Prof. Kirstin Hagelskjær Petersen
kirstin@cornell.edu

Fast Robots

Probability and Bayes Theorem

Recap from ECE 3100 Intro to Probability and Inference

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

- Conditional probability
  - $p(x|y) = \frac{p(x,y)}{p(y)}$

Conditional probability

- $p(x|y) = \frac{p(x,y)}{p(y)}$

- Robot/sensor example

- Exercise
  - You meet a guy, and he says he has a sibling, what is the probability that the sibling is female?
    - guy/girl
    - guy/guy
    - girl/guy
    - girl/girl (<ruled out)
    - 33%
  - Independent variables

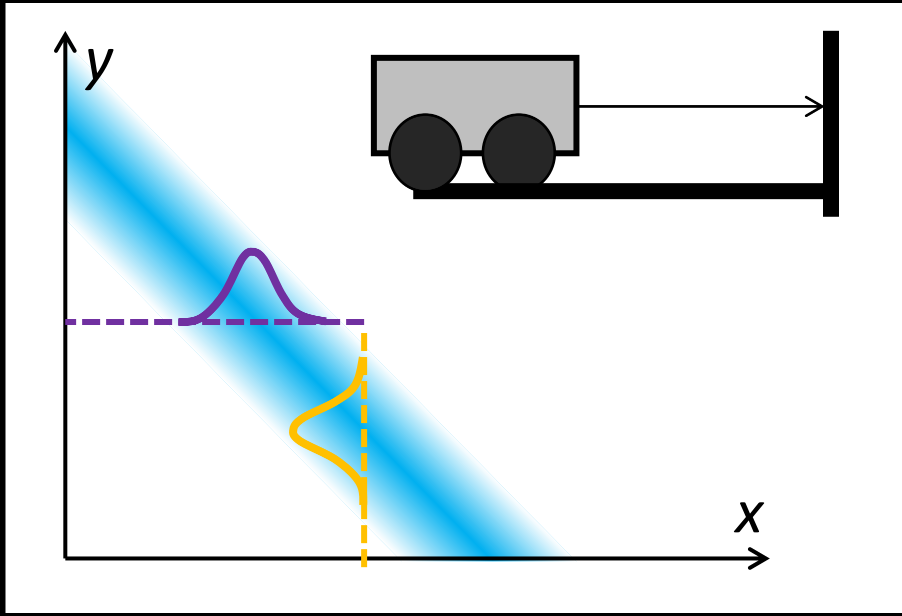

Recap from ECE 3100 Intro to Probability and Inference

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

Why bother considering uncertainty?

- Uncertainty is inherent in the world
- Five major factors
  - Unpredictable environments
  - Sensors
    - Subject to physical laws
    - Signal to noise ratio
  - Robot motion
    - Noise, wear and tear, battery state, etc.
    - Accuracy versus cost
  - Models
    - Abstractions of the real world
  - Computation
    - Real time systems
    - Timely response versus accuracy

Exercise

- Is this dress black and royal blue, or white and gold?
- Where does the uncertainty come from?
  - blue and black under a yellow-tinted illumination (left)
  - white and gold under a blue-tinted illumination (right)

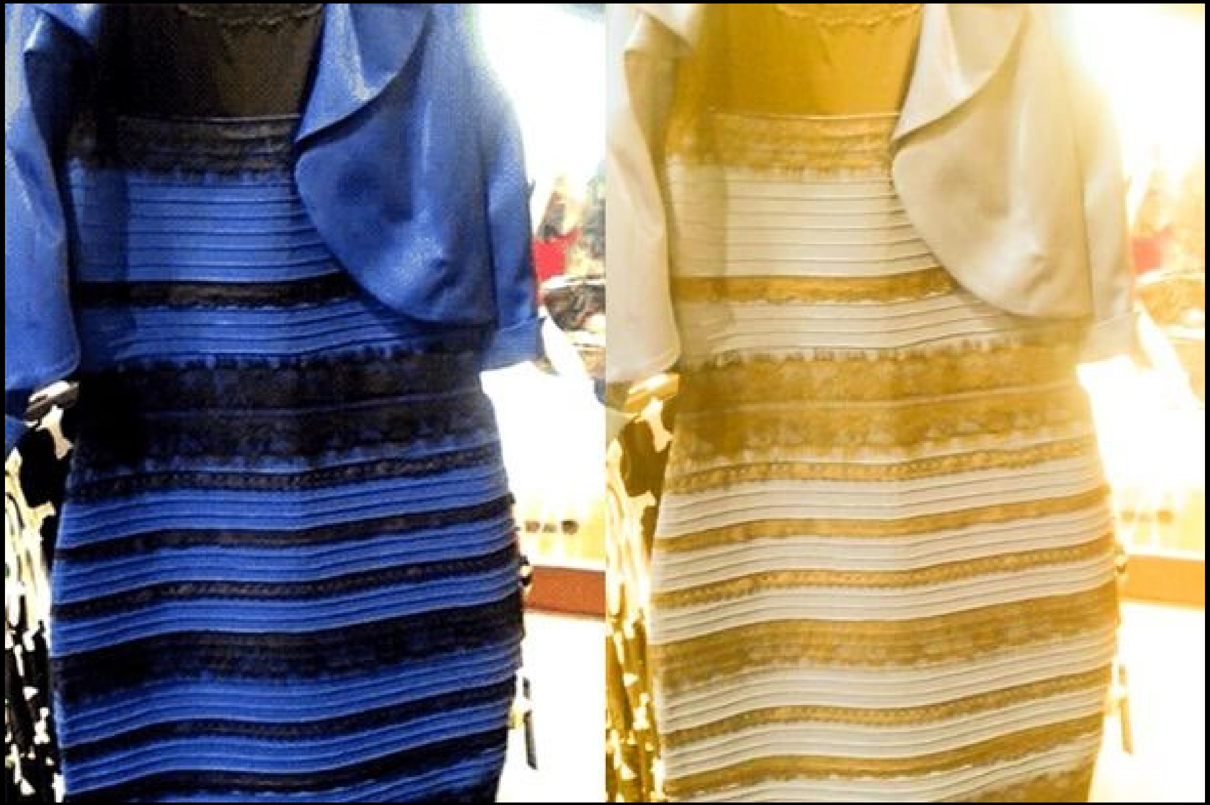

Probabilistic Approach

"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

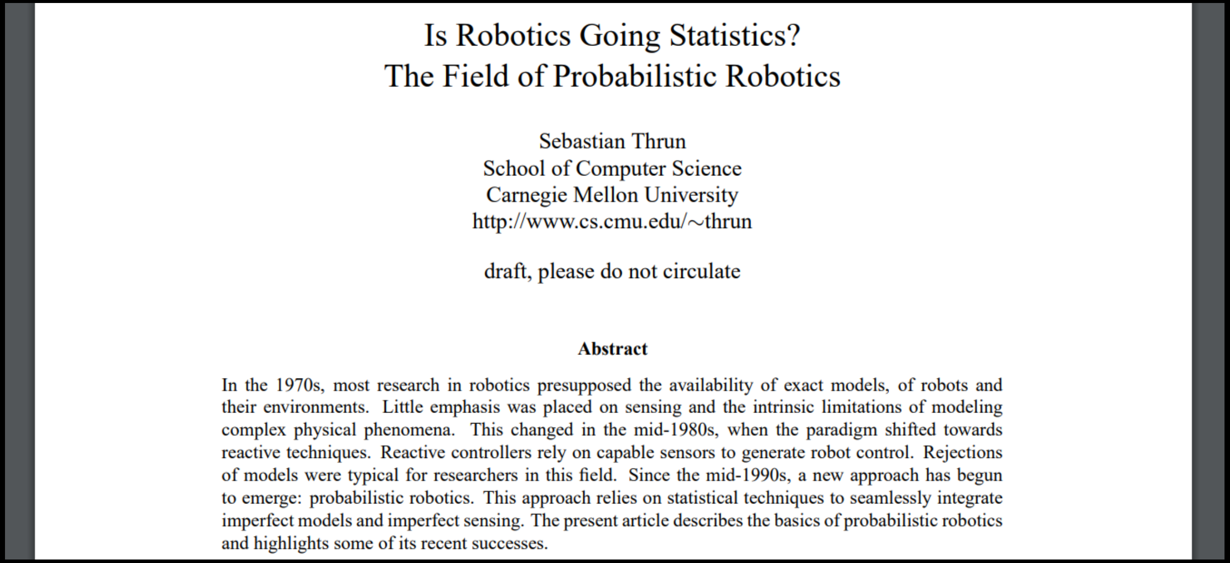

Is Robotics Going Statistics?
The Field of Probabilistic Robotics

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 the uncertainty using probability theory
+ Can accommodate inaccurate models
+ Can accommodate imperfect sensors
+ Robust in real-world applications
+ Best known approach to many hard robotics problems

- Computationally demanding
- Need to approximate
- 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
  - *EdDiscussion:* "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?

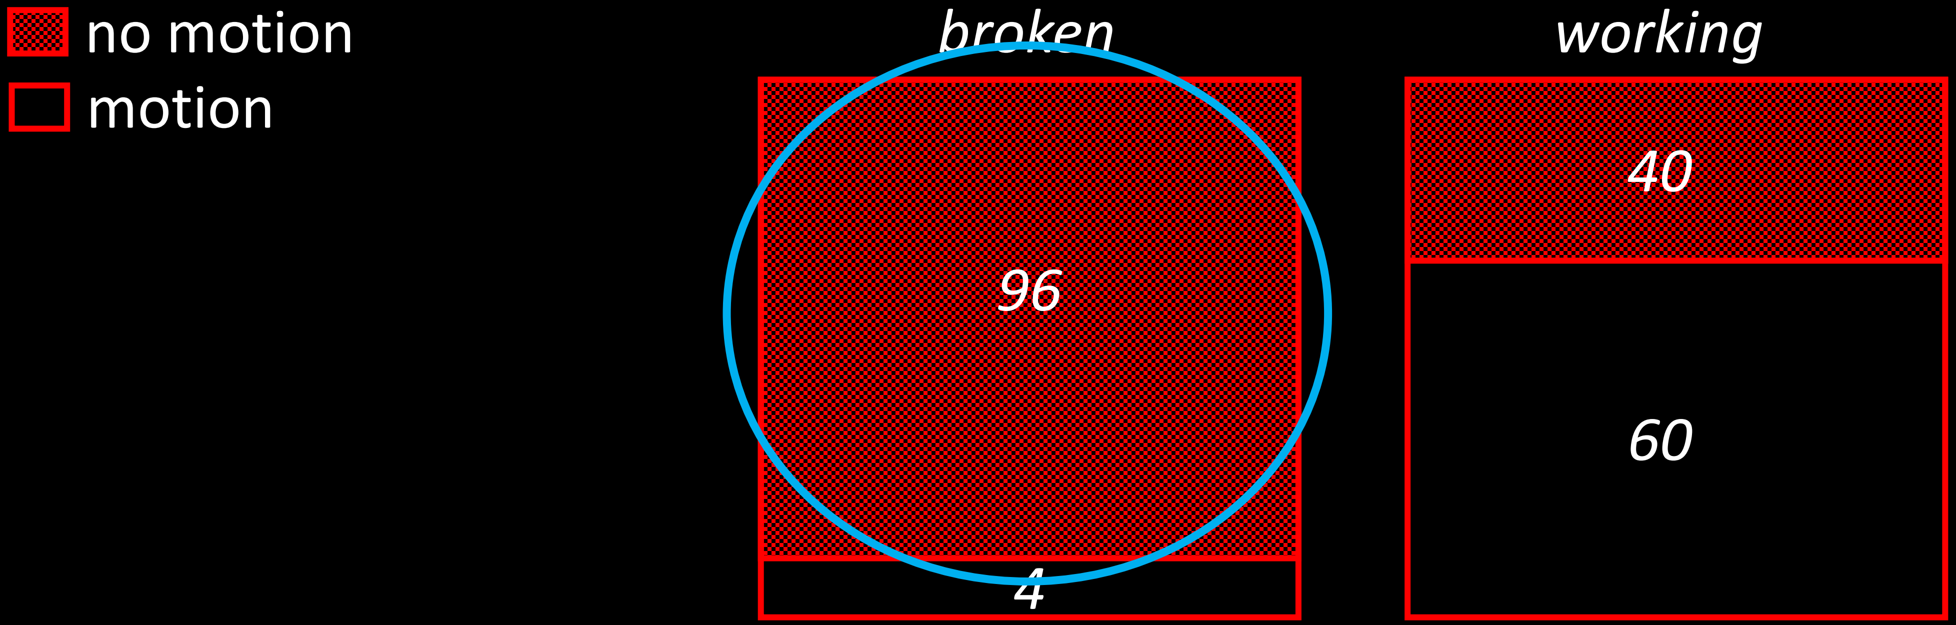

Bayesian Inference

- Inference = educated guessing
- Bayesian inference = guessing in the style of Bayes
- Example
  - *EdDiscussion:* "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?

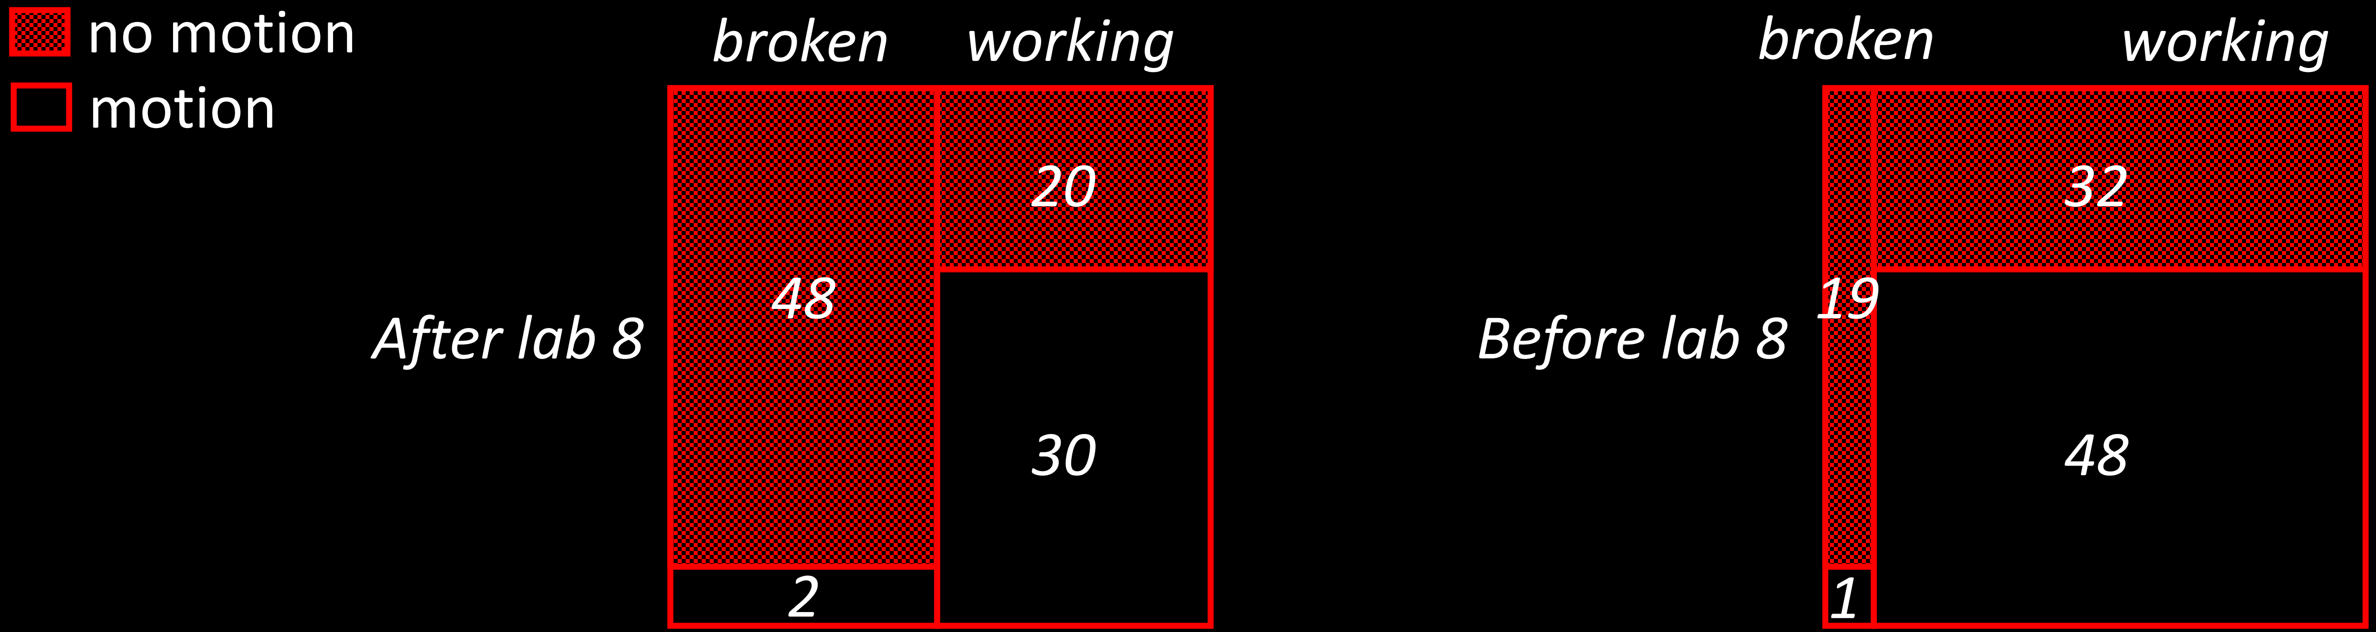

Bayesian Inference

- Inference = educated guessing
- Bayesian inference = guessing in the style of Bayes
- Example
  - *EdDiscussion:* "My robot stopped moving, the hardware is broken, send me new parts"
- **Translate to math**
  - P(something) = #something / #everything
  - Before lab 8:
    - P(broken) = #broken / #kits = 20 / 100 = 0.2
    - P(working) = #working / #kits = 80 / 100 = 0.8
  - After lab 8:
    - P(broken) = #broken / #kits = 50 / 100 = 0.5
    - P(working) = #working / #kits = 50 / 100 = 0.5

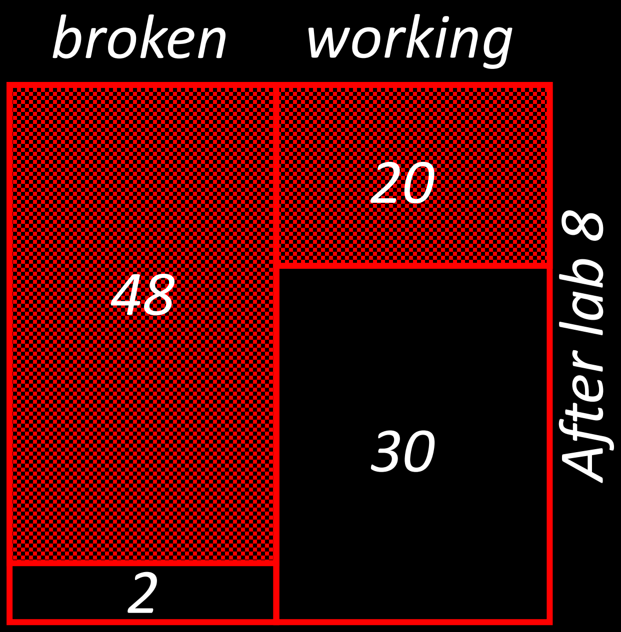

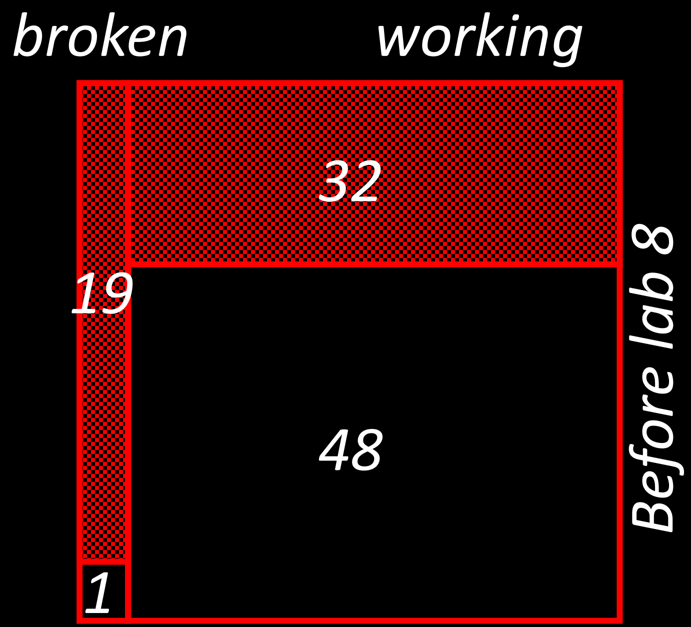

Bayesian Inference

- Inference = educated guessing
- Bayesian inference = guessing in the style of Bayes
- Example
  - *EdDiscussion:* "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?

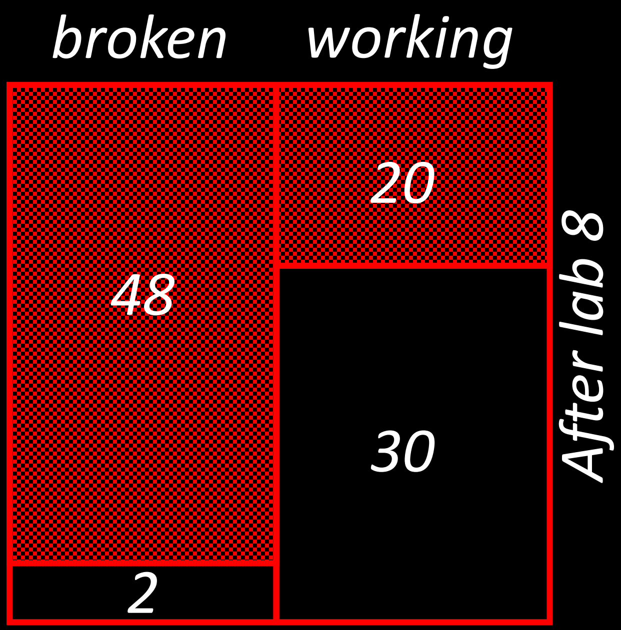

- **Conditional Probability**
  - If you know that the robot is broken, what is the probability that it stopped moving?
  - P(no motion | broken) = #broken and no motion / #broken
  - After lab 8 = 48/50 = 0.96
  - P(no motion | working) = #working and no motion / #working
  - After lab 8 = 20/50 = 0.40

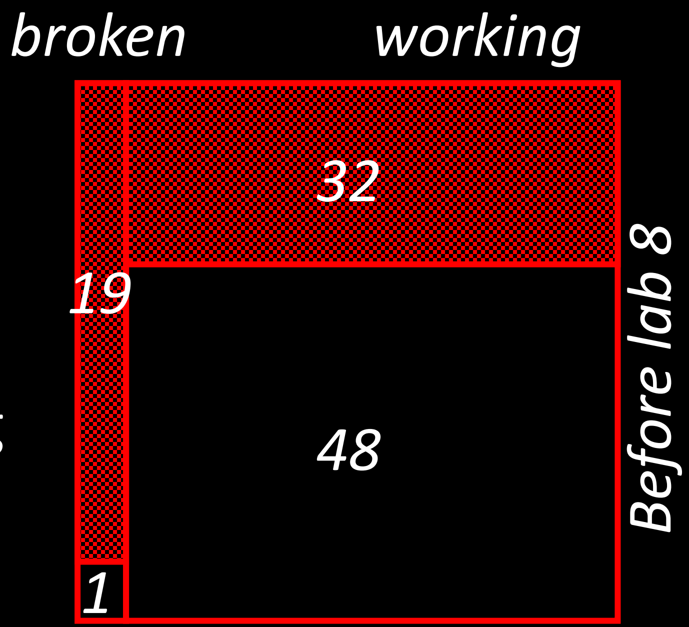

Bayesian Inference

- Inference = educated guessing
- Bayesian inference = guessing in the style of Bayes
- Example
  - *EdDiscussion:* "My robot stopped moving, the hardware is broken, send me new parts"

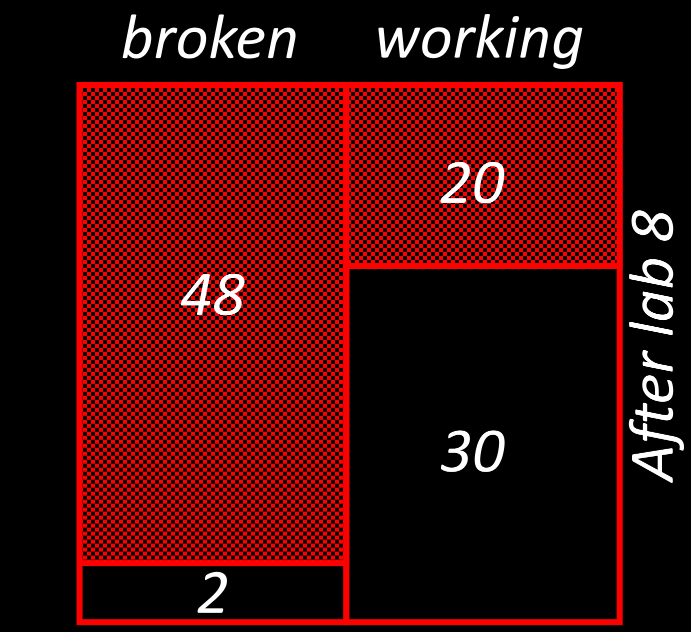

- **Conditional Probability**
  - If you know that the robot is broken, what is the probability that it stopped moving?
  - P(no motion | broken) = #broken and no motion / #broken
  - Before lab 8 = 19/20 = 0.96
  - P(no motion | working) = #working and no motion / #working
  - Before lab 8 = 32/80 = 0.40

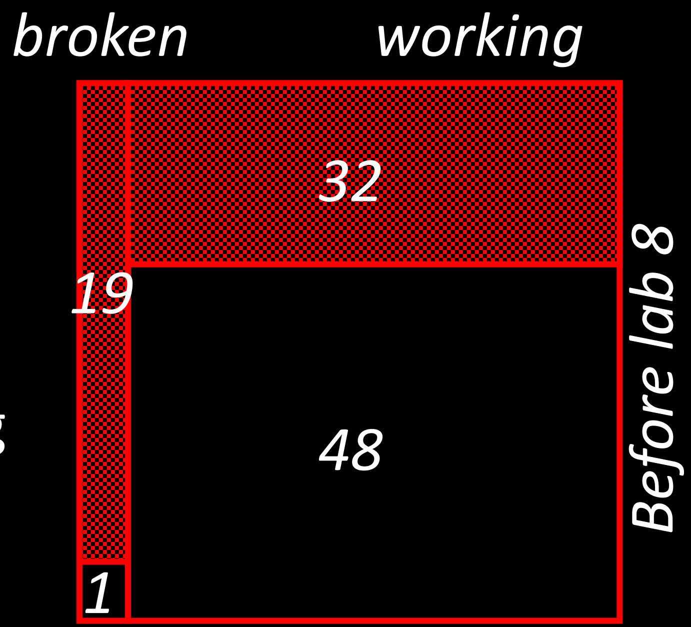

Bayesian Inference

- Inference = educated guessing
- Bayesian inference = guessing in the style of Bayes
- Example
  - *EdDiscussion:* "My robot stopped moving, the hardware is broken, send me new parts"

- **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)

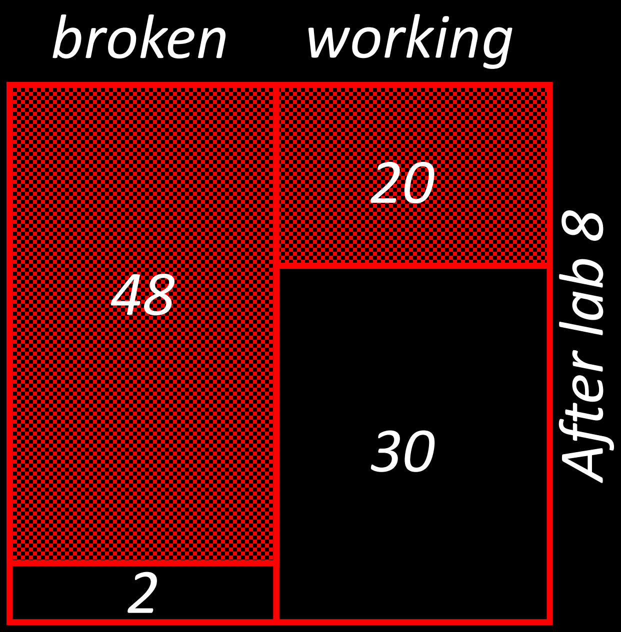

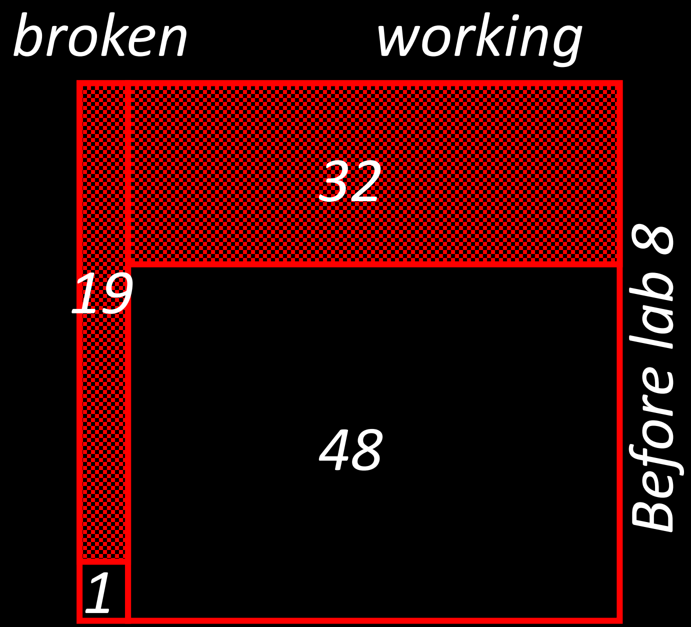

Bayesian Inference

- Inference = educated guessing
- Bayesian inference = guessing in the style of Bayes
- Example
  - *EdDiscussion:* "My robot stopped moving, the hardware is broken, send me new parts"

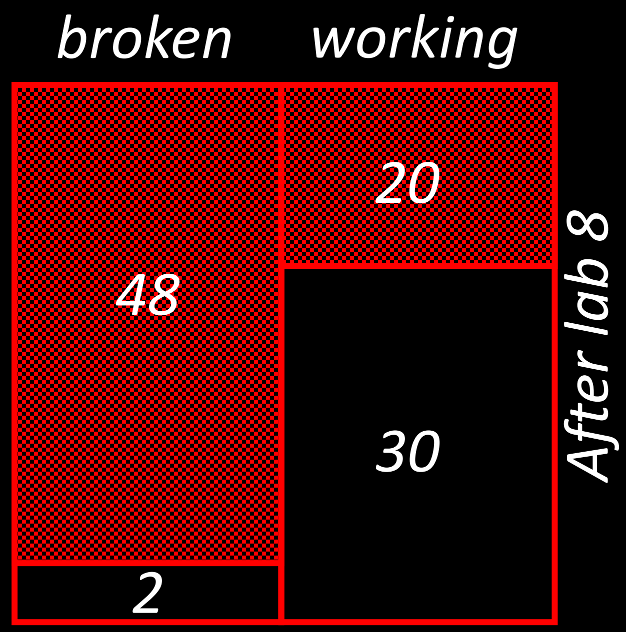

- **Joint Probability**
  - What is the probability that the robot is both broken and not moving?
  - P(broken and not moving)
    = P(broken)*P(not moving | broken)
    = 0.5 * 0.96 = 0.48

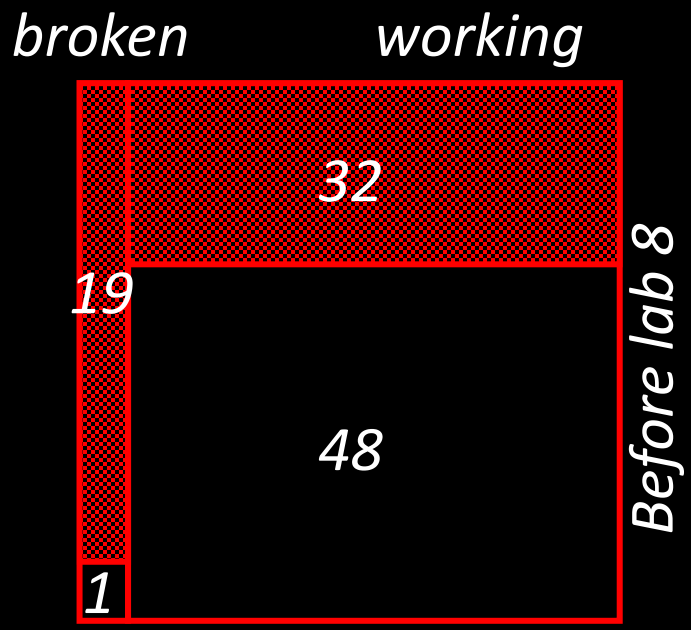

Bayesian Inference

- Inference = educated guessing
- Bayesian inference = guessing in the style of Bayes
- Example
  - *EdDiscussion:* "My robot stopped moving, the hardware is broken, send me new parts"

- **Joint Probability**
  - What is the probability that the robot is both broken and not moving?
  - P(broken and not moving)
    = P(broken)*P(not moving | broken)
    = 0.20 * 0.96 = 0.192
  - P(working and moving)
    = P(working)*P(moving | working)
    = 0.80 * 0.60 = 0.48

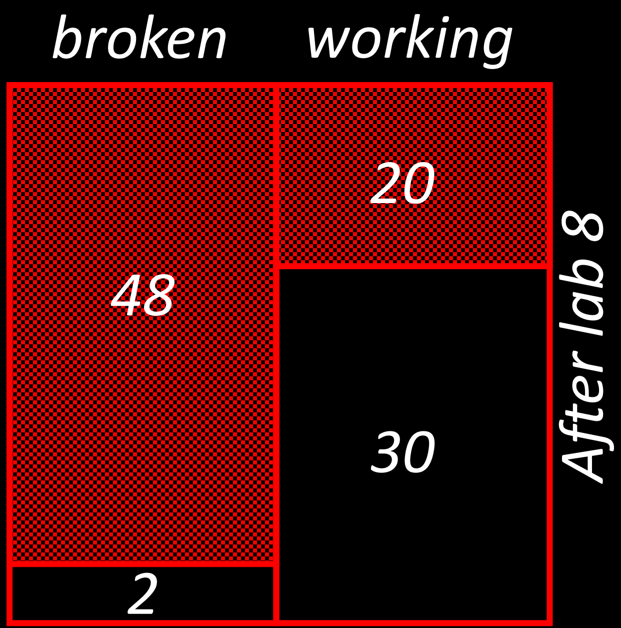

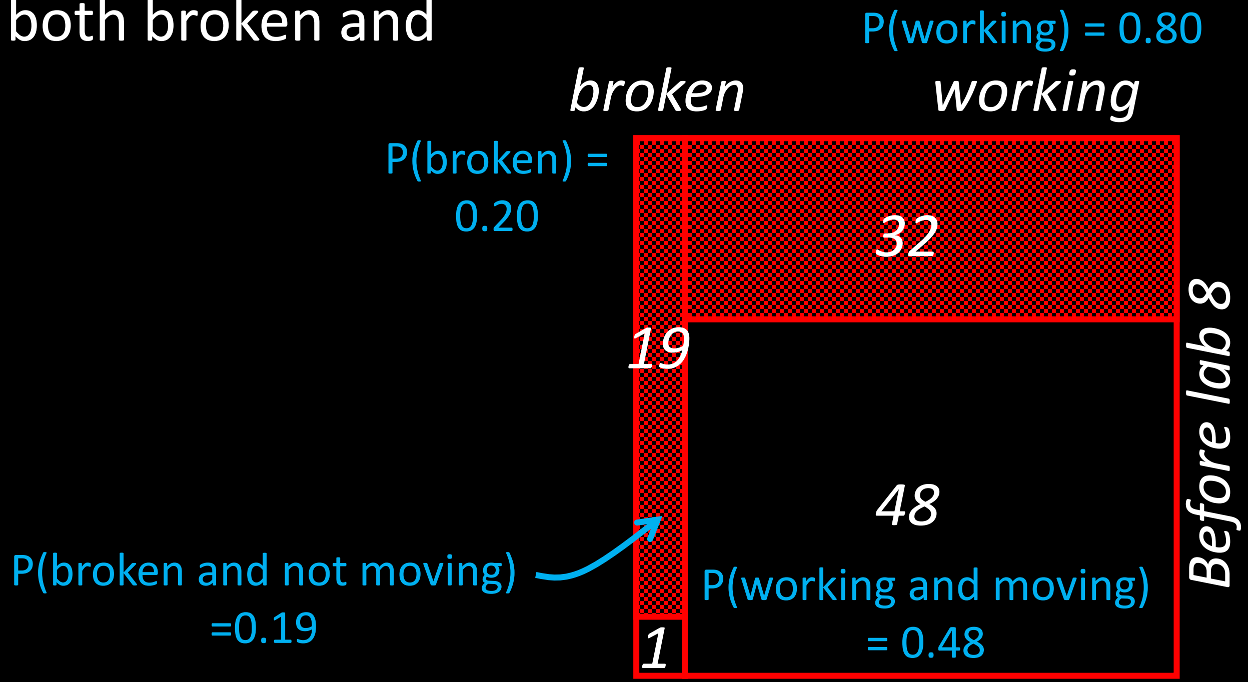

Bayesian Inference

- Inference = educated guessing
- Bayesian inference = guessing in the style of Bayes
- Example
  - *EdDiscussion:* "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?
- **Joint Probability**
  - What is the probability that the robot is both 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)

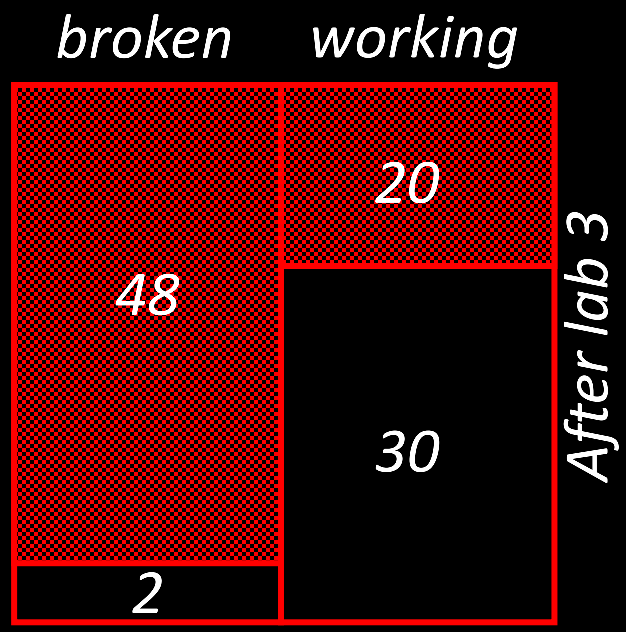

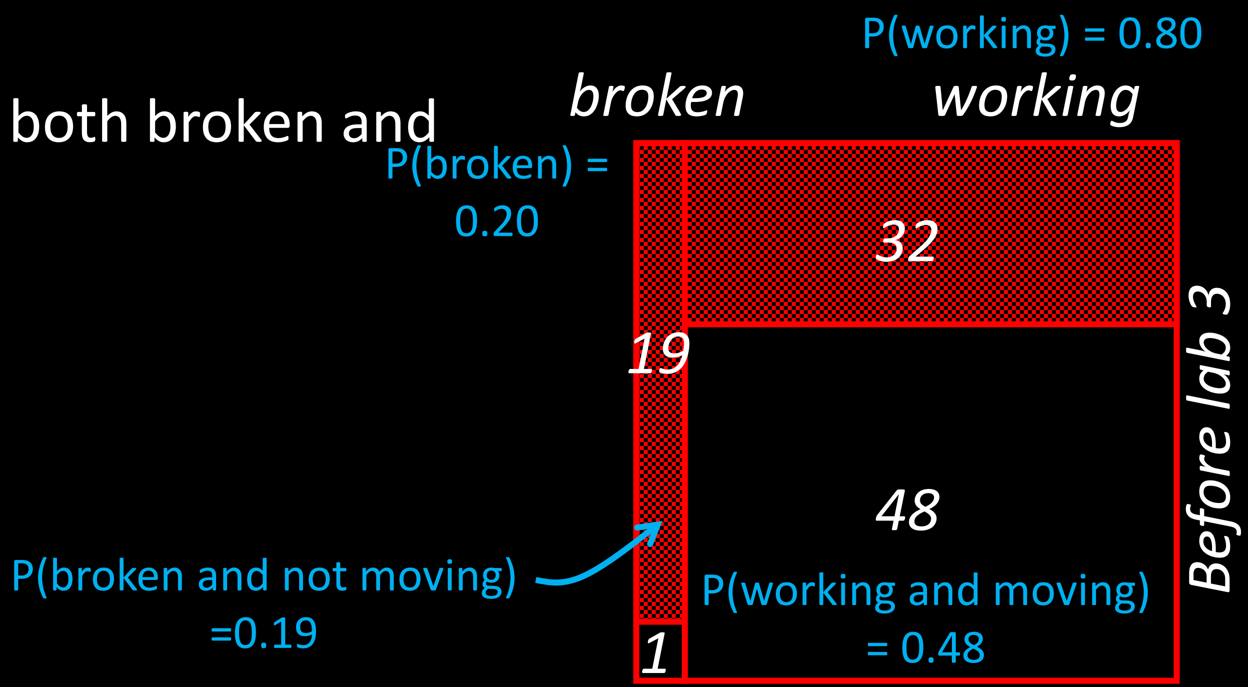

Bayesian Inference

- Inference = educated guessing
- Bayesian inference = guessing in the style of Bayes
- Example
  - *EdDiscussion:* "My robot stopped moving, the hardware is broken, send me new parts"

- **Marginal Probability**
  - P(moving)
    = P(broken and moving) + P(working and moving)
    = 1/100 + 48/100 = 0.49
  - P(not moving)
    = 19/100 + 32/100 = 0.51

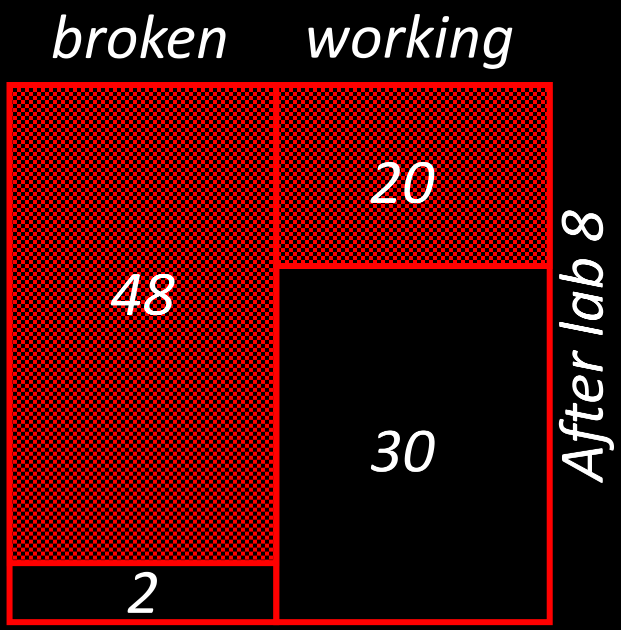

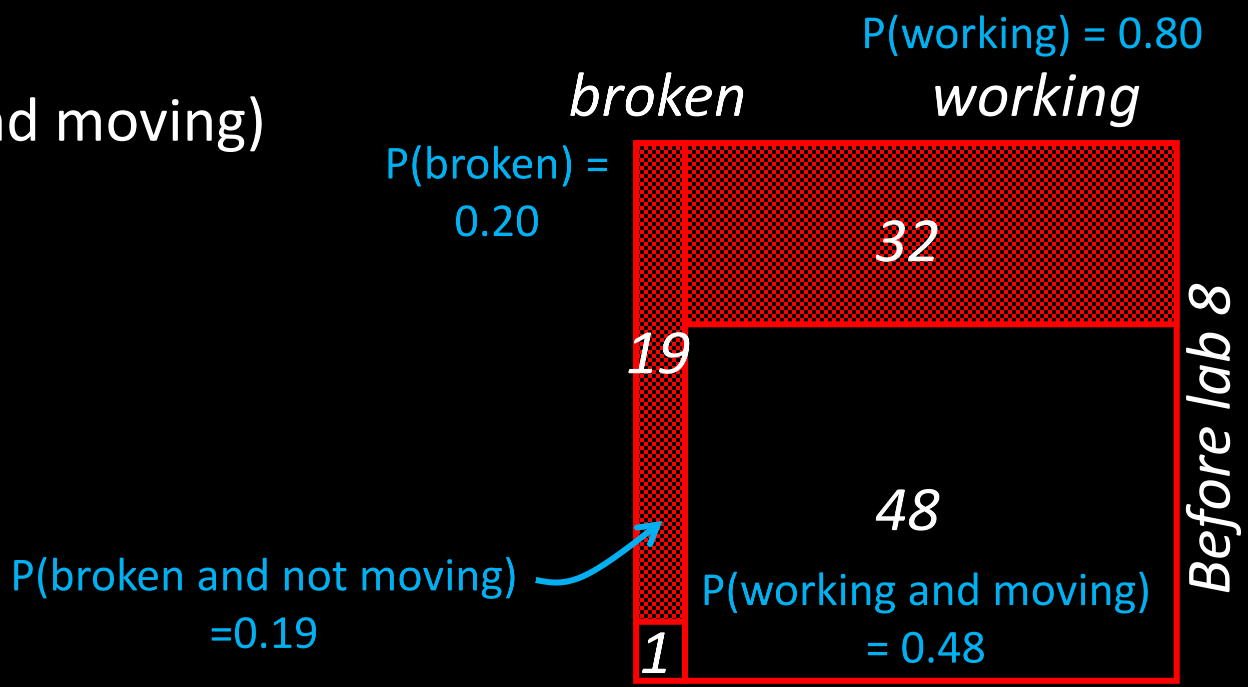

Bayesian Inference

- Inference = educated guessing
- Bayesian inference = guessing in the style of Bayes
- Example
  - *EdDiscussion:* “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 | not moving) = ???
- P(broken and not moving)
  = P(not moving)*P(broken|not moving)
- P(not moving and broken)
  = P(broken)*P(not moving|broken)
- P(broken|not moving) = $\frac{\text{P(broken)*P(not moving|broken)}}{\text{P(not moving)}}$
- Before lab 8 = 0.2*0.96 / 0.51 = 0.38
- After lab 8 = 0.5*0.96 / 0.68 = 0.71

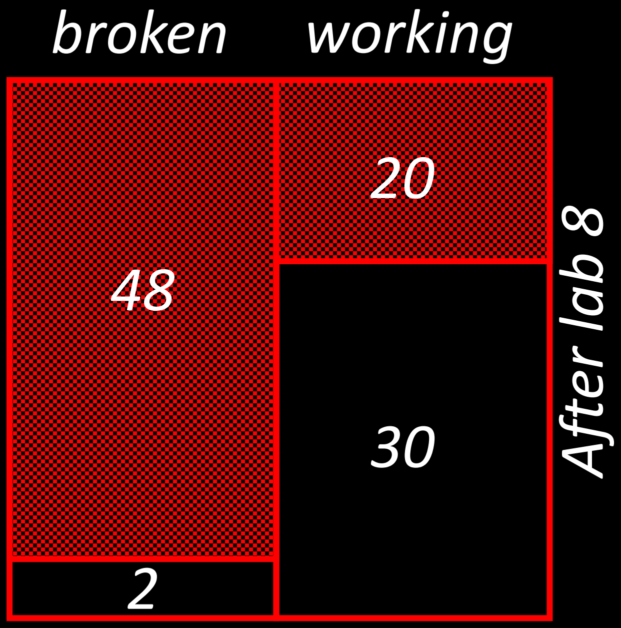

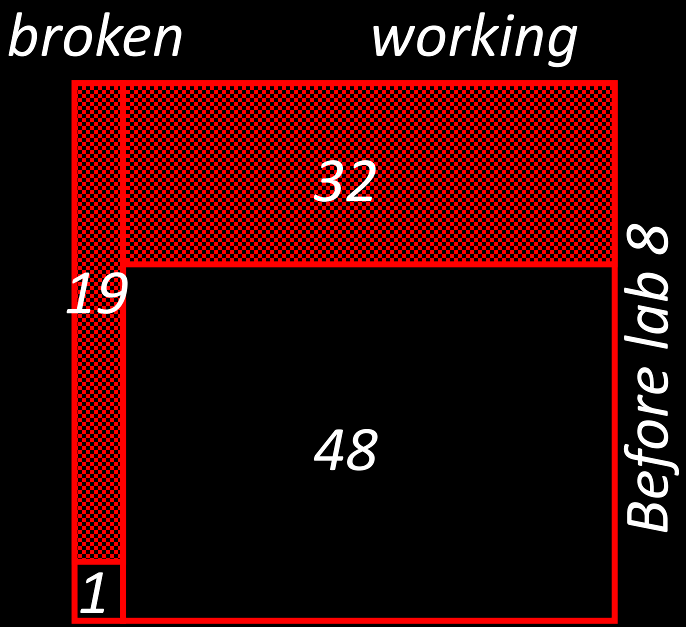

Bayesian Inference

- Bayesian inference = guessing in the style of Bayes

likelihood

prior

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

conditional probability

posterior

marginal likelihood

(constant)

Bayesian Inference

likelihood *prior*

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

conditional probability
posterior

marginal likelihood
(constant)

- y = Sensor data
- x = Robot state/ location

Bayesian Inference

- Lost robot example
  - Initially p(x) is the same for all states
  - If y=1, where are you most likely to be?
  - p(x|y) can be hard to compute
  - What is p(y|x)?

likelihood — prior

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

conditional probability
posterior

marginal likelihood
(constant)

- y = Sensor data
- x = Robot state/ location

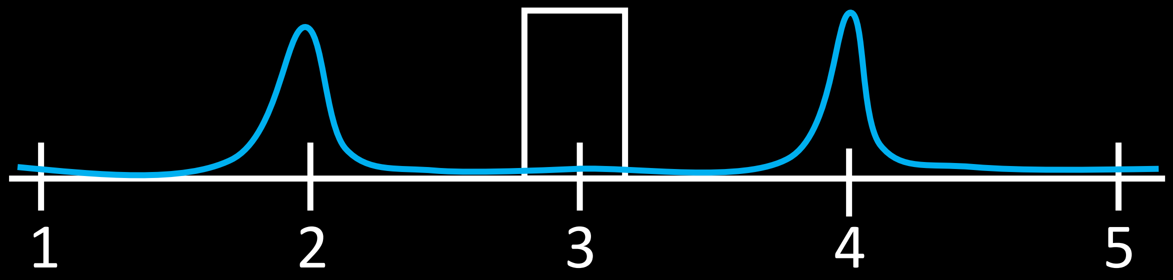

Bayesian Inference

likelihood　　prior

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

conditional probability
posterior

marginal likelihood
(constant)

$$p(x|y) = \frac{p(y|x)p(x)}{\sum_{x'} p(y|x')p(x')}$$

$$p(x|y) = \eta\, p(y|x)p(x)$$

Probability Distribution

- Beliefs

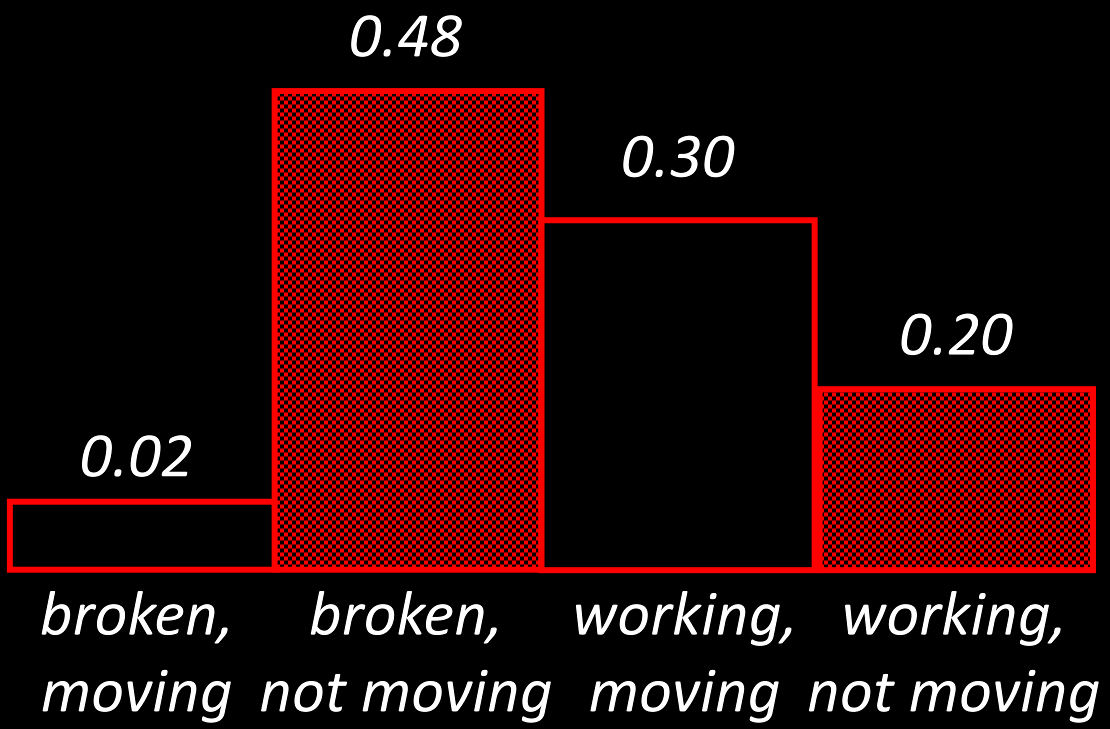

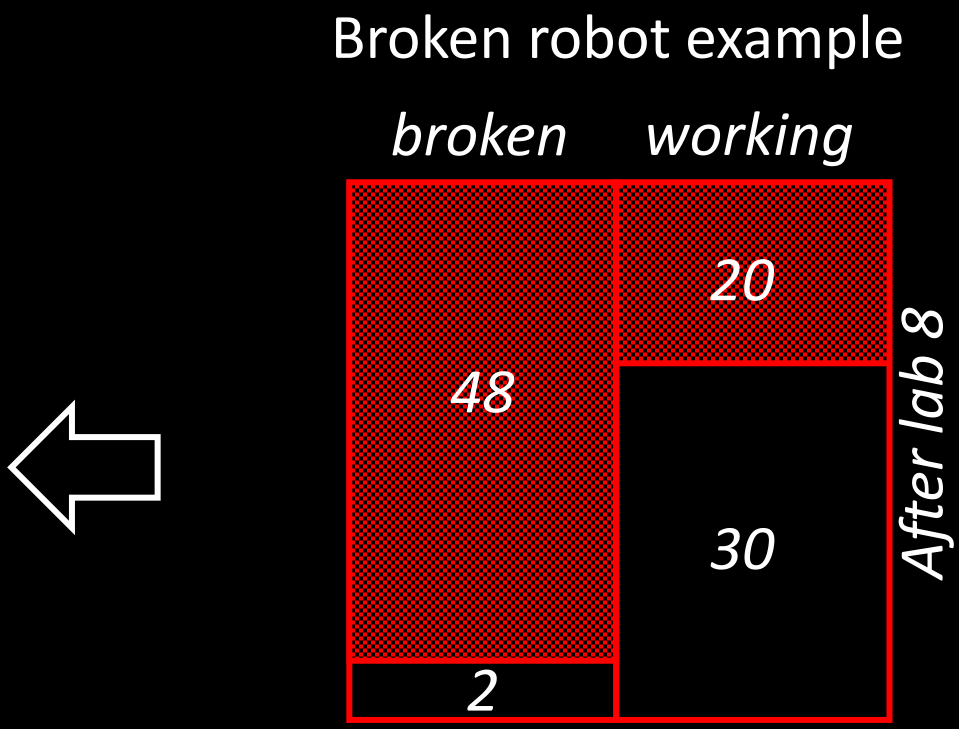

Probability Distribution

- Beliefs

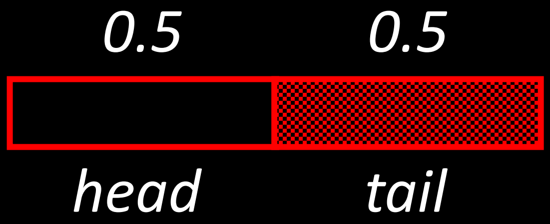

Probability Distribution

- Beliefs

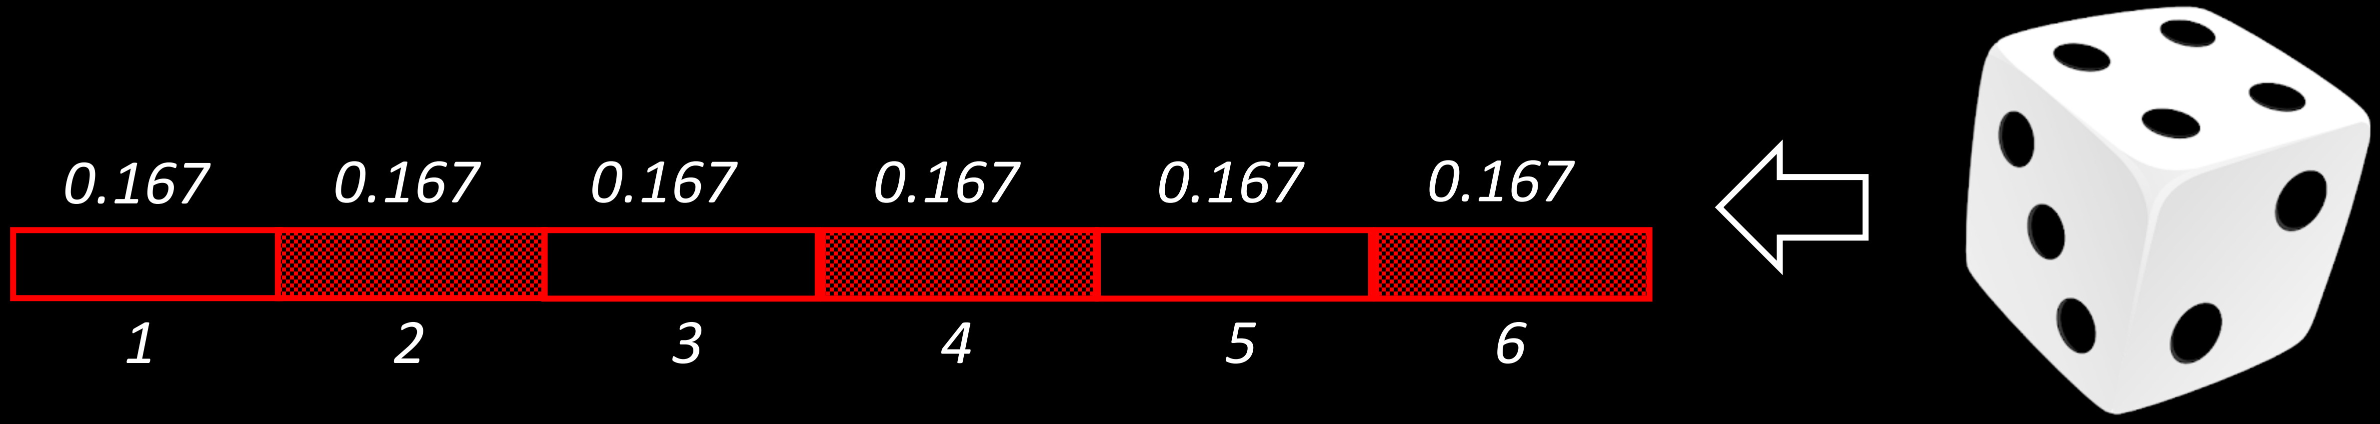

Probability Distribution

- Beliefs

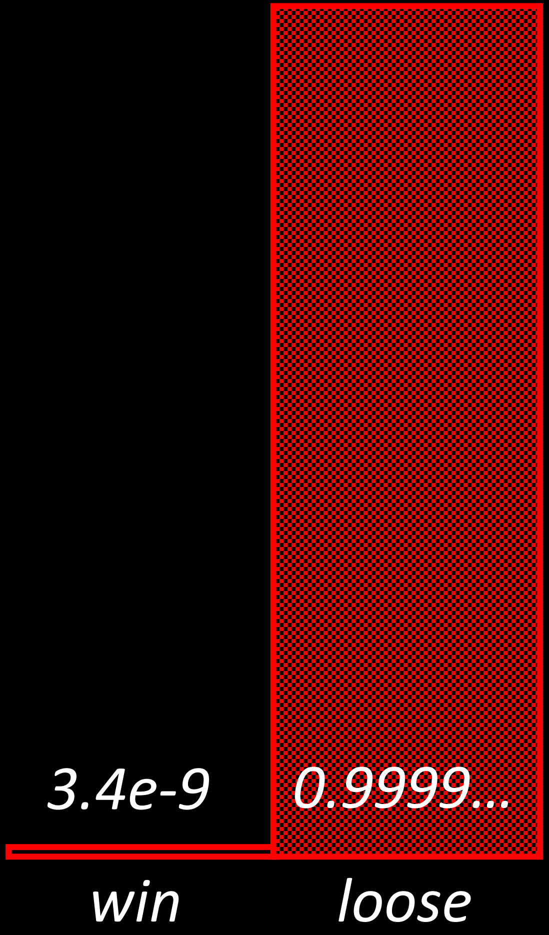

Probability Distribution

- Beliefs
- Discrete -> continuous *probability distribution*
  - Mean, median, most common value, etc.

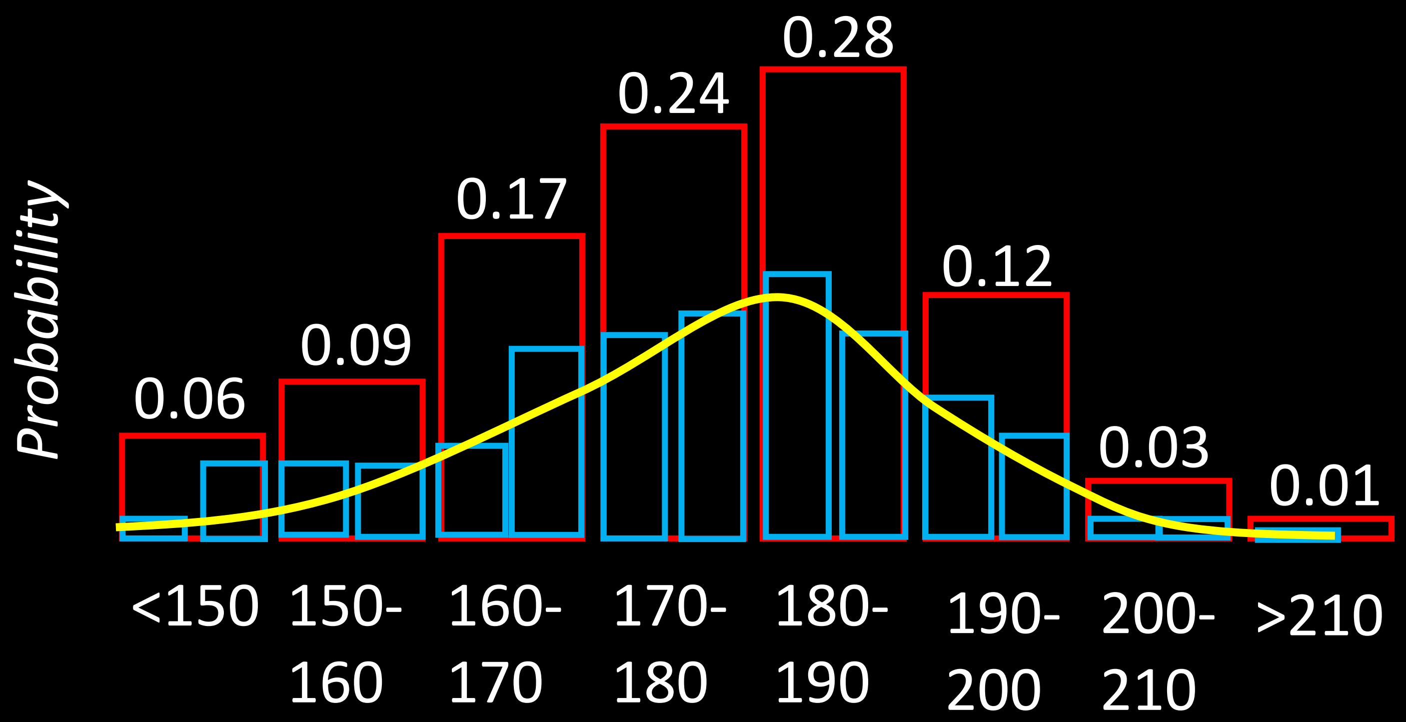

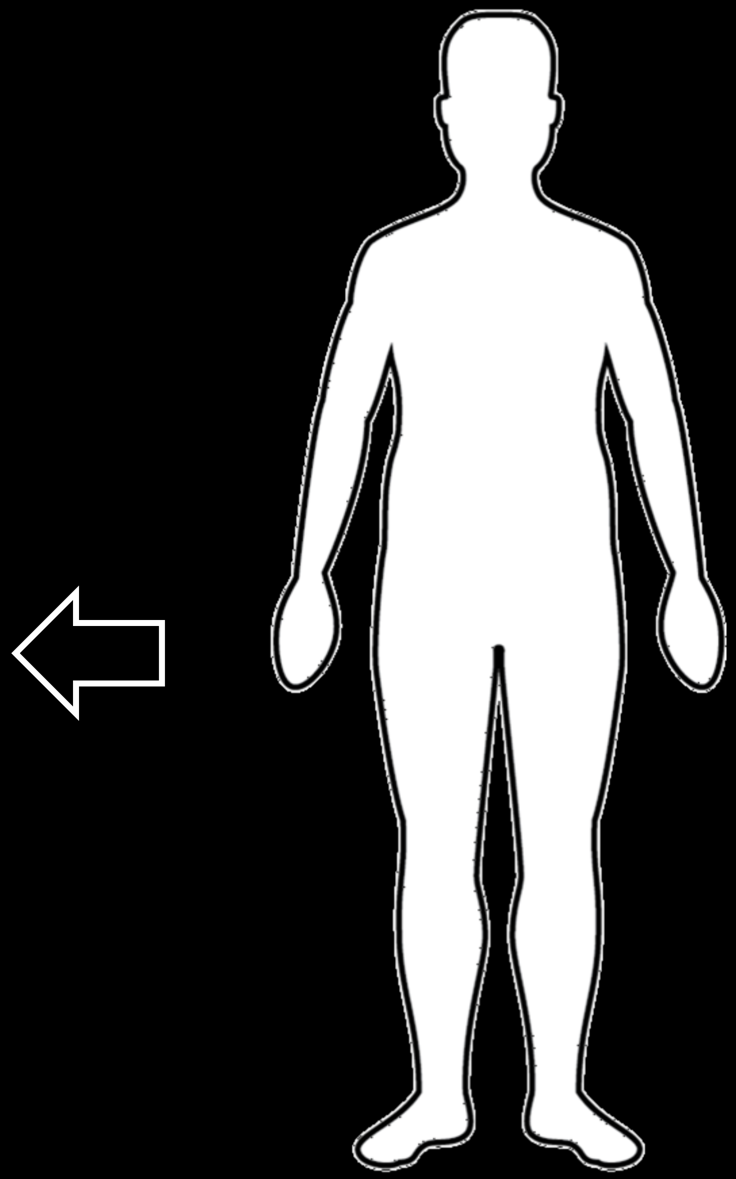

Probability Distributions

- 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: $\sigma^2 = ((8.8-7.91)^2 + (6.6-7.92)^2 + (8.33-7.91)^2)/(3-1) = 1.35$ ft/s
  - Standard deviation: $\sigma = \text{sqrt}(\sigma^2) = 1.16$ ft/s
  - Standard error: $\sigma / \text{sqrt}(3) = 0.67$ ft/s
- Bayesian Statistics
  - Probably 7.91ft/s…

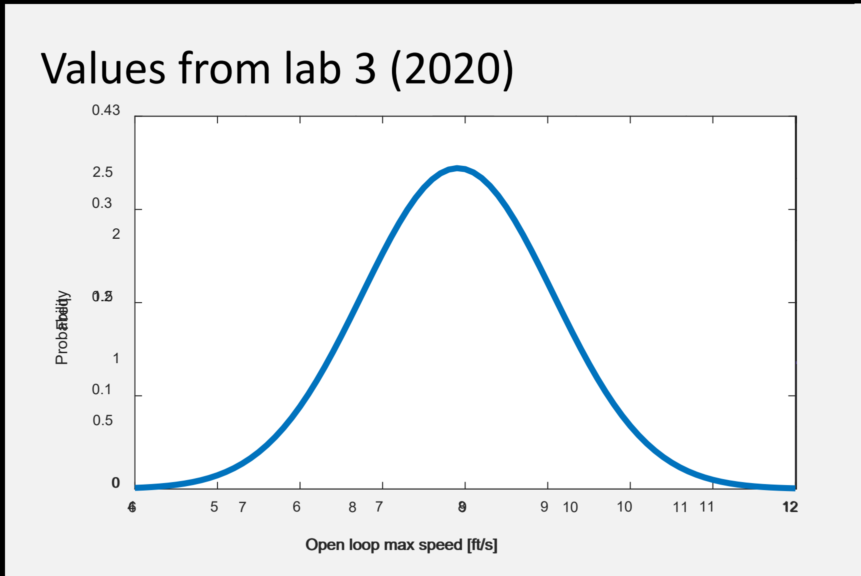

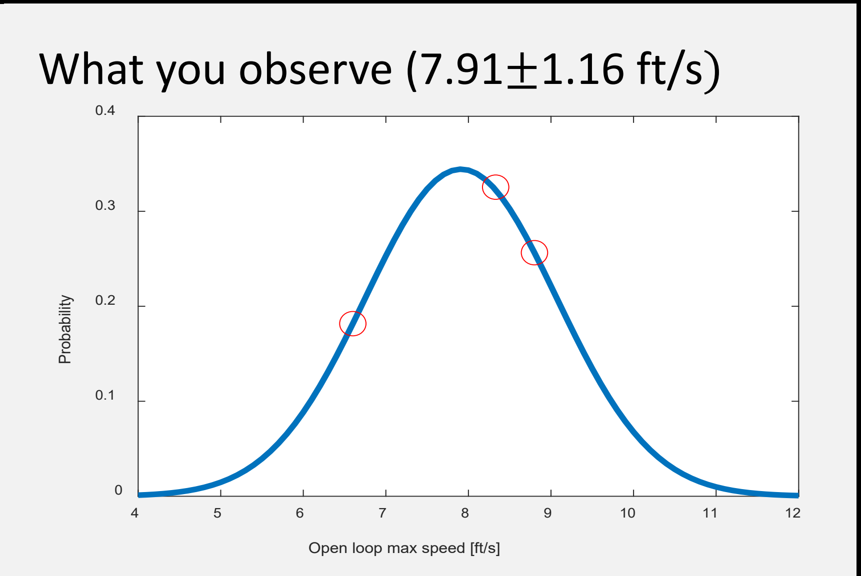

Probability Distributions

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

- Use Bayes theorem
- Instead of events x and y
  - 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

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

- Use Bayes theorem
- Start by assuming nothing
  - P(s) = uniform
  - P(s|m) = P(m|s) * c_1/c_2
  - 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)

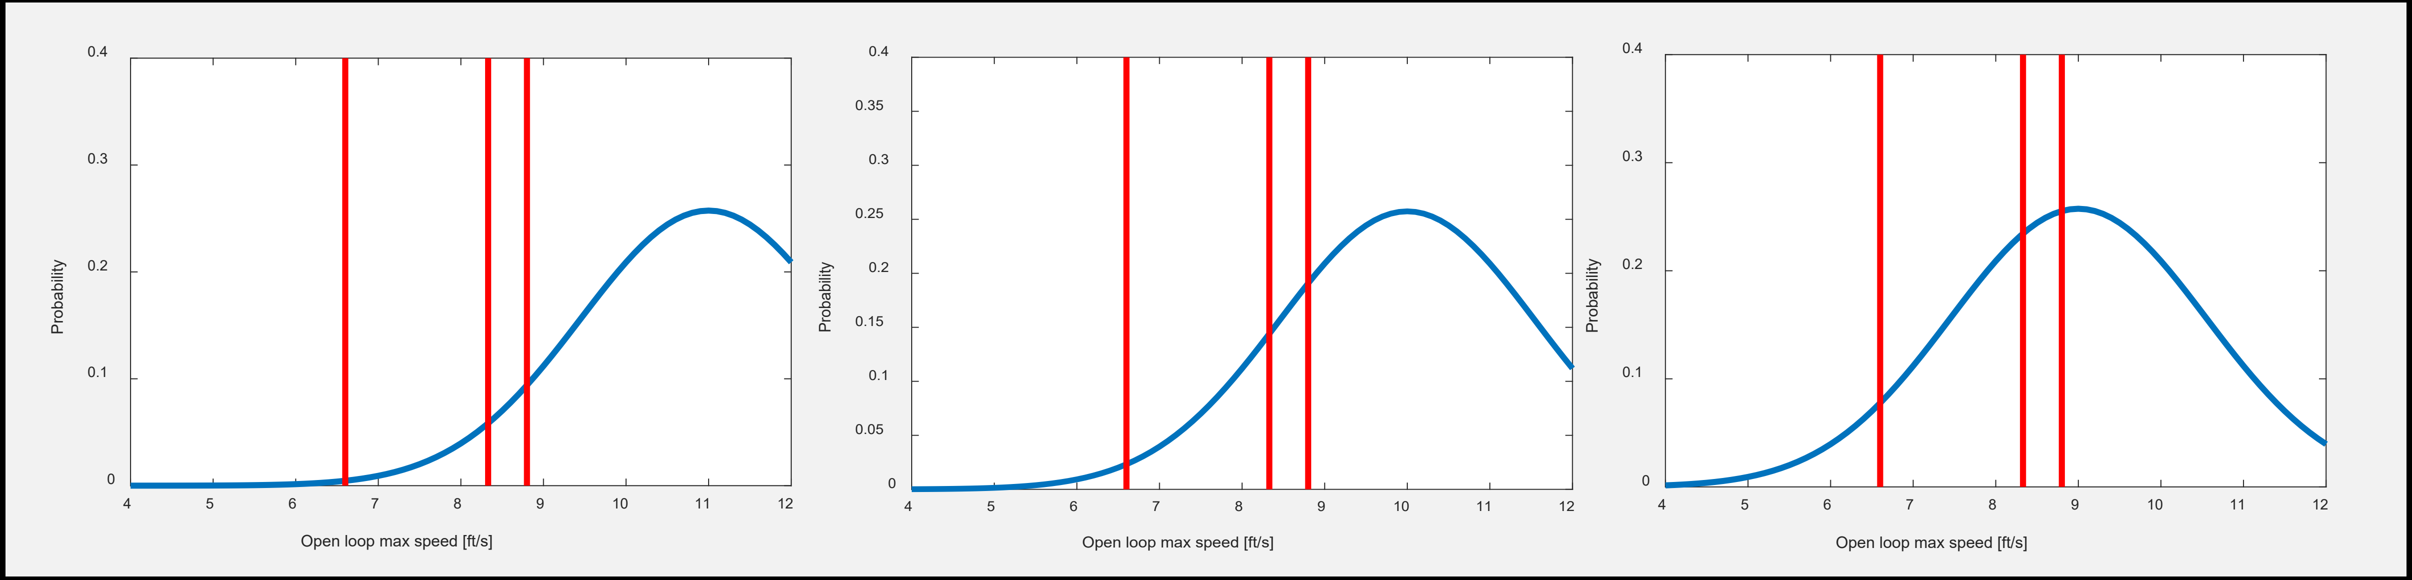

Probability Distributions

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

- Use Bayes theorem
- Start by assuming nothing
  - P(s) = uniform
  - P(s|m) = P(m|s) * c_1/c_2
  - 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)

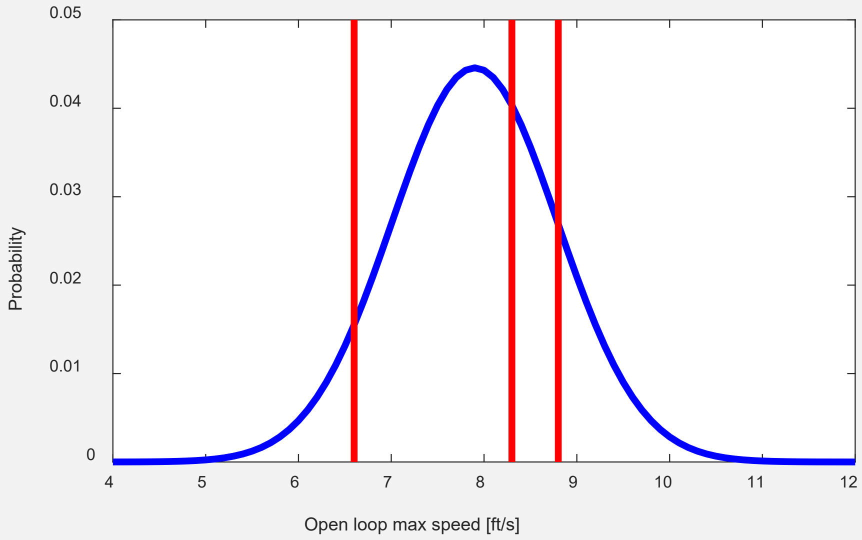

No prior:
Maximum Likelihood Estimate
(MLE)

Probability Distributions

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

- Use Bayes theorem
- Add a prior!
  - You know yesterday's speed, and you can kind of judge the current speed by eye
    - Prior: 7.91 ft/s $\pm$ 1.16ft/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!

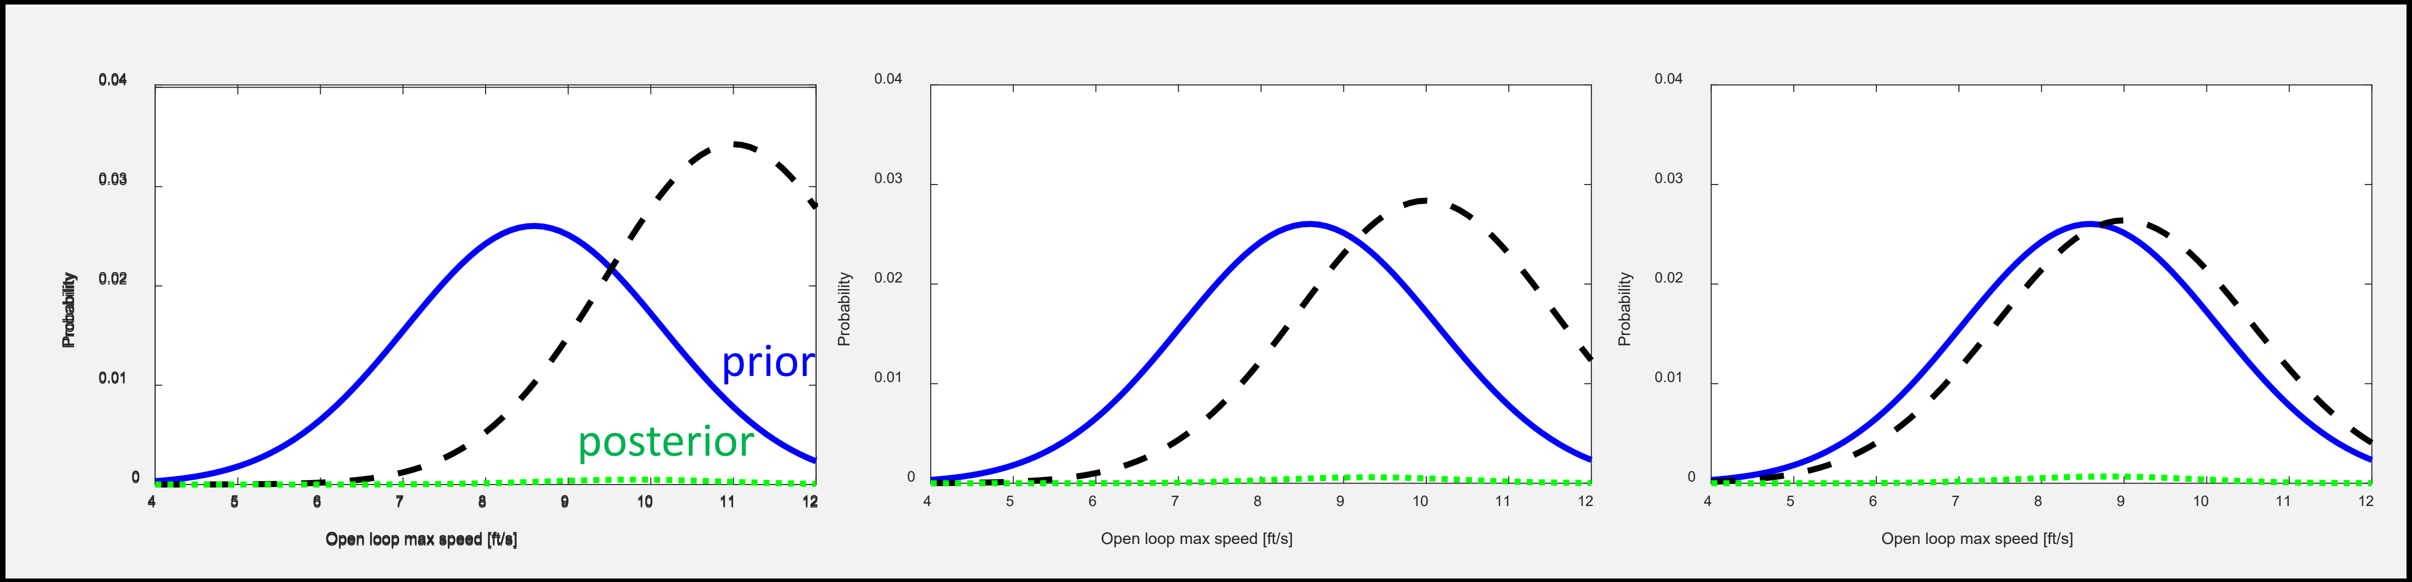

Probability Distributions

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

- Use Bayes theorem
- Add a prior!
  - You know yesterday's speed, and you can kind of judge the current speed by eye
    - Prior: 7.91 ft/s $\pm$ 1.16ft/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

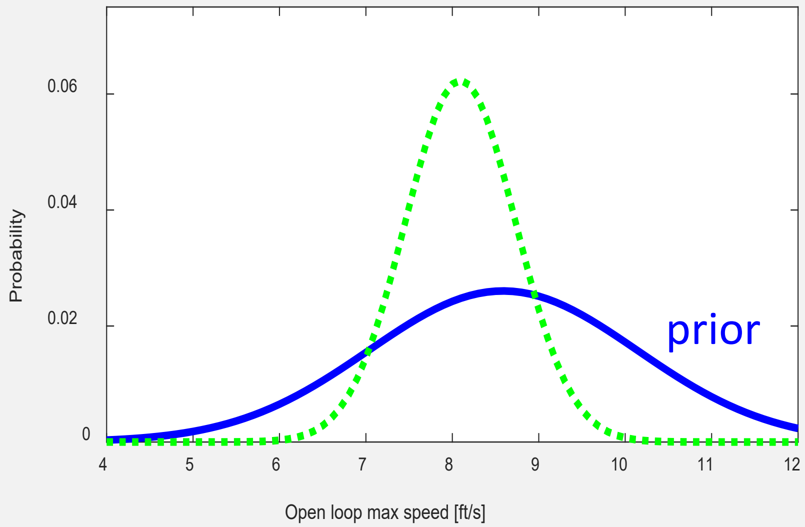

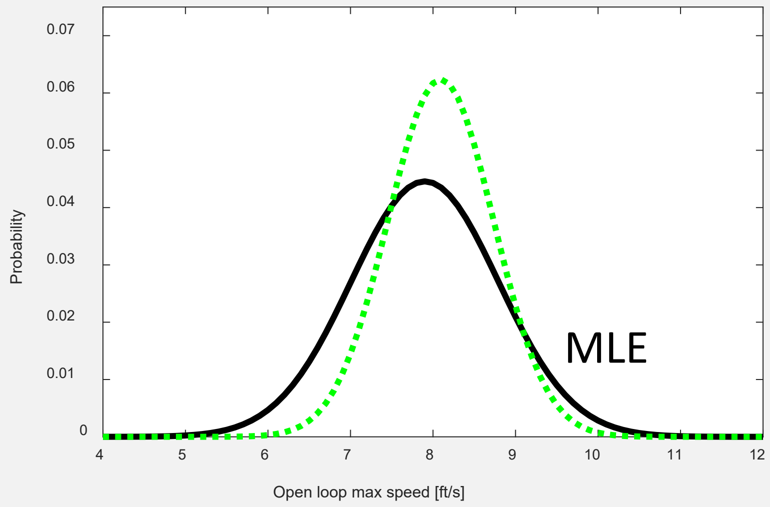

Maximum A Posteriori (MAP)

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.
- “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.”
- “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’s adventures in wonderland

References

- Probabilistic Robotics, book by *Dieter Fox, Sebastian Thrun, and Wolfram Burgard*
- How Bayes Theorem works (Youtube), by Brandon Rohrer