

CSC380: Principles of Data Science

Probability 4

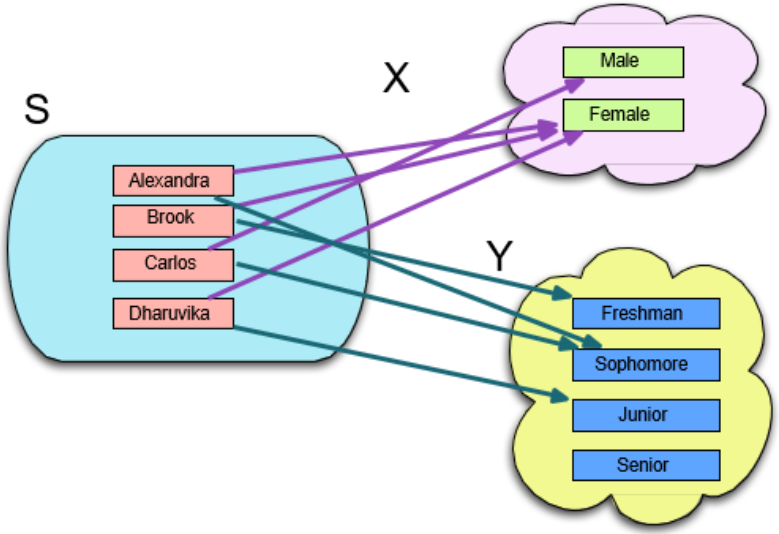
Xinchen Yu

Outline

- Multivariate Random Variables
  - Joint distribution vs. Marginal distribution
  - Independence of RVs
- Expectation and Variance Revisited
  - Covariance, correlation
- Example multivariate RVs

Multivariate Random Variables

Multivariate RVs: example



- X: people -> their genders
- Y: people -> their class year
- We'd like to answer questions such as: does X and Y have a correlation?
  - I.e., is a student in higher class year more likely to be male?
- We call (X, Y) a random vector, or a multivariate RV, and will study its *joint* distribution

Joint distribution of discrete RVs

- The joint PMF (probability mass function) of discrete random variables X, Y:

$$f(x, y) = P(X = x, Y = y)$$

Examples

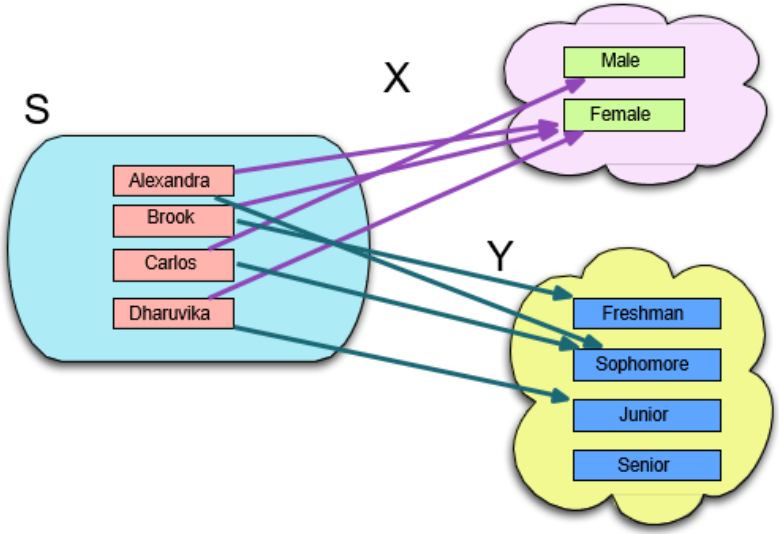
Alexandra

$$P(X = \text{Fem}, Y = \text{Soph}) = \frac{1}{4}$$

Dharuvika

$$P(X = \text{Fem}, Y = \text{Jun}) = \frac{1}{4}$$

…



Joint distribution of discrete RVs

- X: # of cars owned by a randomly selected household
- Y: # of computers owned by the same household
- Joint pmf shown with a table

	y			
x	1	2	3	4
1	0.1	0	0.1	0
2	0.3	0	0.1	0.2
3	0	0.2	0	0

- Probability that a randomly selected household has ≥ 2 cars and ≥ 2 computers?
  - $P(X \geq 2, Y \geq 2) = 0.5$

Marginal distributions

Given joint distribution of (X, Y), need distribution of one of them, say X.

- Named the *marginal distribution* of X.

	y			
x	1	2	3	4
1	0.1	0	0.1	0
2	0.3	0	0.1	0.2
3	0	0.2	0	0

- How to find $P(X = x)$?
- Using law of total probability:

$$f_1(x) = \sum_y f(x, y)$$

- This operation is called *marginalization* (‘marginalizing out variable Y’, or variable elimination)

Marginal distributions

	y				
x	1	2	3	4	Total
1	0.1	0	0.1	0	0.2
2	0.3	0	0.1	0.2	0.6
3	0	0.2	0	0	0.2
Total	0.4	0.2	0.2	0.2	1.0

f_1: *marginal distribution* of X

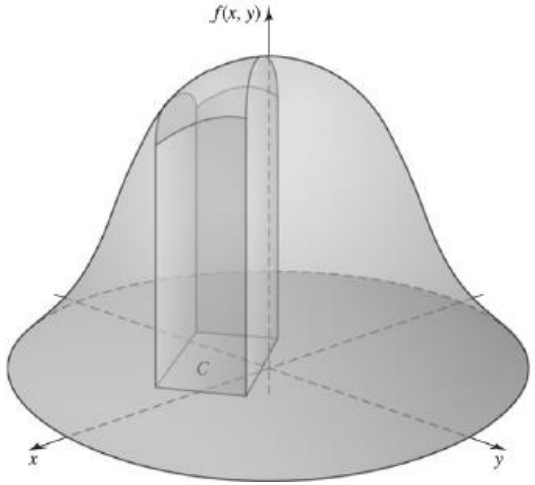
f_2: *marginal distribution* of Y

$$f_1(X = 1) = \sum_y f(1, y) = 0.1 + 0 + 0.1 + 0 = 0.2$$

Joint distribution of continuous RVs

- Any continuous random vector (X,Y) has a *joint probability density function* (PDF) $f(x,y)$, such that for all C,

$$P\big((X,Y) \in C\big) = \iint_C f(x,y)\,dx\,dy$$



$f(x,y)$: represent a 2D surface

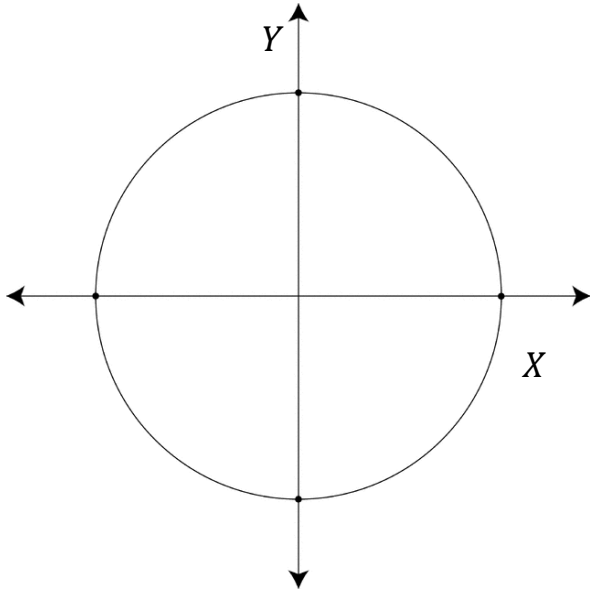
double integral: the *volume* under the surface

Properties:

- f is nonnegative
- $\iint_{R^2} f(x,y)\,dx\,dy = 1$ (R^2 = the whole x-y plane)
  - $P\big((X,Y) \in R^2\big) = 1$

Example: dartboard

- Dartboard with center (0,0) and radius 1; dart lands uniformly at random on the board



- What is the joint PDF of (X, Y)?
- Fact: the PDF is

$$f(x,y) = \begin{cases} c, x^2 + y^2 \leq 1 \\ 0, \text{otherwise} \end{cases}$$

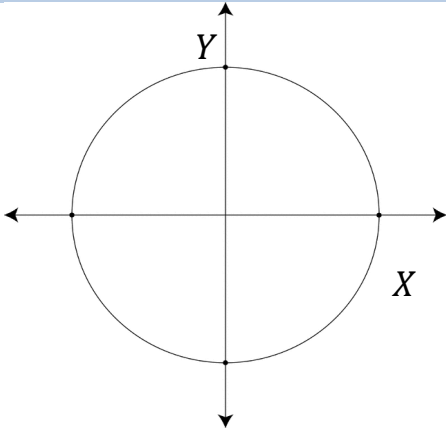
- This is called “the Uniform distribution over the unit disk”

Example: dartboard

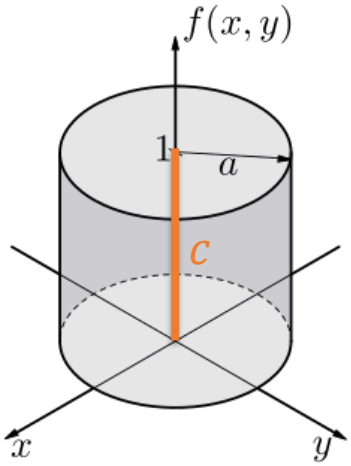
The PDF of X, Y is

$$f(x,y) = \begin{cases} c, x^2 + y^2 \leq 1 \\ 0, \text{otherwise} \end{cases}$$

Can we find c?



Observe: volume under $f(x,y)$ is πc (cylinder) which must also be 1



Therefore, $c = 1/\pi$

Marginal distribution of continuous RV

Given joint distribution of continuous RV (X, Y), how to find X's PDF f_1?

Fact (marginalization) $f_1(x) = \int_R f(x, y)\, dy$

Replacing summation with integration in the continuous case ('marginalizing / integrating out variable Y')

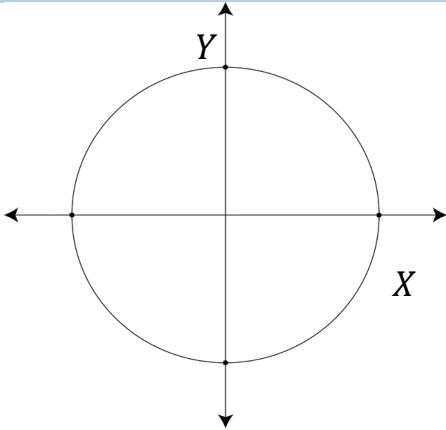
How about Y's PDF f_2?

- Marginalize out X

Example: dartboard

The PDF of X, Y is

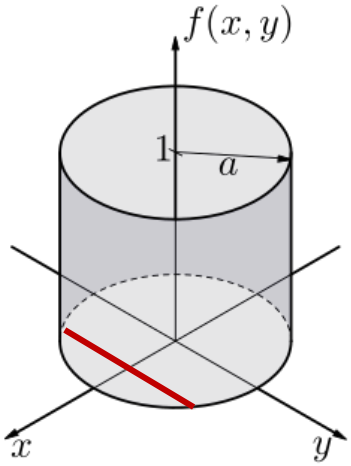
$$f(x,y) = \begin{cases} \frac{1}{\pi}, x^2 + y^2 \leq 1 \\ 0, \text{otherwise} \end{cases}$$



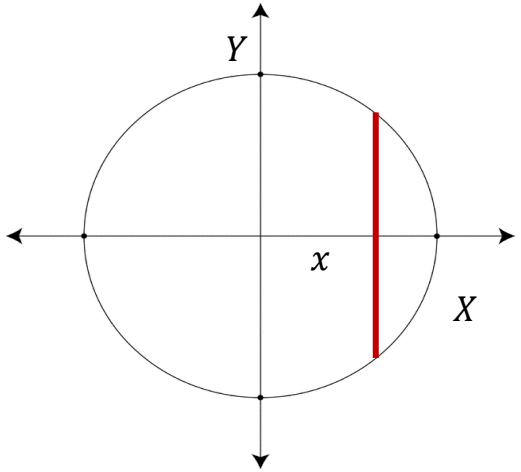
What is the marginal distribution over X?

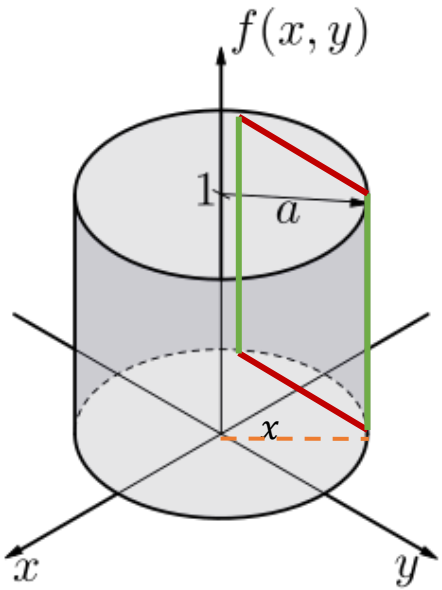
$$f_1(x) = \int_{-\infty}^{+\infty} f(x,y)\, dy$$

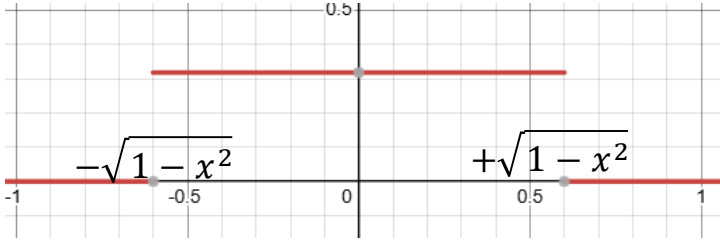
How to find this integral?



Example: dartboard



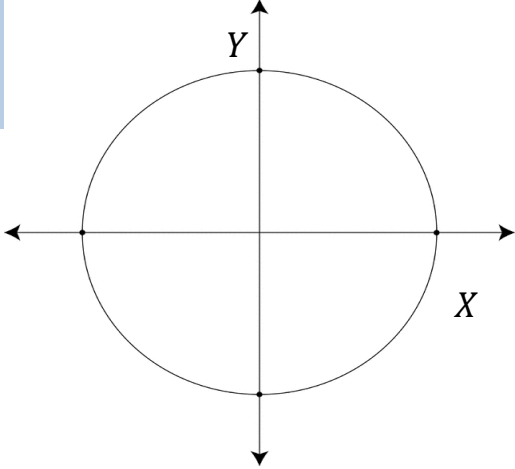




For a fixed $x \in [-1, 1]$, we can think of $f(x)$ is the area of the slice:

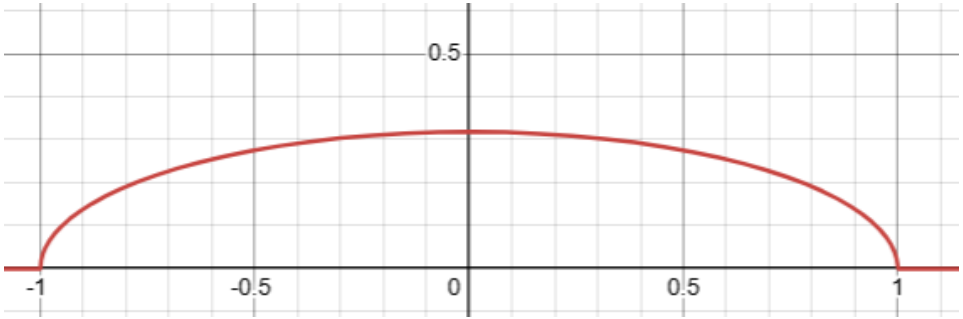
- height: $\frac{1}{\pi}$, width: $2 \cdot \sqrt{1-x^2}$
- $f_1(x) = \frac{2}{\pi} \cdot \sqrt{1-x^2}$

Example: dartboard



- In summary,

$$f(x) = \begin{cases} \frac{2}{\pi} \cdot \sqrt{1 - x^2}, x \in [-1,1] \\ 0, \quad \text{otherwise} \end{cases}$$



X's distribution is NOT $\text{Uniform}([-1,1])$!

Actually makes sense: X closer to 1 is harder to be hit

Joint distribution of more than 3 RVs

- We can consider the joint distribution of more than 3 random variables,
  - E.g. (A,B,C), A = gender, B = class year, C = blood type
- Discrete RVs: can still define joint PMFs

a	b	c	$P(A=a, B=b, C=c)$
0	0	0	0.06
0	0	1	0.09
0	1	0	0.08
0	1	1	0.12
1	0	0	0.06
1	0	1	0.24
1	1	0	0.10
1	1	1	0.25

Marginalization

a	b	c	$P(A=a, B=b, C=c)$
0	0	0	0.06
0	0	1	0.09
0	1	0	0.08
0	1	1	0.12
1	0	0	0.06
1	0	1	0.24
1	1	0	0.10
1	1	1	0.25

Given the joint distribution of (A, B, C)

- What is the distribution of A?
  - Need to find $P(A = 0)$ and $P(A = 1)$

$$P(A = 0) = \sum_{b,c} P(A = 0, B = b, C = c)$$

Marginalization: summing over irrelevant variables

- What is the joint distribution of (A, B)?
  - Need to find $P(A = 0, B = 0), \ldots, P(A = 1, B = 1)$

$$P(A = 0, B = 0) = \sum_{c} P(A = 0, B = 0, C = c)$$

Marginalization for continuous RVs

Suppose joint PDF of (A, B, C) is $f(a, b, c)$

- What is the PDF of A?

$$f_A(a) = \iint_{R^2} f(a, b, c)\ db\ dc$$

- What is the joint PDF of (A, B)?

$$f_{A,B}(a, b) = \int_R f(a, b, c) dc$$

Marginalization: summing over irrelevant variables

- These operations generalize to joint PDFs of more RVs..

Plan

- Multivariate RVs
  - $f_1(x) = \sum_y f(x, y)$ for discrete X, Y
  - $f_1(x) = \int_R f(x, y)\, dy$ for continuous X, Y
- Independence of RVs
- Conditional distribution of RVs
- Mean of conditional distribution
- Finding distribution of X + Y when they are independent

Independence of RVs

Independence of two RVs

- RVs X, Y are independent (denoted by $X \perp\!\!\!\perp Y$) if

$$f(x,y) = f_1(x) \cdot f_2(y), \quad \textit{for all } x, y$$

PMF or PDF Marginal of X Marginal of Y

- E.g. for discrete X, Y,

$$P(X = 3, Y = 4) = P(X = 3) \cdot P(Y = 4)$$

Therefore, $\{X = 3\}$ and $\{Y = 4\}$ are independent events

Independence is invariant under transformations

Fact If X, Y are independent, then $f(X), g(Y)$ are also independent

E.g. X = tomorrow's temperature (in Celsius); Y = tomorrow's NVIDIA stock price (in $)

f(X) = tomorrow's temperature (in Fahrenheit); g(Y) = tomorrow's NVIDIA stock price (in cents)

Independence of more than two RVs

- RVs $X_1, \ldots, X_n$ are independent if their joint PMF or PDF satisfy

$$f(x_1, x_2, \ldots, x_n) = f_1(x_1) f_2(x_2) \ldots f_n(x_n),$$

PMFs or PDFs Marginal for X_1 Marginal for X_n

for all $x_1, \ldots, x_n$

This captures many real-world applications:

- Independent trials: each X_i is $\text{Bernoulli}(p)$
  - Flip 10 coins: $x_1, x_2, \ldots, x_{10}$

Independence of more than two RVs

Fact If $X_1, \ldots, X_n$ are independent, then

- any subset $X_{i_1}, \ldots, X_{i_p}$ are independent
  - E.g. X_1, X_3, X_7 are independent
- any disjoint subset $(X_{i_1}, \ldots, X_{i_m})$, $(X_{j_1}, \ldots, X_{j_l})$ are independent
  - E.g. (X_1, X_2) is independent of X_3
  - (X_1, X_3) is independent of (X_2, X_4)

Conditional distributions of RVs

Conditional distributions (discrete)

- X, Y have joint PMF f. Y has marginal PMF f_2

- Conditional PMF of X given $Y = y$:

$$g_1(x|y) = \frac{f(x,y)}{f_2(y)}$$

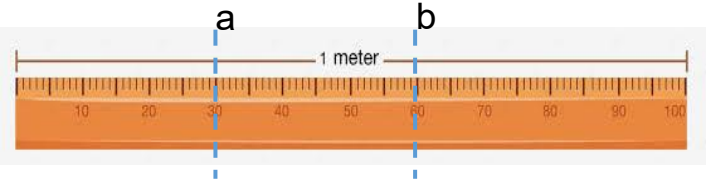
Same as $\frac{P(X=x,Y=y)}{P(Y=y)} = P(X = x \mid Y = y)$

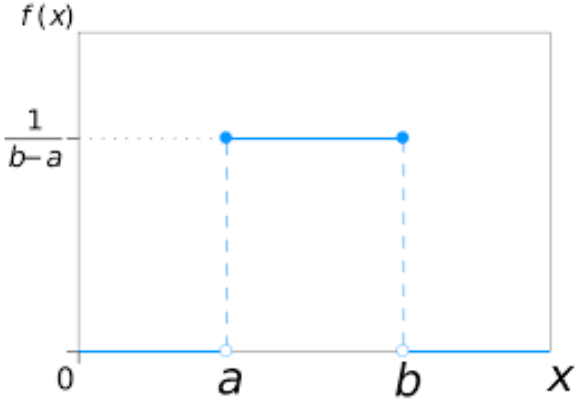
- $g_1(x|y)$ is viewed as a function of x: "the conditional distribution of X given $Y = y$"

Recap: Uniform Distribution

- $X \sim \text{Uniform}([a, b])$

$$f(x) = \begin{cases} 0, & y < a \\ \frac{1}{b-a}, & y \in [a, b] \\ 0, & y > b \end{cases}$$





Conditional distributions & independence

Fact X,Y are independent

$\Leftrightarrow$ for all y, $g(x|y)$ are all equal to $f(x)$

Here, g, f are PMF or PDF depending on the types of X,Y

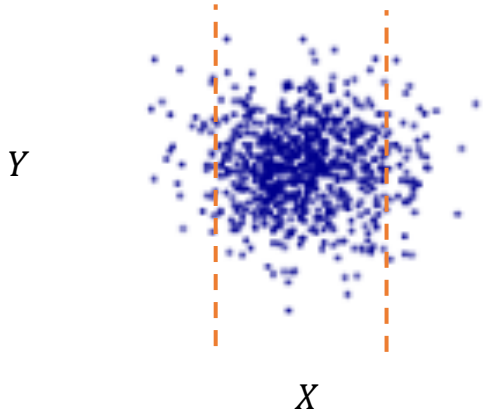
Assume Y can only take the value 1, 2, and 3. We say X,Y are independent when

- $f(X = x) = g(X = x|Y = 1)$, and
- $f(X = x) = g(X = x|Y = 2)$, and
- $f(X = x) = g(X = x|Y = 3)$

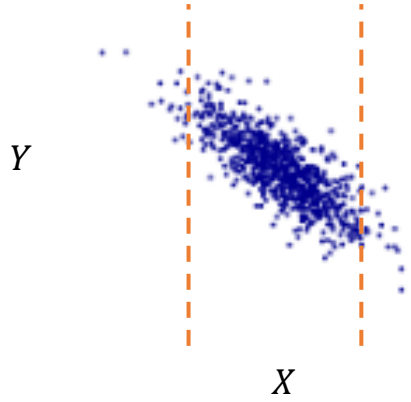
In other words, knowing Y does not change our belief on X

Independence: visualization

- Left: X, Y independent; Right: X, Y not independent



$g_1(y|x=-1)$ $g_1(y|x=+1)$



$g_1(y|x=-1)$ $g_1(y|x=+1)$

Conditional expectation

Definition The mean of the conditional distribution of X given $Y = y$, is called the *conditional expectation* of X given $Y = y$, denoted as $\mathrm{E}[\,X \mid Y = y\,]$.

$\mathrm{E}[\,X \mid Y = y\,]$ can be found by:

- $\sum_x x \cdot g(x|y)$, if X is discrete

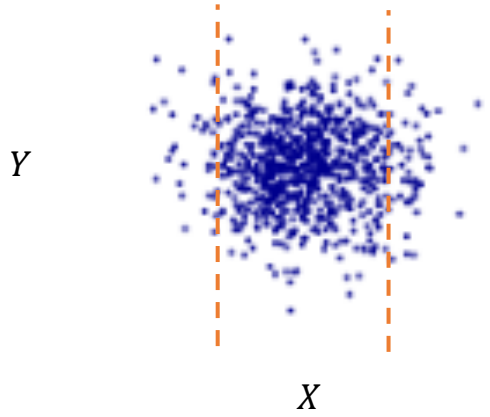
  Conditional PMF

- $\int_{-\infty}^{+\infty} x \cdot g(x|y)\, dx$, if X is continuous

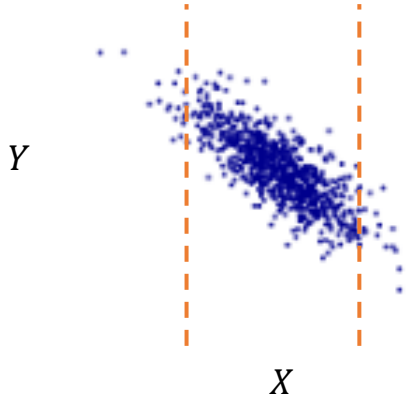
  Conditional PDF

Independence: visualization

- Left: X, Y independent; Right: X, Y not independent



$g_1(y|x=-1)$ $g_1(y|x=+1)$



$g_1(y|x=-1)$ $g_1(y|x=+1)$

Which one is larger, $\mathrm{E}[Y|X=-1]$ or $\mathrm{E}[Y|X=+1]$?
The former

Conditional expectation

Example Roll 2 fair dice. Expected value of die 1 given that their sum is 5?

Solution X: outcome of die 1; Y: sum of 2 dice, $E[X \mid Y = 5]$

Let's find the conditional distribution of X given $Y = 5$ first.

$$g_1(x \mid 5) = P(X = x \mid Y = 5)$$
$$= \frac{P(X=x, Y=5)}{P(Y=5)}$$

When is this nonzero?