

**Agenda:**

**0. Probability review**

**1. Naïve Bayes: another classifier (used for, e.g., Spam Assassin)**

**2. How Google does translation**

**3. Beyond naïve Bayes: the maximum entropy stewpot**

**0. Basics of probability (review & pictures)**

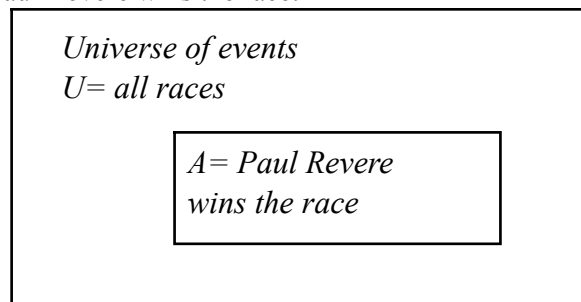
The fundamentals of probability theory: the axioms of probability. Why are these important? The power of the purse: Because while there are *other* attempts to handle the notion of ‘uncertainty’, e.g., ‘fuzzy logic’, ‘3-valued logic’, etc., these axioms are the **only** system with the property that if you **gamble** with them, you **cannot** be unfairly exploited by an opponent who uses some other system (Di Finetti, 1932 theorem).

So, some first concepts.

We say that  $A$  is a **random variable** if  $A$  denotes an event and there is some uncertainty if  $A$  is true.

Typically, we let  $U$  denote the **universe** of all possible events (= all “possible worlds”). Then a subset of  $U$ , call it  $A$ , corresponds to the set of events in which  $A$  is true.

**Example.** Let the universe  $U$  be the set of all horse races. Let *Paul Revere* (abbreviation: P-R) be a horse. Then we can let  $A$  denote the set of racing events in which Paul Revere wins. We can draw this as a picture, where *races* labels the outer square, the universe, and the circle inside is the set of all events where Paul Revere wins the race:



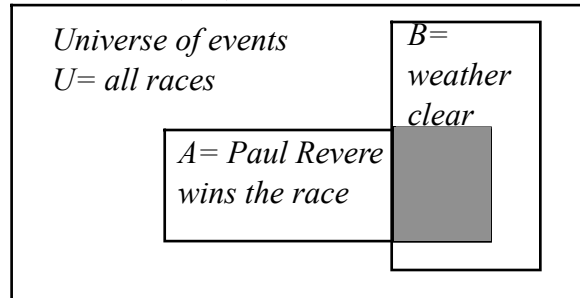
Let us denote by  $P(A)$  the fraction of events (possible worlds in the universe of events) in which  $A$  turns out true. We could spend the next 2 hours on the philosophy of possible worlds and this business. But we won't.

We will compute probabilities using an informal notion of *areas* (formally, we'd use measure theory).

The Universe of all events has total area 1,  $P(U)=1$ , because it denotes all the events that are true.  $P(A)$  then is the area of the smaller rectangle with respect to  $U$  (= the fraction of the total universe in which Paul Revere wins).  $P(\neg A)$  = the races in which Paul Revere does **not** win = the set difference between  $U$  and  $A$ . From this we will posit 3 axioms regarding  $P(A)$ :

- (1)  $0 \leq P(A) \leq 1$  [because: the area of  $A$  cannot be  $< 0$  or  $> 1$  ]
- (2)  $P(\text{true})=1$
- (3)  $P(\text{false})=0$
- (4)  $P(A \vee B) = P(A) + P(B) - P(A, B)$  [where  $\vee$  means “or”, i.e., **either**  $A$  or  $B$  must be true;  $+$  means “add together”, and the comma in  $A, B$  means “and”, i.e., both  $A$  **and**  $B$  must be true]

To see how this last axiom works, let's look at the racing universe with event  $A$ = Paul Revere wins and a second event,  $B$ = the weather is clear. The **shaded area** represents the fraction of events when **both**  $A$  and  $B$  are true, i.e.,  $P(A,B)$ = true:




It should be apparent that in order to figure out the probability of  $A$  **or**  $B$ , we need to add up the areas corresponding to  $A$  and to  $B$ , but then subtract out the shaded area so that it is not counted twice. In this way, we arrive at the formula for the probability of  $A$  **or**  $B$ .

We next turn to the notion of **conditional probability**.

We let  $P(A|B)$  denote the fraction of events/possible worlds in which  $B$  is true, and then *also* have  $A$  true. That is, we 'shrink' the universe from  $U$  down to  $B$ , focusing in on a subset possibly more relevant to our situation, and use *that* as our basis to calculate probabilities.

**Example.** In the figure below, we illustrate the following situation. Let  $H$ = probability that "I have a headache";  $F$ = probability that "I am getting the flu". These are denoted by the rectangles  $H$  and  $F$  in the figure below. Let us assume:

$P(H) = 1/10$ ;  $P(F)=1/40$ . Now let's compute the conditional probability  $P(H|F)$ , i.e., the probability that I have a headache **given** that I have the flu. This is the fraction of flu-events that are also headache events – that is, if we just look at the rectangle  $F$ , what proportion of  $F$  overlaps with  $H$ ? (The answer is  $1/2$ ). Thus,  $P(H|F)=1/2$ .



In other words, to find  $P(H|F)$ , we compute:

(# worlds in which  $H$  **and**  $F$  are true)/(# worlds in which  $F$  is true) or,  
 (area  $H$  **and**  $F$ )/(area of  $F$ ), or  
 $P(H, F)/P(F)$

So this is the **formula for conditional probability**:

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

Note how  $P(B)$  is in the denominator here. Multiplying out, we obtain the important formula called the **chain rule** which we will use in the naïve Bayes classifier:

$$P(A,B) = P(A|B) \cdot P(B)$$

Some other manipulations of conditional probability will be used in what follows. We consider two: (i) simplifications to the *right* of the conditioning bar symbol |; and (ii) simplifications to the *left* of the conditioning bar symbol.

Simplifications to the *right* of the bar:

Suppose we have *lots* of conditions to impose on whether or not Paul Revere wins. For example, this could depend on not only if the weather's clear, but also whether the jockey's brother is a friend of mine, whether Paul Revere won its last race, etc. In other words:

$$P(\text{Paul Revere wins} \mid \text{weather clear, jockey's brother a friend, P-R won last race})$$

With more factors then, we have less *bias*, because we are focusing in on our particular situation, but we will have more *variance*, because it will become harder and harder to measure all these terms perfectly. So, sometimes we will want to reduce the number of factors to the right of the conditioning symbol to those we are more confident we can estimate; this is called *back off*. (We will see this in action soon). There is no problem in simply doing this:

$$P(\text{Paul Revere wins} \mid \text{weather clear, } \cancel{\text{jockey's brother a friend, P-R won last race}})$$

And then of course just having  $P(\text{Paul Revere wins} \mid \text{weather clear})$  remaining. But what about if there are more terms to the *left* of the bar, as in this case:

$$P(\text{Paul Revere wins, Valentine loses, Eitaph loses} \mid \text{weather clear})$$

Note that if we *add* terms to the left the probability should get lower and lower every time we add a new factor. (Why? Think about intersection.) If we just care about Paul Revere, are we then allowed to simply strike out the other two horses, this way?

$$P(\text{Paul Revere wins, } \cancel{\text{Valentine loses, Eitaph loses}} \mid \text{weather clear})$$

The answer is: No! We need to carry out a more complex expansion to isolate Paul Revere on the left. To see how, let's abbreviate Paul Revere wins as  $R$ , Valentine loses as  $V$ , Eitaph loses as  $E$ , and the Weather is clear as  $W$ . Then our conditional probability:

$$P(\text{Paul Revere wins, Valentine loses, Eitaph loses} \mid \text{weather clear})$$

Can be abbreviated as:

$$\frac{P(R,V,E,W)}{P(W)}$$

We can use this formula to derive the **chain rule for conditional probability**:

$$P(\text{Paul Revere wins, Valentine loses, Eitaph loses} \mid \text{weather clear}) = P(\text{Paul Revere wins} \mid \text{Valentine loses, Eitaph loses, weather clear}) \times P(\text{Valentine loses} \mid \text{Eitaph loses, weather clear}) \times P(\text{Eitaph loses} \mid \text{weather clear})$$

Proof. Writing out the 3 terms:

$$\frac{P(R,V,E,W)}{P(W)} = \frac{P(R,V,E,W)}{P(V,E,W)} \times \frac{P(V,E,W)}{P(E,W)} \times \frac{P(E,W)}{P(W)}$$

Now, supposed it is the case that the following simpler expansion holds:


$$P(\text{Paul Revere wins, Valentine loses, Eitaph loses} \mid \text{weather clear}) = P(\text{Paul Revere wins} \mid \cancel{\text{Valentine loses, Eitaph loses}}, \text{weather clear}) \times P(\text{Valentine loses} \mid \cancel{\text{Eitaph loses}}, \text{weather clear}) \times P(\text{Eitaph loses} \mid \text{weather clear})$$

In this case, whether Paul Revere wins or not depends *only* on whether the weather's clear...and not on what the other two horses do. They are irrelevant factors, so we can strike them out. In this case, when the probability is *unchanged* when we drop out conditioning factors, we say that the probability is **conditionally independent** (independent of the other horses, but still conditioned on the weather). More generally, if there are  $n$  factors  $f$ , and each factor is independent of the other, but still dependent on a condition  $c$ , we can write the following, which will be another key ingredient in our naïve Bayes classifier model:


$$P(f_1, \dots, f_n | c) = P(f_1 | c) \times \dots \times P(f_n | c)$$

That is, we can just write out the probability as the product of the  $n$  factors, **assuming they are independent from one another** (the outcomes of these events do not affect the outcomes of one another); note the factors are still dependent on the outcome of event  $c$ .

OK, we come to the last ingredient we shall need, **Bayes' Law**. Again we can illustrate this with the simple picture of headache and flu as before. Recall  $P(H)=1/10$ ;  $P(F)=1/40$ ,  $P(H|F)=1/2$ .



Now we will **label** each of the distinct regions in this diagram,  $A$ ,  $B$ , and  $C$ , as follows.  $A+B$ =area of  $F$ ;  $B+C$ = area of  $H$ :



By the definition of conditional probability,  $P(H|F) = P(H, F) / P(F) = B / (A+B)$ .

Now consider this reasoning: one day you wake up with a headache, and you think, OMG, 50% of flus are associated with headaches, so now I have a 50-50 chance of getting the flu.” Is this reasoning correct?

What we *want* to compute is:  $P(H|F)$ . We already know the *other* conditional probability, that of headache given the flu. Further, by the definition of conditional probability, in terms of the regions  $A$ ,  $B$ , and  $C$ , we have that:  $P(F|H) = B / (B+C)$ . To find this last ratio of regions, we can take the conditional probability  $P(F|H) = B / (A+B)$ , and multiply it by  $(A+B) / (B+C)$ , as follows:

$$\frac{B}{B+C} = \frac{B}{A+B} \cdot \frac{A+B}{B+C} \text{ i.e.,}$$

$$P(F|H) = P(H|F) \cdot \frac{P(F)}{P(H)}$$

$$\text{in our example, } \frac{1/2 \times 1/40}{1/10} = \frac{1/80}{1/10} = \frac{1}{8}$$

The term  $P(F)$  is called the **prior probability** (of getting the flu); the term  $P(H|F)$  is called the **likelihood**; the term  $P(H)$  is the **evidence** (e.g., that you have a headache); and the term  $P(F|H)$  is

called the **posterior probability** of getting the flu (given that you have a headache). So this updated probability is a kind of learning: given the fact (data) that you indeed have a headache, how does the probability of getting the flu change? (It increases from 1/40 to 1/8.) Inverting from  $P(H|F)$  to  $P(F|H)$  is called **Bayes' Law**. It follows from a very simple manipulation of the definition of conditional probability and then application of the chain rule, i.e., that  $P(A,B)=P(A|B)\times P(B)$ :

$$P(B|A) = \frac{P(A,B)}{P(A)} \text{ (by dfn of conditional probability)}$$


$$= \frac{P(A|B) \cdot P(B)}{P(A)} \text{ (by chain rule, replacing } P(A,B)\text{)}$$

Or in words we can say this:

$$\text{posterior} = \frac{\text{likelihood} \times \text{prior}}{\text{evidence}}$$

Now let's put this all to work to build a classifier called **Naïve Bayes**. Like k-means and ID-trees, and Boosting, etc., this will take as input the values of some **features** and then output a classification **label**.

As our example, we will use the common, but valuable task of classifying email into 1 of 2 categories: either good email ("ham") or bad email ("spam"). The underlying probability model follows what is called a **Bayes' net**. We can imagine the following generative process: we pick a label, e.g., "ham", and given this label, email documents of this type will have a certain distribution of feature values  $f_1, \dots, f_n$ . If we pick the other label, "spam", we will get another distribution for the feature values (hopefully distinct). So the picture looks like this, and the idea of course is that **given a new email**, we would like to figure out whether it is ham or spam:



Crucially, we assume that **the features are independent from one another**. (This is the "naïve" part of Naïve Bayes.) Their values depend on (are conditioned on) **only** the value of the label. That is why we draw the networks as above, with **no links** between the features, only from the label directed down to the features.

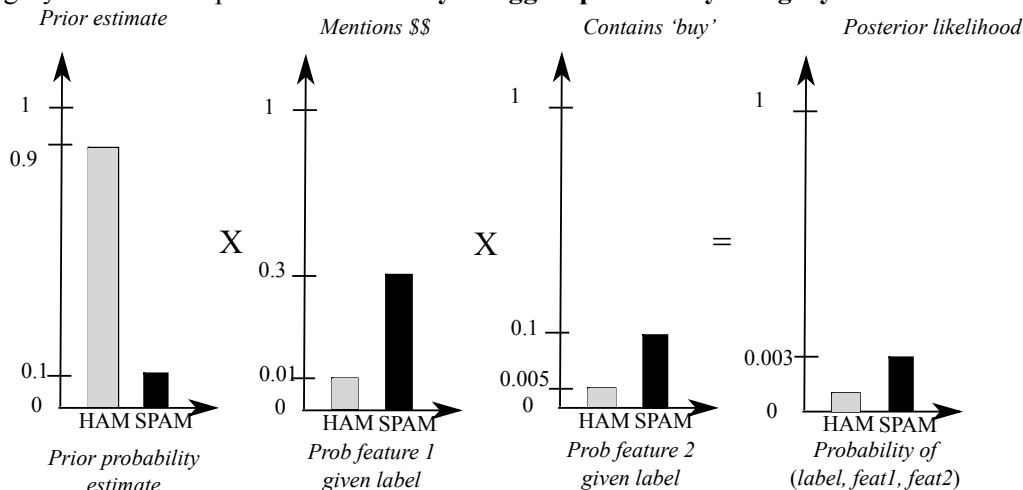
Now here's the idea behind the classification.. Suppose we have estimated that 90% of our email is "ham" (OK), and that 10% is "spam". This gives us our **prior probability estimates**  $P(\text{label}=\text{ham})=0.9$  and  $P(\text{label}=\text{spam})=0.1$ . That's what we can say about any new email **without any additional information**. (We'll see below how we get these estimates.)

Now, when we get a new email, we will get the values of its **features** and use these to adjust the prior probabilities, as with our headache example. (In our example, to keep things simple, we will use only two features.)

So, this new email comes along: *"Buy this amazing new Ginsu knife for only \$39....."*  
Is this ham or spam? We'll assume that we use the following 2 features:

- Feature 1: The email mentions money; this occurs in 30% of spam, and in 1% of ham
- Feature 2: The email contains the word 'buy'; this occurs in 10% of spam, and in 0.5% of ham

We can picture our calculation as follows: our initial prior probabilities for each category are adjusted by **multiplying** the contribution **each feature** ‘votes’ (independently) as to how likely each category is. Then we pick the **most likely = biggest probability category** at the end:



So, in this case, our new email is classified as “spam” because this yields the largest posterior likelihood. Note how we got this value. It is simply this:

$$P(\text{label}) \times P(f_1 | \text{label}) \times P(f_2 | \text{label}) = P(\text{label}, f_1, f_2) \quad [\text{recall from dfn of conditional prob that:}]$$

$$\frac{P(\text{label}, f_1, f_2)}{P(\text{label})} = P(f_1 | \text{label}) \times P(f_2 | \text{label}) \quad \text{IF } f_1, f_2 \text{ are independent of one another}$$

In other words, we multiply the following out to find the label likelihood, and pick the biggest likelihood:

$$\text{Prior probability of a label} \times \text{Probability of feature contributions} = \text{Posterior label likelihood}$$

In our case, for the two labels “ham” and “spam”:

	Prior	×	Pr(featl (\$)  lable)	×	Pr(featl 2 ('buy') label)	=	Label likelihood
Ham:	0.9	×	0.01	×	0.005	=	0.000045 (log of this likelihood: -4.34)
Spam:	0.1	×	0.30	×	0.10	=	0.00303 (log of this likelihood: -2.52)

So, our email is more likely to be spam than ham. In fact, taking the ratios of the log likelihoods, -2.52/-4.32, the email is about 2 orders of magnitude (100x) more likely to be spam than ham.

Recall that: (1) the features **must** be independent of one another; (2) we can add other features, of course...this is what a program like Spam Assassin can do, by training; and (3) one can use this method with lots more categories to **classify** documents (see the end of the handout).

Let’s turn to justifying this approach probabilistically, as well as how we actually estimate the probability values above, via training, and highlighting some pitfalls.

First, why is this justified? We are computing the **maximum** probability that an input email will have a particular label (category), **given** that it has a particular set of features. We pick the label that maximizes:  $P(l=\text{value} | \text{observed features})$ . Let’s follow out this logic. We are maximizing the following quantity over label values:

$$\max P(\text{label} | \text{features}) = \max \frac{P(\text{features}, \text{label})}{P(\text{features})} \quad [\text{by dfn of conditional probability}]$$

But note that the denominator in the expression above,  $P(\text{features}) = P(f_1, \dots, f_n)$  is *constant* no matter what our choice of label value. So, to maximize the above quantity, it suffices to maximize the numerator:

$$\max P(\text{features}, \text{label}) = P(f_1, \dots, f_n, \text{label})$$

By the chain rule, this quantity in turn is just:

$$\max P(\text{label}) \times P(f_1, \dots, f_n | \text{label})$$

But given that the features are all independent of one another, this is the same as (recall our Paul Revere example!):

$$\begin{aligned} & \max P(\text{label}) \times P(f_1 | \text{label}) \times \dots \times P(f_n | \text{label}) \\ & \max_{\text{prior}} \times \text{'vote' } f_1 \quad \times \dots \times \text{'vote' } f_n \end{aligned}$$

Putting this down as a formula, we have:

$$\arg \max_C P(C | f_1, \dots, f_n) = \arg \max_C \frac{P(C) \prod_{i=1}^n P(f_i | C)}{P(f_1, \dots, f_n)} = \arg \max_C P(C) \prod_{i=1}^n P(f_i | C)$$

This is exactly the computation we have carried out. It remains to figure out how we ‘train’ our classifier – that is, how do we get the various estimates of the probabilities above? The simplest thing is just to estimate them from counts in training text, that is, known examples of ham and spam emails. These are the so-called *maximum likelihood estimates*:

$$P(\text{label} = \text{ham}) = \frac{\text{count}(\# \text{ ham emails})}{\text{count}(\text{total} \# \text{ emails})} \quad P(\text{label} = \text{spam}) = \frac{\text{count}(\# \text{ spam emails})}{\text{count}(\text{total} \# \text{ emails})}$$

$$P(f_1 | \text{label} = \text{ham}) = \frac{\text{count}(\# \text{ ham emails mention } \$)}{\text{count}(\text{total} \# \text{ ham emails})}$$

$$P(f_1 | \text{label} = \text{spam}) = \frac{\text{count}(\# \text{ spam emails mention } \$)}{\text{count}(\text{total} \# \text{ spam emails})}$$

$$P(f_2 | \text{label} = \text{ham}) = \frac{\text{count}(\# \text{ ham emails contain 'buy'})}{\text{count}(\text{total} \# \text{ ham emails})}$$

$$P(f_2 | \text{label} = \text{spam}) = \frac{\text{count}(\# \text{ spam emails contain 'buy'})}{\text{count}(\text{total} \# \text{ spam emails})}$$

So this is how we get the estimates. For example, if we have 1000 emails, 900/1000 are ham, and 100/1000 are spam. Of the 100 spam emails, 30/100 mention money, and 1/100 contain ‘buy’. For ham emails, 1/100 mention money and 5/1000 contain ‘buy’.

Note that as the # of data samples (amount of training data) increases, then our estimates should get better; one of the properties of the maximum likelihood estimates is that they will converge to the ‘true’ values as the amount of data goes to infinity. (The mean approaches the true average.) But, if the # of training examples is small, our estimate will be very lousy, and have more noise (variance); there are a variety of things we can do to improve this, but that’s for a machine learning course.

### A second worked example:

MIT decides to use surveys to determine how to sort students into dorms. They decide to use Naive Bayes and survey data from current residents to classify where to put future students.

To collect this “training data”, they surveyed 30 random students.

Each surveyed student is asked to fill out a simple questionnaire with 3 true/false questions.

0. Which dorm do you live in: {East Campus, West Campus, or FSILG}
1. Are a Pyro – i.e. do you enjoy performing feats with fire (or inadvertently trigger fire alarms)?
2. Are you a foreign student or do you like studying foreign languages?
3. Are you in Good shape?

Here are the results. It turns out that our random survey gave us exactly 10 students from each dorm group.

	Pyro	ForeignLang	GoodShape	# surveyed
East Campus	8/10	1/10	3/10	10 10/30
West Campus	3/10	6/10	3/10	10 10/30
FSILG	1/10	3/10	8/10	10 10/30

What can you do this data? We can use these counts to make estimates of the following probabilities:

$P(C)$  (the prior probability of being in any dorm)

$P(f_i | C)$  (the likelihood of having one of the 3 features **given** being in a particular dorm)

E.g.  $P(\text{Pyro}=\text{True} | C=\text{East Campus}) = 8 / 10$      $P(\text{Language} = \text{True} | C= \text{FSILG}) = 3/10$

Now we can use these probability estimates to classify new students by applying Bayes rule, i.e., our formula:

$$\arg \max_C P(C | f_1, \dots, f_n) = \arg \max_C P(C) \prod_{i=1}^n P(f_i | C)$$

Question 1: where would a new student who loves foreign languages most likely be classified if they filled in their incoming survey as follows:

Pyro = *True*

ForeignLang = *False*

GoodShape = *False*

To do this, we compute  $P(C_i | \text{Pyro}=\text{True}, \text{ForeignLang}=\text{False}, \text{Goodshape}=\text{False})$  for all three possible campuses, and find the largest one! (That is what the “arg max” part means.)

For  $C= \text{East campus}$ :

$$\begin{aligned} & \operatorname{argmax} P(C=\text{East} | P=T, F=F, G=F) \\ &= \operatorname{argmax} P(C=\text{East}) * [P(P=T | C=\text{East}) P(L=F|C=\text{East}) P(G=F|C=\text{East}) ] \\ &= (10/30) * [(8/10) (1-1/10)(1-3/10) ] = 1/3 * [(8*9*7)/1000] = 1/3 * [1000] \\ &= 0.072000 \end{aligned}$$

For  $C= \text{West campus}$ :

$$\begin{aligned} & \operatorname{argmax} P(C=\text{West} | P=T, F=F, G=F) \\ &= \operatorname{argmax} P(C=\text{West}) * [P(P=T | C=\text{West}) P(L=F|C=\text{West}) P(G=F|C=\text{West}) ] \\ &= (10/30) * [ (3/10) (1-6/10)(1-3/10) ] \\ &= 1/3 * [ 3*4*7/1000 ] = 1/3 * [84/1000] \\ &= 0.028000 \end{aligned}$$

For  $C= \text{FSILG}$ :

$$\begin{aligned} & \operatorname{argmax} P(C=\text{FSILG} | P=T, F=F, G=F) \\ &= P(C=\text{FSILG}) * [P(P=T|C=\text{FSILG}) P(L=F|C=\text{FSILG}) P(G=F|C=\text{FSILG}) ] \\ &= (10/30) * [(1/10) (1-3/10)(1-8/10)] \\ &= 1/3 * [ (1*7*2*)/1000 ] = 1/3 * [14/1000] = 1/3[14/1000] \\ &= 0.004667 \end{aligned}$$



The largest value for such a student (Pyros *true*, all other attributes, *false*) is **East Campus**.

Question 2. What about an all-round student who checks all the boxes in the incoming survey?  
 $P(C=? \mid \text{Pyro} = \text{True}, \text{ForeignLang} = \text{True}, \text{GoodShape} = \text{True})$

$$\begin{aligned} P(C=\text{East} \mid P=T, F=T, G=T) &= P(C=\text{East}) * [P(P=T \mid C=\text{East})P(L=T \mid C=\text{East})P(G=T \mid C=\text{East})] \\ &= 10/30 * [(8/10)(1/10)(3/10)] \\ &= 1/3 * [8*1*3/1000] = 1/3 * [24/1000] \\ &= 0.008000 \end{aligned}$$

$$\begin{aligned} P(C=\text{West Campus} \mid P=T, F=T, G=T) &= P(C=\text{West}) * [P(P=F \mid C=\text{West})P(L=T \mid C=\text{West})P(G=T \mid C=\text{West})] \\ &= 10/30 * [(3/10)(6/10)(3/10)] \\ &= 1/3 * [3*6*3/1000] = 1/3 * [54/1000] \\ &= \end{aligned}$$

$$\begin{aligned} P(C=\text{FSILG} \mid P=T, F=T, G=T) &= P(C=\text{FSILG}) * [P(P=T \mid C=\text{FSILG})P(L=T \mid C=\text{FSILG})P(G=T \mid C=\text{FSILG})] \\ &= (1/3) * [(1/10)(3/10)(8/10)] \\ &= 1/3 * [1*3*8/1000] = 1/3 * [24/1000] \\ &= 0.008000 \end{aligned}$$

The maximum  $C$  is **West Campus**.

In Naive Bayes, the  $P(C=\text{some value})$  is also known as the “prior”. Knowledge about the prior probabilities can help us distinguish what proportion to assign to each class. In our case we got lucky and it just happened that each campus got 10 students, so the prior in this case is *Uniform*.

### Estimation & its discontents

There is at least one particular case about estimating the probabilities from data counts that we should note. Suppose a particular count is actually 0 – that is, we *never* observe a particular feature associated with a particular label – this will happen especially if we keep adding more and more features. In this case, note that the entire probability product to find the likelihood will *all* be zero, just because one of the estimates is 0. So this is very bad!

There is a whole cottage industry devoted to fixing this problem, and it is called *smoothing*. It is basically the Robin Hood strategy: we rob probability mass from the rich and give it to the poor. In particular, the *simplest* smoothing strategy, invented by Laplace, is called *add-1 smoothing*: if a count is 0, we add 1 to it, so that, e.g., 0/100 goes to 1/100. (We must also *subtract* the appropriate probability mass, i.e., counts, from the *rest* of our estimates, so that the probabilities still add up to 1 in all.)

A second method of smoothing (probability mass redistribution) is due to Alan Turing. He figured this out when he was developing probability formulas for estimating the likelihood of finding German submarines in particular areas of the ocean. What if a submarine had *never* been observed in a particular spot? (Something that’s actually quite likely!) Turing reasoned that a fairly good probability estimate of ‘things never seen’ would be quite close to the estimate of ‘things seen *exactly* once’. This method, now called Good-Turing smoothing (only published until decades after WWII), works well but is finicky. There are whole books devoted to this subject, for machine learning and especially in natural language processing, where we quickly get word sequences never seen before.

One more thing. You may note that in our calculation we multiply together a (possibly long) string of probabilities, one for each feature. With a 1000 features, this value will quickly get very,

very small. So, the usual method is to operate in log space, where multiplication is just addition, so we can maintain accuracy. (That’s why we used log likelihoods above.)

### Beyond Naïve Bayes (Optional)

OK, this method is fine so far as it goes, but it can be improved enormously. Here we will just sketch one method, known as **maximum entropy classification** that can gobble down any set of features, even if they are not independent. Yet remarkably, as first shown by Jaynes (1957), it is the most probabilistically sound method of **combining diverse features**. It rationalizes the general notion of just ‘scoring’ features and adding them up. We won’t prove this here, but just indicate the general approach, which is now broadly used in, e.g., figuring out the part of speech labels in text. (For instance, in the sentence, *police police police*, is the first *police* a Noun or a Verb?)

1. To begin, let’s assume there are now 10 labels for documents, with categories *A, B, C, D, E, F, G, H, I, J*. (So, e.g., category *A* could be travel; *B* sports; *C* business; etc.) If we know this, and **no other information** then given an email *m*, what is our best guess for category *C* (business) given this email, i.e.,  $P(C | m)$ ?

The maximum entropy approach would claim it is 1/10: that is, we maximize the quantity in each of the 10 bins, uniformly, by spreading out the total probability mass of 1 among 10 bins.

2. Now suppose I tell you that 55% of all emails are in category *A*, travel? Now what is the quantity  $P(C | m)$ ? I think it should not be too hard to see that *A* gobbles up 0.55 of the probability mass, leaving 0.45 to be distributed evenly over the remaining 9 categories, or 0.05 for each of the remaining categories, including category *C*, business. So the maximum entropy estimate for  $P(C | m)$  is 0.05.

3. Now suppose I add *another* constraint: that *in addition* to the fact in (2), we know that 10% of all emails contain the word ‘buy’. What is  $P(C | m)$  now? This gets harder to visualize, so we’ll write it out as a table, where the first row is the probability of containing ‘buy’ (which thus must add up to 0.1 of all emails), and second row is the probability of not containing ‘buy’, which we have labeled *other* (which thus must add up to 0.9). Once again following the maximum entropy idea, since we don’t know anything else about the ‘contains buy’ row, we should distribute its 0.1 total *evenly* among the 10 bins, thus giving 0.01 to each. Next, since *all* of category *A* must add up to 0.55, and since the ‘contains buy’ cell holds 0.01, it must be that the cell in the row labeled *other* and in column *A* must have the value 0.54 (so that the column total is 0.55). That leaves  $0.9 - 0.54 = 0.36$  for the rest of the 9 bins in the *other* row. Once again, spreading this evenly, we get  $0.36/9 = 0.04$  for each of these bins (so that each column here adds to 0.05). Thus we have the following table:

	1	2	3	4	5	6	7	8	9	10
	A	B	C	D	E	F	G	H	I	J
<i>buy</i>	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
<i>other</i>	0.54	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04

So, *why* is this called *maximum entropy*? You should realize that by spreading out the values evenly, we are *maximizing the entropy of the cell values*:  $-p \log p$  summed over all entries is at a maximum. (Below we indicate *why* this is a good thing to do.) In any case, we are maximizing the entropy *subject to the constraints specified*. (We have two so far.)

4. So let’s add one more constraint. Suppose that in addition, 80% of the ‘buy’ emails are in *either* category *A* or category *C*. Now we want to figure out  $P(C | m)$ . Gulp! This one is much harder to figure out – in fact, in general to do this, it is like spreadsheets, but we can indicate what has to be true in our table now: the probability of the *buy* row, column *A*, plus the probability in the *buy* row, column *C*, must add up to 0.08 (80% of the 10%). That turns out to be the values 0.051 and 0.029. Since that leaves 0.020 for the rest of the bins in the *buy* row, these must be  $0.020/8=0.0025$ . Since column *A* must still add up to 0.55, then that leaves 0.499 for row *other*, column *A*. Since

the *other* row must still sum to 0.9, we have  $0.9 - 0.499 = 0.401$  to distribute evenly over the rest of the *other* bins, so this is  $0.401/9 = 0.0446$ . If we impose these constraints, you'll see that this is the answer (we don't say how we figured it out!)

	1	2	3	4	5	6	7	8	9	10
	A	B	C	D	E	F	G	H	I	J
<i>buy</i>	0.051	.0025	0.029	.0025	.0025	.0025	.0025	.0025	.0025	.0025
<i>other</i>	0.499	.0446	.0446	.0446	.0446	.0446	.0446	.0446	.0446	.0446

Now we know that  $P(\text{buy}, C) = 0.029$ ;  $P(C | \text{buy}) = 0.29 (= 0.029/0.1)$ ;  $P(A | \text{buy}) = 0.51$ .

This is our classifier, a *maximum entropy* classifier.

The punchline. While there are many possible distributions that could yield the three observed constraints, that 55% of the emails are in category *A*, that 10% of the emails contain *buy*, and that of these 10%, 80% are in category *A* or *C*, the **one distribution** that we picked, where we have **maximized** the entropy of the probability mass subject to these constraints, turns out to be the **only one** having the following two properties, the second one quite remarkable:

1. This distribution follows the form:  $P(\text{email}, \text{label}) = \frac{1}{Z(\lambda)} \exp \sum_i \lambda_i f_i(\text{email}, \text{label})$  where the lambdas are the weights associated with each feature  $f_i$ ; the function  $f_i$  returns 1 if the feature is in the email, and 0 otherwise; and  $Z$  is a normalizing constant to make sure the probabilities all add up to 1.
2. This distribution **maximizes the probability of the training data**,  $\prod_j P(\text{email}_j, \text{label}_j)$

**This is what justifies the method!**