(1,2)	(1,3)	(1,3) w=.1	(1,3)
(3,3)	(2,3)	(2,3) w=.2	(2,3)
(3,3)	(3,2)	(3,2) w=.9	(3,2)
(2,3)	(2,2)	(2,2) w=.4	(3,2)

Example