SDM5008 Advanced Control for Robotics

Lecture Note 9: Probability Review for Reinforcement Learning

- linear algebra

- probability (stochastic)

conditional probability

& expectation

- optimization

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? $\Rightarrow$ modeling uncertainty

world is deterministic! uncertainty is primarily due to lack of information.

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

there is no right or wrong probability.

- Formalism: Probability Space $(\Omega, \mathcal{F}, P)$
  - 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

e.g. toss a coin. $\Omega = \{0, 1\}$, $\mathcal{F} = \{\phi, \{0,1\}, \{0\}, \{1\}\}$

toss a die: $\Omega = \{1, \cdots, 6\}$, $\mathcal{F} = \{\phi, \{1\} \cdots \{6\}, \{1,2\}, \{1,3\} \cdots \{1 \cdots 6\}\}$

Event $A \in \mathcal{F}$. $\Rightarrow$ prob(A)

$A = \phi$ "even", $A = \{4, 5, 6\}$

prob(A) $= \begin{cases} 0.5 & ✓ \\ 0.8 & ✓ \end{cases}$

correct if accurately captures

our uncertainty

prob({1, 2, 3}) = 0.5

P({1, 3}) = 0.6

0.5 = prob({1, 2, 3})

= prob({1, 3}) + prob({2, 3})

= 0.6 +

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:**
  - $P(\emptyset) = 0$

    $P(\Omega) = 1 \Rightarrow P(\Omega \cup \phi) = 1 = P(\Omega) + P(\phi) = 1 \Rightarrow P(\phi) = 0$
  - Law of total probability: $P(B) = \sum_i^n P(B \cap A_i)$, for any partitions $\{A_i\}$ of Ω

divide & conquer.

- Recall a collection of sets $A_1, \ldots, A_n$ is called a partition of Ω if
  - $A_i \cap A_j = \emptyset$, for all $i \neq j$ (mutually exclusive)
  - $A_1 \cup A_2 \cdots \cup A_n = \Omega$

prob(B) = P(B ∩ A1) + P( ) + P( )

Conditional probability

conditional probability is probability, with its own probability space

- Probability of event A happens given that event B has already occurred
  - $P(A|B) \overset{\triangle}{=} \frac{P(A\cap B)}{P(B)}$ conditional probability measure
- 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, $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{P})$ the is derived from an original probability space $(\Omega, \mathcal{F}, P)$ given some event has occurred**
  - After B occurred we are uncertain only about the outcomes inside B

start with $(\Omega, \mathcal{F}, P)$ original probability space : (e.g. toss a die, $\Omega = \{1, \ldots, 6\}$, ...

e.g. $B = \{1, 3, 5\}$

$P(B) = 0.5$)

$\Rightarrow$ B event has occurred $\Rightarrow$ new probability space $\underline{(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{P})}$

- $\tilde{\Omega} = B$, i.e. all the possible outcomes are in B
- $\tilde{\mathcal{F}}$ = all subsets of B (e.g. $A = \{2, 6, 5\}$. $A \notin \tilde{\mathcal{F}}$, but all $\underline{A \cap B} \subseteq \tilde{\mathcal{F}}$

– what is $\tilde{P}$? we want to be consistent with the original P. e.g. suppose ① $C \subseteq B$

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

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

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

$$\Rightarrow P(B \cap A) = P(B|A) \cdot P(A) \Leftarrow$$

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

$\Rightarrow C \in \tilde{\mathcal{F}}$, $\tilde{P}(C) \triangleq ?$

$\tilde{P}(C) \overset{?}{=} P(C)$ ✗

$\Rightarrow C = B$, $\tilde{P}(B) = \tilde{P}(\tilde{\Omega}) = 1$

but $P(C) < 1$ ✗

– One possible updated measure is $\tilde{P}(C) \triangleq \frac{P(C)}{P(B)}$

$\tilde{P}(\tilde{\Omega}) = 1$

- Events A and B are called (statistically) independent if

definition.

  - $P(A|B) = P(A)$

~~indep~~ independent

  - Or equivalently: $P(A \cap B) = P(A)P(B)$

If $C \nsubseteq B$, $C \cap B \subseteq B$

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

$$\Rightarrow P(A \cap B) = P(A)P(B)$$

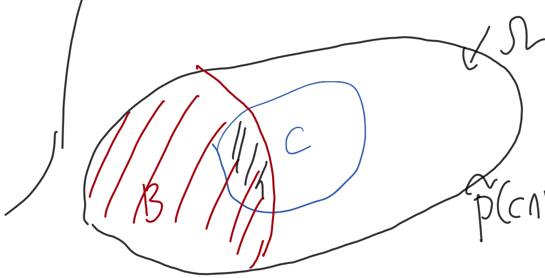

$$\frac{P(C)}{\;}$$

$$\tilde{P}(C \cap B) = P(C|B) = \frac{P(C \cap B)}{P(B)}$$

$$P(\cdot|B) \neq P(\cdot)$$

- **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:

$$A = \text{the draw of a red or a blue chip}$$

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

$\cdot (\Omega, \mathcal{F}, P)$ sampling space $\Omega = \{r_1\, r_2 \cdots r_5,\ w_1\ w_2\ w_3,\ b_1, b_2\}$

$A = \{r_1, \cdots, r_5, b_1, b_2\}$ $B = \{r_1\, r_2 \cdots r_5,\ w_1, w_2, w_3\}$

$$P(A|B) = \frac{P(A \cap B)}{P(B)} = \frac{\frac{1}{2}}{\frac{8}{10}} = \frac{5}{8}$$

$P(A) = \frac{7}{10}$, $P(B) = \frac{8}{10}$ $P(A \cap B) = \frac{1}{2}$

Outline

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

- **What is random variable and random vector?**
  - Deterministic variable:

· e.g. z is a deterministic variable, mean z can take only one value (single-valued variable), which may or may not be known

  - Random variable:

· e.g. z is a random variable (multi-valued variable) can take multiple (or even infinite) possible values, each value occurs with certain probability.

e.g. $z = \{0, 1\}$ with probabilities $\frac{1}{2}$, $\frac{1}{2}$

$z = \left\{ \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ 1 \end{bmatrix} \right\}$ with probabilities $\frac{1}{3}$, $\frac{1}{6}$, $\frac{3}{6}$

Continuous: $z \in [0, 1]$

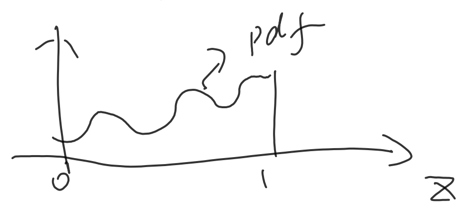

Let a be deterministic variable $\in \mathbb{R}$

let b be random variable.

How to specify probability measure

- Discrete random variable: probability mass function (pmf)

  e.g. toss a coin or die

$\Omega = \{0, 1\}$

$p(0) = 0.4$

$p(1) = 0.6$

$\Omega = \{1, \cdots, N\}$, $p(i) = \text{prob}(X = i)$

$\mathcal{F}$ = subsets of Ω

$A \in \mathcal{F}$ $\quad P(A) = \sum_{i \in A} p(i)$

$A = \{1, 6, 8\} = \{1\} \cup \{6\} \cup \{8\}$

- Continuous random variable: probability density function (pdf)

  e.g. temperature density

$X \in \mathbb{R}$, $\Omega = \mathbb{R}$ $\quad$ pdf: $f_X(x) =$ "prob" of $X = x$

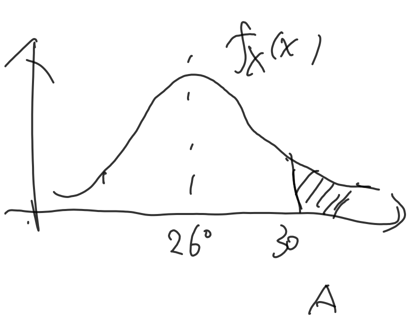

For any event $A \subseteq \mathbb{R}$ (e.g. $A = [0, 1) \cup (\frac{1}{2}, 8]$)

$$P(A) = \int_{x \in A} f_X(x)\, dx$$

How to specify probability measure

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

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

X: random variable in $\mathbb{R}^n$

$X \in \mathbb{R}^2$

$x = \begin{bmatrix} 1 \\ 2 \end{bmatrix}$

$X = x$

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

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

$x = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}$

$\Rightarrow$ short hand notation for $f(x_1, x_2, \cdots, x_n)$

$f_X(x) \approx$ prob $X = x$

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

$\int f(x) dx_1 dx_2 \cdots$

multi-dim integration

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

Continuous random vector: $\boldsymbol{E(X) \triangleq \int_{R^n} x f(x) dx}$

Discrete random vector: $\boldsymbol{E(X) \triangleq \sum_x x \cdot Prob(X = x)}$

$E(X)$ → integral, sum

$X = \begin{bmatrix} X_1 \\ X_2 \end{bmatrix}$ by definition

$E(X) = \int \underbrace{\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}}_{\in \mathbb{R}^2} \cdot \underbrace{f(x_1, x_2)}_{\text{scalar}} dx_1 dx_2 = \begin{bmatrix} \iint x_1 f(x_1,x_2) dx_1 dx_2 \\ \iint x_2 f(x_1,x_2) dx_1 dx_2 \end{bmatrix}$

$\to \underbrace{\int x_1 \overbrace{\int f(x_1,x_2) dx_2}^{f(x_1)} dx_1}_{E(x_1)}$

- ① Expectation: $E(X) = \begin{bmatrix} E(X_1) \\ E(X_2) \\ \vdots \\ E(X_n) \end{bmatrix}$

$E_{X_1,\dots,X_n}(X) = \begin{bmatrix} E_{x_1}(X_1) \\ \\ \end{bmatrix}$

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

$E(X) = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \cdot \frac{1}{2} + \begin{bmatrix} 1 \\ 2 \end{bmatrix} \frac{1}{3} + \begin{bmatrix} -1 \\ 1 \end{bmatrix} \frac{1}{6} = \begin{bmatrix} \frac{1}{6} \\ \frac{4}{3} \end{bmatrix}$

X_1:	0	1	-1
$Prob(X_1)$:	$\frac{1}{2}$	$\frac{1}{3}$	$\frac{1}{6}$

$E(X) = \begin{bmatrix} E(X_1) \\ E(X_2) \end{bmatrix}$,

$E(X_1) = 0 \cdot \frac{1}{2} + 1 \cdot \frac{1}{3} + (-1) \frac{1}{6} = \frac{1}{6}$

$E(X_2) = 1 \cdot \frac{2}{3} + 2 \cdot \frac{1}{3} = \frac{4}{3}$

X_2:	1	2
$prob(\cdot)$:	$\frac{2}{3}$	$\frac{1}{3}$

Linearity of Expectation:

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

$$\boldsymbol{E(AX) = AE(X)}$$

why ?

$$E(AX) = \int (Ax) \cdot f_X(x)\, dx = A \int x f_X(x) dx = AE(X)$$

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

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

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

$$= \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0.5 \\ 0.25 \end{bmatrix} + \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0.1 \\ 0.2 \\ 0.3 \end{bmatrix}$$

$$= \begin{bmatrix} 0.85 \\ 0.55 \end{bmatrix}$$

Outline

- Probability and Conditional Probability

- Random Variables and Random Vectors

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

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

$Z = \begin{bmatrix} X \\ Y \end{bmatrix} \in \mathbb{R}^{n+m} \sim f_Z(z)$

- Completely determined by joint density (mass) function: $(X, Y) \sim f_{XY}(x, y)$ $\approx$ "Prob" of $X = x, Y = y$

  Compute probability:

  $$P((X,Y) \in A) = \int_A f_{XY}(x,y)\,dx\,dy$$

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

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

  $$\triangleq \sum_{\text{all possible } y} \text{prob}(X = x, Y = y)$$

- Example: $X = \begin{bmatrix} X_1 \\ X_2 \end{bmatrix}, Prob\left(X = \begin{bmatrix} 0 \\ 1 \end{bmatrix}\right) = \frac{1}{2}, Prob\left(X = \begin{bmatrix} 1 \\ 2 \end{bmatrix}\right) = \frac{1}{3}, Prob\left(X = \begin{bmatrix} -1 \\ 1 \end{bmatrix}\right) = \frac{1}{6}$
  - This is joint distribution for X_1, X_2

marginal of X_2:

X_2	1	2
prob(X_2)	$\frac{2}{3}$	$\frac{1}{3}$

- 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)

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

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

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

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

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

$$P(X=i) = \sum_j P(X=i, Y=j)$$

$$= \sum_j P(X=i \mid Y=j) \cdot P(Y=j)$$

- ***X* is independent of *Y***, denoted by $X \perp Y$,

  if and only if $f_{XY}(x, y) = f_X(x) f_Y(y)$

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

is expectation.

- **Conditional expectation:**
  - The conditional mean of $X|Y=y$ is

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

z

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

- Example 1:
  - $E(X|Y=1) = 2\cdot\frac{1}{2} + 3\cdot\frac{1}{4} + 4\cdot\frac{1}{4} = \frac{11}{4}$

$X\|Y=1$	2	3	4
$prob(X\|Y=1)$	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{4}$

		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

  - $E(X|Y=2) = 3\cdot\frac{1}{2} + 4\cdot\frac{1}{4} + 5\frac{1}{4} = \frac{15}{4}$

- E(X | Y=3) = ?

$$f_{XY}(x,y) = \begin{cases} c & \text{if } x,y \in S \\ 0 & \text{ow} \end{cases}$$

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

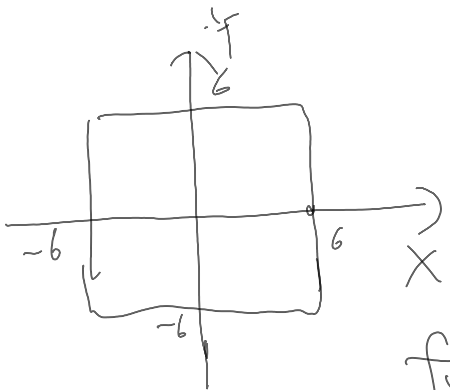

$$f_{XY}(x,y) = \begin{cases} \frac{1}{144}, & x,y \in S \\ 0, & \text{otherwise} \end{cases}$$

We need to compute $E(Y|X=x) = \int y\, f_{Y|X}(y|x)\, dy$

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

$$f_X(x) = \int f_{XY}(x,y)\, dy = \begin{cases} 0, & x < -6, x > 6 \\ \underbrace{\int_{-6}^{6} \frac{1}{144} dy}_{} & x \in [-6,6] \end{cases} = \begin{cases} \frac{1}{12} & x \in [-6,6] \\ 0, & \text{otherwise} \end{cases}$$

$$f_{Y|X}(y|x) = \begin{cases} \frac{\frac{1}{144}}{\frac{1}{12}} = \frac{1}{12} & \text{if } x,y \in S \\ 0 & \text{otherwise} \end{cases}$$

$$E(Y|X=x) = \int_{-\infty}^{+\infty} y\, f_{Y|X}(y|x)\, dy = 0$$

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

$E(X) = E(E(X|Y))$

discrete case $X = \{1 \cdots\}$

$E(X) = \sum_i i \cdot p(X=i) = \sum_i i\left(\sum_j p(X=i, Y=j)\right)$

$= \sum_i i\left(\sum_j p(X=i|Y=j) \cdot p(Y=j)\right)$

$= \sum_j \left(\sum_i \left(i\, p(X=i|Y=j)\right)\right) \cdot p(Y=j)$

$\Downarrow$ $E(X|Y=j)$

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

- Continue Example 1:

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

- Compute $E(X)$

Method 1: by definition of $E(X)$

X	2	3	4	5	6
prob(·)	$\frac{1}{4}$	$\frac{7}{24}$	$\frac{7}{24}$	$\frac{3}{24}$	$\frac{1}{24}$

$\Rightarrow E(X) = 2 \cdot \frac{1}{4} + 3 \cdot \frac{7}{24} + \cdots = \#$

equal.

- Method 2 $E(X) = \underbrace{E(X \mid Y=1)} \cdot \text{prob}(Y=1) + E(X \mid Y=2) \cdot \text{prob}(Y=2) + E(X \mid Y=3)\, \text{prob}(Y=3)$

$= \frac{11}{4} \cdot \frac{1}{2} + \frac{15}{4} \cdot \frac{1}{3} + ? \cdot \frac{1}{6} = \text{Same } \#$

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

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

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

· towering property: $E(X^2Y^2) = E_Y\left(E_{X|Y}(X^2Y^2 \mid Y)\right)$

$= \text{prob}(Y=-1) \cdot E(X^2Y^2 \mid Y=-1) + \text{prob}(Y=0)\, \cancel{E(X^2Y^2 \mid Y=0)}$ 0

$+ \text{prob}(Y=1)\, E(X^2Y^2 \mid Y=1)$

$E(X^2Y^2 \mid Y=-1) = \cancel{0^2 \cdot (-1)^2\, \text{prob}(X=0 \mid Y=-1)} + 1^2 (-1)^2\, \text{prob}(X=1 \mid Y=-1) + \cancel{2^2 \cdot (-1)\, \text{prob}(X=2 \mid Y=-1)}$

$= 1$

$E(X^2Y^2 \mid Y=1) = 1^2 (1)^2 \cdot \frac{1}{2} + 2^2 \cdot 1^2 \cdot \frac{1}{2} = \frac{5}{2}$

$E(X^2Y^2) = 1 \cdot \frac{1}{6} + \frac{5}{2} \times \frac{1}{3} = 1$

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

- More discussions