

MULTISCALE COMMUNITY MINING IN NETWORKS USING SPECTRAL GRAPH WAVELETS

Nicolas Tremblay, Pierre Borgnat

CNRS, Laboratoire de Physique (UMR 5672), École Normale Supérieure de Lyon, France

ABSTRACT

For data represented by networks, the community structure of the underlying graph is of great interest. A classical clustering problem is to uncover the overall “best” partition of nodes in communities. Here, a more elaborate description is proposed in which community structures are identified at different scales. To this end, we take advantage of the local and scale-dependent information encoded in graph wavelets. After new developments for the practical use of graph wavelets, studying proper scale boundaries and parameters and introducing scaling functions, we propose a method to mine for communities in complex networks in a scale-dependent manner. It relies on classifying nodes according to their wavelets or scaling functions, using a scale-dependent modularity function. An example on a graph benchmark having hierarchical communities shows that we estimate successfully its multiscale structure.

Index Terms— Graph wavelets, community mining, multiscale community, spectral clustering

1. INTRODUCTION

In a large number of applications, data are naturally represented as networks (or weighted graphs): social networks, sensor networks, Internet networks, neuronal networks, transportation networks, biological networks,... Analyzing such networks has been emerging in the last decade as a central issue in the study of complex networks [1, 2]. A striking property of many networks, and a common way of simplifying the network’s analysis, is their modular structure, i.e., there exists groups of nodes that are more connected with themselves than with the rest of the network; these groups are called communities [3]. As nodes in a same community tend to share common properties, community mining provides both a sketch of the structure of a network, and some insight on nodes’ properties. One issue in community mining is defining the scale at which one wants to analyze the network. Many algorithms (see the review [3]) are based on the optimisation of appropriate evaluation functions such as the popular modularity [4]

and generally discard this question or propose only *ad-hoc* discussions. Modularity for instance is known to favour an intrinsic scale of description [5, 6].

The present work develops a scale-dependent procedure which identifies community structures, i.e., that classifies nodes according to their community, at different scales. After one decides on a scale of interest (or a collection of scales), the objective is to mine for the communities at this(ese) scale(s). Using multiple scales will provide a fully multiscale community description of the network. Some authors have proposed multiscale community mining either based on random walk processes [7, 8], or on definitions of parametric modularities [9, 10]. Our proposition for community mining is to rely on the recent construction of graph wavelets based on spectral graph theory [11].

By nature, the wavelet associated to a node a and a scale s is local. It is centered around this node and spreads on its neighbourhood so that the larger is s , the larger is the spanned neighbourhood. In some sense, wavelets give an “egocentered” view of how a node “sees” the network at that scale. Taking advantage of this local information encoded in wavelets, we develop an approach that clusters together nodes whose local environments are similar, i.e., whose associated wavelets are correlated. Then, to uncover the community structure at a given scale, we circumvent the intrinsic scale issue of the classical modularity by defining a scale-dependent modularity (using wavelets) which leads to this scale’s community structure when it is maximized. While graph wavelets have been applied now to classical problems in signal processing (e.g., source estimation for EEG [12]) or image processing [13], it has never been used for the detection of community in networks.

Section 2 recalls useful background material: an example of network with multiscale communities, and elements of spectral graph theory and wavelets. Some contributions to the use of spectral graph wavelets are presented in section 3: in 3.1, we discuss that once a network is given, a proper choice of scale boundaries ends up with parameters for the band-pass filter defining the wavelets that are different from [11]. In 3.2, scaling functions are defined for graph wavelets, as they will prove more robust than wavelets for community mining (up to our knowledge, this is the first introduction of scaling functions for graph wavelets). In section 4, the algorithm for

This work was partially supported by the GDR ISIS and the CNRS/JSPS Grants “Benchmarking and statistical analysis tools for modern Internet and sensor network traffics,” 2010 and 2012.

multiscale community mining is described and it is applied on the benchmark in section 5. We conclude in section 6.

2. BACKGROUND MATERIAL

The present work relies on one side on the notion of multiscale communities in graphs, as discussed for instance in [8, 14], and on spectral graph wavelets as introduced in [11] on another. The reader is supposed to have some familiarity with both works. Still, this section recalls useful elements.

2.1. Model of network with multiscale communities

Following [8], the model of graph defined in [14], for which the multiscale community structure is known, is adopted as a benchmark to test multiscale community mining. The global density of links is controlled by a first parameter \bar{k} (which is the mean degree of the nodes). Then, the intra-community and inter-community relative density of links is fixed by a second parameter ρ . As an example, we choose $\bar{k} = 16$ and $\rho = 1$. We consider graphs of 640 nodes, divided in three hierarchical levels: there are 64 small communities of 10 nodes each (the finest scale) embedded in 16 communities of 40 nodes each (the intermediate scale), themselves embedded in 4 communities of 160 nodes each (the coarsest scale). A realization of the network is visualized in Fig. 1.

2.2. Spectral Graph Theory and Fourier Modes

Let $\mathcal{G} = (V, E, A)$ be a undirected weighted graph with V the set of nodes, E the set of edges, and A the weighted adjacency matrix such that $A_{ij} = A_{ji} \geq 0$ is the weight of the edge between nodes i and j . Note N the total number of nodes.

Let us define the graph's Laplacian matrix $L = D - A$ where D is a diagonal matrix with $D_{ii} = d_i = \sum_{j \neq i} A_{ij}$ the strength of node i . L is real symmetric, therefore diagonalisable: its spectrum is composed of its sorted eigenvalues $(\lambda_l)_{l=0 \dots N-1}$, so that $\lambda_0 \leq \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_{N-1}$; and of the matrix χ of its normalized eigenvectors: $\chi = (\chi_0 | \chi_1 | \dots | \chi_{N-1})$. Considering only connected graphs, the multiplicity of eigenvalue $\lambda_0 = 0$ is 1.

By analogy to the continuous Laplacian operator whose eigenfunctions are the continuous Fourier modes and eigenvalues their squared frequencies, χ is considered as the matrix of the graph's Fourier modes, and $(\sqrt{\lambda_l})_{l=0 \dots N-1}$ its set of associated "frequencies". More details are found also in [11].

2.3. The Graph Wavelets

Graph wavelets were defined in [11] using the graph Fourier modes. Let us note $\psi_{s,a}$ the wavelet centered around node a . Its construction is based on band-pass filters defined in the graph Fourier domain, generated by stretching a band-pass filter kernel $g(\cdot)$ by a scale parameter $s > 0$. The stretched filter has matrix representation $\hat{G}_s = \text{diag}(g(s\lambda_0), \dots, g(s\lambda_{N-1}))$

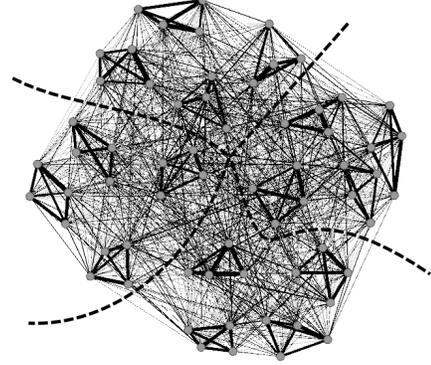


Fig. 1. Sketch of the graph discussed in 2.1: each node displayed is actually a community of 10 nodes. The thickness of each link is proportional to the total number of links between the two corresponding communities.

that is diagonal on the Fourier modes (the N eigenvectors of L). Hence, the wavelet basis at scale s reads as:

$$\Psi_s = (\psi_{s,0} | \psi_{s,1} | \dots | \psi_{s,N-1}) = \chi \hat{G}_s \chi^\top.$$

Here, our main use is of the localized wavelets $\psi_{s,a}$ themselves. Note that the wavelet transform at scale s of a signal f would be obtained by decomposing f on Ψ_s .

The intuition behind this definition of wavelets on graphs is that, at small scales (small scale parameter s), the filter $g(s \cdot)$ is stretched out, thus letting through high frequency modes essential to good localization: corresponding wavelets extends only to their close neighbourhood in the graph. At large scales (large s) the filter function is compressed around low frequency modes and this creates wavelets encoding a coarser description of the local environment.

We use the band-pass filter kernel g proposed in [11]:

$$g(x; \alpha, \beta, x_1, x_2) = \begin{cases} x_1^{-\alpha} x^\alpha & \text{for } x < x_1 \\ p(x) & \text{for } x_1 \leq x \leq x_2 \\ x_2^\beta x^{-\beta} & \text{for } x > x_2. \end{cases} \quad (1)$$

$p(x)$ is taken as the unique cubic polynomial interpolation that respects the continuity of g and its derivative g' .

We propose in the following a new way to set the parameters different from [11], by studying in 3.1 the scale boundaries that are relevant for community mining.

3. GRAPH WAVELET PARAMETERS AND GRAPH SCALING FUNCTION FOR COMMUNITY MINING

We first propose a new way to specify the parameters of g defining the graph wavelets, by studying the range for the scales s relevant for community mining. Then, we introduce graph scaling functions associated to graph wavelets.

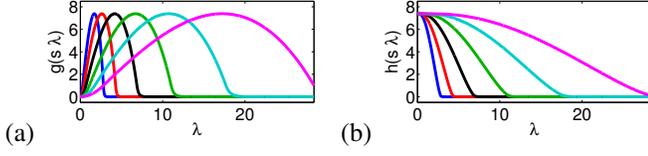


Fig. 2. (a) Band-pass filter functions g and (b) their associated low-pass filter functions h for six different scales within the scale boundaries: $s = 3.6 (s_{max}), 2.4, 1.5, 0.9, 0.6$ and $0.4 (s_{min})$; and the parameters from 3.1: $x_1 = 1, x_2 = 10, \alpha = 2$ and $\beta = 35$. Filters computed for the network in Fig.1.

3.1. Band-pass filter parameters and scale range

For the band-pass filter kernel g of eq. (1), parameters α, β, x_1 and x_2 have to be carefully chosen so as to generate appropriate wavelets, as well as the scales s to be used. Instead of following [11], the parameters are based here on an argument of spectral clustering of graphs [3]: the eigenvector χ_1 (associated to the smallest non-zero eigenvalue λ_1) is the first in importance for community mining because it contains information of the coarsest description of the graph.

A first consequence is that the maximum scale parameter s_{max} is set so that the filter function $g(s_{max} x)$ starts decaying as a power law only after $x = \lambda_1$, hence $s_{max} = x_2 / \lambda_1$. We require also that the filter at the maximum scale is highly selective around λ_1 ; for that, all other eigenmodes (especially λ_2) have to be attenuated. Choosing an attenuation by a factor 10, this leads us to: $g(s_{max} \lambda_1) = 10 g(s_{max} \lambda_2)$, hence $\beta = 1 / \log_{10} \left(\frac{\lambda_2}{\lambda_1} \right)$.

Second, we need to keep a part of χ_1 in the wavelets of every scale, so that all wavelets are sensitive to large scale community structure. We propose as minimum scale s_{min} the one for which $g(s_{min} \lambda_1)$ becomes smaller than 1. Using eq. (1), this gives $s_{min} = x_1 / \lambda_1$. Imposing also that $g(s_{min} \cdot)$ spans the whole range of eigenvalues (so that the wavelet basis is not blind to any eigenvalue), we should have $s_{min} \lambda_{N-1} = x_2$.

This argumentation gives us a value for β and three equations linking x_1, x_2, s_{min} and s_{max} . For x_1 and α describing the cut-off at low frequency, there is no argument from spectral clustering to set them. Following [11], we set $x_1 = 1$ and $\alpha = 2$. This in turn sets $x_2 = \lambda_{N-1} / \lambda_1$ and, thereby, $s_{max} = \lambda_{N-1} / \lambda_1^2$ and $s_{min} = \frac{1}{\lambda_1}$. Fig. 2.(a) shows examples of band-pass filters $g(s \cdot)$ when using the proposed range of scale and parameters, computed for the graph of Fig. 1.

3.2. Graph scaling functions

Let us introduce scale dependent scaling functions on graphs by analogy to the case of the continuous wavelet transform [15]. For that, we define the scale-dependent low-pass filters, noted \hat{H}_s , derived from a unique low-pass filter kernel h

stretched by the scale parameter s :

$$\hat{H}_s = \text{diag}(h(s\lambda_0), h(s\lambda_1), \dots, h(s\lambda_{N-1})).$$

For consistency with classical wavelet theory [15], we impose

$$h(x) = \left(\int_x^\infty \frac{|g(x')|^2}{x'} dx' \right)^{1/2}.$$

Then, the columns $\phi_{s,a}$ of $\Phi_s = \chi \hat{H}_s \chi^\top$ may be understood as scaling functions at scale s . Fig. 2.(b) shows the low-pass filters corresponding to these scaling functions. These graph scaling functions will be useful in the next sections, as they will lead to better results.

4. MULTISCALE COMMUNITY MINING

4.1. Elements of the method

Clustering of nodes according to their wavelets. For community mining, we aim at grouping together nodes whose environments are similar. As the local environment at a given scale s is encoded in the graph wavelets at each node a for this scale, we use $\psi_{s,a}$ as a feature vector for the local view from node a . A comparison of the views from two nodes a and b is obtained by taking the correlation distance (equal to 1 minus the correlation coefficient) between wavelets:

$$d_s^g(a, b) = 1 - (\tilde{\psi}_{s,a})^T \tilde{\psi}_{s,b}$$

where $\tilde{\psi}_{s,a}$ is the wavelet after normalization in energy.

In order to group nodes together, a hierarchical complete-linkage clustering algorithm [16] is used, with an additional connectivity constraint [17]: a node cannot be clustered in a group of nodes to which it has no path in the original network. This algorithm outputs a dendrogram. As we do not know beforehand how many clusters there are in the network, we have to evaluate each possible subdivision of the dendrogram and estimate which one is relevant.

Estimation of the number of communities with a scale-dependent modularity. As an evaluation function of the relevance of a given clustering, we introduce a scale-dependent definition of the modularity. Let S be a matrix of size $N \times J$ coding for a community clustering: $S = (\mathbb{1}_{C_1} | \mathbb{1}_{C_2} | \dots | \mathbb{1}_{C_J})$ where $\mathbb{1}_{C_j}$ is the binary indicator function of community C_j (i.e., $\mathbb{1}_{C_j i} = 1$ if node i is in C_j , else 0). The classical modularity matrix B is defined as [4]:

$$B = \frac{1}{2m} \left(A - \frac{dd^\top}{2m} \right)$$

where m is the sum of all weights in the graph and d the vector of strengths of the nodes. The usual modularity is then computed as: $Q = \text{tr}(S^\top B S)$. Here, we introduce a filtered version of B at each scale s :

$$B_s^g = \Phi_s^\top B = \chi \hat{G}_s \chi^\top B.$$

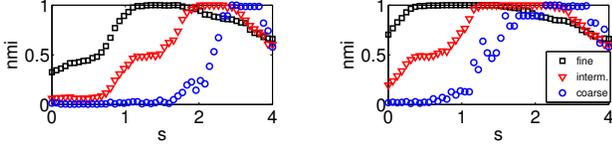


Fig. 3. Normalized mutual information (nmi) between the theoretical partition and the empirical one obtained when cutting the dendrogram with theoretical number of clusters. Left: using wavelets. Right: using scaling functions.

Then, a scale-dependent modularity is $Q_s^g = \text{tr}(S^\top B_s^g S)$.

Choosing the subdivision of the dendrogram (from the clustering of nodes) that maximizes the scale-dependent modularity leads to a community structure at scale s . We thereby circumvent the existing problem of classical modularity which imposes its own uncontrolled scale [5, 6].

In this method, we may replace wavelet functions (with filter g) by scaling functions (with filter h). In that case, the average of the scaling functions is removed to compute the correlation distance d_s^h . Results with both choices will be shown and we discuss afterwards which one is to be preferred.

4.2. Proposed method for multiscale community mining

The proposed method is summarized as follows.

1. Choose one scale s in the range of scales $[s_{min}, s_{max}]$ relevant to study the community structure (see 3.1).
2. Compute the graph wavelets (resp. the scaling functions) at s for each node (as described in 2.3, resp. 3.2).
3. Cluster the nodes according to the correlation distance d_s^g between wavelets (resp. d_s^h between scaling functions), using a hierarchical complete-linkage clustering. The output is a dendrogram.
4. Evaluate each subdivision of the dendrogram with the filtered modularity, and keep the one yielding the maximum filtered modularity Q_s^g (resp. Q_s^h).

This outputs the community structure at scale s . To obtain a full multiscale description, repeat the steps with a sampling of the scales between the scale boundaries.

5. TEST OF THE METHOD ON A HIERARCHICAL BENCHMARK NETWORK

We test the algorithm on the model of network described in 2.1. Following the method of 4.2, we obtain in step 3 one dendrogram per scale. Let us first assess the validity of the dendrograms. For each of the three levels of description, we compute the normalized mutual information (nmi) [18] between the known theoretical partition and the empirical one obtained when cutting the dendrogram with *a priori* knowledge of the total number of clusters. We plot it with respect to the scale parameter in Fig. 3. The black squares stands for the

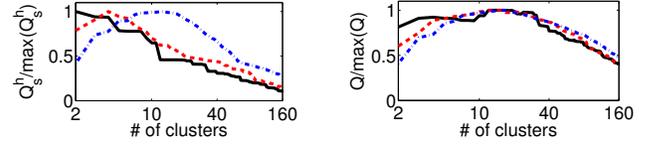


Fig. 4. Modularity: scale-dependent Q_s^h (left) or usual Q (right) with respect to the number of clusters for three different scales s ($s = 1.2$ in dot-dashed blue, 2.4 in dashed red, 3.3 in black line). The maximum should estimate the number of communities at each scale; it is correct only for Q_s^h .

fine level of description, the red triangles for the intermediate one, and the blue circles for the coarsest one. A nmi value of 1 means that both partition are the same. Results are obtained using wavelets (left), or scaling functions (right). We see that there always exists an interval of scales for which one may recover the theoretical partition from the dendrogram. When using scaling functions, the intervals at value 1 are larger, especially for small scale. This validates the first three steps of our method: the dendrograms do contain the relevant scale-dependent information.

A next question is: if we don't use our *a priori* knowledge of the theoretical number of clusters, does the maximisation of the proposed scale-dependent modularity estimate the correct number of clusters? Fig. 4 displays the scale-dependent (left) and classical (right) modularity with respect to the number of clusters kept in the dendrogram, for three different scales s . The maxima that define the best partition at each scale are significantly different. For classical modularity, it is always near the same number of clusters (hence pointing to communities having always the same mean size). For the scale-dependent modularity, it changes with s : the smaller s points to a higher number of clusters, hence smaller clusters. Here, we used Q_s^h but results are similar with wavelets. This illustrates why we need a scale-dependent modularity.

Quantitative results for the method are obtained by creating bootstrap samples of the network model of 2.1, by randomly adding $\pm 10\%$ of the weight of each link as in [8, 19]. We plot in Fig. 5 the average (over 100 bootstrap samples) of the estimated number of communities (top), and the average size of the uncovered communities (bottom). The dotted horizontal lines correspond to the three theoretical community structures. The method successfully recovers these different community structures, both in term of number of communities and size. Using scaling functions lead to better results, with larger and better defined intervals where these numbers and sizes are correctly estimated.

6. CONCLUSION

We proposed a method for multiscale community mining in graphs, for which no parameter needs to be adjusted. The

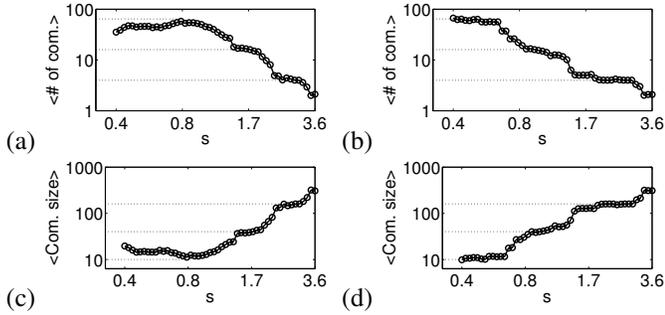


Fig. 5. (a), (b) Average (over 100 bootstrap samples of the graph) of the estimated number of communities. (c), (d) Average size of uncovered communities. Left: using wavelets. Right: using scaling functions.

mathematical soundness of graph wavelets, on which this method is based, is one of its great asset. The first aspect of our work was to complement certain aspects of graph wavelets (range of scale, parameters of the band-pass filter kernel, and corresponding graph scaling functions) for their application to community detection in graphs. For instance, scaling functions associated to graph wavelets provide slightly better results for scale-dependent community mining.

The weakness of the method is the computational cost. Two steps are costly: the diagonalisation of the Laplacian, and the evaluation of each possible subdivision of the dendrograms. Regarding the first problem, Hammond et. al [11] proposed a fast wavelet transform based on Chebyshev polynomial approximation that does not require the Laplacian's diagonalisation. This approximation will be easily used in our method. For the dendrogram mining, instead of looking at all possible subdivisions, future work may look only with subdivisions given by the largest gaps of the dendrograms, or, as the scale-dependent modularity has a bell shape curve (see Fig. 4), one could search only for a large local maximum.

7. REFERENCES

- [1] F.R.K. Chung and L. Lu, *Complex graphs and networks*, Number 107. Amer Mathematical Society, 2006.
- [2] A. Barrat, M. Barthlemy, and A. Vespignani, *Dynamical processes on complex networks*, Cambridge University Press, 2008.
- [3] S. Fortunato, "Community detection in graphs," *Physics Reports*, vol. 486, no. 3-5, pp. 75–174, 2010.
- [4] M.E.J. Newman, "Modularity and community structure in networks," *PNAS*, vol. 103, no. 23, pp. 8577, 2006.
- [5] S. Fortunato and M. Barthelemy, "Resolution limit in community detection," *PNAS*, vol. 104, no. 1, pp. 36, 2007.
- [6] J.M. Kumpula, J. Saramäki, K. Kaski, and J. Kertesz, "Limited resolution in complex network community detection with potts model approach," *Eur. Phys. J. B*, vol. 56, no. 1, pp. 41–45, 2007.
- [7] M.T. Schaub, J.C. Delvenne, S.N. Yaliraki, and M. Barahona, "Markov dynamics as a zooming lens for multiscale community detection: non clique-like communities and the field-of-view limit," *PloS one*, vol. 7, no. 2, pp. e32210, 2012.
- [8] R. Lambiotte, "Multi-scale modularity in complex networks," in *Proc. 8th Int. Symp. WiOpt*, 2010, pp. 546–553.
- [9] J. Reichardt and S. Bornholdt, "Statistical mechanics of community detection," *Physical Review E*, vol. 74, no. 1, pp. 016110, 2006.
- [10] A. Arenas, A. Fernandez, and S. Gomez, "Analysis of the structure of complex networks at different resolution levels," *New Journal of Physics*, vol. 10, no. 5, pp. 053039, 2008.
- [11] D.K. Hammond, P. Vandergheynst, and R. Gribonval, "Wavelets on graphs via spectral graph theory," *Applied and Computational Harmonic Analysis*, vol. 30, no. 2, pp. 129–150, 2011.
- [12] D.K. Hammond, B. Scherrer, and A. Malony, "Incorporating anatomical connectivity into eeg source estimation via sparse approximation with cortical graph wavelets," in *IEEE ICASSP*, 2012, pp. 573–576.
- [13] S.K. Narang, Y.H. Chao, and A. Ortega, "Graph-wavelet filterbanks for edge-aware image processing," in *IEEE SSP Workshop*, 2012, pp. 141–144.
- [14] M. Sales-Pardo, R. Guimera, A.A. Moreira, and L.A.N. Amaral, "Extracting the hierarchical organization of complex systems," *PNAS*, vol. 104, no. 39, pp. 15224–15229, 2007.
- [15] R. Carmona, W.L. Hwang, and B. Torrsani, *Practical Time-Frequency analysis*, Academic Press, 1998.
- [16] A.K. Jain, M.N. Murty, and P.J. Flynn, "Data clustering: a review," *ACM computing surveys (CSUR)*, vol. 31, no. 3, pp. 264–323, 1999.
- [17] L. Donetti and M.A. Munoz, "Detecting network communities: a new systematic and efficient algorithm," *J. Stat. Mech.*, vol. 2004, pp. P10012, 2004.
- [18] L. Danon, A. Diaz-Guilera, J. Duch, and A. Arenas, "Comparing community structure identification," *J. Stat. Mech.*, vol. 2005, pp. P09008, 2005.
- [19] M. Rosvall and C.T. Bergstrom, "Mapping change in large networks," *PloS one*, vol. 5, no. 1, pp. e8694, 2010.