Probability Background and Notation Introduction for Robot State Estimation

Sarah Sebo

April 2021

1 Probability Background

1.1 Random Variables

A random variable is "described informally as a variable whose values depend on outcomes of a random phenomenon"[1]. In this course, we will represent robot states, measurements, and controls as random variables.

For example X could be a random variable representing the outcome of drawing a marble out of a bag that contains 3 blue marbles and 2 red marbles. x then represents a specific outcome (e.g., drawing a blue marble out of the bag). If X has a discrete number of possible outcomes, we can represent the probability of a specific marble draw outcome as

$$p(X = x)$$

In our specific example of a bag containing 3 blue marbles and 2 red marbles, $p(X = \text{blue marble}) = 0.6$ and $p(X = \text{blue marble}) = 0.4$. It is also must be true that the probability of all possible outcomes sum to 1:

$$\sum_x p(X = x) = 1$$

Random variables we consider in this course (e.g., the location of the robot) may also change over time. In order to capture the value of a random variable a specific point in time, we use the notation X_t. If we want to represent the random variable X_t at $t = 1$, we would write that as $X_{t=1}$ or simply X_1.

1.2 Joint Probability Distributions

If we have two random variables, X and Y, we may be interested in their joint probability distribution $p(X = x \text{ and } Y = y)$, another way to write that is $p(X = x, Y = y)$. If X and Y are independent, then it is true that

[1] https://en.wikipedia.org/wiki/Random_variable

$$p(X = x, Y = y) = p(X = x) \cdot p(Y = y)$$

For example, if X represents the result of flipping a first fair coin and Y represents the result of flipping a second fair coin, the probability of both coins getting 'heads' is:

$$\begin{aligned} p(X = \text{heads}, Y = \text{heads}) &= p(X = \text{heads}) \cdot p(Y = \text{heads}) \\ &= 0.5 \cdot 0.5 \\ &= 0.25 \end{aligned}$$

1.3 Conditional Probability

It may be the case that the probability of a particular random variable may be informed by the outcome of another random variable. For example, the probability that I am hungry ($p(X = \text{hungry})$) will likely be conditioned upon whether or not I've already eaten lunch ($p(Y = \text{eaten lunch})$). The conditional probability would be represented as

$$p(\text{hungry} \mid \text{eaten lunch}) = p(X = \text{hungry} \mid Y = \text{eaten lunch})$$

Additionally, the conditional probability can be represented as

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

Following from the definition of conditional probability, the following represents the theorem of total probability:

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

Additionally, Bayes' theorem relates the conditional probability $p(x|y)$ to it's reverse, $p(y|x)$:

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

2 Robotics Notation for this Course

We will represent **the robot's state** with the random variable X. A specific robot state is represented by x, a specific state at a particular point in time is represented by x_t. Possible robot states may include:

- The robot's location or pose
- The robot's linear & angular velocities

- The robot's joint angles (think for a robot arm)

We will represent **sensor measurements** with the random variable Z, where a specific sensor measurement is represented by z or z_t. Robot sensors may include:

- Laser range finders (LiDAR)
- RGB cameras
- Sonar sensors
- Infared sensors
- Temperature sensors

Finally, we will represent **control information** with the random variable U, where a specific set of robot controls is represented by u or u_t. Robot controls usually involve commands to motors (e.g., controlling joint angle positions for a robot arm, controlling wheel motor velocities).

It's helpful to think about robot behavior in the following steps:

1. The robot starts out in state x_t
2. The robot receives a measurement from its sensor(s) z_t
3. After receiving that sensor measurement, the robot decides to take action u_t
4. After taking action u_t, the robot's state updates to x_{t+1}

3 Belief Distributions

One essential concept in this course is probabilistic *belief.* Belief in this context represents the robot's knowledge of its state. For many reasons, a robot might be uncertain about its state, e.g., sensors might be noisy, sometimes robot control actions may not be executed as the robot intends (e.g., Turtlebot3 'drift' in Gazebo).

We will represent belief as a probability distribution across all of the possible states. We can inform our calculation of belief by considering the robot's past measurements ($z_{1:t}$) and actions ($u_{1:t}$) from time $t1$ to t within its environment:

$$bel(x_t) = p(x_t|z_{1:t}, u_{1:t})$$

For an example, let's say that a robot lives within the following grid world with four horizontal grid cells representing the possible states: A, B, C, D:

A	B	C	D

If the robot has no idea where it is at time t, then the robot's belief of which state it is in is evenly distributed across the 4 possible states:

$$bel(x_t = A) = 0.25$$
$$bel(x_t = B) = 0.25$$
$$bel(x_t = C) = 0.25$$
$$bel(x_t = D) = 0.25$$

Now, let's say that a robot receives a measurement z_t indicating that the robot is likely at one of the ends of the grid world. The robot may update its belief so that it looks something like:

$$bel(x_t = A) = 0.40$$
$$bel(x_t = B) = 0.10$$
$$bel(x_t = C) = 0.10$$
$$bel(x_t = D) = 0.40$$

This indicates that the robot has higher belief that it is in either state A or D. There is still a small probability that the robot is in state B or C, but that is less likely given the recent sensor measurement.

The belief can also be updated based on the movement of the robot. Let's imagine that the robot were to move 1 grid cell to the right (and let's assume that if the robot moved to the right in cell D, it would end up in cell A). After moving 1 grid cell to the right, the belief would update and might look something like:

$$bel(x_{t+1} = A) = 0.38$$
$$bel(x_{t+1} = B) = 0.38$$
$$bel(x_{t+1} = C) = 0.12$$
$$bel(x_{t+1} = D) = 0.12$$

We see here that the robot has the highest belief that it is in either state A or B. Additionally, The probabilities are not the same as they were in the last step to account for the fact that there is a small chance that the robot didn't actually move as it intended to.