



CARE: Large Precision Matrix Estimation for Compositional Data

Shucong Zhang, Huiyuan Wang & Wei Lin

To cite this article: Shucong Zhang, Huiyuan Wang & Wei Lin (29 Apr 2024): CARE: Large Precision Matrix Estimation for Compositional Data, Journal of the American Statistical Association, DOI: [10.1080/01621459.2024.2335586](https://doi.org/10.1080/01621459.2024.2335586)

To link to this article: <https://doi.org/10.1080/01621459.2024.2335586>



View supplementary material [↗](#)



Published online: 29 Apr 2024.



Submit your article to this journal [↗](#)



Article views: 647





View related articles [↗](#)



View Crossmark data [↗](#)



CARE: Large Precision Matrix Estimation for Compositional Data

Shucong Zhang^a, Huiyuan Wang^b , and Wei Lin^c 

^aSchool of Statistics, University of International Business and Economics, Beijing, China; ^bDepartment of Biostatistics, Epidemiology and Informatics, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA; ^cSchool of Mathematical Sciences and Center for Statistical Science, Peking University, Beijing, China

ABSTRACT

High-dimensional compositional data are prevalent in many applications. The simplex constraint poses intrinsic challenges to inferring the conditional dependence relationships among the components forming a composition, as encoded by a large precision matrix. We introduce a precise specification of the compositional precision matrix and relate it to its basis counterpart, which is shown to be asymptotically identifiable under suitable sparsity assumptions. By exploiting this connection, we propose a composition adaptive regularized estimation (CARE) method for estimating the sparse basis precision matrix. We derive rates of convergence for the estimator and provide theoretical guarantees on support recovery and data-driven parameter tuning. Our theory reveals an intriguing tradeoff between identification and estimation, thereby highlighting the blessing of dimensionality in compositional data analysis. In particular, in sufficiently high dimensions, the CARE estimator achieves minimax optimality and performs as well as if the basis were observed. We further discuss how our framework can be extended to handle data containing zeros, including sampling zeros and structural zeros. The advantages of CARE over existing methods are illustrated by simulation studies and an application to inferring microbial ecological networks in the human gut. Supplementary materials for this article are available online, including a standardized description of the materials available for reproducing the work.

ARTICLE HISTORY

Received March 2021
Accepted March 2024

KEYWORDS

Blessing of dimensionality;
Graphical modeling;
High-dimensional data;
Identifiability; Microbiome

1. Introduction

High-dimensional compositional data, which lie in a high-dimensional unit simplex, arise in a wide spectrum of contemporary applications. For example, in metagenomic studies, the composition of microbial communities can be comprehensively quantified using 16S rRNA gene or shotgun metagenomic sequencing, which yields the relative abundances of thousands of bacterial taxa (Li 2015). Other prominent examples include the chemical composition of particulate matter in environmental studies (van Donkelaar et al. 2019), market shares of industries in economics (Sutton 2007), and word frequencies of texts in machine learning (Blei and Lafferty 2007). A central task in such applications is to infer the conditional dependence relationships among the components forming a composition, which are often represented by an undirected graphical model and encoded by the inverse covariance or precision matrix (Lauritzen 1996).

The simplex constraint creates intrinsic challenges to inference for high-dimensional graphical models and precision matrices. A conventional method starts with applying the log-ratio transformation (Aitchison 1982, 2003) to the compositional data, and proceeds with inference as if the transformed data lie in a Euclidean space. This approach, however, is unsatisfactory for our problem in the following respects.

1. The simplex constraint induces negative correlations between originally independent components, and hence a “null” structure of the compositional precision matrix away from a diagonal

matrix. It is therefore not clear how to impose the sparsity structure as in the case of Euclidean data.

2. Among the conventional matrix specifications of compositional covariance structures (Aitchison 2003), the log-ratio and isometric log-ratio covariance matrices treat the components asymmetrically, while the variation matrix is not a covariance matrix and the centered log-ratio covariance matrix is not invertible. Thus, none of these specifications give rise to a compositional precision matrix that is permutation invariant under sparsity assumptions.
3. Since each log-ratio involves more than one component, conditional dependence relationships among the log-ratios are more difficult to interpret, especially when model selection or support recovery is the primary purpose.

A potential way to resolve these difficulties is to return to the basis precision matrix Ω_0 , that is, the precision matrix of the log-basis variables that generate a composition; see Section 2 for the formal definition. Since the log-basis variables are free of constraints, Ω_0 can be assumed to be sparse and interpreted in the usual manner. However, owing to the many-to-one relationship between the basis and the composition, Ω_0 is unidentifiable from compositional data in any finite dimension, making it a seemingly unrealistic goal.

A closely related question was addressed by Cao, Lin, and Li (2019). They showed that the basis covariance matrix $\Sigma_0 = \Omega_0^{-1}$ belonging to a suitable sparsity class is approximately identifiable and indistinguishable from the centered log-ratio covariance matrix Σ_c as the dimensionality increases. Their idea relies on a

fundamental decomposition, originally due to Aitchison (2003), that relates the basis and compositional covariance structures. On the other hand, such a relationship for the precision matrices has long been lacking in the literature, preventing the problem of precision matrix estimation from being solved. Any attempt to apply Cao, Lin, and Li's (2019) argument directly to the precision matrix setting would face the necessity of imposing sparsity assumptions simultaneously on the covariance and precision matrices, which is unreasonable and impractical.

This article bridges the conceptual gap by introducing a precise specification of the compositional precision matrix and establishing a transparent relationship between the basis and compositional precision matrices. This novel connection inherits a similar low-rank plus sparse form from the covariance decomposition mentioned above, allowing us to prove the approximate identifiability of Ω_0 under suitable sparsity assumptions. These results are of independent interest and would be useful for other inference problems for compositional data where precision matrices play a role. Building on our precision matrix specification and its nice properties, we adopt the idea of constrained ℓ_1 -minimization (Cai, Liu, and Luo 2011) to develop a composition adaptive regularized estimation (CARE) method for estimating the sparse basis precision matrix. We further derive rates of convergence for the resulting estimator under various matrix norms, and provide theoretical guarantees on support recovery and data-driven parameter tuning.

Our theoretical results reveal an intriguing tradeoff between identification and estimation. More specifically, as the dimension p grows, the identification error decreases, while the estimation error increases. Remarkably, when $p \log p \gg n$ with n being the sample size, the estimation error dominates the identification error and the CARE estimator achieves the minimax optimal rates for unconstrained data in Cai, Liu, and Zhou (2016). This entails that, in sufficiently high dimensions, our method performs as well as if the basis were observed. This blessing of dimensionality has direct implications for practical data analysis. For example, in microbiome studies, the number of bacterial taxa depends on the taxonomic level under examination, which should be judiciously chosen to facilitate inference and interpretation.

1.1. Related Work

There is a vast and rapidly growing literature on high-dimensional graphical models and precision matrix estimation. The existing methods fall roughly into three categories. The neighborhood-based approach exploits a nodewise regression formulation of Gaussian graphical models, and fits the high-dimensional regressions using regularization methods such as the lasso, Dantzig selector, and scaled lasso (Meinshausen and Bühlmann 2006; Yuan 2010; Ren et al. 2015; Fan and Lv 2016). Instead, the likelihood-based approach estimates the precision matrix as a whole by directly penalizing the joint Gaussian likelihood with the graphical lasso (Yuan and Lin 2007; Friedman, Hastie, and Tibshirani 2008) or nonconvex penalties (Lam and Fan 2009). Finally, the CLIME-related methods are motivated by modified score functions and solve constrained ℓ_1 -minimization or penalized empirical risk minimization problems (Cai, Liu, and Luo 2011; Zhang and Zou 2014; Liu and

Luo 2015; Cai, Liu, and Zhou 2016). Their ideas are to simplify the loss function and yield theoretical and computational advantages.

Extensions of the problem to compositional data have attracted considerable attention in the computational biology literature. The SPIEC-EASI procedure proposed by Kurtz et al. (2015) treats the centered log-ratio transformed data as Euclidean data and directly applies the usual neighborhood selection and graphical lasso methods. Fang et al. (2017) developed an ℓ_1 -penalized likelihood method called gCoda by parameterizing the logistic normal likelihood with Ω_0 . The identifiability issue under the parameterization was noted but not rigorously addressed. A related method using the Bayesian graphical lasso was considered by Schwager et al. (2017). Yuan, He, and Deng (2019) introduced an ℓ_1 -penalized CD-trace approach by borrowing the idea of D-trace loss from Zhang and Zou (2014). The derivation, however, relies on an exchangeability condition, which is equivalent to the row sums of Σ_0 being all equal. Zhang and He (2019) employed the sparse columnwise inverse operator (SCIO) method of Liu and Luo (2015) to find an approximate inverse of the sample centered log-ratio covariance matrix. Their method is based on the same heuristic argument as that of SPIEC-EASI, and would suffer from similar conceptual issues.

The difficulties and ambiguities in deriving and lending support to the aforementioned methods showcase the lack of convenient mathematical tools for the study of compositional precision matrices. As commented by Aitchison (2003, p. 64), before introducing the tools of compositional covariance matrices,

we would not expect that excellent tool of the wide open spaces (or \mathbb{R}^d) of North America, namely the barbecue, necessarily to be an appropriate concept for cooking in the confined space (or \mathcal{S}^d) of a low-cost housing flatlet in Hong Kong. If our concepts fail to serve us in new situations we must invent new concepts.

Similarly, we would not expect the oven for roasting Cantonese ducks (or compositional covariance matrices) to be perfect for roasting Peking ducks (or compositional precision matrices). We shall therefore suggest a new concept, whose mathematical form was first defined but not explored for inferential use by Aitchison (2003), along with the associated tools that are more convenient and powerful for our purpose.

1.2. Organization of the Article

In Section 2, we present the precision matrix specification and its properties, thereby addressing the identifiability problem. Section 3 motivates and describes the CARE methodology. Section 4 establishes theoretical guarantees for our proposed estimator in terms of rates of convergence, support recovery, and data-driven choice of tuning parameters. Extensions to data containing zeros and a perturbation theory are developed in Section 5. Simulation studies and an application to human gut microbiome data are presented in Sections 6 and 7, respectively. Section 8 contains some final discussion. Proofs and additional discussion and numerical results are provided in the supplementary materials.

We close this section by introducing some notation to be used throughout the article. For any vector \mathbf{a} , matrix \mathbf{A} , and $1 \leq r \leq \infty$, we denote by $\|\mathbf{a}\|_r$ the vector ℓ_r -norm, and $\|\mathbf{A}\|_{\max}$, $\|\mathbf{A}\|_{L_r}$, and $\|\mathbf{A}\|_F$ the entrywise ℓ_∞ -norm, matrix ℓ_r -norm, and Frobenius norm, respectively. Also, $\lambda_{\min}(\mathbf{A})$ and $\lambda_{\max}(\mathbf{A})$ denote the smallest and largest eigenvalues of \mathbf{A} , respectively, and $\mathbf{A} \succ 0$ stands for \mathbf{A} being symmetric and positive definite. Finally, we write \mathbf{I}_p for the $p \times p$ identity matrix, $\mathbf{1}_p$ the p -vector of 1s, and \mathbf{e}_j the vector with 1 in the j th component and 0 elsewhere.

2. Precision Matrix Specification and Identifiability

Consider a composition $\mathbf{X} = (X_1, \dots, X_p)^T$ that takes values in the $(p-1)$ -dimensional unit simplex $S^{p-1} = \{(x_1, \dots, x_p) : x_j > 0 \text{ for } j = 1, \dots, p, \sum_{j=1}^p x_j = 1\}$. Let $\mathbf{W} = (W_1, \dots, W_p)^T$ with $W_j > 0$ for all j be a latent vector, called the basis, that generates the observed composition via the normalization

$$X_j = \frac{W_j}{\sum_{i=1}^p W_i}, \quad j = 1, \dots, p.$$

To specify a basis covariance structure compatible with the log-ratio form of compositional covariance structure, we define the log-basis vector $\mathbf{Y} = (Y_1, \dots, Y_p)^T$ with $Y_j = \log W_j$ for all j . Denote by $\mathbf{\Omega}_0 = (\omega_{ij}^0)$ the precision matrix of \mathbf{Y} , which we call the basis precision matrix. In the important case where \mathbf{Y} is multivariate normal, so that \mathbf{X} is logistic normal (Aitchison 2003, Property 6.1), it follows from nonparanormal graphical model theory (Liu, Lafferty, and Wasserman 2009) that W_i and W_j are conditionally independent given the other $p-2$ variables if and only if $\omega_{ij}^0 = 0$. The graphical model interpretation of the precision matrix has also been extended to certain discrete and non-Gaussian cases (Loh and Wainwright 2013; Morrison, Baptista, and Basor 2022). Hence, recovering the graphical model structure among the p basis variables amounts to estimating the support of $\mathbf{\Omega}_0$. In microbiome studies, \mathbf{X} represents the relative abundances of bacterial taxa measured by 16S rRNA gene or shotgun metagenomic sequencing, while \mathbf{W} represents the absolute abundances that are of primary interest but not directly observed. It has been found that sample abundance measurements after accounting for undersampling are well approximated by a log-normal distribution (Limpert, Stahel, and Abbt 2001; Paulson et al. 2013), and thus the nonparanormal graphical model seems plausible. In such applications, our goal is to estimate the basis precision matrix $\mathbf{\Omega}_0$ from the observed data on \mathbf{X} .

Denote by $\mathbf{\Sigma}_0 = (\sigma_{ij}^0) = \mathbf{\Omega}_0^{-1}$ the basis covariance matrix, that is, the covariance matrix of \mathbf{Y} . The centered log-ratio transformation (Aitchison 2003) gives the transformed data $\mathbf{Z} = (Z_1, \dots, Z_p)^T$ with $Z_j = \log\{X_j/g(\mathbf{X})\}$, where $g(\mathbf{X}) = (\prod_{j=1}^p X_j)^{1/p}$ is the geometric mean of \mathbf{X} . Accordingly, the centered log-ratio covariance matrix $\mathbf{\Sigma}_c = (\sigma_{ij}^c)$ is the covariance matrix of \mathbf{Z} . Although $\mathbf{\Sigma}_0$ is generally unidentifiable, Cao, Lin, and Li (2019) showed that it is close to its compositional counterpart $\mathbf{\Sigma}_c$ in the sense that

$$\|\mathbf{\Sigma}_0 - \mathbf{\Sigma}_c\|_{\max} = O\left(\frac{\|\mathbf{\Sigma}_0\|_{L_1}}{p}\right).$$

This identifiability bound derives from the representation of $\mathbf{\Sigma}_c$ in terms of $\mathbf{\Sigma}_0$:

$$\mathbf{\Sigma}_c = \mathbf{\Sigma}_0 - \mathbf{v}_p \mathbf{v}_p^T \mathbf{\Sigma}_0 - \mathbf{\Sigma}_0 \mathbf{v}_p \mathbf{v}_p^T + \mathbf{v}_p \mathbf{v}_p^T \mathbf{\Sigma}_0 \mathbf{v}_p \mathbf{v}_p^T, \quad (1)$$

where $\mathbf{v}_p = \mathbf{1}_p / \sqrt{p}$. As a result, under sparsity assumptions on $\mathbf{\Sigma}_0$ such that $\|\mathbf{\Sigma}_0\|_{L_1} = o(p)$, $\mathbf{\Sigma}_0$ is asymptotically identifiable and indistinguishable with $\mathbf{\Sigma}_c$ as $p \rightarrow \infty$. This fact entails that $\mathbf{\Sigma}_c$ serves as a good proxy for $\mathbf{\Sigma}_0$ in sparse covariance estimation. Intuitively, we would expect a well-defined inverse of $\mathbf{\Sigma}_c$ to be a good approximation of $\mathbf{\Omega}_0$. However, owing to the zero-sum constraint on \mathbf{Z} , $\mathbf{\Sigma}_c$ has an eigenvalue 0 and the corresponding eigenvector \mathbf{v}_p (Aitchison 2003, Theorem 4.6).

A conceptually simple remedy for the singularity of $\mathbf{\Sigma}_c$ is to find a generalized or pseudo-inverse. Let $\mathbf{\Sigma}_c = \mathbf{V} \mathbf{D}_c \mathbf{V}^T$ be the spectral decomposition of $\mathbf{\Sigma}_c$, where \mathbf{V} is a $p \times (p-1)$ matrix of eigenvectors and $\mathbf{D}_c = \text{diag}(d_1, \dots, d_{p-1})$ with $d_1 \geq \dots \geq d_{p-1} > 0$. The rank-deficient matrix $\mathbf{\Sigma}_c$ can be made invertible by a rank-1 correction. More precisely, for any constant $\rho > 0$, $\mathbf{\Sigma}_c + \rho \mathbf{v}_p \mathbf{v}_p^T$ is nonsingular and has the spectral decomposition

$$\mathbf{\Sigma}_c + \rho \mathbf{v}_p \mathbf{v}_p^T = (\mathbf{V}, \mathbf{v}_p) \begin{pmatrix} \mathbf{D}_c & \mathbf{0} \\ \mathbf{0} & \rho \end{pmatrix} (\mathbf{V}, \mathbf{v}_p)^T,$$

whose inverse is given by

$$(\mathbf{\Sigma}_c + \rho \mathbf{v}_p \mathbf{v}_p^T)^{-1} = (\mathbf{V}, \mathbf{v}_p) \begin{pmatrix} \mathbf{D}_c^{-1} & \mathbf{0} \\ \mathbf{0} & \rho^{-1} \end{pmatrix} (\mathbf{V}, \mathbf{v}_p)^T.$$

The Moore–Penrose inverse of $\mathbf{\Sigma}_c$ can then be expressed as

$$\mathbf{\Sigma}_c^+ = \mathbf{V} \mathbf{D}_c^{-1} \mathbf{V}^T = (\mathbf{\Sigma}_c + \rho \mathbf{v}_p \mathbf{v}_p^T)^{-1} - \frac{1}{\rho} \mathbf{v}_p \mathbf{v}_p^T. \quad (2)$$

We now define $\mathbf{\Omega}_c = (\omega_{ij}^c) = \mathbf{\Sigma}_c^+$ as the *compositional precision matrix*, since it provides a full parameterization of compositional inverse covariance structure and is identifiable from the observed compositional data. The definition of $\mathbf{\Sigma}_c^+$ and expression (2) were mentioned by Aitchison (2003, Sec. 5.6), but its use for modeling and inference was not pursued therein. Indeed, since $\mathbf{\Omega}_c$ is indirectly defined through $\mathbf{\Sigma}_c$, it is less interpretable and may not be very useful without an explicit link to the basis precision matrix $\mathbf{\Omega}_0$. Fortunately, by substituting (1) into (2) and using the Sherman–Morrison formula, we can establish the following relationship between $\mathbf{\Omega}_c$ and $\mathbf{\Omega}_0$.

Theorem 1. The following relationship between $\mathbf{\Omega}_c$ and $\mathbf{\Omega}_0$ holds:

$$\mathbf{\Omega}_c = \mathbf{\Omega}_0 - \frac{\mathbf{\Omega}_0 \mathbf{1}_p \mathbf{1}_p^T \mathbf{\Omega}_0}{\mathbf{1}_p^T \mathbf{\Omega}_0 \mathbf{1}_p}.$$

While there can be many possible ways to specify the compositional precision matrix, Theorem 1 illustrates a major advantage of our particular specification: it admits a representation as the sum of its basis counterpart $\mathbf{\Omega}_0$ and a low-rank correction, which inherits a similar form from the decomposition of $\mathbf{\Sigma}_c$ in (1). Here, the rank-1 component reflects a global effect of compositionality across all rows and columns of $\mathbf{\Omega}_0$, and is related to the information loss due to the normalization operation. In light of this decomposition, statistical wisdom learned from the literature on low-rank plus sparse matrix recovery problems

(Candès et al. 2011; Chandrasekaran, Parrilo, and Willsky 2012; Fan, Liao, and Mincheva 2013) suggests that a suitably sparse $\mathbf{\Omega}_0$ can be separated from the compositionality effect. The following proposition collects some useful properties of $\mathbf{\Omega}_c$.

Proposition 1. The matrix $\mathbf{\Omega}_c$ satisfies the following properties:

- (a) $\mathbf{\Sigma}_c \mathbf{\Omega}_c = \mathbf{G}$, where $\mathbf{G} = \mathbf{I}_p - p^{-1} \mathbf{1}_p \mathbf{1}_p^T$;
- (b) $\mathbf{G} \mathbf{\Omega}_c = \mathbf{\Omega}_c$ and $\mathbf{\Omega}_c \mathbf{1}_p = \mathbf{0}$;
- (c) $\lambda_{\min}(\mathbf{\Omega}_0) \leq \lambda_{\min}^+(\mathbf{\Omega}_c) \leq \lambda_{\max}(\mathbf{\Omega}_c) \leq \lambda_{\max}(\mathbf{\Omega}_0)$, where $\lambda_{\min}^+(\mathbf{\Omega}_c)$ denotes the smallest positive eigenvalue of $\mathbf{\Omega}_c$;
- (d) $\sigma_{jj}^2 \omega_{jj}^c \geq (1 - 1/p)^2$ for $j = 1, \dots, p$.

Part (a) in Proposition 1 indicates that $\mathbf{\Omega}_c$ is an approximate inverse of $\mathbf{\Sigma}_c$ with \mathbf{G} playing the role of the identity matrix for centered log-ratio vectors. In fact, any p -vector \mathbf{a} satisfies $\mathbf{a}^T \mathbf{1}_p = 0$ if and only if $\mathbf{G} \mathbf{a} = \mathbf{a}$ (Aitchison 2003, App. A). Part (b) parallels Theorem 4.6 of Aitchison (2003), part (c) confines the range of the positive eigenvalues of $\mathbf{\Omega}_c$ to within that of $\mathbf{\Omega}_0$, and part (d) is analogous to the inequality that $\sigma_{jj}^0 \omega_{jj}^0 \geq 1$. These properties are essential for the development of methodology and theory based on $\mathbf{\Omega}_c$.

To impose sparsity on the target matrix $\mathbf{\Omega}_0$, we consider the class of sparse precision matrices

$$\mathcal{U}_q(s_0(p), M_p) = \left\{ \mathbf{\Omega} = (\omega_{ij}) : \mathbf{\Omega} \succ 0, \max_{1 \leq j \leq p} \sum_{i=1}^p |\omega_{ij}|^q \leq s_0(p), \|\mathbf{\Omega}\|_{L_1} \leq M_p, 1/R \leq \lambda_{\min}(\mathbf{\Omega}) \leq \lambda_{\max}(\mathbf{\Omega}) \leq R \right\},$$

where $0 \leq q < 1$, $R > 1$ is a constant, and $s_0(p)$ and M_p are bounded away from 0 and may diverge with p . This class covers a wide range of sparse precision matrices and is similar to those considered for unconstrained high-dimensional data (Cai, Liu, and Luo 2011; Ren et al. 2015; Cai, Liu, and Zhou 2016). Under the assumption that $\mathbf{\Omega}_0$ belongs to the class $\mathcal{U}_q(s_0(p), M_p)$, $\mathbf{\Omega}_c$ is generally not sparse, but remains close to $\mathbf{\Omega}_0$ and preserves some of its properties, as formalized by the following results.

Proposition 2. Suppose that $\mathbf{\Omega}_0 \in \mathcal{U}_q(s_0(p), M_p)$. Then

$$\frac{R^{-3}}{p} \leq \|\mathbf{\Omega}_0 - \mathbf{\Omega}_c\|_{\max} \leq \frac{RM_p^2}{p}.$$

Proposition 3. Suppose that $\mathbf{\Omega}_0 \in \mathcal{U}_q(s_0(p), M_p)$. Then

$$\|\mathbf{\Omega}_c\|_{L_1} \leq (1 + R^2)M_p, \quad 1/R \leq \lambda_{\min}^+(\mathbf{\Omega}_c) \leq \lambda_{\max}(\mathbf{\Omega}_c) \leq R.$$

Proposition 2 entails that the sparse basis precision matrix $\mathbf{\Omega}_0$ is approximately identifiable provided that $M_p = o(\sqrt{p})$. The identification error due to approximating $\mathbf{\Omega}_0$ by $\mathbf{\Omega}_c$ vanishes asymptotically as $p \rightarrow \infty$. This blessing of dimensionality motivates us to use $\mathbf{\Omega}_c$ as a proxy for $\mathbf{\Omega}_0$ and develop procedures for estimating $\mathbf{\Omega}_0$ through $\mathbf{\Omega}_c$. In particular, when M_p does not grow with p , the identification error decays at the rate of $O(p^{-1})$, which is sharp and cannot be improved. Proposition 3 gives bounds on the matrix ℓ_1 -norm and positive eigenvalues of $\mathbf{\Omega}_c$, which are analogous to those for $\mathbf{\Omega}_0$. Finally, we observe that if the eigenvalue condition in the definition of $\mathcal{U}_q(s_0(p), M_p)$ is replaced by $\lambda_{\max}(\mathbf{\Omega})/\lambda_{\min}(\mathbf{\Omega}) \leq R$ as in Cai, Liu, and

Zhou (2016), then the $\|\cdot\|_{\max}$ -distance between $\mathbf{\Omega}_0$ and $\mathbf{\Omega}_c$ in Proposition 2 can be similarly bounded as

$$\|\mathbf{\Omega}_0 - \mathbf{\Omega}_c\|_{\max} \leq \frac{R\|\mathbf{\Omega}_0\|_{L_1}}{\sqrt{p}} \leq \frac{RM_p}{\sqrt{p}}.$$

In this case, the approximate identifiability of $\mathbf{\Omega}_0$ still holds as long as $M_p = o(\sqrt{p})$.

3. Composition Adaptive Regularized Estimation

Suppose that $(\mathbf{X}_1, \mathbf{Z}_1), \dots, (\mathbf{X}_n, \mathbf{Z}_n)$ are independent realizations of (\mathbf{X}, \mathbf{Z}) . Denote by $\widehat{\mathbf{\Sigma}}_c = (\hat{\sigma}_{ij}^c) = n^{-1} \sum_{k=1}^n (\mathbf{Z}_k - \bar{\mathbf{Z}})(\mathbf{Z}_k - \bar{\mathbf{Z}})^T$ the sample centered log-ratio covariance matrix, where $\bar{\mathbf{Z}} = n^{-1} \sum_{k=1}^n \mathbf{Z}_k$. In view of Proposition 2, we shall develop an estimator of $\mathbf{\Omega}_0$ via the proxy matrix $\mathbf{\Omega}_c$ based on the observed data or $\widehat{\mathbf{\Sigma}}_c$. However, $\mathbf{\Omega}_c$ is in general not sparse, although it is entrywise close to the sparse matrix $\mathbf{\Omega}_0$. Also, we shall not impose restrictive distributional assumptions. Thus, it is unclear whether existing neighborhood or likelihood-based methods can effectively leverage the proxy $\mathbf{\Omega}_c$. A further look at the CLIME estimator reveals that its rate of convergence under the entrywise ℓ_∞ -norm depends crucially on the matrix ℓ_1 -norm bound M_p (Cai, Liu, and Luo 2011, Theorem 6). This, together with the fact from Proposition 2 that the identification error between $\mathbf{\Omega}_0$ and $\mathbf{\Omega}_c$ can be tightly bounded in terms of M_p , motivates us to consider a CLIME-type estimator for $\mathbf{\Omega}_0$.

Without loss of generality, we assume that the log-basis vector \mathbf{Y} has mean $\mathbf{0}$. Since our inference goal is to estimate the precision matrix $\mathbf{\Omega}_0$ of \mathbf{Y} , it is conventional to impose the following two types of tail conditions on \mathbf{Y} .

Condition 1 (Sub-Gaussian tail). Assume that $\log p = o(n)$ and there exist some constants $\eta > 0$ and $K > 0$ such that

$$\sup_{\mathbf{s} \in \mathbb{R}^p, \|\mathbf{s}\|_2=1} E[\exp\{\eta(\mathbf{s}^T \mathbf{Y})^2 / \text{Var}(\mathbf{s}^T \mathbf{Y})\}] \leq K.$$

Condition 2 (Polynomial-type tail). Assume that $p = O(n^\gamma)$ for some constant $\gamma > 0$ and there exist some constants $\varepsilon > 0$ and $K' > 0$ such that

$$\sup_{\mathbf{s} \in \mathbb{R}^p, \|\mathbf{s}\|_2=1} E|\mathbf{s}^T \mathbf{Y}|^{4\gamma+4+\varepsilon} \leq K'.$$

If the log-basis \mathbf{Y} were observable, then $\widehat{\mathbf{\Sigma}} \mathbf{\Omega}_0$ would concentrate around \mathbf{I}_p , where $\widehat{\mathbf{\Sigma}}$ is the sample basis covariance matrix. In our setting, however, \mathbf{Y} is not observable and our estimator should be based on $\widehat{\mathbf{\Sigma}}_c$. Since $\mathbf{\Omega}_0$ and $\mathbf{\Omega}_c$ are asymptotically indistinguishable when $M_p = o(\sqrt{p})$, by part (a) of Proposition 1 we expect that $\widehat{\mathbf{\Sigma}}_c \mathbf{\Omega}_0$, similar to $\widehat{\mathbf{\Sigma}}_c \mathbf{\Omega}_c$, would concentrate around \mathbf{G} . The following lemma shows that this is indeed the case.

Lemma 1. Suppose that $\mathbf{\Omega}_0 \in \mathcal{U}_q(s_0(p), M_p)$ with $M_p = o(\sqrt{p})$. Under Condition 1 (or Condition 2), for any $\xi > 0$, there exists some constant $C_0 > 0$ such that

$$\|\widehat{\mathbf{\Sigma}}_c \mathbf{\Omega}_0 - \mathbf{G}\|_{\max} \leq C_0 \left(\sqrt{\frac{\log p}{n}} + \frac{M_p}{\sqrt{p}} \right)$$

with probability at least $1 - O(p^{-\xi})$ (or $1 - O(p^{-\xi/2} + n^{-\varepsilon/8})$).

Lemma 1 allows us to construct a constrained ℓ_1 -minimization procedure for estimating $\mathbf{\Omega}_0$ in a similar spirit to the CLIME approach of Cai, Liu, and Luo (2011). More specifically, we consider the following optimization problem:

$$\text{minimize } \|\mathbf{\Omega}\|_1 \quad \text{subject to} \quad \|\widehat{\mathbf{\Sigma}}_c \mathbf{\Omega} - \mathbf{G}\|_{\max} \leq \lambda, \quad (3)$$

where $\lambda > 0$ is a tuning parameter and can be easily chosen to ensure the feasibility of $\mathbf{\Omega}_0$ provided that $M_p = o(\sqrt{p})$. It is important to note that the second term M_p/\sqrt{p} of the error bound in **Lemma 1** is due to the departure of $\mathbf{\Omega}_c$ from $\mathbf{\Omega}_0$. The main differences from the case of unconstrained data are that problem (3) incorporates the identifiability gap between $\mathbf{\Omega}_0$ and $\mathbf{\Omega}_c$, and that the matrix \mathbf{G} plays the role of the identity matrix. The overall optimization problem (3) can be further decomposed into p independent vector minimization problems. As a result, the estimator $\widehat{\mathbf{\Omega}} = (\widehat{\omega}_{ij}) = (\widehat{\omega}_1, \dots, \widehat{\omega}_p)$ is defined through the solutions $\widehat{\omega}_j$ to the optimization problems

$$\text{minimize } \|\omega_j\|_1 \quad \text{subject to} \quad \left\| \widehat{\mathbf{\Sigma}}_c \omega_j - \left(\mathbf{e}_j - \frac{\mathbf{1}_p}{p} \right) \right\|_{\max} \leq \lambda_j, \quad (4)$$

where $\lambda_j > 0$ are tuning parameters. To adapt to the variability and sparsity of the columns of $\mathbf{\Omega}_0$, λ_j are allowed to vary from column to column, thus differing from the CLIME method where the tuning parameter is fixed. In practice, λ_j can be chosen by cross-validation, which will be further discussed in **Section 4.2**. Finally, we define our CARE estimator $\widehat{\mathbf{\Omega}} = (\widehat{\omega}_{ij})$ by symmetrizing $\widehat{\mathbf{\Omega}}$:

$$\widehat{\omega}_{ij} = \widehat{\omega}_{ji} = \widehat{\omega}_{ij} I(|\widehat{\omega}_{ij}| \leq |\widehat{\omega}_{ji}|) + \widehat{\omega}_{ji} I(|\widehat{\omega}_{ij}| > |\widehat{\omega}_{ji}|).$$

It is worth mentioning that our methodology differ from those of Cao, Lin, and Li (2019) and Cai, Liu, and Zhou (2016) in some important ways. First, while the COAT method is optimization-free, our method requires solving an optimization problem, whose dependence on the basis-composition relationship is implicit. This would make the theoretical development technically more challenging. Second, our method consists of a single-step constrained ℓ_1 -minimization procedure with column-specific tuning parameters, which departs markedly from the two-step ACLIME procedure and adapts better to the identification error in our problem. An alternative method that follows the ACLIME approach more closely is discussed in Supplementary Section S.2.

4. Theoretical Properties

We investigate the theoretical properties of the CARE estimator. In **Section 4.1**, we derive rates of convergence under various matrix norms and provide a support recovery guarantee for the estimator $\widehat{\mathbf{\Omega}}$. Data-driven choice of the tuning parameters λ_j is justified in **Section 4.2**. In deriving the rates of convergence, we explicitly characterize the impact of dimensionality on the degree of identifiability and decompose the total error into an estimation error and an identification error. Results of this kind are rare in the literature, but see Fan, Liao, and Mincheva (2013) and Cao, Lin, and Li (2019).

4.1. Rates of Convergence and Support Recovery

We first present the rates of convergence for the estimator $\widehat{\mathbf{\Omega}}$ under different matrix norms.

Theorem 2. Suppose that $\mathbf{\Omega}_0 \in \mathcal{U}_q(s_0(p), M_p)$ with $M_p = o(\sqrt{p})$ and the tuning parameters are chosen as $\lambda_j \asymp \sqrt{(\log p)/n} + M_p/\sqrt{p}$ for all j . Under **Condition 1** (or **Condition 2**), for any $\xi > 0$, there exists some constant $C > 0$ such that the estimator $\widehat{\mathbf{\Omega}}$ satisfies

$$\begin{aligned} \|\widehat{\mathbf{\Omega}} - \mathbf{\Omega}_0\|_{\max} &\leq C \left(M_p \sqrt{\frac{\log p}{n}} + \frac{M_p^2}{\sqrt{p}} \right), \\ \|\widehat{\mathbf{\Omega}} - \mathbf{\Omega}_0\|_{L_1} &\leq C s_0(p) \left(M_p \sqrt{\frac{\log p}{n}} + \frac{M_p^2}{\sqrt{p}} \right)^{1-q}, \\ \frac{1}{p} \|\widehat{\mathbf{\Omega}} - \mathbf{\Omega}_0\|_F^2 &\leq C s_0(p) \left(M_p \sqrt{\frac{\log p}{n}} + \frac{M_p^2}{\sqrt{p}} \right)^{2-q} \end{aligned}$$

with probability at least $1 - O(p^{-\xi})$ (or $1 - O(p^{-\xi/2} + n^{-\varepsilon/8})$).

Note that the error bound under the matrix ℓ_1 -norm also holds under any matrix ℓ_r -norm with $1 \leq r \leq \infty$, since $\widehat{\mathbf{\Omega}}$ and $\mathbf{\Omega}_0$ are symmetric (Cai, Liu, and Zhou 2016, Lemma 7.2). All the error bounds in **Theorem 2** have an appealing form that decomposes into two terms. The first term represents the estimation error due to the estimation of $\mathbf{\Omega}_c$ based on the observed data, while the second term arises from the identification error between $\mathbf{\Omega}_c$ and $\mathbf{\Omega}_0$. Note that the latter is due to the intrinsic difficulty of nonidentifiability in estimating $\mathbf{\Omega}_0$; as a result, it cannot be eliminated by any other methods. Theoretically, this can be seen from **Proposition 2**, where the identification error between $\mathbf{\Omega}_c$ and $\mathbf{\Omega}_0$ is bounded away from zero. Moreover, it is important to note that the dimension p plays opposite roles in these two terms: it contributes a factor of $\log p$ to the former but a factor of $p^{-1/2}$ to the latter. This leads to a clear tradeoff between identification and estimation, or a tradeoff between the blessing and the curse of dimensionality, a striking phenomenon not observed in precision matrix estimation for unconstrained data. The quantity M_p also has a role to play in the tradeoff and has a larger impact on the identification error than on the estimation error. As p becomes sufficiently large, the estimation error will dominate the identification error, as summarized in the following corollary.

Corollary 1. Assume that the conditions of **Theorem 2** hold. If $M_p = o(\sqrt{p(\log p)/n})$, then the estimator $\widehat{\mathbf{\Omega}}$ satisfies

$$\begin{aligned} \|\widehat{\mathbf{\Omega}} - \mathbf{\Omega}_0\|_{\max} &\leq C M_p \sqrt{\frac{\log p}{n}}, \\ \|\widehat{\mathbf{\Omega}} - \mathbf{\Omega}_0\|_{L_1} &\leq C M_p^{1-q} s_0(p) \left(\frac{\log p}{n} \right)^{(1-q)/2}, \\ \frac{1}{p} \|\widehat{\mathbf{\Omega}} - \mathbf{\Omega}_0\|_F^2 &\leq C M_p^{2-q} s_0(p) \left(\frac{\log p}{n} \right)^{1-q/2} \end{aligned}$$

with probability at least $1 - O(p^{-\xi})$ or $1 - O(p^{-\xi/2} + n^{-\varepsilon/8})$.

Corollary 1 requires that $M_p = o(\sqrt{p(\log p)/n})$, which is stronger than the identifiability condition $M_p = o(\sqrt{p})$ and suggests a phase transition phenomenon for optimal estimation.

In particular, the resulting rates of convergence under the matrix ℓ_r -norm with $1 \leq r \leq \infty$ and Frobenius norm match the min-max optimal rates for unconstrained data (Cai, Liu, and Zhou 2016). This indicates that, in sufficiently high dimensions, the CARE estimator performs as well as if the basis were observed, highlighting the blessing of dimensionality in compositional data analysis.

Based on the entrywise ℓ_∞ -error bound in Theorem 2, we define the hard-thresholded estimator $\hat{\Omega}^t = (\hat{\omega}_{ij}^t)$ with $\hat{\omega}_{ij}^t = \hat{\omega}_{ij} I(|\hat{\omega}_{ij}| > \tau_{np})$, where $\tau_{np} = C(M_p \sqrt{(\log p)/n} + M_p^2/\sqrt{p})$. Our next result shows that, under an additional minimum signal assumption, the hard-thresholded estimator $\hat{\Omega}^t$ recovers the support of Ω_0 successfully.

Theorem 3. Assume that the conditions of Theorem 2 hold. If $\min_{(i,j): \omega_{ij}^0 \neq 0} |\omega_{ij}^0| > 2\tau_{np}$, then the estimator $\hat{\Omega}^t$ satisfies

$$\text{sgn}(\hat{\omega}_{ij}^t) = \text{sgn}(\omega_{ij}^0)$$

for all $i, j = 1, \dots, p$ with probability at least $1 - O(p^{-\xi})$ or $1 - O(p^{-\xi/2} + n^{-\varepsilon/8})$.

4.2. Data-Driven Choice of λ_j

We describe the cross-validation procedure for choosing the tuning parameters λ_j in the p vector optimization problems (4). We randomly split the whole dataset into a training set with sample size n_1 and a test set with sample size $n_2 = n - n_1$, where $n_1 \asymp n_2 \asymp n$. To measure the predictive performance of $\tilde{\omega}_j$, we consider the loss function

$$L(\omega_j, \Sigma_c) = \frac{1}{2} \omega_j^T \Sigma_c \omega_j - \left(\mathbf{e}_j - \frac{\mathbf{1}_p}{p} \right)^T \omega_j, \quad (5)$$

whose derivative $\partial L / \partial \omega_j = \Sigma_c \omega_j - (\mathbf{e}_j - \mathbf{1}_p/p)$ corresponds to the constraint in problem (4). In a similar spirit to the D-trace loss of Zhang and Zou (2014), we write the loss function (5) in matrix form as

$$L(\Omega, \Sigma_c) = \frac{1}{2} \text{tr}(\Omega \Sigma_c \Omega) - \text{tr}(\mathbf{G} \Omega).$$

There is a subtle connection between $L(\Omega, \Sigma_c)$ and the CD-trace loss function (Yuan, He, and Deng 2019)

$$L_D(\Omega, \Sigma_0) = \frac{1}{2} \text{tr}(\mathbf{G} \Omega \mathbf{G} \Sigma_0 \mathbf{G} \Omega) - \text{tr}(\mathbf{G} \Omega),$$

where Ω is a positive definite basis precision matrix. This loss function is minimized at Ω_0 only when the exchangeability condition $\mathbf{G} \Sigma_0 = \Sigma_0 \mathbf{G}$ holds. Similarly, if $\mathbf{G} \Omega_0 = \Omega_0 \mathbf{G}$, from $\mathbf{G} \Sigma_0 \mathbf{G} = \Sigma_c$ and $\mathbf{G}^2 = \mathbf{G}$ we see that the two loss functions coincide at Ω_0 .

For $j = 1, \dots, p$ and the b th split, denote by $\hat{\omega}_j^{(1b)}(\lambda_j)$ the estimator with tuning parameter λ_j based on the training set, and by $\hat{\Sigma}_c^{(2b)}$ the sample centered log-ratio covariance matrix based on the test set. We choose the optimal value $\hat{\lambda}_j$ of λ_j by minimizing

$$\text{CV}(\lambda_j) = \frac{1}{B} \sum_{b=1}^B L\{\hat{\omega}_j^{(1b)}(\lambda_j), \hat{\Sigma}_c^{(2b)}\},$$

and compute the estimator $\hat{\omega}_j(\hat{\lambda}_j)$ based on the whole dataset. Let $\hat{\lambda} = (\hat{\lambda}_1, \dots, \hat{\lambda}_p)$ and then the data-driven estimator $\hat{\Omega}(\hat{\lambda})$ is obtained by combining $\hat{\omega}_j(\hat{\lambda}_j)$ for all columns. The procedure searches for the optimal $\hat{\lambda}_j$ through a grid of points $\delta_j \ell / N$, $\ell = 1, \dots, N$, for some sufficiently large $\delta_j > 0$. In practice, δ_j is set to the smallest value such that $\hat{\omega}_j = \mathbf{0}$; it is seen immediately from problem (4) that such a choice is $\delta_j = 1 - p^{-1}$. To establish theoretical guarantees for the data-driven estimator, we consider for simplicity the case of $B = 1$ and omit the superscript b . The rate of convergence for the data-driven estimator $\hat{\Omega}^{(1)}(\hat{\lambda}) = (\hat{\omega}_1^{(1)}(\hat{\lambda}_1), \dots, \hat{\omega}_p^{(1)}(\hat{\lambda}_p))$ is given by the following result.

Theorem 4. Assume that Condition 1 holds, $\Omega_0 \in \mathcal{U}_q(s_0(p), M_p)$ with $M_p = o(\sqrt{p})$, $\log N = o(n)$, and $\min_{1 \leq j \leq p} \delta_j \geq NC_0(\sqrt{(\log p)/n} + M_p/\sqrt{p})$ for $C_0 > 0$ defined in Lemma 1. Then the data-driven estimator $\hat{\Omega}^{(1)}(\hat{\lambda})$ satisfies

$$\frac{1}{p} \|\hat{\Omega}^{(1)}(\hat{\lambda}) - \Omega_0\|_F^2 = O_p \left\{ s_0(p) \left(M_p \sqrt{\frac{\log p}{n}} + \frac{M_p^2}{\sqrt{p}} \right)^{2-q} \right\}.$$

5. Extensions to Data Containing Zeros

We have so far confined our discussion to data lying in the strictly positive simplex. In many applications, however, the compositions are not directly observable but instead estimated from count data. When the count data are sparse, as is the case in microbiome studies, the compositional data obtained by simple normalization may contain an abundance of zeros and thus prevent the direct application of log-ratio transformations. In such situations, it is useful to distinguish between sampling zeros, which are due to undersampling of rare yet present taxa, and structural zeros, which represent truly absent taxa regardless of the sampling or sequencing depth (Agresti 2013, Sec. 10.6). We describe how our framework can be extended to accommodate these two types of zeros in Sections 5.1 and 5.2.

5.1. Sampling Zeros

In the presence of only sampling zeros, the underlying proportions are still positive and can be estimated by many existing methods, ranging from easy-to-implement zero replacement and variable correction procedures (Shi, Zhou, and Zhang 2022) to more sophisticated approaches that borrow information across samples and taxa (Zhang and Lin 2019; Cao, Zhang, and Li 2020). This amounts to observing a noisy version of the compositional data, which are subject to measurement error that should be accounted for in subsequent analysis. Since the optimization problem (3) depends on the data only through $\hat{\Sigma}_c$, it suffices to consider a perturbed covariance matrix $\check{\Sigma}_c = \hat{\Sigma}_c + \mathbf{E}$. For simplicity, consider the CARE estimator $\hat{\Omega} \equiv \hat{\Omega}(\lambda)$ with $\lambda_j = \lambda$ for all j , and let $\check{\Omega}$ be the perturbed version with $\hat{\Sigma}_c$ replaced by $\check{\Sigma}_c$. Let $\lambda_0 = \|\hat{\Sigma}_c \Omega_0 - \mathbf{G}\|_{\max}$. The following perturbation result characterizes the impact of the measurement error \mathbf{E} on the estimator $\hat{\Omega}$.

Theorem 5. Suppose that $\Omega_0 \in \mathcal{U}_q(s_0(p), M_p)$. If $\lambda = \lambda_0 + M_p \|\mathbf{E}\|_{\max}$, then the estimator $\check{\Omega}$ satisfies

$$\|\check{\mathbf{\Omega}} - \widehat{\mathbf{\Omega}}\|_{\max} \leq 2M_p(\lambda_0 + p^{-1}) + 3M_p^2\|\mathbf{E}\|_{\max}.$$

It is readily seen from [Theorem 5](#) that as long as \mathbf{E} is small enough so that $M_p\|\mathbf{E}\|_{\max} = O(\lambda_0)$, the second term in the perturbation bound is small. Furthermore, by [Lemma 1](#) we can choose $\lambda_0 \asymp \sqrt{(\log p)/n} + M_p/\sqrt{p}$ with high probability, so that the term p^{-1} is also negligible. Then the bound becomes

$$\|\check{\mathbf{\Omega}} - \widehat{\mathbf{\Omega}}\|_{\max} = O\left(M_p\sqrt{\frac{\log p}{n}} + \frac{M_p^2}{\sqrt{p}}\right),$$

matching the entrywise ℓ_∞ -error bound in [Theorem 2](#). This indicates that the measurement error would have no notable effect on the performance of the CARE estimator.

5.2. Structural Zeros

Structural zeros may occur when the frequency of zeros in the count data is higher than the sampling model predicts. A popular approach to this problem is zero-inflated models, which are mixtures of a point mass at zero (zero part) and a count data distribution with positive mean (nondegenerate part). One can, in principle, use a zero-inflated model for multivariate count data (e.g., Tang and Chen 2019; Xu, Demmer, and Li 2021; Zeng et al. 2022) to obtain strictly positive composition estimates from the nondegenerate part of the model; our approach can then be applied to these positive estimates as described earlier. The low-rank assumptions underlying some of these models, however, may lead to rank-deficient covariance structures, which render the estimation of a sparse precision matrix impossible. A remedy for the rank degeneracy problem is to construct an unbiased estimator of the centered log-ratio covariance matrix by inverse probability weighting. More specifically, suppose that $X_{kp} > 0$ for all k and let $\mathbf{Z}_k^{(p)} = (Z_{k1}^{(p)}, \dots, Z_{k,p-1}^{(p)})^T$ with $Z_{kj}^{(p)} = \log(X_{kj}/X_{kp})$ be the additive log-ratio transformation of \mathbf{X}_k . Let Δ_{kj} be the latent variable indicating whether X_{kj} comes from the nondegenerate part, and π_j the corresponding probability. In practice, Δ_{kj} and π_j are unknown but can be estimated from the adopted zero-inflated model. The covariance matrix $\mathbf{\Sigma}_{(p)}$ of $\mathbf{Z}_k^{(p)}$ can then be estimated by $\widehat{\mathbf{\Sigma}}_{(p)} = (\hat{\sigma}_{ij}^{(p)})$ with

$$\hat{\sigma}_{ij}^{(p)} = \frac{1}{n} \sum_{k=1}^n \frac{\Delta_{ki}\Delta_{kj}}{\pi_{ij}} Z_{ki}^{(p)} Z_{kj}^{(p)},$$

where $\pi_{ij} = \pi_i\pi_j$ if $i \neq j$ and π_i otherwise. By the relationship between $\mathbf{\Sigma}_c$ and $\mathbf{\Sigma}_{(p)}$ (Aitchison 2003, eq. (4.28)), we obtain an unbiased estimator of $\mathbf{\Sigma}_c$:

$$\check{\mathbf{\Sigma}}_c = \mathbf{F}^T \mathbf{H}^{-1} \widehat{\mathbf{\Sigma}}_{(p)} \mathbf{H}^{-1} \mathbf{F},$$

where $\mathbf{F} = (\mathbf{I}_{p-1}, -\mathbf{1}_{p-1})$ and $\mathbf{H} = \mathbf{I}_{p-1} + \mathbf{1}_{p-1}\mathbf{1}_{p-1}^T$. Now we can substitute $\check{\mathbf{\Sigma}}_c$ for $\widehat{\mathbf{\Sigma}}_c$ in the optimization problem (4). The performance of this extension, of course, hinges on the quality of the estimates of Δ_{kj} and π_j . A detailed investigation is needed but beyond the scope of this article.

5.3. Limitations

Despite the extensions outlined above, important limitations of our method exist. First, since structural zeros cannot be handled by the log transformation per se, our method should be combined with a zero-inflated model to estimate the covariance structure of its nondegenerate part. More generally, the covariance structure of the zero part and the interaction between the zero part and the nondegenerate part are also of interest, whose modeling remains an open problem. Second, our measurement error framework for dealing with sampling zeros requires the perturbation term \mathbf{E} to be sufficiently small, which may not be satisfied when the true proportions are extremely small and the total counts are limited. Developing more robust and efficient methods for these scenarios is an interesting topic.

6. Simulation Studies

In this section, we examine the numerical performance of the CARE estimator through simulation studies under different scenarios and compare it with a variety of existing methods. To implement our procedure, by splitting ω_j into a positive part ω_j^+ and a negative part ω_j^- , we reformulate problem (4) as the following linear programming problem:

$$\begin{aligned} & \text{minimize } \mathbf{1}_p^T (\omega_j^+ + \omega_j^-) \quad \text{subject to} \\ & \begin{pmatrix} \widehat{\mathbf{\Sigma}}_c & -\widehat{\mathbf{\Sigma}}_c \\ -\widehat{\mathbf{\Sigma}}_c & \widehat{\mathbf{\Sigma}}_c \end{pmatrix} \begin{pmatrix} \omega_j^+ \\ \omega_j^- \end{pmatrix} \leq \begin{pmatrix} \lambda_j \mathbf{1}_p + \mathbf{e}_j - \mathbf{1}_p/p \\ \lambda_j \mathbf{1}_p - \mathbf{e}_j + \mathbf{1}_p/p \end{pmatrix}, \\ & \omega_j^+ \geq 0, \quad \omega_j^- \geq 0, \end{aligned}$$

where the inequalities apply componentwise. These problems have the same size as, and hence comparable computational efficiency to, those in the CLIME procedure, and can be solved by many existing algorithms. Here we use an implementation of the parametric simplex method (Vanderbei 2020), which produces the entire solution path by solving the linear programming problem only once. Compared with interior-point and simplex methods, the parametric simplex method is more efficient for large-scale sparse learning problems (Pang et al. 2017).

We compare our estimator to three previously proposed methods: CD-trace (Yuan, He, and Deng 2019), gCoda (Fang et al. 2017), and SPIEC-EASI with the graphical lasso (Kurtz et al. 2015). The oracle estimator, which is obtained by applying the CLIME approach to the basis as it were observed, is also included for comparison. This ideal estimator serves as the benchmark for assessing whether the compositionality effect has been removed.

6.1. Simulation Results for Compositional Data

We first consider the case with positive compositional data and no zero correction is needed. The compositional data were generated as follows. First, we independently sampled $\mathbf{Y}_k = (Y_{k1}, \dots, Y_{kp})^T$, $k = 1, \dots, n$, from a multivariate normal distribution $N_p(\mathbf{0}, \mathbf{\Omega}_0^{-1})$. Then we generated $\mathbf{W}_k = (W_{k1}, \dots, W_{kp})^T$ and $\mathbf{X}_k = (X_{k1}, \dots, X_{kp})^T$, $k = 1, \dots, n$, through the transformations $W_{kj} = \exp(Y_{kj})$ and $X_{kj} = W_{kj} / \sum_{i=1}^p W_{ki}$ for

Table 1. Means and standard errors (in parentheses) of performance measures for different methods in model (a) over 100 replications.

p	Method				
	CARE	Oracle	CD-trace	gCoda	SPIEC-EASI
Spectral norm loss					
50	2.46 (0.12)	2.29 (0.12)	3.21 (0.03)	3.47 (0.05)	3.35 (0.04)
100	2.70 (0.10)	2.59 (0.10)	3.24 (0.02)	3.62 (0.04)	3.44 (0.02)
200	2.96 (0.07)	2.91 (0.08)	3.27 (0.02)	3.70 (0.02)	3.37 (0.02)
400	3.25 (0.03)	3.24 (0.04)	3.34 (0.02)	3.76 (0.02)	3.50 (0.03)
Matrix ℓ_1 -norm loss					
50	3.28 (0.20)	3.05 (0.23)	3.53 (0.06)	3.86 (0.06)	3.93 (0.06)
100	3.45 (0.15)	3.36 (0.16)	3.61 (0.07)	3.98 (0.05)	3.89 (0.03)
200	3.61 (0.12)	3.56 (0.11)	3.64 (0.05)	4.07 (0.03)	3.91 (0.03)
400	3.77 (0.07)	3.76 (0.07)	3.76 (0.07)	4.18 (0.03)	3.99 (0.03)
Frobenius norm loss					
50	6.79 (0.29)	6.31 (0.25)	9.83 (0.08)	10.93 (0.21)	10.37 (0.13)
100	10.30 (0.25)	9.84 (0.27)	14.22 (0.07)	16.35 (0.24)	14.87 (0.04)
200	16.68 (0.19)	16.33 (0.19)	20.21 (0.07)	23.98 (0.18)	20.59 (0.07)
400	27.04 (0.15)	26.85 (0.16)	29.12 (0.12)	34.50 (0.18)	29.71 (0.33)
True positive rate (%)					
50	90.5 (2.6)	94.5 (2.3)	74.9 (2.8)	74.5 (6.4)	54.8 (4.0)
100	89.1 (2.2)	91.8 (2.0)	60.4 (2.4)	49.2 (5.8)	65.6 (2.3)
200	83.7 (1.6)	85.9 (1.4)	60.9 (1.5)	39.6 (3.5)	66.9 (1.6)
400	68.2 (1.4)	69.6 (1.5)	53.0 (1.1)	34.2 (2.2)	65.4 (3.4)
False positive rate (%)					
50	7.4 (0.7)	6.0 (0.7)	3.6 (0.2)	9.3 (1.0)	5.9 (0.4)
100	3.6 (0.3)	3.3 (0.2)	1.1 (0.1)	2.9 (0.4)	2.3 (0.1)
200	1.3 (0.1)	1.2 (0.1)	0.5 (0.0)	1.2 (0.1)	1.2 (0.1)
400	0.4 (0.0)	0.4 (0.0)	0.2 (0.0)	0.5 (0.1)	0.5 (0.1)

Table 2. Means and standard errors (in parentheses) of performance measures for different methods in model (b) over 100 replications.

p	Method				
	CARE	Oracle	CD-trace	gCoda	SPIEC-EASI
Spectral norm loss					
50	2.04 (0.17)	1.91 (0.20)	2.29 (0.11)	3.47 (0.07)	3.68 (0.09)
100	1.99 (0.12)	1.92 (0.14)	2.23 (0.05)	3.77 (0.12)	3.55 (0.12)
200	2.23 (0.15)	2.16 (0.12)	2.83 (0.09)	3.82 (0.06)	3.75 (0.04)
400	2.44 (0.15)	2.36 (0.14)	2.94 (0.07)	4.02 (0.01)	3.43 (0.02)
Matrix ℓ_1 -norm loss					
50	3.60 (0.35)	3.17 (0.43)	3.77 (0.22)	6.03 (0.18)	6.35 (0.16)
100	4.22 (0.47)	3.76 (0.56)	4.38 (0.35)	7.97 (0.30)	7.21 (0.34)
200	4.59 (0.42)	4.35 (0.33)	6.39 (0.23)	7.49 (0.06)	7.05 (0.11)
400	5.69 (0.50)	5.48 (0.44)	7.15 (0.17)	8.69 (0.03)	6.91 (0.09)
Frobenius norm loss					
50	5.49 (0.24)	5.18 (0.23)	6.13 (0.16)	11.27 (0.31)	13.16 (0.33)
100	7.78 (0.23)	7.54 (0.24)	9.03 (0.14)	16.47 (0.45)	17.84 (0.43)
200	11.68 (0.25)	11.45 (0.27)	13.97 (0.13)	25.03 (0.75)	26.94 (0.04)
400	16.84 (0.12)	16.69 (0.14)	20.18 (0.10)	34.12 (0.08)	35.95 (0.02)
True positive rate (%)					
50	87.6 (3.8)	89.4 (3.8)	79.7 (3.3)	56.9 (2.0)	54.4 (4.6)
100	85.8 (3.3)	87.4 (3.0)	65.6 (3.1)	30.5 (2.6)	63.6 (8.9)
200	81.7 (2.4)	83.2 (2.3)	53.4 (2.4)	14.8 (2.2)	71.9 (2.2)
400	82.4 (1.5)	83.1 (1.5)	48.5 (1.1)	10.2 (0.4)	90.4 (1.1)
False positive rate (%)					
50	2.1 (0.5)	1.4 (0.5)	2.1 (0.3)	4.4 (0.4)	10.0 (1.8)
100	1.2 (0.2)	1.0 (0.2)	0.4 (0.1)	1.1 (0.1)	4.5 (0.4)
200	0.7 (0.1)	0.7 (0.1)	0.1 (0.0)	0.3 (0.0)	1.1 (0.1)
400	0.2 (0.0)	0.2 (0.0)	0.0 (0.0)	0.1 (0.0)	0.5 (0.0)

$j = 1, \dots, p$. As a result, \mathbf{W}_k and \mathbf{X}_k follow multivariate log-normal and logistic-normal distributions, respectively. We considered the following four models for generating $\mathbf{\Omega}_0$ with diverse network structures.

- (a) *Band graph*: Let $\mathbf{\Omega}_1 = (\omega_{ij}^1)$ and $\mathbf{\Omega}_0 = \mathbf{\Omega}_1 + (|\lambda_{\min}(\mathbf{\Omega}_1)| + 0.01)\mathbf{I}_p$, where $\omega_{i,i+1}^1 = \omega_{i+1,i}^1 = 0.8$, $\omega_{i,i+2}^1 = \omega_{i+2,i}^1 = 0.5$, $\omega_{ij}^1 = 0$ for $|i - j| \geq 3$, and ω_{ii}^1 were drawn from a uniform distribution $U(1, 2)$.
- (b) *Hub graph*: The p nodes were divided into blocks of size 5. For each block, one hub was selected and connected to the other nodes in the same block, and each edge weight was set to 0.8 or 0.5 with equal probability. The diagonal entries were set as in model (a).
- (c) *Block graph*: The p nodes were equally divided into five blocks. Each pair of nodes in the same block were connected with probability $20/p$, and each edge weight was set as in model (b). The diagonal entries were set large enough as in model (a) so that $\mathbf{\Omega}_0$ was positive definite.
- (d) *Random graph*: Each pair of nodes were connected with probability $4/p$. The edge weights were set as in model (b) and the diagonal entries were set as in model (a).

Throughout the simulations, we took the sample size $n = 200$ and dimension $p = 50, 100, 200, 400$. Five performance measures are adopted: the spectral norm, matrix ℓ_1 -norm and Frobenius norm losses for evaluating the estimation accuracy, and the true positive and false positive rates for assessing the support recovery performance. For the CARE and oracle methods, the tuning parameters λ_j were chosen by 5-fold cross-validation.

The simulation results for models (a) and (b) over 100 replications are reported in Tables 1 and 2, and those for models (c) and (d) in Supplementary Tables S4 and S5. Overall, we see

that the CARE and oracle methods perform nearly equally well and outperform the other three competitors by a large margin in almost all settings. In particular, the performance of CARE tends to be closer to that of the oracle estimator in higher dimensions, confirming the blessing of dimensionality revealed by our theory. Among the three previously proposed methods, CD-trace seems to outperform gCoda and SPIEC-EASI in terms of estimation accuracy. In addition, SPIEC-EASI appears to have some advantages for recovering more edges in high dimensions over CD-trace and gCoda, but it does so at the expense of an inflated false positive rate.

To compare the support recovery performance without using a specific tuning approach, the receiver operating characteristic (ROC) curves for all methods under models (a)–(d) are shown in Figure 1. We observe that the ROC curves of the CARE and oracle methods dominate those of the other three methods uniformly in low to moderate dimensions, as well as in high dimensions when the false positive rates are small. These results embody the superiority of CARE in terms of support recovery. Moreover, SPIEC-EASI performs better for models (a) and (c) than for models (b) and (d) in low dimensions, indicating that its performance may depend more critically on the network topology. The higher right tails of ROC curves for the competing methods in certain high-dimensional settings may not be relevant in practice, since only the sparse regime is of interest.

6.2. Simulation Results for Count Data

To illustrate how our method work with zeros, we next examine the case where the generated data are counts and include many sampling zeros. We adopted the graph structure of model (a) and the same settings as before for generating $(\mathbf{Y}_k, \mathbf{X}_k)$, $k = 1, \dots, n$, except that the mean vector of \mathbf{Y}_k was drawn from a uniform dis-

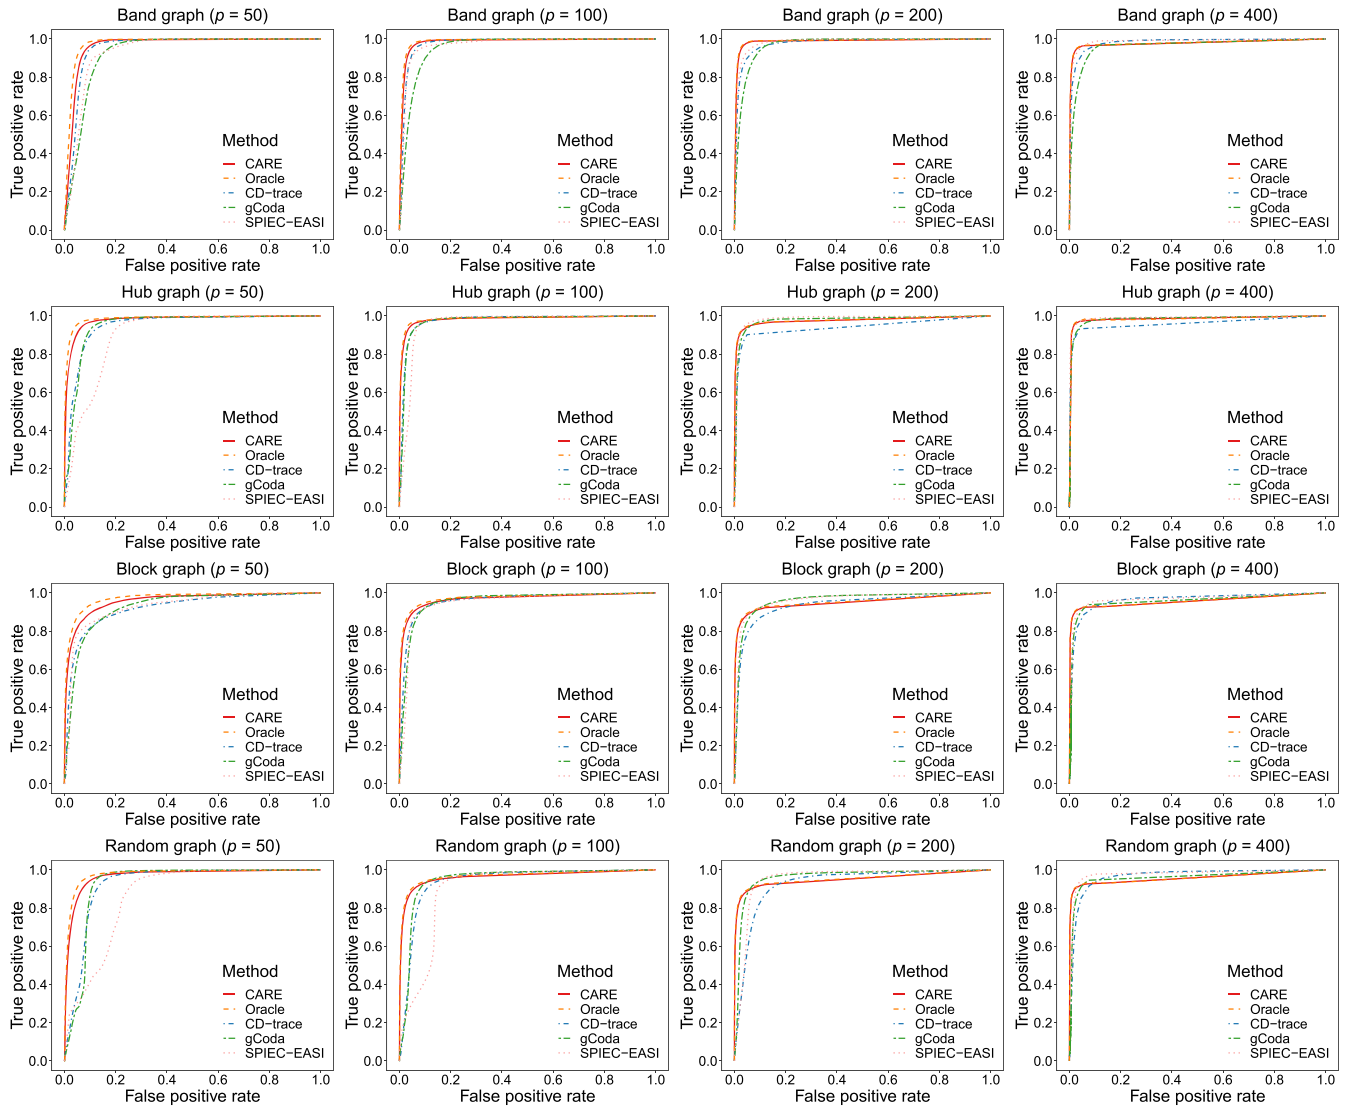


Figure 1. The ROC curves for different methods in models (a)–(d) with $p = 50, 100, 200, 400$.

tribution $U(0, 5)$ to allow a higher diversity. The count data were then generated from a multinomial distribution $\text{Mult}(m_k, \mathbf{X}_k)$, where m_k were sampled uniformly from $15p, \dots, 15p + 500$ and $p, \dots, 2p$, resulting in about 35% and 70% zero counts, respectively, for $p = 50$ and 100.

The following methods are included in our comparisons: CARE using the compositional data directly (Comp+CARE), CARE, CD-trace, and SPIEC-EASI using composition estimates from the logistic normal multinomial model of Zhang and Lin (2019) (LNM+CARE, LNM+CD-trace, and LNM+SPIEC-EASI, respectively), and CARE using the +0.5 variable correction procedure of Shi, Zhou, and Zhang (2022) (VC+CARE). The simulation results reported in Table 3 suggest that the proposed method works reasonably well with count data that contain many zeros, with the performance diminishing as the proportion of zeros increases. Moreover, model-based, more computationally intensive methods for composition estimation such as LNM tend to improve the performance of our method. All combined with LNM, CARE still exhibits a performance gap over CD-trace and SPIEC-EASI by identifying substantially more true positives.

7. Application to Gut Microbiome Data

The gut microbiome is considered to participate in many host physiological processes and have a tremendous impact on human health. Reconstructing ecological networks from metagenomic data holds the potential to reveal complex microbial interaction patterns in the gut. Some studies (e.g., Busiello et al. 2017) have found that species interaction networks are sparse and have related the property to explorability and dynamical robustness, thus meeting the requirement of our method. Here we illustrate our method by applying it to a dataset in Wu et al. (2011), which was previously analyzed by Cao, Lin, and Li (2019) using covariance estimation. In this study, DNA from fecal samples of 98 healthy subjects were quantified by 454/Roche pyrosequencing of 16S rRNA gene segments, yielding 87 genera that appeared in at least one sample. These subjects were divided into a lean group of $n = 63$ samples with body mass index (BMI) < 25 and an obese group of $n = 35$ samples with BMI ≥ 25 . To ensure stable detection of microbial interactions, we filtered out too rare genera and retained $p = 40$ genera that were present in at least four samples in each group. We adopted

Table 3. Simulation results for count data: means and standard errors (in parentheses) of performance measures for different methods in model (a) over 100 replications.

		Method					
p		Zeros	Comp+CARE	LMN+CARE	VC+CARE	LMN+CD-trace	LMN+SPIEC-EASI
Spectral norm loss							
50	35%	2.46 (0.12)	3.21 (0.09)	3.40 (0.08)	3.41 (0.06)	3.42 (0.04)	
	70%	–	3.78 (1.09)	14.78 (3.80)	3.67 (0.79)	3.44 (0.03)	
100	35%	2.91 (0.09)	3.44 (0.04)	3.47 (0.06)	3.47 (0.04)	3.60 (0.02)	
	70%	–	3.82 (0.97)	14.18 (1.59)	3.76 (0.88)	3.60 (0.01)	
Matrix ℓ_1 -norm loss							
50	35%	3.23 (0.19)	3.83 (0.11)	4.07 (0.14)	3.82 (0.09)	3.97 (0.06)	
	70%	–	4.61 (1.24)	17.16 (3.93)	4.44 (1.09)	3.94 (0.08)	
100	35%	3.60 (0.11)	3.99 (0.07)	4.09 (0.10)	3.87 (0.07)	3.98 (0.05)	
	70%	–	4.52 (1.14)	16.59 (1.63)	4.39 (1.18)	3.99 (0.05)	
Frobenius norm loss							
50	35%	6.79 (0.21)	9.43 (0.24)	10.13 (0.13)	10.31 (0.12)	10.74 (0.21)	
	70%	–	10.61 (0.65)	28.74 (5.87)	10.66 (0.50)	10.98 (0.13)	
100	35%	11.16 (0.20)	14.79 (0.20)	14.81 (0.15)	15.18 (0.14)	15.84 (0.10)	
	70%	–	15.58 (0.45)	42.97 (5.17)	15.41 (0.43)	16.04 (0.04)	
True positive rate (%)							
50	35%	91.2 (3.0)	63.8 (5.4)	43.3 (5.0)	42.7 (4.0)	37.9 (7.0)	
	70%	–	24.3 (4.1)	10.3 (2.6)	15.9 (3.2)	17.5 (4.0)	
100	35%	83.3 (2.5)	47.7 (3.7)	32.5 (3.4)	28.7 (4.0)	28.8 (3.6)	
	70%	–	16.8 (2.7)	6.1 (1.9)	9.4 (1.9)	12.1 (2.0)	
False positive rate (%)							
50	35%	7.3 (0.8)	10.8 (1.2)	8.2 (1.0)	4.0 (0.5)	5.1 (0.8)	
	70%	–	9.1 (1.2)	4.1 (0.7)	4.0 (0.8)	3.4 (0.8)	
100	35%	2.8 (0.2)	3.7 (0.4)	2.7 (0.4)	1.0 (0.2)	1.6 (0.2)	
	70%	–	3.6 (0.4)	1.2 (0.2)	1.0 (0.3)	1.4 (0.3)	

Table 4. Numbers of positive and negative edges and stability of networks for different methods applied to the gut microbiome data.

Method	Number of all edges		Number of stable edges		Network stability
	Positive	Negative	Positive	Negative	
Lean group					
CARE	9	19	5	14	0.766
CD-trace	9	10	4	1	0.582
gCoda	12	16	2	5	0.577
SPIEC-EASI	12	18	4	6	0.680
Obese group					
CARE	4	8	2	3	0.708
CD-trace	16	15	1	1	0.538
gCoda	4	6	0	3	0.636
SPIEC-EASI	13	9	1	3	0.551

the multisample approach of Cao, Zhang, and Li (2020) to deal with excess zeros and obtained a positive composition matrix.

For each group, we estimated the basis precision matrix using different methods. The results are represented as microbial interaction networks among the genera. For the CARE method, the tuning parameters were selected by 10-fold cross-validation. To assess the stability of support recovery, we randomly subsampled 80% of the subjects and recorded the proportion of edges that reoccurred in the subsample. We repeated the subsampling procedure 100 times and adopt the average proportion of reproduced edges as a measure of network stability. Finally, only those edges reproduced in at least 80% of the subsamples were retained in the networks. The networks identified by the CARE method are shown in Figure 2, and those by the other methods in Supplementary Figures S3–S5. The numbers of positive and negative edges and stability of networks for all methods are summarized in Table 4.

As depicted in Figure 2, the network structures revealed by the CARE method for the lean and obese groups look markedly

different. The genus–genus interactions for the obese group are substantially fewer and less complex than those for the lean group, agreeing with the previous finding that obese microbiomes induce a less modular metabolic network than lean microbiomes (Greenblum, Turnbaugh, and Borenstein 2012). From Table 4 we see that CARE maintains the highest network stability and detects more stable edges than the other methods. Interestingly, the networks constructed by CARE involve more negative than positive interactions; a similar phenomenon was observed by Cao, Lin, and Li (2019) in microbial correlation networks and tends to be supported by the ecological theory of microbiome stability (Coyte, Schluter, and Foster 2015). More discussion on the identified microbial interactions can be found in Supplementary Section S.3.3.

8. Discussion

The specification of the compositional precision matrix and the associated tools developed in this article establish a natural link to the basis inverse covariance structure. The blessing of dimensionality derived from this inverse covariance relationship formalizes the intuition that the nonidentifiability due to not knowing the magnitude of the basis spreads over all p components and becomes negligible as the dimensionality grows. These insights are of both theoretical and practical importance, opening up the possibility of investigating existing and developing new methodology for estimating the basis precision matrix. The general idea is to use Ω_c as a proxy for Ω_0 and develop procedures for estimating Ω_0 through Ω_c . Although Ω_c is not sparse, its matrix ℓ_1 -norm can be tightly controlled, a fact that can be directly exploited by a compositionally adjusted CLIME procedure. Extensions to neighborhood and likelihood-based methods seem possible but remain to be explored. Despite the similarities we have discussed, our problem also differs from general low-rank plus sparse matrix recovery problems in important ways. In particular, the rank in our problem is known and need not be estimated from the data. Also, the usual spiked eigenvalue assumption (e.g., Fan, Liao, and Mincheva 2013) is not met. These features prevent methods based on nuclear norm regularization or principal components from being applied to our context.

Our conditional dependence modeling approach relies on the log transformation, which is not artificial but natural in at least two ways. First, as a consequence of a large number of independent multiplicative effects, the log-normal distribution is usually a better model than the normal for many real data with small means, large variances, and positive values (Limpert, Stahel, and Abbt 2001). This is, for instance, the case for microbiome data since bacterial abundance generally follows an exponential growth model with additional variability. Second, methods and theory induced by the log transformation, such as multinomial logistic and log-linear models, have been fundamental to categorical data analysis (Agresti 2013). In light of the intimate connections between the two areas, the log transformation is a sensible choice that is compatible with many related methods and theory. Nevertheless, it is not the only viable way to deal with the simplex constraint. Notably, square root and power transformations have been applied to compositional data in regression and principal component analysis (Scealy and Welsh

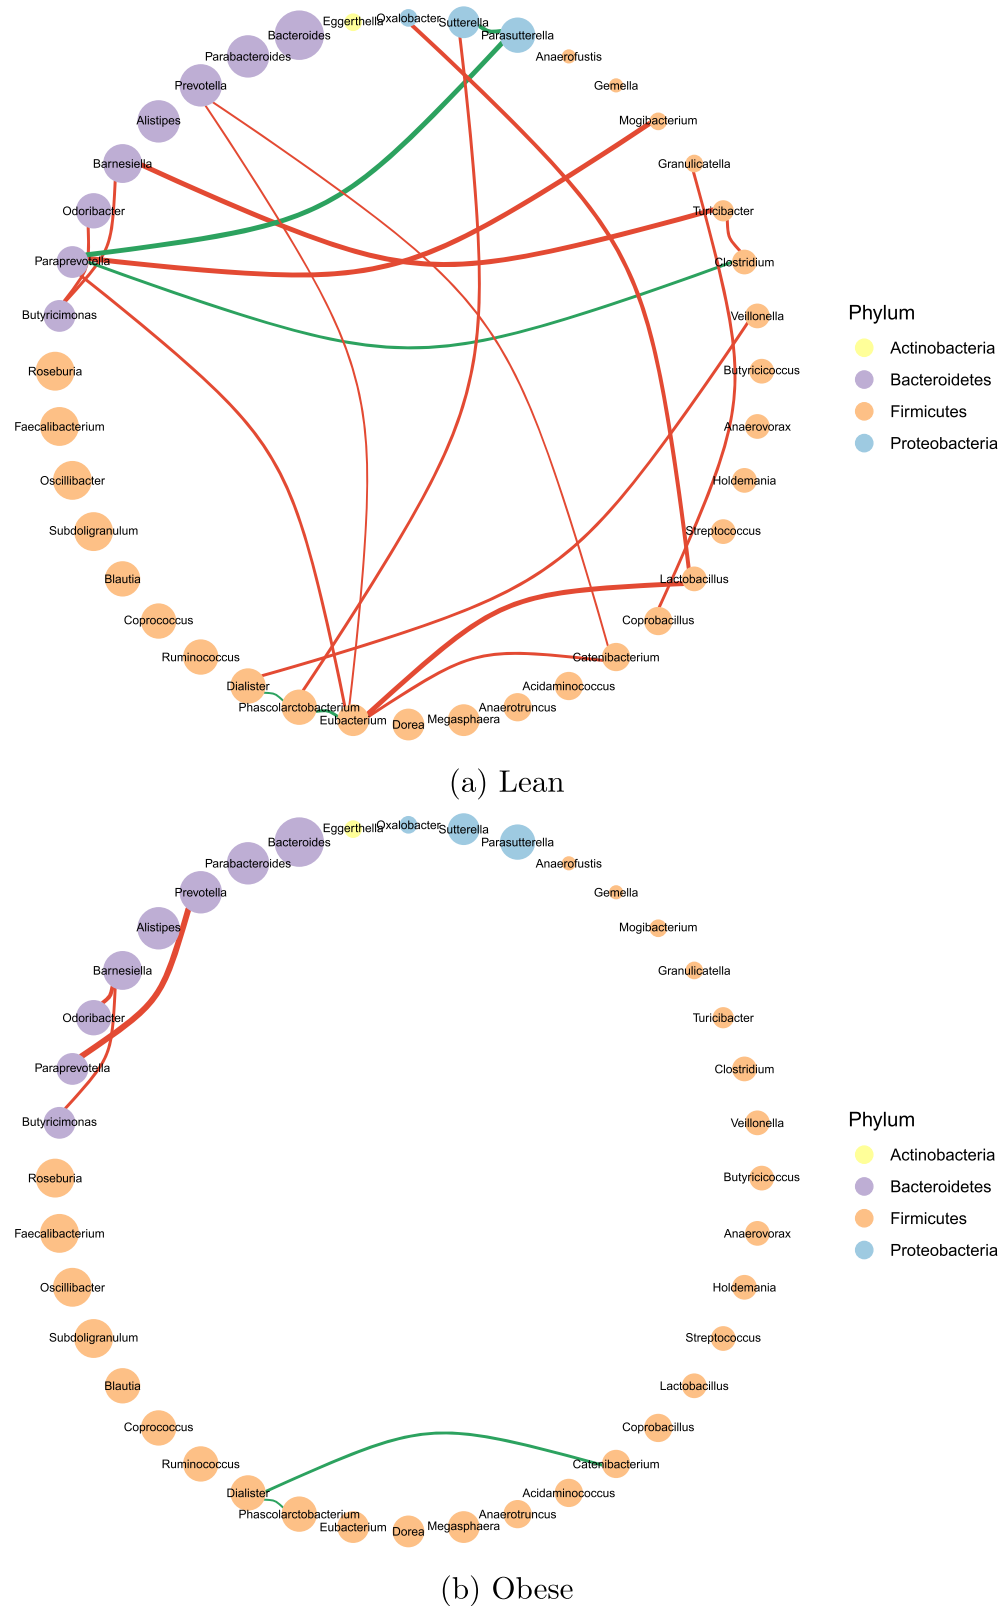


Figure 2. Microbial interaction networks identified by the CARE method for the (a) lean and (b) obese groups in the gut microbiome data. Positive and negative edges are displayed in green and red, respectively, with thicknesses proportional to their strengths. Node sizes are proportional to the relative abundances of genera among all samples.

2011; Scealy et al. 2015). It would be worthwhile to develop notions of conditional dependence and graphical models under these alternative transformations, which we leave for future work.

Supplementary Materials

The supplementary materials contain the proofs of theoretical results, additional discussion and numerical results, and R code and data.

Acknowledgments

The authors thank the Associate Editor and two reviewers for valuable comments that have led to an improved article.

Disclosure Statement

The authors report there are no competing interests to declare.

Funding

Lin's research was supported by National Natural Science Foundation of China (12171012, 12292980, and 12292981). Zhang's research was supported by National Natural Science Foundation of China (12201111 and 12371264) and the Fundamental Research Funds for the Central Universities in UIBE (21QD13).

ORCID

Huiyuan Wang  <http://orcid.org/0009-0004-8796-7376>

Wei Lin  <http://orcid.org/0000-0002-7598-6199>

References

- Agresti, A. (2013), *Categorical Data Analysis* (3rd ed.), Hoboken, NJ: Wiley. [6,10]
- Aitchison, J. (1982), "The Statistical Analysis of Compositional Data" (with discussion), *Journal of the Royal Statistical Society, Series B*, 44, 139–177. [1]
- (2003), *The Statistical Analysis of Compositional Data*, Caldwell: The Blackburn Press. [1,2,3,4,7]
- Blei, D. M., and Lafferty, J. D. (2007), "A Correlated Topic Model of Science," *The Annals of Applied Statistics*, 1, 17–35. [1]
- Busiello, D. M., Suweis, S., Hidalgo, J., and Maritan, A. (2017), "Explorability and the Origin of Network Sparsity in Living Systems," *Scientific Reports*, 7, 12323. [9]
- Cai, T., Liu, W., and Luo, X. (2011), "A Constrained ℓ_1 Minimization Approach to Sparse Precision Matrix Estimation," *Journal of the American Statistical Association*, 106, 594–607. [2,4,5]
- Cai, T. T., Liu, W., and Zhou, H. H. (2016), "Estimating Sparse Precision Matrix: Optimal Rates of Convergence and Adaptive Estimation," *The Annals of Statistics*, 44, 455–488. [2,4,5,6]
- Candès, E. J., Li, X., Ma, Y., and Wright, J. (2011), "Robust Principal Component Analysis?" *Journal of the ACM*, 58, 11. [4]
- Cao, Y., Lin, W., and Li, H. (2019), "Large Covariance Estimation for Compositional Data via Composition-Adjusted Thresholding," *Journal of the American Statistical Association*, 114, 759–772. [1,2,3,5,9,10]
- Cao, Y., Zhang, A., and Li, H. (2020), "Multisample Estimation of Bacterial Composition Matrices in Metagenomics Data," *Biometrika*, 107, 75–92. [6,10]
- Chandrasekaran, V., Parrilo, P. A., and Willsky, A. S. (2012), "Latent Variable Graphical Model Selection via Convex Optimization" (with discussion), *The Annals of Statistics*, 40, 1935–2013. [4]
- Coyte, K. Z., Schluter, J., and Foster, K. R. (2015), "The Ecology of the Microbiome: Networks, Competition, and Stability," *Science*, 350, 663–666. [10]
- Fan, J., Liao, Y., and Mincheva, M. (2013), "Large Covariance Estimation by Thresholding Principal Orthogonal Complements" (with discussion), *Journal of the Royal Statistical Society, Series B*, 75, 603–680. [4,5,10]
- Fan, Y., and Lv, J. (2016), "Innovated Scalable Efficient Estimation in Ultra-Large Gaussian Graphical Models," *The Annals of Statistics*, 44, 2098–2126. [2]
- Fang, H., Huang, C., Zhao, H., and Deng, M. (2017), "gCoda: Conditional Dependence Network Inference for Compositional Data," *Journal of Computational Biology*, 24, 699–708. [2,7]
- Friedman, J., Hastie, T., and Tibshirani, R. (2008), "Sparse Inverse Covariance Estimation With the Graphical Lasso," *Biostatistics*, 9, 432–441. [2]
- Greenblum, S., Turnbaugh, P. J., and Borenstein, E. (2012), "Metagenomic Systems Biology of the Human Gut Microbiome Reveals Topological Shifts Associated With Obesity and Inflammatory Bowel Disease," *Proceedings of the National Academy of Sciences*, 109, 594–599. [10]
- Kurtz, Z. D., Müller, C. L., Miraldi, E. R., Littman, D. R., Blaser, M. J., and Bonneau, R. A. (2015), "Sparse and Compositionally Robust Inference of Microbial Ecological Networks," *PLoS Computational Biology*, 11, e1004226. [2,7]
- Lam, C., and Fan, J. (2009), "Sparsistency and Rates of Convergence in Large Covariance Matrix Estimation," *The Annals of Statistics*, 37, 4254–4278. [2]
- Lauritzen, S. L. (1996), *Graphical Models*, Oxford: Oxford University Press. [1]
- Li, H. (2015), "Microbiome, Metagenomics, and High-Dimensional Compositional Data Analysis," *Annual Review of Statistics and Its Application*, 2, 73–94. [1]
- Limpert, E., Stahel, W. A., and Abbt, M. (2001), "Log-Normal Distributions Across the Sciences: Keys and Clues," *BioScience*, 51, 341–352. [3,10]
- Liu, H., Lafferty, J., and Wasserman, L. (2009), "The Nonparanormal: Semiparametric Estimation of High Dimensional Undirected Graphs," *Journal of Machine Learning Research*, 10, 2295–2328. [3]
- Liu, W., and Luo, X. (2015), "Fast and Adaptive Sparse Precision Matrix Estimation in High Dimensions," *Journal of Multivariate Analysis*, 135, 153–162. [2]
- Loh, P.-L., and Wainwright, M. J. (2013), "Structure Estimation for Discrete Graphical Models: Generalized Covariance Matrices and Their Inverses," *The Annals of Statistics*, 41, 3022–3049. [3]
- Meinshausen, N., and Bühlmann, P. (2006), "High-Dimensional Graphs and Variable Selection With the Lasso," *The Annals of Statistics*, 34, 1436–1462. [2]
- Morrison, R., Baptista, R., and Basor, E. (2022), "Diagonal Nonlinear Transformations Preserve Structure in Covariance and Precision Matrices," *Journal of Multivariate Analysis*, 190, 104983. [3]
- Pang, H., Vanderbei, R., Liu, H., and Zhao, T. (2017), "Parametric Simplex Method for Sparse Learning," in *Advances in Neural Information Processing Systems*, vol. 30, pp. 188–197. [7]
- Paulson, J. N., Stine, O. C., Bravo, H. C., and Pop, M. (2013), "Differential Abundance Analysis for Microbial Marker-Gene Surveys," *Nature Methods*, 10, 1200–1202. [3]
- Ren, Z., Sun, T., Zhang, C.-H., and Zhou, H. H. (2015), "Asymptotic Normality and Optimalities in Estimation of Large Gaussian Graphical Models," *The Annals of Statistics*, 43, 991–1026. [2,4]
- Scealy, J. L., de Caritat, P., Grunsky, E. C., Tsagris, M. T., and Welsh, A. H. (2015), "Robust Principal Component Analysis for Power Transformed Compositional Data," *Journal of the American Statistical Association*, 110, 136–148. [11]
- Scealy, J. L., and Welsh, A. H. (2011), "Regression for Compositional Data by Using Distributions Defined on the Hypersphere," *Journal of the Royal Statistical Society, Series B*, 73, 351–375. [11]
- Schwager, E., Mallick, H., Venz, S., and Huttenhower, C. (2017), "A Bayesian Method for Detecting Pairwise Associations in Compositional Data," *PLoS Computational Biology*, 13, e1005852. [2]
- Shi, P., Zhou, Y., and Zhang, A. R. (2022), "High-Dimensional Log-Error-in-Variable Regression With Applications to Microbial Compositional Data Analysis," *Biometrika*, 109, 405–420. [6,9]
- Sutton, J. (2007), "Market Share Dynamics and the 'Persistence of Leadership' Debate," *American Economic Review*, 97, 222–241. [1]
- Tang, Z.-Z., and Chen, G. (2019), "Zero-Inflated Generalized Dirichlet Multinomial Regression Model for Microbiome Compositional Data Analysis," *Biostatistics*, 20, 698–713. [7]
- van Donkelaar, A., Martin, R. V., Li, C., and Burnett, R. T. (2019), "Regional Estimates of Chemical Composition of Fine Particulate Matter Using A Combined Geoscience-Statistical Method With Information From Satellites, Models, and Monitors," *Environmental Science & Technology*, 53, 2595–2611. [1]
- Vanderbei, R. J. (2020), *Linear Programming: Foundations and Extensions* (5th ed.), Cham: Springer. [7]
- Wu, G. D., Chen, J., Hoffmann, C., Bittinger, K., Chen, Y.-Y., Keilbaugh, S. A., Bewtra, M., Knights, D., Walters, W. A., Knight, R., Sinha, R., Gilroy, E., Gupta, K., Baldassano, R., Nessel, L., Li, H., Bushman, F. D.,

- and Lewis, J. D. (2011), “Linking Long-Term Dietary Patterns With Gut Microbial Enterotypes,” *Science*, 334, 105–108. [9]
- Xu, T., Demmer, R. T., and Li, G. (2021), “Zero-Inflated Poisson Factor Model With Application to Microbiome Read Counts,” *Biometrics*, 77, 91–101. [7]
- Yuan, H., He, S., and Deng, M. (2019), “Compositional Data Network Analysis via Lasso Penalized D-Trace Loss,” *Bioinformatics*, 35, 3404–3411. [2,6,7]
- Yuan, M. (2010), “High Dimensional Inverse Covariance Matrix Estimation via Linear Programming,” *Journal of Machine Learning Research*, 11, 2261–2286. [2]
- Yuan, M., and Lin, Y. (2007), “Model Selection and Estimation in the Gaussian Graphical Model,” *Biometrika*, 94, 19–35. [2]
- Zeng, Y., Pang, D., Zhao, H., and Wang, T. (2022), “A Zero-Inflated Logistic Normal Multinomial Model for Extracting Microbial Compositions,” *Journal of the American Statistical Association*, to appear. [7]
- Zhang, J., and Lin, W. (2019), “Scalable Estimation and Regularization for the Logistic Normal Multinomial Model,” *Biometrics*, 75, 1098–1108. [6,9]
- Zhang, T., and Zou, H. (2014), “Sparse Precision Matrix Estimation via Lasso Penalized D-Trace Loss,” *Biometrika*, 101, 103–120. [2,6]
- Zhang, X., and He, F. (2019), “Large Precision Matrix Estimation for Compositional Data,” *Statistics and Application*, 8, 777–783. [2]