

Estimating a Dirichlet distribution

Thomas P. Minka

2000 (revised 2003, 2009, 2012)

Abstract

The Dirichlet distribution and its compound variant, the Dirichlet-multinomial, are two of the most basic models for proportional data, such as the mix of vocabulary words in a text document. Yet the maximum-likelihood estimate of these distributions is not available in closed-form. This paper describes simple and efficient iterative schemes for obtaining parameter estimates in these models. In each case, a fixed-point iteration and a Newton-Raphson (or generalized Newton-Raphson) iteration is provided.

1 The Dirichlet distribution

The Dirichlet distribution is a model of how proportions vary. Let \mathbf{p} denote a random vector whose elements sum to 1, so that p_k represents the proportion of item k . Under the Dirichlet model with parameter vector $\boldsymbol{\alpha}$, the probability density at \mathbf{p} is

$$p(\mathbf{p}) \sim \mathcal{D}(\alpha_1, \dots, \alpha_K) = \frac{\Gamma(\sum_k \alpha_k)}{\prod_k \Gamma(\alpha_k)} \prod_k p_k^{\alpha_k - 1} \quad (1)$$

$$\text{where } p_k > 0 \quad (2)$$

$$\sum_k p_k = 1 \quad (3)$$

The parameters $\boldsymbol{\alpha}$ can be estimated from a training set of proportions: $D = \{\mathbf{p}_1, \dots, \mathbf{p}_N\}$. The maximum-likelihood estimate of $\boldsymbol{\alpha}$ maximizes $p(D|\boldsymbol{\alpha}) = \prod_i p(\mathbf{p}_i|\boldsymbol{\alpha})$. The log-likelihood can be written

$$\log p(D|\boldsymbol{\alpha}) = N \log \Gamma(\sum_k \alpha_k) - N \sum_k \log \Gamma(\alpha_k) + N \sum_k (\alpha_k - 1) \log \bar{p}_k \quad (4)$$

$$\text{where } \log \bar{p}_k = \frac{1}{N} \sum_i \log p_{ik} \quad (5)$$

This objective is convex in $\boldsymbol{\alpha}$ since Dirichlet is in the exponential family. This implies that the likelihood is unimodal and the maximum can be found by a simple search. A direct convexity proof has also been given by Ronning (1989). The gradient of the log-likelihood with respect to one α_k is

$$g_k = \frac{d \log p(D|\boldsymbol{\alpha})}{d \alpha_k} = N \Psi(\sum_k \alpha_k) - N \Psi(\alpha_k) + N \log \bar{p}_k \quad (6)$$

$$\Psi(x) = \frac{d \log \Gamma(x)}{dx} \quad (7)$$

Ψ is known as the digamma function and is similar to the natural logarithm. As always with the exponential family, when the gradient is zero, the expected sufficient statistics are equal to the observed sufficient statistics. In this case, the expected sufficient statistics are

$$E[\log p_k] = \Psi(\alpha_k) - \Psi(\sum_k \alpha_k) \quad (8)$$

and the observed sufficient statistics are $\log \bar{p}_k$.

A fixed-point iteration for maximizing the likelihood can be derived as follows. Given an initial guess for α , we construct a simple lower bound on the likelihood which is tight at α . The maximum of this bound is computed in closed-form and it becomes the new guess. Such an iteration is guaranteed to converge to a stationary point of the likelihood—in fact it is the same principle behind the EM algorithm (Minka, 1998). For the Dirichlet, the maximum is the only stationary point.

As shown in appendix A, a bound on $\Gamma(\sum_k \alpha_k)$ leads to the following fixed-point iteration:

$$\Psi(\alpha_k^{new}) = \Psi\left(\sum_k \alpha_k^{old}\right) + \log \bar{p}_k \quad (9)$$

This algorithm requires inverting the Ψ function—a procedure which is described in appendix C.

Another approach to finding a stationary point is Newton iteration. The second-derivatives, i.e. Hessian matrix, of the log-likelihood are given by

$$\frac{d \log p(D|\alpha)}{d\alpha_k^2} = N\Psi'\left(\sum_k \alpha_k\right) - N\Psi'(\alpha_k) \quad (10)$$

$$\frac{d \log p(D|\alpha)}{d\alpha_k d\alpha_j} = N\Psi'\left(\sum_k \alpha_k\right) \quad (k \neq j) \quad (11)$$

Ψ' is known as the trigamma function. The Hessian can be written in matrix form as

$$\mathbf{H} = \mathbf{Q} + \mathbf{1}\mathbf{1}^T z \quad (12)$$

$$q_{jk} = -N\Psi'(\alpha_k)\delta(j-k) \quad (13)$$

$$z = N\Psi'\left(\sum_k \alpha_k\right) \quad (14)$$

One Newton step is therefore

$$\alpha^{new} = \alpha^{old} - \mathbf{H}^{-1} \mathbf{g} \quad (15)$$

$$\mathbf{H}^{-1} = \mathbf{Q}^{-1} - \frac{\mathbf{Q}^{-1}\mathbf{1}\mathbf{1}^T\mathbf{Q}^{-1}}{1/z + \mathbf{1}^T\mathbf{Q}^{-1}\mathbf{1}} \quad (16)$$

$$(\mathbf{H}^{-1}\mathbf{g})_k = \frac{g_k - b}{q_{kk}} \quad (17)$$

$$\text{where } b = \frac{\mathbf{1}^T\mathbf{Q}^{-1}\mathbf{g}}{1/z + \mathbf{1}^T\mathbf{Q}^{-1}\mathbf{1}} = \frac{\sum_j g_j/q_{jj}}{1/z + \sum_j 1/q_{jj}} \quad (18)$$

Unlike some Newton algorithms, this one does not require storing or inverting the Hessian matrix explicitly. The same Newton algorithm was given by Ronning (1989) and by Naryanan (1991). Naryanan also derives a stopping rule for the iteration.

An approximate MLE, useful for initialization, is given by finding the density which matches the moments of the data. The first two moments of the density are

$$E[p_k] = \frac{\alpha_k}{\sum_k \alpha_k} \quad (19)$$

$$E[p_k^2] = E[p_k] \frac{1 + \alpha_k}{1 + \sum_k \alpha_k} \quad (20)$$

$$\sum_k \alpha_k = \frac{E[p_1] - E[p_1^2]}{E[p_1^2] - E[p_1]^2} \quad (21)$$

Multiplying (21) and (19) gives a formula for α_k in terms of moments. Equation (21) uses p_1 , but any other p_k could also be used to estimate $\sum_k \alpha_k$. Ronning (1989) suggests instead using all of the p_k 's via

$$\text{var}(p_k) = \frac{E[p_k](1 - E[p_k])}{1 + \sum_k \alpha_k} \quad (22)$$

$$\log \sum_k \alpha_k = \frac{1}{K-1} \sum_{k=1}^{K-1} \log \left(\frac{E[p_k](1 - E[p_k])}{\text{var}(p_k)} - 1 \right) \quad (23)$$

Another approximate MLE, specifically for the case $K = 2$, is given by Johnson & Kotz (1970):

$$\alpha_1 = \frac{1}{2} \frac{1 - \bar{p}_2}{1 - \bar{p}_1 - \bar{p}_2} \quad (24)$$

$$\alpha_2 = \frac{1}{2} \frac{1 - \bar{p}_1}{1 - \bar{p}_1 - \bar{p}_2} \quad (25)$$

2 Estimating Dirichlet mean and precision separately

The α parameters of the Dirichlet can be understood by considering the following alternative representation:

$$s = \sum_k \alpha_k \quad (26)$$

$$\mathbf{m} = E[p_k] = \alpha/s \quad (27)$$

Here \mathbf{m} is the mean of the distribution for \mathbf{p} and s can be understood as the *precision*. When s is large, \mathbf{p} is likely to be near \mathbf{m} , and when s is small, \mathbf{p} is distributed more diffusely. Interpretation of s and \mathbf{m} suggests situations, such as hierarchical modeling, in which we may want to fix one parameter and only optimize the other. Additionally, s and \mathbf{m} are roughly decoupled in the maximum-likelihood objective, which means we can get simplifications and speedups by optimizing them alternately. Thus, in this section we will reparameterize the distribution with (s, \mathbf{m}) where

$$\alpha_k = sm_k \quad (28)$$

$$\sum_k m_k = 1 \quad (29)$$

2.1 Estimating Dirichlet precision

The likelihood for s alone is

$$p(D|s) \propto \left(\frac{\Gamma(s) \exp(s \sum_k m_k \log \bar{p}_k)}{\prod_k \Gamma(sm_k)} \right)^N \quad (30)$$

whose derivatives are

$$\frac{d \log p(D|s)}{ds} = N\Psi(s) - N \sum_k m_k \Psi(sm_k) + N \sum_k m_k \log \bar{p}_k \quad (31)$$

$$\frac{d^2 \log p(D|s)}{ds^2} = N\Psi'(s) - N \sum_k m_k^2 \Psi'(sm_k) \quad (32)$$

A convergent fixed-point iteration for s is

$$(K-1)/s^{new} = (K-1)/s - \Psi(s) + \sum_k m_k \Psi(sm_k) - \sum_k m_k \log \bar{p}_k \quad (33)$$

Proof Use the bound (see Appendix E)

$$\frac{\Gamma(s)}{\prod_k \Gamma(sm_k)} \geq \exp(sb + (K-1)\log(s) + c) \quad (34)$$

$$b = \Psi(\hat{s}) - \sum_k m_k \Psi(\hat{s}m_k) - (K-1)/\hat{s} \quad (35)$$

to get

$$\log p(D|s) \geq s \sum_k m_k \log \bar{p}_k + sb + (K-1)\log(s) + (\text{const.}) \quad (36)$$

from which (33) follows.

This iteration is only first-order convergent because the bound only matches the first derivative of the likelihood. We can derive a second-order method using the technique of generalized Newton iteration Minka (2000). The idea is to approximate the likelihood by a simpler function, by matching the first two derivatives at the current guess:

$$\frac{\Gamma(s)}{\prod_k \Gamma(sm_k)} \approx \exp(sb + a\log(s) + c) \quad (37)$$

$$a = -\hat{s}^2(\Psi'(\hat{s}) - \sum_k m_k^2 \Psi'(\hat{s}m_k)) \quad (38)$$

$$b = \Psi(\hat{s}) - \sum_k m_k \Psi(\hat{s}m_k) - a/\hat{s} \quad (39)$$

Maximizing the approximation leads to the update

$$\frac{1}{s^{new}} = \frac{1}{s} + \frac{1}{s^2} \left(\frac{d^2 \log p(D|s)}{ds^2} \right)^{-1} \left(\frac{d \log p(D|s)}{ds} \right) \quad (40)$$

This update resembles Newton-Raphson, but converges faster.

For initialization of s , it is useful to derive a closed-form approximate MLE. Stirling's approximation to Γ gives

$$\frac{\Gamma(s) \exp(s \sum_k m_k \log \bar{p}_k)}{\prod_k \Gamma(sm_k)} \approx \left(\frac{s}{2\pi} \right)^{(K-1)/2} \prod_k m_k^{1/2} \exp(s \sum_k m_k \log \frac{\bar{p}_k}{m_k}) \quad (41)$$

$$\hat{s} \approx \frac{(K-1)/2}{-\sum_k m_k \log \frac{\bar{p}_k}{m_k}} \quad (42)$$

2.2 Estimating Dirichlet mean

Now suppose we fix the precision s and want to estimate the mean \mathbf{m} . The likelihood for \mathbf{m} alone is

$$p(D|\mathbf{m}) \propto \left(\prod_k \frac{\exp(sm_k \log \bar{p}_k)}{\Gamma(sm_k)} \right)^N \quad (43)$$

Reparameterize with the unconstrained vector \mathbf{z} , to get the gradient:

$$m_k = \frac{z_k}{\sum_j z_j} \quad (44)$$

$$\frac{d \log p(D|\mathbf{m})}{dz_k} = \frac{Ns}{\sum_j z_j} \left(\log \bar{p}_k - \Psi(sm_k) - \sum_j m_j (\log \bar{p}_j - \Psi(sm_j)) \right) \quad (45)$$

The MLE can be computed by the fixed-point iteration

$$\Psi(\alpha_k) = \log \bar{p}_k - \sum_j m_j^{old} (\log \bar{p}_j - \Psi(sm_j^{old})) \quad (46)$$

$$m_k^{new} = \frac{\alpha_k}{\sum_j \alpha_j} \quad (47)$$

This update converges very quickly.

3 The Dirichlet-multinomial/Polya distribution

The Dirichlet-multinomial distribution is a compound distribution where \mathbf{p} is drawn from a Dirichlet and then a sample of discrete outcomes \mathbf{x} is drawn from a multinomial with probability vector \mathbf{p} . This compounding is essentially a Polya urn scheme, so the Dirichlet-multinomial is also called the *Polya distribution*. Let n_k be the number of times the outcome was k , i.e.

$$n_k = \sum_j \delta(x_j - k) \quad (48)$$

Then the resulting distribution over \mathbf{x} , a vector of outcomes, is

$$p(\mathbf{x}|\alpha) = \int_{\mathbf{p}} p(\mathbf{x}|\mathbf{p})p(\mathbf{p}|\alpha)d\mathbf{p} \quad (49)$$

$$= \frac{\Gamma(\sum_k \alpha_k)}{\Gamma(\sum_k n_k + \alpha_k)} \prod_k \frac{\Gamma(n_k + \alpha_k)}{\Gamma(\alpha_k)} \quad (50)$$

This distribution is also parameterized by α , which can be estimated from a training set of count vectors: $D = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$. The likelihood is

$$n_i = \sum_k n_{ik} \quad (51)$$

$$p(D|\alpha) = \prod_i p(\mathbf{x}_i|\alpha) \quad (52)$$

$$= \prod_i \left(\frac{\Gamma(\sum_k \alpha_k)}{\Gamma(n_i + \sum_k \alpha_k)} \prod_k \frac{\Gamma(n_{ik} + \alpha_k)}{\Gamma(\alpha_k)} \right) \quad (53)$$

The gradient of the log-likelihood is

$$g_k = \frac{d \log p(D|\alpha)}{d\alpha_k} = \sum_i \Psi(\sum_k \alpha_k) - \Psi(n_i + \sum_k \alpha_k) + \Psi(n_{ik} + \alpha_k) - \Psi(\alpha_k) \quad (54)$$

The maximum can be computed via the fixed-point iteration

$$\alpha_k^{new} = \alpha_k \frac{\sum_i \Psi(n_{ik} + \alpha_k) - \Psi(\alpha_k)}{\sum_i \Psi(n_i + \sum_k \alpha_k) - \Psi(\sum_k \alpha_k)} \quad (55)$$

(see appendix B).

Alternatively, there is a simplified Newton iteration as in the Dirichlet case. The Hessian of the log-likelihood is

$$\frac{d \log p(D|\alpha)}{d\alpha_k^2} = \sum_i \Psi'(\sum_k \alpha_k) - \Psi'(n_i + \sum_k \alpha_k) + \Psi'(n_{ik} + \alpha_k) - \Psi'(\alpha_k) \quad (56)$$

$$\frac{d \log p(D|\alpha)}{d\alpha_k d\alpha_j} = \sum_i \Psi'(\sum_k \alpha_k) - \Psi'(n_i + \sum_k \alpha_k) \quad (k \neq j) \quad (57)$$

The Hessian can be written in matrix form as

$$\mathbf{H} = \mathbf{Q} + \mathbf{1}\mathbf{1}^T z \quad (58)$$

$$q_{jk} = \delta(j - k) \sum_i \Psi'(n_{ik} + \alpha_k) - \Psi'(\alpha_k) \quad (59)$$

$$z = \sum_i \Psi'(\sum_k \alpha_k) - \Psi'(n_i + \sum_k \alpha_k) \quad (60)$$

from which a Newton step can be computed as before. The search can be initialized with the moment matching estimate where p_{ik} is approximated by n_{ik}/n_i .

Another approach is to reduce this problem to the previous one via EM; see appendix D.

A different method is to maximize the leave-one-out (LOO) likelihood instead of the true likelihood. The LOO likelihood is the product of the probability of each sample given the remaining data and the parameters. The LOO log-likelihood is

$$f(\alpha) = \sum_{ik} n_{ik} \log \left(\frac{n_{ik} - 1 + \alpha_k}{n_i - 1 + \sum_k \alpha_k} \right) = \sum_{ik} n_{ik} \log(n_{ik} - 1 + \alpha_k) - \sum_i n_i \log(n_i - 1 + \sum_k \alpha_k) \quad (61)$$

Note that it doesn't involve any special functions. The derivatives are

$$\frac{df(\alpha)}{d\alpha_k} = \sum_i \frac{n_{ik}}{n_{ik} - 1 + \alpha_k} - \frac{n_i}{n_i - 1 + \sum_k \alpha_k} \quad (62)$$

$$\frac{df(\alpha)}{d\alpha_k^2} = \sum_i -\frac{n_{ik}}{(n_{ik} - 1 + \alpha_k)^2} + \frac{n_i}{(n_i - 1 + \sum_k \alpha_k)^2} \quad (63)$$

$$\frac{df(\alpha)}{d\alpha_k d\alpha_j} = \sum_i \frac{n_i}{(n_i - 1 + \sum_k \alpha_k)^2} \quad (64)$$

A convergent fixed-point iteration is

$$\alpha_k^{new} = \alpha_k \frac{\sum_i \frac{n_{ik}}{n_{ik} - 1 + \alpha_k}}{\sum_i \frac{n_i}{n_i - 1 + \sum_k \alpha_k}} \quad (65)$$

Proof Use the bounds

$$\log(n + x) \geq q \log x + (1 - q) \log n - q \log q - (1 - q) \log(1 - q) \quad (66)$$

$$q = \frac{\hat{x}}{n + \hat{x}} \quad (67)$$

$$\log(x) \leq ax - 1 + \log \hat{x} \quad (68)$$

$$a = 1/\hat{x} \quad (69)$$

to get

$$f(\alpha) \geq \sum_i n_{ik} q_{ik} \log \alpha_k - n_i a_i \sum_k \alpha_k + (\text{const.}) \quad (70)$$

leading to (65).

The LOO likelihood can be interpreted as the approximation

$$\frac{\Gamma(x+n)}{\Gamma(x)} \approx (x+n-1)^n \quad (71)$$

4 Estimating Polya mean and precision separately

The α parameters of the Polya distribution can be decomposed into mean \mathbf{m} and precision s , just as in the Dirichlet case, and optimization can be done separately on each part. The decomposition also leads to an interesting interpretation of the Polya, discussed in the next subsection.

4.1 A novel interpretation of the Polya distribution

The Polya distribution can be interpreted as a multinomial distribution over a modified set of counts, with a special normalizer. To see why, consider the log-probability of the outcomes \mathbf{x} under the Polya versus multinomial:

$$\log p(\mathbf{x}|\alpha) = \log \Gamma(s) - \log \Gamma(s+n) + \sum_k \log \Gamma(n_k + sm_k) - \log \Gamma(sm_k) \quad (72)$$

$$\log p(\mathbf{x}|\mathbf{p}) = \sum_k n_k \log p_k \quad (73)$$

The multinomial is an exponential family, and the Polya is not. But we can find an approximate exponential family representation of the Polya, by considering derivatives. In the multinomial case, the counts can be recovered from the expression

$$n_k = p_k \frac{d \log p(\mathbf{x}|\mathbf{p})}{dp_k} \quad (74)$$

In the Polya case, the analogous expression is

$$\tilde{n}_k = m_k \frac{d \log p(\mathbf{x}|\alpha)}{dm_k} \quad (75)$$

$$= \alpha_k (\Psi(n_k + \alpha_k) - \Psi(\alpha_k)) \equiv \nu(n_k, \alpha_k) \quad (76)$$

The log-probability of \mathbf{x} under the Polya can thus be approximated by

$$\log p(\mathbf{x}|\alpha) = \sum_k \tilde{n}_k \log m_k \quad (77)$$

When $s \rightarrow \infty$, the Dirichlet-multinomial becomes an ordinary multinomial with $\mathbf{p} = \mathbf{m}$, and therefore the ‘effective’ counts are the same as the ordinary counts:

$$\nu(n_k, \infty) = n_k \quad (78)$$

At the other extreme, as $s \rightarrow 0$, the Dirichlet-multinomial favors extreme proportions, and the effective counts are a binarized version of the original counts:

$$\nu(n_k, 0) = \begin{cases} 0 & \text{if } n_k = 0, \\ 1 & \text{if } n_k > 0. \end{cases} \quad (79)$$

For intermediate values of α , the mapping ν behaves like a logarithm, reducing the influence of large counts on the likelihood (see figure 1). Thus the Polya can be understood as a multinomial with ‘damped’ counts.




Figure 1: The ‘effective’ counts for the Polya distribution, as a function of the original count n_k and the parameter α_k .

This representation of the Polya also arises in the estimation of \mathbf{m} when s is fixed (section 5).

4.2 Estimating Polya precision

The likelihood for s alone is

$$p(D|s) \propto \prod_i \left(\frac{\Gamma(s)}{\Gamma(n_i + s)} \prod_k \frac{\Gamma(n_{ik} + sm_k)}{\Gamma(sm_k)} \right) \quad (80)$$

The derivatives are

$$\frac{d \log p(D|s)}{ds} = \sum_i \Psi(s) - \Psi(n_i + s) + \sum_k m_k \Psi(n_{ik} + sm_k) - m_k \Psi(sm_k) \quad (81)$$

$$\frac{d^2 \log p(D|s)}{ds^2} = \sum_i \Psi'(s) - \Psi'(n_i + s) + \sum_k m_k^2 \Psi'(n_{ik} + sm_k) - m_k^2 \Psi'(sm_k) \quad (82)$$

A convergent fixed-point iteration is

$$s^{new} = s \frac{\sum_{ik} m_k \Psi(n_{ik} + sm_k) - m_k \Psi(sm_k)}{\sum_i \Psi(n_i + s) - \Psi(s)} \quad (83)$$

(the proof is similar to (55)). However, it is very slow. We can get a fast second-order method as follows. When s is small, i.e. the gradient is positive, use the approximation

$$\log p(D|s) \approx a \log(s) + cs + k \quad (84)$$

$$a = -s_0^2 f''(s_0) \quad (85)$$

$$c = f'(s_0) - a/s_0 \quad (86)$$

to get the update

$$s^{new} = -a/c = s / \left(1 + \frac{f'(s)}{s f''(s)}\right) \quad (87)$$

except when $c \geq 0$, in which case the solution is $s = \infty$. When s is large, i.e. the gradient is negative, use the approximation

$$\log p(D|s) \approx \frac{a}{2s^2} + \frac{c}{s} \quad (88)$$

$$a = s_0^3 (s_0 f''(s_0) + 2f'(s_0)) \quad (89)$$

$$c = -(s_0^2 f'(s_0) + a/s_0) \quad (90)$$

to get the update

$$s^{new} = -a/c = s - \frac{f'(s)}{f''(s) + 3f'(s)/s} \quad (91)$$

For large s , the value of a tends to be numerically unstable. If $s f''(s) + 2f'(s)$ is within machine epsilon, then it is better to substitute the limiting value:

$$a \rightarrow \sum_i \frac{n_i(n_i - 1)(2n_i - 1)}{6} - \sum_{ik} \frac{n_{ik}(n_{ik} - 1)(2n_{ik} - 1)}{6m_k^2} \quad (92)$$

A even faster update for large s is possible by using a richer approximation:

$$\log p(D|s) \approx c \log\left(\frac{s}{s+b}\right) + \frac{e}{s+b} \quad (93)$$

$$c = \sum_{ik} \delta(n_{ik} > 0) - \sum_i \delta(n_i > 0) \quad (94)$$

$$e = -\frac{s_0 + b}{s_0} (s_0(s_0 + b)f'(s_0) - cb) \quad (95)$$

$$b = \text{RootOf}(a_2 b^2 + a_1 b + a_0) \quad (96)$$

$$a_2 = s_0^3 (s_0 f''(s_0) + 2f'(s_0)) \quad (97)$$

$$a_1 = 2s_0^2 (s_0 f''(s_0) + f'(s_0)) \quad (98)$$

$$a_0 = s_0^2 f''(s_0) + c \quad (99)$$

The approximation comes from setting $c \log s$ to match $\log p(D|s)$ as $s \rightarrow 0$ and then choosing (b, e) to match the first two derivatives of f at the current s . The resulting update is

$$s^{new} = \frac{cb^2}{e - cb} = \left(\frac{1}{s} - \frac{f'(s)(s+b)^2}{cb^2}\right)^{-1} \quad (100)$$

Note that a_2 is equivalent to a above and should be corrected for stability via the same method.

The case of large dimension

An interesting special case arises when K is very large. The precision can be estimated simply by counting the number of *singleton* elements in each \mathbf{x} . Because the precision acts like a smoothing

parameter on the estimate of \mathbf{m} , this result is reminiscent of smoothing methods in document modeling which are based on counting singletons.

If m_k is roughly uniform and $K \gg 1$, then $\alpha_k \ll 1$ and we can use the approximations

$$\Gamma(n_k + \alpha_k) \approx \Gamma(n_k) \quad (101)$$

$$\Gamma(\alpha_k) \approx 1/\alpha_k \quad (102)$$

$$p(\mathbf{x}|s) \approx \frac{\Gamma(s)}{\Gamma(n+s)} \prod_{n_k > 0} sm_k \Gamma(n_k) \quad (103)$$

$$\propto \frac{\Gamma(s)s^{\hat{K}}}{\Gamma(n+s)} \quad (104)$$

where \hat{K} is the number of unique observations in \mathbf{x} . The approximation does not hold if s is large, which can happen when \mathbf{m} is a good match to the data. But if the dimensionality is large enough, the data will be too sparse for this to happen. The derivatives become

$$\frac{d \log p(D|s)}{ds} \approx \sum_i \Psi(s) - \Psi(n_i + s) + \hat{K}_i/s \quad (105)$$

$$\frac{d^2 \log p(D|s)}{ds^2} \approx \sum_i \Psi'(s) - \Psi'(n_i + s) - \hat{K}_i/s^2 \quad (106)$$

Newton iteration can be used as long as the maximum for s is not on the boundary of $(0, \infty)$. These cases occur when $\hat{K} = 1$ and $\hat{K} = n$.

When the gradient is zero, we have

$$\hat{K} = s(\Psi(n+s) - \Psi(s)) = E[K|s, n] \quad (107)$$

A convergent fixed-point iteration is

$$s^{new} = \frac{\sum_i \hat{K}_i}{\sum_i \Psi(n_i + s^{old}) - \Psi(s^{old})} \quad (108)$$

Proof Use the bound

$$\frac{\Gamma(s)}{\Gamma(n+s)} \geq \frac{\Gamma(\hat{s}) \exp((\hat{s} - s)b)}{\Gamma(n + \hat{s})} \quad (109)$$

$$b = \Psi(n + \hat{s}) - \Psi(\hat{s}) \quad (110)$$

to get

$$p(D|s) \geq -s \sum_i b_i + \sum_i \hat{K}_i \log s + (\text{const.}) \quad (111)$$

leading to (108).

Applying the large K approximation to the LOO likelihood gives

$$t = \sum_{ik} \delta(n_{ik} - 1) \quad (\text{number of singletons}) \quad (112)$$

$$f(s) = t \log s - \sum_i n_i \log(n_i - 1 + s) \quad (113)$$

$$\frac{df(s)}{ds} = \frac{t}{s} - \sum_i \frac{n_i}{n_i - 1 + s} \quad (114)$$

For $N = 1$:

$$s = \frac{t(n-1)}{n-t} \quad (115)$$

$$\frac{s}{s+n} = \frac{t(n-1)}{n^2-t} \approx \frac{t}{n} \quad (116)$$

which is the result we wanted.

5 Estimating Polya mean

The likelihood for \mathbf{m} only is

$$p(D|\mathbf{m}) \propto \prod_{ik} \frac{\Gamma(n_{ik} + sm_k)}{\Gamma(sm_k)} \quad (117)$$

The maximum can be computed by the fixed-point iteration

$$m_k^{new} \propto \sum_i \nu(n_{ik}, sm_k) \quad (118)$$

where ν is defined in (76). This update can be understood intuitively as the maximum-likelihood estimate of a multinomial distribution from effective counts $\tilde{n}_{ik} = \nu(n_{ik}, sm_k)$. The proof of this iteration is similar to (55).

For a Newton-Raphson iteration, reparameterize to get

$$m_K = 1 - \sum_{k=1}^{K-1} m_k \quad (119)$$

$$g_k = \frac{d \log p(D|\mathbf{m})}{dm_k} = s \sum_i \Psi(n_{ik} + sm_k) - \Psi(sm_k) - \Psi(n_{iK} + sm_K) + \Psi(sm_K) \quad (120)$$

$$\frac{d^2 \log p(D|\mathbf{m})}{dm_k^2} = s^2 \sum_i \Psi'(n_{ik} + sm_k) - \Psi'(sm_k) + \Psi'(n_{iK} + sm_K) - \Psi'(sm_K) \quad (121)$$

$$\frac{d^2 \log p(D|\mathbf{m})}{dm_k m_j} = s^2 \sum_i \Psi'(n_{iK} + sm_K) - \Psi'(sm_K) \quad (122)$$

The search should be initialized at $m_k \propto \sum_i n_{ik}$, since for large s this is the exact optimum.

References

- Johnson, N. L., & Kotz, S. (1970). *Distributions in statistics: Continuous univariate distributions*. New York: Houghton Mifflin.
- Minka, T. P. (1998). Expectation-Maximization as lower bound maximization. <http://research.microsoft.com/~minka/papers/em.html>.
- Minka, T. P. (2000). Beyond newton's method. <http://research.microsoft.com/~minka/papers/newton.html>.

Naryanan, A. (1991). Algorithm as 266: Maximum likelihood estimation of the parameters of the dirichlet distribution. *Applied Statistics*, 40, 365–374.
<http://www.psc.edu/~burkardt/dirichlet.html>.

Ronning, G. (1989). Maximum-likelihood estimation of dirichlet distributions. *Journal of Statistical Computation and Simulation*, 32, 215–221.

A Proof of (9)

Use the bound

$$\Gamma(x) \geq \Gamma(\hat{x}) \exp((x - \hat{x})\Psi(\hat{x})) \quad (123)$$

to get

$$\frac{1}{N} \log p(D|\alpha) \geq \left(\sum_k \alpha_k \right) \Psi \left(\sum_k \alpha_k^{old} \right) - \sum_k \log \Gamma(\alpha_k) + \sum_k (\alpha_k - 1) \log \bar{p}_k + (\text{const.}) \quad (124)$$

leading to (9).

B Proof of (55)

Use the bound

$$\frac{\Gamma(x)}{\Gamma(n+x)} \geq \frac{\Gamma(\hat{x}) \exp((\hat{x} - x)b)}{\Gamma(n+\hat{x})} \quad (125)$$

$$b = \Psi(n+\hat{x}) - \Psi(\hat{x}) \quad (126)$$

and the bound

$$\frac{\Gamma(n+x)}{\Gamma(x)} \geq cx^a \quad \text{if } n \geq 1, x \geq 0 \quad (127)$$

$$a = (\Psi(n+\hat{x}) - \Psi(\hat{x}))\hat{x} \quad (128)$$

$$c = \frac{\Gamma(n+\hat{x})}{\Gamma(\hat{x})} \hat{x}^{-a} \quad (129)$$

to get

$$\log p(D|\alpha) \geq - \left(\sum_k \alpha_k - 1 \right) \sum_i b_i + \sum_k a_{ik} \log \alpha_k + (\text{const.}) \quad (130)$$

leading to (55).

Proof of (125): Use Jensen's inequality on the integral definition of the Beta function.

$$\frac{\Gamma(n)\Gamma(x)}{\Gamma(n+x)} = \int_0^1 t^{x-1}(1-t)^{n-1} dt \quad (131)$$

$$\geq \exp \left(\int_0^1 q(t) \log \frac{t^{x-1}(1-t)^{n-1}}{q(t)} dt \right) \quad (132)$$

$$q(t) = \frac{\Gamma(n+\hat{x})}{\Gamma(n)\Gamma(\hat{x})} t^{\hat{x}-1} (1-t)^{n-1} \quad (133)$$

Proof of (127): The bound corresponds to a linear expansion of $\log \Gamma(n+x) - \log \Gamma(x)$ in $\log(x)$. Thus we only need to show that $\log \Gamma(n+x) - \log \Gamma(x)$ is convex in $\log(x)$. By taking derivatives, this amounts to proving that:

$$h(x, n) = x(\Psi(n+x) - \Psi(x)) + x^2(\Psi'(n+x) - \Psi'(x)) \geq 0 \quad \text{if } n \geq 1 \quad (134)$$

The digamma function has the following integral representation:

$$\Psi(x) = \int_0^\infty \left(\frac{e^t}{t} - \frac{e^{-xt}}{1-e^{-t}} \right) dt \quad (135)$$

Substituting this into (134) gives:

$$h(x, n) = \int_0^\infty x(1-xt)e^{-xt} \frac{1-e^{-nt}}{1-e^{-t}} dt \quad (136)$$

Divide the integrand into two parts and apply integration by parts:

$$f(t) = x(1-xt)e^{-xt} \quad (137)$$

$$g(t) = \frac{1-e^{-nt}}{1-e^{-t}} \quad (138)$$

$$F(t) = xte^{-xt} \quad (\text{anti-derivative of } f) \quad (139)$$

$$h(x, n) = \int_0^\infty f(t)g(t)dt \quad (140)$$

$$= 0 - \int_0^\infty F(t)g'(t)dt \quad (141)$$

Since $F(t) \geq 0$ for all $t \geq 0$, we only have to show $g'(t) \leq 0$ for all $t \geq 0$. This amounts to showing that:

$$(1-e^{-t})(ne^{-nt}) \leq (1-e^{-nt})e^{-t} \quad (142)$$

$$e^t - 1 \leq (e^{nt} - 1)/n \quad (143)$$

$$t + t^2/2 + t^3/3! + \dots \leq t + nt^2/2 + n^2t^3/3! + \dots \quad (144)$$

The last line follows from $n \geq 1$ which completes the proof. Note that if $n \leq 1$, the last inequality flips to show that the function is concave, i.e. the right-hand side of (127) becomes an upper bound.

C Inverting the Ψ function

This section describes how to compute a high-accuracy solution to

$$\Psi(x) = y \quad (145)$$

for x given y . Given a starting guess for x , Newton's method can be used to find the root of $\Psi(x) - y = 0$. The Newton update is

$$x^{new} = x^{old} - \frac{\Psi(x) - y}{\Psi'(x)} \quad (146)$$

To start the iteration, use the following asymptotic formulas for $\Psi(x)$:

$$\Psi(x) \approx \begin{cases} \log(x-1/2) & \text{if } x \geq 0.6 \\ -\frac{1}{x} - \gamma & \text{if } x < 0.6 \end{cases} \quad (147)$$

$$\gamma = -\Psi(1) \quad (148)$$

to get

$$\Psi^{-1}(y) \approx \begin{cases} \exp(y) + 1/2 & \text{if } y \geq -2.22 \\ -\frac{1}{y+\gamma} & \text{if } y < -2.22 \end{cases} \quad (149)$$

With this initialization, five Newton iterations are sufficient to reach fourteen digits of precision.

D EM for estimation from counts

Any algorithm for estimation from probability vectors can be turned into an algorithm for estimation from counts, by treating the \mathbf{p}_i as hidden variables in EM. The E-step computes a posterior distribution over \mathbf{p}_i :

$$q(\mathbf{p}_i) \sim \mathcal{D}(n_{ik} + \alpha_k) \quad (150)$$

and the M-step maximizes

$$E\left[\sum_i \log p(\mathbf{p}_i|\alpha)\right] = N \log \Gamma\left(\sum_k \alpha_k\right) - N \sum_k \log \Gamma(\alpha_k) + N \sum_k (\alpha_k - 1) \log \bar{p}_k \quad (151)$$

$$\text{where } \log \bar{p}_k = \frac{1}{N} \sum_i E[\log p_{ik}] \quad (152)$$

$$= \frac{1}{N} \sum_i \Psi(n_{ik} + \alpha_k^{old}) - \Psi(n_i + \sum_k \alpha_k^{old}) \quad (153)$$

This is the same optimization problem as in section 1, with a new definition for \bar{p} . It is not necessary or desirable to reach the exact maximum in the M-step; a single Newton step will do. The Newton step will end up using the old Hessian (10) but the new gradient (54). Compared to the exact Newton algorithm, this uses half as much computation per iteration, but usually requires more than twice the iterations.

E Proof of (34)

We want to prove the bound

$$\frac{\Gamma(s)}{\prod_k \Gamma(sm_k)} \geq \exp(sb + (K-1)\log(s) + c) \quad (154)$$

$$b = \Psi(\hat{s}) - \sum_k m_k \Psi(\hat{s}m_k) - (K-1)/\hat{s} \quad (155)$$

Define the function $f(s)$ via

$$f(s) = \log \Gamma(s+1) - \sum_k \log \Gamma(sm_k + 1) \quad (156)$$

$$= \log \Gamma(s) - \sum_k \log \Gamma(sm_k) - (K-1)s + \text{const.} \quad (157)$$

The bound is equivalent to saying that $f(s)$ can be lower bounded by a linear function in s at any point \hat{s} , i.e. $f(s)$ is convex in s . To show this, take the second derivative of $f(s)$:

$$\frac{df}{ds} = \Psi(s+1) - \sum_k \Psi(sm_k + 1)m_k \quad (158)$$

$$\frac{d^2f}{ds^2} = \Psi'(s+1) - \sum_k \Psi'(sm_k + 1)m_k^2 \quad (159)$$

We need to show that (159) is always positive. Since the function $g(x) = \Psi'(x+1)x$ is increasing, we know that:

$$g(s) \geq g(sm_k) \tag{160}$$

$$\sum_k m_k g(s) \geq \sum_k m_k g(sm_k) \tag{161}$$

$$g(s) \geq \sum_k m_k g(sm_k) \tag{162}$$

$$\Psi'(s+1) \geq \sum_k \Psi'(sm_k+1)m_k^2 \tag{163}$$

thus $f(s)$ is convex and the bound follows.