

Online Barycenter Estimation of Large Weighted Graphs. A multiscale strategy.

Ioana Gavra · Laurent Risser

the date of receipt and acceptance should be inserted later

Abstract In this paper, we propose a new method to compute the barycenter of large weighted graphs endowed with probability measures on their vertex set. We suppose that the edge weights are distances between the nodes and that the probability measure on the nodes is related to events observed there. For instance, a graph can represent a subway network: its edge weights are the distance between two stations, and the observed events at each node are the subway users getting in or leaving the subway network at this station. The probability measure on the vertices does not need to be explicitly known. Our strategy only uses observed node related events to give more or less emphasis to the different nodes. Furthermore, the barycenter estimation can be updated in real time with each new event. We propose a multiscale extension of [8] where the described strategy is valid only for medium-sized graphs due to memory costs. Our multiscale approach is inspired from the geometrical decomposition of the barycenter in a Euclidean space: we apply a heuristic *divide et impera* strategy based on a preliminary clustering. Our strategy is finally assessed on road- and social-networks of up to 10^6 nodes. We show that its results compare well with [8] in terms of accuracy and stability on small graphs, and that it can additionally be used on large graphs even on standard laptops.

Keywords Large graphs analysis · Barycenter estimation · Online statistics

Ioana Gavra
Universit Toulouse 1 Capitole, Institut de Mathématiques de Toulouse
E-mail: Ioana.Gavra@math.univ-toulouse.fr

Laurent Risser
CNRS, Institut de Mathématiques de Toulouse (UMR5219).
E-mail: lrissier@math.univ-toulouse.fr

1 Introduction

1.1 Context

Why graphs? Graph structures can model complex phenomena of high interest in a wide variety of domains and play an important role in various fields of data analysis. Although graphs have been used for quite a while in some fields, *e.g.* in sociology since the 1930's [14], the recent explosion of available data and computational resources boosted the importance of studying and understanding networks. Among the main application fields, one can count computer science (Web understanding [17] and representation), biology (neural or protein networks, genes), social sciences (analysis of citations graphs, social networks [11]), machine learning [9], statistical or quantum physics [5], marketing (consumers preference graphs) and computational linguistics [15].

Barycenter: motivation and applications. Singling out the most influential node or nodes can be seen as a first step to understand the structure of a network. Different notions of *node centrality* have been introduced to measure the influence or the importance of nodes of interest in a network. Centrality notions are sometimes related to the mean distance from each node to all others [1], to the degree of each node [7] or even to the eigenvalues of the graph's adjacency matrix [2]. A rather complete survey can be found in [3].

As far as the authors know, these notions of centrality do not take into account any weight on the nodes (but only on the edges), although there are numerous applications where this would be rather natural. For example, in the case of a metro network, when trying to establish a *central* station, it is quite reasonable to take into account the number of passengers that use each station. In the case of a traffic network, the node-weight can model how many cars pass by a given intersection; in the case of a social network it can model the number of followers (or likes, or posts, *etc.*) of each individual.

To take this kind of information into account, throughout this paper, we interest ourselves to the barycenter of a graph with respect to a probability measure on the node set, as defined in our previous work [8]. As we will see later on, is a natural extension of the expected value on a Euclidean space. Furthermore, our algorithm is developed in an online context: it does not need the exact knowledge of the probability measure (the number of passengers that use each station), but only observations of this random distribution (we can see when a passenger uses a station), and can be easily updated at the arrival of a new observation.

Besides determining a central node, the knowledge of such a barycenter on a graph can be of multiple use. For example, from a statistical point of view, for a fixed graph, the computation of the barycenter using two data sets of observations could be used to determine if the two sets are sampled from the same probability measure (on the vertex set). Such a mean position can also be a preliminary step for a more detailed study, like the one provided by a generalization of a Principal Component Analysis, that could translate the

main statistical fluctuations among the nodes of the network. The barycenter can also be useful in graph representation, since setting the barycenter in a central position can provide an intuitive visualization.

1.2 Online graph barycenter estimation

1.2.1 Graph barycenter definition based on Fréchet means

Since the networks studied in this paper are finite, the node weights can be seen as a probability measure on the nodes set. For a probability measure defined on an Euclidean space, there are two classical notions of centrality: the median and the Euclidean mean. Defining an average or central position in a non-euclidean metric space is not straight forward since the natural addition or averaging operations are not necessarily defined.

Back in 1948, M. Fréchet presented a possible answer to this problem, not only for the median and the mean of a probability measure, but for moments of all orders [6]. He introduced a notion of typical position of order p for a random variable Z defined on a general metric space (\mathcal{E}, d) and distributed according to any probability measure ν . This is now known as the p -Fréchet mean, or simply the p -mean, and is defined as:

$$M_\nu^{(p)} := \arg \min_{x \in \mathcal{E}} \mathbb{E}_{Z \sim \nu} [d^p(x, Z)].$$

This definition might seem counter-intuitive, but one can notice that this variational formulation also holds for real random variables. For example, if Z is a random variable distributed according to a distribution ν on \mathbb{R}^d , its expected value, given by $m_\nu = \int_{\mathbb{R}^d} x d\nu(x)$ is also the point that minimizes:

$$x \mapsto \mathbb{E}_{Z \sim \nu} [|x - Z|^2].$$

Now, let $G = (N, E)$ denote a finite weighted graph, E its edges set and ν a probability measure on its nodes set N . The barycenter of a graph $G = (N, E)$ is then naturally defined as the 2-Fréchet mean, that we simply denote Fréchet mean:

$$M_\nu = \operatorname{argmin}_{x \in N} \sum_{y \in N} d^2(x, y) \nu(y).$$

One should remark that the Frchet mean of a weighed graph is not necessarily unique.

1.2.2 Online estimation framework

We place ourselves in the online estimation framework, in the sense that we suppose that the probability measure ν is unknown. A sequence $(Y_n)_{n \geq 0}$ of i.i.d. random variables distributed according to ν is instead available. For instance, an observation Y_n can be interpreted as the access of a passenger to a given station for subway networks, the passage of a car on a given crossroad for traffic networks, or a paper download in a scientific social network.

2 Barycenter estimation using simulated annealing

In [8], the authors proposed a method to estimate the barycenter of weighted graphs. This algorithm is based on a simulated annealing algorithm with homogenization. We also established, in the same paper [8], its convergence from a theoretical point of view. Since this method is one of the corner stones of our current work, we briefly explain its principles and its main parameters in this section. We also give in Alg. 1 a simplified pseudo-code that explains how it practically works.

Algorithm 1: Graph barycenter estimation algorithm of [8]

Require: Continuous version of $G = (N, E)$, *i.e.* Γ_G .

Require: Observations sequence $Y = (Y_k)_{k \geq 1}$ on the nodes set N .

Require: Increasing inverse temperature $(\beta_t)_{t \geq 0}$ and intensity $(\alpha_t)_{t \geq 0}$.

- 1: Pick $X_0 \in \Gamma_G$ and set $K = \text{len}(Y) - 1$.¹
- 2: $T_0 = 0$.
- 3: **for** $k = 0 : K$ **do**
- 4: Generate T_k according to α_k .
- 5: Generate $\varepsilon_k \sim \mathcal{N}(0, \sqrt{T_k - T_{k-1}})$.
- 6: Randomly move X_k (Brownian motion): $X_k = X_k + \vec{h}_k \varepsilon_k$, where \vec{h}_k is a direction uniformly chosen among the directions departing from X_k , and ε_k is a step size.
- 7: Deterministically move X_k towards Y_{k+1} : $X_{k+1} = X_k + \beta_{T_k} \alpha_{T_k}^{-1} \overrightarrow{X_k Y_{k+1}}$, where $\overrightarrow{X_k Y_{k+1}}$ represents the shortest (geodesic) path from X_k to Y_{k+1} in Γ_G .
- 8: **end for**
- 9: **return** Graph location X_K estimated as the barycenter of Γ_G . We consider the nearest node to X_K in G as its barycenter.

¹ Here $\text{len}(Y)$ represents the length of the sequence $(Y_k)_{k \geq 1}$.

Simulated annealing is an optimization technique based on a gradient descent dynamic to which we add a random perturbation in order to help it escape local traps. The importance of this random perturbation is then decreased progressively in order to cool down the system and let the algorithm converge (or stabilize). This effect is parametrized by a continuous function $(\beta_t)_{t \geq 0}$, that represents the inverse of the so-called temperature schedule: when β_t is small, the system is *hot* and the random noise is quite important with respect to the gradient descent term. Then, when β_t goes to infinity, the random perturbation is negligible.

The convergence of the simulated annealing, to the set of global minimums, is guaranteed from a theoretical point of view for logarithmic evolutions of the temperature, *i.e.* $\beta_t = \beta \log t$. Large values of β increase the convergence rate of the algorithm. However, if its value is too large (above a certain threshold that in our case depends on the graph's structure), the algorithm might converge to a local minimum instead of a global one (see for example [10]). In practice, a linear growth ($\beta_t = \beta t$) is more commonly used, so we have tested both versions.




Fig. 1 (left) Example of discrete graph G , and (right) corresponding continuous version Γ_G . X_t represents the current position of the algorithm in Γ_G . In this example its closest node in G is the node B .

Another important parameter comes from the on-line aspect of the algorithm. In our model, we simulate the arrival times T_n of the observations Y_n by an inhomogeneous Poisson process $(N_t^\alpha)_{t \geq 0}$ ², where $(\alpha_t)_{t \geq 0}$ is a continuous and increasing function that describes the rate at which we use the sequence of observations $(Y_n)_{n \geq 0}$. We refer to $(\alpha_t)_{t \geq 0}$ as the *intensity* of the process. On one hand, and from a theoretical point of view, using more observations improves the algorithm's accuracy and convergence rate, so it may seem natural to use large values for α_t . On the other hand, in practice, observations can be costly and limited, so one would like to limit their use as much as possible.

We also emphasize that Alg1 runs on Γ_G , a continuous version of the discrete graph G , where each edge $e = (u, v)$ of length L_e is seen as an interval $[0, L_e]$ such that an extremity of this segment corresponds to one of the nodes of the edge (see illustration Fig. 2). The process X_t that represents the barycenter estimation at increasing times t , therefore lives on the graph edges and not just its vertices. Nevertheless, a current estimation of a central node can naturally be defined as the closest vertex to the position of X_t .

Note that Alg. 1 is described for K observations. In an on-line context, the algorithm can then be used in three different ways:

1. If more than $K + 1$ observations are known: those used in the algorithm can be randomly picked-up.
2. If less than $K + 1$ observations are known and we will not have access to additional observations:
Iteratively perform
 - (a) randomly shuffle the observations
 - (b) use the shuffled observations
 until Alg. 1 ends. This strategy will be used in our tests.
3. If less than $K + 1$ observations are known and we will have access to additional observations:
 - (a) Use Alg. 1 on currently known observations to have a first guess of the graph barycenter.
 - (b) Wait for new observations to make the barycenter estimation more accurate with additional iterations of Alg. 1.

² T_n is the n -th jumping time of the Poisson process N_t^α , $T_n := \inf\{t : N_t^\alpha = n\}$

The key issue with this strategy on large graphs, which motivates the present paper, is that the deterministic move (line 7 of Alg. 1) requires to compute the shortest path from X_k to Y_{k+1} . To achieve this, we use of a standard Dijkstra’s algorithm which is particularly demanding in terms of computational times, especially when computed $K + 1$ times. The solution of [8] is then to pre-compute once for all the shortest distances between all node pairs and then to use this information for a quick algorithm execution. Computing these distances is $\#N$ times slower than computing the shortest path between two nodes, where $\#N$ is the number of nodes in $G = (N, E)$. This solution then makes sense when $K + 1$ is larger than $\#N$, or when multiple runs of the algorithm will be performed on the same graph, *e.g.* in order to evaluate the barycenters related to different observation sets Y . The major drawback of this strategy is however that it requires to store a $\#N \times \#N$ matrix in memory, which is unrealistic when $\#N$ is large. Moreover, the algorithmic cost of a Dijkstra’s algorithm on our weighted graphs is anyway $\mathcal{O}(\#N^2)$ and therefore does not scale at all to large graphs. We then propose to use the multiscale solution described in the following section.

3 Multiscale barycenter estimation

3.1 General framework

Our method is motivated by a property of geometrical decomposition of the barycenter in Euclidean spaces. We will describe this property Section 3.2 and how we heuristically extend it to graph structures. In practice, an analysed graph $G = (N, E)$ will be clusterized into a collection of sub-graphs $G_i = (C_i, E_i)$ and its barycenter will then be estimated using the *Divide and Conquer* strategy given Alg. 2. The key idea behind this multiscale algorithm is that a first coarse barycenter of G , denoted \tilde{b} , is computed on \tilde{G} which is a subsampled representation of G . This computation can be performed with reasonable resources if \tilde{G} is sufficiently small. This coarse barycenter estimate is then made more accurate by using \tilde{G} , a multiscale representation of G where all nodes of G are used in the sub-graph of \tilde{b} , and only the nodes of \tilde{G} are used in the rest of the domain. One can then finely estimate the barycenter of G in a region where it is likely to belong, again by using a compact representation of G .

This algorithm is directly related to Alg. 1 and the graph partition properties introduced Section 3.2. Computation of the sub-graphs G_i from a partition on large graphs is also discussed Section 3.3. Importantly, items 2 and 4 (of Alg. 2) are however not as obvious as they may appear for two main reasons: (1) They should be scalable on large datasets, and (2) pertinent heuristics have to be used to define \tilde{G} and \hat{G} , the simplified versions of G , so that they lead to accurate barycenter estimates of G . These two items are presented Section 3.4 and Section 3.5, and are the main contribution of this paper in the context of online graph barycenter estimation.

Algorithm 2: Multiscale barycenter estimation**Require:** Graph $G = (N, E)$.

- 1: Partition $G = (N, E)$ is partitioned into I sub-graphs $G_i = (C_i, E_i)$.
- 2: Undersample $G = (N, E)$ in $\tilde{G} = (\tilde{N}, \tilde{E})$, where each node of \tilde{G} represents a compact description of sub-graph G_i .
- 3: Estimate the barycenter \tilde{b} of \tilde{G} using Alg. 1.
- 4: Compute a multiscale graph $\hat{G} = (\hat{C}, \hat{E})$ with the nodes of G in the subgraph of \tilde{b} and the nodes of \tilde{G} elsewhere.
- 5: Estimate the barycenter \bar{b} of \hat{G} using Alg. 1.
- 6: **return** Node \bar{b} estimated as the barycenter of G

3.2 Graph partition

It is well known that for n points, denoted $(A_i)_{i=1, \dots, n}$, of an affine space and an associated sequence of scalars $(a_i)_{i=1, \dots, n}$ of non-null sum, the barycenter is defined as the only point $G = \text{bar}((A_i, a_i))_{i=1, \dots, n}$ such that:

$$\sum_{i=1}^n a_i \overrightarrow{GA_i} = \vec{0}.$$

Suppose now that the points are partitioned into two sets I and J and that the corresponding (a_i) have a non-null sums. We then denote $G_I = \text{bar}((A_i, a_i))_{i \in I}$ and $G_J = \text{bar}((A_i, a_i))_{i \in J}$. The decomposition property states that the barycenter of the n points is the barycenter of the two sub-barycenters, meaning:

$$\text{bar}((A_i, a_i))_{i=1 \dots n} = \text{bar} \left((G_I, \sum_{i \in I} a_i), (G_J, \sum_{i \in J} a_i) \right). \quad (1)$$

This property can be iterated multiple times and still holds for k partitions of this type, $k \leq n$.

Our multiscale graph barycenter estimation strategy is directly inspired by this aggregation property. In order to use a similar method on graphs, we use partitions (or clusters) $(C_i)_{i=1 \dots k}$ of the nodes set N in G . We also define a sub-graph as:

Definition 1 For $C_i \subset N$, a subset of the nodes set, we call associated sub-graph G_i a graph $G_i = (C_i, E_i)$ formed by all edges of the initial graph G , connecting two points of C_i . In other words, the edges set of G_i is:

$$E_i = \{e = (e_-, e_+) \in E \mid e_-, e_+ \in C_i\}. \quad (2)$$

A partition $\mathbf{P} = (C_i)_{i=1 \dots k}$ is called valid and can be used to compute graph barycenters, if it satisfies the following properties:

1. The subsets C_i are disjoint and their union contains all the nodes, *i.e.*

$$N = \bigsqcup_{i=1}^k C_i;$$

2. The weight associated to each subset is non-null: $\forall 0 \leq i \leq k, \nu(C_i) \neq 0$;

3. For each subset C_i , the associated sub-graph G_i is connected.

Note that property (2) is implicit in our framework since ν charges each node, $\nu(y) > 0, \forall y \in N$. Property (3) is very important since the notion of barycenter in Alg. 2 is only defined for connected graphs. Interestingly, a wide variety of established clustering algorithms efficiently define valid partitions of the nodes, even on large graphs. For instance, [13] is a strategy that computes optimal graph partitions with respect to an energy, the modularity [16], using a greedy approach that scales to large graphs. We will therefore not develop this discussion in our paper and focus instead on the definition of \tilde{G} and \hat{G} . First we define neighboring clusters w.r.t. G .

Definition 2 Two disjoint subsets $C_i, C_j \subset N$ of the graph $G = (N, E)$ are neighboring clusters, denoted $C_i \sim C_j$, if there exists a pair of nodes $v_i \in C_i$ and $v_j \in C_j$ that are neighbors in G :

$$C_i \sim C_j \iff \exists v_i \in C_i, v_j \in C_j \text{ such that } (v_i, v_j) \in E. \quad (3)$$

In what follows, the information contained in the sub-graphs $G_i = (C_i, E_i)$ described above will be summarized in the graphs \tilde{G} and \hat{G} (the upscale and multi-scale versions of G). We can remark that the union of the edges in all clusters C_i does not contain all the edges of the initial graph. The remaining edges will however be used to define the edges of \tilde{G} and \hat{G} .

3.3 Computing the sub-graphs $G_i = (C_i, E_i)$

The collection of sub-graphs $(G_i = (C_i, E_i))_{1 \leq i \leq k}$ will be the key to subsample G in \tilde{G} and \hat{G} . Here, we consider as known the partition $(C_i)_{i=1 \dots k}$ of the node set N . We then use Alg. 3 to realistically compute the associated sub-graphs $G_i = (C_i, E_i)$.

At a first sight, the algorithmic cost of Alg. 3 appears to be $\mathcal{O}(\#E)$. Checking the cluster of e_- and e_+ (line 5) can however have an algorithmic cost $\mathcal{O}(\#N)$ if improperly coded. In our program, the node clusters are coded in Python dictionaries, making this task $\mathcal{O}(1)$ in average [12]. Using lower level programming languages, such as C++, the node identifiers could be first replaced by integers between 0 and $\#N - 1$ and then their clusters would be stored in a vector of size $\#N$, the cluster of node i being stored at the i^{th} entry of this vector. To be efficient this pre-treatment requires to sort the node labels; which has typically a cost $\mathcal{O}(\#N \log(\#N))$. This is for instance the case by using the standard C++ function `std::sort`. Once the node labels sorted, checking a node label is finally $\mathcal{O}(1)$, so this strategy also scales well to large graphs.

Finally, note that our algorithm not only computes the sub-graphs G_i but also a compact information of their boundaries B_i . More specifically, each element j of B_i stores pertinent information related to a node at the boundary of G_i and directly connected to a node in another sub-graph of G . This boundary information will help us define a scalable strategy to generate the edges of \tilde{G} and \hat{G} .

Algorithm 3: Sub-graphs $G_i(C_i, E_i)$, $i \in \{1, \dots, k\}$ computation.

Require: Graph $G = (N, E)$.
Require: Nodes partition $(C_i)_{i=1 \dots k}$.

- 1: **for** $i = 1 : k$ **do**
- 2: Create a void edge list E_i
- 3: Create a void list B_i of nodes information at the boundary of G_i
- 4: **end for**
- 5: **for all** $e = (e_-, e_+) \in E$ **do**
- 6: **if** e_- and e_+ are in the same cluster C_i **then**
- 7: Add e to E_i
- 8: **else**
- 9: We consider e_- in cluster i and e_+ in cluster j .
- 10: Add $[e_-, e_+, \text{weight}(e_-, e_+), j]$ to B_i .
- 11: Add $[e_+, e_-, \text{weight}(e_+, e_-), i]$ to B_j .
- 12: **end if**
- 13: **end for**
- 14: **return** The sub-graphs $G_i(C_i, E_i)$, $i \in \{1, \dots, k\}$.
- 15: **return** Sub-graphs boundary information B_i , $i \in \{1, \dots, k\}$.

3.4 Computing the subsampled graph \tilde{G}

We now present how to compute the subsampled graph $\tilde{G} = (\tilde{N}, \tilde{E})$ based on the collection of sub-graphs $G_i = (C_i, E_i)$ and the boundary information B_i , $i \in \{1, \dots, k\}$. In our strategy, each node \tilde{v}_i of \tilde{N} depends on the properties of G_i and each edge of \tilde{E} depends on the B_i . Every cluster C_i is represented in \tilde{N} by a single node v_i . The edge set of the new graph is defined by:

$$\tilde{E} = \{(v_i, v_j) \text{ with } v_i, v_j \in \tilde{N}, v_i \in C_i, v_j \in C_j \text{ and } C_i \sim C_j\}. \quad (4)$$

There exists an edge between two nodes if and only if their respective clusters are neighboring clusters and the length of each new edge is defined as the distance between its extremities in the subgraph $G_{ij} = (C_i \cup C_j, E_{ij})$, associated to $C_i \cup C_j$ (an example is illustrated in Fig. 2).

From a mathematical point of view, the probability associated to each node is the total probability of the cluster that contains it:

$$\forall v \in \tilde{N}, \quad \nu_{\tilde{G}}(v) = \nu(C_i), \quad \text{where } C_i \text{ is such that } v \in C_i. \quad (5)$$

The definition of the associated probability measure for the upscale graph in (5) is the analog of summing the scalars in the affine case in (1). Accessing independent random variables $(Y_n)_{n \geq 0}$ distributed according to ν , we can easily define another sequence $(\tilde{Y}_n)_{n \geq 0}$ of i.i.d. random variable of law $\nu_{\tilde{G}}$:

$$\tilde{Y}_n = c_i \in \tilde{N} \cap C_i \quad \text{if and only if} \quad Y_n \in C_i. \quad (6)$$

From the simulation point of view, when we have access to $(Y_n)_{n \geq 0}$, (6) means that every time a node in a cluster C_i is given by the sequence, we see it as the unique node c_i that represents the cluster in the upscale graph. An example of this procedure is illustrated on a simple graph in Figure 2 and further details are presented in the following subsections.




Fig. 2 On the left hand side we have an initial graph G partitioned in 4 communities (the nodes of each community are represented by different symbols and colors). We consider ν as the uniform probability on its nodes set and all edges have length 1. We take a subset \tilde{N} formed of nodes represented in a larger size. The upscale graph \tilde{G} , is represented on the right hand side, along with the length of the new edges and the new probability corresponding to each node. We have chosen to represent the distribution $\nu_{\tilde{G}}$, instead of a sequence $(\tilde{Y}_n)_{n \geq 0}$, because it is easier to visualize.

3.4.1 Computing the subsampled nodes \tilde{N}

A natural strategy to compute each subsampled node \tilde{v}_i is to define it as a barycenter of $G_i(C_i, E_i)$ using Alg. 1 (we remind the reader that the barycenter is not necessarily unique). This however requires to have a sufficient number of observations Y in C_i . We then instead randomly draw a node C_i to define \tilde{v}_i , which has also the advantage of having a negligible computational cost.

3.4.2 Computing the distance between \tilde{v}_i and the boundaries of G_i

Defining the distances between the nodes \tilde{N} in \tilde{G} is the trickiest step of our strategy. Consider two neighbor sub-graphs G_i and G_j . We will then define in next subsection the distance between \tilde{v}_i and \tilde{v}_j as equal to the shortest distance between \tilde{v}_i and \tilde{v}_j in the union of the sub-graphs of G_i and G_j , plus the edges of G linking these subgraphs that are saved in B_i and B_j .

We now explain how we compute, for each region i , the distance between \tilde{v}_i and all boundary nodes of sub-graph G_i , *i.e.* the $B_i[j][1]$, $j \in \{1, \dots, \#B_i\}$ (see lines 9 to 11 of Alg. 3). This could be done by running a Dijkstra's algorithm as discussed in the end of Section 2 for the whole graph, but it would be too costly. A fundamental remark here is that although this strategy is far too computationally consuming for the whole graph (algorithmic cost is $\mathcal{O}(\#N^2)$), it becomes realistic on much smaller sub-graphs $G_i(C_i, E_i)$. It can also be straightforwardly parallelized on different sub-graphs. After having computed the distances, we add them to the corresponding sub-lists of B_i . To sum up, each element j of the list B_i is finally a sub-list where:

- $B_i[j][1]$ is the node of G_i at the cluster boundary;
- $B_i[j][2]$ is the node outside of G_i linked to $B_i[j][1]$;

- $B_i[j][3]$ is the distance between $B_i[j][1]$ and $B_i[j][2]$ in G ;
- $B_i[j][4]$ is the cluster of $B_i[j][2]$;
- $B_i[j][5]$ is the distance between $B_i[j][1]$ and \tilde{v}_i .

3.4.3 Computing the subsampled edges \tilde{E}

The definition of the edges \tilde{E} only depends on the boundary information B_i , $i \in \{1, \dots, k\}$, as explained in Alg. 4.

Algorithm 4: Subsampled graph edges \tilde{E} computation.

Require: Clusters boundary information B_i , $i \in \{1, \dots, k\}$.

- 1: **for** $i = 1 : k$ **do**
- 2: **for** $j = 1 : \#B_i$ **do**
- 3: **if** $i < B_i[j][4]$ **then**
- 4: $\tilde{i} = B_i[j][4]$
- 5: Identify \tilde{j} so that $B_i[j][1] == B_{\tilde{i}}[\tilde{j}][2]$ and $B_i[j][2] == B_{\tilde{i}}[\tilde{j}][1]$.
- 6: Compute $TmpDist = B_i[j][3] + B_i[j][5] + B_{\tilde{i}}[\tilde{j}][5]$.
- 7: **if** \tilde{E} does not contain the edge $(\tilde{v}_i, \tilde{v}_{\tilde{i}})$ or its length is $> TmpDist$ **then**
- 8: Add or update edge $(\tilde{v}_i, \tilde{v}_{\tilde{i}})$ to \tilde{E} with distance $TmpDist$.
- 9: **end if**
- 10: **end if**
- 11: **end for**
- 12: **end for**
- 13: **return** Subsampled graph edges \tilde{E} .

Remark that the test line 3, of Alg. 4 prevents from avoid having multiple edges linking the same nodes as the graphs are undirected. This algorithm is again computationally reasonable as the main double for loop first depends the number of clusters and the number of nodes at the clusters boundaries. The instructions in this double loop are also reasonable as they are linearly related to limited number of edges and nodes.

3.4.4 Projecting the observations Y from the nodes of G to those of \tilde{G}

We recall that efficient techniques were described in Subsection 3.3 to find the sub-graph G_i associated to each node of N . We use the same method to project the observations Y on N to each node \tilde{v}_i of \tilde{G} . The node \tilde{v}_i indeed represents all the nodes C_i of G_i .

3.5 Multiscale graph

3.5.1 Motivation

A straightforward extension of the decomposition of the barycenter in the Euclidean case, to the context of graphs, can be described as follows: take a valid

partition, compute the barycenter of each cluster, create a new subgraph (as explained in Subsection 3.4) and finally compute its barycenter. This procedure induces thus a notion of centrality on the set of clusters.

Therefore, by choosing each subsampled node \tilde{v}_i as a Fréchet mean of C_i (in Subsection 3.4.1), creating the corresponding subsampled graph \tilde{G} as described above, and then estimating its barycenter, we obtain a central cluster. If the chosen partition has a specific meaning, this procedure can have an interest on its own, allowing us to study some larger scale properties of the graph. For example if each cluster C_i represents a community, this is a natural way of defining a central community.

Independently of the method used to define the starting points \tilde{N} (randomly chosen representatives of each subset C_i or estimated barycenter), since the graph does not have the same properties as the euclidean space, a barycenter \tilde{b} of the subsampled graph \tilde{G} is not necessarily a barycenter of the initial graph. However, for reasonable partitions, one might expect the Fréchet mean of the initial graph not to be far from the central community that contains \tilde{b} . This assumption motivates the next step in our approach: building the multiscale graph as detailed in Subsection 3.5.2

To sum up, if we are interested in a rough notion of centrality, finding a central cluster might be enough. To gain precision, we decided to add an extra-step to our method by reinserting a central cluster in the subsampled graph (line 4 of Alg 2). Finally, we consider as barycenter of the initial graph, the estimated barycenter of this new multiscale version (line 5 of Alg 2).

3.5.2 Definition of a multiscale graph

For a valid partition of the nodes set $P = (C_i)_{i \leq k}$, let $(G_i, \nu_i)_{i \leq k}$ denote the associate sub-graphs with their respective probabilities measures, defined in Subsection 3.3. With the notations introduced in Subsection 3.4, let $\tilde{G} = (\tilde{N}, \tilde{E})$ be an up-scale version of G corresponding to the partition P . In what follows we define $\hat{G} = (\hat{N}, \hat{E})$ the *multi-scale version* of G with respect to (\tilde{G}, C) , where C is an element of P , and $\hat{\nu}$ the corresponding probability measure. An example is illustrated in Fig.3.

The definition of the nodes set and the associated probability are straightforward. \hat{N} contains the nodes of \tilde{G} and C :

$$\hat{N} = \tilde{N} \cup C, \quad (7)$$

and $\hat{\nu}$ redistributes the mass of C over its nodes, while leaving the others values of $\nu_{\tilde{G}}$ unchanged:

$$\hat{\nu}(v) = \begin{cases} \nu(v) & \text{if } v \in C \\ \nu_{\tilde{G}}(v) & \text{if } v \in \tilde{N} \end{cases} \quad (8)$$

The edge set \hat{E} contains the edges of \tilde{G} , except those that were added to c , the node that represents the cluster C in the up-scale graph, and all internal

edges of C . On top of that we add new edges going from boundary of C to the nodes corresponding to its neighboring clusters:

$$\text{BorderEdges}(C, P) = \{(v, c_i) \mid v \in C, \exists v_i \in C_i \text{ with } (v, c_i) \in E\} \quad (9)$$

The length of such an edge is defined as the initial distance between its extremities in the subgraph corresponding to $C \cup C_i$. Now the set of edges \hat{E} can be written as:

$$\hat{E} = E_C \cup \left(\tilde{E} \setminus \{e \mid e \sim c\} \right) \cup \text{BorderEdges}(C, P), \quad (10)$$

where $e \sim c$ means that c is a node of e .

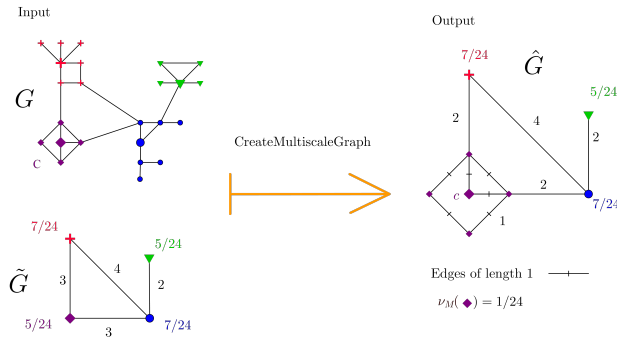


Fig. 3 \hat{G} is the multiscale version of G with respect to the up-scale graph \tilde{G} (shown in Fig. 2) and the cluster C (illustrated by diamond shaped violet nodes).

3.5.3 Computing the multiscale graph \hat{G}

As explained Alg. 2, the multiscale graph $\hat{G} = (\hat{C}, \hat{E})$ has the nodes of G in the subgraph of \hat{b} and the nodes of \tilde{G} elsewhere, where \hat{b} is the estimated the barycenter of \tilde{G} using Alg. 1.

We denote \tilde{i} the label of the sub-graph containing \tilde{b} . The construction of \hat{G} is then the same as the one of \tilde{G} , except in the sub-graph $G_{\tilde{i}}$, where the original nodes and edges are preserved. At the boundary between $G_{\tilde{i}}$ and the subsampled domain the distance given to the edge is slightly different to line 6 of Alg. 4. The distance between a sub-graph representative and the sub-graph boundary (*i.e.* $B_{\tilde{i}}[j][5]$ or $B_{\tilde{i}}[\tilde{j}][5]$) in $G_{\tilde{i}}$ is obviously not considered.

A fine representation of G is then constructed in $G_{\tilde{i}}$, the central cluster, and a coarse representation of G is constructed elsewhere. The goal of this multiscale graph is to make it possible to finely estimate the barycenter of G with reasonable computational resources.

Of course, from a simulation point of view, constructing a sequence of random variables distributed according to $\hat{\nu}$, the probability distribution defined in (8) is straightforward once we have access to $(Y_n)_{n \geq 0}$ of law ν . We simply set:

$$\hat{Y}_n = \begin{cases} Y_n, & \text{if } Y_n \in C \\ v_i, & \text{if } Y_n \in C_i. \end{cases} \quad (11)$$

3.6 Illustration on the Parisian subway

We now illustrate what can be G , \tilde{G} and \hat{G} on the Parisian subway network, which we already used in [8]. Fig. 4 represents the complete Parisian subway network. The graph was downloaded at <http://perso.esiee.fr/~coustyj/EnglishMorphoGraph/PS3.html>, has 296 nodes and 353 edges. The nodes obviously represent the metro stations. Each edge is a connection between two stations and its length is the time needed to go from one station to the other.




Fig. 4 Complete graph of Parisian Metro G . The coloured nodes are the estimated barycenters of the precomputed sub-graphs G_i and are be the nodes \tilde{N} of \tilde{G} (see Fig. 5).

In Fig. 5(top) we represent the subsampled graph \tilde{G} . One can see that the actual barycenter of the initial graph G (Chatelet) is not included in \tilde{G} , and thus can't be estimated as its center. It can however be estimated in the multiscale graph \hat{G} that is show Fig. 5(bottom). This example was designed to illustrate our strategy on a small graph. We explore more thoroughly its behavior on graphs of various sizes (from 10^3 to $> 10^6$ nodes) in section 4.




Fig. 5 (top) Subsampled Parisian Metro graph \tilde{G} , and **(bottom)** multiscale Parisian Metro graph \hat{G} . The width of the edges is inversely proportional to the time needed to go from one station to the other.

4 Results

4.1 Results on small graphs

In order to validate our strategy, we tested it on three small graphs for which we have access to the ground-truth barycenter. The first one is the Parisian metro network described in 3.6. The other two subgraphs of Facebook from the Stanford Large Network Dataset Collection³: **(FB2000)** has 2000 nodes and 37645 edges and **(FB4000)** has 4039 nodes and 88234 edges and fully contains (FB2000).

We performed three type of tests with default parameters:

- 1 Single scale estimation using Alg.1.
- 2 Multi-scale estimation; in the upscale graph each cluster is represented by its barycenters (estimated using Alg.1).

³ <https://snap.stanford.edu/data/>

	1 Single scale	2 Multi-scale	3 Multi-scale random
Paris Network	100%	97 %	97 %
FB2000	100%	100%	100%
FB4000	100%	80%	73 %

Table 1 Success ratio obtained with each strategy on 100 Monte Carlo runs for the Parisian metro network and the Facebook sub-graphs FB2000 and FB4000.

3 Multi-scale random; in the upscale graph each cluster is represented by a node sampled at random (uniformly) among its nodes.

We ran 100 Monte Carlo simulations for each strategy. A run is considered successful if the returned node is the true barycenter of the graph. We sum up the results in Table 1.

As one can see in Table 1, on the first two graphs, Parisian metro network and the FB2000, the performance of the algorithm does not seem influenced by the way we choose the nodes that represent each cluster in the upscale graph. On the third one, FB4000, even though the success ratio decreases slightly, the algorithm still performs rather well when the representative nodes in the preliminary phase are chosen at random. This is what motivated us to apply this strategy on larger graphs in order to reduce the computational cost and the number of used observations. We do not claim however that choosing the representative nodes at random is efficient in any framework. Indeed, depending on the graph’s structure and the initial partition, there exist cases where this first approximation is important and the choice of the representative node of a cluster could have a more significant impact on the quality of the results. In such cases it could be preferable to consider as representative, at least for clusters on which we have access to an important number of observations, their estimated barycenter.

4.2 New York urban area

After measuring the stability and accuracy of our method on small graphs, where the ground-truth barycenter is known, we have chosen to test it on a graph formed by the crossroads in a rather large New York urban area. We referred to it as such by convenience, but the area is not limited to the state of New York, see Figure 6. The graph has 264.346 nodes and 733.846 edges and can be found on the website of the Center for Discrete Mathematics and Theoretical Computer Science ⁴. On the website it is mentioned that some gaps might exist and thus the graph does not necessarily contain all crossroads. The nodes are the GPS coordinates of crossroads and the edges represent the streets between them. The distance considered between two nodes is the physical one, and not the transit time. However, the unit of measurement is not specified on the website. The graph’s diameter, in terms of euclidean distance between the GPS coordinates, is of 138.28 km. The intra-graph distance between the

⁴ <http://www.dis.uniroma1.it/challenge9/download.shtml>

corresponding nodes is 1 515 657. The average length of an edge in the graph is of 1296. From now on, all graph distances will be divided by 10^3 .

Furthermore, the graph is undirected, namely we consider that each street allows travel in both directions (even though in reality this not necessarily the case).




Fig. 6 New York urban area. Image obtained using Cytoscape and ©Google Maps. Purple points represent estimated centers of the GP700 partition. Nodes were located according to their GPS coordinates using the *CoordinatesLayout* plugin of *Cytoscape* (<http://apps.cytoscape.org/apps/coordinateslayout>).

We have performed two types of preliminary clustering. One, based on a bottom up approach, meant to provide clusters of homogeneous size, and another, based on a Markov Clustering algorithm developed by Stijn van Dongen [4]. The graph partition obtained with the first method has 700 clusters (from now on it will be referred to as GP700) and the second one has 1776 (we will refer to it as MCL12). Applying our method on these partitions enables us to estimate the barycenter on much smaller graphs. The corresponding subsampled graphs, as defined in Section 3.4, have 700 nodes (representing 0.2% of $\#N$) and 1682 edges, respectively 1776 (0.6% of $\#N$) nodes and 3916 edges.

In our tests, the multiscale versions, as defined in Section 3.5, for the GP700 partition have generally less than 1200 nodes (0.4% of $\#N$) and less than 2400 edges (0.3% of $\#E$). As for the MCL12 partition, the multiscale graphs have less than 2500 nodes (0.9% of $\#N$) and less than 5200 edges (0.7% of $\#E$).

Technical details In terms of memory cost, this kind of graph can definitely not be handled with the method proposed in [8]. A rough estimation suggests that the associated matrix distance would need around 360 GB of memory, whereas this new method employs far less resources, being of the order of 15GB (or less). The computational time is a bit long, but reasonable. Using the default parameters it takes, in average, 3 h 30 min for the GP700 and 7 h 30 min for the MCL12. This can be easily improved by paralelizing the computation of the subsampled graph \tilde{G} , namely the computation of its sub-sampled nodes (see sub-section 3.4.1) and the informations related to the borders (see subsection 3.4.2), needed for the computation of its edges (in Alg.4). Indeed, the computational times were mostly dedicated to the definition of the sub-sampled nodes \tilde{v}_i and their distance to the sub-graphs boundaries B_i . For instance, for GP700, computing the barycenter of the subsampled and multiscale graphs required about 20 and 50 seconds, respectively. This is almost negligible with respect to the total time of 3 h 30 min. Computational times are much lower if the \tilde{v}_i and the B_i are pre-computed and stored once for all. In a general manner, it is not surprising that the barycenter's estimation on GP700 is faster, because its clusters have a more homogeneous size and are thus easier to handle.

Since the graph is too dense to visualize, we have chosen to use the GP700 partition in order to facilitate Figure 6. To be more precise, we have used the upper scale approximation procedure described in Section 3.4 to form a new graph from the estimated centers of each cluster. A visualization using the GPS coordinates of the nodes was created with the help of the Cytoscape software. This illustration was afterwards overlaid on a map of the area provided by ©Google Maps. The result is shown in Figure 6. The purpose of this figure is to give an idea of the area covered by the complete graph and not to show the exact position of each node in the upper scale graph.

4.2.1 Results on two different partitions

We illustrate the results obtained in 5 Monte Carlo runs on the partitions GP700 and MCL12 in Figure 7. Since we have the GPS coordinates of each node, we used them to represent the estimated barycenters with ©Google Maps.

To measure the variability of the results, we computed the mean distance between each set of barycenters directly on the graph and the mean euclidean distance, using the GPS coordinates. For both partitions, the mean graph distance represents less than 3% of the graphs diameter and the mean euclidean

Partition	Mean distance on graph	Euclidean mean distance
GP700	35	2.44 km
MCL12	44	2.96 km

Table 2 Mean distance between the final centers obtained in 5 runs for each partition.

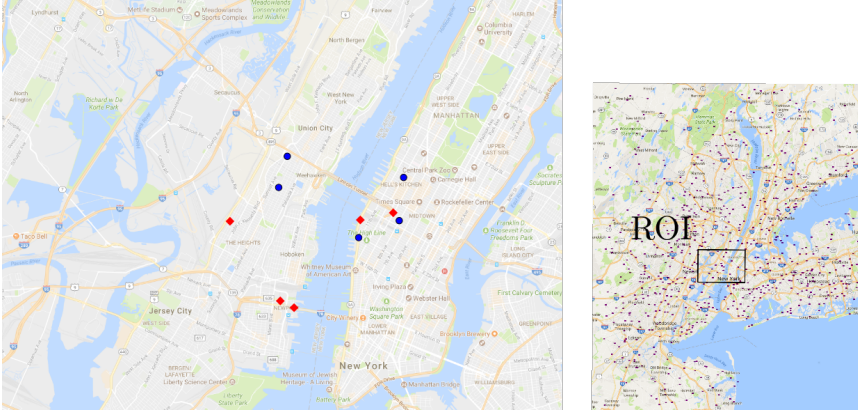


Fig. 7 On the left hand side we have a general view over the complete New York graph represented in Figure 6. On the right hand side, we have the Region Of Interest (ROI). The black round points are the barycenters obtained on the GP700 partition and the red diamond-shaped points are the barycenters obtained on MCL12. The right hand-side image was obtained using ©Geogebra and ©Google Maps.

distance represents approximately 2% of the euclidean diameter. The results are summarized in Table 2.

4.2.2 Parameters influence

In this section we illustrate the influence of the parameters on the results obtained on GP700. As explained in Section 2, the main parameters are: the temperature schedule β , the rate at which the observations are used α and the stopping time T . We consider as default parameters, and denote β^* , α^* and T_{\max}^* , the parameters introduced in [8].

The influence of these parameters is tested on two different strategies:

RANDOM CENTER LINEAR This is the strategy considered by default in this paper: for each cluster, a representative node is chosen uniformly at random among its vertex set and a linear growth of the temperature schedule is used in the estimation of the barycenters.

ESTIMATED CENTERS LOG The representative of each cluster is an estimation of its barycenter obtained with the method described in Alg. 1 and using a logarithmic temperature schedule.

Parameter	Estimated Center Log	Random Center Linear
default	55	35
$0.25\beta^*$	92	32
$0.5\beta^*$	88	21
$2\beta^*$	44	45
$4\beta^*$	44	56

Table 3 Mean distance between the final centers obtained in 5 runs for each temperature schedule.

Parameter	Estimated Center Log	Random Center Linear
default	55	35
$0.25T_{\max}^*$	68	33
$0.5T_{\max}^*$	70	35
$2T_{\max}^*$	33	60
$4T_{\max}^*$	54	104

Table 4 Mean distance between the final centers obtained in 5 runs for each set of parameters.

Influence of the temperature schedule. The temperature schedule β_t is a very important parameter linked to simulated annealing. As explained in Section 2, the convergence of the simulated annealing is guaranteed from a theoretical point of view for logarithmic evolutions of the temperature ($\beta_t = \beta \log t$). This is why this type of growth was used in [8] and therefore was also used to established the value of the constant β^* .

To test the influence of the variation of the constant β for both strategies (linear and logarithmic), We ran our algorithm 5 times for each set of parameters and computed the mean distance between the estimated barycenters in order to measure the stability. As one can see on Table 3, for a logarithmic schedule, increasing the constant β , reduces the variations of the results (because the algorithm converges faster). However for a linear evolution, increasing the value of the constant β makes the algorithm unstable (probably because the algorithm tends to converge to local minimums).

Influence of α and the stopping time. In our tests, the rate at which we use the observations is calibrated with respect to a stopping time. The stopping time, for a barycenter estimation using Algorithm 1, is chosen as a function of the number of nodes in the graph:

$$T_{\max}^* = 0.1\#V + 100.$$

The use of current observations is then distributed at a rate α_t that insures that between T_{\max}^* and T_{\max}^* we use approximately $S^* = 1000$ observations. In theory, a balance between the intensity α_t and the temperature β_t is mandatory for the convergence of the simulated annealing. So increasing the stopping time, without increasing accordingly S^* , reduces the intensity rate α . As shown in Table 4 this can be problematic, especially when we use a linear growth temperature.

Parameter	Estimated Center Log	Random Center Linear
default	55	35
$0.01O^*$	13	62
$0.1O^*$	58	81
$10O^*$	88	58
$100O^*$	82	44

Table 5 Mean distance between the final centers obtained in 5 runs for each set of parameters.

Influence of the observations number. As explained in Section 2, when we have access to a limited number of observations, there are multiple strategies available. For the current tests, we used strategy 2, namely we supposed that less than the necessary number of observations are available, and thus once the list is exhausted, we shuffle it and reuse it. This is equivalent to accurately estimating the barycenter of the multiscale graph, endowed with the probability measure $\hat{\nu}$, corresponding to the known observations.

For our tests, we generated a list of $O^* = 10000$ observations by choosing nodes uniformly at random among the vertex set. We then created other lists with more or less observations to measure the influence of this number on our method. The results are summarised in Table 5. For the linear temperature schedule, the variance seems to be quite stable with respect to the number of observations. The fact that for a logarithmic evolution, the variance increases with the number of observations is not surprising since, less observations imply a more concentrated probability measure $\hat{\nu}$, and thus its barycenter might be easier to estimate. However, the stability of the algorithm should not be regarded as the ultimate guarantee of the quality of its results. The more observations we use, the more we get closer to estimating the barycenter that corresponds to the uniform probability measure on the entire graph. And even though this bias is not very important for a uniform measure, we expected it to be more prominent on heterogeneous probability measures.

4.3 Social network

In order to test our strategy on a large graph representing a social network, we also applied our strategy on the com-Youtube graph of the Standard Large Network Dataset Collection⁵. This graph has 1,134,890 nodes and 2,987,624 edges. Each node represents a user and the edges represent user-to-user links. Of course, the data is anonymized. Each edge is of length 1 and the observations are uniformly sampled from the vertex set. We used the same bottom-up approach as for the New-York road network to clusterize the graph, and computed 5000 clusters. The subsampled network \tilde{G} has then 5000 nodes (0.9% of $\#N$) and about 410000 edges (14% of $\#E$) and the multiscale network \hat{G} has about 25000 nodes (2.2% of $\#N$) and about 460000 edges (16% of $\#E$).

⁵ <https://snap.stanford.edu/data/com-Youtube.html>

It is interesting to remark that due to the structure of this graph, the gains in term of edges is lower than for the New York road network. The clusters in a social network are indeed much more related than in a road network.

Results Using the default parameters, we ran the algorithm 4 times and we obtained three times the node 1072 as its barycenter and one time the node 663931. Remark that the subsampled graph \tilde{G} had 5000 nodes the multiscale graph \hat{G} typically had. These nodes are not directly connected and the distance between them is equal to 2. We could say that this distance represents an average closeness, since it is slightly lower than the mean distance between each of them and all other nodes of the graph, which is approximately 2.95 for 1072 and 3.51 for 663931. These results then appear as stable and thus promising. Computational time for one run of the algorithm was around 64 hours, but most of this was dedicated to the pre-computation of subsampled nodes \tilde{v}_i and their distance to the sub-graphs boundaries B_i . It took indeed about 58 hours to compute the \tilde{v}_i and B_i and about 1:10 and 7:05 hours to estimate the barycenters of \tilde{G} and \hat{G} , respectively. Computational times are then reasonable if the first step is performed once for all. More importantly, it is also doable on a small server with 64GB, which was our initial motivation.

4.4 Conclusion

Memory cost. From a computational point of view, the multiscale approach drastically reduces the memory costs and thus can be used on larger graphs. We indeed reduced the nodes number in the two large graphs we presented Sections 4.2 and 4.3 by factors of about 100 and 50, respectively. As shown in Section 4.2 on the New York road network, the variability of the graph barycenter estimate obtained on different runs is approximately equal to 2% of the graph diameter. This is a particularly satisfactory result with regards to the fact that these estimates were run on a compact representation of the graph that is about 100 times smaller than the original one. It therefore appears that the subsampling heuristics we used made sense, at least on a road-type network. On a social network graph, the results were also promising.

Computational time. Currently, the computational time is rather long, but as mentioned before, it can be reduced by parallelizing some of the intermediate procedures. Moreover, the most costly part of the algorithm is the conception of the subsampled graph and this step does not need to be done at the arrival of a new observation in the online context.

Online update of the barycenter. As mentioned before, the most time costly operation is the creation of the upscale graph. However, the actualization of the barycenter on the multiscale graph at the arrival of a new observation is instantaneous. If the informations regarding the upscale graph are stored, we could even reset the algorithm from this stage and update the barycenter

estimation on the subsampled graph. Assuming that the estimated central cluster would change with this new observation, creating a new multiscale graph and estimating a new barycenter on it (using the default parameters and the corresponding number of observations) would take less than 1 minute for the New York graph (with the partitions we used for our tests) and around 7 hours for the Youtube graph. The time needed for this operation depends a lot on the size of the clusters and not only on the size of the initial graph. For example, in our tests for the Youtube graph, computing the distances on the multiscale graph takes around 6 hours and estimating its barycenter only one.

A Package description

The *LGC_estim* package contains the strategy described in this paper. It is entirely written in Python and was tested using Python 2.7 and 3.4. Outside of Python modules that can be considered as standard (*Numpy*, *sys*, *os*, \dots), the only specific dependence of our package is *NetworkX*. This module is widely used for graphs management and analysis in Python⁶. Note that all our tests were made using the version 1.11 of NetworkX.

There are two ways to use the *LGC_estim* package: It was primarily designed to be used as a script but it can be alternatively used as a Python module. A *README* file at the package root directory explains how to use it in both cases through simple examples. Data files representing the Parisian subway network of section 3.6 are included in the *data* directory to run the examples. Note that by simply executing the command line *python LargeGraphCenterEstimator.py*, as for any Python script, a help message will give instructions to follow to properly estimate graph barycenters. Note finally that using the *LGC_estim* package as a Python module requires to understand the key classes and functions we used in our code, but is doable as shown in the *README* file.

Acknowledgements The authors thank Sébastien Gadat for fruitful discussions and helpful insights.

References

1. Bavelas, A.: Communication patterns in task-oriented groups. *Journal of the Acoustical Society of America* (1950)
2. Bonacich, P.: Factoring and weighting approaches to status scores and clique identification. *The Journal of Mathematical Sociology* **2**(1), 113–120 (1972)
3. Borgatti, S.P.: Centrality and network flow. *Social Networks* **1**, 55–71 (2005)
4. Dongen, S.: Graph Clustering by Flow Simulation. Utrecht (2000). [Http://www.library.uu.nl/digiarchief/dip/diss/1895620/inhoud.htm](http://www.library.uu.nl/digiarchief/dip/diss/1895620/inhoud.htm)
5. Estrada, E.: Introduction to Complex Networks. Structure and Dynamics, chapter of Evolutionary Equations with Applications to Natural Sciences. *Lecture Notes in Mathematics*. Springer (2015)
6. Fréchet, M.: Les éléments aléatoires de nature quelconque dans un espace distancié. *Annales de l'Institut Henri Poincaré (B)* **10**, 215–310 (1948)
7. Freeman, L.C.: Centrality in social networks conceptual clarification. *Social Networks* **1**(3), 215 – 239 (1978)
8. Gadat, S., Gavra, I., Risser, L.: How to calculate the barycenter of a weighted graph (to appear)

⁶ For installation, please go to <https://networkx.github.io/documentation/networkx-1.11/install.html>

9. Goldenberg, A., Zheng, A.X., Fienberg, S., Airoldi, E.M.: A survey of statistical network models. *Found. Trends Mach. Learn.* **2**(2), 129–233 (2010)
10. Hajek, B.: Cooling schedules for optimal annealing. *Mathematics of operations research* **13**(2), 311–329 (1988)
11. Handcock, M., K.Gile.: Modeling social networks from sampled data. *The Annals of Applied Statistics* (2010)
12. Hartley, J.: Time complexity in python (2018). URL <https://wiki.python.org/moin/TimeComplexity>
13. Lozano, E., Seal, M., Jelveh, Z.: Python package agglomcluster 2.0.4 (2018). URL <https://pypi.python.org/pypi/AgglomCluster/1.0.2>
14. Moreno, J.: *Who Shall Survive?* Nervous and Mental Disease Publishing Company (1934)
15. Nastase, V., Szpakowicz, S.: A study of two graph algorithms in topic-driven summarization pp. 29–32 (2006)
16. Newman, M.E.J.: Fast algorithm for detecting community structure in networks. *Phys. Rev. E* **69**, 066,133 (2004)
17. Pastor-Satorras, R., Vespignani, A.: *Evolution and structure of the internet : A statistical physics approach.* Cambridge University Press (2007)