

# Multimodal Network Alignment

Huda Nassar\*

David F. Gleich\*

## Abstract

A multimodal network encodes relationships between the same set of nodes in multiple settings, and network alignment is a powerful tool for transferring information and insight between a pair of networks. We propose a method for multimodal network alignment that computes a matrix which indicates the alignment, but produces the result as a low-rank factorization directly. We then propose new methods to compute approximate maximum weight matchings of low-rank matrices to produce an alignment. We evaluate our approach by applying it on synthetic networks and use it to de-anonymize a multimodal transportation network.

## Keywords

network alignment, low-rank matching, multimodal data

## 1 Introduction & Motivation

Network alignment is a technique to identify related nodes between two distinct networks. The result is a formal alignment or matching between the vertices of each network such that a vertex can only be aligned or matched to a single vertex in the opposite network. This methodology has its roots in the domains of ontology alignment [9], protein-protein network analysis [10, 22], social network de-anonymization [15], and object recognition [5]. The results of network alignment are often used for hypothetical information transfer. For instance, if we understand how a particular experiment disrupts a subnetwork of protein interactions in a mouse and we know the related interactions in a human, we have evidence for the hypothesis that the experiment would disrupt the mapped interactions in a human.

The network alignment problem and methods have traditionally been formulated for a pair of networks. Recently, much of the network data emerging in scientific applications and engineering studies is multimodal and contains separate types of relational information between vertices [23, 4, 1, 20, 19]. For instance, transportation networks are often described as multimodal when they feature multiple interconnecting but distinct networks,

such as when subway lines reach airports and national train systems. In biology, multimodal networks show different types of relationships that can occur between proteins and genes such as protein sequence similarity, co-occurrence in genes, co-occurrence in scientific papers, experimental interaction and more.

This motivates our present manuscript. We seek to produce effective methods to align a pair of multimodal networks. We will be more specific about our notion of a multimodal network in the following section. For now, consider multimodal networks in biology where proteins have many different types of relationships as described above. We may, for instance, be interested in aligning these networks generated from two species to infer related complexes and systems. In this case, the modes of each multimodal network are shared between the two networks to align.

We propose a new method that takes in a pair of multimodal networks with the same set of modes and produces an alignment between the vertices encoded in the network with no apriori knowledge of the relationship between the vertices. Network alignment problems are notoriously difficult from a theoretical perspective as they generalize the subgraph isomorphism problem. In this context, principled heuristic solutions abound [17, 22, 5]. Our method is a principled and effective heuristic that generalizes the ideas utilized in the IsoRank method [22] in concert with the insights from the network similarity decomposition [13]. It is designed for multimodal networks with 10s-1000s of modes with 1,000-100,000 vertices.

On top of this new method, we contribute a new theory about how to take the output from our heuristics and efficiently turn it into a matching between the graphs. This solves one of the key challenges in how to deal with multimodal networks and alignment. Many good ideas result in a polynomial explosion in problem size. This causes methods that have been developed for the pairwise case to be nearly impossible to use for the multimodal setup we propose [2, 11]. More specifically, this is a memory bottleneck: they would need terabytes of memory to handle problems that start off as a megabyte of data. We state this new theory as an independently useful primitive of finding a matching in a low-rank matrix. For this problem, we investigate a

\*Dept. of Computer Science, Purdue University. Email: {hnassar, dgleich}@purdue.edu, Supported by DARPA SIMPLEX and NSF award CCF-1149756, IIS-1422918, IIS-1546488, CCF-093937 and the Sloan Foundation.

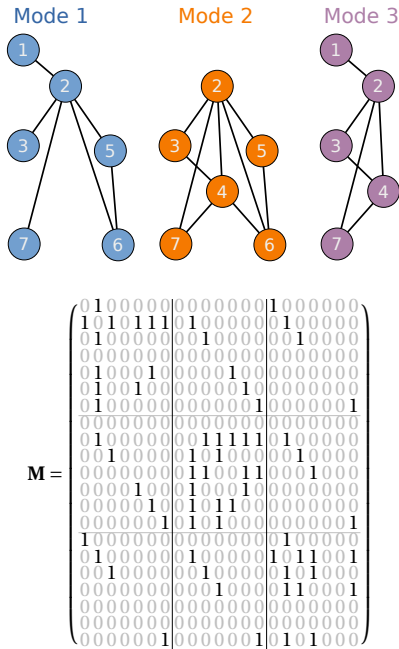


Figure 1: A multimodal network consists of a common set of vertices identifiers (in this case, 1-7) and a set of modes that define edges. Below, we illustrate the multimodal adjacency matrix, which we will use to align the multimodal networks.

highly efficient  $1/k$  approximation for a rank- $k$  matrix that finds the matching by using the low rank factors of the matrix *only*.

One of the difficulties with network alignment is that effective solutions for applications often exploit and use features of the data arising in those applications, which may not be generalizable to other instances. For instance, in biology, using information about the relationship between the protein sequences is vital to attain the best solutions [22, 17]. Thus, when we seek to evaluate our new methods, we will do so in (i) detailed synthetic experiments and (ii) a case-study and demonstration of our technique where we de-anonymize a publicly released transportation network that lacks meaningful vertex identifiers [20]. The goal of the synthetic experiments is to address the hypothesis that using multimodal alignment is better than straightforward generalizations of using existing network alignment methods to pairwise align each mode or a combined network. In the case of the transportation data, only our multimodal alignment method is able to completely map the data between the public data and the original database with labels. Our goal with both of these experiments is to reveal properties of our method and ideas that would be useful in any domain-specific application of network alignment but without the engineering that tends to occur in these applications. Towards that end, we do no tuning or

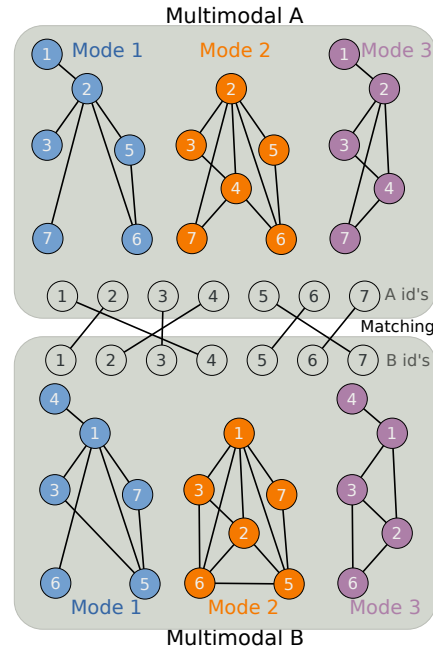


Figure 2: An illustration of the multimodal alignment problem. The goal is to identify the matching illustrated between the node id's given only the multimodal networks  $A$  and  $B$ . We assume the modes have a known alignment. The overlap of this matching is 20.

parameter selection.

In short, the contributions of our manuscript are:

1. A precise statement of a multimodal network alignment problem (Section 2).
2. A new method, multimodal similarity decomposition (MSD), to solve multimodal network alignment problems that uses all the information among the modes (Section 4) and goes beyond baseline methods that align the modes individually (Section 6).
3. Independently useful results about how to approximate a bipartite max-weight matching where the matching matrix is low-rank (Section 5).
4. We find that simple methods for multimodal alignment outperforms strong methods for pairwise network alignment of the individual modes when there are many vertices missing in each mode (Section 7).
5. A case-study with de-anonymizing a publicly released multimodal network of airports and airlines [20] that cannot be de-anonymized using state-of-the-art network alignment methods. We also study the set of most-helpful modes to find this alignment (Section 8).

We also review existing network alignment methods and ideas to contextualize our contribution (Section 3). We will post our codes at [www.cs.purdue.edu/~dgleich/codes/multinsd/](http://www.cs.purdue.edu/~dgleich/codes/multinsd/) for reproducibility and additional sensitivity experiments.

## 2 Multimodal networks & Multimodal network alignment

In the context of this paper, a multimodal network is a common set of vertex labels (the set  $V$ ) and multiple sets of undirected edges over these vertices. Each set of edges represents a mode, so we have  $E^{(1)}, \dots, E^{(m)}$  for  $m$  modes. We illustrate an example in Figure 1 with three modes. Note that only a subset of vertices may be present in each mode.

To manipulate multimodal network data computationally, we use a multimodal adjacency matrix. Our definition is in the spirit of how multislice, multiplex networks, and temporal networks are represented [18], although the details of our specific construction differ. Let  $\mathbf{A}_1, \dots, \mathbf{A}_m$  represent the adjacency matrices of each mode. Then the multimodal adjacency matrix is:

$$(2.1) \quad \mathbf{M} = \begin{bmatrix} \mathbf{A}_1 & \mathbf{C}_{12} & \cdots & \mathbf{C}_{1m} \\ \mathbf{C}_{12}^T & \mathbf{A}_2 & \cdots & \mathbf{C}_{2m} \\ \vdots & \ddots & \ddots & \vdots \\ \mathbf{C}_{1m}^T & \cdots & \cdots & \mathbf{A}_m \end{bmatrix}.$$

Here, the matrix  $\mathbf{C}_{ij}$  represents the cross-modal associations between mode  $i$  and mode  $j$ . This is always a binary diagonal matrix where  $\mathbf{C}_{ij}(k, k) = 1$  if vertex  $k$  is present in both mode  $i$  and mode  $j$ . All other entries are 0. An example is illustrated in Figure 1. Note that  $\mathbf{M}$  is sparse, and the cross-modal edges do not introduce too many new entries beyond the data.

The goal of the multimodal network alignment problem is to produce a matching between the vertices of two multimodal networks that *aligns* the networks by maximizing the number of edges of each network that are preserved under the matching. Recall that a matching is a 1 – 1 relationship between two sets. Let  $V_A$  and  $V_B$  be the vertex sets of multimodal networks  $A$  and  $B$ , respectively. We use  $v \leftrightarrow v'$  to denote a matching between  $v \in V_A$  and  $v' \in V_B$ . Let  $E_A^{(k)}$  and  $E_B^{(k)}$  be the edge sets of the  $k$ th mode of  $A$  and  $B$  as well. Then we seek a matching between  $V_A$  and  $V_B$  that maximizes

$$(2.2) \quad \sum_{k=1}^m \sum_{\substack{(u_A, v_A) \\ \in E_A^{(k)}}} \sum_{\substack{(u_B, v_B) \\ \in E_B^{(k)}}} \begin{cases} 1 & \text{if } u_A \leftrightarrow u_B \text{ and } v_A \leftrightarrow v_B \\ 0 & \text{otherwise.} \end{cases}$$

We call this objective the *overlap* of a matching. Note that this is a single matching between the vertices that is then evaluated over all the modes, which is possible because we assume that the correspondence between each mode is known and mode  $k$  in network  $A$  corresponds to mode  $k$  in network  $B$ . See Figure 2 for an example.

In many cases it is convenient to state the matching as a matrix  $\mathbf{X}$  where the rows are indexed by vertices

in  $V_A$  and columns by vertices in  $V_B$ . Let  $X(v, v') = 1$  if  $v \in V_A$  matches to  $v' \in V_B$  and 0 otherwise. Also, let  $\{\mathbf{A}_k\}$  and  $\{\mathbf{B}_k\}$  be the adjacency matrices for each mode of the multimodal networks  $A$  and  $B$ . Then the multimodal network alignment problem is

$$(2.3) \quad \begin{aligned} & \underset{\mathbf{X}}{\text{maximize}} && \frac{1}{2} \sum_{k=1}^m \sum_{ij} [\mathbf{X}^T \mathbf{A}_k \mathbf{X}]_{ij} [\mathbf{B}_k]_{ij} \\ & \text{subject to} && \sum_i X_{ij} \leq 1 \text{ for all } j \\ & && \sum_j X_{ij} \leq 1 \text{ for all } i \\ & && X_{ij} \in \{0, 1\}. \end{aligned}$$

This is an integer optimization problem over the space of matrices that is extremely challenging to solve exactly. It generalizes the subgraph isomorphism problem and all existing network alignment instances, which correspond to the case where  $m = 1$ .

## 3 Related work & Background

This section is meant to position our work within the broader context of network alignment methods. For recent surveys that explore these dimensions, see [6]. We are considering global network alignment which seeks a single matching that applies to the entire network instead of multiple local alignments between the networks. The global network alignment problem is almost always stated as an attempt to maximize the number of overlapped edges combined with maximizing domain-specific notions of apriori known vertex similarity. For instance, in protein-protein interaction networks, this notion of similarity is the similarity score of the respective genetic sequences [22]. In ontology alignment, this may come from a textual measure on the ontology label. The hardest network alignment problems lack these *hints* about how to align the networks. Our method falls into this hard case where we assume no apriori knowledge of the alignment. Another important feature is scoring the resulting match when the networks have substantially different sizes and edge densities. This aspect has been explored most thoroughly in terms of biological networks [24].

There are two prominent classes of methods: (i) embedding and (ii) integer optimization relaxations and heuristics. Embedding methods seek to compute a feature vector for each node of  $A$  and  $B$  independently. They then use relationships between the feature vectors to generate the matching. Using eigenvectors is common for this task in pattern recognition [12]. Recent methods have proposed graphlet counts [16] and eigenvector histograms [21]. One closely related method to our multimodal proposal is to generate the embeddings based on node and edge types [8]. Obtaining these feature vectors can be extremely expensive in terms of computation, but results in memory efficient methods.

Our method is most strongly related to the integer optimization and relaxation framework that seeks to maximize functions related to (2.3) in the case of one mode, perhaps with an additional term reflecting the vertex similarity. This problem is a specific instance of a quadratic assignment problem [3]. IsoRank was one of the first methods in this class [22]. Subsequent techniques include belief propagation methods [2], Lagrangian relaxations [11], spectral methods [7], and tensor eigenvectors for motif-alignment [17]. Essentially, all of these methods store a dense, real-valued heuristic matrix  $\mathbf{Y}$  of size  $|V_A| \times |V_B|$ . This needs memory that is *quadratic* in the size of the networks and greatly limits scalability. One of the most scalable methods results from a crucial analytical insight into the structure of the IsoRank heuristic. The resulting method – network similarity decomposition [17] – can be considered a hybrid of embedding and integer optimization. Our results show how these ideas enable a seamless generalization to multimodal networks and we return to explain in more depth in the next section.

#### 4 Multimodal Similarity Decomposition: A multi-modal generalization of IsoRank and the Network Similarity Decomposition

We now present our approach to compute a multimodal network alignment. The high level idea is that we are going to run the IsoRank method to align the multimodal adjacency matrices  $\mathbf{M}$  and  $\mathbf{N}$  corresponding to multimodal networks  $A$  and  $B$ . Doing so will involve a number of new insights about how the methods will behave on multimodal networks and exploiting the structure of the methods. The result of this section is a heuristic solution matrix  $\mathbf{Y}_{MN}$  that we describe how to turn into a multimodal alignment in Section 5. We begin by reviewing IsoRank [22] and the network similarity decomposition [13]. We assume the networks are undirected in this derivations, but note that directed generalizations are possible.

**4.1 IsoRank & the network similarity decomposition** Consider aligning two standard networks  $A$  and  $B$  with adjacency matrices  $\mathbf{A}$  and  $\mathbf{B}$ . The IsoRank method generates a heuristic matrix  $\mathbf{Y}$  where  $Y(v, v')$  is large if  $v$  seems like it should match with  $v'$ . To do this, IsoRank computes  $\mathbf{Y}$  as the solution of a PageRank problem on the product graph of  $A$  and  $B$ . In a small surprise, this is a well-motivated idea. For more details about this, we refer the reader to [2]. We can express this PageRank problem in terms of the final matrix  $\mathbf{Y}$  where  $\mathbf{P}$  is the degree-normalized matrix for network  $A$ ,  $P_{ij} = A_{ij}/d_j$ ,  $\mathbf{Q}$  is the degree-normalized matrix for  $B$ , and  $S(v, v')$  is the apriori similarity of node  $v$  in  $A$  and

$v'$  in  $B$ :

$$(4.4) \quad \mathbf{Y} = \alpha \mathbf{P} \mathbf{Y} \mathbf{Q}^T + (1 - \alpha) \mathbf{S}$$

One small note: we assume in these equations that  $A$  and  $B$  have no nodes with zero outdegree. If there are such nodes, then we need to normalize  $\mathbf{Y}$  to be a probability distribution. When there is no apriori similarity known, then all entries  $\mathbf{S}$  are the constant  $|V_A|^{-1}|V_B|^{-1}$ , which reflects uniform similarity.

The insight in the network similarity decomposition (NSD) is that if  $\mathbf{S}$  is rank 1 – as it would be in the case of uniform similarity – then running  $t$  iterations of the power method to compute PageRank will result in a rank  $t + 1$  matrix  $\mathbf{Y}$ . Moreover, the low-rank factorization of  $\mathbf{Y}$  can be easily computed by running the power method for PageRank on network  $A$  and network  $B$  *independently*. (In this sense, this is similar to an embedding method, because we compute the PageRank iterates of each network separately, then combine them.) This is easy to see when  $\mathbf{S} = \mathbf{u}\mathbf{v}^T$  then  $t$  iterations of the power-method for (4.4) produce:

$$\mathbf{Y}^{(t)} = (1 - \alpha) \sum_{k=0}^{t-1} \alpha^k \mathbf{P}^k \mathbf{u}\mathbf{v}^T (\mathbf{Q}^T)^k + \alpha^t \mathbf{P}^t \mathbf{u}\mathbf{v}^T (\mathbf{Q}^T)^t$$

(This matrix is a sum of  $t + 1$  rank-1 matrices of the form  $\text{scalar} \cdot (\mathbf{P}^t \mathbf{u})(\mathbf{Q}^t \mathbf{v})^T$ .) In this way, a low-rank factorization of the IsoRank solution can be computed extremely efficiently. Run  $t$  steps of the power method for PageRank on network  $A$  and  $B$  separately and store the iterations (we will be more precise when we explain our method for multimodal networks next). Reasonable values of  $t$  are between 5 and 25 as the PageRank series converges extremely quickly; for  $\alpha$  we use  $\alpha = 0.9$ .

**4.2 Our Multimodal Similarity Decomposition** Our goals with the multimodal similarity decomposition and multimodal alignment closely mirror those of NSD and IsoRank. In our case, we want to allow alignment information to *flow* between modes of the network to reinforce alignments between vertices that are present in various modes. This intuition suggests that the IsoRank scores for the alignment of multimodal adjacency matrices will be an effective heuristic for our problem. This will flow alignment information between modes via the cross-modal edges  $\mathbf{C}_{ij}$ .

Since we assume that the alignment between modes of the network are already known, we can strengthen our heuristic by customizing the matrix  $\mathbf{S}$ . Recall that  $S_{ij}$  is the apriori similarity of node  $i$  in  $A$  and  $j$  in  $B$ . The multimodal adjacency matrix has  $m$  copies of each vertex identifier – one for each vertex in each of the  $m$

modes. Thus, we use the matrix

$$(4.5) \quad \mathbf{S} = \begin{bmatrix} \gamma \cdot \mathbf{ones} & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & \gamma \cdot \mathbf{ones} \end{bmatrix},$$

where  $\mathbf{ones}$  is a matrix of size  $|V_A| \times |V_B|$  and  $\gamma$  is chosen such that all entries of  $\mathbf{S}$  sum to 1 (that is,  $\gamma = m^{-1}|V_A|^{-1}|V_B|^{-1}$ ). This choice of  $\mathbf{S}$  corresponds to a rank- $m$  matrix, where each mode constitutes a single rank-1 factor suggesting that nodes in mode 1 in multimodal network  $A$  should correspond to nodes in mode 1 of multimodal network  $B$ .

Let  $\mathbf{M}$  and  $\mathbf{N}$  be the multimodal adjacency matrices for multimodal networks  $A$  and  $B$ , respectively. Let  $\mathbf{P}_M$  and  $\mathbf{Q}_N$  correspond to degree-normalization (by columns) of  $\mathbf{M}$  and  $\mathbf{N}$ . Thus, the heuristic  $\mathbf{Y}_{MN}$  we compute solves:

$$(4.6) \quad \mathbf{Y}_{MN} = \alpha \mathbf{P}_M \mathbf{Y}_{MN} \mathbf{Q}_N^T + (1 - \alpha) \mathbf{S}$$

where  $\mathbf{S}$  is the multimodal similarity (4.5). The result  $\mathbf{Y}_{MN}$  gives a score for aligning each node in each mode of  $A$  to each node in each mode  $B$ . This is not the type of output we want, so in subsequent sections, we discuss how to turn this output into the type of matching we expect.

At this point, we can utilize the same methodology underlying the network similarity decomposition to efficiently compute a low-rank decomposition of  $\mathbf{Y}_{MN}$ . Note that the matrix  $\mathbf{S}$  is rank- $m$ . Thus, we can compute  $t$  iterations of the power-method to solve (4.6) and the result is a rank  $(m) \cdot (t + 1)$  matrix  $\mathbf{Y}_{MN}$  following the exact same argument as expression for the network similarity decomposition. The pseudocode is shown in Figure 3. One small detail is that we assumed in the derivation that  $\mathbf{P}_M$  and  $\mathbf{Q}_N$  have no empty rows or columns. If this is not the case, then each iterate needs to be normalized to be a probability distribution as in the pseudocode.

## 5 Matching algorithms & Resolving conflicts

After the multimodal network decomposition (MSD), the next step is to produce a matching between the vertices. A standard technique here is to run a max-weight bipartite matching on the matrix  $\mathbf{Y}$  for the network alignment case [22, 2]. There are two issues with this step. The first is the size of the resulting matrix  $\mathbf{Y}_{MN}$  can be prohibitively large if we realize it as a dense matrix. This is especially true if we have a large multimodal network with many modes. A problem with 100 modes and 5000 vertices would produce a  $500,000 \times 500,000$  dense matrix whereas the low-rank factors for 8 iterations would require two  $500,000 \times 900$  matrices. Thus, we

Input: Multimodal networks  $A, B, \alpha$ , iteration  $t$   
 Output: Matrices  $\mathbf{U}$  and  $\mathbf{V}$  such that  $\mathbf{Y}_{MN}^{(t)} = \mathbf{U}\mathbf{V}^T$   
 $\mathbf{U} = \text{PageRankPowers}(A, \alpha, t)$   
 $\mathbf{V} = \text{PageRankPowers}(B, \alpha, t)$

PageRankPowers( $A, \alpha, t$ )

Set  $\mathbf{M}$  to be the multimodal adjacency of  $A$ .

Set  $\mathbf{P}$  to be the column normalized matrix  $\mathbf{M}$ .

Allocate  $\mathbf{Z}$  as a  $|V| \times (t + 1) \times m$  array

for  $k=1$  to  $m$  (the number of modes)

for all  $i=1$  to  $|V|m$

$\mathbf{Z}[i, 0, k] = \begin{cases} (\sqrt{m}|V|)^{-1} & i \text{ is an index in mode } k \\ 0 & \text{otherwise} \end{cases}$

for  $j=1:t$

$\mathbf{Z}[:, j, k] = \text{normalize}(\mathbf{P} \mathbf{Z}[:, j-1, k])$

$\mathbf{Z}[:, j-1, k] = \sqrt{(1-\alpha)\alpha^{j-1}} \mathbf{Z}[:, j-1, k]$

end

$\mathbf{Z}[:, t, k] = \sqrt{\alpha^t} \mathbf{Z}[:, t, k]$

end

return  $\mathbf{U}$  as  $\mathbf{Z}$  reshaped to  $|V| \times (t + 1)(m)$

Figure 3: Pseudocode for the multimodal similarity decomposition:  $m$  is the number of modes of the multimodal networks  $A$  and  $B$ , and the column-normalization computes:  $P_{ij} = M_{ij} / \sum_{\ell} M_{\ell,j}$ . The `normalize` function divides a vector by its sum. The reshaping occurs by column such that  $U(:, 1) = Z(:, 0, 1)$ . We assume the reshaping happens the same way on  $\mathbf{U}$  and  $\mathbf{V}$  to ensure the low-rank factors are aligned. The result  $\mathbf{U}$  and  $\mathbf{V}$  must still be matched as in Section 5 to solve (2.3).

investigate memory-efficient matching routines that deal with  $\mathbf{Y}_{MN}$  via its low-rank decomposition. The second issue is that this output is a matching on the rows of the multimodal adjacency matrix and not a matching on the vertex set of the multimodal network. Aggregating them can result in conflicting matches and we describe a few ways to resolve these.

### 5.1 Approximate matching via low-rank factors

Bipartite matching is a common subroutine that has a well-known polynomial time algorithm. The problem with this algorithm is that it requires the matrix of weights associated with each edge of the bipartite graph. For our case, this matrix is large and dense, rendering the algorithm infeasible due to the memory required. However, note that the bottleneck is the density of the matrix. Bipartite matching is computationally tractable with *sparse* matrices of exactly the same size because the algorithms utilize the sparsity to lower the memory requirement. For the purposes of this section, let  $\mathbf{Y} = \mathbf{U}\mathbf{V}^T$  be the matrix we wish to run a bipartite matching on and  $\mathbf{U}$  be  $m \times k$ . We note that the low-rank factors from MSD are non-negative, so we assume  $\mathbf{U}, \mathbf{V} \geq 0$ .

We begin with an extremely idealized case. If  $k = 1$ ,

then  $\mathbf{Y} = \mathbf{u}\mathbf{v}^T$  is rank-1, and it is extremely easy to compute a maximum weight matching. This only involves *sorting* the vectors  $\mathbf{u}$  and  $\mathbf{v}$  in decreasing order and taking the largest feasible set of components. (This fact is a simple corollary of the rearrangement inequality, so we omit a formal statement and proof.)

For the remainder of this derivation, we'll need one additional bit of notation: the matrix inner product  $\mathbf{A} \bullet \mathbf{B} = \sum_{ij} A_{ij}B_{ij}$ . Recall that a matching can be expressed as a matrix where  $X_{ij} = 1$  if  $i$  is matched to  $j$  and all other entries are 0. Given any matching  $\mathbf{X}$  the weight of the matching is  $\mathbf{X} \bullet \mathbf{Y}$ . We now show that using the *best matching* from each rank-1 factor of  $\mathbf{Y}$  gives a  $1/k$ -approximation when  $\mathbf{Y}$  has rank- $k$ .

**THEOREM 5.1.** *Let  $\mathbf{Y} = \mathbf{U}\mathbf{V}^T$  be a rank- $k$  non-negative decomposition for  $\mathbf{Y}$ . Let  $\mathbf{X}_i$  be the maximum weight matching corresponding to the rank-1 factor  $\mathbf{u}_i\mathbf{v}_i^T$ , and let  $f_i = \mathbf{X}_i \bullet (\mathbf{u}_i\mathbf{v}_i^T)$  be its weight. If  $f^*$  is the largest value of  $f_i$ , then  $f^*$  is a  $1/k$ -approximation to the maximum weight matching on  $\mathbf{Y}$ .*

*Proof.* Let  $\mathbf{X}^*$  be any optimal maximum weight matching. Then

$$\underbrace{\mathbf{X}^* \bullet \mathbf{Y}}_{\text{optimal matching weight}} = \underbrace{\sum_{i=1}^k (\mathbf{X}^* \bullet \mathbf{u}_i\mathbf{v}_i^T)}_{f_i \text{ is optimal for } \mathbf{u}_i\mathbf{v}_i^T} \leq \sum_{i=1}^k f_i \leq kf^*.$$

While this provides a simple approximation bound, the value of  $k$  for our experiments is often around 1000, rendering it rather pointless in terms of theory. However, we find that a single factor often produces exceptionally good approximation factors – far beyond what the theory from this result would show. We plan to continue investigating this problem and conjecture that the structure of the MSD vectors themselves may suggest a tighter approximation.

In practice, there are a few additional improvements we can make. Two of these apply to general problems and also produce a  $1/k$  approximation. The third improvement is specific to the multimodal network alignment objective. We also discuss a few other ideas suggested in the literature.

**Maximum weight  $1/k$ .** In the first variant, we compute the weight of the matching  $\mathbf{X}_i$  in the full matrix  $\mathbf{Y}$  via  $\mathbf{Y} \bullet \mathbf{X}_i$ . The runtime of this computation is  $nk$  and it can be done entirely with the low-rank factors. Then we select the largest weight among matchings  $\mathbf{X}_i$ . This is still a  $1/k$  approximation because  $\mathbf{Y} \bullet \mathbf{X}_i \geq (\mathbf{u}_i\mathbf{v}_i^T) \bullet \mathbf{X}_i$  due to the non-negativity.

**Union of matchings.** While we cannot always expect to solve the dense bipartite matching problem, we are able to solve sparse bipartite matching problems on the size of the vertex sets easily. In this case, we

Table 1: We compare runtimes of approximate bipartite matching where  $\mathbf{U}$  and  $\mathbf{V}$  are  $n \times k$  for simplicity; all methods use  $O(nk)$  memory. Also,  $\text{BM}(E)$  is the runtime for bipartite matching on  $n$  nodes with  $E$  total entries, and  $|E|$  is the number of edges in the original networks.

Method	Runtime
Simple $1/k$	$nk \log n$
Max Weight $1/k$	$nk \log n + nk^2$
Union $1/k$	$nk \log n + nk^2 + \text{BM}(nk)$
Max Overlap	$nk \log n +  E $
$r$ -Sparsified	$n^2k + \text{BM}(nr)$
Greedy $1/2$	$n^3k$

can take the union of edges produced by the matchings  $\mathbf{X}_1, \dots, \mathbf{X}_k$  and find the best bipartite matching in this sparse subset of the entries of  $\mathbf{Y}$ . Again, this can only make our approximation better and so we still get  $1/k$ .

**Maximum overlap.** The actual goal of our problem is to achieve the best alignment between two multimodal networks. Thus, rather than compute the *weight* of the matching in the heuristic  $\mathbf{Y}$ , we directly evaluate the *overlap* produced by each of the matchings  $\mathbf{X}_i$  between the multimodal adjacency matrices. (Note this is slightly different from the matching we will consider after we resolve conflicts below.) This picks the matching  $\mathbf{X}_i$  that corresponds with the true objective functions.

**Other possibilities.** In ref. [14], they suggest computing  $\mathbf{Y}$  a row at a time and retaining the largest  $r$  entries as edges of a graph. This forms a sparse graph that can be used in a parallel maximum weight bipartite matching. Alternatively, it is feasible to run a greedy  $1/2$ -approximate matching. We found both of these underperformed our ideas in terms of final quality and were no faster.

**Runtime & Memory** All of these schemes keep the matrix  $\mathbf{Y}$  stored as factors  $\mathbf{U}$  and  $\mathbf{V}$ . Each matching takes  $O(n)$  storage, so storing  $k$  matches requires  $O(nk)$  memory, and this exactly mirrors the storage of the factors  $\mathbf{U}$  and  $\mathbf{V}$  themselves. We provide a comparison of runtime in Table 1.

**Practical considerations.** All of the approximations here are extremely fast to compute for values of  $n$  up to a few million. So in our codes, we will usually compute multiple matchings and choose the best based on the overall alignment scores after we resolve conflicts as discussed next.

**5.2 Resolving conflicts** So far, we have presented methods to derive a matching on a low-rank approximation. In the context of multimodal networks, this matches at the level of the row and column identifiers of the multimodal adjacency matrices. This is a problem because each vertex id in the multimodal network

occurs at  $m$  different rows of the multimodal adjacency matrix. For instance, if multimodal network  $A$  consists of 3 modes, vertex 1 can be matched to three different vertices, one for each mode!

To resolve these conflicts, we have two methods and choose the best among them based on the overlap of the alignment they produce. The first method is to greedily examine the multimodal matching, project each match to nodes in  $V_A$  and  $V_B$ , and accept it if it's still feasible in the projected alignment. The second method is to project the multimodal matchings to one bipartite network on  $V_A$  and  $V_B$ . Put another way, we take the union of all matches induced by the single multimodal matching. Each edge is weighted by the value in  $\mathbf{Y}_{MN}$  (and duplicates are summed) that arose from the matching. Once we project all the matchings, we use a bipartite matching to obtain the final matching.

## 6 Baseline methods

Before we describe our experiments, we explain how the multimodal alignment problem can be addressed using existing methods for network alignment. These methods do not take advantage of the multimodality of the data to guide the alignment, but nevertheless provide strong baselines. The first baseline is to smash the networks together and look at an alignment of the unimodal networks with edge set  $\cup_{k=1}^m E_A^{(k)}$  and  $\cup_{k=1}^m E_B^{(k)}$ . We can use any existing network alignment tool for this task. In the following experiments, we use Klau's relaxation [11] because it had the highest performance among a number of methods we evaluated (see online codes). In addition to aligning the smashed unimodal networks, we can also compute a pairwise alignment between each mode. Thus, we align the edges  $E_A^{(k)}$  to  $E_B^{(k)}$ , which provides another matching. Note that each alignment produces exactly the type of output we need and there is no need to resolve conflicts as we saw in the multimodal case. We can then pick the *best* among these  $m + 1$  alignments (1 alignment for each mode and 1 alignment for the smashed network) based on the total overlap they induce in the full multimodal network. We call these methods *pairwise alignment* rather than *multimodal alignment*.

## 7 Synthetic experiments

In our first experiment, we want to understand when multimodal alignment results in better alignment compared with using a strong algorithm for the pairwise case. We use our multimodal method with  $\alpha = 0.9$ , 10 iterations, and then compute the dense matrix  $\mathbf{Y}_{MN}$  to use with bipartite matching because this test case is sufficiently small. We compare it against Klau's method [11] as a pairwise aligner as in Section 6. Klau's algorithm takes




Figure 5: At left, we show how the recovery rate changes as we add modes when  $p = 0.1$  and  $q = 0.2$ ; at right we see  $p = 0.2$  and  $q = 0.1$ . The performance of our multimodal alignment increases as we add modes. The line is the mean over 50 trials and the bands show the 10% and 90% percentiles.

significantly longer than MSD.

We generate a multimodal network alignment problem with 36 nodes and 6 modes as follows. We first create a 12-node Erdős-Rényi graph with average degree 3, and then join copies into a single 36-node network. The purpose of combining the networks is to put a degree of symmetry into the graph to make the alignment harder. (For more details about the networks construction refer to the code associated with this paper). This is our reference graph. Next, we generate a mode by randomly deleting vertices with probability  $p$  and randomly deleting edges with probability  $q/2$ . Then we generate two instances of this modal graph with deleted edges again with probability  $q/2$ . One instance goes to multimodal network  $A$  and the other goes to multimodal network  $B$ . At the end, we have two multimodal networks  $A$  and  $B$  where each mode shares a number of relationships.

In figure 4, we show the fraction of edges aligned using a method over the total number of edges in the networks as we vary  $p$  and  $q$  for both MSD and the best of pairwise alignments (Section 6). We do not use node-based recovery due to many nodes of degree one, which cannot be resolved. The results are the mean over 50 trials, and they show a large regime where multimodal alignment is superior to any pairwise alignment. Specifically, when the vertex deletion probability  $p$  is large, the multimodal alignment is the only method to accurately align the networks. When edge deletion is high, smashing the networks effectively reconstitutes the original network and we see the superior performance of Klau's pairwise method.

Next, in Figure 5, we show how the behavior of the methods change as we vary the number of modes of the multimodal network. As expected, the performance of our method increases with additional modes, whereas the performance of the pairwise method is consistent (or even slightly decreasing).

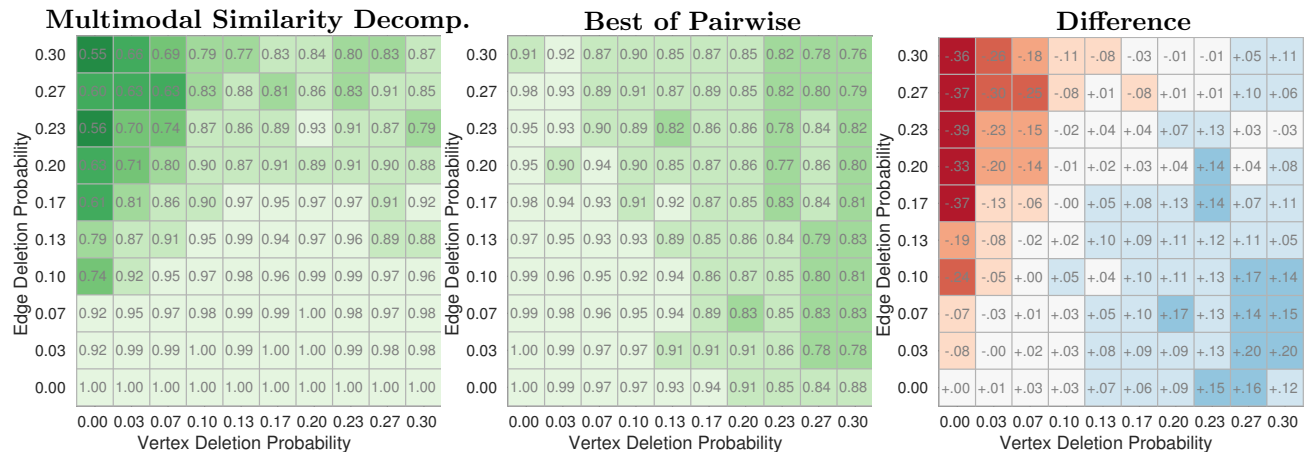


Figure 4: At left, the recovery results for our multimodal similarity decomposition (MSD) show excellent performance even in the high vertex deletion regime. In the middle, the best of any pairwise recovery result shows good results in the high edge deletion, but worse results with high vertex deletion. At right, the difference shows blue when MSD outperforms the best of the pairwise alignments by at least 5% recovery.

## 8 De-anonymizing transportation datasets

In this section we present a case-study with de-anonymizing a publicly released dataset of airlines and airports, where we treat each airline as a mode. We show in this section that only our proposed multimodal alignment method can fully de-anonymize the network and align each edge. We run these experiments with  $\alpha = 0.9, 10$  and illustrate how various low-rank matching procedures perform. The anonymized network is the European air network from ref. [20], which was released with anonymized airport identifiers but with all airline identifiers. The data originally came from the OpenFlights repository (<http://openflights.org>). We wanted to take the original data and use it to restore identifiers to the anonymized data. We selected the May 2013 release based on the publication dates. These networks have 594 airports and 175 airlines and 6468 edges, respectively. Our goal is to align the multimodal network such that each edge is matched.

**8.1 Performance at de-anonymization** We consider an alignment to de-anonymize the network if it is able to overlap all 6468 edges through the matched vertex identifiers. In Table 2, we show the overlap size for our multimodal methods and a number of different low-rank matching techniques from Section 5. This time we are not able to run the full bipartite matching unless we use a computer with 512GB of RAM and so we don't report those results. The best result from the baseline pairwise alignment methods miss 70 edges. All of our low-rank matching approximations achieve the full overlap score.

## 8.2 Which modes help multimodal alignment?

The final experiment seeks to understand *which* modes have the highest impact on the de-anonymization performance. In many applications, collecting additional data incurs a cost and this experiment is designed to provide insight as far as what type of multimodal data would be most helpful to collect. Our goal is to use only a subset of the 175 modes to produce an alignment, and then compute the overlap that results from using this alignment on all 175 modes. We sort the modes based on a number of graph theoretic measures to produce interesting subsets. The per-mode measures are: edge count, unique vertex count, average degree, triangle count, density. Figure 6 shows the results for these measures along with a few random orderings of the modes. The goal is to align the highest fraction of edges using the fewest modes. Most of the graph theoretic measures have similar performance, which suggests that any could be a proxy for which data to collect. Density is a notable outlier as some of the modes consist of a single edge, resulting in a high density, but little information about the global alignment.

## 9 Conclusion & Future Work

This paper demonstrates an advantage to using multimodal features of data: it makes alignment problems easier. Here, we compared our multimodal network decomposition (MSD) against a carefully engineered method in pairwise alignment (Klau's). We also did all of the pairwise experiments reported here with the IsoRank method instead of Klau's and the results were uniformly worse as far as the alignment quality. However, when we use the multimodal extension to IsoRank we have proposed, accurate alignments are easy to ob-



Table 2: The multimodal airport networks have 6468 edges. All the multimodal methods are able to de-anonymize this network. The best pairwise results come from smashing the multimodal structure and aligning the resulting networks, which misses 70 edges compared to the multimodal alignments and takes more time.

Type	Method	Overlap
Multimodal	MSD & Max Weight $1/k$	6468
Multimodal	MSD & Union $1/k$	6468
Multimodal	MSD & Max Overlap	6468
Pairwise	smashed	6398
Pairwise	best mode	3127

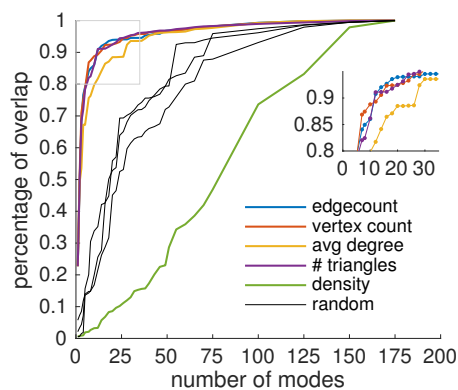


Figure 6: As we use more modes to produce an alignment for the airport data, the performance increases. This experiment shows that using modes that have many edges (edge count), touch many unique vertices (vertex count), have a high average degree (avg degree), or have high triangle count (# triangles) all outperform random selections. Alternatively, density does not work as discussed in the main text.

tain. For instance, note that in Figure 5, MSD used the additional modes to improve the alignment, whereas the pairwise method did not.

Although we have not captured runtimes precisely in this manuscript, let us note that Klau's method takes a great deal more computational effort than our approximation, and results in an extremely high quality pairwise solution for any mode or the smashed networks. To give some sense of the difference in effort required, the results on the synthetic experiment for Figure 4 took *hours* to compute with Klau's method whereas MSD and all the matching took a few minutes.

In future work, we plan to explore using this type of multimodal alignment in the context of protein and gene relationships. We believe that using this new methodology – combined with particular domain specific adaptations – will result in new biological insights.

## References

[1] F. Battiston, V. Nicosia, and V. Latora. Structural measures for multiplex networks. *Phys. Rev. E*, 89:032804, 2014.

[2] M. Bayati, D. F. Gleich, A. Saberi, and Y. Wang. Message-passing algorithms for sparse network alignment. *ACM Trans. Knowl. Discov. Data*, 7(1):3:1–3:31, 2013.

[3] R. Burkard, M. Dell'Amico, and S. Martello. *Assignment Problems*. SIAM, revised edition, 2012.

[4] A. Cardillo, J. Gómez-Gardeñes, M. Zanin, M. Romance, D. Papo, F. d. Pozo, and S. Boccaletti. Emergence of network features from multiplexity. *Sci. Rep.*, 3:1344 EP, 02 2013.

[5] D. Conte, P. P. Foggia, C. Sansone, and M. Vento. Thirty years of graph matching in pattern recognition. *Int. J. Pattern Recogn.*, 18(3):265–298, May 2004.

[6] A. Elmsallati, C. Clark, and J. Kalita. Global alignment of protein-protein interaction networks: A survey. *IEEE/ACM TCCB*, 13(4):689–705, July 2016.

[7] S. Feizi, G. Quon, M. Recamonde-Mendoza, M. Medard, M. Kellis, and A. Jadbabaie. Spectral alignment of networks. *arXiv*, cs.DS:1602.04181, 2016.

[8] C. Fraikin and P. van Dooren. Graph matching with type constraints on nodes and edges. Number 07071 in Dagstuhl Proceedings, 2007.

[9] W. Hu, Y. Qu, and G. Cheng. Matching large ontologies: A divide-and-conquer approach. *Data Knowl. Eng.*, 67(1):140–160, October 2008.

[10] B. P. Kelley, B. Yuan, F. Lewitter, R. Sharan, B. R. Stockwell, and T. Ideker. PathBLAST: a tool for alignment of protein interaction networks. *Nucl. Acids Res.*, 32:W83–88, 2004.

[11] G. Klau. A new graph-based method for pairwise global network alignment. *BMC Bioinform.*, 10(Suppl 1):S59, 2009.

[12] D. Knossow, A. Sharma, D. Mateus, and R. Horaud. Inexact matching of large and sparse graphs using laplacian eigenvectors. Number 5534 in LNCS, pp. 144–153. 2009.

[13] G. Kollias, S. Mohammadi, and A. Grama. Network similarity decomposition (NSD): A fast and scalable approach to network alignment. *IEEE TKDE*, PP(99):1, 2011.

[14] G. Kollias, M. Sathe, O. Schenk, and A. Grama. Fast parallel algorithms for graph similarity and matching. *J. Parallel Dist. Comp.*, 74(5):2400 – 2410, 2014.

[15] N. Korula and S. Lattanzi. An efficient reconciliation algorithm for social networks. *Proc. VLDB Endow.*, 7(5):377–388, Jan. 2014.

[16] O. Kuchaiev, T. Milenković, V. Memišević, W. Hayes, and N. Pržulj. Topological network alignment uncovers biological function and phylogeny. *J. R. Soc. Interface*, 2010.

[17] S. Mohammadi, D. F. Gleich, T. G. Kolda, and A. Grama. Triangular alignment (TAME): A tensor-based approach for higher-order network alignment. *TCCB*, Online:1–14, 2016.

[18] P. J. Mucha, T. Richardson, K. Macon, M. A. Porter, and J.-P. Onnela. Community structure in time-dependent, multiscale, and multiplex networks. *Science*, 328(5980):876–878, 2010.

[19] J. Ni, H. Tong, W. Fan, and X. Zhang. Inside the atoms: Ranking on a network of networks. In *KDD*, pp. 1356–1365, 2014.

[20] V. Nicosia and V. Latora. Measuring and modeling correlations in multiplex networks. *Phys. Rev. E*, 92:032805, Sep 2015.

[21] R. Patro and C. Kingsford. Global network alignment using multiscale spectral signatures. *Bioinformatics*, 28(23):3105–3114, 2012.

[22] R. Singh, J. Xu, and B. Berger. Global alignment of multiple protein interaction networks with application to functional orthology detection. *PNAS*, 105(35):12763–12768, 2008.

[23] M. Szell, R. Lambiotte, and S. Thurner. Multirelational organization of large-scale social networks in an online world. *PNAS*, 107(31):13636–13641, 2010.

[24] V. Vijayan, V. Saraph, and T. Milenković. Magna++: Maximizing accuracy in global network alignment via both node and edge conservation. *Bioinformatics*, 31(14):2409–2411, 2015.