Spectral Measures of Distortion for Change Detection in Dynamic Graphs

Luca Castelli Aleardi[1], Semih Salihoglu[2], Gurprit Singh[3] Maks Ovsjanikov[1]

[1] LIX, Ecole Polytechnique, France, amturing,maks@lix.polytechnique.fr
[2] University of Waterloo, Canada, semih.salihoglu@uwaterloo.ca
[3] Max Planck Institute for Informatics, Germany, gsingh@mpi-inf.mpg.de

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 minor and isolated changes, 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.

Keywords: network visualization, dynamic networks, spectral methods

1 Introduction

Many real-world evolving systems can be conveniently encoded as dynamic graphs (e.g., biological networks, where the connectivity can represent the evolution of protein interactions or collaboration networks, where new links and nodes are added over time [1,2,24]). 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 [13,15,16] have been proposed to produce informative summaries of dynamic graphs. However, as discussed 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 (i.e. disparate and unrelated, local changes), 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 [6,7,10,11,13,15,26,25] make use of a large variety of different graph drawing techniques in order to update the layout while preserving the mental map, which is considered as the main requirement to facilitate graph exploration [3]. For instance, two interesting approaches combine the notion of vertex ages [15] and node pinning weights [13] 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. Spectral methods, which have been used for graph visualization for several decades [18,19], have recently been adapted to deal with dynamic networks [6,10,26]. 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 [10,26,27] that can be combined with arbitrary static layouts.

Distorted Region Detection. 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 [16], 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 [17].

Our Contribution. Our work is inspired by previous techniques [13,15,16] that define various distortion measures aimed at capturing changes between graphs. These distortion measures are then combined either with force-directed methods [14] 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 $\{G_1,\ldots,G_T\}$ where $G_i=(V_i,E_i)$ represents a snapshot at time i.

As in previous works [10,11,25,13,15,16,26], we assume that we are given the correspondence between nodes in two consecutive snapshots. The nodes and edges of each G_i are allowed to be associated with positive weights: we 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* $f:V_i\longrightarrow\mathbb{R}$, 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.

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 $f:V_i\longrightarrow\mathbb{R}$. We then construct *a multi-scale family* $\mathscr{F}$ 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 $\mathscr{F}$ (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 $\mathscr{F}$ and 3) provide theoretical guarantees, making sure that the computed maximizer of $E(G_i,G_{i+1},f)$ is globally optimal.

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 $E(G_i,G_{i+1},f)$

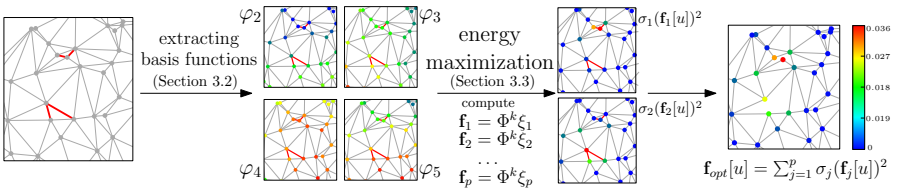

Fig. 1: Algorithm overview: given a graph and its modification, (removed edges are marked in red), we first construct a set of basis functions, then pick a distortion energy and compute its maximizers in the linear span of the basis. Finally, we compute the optimal distortion function that smoothly reflects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

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.

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., $\varphi_j : V_i \longrightarrow \mathbb{R}$ 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 $\{R_1, R_2, \ldots, R_k\}$ we define $\varphi_j : V_i \longrightarrow [0,1]$ 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 $\varphi_j(u) = h(dist(u))$ where $dist(u)$ is the distance from u to the closest modification and $h : \{0,1,\ldots n-1\} \rightarrow [0,1]$ is a decreasing function provided by the user (we define $h(dist(u)) = 1/(1+dist(u))$ 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. We consider as basis the eigenfunctions associated with the *k smallest eigenvalues* of the generalized eigenvalue problem:

$$L_{G_i}\varphi = \lambda D_{G_i}\varphi \tag{2}$$

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 [23]. 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 $f = \sum_{j=1}^{k} a_j \varphi_j$

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

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 $\mathscr{F}$ at that scale. For example, when using the Laplacian basis, compute the k smallest eigenvalues $\lambda_1, \ldots, \lambda_k$ and the corresponding eigenfunctions $\varphi_1, \ldots, \varphi_k$ of the problem $L_{G_i}\varphi = \lambda D_{G_i}\varphi$. 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 $f_{opt} := \arg\max_{f \in \text{span}(\Phi^k)} E_r(G_i, G_{i+1}, f)$. For this, we compute a solution of the eigenproblem: $\max_\sigma S_{i+1}\xi = \sigma\xi$ where the matrix S_{i+1} (of size $k \times k$) is given as:

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

   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 $\xi_{\max}$ (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: $\mathbf{f}_{\text{opt}} = \Phi^k \xi_{\max}$.

The correctness of the algorithm above for computing the optimal distortion function is ensured by the following Lemma (the proof is provided in [8]):

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 $\mathbf{f}_{\text{opt}}$ 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: $\mathbf{f}_{\text{opt}} = \sum_{j \le p} \sigma_j (\mathbf{f}_j[u])^2$, 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 $\mathbf{f}_{\text{opt}}$ to have values between 0 and 1.

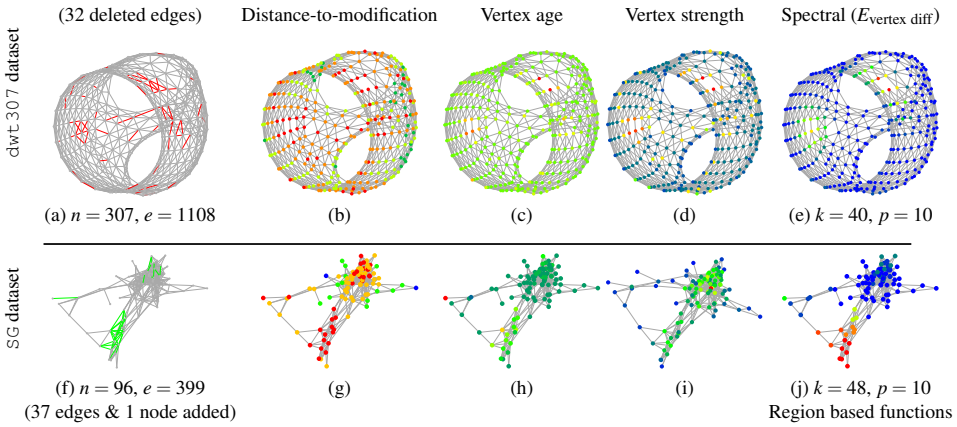

Fig. 2: (Left) Network evolution: the red (resp. green) segments represent the removed (resp. added) edges in a single time step. (Columns 2-5) Qualitative comparison between the distortion measures δ_{DM} (distance-to-modification), δ_{VA} (vertex age) and δ_{VS} (vertex strength) with our spectral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

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.

3 Experimental evaluation

We provide an experimental evaluation of our spectral distortion (denoted δ_{SP}) and compare it to other approaches proposed in the visualization of dynamic graphs.

Choice of the energy. Remark that the choice of energy is application specific. For example, if the network evolution is primarily controlled by change in *vertex degrees* (or weights), then $E_{\text{vertex diff}}$ is more appropriate. This happens, in particular, if the evolution contains *only* edge additions or removals but not both. Since this is the case for the datasets we consider below, we only present results using $E_{\text{vertex diff}}$. We also report in [8] experimental results for the energy $E_{\text{edge diff}}$, that is more suitable for graphs evolving under both edge removals and additions.

Datasets. We perform tests on dynamic networks having different structural properties and evolution behavior: as in most existing works [10,11,13,15], 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.). As in [13] we construct a dynamic sequence of networks with a simple process based on a random edge decimation.

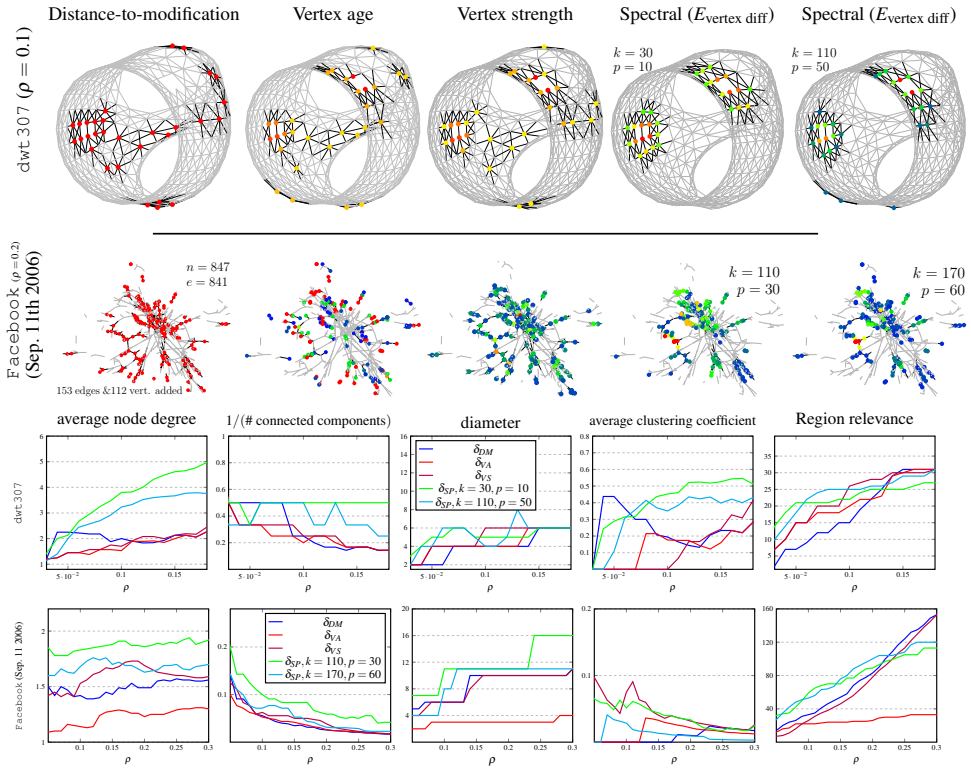

Fig. 3: Threshold filtering approach: in the pictures above we highlight a fraction ρ of vertices with highest distortion. In the charts below we plot the structural properties and the *region relevance* of each filtered graph as a function of ρ.

Complex networks. We consider a sequence of networks (also evaluated in [11]) extracted from the `Facebook-Growth` dataset [4]: this dataset spans an interval from Sep. 2006 to Jan. 2009. We also deal with snapshots of the interaction network analyzed in [9] (referred to as `SG`), where edges describe the face-to-face proximity between individuals visiting the Science Gallery in Dublin. Both networks evolves under a dynamic process where new entities/links are added over time. More extensive experiments including other networks can be found in [8].

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 [13], [15] and [16], respectively. The *distance-to-modification* distortion is the function $\delta_{DM}(u) : V \to [0,1]$ 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

[4] The raw data is available at https://www.eecs.wsu.edu/~yyao/StreamingGraphs.html

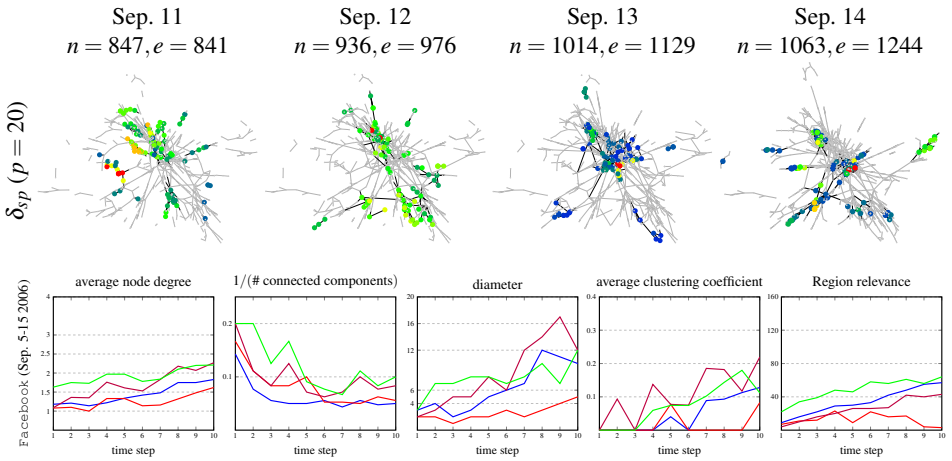

Fig. 4: Dynamic growth of the `Facebook` network: we compare distortion functions evaluating the evolution of structural properties over the first ten days from Sep. 5 to Sep. 15 2006. The pictures above show the graph filtered using spectral distortion over four time steps. All pictures and plots are obtained setting $\rho = 0.1$.

the graph. More precisely, following [13] 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 Gorochowski et al. [15] introduced an age function $age(u,i)$ that quantifies the amount of changes for vertex u at time i (see [15] for the definition). 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 [16], that integrates vertex degrees and makes use of an exponential sliding time-window (we normalize to get values in $[0,1]$). For the sake of completeness a detailed definition of $age(u,i)$ and δ_{VS} can be found in [8].

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 [16], we compute the subgraph induced by non isolated vertices, containing at most $\rho|V_{i+1}|$ vertices (see Figure 3 for an illustration).

Approach 2: BFS traversal. One drawback of the previous approach is that it can result in many disconnected regions, especially when the network 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 (whose size is denoted N_r). This process is repeated taking as seeds a fraction ρ of vertices with largest distortion.

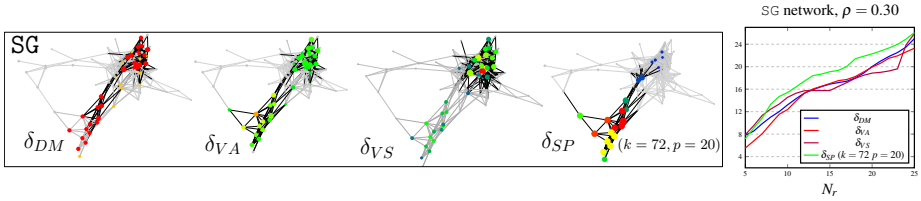

Fig. 5: Region extraction approach based on BFS traversal. (left) for each distortion function (and for a fixed value of N_r), we highlight the *average region* extracted for the SG network. (right) plots of the average region *relevance* (higher is better).

Region quality. We evaluate the extracted regions according to several measures of quality: first, as suggested in [16] we measure the *structural properties* (e.g. average vertex degree, average clustering coefficient, diameter, number of connected components) of the extracted regions: for instance, higher values of the average node degree (or $1/$#connected components) are better, as they correspond to filtered regions which are locally more dense and less fragmented. Second, we evaluate the total number of edge modifications (additions or deletions) involving vertices that lie in a region, referred to as *region relevance* (higher values are better).

3.1 Experimental comparison and discussion

Qualitative comparison. In our first evaluation we plot the distortion functions obtained by different methods across several datasets in Fig. 2. 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).

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.

Filtering. The layouts and plots of Fig. 3 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 (especially δ_{DM} and δ_{VA}), that lead to highly disconnected and sparse sub-graphs, and more drastic fluctuations of structural properties, 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 evolution, while discarding local irrelevant noisy changes (see layouts and plots of region relevance in Fig. 3 and Fig. 5): even for small value of ρ, the spectral distortion allows us to correctly highlight the most relevant changed

regions, while δ_{DM}, δ_{VA} and δ_{VS} lead to much more dispersion, not being able to discarding noisy modifications. Observe that as ρ increases, the capability of δ_{DM}, δ_{VA} and δ_{VS} distortions to capture modifications also increases (see the plot of region relevance in Fig. 3). Above a given threshold the total amount of changes captured by spectral distortion does not increase significantly since irrelevant local changes are ignored: this is the unavoidable price to pay for keeping the filtered graph fragmented as little as possible.

The plots of Fig. 4 show the evolution of structural properties over the first ten days of the `Facebook` network sequence (all results and layouts are obtained setting $\rho = 0.1$): the spectral distortion always captures a larger amount of modifications, while leading in overall to a smoother transition of the filtered graphs, which are more dense and locally connected, between consecutive time steps.

Runtime performance. The most time-consuming step in the calculation of our spectral distortions (Sec. 2.3) is the computation of the basis; steps (2) and (3) have a negligible cost as they involve simple algebraic matrix operations (of size $k \times k$ and $k \times n$, with $k \ll n$). A more detailed discussion can be found in [8].

4 Conclusion, Limitations and Future Work

We propose a novel, multi-scale framework for robustly detecting and visualizing dynamic 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 $G_1, G_2, ..., G_{i+1}$ (as done in [15,16]). 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. It could be interesting to consider other operators than the Laplacian, which are more adapted to deal with the large-scale structure of sparse networks [20,22]. As further application, it could interesting to see whether our approach could apply to the problem of detecting *change-points* [21] and temporal anomalies in evolving networks.

Acknowledgements Parts of this work were supported by the Jean Marjoulet chair from Ecole Polytechnique, a Google Focused Research Award, the ERC Starting Grant No. 758800 (EXPROTEA) and the French ANR GATO (ANR-16-CE40-0009-01).

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., Krasnow, 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. Castelli Aleardi, L., Salihoglu, S., Singh, G., Ovsjanikov, M.: Spectral Measures of Distortion for Change Detection in Dynamic Graphs (extended version). URL https://hal.archives-ouvertes.fr/hal-01864079
9. 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)
10. Che, L., Liang, J., Yuan, X., Shen, J., Xu, J., Li, Y.: Laplacian-based dynamic graph visualization. In: IEEE PacificVis, pp. 69–73 (2015)
11. 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)
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
13. Frishman, Y., Tal, A.: Online dynamic graph drawing. IEEE Trans. Vis. Comput. Graph. **14**(4), 727–740 (2008)
14. Fruchterman, T.M.J., Reingold, E.M.: Graph drawing by force-directed placement. Softw., Pract. Exper. **21**(11), 1129–1164 (1991)
15. Gorochowski, 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)
16. Grabowicz, P.A., Aiello, L.M., Menczer, F.: Fast filtering and animation of large dynamic networks. EPJ Data Science **3**(1), 27 (2014)
17. Hadlak, S., Schulz, H., Schumann, H.: In situ exploration of large dynamic networks. IEEE Trans. Vis. Comput. Graph. **17**(12), 2334–2343 (2011)
18. Hall, K.M.: An r-Dimensional Quadratic Placement Algorithm. Management Science **17**(3) (1970)
19. Koren, Y.: Drawing graphs by eigenvectors: theory and practice. Computers and Mathematics with Applications **49**(11), 1867 – 1888 (2005)
20. Krzakala, F., Moore, C., Mossel, E., Neeman, J., Sly, A., Zdeborová, L., Zhang, P.: Spectral redemption in clustering sparse networks. Proceedings of the National Academy of Sciences **110**(52), 20,935–20,940 (2013)
21. Peel, L., Clauset, A.: Detecting change points in the large-scale structure of evolving networks. In: Proc. of the Twenty-Ninth AAAI Conference on Artificial Intelligence, pp. 2914–2920 (2015)
22. Saade, A., Krzakala, F., Zdeborová, L.: Spectral clustering of graphs with the bethe hessian. In: Annual Conference on Neural Information Processing Systems (NIPS), pp. 406–414 (2014)
23. 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)
24. 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)
25. 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)
26. Xu, K.S., Kliger, M., III, A.O.H.: A regularized graph layout framework for dynamic network visualization. Data Min. Knowl. Discov. **27**(1), 84–116 (2013)
27. 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)