Bayes Filter II

Fast Robots, ECE4160/5160, MAE 4190/5190

E. Farrell Helbling, 4/8/25

Class Action Items

- Lab 9: mapping starts today! The world is setup in the front room of the lab.
  - If you still need to work on Lab 8, there is space in the hallway
  - Remember there are extra points for best stunt and best blooper!
    - We will send a google poll next week after extension period ends.
- Please install the simulator before class on Thursday!
  - Next lecture is a flipped classroom, please bring your laptops! We will help students debug their simulator downloads.
  - We will also help if you have questions about the exercises.

Lab 9: Mapping

- Objective: generate a map using your robot and ToF sensor (range finding)
- Strategy: Place your robot in (at least) 4 marked positions on the floor and spin while taking measurements.
- Control:
  - Open loop
  - PID on orientation (DMP or integrated gyro)
  - PID on angular velocity (gyro or differentiated DMP)
- Sanity check: polar plot, repeated polar plots
- Scatter plot: Use transformation matrices
- Convert to a line-based map

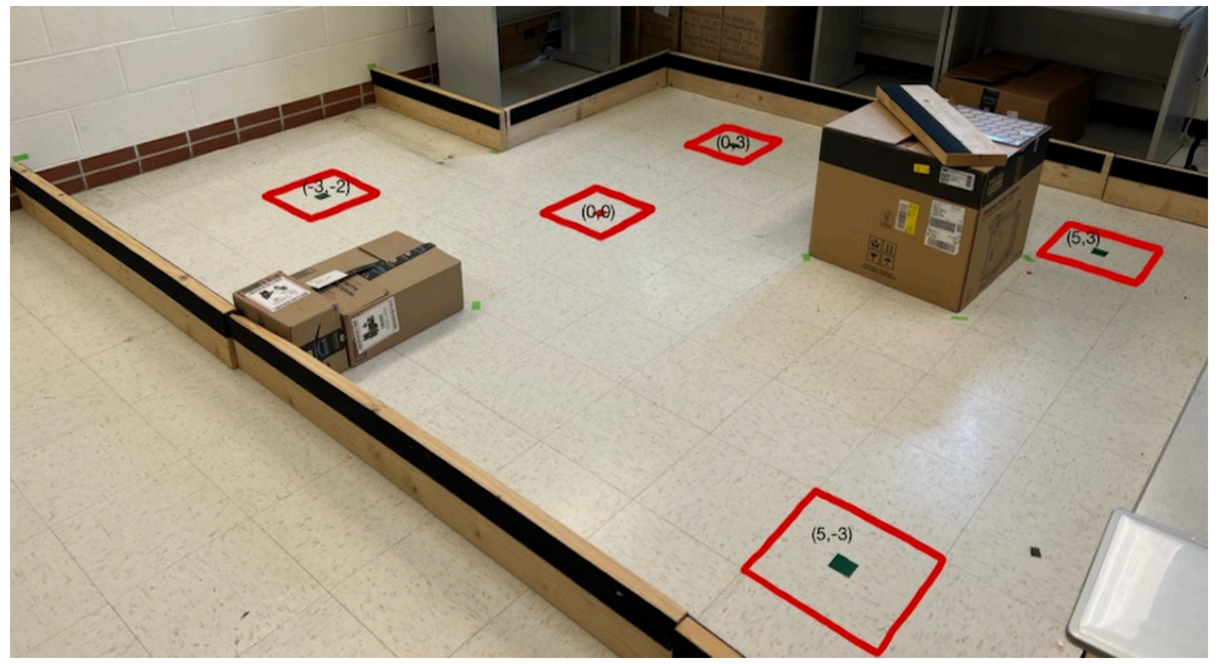

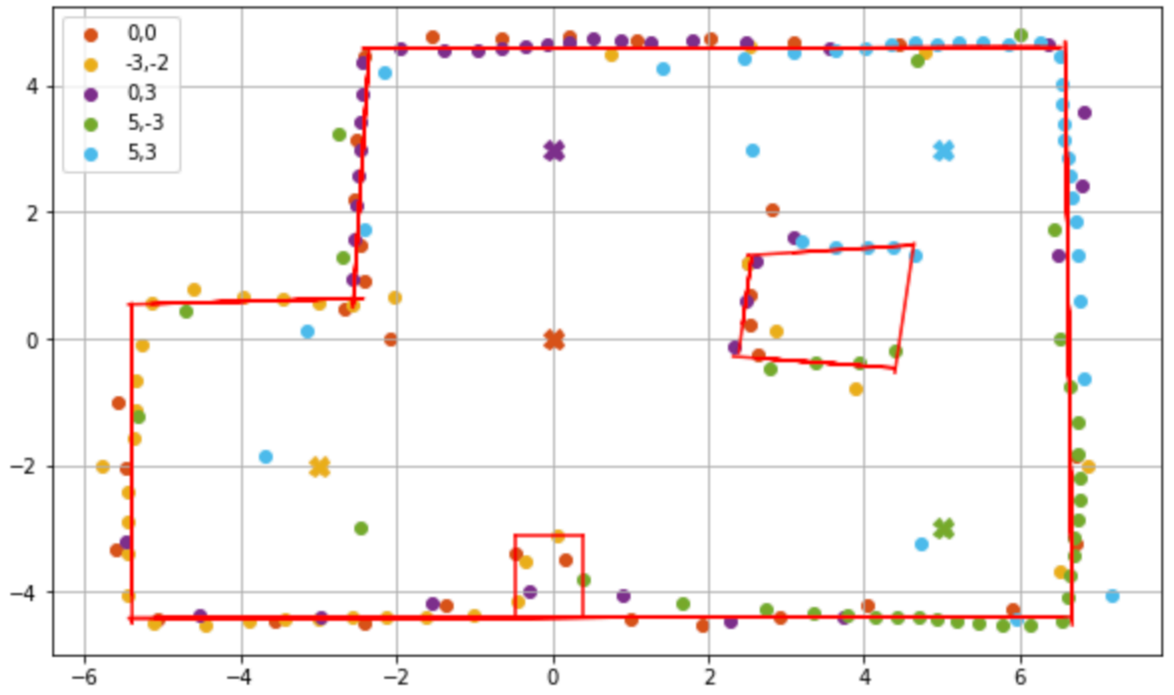

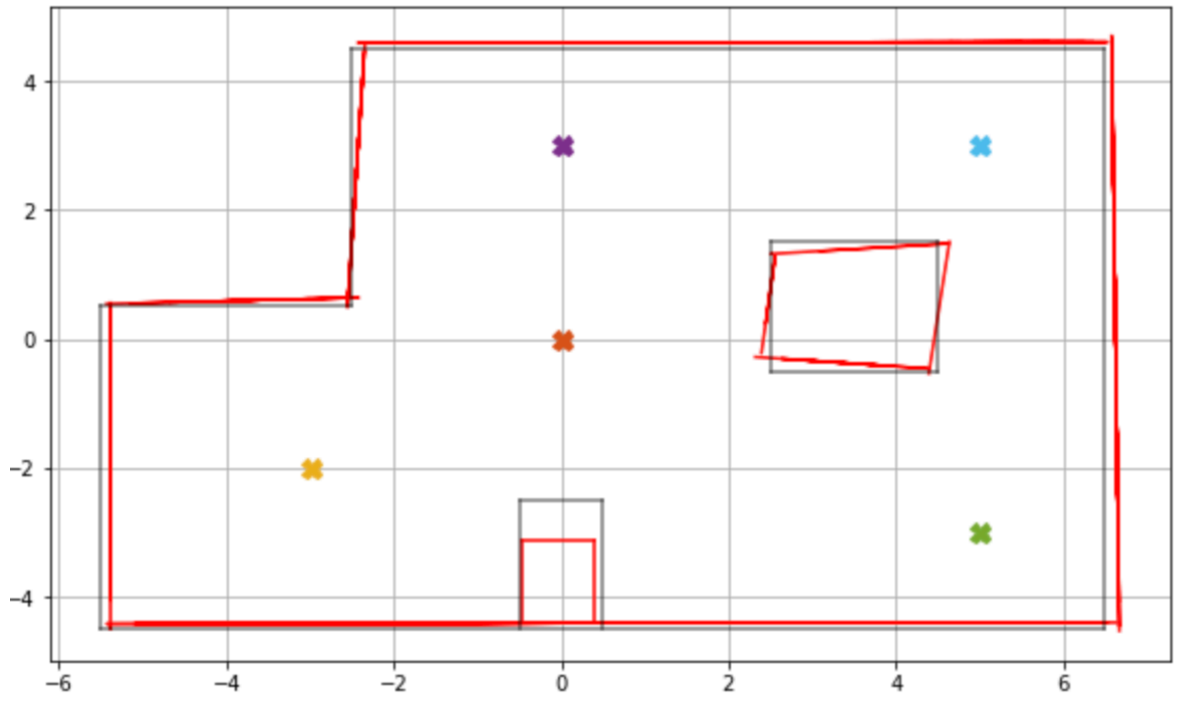

Images from Mikayla Lahr (2024) and Aryaa Pai (2022)

Summary of Bayes Filter

- The robot performs a series of alternating actions/ measurements
- Given:
  - Sensor model: $p(z_t \mid x_t)$
  - Action model: $p(x_t \mid u_t, x_{t-1})$
  - Initial conditions: $p(x_0)$
- Compute:
  - State of dynamic system
  - Posterior of the state (belief): $bel(x_t) = p(x_t \mid u_1, z_1, \ldots, u_t, z_t)$

1. **Algorithm Bayes_Filter** $(bel(x_{t-1}), u_t, z_t)$:
2. **for** all x_t do
3. $\overline{bel}(x_t) = \Sigma_{x_{t-1}}\ p(x_t \mid u_t, x_{t-1})\ bel(x_{t-1})$
4. $bel(x_t) = \eta\ p(z_t \mid x_t)\ \overline{bel}(x_t)$
5. end **for**
6. **return** $bel(x_t)$

Example 1

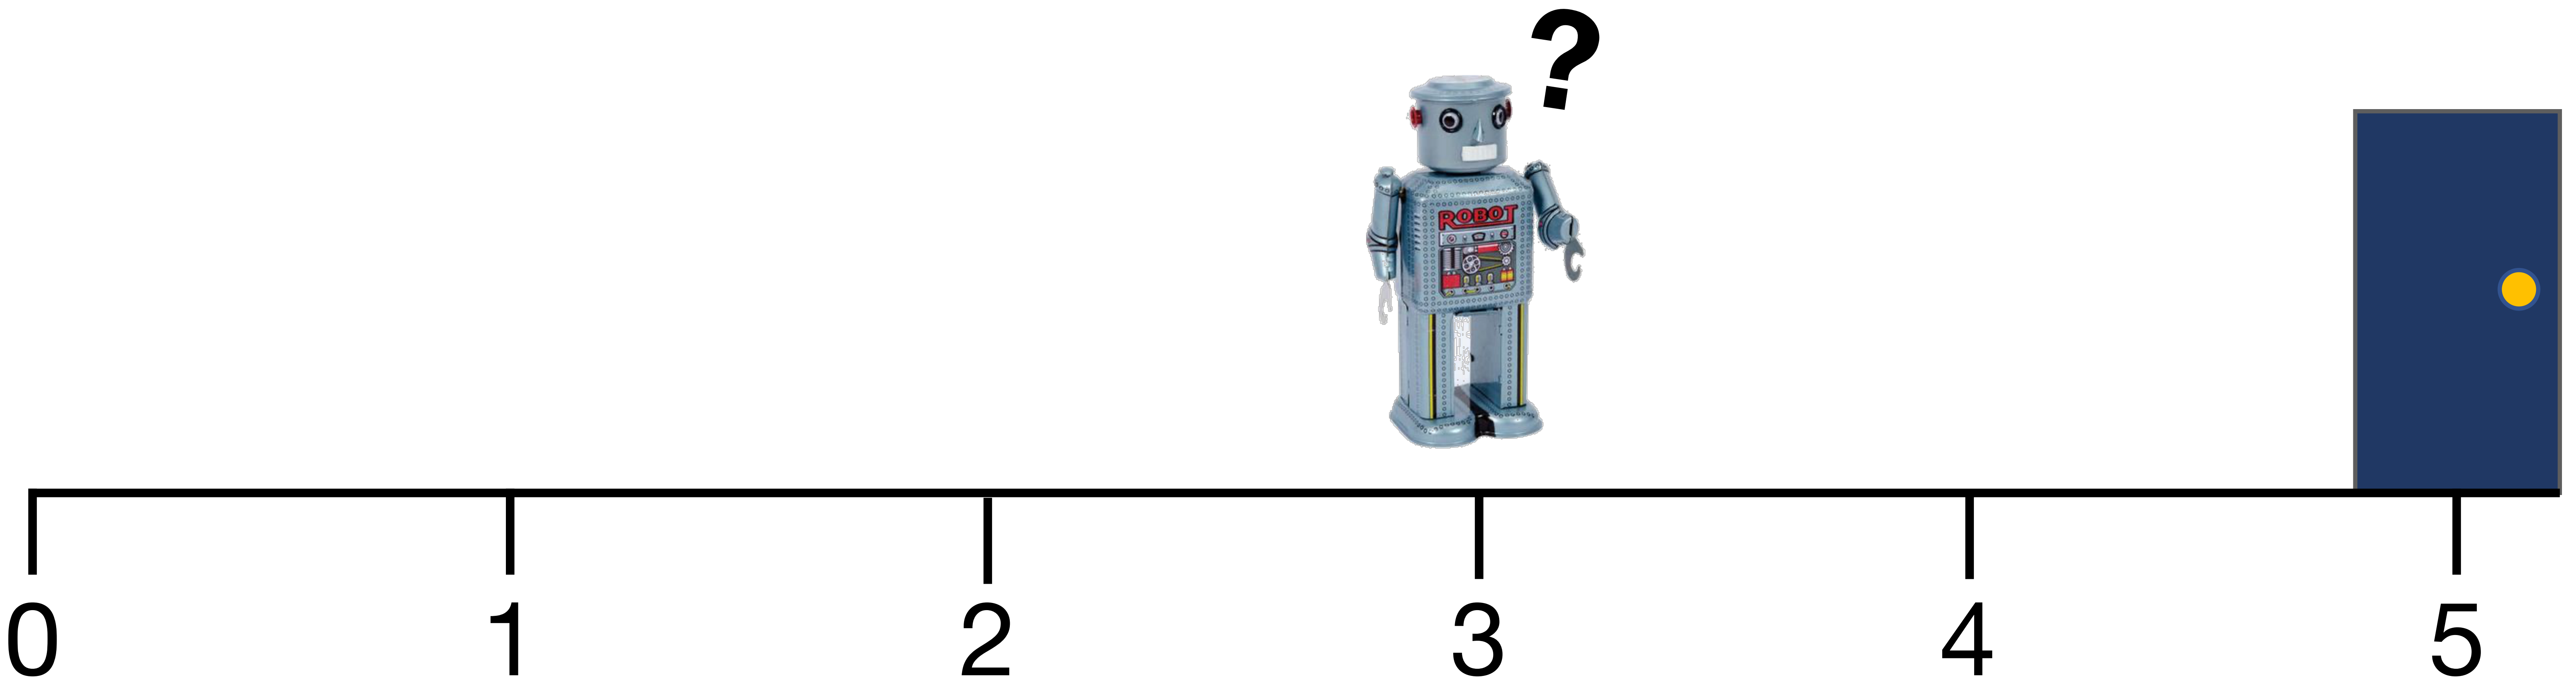

Example 1

- So, what do we need to run the Bayes Filter?

- Motion model

$$p(x+1 \mid x, u=+1) = 0.5$$
$$p(x \mid x, u=+1) = 0.5$$
$$p(x-1 \mid x, u=-1) = 0.5$$
$$p(x \mid x, u=-1) = 0.5$$

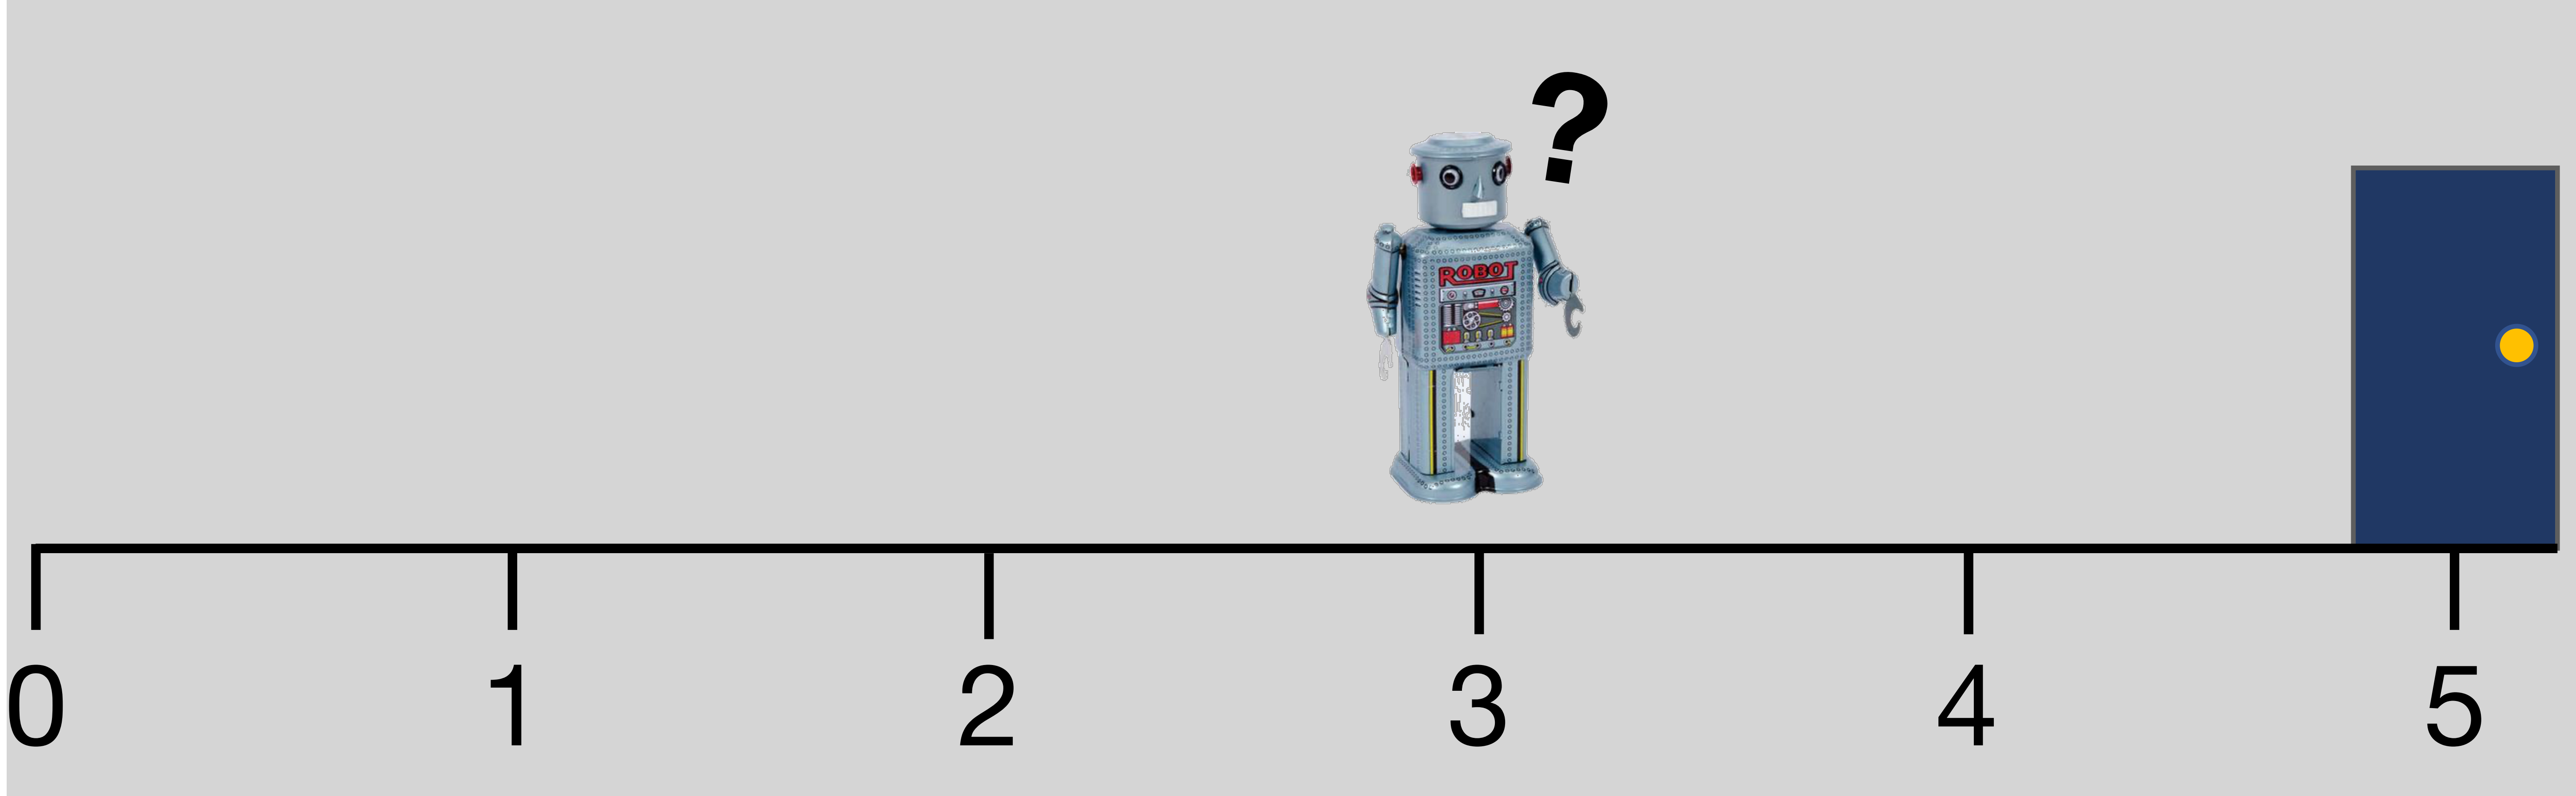

- Measurement model

$$p(Z=\text{door} \mid X=5) = 0.5$$
$$p(Z=\text{door} \mid X=4) = 0.25$$
$$p(Z=\text{door} \mid X=3) = 0$$

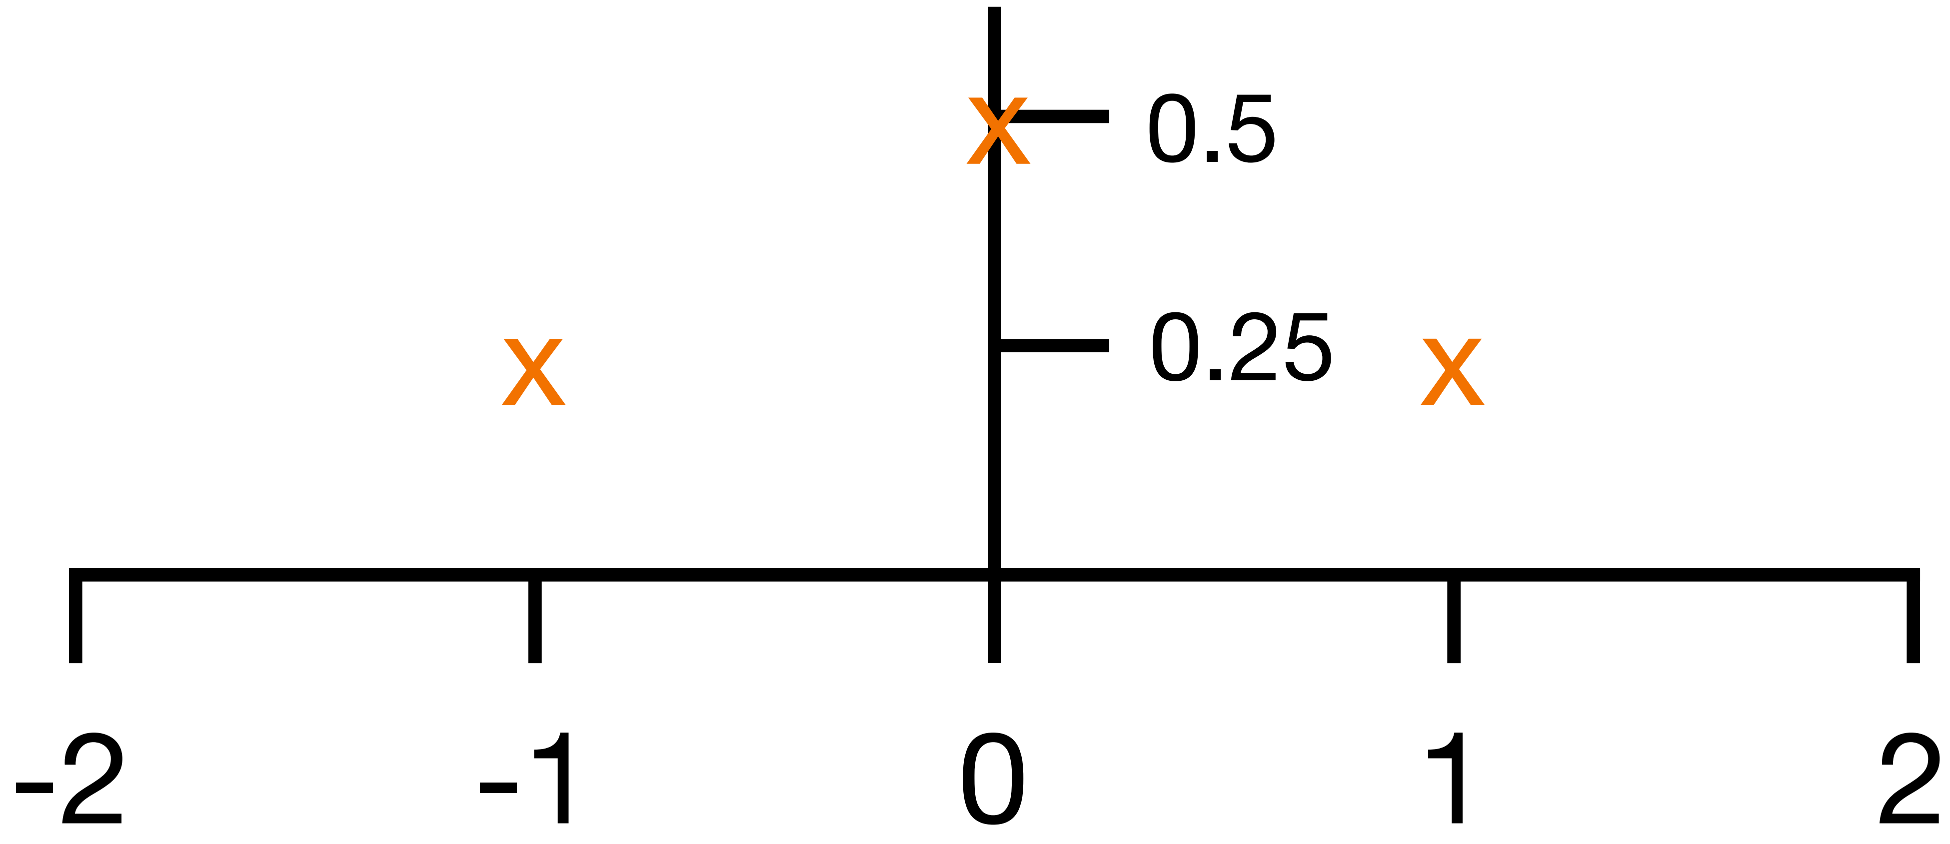

Example 1

At $t = 0$, no information

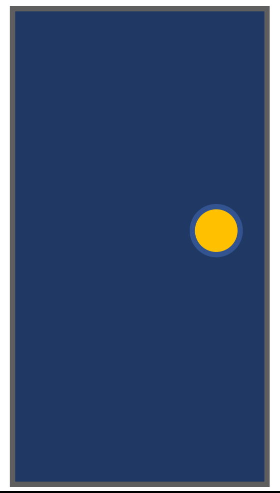

State	0	1	2	3	4	5
$p(x_0)$						

Example 1

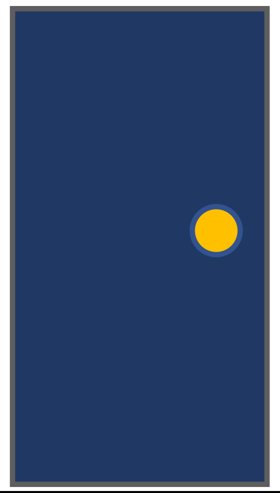

At $t = 0$, no information

State	0	1	2	3	4	5
$p(x_0)$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$

At $t = 1,\ U_1 = \text{NOP},\ Z_1 = \text{door}$

State	0	1	2	3	4	5
$p(x_1)$						

Do we have to do the prediction step?

Do the update step!

Example 1

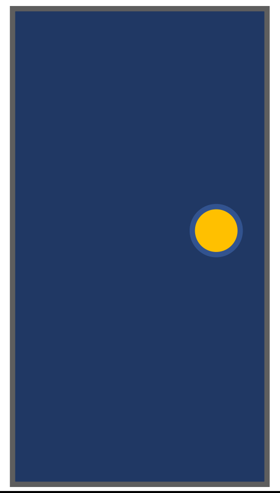

At $t = 0$, no information

State	0	1	2	3	4	5
$p(x_0)$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$

At $t = 1,\ U_1 = \text{NOP},\ Z_1 = \text{door}$

State	0	1	2	3	4	5
$p(x_1)$	0	0	0	0	$\frac{\frac{1}{6} \cdot \frac{1}{4}}{\frac{1}{6} \cdot \frac{1}{4} + \frac{1}{6} \cdot \frac{1}{2}}$	$\frac{\frac{1}{6} \cdot \frac{1}{2}}{\frac{1}{6} \cdot \frac{1}{4} + \frac{1}{6} \cdot \frac{1}{2}}$

Example 1

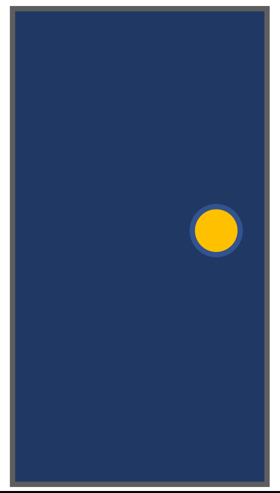

At $t = 1,\ U_1 = \mathrm{NOP},\ Z_1 = \mathrm{door}$

State	0	1	2	3	4	5
$p(x_1)$	0	0	0	0	$\frac{1}{3}$	$\frac{2}{3}$

At $t = 2,\ U_2 = -1$

State	0	1	2	3	4	5
$p(x_2)$						

Example 1

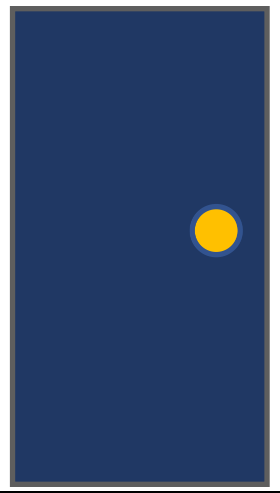

At $t = 1,\ U_1 = \text{NOP},\ Z_1 = \text{door}$

State	0	1	2	3	4	5
$p(x_1)$	0	0	0	0	$\frac{1}{3}$	$\frac{2}{3}$

At $t = 2,\ U_2 = -1$

State	0	1	2	3	4	5
$p(x_2)$	0	0	0	$\frac{1}{3} \cdot \frac{1}{2}$	$\frac{1}{3} \cdot \frac{1}{2} + \frac{2}{3} \cdot \frac{1}{2}$	$\frac{2}{3} \cdot \frac{1}{2}$

Example 1

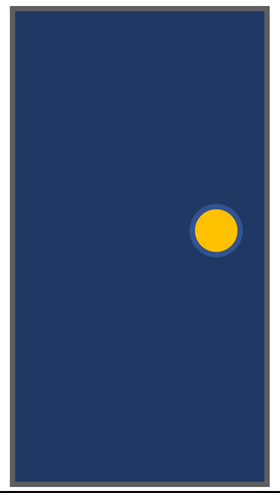

At $t = 2,\ U_2 = -1$

State	0	1	2	3	4	5
$p(x_2)$	0	0	0	$\frac{1}{6}$	$\frac{1}{2}$	$\frac{1}{3}$

At $t = 2,\ U_2 = -1,\ Z_2 = \text{door}$

State	0	1	2	3	4	5
$p(x_2)$						

Example 1

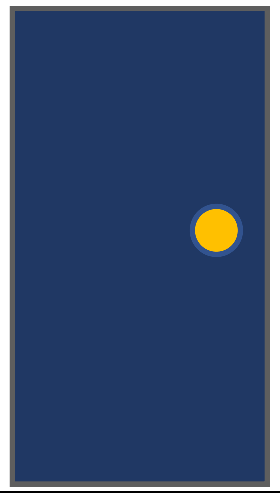

At $t = 2,\ U_2 = -1$

State	0	1	2	3	4	5
$p(x_2)$	0	0	0	$\frac{1}{6}$	$\frac{1}{2}$	$\frac{1}{3}$

At $t = 2,\ U_2 = -1,\ Z_2 = \text{door}$

State	0	1	2	3	4	5
$p(x_2)$	0	0	0	$\frac{1}{6} \cdot 0$	$\dfrac{\frac{1}{2} \cdot \frac{1}{4}}{\frac{1}{2} \cdot \frac{1}{4} + \frac{1}{3} \cdot \frac{1}{2}}$	$\dfrac{\frac{1}{3} \cdot \frac{1}{2}}{\frac{1}{2} \cdot \frac{1}{4} + \frac{1}{3} \cdot \frac{1}{2}}$

Example 1

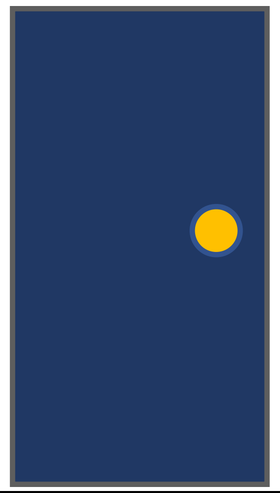

At $t = 0$, we are absolutely certain the robot is at state $X_0 = 0$

State	0	1	2	3	4	5
$p(x_0)$						

Example 1

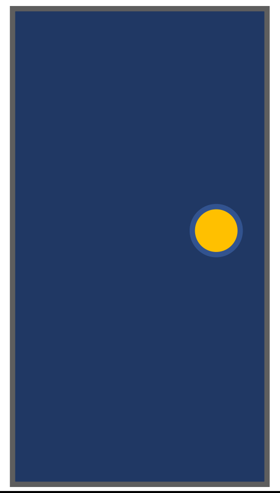

At $t = 0$, we are absolutely certain the robot is at state $X_0 = 0$

State	0	1	2	3	4	5
$p(x_0)$	1	0	0	0	0	0

At $t = 1,\ U_1 = \text{NOP},\ Z_1 = \text{door}$

State	0	1	2	3	4	5
$p(x_1)$						

Example 1

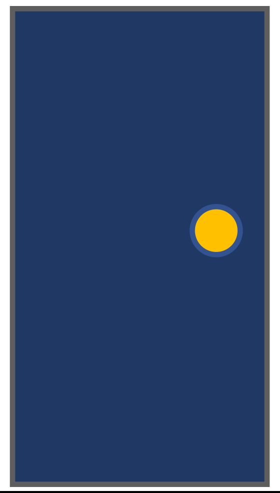

At $t = 0$, we are absolutely certain the robot is at state $X_0 = 0$

State	0	1	2	3	4	5
$p(x_0)$	1	0	0	0	0	0

At $t = 1,\ U_1 = \text{NOP},\ Z_1 = \text{door}$

State	0	1	2	3	4	5
$p(x_1)$	0	0	0	0	0	0

Example 1

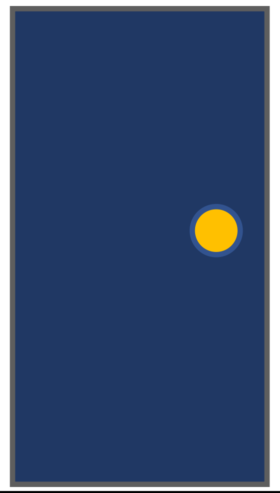

At $t = 0$, we are "absolutely" certain the robot is at state $X_0 = 0$

State	0	1	2	3	4	5
$p(x_0)$	$\frac{19}{20}$	$\frac{1}{100}$	$\frac{1}{100}$	$\frac{1}{100}$	$\frac{1}{100}$	$\frac{1}{100}$

At $t = 1,\ U_1 = \mathrm{NOP},\ Z_1 = \mathrm{door}$

State	0	1	2	3	4	5
$p(x_1)$						

Example 1

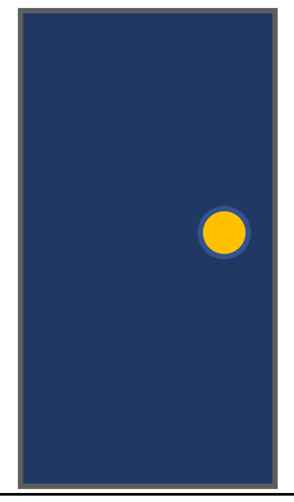

At $t = 0$, we are "absolutely" certain the robot is at state $X_0 = 0$

State	0	1	2	3	4	5
$p(x_0)$	$\frac{19}{20}$	$\frac{1}{100}$	$\frac{1}{100}$	$\frac{1}{100}$	$\frac{1}{100}$	$\frac{1}{100}$

At $t = 1,\ U_1 = \mathrm{NOP},\ Z_1 = \mathrm{door}$

State	0	1	2	3	4	5
$p(x_1)$	0	0	0	0	$\frac{1}{3}$	$\frac{2}{3}$

Always believe, even if just a little, in the improbable!

(deterministic approaches are fragile!)

Example 2

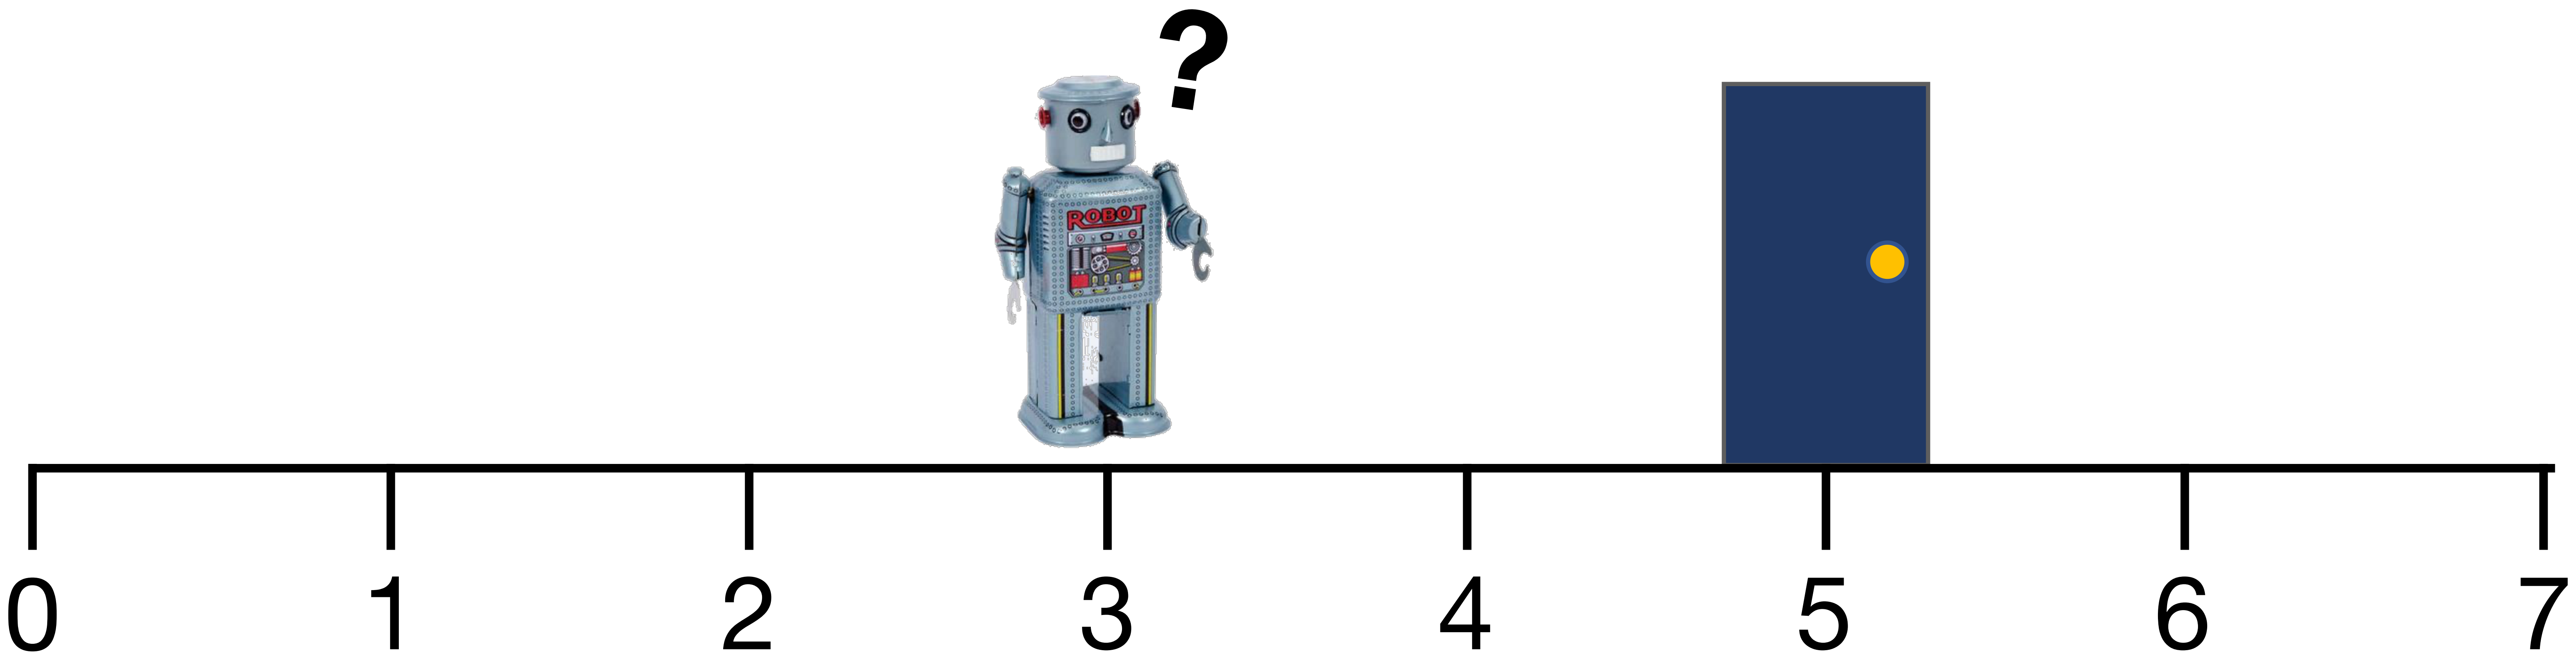

adapted from Prof. Fred Martin at Umass

Example 2

Bayes with beans

- World
  - 1D continuous, 7 states
  - … door at state 5
- Motion model
  - 80% correct, 20% fail
- Sensor model
  - 90% correct, 10% fail
- Initial belief
- Take an action: +1
- Take a sensor reading: door!

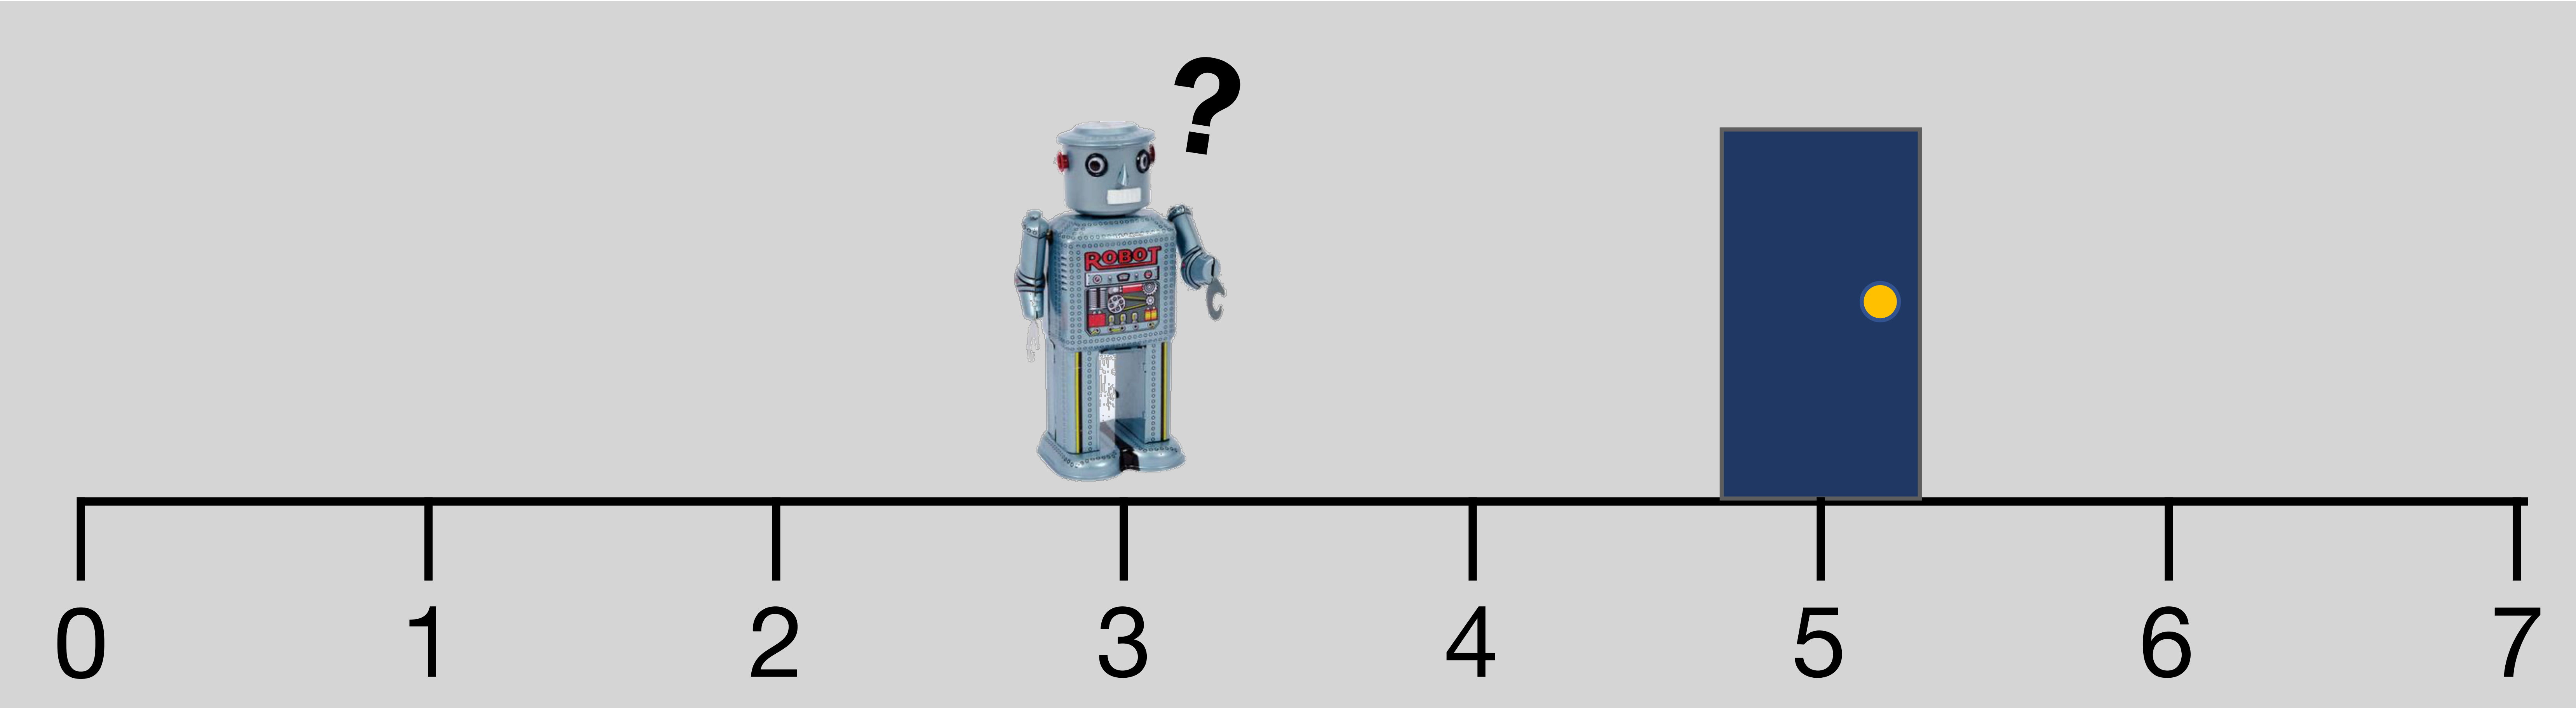

adapted from Prof. Fred Martin at Umass

Example 3

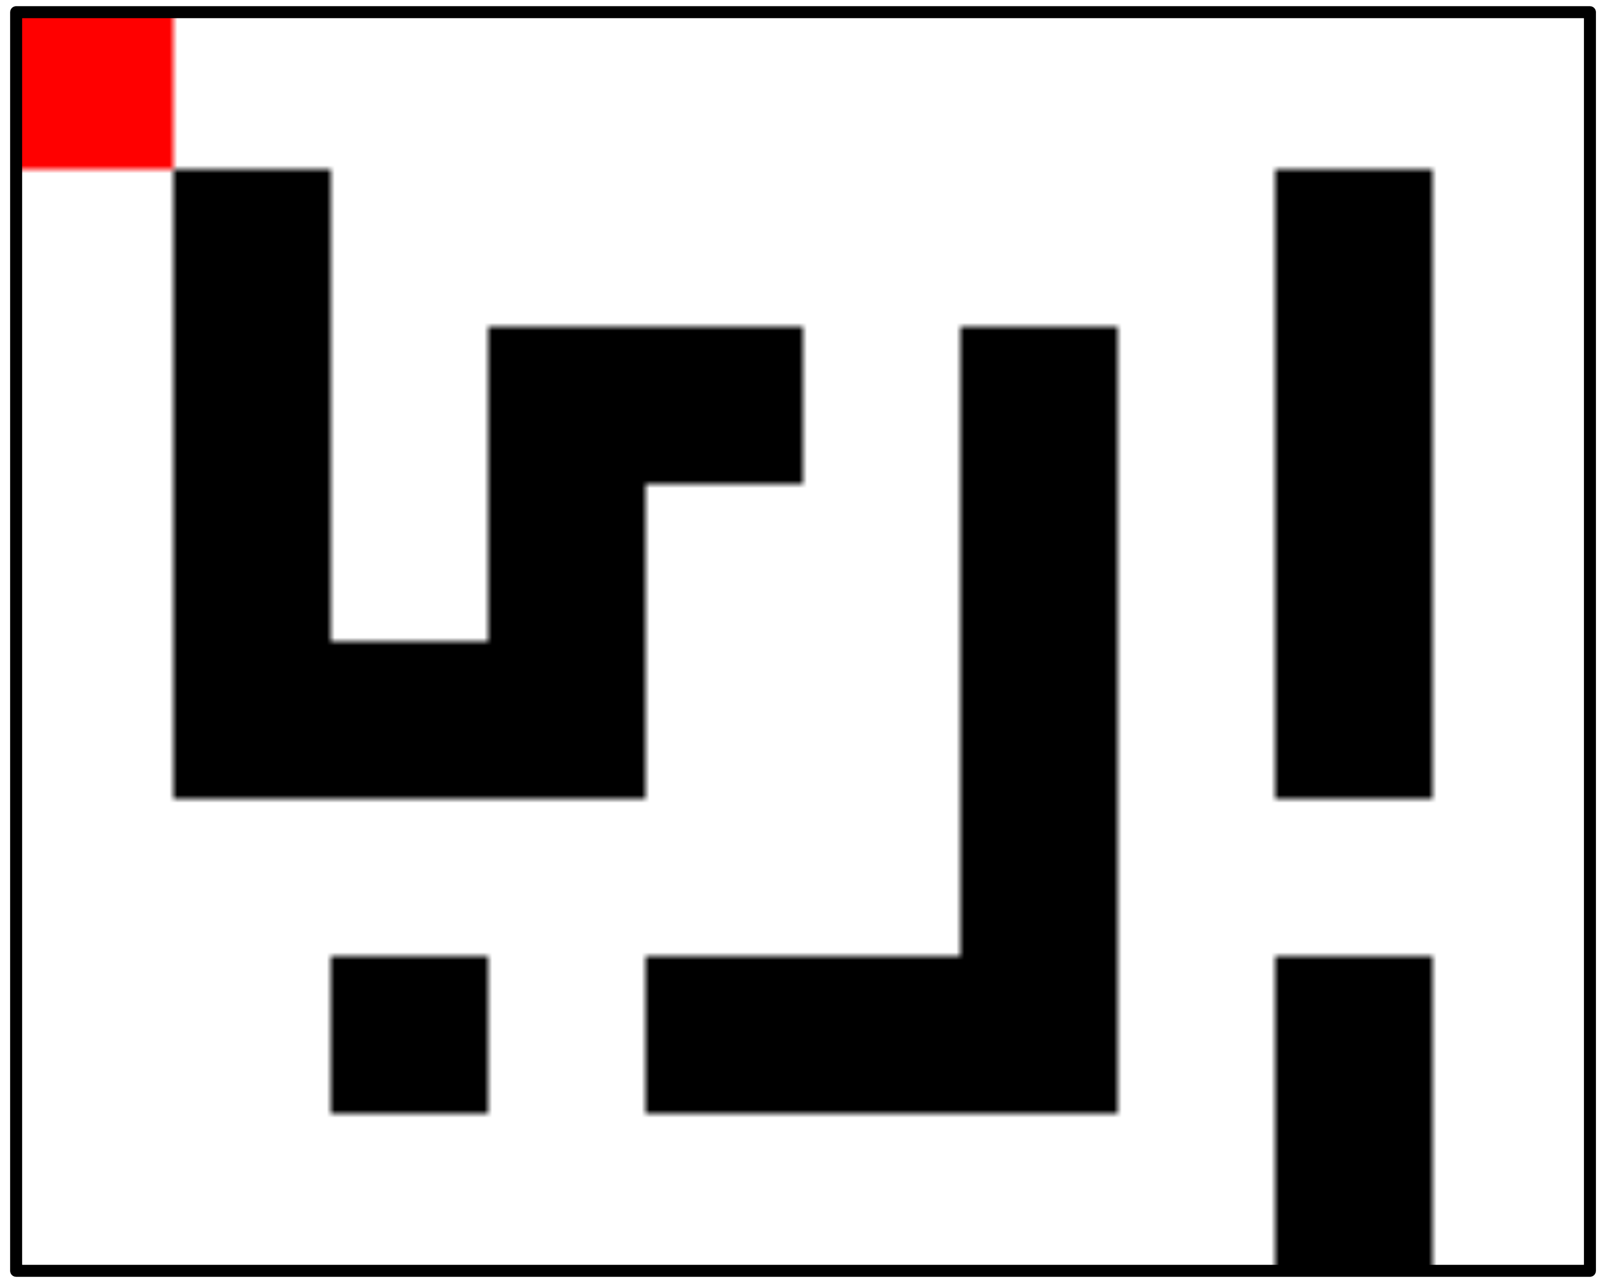

Example 3

- 8x10 discrete world
  - Known map with obstacles and walls
- Robot state
  - Location in the map (no orientation)
  - Initial state is (0,0)

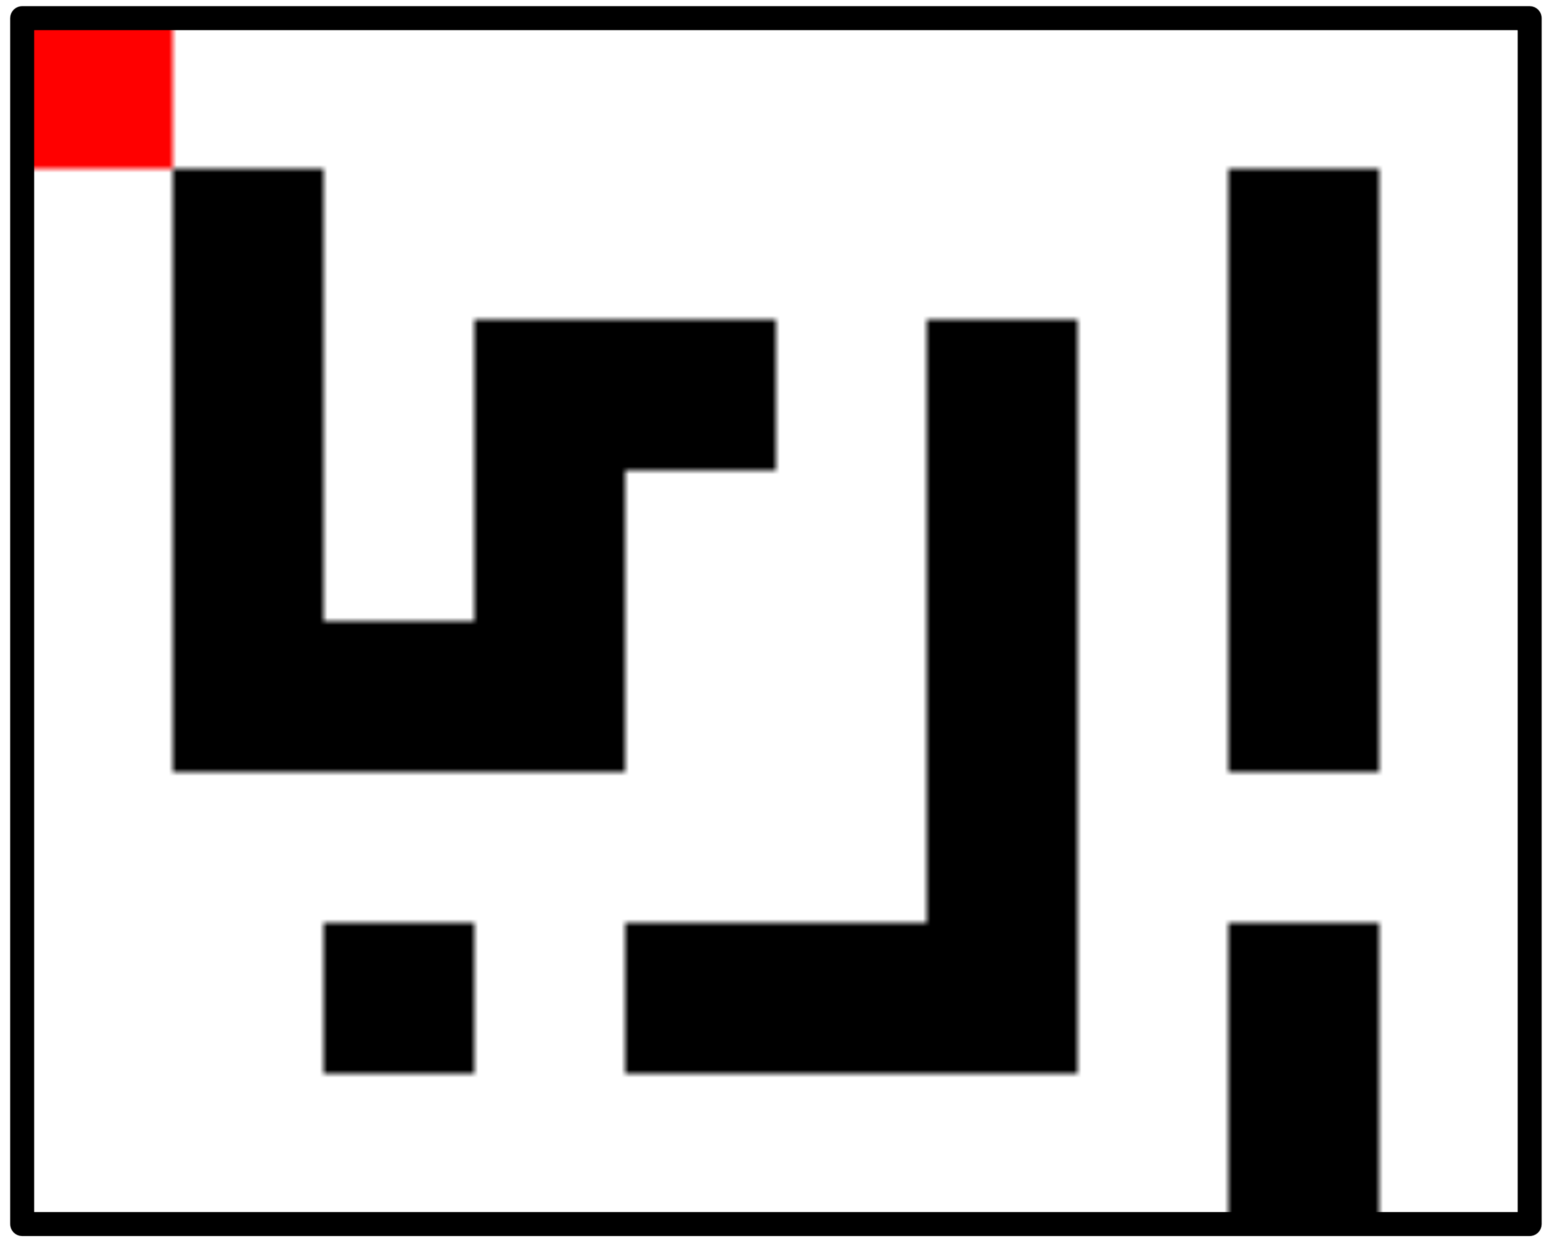

x is the set of possible locations

X is one location

Example 3

Transition model

- No matter what I tell my robot to do, it makes a random move or stays in place!

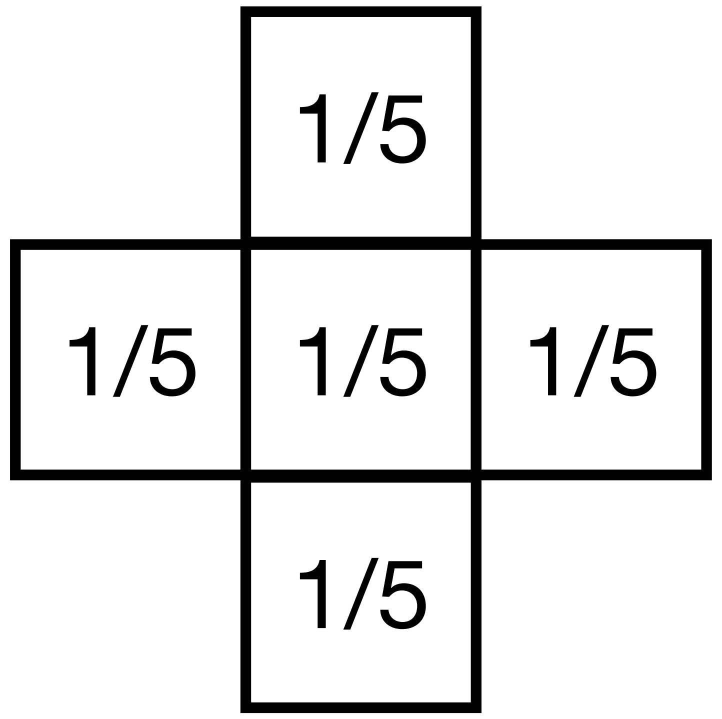

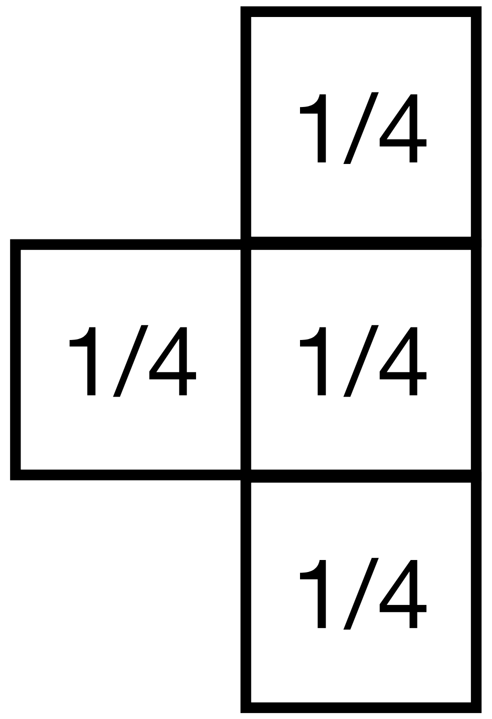

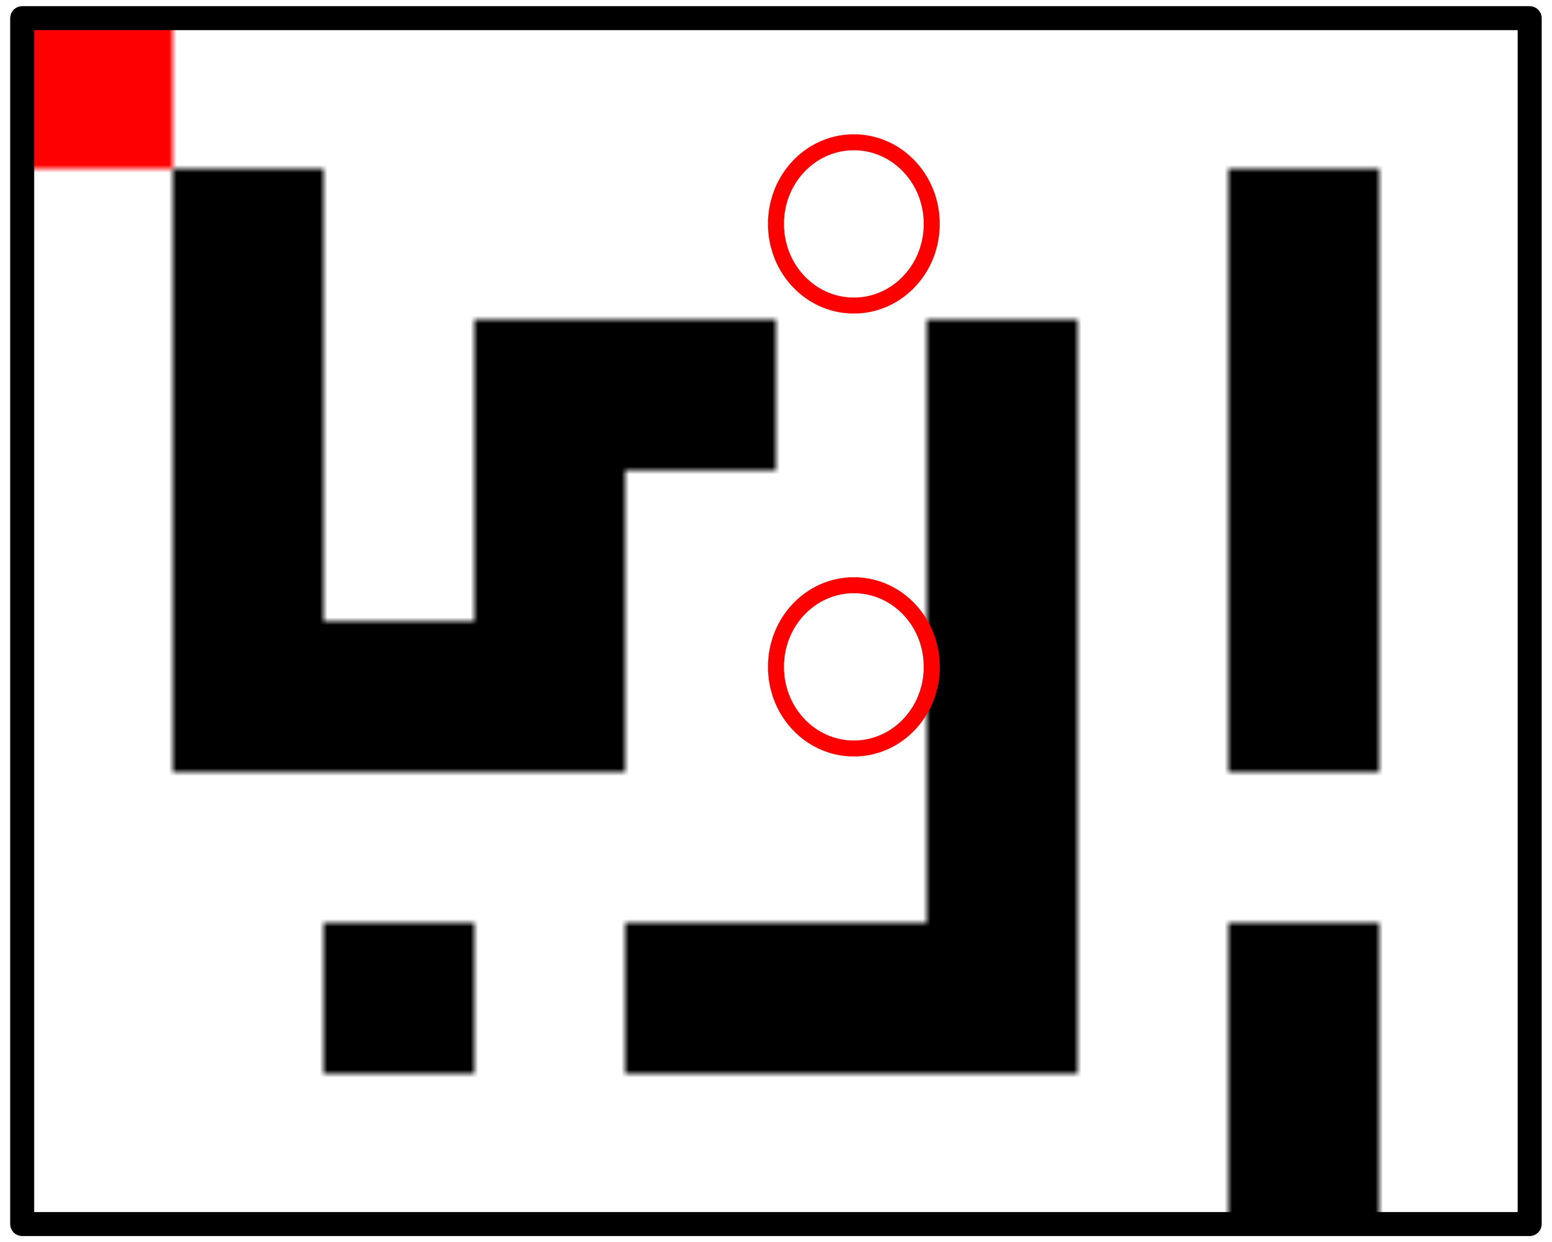

x is the set of possible locations

X is one location

Example 3

Transition model

- No matter what I tell my robot to do, it makes a random move or stays in place!
- Transition matrix, A
  - Probability to move from state *j* to state *i*

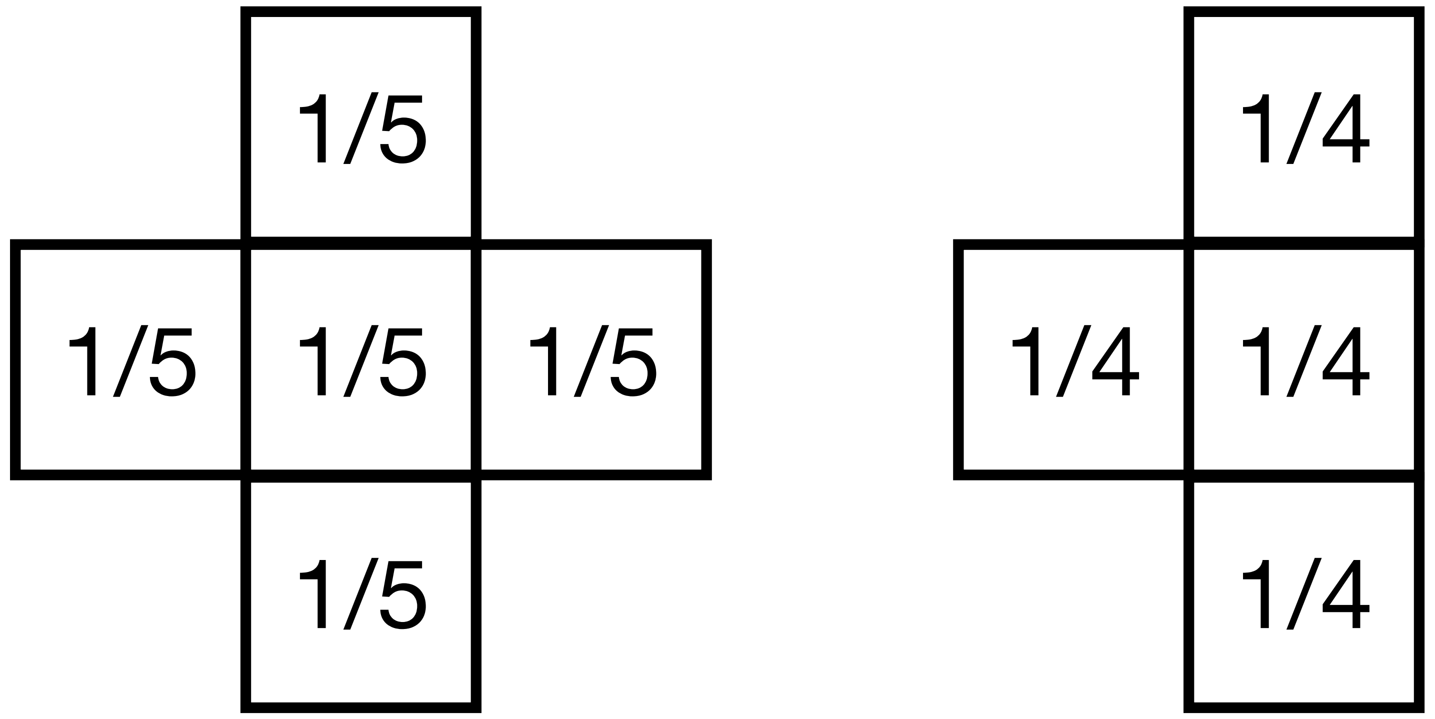

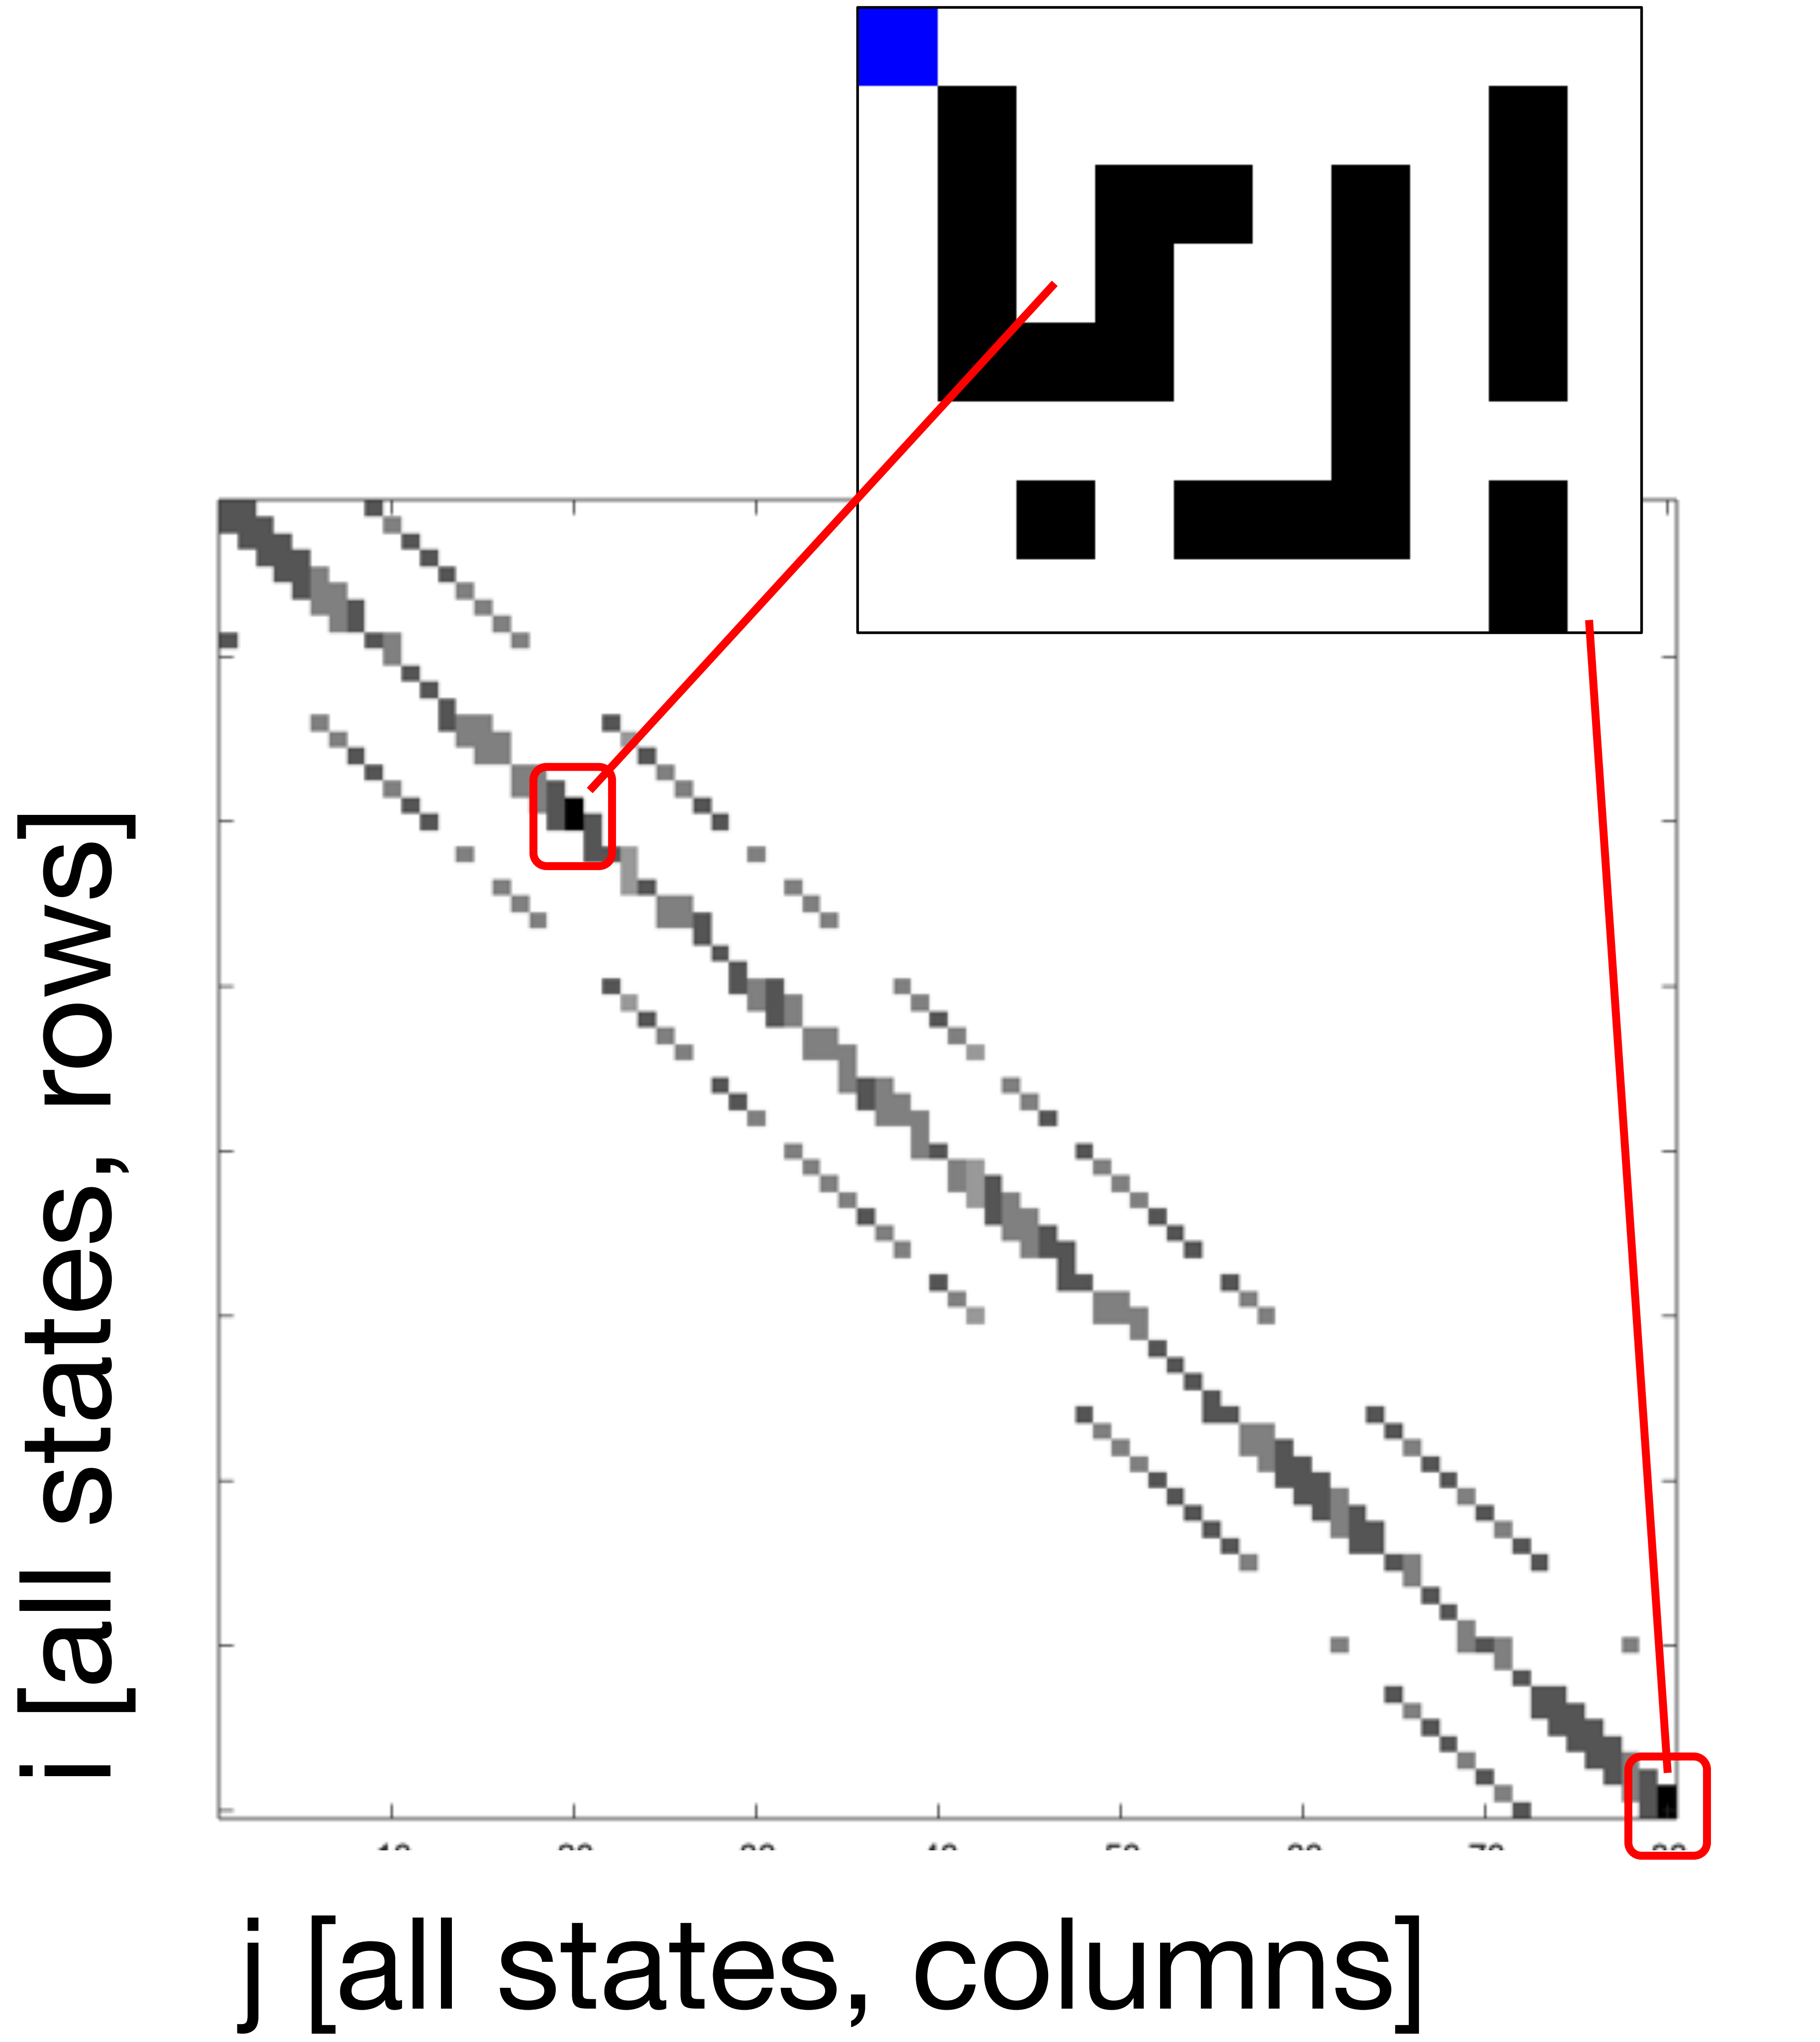

Example 3

Practical implementation

- Set up the world
- Compute the transition matrix, A
- Take actions
  - Cumulative distribution
  - `find(Mtri*A*s >= rand(),1,'first');`

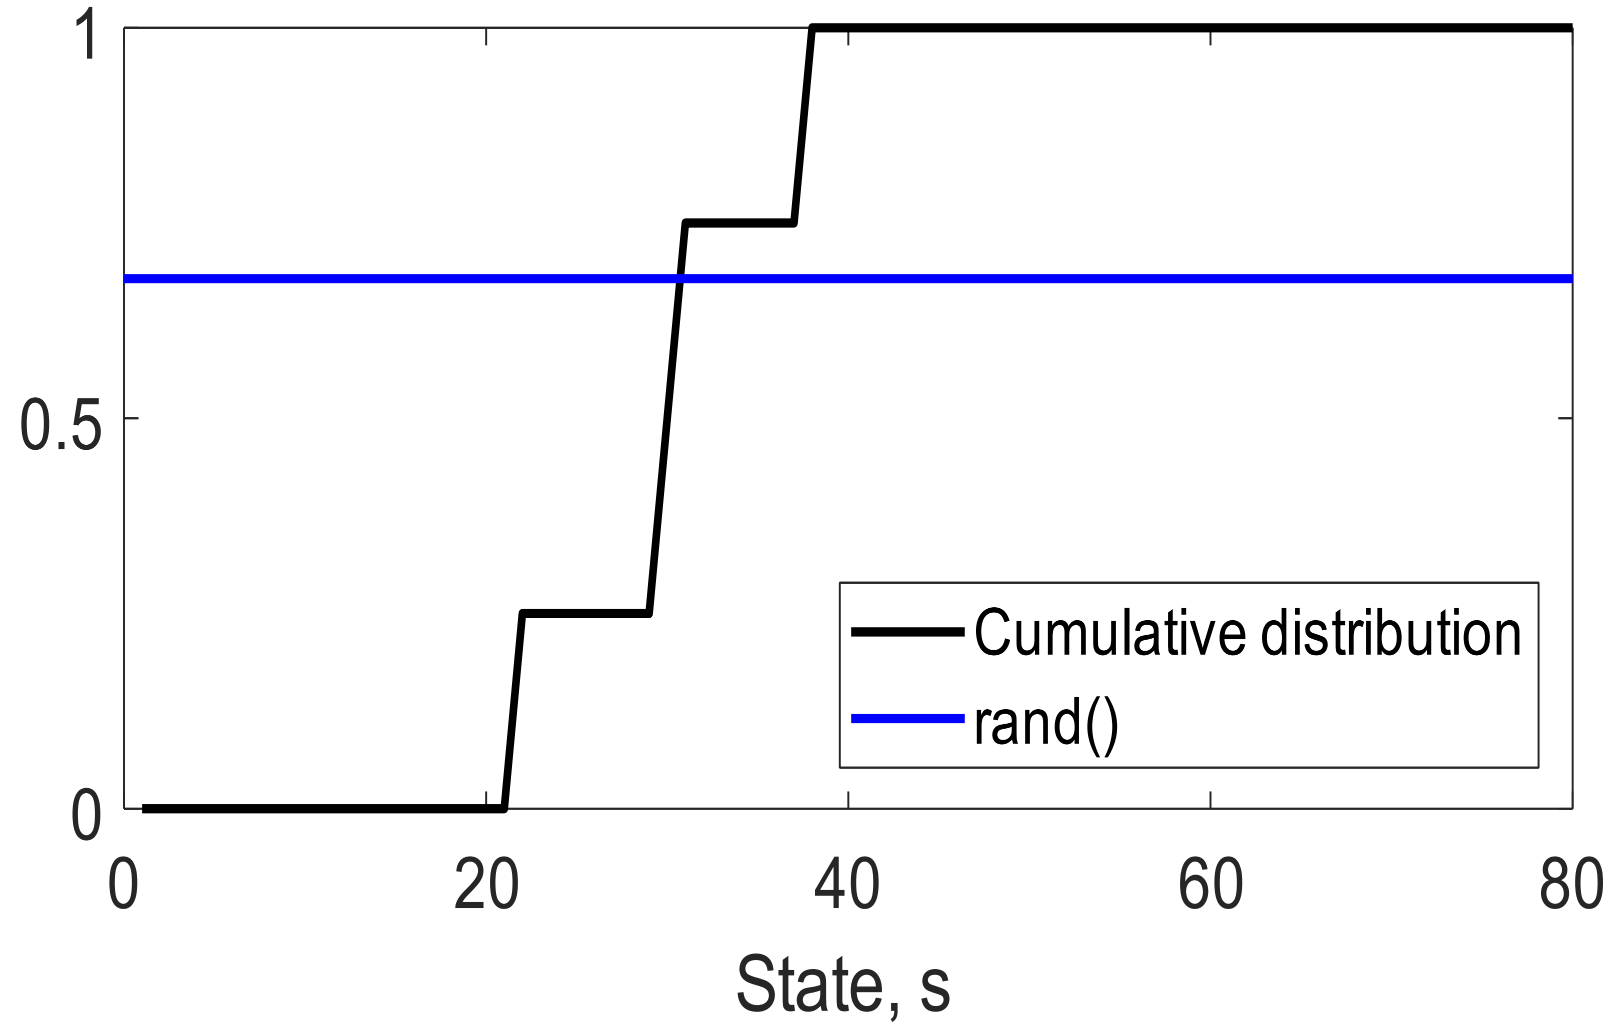

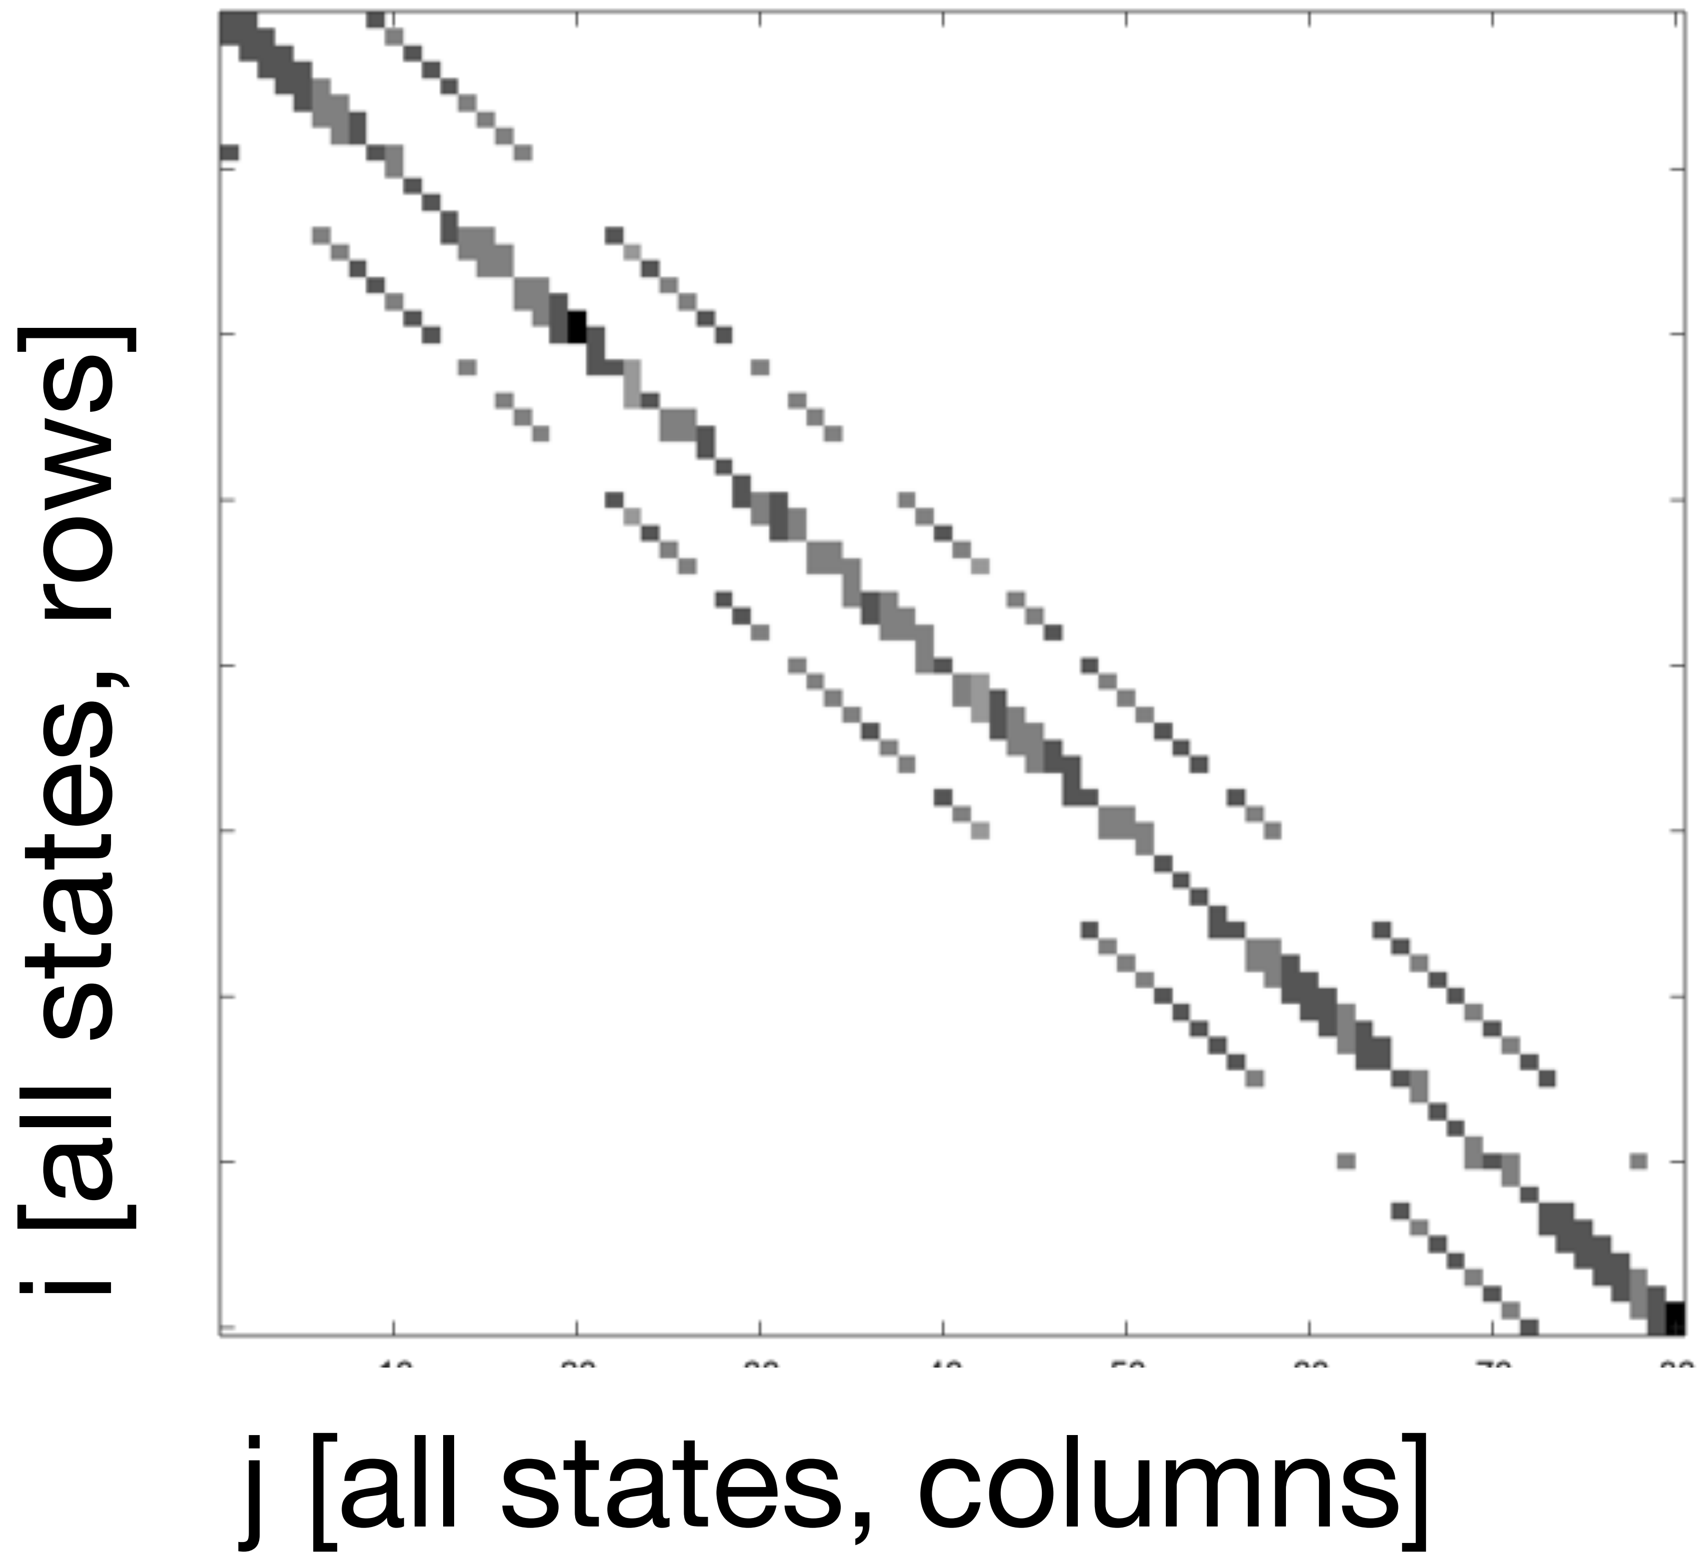

Example 3

Prediction step

1. **Prediction step** $(bel(x_{t-1}), \cancel{u_t})$:
2. **for** all x_t do
3. $\overline{bel}(x_t) = \Sigma_{x_{t-1}} \; p(x_t \mid \cancel{u_t}, x_{t-1}) \; bel(x_{t-1})$
4. end **for**

1. **Matrix implementation**
2. $\overline{bel} = A \cdot bel_{t-1}$

where A is the transition matrix (80x80) and bel is the probability distribution over all states (80x1)

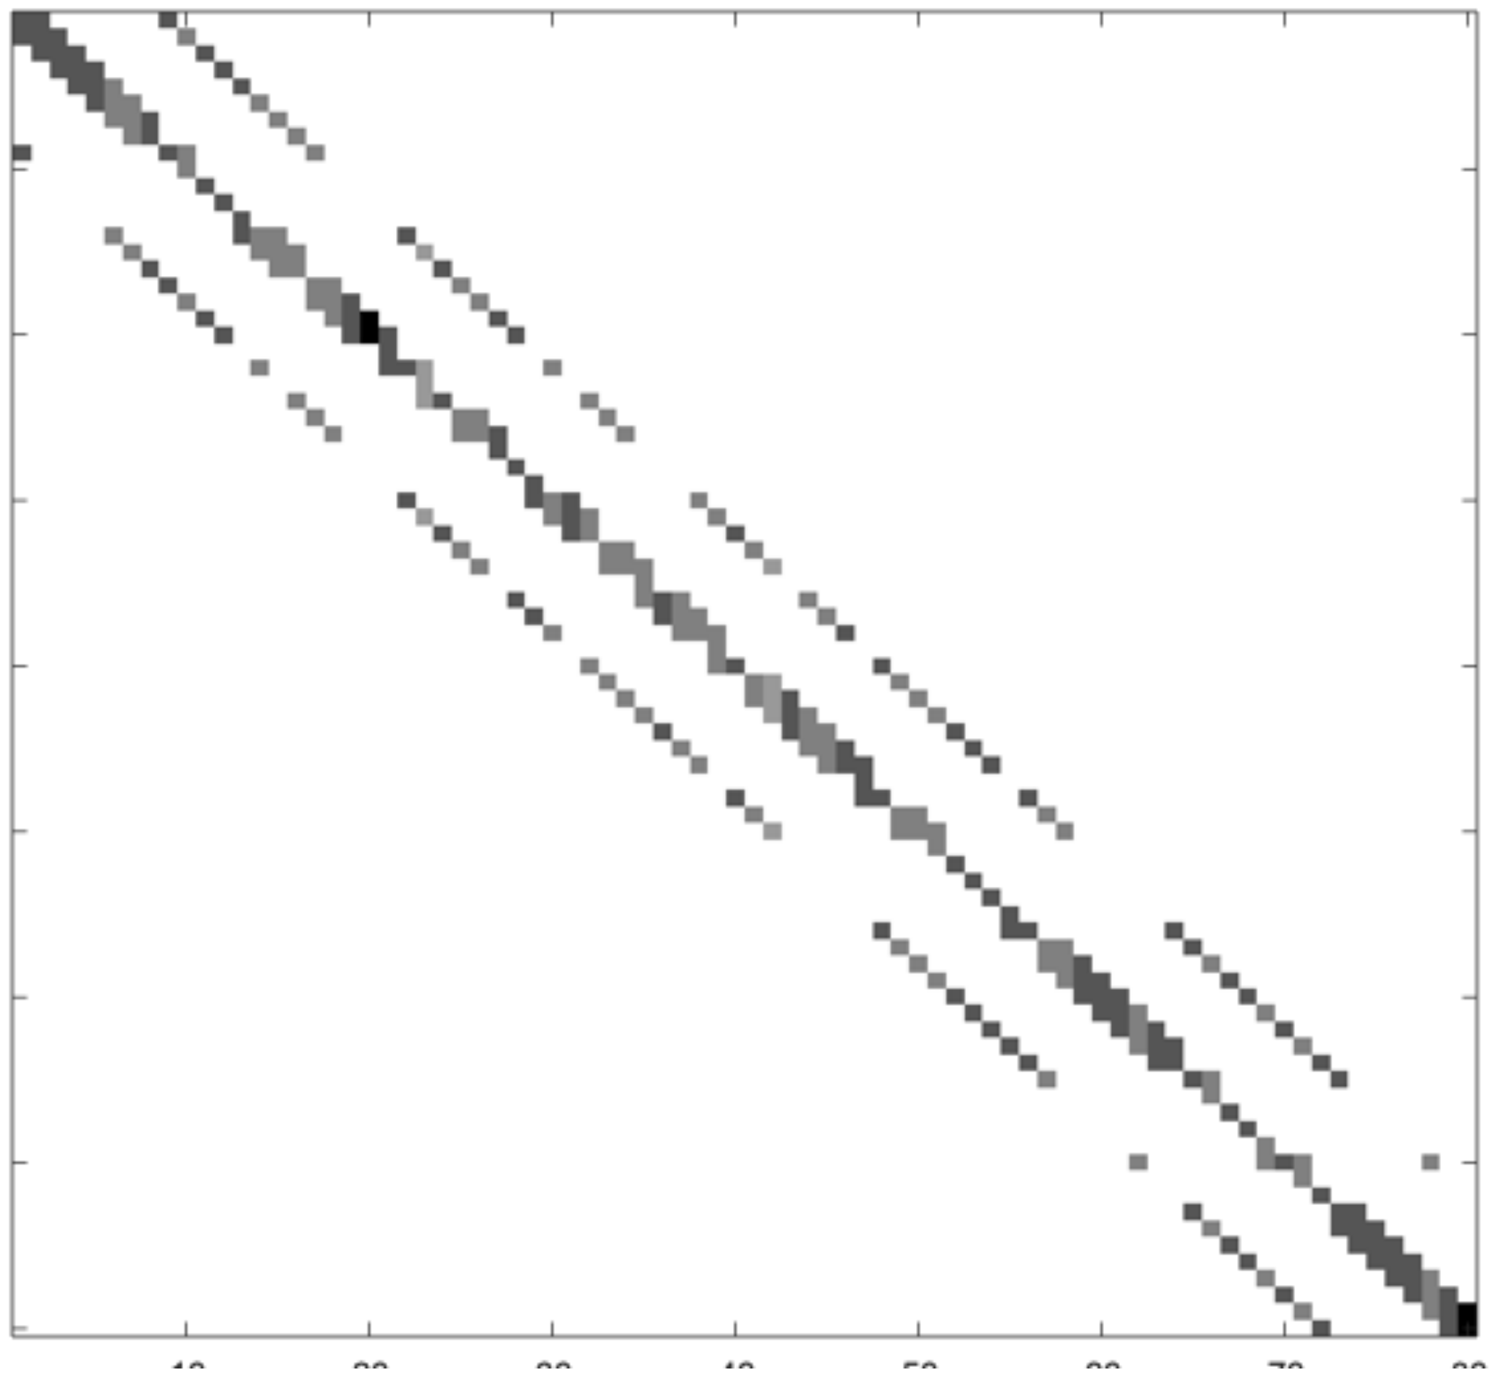

Example 3

Prediction step

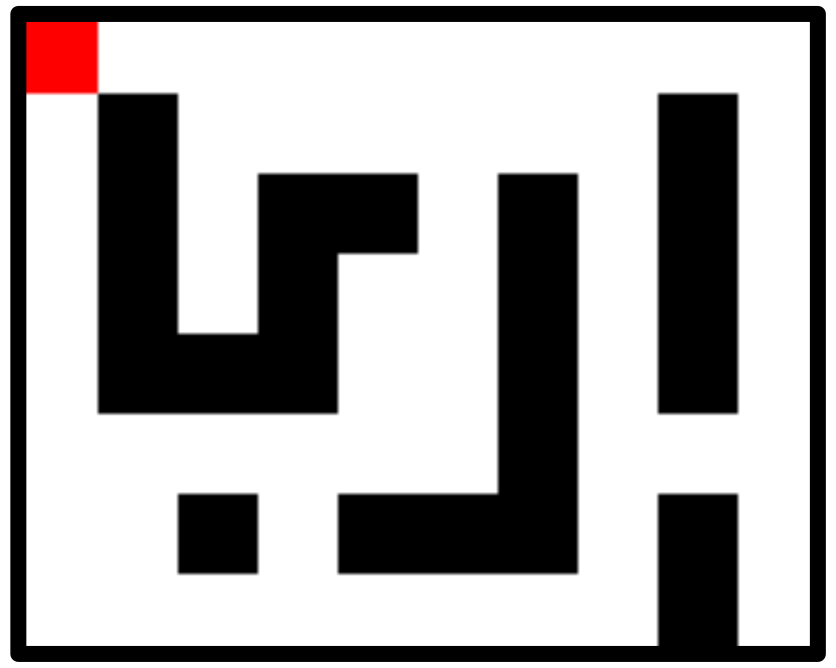

$\overline{bel}_0$

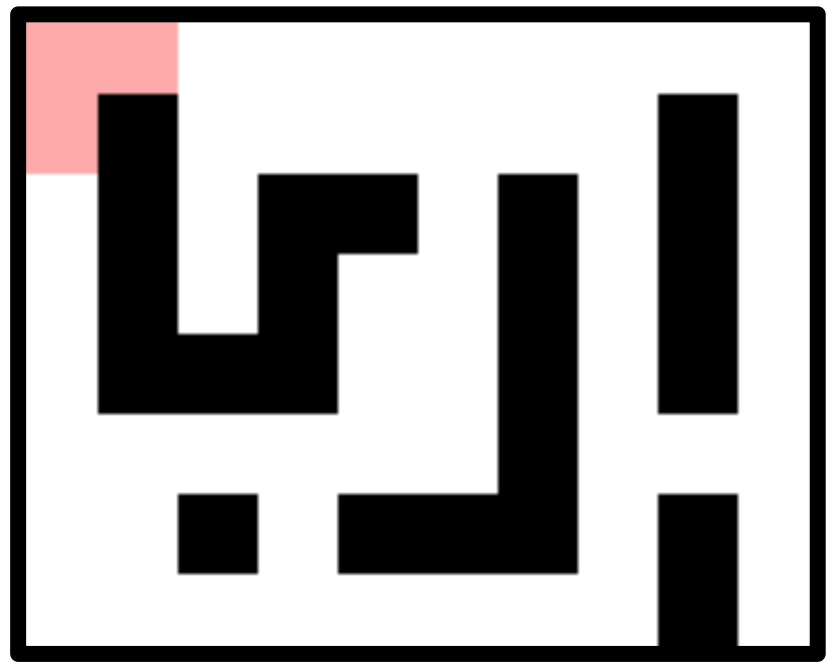

$\overline{bel}_1$

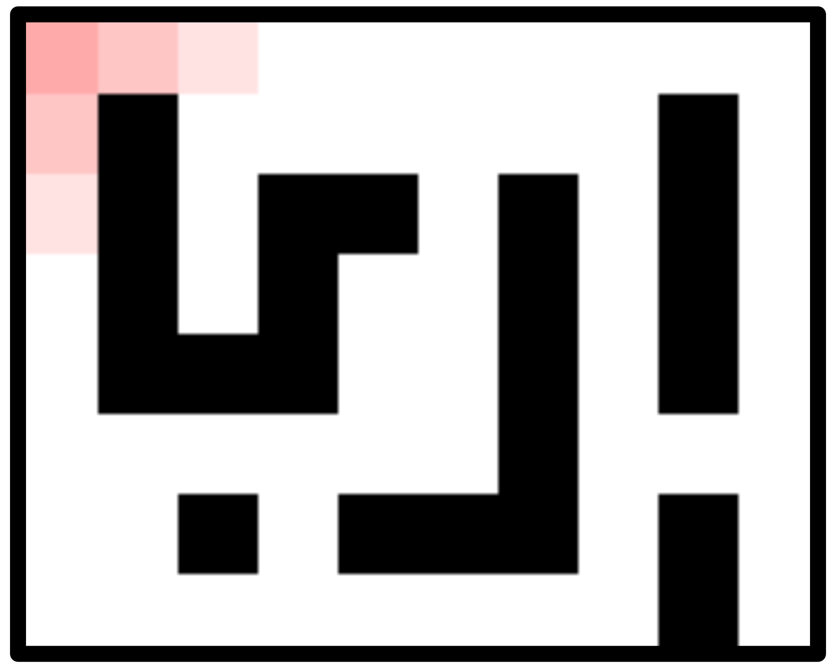

$\overline{bel}_2$

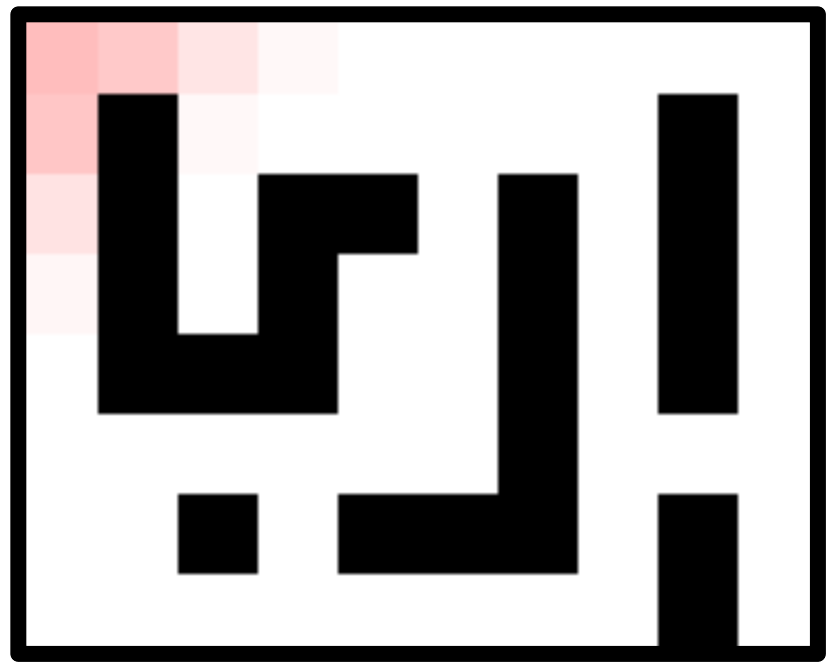

$\overline{bel}_3$

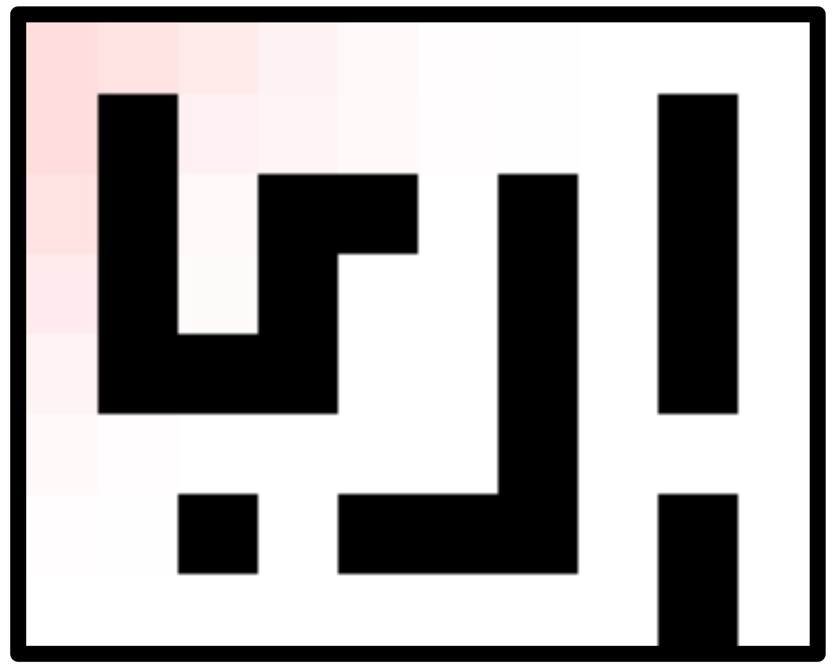

$\ldots\overline{bel}_{10}$

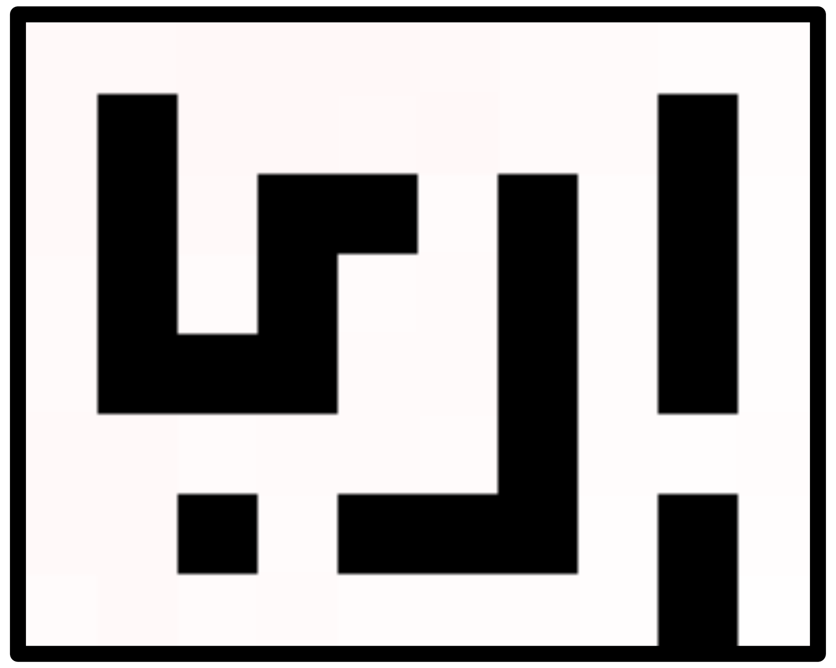

$\ldots\overline{bel}_{100}$

Example 3

Observations

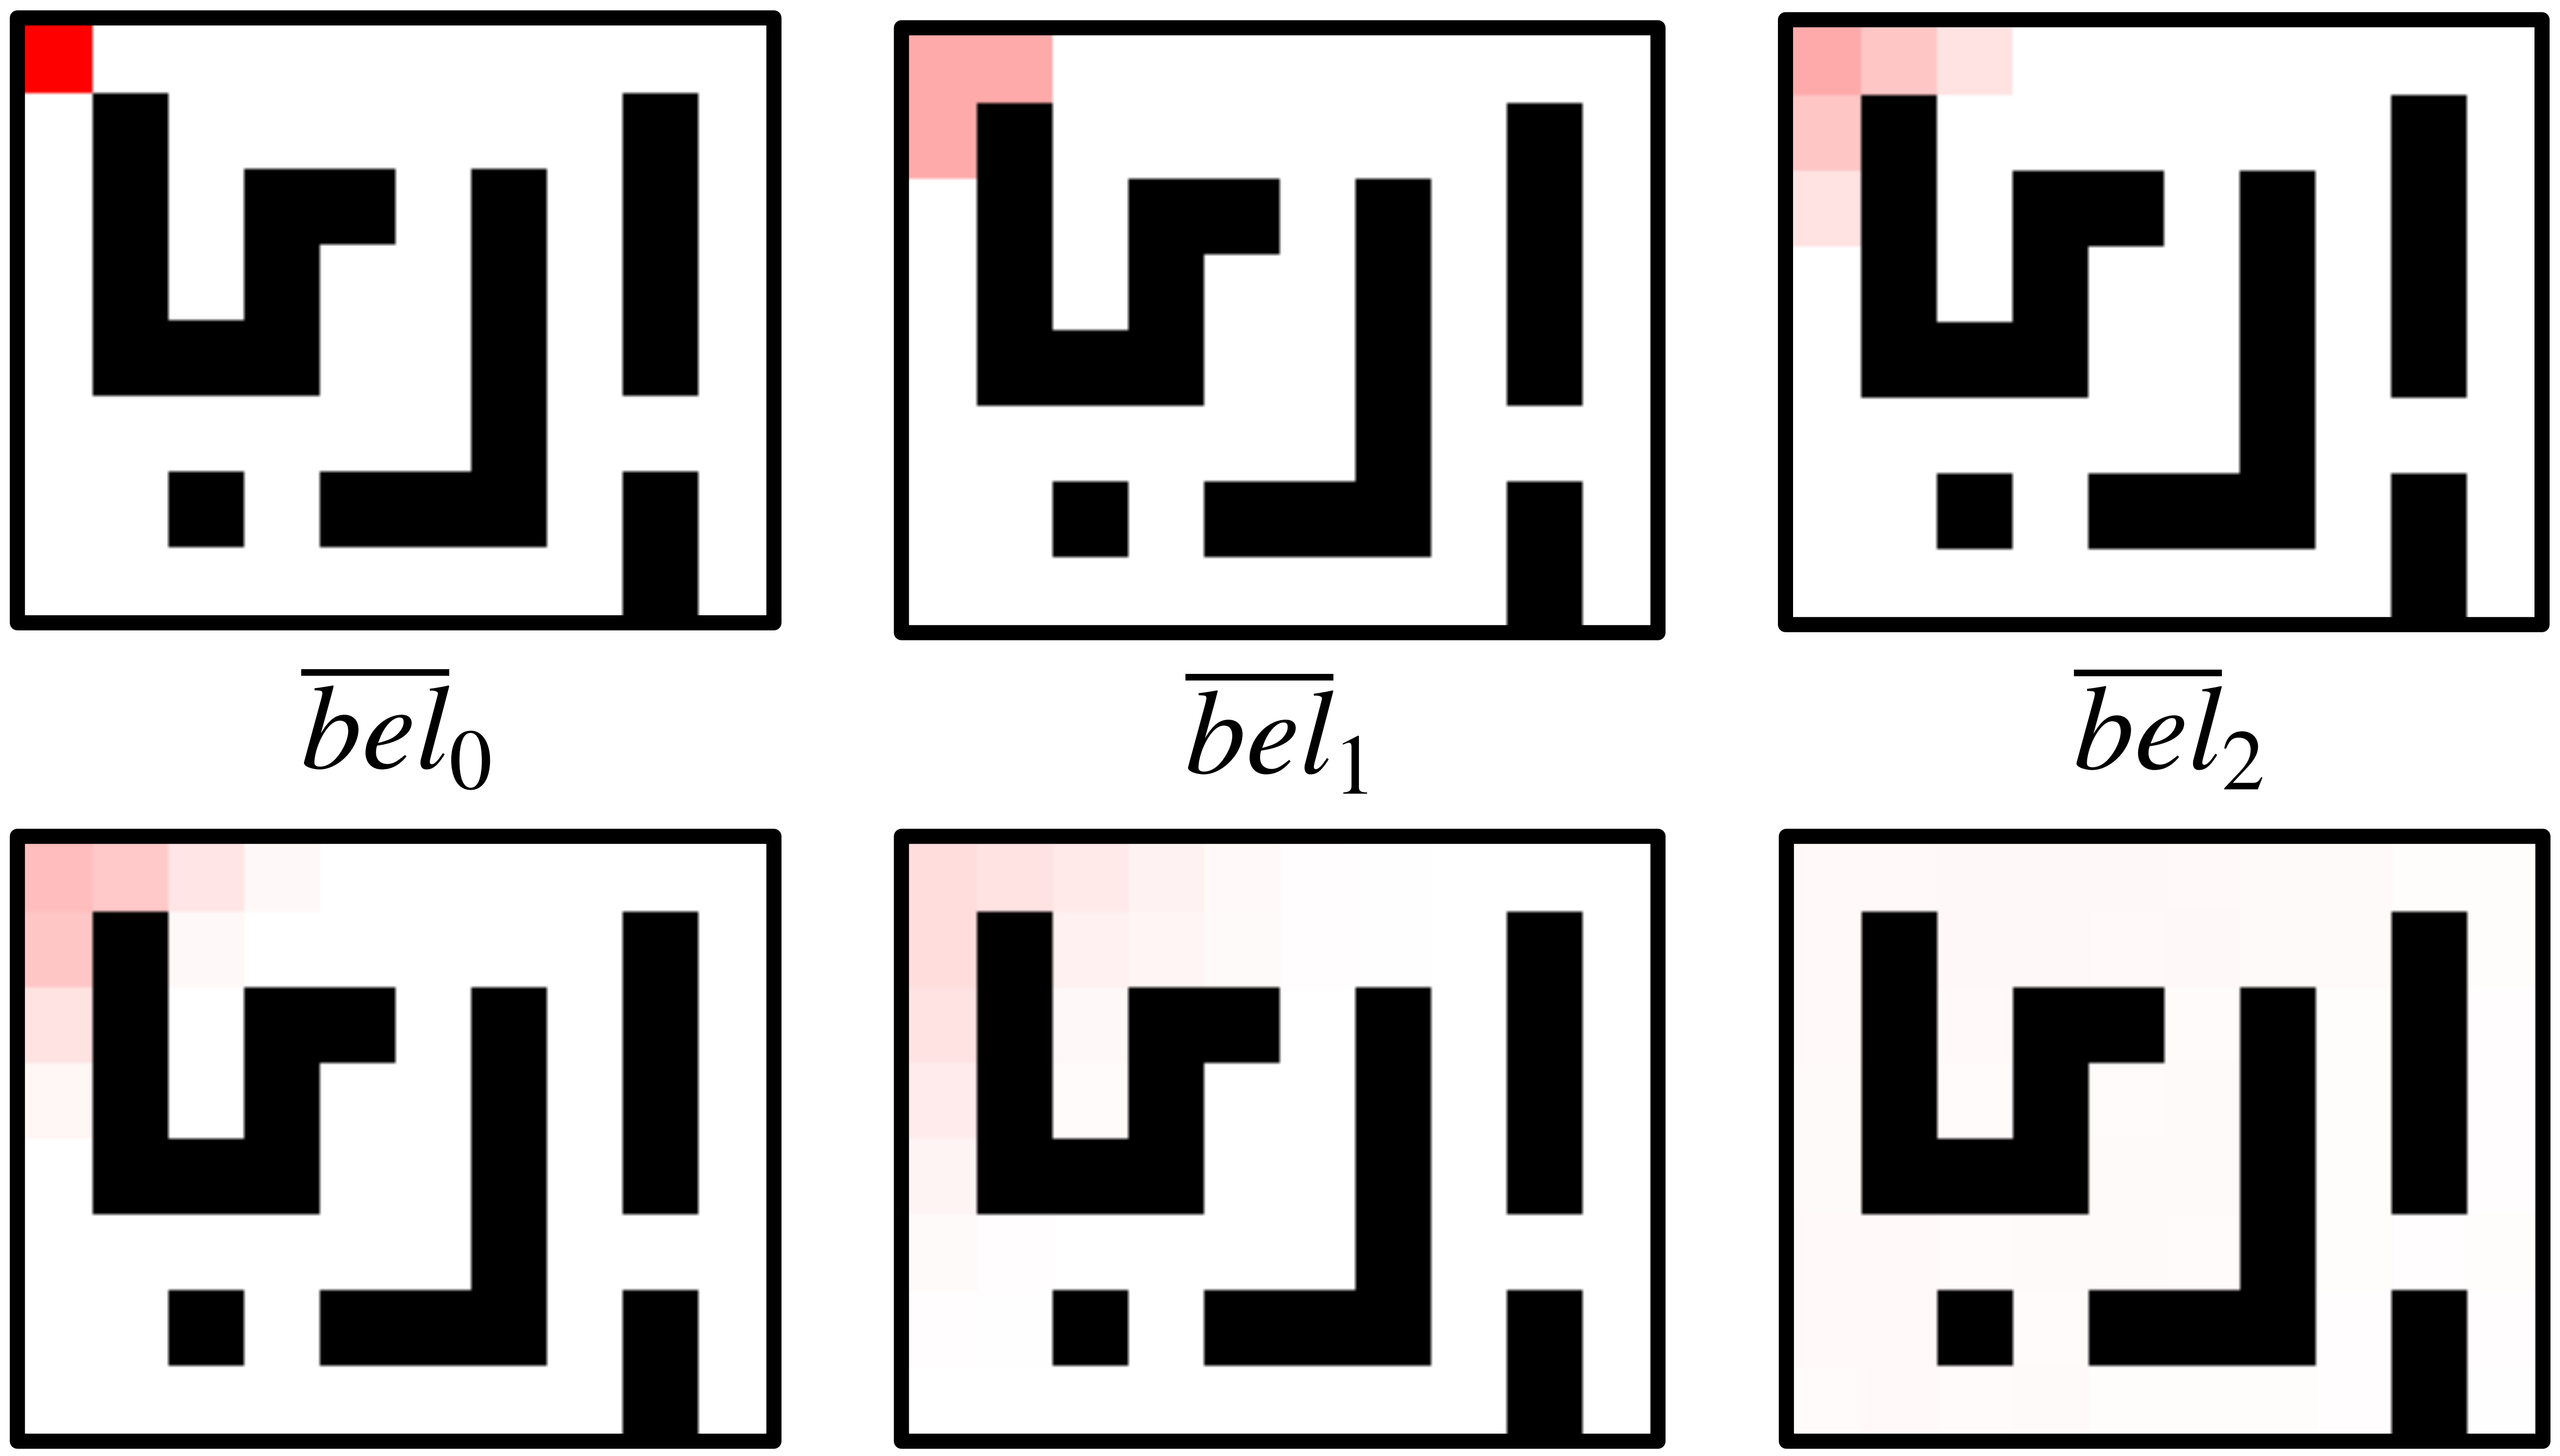

$\overline{bel}_0$ $\overline{bel}_1$ $\overline{bel}_2$

$\overline{bel}_3$ $\ldots\overline{bel}_{10}$ $\ldots\overline{bel}_{100}$

- The robot may not know where it is, but it **does** have a physical state
- It will have observations tied to that state

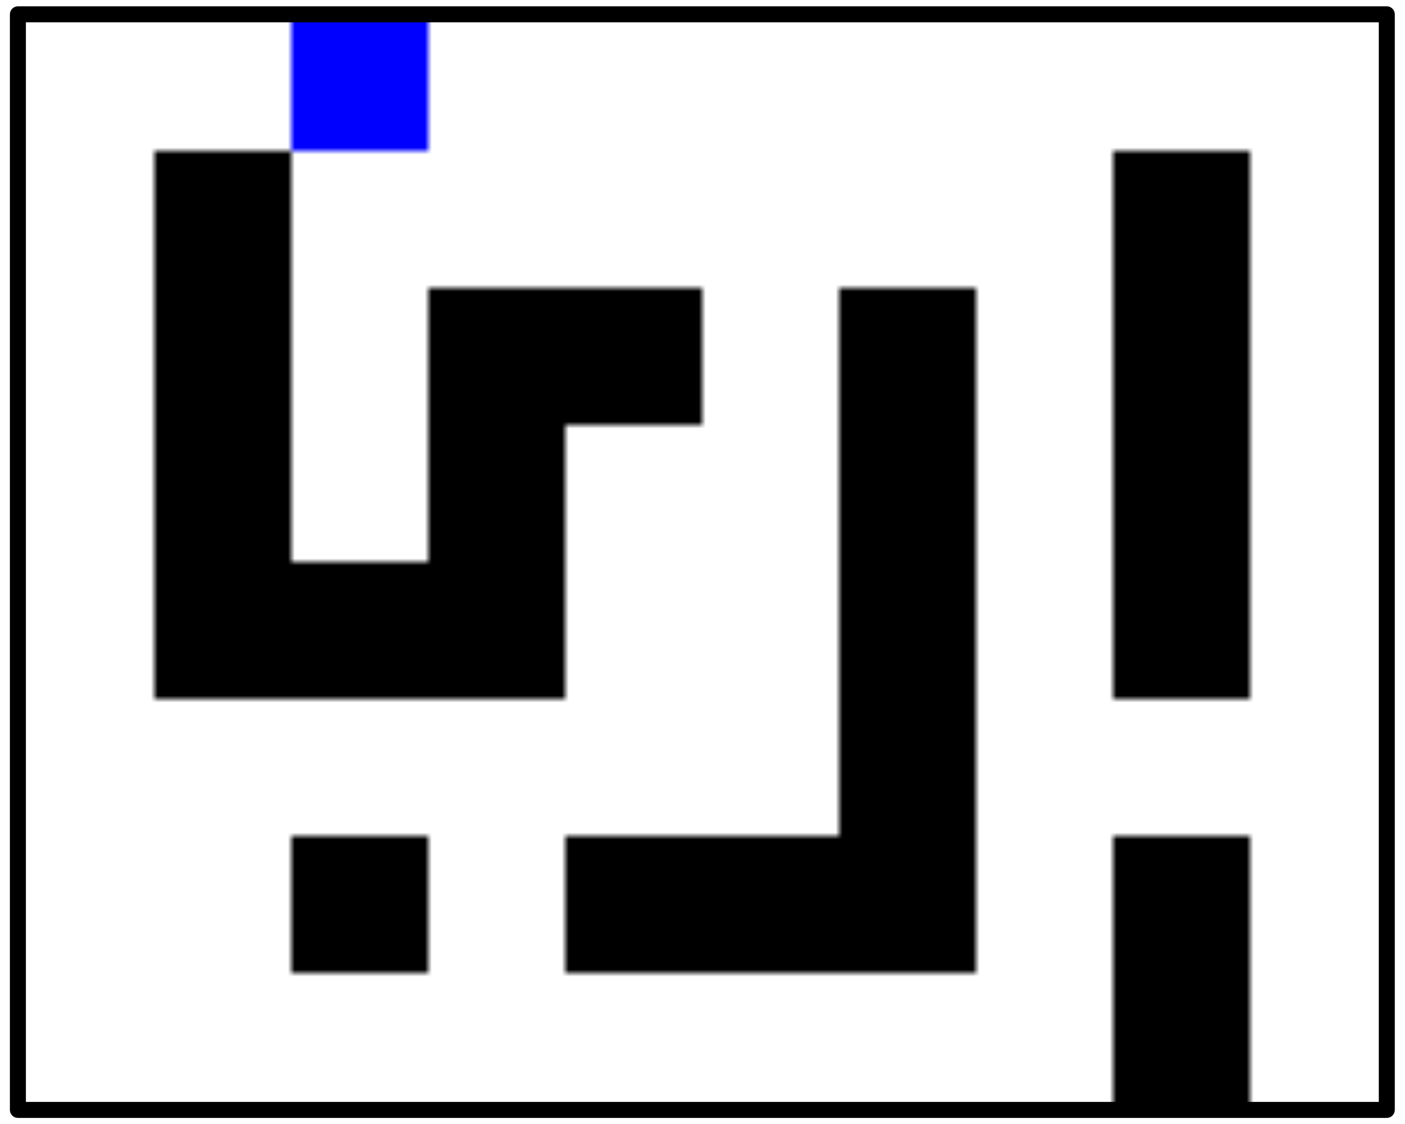

Example 3

Observation model

- In every time step, we sense each of the four neighboring cells (N, E, S, W)
- In z, each reading is independent and correct with 90% probability

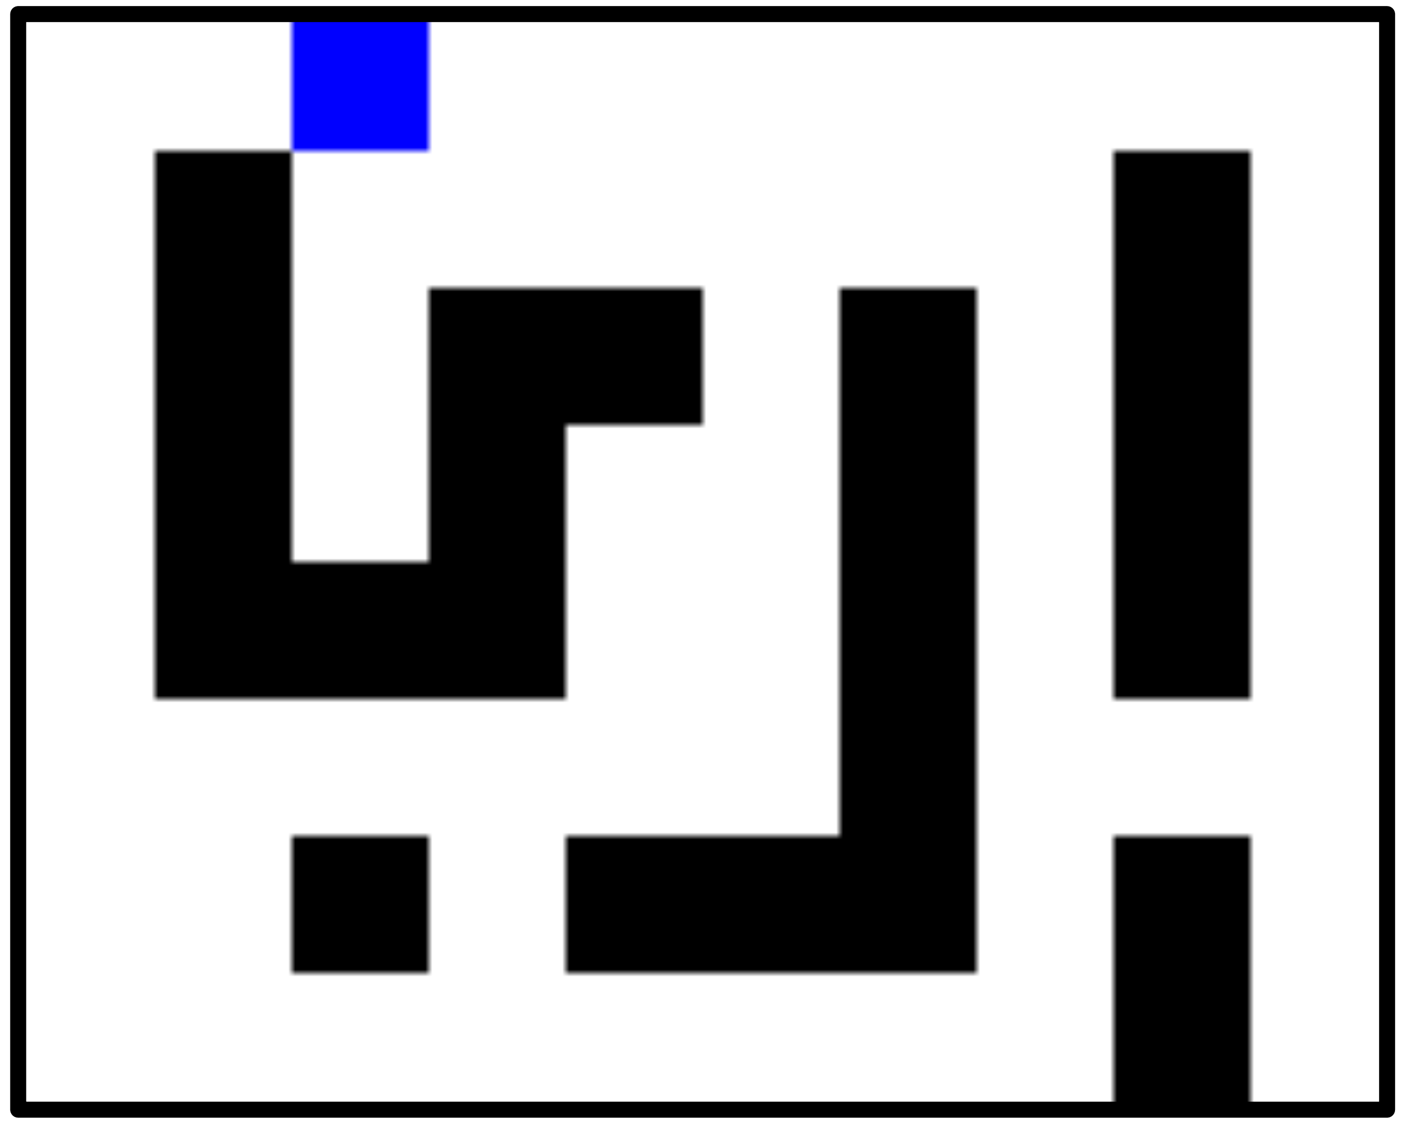

x is the set of possible locations

X is one location

z are the sensor measurements

Example 3

Observation model

- In every time step, we sense each of the four neighboring cells (N, E, S, W)
- In z, each reading is independent and correct with 90% probability

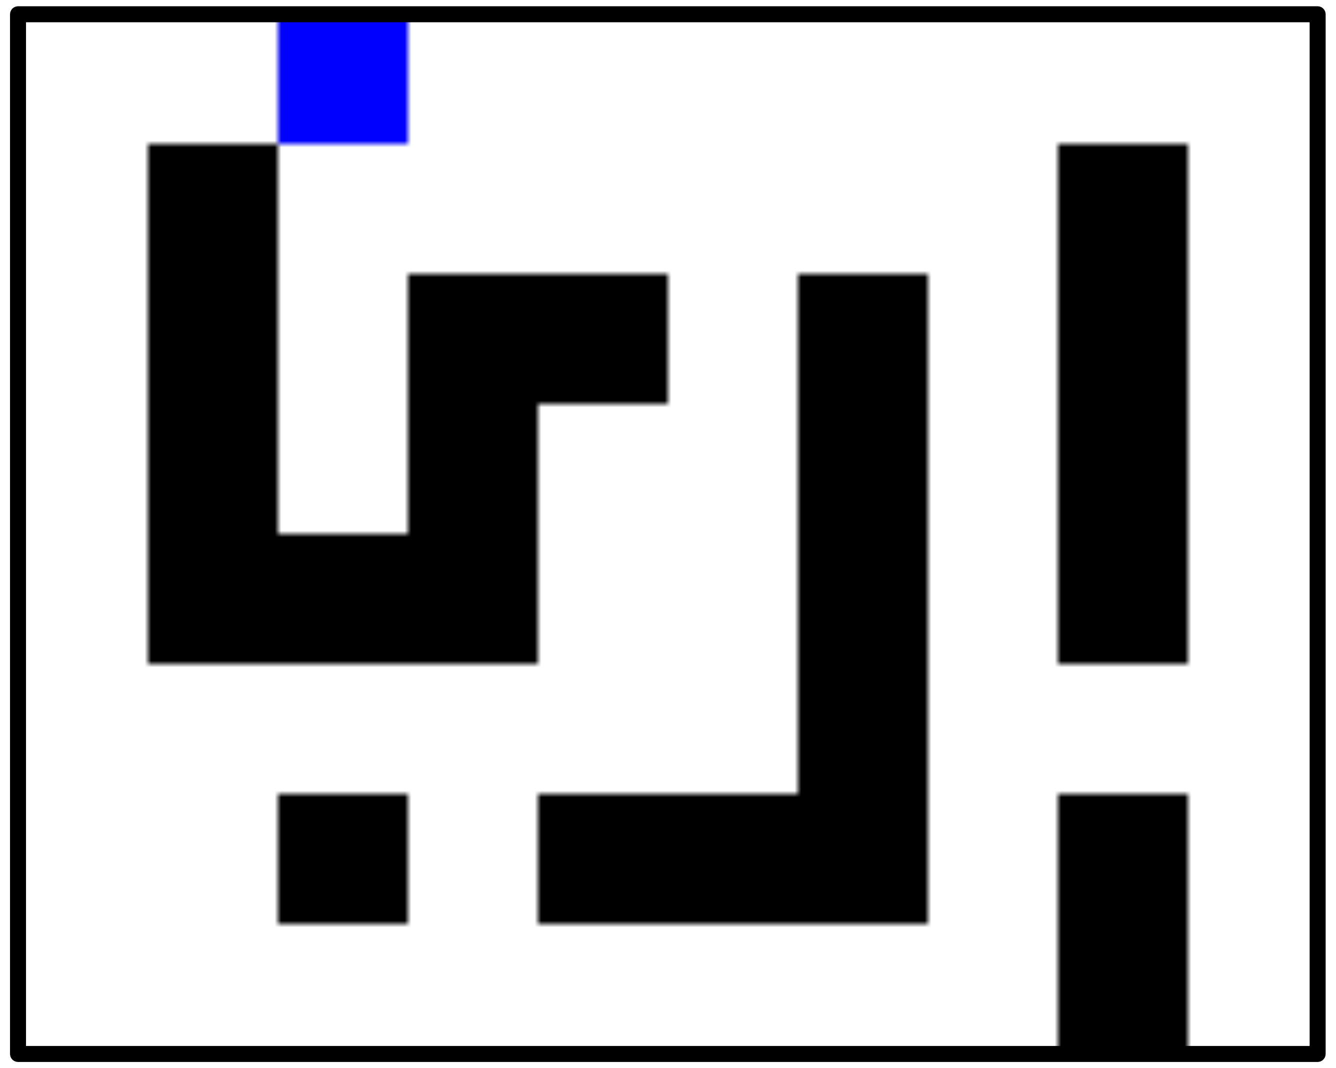

$$p(\text{no walls} \mid x) = 0.1 \cdot 0.9 \cdot 0.9 \cdot 0.9$$

Highest likelihood $\longrightarrow$

$$p(\text{N} \mid x) = 0.9 \cdot 0.9 \cdot 0.9 \cdot 0.9$$

$$p(\text{W} \mid x) = 0.1 \cdot 0.9 \cdot 0.9 \cdot 0.1$$

$$p(\text{S} \mid x) = 0.1 \cdot 0.9 \cdot 0.1 \cdot 0.9$$

$$p(\text{E} \mid x) = 0.1 \cdot 0.1 \cdot 0.9 \cdot 0.9$$

$$\cdots$$

$$p(\text{NW} \mid x) = 0.9 \cdot 0.9 \cdot 0.9 \cdot 0.1$$

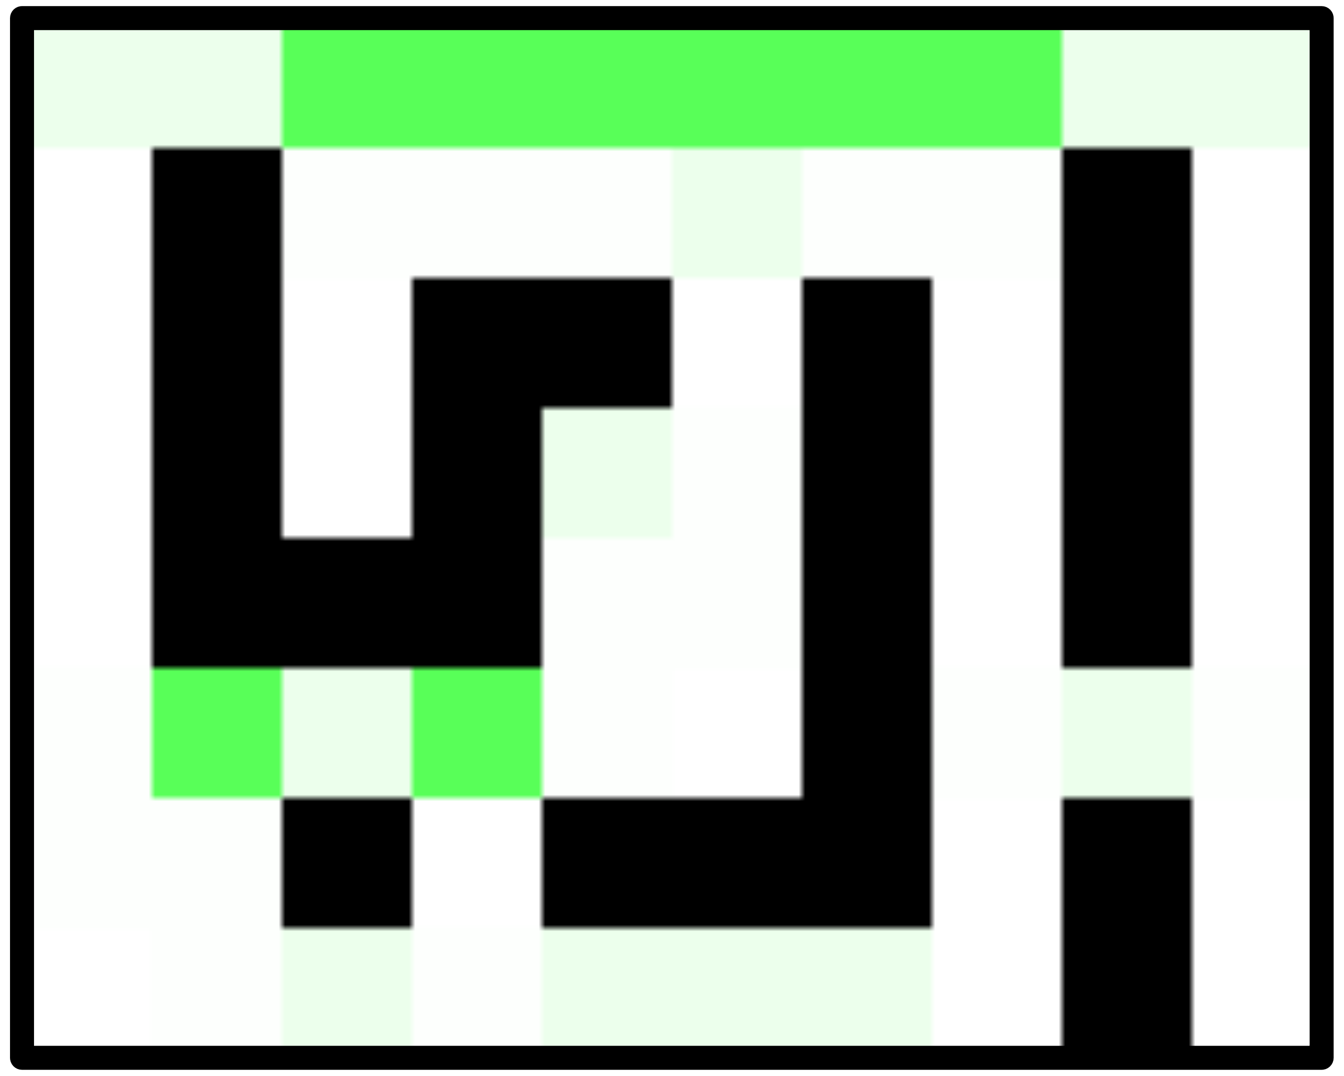

$p(z \mid X)$

How many combinations are there per state?

2^4

Example 3

Observation model

- In every time step, we sense each of the four neighboring cells (N, E, S, W)
- In z, each reading is independent and correct with 90% probability

- If all readings are correct:
  - $\Sigma |z_t - z'_{xt}| = 0$
  - $p_z(x_t) = 0.6561$
- If all readings are incorrect:
  - $\Sigma |z_t - z'_{xt}| = 4$
  - $p_z(x_t) = 0.0001$

1. **Algorithm Bayes_Filter** $(bel(x_{t-1}), u_t, z_t)$:
2. **for** all x_t do
3. $\overline{bel}(x_t) = \Sigma_{x_{t-1}} p(x_t | u_t, x_{t-1})\, bel(x_{t-1})$
4. $bel(x_t) = \eta\, p(z_t | x_t)\, \overline{bel}(x_t)$
5. end **for**
6. **return** $bel(x_t)$

1. **Likelihood of Observations,** p_{zX}**:**
2. **for** all x_t do
3. $p_{zX}(x_t) = 0.9^{4-\Sigma|z_t - z'_{xt}|}\, 0.1^{\Sigma|z_t - z'_{xt}|}$
4. end **for**

where p_{zX} **is a vector (80x1)**

Example 3

Observation model

- In every time step, we sense each of the four neighboring cells (N, E, S, W)
- In z, each reading is independent and correct with 90% probability

1. **Algorithm Bayes_Filter** $(bel(x_{t-1}), u_t, z_t)$:
2. **for** all x_t do
3. $\overline{bel}(x_t) = \Sigma_{x_{t-1}}\ p(x_t \mid u_t, x_{t-1})\ bel(x_{t-1})$
4. $\boxed{bel(x_t) = \eta\ p(z_t \mid x_t)\ \overline{bel}(x_t)}$
5. end **for**
6. **return** $bel(x_t)$

1. **Compute new belief:**
2. $bel_t = \dfrac{p_{zX}\overline{bel}}{\Sigma(p_{zX}\overline{bel})}$

where $\overline{bel}$ is a vector (80x1)
and p_{zX} is a vector (80x1)

Example 3

1. **Algorithm Bayes_Filter** (bel_{t-1}, z_t) :
2. $\overline{bel} = A\ bel_{t-1}$
3. **for** all x_t do
4. $p_{zX}(x_t) = 0.9^{4-\Sigma|z_t - z'_{xt}|}\ 0.1^{\Sigma|z_t - z'_{xt}|}$
5. end **for**
6. $bel_t = \dfrac{p_{zX}\ \overline{bel}}{\Sigma(p_{zX}\ \overline{bel})}$
7. **return** bel_t

Only do this for states with a belief > threshold

Cache and look up

Example 3

Sensed W wall

Sensed N wall

No clue!

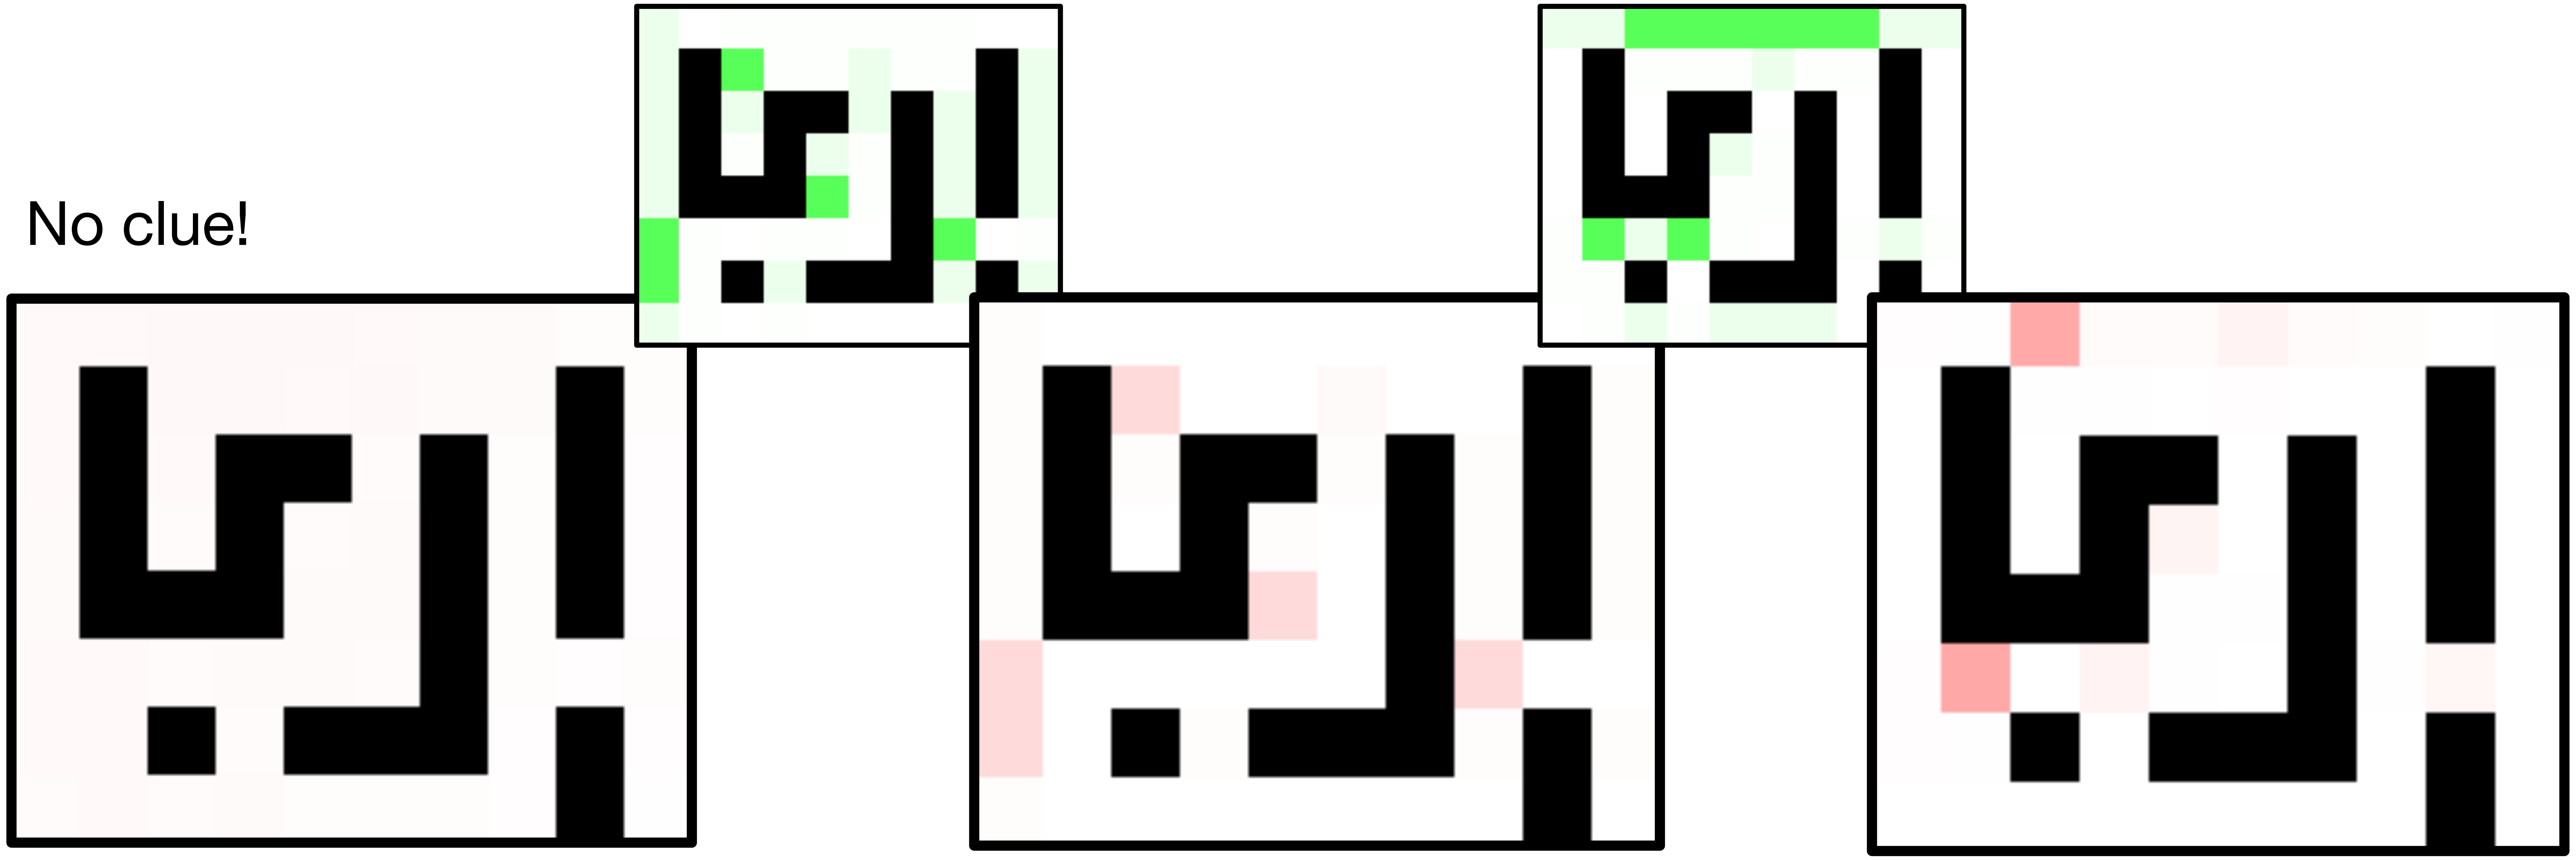

In two steps, we homed in on where we are!

Can we do better?

- Improved transition model
- Deliberately move in directions that give you more information

Today's examples

- Example 1: robot in the 1D world
    - Important to have some belief in all states
- Example 2: Bayes with beans
    - Important to normalize
- Example 3: (x,y) robot in a grid world
    - Important to improve computational efficiency
        - Matrices
        - Pre-cache

1. **Algorithm Bayes_Filter** $(bel(x_{t-1}), u_t, z_t)$:
2. **for** all x_t do
3. $\overline{bel}(x_t) = \Sigma_{x_{t-1}}\ p(x_t \mid u_t, x_{t-1})\ bel(x_{t-1})$
4. $bel(x_t) = \eta\ p(z_t \mid x_t)\ \overline{bel}(x_t)$
5. end **for**
6. **return** $bel(x_t)$

Summary of Bayes Filter

- Use temporal consistency between observations that are poor estimates individually
- Localization can work with…
  - completely random motion
  - noisy sensors
  - Remember to..
    - remain probabilistic
    - normalize
    - improve efficiency

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

    for all x_t do

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

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

    end **for**

return $bel(x_t)$