Sensing 3:
Probability Review

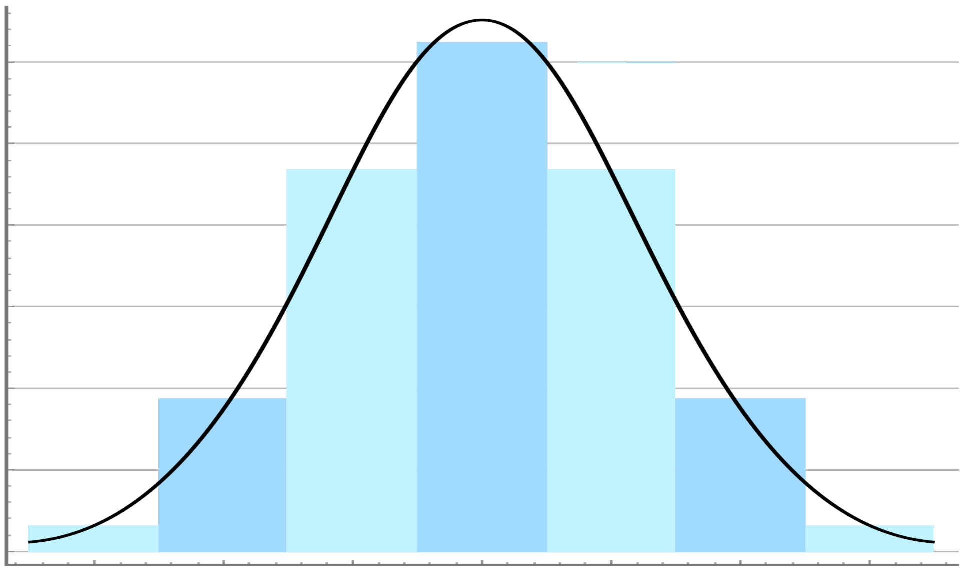

Many slides adapted from:
Siegwart, Nourbakhsh and Scaramuzza, Autonomous Mobile Robots
Thrun, Burgard and Fox, Probabilistic Robotics
Michael S. Lewicki, Probability Theory 2007, Carnegie Mellon
Russell and Norvig, Artificial Intelligence: A Modern Approach

Bookkeeping

- Quiz grades finalized, with regrades
  - E-field detection; omni wheels (up/down top/bottom left/right)
- Grades
  - Check your current grade (But don't panic)
- Assignment 1
  - Blackboard is open, assignment is up
- Projects
  - First deliverable on Tuesday
- Readings posted after class
  - Probability review; will be light reading unless you need it

High Level View

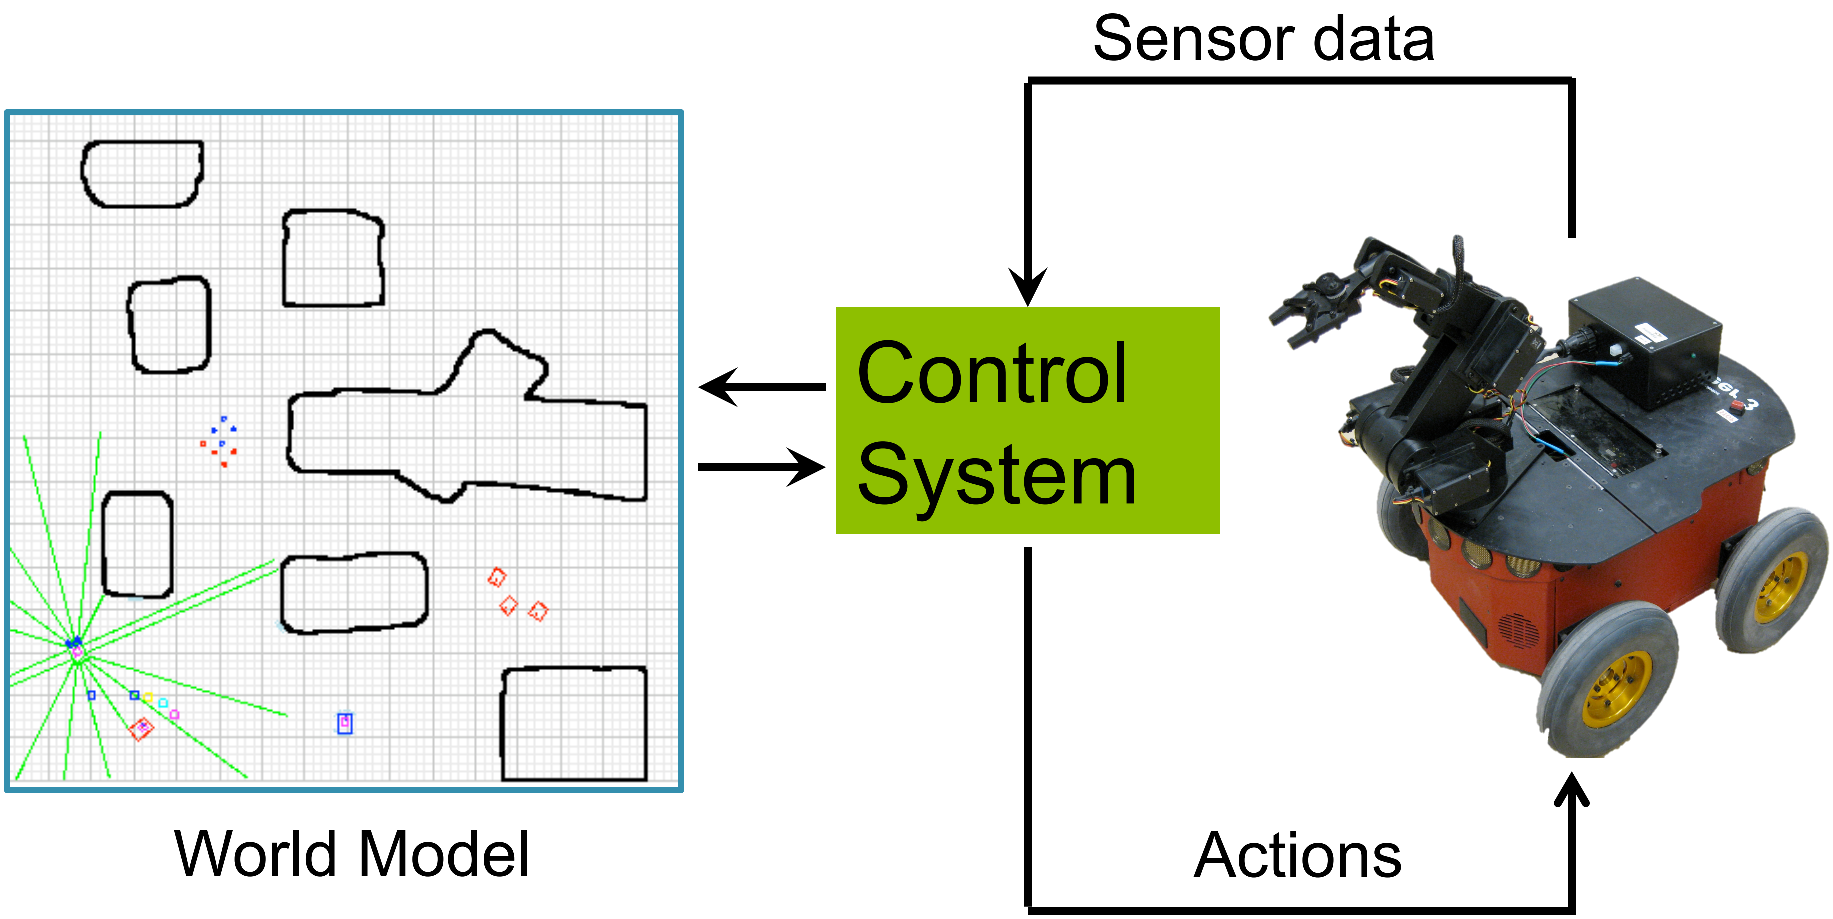

Uncertainty in Robotics

- Fundamentally, models are imperfect.
  - Sensors aren't perfect
  - Actuation isn't either
  - But you have to do *something*

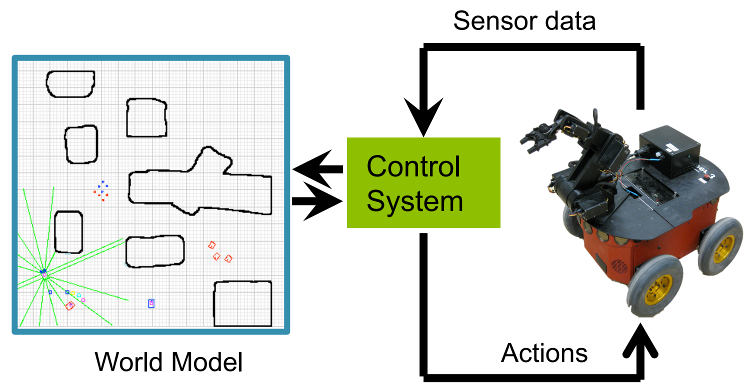

- Probability as uncertainty
  - Probability theory can be applied to these problems
- Key idea: explicit representation of uncertainty using the calculus of probability theory

Perception = state estimation

Action = utility optimization

Error and Uncertainty

- Sensing is *always* related to uncertainty.
- What are the sources of uncertainties?
  - Blown-out camera; iffy rangefinder; skidding wheel; background noise; poor speech model; what else?
- How can uncertainties be represented / quantified
  - Deterministic vs. random error
- How do they propagate?
  - Uncertainty of a function of uncertain values?
  - How do uncertainties combine if different sensor reading are *fused*?

Example: State Estimation

- Is the door open?
  - Camera + edge detection says the door is not at right angles
  - Odometry says I'm 2.0 meters away from door frame
  - Depth sensor says I'm 2.0 meters away from door

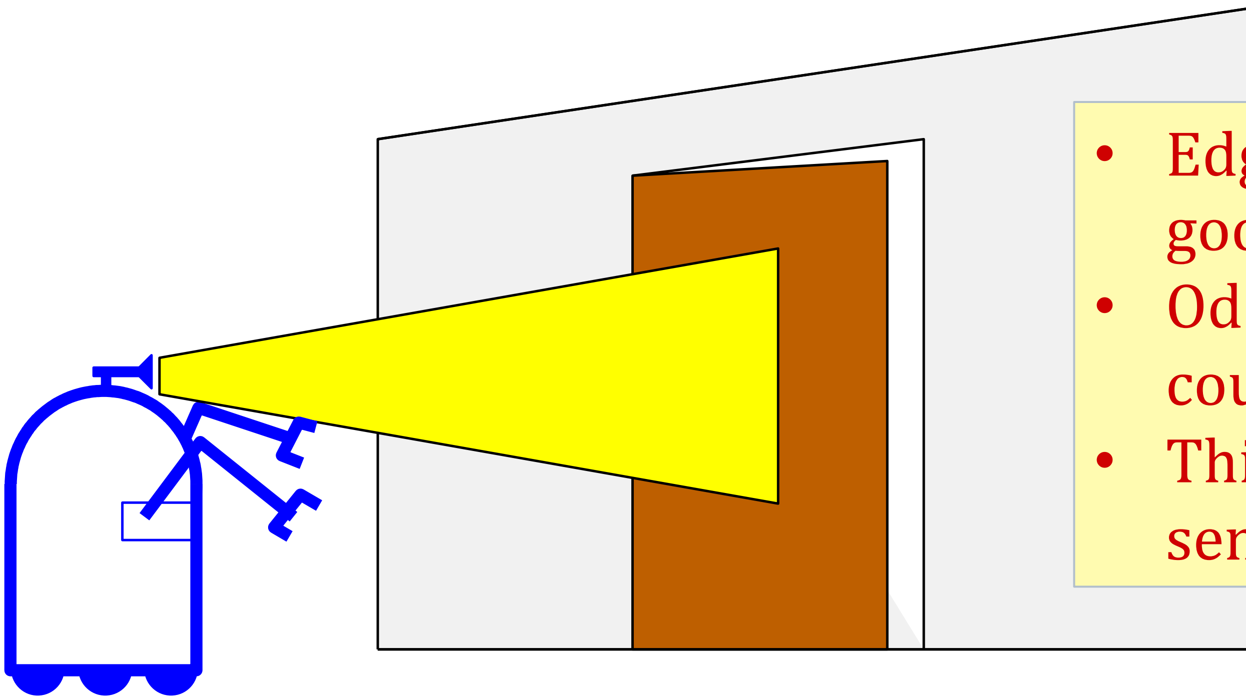

- Edge detection pretty good indoors?
- Odometry very noisy; could be off by 20cm.
- This specific depth sensor is very good

Distributions

- How can a reading be wrong?

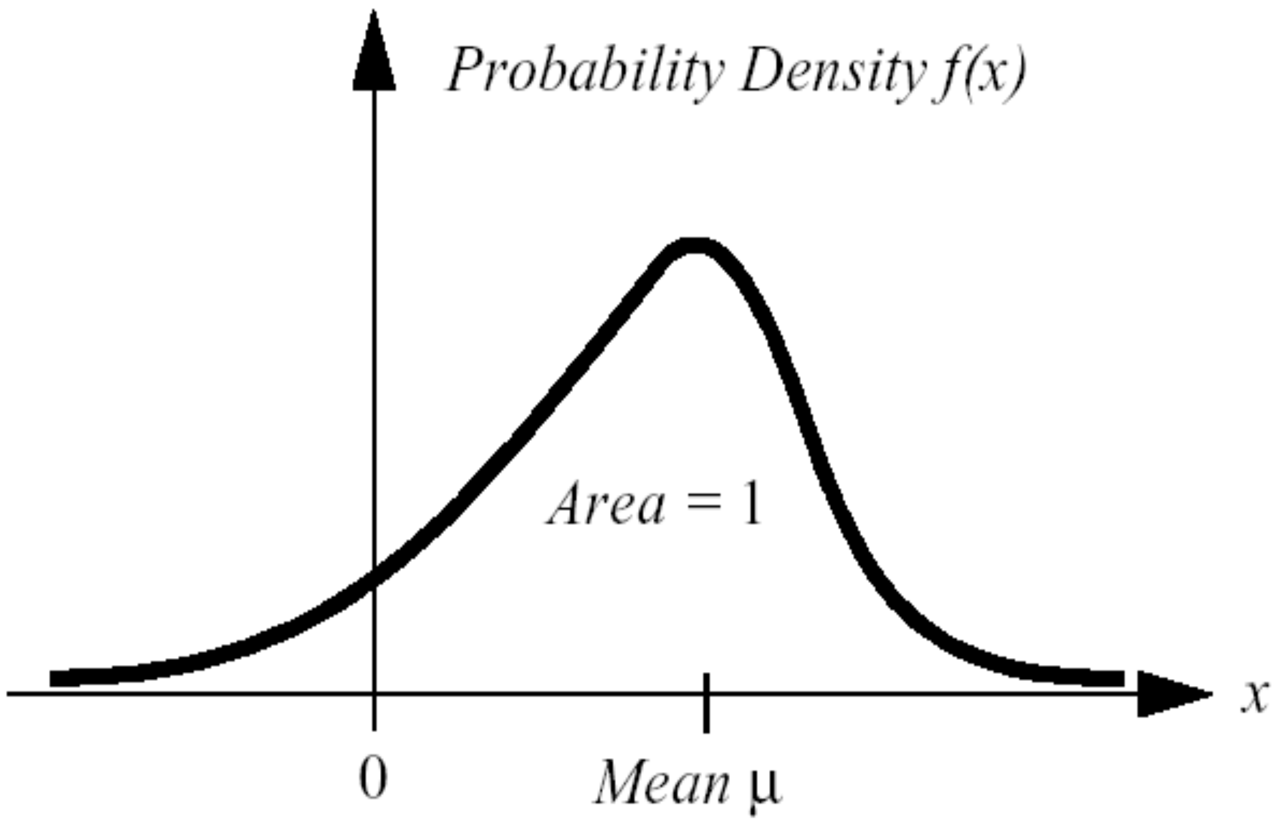

Vision

- What are we looking at?

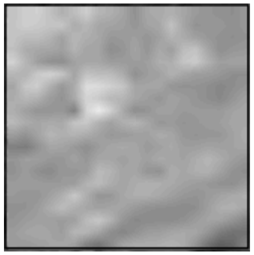

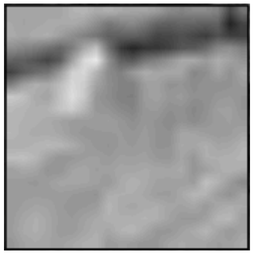

Vision

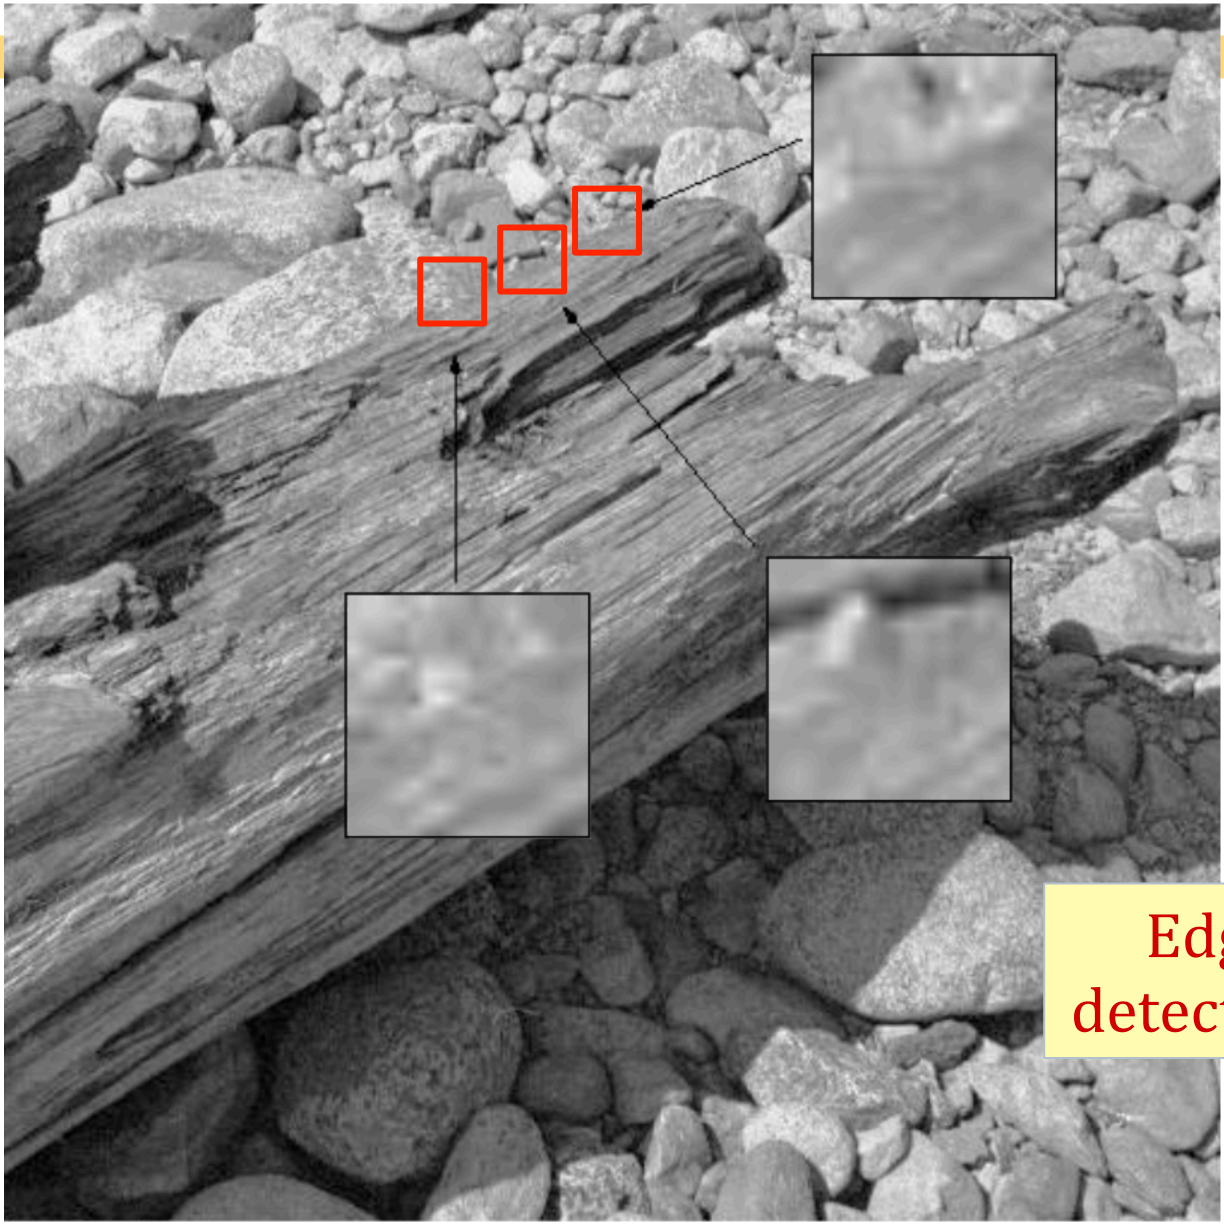

Edge detection?

Using Probability

- Making rational decisions under uncertainty
  - *Probability*
    - the precise representation of knowledge and uncertainty
  - *Probability theory*
    - How to optimally update your knowledge based on new information
  - *Decision theory: probability theory + utility theory*
    - How to use this information to achieve maximum expected utility

- Consider a bus schedule. What's the utility function?
  - A schedule says the bus comes at 8:05.

    Situation A: You have a class at 8:30.

    Situation B: You have a class at 8:30, and it's cold and raining.

    Situation C: You have a final exam at 8:30.

Discrete Random Variables

- X denotes a random variable.
- X can take countable number of values in $\{x_1, x_2, \ldots, x_n\}$
- $P(X=x_i)$ or $P(x_i)$ or $Pr(x_i)$ is the probability that the random variable X takes on value x_i.
- $P(\cdot)$ is called its probability mass function.
- E.g.

$$P(RoomType) = \langle 0.7, 0.2, 0.08, 0.02 \rangle$$

Continuous Random Variables

- X takes on values in the continuum.
- $p(X=x)$, or $p(x)$, is a probability density function.

$$\Pr(x \in (a,b)) = \int_a^b p(x)\,dx$$

- E.g.

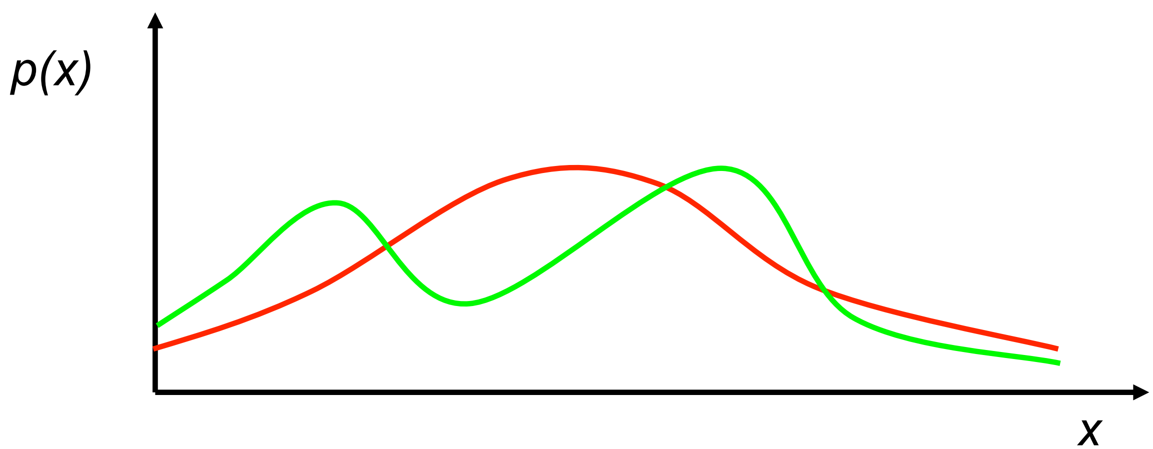

Axioms of Probability

- Pr(*A*) denotes probability that proposition *A* is true.
- Axioms (Kolmogorov):

$$0 \le \mathrm{P}(A) \le 1$$

$$\mathrm{P}(True) = 1 \qquad \mathrm{P}(False) = 0$$

$$\mathrm{P}(A \vee B) = \mathrm{P}(A) + P(B) - \mathrm{P}(A \wedge B)$$

- Corollaries:
  - A single random variable must sum to one: $\sum_{i=1}^{n} P(D = d_i) = 1$
  - The joint probability of a set of variables must also sum to 1
- If A and B are mutually exclusive: $P(A \vee B) = P(A) + P(B)$

Conditionality

- $P(B|A)$
  - Probability of event B given Event A

    - *aka* -

  - Event A has already happened,

    Now what is the chance of event B?

- $P(B \mid A)$ is the "Conditional Probability" of B given A

Rules of Probability

- Conditional probability

$$P(A \mid B) = \frac{P(A \wedge B)}{P(B)}, \qquad P(B) > 0$$

- Corollary: Bayes Law

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

$$\Rightarrow \quad P(B|A) = \frac{P(A \mid B)\,P(B)}{P(A)} = \frac{\text{likelihood} \bullet \text{prior}}{\text{evidence}}$$

Probability of an event based on conditions that may relate to that event

Bayes Bayes Bayes!

- Probability of an event based on conditions that may relate to that event

$$P(x,y) = P(x \mid y)P(y) = P(y \mid x)P(x)$$

$$\Rightarrow$$

$$P(x \mid y) = \frac{P(y \mid x)\, P(x)}{P(y)} = \frac{\text{likelihood} \cdot \text{prior}}{\text{evidence}}$$

Independence

- Two variables X, Y are independent when the probability of X is not related to the probability Y:

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

and

$$P(x \text{ and } y) = P(x) \cdot P(y)$$

for all values of X and Y

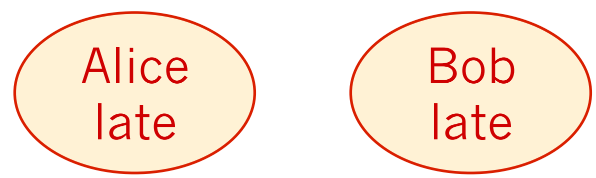

- Is Alice late to work? Is Bob late to work?

Conditional Dependence

- Two variables X, Y are *conditionally dependent* when P(X) and P(Y) each depend on a third factor, P(Z):

$$P(x,y|z)=P(x|z)P(y|z)$$

$$\Leftrightarrow$$

$$P(x|z)=P(x|z,y)$$

and

$$P(y|z)=P(y|z,x)$$

for all values of X and Y

- Alice late / Bob late / Snowing

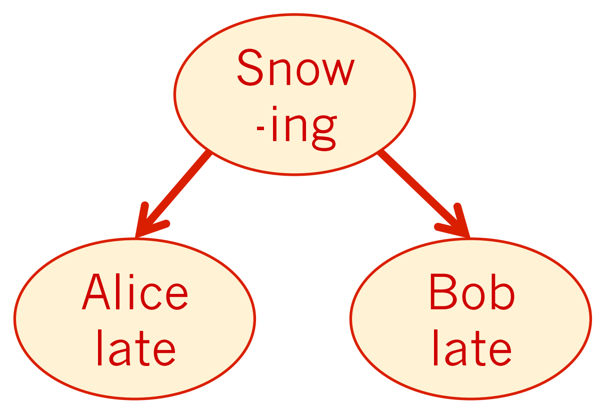

Bayes + Background Knowledge

- Probability of an event based on conditions that may relate to that event
- Example: Does Alice have cancer?
  - Alice is 65
  - *If cancer is related to age*, we can use that knowledge to improve accuracy of our assessment using Bayes

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

Conditioning

- Total probability:

$$P(x) = \int P(x, z)dz$$

$$P(x) = \int P(x \mid z)P(z)dz$$

$$P(x \mid y) = \int P(x \mid y, z)\, P(z)\, dz$$

State Estimation

- Suppose a robot obtains measurement z
  - Z = vision + edge detection
- What is $P(open|z)$?

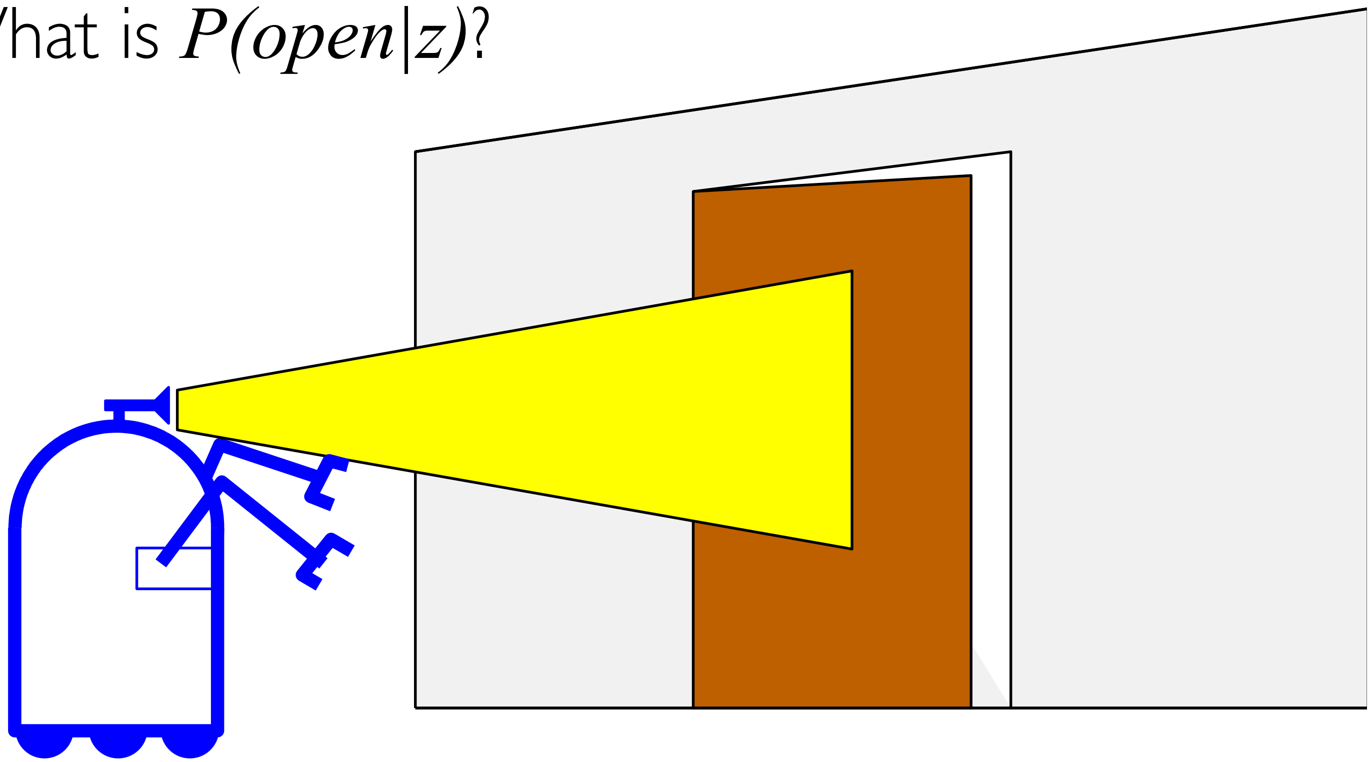

Casual (Observed) Priors

- $P(open|z)$ is diagnostic.
- $P(z|open)$ is causal.
- Often causal knowledge is easier to obtain.
- Bayes rule allows us to use causal knowledge:

count frequencies!

$$P(open \mid z) = \frac{P(z \mid open)P(open)}{P(z)}$$

Example

- $P(z|open) = 0.6$ $P(z|\neg open) = 0.3$
- $P(open) = P(\neg open) = 0.5$

$$P(open \mid z) = \frac{P(z \mid open)P(open)}{P(z \mid open)p(open) + P(z \mid \neg open)p(\neg open)}$$

$$P(open \mid z) = \frac{0.6 \cdot 0.5}{0.6 \cdot 0.5 + 0.3 \cdot 0.5} = \frac{2}{3} = 0.67$$

Measurement *z* increases probability that the door is open.

Combining Evidence

- Suppose our robot obtains another observation z_2.
- How can we integrate this new information?
- More generally, how can we estimate $P(x| z_1...z_n)$?

Recursive Bayesian Updating

$$P(x \mid z_1, \ldots, z_n) = \frac{P(z_n \mid x, z_1, \ldots, z_{n-1})\, P(x \mid z_1, \ldots, z_{n-1})}{P(z_n \mid z_1, \ldots, z_{n-1})}$$

Markov assumption: z_n is independent of $z_1, \ldots, z_{n-1}$ if we know x

$$P(x \mid z_1, \ldots, z_n) = \frac{P(z_n \mid x)\, P(x \mid z_1, \ldots, z_{n-1})}{P(z_n \mid z_1, \ldots, z_{n-1})}$$

$$= \eta\, P(z_n \mid x)\, P(x \mid z_1, \ldots, z_{n-1})$$

$$= \eta_{1 \ldots n} \prod_{i=1 \ldots n} P(z_i \mid x)\, P(x)$$

$P(B|A)$: probability of B given A

Second Measurement

- $P(z_2|open) = 0.5$ $\qquad P(z_2|\neg open) = 0.6$
- $P(open|z_1)=2/3$

$$P(open \mid z_2, z_1) = \frac{P(z_2 \mid open)\, P(open \mid z_1)}{P(z_2 \mid open)\, P(open \mid z_1) + P(z_2 \mid \neg open)\, P(\neg open \mid z_1)}$$

$$= \frac{\frac{1}{2} \cdot \frac{2}{3}}{\frac{1}{2} \cdot \frac{2}{3} + \frac{3}{5} \cdot \frac{1}{3}} = \frac{5}{8} = 0.625$$

Measurement z decreases probability that the door is open.