Prof. Kirstin Hagelskjær Petersen
kirstin@cornell.edu

Fast Robots

Bayes Filter
Examples

Bayes Filter II

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

1. for all x_t do
2. $\overline{bel}(x_t) = \sum_{x_{t-1}} p(x_t|u_t, x_{t-1})\ bel(x_{t-1})$ [Prediction Step]
3. $bel(x_t) = \eta\, p(z_t|x_t)\ \overline{bel}(x_t)$ [Update/Measurement Step]
4. endfor
5. return $bel(x_t)$

- Example 1
  - Robot in a 1D world
  - The importance of having some belief in all states
- Example 2
  - Bayes with beans
  - Remember to normalize!
- Example 3
  - (x,y)-robot in a grid world
  - Computational efficiency
    - Matrices
    - Pre-cache observations

Bayes Filter - Example 1

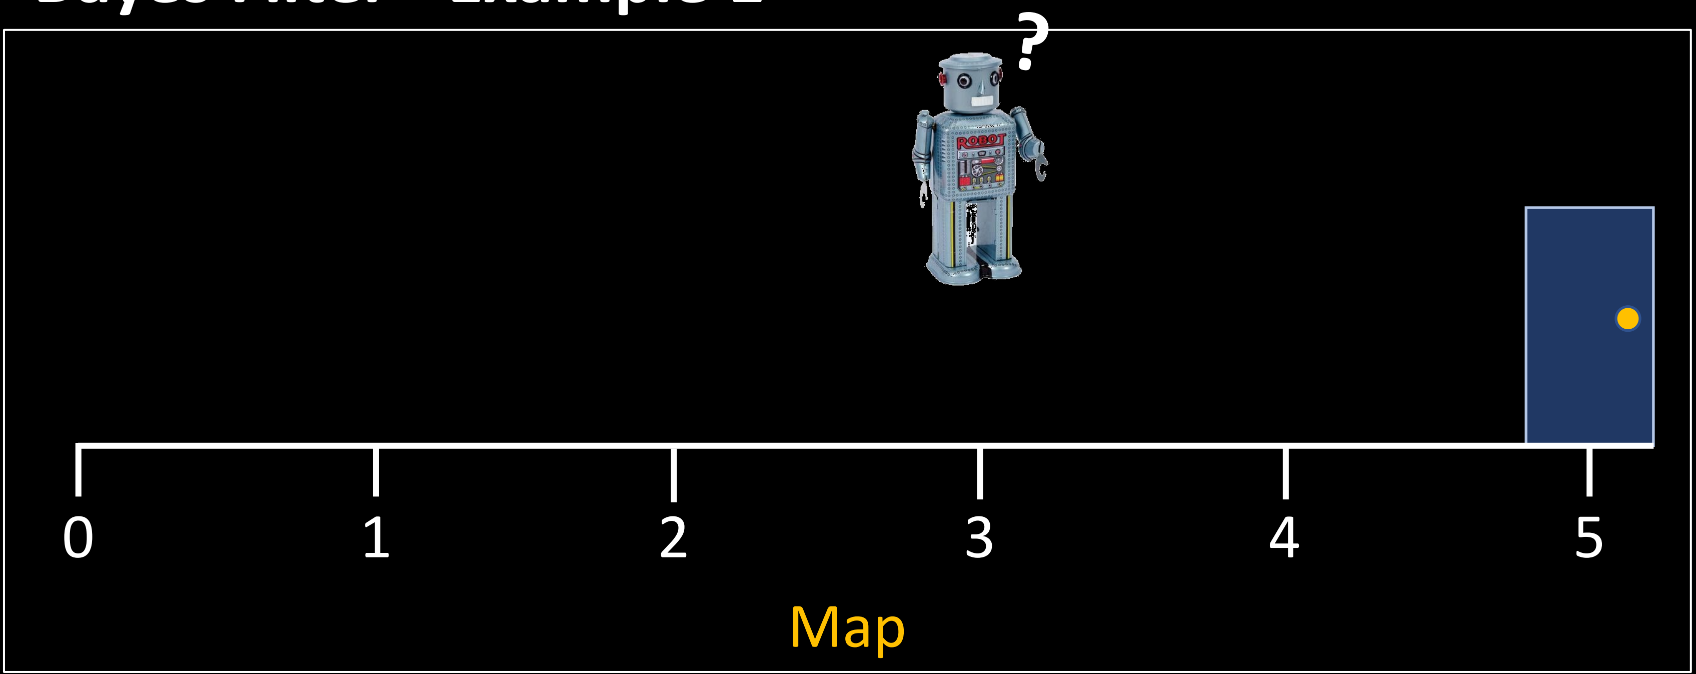

- ***What do we need to run the Bayes filter?***

$p(z|x=5) = ?$

$P(z=door, x=5) = 0.5$
$P(z=door, x=4) = 0.25$
$P(z=door, x=3) = 0$

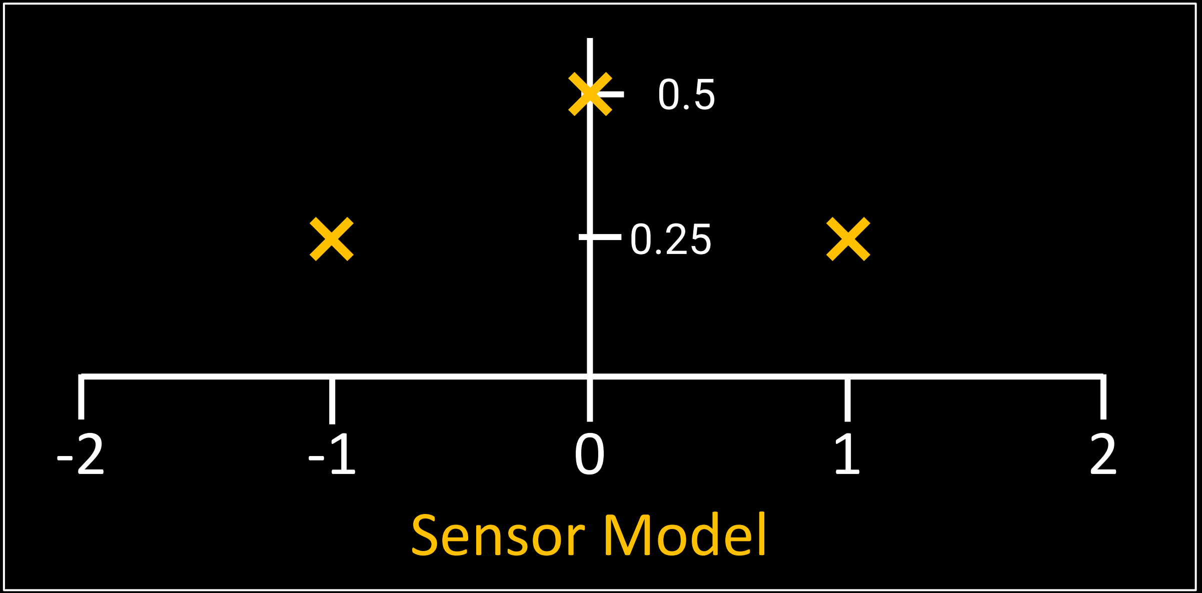

$$p(x+1|x,u=+1) = 0.5$$
$$p(x|x,u=+1) = 0.5$$

$$p(x-1|x,u=-1) = 0.5$$
$$p(x|x,u=-1) = 0.5$$

Motion Model

Bayes Filter - Example 1

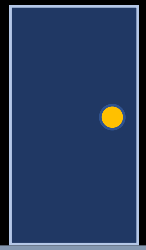

At $t = 0$, no information

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

Bayes Filter - Example 1

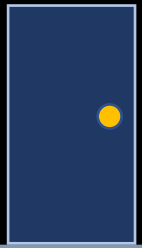

At $t = 0$, no information

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 $t = 1$, $U_1 = \text{do\_nothing}$, $Z_1 = \text{door}$

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

Do we have to do the prediction step?

Do the update step!

Bayes Filter - Example 1

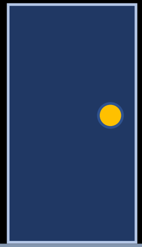

At $t = 0$, no information

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 $t = 1$, $U_1 = \text{do\_nothing}$, $Z_1 = \text{door}$

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

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

Bayes Filter - Example 1

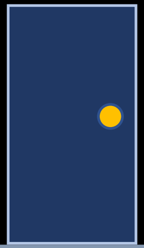

At $t = 1$, $U_1 = \text{do\_nothing}$, $Z_1 = \text{door}$

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
$p(x_2)$						

Bayes Filter - Example 1

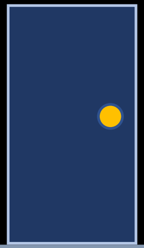

At $t = 1$, $U_1 = \text{do\_nothing}$, $Z_1 = \text{door}$

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	$\frac{1}{3} \times \frac{1}{2}$	$\frac{1}{3} \times \frac{1}{2} + \frac{2}{3} \times \frac{1}{2}$	$\frac{2}{3} \times \frac{1}{2}$

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

Bayes Filter - Example 1

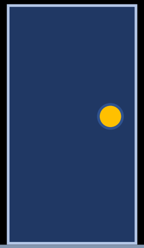

At $t = 1$, U_1 = do_nothing, Z_1 = door

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

At $t = 2$, $U_2 =$ -1, Z_2 = door

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

Bayes Filter - Example 1

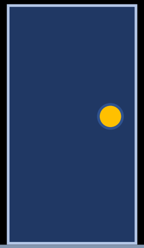

At $t = 1$, $U_1 = \text{do\_nothing}$, $Z_1 = \text{door}$

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

State	0	1	2	3	4	5
$p(x_2)$	0	0	0	0	$\frac{3}{7}$	$\frac{4}{7}$

Bayes Filter - Example 1 (initial conditions 1)

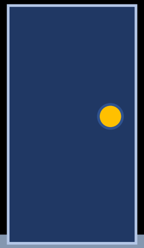

At t=0, we are absolutely certain the robot is at state $X_0 = 0$

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

Bayes Filter - Example 1 (initial conditions 1)

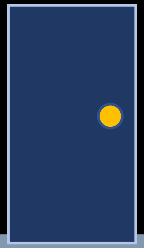

At t=0, we are absolutely certain the robot is at state $X_0 = 0$

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

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

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

Bayes Filter - Example 1 (initial conditions 1)

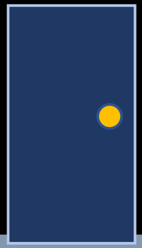

At t=0, we are absolutely certain the robot is at state $X_0 = 0$

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

At t=1, U_1 = do_nothing , Z_1 = door

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

Bayes Filter - Example 1 (initial conditions 2)

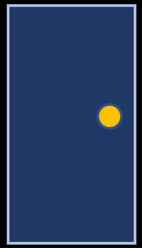

At t=0, we are "absolutely" certain the robot is at state $X_0 = 0$

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, $U_1 = \text{do\_nothing}$, $Z_1 = \text{door}$

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

Bayes Filter - Example 1 (initial conditions 2)

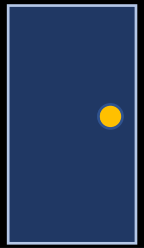

At t=0, we are “absolutely” certain the robot is at state $X_0 = 0$

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, $U_1 = \text{do\_nothing}$, $Z_1 = \text{door}$

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

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

(deterministic approaches are fragile!)

Bayes Filter – Example 2

- Bayes with beans
  - World
    - 1D continuous robot world
    - Discretized into 7 states
    - …with a door at state 5
  - Motion model
    - 80% correct, 20% fails
  - Sensor model
    - 90% correct, 10% fails
  - Initial belief
  - Take an action: +1
  - Take a sensor reading: door!

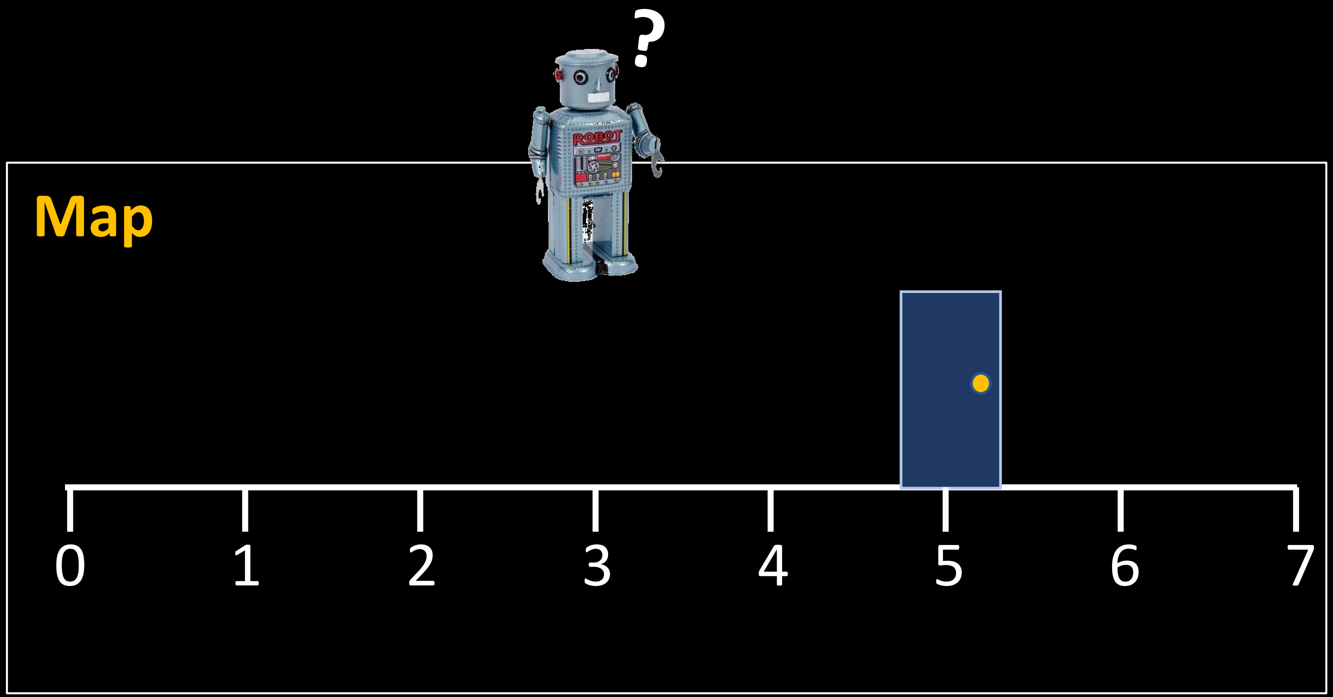

Bayes Filter – Example 3

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

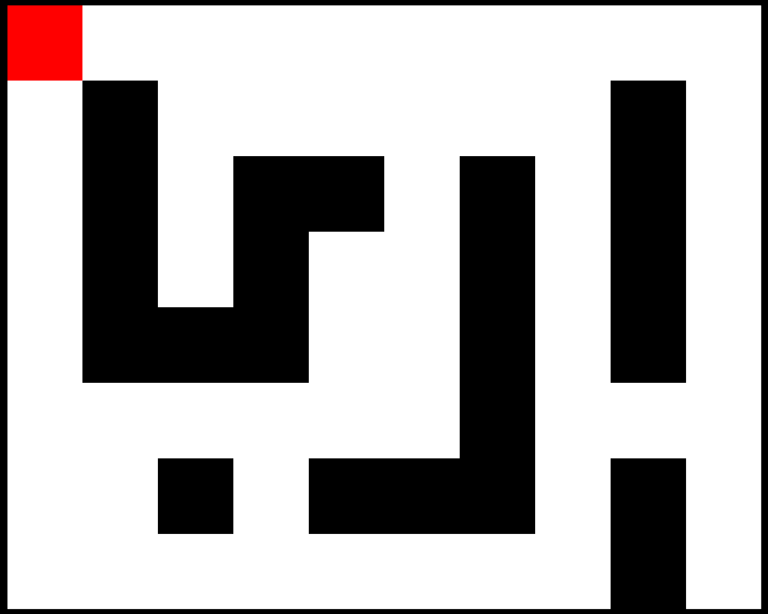

X is the set of possible locations
x is one of these locations

Bayes Filter – Example 3

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

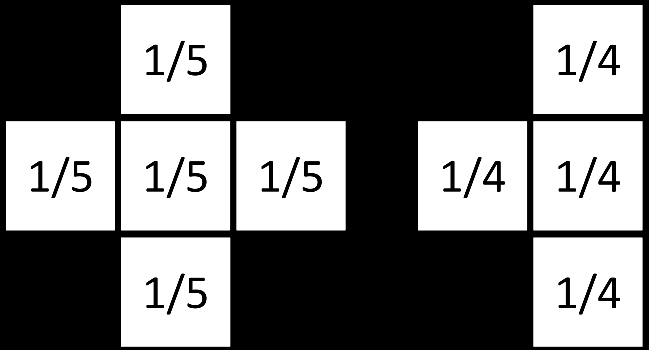

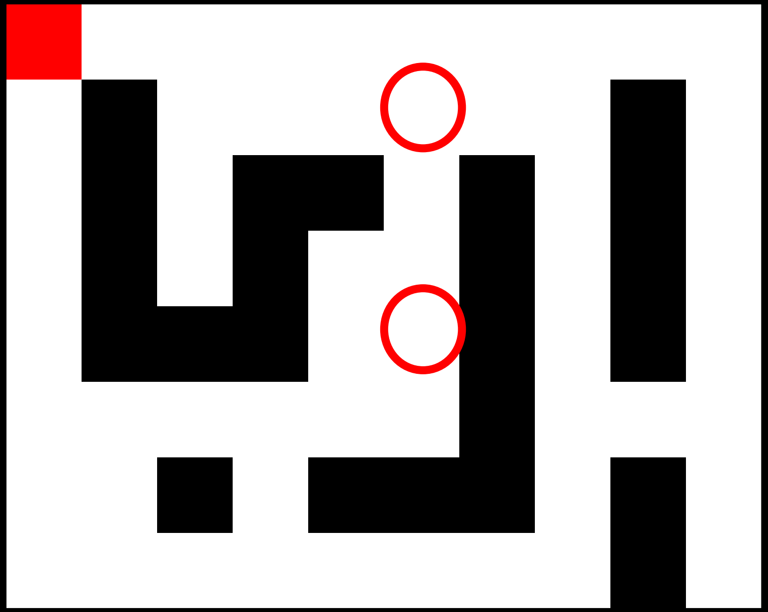

X is the set of possible locations
x is one of these locations

Bayes Filter – 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*

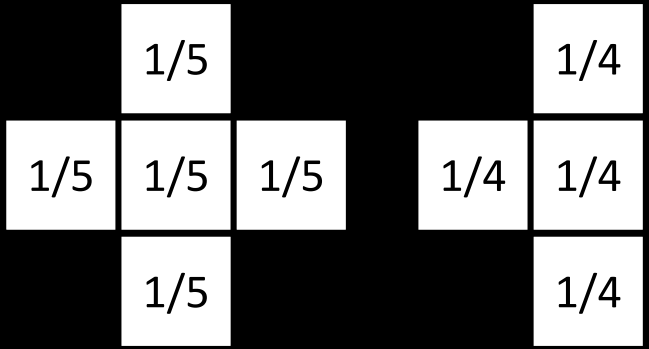

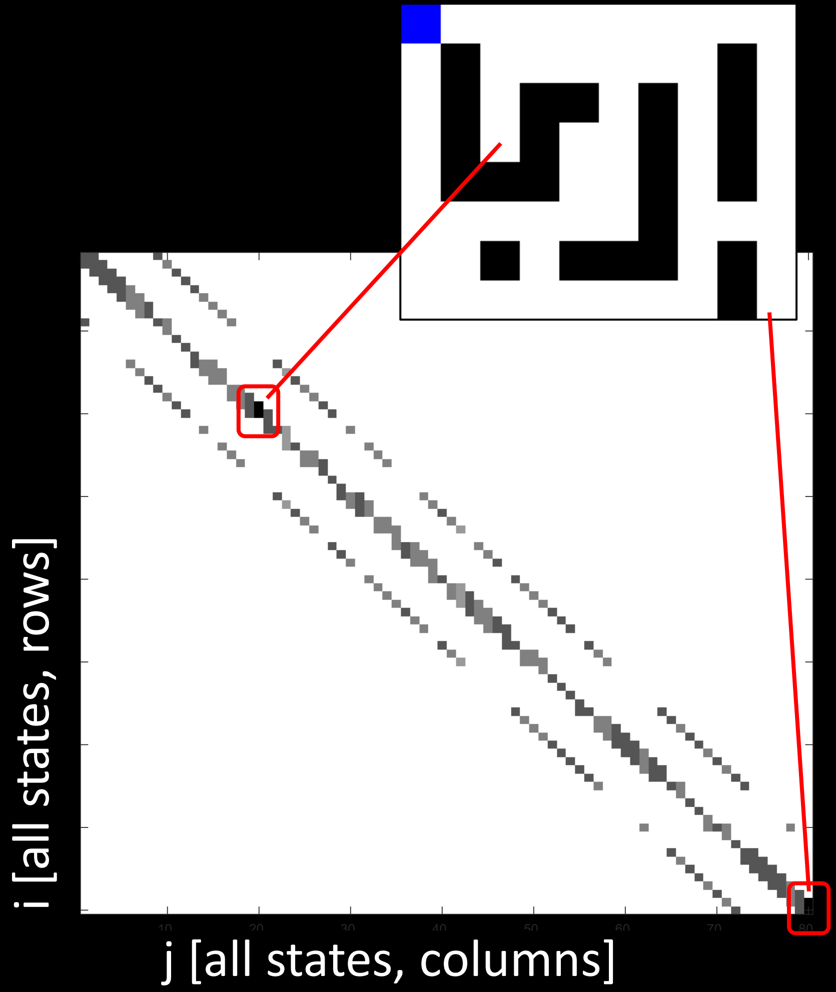

Bayes Filter – Example 3

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

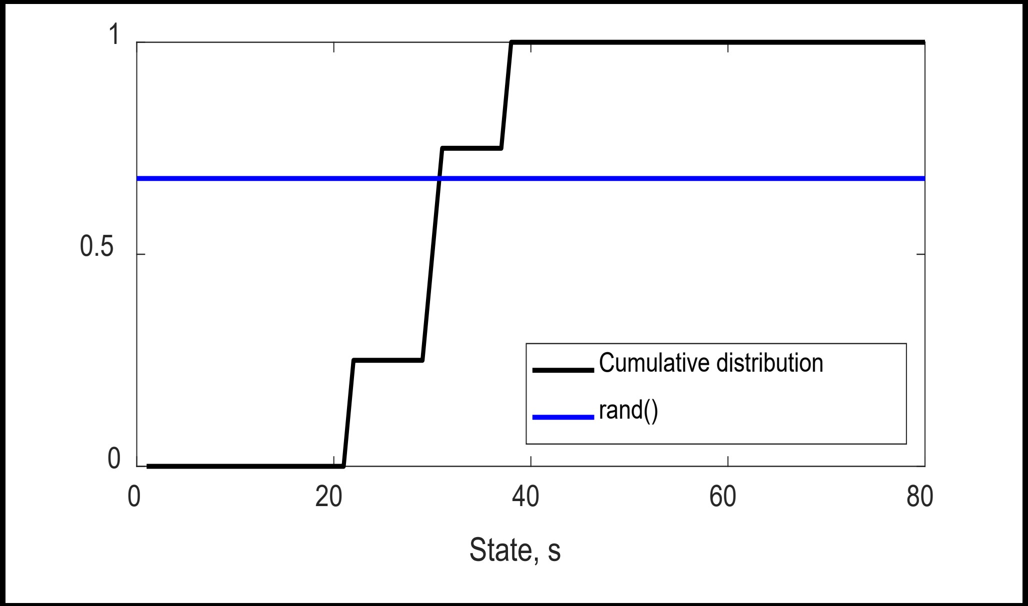

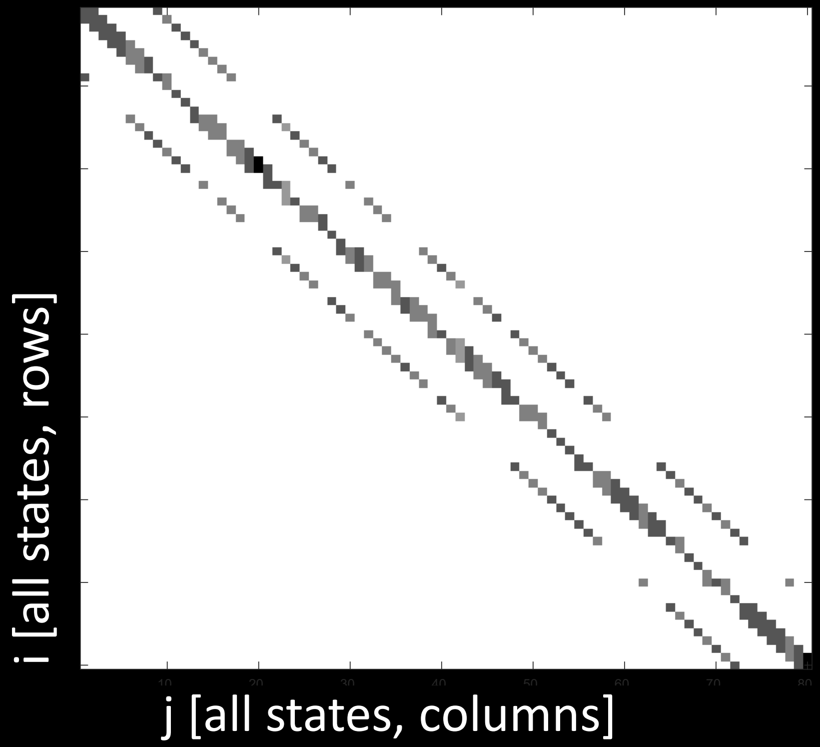

Bayes Filter – Example 3

- Prediction step

Prediction step $(bel(x_{t-1}), \cancel{u_t})$:
1. for all x_t do
2. $\overline{bel}(x_t) = \sum_{x_{t-1}} p(x_t | \cancel{u_t}, x_{t-1}) \; bel(x_{t-1})$
3. endfor

Matrix implementation

1. $\overline{bel} = A \; bel_{t-1}$

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

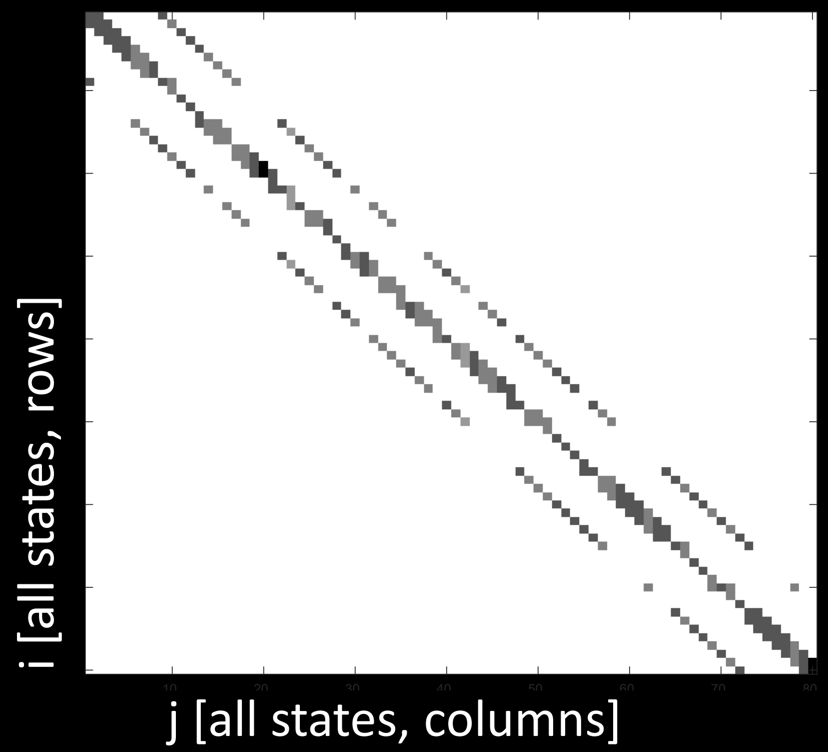

Bayes Filter – Example 3

- Prediction step

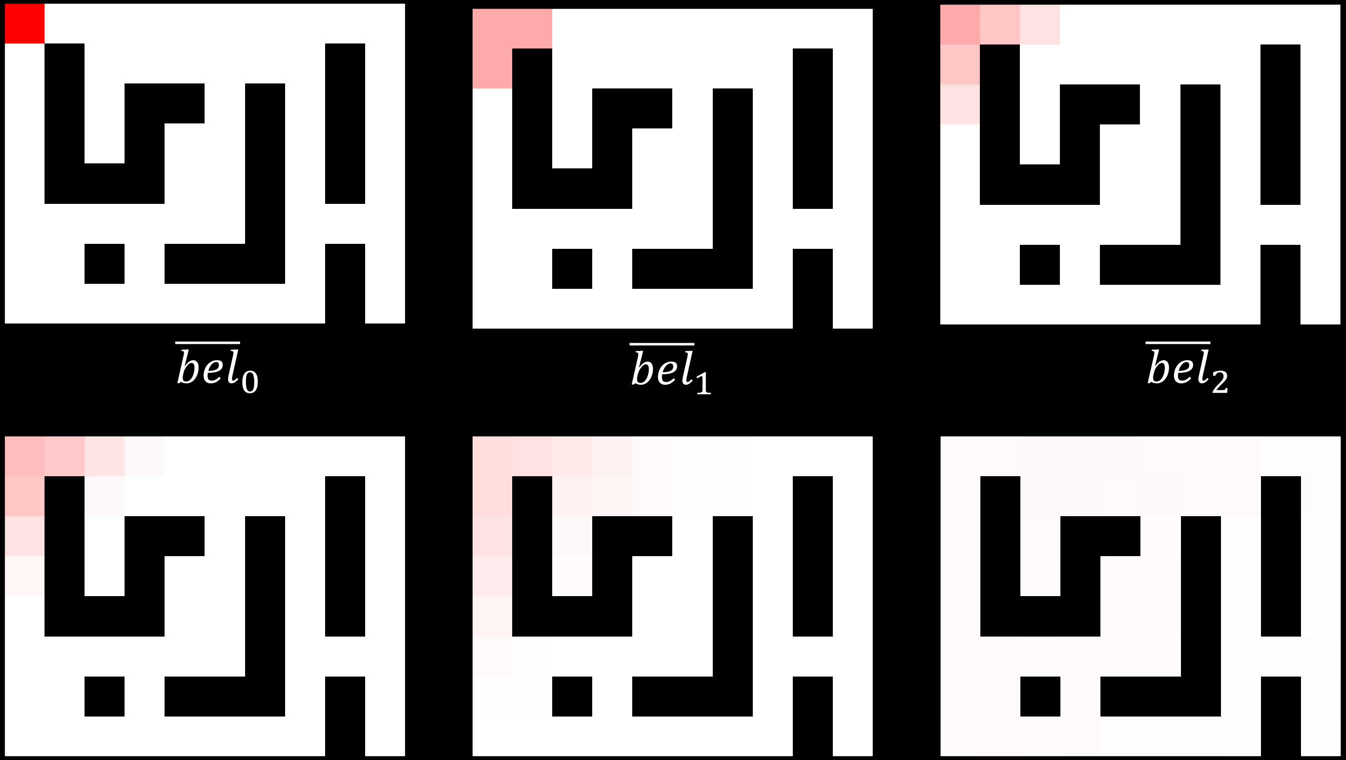

$\overline{bel}_0$ $\overline{bel}_1$ $\overline{bel}_2$

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

Bayes Filter – Example 3

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

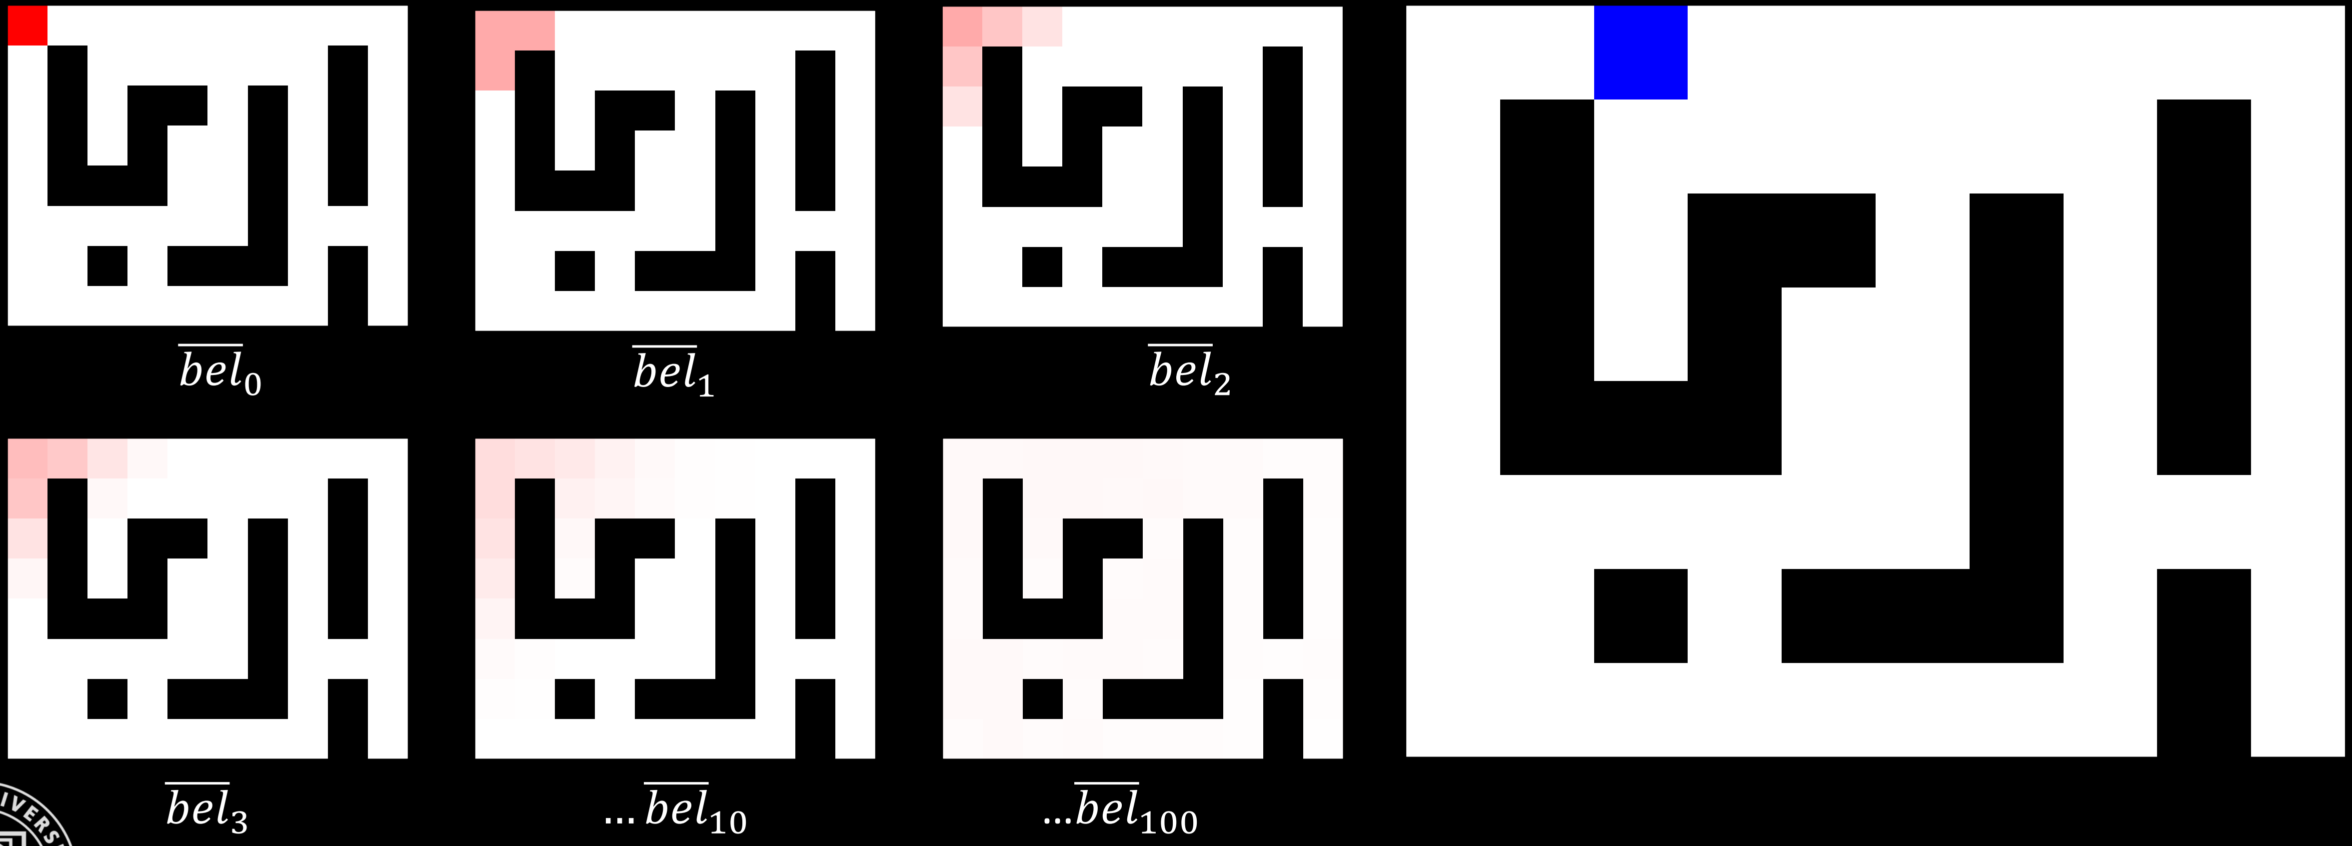

Bayes Filter – Example 3

- 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

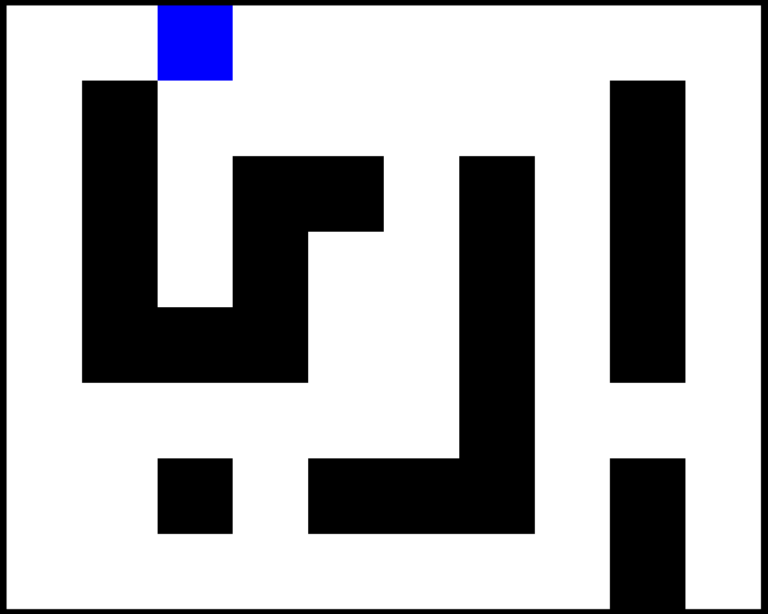

X is the set of possible locations
x is one of these locations
z are the sensor measurements

Bayes Filter – 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(no walls | x) = 0.1*0.9*0.9*0.9
- P(N| x) = 0.9*0.9*0.9*0.9 ← *highest likelihood*
- P(W| x) = 0.1*0.9*0.9*0.1
- P(S | x) = 0.1*0.9*0.1*0.9
- P(E | x) = 0.1*0.1*0.9*0.9
- …
- P(NW | x) = 0.9*0.9*0.9*0.1
  - How many combinations are there per state?
    - 2^4

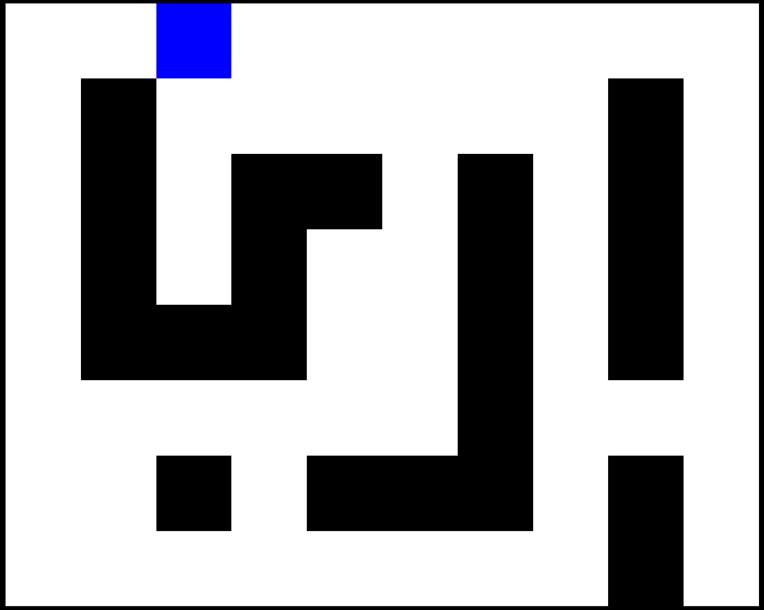

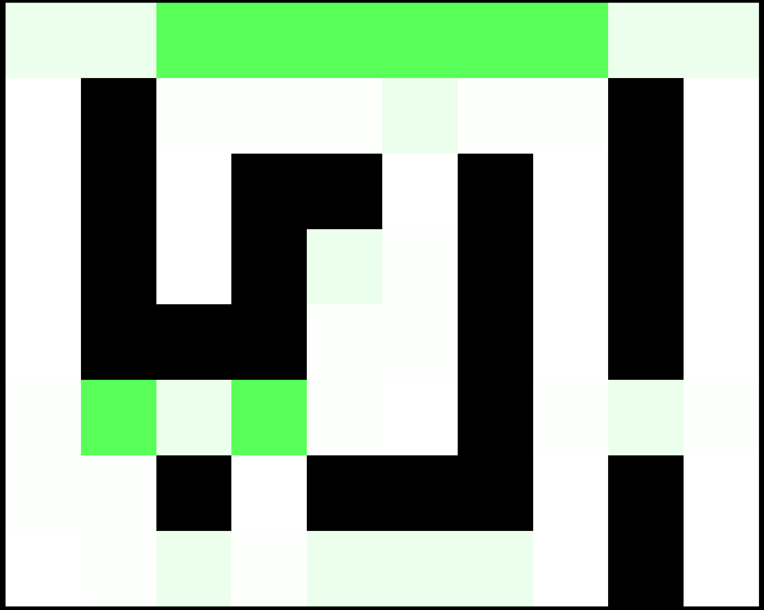

P(z|X)

Bayes Filter – 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

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

1. for all x_t do
2. $\overline{bel}(x_t) = \sum_{x_{t-1}} p(x_t | u_t, x_{t-1})\ bel(x_{t-1})$
3. $bel(x_t) = \eta\ p(z_t | x_t)\ \overline{bel}(x_t)$
4. endfor
5. return $bel(x_t)$

- If all readings match:
  - $\sum |z_t - z'_{xt}| = 0$
  - $p_z(x_t) = 0.6561$
- If all readings differ:
  - $\sum |z_t - z'_{xt}| = 4$
  - $p_z(x_t) = 0.0001$

Compute likelihood of observations, p_{zX}

1. for all x_t do
2. $pzX(x_t) = 0.9^{4 - \sum |z_t - z'_{xt}|}\ 0.1^{\sum |z_t - z'_{xt}|}$
3. Endfor

…where p_{zX} is a vector (80x1)

Bayes Filter – 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

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

1. for all x_t do
2. $\overline{bel}(x_t) = \sum_{x_{t-1}} p(x_t | u_t, x_{t-1}) \; bel(x_{t-1})$
3. $bel(x_t) = \eta \, p(z_t | x_t) \, \overline{bel}(x_t)$
4. endfor
5. return $bel(x_t)$

Compute new belief

1. $bel_t = p_{zX} \, \overline{bel} \, / \sum(p_{zX} \, \overline{bel})$

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

Bayes Filter – Example 3

- Bayes Filter

Algorithm Bayes_Filter(bel_{t-1}, z_t):

1. $\overline{bel} = A\, bel_{t-1}$
2. for all x_t do
3. $p_{zX}(x_t) = 0.9^{4-\sum|z_t - z'_{xt}|}\, 0.1^{\sum|z_t - z'_{xt}|}$
4. Endfor
5. $bel_t = \overline{bel}\, pz_X \,/\, \sum(\overline{bel}\, pz_X)$

Only do computations for states with a belt-1 > threshold

Precache and look up for faster operation

Bayes Filter – Example 3

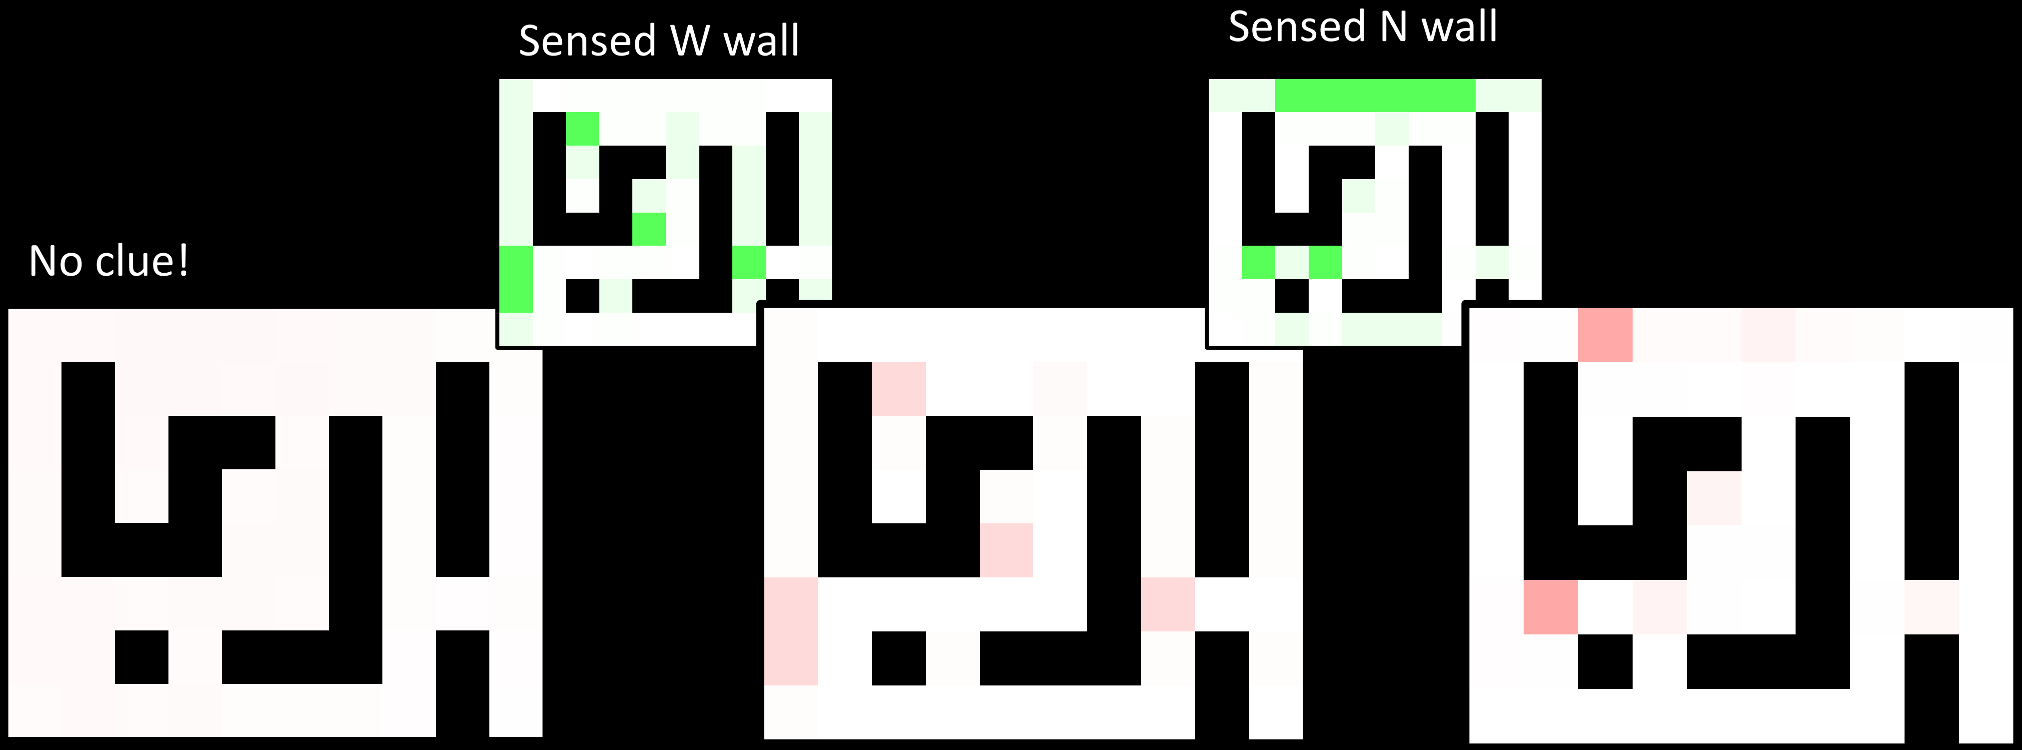

In two steps, we homed in on where we are!
...

- ***How good is the Bayes Filter?***
- ***Can you do better?***
  - Improved transition model
  - Deliberately move in directions that give you more information

Bayes Filter II

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

1. for all x_t do
2. $\overline{bel}(x_t) = \sum_{x_{t-1}} p(x_t | u_t, x_{t-1})\ bel(x_{t-1})$ [Prediction Step]
3. $bel(x_t) = \eta\, p(z_t | x_t)\ \overline{bel}(x_t)$ [Update/Measurement Step]
4. endfor
5. return $bel(x_t)$

- Example 1
  - Robot in a 1D world
  - The importance of having some belief in all states
- Example 2
  - Bayes with beans
  - The importance of normalization
- Example 3
  - (x,y)-robot in a grid world
  - Computational efficiency
    - Matrices
    - Pre-cache observations

Summary

- Use temporal consistency between observations that are poor estimates individually
- Localization can work with…
  - …completely random motion
  - …noisy sensors
  - Remember to…
    - Don't be deterministic
    - Normalize
    - Efficient computation

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

1. for all x_t do
2. $\overline{bel}(x_t) = \sum_{x_{t-1}} p(x_t | u_t, x_{t-1})\, bel(x_{t-1})$ (Motion models: $p(x_t | u_t, x_{t-1})$)
3. $bel(x_t) = \eta\, p(z_t | x_t)\, \overline{bel}(x_t)$ (Sensor models: $p(z_t | x_t)$)
4. endfor
5. return $bel(x_t)$

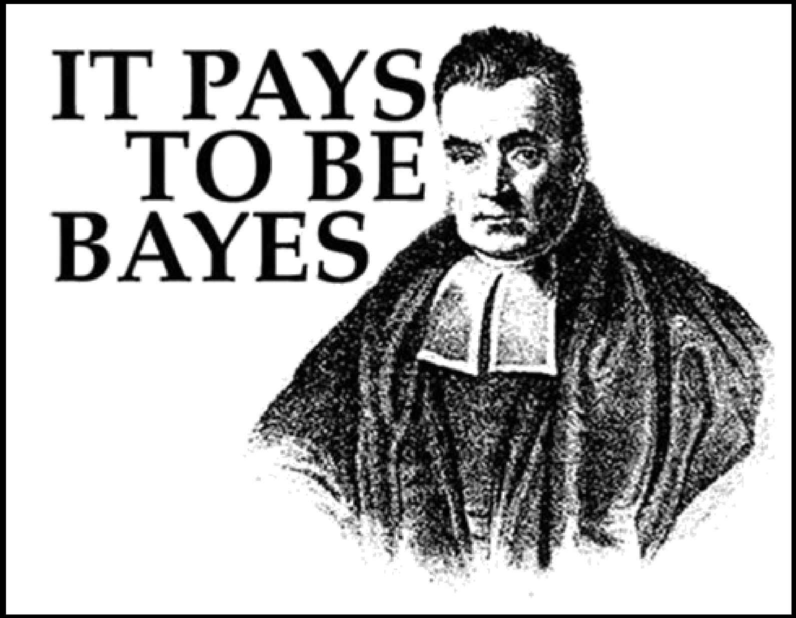

Logistics

- **Lab 8 – Stunts**
  - Voting can start ~~Monday~~ *next* Monday 18th
  - Please submit your votes by Friday April 22nd
    - https://tinyurl.com/vp5wrten
    - 10 points for best stunt
    - 1 point for best blooper
- ***Teammates??***
  - https://forms.gle/HQr6hRK8AER5JcpD6
- **Lab 9 – Mapping**
  - Yes! You still have an opportunity to get your map this week
    - We plan to start grading Friday
- **Lab 10 – Simulation**
  - If you finish early, we strongly encourage you to get a head start on the Lab 11 documentation!

Prof. Kirstin Hagelskjær Petersen
kirstin@cornell.edu

Fast Robots

Upcoming labs

Lab 10 – Simulation Software

- Multiple processes
  - Simulator
    - Robot
      - Motion
      - Ground truth
    - YAML (map and other parameters)
  - Plotter
  - Controller
    - Get odometry pose, get and plot sensor data, move the robot, etc.

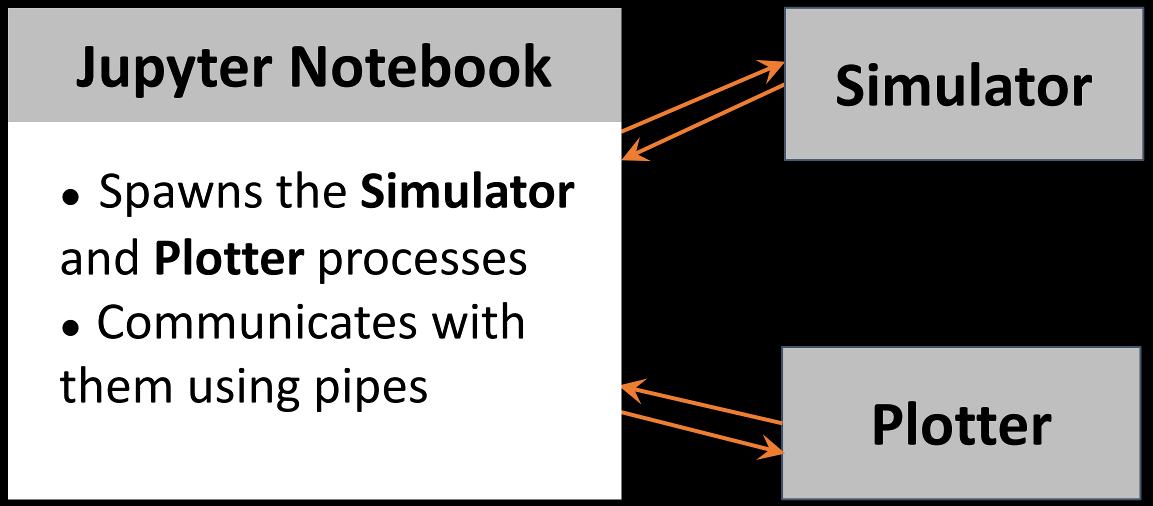

Lab 10 - Simulation

Commander

Functions to interact with the simulator and the plotter

Lab 11 - Localization on the virtual robot

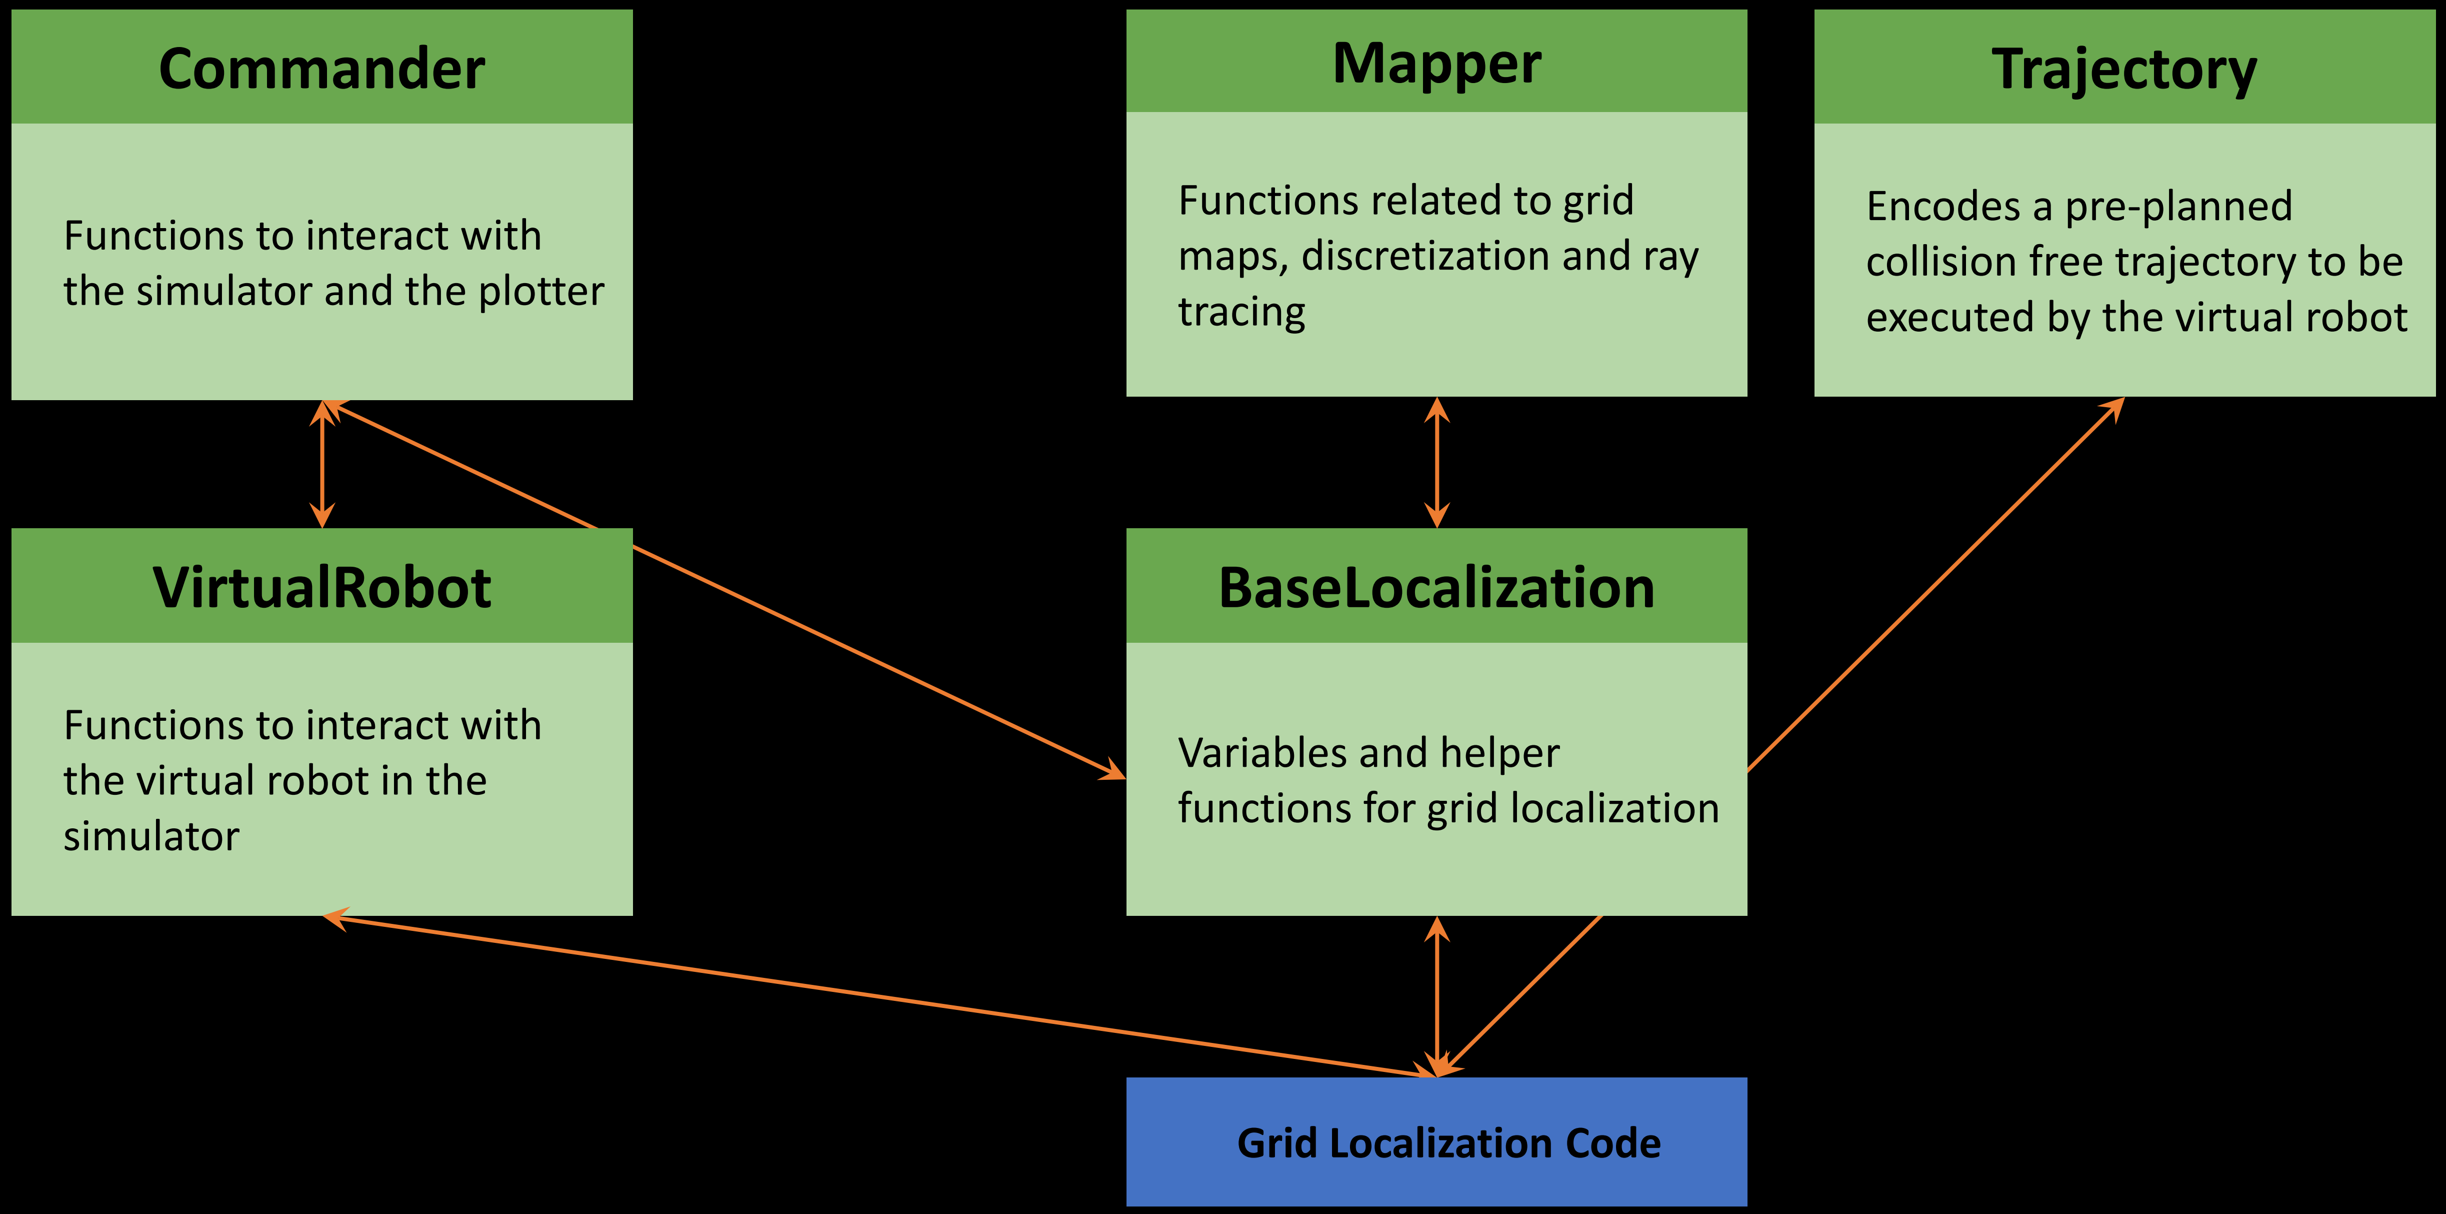

Lab 12 - Pure localization on the virtual robot

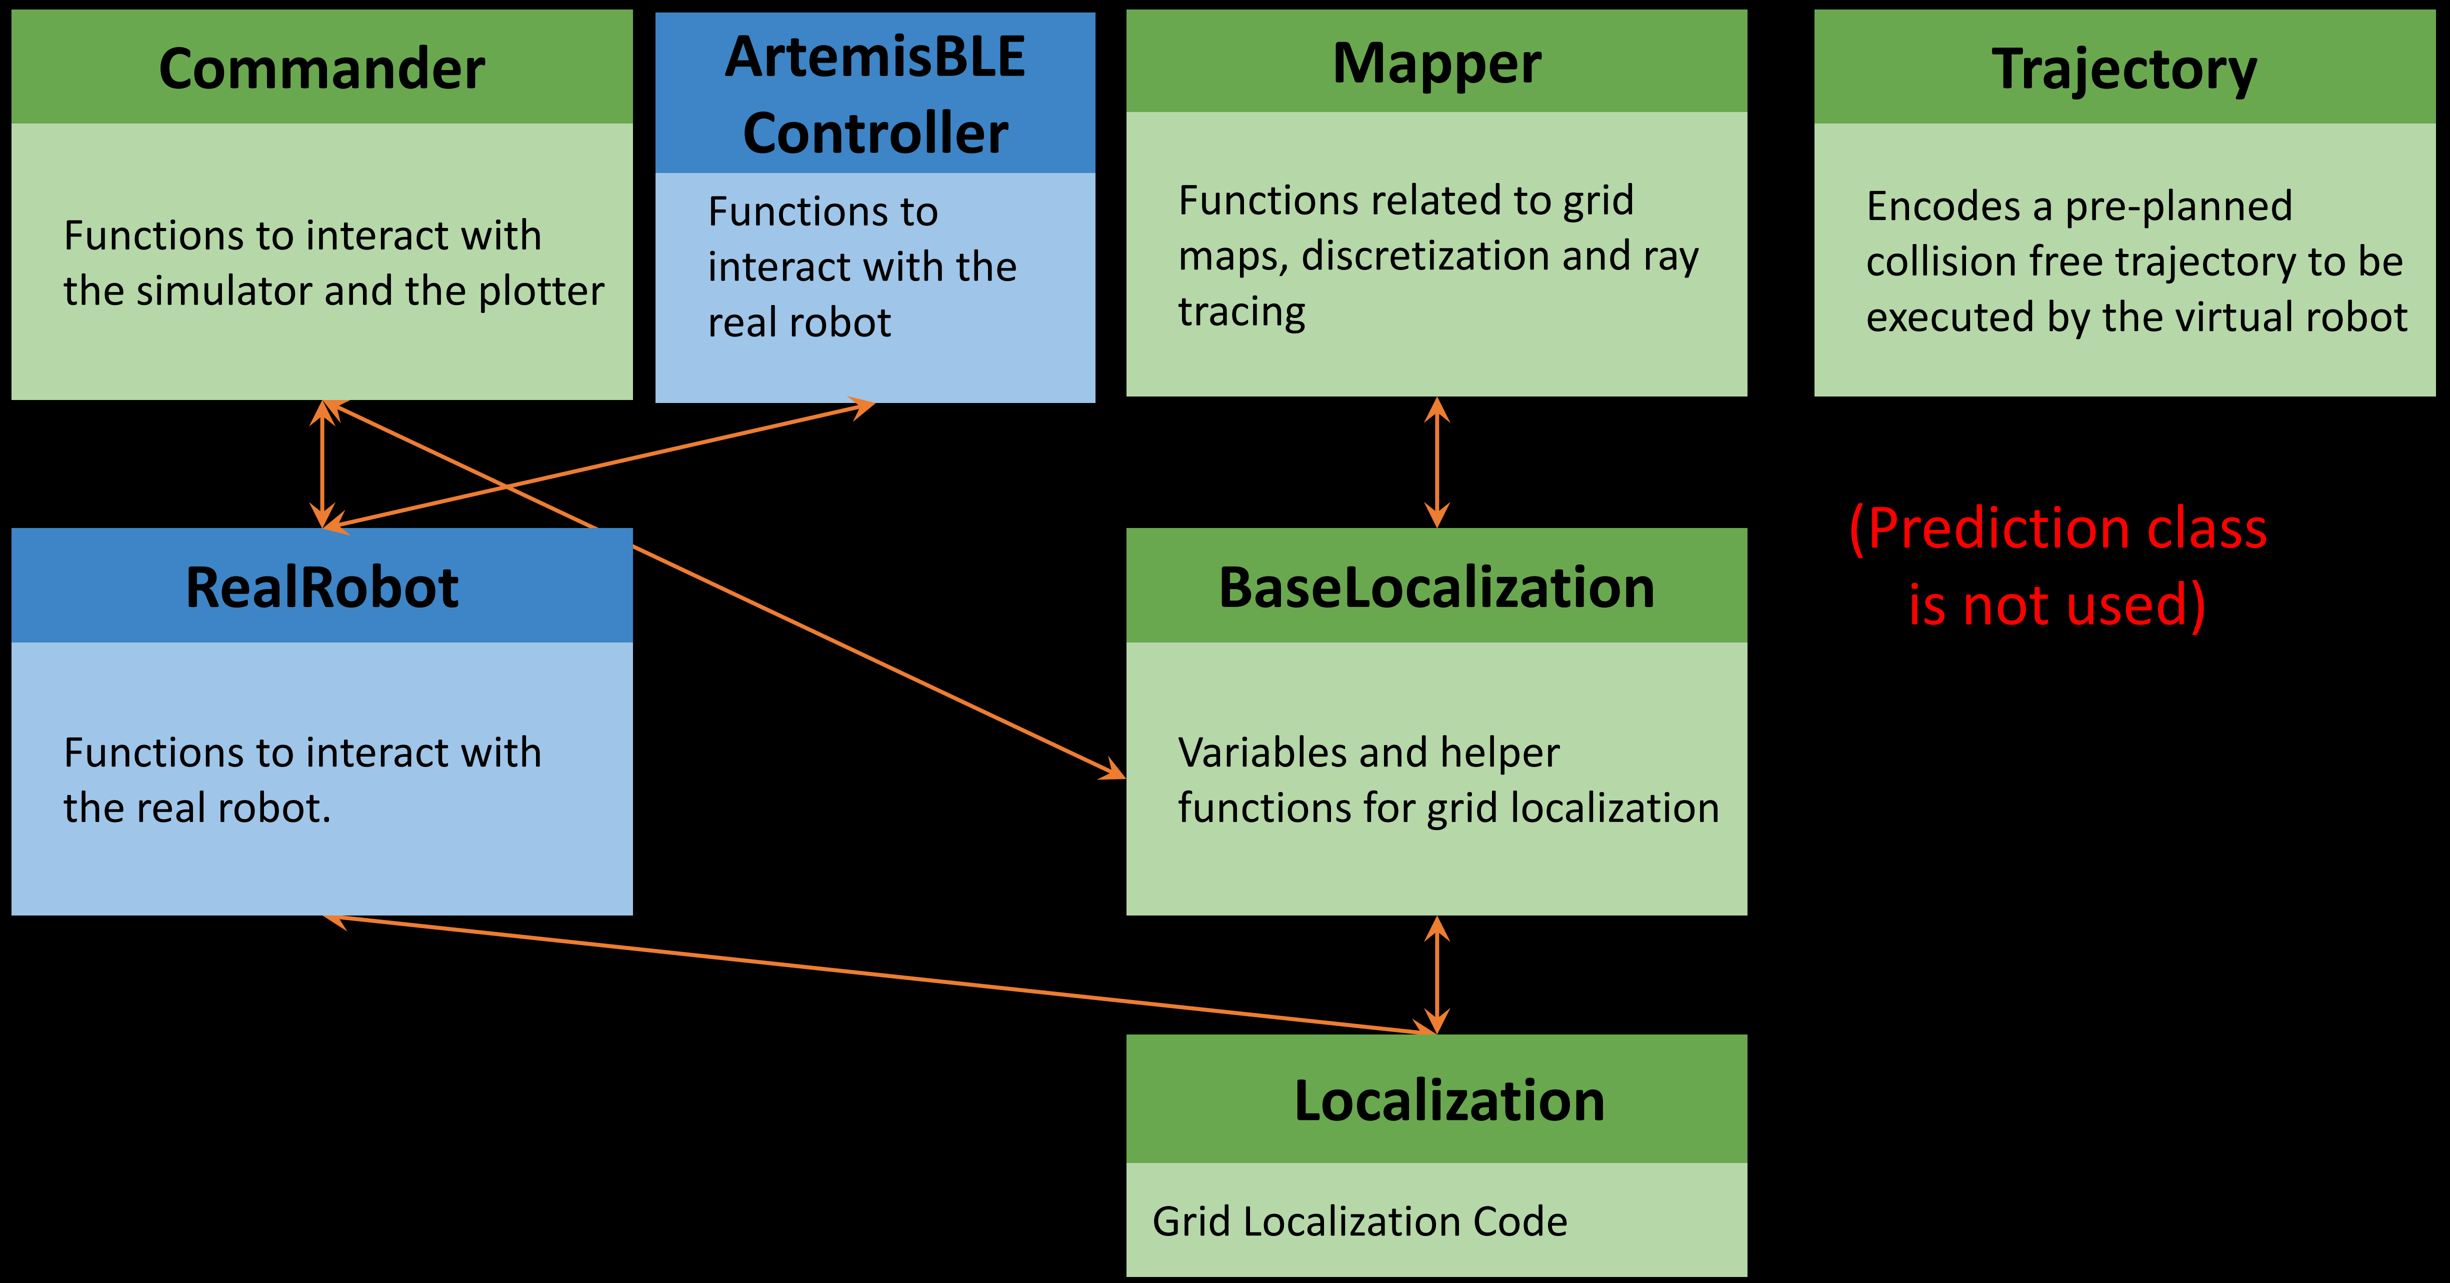

Lab 13 - Localization and planning on the virtual robot