

The Maximum Entropy Principle For Compositional Data

Corey Weistuch¹, Jiening Zhu², Joseph O. Deasy¹, Allen R. Tannenbaum^{2,3}

1 Department of Medical Physics, Memorial Sloan Kettering Cancer Center, NY, United States

2 Department of Applied Mathematics & Statistics, Stony Brook University, NY, United States

3 Department of Computer Science, Stony Brook University, NY, United States

* allen.tannenbaum@stonybrook.edu

Abstract

In this work, we provide a general method for inferring the stochastic behavior of compositional systems. Our approach is guided by the principle of maximum entropy, a data-driven modeling technique. In particular, we show that our method can accurately capture stochastic, inter-species relationships with minimal model parameters. We provide two proofs of principle. First, we measure the relative abundances of different bacteria and infer how they interact. Second, we show that our method outperforms a common alternative for the extraction of gene-gene interactions in triple-negative breast cancer.

Author summary

Compositional systems, represented as proportions of some whole, are ubiquitous. They encompass the abundances of proteins in a cell, the distribution of organisms in nature, and the stoichiometry of the most basic chemical reactions. Thus, a central goal is to understand how such processes emerge from the behaviors of their components and their pairwise interactions. Such a study, however, is challenging for two key reasons. Firstly, such systems are complex and depend, often stochastically, on their constituent parts. Secondly, the data lie on a simplex which influences their correlations. We provide a general and data-driven modeling tool for compositional systems to resolve both of these issues. We achieve this through the principle of maximum entropy, which requires only minimal assumptions and limited experimental data in contrast to current alternatives. We show that our approach provides novel and biologically-intuitive insights and is promising as a comprehensive quantitative framework for compositional data.

Introduction

Describing the compositions of physical systems, such as in mixtures of industrial chemical reactions, across bacteria in the microbiome, or relative influences in cancer networks is of significant practical importance. In the present work, these systems are modeled as networks of components (or nodes) and their unknown node-node interactions. However, the challenge of inferring these interactions lies in incorporating the defining feature of such compositions: the total proportion across components must always sum to one (or 100%).

Much recent interest has been devoted to improving the statistical analysis of compositional data [1]. The typical strategies that have been employed broadly fall into two categories. First, many apply traditional statistics (such as correlational analyses). Applied to compositional data, however, such tools are known to generate highly inaccurate results [1–3]. The second approach is to consider the simplicial geometry of compositional data more carefully. Here, a simplex describes the set of points that sum to a fixed quantity. Such simplices (triangles, tetrahedrons, etc.) are common in mathematics but less studied in statistics [4, 5]. Current methods for utilizing this information first modify the data (such as through log ratio transformations), but this can introduce unwarranted biases into downstream analysis [1, 6, 7]. What is needed is a more general and principled approach for describing compositional data.

In contrast to previous approaches, we aim to infer a general model for compositional data from first principles. The natural method for this is the principle of maximum entropy or Max Ent [8–11]. Here, one provides constraints, such as means, variances, and even the geometry of the data itself, and Max Ent provides the model. The advantage of this approach is twofold. First, as opposed to other modeling approaches, Max Ent makes minimal assumptions that are not warranted by the data itself [12]. Second, Max Ent is a widely and successfully utilized modeling framework for complex biological systems [13–18]. We provide theory and practical demonstrations of our new approach in the present work.

The model

Suppose one is given several stochastic observations of the relative abundances of N different components. Each of these observations may be represented as a vector $\Gamma = \{s_1, s_2, \dots, s_N\}$. Our goal is to infer the most likely and least-biased inter-component relationships that give rise to these observations (see Fig 1). The unique model with this property is provided by the *principle of maximum entropy*, which selects the model P that both maximizes the entropy $S = -\sum_{\Gamma} P_{\Gamma} \log P_{\Gamma}$ and satisfies known constraints from the data. Here, the standard constraints are the estimated first and second moments, $M_i = \langle s_i \rangle$ and $\chi_{ij} = \langle s_i s_j \rangle$ [19], as well as the special compositional constraint, $\sum_i s_i = 1$ (or 100%). The resulting solution P^* , obtained through the method of Lagrange multipliers, is given by:


$$P_{\Gamma}^* = Z^{-1} \exp \left[\sum_i \left(h_i + \frac{1}{2} \sum_{j \neq i} K_{ij} s_j \right) s_i \right],$$
$$Z = \int_{\sum_i^N s_i = 1} \exp \left[\sum_i \left(h_i + \frac{1}{2} \sum_{j \neq i} K_{ij} s_j \right) s_i \right] d\vec{s} \quad (1)$$

Here h_i and K_{ij} enforce, respectively, the means M_i and the covariances $\chi_{ij} - M_i M_j$. The normalizing constant Z is defined by an intractable integral over the simplex. Thus, the model parameters are found using an adapted pseudolikelihood approximation (see Methods: The simplex pseudolikelihood method). Finally, as $\sum_i s_i = 1$, several constraints are redundant. Thus, we set $h_N = 0$ and K_{ii} ($i = 1, 2, \dots, N$) to 0 (see Methods: Refining the maximum entropy parameters).


In summary, Eq 1 provides the Compositional Maximum Entropy model (CME) subject to known means and covariances. The CME method provides interpretable influence weights h_i for each component node i as well as the interaction strengths K_{ij} between each pair of components (i and j). Below, we provide two proofs of principle of the method: in a model of the abundances of co-evolving species and the analysis of gene expression data in cancer.

Fig 1. The Compositional Maximum Entropy (CME) approach. (a). Through maximum entropy, CME infers the unknown generative model of the observed component abundances. (b). h_i embodies the influence of each (i) component. Components with large h_i tend to have higher abundances than those with small h_i . K_{ij} embodies the interaction between pairs of components. Pairs with $K_{ij} > 0$ tend to coexist, while pairs with $K_{ij} < 0$ tend to be mutually exclusive.

(a)



(b)



Quantifying competition among co-evolving species

The quantification of competition among bacteria in the gut, market forces in the economy (or even among scientists) is of course of great interest. A simple and widely-used mechanism is provided by the competitive Lotka-Volterra model (cLV), which describes the population dynamics (i.e., the abundances) of different species vying for a shared resource [20–22]. The population (\tilde{s}_i) of each species i depends on its growth rate r_i and interaction α_{ij} with each other species j . Furthermore, the population of each type stops growing as it nears its carrying capacity κ_i , representing the complete exhaustion of resources.

$$\frac{d\tilde{s}_i}{dt} = r_i \tilde{s}_i \times \left(1 - \frac{\sum_j \alpha_{ij} \tilde{s}_j}{\kappa_i}\right) \quad (2)$$

While cLV remains a powerful model for predicting population dynamics, several challenges remain in calibrating it to experimental data. First, we are often only provided with relative (normalized) species abundances. Tools handling both this information loss and the resulting compositional data remain problematic [2, 23]. In addition, we rarely have access to the full time series. Bacterial abundances, for example, are typically measured sparsely but across many different conditions and environments [2].

Here we show that CME can provide accurate quantitative estimates of inter-species interactions, as predicted by cLV, using only available experimental information. The

simulated cLV abundances \tilde{s}_i are first normalized to resemble experimental data:


$$s_i = \frac{\tilde{s}_i}{\sum_j \tilde{s}_j}, \quad i = 1, 2, \dots, N \quad (3)$$

The time-evolving relative abundances $s_i(t)$ are then randomly sampled to apply CME. Compared to the cLV model, our proposed approach requires fewer parameters that are thus more resolvable from the limited available data.

cLV models exhibit three broad classes of stable inter-species behaviors: mutualism (they coexist), neutralism (they ignore each other), and competition (only one type can exist at a time) [23]. To illustrate these behaviors, we consider a cLV model of three different species with equal interactions $\alpha_{ij} = \alpha$. Figure 2 shows the dynamics and abundance distributions for each of three different regimes: $\alpha = 0.6$ (mutualism, Fig. 2a), $\alpha = 1.2$ (neutralism, Fig. 2b), and $\alpha = 4.0$ (competition, Fig. 2c). For simplicity, r_i , κ_i , and the self-interactions α_{ii} are fixed at 1. Gaussian noise was then added to the simulated dynamics to introduce additional inter-sample variability.

The cLV model exhibits a sharp qualitative transformation in its abundance distribution, from a unimodal (Fig. 2a) to a trimodal (Fig. 2c) behavior, as α is increased above a critical value ($\alpha \approx 1.2$, Fig. 2b). Despite requiring fewer parameters ($h = 0$ and K , compared to the original four of cLV), CME (right) captures the cLV model behavior (center) across this transformation; $K > 0$ describes mutualism while $K < 0$ describes competition.

Fig 2. Simulated abundances of three co-evolving species under mutualism (a), neutralism (b), and competition (c). Left, the cLV simulated abundances of each of the three interacting species over time. Center, the corresponding abundance distribution (cLV). Right, the best fit maximum entropy distribution (CME).



In summary, our model provides a simple, data-driven framework for modeling inter-species relationships from limited experimental information. We next consider the more complex case involving heterogeneous interactions from gene expression data in cancer.

Revealing driving interactions in cancer networks

Cancer is a heterogeneous disease involving complex molecular interactions between many genes. Despite the wealth of information provided by modern experimental tools, the application of such molecular data, including gene expression, to identify effective drug targets continues to face two significant obstacles. First, the accuracy of experimental expression profiles differs between genes [24]. Thus influences from biologically critical but more poorly resolved genes may be overlooked. Second, genes of typical interest often interact, and their effects overlap [25].


Novel network analysis techniques have been developed to refine the genetic signatures of critical genes in cancer. These approaches have been utilized to discover feedback structures in gene interaction networks, identify hubs and bridges, and define measures of robustness and fragility [26–28]. The Wasserstein distance from optimal transport lies as the basis for such methodologies, and in addition to the above references has been directly applied to the stationary (normalized) measures of the networks in question to derive biological information, e.g. showing that pediatric sarcoma data forms a unique cluster [29]. We will now show that CME may provide an important tool for such problems and help point to potential driver genes and their most important interactions.

To test our method, we analyze whole-genome expression data of triple-negative breast tumors, a highly aggressive and complex type of cancer. While many genes are known to be dysregulated in this disease, the relative influence of individual genes is far from established [30]. The data consist of expression profiles from 299 disease samples in METABRIC (Methods: The METABRIC dataset) [31]. We obtained normalized weights for each of $N = 3147$ genes using the Human Protein Reference Database (HPRD) for each sample (Methods: Network identification) [32]. As most of these genes provide no signal in the data, we renormalized these weights after considering only the top 17 highest variability genes with known relevance to cancer (according to OncoKB, see [33] and Methods: Data preparation). Figure 3 illustrates the known connectivity of these genes, but with node size and color proportional to their inferred maximum entropy node weights (h_i). We immediately notice two key details. First, our genes of interest all form a tightly connected network. Second, despite being highly correlated with each other (as the topology would suggest), these genes have unequal influences on the data. The highest-ranked genes, SRC and TP53, are also known master regulators of cancer [34,35].

A major strength of maximum entropy methods is identifying key node-node interactions underlying the more complex covariances measured from data. This is illustrated in Figure 4, which compares the maximum entropy pairwise interactions K_{ij} to those inferred from a widely-used alternative statistical model, the logit-normal distribution (Methods: Implementation of the logit-normal distribution); there is an identifiable mapping between the strongest magnitude maximum entropy interactions (Fig. 4a), in contrast to these obtained from the logit-normal (Fig. 4b), and their corresponding gene-gene covariances (Fig. 4c).

We also note that the two top maximum entropy interactions alone (SRC/TP53 and BRCA1/PTPN11) provide an intuitive explanation for some of the key features of the data. SRC and TP53 maintain the critical balance between growth (SRC) and damage repair (TP53): enhanced SRC (or repressed TP53) promotes cell survival, growth, and

Fig 3. Maximum entropy ranking of key genes in triple-negative breast cancer. Edges correspond to protein-protein interactions obtained from HPRD. Node color and size correspond to their influence (h_i).



metastasis, while the reverse leads to accelerated aging [34–36]. This known and critical negative interaction between SRC and TP53 separates most of the 17 genes into two distinct (and negatively covarying) clusters. Thus, since BRCA1 and PTPN11 belong to opposing clusters, their corrected interaction, as revealed by both maximum entropy and logit-normal modeling, is much larger than expected from their weak, positive covariance. Interestingly, both BRCA1 and PTPN11, along with SRC and TP53, are involved in the JAK-STAT pathway [37, 38]. Thus, these genes may have a general and synergistic role in cancer that remains to be explored.

Yet, while the logit-normal model does appear to resolve some features (such as the subtle covariance between AKT1 and EP300) that CME neglects, the interactions predicted by this method generally appear difficult to interpret in the context of the original covariance matrix: it predicts many interactions between uncorrelated genes and fails to resolve, among others, the clear negative covariance between SRC and TP53. Overall, the CME method provides a parsimonious biological mechanism, involving known cancer drivers and only a few of their interactions, for the genetic variability in this poorly understood disease.

Discussion

We have provided CME, a probabilistic framework for inferring the behaviors of compositional systems from data. Typically, models are deduced bottom-up, starting from mathematical relationships between individual components and combining them often in a complex, nonlinear way. However, as we have described for the Lotka-Volterra model, these interactions can rarely be resolved from the available experimental data. CME, instead, takes a top-down approach – starting from the data and learning the most parsimonious model for it. As evidenced by our breast cancer analysis, CME may also provide more interpretable insights into the organization of compositional systems.

For simplicity, we have considered only small networks; however, our method can be easily extended to much larger networks. First, the pseudolikelihood approach at the core of our method has been successfully applied, with the proper regularization, to networks consisting of thousands of nodes [39]. Second, the implementation of our algorithm uses a scalable L-BFGS algorithm and is fully parallelized across multiple CPU cores.

Similar to partial correlational analysis [40–42], maximum entropy computes direct pairwise interactions by controlling for the confounding indirect effects of the other nodes. Despite being widely used in data analysis and machine learning, partial correlations are only appropriate for linear associations or Gaussian-like data [40]. Maximum entropy methods, such as our application to compositional data [41], are, by contrast, much more general.


Our approach can incorporate more general model constraints as well. The compositional simplex constraint is enforced using the method of Lagrange multipliers. Other geometries [43, 44], even higher-order moments, can be included simply by including new Lagrange multipliers. Finally, as a maximum entropy model, CME is naturally compatible with constraints on other types of data as well [45].

Conclusion


We proposed CME, a data-driven framework for modeling compositions in multi-species networks. We utilize maximum entropy, a first-principles modeling approach, to learn influential nodes and their network connections using only the available experimental information. Our method requires minimal assumptions and no modifications of the

Fig 4. Comparison between three breast cancer network analyses: CME (a), logit-normal (b), and the data covariances (c). Maximum entropy and logit-normal results are shown on a log-scale to reveal the most influential positive (red) and negative (blue) interactions.


(a)



(b)



(c)



experimental data. Furthermore, the method can be easily generalized to incorporate new types of constraints and data that may emerge. 197
198

Methods

The simplex pseudolikelihood method

Fitting maximum entropy models to data is generally computationally intractable. 201
Thus, to fit CME, we will adapt the widely-used *pseudolikelihood* approximation [39]. 202
This method requires two pieces of information. First, we need a formula to compute 203
the conditional distribution $P(s_i|s_{\sim i})$, where $s_{\sim i}$ represents all of the variables of 204
interest s_j ($j = 1, 2, \dots, N-1$) excluding s_i and $s_N = 1 - \sum_{i=1}^{N-1} s_i$. For the simplex 205
model, we have: 206

$$P(s_i|s_{\sim i}) = \tilde{Z}_i(s_{\sim i})^{-1} \exp \left[\left(\tilde{h}_i + \frac{1}{2} \tilde{K}_{ii} s_i + \sum_{j \neq i} \tilde{K}_{ij} s_j \right) s_i \right] \quad (4)$$

$$\tilde{Z}_i(s_{\sim i}) = \int_0^{1 - \sum_{j \neq i}^{N-1} s_j} \exp \left[\left(\tilde{h}_i + \frac{1}{2} \tilde{K}_{ii} s_i + \sum_{j \neq i} \tilde{K}_{ij} s_j \right) s_i \right] ds_i \quad (5)$$

$$\tilde{h}_i = h_i + \frac{1}{2} (K_{iN} + K_{Ni}), \quad \tilde{K}_{ij} = K_{ij} - K_{iN} - K_{Ni} \quad (6)$$

Unlike that of Eq 1, \tilde{Z}_i is a tractable Gaussian-like integral. However, its value is 207
sample dependent. Thus, the second required piece of information is the actual samples 208
of the agent proportions s_i^d ($d = 1, 2, \dots, D$) rather than simply the summary means and 209
covariances. Together these enable the maximization of the *pseudolikelihood* functions 210
 ℓ_{PL}^i (see Methods: Model implementation): 211

$$\ell_{PL}^i = \tilde{h}_i M_i + \frac{1}{2} \tilde{K}_{ii} \chi_{ii} + \sum_{j \neq i}^{N-1} \tilde{K}_{ij} \chi_{ij} - D^{-1} \sum_{d=1}^D \log \tilde{Z}_i(s_{\sim i}^d) \quad (7)$$

Refining the maximum entropy parameters


One challenge in modeling compositional data is handling the parameter redundancies 213
induced by the compositional constraint $\sum_i s_i = 1$. Specifically, M_N , χ_{iN} , and χ_{Ni} are 214
entirely determined from the other data constraints. We could set the associated 215
Lagrange multipliers to 0, but this would hide information about node N (as all of its 216
connections would be forced to 0). 217

Instead, we recover interpretable model parameters with the following 218
transformations: 219

$$K_{ij} = \frac{1}{2} (\tilde{K}_{ij} + \tilde{K}_{ji} - \tilde{K}_{ii} - \tilde{K}_{jj}), \quad h_i = \tilde{h}_i - K_{iN} \quad (8)$$

By forcing K_{ii} to be 0 in Eq 1, we can resolve the interaction strengths between all 220
pairs of nodes in the data. For simplicity, we have defined $h_N = 0$. However, we can 221
increase or decrease all h_i by any constant and still have an equally good fit. Thus we 222
introduce another transformation to facilitate intra-model comparison of these node 223
weights: 224

$$Q_i = \frac{e^{h_i}}{\sum_{i=1}^N e^{h_i}} \quad (9)$$



S1 Fig. Comparison of CME model covariances to the sample covariances of the cLV model. We observe complete agreement between our model and the data (see Figure 2), confirming the correctness of our maximum entropy fitting algorithm. The means, not shown, were all equal to 1/3 as expected.

Conceptually, the quotient Q_i compares the relative probability of observing a network configuration with influence dominated ($P^*(s_i = 1)$) by node i . We posit this as a useful comparison metric for future studies of compositional systems modeled under different conditions.

Model implementation

To provide a high-accuracy, low overhead approximate maximum of the CME log pseudolikelihood functions, we performed convex optimization using L-BFGS [46] augmented by automatic differentiation. To validate our method, we also designed a custom Monte-Carlo scheme to simulate from CME models. This scheme considers the fitted h_i and K_{ij} parameters and numerically estimates the corresponding means M_i and covariances $\Sigma_{ij} = \chi_{ij} - M_i M_j$. In contrast to CME, such simulation is prohibitively expensive for even moderately-sized, strongly-interacting networks. However, it enabled us to confirm the high accuracy of our model on our Lotka-Volterra simulations (see S1 Figure).

The METABRIC dataset

Microarray gene expression data for METABRIC were downloaded from the cBioPortal database [47, 48]. The METABRIC dataset, containing 1904 samples, is one of the most extensive publicly-available breast cancer studies [31]. We utilized microarray gene expression data containing 24368 genes from the 299 triple-negative samples.

Network identification

To quantify the (normalized) influence of genes relevant to triple-negative breast cancer, we utilized the method of network Markov chains [26–28]. The Human Protein Reference Database (HPRD) provides a curated interaction network of most human proteins [32]. Thus, to perform our analysis, we utilized the largest connected component, consisting of 3147 genes, obtained from the intersection of HPRD with the METABRIC gene list. We then performed network analysis as in [28] using the subset of 288 genes annotated in OncoKB, a curated database of prominent cancer genes [33].

Data preparation

For each sample, we obtain a measure of the relative influence of each of 288 genes. To identify potential drivers of the variability of these influences across the data, we computed their inter-sample Pearson correlations. We identified two distinct clusters of highly correlated genes: one containing a small number of immune-adjacent genes and the other, a much larger component, containing prominent breast cancer genes such as TP53 and BRCA1. Thus, we utilized only this second component for our analysis.

Our primary goal is to identify genes and their interactions that potentially drive the variability in treatment responses observed in triple-negative breast cancer [30]. Likely genes include only those with large influence and inter-subject variability. Upon computing the variance in the influence of each gene, we found 17 candidates with markedly higher variance than the remaining bulk. We thus renormalized node influence across these 17 prime candidates before performing our maximum entropy analysis.

Implementation of the logit-normal distribution

An alternative to CME, the logit-normal distribution is given by [1]:

$$P_{\Gamma} = Z_{LN}^{-1} \frac{1}{\prod_{i=1}^N s_i} e^{-\frac{1}{2} \left\{ \log \left(\frac{s_{\tilde{N}}}{s_N} \right) - M_{LN} \right\}^T \Sigma_{LN}^{-1} \left\{ \log \left(\frac{s_{\tilde{N}}}{s_N} \right) - M_{LN} \right\}} \quad (10)$$

where M_{LN} and Σ_{LN} are the means and covariances of the transformed data:

$y = \left[\log\left(\frac{s_1}{s_N}\right), \dots, \log\left(\frac{s_{N-1}}{s_N}\right) \right]$. Here, the feature of interest is the precision matrix

$K_{LN}^* = -\Sigma_{LN}^{-1}$ which, under fairly general circumstances, has been shown to approximate maximum entropy interactions [13]. As with CME, we then utilized Eq 8 to define symmetric interactions between all pairs of nodes rather than simply the first $N - 1$.

Funding

C.W. is funded by the Marie-Josée Kravis Fellowship in Quantitative Biology. A.R.T. is funded by AFOSR grants FA9550-17-1-0435, FA9550-20-1-0029, NIH grant R01AT011419. J.O.D. is funded by NIH grant R21-CA234752 and NIH/NCI Cancer Center Support grant P30 CA008748. J.O.D. and A.R.T. are funded by Breast Cancer Research Foundation Grant BCRF-17-193. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Code availability

The source code and data used to produce the results and analyses presented in this manuscript are available from Github Git repository:
https://github.com/Corey651/Compositional_Maximum_Entropy.

References

1. Aitchison J. The statistical analysis of compositional data. Journal of the Royal Statistical Society: Series B (Methodological). 1982;44(2):139–160.
2. Carr A, Diener C, Baliga NS, Gibbons SM. Use and abuse of correlation analyses in microbial ecology. The ISME journal. 2019;13(11):2647–2655.

3. Gloor GB, Macklaim JM, Pawlowsky-Glahn V, Egoscue JJ. Microbiome datasets are compositional: and this is not optional. *Frontiers in microbiology*. 2017;8:2224.
4. Zuev K, Eisenberg O, Krioukov D. Exponential random simplicial complexes. *Journal of Physics A: Mathematical and Theoretical*. 2015;48(46):465002.
5. Nelder JA, Mead R. A simplex method for function minimization. *The computer journal*. 1965;7(4):308–313.
6. Weiss S, Xu ZZ, Peddada S, Amir A, Bittinger K, Gonzalez A, et al. Normalization and microbial differential abundance strategies depend upon data characteristics. *Microbiome*. 2017;5(1):1–18.
7. Calle ML. Statistical analysis of metagenomics data. *Genomics & informatics*. 2019;17(1).
8. Jaynes E, et al. The maximum entropy formalism. Ed Levine, RD, Tribus, M, Where do we stand. 1979;
9. Jaynes ET. On the rationale of maximum-entropy methods. *Proceedings of the IEEE*. 1982;70(9):939–952.
10. Pressé S, Ghosh K, Lee J, Dill KA. Principles of maximum entropy and maximum caliber in statistical physics. *Reviews of Modern Physics*. 2013;85(3):1115.
11. Dixit PD, Wagoner J, Weistuch C, Pressé S, Ghosh K, Dill KA. Perspective: Maximum caliber is a general variational principle for dynamical systems. *The Journal of chemical physics*. 2018;148(1):010901.
12. Shore J, Johnson R. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on information theory*. 1980;26(1):26–37.
13. Weistuch C, Agozzino L, Mujica-Parodi LR, Dill KA. Inferring a network from dynamical signals at its nodes. *PLoS computational biology*. 2020;16(11):e1008435.
14. Weistuch C, Mujica-Parodi LR, Razban RM, Antal B, van Nieuwenhuizen H, Amgalan A, et al. Metabolism modulates network synchrony in the aging brain. *Proceedings of the National Academy of Sciences*. 2021;118(40).
15. Weistuch C, Mujica-Parodi LR, Dill K. The refractory period matters: unifying mechanisms of macroscopic brain waves. *Neural Computation*. 2021;33(5):1145–1163.
16. Schneidman E, Berry MJ, Segev R, Bialek W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*. 2006;440(7087):1007–1012.
17. Morcos F, Pagnani A, Lunt B, Bertolino A, Marks DS, Sander C, et al. Direct-coupling analysis of residue coevolution captures native contacts across many protein families. *Proceedings of the National Academy of Sciences*. 2011;108(49):E1293–E1301.
18. Ghosh K, Dixit PD, Agozzino L, Dill KA. The maximum caliber variational principle for nonequilibria. *Annu Rev Phys Chem*. 2020;71:213–238.

19. Merchan L, Nemenman I. On the sufficiency of pairwise interactions in maximum entropy models of networks. *Journal of Statistical Physics*. 2016;162(5):1294–1308.
20. Berry D, Widder S. Deciphering microbial interactions and detecting keystone species with co-occurrence networks. *Frontiers in microbiology*. 2014;5:219.
21. Marasco A, Picucci A, Romano A. Market share dynamics using Lotka–Volterra models. *Technological forecasting and social change*. 2016;105:49–62.
22. Stein RR, Bucci V, Toussaint NC, Buffie CG, Rätsch G, Pamer EG, et al. Ecological modeling from time-series inference: insight into dynamics and stability of intestinal microbiota. *PLoS computational biology*. 2013;9(12):e1003388.
23. Faust K, Raes J. Microbial interactions: from networks to models. *Nature Reviews Microbiology*. 2012;10(8):538–550.
24. Consortium S, et al. A comprehensive assessment of RNA-seq accuracy, reproducibility and information content by the Sequencing Quality Control Consortium. *Nature biotechnology*. 2014;32(9):903.
25. Paull EO, Aytes A, Jones SJ, Subramaniam PS, Giorgi FM, Douglass EF, et al. A modular master regulator landscape controls cancer transcriptional identity. *Cell*. 2021;184(2):334–351.
26. Sandhu R, Georgiou T, Reznik E, Zhu L, Kolesov I, Senbabaoglu Y, et al. Graph curvature for differentiating cancer networks. *Scientific reports*. 2015;5(1):1–13.
27. West J, Bianconi G, Severini S, Teschendorff AE. Differential network entropy reveals cancer system hallmarks. *Scientific reports*. 2012;2(1):1–8.
28. Zhu J, Oh JH, Deasy JO, Tannenbaum AR. vWCluster: Vector-valued optimal transport for network based clustering using multi-omics data in breast cancer. *Plos one*. 2022;17(3):e0265150.
29. Chen Y, Cruz FD, Sandhu R, Kung AL, Mundi P, Deasy JO, et al. Pediatric sarcoma data forms a unique cluster measured via the earth mover’s distance. *Scientific reports*. 2017;7(1):1–9.
30. Shah SP, Roth A, Goya R, Oloumi A, Ha G, Zhao Y, et al. The clonal and mutational evolution spectrum of primary triple-negative breast cancers. *Nature*. 2012;486(7403):395–399.
31. Curtis C, Shah SP, Chin SF, Turashvili G, Rueda OM, Dunning MJ, et al. The genomic and transcriptomic architecture of 2,000 breast tumours reveals novel subgroups. *Nature*. 2012;486(7403):346–352.
32. Peri S, Navarro JD, Amanchy R, Kristiansen TZ, Jonnalagadda CK, Surendranath V, et al. Development of human protein reference database as an initial platform for approaching systems biology in humans. *Genome research*. 2003;13(10):2363–2371.
33. Chakravarty D, Gao J, Phillips S, Kundra R, Zhang H, Wang J, et al. OncoKB: a precision oncology knowledge base. *JCO precision oncology*. 2017;1:1–16.
34. Irby RB, Yeatman TJ. Role of Src expression and activation in human cancer. *Oncogene*. 2000;19(49):5636–5642.

35. Finlay CA, Hinds PW, Levine AJ. The p53 proto-oncogene can act as a suppressor of transformation. *Cell*. 1989;57(7):1083–1093.
36. Niu G, Wright KL, Ma Y, Wright GM, Huang M, Irby R, et al. Role of Stat3 in regulating p53 expression and function. *Molecular and cellular biology*. 2005;25(17):7432–7440.
37. Gao B, Shen X, Kunos G, Meng Q, Goldberg ID, Rosen EM, et al. Constitutive activation of JAK–STAT3 signaling by BRCA1 in human prostate cancer cells. *FEBS letters*. 2001;488(3):179–184.
38. Liu X, Qu CK. Protein tyrosine phosphatase SHP-2 (PTPN11) in hematopoiesis and leukemogenesis. *Journal of signal transduction*. 2011;2011.
39. Stein RR, Marks DS, Sander C. Inferring pairwise interactions from biological data using maximum-entropy probability models. *PLoS computational biology*. 2015;11(7):e1004182.
40. Baba K, Shibata R, Sibuya M. Partial correlation and conditional correlation as measures of conditional independence. *Australian & New Zealand Journal of Statistics*. 2004;46(4):657–664.
41. Erb I. Partial correlations in compositional data analysis. *Applied Computing and Geosciences*. 2020;6:100026.
42. Williams DR, Rast P. Back to the basics: Rethinking partial correlation network methodology. *British Journal of Mathematical and Statistical Psychology*. 2020;73(2):187–212.
43. Aneiros G, Cao R, Fraiman R, Genest C, Vieu P. Recent advances in functional data analysis and high-dimensional statistics. *Journal of Multivariate Analysis*. 2019;170:3–9.
44. Mardia KV, Jupp PE, Mardia K. *Directional statistics*. vol. 2. Wiley Online Library; 2000.
45. Mo Q, Wang S, Seshan VE, Olshen AB, Schultz N, Sander C, et al. Pattern discovery and cancer gene identification in integrated cancer genomic data. *Proceedings of the National Academy of Sciences*. 2013;110(11):4245–4250.
46. Malouf R. A comparison of algorithms for maximum entropy parameter estimation. In: *COLING-02: The 6th Conference on Natural Language Learning 2002 (CoNLL-2002)*; 2002.
47. Cerami E, Gao J, Dogrusoz U, Gross BE, Sumer SO, Aksoy BA, et al.. The cBio cancer genomics portal: an open platform for exploring multidimensional cancer genomics data; 2012.
48. Gao J, Aksoy BA, Dogrusoz U, Dresdner G, Gross B, Sumer SO, et al. Integrative analysis of complex cancer genomics and clinical profiles using the cBioPortal. *Science signaling*. 2013;6(269):pl1–pl1.