Autonomous Robotics

Probability Review and Interaction Model

Josiah Hanna
University of Wisconsin — Madison

Credit: the many AI faculty who developed some of these slides for CS 540 at UW — Madison.

Learning Outcomes

After today's lecture, you will:

- Understand the foundational topics in probability necessary for this course.
- Be able to describe the fundamental parts of a general model of robot-environment interaction.

Why probability?

- Represent uncertainty in the world.
- Represent beliefs about the state of the world.

Probability in robotics

- Represent beliefs about the true state of the world.
- Represent uncertainty about the effects of actions.
- Represent uncertainty about what observation is produced in different states.

Discrete Random Variables

- Let X be a random variable that takes on a value $x \in \mathcal{X}$, where $\mathcal{X}$ is a set with a finite number of elements.
    - Example: the result of rolling a single dice.
- $p(X = x)$ is the probability that X takes on the value x.
- $\sum_x p(X = x) = 1$ and $\forall x, 0 \leq p(X = x) \leq 1$.
- For compactness, write $p(x)$.
- p is a probability mass function.

Continuous Random Variables

- Or X could be a random variable that takes on a value $x \in \mathcal{X}$, where $\mathcal{X}$ is a continuous set.
  - Example: the height of the first person you see after leaving this classroom.
- $p(X = x)$ is the probability that X takes on the value x.
- $\int_x p(X = x)dx = 1$ and $\forall x, p(X = x) \geq 0$.
- p is a probability density function.

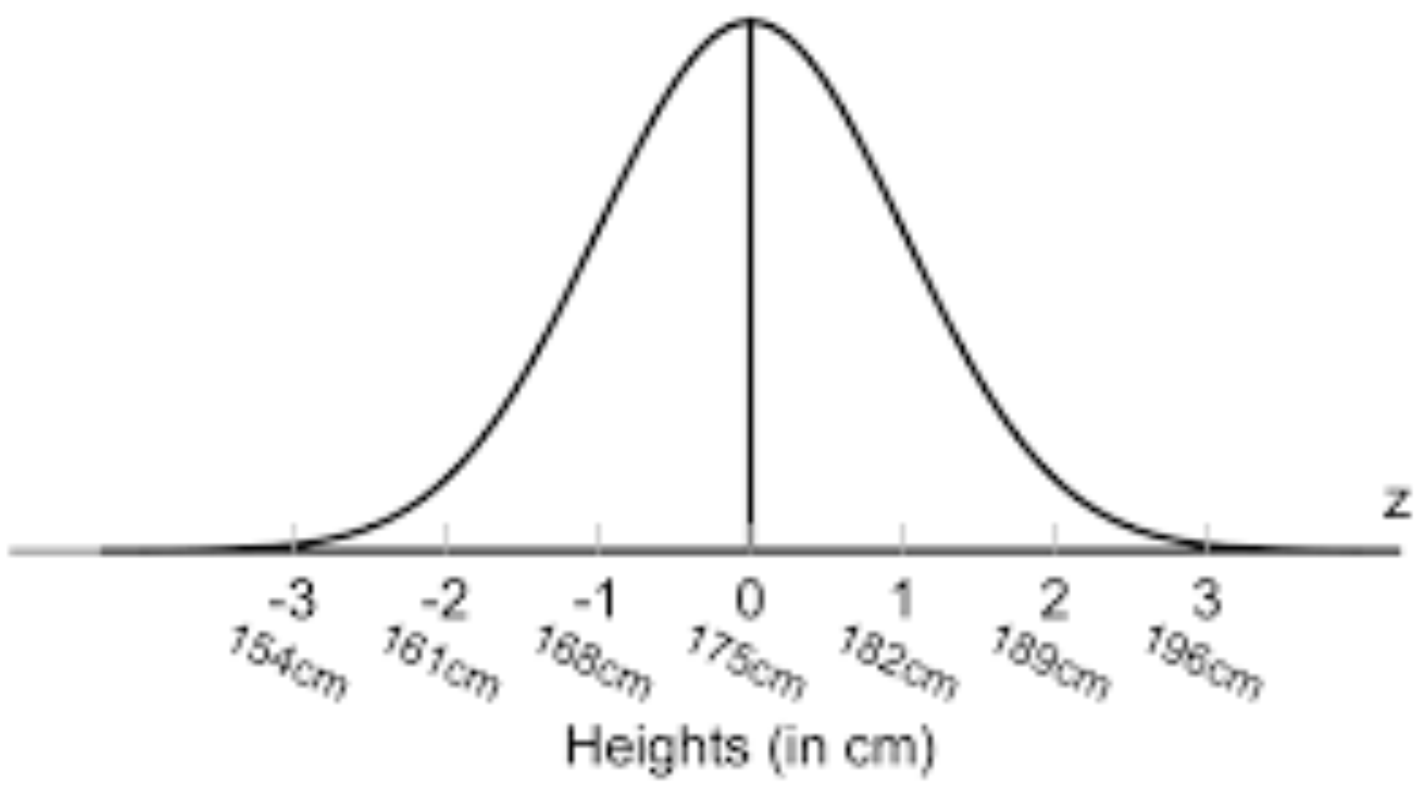

Random Sampling

- Sampling is assigning a value to a random variable according to some probability distribution (either a pmf or a pdf).
  - Example: roll a dice and observe the outcome.
- Write $X \sim p$ to denote that variable X has value distributed according to p.

Joint Distributions

- Move from one variable to multiple variables.

- Joint distribution of X and Y: $p(X = a, Y = b)$

  - Why? Work with multiple types of uncertainty and model interactions.

Marginal Distributions

- Given a joint distribution: $P(X = a, Y = b)$
  - Get the distribution in just one variable:

$$P(X = a) = \sum_b P(X = a, Y = b)$$

  - This is the "marginal" distribution of X.
  - "Marginalize out" the other variable, Y.

Basics: **Marginal** Probability

$$P(X = a) = \sum_b P(X = a, Y = b)$$

	Sunny	Cloudy	Rainy
hot	150/365	40/365	5/365
cold	50/365	60/365	60/365

$$[P(\text{hot}), P(\text{cold})] = [\tfrac{195}{365}, \tfrac{170}{365}]$$

Probability Tables

- Write our distributions as tables

	Sunny	Cloudy	Rainy
hot	150/365	40/365	5/365
cold	50/365	60/365	60/365

- # of entries? 6.
  - If we have n variables with k values, we get k^n entries
  - **Big!** For a 1080p screen, 12 bit color, size of table: $10^{7490589}$
  - No way of writing down all the terms.

Independence

- Two random variables are independent if

$$P(X,Y) = P(X)P(Y)$$

- Why useful? Go from k^n entries in a table to $\approx kn$.
- Collapse joint distribution into the **product** of marginals.

Independence Example

$$p(X = a, Y = b) = p(X = a)p(Y = b)$$

Conditional Probability

- Express how knowledge of one variable changes belief about another variable,

$$P(X = a|Y = b) = \frac{P(X = a, Y = b)}{P(Y = b)}$$

- Variables can be conditionally independent:

$$P(X, Y|Z) = P(X|Z)P(Y|Z)$$

Conditional Independence Example

Bayesian Network: nodes are variables and edges show probabilistic dependencies.

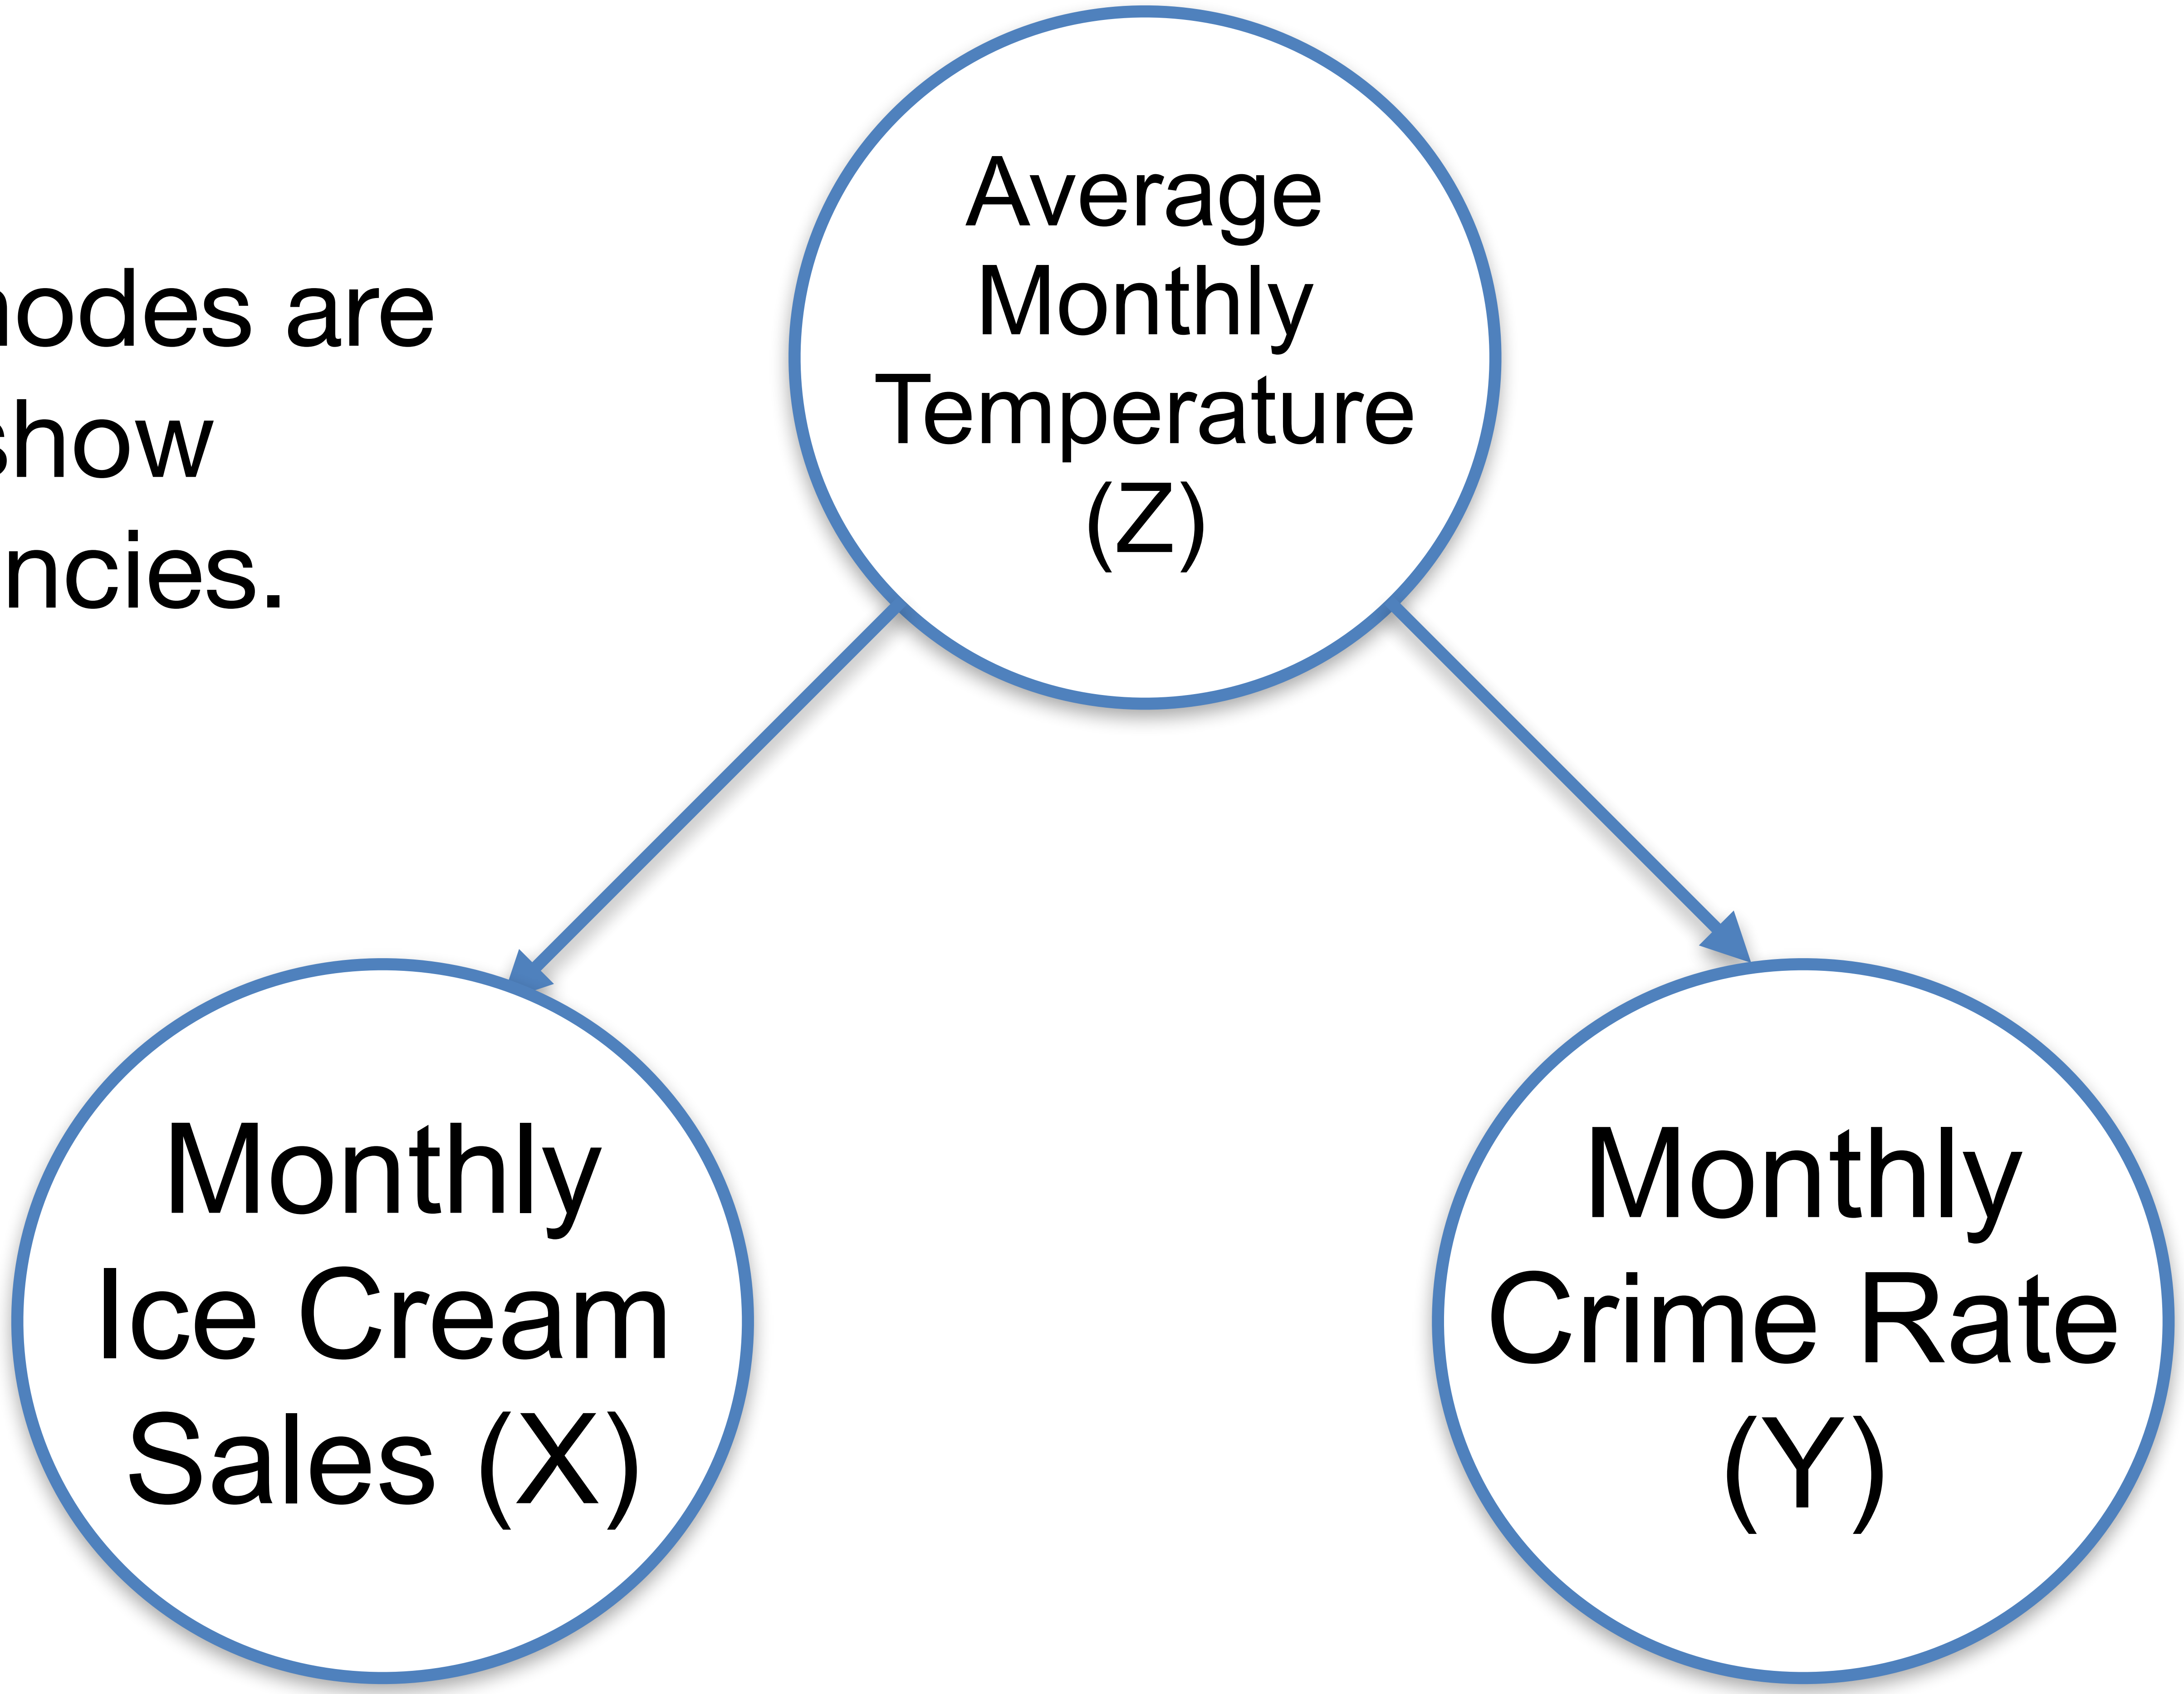

$$p(X = a, Y = b) \neq p(X = a)p(Y = b)$$

Conditional Independence Example

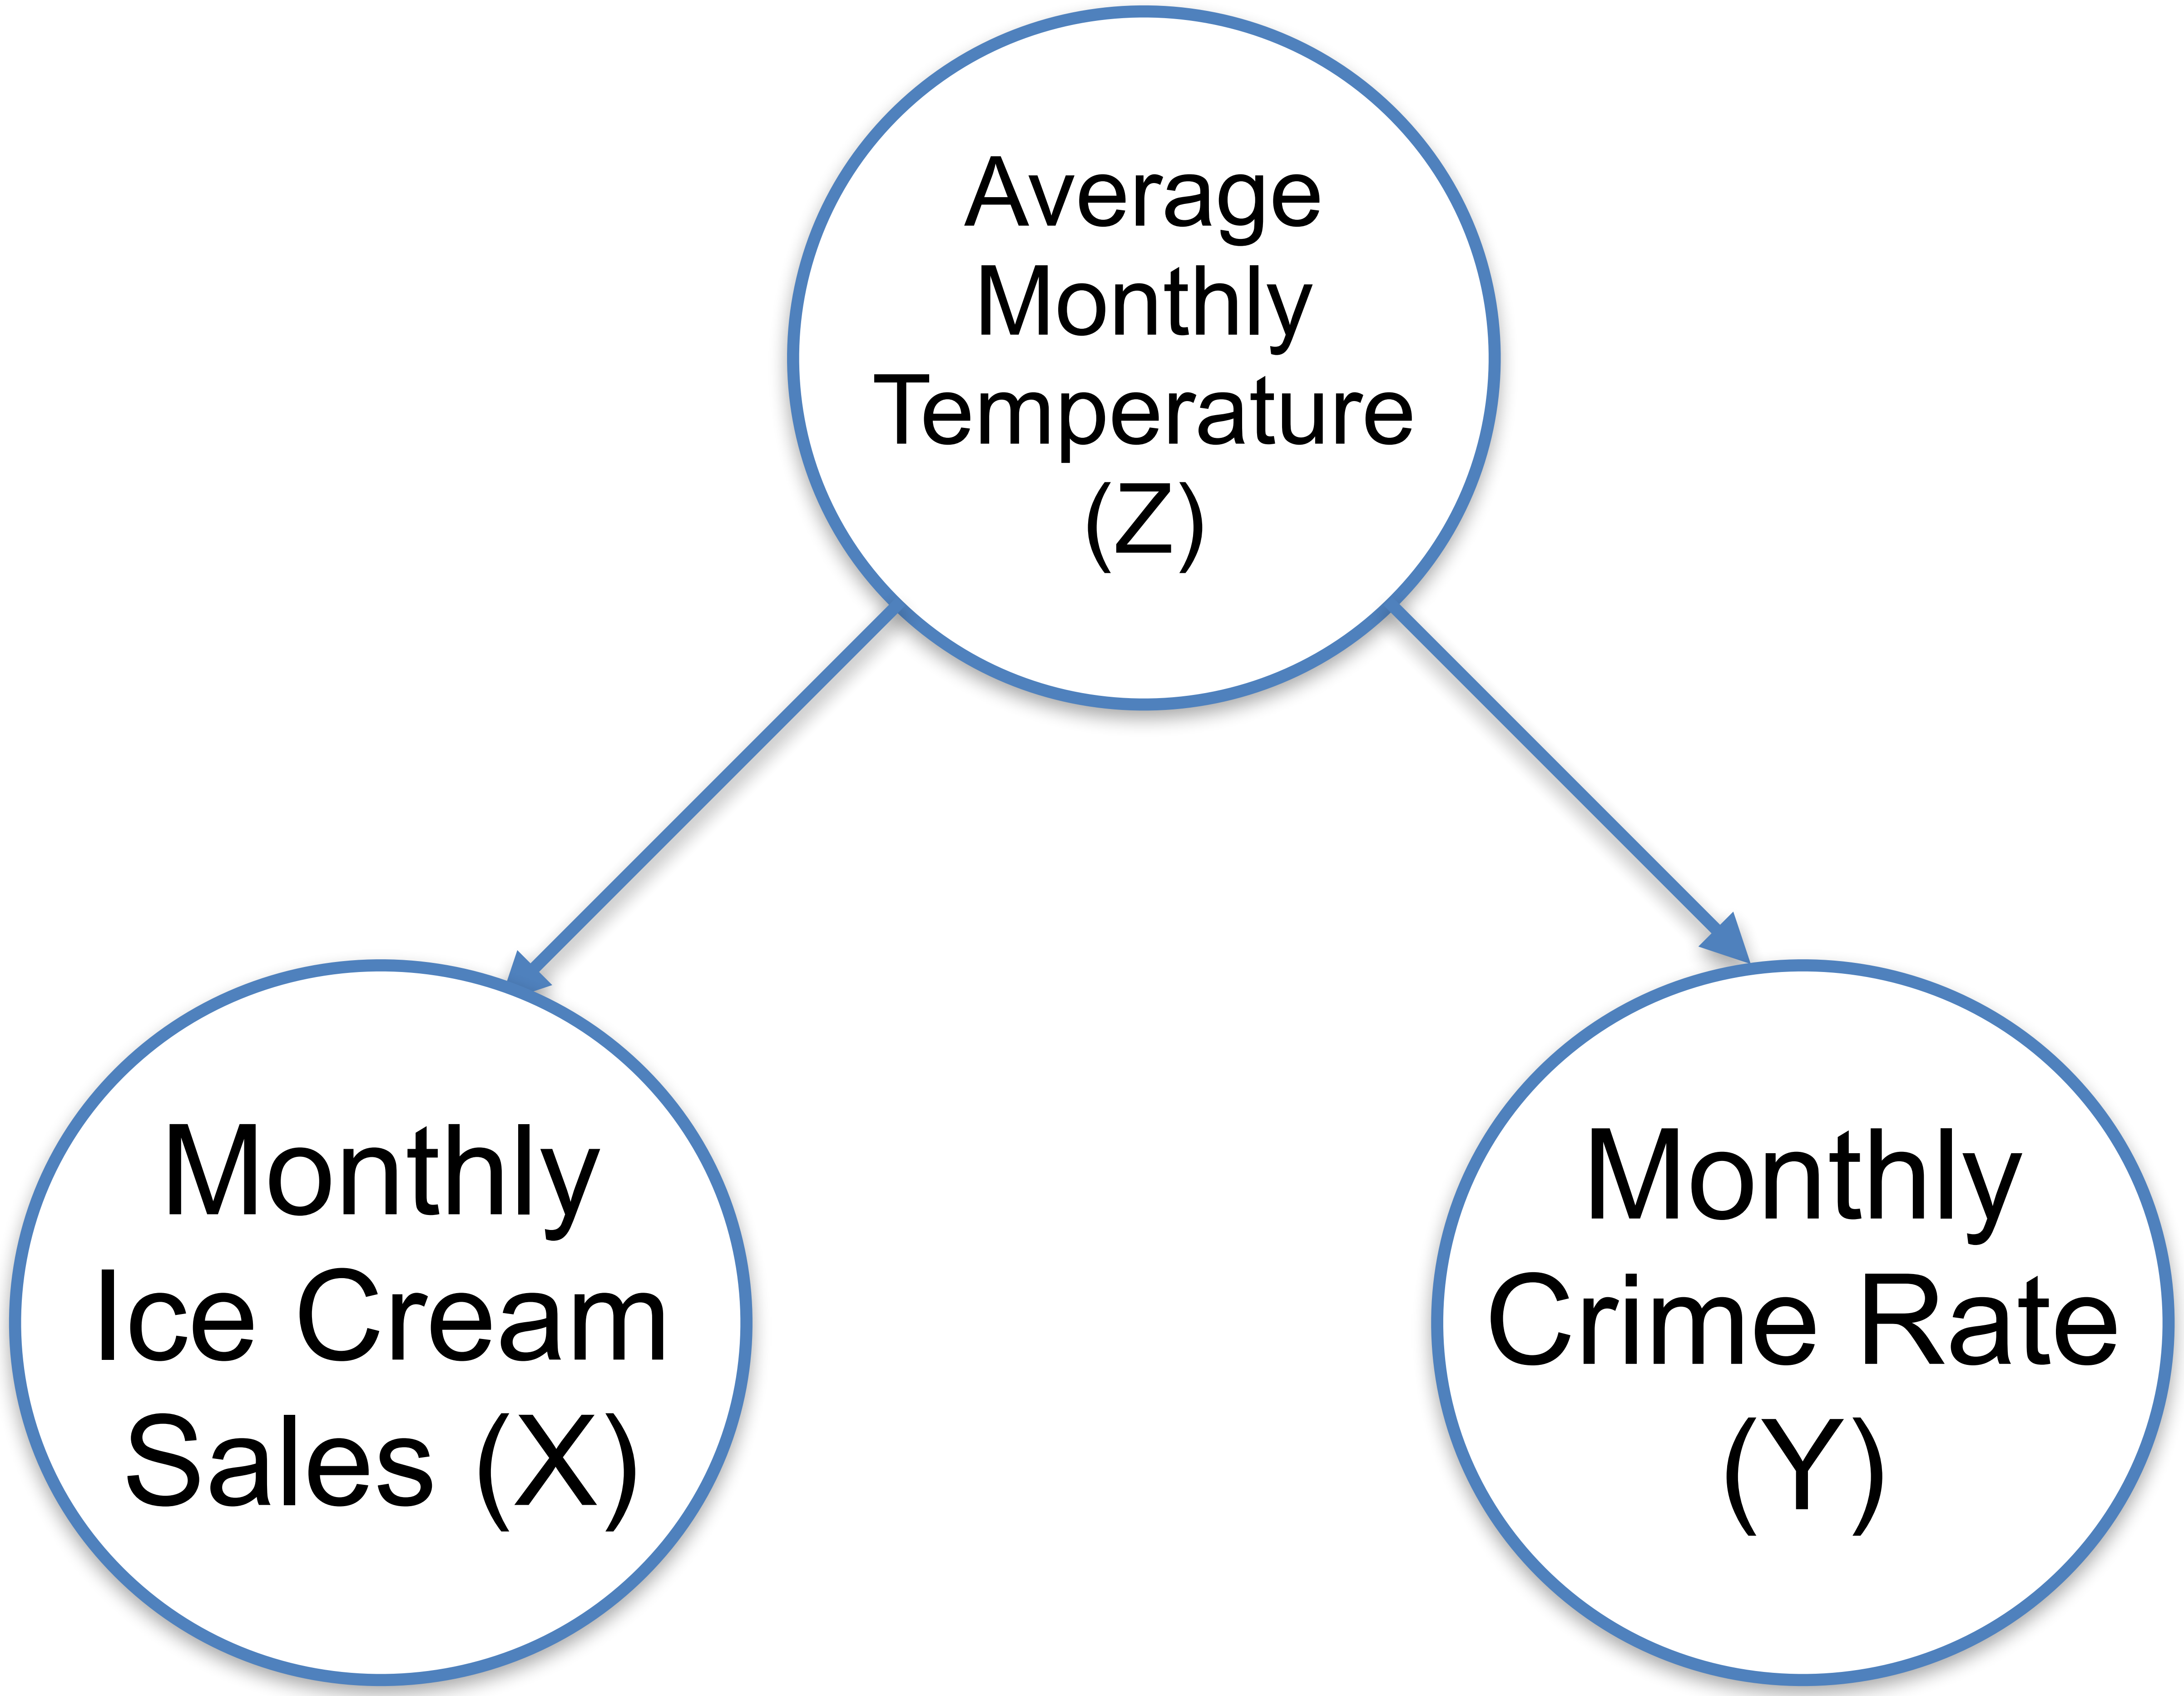

$$p(X = a, Y = b \mid Z = c) = p(X = a \mid Z = c)p(Y = b \mid Z = c)$$

Reasoning With Conditional Distributions

- Evaluating probabilities:
  - Wake up with a sore throat.
  - Do I have the flu?

- Logic approach: $S \rightarrow F$
  - Too strong.
- **Inference**: compute probability given evidence $P(F|S)$
  - Can be much more complex!

Using **Bayes' Rule**

- Want: $P(F|S)$
- **Bayes' Rule:** $P(F|S) = \frac{P(F,S)}{P(S)} = \frac{P(S|F)P(F)}{P(S)}$
- Parts:
  - $P(S) = 0.1$ Sore throat rate
  - $P(F) = 0.01$ Flu rate
  - $P(S|F) = 0.9$ Sore throat rate among flu sufferers

So: $P(F|S) = 0.09$

Using Bayes' Rule

- Interpretation $P(F|S) = 0.09$
  - Much higher chance of flu than normal rate (0.01).
  - Very different from $P(S|F) = 0.9$
    - 90% of folks with flu have a sore throat.
    - But, only 9% of folks with a sore throat have flu.

- Idea: **update** probabilities from evidence

Bayesian **Inference**

- Fancy name for what we just did. Terminology:

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)}$$

- *H* is the hypothesis
- *E* is the evidence

Bayesian Inference

- Terminology:

$$P(H|E) = \frac{P(E|H)\textcolor{red}{P(H)}}{P(E)} \longleftarrow \textcolor{red}{\textbf{Prior}}$$

- Prior: estimate of the probability **without** evidence

Bayesian Inference

- Terminology:

$$P(H|E) = \frac{\overbrace{P(E|H)}^{\text{Likelihood}}P(H)}{P(E)}$$

- Likelihood: probability of evidence **given a hypothesis**.

Bayesian Inference

- Terminology:

$$\underbrace{P(H|E)}_{\text{Posterior}} = \frac{P(E|H)P(H)}{P(E)}$$

- Posterior: probability of hypothesis **given evidence**.

Quick Quiz

A robot is placed in front of a door that has an equal probability of being open or closed. The robot has a sensor that measures if the door is open or closed. With probability 0.75, the sensor gives the correct measurement and otherwise gives the opposite response. The robot receives two independent sensor readings that both indicate the door is open. What should a Bayesian robot believe about the true state of the door?

$$p(D \mid O_1, O_2) = \frac{p(D, O_1, O_2)}{p(O_1, O_2)}$$

$$= \frac{p(D)p(O_1 \mid D)p(O_2 \mid D)}{p(D)p(O_1 \mid D)p(O_2 \mid D) + p(\overline{D})p(O_1 \mid \overline{D})p(O_2 \mid \overline{D})}$$

$$= \frac{\frac{1}{2}\frac{3}{4}\frac{3}{4}}{\frac{1}{2}\frac{3}{4}\frac{3}{4} + \frac{1}{2}\frac{1}{4}\frac{1}{4}} = \frac{9}{10}$$

Quick Quiz

A robot is placed in front of a door that has an equal probability of being open or closed. The robot has a sensor that measures if the door is open or closed. With probability 0.75, the sensor gives the correct measurement and otherwise gives the opposite response. The robot receives two independent sensor readings that both indicate the door is open. What should a Bayesian robot believe about the true state of the door?

First update:

$$p(D \mid O_1) \propto p(D)p(O_1 \mid D) = (\frac{1}{2})(\frac{3}{4}) = 3/8$$

$$p(\overline{D} \mid O_1) \propto p(\overline{D})p(O_1 \mid \overline{D}) = (\frac{1}{2})(\frac{1}{4}) = 1/8$$

$$\Longrightarrow \quad p(D \mid O_1) = 3/4$$

Quick Quiz

A robot is placed in front of a door that has an equal probability of being open or closed. The robot has a sensor that measures if the door is open or closed. With probability 0.75, the sensor gives the correct measurement and otherwise gives the opposite response. The robot receives two independent sensor readings that both indicate the door is open. What should a Bayesian robot believe about the true state of the door?

Note: $p(O_2 \mid D, O_1) = p(O_2 \mid D)$ by independence assumption.

Second update:

$$p(D \mid O_1, O_2) \propto p(D \mid O_1)p(O_2 \mid D, O_1) = (\frac{3}{4})(\frac{3}{4}) = 9/16$$

$$p(\overline{D} \mid O_1, O_2) \propto p(\overline{D} \mid O_1)p(O_2 \mid \overline{D}, O_1) = (\frac{1}{4})(\frac{1}{4}) = 1/16$$

$$\Longrightarrow \quad p(D \mid O_1, O_2) = 9/10$$

(Bayes Filter (Week 3))

Learning Outcomes

After today's lecture, you will:

- Understand the foundational topics in probability necessary for this course.

- **Be able to describe the fundamental parts of a general model of robot-environment interaction.**

States

- From *Probabilistic Robotics*: the collection of all aspects of the robot and its environment that can impact the future.

- Examples: robot pose, battery life, location of people, velocity

- Aspects that change (dynamic state) vs aspects that don't (static state).

- State variables often take on continuous values.

  - Example: The pose of a robot in a plane is a point in $\mathbb{R}^3$.

- Write x_t to define the state at time t.

Observations

- Information about the state of the environment at a moment in time.
- Perceived by the robot through its sensors.
    - Also called measurements or percepts.
- Typically, do NOT fully reveal the state.
    - Observations can be noisy. Example: lidar scan returns noisy distance readings.
    - Observations can be partial. Example: occlusion hides some aspects of state.
- Write z_t or (sometimes) y_t to denote the observation at time t.

Markov Assumption

- We will assume that the state is defined in a way that is sufficient for predicting the future.
  - Call this a *complete* or *Markov* state.

Control at time t.

- Formally, we say that $p(x_{t+1} \mid x_t, u_t) = p(x_{t+1} \mid x_{0:t}, u_{0:t})$.
  - Knowing the past does not help you predict the next state any better.
- This assumption is for developing tractable algorithms and often only holds approximately in practice.

Probabilistic Interaction Model

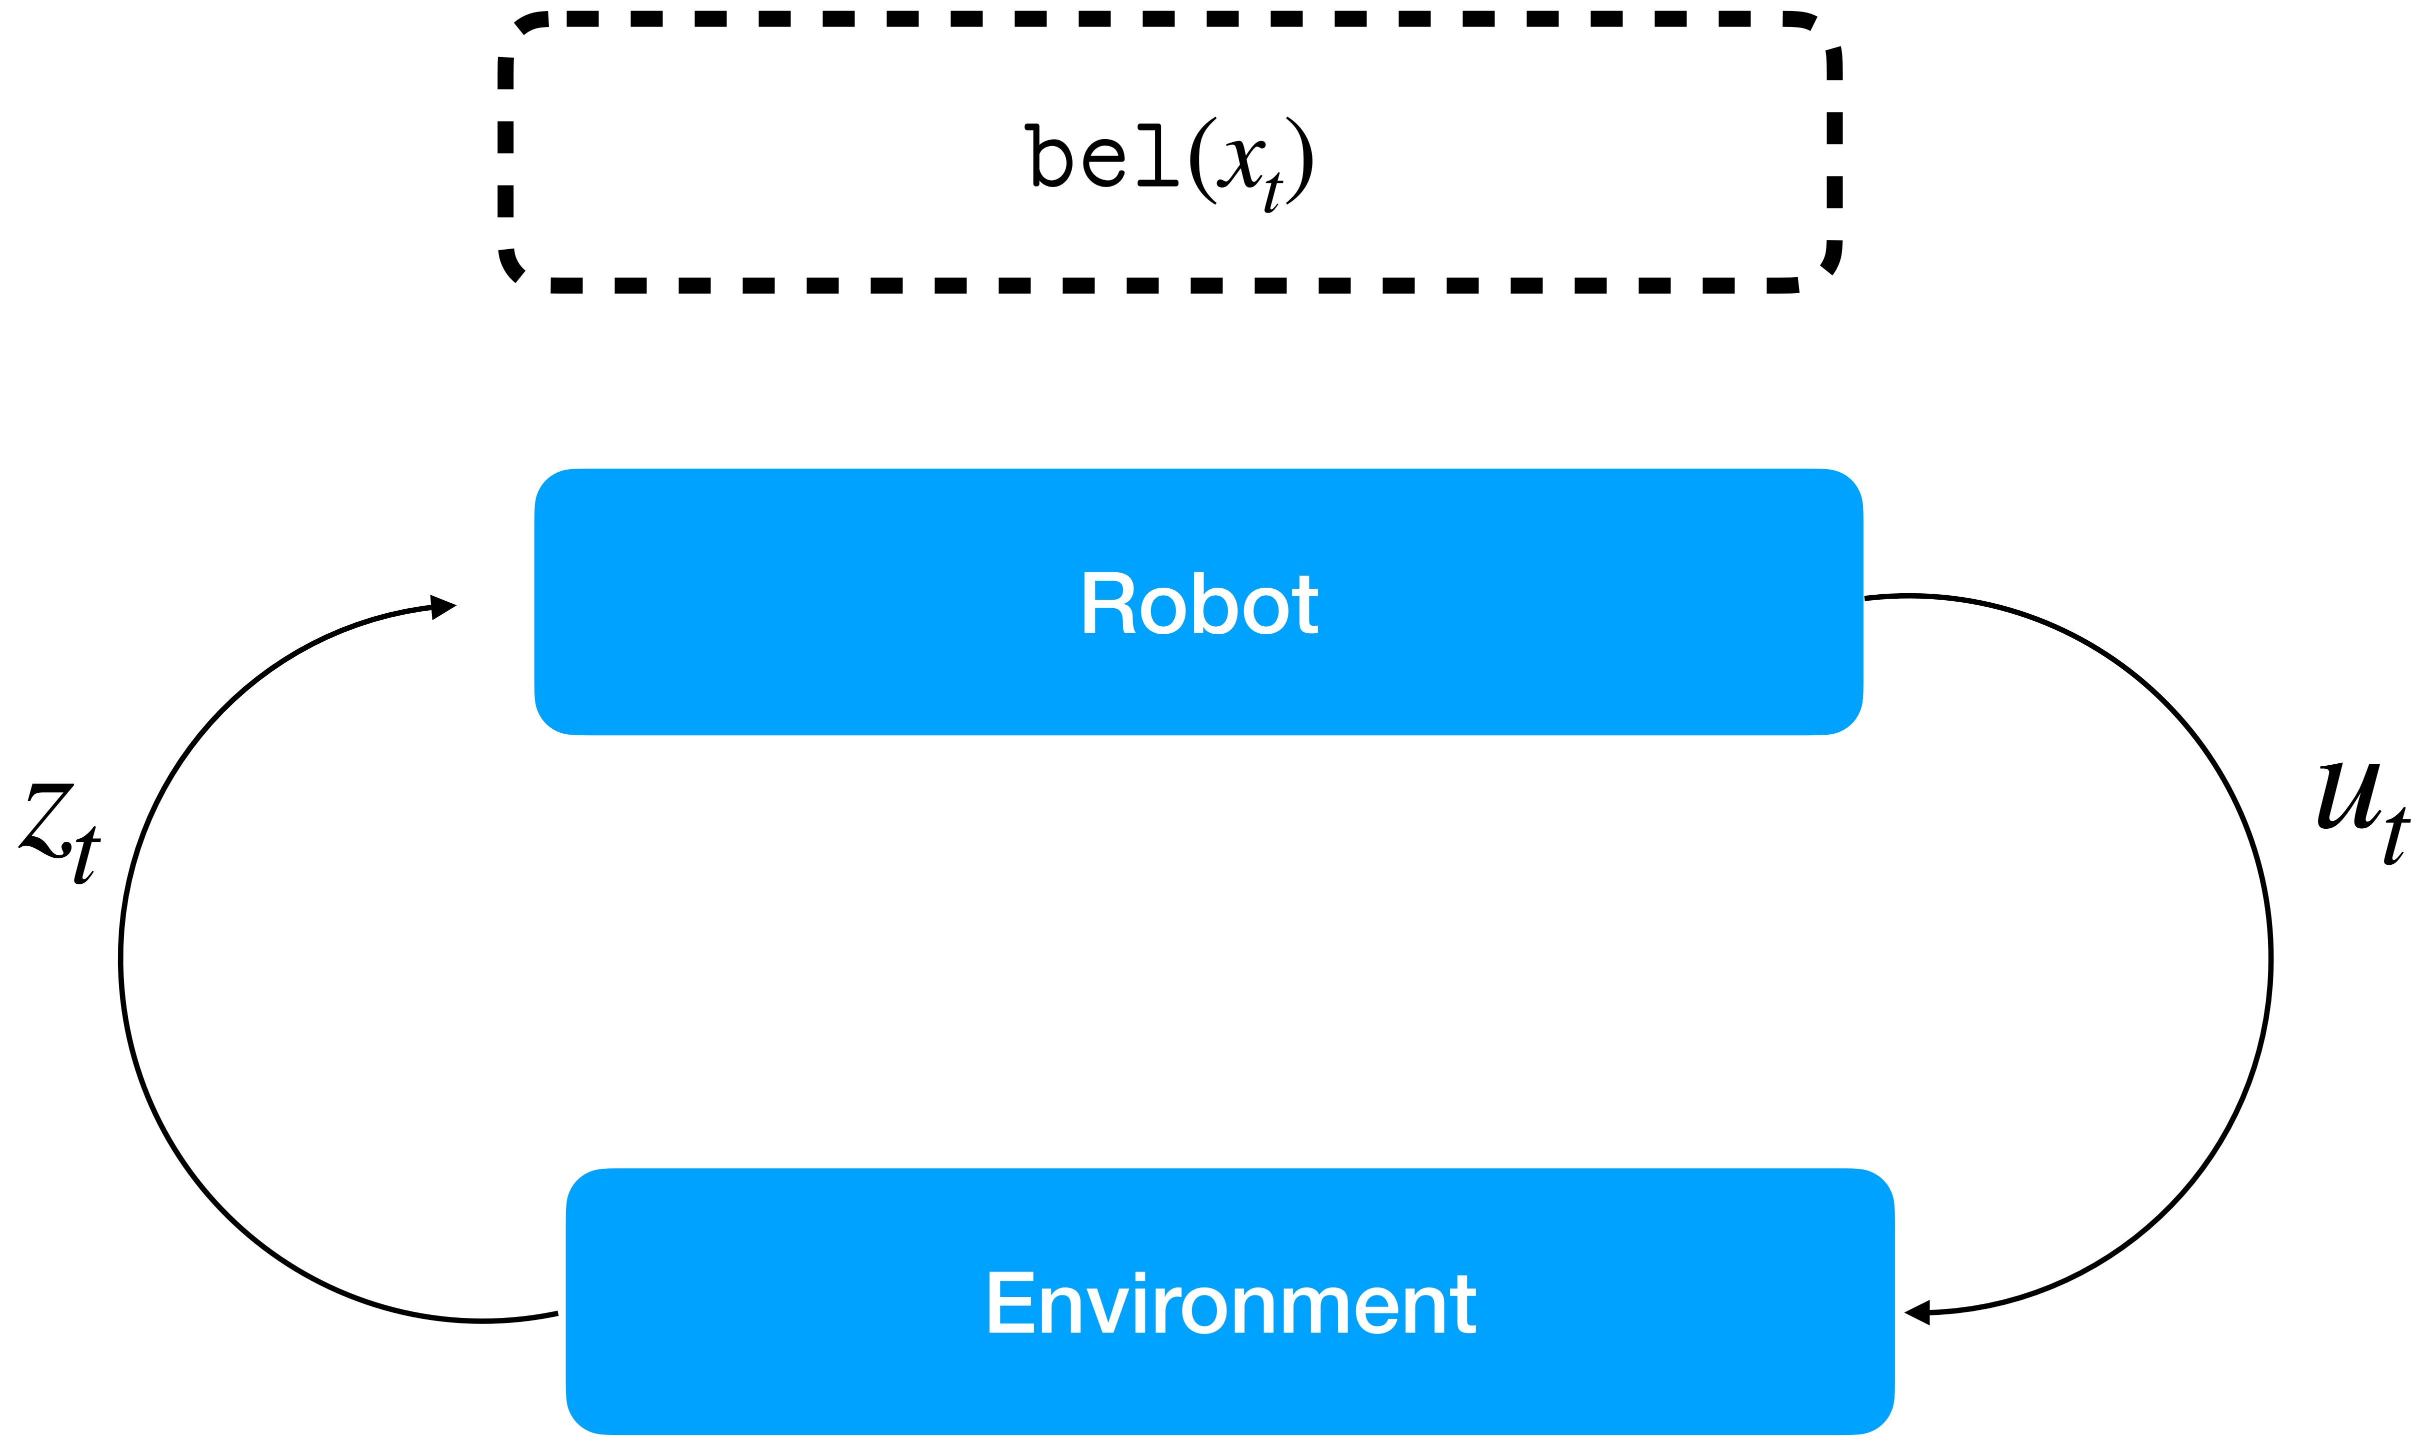

$$x_{t+1} \sim p(x_t, u_t)$$

$$z_{t+1} \sim g(x_{t+1})$$

Deterministic Interaction Model

$$u_t \leftarrow H(y_t)$$

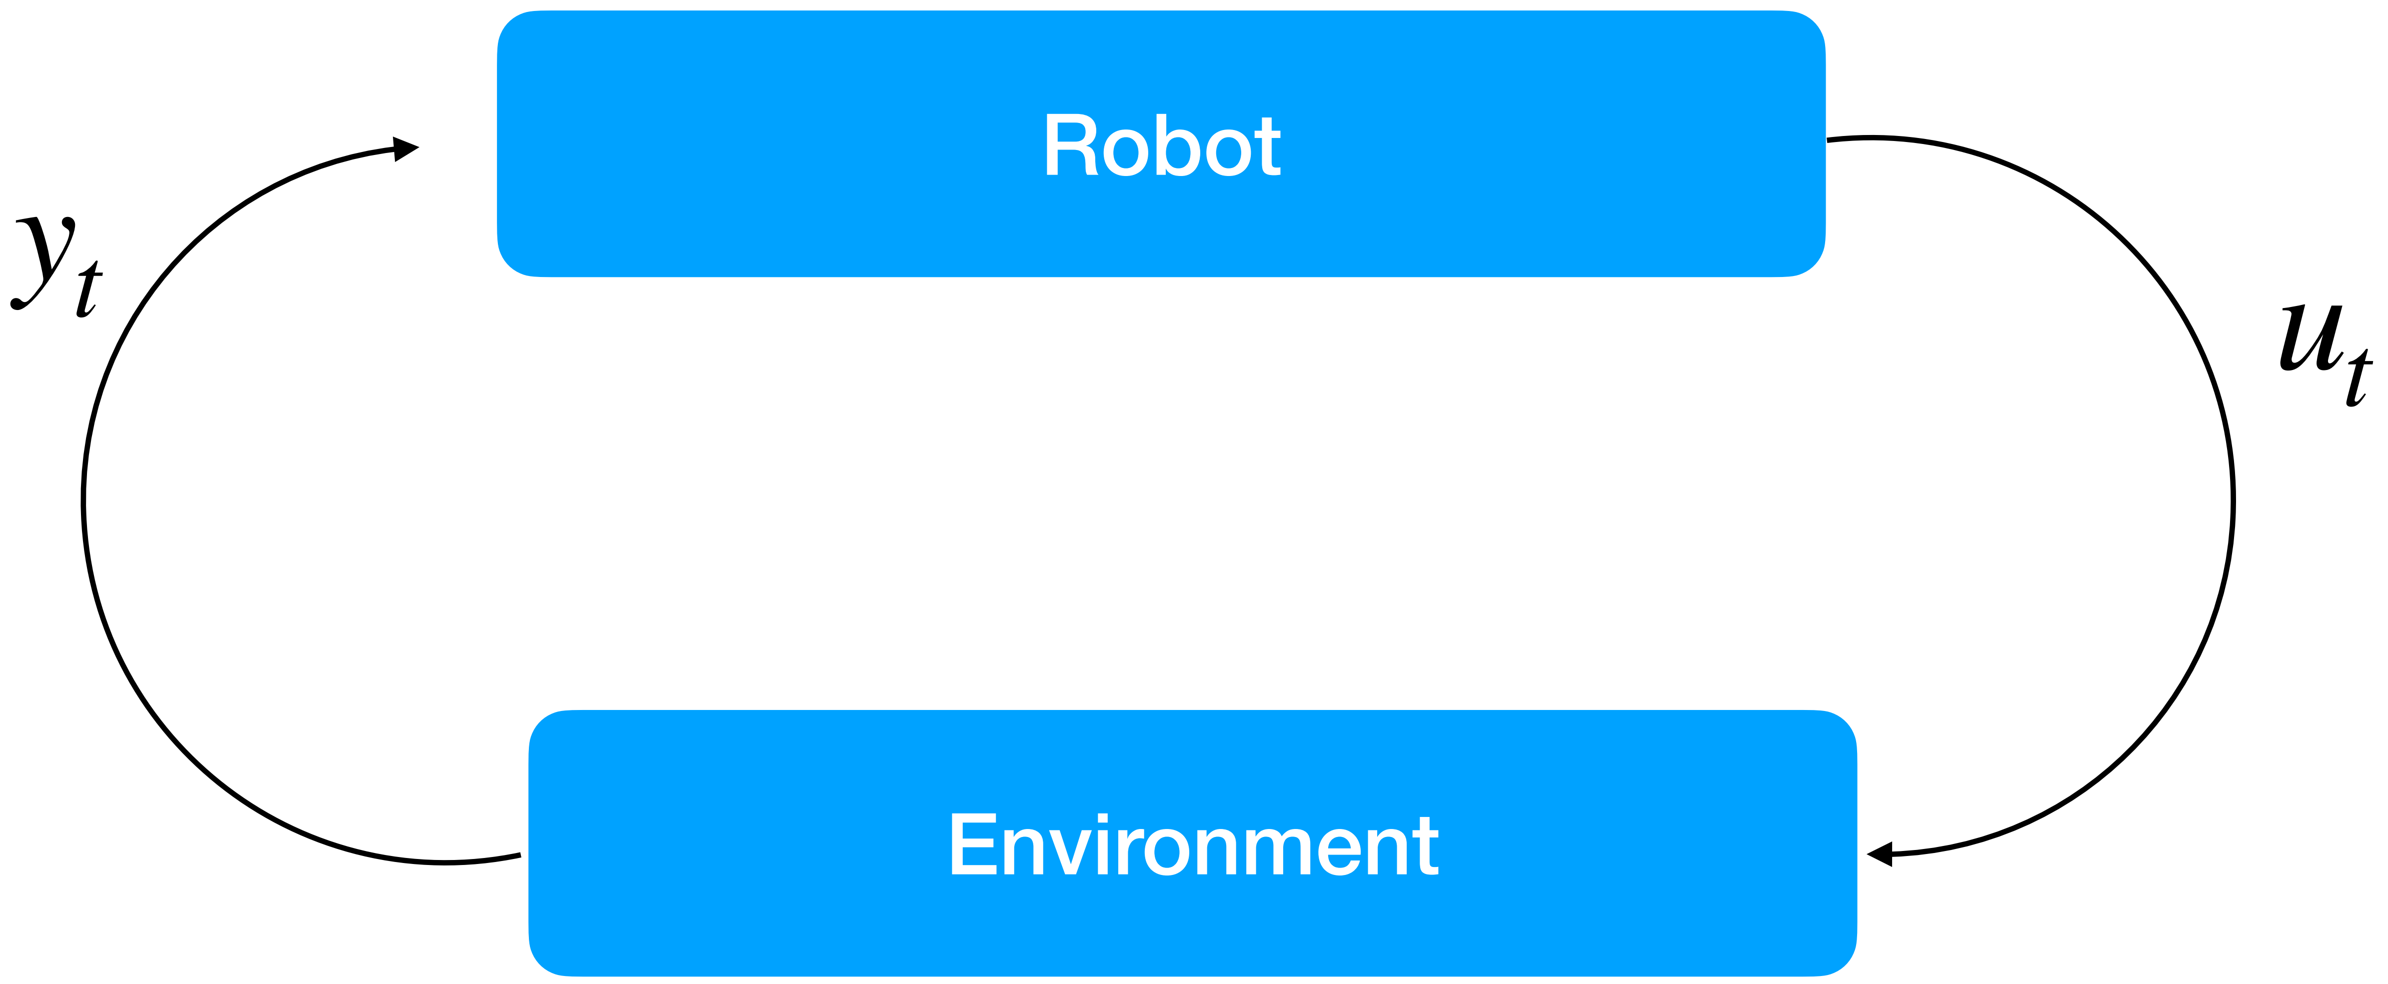

$$x_{t+1} \leftarrow F(x_t, u_t)$$

Or

$$\dot{x}_t \leftarrow F(x_t, u_t)$$
$$x_{t+1} \leftarrow x_t + \dot{x}\Delta t$$

$$y_{t+1} \leftarrow G(x_{t+1})$$

Summary

- Review of key probability concepts: distributions, joint, marginal, conditional, conditional independence, Bayes rule, Bayesian inference.
- Introduced robot-environment interaction model:
    - States, observations.
    - Markov assumption.
    - Interaction models.

Action Items

- Join Piazza and Gradescope.
- Complete the background survey: https://forms.gle/d8hmnQGWQc9SMVcN6
- [Optional but encouraged] Download Webots and complete a tutorial.
- Send a reading response by 12pm on Monday.