Happy and Immersive Clustering Segmentations of Biological Co-Expression Patterns

Richard Tjörnhammar[a,b]

[a]KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden
[b]SciLifeLab, Tomtebodavägen 23, SE-171 65 Solna, Sweden

Abstract

In this work, we present an approach for evaluating segmentation strategies and solving the biological problem of creating robust interpretable maps of biological data by employing wards agglomerative hierarchical clustering applied to coexpression coordinates to deduce a faithful representation of the input.

We adopt and quantify two analyte-centric metrics named happiness and immersiveness, one for describing the suitability of a single analyte concerning the segmentation as well as a second metric for describing how well the segmentation catches the underlying data variation. We show that these two functions drive aggregation and segregation of segmentation respectively and can produce trustworthy segmentation solutions. We discover that the immersiveness metric exhibits higher-order phase transition properties in its derivative to cluster numbers.

Finally, we find that the cluster representations and label annotations, in the case with clusters of high immersiveness, correspond to compositionally inferred labels with the highest specificity. The interconnectedness mirrors the potential relationships between cluster representations, label annotations, and inferred labels, emphasizing the intricate nature of biology and the representation of the specific expressions of gene products.

Keywords: Agglomerative Hierarchical Clustering, Internal Metrics, Statistical Learning, Co-expression analysis, Compositional Analysis, Order-Disorder Transitions

Email address: `richardt@kth.se` (Richard Tjörnhammar)

Introduction

Background

One key feature of modern biotechnology platforms is the excessive amount of viable information present for a specific sample. In the fields of proteomics and transcriptomics, genetic products are quantified to assess the emergent properties of the studied systems. This can include both protein-coding as well as non-protein-coding gene products, where the former has traditionally been the main interest of many research endeavours [33, 32].

A large number of measured analytes, used to describe a single system, effectively means that the problem is overdetermined. Many techniques exists and are used to conduct dimensionality reduction while conserving the amount of useful information for describing the system state [2, 21, 36, 34, 18, 11]. These include a large array of, what today is thought of as machine learning, techniques. Some of the most widely used in bioinformatics that relates to our approach include, but are not limited to clustering techniques [7, 22], regression models or factorisation-based models [10] as well as classifiers such random forest models [5, 37] or can make use of batch correction approaches alongside clustering [20].

Our distance dependant formalism is thereby highly reminiscent of co-expression clustering [17] and other Principle Component Analysis (PCA) based methods [39]. We are not deducing a scale-free representation [3] in a directed fashion of the clustering solution such as is potentially done in Weighted Gene Co-Expression Network Analysis (WGCNA) methods [38, 25]. The main reason is that the hierarchical clustering approach, in general, cannot be assumed to describe free-scaling network modules.

In order not to calculate the pseudo-inverse of the expressions and use it to solve the biological problem for the sample encodings, we adopt a factorisation scheme. By diagonalisation and utilizing the insight that the factorised matrices are globally aligned component matrices of the full data we may directly attribute features with sample group properties. This can be realised by studying the SVD [11] method and realizing that a mean grouping of component factors in sample space corresponds to the mean grouped sample value before diagonalisation and that they are aligned with the feature components post-diagonalisation. From the decomposition, one can also deduce an adequate distance matrix representation that can be used for clustering.

Clustering is a rich field with many popular approaches that also include neighbor-based approaches such as the ones developed by Seurat [15], Cluster Finder [22] and Louvain or Leiden [31] clustering. Its long history stems from the need to coarse grain data and uncover hidden groups and data variation within a data

set [10]. Unsurprisingly, clustering has not only a long history but also a multi-faceted epistemological heritage with ties to many disciplines with applications in biological sciences using physical approaches [19, 1]. Most clustering approaches are employed to describe nonlinear effects and variations in data and thereby share a fundamental relationship with nonlinear PCA [26], and while being well known is still being actively studied [28].

In this work, we adopt Wards Agglomerative Hierarchical Clustering (AHC) [35] when working with real-world data. We have chosen Wards method since it is believed to produce segmentations with minimal spatial distortion and minimized inter-cluster variances [8] to explore the coarse-graining properties of the component distance hierarchies.

Clustering

For a general clustering problem, we want to deduce a segmentation strategy of K analytes into M cluster label segments.

To make a connection between system fragmentation and system information in a physical sense for clustering we define a system constituent state, x_i to have the probability p_i with the following relation

$$\sum_i p_i = 1 \tag{1}$$

so that we can form averages as

$$f_{avg}(\mathbf{x}) \quad = \quad <f> \quad = \quad \sum_i p_i f(x_i) \tag{2}$$

and define

$$\mathcal{Z}(\beta) \equiv \sum_i e^{-\beta f(x_i)} \tag{3}$$

as the normalisation criterion when the system is subject to an inverse fractionation probability $\beta \propto T^{-1}$ in equilibrium [9] and it is clear that this choice maximizes the entropy of system segmentation defined in the information space volume $\mathcal{I}$ [24]. We recognize the stationary system information energy score as

$$E = -\left(\frac{\partial}{\partial \beta} \log \mathcal{Z}(\beta)\right)_{\mathcal{I}} \tag{4}$$

thereby via

$$C_{\mathcal{I}} \equiv \left(\frac{\partial E}{\partial T}\right)_{\mathcal{I}} \tag{5}$$
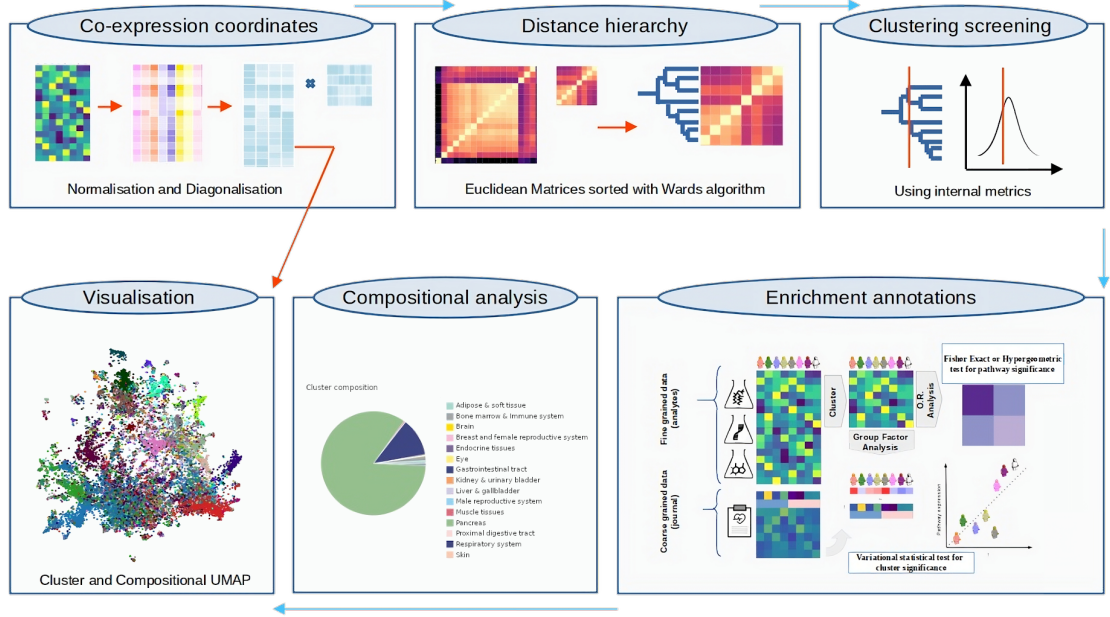
3

Figure 1: The publically available Biocartpgraph pipeline aims at delivering annotated clustering solutions of biological analyte data. This work is focusing on the screening evaluation section of the biocartographer

and using the fact that $\frac{\partial}{\partial T} = -\beta^2 \frac{\partial}{\partial \beta}$ yields

$$
\begin{aligned}
C_{\mathcal{I}} &= \beta^2 \frac{\partial^2}{\partial \beta^2} \log \mathcal{Z}(\beta) & (6) \\
&= \beta^2 \left(<f^2> - <f>^2 \right) & (7) \\
&= \beta^2 \sigma_f^2 & (8)
\end{aligned}
$$

which is recognizable as an information capacity for f times the fractionation contribution so that we can define

$$
C_f = \sigma_f^2 \qquad (9)
$$

We are assessing the success of the clustering segmentation strategy based on how well the cluster label inferences describe the underlying distance matrix, such as in [14]. An assumption in clustering Area Under the Reciever Operator Curve (AUROC) calculations is that the label inferences are related via a similarity measure, such as the covariational structure, of data. This is not necessarily true for an arbitrary segmentation strategy but is the standard assumption made in co-expression-based clustering. Here we employ distance measures instead of

4

similarity and the assessment approach in [14] takes the form of Figure 2. The covariation as well as the clustering label inference problem can be modelled using several different approaches and we define the coexpression coordinates to be calculated from the PCA scores using a Euclidean distance metric. The useful dimension of the PCA coordinates correspond to the minimal extension of the input data ($\mathbf{X}$) dimensions so that the number of components becomes $n_C = \min \dim \mathbf{X} - 1$. Which, for our datasets, are determined by the number of samples and sufficient to model all the information present in the datasets [27].

For our expression data matrix $\mathbf{X}$, containing K features and L samples, we define the standardised values $\mathbf{Z} = \frac{\mathbf{X} - \mu_i}{\sigma_i}$ such that we assume $Z_{ij} \in \mathcal{N}(0, 1)$, implying that μ_i, σ_i are the expression sample means and standard deviation values respectively. Following $\mathbf{Z}$ our diagonalised component matrices becomes

$$\mathbf{U}\mathbf{S}\mathbf{V}^T = \mathbf{Z} \tag{10}$$

where $\mathbf{P}$ defines the absolute coordinates of the covariation matrix as in equation (11)

$$\mathbf{P} = \mathbf{U}\mathbf{S} \tag{11}$$

From this, it is apparent that the covariation matrix can be calculated as

$$\mathbf{C} = \frac{1}{N - 1}\mathbf{P}\mathbf{P}^T \tag{12}$$

that we employ to calculate a covariation distance given by the Euclidean distance between covariation coordinates

$$D_{ij}^2 = \sum_{k=0}^{n_C}(\mathbf{P}_{ik} - \mathbf{P}_{jk})^2 \tag{13}$$

This distance matrix is then further employed in an AHC step that we screen to deduce a segmentation cut for the clustering representation.

Various dimensionality reduction techniques can be employed to coarse-grain the system and reduce the complexity of the studied problem [2, 23, 18]. Here we employ Uniform Manifold Approximation and Projection (UMAP) transformations to the co-expression coordinates before clustering real-world data. The clustering approach is chosen to be a hierarchical approach in that it creates well-connected segmentations. For any cut through the hierarchy, we can directly determine the compositions and AUROC metrics. This implies that once a segmentation strategy has been chosen we can assess when the underlying variation is no longer markedly improved by additional clusters. This sets the number of cluster states in our model.
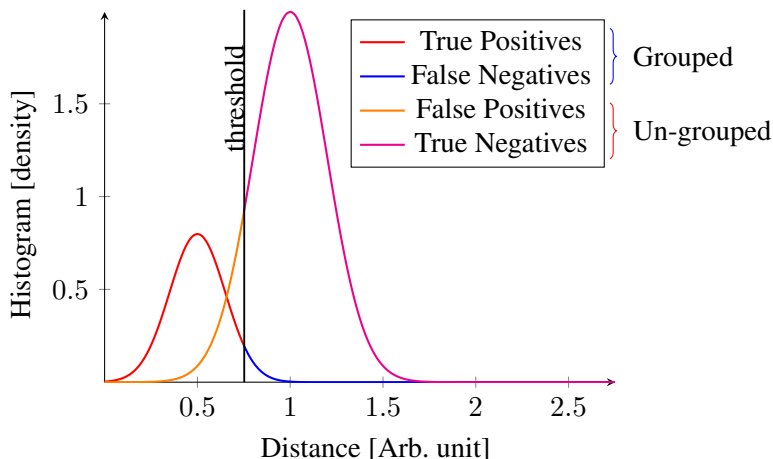
5

Figure 2: Schematic of unsupervised multilabel inference evaluation

Since the chosen segmentation across analytes can be compositionally studied we may deduce what compositional component is in majority as well as its compositional specificity. This forms a foundation for compositionally auto-annotating the analyte feature. In the case where sample label groups are good approximatives of the underlying covariation this annotation corresponds to a compositionally inferred attribution, with high AUROC values. We have chosen to employ the Gini coefficient, in place of many other available compositional coefficients [16, 12], mainly because of its broad usage but also because it is believed to share an approximate relationship with the clustering-based Area Under the Curve (AUC) $\approx \frac{gini+1}{2}$ [14].

Now we present the two new validation metrics for the variationally linked segmentation strategies for the attribution of high-dimensional biological data. Throughout this text, the input data $\mathbf{X}$ corresponds to publicly available Ribonucleic Acid sequencing (RNA-seq) expression data coming from any of five datasets that are explicitly stated together with figures belonging to them.

1. Method

The author developed the software package Biocartograph [1] [30], see Figure 1 for evaluation and automatic annotation of clustering solutions of biological data.

Since we employ AHC on diagonalised expression data the critical assumption in all the methodologies is the conceptual dependence on the ability of the coordinates to recreate a covariation distance matrix. This distance matrix is employed in an agglomerative hierarchical clustering step, where the full hierarchy is determined and screened to deduce a viable starting position for the clustering representation. We employ a screening function, similar to what was done in previous work [29], but forsake the optimisation and use the formalism to fully evaluate the hierarchical agglomerative clustering solution space by offsetting the internal metrics of the analytes versus the clusters. We propose two new metrics (called happiness and immersiveness) for assessing per analyte relevances in the following two sections.

Happiness

For any analyte belonging to any cluster in the segmentation, we can deduce the cluster attribution of its N_e closest nearest neighbours. We postulate that the happiness of an analyte is defined in regards to how homogenous the neighbour cluster attribution of those closest neighbours is. We define the attribution function as $\nu(i)$, which returns the cluster label of an analyte i. The neighbour function $\eta(i,j)$ defines the neighbour j of analyte i. We also define the unit Dirac delta function as $\delta_{a,b}$. Then we can formulate the happiness of the analyte i as:

$$h(i) = \frac{1}{N_e} \sum_{j=1}^{N_e} \delta_{\nu(i),\nu(\eta(i,j))} \tag{14}$$

The value of N_e represents the closest environment and is chosen to be 20 in this text. To conduct the calculation over all neighbours the sum in Eq. 14 becomes the sum over all j in $\delta_{\nu(i),\nu(j)}$. The number of neighbours is always assumed to include itself in this work. As a result, we immediately understand that the happiness metric can be seen as an interior volume function of the clustering solution and will as a result minimize interfacial area and promote forming a large homogenous solution. The metric is complementary to the immersiveness and favours aggregation. The average happiness is defined as:

$$H = \frac{1}{K} \sum_{i=1}^{K} h(i) \tag{15}$$

---

[1] https://github.com/rictjo/biocarta

7

Immersiveness

The clustering segmentation AUROC is calculated in the same fashion as has been done in previous works [14, 13], but also extended into the immersion score by calculating an analogous AUC per analyte towards all others in the distance matrix. For analyte i, the i:th row of the distance matrix is used to construct the scores instead of the entire matrix. The full clustering AUROC of the entire solution is defined in terms of all analyte distances and calculated as in [13]. The only distinguishing property in the distance matrix is whether the distance pairs are grouped in a cluster or not. Thereby the immersion is defined as one analyte to all other similarity values but is similar to the full AUC calculation. We denote this immersion score, or per analyte (i) score, non-rank corrected AUC, $\iota(i)$. We can readily understand that any AUC function promotes many cluster labels as the cluster-id factors approach the underlying covariational structure that the distance matrix is calculated from. Thereby forming a monotonically increasing function versus the number of clusters in the segmentation solution (M) and promoting segregation. We define the average immersion as the immersiveness of the solution as:

$$I = \frac{1}{K} \sum_{i=1}^{K} \iota(i) \tag{16}$$

The immersiveness can be understood as a restricted AUC metric where the distance matrix row-to-row interactions are included as an average interaction.

Here we elaborate on the calculation of the $\iota(i)$. To describe the immersion $\iota(i)$ we make use of the finite unit Dirac delta function ($\delta_{ij} = \delta_{\nu(i),\nu(j)}$) expressed as an incidence matrix. We note that this matrix describes when any pair is grouped while $1 - \delta$ describes all ungrouped pairs. Now we have

$$\Delta_{ij}^{g} \quad = D_{i,j}^{grouped} \quad = D_{ij}\delta_{ij} \tag{17}$$

$$\Delta_{ij}^{u} \quad = D_{i,j}^{ungrouped} \quad = D_{ij}(1 - \delta_{ij}) \tag{18}$$

$$\tag{19}$$

as the distance decomposition for our multilabel clustering solution. To construct the evaluation of the case discrimination, as in Fig. 2, we employ the standard heaviside function $\Theta(x)$. This results in an expression for the True Positive Rate (TPR)

$$g(i,d) \quad = 1 + \sum_{j=1}^{K} \Theta(d - \Delta_{ij}^{g}) \tag{20}$$

$$TPR(i,d) \quad = g(i,d)/\int_{0}^{\infty} g(i,x)dx \tag{21}$$

8

as well as one for the False Positive Rate (FPR)

$$u(i,d) \quad = 1 + \sum_{j=1}^{K} \Theta(d - \Delta_{ij}^{u}) \tag{22}$$

$$FPR(i,d) \quad = u(i,d) / \int_{0}^{\infty} u(i,x)dx \tag{23}$$

Since these complementary metrics are expressed for all i and evaluated at the same distance we can form

$$\iota = \int_{0}^{1} TPR(FPR)d(FPR) \tag{24}$$

and note that the integrand ROC curve, in equation (24), is our sought information functional f for use in determining the information capacity, see equation (9).

Furthermore, we define that H is the average happiness of all analytes, at a hierarchy level ξ, while the I function is the average of the immersions of the segmentation analytes, $\partial H/\partial \xi \leq 0 \quad \forall \xi$ and that $\partial I/\partial \xi \geq 0 \quad \forall \xi$ with $\xi \propto T$. We thereby employ a weak meaning of the terms monotonically increasing or decreasing when assuming that the above derivative relationships are true respectively. The average immersiveness or AUC both favour segregation and a straightforward way of assessing the fitness of the clustering segmentation is via the AUC.

A unified segmentation score can be defined as:

$$\hat{G}(\xi,p,q) = \frac{\hat{H}(\xi)^{p}\hat{I}(\xi)^{q}}{\hat{H}(\xi)^{p} + \hat{I}(\xi)^{q}} \qquad \forall \xi, p, q > 0 \tag{25}$$

Where $\hat{H}$ and $\hat{I}$ are the max-min range normalized happiness and immersiveness functions respectively. The un-normalized functions are not guaranteed to form a unimodal $\hat{G}$ function, why it can only serve as a screening function once the full hierarchy solution has been obtained. Since the AUC computation is computationally expensive it also does not serve as a good optimisation goal function.

Compositional coefficients

In our context, the sample labels in the labelled data do not refer directly to the clustering inferred labels. It is important to note that compositional coefficients can be calculated for both. In the first case, compositional information from the samples can be attributed to cluster labels. In the latter, cluster attribution of sample labels is inferred because compositional variance is segmented on the cluster-ID. The same compositional metric can either refer to the specificity for each cluster versus sample attribution or the cluster attribution of each unique sample label. We need to distinguish between sample groupings and clustering labels in a compositional context. One pertains to the analyte attribution and the second to the sample attribution.

9

Aggregating samples on common labels and retaining the summed expressions, the mean expressions, or the normalized expression of those samples across labels results in markedly different fractional composition information for those labels. Here we make the further distinction between compositional coefficients calculated for different qualities of the groupings.

We perform the cluster composition calculation in three steps. The first is to aggregate from samples to labels via direct summation to obtain analyte expressions for the composition of sample labels. Secondly, we normalize the expressions for each analyte to make analyte contributions comparable. This is achieved by sum-normalization which ensures that highly and lowly expressed analytes retain their expression profile across the sample labels and their profiles contribute equally. This is done before finally aggregating a typical cluster profile via summation of analytes onto cluster labels.

2. Results

The overarching goal has been to find an unbiased rationale for representing the data well. We will study this in the context of hierarchical clustering and the segmentations produced by it. All of the level cuts through all the hierarchical segmentation strategies for five different data sets were evaluated. A suitable representation of the covariational structure was determined by studying both the information capacity as well as the unimodal cluster size-dependent function, in the form of equation (25). We will report on the the interplay of those metrics, the phase transition point and the screening function as well as the compositional impact, in this section.

We first focus on the derived information capacity calculated in accord with equation (9) in order to make an initial assessment of our modelling assumptions. This was conducted for several AHC linkage approaches using simulated data, see Fig. 3 for the complete linkage AHC example. In Fig. 3 A we depict the parts of a single immersion calculation as well as the full immersiveness decomposition. Our formalism allows us to directly calculate the standard deviation of both the TPR and FPR, which we can further employ to calculate the error associated with each AUROC value, depicted in Fig. 3 B. The resulting standard deviation of the immersiveness or average standard deviation of the immersions can via equation (9) be interpreted as an information capacity that obtains max value during the information transition. This point coalesces with the maximum TPR change subject to minimal change of the FPR. From Fig. 3 it is clear that our approach works for simulated data, see supplementary material for additional graphs and information. This approach was thereby further employed to study real-world datasets.
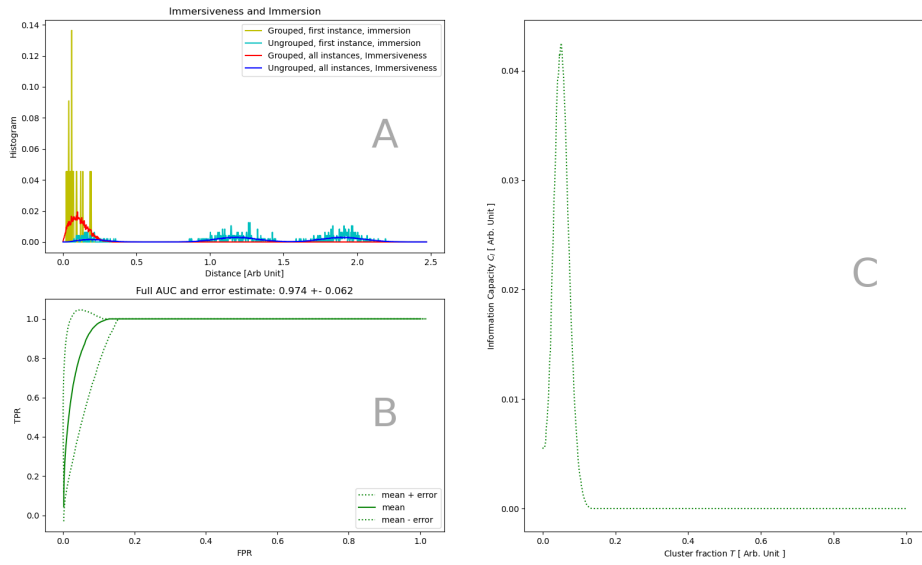
Figure 3: Synthetic data example employing agglomerative hierarchical clustering and complete linkage (maximum distance) method for cluster distance attribution. A) Depicting a typical clustering solution immersiveness together with the first $\iota(i = 1)$. B) From the standard deviation of the immersiveness from the ι we estimate the error bounds for the AUC. C) The average ι fluctuations σ_ι can be interpreted as an information capacity across a domain where 0 translates to a single cluster and 1 equates to a K cluster solution. We call this domain measure, which measures the degree of system fractionation, the temperature of the clustering solution.
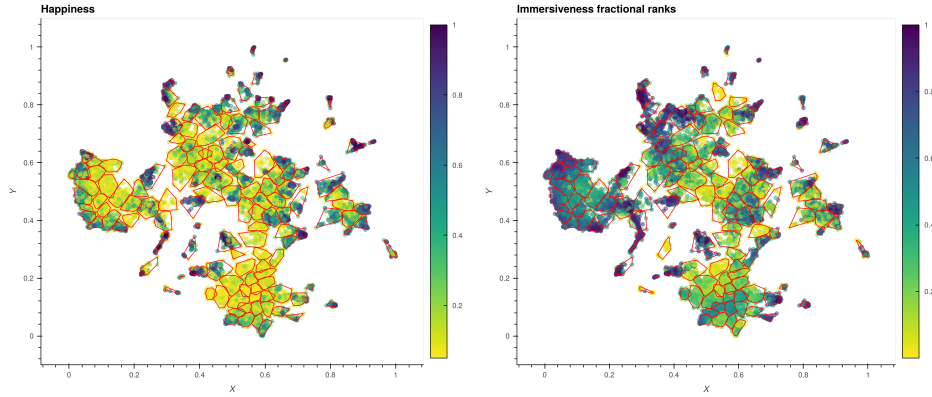
Figure 4: Single-cell data exploration example. The clustering UMAPS is colored by happiness and immersions to show their distributions respectively with convex hull borders of the clusters marked in red

From Fig. 4 it is clear that both the h and ι metrics describe a similar underlying property of the data, namely how well it serves to describe covariational distances. In either case, denser regions correspond to more distinct attribution labels and higher happiness and immersiveness metrics. The difference between both metrics resides in their respective ranges as well as scaling with segmentation. The averages (H, I) are determined as a function of the underlying analyte parameters (h, ι), but since both are dependent on the number of clusters in the segmentation (M) the scaling against the total amount of clusters can be expected to be different. In Fig. 4 we have chosen a cut through the distance hierarchy that yields $M = 139$ clusters (for the single-cell data set). The choice of M is not arbitrary, but note that only the full hierarchy preserves the distance- and topologic information of the full dataset. The value comes from the amount of clusters that will balance the amount of clusters against their mean sizes, see Fig. 6: Bottom-Right [29]. Since the happiness values are more distinct than the individualized immersion values we have chosen to calculate and show the rank normalized immersions for visualisation purposes, in Fig. 4. This shows that the happiness and the immersion scores model a similar data property but that the quality of the property is not apparent for a given segmentation representation. In this section, we want to address the quality of the metrics. The immersiveness is expressed as the mean of all immersions and therefore does not treat all pairs on an equal footing during the rank assessment of similarity (or distance) values. This implies that the immersiveness is limited in resolution as compared to the full AUROC calculation. In Fig. 5: A we see that the same behaviour is captured by both metrics, but that the immersiveness value is lower and smoother as compared to the AUROC. In Fig. 5: B) we can readily
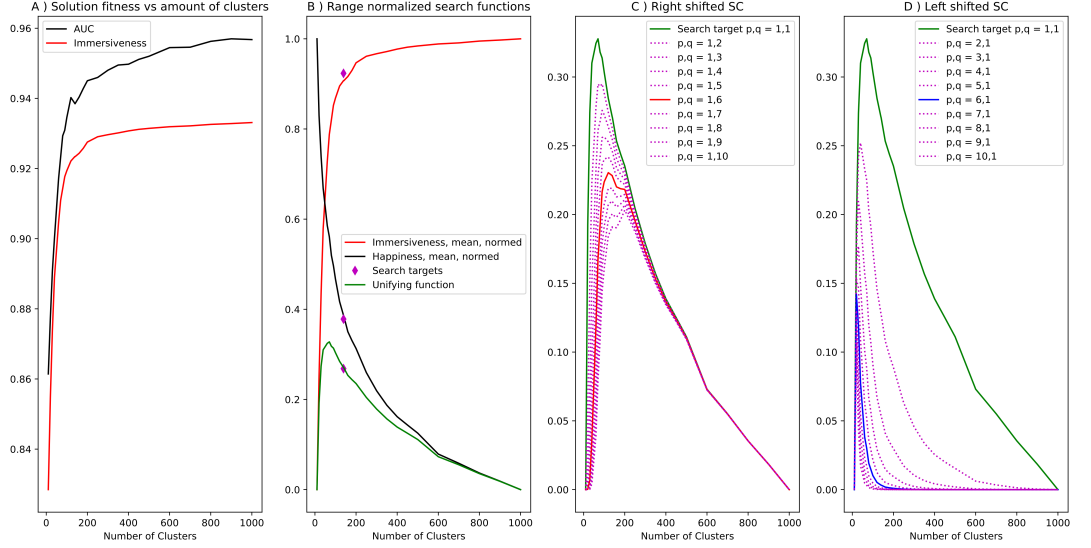
12

Figure 5: Single-cell data exploration example. A) Showing the AUC and the immersiveness (average immersion) for the different clustering solutions. B) Showing the range normalized mean happiness and mean immersion functions together with their unifying $\hat{G}$ function. The search targets are the chosen cut with $M = 139$ clusters. C) The right-shifted $\hat{G}$ functions under increasing powers of q. D) The left shifted $\hat{G}$ function under increasing powers of p.

see the scaling behaviour of the two suggested metrics in relation to the number of clusters in the segmentation solution. The H is monotonically decreasing with M while the I function is monotonically increasing. We can also see that the range normalized functions can be used to devise a unimodal search function but that the naive assumptions underlying such a formulation do not correspond to a high AUROC score. We can however use the p, q parameters to shift the peak to either emphasise happiness or immersiveness, see Fig. 5 C and Fig. 5 D.

Since the immersiveness is monotonically increasing but exhibits a stark derivative magnitude dependence versus the number of clusters one might expect a phase transition between the high and low AUC cluster values. We can use the average of the derivatives of the immersions to obtain a better numerical sampling of the immersiveness derivative to capture this behaviour. The behaviour of the average of the immersion functions derivatives is depicted in Fig. 6. The transition point coalesces with a screening function that uses equation (25), with $p, q = 1, 1$, but offsets the number of clusters against the mean size of those clusters, as done

13

in [29]. This tendency is depicted in Fig. 6 where we have assessed how the average derivative of the immersion values scales with respect to the number of clusters. It is clear from Fig. 6 that there is an abrupt change in the individualized analyte AUC values and that this peak position coalesces with one found by using a unimodal search function that offsets the size of clusters versus the number of clusters in the segmentation cut through an agglomerative hierarchical clustering solution that employs Ward's partitioning strategy. We can see in Fig. 5 that both the AUC and immersiveness values are monotonically increasing functions but that the gain over increasing cluster numbers drops at higher cluster numbers. These two regimes, where the gain is large versus small, are well separated by both the numerical behaviour of the average immersion derivatives as well as the size-dependent G function. To conclude we can see that the direct numerical derivative disruption of the immersiveness follows the position of the information capacity as well as the unimodal size-dependent search function. We can observe that the immersiveness and happiness metrics drive segregation and aggregation respectively. This observation is aligned with our theoretical understanding of the metrics.

Composition

In this section, we want to address whether the stated relationship between the clustering AUROC measure and the compositional Gini coefficient holds for our labelled data. Furthermore, the compositional dependence can be explored for the different data sets to increase our understanding of the connections between composition and covariation. In Fig. 7 we can see that there is no strong direct relationship between the absolute compositional coefficients [16] and the immersion values. We can however see that higher compositional coefficients, on average, correspond to higher immersion values. This weak dependence could be understood in that the covariational distances are relative metrics calculated for a clustering solution. All the compositional coefficients are calculated atomically for each label grouping per analyte without dependence on the segmentation so the result is expected. From Fig. 7 we find that many of the clusters analytes belonging to clusters with a high gini coefficient contain analytes that are close in absolute compositions label profiles. The clusters with the highest gini coefficients ($\approx 5\%$ of clusters) have absolute, fractional and sum-normed compositions that are similar in group label profiles (sample qualities). This disproves the notion that clusters with high immersion scores also obtain high Gini coefficients. This was explored and found to be a consistent observation for all datasets and we show the typical coefficient fall-off behaviour and compositional make-up in Fig. 8. For each cluster, we have chosen to calculate the typical compositional character of each cluster. The depiction in Fig. 8 illustrates what the compositional profile of a randomly drawn analyte would look like. The fractions thereby do not relate directly to an absolute
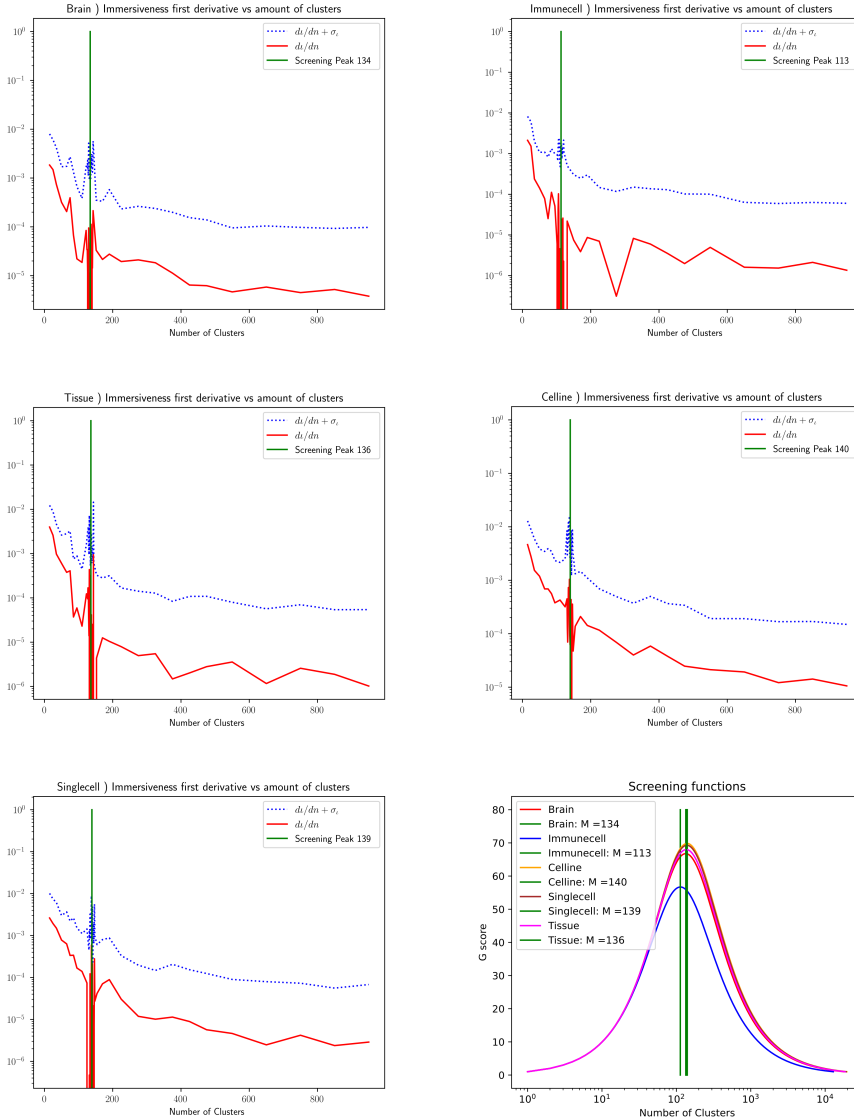
Figure 6: X) The average of the immersion derivatives and the average with the first standard deviation added in ($X \in \{$Brain, Immune-cells, Single-cells, Cell-lines, Tissues$\}$). The largest derivative numerical instability occurs around $M = 139$ clusters for the Single-cell dataset. Screening functions) The set of employed full unimodal, cluster size dependent, screening functions. These functions only assess the size of clusters and the number of clusters. Peak max value occurring for the Single-cell data at $M = 139$ clusters.
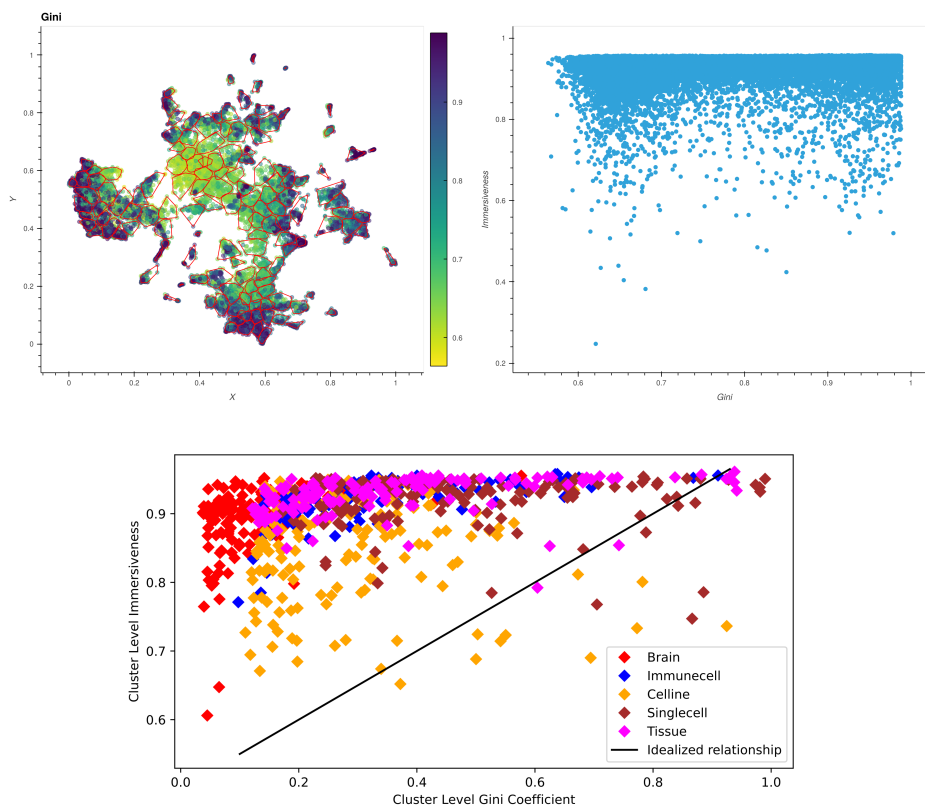
Figure 7: Single-cell data exploration example. Left) The Gini coefficient of each analyte's absolute composition in the input. Right) The Gini coefficient versus immersion shows no direct relationship. Bottom) The Cluster level aggregated comparisons between Immeriveness values and their respective Gini coefficients for the five data-set screened peak max position segmentation solutions.

number of label types for each cluster but rather what the most common expression distribution for any analyte belonging to that cluster would look like.

3. Discussion

The main finding of this article is the transition state behaviour of the immersiveness metric that is intimately related to the traditional AUROC of the clustering solution. We will elaborate on some of the details in this section.

Since the sample labels are not necessarily strongly linked with the underlying covariation one should not expect a strong relationship between the Gini coefficient and the immersion values for a general clustering approach on labelled data. This

is evident as being the case in our single-cell data, see Fig. 7 (Right) as well as Fig. 8.

The ι metric is costly to calculate for any segmentation step in the family of segmentations belonging to the full hierarchy. Generally, algorithmic scaling is similar to or worse for the immersion calculation than the efficient hierarchical construction. Furthermore, since numerical derivatives are inherently noisy, we could not use the average $d\iota/dn$ directly to efficiently construct a search function concerning AUC completeness during segmentation construction, however, the σ_ι^2 works well for this purpose. The disruption of the immersion derivatives coalescence with the cluster size screening implies that it might be a better candidate for direct optimisation strategies [29]. We note that the immersiveness calculation is more efficient by employing our formalism than calculating related measures such as the Davies-Bouldin score [6] or the Calinski-Harabasz index [4].

We note that the observed transition only applies to the type of segmentation strategy where the clustering is formed in a systematically connected fashion. Where the search function is guaranteed analytic by the agglomerative hierarchical approach. This is not necessarily the case for other popular clustering approaches [31]. The coalescence of the numerical disruption of the immersiveness derivatives and the size-dependent G function is expected to hold for other agglomerative hierarchical approaches than Wards, as explored in the supplementary information. The ι transition property was explored for both real-world and simulated data using several AHC approaches and is shown to work as an assessment for hierarchical clustering solutions.

The immersiveness or immersion score lets us decide on cluster numbers using the covariation structure of the data. A comparison with other internal metrics exists in the supplementary where compositionally inferred labels are compared to cluster segmentation labels. We could not use label set scores to decide on cluster numbers in a fashion directly connected to the covariational structure. Metrics such as mutual information and ARI score etc are strictly dependent on comparing label sets and not those labels' ability to model covariation thereby rendering them useless in our context. The Davies-Bouldin score and Calinski-Harabasz index, but not the Silhouette score or Dunn index, are expected to exhibit a similar transition state as the immersiveness evaluated here. The comparative study of the transition behaviour of those metrics could form the basis for a future study.

4. Conclusions

Comparing Fig. 7, Fig. 4 and Fig. 8 it is clear that the Gini coefficient describes the typical cluster composition well but cannot generally be used to re-create the AUC value belonging to the cluster group analytes.

17

For our clustering approach, we can find a transition point of the AUC that relates to the number of clusters we need to describe the data well. We see that this behaviour occurs in both synthetic as well as real-world data using a plethora of AHC linkage approaches. This has a bearing on the modelling of biological function since the minimal set of uniquely covarying groups of analytes forms a basis for solving the encoding problem while discarding tightly covarying contributions during solution construction. We can show that the coarse graning transition point coalesces with the strategy of finding a size distribution offset between cluster content and amount of clusters. Thereby potentially facilitating a faster and more efficient strategy for finding a meaningful representation of the data under study.

References

[1] H. Agrawal and E. Domany. Potts Ferromagnets on Coexpressed Gene Networks: Identifying Maximally Stable Partitions. *Physical Review Letters*, 90(15):158102, Apr. 2003. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.90.158102. URL https://link.aps.org/doi/10.1103/PhysRevLett.90.158102.

[2] E. Becht, L. McInnes, J. Healy, C.-A. Dutertre, I. W. H. Kwok, L. G. Ng, F. Ginhoux, and E. W. Newell. Dimensionality reduction for visualizing single-cell data using UMAP. *Nature Biotechnology*, 37(1):38–44, Jan. 2019. ISSN 1087-0156, 1546-1696. doi: 10.1038/nbt.4314. URL http://www.nature.com/articles/nbt.4314.

[3] A. D. Broido and A. Clauset. Scale-free networks are rare. *Nature Communications*, 10(1), Dec. 2019. ISSN 2041-1723. doi: 10.1038/s41467-019-08746-5. URL http://www.nature.com/articles/s41467-019-08746-5.

[4] T. Caliński and J. Harabasz. A dendrite method for cluster analysis. *Communications in Statistics*, 3(1):1–27, 1974. doi: 10.1080/03610927408827101.

[5] T. Claeys, M. Menu, R. Bouwmeester, K. Gevaert, and L. Martens. Machine learning on large-scale proteomics data identifies tissue and cell-type specific proteins. *Journal of Proteome Research*, 22(4):1181–1192, 2023. doi: 10.1021/acs.jproteome.2c00644. URL https://doi.org/10.1021/acs.jproteome.2c00644. PMID: 36963412.

[6] D. L. Davies and D. W. Bouldin. A cluster separation measure. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-1(2):224–227, 1979. doi: 10.1109/TPAMI.1979.4766909.

[7] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, KDD'96, page 226–231. AAAI Press, 1996.

[8] A. Fernández and S. Gómez. Versatile Linkage: a Family of Space-Conserving Strategies for Agglomerative Hierarchical Clustering. *Journal of Classification*, 37(3):584–597, Oct. 2020. ISSN 0176-4268, 1432-1343. doi: 10.1007/s00357-019-09339-z. URL http://link.springer.com/10.1007/s00357-019-09339-z.

[9] A. Ghavasieh and M. De Domenico. Diversity of information pathways drives sparsity in real-world networks. *Nature Physics*, Jan. 2024. ISSN 1745-2473, 1745-2481. doi: 10.1038/s41567-023-02330-x. URL `https://www.nature.com/articles/s41567-023-02330-x`.

[10] M. Girolami. Mercer kernel-based clustering in feature space. *IEEE Transactions on Neural Networks*, 13(3):780–784, May 2002. ISSN 1045-9227. doi: 10.1109/TNN.2002.1000150. URL `http://ieeexplore.ieee.org/document/1000150/`.

[11] G. H. Golub and C. Reinsch. Singular value decomposition and least squares solutions. *Numerische Mathematik*, 14(5):403–420, Apr. 1970. ISSN 0029-599X, 0945-3245. doi: 10.1007/BF02163027. URL `http://link.springer.com/10.1007/BF02163027`.

[12] J. Górecki, M. Hofert, and M. Holeňa. Kendall's Tau and Agglomerative Clustering for Structure Determination of Hierarchical Archimedean Copulas. *Dependence Modeling*, 5(1):75–87, Jan. 2017. ISSN 2300-2298. doi: 10.1515/demo-2017-0005. URL `http://www.degruyter.com/view/j/demo.2017.5.issue-1/demo-2017-0005/demo-2017-0005.xml`. Number: 1.

[13] J. A. Hanley and B. J. McNeil. The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology*, 143(1):29–36, Apr. 1982. ISSN 0033-8419, 1527-1315. doi: 10.1148/radiology.143.1.7063747. URL `http://pubs.rsna.org/doi/10.1148/radiology.143.1.7063747`.

[14] P. A. Jaskowiak, I. G. Costa, and R. J. G. B. Campello. The area under the ROC curve as a measure of clustering quality. *Data Mining and Knowledge Discovery*, 36(3):1219–1245, May 2022. ISSN 1384-5810, 1573-756X. doi: 10.1007/s10618-022-00829-0. URL `https://link.springer.com/10.1007/s10618-022-00829-0`.

[15] V. Y. Kiselev, K. Kirschner, M. T. Schaub, T. Andrews, A. Yiu, T. Chandra, K. N. Natarajan, W. Reik, M. Barahona, A. R. Green, and M. Hemberg. SC3: consensus clustering of single-cell RNA-seq data. *Nature Methods*, 14(5): 483–486, May 2017. ISSN 1548-7091, 1548-7105. doi: 10.1038/nmeth.4236. URL `http://www.nature.com/articles/nmeth.4236`.

[16] N. Kryuchkova-Mostacci and M. Robinson-Rechavi. A benchmark of gene expression tissue-specificity metrics. *Briefings in Bioinformatics*, page bbw008, Feb. 2016. ISSN 1467-5463, 1477-4054. doi:

10.1093/bib/bbw008. URL `https://academic.oup.com/bib/article-lookup/doi/10.1093/bib/bbw008`.

[17] G. G. Lemoine, M.-P. Scott-Boyer, B. Ambroise, O. Périn, and A. Droit. GWENA: gene co-expression networks analysis and extended modules characterization in a single Bioconductor package. *BMC Bioinformatics*, 22 (1):267, Dec. 2021. ISSN 1471-2105. doi: 10.1186/s12859-021-04179-4. URL `https://bmcbioinformatics.biomedcentral.com/articles/10.1186/s12859-021-04179-4`.

[18] L. McInnes, J. Healy, and J. Melville. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. 2018. doi: 10.48550/ARXIV.1802.03426. URL `https://arxiv.org/abs/1802.03426`. Publisher: arXiv Version Number: 3.

[19] A. Murua, L. Stanberry, and W. Stuetzle. On potts model clustering, kernel k-means, and density estimation. *Journal of Computational and Graphical Statistics*, 17(3):629–658, 2008. ISSN 10618600. URL `http://www.jstor.org/stable/27594329`.

[20] A. F. Møller and J. G. S. Madsen. JOINTLY: interpretable joint clustering of single-cell transcriptomes. *Nature Communications*, 14(1):8473, Dec. 2023. ISSN 2041-1723. doi: 10.1038/s41467-023-44279-8. URL `https://www.nature.com/articles/s41467-023-44279-8`.

[21] A. Narayan, B. Berger, and H. Cho. Assessing single-cell transcriptomic variability through density-preserving data visualization. *Nature Biotechnology*, 39(6):765–774, June 2021. ISSN 1087-0156, 1546-1696. doi: 10.1038/s41587-020-00801-7. URL `https://www.nature.com/articles/s41587-020-00801-7`.

[22] G. Palla, I. Derényi, I. Farkas, and T. Vicsek. Uncovering the overlapping community structure of complex networks in nature and society. *Nature*, 435(7043):814–818, June 2005. ISSN 0028-0836, 1476-4687. doi: 10.1038/nature03607. URL `http://www.nature.com/articles/nature03607`.

[23] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, pages 2825–2830, 2011. URL `https://scikit-learn.org/`.

[24] M. Plischke and B. Bergensen. *Equilibrium Statistical Physics (2nd Edition)*, pages 48–51. World Scientific Publishing Company, 1994. ISBN 9810216416.

[25] P. Schlosser, J. Knaus, M. Schmutz, K. Dohner, C. Plass, L. Bullinger, R. Claus, H. Binder, M. Lubbert, and M. Schumacher. Netboost: Boosting-Supported Network Analysis Improves High-Dimensional Omics Prediction in Acute Myeloid Leukemia and Huntington's Disease. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 18(6):2635–2648, Nov. 2021. ISSN 1545-5963, 1557-9964, 2374-0043. doi: 10.1109/TCBB.2020.2983010. URL https://ieeexplore.ieee.org/document/9084262/.

[26] B. Schölkopf, A. Smola, and K.-R. Müller. Nonlinear Component Analysis as a Kernel Eigenvalue Problem. *Neural Computation*, 10(5): 1299–1319, July 1998. ISSN 0899-7667, 1530-888X. doi: 10.1162/089976698300017467. URL https://direct.mit.edu/neco/article/10/5/1299-1319/6193.

[27] V. Thibeault, A. Allard, and P. Desrosiers. The low-rank hypothesis of complex systems. *Nature Physics*, Jan. 2024. ISSN 1745-2473, 1745-2481. doi: 10.1038/s41567-023-02303-0. URL https://www.nature.com/articles/s41567-023-02303-0.

[28] A. Tirelli, D. O. Carvalho, L. A. Oliveira, J. P. De Lima, N. C. Costa, and R. R. Dos Santos. Unsupervised machine learning approaches to the q-state Potts model. *The European Physical Journal B*, 95(11):189, Nov. 2022. ISSN 1434-6028, 1434-6036. doi: 10.1140/epjb/s10051-022-00453-3. URL https://link.springer.com/10.1140/epjb/s10051-022-00453-3.

[29] R. Tjörnhammar. Clustering optimisation method for highly connected biological data, Aug. 2022. URL https://doi.org/10.48550/arXiv.2208.04720.

[30] R. Tjörnhammar. Biocartograph: Patecatl djungleclassifier for medical multilabel inferences, Dec 2023. URL https://doi.org/10.5281/zenodo.10420490.

[31] V. A. Traag, L. Waltman, and N. J. van Eck. From Louvain to Leiden: guaranteeing well-connected communities. *Scientific Reports*, 9(1):5233, Dec. 2019. ISSN 2045-2322. doi: 10.1038/s41598-019-41695-z. URL http://www.nature.com/articles/s41598-019-41695-z.

[32] M. Uhlen, C. Zhang, S. Lee, E. Sjöstedt, L. Fagerberg, G. Bidkhori, R. Benfeitas, M. Arif, Z. Liu, F. Edfors, K. Sanli, K. von Feilitzen, P. Oksvold, E. Lundberg, S. Hober, P. Nilsson, J. Mattsson, J. M. Schwenk, H. Brunnström, B. Glimelius, T. Sjöblom, P.-H. Edqvist, D. Djureinovic, P. Micke, C. Lindskog, A. Mardinoglu, and F. Ponten. A pathology atlas of the human cancer transcriptome. *Science*, 357(6352):eaan2507, Aug. 2017. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.aan2507. URL `https://www.science.org/doi/10.1126/science.aan2507`.

[33] M. Uhlén, L. Fagerberg, B. M. Hallström, C. Lindskog, P. Oksvold, A. Mardinoglu, A. Sivertsson, C. Kampf, E. Sjöstedt, A. Asplund, I. Olsson, K. Edlund, E. Lundberg, S. Navani, C. A.-K. Szigyarto, J. Odeberg, D. Djureinovic, J. O. Takanen, S. Hober, T. Alm, P.-H. Edqvist, H. Berling, H. Tegel, J. Mulder, J. Rockberg, P. Nilsson, J. M. Schwenk, M. Hamsten, K. von Feilitzen, M. Forsberg, L. Persson, F. Johansson, M. Zwahlen, G. von Heijne, J. Nielsen, and F. Pontén. Tissue-based map of the human proteome. *Science*, 347(6220):1260419, Jan. 2015. doi: 10.1126/science.1260419. URL `http://science.sciencemag.org/content/347/6220/1260419.abstract`. Number: 6220.

[34] H. J. van Veen, N. Saul, D. Eargle, and S. W. Mangham. Kepler mapper: A flexible python implementation of the mapper algorithm. *Journal of Open Source Software*, 4(42):1315, 2019. doi: 10.21105/joss.01315. URL `https://doi.org/10.21105/joss.01315`.

[35] J. H. Ward. Hierarchical Grouping to Optimize an Objective Function. *Journal of the American Statistical Association*, 58(301):236–244, Mar. 1963. ISSN 0162-1459, 1537-274X. doi: 10.1080/01621459.1963.10500845. URL `http://www.tandfonline.com/doi/abs/10.1080/01621459.1963.10500845`.

[36] L. Wasserman. Topological data analysis. *Annual Review of Statistics and Its Application*, 5(1):501–532, 2018. doi: 10.1146/annurev-statistics-031017-100045. URL `https://doi.org/10.1146/annurev-statistics-031017-100045`.

[37] C. Xiong, D. M. Johnson, R. Xu, and J. J. Corso. Random forests for metric learning with implicit pairwise position dependence. In *Knowledge Discovery and Data Mining*, 2012. URL `https://api.semanticscholar.org/CorpusID:8702888`.

[38] B. Zhang and S. Horvath. A general framework for weighted gene co-expression network analysis. *Statistical Applications in Genetics and Molecular Biology*, 4(1), 2005. doi: doi:10.2202/1544-6115.1128. URL `https://doi.org/10.2202/1544-6115.1128`.

[39] K. Zhang, N. R. Zemke, E. J. Armand, and B. Ren. A fast, scalable and versatile tool for analysis of single-cell omics data. *Nature Methods*, Jan. 2024. ISSN 1548-7091, 1548-7105. doi: 10.1038/s41592-023-02139-9. URL `https://www.nature.com/articles/s41592-023-02139-9`.

5. Supplementary information

Software

Environment

versioned R and Python nix environment definitions

Enter the environment using the nix package manager:

```
$ nix-shell versioned_R_and_Python.nix
```

Main package

```
https://github.com/rictjo/biocarta
or
https://github.com/richardtjornhammar/biocartograph
```

Installing

If you are not using nix then install the *Biocartograph* software using pip:

```
$ pip install biocartograph
```

The Biocartograph Single-cell Solution

Interactive Clustering Graphs

The Immersiveness, Happiness and Gini Cluster configuration

Graphs of Coefficients

Immersion derivatives

Single cell derivatives

All data-set derivatives

Screening functions Immersiveness (AUC + error estimates), several linkage methods

Information Capacity, several linkage methods

Toy data generation example

Cluster Composition Gini coefficients

Specificity values of single cell clusters

Cluster level Gini coefficients

Data availablity

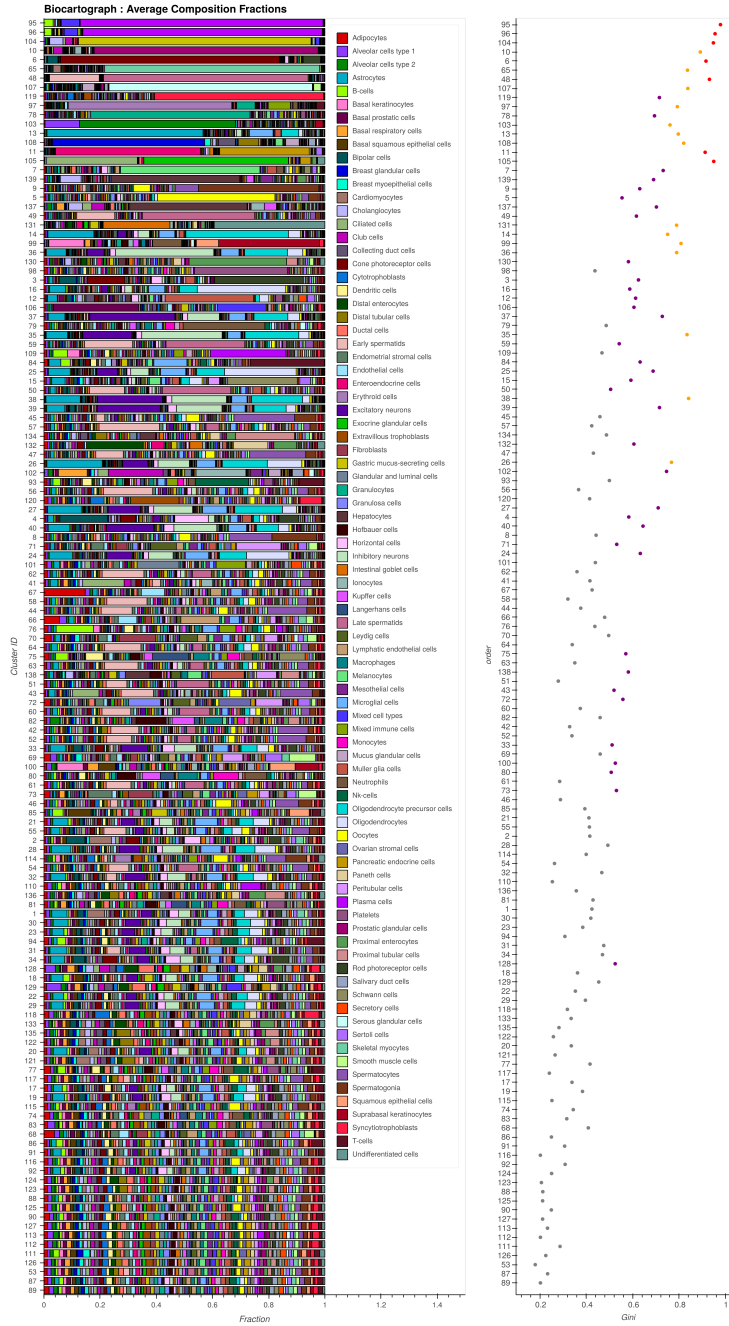Data can be obtained upon reasonable request to the author

Figure 8: Single-cell data exploration example. Left) Cluster average compositions. Right) Cluster Gini coefficients
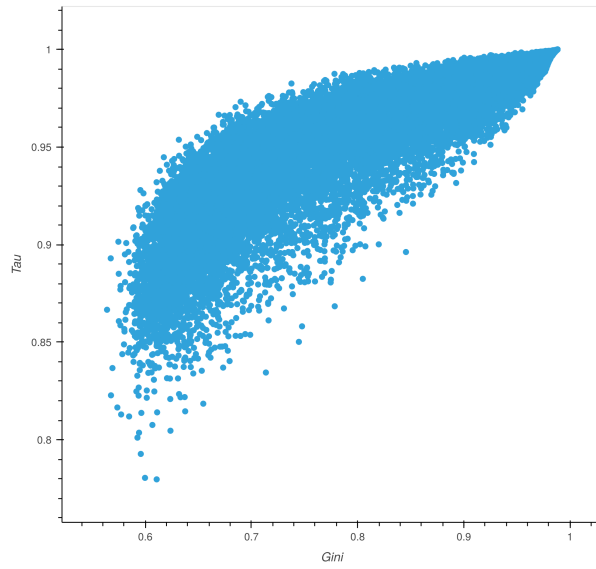
Figure 9: Single-cell data exploration example. The relationship between the (absolute) compositional Gini specificities and Tau specificities