Published in final edited form as:

IEEE Trans Pattern Anal Mach Intell. 2020 November ; 42(11): 2887–2900. doi:10.1109/TPAMI.2019.2914651.

Matched Filters for Noisy Induced Subgraph Detection

Daniel L. Sussman,

Department of Mathematics & Statistics, Boston University, Boston, MA

Vince Lyzinski,

Department of Mathematics, University of Massachusetts, Amherst, MA

Youngser Park,

Department of Applied Math & Statistics, The Johns Hopkins University, Baltimore, MD

Carey E. Priebe [Senior Member, IEEE]

Department of Applied Math & Statistics, The Johns Hopkins University, Baltimore, MD

Abstract

The problem of finding the vertex correspondence between two noisy graphs with different number of vertices where the smaller graph is still large has many applications in social networks, neuroscience, and computer vision. We propose a solution to this problem via a graph matching matched filter: centering and padding the smaller adjacency matrix and applying graph matching methods to align it to the larger network. The centering and padding schemes can be incorporated into any algorithm that matches using adjacency matrices. Under a statistical model for correlated pairs of graphs, which yields a noisy copy of the small graph within the larger graph, the resulting optimization problem can be guaranteed to recover the true vertex correspondence between the networks. However, there are currently no efficient algorithms for solving this problem. To illustrate the possibilities and challenges of such problems, we use an algorithm that can exploit a partially known correspondence and show via varied simulations and applications to *Drosophila* and human connectomes that this approach can achieve good performance.

Keywords

multiple graph inference; subgraph detection; graph matching

1 INTRODUCTION

In many settings, we often want to quantify how multiple networks relate to each other in order to study how actors jointly use these networks. This may arise from multiple modalities, such as communications networks, delivery networks, financial networks, and social networks, or from a time dynamic setting. Similarly, in neuroscience or biology we may seek to compare brain networks or protein networks of different individuals or species. In computer vision, we may seek to determine the relationships between the shapes and objects in two images as represented by graphs. Often, these networks are on different

unmatched sets of vertices that are not the same cardinality. This limits the set of available tools as the adjacency matrices and other graph statistics may not be directly comparable.

In a related fashion, detection and location of particular induced subgraphs that correspond to a given activity or structure of interest are also paramount problems in various domains. Examples include determining whether a particular activity pattern is present in a social network, detecting certain shapes in an image, or discovering motifs in brain networks.

When the induced subgraphs are not too large, with up to ≈ 10 – 20 vertices, there are a number of approaches including tree search approaches [1] and backtracking algorithms [2]. Parallel algorithms have also been proposed to approximately enumerate the number of isomorphic copies of a subgraph appearing in a large network [3]. However, it may be the case that these subgraphs contain $\gg 20$ vertices. In this case, many existing approaches will either fail to find the subgraphs of interest or be computationally intractable. While state space representations [4] have been used to identify these larger subgraphs, these also have limitations in terms of the applicable graph structure. As a simple example of the challenge of this problem, consider finding or enumerating cliques. Finding cliques and estimating the number of cliques of small sizes can be achieved via fast algorithms [5], but the problem of finding the maximum clique is a notoriously challenging problem [6].

Furthermore, while even the subgraph isomorphism is itself NP-complete [7], we often expect the subgraphs may not appear exactly in the graph but rather only approximately. This may be due to errors and missingness in the observation of one or both graphs, so finding an exact subgraph might not be possible.

We view both of these as different points on a continuum of related problems that are all examples of graph matching [8], [9]. In a simpler setting, when the two graphs adjacency matrices $A, B \in \{0, 1\}^{n \times n}$ are of equal size, the graph matching problem (GMP) seeks the permutation matrix

$$\operatorname{argmin}_{P \in \mathcal{P}} \|A - PBP^T\| = \operatorname{argmax}_{P \in \mathcal{P}} \operatorname{trace}(APBP^T), \quad (1)$$

where $\mathcal{P}$ is the set of $n \times n$ permutation matrices. While this problem is NP-hard in general, there are numerous approaches in the literature that attempt to approximately solve the GMP. [8] and [10] provide reviews of the prescient literature. In particular, when prior knowledge about the correspondence between vertices the can be incorporated into the algorithms, the GMP can be approximately solved efficiently for graphs with more than 10^5 vertices [11], [12] without the need for sophisticated modern parallel computing to be brought to bear.

A more challenging problem that we will not consider is to detect anomalous subgraphs within a collection of graphs [13]. In the anomaly detection setting, the structure of the anomalous subgraph may be only known up to certain graph characteristics or deviations from the structure of the remaining graph. In our setting, the approximate topology of the graph to be detected is known explicitly.

To approximately solve the noisy induced subgraph problem, we will use the machinery of graph matching to construct a *graph matching matched filter*, aligning the smaller network to its closest subgraph within the larger network. Our main goals in this manuscript are to investigate the theoretical limitations of this noisy graph matching framework when the graphs may be of very different sizes. To match graphs of radically different sizes in Eq. (1), we consider padding the smaller matrix to render graph matching an appropriate tool for the problem [14]. Under a statistical model for noisily implanting an induced subgraph into a larger network, we show that the true induced subgraph will be found by brute force GM algorithm using an appropriate padding scheme, provided graph correlations and probabilities satisfy certain parameter assumptions.

We further demonstrate the effectiveness of these strategies when the vertex correspondence is partially known. By relaxing the graph matching problem constraints to the set of doubly stochastic matrices, a gradient descent approach can be put to use [15]. The problem is non-convex [16] and hence the initialization can incorporate the partial correspondence and random restarts to seek multiple local minimums. By choosing the optimal of these local minima, an exact or approximate recovery of the vertex correspondence can often be found.

2 BACKGROUND FOR GRAPH MATCHING

In this section we provide a brief background on graph matching, some methods to incorporate prior information, and a statistical model for correlated graphs. Throughout the remainder of this article we will use the following notation.

Let $[n] = \{1, 2, \dots, n\}$. Let $\mathcal{P}_n$ and $\mathcal{D}_n$ denote the set of $n \times n$ permutation matrices and doubly stochastic matrices, respectively. Let $\mathbf{J}_n$ and $\mathbf{0}_n$ denote the $n \times n$ all-ones and all-zeros matrices, respectively, and $\mathbf{J}_{nm}$ and $\mathbf{0}_{nm}$ for $n \times m$ rectangular versions. Let $\mathcal{A}_n$ denote the set of adjacency matrices corresponding to simple undirected graphs, i.e. the set of symmetric hollow $\{0, 1\}$ -valued $n \times n$ matrices. When clear from context, we may omit subscripts indicating the sizes of these matrices. Finally, let $\oplus$ denote the direct sum of two matrices,

$$M \oplus M' = \begin{pmatrix} M & \mathbf{0} \\ \mathbf{0} & M' \end{pmatrix}.$$

2.1 Contrasts with Subgraph Isomorphism

Before describing our noisy matched filter and related procedures, we would like to distinguish our approach and goals from that of exact subgraph isomorphism. The problem of exact subgraph isomorphism is to find an induced subgraph of the larger graph which is isomorphic to the smaller graph. Frequently, the goal is to find all isomorphic copies of the smaller graph within the larger graph. Examples of this could be counting triangle, k -cliques, and k -cycles, some of which have fast algorithms associated with them while in general the problem is NP-complete [7].

The subgraph isomorphism problem is an area of active research [17], especially in areas such as computer vision and pattern recognition. Recent algorithm developments include the VF3 algorithm [1], which builds on the VF2 algorithm [18], and both employ a tree-search

approach. Other popular and recent algorithms include the LAD algorithm [19] and the RI algorithm [20].

The problem of noisy matched filters or inexact subgraph isomorphism differs from exact isomorphism in that only the “closest” induced subgraph is sought after. Indeed, for most of the problems we consider below, an isomorphic copy of the subgraph is likely to not even exist in the larger graph (except for our correlation $\rho = 1$ settings). For this reason, exact subgraph isomorphism approaches cannot succeed. Research for inexact subgraph isomorphism is even more nascent with only a few recent approaches having been developed [21], [22], [23], each of which has a distinct approach and objective including spectral and higher-order contextual approaches.

A second differing emphasis of this work is that the overall goal is to understand when it is possible to find a specific set of vertices within the larger graph that correspond to the vertices in the small graph. When many isomorphic copies of the small graph are present in the large graph, it will be impossible to distinguish which, if any, of these copies is the “true” vertex-correspondence.

As such, since our goals are different from those of exact subgraph isomorphism in two ways, comparisons to exact subgraph isomorphism algorithms are not made. Additionally, our theoretical results are algorithm-independent, applying to any approach which attempts to optimize the given objective functions.

2.2 Graph Matching Algorithms

Solving the GMP problem exactly is NP-complete in general, but there are a number of approaches which provide good approximate solutions [8], [10]. Some approaches rely on tree-based methods for exact branch and bound algorithms for integer programs [24]. While these approaches are guaranteed to find a global minimum, their computational scaling is very poor. For larger graphs, the constraints for the GMP are often relaxed so that continuous optimization machinery can be deployed [15], [25], [26]. The relaxed solutions are then projected back onto $\mathcal{P}$ yielding an approximate solution to the graph matching problem. Some relaxations involve applying spectral methods which enables the application of fast linear algebra techniques [27]. A convex approach has also been proposed but requires the graph to have certain properties [28], which do not frequently hold in practice.

For the computational experiments in this manuscript we will rely on a principled indefinite relaxation of the GMP constraints [16] to the set of doubly stochastic matrices $\mathcal{D}$, the convex hull of $\mathcal{P}$. Details of this algorithm are discussed in Section 3.2. This approach has been shown to perform very well in practice on a number of benchmark problems, especially when exploiting multiple random restarts to avoid poor performing local minima [15]. We expect that many of the phenomena observed in our experiments would appear in similar fashions for other algorithms.

Frequently, these approaches can exploit *seeded vertices*, i.e., a partial list of known correspondences between the two vertex sets. When sufficiently many seeds are present, these algorithms, which often have few guarantees without seeds, can be solved efficiently

and the resulting match can be guaranteed to be correct [14], [29] asymptotically almost surely for relatively general random graph models. While the theory we discuss below does not require seeds, our algorithms will use seeds and the algorithmic performance relies heavily on a sufficient number of a priori available seeds.

In this work, we employ a soft seeding approach [30] which uses the prior information to initialize the algorithm but does not insist that the correspondence for seeds be fixed; i.e., if the seeded vertices are $[m]$ in both networks, we initialize the gradient descent graph matching algorithm at $I_m \oplus D$ for suitable D in $\mathcal{D}_{n-m}$. Conversely, in the hard-seeded setting we optimize over Π_{n-m} with the global solution being then of the form $I_m \oplus P$ for $P \in \Pi_{n-m}$. This allows for the possibility that seeds may be noisy, allowing for the algorithm to potentially correct for incorrect initial seeds, a scenario that we do not explore further here.

2.3 Statistical Models

In order to study the applicability and limitations of a graph matching approach for subgraph detections, we analyze the problem from a statistical viewpoint in the setting of the correlated heterogeneous Erdős-Rényi model [16]. The algorithmic approaches we use in our empirical investigations are only partially inspired by this model. The model enforces that the objective functions considered are not unreasonable but, except for the oracle approach, the objective functions are based primarily on topological considerations.

The following definition provides a distribution for pairs of random graphs with different numbers of vertices where there is a positive correlation between corresponding vertex-pairs. In particular, the marginal probabilities of edges in each graph can vary and the presence of an edge in one graph increases the conditional probability of the corresponding edge in the second graph. In a social network setting, this can be interpreted as two people who are related in one graph are more likely to be related in the other graph. Without loss of generality, we assume that the smaller graph (of order n_c) is denoted A and that the corresponding vertices to A in B are the first n_c vertices.

Definition 1 (Correlated Erdős-Rényi). Suppose $\Lambda \in [0, 1]^{n \times n}$ and $R \in [0, 1]^{n_c \times n_c}$ for $0 < n_c < n$. Denote by Λ^c , the order n_c principal submatrix of Λ . A pair of adjacency matrices $(A, B) \sim \text{CorrER}(\Lambda, R)$ if $A \in \mathcal{A}_n$, $B \in \mathcal{A}_n$, and for each $u < v$, B_{uv} are independent with $B_{uv} \sim \text{Bernoulli}(\Lambda_{uv})$ and for each $u < v < n_c$, A_{uv} are independent with $A_{uv} \sim \text{Bernoulli}(\Lambda_{uv})$. Additionally, for each $u, v, u', v' \in [n_c]$ with $u < v$, $u' < v'$, $u \neq u'$, $v \neq v'$, $B_{u,v}$ and $A_{u',v'}$ are independent except that the Pearson correlation $\text{corr}(A_{uv}, B_{uv}) = R_{uv}$.

When $n_c = n$, it can be shown that the solution to the GMP will asymptotically almost surely yield the correct vertex correspondence under mild assumptions on Λ and R , i.e., the unique element in the argmin in Eq. (1) is the identity matrix I [29], [16].

Note that other authors have considered equivalent reparameterized versions of the model above. In particular, a common viewpoint is that each observed graph is a random subgraph of some common underlying graph, which also arises from independent edge sampling [12]. By varying the parameters for the underlying graph and the parameters for the random

subgraph, any correlated Erdős-Rényi can be represented in this way, provided all correlations are non-negative. The reverse is also true, and more distributions can be constructed by allowing for negative correlations and non-identical distributions. Another closely related possibility is that one graph is a noisy observation of an induced subgraph from the other graph, where there may be extra or missing edges [16]. This can also be related to the correlated Erdős-Rényi by an appropriate reparameterization.

The marginal distribution of each graph is an independent edge graph which is closely related to ideas of exchangeable random graphs and latent position models which have been used extensively in various application areas [31], [32]. Any random graph, including exchangeable random graphs, can be viewed as a mixture of independent edge random graphs [33]. Viewed from this perspective, much of the theory we consider below can be extended to such mixtures presuming that the assumptions hold with sufficiently high probability across mixture components.

3 PADDING APPROACHES

In order to match pairs of graphs with differing numbers of vertices, we propose to pad the smaller graph/matrix with enough rows and columns to match the order of the larger matrix. We will consider a trio of padding schemes which will result in differing qualities for the resulting match [14]: *Naive Padding* Let $\tilde{A} = A \oplus \mathbf{0}_{n_j}$ and match $\tilde{A}$ and B ; *Centered Padding* Let $\tilde{A} = (2A - \mathbf{J}) \oplus \mathbf{0}_{n_j}$, let $\tilde{B} = 2B - \mathbf{J}$, and match $\tilde{A}$ and $\tilde{B}$; *Oracle Padding* Let $\tilde{A} = (A - \Lambda_{n_c}) \oplus \mathbf{0}$, let $\tilde{B} = B - \Lambda$, and match $\tilde{A}$ and $\tilde{B}$.

As we will see in the next section, the naive padding scheme—which finds the best fitting subgraph of B to match with A —will not be guaranteed to find the true correspondence between nodes. In contrast, the centered padding scheme—which finds the best fitting induced subgraph of B to match with A —and the oracle padding scheme are guaranteed to succeed under mild model conditions, even in the presence of an exponentially small (in terms of the size of B) subgraph A . In general, the oracle padding scheme will be inaccessible as Λ is unknown, but using various methods to estimate Λ [34], [35], we can approximate Λ in ways that can improve matching performance. Importantly, the padding and centering schemes are not tied to the algorithms described Section 3.2, and can be used with any graph matching approach which uses the adjacency matrices.

Note that each of these padding approaches corresponds to a particular graph detection problem. With the naive padding scheme, the optimization program is equivalent to the problem $\arg\max_{\sigma: [n_c] \rightarrow [n]} \sum_{u,v} A_{\sigma(u)\sigma(v)} B_{uv}$, where $\rightarrow$ denotes that σ is injective. Hence, the global optimum reveals the largest common subgraph of B and A , as only edge agreements are rewarded whereas edge disagreements are not penalized. Note that if B has a large clique, then the naive approach will tend to match A to this large clique as all edges in A are present in the large clique. As we will see below, the naive padding approach can be guaranteed with high probability to yield an incorrect vertex correspondence under only mild assumptions on R and Λ .

If the graphs have the same number of nodes, both the centered padding scheme and the naive padding scheme are equivalent [36], while they are not equivalent for graphs of different orders. With the centered padding scheme, the optimization program is equivalent to the problem $\operatorname{argmax}_{\sigma: [n_c] \rightarrow [n]} \sum_{u,v} |A_{\sigma(u)\sigma(v)} - B_{uv}|$. In the case, the global optimum reveals the induced subgraph of B which is nearest to A in terms of edge-edit distance [36]. As will be seen below, the optimum to this program can be guaranteed to correspond the true vertex correspondence provided that the graphs are not too sparse and the correlation is sufficiently large.

The oracle padding scheme corresponds to the problem,

$$\operatorname{argmax}_{\sigma: [n_c] \rightarrow [n]} \sum_{u,v} \widehat{\operatorname{Cov}}(A_{\sigma(u)\sigma(v)} B_{uv}),$$

where $\widehat{\operatorname{Cov}}$ denotes the empirical covariance when both means are known, $\widehat{\operatorname{Cov}}(A_{ij}, B_{uv}) = (A_{ij} - \Lambda_{ij})(B_{uv} - \Lambda_{uv})$. As this approach eliminates any undue rewards from matching likely edges to even more likely edges or unlikely edges to even more unlikely edges, the global optimum is guaranteed with high probability to be the true correspondence even under mild assumptions on Λ and R . Indeed, even if A and B were non-identically distributed, with means Λ_A and Λ_B respectively, by matching $A - \Lambda_A$ and $B - \Lambda_B$, the theoretical results will still hold [37].

While the theory below shows that under the correlated heterogeneous Erdős-Rényi model the oracle padding yields stronger guarantees than the centered padding method and the naive padding method can have very bad performance, this will not necessarily translate in practice. This can occur especially due to the fact that we do not find the true global optimum of the GMP but rather a local optimum near our initialization. Our simulations and real data examples illustrate how centering or naive may be preferable in certain scenarios, and biased estimates of the oracle may also improve performance. In practice, one may try different paddings and even use the solution of one padding approach to initialize the optimization for another padding approach. See Section 4 for examples where various approaches, including matching for a second time, may provide more or less ideal performance.

3.1 Theory

For each of the padding scenarios, we will consider n_c and n as tending to ∞ in order to understand the ability of these optimization programs to detect progressively larger subgraphs. Note that we will require that the number of vertices n_c is growing with n . Indeed, if n_c is fixed and n grows then eventually every subgraph of size n_c will likely appear multiple times as an induced subgraph in B just by chance [38].

For each padding scenario let $P^* \in \{0, 1\}^{n_c \times n_c}$ denote the order n_c principal submatrix of the solution to the corresponding graph matching problem. The proofs of each theorem below can be found in Appendix B

The first theorem is a negative result that shows that under weak assumptions on Λ^c , one can construct a Λ under which the naive padding scheme is almost surely guaranteed to not detect the errorful version of A in B . Indeed, it can be guaranteed to recover zero true vertex correspondences with high probability.

Theorem 2. Suppose $R \in [0, 1]^{n_c \times n_c}$ satisfies $R < \rho$ entry-wise for some scalar $\rho \in (0, 1)$. Suppose $n_c < n/2$, $\beta \in (0, 1)$, and entry-wise $\Lambda^c \leq \beta$ and $\Lambda^c = \omega(n_c^{-1} \log n_c)$. Then there exists Λ such that for $A, B \sim \text{CorrER}(\Lambda, R)$, using the naive padding scheme,

$$\mathbb{P}[P^* \neq I] \geq 1 - \exp\{-C\epsilon^2(\log n_c)^2\}, \quad (2)$$

for some universal constant $C > 0$ and $\epsilon \in (0, 1 - \beta - (1 - \beta)\rho)$.

This result hinges on the fact that the naive padding scheme finds the best matching subgraph, rather than the best matching induced subgraph; indeed, there is no penalty for matching non-edges in A to edges in B . Hence, if B has a dense substructure of size n_c which does not correspond to A then the naive padding scheme will match A to that dense structure with high probability, regardless of the underlying correlation. The proof demonstrates an extreme case where all vertices in A are matched (incorrectly) to vertices in B without true matches. Deviations on the proof will yield that any number of vertices may be mismatched with only a mildly denser substructure in B . Recall that since the centered padding scheme and the naive padding scheme are equivalent when $n = n_c$, the requirement that n_c is sufficiently small is necessary.

On the other hand, the centered padding scheme can be guaranteed to result in the correct detection of the subgraph even when the number of vertices in B is exponentially larger than the number of vertices in A provided $R > 1/2 + \epsilon$.

Theorem 3. Suppose that $A, B \sim \text{CorrER}(\Lambda, R)$ with $R_{uv} \in [1/2 + \epsilon, 1]$ and $\Lambda_{uv} \in [\alpha, 1 - \alpha]$ for $\alpha, \epsilon \in (0, 1/2)$. It holds that $\frac{\log(n)}{\epsilon^2 n_c \alpha (1 - \alpha)^2} = o(1)$ implies that using the centered padding scheme

$$\mathbb{P}[P^* \neq I] \leq 2\exp\left\{-\Theta\left(n_c(\epsilon\alpha(1 - \alpha))^2\right)\right\}.$$

Stated simply, this theorem only requires that the correlations are sufficiently large to guarantee that large subgraphs of logarithmic size can be found via an oracle GM algorithm.

Analogous to the naive padding scheme, the centered padding scheme does have its limitations. The Theorem below indicates that if n_c is too small then with high probability, the centered padding scheme will fail to recover the true correspondence for certain Λ, R . Importantly, for the centered padding scheme this is a much weaker limitation than for the naive scheme.

Theorem 4. Suppose $R, \Lambda_c \in [\alpha, 1 - \alpha]^{n_c \times n_c}$ for some $\alpha \in (0, 1/2)$. There exists a constants $\xi, \epsilon > 0$ such that if $n_c < \xi \log n$ then there exists a $\Lambda \in [0, 1]^{n \times n}$ such that

$$\Lambda_{uv} (1 - \Lambda_{uv}) R_{uv} \geq 1/2 + \epsilon$$

for all $u, v \in [n_c]$, and for $A, B \sim \text{CorrER}(\Lambda, R)$, using the centered padding scheme,

$$\mathbb{P}[P^* \neq I] = 1 - o(1).$$

Note, that the $O(\log n)$ rate corresponds to the information theoretic limit for detecting a planted clique within a random graph, where detecting a clique of size less than $2 \log_2 n$ is impossible [39].

Finally, while the oracle padding is inaccessible for general Λ , it represents the optimal padding scheme as it eliminates any empirical correlations introduced by Λ leaving only the theoretical correlations from R .

Theorem 5. Suppose that $A, B \sim \text{CorrER}(\Lambda, R)$ with $R_{uv} \in [\rho, 1]$ and $\Lambda_{uv} \in [\alpha, 1 - \alpha]$, for some $\alpha, \rho \in (0, 1)$. Then, using the centered padding scheme,

$$\mathbb{P}[P^* \neq I] \leq 2 \exp \left\{ - \Theta \left(n_c (\rho \alpha (1 - \alpha))^2 \right) \right\}.$$

Remark 6. Note that while the padding schemes are still not equivalent in the homogeneous Erdős-Rényi case, the distribution of the matrix P^* are equal under all three padding schemes and hence all are guaranteed to find the correct solution with high probability under identical scenarios.

3.2 Computation

Our approach to solve the graph matching problem in this setting will be identical to an approach for graphs of equal sizes; see the FAQ algorithm of [15] for full detail. In particular, we will relax the GMP constraints from $\mathcal{P}$ to $\mathcal{D}$ and use gradient descent starting at a given D_0 . We will also incorporate seeds by appropriately modifying D_0 . This gradient ascent approach is then given by Algorithm 1.

```

Data:  $A, B \in \mathcal{A}, D_0 \in \mathcal{D}, k = 0$ 
while not converged do
1    $P_k \leftarrow \text{argmin}_{P \in \mathcal{P}} -\text{trace}(\tilde{A} D_k \tilde{B} P);$ 
2    $\alpha_k \leftarrow \text{argmin}_{\alpha \in [0,1]} -\text{trace}(\tilde{A} D_\alpha \tilde{B} D_\alpha),$  where
       $D_\alpha = \alpha D_k + (1 - \alpha) P_k;$ 
3    $D_{k+1} \leftarrow D_{\alpha_k}$  and  $k \leftarrow k + 1;$ 
end
4 Project  $D_k$  onto  $\mathcal{P}$  yielding  $P^*$ ;
5 Return  $P^*$ 

```

Algorithm 1: Fast approximate quadratic assign program (FAQ) algorithm [15] for graph matching.

Note that using any of the padding schemes, we do not need to store or compute the entire matrices D_k or P_k as the objective function only depends on their first n_c rows. Hence, lines 1 and 4 can be simplified and accomplished by searching over the set of $n_c \times n$ matrices corresponding to injections from $[n_c]$ to $[n]$, or equivalently the first n_c rows of permutation matrices. In this way, lines 1 and 4 can be solved effectively by variants of the Hungarian algorithm for non-square matrices [40]. Line 2 is a quadratic equation in α and is easily solved. Furthermore, by exploiting the fact that $\tilde{A}$ and $\tilde{B}$ are likely formed by the difference between a sparse and low-rank matrix, fast matrix multiplication and storage can also be exploited for further runtime improvements.

The computational complexity of each iteration is that of solving a rectangular linear assignment problem which can be solved in $O(n^2 n_c)$. The memory/storage complexity for the graphs is n^2 for dense matrices and on the order of the number of edges for sparse graphs. For the doubly stochastic matrix, this can be stored in $O(n_c n)$ space. However, as $D_k = \left(\prod_{j=1}^k \alpha_j\right) D_0 + \sum_{i=1}^k \left(\prod_{j=1}^i \alpha_j\right) P_i$ memory (and computational) gains can be made as long as D_0 has a sparse plus low-rank structure. As an example, below we consider $D_0 = \alpha_0 \mathbf{I}/n + (1 - \alpha_0)P$ for some permutation matrix P . In this case D_0 can be stored effectively with $O(n_c)$ memory and D_k can be stored with at most $O(k n_c)$ memory. The most memory intensive aspect of the problem is storing the gradient, which in general will not be sparse, however it may also have a sparse plus low-rank structure.

The convergence criterion is generally easy to check as the optimal doubly stochastic matrix is frequently itself a permutation matrix, which also means the final projection step can be omitted. While this algorithm is not guaranteed to converge to a global optimum, if there are enough seeds or if the matrix D_0 is sufficiently close to the identity, the local maximum which this procedure converges to will likely be the identity.

4 EXPERIMENTS

In this section we demonstrate the effectiveness of a graph matching matched filter for errorful subgraph detection in both synthetic and real data settings. Without loss of generality, we will assume that the seeds are the first s vertices for some $s \leq n_c$. The matched filter algorithm is given in Algorithm 2

Data: Template A , network B , number of seeds s ,
number of MC replicates M , padding scheme
Result: Matchings $\{B_1, B_2, \dots, B_M\}$

```

1  $\tilde{A}, \tilde{B} \leftarrow$  padded and centered  $A, B$  according to
padding scheme;
for  $i \leftarrow 1$  to  $M$  do
2   Generate a random doubly stochastic matrix

$$D_0 \leftarrow I_s \oplus \left( \alpha P + \frac{1-\alpha}{n-s} \mathbf{J}_{n-s} \right),$$

where  $P$  is a random permutation matrix and
 $\alpha \sim \text{Unif}(0, 1)$ ;
3   Initialize FAQ at  $D_0$ ; match  $\tilde{A}$  to  $\tilde{B}$ ;
4   FAQ output is  $P^*$ , assignment of vertices in  $B$  to
vertices in  $A$ 
end

```

Algorithm 2: Matched filters algorithm via random restarts, soft seeding, and gradient descent.

As described in Line 2, we initialize the FAQ algorithm at a random start point which initially aligns the seeded vertices across graphs. In particular, our random starting point has the identity as its principal $s \times s$ submatrix, corresponding to the seeds. The remaining matrix is sampled as a random convex combination of the the baricenter matrix $\frac{1}{n-s} \mathbf{J}_{n-s}$ and a random permutation matrix sampled uniformly from $\mathcal{P}_{n-s}$. Repeating this process M times, we output M potential matches for the subgraph A in B .

To explore the effectiveness of this procedure, we consider applying our matched filter in the following scenarios: pairs of homogeneous CorrER graphs where we vary the uniform correlation, a planted partition model where the goal is to find the dense partition, a heterogeneous model with two different subgraph sampling schemes, and an application to *Drosophila* and Human connectomes.

4.1 Correlated Homogeneous Erdős-Rényi

As a synthetic data example, we will consider subgraph detection in the CorrER($\lambda, \mathbf{J}, \rho \mathbf{J}$) with $\lambda = 0.5$ (i.e., the maximum entropy ER model) for $\rho = 0.7, 0.8, 0.9$, and 1. When $\rho = 1$, the induced subgraph of B corresponding to A is exactly isomorphic to A . We consider $n = 500$, $n_c = 40$, and $M = 200$ and we replicated the entire experiment for 200 sampled graph pairs. For each experiment, we rank the objective function values, computed as $\sum_{i,j < n_c} (\tilde{A} - P \tilde{B} P^T)_{ij}^2$, among the $M = 200$ random restarts, from smallest to largest.

Figure 1a shows the average performance, using the centered padding scheme, across the replicates at each rank. As is evident, without seeds very few vertices are matched correctly, even when the corresponding graphs are isomorphic. However, as both the number of seeds increase and as the correlations increase, the performance of the procedure becomes very good. Note that when $\rho = 1$, we would expect exact subgraph matching procedures to perfectly recover the correspondence, while they would fail for all lower correlations. These

exact procedures do not require seeds but generally scale poorly with the number of vertices in each graph.

Interpreting the correlation 0.9 panel, if 15 seeds are used then nearly all restarts will yield a perfect recovery of the original correspondence. Hence, only one restart is required. Using 7 seeds, most restarts yield suboptimal performance, with only slightly more than the original 7 seeds recovered. However, if we perform 200 restarts and choose the best among them then this will often yield a near perfect recovery of the original correspondence. Similar interpretations can be made for the other correlations.

Figure 1b gives a complementary view of this same simulation, showing the average value of the objective function as a function of the number of matched vertices. Each point is scaled according to the average number of restarts (out of 200) which achieve the given number of matched vertices. For $\rho = 0.7$, the objective function decreases smoothly with the number of matched vertices, and the number of matched vertices is relatively evenly spread, especially for 15 seeds. On the other hand, for higher correlations and for 15 seeds, there is a substantial gap in both the objective function and the matched vertices.

This gap is highly important for practical applications. Indeed, the presence of a gap in the objective function would indicate that the correct matching was likely found. Without the gap, the guarantees that the best match in terms of the objective function corresponds to a correct match is less certain (at least to practically ascertain). Knowledge of the size of the gap would also allow for a fast and adaptive algorithm which performs random restarts until a sufficient objective function gap is observed. This is discussed further in Section 5.

Finally, table one shows average run times, on a standard laptop, for each restart under each parameterization. Even under the worst of the settings, 200 restarts can be performed in less than 10 minutes. It is worth noting that better algorithmic performance correlates to faster algorithmic runtime, a feature which was observed in previous work using the FAQ algorithm [14].

4.2 Planted Partition Correlated Erdős-Rényi graphs

As another scenario, we consider the problem where the larger graph has a denser subgraph which corresponds to the smaller graph. We take $n = 500$, $n_c = 40$, $A, B \sim \text{CorrER}(\Lambda, 0.9\mathbf{J})$ with

$$\Lambda = \begin{pmatrix} q\mathbf{J}_{n_c} & p\mathbf{J}_{n - n_c, n} \\ p\mathbf{J}_{n, n - n_c} & p\mathbf{J}_{n - n_c} \end{pmatrix},$$

where $p = 0.25$ and $q \in \{0.25, 0.3, \dots, 0.5\}$. For 100 random restarts, 7 seeds and using the centered padding, we replicated each setting 200 times. These parameters were chosen as illustrative examples that demonstrate the general phenomena described below.

Figure 2 shows the average number of correctly matched vertices, using the centered scheme, at each objective function rank for each value of q . As would be expected, the

problem becomes significantly easier as q increases. A similar gap phenomenon is observed in this setting, as long as q is sufficiently large, though we do not illustrate it here.

4.3 Random Dot Product Graph

In order to better evaluate the performance of the procedure under heterogeneous graph models, we will now consider graphs from the random dot product graph (RDPG) model [41], [42]. The RDPG model provides a generalization of the stochastic blockmodel, including mixed-membership [43] and degree-correction [44], allowing for multiple forms heterogeneity in vertex behavior. As we show below, this heterogeneity allows for the exploration of how the latent structure of the induced subgraph influences performance for our methods. For this model, the random adjacency matrix A has a low-rank expectation $\mathbb{E}[A]$ which is conducive to approximation using spectral methods. Below we center by a low-rank approximation of the adjacency matrix, which is probabilistically guaranteed to be an accurate estimate of $\mathbb{E}[A]$ [45]. This serves to approximate the oracle procedure.

Each vertex v is associated with a latent position $X_v \in \mathbb{R}^2$, which is sampled uniformly from the set $\{(x, y) \in (0, 1)^2 : x + y < 1\}$. In this case the matrix Λ is given by $\Lambda_{uv} = X_u^T X_v$. We again take $n = 500$ and $n_c = 40$ and sample $A, B \sim \text{CorrER}(\Lambda, 0.9\mathbf{J})$. The matched subgraphs are sampled in two different ways. Under one sampling design, the matched nodes simply correspond to a random sample of the latent positions. Under the “max-angle” design, the matched nodes are selected to be those with the largest ratio between the second and first elements of the latent positions. Figure 3 shows a representative set of latent positions corresponding to this model. The red circles correspond to the max-angle sampling and the points in blue correspond to a possible random sampling (these do not overlap for the purposes of visualization).

We sampled 200 distinct graphs under each subgraph sampling setting and performed 100 random restarts for each sampled graphs. For each restart, the matchings were performed with the naive, centered, and two “oracle”-type padding schemes. As the model above is a low-rank model, we centered each matrix according to its best rank-1 and rank-2 approximations to approximate the oracle padding scheme. Letting,

$$\hat{\Lambda}_r^{(M)} = \underset{\Lambda : \text{rank}(\Lambda) = r}{\text{argmin}} \|M - \Lambda\|_F^2,$$

for each of $M = A$ and B , these schemes match $(A - \hat{\Lambda}_r^{(A)}) \oplus \mathbf{0}_{n-n_c}$ to $B - \hat{\Lambda}_r^{(B)}$ for $r = 1, 2$.

Figure 4a shows the average accuracy at each rank of the objective function. First, the max-angle subgraph is clearly much easier to discover than the random subgraph. The center scheme had by far the best performance for the max-angle case and in the random case it also performed well. The rank-2 centering scheme also performed well in the max-angle case and had very similar performance to the center scheme. Both the rank-1 and naive padding schemes failed to display good performance in either case.

Under either scenario, the rank-one approximation eliminates most degree based differences between vertices. In some cases this may be helpful if the degree information is not informative about the match, but in this case it appeared to hurt performance in both cases. Rank-two centering in this case seems to restore some of the performance losses from rank-1 centering. This also is somewhat unexpected since the heterogeneity from both principal directions can be used to help identify the vertices. The reasons that the rank-2 centering is better than the rank-1 centering are unclear and an area of current investigation. Note that constant centering maintains all differential vertex behavior which is a strong indicator of which vertices to match to. The naive padding fails here as there are always sets of vertices with similar relative propensities for adjacencies to other vertices but with strictly larger absolute probabilities of adjacencies.

To understand why the max-angle problem is so much easier, note that, under either subgraph sampling scheme, there will be many nodes which have similar behavior to one or more of the sampled nodes due to having nearby latent positions. Under random sampling, most of these similar nodes are likely to correspond to nodes that do not have a match and hence will introduce errors. However, under max-angle sampling the similar nodes have a better chance of being other matchable nodes which will introduce fewer errors as the theoretical correlation will strongly distinguish these vertices.

Figure 4b shows the average value of the objective function as a function of the number of matched vertices. Since the different centering schemes have different objective functions which will have slightly different scales, we rescaled objectives functions for each scheme so that the largest value for each was set to one and the smallest was set to zero. The original objective functions for the rank-1, rank-2, and center schemes ranged from 7–23, 8–18, and 11–36 for the max angle setting and 15–25, 9–17, and 9–31 for the random setting, respectively. Each point is sized according to the average number of restarts (out of 100) which achieve the given number of matched vertices. Note that we again observe a gap in the objective function between the bulk of solutions with less than 20 correct matches and those with more than 35 correct matches.

4.4 Finding Kenyon cells in a *Drosophila* connectome

As an application to brain networks, we consider using the matched filter to locate an induced subgraph of Kenyon (K) cells in the fully reconstructed *Drosophila* mushroom body of [46]. Using the induced subgraph of the K-cells in the right hemisphere of the mushroom body (i.e., as *A*), we seek to find the K-cells on the left hemisphere (i.e., as *B*). Although in this example, the K-cells are identified across both hemispheres, this was achieved only with great effort and expenditure. Being able to use one hemisphere to locate structure in the other hemisphere could potentially allow for faster, cheaper neuron identification in future connectomes. In particular, this can be extended to finding similar structures across connectomes collected from multiple individuals.

After initial data preprocessing, there are $n_c = 100$ K-cells in each hemisphere and $n = 209$ vertices total in the right hemisphere. The pair of graphs are illustrated in Figure 5, with the first 100 nodes in the left image being the K-cells of the left hemisphere.

As the true correspondence between these pair of graphs is unknown, the accuracy is measured by the number of cells in the right K graph which are matched to K -cells in the left graph, rather than the number of correctly matched pairs.

We consider $s = 0$ and 10 soft seeds and $M = 50,000$ random restarts for our matched filter. The seeds were determined using cells which were given the same name across the two hemispheres.

In addition to performing the matching in our previous examples using naive, centered, and low-rank centered schemes, we also consider a second “re-matching”. Letting D_1^* denote the doubly stochastic solution found by the first matching (prior to projecting to the permutations), we consider rematching starting at D_1^* using a (different) low-rank centering.

In this way the initial matching may find gross similarities which are then refined by removing more empirical correlation between the graphs. Whether this procedure is effective will of course depend on the particular choices for the first and second matching.

Figure 6 shows the number of correctly matched vertices at each rank of the objective function, after the first match. The ranks were linearly rescaled to be between 0 and 1. Importantly, only the ranking of the initial ranking is used, (not of the second ranking).

In this instance we smoothed the performance across ranks using a Gaussian kernel smoother with bandwidth 0.02. The performance of each matching scheme and the scheme used for the second matching can all have substantial impacts on the overall performance. For example, while the centered scheme alone has very poor performance, if it is followed by the rank-1 scheme then the performance becomes better than the naive.

The overall best performance, according to this smoothed performances estimate, is achieved by using the rank-1 scheme followed by the rank-3 scheme. As these graphs have many matches with similar performance, each of which are local minima for the various objective functions, we fail to observe the substantial gap between accurate and inaccurate solutions. Indeed, most solutions under the rank-1 scheme achieve between 80 and 91 correct matches.

We note that in this instance the number of seeds barely impacts performance. This is likely due to two reasons. These seeds are merely guesses at some true correspondence and the best matchings often does not maintain all of the original seeds. Indeed, using 10 seeds, nearly every match does not include at least one of the original seed pairs. Also, as can be seen in Fig. 6 the graphs have a very distinctive structure where certain nodes, such as those with indices greater than 125, are very unlikely to be matched because of the directed nature of the graph. These nodes can be excluded, which is the opposite side of the coin of having seeds, and also makes the matching problem easier.

In order to visualize matches across restarts, for each pair of vertices across graphs, we can count the number of restarts where the pair was matched. Fig. 5c shows a heat map for these counts for the matches achieved using the rank-1 centering followed by the rank-3 centering

using 0 seeds. The total number of restarts was 5000, corresponding to the column sums of the corresponding matrix.

Most of the matchings occur between Kenyon cells and even within these matching there are clear patterns of certain vertices frequently being matched to each other and groups of vertices that appear to be mapped together. Note also that many match pairs occur infrequently, with 32% of pairs never being matched. 58% of pairs which were matched were matched less than 100 times. This visualization and the associated matrix could easily be further analyzed to discover distinct patterns and grouping of vertices.

Table 2 shows the matching performance (out of 100) for the solution which achieved the smallest objective function (after the first matching). Using this simple method for selecting a match, we achieve 89 correct matches. The best overall matches were able to find 94 K-cells but these matches did not appear among the top 1000 objective functions values.

One might imagine that better performance could be achieved by sorting according to the objective function value for the second match. However, this actually results in adverse behavior where lower objective functions result in poorer performance. While we do not exactly understand this phenomenon, if we view the second match as only a refinement, then we postulate that lower objective function values have strayed farther from the initial match, which may reduce performance if the initial match is good.

4.5 Large-scale DTMRI Graphs

Diffusion tensor magnetic resonance imaging (DTMRI) is a form of MRI which allows for the estimation of coarse scale approximations of the networks of a human brain [47], [48], so called connectomes. While the estimated graphs frequently only have 10s or 100s of vertices, recent tractography methods and parcellation techniques have allowed the creation of connectomes with tens of thousands of vertices. To illustrate the ability of algorithm to scale to larger problems, we considered a matching problem on one such graph.

In particular, we used the DS72784 atlas from the neurodata repository [47], and considered a problem analogous to finding Kenyon cells in the *Drosophila* connectome. The graph has ≈ 40 thousand nodes. Using the coarser Desikan atlas, each vertex can be assigned to one of 70 regions, 35 in each hemisphere. Corresponding regions are numbered by r and $r + 35$, for $r = 1, \dots, 35$.

As in the *Drosophila* example, we considered the problem of matching the subgraph corresponding to a region in one hemisphere, to the other hemisphere. In particular, we considered the matching of induced subgraphs of regions 7, 20, 33 and 34 in the left hemisphere to the right hemisphere, and matching regions 42, 55, 68, and 69, in the right hemisphere to the left hemisphere. These regions were picked because they were relatively close in number of vertices and edges and not too large. Details, for the numbers of vertices and edges in each graph are in the first three columns Table 3.

For these experiments we used the same randomized initialization procedure and considered using 0, 10 and 20 seeds. As the true vertex correspondence is unknown, for seeds we randomly selected pairs of vertices from the matched regions. Each of the eight region-to-

hemisphere matchings was repeated 400 times for each seeding level. While the number of iterations required per repetition varied, the time per iteration ranged from 5.6 ± 1.5 seconds for the smallest graphs to 8.1 ± 2.6 seconds for the largest graphs.

As the goal for these experiments was primarily to show the ability of our methods to scale to larger problems, we only considered using the centered padding scheme. The results in terms of the number of vertices correctly identified as belonging to the corresponding region indicate that using this procedure we were unable to accurately locate the corresponding region.

The rightmost 6 columns of Table 3 show the number of correctly matched vertices and the number of edge errors for the best edge error match, with the fewest number of edge errors, and the best region match, the match with the most vertices matched to the correct region. We also indicate the number of seeds used for each of these matchings. Note that the fewest edge error matching all used 0 seeds, had less than half the number of edge errors as the best region match, and generally did not find any vertices in the correct region.

Figure 7 shows an example of the two matchings for matching region 7 in the left hemisphere. The matching on the left is the best matching achieved in terms of minimizing the objective function and the matching on the right is the best matching achieved in terms of vertices correctly identified as belonging to the corresponding region. The central figure shows the adjacency matrix for the original small graph and the top left and right show the adjacency matrices for the induced subgraphs that were found by the procedure. The bottom left and right show the difference between the original graph and the matched graph, with missing edges shown in blue and extra edges shown in red. Both matches appear to share the overall structure of the center plot but the left match matches many finer details that the right match does not.

5 DISCUSSION

In this manuscript we have proposed a number of padding methods to transform the noisy subgraph detection problem into a graph matching problem. Each padding scheme emphasizes a different aspects of common graph structure. Theoretically, the centered and oracle padding schemes are guaranteed to recover the original vertex correspondence in the correlated Erdős-Rényi statistical model, provided correlations and edge probabilities do not decay too quickly and the order of the smaller graph is at least logarithmic in the order of the larger graph.

By using a partially known vertex correspondence and random restarts, we are frequently able to recover the full correspondence accurately in a number of synthetic data settings and in a connectomics setting, using a gradient ascent approach on a relaxation of the original problem. Importantly, we frequently observe a gap in the objective function value between a restart that achieves good performance and those with poor performance [49], [50]. If the size of the gap can be predicted, it can be used to detect the fact that the correct alignment has been found rather than a spurious local minimum, and additionally can provide an adaptive stopping criterion for performing random restarts.

Rigorous results for the expected gap size in the correlated Erdős-Rényi model are currently unknown. Two recent results for graphs with the same number of vertices do provide some hints. [51] investigates the ratio of $\|A - B\|_F^2$ and $\mathbb{E}_P \|A - PBP\|^T$, where the expectation is taken with respect to a uniform distribution over all permutation matrices. The authors show that this quantity will converge to $1 - \rho_T$, where ρ_T is the “total correlation” of the graph pair, incorporating the correlation of the edge probabilities and the correlation of the edge presence random variables. This result implies that $\|A - B\|_F^2 \approx (1 - \rho_T) \mathbb{E}_P [\|A - PBP\|_F^2]$.

Of course, the error for a local optimum will not necessarily be close to the error for a random permutation, so this result is not directly applicable. Another related result is in [30], which shows that under certain conditions there will be no local optimum which correctly aligns more than $\theta(\sqrt{n})$ vertices. This gap in the number of aligned vertices presumably leads to a gap in the objective function. Another related avenue is that for quadratic assignment programs for independent homogeneous Erdős-Rényi graph, the deviation of the objective function at the solution compared to the objective at any fixed permutation has been well studied [50], [49]. Further explorations along these lines are active areas of research.

As the best objective function does not always correspond to the best matching performance, it can be difficult to evaluate performance when ground truth is not present. One way to quantify uncertainty is to study the proportion of times that two nodes are matched across graphs among the random restarts. By essentially taking the average of the permutations, one could determine which matches are most probable and which have higher uncertainty. We demonstrate this idea in Figure 5c where we see patterns of matches indicating that certain groups of vertices are more similar within the set of K-cells.

We have assumed that the induced subgraphs are identically distributed and positively correlated. These assumptions are important for our proof. Indeed, if negative correlation is allowed, the objective functions are seeking the wrong solution. Non-identical graphs can be more easily dealt with by using separate oracle estimates, as was done in the random dot product and Drosophila connectome examples. Further theoretical issues such as non-independent edge graph may also be possible to deal with by exploiting the possibility that the graph pair is constructed out of a still large number of independent Bernoulli trials, none of which can drastically change graph structure. Such an assumption may allow for the application of Proposition 7 as in our proofs below.

The general question of which pairs of graphs are matchable, by which we mean the “true” correspondence is recoverable, remains open. One of the key aspects that make it possible in the correlated Erdős-Rényi model is that the noise introduced by the independent edges introduces asymmetries in the observed graphs, so that there are few or no automorphisms or even near automorphisms. Hence, under the correlated setting, all but the true correspondence lead to an increase in the number of edge discrepancies. For two arbitrary graphs, it is reasonable to expect that a similar conditions would be required, namely that there is a sufficient agreement between the graphs at the true correspondence, and that each graph is sufficiently asymmetric so that no other correspondence also has a small number of edge disagreements.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This material is based on research sponsored by the Air Force Research Laboratory and DARPA under agreement number FA8750-18-2-0066. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Research Laboratory and DARPA or the U.S. Government. This work is also partially supported by a grant from MIT Lincoln Labs.

Biographies

Daniel L. Sussman received the BA degree in mathematics from Cornell University, in 2008, and the PhD degree in Applied Mathematics and Statistics from JHU, in 2014. From 2014 to 2016 he worked as a postdoctoral fellow at the Harvard University Statistics Department. He is currently an Assistant Professor in the Mathematics and Statistics Department at Boston University. His research interests include spectral methods for graph inference, multiple graph inference, graph matching, connectomics, and causal inference under interference.

Vince Lyzinski received the B.Sc. degree in mathematics from University of Notre Dame in 2006, the MS degree in mathematics from John Hopkins University (JHU) in 2007, the MSE degree in applied mathematics and statistics from JHU in 2010, and the PhD degree in applied mathematics and statistics from JHU in 2013. From 2013–2014 he was a postdoctoral fellow in the Applied Mathematics and Statistics (AMS) Department at JHU. During 2014–2017, he was a Senior Research Scientist at the JHU HLTCOE and an Assistant Research Professor in the AMS Department at JHU. Starting July 2017, he has been an Assistant Professor at the University of Massachusetts Amherst in the Department of Mathematics and Statistics. His research interests include graph matching, statistical inference on random graphs, pattern recognition, dimensionality reduction, stochastic processes, and high-dimensional data analysis.

Youngser Park received the B.E. degree in electrical engineering from Inha University in Seoul, Korea in 1985, the M.S. and Ph.D. degrees in computer science from The George Washington University in 1991 and 2011 respectively. From 1998 to 2000 he worked at the Johns Hopkins Medical Institutes as a senior research engineer. From 2003 until 2011 he worked as a senior research analyst, and has been an associate research scientist since 2011 in the Center for Imaging Science at the Johns Hopkins University. At Johns Hopkins, he holds joint appointments in the Department of Applied Mathematics and Statistics and the Human Language Technology Center of Excellence. His current research interests are clustering algorithms, pattern classification, and data mining for high-dimensional and graph data.

Carey E. Priebe received the BS degree in mathematics from Purdue University, in 1984, the MS degree in computer science from San Diego State University, in 1988, and the PhD degree in information technology (computational statistics) from George Mason University, in 1993. He is currently a professor in the Department of Applied Mathematics and Statistics, JHU. His research interests include computational statistics, statistical pattern recognition, and statistical inference for high-dimensional and graph data. He is a Senior Member of the IEEE.

REFERENCES

- [1]. Carletti V, Foggia P, Saggese A, and Vento M, "Challenging the time complexity of exact subgraph isomorphism for huge and dense graphs with VF3," *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 4, pp. 804–818, 4 2018. [PubMed: 28436848]
- [2]. Kuramochi M and Karypis G, "Frequent subgraph discovery," in *Proceedings 2001 IEEE International Conference on Data Mining*, 2001, pp. 313–320.
- [3]. Slota GM and Madduri K, "Fast approximate subgraph counting and enumeration," in *2013 42nd International Conference on Parallel Processing*, Oct. 2013, pp. 210–219.
- [4]. Cordella LP, Foggia P, Sansone C, and Vento M, "A (sub)graph isomorphism algorithm for matching large graphs," *IEEE transactions on pattern analysis and machine intelligence*, vol. 26, no. 10, pp. 1367–1372, 10 2004. [PubMed: 15641723]
- [5]. Schmidt MC, Samatova NF, Thomas K, and Park B-H, "A scalable, parallel algorithm for maximal clique enumeration," *Journal of parallel and distributed computing*, vol. 69, no. 4, pp. 417–428, 4 2009.
- [6]. Bomze IM, Budinich M, Pardalos PM, and Pelillo M, "The maximum clique problem," in *Handbook of Combinatorial Optimization: Supplement Volume A*, Du D-Z and Pardalos PM, Eds. Boston, MA: Springer US, 1999, pp. 1–74.
- [7]. Ullmann JR, "An algorithm for subgraph isomorphism," *Journal of the ACM*, vol. 23, no. 1, pp. 31–42, 1 1976.

- [8]. Conte D, Foggia P, Sansone C, and Vento M, "Thirty years of graph matching in pattern recognition," *International Journal of Pattern Recognition and Artificial Intelligence*, vol. 18, no. 03, pp. 265–298, 2004.
- [9]. Foggia P, Percannella G, and Vento M, "Graph matching and learning in pattern recognition in the last 10 years," *International Journal of Pattern Recognition and Artificial Intelligence*, vol. 28, no. 01, p. 1450001, 2014.
- [10]. Emmert-Streib F, Dehmer M, and Shi Y, "Fifty years of graph matching, network alignment and network comparison," *Information sciences*, vol. 346–347, pp. 180–197, 2016.
- [11]. Lyzinski V, Sussman DL, Fishkind DE, Pao H, Chen L, Vogelstein JT, Park Y, and Priebe CE, "Spectral clustering for divide-and-conquer graph matching," *Parallel Computing*, vol. 47, pp. 70–87, 2015.
- [12]. Yartseva L and Grossglauser M, "On the performance of percolation graph matching," in *Proceedings of the first ACM conference on Online social networks*. ACM, 2013, pp. 119–130.
- [13]. Akoglu L, Tong H, and Koutra D, "Graph based anomaly detection and description: a survey," *Data mining and knowledge discovery*, vol. 29, no. 3, pp. 626–688, 5 2015.
- [14]. Fishkind D, Adali S, Patsolic HG, Meng L, Lyzinski V, and Priebe C, "Seeded graph matching," *arXiv preprint arXiv:1209.0367*, 2017.
- [15]. Vogelstein JT, Conroy JM, Lyzinski V, Podrazik LJ, Kratzer SG, Harley ET, Fishkind DE, Vogelstein RJ, and Priebe CE, "Fast Approximate Quadratic Programming for Graph Matching," *PLoS ONE*, vol. 10, no. 04, 2014.
- [16]. Lyzinski V, Fishkind DE, Fiori M, Vogelstein JT, Priebe CE, and Sapiro G, "Graph matching: Relax at your own risk," *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 1, pp. 60–73, 2016. [PubMed: 26656578]
- [17]. Samsi S, Gadepally V, Hurley M, Jones M, Kao E, Mohindra S, Monticciolo P, Reuther A, Smith S, Song W, Staheli D, and Kepner J, "Static graph challenge: Subgraph isomorphism," in *2017 IEEE High Performance Extreme Computing Conference (HPEC)*, 2017, pp. 1–6.
- [18]. Cordella LP, Foggia P, Sansone C, and Vento M, "A (sub)graph isomorphism algorithm for matching large graphs," *IEEE transactions on pattern analysis and machine intelligence*, vol. 26, no. 10, pp. 1367–1372, 10 2004. [PubMed: 15641723]
- [19]. Solnon C, "Alldifferent-based filtering for subgraph isomorphism," *Artificial Intelligence*, vol. 174, no. 12, pp. 850–864, 2010.
- [20]. Bonnici V, Giugno R, Pulvirenti A, Shasha D, and Ferro A, "A subgraph isomorphism algorithm and its application to biochemical data," *BMC bioinformatics*, vol. 14 Suppl 7, p. S13, 4 2013.
- [21]. Caelli T and Kosinov S, "An eigenspace projection clustering method for inexact graph matching," *IEEE transactions on pattern analysis and machine intelligence*, vol. 26, no. 4, pp. 515–519, 4 2004. [PubMed: 15382655]
- [22]. Dutta A, Lladós J, Bunke H, and Pal U, "Product graph-based higher order contextual similarities for inexact subgraph matching," 2 2017.
- [23]. Carletti V, "Exact and inexact methods for graph similarity in structural pattern recognition PhD thesis of vincenzo carletti," Ph.D. dissertation, Université de Caen; Università degli studi di Salerno, 2016.
- [24]. Ghahraman DE, Wong AKC, and Au T, "Graph optimal monomorphism algorithms," *IEEE transactions on systems, man, and cybernetics*, vol. 10, no. 4, pp. 181–188, 4 1980.
- [25]. Zaslavskiy M, Bach F, and Vert J-P, "A path following algorithm for the graph matching problem," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 12, pp. 2227–2242, 2009. [PubMed: 19834143]
- [26]. Fiori M, Sprechmann P, Vogelstein J, Musé P, and Sapiro G, "Robust multimodal graph matching: Sparse coding meets graph matching," *Advances in Neural Information Processing Systems*, pp. 127–135, 2013.
- [27]. Egozi A, Keller Y, and Guterman H, "A probabilistic approach to spectral graph matching," *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 1, pp. 18–27, 1 2013. [PubMed: 22350163]
- [28]. Aflalo Y, Bronstein A, and Kimmel R, "On convex relaxation of graph isomorphism," *Proceedings of the National Academy of Sciences*, vol. 112, no. 10, pp. 2942–2947, 2015.

- [29]. Lyzinski V, Fishkind DE, and Priebe CE, “Seeded graph matching for correlated Erdos-Renyi graphs,” *Journal of Machine Learning Research*, vol. 15, pp. 3513–3540, 2014.
- [30]. Fang F, Sussman DL, and Lyzinski V, “Tractable graph matching via soft seeding,” *arXiv preprint arXiv:1807.09299*, 2018.
- [31]. Hoff PD, “Multiplicative latent factor models for description and prediction of social networks,” *Computational & Mathematical Organization Theory*, vol. 15, no. 4, pp. 261–272, 2009.
- [32]. Hoff PD, Raftery AE, and Handcock MS, “Latent space approaches to social network analysis,” *Journal of the American Statistical Association*, vol. 97, no. 460, pp. 1090–1098, 2002.
- [33]. Durante D, Dunson DB, and Vogelstein JT, “Nonparametric bayes modeling of populations of networks,” *Journal of the American Statistical Association*, pp. 1–15, 8 2016.
- [34]. Chatterjee S, “Matrix estimation by universal singular value thresholding,” *The Annals of Statistics*, vol. 43, no. 1, pp. 177–214, 2014.
- [35]. Davenport MA, Plan Y, van den Berg E, and Wootters M, “1-bit matrix completion,” *Information and Inference: A Journal of the IMA*, vol. 3, no. 3, pp. 189–223, 9 2014.
- [36]. Bunke H, “On a relation between graph edit distance and maximum common subgraph,” *Pattern recognition letters*, vol. 18, no. 8, pp. 689–694, 8 1997.
- [37]. Lyzinski V and Sussman DL, “Matchability of heterogeneous networks pairs,” *arXiv preprint, arXiv:1705.02294*, 2017.
- [38]. Lovász L, *Large networks and graph limits*. American Mathematical Society Providence, 2012, vol. 60.
- [39]. Alon N, Yuster R, and Zwick U, “Color-coding,” *Journal of the ACM (JACM)*, vol. 42, no. 4, pp. 844–856, 1995.
- [40]. Munkres J, “Algorithms for the assignment and transportation problems,” *Journal of the Society for Industrial and Applied Mathematics*, vol. 5, no. 1, pp. 32–38, 1957.
- [41]. Young S and Scheinerman E, “Random dot product graph models for social networks,” in *Proceedings of the 5th international conference on algorithms and models for the web-graph*, 2007, pp. 138–149.
- [42]. Nickel CLM, “Random dot product graphs: A model for social networks,” *Ph.D. dissertation*, Johns Hopkins University, 2006.
- [43]. Airoldi EM, Blei DM, Fienberg SE, and Xing EP, “Mixed membership stochastic blockmodels,” *Journal of machine learning research: JMLR*, vol. 9, pp. 1981–2014, 9 2008. [PubMed: 21701698]
- [44]. Karrer B and Newman MEJ, “Stochastic blockmodels and community structure in networks,” *Physical Review E*, vol. 83, no. 1, p. 016107, 2011.
- [45]. Athreya A, Fishkind DE, Levin K, Lyzinski V, Park Y, Qin Y, Sussman DL, Tang M, Vogelstein JT, and Priebe CE, “Statistical inference on random dot product graphs: a survey,” 9 2017.
- [46]. Eichler K, Li F, Litwin-Kumar A, Park Y, Andrade I, Schneider-Mizell CM, Saumweber T, Huser A, Eschbach C, Gerber B et al., “The complete connectome of a learning and memory centre in an insect brain,” *Nature*, vol. 548, no. 7666, 2017.
- [47]. Kiar G, Bridgeford E, Roncal WG, Chandrashekhkar V, Mhembere D, Burns R, and Vogelstein J, “neurodata/ndmg: Stable ndmg-DWI pipeline release,” 1 2018.
- [48]. Kiar G, Bridgeford E, Roncal WG, Consortium for Reliability and Reproducibility (CoRR), Chandrashekhkar V, Mhembere D, Ryman S, Zuo X-N, Marguiles DS, Cameron Craddock R, Priebe CE, Jung R, Calhoun V, Caffo B, Burns R, Milham MP, and Vogelstein J, “A High-Throughput pipeline identifies robust connectomes but troublesome variability,” 4 2018.
- [49]. Finke G, Burkard RE, and Rendl F, “Quadratic assignment problems,” *North-Holland Mathematics Studies*, vol. 132, pp. 61–82, 1987.
- [50]. Burkard RE, Çela E, Pardalos PM, and Pitsoulis LS, “The quadratic assignment problem,” in *Handbook of Combinatorial Optimization: Volume 1–3*, Du D-Z and Pardalos PM, Eds. Boston, MA: Springer US, 1999, pp. 1713–1809.
- [51]. Fishkind DE, Meng L, Sun A, Priebe CE, and Lyzinski V, “Alignment strength and correlation for graphs,” 8 2018.

(a) We plot the average number of matched vertices, out of 40, which were correctly matched according to the rank of the objective function. The different colors correspond to using either 0, 7, or 15 soft seeds and each panel corresponds to the homogeneous Erdős-Rényi, as described in Section 4.1.

(b) Each point denotes the conditional mean of the objective function given the number of matched vertices. Each point is scaled according to the average number of restarts that achieve the given number of matched vertices. Note that for correlations ≥ 0.8 , there is a large gap in the expected objective function value for when there are ≈ 40 correct matches as compared to < 40 .

Fig. 1:
Results for homogeneous Erdős-Rényi experiments

Fig. 2: Number of correctly matched vertices as a function of the objective function rank for the planted partition model described in Section 4.2. Each line corresponds to the density of the induced subgraph, ranging from $q = 0.25$ to 0.5 .

Fig. 3:

A representative set of latent positions for the random dot product graph model. The red circles correspond to the max-angle sampling and the points in blue correspond to a possible random sampling (these do not overlap for the purposes of visualization).

(a) The number of correctly matched vertices out of 40 according the ranks of the objective function across 100 random restarts. The left panel corresponds to the scenario where the smaller graph is correlated with a subgraph determined by underlying vertex attributes and the right panel corresponds to a random subgraph.

(b) Each point denotes the conditional mean of the objective function given the number of matched vertices. Each point is scaled according the average number of restarts that achieve the given number of matched vertices. Note that despite the fact that there are fewer high accuracy solutions for the random subgraph, in both situations there is a substantial gap between the objective function value for high accuracy solutions and those with low accuracy.

Fig. 4:
Random dot product graph matching results.

Fig. 5:
Connectomes and matching analysis of the *Drosophila* connectomes.

Fig. 6:

These panels show the accuracy for detecting the 100 K-cells in the left hemisphere using the K-cells subgraph for the right hemisphere. The four columns correspond to the 4 different centering options and each curve corresponds to the centering that was used for the second matching. Kernel smoothing with bandwidth 0.02 was used to smooth the average accuracy across the normalized objective function ranks (for the first match). The two rows correspond to using 0 and 10 seeds.

Fig. 7:
 Example matches from human DTMRI graphs.

TABLE 1:

Runtimes in seconds per restart for each setting in terms of the number of seeds and the correlation ρ in the homogeneous Erdős-Rényi setting.

Seeds	runtime in seconds			
	$\rho = 0.7$	$\rho = 0.8$	$\rho = 0.9$	$\rho = 1$
0	1.96 ± 0.565	1.92 ± 0.478	1.9 ± 0.479	1.91 ± 0.48
7	1.85 ± 0.476	1.8 ± 0.466	1.74 ± 0.484	1.58 ± 0.546
15	1.5 ± 0.467	1.03 ± 0.395	0.73 ± 0.22	0.61 ± 0.164

TABLE 2:

The number of K-cells in the left hemisphere detected by the match with the best objective function after the first match. The best among these, highlighted in bold, is 89 which was achieved by first performing a rank-1 centered match and then performing a rank-4 centered match, or a rank-2 centered match in the case of zero seeds.

seeds	scheme	second match				
		none	rank-1	rank-2	rank-3	rank-4
0	rank-1	86		89	88	89
	rank-2	82	82		83	85
	center	77	86	82	80	85
	naive	84	87	83	86	86
10	rank-1	84		85	86	89
	rank-2	76	84		84	81
	center	78	86	82	83	83
	naive	84	88	85	83	86

TABLE 3:

Performance comparisons for the large scale DTMRI examples.

hemisphere/region	# nodes	# edges	best matches				most correct vertices			
			fewest edge errors		# correct matches		# edge errors		# correct matches	
left	20,412	922,793	# correct matches	# edge errors	# seeds	# correct matches	# edge errors	# seeds	# correct matches	# seeds
right	20,401	921,865								
left 7	186	611	9	49	0	56	121	20		
right 7	173	646	0	81	0	65	576	20		
left 20	136	816	1	121	0	78	324	10		
right 20	192	1077	2	144	0	56	256	20		
left 33	70	369	0	36	0	45	289	20		
right 33	81	355	0	49	0	30	121	10		
left 34	206	1528	0	324	0	78	729	10		
right 34	173	1363	0	256	0	95	625	10		