

SDM5008 Advanced Control for Robotics

**Lecture Note 9: Probability Review for
Reinforcement Learning**

**Prof. Wei Zhang
Southern University of Science and Technology**

zhangw3@sustech.edu.cn
<https://www.wzhanglab.site/>

Outline

- **Probability and Conditional Probability**
- Random Variables and Random Vectors
- Jointly Distributed Random Vectors and Conditional Expectation

What is probability?

⇒ modeling uncertainty

world is deterministic!
uncertainty is primarily

- A formal way to quantify the uncertainty of 'our' knowledge about the physical world

there is no right or wrong probability

due to
(lack of
information)

- Formalism: Probability Space (Ω, \mathcal{F}, P)
 - Ω : **sampling space**: a set of all possible outcomes (maybe infinite)
 - \mathcal{F} : **event space**: collection of events of interest (event is a subset of Ω)
 - $P: \mathcal{F} \rightarrow [0,1]$ probability measure: assign event in \mathcal{F} to a real number between 0 and 1

eg. toss a coin · $\Omega =$

Axioms of probability:

- $P(A) \geq 0$
- $P(\Omega) = 1$
- $A \cap B = \emptyset \Rightarrow P(A \cup B) = P(A) + P(B)$
- **Important consequences:**
 - $P(\emptyset) = 0$
 - Law of total probability: $P(B) = \sum_i^n P(B \cap A_i)$, for any partitions $\{A_i\}$ of Ω
 - Recall a collection of sets A_1, \dots, A_n is called a partition of Ω if
 - $A_i \cap A_j = \emptyset$, for all $i \neq j$ (mutually exclusive)
 - $A_1 \cup A_2 \cdots \cup A_n = \Omega$

Conditional probability

- Probability of event A happens given that event B has already occurred

$$\bullet P(A|B) = \frac{P(A \cap B)}{P(B)}$$

- We assume $P(B) > 0$ in the above definition
- **What does it mean?**
 - Conditional probability is a probability: $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{P})$
 - **“Conditional” means, $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{P})$ is derived from an original probability space (Ω, \mathcal{F}, P) given some event has occurred**
 - After B occurred we are uncertain only about the outcomes inside B

- Bayes rule: relate $P(A | B)$ to $P(B | A)$

$$P(A | B) = \frac{P(B|A)P(A)}{P(B)}$$

- Events A and B are called (statistically) independent if
 - $P(A|B) = P(A)$
 - Or equivalently: $P(A \cap B) = P(A)P(B)$

- **Example of conditional probability:** A bowl contains 10 chips of equal size: 5 red, 3 white, and 2 blue. We draw a chip at random and define the event:

A = the draw of a red or a blue chip

Suppose you are told the chip drawn is not blue, what is the new probability of A

Outline

- Probability and Conditional Probability
- **Random Variables and Random Vectors**
- Jointly Distributed Random Vectors and Conditional Expectation

How to specify probability measure

- Discrete random variable: probability mass function (pmf)
e.g. toss a coin or die

- Continuous random variable: probability density function (pdf)
e.g. temperature density

How to specify probability measure

- Random vector: scalar random variables listed according to certain order

- n-dimensional random vector: $X = \begin{bmatrix} X_1 \\ X_2 \\ \vdots \\ X_n \end{bmatrix}$

- Notation: We typically use capital to denote random variables (vectors) and lower case letter to denote specific values the random variable takes

- density function: $f(x), x \in \mathbb{R}^n$

- probability evaluation: $P(X \in A) = \int_A f(x)dx$

Expectation of a random vector $X \in R^n$:

Continuous random vector: $E(X) = \int_{R^n} \mathbf{x} f(\mathbf{x}) d\mathbf{x}$

Discrete random vector: $E(X) = \sum_{\mathbf{x}} \mathbf{x} \cdot \text{Prob}(X = \mathbf{x})$

- Expectation: $E(X) = \begin{bmatrix} E(X_1) \\ E(X_2) \\ \vdots \\ E(X_n) \end{bmatrix}$

- Examples: Let $X \in R^2$ be discrete random variable with $\text{Prob}\left(X = \begin{bmatrix} 0 \\ 1 \end{bmatrix}\right) = \frac{1}{2}$, $\text{Prob}\left(X = \begin{bmatrix} 1 \\ 2 \end{bmatrix}\right) = \frac{1}{3}$, $\text{Prob}\left(X = \begin{bmatrix} -1 \\ 1 \end{bmatrix}\right) = \frac{1}{6}$. Compute $E(X)$

Linearity of Expectation:

- Expectation of AX with deterministic constant $A \in R^{m \times n}$ matrix:

$$E(AX) = AE(X)$$

- More generally, $E(AX + BY) = AE(X) + BE(Y)$

- Example: Suppose $X \in R^2, Y \in R^3$, with $E(X) = \begin{bmatrix} 0.5 \\ 0.25 \end{bmatrix}$, $E(Y) = \begin{bmatrix} 0.1 \\ 0.2 \\ 0.3 \end{bmatrix}$,

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, B = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \text{ Compute } E(AX + BY)$$

Outline

- Probability and Conditional Probability
- Random Variables and Random Vectors
- **Jointly Distributed Random Vectors and Conditional Expectation**

Jointly distributed random vectors: $X \in \mathbb{R}^n, Y \in \mathbb{R}^m$

- Completely determined by joint density (mass) function:

$$(X, Y) \sim f_{XY}(x, y)$$

Compute probability:

- marginal density: $X \sim f_X(x), Y \sim f_Y(y)$, where

$$f_X(x) = \int_{\mathbb{R}^m} f_{XY}(x, y) dy, \quad f_Y(y) = \int_{\mathbb{R}^n} f_{XY}(x, y) dx,$$

- Example: $x = \begin{bmatrix} X_1 \\ X_2 \end{bmatrix}, \text{Prob}\left(X = \begin{bmatrix} 0 \\ 1 \end{bmatrix}\right) = \frac{1}{2}, \text{Prob}\left(X = \begin{bmatrix} 1 \\ 2 \end{bmatrix}\right) = \frac{1}{3}, \text{Prob}\left(X = \begin{bmatrix} -1 \\ 1 \end{bmatrix}\right) = \frac{1}{6}$

- This is joint distribution for X_1, X_2

- The conditional density: $(X, Y) \sim f_{XY}(x, y)$
 - Quantify how the observation of a value of Y , $Y = y$, affects your belief about the density of X
 - The conditional probability definition implies (nontrivially)

$$P(A | B) = P(A \cap B) / P(B) \Rightarrow p_{X|Y}(X = i | Y = j) = \frac{p_{XY}(X=i, Y=j)}{\sum_i p_{XY}(X=i, Y=j)}$$

$$f_{X|Y}(x|y) = \frac{f_{XY}(x, y)}{f_Y(y)}$$

- Law of total probability: $P(A) = \sum_{i=1}^n P(A \cap B_i) = \sum_{i=1}^n P(A|B_i)P(B_i)$

$$f_X(x) = \int_{R^m} f_{X|Y}(x|y) f_Y(y) dy$$

$$f_Y(y) = \int_{R^n} f_{Y|X}(y|x) f_X(x) dx$$

- **X is independent of Y** , denoted by $X \perp Y$,
if and only if $f_{XY}(x, y) = f_X(x) f_Y(y)$

- **Conditional expectation:**

- The conditional mean of $X|Y = y$ is

$$E(X|Y = y) = \int_{R^n} x f_{X|Y}(x|y) dx$$

$$E(X|Y = y) = \sum_i i \cdot \text{Prob}(X = i|Y = y)$$

- Example 1:

- $E(X|Y = 1)$

		X				
		2	3	4	5	6
Y	1	1/4	1/8	1/8		
	2		1/6	1/12	1/12	
	3			1/12	1/24	1/24

- $E(X|Y = 2)$

- $E(X|Y=3)$
- **Example 2:** Suppose that (X, Y) is uniformly distributed on the square $S = \{(x, y) : -6 < x < 6, -6 < y < 6\}$. Find $E(Y | X = x)$.

- Law of total probability implies:
 - $E(X) = \sum_y E(X|Y = y) \cdot p_Y(Y = y)$

- $E(g(X, Y)) = \sum_y E(g(X, Y)|Y = y) \cdot p_Y(Y = y)$

- Continue Example 1:

		<i>X</i>				
		2	3	4	5	6
<i>Y</i>	1	1/4	1/8	1/8		
	2		1/6	1/12	1/12	
	3			1/12	1/24	1/24

- Example 3.: outcomes with equal chance: $(1,1)$, $(2, 0)$, $(2,1)$, $(1,0)$, $(1,-1)$, $(0,0)$, with $g(X, Y) = X^2Y^2$

Method 1: $E(g(X, Y)) = E(X^2Y^2) = 1^2 \cdot (-1)^2 \cdot \frac{1}{6} + 1^2 \cdot 1^2 \cdot \frac{1}{6} + 2^2 \cdot 1^2 \cdot \frac{1}{6} = 1$

Method 2: conditioning on values of $Y = -1, 0, 1$

		X		
		0	1	2
Y	-1	0	1/6	0
	0	1/6	1/6	1/6
	1	0	1/6	1/6

- More discussions