



General Inference

CS109, Stanford University

Light Midterm Reflection

Learning Goals

1. Finish conversation on correlations
2. Learn rejection sampling
3. Maybe get to the beta? (What's a beta?)



BAYES NETS!

Where do models come from?

Multiple Random Variables. Start of Digital Revolution



Multiple Random Variables. Start of Digital Revolution

Conditions that match your symptoms

UNDERSTANDING YOUR RESULTS [i](#)

Migraine headache (adult)



Moderate match



Acute Sinusitis



Fair match



Stroke



Fair match



Gender **Male**

Age **30**


[Edit](#)

My Symptoms

[Edit](#)

dizziness, one sided headache

Constructing a Bayesian Network



$$\begin{aligned} P(T = 1 | F_{lu} = 0, U = 0) \\ P(T = 1 | F_{lu} = 0, U = 1) \\ P(T = 1 | F_{lu} = 1, U = 0) \\ P(T = 1 | F_{lu} = 1, U = 1) \end{aligned}$$

In a Bayesian Network,
Each random variable is caused by
its **parents**. Def $P(\text{node} \mid \text{parents})$

- Node: random variable
- Directed edge: causality

Examples:

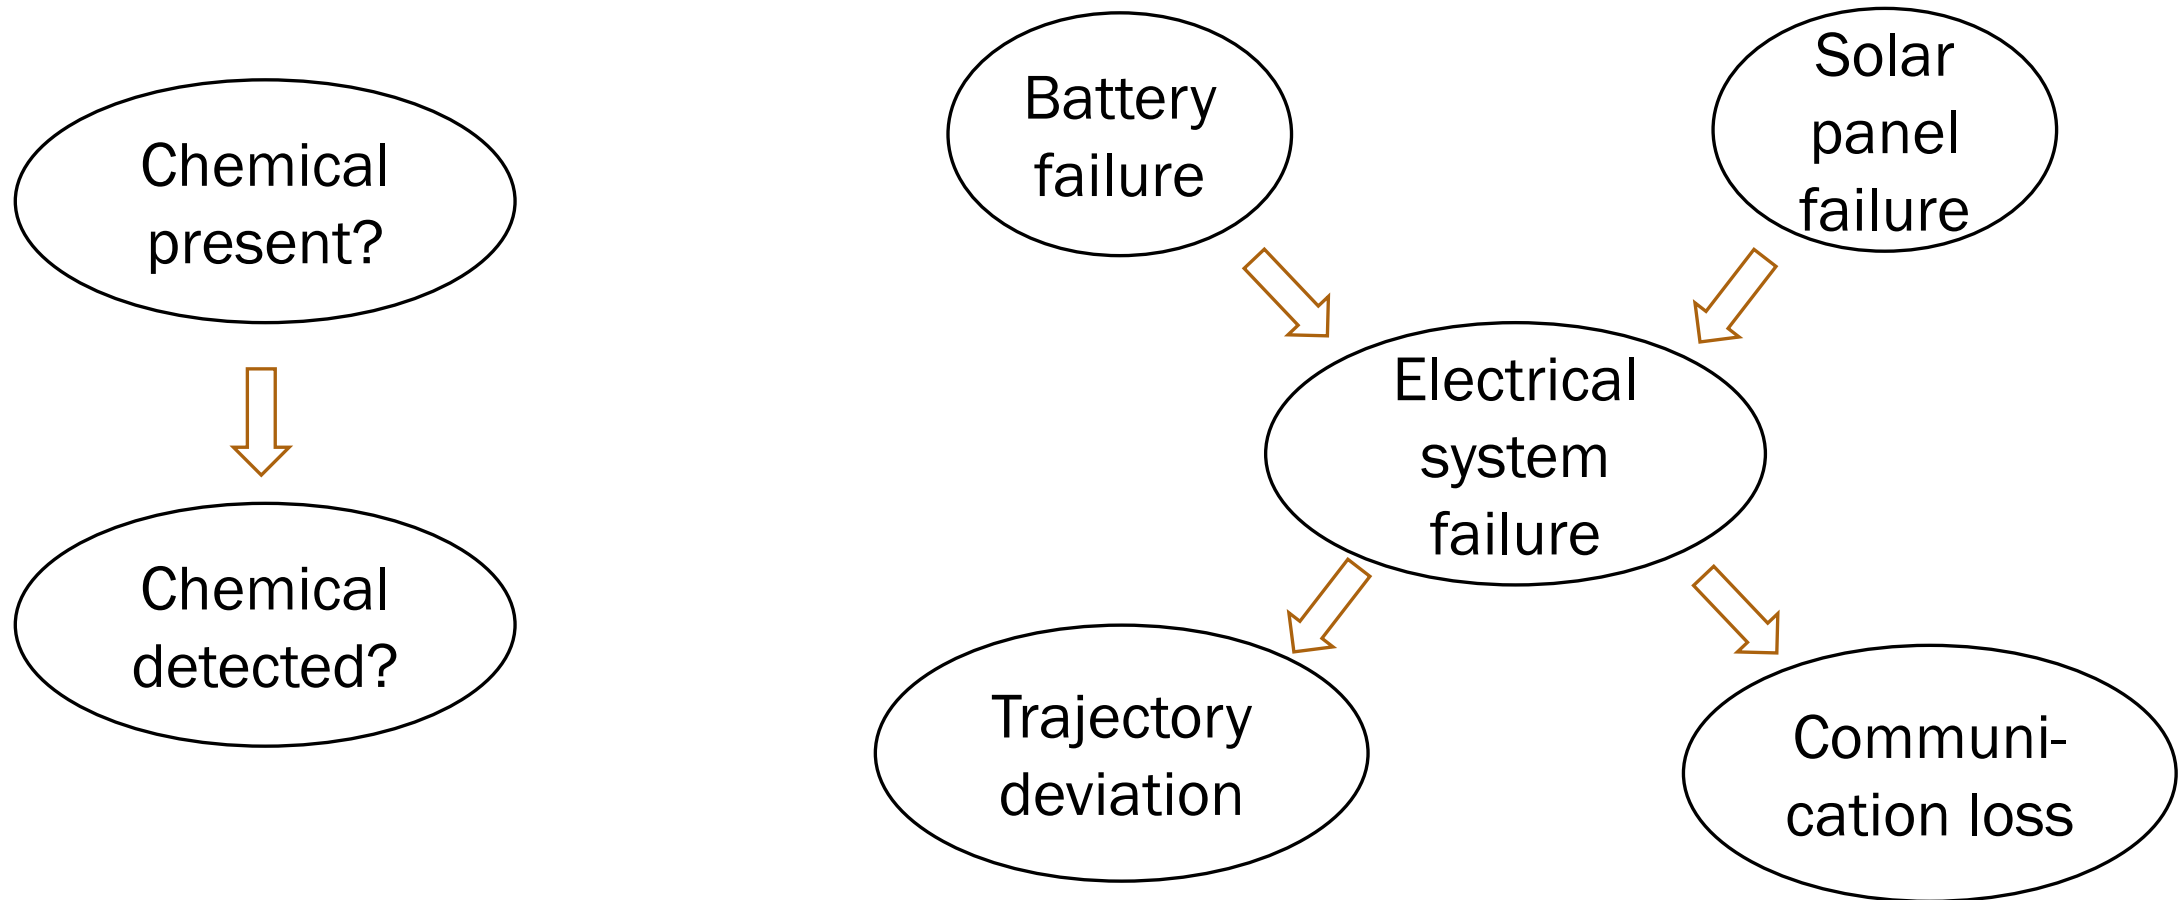
- $P(F_{lu} = 1)$
- $P(U = 0)$
- $P(F_{ev} = 1 | F_{lu} = 1), P(F_{ev} = 1 | F_{lu} = 0)$
- $P(T = 1 | F_{lu} = 0, U = 0) \dots$

Make a *Generative* Model




A good probabilistic model is **generative**. It explains the process through which the joint is **created**.

Other applications



Bayesian Network

Simple Disease Model



```
def get_prob_Xi(x, parents):
```

```
    # what is the probability that Xi = x
```

```
    # given the list parents of assignments to
```

```
    # the parents variables Xi
```

$$P(\text{Uni} = 1) = 0.8$$

$$P(\text{Influenza} = 1 | \text{Uni} = 1) = 0.2$$

$$P(\text{Influenza} = 1 | \text{Uni} = 0) = 0.1$$

$$P(\text{Tired} = 1 | \text{Uni} = 0, \text{Influenza} = 0) = 0.1$$

$$P(\text{Tired} = 1 | \text{Uni} = 1, \text{Influenza} = 0) = 0.8$$

$$P(\text{Fever} = 1 | \text{Influenza} = 1) = 0.9$$


$$P(\text{Fever} = 1 | \text{Influenza} = 0) = 0.05$$

$$P(\text{Tired} = 1 | \text{Uni} = 0, \text{Influenza} = 1) = 0.9$$





$$P(\text{Tired} = 1 | \text{Uni} = 1, \text{Influenza} = 1) = 1.0$$

Bayesian Network Assumption

Simple Disease Model



Order nodes by ancestry



$$P(\text{Joint}) = \prod_i P(x_i | x_{i-1}, \dots, x_1)$$


←

$$= \prod_i P(x_i | \text{Values of parents of } X_i)$$

Assume: Once you know the value of the parents of a variable in your network, X_i , any further information about non-descendants will not change your belief in X_i .

Bayesian Network

Simple Disease Model



Let's go ahead and figure this one out together!

$$P(\text{Fever} = 0, \text{Influenza} = 1, \text{Uni} = 0, \text{Tired} = 0)$$

$$P(\text{Uni} = 1) = 0.8$$

$$P(\text{Influenza} = 1 | \text{Uni} = 1) = 0.2$$

$$P(\text{Influenza} = 1 | \text{Uni} = 0) = 0.1$$

$$P(\text{Tired} = 1 | \text{Uni} = 0, \text{Influenza} = 0) = 0.1$$

$$P(\text{Tired} = 1 | \text{Uni} = 1, \text{Influenza} = 0) = 0.8$$

$$P(\text{Fever} = 1 | \text{Influenza} = 1) = 0.9$$

$$P(\text{Fever} = 1 | \text{Influenza} = 0) = 0.05$$

$$P(\text{Tired} = 1 | \text{Uni} = 0, \text{Influenza} = 1) = 0.9$$

$$P(\text{Tired} = 1 | \text{Uni} = 1, \text{Influenza} = 1) = 1.0$$


End Review

How do people design these?

The art of modelling

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

$$P(U = 1) = 0.8$$



1. Design this

2. Also design this.
Later in CS109: learn
this from data

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

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

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

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


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

$$P(T = 1 | F_{lu} = 1, U = 1) = 1.0$$

Calculate the Covariance / Correlation (new stat!)

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

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



Covariance of Zero Does Not Mean Independence!

X and Y are random variables:

X is -1, 0 or 1 with equal probability

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

Covariance of Zero Does Not Mean Independence!

X and Y are random variables with PMF:

$Y \backslash X$	-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 $XY = 0$, $E[XY] = 0$
- $\text{Cov}(X, Y) = E[XY] - E[X]E[Y] = 0 - 0 = 0$

But, X and Y are clearly dependent!

What is Wrong With This?

Consider the following data:

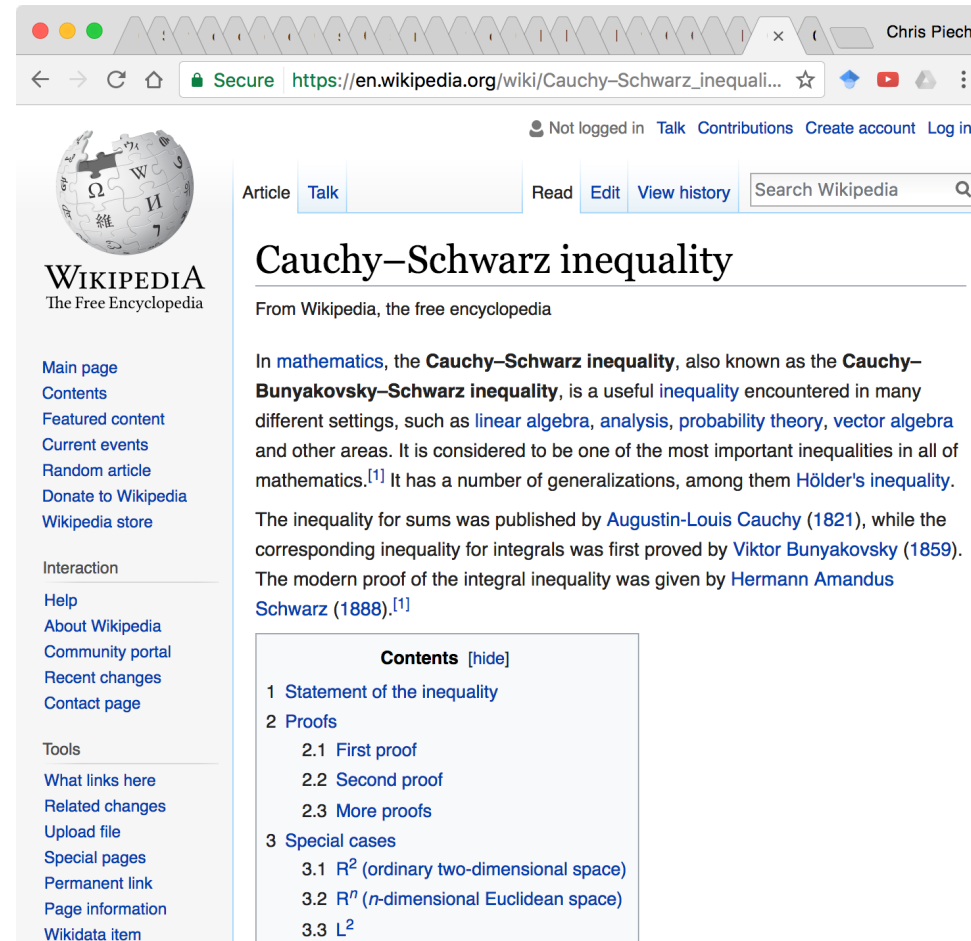
Weight	Height	Weight * Height
64	57	3648
71	59	4189
53	49	2597
67	62	4154
55	51	2805
58	50	2900
77	55	4235
57	48	2736
56	42	2352
51	42	2142
76	61	4636
68	57	3876



$$\begin{array}{lll} E[W] & E[H] & E[W*H] \\ = 62.75 & = 52.75 & = 3355.83 \end{array}$$

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

Cauchy Schwarz, a great way to normalize!



The screenshot shows a web browser window displaying the Wikipedia article for the Cauchy–Schwarz inequality. The browser's address bar shows the URL https://en.wikipedia.org/wiki/Cauchy–Schwarz_inequality. The page header includes the Wikipedia logo and the text "WIKIPEDIA The Free Encyclopedia". The article title is "Cauchy–Schwarz inequality". The main text of the article states: "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]" Below the main text is a "Contents" section with the following items: 1 Statement of the inequality, 2 Proofs (with sub-items 2.1 First proof, 2.2 Second proof, 2.3 More proofs), and 3 Special cases (with sub-items 3.1 R² (ordinary two-dimensional space), 3.2 Rⁿ (n-dimensional Euclidean space), 3.3 L²).

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

Just in case you wanted to see the proof...

Probability theory [\[edit\]](#)

Let X and Y be [random variables](#). Then the covariance inequality^{[17][18]} is given by:

$$\text{Var}(X) \geq \frac{\text{Cov}(X, Y)^2}{\text{Var}(Y)}.$$

After defining an inner product on the set of random variables using the expectation of their product,

$$\langle X, Y \rangle := \mathbf{E}(XY),$$

the Cauchy–Schwarz inequality becomes

$$|\mathbf{E}(XY)|^2 \leq \mathbf{E}(X^2) \mathbf{E}(Y^2).$$

To prove the covariance inequality using the Cauchy–Schwarz inequality, let $\mu = \mathbf{E}(X)$ and $\nu = \mathbf{E}(Y)$, then

$$\begin{aligned} |\text{Cov}(X, Y)|^2 &= |\mathbf{E}((X - \mu)(Y - \nu))|^2 \\ &= |\langle X - \mu, Y - \nu \rangle|^2 \\ &\leq \langle X - \mu, X - \mu \rangle \langle Y - \nu, Y - \nu \rangle \\ &= \mathbf{E}((X - \mu)^2) \mathbf{E}((Y - \nu)^2) \\ &= \text{Var}(X) \text{Var}(Y), \end{aligned}$$

where Var denotes [variance](#) and Cov denotes [covariance](#).

Correlation is just normalized Covariance



Correlation

Covariance

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

The equation is annotated with purple arrows. One arrow points from the word 'Correlation' above to the symbol $\rho(X, Y)$. Another arrow points from the word 'Covariance' above to the $\text{Cov}(X, Y)$ term in the numerator.


It is always true that

$$\text{Cov}(X, Y) < \sqrt{\text{Var}(X)\text{Var}(Y)}$$

$$\text{Cov}(X, Y) > -\sqrt{\text{Var}(X)\text{Var}(Y)}$$

We have models. Need to solve
problems

Inference: Algebra




In a Bayesian Network,
Each random variable is **conditionally independent** of its non-descendants, **given its parents**.

- Node: random variable
- Directed edge: conditional dependency

Inference via math

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

$$P(U = 1) = 0.8$$



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

Compute joint probabilities using chain rule.


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

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

Inference via math

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

$$P(U = 1) = 0.8$$



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

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

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

1. Compute joint probabilities

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

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

2. Definition of conditional probability


$$\frac{P(F_{lu} = 1, F_{ev} = 0, U = 0, T = 1)}{\sum_x P(F_{lu} = x, F_{ev} = 0, U = 0, T = 1)}$$

$$= 0.095$$

Inference via math

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

$$P(U = 1) = 0.8$$



3. $P(F_{lu} = 1 | U = 1, T = 1)$?


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

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

Inference via math

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

$$P(U = 1) = 0.8$$



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

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

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

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

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

$$P(T = 1|F_{lu} = 1, U = 1) = 1.0$$

3. $P(F_{lu} = 1|U = 1, T = 1)$?

1. Compute joint probabilities

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

...

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

2. Definition of conditional probability

$$\frac{\sum_y P(F_{lu} = 1, U = 1, F_{ev} = y, T = 1)}{\sum_x \sum_y P(F_{lu} = x, U = 1, F_{ev} = y, T = 1)}$$


$$= 0.122$$

Rejection sampling algorithm

Step 0:
Have a fully specified
Bayesian Network

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

$$P(U = 1) = 0.8$$



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

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

Alg #0: Straight Math


Too many possible **inference** questions one could ask...

Alg #1: Rejection Sampling

```
3 N_SAMPLES = 100000
4
5 # Program: Joint Sample
6 # -----
7 # we can answer any probability question
8 # with multivariate samples from the joint,
9 # where conditioned variables match
10 def main():
11     obs = getObservation()
12     print 'Observation = ', obs
13
14     samples = sampleATon()
15     prob = probFluGivenObs(samples, obs)
16     print 'Pr(Flu) = ', prob
```

```
71 # Method: Sample A Ton
72 # -----
73 # chose N_SAMPLES with likelihood proportional
74 # to the joint distribution
75 def sampleATon():
76     samples = []
77     for i in range(N_SAMPLES):
78         sample = makeSample()
79         samples.append(sample)
80     return samples
```

Recall: Probabilistic Model



```
82 # Method: Make Sample
83 # -----
84 # chose a single sample from the joint distribut
85 # based on the medical "Probabilistic Graphical
86 def makeSample():
87     # prior on causal factors
88     flu = bern(0.1)
89     und = bern(0.8)
90
91     # choose fever based on flue
92     if flu == 1: fev = bern(0.9)
93     else:       fev = bern(0.05)
94
95     # choose tired based on (undergrade and flu)
96     if und == 1 and flu == 1:   tir = bern(1.0)
97     elif und == 1 and flu == 0: tir = bern(0.8)
98     elif und == 0 and flu == 1: tir = bern(0.9)
99     else:                       tir = bern(0.1)
100
101     # a sample from the joint has an
102     # assignment to *all* random variables
103     return [flu, und, fev, tir]
```

```
82 # Method: Make Sample
83 # -----
84 # chose a single sample from the joint distribut
85 # based on the medical "Probabilistic Graphical
86 def makeSample():
87     # prior on causal factors
88     flu = bern(0.1)
89     und = bern(0.8)
90
91     # choose fever based on flue
92     if flu == 1: fev = bern(0.9)
93     else:       fev = bern(0.05)
94
95     # choose tired based on (undergrade and flu)
96     if und == 1 and flu == 1:   tir = bern(1.0)
97     elif und == 1 and flu == 0: tir = bern(0.8)
98     elif und == 0 and flu == 1: tir = bern(0.9)
99     else:                       tir = bern(0.1)
100
101     # a sample from the joint has an
102     # assignment to *all* random variables
103     return [flu, und, fev, tir]
```

```

82 # Method: Make Sample
83 # -----
84 # chose a single sample from the joint distribut
85 # based on the medical "Probabilistic Graphical
86 def makeSample():
87     # prior on causal factors
88     flu = bern(0.1)
89     und = bern(0.8)
90
91     # choose fever based on flue
92     if flu == 1: fev = bern(0.9)
93     else:         fev = bern(0.05)
94
95     # choose tired based on (undergrade and flu)
96     if und == 1 and flu == 1:   tir = bern(1.0)
97     elif und == 1 and flu == 0: tir = bern(0.8)
98     elif und == 0 and flu == 1: tir = bern(0.9)
99     else:                       tir = bern(0.1)
100
101     # a sample from the joint has an
102     # assignment to *all* random variables
103     return [flu, und, fev, tir]
104

```

```
82 # Method: Make Sample
83 # -----
84 # chose a single sample from the joint distribut
85 # based on the medical "Probabilistic Graphical
86 def makeSample():
87     # prior on causal factors
88     flu = bern(0.1)
89     und = bern(0.8)
90
91     # choose fever based on flue
92     if flu == 1: fev = bern(0.9)
93     else:         fev = bern(0.05)
94
95     # choose tired based on (undergrade and flu)
96     if und == 1 and flu == 1:   tir = bern(1.0)
97     elif und == 1 and flu == 0: tir = bern(0.8)
98     elif und == 0 and flu == 1: tir = bern(0.9)
99     else:                       tir = bern(0.1)
100
101     # a sample from the joint has an
102     # assignment to *all* random variables
103     return [flu, und, fev, tir]
```

Alg #1: Rejection Sampling

```
1
2
3 N_SAMPLES = 100000
4
5 # Program: Joint Sample
6 # -----
7 # we can answer any pro
8 # with multivariate sam
9 # where conditioned var
10 def main():
11     obs = getObservatio
12     print 'Observation
13
14     samples = sampleATo
15     prob = probFluGiven
16     print 'Pr(Flu) = ',
17
```

```
webMd — -bash — 30x20
[0, 1, 0, 1]
[1, 1, 1, 1]
[0, 1, 0, 1]
[0, 1, 0, 0]
[0, 1, 0, 0]
[0, 1, 0, 1]
[0, 1, 0, 1]
[0, 1, 0, 1]
[0, 0, 0, 0]
[0, 0, 0, 0]
[0, 1, 0, 1]
[0, 1, 0, 1]
[0, 1, 0, 1]
[0, 1, 0, 1]
[0, 1, 0, 1]
[0, 1, 0, 1]
[0, 1, 0, 0]
[0, 0, 0, 0]
[0, 1, 0, 1]
[0, 1, 0, 1]
[0, 1, 0, 1]
[1, 1, 0, 1]
```

Alg #1: Rejection Sampling

```
3 N_SAMPLES = 100000
4
5 # Program: Joint Sample
6 # -----
7 # we can answer any probability question
8 # with multivariate samples from the joint,
9 # where conditioned variables match
10 def main():
11     obs = getObservation()
12     print 'Observation = ', obs
13
14     samples = sampleATon()
15     prob = probFluGivenObs(samples, obs)
16     print 'Pr(Flu) = ', prob
```

```
25 # Method: Probability of Flu Given Observation
26 # -----
27 # Calculate the probability of flu given many
28 # samples from the joint distribution and a set
29 # of observations to condition on.
30 def probFluGivenObs(samples, obs):
31     # reject all samples which don't align
32     # with condition
33     keepSamples = []
34     for sample in samples:
35         if checkObsMatch(sample, obs):
36             keepSamples.append(sample)
37
38     # from remaining, simply count...
39     fluCount = 0
40     for sample in keepSamples:
41         [flu, und, fev, tir] = sample
42         if flu == 1:
43             fluCount += 1
44
45     # counting can be so sweet...
46     return float(fluCount) / len(keepSamples)
```

```

25 # Method: Probability of Flu Given Observation
26 # -----
27 # Calculate the probability of flu given many
28 # samples from the joint distribution and a set
29 # of observations to condition on.
30 def probFluGivenObs(samples, obs):
31     # reject all samples which don't align
32     # with condition
33     keepSamples = []
34     for sample in samples:
35         if checkObsMatch(sample, obs):
36             keepSamples.append(sample)
37
38     # from remaining, simply count...
39     fluCount = 0
40     for sample in keepSamples:
41         [flu, und, fev, tir] = sample
42         if flu == 1:
43             fluCount += 1
44
45     # counting can be so sweet...
46     return float(fluCount) / len(keepSamples)

```

```
25 # Method: Probability of Flu Given Observation
26 # -----
27 # Calculate the probability of flu given many
28 # samples from the joint distribution and a set
29 # of observations to condition on.
30 def probFluGivenObs(samples, obs):
31     # reject all samples which don't align
32     # with condition
33     keepSamples = []
34     for sample in samples:
35         if checkObsMatch(sample, obs):
36             keepSamples.append(sample)
37
38     # from remaining, simply count...
39     fluCount = 0
40     for sample in keepSamples:
41         [flu, und, fev, tir] = sample
42         if flu == 1:
43             fluCount += 1
44
45     # counting can be so sweet...
46     return float(fluCount) / len(keepSamples)
```


Lets try it!

BACK ←
TO **CODE**
THE

Rejection sampling algorithm

Inference question: What is $P(F_{lu} = 1 | U = 1, T = 1)$?

$$\text{probability} \approx \frac{\# \text{ samples with } (F_{lu} = 1, U = 1, T = 1)}{\# \text{ samples with } (U = 1, T = 1)}$$

Why would this definition of approximate probability make sense?



Why would this approximate probability make sense?


Inference question: What is $P(F_{lu} = 1 | U = 1, T = 1)$?

$$\text{probability} \approx \frac{\# \text{ samples with } (F_{lu} = 1, U = 1, T = 1)}{\# \text{ samples with } (U = 1, T = 1)}$$


Recall our definition of probability as a frequency: $P(E) = \lim_{n \rightarrow \infty} \frac{n(E)}{n}$ $n = \#$ of total trials
 $n(E) = \#$ trials where E occurs

```
webMd — -bash — 39x20
[0, 1, 1, 0]
[1, 0, 1, 1]
[0, 1, 0, 1]
[0, 1, 0, 0]
[0, 1, 0, 0]
[0, 1, 1, 0]
[1, 1, 1, 1]
[0, 1, 0, 0]
[0, 0, 0, 1]
[0, 1, 0, 1]
[0, 1, 0, 1]
[0, 1, 0, 1]
[0, 1, 0, 0]
[0, 1, 0, 1]
[0, 1, 0, 0]
[0, 0, 0, 0]
[0, 0, 0, 1]
Observation = [None, None, None, None]
Pr(Flu | Obs) = 0.10164
>
```

If you can sample enough from the joint distribution, you can answer any probability question



Each one of these is one joint sample:
[Flu, Undergrad, Fever, Tired]



Where are we in CS109?

Overview of Topics



Counting
Theory



Core
Probability



Random
Variables



Probabilistic
Models



Uncertainty
Theory



Machine
Learning



Which video are you more likely to like?

Davie504




👍 10,000 👎 50

Not Davie504



👍 10 👎 0



Which drug should you give if you are uncertain about p ?

Drug A



Drug B



Which one do you give to a patient?

Philosophical Ponderings:

You ask about the probability of rain tomorrow.

Person A: My leg itches when it rains and its kind of itchy.... Uh, $p = .80$

Person B: I have done complex calculations and have seen 10,451 days like tomorrow... $p = 0.80$

What is the difference between the two estimates?

*“Those who are able to
represent what they do not
know make better decisions”
- CS109*

Today we are going to learn
something unintuitive, beautiful and
useful

Let's play a game!

Flip a plate 5 times. If you get heads 3 times you win



Credit: Rembrandt via Dall E

$$\begin{aligned} P(X = 3) &= \binom{5}{3} \cdot \frac{1}{2}^3 \cdot \frac{1}{2}^2 \\ &= 0.3125 \end{aligned}$$

What if you don't know a probability?



What if you don't know a probability?




What is your belief that you flip a heads
on my coin?



The parameter p to a binomial can be a random variable

Coffee out of 10 Bevs.....What is your Belief of the $P(\text{Coffee})$?



9 Heads out of 10 Flips. What is your Belief in p ?



9 Heads out of 10 Flips. What is your Belief in p ?

Let X be our belief about the probability of heads:

Binomial \rightarrow

$$\cancel{P}(X = x | H = 9, T = 1) = \frac{P(H = 9, T = 1 | X = x) f(X = x)}{P(H = 9, T = 1)}$$

Uniform?

9 Heads out of 10 Flips. What is your Belief in p ?

Let X be our belief about the probability of heads:

Binomial \rightarrow


$$\begin{aligned} & \cancel{P}(X = x | H = 9, T = 1) \\ &= \frac{P(H = 9, T = 1 | X = x) f(X = x)}{P(H = 9, T = 1)} \quad \leftarrow \text{Uniform?} \\ &= \frac{\binom{10}{9} x^9 (1 - x)^1}{P(H = 9, T = 1)} \end{aligned}$$

9 Heads out of 10 Flips. What is your Belief in p ?

Let X be our belief about the probability of heads:


Binomial \rightarrow

$$\begin{aligned} & \cancel{P}(X = x | H = 9, T = 1) \\ &= \frac{P(H = 9, T = 1 | X = x) f(X = x)}{P(H = 9, T = 1)} \quad \leftarrow \text{Uniform?} \\ &= \frac{\binom{10}{9} x^9 (1 - x)^1}{P(H = 9, T = 1)} \\ &= K \cdot x^9 (1 - x)^1 \end{aligned}$$



9 Heads out of 10 Flips. What is your Belief in p ?

$$f(X = x | H = 9, T = 1)$$



Flip a coin with unknown probability

Flip a coin ($n + m$) times, comes up with n heads

- We don't know probability X that coin comes up heads

Frequentist (never prior)

$$X = \lim_{n+m \rightarrow \infty} \frac{n}{n+m}$$
$$\approx \frac{n}{n+m}$$

X is (often) a single value

Bayesian (prior is great)

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

X is a random variable. Leads to a belief distribution which captures confidence

Flip a coin with unknown probability!

Flip a coin ($n + m$) times, comes up with n heads

- We don't know probability X that coin comes up heads
- Our belief before flipping coins is that: $X \sim \text{Uni}(0, 1)$
- Let N = number of heads
- Given $X = x$, coin flips independent: $(N \mid X) \sim \text{Bin}(n + m, x)$

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

Bayesian
"posterior"
probability distribution

Bayesian "prior"
probability distribution

Flip a coin with unknown probability!

Flip a coin $(n + m)$ times, comes up with n heads

- We don't know probability X that coin comes up heads
- Our belief before flipping coins is that: $X \sim \text{Uni}(0, 1)$
- Let N = number of heads
- Given $X = x$, coin flips independent: $(N | X) \sim \text{Bin}(n + m, x)$

$$\begin{aligned} f_{X|N}(x|n) &= \frac{P(N = n | X = x) f_X(x)}{P(N = n)} \quad 1 \\ &= \frac{\binom{n+m}{n} x^n (1-x)^m}{P(N = n)} \\ &= \frac{\binom{n+m}{n}}{P(N = n)} x^n (1-x)^m \\ &= \frac{1}{c} \cdot x^n (1-x)^m \quad \text{where } c = \int_0^1 x^n (1-x)^m dx \end{aligned}$$

Binomial

Move terms around

Flip a coin with unknown probability!



If you start with a $X \sim \text{Uni}(0, 1)$ prior over probability, and observe:

n “successes” and
 m “failures”...

Your new belief about the probability is:


$$f_X(x) = \frac{1}{c} \cdot x^n (1 - x)^m$$

where $c = \int_0^1 x^n (1 - x)^m$

Belief after 7 success and 1 fail

$$f_X(x) = \frac{1}{c} \cdot x^n (1-x)^m$$

$n=7$ $m=1$



Equivalently!



If you start with a $X \sim \text{Uni}(0, 1)$ prior over probability, and observe:

let $a = \text{num "successes"} + 1$

let $b = \text{num "failures"} + 1$

Your new belief about the probability is:

$$f_X(x) = \frac{1}{c} \cdot x^{a-1} (1-x)^{b-1}$$

where $c = \int_0^1 x^{a-1} (1-x)^{b-1}$




Beta Random Variable

X is a **Beta Random Variable**: $X \sim \text{Beta}(a, b)$

- Probability Density Function (PDF): (where $a, b > 0$)

$$f(x) = \begin{cases} \frac{1}{B(a,b)} x^{a-1} (1-x)^{b-1} & 0 < x < 1 \\ 0 & \text{otherwise} \end{cases}$$

$$B(a,b) = \int_0^1 x^{a-1} (1-x)^{b-1} dx$$




- Symmetric when $a = b$


$$E[X] = \frac{a}{a+b}$$

$$\text{Var}(X) = \frac{ab}{(a+b)^2(a+b+1)}$$

Beta is the Random Variable for Probabilities



Used to represent a distributed belief of a probability





Beta Parameters *can*
come from experiments:

$$a = \text{“successes”} + 1$$

$$b = \text{“failures”} + 1$$

And with that, we'll end for today!

See you Friday!