CS 3630!

Lecture 3:
A Trash Sorting Robot:
states and actions

Lecture 2 Recap

A Taxonomy of Robotics Topics

For each module in this class, we'll consider six distinct aspects of robotics:

1. **State:** How does the robot represent its world, and itself?
2. **Actions:** What can the robot do, and how to represent this?
3. **Sensors:** What information about the world can be ascertained via sensing, and how do we model this process?
4. **Perception:** How can we combine sensor data with contextual knowledge to understand the current state?
5. **Planning:** What actions should the robot execute to transform the state of the world into a desired goal state?
6. **Learning:** How can the robot improve its knowledge over time, using information that it acquires during operation?

Each chapter of the book includes six sections, corresponding to these topics.

Robots in the real world

For specific applications, these topics correspond to specific problems that robots must solve to operate effectively.

For example, a museum guide robot might need to solve the following problems:

- State: where is the robot, and where are the humans to be guided?
- Actions: move from room to room
- Sensors: cameras
- Perception: use computer vision to understand human intention, and to localize
- Planning: what path to take in order to guide humans to their desired exhibit
- Learning: which parts of the museum are crowded, and when to avoid these?

How do robots function in the world

When they are deployed in the world, most robots use the so-called ***Sense-Think-Act*** paradigm of operation.

This can be viewed as an overall control structure, in which state, actions, sensors, perception, planning, and learning play specific roles.

Example: Navigation in a Known Environment

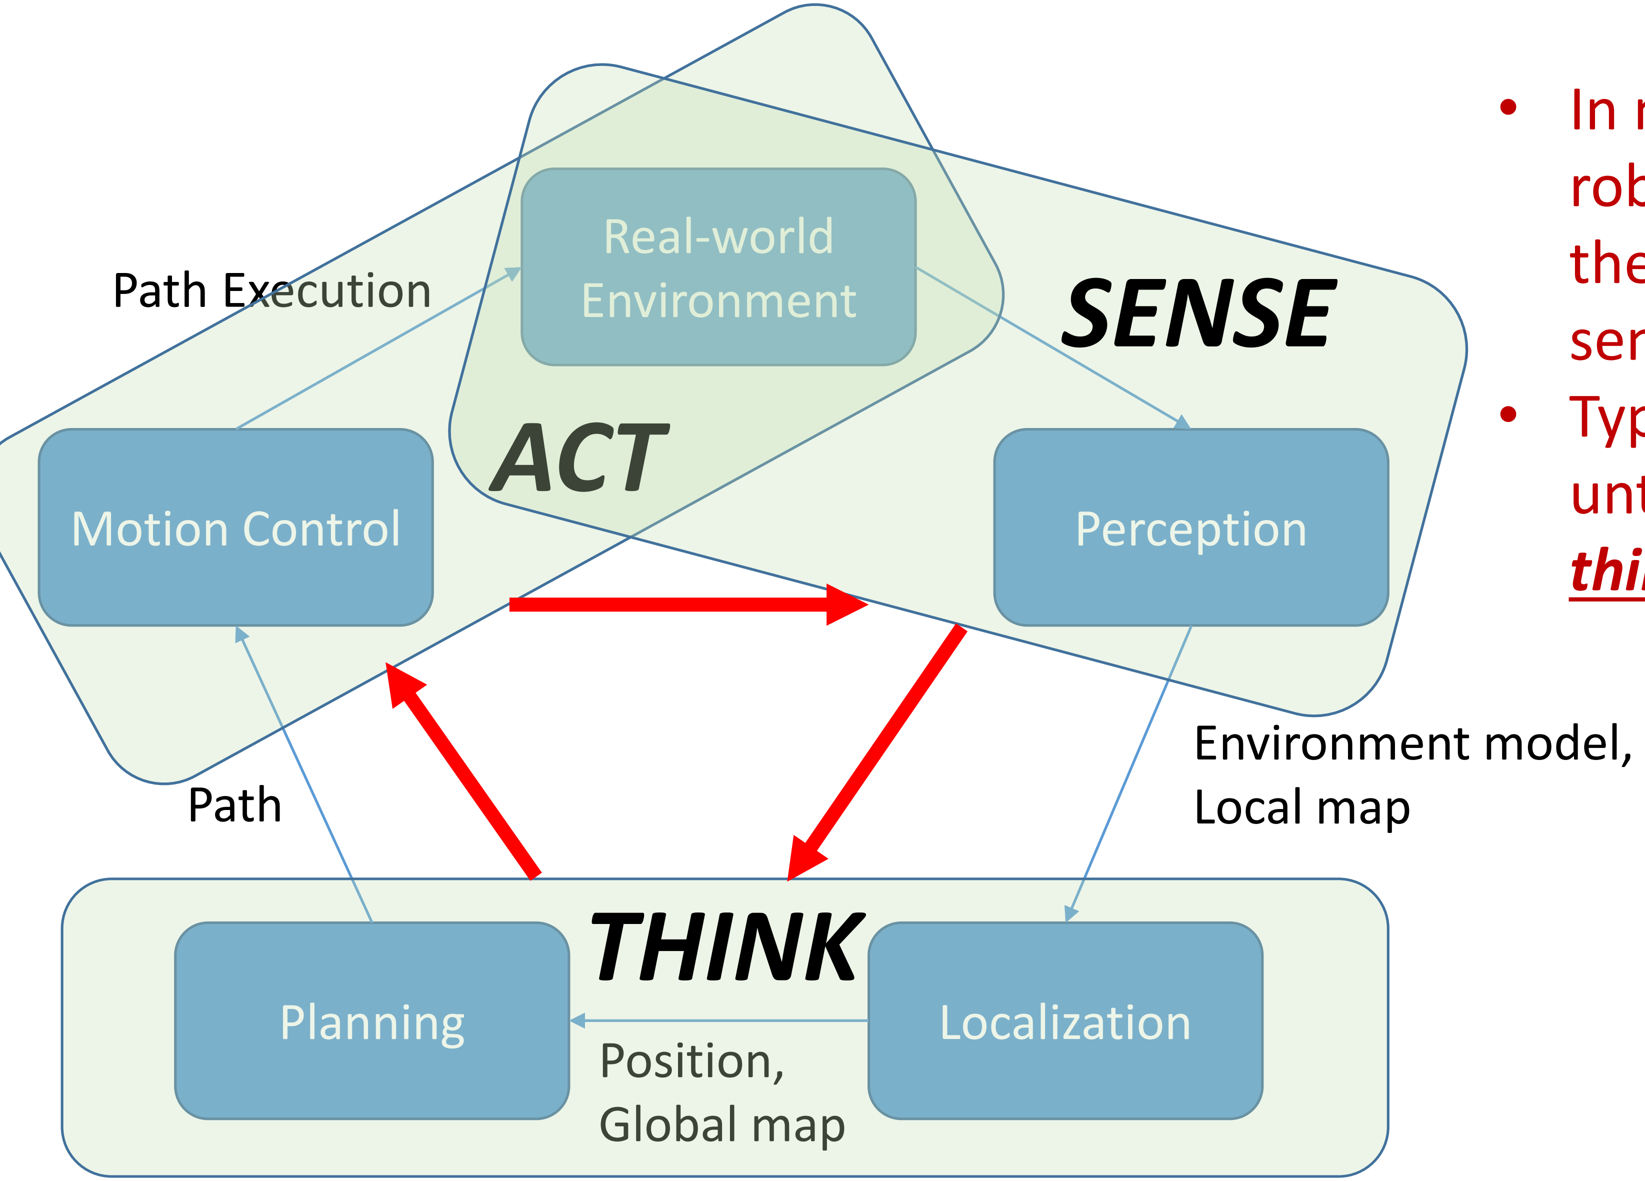

- In most robotics applications, the robot does not succeed to perform the task using a single episode of sense, think, act.
- Typically, these stages are repeated until the task is achieved: the ***sense, think, act loop.***

Sense, Think, Act at Different Time Scales

The time to complete one cycle of this loop depends on the task:

- Playing chess: minutes
- Hand-eye coordination: 30 Hz
- Force controlled robot: Order of KHz

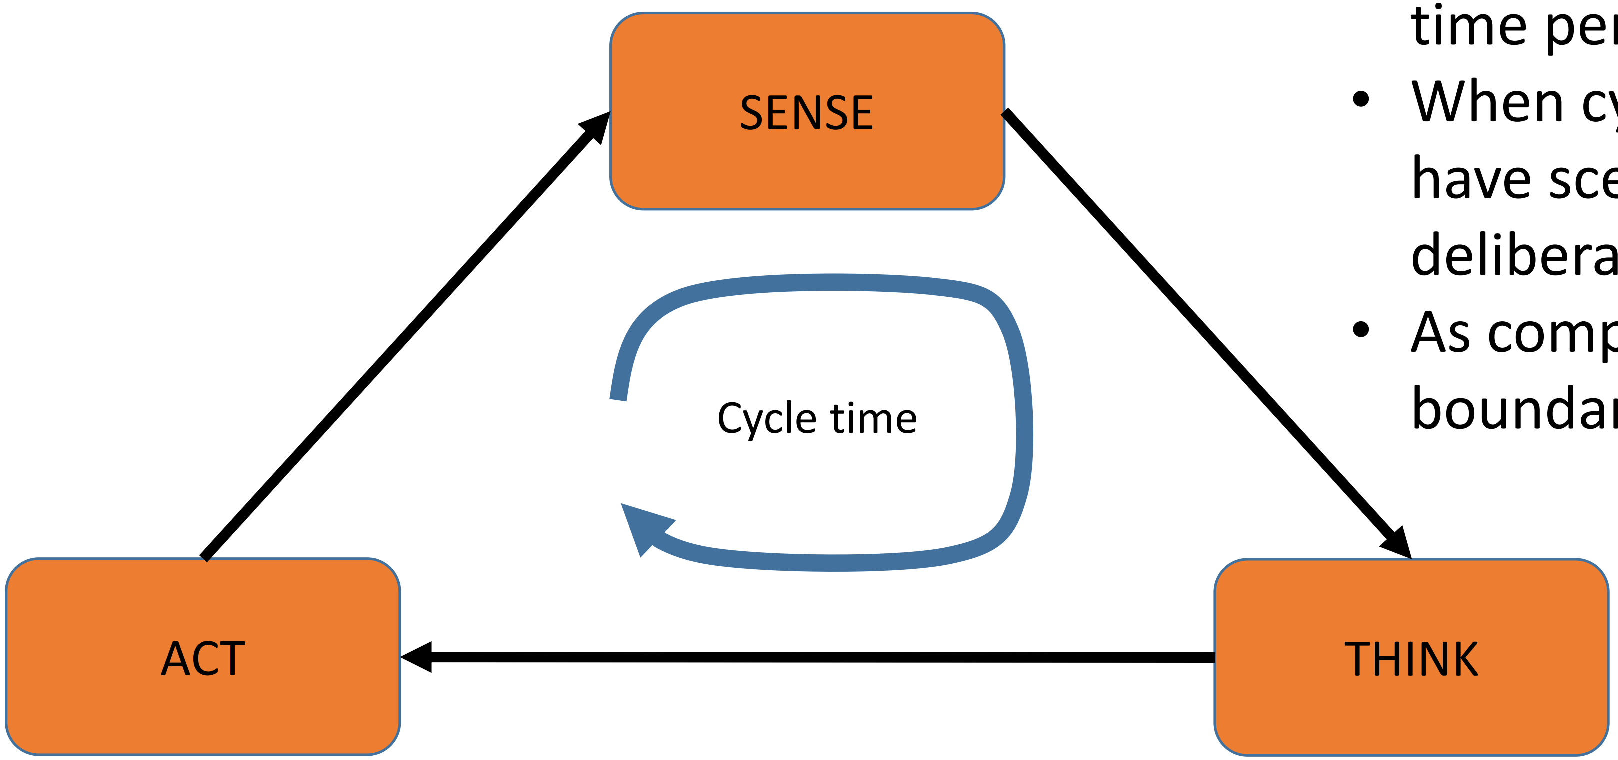

- When cycle time is very fast, we use tools from control theory, and model systems using differential equations (continuous time performance).
- When cycle time is very slow, we might have scene understanding and deliberative planning.
- As computers become faster, the boundary between these begins to blur.

A Trash Sorting Robot

Chapter 2

A Trash Sorting Robot

Our first example is a trash sorting robot.

Individual pieces of trash arrive to the robot's work cell on a conveyor belt.

The robot's task is to place each piece of trash in an appropriate bin:

- Glass
- Mixed paper
- Metal
- Nop

Sensors measure various characteristics of the trash, which are used to make inferences about the object type (perception).

We assume sensor uncertainty, but perfect execution of actions.

Over time, sensor models can be refined using machine learning methods.

Modeling the World State

For this problem, the only interesting aspect of the world state is the specific material composition of the item of trash that is currently in the robot’s work cell.

We consider five possibilities:

- Cardboard
- Paper
- Cans
- Scrap Metal
- Bottles

For this chapter, we assume that there are no other possibilities.

You should probably just memorize these now, because they’re going to be used a lot in this chapter.

Modeling Uncertainty

We assume that there is uncertainty in sensing, and therefore, it is not possible to know with certainty the world state.

We consider the state to be a random quantity, with five possible outcomes:

$$\Omega = \{\text{cardboard, paper, cans, scrap metal, bottle}\}$$

In probability theory,

- The set Ω is called the sample space.
- Each $\omega \in \Omega$ is called an outcome.
- A subset $A \subset \Omega$ is called an event.

Denote by $\mathfrak{B} = \{A | A \subset \Omega\}$ the set of all events.

Probability distributions map events to probabilities, $\boldsymbol{P}\colon \mathfrak{B} \rightarrow [\mathbf{0}, \mathbf{1}]$

Examples

Suppose the probabilities associated with the five outcomes are given as:

Category (ω)	$P(\{\omega\})$
Cardboard	0.20
Paper	0.30
Cans	0.25
Scrap Metal	0.20
Bottle	0.05

Compute the following:

- The probability that an item is a paper product: $\boldsymbol{P(\{A_1\})}$
- The probability that an item is a metal product: $\boldsymbol{P(\{A_2\})}$

Answers:

- $\boldsymbol{P(\{A_1\}) = P(\{cardboard\}) + P(\{paper\}) = 0.5}$
- $\boldsymbol{P(\{A_2\}) = P(\{cans\}) + P(\{scrap\ metal\}) = 0.45}$

Define three events

A_1 = {cardboard, paper}

A_2 = {cans, scrap metal}

Some properties of probability distributions

Three Axioms of Probability Theory:

1. For A $\subset \Omega$, P(A) ≥ 0
   - *There's no such thing as negative probability.*
2. P(Ω) =1
   - *The probability that something happened is 1.*
3. For $\mathrm{A_i}, \mathrm{A_j} \subset \Omega$, if $\mathrm{A_i} \cap \mathrm{A_j} = \emptyset$, then $P(A_i \cup A_j) = P(A_i) + P(A_j)$
   - *If two events are disjoint (aka mutually exclusive), then the probability that one of the two events occurred equals the sum of the probabilities for the two events.*
   - *The second and third axiom immediately imply that* P($\emptyset$) =0.

A handy relationship:

Since $\boldsymbol{A \cap \bar{A} = \emptyset}$ ***and*** $\boldsymbol{A \cup \bar{A} = \Omega}$***, we can conclude that***

$$\boldsymbol{P(A) + P(\bar{A}) = 1}$$

which implies

$$\boldsymbol{P(\bar{A}) = 1 - P(A)}$$

Proof:

1. $\boldsymbol{A \cap \bar{A} = \emptyset}$ **implies** $\mathbf{P}(\boldsymbol{A \cup \bar{A}}) = \mathbf{P}(\mathbf{A}) + \mathbf{P}\overline{(\boldsymbol{A})}$ **[Axiom 3]**
2. $\boldsymbol{A \cup \bar{A} = \Omega}$ **impies** $\mathbf{P}(\boldsymbol{A \cup \bar{A}}) = \boldsymbol{P(\Omega)} = \mathbf{1}$ **[Axiom 2]**
3. **Together, 1 and 2 imply** $\mathbf{P}(\mathbf{A}) + \mathbf{P}\overline{(\boldsymbol{A})} = \mathbf{1}$

Examples

Suppose the probabilities associated with the five outcomes are given as:

Category (ω)	$P(\{\omega\})$
Cardboard	0.20
Paper	0.30
Cans	0.25
Scrap Metal	0.20
Bottle	0.05

Compute the following:

- The probability that an item is a paper product: $\boldsymbol{P(\{A_1\})}$
- The probability that an item is a metal product: $\boldsymbol{P(\{A_2\})}$
- The probability that an item is not a paper product $\boldsymbol{P(\{\bar{A}_1\})}$

Answers:

- $\boldsymbol{P(\{A_1\}) = P(\{cardboard\}) + P(\{paper\}) = 0.5}$
- $\boldsymbol{P(\{A_2\}) = P(\{cans\}) + P(\{scrap\ metal\}) = 0.45}$
- $\boldsymbol{P(\{\bar{A}_1\}) = P(\Omega - \mathrm{A}_1) = \mathrm{P}(\{\mathrm{cans}, \mathrm{scrap\ metal}, \mathrm{bottle}\}) = 0.5}$

Define two events

A_1 = {cardboard, paper}

A_2 = {cans, scrap metal}

Prior Probability Distributions

What can we say about the probabilities of various outcome before we even invoke the robot's sensors?

- Our beliefs about the probabilities of various outcomes can be encoded in a ***prior distribution*** --- i.e., the a priori belief about the world.
- Priors can be estimated using data, or can be inferred using domain knowledge (e.g., a fair coin should land on heads 50% of the time).

In the book, we estimate prior probabilities using *observed data*:

- Cardboard occurs about 200 times for each 1000 item of trash.
- Paper occurs about 300 times for each 1000 item of trash.
- Cans occur about 250 times for each 1000 item of trash.
- Scrap Metal occurs about 200 times for each 1000 item of trash.
- Bottles occur about 50 times for each 1000 item of trash.

Is there any reason to believe that this approach should work in practice?

Borel's law of large numbers

- Let $A \subset \Omega$ be an event with probability $\mathrm{P}(A) = p$.
- Suppose we run our experiment n times, and we observe that event A occurs $N_n(A)$ times.
- Then, with probability one

$$\frac{N_n(A)}{n} \to p \text{ as } n \to \infty$$

➢ ***As the number of trials goes to infinity, the proportion of times that an event occurs approaches the probability of that event.***

➢ ***If we make enough observations, we can start to trust that we have good estimates of prior probabilities!***

Shamelessly taken from Wikipedia.

Machine Learning

In fact, we have just seen a first, simple example of machine learning:

1. ***Count the number of occurrences of each category.***
2. ***Use their relative proportions as an estimate of the prior probability distribution.***

We'll go a bit deeper later in this chapter.

Simulation by sampling

- Often useful to simulate robot systems. In our case, we might like to simulate the arrival of trash to our sorting system, such that it accurately reflects the prior distribution?
- How can we generate a sequence of samples, say $\omega_1, \omega_2, \ldots, \omega_n$, such that $\omega_i = cardboard$ for approximately 20% of the samples, $\omega_i = paper$ for approximately 30% of the samples, etc.?

Image from here

Random Variables

A ***random variable*** is a mapping from outcomes to real numbers, $X: \Omega \rightarrow \mathbb{R}$.

For example, we can map our categories to integers:

- Cardboard $\longrightarrow$ 0
- Paper $\longrightarrow$ 1
- Can $\longrightarrow$ 2
- Scrap Metal $\longrightarrow$ 3
- Bottle $\longrightarrow$ 4

- We typically use upper case letters, e.g., X, to denote a random variable, and lower-case letters, e.g., x_i, to denote the values taken by X.
- In our example, $X \in \{0,1,2,3,4\}$ indicates that X is a random variable that can take values from the set $\{0,1,2,3,4\}$.

Probability Mass Functions (pmf's)

- When a random variable takes its values from a finite (or possibly countably infinite) set, it is called a ***discrete random variable***.
- The probability distribution for a discrete random variable is typically defined as a ***probability mass function (pmf)***.
- For random variable X, the pmf is defined as

$$p_X(x) \triangleq P(X = x)$$

For our example,

- $p_X(0) = 0.20 \quad \textit{cardboard}$
- $p_X(1) = 0.30 \quad \textit{paper}$
- $p_X(2) = 0.25 \quad \textit{can}$
- $p_X(3) = 0.20 \quad \textit{scrap metal}$
- $p_X(4) = 0.05 \quad \textit{bottle}$

As we will soon see, random variables can be very useful for outcomes that are naturally associated to real numbers, e.g., roll of a die, weight of a person, or, in our case, cost of applying an action.

Using pmf's

Even for this example, where categories don't naturally have numerical semantics, we can use the pmf to answer interesting questions.

For example, what is the probability that an object is a paper product?

- Paper products correspond to paper and cardboard, $X \in \{0,1\}$:

$$P(X \in \{0,1\}) = p_X(0) + p_X(1) = 0.5$$

Alternatively, we could write:

$$P(X \in \{0,1\}) = P(X \leq 1)$$

This form, $P(X \leq \alpha)$ turns out to be very useful.

Cumulative Distribution Function

The Cumulative Distribution Function (CDF) is defined as

$$F_X(\alpha) = P(X \leq \alpha) = \sum_{x_i \leq \alpha} p_X(x_i)$$

If we order the x_i's, such that $x_0 < x_2 \ldots < x_n$ we can write this as:

$$F_X(\alpha) = P(X \leq \alpha) = \sum_{i=0}^{k-1} p_X(x_i)$$

when we choose k such that $x_{k-1} \leq \alpha < x_k$.

CDF for our trash categories

It is straightforward to compute the CDF for the r.v. associated to various trash categories:

$$F_X(\alpha) = P(X \le \alpha) = \sum_{i=0}^{k-1} p_X(x_i)$$

r.v. $\boldsymbol{x}$	$\boldsymbol{p_X(x)}$
0	0.20
1	0.30
2	0.25
3	0.20
4	0.05

Category ($\boldsymbol{\omega}$)	**r.v. $\boldsymbol{x}$**	$\boldsymbol{F_X(\alpha)}$
Cardboard	0	$P(X \le 0) = 0.20$, $\boldsymbol{\alpha} = 0$
Paper	1	$P(X \le 1) = 0.50$, $\boldsymbol{\alpha} = 1$
Cans	2	$P(X \le 2) = 0.75$, $\boldsymbol{\alpha} = 2$
Scrap Metal	3	$P(X \le 3) = 0.90$, $\boldsymbol{\alpha} = 3$
Bottle	4	$P(X \le 4) = 1.00$, $\boldsymbol{\alpha} = 4$

Simulation by sampling

- So, how can we generate a sequence of samples, say $\omega_1, \omega_2, \ldots, \omega_n$, such that $\omega_i = cardboard$ for approximately 20% of the samples, $\omega_i = paper$ for approximately 30% of the samples, etc.?
- Sadly, most programming languages do not include library functions to sample from arbitrary probability distributions.
- Happily, the is almost always a random number generator that generates a random sample from the unit interval, $x \sim U(0,1)$.
- The notation $x \sim U(0,1)$ indicates that x is a number chosen at random from the interval $[0,1]$, and that all possible outcomes are equally likely.

➢ Let's see how to use this…

Simulation by sampling

Suppose we generate the samples $s_1 = 0.97$ and $s_2 = 0.29$

r.v. x	$p_X(x)$	$F_X(\alpha)$, $\alpha = 0, 1, 2, 3, 4$
0	0.20	0.20
1	0.30	0.50
2	0.25	0.75
3	0.20	0.95
4	0.05	1.00

- Note that $0.95 < (s\_1 = 0.97) \leq 1.$
- The probability that this occurs is exactly 0.05, since the probability of $x \in [a, b] = (b - a)$ for the uniform distribution on $[0,1]$.
- $P(bottle) = 0.05$ …. Return category bottle.

- Similarly, $0.20 < (s_2 = 0.29) \leq 0.50$
- The probability that this occurs is exactly 0.30.
- $P(paper) = 0.30$ … Return category paper.

We can generalize this to develop an algorithm that draws a sample from an arbitrary distribution.

1. Generate a sample $x \sim U(0,1)$.
2. Determine k such that $F_X(x_{k-1}) < x \leq F_X(x_k)$.
3. Select category ω_k

Use the book!

Now is the time to visit the online book, explore the concepts, and play with the code to ensure that you understand what we have just discussed.

- *Try different prior distributions, and build the corresponding CDF.*
- *Be sure that your hand calculations match the results from the code.*
- *Generate many samples. Compare the sample distribution (i.e., the proportion of occurrences of each category) to the true prior.*
- *Increase the number of samples. You should notice that the sampling distribution becomes increasingly similar to the true prior as you increase the number of samples.*

Actions

For this problem, the robot either places an item of trash into one of three bins, or lets the item pass through the work cell.

This gives four possible actions:

- a_1: Glass Bin
- a_2: Metal Bin
- a_3: Paper Bin
- a_4: Nop (let object pass through the workcell)

For this chapter, we assume that actions are executed without error, every time.

However, since we don't know with certainty the category for an item of trash in the work cell, the efficacy of an action is also uncertain.

Assessing Risk

- Because there is uncertainty in the category of a piece of trash, the robot risks making mistakes when choosing actions.
- Different mistakes have different costs.
  - Placing metal in the paper bin might seriously damage paper processing equipment.
  - Placing paper in the metal bin is unlikely to cause much harm.

COST	**cardboard**	**paper**	**can**	**scrap metal**	**bottle**
glass bin	2	2	4	6	0
metal bin	1	1	0	0	2
paper bin	0	0	5	10	3
nop	1	1	1	1	1

To account for these variations, we can define a table of costs for applying each action (rows) to each category (columns).

Assessing Risk

- Because there is uncertainty in the category of a piece of trash, the robot risks making mistakes when choosing actions.
- Different mistakes have different costs.
  - Placing metal in the paper bin might seriously damage paper processing equipment.
  - Placing paper in the metal bin is unlikely to cause much harm.

COST	cardboard	paper	can	scrap metal	bottle
glass bin	2	2	4	6	0
metal bin	1	1	0	0	2
paper bin	0	0	5	10	3
nop	1	1	1	1	1

We assign zero costs to correct actions.

Assessing Risk

- Because there is uncertainty in the category of a piece of trash, the robot risks making mistakes when choosing actions.
- Different mistakes have different costs.
  - Placing metal in the paper bin could cause serious damage of paper processing equipment.
  - Placing paper in the metal bin is unlikely to cause much harm.

COST	cardboard	paper	can	scrap metal	bottle
glass bin	2	2	4	6	0
metal bin	1	1	0	0	2
paper bin	0	0	5	10	3
nop	1	1	1	1	1

We assign zero costs to correct actions.

The cost of Nop is due to the need for human labor to sort the item of trash.

Cost as a Random Variable

Since we only have probabilistic knowledge of an item's category, we can regard the cost of executing an action as a discrete random variable.

Consider action a_3, *place the item in the mixed paper bin*.

- Let X be the r.v. that denotes the cost of applying action a_3.
- From the table of costs, we see that $X \in \{0,5,10,3\}$, since these are the only possible costs for this action.

➢ What can we say about the probability distribution for X?

Computing pmf's

To compute the pmf, recall that the random variable is a mapping from outcomes to real numbers.

There are five possible outcomes. The object must be from one of five categories, each of which has a cost.

- Compute $p_X(x)$ for each x.

Category	P(C)	Cost
cardboard	0.20	0
paper	0.30	0
can	0.25	5
scrap metal	0.20	10
bottle	0.05	3

- $X = 0$ for cardboard and paper.
- $P(\{cardboard, paper\}) = P(\{cardboard\}) + P(\{paper\}) = 0.5$
- $p_X(0) = 0.5$

- $X = 5$ for can.
- $P(\{can\}) = 0.25$
- $p_X(5) = 0.25$

- $X = 3$ for bottle.
- $P(\{bottle\}) = 0.05$
- $p_X(3) = 0.05$

- $X = 10$ for scrap metal.
- $P(\{scrap\ metal\}) = 0.20$
- $p_X(10) = 0.20$

Expectation

- Probabilities tell us something about a single outcome, but this isn't really very useful. Gamblers who make one-time bets based on probabilities can lose a lot of money.
- Most robots operate for prolonged periods of time.
- The notion of average cost over many trials seems like a useful thing to know.

➢ ***This is exactly the concept of expectation in probability theory.***

Expectation

If a r.v. X takes its values from a finite set, $X \in \{x_1, \ldots, x_n\}$, the ***expected value*** of X, denoted $\boldsymbol{E[X]}$, is defined by:

$$E[X] = \sum_{i=1}^{n} x_i \, p_X(x_i)$$

- Expectation is a property of a probability distribution
- **$E[X]$ is *not* the value you should expect to see for any specific outcome!!**

Examples

Let $X \in \{x_1, \ldots, x_n\}$ be a discrete r.v. that corresponds to the number of dots shown on a fair die.

- $X \in \{1,2,3,4,5,6\}$ and $p_X(x_i) = \frac{1}{6}$ for all i

➢ Compute $E[X]$.

$$E[X] = \sum_{i=1}^{n} x_i\, p_X(x_i) = \sum_{i=1}^{6} \frac{1}{6} i = \frac{1}{6} + \frac{2}{6} + \frac{3}{6} + \frac{4}{6} + \frac{5}{6} + \frac{6}{6} = \frac{21}{6} = 3.5$$

Trash Sorting…

We can now easily evaluate the expected cost for each action under the prior probability distribution.

COST	Card board	paper	can	scrap metal	bottle	Expected Cost	
glass bin	2	2	4	6	0	3.2	$2 \times 0.5 + 4 \times 0.25 + 6 \times 0.2 = 3.2$
metal bin	1	1	0	0	2	0.6	$1 \times 0.5 + 2 \times 0.05 = 0.6$
paper bin	0	0	5	10	3	3.4	$5 \times 0.25 + 10 \times 0.2 + 3 \times 0.05 = 3.4$
nop	1	1	1	1	1	1.0	$1 \times 0.5 + 1 \times 0.25 + 1 \times 0.2 + 1 \times 0.05 = 1.0$
$P(\omega)$	0.20	0.30	0.25	0.20	0.05		

Simulation by sampling

Earlier, we simulated our trash sorting system using a sampling algorithm. Let's apply those ideas here.

1. Generate N samples from the prior distribution on categories.
2. Compute the cost c_i for each sample for action a_k.
3. Compute the average cost as:

$$\overline{cost_k} = \frac{1}{N} \sum_{i=1}^{N} c_i$$

4. Compare $\overline{cost_k}$ to $E[X]$ for action a_k (where X is the r.v. for cost).

Probability *vs* Statistics

- **Probability theory** is the study of a certain class of mathematical functions (probability distributions).

- A **statistic** is any function of data (including the identity function), and statistics is the study of such functions.

$$E[X] = \sum_{i=1}^{n} x_i \, p_X(x_i)$$

$E[X]$ is a property of $p_X(x_i)$

➢ ***Probability Theory***

$$\overline{cost_k} = \frac{1}{N} \sum_{i=1}^{N} c_i$$

$\overline{cost_k}$ is a function of data, c_i

➢ ***Statistics***

Probability Theory *and* Statistics

If it happens that certain probability distributions do a good job of describing how the world behaves, then probability theory can provide a rigorous basis for a system of inference about data.

The Weak Law of Large Numbers:

Consider a data set drawn from probability distribution p_X, with expected value $E[X] = \mu$.
For any $\epsilon > 0$, if $\bar{x}_N$ denotes the average of a data set of size N , then

$$\lim_{n\to\infty} P(|\bar{x}_N - \mu| < \epsilon) = 1$$

As the size of the data set increases, with probability one the average is arbitrarily close to the mean.

Probability Theory and Statistics

The connections between probability theory and statistics are often formalized by theorems that express variations on a simple concept:

As the size of a data set becomes large, the statistics of that data set will become increasingly good approximations for various properties of the underlying probability distribution from which the data set was generated.

- *This is one of the reasons simulation by sampling works.*
- *These theorems are important for statistical inference, machine learning, and many other problems that involve data drawn from stochastic systems.*

Next Lecture: Sensing and Perception

- Conditional probability:
  - How do sensor observations affect our beliefs about the world?
  - A key tool for data-based inference

- Continuous random variables:
  - Unlike our five categories of trash, some things are best described along a continuum.
  - Things like weight, distance are described using continuous measurements.
  - Gaussian Distributions

- Maximum likelihood inference
  - Making decisions using conditional probabilities
  - Combining information from multiple sensors