

Chapter 3

Generative models for discrete data

3.1 Generative classifier

$$p(y = c|\mathbf{x}, \boldsymbol{\theta}) = \frac{p(y = c|\boldsymbol{\theta})p(\mathbf{x}|y = c, \boldsymbol{\theta})}{\sum_{c'} p(y = c'|\boldsymbol{\theta})p(\mathbf{x}|y = c', \boldsymbol{\theta})} \quad (3.1)$$

This is called a **generative classifier**, since it specifies how to generate the data using the **class conditional density** $p(\mathbf{x}|y = c)$ and the class prior $p(y = c)$. An alternative approach is to directly fit the class posterior, $p(y = c|\mathbf{x})$; this is known as a **discriminative classifier**.

3.2 Bayesian concept learning

Psychological research has shown that people can learn concepts from positive examples alone (Xu and Tenenbaum 2007).

We can think of learning the meaning of a word as equivalent to **concept learning**, which in turn is equivalent to binary classification. To see this, define $f(\mathbf{x}) = 1$ if \mathbf{x} is an example of the concept C , and $f(\mathbf{x}) = 0$ otherwise. Then the goal is to learn the indicator function f , which just defines which elements are in the set C .

3.2.1 Likelihood

$$p(\mathcal{D}|h) \triangleq \left(\frac{1}{\text{size}(h)}\right)^N = \left(\frac{1}{|h|}\right)^N \quad (3.2)$$

This crucial equation embodies what Tenenbaum calls the **size principle**, which means the model favours the simplest (smallest) hypothesis consistent with the data. This is more commonly known as **Occams razor**¹⁴.

3.2.2 Prior

The prior is decided by human, not machines, so it is subjective. The subjectivity of the prior is controversial. For example, that a child and a math professor will reach dif-

ferent answers. In fact, they presumably not only have different priors, but also different hypothesis spaces. However, we can finesse that by defining the hypothesis space of the child and the math professor to be the same, and then setting the child's prior weight to be zero on certain advanced concepts. Thus there is no sharp distinction between the prior and the hypothesis space.

However, the prior is the mechanism by which background knowledge can be brought to bear on a problem. Without this, rapid learning (i.e., from small samples sizes) is impossible.

3.2.3 Posterior

The posterior is simply the likelihood times the prior, normalized.

$$p(h|\mathcal{D}) \triangleq \frac{p(\mathcal{D}|h)p(h)}{\sum_{h' \in \mathcal{H}} p(\mathcal{D}|h')p(h')} = \frac{\mathbb{I}(\mathcal{D} \in h)p(h)}{\sum_{h' \in \mathcal{H}} \mathbb{I}(\mathcal{D} \in h')p(h')} \quad (3.3)$$

where $\mathbb{I}(\mathcal{D} \in h)p(h)$ is 1 **iff** (iff and only if) all the data are in the extension of the hypothesis h .

In general, when we have enough data, the posterior $p(h|\mathcal{D})$ becomes peaked on a single concept, namely the MAP estimate, i.e.,

$$p(h|\mathcal{D}) \rightarrow \hat{h}^{MAP} \quad (3.4)$$

where \hat{h}^{MAP} is the posterior mode,

$$\begin{aligned} \hat{h}^{MAP} &\triangleq \arg \max_h p(h|\mathcal{D}) = \arg \max_h p(\mathcal{D}|h)p(h) \\ &= \arg \max_h [\log p(\mathcal{D}|h) + \log p(h)] \end{aligned} \quad (3.5)$$

Since the likelihood term depends exponentially on N , and the prior stays constant, as we get more and more data, the MAP estimate converges towards the **maximum likelihood estimate** or **MLE**:

$$\hat{h}^{MLE} \triangleq \arg \max_h p(\mathcal{D}|h) = \arg \max_h \log p(\mathcal{D}|h) \quad (3.6)$$

In other words, if we have enough data, we see that the **data overwhelms the prior**.

¹⁴ http://en.wikipedia.org/wiki/Occam%27s_razor

3.2.4 Posterior predictive distribution

The concept of **posterior predictive distribution**¹⁵ is normally used in a Bayesian context, where it makes use of the entire posterior distribution of the parameters given the observed data to yield a probability distribution over an interval rather than simply a point estimate.

$$p(\tilde{x}|\mathcal{D}) \triangleq \mathbb{E}_{h|\mathcal{D}}[p(\tilde{x}|h)] = \begin{cases} \sum_h p(\tilde{x}|h)p(h|\mathcal{D}) \\ \int p(\tilde{x}|h)p(h|\mathcal{D})dh \end{cases} \quad (3.7)$$

This is just a weighted average of the predictions of each individual hypothesis and is called **Bayes model averaging**(Hoeting et al. 1999).

3.3 The beta-binomial model

3.3.1 Likelihood

Given $X \sim \text{Bin}(\theta)$, the likelihood of \mathcal{D} is given by

$$p(\mathcal{D}|\theta) = \text{Bin}(N_1|N, \theta) \quad (3.8)$$

3.3.2 Prior

$$\text{Beta}(\theta|a, b) \propto \theta^{a-1}(1-\theta)^{b-1} \quad (3.9)$$

The parameters of the prior are called **hyper-parameters**.

3.3.3 Posterior

$$\begin{aligned} p(\theta|\mathcal{D}) &\propto \text{Bin}(N_1|N_1 + N_0, \theta)\text{Beta}(\theta|a, b) \\ &= \text{Beta}(\theta|N_1 + a, N_0 + b) \end{aligned} \quad (3.10)$$

Note that updating the posterior sequentially is equivalent to updating in a single batch. To see this, suppose we have two data sets \mathcal{D}_a and \mathcal{D}_b with sufficient statistics N_1^a, N_0^a and N_1^b, N_0^b . Let $N_1 = N_1^a + N_1^b$ and $N_0 = N_0^a + N_0^b$ be the sufficient statistics of the combined datasets. In batch mode we have

$$\begin{aligned} p(\theta|\mathcal{D}_a, \mathcal{D}_b) &= p(\theta, \mathcal{D}_b|\mathcal{D}_a)p(\mathcal{D}_a) \\ &\propto p(\theta, \mathcal{D}_b|\mathcal{D}_a) \\ &= p(\mathcal{D}_b, \theta|\mathcal{D}_a) \\ &= p(\mathcal{D}_b|\theta)p(\theta|\mathcal{D}_a) \end{aligned}$$

Combine Equation 3.10 and 2.19

$$\begin{aligned} &= \text{Bin}(N_1^b|\theta, N_1^b + N_0^b)\text{Beta}(\theta|N_1^a + a, N_0^a + b) \\ &= \text{Beta}(\theta|N_1^a + N_1^b + a, N_0^a + N_0^b + b) \end{aligned}$$

This makes Bayesian inference particularly well-suited to **online learning**, as we will see later.

3.3.3.1 Posterior mean and mode

From Table 2.7, the posterior mean is given by

$$\bar{\theta} = \frac{a + N_1}{a + b + N} \quad (3.11)$$

The mode is given by

$$\hat{\theta}_{MAP} = \frac{a + N_1 - 1}{a + b + N - 2} \quad (3.12)$$

If we use a uniform prior, then the MAP estimate reduces to the MLE,

$$\hat{\theta}_{MLE} = \frac{N_1}{N} \quad (3.13)$$

We will now show that the posterior mean is convex combination of the prior mean and the MLE, which captures the notion that the posterior is a compromise between what we previously believed and what the data is telling us.

3.3.3.2 Posterior variance

The mean and mode are point estimates, but it is useful to know how much we can trust them. The variance of the posterior is one way to measure this. The variance of the Beta posterior is given by

$$\text{var}(\theta|\mathcal{D}) = \frac{(a + N_1)(b + N_0)}{(a + N_1 + b + N_0)^2(a + N_1 + b + N_0 + 1)} \quad (3.14)$$

We can simplify this formidable expression in the case that $N \gg a, b$, to get

$$\text{var}(\theta|\mathcal{D}) \approx \frac{N_1 N_0}{NNN} = \frac{\hat{\theta}_{MLE}(1 - \hat{\theta}_{MLE})}{N} \quad (3.15)$$

¹⁵ http://en.wikipedia.org/wiki/Posterior_predictive_distribution

3.3.4 Posterior predictive distribution

So far, we have been focusing on inference of the unknown parameter(s). Let us now turn our attention to prediction of future observable data.

Consider predicting the probability of heads in a single future trial under a $\text{Beta}(a, b)$ posterior. We have

$$\begin{aligned} p(\tilde{x}|\mathcal{D}) &= \int_0^1 p(\tilde{x}|\theta)p(\theta|\mathcal{D})d\theta \\ &= \int_0^1 \theta \text{Beta}(\theta|a, b)d\theta \\ &= \mathbb{E}[\theta|\mathcal{D}] = \frac{a}{a+b} \end{aligned} \quad (3.16)$$

3.3.4.1 Overfitting and the black swan paradox

Let us now derive a simple Bayesian solution to the problem. We will use a uniform prior, so $a = b = 1$. In this case, plugging in the posterior mean gives **Laplace's rule of succession**

$$p(\tilde{x}|\mathcal{D}) = \frac{N_1 + 1}{N_0 + N_1 + 1} \quad (3.17)$$

This justifies the common practice of adding 1 to the empirical counts, normalizing and then plugging them in, a technique known as **add-one smoothing**. (Note that plugging in the MAP parameters would not have this smoothing effect, since the mode becomes the MLE if $a = b = 1$, see Section 3.3.3.1.)

3.3.4.2 Predicting the outcome of multiple future trials

Suppose now we were interested in predicting the number of heads, \tilde{x} , in M future trials. This is given by

$$p(\tilde{x}|\mathcal{D}) = \int_0^1 \text{Bin}(\tilde{x}|M, \theta)\text{Beta}(\theta|a, b)d\theta \quad (3.18)$$

$$= \binom{M}{\tilde{x}} \frac{1}{B(a, b)} \int_0^1 \theta^{\tilde{x}}(1-\theta)^{M-\tilde{x}}\theta^{a-1}(1-\theta)^{b-1}d\theta \quad (3.19)$$

We recognize the integral as the normalization constant for a $\text{Beta}(a + \tilde{x}, M\tilde{x} + b)$ distribution. Hence

$$\int_0^1 \theta^{\tilde{x}}(1-\theta)^{M-\tilde{x}}\theta^{a-1}(1-\theta)^{b-1}d\theta = B(\tilde{x} + a, M - \tilde{x} + b) \quad (3.20)$$

Thus we find that the posterior predictive is given by the following, known as the (compound) **beta-binomial distribution**:

$$Bb(x|a, b, M) \triangleq \binom{M}{x} \frac{B(x+a, M-x+b)}{B(a, b)} \quad (3.21)$$

This distribution has the following mean and variance

$$\text{mean} = M \frac{a}{a+b}, \text{var} = \frac{Mab}{(a+b)^2} \frac{a+b+M}{a+b+1} \quad (3.22)$$

This process is illustrated in Figure 3.1. We start with a $\text{Beta}(2, 2)$ prior, and plot the posterior predictive density after seeing $N_1 = 3$ heads and $N_0 = 17$ tails. Figure 3.1(b) plots a plug-in approximation using a MAP estimate. We see that the Bayesian prediction has longer tails, spreading its probability mass more widely, and is therefore less prone to overfitting and blackswan type paradoxes.

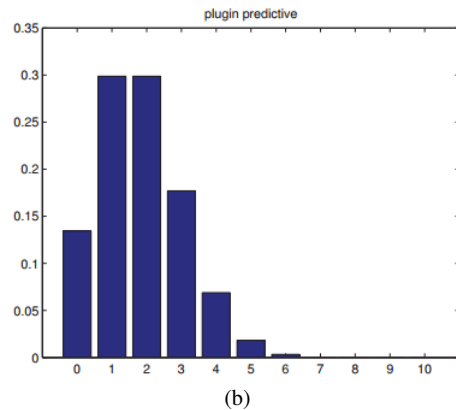



Fig. 3.1: (a) Posterior predictive distributions after seeing $N_1 = 3, N_0 = 17$. (b) MAP estimation.

3.4 The Dirichlet-multinomial model

In the previous section, we discussed how to infer the probability that a coin comes up heads. In this section,

we generalize these results to infer the probability that a dice with K sides comes up as face k .

3.4.1 Likelihood

Suppose we observe N dice rolls, $\mathcal{D} = \{x_1, x_2, \dots, x_N\}$, where $x_i \in \{1, 2, \dots, K\}$. The likelihood has the form

$$p(\mathcal{D}|\boldsymbol{\theta}) = \binom{N}{N_1 \dots N_K} \prod_{k=1}^K \theta_k^{N_k} \quad \text{where } N_k = \sum_{i=1}^N \mathbb{I}(y_i = k) \quad (3.23)$$

almost the same as Equation 2.21.

3.4.2 Prior

$$\text{Dir}(\boldsymbol{\theta}|\boldsymbol{\alpha}) = \frac{1}{B(\boldsymbol{\alpha})} \prod_{k=1}^K \theta_k^{\alpha_k - 1} \mathbb{I}(\boldsymbol{\theta} \in S_K) \quad (3.24)$$

3.4.3 Posterior

$$p(\boldsymbol{\theta}|\mathcal{D}) \propto p(\mathcal{D}|\boldsymbol{\theta})p(\boldsymbol{\theta}) \quad (3.25)$$

$$\propto \prod_{k=1}^K \theta_k^{N_k} \theta_k^{\alpha_k - 1} = \prod_{k=1}^K \theta_k^{N_k + \alpha_k - 1} \quad (3.26)$$

$$= \text{Dir}(\boldsymbol{\theta}|\boldsymbol{\alpha}_1 + N_1, \dots, \boldsymbol{\alpha}_K + N_K) \quad (3.27)$$

From Equation 2.38, the MAP estimate is given by

$$\hat{\theta}_k = \frac{N_k + \alpha_k - 1}{N + \alpha_0 - K} \quad (3.28)$$

If we use a uniform prior, $\alpha_k = 1$, we recover the MLE:

$$\hat{\theta}_k = \frac{N_k}{N} \quad (3.29)$$

3.4.4 Posterior predictive distribution

The posterior predictive distribution for a single multinoulli trial is given by the following expression:

$$p(X = j|\mathcal{D}) = \int p(X = j|\boldsymbol{\theta})p(\boldsymbol{\theta}|\mathcal{D})d\boldsymbol{\theta} \quad (3.30)$$

$$= \int p(X = j|\boldsymbol{\theta}_j) \left[\int p(\boldsymbol{\theta}_{-j}, \boldsymbol{\theta}_j|\mathcal{D})d\boldsymbol{\theta}_{-j} \right] d\boldsymbol{\theta}_j \quad (3.31)$$

$$= \int \boldsymbol{\theta}_j p(\boldsymbol{\theta}_j|\mathcal{D})d\boldsymbol{\theta}_j = \mathbb{E}[\boldsymbol{\theta}_j|\mathcal{D}] = \frac{\boldsymbol{\alpha}_j + N_j}{\boldsymbol{\alpha}_0 + N} \quad (3.32)$$

where $\boldsymbol{\theta}_{-j}$ are all the components of $\boldsymbol{\theta}$ except $\boldsymbol{\theta}_j$.

The above expression avoids the zero-count problem. In fact, this form of Bayesian smoothing is even more important in the multinomial case than the binary case, since the likelihood of data sparsity increases once we start partitioning the data into many categories.

3.5 Naive Bayes classifiers

Assume the features are **conditionally independent** given the class label, then the class conditional density has the following form

$$p(\mathbf{x}|y = c, \boldsymbol{\theta}) = \prod_{j=1}^D p(x_j|y = c, \boldsymbol{\theta}_{jc}) \quad (3.33)$$

The resulting model is called a **naive Bayes classifier**(NBC).

The form of the class-conditional density depends on the type of each feature. We give some possibilities below:

- In the case of real-valued features, we can use the Gaussian distribution: $p(\mathbf{x}|y, \boldsymbol{\theta}) = \prod_{j=1}^D \mathcal{N}(x_j|\boldsymbol{\mu}_{jc}, \boldsymbol{\sigma}_{jc}^2)$, where $\boldsymbol{\mu}_{jc}$ is the mean of feature j in objects of class c , and $\boldsymbol{\sigma}_{jc}^2$ is its variance.
- In the case of binary features, $x_j \in \{0, 1\}$, we can use the Bernoulli distribution: $p(\mathbf{x}|y, \boldsymbol{\theta}) = \prod_{j=1}^D \text{Ber}(x_j|\boldsymbol{\mu}_{jc})$, where $\boldsymbol{\mu}_{jc}$ is the probability that feature j occurs in class c . This is sometimes called the **multivariate Bernoulli naive Bayes** model. We will see an application of this below.
- In the case of categorical features, $x_j \in \{a_{j1}, a_{j2}, \dots, a_{jS_j}\}$, we can use the multinoulli distribution: $p(\mathbf{x}|y, \boldsymbol{\theta}) = \prod_{j=1}^D \text{Cat}(x_j|\boldsymbol{\mu}_{jc})$, where $\boldsymbol{\mu}_{jc}$ is a histogram over the K possible values for x_j in class c .

Obviously we can handle other kinds of features, or use different distributional assumptions. Also, it is easy to mix and match features of different types.

3.5.1 Optimization

We now discuss how to train a naive Bayes classifier. This usually means computing the MLE or the MAP estimate for the parameters. However, we will also discuss how to compute the full posterior, $p(\theta|\mathcal{D})$.

3.5.1.1 MLE for NBC

The probability for a single data case is given by

$$\begin{aligned} p(\mathbf{x}_i, y_i | \theta) &= p(y_i | \pi) \prod_j p(x_{ij} | \theta_j) \\ &= \prod_c \pi_c^{\mathbb{I}(y_i=c)} \prod_j \prod_c p(x_{ij} | \theta_{jc})^{\mathbb{I}(y_i=c)} \end{aligned} \quad (3.34)$$

Hence the log-likelihood is given by

$$p(\mathcal{D} | \theta) = \sum_{c=1}^C N_c \log \pi_c + \sum_{j=1}^D \sum_{c=1}^C \sum_{i: y_i=c} \log p(x_{ij} | \theta_{jc}) \quad (3.35)$$

where $N_c \triangleq \sum_i \mathbb{I}(y_i = c)$ is the number of feature vectors in class c .

We see that this expression decomposes into a series of terms, one concerning π , and DC terms containing the θ_{jc} s. Hence we can optimize all these parameters separately.

From Equation 3.29, the MLE for the class prior is given by

$$\hat{\pi}_c = \frac{N_c}{N} \quad (3.36)$$

The MLE for θ_{jc} s depends on the type of distribution we choose to use for each feature.

In the case of binary features, $x_j \in \{0, 1\}$, $x_j | y = c \sim \text{Ber}(\theta_{jc})$, hence

$$\hat{\theta}_{jc} = \frac{N_{jc}}{N_c} \quad (3.37)$$

where $N_{jc} \triangleq \sum_{i: y_i=c} \mathbb{I}(x_{ij} = 1)$ is the number that feature j occurs in class c .

In the case of categorical features, $x_j \in \{a_{j1}, a_{j2}, \dots, a_{jS_j}\}$, $x_j | y = c \sim \text{Cat}(\theta_{jc})$, hence

$$\hat{\theta}_{jc} = \left(\frac{N_{j1c}}{N_c}, \frac{N_{j2c}}{N_c}, \dots, \frac{N_{jS_j c}}{N_c} \right)^T \quad (3.38)$$

where $N_{jkc} \triangleq \sum_{i=1}^N \mathbb{I}(x_{ij} = a_{jk}, y_i = c)$ is the number that feature $x_j = a_{jk}$ occurs in class c .

3.5.1.2 Bayesian naive Bayes

Use a $\text{Dir}(\alpha)$ prior for π .

In the case of binary features, use a $\text{Beta}(\beta_0, \beta_1)$ prior for each θ_{jc} ; in the case of categorical features, use a $\text{Dir}(\alpha)$ prior for each θ_{jc} . Often we just take $\alpha = \mathbf{1}$ and $\beta = \mathbf{1}$, corresponding to **add-one** or **Laplace smoothing**.

3.5.2 Using the model for prediction

The goal is to compute

$$\begin{aligned} y = f(\mathbf{x}) &= \arg \max_c P(y = c | \mathbf{x}, \theta) \\ &= P(y = c | \theta) \prod_{j=1}^D P(x_j | y = c, \theta) \end{aligned} \quad (3.39)$$

We can estimate parameters using MLE or MAP, then the posterior predictive density is obtained by simply plugging in the parameters $\bar{\theta}$ (MLE) or $\hat{\theta}$ (MAP).

Or we can use BMA, just integrate out the unknown parameters.

3.5.3 The log-sum-exp trick

when using generative classifiers of any kind, computing the posterior over class labels using Equation 3.1 can fail due to **numerical underflow**. The problem is that $p(\mathbf{x} | y = c)$ is often a very small number, especially if \mathbf{x} is a high-dimensional vector. This is because we require that $\sum_{\mathbf{x}} p(\mathbf{x} | y) = 1$, so the probability of observing any particular high-dimensional vector is small. The obvious solution is to take logs when applying Bayes rule, as follows:

$$\log p(y = c | \mathbf{x}, \theta) = b_c - \log \left(\sum_{c'} e^{b_{c'}} \right) \quad (3.40)$$

where $b_c \triangleq \log p(\mathbf{x} | y = c, \theta) + \log p(y = c | \theta)$.

We can factor out the largest term, and just represent the remaining numbers relative to that. For example,

$$\begin{aligned} \log(e^{-120} + e^{-121}) &= \log(e^{-120}(1 + e^{-1})) \\ &= \log(1 + e^{-1}) - 120 \end{aligned} \quad (3.41)$$

In general, we have

$$\sum_c e^{b_c} = \log \left[\left(\sum e^{b_c - B} \right) e^B \right] = \log \left(\sum e^{b_c - B} \right) + B \quad (3.42)$$

where $B \triangleq \max\{b_c\}$.

This is called the **log-sum-exp** trick, and is widely used.

3.5.4 Feature selection using mutual information

Since an NBC is fitting a joint distribution over potentially many features, it can suffer from overfitting. In addition, the run-time cost is $O(D)$, which may be too high for some applications.

One common approach to tackling both of these problems is to perform **feature selection**, to remove irrelevant features that do not help much with the classification problem. The simplest approach to feature selection is to evaluate the relevance of each feature separately, and then take the top K , where K is chosen based on some tradeoff between accuracy and complexity. This approach is known as **variable ranking, filtering, or screening**.

One way to measure relevance is to use mutual information (Section 2.8.3) between feature X_j and the class label Y

$$\mathbb{I}(X_j, Y) = \sum_x \sum_y p(x_j, y) \log \frac{p(x_j, y)}{p(x_j)p(y)} \quad (3.43)$$

If the features are binary, it is easy to show that the MI can be computed as follows

$$\mathbb{I}_j = \sum_c \left[\theta_{jc} \pi_c \log \frac{\theta_{jc}}{\theta_j} + (1 - \theta_{jc}) \pi_c \log \frac{1 - \theta_{jc}}{1 - \theta_j} \right] \quad (3.44)$$

where $\pi_c = p(y = c)$, $\theta_{jc} = p(x_j = 1 | y = c)$, and $\theta_j = p(x_j = 1) = \sum_c \pi_c \theta_{jc}$.

3.5.5 Classifying documents using bag of words

Document classification is the problem of classifying text documents into different categories.

3.5.5.1 Bernoulli product model

One simple approach is to represent each document as a binary vector, which records whether each word is present or not, so $x_{ij} = 1$ iff word j occurs in document i , otherwise $x_{ij} = 0$. We can then use the following class conditional density:

$$\begin{aligned} p(\mathbf{x}_i | y_i = c, \boldsymbol{\theta}) &= \prod_{j=1}^D \text{Ber}(x_{ij} | \theta_{jc}) \\ &= \prod_{j=1}^D \theta_{jc}^{x_{ij}} (1 - \theta_{jc})^{1-x_{ij}} \end{aligned} \quad (3.45)$$

This is called the **Bernoulli product model**, or the **binary independence model**.

3.5.5.2 Multinomial document classifier

However, ignoring the number of times each word occurs in a document loses some information (McCallum and Nigam 1998). A more accurate representation counts the number of occurrences of each word. Specifically, let \mathbf{x}_i be a vector of counts for document i , so $x_{ij} \in \{0, 1, \dots, N_i\}$, where N_i is the number of terms in document i (so $\sum_{j=1}^D x_{ij} = N_i$). For the class conditional densities, we can use a multinomial distribution:

$$p(\mathbf{x}_i | y_i = c, \boldsymbol{\theta}) = \text{Mu}(\mathbf{x}_i | N_i, \boldsymbol{\theta}_c) = \frac{N_i!}{\prod_{j=1}^D x_{ij}!} \prod_{j=1}^D \theta_{jc}^{x_{ij}} \quad (3.46)$$

where we have implicitly assumed that the document length N_i is independent of the class. Here θ_{jc} is the probability of generating word j in documents of class c ; these parameters satisfy the constraint that $\sum_{j=1}^D \theta_{jc} = 1$ for each class c .

Although the multinomial classifier is easy to train and easy to use at test time, it does not work particularly well for document classification. One reason for this is that it does not take into account the **burstiness** of word usage. This refers to the phenomenon that most words never appear in any given document, but if they do appear once, they are likely to appear more than once, i.e., words occur in bursts.

The multinomial model cannot capture the burstiness phenomenon. To see why, note that Equation 3.46 has the form $\theta_{jc}^{x_{ij}}$, and since $\theta_{jc} \ll 1$ for rare words, it becomes increasingly unlikely to generate many of them. For more frequent words, the decay rate is not as fast. To see why intuitively, note that the most frequent words are function words which are not specific to the class, such as and, the, and but; the chance of the word and occurring is pretty much the same no matter how many times it has previously occurred (modulo document length), so the independence assumption is more reasonable for common words. However, since rare words are the ones that matter most for classification purposes, these are the ones we want to model the most carefully.

3.5.5.3 DCM model

Various ad hoc heuristics have been proposed to improve the performance of the multinomial document classifier (Rennie et al. 2003). We now present an alternative class conditional density that performs as well as these ad hoc methods, yet is probabilistically sound (Madsen et al. 2005).

Suppose we simply replace the multinomial class conditional density with the **Dirichlet Compound Multinomial** or **DCM** density, defined as follows:

$$\begin{aligned} p(\mathbf{x}_i|y_i = c, \boldsymbol{\alpha}) &= \int \text{Mu}(\mathbf{x}_i|N_i, \boldsymbol{\theta}_c) \text{Dir}(\boldsymbol{\theta}_c|\boldsymbol{\alpha}_c) \\ &= \frac{N_i!}{\prod_{j=1}^D x_{ij}!} \prod_{j=1}^D \frac{B(\mathbf{x}_i + \boldsymbol{\alpha}_c)}{B(\boldsymbol{\alpha}_c)} \end{aligned} \quad (3.47)$$

(This equation is derived in Equation TODO.) Surprisingly this simple change is all that is needed to capture the burstiness phenomenon. The intuitive reason for this is as follows: After seeing one occurrence of a word, say word $_j$, the posterior counts on j gets updated, making another occurrence of word $_j$ more likely. By contrast, if j is fixed, then the occurrences of each word are independent. The multinomial model corresponds to drawing a ball from an urn with K colors of ball, recording its color, and then replacing it. By contrast, the DCM model corresponds to drawing a ball, recording its color, and then replacing it with one additional copy; this is called the **Polya urn**.

Using the DCM as the class conditional density gives much better results than using the multinomial, and has performance comparable to state of the art methods, as described in (Madsen et al. 2005). The only disadvantage is that fitting the DCM model is more complex; see (Minka 2000e; Elkan 2006) for the details.