



HAL
open science

Spectral Measures of Distortion for Change Detection in Dynamic Graphs

Luca Castelli Aleardi, Semih Salihoglu, Gurprit Singh, Maks Ovsjanikov

► **To cite this version:**

Luca Castelli Aleardi, Semih Salihoglu, Gurprit Singh, Maks Ovsjanikov. Spectral Measures of Distortion for Change Detection in Dynamic Graphs. 2018. hal-01864079v1

HAL Id: hal-01864079

<https://hal.science/hal-01864079v1>

Preprint submitted on 29 Aug 2018 (v1), last revised 17 Sep 2018 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Spectral Measures of Distortion for Change Detection in Dynamic Graphs

Luca Castelli Aleardi¹, Semih Salihoglu², Gurprit Singh³ Maks Ovsjanikov¹

Abstract We propose a novel framework for detecting, quantifying and visualizing changes between two snapshots of a dynamic network. Unlike existing approaches, which can be sensitive to noise, and are often based on heuristics, we show how a theoretically-justified, inherently multi-scale notion of change, or distortion, can be defined and computed using spectral graph-theoretic tools. Our primary observation is that informative, robust and multi-scale measures of change can be obtained by computing a real-valued function (which we call *the distortion function*) on the nodes of the input graph, via the optimization of a pre-defined *distortion energy* in a provably optimal way. Based on extensive tests on a wide variety of networks, we demonstrate the ability of our approach to highlight the evolution of the network in an informative and multi-scale manner.

1 Introduction

Many real-world evolving systems can be conveniently encoded as dynamic graphs (e.g., collaboration networks, where new links and nodes are added over time or biological networks, where the connectivity can represent the evolution of protein interactions [1, 2, 26]). A key challenge in the visualization and analysis of dynamic networks is capturing the structural changes in the graph in a robust and efficient way. In particular, a fundamental problem is to define principled measures of change or difference between graphs that can be used to highlight the modified regions, while not being sensitive to noise. Unfortunately, in many cases these two objectives are contradictory, especially in the presence of large networks with many structural changes, where simple measures result in highly noisy highlighted regions that are difficult to interpret. In this context, several methods [14, 16, 17] have been proposed to produce informative summaries of dynamic graphs. However, as dis-

¹LIX, Ecole Polytechnique, France · ²University of Waterloo, Canada · ³Max Planck Institute for Informatics, Germany.

cussed in Section 3, they typically suffer from the lack of precise control over the type of changes that are considered and the *scale* at which they are computed, e.g., local changes of individual vertices vs. global distortions to the structure of entire regions in the graph. Furthermore, these existing techniques typically do not allow to capture and highlight only *the primary* areas of change, and can result in cluttered visualizations. Finally, most existing methods are based on heuristics and often lack formal guarantees of global optimality.

In this paper, we propose an efficient, flexible and multi-scale framework for detecting and analyzing changes in dynamic graphs, in the online setting, where only the current and the previous timestamp information are known. Given a pair of graphs, our framework is based on defining and computing an optimal *node distortion function*, which associates to each vertex a real value that quantifies the change associated with this vertex across the two graphs (with higher values corresponding to larger changes). We use the term *distortion* to emphasize the fact that our function should measure the most dramatic changes in graph structure, at a given scale. To this end, we leverage concepts from spectral graph theory to both ensure robustness against noise (by controlling the scale at which the changes are computed), and provide theoretical guarantees on the global optimality of the highlighted changes.

1.1 Related Work

The problem of capturing changes in dynamic graphs has recently attracted a lot of attention and a variety of different techniques have been proposed, including the use of animation (time-to-time mapping) or timeline (time-to-space mapping) methods (see [4], for a taxonomy of methods for dynamic graph visualization).

Network visualization. Many existing works [6, 7, 9, 10, 14, 16, 22, 29, 28] make use of a large variety of different graph drawing techniques in order to update the layout while preserving the mental map [22], which is considered as the main requirement to obtain good readability and facilitate graph exploration [3]. For instance, two interesting approaches combine the notion of vertex ages [16] and node pinning weights [14] with force-directed layouts to produce visualizations of dynamic graphs: the main idea consists to reduce node displacement via a mechanism based on the associated distortion at a vertex (which is close to simulated annealing [11]). The approach described in [14], where node pinning weights are associated with a distance-to-modification measure, is more sophisticated: it integrates force-directed layouts with a coarsening phase, and can exploit GPU parallelism for efficiency. Spectral methods, which have been extensively used for graph visualization for several decades [19, 21, 24], have recently been adapted to deal with dynamic networks [6, 9, 29]. While one approach is to extend the classic spectral layout to the case of dynamic graphs [6], another possible use of spectral methods is in the layout post-processing [9, 29, 30] that can be combined with arbitrary static layouts. In this case, node positions or even additional data such as grouping or temporal penalties can reduce node displacements.

Distorted Region Detection. As far as we know there are relatively few works, in the context of network visualization, that consider the problem of detecting and highlighting the regions with the most relevant changes and evolution. A notable exception is [17], where the authors address this problem by introducing a measure of relevance for vertices and weighted edges (called *strength*), and makes use of a threshold filtering based on a sliding time-window in order to efficiently visualize the most relevant evolving regions in the graph. Detecting relevant regions is a crucial ingredient that can be combined with other tools for the visualization of large networks [18].

Our Motivation and Contribution. Our work is inspired by previous techniques [14, 16, 17] that define various distortion measures aimed at capturing changes between graphs. These distortion measures are then combined either with force-directed methods [15, 20, 27] for graph visualization, or integrated in detecting most *distorted regions*. In context of the former, a distortion function allows controlling the displacement of vertices by manipulating the forces acting on them, while in the latter it can be used to filter vertices by, e.g., thresholding their values. The main distinguishing characteristics of our framework is that it provides precise control over both the type of distortions and the scale at which changes are computed, with theoretical optimality guarantees for the computed distortion function. None of the existing methods enable these features in a single coherent framework.

2 Proposed Framework

Preliminaries: A dynamic network is defined by a sequence of graphs $\{G_1, \dots, G_T\}$ where $G_i = (V_i, E_i)$ represents a snapshot at time i . As in previous works [9, 10, 28, 14, 16, 17, 29], we assume that we are given the correspondence between nodes in two consecutive snapshots. We further assume that the nodes and edges of each G_i are associated with positive weights, and denote by $d_i(u), w_i(u, v)$ the weight of vertex u and of edge (u, v) in G_i , respectively. In the simplest case, $d_i(u)$ is the degree of vertex u in G_i , and, for unweighted graphs, $w_i(u, v) = 1$ if $(u, v) \in E_i$ and 0 otherwise. We then define a diagonal matrix D_i , s.t. $D_i(u, u) = d_i(u)$ and a symmetric matrix W_i , s.t. $W_i(u, v) = w_i(u, v) = w_i(v, u)$. Finally, we define the weighted Laplacian matrix L_{G_i} as: $L_{G_i} = U_i - W_i$, where U_i is a diagonal matrix, s.t. $U_i(u, u) = \sum_v w_i(u, v)$.

Our framework aims at detecting the structural changes between two instances G_i and G_{i+1} of a dynamic network (graph) by computing a *distortion function* $f : V_i \rightarrow \mathbb{R}$, such that $f(v)$ should: 1) quantify the changes at vertex v between G_i and G_{i+1} , with higher values corresponding to larger changes and $f(v) = 0$ corresponding to no changes, 2) be multi-scale to reflect changes in the neighborhood of v at any given scale, 3) be efficiently computable in practice. The second property is especially important to detect and highlight evolving *regions* in graphs and to gain resilience to noisy, dispersed changes throughout the graph.




Fig. 1: Algorithm overview: given a graph and its modification, (removed edges are marked in red), we first construct a set of basis functions, then pick a distortion energy and compute its maximizers in the linear span of the basis. Finally, we compute the optimal distortion function that smoothly reflects the changes at each vertex from very high (in red) to least affected vertices (in blue).

General overview: To compute the *distortion function* that satisfies the above three criteria, we first define a *distortion energy* $E(G_i, G_{i+1}, f)$ that assigns a scalar score to an arbitrary real-valued function $f : V_i \rightarrow \mathbb{R}$. We then construct a *multi-scale family* \mathcal{F} of functions. Finally, we compute the *optimal distortion function* f_{opt} which corresponds to the *maximizer* of $E(G_i, G_{i+1}, f)$ for given G_i, G_{i+1} , and s.t. f lies within the family \mathcal{F} (see Fig. 1 for an illustration). We then use f_{opt} to visualize the changes across G_i, G_{i+1} , by color-coding the vertices, and to detect the most distorted regions as described in Section 3.

Our motivation for adopting this pipeline is that it allows us to 1) precisely control the types of changes the approach should be sensitive to via the choice of distortion energy, 2) control the *scale* of the computed changes via the choice of family \mathcal{F} and 3) provide theoretical guarantees, making sure that the computed maximizer of $E(G_i, G_{i+1}, f)$ is globally optimal. Moreover, this pipeline allows us to rigorously formalize the notion *distance-to-modification* behavior, that was sought in previous works [16], without a formal definition or rigorous guarantees.

2.1 Distortion Energies

The main considerations when choosing a distortion energy are first to ensure that it is sensitive to the changes that are meaningful in the context of the evolution of G_i and second, it should be easy to optimize, so that computing the maximizer f of $E(G_i, G_{i+1}, f)$ w.r.t. f can be done efficiently in practice. To this end, we consider the following two distortion energies (notations from Sec. 2):

$$\begin{aligned}
 E_{\text{vertex diff}}(G_i, G_{i+1}, f) &= \frac{\sum_u f(u)^2 (d_i(u) - d_{i+1}(u))^2}{\sum_u f(u)^2 d_i(u)}, \\
 E_{\text{edge diff}}(G_i, G_{i+1}, f) &= \frac{\sum_{u,v} (f(u) - f(v))^2 (w_i(u,v) - w_{i+1}(u,v))^2}{\sum_u f(u)^2 d_i(u)}.
 \end{aligned} \tag{1}$$

Intuitively, these two energies, $E_{\text{vertex diff}}$ and $E_{\text{edge diff}}$, are sensitive to the absolute *changes* (increase or decrease) of the weights of the vertices and the edges, respectively. These energies associate a scalar score to any real-valued function f , which can then be used to find the optimal distortion function by computing $f_{\text{opt}} = \arg \max_f E(G_i, G_{i+1}, f)$. In each case, $E(G_i, G_{i+1}, f)$ is large when f is supported on the regions where these weights change the most. Note also that both $E_{\text{vertex diff}}(f)$, and $E_{\text{edge diff}}(f)$ are defined such that the change is scaled by the vertex weight $d_i(u)$ in the denominator. This is because a single edge addition or deletion is typically less important for a vertex with a large weight (or degree) than for a vertex with a small weight. However, this choice might also be application-specific. If this normalization is not necessary, it can easily be removed from all of the derivations below.

Remark that the choice of energy is application specific. For example, if the network evolution is primarily controlled by change in *vertex weights*, then $E_{\text{vertex diff}}$ is more appropriate. In most cases below, however, the graphs evolve due to edge modifications and we use $E_{\text{edge diff}}$ as the main distortion energy, although we report results for both.

2.2 Choice of Scale via Reduced Functional Space

After selecting a distortion energy $E(G_i, G_{i+1}, f)$ our goal is to find the optimal distortion function f_{opt} by maximizing E for given G_i, G_{i+1} . Moreover, as mentioned in Section 2, we would like to be able to control the *scale* of the solution, in order to gain robustness to disparate, possibly noisy local changes, and to detect the areas or regions where the most important changes occur.

Thus, instead of maximizing $E(G_i, G_{i+1}, f)$ across all choices of f , we propose to consider the function f that lies in the appropriate functional subspace. More concretely, we enforce the function f to lie in the *linear subspace* spanned by some “desirable” functions φ , i.e., we force: $f = \sum_{j=1}^k a_j \varphi_j$ where φ_j are some k fixed pre-defined functions, i.e., $\varphi_j : V_i \rightarrow \mathbb{R}$ and a_j are the unknown scalar coefficients. This way, computing the optimal f amounts to finding the coefficients $\{a_k\}$ such that $E(G_i, G_{i+1}, f) = E(G_i, G_{i+1}, \sum_j a_j \varphi_j)$ is maximized.

In practice, we control the *scale* of the solution via the choice of k , which corresponds to the dimensionality of the functional space. A small value of k corresponds to *global scale* as it enforces f to be chosen in the space spanned by a small set (of potentially globally supported) functions, whereas larger values of k provide more freedom for selecting the optimal distortion function f . In the limit, when k equals the number of vertices in the graph, and φ are linearly independent, then f can be chosen to be an arbitrary function, including an indicator function of a single vertex. In this paper, we consider the following two families of functions:

Option 1: Region-based functions. Perhaps the most intuitive choice of a functional family corresponds to simply taking φ_j to represent indicator (characteristic)

functions of some regions on the graph G_i . In the simplest case, each such function can represent a neighborhood of some fixed vertex, of a given size. More precisely, given a partition of the vertex set V_i into k distinct regions $\{R_1, R_2, \dots, R_k\}$ we define $\varphi_j : V_i \rightarrow [0, 1]$ as the indicator function of R_j : $\varphi_j(u) = 1$ if $u \in R_j$, and $\varphi_j(u) = 0$ otherwise. We can also incorporate a distance-to-modification behavior by simply defining $\varphi_j(u) = 1$ if $u \in R_j$ and u has a modified neighborhood, and $\varphi_j(u) = h(\text{dist}(u))$ where $\text{dist}(u)$ is the distance from u to the closest modification and $h : \{0, 1, \dots, n-1\} \rightarrow [0, 1]$ is a decreasing function provided by the user (we define $h(\text{dist}(u)) = 1/(1 + \text{dist}(u))$ in the experiments reported in Section 3). In practice we use the partitions computed by the *Louvain algorithm* [5], which is a hierarchical method for community detection based on modularity optimization. Such a clustering method is especially appropriate in our setting since it provides a *resolution* parameter that controls the desired level in the clustering hierarchy, which naturally corresponds to the scale at which we analyze the graph.

Option 2: Laplacian eigen-basis. Another basis that we consider is given by the eigenfunctions associated with the k *smallest eigenvalues* of the following generalized eigenvalue problem:

$$L_{G_i} \varphi = \lambda D_{G_i} \varphi \quad (2)$$

where L_{G_i} is the Laplacian matrix of G_i and D_{G_i} is the diagonal matrix of vertex weights. We choose this basis because the eigenfunctions of the Laplacian naturally have a “multi-scale” property, which intuitively corresponds to the equivalent of Fourier bases and which has been used extensively in the context of signal processing on graphs [25]. In our context, remark that each φ_j is associated with a non-negative eigenvalue $\lambda_j \geq 0$ (since L_{G_i} and D_{G_i} are symmetric positive semi-definite). Moreover, a simple calculation shows that for any $f = \sum_{j=1}^k a_j \varphi_j$

$$\frac{f^T L_{G_i} f}{f^T D_{G_i} f} = \frac{\sum_{u,v} (f(u) - f(v))^2 w_i(u,v)}{\sum_u f(u)^2 d_i(u)} \leq \lambda_k \quad (3)$$

Now, the quantity $\sum_{u,v} (f(u) - f(v))^2 w_i(u,v)$ can naturally be interpreted as the “smoothness” of the function, since it captures the sum of the squared differences of f along the edges of the graph. This means that if a function f lies in the span of the eigenfunctions corresponding to the k smallest eigenvalues of the problem in Eq. (2), then the smoothness of f , as defined in Eq. (3), is bounded by λ_k .

Note that for a connected graph, when $k = 1$, then f must be a constant function, since $\lambda_1 = 0$, and φ_1 must be constant. Conversely, if $k = n$ then f can be an arbitrary function, since φ are linearly independent. Thus, we interpret k as controlling the scale of the solution, where small k corresponds to global scale (very smooth functions), and large k corresponds to local scale (possibly arbitrarily irregular, or concentrated functions).

2.3 An Algorithm to Compute the Spectral Distortion

We can now design a simple procedure for computing the optimal distortion function, where all steps can be expressed in terms of linear algebraic computations. Assuming we want to maximize the energy $E_r(G_i, G_{i+1}, f)$, we compute the optimal distortion function f_{opt} with the following three steps:

1. Choose the value for the parameter k and compute the family \mathcal{F} at that scale. For example, when using the Laplacian basis, compute the k smallest eigenvalues $\lambda_1, \dots, \lambda_k$ and the corresponding eigenfunctions ϕ_1, \dots, ϕ_k of the problem $L_{G_i} \phi = \lambda D_{G_i} \phi$. Store these functions as columns of the matrix Φ^k .
2. Given the matrix Φ^k and an energy E_r , compute the optimum function $f_{\text{opt}} := \arg \max_{f \in \text{span}(\Phi^k)} E_r(G_i, G_{i+1}, f)$. For this, we solve the eigenproblem: $\max_{\sigma} S_{i+1} \xi = \sigma \xi$ where the matrix S_{i+1} (of size $k \times k$) is given as:

$$\begin{aligned} E_{\text{vertex diff}} : S_{i+1} &:= (\Phi^k)^T (D_{G_i} - D_{G_{i+1}})^2 \Phi^k \\ E_{\text{edge diff}} : S_{i+1} &:= (\Phi^k)^T L_{G_i, G_{i+1}}^- \Phi^k \end{aligned} \quad (4)$$

- Here the diagonal matrices D and the Laplacian matrices L follow the definitions given in the beginning of Section 2, and the matrix L^- is defined such that $L^-(u, v) = -(w_i(u, v) - w_{i+1}(u, v))^2$, and $L^-(u, u) = \sum_v (w_i(u, v) - w_{i+1}(u, v))^2$.
3. Finally, the optimal distortion function corresponds to the eigenvector ξ_{max} (having size k) associated with the largest eigenvalue σ_{max} of the eigenproblem from step 2. The function f on the vertices of G_i can be computed via the matrix product: $\mathbf{f}_{\text{opt}} = \Phi^k \xi_{\text{max}}$.

The correctness of the algorithm above for computing the optimal distortion function is ensured by the following Lemma (see the Appendix for the proof):

Lemma 1. *The algorithm described in the three steps above is guaranteed to result in a distortion function that maximizes the given energy, while remaining within the subspace spanned by the eigenfunctions Φ^k .*

Note that the pipeline above is designed to compute a single optimal distortion function \mathbf{f}_{opt} at the given scale k . In practice, this typically corresponds to detecting a single most distorted area or region of the graph. In order to compute the top p most distorted regions, it can be convenient to consider more than one eigenvector of S_{i+1} in step 3 above. In this case we take the linear combination of the squares of eigenvectors corresponding to the p largest eigenvalues of S_{i+1} described in step 2. In other words we compute: $\mathbf{f}_{\text{opt}} = \sum_{j \leq p} \sigma_j (\mathbf{f}_j[u])^2$, where we define $\mathbf{f}_j = \Phi^k \xi_j$ and σ_j is the j^{th} largest eigenvalue of S_{i+1} . Finally we normalize \mathbf{f}_{opt} to have values between 0 and 1.

Key parameters: The key parameters in our framework include: the choice of the basis, the distortion energy function, the choice of “scale” or smoothness parameter k , and the choice of p corresponding to the number of largest eigenvalues, related to the number of highest-distorted parts that are considered.

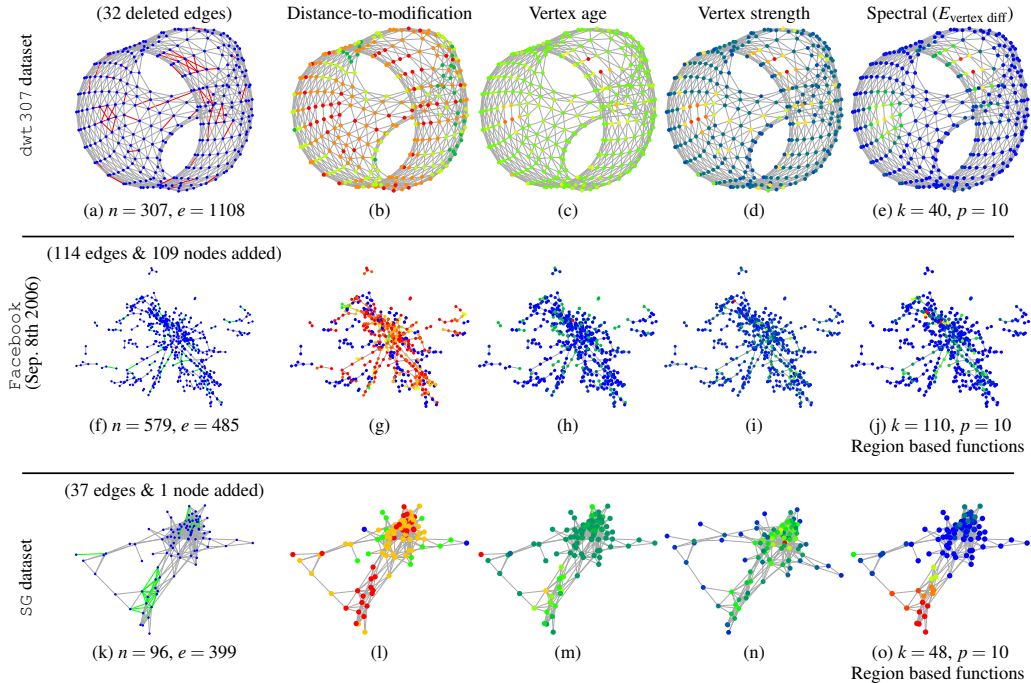


Fig. 2: (Left) Network evolution: the red (resp. green) segments represent the removed (resp. added) edges in a single time step. (Columns 2-5) Qualitative comparison between the distortion measures δ_{DM} (distance-to-modification), δ_{VA} (vertex age) and δ_{VS} (vertex strength) with our spectral distortion.

3 Experimental evaluation

We provide an experimental evaluation of our spectral of distortion and compare it to other approaches previously proposed in the visualization of dynamic graphs.

3.1 Datasets

We perform tests on dynamic networks having different structural properties and evolution behavior: as in most existing works [9, 10, 14, 16], all of these graphs are undirected and unweighted.

Real-world networks. We consider real-world networks from the SuiteSparse Matrix Collection [12] (`3elt`, `dwt307`, etc.). The sizes of these graphs range from hundreds to thousands of nodes. As in [14] we construct a dynamic sequence of networks with a simple process based on a random edge decimation.

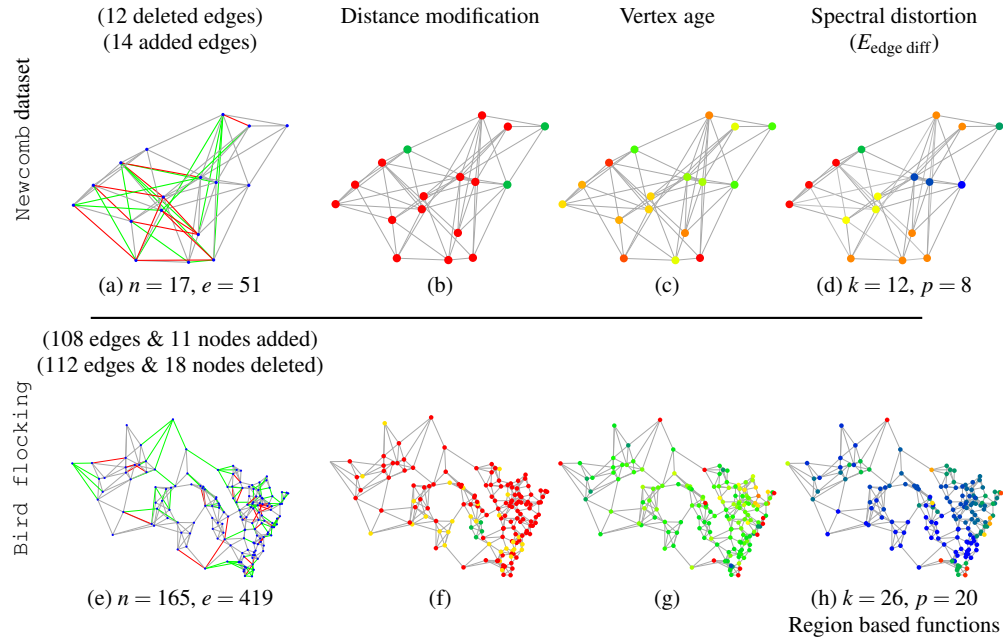


Fig. 3: Qualitative comparison of distortion functions for the `Newcomb` and `Bird flocking` datasets. We compare different distortion measures δ_{DM} (distance-to-modification) and δ_{VA} (vertex age) with our spectral distortion (obtained optimizing the energy $E_{\text{edge diff}}$).

Social networks. We deal with networks having a different underlying structure. For instance we consider networks extracted from the `Facebook-Growth` dataset¹ (also evaluated in [10]): this dataset spans an interval starting from Sep. 5th 2006 and evolves under a dynamic process where new entities/links are added over time. we extract a sequence of dynamic networks following the evolution of the largest connected component over the first 15 days (each time step correspond to the growth during a single day).

Interaction networks. We consider the sequence of aggregated networks analyzed in [8] (referred to as `SG`), where nodes are individuals visiting the Science Gallery in Dublin and edges describe the face-to-face proximity between individuals over a daily time window: these networks have a few hundreds of nodes and evolve under an aggregation process (only vertex and edge additions occur).

Newcomb network. We consider the `Newcomb's fraternity` dataset (referred to as `newcomb` in previous works [9, 14, 16]), whose evolution is driven by a dramatic and global rewiring process.

¹ The dataset is available at <https://www.eecs.wsu.edu/~yyao/StreamingGraphs.html>

Proximity networks. We also evaluate a second dataset whose growth is based on a dynamic process where new entities and links are added or removed over time. More precisely we consider a collection of networks (referred to as `bird flocking`), describing the collective motion of birds. Given the 3D locations of birds (from the dataset [13]) we construct a sequence of networks by computing the k -nearest neighbor graph at each time step (this choice is motivated by the behavior of interactions, which are independent of the distance between birds). This network sequence is particularly interesting since the evolution is driven by a dramatic edge rewiring process, including vertex additions and removals (which occur because of occlusion issues during the tracking phase).

3.2 Baselines

We compare our approach to three other notions of distortion: *distance-to-modification*, *vertex age* and *vertex strength* distortions based on the methods proposed in [14], [16] and [17] respectively.

The *Distance-to-modification* distortion is the function $\delta_{DM}(u) : V \rightarrow [0, 1]$ defined by $\delta_{DM}(u) = h(dist(u))$ where h is a decreasing function, and $dist(u)$ is the graph distance between u and the closest modification in the graph. More precisely, following [14] we define $\delta_{DM}(u) = 1 - \alpha^{(1 - \frac{dist(u)}{R_c})}$ where α (scale parameter) and R_c (cutoff distance) are user-supplied parameters.

To better distinguish between major and minor changes Gorochoowski et al. [16] introduced an age function $age(u, i)$ that quantifies the amount of changes for vertex u at time i . We then take a *vertex age distortion* $\delta_{VA}(u) = e^{-\beta age(u, i)}$, where $\beta \in \mathbb{R}^+$ is a user supplied parameter.

Finally, we define the distortion δ_{VS} based on the notion of *vertex strength* following the approach proposed in [17], that integrates vertex degrees and makes use of an exponential sliding time-window (we normalize to get values in $[0, 1]$).

For completeness we provide the definition of $age(u, i)$ in the appendix.

3.3 Evaluation: Region extraction

We evaluate different distortion measures by their ability to highlight relevant *regions* of the graph that undergo the biggest changes. Thus, we use each measure to first extract the regions on the graph as follows:

Approach 1: threshold filtering. In the simplest case, we simply keep a small fraction $\rho \in [0, 1]$ of the vertices having the largest distortion: as suggested in [17], we compute the sub-graph induced by non isolated vertices, containing at most $\rho|V_{i+1}|$ vertices (see Figure 4 for an illustration).

Approach 2: distortion-driven BFS traversal. One drawback of the previous approach is that it can result in many disconnected regions, especially when the net-

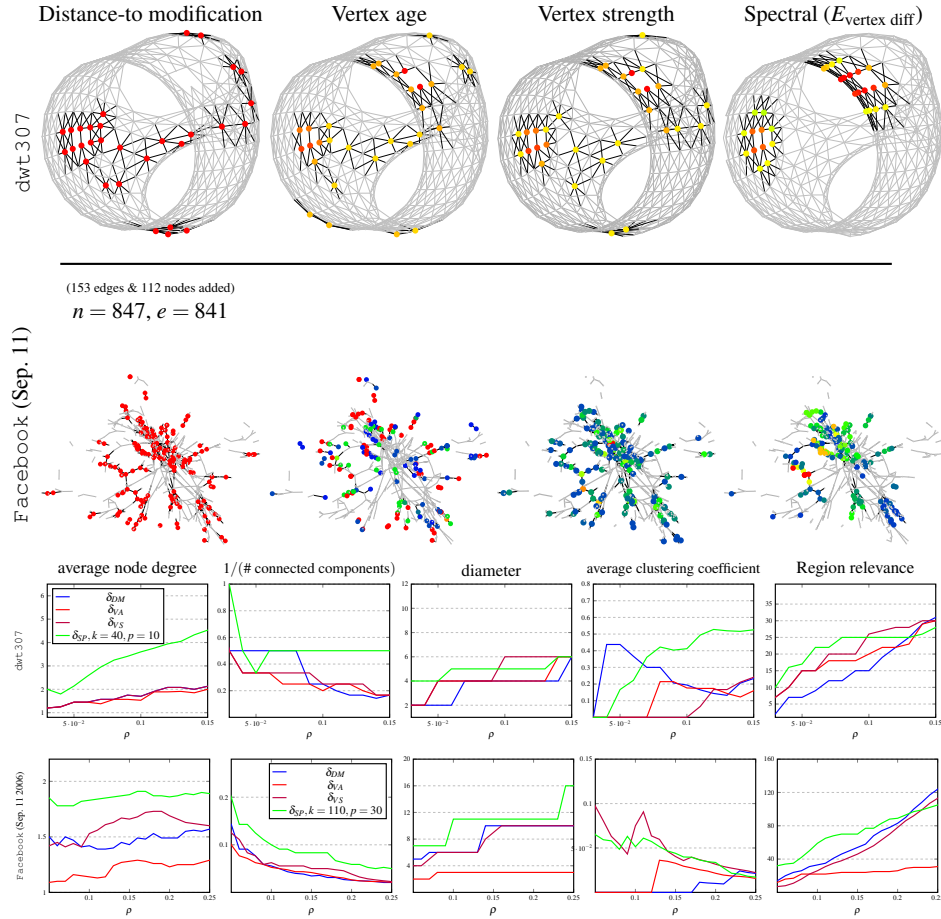


Fig. 4: Threshold filtering approach: in the pictures above we highlight a fraction ρ of vertices with highest distortion, setting $\rho = 0.1$ for dwt 307 and $\rho = 0.2$ for Facebook. In the charts below we plot the structural properties (*average node degree, average clustering coefficient, ...*) and the *region relevance* of each filtered graph as a function of ρ .

work evolution is dramatic. To counter this, we also use each distortion function to perform a BFS traversal (vertices with large distortion are visited first) starting from some initial seed vertex, which guarantees a single connected region (let N_r denote the number of vertices of such a region). For fairness, we repeat this process taking as seeds a fraction ρ of vertices with largest distortion.

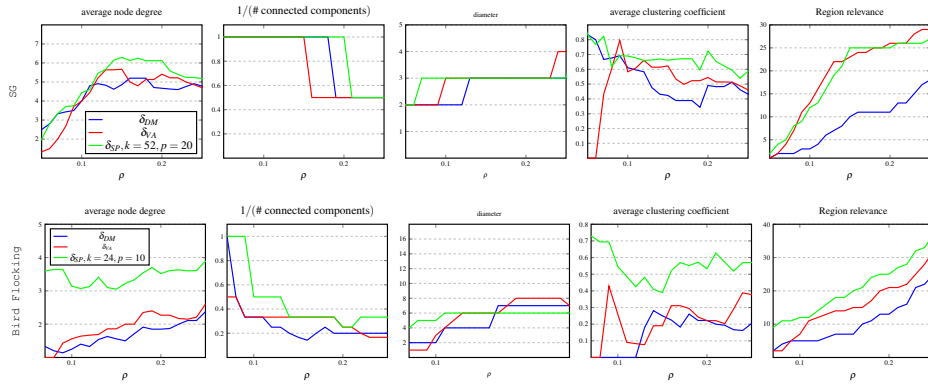


Fig. 5: Statistics concerning the threshold filtering approach for the SG and Bird flocking networks. We plot the structural properties and the *region relevance* as a function of ρ (fraction of vertices with highest distortion).

Region quality We evaluate the extracted regions according several measures of quality: first, as suggested in [17] we measure the *structural properties*² of the extracted regions (e.g., average node degree, number of connected components, average clustering coefficient, ...). Second, we define the *relevance* of a region as the total number of edge modifications (additions or deletions) involving vertices that lie in it.

3.4 Experimental comparison and discussion

Qualitative comparison

In our first, qualitative, evaluation we plot the distortion functions obtained by different methods across several datasets in Fig. 2 (for the energy $E_{\text{vertex diff}}$) and Fig. 3

² We make use of the GraphStream library [23] for the computation of structural network parameters.

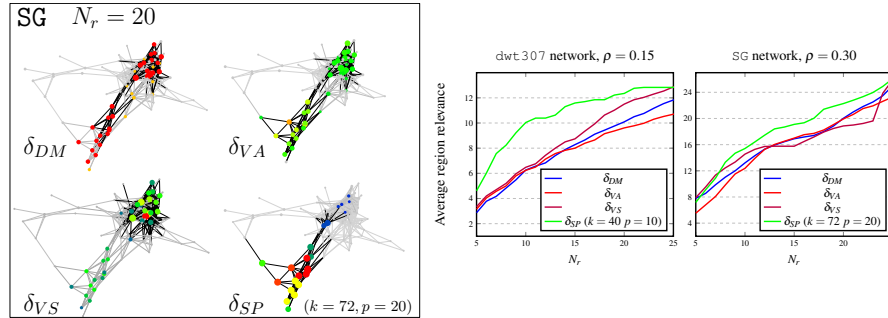


Fig. 6: We compare the region extraction approach based on BFS traversal. (left) for each distortion function (and for a fixed value of N_r), we highlight the *average region* extracted with the BFS traversal. (right) plots of the average region *relevance* (higher is better).

(for the energy $E_{\text{edge diff}}$). In all our examples we use the heat map color-scale representing distortion values in the range $[0 \dots \max_{u \in V} \delta(u)]$, from blue to green to red. As can be seen in Fig. 2, previous methods can be either *too* sensitive, highlighting a large portion of the graph (e.g., in the case of the δ_{DM} distortion) or fail to distinguish *regions* or parts of the graph that undergo the changes (in the case of the δ_{VA} distortion for the SG, for example). Baseline approaches clearly fail to distinguish between major and negligible modifications, and are not very informative when dealing with dramatic changes as the for the Newcomb and bird flocking networks (see Fig. 3).

On the other hand, our method is precise and at the same time captures the regions that undergo the most changes in a multi-scale manner.

Threshold filtering

As done in [17], we compare distortion functions by evaluating the evolution of the structural properties, performing the extraction of relevant regions with the threshold filtering approach. The layouts and plots of Fig. 4 and Fig. 5 show that the regions detected with our spectral approach do not lead to dramatic fluctuations of the main structural properties: the filtered vertices are likely to define connected subgraphs, whose structural properties vary in a smooth way, even for small values of ρ . This does not hold for basic distortion measures (δ_{DM} and δ_{VA}), that lead to highly disconnected sub-graphs, and more drastic fluctuations of modularity and clustering coefficients, not being able to distinguish the most relevant changes from local noise. Observe that our spectral approach offers to the user the capability of choosing the parameters k and p in order to reach the right scale (tuning k) and only keep the highest distorted regions (tuning p). This feature helps the user to select, depending on the application and desirable goals, the regions with the most relevant

Network	vertices	edges	Distortion vertex age	Spectral distortion			Force-directed layout one iteration	
				Step 1 ($k = 40$)		Steps 2 and 3		
				eigen-basis	region functions	$k = 40$	FR91 [15]	with octrees
SG	96	399	0.0003	0.037	0.002	0.0003	0.002	0.001
dwt307	307	1.1K	0.0015	0.048	0.004	0.0007	0.009	0.003
3elt	4720	13.7K	0.009	0.22	0.075	0.003	0.73	0.11
barth5	15606	61.4K	0.026	0.84	0.32	0.011	15.2	0.57

Table 1: This table reports the average runtime performance of all steps involved in the computation of graph distortions and force-directed layouts (for a single evolutionary step). The complexity of the vertex age distortion (implemented in Java) is compared to the computation of spectral distortion setting $k = 40$. The Laplacian eigen basis is computed using Matlab 14 (option 2), while the region-based functions (option 1) are computed with a Java implementation of the Louvain algorithm. The last two columns reports the time required to run a single iteration of a force-directed layout: we implement the standard FR91 layout, as well as the fast version with approximate calculation of repulsive forces using octrees. All performances are expressed in seconds.

evolution, while discarding local irrelevant noisy changes (see e.g. Fig. 4): even for small value of ρ , the spectral distortion allows us to correctly highlight the most relevant changed regions, while δ_{DM} and δ_{VA} lead to much more dispersion, not being able to discarding noisy modifications.

BFS-driven region extraction

For a fixed value of ρ , we plot in Fig. 6 the average relevance Rel_{av} , varying the size N_r . As confirmed by our experiments, the spectral distortion always allows us to detect regions with higher values of the region relevance. This reflects the ability of spectral distortion to distinguish between regions, depending on the amount of local changes: a region strongly highlighted by spectral distortion will be visited first, leading to a significant contribution to the average energy. While small regions with a negligible amount of changes will be discarded during the BFS traversal, since they are unlikely to be detected by the spectral distortion.

3.5 Time Complexity and Runtime Performance

Below we provide a discussion on the theoretical and practical complexity of the algorithms described in this work. We have developed a pure Java implementation of all algorithms presented in this work; for the calculation of the Laplacian eigen-basis we use the eigen solver provided by Matlab (version 14). All our tests are run on a HP EliteBook, equipped with an Intel Core i7 (2.66GHz) and 8GB of RAM, using Java 1.8.

First observe that the computation of the spectral distortion in a single evolutionary step (recall the procedure in Section 2.3) is expressed in terms of linear algebra computations, which makes our algorithm very easy to implement and quite efficient in practice. To simplify the notation and discussion about practical performance, we assume that the number n of nodes is roughly the same in G_i and G_{i+1} .

Computation of basis functions.

The most expensive step in our distortion estimation algorithm is the computation of the basis functions.

When using option 2 (Laplacian eigen-basis) the extraction of the functions φ_j involves computing the first k eigenvalues of a generalized eigenproblem of size $n \times n$ for a graph with n nodes. In practice we use an iterative Arnoldi-Lanczos algorithm implemented in Matlab's `eigs` function³. The running time strongly depends on the sparsity structure of the input matrices and in practice scales approximately linearly with respect to the number of edges and quadratically with respect to k for small values of k .

When using option 1 (region-based functions) the runtime performances are much better: the timing cost is dominated by the computation of a clustering partition of the desirable size. For this we make use of the Louvain algorithm based on modularity optimization, whose complexity is assumed to be $O(n \log n)$ in practice, which allows processing large networks efficiently.

The runtime performances reported in Table 1 (columns 5 and 6) correspond to the computation of basis functions for $k = 40$.

Energy optimization.

The steps (2) and (3) of our algorithm are much less expensive, as the matrices involved have sizes $k \times k$ and $k \times n$ respectively (k and p being usually much smaller than n), which leads to approximate complexity of $O(k^3 + k^2n + pn)$: in practice the runtime performance of these steps are negligible and comparable to the ones of the computation of the vertex age distortion (see Table 1).

Finally, observe that according to the results reported in Table 1 (see last two columns), the overall computational cost of our spectral distortion is, for sufficiently large values of n , smaller than the timing cost required to run a single iteration of a standard force-directed method [15], which requires $O(|E_i| + |V_i|^2)$ time in the basic setting. When using a fast approximation of repulsive forces (using octrees), the timing cost of spectral distortion is comparable to the one required to run a single iteration.

³ We also perform tests using several Java libraries for linear algebra (Jama, Colt, Parallel Colt, MTJ, ...): as far as we can see there are no Java libraries providing efficient eigen solvers (for the generalized eigen problem) for processing large sparse linear systems.

4 Conclusion, Limitations and Future Work

We propose a novel, multi-scale framework for robustly detecting and visualizing changes in networks across two different time stamps (G_i and G_{i+1}). However, it is possible to integrate a time term in the computation of spectral distortion, that would take into account a full sequence G_1, G_2, \dots, G_{i+1} (as done in [16, 17]). Our framework described in Section 2.1, although illustrated on unweighted graphs, is rather general and allows to consider graphs with both edge and vertex weights.

References

1. Ahmed, A., Xing, E.P.: Recovering time-varying networks of dependencies in social and biological studies. *PNAS* **106**(29), 11,878–11,883 (2009)
2. Arbeitman, M.N., Furlong, E.E., Imam, F., Johnson, E., Null, B.H., Baker, B.S., Kravnow, M.A., Scott, M.P., Davis, R.W., White, K.P.: Gene expression during the life cycle of *drosophila melanogaster*. *Science* **297**(5590), 2270–2275 (2002)
3. Archambault, D.W., Purchase, H.C.: The "map" in the mental map: Experimental results in dynamic graph drawing. *Int. J. Hum.-Comput. Stud.* **71**(11), 1044–1055 (2013)
4. Beck, F., Burch, M., Diehl, S., Weiskopf, D.: The State of the Art in Visualizing Dynamic Graphs. In: R. Borgo, R. Maciejewski, I. Viola (eds.) *EuroVis - STARs* (2014)
5. Blondel, V., Guillaume, J.L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. *Journal Of Statistical Mechanics: Theory And Experiment* (2008)
6. Brandes, U., Fleischer, D., Puppe, T.: Dynamic spectral layout with an application to small worlds. *J. Graph Algorithms Appl.* **11**(2), 325–343 (2007)
7. Brandes, U., Wagner, D.: A bayesian paradigm for dynamic graph layout. In: *Graph Drawing, 5th Intern. Symp., GD '97*, pp. 236–247 (1997)
8. Cattuto, C., Van den Broeck, W., Barrat, A., Colizza, V., Pinton, J., Vespignani, A.: Dynamics of person-to-person interactions from distributed rfid sensor networks. *PLOS ONE* **5**(7), e11,596 (2010)
9. Che, L., Liang, J., Yuan, X., Shen, J., Xu, J., Li, Y.: Laplacian-based dynamic graph visualization. In: *IEEE PacificVis*, pp. 69–73 (2015)
10. Crnovrsanin, T., Chu, J., Ma, K.: An incremental layout method for visualizing online dynamic graphs. In: *Graph Drawing and Network Visualization*, pp. 16–29 (2015)
11. Davidson, R., Harel, D.: Drawing graphs nicely using simulated annealing. *ACM Trans. Graph.* **15**(4), 301–331 (1996)
12. Davis, T.A., Hu, Y.: The university of florida sparse matrix collection. *ACM Trans. Math. Softw.* **38**(1), 1 (2011). [Http://www.cise.ufl.edu/research/sparse/matrices](http://www.cise.ufl.edu/research/sparse/matrices)
13. Evangelista, D., Ray, D., Raja, S., Hedrick, T.: Three-dimensional trajectories and network analyses of group behaviour within chimney swift flocks during approaches to the roost. *Proceedings of the Royal Society B* **284**(1849) (2017)
14. Frishman, Y., Tal, A.: Online dynamic graph drawing. *IEEE Trans. Vis. Comput. Graph.* **14**(4), 727–740 (2008)
15. Fruchterman, T.M.J., Reingold, E.M.: Graph drawing by force-directed placement. *Softw., Pract. Exper.* **21**(11), 1129–1164 (1991)
16. Gorochoowski, T.E., di Bernardo, M., Grierson, C.S.: Using aging to visually uncover evolutionary processes on networks. *IEEE Trans. Vis. Comput. Graph.* **18**(8), 1343–1352 (2012)
17. Grabowicz, P.A., Aiello, L.M., Menczer, F.: Fast filtering and animation of large dynamic networks. *EPJ Data Science* **3**(1), 27 (2014)
18. Hadlak, S., Schulz, H., Schumann, H.: In situ exploration of large dynamic networks. *IEEE Trans. Vis. Comput. Graph.* **17**(12), 2334–2343 (2011)

19. Hall, K.M.: An r-Dimensional Quadratic Placement Algorithm. *Management Science* **17**(3) (1970)
20. Kamada, T., Kawai, S.: An algorithm for drawing general undirected graphs. *Inf. Process. Lett.* **31**(1), 7–15 (1989)
21. Koren, Y.: Drawing graphs by eigenvectors: theory and practice. *Computers and Mathematics with Applications* **49**(11), 1867 – 1888 (2005)
22. Misue, K., Eades, P., Lai, W., Sugiyama, K.: Layout adjustment and the mental map. *J. Vis. Lang. Comput.* **6**(2), 183–210 (1995)
23. Pigné, Y., Dutot, A., Guinand, F., Olivier, D.: Graphstream: A tool for bridging the gap between complex systems and dynamic graphs. *Emergent Properties in Natural and Artificial Complex Systems* (2007)
24. Pisanski, T., Shawe-Taylor, J.: Characterizing graph drawing with eigenvectors. *J. Chem. Inf. Comput. Sci.* **40**(3), 567–571 (2000)
25. Shuman, D.I., Narang, S.K., Frossard, P., Ortega, A., Vandergheynst, P.: The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Processing Magazine* **30**(3), 83–98 (2013)
26. Tang, J., Scellato, S., Musolesi, M., Mascolo, C., Latora, V.: Small-world behavior in time-varying graphs. *Physical Review E* **81**(5), 055,101 (2010)
27. Walshaw, C.: A multilevel algorithm for force-directed graph-drawing. *J. Graph Algorithms Appl.* **7**(3), 253–285 (2003)
28. X. Du, Y.W., Wu, L.: A multi-constraint layout algorithm for dynamic network visualization. In: *IEEE 2nd International Conference on Big Data Analysis (ICBDA)*, pp. 832–836 (2017)
29. Xu, K.S., Klinger, M., III, A.O.H.: A regularized graph layout framework for dynamic network visualization. *Data Min. Knowl. Discov.* **27**(1), 84–116 (2013)
30. Yuan, X., Che, L., Hu, Y., Zhang, X.: Intelligent graph layout using many users’ input. *IEEE Trans. Vis. Comput. Graph.* **18**(12), 2699–2708 (2012)

5 Appendix

5.1 Proof of Lemma 1

The validity of Lemma 1 follows from the following statement:

Lemma 2. *The maximizer of $E_{vertex\ diff}$ and $E_{edge\ diff}$ among all functions in the span of Φ^k is given by $\Phi^k \xi_{\max}$ where ξ_{\max} is the eigenvector corresponding to the largest eigenvalue of S_{i+1} , given in Eq. (4) in Section 2.3.*

Proof. First note that any function f in the span of Φ^k can be written as $f = \Phi^k \mathbf{f}$, where \mathbf{f} is some vector of coefficients. Moreover, for any function f , we can write the energies in matrix-vector form as follows:

$$\begin{aligned}
E_{\text{vertex diff}}(f) &= \frac{\sum_u f(u)^2 (d_i(u) - d_{i+1}(u))^2}{\sum_u f(u)^2 d_i(u)} \\
&= \frac{f^T (D_{G_i} - D_{G_{i+1}})^2 f}{f^T D_{G_i} f}, \text{ while} \\
E_{\text{edge diff}}(f) &= \frac{\sum_{u,v} (f(u) - f(v))^2 (w_i(u,v) - w_{i+1}(u,v))^2}{\sum_u f(u)^2 d_i(u)} \\
&= \frac{f^T L_{G_i, G_{i+1}}^- f}{f^T D_{G_i} f}.
\end{aligned}$$

The last equality holds since for any matrix W , s.t., $W(u,u) = -\sum_{v \neq u} W(u,v)$, we have: $f^T W f = -\sum_{u,v} W(u,v) (f(u) - f(v))^2$. Now note that $L_{G_i, G_{i+1}}^-$ satisfies this property and $L^-(u,v) = -(w_i(u,v) - w_{i+1}(u,v))^2$. This implies that if $f = \Phi^k \mathbf{f}$, then for both energies we have:

$$E(f) = \frac{\mathbf{f}^T S_{i+1} \mathbf{f}}{\mathbf{f}^T \Phi^k D_{G_i} \Phi^k \mathbf{f}},$$

where S_{i+1} is defined appropriately (as in Eq. (4)). Furthermore since Φ^k are given as the eigenvectors corresponding to the generalized eigenproblem $L_{G_i} \phi = \lambda D_{G_i}$, they are orthonormal with respect to the matrix D_{G_i} , so that $\Phi^k D_{G_i} \Phi^k = Id$. Therefore we have: $E(f) = \frac{\mathbf{f}^T S_{i+1} \mathbf{f}}{\mathbf{f}^T \mathbf{f}}$. By the standard min-max theorem, the vector that maximizes this ratio (the Rayleigh-Ritz quotient) must be the eigenvector corresponding to the largest eigenvalue of S_i , which implies the result of the lemma.