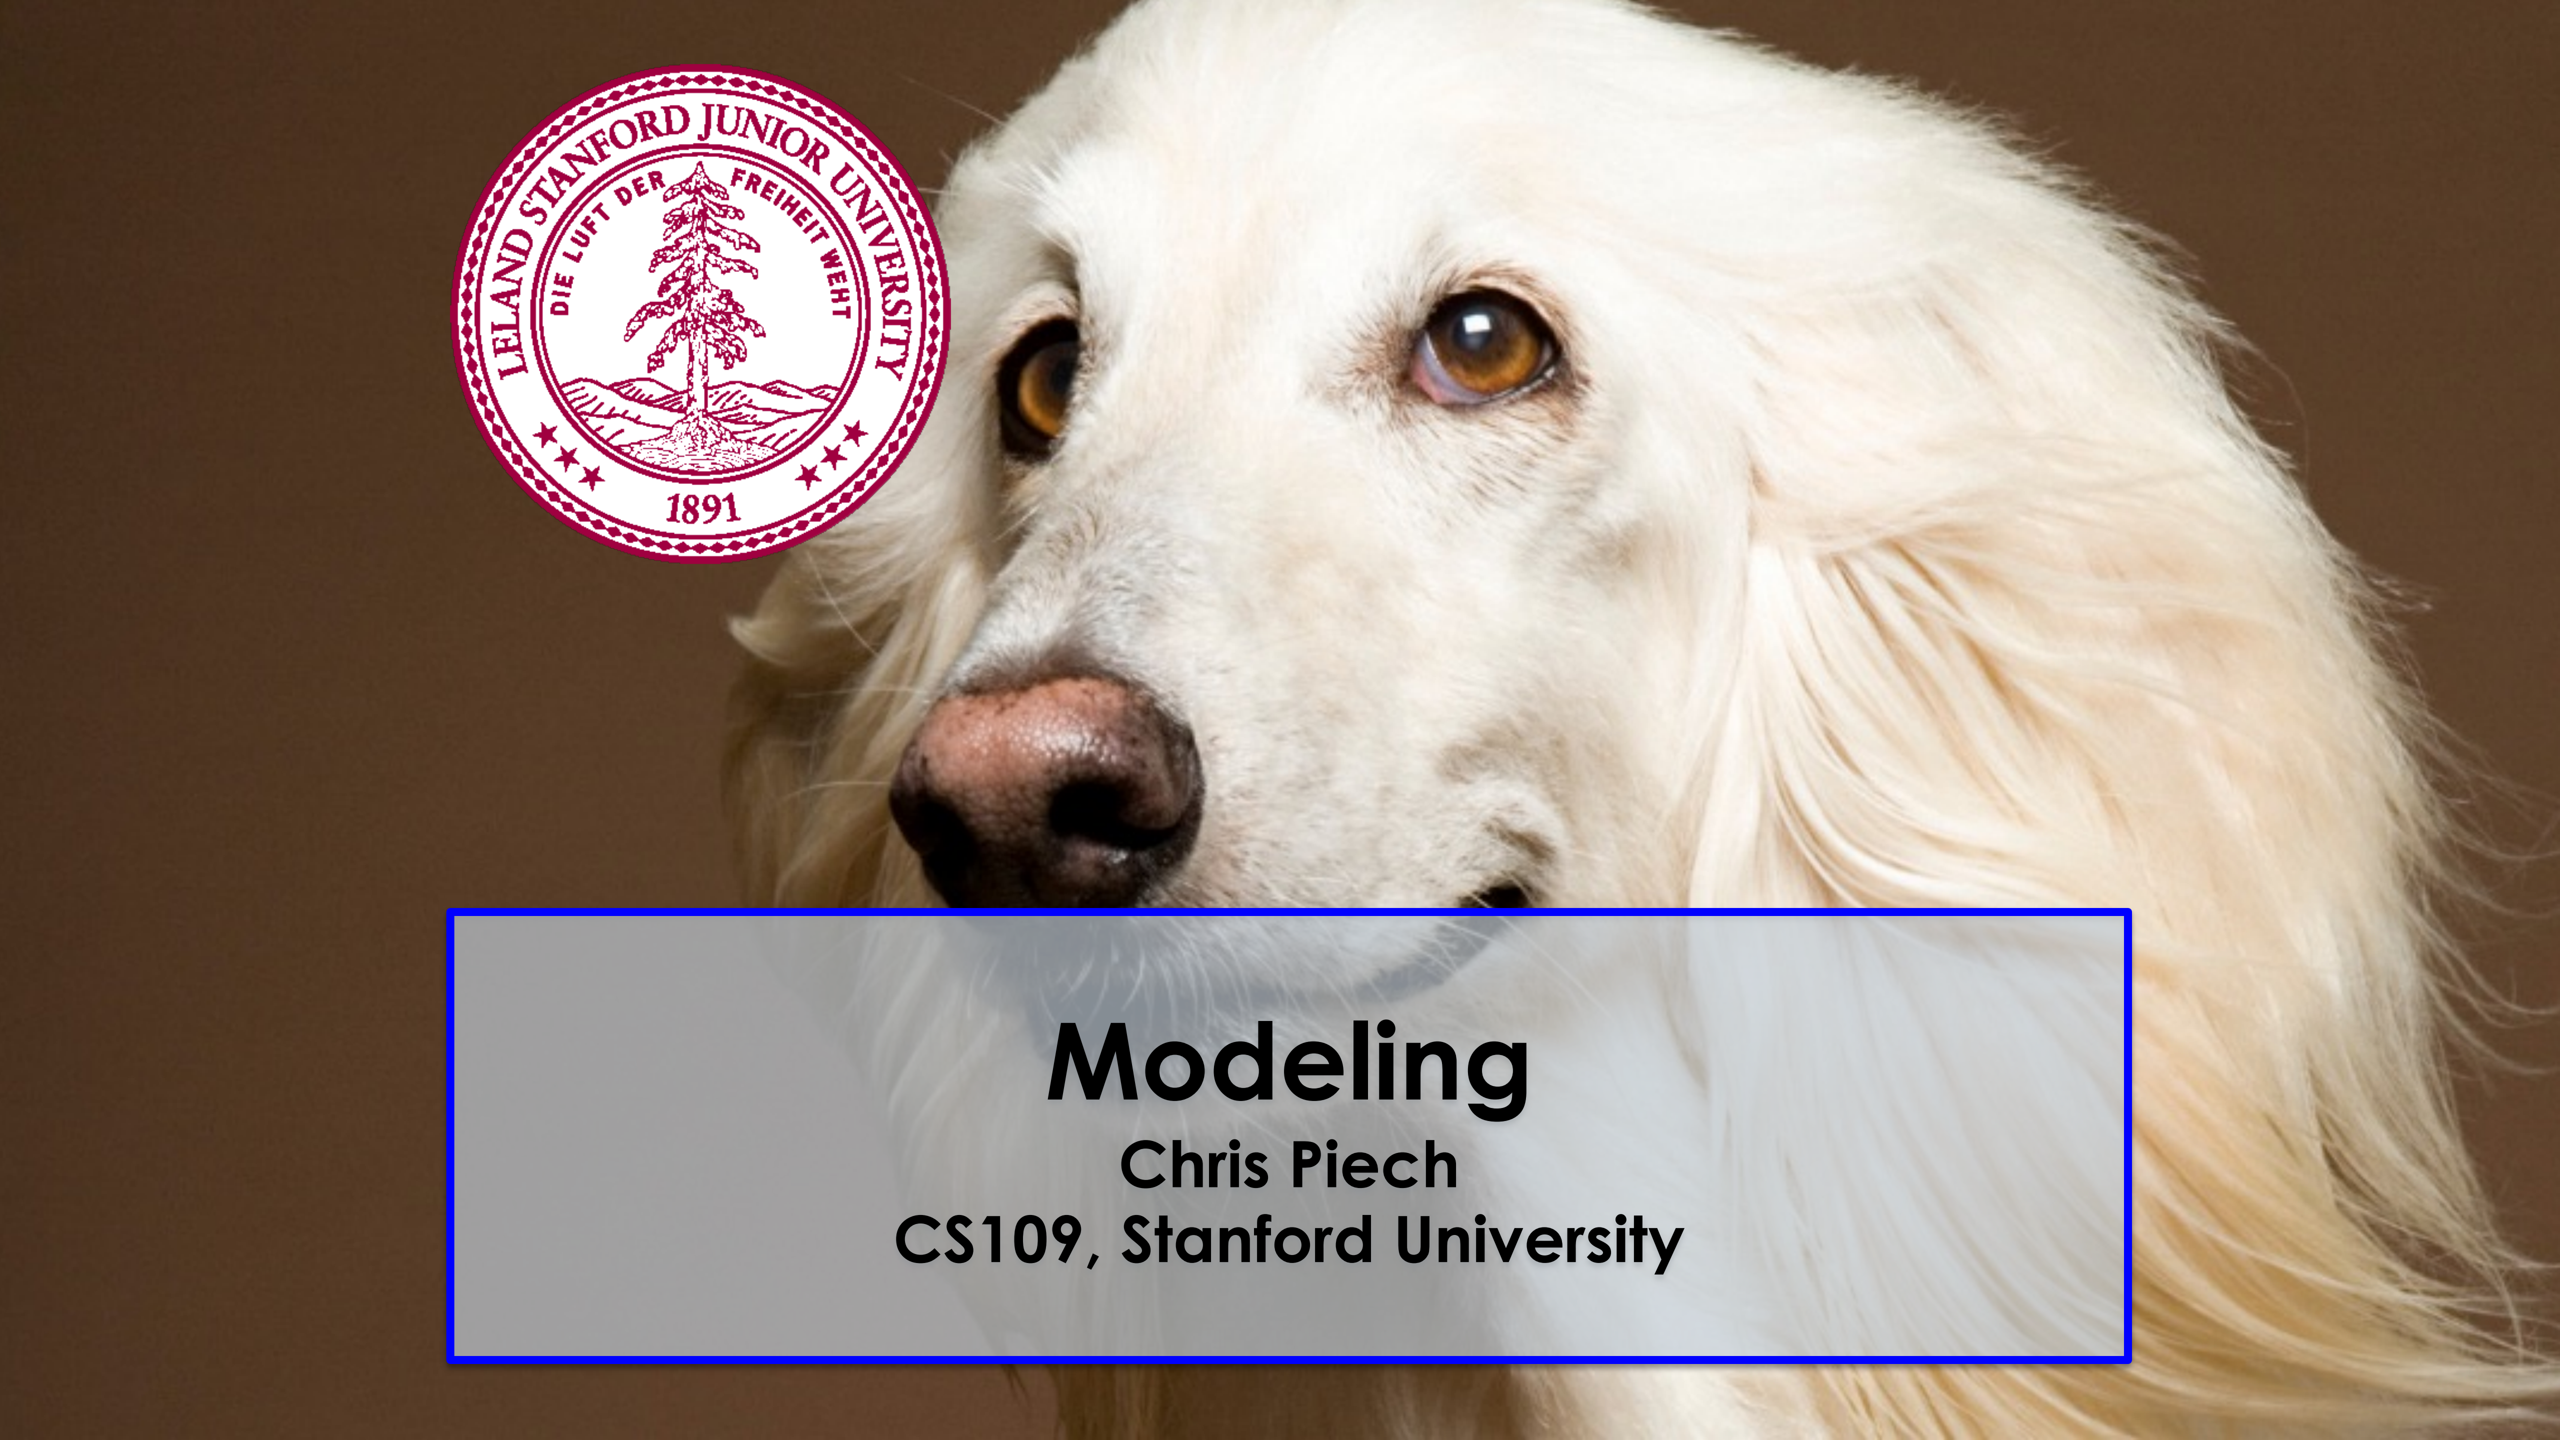

Modeling

Chris Piech
CS109, Stanford University

Review

Last Week: Joint Distributions

Joint Distribution *noun*

The probability of a simultaneous assignment to ***all*** the random variables in a probabilistic model.

Eg:

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

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

$$P(X = x, Y = y, \cdots, Z = z)$$

Last Friday, Monday: Inference

Inference *noun*

An updated belief about a random variable (or multiple) based on conditional knowledge regarding another random variable (or multiple) in a probabilistic model.

TLDR: conditional probability with random variables.

Inference with Continuous

Q: At birth, girl elephant weights are distributed as a Gaussian with mean = 160kg, std = 7kg. At birth, boy elephant weights are distributed as a Gaussian with mean = 165kg, std = 3kg. All you know about a newborn elephant is that it is 163kg. What is the probability that it is a girl?

Model Shown Graphically

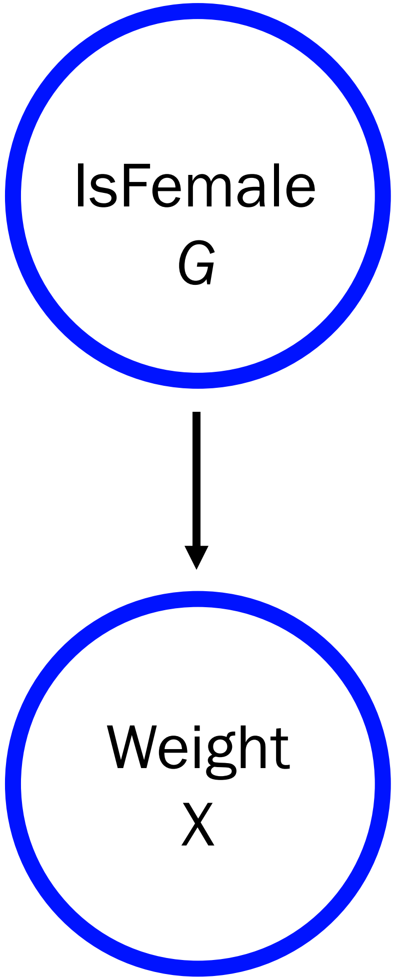

G = 1 is Bern(p = 0.5)

X | G = 1 is N(μ = 160, $\sigma^2 = 7^2$)
X | G = 0 is N(μ = 165, $\sigma^2 = 3^2$)

Does this define the joint?

$$f(G = g, X = x)$$
$$= f(X = x|G = g)P(G = g)$$

Q: What is P(G = 1 | X = 163)

All the Bayes Belong to Us

M,N are discrete. X, Y are continuous

OG Bayes

$$p_{M|N}(m|n) = \frac{P_{N|M}(n|m)p_M(m)}{p_N(n)}$$

Mix Bayes #1

$$f_{X|N}(x|n) = \frac{P_{N|X}(n|x)f_X(x)}{P_N(n)}$$

Mix Bayes #2

$$p_{N|X}(n|x) = \frac{f_{X|N}(x|n)p_N(n)}{f_X(x)}$$

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

Learning Goals

1. Perspective on the artform of how to design probabilistic models
2. How to calculate Correlations
3. Use and verify Independence with Random Variables

Lets talk about how to make a model

Model Version #1: Python That Outputs a **Joint** Sample

Sample Baby Elephant

Sex: Female
Weight: 161kg

Model Version #2: Bayesian Network

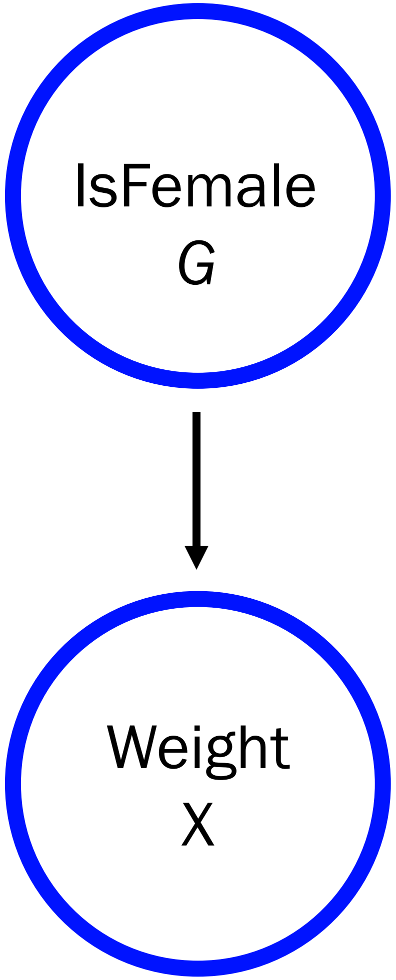

G = 1 is Bern(p = 0.5)

X | G = 1 is N(μ = 160, $\sigma^2 = 7^2$)

X | G = 0 is N(μ = 165, $\sigma^2 = 3^2$)

Does this define the joint?

$$f(G = g, X = x)$$

$$= f(X = x | G = g) P(G = g)$$

Why You Need a Model

WebMD®

Inference

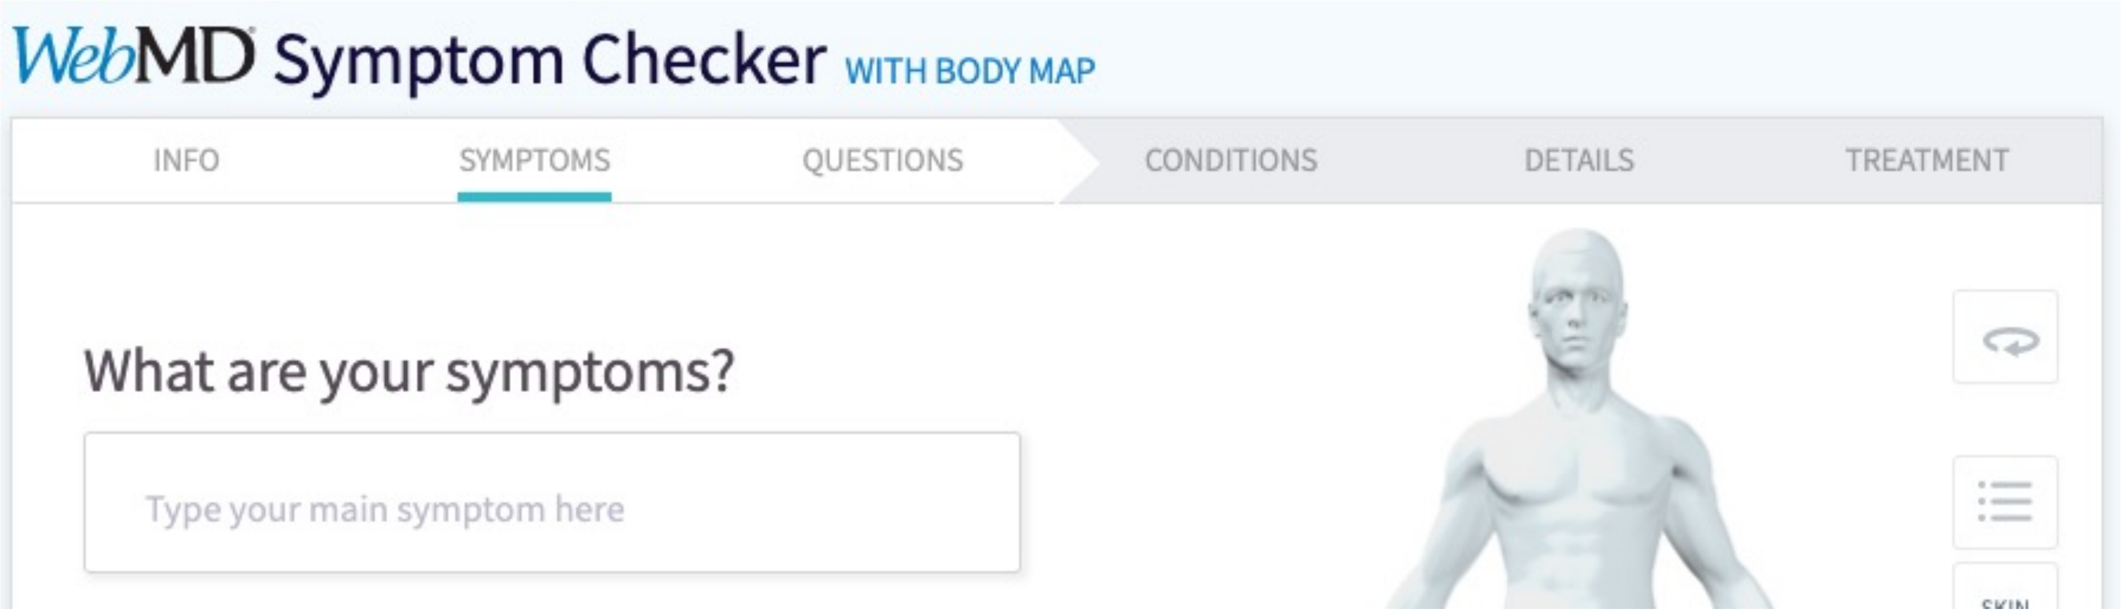
WebMD Symptom Checker WITH BODY MAP
INFO
SYMPTOMS
QUESTIONS
CONDITIONS
DETAILS
TREATMENT
What are your symptoms?
Type your main symptom here

Inference

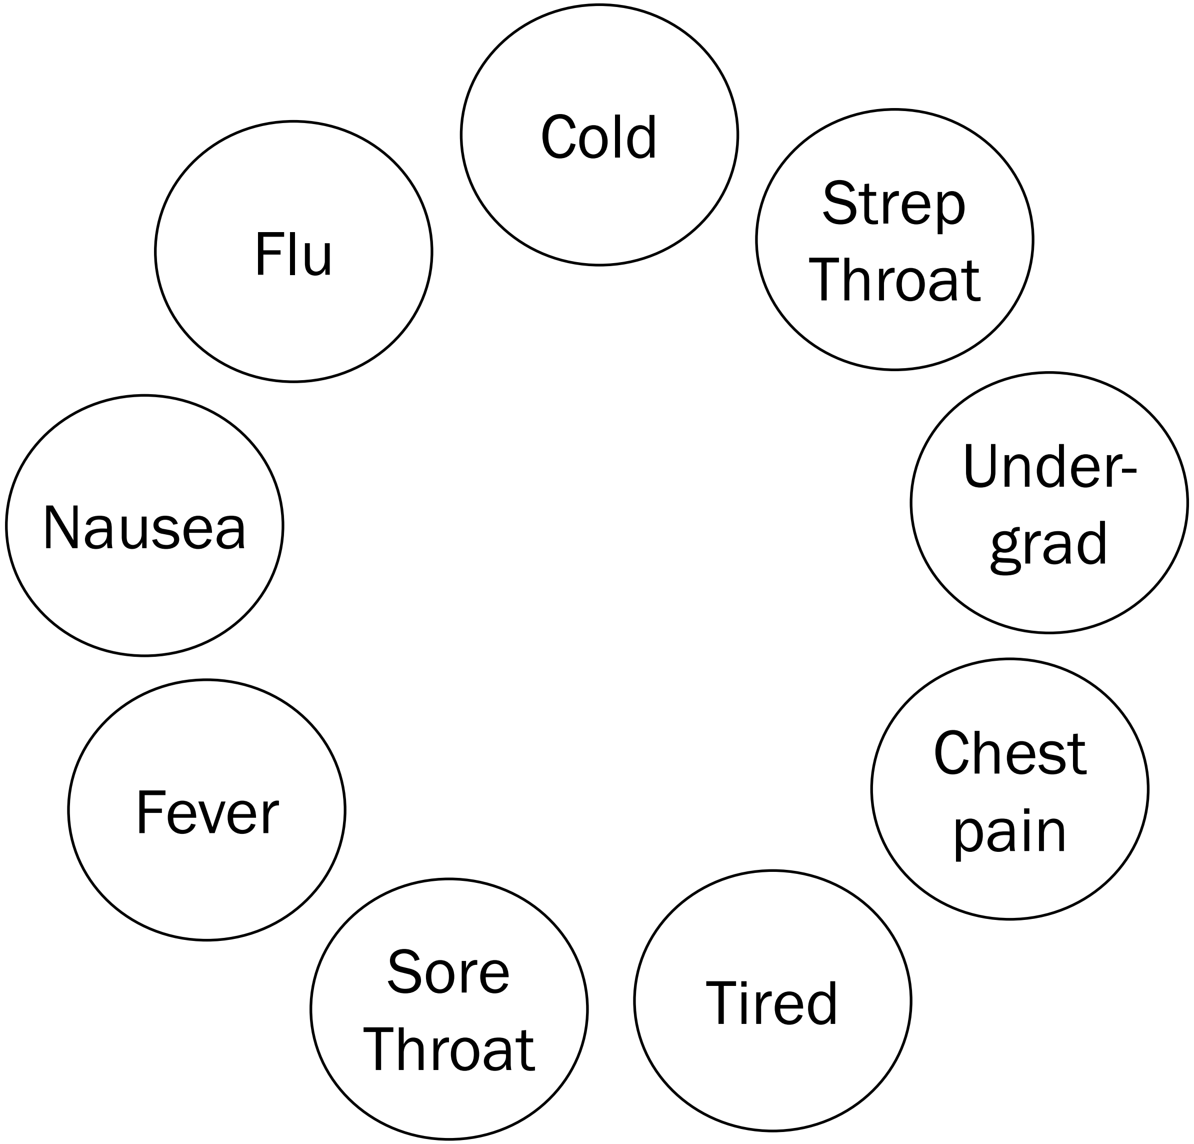

Inference question:

Given the values of some random variables, what are the conditional distributions of some other random variables?

Inference

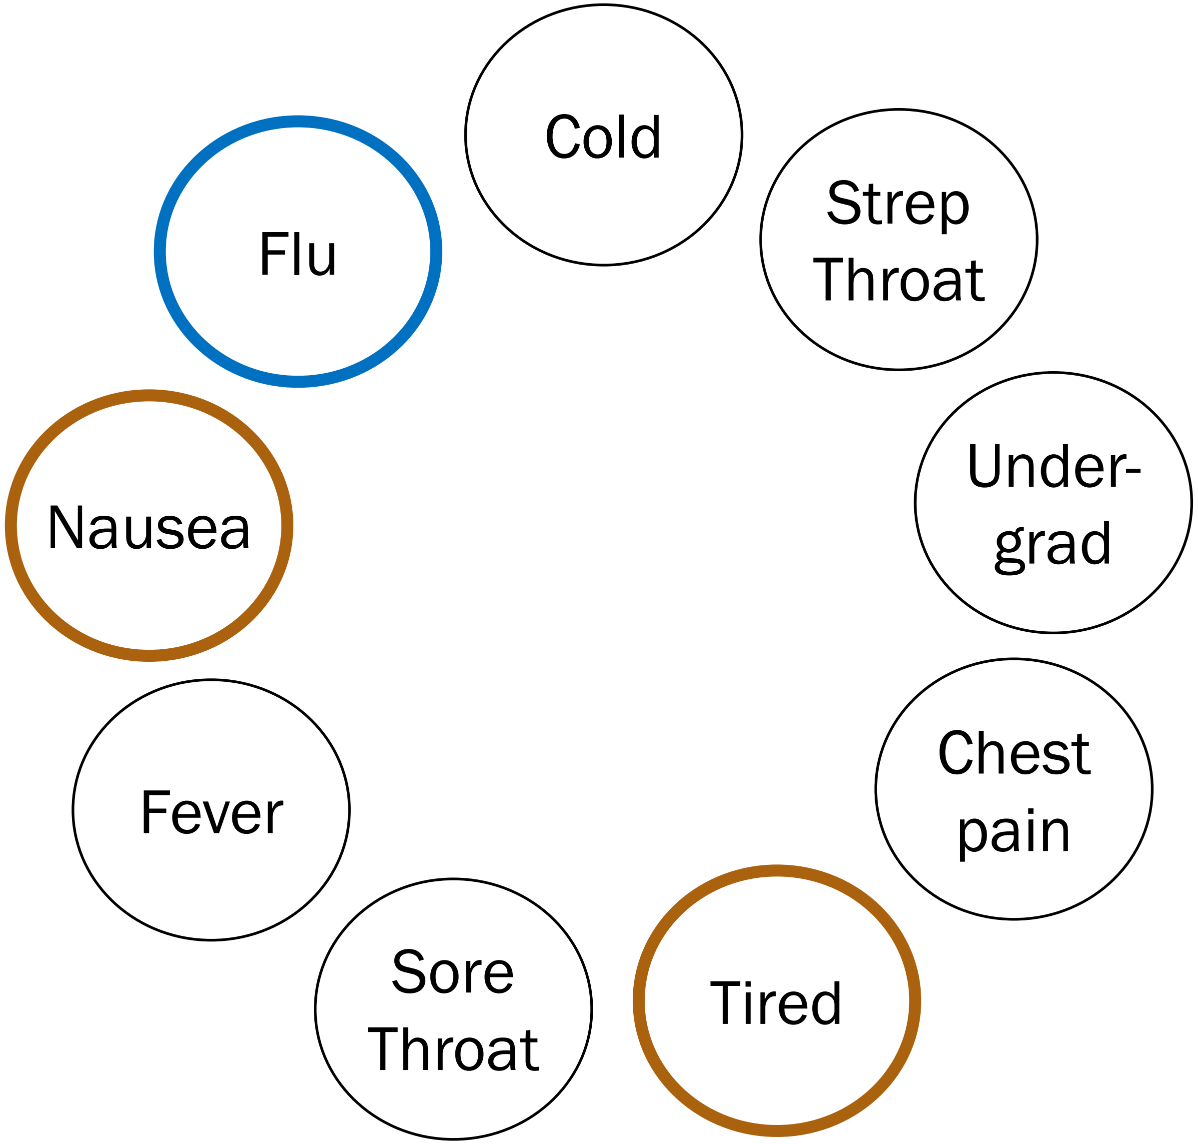

One inference question:

$$P(F = 1 | N = 1, T = 1)$$

$$= \frac{P(F = 1, N = 1, T = 1)}{P(N = 1, T = 1)}$$

Inference

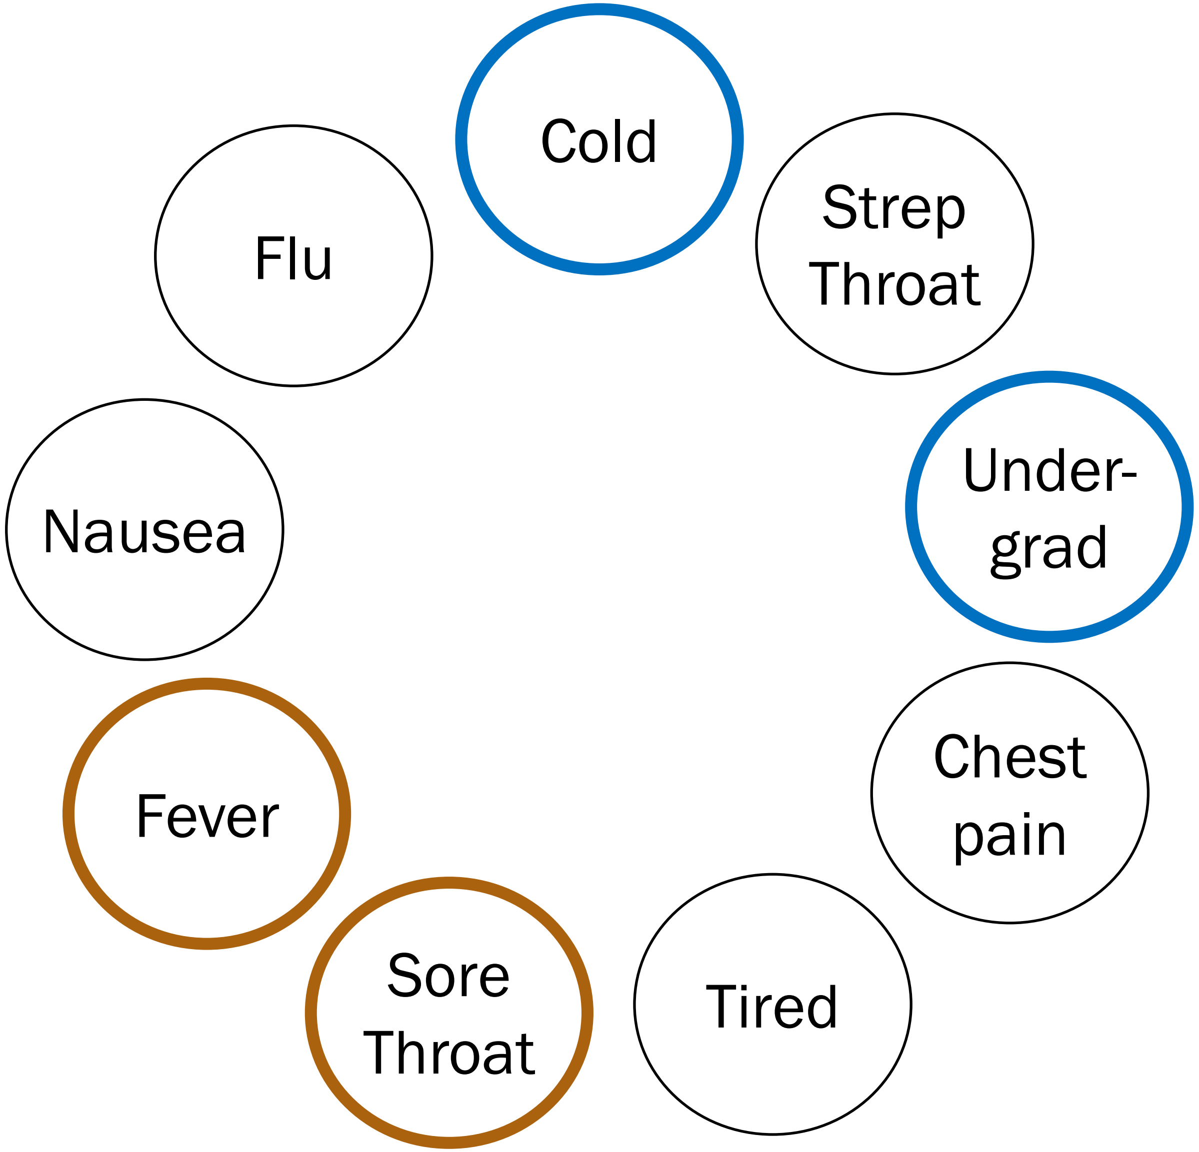

Another inference question:

$$P(C_o = 1, U = 1 | S = 0, F_e = 0)$$

$$= \frac{P(C_o = 1, U = 1, S = 0, F_e = 0)}{P(S = 0, F_e = 0)}$$

Inference

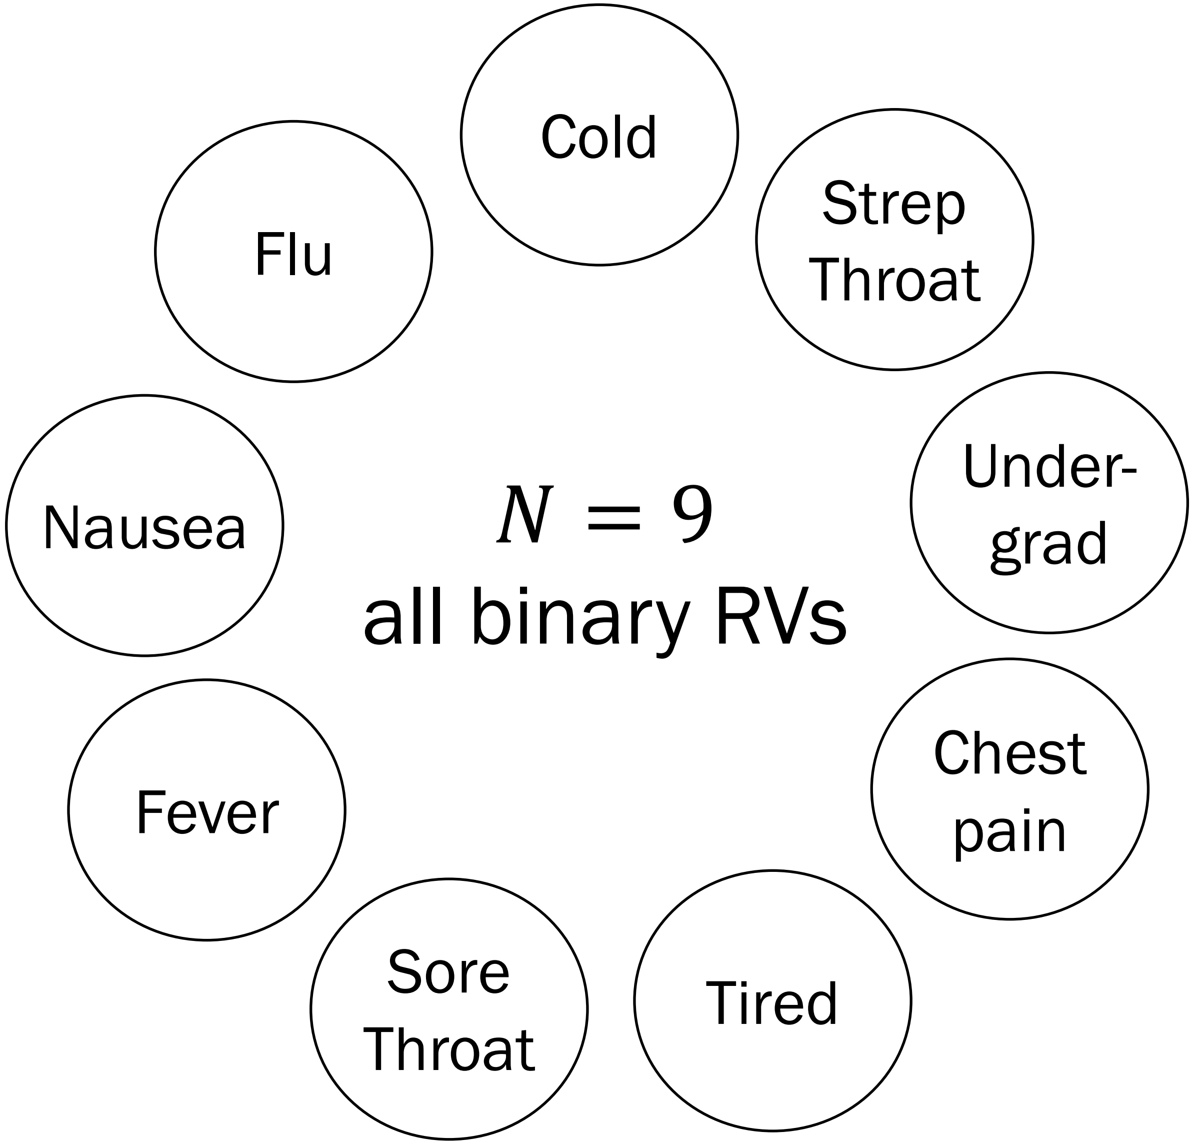

If we knew the **joint distribution**, we can answer all probabilistic inference questions.

What is the size of the joint probability table?

A. 2^{N-1} entries
B. N^2 entries
C. 2^N entries
D. None/other/don't know

Inference

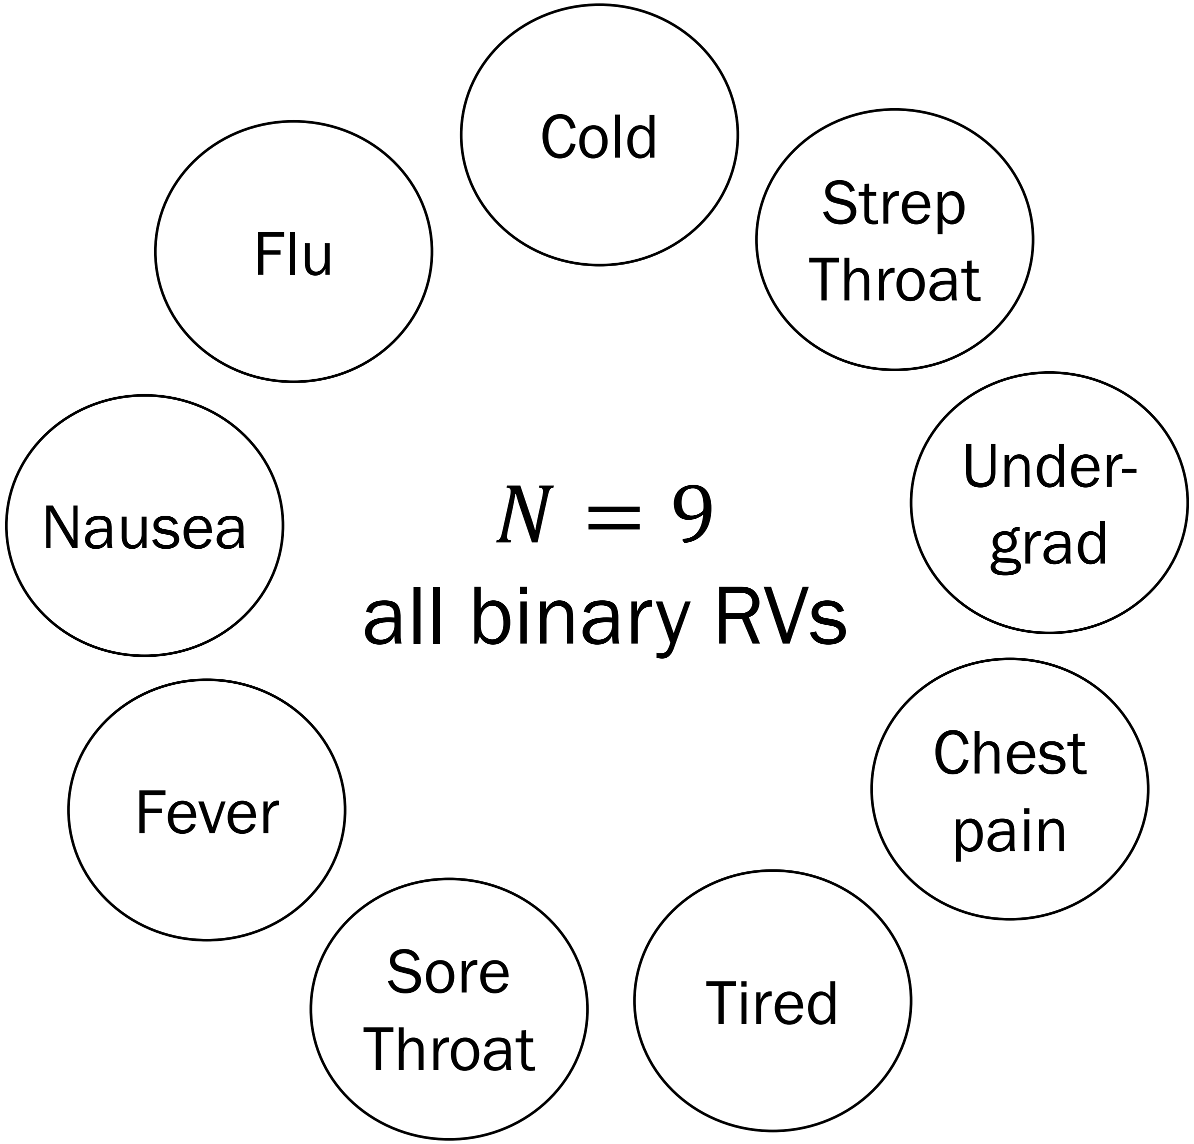

If we knew the **joint distribution**, we can answer all probabilistic inference questions.

What is the size of the joint probability table?

A. 2^{N-1} entries
B. N^2 entries
C. 2^N entries
D. None/other/don't know

Naively specifying a joint distribution is, in general, intractable.

N can be large…

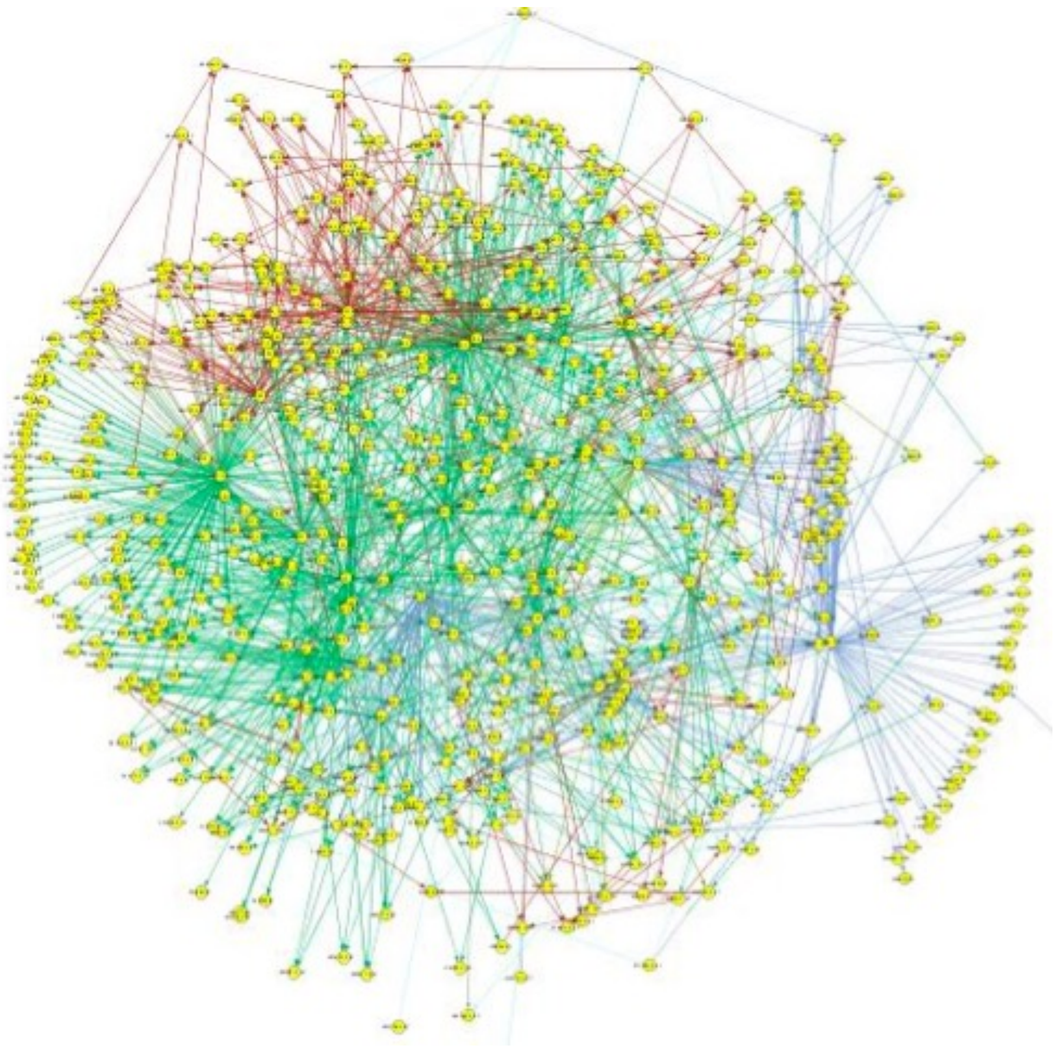

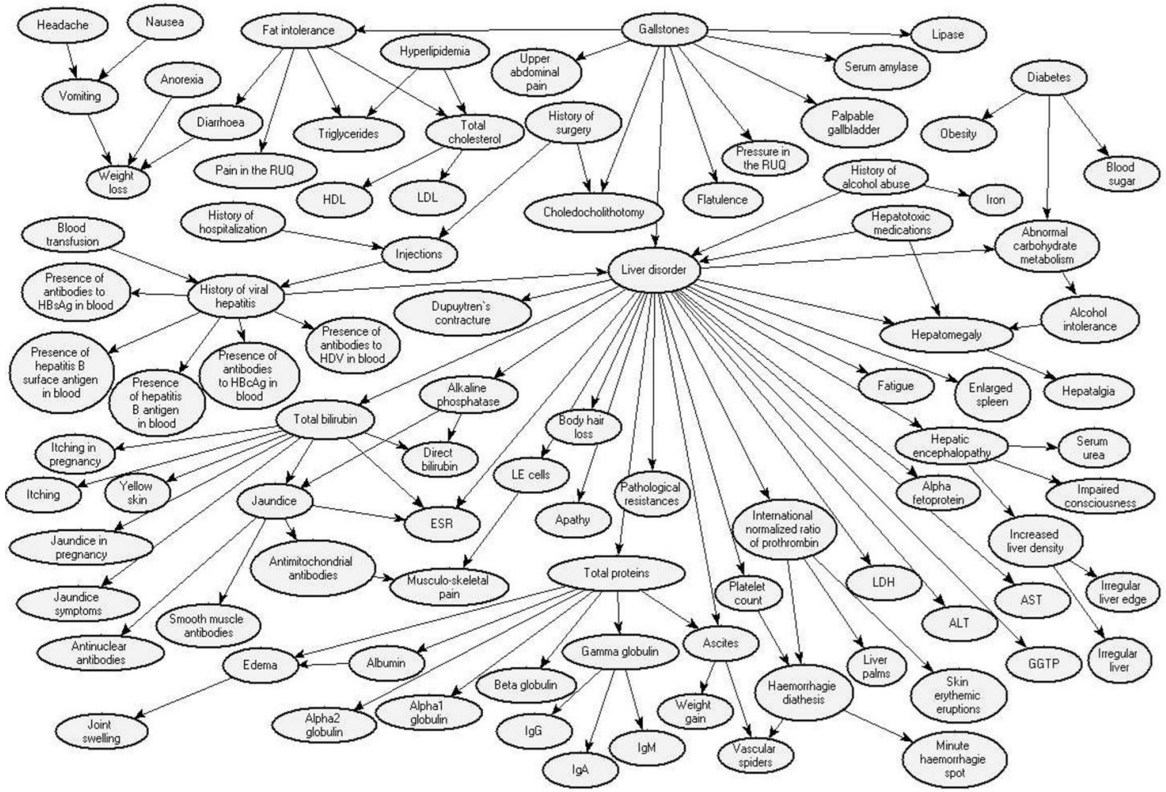

Bayesian Networks

A simpler WebMD

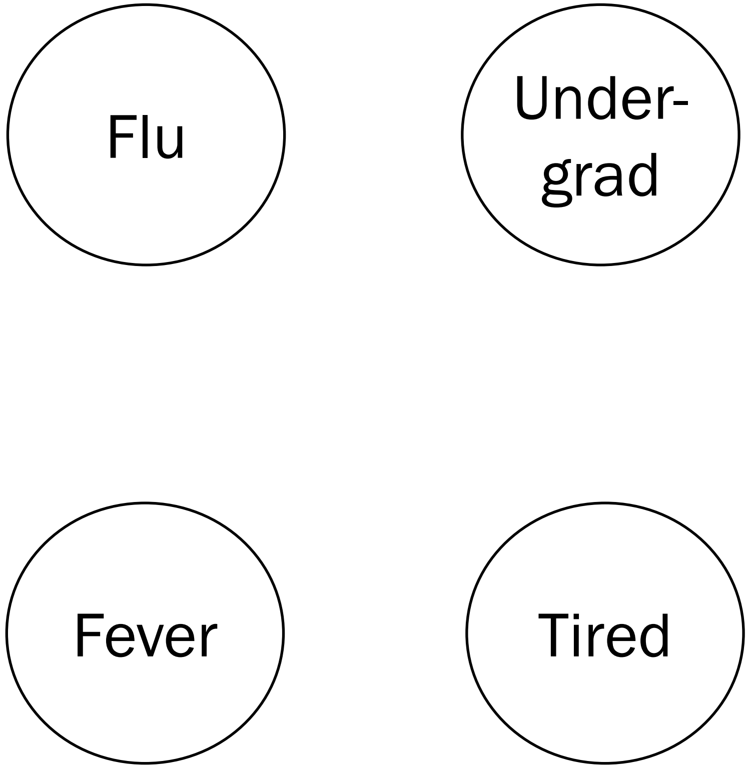

Great! Just specify $2^4 = 16$ joint probabilities…?

$$P(F_{lu} = a, F_{ev} = b, U = c, T = d)$$

What would a Stanford flu expert do?

Describe the joint distribution using causality!

Constructing a Bayesian Network

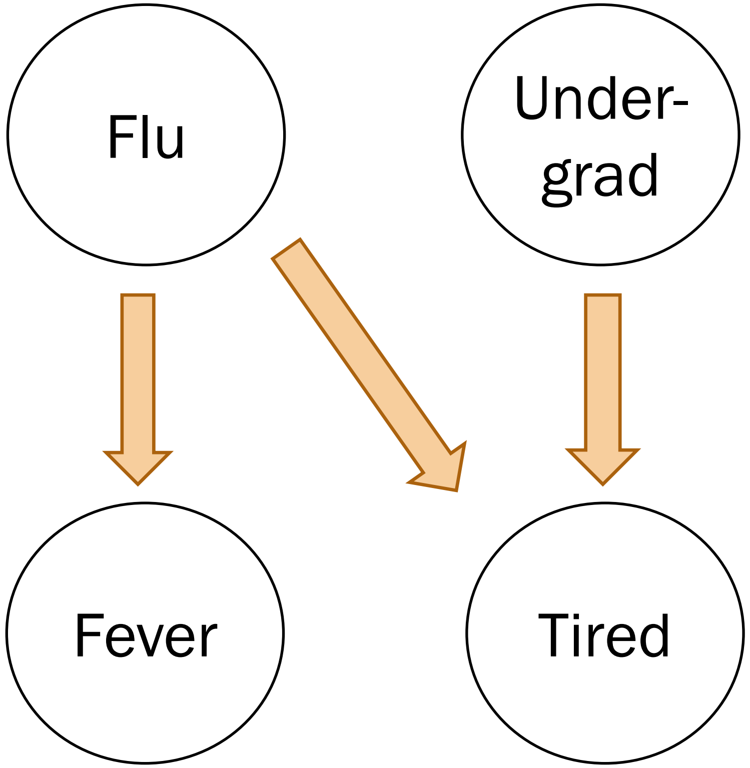

What would a Stanford flu expert do?

✅ 1. Describe the joint distribution using causality.
2. Provide P(values|parents) for each random variable
3. Implicitly assumes independences.

Constructing a Bayesian Network

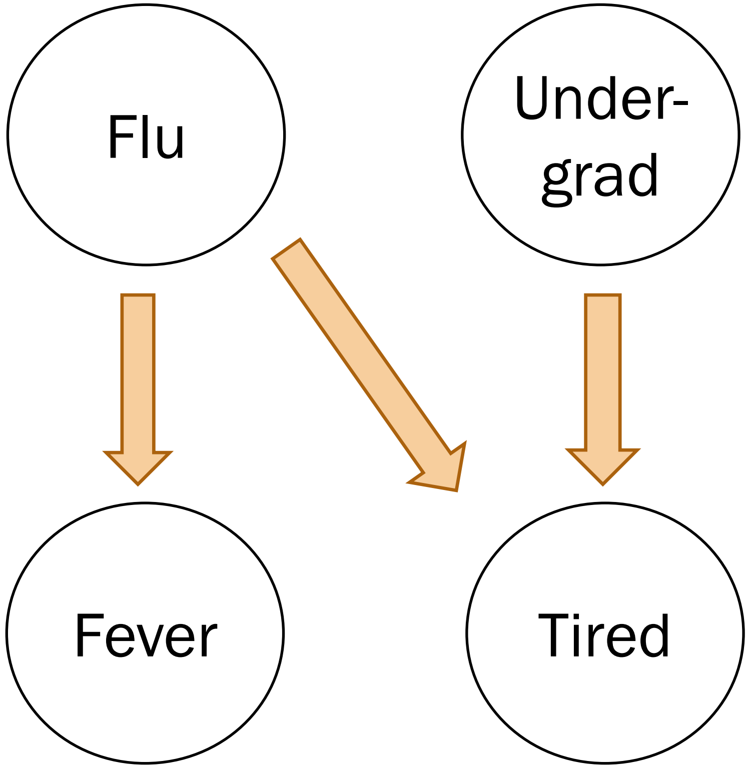

What would a Stanford flu expert do?

✅ 1. Describe the joint distribution using causality.
2. Provide P(values|parents) for each random variable
3. Implicitly assumes independences.

Constructing a Bayesian Network

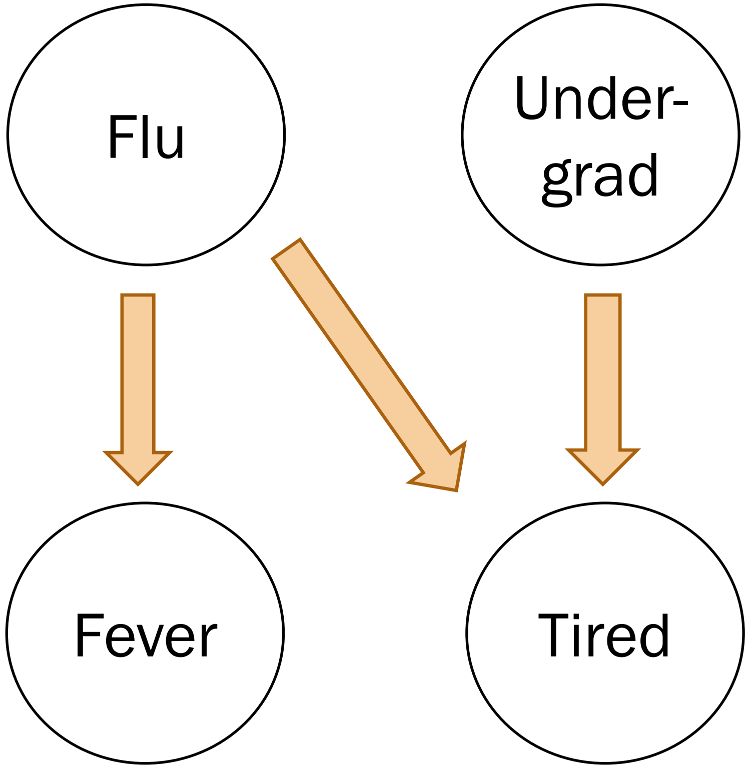

$P(T = 1|F_{lu} = 0, U = 0)$
$P(T = 1|F_{lu} = 0, U = 1)$
$P(T = 1|F_{lu} = 1, U = 0)$
$P(T = 1|F_{lu} = 1, U = 1)$

In a Bayesian Network,

Each random variable is caused by its parents. Def P(node | parents)

- Node: random variable
- Directed edge: causality

Examples:

- $P(F_{lu} = 1)$
- $P(U = 0)$
- $P(F_{ev} = 1|F_{lu} = 1)$, $P(F_{ev} = 1|F_{lu} = 0)$
- $P(T = 1|F_{lu} = 0, U = 0)$ …

Constructing a Bayesian Network

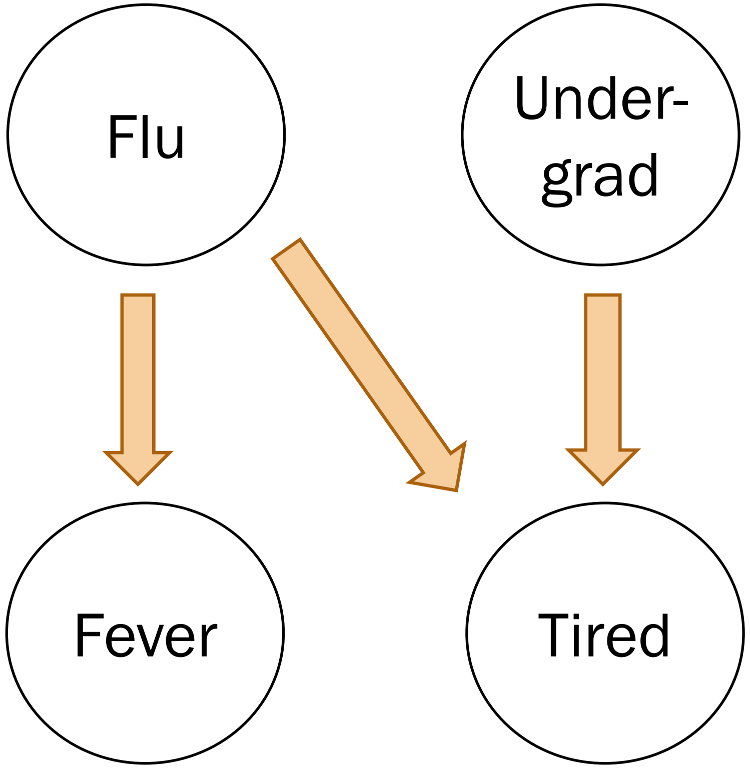

What would a Stanford flu expert do?

✅ 1. Describe the joint distribution using causality.

✅ 2. Provide P(values|parents) for each random variable

✅ 3. Implicitly assumes independences.

Constructing a Bayesian Network

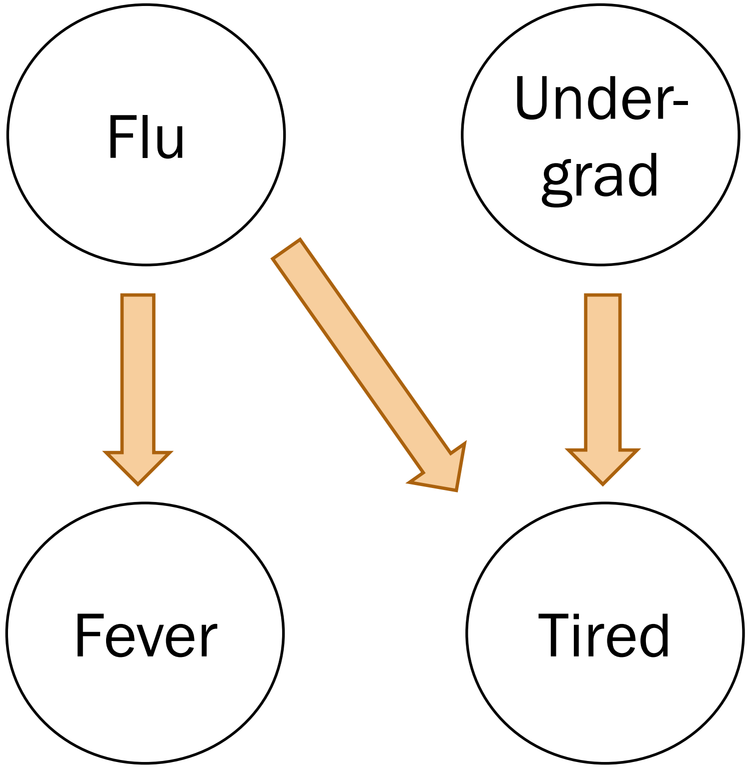

This model assumes that Flu and being an Undergraduate are independent.

Neat trick: it also assumes that fever and tired are conditionally independent given Flu.

You need to tell a generative story. You do not need to be able to reason about all the implied independencies

Bug: Constructing a Bayesian Network

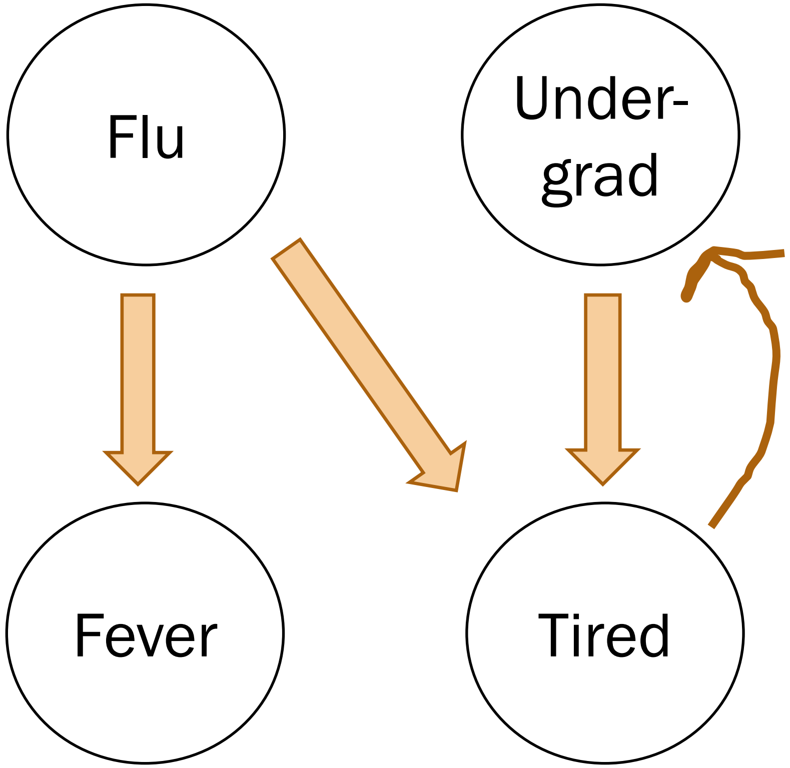

Must by acyclic!

Bayesian Network Assumption:

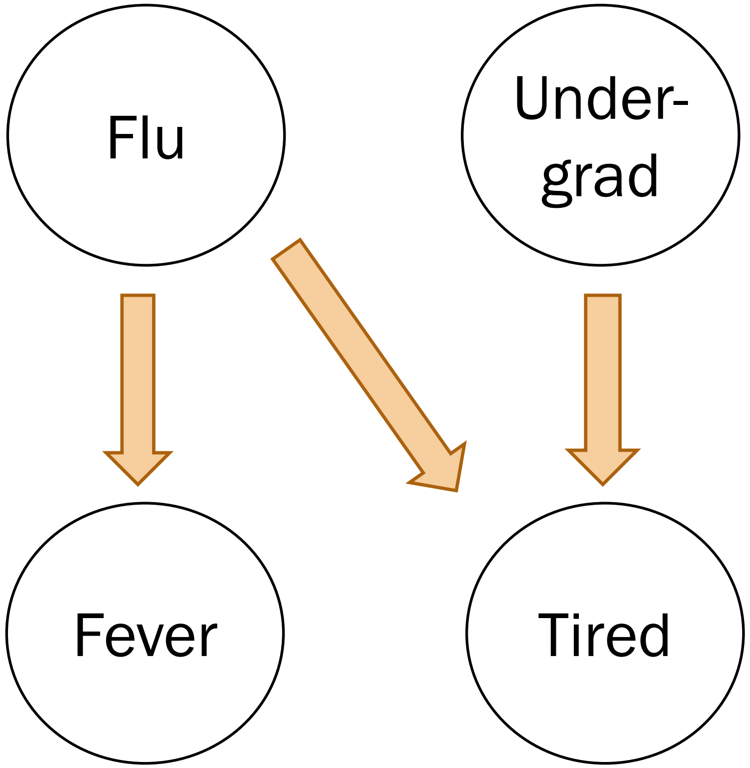

$$P(\text{Joint}) = P(X_1 = x_1 \dots X_n = x_n)$$
$$= \prod_i P(X_i = x_i | \text{parents of } X_i)$$

Model Version #1: Python That Outputs a **Joint** Sample

Sample Patient

Flu: 1
Undergrad: 0
Fever: 1
Tired: 0

Bayes Nets tell a generative story.

This leads to many independence assumptions

Makes it **tractable** to represent the joint

Inference via math

$P(F_{lu} = 1) = 0.1$ $P(U = 1) = 0.8$

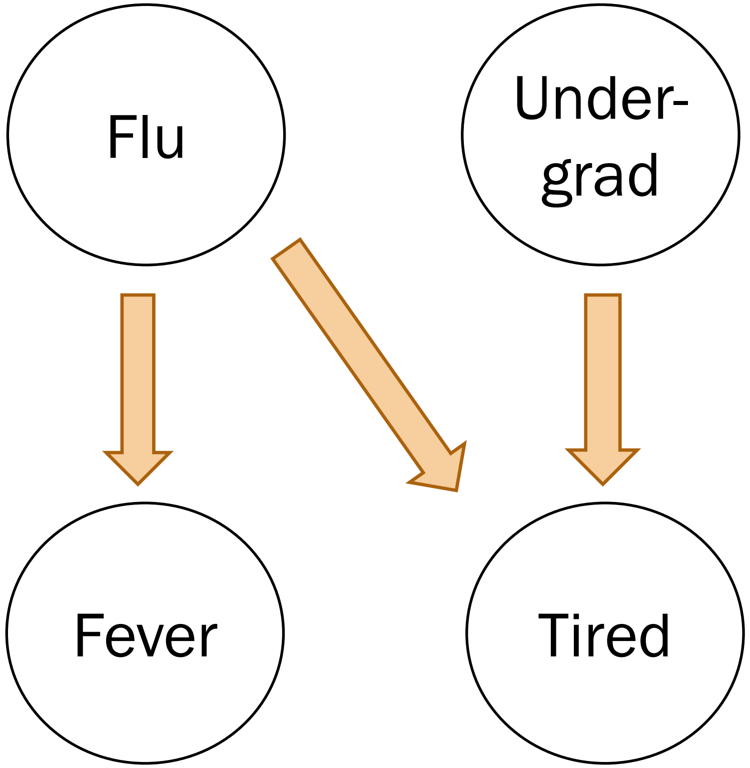

$P(F_{ev} = 1 | F_{lu} = 1) = 0.9$
$P(F_{ev} = 1 | F_{lu} = 0) = 0.05$

$P(T = 1 | F_{lu} = 0, U = 0) = 0.1$
$P(T = 1 | F_{lu} = 0, U = 1) = 0.8$
$P(T = 1 | F_{lu} = 1, U = 0) = 0.9$
$P(T = 1 | F_{lu} = 1, U = 1) = 1.0$

1. $P(F_{lu} = 0, U = 1, F_{ev} = 0, T = 1)$?

   Compute joint probabilities using chain rule.

Independence of RVs

Independent discrete RVs

Recall the definition of independent events E and F:

$$P(EF) = P(E)P(F)$$

Two discrete random variables X and Y are **independent** if:

for all x, y:

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

Different notation, same idea:

$$p_{X,Y}(x, y) = p_X(x)p_Y(y)$$

- Intuitively: knowing value of X tells us nothing about the distribution of Y (and vice versa)
- If two variables are not independent, they are called **dependent**.

Dice (after all this time, still our friends)

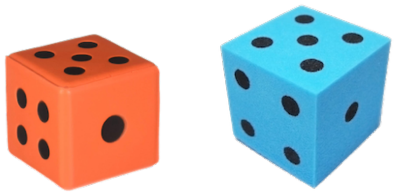

Let: D_1 and D_2 be the outcomes of two rolls
$S = D_1 + D_2$, the sum of two rolls

- Each roll of a fair, 6-sided die is an independent trial.
- Random variables D_1 and D_2 are independent.

1. Are events $(D_1 = 1)$ and $(S = 7)$ independent?
2. Are events $(D_1 = 1)$ and $(S = 5)$ independent?
3. Are random variables D_1 and S independent?

Dice (after all this time, still our friends)

Let: D_1 and D_2 be the outcomes of two rolls
$S = D_1 + D_2$, the sum of two rolls

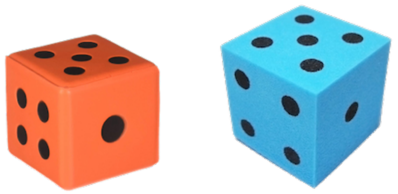

- Each roll of a 6-sided die is an independent trial.
- Random variables D_1 and D_2 are independent.

1. Are events $(D_1 = 1)$ and $(S = 7)$ independent? ✅
2. Are events $(D_1 = 1)$ and $(S = 5)$ independent? ❌
3. Are random variables D_1 and S independent? ❌

All events $(X = x, Y = y)$ must be independent for X, Y to be independent RVs.

Can I discover independence from data?

C15 fx 3

	A	B	C	D	E	F	G	H
1	Music	Dance	Folk	Country	Classical music	Musical	Pop	Rock
2	5	2	1	2	2	1	5	5
3	4	2	1	1	1	2	3	5
4	5	2	2	3	4	5	3	5
5	5	2	1	1	1	1	2	2
6	5	4	3	2	4	3	5	3
7	5	2	3	2	3	3	2	5
8	5	5	3	1	2	2	5	3
9	5	3	2	1	2	2	4	5
10	5	3	1	1	2	4	3	5
11	5	2	5	2	2	5	3	5
12	5	3	2	1	2	3	4	3
13	5	1	1	1	4	1	2	5
14	5	1	2	1	4	3	3	5
15	5	5	3	2	1	5	5	2
16	5	2	1	1	2	3	4	5
17	1	2	2	3	4	3	3	5
18	5	3	1	1	1	2	4	4
19	5	3	3	3	2	2	4	4
20	5	5	4	3	4	5	5	4
21	5	3	3	2	4	2	2	4
22	5	3	2	3	4	3	2	5
23	5	1	1	3	2	2	2	5
24	5	3	2	3	3	3	4	
25	5	4	2	2	2	4	4	5
26	5	3	1	1	4	3	3	5
27	5	4	2	1	2	3	5	1
28	5	5	5	4	5	3	4	4
29	4	3	4	1	3	2	2	4
30	5	5	1	1	1	1	3	4
31	5	3	4	2	3	3	3	4
32	4	4	3	3	3	3	4	4
33	4	4	1	3	2	3	5	3
34	5	3	1	3	2	3	3	4
35	5	2	2	3	4	5	4	3

music

Why is it harder to find independences here than for bat DNA expression?

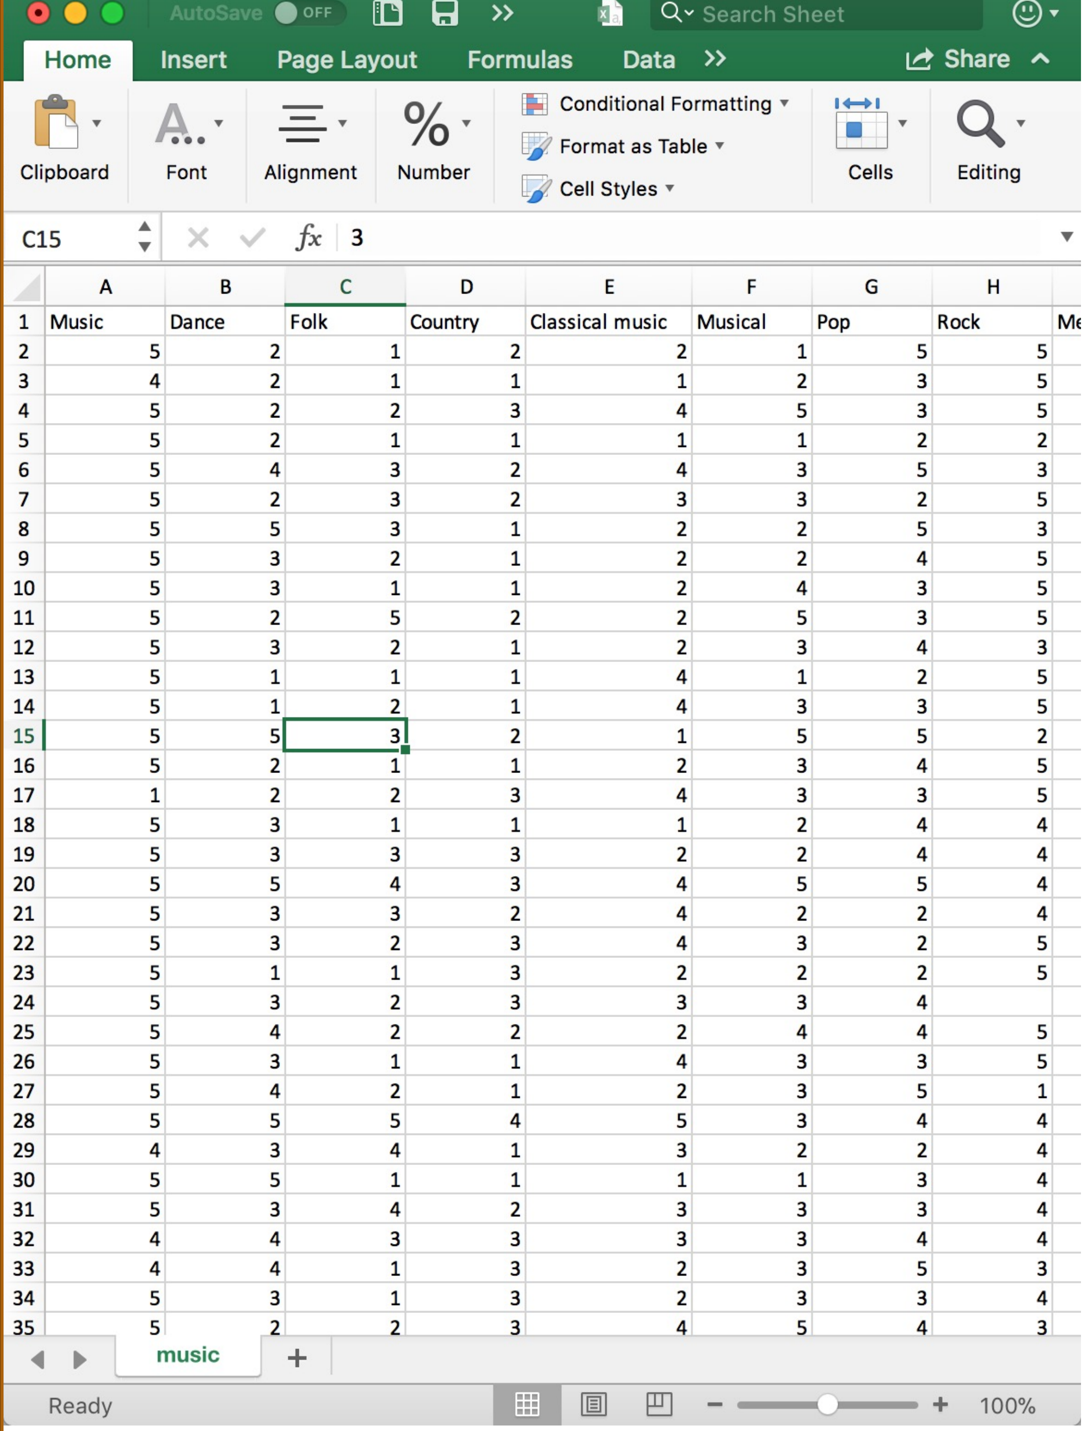

	A	B	C	D	E	F	G	H
1	Music	Dance	Folk	Country	Classical music	Musical	Pop	Rock
2	5	2	1	2	2	1	5	5
3	4	2	1	1	1	2	3	5
4	5	2	2	3	4	5	3	5
5	5	2	1	1	1	1	2	2
6	5	4	3	2	4	3	5	3
7	5	2	3	2	3	3	2	5
8	5	5	3	1	2	2	5	3
9	5	3	2	1	2	2	4	5
10	5	3	1	1	2	4	3	5
11	5	2	5	2	2	5	3	5
12	5	3	2	1	2	3	4	3
13	5	1	1	1	4	1	2	5
14	5	1	2	1	4	3	3	5
15	5	5	3	2	1	5	5	2
16	5	2	1	1	2	3	4	5
17	1	2	2	3	4	3	3	5
18	5	3	1	1	1	2	4	4
19	5	3	3	3	2	2	4	4
20	5	5	4	3	4	5	5	4
21	5	3	3	2	4	2	2	4
22	5	3	2	3	4	3	2	5
23	5	1	1	3	2	2	2	5
24	5	3	2	3	3	3	4	
25	5	4	2	2	2	4	4	5
26	5	3	1	1	4	3	3	5
27	5	4	2	1	2	3	5	1
28	5	5	5	4	5	3	4	4
29	4	3	4	1	3	2	2	4
30	5	5	1	1	1	1	3	4
31	5	3	4	2	3	3	3	4
32	4	4	3	3	3	3	4	4
33	4	4	1	3	2	3	5	3
34	5	3	1	3	2	3	3	4
35	5	2	2	3	4	5	4	3

Dance of the Covariance

Recall our Ebola Bats

Bat Data

Gene1	Gene2	Gene3	Gene4	Gene5	Trait
TRUE	FALSE	TRUE	TRUE	FALSE	FALSE
FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
TRUE	FALSE	TRUE	FALSE	FALSE	FALSE
TRUE	FALSE	TRUE	TRUE	TRUE	FALSE
FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
FALSE	FALSE	FALSE	TRUE	FALSE	FALSE
TRUE	FALSE	FALSE	TRUE	FALSE	FALSE
TRUE	FALSE	FALSE	TRUE	FALSE	FALSE
TRUE	FALSE	TRUE	FALSE	FALSE	FALSE
FALSE	TRUE	FALSE	TRUE	FALSE	FALSE
TRUE	TRUE	FALSE	TRUE	FALSE	FALSE
TRUE	FALSE	FALSE	TRUE	FALSE	FALSE
TRUE	FALSE	TRUE	TRUE	TRUE	FALSE
FALSE	FALSE	TRUE	TRUE	FALSE	FALSE
TRUE	FALSE	FALSE	TRUE	FALSE	FALSE
TRUE	FALSE	FALSE	TRUE	FALSE	FALSE
		…			
TRUE	FALSE	FALSE	TRUE	FALSE	FALSE

Expression Amount

Gene5	Trait
0.76	0.83
0.94	0.85
0.82	0.03
0.94	0.32
0.50	0.10
0.40	0.53
0.90	0.67
0.29	0.71
0.72	0.25
0.15	0.24
0.79	0.98
0.68	0.77
0.71	0.37
0.36	0.18
0.62	0.08
0.59	0.38
0.82	0.76

Spot The Difference

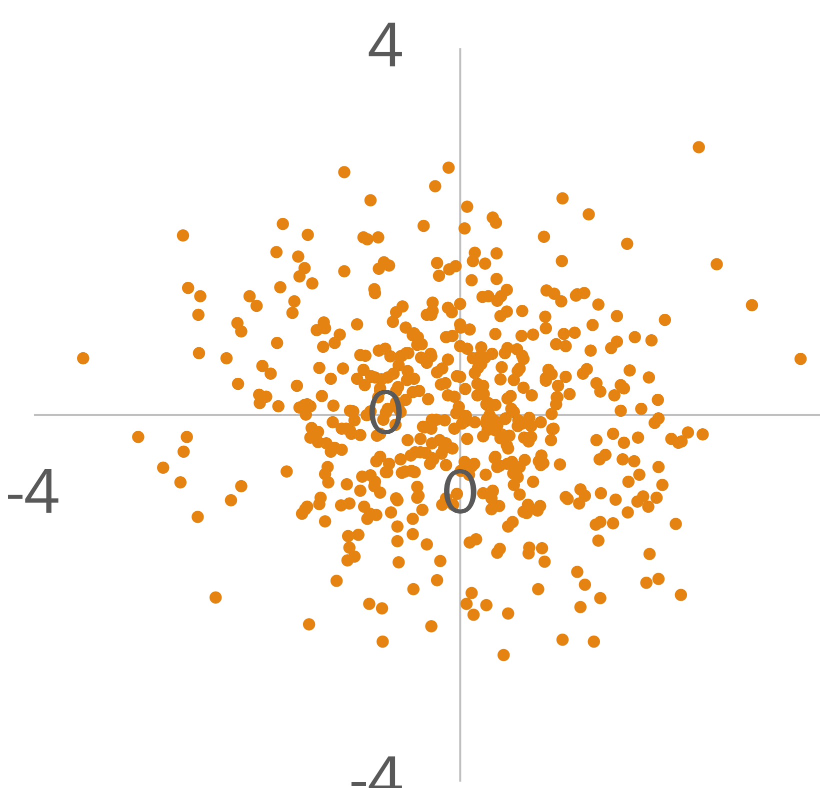

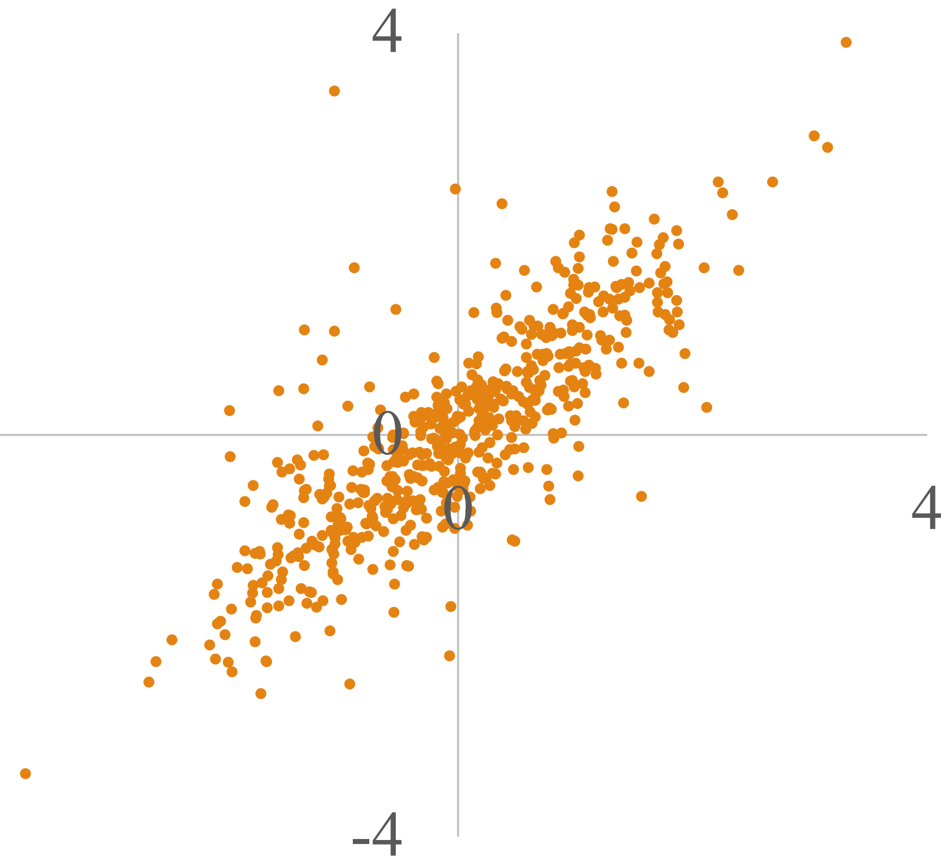

Spot The Difference

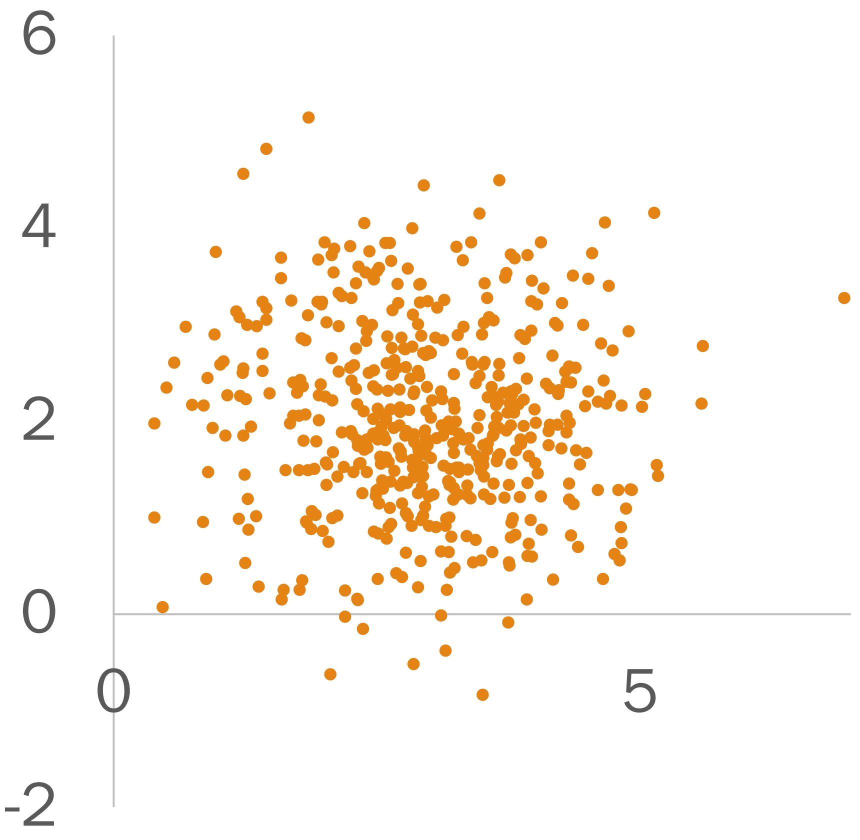

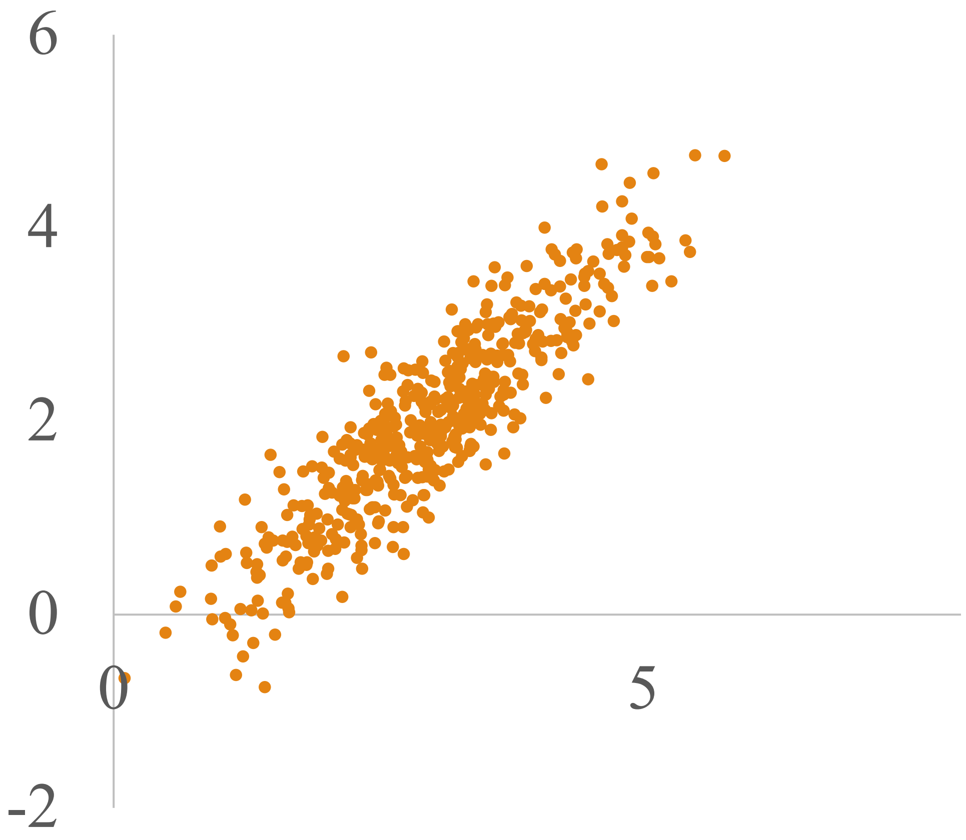

Vary Together

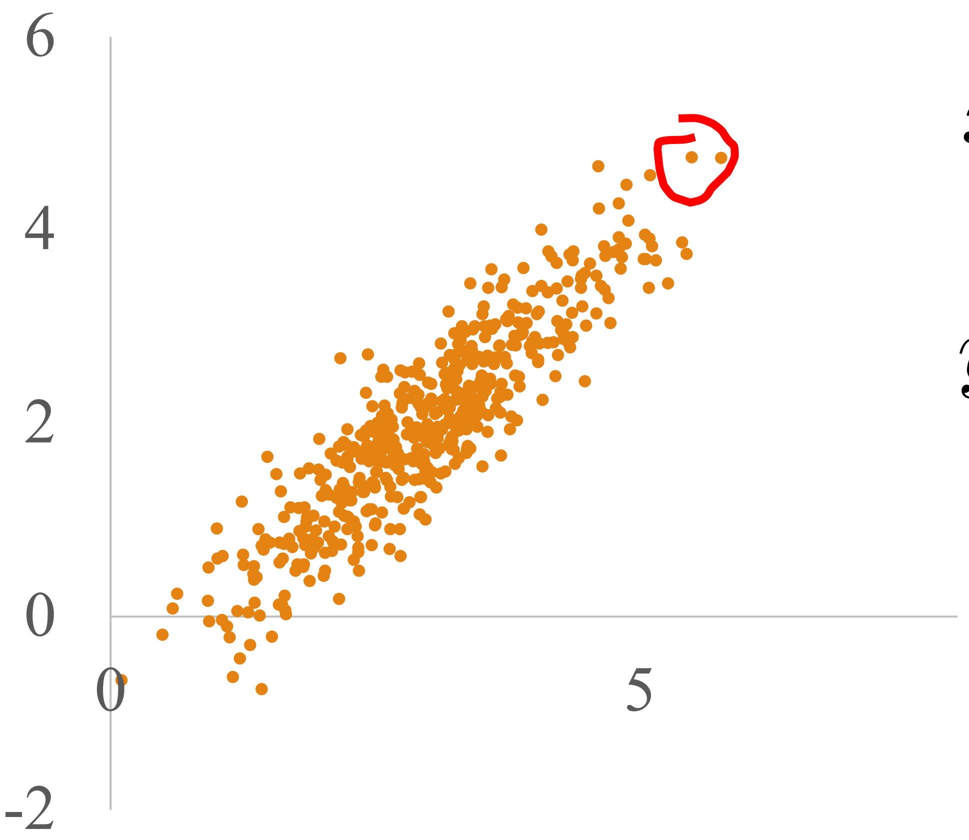

$$x - E[x] = 3$$

$$y - E[y] = 2.6$$

$$(x - E[x])(y - E[y]) = 7.8$$

Vary Together

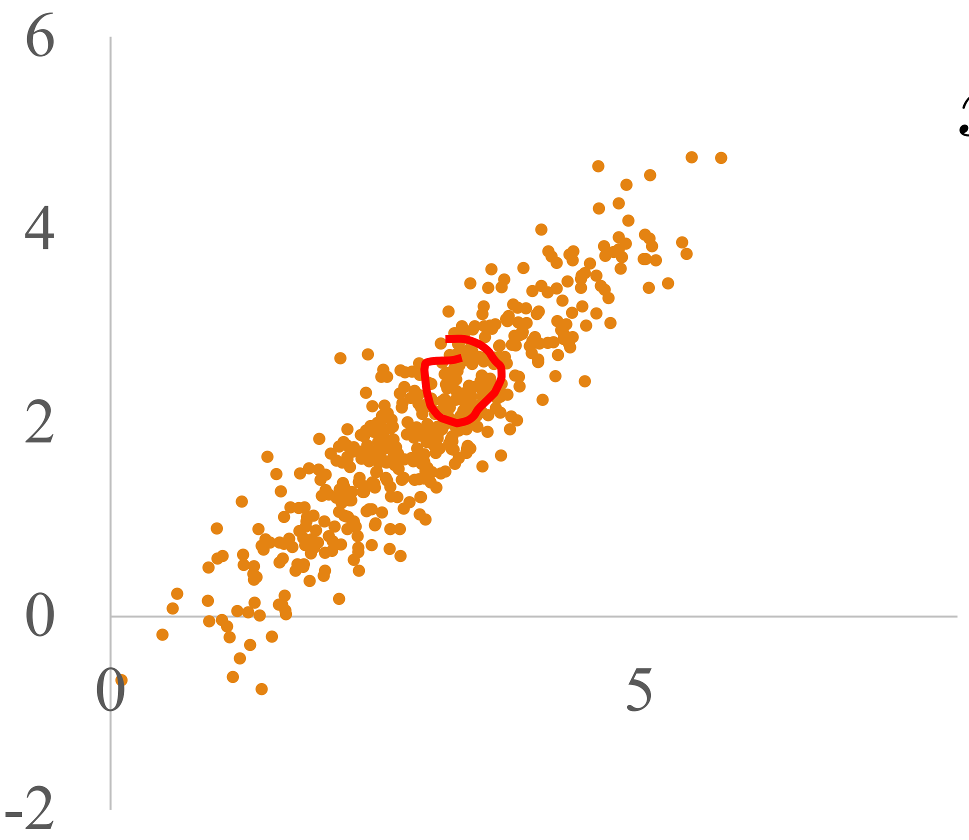

$$x - E[x] \approx 0$$

$$y - E[y] \approx 0$$

$$(x - E[x])(y - E[y]) = 0$$

Vary Together

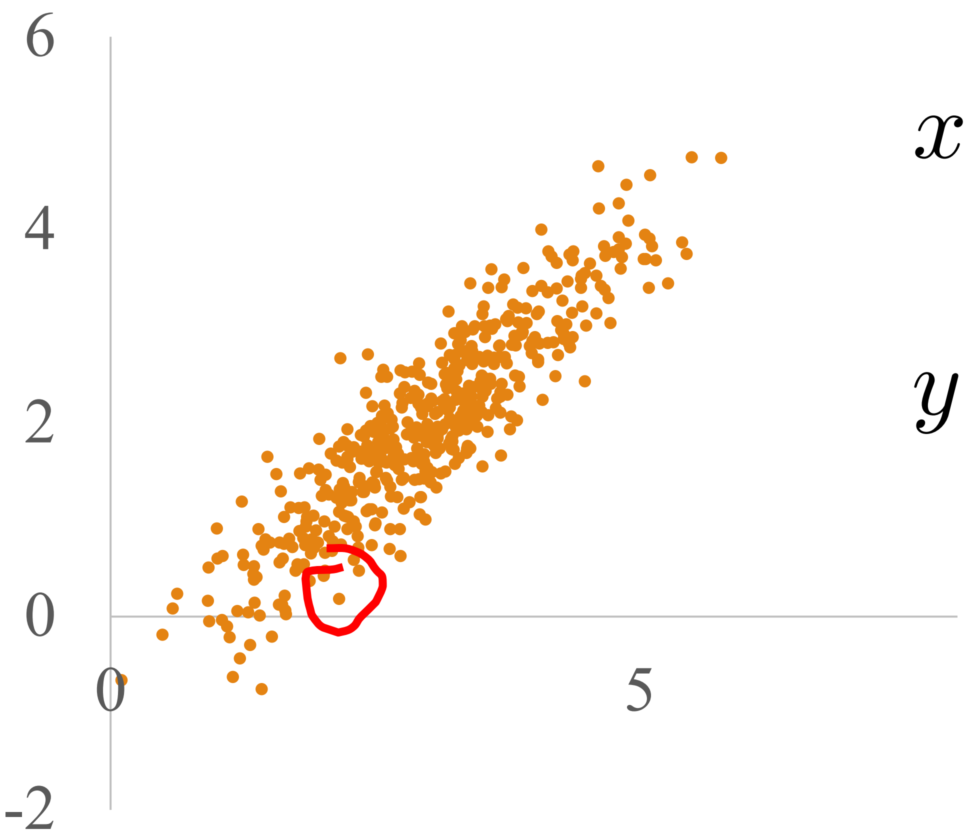

$$x - E[x] = -1.1$$

$$y - E[y] = -2.8$$

$$(x - E[x])(y - E[y]) \approx 3.1$$

Understanding Covariance

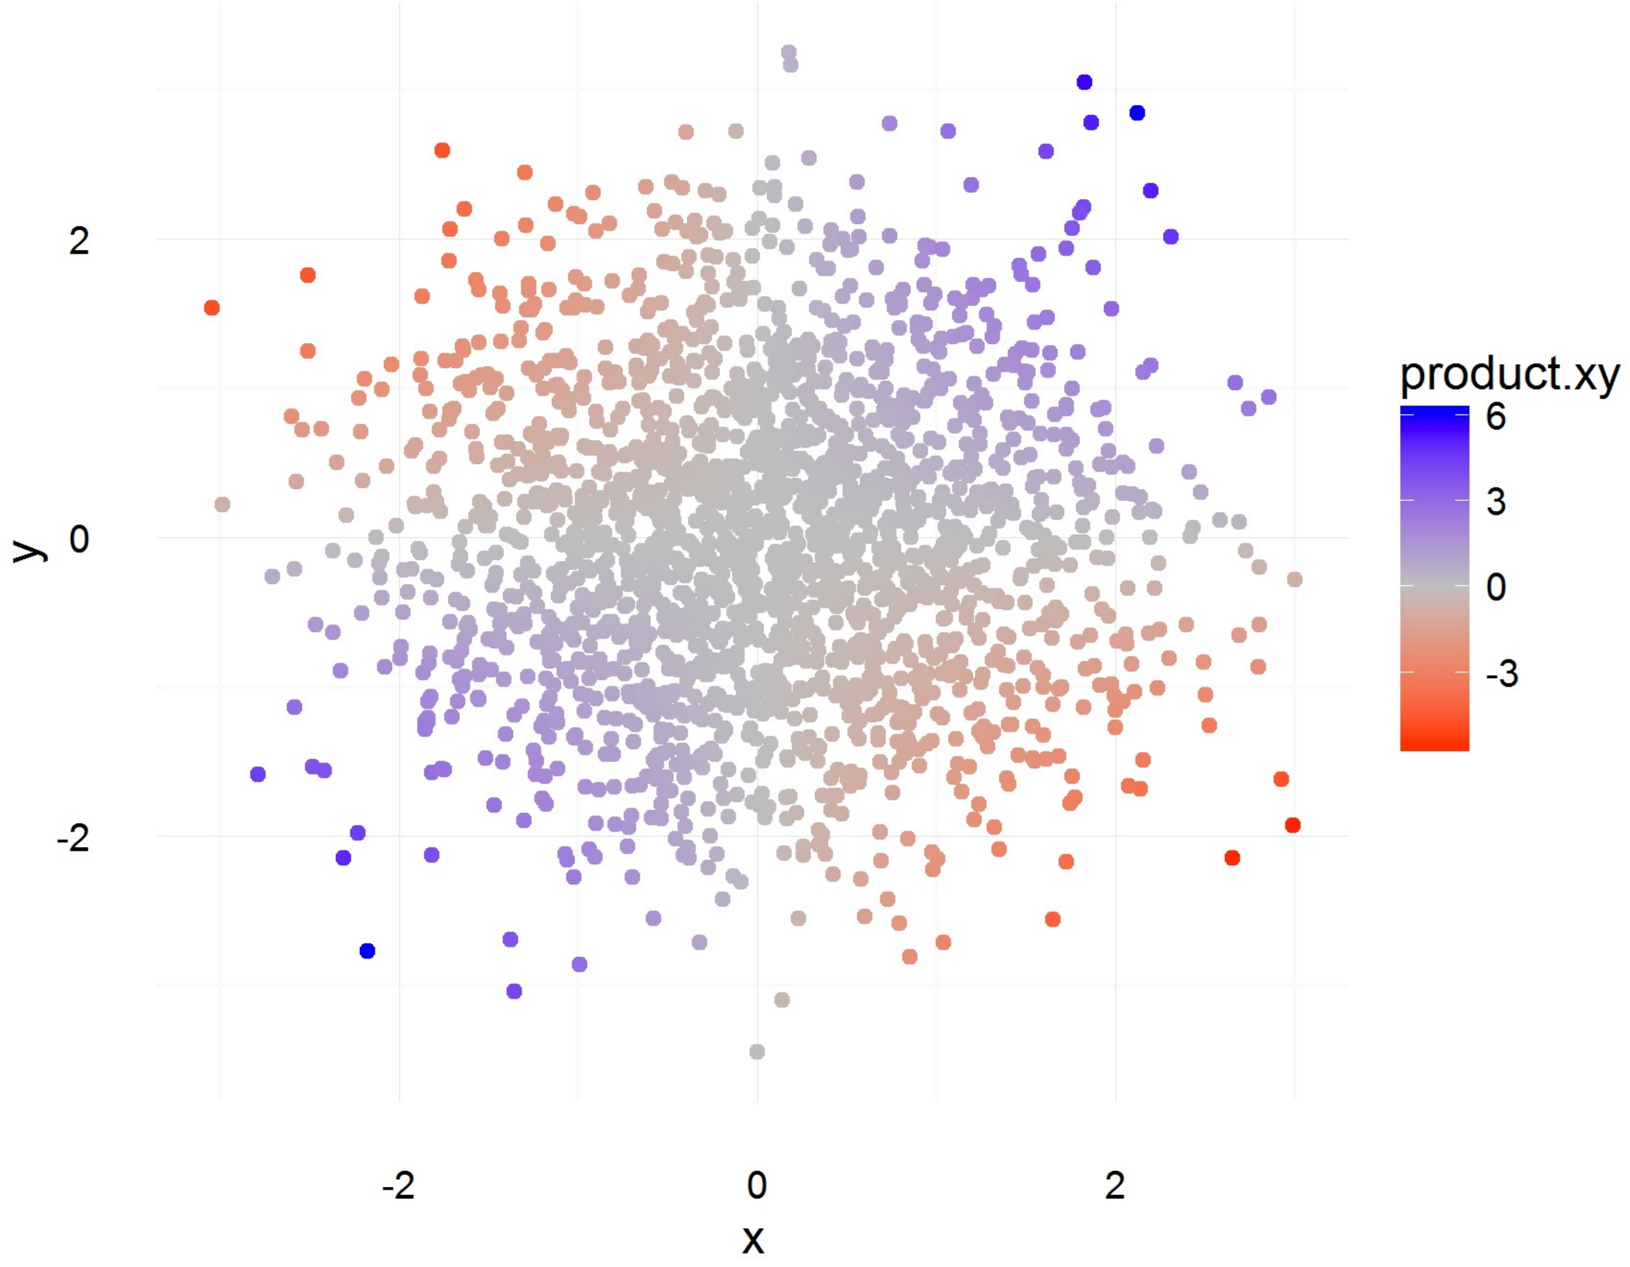

The Dance of the Covariance

Say X and Y are arbitrary random variables

Covariance of X and Y:

$$\text{Cov}(X,Y) = E[(X - E[X])(Y - E[Y])]$$

x	y	(x – E[X])(y – E[Y])p(x,y)
Above mean	Above mean	Positive
Bellow mean	Bellow mean	Positive
Bellow mean	Above mean	Negative
Above mean	Bellow mean	Negative

The Dance of the Covariance

Say X and Y are arbitrary random variables

Covariance of X and Y:

$$\text{Cov}(X,Y) = E[(X - E[X])(Y - E[Y])]$$

Equivalently:

$$\begin{aligned}\text{Cov}(X,Y) &= E[XY - E[X]Y - XE[Y] + E[Y]E[X]] \\ &= E[XY] - E[X]E[Y] - E[X]E[Y] + E[X]E[Y] \\ &= E[XY] - E[X]E[Y]\end{aligned}$$

- X and Y independent, E[XY] = E[X]E[Y] → Cov(X,Y) = 0
- But Cov(X,Y) = 0 does **not** imply X and Y independent!

Covariance and Data

Consider the following data:

Weight	Height	Weight * Height
64	57	3648
71	59	4189
53	49	2597
67	62	4154
55	51	2805
58	50	2900
77	55	4235
57	48	2736
56	42	2352
51	42	2142
76	61	4636
68	57	3876

E[W] = 62.75   E[H] = 52.75   E[W*H] = 3355.83

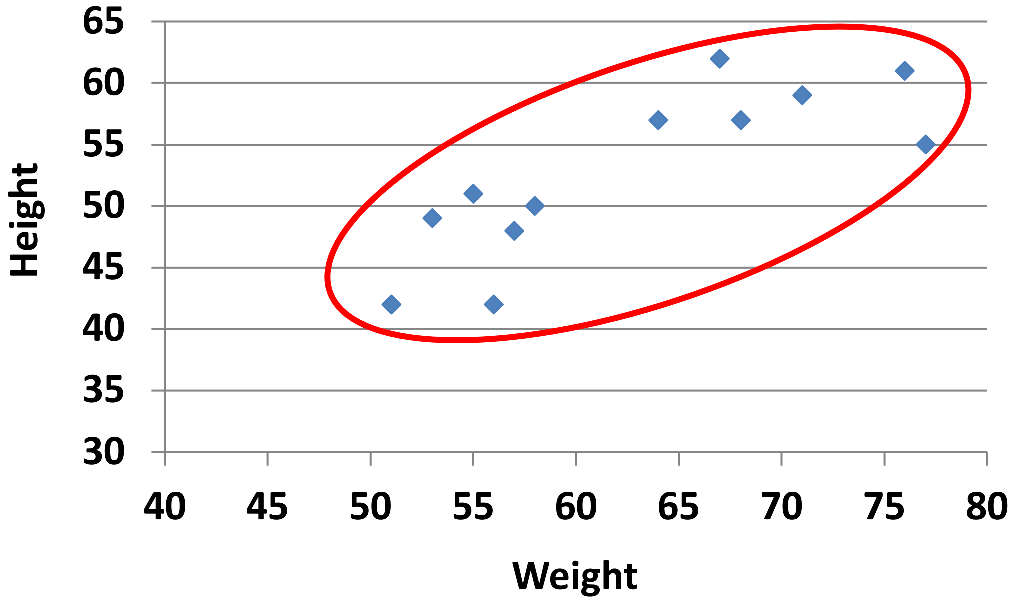

$$\text{Cov}(W, H) = E[W*H] - E[W]E[H]$$
$$= 3355.83 - (62.75)(52.75)$$
$$= 45.77$$

Covariance

Poll: (a) positive, (b) negative, (c) zero

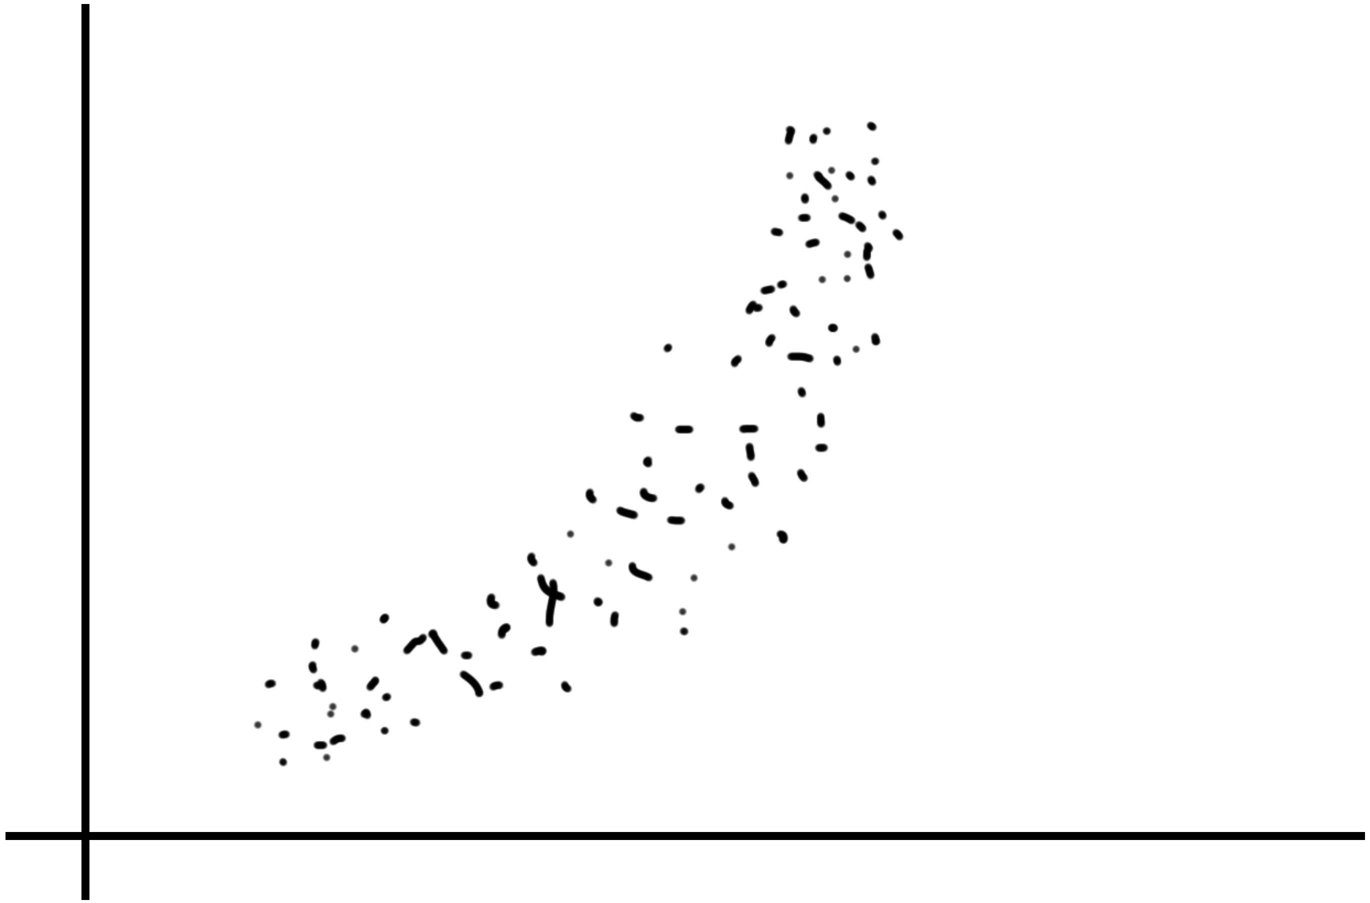

Covariance

Is the Covariance: (a) positive, (b) negative, (c) zero

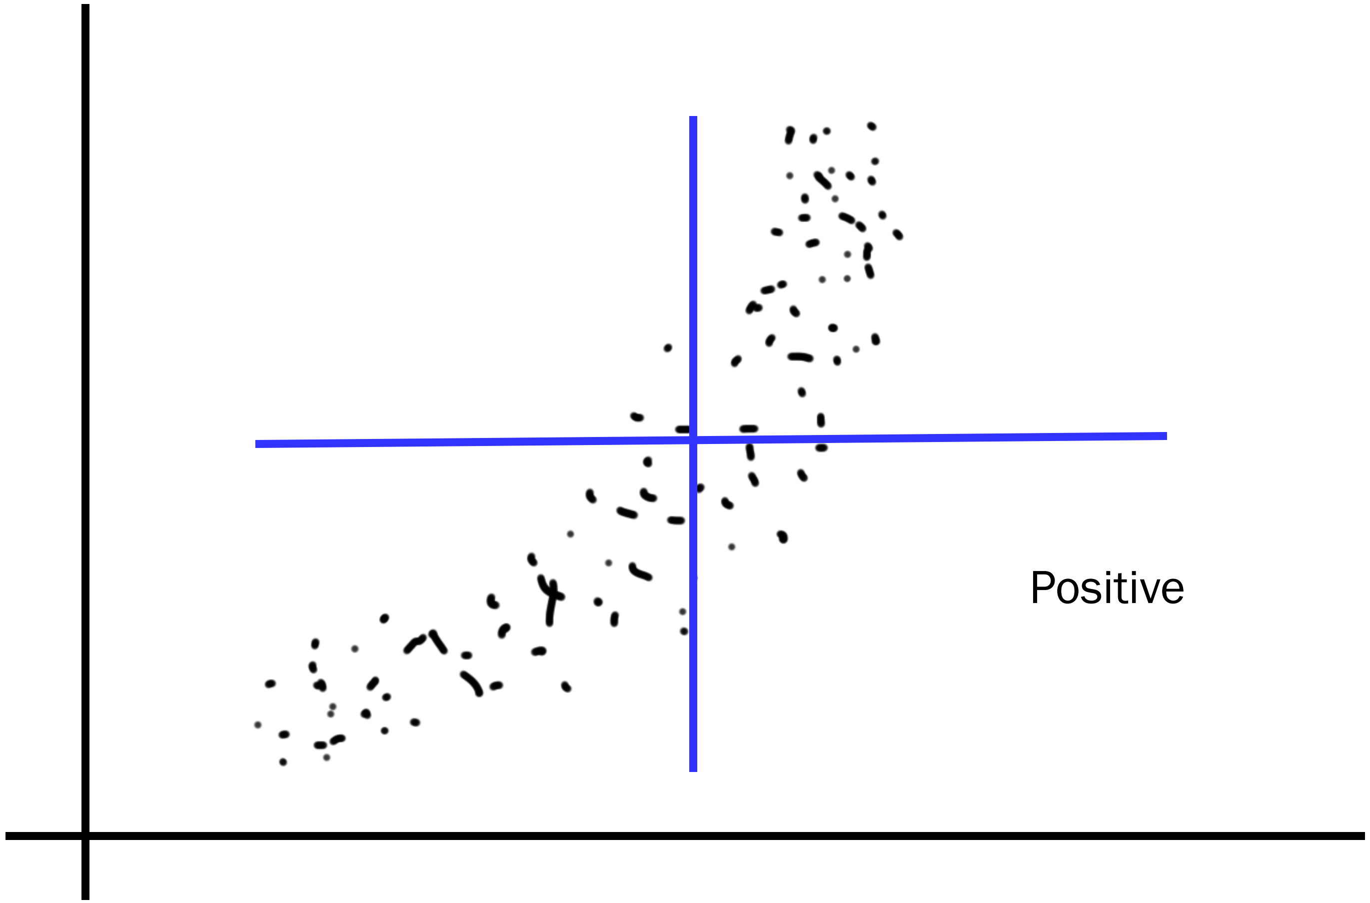

Positive

Covariance

Is the Covariance: (a) positive, (b) negative, (c) zero

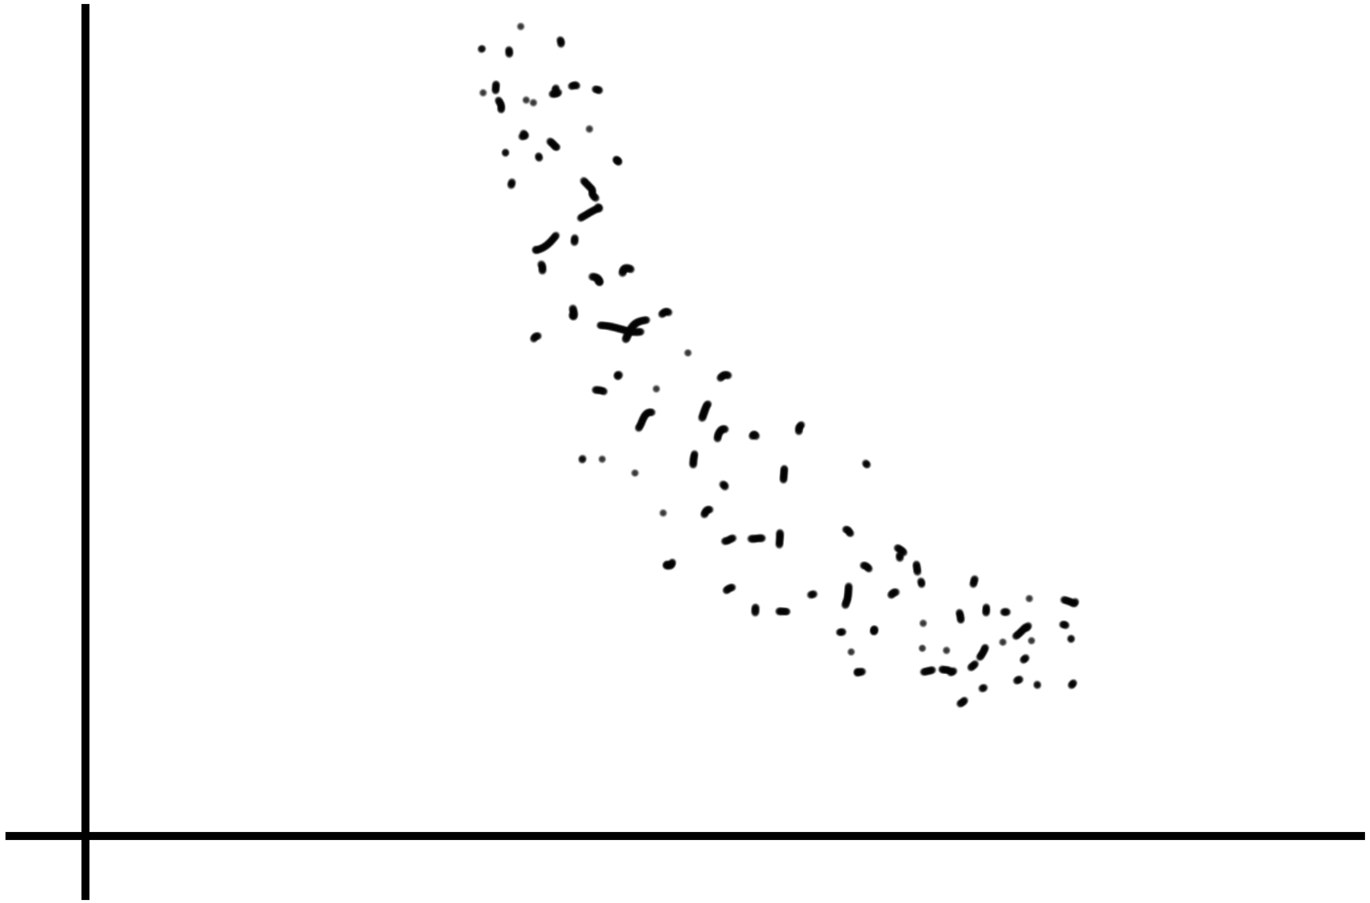

Covariance

Is the Covariance: (a) positive, (b) negative, (c) zero

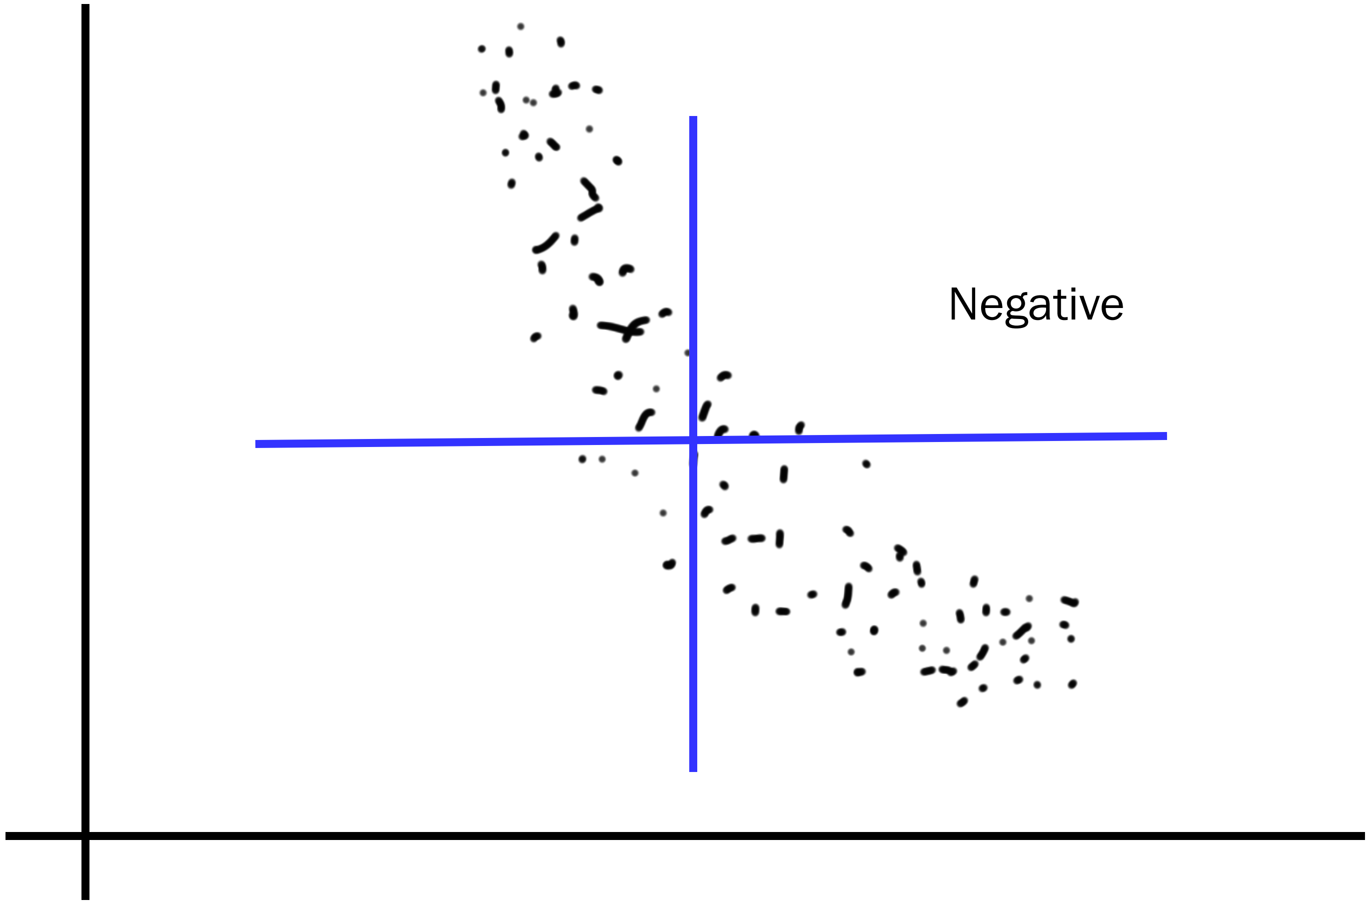

Covariance

Is the Covariance: (a) positive, (b) negative, (c) zero

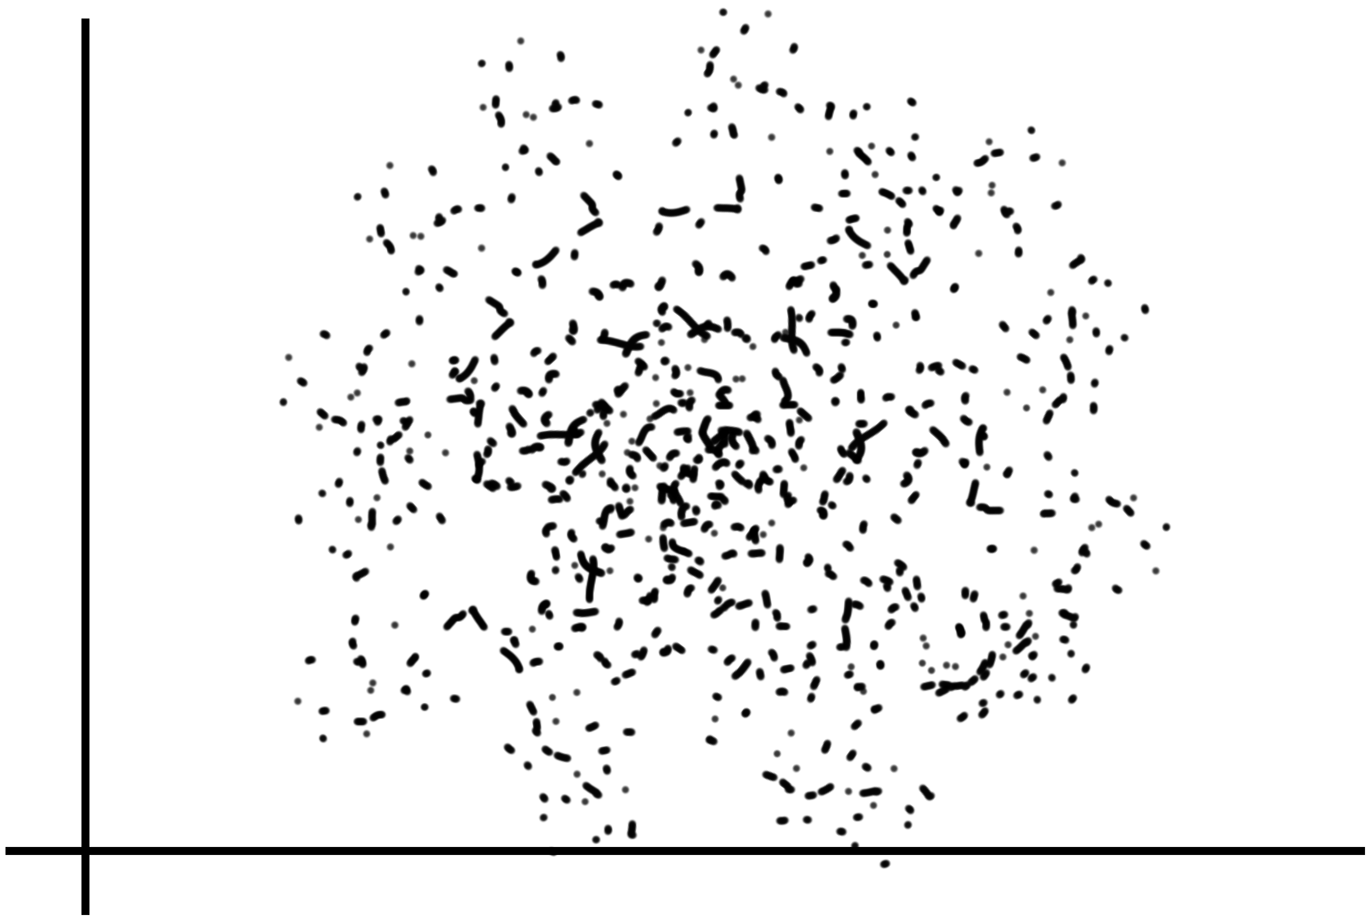

Covariance

Is the Covariance: (a) positive, (b) negative, (c) zero

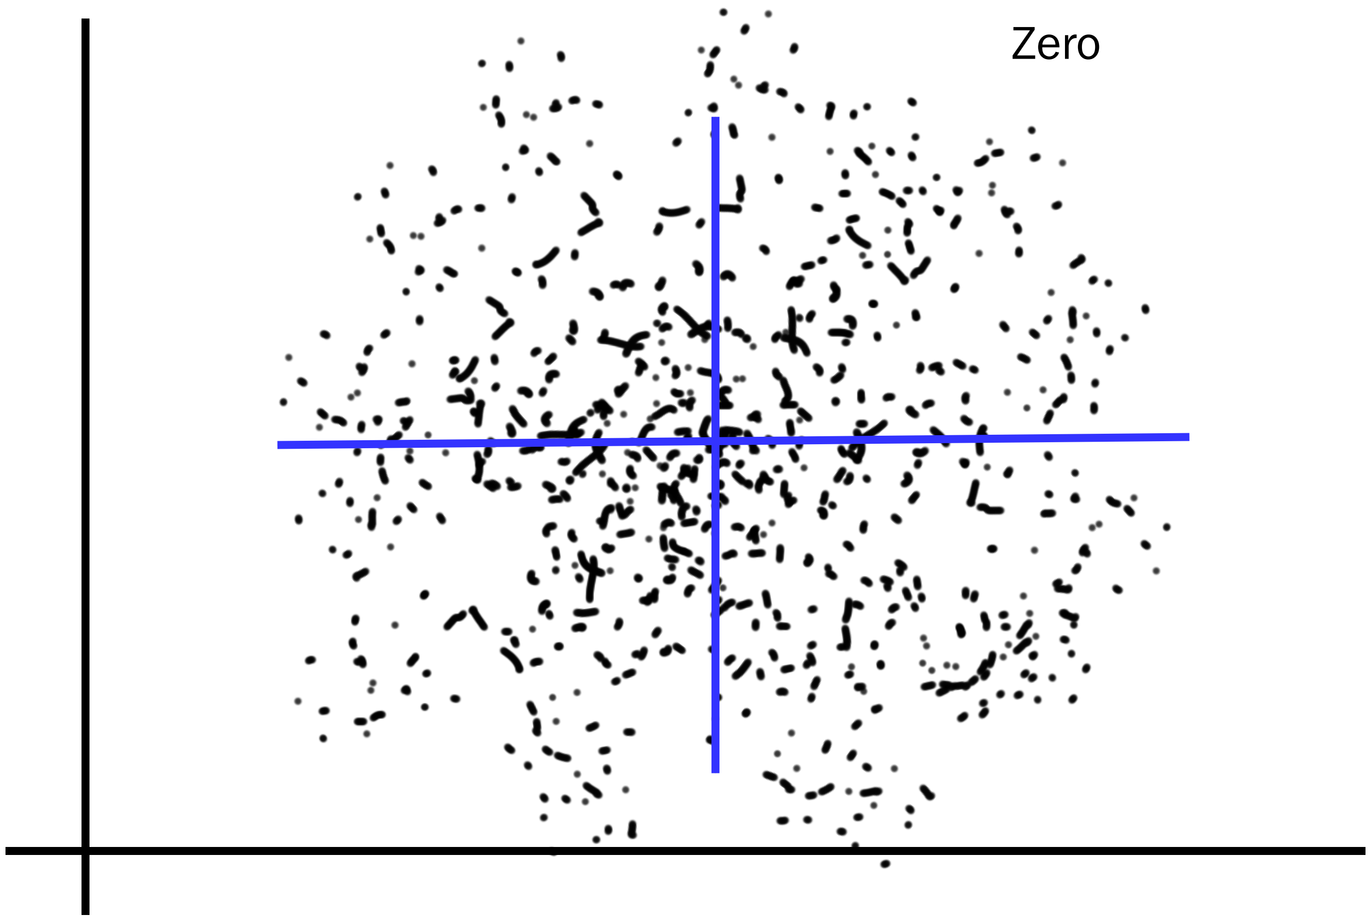

Independence and Covariance

X and Y are random variables with PMF:

Y \ X	-1	0	1	$p_Y(y)$
0	1/3	0	1/3	2/3
1	0	1/3	0	1/3
$p_X(x)$	1/3	1/3	1/3	1

$$Y = \begin{cases} 0 & \text{if } X \neq 0 \\ 1 & \text{otherwise} \end{cases}$$

- E[X] = -1(1/3) + 0(1/3) + 1(1/3) = 0
- E[Y] = 0(2/3) + 1(1/3) = 1/3
- Since XY = 0, E[XY] = 0
- Cov(X, Y) = E[XY] – E[X]E[Y] = 0 – 0 = 0

But, X and Y are clearly dependent!

Properties of Covariance

Say X and Y are arbitrary random variables

- $\text{Cov}(X,Y) = \text{Cov}(Y,X)$
- $\text{Cov}(X,X) = E[X^2] - E[X]E[X] = \text{Var}(X)$
- $\text{Cov}(aX + b, Y) = a\text{Cov}(X,Y)$

Correlation

What is Wrong With This?

Consider the following data:

Weight	Height	Weight * Height
64	57	3648
71	59	4189
53	49	2597
67	62	4154
55	51	2805
58	50	2900
77	55	4235
57	48	2736
56	42	2352
51	42	2142
76	61	4636
68	57	3876

E[W]	E[H]	E[W*H]
= 62.75	= 52.75	= 3355.83

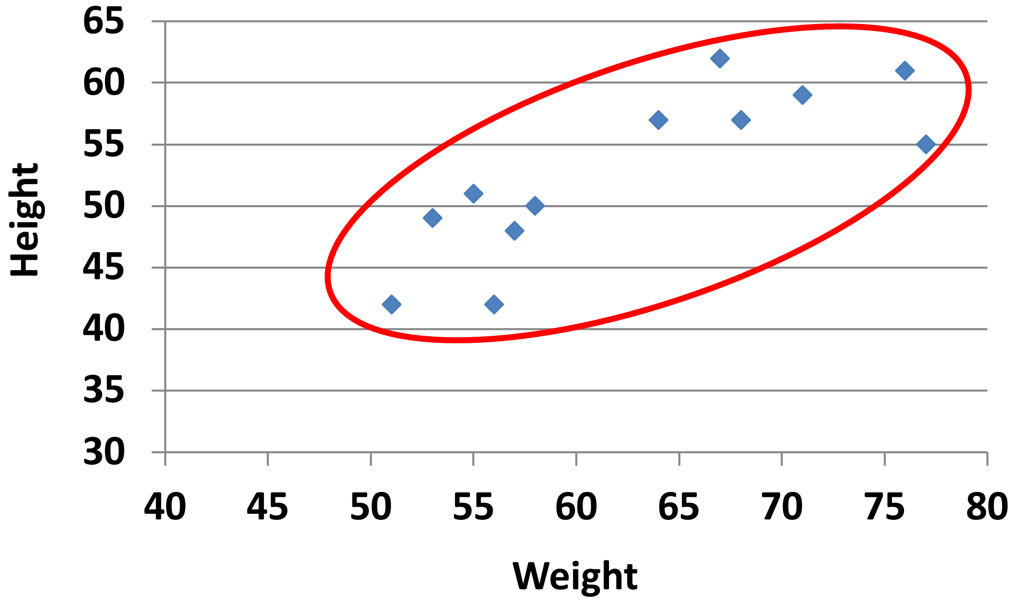

$$\text{Cov}(W, H) = E[W*H] - E[W]E[H]$$
$$= 3355.83 - (62.75)(52.75)$$
$$= 45.77$$

Cauchy Schwarz, a great way to normalize!

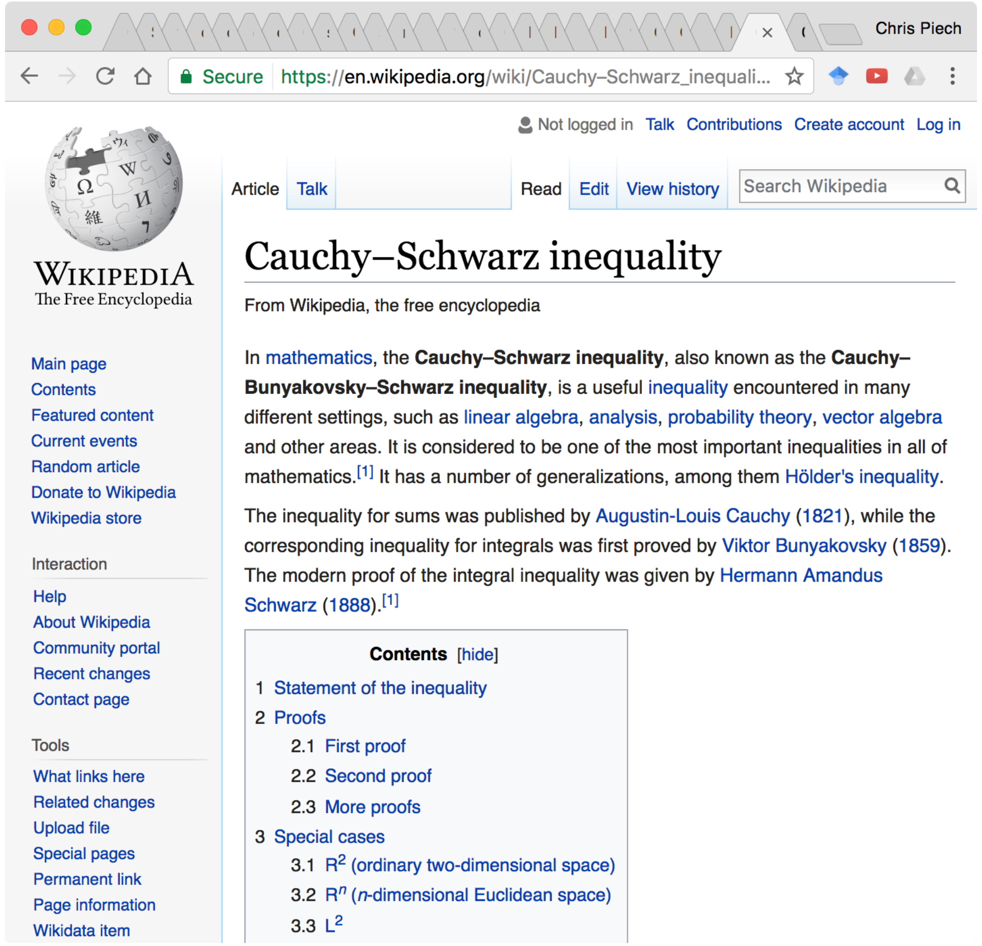
Chris Piech

Secure https://en.wikipedia.org/wiki/Cauchy–Schwarz_inequali...

Not logged in Talk Contributions Create account Log in

WIKIPEDIA
The Free Encyclopedia

Main page
Contents
Featured content
Current events
Random article
Donate to Wikipedia
Wikipedia store

Interaction

Help
About Wikipedia
Community portal
Recent changes
Contact page

Tools

What links here
Related changes
Upload file
Special pages
Permanent link
Page information
Wikidata item

Article Talk Read Edit View history Search Wikipedia

Cauchy–Schwarz inequality

From Wikipedia, the free encyclopedia

In mathematics, the **Cauchy–Schwarz inequality**, also known as the **Cauchy–Bunyakovsky–Schwarz inequality**, is a useful inequality encountered in many different settings, such as linear algebra, analysis, probability theory, vector algebra and other areas. It is considered to be one of the most important inequalities in all of mathematics.[1] It has a number of generalizations, among them Hölder's inequality.

The inequality for sums was published by Augustin-Louis Cauchy (1821), while the corresponding inequality for integrals was first proved by Viktor Bunyakovsky (1859). The modern proof of the integral inequality was given by Hermann Amandus Schwarz (1888).[1]

Contents [hide]

1 Statement of the inequality
2 Proofs
2.1 First proof
2.2 Second proof
2.3 More proofs
3 Special cases
3.1 $\mathbb{R}^2$ (ordinary two-dimensional space)
3.2 $\mathbb{R}^n$ (n-dimensional Euclidean space)
3.3 L^2

$$-\text{Std}(X)\text{Std}(Y) \leq \text{Cov}(X,Y) \leq \text{Std}(X)\text{Std}(Y)$$

Viva La Correlatión

Say X and Y are arbitrary random variables

- Correlation of X and Y, denoted $\rho(X, Y)$:

$$\rho(X,Y) = \frac{\mathrm{Cov}(X,Y)}{\sqrt{\mathrm{Var(X)Var(Y)}}}$$

- Note: $-1 \leq \rho(X, Y) \leq 1$
- Correlation measures <u>linearity</u> between X and Y
- $\rho(X, Y) = 1 \quad \Rightarrow \quad Y = aX + b$ where $a = \sigma_y/\sigma_x$
- $\rho(X, Y) = -1 \quad \Rightarrow \quad Y = aX + b$ where $a = -\sigma_y/\sigma_x$
- $\rho(X, Y) = 0 \quad \Rightarrow$ absence of <u>linear</u> relationship
  - But, X and Y can still be related in some other way!
- If $\rho(X, Y) = 0$, we say X and Y are "uncorrelated"
  - Note: Independence implies uncorrelated, but **<u>not</u>** vice versa!

Viva La Correlatión

Say X and Y are arbitrary random variables

- Correlation of X and Y, denoted $\rho(X, Y)$:

$$\rho(X,Y) = \frac{\text{Cov}(X,Y)}{\sqrt{\text{Var(X)Var(Y)}}}$$

Say Y = cX. Correlation should be 1.

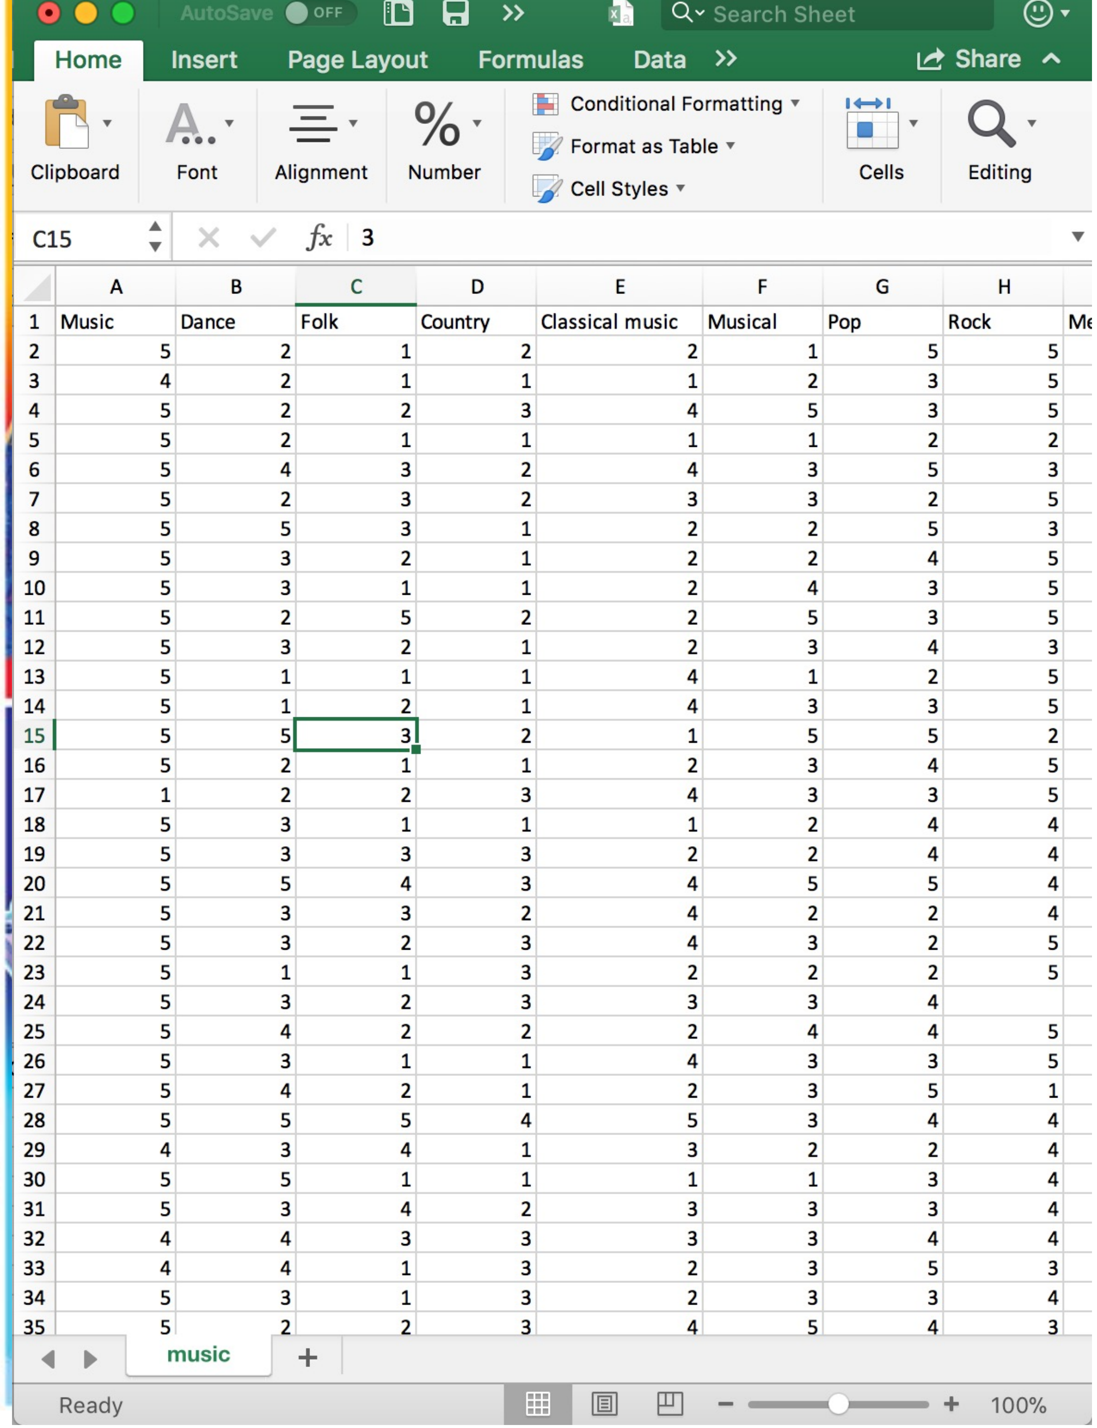

	A	B	C	D	E	F	G	H	
1	Music	Dance	Folk	Country	Classical music	Musical	Pop	Rock	Me
2	5	2	1	2	2	1	5	5	
3	4	2	1	1	1	2	3	5	
4	5	2	2	3	4	5	3	5	
5	5	2	1	1	1	1	2	2	
6	5	4	3	2	4	3	5	3	
7	5	2	3	2	3	3	2	5	
8	5	5	3	1	2	2	5	3	
9	5	3	2	1	2	2	4	5	
10	5	3	1	1	2	4	3	5	
11	5	2	5	2	2	5	3	5	
12	5	3	2	1	2	3	4	3	
13	5	1	1	1	4	1	2	5	
14	5	1	2	1	4	3	3	5	
15	5	5	3	2	1	5	5	2	
16	5	2	1	1	2	3	4	5	
17	1	2	2	3	4	3	3	5	
18	5	3	1	1	1	2	4	4	
19	5	3	3	3	2	2	4	4	
20	5	5	4	3	4	5	5	4	
21	5	3	3	2	4	2	2	4	
22	5	3	2	3	4	3	2	5	
23	5	1	1	3	2	2	2	5	
24	5	3	2	3	3	3	4		
25	5	4	2	2	2	4	4	5	
26	5	3	1	1	4	3	3	5	
27	5	4	2	1	2	3	5	1	
28	5	5	5	4	5	3	4	4	
29	4	3	4	1	3	2	2	4	
30	5	5	1	1	1	1	3	4	
31	5	3	4	2	3	3	3	4	
32	4	4	3	3	3	3	4	4	
33	4	4	1	3	2	3	5	3	
34	5	3	1	3	2	3	3	4	
35	5	2	2	3	4	5	4	3	

See You on Friday

Extra Slides

Do Indicators Correlate?

Let I_A and I_B be indicators for events A and B

$$I_A = \begin{cases} 1 & \text{if } A \text{ occurs} \\ 0 & \text{otherwise} \end{cases} \qquad I_B = \begin{cases} 1 & \text{if } B \text{ occurs} \\ 0 & \text{otherwise} \end{cases}$$

- $E[I_A] = P(A)$, $E[I_B] = P(B)$, $E[I_A I_B] = P(AB)$
- $\text{Cov}(I_A, I_B) = E[I_A I_B] - E[I_A]\,E[I_B]$
  $= P(AB) - P(A)P(B)$
  $= P(A \mid B)P(B) - P(A)P(B)$
  $= P(B)[P(A \mid B) - P(A)]$
- $\text{Cov}(I_A, I_B)$ determined by $P(A \mid B) - P(A)$
- $P(A \mid B) > P(A) \Rightarrow \rho(I_A, I_B) > 0$
- $P(A \mid B) = P(A) \Rightarrow \rho(I_A, I_B) = 0$ (and $\text{Cov}(I_A, I_B) = 0$)
- $P(A \mid B) < P(A) \Rightarrow \rho(I_A, I_B) < 0$

Correlation of Music Tastes

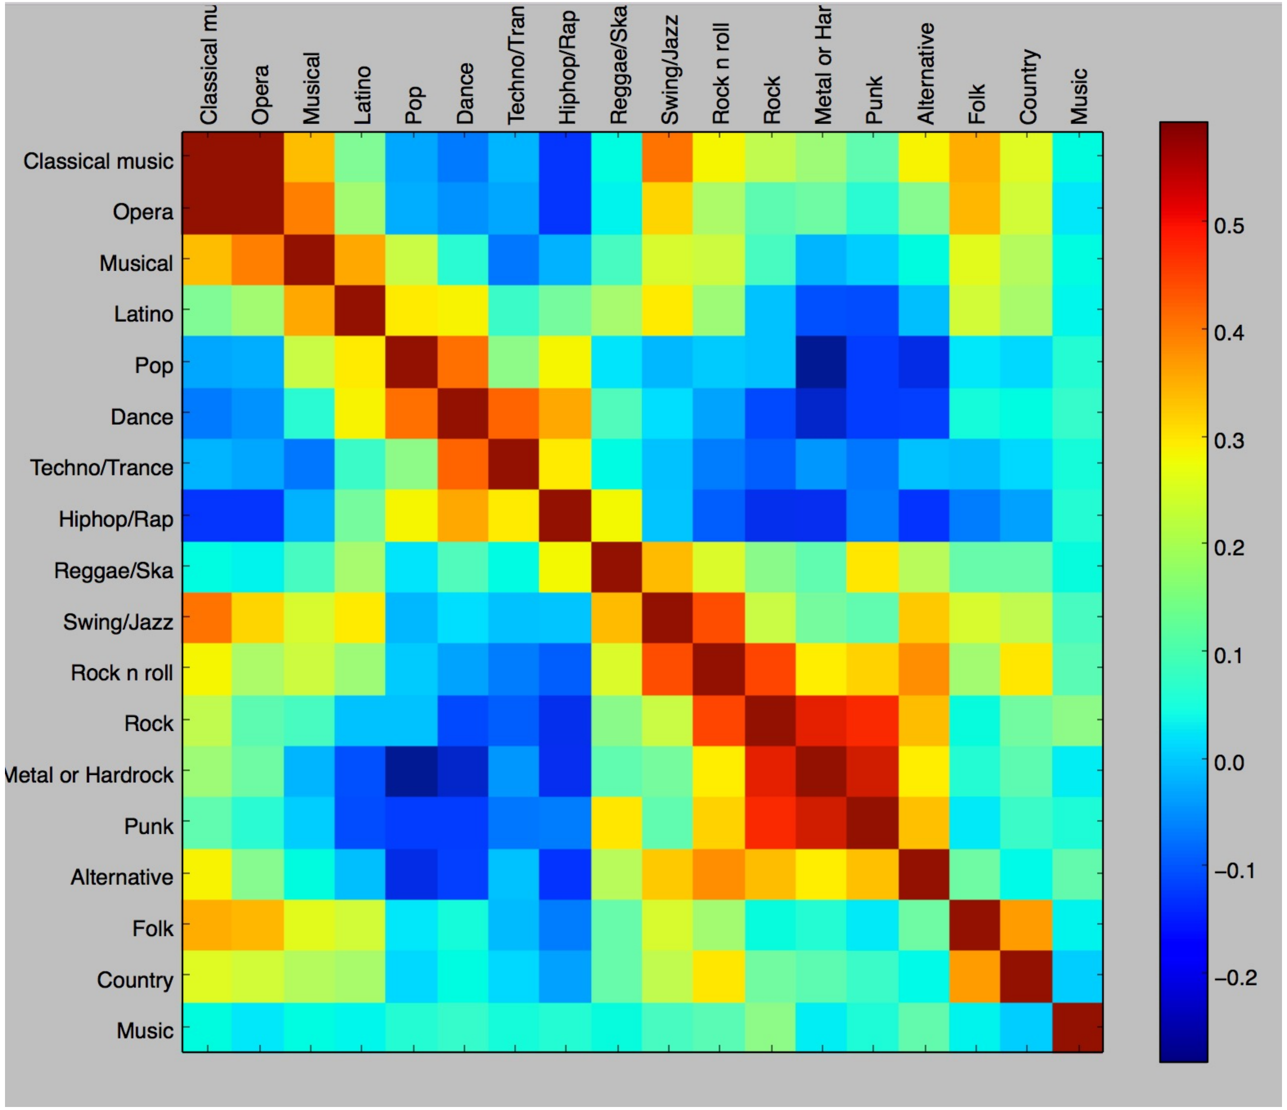

Correlation of Music Tastes

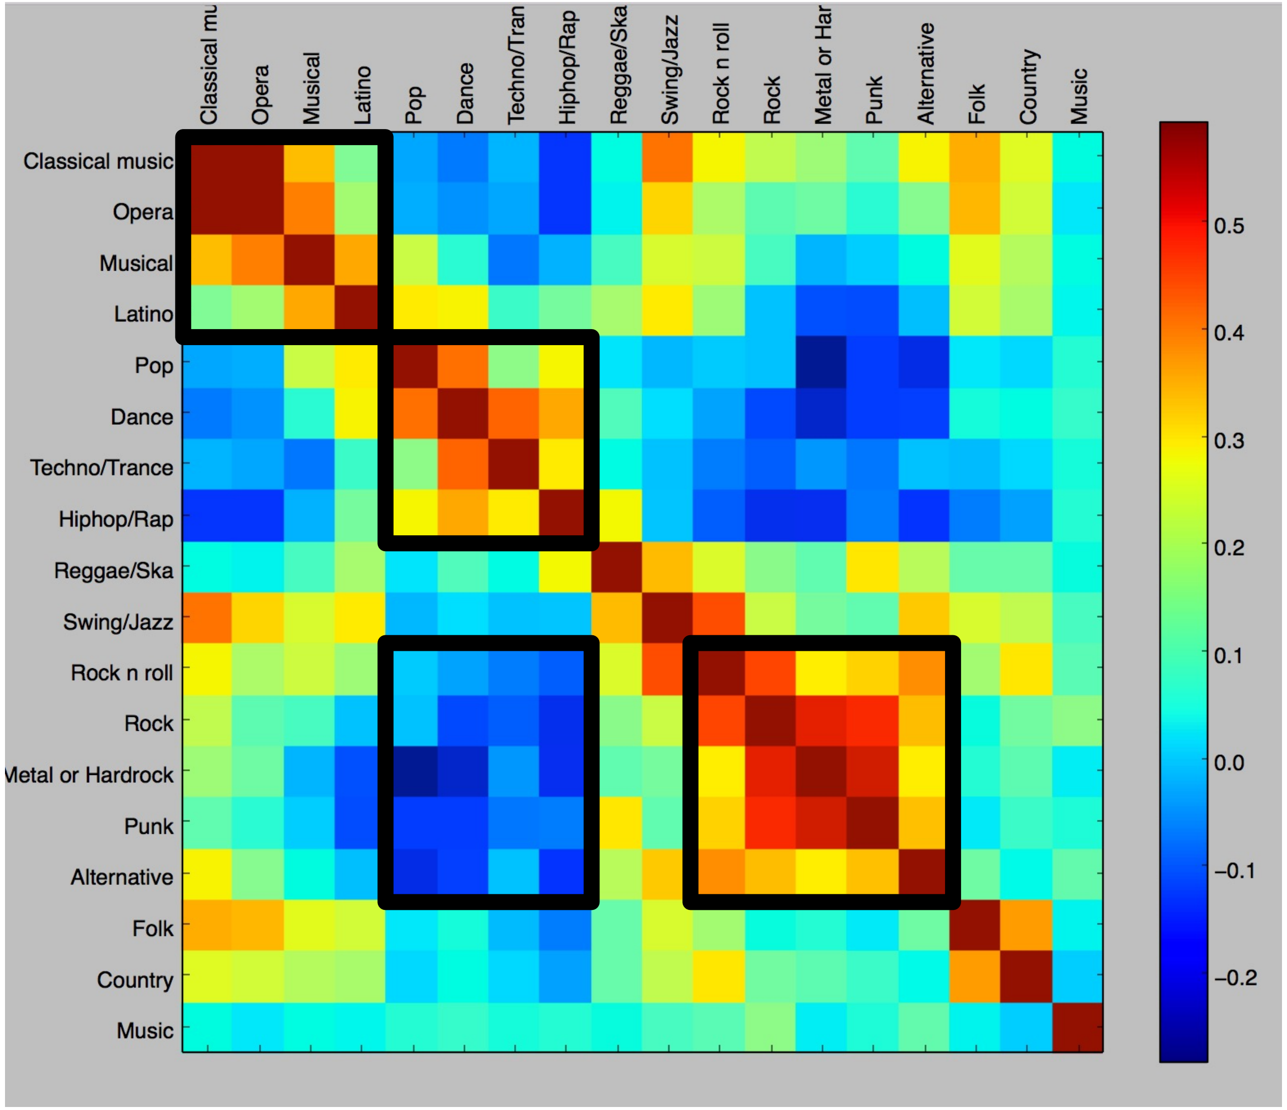

From Correlation to Bayes Net. Alternative!

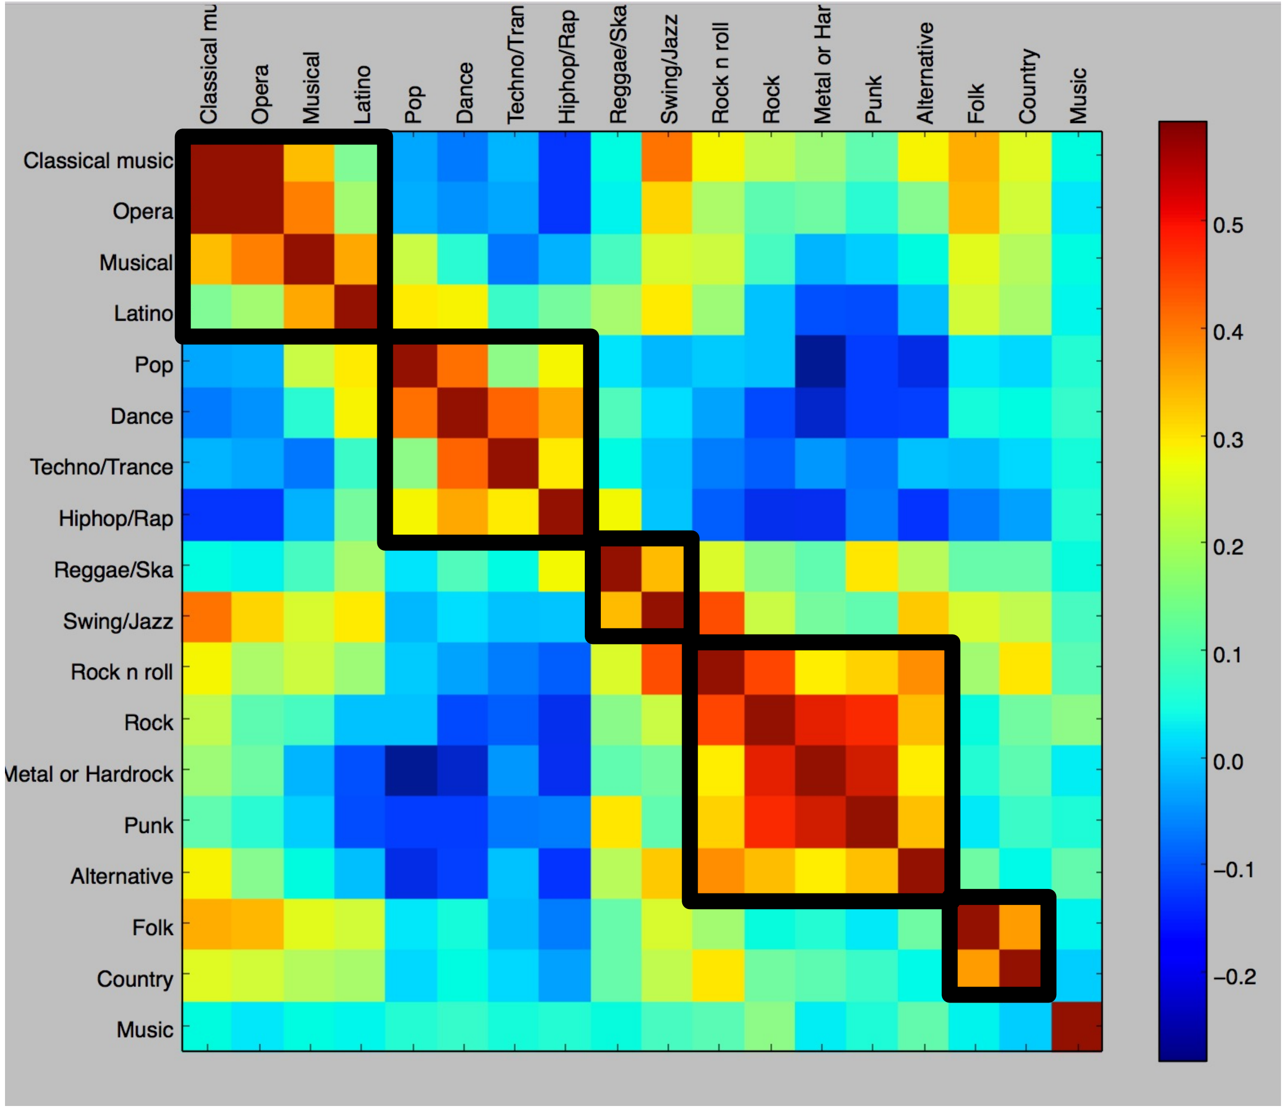

From Correlation to Bayes Net. Alternative!

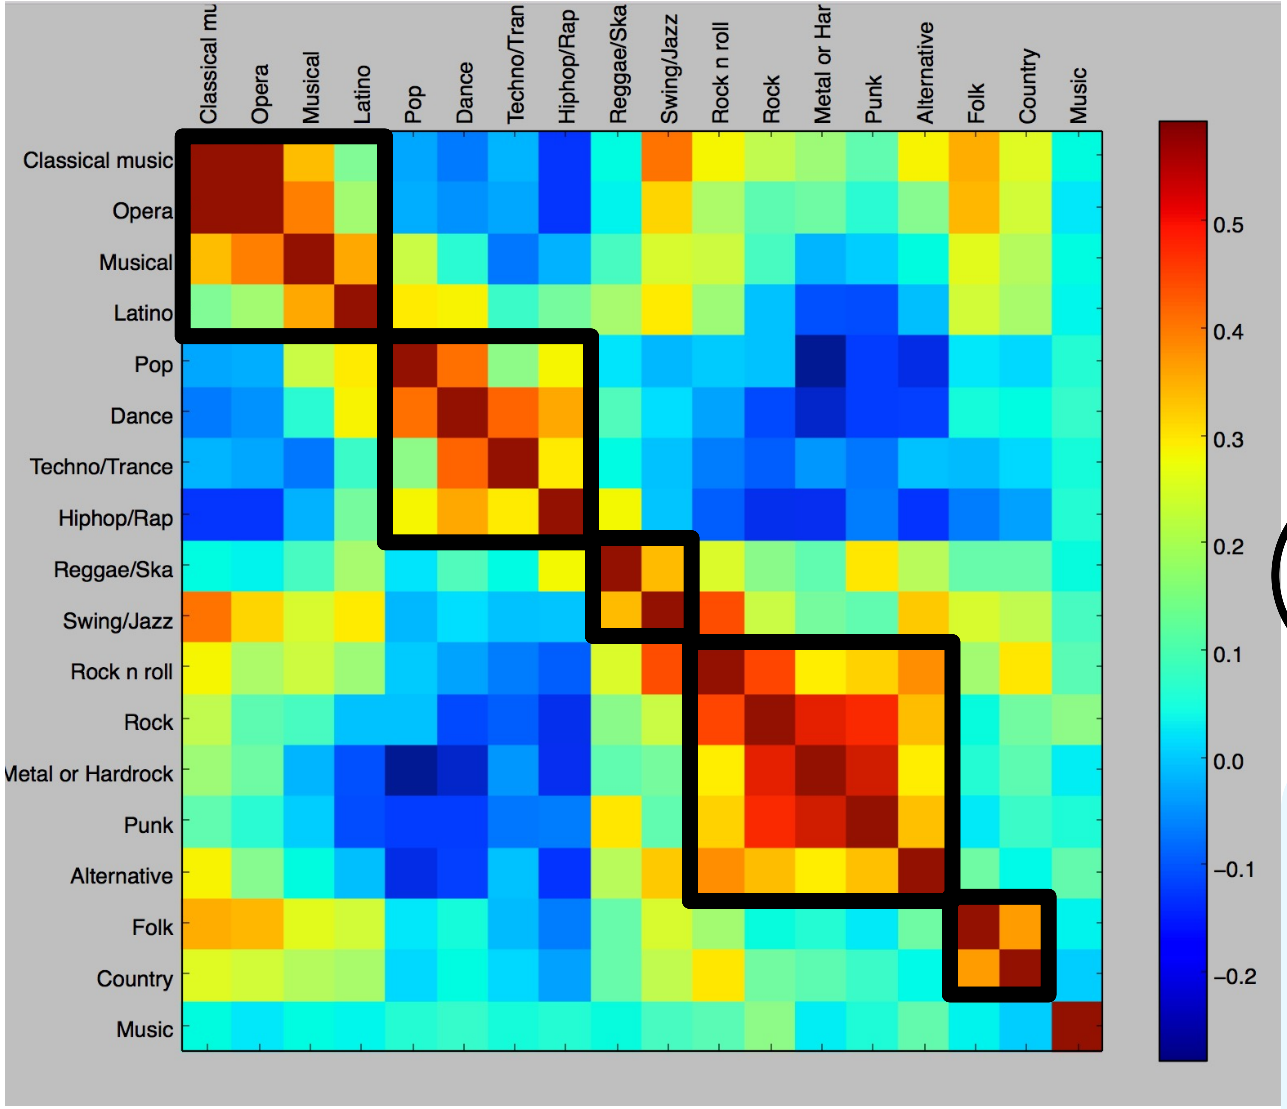

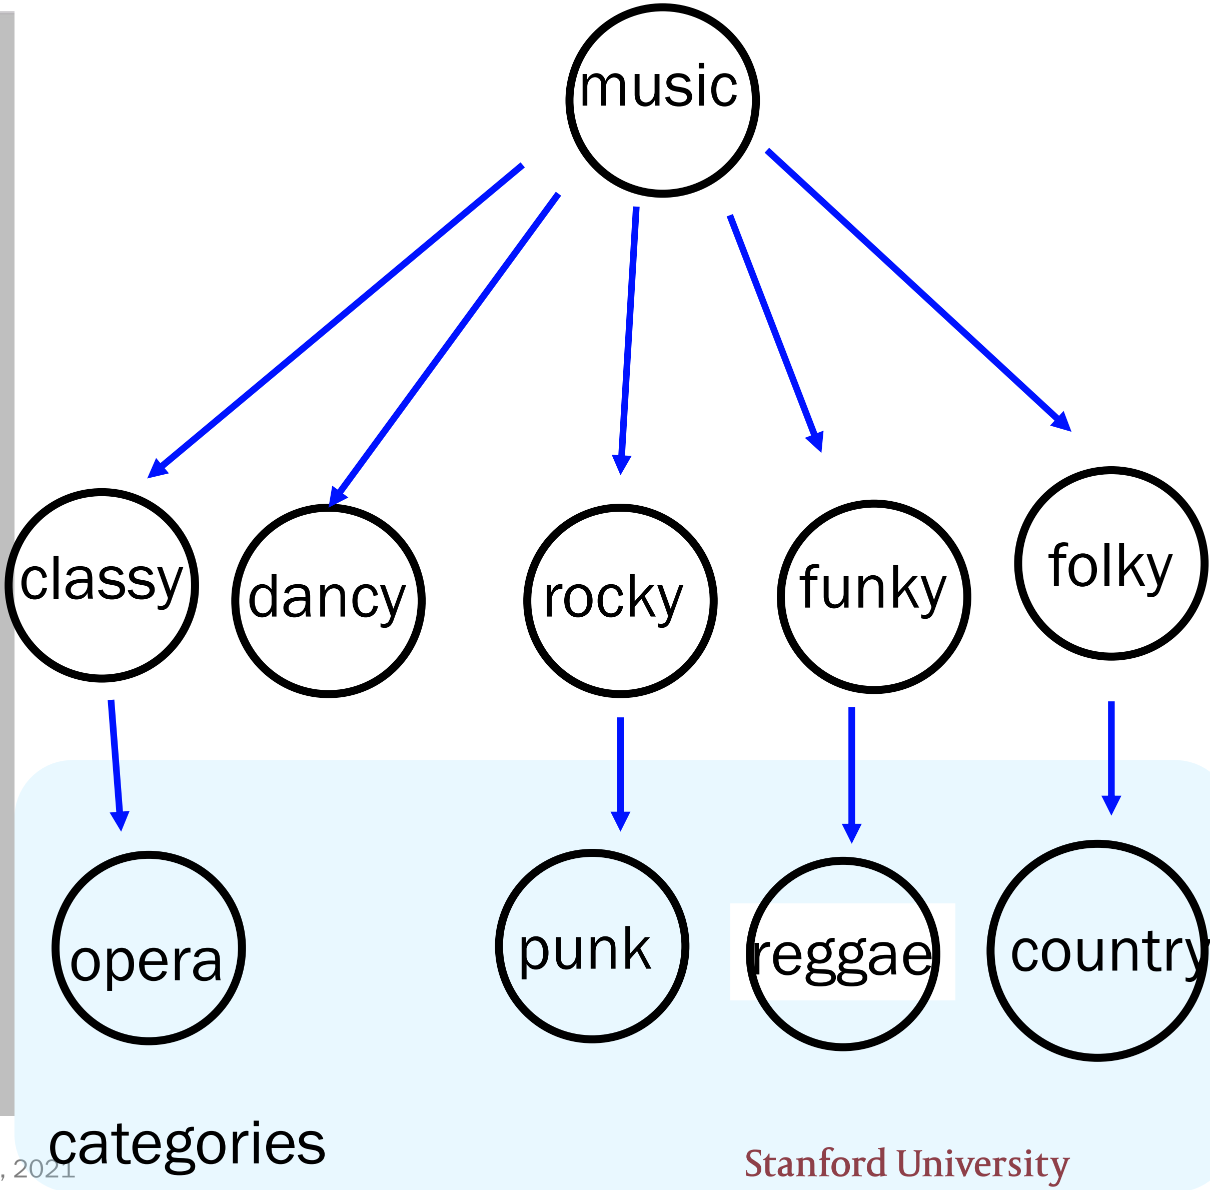

From Correlation to Bayes Net. Alternative!

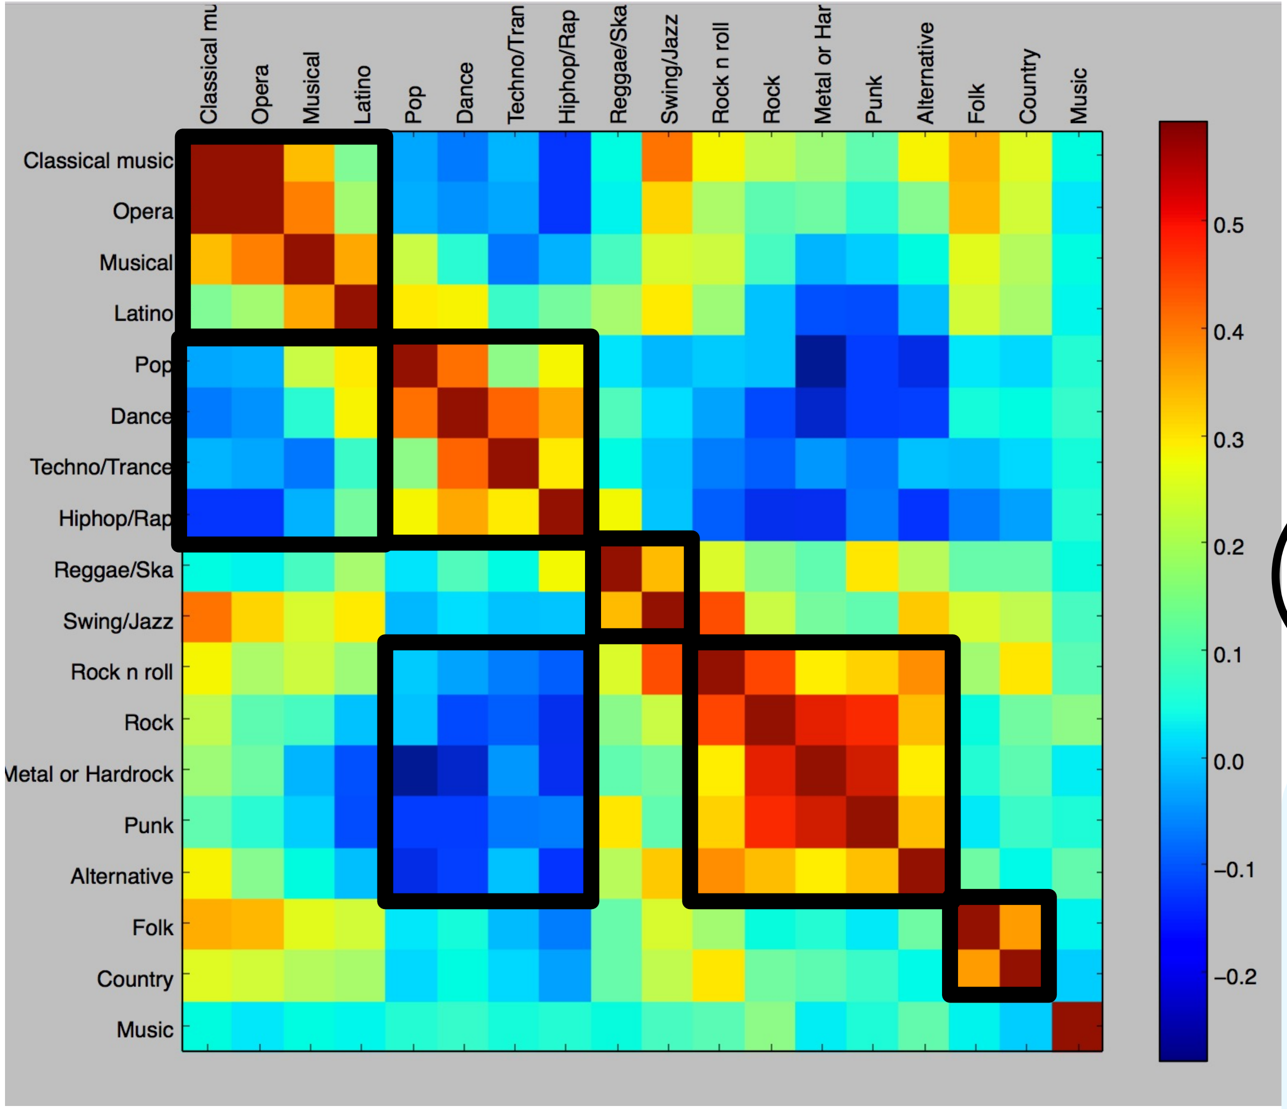

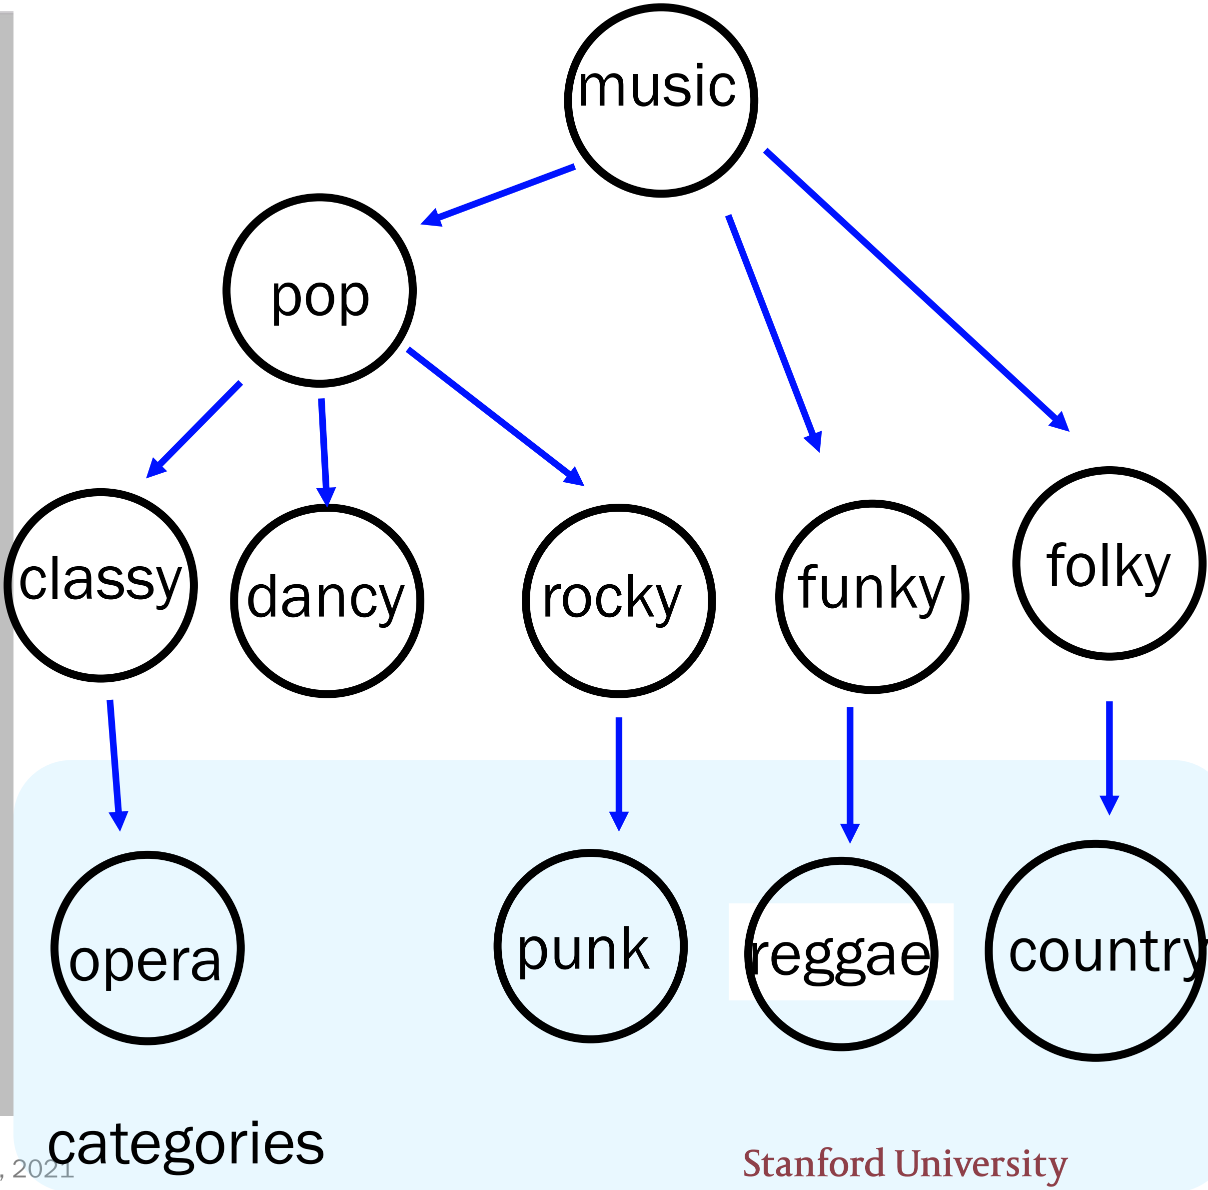

From Correlation to Bayes Net. Alternative!

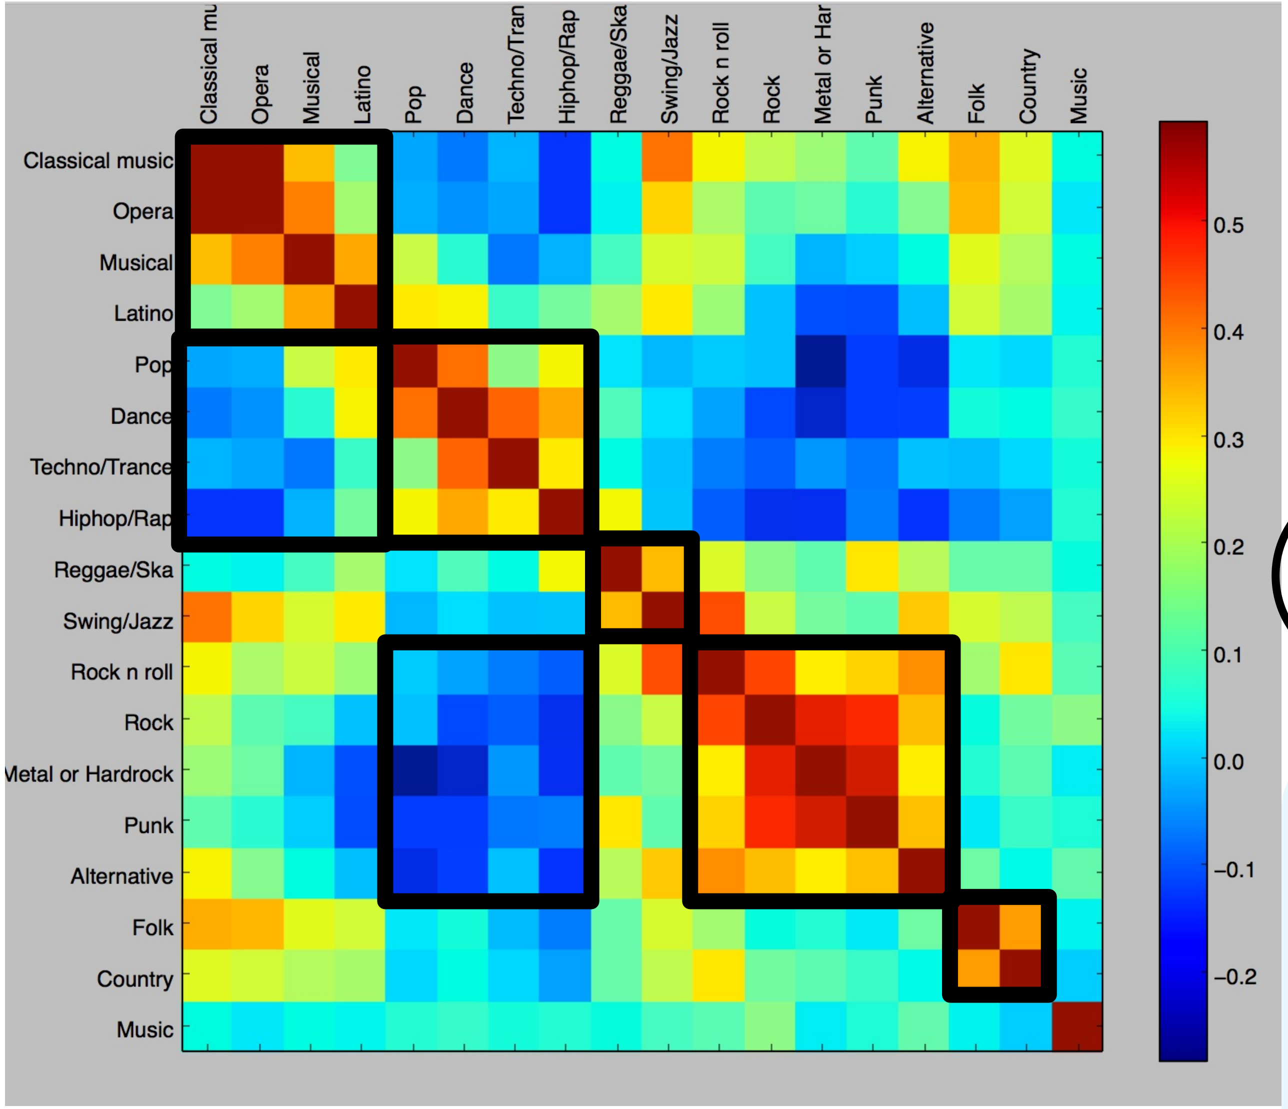

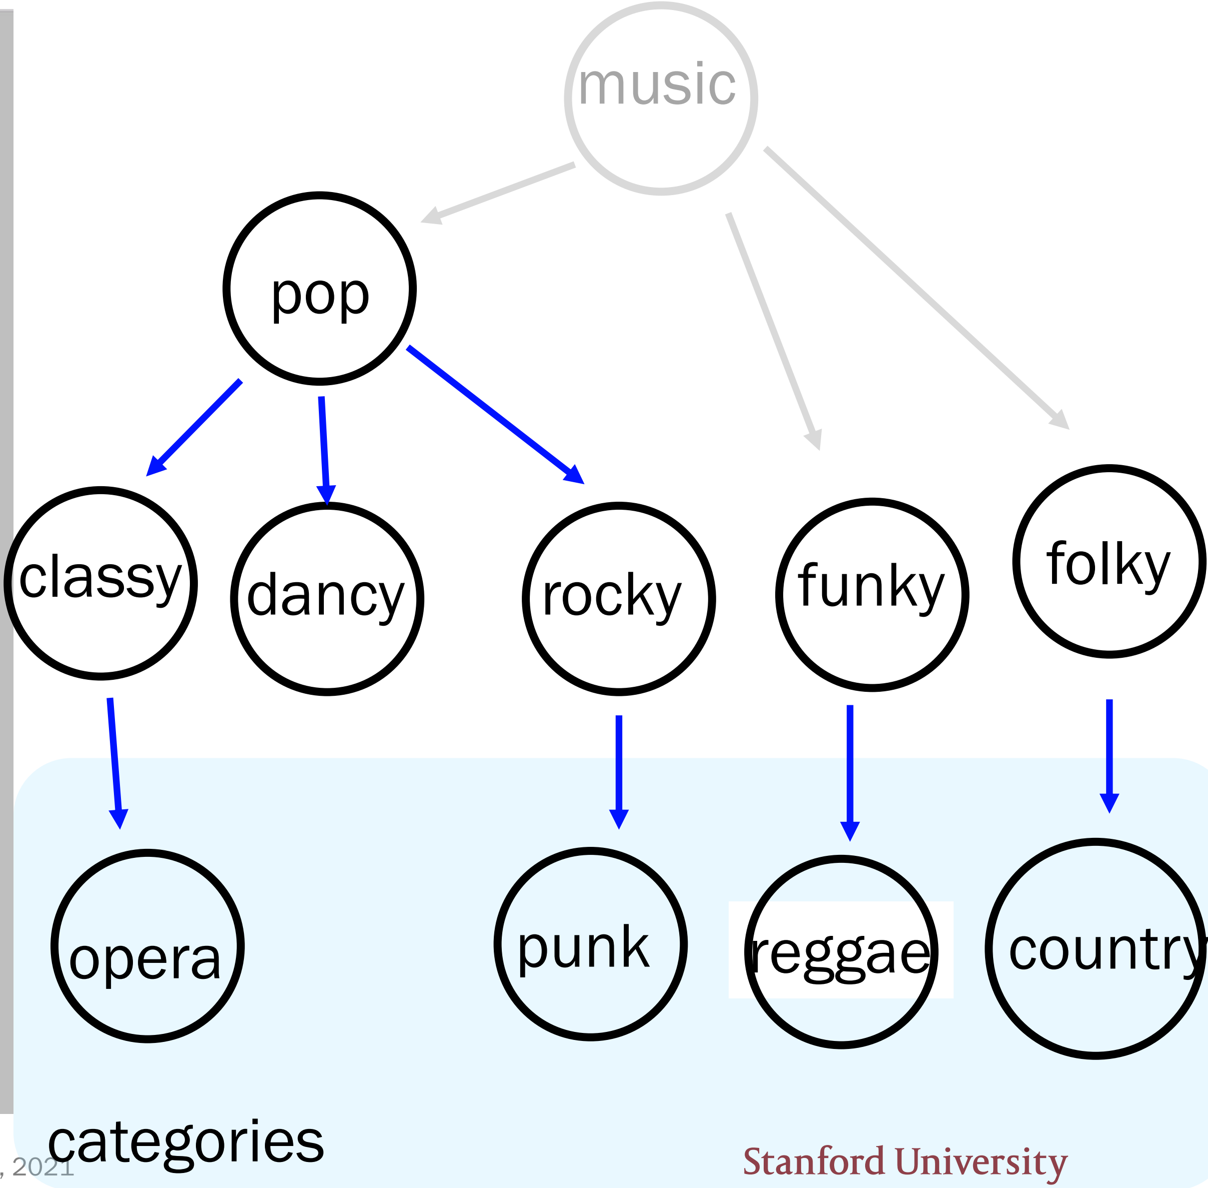

How do you know if your model is good?

Answer: it is accurate at inference (especially tasks you care about)

Rock Music Vs Oil?

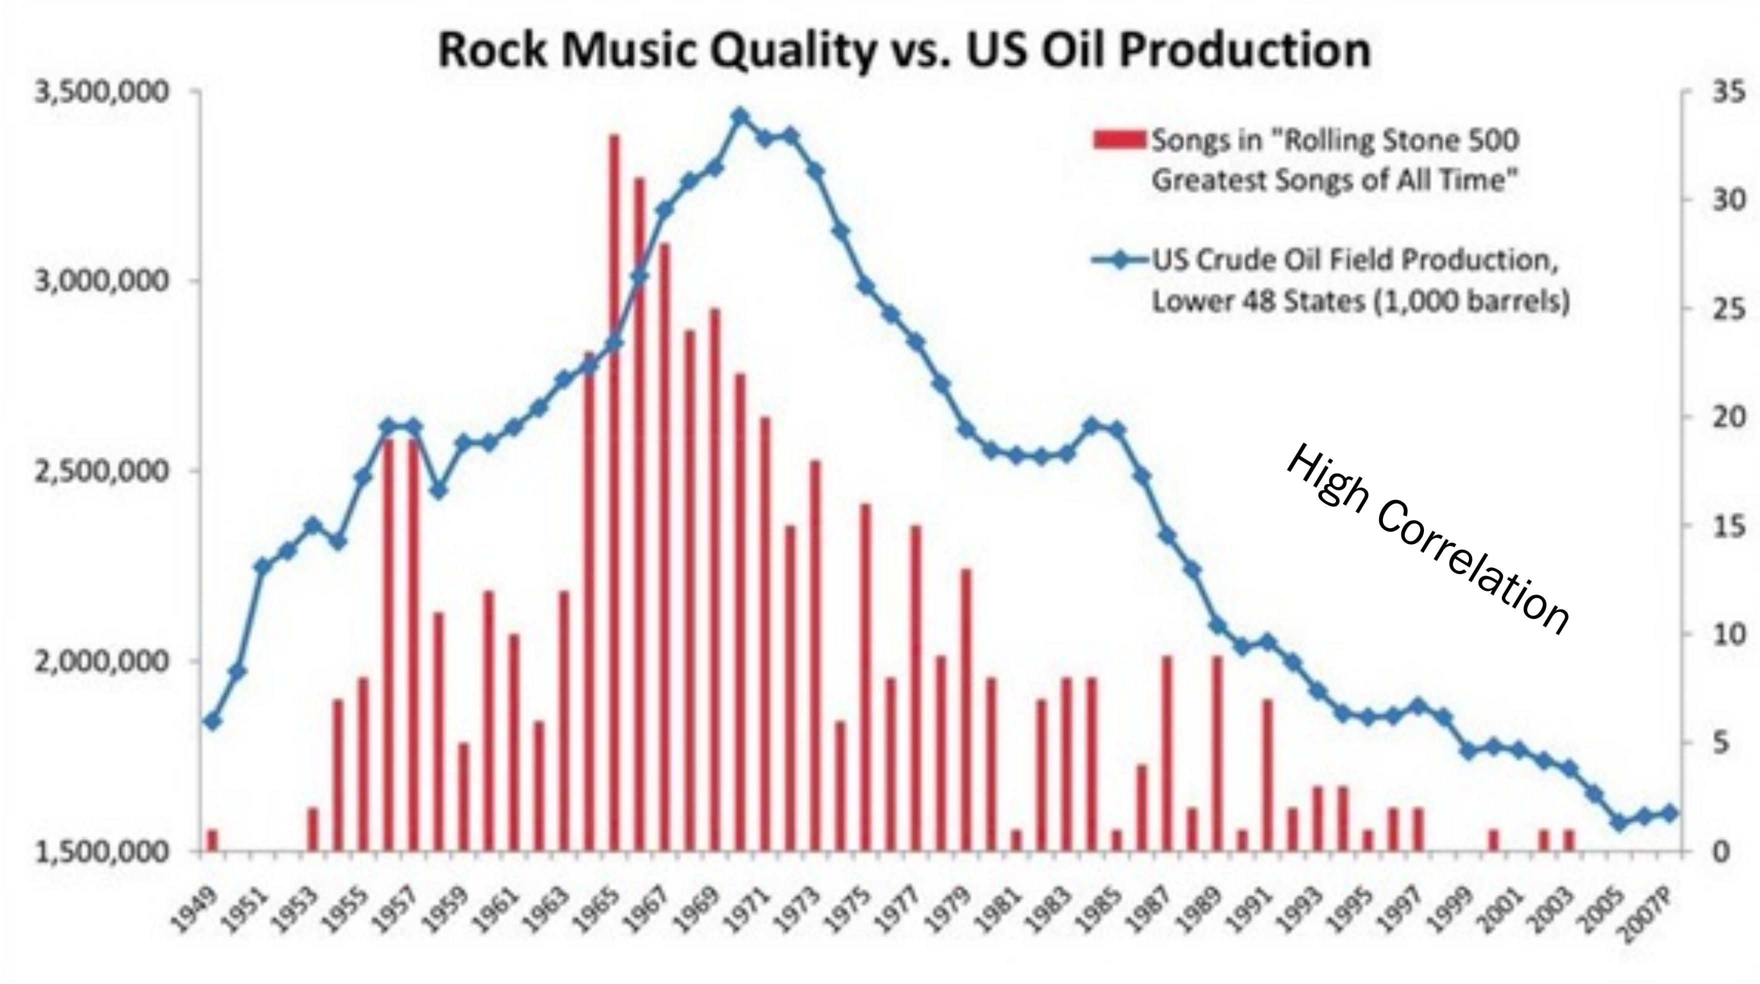

Hubbert Peak Theory

http://www.aei.org/publication/blog/

Tell your friends!

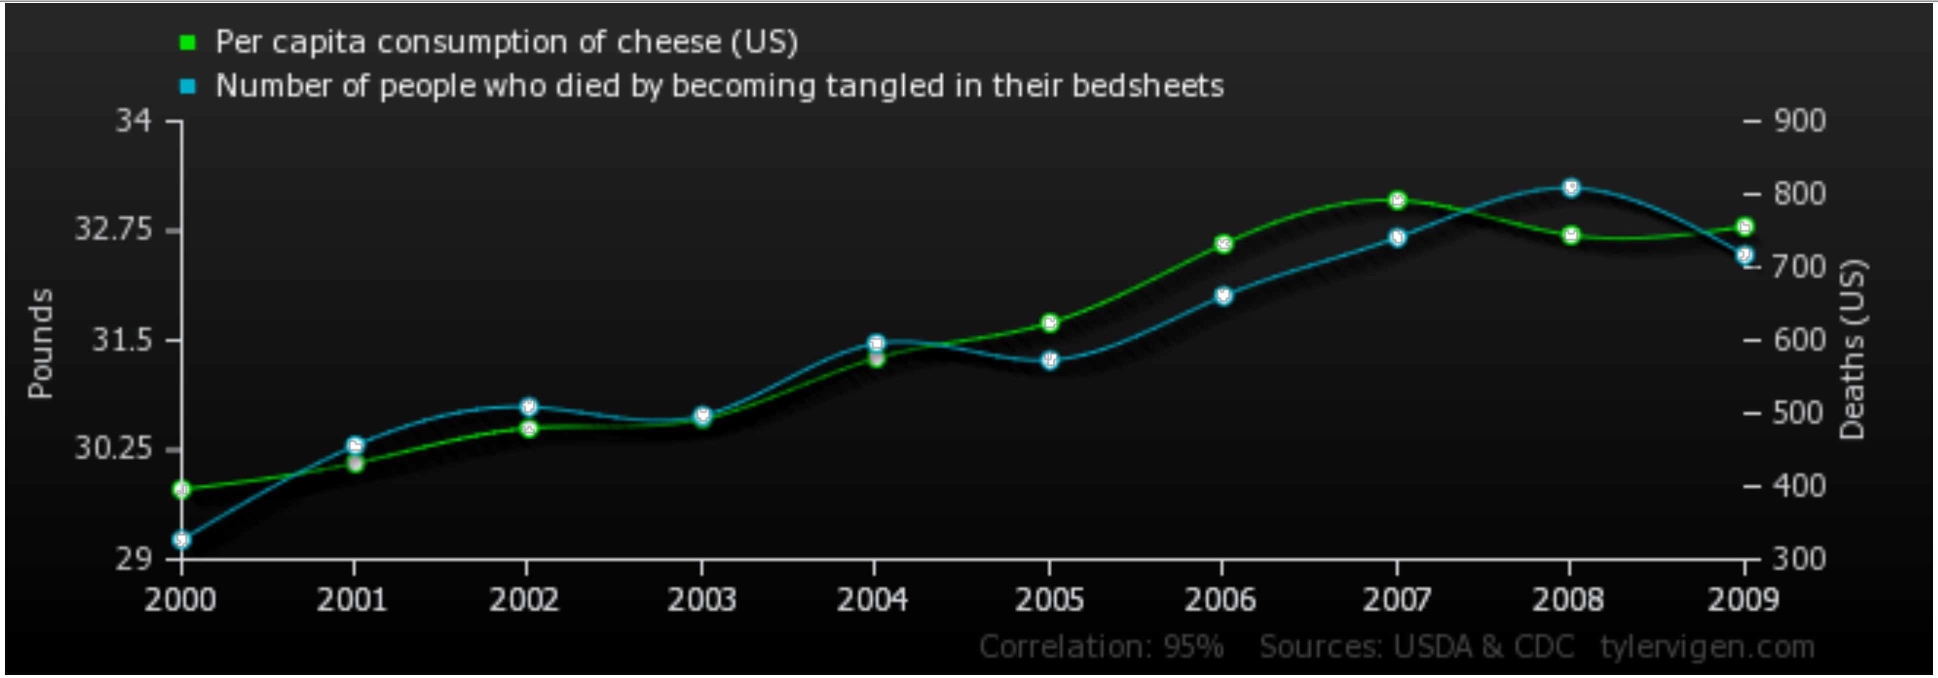

	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009
Per capita consumption of cheese (US) *Pounds (USDA)*	29.8	30.1	30.5	30.6	31.3	31.7	32.6	33.1	32.7	32.8
Number of people who died by becoming tangled in their bedsheets *Deaths (US) (CDC)*	327	456	509	497	596	573	661	741	809	717

Correlation: 0.947091

Divorce Vs Butter?

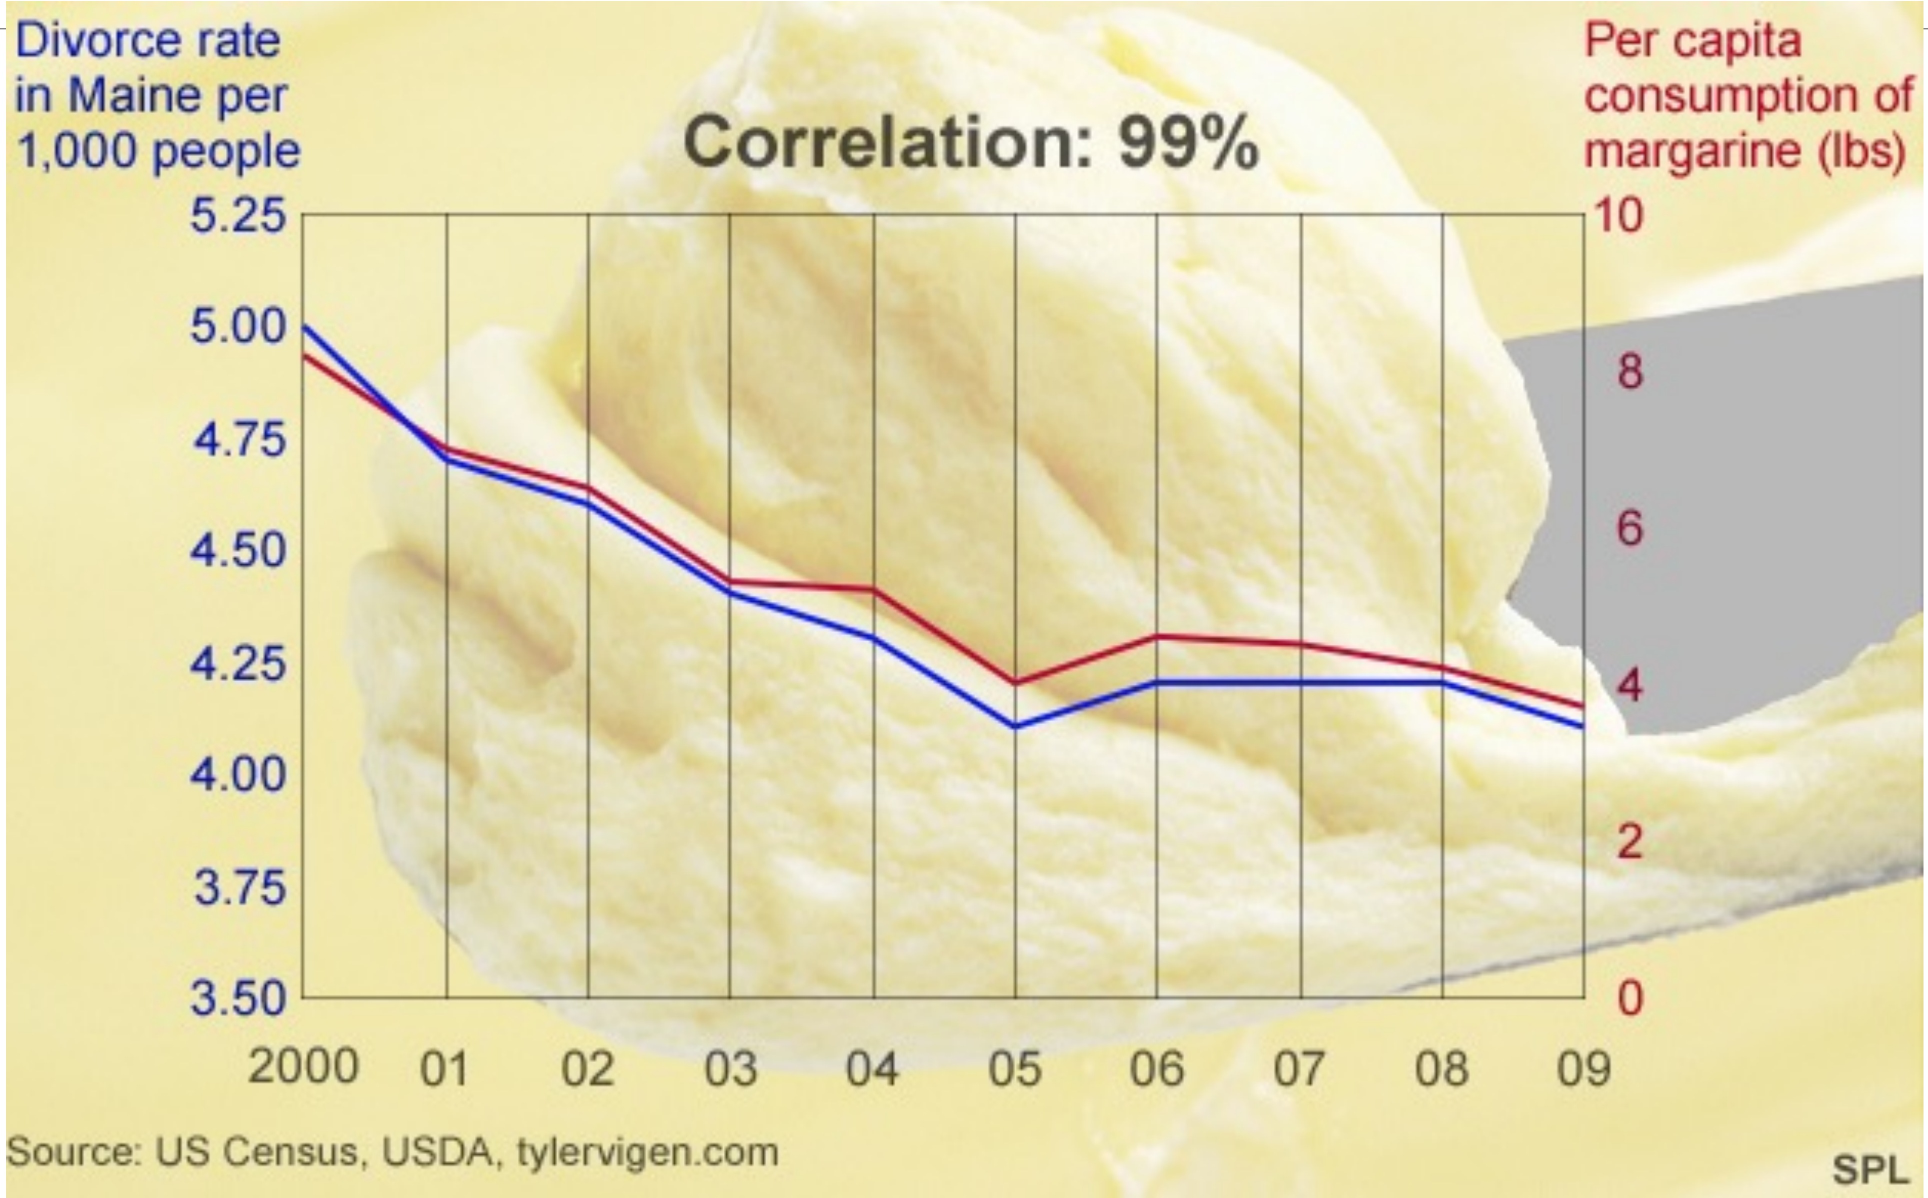

http://www.bbc.com/news/magazine-27537142