

Quiver Laplacians and Feature Selection

Otto Sumray^{1,2}, Heather A. Harrington^{2,3,4,5}, Vidit Nanda²

ABSTRACT. The challenge of selecting the most relevant features of a given dataset arises ubiquitously in data analysis and dimensionality reduction. However, features found to be of high importance for the entire dataset may not be relevant to subsets of interest, and vice versa. Given a feature selector and a fixed decomposition of the data into subsets, we describe a method for identifying selected features which are compatible with the decomposition into subsets. We achieve this by re-framing the problem of finding compatible features to one of finding sections of a suitable quiver representation. In order to approximate such sections, we then introduce a Laplacian operator for quiver representations valued in Hilbert spaces. We provide explicit bounds on how the spectrum of a quiver Laplacian changes when the representation and the underlying quiver are modified in certain natural ways. Finally, we apply this machinery to the study of peak-calling algorithms which measure chromatin accessibility in single-cell data. We demonstrate that eigenvectors of the associated quiver Laplacian yield locally and globally compatible features.

1. Introduction

When working with large quantities of unstructured data, it is often convenient – and on occasion, indispensable – to employ a family of real-valued *features* defined on the dataset. It has, for instance, become standard practice within natural language processing to embed words into Euclidean space based on their relative proximity to each other in a substantial corpus of text [23]. In this case, the coordinates serve as features. Regardless of how such an embedding has been obtained, one immediately seeks to simplify matters by identifying those features which are most relevant to the task at hand. There are several ways of quantifying the relevance of a given feature — one may use domain-specific context, optimise a statistical quantity such as variance, or exploit the presence of labelled training data, etc. This process of *feature selection* forms a cornerstone of data analysis. Examples of well-known feature selectors include principal components [26], LASSO [32] and recursive feature elimination [15], to name but a few.

Email of corresponding author: otto.sumray@maths.ox.ac.uk

¹LUDWIG CANCER RESEARCH, UNIVERSITY OF OXFORD

²MATHEMATICAL INSTITUTE, UNIVERSITY OF OXFORD

³MAX PLANCK INSTITUTE OF MOLECULAR CELL BIOLOGY AND GENETICS

⁴CENTRE FOR SYSTEMS BIOLOGY DRESDEN

⁵FACULTY OF MATHEMATICS, TECHNISCHE UNIVERSITÄT DRESDEN

We focus here on feature selection in the presence of an auxiliary decomposition of the given dataset. As a toy example of the situation which we have in mind, let us consider the plots in Figure 1 of two features defined on datasets which have been decomposed into two overlapping subsets A and B .




FIGURE 1. Two features defined on a dataset equipped with a decomposition into two overlapping subsets.

Suppose we rank the relevance of features by their variance. In Figure 1.I, feature 1 is most relevant for subsets A , B , and $A \cap B$, but would not be more relevant than feature 2 on their union $A \cup B$. On the other hand in Figure 1.II, feature 1 is most relevant for A , and feature 2 for B , both are equally relevant for $A \cap B$ and $A \cup B$. Such unfortunate inconsistencies arise quite frequently, especially in biological contexts — we are often interested in features defined on datasets that are naturally organised by types, sub-types and so forth, such as disease or cell types. For such data, any feature selection process (which remains agnostic to the decomposition) is liable to select features that are irrelevant to certain subsets while ignoring features that are relevant to certain subsets.

This Paper. Here we present a principled framework for selecting features on a finite set X that are consistent with respect to a decomposition of X into a *cover* \mathcal{U} . Crucially, the subsets of X which comprise \mathcal{U} are allowed to admit arbitrary overlaps. Given the immense diversity of available feature selectors and their inner workings, we adopt an abstract approach. Let \mathbb{K} denote either the field \mathbb{R} of real numbers or the field \mathbb{C} of complex numbers. For our purposes, a feature selector \mathbf{S} is any deterministic process that

- (1) accepts as input some finite collection \mathcal{F} of functions $X \rightarrow \mathbb{K}$, and
- (2) outputs a subspace $\mathbf{S}_X^{\mathcal{F}}$ of the \mathbb{K} -linear span $\mathbf{V}(\mathcal{F})$ of \mathcal{F} .

We assume an inner product structure on $\mathbf{V}(\mathcal{F})$, but impose no further constraints on \mathbf{S} ; as such, the results and constructions of this paper may be applied verbatim to any of the standard feature selectors. Similarly, there are no constraints on the cover \mathcal{U} besides the requirement that $X = \bigcup_{U \in \mathcal{U}} U$.

For the purposes of these introductory remarks, let us call a nonempty subset $\sigma \subset \mathcal{U}$ *admissible* whenever the intersection $|\sigma| := \bigcap_{U \in \sigma} U$ is nonempty. Our quest here is to isolate the largest subspace of $\mathbf{S}_X^{\mathcal{F}}$ that is compatible with respect to the decomposition of X into admissible subsets. Two forms of compatibility are studied here: *local* and *global*.

Given an admissible subset $\sigma \in \mathcal{U}$, a feature $v : |\sigma| \rightarrow \mathbb{R}$ in $\mathbf{S}_{|\sigma|}^{\mathcal{F}}$ is locally compatible with \mathcal{U} if for every admissible $\tau \in \mathcal{U}$ with $|\tau| \supset |\sigma|$, there is some feature $v^+ : |\tau| \rightarrow \mathbb{R}$ in $\mathbf{S}_{|\tau|}^{\mathcal{F}}$ whose restriction to $|\sigma|$ coincides with v . In other words, a locally compatible feature is one which continues to be selected as we pass to larger admissible subsets. Conversely, consider a collection of features $\{v_\sigma\}$, one for each admissible σ . Such a collection is globally compatible with \mathcal{U} if for each pair $|\sigma| \subset |\tau|$ of admissible subsets, the orthogonal projection $\pi_\tau^{\mathcal{F}} : \mathbf{V}(\mathcal{F}) \rightarrow \mathbf{S}_{|\tau|}^{\mathcal{F}}$ sends v_σ to v_τ . Thus, every globally compatible feature is generated by functions $\{v_U : U \rightarrow \mathbb{R} \mid U \in \mathcal{U}\}$ such that the orthogonality relation

$$(v_U - v_{U'}) \perp \mathbf{S}_{U \cap U'}^{\mathcal{F}}$$

holds in $\mathbf{V}(\mathcal{F})$ for every $U, U' \in \mathcal{U}$.

The first question addressed in our work here is:

Fix a feature selector \mathbf{S} and a cover \mathcal{U} of X . Given an input \mathcal{F} family of functions $X \rightarrow \mathbb{R}$, to what extent can we efficiently discover all of the locally and globally compatible features output by \mathbf{S} ?

We approach this problem by recasting it in the language of quiver representations. The underlying quiver $Q^{\mathcal{U}}$ is a directed graph whose vertices are admissible sets, with a single edge $\sigma \rightarrow \tau$ whenever $|\sigma| \subset |\tau|$. The output of \mathbf{S} with input \mathcal{F} induces a representation $\mathbf{A}_\bullet = \mathbf{A}_\bullet^{\mathbf{S}, \mathcal{F}}$ of Q valued in the category of finite dimensional Hilbert spaces. Explicitly, each vertex σ is assigned the Hilbert space $\mathbf{A}_\sigma := \mathbf{S}_{|\sigma|}^{\mathcal{F}}$, while each edge $\sigma \rightarrow \tau$ is assigned the linear map $\mathbf{A}_{\sigma \rightarrow \tau} : \mathbf{A}_\sigma \rightarrow \mathbf{A}_\tau$ obtained by restricting $\pi_\tau^{\mathcal{F}}$ to $\mathbf{S}_{|\sigma|}^{\mathcal{F}}$. We show that the **sections** of this representation (as introduced in [30] and generalised to the multilinear setting in [24]), recover all the globally compatible features. Similarly, sections of this representation over certain augmented sub-quivers yield all the locally compatible features of \mathbf{S} .

The presence of noise severely impairs the compatibility framework described above. The main difficulty is that one never obtains an exact equality between (projections of) features selected over an admissible set $|\tau|$ and features selected over an admissible subset $|\sigma|$. The best-case scenario is that features in \mathbf{A}_σ lie *near*¹ the restriction of some feature in \mathbf{A}_τ . In fact, exact equality may not even be desirable — often, features are highly correlated, and collapsing together such features will further reduce the dimension of the space of selected features. Thus, there are many compelling reasons to seek a less rigid compatibility framework built around approximate sections of quiver representations. With this consideration in mind, the second question addressed in this paper is:

To what extent can we define and efficiently compute the approximate sections of a quiver representation?

We address this challenge by introducing a new Laplacian operator.

The **quiver Laplacian** L introduced here is a (Hermitian, positive semidefinite) endomorphism of the *total space* $\text{Tot}(\mathbf{A}_\bullet) := \prod_\sigma \mathbf{A}_\sigma$. In the special case where every \mathbf{A}_σ is

¹Here proximity is determined by the inner product metric on $\mathbf{V}(\mathcal{F})$.

one-dimensional and all orthogonal projection maps are identities, L coincides with the ordinary graph Laplacian of $Q^{\mathcal{U}}$. We note that Laplacians have been defined for several discrete algebraic/combinatorial objects: from the original case of graphs [3], to simplicial complexes [12, 17], hypergraphs [2] and principal bundles on graphs [10]. Most closely related to our quiver Laplacian are the analogous operators defined for cellular sheaves in [16] (but see also [11, 25]). The space of sections of a quiver representation coincides precisely with the kernel of the affiliated quiver Laplacian. Armed with the quiver Laplacian, we are able to define, for every $\epsilon > 0$, the ϵ -approximate sections of a quiver representation as those normalised vectors $x \in \text{Tot}(\mathbf{A}_{\bullet}^{\mathcal{S}, \mathcal{F}})$ which satisfy $\langle Lx, x \rangle \leq \epsilon$. Thus, the spans of eigenvectors of L associated to small nonzero eigenvalues of L delineate special families of approximate sections. Motivated by these considerations, we establish a host of fundamental properties for quiver Laplacians and their spectra.

Finally, we use the spectrum of an appropriate quiver Laplacian to study feature selection (as given by a *peak-calling* algorithm) for chromatin accessibility² data from [28]. Our dataset comes from single-cell sequencing. This type of data has several properties which make the Laplacian-driven approach particularly appropriate:

- (1) The dimensionality is massive — a naïve estimate assumes that every genomic position might be accessible, giving around 3 billion features. Therefore, feature selection (via peak-calling) becomes a crucial intermediate step towards making the analysis tractable.
- (2) Different cell types have different relevant features, and it is important to make the feature selection process compatible with the types of cells under consideration. This makes it necessary to consider peaks not only for the entire dataset, but also for certain distinguished subsets.
- (3) In many relevant cases where the data arise from diseases such as cancer or from the immune system (e.g. T cells), the cells might exhibit plasticity and their type is therefore not well-defined. In such cases, our distinguished subsets might overlap and should be modelled by a cover rather than a partition.
- (4) There is some stochasticity in the output of peak-calling algorithms — in particular, the accessible genomic positions are only known approximately. Thus, the accessibility status of nearby positions is not independent; we capture this phenomenon via the inner product structure on feature space.

Summary of Results. Let Q be a finite quiver with vertex set Q_0 , edge set Q_1 and maps $s, t : Q_1 \rightarrow Q_0$ which specify the source and target vertices of each edge. We consider representations of Q — each such representation \mathbf{A}_{\bullet} assigns a finite-dimensional (real or complex) Hilbert space \mathbf{A}_v to every vertex $v \in Q_0$ and a linear map $\mathbf{A}_e : \mathbf{A}_{s(e)} \rightarrow \mathbf{A}_{t(e)}$ to every edge $e \in Q_1$. We write $L_{\mathbf{A}}$ for the Laplacian of \mathbf{A}_{\bullet} (see Definition 4.2 below

²Chromatin accessibility measures the extent to which regions within the genome (among a specific family of cells) are reachable by cellular machinery, such as transcription factors.

for details), and arrange its eigenvalues as

$$0 \leq \lambda_1(\mathbf{A}) \leq \lambda_2(\mathbf{A}) \leq \cdots \leq \lambda_n(\mathbf{A}),$$

where $n := \dim \text{Tot}(\mathbf{A}_\bullet)$ is the total dimension of \mathbf{A}_\bullet .

Here is a simplified version of our first main result, which provides upper bounds on Laplacian eigenvalues of a representation \mathbf{A}'_\bullet in terms of eigenvalues of \mathbf{A}_\bullet whenever there exists a family of vertex-indexed linear maps $\mathbf{A}_v \rightarrow \mathbf{A}'_v$.

THEOREM (I). *Let \mathbf{A}_\bullet and \mathbf{A}'_\bullet be two representations of Q , and consider any vertex-indexed collection $\tau = \{\tau_v : \mathbf{A}_v \rightarrow \mathbf{A}'_v \mid v \in Q_0\}$ of linear maps. Write q for the dimension of $\ker \tau$ viewed as a (block diagonal) map $\text{Tot}(\mathbf{A}_\bullet) \rightarrow \text{Tot}(\mathbf{A}'_\bullet)$. There exist positive constants α, β, γ – each dependent on τ – such that the inequality*

$$\lambda_k(\mathbf{A}') \leq \alpha \cdot \lambda_{k+q}(\mathbf{A}) + \beta \cdot \sqrt{\lambda_{k+q}(\mathbf{A})} + \gamma$$

holds for every k in $\{1, 2, \dots, n - q\}$.

The full statement of this result, along with explicit formulas for α, β and γ , is recorded in Theorem 5.2. If τ is a morphism of quiver representations $\mathbf{A}_\bullet \rightarrow \mathbf{A}'_\bullet$, then both β and γ vanish and we are left with a much simpler bound $\lambda_k(\mathbf{A}') \leq \alpha \cdot \lambda_{k+q}(\mathbf{A})$.

Our next main result takes the form of a *spectral stability* result for feature selectors. The main challenge here is that the total spaces of two representations $\mathbf{A}_\bullet^{\mathbf{S}, \mathcal{F}}$ and $\mathbf{A}_\bullet^{\mathbf{T}, \mathcal{F}}$ may have different dimensions, so even the number of eigenvalues of the corresponding Laplacians could be different. Nevertheless, let $\mu_{\mathbf{S}}$ and $\mu_{\mathbf{T}}$ be the probability measures on \mathbb{R} given by taking the average of Dirac measures concentrated at the eigenvalues $\{\lambda_i(\mathbf{A}_\bullet^{\mathbf{S}, \mathcal{F}})\}$ and $\{\lambda_j(\mathbf{A}_\bullet^{\mathbf{T}, \mathcal{F}})\}$, respectively. We obtain, in Corollary 5.5, the following bound on the 1-Wasserstein distance $W_1(\mu_{\mathbf{S}}, \mu_{\mathbf{T}})$ between these two measures.

THEOREM (II). *Let \mathbf{S} and \mathbf{T} be two feature selectors defined on a finite set X which is equipped with a cover \mathcal{U} . Fix a finite collection \mathcal{F} of functions $X \rightarrow \mathbb{K}$. For each admissible set σ , define*

$$c_\sigma := \left| \dim \mathbf{S}_{|\sigma|}^{\mathcal{F}} - \dim \mathbf{T}_{|\sigma|}^{\mathcal{F}} \right| \quad \text{and} \quad \epsilon_\sigma := \text{dist}_{\text{Gr}} \left(\mathbf{S}_{|\sigma|}^{\mathcal{F}}, \mathbf{T}_{|\sigma|}^{\mathcal{F}} \right),$$

where the latter expression invokes the Grassmannian distance between subspaces of $\mathbf{V}(\mathcal{F})$. Define $c := \sum_\sigma c_\sigma$ and $\epsilon := \max_\sigma \epsilon_\sigma$. There exist trigonometric polynomials f and g such that

$$W_1(\mu_{\mathbf{S}}, \mu_{\mathbf{T}}) \leq \frac{|Q_1^{\mathcal{U}}|}{m} \left[f(\epsilon) \cdot c + g(\epsilon) \cdot (m - c) \right],$$

where $|Q_1^{\mathcal{U}}|$ is the number of edges in $Q^{\mathcal{U}}$ and $m = \max\{\dim \text{Tot}(\mathbf{A}_\bullet^{\mathbf{S}, \mathcal{F}}), \dim \text{Tot}(\mathbf{A}_\bullet^{\mathbf{T}, \mathcal{F}})\}$.

We note that $c = 0$ if and only if the two feature selectors assign equidimensional spaces to each admissible set. Explicit formulas for f and g are provided in Section 5. Crucially, we have $g(0) = 0$; in the equidimensional case, the first term in our bound vanishes and the second term approaches 0 as $\epsilon \rightarrow 0$.

Our next main result provides several flavours of eigenvalue interlacing inequalities for quiver Laplacians which hold when quivers (and the overlaid representations) are modified in certain natural ways.

THEOREM (III). *Let \mathbf{A}_\bullet be a representation of Q . The following inequalities hold for appropriate integers i . The constants k, r, w_1, w_2 which appear below depend on \mathbf{A}_\bullet as well as the modification performed to create a new representation \mathbf{A}'_\bullet of a new quiver Q' .*

(1) *If Q' is obtained by removing edges from Q and \mathbf{A}'_\bullet is the restriction of \mathbf{A}_\bullet to Q' , then*

$$\lambda_{k+i}(\mathbf{A}_\bullet) \leq \lambda_{k+i}(\mathbf{A}'_\bullet) \leq \lambda_{k+r+i}(\mathbf{A}_\bullet).$$

(2) *If Q' is obtained by removing vertices (and incident edges) from Q and \mathbf{A}'_\bullet is the restriction of \mathbf{A}_\bullet to Q' then*

$$\lambda_i(\mathbf{A}_\bullet) - w_1 \leq \lambda_i(\mathbf{A}'_\bullet) \leq \lambda_{k+i}(\mathbf{A}_\bullet) - w_2.$$

(3) *If Q' and \mathbf{A}'_\bullet are the result of performing an admissible homotopy of Q , then*

$$\varphi^{-2}\lambda_i(\mathbf{A}_\bullet) \leq \lambda_i(\mathbf{A}'_\bullet) \leq \varphi^2\lambda_i(\mathbf{A}_\bullet)$$

where φ is the golden ratio.

(4) *If Q' and \mathbf{A}'_\bullet are the result of performing an admissible Kron reduction of \mathbf{A}_\bullet , then*

$$\lambda_i(\mathbf{A}_\bullet) \leq \lambda_i(\mathbf{A}'_\bullet) \leq \lambda_{i+r}(\mathbf{A}_\bullet).$$

Detailed versions of these results, including explicit formulas for various constants as well as precise definitions of admissible homotopy and Kron reduction, are located in Section 6. In order to state and prove these results in an appropriately general setting, we define *sheaves* on quivers — these are almost identical to cellular sheaves [6, 5, 16] on graphs, except that our graphs are allowed to have self-loops.

As mentioned above, much of our work here has been motivated by the desire to better understand and improve peak-calling algorithms on datasets X which come naturally equipped with a decomposition \mathcal{U} into subsets. For this purpose, it is convenient to combine the quiver representations whose sections give locally and globally compatible features of a given feature selector \mathbf{S} into a single representation of a larger quiver. This *combined representation*, denoted $\mathbf{N}_\bullet^{\mathbf{S}, \mathcal{F}}$, is described in Section 3.1. By design, its sections correspond to features which are both locally and globally compatible with \mathcal{U} along \mathcal{F} , and we naturally seek its approximate sections. Unfortunately, extracting spectral data for its Laplacian presents a serious computational challenge — both the underlying quiver and the total dimension are typically much larger than $Q^{\mathcal{U}}$ and $\mathbf{A}_\bullet^{\mathbf{S}, \mathcal{F}}$. The third main result of this paper is a remedy which takes the form of a *reduction theorem*.

THEOREM (IV). *Let \mathbf{S} be a feature selector on a set X equipped with a cover \mathcal{U} such that $Q^{\mathcal{U}}$ is weakly connected. Let $\mathbf{N}_\bullet^{\mathbf{S}, \mathcal{F}}$ be the combined representation associated to a finite collection \mathcal{F} of functions $X \rightarrow \mathbb{K}$. Fix a vertex $\sigma \in Q_0^{\mathcal{U}}$. There exists a new quiver Q' with a single vertex and $|Q_0^{\mathcal{U}}|$ edges along with a representation \mathbf{A}'_\bullet of Q' satisfying two properties:*

- (1) For each vertex τ of $Q^{\mathcal{U}}$, let id_τ be the identity map on $\mathbf{S}_{|\tau|}^{\mathcal{F}}$ and let $\iota_\tau^{\mathcal{F}} : \mathbf{S}_{|\tau|}^{\mathcal{F}} \hookrightarrow \mathbf{V}(\mathcal{F})$ denote the inclusion map (which is adjoint to $\pi_\tau^{\mathcal{F}}$). Then, the Laplacian of \mathbf{A}'_\bullet is

$$L_{\mathbf{A}'} = \sum_{\tau \in Q_0^{\mathcal{U}}} \left[\text{id}_\sigma - \pi_\sigma^{\mathcal{F}} \iota_\tau^{\mathcal{F}} \pi_\tau^{\mathcal{F}} \iota_\sigma^{\mathcal{F}} \right]$$

- (2) Let n be the total dimension of $\mathbf{N}_{\bullet}^{\mathbf{S}, \mathcal{F}}$. There exist positive constants C_1 and C_2 , which depend only on $\mathbf{A}_{\bullet}^{\mathbf{S}, \mathcal{F}}$, such that the inequalities

$$C_1 \cdot \lambda_i(\mathbf{N}_{\bullet}^{\mathbf{S}, \mathcal{F}}) \leq \lambda_i(\mathbf{A}'_\bullet) \leq C_2 \cdot \lambda_{i+k}(\mathbf{N}_{\bullet}^{\mathbf{S}, \mathcal{F}})$$

hold for every i in $\{1, 2, \dots, n\}$, with $k := n - \dim \mathbf{S}_{|\sigma|}^{\mathcal{F}}$.

The full statement of this result is recorded in Theorem 7.1.

Finally, we examined a single-cell chromatin accessibility (ATAC-seq) dataset from [28]; here X is a dataset consisting of around 3×10^4 tumour-infiltrating T cells, while the input \mathcal{F} consists of 20,000 binary-valued features $X \rightarrow \mathbb{R}$. To build a cover \mathcal{U} of X , we used a standard notion of genomic proximity between cells, constructed a 15-nearest neighbour graph, and expanded the communities of that graph by one hop. This produces a quiver $Q^{\mathcal{U}}$ with 87 vertices. We then ran the MACS2 peak-calling algorithm [39] to find the 1,000 most accessible regions for each admissible set and built a quiver Laplacian for a variant of the combined representation of \mathbf{S} (as described in Theorem 7.2).

OBSERVATION (V). *The first 3,000 eigenvectors of the quiver Laplacian described above had the following properties:*

- (1) *The majority of eigenvectors displayed a tight grouping of genomic positions (around 150 consecutive base-pairs out of approximately 3 billion) across the vertices.*
- (2) *Most eigenvectors were supported in most of the vector spaces assigned to vertices of the quiver. In particular, one such eigenvector that was supported on every vertex, lay near the gene FAM72D, which is involved in the cell cycle and hence active for every type of T cell.*
- (3) *Conversely, the support of a handful of eigenvectors was localised to a very small subset of vertices. One such eigenvector lay near the gene STAT3, which regulates Th1, Th17 and Treg cells; this eigenvector was supported on just 3 vertices, each of which corresponds to admissible subpopulations particularly rich in these three types of T cells.*

Thus, eigenvectors of the quiver Laplacian were able to isolate relevant and consistent meta-features across different scales within the given dataset.

Outline. The paper is structured as follows: §2 describes how a feature selector together with a cover of the data yields a quiver representation. In §3 we describe both local and global compatibility of features with respect to the cover. We also describe how these features can be computed simultaneously as sections of a larger quiver representation. §4 introduces the quiver Laplacian and how its eigenvectors may be used to approximate the space of sections. §5 studies how a transformation of quiver representations relates

the spectra of the respective representations and bounds their spectral distance. This then gives a stability of the spectrum of the representation associated to a feature selector. §6 bounds the changes in the spectrum of a quiver Laplacian when the underlying quiver undergoes various operations. §7 then uses these operations to simplify the computation of approximate sections of a feature selector. Finally, in §8 we use quiver Laplacians to extract (approximately) compatible peaks in single-cell chromatin accessibility data.

Acknowledgements

V.T. and H.A.H are grateful for the support provided by the UK Centre for Topological Data Analysis EPSRC grant EP/R018472/1. H.A.H. gratefully acknowledges funding from the Royal Society. VN is partially supported by US AFOSR grant FA9550-22-1-0462. OS would like to thank Renee Hoekzema, Ka Man (Ambrose) Yim, Phil Xie, and Xin Lu for many helpful conversations. For the purpose of Open Access, the authors have applied a CC BY public copyright licence to any Author Accepted Manuscript (AAM) version arising from this submission.

2. From Feature Selection to Quiver Representations

Given a finite set X and a field \mathbb{K} , let X^* be the \mathbb{K} -vector space which consists of all functions $X \rightarrow \mathbb{K}$. Denote by $\text{Sub}(X^*)$ the collection of all subsets of X^* (not necessarily subspaces), and for each such subset \mathcal{F} write $\mathbf{V}(\mathcal{F})$ to indicate the subspace of X^* given by the \mathbb{K} -linear span of \mathcal{F} . Let $\mathbf{Gr}_k(X^*)$ be the Grassmannian of k -dimensional subspaces of X^* for each $k \geq 0$, and consider the disjoint union

$$\mathbf{Gr}_\infty(X^*) := \coprod_{k \geq 0} \mathbf{Gr}_k(X^*).$$

Here is the main object of study in this work.

DEFINITION 2.1. A **feature selector** on X is any map $\mathbf{S} : \text{Sub}(X^*) \rightarrow \mathbf{Gr}_\infty(X^*)$, such that for each input $\mathcal{F} \subset X^*$, the corresponding output $\mathbf{S}_X^{\mathcal{F}}$ is a subspace of $\mathbf{V}(\mathcal{F})$.

We now fix a feature selector \mathbf{S} on X ; readers may wish to keep the following concrete example in mind throughout the remainder of §2.

EXAMPLE 2.2. Among the most well-known and ubiquitous feature selectors is *principal components analysis* [19] — here,

- X is a finite origin-centered subset of some real Euclidean space \mathbb{R}^n ,
- \mathbb{K} is the field \mathbb{R} of real numbers, and
- \mathcal{F} is the set $\{x \mapsto \langle x, p \rangle\}$ of maps given by taking inner products with the unit vectors p .

For a fixed $k \ll n$, the output $\mathbf{S}_X^{\mathcal{F}}$ is the vector space spanned by the k inner product functions (or equivalently, the k unit vectors) along which the variance in X is maximised.

Our goal here is to find a principled framework for relating the vector subspaces $\mathbf{S}_Y^{\mathcal{F}}$ to each other for a fixed \mathcal{F} as Y ranges over subsets of X . Already for principal components analysis, it is clear that there need not be any coherent relationship between $\mathbf{S}_X^{\mathcal{F}}$ and $\mathbf{S}_Y^{\mathcal{F}}$ for arbitrary $Y \subset X$ — for instance, neither one is guaranteed to be a subspace of the other within $\mathbf{V}(\mathcal{F})$. It therefore becomes necessary to constrain the family of subsets under consideration. Here we will describe how to relate the subspaces $\mathbf{S}_Y^{\mathcal{F}}$ for Y ranging over the lattice of subsets generated from a chosen *cover* \mathcal{U} of X . We recall that any such \mathcal{U} is a finite collection of nonempty subsets satisfying $X = \bigcup_{U \in \mathcal{U}} U$.

DEFINITION 2.3. The **quiver associated to a cover** \mathcal{U} of X is the directed graph $Q^{\mathcal{U}}$ whose

- (1) vertices are subsets $\sigma \subset \mathcal{U}$ whose *support* $|\sigma| := \bigcap_{U \in \sigma} U$ is nonempty; and,
- (2) there is a unique directed edge $\sigma \rightarrow \tau$ whenever σ properly contains τ (and hence $|\sigma| \subset |\tau|$).

We now re-assemble the algebraic data of \mathbf{S} to produce a *representation* of $Q^{\mathcal{U}}$, i.e., an assignment of vector spaces to vertices and linear maps to edges [29]. Implicit in the construction below is the choice of an inner product structure on X^* , which allows us to define adjoints of linear maps, and hence, orthogonal projections onto subspaces. For each vertex σ of $Q^{\mathcal{U}}$, let $\iota_{\sigma}^{\mathcal{F}} : \mathbf{S}_{|\sigma|}^{\mathcal{F}} \hookrightarrow \mathbf{V}(\mathcal{F})$ denote the inclusion map and let $\pi_{\sigma}^{\mathcal{F}} : \mathbf{V}(\mathcal{F}) \rightarrow \mathbf{S}_{|\sigma|}^{\mathcal{F}}$ denote its adjoint, i.e., the orthogonal projection in X^* with respect to our chosen inner product.

DEFINITION 2.4. For each subset $\mathcal{F} \subset X^*$, the **S-representation** of $Q^{\mathcal{U}}$ along \mathcal{F} consists of the following assignments $\mathbf{A}_{\bullet} = \mathbf{A}_{\bullet}^{\mathbf{S}, \mathcal{F}}$:

- (1) every vertex σ of $Q^{\mathcal{U}}$ is assigned the vector space $\mathbf{A}_{\sigma} := \mathbf{S}_{|\sigma|}^{\mathcal{F}}$, which is a subspace of $\mathbf{V}(\mathcal{F}|_{|\sigma|})$ and hence³ of $\mathbf{V}(\mathcal{F})$; moreover,
- (2) every edge $\sigma \rightarrow \tau$ of $Q^{\mathcal{U}}$ is assigned the \mathbb{K} -linear map $\mathbf{A}_{\sigma \rightarrow \tau} : \mathbf{A}_{\sigma} \rightarrow \mathbf{A}_{\tau}$ given by $\pi_{\tau}^{\mathcal{F}} \circ \iota_{\sigma}^{\mathcal{F}}$.

It may be worth noting that the vertices of $Q^{\mathcal{U}}$ are simplices in the *nerve* of the cover \mathcal{U} [8], and every edge $\sigma \rightarrow \tau$ in $Q^{\mathcal{U}}$ corresponds to a face relation (where the simplex τ lies in the boundary of the simplex σ). In particular, it follows from the definition of $Q^{\mathcal{U}}$ above that the existence of adjacent edges $\sigma \rightarrow \tau \rightarrow \gamma$ forces the existence of the edge $\sigma \rightarrow \gamma$. We call \mathbf{A}_{\bullet} a **sheaf** on the nerve of \mathcal{U} if it satisfies the following *associativity criterion*: for every such pair of adjacent edges, the map $\mathbf{A}_{\sigma \rightarrow \gamma}$ must equal the composite $\mathbf{A}_{\tau \rightarrow \gamma} \circ \mathbf{A}_{\sigma \rightarrow \tau}$. In general, we do not expect this associative property to hold when \mathbf{A}_{\bullet} arises from a feature selector. The next section contains our attempt to rectify this shortcoming.

³Here we adopt the convention that any function $f : |\sigma| \rightarrow \mathbb{K}$ extends to all of X by setting $f = 0$ on $X \setminus |\sigma|$. Conversely, any function $X \rightarrow \mathbb{K}$ automatically restricts to a function on $|\sigma|$.

3. Compatibility and Sections

Here we use the \mathbf{S} -representation to define and study two notions of compatibility for a feature selector \mathbf{S} on X against a cover \mathcal{U} when invoked with input $\mathcal{F} \subset X^*$:

- (1) Fix a vertex $\sigma \in Q_0^{\mathcal{U}}$. An element $v_\sigma \in \mathbf{A}_\sigma^{\mathcal{S}, \mathcal{F}}$ is *locally compatible* with \mathcal{U} if

$$\iota_\sigma^{\mathcal{F}}(v_\sigma) = \iota_\tau^{\mathcal{F}} \circ \mathbf{A}_{\sigma \rightarrow \tau}(v_\sigma)$$

holds for every edge $\sigma \rightarrow \tau$ in Q_1 .

- (2) A collection $\{v_\sigma \in \mathbf{A}_\sigma \mid \sigma \in Q_0^{\mathcal{U}}\}$ is *globally compatible* with \mathcal{U} if the equality

$$\mathbf{A}_{\sigma \rightarrow \tau}(v_\sigma) = v_\tau$$

holds for every edge $\sigma \rightarrow \tau$ in $Q^{\mathcal{U}}$.

This notion of local compatibility checks whether or not (the orthogonal projections of) a given feature v_σ which has been selected by \mathbf{S} over $|\sigma|$ continue to be selected over all greater subsets $|\tau| \supset |\sigma|$ generated by the cover \mathcal{U} . In this sense, local compatibility tests the robustness of selected features as we decrease the depth of the cover. Already for principal component analysis, it is clear that not every feature will be locally compatible, i.e., in general $\mathbf{A}_{\sigma \rightarrow \tau}$ has a non-trivial kernel. Global compatibility, on the other hand, tests features horizontally across the quiver — to be globally compatible, the vectors of a family $\{v_\sigma\}$ must map coherently onto each other under the linear maps of \mathbf{A}_\bullet as we traverse zigzag paths in $Q^{\mathcal{U}}$ of the form

$$\sigma_1 \rightarrow \tau_1 \leftarrow \sigma_2 \rightarrow \tau_2 \leftarrow \cdots \leftarrow \sigma_k \rightarrow \tau_k.$$

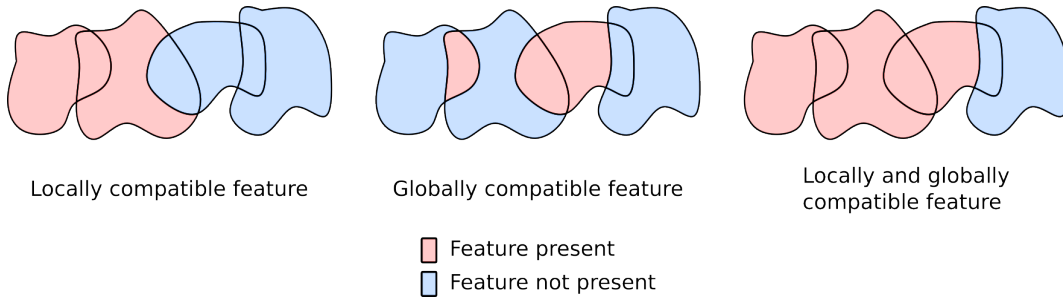


FIGURE 2. Example of different compatibility conditions.

We defer the study of local compatibility for now, and will focus on global compatibility.

3.1. Global Compatibility and Sections. Let $Q = (s, t : Q_1 \rightarrow Q_0)$ be a finite quiver and let \mathbf{A}_\bullet be a representation of Q valued in the category $\mathbf{FdHilb}(\mathbb{K})$ of finite-dimensional Hilbert spaces over the field $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$. The *total space* of \mathbf{A}_\bullet is defined as the product

$$\text{Tot}(\mathbf{A}_\bullet) := \prod_{i \in Q_0} \mathbf{A}_i. \quad (1)$$

Since Q is finite, $\text{Tot}(\mathbf{A}_\bullet)$ inherits a finite-dimensional Hilbert space structure from its \mathbf{A}_i factors. Motivated by the global compatibility criterion from Definition 3.2, we highlight a relevant subspace of the total space below.

DEFINITION 3.1. A tuple $\gamma := (\gamma_i \in \mathbf{A}_i \mid i \in Q_0)$ in $\text{Tot}(\mathbf{A}_\bullet)$ is called a **section** of \mathbf{A}_\bullet if for every edge e in Q_1 we have $\mathbf{A}_e(\gamma_{s(e)}) = \gamma_{t(e)}$.

It is straightforward to check that sums and scalar multiples of sections are also sections; we write $\Gamma(\mathbf{A}_\bullet)$ to denote the vector space of sections of \mathbf{A}_\bullet . This space was introduced in [30] along with an effective algorithm for its computation — see [30, Sec 5.2]. Sections are closely related to the compatibility criteria from Definition 3.2. In particular, if \mathbf{S} is a feature selector on a set X , then the sections of $\mathbf{A}_\bullet^{\mathbf{S}, \mathcal{F}}$ are exactly globally compatible features of \mathbf{S} against the cover \mathcal{U} . Sections can also be used to compute locally compatible features, as we will now describe.

3.2. Local Compatibility.

DEFINITION 3.2. A feature selector \mathbf{S} on X is **locally compatible** with \mathcal{U} on $\mathcal{F} \subset X^*$ if the following triangle of vector spaces commutes for each edge $\sigma \rightarrow \tau$ of $Q^{\mathcal{U}}$:

$$\begin{array}{ccc} \mathbf{A}_\sigma^{\mathbf{S}, \mathcal{F}} & \xrightarrow{\mathbf{A}_{\sigma \rightarrow \tau}^{\mathbf{S}, \mathcal{F}}} & \mathbf{A}_\tau^{\mathbf{S}, \mathcal{F}} \\ & \searrow \iota_\sigma^{\mathcal{F}} & \swarrow \iota_\tau^{\mathcal{F}} \\ & \mathbf{V}(\mathcal{F}) & \end{array}$$

Equivalently, the identity $\iota_\sigma^{\mathcal{F}} = \iota_\tau^{\mathcal{F}} \circ \pi_\tau^{\mathcal{F}} \circ \iota_\sigma^{\mathcal{F}}$ holds for every edge $\sigma \rightarrow \tau$ of $Q^{\mathcal{U}}$.

It will be useful to rephrase this compatibility condition entirely within the realm of quiver representations by blowing up the above triangle to a square. To this end, we introduce another representation $\mathbf{C}_\bullet = \mathbf{C}_\bullet^{\mathcal{F}}$ of $Q^{\mathcal{U}}$ which may be associated to every $\mathcal{F} \subset X^*$ (and which does not depend on \mathbf{S}). This *constant representation* assigns $\mathbf{C}_\sigma := \mathbf{V}(\mathcal{F})$ to every vertex σ and the identity map to each edge $\sigma \rightarrow \tau$. For each vertex σ , there is an evident inclusion map

$$\iota_\sigma^{\mathcal{F}} : \mathbf{A}_\sigma^{\mathbf{S}, \mathcal{F}} \hookrightarrow \mathbf{C}_\sigma^{\mathcal{F}},$$

since the codomain is $\mathbf{V}(\mathcal{F})$ and the domain is its subspace $\mathbf{S}_{|\sigma|}^{\mathcal{F}}$. Now, for every edge $\sigma \rightarrow \tau$ we have a diagram of four linear maps:

$$\begin{array}{ccc} \mathbf{A}_\sigma^{\mathbf{S}, \mathcal{F}} & \xrightarrow{\mathbf{A}_{\sigma \rightarrow \tau}^{\mathbf{S}, \mathcal{F}}} & \mathbf{A}_\tau^{\mathbf{S}, \mathcal{F}} \\ \iota_\sigma^{\mathcal{F}} \downarrow & & \downarrow \iota_\tau^{\mathcal{F}} \\ \mathbf{C}_\tau^{\mathcal{F}} & \xrightarrow{\mathbf{C}_{\sigma \rightarrow \tau}^{\mathcal{F}}} & \mathbf{C}_\sigma^{\mathcal{F}} \end{array} \quad (2)$$

The vertical maps $\iota^{\mathcal{F}}$ prescribe a *morphism of quiver representations* from $\mathbf{A}_\bullet^{\mathbf{S}, \mathcal{F}}$ to $\mathbf{C}_\bullet^{\mathcal{F}}$ if and only if the above diagram commutes for every edge $\sigma \rightarrow \tau$ in $Q^{\mathcal{U}}$. By definition of

$\mathbf{C}_{\bullet}^{\mathcal{F}}$, the bottom edge of this diagram is precisely $\mathbf{V}(\mathcal{F}) \xrightarrow{=} \mathbf{V}(\mathcal{F})$, so this square is an elementary modification of the triangle from Definition 3.2. Thus, we have arrived at the following straightforward result.

PROPOSITION 3.3. *The feature selector \mathbf{S} is locally compatible with \mathcal{U} on $\mathcal{F} \subset X^*$ if and only if the collection of maps*

$$\left\{ \iota_{\sigma}^{\mathcal{F}} : \mathbf{A}_{\sigma}^{\mathbf{S}, \mathcal{F}} \hookrightarrow \mathbf{C}_{\sigma}^{\mathcal{F}} \right\}$$

indexed by vertices σ of $Q^{\mathcal{U}}$ prescribe a morphism of quiver representations $\mathbf{A}_{\bullet}^{\mathbf{S}, \mathcal{F}} \rightarrow \mathbf{C}_{\bullet}^{\mathcal{F}}$.

As described earlier, a morphism τ between quiver representations \mathbf{A}_{\bullet} and \mathbf{A}'_{\bullet} is a collection

$$\tau = \{ \tau_v : \mathbf{A}_v \rightarrow \mathbf{A}'_v \mid v \in Q_0 \}$$

of linear maps τ_v such that

$$\mathbf{A}'_e \circ \tau_{s(e)} = \tau_{t(e)} \circ \mathbf{A}_{t(e)} \quad (3)$$

for each edge $e \in Q_1$. The commutivity condition in (3) can be too strong for quiver representations obtained from data, so we will define a *transformation* τ between quiver representations as a collection of linear maps

$$\tau = \{ \tau_v : \mathbf{A}_v \rightarrow \mathbf{A}'_v \mid v \in Q_0 \}$$

without the commutivity condition.

DEFINITION 3.4. Let Q be a quiver. For each vertex $u \in Q_0$ define the **floret** of based at u to be the following quiver $F^{Q,u}$:

- (1) the set of vertices is the disjoint union

$$F_0^{Q,u} = \{u\} \sqcup \{v_e \mid (e : u \rightarrow v) \in Q_1\}.$$

In other words, there is a distinguished copy of u , which we label u_0 henceforth, and an additional object denoted v_e for every edge in Q_1 from u to v .

- (2) Define the set of edges of $F^{Q,u}$ as

$$F_1^{Q,u} = \{e_{LD}, e_{DL} : u_0 \rightarrow v_e \mid (e : u \rightarrow v) \in Q_1\}.$$

In other words, for each edge $e : u \rightarrow v$ in Q there are two distinct edges e_{LD} and e_{DL} from u_0 to v_e in $F^{Q,u}$.

Now suppose \mathbf{A}_{\bullet} and \mathbf{A}'_{\bullet} are representations of the quiver Q with a transformation $\tau : \mathbf{A}_{\bullet} \rightarrow \mathbf{A}'_{\bullet}$. Define the following representation $\mathbf{F}_{\bullet}^{\tau,u}$ of $F^{Q,u}$:

$$\begin{aligned} \mathbf{F}_{u_0}^{\tau,u} &= \mathbf{A}_u & \mathbf{F}_{v_e}^{\tau,u} &= \mathbf{A}'_v \\ \mathbf{F}_{e_{LD}}^{\tau,u} &= \mathbf{A}'_e \circ \tau_u & \mathbf{F}_{e_{DL}}^{\tau,u} &= \tau_{t(e)} \circ \mathbf{A}_e. \end{aligned}$$

See Figure 3.4 for an illustration.

Constructing the floret now allows us to compute locally compatible features: in the case where the transformation is $\iota_{\sigma}^{\mathcal{F}} : \mathbf{A}_{\sigma}^{\mathbf{S}, \mathcal{F}} \rightarrow \mathbf{C}_{\sigma}^{\mathcal{F}}$ as in Proposition 3.3, the sections $\Gamma(\mathbf{F}_{\bullet}^{\iota_{\sigma}^{\mathcal{F}}})$ are exactly the locally compatible features at σ .

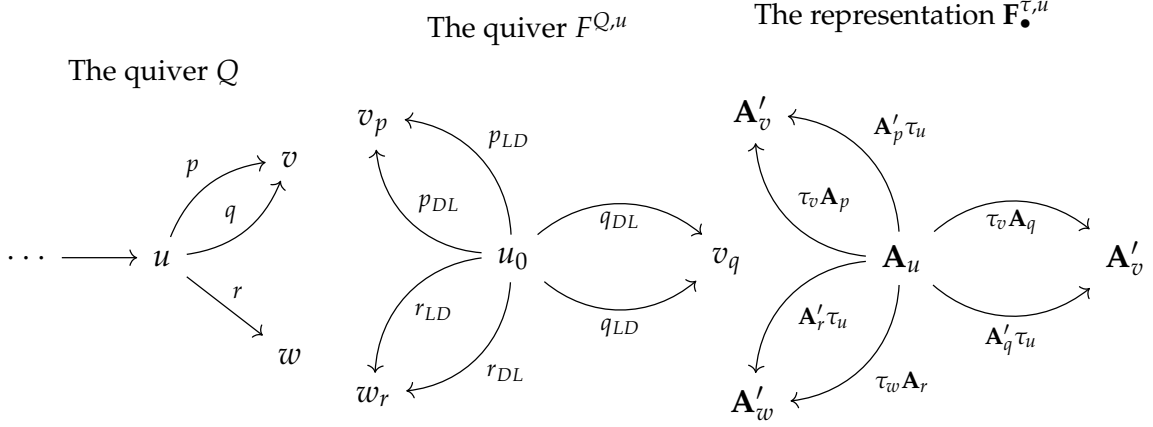


FIGURE 3. An illustration of the floret based at u and its representation.

3.3. Bicompatible features. In Section 4 below, we will consider the problem of discovering features which satisfy approximate versions of the local and global compatibility criteria. For this purpose, it will be helpful to construct a single quiver whose sections correspond to features which are both locally and globally compatible. The first step in this construction is a method for gluing quivers and their representations. Let $Q = (s, t : Q_1 \rightarrow Q_0)$ and $Q' = (s', t' : Q'_1 \rightarrow Q'_0)$ be two quivers.

DEFINITION 3.5. For each subset $R \subset Q_0 \times Q'_0$, let $Q \cup_R Q'$ be the **merged quiver** whose

- (1) vertex set is the quotient $V := (Q_0 \sqcup Q'_0) / R$,
- (2) edge set is the disjoint union $E := Q_1 \sqcup Q'_1$,

and source/target maps are given as follows. Let ρ be the composite $Q_0 \hookrightarrow Q_0 \sqcup Q'_0 \twoheadrightarrow V$, and define $\rho' : Q'_0 \rightarrow V$ similarly. If $e \in E$ lies in Q_1 , then its source and target vertices are $\rho \circ s(e)$ and $\rho \circ t(e)$; and if e lies in Q'_1 , then its source and target vertices are $\rho' \circ s'(e)$ and $\rho' \circ t'(e)$ respectively.

Fix representations \mathbf{A}_\bullet of a quiver Q and \mathbf{A}'_\bullet of another quiver Q' .

DEFINITION 3.6. Let $R \subset Q_0 \times Q'_0$ be any subset for which we have $\mathbf{A}_i = \mathbf{A}'_j$ whenever (i, j) lies in R . The **merged representation** of \mathbf{A}_\bullet and \mathbf{A}'_\bullet , denoted $(\mathbf{A} \cup_R \mathbf{A}')_\bullet$, is the following representation of $Q \cup_R Q'$. To vertices, it assigns

$$(\mathbf{A} \cup_R \mathbf{A}')_i := \begin{cases} \mathbf{A}_i & i \in Q_0 \text{ and } (i, j) \notin R \text{ for any } j \in Q'_0 \\ \mathbf{A}'_i & i \in Q'_0 \text{ and } (j, i) \notin R \text{ for any } j \in Q_0 \\ \mathbf{A}_i = \mathbf{A}'_j & \text{if } (i, j) \in R \\ \mathbf{A}_j = \mathbf{A}'_i & \text{if } (j, i) \in R \end{cases}$$

And on edges, we have

$$(\mathbf{A} \cup_R \mathbf{A}')_e := \begin{cases} \mathbf{A}_e & \text{if } e \in Q_1, \\ \mathbf{A}'_e & \text{if } e \in Q'_1. \end{cases}$$

One may embed both $\text{Tot}(\mathbf{A}_\bullet)$ and $\text{Tot}(\mathbf{A}'_\bullet)$ within the total space of the merged representation in the natural way, and it is readily seen that the space of sections of $(\mathbf{A} \cup_R \mathbf{A}')_\bullet$ is (isomorphic to) the intersection $\Gamma(\mathbf{A}_\bullet) \cap \Gamma(\mathbf{A}'_\bullet)$.

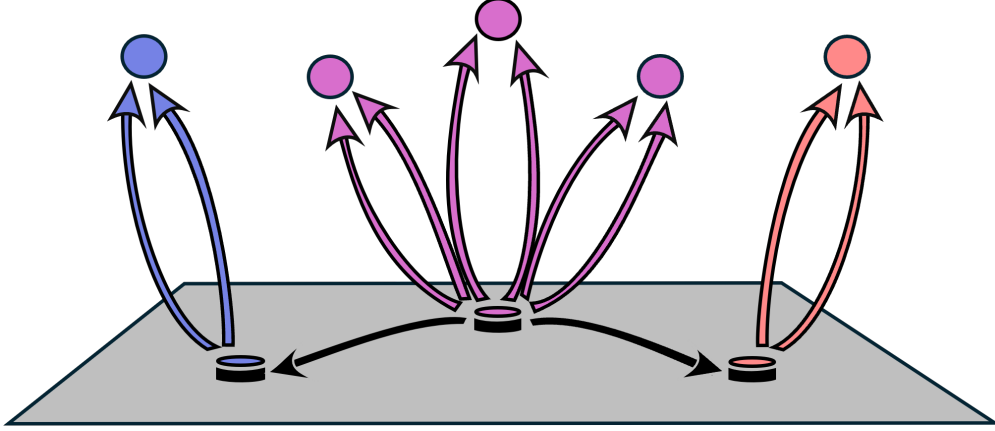


FIGURE 4. Illustration of the merged quiver

Returning to the case of interest, let \mathbf{S} be a feature selector on a finite set X ; consider a cover \mathcal{U} of X and a subset $\mathcal{F} \subset X^*$. Let $\widehat{Q}^{\mathcal{U}}$ be the quiver given by the disjoint union of florets \mathcal{I}^σ indexed by vertices $\sigma \in Q_0^{\mathcal{U}}$, and let $\widehat{\mathbf{A}}^{\mathbf{S}, \mathcal{F}}$ be the representation of $\widehat{Q}^{\mathcal{U}}$ given by the floret functors \mathcal{B}^σ . We define $R \subset (Q^{\mathcal{U}})_0 \times \widehat{Q}^{\mathcal{U}}$ as the collection

$$R := \left\{ (\sigma, \sigma_0) \mid \sigma \in Q_0^{\mathcal{U}} \right\},$$

where σ_0 is the central vertex of \mathcal{I}^σ .

DEFINITION 3.7. The **combined \mathbf{S} -representation** (of $Q^{\mathcal{U}} \cup_R \widehat{Q}^{\mathcal{U}}$, along \mathcal{F}) is the merged representation $\mathbf{N}_\bullet^{\mathbf{S}, \mathcal{F}} := \left(\mathbf{A}^{\mathbf{S}, \mathcal{F}} \cup_R \widehat{\mathbf{A}}^{\mathbf{S}, \mathcal{F}} \right)_\bullet$.

By design, the sections of $\mathbf{N}_\bullet^{\mathbf{S}, \mathcal{F}}$ correspond to relevant features which are both locally and globally compatible with \mathcal{U} . We call these sections the **bicompatible features** of \mathbf{S} .

EXAMPLE 3.8. Suppose $\mathcal{U} = \{U_1, U_2\}$ is a cover of X , such that $U_1 \cap U_2 \neq \emptyset$. Define $\tau_1 = \{U_1\}$, $\tau_2 = \{U_2\}$ and $\sigma = \{U_1, U_2\}$. The quiver $Q^{\mathcal{U}}$ is

$$\tau_1 \longleftarrow \sigma \longrightarrow \tau_2.$$

Consider a bicompatible feature, i.e., a section $x \in \Gamma(\mathbf{N}_\bullet^{\mathbf{S}, \mathcal{F}})$. By definition, we have $\iota_{\tau_2}^{\mathcal{F}}(x_{\tau_2}) = \iota_\sigma^{\mathcal{F}}(x_\sigma) = \iota_{\tau_1}^{\mathcal{F}}(x_{\tau_1})$, hence

$$\Gamma(\mathbf{N}_\bullet^{\mathbf{S}, \mathcal{F}}) \cong \iota^{\mathcal{F}}(\mathbf{S}_{|\tau_1}^{\mathcal{F}}) \cap \iota^{\mathcal{F}}(\mathbf{S}_{|\sigma}^{\mathcal{F}}) \cap \iota^{\mathcal{F}}(\mathbf{S}_{|\tau_2}^{\mathcal{F}}).$$

Thus, bicompatible features are the ones which are relevant for every $U \in \mathcal{U}$.

3.4. Dual Representations. If one is also interested in features which are relevant for some – but not all – $U \in \mathcal{U}$, then it becomes necessary to consider a different notion of bicompatibility. Given a quiver $Q = (s, t : Q_1 \rightarrow Q_0)$, its *dual* Q^* is the quiver with the same vertices and edges but with source and target maps interchanged, i.e., $s^* := t$ and $t^* := s$. Every representation \mathbf{A}_\bullet of Q valued in $\mathbf{FdHilb}(\mathbb{K})$ induces a unique *dual representation* \mathbf{A}_\bullet^* of Q^* , which is obtained by taking the adjoint of every edge map: namely, $(\mathbf{A}^*)_e := (\mathbf{A}_e)^*$ for each $e \in Q_1$.

When \mathbf{A}_\bullet is the \mathbf{S} -representation (from Definition 2.4) of a feature selector \mathbf{S}_\bullet along some input \mathcal{F} , the fact that $\iota_\sigma^\mathcal{F}$ and $\pi_\sigma^\mathcal{F}$ form an adjoint pair for each vertex $\sigma \in Q_0^\mathcal{U}$ implies that for every edge $\sigma \rightarrow \tau$ in $Q^\mathcal{U}$, we have

$$(\mathbf{A}^{\mathbf{S}, \mathcal{F}})_{\tau \rightarrow \sigma}^* = (\pi_\tau^\mathcal{F} \circ \iota_\sigma^\mathcal{F})^* = \pi_\sigma^\mathcal{F} \circ \iota_\tau^\mathcal{F}.$$

DEFINITION 3.9. The **mixed \mathbf{S} -representation** (of $(Q^\mathcal{U})^* \cup_R \widehat{Q}^\mathcal{U}$, along \mathcal{F}) is the merged representation $\mathbf{M}_\bullet^{\mathbf{S}, \mathcal{F}} := \left((\mathbf{A}^{\mathbf{S}, \mathcal{F}})^* \cup_R \widehat{\mathbf{A}}^{\mathbf{S}, \mathcal{F}} \right)_\bullet$.

Let us revisit Example 3.8 with a view towards understanding the sections of $\mathbf{M}_\bullet^{\mathbf{S}, \mathcal{F}}$. These sections are given by the direct sum

$$\Gamma(\mathbf{M}_\bullet^{\mathbf{S}, \mathcal{F}}) \cong (\iota_{\tau_1}(\mathbf{S}_{|\tau_1|}^\mathcal{F}) \cap \iota_\sigma(\mathbf{S}_{|\sigma|}^\mathcal{F})^\perp) \oplus (\iota_{\tau_2}(\mathbf{S}_{|\tau_2|}^\mathcal{F}) \cap \iota_\sigma(\mathbf{S}_{|\sigma|}^\mathcal{F})^\perp) \oplus \Gamma(\mathbf{N}_\bullet^{\mathbf{S}, \mathcal{F}}).$$

In this case, $\iota_{\tau_1}(\mathbf{S}_{|\tau_1|}^\mathcal{F}) \cap \iota_\sigma(\mathbf{S}_{|\sigma|}^\mathcal{F})^\perp$ corresponds to features that are relevant to U_1 but not to $U_1 \cap U_2$. Note that

$$(\iota_{\tau_1}(\mathbf{S}_{|\tau_1|}^\mathcal{F}) \cap \iota_\sigma(\mathbf{S}_{|\sigma|}^\mathcal{F})^\perp) \cap (\iota_{\tau_2}(\mathbf{S}_{|\tau_2|}^\mathcal{F}) \cap \iota_\sigma(\mathbf{S}_{|\sigma|}^\mathcal{F})^\perp)$$

might not be trivial, i.e. there could be a feature relevant to both U_1 and U_2 but not to $U_1 \cap U_2$, and hence this feature corresponds to two independent sections.

4. The Quiver Laplacian and Approximate Sections

So far, we have reframed compatibility-testing of feature selectors in terms of finding the sections of a quiver representation. However, the space of sections might be trivial [30, Sec 5.1]. Rather than seeking exact solutions, we endeavour to find approximate sections.

4.1. The Quiver Laplacian. Let $\mathbf{FdHilb}(\mathbb{K})$ denote the category of finite-dimensional Hilbert spaces over the field $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$; every morphism $A : U \rightarrow V$ in this category admits an *adjoint* morphism $A^* : V \rightarrow U$ characterised by $\langle Au, v \rangle = \langle u, A^*v \rangle$ for all (u, v) in $U \times V$. Here we fix a finite quiver $Q = (s, t : Q_1 \rightarrow Q_0)$ and consider a representation \mathbf{A}_\bullet of Q valued in $\mathbf{FdHilb}(\mathbb{K})$. We recall the total space from (1) and define *target space* of \mathbf{A}_\bullet as the direct product

$$\text{Tar}(\mathbf{A}_\bullet) := \prod_{e \in Q_1} \mathbf{A}_{t(e)}.$$

Since we have assumed that Q is finite, both these spaces inherit a Hilbert space structure from the individual factors of the form \mathbf{A}_v . We will denote the identity map on \mathbf{A}_v by id_v rather than $\text{id}_{\mathbf{A}_v}$.

DEFINITION 4.1. Given an $\mathbf{FdHilb}(\mathbb{K})$ -valued quiver representation \mathbf{A}_\bullet of Q , its **boundary operator**

$$B_{\mathbf{A}} : \text{Tot}(\mathbf{A}_\bullet) \rightarrow \text{Tar}(\mathbf{A}_\bullet)$$

is the linear map given in component form by

$$(B_{\mathbf{A}})_{e,v} = \begin{cases} \mathbf{A}_e - \text{id}_v & \text{if } s(e) = t(e) = v, \\ \mathbf{A}_e & \text{if } s(e) = v \text{ and } t(e) \neq v, \\ -\text{id}_v & \text{if } s(e) \neq v \text{ and } t(e) = v, \\ \mathbf{0} & \text{otherwise.} \end{cases}$$

If the quiver Q is finite, then it follows from a simple calculation that $\ker B_{\mathbf{A}}$ is isomorphic (as a Hilbert space) to the space $\Gamma(Q; \mathbf{A}_\bullet)$ of \mathbf{A}_\bullet 's sections. From the boundary operator of a quiver representation, we can construct another associated operator, the Laplacian, which will allow us to compute approximate sections.

DEFINITION 4.2. The **Laplacian** of \mathbf{A}_\bullet is the endomorphism $L_{\mathbf{A}} : \text{Tot}(\mathbf{A}_\bullet) \rightarrow \text{Tot}(\mathbf{A}_\bullet)$ given by

$$L_{\mathbf{A}} = B_{\mathbf{A}}^* B_{\mathbf{A}}.$$

In component form, we have

$$(L_{\mathbf{A}})_{v,v} = \sum_{\substack{e \in Q_1 \\ s(e)=v}} \mathbf{A}_e^* \mathbf{A}_e + \sum_{\substack{e \in Q_1 \\ t(e)=v}} \text{id}_v - \sum_{\substack{e \in Q_1 \\ s(e)=t(e)=v}} [\mathbf{A}_e^* + \mathbf{A}_e]$$

and for $u \neq v$

$$(L_{\mathbf{A}})_{u,v} = - \sum_{\substack{e \in Q_1 \\ s(e)=u \\ t(e)=v}} \mathbf{A}_e^* - \sum_{\substack{e \in Q_1 \\ s(e)=v \\ t(e)=u}} \mathbf{A}_e.$$

Since $\ker L_{\mathbf{A}} = \ker B_{\mathbf{A}}$, the kernel of $L_{\mathbf{A}}$ also computes the sections of \mathbf{A}_\bullet . The main advantage of considering $L_{\mathbf{A}}$ rather than $B_{\mathbf{A}}$ in this context is that it is a Hermitian and positive semi-definite matrix. As such, it enjoys favourable spectral properties; in particular, all of its eigenvalues are all real and non-negative. We shall order these eigenvalues in increasing fashion, with multiplicity:

$$\lambda_1(L_{\mathbf{A}}) \leq \cdots \leq \lambda_n(L_{\mathbf{A}})$$

where $n := \dim \text{Tot}(\mathbf{A}_\bullet)$. If $\varphi : \mathbf{A}_\bullet^1 \rightarrow \mathbf{A}_\bullet^2$ is an isomorphism of quiver representations with each φ_v a unitary map then

$$L_{\mathbf{A}^1} = \Phi \circ L_{\mathbf{A}^2} \circ \Phi^*,$$

where $\Phi = \prod_{v \in Q_0} \varphi_v$, thus we may speak unambiguously of the eigenvalues $\lambda_i(\mathbf{A}_\bullet) := \lambda_i(L_{\mathbf{A}})$ of a quiver representation \mathbf{A}_\bullet .

4.2. Approximate Sections. To measure the distance in $\text{Tot}(\mathbf{A}_\bullet)$ between an arbitrary vector x and the space $\Gamma(\mathbf{A}_\bullet)$ of sections, we consider the **Dirichlet energy**:

$$\mathcal{E}_\mathbf{A}(x) := \langle x, L_\mathbf{A}x \rangle = \sum_{e \in Q_1} \|\mathbf{A}_e(x_{s(e)} - x_{t(e)})\|_{\mathbf{A}_{t(e)}}^2 \in \mathbb{R}_{\geq 0}$$

Here $\langle -, - \rangle$ is the induced inner product on $\text{Tot}(\mathbf{A}_\bullet)$, while x_v denotes the component of x in \mathbf{A}_v , and $\| - \|$ denotes the norm induced by the inner product. As $L_\mathbf{A}$ is Hermitian, we can form an orthonormal eigenbasis e_1, \dots, e_n where e_i corresponds to $\lambda_i(\mathbf{A}_\bullet)$. Writing x in this basis, we have

$$\mathcal{E}_\mathbf{A}(x) = \mathcal{E}_\mathbf{A} \left(\sum_i x_i e_i \right) = \sum_i \lambda_i(\mathbf{A}_\bullet) x_i^2.$$

Thus, $\sqrt{\mathcal{E}_\mathbf{A}(x)}$ is the λ_\bullet -weighted distance from x to $\Gamma(\mathbf{A}_\bullet)$; and moreover,

$$\sum_{i: \lambda_i(\mathbf{A}_\bullet)=0} x_i e_i$$

is the closest vector to x in $\Gamma(\mathbf{A}_\bullet)$. This motivates the following definition. (A similar notion for cellular sheaves was introduced in [20]).

DEFINITION 4.3. For a quiver Q and an $\mathbf{FdHilb}(\mathbb{K})$ -valued representation \mathbf{A}_\bullet of Q , and ε -approximate section of \mathbf{A}_\bullet is an $x \in \text{Tot}(\mathbf{A}_\bullet)$ such that $\|x\| = 1$ and $\mathcal{E}_\mathbf{A}(x) \leq \varepsilon$.

By Theorem B.1, if $x \in \text{span}\{e_1, \dots, e_k\}$ and $\|x\| = 1$, then x is a λ_k -approximate section. Note this is not exhaustive: there can exist λ_k -approximate sections of \mathbf{A}_\bullet that are not in the span of the first k eigenvectors, e.g. $\lambda_1 x_1 e_1 + \lambda_{k+1} x_{k+1} e_{k+1}$ for small enough x_{k+1} . Regardless, to compute approximate sections of a quiver representation one can compute eigenvectors of $L_\mathbf{A}$ with small eigenvalues. As we have connected approximate sections with the spectra of the quiver representation, we now want to understand how the spectra changes if we perturb the quiver representation (§5), or the underlying quiver (§6).

5. Spectral Stability of Feature Selectors

Let Q be a quiver and suppose \mathbf{A}_\bullet and $\widehat{\mathbf{A}}_\bullet$ are two representations of Q . Suppose τ is a collection $\{(\tau_v : \mathbf{A}_v \rightarrow \widehat{\mathbf{A}}_v) : v \in Q_0\}$ of linear maps. Define the *defect* of τ at an edge $e \in Q_1$ to be

$$\partial(\tau, e) = \left\| \widehat{\mathbf{A}}_e \tau_{s(e)} - \tau_{t(e)} \mathbf{A}_e \right\|,$$

Here, and henceforth, the expression $\| - \|$ when applied to a linear map between Hilbert spaces will always indicate the operator norm — for $M : V \rightarrow W$, the norm $\|M\|$ equals the supremum of $\|Mx\|_W$ over unit vectors $x \in V$. Now the *total defect* of τ is:

$$\partial(\tau) = \left(\sum_{e \in Q_1} \partial(\tau, e)^2 \right)^{\frac{1}{2}}.$$

The defect of τ is zero if and only if when τ prescribes a morphism of quiver representations. Given bases of $\text{Tot}(\mathbf{A}_\bullet)$ and $\text{Tot}(\widehat{\mathbf{A}}_\bullet)$, we view τ as a block diagonal matrix. Let

τ^+ denote its Moore-Penrose pseudoinverse. The values $\|\tau\|$ and $\|\tau^+\|$ are unique up to unitary transforms of $\text{Tot}(\mathbf{A}_\bullet)$ and $\text{Tot}(\widehat{\mathbf{A}}_\bullet)$. Finally, we let

$$\kappa(\tau) := \|\tau\| \|\tau^+\|.$$

be the *generalised condition number* of τ .

We seek to describe how the existence of τ forces a relationship between the Laplacian spectra of \mathbf{A}_\bullet and $\widehat{\mathbf{A}}_\bullet$. The first step in this direction is the following lemma.

LEMMA 5.1. *For a matrix τ , if $x \in (\ker \tau)^\perp$ then*

$$\|\tau x\| \geq \|\tau^+\|^{-1} \|x\|.$$

PROOF. Let u_1, \dots, u_m and v_1, \dots, v_n be left and right singular vectors of τ respectively. Let $x = \sum_{i=1}^n x_i v_i$. Then $\tau x = \sum_{i=1}^{\min\{n,m\}} \sigma_i x_i u_i$, but as $x \in \ker \tau^\perp$ we have that $x_j = 0$ when $\sigma_j = 0$ or $j > \min\{m, n\}$. Thus

$$\|\tau x\| \geq \left\| \sum_{i=1}^{\min\{n,m\}} \sigma_+ x_i u_i \right\| = \sigma_+ \|x\| = \|\tau^+\|^{-1} \|x\|$$

where σ_+ is the smallest non-zero singular value of τ . □

The next theorem describes how the existence of a transformation between quiver representations constrains their eigenvalues. In the statement below (and henceforth), we use $\text{null}(f)$ as a shorthand for the dimension of the kernel of a linear map $f : U \rightarrow V$ of finite-dimensional Hilbert spaces.

THEOREM 5.2. *Suppose \mathbf{A}_\bullet and $\widehat{\mathbf{A}}_\bullet$ are two representations of a quiver Q whose total spaces have dimensions n and m respectively. Let $\tau = \{\tau_v : \mathbf{A}_v \rightarrow \widehat{\mathbf{A}}_v \mid v \in Q_0\}$ be a collection of linear maps, viewed as a single map $\text{Tot}(\mathbf{A}_\bullet) \rightarrow \text{Tot}(\widehat{\mathbf{A}}_\bullet)$. Set $q := \text{null}(\tau)$; for every integer k such that $1 \leq k \leq n - q$, we have*

$$\lambda_k(\widehat{\mathbf{A}}) \leq \kappa(\tau)^2 \lambda_{k+q}(\mathbf{A}) + \partial(\tau) \|\tau^+\| \left[2\kappa(\tau) \lambda_{k+q}(\mathbf{A})^{\frac{1}{2}} + \partial(\tau) \|\tau^+\| \right].$$

In particular, if τ is a morphism of quiver representations, then $\partial(\tau)$ vanishes and hence

$$\lambda_k(\widehat{\mathbf{A}}) \leq \kappa(\tau)^2 \lambda_{k+q}(\mathbf{A}).$$

PROOF. Let x_1, \dots, x_n and y_1, \dots, y_m be eigenvectors of \mathbf{A}_\bullet and $\widehat{\mathbf{A}}_\bullet$ respectively, ordered by increasing eigenvalue. Suppose k is an integer such that $1 \leq k \leq n - q$. Let $l = k + q$ and define $X_l = \text{span}\{x_1, \dots, x_l\}$ and $Y^k = \text{span}\{y_k, \dots, y_m\}$. As $l = k + q$ we have that $\dim \tau(X_l) \geq k$ thus $\dim \tau(X_l) + \dim Y^k$ is at least $m + 1$ hence $\tau(X_l)$ and Y^k intersect non-trivially. Thus there exists some non-zero $y \in \tau(X_l) \cap Y^k$, and as $y \in Y^k$, by Theorem B.1 it holds that

$$\lambda_k(\widehat{\mathbf{A}}) \leq \frac{1}{\|y\|^2} \sum_{e \in Q_1} \left\| \widehat{\mathbf{A}}_e y_{s(e)} - y_{t(e)} \right\|^2.$$

Now we shall bound the right-hand side of the above. Since $y \in \tau(X_l)$ there exists some $x \in X_l \cap (\ker \tau)^\perp$ such that $\tau x = y$. Using Lemma 5.1 have that

$$\frac{1}{\|\tau x\|^2} \sum_{e \in Q_1} \left\| \widehat{\mathbf{A}}_e \tau_{s(e)} x_{s(e)} - \tau_{t(e)} x_{t(e)} \right\|^2 \leq \frac{\|\tau^+\|^2}{\|x\|^2} \sum_{e \in Q_1} \left\| \widehat{\mathbf{A}}_e \tau_{s(e)} x_{s(e)} - \tau_{t(e)} x_{t(e)} \right\|^2.$$

Then applying the triangle inequality gives us

$$\begin{aligned} & \frac{1}{\|\tau x\|^2} \sum_{e \in Q_1} \left\| \widehat{\mathbf{A}}_e \tau_{s(e)} x_{s(e)} - \tau_{t(e)} x_{t(e)} \right\|^2 \\ & \leq \frac{\|\tau^+\|^2}{\|x\|^2} \sum_{e \in Q_1} \left\| \tau_{t(e)} \mathbf{A}_e x_{s(e)} - \tau_{t(e)} x_{t(e)} \right\|^2 \\ & \quad + 2 \frac{\|\tau^+\|^2}{\|x\|^2} \sum_{e \in Q_1} \left\| \tau_{t(e)} \mathbf{A}_e x_{s(e)} - \tau_{t(e)} x_{t(e)} \right\| \left\| \widehat{\mathbf{A}}_e \tau_{s(e)} x_{s(e)} - \tau_{t(e)} \mathbf{A}_e x_{s(e)} \right\| \\ & \quad + \frac{\|\tau^+\|^2}{\|x\|^2} \sum_{e \in Q_1} \left\| \widehat{\mathbf{A}}_e \tau_{s(e)} x_{s(e)} - \tau_{t(e)} \mathbf{A}_e x_{s(e)} \right\|^2 \end{aligned}$$

and then with the Cauchy-Schwarz inequality we have

$$\begin{aligned} & \frac{1}{\|\tau x\|^2} \sum_{e \in Q_1} \left\| \widehat{\mathbf{A}}_e \tau_{s(e)} x_{s(e)} - \tau_{t(e)} x_{t(e)} \right\|^2 \\ & \leq \kappa(\tau)^2 \frac{1}{\|x\|^2} \sum_{e \in Q_1} \left\| \mathbf{A}_e x_{s(e)} - \tau_{t(e)} x_{t(e)} \right\|^2 \\ & \quad + \kappa(\tau) \|\tau^+\| \left[\frac{1}{\|x\|^2} \sum_{e \in Q_1} \left\| \mathbf{A}_e x_{s(e)} - \tau_{t(e)} x_{t(e)} \right\|^2 \right]^{\frac{1}{2}} \partial(\tau) + \|\tau^+\|^2 \partial(\tau)^2 \end{aligned}$$

and finally since $x \in X_l$, we may apply Theorem B.1 to obtain

$$\begin{aligned} & \frac{1}{\|\tau x\|^2} \sum_{e \in Q_1} \left\| \widehat{\mathbf{A}}_e \tau_{s(e)} x_{s(e)} - \tau_{t(e)} x_{t(e)} \right\|^2 \\ & \leq \kappa(\tau)^2 \lambda_l(\mathbf{A}) + 2\kappa(\tau) \|\tau^+\| \lambda_l(\mathbf{A})^{\frac{1}{2}} \partial(\tau) + \|\tau^+\|^2 \partial(\tau)^2 \end{aligned}$$

completing the proof. \square

We shall now describe how Theorem 5.2 gives us a bound on the spectral pseudo-metric between quiver representations. If the total spaces of \mathbf{A} and $\widehat{\mathbf{A}}_\bullet$ have the same dimension, then we can directly compare the associated Laplacian spectra. Following [14], we use the Wasserstein metric to perform such comparisons in order to account for the case where total spaces have different dimensions.

DEFINITION 5.3. For each real number $r \geq 0$, let $B(Q, r)$ be the set of representations \mathbf{A}_\bullet of Q which satisfy $\max_{e \in Q_1} \|A_e\| \leq r$.

By design, if $\mathbf{A}_\bullet \in B(Q, r)$, then its eigenvalues are bounded by $|Q_1|(r+1)^2$.

Given a representation \mathbf{A}_\bullet of Q with total dimension n , its *spectral measure* is

$$\mu_{\mathbf{A}} = \sum_{i=1}^n \delta_{\lambda_i(\mathbf{A})}$$

where $\delta_{\lambda_i(\mathbf{A})}$ is the Dirac measure concentrated at $\lambda_i(\mathbf{A})$. Given two probability measures μ_1, μ_2 on $\mathbb{R}_{\geq 0}$, a *coupling* γ of μ_1 and μ_2 is any probability measure on $\mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}$ whose marginals on each factor are μ_1 and μ_2 respectively. Denote by $\Gamma(\mu_1, \mu_2)$ the set of all such couplings. Then the Wasserstein p -metric between μ_1 and μ_2 is defined as

$$W_p(\mu_1, \mu_2) = \inf_{\gamma \in \Gamma(\mu_1, \mu_2)} \left(\int_{\mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}} \|x - y\|^p d\gamma(x, y) \right)^{\frac{1}{p}}$$

Adapting the proof of Theorem 6.3 in [14] we have the following:

THEOREM 5.4. *Suppose $\mathbf{A}_\bullet, \widehat{\mathbf{A}}_\bullet \in B(Q, r)$ with total dimensions n and m respectively, where $n \geq m$. Suppose τ is a non-zero transformation from \mathbf{A}_\bullet to $\widehat{\mathbf{A}}_\bullet$. If $\text{rank}(\eta) \geq \text{null}(\eta^*)$ then*

$$W_1(\mu_{\mathbf{A}}, \mu_{\widehat{\mathbf{A}}}) \leq C$$

where C depends on τ, n and r . Moreover, if $\|\tau\|, \|\tau^+\| \leq 1 + \varepsilon$ and $\partial(\tau), \partial(\tau^*) \leq \varepsilon$ for some $\varepsilon > 0$, then we have

$$\lim_{\varepsilon \rightarrow 0} C = \frac{2 \text{null}(\tau) + 2 \text{null}(\tau^*)}{n} |Q_1|(r+1)^2,$$

which, if $\text{null}(\tau) + \text{null}(\tau^*) < n/2$, will be less than naïve bound of $|Q_1|(r+1)^2$.

PROOF. Define $c_\tau = \text{null}(\tau) + \text{null}(\tau^*)$. We will now define a coupling γ between $\mu_{\mathbf{A}}$ and $\mu_{\widehat{\mathbf{A}}}$. For each positive integer p , let $[p]$ denote the set $\{1, 2, \dots, p-1, p\}$. Setting $\ell = \text{null}(\tau^*)$ and $u = n - \text{null}(\tau)$, define the following partition of $[n] \times [m]$:

$$\begin{aligned} P_1 &= \{(i, i) : i \in [u] \setminus [\ell]\} \\ P_2 &= \{(i, j) : i \in [\ell] \text{ and } j \in [\ell]\} \\ &\quad \cup \{(i, j) : i \in [\ell] \text{ and } j \in [m] \setminus [u]\} \\ &\quad \cup \{(i, j) : i \in [n] \setminus [u] \text{ and } j \in [\ell]\} \\ &\quad \cup \{(i, j) : i \in [n] \setminus [u] \text{ and } j \in [m] \setminus [u]\} \\ P_3 &= \{(i, j) : i \in [\ell] \text{ and } j \in [u] \setminus [\ell]\} \\ &\quad \cup \{(i, j) : i \in [n] \setminus [u] \text{ and } j \in [u] \setminus [\ell]\} \\ P_4 &= ([n] \times [m]) \setminus (P_1 \cup P_2 \cup P_3). \end{aligned}$$

Now define the desired coupling γ as

$$\gamma = \sum_{\substack{1 \leq i \leq n \\ 1 \leq j \leq m}} w_{ij} \delta_{(\lambda_i(\mathbf{A}), \lambda_j(\widehat{\mathbf{A}}))} \quad \text{where} \quad w_{ij} = \begin{cases} \frac{1}{n} & (i, j) \in P_1, \\ \frac{1}{m(n-c_\tau)} & (i, j) \in P_2, \\ \left(\frac{1}{m} - \frac{1}{n}\right) \frac{1}{n-c_\tau} & (i, j) \in P_3, \\ 0 & (i, j) \in P_4. \end{cases}$$

Set $R = |Q_1|(r+1)^2$, which is an upper bound for the eigenvalues of \mathbf{A}_\bullet . Then using Theorem 5.2 for τ and τ^* and substituting R gives us

$$\kappa(\tau)^{-2} \lambda_{i-\text{null}(\tau^*)}(\mathbf{A}_\bullet) - a_{\tau^*, R} \leq \lambda_i(\widehat{\mathbf{A}}_\bullet) \leq \kappa(\tau)^2 \lambda_{i+\text{null}(\tau)}(\mathbf{A}_\bullet) + b_{\tau, R}$$

with the constants

$$a_{\tau^*, R} = 2\kappa(\tau)^{-1} \|\tau^+\| \|\partial(\tau^*)\| \sqrt{R} + \kappa(\tau)^{-2} \|\tau^+\|^2 \|\partial(\tau^*)\|^2, \text{ and}$$

$$b_{\tau, R} = 2\kappa(\tau) \|\tau^+\| \|\partial(\tau)\| \sqrt{R} + \|\tau^+\|^2 \|\partial(\tau)\|^2.$$

Computing cost of γ gives us

$$\begin{aligned} W_1(\mu_{\mathbf{A}}, \mu_{\widehat{\mathbf{A}}}) &\leq \frac{1}{n} \left[\sum_{k=\text{null}(\tau^*)+1}^{n-\text{null}(\tau)} |\lambda_k(\mathbf{A}_\bullet) - \lambda_k(\widehat{\mathbf{A}}_\bullet)| + c_\tau R \right] \\ &\leq \frac{1}{n} \left[\sum_{k=\text{null}(\tau^*)+1}^{n-\text{null}(\tau)} (\kappa(\tau)^2 \lambda_{k+\text{null}(\tau)}(\mathbf{A}_\bullet) - \kappa(\tau)^{-2} \lambda_{k-\text{null}(\tau^*)}(\mathbf{A}_\bullet)) \right. \\ &\quad \left. + (n - c_\tau)(a_{\tau^*, R} + b_{\tau, R}) + c_\tau R \right]. \end{aligned}$$

Rearranging the sum on the right side yields

$$\frac{1}{n} \left[c_\tau \kappa(\tau)^2 R + (n - 2c_\tau)(\kappa(\tau)^2 - \kappa(\tau)^{-2})R + (n - c_\tau)(a_{\tau^*, R} + b_{\tau, R}) + c_\tau R \right],$$

which is the desired C . \square

Suppose A, B are a pair of orthogonal matrices with dimensions $n \times k$ and $n \times l$ respectively. Then the matrix A^*B is an orthogonal projection $\text{img } B \rightarrow \text{img } A$, where $\text{img } A$ denotes the image of A . Let the singular values of A^*B be $\sigma_1, \dots, \sigma_r$ where $r = \min\{k, l\}$. Then the principal angles between $\text{img } A$ and $\text{img } B$ are defined by $\cos \theta_i = \sigma_i$ and the Grassmannian metric between $\text{img } A$ and $\text{img } B$ as defined by [37] is

$$d_{\mathbf{Gr}_\infty}(\text{img } A, \text{img } B) := \left(\sum_{i=1}^r \theta_i^2 \right)^{\frac{1}{2}}.$$

COROLLARY 5.5. *Suppose $\mathbf{S}, \mathbf{T} : \text{Sub}(X^*) \rightarrow \mathbf{Gr}_\infty(X^*)$ are two feature selectors on X and suppose \mathcal{U} is a cover of X and $\mathcal{F} \subset X^*$. Suppose that*

$$\max_{\sigma \in Q_0^{\mathcal{U}}} d_{\mathbf{Gr}_\infty}(\mathbf{S}_{|\sigma|}, \mathbf{T}_{|\sigma|}) \leq \varepsilon < \frac{\pi}{2}$$

for some $\varepsilon > 0$, and define

$$c = \sum_{\sigma \in Q_0^{\mathcal{Z}}} |\dim \mathbf{S}_{|\sigma|}^{\mathcal{F}} - \dim \mathbf{T}_{|\sigma|}^{\mathcal{F}}|.$$

Then, $W_1(\mu_{\mathbf{A}\mathbf{S}, \mathcal{F}}, \mu_{\mathbf{A}\mathbf{T}, \mathcal{F}})$ is no larger than

$$\frac{4|Q_1|}{n} \left[(1 + \cos^2 \varepsilon)c + \left(2 \sec \varepsilon \tan \varepsilon + (1 + \sec^2 \varepsilon) \tan^2 \varepsilon + \sec^2 \varepsilon - \cos^2 \varepsilon \right) (n - c) \right]$$

where $n = \max\{\dim \text{Tot}(\mathbf{A}_{\bullet}^{\mathbf{S}, \mathcal{F}}), \dim \text{Tot}(\mathbf{A}_{\bullet}^{\mathbf{T}, \mathcal{F}})\}$.

PROOF. Assume that $\dim \mathbf{S}_{\bullet}^{\mathcal{F}} \geq \dim \mathbf{T}_{\bullet}^{\mathcal{F}}$. Define the transformation η with components $\eta_{\sigma} = \pi_{\sigma}^{\mathbf{T}} \iota_{\sigma}^{\mathbf{S}}$. Then $\cos \varepsilon \leq \|\eta\| \leq 1$ and $1 \leq \|\eta^+\| \leq (\cos \varepsilon)^{-1}$. In order to bound $\partial(\eta)$ note that

$$\|\iota_{\sigma}^{\mathbf{S}} \pi_{\sigma}^{\mathbf{S}} - \iota_{\sigma}^{\mathbf{T}} \pi_{\sigma}^{\mathbf{T}}\|_2 = \sin \theta_{\sigma, \max} \leq \sin \varepsilon,$$

where $\theta_{\sigma, \max}$ is the largest principal angle between $\mathbf{S}_{|\sigma|}$ and $\mathbf{T}_{|\sigma|}$. This is known as the projection metric [37] where the equality can be seen from Theorem 2.6.1 in [13]. This then implies that, for each edge $\sigma \rightarrow \tau$, we have $\partial(\eta, \sigma \rightarrow \tau) \leq 2 \sin \varepsilon$ thus $\partial(\eta) \leq 2\sqrt{|Q_1|} \sin \varepsilon$. Similarly, $\partial(\eta^*) \leq 2\sqrt{|Q_1|} \sin \varepsilon$. Applying Theorem 5.4 together with these bounds provide the result. \square

6. Eigenvalue Inequalities for Quiver Operations

So far, we have described our problem as computing sections of a quiver representation. Here we seek to simplify the representation and the underlying quiver whilst preserving its space of sections and bounding the change in its eigenvalues. Here we will establish similar eigenvalue inequalities when the quiver itself is deformed via certain natural operations. The most convenient framework for extracting such inequalities is to upgrade the given quiver representation to a sheaf, as described below.

6.1. Sheaves over Quivers. Fix a quiver $Q = (s, t : Q_1 \rightarrow Q_0)$.

DEFINITION 6.1. An $\mathbf{FdHilb}(\mathbb{K})$ -valued **sheaf** \mathcal{A}_{\bullet} over Q consists of:

- A finite-dimensional Hilbert space \mathcal{A}_v for each vertex $v \in Q_0$,
- a finite-dimensional Hilbert space \mathcal{A}_e for each edge $e \in Q_1$, and
- for each for each edge $e \in Q_1$, a pair of linear maps

$$\mathcal{A}_{s(e) \rightarrow e} : \mathcal{A}_{s(e)} \rightarrow \mathcal{A}_e \quad \text{and} \quad \mathcal{A}_{e \leftarrow t(e)} : \mathcal{A}_{t(e)} \rightarrow \mathcal{A}_e.$$

Such sheaves are closely related to representations of quivers. Every representation \mathbf{A}_{\bullet} over Q induces a sheaf \mathcal{A}_{\bullet} over Q by setting $\mathcal{A}_{s(e) \rightarrow e} = \mathbf{A}_e$ and $\mathcal{A}_{e \leftarrow t(e)} = \text{id}_{t(e)}$. Conversely, every sheaf evidently induces a representation on the subdivision of Q where every edge $e : u \rightarrow v$ is decomposed into the zigzag $u \rightarrow e \leftarrow v$. Thus, there is a natural notion of Laplacian for sheaves over a quiver, a special case⁴ of which appears in [16].

⁴The only difference between our quiver sheaf Laplacian and the sheaf Laplacians of [16] over one-dimensional cell complexes is that we allow edges to have the same source and target vertex. On the other hand, cellular sheaves and their Laplacians are also defined over regular cell complexes of dimension > 1 .

DEFINITION 6.2. Suppose \mathcal{A}_\bullet is a $\mathbf{FdHilb}(\mathbb{K})$ -valued sheaf over a quiver Q . Define $\text{Tot}(\mathcal{A}_\bullet) = \prod_{v \in Q_0} \mathcal{A}_v$ and $\text{Tar}(\mathcal{A}_\bullet) = \prod_{e \in Q_1} \mathcal{A}_e$.

- (1) The **boundary operator** of \mathbf{A}_\bullet is the linear map $B_{\mathcal{A}} : \text{Tot}(\mathcal{A}_\bullet) \rightarrow \text{Tar}(\mathcal{A}_\bullet)$ of \mathcal{A}_\bullet given in component form by

$$(B_{\mathcal{A}})_{e,v} = \begin{cases} \mathcal{A}_{v \rightarrow e} - \mathcal{A}_{e \leftarrow v} & \text{if } s(e) = t(e) = v, \\ \mathcal{A}_{v \rightarrow e} & \text{if } s(e) = v \text{ and } t(e) \neq v, \\ -\mathcal{A}_{e \leftarrow v} & \text{if } s(e) \neq v \text{ and } t(e) = v, \\ \mathbf{0} & \text{otherwise.} \end{cases}$$

- (2) The **Laplacian** of \mathcal{A}_\bullet is then defined as

$$L_{\mathcal{A}} = B_{\mathcal{A}}^* B_{\mathcal{A}}.$$


- (3) Finally, the **Dirichlet energy** for \mathcal{A} is the function $\mathcal{E}_{\mathcal{A}} : \text{Tot}(\mathcal{A}_\bullet) \rightarrow \mathbb{R}$ given by

$$\mathcal{E}_{\mathcal{A}}(x) = \sum_{e \in Q_1} \left\| \mathcal{A}_{s(e) \rightarrow e}(x_{s(e)}) - \mathcal{A}_{e \leftarrow t(e)}(x_{t(e)}) \right\|_{\mathcal{A}_e}^2.$$


If one constructs a sheaf \mathcal{A}_\bullet from a quiver representation \mathbf{A}_\bullet of Q as described above, then the sheaf boundary operator $B_{\mathcal{A}}$ coincides exactly with the quiver boundary operator $B_{\mathbf{A}}$; therefore the Laplacians of \mathbf{A} and \mathcal{A} , as well as their spectra, also coincide in this case.

6.2. Operations on the Quiver. This section describes how the spectrum of a sheaf, or indeed a representation, over a quiver changes under various combinatorial operations on the quiver. Removing an edge or a vertex gives bounds on the change in spectrum, as Proposition A.1 describes, however these may reduce the dimension of the space of sections. Propositions 6.5, 6.7, and 6.9 describe operations on the quiver, dependent on the sheaf, that preserve the space of sections. These are described for sheaves over a quiver but all these operations also preserve quiver representations. These intermediate results are used to simplify quivers (while bounding the change in eigenvalues) in our Proofs of Theorems 7.1 and 7.2.

DEFINITION 6.3. For a quiver Q , define an *edge pairing* as a collection P of pairs of distinct neighbouring edges $(e_1^1, e_2^1), \dots, (e_1^m, e_2^m)$ of Q such no two edges are repeated, i.e. $e_j^i \neq e_l^k$ unless $i = k$ and $j = l$. To simplify notation assume that each pair (e_1^i, e_2^i) , the edges e_1^i and e_2^i are directed away from their common vertex u^i . Label the other vertices for e_1^i and e_2^i , possibly the same vertex, as v^i and w^i respectively.



We then define Q^P as the P -homotopic quiver to Q where Q^P has the same nodes and edges as Q except with each edge e_2^i replaced by an edge e_3^i between v^i and w^i .



For \mathcal{A}_\bullet is a sheaf over Q and a pairing P of Q , say that \mathcal{A}_\bullet is compatible with P if for each pair $(e_1^i, e_2^i) \in P$ we have that $\mathcal{A}_{e_1^i} = \mathcal{A}_{e_2^i}$ and $\mathcal{A}_{u^i \rightarrow e_1^i} = \mathcal{A}_{u^i \rightarrow e_2^i}$. Such a compatible sheaf induces a sheaf on Q^P as follows: define the sheaf \mathcal{A}_\bullet^P on Q^P as the same as \mathcal{A}_\bullet on the common edges of Q and Q^P and otherwise by $\mathcal{A}_{e_3^i}^P = \mathcal{A}_{e_2^i}$, $\mathcal{A}_{e_3^i \leftarrow e_3^i}^P = \mathcal{A}_{e_1^i \leftarrow e_1^i}$ and $\mathcal{A}_{w^i \rightarrow e_3^i}^P = \mathcal{A}_{e_2^i \leftarrow w^i}$ for each i . The next proposition shows the i -th eigenvalue of \mathcal{A}_\bullet^P is bounded by a scalar multiple of the i -th eigenvalue of \mathcal{A}_\bullet . First we will need a lemma.

LEMMA 6.4. *Suppose $X, Y, Z \in \mathbb{R}^n$. Then*

$$\|X - Z\|^2 + \|Z - Y\|^2 \leq \varphi^2(\|X - Y\|^2 + \|Y - Z\|^2)$$

where φ is the golden ratio $\varphi = \frac{1+\sqrt{5}}{2}$. This is also true for any metric space (M, d) and can be written as

$$d(X, Z)^2 \leq \varphi^2 d(X, Y)^2 + \varphi d(Y, Z)^2.$$

PROOF. Let us abbreviate

$$x = \|Y - Z\| \quad y = \|X - Z\| \quad z = \|X - Y\|.$$

If any of x, y or z are zero, then the proposition is immediately true, so assume that x, y and z are strictly positive. By the triangle inequality, we have

$$\frac{x^2 + y^2}{x^2 + z^2} \leq \frac{x^2 + (x + z)^2}{x^2 + z^2}.$$

We then want to know when the expression on the right side is maximised. As $x, z > 0$, there exists $u > 0$ for which $x = uz$. Substituting this into the right side gives

$$\frac{2u^2 + 2u + 1}{u^2 + 1},$$

which is maximised when $u = \varphi$. Thus, the maximum value of the ratio is φ^2 . If X, Y, Z are colinear with Y between X and Z and $\|Y - Z\| = \varphi\|X - Y\|$, then the maximum is attained, which shows that this inequality is sharp. \square

PROPOSITION 6.5. *For Q a quiver and P an edge pairing of Q , suppose that \mathbf{A}_\bullet is a sheaf on Q compatible with P . Then the eigenvalues of \mathcal{A}_\bullet and \mathcal{A}_\bullet^P are related by*

$$\varphi^{-2} \lambda_j(\mathcal{A}_\bullet) \leq \lambda_j(\mathcal{A}_\bullet^P) \leq \varphi^2 \lambda_j(\mathcal{A}_\bullet)$$

and vice-versa for $j = 1, \dots, n$ where n is the total dimension of \mathbf{A}_\bullet .

PROOF. Using Lemma 6.4 we have

$$\begin{aligned}
 \mathcal{E}_{\mathcal{A}}(x) &= \sum_{(e_1^i, e_2^i) \in P} \left[\left\| \mathcal{A}_{e_1^i \leftarrow v^i} x_{v^i} - \mathcal{A}_{u^i \rightarrow e_1^i} x_{u^i} \right\|_{\mathcal{A}_{e_1^i}}^2 + \left\| \mathcal{A}_{u^i \rightarrow e_2^i} x_{u^i} - \mathcal{A}_{e_2^i \leftarrow w^i} x_{w^i} \right\|_{\mathcal{A}_{e_2^i}}^2 \right] \\
 &\quad + \sum_{\text{unpaired } e \in Q_1} \left\| \mathcal{A}_{s(e) \rightarrow e} x_{s(e)} - \mathcal{A}_{e \leftarrow t(e)} x_{t(e)} \right\|_{\mathcal{A}_{s(e)}}^2 \\
 &\leq \varphi^2 \sum_{(e_1^i, e_2^i) \in P} \left[\left\| \mathcal{A}_{u^i \rightarrow e_1^i} x_{u^i} - \mathcal{A}_{e_1^i \leftarrow v^i} x_{v^i} \right\|_{\mathcal{A}_{e_1^i}}^2 + \left\| \mathcal{A}_{e_3^i \leftarrow v^i} x_{v^i} - \mathcal{A}_{w^i \rightarrow e_3^i} x_{w^i} \right\|_{\mathcal{A}_{e_3^i}}^2 \right] \\
 &\quad + \sum_{\text{unpaired } e \in Q_1} \left\| \mathcal{A}_{s(e) \rightarrow e} x_{s(e)} - \mathcal{A}_{e \leftarrow t(e)} x_{t(e)} \right\|_{\mathcal{A}_{s(e)}}^2 \\
 &\leq \varphi^2 \mathcal{E}_{\mathcal{A}^P}(x).
 \end{aligned}$$

Then by Theorem B.2

$$\lambda_i(\mathcal{A}_{\bullet}) = \min_{\dim V=i} \left[\max_{x \in V \setminus 0} \frac{\mathcal{E}_{\mathcal{A}}(x)}{\langle x, x \rangle} \right] \leq \min_{\dim V=i} \left[\max_{x \in V \setminus 0} \frac{\varphi^2 \mathcal{E}_{\mathcal{A}^P}(x)}{\langle x, x \rangle} \right] = \varphi^2 \lambda_i(\mathcal{A}_{\bullet}^P).$$

proving the proposition. \square

We note that Proposition 6.5 directly applies to graphs. Although we make no use of the following corollary, we hope that it will be of independent interest to spectral graph theorists.

COROLLARY 6.6. *Suppose $G = (V, E)$ is a multigraph with n vertices. Suppose we have an edge pairing $P = (e_1^1, e_2^1), \dots, (e_1^m, e_2^m)$ on G as in Definition 6.3. Let \widehat{G} with the same nodes and edges as G except with each edge e_2^i replaced by an edge e_3^i between v^i and w^i . Then the eigenvalues of $L_{\widehat{G}}$ and L_G are related by*

$$\varphi^{-2} \lambda_j(L_G) \leq \lambda_j(L_{\widehat{G}}) \leq \varphi^2 \lambda_j(L_G)$$

and vice-versa for $j = 1, \dots, n$.

PROPOSITION 6.7. *Suppose \mathcal{A}_{\bullet} is a sheaf over a quiver Q with total dimension n . Given a subquiver Q' , let \mathcal{A}'_{\bullet} denote the restriction of \mathcal{A}_{\bullet} to Q' .*

6.7(a) *If $e \in Q_1$ is an edge such that $s(e) = t(e)$ and $\mathcal{A}_{s(e)} = \mathcal{A}_{t(e)} = \text{id}$ then $\lambda_i(\mathcal{A}'_{\bullet}) = \lambda_i(\mathcal{A}_{\bullet})$ for $i = 1, \dots, n$ where \mathcal{A}'_{\bullet} is \mathcal{A}_{\bullet} restricted to the subquiver Q' of Q with the edge e removed.*

6.7(b) *Suppose we have an edge pairing $P = (e_1^1, e_2^1), \dots, (e_1^m, e_2^m)$ of Q such that for each i , the edges e_1^i, e_2^i have the same source and target vertices. Assume that for each i , there exists a linear map $M_i : \mathcal{A}_{e_2^i} \rightarrow \mathcal{A}_{e_1^i}$ satisfying both*

$$\mathcal{A}_{s(e_1^i) \rightarrow e_1^i} = M_i \mathcal{A}_{s(e_2^i) \rightarrow e_2^i} \quad \text{and} \quad \mathcal{A}_{e_1^i \leftarrow t(e_1^i)} = M_i \mathcal{A}_{e_2^i \leftarrow t(e_2^i)}$$

Let Q' be the subquiver of Q with all the edges e_1^1, \dots, e_2^n removed, and let \mathcal{A}'_\bullet be the restriction of \mathcal{A}_\bullet to Q' . Then, we have

$$(1 + \Gamma_P)^{-1} \lambda_i(\mathcal{A}_\bullet) \leq \lambda_i(\mathcal{A}'_\bullet) \leq \lambda_i(\mathcal{A}_\bullet)$$

for all $i = 1, \dots, n$, where $\Gamma_P := \max_i \{\|M_i\|^2\}$.

PROOF. For 6.7(a), observe that $\mathcal{E}_{\mathcal{A}}(x) = \mathcal{E}_{\mathcal{A}'_\bullet}(x)$ and thus the result follows from Theorem B.2. For 6.7(b), we have

$$\begin{aligned} \mathcal{E}_{\mathcal{A}}(x) &\leq \sum_{(e_1^i, e_2^i) \in P} (1 + \|M_i\|^2) \left\| \mathcal{A}_{s(e_1^i) \rightarrow e_1^i} x_{s(e_1^i)} - \mathcal{A}_{e_1^i \leftarrow t(e_1^i)} x_{t(e_1^i)} \right\|^2 \\ &\quad + \sum_{\text{unpaired } e \in Q_1} \left\| \mathcal{A}_{s(e) \rightarrow e} x_{s(e)} - \mathcal{A}_{e \leftarrow t(e)} x_{t(e)} \right\|^2_{\mathcal{A}_{s(e)}} \\ &\leq (1 + \Gamma_P) \mathcal{E}_{\mathcal{A}'_\bullet}(x). \end{aligned}$$


and the result follows from Theorem B.2. \square

In the following, at the risk of causing ambiguity with self-loops, we will ignore the direction of edges in and use the relation $v \sim e$ if either $s(e) = v$ or $t(e) = v$.

DEFINITION 6.8. Let Q be a quiver and let V be a subset of vertices such that there are no edges in Q between vertices of V . The **V-reduced quiver** Q^V is defined as follows. For each vertex $v \in V$,

- (1) first remove v and any associated edges from Q ;
- (2) then, for each pair of edges (p, q) in Q such that $u \sim p \sim v \sim q \sim w$ for vertices $u, w \in Q_0 \setminus V$, add a new edge pq between u and w .

(We note in the second step above that if $u = w$, then pq forms a self-loop.) The resulting quiver is the desired Q^V .



We now describe an analogue of the Kron reduction of a graph [7] for sheaves over a quiver. As discussed in [16], Kron reduction does not work in general for sheaves over a (loopless) graph. However, as the next proposition shows, if one restricts the type of

vertex removed, dependent on the sheaf, and moreover allows for loops in the underlying quiver, then Kron reduction is possible.

PROPOSITION 6.9. *Given a sheaf \mathcal{A}_\bullet over a quiver Q , assume that $V \subset Q_0$ is a subset of vertices such that there are no edges in Q between vertices of V . Further suppose that for each $v \in V$ there exists some real number $\alpha_v > 0$ such that for each edge $e \sim v$, we have*

$$\text{either } \mathcal{A}_{v \rightarrow e}^* \mathcal{A}_{v \rightarrow e} = \alpha_v \text{id}_v \quad \text{or} \quad \mathcal{A}_{e \leftarrow v}^* \mathcal{A}_{e \leftarrow v} = \alpha_v \text{id}_v$$

as appropriate. Then there exists a sheaf \mathcal{A}_\bullet^V over Q^V such that $\Gamma(\mathcal{A}_\bullet^V) \cong \Gamma(\mathcal{A}_\bullet)$ and

$$\lambda_i(\mathcal{A}_\bullet) \leq \lambda_i(\mathcal{A}_\bullet^V) \leq \lambda_{i+r}(\mathcal{A}_\bullet)$$

for $i = 1, \dots, n - r$ where n is the total dimension of \mathcal{A}_\bullet and $r = \sum_{v \in V} \dim \mathcal{A}_v$.

PROOF. Consider the Laplacian $L_{\mathcal{A}}$. It can be written in a block form as

$$L_{\mathcal{A}} = \begin{pmatrix} Q_0 \setminus V & V \\ X & Y^* \\ Y & D \end{pmatrix} \begin{matrix} Q_0 \setminus V \\ V \end{matrix}.$$

As there are no edges between vertices of V , D is a diagonal matrix where the (v, v) block is $d_v \alpha_v \text{id}$ where d_v is the degree of v . Thus D is invertible, and we can form the Schur complement

$$L_{\mathcal{A}}/D := X - Y^* D^{-1} Y.$$

A vector $[x \ z]^T$ satisfies

$$\begin{bmatrix} X & Y^* \\ Y & D \end{bmatrix} \begin{bmatrix} x \\ z \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

if and only if $Xx + Y^*z = 0$ and $Yx + Dz = 0$. As D is invertible this can be rearranged to $Xx - YD^{-1}Y^*x = 0$ hence $\ker(L_{\mathcal{A}}/D) \cong \ker L_{\mathcal{A}}$. For the interlacing, we can directly apply Theorem B.5.

We will now define the sheaf \mathcal{A}_\bullet^V . As in Definition 6.8, we will ignore orientation of the edges and write $u \sim e$ if either $s(e) = v$ or $t(e) = v$ and $\mathcal{A}_{v \sim e}$ for the corresponding assigned linear map. If $p \neq q$, define $\mathcal{A}_{u \sim pq}^V = (d_v \alpha_v)^{-1/2} \mathcal{A}_{v \sim p}^* \mathcal{A}_{u \sim p}$ and likewise for w . Otherwise, if $p = q$, define $\mathcal{A}_{u \rightarrow pp}^V = Z_p \mathcal{A}_{u \sim p}$ and $\mathcal{A}_{pp \leftarrow u}^V = \mathbf{0}$ for a linear map Z_p which we will now define. First observe that for every non-zero vector x , we have

$$\frac{1}{\alpha_v} x^* \mathcal{A}_{v \sim p} \mathcal{A}_{v \sim p}^* x \leq \frac{1}{\alpha_v} \|\mathcal{A}_{v \sim p}^*\|^2 \|x\|^2 = \frac{1}{\alpha_v} \|\mathcal{A}_{v \sim p}^* \mathcal{A}_{v \sim p}\| \|x\|^2 = \|x\|^2$$

thus there exists some Z_p such that

$$\text{id} - \frac{1}{\alpha_v} \mathcal{A}_{v \sim p} \mathcal{A}_{v \sim p}^* = Z_p^* Z_p.$$

since the left-hand side is positive semi-definite.

Now we claim that $L_{\mathcal{A}^V} = L_{\mathcal{A}}/D$. To do so, consider vertices $u \in Q_0 \setminus V$ and $v \in V$. For each edge $u \sim p \sim v$, the contribution to the diagonal block indexed by (u, u) of $L_{\mathcal{A}}/D$ is

$$\mathcal{A}_{u \sim p}^* \mathcal{A}_{u \sim p} - \frac{1}{d_v \alpha_v} \sum_{u \sim q \sim v} \mathcal{A}_{u \sim p}^* \mathcal{A}_{v \sim p} \mathcal{A}_{v \sim q}^* \mathcal{A}_{u \sim q}.$$

We need to show this is equal to the contribution from the new edges of \mathcal{A}_{\bullet}^Q . For each vertex $u \in Q_0 \setminus V$, the contribution to the diagonal block is

$$\begin{aligned} & \sum_{v \in V} \sum_{u \sim p \sim v} \left[\mathcal{A}_{u \sim p}^* \mathcal{A}_{u \sim p} - \frac{1}{d_v \alpha_v} \sum_{u \sim q \sim v} \mathcal{A}_{u \sim p}^* \mathcal{A}_{v \sim p} \mathcal{A}_{v \sim q}^* \mathcal{A}_{u \sim q} \right] \\ &= \sum_{v \in V} \left[\sum_{u \sim p \sim v} \left[\mathcal{A}_{u \sim p}^* \mathcal{A}_{u \sim p} - \frac{d_v^u}{d_v \alpha_v} \mathcal{A}_{u \sim p}^* \mathcal{A}_{v \sim p} \mathcal{A}_{v \sim p}^* \mathcal{A}_{u \sim p} \right] \right. \\ & \quad \left. + \sum_{\substack{u \sim p \sim v \\ u \sim q \sim v \\ p \neq q}} \frac{1}{d_v \alpha_v} (\mathcal{A}_{v \sim p}^* \mathcal{A}_{u \sim p} - \mathcal{A}_{v \sim q}^* \mathcal{A}_{u \sim q})^* (\mathcal{A}_{v \sim p}^* \mathcal{A}_{u \sim p} - \mathcal{A}_{v \sim q}^* \mathcal{A}_{u \sim q}) \right] \end{aligned}$$

where d_v^u is the number of edges between u and v . The term on the last line corresponds to the new self-loops on u given by v . By the defining property of Z_p , the term

$$\sum_{u \sim p \sim v} \left[\mathcal{A}_{u \sim p}^* \mathcal{A}_{u \sim p} - \frac{d_v^u}{d_v \alpha_v} \mathcal{A}_{u \sim p}^* \mathcal{A}_{v \sim p} \mathcal{A}_{v \sim p}^* \mathcal{A}_{u \sim p} \right]$$

equals

$$\sum_{u \sim p \sim v} \left[\frac{d_v - d_v^u}{d_v \alpha_v} \mathcal{A}_{u \sim p}^* \mathcal{A}_{v \sim p} \mathcal{A}_{v \sim p}^* \mathcal{A}_{u \sim p} + \mathcal{A}_{u \sim p}^* Z_p^* Z_p \mathcal{A}_{u \sim p} \right]$$

and thus corresponds with the contribution of the self-loop pp and the new edges to different vertices, as required. \square

7. Reduction of the Feature Selector Problem

Using the quiver operations in §6, we can reduce the number of vertices in the quiver whilst preserving the space of sections; this allows us to reduce the difficulty of computing approximate sections.

THEOREM 7.1. *Let \mathbf{S} be a feature selector on X , \mathcal{U} a cover of X such that $Q^{\mathcal{U}}$ is (weakly) connected, and $\mathcal{F} \subset X^*$. Fix a vertex $\sigma_0 \in Q_0^{\mathcal{F}}$. Then there exists a quiver Q' and a representation \mathbf{A}'_{\bullet} of Q' satisfying the following properties:*

- (1) *The quiver Q' has only one vertex and $|Q_0^{\mathcal{F}}|$ edges.*
- (2) *The Laplacian of \mathbf{A}'_{\bullet} is:*

$$\sum_{\tau \in Q_0^{\mathcal{F}}} \left[\text{id}_{\sigma_0} - \pi_{\sigma_0}^{\mathcal{F}} l_{\tau}^{\mathcal{F}} \pi_{\tau}^{\mathcal{F}} l_{\sigma_0}^{\mathcal{F}} \right]$$

(3) There exist constants $C_1, C_2 > 0$ that depend on \mathbf{S}, \mathcal{U} and \mathcal{F} such that

$$C_1 \lambda_i(\mathbf{N}_{\bullet}^{\mathbf{S}, \mathcal{F}}) \leq \lambda_i(\mathbf{A}_{\bullet}') \leq C_2 \lambda_{i+k}(\mathbf{N}_{\bullet}^{\mathbf{S}, \mathcal{F}})$$

for $i = 1, \dots, n - k$.

Here n is the total dimension of $\mathbf{N}_{\bullet}^{\mathbf{S}, \mathcal{F}}$ and $k := n - \dim \mathbf{A}_{\sigma_0}^{\mathbf{S}, \mathcal{F}}$.

PROOF. We view $\mathbf{N}_{\bullet}^{\mathbf{S}, \mathcal{F}}$ as a sheaf \mathcal{A}_{\bullet} over a quiver Q , defined as follows:

$$\begin{aligned} Q_0 &= \left\{ v_{\sigma} : \sigma \in Q_0^{\mathcal{U}} \right\} \cup \left\{ v_{\sigma\tau} : (\sigma \rightarrow \tau) \in Q_1^{\mathcal{U}} \right\}, \\ Q_1 &= \left\{ e_{\sigma\tau}, r_{\sigma\tau}^{DL}, r_{\sigma\tau}^{LD} : (\sigma \rightarrow \tau) \in Q_1^{\mathcal{U}} \right\}, \\ s(e_{\sigma\tau}) &= v_{\sigma}, \quad t(e_{\sigma\tau}) = v_{\tau}, \\ s(r_{\sigma\tau}^{DL}) &= s(r_{\sigma\tau}^{LD}) = v_{\sigma}, \quad t(r_{\sigma\tau}^{DL}) = t(r_{\sigma\tau}^{LD}) = v_{\sigma\tau}. \end{aligned}$$

This is precisely the merged quiver on which $\mathbf{N}_{\bullet}^{\mathbf{S}, \mathcal{F}}$ is defined, except that we have dispensed with identity self-loops in light of Proposition 6.7(a). The desired sheaf \mathcal{A}_{\bullet} assigns the following data to vertices and edges of Q for each $\sigma \in Q_0^{\mathcal{U}}$:

$$\begin{aligned} \mathcal{A}_{v_{\sigma} \rightarrow e_{\sigma\tau}} &= \pi_{\tau}^{\mathcal{F}} \circ \iota_{\sigma}^{\mathcal{F}}, & \mathcal{A}_{e_{\sigma\tau} \leftarrow v_{\tau}} &= \text{id}, & \mathcal{A}_{v_{\sigma} \rightarrow r_{\sigma\tau}^{DL}} &= \iota_{\tau}^{\mathcal{F}} \circ \pi_{\tau}^{\mathcal{F}} \circ \iota_{\sigma}^{\mathcal{F}}, \\ \mathcal{A}_{r_{\sigma\tau}^{DL} \leftarrow v_{\sigma}} &= \text{id}, & \mathcal{A}_{v_{\sigma} \rightarrow r_{\sigma\tau}^{LD}} &= \iota_{\sigma}^{\mathcal{F}}, & \mathcal{A}_{r_{\sigma\tau}^{DL} \leftarrow v_{\sigma}} &= \text{id}. \end{aligned}$$

Use Proposition 6.7(b) to get a new copy $e_{\sigma\tau}^2$ of the edge $e_{\sigma\tau}$ with the maps multiplied by $\iota_{\tau}^{\mathcal{F}}$. Use this edge with Proposition 6.5 to move the edge $r_{\sigma\tau}^{DL}$ to an edge between v_{τ} and $v_{\sigma\tau}$, then using the edge $r_{\sigma\tau}^{LD}$ with Proposition 6.5 move this edge between v_{τ} and $v_{\sigma\tau}$ to an edge between v_{σ} and v_{τ} which has maps ι_{σ} and ι_{τ} respectively. Call this edge $p_{\sigma\tau}$. Using Proposition 6.7(b) can remove both $e_{\sigma\tau}$ and $e_{\sigma\tau}^2$ since $\pi_{\tau}^{\mathcal{F}}$ is the left inverse of $\iota_{\tau}^{\mathcal{F}}$. Finally, we can remove the edge $r_{\sigma\tau}^{LD}$ with Proposition 6.9 as $(\iota_{\sigma}^{\mathcal{F}})^* \iota_{\sigma}^{\mathcal{F}} = \text{id}$. Now choose some vertex σ_0 . As the nerve is connected, and each edge from each vertex has the same restriction maps to that vertex we can keep applying Proposition 6.5 until all edges have σ_0 as a boundary. There might be double edges, but these can be removed with Proposition 6.7(b). Then we apply Proposition 6.9 to reduce to just the vertex σ_0 . The Laplacian is now

$$\sum_{\tau \in Q_0^{\mathcal{U}}} \left[\text{id}_{\sigma_0} - \pi_{\sigma_0}^{\mathcal{F}} \iota_{\tau}^{\mathcal{F}} \pi_{\tau}^{\mathcal{F}} \iota_{\sigma_0}^{\mathcal{F}} \right]$$

and the eigenvalue relationship follows by tracking all the uses of Propositions 6.7(a), 6.7(b), 6.5, 6.9. \square

THEOREM 7.2. *Let \mathbf{S} be a feature selector on a set X equipped with a cover \mathcal{U} . For each subset $\mathcal{F} \subset X^*$, here exists a sheaf \mathcal{A}' over a quiver Q' with Laplacian*

$$(L_{\mathcal{A}'})_{\sigma, \tau} = \begin{cases} \sum_{\tau \subset \sigma} (2\text{id}_{\sigma} - \pi_{\sigma}^{\mathcal{F}} \iota_{\tau}^{\mathcal{F}} \pi_{\tau}^{\mathcal{F}} \iota_{\sigma}^{\mathcal{F}}) + \sum_{\sigma \subset \gamma} \pi_{\sigma}^{\mathcal{F}} \iota_{\gamma}^{\mathcal{F}} \pi_{\gamma}^{\mathcal{F}} \iota_{\sigma}^{\mathcal{F}} & \text{if } \sigma = \tau, \\ -\pi_{\sigma}^{\mathcal{F}} \iota_{\tau}^{\mathcal{F}} & \text{if } \tau \subset \sigma \text{ or } \sigma \subset \tau, \\ 0 & \text{otherwise.} \end{cases}$$

such that there is an isomorphism $\Gamma(\mathcal{A}'_\bullet) \cong \Gamma(\mathbf{M}_\bullet^{\mathcal{S}, \mathcal{F}})$; and moreover, we have

$$\lambda_i(\mathbf{M}_\bullet^{\mathcal{S}, \mathcal{F}}) \leq \lambda_i(\mathcal{A}'_\bullet) \leq \lambda_{i+k}(\mathbf{N}_\bullet^{\mathcal{S}, \mathcal{F}})$$

for all $i \in \{1, \dots, n-k\}$. Here $\mathbf{M}_\bullet^{\mathcal{S}, \mathcal{F}}$ is the mixed representation of Definition 3.9, while n is its total dimension and $k := n - \sum_{\sigma \in Q_0} \dim \mathbf{A}_\sigma^{\mathcal{S}, \mathcal{F}}$.

PROOF. Consider the representation $(\mathbf{A}_\bullet^{\mathcal{S}, \mathcal{F}})^* \sqcup_R \widehat{\mathbf{A}}_\bullet^{\mathcal{S}, \mathcal{F}}$ as defined in Definition 3.9, and view it as a sheaf \mathcal{A}_\bullet over a quiver Q . We can immediately remove identity self-loops using Proposition 6.7(a). The quiver Q is as follows:

$$\begin{aligned} Q_0 &= \left\{ v_\sigma : \sigma \in Q_0^{\mathcal{U}} \right\} \cup \left\{ v_{\sigma\tau} : (\sigma \rightarrow \tau) \in Q_1^{\mathcal{U}} \right\}, \\ Q_1 &= \left\{ e_{\sigma\tau}, r_{\sigma\tau}^{DL}, r_{\sigma\tau}^{LD} : (\sigma \rightarrow \tau) \in Q_1^{\mathcal{U}} \right\}, \\ s(e_{\sigma\tau}) &= v_\tau, \quad t(e_{\sigma\tau}) = v_\sigma, \\ s(r_{\sigma\tau}^{DL}) &= s(r_{\sigma\tau}^{LD}) = v_\sigma, \quad t(r_{\sigma\tau}^{DL}) = t(r_{\sigma\tau}^{LD}) = v_{\sigma\tau}. \end{aligned}$$

The sheaf \mathcal{A}_\bullet consists of the following for each $\sigma \in Q^{\mathcal{U}}$:

$$\begin{aligned} \mathcal{A}_{v_\tau \rightarrow e_{\sigma\tau}} &= \pi_\sigma^{\mathcal{F}} \circ l_\tau^{\mathcal{F}}, & \mathcal{A}_{e_{\sigma\tau} \leftarrow v_\sigma} &= \text{id}, \\ \mathcal{A}_{v_\sigma \rightarrow r_{\sigma\tau}^{DL}} &= l_\tau^{\mathcal{F}} \circ \pi_\tau^{\mathcal{F}} \circ l_\sigma^{\mathcal{F}}, & \mathcal{A}_{r_{\sigma\tau}^{DL} \leftarrow v_\sigma} &= \text{id}, \\ \mathcal{A}_{v_\sigma \rightarrow r_{\sigma\tau}^{LD}} &= l_\sigma^{\mathcal{F}}, & \mathcal{A}_{r_{\sigma\tau}^{LD} \leftarrow v_\sigma} &= \text{id}. \end{aligned}$$

We can use Proposition 6.9 to remove each $v_{\sigma\tau}$. In doing so, we add three new self-loops to v_σ ; two identities and a self loop with maps $l_\sigma^{\mathcal{F}}$ and $l_\tau^{\mathcal{F}} \circ \pi_\tau^{\mathcal{F}} \circ l_\sigma^{\mathcal{F}}$. The two identity self loops can be removed with Proposition 6.7(a). \square

7.1. Remarks on computation. In Theorems 7.1 and 7.2, we have assumed that L is Hermitian with respect to some orthonormal basis. Often, however, the bases we would prefer to use are not orthogonal with respect to our chosen inner product. We will see such an example in §8, where computing inner products is straightforward whilst finding an orthonormal basis is computationally infeasible. It is then advantageous then to compute eigenvectors of a Laplacian without such a change of basis. Suppose, for each $\sigma \in \mathcal{U}$, we have some basis \mathcal{B}_σ of $\mathbf{S}_{|\sigma|}^{\mathcal{F}}$ and a Hermitian positive definite matrix $M_\sigma^{\mathcal{S}, \mathcal{F}}$ representing the inner product in this basis. As $M_\sigma^{\mathcal{S}, \mathcal{F}}$ is positive definite there exists an invertible W_σ such that $M_\sigma^{\mathcal{S}, \mathcal{F}} = W_\sigma^* W_\sigma$. Then, in this basis, we now have for $\sigma, \tau \in \mathcal{U}$

$$\begin{aligned} \pi_\sigma^{\mathcal{F}} l_\tau^{\mathcal{F}} &= W_\sigma^{-*} M_{\sigma\tau}^{\mathcal{S}, \mathcal{F}} W_\tau^{-1}, \\ \pi_\sigma^{\mathcal{F}} l_\tau^{\mathcal{F}} \pi_\tau^{\mathcal{F}} l_\sigma^{\mathcal{F}} &= W_\sigma^{-*} M_{\sigma\tau}^{\mathcal{S}, \mathcal{F}} (M_\tau^{\mathcal{S}, \mathcal{F}})^{-1} M_{\tau\sigma}^{\mathcal{S}, \mathcal{F}} W_\sigma^{-1} \end{aligned}$$

where $M_{\sigma\tau}^{\mathcal{S}, \mathcal{F}}$ is the matrix of inner products between \mathcal{B}_σ and \mathcal{B}_τ . Now computing eigenvectors of a Laplacian L in this basis can be realised as the generalised eigenvector problem

$$Lx = \lambda M_{\mathcal{U}}^{\mathcal{S}, \mathcal{F}} x$$

where $M_{\mathcal{U}}^{\mathcal{S}, \mathcal{F}}$ is the block diagonal matrix with blocks $M_\sigma^{\mathcal{S}, \mathcal{F}}$ for $\sigma \in \mathcal{U}$.

As the matrices $M_{\tau}^{\mathbf{S}, \mathcal{F}}$ are Hermitian and positive-definite, the products

$$M_{\sigma\tau}^{\mathbf{S}, \mathcal{F}} (M_{\tau}^{\mathbf{S}, \mathcal{F}})^{-1} M_{\tau\sigma}^{\mathbf{S}, \mathcal{F}}$$

can be quickly computed using the Cholesky decomposition of $M_{\tau}^{\mathbf{S}, \mathcal{F}}$. Furthermore, as L is Hermitian and usually sparse, there exist computationally efficient methods for calculating its smallest eigenvalues and their eigenvectors, e.g. variations on the Lanczos method [21], combined with the Cholesky decomposition (see e.g. [18]) of $L + \alpha\mathbf{I}$ for some small $\alpha > 0$.

8. Single-Cell Chromatin Accessibility

Chromatin is a macromolecule consisting of DNA and certain proteins, primarily histones [4]. Histones facilitate the dense compaction of DNA, which allows it to fit inside limited space within cells. The fundamental unit of chromatin is a *nucleosome*, which consists of approximately 147 *base-pairs* (namely, the usual A, C, G and T nucleotides) wound around histone molecules. The density of nucleosomes varies considerably across the genome; low density regions are more accessible to transcription factors and other binding proteins. Thus, understanding the relative accessibility of various parts of the genome is important in the study of gene regulation and its variation across different cell types.

Among the most promising experimental frameworks for measuring chromatin accessibility is **ATAC-seq** [1], which has been extended to single cell sequencing **scATAC-seq** (see e.g. [28]). The output of scATAC-seq may be viewed from a mathematical perspective as follows. Consider a set X of cells $\{x_1, \dots, x_N\}$, with N being on the order of thousands. Each cell x_i has a genome consisting of M base pairs, where M is approximately three billion. The output of ATAC-seq on X is an $N \times M$ integer matrix. For each $1 \leq i \leq N$ and $1 \leq j \leq M$, the entry $\alpha_{i,j}$ is a number in $\{0, 1, 2\}$ representing the accessibility (for cell x_i at the j -th genomic position). Such data is high-dimensional, noisy and sparse, and hence warrants feature selection. This is accomplished via a class of techniques called **peak-calling** algorithms [36]. The goal of these methods is to identify specific locations along the genome, called *summits*, which are significantly accessible across the cells in X — see for instance the MACS2 algorithm [39].

REMARK 8.1. We are often interested in quantifying accessibility when one restricts to various sub-populations $Y \subset X$ of cell types. However, the peaks of various Y may be lost when calling peaks on X , particularly when the cardinalities of the X under consideration are unbalanced. The standard remedy is to first cluster the data and then call peaks on each cluster [28]. Subsequently, peaks from all the clusters are combined, with various rules being applied for grouping and extending peaks which overlap. We show below how the quiver Laplacian provides a different mechanism for achieving similar results.

8.1. Pre-processing the data. We use the single-cell ATAC-seq data of [28], which consist of 28,274 tumour-infiltrating T cells. The pre-processing steps are the same as those of *ibid*:

- (1) We align fragments to the GRCH37/HG19 assembly.
- (2) We filter sample profiles to keep those with at least 1,000 fragments and a transcription start site (TSS) enrichment score of 8.
- (3) We create a count matrix by tiling the genome with 2,500 base-pair windows and counting overlaps of endpoints of fragments with the windows.
- (4) Next, we binarise this matrix and retain only the top 20,000 most accessible sites; call the resulting $28,724 \times 20,000$ boolean matrix B .
- (5) We then create a frequency-inverse document frequency matrix, apply truncated singular value decomposition to obtain the 2nd to 25th singular vectors. This results in a reduced matrix E of size 28724×24 .

8.2. Generating the quiver. We build an undirected graph G from the reduced matrix E as follows. Its vertices are the cells C_i , which correspond to rows of E . Treating each such row as a vector in \mathbb{R}^{24} , let $\delta_{i,j} \geq 0$ be the Euclidean distance between rows i and j for each $i < j$. We construct the 15 nearest neighbour graph generated by the $\delta_{i,j}$ distances. Our analysis now deviates from [28]. Using the Leiden algorithm [33] the vertex set V of G is partitioned into finitely many communities, $V = \coprod_i V_i$. As in *ibid*, communities of size < 200 are excluded. We then generate a cover \mathcal{U} of X from the surviving communities as follows: each cover element $U_i \in \mathcal{U}$ has the form $V_i \cup W_i$, where $W_i \subset (V - V_i)$ consists of those vertices which share an edge with some vertex of V_i . The resulting quiver $Q^{\mathcal{U}}$ has 87 vertices, and is depicted in Figure 6A. Similarly, Figure 6B displays the contents of each vertex relative to the overall distribution, using the cell types identified in [28].

8.3. Building the representation. Having constructed $Q^{\mathcal{U}}$, we apply our feature selector \mathbf{S} – which in this case is the MACS2 peak-calling algorithm [39] – to the fragment profiles of cells corresponding to each individual vertex. In the notation of Section 2, the input \mathcal{F} is given by the columns of the binarised matrix B ; explicitly, the map $f_j : X \rightarrow \mathbb{R}$ associated to the j -th column of E sends each x_i to the entry $B_{ij} \in \{0, 1\}$. Calling peaks on this set of features produces a vector space $\mathbf{S}_{|\sigma|} \subset \mathbf{V}(\mathcal{F})$ for every vertex σ of $Q^{\mathcal{U}}$. We used the same parameters as in [28], and obtained several hundred thousand summits for each vertex as a result. For the purposes of this demonstration, we retained only the most significant 1,000 summits per vertex. Thus the total dimension of the selected vector spaces is 87,000. Since peak-calling on different sub-populations may change the true summit location on the genome slightly, we follow [38] and use a Gaussian kernel between summits (with bandwidth equal to 1,000 base-pairs) as the inner product on $\mathbf{V}(\mathcal{F})$. This allows us to build the mixed representation $\mathbf{M}_{\bullet}^{\mathbf{S}, \mathcal{F}}$ (see Definition 3.9).

8.4. Results. We computed the first 3,000 eigenvectors of the Laplacian of $\mathbf{M}_{\bullet}^{\mathbf{S}, \mathcal{F}}$. Since the total space admits a distinguished basis whose vectors correspond to summits (i.e., certain genomic positions $\{p_k\}$), we may associate to each eigenvector $v = \sum_k \alpha_k p_k$ the subset of supported genomic positions $P_v := \{k \mid \alpha_k \neq 0\}$ ⁵. Figure 6D shows that for

⁵Or more precisely, $|\alpha_k| > 10^{-5}$, the tolerance parameter for ARPACK-NG [22] in our experiment.

an overwhelming majority of eigenvectors v , the diameter of P_v for most v approximately equals 147 base-pairs, which is the average length of DNA around a nucleosome. A small fraction of the eigenspaces had dimension > 1 ; this resulted in a mixing of genomic positions, as shown by the blue outliers in Figure 6D. Nevertheless, in almost all cases, we observed that **the genomic supports of eigenvectors are tightly aligned across the vertices of $Q^{\mathcal{Z}}$** .

We then performed a more detailed analysis of the sparsity pattern of the 3,000 eigenvectors. The outcome of this analysis has been summarised in Figure 6E, which plots the 3,000 eigenvectors against all 87 vertices of $Q^{\mathcal{Z}}$ and highlights those vertices along which a given eigenvector has a non-zero component. We observe that the eigenvectors fall broadly into two classes. The first class, which is considerably larger, consists of those **eigenvectors whose genomic support contains summits which are shared by most, if not all vertices**: see, for instance, the components of eigenvector 11 in Figure 7. The second class contains **eigenvectors whose genomic support contains summits which are shared by a small fraction of the vertices**. A prime example of this type is eigenvector 2702, which has nonzero components along only 3 of the 87 vertices, and has also been depicted in Figure 7.

Finally, we note that (at least for these two examples) **the size of the support of the eigenvector relates to biological function**. Eigenvector 11, for instance, is concentrated in a regulatory region a few hundred base pairs downstream of the gene *FAM72D*. This gene is involved in the cell cycle [9], and is therefore relevant to most subpopulations of cells. On the other hand, eigenvector 2702 lies just downstream of the gene *STAT3*, which is a transcription factor known to be an important regulator for differentiation and proliferation in Th1, Th17, and Treg cells [27]. We also note that this eigenvector was only supported on the vertices 8, (8, 11), and 11 (Figure 7). These vertices are particularly enriched in Th1, Th17 and Treg cells (see Figure 6B). Indeed, Figure 6E reveals that many other eigenvectors are also uniquely supported on the vertices 8, (8, 11), 11 and (1, 11). Since vertex 11 is relatively disconnected from the other vertices, our analysis highlights it as an unusual sub-population of X .

Code and data availability

Code produced for the analysis will be available at

https://github.com/osumray/harmonic_feature_selection.

The data from [28] is available from the Gene Expression Omnibus with accession number [GSE129785](#).

Appendix A. Removing Edges and Vertices

The following extends known results from graphs, e.g. see [34] and [35] respectively.

PROPOSITION A.1. *Let \mathcal{A}_\bullet be a sheaf over a quiver Q , and let n be its total dimension.*

A.1(a) Let $R \subset Q_1$ be a subset of edges of Q . Let Q' be the quiver with edges $Q_1 \setminus R$ and \mathcal{A}' the restriction of \mathcal{A} to Q' . Then

$$\lambda_{k+i}(\mathcal{A}_\bullet) \leq \lambda_{k+i}(\mathcal{A}'_\bullet) \leq \lambda_{k+r+i}(\mathcal{A}_\bullet)$$

for $i = 1, \dots, n - k - r$ where $k = \dim \Gamma(\mathcal{A}_\bullet)$ and $r = \sum_{e \in R} \dim \mathcal{A}_e$.

A.1(b) Suppose $V \subset Q_0$ is a subset of vertices of Q . Let Q' be the quiver with vertices $Q_0 \setminus V$ and all edges $e \in Q_1$ such that either $s(e) \in V$ or $t(e) \in V$, and \mathcal{A}'_\bullet the restriction of \mathcal{A}_\bullet to Q' . Then

$$\lambda_i(\mathcal{A}_\bullet) - w_1 \leq \lambda_i(\mathcal{A}'_\bullet) \leq \lambda_{k+i}(\mathcal{A}_\bullet) - w_2$$

for $i = 1, \dots, n - k$ where $k = \sum_{u \in Q_0 \setminus V} \dim \mathcal{A}_u$ and

$$w_1 = \max_{v \in V} \left\{ \sum_{\substack{s(e)=v \\ t(e) \notin V}} \sigma_{\max}(\mathcal{A}_{v \rightarrow e})^2 + \sum_{\substack{t(e)=v \\ s(e) \notin V}} \sigma_{\max}(\mathcal{A}_{e \leftarrow v})^2 \right\},$$

$$w_2 = \min_{v \in V} \left\{ \sum_{\substack{s(e)=v \\ t(e) \notin V}} \sigma_{\min}(\mathcal{A}_{v \rightarrow e})^2 + \sum_{\substack{t(e)=v \\ s(e) \notin V}} \sigma_{\min}(\mathcal{A}_{e \leftarrow v})^2 \right\}.$$

PROOF OF PROPOSITION A.1(A). Define $L_{\mathcal{A}, \text{op}} = B_{\mathcal{A}} B_{\mathcal{A}}^*$. Let $k = \dim \ker L_{\mathcal{A}}$ and $c = \dim \ker L_{\mathcal{A}, \text{op}}$. Then for $i = 1, \dots, n - k$, we have that

$$\lambda_{k+i}(L_{\mathcal{A}}) = \lambda_{c+i}(L_{\mathcal{A}, \text{op}}).$$

Now, removing the set of edges $R \subset Q_1$ corresponds to the $r \times r$ principal submatrix of $L_{\mathcal{A}, \text{op}}$, where $r := \sum_{e \in R} \dim \mathcal{A}_e$. By Theorem B.3 we have

$$\lambda_i(L_{\mathcal{A}, \text{op}}) \leq \lambda_i(L_{\mathcal{A}', \text{op}}) \leq \lambda_{i+r}(L_{\mathcal{A}, \text{op}})$$

where $i \in \{1, \dots, m - r\}$ thus

$$\lambda_{k+i}(\mathcal{A}_\bullet) \leq \lambda_{k+i}(\mathcal{A}'_\bullet) \leq \lambda_{k+r+i}(\mathcal{A}_\bullet)$$

for $i \in \{1, \dots, n - k - r\}$. □

PROOF OF PROPOSITION A.1(B). The boundary matrix $B_{\mathcal{A}}$ has block form

$$\begin{bmatrix} B_{11} & \mathbf{0} \\ B_{21} & B_{22} \end{bmatrix}$$

where $B_{11} = B_{\mathcal{A}'}$. Here, $B_{\mathcal{A}}$ has row blocks indexed by edges between vertices in V and column blocks indexed by vertices in V . Then Laplacian $L_{\mathcal{A}} = B_{\mathcal{A}}^* B_{\mathcal{A}}$ has block form

$$\begin{bmatrix} B_{11}^* B_{11} + B_{21}^* B_{21} & B_{21}^* B_{22} \\ B_{22}^* B_{21} & B_{22}^* B_{22} \end{bmatrix}.$$

Define $L_{Q^0 \setminus V} = B_{11}^* B_{11} + B_{21}^* B_{21}$ which is a principal submatrix of $L_{\mathcal{A}}$. Note that $B_{21}^* B_{21}$ is a diagonal matrix: The v, v block is

$$\sum_{\substack{s(e)=v \\ t(e) \neq V}} \mathcal{A}_{v \rightarrow e}^* \mathcal{A}_{v \rightarrow e} + \sum_{\substack{t(e)=v \\ s(e) \neq V}} \mathcal{A}_{e \leftarrow v}^* \mathcal{A}_{e \leftarrow v}.$$

We now have that

$$\begin{aligned}
 \lambda_i(\mathcal{A}') &= \lambda_i(B_{11}^* B_{11}) = \lambda_i(L_{Q_0 \setminus V} - B_{21}^* B_{21}) \leq \lambda_i(L_{Q_0 \setminus V}) + \lambda_{n-k}(-B_{21}^* B_{21}) \\
 &= \lambda_i(L_{Q_0 \setminus V}) - \lambda_1(B_{21}^* B_{21}) \\
 &= \lambda_i(L_{Q_0 \setminus V}) - \min_{v \in V} \lambda_1 \left(\sum_{\substack{s(e)=v \\ t(e) \neq V}} \mathcal{A}_{v \rightarrow e}^* \mathcal{A}_{v \rightarrow e} + \sum_{\substack{t(e)=v \\ s(e) \neq V}} \mathcal{A}_{e \leftarrow v}^* \mathcal{A}_{e \leftarrow v} \right) \\
 &\leq \lambda_i(L_{Q_0 \setminus V}) - \min_{v \in V} \left\{ \sum_{\substack{s(e)=v \\ t(e) \neq V}} \lambda_1(\mathcal{A}_{v \rightarrow e}^* \mathcal{A}_{v \rightarrow e}) + \sum_{\substack{t(e)=v \\ s(e) \neq V}} \lambda_1(\mathcal{A}_{e \leftarrow v}^* \mathcal{A}_{e \leftarrow v}) \right\} \\
 &= \lambda_i(L_{Q_0 \setminus V}) - \min_{v \in V} \left\{ \sum_{\substack{s(e)=v \\ t(e) \neq V}} \sigma_{\min}(\mathcal{A}_{v \rightarrow e})^2 + \sum_{\substack{t(e)=v \\ s(e) \neq V}} \sigma_{\min}(\mathcal{A}_{e \leftarrow v})^2 \right\} \\
 &\leq \lambda_{i+k}(L_{\mathcal{A}}) - \min_{v \in V} \left\{ \sum_{\substack{s(e)=v \\ t(e) \neq V}} \sigma_{\min}(\mathcal{A}_{v \rightarrow e})^2 + \sum_{\substack{t(e)=v \\ s(e) \neq V}} \sigma_{\min}(\mathcal{A}_{e \leftarrow v})^2 \right\}.
 \end{aligned}$$

The first two inequalities are from Theorem B.4 and the final inequality is due to Theorem B.3. The right inequality is proven similarly. \square

Appendix B. Classical Results on Eigenvalues of Hermitian Matrices

In the following, for an $n \times n$ Hermitian matrix A , let $\lambda_1(A), \dots, \lambda_n(A)$ be the eigenvalues ordered in increasing fashion.

THEOREM B.1 (Rayleigh, see [18] Theorem 4.2.2). *Let A be a Hermitian matrix. Then for integers $1 \leq i_1, \dots, i_k \leq n$ let x_{i_1}, \dots, x_{i_k} be orthonormal, where x_{i_j} is an eigenvector corresponding to $\lambda_{i_j}(A)$. Define $S = \text{span}\{x_{i_1}, \dots, x_{i_k}\}$. Then for any $x \in S$*

$$\lambda_{i_1}(A) \leq \frac{x^* A x}{\|x\|^2} \leq \lambda_{i_k}(A).$$

THEOREM B.2 (Courant-Fischer, see [18] Theorem 4.2.6). *Let A be Hermitian matrix. Then*

$$\lambda_k(A) = \min_{\dim V=k} \left[\max_{x \in V \setminus 0} \frac{\langle x, Ax \rangle}{\langle x, x \rangle} \right].$$

THEOREM B.3 (Cauchy's interlacing theorem, see [18] Theorem 4.3.17). *Let A be a Hermitian matrix of size $n \times n$ and B an $m \times m$ principal submatrix of A . Then*

$$\lambda_i(A) \leq \lambda_i(B) \leq \lambda_{i+n-m}(A)$$

for $i = 1, \dots, m$.

THEOREM B.4 (Weyl's inequalities, see [18] Theorem 4.3.1). *For A, B Hermitian matrices of size $n \times n$ the following inequalities hold:*

$$\lambda_i(A + B) \leq \lambda_{i+j}(A) + \lambda_{n-j}(B)$$

for $j = 0, \dots, n - 1$ and

$$\lambda_i(A + B) \geq \lambda_{i-j+1}(A) + \lambda_j(B)$$

for $j = 1, \dots, i$.

THEOREM B.5 (See [31] Theorem 5). *For H a Hermitian matrix of dimension n given in block form as*

$$H = \begin{bmatrix} A & B^* \\ B & D \end{bmatrix}$$

where D is non-singular let $H/D = A - BD^{-1}B^*$ be the Schur complement of D . Then the eigenvalues of H and H/D are interlaced:

$$\lambda_i(H) \leq \lambda_i(H/D) \leq \lambda_{i+r}(H)$$

for $i = 1, \dots, n - r$ where r is the rank of D .

References

- [1] J. D. Buenrostro, P. G. Giresi, L. C. Zaba, H. Y. Chang, and W. J. Greenleaf. Transposition of native chromatin for fast and sensitive epigenomic profiling of open chromatin, DNA-binding proteins and nucleosome position. *Nature methods*, 10(12):1213–1218, 2013.
- [2] F. R. Chung. The Laplacian of a hypergraph. *Expanding graphs*, 10:21–36, 1992.
- [3] F. R. Chung. *Spectral graph theory*, volume 92. American Mathematical Soc., 1997.
- [4] G. Cooper and K. Adams. *The cell: a molecular approach*. Oxford University Press, 2022.
- [5] J. Curry, R. Ghrist, and V. Nanda. Discrete Morse theory for computing cellular sheaf cohomology. *Foundations of Computational Mathematics*, 16:875–897, 2016.
- [6] J. M. Curry. *Sheaves, cosheaves and applications*. University of Pennsylvania, 2014.
- [7] F. Dorfler and F. Bullo. Synchronization of power networks: Network reduction and effective resistance. *IFAC Proceedings Volumes*, 43(19):197–202, 2010.
- [8] S. Eilenberg and N. Steenrod. *Foundations of algebraic topology*, volume 2193. Princeton University Press, 2015.
- [9] M. Fischer, P. Grossmann, M. Padi, and J. A. DeCaprio. Integration of TP53, DREAM, MMB-FOXM1 and RB-E2F target gene analyses identifies cell cycle gene regulatory networks. *Nucleic acids research*, 44(13):6070–6086, 2016.
- [10] T. Gao, J. Brodzki, and S. Mukherjee. The geometry of synchronization problems and learning group actions. *Discrete & Computational Geometry*, 65(1):150–211, 2021.
- [11] R. Ghrist and H. Riess. Cellular sheaves of lattices and the Tarski Laplacian. *arXiv preprint arXiv:2007.04099*, 2020.
- [12] T. E. Goldberg. *Combinatorial Laplacians of Simplicial Complexes*. Senior thesis, Bard College, 2002.
- [13] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, 3rd ed edition, 1996.
- [14] J. Gu, B. Hua, and S. Liu. Spectral distances on graphs. *Discrete Applied Mathematics*, 190:56–74, 2015.
- [15] I. Guyon, J. Weston, S. Barnhill, and V. Vapnik. Gene selection for cancer classification using support vector machines. *Machine learning*, 46:389–422, 2002.
- [16] J. Hansen and R. Ghrist. Toward a Spectral Theory of Cellular Sheaves. *Journal of Applied and Computational Topology*, 3(4), 2019.
- [17] A. N. Hirani. *Discrete exterior calculus*. California Institute of Technology, 2003.
- [18] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, second edition, corrected reprint edition, 2017.

- [19] I. T. Jolliffe. *Principal Components Analysis, 2nd Ed.* Springer, 2002.
- [20] C. A. Joslyn, L. Charles, C. DePerno, N. Gould, K. Nowak, B. Praggastis, E. Purvine, M. Robinson, J. Strules, and P. Whitney. A sheaf theoretical approach to uncertainty quantification of heterogeneous geolocation information. *Sensors*, 20(12):3418, 2020.
- [21] C. Lanczos. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *United States Governm. Press Office Los Angeles, CA*, 1950.
- [22] R. Lehoucq, K. Maschhoff, D. Sorensen, C. Yang, F. Houssen, S. Ledru, et al. ARPACK-NG: Large scale eigenvalue problem solver. *Astrophysics Source Code Library*, pages ascl-2306, 2023.
- [23] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- [24] T. Muller, A. Seigal, and V. Nanda. Multilinear hyperquiver representations. *arXiv:2305.05622v2 [math.AG]*, 2023.
- [25] A. Parada-Mayorga, H. Riess, A. Ribeiro, and R. Ghrist. Quiver signal processing (QSP). *arXiv preprint arXiv:2010.11525*, 2020.
- [26] K. Pearson F.R.S. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, 1901.
- [27] C. Rébé and F. Ghiringhelli. STAT3, a master regulator of anti-tumor immune response. *Cancers*, 11(9):1280, 2019.
- [28] A. T. Satpathy, J. M. Granja, K. E. Yost, Y. Qi, F. Meschi, G. P. McDermott, B. N. Olsen, M. R. Mumbach, S. E. Pierce, M. R. Corces, P. Shah, J. C. Bell, D. Jhutti, C. M. Nemeč, J. Wang, L. Wang, Y. Yin, P. G. Giresi, A. L. S. Chang, G. X. Y. Zheng, W. J. Greenleaf, and H. Y. Chang. Massively parallel single-cell chromatin landscapes of human immune cell development and intratumoral T cell exhaustion. *Nature Biotechnology*, 37(8):925–936, 2019.
- [29] R. Schiffler. *Quiver Representations*. Number 184 in CMS Books in Mathematics. Springer, 2014.
- [30] A. Seigal, H. A. Harrington, and V. Nanda. Principal components along quiver representations. *Foundations of Computational Mathematics*, 23(4):1129–1165, 2023.
- [31] R. L. Smith. Some interlacing properties of the Schur complement of a Hermitian matrix. *Linear Algebra and its Applications*, 177:137–144, 1992.
- [32] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58(1):267–288, 1996.
- [33] V. A. Traag, L. Waltman, and N. J. Van Eck. From Louvain to Leiden: guaranteeing well-connected communities. *Scientific reports*, 9(1):5233, 2019.
- [34] J. Van Den Heuvel. Hamilton cycles and eigenvalues of graphs. *Linear algebra and its applications*, 226:723–730, 1995.
- [35] B. Wu, J. Shao, and X. Yuan. Deleting vertices and interlacing Laplacian eigenvalues. *Chinese Annals of Mathematics, Series B*, 31(2):231–236, 2010.
- [36] F. Yan, D. R. Powell, D. J. Curtis, and N. C. Wong. From reads to insight: a hitchhiker’s guide to ATAC-seq data analysis. *Genome biology*, 21:1–16, 2020.
- [37] K. Ye and L.-H. Lim. Schubert varieties and distances between subspaces of different dimensions. *SIAM Journal on Matrix Analysis and Applications*, 37(3):1176–1197, 2016.
- [38] J. Yu, D. Sun, Z. Hou, and L.-Y. Wu. Single-cell ATAC-seq analysis via network refinement with peaks location information. *bioRxiv*, page 2022.11.18.517159, 2022.
- [39] Y. Zhang, T. Liu, C. A. Meyer, J. Eeckhoute, D. S. Johnson, B. E. Bernstein, C. Nusbaum, R. M. Myers, M. Brown, W. Li, et al. Model-based analysis of ChIP-Seq (MACS). *Genome biology*, 9(9):1–9, 2008.

FIGURE 6. Overview of Laplacian from 7.2 applied to tumour-infiltrating T cells. (A) Quiver constructed from cover, with vertex size corresponding to the cardinality of the covering set. (B) Log likelihood of each cell type label within each vertex. (C) First 3,000 eigenvalues. (D) Range of genomic positions on each chromosome corresponding to non-zero components of the first 3,000 eigenvectors. Blue points correspond to positions across multiple chromosomes. (E) Binary matrix corresponding to the existence of a non-zero component in each vertex for the first 3,000 eigenvectors.

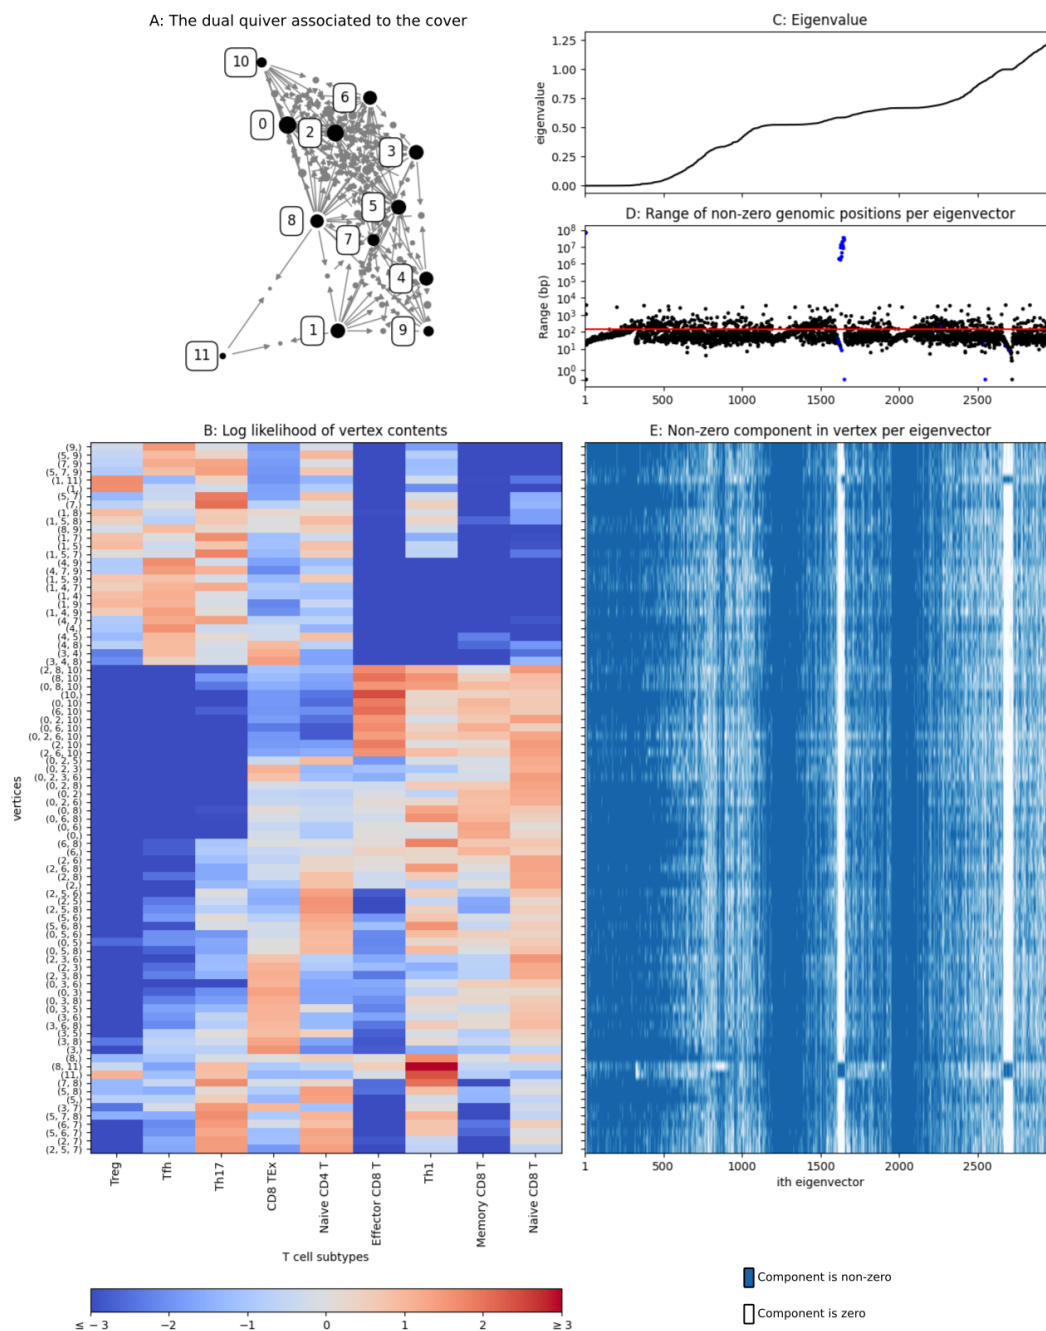


FIGURE 7. Examples of eigenvectors. Eigenvector #11 has a non-zero component at each vertex, whilst eigenvector #2702 has only non-zero components in vertices (8,), (11,) and (8, 11).

