Probability and Bayes Theorem

Fast Robots, ECE 4160/5160, MAE 4190/5190

E. Farrell Helbling, 3/3/26
Slides adapted from Prof. Kirstin Petersen

Lab 5 - Linear PID

- Control your car to end a specific distance from the wall, reliably
  - My advice, go really slowly to start
  - Let’s just think about a good guess for kp to start

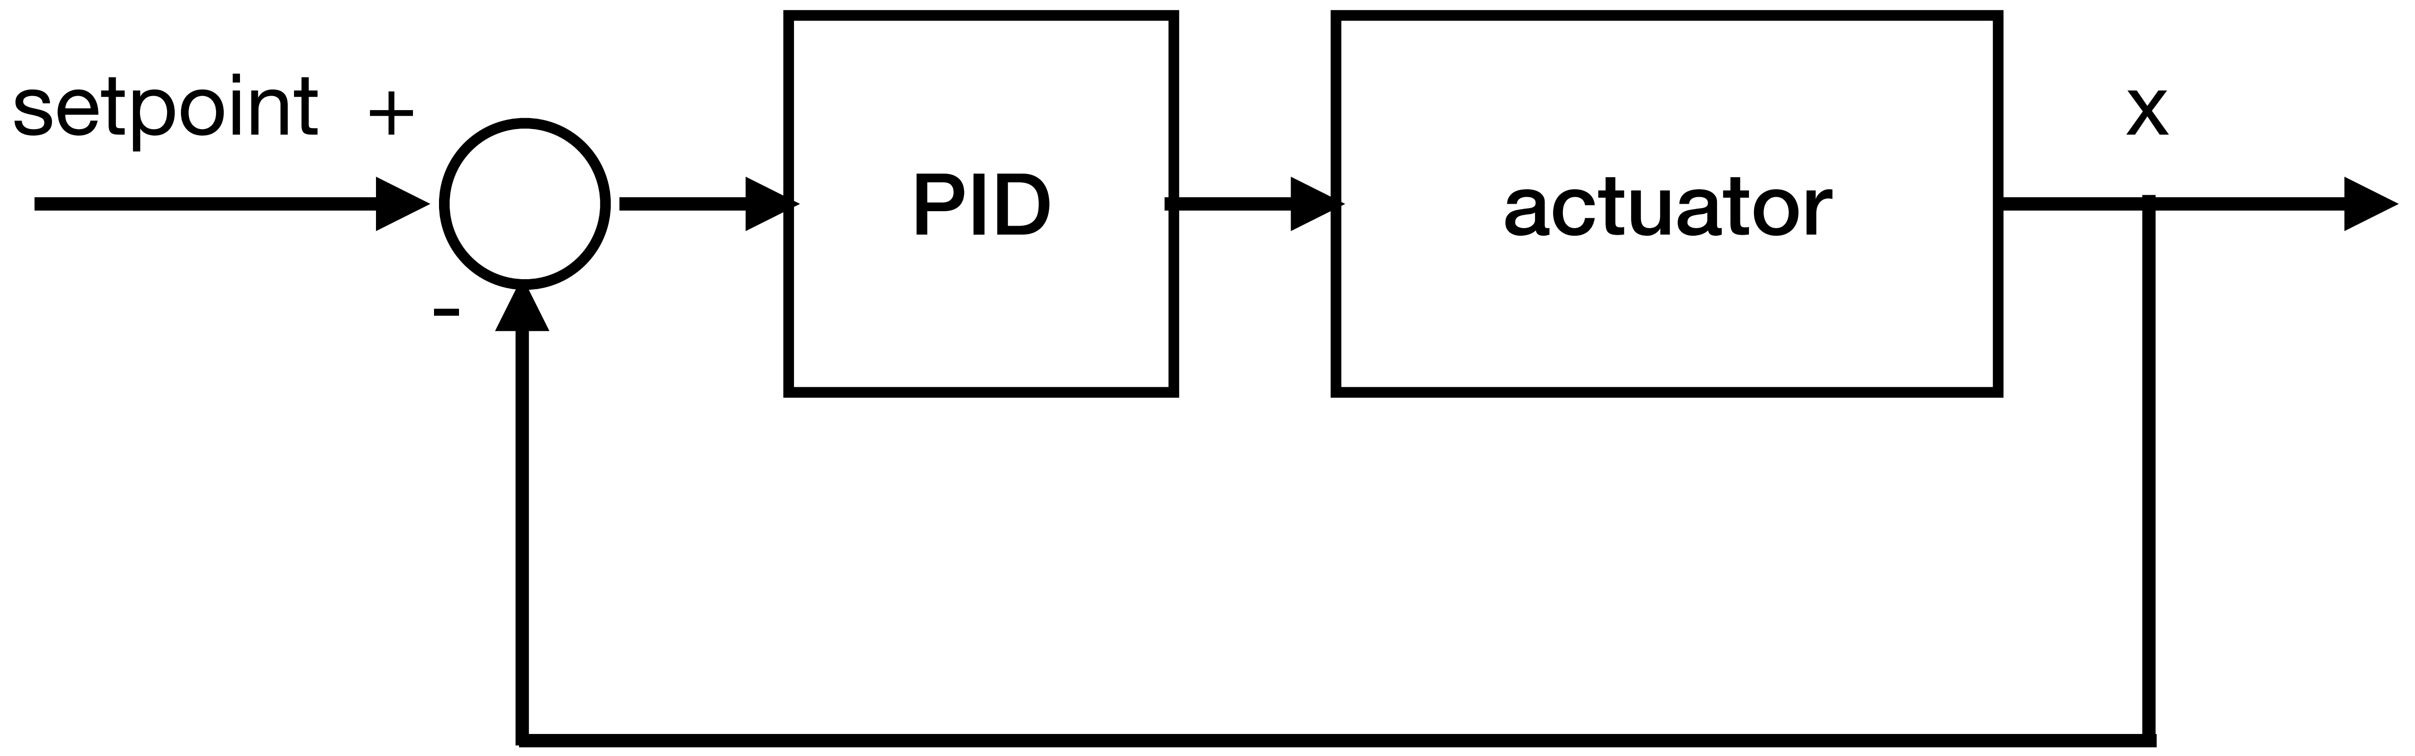

Lab 5 - Linear PID

- Great example from two years ago: https://fast.synthghost.com/lab-5-linear-pid-control/ from Stephan Wagner. You can breeze past his program organization and just get to the lab tasks.

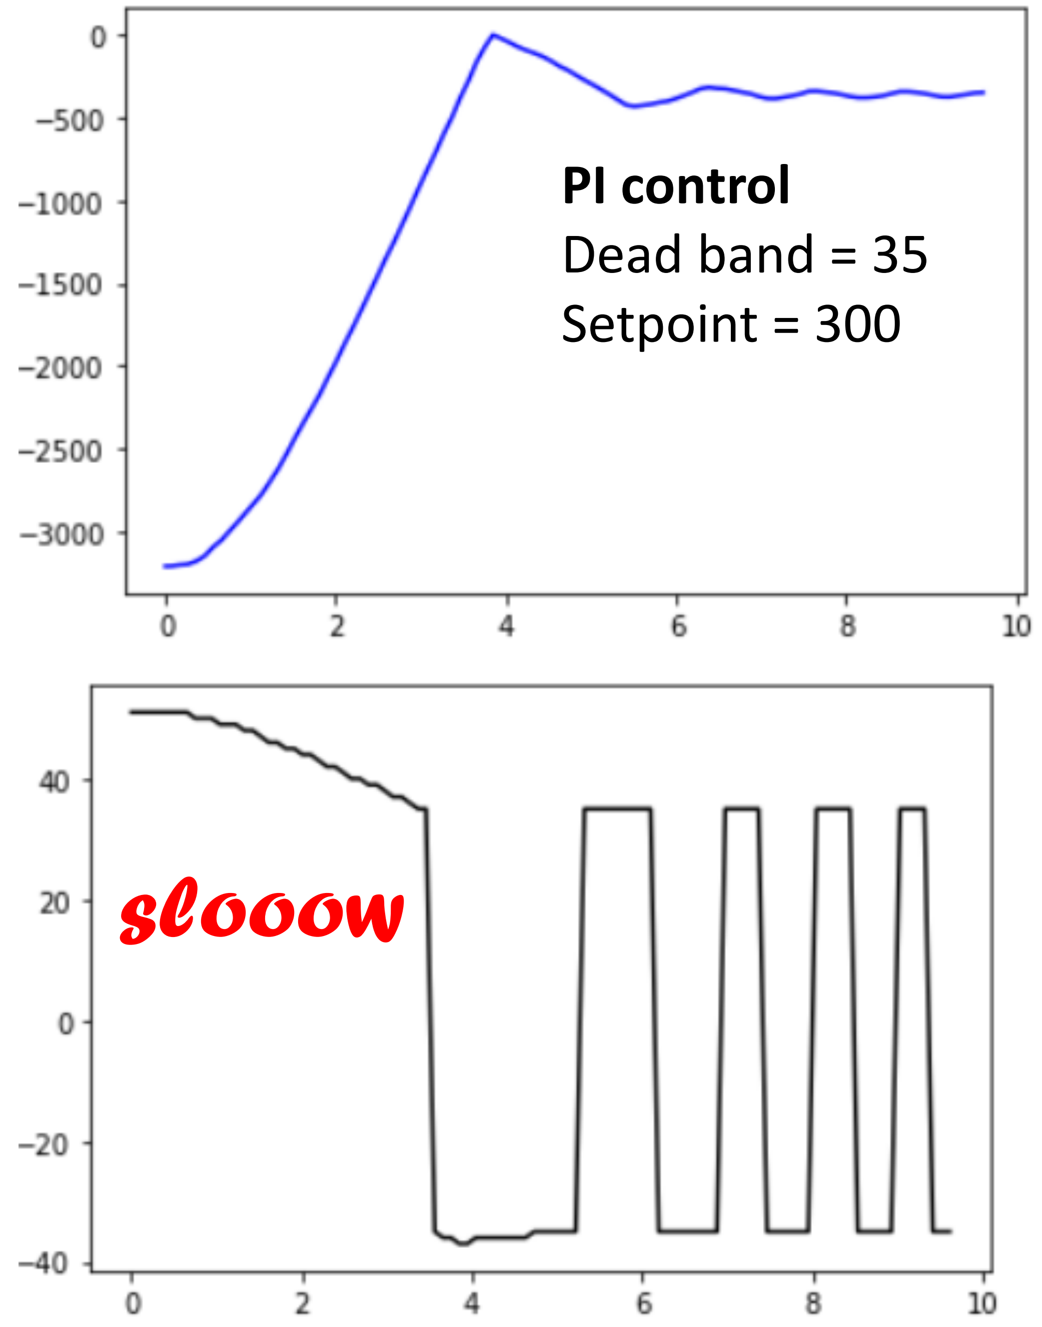

Lab 5 — Linear PID

- Do not write your code such that you approach the wall and stop your motors once you get within a certain distance, this is not really feedback control.
- A good test is if i push your car closer to/ further from the wall, will it return to the set point. **This year we ask you to do this.**
- Watch out for deadband
- Watch out for max PWM limits

Recap

ECE 3100 Intro to Probability and Inference

- Random variable
  - $X : \Omega \rightarrow \mathbb{R}$
- The probability that the random variable X has value x
  - $P(X = x)$ or $p(x)$
- Probabilities sum to 1
  - $\Sigma_x P(X = x) = 1$
- Probabilities are always greater than 0
  - $P(X = x) \geq 0$
- Joint distribution Y
  - $p(x, y) = P(X = x \text{ and } Y = y)$
- Conditional probability
  - $p(x \mid y) = \dfrac{p(x, y)}{p(y)}$

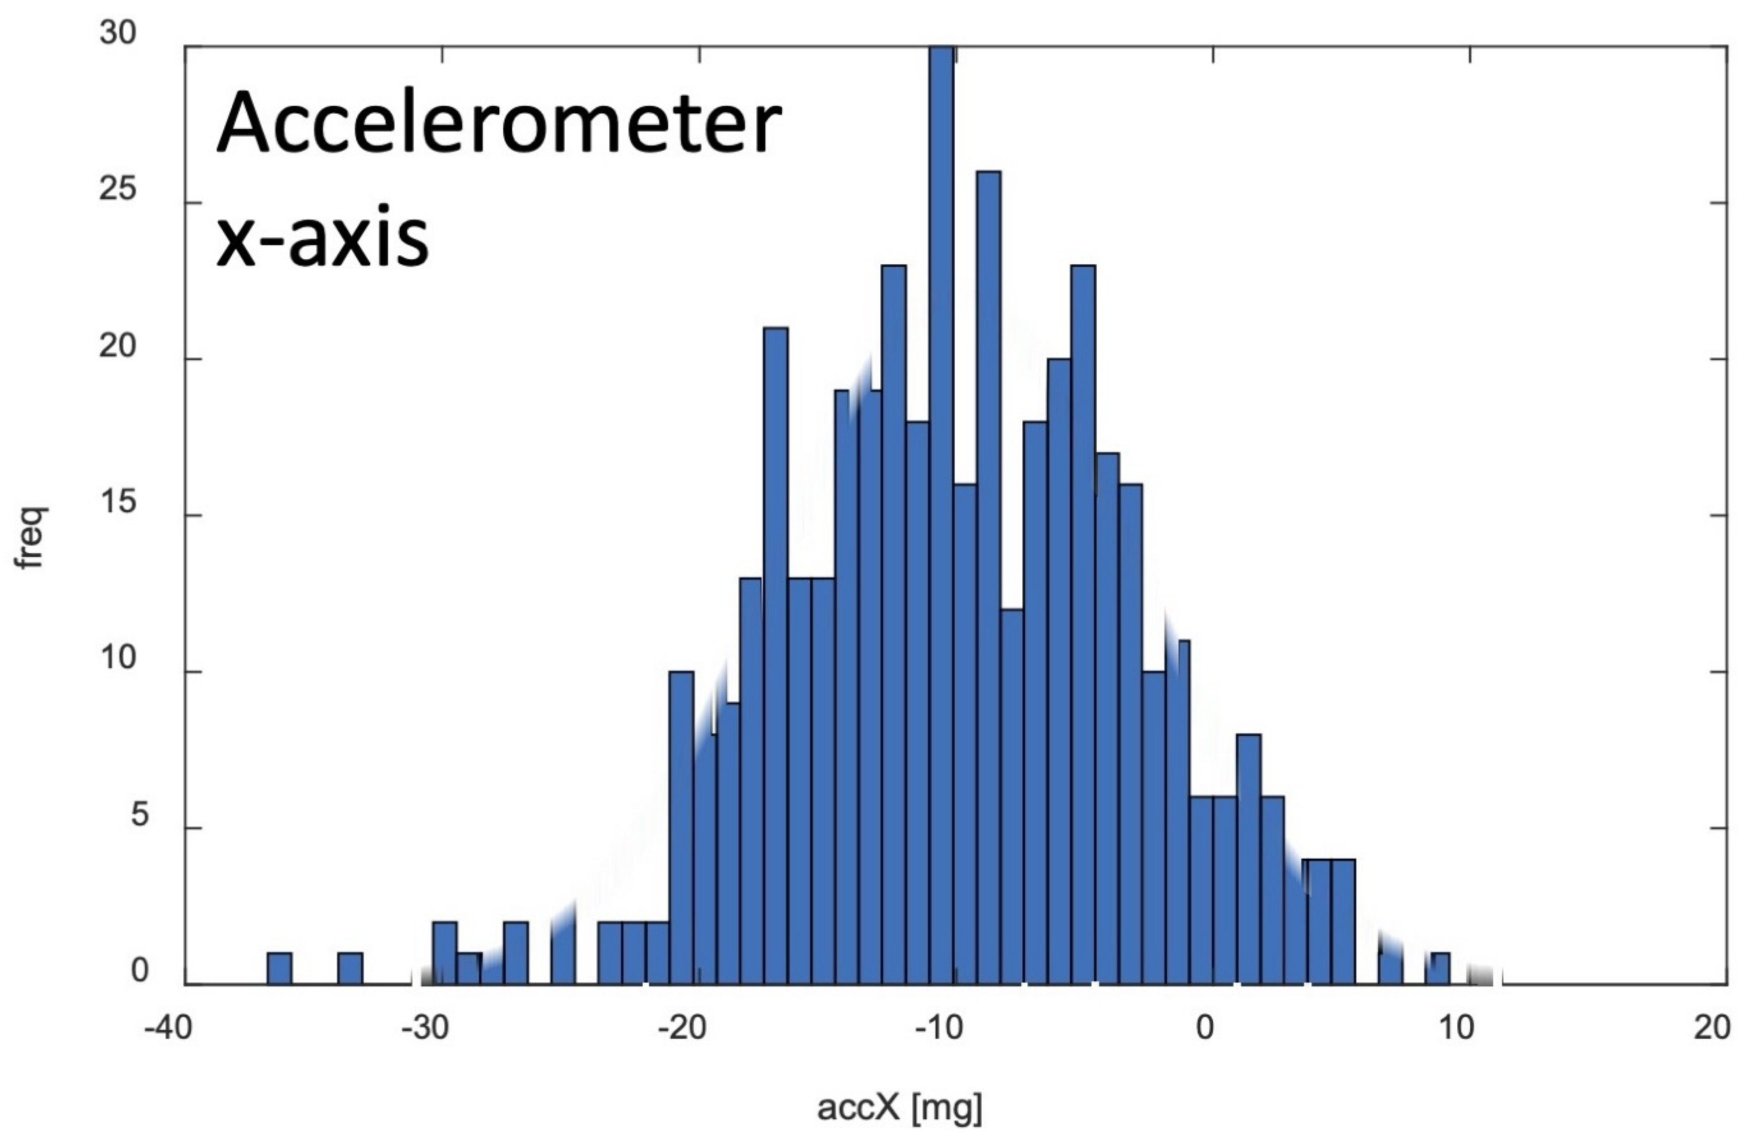

- Mean
  - $(\mu = -9.97306\text{mg})$
- Std Dev
  - $(\sigma = 7.0318\text{mg})$
- Variance (σ^2)
- Gaussian Distributions
  - $[\mu \pm \sigma]$
  - Symmetric
  - Unimodal
  - Sum to “unity”

Conditional Probability

- $p(x \mid y) = \dfrac{p(x, y)}{p(y)}$
- Robot/ sensor example
- Exercise
  - Left-handedness has a probability of 10%. A family has two children, the first is left-handed, what is the probability that the second is left-handed?
    - 10%
  - The family has two children Given that at least one child is left handed, what is the probability that both are?
    - 5% p(both)/p(at least one)

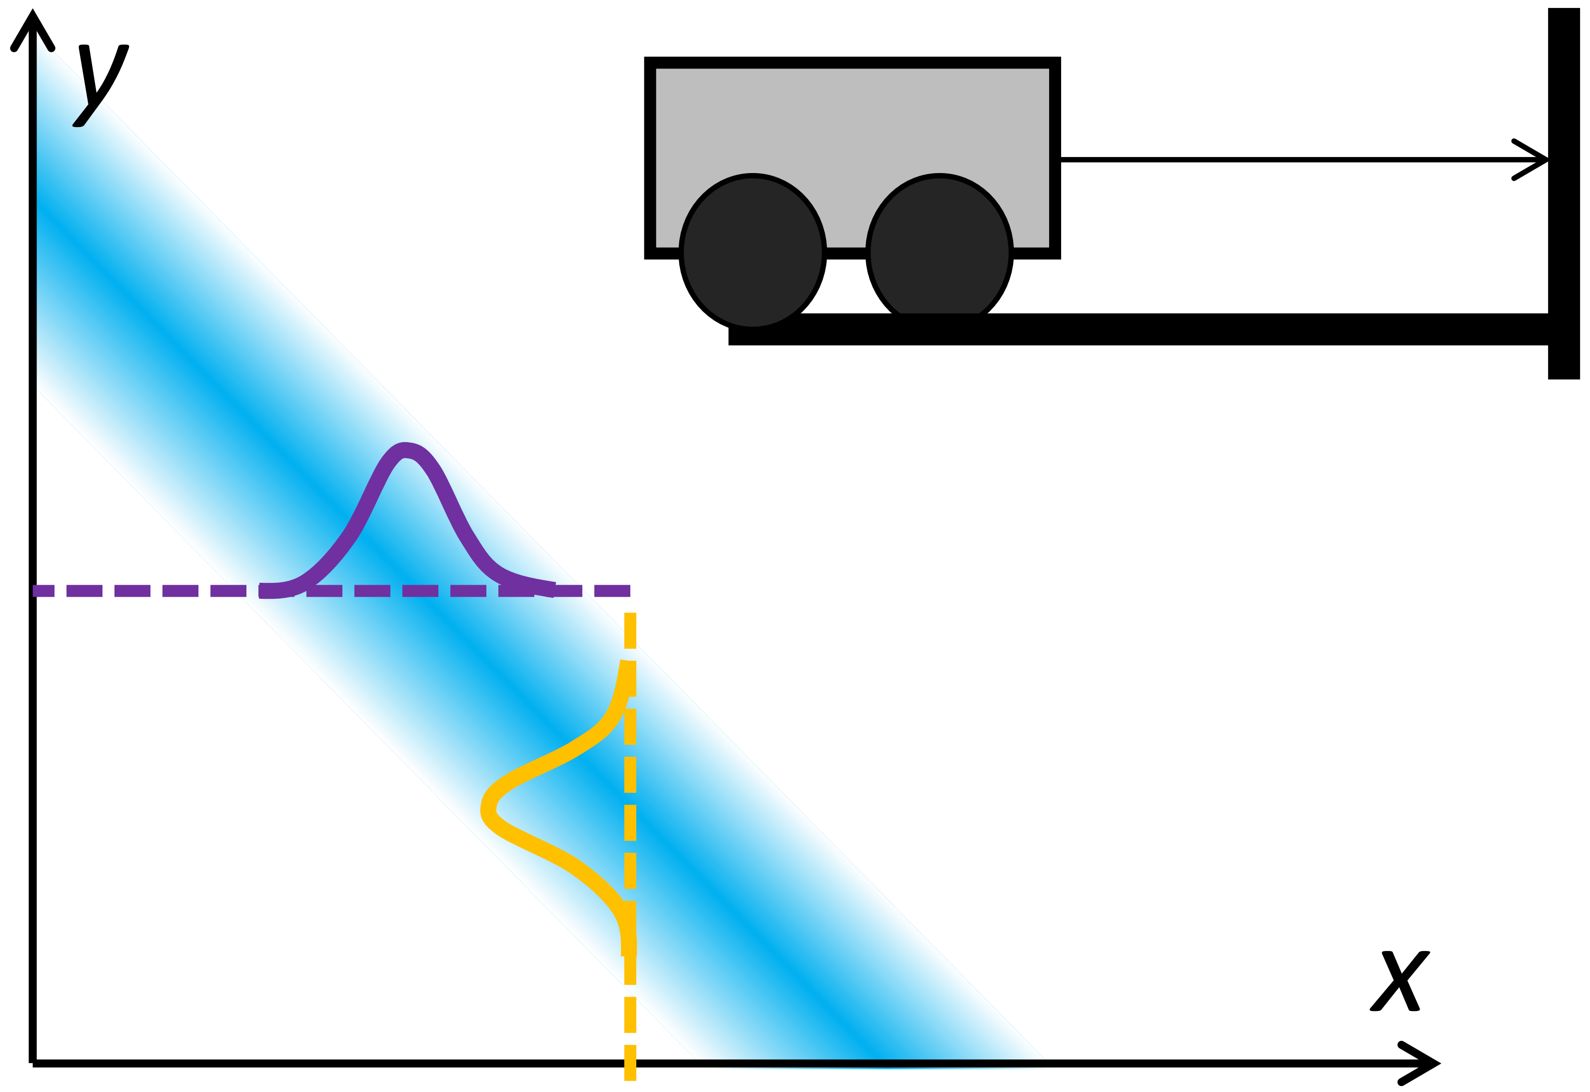

Recap

ECE 3100 Intro to Probability and Inference

- Random variable
  - $X : \Omega \rightarrow \mathbb{R}$
- The probability that the random variable X has value x
  - $P(X = x)$ or $p(x)$
- Probabilities sum to 1
  - $\Sigma_x P(X = x) = 1$
- Probabilities are always greater than 0
  - $P(X = x) \geq 0$
- Joint distribution Y
  - $p(x, y) = P(X = x \text{ and } Y = y)$
- Conditional probability
  - $p(x \mid y) = \dfrac{p(x, y)}{p(y)}$
- Independence
  - $p(x, y) = p(x)p(y)$
  - $p(x \mid y) = p(x) = \dfrac{p(x, y)}{p(y)}$
- If X and Y are conditionally independent given Z = z
  - $p(x, y \mid z) = p(x \mid z)p(y \mid z)$
- Marginal Probability
  - $p(x) = \Sigma_y p(x \mid y)p(y)$

Why consider uncertainty?

Five major factors

1. Unpredictable environments
2. Sensors
3. Robot Motion
4. Models
5. Computation

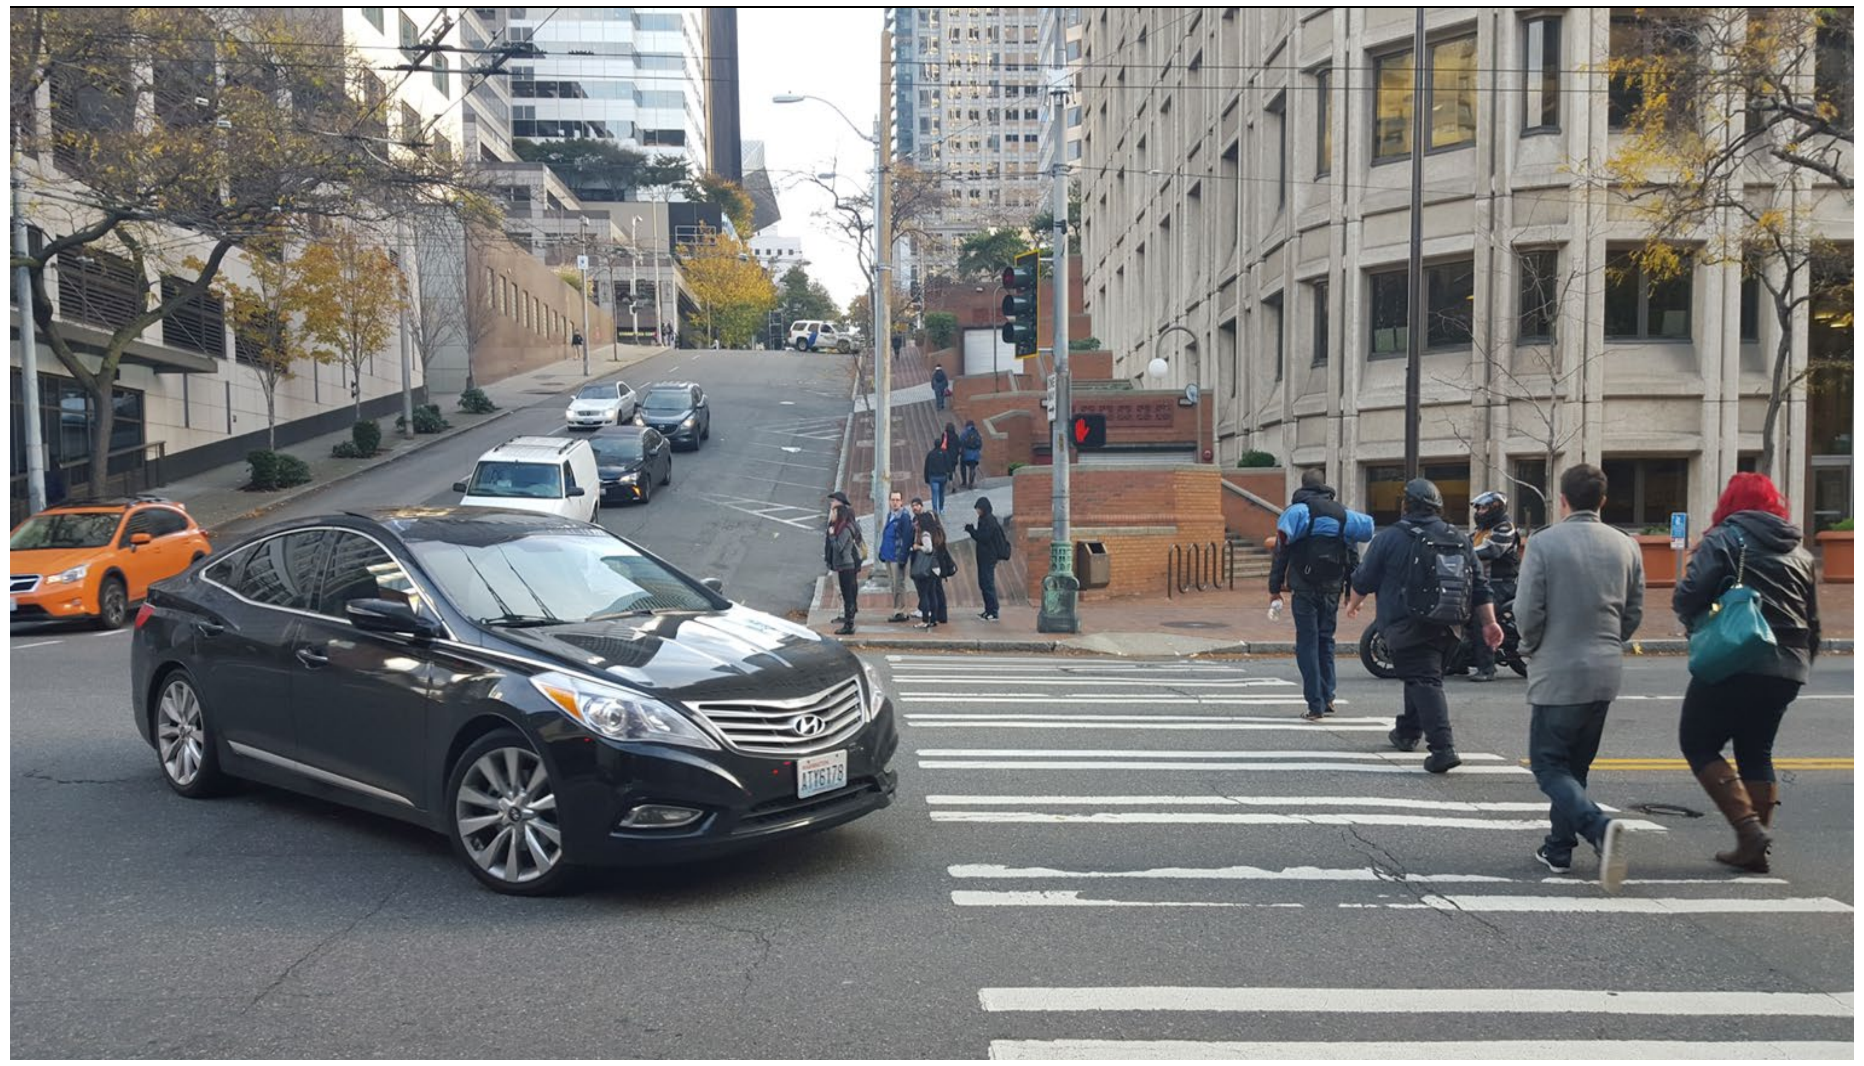

Exercise

The debate of 2015: what color is this dress?

- Is this dress black and blue or white and gold?
- Where does the uncertainty come from?
  - Blue and black under yellow-tinted lighting
  - White and gold under blue-tinted lighting

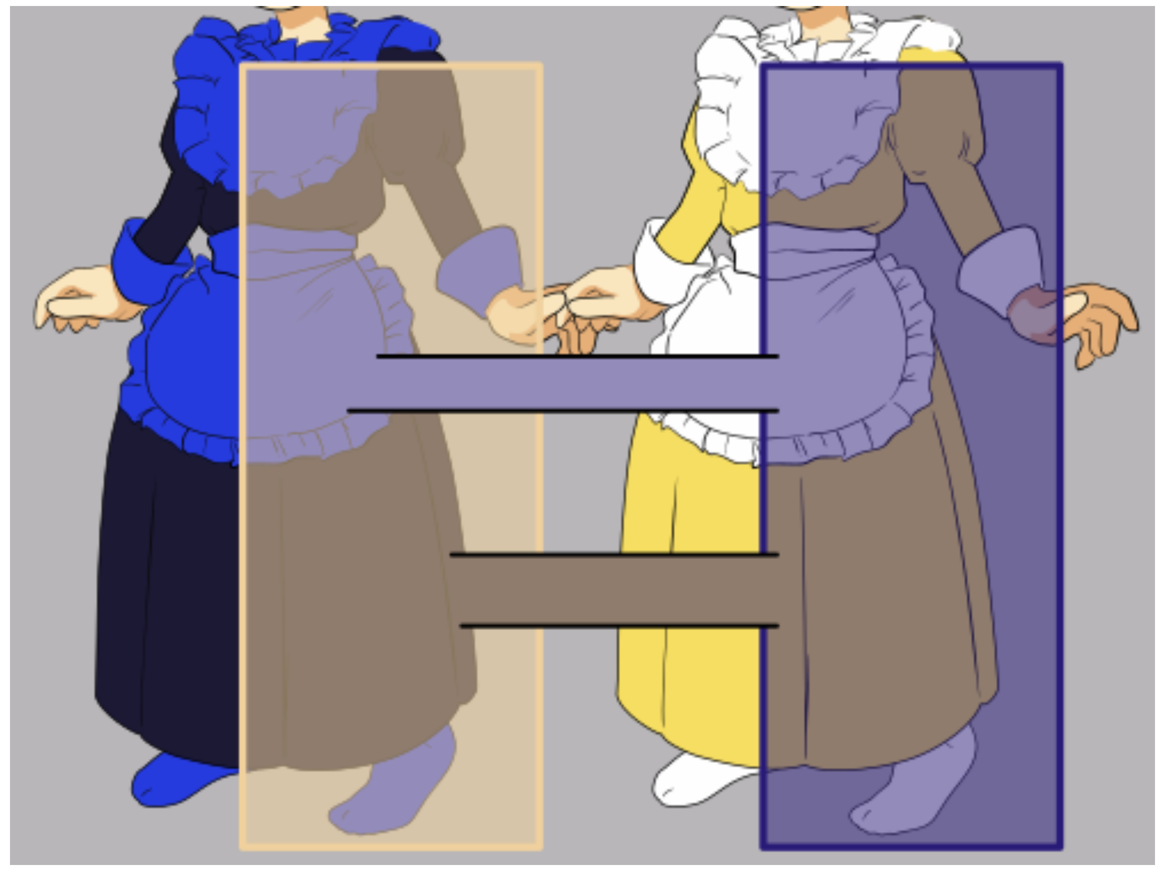

Robot-Environment Interaction

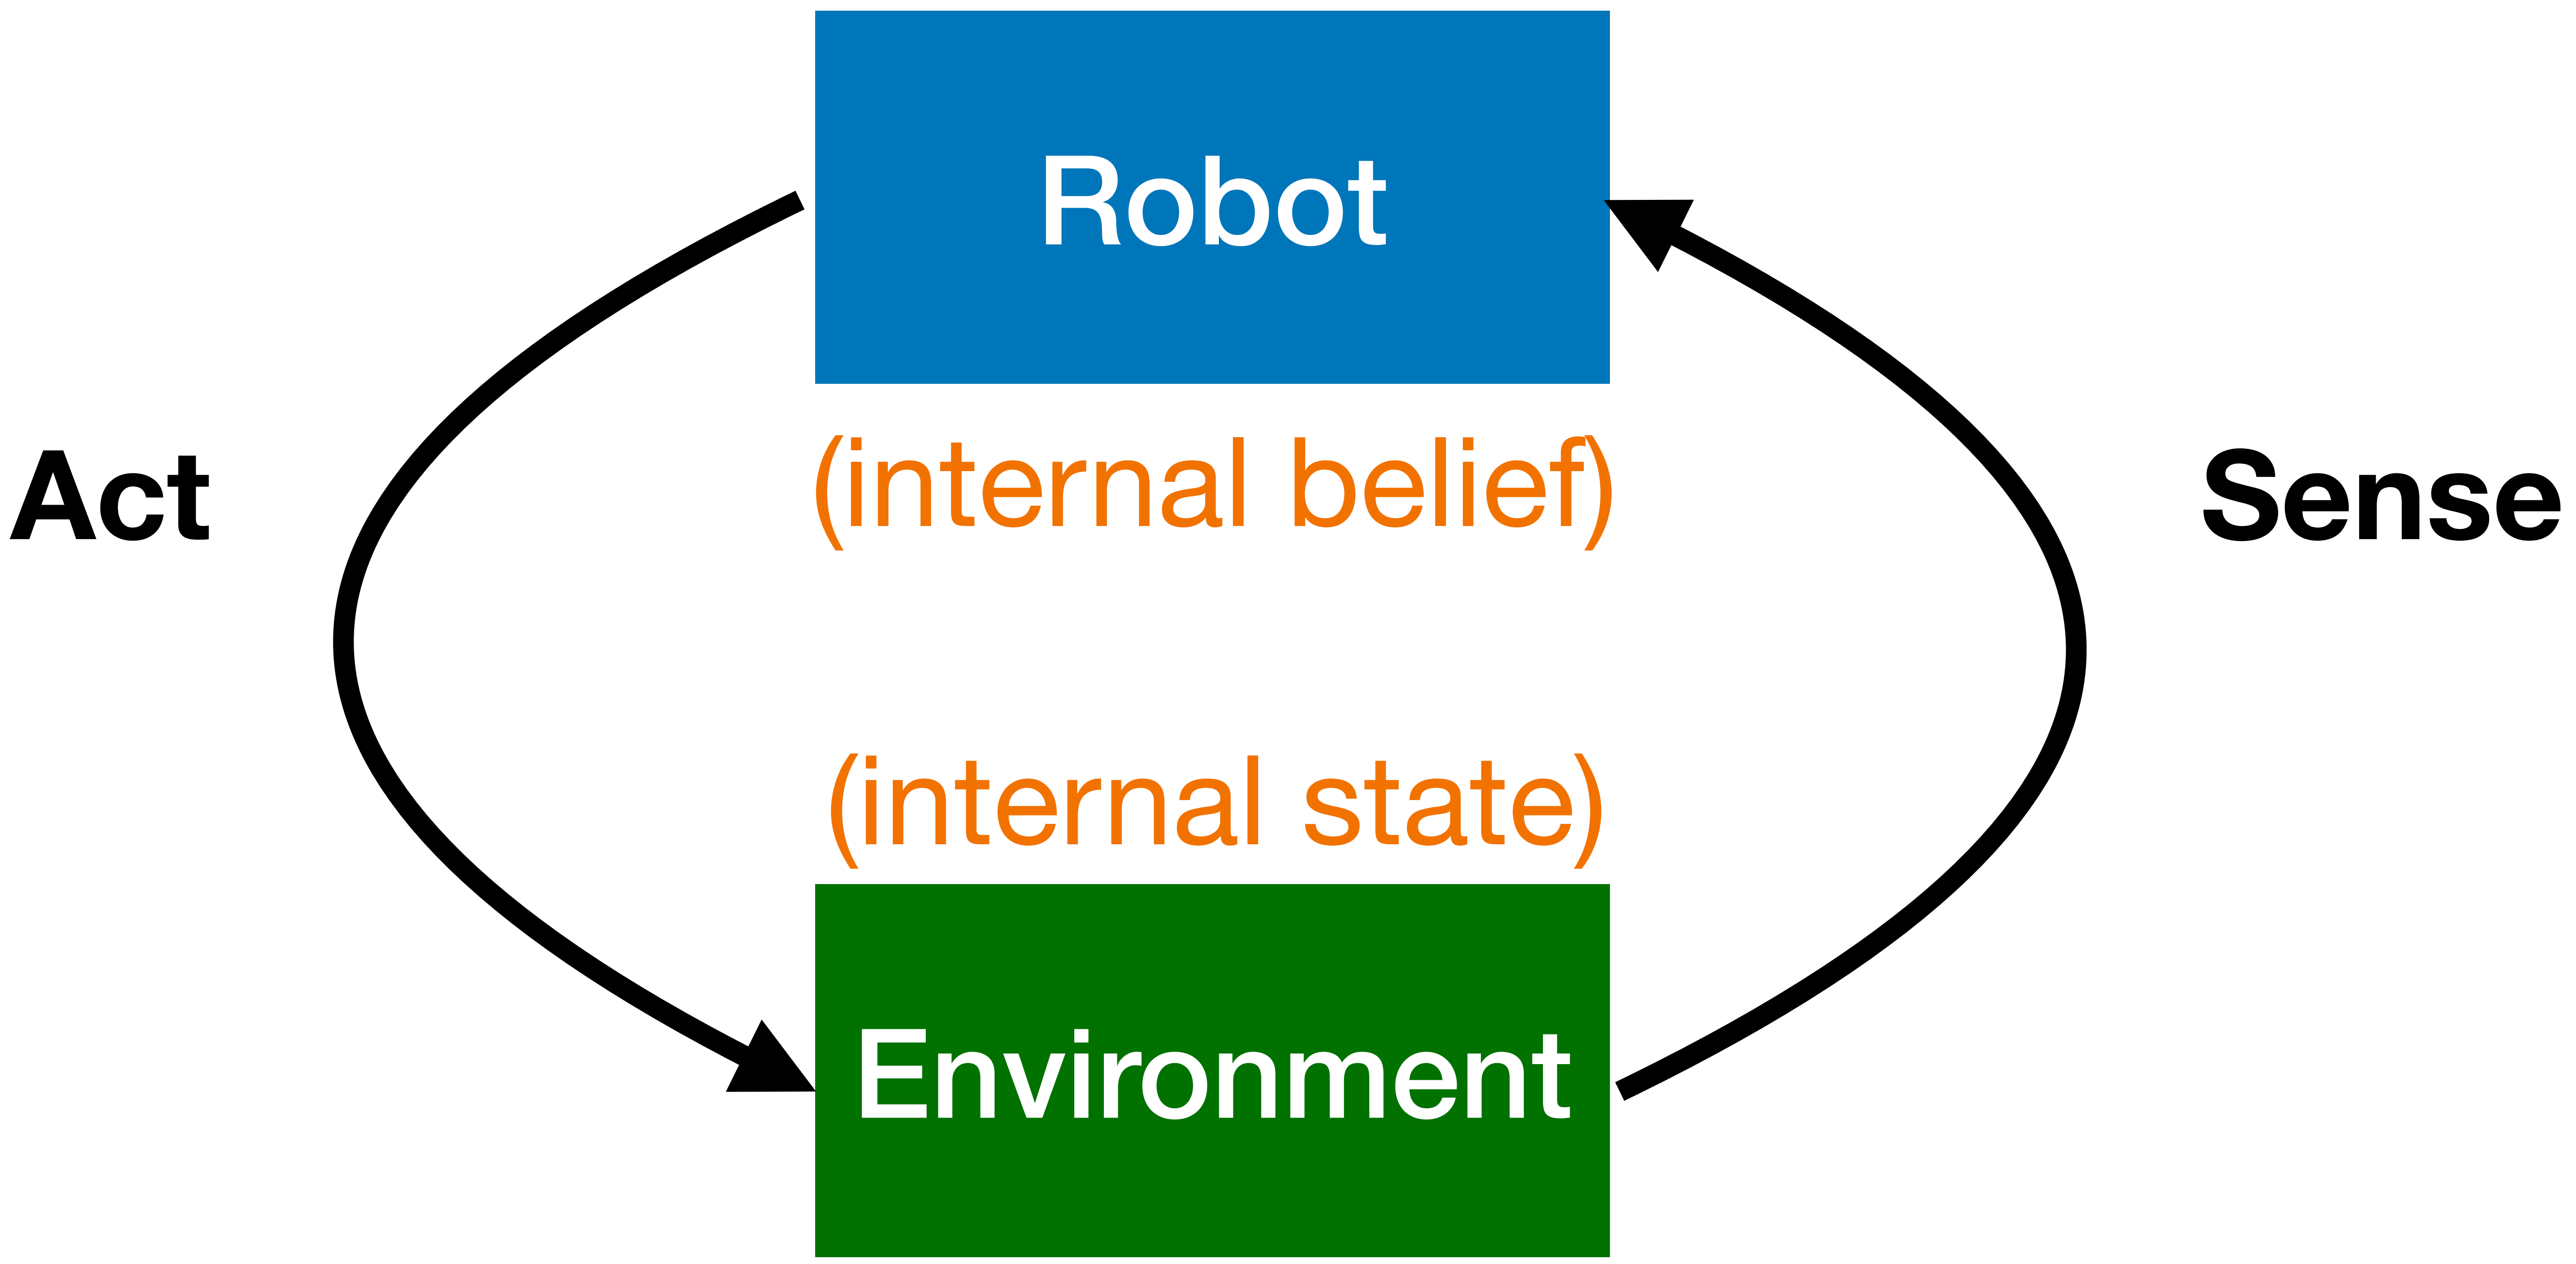

- Two fundamental types of interaction between a robot and its environment:
    - Sensor measurements/ observations
    - Control actions

Probabilistic Approach

"A robot that carries a notion of its own uncertainty and that acts accordingly is superior to one that does not" - *Probabilistic Robotics* by Thrun, Burgard, Fox

- Probabilistic approaches in contrast to traditional model-based motion planning techniques or reactive behavior-based motion:
  - Tend to be more robust to sensor and model limitations
  - Weaker requirements on the accuracy of the robot's models

Is Robotics Going Statistics?
The Field of Probabilistic Robotics

Sebastian Thrun
School of Computer Science
Carnegie Mellon University
http://www.cs.cmu.edu/∼thrun

draft, please do not circulate

Abstract

In the 1970s, most research in robotics presupposed the availability of exact models, of robots and their environments. Little emphasis was placed on sensing and the intrinsic limitations of modeling complex physical phenomena. This changed in the mid-1980s, when the paradigm shifted towards reactive techniques. Reactive controllers rely on capable sensors to generate robot control. Rejections of models were typical for researchers in this field. Since the mid-1990s, a new approach has begun to emerge: probabilistic robotics. This approach relies on statistical techniques to seamlessly integrate imperfect models and imperfect sensing. The present article describes the basics of probabilistic robotics and highlights some of its recent successes.

Probabilistic Approach

✓ Explicitly represent uncertainty using probabilities

✓ Accommodate inaccurate models

✓ Accommodate imperfect sensors

✓ Robust in real-world applications

✓ Best known approach to many hard robotics problems

- Computationally demanding
- Approximations
- False Assumptions

Is Robotics Going Statistics?
The Field of Probabilistic Robotics

Sebastian Thrun
School of Computer Science
Carnegie Mellon University
http://www.cs.cmu.edu/~thrun

draft, please do not circulate

Abstract

In the 1970s, most research in robotics presupposed the availability of exact models, of robots and their environments. Little emphasis was placed on sensing and the intrinsic limitations of modeling complex physical phenomena. This changed in the mid-1980s, when the paradigm shifted towards reactive techniques. Reactive controllers rely on capable sensors to generate robot control. Rejections of models were typical for researchers in this field. Since the mid-1990s, a new approach has begun to emerge: probabilistic robotics. This approach relies on statistical techniques to seamlessly integrate imperfect models and imperfect sensing. The present article describes the basics of probabilistic robotics and highlights some of its recent successes.

Bayesian Inference

- Inference = educated guessing
- Bayesian inference = guessing in the style of Bayes
- Example from Ed discussion:
  - "My robot stopped moving, the hardware is broken, send me new parts"
  - What is the probability that the robot is broken given that it stopped moving?

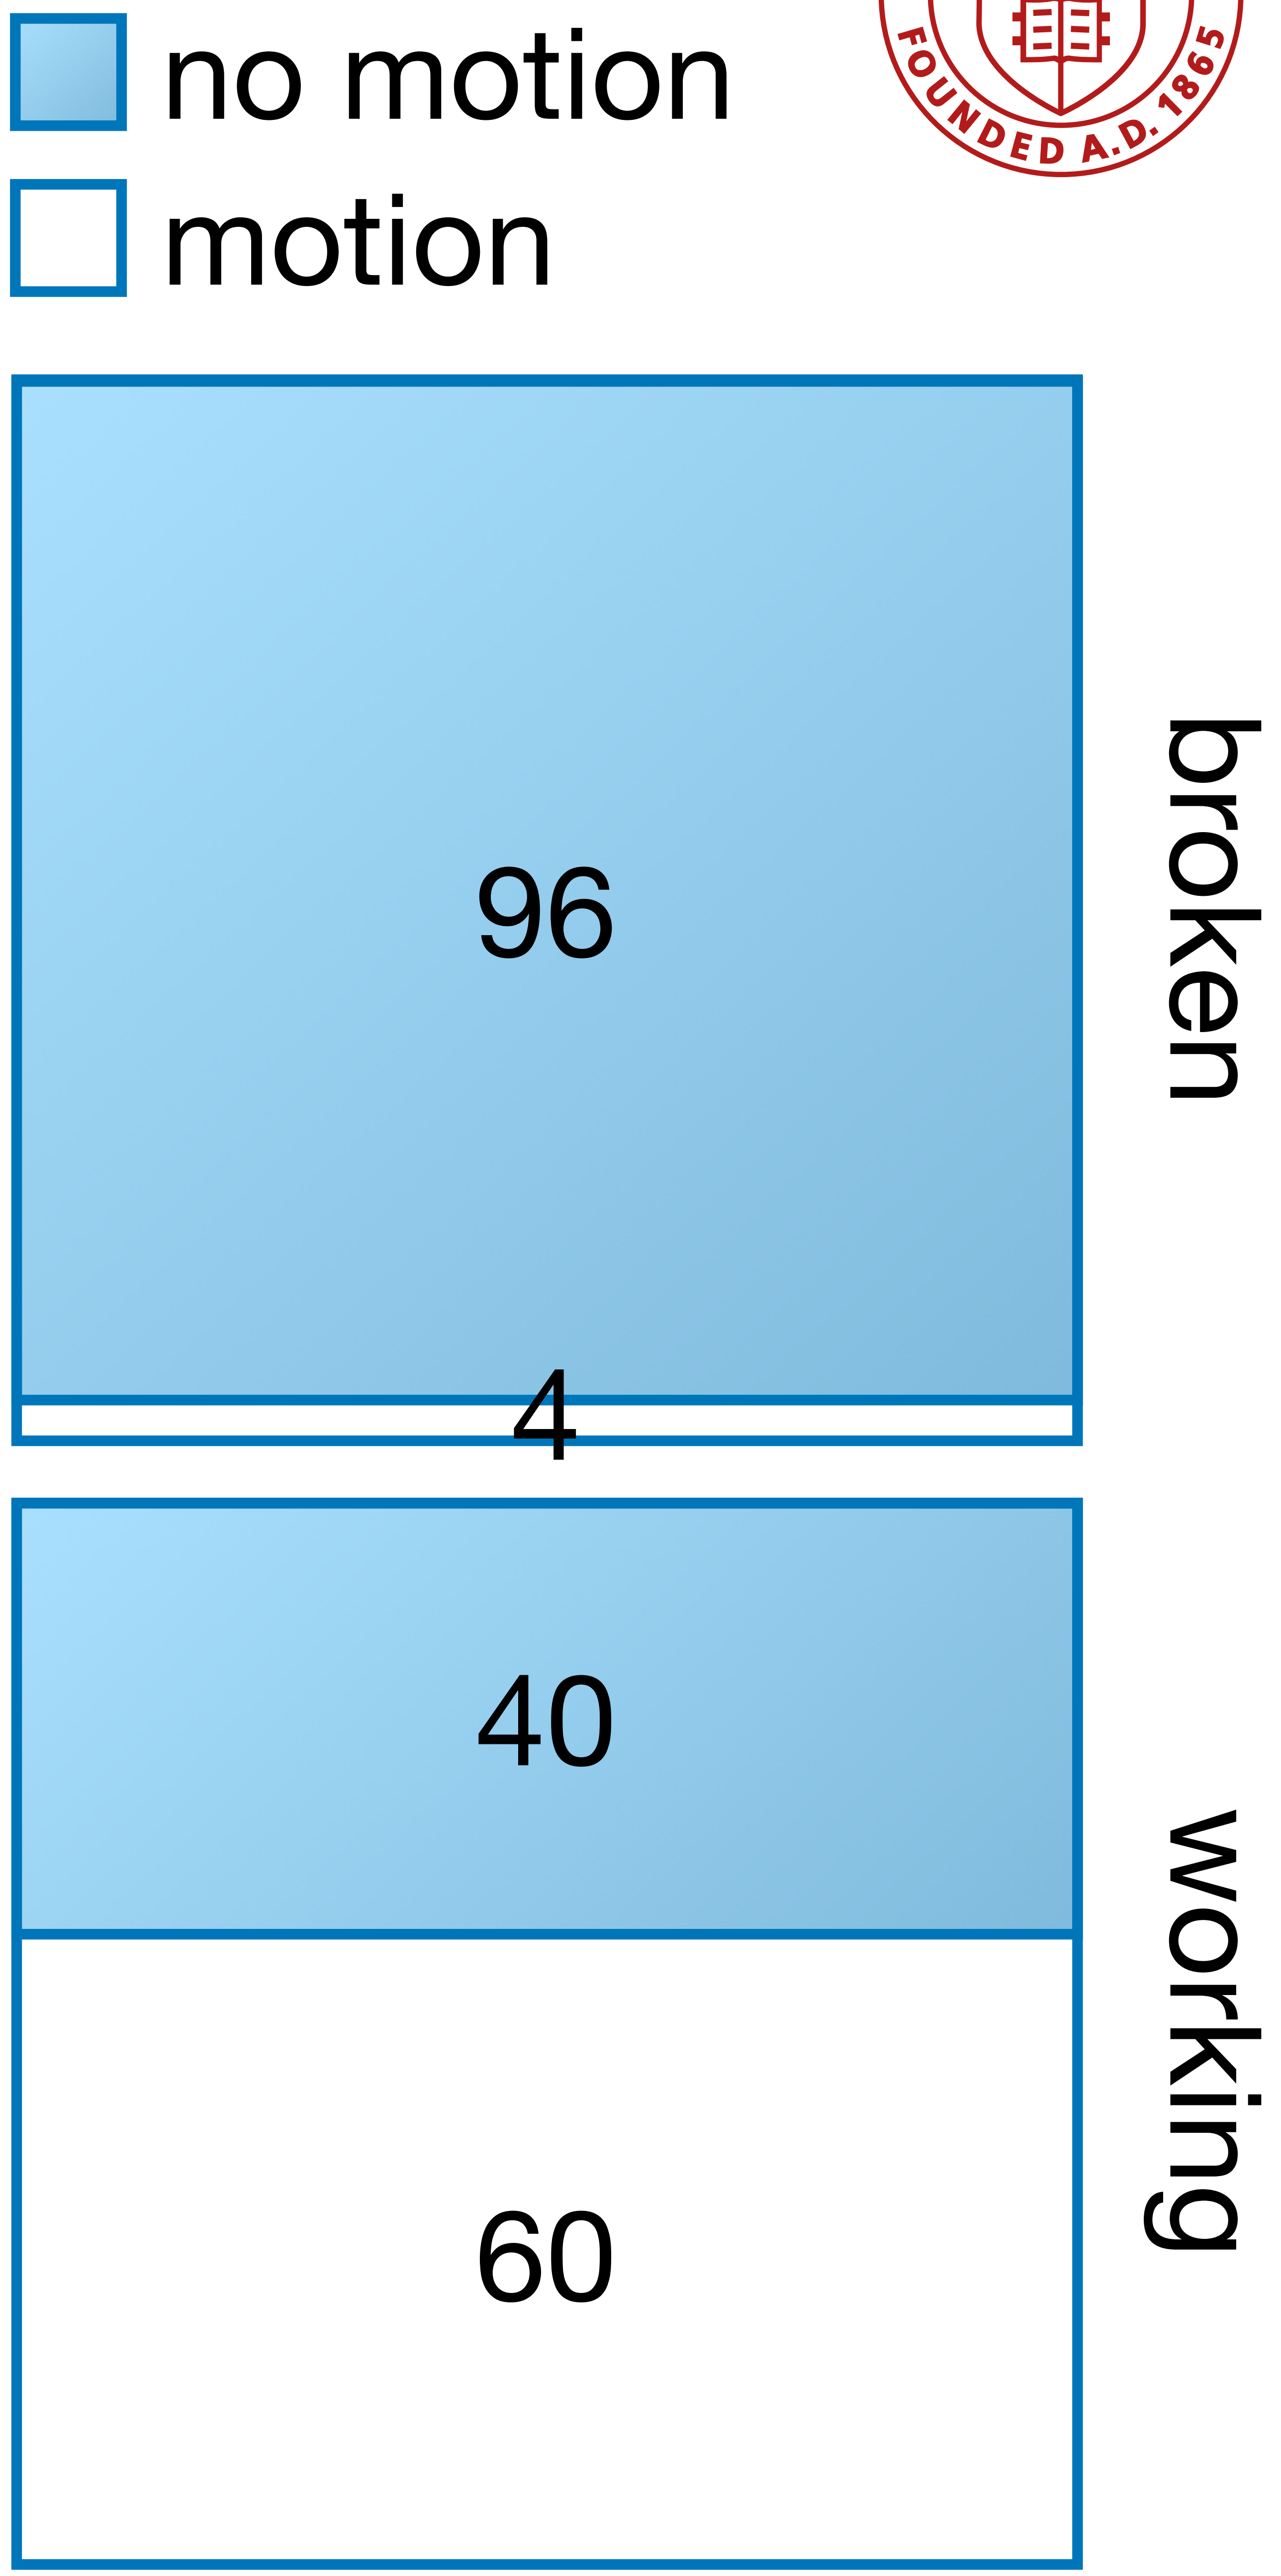

Bayesian Inference

- Inference = educated guessing
- Bayesian inference = guessing in the style of Bayes
- Example
  - “My robot stopped moving, the hardware is broken, send me new parts”
  - What is the probability that the robot is broken given that it stopped moving?

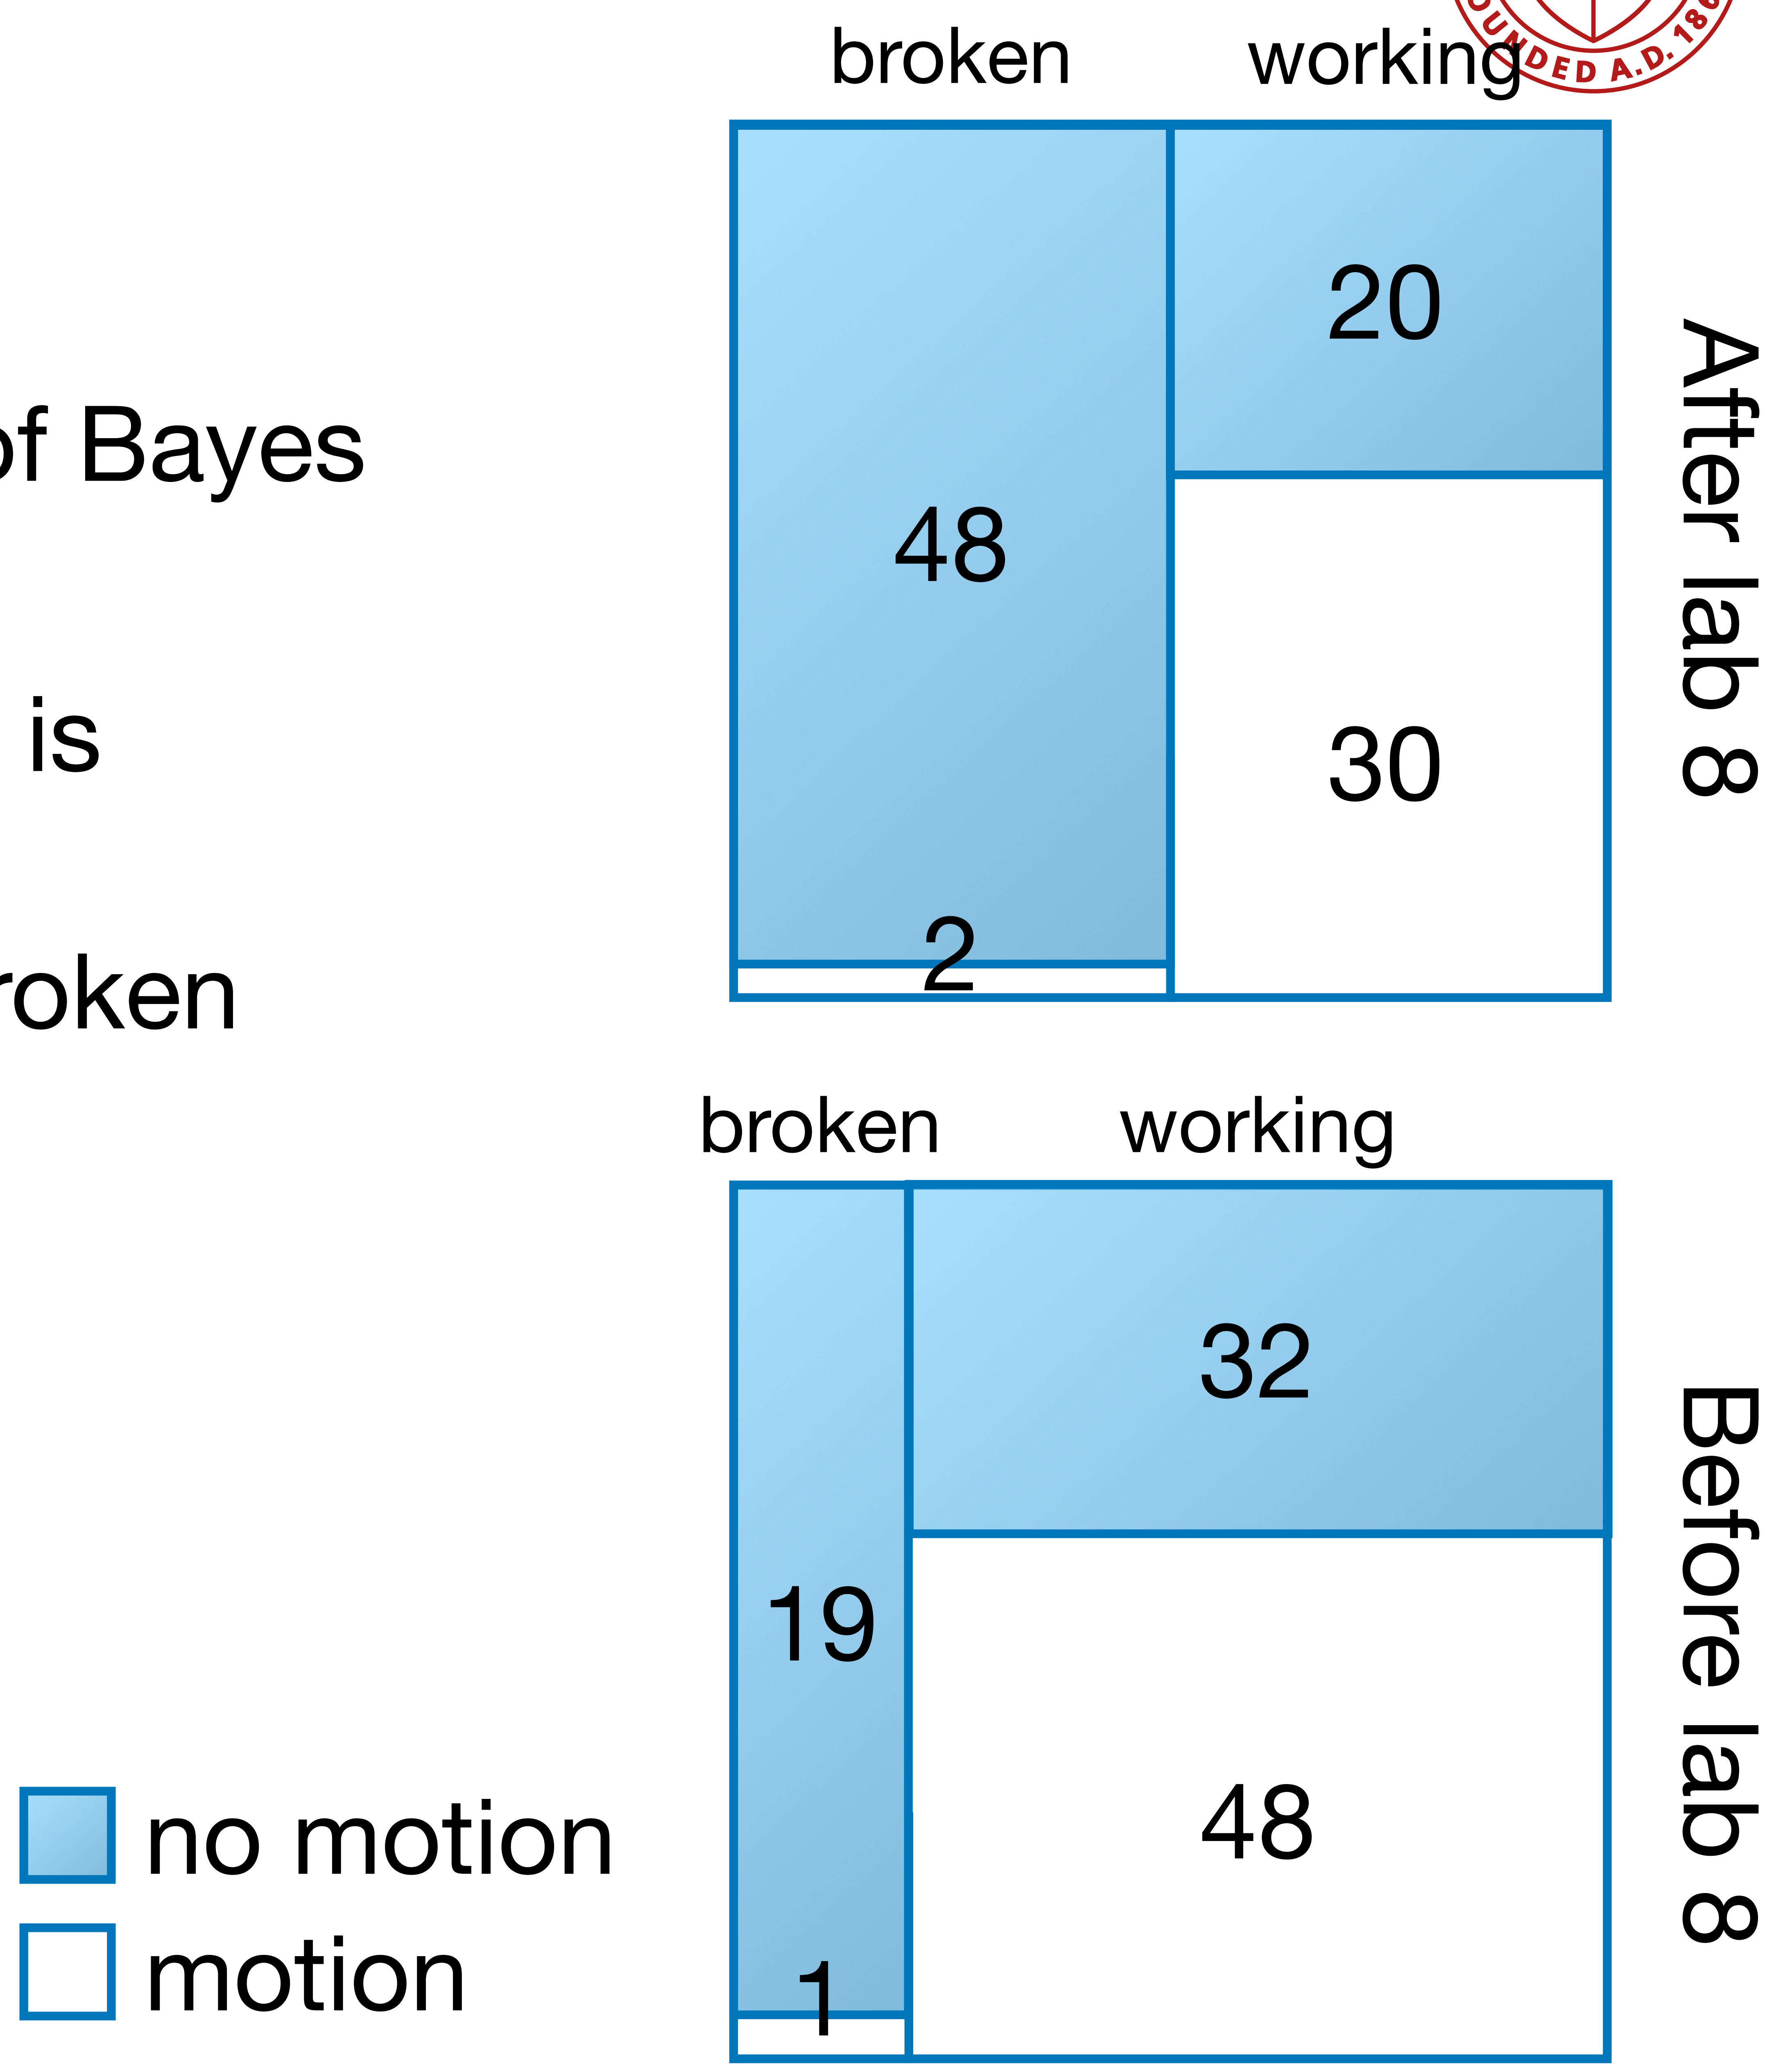

Bayesian Inference

Translate to probability

- P(something) = #something/ #everything
- Before lab 8:
  - P(broken) =
  - P(working) =
- After lab 8:
  - P(broken) =
  - P(working) =

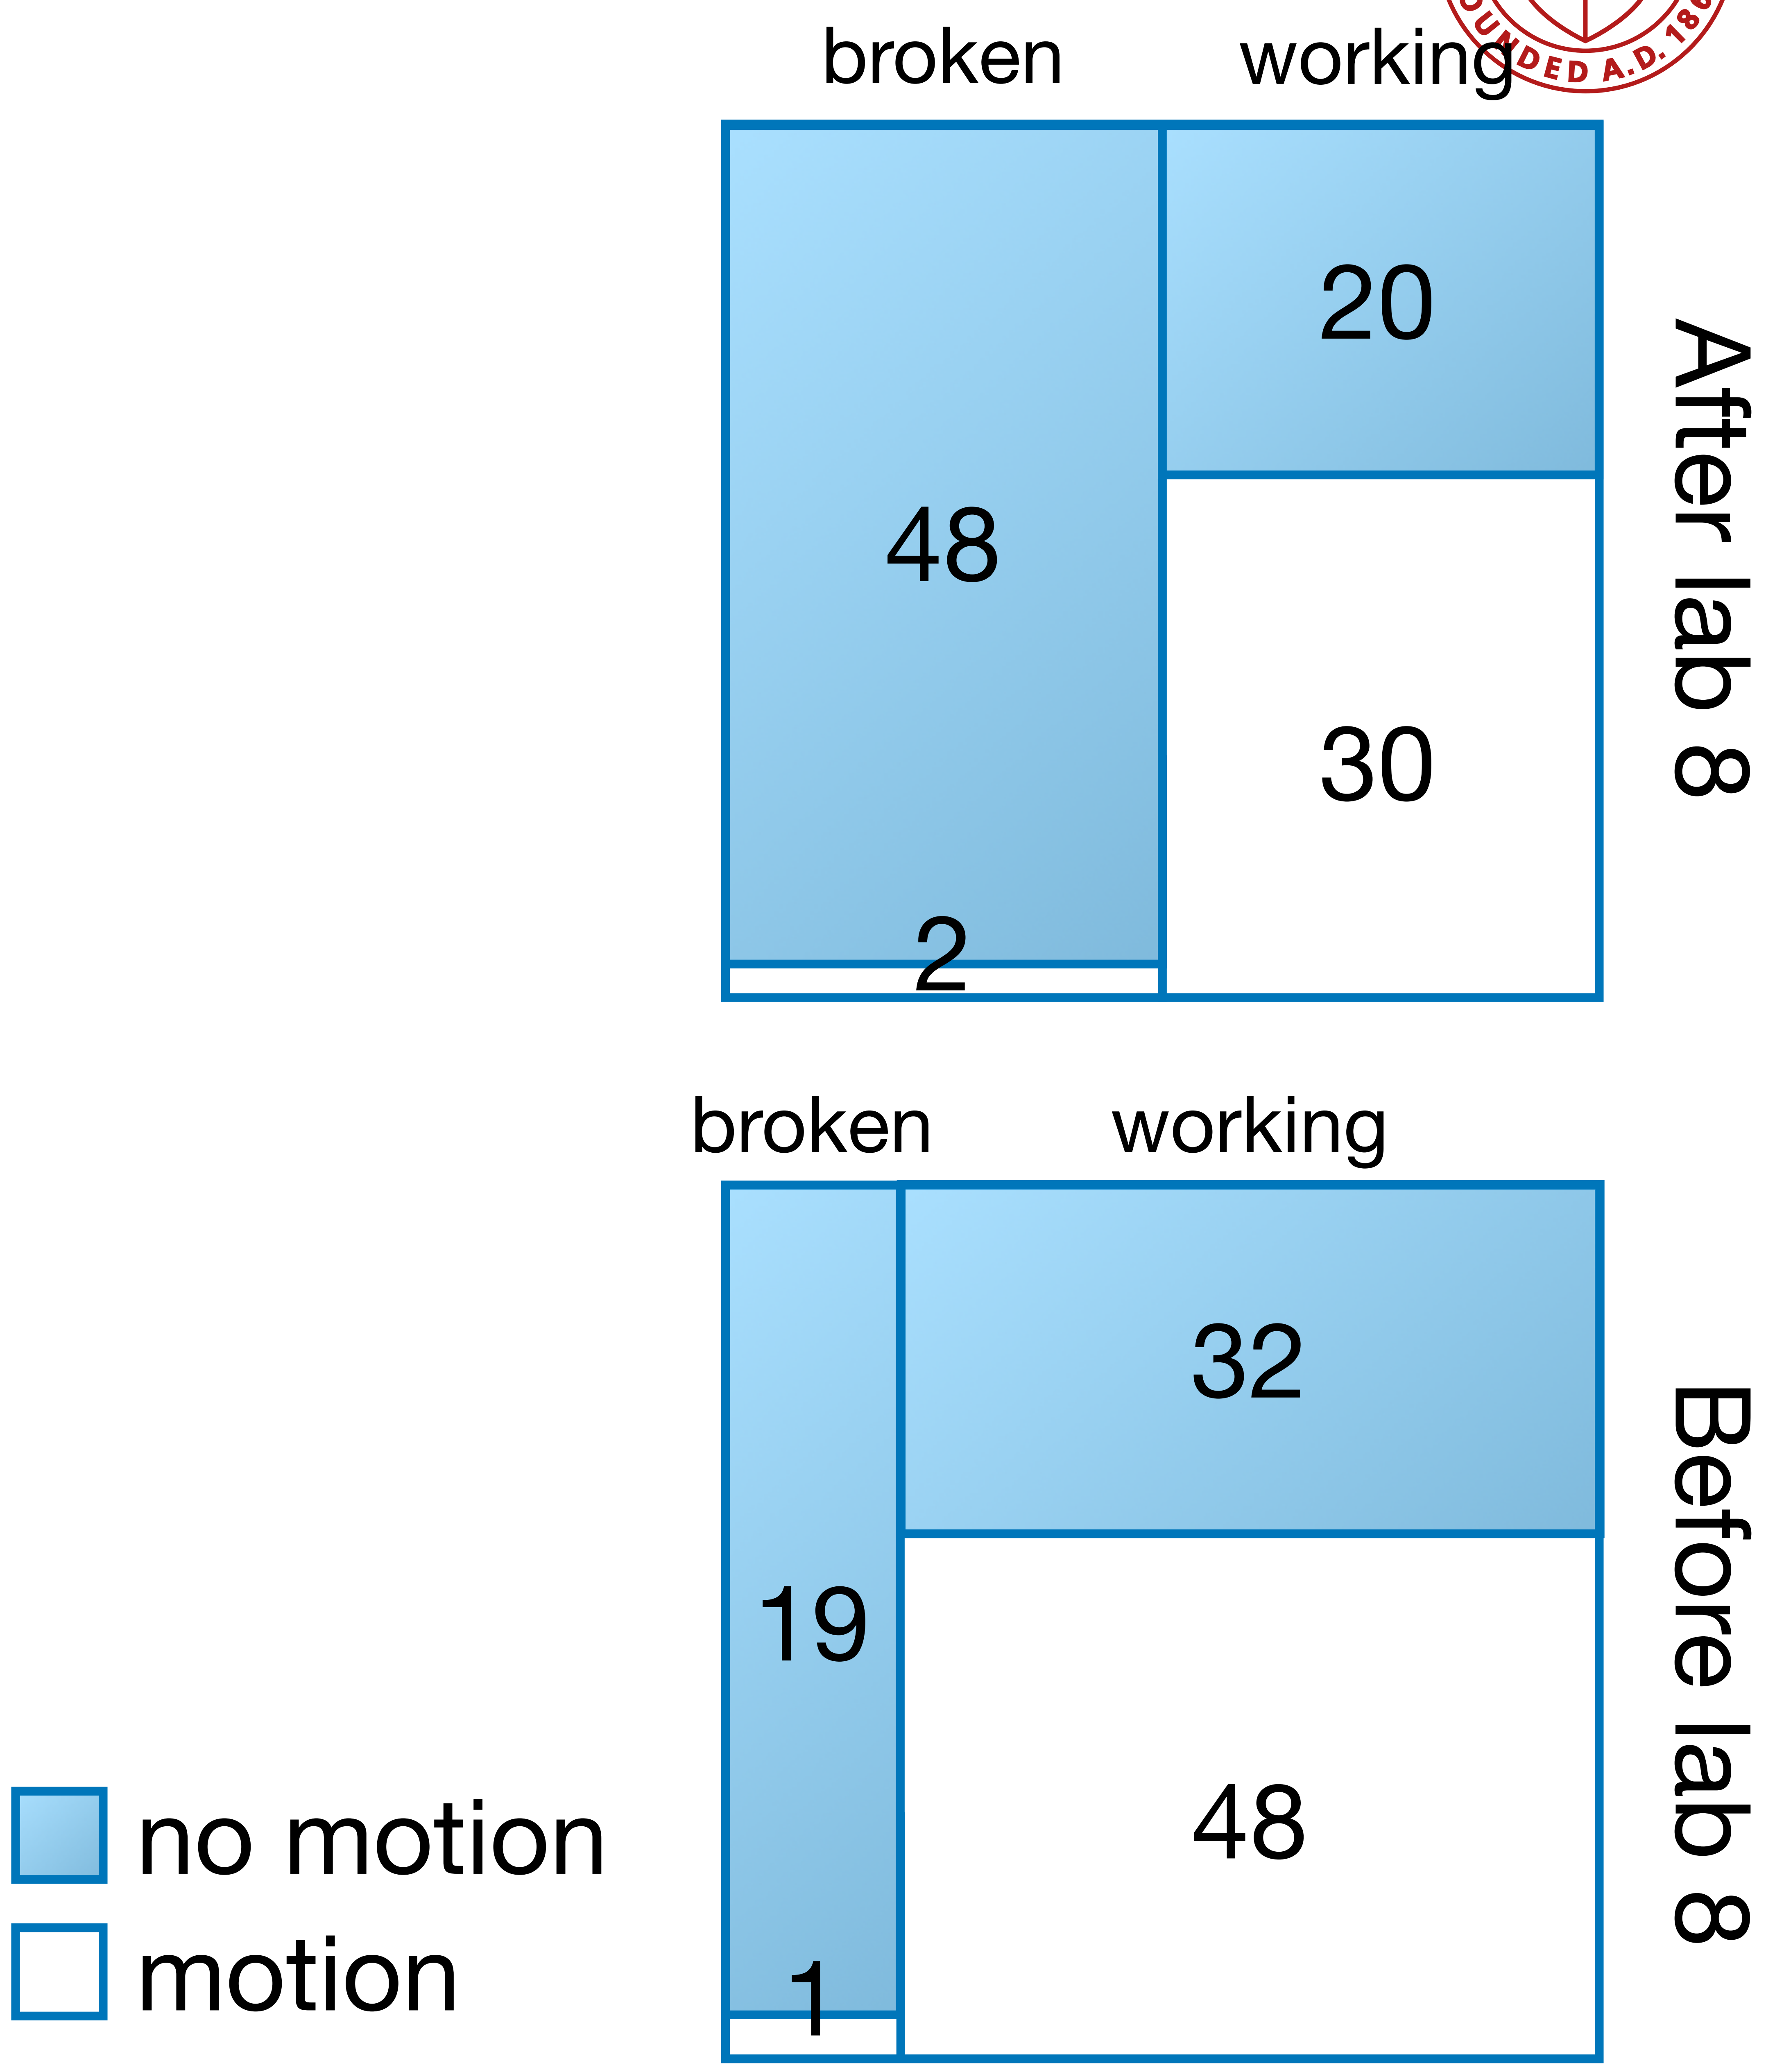

Bayesian Inference

Conditional Probability

- If you know that the robot is broken, what is the probability that it stopped moving?
- P(no motion | broken) =

- P(no motion | working) =

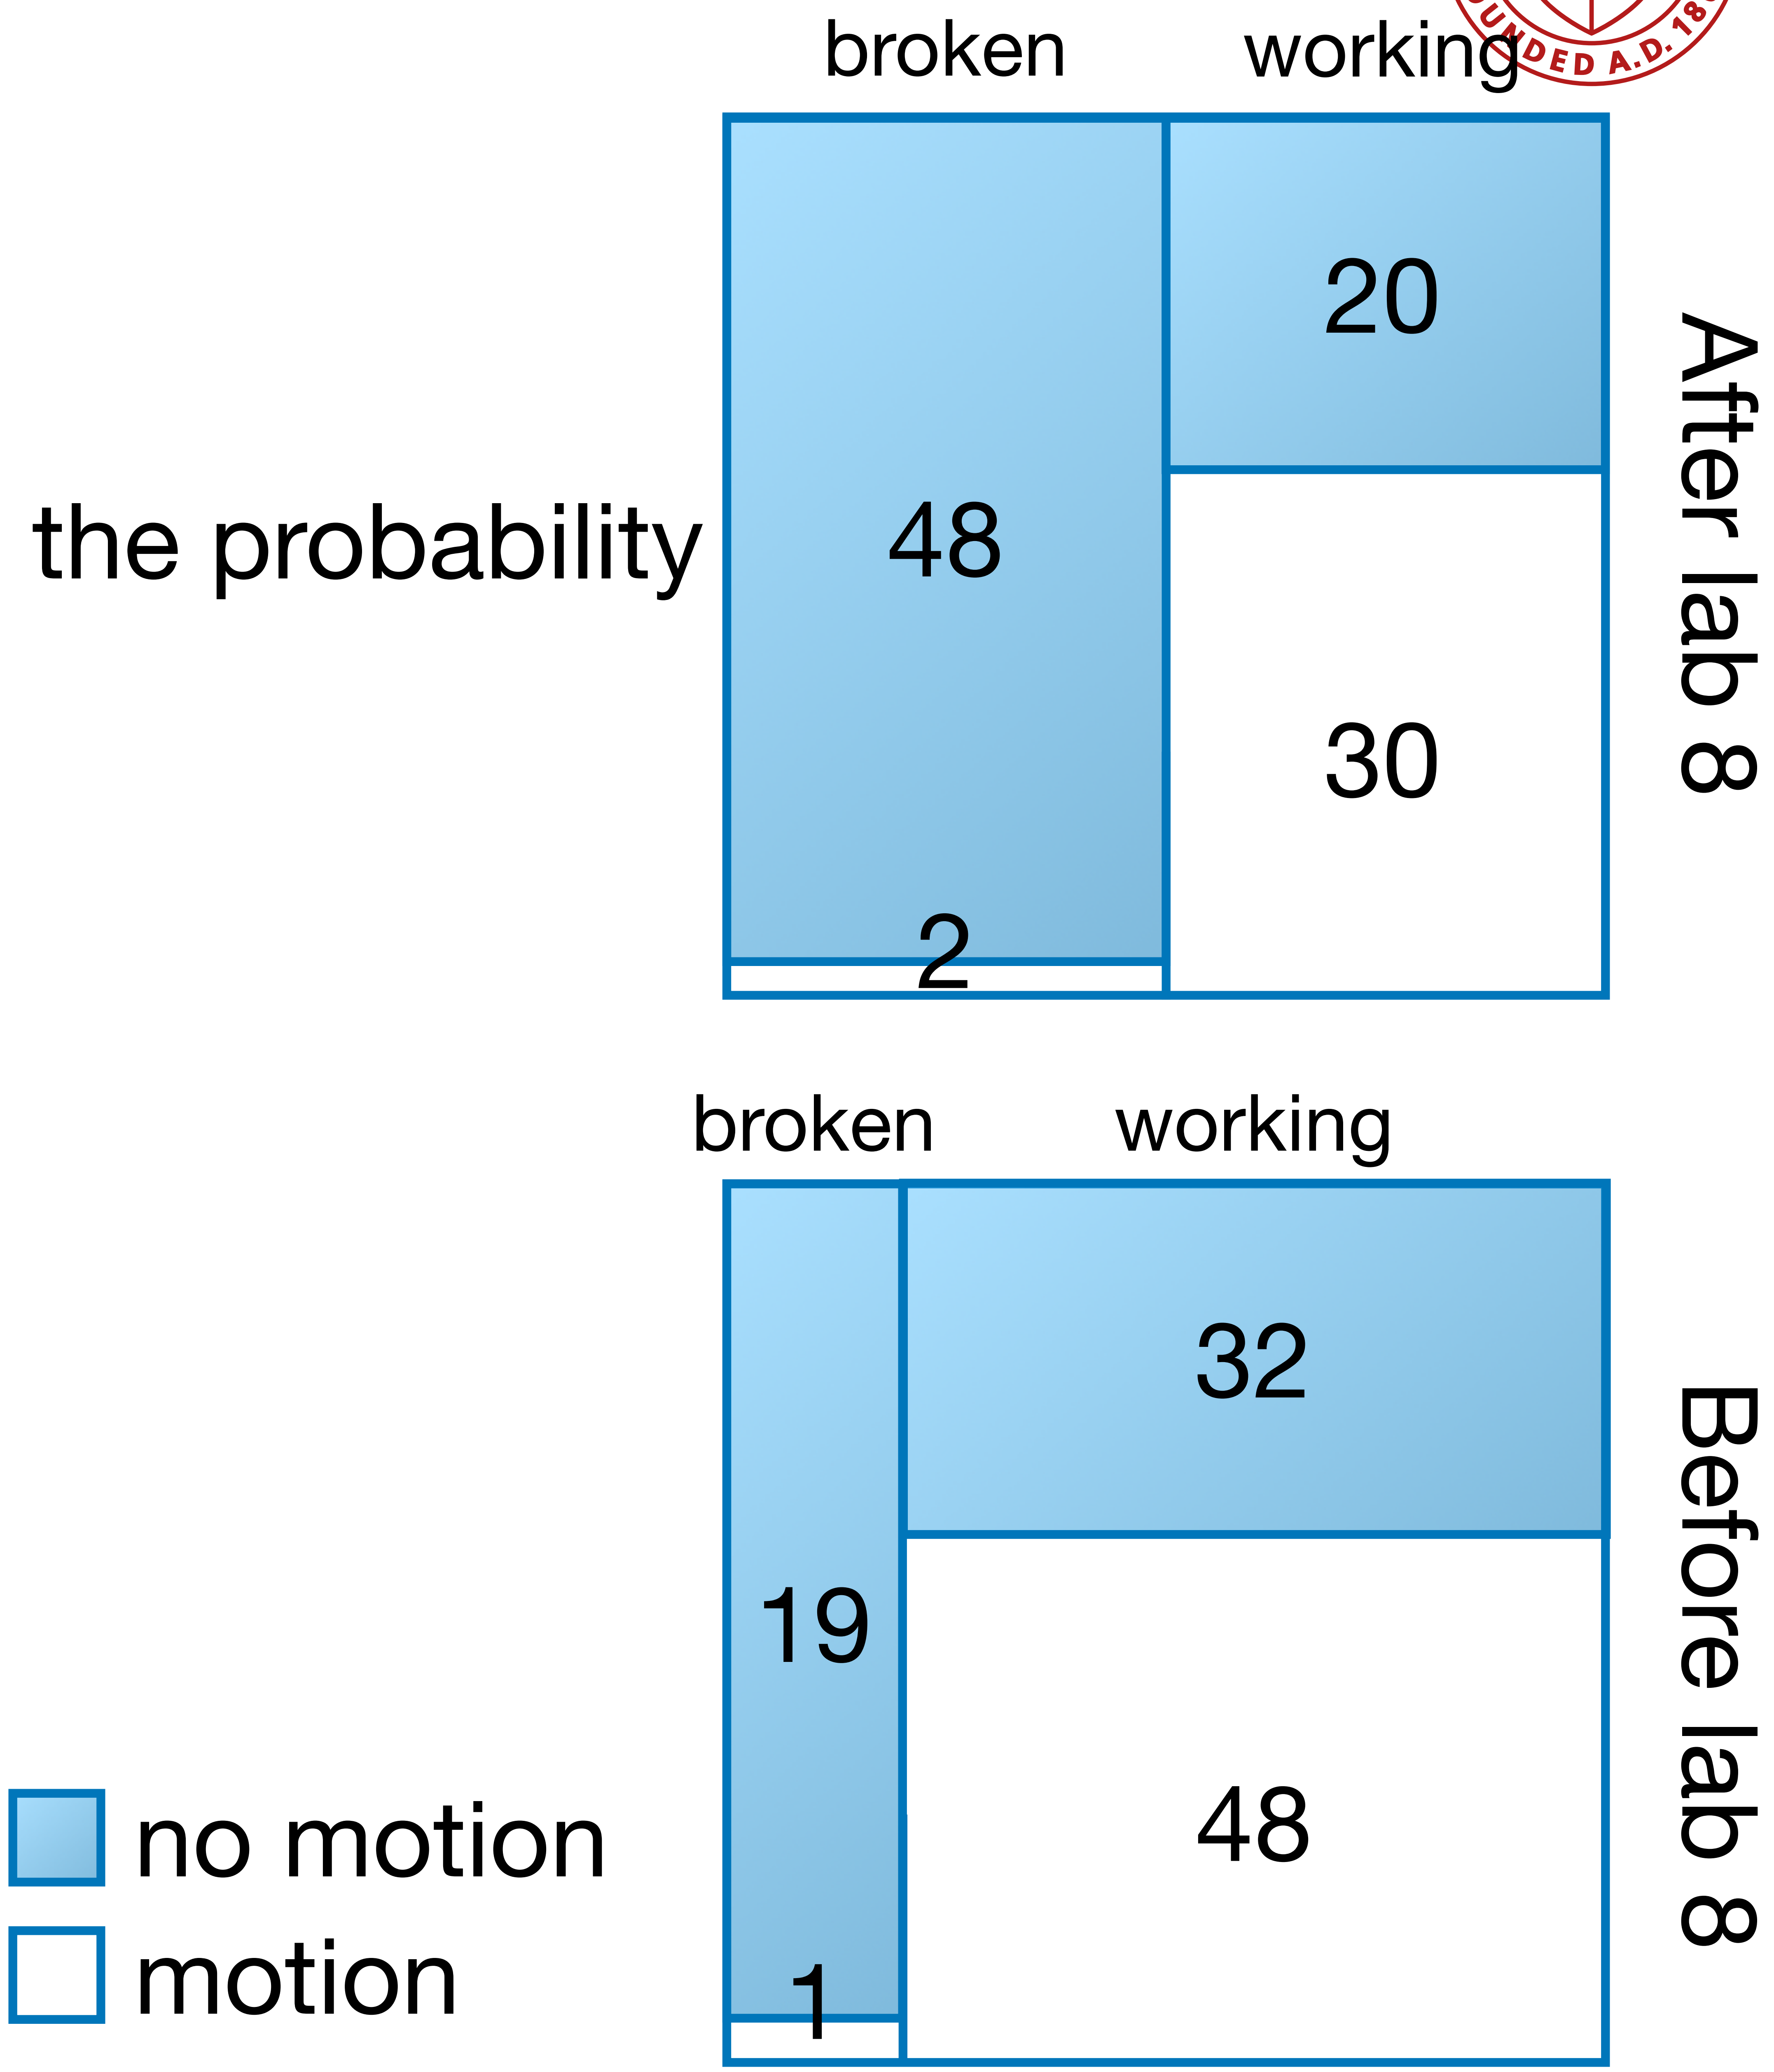

Bayesian Inference

Conditional Probability

- If you know that the robot is broken, what is the probability that it stopped moving?
- P(no motion | broken) =
  #(broken and no motion)/#broken
  - Before lab 8 = 19/20 = 0.95
- P(no motion | working) =
  #(working and no motion)/#working
  - Before lab 8 = 32/80 = 0.40

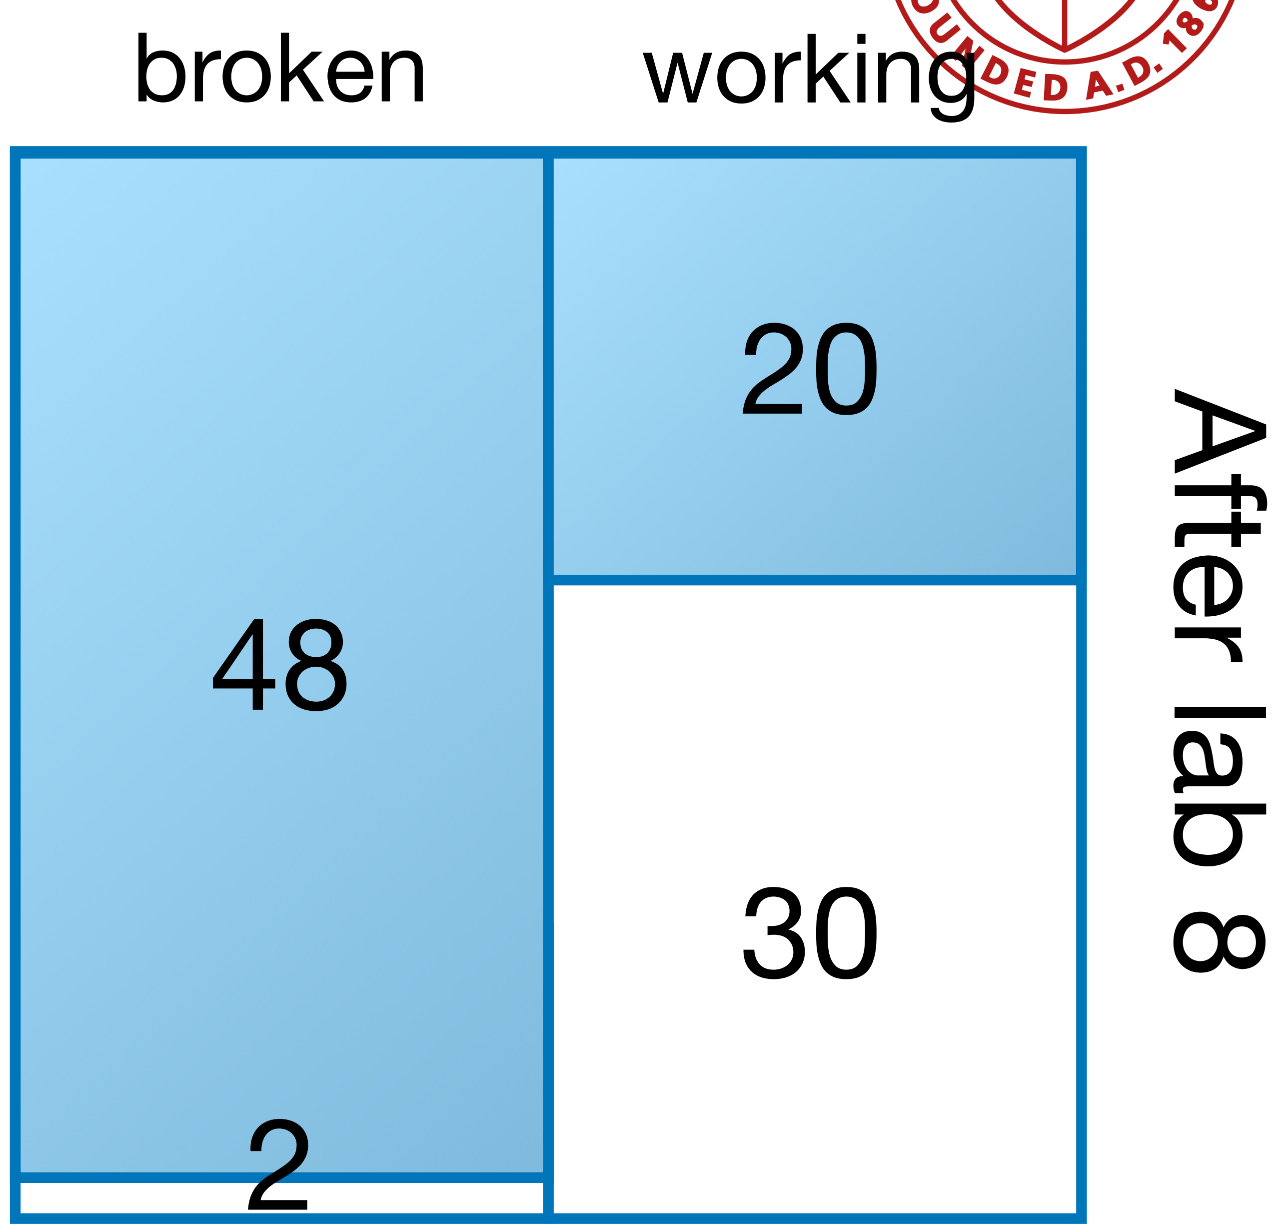

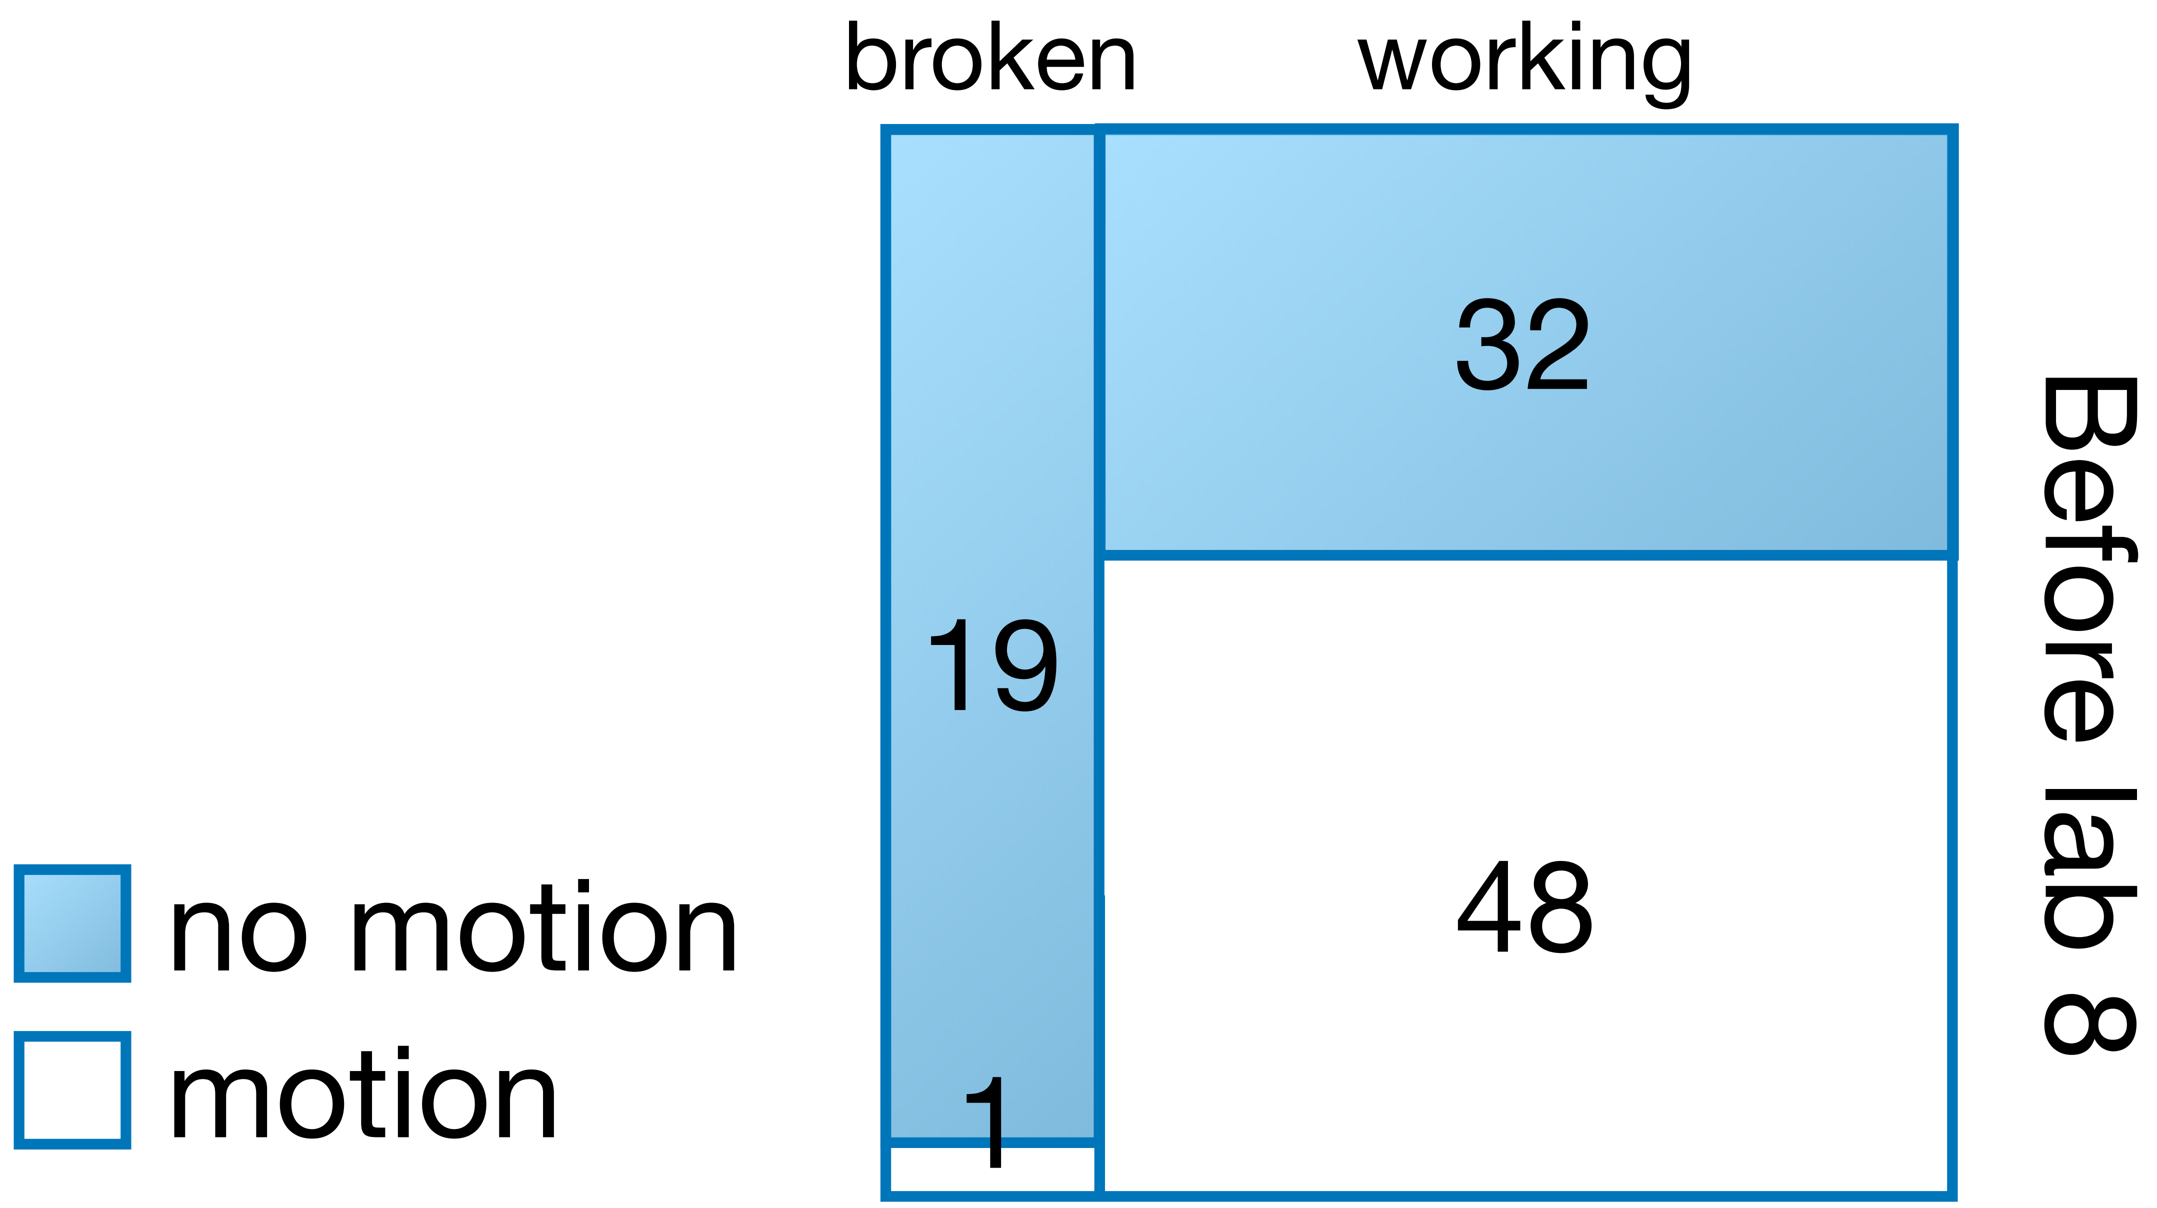

Bayesian Inference

Conditional Probability

- If you know that the robot is broken, what is the probability that it stopped moving?
- P(A|B) is the probability of A, given B
- Note: P(A|B) is not equal to P(B|A)
  - P(cute | puppy) ≠ P(puppy | cute)

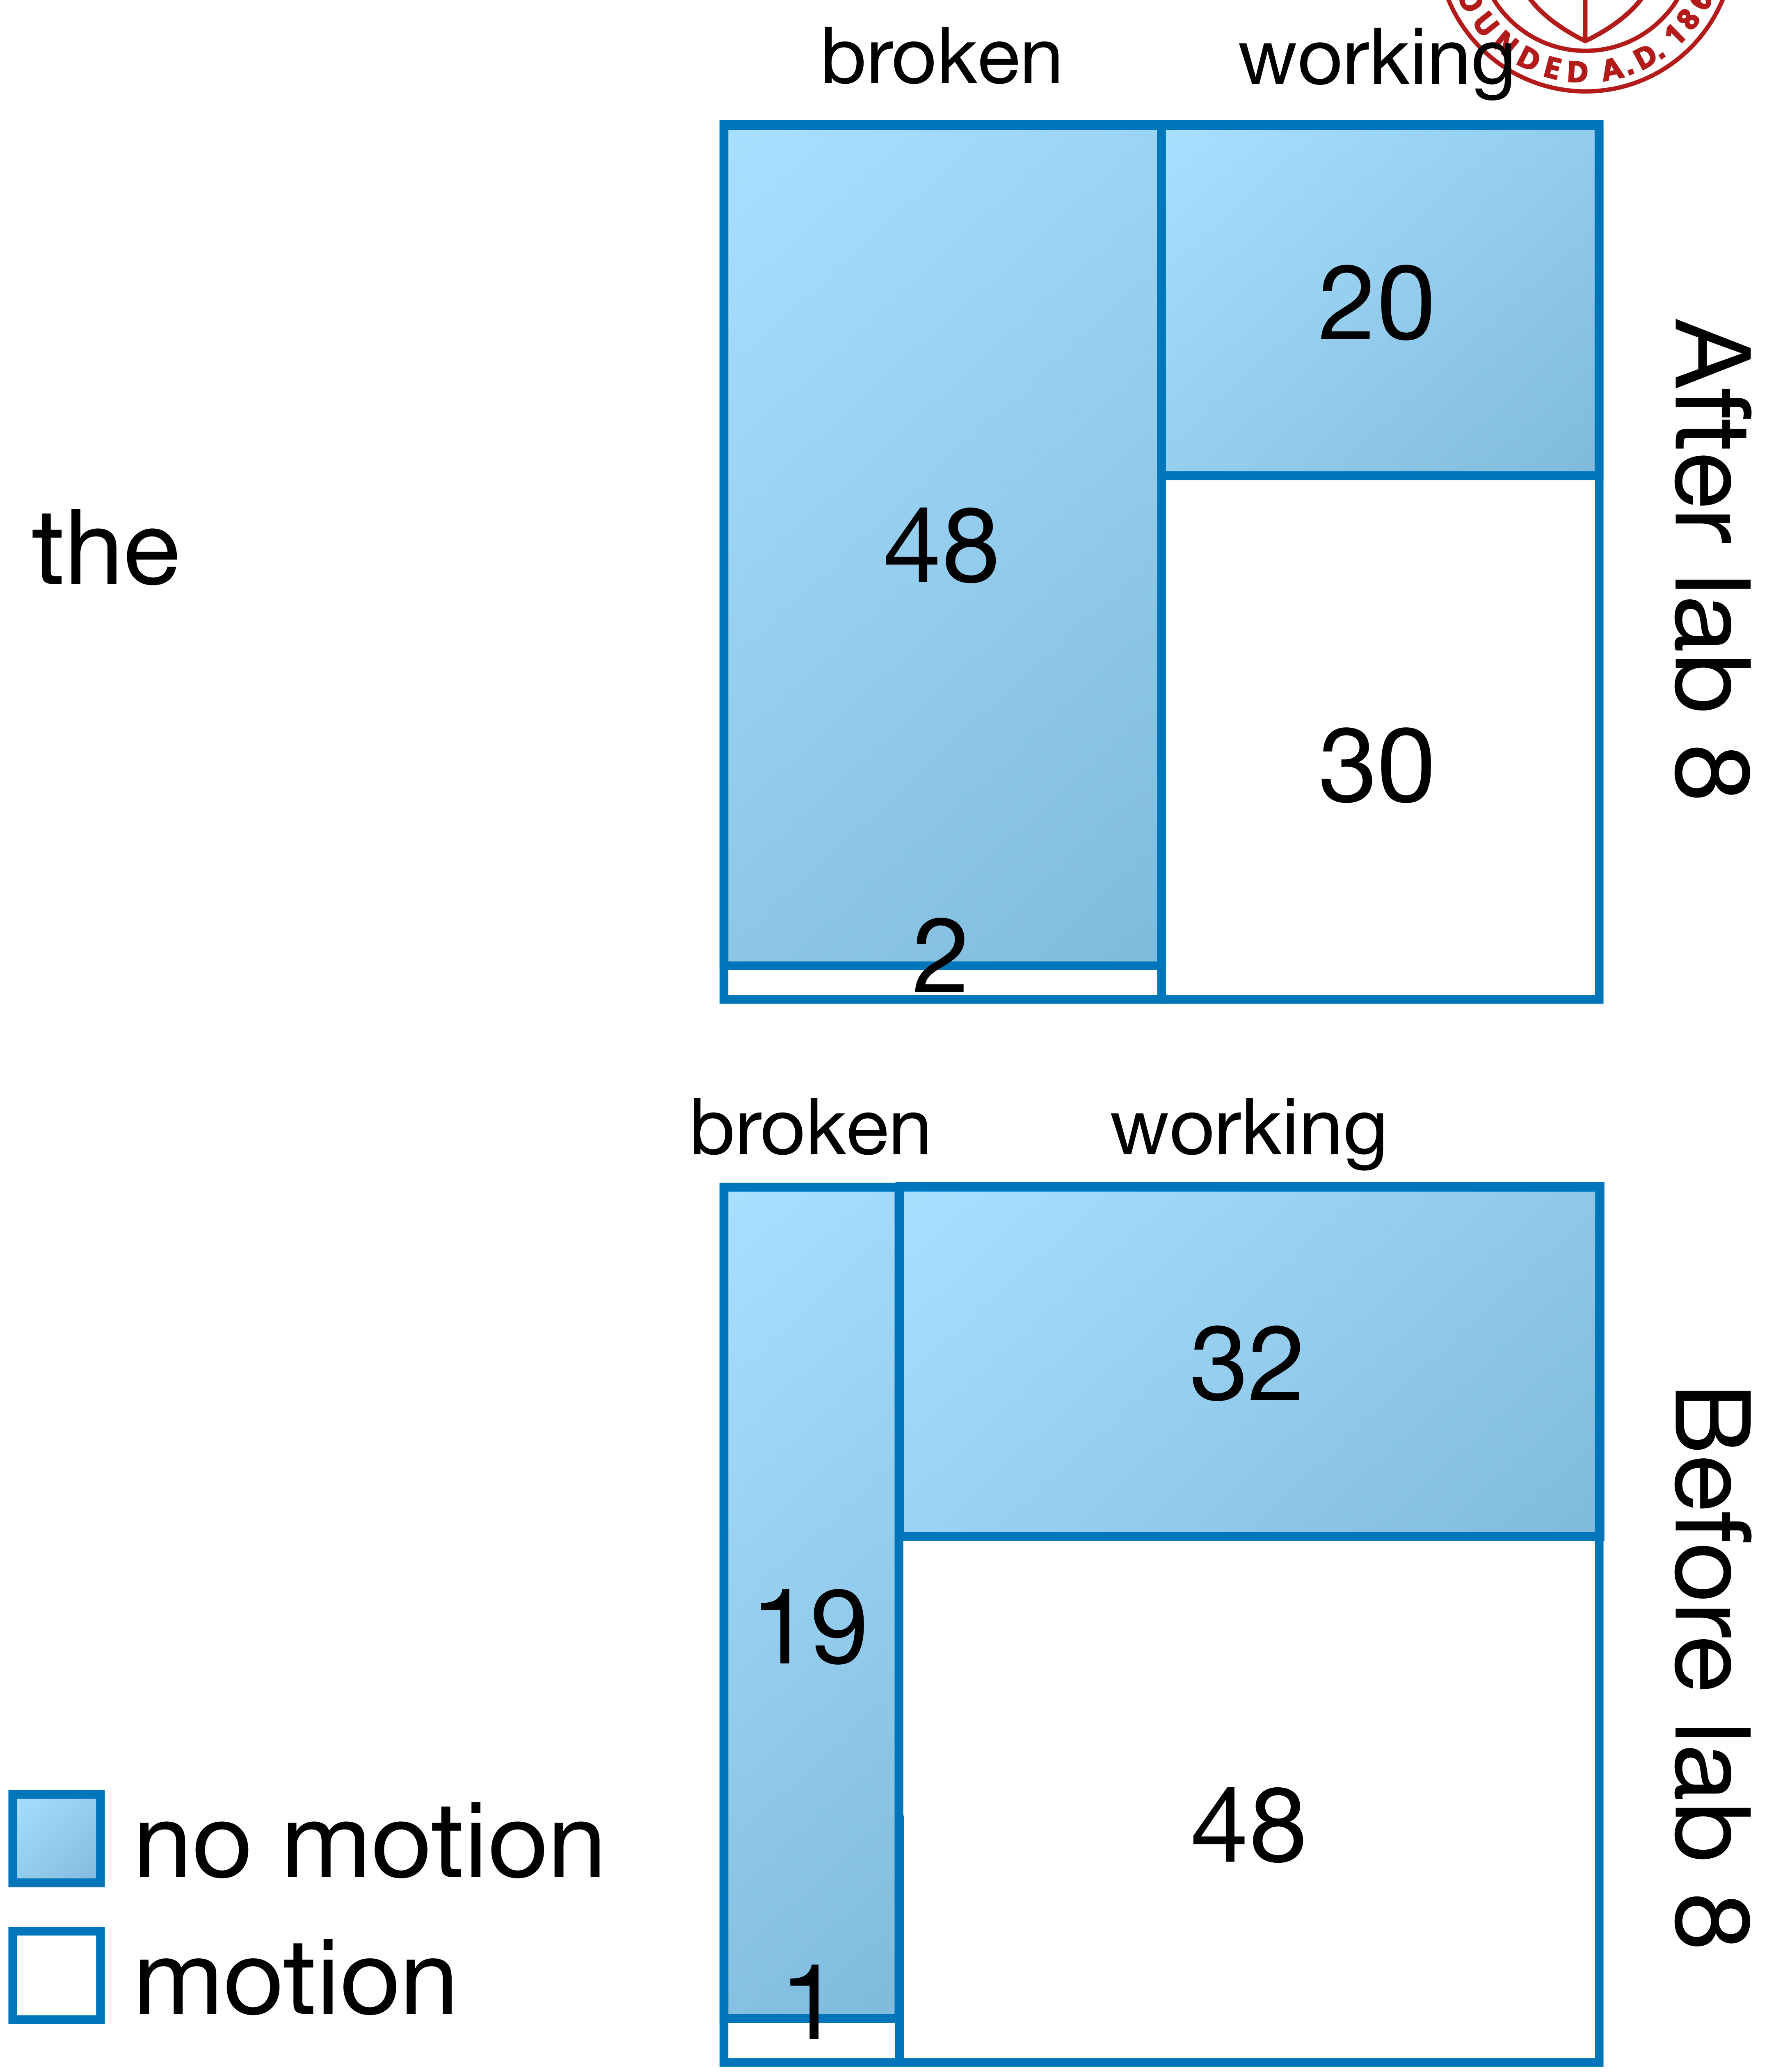

Bayesian Inference

Joint Probability

- What is the probability that the robot is both broken and not moving?
- P(broken & no motion)
  = P(broken) * P(no motion | broken)
    - After lab 8 = 0.5*0.96 = 0.48

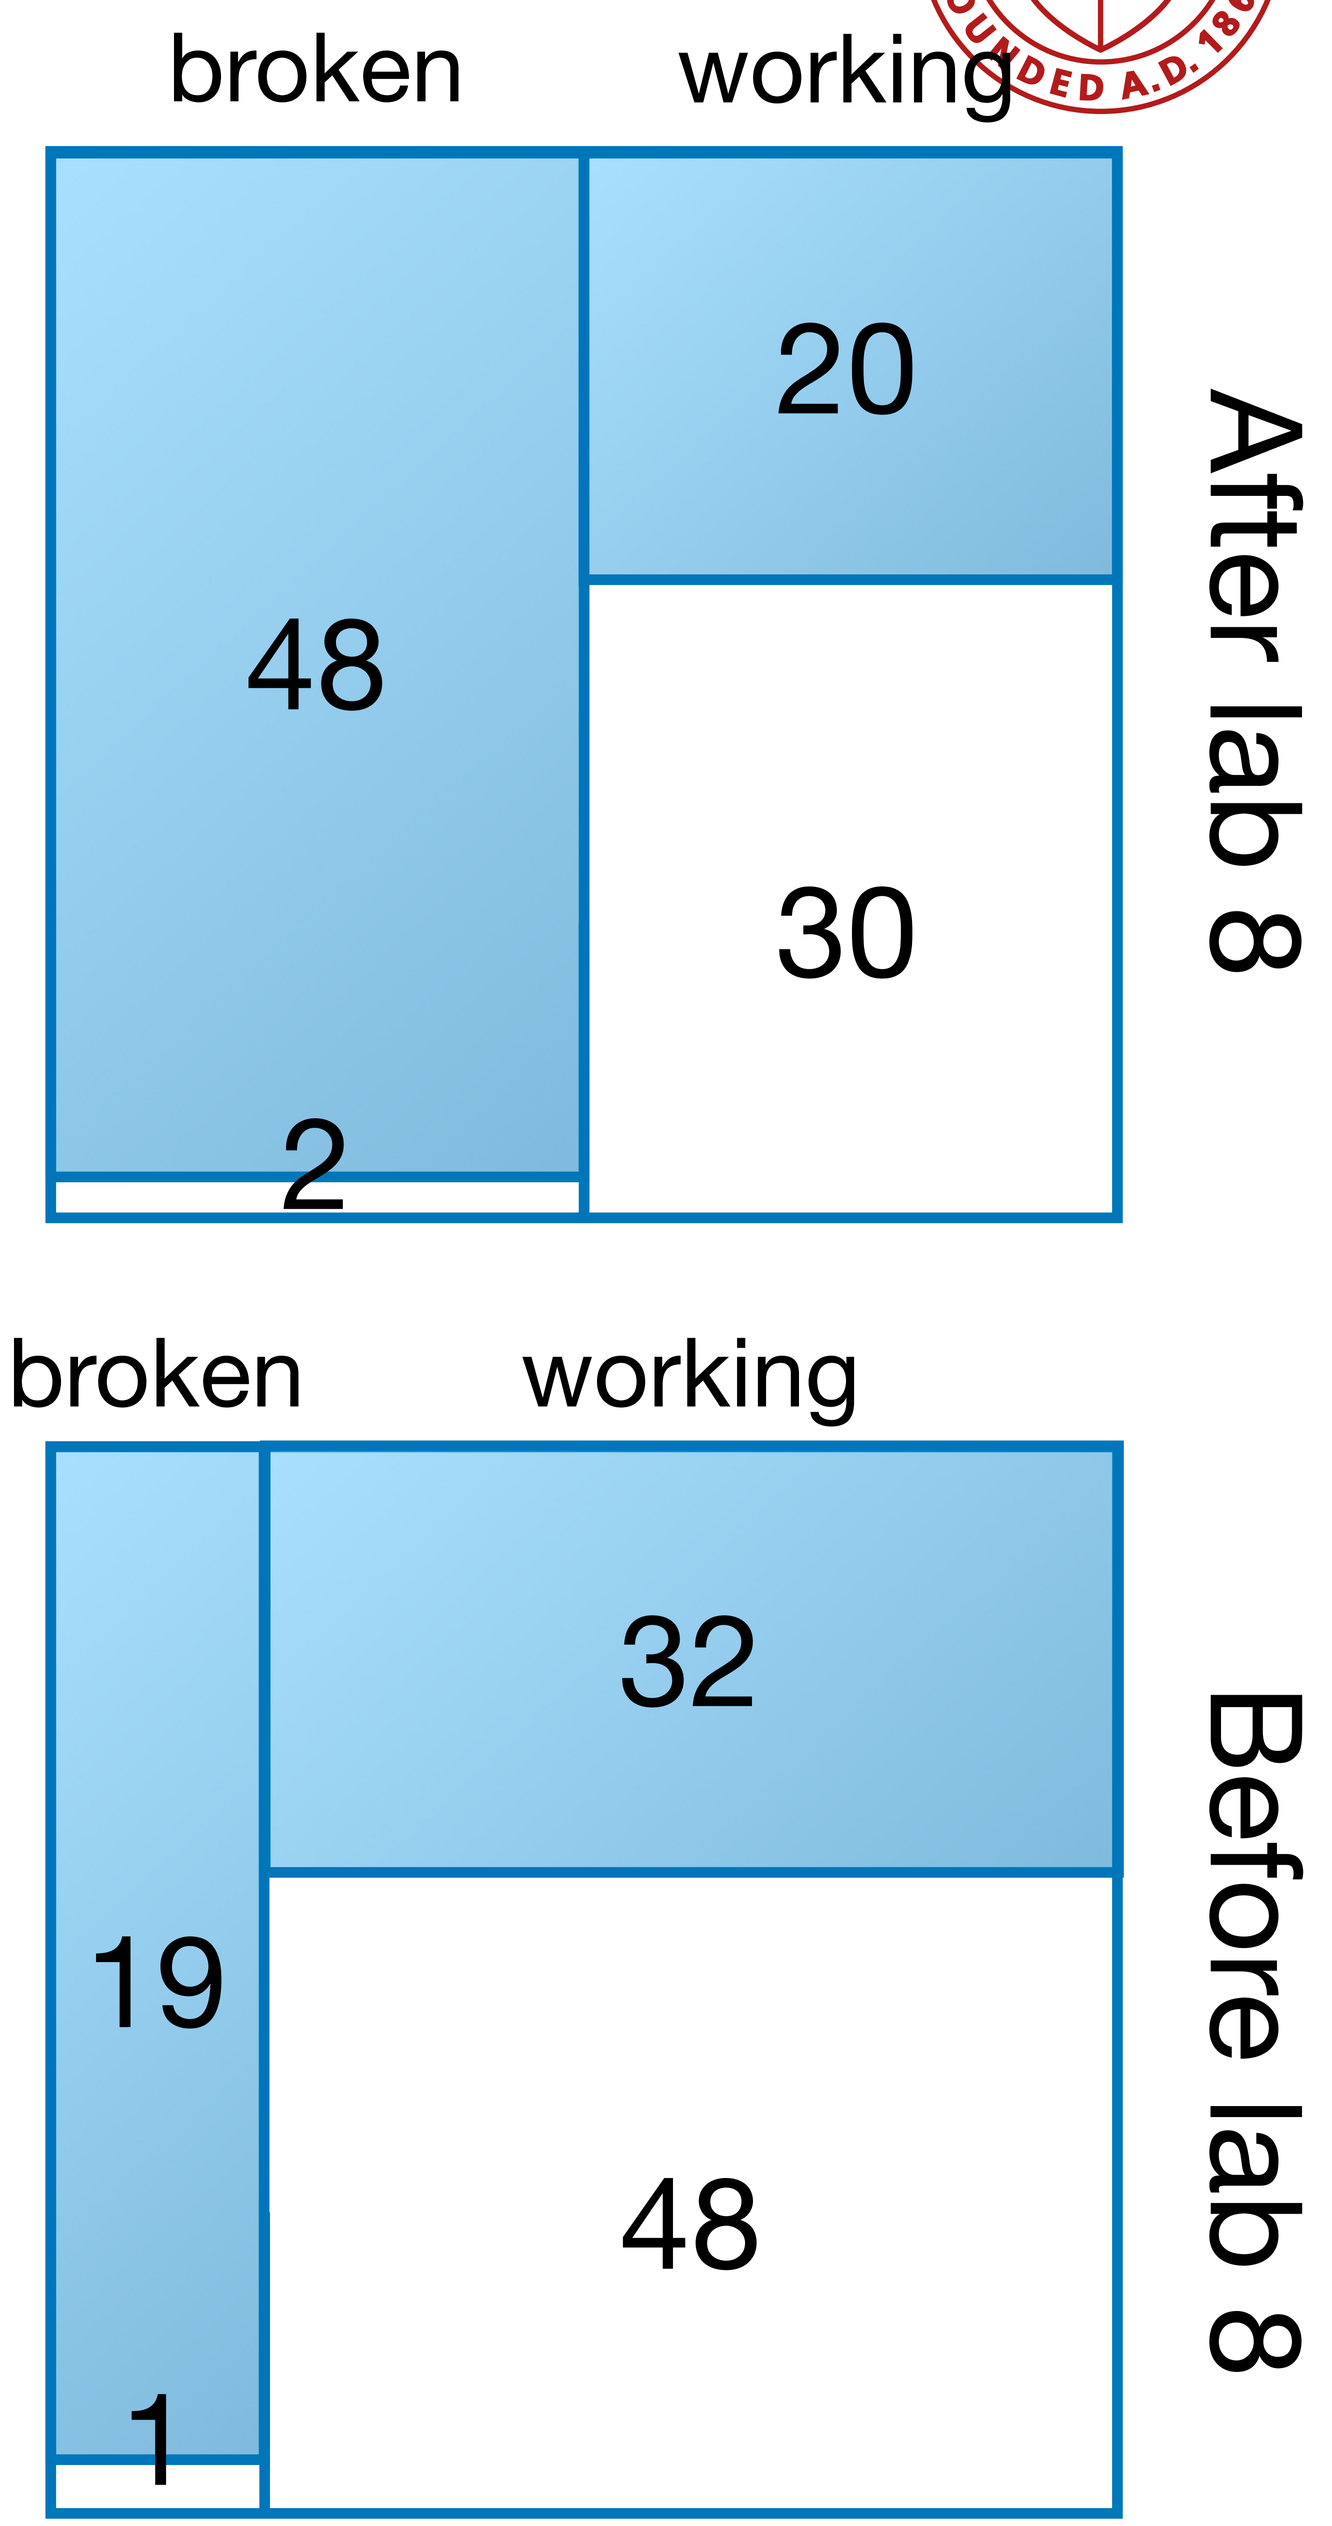

Bayesian Inference

Joint Probability

- What is the probability that the robot is both broken and not moving?
- P(broken & no motion)
  = P(broken) * P(no motion | broken)
  - Before lab 8 =
- P(working & motion)
  = P(working) * P(motion | working)
  - Before lab 8 = 0.8*0.6 = 0.48

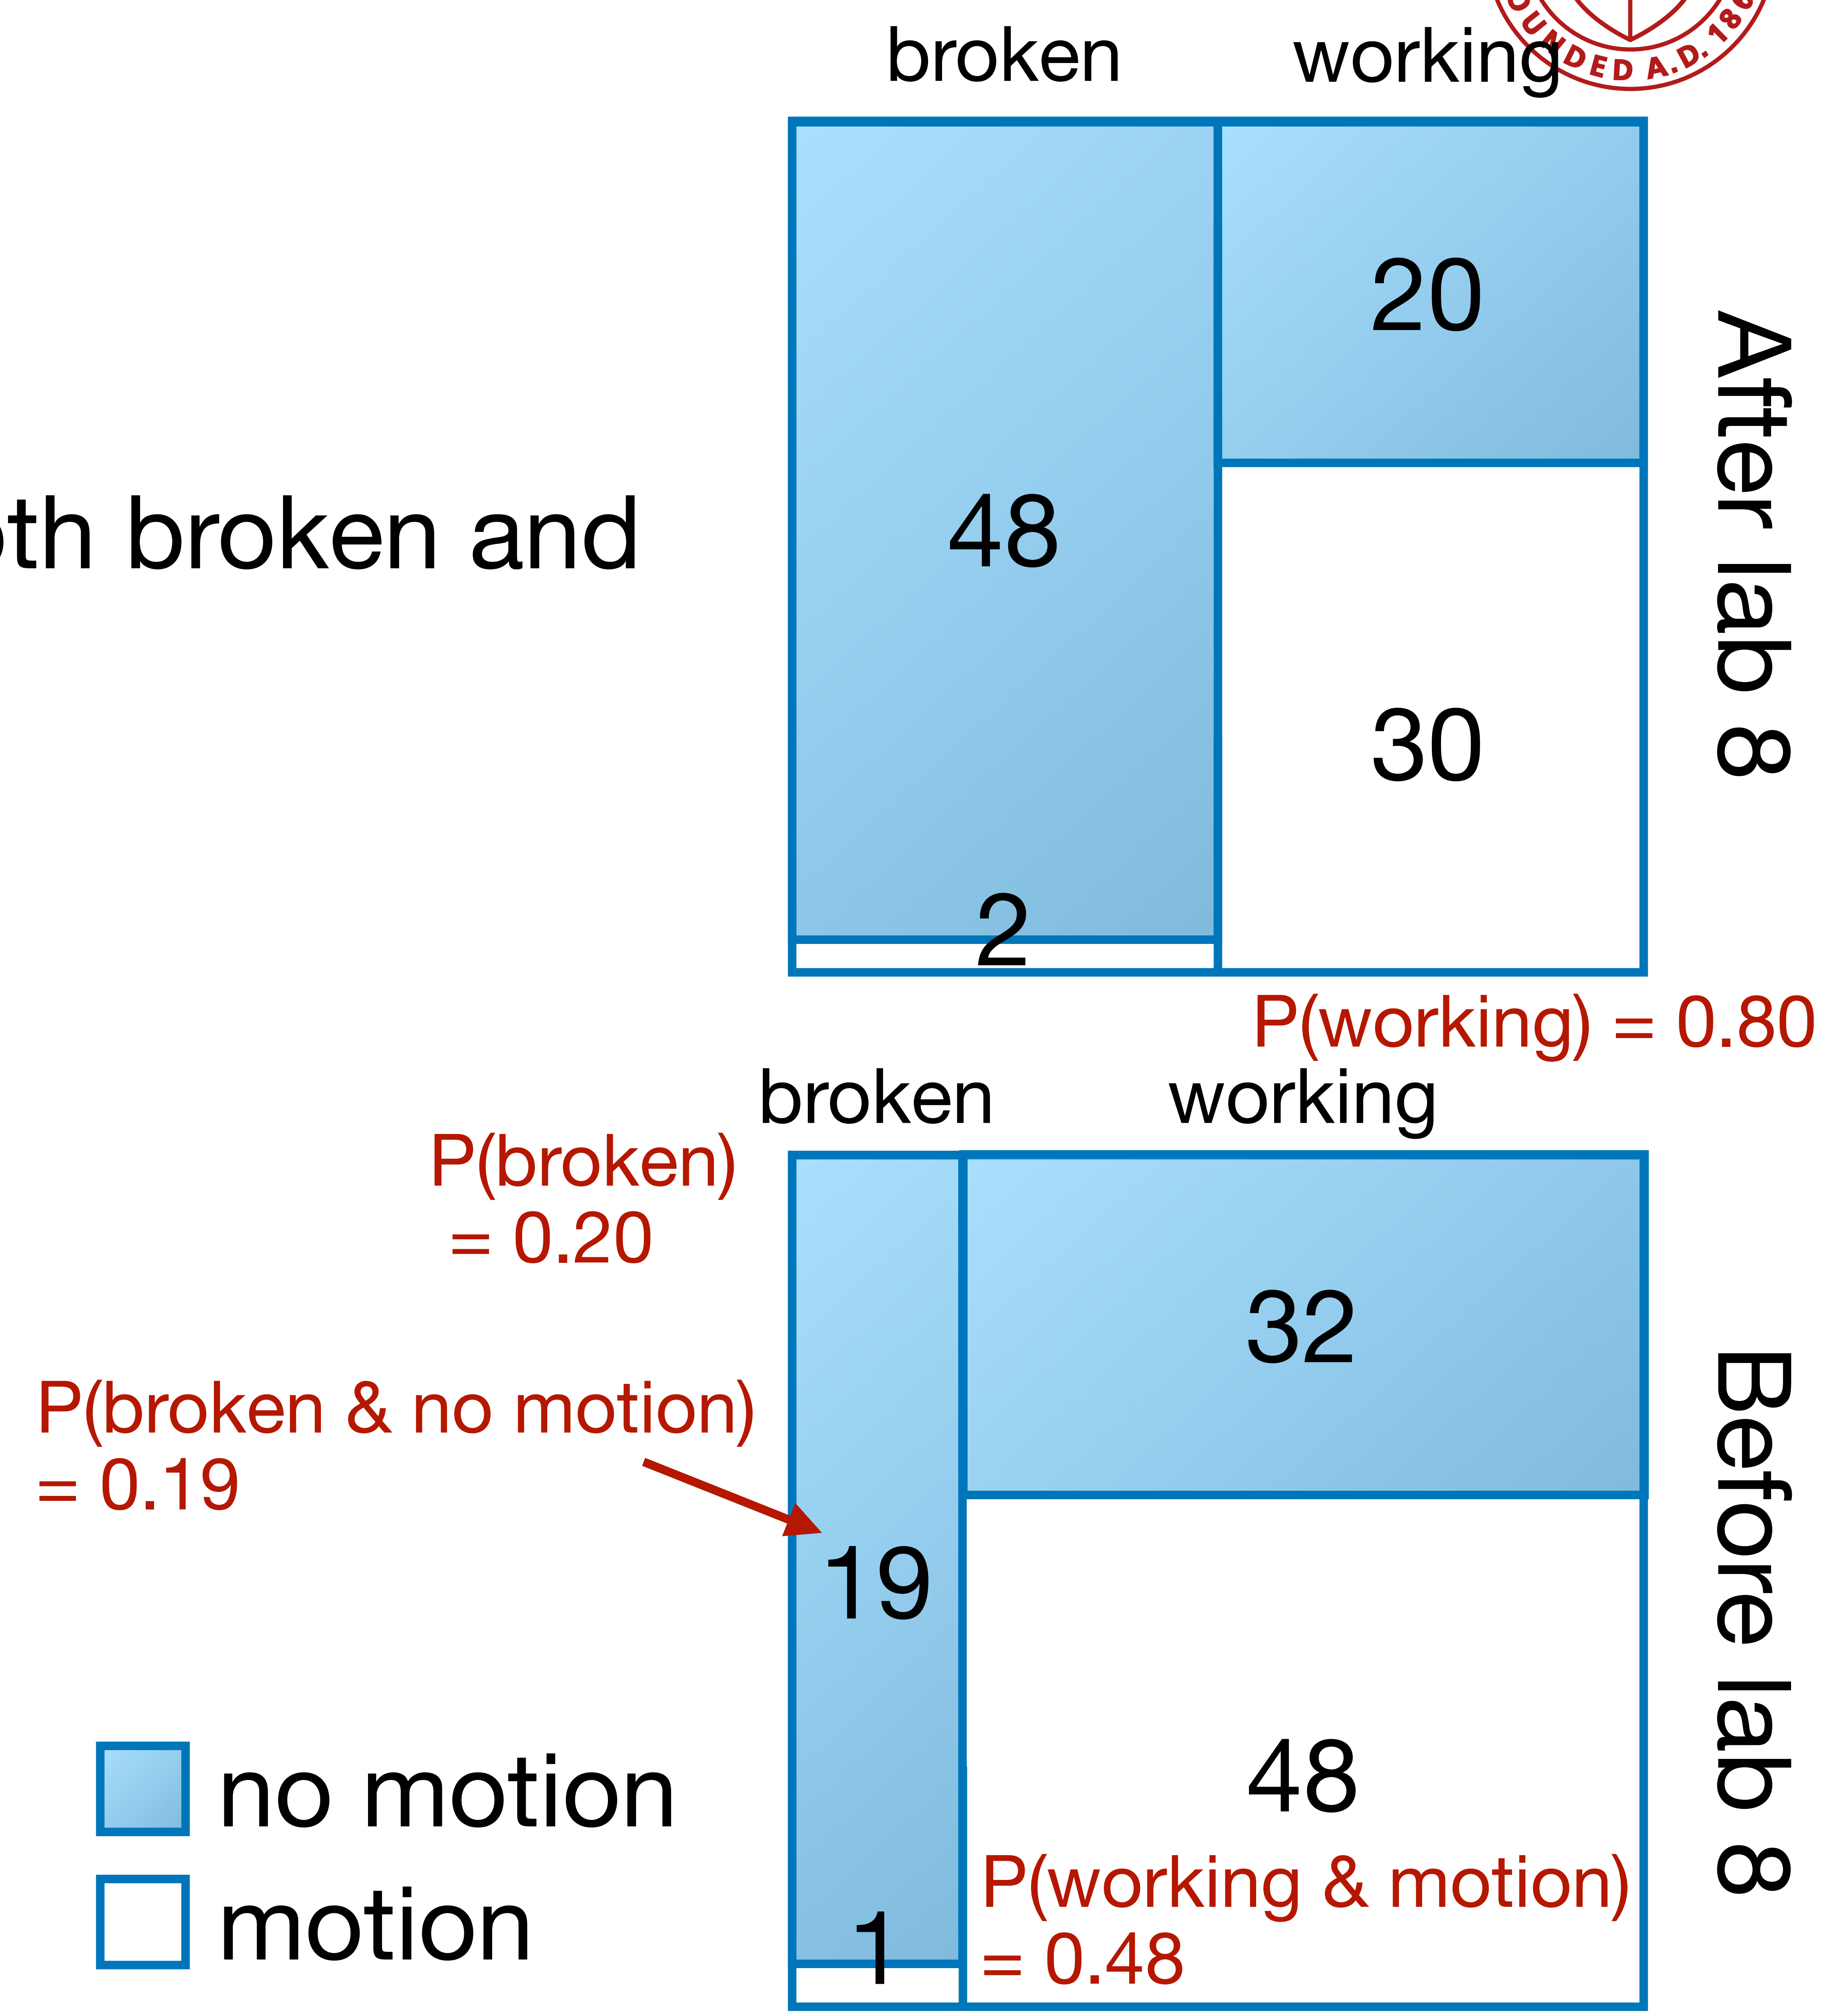

Bayesian Inference

Joint Probability

- What is the probability that the robot is broken and not moving?
- P(A, B) = P(A∩B) = P(A and B)
- P(A, B) = P(A)*P(B|A)
- P(A, B) = P(B, A)

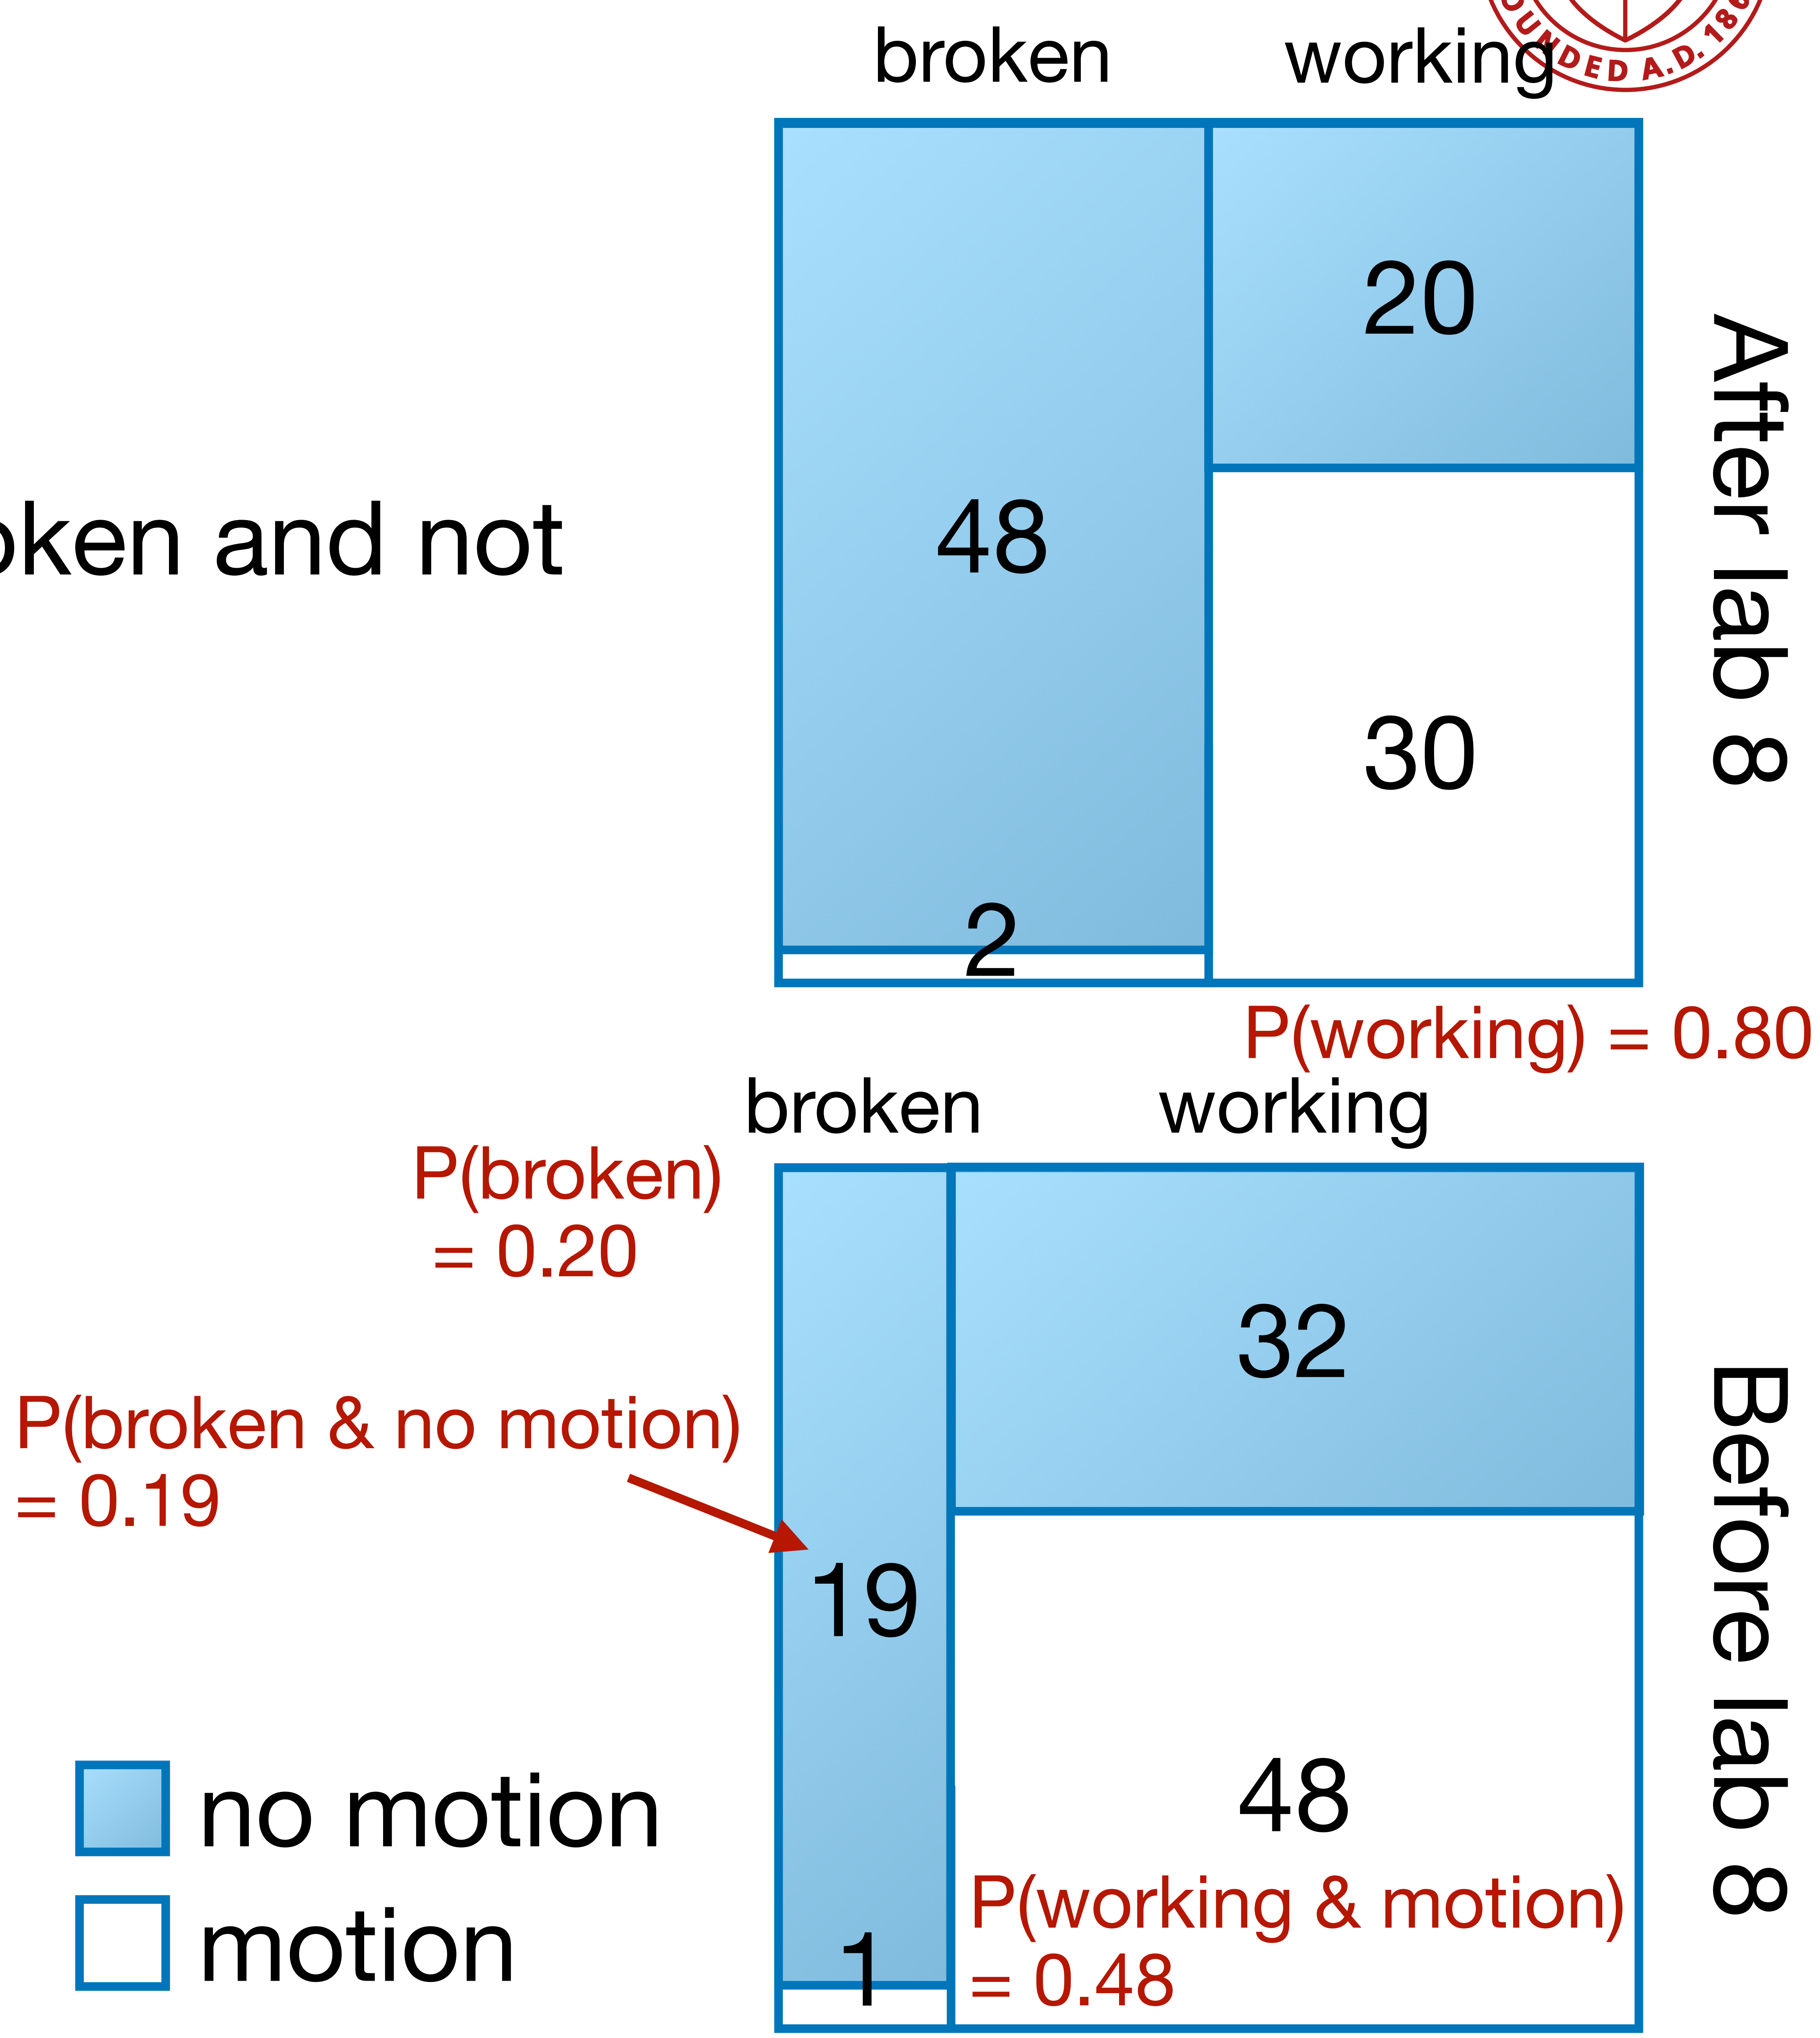

Bayesian Inference

Marginal Probability

- P(motion)
- P(no motion)

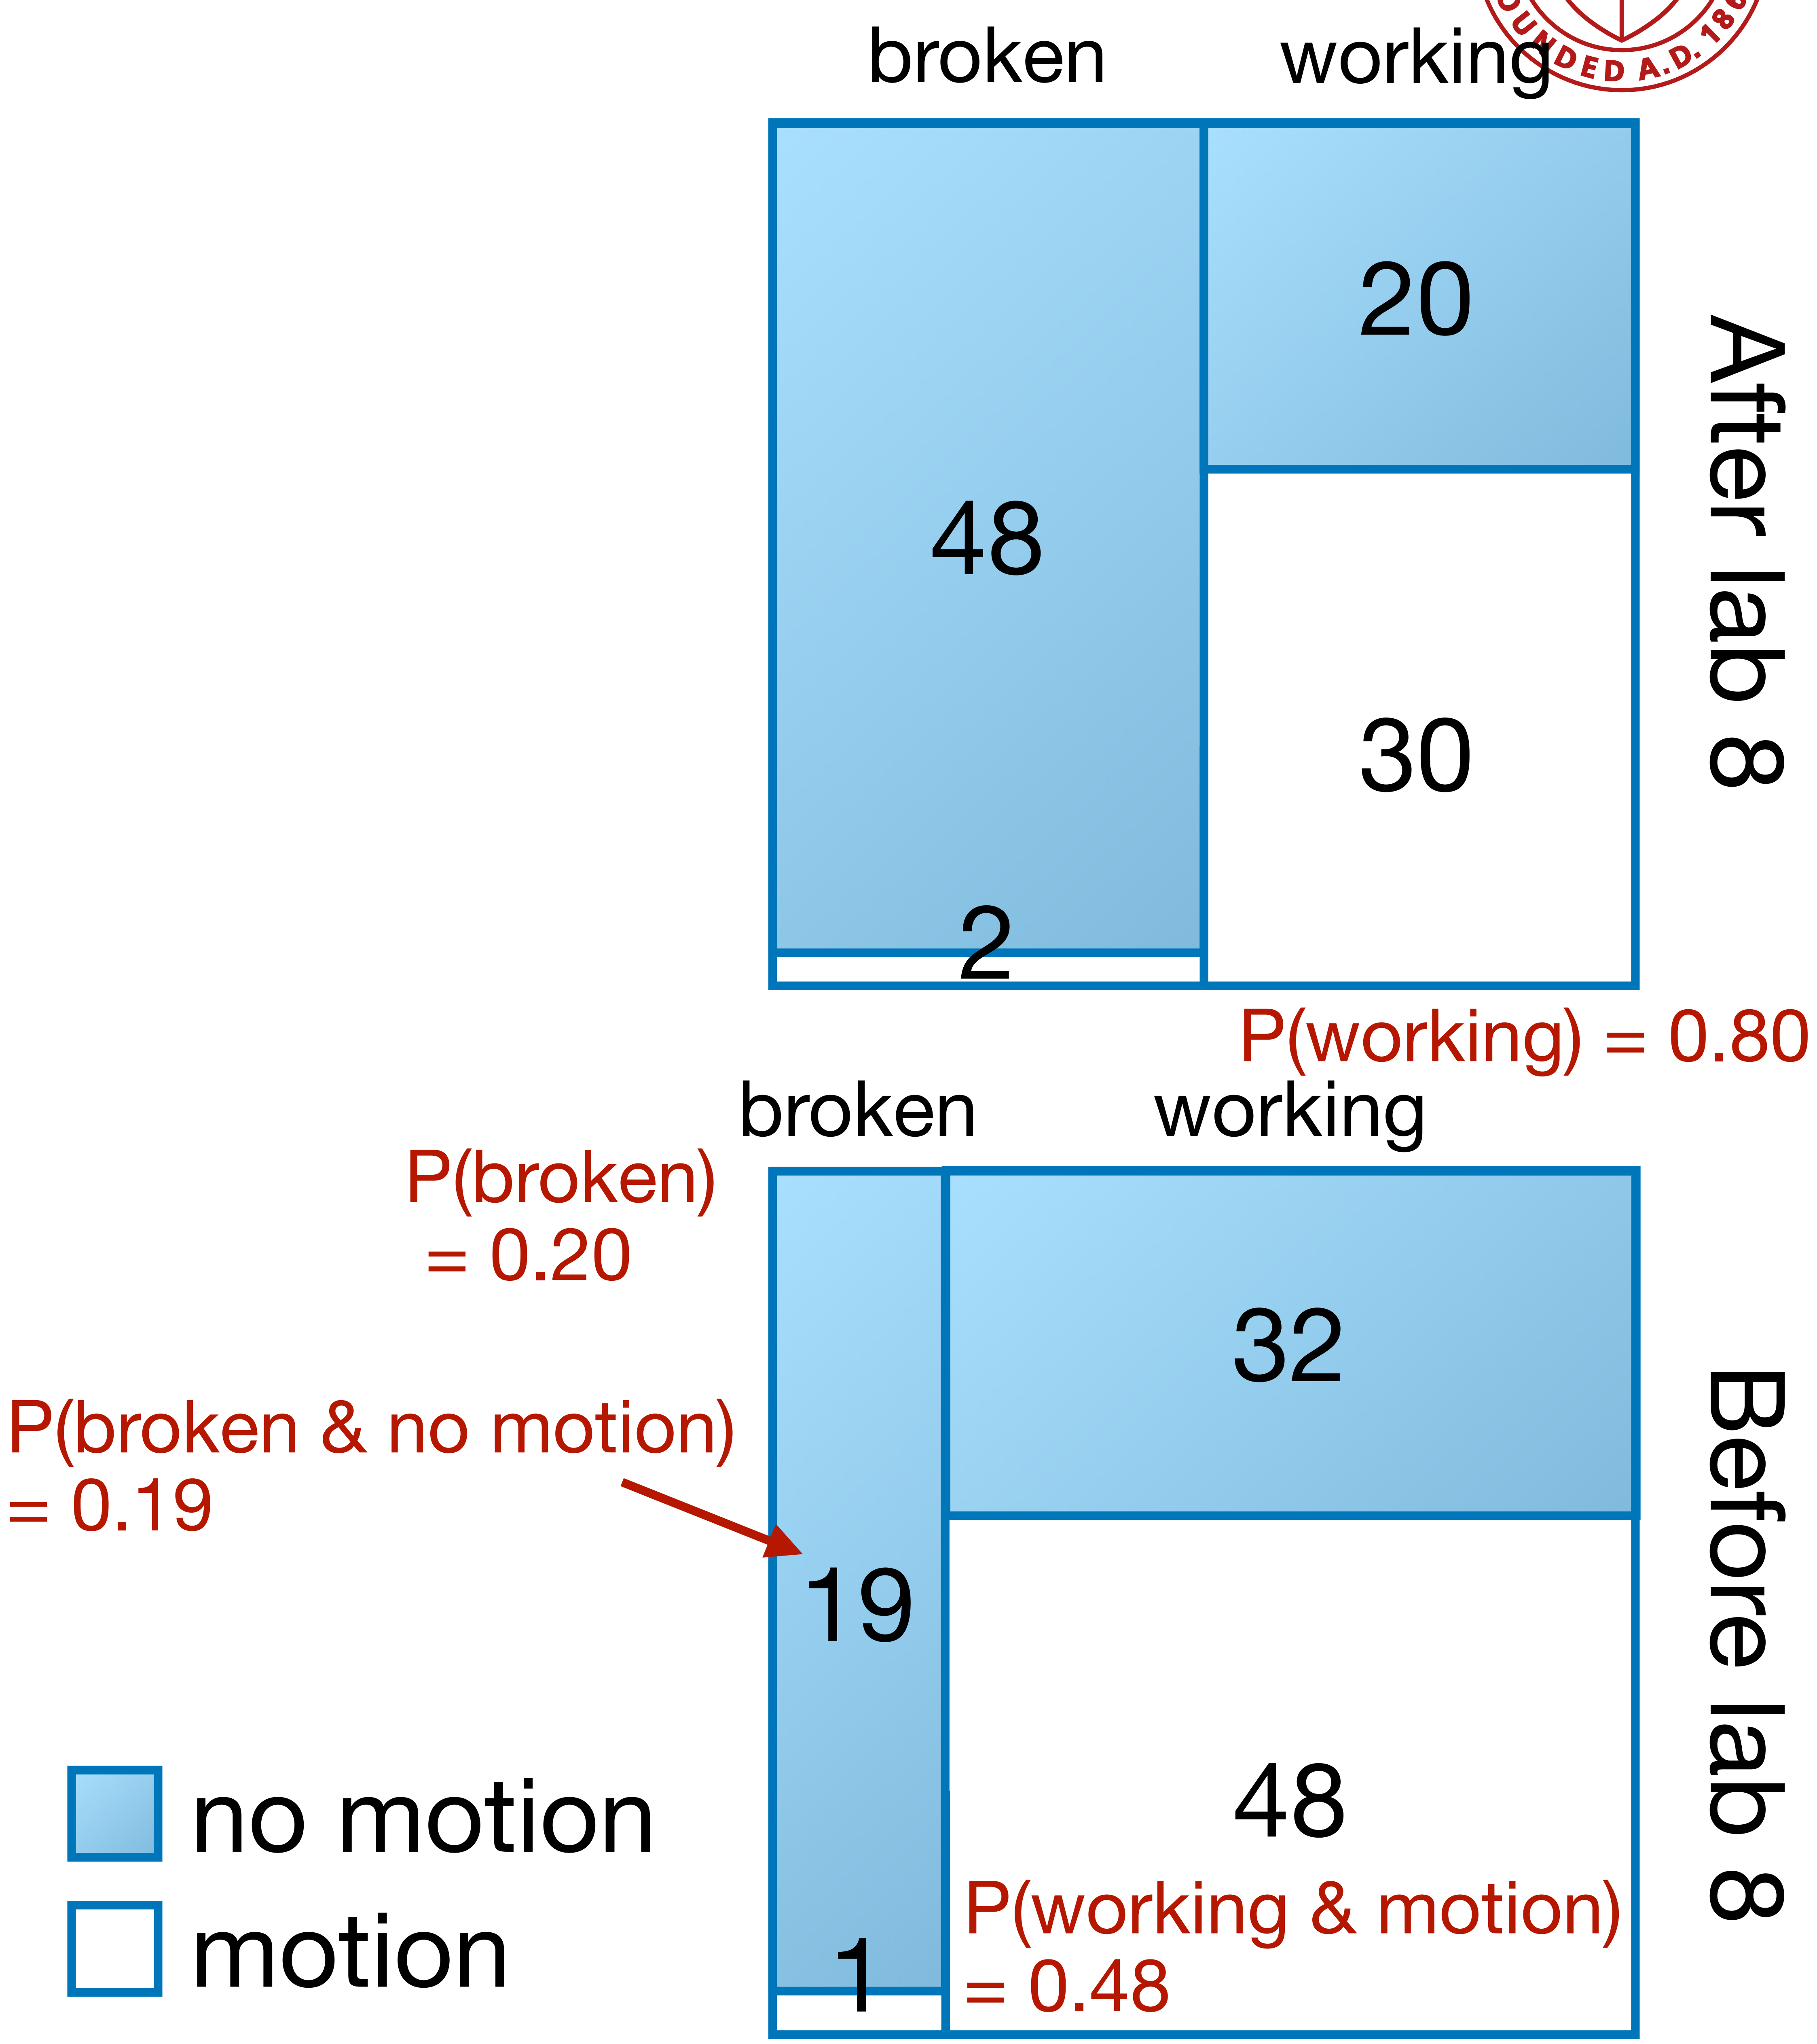

Bayesian Inference

- “My robot stopped moving, the hardware is broken, send me new parts”
- What is the probability that the robot is broken given that it stopped moving?
- P(broken | no motion)
- P(broken & no motion)

- P(no motion & broken)

- $P(\text{broken} \mid \text{no motion}) = \dfrac{P(\text{broken})^{*}\, P(\text{no motion} \mid \text{broken})}{P(\text{no motion})}$

- Before lab 8 = 0.2*0.96/0.51 = 0.38
- After lab 8 = 0.5*0.96/0.68 = 0.71

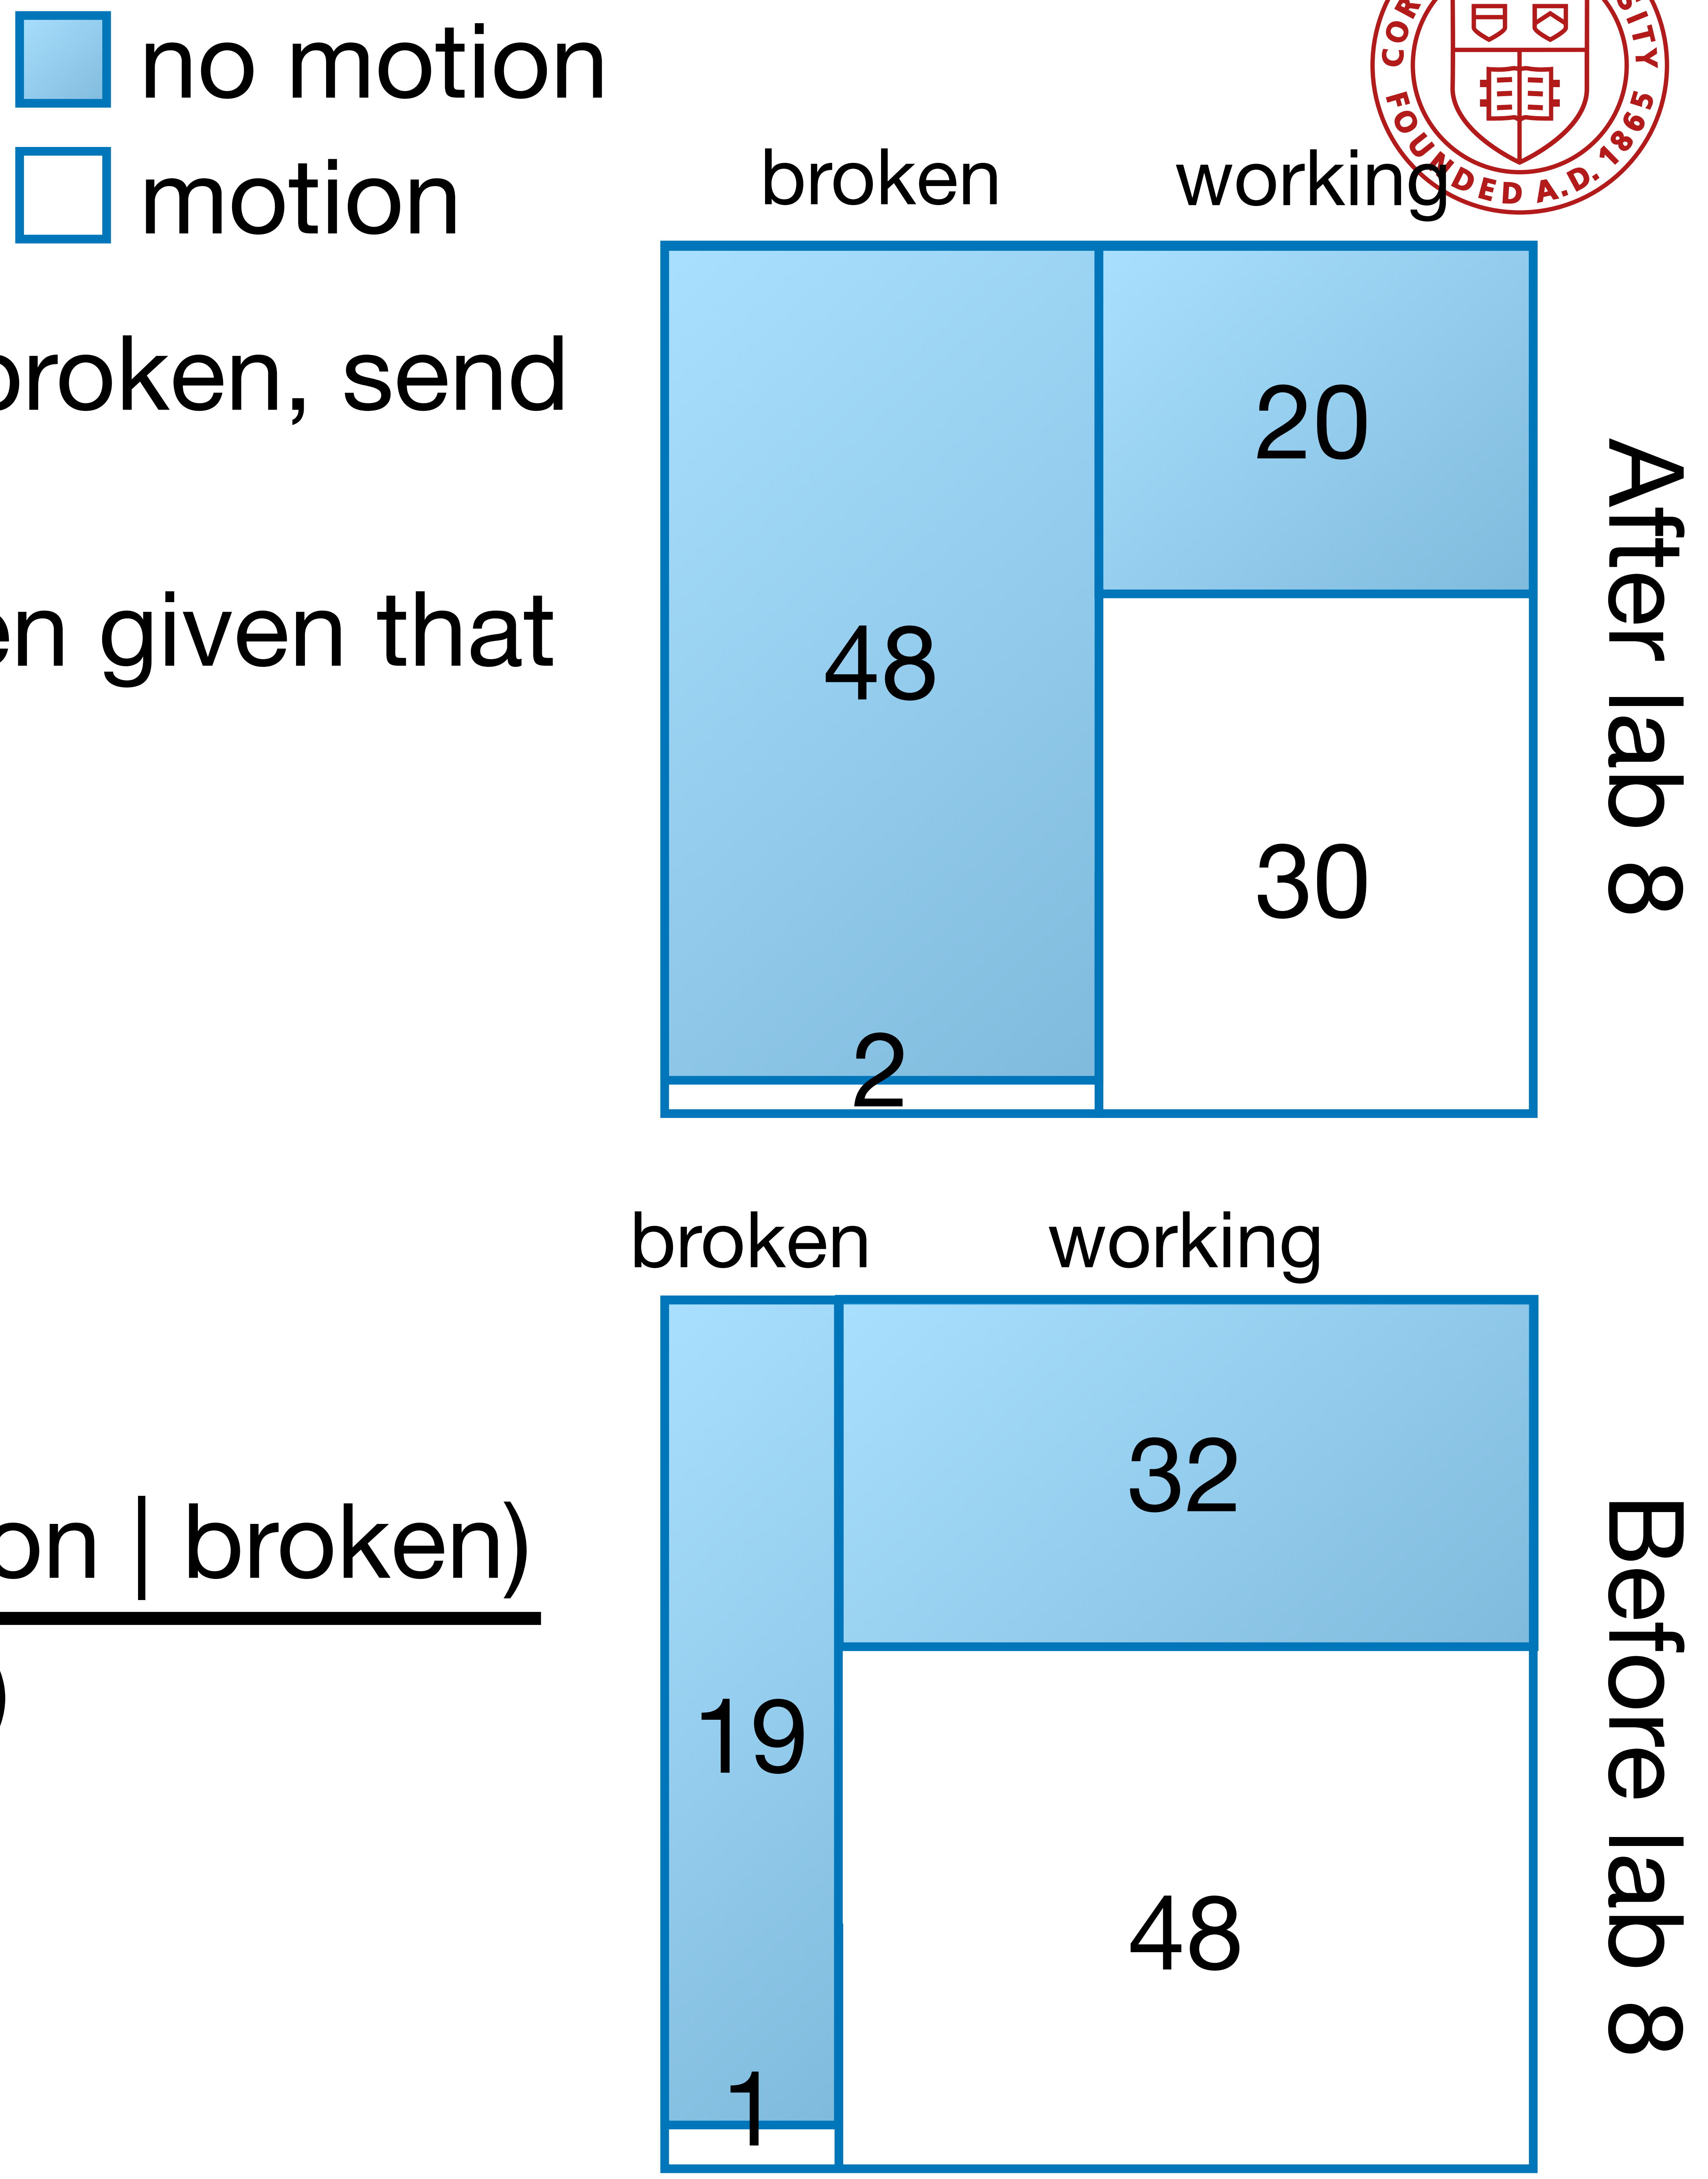

Bayesian Inference

likelihood — prior

$$P(x|z) = \frac{P(z|x)P(x)}{P(z)}$$

posterior

marginal likelihood (constant)

z = Sensor data

x = Robot state /location

Bayesian Inference

Beliefs

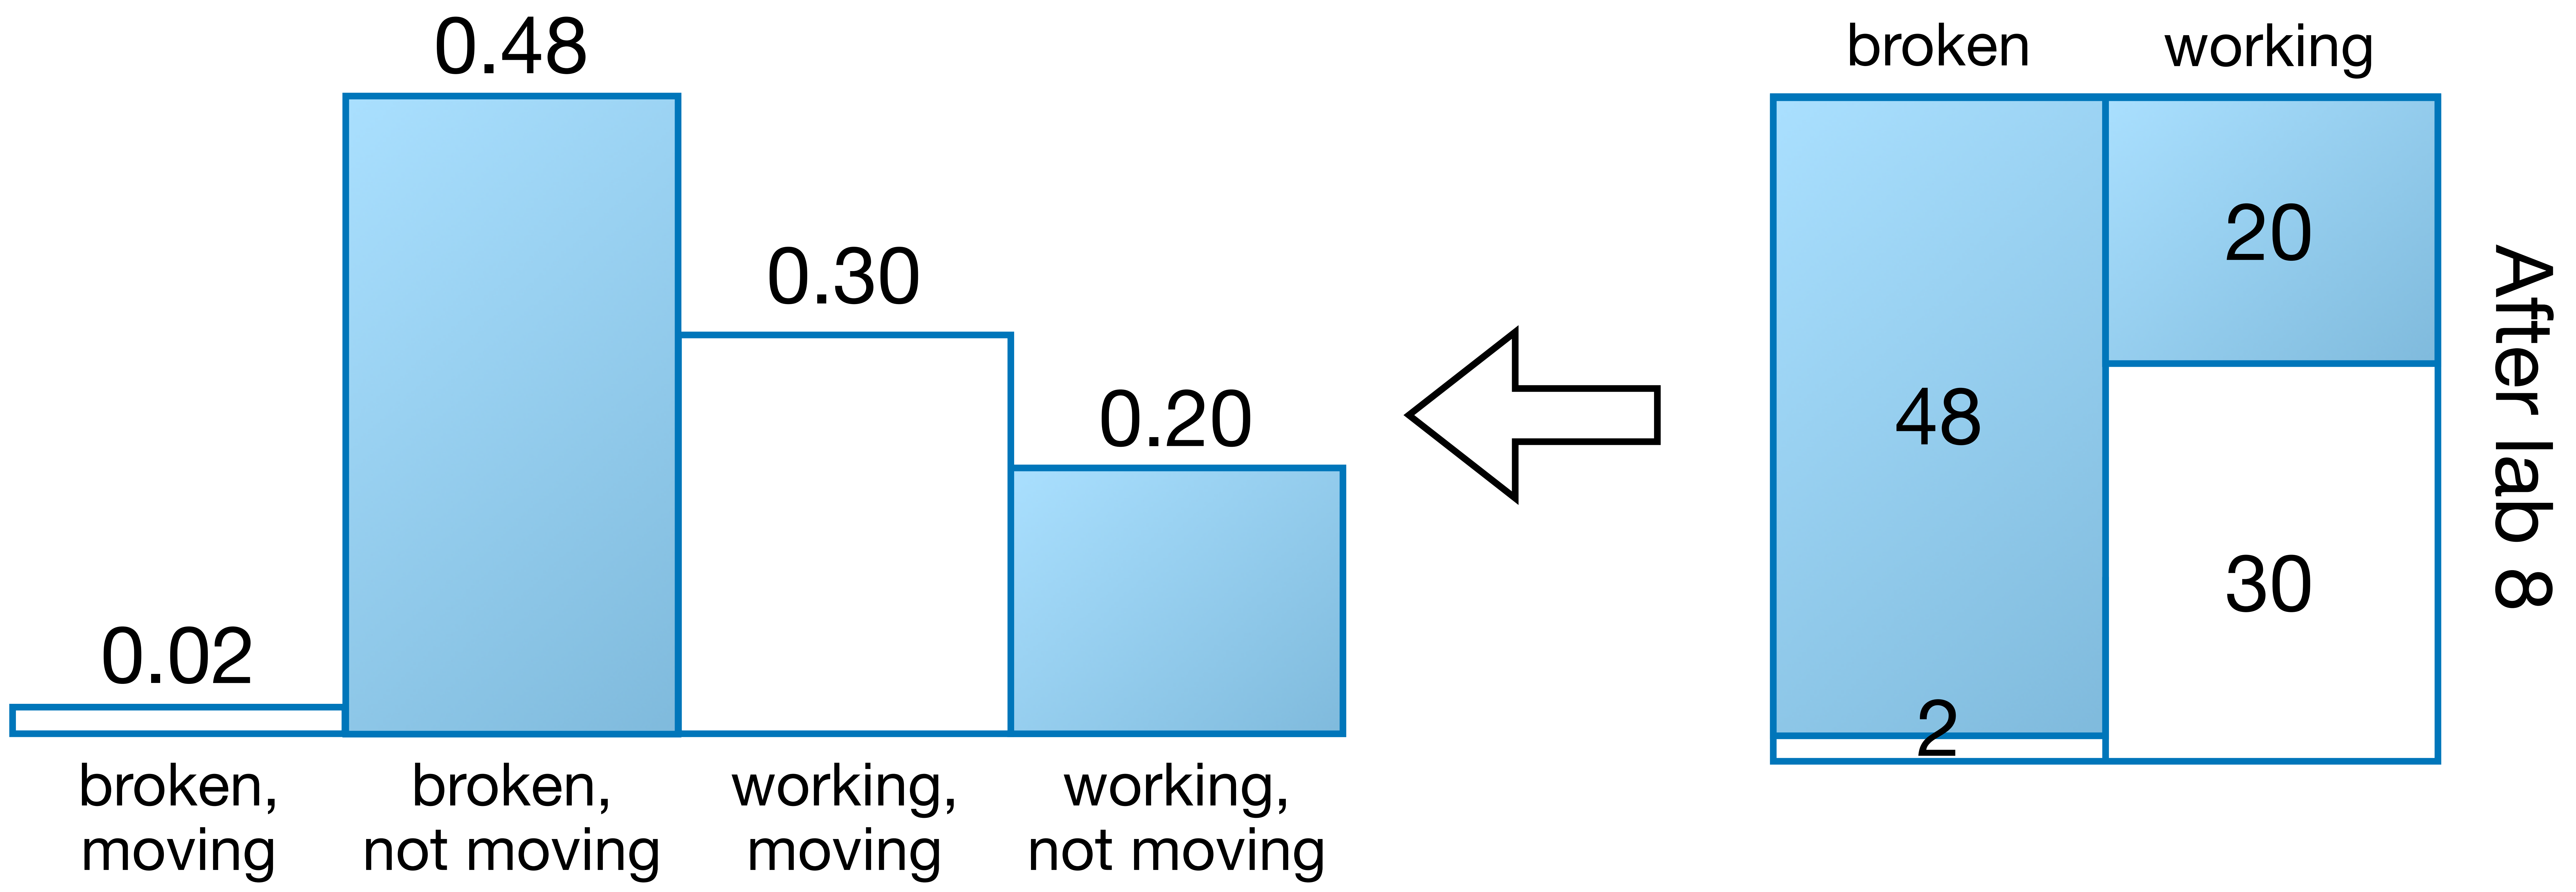

Bayesian Inference

Beliefs

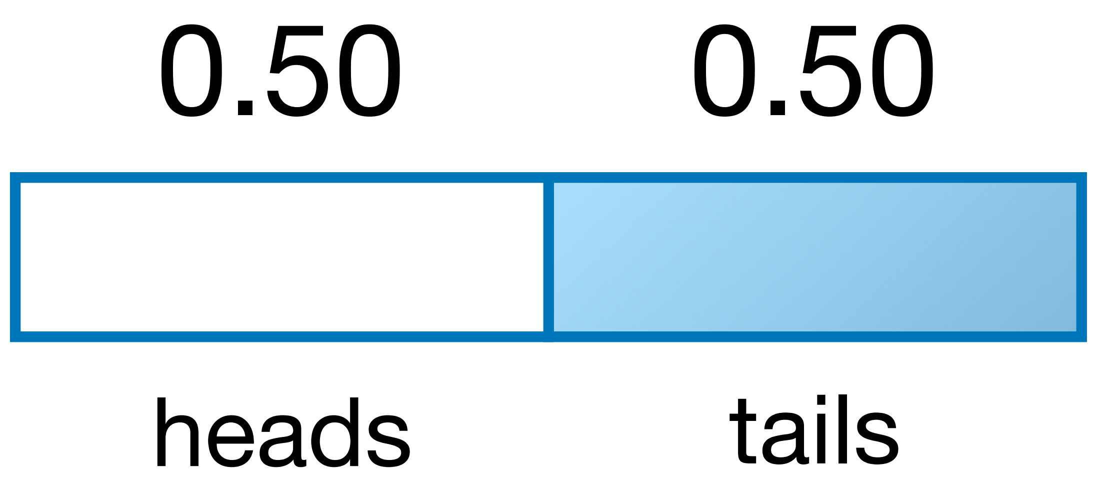

Bayesian Inference

Beliefs

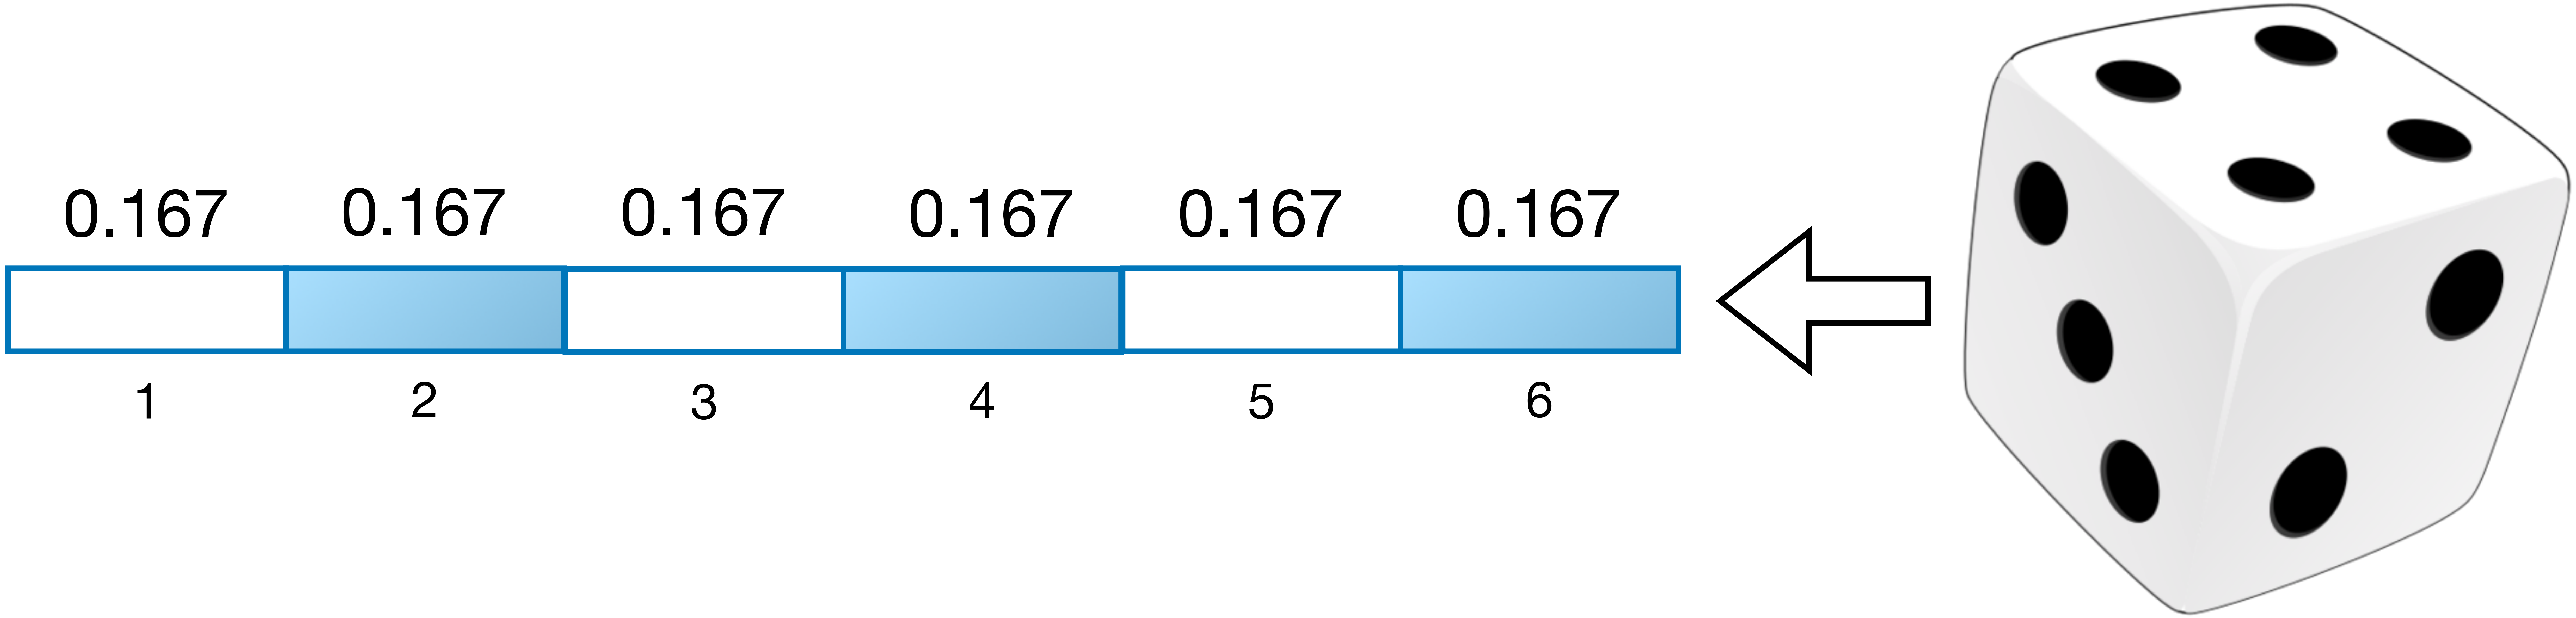

Bayesian Inference

Beliefs

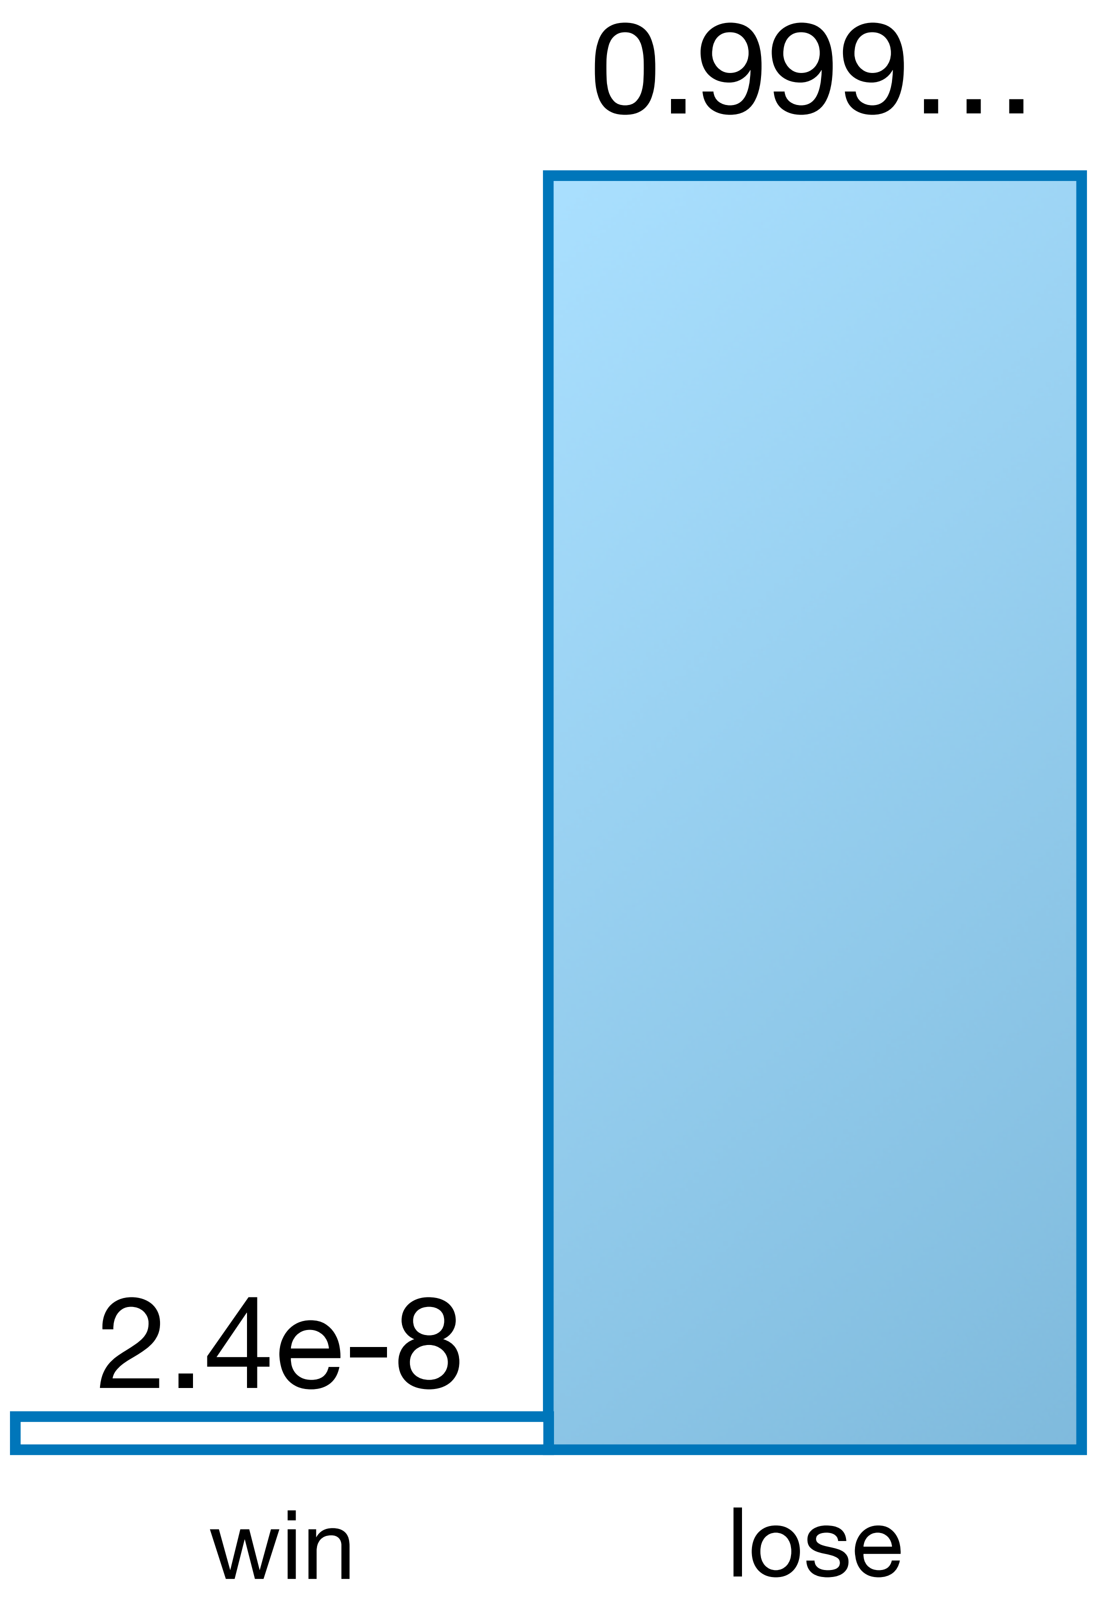

Bayesian Inference

Beliefs

- Discrete —> continuous **probability distributions**
  - Mean, median, most common value, etc.

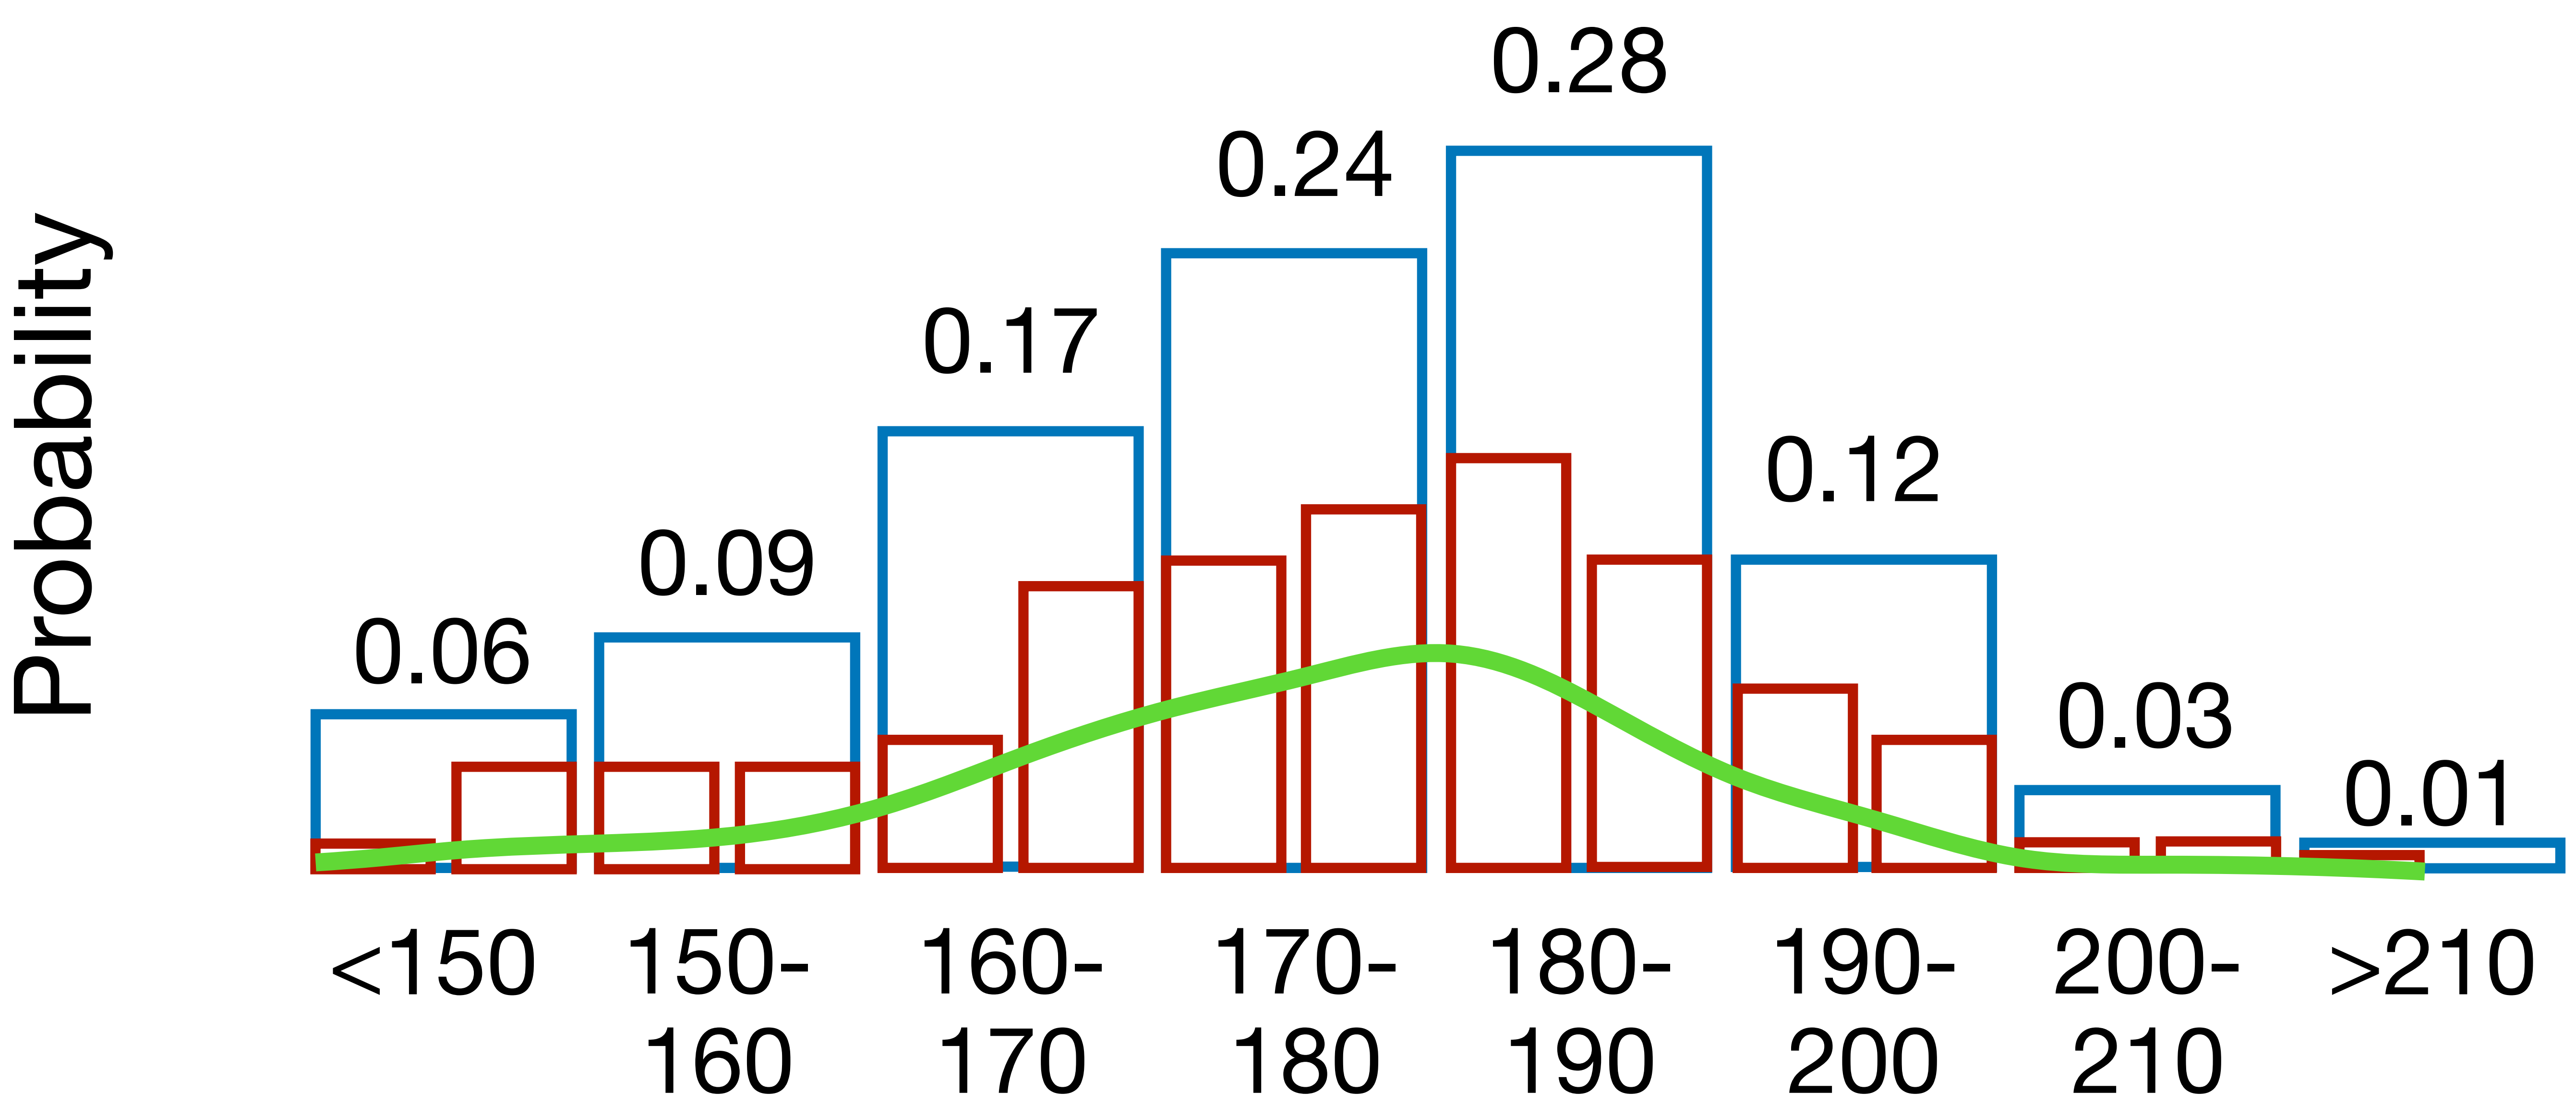

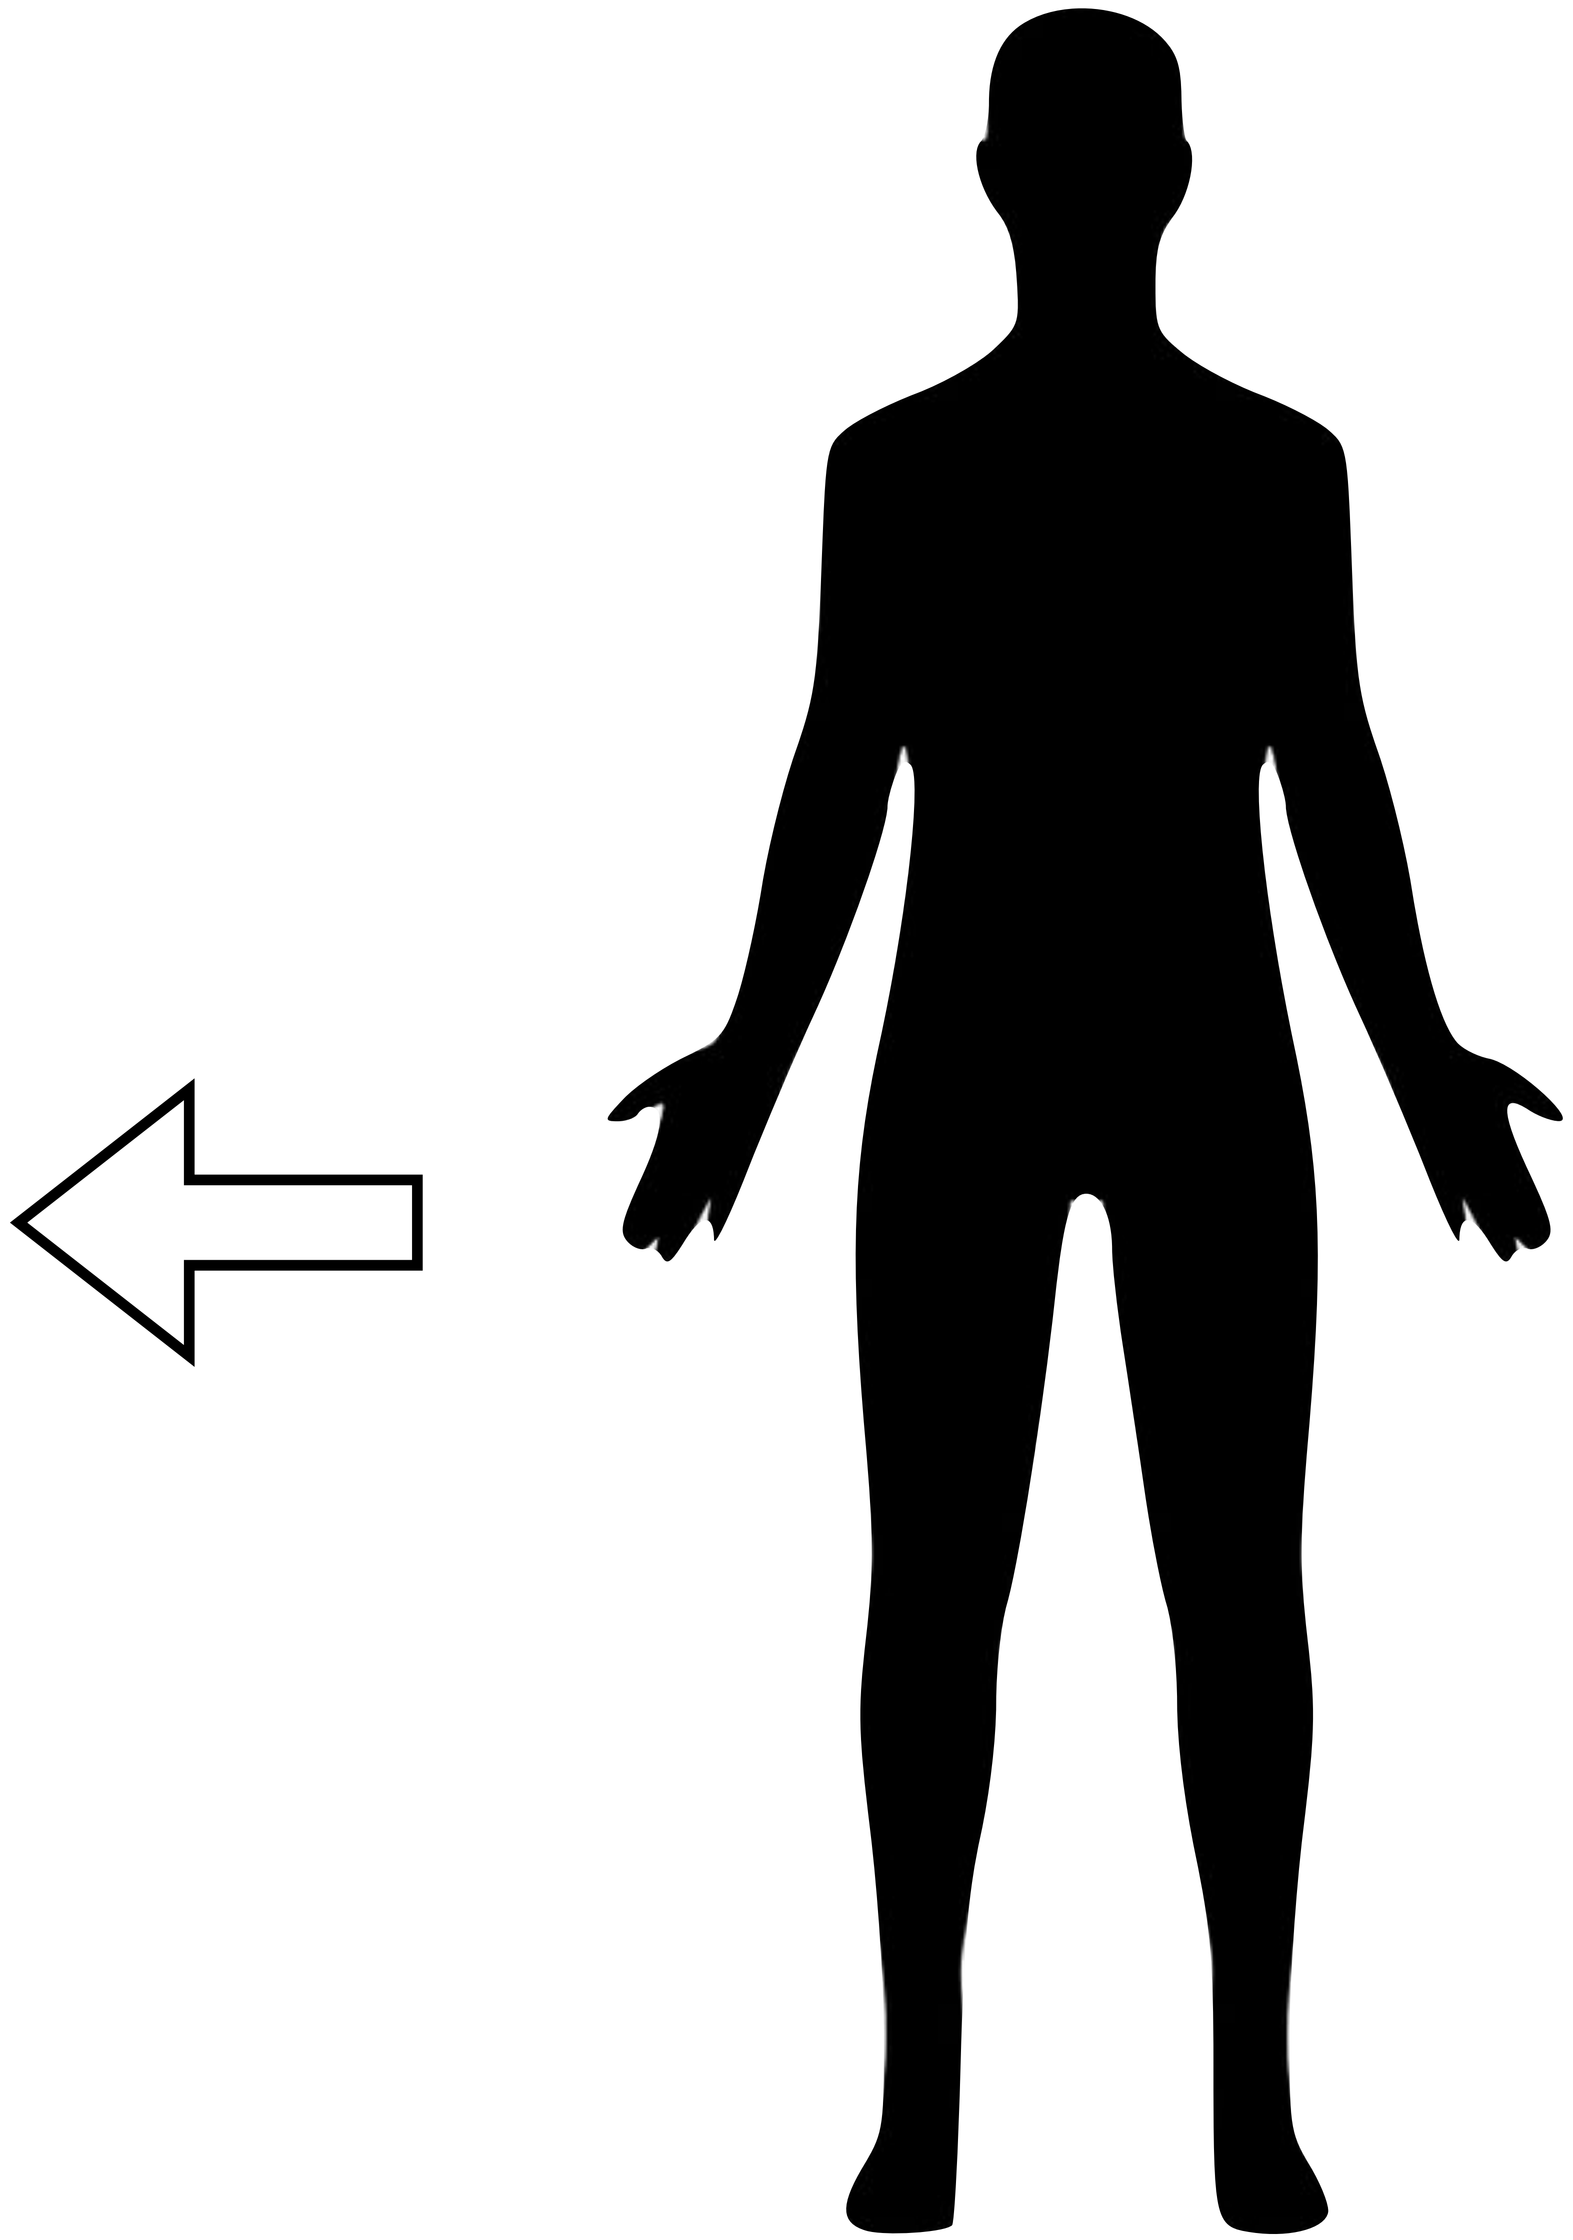

Probability Distributions

- What is the maximum speed of your robot?
  - Your speed is 8.8 ft/s, 6.6 ft/s, 8.33 ft/s, but what is the actual value?
- Frequentist statistics
  - Mean: $\mu = (8.8+6.6+8.33)/3 = 7.91$ ft/s
  - Variance: $\sigma^2 = ((8.8-7.91)^2+(6.6-7.91)^2+(8.33-7.91)^2)/(3-1) = 1.35$ ft/s
  - Standard deviation: $\sigma = \sqrt{\sigma^2} = 1.16$ ft/s
  - Standard error: $\sigma/\sqrt{3} = 0.67$ ft/s
- Bayesian statistics
  - Probably 7.91 ft/s…

Values from lab 3 (2020)

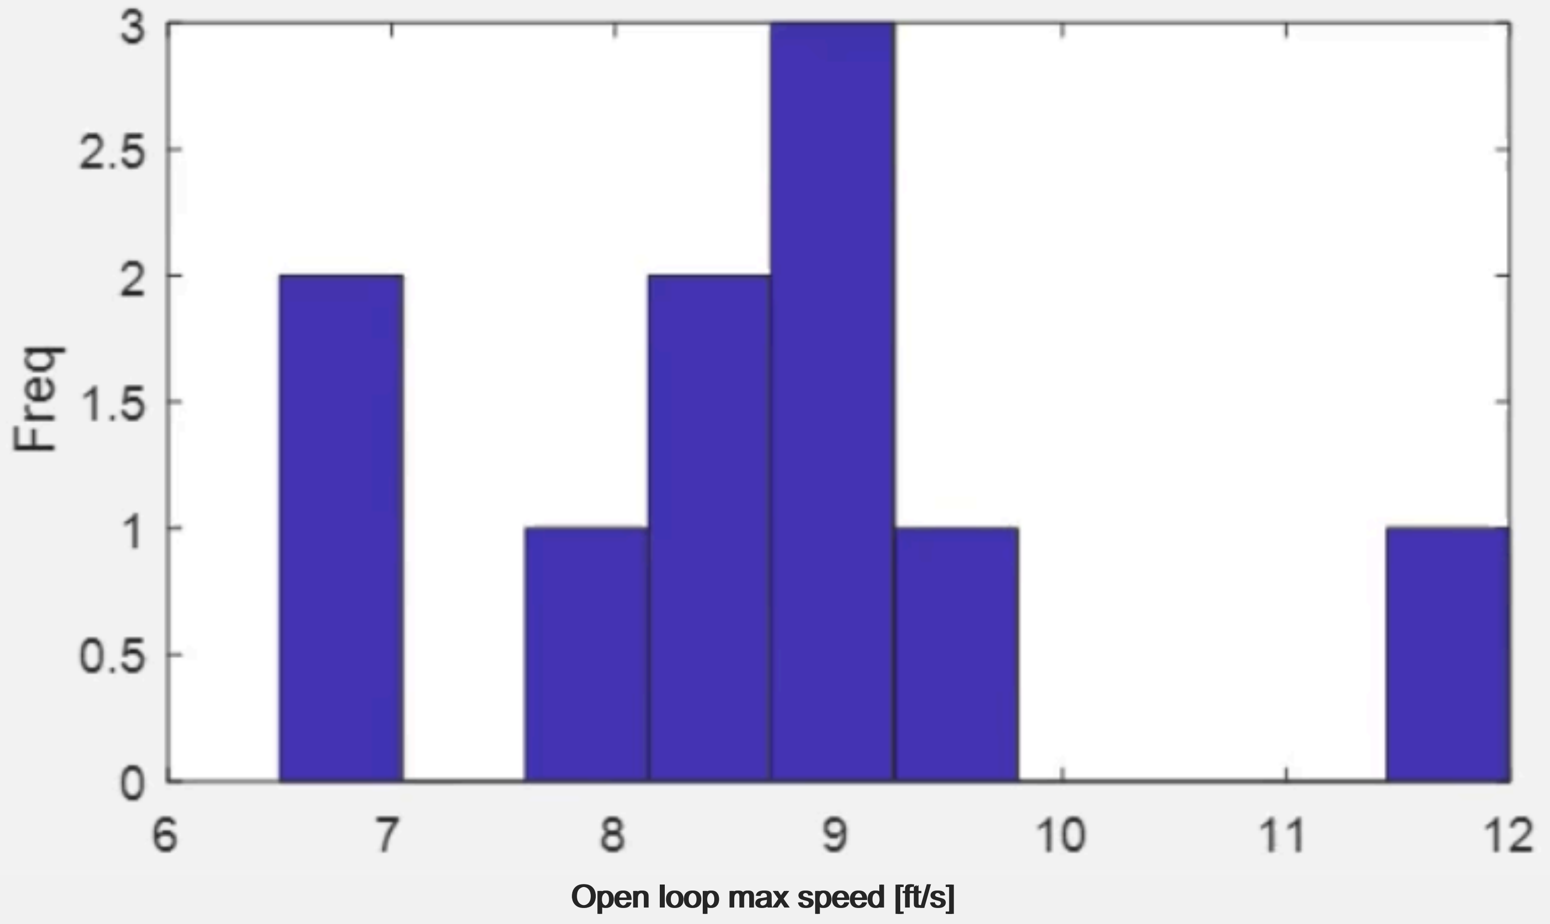

What you observe

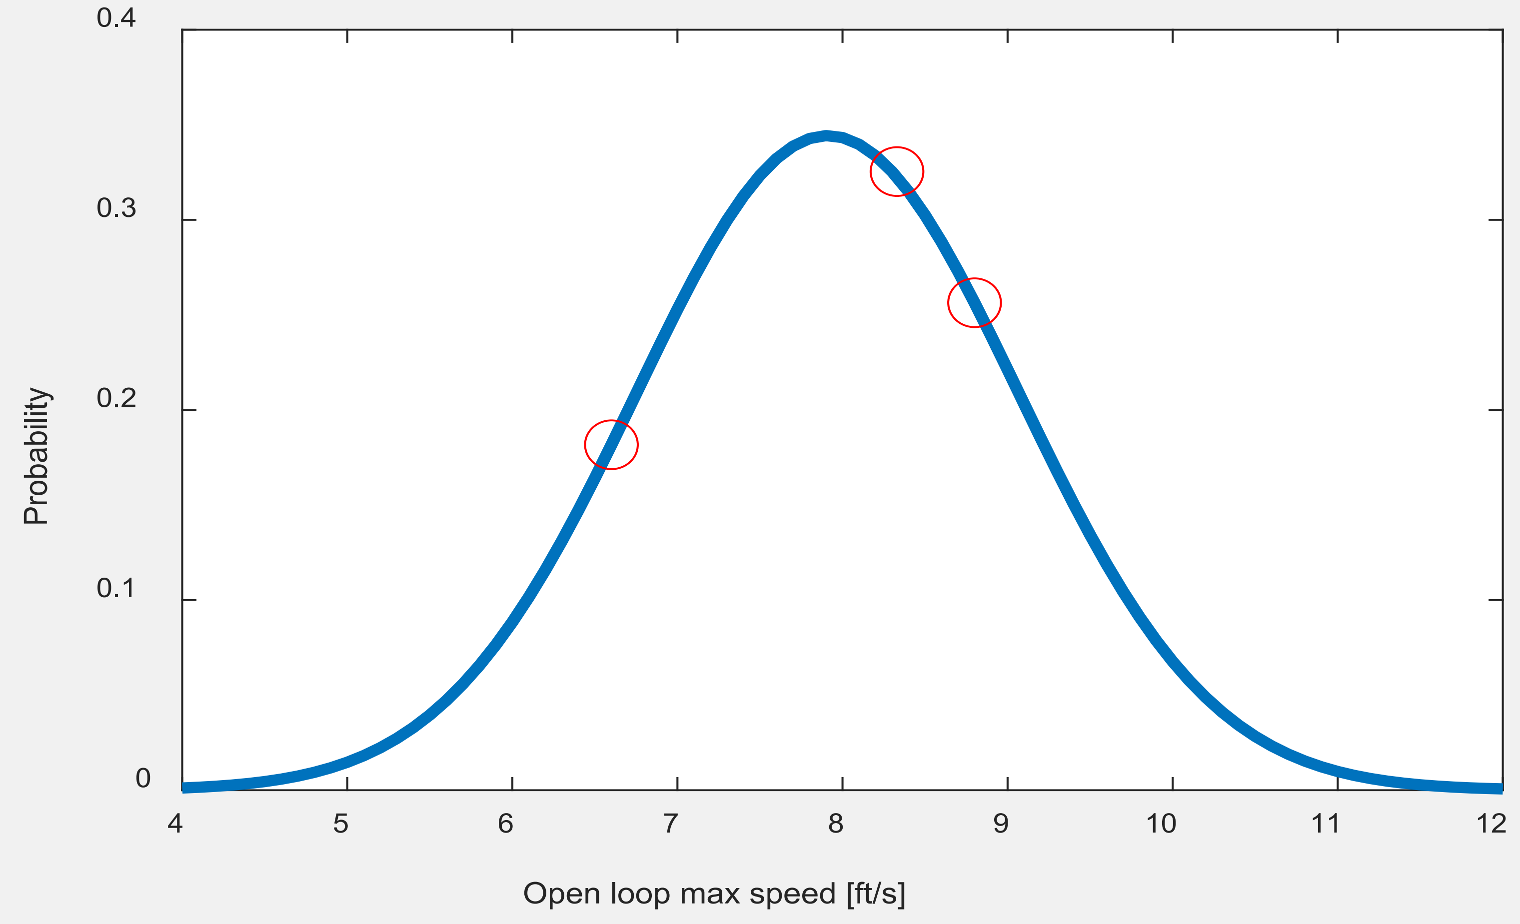

Probability Distributions

$$P(x|z) = \frac{P(z|x)P(x)}{P(z)}$$

- Use Bayes Theorem
- Instead of events x and z
  - Substitute "s" for the actual speed
  - Substitute "m" for the measurements
- $P(s)$ is our prior
- $P(m|s)$ is the likelihood associated with those measurements
- $P(s|m)$ is what we believe about the speed given those measurements
- $P(m)$ is the marginal likelihood
- Procedure:
  - Start with a belief
  - Update it
  - End up with a new belief!

Probability Distributions

$$P(s|m) = \frac{P(m|s)P(s)}{P(m)}$$

- Start by assuming nothing
  - P(s) uniform
  - P(s|m) = P(m|s)* c1/c2
  - Simplified: P(s|m) = P(m|s)
    - *Guess!* What if the actual max speed is 11 ft/s?
    - P(s = 11 | m = [6.6, 8.33, 8.8]) = P(m = [6.6, 8.33, 8.8] | s=11)
    - P(m = 6.6| s = 11)* P(m = 8.33 | s = 11)* P(m = 8.8 | s = 11)

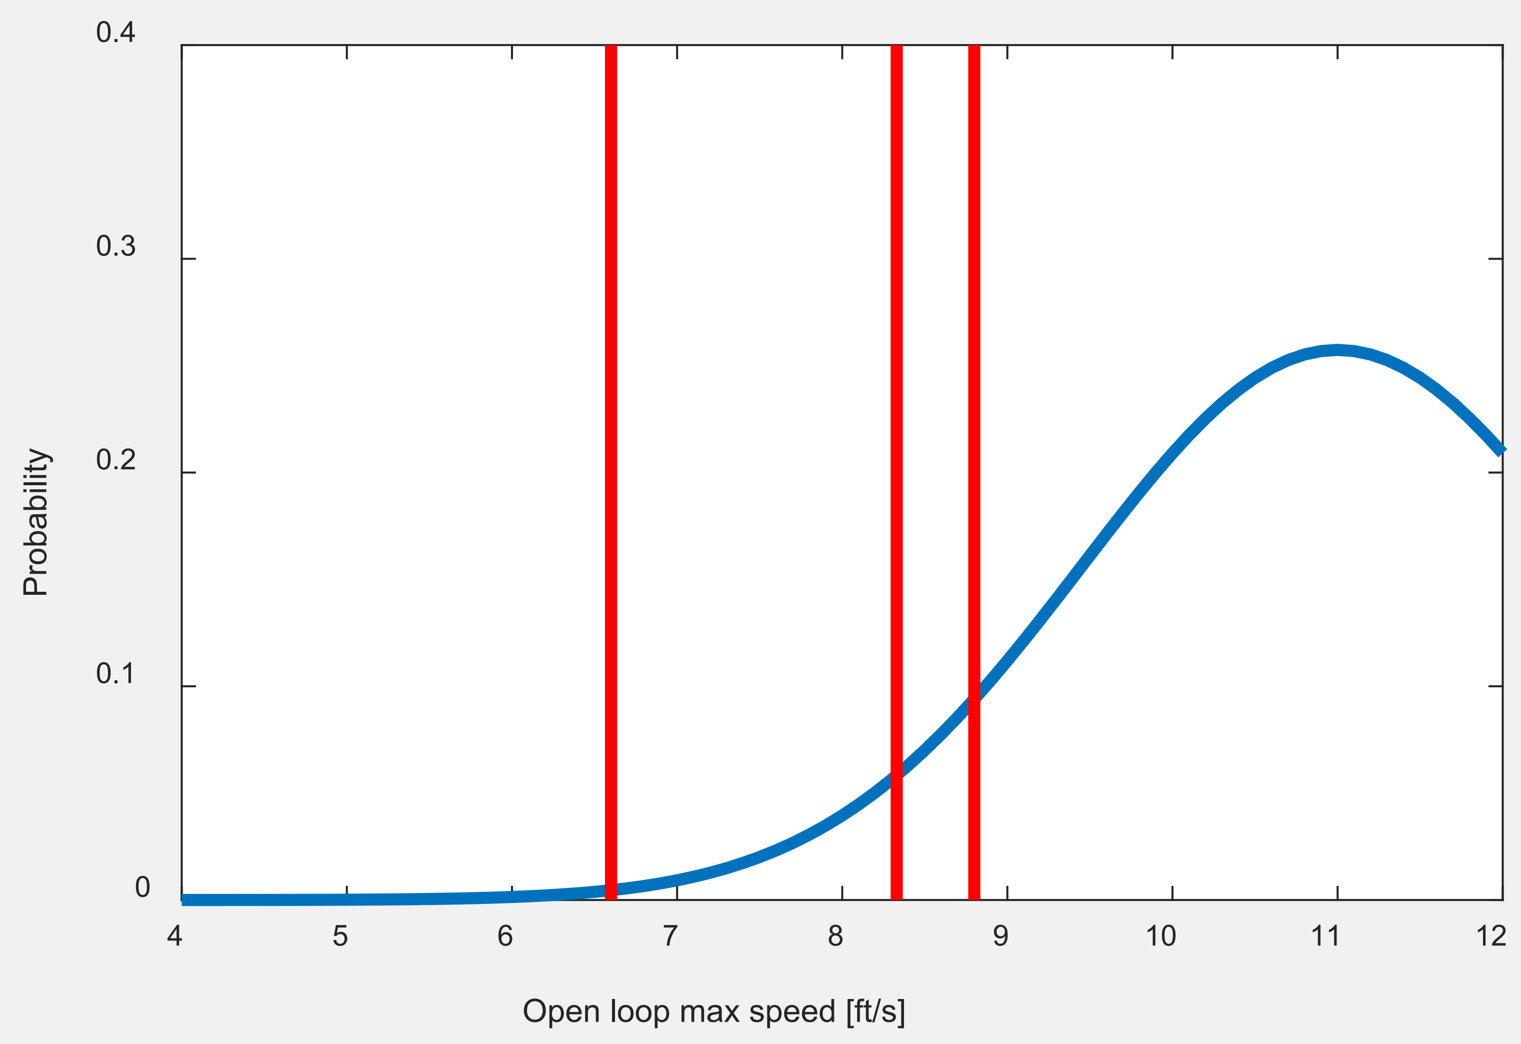

Probability Distributions

$$P(x|z) = \frac{P(z|x)P(x)}{P(z)}$$

- Start by assuming nothing
  - P(s) uniform
  - P(s|m) = P(m|s)* c1/c2
  - Simplified: P(s|m) = P(m|s)
    - *Guess!* What if the actual max speed is 11 ft/s?
    - P(s = 11 | m = [6.6, 8.33, 8.8]) = P(m = [6.6, 8.33, 8.8] | s=11)
    - P(m = 6.6| s = 11)* P(m = 8.33 | s = 11)* P(m = 8.8 | s = 11)

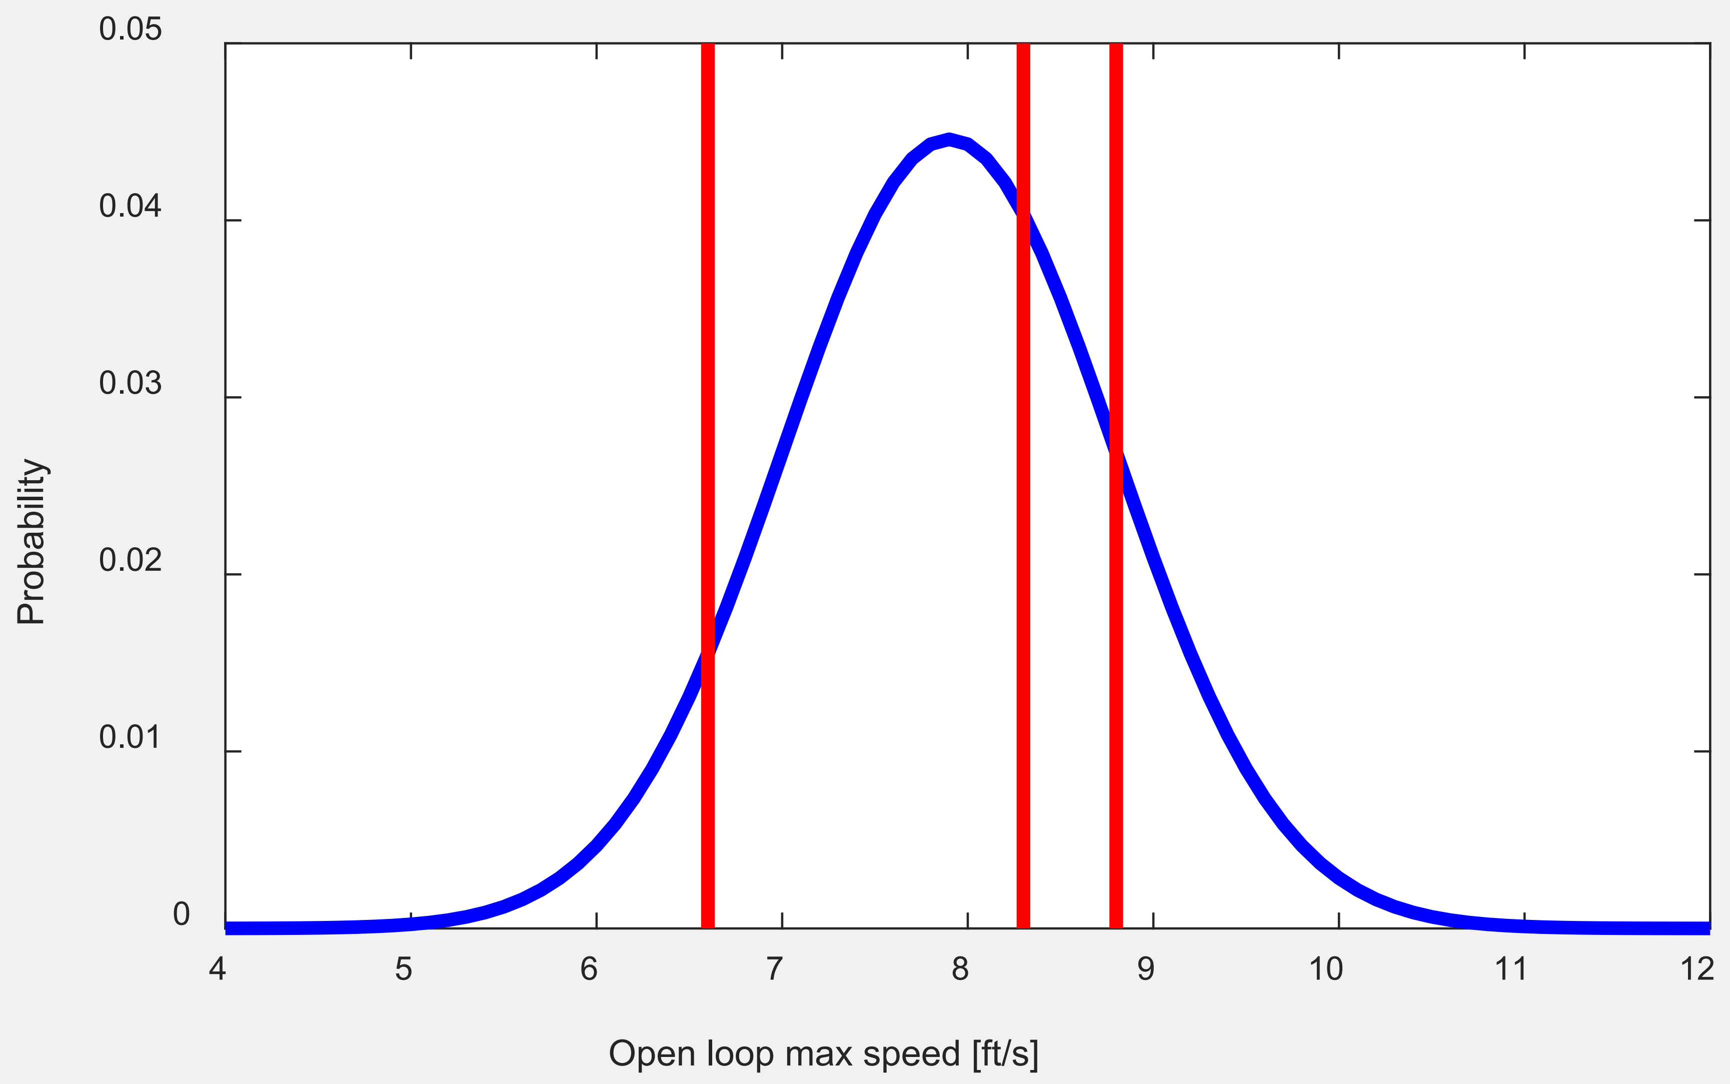

Probability Distributions

$$P(x|z) = \frac{P(z|x)P(x)}{P(z)}$$

- Add a prior!
  - You know yesterday's speed, and you can judge the current speed
    - Prior: 7.91± 1.16 ft/s
  - P(s = 11| m = [6.6, 8.33, 8.8]) = P(m = [6.6, 8.33, 8.8]| s =11)*P(s = 11)

  - Repeat the process!

Probability Distributions

$$P(x|z) = \frac{P(z|x)P(x)}{P(z)}$$

- Add a prior!
  - You know yesterday's speed, and you can judge the current speed
    - Prior: 7.91± 1.16 ft/s
  - P(s = 11| m = [6.6, 8.33, 8.8]) = P(m = [6.6, 8.33, 8.8]| s =11)*P(s = 11)
    = P(m=6.6| s=11)* P(s=11)*P(m=8.33| s=11)* P(s=11)* P(m=8.8| s=11)* P(s=11
  - Repeat the process!
  - Add everything up to get the posterior distribution!

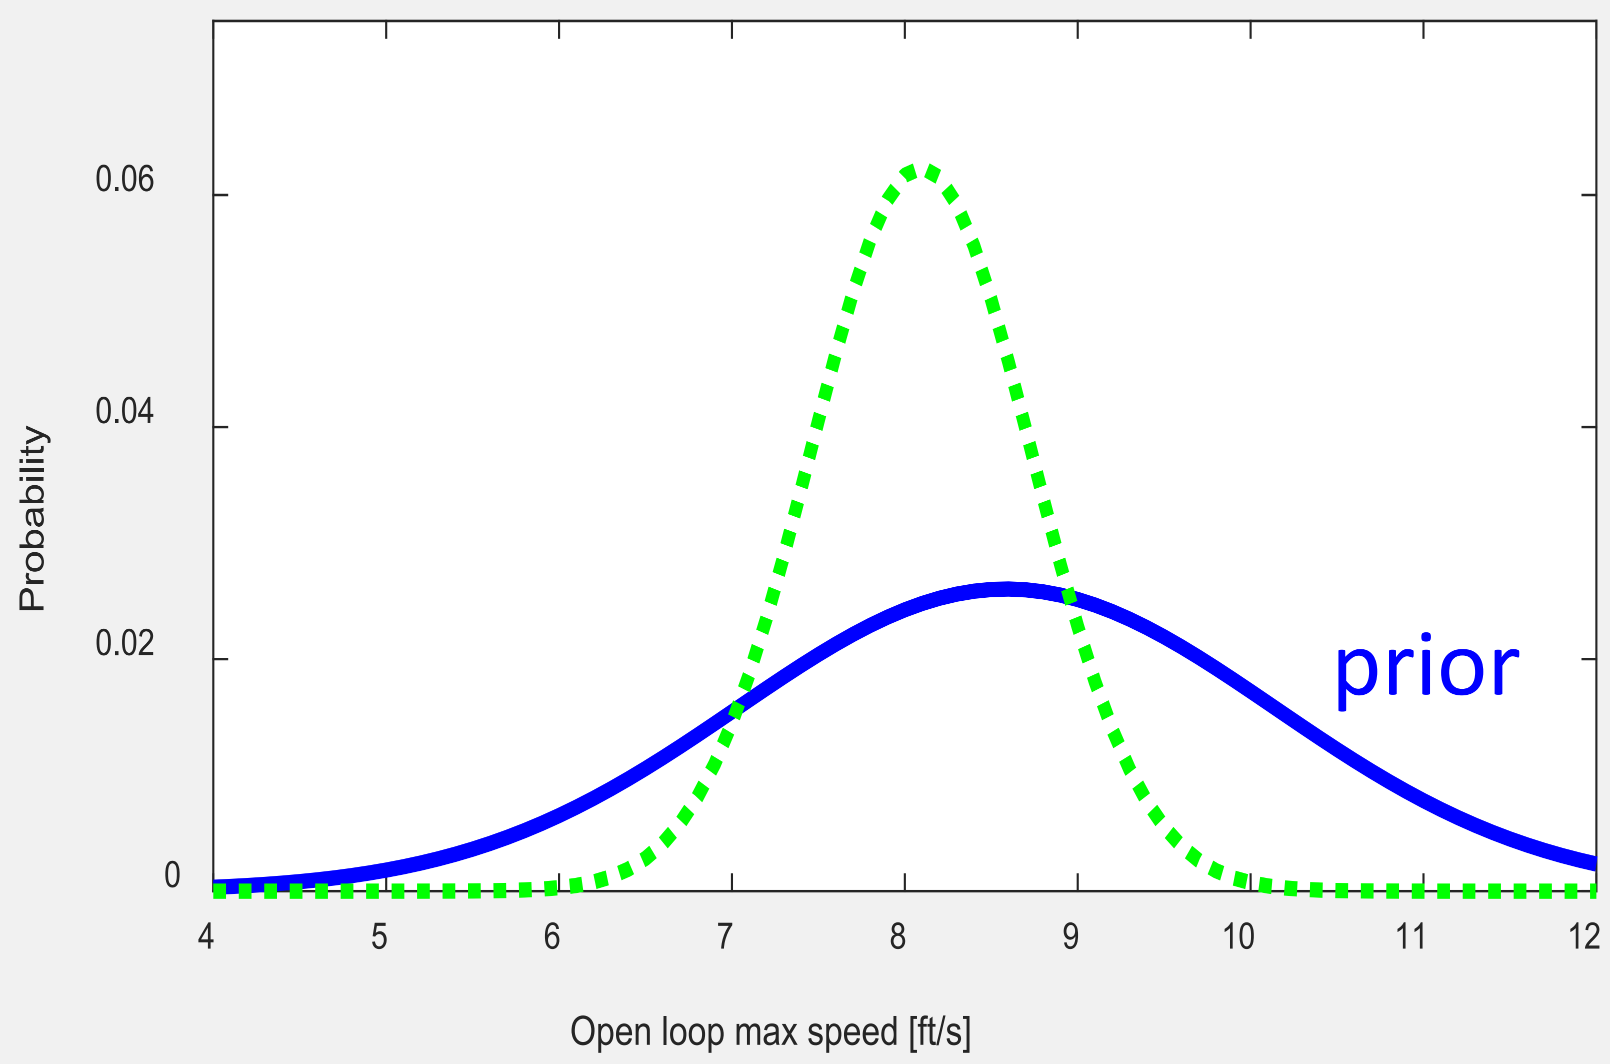

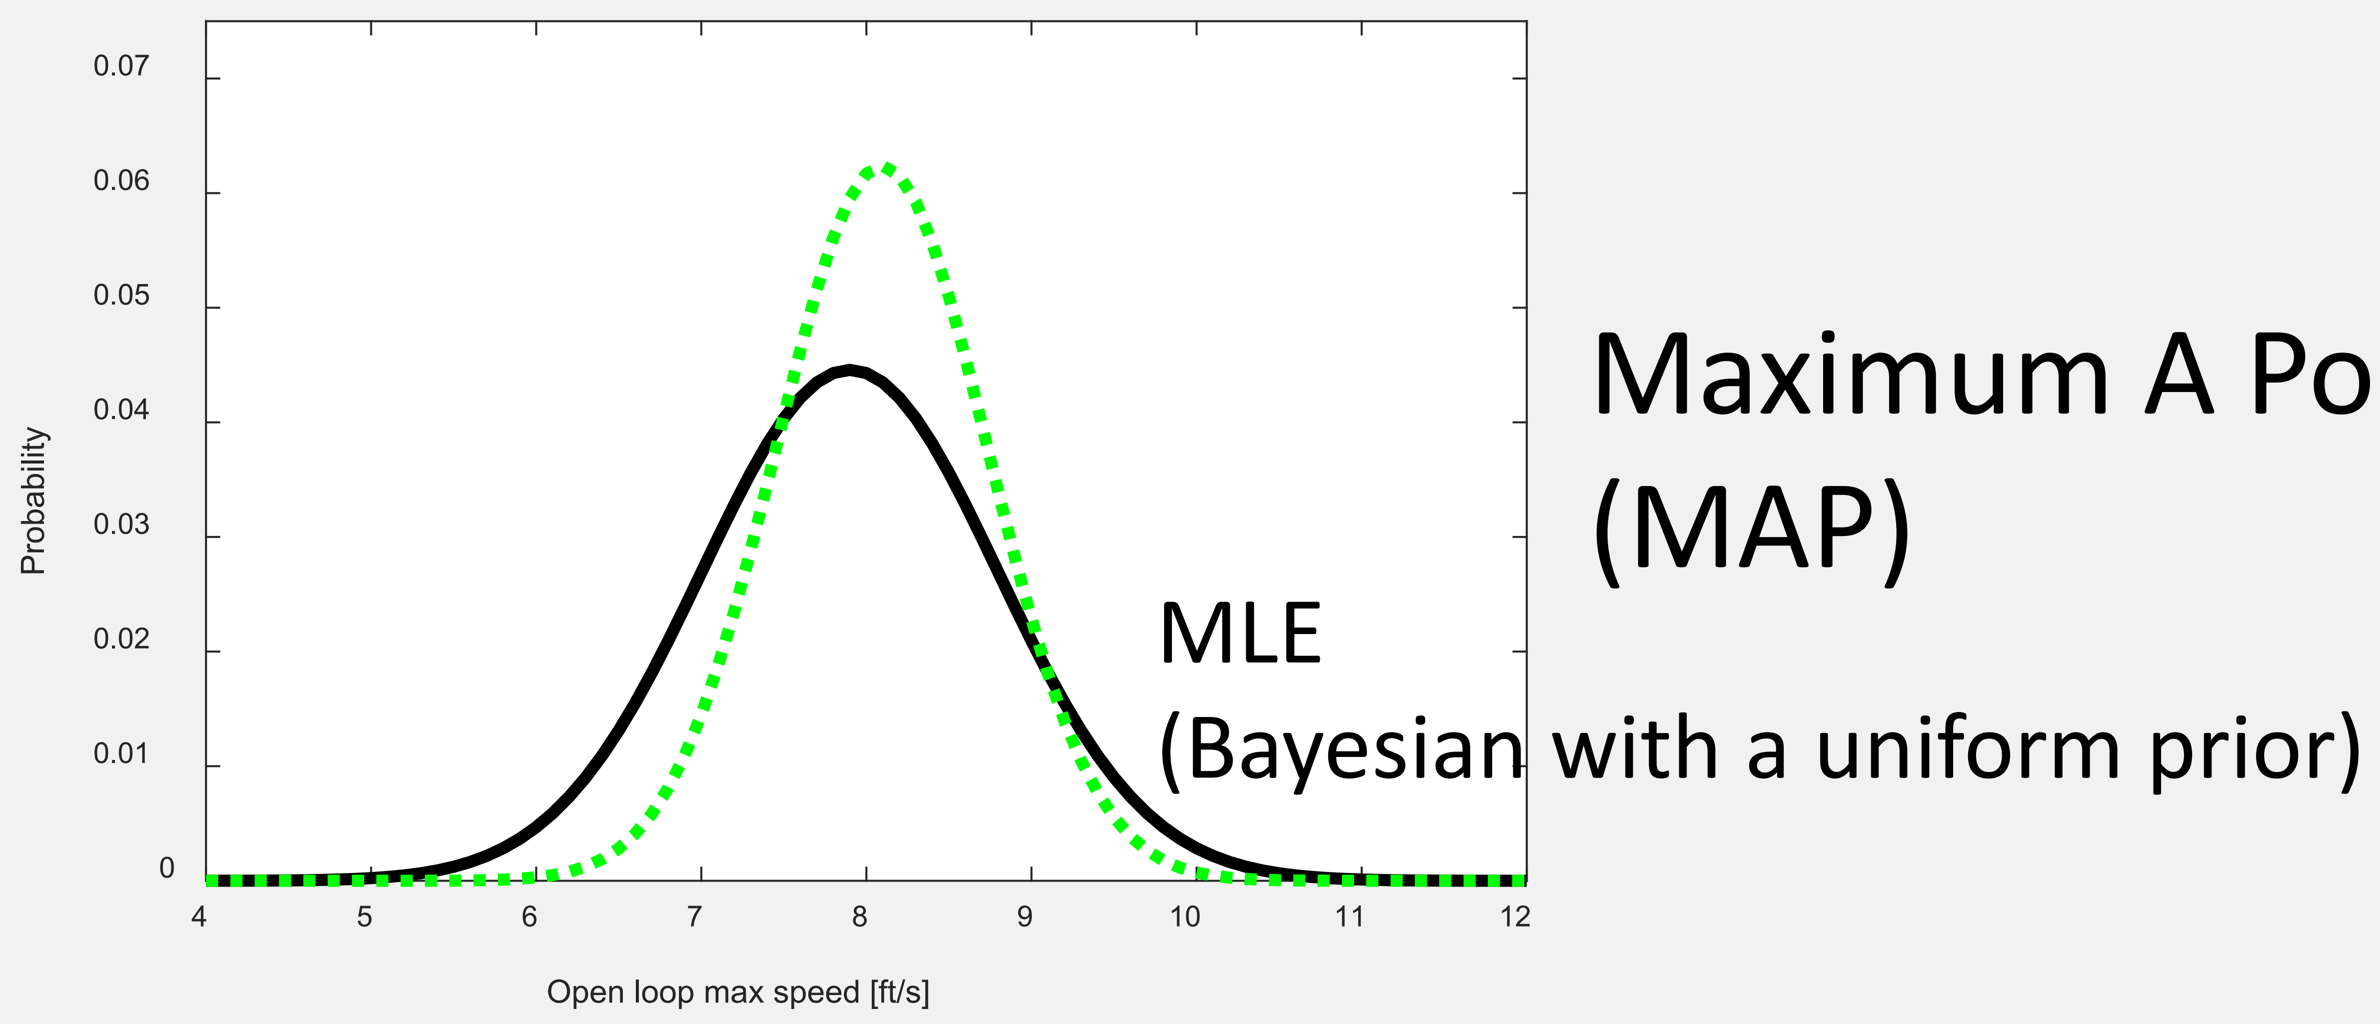

Maximum A Posteriori (MAP)

Bayesian Inference

likelihood prior

$$P(x|z) = \frac{P(z|x)P(x)}{P(z)}$$

z = Sensor data
x = Robot state /location

conditional probability
posterior

marginal likelihood
(constant)

Probabilistic Approach

- ✓ Explicitly represent uncertainty using probabilities
- ✓ Accomodate inaccurate models
- ✓ Accomodate imperfect sensors
- ✓ Robust in real-world applications
- ✓ Best known approach to many hard robotics problems

- \- Computationally demanding
- \- Approximations
- \- False Assumptions

Is Robotics Going Statistics?
The Field of Probabilistic Robotics

Sebastian Thrun
School of Computer Science
Carnegie Mellon University
http://www.cs.cmu.edu/~thrun

draft, please do not circulate

Abstract

In the 1970s, most research in robotics presupposed the availability of exact models, of robots and their environments. Little emphasis was placed on sensing and the intrinsic limitations of modeling complex physical phenomena. This changed in the mid-1980s, when the paradigm shifted towards reactive techniques. Reactive controllers rely on capable sensors to generate robot control. Rejections of models were typical for researchers in this field. Since the mid-1990s, a new approach has begun to emerge: probabilistic robotics. This approach relies on statistical techniques to seamlessly integrate imperfect models and imperfect sensing. The present article describes the basics of probabilistic robotics and highlights some of its recent successes.

Probability Distributions

- Always believe the impossible, at least a little bit!
- Leave room for believing the unlikely. Leave a non-zero probability unless you are **absolutely** certain.
- “It ain’t what you don’t know that gets you into trouble. It’s what you know for sure that just ain’t so.” — Mark Twain
- “Alice laughed ‘there’s no use trying,’ she said, ‘one can’t believe impossible things.’ ‘I daresay you haven’t had much practice’ said the Queen. ‘When I was younger, I always did it for half an hour a day. Why sometimes, I’ve believed as many as six impossible things before breakfast.’”

Alice’s Adventures in Wonderland

References

- "Probabilistic Robotics," Dieter Fox, Sebastian Thrun, and Wolfram Burgard
- How Bayes Theorem works (YouTube), Brandon Rohrer