ECE3400: Intelligent Physical Systems

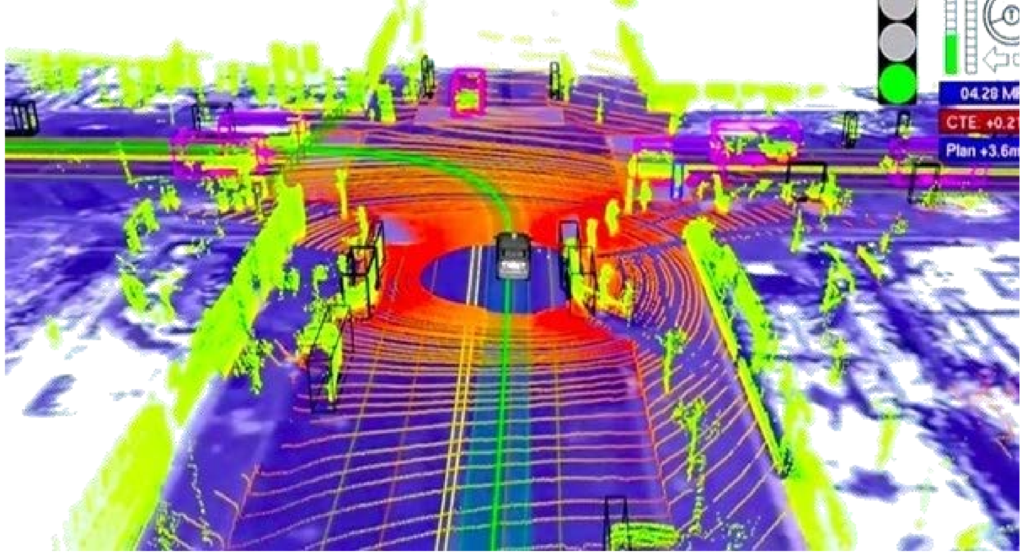

Probability and Localization

Classes of Interest:
ECE 3100: Intro to Probability and Inference
ECE 5412: Bayesian Estimation and Learning
MAE 4180/5780, CS3758: Autonomous Mobile Robots

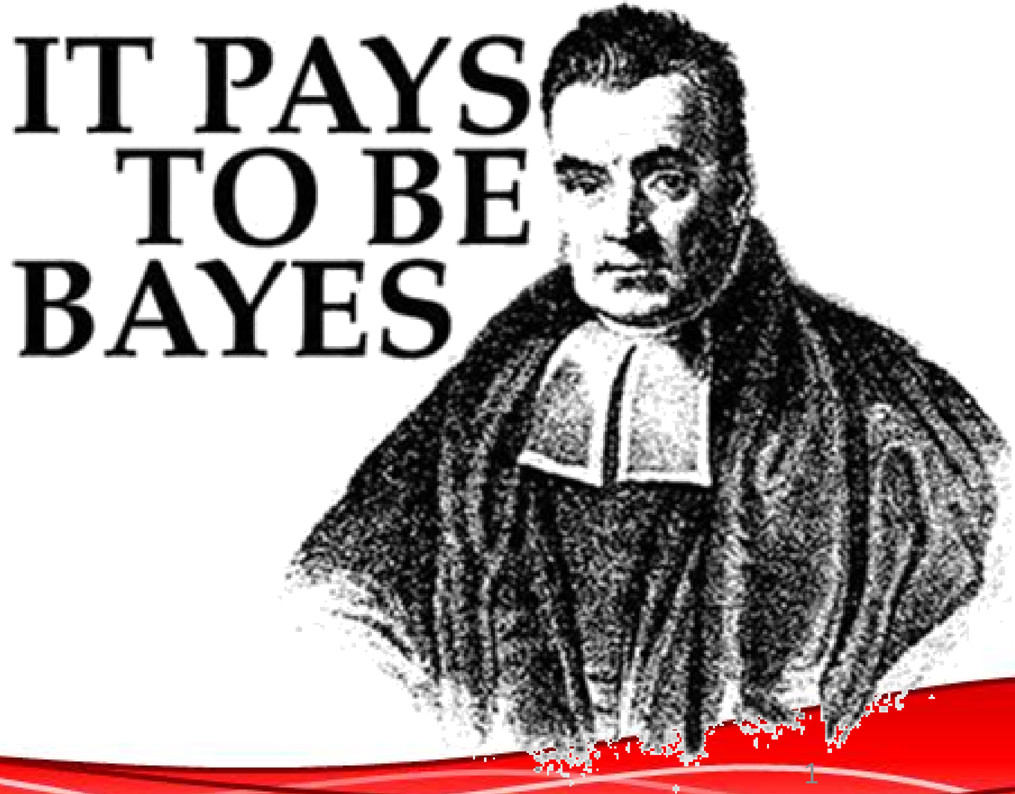

Reliability

- *How well does your robot go straight?*
- *How well does your robot turn?*
- *How well does your wall sensor detect walls?*
    - *...Are you really sure it works perfectly?*

Nothing's prefect!

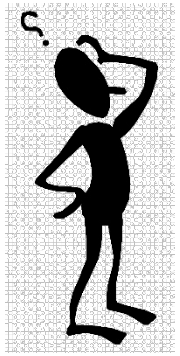

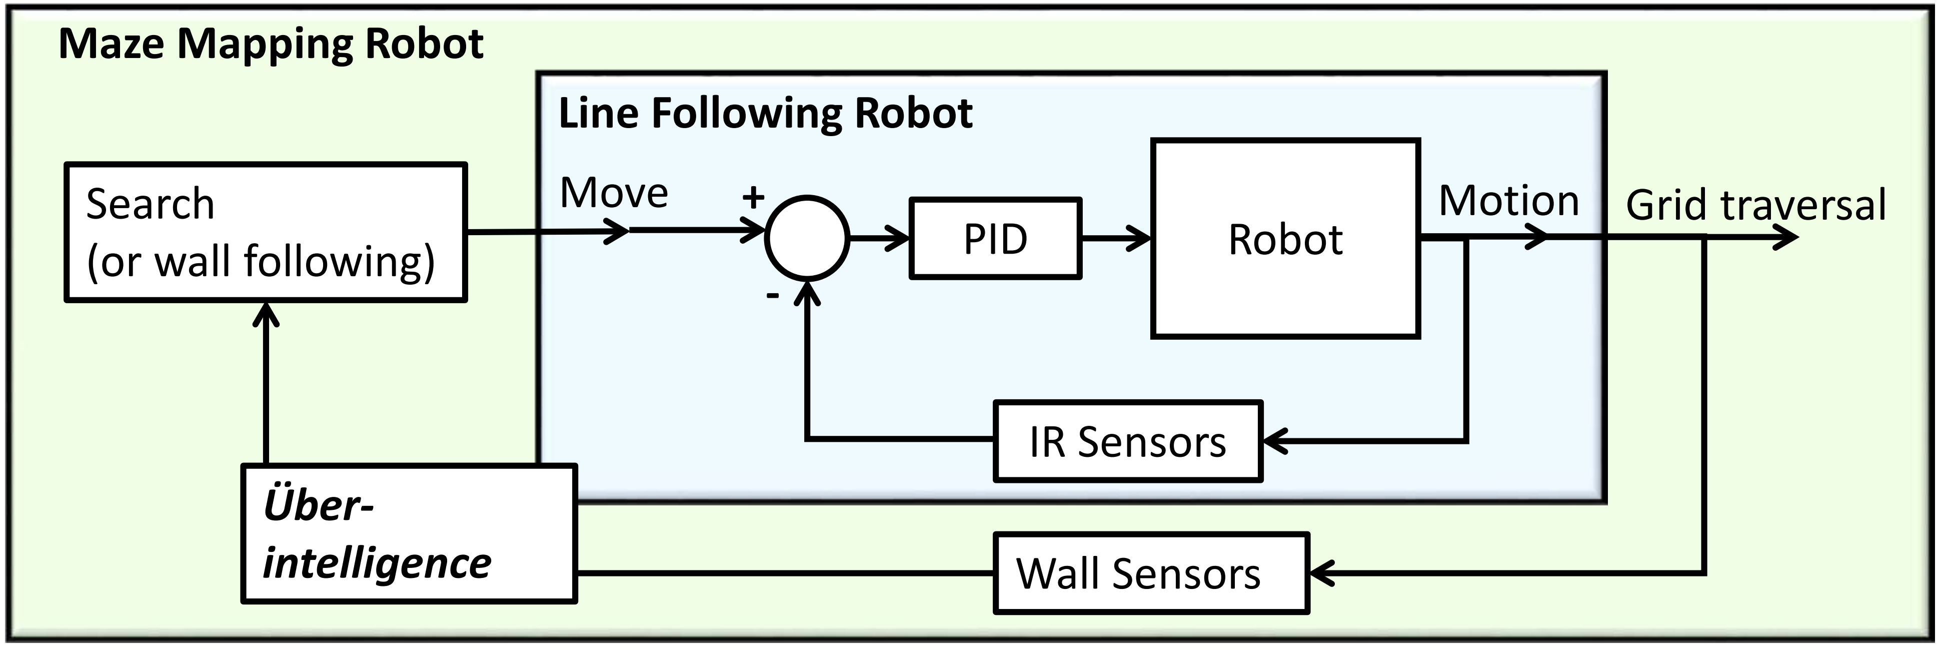

Simultaneous Localization And Mapping

Reliability

- *How well does the robot go straight?*
  - Describe your setup:
    - Move straight between two junctions 50 times in a row
  - Results:
    - It went straight 46 out of 50 times
    - Twice it went left/right instead
    - Once it overshot
    - Once it stayed put
    - *GoStraight() is 92% reliable.*
  - Results:
    - Mean ± standard deviation
    - Max overshoot

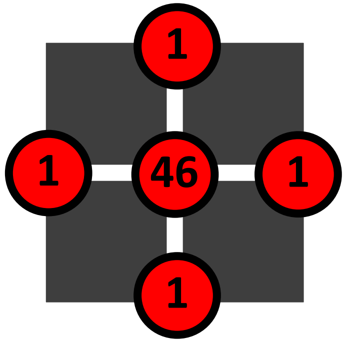

Frequentist Statistics

- *What is the issue with this?*
  - Depends on the number of trials

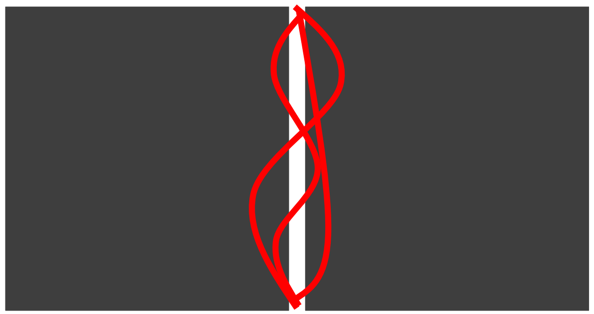

Bayesian Statistics

- Educated guesses based on probability distributions, to update beliefs in the evidence of new data
- *Representing all parameters as probability distributions*

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

likelihood: $P(B|A)$

prior: $P(A)$

conditional probability / *posterior*: $P(A|B)$

marginal likelihood (constant): $P(B)$

Bayesian Statistics

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

I saw a wall to the north, am I to the right or the left of the red line?

- Scenario #1:
  - Left 6 squares: 3 north wall, 3 have no north wall
  - Right 6 squares: 2 north wall, 4 have no north wall
- ***What is the best guess?***
  - Left
- Scenario #2:
  - Left 6 squares: 3 north walls
  - Right 12 squares: 4 north walls
- ***What is the best guess?***
  - Right

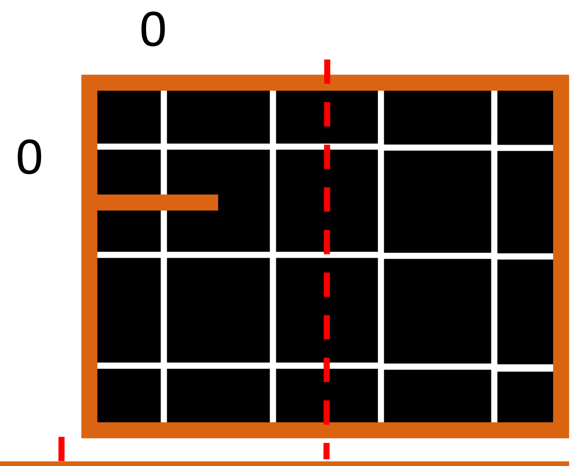

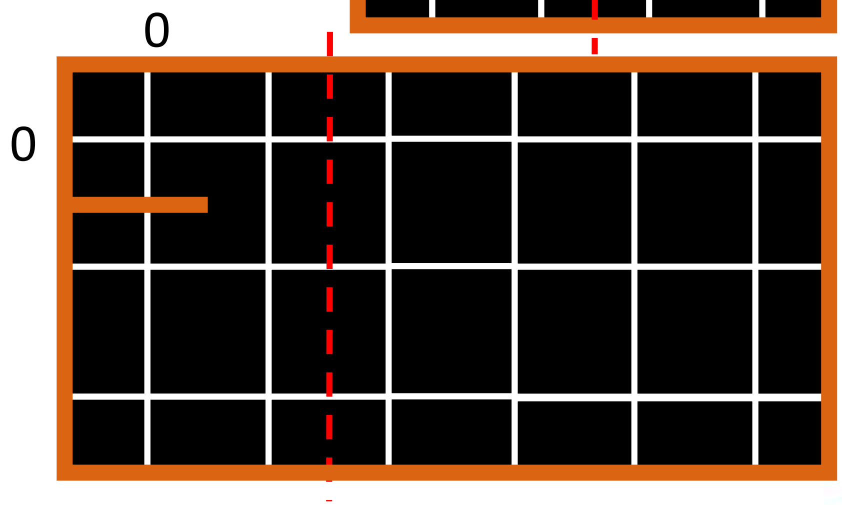

Bayesian Statistics

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

I know which side I'm on, what is the probability there is a wall to the north?

Conditional Probability

- Scenario #1:
  - $P(wall_{NORTH}|left) = 3/6 = 0.5$
  - $P(wall_{NORTH}|right) = 2/6 = 0.33$
- Scenario #2
  - $P(wall_{NORTH}|left) = 3/6 = 0.5$
  - $P(wall_{NORTH}|right) = 4/12 = 0.33$

- *Not interchangeable:
  - $P(A|B) \neq P(B|A)$

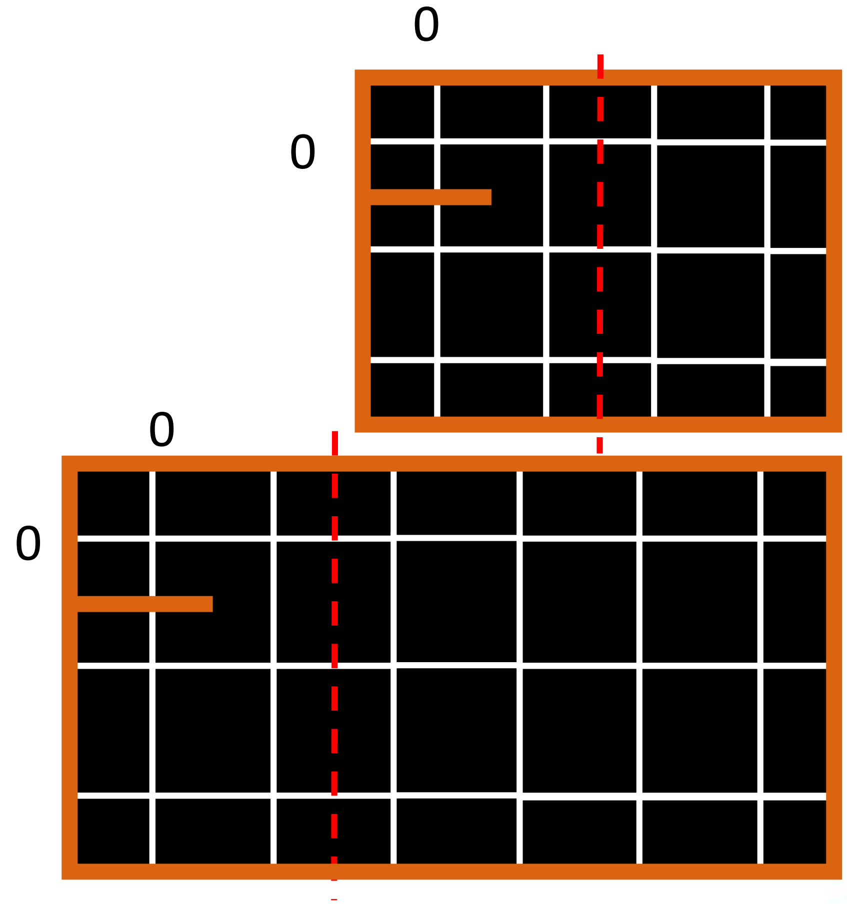

Bayesian Statistics

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

What is the probability that I am on the left side and there is a wall to my north?

- **Joint Probability**
- Scenario #1
  - $P(\text{left} \cap \text{wall}_{NORTH})$
    $= P(\text{left}) * P(\text{wall}_{NORTH}|\text{left})$
    $= 0.5*0.5 = 0.25$
  - $P(\text{right} \cap \text{wall}_{NORTH})$
    $= P(\text{right}) * P(\text{wall}_{NORTH}|\text{right})$
    $= 0.5*0.33 = 0.17$

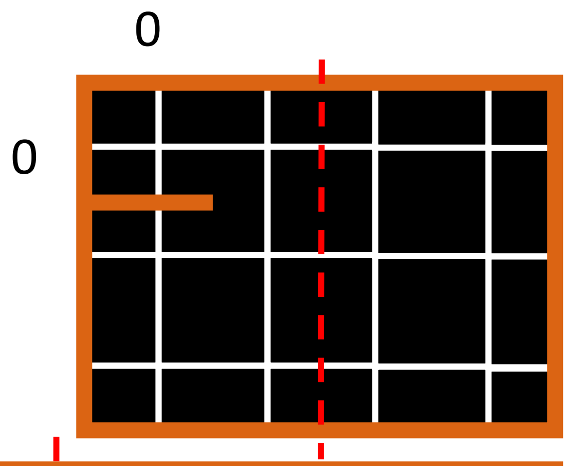

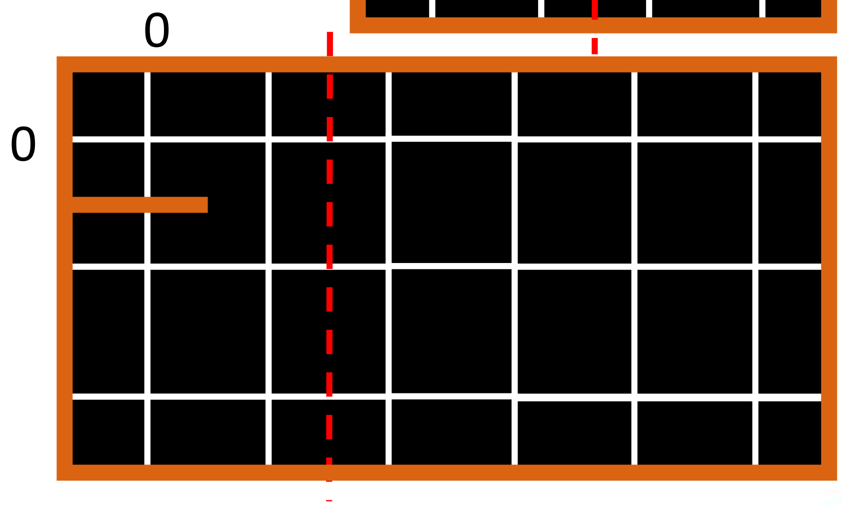

Bayesian Statistics

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

What is the probability that I am on the left side and there is a wall to my north?

- **Joint Probability**
- Scenario #1
  - $P(\text{left} \cap \text{wall}_{NORTH})$
    $= P(\text{left}) * P(\text{wall}_{NORTH}|\text{left})$
    $= 0.5*0.5 = 0.25$
- Scenario #2
  - $P(\text{right} \cap \text{wall}_{NORTH})$
    $= P(\text{right}) * P(\text{wall}_{NORTH}|\text{right})$
    $= 0.67*0.33 = 0.22$
- **Interchangeable:*
  - $P(A \cap B) = P(B \cap A)$

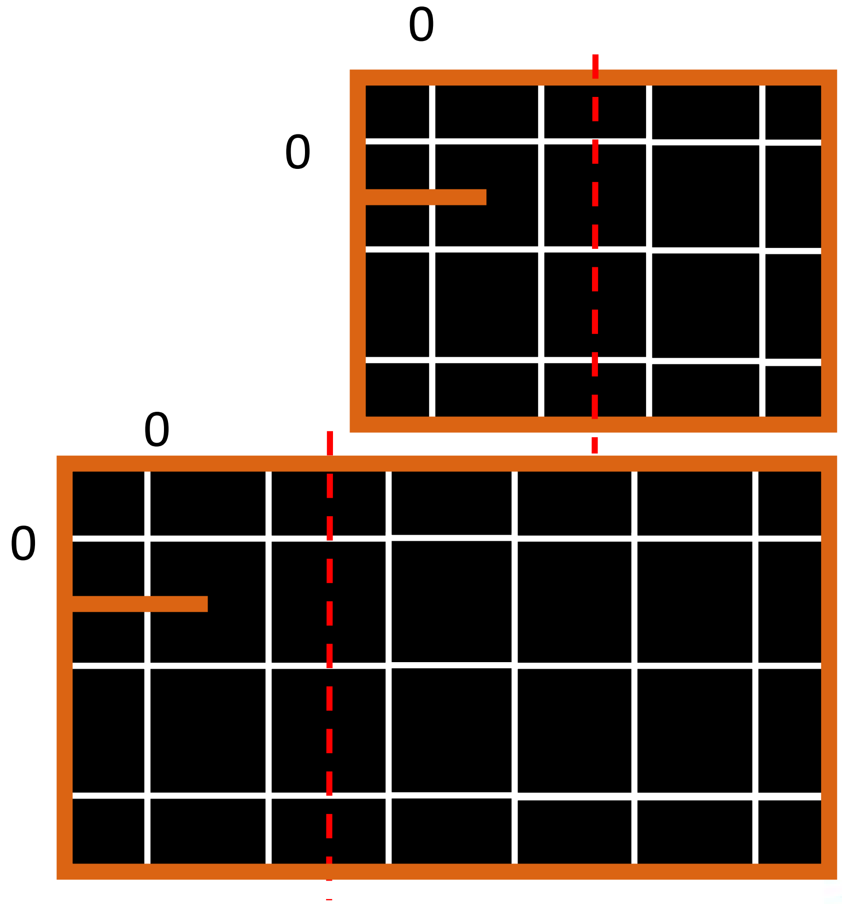

Bayesian Statistics

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

What is the probability that there is a wall to the north?

- **Marginal Likelihood**
- Scenario #1
  - $P(wall_{NORTH})$
    $= P(left \cap wall_{NORTH}) + P(right \cap wall_{NORTH})$
    $= 0.25 + 0.17 = 0.42$
- Scenario #2
  - $P(wall_{NORTH})$
    $= P(left \cap wall_{NORTH}) + P(right \cap wall_{NORTH})$
    $= 0.17 + 0.22 = 0.39$

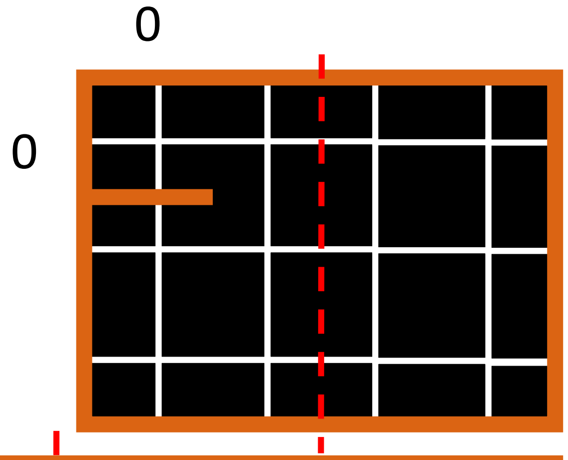

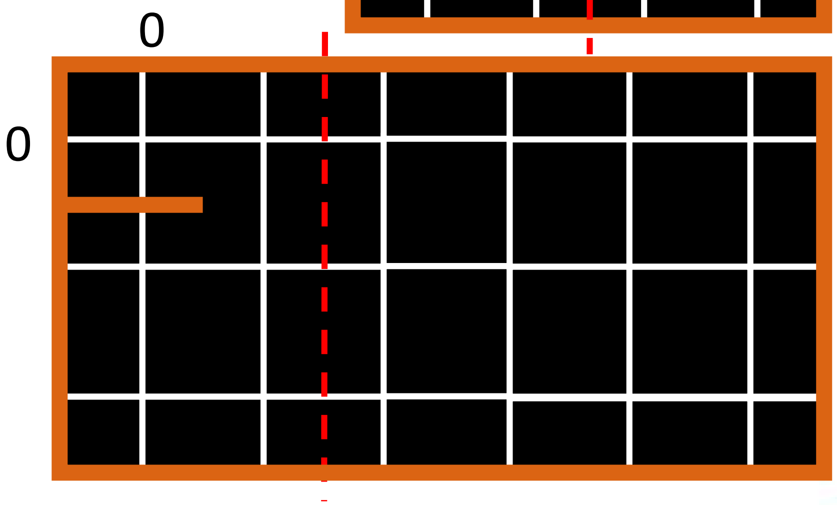

Bayesian Statistics

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

I saw a wall to the north, am I to the right or the left of the red line?

- Scenario #1:
- $P(right|wall_{NORTH}) = ??$
- Known:
  - Joint probability that I am on the right and there is a north wall:
    - $P(wall_{NORTH} \cap right) = P(right) * P(wall_{NORTH} \mid right)$
  - The opposite:
    - $P(right \cap wall_{NORTH}) = P(wall_{NORTH}) * P(right \mid wall_{NORTH})$
  - $P(right \cap wall_{NORTH}) = P(wall_{NORTH} \cap right)$
  - $P(right \mid wall_{NORTH}) = P(right) * P(wall_{NORTH} \mid right) / P(wall_{NORTH})$
  - $P(right \mid wall_{NORTH}) = 0.5*0.33/0.42 = 0.4$
- Scenario #2: $p(right \mid wall_{NORTH}) = 0.67*0.33/0.42 = 0.53$

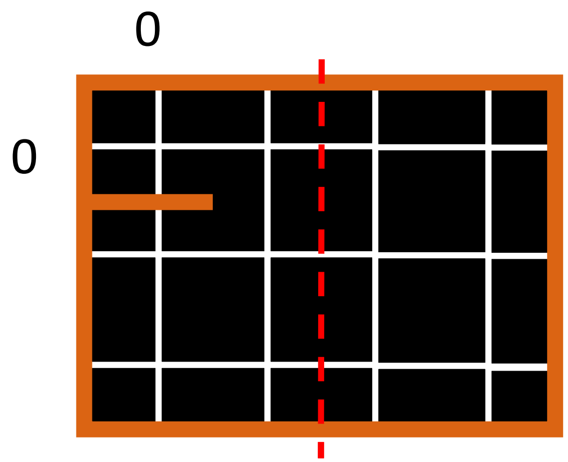

Bayesian Statistics

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

I saw a wall to the north, am I to the right or the left of the red line?

- Scenario #1:
  - $P(\text{right} \mid \text{wall}_{NORTH}) = P(\text{right}) * P(\text{wall}_{NORTH} \mid \text{right}) / P(\text{wall}_{NORTH})$

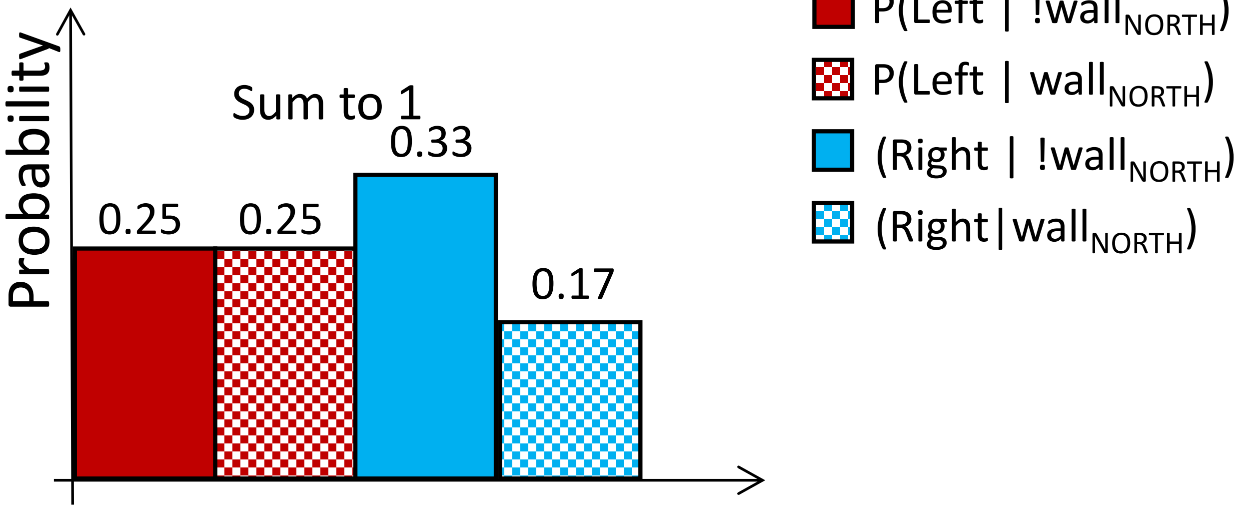

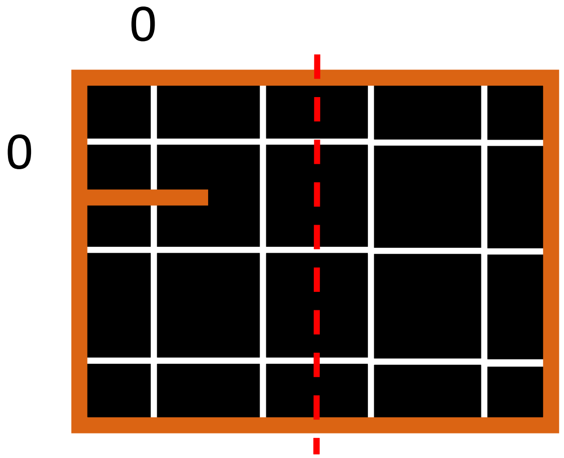

Bayesian Statistics

- Educated guesses based on probability distributions, to update beliefs in the evidence of new data
- Representing all parameters as probability distributions

likelihood

prior

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

conditional probability
posterior

marginal likelihood
(constant)

Bayesian Statistics

- Educated guesses based on probability distributions, to update beliefs in the evidence of new data
- Representing all parameters as probability distributions

What we want — *what we have* — *probability of being in this location*

$$P(x|y) = \frac{P(y|x) \cdot P(x)}{P(y)}$$

how likely the measurement is

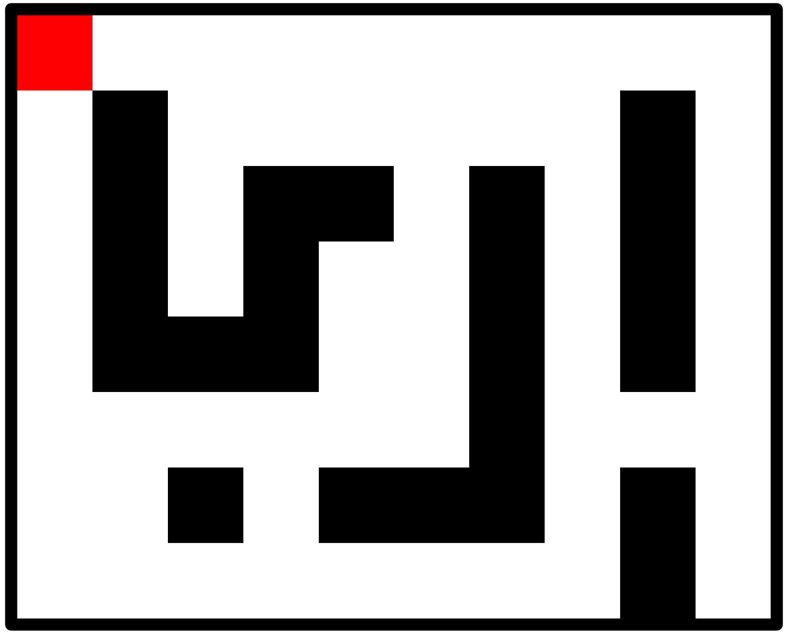

X is the set of possible locations
x is one of these locations
y is the sensor measurement

Bayesian Statistics

- Educated guesses based on probability distributions, to update beliefs in the evidence of new data
- Representing all parameters as probability distributions

What we want — what we have — probability of being in each location

$$P(X|y) = \frac{P(y|X) \cdot P(X)}{P(y)}$$

how likely the measurement is

$\rightarrow 1/P(y) = \alpha$

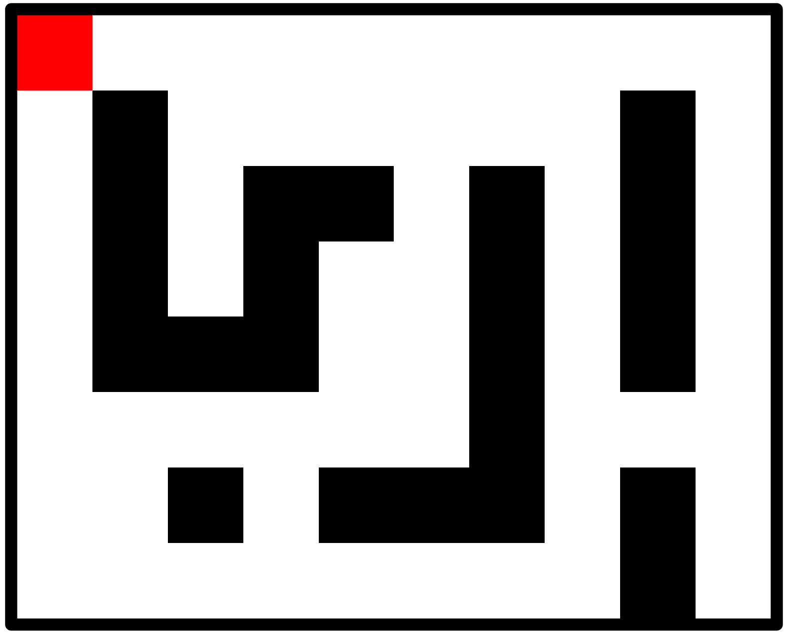

X is the set of possible locations

x is one of these locations

y is the sensor measurement

Robot Motion

- **Transition model**
  - No matter what I tell my robot to do, it makes a random move or stays!

Transition Matrix:

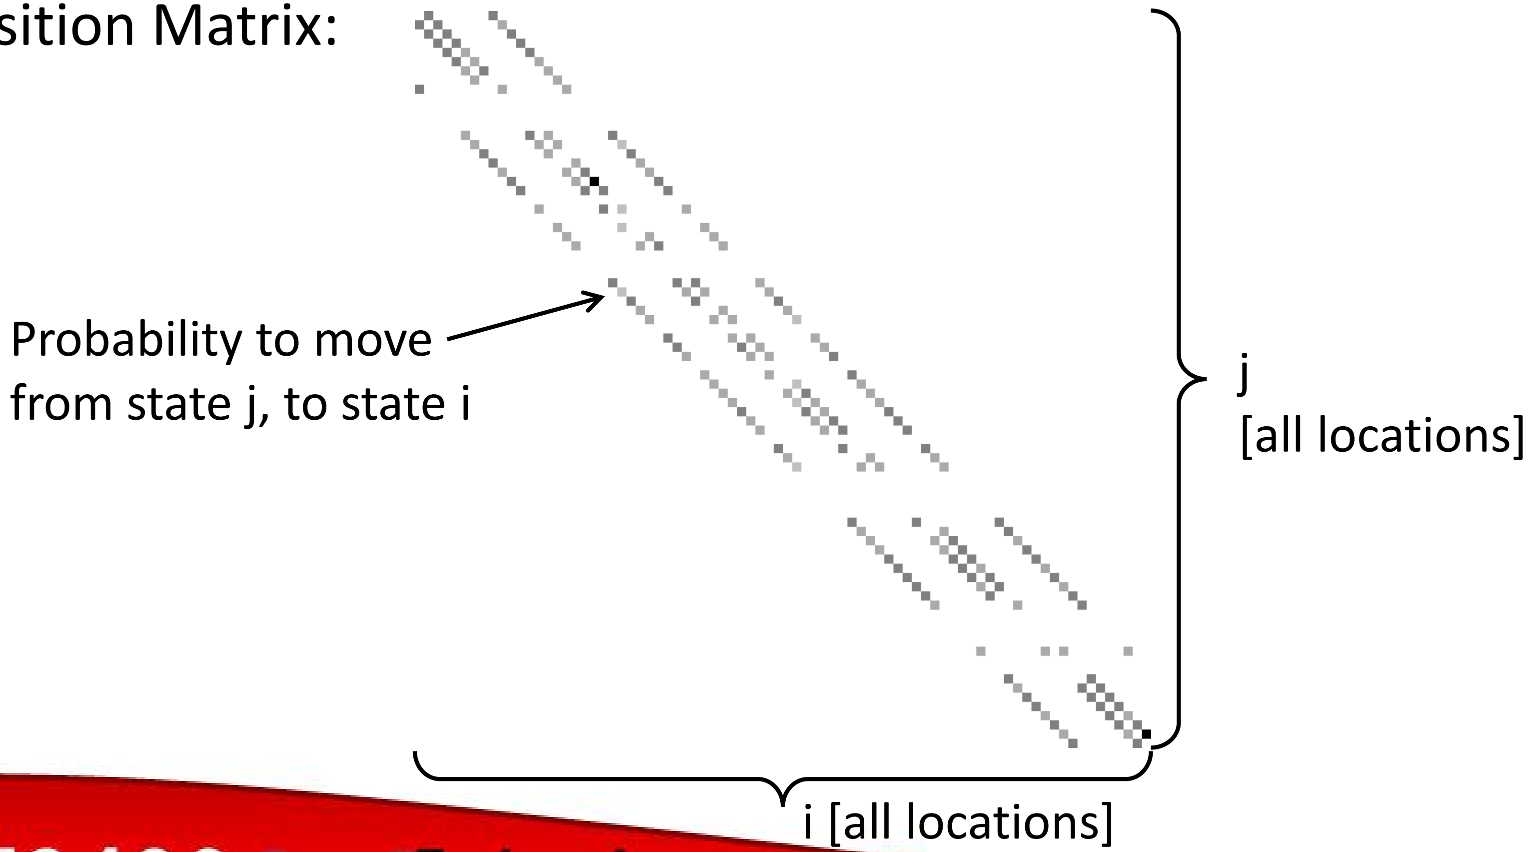

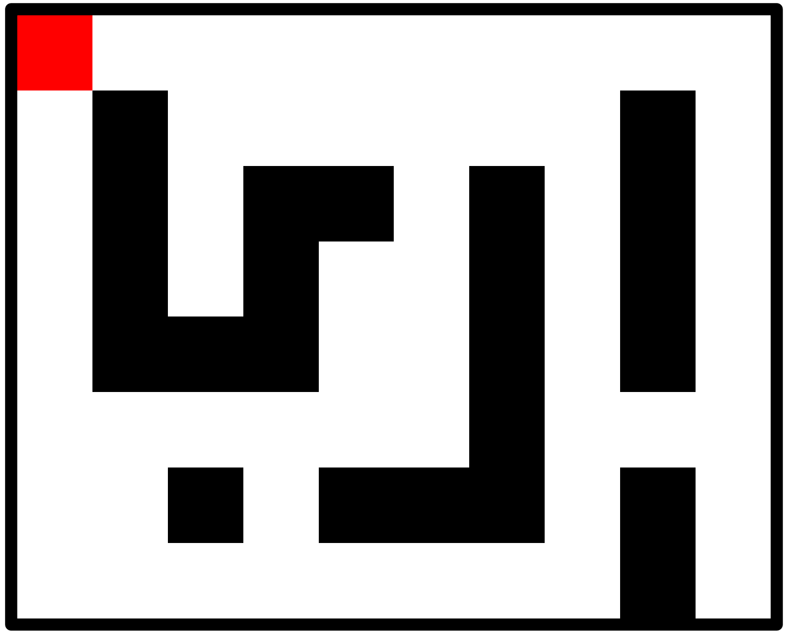

X is the set of possible locations
x is one of these locations
y is the sensor measurement

Robot Motion

- **Transition model**
  - No matter what I tell my robot to do, it makes a random move or stays!

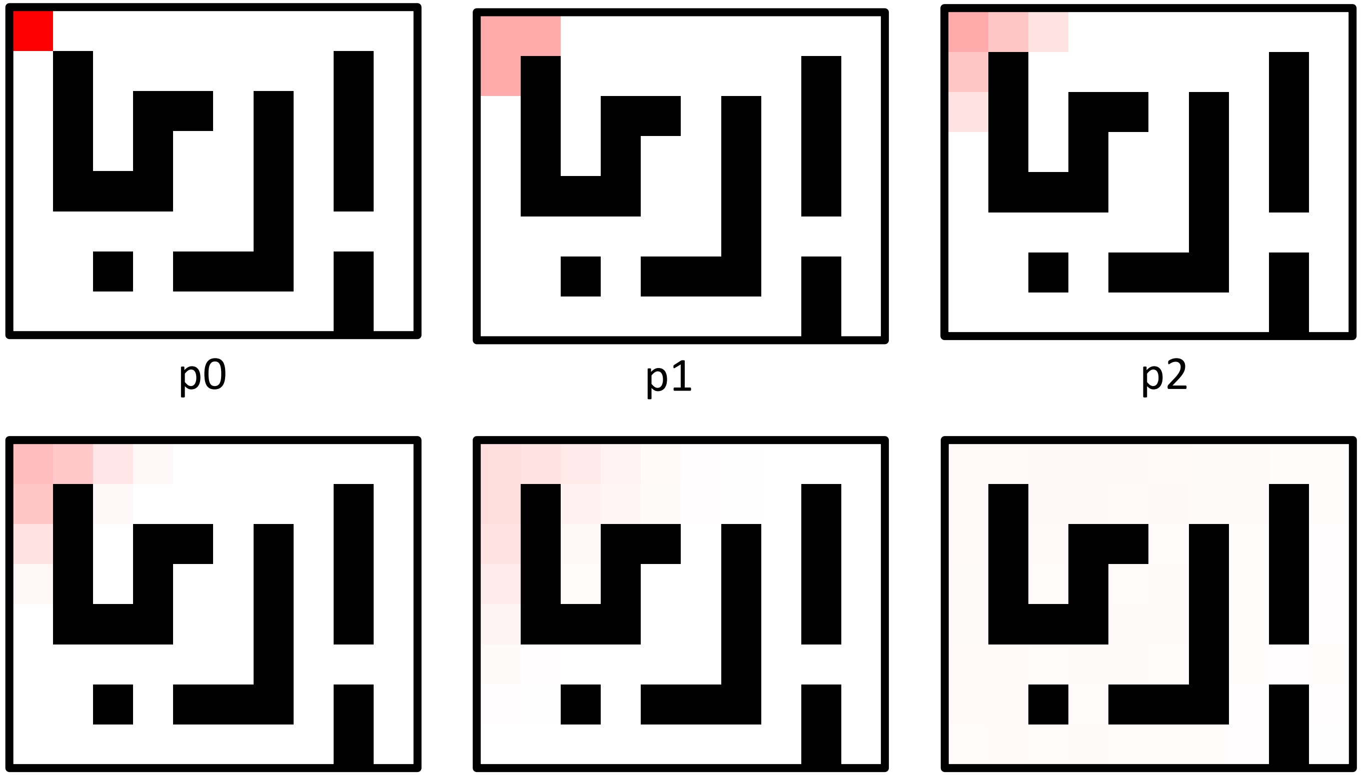

p0 p1 p2 p3 … p10 … p100

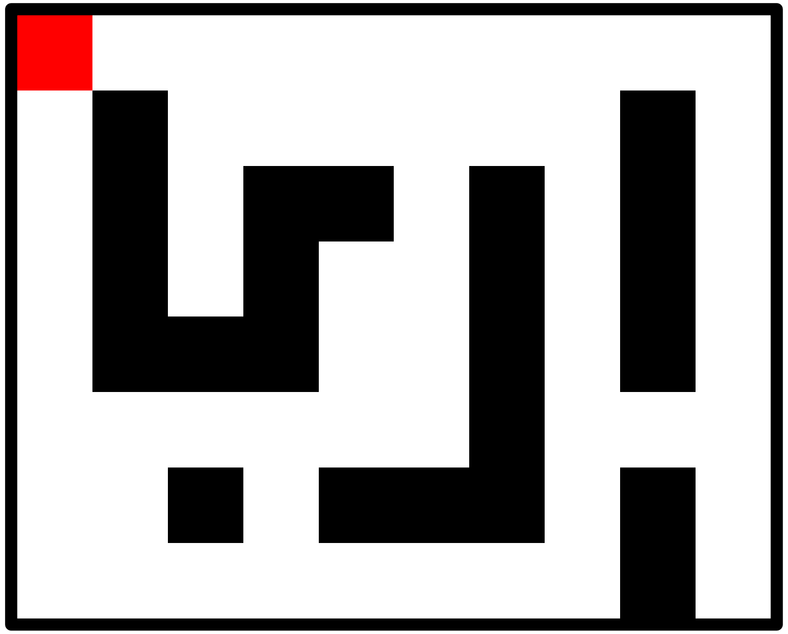

X is the set of possible locations
x is one of these locations
y is the sensor measurement

Robot Motion

- **Transition model**
  - The robot may not know where it is, but it *does* have a physical state

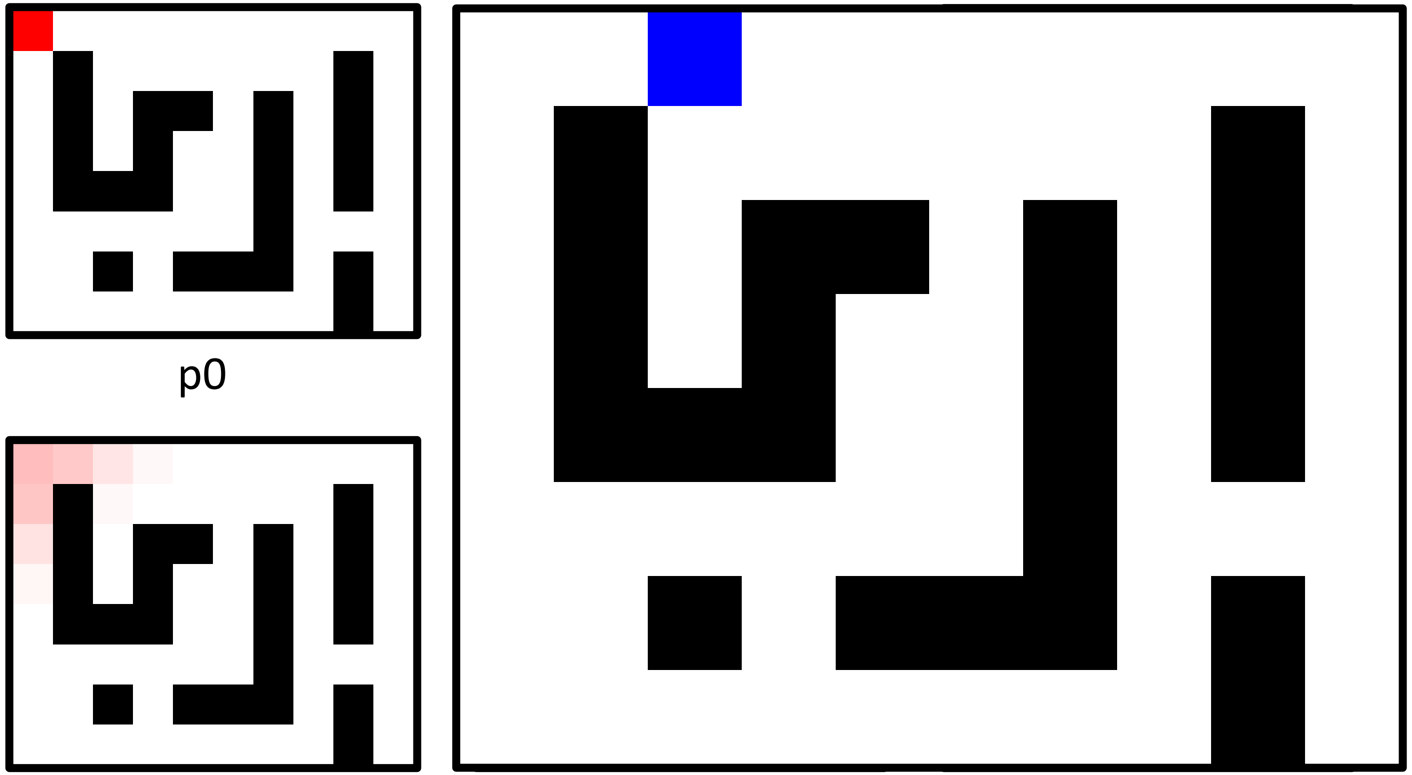

p0 p3 … p10 … p100

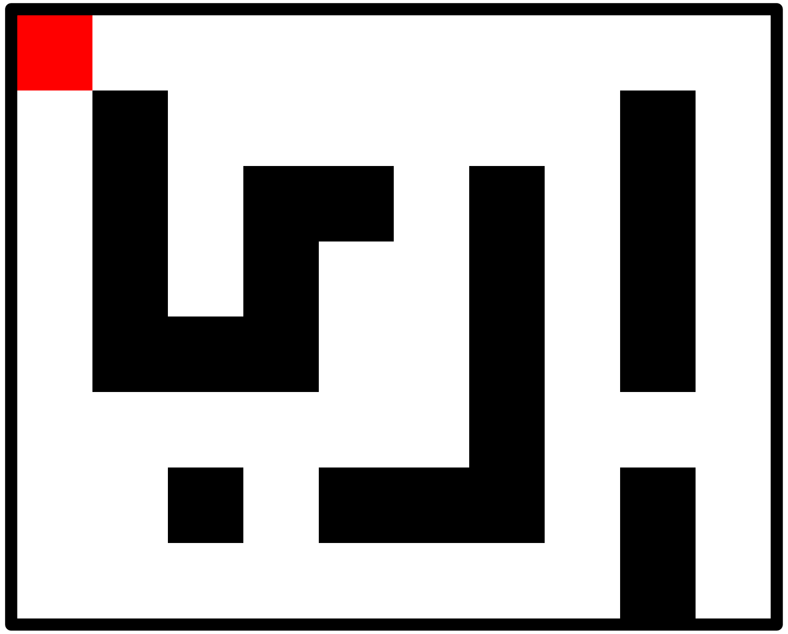

X is the set of possible locations
x is one of these locations
y is the sensor measurement

Robot Motion

- **Transition model**
  - The robot may not know where it is, but it *does* have a physical state

- North = Wall
- East = No Wall
- South = No Wall
- West = No Wall

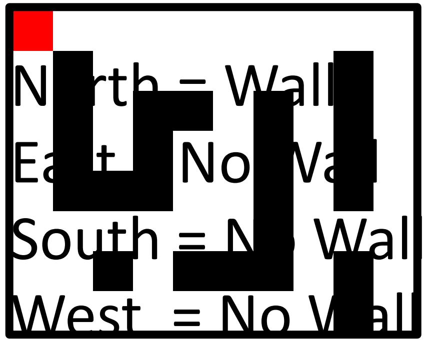

p0

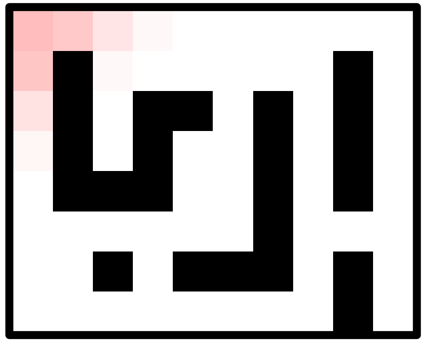

p3

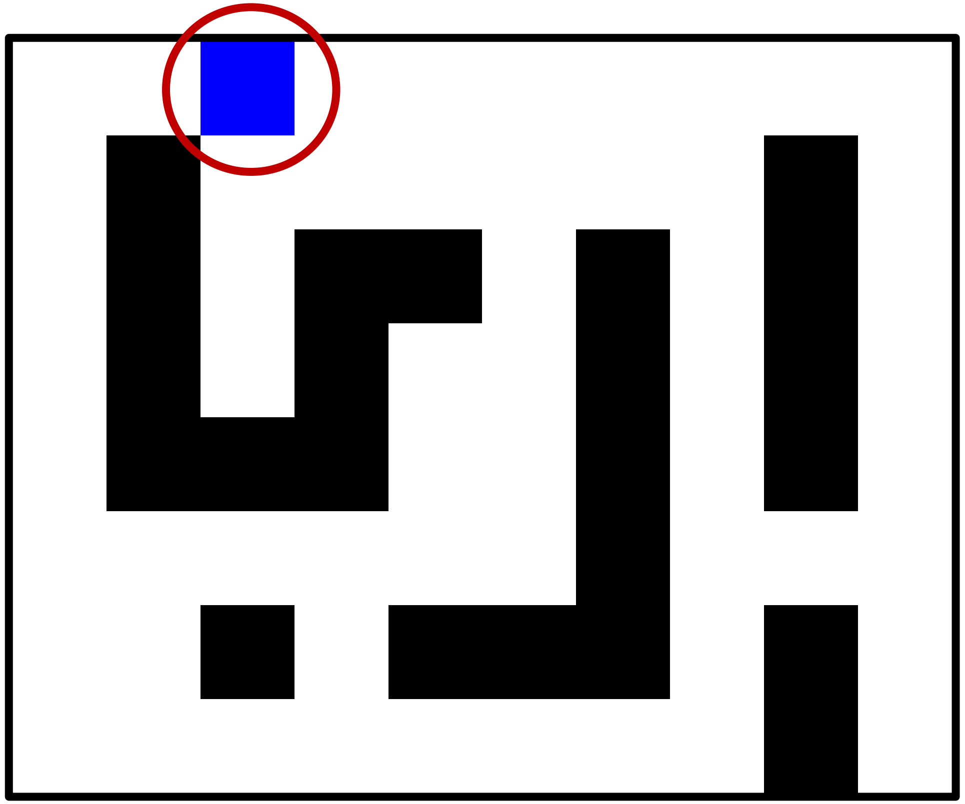

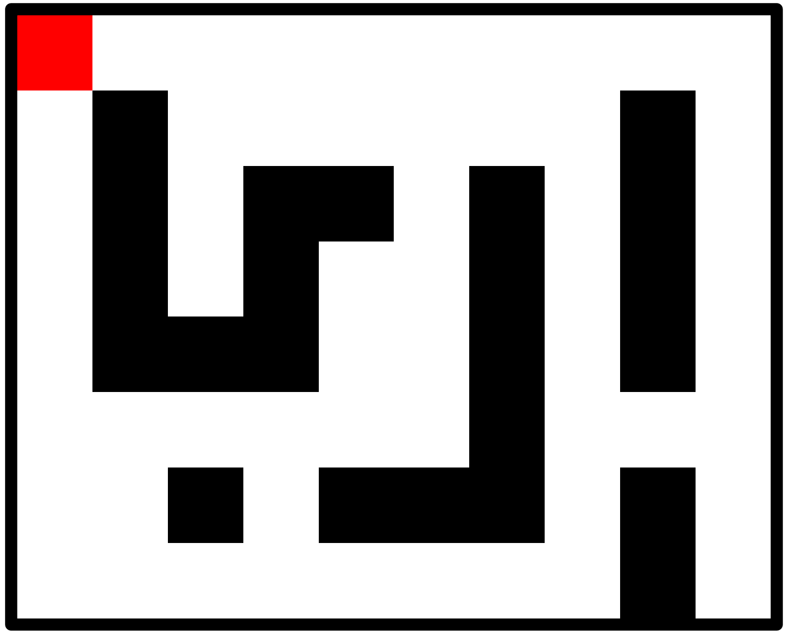

X is the set of possible locations

x is one of these locations

y is the sensor measurement

Wall Sensor

- **Sensor model**
  - Correct 90% of the time
  - (10% misses walls; 10% sees walls that aren't there)

- Sensor output:
- [North, East, West, South]

- P(no walls | x) = 0.1*0.9*0.9*0.1
- P(N | x) = 0.9*0.9*0.9*0.1
- P(W | x) = 0.1*0.9*0.9*0.9
- P(S | x) = 0.1*0.9*0.1*0.1
- P(E | x) = 0.1*0.9*0.1*0.1
- …
- P(NW | x) = 0.9*0.9*0.9*0.9

highest probability

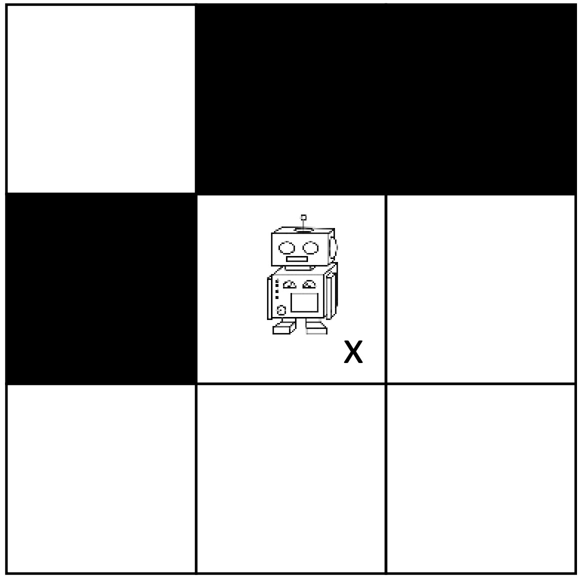

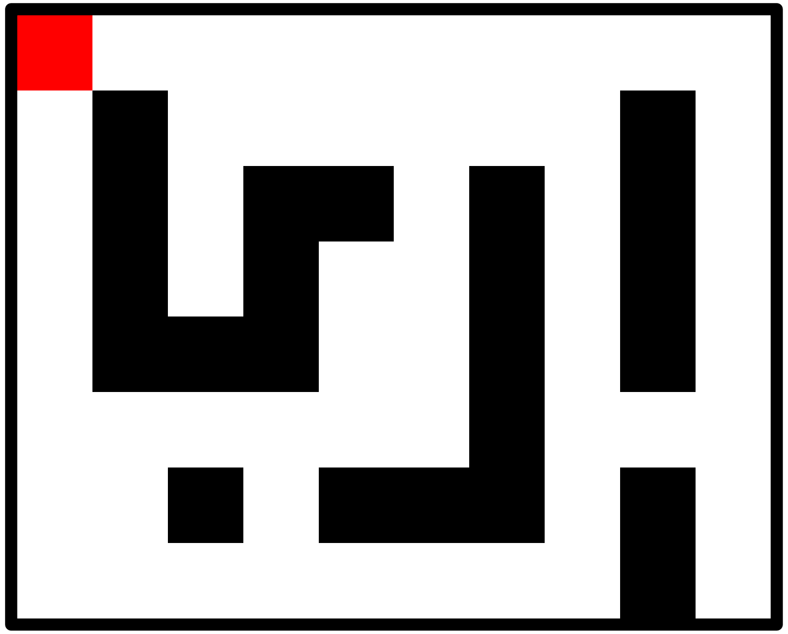

X is the set of possible locations
x is one of these locations
y is the sensor measurement

Wall Sensor

- **Sensor model**
  - Correct 90% of the time
  - …Compute the likelihood of an observation from each state, P(y|X)

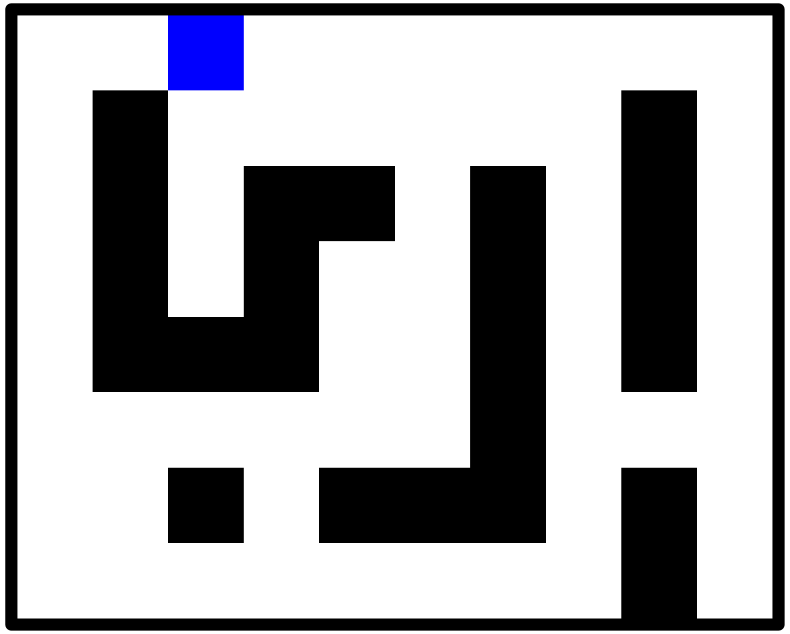

a robot state (x)

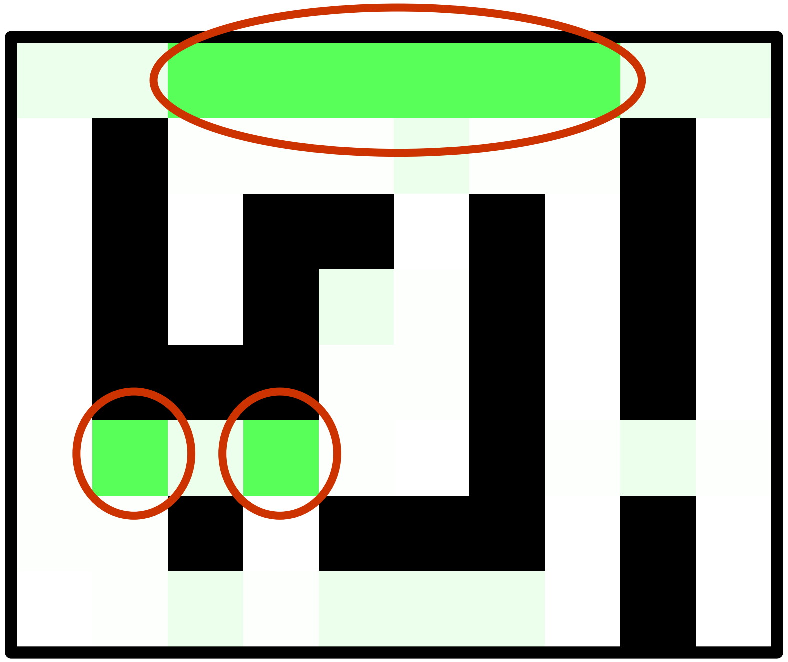

P(y|X)

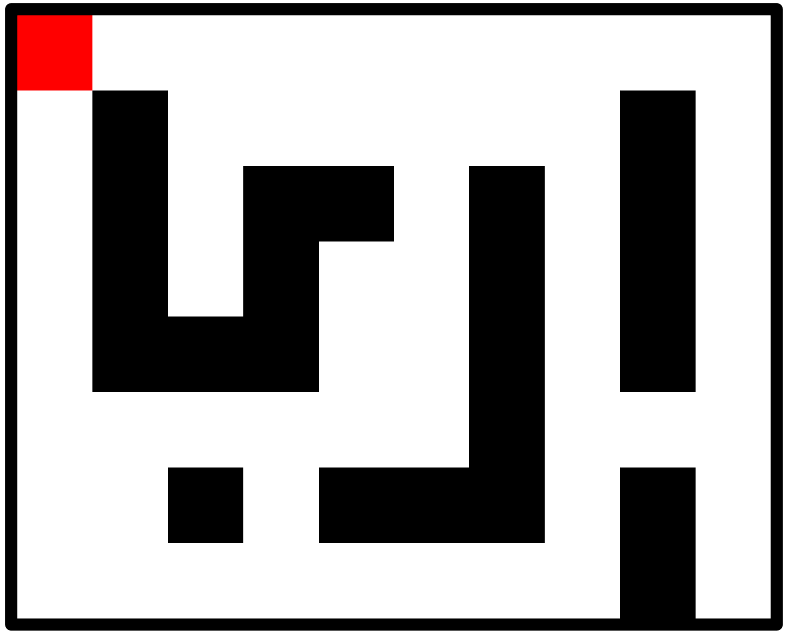

X is the set of possible locations
x is one of these locations
y is the sensor measurement

Combine the Motion and Sensor Models

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

- α: Normalize if necessary
- $P(y_{t+1} | X_{t+1})$: **Correct prediction** by weighing in the measurement
- $P(X_{t+1} | x_t)$: **Transition Model** Predict probability distribution after action
- $P(x_t | y_{1:t})$: **Previous State Estimate**

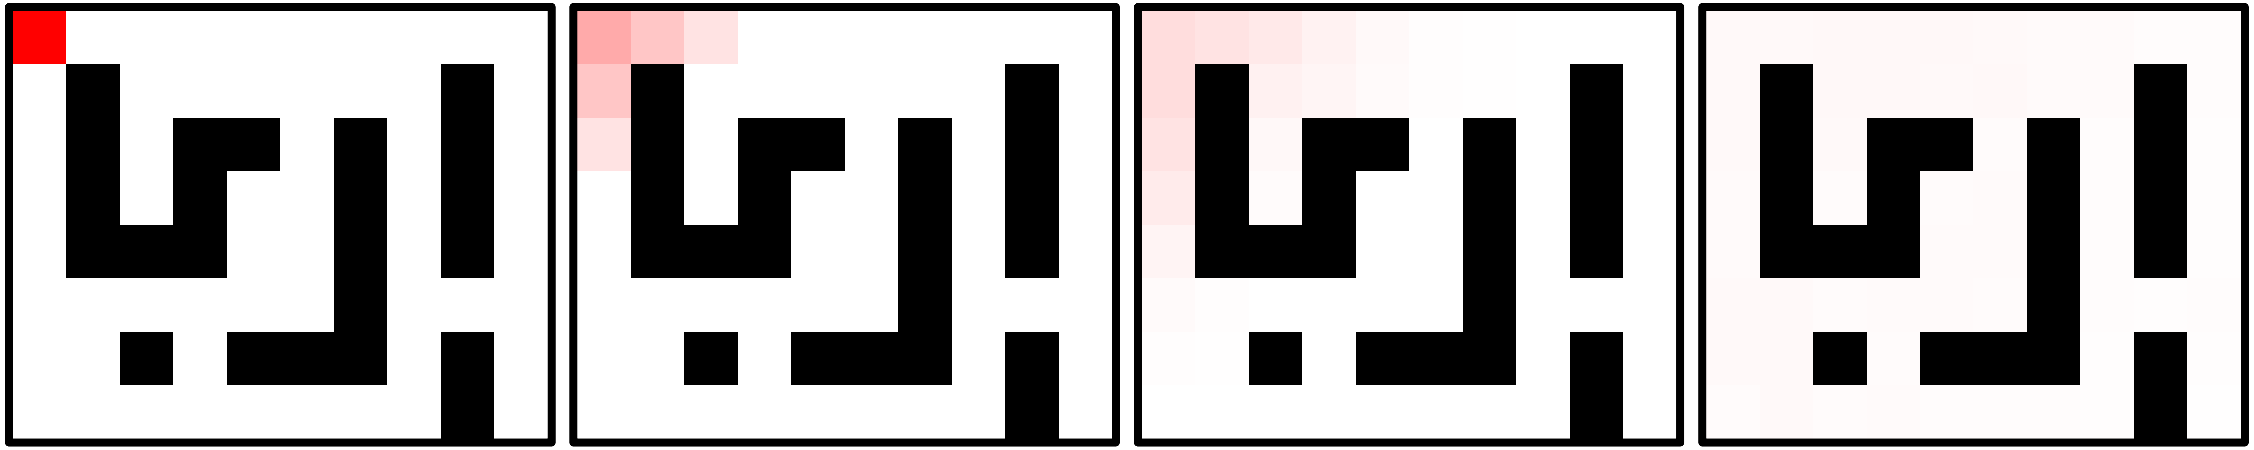

No Observations

Combine the Motion and Sensor Models

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

Normalize if necessary

Correct prediction by weighin measure

Transition Model redict probability ribution after action

Previous State Estimate

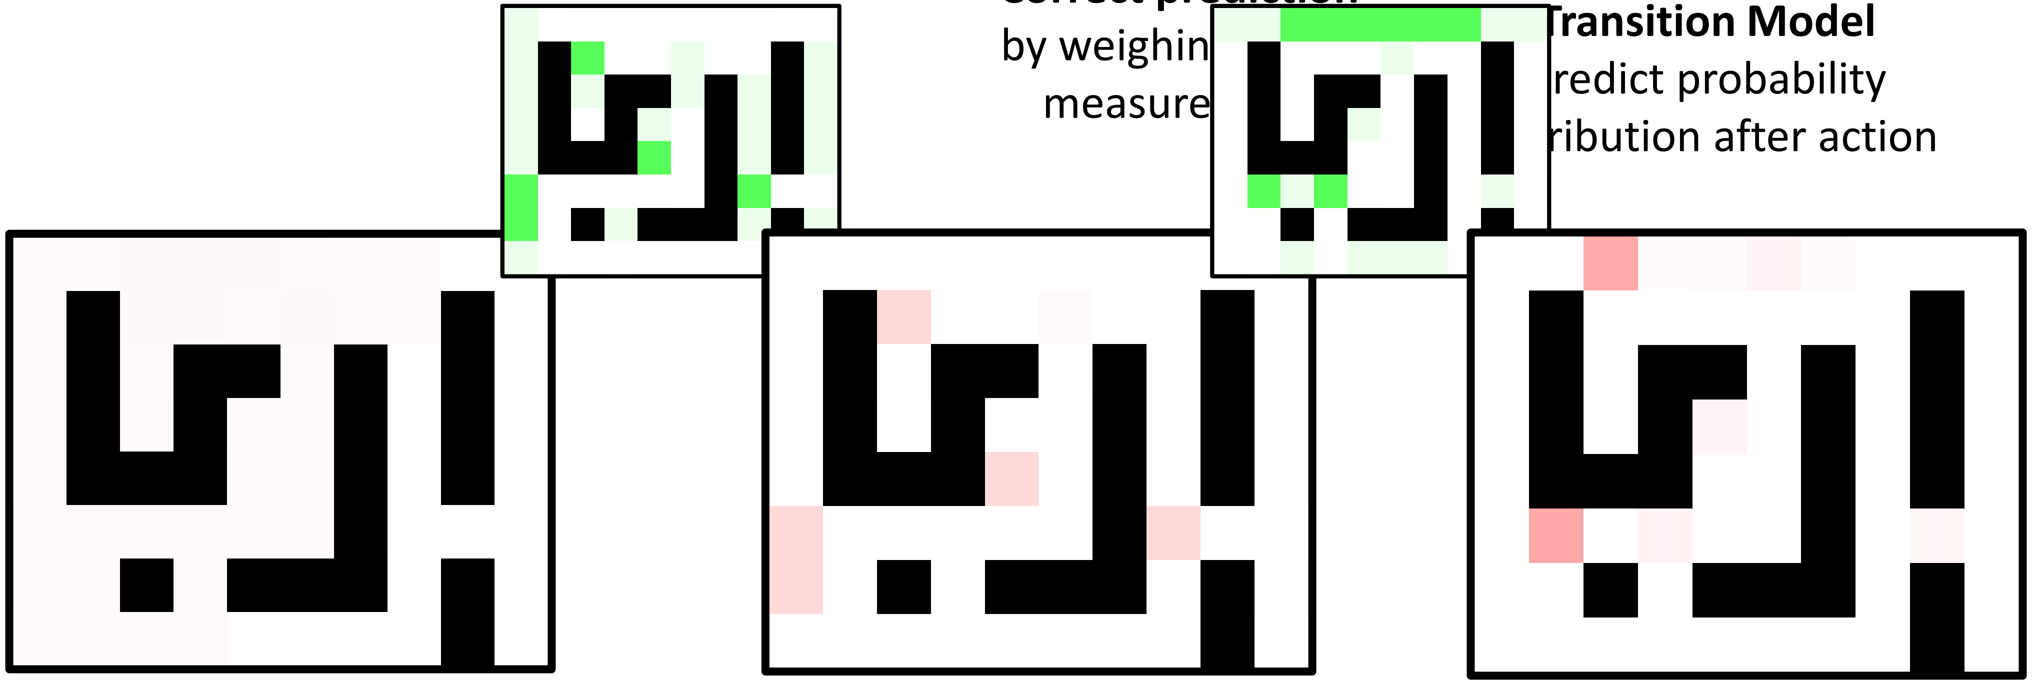

In two steps, we homed in on where we are!

...

Observations

Combine the Motion and Sensor Models

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

- α: Normalize if necessary
- $P(y_{t+1}|X_{t+1})$: **Correct prediction** by weighing in the measurement
- $P(X_{t+1}\,|x_t)$: **Transition Model** Predict probability distribution after action
- $P(x_t\,|y_{1:t})$: **Previous State Estimate**

Can you do better?

Improved Transition Model

$$P(X_{t+1}|\, y_{1:t+1}) = \alpha\, P(y_{t+1}|X_{t+1}) \sum_{xt} P(X_{t+1}\,|x_t, u_{t-1})\, P(x_t\,|y_{1:t})$$

- u_{t-1}: **Factor in Input** Predict probability distribution after deliberate action

How would you actually do that?

Combine the Motion and Sensor Models

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

α: Normalize if necessary

$P(y_{t+1}|X_{t+1})$: **Correct prediction** by weighing in the measurement

$P(X_{t+1}\,|x_t)$: **Transition Model** Predict probability distribution after action

$P(x_t\,|y_{1:t})$: **Previous State Estimate**

Improved Transition Model

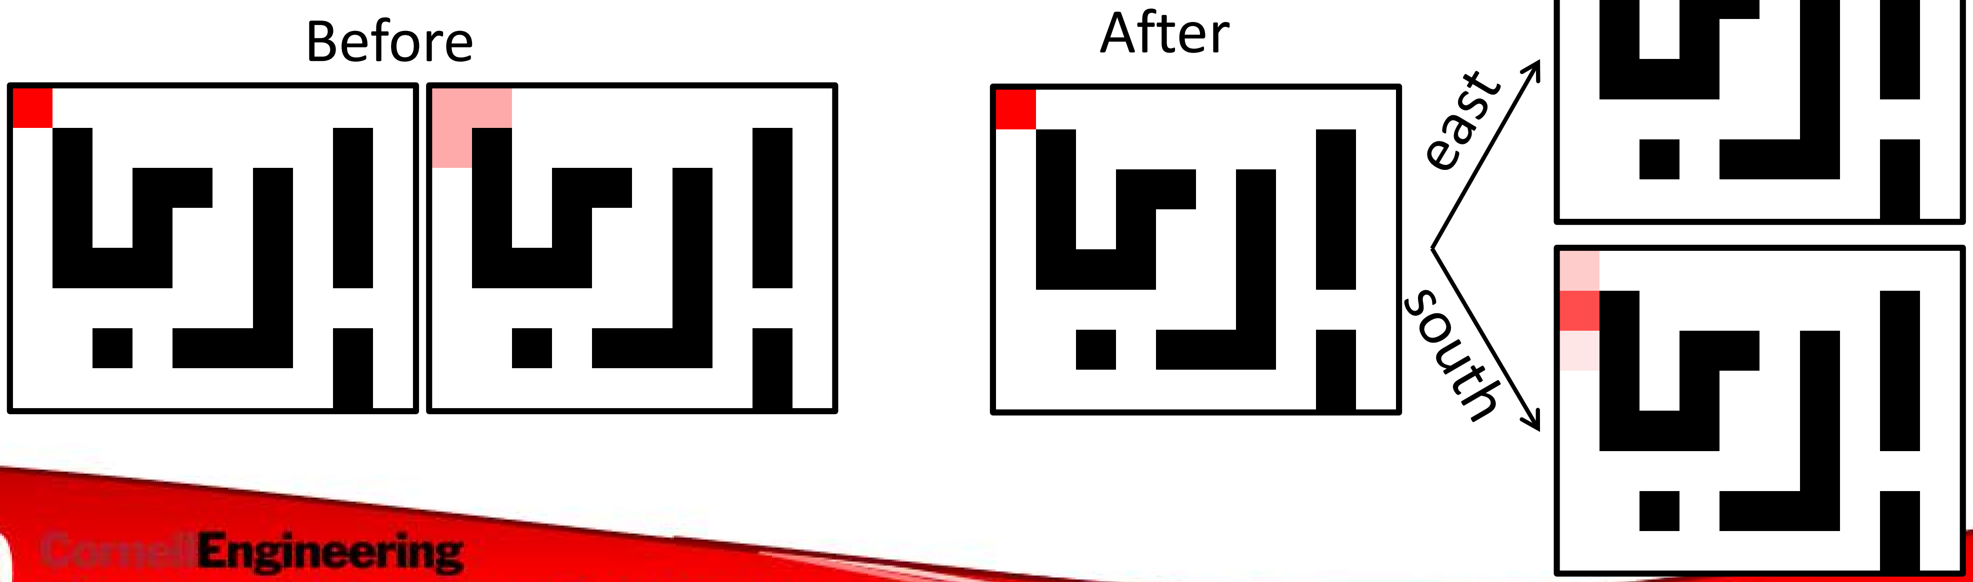

Combine the Motion and Sensor Models

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

Normalize if necessary (α)

Correct prediction by weighing in the measurement ($P(y_{t+1}|X_{t+1})$)

Transition Model Predict probability distribution after action ($P(X_{t+1} | x_t)$)

Previous State Estimate ($P(x_t | y_{1:t})$)

Improved Transition Model

What else could you do to localize faster?

- Deliberately move in directions that give you more information

Summary

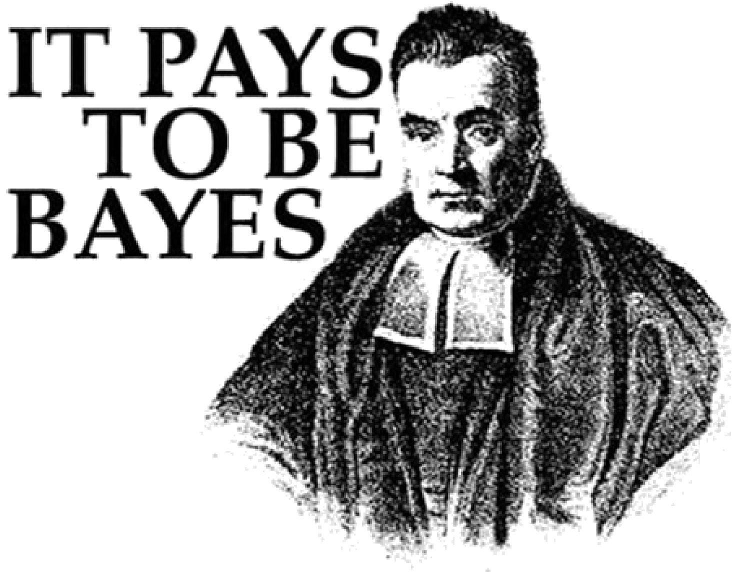

- Use temporal consistency between observations that are poor estimates individually
- Localization can work with…
    - …completely random motion
    - …noisy sensors
    - (returns a probability of where you are)
- This general approach works with more complicated, states, observation models, and transition models.

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

New estimate = Factor in observation Predict Old estimate

Go Build Robots!