



Power transformations of relative count data as a shrinkage problem

Ionas Erb¹

Received: 25 May 2022 / Revised: 13 February 2023 / Accepted: 14 March 2023 /

Published online: 13 April 2023

© The Author(s), under exclusive licence to Springer Nature Singapore Pte Ltd. 2023

Abstract

Here we show an application of our recently proposed information-geometric approach to compositional data analysis (CoDA). This application regards relative count data, which are, e.g., obtained from sequencing experiments. First we review in some detail a variety of necessary concepts ranging from basic count distributions and their information-geometric description over the link between Bayesian statistics and shrinkage to the use of power transformations in CoDA. We then show that *powering*, i.e., the equivalent to scalar multiplication on the simplex, can be understood as a shrinkage problem on the tangent space of the simplex. In information-geometric terms, traditional shrinkage corresponds to an optimization along a mixture (or m -) geodesic, while powering (or, as we call it, *exponential shrinkage*) can be optimized along an exponential (or e -) geodesic. While the m -geodesic corresponds to the posterior mean of the multinomial counts using a conjugate prior, the e -geodesic corresponds to an alternative parametrization of the posterior where prior and data contributions are weighted by geometric rather than arithmetic means. To optimize the exponential shrinkage parameter, we use mean-squared error as a cost function on the tangent space. This is just the expected squared Aitchison distance from the true parameter. We derive an analytic solution for its minimum based on the delta method and test it via simulations. We also discuss exponential shrinkage as an alternative to zero imputation for dimension reduction and data normalization.

Keywords Compositional data · Information geometry · Dual geodesics · Multinomial distribution · Box–Cox transformation · Zero handling · James–Stein shrinkage · Empirical Bayes

Communicated by Nihat Ay.

✉ Ionas Erb
ionas.erb@crg.eu

¹ The Barcelona Institute of Science and Technology, Centre for Genomic Regulation (CRG), C/ Dr Aiguader, 88, 08003 Barcelona, Spain

1 Introduction

Counting discrete events seems one of the simplest ways of collecting data, but compositional bias when directly comparing such counts in varying contexts can lead intuition astray. Often, the lack of a common scale in samples taken from different environments or experimental conditions makes direct comparisons between counts meaningless. We need to gauge by internal references before we can make external comparisons. Compositional data analysis (CoDA, e.g. [1]) uses scale-free methods on data occurring in form of percentages, and its log-ratio methodology [2] has been applied to relative counts as well. While the sample spaces [3] of both data types are certainly not the same, the underlying problematic is identical: direct comparisons across samples can have paradoxical effects due to the lack of a common scale [4]. We have recently proposed to make use of information geometry [5] to analyse compositional data [6]. The information-geometric approach is even more natural for relative count data, and simple count distributions like the categorical or multinomial have served as examples to illustrate basic concepts in information geometry. Here we aim to demonstrate the usefulness of information-geometric concepts for the analysis of count data that are compositional in a well-defined sense.

Let us quickly sketch the main idea of this contribution. Consider a vector of counts $(n_i)_{i=1}^D$ that were produced by some process with unknown independent count probabilities q_i . It is well known that the empirical estimator for such multinomial probabilities

$$\hat{q}_i = \frac{n_i}{\sum_{k=1}^D n_k} \quad (1)$$


(although it is the one that maximizes the likelihood of the data) can be much improved upon when the denominator is not large compared with D . In this case, a better alternative is the convex combination

$$\hat{q}_i^{\text{sh}} = \lambda \frac{1}{D} + (1 - \lambda) \hat{q}_i \quad (2)$$

of the estimator with the equidistribution, for an optimized value of the parameter $0 \leq \lambda \leq 1$. This is an example of what is known as *shrinkage* of \hat{q}_i toward the target $1/D$. The reason why this works can be understood from a Bayesian perspective. The shrinkage estimator (2), instead of maximizing the likelihood of the data, maximizes the posterior probability of a suitable parameter of the multinomial (assuming a simple conjugate prior). Optimization of λ corresponds to adjusting the weight that the prior will have compared with the weight that will be assumed for the data. But why is \hat{q}_i^{sh} a good approximation of q_i ? It turns out that maximizing the posterior probability corresponds to minimizing the divergence of \hat{q}_i^{sh} from q_i .

As the parameters (and estimators) we are dealing with are probabilities themselves, they can be understood as points in a finite simplex (which happens to be the CoDA sample space). From an information-geometric point of view, the shrinkage estimator is optimized along the *mixture geodesic* (or *m-geodesic*) between the equidistribution

Fig. 1 Exponential (curved orange line) and mixture (blue straight line) geodesics between the equidistribution $(1/3, 1/3, 1/3)$ and an observed point $(n_1/n, n_2/n, n_3/n)$ in the 3-part simplex



and the observed point $(\hat{q}_i)_{i=1}^D$ (see the blue line in Fig. 1). Geodesics provide intuition, e.g., a generalized Pythagorean theorem makes use of them. Unlike in Euclidean geometry, however, we need *two* types of geodesics for Pythagoras to work. The natural counterparts to m -geodesics are the *exponential geodesics* (or e -geodesics). These are convex combinations of points in exponential coordinates, which are dual to the mixture coordinates (via the Legendre duality that underlies information geometry). Let us now consider the e -geodesic between the two points in question (see the orange curve in Fig. 1).

It turns out that the e -geodesic corresponds to an alternative parametrization of the posterior probability, where the prior and likelihood contribute via weighted geometric means. A point on the e -geodesic is just another estimator of the posterior mean that uses this alternative parametrization. When back-transforming exponential coordinates to the original parameter, this geodesic can be written as

$$\hat{q}_i^{es} = \frac{\hat{q}_i^\beta}{\sum_{k=1}^D \hat{q}_k^\beta}, \tag{3}$$

with $0 \leq \beta \leq 1$. This kind of exponential scaling is well known in statistical physics, where β is the inverse temperature. It is also used when Box–Cox transforming data to reduce skew or to replace logarithms by approximate expressions when zeros are involved. In the CoDA context, β can be used to mediate between χ -squared distance and Aitchison distance and thus makes a connection between log-ratio analysis and Correspondence Analysis (CA) [7]. The latter can handle zeros while the former needs to impute them.

For finding the optimal value of the shrinkage parameter λ , a simple analytic solution for minimization of the mean squared error (MSE) with respect to the true parameter can be found [8, 9]. To use the same strategy for the β -parameter of the e -geodesic, we propose to use an MSE on the tangent space. This is just expected Aitchison dis-

tance between the estimator and the true parameter. We derive an analytic solution that approximates an optimal β based on the Delta method (i.e., via Taylor expansion). This is computationally inexpensive and can, e.g., be used as a data preprocessing for dimension reduction techniques like CA. Simulations show that this approach holds promise for data with many essential zeros. We discuss the exponential shrinkage estimator as an additional tool that avoids the pseudocounts of current procedures in contexts where zero imputation may be inappropriate. On a theoretical level, this contribution aims to unify power transformations with shrinkage under the same conceptual framework.

Section 2 presents essentially review material, with the first two paragraphs dedicated to some very general statistical motivation. We then introduce the information geometric formulation of the multinomial likelihood and posterior and make some methodological excursions of a more technical nature in paragraphs 2.6 and 2.8. In these paragraphs, we reformulate known minimizations of relative entropy and of expected quadratic loss in form of propositions that will serve us in the subsequent application. Section 3 is then dedicated to the application of the material presented. It includes the definition of an alternative shrinkage estimator and its optimisation along the exponential geodesic as well as a benchmark of it using simulations. All the proofs and some of the more lengthy algebraic derivations are deferred to the Appendix.

2 Preliminaries

2.1 Sequencing data are relative

Let us first discuss the practical relevance of relative counts for contemporary biomedical data. While it is usually acknowledged that data produced by DNA sequencing instruments are relative [10], a number of arguments for the current dominance of absolute approaches have been put forward. We will discuss one of these arguments here: The constraint on the counts does not hold strictly, i.e., it is itself a fluctuating quantity [11].

Counting the times n_j a specific event j occurs within a fixed time interval, under very general assumptions (i.e., independence of events from previous occurrences, fixed average rate of occurrence, no simultaneous occurrences), the resulting data will be distributed according to a Poisson distribution:

$$p_P(n_j | \lambda_j) = \frac{\lambda_j^{n_j}}{n_j!} e^{-\lambda_j}. \quad (4)$$

Here, λ_j denotes the average occurrence rate¹ of an event j . When considering D such events now, and assuming they don't influence each other, we can write the overall probability of the D -dimensional vector of counts \mathbf{n} simply as a product of D such distributions.

¹ Recall that the λ parameter coincides with the expected counts and also their variance. In practice, this could, e.g., be gene-transcriptional activities [12].

Consider now a modification of this scenario where we observe these D events taking place but instead of fixing a time interval, we will simply stop counting after we have observed n events. The resulting distribution is a multinomial

$$p_n(\mathbf{n} \mid \mathbf{q}) = \frac{n!}{\prod_{j=1}^D n_j!} \prod_{j=1}^D q_j^{n_j}, \tag{5}$$

where $\mathbf{q} = (q_j)_{j=1}^D$ is the vector of individual event probabilities.² The multinomial encodes a constraint on \mathbf{n} that leads to a mutual dependence between the parts. In this sense, it models a composition of counts.

To see the connection between these two scenarios, let us come back to the independent Poisson distribution. It can be written as

$$\begin{aligned} p_P(\mathbf{n} \mid \boldsymbol{\lambda}) &= \prod_{j=1}^D \frac{\lambda_j^{n_j}}{n_j!} e^{-\lambda_j} \\ &= \frac{\lambda^n}{n!} e^{-\lambda} \frac{n!}{\prod_{j=1}^D n_j!} \prod_{j=1}^D \left(\frac{\lambda_j}{\lambda}\right)^{n_j} = p_P(n \mid \lambda) p_n(\mathbf{n} \mid \mathbf{q}). \end{aligned} \tag{6}$$

Here λ denotes the sum over the components of $\boldsymbol{\lambda}$, and $\mathbf{q} = \boldsymbol{\lambda}/\lambda$. We see that the independent Poisson distributions factorize into a univariate Poisson of n with parameter λ as well as a multinomial distribution p_n that has n and \mathbf{q} as parameters. This well-known relationship between the Poisson and the multinomial is interesting when discussing the argument against compositionality above. First we note that a variation in the constraining variable n can only be used for a correct estimate of the rate parameters λ_j of the D Poisson processes if the overall rate λ is exactly their sum. Modelling by a multinomial can thus be perfectly justified for a stochastic n whose rate γ is of no interest to the analyst because it is decoupled from the $\boldsymbol{\lambda}$, in the sense that $\gamma \neq \lambda$. For sequencing data, the constraint on n is imposed by the capacity of the sequencing instrument while the variation in n can be caused by other aspects of the protocol (e.g., the subsequent read mapping). The practical effects of the constraint are well documented [13, 14] and aren't invalidated by the stochastic nature of n .

For an applicaton of the multinomial to single-cell sequencing data, see [15]. A pragmatic approach is taken in [12], where it is acknowledged that the q_j , not the λ_j should be the modelling objective, but (for practical reasons) their modelling is done by an independent Poisson that is reparametrized as $p_P(\mathbf{n} \mid \lambda, \mathbf{q})$. The Poisson can serve as an approximation whenever there are no dominant parts for which q_j becomes too large. The modelling gets complicated again as soon as co-variation of parts across samples are taken into account.

² Note that we chose to put the auxiliary parameter n as a subscript for a more compact notation.

2.2 Variation across samples, Bayes

According to the Bayesian paradigm, probabilities are subjective in the sense that they quantify degrees of knowledge [16]. This quantification involves both data and model parameters, and both can be arguments to probability functions. While we assume a fixed parameter when considering a single sample \mathbf{n} , it makes sense to let the parameter vary according to some distribution when considering many samples that were obtained under different conditions. This is typically the case when we have a data matrix where counts for D variables (or compositional parts) indexed by the columns are collected in N samples indexed by the rows.

As an example, consider the special case of the multinomial p_n . Our choice of the prior π quantifying the probability of the parameter \mathbf{q} will determine the functional form of the joint distribution and thus affect our ability to capture the variability across samples:

$$p_n(\mathbf{n}, \mathbf{q}) = p_n(\mathbf{n} \mid \mathbf{q})\pi(\mathbf{q}). \quad (7)$$

Integrating the joint probability³ over the parameter \mathbf{q} would leave us again with \mathbf{n} as the only argument. The resulting marginal distribution will depend on the hyperparameters of the prior (which we left out in the formula above).⁴ If we divide (7) by it, we renormalize and obtain the posterior probability of the parameter \mathbf{q} , giving us Bayes' theorem.

An excellent choice for π would be a $D - 1$ -dimensional multivariate normal of the log-ratios $\log(q_i/q_D)$. This allows for a compositional modelling of the second-order interactions between parts that captures the over-dispersion often observed in real-world data [17, 18]. While this logistic-normal multinomial model has no analytic solution, Markov-Chain Monte Carlo can be used, like in a recent application to differential association networks in microbiome data [19]. Note that the interest is now in the hyperparameters of the prior, especially in the covariance matrix of the log-ratios of \mathbf{q} .

A less realistic but more tractable solution is obtained when simply choosing the conjugate prior to the multinomial, i.e., the Dirichlet distribution. While we will later describe it in more detail, let us here point out that this model implies that all interaction between parts comes from the constraint that counts have to add to n . It is thus the model with the greatest degree of independence that can be achieved for compositions [2].

2.3 Dual coordinates for count distributions

We have recently proposed to treat compositional data with the methods of information geometry [6]. The fact that the geometric structure of the discrete probability simplex

³ While we use the convention to denote it by the same symbol as the likelihood, this is generally not a multinomial.

⁴ An example that concerns much of the current modelling of sequencing data is going from the Poisson distribution to the (overdispersed) negative binomial distribution when integrating out the original λ_j parameter with a conjugate gamma prior.

can be exploited for the analysis of compositional data has been observed before, e.g. [20]. Compositions \mathbf{q} can be described as categorical distributions that live on a finite dimensional open⁵ simplex

$$\mathcal{S}^D = \left\{ (q_1, \dots, q_D)^T \in \mathbb{R}^D : q_i > 0, i = 1, \dots, D, \sum_i^D q_i = 1 \right\}. \tag{8}$$

The finite version of information geometry contains already all its important concepts but often provides a more intuitive approach, see [5, 21]. For a comprehensive treatment of the finite case, see Chapter 2 of [22]. We are now showing a concrete example of an application to CoDA that slightly extends our framework in [6] to deal with relative count data.

To briefly recapitulate, we start from the two natural coordinate systems used in information geometry: the expectation parameters $\boldsymbol{\eta}$ (whose components carry lower indices) and the exponential parameters $\boldsymbol{\theta}$ (with upper indices). Consider again the case where the occurrence of D discrete events is encoded by a random variable $R = r \in \{1, \dots, D\}$ with occurrence probabilities \mathbf{q} . The $D - 1$ -dimensional vector of expectation parameters $\boldsymbol{\eta}$ consists simply of those probabilities that can vary freely (while all of them have to sum to 1). The probability of an event in terms of $\boldsymbol{\eta}$ can then be written as

$$p(r | \boldsymbol{\eta}) = \begin{cases} \eta_r & \text{if } r \leq D - 1, \\ 1 - \sum_{i=1}^{D-1} \eta_i & \text{if } r = D, \end{cases} \quad r = 1, \dots, D. \tag{9}$$

Alternatively, this distribution can be parametrized using what is known as the alr-transformation in CoDA:

$$\theta^j = \log \frac{q_j}{q_D}, \quad j = 1, \dots, D - 1. \tag{10}$$

Note that we are not (as often done in CoDA) log-ratio transforming the data themselves, but their underlying parameters \mathbf{q} . With this, we can write our distribution in the form

$$p(r | \boldsymbol{\theta}) = \exp \left(\sum_{k=1}^{D-1} \theta^k \mathbb{1}_k(r) - \psi(\boldsymbol{\theta}) \right), \quad r = 1, \dots, D, \tag{11}$$

where $\mathbb{1}_k(r) = 1$ if $r = k$, and $\mathbb{1}_k(r) = 0$ otherwise. The function ψ ensures normalization and is known as the free energy. It is given by

$$\psi(\boldsymbol{\theta}) = \log \left(1 + \sum_{i=1}^{D-1} e^{\theta^i} \right) = -\log q_D. \tag{12}$$

⁵ This is a technical requirement so we can use logarithms. More often than not, compositional data will fall on a closed simplex [1].

How do we get from a single outcome r to the multinomial counts \mathbf{n} ? Let us first consider n outcomes $\mathbf{r} = (r_1, \dots, r_n)$. Their probability is simply the product over (11):

$$\begin{aligned} p(\mathbf{r} \mid n, \boldsymbol{\theta}) &= \prod_{i=1}^n p(r_i \mid \boldsymbol{\theta}) \\ &= \exp \sum_{i=1}^n \left(\sum_{k=1}^{D-1} \theta^k \mathbb{1}_k(r_i) - \psi(\boldsymbol{\theta}) \right), \\ &= \exp \left(\sum_{k=1}^{D-1} \theta^k n_k(\mathbf{r}) - n\psi(\boldsymbol{\theta}) \right), \end{aligned} \quad (13)$$

where $n_k(\mathbf{r}) := \sum_{i=1}^n \mathbb{1}_k(r_i)$. This latter expression encodes the D components of our relative counts \mathbf{n} . To obtain their probability of occurrence, we note that many outcomes \mathbf{r} lead to the same outcomes of counts. Counting these leads to a factor given by the multinomial coefficient:

$$p_0(\mathbf{n} \mid n) = \frac{n!}{n_1! \dots n_D!} = \binom{n}{n_1 \dots n_D}. \quad (14)$$

With this base measure, we can finally write our multinomial (5) in form of an exponential family

$$p_n(\mathbf{n} \mid \boldsymbol{\theta}) = p_0(\mathbf{n} \mid n) \exp \left(\sum_{k=1}^{D-1} \theta^k n_k - n\psi(\boldsymbol{\theta}) \right). \quad (15)$$

We see that the exponential coordinates remain the same regardless of the number of observations. It is often convenient to drop the base measure and, changing the random variable, resort to the expression (13). Also, as we can see from (15), to obtain the multi-event versions of $\boldsymbol{\eta}$ and $\psi(\boldsymbol{\theta})$, we just need to multiply by n . Due to the Legendre duality of the natural coordinates, we can obtain the multi-event expectation coordinates by taking partial derivatives

$$n\eta_j = \frac{\partial}{\partial \theta^j} n\psi(\boldsymbol{\theta}) = \mathbb{E}_{p_n}(n_j) = nq_j, \quad j = 1, \dots, D-1. \quad (16)$$

Finally, the potential that is dual to the multi-event free energy $n\psi(\boldsymbol{\theta})$, i.e., the negative Shannon entropy of (13), is given by $n\phi(\boldsymbol{\eta})$, where

$$\phi(\boldsymbol{\eta}) = \sum_{k=1}^{D-1} \eta_k \log \eta_k + \left(1 - \sum_{k=1}^{D-1} \eta_k \right) \log \left(1 - \sum_{k=1}^{D-1} \eta_k \right). \quad (17)$$

2.4 Parameter divergence from observed points

In the previous section, we have derived expressions for probabilities of data given some model parameters. These parameters happen to be compositions, and as such they can be depicted as points in a simplex. When normalizing a sample of count data by their total, we can also represent it as a so-called *observed point* [5] in the simplex:

$$\hat{\mathbf{q}} = \left(\frac{n_1}{n}, \dots, \frac{n_D}{n} \right)^T. \tag{18}$$

This is the empirical estimate of the parameter \mathbf{q} . The empirical estimate is also known as the *type* of a sequence \mathbf{r} of independent random variables. Our dual coordinates associated with the observed point are

$$\hat{\boldsymbol{\theta}} = \left(\log \frac{n_1}{n_D}, \dots, \log \frac{n_{D-1}}{n_D} \right)^T, \tag{19}$$

$$n\hat{\boldsymbol{\eta}} = (n_1, \dots, n_{D-1})^T. \tag{20}$$

One of the fundamental results of the method of types (e.g., [23]) is an equality relating the true distribution to the observed point:

$$p(\mathbf{r} \mid n, \boldsymbol{\theta}) = \exp(n\phi(\hat{\boldsymbol{\eta}}) - nD_\phi(\hat{\mathbf{q}} \parallel \mathbf{q})), \tag{21}$$

where

$$D_\phi(\hat{\mathbf{q}} \parallel \mathbf{q}) = \sum_{j=1}^D \frac{n_j}{n} \log \frac{n_j}{nq_j} \tag{22}$$

is the relative entropy, or Kullback–Leibler (KL) divergence, between the empirical and the true parameter compositions. The expression (21) can be easily derived by simple algebraic rearrangement of (13) using the expressions for ϕ and D_ϕ . With (21), it is clear that we can write the multi-event version of our divergence as

$$nD_\phi(\hat{\mathbf{q}} \parallel \mathbf{q}) = n\phi(\hat{\boldsymbol{\eta}}) - \log p(\mathbf{r} \mid n, \boldsymbol{\theta}). \tag{23}$$

As the first term does not depend on $\boldsymbol{\theta}$, this shows why taking the maximum of the likelihood $p(\mathbf{r} \mid n, \boldsymbol{\theta})$ over $\boldsymbol{\theta}$ is equivalent to minimizing the KL-divergence between the estimated and the true parameter composition.

More general relationships of this kind can be derived from a fundamental information-geometric equality that is due to the Legendre duality between ψ and ϕ :

$$D_\phi(\hat{\mathbf{q}} \parallel \mathbf{q}) = \phi(\hat{\boldsymbol{\eta}}) + \psi(\boldsymbol{\theta}) - \boldsymbol{\theta}^T \hat{\boldsymbol{\eta}}. \tag{24}$$

Minimizing a dissimilarity between distributions can be understood as a projection. Here we project the observed point onto the manifold of distributions parametrized by

θ . In information geometry, this minimization of the KL-divergence is known under the name of m -projection, see [5]. In Sect. 2.6, we will show a result that is more general than (23) in the sense that it does not only hold for the likelihood but also for prior and posterior probability.

2.5 Posterior probability of the parameter

For the Bayesian estimation of a parameter we have to construct a posterior distribution of the parameter that also takes into account its prior distribution π , which itself can depend on a vector of hyperparameters α . For a review of Bayesian inference for categorical data see [25]. The posterior probability density of the parameter in terms of the exponential parameter θ is

$$p(\theta | \mathbf{r}, n, \alpha) = \frac{p(\mathbf{r} | n, \theta)\pi(\theta | \alpha)}{\int d\theta' p(\mathbf{r} | n, \theta')\pi(\theta' | \alpha)}. \quad (25)$$

Instead of maximizing the likelihood over θ , we can now maximize the posterior to obtain the best parameter estimate.⁶ Inserting (13), the posterior (25) evaluates to

$$p(\theta | \mathbf{r}, n, \alpha) = \pi(\theta | \alpha) \exp \left(\sum_{k=1}^{D-1} \theta^k n_k(\mathbf{r}) - n\psi(\theta) - \log p(\mathbf{r} | \alpha) \right). \quad (26)$$

where $p(\mathbf{r} | \alpha)$ is the normalizing integral in the denominator of (25). Seeing this as an exponential family, we note that the parameter and the random variables have exchanged their roles. The prior can be written as a new base measure now, while the new free energy is given by $\log p(\mathbf{r} | \alpha)$.⁷

A prior that has the same functional form as the resulting posterior is called a conjugate prior. Using a conjugate prior makes closed-form solutions of the posterior possible. The general form of the conjugate prior for an exponential family is well known [24], but it is instructive to obtain it as follows. We copy the functional form of (26) and obtain a D -parameter conjugate prior as

$$\pi(\theta | \alpha) = \pi_0(\theta) \exp \left(\sum_{k=1}^{D-1} \theta^k f_k(\alpha) - \left[\sum_{k=1}^D f_k(\alpha) \right] \psi(\theta) - \chi(\alpha) \right), \quad (27)$$

⁶ Alternatively, we could take the expectation value of θ with respect to its posterior.

⁷ To explain the extra term $-n\psi(\theta)$ in this picture, n and $-\psi(\theta)$ can be considered extra components of the vectors \mathbf{n} and θ , respectively.

where π_0 is a base measure, f_k is a sufficient statistic of the k -th hyperparameter, and χ the normalization. With this, the posterior (26) becomes

$$p(\boldsymbol{\theta} \mid \mathbf{r}, n, \boldsymbol{\alpha}) = \pi_0(\boldsymbol{\theta}) \times \exp \left(\sum_{k=1}^{D-1} \theta^k (n_k(\mathbf{r}) + f_k(\boldsymbol{\alpha})) - \left[n + \sum_{k=1}^D f_k(\boldsymbol{\alpha}) \right] \psi(\boldsymbol{\theta}) - \chi(\boldsymbol{\alpha}) - \log p(\mathbf{r} \mid \boldsymbol{\alpha}) \right). \tag{28}$$

In our categorical case it is well known [25] that the conjugate prior is a Dirichlet distribution with parameters $\boldsymbol{\alpha}$. The expressions involved evaluate to

$$f_k(\boldsymbol{\alpha}) = \alpha_k, \tag{29}$$

$$\pi_0(\boldsymbol{\theta}) = 1, \tag{30}$$

$$\chi(\boldsymbol{\alpha}) = \log B(\boldsymbol{\alpha}), \tag{31}$$

$$p(\mathbf{r} \mid \boldsymbol{\alpha}) = \frac{B((n_k(\mathbf{r}) + \alpha_k)_{k=1}^D)}{B(\boldsymbol{\alpha})}, \tag{32}$$

where B denotes the multivariate beta function. (For clarity, we give a short derivation for $p(\mathbf{r} \mid \boldsymbol{\alpha})$ in the Appendix.) With these expressions, the posterior simplifies to

$$p(\boldsymbol{\theta} \mid \mathbf{r}, n, \boldsymbol{\alpha}) = \exp \left(\sum_{k=1}^{D-1} \theta^k (n_k(\mathbf{r}) + \alpha_k) - \left[n + \sum_{k=1}^D \alpha_k \right] \psi(\boldsymbol{\theta}) - \log B(\mathbf{n}(\mathbf{r}) + \boldsymbol{\alpha}) \right). \tag{33}$$

We can see here the widely-used result that the posterior is obtained from the likelihood by simply adding the conjugate prior parameters as pseudo counts to the respective event counts and then renormalizing.

2.6 Parameter divergence from general estimators

The similarity between the likelihood and our expression for the posterior suggests that we can maximize the posterior similarly to the likelihood by minimizing a certain KL-divergence. Indeed, the following proposition shows that maximizing prior, likelihood, or posterior always corresponds to a minimization of KL-divergence between a suitable estimator and \mathbf{q} :

Proposition 1 *Let \mathbf{q} be a parameter of probabilities with exponential coordinates $\boldsymbol{\theta}$ via $p(\mathbf{r} \mid \boldsymbol{\theta})$ with free energy $\psi(\boldsymbol{\theta})$ as defined in (10)–(12). Further, let the function $f : \mathcal{S}^D \times \mathbb{R}_+ \times \mathbb{R}^{D-1} \rightarrow \mathbb{R}_+$ be given by*

$$f(\tilde{\mathbf{q}}, \tilde{n}, \boldsymbol{\theta}) = Z(\tilde{n}, \tilde{\mathbf{q}}) \exp \left\{ \tilde{n} \left(\boldsymbol{\theta}^T \tilde{\boldsymbol{\eta}} - \psi(\boldsymbol{\theta}) \right) \right\},$$

where $\tilde{\mathbf{q}}$ is an estimator of \mathbf{q} with expectation coordinates $\tilde{\boldsymbol{\eta}}$, \tilde{n} denotes a positive real, and Z a positive function. We then have

$$\tilde{n}D_\phi(\tilde{\mathbf{q}} \parallel \mathbf{q}) = \tilde{n}\phi(\tilde{\boldsymbol{\eta}}) + Z(\tilde{n}, \tilde{\mathbf{q}}) - \log f(\tilde{\mathbf{q}}, \tilde{n}, \boldsymbol{\theta}),$$

with ϕ the Lagrange dual to ψ as defined in (17) and D_ϕ the KL-divergence.

The proof makes use of (24) and otherwise consists in a simple rearrangement of terms (see Appendix).

Corollary 1 Maximization of $\log f(\tilde{\mathbf{q}}, \tilde{n}, \boldsymbol{\theta})$ as a function of $\boldsymbol{\theta}$ minimizes $D_\phi(\tilde{\mathbf{q}} \parallel \mathbf{q})$ as a function of \mathbf{q} .

This is clear because the other (data-dependent) terms do not depend on the parameter.

Example 1 Shrinkage estimator:

We use as our estimator $\tilde{\mathbf{q}}$ the expected value of \mathbf{q} under the posterior (33), the so-called shrinkage estimator $\hat{\mathbf{q}}^{\text{sh}}$

$$\tilde{\mathbf{q}} = \hat{\mathbf{q}}^{\text{sh}} := \mathbb{E}_\theta(\mathbf{q} \mid \mathbf{r}, n, \boldsymbol{\alpha}) = \frac{\mathbf{n} + \boldsymbol{\alpha}}{n + \sum_{k=1}^D \alpha_k}, \tag{34}$$

and set $\tilde{n} = \hat{n} := n + \sum_{k=1}^D \alpha_k$. This allows us to reparametrize the posterior in the required form

$$p(\boldsymbol{\theta} \mid \hat{\mathbf{q}}^{\text{sh}}, \hat{n}) = \exp\left(\hat{n} \left[\sum_{k=1}^{D-1} \theta^k \hat{q}_k^{\text{sh}} - \psi(\boldsymbol{\theta}) \right] - \log B(\hat{n}\hat{\mathbf{q}}^{\text{sh}})\right), \tag{35}$$

and thus $f(\tilde{\mathbf{q}}, \tilde{n}, \boldsymbol{\theta}) = p(\boldsymbol{\theta} \mid \mathbf{r}, n, \boldsymbol{\alpha})$ and $Z(\tilde{n}, \tilde{\mathbf{q}}) = 1/B(\hat{n}\hat{\mathbf{q}}^{\text{sh}})$. With this, the proposition gives

$$\hat{n}D_\phi(\hat{\mathbf{q}}^{\text{sh}} \parallel \mathbf{q}) = \hat{n}\phi(\hat{\boldsymbol{\eta}}^{\text{sh}}) - \log B(\hat{n}\hat{\mathbf{q}}^{\text{sh}}) - \log p(\boldsymbol{\theta} \mid \hat{\mathbf{q}}^{\text{sh}}, \hat{n}). \tag{36}$$

Thus finding the $\boldsymbol{\theta}$ that maximizes the posterior is equivalent to minimizing the KL-divergence between the shrinkage estimator and the true parameter \mathbf{q} .

Example 2 Empirical estimator:

The empirical estimator of the multinomial distribution is a straightforward application: $\tilde{\mathbf{q}} = \hat{\mathbf{q}} := \mathbf{n}/n$, $\tilde{n} = n$, and $f(\tilde{\mathbf{q}}, \tilde{n}, \boldsymbol{\theta}) = p_n(\mathbf{n} \mid \boldsymbol{\theta})$ as given by (15), so $Z(\tilde{n}, \tilde{\mathbf{q}})$ is the multinomial coefficient $p_0(\mathbf{n} \mid n)$. The proposition gives (23) with an additional subtraction of the $\log p_0$ term.

Clearly, another example consists in maximizing the prior probability of $\boldsymbol{\theta}$ to minimize the divergence between $\boldsymbol{\alpha}/\sum_k \alpha_k$ and \mathbf{q} . In Sect. 3 we will define another version of the shrinkage estimator, which will provide us with yet another application of the proposition. Note that $f(\tilde{\mathbf{q}}, \tilde{n}, \boldsymbol{\theta})$ has the general form of a conjugate prior of an exponential family, so Proposition 1 holds for exponential families in general. A more general treatment than the one presented here can be found in [26].

2.7 Decision-theoretic risk

Decision theory (e.g., [27]) provides a foundational framework for statistics. While it is closely linked with Bayesian analysis, it can also be formulated from a frequentist point of view. In any case, it implies the construction of a loss function that incorporates statistical knowledge in order to quantify the risk of a wrong decision. Such a loss function L has the “true state of nature” and an action (based on some knowledge) as its arguments. Perhaps the most important example for these arguments would be the true parameter \mathbf{q} of a distribution and some estimator $\hat{\mathbf{q}}$, where the latter would be identified with the action based on it. Given some loss $L(\mathbf{q}, \hat{\mathbf{q}})$, the risk we incur when basing our decision on the estimator is then some expected value

$$R(\hat{\mathbf{q}}) = \mathbb{E}L(\mathbf{q}, \hat{\mathbf{q}}). \tag{37}$$

Bayesian and frequentist schools disagree on the type of expectation that should be taken here. While for the Bayesian the expectation is taken with respect to the posterior probability⁸ of the parameter \mathbf{q} , the frequentist averages over all instances of the random variables (which follow a distribution parametrized by \mathbf{q})⁹. As a consequence, the risk remains a function of \mathbf{q} . A frequentist then calls an estimator $\hat{\mathbf{q}}_1$ R -better than $\hat{\mathbf{q}}_2$ when $R_{\mathbf{q}}(\hat{\mathbf{q}}_1) \leq R_{\mathbf{q}}(\hat{\mathbf{q}}_2)$ for all \mathbf{q} , with strict inequality for some of them. An estimator is called *inadmissible* if there exists an R -better estimator.

Often, for pragmatic reasons, a quadratic loss leading to a mean squared error (MSE) risk function is assumed. Beside its simplicity, one benefit is that for unbiased estimators, the (frequentist) risk is simply the variance of the estimator:

$$R_{\mathbf{q}}(\hat{\mathbf{q}}) = \mathbb{E} \left[(\hat{\mathbf{q}} - \mathbf{q})^2 \right] = \sum_{j=1}^D \left[\text{var}(\hat{q}_j - q_j) + \mathbb{E}^2(\hat{q}_j - q_j) \right] = \sum_{j=1}^D \text{var}(\hat{q}_j). \tag{38}$$

Here, the bias-variance decomposition of the MSE was used, and the last equality follows from the facts that q_j is not stochastic and that the bias $\mathbb{E}[\hat{\mathbf{q}} - \mathbf{q}]$ vanishes. Note that here we do not have to know the true value of \mathbf{q} to evaluate its risk because in practice, to evaluate the variance of the estimator, its empirical estimate is used. As an example, for the empirical estimator (18), the variance components would be estimated by $\hat{q}_j(1 - \hat{q}_j)/(n - 1)$.

2.8 James–Stein shrinkage and regularization

The empirical estimator $\hat{\mathbf{q}}$ is (unlike the empirical estimator of the multivariate normal mean) known to be admissible under quadratic loss [28], so there is no “Stein effect” [29] for the multinomial. While the Bayesian estimator (34) isn’t uniformly better

⁸ In a data-free context, it can also be taken with respect to the prior probability.

⁹ An example of such a risk function is $D_{\phi}(\hat{\mathbf{q}} \parallel \mathbf{q})$.

than the empirical estimator for all parameter values,¹⁰ its flattening of the data can result in much smaller mean squared error than with the empirical estimate. This will be made plausible in the following. Let us rewrite (34) as a convex combination

$$\hat{q}^{\text{sh}} = \lambda \tau + (1 - \lambda) \hat{q} \tag{39}$$

between the *target distribution* τ and the empirical estimator \hat{q} . That this is equivalent to (34) can be seen when defining

$$\lambda := \frac{\sum_{k=1}^D \alpha_k}{n + \sum_{k=1}^D \alpha_k}, \tag{40}$$

$$\tau_j := \frac{\alpha_j}{\sum_{k=1}^D \alpha_k}, \quad j = 1, \dots, D. \tag{41}$$

\hat{q}^{sh} is called a James-Stein type [30] shrinkage estimator of q , see also [31] as well as the discussion in [9]. Choosing the maximum-entropy target, i.e., the equidistribution $\tau_j = 1/D$ for all $j = 1, \dots, D$, the target term can be understood as a regularization of the empirical estimator.

Remember that \hat{q}^{sh} is the posterior expected value of q . The fact that the posterior expected value of a random variable is a linear function of its empirical estimate is equivalent to the use of a conjugate prior. This is a result that holds for exponential families in general [24].

This linearity is helpful for evaluating the accuracy of the shrinkage estimator, again using the expected quadratic loss as a risk function. We shall give a result that is slightly more general than necessary for this estimator because we will again need it in Sect. 3:

Proposition 2 *Let $f_j, j = 1, \dots, D$ be the components of a function $f : \mathcal{S}^D \rightarrow \mathbb{R}^D$ acting on a vector of probabilities. Let τ be a D -dimensional probability parameter and \hat{q} the multinomial empirical estimator. Then, for $0 \leq \lambda \leq 1$, the convexly combined estimator $f(\tilde{q})$ of $f(q)$ given by its components*

$$f_j(\tilde{q}) := \lambda f_j(\tau) + (1 - \lambda) f_j(\hat{q}), \quad j = 1, \dots, D$$

(i) *has a quadratic risk with respect to $f(q)$ given by*

$$R_q(\tilde{q}) = (1 - \lambda)^2 \sum_{j=1}^D \text{var}(f_j(\hat{q})) + \sum_{j=1}^D \left[\mathbb{E} f_j(\hat{q}) - f_j(q) - \lambda(\mathbb{E} f_j(\hat{q}) - f_j(\tau)) \right]^2.$$

¹⁰ An example where the empirical estimator gives a better value for q_j is the case where $n_j = 0$ and the prior value of the Bayesian estimator is further away from q_j than q_j is from zero.

(ii) *The minimum risk is attained for*

$$\lambda^* = \frac{\sum_{j=1}^D \left[\text{var}(f_j(\hat{q})) + (\mathbb{E}f_j(\hat{q}) - f_j(\mathbf{q}))(\mathbb{E}f_j(\hat{q}) - f_j(\boldsymbol{\tau})) \right]}{\sum_{j=1}^D \mathbb{E}[f_j(\hat{q}) - f_j(\boldsymbol{\tau})]^2}.$$

The proof is provided in the Appendix. This is a slight modification of the lemma shown in [8], see also the derivation in [32] and the application to the multinomial in [9]. To apply the proposition to \hat{q}^{sh} , we observe that f_j simply corresponds to taking the j -th component and simplifications occur because the bias of \hat{q} vanishes: $\mathbb{E}f_j(\hat{q}) - f_j(\mathbf{q}) = \mathbb{E}\hat{q}_j - q_j = 0$. We obtain

$$R_{\mathbf{q}}(\hat{q}^{\text{sh}}) = (1 - \lambda)^2 \sum_{j=1}^D \text{var}(\hat{q}_j) + \lambda^2 \sum_{j=1}^D \mathbb{E}^2(\hat{q}_j - \tau_j), \tag{42}$$

with minimum risk at

$$\lambda^* = \frac{\sum_{j=1}^D \text{var}(\hat{q}_j)}{\sum_{j=1}^D \mathbb{E}[(\hat{q}_j - \tau_j)^2]}. \tag{43}$$

We can see that the risk function is a weighted average over the risk of the empirical estimator and an additional term that punishes expected difference from the target. Tuning the size of λ , we can trade off the bias of the target against the variance of the empirical estimate to obtain a smaller risk than (38). Estimators based on small sample data will generalize better to new data when flattening the data to a well-specified extent using an uninformative, maximum-entropy model. The amount of flattening depends on the data at hand and is optimized via the weight λ of the target. Note that the relationships (40) and (41) imply that this is similar to an empirical Bayes procedure where we tune the size of the pseudocounts α_j and by this, adjust the a-priori sample size $\sum \alpha_k = n\lambda/(1 - \lambda)$. To evaluate (43), the empirical estimates for variance and expectation are used in practice.

2.9 Power-transformed compositions and their Euclidean distance in ordination

Power transformations [33] have traditionally been applied to data in order to fulfill certain distributional assumptions. For instance, a suitable power transformation can reduce skew so data appear approximately normal. In the case of Poisson counts, where variance equals the mean, the square root transformation is a common choice to “stabilize” the variance (i.e., make it approximately constant independently of the mean). More generally, power transformations can appear through the link functions of generalized linear models [35] and then enable a fit of the data to a true underlying distribution.

Methods for dimension reduction and data visualization (a.k.a. *ordination*) such as Principal Component Analysis (PCA) often use some version of Euclidean distance

between multivariate samples:

$$d^2(\hat{\mathbf{q}}_1, \hat{\mathbf{q}}_2) = \sum_{j=1}^D \omega_j (\hat{q}_{1j} - \hat{q}_{2j})^2, \tag{44}$$

where the ω_j are suitable weights. Here, for the data, we used the empirical parameter estimates of the count distribution $\hat{\mathbf{q}}$ instead of the counts \mathbf{n} themselves. In the case of relative counts, where the total of each sample is not of direct interest, this seems a good idea because we want to visualize the “shape” of the data without their “size” [34]. There are two main ordination methods that are relational in the sense that they visualize shape only [35], Correspondence Analysis (CA) and log-ratio analysis (LRA). CA uses a weighting scheme that involves row and column totals of the data matrix. In this way, it takes into account the data size indirectly to account for the precision of the shape estimates. LRA, in contrast, is a PCA of data that are log-transformed and double-centred. Here, relationships between parts remain invariant under taking subsets of the data,¹¹ and it is better suited for true compositions. It was shown [7] that via the following limit of the Box–Cox family [36] of power transformations

$$\lim_{\beta \rightarrow 0} \frac{x^\beta - 1}{\beta} = \log(x), \tag{45}$$

CA on power-transformed data converges to LRA. CA and LRA are thus special cases of a more general family of ordination methods. To make this more precise in the case of unweighted LRA, consider the following transformation of our empirical estimates:

$$f_\beta(\hat{\mathbf{q}}) = \left(\frac{\hat{q}_1^\beta}{\sum_{k=1}^D \hat{q}_k^\beta}, \dots, \frac{\hat{q}_D^\beta}{\sum_{k=1}^D \hat{q}_k^\beta} \right)^T. \tag{46}$$

When now using uniform weights $\omega_j = D^2$, the limit

$$\lim_{\beta \rightarrow 0} \frac{1}{\beta^2} d^2(f_\beta(\hat{\mathbf{q}}_1), f_\beta(\hat{\mathbf{q}}_2)) \tag{47}$$

is the squared Aitchison distance

$$d_A^2(\hat{\mathbf{q}}_1, \hat{\mathbf{q}}_2) = \frac{1}{D} \sum_{i=1}^D \sum_{j<i}^D \left(\log \frac{\hat{q}_{1i}}{\hat{q}_{1j}} - \log \frac{\hat{q}_{2i}}{\hat{q}_{2j}} \right)^2 \tag{48}$$

(see [6] for a proof). Aitchison (or log-ratio) distance is the metric underlying LRA. Using the transformation f_β before evaluating Euclidean distance induces a parametrized class of distance measures that include the ones used in CA ($\beta = 1$)

¹¹ This property is known as subcompositional coherence.

and LRA ($\beta = 0$) as special cases.¹² When using finite, “small enough” values of the power parameter β , the subcompositional coherence of LRA remains approximately satisfied while there is no need for zero imputation (as CA does not involve logarithms). One can obtain an optimal value of the power parameter in the sense that it maximizes the Procrustes correlation between the log-ratio transformed data (using zero imputation) and the coordinates from the power-transformed CA (keeping the zeros) [37].

3 Exponential shrinkage

In this section we want to define and test an estimator based on the power transformation (46). The justification of this estimator comes from a formal analogy with \hat{q}^{sh} . This analogy is more apparent when introducing the generalized notions of addition (a.k.a. *perturbation*) and scalar multiplication (a.k.a. *powering*) that equip the simplex with a linear structure. For $\mathbf{q}, \mathbf{p} \in \mathcal{S}^D$, and some $\beta \in \mathbb{R}$, they are defined as the vectors

$$\mathbf{q} \oplus \mathbf{p} := \mathcal{C}(q_1 p_1, \dots, q_D p_D)^T, \tag{49}$$

$$\beta \odot \mathbf{q} := \mathcal{C}(q_1^\beta, \dots, q_D^\beta)^T, \tag{50}$$

where \mathcal{C} denotes the closure operation $\mathcal{C}\mathbf{q} := \mathbf{q} / \sum_i q_i$. An inverse perturbation is given by $\ominus \mathbf{q} := \oplus(-1) \odot \mathbf{q}$.

3.1 Power transformed compositions as convex combinations, dual geodesics

The shrinkage estimator (39) is a weighted mean of the target and the observed point. This convex combination is an example for what is known as a mixture geodesic (or m -geodesic) in information geometry. Consider now a similar structure using the operations of perturbation and powering introduced above:

$$\tilde{\mathbf{q}} = \lambda \odot \boldsymbol{\tau} \oplus (1 - \lambda) \odot \hat{\mathbf{q}}. \tag{51}$$

This describes a so-called *exponential* geodesic (or e -geodesic).¹³ Usually [5], both types of geodesics are written in terms of their dual coordinates:

$$\boldsymbol{\eta}(\lambda) = \lambda \boldsymbol{\eta}_\tau + (1 - \lambda) \boldsymbol{\eta}_{\hat{\mathbf{q}}}, \tag{52}$$

$$\boldsymbol{\theta}(\lambda) = \lambda \boldsymbol{\theta}_\tau + (1 - \lambda) \boldsymbol{\theta}_{\hat{\mathbf{q}}}, \tag{53}$$

where we used subscripts to indicate at which points the coordinates are evaluated. Coming back to the power-transformation (46), we can easily see that it is described by the exponential geodesic between the observed point and the uniform target: Evaluating

¹² Note that the row weights are assumed to be uniform for the special case of compositional data.

¹³ This is also known as the *Hellinger arc* connecting two distributions.

the exponential coordinates at $f_\beta(\hat{q})$, we have

$$\theta_{f_\beta(\hat{q})} = \left(\log \frac{\hat{q}_1^\beta}{\hat{q}_D^\beta}, \dots, \log \frac{\hat{q}_{D-1}^\beta}{\hat{q}_D^\beta} \right)^T = \beta \theta_{\hat{q}}. \tag{54}$$

We also notice that for $\tau = (1/D)_{i=1}^D, \theta_\tau$ vanishes. Setting $\beta = 1 - \lambda$, we immediately obtain (53). When evaluating (53) for a general target, we can use the form (51) to obtain a generalized power transformation in terms of the original parameters:

$$\hat{q}^{\text{es}} := \left(\frac{\tau_1^{1-\beta} \hat{q}_1^\beta}{\sum_{k=1}^D \tau_k^{1-\beta} \hat{q}_k^\beta}, \dots, \frac{\tau_D^{1-\beta} \hat{q}_D^\beta}{\sum_{k=1}^D \tau_k^{1-\beta} \hat{q}_k^\beta} \right)^T. \tag{55}$$

Comparing \hat{q}^{es} with the shrinkage estimator (34), we see that instead of a weighted arithmetic mean between the target and the empirical estimator, here we evaluate a weighted geometric mean between them.

3.2 Another reparametrization of the posterior

Since the generalized power transformation (55) can be described as a convex combination in exponential coordinates, it shares a structural similarity with the shrinkage estimator (34), which is obtained from a convex combination of expectation (a.k.a. mixture) coordinates. To make this a shrinkage problem, however, we need the resulting quantity \hat{q}^{es} to be interpreted as an estimator. Here we argue that \hat{q}^{es} is simply a reparametrization of \hat{q}^{sh} similar to (39). There, we went from $\mathcal{C}(n + \alpha)$ to an expression involving λ, τ , and \hat{q} . We also showed a simple reparametrization of the posterior of θ in terms of \hat{q}^{sh} together with the posterior sample size \hat{n} , see (35). Such alternative ways of writing posterior and posterior expectation can be obtained using \hat{q}^{es} as well, as we will show in the following.

As we have seen in the previous section, an alternative parameter β can be used to define a geometric mean between target and observed point. Defining $\tilde{n} := \sum_{k=1}^D \tau_k^{1-\beta} n_k^\beta$, in the expression for the posterior (35) we can simply replace $\hat{n} \hat{q}^{\text{sh}}$ by new Dirichlet parameters $\tilde{n} \hat{q}^{\text{es}}$ to obtain the following expression of the posterior:

$$p(\theta \mid \hat{q}^{\text{es}}, \tilde{n}) = \exp \left(\tilde{n} \left[\sum_{k=1}^{D-1} \theta^k \hat{q}_k^{\text{es}} - \psi(\theta) \right] - \log B(\tilde{n} \hat{q}^{\text{es}}) \right). \tag{56}$$

This provides us with another example for Proposition 1. Maximizing the posterior thus corresponds to a minimization of the KL-divergence between \hat{q}^{es} and the true parameter. Furthermore, the derivation of (32) given in the Appendix also shows that $B(\tilde{n} \hat{q}^{\text{es}})$ normalizes (35).¹⁴ Note that this also implies that the posterior expectation of q can be written equally valid as either the shrinkage estimator \hat{q}^{sh} or as the *exponential*

¹⁴ As $\hat{n} \hat{q}^{\text{sh}} = n + \alpha$, and $\tilde{n} \hat{q}^{\text{es}}$ has the exact same form.

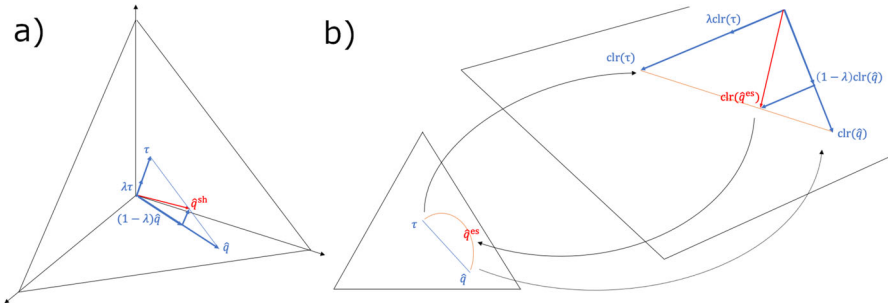


Fig. 2 **a** The shrinkage estimator \hat{q}^{sh} (in red) obtained by an addition of scaled vectors (in blue) ending in the unit simplex (shown in black). The m -geodesic connecting τ and \hat{q} is shown as a thin blue line. **b** The exponential shrinkage estimator \hat{q}^{es} (in red) obtained by vector addition in the tangent space. The e -geodesic is shown as a curved orange line in the simplex and a straight orange line in the tangent space

shrinkage estimator \hat{q}^{es} . This means that the exponential shrinkage estimator is nothing but the reparametrized posterior expectation of q .

3.3 Quadratic risk on the tangent space

To evaluate the accuracy of the exponential shrinkage estimator, we would like a simple risk function like the MSE. We saw previously that with this risk function, an analytic estimate of the optimal prior weight was essentially possible because of the linearity of the shrinkage estimator. However, a generalized notion of linearity is now needed: While m -geodesics are straight lines in the simplex, e -geodesics are straight lines in its tangent space

$$\mathcal{T}^D = \left\{ v \in \mathbb{R}^D : \sum_{i=1}^D v_i = 0 \right\}. \tag{57}$$

A mapping from the simplex to \mathcal{T}^D (a.k.a. *clr plane* in CoDA) is known as the *clr* transformation

$$\text{clr}(q) = \left(\log \frac{q_1}{g(q)}, \dots, \log \frac{q_D}{g(q)} \right)^T, \tag{58}$$

where g denotes the geometric mean $g(x) = \left(\prod_{i=1}^D q_i \right)^{1/D}$. This mapping is fundamental in both information geometry and CoDA. The constraint that the *clr* components sum to zero means that the points on an exponential geodesic retain their normalization on the simplex.

With this, a quadratic loss function in analogy to the one on the simplex can be obtained by first mapping the compositions in question to the tangent space and then using squared Euclidean distance again (see Fig. 2).

Let us first define the loss function on the tangent space for the empirical estimator:

$$L_A(\mathbf{q}, \hat{\mathbf{q}}) = \sum_{j=1}^D (\text{clr}_j(\hat{\mathbf{q}}) - \text{clr}_j(\mathbf{q}))^2. \tag{59}$$

This is the (squared) Aitchison distance, i.e., an alternative expression of (48). Via the mapping of the simplex to \mathcal{T}^D , the expression $\text{clr}(\hat{\mathbf{q}}) - \text{clr}(\mathbf{q})$ can be interpreted as a difference vector between compositions [6]. One can write this in form of a perturbation with the notation $\hat{\mathbf{q}} \ominus \mathbf{q}$, which makes the analogy with (38) even more compelling. The “exponential” analogue to the MSE of Sect. 2.7 is the risk function associated with the squared Aitchison loss, i.e. the expectation

$$\tilde{R}_q(\hat{\mathbf{q}}) = \mathbb{E}L_A(\mathbf{q}, \hat{\mathbf{q}}) = \sum_{j=1}^D \left[\text{var}(\text{clr}_j(\hat{\mathbf{q}})) + \mathbb{E}^2(\text{clr}_j(\hat{\mathbf{q}}) - \text{clr}_j(\mathbf{q})) \right]. \tag{60}$$

Unfortunately, in this case the bias term does not vanish for the empirical estimator, and we shall need an approximation to evaluate it.

3.4 Optimization along the exponential geodesic

We can now use our modified risk function on the exponential shrinkage estimator, in analogy to (42), to minimize it with respect to $\lambda = 1 - \beta$. Using Proposition 2 with $f_j(\cdot) = \text{clr}_j(\cdot)$, and $\lambda = 1 - \beta$, for the MSE of $\text{clr}(\hat{\mathbf{q}}^{\text{es}})$ we obtain

$$R_q(\hat{\mathbf{q}}^{\text{es}}) = (1 - \lambda)^2 \sum_{j=1}^D \text{var}(\text{clr}_j(\hat{\mathbf{q}})) + \sum_{j=1}^D \left[\lambda \mathbb{E}(\text{clr}_j(\boldsymbol{\tau}) - \text{clr}_j(\hat{\mathbf{q}})) + \mathbb{E} \text{clr}_j(\hat{\mathbf{q}}) - \text{clr}_j(\mathbf{q}) \right]^2. \tag{61}$$

A solution for the minimum can be found at

$$\lambda_{\min} = \frac{\sum_{j=1}^D \left[\text{var}(\text{clr}_j(\hat{\mathbf{q}})) - \mathbb{E}(\text{clr}_j(\boldsymbol{\tau}) - \text{clr}_j(\hat{\mathbf{q}})) (\mathbb{E} \text{clr}_j(\hat{\mathbf{q}}) - \text{clr}_j(\mathbf{q})) \right]}{\sum_{j=1}^D \mathbb{E} \left[(\text{clr}_j(\boldsymbol{\tau}) - \text{clr}_j(\hat{\mathbf{q}}))^2 \right]} \tag{62}$$

Again, this can be evaluated in practice by replacing \mathbf{q} by the best estimator available. To estimate the variance and the expectation terms of the clr-transformed empirical estimator, we resort to Taylor expansion. While the expressions become a bit more

unwieldy compared with the ones on the m -geodesic, we can still evaluate them explicitly. For the mean we get

$$\mathbb{E}\text{clr}_j(\hat{\mathbf{q}}) \approx E_j := \text{clr}_j(\mathbf{q}) - \frac{1 - q_j}{2q_j n} + \frac{1}{2D} \sum_{k=1}^D \frac{1 - q_k}{q_k n}, \tag{63}$$

and for the variance (where this approximation is known as the Delta method)

$$\begin{aligned} \text{var}(\text{clr}_j(\hat{\mathbf{q}})) &\approx \\ V_j &:= \left(1 - \frac{2}{D}\right) \frac{1 - q_j}{q_j n} + \frac{1}{D^2} \sum_{k=1}^D \frac{1 - q_k}{q_k n} - \frac{1}{n} \left(3 - \frac{7}{D} + \frac{4}{D^2}\right) \end{aligned} \tag{64}$$

(see Appendix for a derivation). In the case of the maximum-entropy target, the $\text{clr}_j(\boldsymbol{\tau})$ terms in (62) vanish, and an estimator of the optimal power can be obtained by

$$\beta^* = 1 - \frac{\sum_{k=1}^D [V_k - E_k(E_k - \text{clr}_k(\mathbf{q}))]}{\sum_{k=1}^D [V_k + E_k^2]}. \tag{65}$$

3.5 Performance on simulated data

We can now test how well we can infer true frequencies from simulated data using the exponential shrinkage estimator. For this, we use the equidistribution as the target and optimize the β parameter as described before.

This should not be understood as an intent at a comprehensive benchmark but rather as a proof of concept. We test performance on multinomial counts only. The three different multinomial distributions ($D=100$) shown in Fig. 3 were obtained by sampling from Dirichlet distributions with three different choices for the hyper parameters. These were chosen to obtain multinomial parameters that are far from equidistributed and have an increasing number of essential zeros. As a measure of performance, we chose MSE as in [9]. Beside being simple and intuitive, MSE has the advantage that zeros are not problematic as there are no logarithms involved. Both zeros as obtained from undersampling (i.e., *count* zeros) as well as those that occur because parameters are truly (or almost) zero (so-called *essential* zeros) will have the effect that the observed point $\hat{\mathbf{q}}$ falls on the boundary of the simplex. This is not a problem for the shrinkage estimator, as m -geodesics can go from the centre to the boundary. However, e -geodesics are only defined inside the simplex, and we have to redefine the observed point as its projection to the nonzero parts, with a subsequent change in the dimension D . In any case, it is only the nonzero parts that can be modified by the exponential shrinkage estimator. As an approximation of the true parameter in the expressions (63) and (64), we use the shrinkage estimator $\hat{\mathbf{q}}^{\text{sh}}$. The exponential shrinkage estimator is optimized over the nonzero parts only. The results show that exponential shrinkage outperforms the empirical estimator but cannot compete with the shrinkage estimator

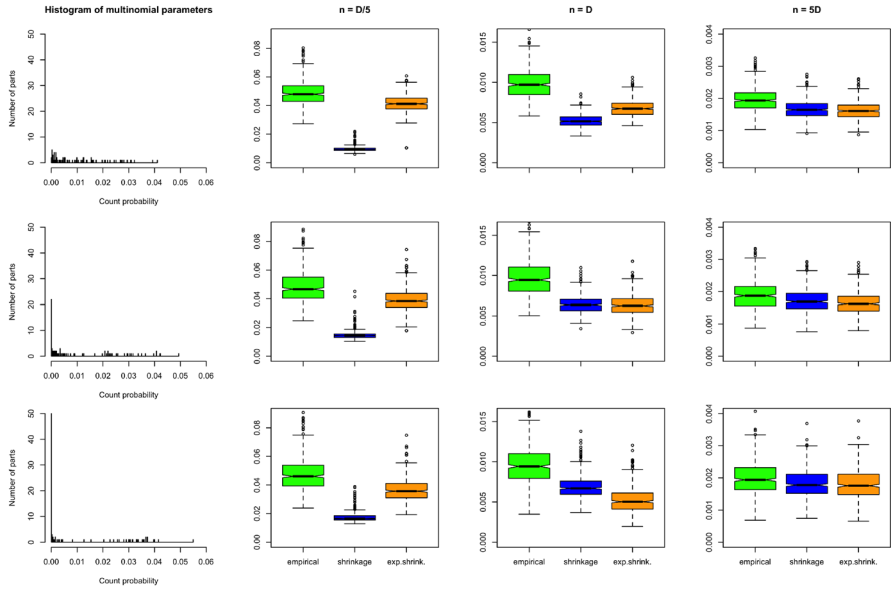


Fig. 3 Mean squared error (MSE) of the empirical estimator (green), the shrinkage estimator (blue), and the exponential shrinkage estimator (orange). Data are sampled from multinomial distributions with increasing sparsity. Boxplots in each row show the MSEs of 500 samples from the multinomial whose histogram is shown in the first column. Sample size increases from left to right ($n = D/5, D, 5D$), while sparsity increases from top down. As $D = 100$, the vertical axis in the histograms can be read as a percentage. Note that the vertical boxplot axes change their range between columns

if the data are severely undersampled (first column in Fig. 3). There is a sweet spot of performance when many essential zeros are present and the data are sampled at reasonable depth (middle column). In this case, the exponential shrinkage estimator can outperform the shrinkage estimator. Clearly, it is “already correct” for the unobserved values, while the shrinkage estimator imputes them. Further increasing sample size essentially equalizes the performance of all estimators (right column). Note that the presence of zeros in the multinomial parameters effectively increases the sample size as the same counts are now distributed over fewer parts. The two factors studied in Fig. 3, sample size and sparsity, are thus not independent of each other in their effects.

3.6 Discussion

We have shown that power transformations of relative count data can be understood as a shrinkage problem. An analytic solution for the optimal power for given data can be obtained in a way that is analogous to what was proposed for finding an optimal flattening constant. We find the underlying information-geometric structure intriguing: Both types of geodesics between the empirical estimate and the maximum-entropy estimate give rise to their own shrinkage problem. But we think that there are also practical implications for data analysis. In the context of compositional data visualization, power transformations have been proposed as an approximation to log-

ratio transformations, which require zero imputation. Correspondence Analysis (CA), one of the best methods for visualizing two-way tables containing counts, can be made more suitable for relative count data when applying such a transformation. It then approximates log-ratio analysis (LRA), whose visualization appeals more to our Euclidean intuition but whose zero imputed data may be suboptimal or even impossible for very sparse data sets. For side-by-side visualizations of geochemical and single-cell data using both methods, see [37]. While CA is a visualization of the stretched out (weighted) simplex, LRA is a PCA on its tangent space (the clr plane). When using the hybrid approach of CA with power transformed counts, currently a uniform power parameter is applied to an entire data matrix that could contain rows with heterogeneous sample sizes. As we have seen, in terms of an optimal approximation to the underlying parameters in each row, this would work best if samples follow the same distribution and the sample sizes are not too different. On the other hand, we could argue that, from a modelling perspective, it would be better to find the best power for each row in the data matrix separately. While the deformation with respect to LRA would now be heterogeneous among samples, the fit with underlying population parameters would be better. The shrinkage approach is of course applicable beyond data visualization, and we think that applying it as a kind of data normalization holds some promise for very sparse data sets as occurring in microbiome analysis or single-cell genomics. Not all of these zeros are essential zeros, but many of them may be caused by truly small occurrence probabilities. If so, the commonly applied log transform with a uniform pseudocount would almost certainly be less suitable than a data-driven power transformation as proposed here. While this approach may still appear overly simplistic, given today’s highly complex data acquisition protocols where effects of statistical and engineering decisions are hard to disentangle, simple approaches often perform similarly well as highly complex ones [38].

Appendix

Derivation of Equation 32

Inserting the expressions (29–31) into the general conjugate prior (27), we obtain

$$\pi(\boldsymbol{\theta} \mid \boldsymbol{\alpha}) = \exp\left(\sum_{k=1}^{D-1} \alpha_k \theta^k - \psi(\boldsymbol{\theta}) \sum_{k=1}^D \alpha_k - \log B(\boldsymbol{\alpha})\right). \tag{66}$$

Together with (13), the denominator in (25) becomes

$$p(\mathbf{r} \mid \boldsymbol{\alpha}) = \frac{1}{B(\boldsymbol{\alpha})} \int d\boldsymbol{\theta} \exp\left(\sum_{k=1}^{D-1} \theta^k (n_k(\mathbf{r}) + \alpha_k) - \left(n + \sum_{k=1}^D \alpha_k\right) \psi(\boldsymbol{\theta})\right). \tag{67}$$

Now a variable transformation to the original parameter \mathbf{q} with Jacobian $\det(\partial\theta_j/\partial q_j)_{j=1}^{D-1} = \prod_{k=1}^D q_k^{-1}$ gives for the integral

$$\begin{aligned}
 & B(\boldsymbol{\alpha})p(\mathbf{r} \mid \boldsymbol{\alpha}) \\
 &= \int \frac{d\mathbf{q}}{\prod_{k=1}^D q_k} \exp\left(\sum_{k=1}^{D-1} (n_k + \alpha_k) \log \frac{q_k}{q_D} + \left(n + \sum_{k=1}^D \alpha_k\right) \log q_D\right) \\
 &= \int \frac{d\mathbf{q}}{\prod_{k=1}^D q_k} \exp\left(\sum_{k=1}^{D-1} (n_k + \alpha_k) \log q_k + (n_D + \alpha_D) \log q_D\right) \\
 &= \int d\mathbf{q} \prod_{k=1}^D q_k^{n_k + \alpha_k - 1} = B(\mathbf{n} + \boldsymbol{\alpha}), \tag{68}
 \end{aligned}$$

by definition of the multivariate beta function. We shortened $n_k(\mathbf{r})$ to n_k here.

Proof of Proposition 1

By definition of f we have

$$-\log f(\tilde{\mathbf{q}}, \tilde{n}, \boldsymbol{\theta}) = -\log Z(\tilde{n}, \tilde{\mathbf{q}}) - \{\tilde{n}(\boldsymbol{\theta}\tilde{\boldsymbol{\eta}} - \boldsymbol{\psi}(\boldsymbol{\theta}))\}. \tag{69}$$

Using (24), the negative curly brackets can be replaced by $\tilde{n}(D_\phi - \phi)$, so we obtain

$$-\log f(\tilde{\mathbf{q}}, \tilde{n}, \boldsymbol{\theta}) = -\log Z(\tilde{n}, \tilde{\mathbf{q}}) + \tilde{n}(D(\tilde{\mathbf{q}} \parallel \mathbf{q}) - \phi(\tilde{\boldsymbol{\eta}})). \tag{70}$$

Rearranging terms, we obtain the proposition:

$$\tilde{n}D(\tilde{\mathbf{q}} \parallel \mathbf{q}) = \tilde{n}\phi(\tilde{\boldsymbol{\eta}}) + \log Z(\tilde{n}, \tilde{\mathbf{q}}) - \log f(\tilde{\mathbf{q}}, \tilde{n}, \boldsymbol{\theta}). \tag{71}$$

Proof of Proposition 2

(i) Using the bias-variance decomposition as in (38), for the quadratic risk of $\tilde{\mathbf{q}}$ we obtain

$$\begin{aligned}
 R_q(\tilde{\mathbf{q}}) &= \mathbb{E}(\tilde{\mathbf{q}} - f(\mathbf{q}))^2 \\
 &= \sum_{j=1}^D \text{var}(\lambda f_j(\boldsymbol{\tau}) + (1 - \lambda)f_j(\hat{\mathbf{q}}) - f_j(\mathbf{q})) \\
 &\quad + \sum_{j=1}^D \mathbb{E}^2(\lambda f_j(\boldsymbol{\tau}) + (1 - \lambda)f_j(\hat{\mathbf{q}}) - f_j(\mathbf{q})) \\
 &= \sum_{j=1}^D \left[\lambda^2 \text{var}(f_j(\boldsymbol{\tau})) + (1 - \lambda)^2 \text{var}(f_j(\hat{\mathbf{q}})) + 2\lambda(1 - \lambda) \text{cov}(f_j(\boldsymbol{\tau}), f_j(\hat{\mathbf{q}})) \right] \\
 &\quad + \sum_{j=1}^D \mathbb{E}^2(\lambda(f_j(\boldsymbol{\tau}) - f_j(\hat{\mathbf{q}})) + f_j(\hat{\mathbf{q}}) - f_j(\mathbf{q}))
 \end{aligned}$$

$$= (1 - \lambda)^2 \sum_{j=1}^D \text{var}(f_j(\hat{\mathbf{q}})) + \sum_{j=1}^D \left(\mathbb{E}f_j(\hat{\mathbf{q}}) - f_j(\mathbf{q}) - \lambda(\mathbb{E}f_j(\hat{\mathbf{q}}) - f_j(\boldsymbol{\tau})) \right)^2. \tag{72}$$

For the first equality, the variance of the sum is evaluated in the usual way as a quadratic form. We can ignore the $f_j(\mathbf{q})$ term because it is constant. Similarly, the last equality uses the fact that the $f_j(\boldsymbol{\tau})$ are fixed parameters, so their variance and covariance terms vanish, showing the first part of the proposition.

(ii) To obtain the minimum of the cost function, we derive by λ and set the result zero (while the second derivative is always greater 0):

$$\begin{aligned} \frac{dR_{\mathbf{q}}(\tilde{\mathbf{q}})}{d\lambda} &= -2(1 - \lambda) \sum_{j=1}^D \text{var}(f_j(\hat{\mathbf{q}})) \\ -2 \sum_{j=1}^D (\mathbb{E}f_j(\hat{\mathbf{q}}) - f_j(\boldsymbol{\tau})) &\left(\mathbb{E}f_j(\hat{\mathbf{q}}) - f_j(\mathbf{q}) - \lambda(\mathbb{E}f_j(\hat{\mathbf{q}}) - f_j(\boldsymbol{\tau})) \right) = 0. \end{aligned} \tag{73}$$

From this it follows that

$$\begin{aligned} \sum_{j=1}^D \text{var}(f_j(\hat{\mathbf{q}})) + \sum_{j=1}^D (\mathbb{E}f_j(\hat{\mathbf{q}}) - f_j(\boldsymbol{\tau}))(\mathbb{E}f_j(\hat{\mathbf{q}}) - f_j(\mathbf{q})) \\ = \lambda \sum_{j=1}^D \text{var}(f_j(\hat{\mathbf{q}})) + \lambda \sum_{j=1}^D (\mathbb{E}f_j(\hat{\mathbf{q}}) - f_j(\boldsymbol{\tau}))^2. \end{aligned} \tag{74}$$

Finally, using the fact that $\text{var}(f_j(\hat{\mathbf{q}})) + (\mathbb{E}f_j(\hat{\mathbf{q}}) - f_j(\boldsymbol{\tau}))^2 = \mathbb{E}[(f_j(\hat{\mathbf{q}}) - f_j(\boldsymbol{\tau}))^2]$, we obtain (ii), concluding the proof.

Expectation and variance of the CLR-transformed empirical estimator

Consider the Taylor expansion of the j -th component of $\text{clr}(\hat{\mathbf{q}})$ around \mathbf{q} up to second order terms

$$\begin{aligned} \text{clr}_j(\hat{\mathbf{q}}) &\approx \text{clr}_j(\mathbf{q}) \\ &+ \sum_{k=1}^D \frac{\partial \text{clr}_j(\mathbf{q})}{\partial q_k} (\hat{q}_k - q_k) + \frac{1}{2} \sum_{k,l} \frac{\partial^2 \text{clr}_j(\mathbf{q})}{\partial q_k \partial q_l} (\hat{q}_k - q_k)(\hat{q}_l - q_l). \end{aligned} \tag{75}$$

The first derivatives evaluate to

$$\frac{\partial \text{clr}_j(\mathbf{q})}{\partial q_k} = \begin{cases} \frac{1-1/D}{q_j} & \text{if } j = k, \\ \frac{1}{Dq_k} & \text{if } j \neq k, \end{cases} \tag{76}$$

and the second derivatives are

$$\frac{\partial^2 \text{clr}_j(\mathbf{q})}{\partial q_k \partial q_l} = \begin{cases} -\frac{1-1/D}{q_j^2} & \text{if } j = k = l, \\ \frac{1}{Dq_k^2} & \text{if } j \neq k = l, \\ 0 & \text{else.} \end{cases} \tag{77}$$

When taking the expectation of (75), the first-order terms vanish due to the linearity of expectation. In the second-order terms, only those where $k = l$ remain. We thus obtain

$$\mathbb{E} \text{clr}_j(\hat{\mathbf{q}}) \approx \text{clr}_j(\mathbf{q}) - \frac{\mathbb{E}(\hat{q}_j - q_j)^2}{2q_j^2} + \frac{1}{2D} \sum_{k=1}^D \frac{\mathbb{E}(\hat{q}_k - q_k)^2}{q_k^2}. \tag{78}$$

Now using the bias-variance decomposition (38), we have

$$\mathbb{E}(\hat{q}_j - q_j)^2 = \text{var}(\hat{q}_j) = \frac{1}{n^2} \text{var}(n_j) = \frac{q_j(1 - q_j)}{n}. \tag{79}$$

Inserting this into (78), we obtain (63). Similarly, for the variance of the clr-transformed empirical estimator, we evaluate the variance of (75). The 0-th order does not contribute because it is non-stochastic, and we ignore the second-order terms as commonly done using the Delta method. The variance V_j of the first order terms evaluates to

$$V_j = \text{var} \left(\sum_{k=1}^D \frac{\partial \text{clr}_j(\mathbf{q})}{\partial q_k} \hat{q}_k \right) = \sum_{k,l} \left(\frac{\partial \text{clr}_j(\mathbf{q})}{\partial q_k} \right) \left(\frac{\partial \text{clr}_j(\mathbf{q})}{\partial q_l} \right) \text{cov}(\hat{q}_k, \hat{q}_l), \tag{80}$$

by evaluating the square and using bilinearity of covariance. The covariance elements for equal indices are given in (79). The off-diagonal terms are

$$\text{cov}(\hat{q}_k, \hat{q}_l) = \text{cov} \left(\frac{n_k}{n}, \frac{n_l}{n} \right) = \frac{1}{n^2} \text{cov}(n_k, n_l) \stackrel{k \neq l}{=} \frac{-q_k q_l}{n}, \tag{81}$$

by the well-known expression in the multinomial case. We now collect the respective covariance terms and first derivatives to evaluate (80). The double sum decomposes into four terms that correspond to the cases where the indices are not equal and do not contain j , are equal and don't contain j , are not equal and one of them is j , and are both equal to j , respectively:

$$\begin{aligned} V_j &= \sum_{k \neq j} \sum_{\substack{l \neq j, \\ l \neq k}} \frac{-q_k q_l}{D^2 q_k q_l n} \\ &+ \sum_{l \neq j} \frac{q_l(1 - q_l)}{D^2 q_l^2 n} + 2 \sum_{k \neq j} \frac{(1 - 1/D)(-q_k q_j)}{D q_k q_j n} + \frac{(1 - 1/D)^2 q_j(1 - q_j)}{q_j^2 n} \end{aligned}$$

$$\begin{aligned}
 &= \frac{-(D-1)(D-2)}{D^2n} + \sum_{l \neq j} \frac{1-q_l}{D^2q_l n} - 2 \frac{(1-1/D)(D-1)}{Dn} \\
 &\quad + \frac{(1-1/D)^2(1-q_j)}{q_j n}. \tag{82}
 \end{aligned}$$

This can be further simplified including a part of the last term in the summation of the second term (getting rid of $l \neq j$) and joining the two terms independent of q into a single expression. After this we obtain

$$\begin{aligned}
 V_j &= (1-2/D) \frac{1-q_j}{q_j n} + \frac{1}{D^2} \sum_{l=1}^D \frac{1-q_l}{q_l n} \\
 &\quad - \frac{1}{n} \left(\frac{(D-1)(D-2)}{D^2} + 2 \frac{(1-1/D)(D-1)}{D} \right). \tag{83}
 \end{aligned}$$

With further simplification of the last term, this is (64).

Acknowledgements I thank Nihat Ay for helpful comments on an early version of the manuscript. An anonymous reviewer’s suggestions led to further improvements.

Data availability R scripts for synthetic data and analysis can be downloaded from the GitHub repository [ionase/exshrink](https://github.com/nihaia/exshrink).

Declaration

Conflict of interest I declare that there is no conflict of interest.

References

1. Greenacre, M.: Compositional data analysis. *Annu. Rev. Stat. Appl.* **8**(1), 271–299 (2021)
2. Aitchison, J.: *The Statistical Analysis of Compositional Data*. Chapman and Hall, London (1986)
3. Egozcue, J.J., Pawłowsky-Glahn, V.: Compositional data: the sample space and its structure. *TEST* **28**(3), 599–638 (2019)
4. Erb, I., Gloor, G.B., Quinn, T.P.: Editorial: Compositional data analysis and related methods applied to genomics—a first special issue from NAR Genomics and Bioinformatics. *NAR Genom Bioinform* **2**(4), lqaa103 (2020)
5. Amari, S.: *Information Geometry and Its Applications*. Applied Mathematical Sciences, vol. 194. Springer, Berlin (2016)
6. Erb, I., Ay, N.: The information-geometric perspective of compositional data analysis. In: Filzmoser, P., Hron, K., Martín-Fernández, J.A., Palarea-Albaladejo, J. (eds.) *Advances in Compositional Data Analysis*, pp. 21–43. Springer, New York (2021)
7. Greenacre, M.: Log-ratio analysis is a limiting case of correspondence analysis. *Math. Geosci.* **42**, 129 (2010)
8. Ledoit, O., Wolf, M.: Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *J. Empir. Financ.* **10**, 603–621 (2003)
9. Hausser, J., Strimmer, K.: Entropy inference and the James–Stein estimator, with application to non-linear gene association networks. *J. Mach. Learn. Res.* **10**, 1469–1484 (2009)
10. Quinn, T.P., Erb, I., Richardson, M.F., Crowley, T.M.: Understanding sequencing data as compositions: an outlook and review. *Bioinformatics* **34**(16), 2870–2878 (2018)
11. Jeganathan, P., Holmes, S.P.: A statistical perspective on the challenges in molecular microbial biology. *J. Agric. Biol. Environ. Stat.* **26**, 131–160 (2021)

12. Breda, J., Zavolan, M., van Nimwegen, E.: Bayesian inference of gene expression states from single-cell RNA-seq data. *Nat. Biotechnol.* **39**, 1008–1016 (2021)
13. Robinson, M.D., Oshlack, A.: A scaling normalization method for differential expression analysis of RNA-seq data. *Genome Biol.* **11**, R25 (2010)
14. Lovén, J., Orlando, D.A., Sigova, A.A., Lin, C.Y., Rahl, P.B., Burge, C.B., Levens, D.L., Lee, T.I., Young, R.A.: Revisiting global gene expression analysis. *Cell* **151**, 476–482 (2012)
15. Townes, F.W., Hicks, S.C., Aryee, M.J., Irizarry, R.A.: Feature selection and dimension reduction for single-cell RNA-Seq based on a multinomial model. *Genome Biol.* **20**, 295 (2019)
16. de Finetti, B.: *Theory of Probability, A critical Introductory Treatment*. Wiley, Oxford (2017)
17. Billheimer, D., Guttorp, P., Fagan, W.F.: Statistical interpretation of species composition. *J. Am. Stat. Assoc.* **96**, 1205–1214 (2001)
18. Xia, F., Chen, J., Fung, W.K., Li, H.: A logistic normal multinomial regression model for microbiome compositional data analysis. *Biometrics* **69**, 1053–1063 (2013)
19. McGregor, K., Labbe, A., Greenwood, C.M.T.: MDiNE: a model to estimate differential co-occurrence networks in microbiome studies. *Bioinformatics* **36**(6), 1840–1847 (2020)
20. Avalos, M., Nock, R., Ong, C. S., Rouar, J., Sun, K.: Representation learning of compositional data. *Adv. Neural Inf. Process. Syst.* **31** (2018)
21. Gzyl, H., Nielsen, F.: Geometry of the probability simplex and its connection to the maximum entropy method. *J. Appl. Math. Stat. Inform.* **16**(1), 25–35 (2020)
22. Ay, N., Jost, J., Le, H.V., Schwachhöfer, L.: *Information Geometry. A Series of Modern Surveys in Mathematics*, vol. 64. Springer, Berlin (2017)
23. Cover, T.M., Thomas, J.A.: *Elements of Information Theory*. Wiley, Oxford (2006)
24. Diaconis, P., Ylvisaker, D.: Conjugate priors for exponential families. *Ann. Stat.* **7**(2), 269–281 (1979)
25. Agresti, A., Hitchcock, D.B.: Bayesian inference for categorical data analysis. *Stat. Methods Appl.* **14**, 297–330 (2005)
26. Agarwal, A., Daumé, I.I.H.: A geometric view of conjugate priors. *Mach. Learn.* **81**, 99–113 (2010)
27. Berger, J.O.: *Statistical Decision Theory and Bayesian Analysis*. Springer, Berlin (1985)
28. Johnson, B.M.: On the admissible estimators for certain fixed sample binomial problems. *Ann. Math. Stat.* **42**(5), 1579–1587 (1971)
29. Stein, C.: Inadmissibility of the usual estimator for the mean of a multivariate distribution. In: *Proc. Third Berkeley Symp. Math. Statist. Probab.*, vol. 1. Univ. California Press, pp. 197–206 (1956)
30. James, W, Stein, C.: Estimation with quadratic loss. In: *Proc. Fourth Berkeley Symp. Math. Statist. Probab.*, vol. 1. Univ. California Press, pp. 361–379 (1961)
31. Efron, B., Morris, C.: Stein’s estimation rule and its competitors—an empirical Bayes approach. *J. Am. Stat. Assoc.* **68**(341), 117–130 (1973)
32. Schäfer, J., Strimmer, K.: A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics. *Stat. Appl. Genet. Mol. Biol.* **4**(1), 32 (2005)
33. Greenacre, M.: Power transformations in correspondence analysis. *Comput. Stat. Data Anal.* **53**(8), 3107–3116 (2009)
34. Greenacre, M.: ‘Size’ and ‘shape’ in the measurement of multivariate proximity. *Methods Ecol. Evol.* **8**(11), 1415–1424 (2017)
35. Greenacre, M.: *Biplots in Practice*. Fundación BBVA (2010)
36. Box, G.E.P., Cox, D.R.: An analysis of transformations. *J. R. Stat. Soc. B* **26**(2), 211–252 (1964)
37. Greenacre, M., Grunsky, E., Bacon-Shone, J., Erb, I., Quinn, T.: Aitchison’s Compositional Data Analysis 40 years On: A Reappraisal. *Stat. Sci. Advance Publication* 1–25 (2023). <https://doi.org/10.1214/22-STS880>
38. Boeshaghi, A.S., Hallgrímisdóttir, I.B., Gálvez-Merchán, A., Pachter, L.: Depth normalization for single-cell genomics count data. *bioRxiv* 2022.05.06.490859 (2022)

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.