

Magnetic Eigenmaps for the Visualization of Directed Networks

Michaël Fanuel, Carlos M. Alaíz, Ángela Fernández, Johan A. K. Suykens

KU Leuven, Department of Electrical Engineering (ESAT), Kasteelpark Arenberg 10, B-3001 Leuven, Belgium

Abstract

We propose a framework for the visualization of directed networks relying on the eigenfunctions of the magnetic Laplacian, called here Magnetic Eigenmaps. The magnetic Laplacian is a complex deformation of the well-known combinatorial Laplacian. Features such as density of links and directionality patterns are revealed by plotting the phases of the first magnetic eigenvectors. An interpretation of the magnetic eigenvectors is given in connection with the angular synchronization problem. Illustrations of our method are given for both artificial and real networks.

Keywords: Magnetic Eigenmaps, Magnetic Laplacian, Directed Graph, Data Visualization

1. Introduction

Many problems in neuroscience, biology, social or computer science are phrased in terms of networks and graphs. The embedding of data points forming undirected graphs can be performed using manifold learning methods, among which are the so-called Laplacian Eigenmaps [1] and Diffusion Maps [2]. In the same spirit, the embedding of a directed graph originating from the sampling of a vector field on a manifold was studied in [3]. A Laplacian for strongly connected and aperiodic directed networks was introduced by Chung [4] in relation with a random walk process, which was used for visualization e.g. in [5]. Actually, Laplacians are very useful tools for community detection and data visualization. A common feature of these approaches is the relevance of the discrete or combinatorial Laplacian, and its normalized versions.

In this letter, no assumption on the origin of directed networks is needed, so we could deal, for example, with networks of webpages which are not embedded in any vector space. In particular, we propose here the use of another Laplacian which naturally exists for a general connected directed network, called the magnetic Laplacian. This operator is actually a vector bundle Laplacian as described in [6, 7] and a Connection Laplacian [8]. Interestingly, the magnetic Laplacian can be interpreted as a discrete quantum mechanical Hamiltonian of a charged particle on a network, influenced by a magnetic flux [9–11]. The method that we describe assigns a complex rotation, i.e., an element of $U(1)$, to each directed link, and the orientation of the link determines the direction of the rotation [12].

This letter is organized as follows. In Section 2 the magnetic Laplacian and its eigenvectors are introduced. A method using the complex phase of these eigenvectors for visualizing directed graphs is proposed in Section 3. Some examples are shown in Section 4, and the letter ends with some conclusions in Section 5.

2. Magnetic Laplacian and Eigenmaps

Consider a connected graph $\mathcal{G} = (V, E)$ with a set of N nodes V and a set of undirected edges E . In the case of an undirected graph, a symmetric weight matrix $W^{(s)}$ is given with elements $[W^{(s)}]_{ij} = w_{ij}^{(s)} \geq 0$ for all i and $j \in V$. The Laplacian Eigenmaps are the eigenvectors of the combinatorial Laplacian $L^{(0)} = D - W^{(s)}$, where D is the diagonal degree matrix with matrix elements $[D]_{ii} = d_i = \sum_{j \in V} w_{ij}^{(s)}$ for all $i \in V$. The volume of a subgraph \mathcal{S}_A of \mathcal{G} with node set A is simply $\text{vol}(\mathcal{S}_A) = \sum_{i \in A} d_i$. In the case of directed networks, the graph is given by an asymmetric weight matrix W with elements $[W]_{ij} = w_{ij} \geq 0$. For simplicity, the weights are chosen to be binary, i.e., $w_{ij} = 1$ if there is a link from i to j and $w_{ij} = 0$ otherwise. The weight

matrix W can be decomposed into a symmetric term $w_{ij}^{(s)} = (w_{ij} + w_{ji})/2$, indicating that $\{i, j\} \in E$, and a skew-symmetric term, the edge flow $a_{ij} = -a_{ji}$ encoding the direction of the link. For all $\{i, j\} \in E$, we have $a_{ij} = 1$ if the link points from i to j , and $a_{ij} = 0$ if $\{i, j\}$ is not directed.

In this letter, the magnetic Laplacian is defined as the self-adjoint, positive semi-definite operator $L^{(g)} = D - T^{(g)} \odot W^{(s)}$, where D is the degree matrix associated with the symmetrized weight matrix, $0 \leq g < 1/2$ is an electric charge parameter, and $[T^{(g)} \odot W^{(s)}]_{ij} = \exp(i2\pi g a_{ji}) w_{ij}^{(s)}$ (notice the Hadamard product \odot). The solutions of the generalized eigenvalue problem $L^{(g)}\phi = \lambda D\phi$ are the *Magnetic Eigenmaps* $\phi_k^{(g)}$ associated with the eigenvalues $\lambda_k^{(g)} \geq 0$ for $k \in \{0, \dots, N-1\}$ [13] (we assume $\lambda_0^{(g)} \leq \lambda_1^{(g)} \leq \dots \leq \lambda_{N-1}^{(g)}$).

2.1. Interpretation of the first eigenvectors

While the calculation of the *second* eigenvector of the normalized combinatorial Laplacian is a relaxation of the (normalized) cut problem, the calculation of the *first* eigenvector of the normalized magnetic Laplacian is a relaxation of the angular synchronization problem [14]. Given a subgraph \mathcal{S} of \mathcal{G} (in general, one can choose $\mathcal{S} = \mathcal{G}$), the angular synchronization problem consists in finding the angles $\boldsymbol{\theta}^* = (\theta_1^*, \dots, \theta_N^*)^\top \in \mathbb{U}(1)^N$ given by $\boldsymbol{\theta}^* \in \arg \min_{\boldsymbol{\theta}} \eta_{\mathcal{S}}(\boldsymbol{\theta})$ where the frustration [15] is defined by

$$\eta_{\mathcal{S}}(\boldsymbol{\theta}) = \frac{1}{2} \frac{\sum_{i,j \in \mathcal{S}} w_{ij}^{(s)} |e^{i\theta_i} - e^{i\theta_{ij}} e^{i\theta_j}|^2}{\sum_{i \in V} d_i},$$

with $\theta_{ij} = 2\pi g a_{ji}$ for all $i, j \in \mathcal{S}$ such that $w_{ij}^{(s)} \neq 0$. Notice that $\sum_{i \in V} d_i = \text{vol}(\mathcal{G})$. The lowest eigenvector of the normalized magnetic Laplacian $\phi_0^{(g)}$ is the solution of the spectral problem relaxing $\min_{\boldsymbol{\theta}} \eta_{\mathcal{S}}(\boldsymbol{\theta})$. Our first conclusion is that computing the complex phase of $\phi_0^{(g)}$ yields an approximation of $\boldsymbol{\theta}^*$ that we propose to choose as the first visualization coordinate. In [12], the solution of the angular synchronization problem is shown to provide a ranking of the nodes in directed graphs, although a slightly different eigenvector problem is considered.

The performance of the spectral relaxation of the cut problem can be studied using a classical result of spectral graph theory, the Cheeger inequality, which relates the Cheeger constant to the *second* smallest eigenvalue of the combinatorial Laplacian, providing the worst case performance for the spectral clustering method. Analogous results relate the *first* smallest eigenvalue of the Connection Laplacian [15] and the magnetic Laplacian [16] to a frustration quantifying the amount of inconsistency in the connection graph. In particular, the performance naturally depends on the inverse of the spectral gap $1/\lambda_1^{(0)}$ of the undirected measurement graph. Indeed, the quality of the synchronization will benefit from a good connectivity of the nodes in the measurement graph. On the contrary, if $\lambda_1^{(0)}$ is small, it could be instructive to find the subgraphs of \mathcal{G} where the frustration is minimal by cutting edges where the angular synchronization is not accurate. Suppose that $A \subset V$ is the vertex set of a subgraph \mathcal{S}_A of \mathcal{G} and let \bar{A} be its complement in V . A combinatorial graph partitioning problem is proposed, that is, $\min_A E_{A, \bar{A}}(\boldsymbol{\theta}^*)$ with

$$E_{A, \bar{A}}(\boldsymbol{\theta}^*) = \frac{\text{vol}(\mathcal{S}_{\bar{A}})}{\text{vol}(\mathcal{G})} \eta_A(\boldsymbol{\theta}^*) + \frac{\text{vol}(\mathcal{S}_A)}{\text{vol}(\mathcal{G})} \eta_{\bar{A}}(\boldsymbol{\theta}^*) + \left(\frac{c_{A, \bar{A}}}{\text{vol}(\mathcal{S}_A)} + \frac{c_{\bar{A}, A}}{\text{vol}(\mathcal{S}_{\bar{A}})} \right) + \frac{\gamma_{A, \bar{A}}(\boldsymbol{\theta}^*)}{\text{vol}(\mathcal{G})}, \quad (1)$$

where $\text{vol}(\mathcal{S}_A) = \sum_{i \in A} d_i$ is the volume of \mathcal{S}_A , whereas the cut is $c_{A, \bar{A}} = c_{\bar{A}, A} = \sum_{i \in A, j \in \bar{A}} w_{ij}^{(s)}$. We have defined the generalized cut

$$\gamma_{A, \bar{A}}(\boldsymbol{\theta}^*) = -4 \sum_{i \in A, j \in \bar{A}} w_{ij}^{(s)} \sin^2 \left(\frac{\theta_i^* - \theta_j^* - \theta_{ij}}{2} \right). \quad (2)$$

Notice that each term of (2) is minimized when $|\theta_i^* - \theta_j^* - \theta_{ij}| = \pi$, i.e., when the error made in the synchronization of rotations along the link $\{i, j\}$ is large. Problem (1) generalizes the normalized cut problem aiming to find A and \bar{A} so that a combination of the frustration of both subgraphs, the normalized

cut and the generalized cut (2) is minimal, while the partition is balanced. In order to construct a relaxation of this problem, we define

$$(f_{A,\bar{A}})_i = \begin{cases} \sqrt{\text{vol}(\mathcal{S}_{\bar{A}})/\text{vol}(\mathcal{S}_A)} e^{i\theta_i^*} & \text{if } i \in A, \\ -\sqrt{\text{vol}(\mathcal{S}_A)/\text{vol}(\mathcal{S}_{\bar{A}})} e^{i\theta_i^*} & \text{if } i \in \bar{A}. \end{cases} \quad (3)$$

Then, we have $E_{A,\bar{A}}(\theta^*) = (f_{A,\bar{A}}^\dagger L^{(g)} f_{A,\bar{A}}) / (f_{A,\bar{A}}^\dagger D f_{A,\bar{A}})$ as well as the relations $\sum_{i \in V} d_i e^{-i\theta_i^*} (f_{A,\bar{A}})_i = 0$ and $\sum_{i \in V} d_i |(f_{A,\bar{A}})_i|^2 = \text{vol}(\mathcal{G})$. As a consequence, a spectral relaxation of the combinatorial problem is given by

$$\min_{f \in \mathbb{C}_0^N} \frac{f^\dagger L^{(g)} f}{f^\dagger D f} \text{ s.t. } f^\dagger D \phi_0^{(g)} = 0, \quad (4)$$

which corresponds to the generalized eigenvalue problem for the second least eigenvector $\phi_1^{(g)}$. In view of (3), we conclude that $\text{phase}(\phi_0^{(g)})$ will instruct us about the partition of the graph minimizing (1). Indeed, we have $\text{phase}((f_{A,\bar{A}})_i) = \theta_i^*$ if $i \in A$ and $\text{phase}((f_{A,\bar{A}})_i) = \theta_i^* + \pi$ if $i \in \bar{A}$. The eigenvector $\phi_1^{(g)}$ being the relaxed version of $(f_{A,\bar{A}})_i$, we expect that, by embedding the graph with $\text{phase}(\phi_0^{(g)})$ as first coordinate and $\text{phase}(\phi_1^{(g)})$ as second coordinate, we obtain two parallel groups distant by π . By analogy with the Laplacian eigenmaps, similar results are expected for the next eigenvalues.

2.2. Relation with the combinatorial Laplacian

When the edge flow of the network is given exactly by a certain potential h , i.e., $a_{ij} = h_j - h_i$ for all $\{i, j\} \in E$, the spectrum of the magnetic Laplacian corresponds to the spectrum of the combinatorial Laplacian, and the eigenvectors are related by $\phi_{k,i}^{(g)} = \exp(i2\pi g h_i) \phi_{k,i}^{(0)}$ for all $i \in V$ and all $k \in \{0, \dots, N-1\}$. This particular case is characterized by the first eigenvalue, as stated in the following proposition, which can be seen as a consequence of [15, Theorem 2.6] and of [10, 11] in the context of mathematical physics.

Proposition 1. *Consider a connected graph \mathcal{G} . The magnetic Laplacian $L^{(g)}$ has a zero eigenvalue iff there exists a function h satisfying, for any link $\{i, j\} \in E$, $a_{ij} = h_j - h_i$.*

Nonetheless, since we assume here that $a_{ij} \in \{-1, 0, 1\}$ for all $\{i, j\} \in E$, this situation only happens in our context if the graph is a tree. The relation in the general case can be characterized by Lemma 1. First let us define the parallel transporter $t_P^{(g)} \in \text{U}(1)$ over a finite path $P = \{i_1, i_2, \dots, i_n\}$ of length $n > 1$ in the graph \mathcal{G} as $t_P^{(g)} = \exp(i2\pi g (a_{i_1 i_2} + \dots + a_{i_{n-1} i_n}))$. Moreover, following [17] we call a directed graph ϵ -consistent if, for every simple cycle $\{i_1, i_2, \dots, i_n, i_{n+1} = i_1\}$, we have $|t_C^{(g)} - 1| \leq \epsilon$, with $t_C^{(g)} = \exp(i2\pi g \oint_C a)$, where the magnetic flux is defined by the discrete line integral $\oint_C a \triangleq a_{i_1 i_2} + \dots + a_{i_n i_1}$. We now state an elementary result given in [11].

Lemma 1. *Consider a connected directed graph \mathcal{G} , and let T be any spanning tree of \mathcal{G} , and $\bar{T} = E \setminus T$. For all $0 \leq g \leq 1/2$, the magnetic Laplacian $L^{(g)}$ is unitarily equivalent to the operator given by*

$$\left(\tilde{L}_{(T)}^{(g)} f \right)_i = \sum_{\{j|\{i,j\} \in T\}} w_{ij}^{(s)} (f_i - f_j) + \sum_{\{j|\{i,j\} \in \bar{T}\}} w_{ij}^{(s)} \left(f_i - t_{\circ,ij}^{(g)} f_j \right),$$

for all $f_i \in \mathbb{C}$ and all $i \in V$. The operator $t_{\circ,ij}^{(g)}$ is defined as $t_{\circ,ij}^{(g)} = t_{i_0 \rightarrow i}^{(g)} t_{ij}^{(g)} t_{j \rightarrow i_0}^{(g)}$ for any $i_0 \in V$, and $i_0 \rightarrow k$ denotes the unique path in the tree T going from $i_0 \in T$ to $k \in T$. The unitary transformation, realizing $\tilde{L}_{(T)}^{(g)} = U_{(T)}^\dagger L^{(g)} U_{(T)}$, is given by the diagonal matrix $U_{(T)}$, with elements $[U_{(T)}]_{ii} = t_{i_0 \rightarrow i}^{(g)}$ for all $i \in T$.

Notice that the accumulated complex phase $t_{\circ,ij}^{(g)}$ (holonomy or magnetic flux) in the contour integral does not depend on the choice of i_0 . Its value is equal to the contour integral on the smallest loop including the link $\{i, j\}$. Moreover, as a corollary, we have that the normalized magnetic Laplacian $L_N^{(g)} = D^{-1/2} L^{(g)} D^{-1/2}$ is unitarily equivalent to $D^{-1/2} \tilde{L}_{(T)}^{(g)} D^{-1/2}$. Hence, if the graph \mathcal{G} is a tree, the spectrum of the magnetic Laplacian is exactly the spectrum of the combinatorial Laplacian and their eigenvectors are in one to one correspondence, thanks to diagonal unitary transformation.




Figure 1: (a): Representation of the 2-torus as the square $[0, 2\pi] \times [0, 2\pi]$ with identified sides; the position of the cuts is arbitrary and can be adapted to each particular dataset. (b): Example of the magnetic eigenmaps plotted over the 3-dimensional 2-torus for a graph with two communities, corresponding to the dataset “Political Blogosphere” explained in detail in Section 4.

3. Visualization of density and directionality

The eigenmaps can be calculated by computing the eigenvectors of the normalized magnetic Laplacian $L_N^{(g)} = D^{-1/2} L^{(g)} D^{-1/2}$ and, therefore, these eigenvectors are obtained by calculating the largest eigenvalues of $D^{-1/2} T^{(g)} \odot W^{(s)} D^{-1/2}$. Once the Laplacian is normalized, we propose to embed the network by the mapping $i \mapsto (\text{phase}(\phi_{0,i}^{(g)}), \text{phase}(\phi_{1,i}^{(g)}), \dots, \text{phase}(\phi_{n,i}^{(g)}))^\top$. Notice that the phase operator identifies the angles that differ by 2π , which means that the geometrical representation, for the 1-dimensional case, is just a circle, whereas for the general n -dimensional case is an n -torus. Hence, for visualization purposes, the directed network will be embedded on a 2-torus represented as the square $[0, 2\pi] \times [0, 2\pi]$ with opposite sides identified, as shown in Figure 1. Therefore, the visualization will be symmetric if an eigenvector undergoes a global rotation in the complex plane. Notice that given an eigenvector $\phi_k^{(g)}$ of the magnetic Laplacian, another eigenvector of the same eigenvalue is given by $e^{i\alpha} \phi_k^{(g)}$. We will show empirically that this low dimensional embedding is able to visualize at the same time dense regions of links, revealed by the y -axis, and patterns determined by the link directions, given by the x -axis.

3.1. Consistency of the mapping

The visualization method relies only on the phases of the magnetic eigenmaps, which contain most of the information if the magnetic Laplacian is unitarily equivalent to the combinatorial Laplacian. We will characterize now the error made if this is not the case. A bound on the variance of $|\phi_0^{(g)}|$ is obtained from [15, Lemma 3.3] that is adapted below to the setting of this paper. For all $z \in \mathbb{C}^N$, we define $\tilde{z} \in \mathbb{C}^N$ such that $\tilde{z}_i = z_i/|z_i|$ if $z_i \neq 0$ and $\tilde{z}_i = 0$ if $z_i = 0$.

Lemma 2 (Bandeira *et al.* [15, Lemma 3.3]). *For all $z \in \mathbb{C}^N$, we have $\sum_i d_i |z_i - \mu \tilde{z}_i|^2 \leq \eta(z) \sum_i d_i |z_i|^2 / \lambda_1^{(0)}$, with $\mu = (1/\text{vol}(\mathcal{G})) \sum_j d_j |z_j|$ and $\eta(z) = z^\dagger L^{(g)} z / (z^\dagger D z)$.*

Hence, by taking $z = \phi_0^{(g)}$, we have a bound on the variability of $|\phi_0^{(g)}|$. More precisely, the variation $|\phi_0^{(g)}|$ with respect to its mean value is constrained as follows:

$$\frac{\sum_i d_i \left| |\phi_{0,i}^{(g)}| - \mu_0 \right|^2}{\sum_i d_i |\phi_{0,i}^{(g)}|^2} \leq \frac{\lambda_0^{(g)}}{\lambda_1^{(0)}}, \quad \text{with } \mu_0 = \frac{1}{\text{vol}(\mathcal{G})} \sum_j d_j |\phi_{0,j}^{(g)}|. \quad (5)$$

Therefore, since the spectral gap $\lambda_1^{(0)}$ is only determined by the density of the undirected graph, the previous bound on the variability of the modulus $|\phi_0^{(g)}|$ can only be reduced if the eigenvalue $\lambda_0^{(g)}$ is made smaller. In particular, Theorem 1 in [17] can be adapted to the case of the magnetic Laplacian, so that the smallest eigenvalue of the magnetic Laplacian of an ϵ -consistent graph satisfies $\lambda_0^{(g)} \leq \epsilon^2/2$. In our particular case,

this bound can be improved by relating the inconsistency to a topological property of the graph leading to possible inconsistencies: the first Betti number $\beta_1 = |E| - |V| + 1$, i.e., the number of simple cycles in the graph. Notice that there are exactly β_1 edges to remove to the graph in order to have a tree.

Lemma 3. *Consider a connected directed graph \mathcal{G} , and let T be any spanning tree of \mathcal{G} , and $\bar{T} = E \setminus T$. We have the following inequality*

$$\lambda_0^{(g)} \leq \epsilon_g^2 \frac{\sum_{\{i,j\} \in \bar{T}} w_{ij}^{(s)}}{\text{vol}(\mathcal{G})}, \quad (6)$$

where the summation includes only β_1 terms (associated with the number of simple cycles). In particular, if the weights are binary $\lambda_0^{(g)} \leq \frac{\epsilon_g^2 \beta_1}{2|E|}$.

Proof. Fix a spanning tree T of \mathcal{G} . Using Lemma 1, we can consider the smallest eigenvalue of the unitary equivalent operator $\tilde{L}_{(T)}^{(g)}$, which is the minimum of the Rayleigh quotient $R(f) = \sum_i f_i^\dagger (\tilde{L}_{(T)}^{(g)} f)_i / \sum_i d_i |f_i|^2$. Specifically, choosing $f_i = 1$ for all $i \in V$, we have

$$\lambda_0^{(g)} \leq R(f) = \frac{\sum_{\{i,j\} \in \bar{T}} w_{ij}^{(s)} |f_i - t_{o,ij}^{(g)} f_j|^2}{\sum_{i \in V} d_i} = \frac{\sum_{\{i,j\} \in \bar{T}} w_{ij}^{(s)} |1 - t_{o,ij}^{(g)}|^2}{\text{vol}(\mathcal{G})} \leq \epsilon_g^2 \frac{\sum_{\{i,j\} \in \bar{T}} w_{ij}^{(s)}}{\text{vol}(\mathcal{G})}.$$

□

The bound given in (5) suggests to reduce the variability of the modulus (the information lost when considering only the phases) by making $\lambda_0^{(g)}$ smaller. According to Lemma 3, this can be done looking for the g that minimizes the inconsistency, which is a task dependent on each particular graph. Moreover, the results above are only upper bounds, so there is no guarantee that such a value of g is the best possible choice. Finally, notice that we are interested in maximizing the information included in the phase of the eigenvector, not only in minimizing the information lost while ignoring the modulus. As a trivial counterexample, if we take $g \rightarrow 0$ then we will recover the combinatorial Laplacian, and hence $\lambda_0^{(g)} \rightarrow 0$. Although in that case the eigenvector will be constant in modulus, and no information will be lost, its phase will also be constant, and thus it will provide no information.

3.2. Selection of g

The choice of the electric charge parameter g influences the visualization method. As stated above, there is not an established method to select it. In general, we propose to choose a quantized charge $g = k/m$ with $k \notin m\mathbb{Z}$. The particular value $g = 1/3$ is suited in the presence of directed triangles, while $g = 1/4$ is relevant in the presence of directed 4-cycles. For $g = 2/5$, the magnetic Laplacian becomes quite similar with the signed Laplacian, while for $g = 1/2$ it is a signed Laplacian associated with the same graph where undirected edges are labelled as (+) and directed edges as (-). For more details we refer to [13]. Notice that choosing $g > 1/2$ would be equivalent to a flip of all link directions.

3.3. Connection with Vector Diffusion Maps

The computation of the eigenvectors of the normalized magnetic Laplacian can resemble the Vector Diffusion Maps of Singer and Wu [8], although in our case we work with a complex and unitary transporter in $U(1)$, instead of with an orthogonal transporter. In particular, the main difference between both approaches, apart from the working space (\mathbb{R}^n and \mathbb{C} , respectively), resides in the transport term. In the case of Vector Diffusion Maps, it is an element of $SO(n)$ determined by Local PCA, and for Magnetic Eigenmaps it is an element of $U(1)$ which can be tuned by the user through the parameter g in order to highlight certain properties of the graph. Moreover, the methodology of both approaches differs in the way they map the data. On the one hand, Vector Diffusion Maps follows a natural extension of Diffusion Maps [2], so that each point is mapped to a matrix defined in terms of the eigenvalues and eigenvectors of the transition matrix. On the other hand, the proposed magnetic eigenmaps map the points to the phases of the first eigenvectors of the corresponding Laplacian. Therefore, although both methods share some similarities, they are essentially different, and none of them can be seen as a particular case of each other.

Algorithm 1 Magnetic Eigenmaps Visualization

procedure MEIGENMAPS(W, g) $W^{(s)} \leftarrow (W + W^\top) / 2$ $A \leftarrow W - W^\top$ $d_{ii} \leftarrow \sum_j w_{ij}^{(s)}$ $t_{ij}^{(g)} \leftarrow e^{i2\pi g a_{ji}}$ $L^{(g)} \leftarrow D - W^{(s)} \odot T^{(g)}$ $L_N^{(g)} \leftarrow D^{-1/2} L^{(g)} D^{-1/2}$ $\phi_0^{(g)}, \phi_1^{(g)}, \dots \leftarrow \text{EIGS}(L_N^{(g)})$ return $\left\{ \text{phase}(\phi_m^{(g)}) \right\}_{m=0}^n$	<ul style="list-style-type: none"> ▶ Symmetric weights. ▶ Edge flow. ▶ Degree matrix. ▶ Transporter. ▶ Magnetic Laplacian. ▶ Normalized Laplacian. ▶ Eigenvectors. ▶ Phases.
--	--

4. Applications

We will illustrate now how Magnetic Eigenmaps can be applied to the visualization of directed graphs on two synthetic and two real networks. For completeness, we include the procedure to compute the magnetic eigenmaps for a directed graph with binary weights in Algorithm 1. In all the examples we will depict the aspect of the original network as a baseline, using for this purpose expert knowledge or a force-directed layout (based on using attractive forces between adjacent nodes and repulsive forces between distant nodes). We will also compare with Diffusion Maps for the real examples. In this context, the embedding is obtained by computing the algorithm with $g = 0$, and plotting the first eigenvectors of the corresponding Laplacian (instead of their phases). For the real networks we will also depict the spectrum.

4.1. Artificial networks

We propose to visualize first the artificial network with a running flow of Figure 2a, where the coordinates of the nodes in the real plane have been chosen according to our knowledge about the underlying groups. This network, constructed according to [18], consists of three groups of ten nodes (A , B and C). Two nodes in the same group are linked with a probability 0.5. Any node has also a probability 0.5 to be connected to a node from another group. Furthermore 90 percent of these interconnections are directed in the direction of the flow, i.e., $A \rightarrow B$, $B \rightarrow C$ and $C \rightarrow A$. Plotting the real and imaginary parts of the first eigenvector of the magnetic Laplacian can indicate the presence of a running flow in the network, as illustrated in Figure 2b. However there could also be dense clusters in the network that are not revealed using only the phase of the first eigenfunction. In order to actually visualize the three groups and the density information, we use our proposed method of plotting the complex phase of the first eigenfunction versus the phase of the second eigenfunction of the magnetic Laplacian, as illustrated in Figure 2c. Notice that the phase of the second eigenfunction does not distinguish specific dense clusters in the network, while the phase of the first eigenfunction (corresponding to directionality) is able to separate the three groups.

We are going to consider now an example of a network with a small number of nodes playing a particular role and then a clear structure with two dense clusters. In [19], an artificial network of 32 nodes is built as follows: it consists of two dense groups of 14 nodes with a few interconnecting links and two pairs of nodes are connected to the whole network. The first pair has only in-coming links, while the second pair has only out-going links. An illustration using the force layout is given in Figure 3a. Plotting the real and imaginary parts of the first eigenfunction allows to distinguish the two pairs from the rest of the network, as showed in Figure 3b. However, it is more instructive to visualize the network using the phase of the two first eigenfunctions of the magnetic Laplacian. Indeed, the two groups and the two pairs are easily separated in Figure 3c, where the phase of the first eigenfunction (directionality) is able to separate the two pairs of disconnected points from the rest of the set, defining three directionality-groups: green points, the yellow points and the blue and red points together. On the other side, the phase of the second eigenfunction (corresponding to density information) shows two groups: the blue and green points versus the red and yellow points. Combining the information given by the two phases we are able to easily separate visually the four groups that we were looking for.

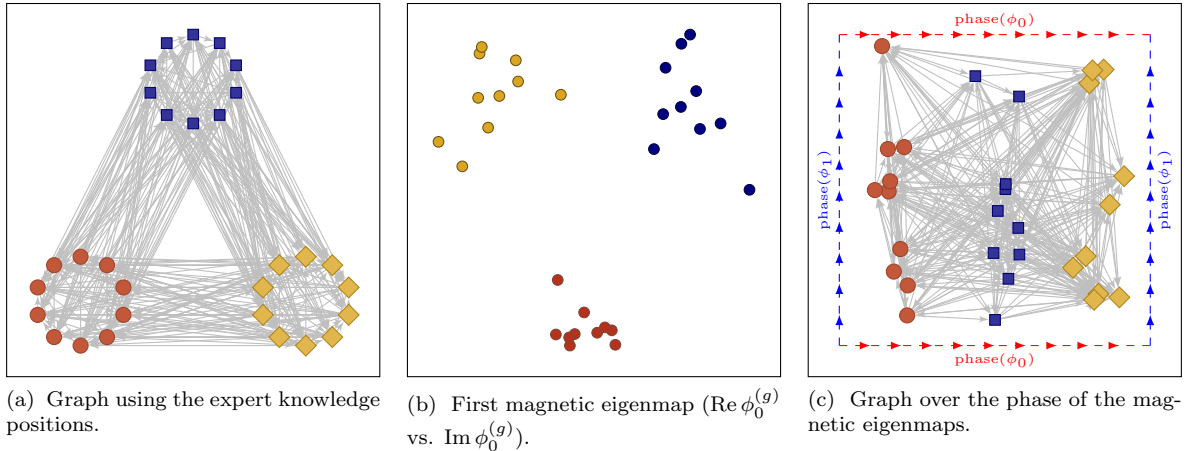


Figure 2: Artificial network with a running flow. The colours indicate the three groups, and the magnetic eigenmaps correspond to $g = 1/4$.

4.2. Directed networks from real data

In the previous section, we considered directed networks having known structures either in terms of link directions or link density. Indeed, Magnetic Eigenmaps is able to provide simultaneously information about direction and density, as we illustrate now also on real directed networks where these two aspects are relevant. Let us emphasize that we do not pretend that Magnetic Eigenmaps will give the best possible result on each dataset. Knowing the nature of the datasets, an *ad hoc* visualization method will certainly give a better result. We will however assume that the origin of the dataset is unknown in order to show that Magnetic Eigenmaps indeed reveals relevant features.

The network used represents the *common adjective and noun adjacencies* for the novel “David Copperfield” by Charles Dickens [20]. This directed graph has 112 nodes that represent the most commonly occurring adjectives and nouns in the book. Edges connect any pair of words that occur in adjacent position in the text of the book. From the structure of English, a certain directional structure can be anticipated, i.e., adjectives are expected to be found before nouns. The graph representation of this dataset, using a force layout, is shown in Figure 4a, where the structure can hardly be guessed. However, considering the phase of the two first magnetic eigenmaps (we have used $g = 2/5$ since the graph is almost bipartite), it is possible to visualize the presence of two groups, as illustrated in Figure 4b (indeed, the information is provided mostly by the first coordinate, corresponding to directionality). Finally, we show in Figure 4c the first embedding coordinates in the Diffusion Maps case (using the normalized Laplacian), more concretely the second and third eigenvectors (the first one is discarded because it is constant). In this case both classes appeared mixed, so these two diffusion coordinates are not able to reveal the structure of the data. The reason is that Diffusion Maps captures the density of the links but not the directionality, while Magnetic Eigenmaps relates both. Additionally, Figure 5a represents the first eigenvalues of the combinatorial ($g = 0$) and magnetic Laplacian ($g = 2/5$). The increase of these eigenvalues is smooth, presenting just an eigengap between the first eigenvalue and the second one. This remark motivates the choice of the first and second eigenvectors for this example as visualization coordinates. Moreover, the differences between the two spectra, and the nonzero initial eigenvalue, show that the structure of the graph is not trivial (see Proposition 1).

The second real dataset used in these experiments represents the *political blogosphere* in February of 2005 [21]. This directed graph is composed of 1222 nodes that indicate the political leaning, meaning left or liberal and right or conservative (disconnected points were removed). The data on political leaning comes from blog directories and some of the blogs were labelled manually, based on incoming and outgoing links and posts around the time of the 2004 presidential election in the USA. The links between blogs were automatically extracted from a crawl of the front page of the blog. From Figure 6a, where the network

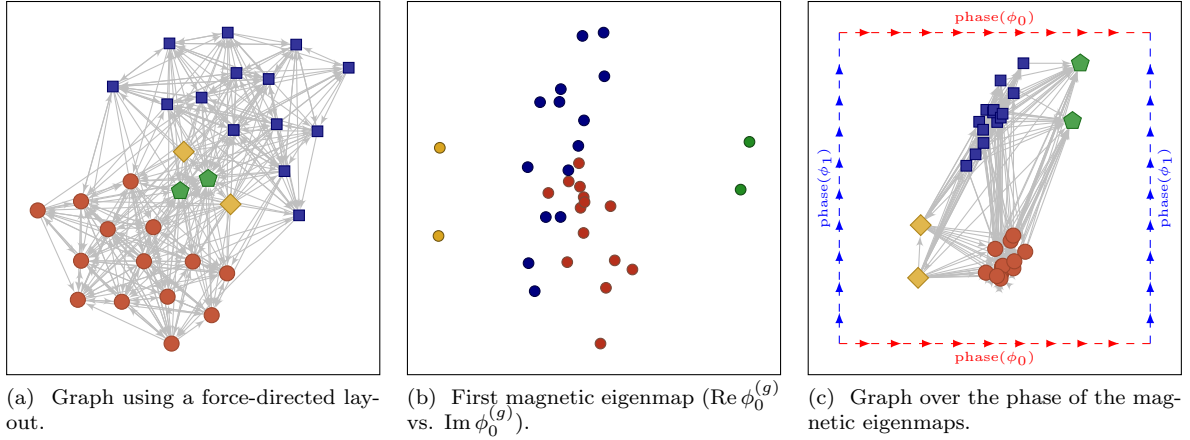


Figure 3: Artificial network with two dense clusters and two pairs of nodes with a specific role. The colours indicate the dense clusters and the hub pairs, and the magnetic eigenmaps correspond to $g = 1/4$.

has been depicted using the force layout, it is already possible to guess the presence of two dense groups of webpages. The first eigenvalues of the magnetic Laplacian in Figure 5b instruct us to consider the two first pairs of eigenvalues in order to visualize two different structures. In Figures 6b and 6c, the magnetic eigenmaps do not distinguish the two classes of nodes, however we observe that some webpages are less connected to the rest of the network, whereas in Figure 6d we see the two classes clearly separated. The latter mapping is also shown in \mathbb{R}^3 over the torus as an illustration in Figure 1b. For comparison purposes, we show also in this case the first embedding diffusion coordinates in Figures 6e and 6f. In this example, where the graph structure is clearer than in the previous dataset, Diffusion Maps is able to condense the two classes separately just using the density of the graph. We would like to highlight that the distinction between both classes is not very clear when we select the second and third eigenvectors (the first eigenvector is again discarded), whereas a cleaner classification structure is obtained when the fourth and fifth eigenvectors are depicted. Nevertheless, Magnetic Eigenmaps represents well both the connectivity and density structure.

5. Conclusions

In this letter, we have proposed the use of the eigenvectors of the magnetic Laplacian, called here *Magnetic Eigenmaps*, for the visualization of directed networks. Our work is a natural extension of the Laplacian Eigenmaps. Computationally, the method reduces to the calculation of the eigenvectors of maximal eigenvalues of a Hermitian matrix, which can be conveniently performed thanks to e.g. the power method. The advantages of this approach were illustrated on artificial and real datasets, showing that our method is able to reveal both the directionality and connectivity patterns of the networks.

Acknowledgements

The authors thank the following organizations. • EU: The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) / ERC AdG A-DATADRIVE-B (290923). This paper reflects only the authors’ views, the Union is not liable for any use that may be made of the contained information. • Research Council KUL: GOA/10/09 MaNet, CoE PFV/10/002 (OPTEC), BIL12/11T; PhD/Postdoc grants. • Flemish Government: – FWO: G.0377.12 (Structured systems), G.088114N (Tensor based data similarity); PhD/Postdoc grants. – IWT: SBO POM (100031); PhD/Postdoc grants. • iMinds Medical Information Technologies SBO 2014. • Belgian Federal Science Policy Office: IUAP P7/19 (DYSCO, Dynamical systems, control and optimization, 2012-2017).

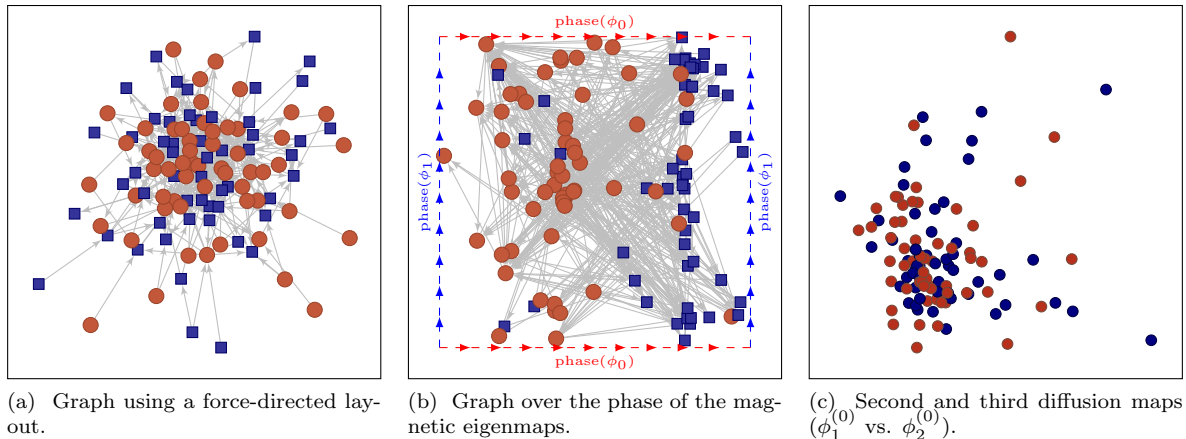


Figure 4: Word adjacencies example. The colours indicate the class labels, nouns (\bullet) and adjectives (\bullet), and the magnetic eigenmaps correspond to $g = 2/5$.




Figure 5: First eigenvalues of the combinatorial and normalized magnetic Laplacian for the real networks.

References

- [1] M. Belkin, P. Niyogi, Laplacian Eigenmaps for dimensionality reduction and data representation, *Neural Comput.* 15 (6) (2003) 1373–1396.
- [2] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, F. Warner, S. Zucker, Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps, in: *Proceedings of the National Academy of Sciences*, 2005, pp. 7426–7431.
- [3] D. C. Perrault-Joncas, M. Meila, Directed graph embedding: an algorithm based on continuous limits of Laplacian-type operators, in: *Advances in Neural Information Processing Systems*, 2011, pp. 990–998.
- [4] F. Chung, Laplacians and the Cheeger inequality for directed graphs, *Annals of Combinatorics* 9 (2005) 1–19.
- [5] Q. Zheng, D. B. Skillicorn, Spectral embedding of directed networks, in: *2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 2015, pp. 432–439.
- [6] R. Kenyon, Spanning forests and the vector bundle Laplacian, *The Annals of Probability* 39 (5) (2011) 1983–2017.
- [7] R. Forman, Determinants of Laplacians on graphs, *Topology* 32 (1) (1993) 35–46.
- [8] A. Singer, H. Wu, Vector Diffusion Maps and the connection Laplacian, *Commun Pure Appl Math.* 65 (8) (2012) 1067–1144.
- [9] M. A. Shubin, Discrete magnetic Laplacian, *Comm. Math. Phys.* 164 (1994) 259–275.
- [10] Y. Colin de Verdière, Magnetic interpretation of the nodal defect on graphs, *Analysis and PDE* 6 (5) (2013) 1235–1242.
- [11] G. Berkolaiko, Nodal count of graph eigenfunctions via magnetic perturbations, *Analysis and PDE* 6 (5) (2013) 1213–1233.
- [12] M. Cucuringu, Sync-Rank: Robust ranking, constrained ranking and rank aggregation via eigenvector and SDP synchronization, *IEEE Transactions on Network Science and Engineering* 3 (1) (2016) 58–79. doi:10.1109/TNSE.2016.2523761.
- [13] M. Fanuel, C. M. Alaíz, J. A. K. Suykens, Magnetic eigenmaps for community detection in directed networks, *ArXiv e-prints* arXiv:1606.07359.
- [14] A. Singer, Angular synchronization by eigenvectors and semidefinite programming, *Applied and computational harmonic analysis* 30 (1) (2011) 20–36.
- [15] A. S. Bandeira, A. Singer, D. A. Spielman, A Cheeger inequality for the graph connection Laplacian, *SIAM Journal on Matrix Analysis and Applications (SIMAX)* 34 (4) (2013) 1611–1630.

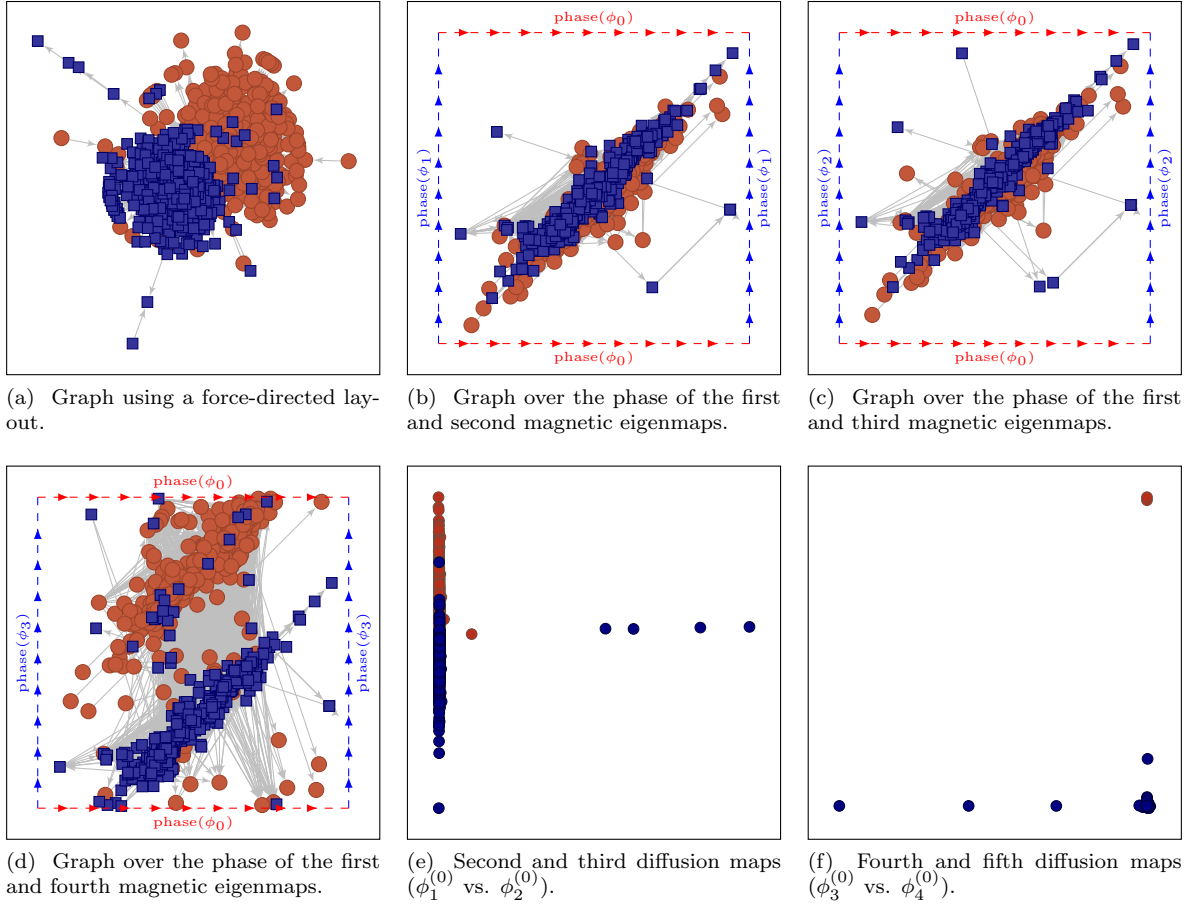


Figure 6: Political blogosphere example. The colours indicate the class labels, left leaning (\bullet) and right leaning (\bullet), and the magnetic eigenmaps correspond to $g = 1/4$.

- [16] C. Lange, S. Liu, N. Peyerimhoff, O. Post, Frustration index and Cheeger inequalities for discrete and continuous magnetic Laplacians, *Calculus of Variations and Partial Differential Equations* 54 (4) (2015) 4165–4196.
- [17] F. Chung, M. Kempton, *A Local Clustering Algorithm for Connection Graphs*, Springer International Publishing, 2013, pp. 26–43.
- [18] A. Lancichinetti, S. Fortunato, Benchmarks for testing community detection algorithms on directed and weighted graphs with overlapping communities, *Phys. Rev. E* 80 (2009) 016118.
- [19] E. A. Leicht, M. E. J. Newman, Community structure in directed networks, *Phys. Rev. Lett.* 100 (118703).
- [20] M. E. J. Newman, Finding community structure in networks using the eigenvectors of matrices, *Phys. Rev. E* 74 (2006) 036104.
- [21] L. Adamic, N. Glance, The political blogosphere and the 2004 US election,, in *Proceedings of the WWW-2005 Workshop on the Weblogging Ecosystem* (2005).