

RESEARCH

Open Access



# Motif-based spectral clustering of weighted directed networks

William G. Underwood<sup>1,2\*</sup> , Andrew Elliott<sup>3,4</sup> and Mihai Cucuringu<sup>1,3\*</sup>

\*Correspondence:

wgu2@princeton.edu;

mihai.cucuringu@stats.ox.ac.uk

<sup>1</sup>Department of Statistics, University of Oxford, 24–29 St. Giles, Oxford, OX1 3LB, UK

<sup>3</sup>The Alan Turing Institute, British Library, 96 Euston Road, London NW1 2DB, UK

Full list of author information is available at the end of the article

## Abstract

Clustering is an essential technique for network analysis, with applications in a diverse range of fields. Although spectral clustering is a popular and effective method, it fails to consider higher-order structure and can perform poorly on directed networks. One approach is to capture and cluster higher-order structures using motif adjacency matrices. However, current formulations fail to take edge weights into account, and thus are somewhat limited when weight is a key component of the network under study. We address these shortcomings by exploring motif-based weighted spectral clustering methods. We present new and computationally useful matrix formulae for motif adjacency matrices on weighted networks, which can be used to construct efficient algorithms for any anchored or non-anchored motif on three nodes. In a very sparse regime, our proposed method can handle graphs with a million nodes and tens of millions of edges. We further use our framework to construct a motif-based approach for clustering bipartite networks.

We provide comprehensive experimental results, demonstrating (i) the scalability of our approach, (ii) advantages of higher-order clustering on synthetic examples, and (iii) the effectiveness of our techniques on a variety of real world data sets; and compare against several techniques from the literature. We conclude that motif-based spectral clustering is a valuable tool for analysis of directed and bipartite weighted networks, which is also scalable and easy to implement.

**Keywords:** Motif, Spectral clustering, Weighted network, Directed network, Community detection, Graph Laplacian, Bipartite network

## 1 Introduction

Networks are ubiquitous in modern society; from the internet and online blogs to protein interactions and human migration, we are surrounded by inherently connected structures (Kolaczyk and Csárdi 2014). The mathematical and statistical analysis of networks is therefore an important area of modern research, with applications in a diverse range of fields, including biology (Albert 2005), chemistry (Jacob and Lapkin 2018), physics (Newman 2008) and sociology (Adamic and Glance 2005).

A common task in network analysis is that of *clustering* (Schaeffer 2007). Network clustering refers to the partitioning of a network into “clusters”, so that nodes in each cluster are similar (in some sense), while nodes in different clusters are dissimilar. For a review

of approaches, see for example (Fortunato 2010) or (Fortunato and Hric 2016). Spectral methods for network clustering have a long and successful history (Cheeger 1969; Donath and Hoffman 1972; Guattery and Miller 1995), and have become increasingly popular in recent years. These techniques exhibit many attractive properties including generality, ease of implementation and scalability (Von Luxburg 2007); in addition to often being amenable to theoretical analysis by using tools from matrix perturbation theory (Stewart and Sun 1990). However, traditional spectral methods have shortcomings, particularly involving their inability to consider higher-order network structures (organizations above the level of individual nodes and edges), which have become of increasing interest in recent years (Rosvall et al. 2014; Benson et al. 2016; Benson et al. 2018); and their insensitivity to edge directions<sup>1</sup> (Cucuringu et al. 2019b). Such weaknesses can lead to unsatisfactory results, especially when considering directed networks. Motif-based spectral methods have proven more effective for clustering directed networks on the basis of higher-order structures (Tsourakakis et al. 2017), with the introduction of the *motif adjacency matrix* (MAM) (Benson et al. 2016). While weights can be important in network clustering (Newman 2004), to the best of our knowledge, these motif-based methods have not been comprehensively investigated on weighted networks. Thus, in this paper, we focus on extending these methods to the family of weighted directed networks.

**Contribution** In this paper, we explore motif-based spectral clustering methods with a focus on addressing these shortcomings by generalizing motif-based spectral methods to *weighted* directed networks. Our main contributions include a collection of new matrix-based formulae for MAMs on weighted directed networks, and a motif-based approach for clustering bipartite networks. We also provide computational analysis of our approaches, demonstrations of scalability and comprehensive experimental results, both from synthetic data (variants of stochastic block models) and from real world network data. Finally, we provide a thoroughly tested, scalable implementation of our proposed matrix-based MAM formulae in both Python and R, which can be found at <https://github.com/wgunderwood/motifcluster>.

**Paper layout** Our paper is organized as follows. In Section 2, we describe our graph-theoretic framework which provides a natural model for real world weighted directed networks and weighted bipartite networks. In Section 3 we develop our methodology, and state and prove new matrix-based formulae for MAMs. We explore the computational complexity of our approaches and demonstrate their scalability on sparse graphs. In Section 4 we explore the performance of our approaches on several synthetic examples. We demonstrate the utility of considering weight and higher-order structure, and compare against non-weighted and non-higher-order methods. In Section 5, we apply our methods to real world data sets, demonstrating that they can uncover interesting divisions in weighted directed networks and in weighted bipartite networks. We compare our performance to standard methods and highlight our ability to avoid misclassification. Finally, in Section 6 we present our conclusions and discuss future work.

---

<sup>1</sup>Although there are some spectral approaches which do consider edge direction e.g. (Rohe et al. 2016; Satuluri and Parthasarathy 2011).

## 2 Framework

In this section, we give notation and definitions for our graph-theoretic framework, with the aim of being able to define our weighted generalizations of motif adjacency matrices in Section 3. The motif adjacency matrix (Benson et al. 2016) is the central object in motif-based spectral clustering, and serves as a similarity matrix for spectral clustering. In an unweighted MAM  $M$ , the entry  $M_{ij}$  is proportional to the number of motifs of a given type that include both of the vertices  $i$  and  $j$ .

**Notation** Graph notation is notoriously inconsistent in the literature; in this work, a *graph* is a triple  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$  where  $\mathcal{V}$  is the *vertex set*,  $\mathcal{E} \subseteq \{(i, j) : i, j \in \mathcal{V}, i \neq j\}$  is the *edge set* and  $W: \mathcal{E} \rightarrow (0, \infty)$  is the *weight map*. A graph  $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$  is a *subgraph* of a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  (write  $\mathcal{G}' \leq \mathcal{G}$ ) if  $\mathcal{V}' \subseteq \mathcal{V}$  and  $\mathcal{E}' \subseteq \mathcal{E}$ . It is an *induced subgraph* (write  $\mathcal{G}' < \mathcal{G}$ ) if further  $\mathcal{E}' = \mathcal{E} \cap (\mathcal{V}' \times \mathcal{V}')$ . A graph  $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$  is *isomorphic* to a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  (write  $\mathcal{G}' \cong \mathcal{G}$ ) if there exists a bijection  $\phi: \mathcal{V}' \rightarrow \mathcal{V}$  with  $(u, v) \in \mathcal{E}' \iff (\phi(u), \phi(v)) \in \mathcal{E}$ . An isomorphism from a graph to itself is called an *automorphism*. Where it is not relevant or understood from the context, we may sometimes omit the weight map  $W$ .

As we are considering directed weighted graphs, it is convenient to consider five indicator matrices that capture the different possible relationships between pairs of nodes, namely the directed indicator matrix  $J$ , the single-edge indicator matrix  $J_s$ , the double-edge indicator matrix  $J_d$ , the missing-edge indicator matrix  $J_0$ , and the vertex-distinct indicator matrix  $J_n$ :

$$\begin{aligned} J_{ij} &:= \mathbb{I}\{(i, j) \in \mathcal{E}\}, \\ (J_s)_{ij} &:= \mathbb{I}\{(i, j) \in \mathcal{E} \text{ and } (j, i) \notin \mathcal{E}\}, \\ (J_d)_{ij} &:= \mathbb{I}\{(i, j) \in \mathcal{E} \text{ and } (j, i) \in \mathcal{E}\}, \\ (J_0)_{ij} &:= \mathbb{I}\{(i, j) \notin \mathcal{E} \text{ and } (j, i) \notin \mathcal{E} \text{ and } i \neq j\}, \\ (J_n)_{ij} &:= \mathbb{I}\{i \neq j\}, \end{aligned}$$

where  $\mathbb{I}$  is an indicator function. Furthermore, we consider the following three weighted adjacency matrices, corresponding to directed edges, single edges, and double edges, respectively:

$$\begin{aligned} G_{ij} &:= W((i, j)) \mathbb{I}\{(i, j) \in \mathcal{E}\}, \\ (G_s)_{ij} &:= W((i, j)) \mathbb{I}\{(i, j) \in \mathcal{E} \text{ and } (j, i) \notin \mathcal{E}\}, \\ (G_d)_{ij} &:= (W((i, j)) + W((j, i))) \mathbb{I}\{(i, j) \in \mathcal{E} \text{ and } (j, i) \in \mathcal{E}\}. \end{aligned}$$

We can extend to the setting of undirected graphs by constructing a new graph, where each undirected edge is replaced by a bi-directional edge.

**Definition 1** (Motifs and anchor sets) *A motif is a pair  $(\mathcal{M}, \mathcal{A})$  where  $\mathcal{M} = (\mathcal{V}_{\mathcal{M}}, \mathcal{E}_{\mathcal{M}})$  is a (weakly) connected graph with  $\mathcal{V}_{\mathcal{M}} = \{1, \dots, m\}$  for some small  $m \geq 2$ , and an anchor set is  $\mathcal{A} \subseteq \mathcal{V}_{\mathcal{M}}$  with  $|\mathcal{A}| \geq 2$ . If  $\mathcal{A} \neq \mathcal{V}_{\mathcal{M}}$ , we say the motif is anchored, and if  $\mathcal{A} = \mathcal{V}_{\mathcal{M}}$  we say it is simple. We say that  $\mathcal{H}$  is a functional instance of  $\mathcal{M}$  in  $\mathcal{G}$  if  $\mathcal{M} \cong \mathcal{H} \leq \mathcal{G}$ , and we say that  $\mathcal{H}$  is a structural instance of  $\mathcal{M}$  in  $\mathcal{G}$  if  $\mathcal{M} \cong \mathcal{H} < \mathcal{G}$ . When an anchor set is*

not given, it is assumed that the motif is simple. Figure 1 shows all the simple motifs (up to isomorphism) on at most three vertices.

Structural instances occur when an exact copy of the motif is present in the graph: i.e. edges present (resp. not present) in the motif are present (resp. not present) in the graph. Functional instances are less restrictive, and occur when the motif is present in a graph, but potentially with extra edges.

Anchor sets (Benson et al. 2016) can be thought of as the set of locations in a motif that we consider important for the motif structure. For example we could consider the 2-path motif  $\mathcal{M}_9$ , and try to cluster together nodes which appear at the start or end of a 2-path, but maybe not in the middle. Then we can define the anchor set as the subset of motif vertices which should not be separated. In  $\mathcal{M}_9$ , the start and end vertices would correspond to  $\mathcal{A} = \{1, 3\}$  (Fig. 1). Anchor sets are crucial for defining the collider and expander motifs given in Section 3.2.

Finally, we require one additional definition before we can state our generalization of MAMs to weighted networks.


**Definition 2** (Anchored pairs) *Let  $\mathcal{G}$  be a graph and  $(\mathcal{M}, \mathcal{A})$  a motif. Suppose  $\mathcal{H}$  is an instance of  $\mathcal{M}$  in  $\mathcal{G}$ . Define the anchored pairs of the instance  $\mathcal{H}$  as*

$$\mathcal{A}(\mathcal{H}) := \{\{\phi(i), \phi(j)\} : i, j \in \mathcal{A}, i \neq j, \phi \text{ is an isomorphism from } \mathcal{M} \text{ to } \mathcal{H}\}.$$

*This is the set of pairs of vertices for which both vertices lie in the image of the motif's anchor set, under some isomorphism from the motif to the instance.*

### 3 Methodology

In this section we detail our methods for motif-based spectral clustering of weighted directed networks. Firstly we define our weighted generalizations of motif adjacency matrices (MAMs) (Definition 3.1). We further provide computationally useful formulae for weighted MAMs (Proposition 3.1), and discuss their applications to clustering



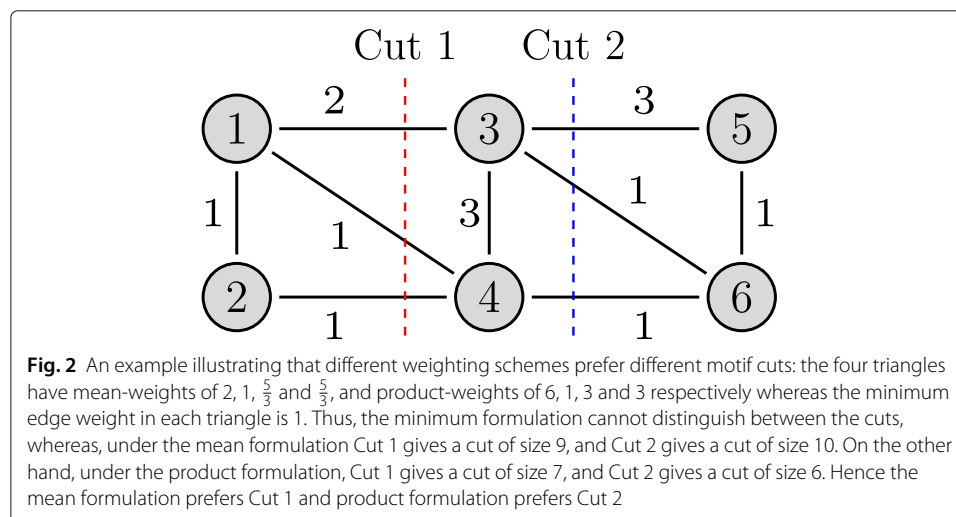
(Section 3.3). Finally we present a complexity analysis of our method for computing weighted MAMs (Proposition 3.3), and empirically demonstrate the scalability of our approach (Section 3.4).

### 3.1 Weighted motif adjacency matrices

Generalizing unweighted measures to weighted networks is non-trivial and application-dependent, and there are typically many possible valid choices. For example, see the approaches for motif weighting proposed by (Onnela et al. 2005) and (Benson et al. 2018). It is often helpful to first consider the generalization of the measure to multi-edges (Newman 2004). We can then view positive integer-weighted edges as multi-edges, and extend to positive real-weighted edges in a natural way.

Thus, we begin by considering the weight to place on a motif which contains multi-edges. A first option might be to consider the minimum number of multi-edges lying on any edge of the motif. This would capture the number of fully edge-disjoint instances of the motif. Another option would be to count the number of unique instances of the motif (possibly counting individual edges more than once), giving the total weight as the product of the number of multi-edges between each pair of nodes in the instance. Alternatively, as a compromise between these two schemes, we might consider the number of distinct motifs present, were the multi-edges distributed evenly among all the edges in the motif (allowing fractional edges for simplicity). This gives the arithmetic mean weighting scheme which is used in Onnela et al. (2005); Benson et al. (2018); Mora et al. (2018) and Simmons et al. (2019).

As an illustrative example of how this choice can affect clustering, Fig. 2 shows how different motif weighting schemes prefer to divide a network in different ways. Here, we treat the undirected edges as bi-directional edges and consider the fully-connected triangle motif  $\mathcal{M}_4$ . The “minimum” approach has no preference between Cut 1 and Cut 2, with all motifs having a weight of 1, and in this example is equivalent to the simple unweighted case. However the “mean” weighting approach (which, up to a scaling factor, is equivalent to summing the edge weights) prefers Cut 1 (cutting motifs of mean-weight 2 and 1 rather



than  $\frac{5}{3}$  and  $\frac{5}{3}$ ), while the “product” approach prefers Cut 2 (cutting motifs of product-weight 3 and 3 rather than 6 and 1). Further, although they ostensibly concern bipartite networks, Section 4.5 exhibits some of the effects of different weighting schemes on the performance of motif-based clustering, and Section 5.3 provides real world motivation for using mean-weighted motifs.

Each of these approaches has merits. The “minimum” approach might be natural when dealing with flow networks, where low-count multi-edges could indicate bottlenecks or points of unreliability. The “product” approach is useful if we want to consider each possible set of edges to be a separate entity, while the “mean” approach is appropriate if we consider the motif to be a single object and wish to count the total number of edges it contains, perhaps as some notion of capacity.

When considering weighted edges however, these approaches have some mathematical differences, particularly in how they handle the presence of large weights. For example, suppose that a graph contains just a few heavy edges. Were two of these heavy edges to appear together in a motif instance, the product formulation would assign a very large total weight to that instance, possibly completely dominating any other instances. On the other hand, using the mean formulation ensures that the total weight of any instance remains on the same scale as the individual edge weights. By taking the minimum edge weight, it may be that the heavy edges do not contribute to any motif weights at all.

With these points in mind, and following (Mora et al. 2018) and (Simmons et al. 2019), we use the “mean-weighted” approach in this paper, defining the weight of a motif instance by its average (mean) edge weight. We acknowledge that other weighting schemes may be more appropriate in some circumstances<sup>2</sup>. For example, (Onnela et al. 2005) introduces several weighting schemes related to the geometric and arithmetic means of the edge weights. Further, these are special cases of the generalized  $p$ -means weighting scheme detailed by Benson et al. (2018).

Now that we have chosen a motif weighting scheme, we can define our weighted generalizations of motif adjacency matrices. We note that while a weighted extension was considered by Wang et al. (2018), in this paper we provide the first thorough exploration, including computational analysis, fast software implementations and comprehensive real world and synthetic experiments.

**Definition 3.1** (Weighted motif adjacency matrices) *Let  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$  be a weighted graph on  $n$  vertices and let  $(\mathcal{M}, \mathcal{A})$  be a motif. The functional and structural (mean-weighted) motif adjacency matrices (MAMs) of size  $n \times n$  of  $(\mathcal{M}, \mathcal{A})$  in  $\mathcal{G}$  are given by*

$$M_{ij}^{func} := \frac{1}{|\mathcal{E}_{\mathcal{M}}|} \sum_{\mathcal{M} \cong \mathcal{H} \leq \mathcal{G}} \mathbb{I}\{\{i, j\} \in \mathcal{A}(\mathcal{H})\} \sum_{e \in \mathcal{E}_{\mathcal{H}}} W(e),$$

$$M_{ij}^{struc} := \frac{1}{|\mathcal{E}_{\mathcal{M}}|} \sum_{\mathcal{M} \cong \mathcal{H} < \mathcal{G}} \mathbb{I}\{\{i, j\} \in \mathcal{A}(\mathcal{H})\} \sum_{e \in \mathcal{E}_{\mathcal{H}}} W(e).$$

When  $W \equiv 1$  and  $\mathcal{M}$  is simple, the (functional or structural) MAM entry  $M_{ij}$  ( $i \neq j$ ) simply counts the (functional or structural) instances of  $\mathcal{M}$  in  $\mathcal{G}$  containing  $i$  and  $j$ . When  $\mathcal{M}$  is not simple,  $M_{ij}$  counts only those instances with anchor sets containing both  $i$  and  $j$ .

<sup>2</sup>For convenience our software package implements three weighting schemes: product, mean and a scheme which ignores the weights.

MAMs are always symmetric, since the only dependency on  $(i, j)$  is via the unordered set  $\{i, j\}$ . In order to state Proposition 3.1, we need one more definition.

**Definition 3.2** (Anchored automorphism classes) *Let  $(\mathcal{M}, \mathcal{A})$  be a motif. Let  $S_{\mathcal{M}}$  be the set of permutations on  $\mathcal{V}_{\mathcal{M}} = \{1, \dots, m\}$  and define the anchor-preserving permutations  $S_{\mathcal{M}, \mathcal{A}} = \{\sigma \in S_{\mathcal{M}} : \{1, m\} \subseteq \sigma(\mathcal{A})\}$ . Let  $\sim$  be the equivalence relation defined on  $S_{\mathcal{M}, \mathcal{A}}$  by:  $\sigma \sim \tau \iff \tau^{-1}\sigma$  is an automorphism of  $\mathcal{M}$ . Finally the anchored automorphism classes are the quotient set  $S_{\mathcal{M}, \mathcal{A}}^{\sim} := S_{\mathcal{M}, \mathcal{A}} / \sim$ .*

The motivation for this definition of anchored automorphism classes is as follows: suppose we are looking for instances of  $(\mathcal{M}, \mathcal{A})$  in  $\mathcal{G}$  which contain nodes  $i$  and  $j$ . We set  $k_1 = i$  and  $k_m = j$  (where  $m$  is the number of nodes in the motif – note we could use any two fixed indices instead of 1 and  $m$  here), and choose some other nodes  $\{k_2, \dots, k_{m-1}\}$ . We want to find all mappings of vertices in  $\mathcal{M}$  to our chosen vertices in  $\mathcal{G}$  by  $u \mapsto k_{\sigma u}$  such that (i) vertices in  $\mathcal{A}$  are mapped to  $i$  and  $j$ , and (ii) mappings which correspond to the same instance are not counted more than once. These conditions are precisely the same as requiring  $\sigma \in S_{\mathcal{M}, \mathcal{A}}^{\sim}$ .

**Proposition 3.1** (MAM formula) *Let  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$  be a graph with vertex set  $\mathcal{V} = \{1, \dots, n\}$  and let  $(\mathcal{M}, \mathcal{A})$  be a motif on  $m$  vertices. For any  $i, j \in \mathcal{V}$  and with  $k_1 = i, k_m = j$ , the functional and structural MAMs of  $(\mathcal{M}, \mathcal{A})$  in  $\mathcal{G}$  are given by*

$$M_{ij}^{func} = \frac{1}{|\mathcal{E}_{\mathcal{M}}|} \sum_{\sigma \in S_{\mathcal{M}, \mathcal{A}}^{\sim}} \sum_{\{k_2, \dots, k_{m-1}\} \subseteq \mathcal{V}} J_{\mathbf{k}, \sigma}^{func} G_{\mathbf{k}, \sigma}^{func}, \tag{1}$$

$$M_{ij}^{struc} = \frac{1}{|\mathcal{E}_{\mathcal{M}}|} \sum_{\sigma \in S_{\mathcal{M}, \mathcal{A}}^{\sim}} \sum_{\{k_2, \dots, k_{m-1}\} \subseteq \mathcal{V}} J_{\mathbf{k}, \sigma}^{struc} G_{\mathbf{k}, \sigma}^{struc}, \tag{2}$$

where  $J_{\mathbf{k}, \sigma}^{func}$  is equal to one (and zero otherwise) if  $\mathcal{M}$  appears as a functional (likewise for structural) instance on the  $m$ -tuple of distinct vertices  $\mathbf{k} = (k_1, \dots, k_m)$  in  $\mathcal{G}$ , under the mapping  $u \mapsto k_{\sigma u}$ ; and in that case,  $G_{\mathbf{k}, \sigma}^{func}$  is the average edge weight of that instance:

$$\begin{aligned} J_{\mathbf{k}, \sigma}^{func} &:= \prod_{\mathcal{E}_{\mathcal{M}}^0} (J_n)_{k_{\sigma u}, k_{\sigma v}} \prod_{\mathcal{E}_{\mathcal{M}}^s} J_{k_{\sigma u}, k_{\sigma v}} \prod_{\mathcal{E}_{\mathcal{M}}^d} (J_d)_{k_{\sigma u}, k_{\sigma v}}, \\ G_{\mathbf{k}, \sigma}^{func} &:= \sum_{\mathcal{E}_{\mathcal{M}}^s} G_{k_{\sigma u}, k_{\sigma v}} + \sum_{\mathcal{E}_{\mathcal{M}}^d} (G_d)_{k_{\sigma u}, k_{\sigma v}}, \\ J_{\mathbf{k}, \sigma}^{struc} &:= \prod_{\mathcal{E}_{\mathcal{M}}^0} (J_0)_{k_{\sigma u}, k_{\sigma v}} \prod_{\mathcal{E}_{\mathcal{M}}^s} (J_s)_{k_{\sigma u}, k_{\sigma v}} \prod_{\mathcal{E}_{\mathcal{M}}^d} (J_d)_{k_{\sigma u}, k_{\sigma v}}, \\ G_{\mathbf{k}, \sigma}^{struc} &:= \sum_{\mathcal{E}_{\mathcal{M}}^s} (G_s)_{k_{\sigma u}, k_{\sigma v}} + \sum_{\mathcal{E}_{\mathcal{M}}^d} (G_d)_{k_{\sigma u}, k_{\sigma v}}, \end{aligned}$$

and the summations and products are over the missing edges, single edges and double edges of  $\mathcal{M}$  as follows:

$$\begin{aligned} \mathcal{E}_{\mathcal{M}}^0 &:= \{(u, v) : 1 \leq u < v \leq m : (u, v) \notin \mathcal{E}_{\mathcal{M}}, (v, u) \notin \mathcal{E}_{\mathcal{M}}\}, \\ \mathcal{E}_{\mathcal{M}}^s &:= \{(u, v) : 1 \leq u < v \leq m : (u, v) \in \mathcal{E}_{\mathcal{M}}, (v, u) \notin \mathcal{E}_{\mathcal{M}}\}, \\ \mathcal{E}_{\mathcal{M}}^d &:= \{(u, v) : 1 \leq u < v \leq m : (u, v) \in \mathcal{E}_{\mathcal{M}}, (v, u) \in \mathcal{E}_{\mathcal{M}}\}. \end{aligned}$$

*Proof* See Proof B.1. □

### 3.2 Motif adjacency matrices for bipartite graphs

We also extend our formulation to weighted bipartite networks, by considering certain 3-node anchored motifs. These are used to create separate similarity matrices for each part of a bipartite graph.

**Definition 3.3** A bipartite graph is a directed graph where the vertices can be partitioned as  $\mathcal{V} = \mathcal{S} \sqcup \mathcal{D}$ , such that every edge starts in  $\mathcal{S}$  and ends in  $\mathcal{D}$ . We refer to  $\mathcal{S}$  as the source vertices and to  $\mathcal{D}$  as the destination vertices.

Our method for clustering bipartite graphs uses two *anchored* motifs; the *collider* and the *expander* (Fig. 3). For both motifs the anchor set is  $\mathcal{A} = \{1, 3\}$ . These motifs are useful for bipartite clustering because when restricted to the source or destination vertices, their MAMs are the adjacency matrices of the weighted projections (Chessa et al. 2014) of the graph  $\mathcal{G}$  (Proposition 3.2). In particular they can be used as similarity matrices for the source and destination vertices respectively. Note that if a bipartite graph is connected, then so are the projections onto its source or destination vertices.

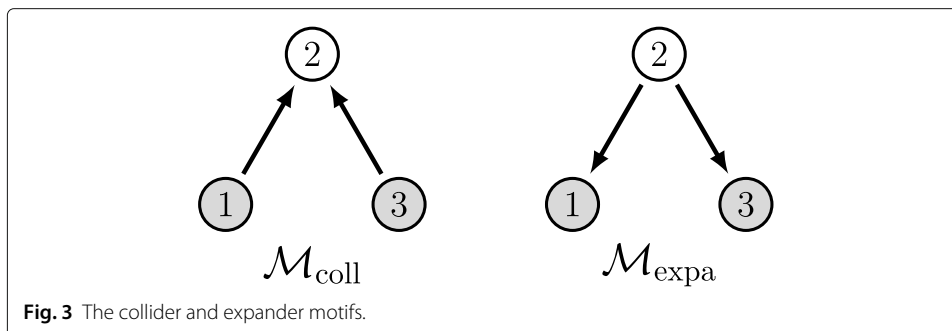
**Proposition 3.2** (Colliders and expanders in bipartite graphs) Let  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$  be a directed bipartite graph. Let  $M_{coll}$  and  $M_{expa}$  be the structural or functional MAMs of  $\mathcal{M}_{coll}$  and  $\mathcal{M}_{expa}$  respectively in  $\mathcal{G}$ . Then

$$(M_{coll})_{ij} = \mathbb{I}\{i \neq j\} \sum_{\substack{k \in \mathcal{D} \\ (i,k), (j,k) \in \mathcal{E}}} \frac{1}{2} [W((i,k)) + W((j,k))], \quad (3)$$

$$(M_{expa})_{ij} = \mathbb{I}\{i \neq j\} \sum_{\substack{k \in \mathcal{S} \\ (k,i), (k,j) \in \mathcal{E}}} \frac{1}{2} [W((k,i)) + W((k,j))]. \quad (4)$$

*Proof* See Proof B.2. □

We note that there are other options available for constructing projections of weighted bipartite graphs, such as the approach in (Stram et al. 2017), which, similarly to our framework, uses sums of edge weights over shared neighbors. Another relevant line of work is that of (Zha et al. 2001), who proposed a certain minimization problem on the bipartite graph, showing that an approximation solution could be obtained via a partial singular value decomposition of a suitably scaled edge weight matrix.





### 3.3 Clustering the motif adjacency matrix

A motif adjacency matrix can be construed as a general pairwise similarity measure, and as such, can be analyzed directly as a new weighted *undirected* graph. Following (Benson et al. 2016), one of the most interesting applications is identifying higher-order clusters with spectral methods, leveraging the fact that the similarity is based on motifs.

To extract clusters, we use a standard approach from the literature on spectral clustering, based on the spectrum of the random-walk Laplacian, followed by  $k$ -means++ (see Appendix C). To this end, we need to define two parameters:  $l$  is the number of random-walk Laplacian eigenvectors to use, and  $k$  is the number of clusters for  $k$ -means++ (see Algorithm 1). We detail our approach for general directed weighted networks in Algorithm 2, and for weighted bipartite networks in Algorithm 3.

#### 3.3.1 Cluster evaluation

When ground-truth clustering is available, we compare it to our recovered clustering using the *Adjusted Rand Index* (ARI) (Hubert and Arabie 1985). The ARI between two clusterings has expected value 0 under random cluster assignment, and maximum value 1 denoting perfect agreement between the clusterings. A larger ARI indicates a more similar clustering, and hence closer to the ground truth.

#### 3.3.2 Connected components

In order for spectral clustering to produce nontrivial clusters from a graph, it is necessary to restrict the graph to its largest connected component. When forming MAMs, even if the original graph is (weakly or strongly) connected, there is no guarantee that the (symmetric) MAM is connected too. Hence we restrict the MAM to its largest connected component before spectral clustering is applied. While this may initially seem to be a flaw with motif-based spectral clustering (since some vertices may not be assigned to any cluster), in fact it can be useful: we only attempt to cluster vertices which are in some sense “well connected” to the rest of the graph. This can result in fewer misclassifications than with traditional spectral clustering, as seen in Section 5.2, where we also investigate MAM regularization as an alternative strategy for dealing with disconnected MAMs.

#### 3.3.3 Motif choice

The selection of a motif is essentially equivalent to selecting a similarity measure between nodes in the network, as the  $(i, j)$ th entry in the MAM corresponds to the similarity between nodes  $i$  and  $j$  used for our clustering procedure. This selection can be important, as we will see in Section 4.4, where the choice of motif can have a significant impact on the clusters obtained. For motif selection, considerations for weighted networks are largely similar to those for unweighted networks. Thus much of this discussion will mirror that of (Benson et al. 2016), although we will highlight the areas where weight may cause deviations.

The first criterion for motif selection is the application domain. This may involve choosing motifs which are related to specific features of the application domain, essentially specializing the similarity measure to the task at hand. For example, in several fields, some motifs may be more relevant or have very different properties; e.g. the feed forward structure in biological networks (Mangan and Alon 2003), various triadic structures in sociology (Wasserman et al. 1994), or the motif  $\mathcal{M}_5$  in food web networks (Benson et al. 2016). For weighted networks the procedure is very similar, but care must

be taken to simultaneously select a motif and a weighting scheme (Section 3.1), in order to capture the similarity measure of interest.

Again following (Benson et al. 2016), in the case where there is no application domain guidance, more principled methods for motif choice are also available. For example, sweep profiles (Shi and Malik 2000) or their motif-coherence counterparts (Benson et al. 2016) can be used to identify motifs which give more clearly distinguished clusters. In the weighted case, these methods require weighted generalizations of the appropriate quantities. We give an example of the application of weighted generalizations of these sweep profiles to real world data in Section 5.2. Finally, one could consider every motif (or a subset of motifs) and then explore each in turn. We demonstrate this approach in our migration experiment in Section 5.1, where we plot the geographic spread of our clusters and then validate on a subset of motifs by considering the cut imbalance ratio scores.

### 3.3.4 Functional vs. structural MAMs

There is also a choice of whether to use functional or structural MAMs for motif-based clustering, and their different properties make them suitable for different circumstances. Firstly, note that  $0 \leq M_{ij}^{\text{struc}} \leq M_{ij}^{\text{func}}$  for all  $i, j \in \mathcal{V}$ . This implies that the largest connected component of  $M^{\text{func}}$  is always at least as large as that of  $M^{\text{struc}}$ , meaning that sometimes more vertices can be assigned to a cluster by using functional MAMs. However, structural MAMs are more discerning about finding motifs, since they require both “existence” and “non-existence” of edges.

## 3.4 Computational analysis

For motifs on at most three vertices, we provide two fast and potentially parallelizable matrix-based procedures for computing weighted MAMs: one for dense regimes and one for sparse regimes. In this section, we explore and demonstrate the scalability of these two approaches.

### 3.4.1 Dense approach

First, Proposition 3.3 bounds the number of matrix operations required to compute a weighted MAM, for a motif on at most three vertices, using our dense formulation. In practice, additional symmetries of the motif often allow computation with even fewer matrix operations (Appendix A).

**Proposition 3.3** (Complexity of MAM formula) *Let  $\mathcal{G}$  be a (weighted, directed) graph on  $n$  vertices, and suppose that the  $n \times n$  directed adjacency matrix  $G$  of  $\mathcal{G}$  is known. Then, computing the adjacency and indicator matrices and calculating a MAM, for a motif on at most three vertices, using Eqs. (1) and (2) in Proposition 3.1, involves at most 18 matrix multiplications, 22 entry-wise multiplications and 21 additions of  $n \times n$  matrices.*

*Proof* See Proof B.3. □

Thus as we have a fixed bound on the number of each operation for any motif, and as each operation is performed on a matrix of the same size, our approach scales with the largest complexity of these operations. In dense matrices, element-wise products and additions are  $\mathcal{O}(n^2)$  and naive matrix multiplication is  $\mathcal{O}(n^3)$ . Therefore, in a naive implementation for a graph on  $n$  vertices, the overall complexity of our dense approach is

$\mathcal{O}(n^3)$ , and the memory requirement is  $\mathcal{O}(n^2)$ . We note that somewhat faster algorithms are available for multiplication of large dense matrices, such as the  $\mathcal{O}(n^{2.81})$  algorithm given by Strassen (1969).

A list of functional MAM formulae for our dense approach, for all simple motifs on at most three vertices, as well as for the collider and expander motifs (used in Section 3.2), is given in Table 5 in Appendix 6. These formulae are generalizations of those stated in Table S6 in the supplementary materials of (Benson et al. 2016) (a list of structural MAMs for triangular motifs in unweighted graphs). Note that the functional MAM formula for the two-vertex motif  $\mathcal{M}_s$  yields the symmetrized adjacency matrix  $M = G + G^\top$ , which can be used for traditional spectral clustering (Appendix C1), as in (Meilă and Pentney 2007).

### 3.4.2 Sparse approach

Real world networks are often sparse (Leskovec and Krevl 2014), and operations with sparse matrices are significantly faster: in sparse matrices with  $b$  non-zero entries, element-wise products and additions are  $\mathcal{O}(b)$ , and matrix multiplications are  $\mathcal{O}(bn)$ . Therefore, we give a slightly different approach with a better computational running time (Section 3.4.3) on sparse graphs, which can extend to graph sizes that are infeasible for our dense approach.

For motifs on at most three vertices, the formulae in Proposition 3.1 are simply sums of terms of the form  $A \circ (BC)$ , where  $A, B$  and  $C$  are adjacency or indicator matrices, and  $\circ$  represents the entry-wise (Hadamard) product. If the directed graph has  $n$  vertices and  $b$  edges, then most of the adjacency and indicator matrices have at most  $b$  non-zero entries. For these matrices, we can compute the (possibly dense) matrix  $BC$  in  $\mathcal{O}(bn)$  time, and then the matrix  $A \circ (BC)$  in  $\mathcal{O}(n^2)$  time. Summing them together is also done in  $\mathcal{O}(n^2)$  time.<sup>3</sup>

However, when considering motifs with missing edges, we use the *dense* indicator matrices  $J_0$  and  $J_n$ . This is problematic as now the term  $BC$  can contain a dense matrix, which can be an issue both for computational reasons<sup>4</sup> and more importantly, for memory requirements. To address this, we rewrite these two matrices as

$$\begin{aligned} J_n &= \mathbf{1} - I, \\ J_0 &= \mathbf{1} - (I + J_s + J_s^\top + J_d) \\ &= \mathbf{1} - (I + \tilde{J}), \end{aligned}$$

where  $\mathbf{1}$  is a matrix of 1s, and expand out the resulting formulae. Element-wise and matrix products with  $\mathbf{1}$  and  $I$  are at most  $\mathcal{O}(n^2)$  so are computationally simple, and do not require generating dense matrices. This allows scaling simply using standard linear algebra libraries.

For sparse graphs the key limitation of this approach is in fact not CPU, but memory (RAM) as while  $B$  and  $C$  might be sparse,  $BC$  may not be. The worst case density of  $BC$  for  $B$  and  $C$  indicator matrices of a connected graph is  $n^2$  (e.g. a star graph).<sup>5</sup> However,

<sup>3</sup>By exploiting sparsity the sums and element-wise products could be faster than  $\mathcal{O}(n^2)$ , but as the complexity is dominated by the matrix multiplication, we do not explore this further.

<sup>4</sup>Although there are sparse-dense matrix multiplication algorithms that address this.

<sup>5</sup>See <https://math.stackexchange.com/questions/1042096/bounds-of-sparse-matrix-multiplication>.

in practice the density is often substantially lower than this, allowing our approaches to scale to very large graphs (Section 3.4.3).

### 3.4.3 Empirical computational speed

In this section, we demonstrate the scalability of our two approaches to computing MAMs. We showcase the running times of four different random graph ensembles. For the first two ensembles, we use the directed Erdős-Rényi (ER) model (Erdős et al. 1959) in which each directed edge between  $n$  nodes exists independently with probability  $p$ . We consider two regimes, a very sparse regime ( $p = \frac{10}{n}$ ) and a less sparse regime ( $p = \frac{100}{n}$ ). For the last two ensembles, we use the undirected Barabási-Albert (BA) model (Barabási and Albert 1999). In this model, the graph is initialized with  $m$  nodes, and  $n - m$  nodes are then added sequentially, with each node connecting to  $m$  previously placed nodes, and with probability of connection proportional to the current degree of each previously placed node, resulting in a graph with a skewed (power-law) degree distribution. We again consider two regimes, a very sparse regime ( $m = 10$ ) and a less sparse regime ( $m = 100$ ). In each of these four ensembles, the expected number of edges in the graph scales as  $\mathcal{O}(n)$ .

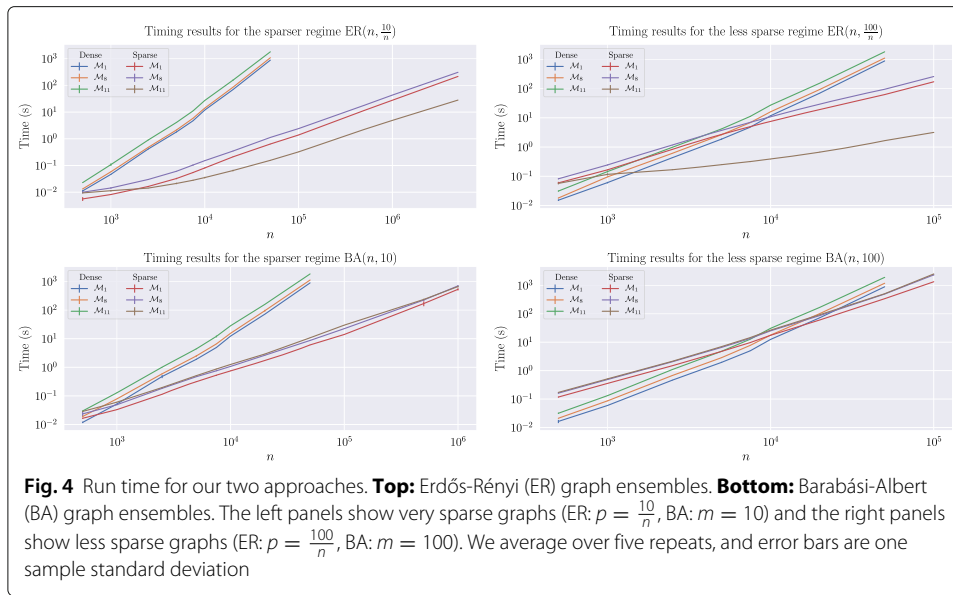
We measure the performance of computing a functional MAM for a representative sample of motifs: one triangle motif  $\mathcal{M}_1$ , the 2-star  $\mathcal{M}_8$ , and a motif with a bi-directional edge,  $\mathcal{M}_{11}$ . We do not include the time taken to generate the graph in either the sparse or the dense matrix form. To make this a fair test, we use an optimized Python environment from the Data Science Virtual Machine image on an E64s v3 Azure virtual machine with 64 vCPUs and 432 GiB of RAM.

The results are summarized in Fig. 4. In the denser graphs (ER:  $p = \frac{100}{n}$ , BA:  $m = 100$ ), we note that for smaller graphs (i.e for  $n$  less than  $\approx 10^4$ ) the dense approach tends to perform as well as, and in many cases exceeds, the performance of the sparse approach, highlighting the advantage of using this approach in denser graphs. For graphs above this size and graphs in the sparser regime (ER:  $p = \frac{10}{n}$ , BA:  $m = 10$ ), the sparse approach outperforms, handling graphs with over a million nodes and tens of millions of edges. In the less sparse regime (ER:  $p = \frac{100}{n}$ , BA:  $m = 100$ ), we consider graphs of size up to  $n = 10^5$  nodes and of order  $10^7$  edges, for which our sparse approach takes less than  $10^2$  seconds for the ER model, and around  $10^3$  seconds for the BA model.

We note that the time required for each method is roughly constant across most of the tested motifs. The exception is the sparse approach for  $\mathcal{M}_{11}$ , which takes substantially less time than the other motifs in both of the ER regimes. We believe this to be related to the bi-directional edge in  $\mathcal{M}_{11}$  (not present in  $\mathcal{M}_1$  or  $\mathcal{M}_8$ ), which is rare in sparse ER graphs, heavily reducing the amount of computation required. This is not seen in the undirected BA graphs, where bi-directional edges are much more common, resulting in a similar level of performance across all tested motifs in this regime. The various scenarios we experimented with highlight the fact that our approach is highly scalable to very large sparse graphs, including those with skewed degree sequences which often arise in real world applications.

## 4 Applications to synthetic data

To validate our method, we consider three synthetic data sets as examples. We selected each example to highlight a different aspect of our approach and to demonstrate the advantages of considering a weighted higher-order clustering measure. Example 1



**Fig. 4** Run time for our two approaches. **Top:** Erdős-Rényi (ER) graph ensembles. **Bottom:** Barabási-Albert (BA) graph ensembles. The left panels show very sparse graphs (ER:  $p = \frac{10}{n}$ , BA:  $m = 10$ ) and the right panels show less sparse graphs (ER:  $p = \frac{100}{n}$ , BA:  $m = 100$ ). We average over five repeats, and error bars are one sample standard deviation

demonstrates the importance of taking the edge weights into account when performing higher-order clustering; Example 2 shows the value of higher-order clustering, demonstrating that clustering using different motifs can yield different insights into the same data; and finally, Example 3 demonstrates the value of our bipartite clustering scheme for detecting structures in weighted bipartite networks.

#### 4.1 Directed stochastic block models

For these tests we use weighted and unweighted *directed stochastic block models* (DSBMs), a broad class of generative models for directed graphs (Nowicki and Snijders 2001). An unweighted DSBM is characterized by a block count  $k$ , a list of block sizes  $(n_i)_{i=1}^k$ , and a connection matrix  $F \in [0, 1]^{k \times k}$ . We define the cumulative block sizes  $n_i^c = \sum_{j=1}^i n_j$ , and the total graph size  $n = n_k^c$ . These are used to construct the group allocations  $g_i = \min\{r : n_r^c \geq i\}$ , and finally a graph  $\mathcal{G}$  is generated with adjacency matrix entries

$$\tilde{G}_{ij} = \text{Ber}(F_{g_i, g_j}) \cdot \mathbb{I}\{i \neq j\}, \quad (5)$$

with all Bernoulli random variables sampled independently. A weighted DSBM is constructed in a similar manner (see for example (Mariadassou et al. 2010) and (Aicher et al. 2014)), but also requires a weight matrix  $\Lambda \in [0, \infty)^{k \times k}$ . In this case, the weighted adjacency matrix entries are generated by the following mixture model

$$G_{ij} = \text{Ber}(F_{g_i, g_j}) \cdot \text{Poi}(\Lambda_{g_i, g_j}) \cdot \mathbb{I}\{i \neq j\}, \quad (6)$$

with all Bernoulli and Poisson random variables sampled independently. Note that the Bernoulli variable now no longer directly corresponds to edge existence, since there is always a non-zero chance of the Poisson variable being zero, which sets  $G_{ij} = 0$  and thus removes the edge. We assume a DSBM is unweighted unless stated otherwise. We evaluate performance using the Adjusted Rand Index (ARI) (Rand 1971), as in Section 3.3.1.

#### 4.2 Bipartite stochastic block models

We define the unweighted *bipartite stochastic block model* (BSBM) with source block count  $k_S$ , destination block count  $k_D$ , source block sizes  $(n_S^i)_{i=1}^{k_S}$ , destination block sizes  $(n_D^i)_{i=1}^{k_D}$ , and bipartite connection matrix  $F_b \in [0, 1]^{k_S \times k_D}$  as the unweighted DSBM with block count  $k = k_S + k_D$ , block sizes  $(n_i)_{i=1}^k = \left( (n_S^i)_{i=1}^{k_S}, (n_D^i)_{i=1}^{k_D} \right)$ , and connection matrix  $F = \begin{pmatrix} 0 & F_b \\ 0 & 0 \end{pmatrix}$ .

A weighted BSBM with bipartite weight matrix  $\Lambda_b \in [0, \infty)^{k_S \times k_D}$  can similarly be constructed by using a weighted DSBM with weight matrix  $\Lambda = \begin{pmatrix} 0 & \Lambda_b \\ 0 & 0 \end{pmatrix}$ . Note that this is a generalization of the model in (Florescu and Perkins 2016).

#### 4.3 Example 1

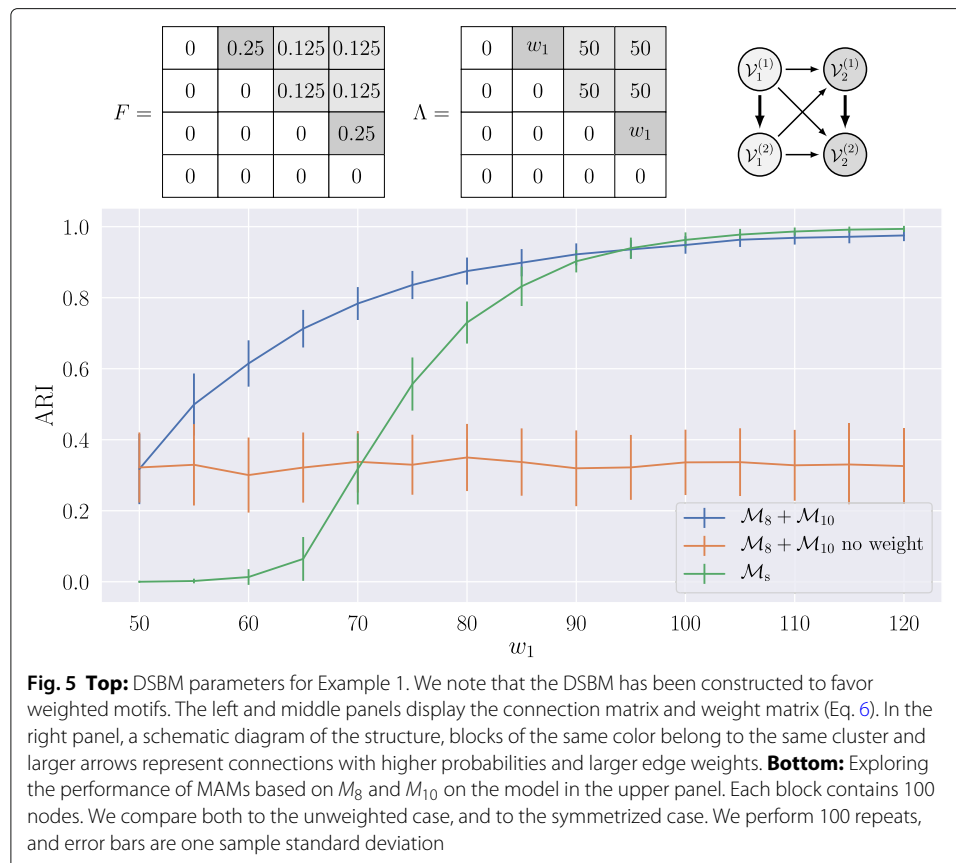
In this example we demonstrate the advantages of taking edge weights into account when detecting higher-order structures. We compare Algorithm 2 with two clusters ( $k = 2$ ) and two eigenvectors ( $l = 2$ ), against two standard approaches. The first comparison is random-walk spectral clustering using a symmetrized weighted matrix, which captures the ability of non-motif-based methods to uncover the underlying structure. This is equivalent to Algorithm 2 with motif  $\mathcal{M}_s$ . Secondly, we compare against our own motif-based approaches, but with all edge weights set to 1, similar to the formulation of (Benson et al. 2016).

To demonstrate the advantages of our approach, we consider an example for which (i) a higher-order structure is present, and (ii) weight is important in the structure. Constructing higher-order structures in a stochastic block model is challenging, as by definition all the edges are independent. Thus the block structure in a DSBM with strongly connected blocks is captured by density rather than by the existence of motifs (although these are correlated).

To this end, we construct clusters that consist of several blocks, and introduce a higher-order structure between blocks of the same cluster. For this example, we use two clusters (Fig. 5 (upper panel)), each one consisting of two blocks each of size 100, for a total of  $n = 400$  nodes. Each cluster  $(\{\mathcal{V}_1^{(1)}, \mathcal{V}_1^{(2)}\}$  and  $\{\mathcal{V}_2^{(1)}, \mathcal{V}_2^{(2)}\})$  consists of two blocks with strong uni-directional (probability  $p = 0.25$ ) connections with potentially large weights (Poisson with mean  $w_1$ ). The two clusters are then linked by weaker inter-block connections with potentially smaller weights (Poisson with mean 50). By design, this model has strong heavy-weighted uni-directional structures, and is thus well captured by  $\mathcal{M}_8$  and  $\mathcal{M}_{10}$ , both of which capture uni-directional structure. Thus, following (Benson et al. 2016) rather than focusing on an individual motif, we use the sum of both MAMs.

The lower panel of Fig. 5 displays the results for functional motifs (structural motifs in Appendix D). As our procedure only clusters the largest connected MAM component, we compute ARI over this component. For  $w_1 = 50$ , all edges have the same expected weight, and thus the performances of weighted and non-weighted motifs are equal. In the mid-range of weights ( $50 < w_1 \leq 90$ ), our approach outperforms both of the others, indicating the advantages of accounting for weights.

Finally, for large weights ( $w_1 > 90$ ), we are (on average) slightly outperformed by  $\mathcal{M}_s$ , the symmetrized weighted adjacency matrix, although the error bars are overlapping. One



possible reason for this is that while using  $\mathcal{M}_8$  and  $\mathcal{M}_{10}$  gives a strong signal, it also introduces noise when motifs with heavy and non-heavy weighted edges span clusters.

The pattern on structural motifs is similar (Fig. 16 in Appendix D), with two key differences: first, the non-weighted motifs have a larger ARI ( $\approx 0.7$  vs.  $\approx 0.3$ ), and second, the weighted motif equals or outperforms  $\mathcal{M}_s$ . We hypothesize the following possible reason for this phenomenon: in this model, motifs on three nodes which span clusters can form triangles, whereas those within clusters cannot (due to the probability-0 connections). Hence versions of the non-triangle motifs  $\mathcal{M}_8$  and  $\mathcal{M}_{10}$  are able to filter out some of the noise described in the previous paragraph, leading to better performance.

Finally, we note that this example has been designed to highlight the advantages of our motifs; and several other structures were considered which do not have this property. When applying this to real world data, it is important to consider the correct motif to use, as discussed in Section 3.3.3.

#### 4.4 Example 2

In the second example, we explore the value of higher-order clustering, by demonstrating that clustering using different motifs can give different albeit potentially equally valuable insights. We showcase this with two experiments. First, we consider the case where, due to the construction of the network, certain motifs are better at detecting the underlying structure in different parameterization regimes. Second, expanding on our first example, we consider a DSBM where, by construction, the network does not have strong

densely connected blocks, and thus different motifs highlight distinct and equally relevant groupings in the network.

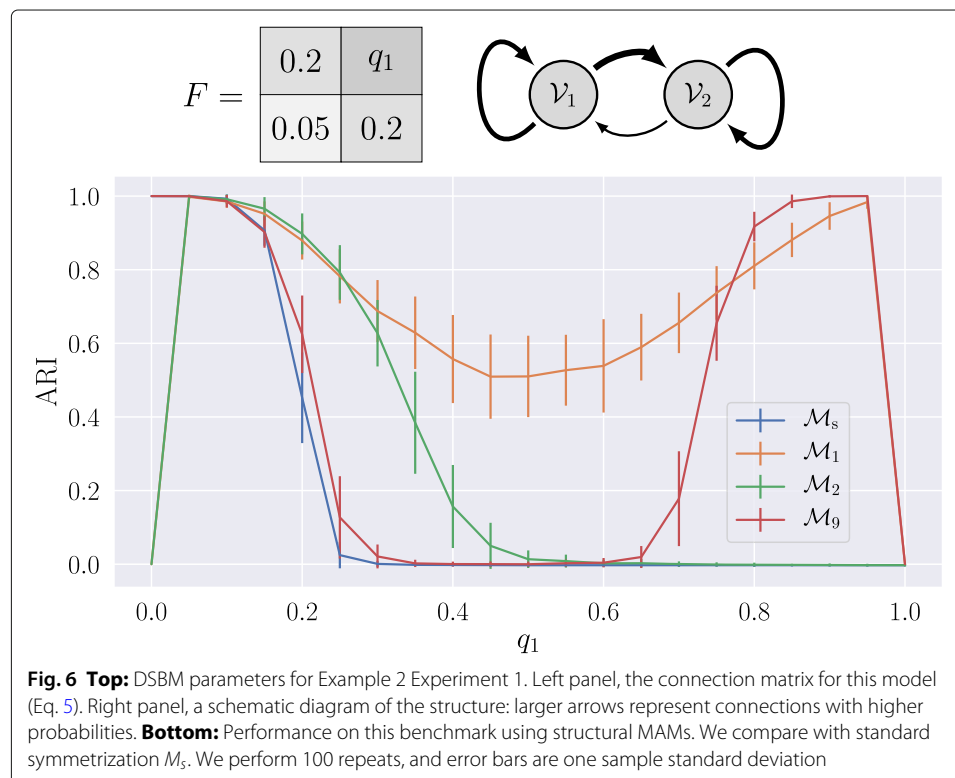
#### 4.4.1 Experiment 1

For the first experiment, we use the DSBM with  $k = 2$ ,  $n_1 = n_2 = 100$  (so  $n = 200$ ) and  $F = \begin{pmatrix} 0.2 & q_1 \\ 0.05 & 0.2 \end{pmatrix}$ , as depicted in the upper panel of Fig. 6.

We test the performance of Algorithm 2 across a few structural motifs with parameters  $k = l = 2$  on this model (functional motifs in Appendix D). It can be seen that the motifs perform well in different regions, with  $\mathcal{M}_1$  outperforming other methods in the regime  $0.3 \leq q_1 \leq 0.7$ , and also performing well outside this range. For large values of  $q_1$ , motif  $\mathcal{M}_9$  performs best. For lower values of  $q_1$ ,  $\mathcal{M}_2$  has a slightly higher average ARI, although the difference is within the standard errors. This altogether demonstrates the importance of selecting the right motif. Furthermore, we compare against  $\mathcal{M}_s$ , the symmetrized weighted adjacency matrix, and each presented motif outperforms this baseline.

We also observe an artifact in this plot: for certain motifs, the performance drops away at  $q_1 = 0$  (triangles) and at  $q_1 = 1$  (all motifs). For these parameter values, the MAM becomes entirely disconnected, with each of the two groups in its own connected component. As our clustering scheme only considers the largest connected component (Section 3.3.2), we then obtain an ARI of 0. In real world graphs, we would recommend investigating all reasonably large connected components or, depending on the application, employing some form of regularization (see Section 5.2).

We note that there is a bi-modality with certain motifs, notable with  $\mathcal{M}_9$  performing well for small and large values of  $q_1$ , a feature not present within the functional motifs





(Fig. 17). We believe this is due to the fact that structural motifs act as a filter: with high values of  $q_1$ , it is difficult to form 2-paths ( $\mathcal{M}_9$ ) between groups without also forming triangular motifs, which would not contribute towards a structural MAM.

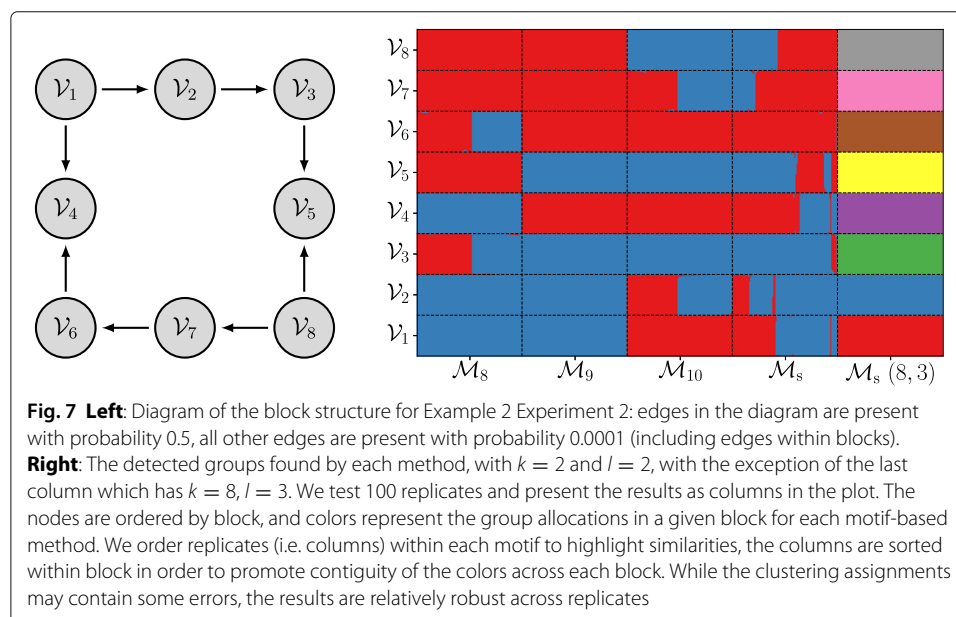
#### 4.4.2 Experiment 2

In this experiment, we demonstrate the ability of MAMs to uncover different structures. We construct a DSBM with 8 groups of size 100 ( $n = 800$ ) with block structure given by the diagram in the left panel of Fig. 7: we place edges which match the block structure with probability 0.5 (i.e.  $F_{ab} = 0.5$  if there is an arrow from  $a$  to  $b$  in the diagram), and 0.0001 otherwise, in order to maintain connectivity.

Following Example 1, in this DSBM we do not have the usual strongly connected groups: in fact, each group has a close to zero probability (0.0001) of within-group connections. Thus, while there are clearly 8 groups with unique connection patterns, there does not exist a “correct” division of the nodes into densely connected groups, but rather, different motifs highlight different structures.

We display the results of using functional motifs with Algorithm 2 with  $k = l = 2$  in Fig. 7 (structural results are similar in Fig. 18). Each column represents the structure uncovered by one of the motifs. For robustness, we perform the experiment on 100 replicates, and place the resulting partitions side by side in each column. For each motif, the columns are sorted within block in order to promote contiguity of the colors across each block.

We observe that (by design) this graph has 3 different clusterings, highlighted by different motifs. First,  $\mathcal{M}_8$  (a 2-star with the edges pointing outwards (Fig. 1)). Each pair of connected blocks is connected by this motif. However, there is a much stronger connection when this motif is part of the higher-order structure. Thus, when we cluster using  $\mathcal{M}_8$ , we obtain two consistent clusters consisting of  $\{\mathcal{V}_1, \mathcal{V}_2, \mathcal{V}_4\}$  and  $\{\mathcal{V}_5, \mathcal{V}_7, \mathcal{V}_8\}$ , each of which is united by this motif at a block level. The remaining blocks without this block



level structure ( $\mathcal{V}_3$  and  $\mathcal{V}_6$ ) are then essentially randomly placed into one of the two clusters, giving the behavior observed in Fig. 7. We observe a similar structure with the other motifs:  $\mathcal{M}_9$  (a 2-path) splits the graph into the two groups characterized by their 2-paths,  $\{\mathcal{V}_1, \mathcal{V}_2, \mathcal{V}_3, \mathcal{V}_5\}$  and  $\{\mathcal{V}_4, \mathcal{V}_6, \mathcal{V}_7, \mathcal{V}_8\}$ ; and finally  $\mathcal{M}_{10}$  (a 2-star with the edges pointing inwards) splits the graph into the two clusters based on this structure, namely,  $\{\mathcal{V}_1, \mathcal{V}_4, \mathcal{V}_6\}$  and  $\{\mathcal{V}_3, \mathcal{V}_5, \mathcal{V}_8\}$  (with a similar behavior to  $\mathcal{M}_8$  for the remaining blocks). Thus, we conclude that our MAMs can obtain very different and equally valid structures in the same graph.

Finally, we compare with  $\mathcal{M}_s$  which is equivalent to clustering with  $G + G^\top$ . Under symmetrization, this graph is a ring of indistinguishable blocks. Therefore, the best possible division is to arbitrarily divide the ring into two roughly equal-sized pieces. Considering Fig. 7, we observe this behavior with many different divisions, and no pair of blocks is consistently placed within the same cluster. For completeness we also compare against  $\mathcal{M}_s$  with  $k = 8$  and  $l = 3$ , which divides the network into the 8 blocks of the underlying DSBM. Although this is a specially constructed example, it highlights the different equally valid structures that can be uncovered, emphasizing the importance of considering the correct motif (see Section 3.3.3) and demonstrating the value of this procedure above standard spectral clustering.

#### 4.5 Example 3

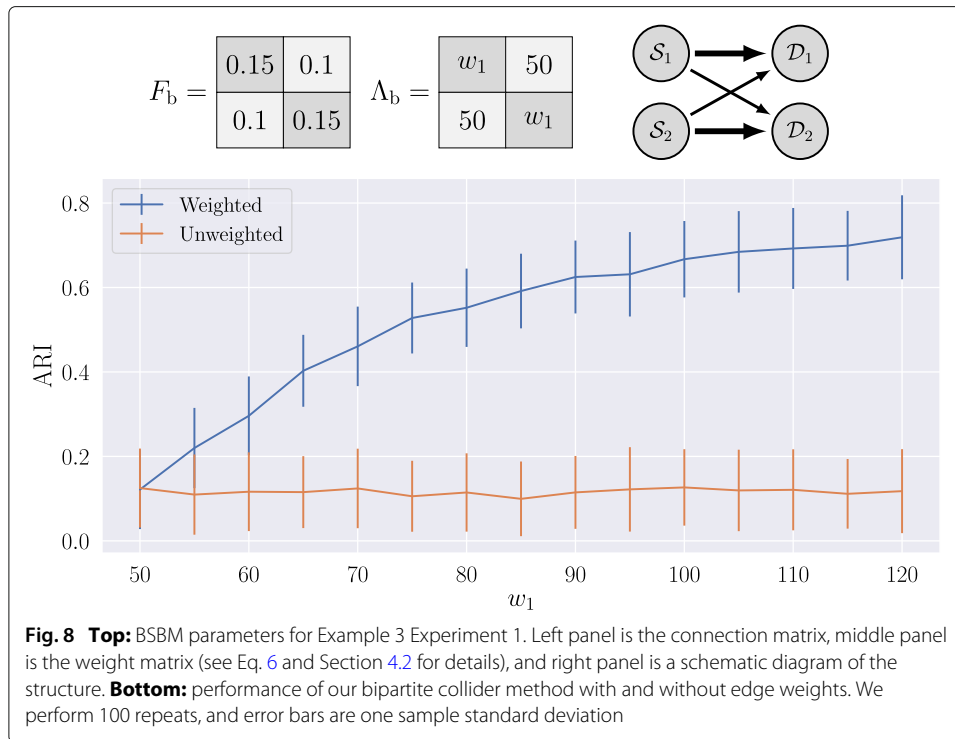
In our final example, we demonstrate clustering the source vertices of bipartite networks, using the collider motif as in Algorithm 3. Clustering the destination vertices with the expander motif is exactly analogous (by simply reversing edge directions), so we do not demonstrate it explicitly. Note also that in bipartite graphs, all instances of the collider or expander motifs are structural, so we need not compare functional and structural MAMs.

We illustrate this with two experiments. In the first one, we show that in certain networks, edge weights are important for our collider method to perform well. In the second experiment, we compare against an alternative method for clustering weighted bipartite networks, showing that under certain conditions, our collider method performs better.

##### 4.5.1 Experiment 1

This example illustrates the advantages of using weights for bipartite clustering. We use a weighted BSBM with block counts  $k_S = k_D = 2$ , block sizes  $n_S^1 = n_S^2 = n_D^1 = n_D^2 = 100$ , bipartite connection matrix  $F_b = \begin{pmatrix} 0.15 & 0.1 \\ 0.1 & 0.15 \end{pmatrix}$ , and bipartite weight matrix  $\Lambda_b = \begin{pmatrix} w_1 & 50 \\ 50 & w_1 \end{pmatrix}$ , where  $w_1$  is a varying parameter. This network has been constructed so that vertices in  $\mathcal{S}_1$  are slightly more likely to connect to vertices in  $\mathcal{D}_1$  than to vertices in  $\mathcal{D}_2$ , and vice versa for  $\mathcal{S}_2$ . This makes the source vertex clusters  $\mathcal{S}_1$  and  $\mathcal{S}_2$  weakly distinguishable when not considering edge weights. However, the edge weights exhibit the same preferences, and this effect becomes stronger as  $w_1$  increases.

We test the performance of Algorithm 3 for clustering the source vertices (using the collider motif), with parameters  $k_S = l_S = 2$  on this model. We compare the results against those obtained by ignoring the edge weights. Fig. 8 shows that the weights in this model allow the structure to be recovered well when the weighting effect is large enough.

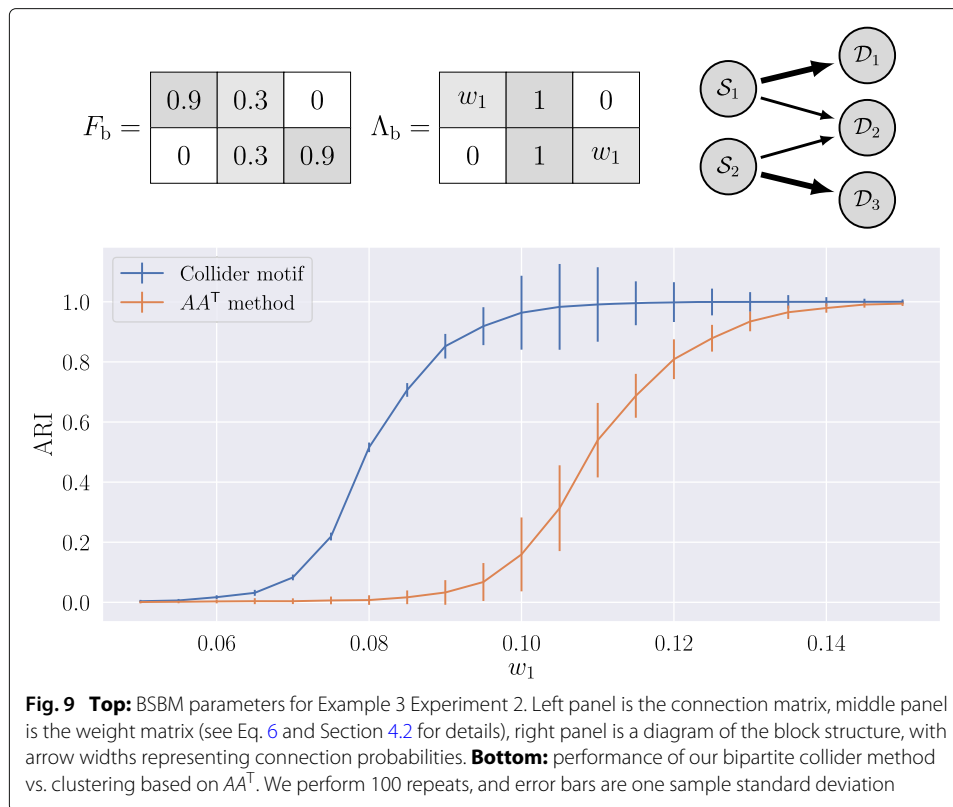


### 4.5.2 Experiment 2

In this final experiment, we show the advantages of using the collider motif over simply clustering using the matrix  $AA^T$ , where  $A \in [0, \infty)^{|S| \times |D|}$  is the weighted bipartite adjacency matrix of the graph (Chessa et al. 2014). We use a weighted BSBM with block counts  $k_S = 2, k_D = 3$ , block sizes  $n_S^1 = n_S^2 = n_D^1 = n_D^2 = n_D^3 = 200$ , bipartite connection matrix  $F_b = \begin{pmatrix} 0.9 & 0.3 & 0 \\ 0 & 0.3 & 0.9 \end{pmatrix}$ , and bipartite weight matrix  $\Lambda_b = \begin{pmatrix} w_1 & 1 & 0 \\ 0 & 1 & w_1 \end{pmatrix}$ , where  $w_1$  is a varying parameter.

This model has been constructed such that vertices in  $S_1$  are more likely to be connected to  $D_1$  than to  $D_2$ , and similarly vertices in  $S_2$  are more likely to be connected to  $D_3$  than to  $D_2$ . However, the weights show a reverse preference (for  $w_1 < 1$ , as in our model), with larger weights assigned to the lower-probability edges. Note that in this regime where weights are small, there is a significant probability of the Poisson weight variables being zero, removing edges which might otherwise have had a high probability of existing.

The results are shown in Fig. 9, which illustrates that our collider method (Algorithm 3 with  $k_S = l_S = 2$ ) is better at distinguishing the source vertex clusters  $S_1$  and  $S_2$  compared to the method based on  $AA^T$ . Although it must be made clear that the model has been specifically constructed to do this, and that this effect occurs only for a limited parameter set, it gives an example of the different behavior observed when using a product-based weight formulation (Section 3.1). We explore this further in real world data in Section 5.3. For our collider method, the expected similarity between two source vertices in the same cluster directly depends on  $w_1$ , while the expected similarity between two source vertices in different clusters is fixed. Hence for large enough values of  $w_1$ , our collider method performs well. However, for the  $AA^T$  method, the expected similarity between two source vertices in the same cluster depends on  $w_1^2$  (since  $AA^T$  contains



squared edge weights), which in this model is small (as  $w_1 < 1$ ). Hence the  $AA^T$  method does not perform so well on this model.

## 5 Applications to real world data

This section details the results of our proposed methodology on a number of directed networks arising from real world data sets. We show that weighted edges and motif-based clustering allow various structures to be uncovered in migration data with the US-MIGRATION network, while the US-POLITICAL-BLOGS network demonstrates how motif-based methods can control misclassification error for weakly-connected vertices, how sweep profiles can be used to select a motif, and how regularization can affect clustering. Finally, we consider the bipartite territory-to-language network, UNICODE-LANGUAGES, and show that when using our bipartite clustering method, mean-weighted motifs can produce more desirable clusters than their unweighted or product-weighted counterparts.

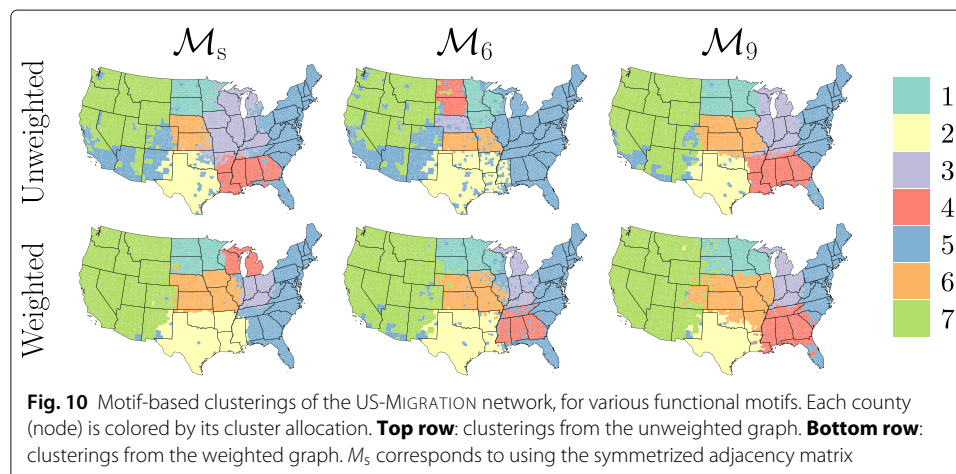
### 5.1 US-MIGRATION network

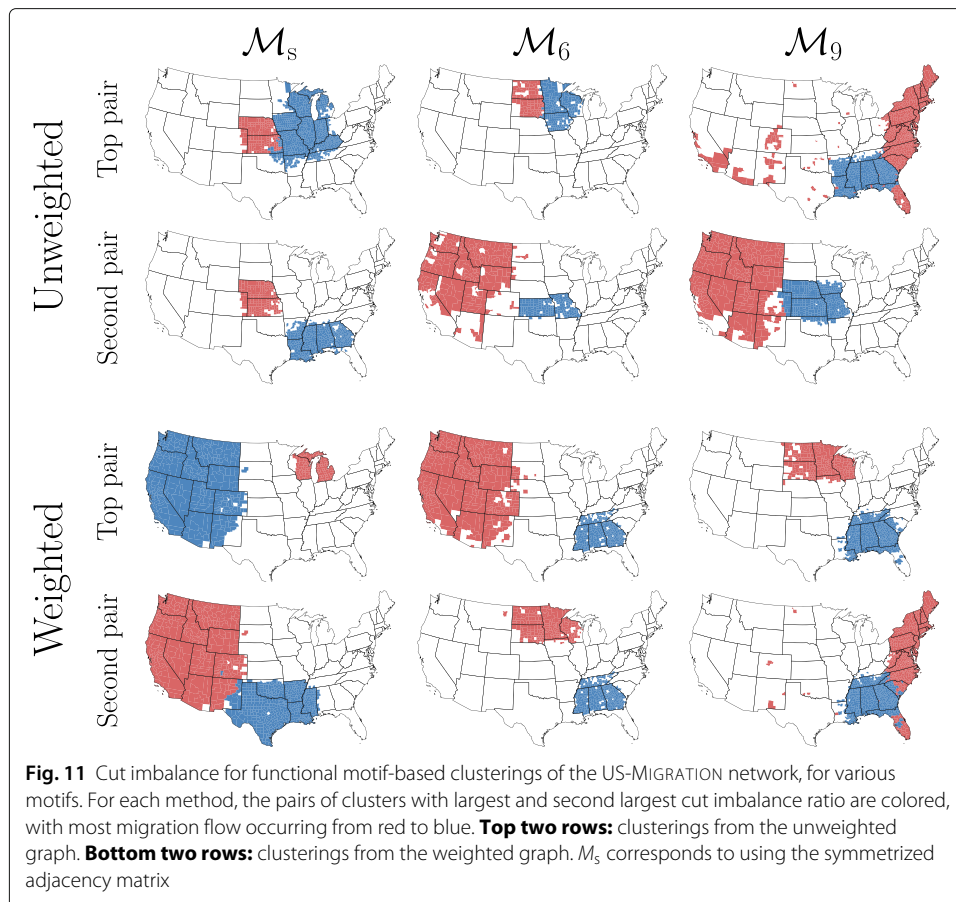
The first data set is the US-MIGRATION network (U.S. Census Bureau 2002), consisting of data collected during the US Census in 2000. Vertices represent the 3075 counties in 49 contiguous states (excluding Alaska and Hawaii, and including the District of Columbia). The 721 432 weighted directed edges represent the number of people migrating from county to county, capped at 10 000 (the 99.9th percentile) to control large entries, as in (Cucuringu et al. 2019b).

We test the performance of Algorithm 2 with three selected motifs:  $\mathcal{M}_s$ ,  $\mathcal{M}_6$  and  $\mathcal{M}_9$  (see Fig. 1).  $\mathcal{M}_s$  gives a spectral clustering method with naive symmetrization.  $\mathcal{M}_6$  represents a pair of counties exchanging migrants, with both also receiving migrants from a third.  $\mathcal{M}_9$  is a path of length two, allowing counties to be deemed similar if they both lie on a heavy-weighted 2-path in the migration network. These motifs give a representative sample, including a triangle and a non-triangle motif, and results with the other motifs (both functional and structural) can be found in Figs. 19 and 20 in Appendix E.

Figure 10 plots maps of the US, with counties colored by the clustering obtained using Algorithm 2 with  $k = l = 7$ . The top row shows clusterings from the unweighted graph, while those in the bottom row are from the weighted graph. The advantage of using weighted graphs is clear, with the clustering showing better spatial coherence (compared with the fragmented blue clusters produced by the unweighted graph). Furthermore, as suggested in Section 3.3.3, the different motifs induce different similarity measures, and thus uncover distinct structures. For example, weighted  $\mathcal{M}_6$  and  $\mathcal{M}_9$  both allocate a cluster to the southern states of Mississippi, Alabama, Georgia and Tennessee, while weighted  $\mathcal{M}_s$  identifies the northern states of Michigan and Wisconsin. Also, weighted  $\mathcal{M}_9$  favors a larger “central” region, which includes significant parts of Colorado, Oklahoma, Arkansas and Illinois.

It is also interesting to investigate the cut imbalance ratio (Appendix C4) (Cucuringu et al. 2019b) associated with these clusters. For each pair of clusters, it measures the imbalance of migration flow from one to the other. Figure 11 indicates, for each method, the first (resp. second) pair of clusters which exhibit the largest (resp. second largest) cut imbalance ratio, with migration mostly occurring from the red counties to the blue counties. We note that the domestic migration report from this census (U.S. Census Bureau 2003) states that the South had the most in-migration and out-migration, with in-migration accounting for over 60% of the region’s total. In Fig. 11, this is only consistently observed (by coloring the South blue) when using weighted three-node motifs. The report also states that out-migration accounts for over 64% of migration involving the Northeast, and this is witnessed only by the motif  $\mathcal{M}_9$  (by coloring the Northeast red). It is also worth mentioning that the pattern of flow from Northeast to South uncovered by





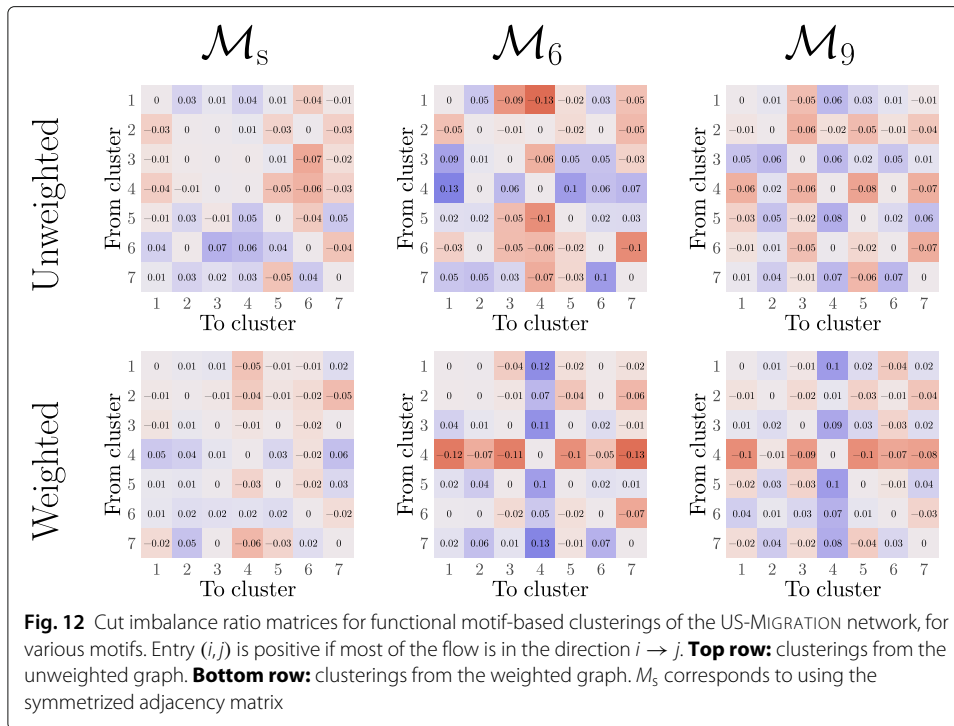
$M_9$  on the unweighted graph (top pair) and the weighted graph (second pair) bears significant resemblance to the topmost imbalanced structure uncovered by the HERM-SYM algorithm in Fig. 16 within (Cucuringu et al. 2019b).

Figure 12 displays the full cut imbalance ratio matrices (Appendix C4) associated with each method. The  $(i, j)$ th matrix entry is positive if there is more migration flow from cluster  $i$  to cluster  $j$  than from cluster  $j$  to cluster  $i$  (yielding an antisymmetric matrix), and a value of  $\pm 1/2$  indicates uni-directional flow. We see that weighted  $M_6$  and  $M_9$  produce clear “stripes” of blue and red in the matrices, where they identify clusters which have more imbalanced migration. This again highlights the ability of weighted motifs to uncover hidden structures (in this case corresponding to large-scale migration patterns).

## 5.2 US-POLITICAL-BLOGS network

The US-POLITICAL-BLOGS network (Adamic and Glance 2005) consists of data collected two months before the 2004 US election. Vertices represent blogs, and are labeled by their political leaning (“liberal” or “conservative”). Weighted directed edges represent the number of citations from one blog to another. After restricting to the largest connected component, there are 536 liberal blogs, 636 conservative blogs (total 1222) and 19024 edges. The network is plotted in Fig. 13.

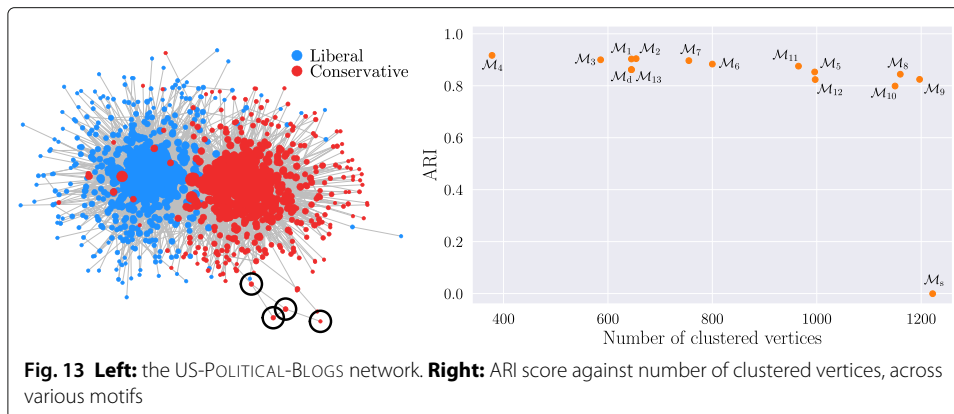
Firstly we use sweep profiles (Shi and Malik 2000) with a few selected motifs to motivate the use of motif-based clustering on this network. To construct a sweep profile, we



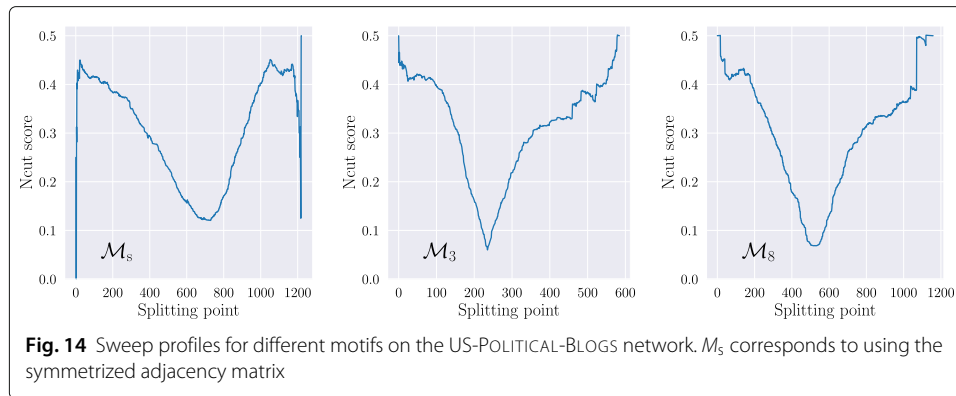
**Fig. 12** Cut imbalance ratio matrices for functional motif-based clusterings of the US-MIGRATION network, for various motifs. Entry  $(i, j)$  is positive if most of the flow is in the direction  $i \rightarrow j$ . **Top row:** clusterings from the unweighted graph. **Bottom row:** clusterings from the weighted graph.  $M_s$  corresponds to using the symmetrized adjacency matrix

order the vertices according to the first non-trivial eigenvector of the random-walk Laplacian, and select a splitting point based on minimizing an objective function; in this case the Ncut score (see Appendix C). Figure 14 exhibits visibly sharper minima for motifs  $M_3$  and  $M_8$  than for  $M_s$  (which corresponds to traditional spectral clustering), indicating that motif-based methods can produce more well-defined clusters (for the vertices contained in the largest connected component of the MAM) than traditional methods on this network.

In fact, traditional two-cluster spectral clustering performs very poorly on this network, with one of the clusters containing just four vertices (circled in Fig. 13), which are very weakly connected to the rest of the graph. However, motif-based clustering performs significantly better. As discussed in Section 3.3.2, our framework tries to reduce the misclassification error by only clustering the nodes which are in the largest connected component of the MAM, and are therefore in some sense well connected to the graph. In



**Fig. 13** **Left:** the US-POLITICAL-BLOGS network. **Right:** ARI score against number of clustered vertices, across various motifs



this dataset, the choice of motif determines the trade-off between the number of vertices assigned to clusters, and the accuracy of those clusters (see Fig. 13 for a plot of ARI score against number of vertices clustered). For example, motif  $M_9$  clusters 1197 vertices with an ARI of 0.82, while the more strongly connected  $M_4$  only clusters 378 vertices, with an improved ARI of 0.92. This is because the largest connected component of the MAM will be smaller for more strongly-connected motifs such as  $M_4$ .

To explore the classification performance of our reduced clustering method, we compare its effectiveness with other approaches from the literature. Table 1 gives the ARI, number of clustered vertices, normalized mutual information (NMI) (Nguyen et al. 2009), and percentage error for various motifs on this network. We note that both of the motifs  $M_3$  and  $M_8$  outperform the NMI score of 0.72 given by the degree-corrected stochastic blockmodel in (Karrer and Newman 2011), and improve on the percentage error of 4.66% given by the normalized spectral clustering procedure in (Li 2017), at the expense of not assigning every vertex to a cluster.

We further compare our proposed method to an alternative approach for dealing with disconnected vertices; namely regularization. We follow the method suggested by (Joseph et al. 2016) and (Zhang and Rohe 2018) using the regularized Laplacian  $L_{rw}^\tau = I - (D + \tau I)^{-1}(G + \tau \mathbf{1}/n)$ , where  $\mathbf{1}$  is a matrix of all ones and  $\tau$  is a tuning parameter. Table 2 shows how using this regularized Laplacian with  $\tau \sim \bar{d}$  (the mean weighted vertex degree), can improve the ARI performance of traditional (top row) and motif-based (second and third rows) spectral clustering while still assigning all of the vertices to clusters. However this regularization method is unable to reach the ARI scores of over 0.9 achieved by our proposed method of motif-based clustering on the largest connected component of the MAM (Table 1). Furthermore, regularization is sensitive to the choice of tuning parameter  $\tau$ , and the default setting of  $\tau = \bar{d}$  suggested by (Zhang and Rohe 2018) and (Qin and Rohe 2013) performs rather poorly here.

**Table 1** Performance of motif-based clustering on the US-POLITICAL-BLOGS network

Motif	ARI	Number of clustered vertices	NMI	Percentage error
$M_5$	0.00	1222	0.00	48.2%
$M_3$	0.90	586	0.83	2.6%
$M_8$	0.84	1160	0.75	4.1%



**Table 2** ARI scores for regularized traditional and motif-based spectral clustering on the US-POLITICAL-BLOGS network

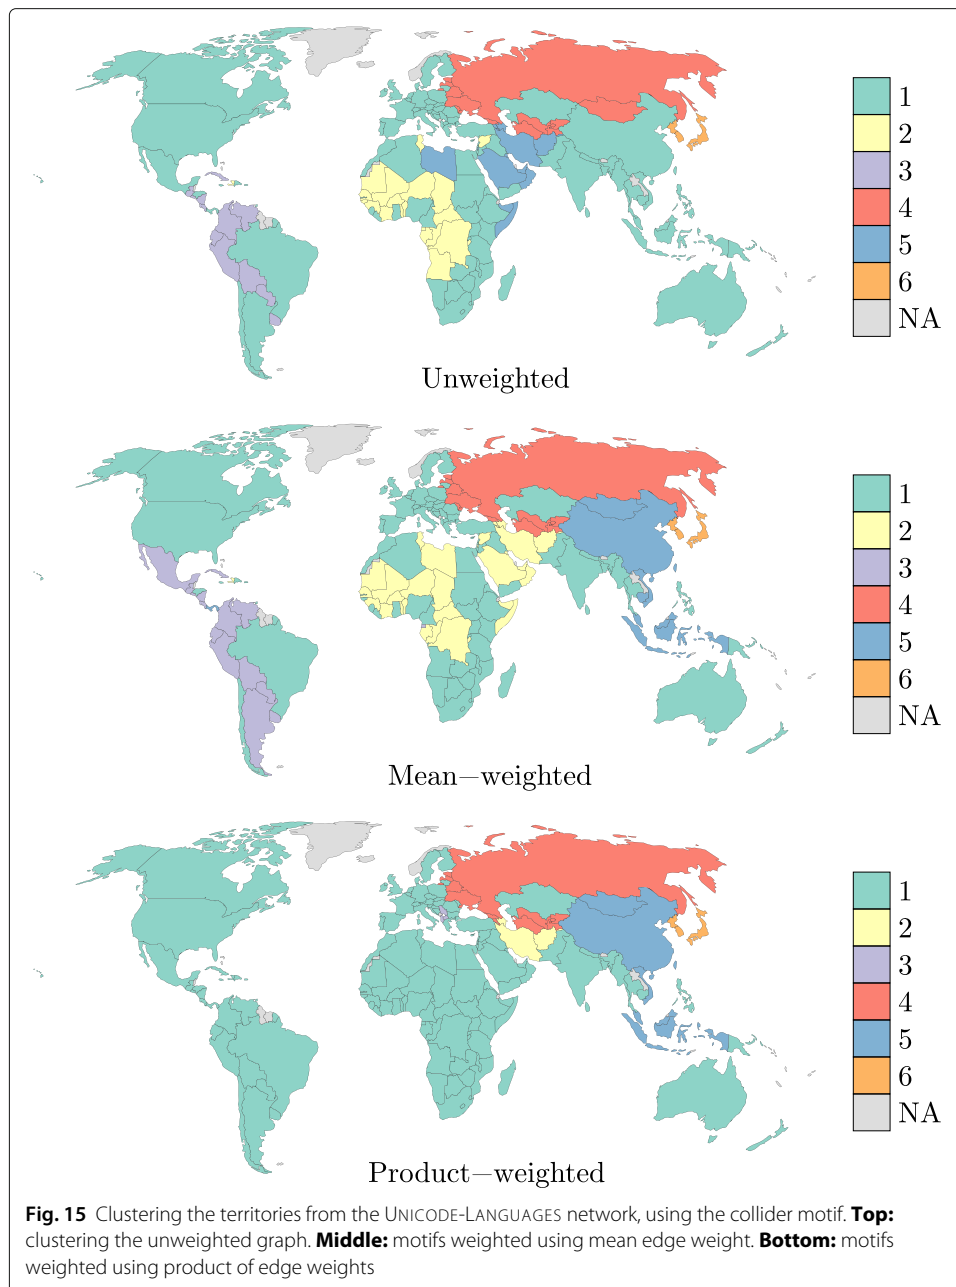
$\frac{\tau}{d}$	10	5	2	1	0.5	0.2	0.1	0.05	0.02	0.01	0.005	0.002	0.001
$\mathcal{M}_5$	0.08	0.14	0.28	0.36	0.49	0.81	0.80	0.81	0.81	0.82	0.82	0.00	0.00
$\mathcal{M}_3$	0.00	0.00	0.01	0.02	0.04	0.10	0.14	0.17	0.24	0.25	0.25	0.00	0.00
$\mathcal{M}_8$	0.05	0.07	0.13	0.28	0.49	0.58	0.62	0.65	0.71	0.75	0.80	0.80	0.80

### 5.3 UNICODE-LANGUAGES bipartite network

The UNICODE-LANGUAGES network (KONECT: The Koblenz Network Collection 2019) is a bipartite network consisting of data collected in 2014 on languages spoken around the world. Source vertices are territories, and destination vertices are languages. Weighted edges from territory to language indicate the number of inhabitants (GeoNames 2019) in that territory who speak the specified language. After removing territories with fewer than one million inhabitants, and languages with fewer than one million speakers, there are 155 territories, 270 languages and 705 edges remaining. We then restrict the network to its largest connected component, and in doing so remove the territories of Laos, Norway and Timor-Leste, and the languages of Lao, Norwegian Bokmål and Norwegian Nynorsk. We test Algorithm 3 with parameters  $k_S = l_S = k_D = l_D = 6$  on this network (such that each cluster in Table 3 contains at least 1% of the world's population), comparing three different motif weighting schemes: unweighted motifs, mean-weighted motifs and product-weighted motifs, as in Section 3.1. For the source vertices, Fig. 15 plots maps of the world with territories colored by the recovered clustering.

**Table 3** Clustering the territories from the UNICODE-LANGUAGES network using the mean-weighted collider motif

Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
India	Iran	Mexico	Russia	China	Japan
United States	DR Congo	Colombia	Ukraine	Indonesia	S. Korea
Brazil	Afghanistan	Argentina	Uzbekistan	Vietnam	N. Korea
Pakistan	Saudi Arabia	Peru	Belarus	Malaysia	
Bangladesh	Syria	Venezuela	Tajikistan	Taiwan	
Nigeria	Côte d'Ivoire	Ecuador	Kyrgyzstan	Cambodia	
Philippines	Burkina Faso	Guatemala	Turkmenistan	Hong Kong	
Ethiopia	Niger	Cuba	Georgia	Singapore	
Germany	Mali	Bolivia	Moldova	Panama	
Egypt	Senegal	Paraguay	Latvia	Mongolia	
Turkey	Tunisia	El Salvador	Estonia		
Thailand	Chad	Nicaragua			
France	Guinea	Costa Rica			
United Kingdom	Somalia	Uruguay			
Italy	Burundi	Eq. Guinea			
Myanmar	Haiti				
South Africa	Benin				
Spain	Azerbaijan				
Tanzania	Togo				
Kenya	Libya				
...	...				
Cluster 1  = 87	Cluster 2  = 29	Cluster 3  = 15	Cluster 4  = 11	Cluster 5  = 10	Cluster 6  = 3



Focusing first purely on the clusters produced by mean-weighted motifs, we observe balanced clusters with good spatial coherence (middle panel in Fig. 15), from which some language groups can be easily recognized. The top 20 territories (by population) in each cluster for this mean-weighted scheme are given in Table 3. Cluster 1 is by far the largest cluster, and includes a wide variety of territories. Cluster 2 contains several Persian-speaking and African French-speaking territories, while Cluster 3 mostly captures Spanish-speaking territories in the Americas. Cluster 4 includes the Slavic territories of Russia and some of its neighbors. Cluster 5 covers China, Hong Kong, Mongolia and some of South-East Asia. Cluster 6 is the smallest cluster and contains only Japan and the Koreans.

For the destination vertices, we present the six clusters obtained by Algorithm 3 with mean-weighted motifs. Table 4 contains the top 20 languages (by number of speakers) in each cluster. Cluster 1 is the largest cluster and contains some European languages and dialects of Arabic. Cluster 2 is also large and includes English, as well as several South Asian languages. Cluster 3 consists of many indigenous African languages. Cluster 4 captures languages from South-East Asia, mostly spoken in Indonesia and Malaysia. Cluster 5 identifies several varieties of Chinese and a few other Central and East Asian languages. Cluster 6 captures more South-East Asian languages, this time from Thailand, Myanmar and Cambodia.

Next, we return to the source vertices (territories) and compare the mean-weighted scheme with the unweighted and product-weighted schemes, as illustrated in Fig. 15. Unlike the mean-weighted approach, the unweighted partition fails to identify the Chinese-speaking territories, which is likely to be related to the fact that this scheme is unable to account for the large number of speakers in China. Further, it does not identify Mexico and Argentina with other Spanish-speaking territories in the Americas. Finally, the product-weighted motifs make no distinctions at all between territories in the Americas, Africa and Western Europe, and instead identify a few much smaller clusters.

## 6 Conclusion and future work

**Contribution** We have introduced generalizations of MAMs to weighted directed graphs and presented new matrix-based formulae for calculating them, which are easy to implement and scalable. By leveraging the popular random-walk spectral clustering algorithm, we have proposed motif-based techniques for clustering both weighted directed

**Table 4** Clustering the languages from the UNICODE-LANGUAGES network using the mean-weighted expander motif

Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
Spanish	English	Swahili	Indonesian	Chinese	Thai
Arabic	Hindi	Kinyarwanda	Javanese	Wu Chinese	N.E. Thai
Portuguese	Bengali	Somali	Malay	Korean	Khmer
French	Urdu	Luba-Lulua	Sundanese	Xiang Chinese	N. Thai
Russian	Punjabi	Kikuyu	Madurese	Hakka Chinese	S. Thai
Japanese	Telugu	Congo Swahili	Minangkabau	Minnan Chinese	Shan
German	Marathi	Luyia	Betawi	Gan Chinese	Pattani Malay
Turkish	Vietnamese	Ganda	Balinese	Kazakh	
Persian	Tamil	Luo	Buginese	Uighur	
Italian	Lahnda	Sukuma	Banjar	Sichuan Yi	
Egyptian Arabic	Filipino	Kalenjin	Achinese	Mongolian	
Polish	Gujarati	Lingala	Sasak	Zhuang	
Nigerian Pidgin	Kannada	Nyankole	Makasar	Tibetan	
Ukrainian	Pushto	Gusii	Lampung Api		
Dutch	Malayalam	Kiga	Rejang		
Algerian Arabic	Oriya	Soga			
Moroccan Arabic	Burmese	Luba-Katanga			
Hausa	Bhojpuri	Meru			
Azerbaijani	Amharic	Teso			
Uzbek	Oromo	Nyamwezi			
...	...	...			
Cluster 1  = 120	Cluster 2  = 90	Cluster 3  = 25	Cluster 4  = 15	Cluster 5  = 13	Cluster 6  = 7

graphs and weighted bipartite graphs. We demonstrated using synthetic data examples that accounting for edge weights and higher-order structure can be essential for good cluster recovery in directed and bipartite graphs, and that different motifs can uncover different but equally insightful clusters. Applications to real world networks have shown that weighted motif-based clustering can reveal a variety of different structures, and that it can reduce misclassification rate, performing favorably in comparison with other methods from the literature.

**Future work** There are limitations to our proposed methodology, which raise a number of research avenues to explore. While our matrix-based formulae for MAMs are simple to implement and scalable, there is scope for further computational improvement. As mentioned in (Benson et al. 2016), fast triangle enumeration algorithms (Demeyer et al. 2013; Wernicke 2006; Wernicke and Rasche 2006) may offer superior performance at scale for triangular motifs, but MAMs for motifs with missing edges are inherently denser. The implementation could also be optimized to further exploit any sparsity structures in the network, rather than relying on general-purpose linear algebra libraries. More recent work on combinatorial graphlet counting (Ahmed et al. 2015) or vertex orbit counting (Pashanasangi and Seshadhri 2020) could also potentially yield further computational improvements. Another shortcoming of our matrix-based formulae is that, unlike motif detection algorithms such as FANMOD (Wernicke and Rasche 2006), they do not easily extend to motifs on four or more vertices. While we believe that this extension is in principle possible using tensor operations rather than matrix operations, it is currently beyond the scope of the work.

We have seen that choosing the correct motif for clustering a network can play a crucial role (Section 4.4), and while we have suggested some methods for doing so (Section 3.3.3), it may be that other techniques can also be used to guide the motif selection process. We would also be interested in seeing more approaches to dealing with disconnected MAMs, perhaps involving some criterion for when a connected component is “large enough” to be worth keeping, or a deeper exploration of regularization methods for motif-based spectral clustering. Extensions to our methodology could involve further analysis of the differences between functional and structural MAMs, between different types of Laplacians (Von Luxburg 2007), or between different weighted stochastic graph models, perhaps following an exponential family method (Aicher et al. 2013). Also, connections between motif-based clustering and hypergraph partitioning (Li and Milenkovic 2017; Li and Milenkovic 2018; Veldt et al. 2020) would be worth further investigation.

Further experimental work would also be interesting. We would like to explore the utility of our framework on additional real data, and suggest that collaboration networks such as (Leskovec and Krevl 2007), transportation networks such as the airports network in (Frey and Dueck 2007), and bipartite preference networks such as (Clauset et al. 2007) could be interesting. Direct comparison of results with other state-of-the-art clustering methods from the literature could also be insightful; suitable benchmarks for performance include the Hermitian matrix-based clustering method in (Cucuringu et al. 2019b), the Motif-based Approximate Personalized PageRank (MAPPR) in (Yin et al. 2017), and TECTONIC from (Tsourakakis et al. 2017). Other established and fast methods from the literature on directed networks include SAPA (Satuluri and Parthasarathy 2011) and DI-SIM (Rohe et al. 2016), which respectively perform a degree-discounted symmetrization;

and a combined row and column clustering into four sets using the concatenation of the left and right singular vectors.

Further potential extensions of our work pertain to the detection of overlapping clusters, allowing nodes to belong to multiple groups (Li et al. 2016), to the detection of core–periphery structures in directed graphs (Elliott et al. 2019a), as well as implications for existing methods on network comparison (Wegner et al. 2018), anomaly detection in directed networks (Elliott et al. 2019b), and ranking from a sparse set of noisy pairwise comparisons (Cucuringu 2016). Furthermore, incorporating our pipeline into deep learning architectures is another avenue worth exploring, building on the recent MOTIFNET architecture (Monti et al. 2018), a graph convolutional neural network for directed graphs, or SIGAT (Huang et al. 2019), another recent graph neural network which incorporates graph motifs in the context of signed networks which contain both positive and negative links (Cucuringu et al. 2019a).

### Appendix A: Motif adjacency matrix formulae

We give explicit matrix-based formulae for mean-weighted motif adjacency matrices  $M$  for all simple motifs  $\mathcal{M}$  on at most three vertices, along with the anchored motifs  $\mathcal{M}_{\text{coll}}$  and  $\mathcal{M}_{\text{expa}}$ . Entry-wise (Hadamard) products are denoted by  $\circ$ . Table 5 gives our dense formulation of functional MAMs. For the dense formulation of structural MAMs, simply replace  $J_n, J$ , and  $G$  with  $J_0, J_s$ , and  $G_s$  respectively. For the sparse formulation of functional MAMs, replace  $J_n$  with  $\mathbf{1} - J$ , and expand and simplify the expressions as in Section 3.4. For the sparse formulation of structural MAMs, replace  $J$  and  $G$  with  $J_s$  and  $G_s$  respectively. Also replace  $J_0$  with  $\mathbf{1} - (I + \tilde{J})$ , where  $\tilde{J} = J_s + J_s^T + J_d$ , and expand and simplify the expressions.

### Appendix B: Proofs

**Proof B.1** (Proposition 3.1, MAM formula) Consider Eq. 1. We sum over functional instances  $\mathcal{M} \cong \mathcal{H} \leq \mathcal{G}$  such that  $\{i, j\} \in \mathcal{A}(\mathcal{H})$ . This is equivalent to summing over  $\{k_2, \dots, k_{m-1}\} \subseteq \mathcal{V}$  and  $\sigma \in S_{\tilde{\mathcal{M}}, \mathcal{A}}$ , such that  $k_u$  are all distinct and

$$(u, v) \in \mathcal{E}_{\mathcal{M}} \implies (k_{\sigma u}, k_{\sigma v}) \in \mathcal{E}. \quad (\dagger) \quad (7)$$

This is because the vertex set  $\{k_2, \dots, k_{m-1}\} \subseteq \mathcal{V}$  indicates which vertices are present in the instance  $\mathcal{H}$ , and  $\sigma$  describes the mapping from  $\mathcal{V}_{\mathcal{M}}$  onto those vertices:  $u \mapsto k_{\sigma u}$ . We take  $\sigma \in S_{\tilde{\mathcal{M}}, \mathcal{A}}$  to ensure that  $\{i, j\} \in \mathcal{A}(\mathcal{H})$  (since  $i = k_1, j = k_m$ ), and that instances are counted exactly once. The condition  $(\dagger)$  is to check that  $\mathcal{H}$  is a functional instance of  $\mathcal{M}$  in  $\mathcal{G}$ . Hence

$$\begin{aligned} M_{ij}^{\text{func}} &= \frac{1}{|\mathcal{E}_{\mathcal{M}}|} \sum_{\mathcal{M} \cong \mathcal{H} \leq \mathcal{G}} \mathbb{I}\{\{i, j\} \in \mathcal{A}(\mathcal{H})\} \sum_{e \in \mathcal{E}_{\mathcal{H}}} W(e) \\ &= \frac{1}{|\mathcal{E}_{\mathcal{M}}|} \sum_{\{k_2, \dots, k_{m-1}\}} \sum_{\sigma \in S_{\tilde{\mathcal{M}}, \mathcal{A}}} \mathbb{I}\{k_u \text{ all distinct, } (\dagger)\} \sum_{e \in \mathcal{E}_{\mathcal{H}}} W(e). \end{aligned}$$

For the first term, by conditioning on the types of edge in  $\mathcal{E}_{\mathcal{M}}$ :

**Table 5** Dense formulation of mean-weighted functional motif adjacency matrices

Motif	$C$	$C'$	$M^{\text{func}}$
$\mathcal{M}_5$			$G + G^T$
$\mathcal{M}_d$			$\frac{1}{2}G_d$
$\mathcal{M}_1$	$J^T \circ (JG) + J^T \circ (GJ) + G^T \circ (JJ)$ $J^T \circ (J_dG) + J^T \circ (G_dJ) + G^T \circ (J_dJ)$		$\frac{1}{3}(C + C^T)$
$\mathcal{M}_2$	$+ J^T \circ (JG_d) + J^T \circ (GJ_d) + G^T \circ (JJ_d)$ $+ J_d \circ (JG) + J_d \circ (GJ) + G_d \circ (JJ)$ $J \circ (J_dG_d) + J \circ (G_dJ_d) + G \circ (J_dJ_d)$		$\frac{1}{4}(C + C^T)$
$\mathcal{M}_3$	$+ J_d \circ (J_dG) + J_d \circ (G_dJ) + G_d \circ (J_dJ)$ $+ J_d \circ (JG_d) + J_d \circ (GJ_d) + G_d \circ (JJ_d)$		$\frac{1}{5}(C + C^T)$
$\mathcal{M}_4$	$J_d \circ (J_dG_d) + J_d \circ (G_dJ_d) + G_d \circ (J_dJ_d)$ $J \circ (JG) + J \circ (GJ) + G \circ (JJ)$		$\frac{1}{6}C$
$\mathcal{M}_5$	$+ J \circ (JG^T) + J \circ (GJ^T) + G \circ (JJ^T)$ $+ J \circ (J^TG) + J \circ (G^TJ) + G \circ (J^TJ)$		$\frac{1}{3}(C + C^T)$
$\mathcal{M}_6$	$J \circ (JG_d) + J \circ (GJ_d) + G \circ (JJ_d) + J_d \circ (J^TG)$	$G_d \circ (J^TJ)$	$\frac{1}{4}(C + C^T + C')$
$\mathcal{M}_7$	$J \circ (J_dG) + J \circ (G_dJ) + G \circ (J_dJ)$	$J_d \circ (JG^T) + J_d \circ (GJ^T)$ $+ G_d \circ (JJ^T)$	$\frac{1}{4}(C + C^T + C')$
$\mathcal{M}_8$	$J \circ (GJ_n) + G \circ (JJ_n)$ $J \circ (J_nG^T) + G \circ (J_nJ^T) + J_n \circ (JG)$	$J_n \circ (J^TG) + J_n \circ (G^TJ)$	$\frac{1}{2}(C + C^T + C')$
$\mathcal{M}_9$	$+ J_n \circ (GJ) + J \circ (G^TJ_n) + G \circ (J^TJ_n)$		$\frac{1}{2}(C + C^T)$
$\mathcal{M}_{10}$	$J \circ (J_nG) + G \circ (J_nJ)$	$J_n \circ (JG^T) + J_n \circ (GJ^T)$	$\frac{1}{2}(C + C^T + C')$
$\mathcal{M}_{11}$	$J_d \circ (GJ_n) + G_d \circ (JJ_n) + J_n \circ (J_dG)$ $+ J_n \circ (G_dJ) + J \circ (G_dJ_n) + G \circ (J_dJ_n)$ $J_d \circ (J_nG) + G_d \circ (J_nJ) + J_n \circ (JG_d)$		$\frac{1}{3}(C + C^T)$
$\mathcal{M}_{12}$	$+ J_n \circ (GJ_d) + J \circ (J_nG_d) + G \circ (J_nJ_d)$		$\frac{1}{3}(C + C^T)$
$\mathcal{M}_{13}$	$J_d \circ (G_dJ_n) + G_d \circ (J_dJ_n) + J_n \circ (J_dG_d)$		$\frac{1}{4}(C + C^T)$
$\mathcal{M}_{\text{coll}}$	$J_n \circ (JG^T)$		$\frac{1}{2}(C + C^T)$
$\mathcal{M}_{\text{expa}}$	$J_n \circ (J^TG)$		$\frac{1}{2}(C + C^T)$

$$\begin{aligned}
\mathbb{I}\{k_u \text{ all distinct, } (\dagger)\} &= \prod_{\mathcal{E}_{\mathcal{M}}^0} \mathbb{I}\{k_{\sigma u} \neq k_{\sigma v}\} \\
&\quad \times \prod_{\mathcal{E}_{\mathcal{M}}^s} \mathbb{I}\{(k_{\sigma u}, k_{\sigma v}) \in \mathcal{E}\} \\
&\quad \times \prod_{\mathcal{E}_{\mathcal{M}}^d} \mathbb{I}\{(k_{\sigma u}, k_{\sigma v}) \in \mathcal{E} \text{ and } (k_{\sigma v}, k_{\sigma u}) \in \mathcal{E}\} \\
&= \prod_{\mathcal{E}_{\mathcal{M}}^0} (J_n)_{k_{\sigma u}, k_{\sigma v}} \prod_{\mathcal{E}_{\mathcal{M}}^s} J_{k_{\sigma u}, k_{\sigma v}} \prod_{\mathcal{E}_{\mathcal{M}}^d} (J_d)_{k_{\sigma u}, k_{\sigma v}} \\
&= J_{\mathbf{k}, \sigma}^{\text{func}}.
\end{aligned}$$

Assuming  $\{k_u \text{ all distinct, } (\dagger)\}$ , the second term is

$$\begin{aligned}
\sum_{e \in \mathcal{E}_{\mathcal{H}}} W(e) &= \sum_{\mathcal{E}_{\mathcal{M}}^s} W((k_{\sigma u}, k_{\sigma v})) + \sum_{\mathcal{E}_{\mathcal{M}}^d} (W((k_{\sigma u}, k_{\sigma v})) + W((k_{\sigma v}, k_{\sigma u}))) \\
&= \sum_{\mathcal{E}_{\mathcal{M}}^s} G_{k_{\sigma u}, k_{\sigma v}} + \sum_{\mathcal{E}_{\mathcal{M}}^d} (G_d)_{k_{\sigma u}, k_{\sigma v}} \\
&= G_{\mathbf{k}, \sigma}^{\text{func}}
\end{aligned}$$

as required. For Eq. 2, we simply change  $(\dagger)$  to  $(\ddagger)$  to check that an instance is a *structural* instance

$$(u, v) \in \mathcal{E}_{\mathcal{M}} \iff (k_{\sigma u}, k_{\sigma v}) \in \mathcal{E} \quad (\ddagger)$$

Now for the first term, the following holds true

$$\begin{aligned}
\mathbb{I}\{k_u \text{ all distinct, } (\ddagger)\} &= \prod_{\mathcal{E}_{\mathcal{M}}^0} \mathbb{I}\{(k_{\sigma u}, k_{\sigma v}) \notin \mathcal{E} \text{ and } (k_{\sigma v}, k_{\sigma u}) \notin \mathcal{E}\} \\
&\quad \times \prod_{\mathcal{E}_{\mathcal{M}}^s} \mathbb{I}\{(k_{\sigma u}, k_{\sigma v}) \in \mathcal{E} \text{ and } (k_{\sigma v}, k_{\sigma u}) \notin \mathcal{E}\} \\
&\quad \times \prod_{\mathcal{E}_{\mathcal{M}}^d} \mathbb{I}\{(k_{\sigma u}, k_{\sigma v}) \in \mathcal{E} \text{ and } (k_{\sigma v}, k_{\sigma u}) \in \mathcal{E}\} \\
&= \prod_{\mathcal{E}_{\mathcal{M}}^0} (J_0)_{k_{\sigma u}, k_{\sigma v}} \prod_{\mathcal{E}_{\mathcal{M}}^s} (J_s)_{k_{\sigma u}, k_{\sigma v}} \prod_{\mathcal{E}_{\mathcal{M}}^d} (J_d)_{k_{\sigma u}, k_{\sigma v}} \\
&= J_{\mathbf{k}, \sigma}^{\text{struc}}.
\end{aligned}$$

Assuming  $\{k_u \text{ all distinct, } (\ddagger)\}$ , the second term is given by

$$\begin{aligned}
\sum_{e \in \mathcal{E}_{\mathcal{H}}} W(e) &= \sum_{\mathcal{E}_{\mathcal{M}}^s} W((k_{\sigma u}, k_{\sigma v})) + \sum_{\mathcal{E}_{\mathcal{M}}^d} (W((k_{\sigma u}, k_{\sigma v})) + W((k_{\sigma v}, k_{\sigma u}))) \\
&= \sum_{\mathcal{E}_{\mathcal{M}}^s} (G_s)_{k_{\sigma u}, k_{\sigma v}} + \sum_{\mathcal{E}_{\mathcal{M}}^d} (G_d)_{k_{\sigma u}, k_{\sigma v}} \\
&= G_{\mathbf{k}, \sigma}^{\text{struc}}.
\end{aligned}$$

**Proof B.2** (Proposition 3.2, Colliders and expanders in bipartite graphs) Consider Eq. 3 and the collider motif  $\mathcal{M}_{\text{coll}}$ . Since  $\mathcal{G}$  is bipartite,  $M_{\text{coll}}^{\text{func}} = M_{\text{coll}}^{\text{struc}} =: M_{\text{coll}}$ , and by Table 5,  $M_{\text{coll}} = \frac{1}{2} J_n \circ (JG^T + GJ^T)$ . Hence

$$\begin{aligned}
(M_{\text{coll}})_{ij} &= \frac{1}{2} (J_n)_{ij} (JG^T + GJ^T)_{ij} \\
&= \mathbb{I}\{i \neq j\} \sum_{k \in \mathcal{V}} \frac{1}{2} (J_{ik} G_{jk} + G_{ik} J_{jk}) \\
&= \mathbb{I}\{i \neq j\} \sum_{k \in \mathcal{V}} \frac{1}{2} \mathbb{I}\{(i, k), (j, k) \in \mathcal{E}\} [W((i, k)) + W((j, k))] \\
&= \mathbb{I}\{i \neq j\} \sum_{\substack{k \in \mathcal{D} \\ (i, k), (j, k) \in \mathcal{E}}} \frac{1}{2} [W((i, k)) + W((j, k))].
\end{aligned}$$

Similarly for the expander motif  $M_{\text{expa}} = \frac{1}{2} J_n \circ (J^T G + G^T J)$ , which yields Eq. (4):

$$\begin{aligned} (M_{\text{expa}})_{ij} &= \frac{1}{2} (J_n)_{ij} (J^T G + G^T J)_{ij} \\ &= \mathbb{I}\{i \neq j\} \sum_{\substack{k \in \mathcal{S} \\ (k,i), (k,j) \in \mathcal{E}}} \frac{1}{2} [W((k,i)) + W((k,j))]. \end{aligned}$$

**Proof B.3** (Proposition 3.3, Complexity of MAM formula) Suppose  $m \leq 3$  and consider  $M^{\text{func}}$ . The adjacency and indicator matrices of  $\mathcal{G}$  are given by

$$\begin{aligned} (1) \quad J &= \mathbb{I}\{G > 0\}, & (5) \quad J_n &= \mathbb{I}\{I_{n \times n} = 0\}, \\ (2) \quad J_0 &= \mathbb{I}\{G + G^T = 0\} \circ J_n, & (6) \quad J_d &= J \circ J^T, \\ (3) \quad J_s &= J - J_d, & (7) \quad G_s &= G \circ J_s, \\ (4) \quad G_d &= (G + G^T) \circ J_d, \end{aligned}$$

and computed using four additions and four element-wise multiplications.  $J_{k,\sigma}^{\text{func}}$  is a product of at most three factors, and  $G_{k,\sigma}^{\text{func}}$  contains at most three summands, thus

$$\sum_{k_2 \in \mathcal{V}} J_{k,\sigma}^{\text{func}} G_{k,\sigma}^{\text{func}}$$

is expressible as a sum of at most three matrices, each of which is constructed with at most one matrix multiplication (where  $\{k_{\sigma r}, k_{\sigma s}\} \neq \{i, j\}$ ) and one entry-wise multiplication (where  $\{k_{\sigma r}, k_{\sigma s}\} = \{i, j\}$ ). This is repeated for each  $\sigma \in \tilde{S}_{\mathcal{M}, \mathcal{A}}$  (at most six times) and the results are summed. Calculations are identical for  $M^{\text{struc}}$ .

### Appendix C: Spectral clustering

We provide a summary of traditional random-walk spectral clustering and show how it applies to motif-based clustering.

#### Overview of spectral clustering

For an undirected graph, the objects to be clustered are the vertices of the graph, and a similarity matrix is provided by the graph's adjacency matrix  $G$ . To cluster directed graphs, the adjacency matrix must first be symmetrized, perhaps by considering  $G + G^T$  (Meilă and Pentney 2007). This symmetrization ignores information about edge direction and higher-order structures, and can lead to poor performance (Benson et al. 2016). Spectral clustering consists of two steps. Firstly, eigendecomposition of a Laplacian matrix embeds the vertices into  $\mathbb{R}^l$ . The  $k$  clusters are then extracted from this space.

#### Graph laplacians

The Laplacians of an undirected graph are a family of matrices which play a central role in spectral clustering. While many different graph Laplacians are available, we focus on just the *random-walk Laplacian*, for reasons concerning objective functions, consistency and computation (Von Luxburg et al. 2004; Von Luxburg 2007). The random-walk Laplacian of an undirected graph  $\mathcal{G}$  with (symmetric) adjacency matrix  $G$  is given by

$$L_{\text{rw}} := I - D^{-1}G,$$

where  $I$  is the identity and  $D_{ii} := \sum_j G_{ij}$  is the diagonal matrix of weighted degrees.



### Graph cuts

Graph cuts provide objective functions which we seek to minimize while clustering the vertices of a graph.

**Definition 3.4** Let  $\mathcal{G}$  be a graph. Let  $\mathcal{P}_1, \dots, \mathcal{P}_k$  be a partition of  $\mathcal{V}$ . Then the normalized cut (Shi and Malik 2000) of  $\mathcal{G}$  with respect to  $\mathcal{P}_1, \dots, \mathcal{P}_k$  is

$$Ncut_{\mathcal{G}}(\mathcal{P}_1, \dots, \mathcal{P}_k) := \frac{1}{2} \sum_{i=1}^k \frac{cut(\mathcal{P}_i, \bar{\mathcal{P}}_i)}{vol(\mathcal{P}_i)},$$

where  $cut(\mathcal{P}_i, \bar{\mathcal{P}}_i) := \sum_{u \in \mathcal{P}_i, v \in \mathcal{V} \setminus \mathcal{P}_i} G_{uv}$  and  $vol(\mathcal{P}_i) := \sum_{u \in \mathcal{P}_i} D_{uu}$ .

Note that more desirable partitions have a lower Ncut value; the numerators penalize partitions which cut a large number of heavily weighted edges, and the denominators penalize partitions which have highly imbalanced cluster sizes. It can be shown (Von Luxburg 2007) that minimizing Ncut over partitions  $\mathcal{P}_1, \dots, \mathcal{P}_k$  is equivalent to finding the cluster indicator matrix  $H \in \mathbb{R}^{n \times k}$  minimizing  $\text{Tr}(H^T(D - G)H)$  subject to  $H_{ij} = vol(\mathcal{P}_j)^{-\frac{1}{2}} \mathbb{I}\{v_i \in \mathcal{P}_j\}$  ( $\dagger$ ), and  $H^T D H = I$ . Solving this problem is in general NP-hard (Wagner and Wagner 1993). However, by dropping the constraint ( $\dagger$ ) and applying the Rayleigh Principle (Lütkepohl 1996), we find that the solution to this relaxed problem is that  $H$  contains the first  $k$  eigenvectors of  $L_{rw}$  as columns (Von Luxburg 2007). In practice, to find  $k$  clusters it is often sufficient to use only the first  $l < k$  eigenvectors of  $L_{rw}$ .

### Cut imbalance ratio

Given a partition of a directed graph, cut imbalance ratio gives a notion of flow imbalance between each pair of clusters.

**Definition 3.5** Let  $\mathcal{G}$  be a graph. Let  $\mathcal{P}_1, \dots, \mathcal{P}_k$  be a partition of  $\mathcal{V}$ . Then the cut imbalance ratio (Cucuringu et al. 2019b) of  $\mathcal{G}$  from  $\mathcal{P}_i$  to  $\mathcal{P}_j$  is

$$CIR_{\mathcal{G}}(\mathcal{P}_i, \mathcal{P}_j) := \frac{1}{2} \frac{cut(\mathcal{P}_i, \mathcal{P}_j) - cut(\mathcal{P}_j, \mathcal{P}_i)}{cut(\mathcal{P}_i, \mathcal{P}_j) + cut(\mathcal{P}_j, \mathcal{P}_i)},$$

where  $cut(\mathcal{P}_i, \mathcal{P}_j) := \sum_{u \in \mathcal{P}_i, v \in \mathcal{P}_j} G_{uv}$ .

The cut imbalance ratio takes values in  $[-\frac{1}{2}, \frac{1}{2}]$ , with positive values indicating more flow from  $\mathcal{P}_i$  to  $\mathcal{P}_j$  and negative values indicating more flow from  $\mathcal{P}_j$  to  $\mathcal{P}_i$ . The *cut imbalance ratio matrix* is the antisymmetric  $k \times k$  matrix with  $(i, j)$ th entry equal to  $CIR_{\mathcal{G}}(\mathcal{P}_i, \mathcal{P}_j)$ .

### Cluster extraction

Once Laplacian eigendecomposition has been used to embed the data into  $\mathbb{R}^l$ , the clusters may be extracted using a variety of methods. We propose  $k$ -means++ (Arthur and Vassilvitskii 2007), a popular clustering algorithm for data in  $\mathbb{R}^l$ , as an appropriate technique. It aims to minimize the within-cluster sum of squares, based on the standard Euclidean metric on  $\mathbb{R}^l$ . This makes it a reasonable candidate for clustering spectral data, since the Euclidean metric corresponds to notions of “diffusion distance” in the original graph (Nadler et al. 2006).

### Random-walk spectral clustering

Algorithm 1 gives random-walk spectral clustering (Von Luxburg 2007), which takes a symmetric connected adjacency matrix as input. We drop the first column of  $H$  (the first eigenvector of  $L_{rw}$ ) since although it should be constant and uninformative, numerical imprecision may give unwanted artifacts. It is worth noting that although the relaxation used in Appendix C.3 is reasonable and often leads to good approximate solutions of the Ncut problem, there are cases where it performs poorly (Guattery and Miller 1998). The Cheeger inequality (Chung 2005) gives a bound on the error introduced by this relaxation.

### Motif-based random-walk spectral clustering

Algorithm 2 gives motif-based random-walk spectral clustering. The algorithm forms a motif adjacency matrix, restricts it to its largest connected component, and applies random-walk spectral clustering (Algorithm 1) to produce clusters. The most computationally expensive part of Algorithm 2 is the calculation of the MAM using a formula from Table 5, as noted by (Benson et al. 2016) for their unweighted MAMs. The complexity of this is analyzed in Section 3.4.

### Bipartite spectral clustering

Algorithm 3 gives our procedure for clustering a bipartite graph. The algorithm uses the collider and expander motifs to create similarity matrices for the source and destination vertices respectively (as in Section 3.2), and then applies random-walk spectral clustering (Algorithm 1) to produce the partitions.

#### Algorithm 1: Random-walk spectral clustering

**Input:** Symmetric adjacency matrix  $G$ , number of clusters  $k$ , dimension  $l$

**Output:** Partition  $\mathcal{P}_1, \dots, \mathcal{P}_k$

```

1 function RWSpectClust ( $G, k, l$ ):
2   Construct the weighted degree matrix  $D_{ii} \leftarrow \sum_j G_{ij}$ 
3   Construct the random-walk Laplacian matrix  $L_{rw} \leftarrow I - D^{-1}G$ 
4   Let  $H$  have the first  $l$  eigenvectors of  $L_{rw}$  as columns
5   Drop the first column of  $H$ 
6   Run  $k$ -means++ on the rows of  $H$  with  $k$  clusters to produce  $\mathcal{P}_1, \dots, \mathcal{P}_k$ 
7   return  $\mathcal{P}_1, \dots, \mathcal{P}_k$ 

```

#### Algorithm 2: Motif-based random-walk spectral clustering


**Input:** Graph  $\mathcal{G}$ , motif  $\mathcal{M}$ , number of clusters  $k$ , dimension  $l$

**Output:** Partition  $\mathcal{P}_1, \dots, \mathcal{P}_k$

```

1 function MotifRWSpectClust ( $\mathcal{G}, \mathcal{M}, k, l$ ):
2   Construct the motif adjacency matrix  $M$  of the graph  $\mathcal{G}$  with motif  $\mathcal{M}$ 
3   Let  $\tilde{M}$  be  $M$  restricted to its largest connected component,  $C$ 
4    $\mathcal{P}_1, \dots, \mathcal{P}_k \leftarrow$  RWSpectClust ( $\tilde{M}, k, l$ )
5   return  $\mathcal{P}_1, \dots, \mathcal{P}_k$ 

```



**Fig. 16** Exploring the performance of structural MAMs based on  $\mathcal{M}_8$  and  $\mathcal{M}_{10}$  on Example 1 (see Fig. 5). Each block contains 100 nodes. We compare both to the unweighted case, and to the symmetrized case. We perform 100 repeats, and error bars are one sample standard deviation

#### Appendix D: Additional synthetic results

In this section we present additional results for some of our synthetic experiments.

**Example 1** In Fig. 16 we present the results for structural motifs, to complement the functional motifs presented in the main paper. We note that when we consider structural motifs, the unweighted motifs perform better but are still outperformed by our weighted method.

**Example 2 – Experiment 1** In Fig. 17 we present the results for the functional motifs for Example 2 Experiment 1. As noted in the main paper, our approaches outperform the baseline, but we do not observe the same bi-modal structure, with  $\mathcal{M}_1$  outperforming all approaches for  $q_1 < 0.6$ , and without the recovery in performance for high values of  $q_1$ .

**Example 2 – Experiment 2** In Fig. 18 we present the structural version of Example 2 Experiment 2. We observe the same structures as we observe for the functional version,

#### Algorithm 3: Bipartite random-walk spectral clustering

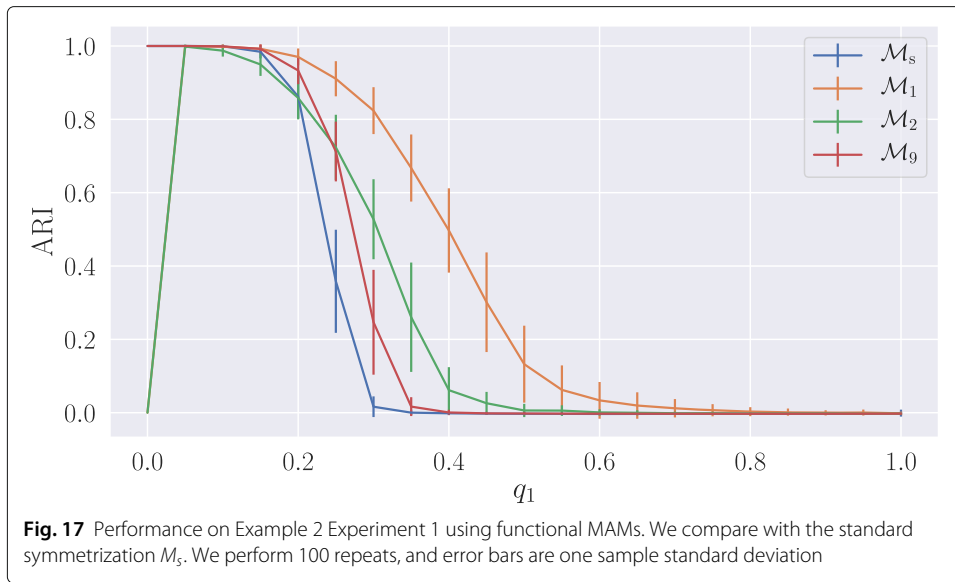
**Input:** Bipartite graph  $\mathcal{G}$ , source clusters  $k_S$ , destination clusters  $k_D$ , source dimension  $l_S$ , destination dimension  $l_D$

**Output:** Source partition  $\mathcal{S}_1, \dots, \mathcal{S}_{k_S}$ , destination partition  $\mathcal{D}_1, \dots, \mathcal{D}_{k_D}$

```

1 function BipartiteRWSpectClust ( $\mathcal{G}, k_S, k_D, l_S, l_D$ ):
2   Construct the collider motif adjacency matrix  $M_{\text{coll}}$  of the graph  $\mathcal{G}$ 
3   Construct the expander motif adjacency matrix  $M_{\text{expa}}$  of the graph  $\mathcal{G}$ 
4    $M_{\text{coll}} \leftarrow M_{\text{coll}}[\mathcal{S}, \mathcal{S}]$ ;           ▷ restrict rows and columns of  $M_{\text{coll}}$  to  $\mathcal{S}$ 
5    $M_{\text{expa}} \leftarrow M_{\text{expa}}[\mathcal{D}, \mathcal{D}]$ ;       ▷ restrict rows and columns of  $M_{\text{expa}}$  to  $\mathcal{D}$ 
6    $\mathcal{S}_1, \dots, \mathcal{S}_{k_S} \leftarrow \text{RWSpectClust}(M_{\text{coll}}, k_S, l_S)$ 
7    $\mathcal{D}_1, \dots, \mathcal{D}_{k_D} \leftarrow \text{RWSpectClust}(M_{\text{expa}}, k_D, l_D)$ 
8   return  $\mathcal{S}_1, \dots, \mathcal{S}_{k_S}$  and  $\mathcal{D}_1, \dots, \mathcal{D}_{k_D}$ 

```

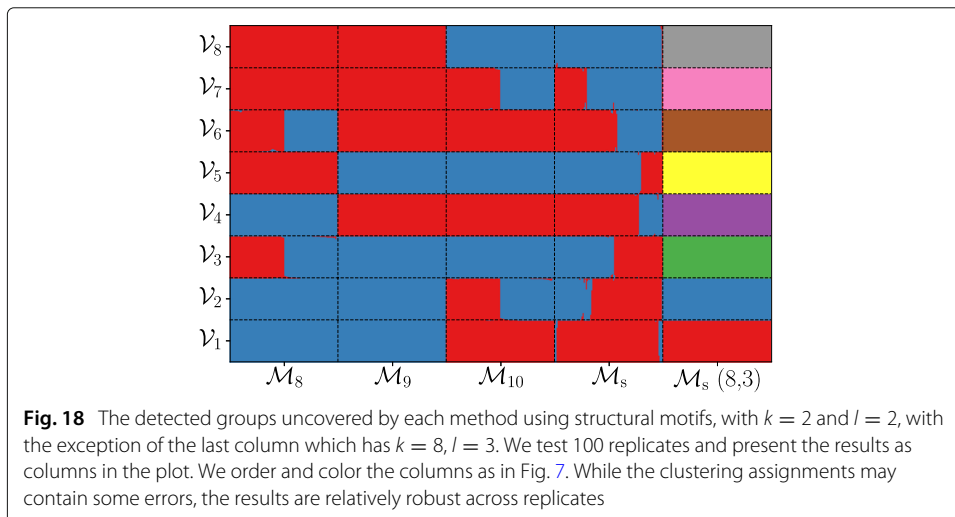


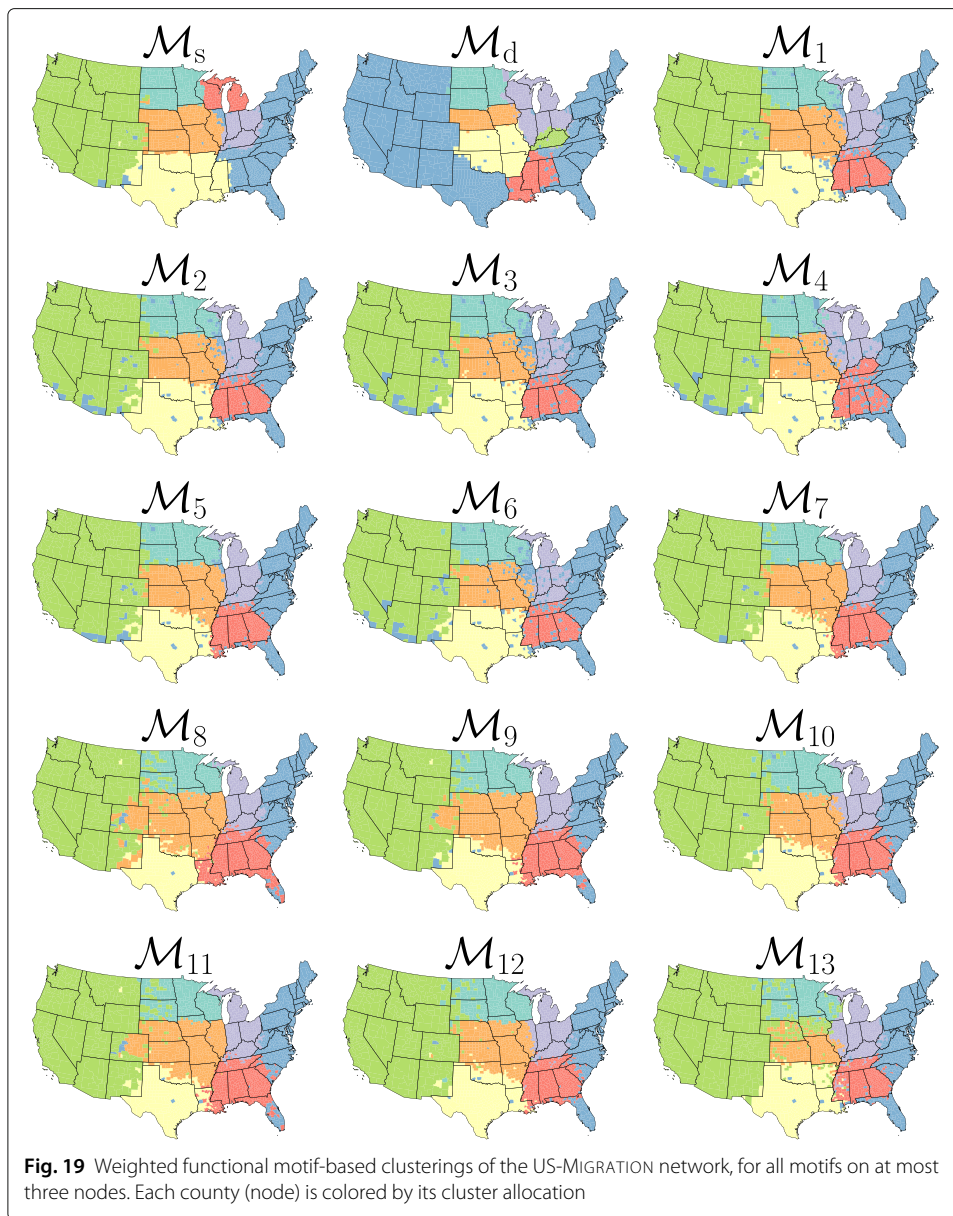
with motifs each highlighting a different but equally relevant structure. Finally the symmetrized adjacency matrix ( $M_s$ ), for which each of the blocks are indistinguishable, does not robustly identify the higher-order structures for  $k = 2$ , but can uncover the blocks when clustering with  $k = 8$  and  $l = 3$ .

**Appendix E: Additional real world results**

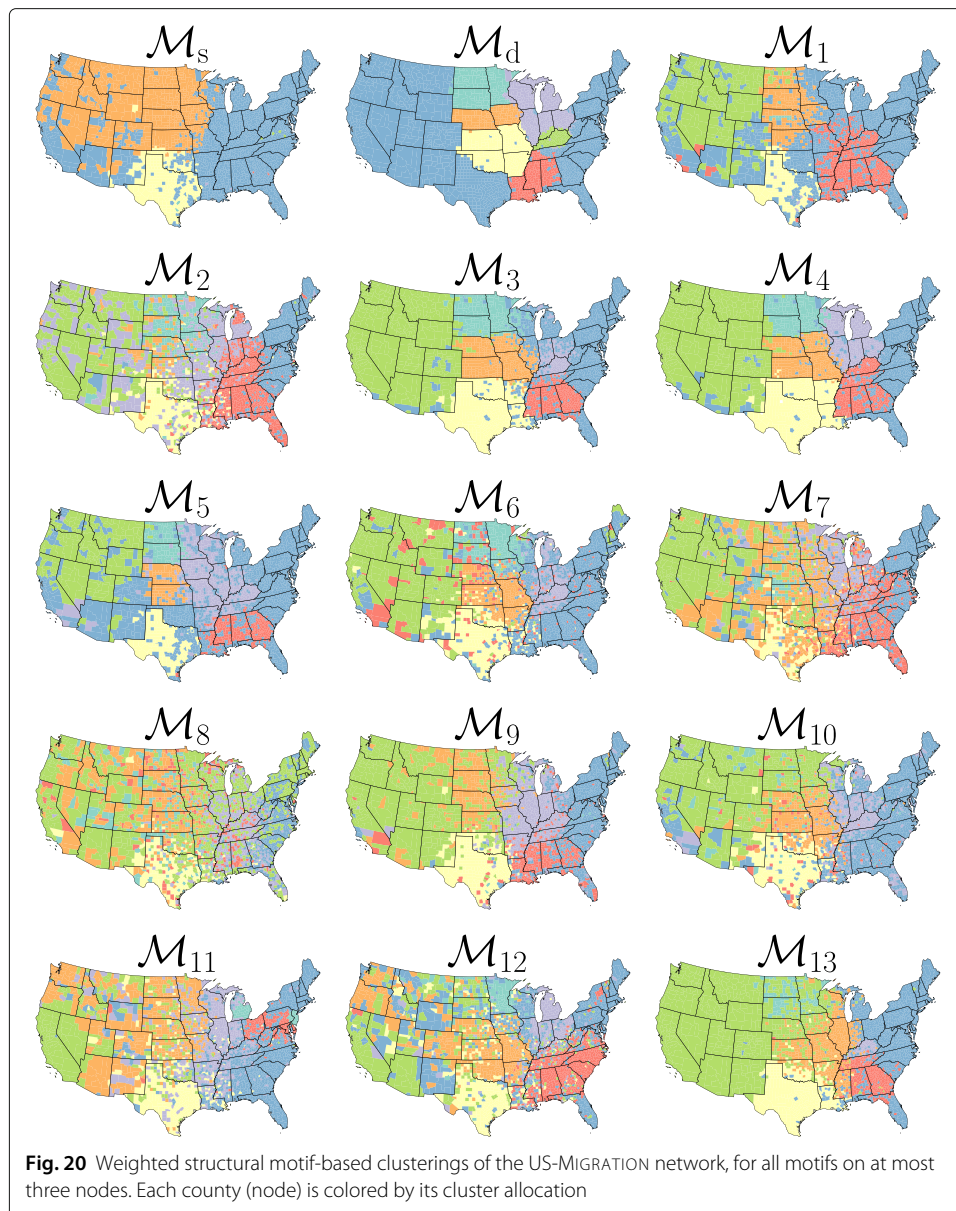
In this section we present additional results for some of our real world experiments.

**US-MIGRATION network** Fig. 19 plots clusterings obtained from the US-MIGRATION network using all 15 weighted functional motifs. Note how the different motifs uncover a variety of different clusterings, and also that they all display good spatial coherence. For completeness, Fig. 20 shows the equivalent plots for structural motifs. Note that motifs  $M_d$  and  $M_4$  are bi-directional cliques on two and three vertices respectively,





so their functional and structural MAMs are identical. The spatial coherence deteriorates for certain structural motifs, and cluster sizes are occasionally less balanced. We believe this is to do with the high local density of the migration network, preventing some motifs from appearing as structural instances in certain places (see Section 4.4 for another example of structural motifs exhibiting this filtering property, which may or may not be desirable, depending on the application). Note that this effect is most pronounced for the motifs containing single edges (i.e. excluding  $\mathcal{M}_d$ ,  $\mathcal{M}_4$  and  $\mathcal{M}_{13}$ ), because structural instances of single edges require an edge in precisely one direction, filtering out pairs of nodes which exhibit either no mutual migration or bi-directional mutual migration.



#### Abbreviations

ARI: Adjusted Rand index; BSBM: Bipartite stochastic block model; DSBM: Directed stochastic block model; ER: Erdős-Rényi; MAM: Motif adjacency matrix

#### Acknowledgments

We would like to thank Lyuba Bozhilova for her feedback on this work. We would also like to thank Microsoft for their generous donation of Azure credits to The Alan Turing Institute which made this work possible.

#### Authors' contributions

WGU developed the methodology, wrote the R code, performed the bipartite synthetic experiments and analyzed the real world data. AE wrote the Python code, contributed to the sparse methodology, and performed the directed synthetic experiments. MC designed and supervised the study. All authors wrote, edited and approved the manuscript.

#### Funding

Andrew Elliott and Mihai Cucuringu acknowledge support from the EPSRC grant EP/N510129/1 at The Alan Turing Institute and Accenture PLC.

**Availability of data and materials**

All data sets are cited and are publicly available online. Source code in R and Python is available at <https://github.com/wgunderwood/motifcluster>.

**Competing interests**

The authors declare that they have no competing interests.

**Author details**

<sup>1</sup>Department of Statistics, University of Oxford, 24–29 St. Giles, Oxford, OX1 3LB, UK. <sup>2</sup>Department of Operations Research and Financial Engineering, Princeton University, Sherrerd Hall, Charlton Street, Princeton 08544, NJ, USA. <sup>3</sup>The Alan Turing Institute, British Library, 96 Euston Road, London NW1 2DB, UK. <sup>4</sup>School of Mathematics and Statistics, University of Glasgow, Glasgow, GL12 8QQ UK.

Received: 23 March 2020 Accepted: 26 July 2020

Published online: 10 September 2020

**References**

- Adamic LA, Glance N (2005) The political blogosphere and the 2004 US election: Divided they blog. In: Proc. of the 3rd Intl. Workshop on Link Discovery. ACM, New York. pp 36–43
- Ahmed NK, Neville J, Rossi RA, Duffield N (2015) Efficient graphlet counting for large networks. In: 2015 IEEE International Conference on Data Mining. IEEE, New York. pp 1–10
- Aicher C, Jacobs AZ, Clauset A (2013) Adapting the Stochastic Block Model to Edge-Weighted Networks. ArXiv preprint. <https://arxiv.org/abs/1305.5782>. Accessed 11 Feb 2020
- Aicher C, Jacobs AZ, Clauset A (2014) Learning latent block structure in weighted networks. *J Compl Netw* 3(2):221–248. <https://doi.org/10.1093/comnet/cnu026>
- Albert R (2005) Scale-free networks in cell biology. *J Cell Sci* 118(21):4947–4957
- Arthur D, Vassilvitskii S (2007) *K*-means++: The advantages of careful seeding. In: Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms. Society for Industrial and Applied Mathematics, Philadelphia. pp 1027–1035
- Benson AR, Abebe R, Schaub MT, Jadbabaie A, Kleinberg J (2018) Simplicial closure and higher-order link prediction. *Proc Natl Acad Sci* 115(48):11221–11230. <https://doi.org/10.1073/pnas.1800683115>
- Barabási A.-L., Albert R (1999) Emergence of scaling in random networks. *Science* 286(5439):509–512
- Benson AR, Gleich DF, Leskovec J (2016) Higher-order organization of complex networks. *Science* 353(6295):163–166
- Cheeger J (1969) A lower bound for the smallest eigenvalue of the Laplacian. In: Proceedings of the Princeton Conference in Honor of Professor S. Bochner. Princeton University Press, Princeton
- Chessa A, Crimaldi I, Riccaboni M, Trapin L (2014) Cluster analysis of weighted bipartite networks: a new copula-based approach. *PLoS ONE* 9(10):1–12
- Chung F (2005) Laplacians and the Cheeger inequality for directed graphs. *Ann Comb* 9(1):1–19
- Clauset A, Tucker E, Sainz M (2007) Filmtipset user movie ratings. *Colo Index Compl Netw*. <https://icon.colorado.edu/>. Accessed 15 Apr 2019
- Cucuringu M (2016) Sync-rank: Robust ranking, constrained ranking and rank aggregation via eigenvector and SDP synchronization. *IEEE Trans Netw Sci Eng* 3(1):58–79
- Cucuringu M, Davies P, Glielmo A, Tyagi H (2019a) SPONGE: A generalized eigenproblem for clustering signed networks. In: AISTATS 2019. PMLR
- Cucuringu M, Li H, Sun H, Zanetti L (2019b) Hermitian matrices for clustering directed graphs: insights and applications. ArXiv preprint. <https://arxiv.org/abs/1908.02096>. Accessed 19 Feb 2020
- Demeyer S, Michoel T, Fostier J, Audenaert P, Pickavet M, Demeester P (2013) The index-based subgraph matching algorithm (ISMA): Fast subgraph enumeration in large networks using optimized search trees. *PLoS ONE* 8(4):1–15. <https://doi.org/10.1371/journal.pone.0061183>
- Donath WE, Hoffman AJ (1972) Algorithms for partitioning of graphs and computer logic based on eigenvectors of connection matrices. *IBM Techn Discl Bull* 15(3):938–944
- Elliott A, Chiu A, Bazzi M, Reinert G, Cucuringu M (2019a) Core-Periphery Structure in Directed Networks. ArXiv preprint. <https://arxiv.org/abs/1912.00984>. Accessed 22 Mar 2020
- Elliott A, Cucuringu M, Luaces MM, Reidy P, Reinert G (2019b) Anomaly Detection in Networks with Application to Financial Transaction Networks. ArXiv preprint. <https://arxiv.org/abs/1901.00402>. Accessed 22 Mar 2020
- Erdős P, Rényi A, et al. (1959) On random graphs. *Publ Math* 6(26):290–297
- Florescu L, Perkins W (2016) Spectral thresholds in the bipartite stochastic block model. In: Conference on Learning Theory. PMLR. pp 943–959
- Fortunato S (2010) Community detection in graphs. *Phys Rep* 486(3–5):75–174
- Fortunato S, Hric D (2016) Community detection in networks: A user guide. *Phys Rep* 659:1–44
- Frey BJ, Dueck D (2007) Clustering by passing messages between data points. *Science* 315(5814):972–976
- GeoNames (2019) GeoNames. <https://www.geonames.org/>. Creative Commons, Accessed 24 Mar 2019
- Guattery S, Miller GL (1995) On the performance of spectral graph partitioning methods Vol. 95. pp 233–242
- Guattery S, Miller GL (1998) On the quality of spectral separators. *SIAM J Matrix Anal Appl* 19(3):701–719
- Huang J, Shen H, Hou L, Cheng X (2019) Signed graph attention networks. In: International Conference on Artificial Neural Networks. Springer, Cham. pp 566–577
- Hubert L, Arabie P (1985) Comparing partitions. *J Classif* 2(1):193–218
- Jacob P-M, Lapkin A (2018) Statistics of the network of organic chemistry. *React Chem Eng* 3(1):102–118
- Joseph A, Yu B, et al (2016) Impact of regularization on spectral clustering. *Ann Stat* 44(4):1765–1791
- Karrer B, Newman ME (2011) Stochastic blockmodels and community structure in networks. *Phys Rev E* 83(1):016107
- Kolaczyk ED, Csárdi G (2014) *Statistical Analysis of Network Data with R*, vol. 65. Springer, New York

- KONECT: The Koblenz Network Collection (2019) Unicode Languages network dataset. <http://konect.cc/networks/unicodelang>. Accessed 24 Mar 2019
- Leskovec J, Krevl A (2007) Astrophysics collaboration network, SNAP Datasets: Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data/ca-AstroPh.html>. Accessed 15 Apr 2019
- Leskovec J, Krevl A (2014) SNAP Datasets: Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data>. Accessed 21 Mar 2020
- Li GX (2017) Divided We Tweet: Community Detection in Political Networks. Final Report, Bachelor of Science in Engineering, Department of Engineering, Princeton University
- Li P, Milenkovic O (2017) Inhomogeneous hypergraph clustering with applications. In: Advances in Neural Information Processing Systems. Curran Associates, Inc., New York. pp 2308–2318
- Li P, Milenkovic O (2018) Submodular hypergraphs: p-laplacians, Cheeger inequalities and spectral clustering. ArXiv preprint. <https://arxiv.org/abs/1803.03833>. Accessed 24 June 2020
- Li P, Dau H, Puleo G, Milenkovic O (2016) Motif Clustering and Overlapping Clustering for Social Network Analysis. ArXiv preprint. <https://arxiv.org/abs/1612.00895>. Accessed 22 Mar 2020
- Lütkepohl H (1996) Handbook of Matrices, vol. 1. Wiley, Chichester
- Mangan S, Alon U (2003) Structure and function of the feed-forward loop network motif. *Proc Natl Acad Sci* 100(21):11980–11985. <https://doi.org/10.1073/pnas.2133841100>
- Mariadassou M, Robin S, Vacher C, et al (2010) Uncovering latent structure in valued graphs: a variational approach. *Ann Appl Stat* 4(2):715–742
- Meilä M, Pentney W (2007) Clustering by weighted cuts in directed graphs. In: Proceedings of the 2007 SIAM International Conference on Data Mining. SIAM, Philadelphia. pp 135–144
- Monti F, Otness K, Bronstein MM (2018) MotifNet: a motif-based Graph Convolutional Network for directed graphs. ArXiv preprint. <https://arxiv.org/abs/1802.01572>. Accessed 22 Mar 2020
- Mora BB, Cirtwill AR, Stouffer DB (2018) pymfinder: a tool for the motif analysis of binary and quantitative complex networks. bioRxiv. <https://doi.org/10.1101/364703>
- Nadler B, Lafon S, Kevrekidis I, Coifman RR (2006) Diffusion maps, spectral clustering and eigenfunctions of Fokker–Planck operators. In: Advances in Neural Information Processing Systems. MIT Press, Cambridge. pp 955–962
- Newman ME (2004) Analysis of weighted networks. *Phys Rev E* 70(5):056131
- Newman M (2008) The physics of networks. *Phys Today* 61(11):33–38
- Nguyen V, Epps J, Bailey J (2009) Information theoretic measures for clusterings comparison: is a correction for chance necessary? In: International Conference on Machine Learning 2009. Association for Computing Machinery (ACM), New York. pp 1073–1080
- Nowicki K, Snijders TAB (2001) Estimation and prediction for stochastic blockstructures. *J Am Stat Assoc* 96(455):1077–1087
- Onnela J-P, Saramäki J, Kertész J, Kaski K (2005) Intensity and coherence of motifs in weighted complex networks. *Phys Rev E* 71(6):065103
- Pashanasangi N, Seshadhri C (2020) Efficiently counting vertex orbits of all 5-vertex subgraphs, by EVOKE. In: Proceedings of the 13th International Conference on Web Search and Data Mining. ACM, New York. pp 447–455
- Qin T, Rohe K (2013) Regularized spectral clustering under the degree-corrected stochastic blockmodel. In: Advances in Neural Information Processing Systems. Curran Associates Inc., New York. pp 3120–3128
- Rand WM (1971) Objective criteria for the evaluation of clustering methods. *J Am Stat Assoc* 66(336):846–850
- Rohe K, Qin T, Yu B (2016) Co-clustering directed graphs to discover asymmetries and directional communities. *Proc Natl Acad Sci* 113(45):12679–12684
- Rosvall M, Esquivel AV, Lancichinetti A, West JD, Lambiotte R (2014) Memory in network flows and its effects on spreading dynamics and community detection. *Nat Commun* 5(1):1–13
- Satuluri V, Parthasarathy S (2011) Symmetrizations for clustering directed graphs. In: Proceedings of the 14th International Conference on Extending Database Technology. ACM, New York. pp 343–354
- Schaeffer SE (2007) Graph clustering. *Comput Sci Rev* 1(1):27–64
- Simmons BI, Sweering MJM, Schilling M, Dicks LV, Sutherland WJ, Clemente RD (2019) bmotif: A package for motif analyses of bipartite networks. *Methods Ecol Evol*. <https://doi.org/10.1111/2041-210X.13149>
- Shi J, Malik J (2000) Normalized cuts and image segmentation. *Departmental Papers (CIS)* 107:888–905
- Stewart GW, Sun J-G (1990) Matrix Perturbation Theory. Academic Press, Boston
- Stram R, Reuss P, Althoff K-D (2017) Weighted one mode projection of a bipartite graph as a local similarity measure. In: International Conference on Case-Based Reasoning. Springer, Cham. pp 375–389
- Strassen V (1969) Gaussian elimination is not optimal. *Numer Math* 13(4):354–356
- Tsourakakis CE, Pachocki J, Mitzenmacher M (2017) Scalable motif-aware graph clustering. In: Proc. of the 26th Intl. Conference on World Wide Web. International World Wide Web Conferences Steering Committee, Geneva. pp 1451–1460
- U.S. Census Bureau (2002) County-to-county migration flow files. <https://www.census.gov/population/www/cen2000/citytoctyflow/index.html>. Accessed 02 Mar 2019
- U.S. Census Bureau (2003) Domestic Migration Across Regions, Divisions, and States: 1995 to 2000. <https://www.census.gov/population/www/cen2000/migration>. Accessed 27 June 2020, Census 2000 Special Reports
- Veldt N, Benson AR, Kleinberg J (2020) Minimizing Localized Ratio Cut Objectives in Hypergraphs. ArXiv preprint. <https://arxiv.org/abs/2002.09441>. Accessed 06 July 2020
- Von Luxburg U (2007) A tutorial on spectral clustering. *Stat Comput* 17(4):395–416
- Von Luxburg U, Bousquet O, Belkin M (2004) On the convergence of spectral clustering on random samples: The normalized case. In: Learning Theory. Springer, Berlin. pp 457–471
- Wagner D, Wagner F (1993) Between min cut and graph bisection. In: International Symposium on Mathematical Foundations of Computer Science. Springer, Berlin, Heidelberg. pp 744–750
- Wang Y, Wang H, Zhang S (2018) A weighted higher-order network analysis of fine particulate matter (PM<sub>2.5</sub>) transport in Yangtze River Delta. *Physica A: Statistical Mechanics and its Applications* 496:654–662



- Wasserman S, Faust K, et al. (1994) *Social Network Analysis: Methods and Applications*, vol. 8. Cambridge university press, Cambridge
- Wegner AE, Ospina-Forero L, Gaunt RE, Deane CM, Reinert G (2018) Identifying networks with common organizational principles. *J Compl Netw* 6(6):887–913. <https://doi.org/10.1093/comnet/cny003>
- Wernicke S (2006) *IEEE/ACM Trans Comput Biol Bioinforma (TCBB)* 3(4):347–359. <https://doi.org/10.1109/TCBB.2006.51>
- Wernicke S, Rasche F (2006) FANMOD: A tool for fast network motif detection. *Bioinformatics* 22(9):1152–1153
- Yin H, Benson AR, Leskovec J, Gleich DF (2017) Local higher-order graph clustering. In: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, New York. pp 555–564
- Zha H, He X, Ding C, Simon H, Gu M (2001) Bipartite graph partitioning and data clustering. In: *Proceedings of the Tenth International Conference on Information and Knowledge Management*. ACM, New York. pp 25–32
- Zhang Y, Rohe K (2018) Understanding regularized spectral clustering via graph conductance. In: *Advances in Neural Information Processing Systems*. Curran Associates, Inc., New York. pp 10631–10640

### **Publisher's Note**

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Submit your manuscript to a SpringerOpen<sup>®</sup> journal and benefit from:**

- ▶ Convenient online submission
- ▶ Rigorous peer review
- ▶ Open access: articles freely available online
- ▶ High visibility within the field
- ▶ Retaining the copyright to your article

---

Submit your next manuscript at ▶ [springeropen.com](https://www.springeropen.com)

---