



**HAL**  
open science

# Fast unfolding of communities in large networks: 15 years later

Vincent Blondel, Jean-Loup Guillaume, Renaud Lambiotte

► **To cite this version:**

Vincent Blondel, Jean-Loup Guillaume, Renaud Lambiotte. Fast unfolding of communities in large networks: 15 years later. *Journal of Statistical Mechanics: Theory and Experiment*, 2024, 2024 (10), pp.10R001. 10.1088/1742-5468/ad6139 . hal-04957823

**HAL Id: hal-04957823**

**<https://hal.science/hal-04957823v1>**

Submitted on 19 Feb 2025

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

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



Distributed under a Creative Commons Attribution 4.0 International License

# Fast unfolding of communities in large networks: 15 years later

Vincent Blondel<sup>1</sup>, Jean-Loup Guillaume<sup>2</sup> and Renaud Lambiotte<sup>3</sup>

<sup>1</sup>UCLouvain, Belgium, vincent.blondel@uclouvain.be

<sup>2</sup>L3i, La Rochelle Université, France, jean-loup.guillaume@univ-lr.fr

<sup>3</sup>Turing Institute, University of Oxford, London, UK, renaud.lambiotte@maths.ox.ac.uk

## Abstract

The Louvain method was proposed 15 years ago as a heuristic method for the fast detection of communities in large networks. During this period, it has emerged as one of the most popular methods for community detection, the task of partitioning vertices of a network into dense groups, usually called communities or clusters. Here, after a short introduction to the method, we give an overview of the different generalizations and modifications that have been proposed in the literature, and also survey the quality functions, beyond modularity, for which it has been implemented.

*Keywords:* Networks, Community Detection, Optimization, Louvain method

## 1 Introduction

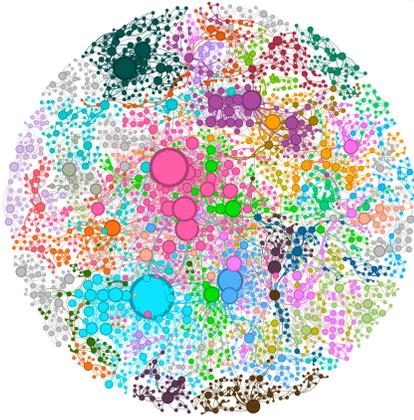
Networks provide a powerful language to model and analyse interacting systems, in a broad range of disciplines [1]. Facebook and food-webs, the brain and the job market are all systems where connectivity is critical, and where elements – here users, species, neurons and workers – are represented by vertices, while pairwise interactions – here friendship, predation, synchronization and job mobility – are represented by edges. Over the last two decades, several tools have been developed to extract information from the myriads of

interactions composing large networks. Among those, community detection has played a central role. Community detection aims at uncovering the community structure of a network, that is its organization into groups where vertices in the same group are strongly connected with each other, and weakly connected with vertices outside the group [2, 3]. Community detection is usually considered as an unsupervised clustering problem where the purpose is to find the best – according to a given definition – division into groups together with the most appropriate number of groups. Community detection is a rich and challenging problem, that has been considered from a variety of perspectives and principles [4] and finds many applications, from network visualization [5] to vertex classification [6]. In its most popular form, the communities are assumed to form a partition of the network, so that each vertex belongs to one single community. Partitioning methods have the advantage that their outcome is easier to interpret than that of overlapping methods, but also to involve less parameters to estimate and to have sound mathematical foundations. Even within the family of partitioning methods, several complementary approaches have been developed, leading to different functions to measure the quality of a partition defined in terms of e.g. the density of links inside communities [7], the number of links between communities [8], the confinement of random walkers inside communities [9, 10] or its likelihood of having generated the observed network [11].

A common challenge across methods is the need for efficient algorithms to optimize their respective quality function, that is finding the partition of the network with an optimal value. This problem is not new and finds its root in graph partitioning, where different algorithms have been proposed since the 1970s. It is more recently, arguably since the seminal works of Newman [7] that the research community started focusing on the related problem of community detection. It rapidly gained popularity due to the increasing access to large network data whose analysis required efficient clustering methods. The most popular quality function is the so-called Newman-Girvan modularity, and community detection is often considered - in practice - synonymous to modularity optimization, despite its limitations and the availability of alternative approaches. Until 2007, the most efficient methods for community detection were either spectral, i.e. defining modules based on dominant eigenvectors of matrices associated to the graph, or greedy, with the early works of Clauset et al. [12] and Wakita and Tsurumi [13]. However, the former are by construction limited to networks of the order of  $10^4$  vertices, orders of magnitude smaller than large-scale social and information networks, whereas

the proposed greedy methods did not perform sufficiently well in terms of accuracy, due to the ruggedness of the optimization landscape.

On March 4 2008, we posted on the arXiv an article entitled “Fast unfolding of communities in large networks ” proposing a greedy method for modularity optimization, and made available an efficient C++ code. A workshop was organized a few days later, on March 8, entitled “Detection and visualization of communities in large complex networks”, where we invited the leading researchers in the field to our research group in Louvain-la-Neuve. After a desk rejection from Nature and a swift rejection from PNAS, the paper was submitted to JSTAT, accepted on September 3 and published on October 9 [14]. 15 years later, the paper has been cited just over 20,000 times on Google Scholar, this is one new citation every 6 hours for more than 15 years. The method had no name in the original article, and was for a time named BGLL after the authors’ initials. The term ”Louvain method” appears for the first time in an article published in September 2009 [15]. The Louvain method, implemented in standard libraries for graph mining (see Table 1), is at the core of the network science ecosystem, with applications in several fields of science (see Figure 1). The purpose of this perspective paper is take a step back and to consider the legacy of the paper, by giving an overview of the research surrounding it, from improvements of the original method to its applications to other quality functions.



Lecture Notes in Computer Science	428
PLoS ONE	263
Scientific Reports	229
Physica A	215
ACM International Conference Proceeding Series	138
Studies in Computational Intelligence	135
IEEE Access	127
Physical Review E	114
Nature Communications	110
NeuroImage	104
Scientometrics	100
Applied Network Science	98
Social Network Analysis and Mining	96
Communications in Computer and Information Science	78
CEUR Workshop Proceedings	69
Sustainability	65
Expert Systems with Applications	60
PLoS Computational Biology	57
Information Sciences	54
PNAS	50

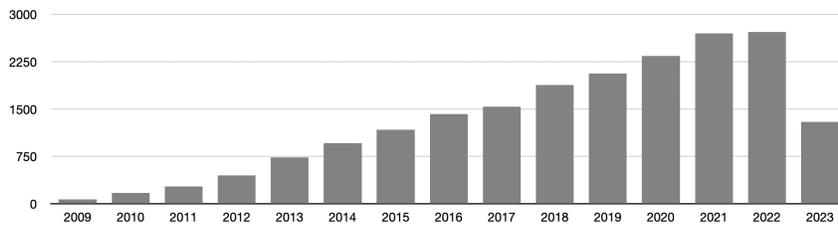


Figure 1: (Top left) Collaboration of researchers having cited the Louvain method. All articles citing [14] have been downloaded from Scopus on the 6th of April 2023. A bipartite network has been constructed relating author IDs, as identified by Scopus, and papers. The collaboration network was then constructed by first projecting the network according to [16], removing edges with a weight smaller than 0.2 and keeping the largest connected component. The resulting network is composed of 2490 vertices. Running the Louvain method identifies 39 communities, with a modularity of 0.926. The size of the vertices is proportional to their weighted degree. (Top right) Originally published in a Physics journal, the Louvain method has found applications in a variety of scientific domains, as shown by the list of the top 20 journals citing [14]. Further down in the list, one can find journals as diverse as Poetics (9 articles), Journal of Transport Geography (8 articles) and Nature Immunology (5 articles). (Bottom) Histogram of the total number of citations per year according to Google Scholar (20,054 citations in total on July 6 2023).

## Libraries

Language	Link	Objective function
Python	<a href="https://networkx.org/">https://networkx.org/</a>	mutiresolution $Q$ , directed
Python	<a href="https://networkit.github.io/">https://networkit.github.io/</a>	mutiresolution $Q$ , parallel
Python	<a href="https://github.com/vtraag/leidenalg">https://github.com/vtraag/leidenalg</a>	several functions, Leiden
Julia	<a href="https://github.com/afternone/CommunityDetection.jl">https://github.com/afternone/CommunityDetection.jl</a>	several functions
Javascript	<a href="https://graphology.github.io/">https://graphology.github.io/</a>	mutiresolution $Q$
Matlab	<a href="https://github.com/GenLouvain/GenLouvain">https://github.com/GenLouvain/GenLouvain</a>	several functions, multilayer
Java	<a href="https://github.com/jcnguyen/research-2015">https://github.com/jcnguyen/research-2015</a>	several functions
Multiple	<a href="https://igraph.org/">https://igraph.org/</a>	mutiresolution $Q$

5

## Softwares

Name	Objective	Link
Cytoscape (plugins)	Analysis and visualisation	<a href="https://cytoscape.org/">https://cytoscape.org/</a>
Gephi	Analysis and visualisation	<a href="https://gephi.org/">https://gephi.org/</a>
Tulip	Analysis and visualisation	<a href="https://tulip.labri.fr/site/">https://tulip.labri.fr/site/</a>
Pajek	Analysis and visualisation	<a href="http://mrvar.fdv.uni-lj.si/pajek/">http://mrvar.fdv.uni-lj.si/pajek/</a>
SAS Viya	Analytics	<a href="https://www.sas.com/en_us/home.html">https://www.sas.com/en_us/home.html</a>
Neo4j	Database	<a href="https://neo4j.com/fr/">https://neo4j.com/fr/</a>
Memgraph	Database	<a href="https://memgraph.com/">https://memgraph.com/</a>

Table 1: List of implementations of the Louvain method.

## 2 Basics on networks and modularity

A network is represented mathematically by a graph  $\mathcal{G}(V, E)$ , made of a set of vertices  $V$  of cardinality  $n \equiv |V|$ , and a set of edges  $E = \{\{i, j\} \mid i, j \in V\}$  of cardinality  $m \equiv |E|$ . Without loss of generality we will identify the vertex set  $V$  with the set  $\{0, \dots, n-1\}$ . In this introductory section, we will assume for the sake of simplicity that the network is undirected, so that edges do not have a direction, and unweighted, so that the graph only encodes the presence or absence of an edge between two vertices. Note that the concepts described in this section, and in particular expressions for modularity, have been generalized to more general settings, as we discuss in detail in Section 4.3. A useful way to encode and investigate the structure of a graph is its adjacency matrix. The adjacency matrix  $\mathbf{A}$  is an  $n \times n$  matrix encoding the presence or absence of an edge between any pair of vertices in the graph. Its entries  $A_{ij}$  are equal to 1 if  $i$  is adjacent to  $j$ , and 0 otherwise. For undirected graphs we have  $A_{ij} = A_{ji}$ , i.e., the adjacency matrix is symmetric ( $\mathbf{A} = \mathbf{A}^\top$ ). Several properties of a graph can be readily calculated from the adjacency matrix, including the degree of each vertex  $i$  as  $k_i = \sum_j A_{ij}$ , but above all, this representation makes it possible to exploit concepts and tools from linear algebra.

Since the seminal works of Watts and Strogatz [17], it has been recognised that real-world networks exhibit a mix of random and non-random patterns. Important types of structures, found in a wide range of empirical data, include clustering, here understood in terms of the density of triangles, fat-tailed degree distribution, power-law or not [18], and the presence of communities in a network [19], that is the presence of dense clusters that are weakly connected with one another. Community detection is an algorithmic problem aiming at finding the community structure of a network. As such, it can be seen as an unsupervised clustering problem, and some methods actually consist in first embedding the vertices in a space, before applying classical clustering methods, such as  $k$ -means [20]. In this section, we describe one of the most popular methods for community detection, despite some limitations described below, based on the notion of Newman-Girvan modularity [7], modularity in short.

Modularity is a measure of the quality of the partition of network into groups. Let us first consider a group of vertices defined by a set  $\alpha$ . The intuitive idea behind modularity is to compare the number of links inside  $\alpha$  with an expectation of this number under a random null model. The choice

of null model is in principle open and should depend on the mechanisms that constrain the formation of edges [21]. The most popular choice is, by far, the soft configuration or Chung-Lu model [22], a random graph model fixing, on expectation, the degrees of each vertex. For this choice, the expected number of edges between two vertices  $i$  and  $j$  is given by  $\frac{k_i k_j}{2m}$ , and the contribution of community  $\alpha$  to modularity is given by

$$\sum_{i,j \in \alpha} \left( A_{ij} - \frac{k_i k_j}{2m} \right). \quad (1)$$

Note that each term in the sum may either be positive or negative depending on the presence or absence of a link between two vertices. Intuitively, a good, dense community will be one with a high value for this expression. Note also that the null model penalises more pairs of vertices with a higher degree, which tends to favor configurations where high degree vertices are placed in different communities [10].

The Newman-Girvan modularity does not focus on one single community, but on the partition of the vertices into  $C$  communities  $\mathcal{P} = \{\alpha_1, \alpha_2, \dots, \alpha_C\}$ . The quality of a partition is obtained from the sum of Eq.(1) over the communities, leading to the expression

$$Q = \frac{1}{2m} \sum_{\alpha \in \mathcal{P}} \sum_{i,j \in \alpha} \left( A_{ij} - \frac{k_i k_j}{2m} \right), \quad (2)$$

where the additional prefactor  $1/(2m)$  ensures that modularity is in the interval  $[-1, 1]$ . Note that modularity can also be interpreted in terms of the assortativity coefficient, a correlation measure of vertex attributes, here the community labels, at the endpoints of edges in the network [23], in terms of a covariance defined on the graph [24] and in terms of the clustered covariance matrix from the random walk exploration of the graph [25]. Intuitively, modularity gives a high value to a partition if its communities concentrate an unexpectedly large number of edges inside its communities. In contrast with classical concepts such as the cut size, modularity allows to compare partitions made of different numbers of communities, and the partition maximising modularity is usually non-trivial, i.e. neither made of one large community or of  $n$  singletons. Modularity has become an essential element of several community detection methods. Modularity optimization methods aim at optimising modularity, that is finding the “best” partition of a network having the highest value of modularity. As modularity optimization is NP-hard [26],

several heuristics have been proposed for modularity optimization. Popular techniques include spectral methods [27] but large networks require faster, greedy methods proceeding by agglomerating groups of vertices into larger ones [28].

Before going further, note that the contribution of vertex  $i$ , assigned to a certain modularity  $\alpha$ , to the modularity of a partition is given by

$$Q_i = \frac{1}{2m} \left( \sum_{j \in \alpha} A_{ij} - \frac{k_i k_\alpha}{2m} \right), \quad (3)$$

where  $k_\alpha$  is the degree of community  $\alpha$ , defined as the sum of the degrees of its vertices. This expression shows that the contribution of vertex  $i$  requires only local information, its number of neighbors inside the community, its degree and the degrees inside the community. Locality is essential to design efficient greedy algorithm for modularity optimization.

Modularity, as a quality function, suffers from some limitations, which are inherited by any modularity optimization algorithm. Among those, let us note its tendency to over-fit, finding communities even in the case of random graphs [29], and to have, implicitly, a characteristic scale making it search for communities with a size that might not be compatible with the underlying structure [30]. Generalizations of modularity have been proposed to alleviate these limitations, for instance by incorporating a resolution parameter allowing to tune the characteristic size of the communities [31, 25, 32].

### 3 The Louvain method

The Louvain method is a greedy optimization heuristic originally proposed to optimize modularity, but its flexibility allows to optimize other quality functions as well, as described below. The approach consists in two phases that are iteratively repeated, until a maximum of modularity is obtained. The first phase consists in moving vertices from their community to that of their neighbors until a local maximum is reached. The second phase then consists in forming a new graph whose vertices are obtained by aggregating the vertices attached to the same community in the previous phase. A combination of the two phases is called a "pass". Starting from a network composed of  $n$  vertices, the number of vertices decreases at each pass, until no improvement of modularity can be obtained.

**First phase: vertex mover (VM).** The first phase begins with an undirected weighted graph having  $C$  vertices to which an index between 0 and  $C - 1$  has been randomly assigned.  $C$  is the number of vertices in the first pass, and the number of communities found in previous pass otherwise. One starts by placing each vertex into its own community, that is we start with a partition made of  $C$  communities. We then consider the first vertex, i.e. with index 0, and evaluate the change of modularity by removing 0 from its community and placing it in the community of each of its neighbors. Importantly for the speed of the method, this step only requires local information – See Eq.(3) –, and its number of operations is proportional to the degree of the vertex. Vertex 0 is then assigned to the community where the increase is maximum, if this maximum increase is positive. Otherwise, vertex 0 is left in its original community – See Algorithm 1 (all algorithms are adapted from [33]). This step is repeated sequentially to all the vertices based on their index. When reaching vertex  $C - 1$ , the process restarts at vertex 0, and the iteration continues until no vertex is moved during a complete iteration over the  $C$  vertices – See Algorithm 2. The phase ends in a local maximum, where greedy move of the vertices between community do not allow for an increase of the modularity. At this moment, a new partition of the vertices into  $C'$  communities has been produced. If  $C' \neq C$ , one proceeds to the next phase, and update the value of the number of communities. Else, the algorithm is finished and returns the partition as its estimate for optimum of modularity.

**Second phase: aggregation.** The second phase consists in building a new weighted graph whose vertices are the  $C$  communities found during the first phase. The weight of the link between two communities is given by the sum of the weights of the links between the vertices of these two communities. Moreover, links between the vertices of a same community lead to self-loops in the aggregated network. This operation is performed to ensure that the modularity of a partition in the aggregated graph has the same value as that of the same partition on the original graph. Once the second phase is finished, the first phase is reapplied on the aggregated network.

The output of the algorithm is the partition uncovered in the last pass, as it is the best partition found in terms of modularity – See Figure 2 and Algorithm 3. The method can also provide intermediary steps, that is the set of partitions of each pass, which provides a hierarchical representation of the network, each community at a pass being composed of a combination of combinations at the previous pass. Since its inception, the method has consistently been shown to provide a good balance between accuracy and

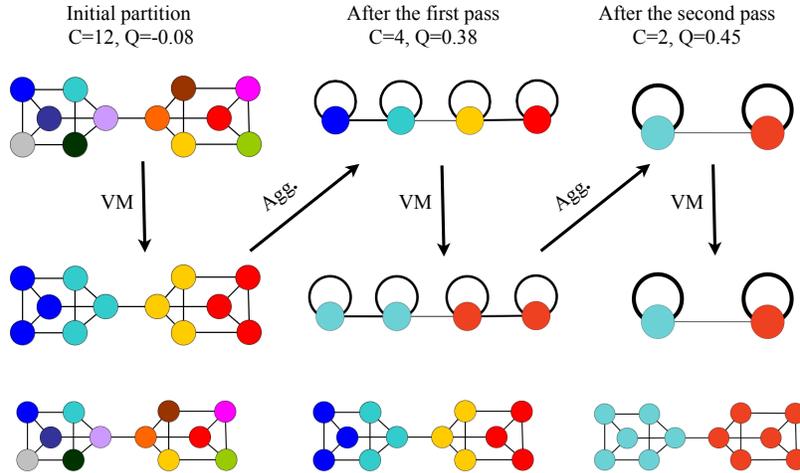


Figure 2: Each pass of the Louvain method is made of two phases. The first phase (Vertex Mover) consists in sequentially moving vertices to neighboring communities leading to a maximum increase of modularity. The second phase aggregates vertices and constructs a meta-graph whose vertices are the communities found after the first phase. The two phases are repeated until no improvement of modularity is observed. In this illustration, the algorithm converges after two passes, uncovering an optimal partition made of 2 communities, and a value of  $Q = 0.45$ . In this example, it is easy to show that aggregating further the two communities into one big community would result in a decline of modularity, as  $Q = 0$  in that case. Figure adapted from [34].

speed of execution, and to have a complexity close to linear in benchmarks and empirical networks. The speed of the method is ensured by the locality of each operation, as vertices only move to neighboring communities. Its accuracy arises from its flexibility, as vertices may be removed from their community and re-assigned to others in the process, and its multi-scale nature, as the community exploration is first done locally, and then over longer distance as the vertices are aggregated. Combined with its simplicity, these strengths have made the Louvain method a popular choice to optimize not only modularity, but also other quality functions, including the Map Equation [9] and Markov stability [32], as we discuss more in detail below. When applied on modularity, as is usually the case, it is important to remember

that the Louvain method inherits the limitations of modularity, including its resolution limit.

---

**Algorithm 1:** SingleVertexMover

---

**Require:**  $G = (V, E, w)$  a weighted graph

**Require:**  $\mathcal{P}$  a partition of  $V$  ( $\mathcal{P}[i]$  = community of vertex  $i$ )

**Ensure:**  $\Delta Q$ , the gain of modularity

**begin**

$c_{old} \leftarrow \mathcal{P}[i]$

    REMOVE( $i, c_{old}$ )

$\mathcal{C} \leftarrow \{\mathcal{P}[j] \mid (i, j) \in E\} \cup \{c_{old}\}$

$c_{new} \leftarrow \arg \max_{c \in \mathcal{C}} \{\text{GAIN}(i, c)\}$

**if**  $\text{GAIN}(i, c_{new}) - \text{LOSS}(i, c_{old}) > 0$  **then**

        INSERT( $i, c_{new}$ )

        return  $\text{GAIN}(i, c_{new}) - \text{LOSS}(i, c_{old})$

**else**

        INSERT( $i, c_{old}$ )

        return 0

---

## 4 Louvain++, a selection of enhancements

Because of its effectiveness and simplicity, the Louvain method has been extensively studied, and numerous modifications have been proposed, whether to improve results, further increase calculation speed or generalize its scope of use. We present these modifications in four different directions, each detailed in its respective subsection:

- the modification of the algorithm’s core, including its initialization, stopping criterion, vertex mover, etc;
- the use of quality functions other than modularity, functions which may be similar to modularity but with other parameters, or functions which may be completely different;

---

**Algorithm 2:** VertexMoverIteration

---

**Require:**  $G = (V, E, w)$  a weighted graph

**Require:**  $\epsilon$  a stopping criterion

**Ensure:** a partition  $\mathcal{P}$  of  $V$

**begin**

  INIT( $\mathcal{P}$ )

**do**

$increase \leftarrow 0$

**forall** vertices  $i$  **do**

$increase \leftarrow increase + \text{SingleVertexMover}(i)$

**while**  $increase > \epsilon$

  return  $\mathcal{P}$

---

---

**Algorithm 3:** Louvain algorithm

---

**Require:**  $G = (V, E, w)$  a weighted graph

**Ensure:** a partition  $\mathcal{P}$  of  $V$

**begin**

**repeat**

$\mathcal{P} \leftarrow \text{VertexMoverIteration}(G)$

$G \leftarrow \text{PartitionToGraph}(\mathcal{P}, G)$

**until** no improvement is possible

---

- generalizations of Louvain to deal with more than simple graphs. These may include weighted, oriented, spatially constrained, temporal networks, etc;
- and finally the use of parallelism to speed up computation time.

## 4.1 Modification of the algorithm

The original article and an extended version of it [34] (published in French in 2011 and in English in 2013) already described several improvements that have been proposed independently and studied in greater details by other authors:

- a partial optimization that consists in stopping the optimization when the modularity gain is below a given threshold (both for iterations or aggregation). This very simple optimization was already implemented in the initial version of Louvain;
- removing leaf vertices, which consists of deleting the leaves since they can only be placed in the community of their single neighbor;
- a VM that consists in not considering all vertices at each iteration;
- returning to lower levels to move again vertices;
- a simulated annealing-inspired VM allowing for non-optimal moves.

In the following, we describe modifications proposed in the different steps of the algorithm, that is its initialization, the vertex mover step and the aggregation step.

**Initialization:** Several types of optimization have targeted the initialization phase. The first one is to reduce the size of the graph before computation by removing a number of vertices that have little or no impact on the results. Based on the observation that a leaf (any vertex with a single neighbor) will inevitably be placed in its neighbor’s community, several authors propose deleting these leaves [34, 35] or, more generally, all tree structures, i.e. the initial leaves and all those that will appear when deleting leaves [36]. The resulting graph corresponds to the 2-core of the original graph and the leaves are added back once the communities are found. A generalization has been proposed in [37] where a k-core decomposition is performed and communities are only computed on the k-core. All the removed vertices are then added back and assigned to a community. The authors propose to use label propagation or any modularity maximization algorithm for this task.

Other approaches have modified the initial partition by not starting with communities containing a single vertex. The authors of [38] identify cliques containing at least 3 vertices that serve as initial communities. However, the way cliques are identified and the management of overlaps between these cliques are not made explicit.

**Vertex mover:** In the Louvain method, the best choice is always preferred during the vertex mover operation. Based on the principle of simulated annealing, it was proposed to choose a non optimal move, even with a negative gain, with a low probability rather than the optimal one [39].

Several articles propose to avoid considering all vertices at each iteration. The simplest version is defined in [40]. If a vertex does not move for several iterations (with a fixed parameter) then it will not be considered anymore. In [36], the authors show that if two communities  $A$  and  $B$  have not changed in the last iteration, then no vertex can move from  $A$  to  $B$  or from  $B$  to  $A$  in the current iteration. A similar idea is used in [41] where the author identify vertices that are more likely to move in the next iteration and only consider these vertices. They propose four situations, the simplest of which considers that if a vertex  $i$  has just moved from the community  $A$  to the community  $B$ , then all  $i$ 's neighbors who are not in  $B$  are likely to join  $i$  at the next iteration. The other three situations are a little more complex and are not used in practice. Limiting oneself to considering only a few vertices at the next iteration could lead to iterations being stopped prematurely and thus to a loss of quality, but in practice the impact is very slight.

In a different fashion, authors of [35] use a concept of seed vertices which are vertices more central than the average (for a given centrality measure defined in the paper). Non-seed vertices are then moved in priority to the neighboring community that contains a seed-vertex and maximizes the (positive) gain of modularity. If there is no such neighboring community or if no gain is positive, then the classical vertex mover is used.

In [42], instead of considering all neighbor communities and pick the best, one neighbor is picked at random and only its community is considered – other random selection procedure have also been considered, e.g. a community is selected proportionally to its size. Even though random selection will not always make the best choice, if a vertex has many neighbors in a given community then this community is more likely to be picked. The author notes that this operation accelerates convergence. This principle has been reused in the Leiden method [43] with other elements.

In most implementations of Louvain, the order in which vertices are considered is random. Although this can be seen as a problem, it does allow us to explore more possible solutions. If the aim is to obtain an identical partition, any fixed order can be used. In [44] the authors treat the vertices by decreasing centrality. This choice appears to improve results slightly, but above all, it speeds up the algorithm.

Other improvements have been proposed to simplify VM in a distributed computing context. We describe them in Section 4.4

**Aggregation and refinement:** A key element of the Louvain method is that the vertex mover operation is not greedy since any given vertex can

be moved several times until convergence is reached. As this step is not irreversible, it is less likely to be penalized by poor initial choices. However, the aggregation phases are irreversible, and the original method does not allow to go back once the graph has been aggregated. For this reason, a number of papers have proposed to challenge this principle and allow to move lower-level vertices even after aggregation. This procedure is called partition refinement.

In the simpler versions [45], a single refinement of the partition is performed on the first level, i.e., the original graph where all vertices are subject to the vertex mover procedure until convergence. In [46, 47], a multi-level refinement is used where, starting from the top level of the partition, the graph is unaggregated once and vertices are allowed to move. This process is repeated moving down the hierarchy until they are back on the initial graph. A different refinement is proposed in [48, 49] where the subgraph induced by each community is partitioned using one level or the complete Louvain method. All these parts of each communities are used to build the aggregated network. A more complex scheme is used in the Leiden method [43] where the vertex mover during refinement is limited and a move is accepted if it increases the modularity even though it is not the best move. Finally, in [50], if small (less than 4 vertices) or weak (that score low using their quality function) communities are found, these communities are broken into individual vertices and a VM step is applied once.

The Louvain method and its evolutions do not impose a priori constraints on the number or size of communities. If the number of communities is known to be  $k$ , then [51] propose to stop the aggregation in Louvain at the highest level that contains more than  $k$  communities (i.e. the next aggregation would result in fewer than  $k$  communities). They then use a greedy method to merge the communities that maximize the modularity gain to obtain exactly  $k$  communities. Finally, VM iterations are done until convergence.

## 4.2 Modification of the quality function

Despite the shortcomings of modularity, several quality functions have been inspired by it and used as a replacement in the Louvain method. For those, the algorithm remains unchanged after adapting the expression for the modularity change during the VM step. These quality functions mostly consider a positive contribution for intra-community edges and a negative contribution for inter-community edges or in the absence of edges.

- Several variations of modularity are using a resolution parameter, and can be shown to be equivalent, up to a trivial map, to the so-called Potts modularity [31]

$$Q_\gamma = \frac{1}{2m} \sum_{\alpha \in \mathcal{P}} \sum_{i,j \in \alpha} \left( A_{ij} - \gamma \frac{k_i k_j}{2m} \right), \quad (4)$$

where the resolution parameter  $\gamma$  allows to give more or less importance to the negative term, and hence to modulate the size of the communities. This expression has the same optimum as the so-called Markov stability, where the quality of a partition is estimated by its tendency to capture random walkers for long times inside communities [25, 32] and where time plays the role of the resolution parameter. Most implementations of Louvain have been generalized for one (or both) of these equivalent variations.

- The balanced modularity [52, 33] uses a symmetric version that also considers the absence of links

$$\frac{1}{2m} \sum_{i,j \in V} \left[ \left( a_{ij} - \frac{k_i k_j}{2m} \right) x_{ij} - \left( \bar{a}_{ij} - \frac{(n - k_i) \cdot (n - k_j)}{n^2 - 2m} \right) \bar{x}_{ij} \right].$$

- The authors of [39] use a function that considers both edges within communities and edges between communities. The linear version of their function is

$$\frac{1}{2m} \sum_{i,j \in V} \left[ \left( a_{ij} - \frac{k_i k_j}{2m} \right) x_{ij} - \beta \left( a_{ij} - \frac{k_i k_j}{2m} \right)^\alpha \bar{x}_{ij} \right],$$

where  $\alpha$  and  $\beta$  are resolution parameters that can be tuned to adjust the size of communities.

- In [53], the authors propose Jaccard Cosine Share Measure (JCSM). For two vertices  $i$  and  $j$ , the main term combines Jaccard (size of the intersection of neighborhoods of  $i$  and  $j$  divided by the size of their union), Cosine (normalized dot product of rows  $i$  and  $j$  of the adjacency matrix) and classical modularity ( $a_{ij} - k_i k_j / 2m$ ).

- Other functions mostly based on the number of connections within or between communities have been proposed. For instance [50] propose a new function defined as the sum for each community of

$$F2(C_i) = \frac{[d_{in}(C_i)]^2}{[d_{in}(C_i) + d_{out}(C_i)]^2}.$$

- Based on other principles than modularity, the map equation introduced in [9] uses the length of description of random walks in a network. The equation uses a two levels description to encode both the modules and the vertices. vertices from different modules can use the same code (the authors use the analogy of street names, which may be similar in different cities). The optimization of the map equation follows closely the steps of the Louvain method. Note that the map equation has also been generalized to uncover a hierarchy of communities [49].

More generally, works have explored the conditions that need to be satisfied by a quality function so that it can be optimized by the Louvain method. In that direction, the authors of [33] have used relational coding to describe several quality function using the same formalism. A quality function  $F$  is said to be linear if  $F(X, A) = \sum_{i,j \in V} \phi(a_{ij})x_{ij} + K$ , where  $a_{ij} = 1$  if  $i$  and  $j$  are connected (0 otherwise),  $x_{ij} = 1$  if  $i$  and  $j$  are in the same community (0 otherwise),  $\phi$  and  $K$  are respectively a function and a constant that only depends on the original data. For example, modularity can be expressed as  $Q = \sum_{i,j \in V} (a_{ij} - k_i k_j / 2m) x_{ij}$ . If a quality function is linear then the gain obtained using VM can be computed locally, henceforth without loss of efficiency. The authors also study separable quality function, i.e. that can be written as  $F(X, A) = \sum_{i,j \in V} \phi(a_{ij})\psi(x_{ij}) + K$ , where  $\psi$  only depends on  $X$ . If the function  $\psi$  only depends on the community of interest then the corresponding quality function can also be used efficiently in Louvain. The authors study five non-linear but separable function and clarify why only some can be optimized efficiently. A different formulation is given in [54]. The Louvain method can be used to optimize any quality function of the form  $trace H^T [F - ab^T] H$ , where  $H$  is the partition indicator matrix,  $F$  is a general matrix derived from the network, and  $a$  and  $b$  are two  $n$  dimensional vectors.

### 4.3 Louvain for non-simple graphs

Modularity has originally been defined for undirected, unweighted networks. Since then, several generalizations have been proposed to allow for more complex graph structures. As compared to the expression Eq.(3), the first term often remains unchanged, to count the total number, or total weight, of edges inside communities, but the second term is modified to take into account the additional constraints induced by the graph structure. For instance, in the case of directed networks, the null model is now chosen according to the directed configuration model, to properly estimate the expected number of edges between two vertices, given their respective in-degree and out-degree [55]. The Louvain method can then be directly generalized with an adapted modularity gain during VM [56]. Other examples include the optimization of modularity for spatial networks where, again, the null model of the quality function can be modified in order to account for spatial constraints [21]. Alternative approaches keep standard expressions for modularity but modify instead the Louvain method to impose a spatial contiguity of the clusters in the uncovered partition [57].

Community detection is also an active field of study for temporal networks [58, 59]. Compared to its static counterpart, communities are now dynamical objects that may encounter different types of change, such as growth or splitting, and algorithms have the additional task to characterize these events. Two different types of approach build on community optimization. First, communities can be uncovered in different snapshots. In that case, some constraints on the communities in adjacent snapshots may be incorporated in order to ensure their continuity in time, such as in modularity maximization under estrangement constraint [60]. After the communities have been found in each snapshot, additional methods determine, with a certain statistical accuracy, when and how a reorganization takes place [61]. A second family of methods aims at finding a decomposition of the network in an extended network in which snapshots are concatenated, e.g. by connecting the same vertex in consecutive snapshots to favor continuity of communities across time. Here, instead of searching communities in each time window, the community detection is performed in one single operation, which can be formulated as a modularity maximization, on a larger, multilayer network [62] and again optimized by the Louvain method [63].

Modularity has also been generalized to quantify the quality of the partition of signed networks. Here, the intuition is place together vertices such

that positive edges are concentrated within the modules and negative edges between the modules. In its simplest setting, the signed modularity consists in dividing the signed graph into two graphs, one defined by the positive interactions and another one defined by the negative interactions, and to search to maximize the difference of the modularities of the positive and negative graph [64]. The Louvain method can be directly generalized for this quality function, and the only step that needs to be considered with care is VM. Consider a vertex  $i$ . One expects to find other vertices of the community of  $i$  in the neighborhood of  $i$  only for the positive edges but its negative neighbors should instead be placed in other communities. Ways to address this complication are either to consider all the communities in the network in the VM step - and not only the neighboring ones - , but this involves a massive slow down of the algorithm, or to consider second neighbors, that is neighbors of neighbors, in the case of negative edges [65].

Finally, some works have focused on attributed networks where nodes carry labels. The aim is to create well-connected communities that are also homogeneous with regard to labels, in order to meet both structural and homophilic criteria. Three families of approaches are proposed in [66]: early fusion where a new graph is built to encode both topology and attributes, simultaneous fusion where the clustering is performed simultaneously on both domains and late fusion where both clusterings are computed independently before being joined. Among those, SAC2 [67] is an early fusion method where each node retains only its  $k$ -nearest neighbors based on a weighted combination of topology and similarity between attributes. Louvain is then used on this construction. Simultaneous fusion generally involves a composite function of modularity for the structural part and a similarity function for the attributes. This is the case in [68] using purity for attributes. Relatedly, SAC1 [67] uses different similarities for attributes, and [69] uses inertia. In all cases Louvain is used to optimize these functions. A late fusion is proposed in [70] where structural communities are found with Louvain and attributed ones with  $k$ -means. The structural communities are only used if they are sufficiently good, otherwise the attribute ones are used. More generally any consensus-based solution can be used to merge several partitions.

#### 4.4 Distributed/parallel/GPU solutions

Many efforts have been made to parallelize the Louvain method, the main challenge being to perform the VM step on several CPUs simultaneously.

There are two main paradigms: shared-memory implementations, where processes share the same memory and must therefore synchronize for writing, and distributed-memory implementations, where processors communicate via messages to maintain a consistent view. The former is limited by the number of physical processing units that can be used, and the latter by the quantity of messages sent. In most distributed implementations, if not all, the input network is partitioned so that vertices and edges are distributed evenly among the processing units. Evaluating the correctness of a parallel implementation can be complex. Because of the impact of the order in which vertices are processed, one does not expect to obtain the same results from one run to the next, or when compared with a sequential run.

**Shared-memory implementations:** In shared-memory implementations several vertices evaluate simultaneously and independently the best VM. They all base their choice on a structure that may no longer be valid at the time of the decision. The assignment of vertices to communities is protected and can only happen in a sequential way so that no inconsistency can occur but a vertex  $i$  may have been assigned between the time a vertex  $j$  made its choice and the time it validated it. Hence convergence might be slower or even not attainable. Several approaches have been considered.

- To keep the VM sequential, [71] propose an implementation by identifying the loops in the code that can be executed in parallel. This includes initialization of the algorithm, identifying all the neighboring communities of a vertex and finally identifying the neighboring community with the best gain. Community merging can also be partially parallelized.
- To avoid conflicts that may result in a slower convergence one solution would be to find vertices that do not conflict with each other and parallelize them. While the extra cost is generally considered too high [72], some authors propose to identify isolate sets of vertices, i.e. vertices that are not neighbors of each other, and to process them in parallel [73]. The authors of [72] propose to neglect this risk of non-convergence but limit the number of iterations. They also compare asynchronous moves with synchronous moves, where all parallelly processed vertices are moved simultaneously.
- The authors of [74] propose to adapt the number of processes dynamically depending on the number of neighbors of a vertex and available

resources.

- In [75], the authors use techniques already presented to avoid considering all vertices at every iteration. In the first iterations they use a pull method, which is the classic method where neighboring communities are recalculated at each iteration for every vertex. In the later iterations there are fewer vertex moves, they switch to a push method where neighboring communities are simply updated (pushed) with each vertex move. This optimization is combined with classic code parallelization, using locks for updates of neighboring communities.

**Distributed memory implementations:** In distributed memory implementations, vertices are distributed among the various processes. Each process will therefore store a given number of vertices. A number of problems arise that require algorithmic adaptation of the Louvain method. First, vertices or edges must be distributed evenly between processes to improve overall calculation time. In complex networks, the presence of very large degree vertices (or hubs) makes vertices/edges distribution more complex. Second, if the ends of an edge are not on the same process, messages will have to be sent to synchronize them and the total number of messages must be kept as low as possible. Finally, since calculations are made independently, it is necessary to be able to question local choices. Several works have looked at, and addressed, these issues.

- In [76], the authors use three levels of parallelism: (1) the network is divided in subgraphs and edges between these subgraphs are ignored. Louvain is applied on each subgraph. Then the individual results are joined, reintroducing the missing edges, and Louvain is reapplied. (2) vertices of each subnetwork are considered in batches and vertices of a given batch are processed in parallel. Potential conflicts if two vertices of a batch are neighbors are not considered but the update of vertex's community are still performed sequentially to avoid inconsistency. (3) Modularity gain for neighboring communities is computed in parallel.
- Processes can store knowledge of the neighbors of the vertices it owns even though such vertices are stored on another device. Such vertices are called "ghost vertices". Similarly, a process can store information about "ghost" communities adjacent to the ones it owns. Ghost vertices and communities are used in [77]. Initially, contiguous block of

vertices are distributed trying to balance the number of edges received by each process. At each iteration, some messages are exchanged between processes owning ghost vertices to update the information about community assignment. Graph coarsening is also performed in a distributed way so as to determine new global ids of communities and to exchange partial edge lists of neighbors communities so that each process can have a complete edge list. A similar procedure adapted for GPUs is detailed in [78]

- The authors of [79] use two hash tables to store the graph and also propose to identify the vertices that will provide the largest improvement of modularity and to only consider such vertices to accelerate the VM step. They show that the amount of such vertices decays exponentially fast so that only the first few iterations are costly.
- The initial distribution of vertices can have an impact on the workload but also on the number of ghost vertices and therefore on the quantity of messages to be transmitted. This is particularly the case for hubs so it is worth choosing the initial partition carefully. In [76], the authors distribute vertices evenly based on their id. This is studied in more details in [80] where the authors show that the degree distribution of such a partition is close to the degree distribution of the whole network, which ensures a fair workload. Then they propose an enhancement that also considers the degree of ghost vertices to improve workload fairness.
- To account for hubs it is also possible to duplicate them. Each duplicate (called a delegate) retains the neighbors belonging to its partition by default, but this can be modified to ensure a better balance of edges between partitions [81]. Then, delegates communicate with each other to ensure synchronization. This general principle is used in [82] as a preprocessing step for a distributed Louvain implementation. The authors also take delegates into account when calculating the modularity gain.
- The authors of [83] use PMETIS, a parallel implementation of METIS [84] to obtain a partition of the network and distribute these parts. Each one is then processed independently, then aggregation is performed and the rest of the algorithm is executed sequentially.

- Independent calculations on each process can result in incorrect communities if vertices are incorrectly assigned to a process. This problem is studied in [85] with the identification of doubtful vertices. The relative commitment for a vertex is calculated by computing the ratio between its internal degree and the maximum internal degree in its community, and a vertex’s affinity for a community is calculated by taking into account the relative commitment of the vertex and its neighbors. Vertices with low commitment are considered doubtful and can be moved to another device if they have many connections with vertices on that device.

## 5 Perspectives

In this article, we have overviewed the different ways in which the Louvain method has been used, generalised and improved over the last 15 years. As we have shown, a wide variety of works have built on its original design, exploring how to use it for other quality functions than modularity, and adapting its different steps for faster, more accurate and/or distributed solutions. A great deal is known about the strengths and weaknesses of the methods as well as its theoretical foundations, yet it remains an active source of inspiration as well as a standard method to cluster large-scale networks in practice. Never did we imagine the impact of the method when it was originally released and it is equally difficult for us to foresee now what will become of it in the years to come. The method has accompanied the different trends that have emerged in network science, from multiplex networks to embeddings and temporal networks. For this reason, it is likely that the Louvain method will find applications and generalisations in the increasingly popular field of higher-order networks [86, 87, 88, 89], but also in continuous generalisations of networks like graphons [90] and in extensions of graphs to non-real weights [91, 92]. But we are even more sure that we will be surprised, and we look forward to sitting down in 15 years, discussing together, discovering in awe the many ways the research community has modified our original idea and, who knows, sharing it with you [93].

## Acknowledgments

Our special thanks go to Etienne Lefebvre, without whom the Louvain method would never have existed. Only one publication as a master student, but what a success! We also thank John Pougué Biyong for his help to analyse the Scopus data and the many colleagues with whom we have collaborated on this topic over the year, in particular Thomas Aynaud, Mauricio Barahona, Romain Campigotto, Patricia Conde-Céspedes, Jean-Charles Delvenne, Mason Porter, Martin Rosvall and Michael Schaub. Finally, many thanks to Sébastien Amoury and Shazia Babul for carefully reading this manuscript.

The work of RL was supported by the EPSRC grants EP/V013068/1 and EP/V03474X/1. The work of JLG was supported by the ANR MITIK project, French National Research Agency (ANR), PRC AAPG2019.

## References

- [1] Newman M 2018 *Networks* (Oxford university press)
- [2] Fortunato S 2010 *Physics reports* **486** 75–174
- [3] Fortunato S and Newman M E 2022 *Nature Physics* **18** 848–850
- [4] Schaub M T, Delvenne J C, Rosvall M and Lambiotte R 2017 *Applied network science* **2** 1–13
- [5] Garza S E and Schaeffer S E 2019 *Physica A: Statistical Mechanics and its Applications* **534** 122058
- [6] Devooght R, Mantrach A, Kivimäki I, Bersini H, Jaimes A and Saerens M 2014 Random walks based modularity: application to semi-supervised learning *Proceedings of the 23rd international conference on World wide web* pp 213–224
- [7] Newman M E and Girvan M 2004 *Physical review E* **69** 026113
- [8] Von Luxburg U 2007 *Statistics and computing* **17** 395–416
- [9] Rosvall M and Bergstrom C 2008 *Proc. Natl. Acad. Sci. USA* **105** 1118–1123

- [10] Lambiotte R, Delvenne J C and Barahona M 2014 *IEEE Transactions on Network Science and Engineering* **1** 76–90
- [11] Peixoto T P 2013 *Physical review letters* **110** 148701
- [12] Clauset A, Newman M E and Moore C 2004 *Physical review E* **70** 066111
- [13] Wakita K and Tsurumi T 2007 Finding community structure in mega-scale social networks *Proceedings of the 16th international conference on World Wide Web* pp 1275–1276
- [14] Blondel V D, Guillaume J L, Lambiotte R and Lefebvre E 2008 *Journal of statistical mechanics: theory and experiment* **2008** P10008
- [15] Lambiotte R and Panzarasa P 2009 *Journal of Informetrics* **3** 180–190 ISSN 1751-1577 science of Science: Conceptualizations and Models of Science URL <https://www.sciencedirect.com/science/article/pii/S1751157709000248>
- [16] Newman M E 2001 *Physical review E* **64** 016132
- [17] Watts D J and Strogatz S H 1998 *nature* **393** 440–442
- [18] Broido A D and Clauset A 2019 *Nature communications* **10** 1017
- [19] Simon H A 1977 *Models of discovery: And other topics in the methods of science* 245–261
- [20] MacQueen J B 1967 Some methods for classification and analysis of multivariate observations *Proc. of the fifth Berkeley Symposium on Mathematical Statistics and Probability* vol 1 ed Cam L M L and Neyman J (University of California Press) pp 281–297
- [21] Expert P, Evans T S, Blondel V D and Lambiotte R 2011 *Proceedings of the National Academy of Sciences* **108** 7663–7668
- [22] Chung F and Lu L 2002 *Annals of combinatorics* **6** 125–145
- [23] Peel L, Delvenne J C and Lambiotte R 2018 *Proceedings of the National Academy of Sciences* **115** 4057–4062
- [24] Devriendt K, Martin-Gutierrez S and Lambiotte R 2022 *SIAM Review* **64** 343–359

- [25] Delvenne J C, Yaliraki S N and Barahona M 2010 *Proceedings of the national academy of sciences* **107** 12755–12760
- [26] Brandes U, Delling D, Gaertler M, Gorke R, Hoefer M, Nikoloski Z and Wagner D 2007 *IEEE transactions on knowledge and data engineering* **20** 172–188
- [27] Newman M E J 2013 *Physical Review E* **88** 042822
- [28] Fortunato S 2010 *Physics Reports* **486** 75–174 ISSN 0370-1573
- [29] Guimera R, Sales-Pardo M and Amaral L A N 2004 *Physical Review E* **70** 025101
- [30] Ghasemian A, Hosseinmardi H and Clauset A 2019 *IEEE Transactions on Knowledge and Data Engineering* **32** 1722–1735
- [31] Reichardt J and Bornholdt S 2006 *Physical review E* **74** 016110
- [32] Lambiotte R, Delvenne J C and Barahona M 2008 Laplacian dynamics and multiscale modular structure in networks arXiv 0812.1770
- [33] Campigotto R, Conde Céspedes P and Guillaume J L 2014 *arXiv preprint arXiv:1406.2518*
- [34] Aynaud T, Blondel V D, Guillaume J L and Lambiotte R 2013 *Multilevel Local Optimization of Modularity* (John Wiley & Sons, Ltd) chap 13, pp 315–345 ISBN 9781118601181 (*Preprint* <https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781118601181.ch13>) URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118601181.ch13>
- [35] Lui D, Huang K, Zhang C, Wu D and Wu S 2021 *Scientific Programming*
- [36] Zhang J, Fei J, Song X and Feng J 2021 *Mathematical Problems in Engineering*
- [37] Peng C, Kolda T G and Pinar A 2014 Accelerating community detection by using k-core subgraphs (*Preprint* 1403.2226)
- [38] Abbas E A and Nawaf H N 2020 Improving louvain algorithm by leveraging cliques for community detection *2020 International Conference on Computer Science and Software Engineering (CSASE)* pp 244–248

- [39] Zhou Z, Wang W and Wang L 2012 Community detection based on an improved modularity *Pattern Recognition* ed Liu C L, Zhang C and Wang L (Berlin, Heidelberg: Springer Berlin Heidelberg) pp 638–645 ISBN 978-3-642-33506-8
- [40] Ryu S and Kim D 2016 Quick community detection of big graph data using modified louvain algorithm *2016 IEEE 18th International Conference on High Performance Computing and Communications; IEEE 14th International Conference on Smart City; IEEE 2nd International Conference on Data Science and Systems (HPCC/SmartCity/DSS)* pp 1442–1445
- [41] Ozaki N, Tezuka H and Inaba M 2016 *International Journal of Computer and Electrical Engineering* **8** 207–218
- [42] Traag V A 2015 *Phys. Rev. E* **92**(3) 032801 URL <https://link.aps.org/doi/10.1103/PhysRevE.92.032801>
- [43] Traag V A, Waltman L and van Eck N J 2019 *Scientific Reports* **9** 5233 ISSN 2045-2322 URL <https://doi.org/10.1038/s41598-019-41695-z>
- [44] Aldabobo A, Sharieh A and Riad J 2022 *Journal of Theoretical and Applied Information Technology* **100**
- [45] Du W and He X 2016 A common strategy to improve community detection performance based on the nodes' property *Bio-inspired Computing – Theories and Applications* ed Gong M, Pan L, Song T and Zhang G (Singapore: Springer Singapore) pp 355–361 ISBN 978-981-10-3614-9
- [46] Gach O and Hao J K 2014 Improving the louvain algorithm for community detection with modularity maximization *Artificial Evolution: 11th International Conference, Evolution Artificielle, EA 2013, Bordeaux, France, October 21-23, 2013. Revised Selected Papers 11* (Springer) pp 145–156
- [47] Rotta R and Noack A 2011 *ACM J. Exp. Algorithmics* **16** ISSN 1084-6654 URL <https://doi.org/10.1145/1963190.1970376>
- [48] Waltman L and van Eck N J 2013 *The European Physical Journal B* **86**

- [49] Rosvall M and Bergstrom C T 2011 *PLoS ONE* **6** e18209 (*Preprint* 1010.0431)
- [50] Yao B, Zhu J, Ma P, Gao K and Ren X 2023 *Applied Sciences* **13** ISSN 2076-3417 URL <https://www.mdpi.com/2076-3417/13/6/4045>
- [51] Darst R K, Nussinov Z and Fortunato S 2014 *Phys. Rev. E* **89**(3) 032809 URL <https://link.aps.org/doi/10.1103/PhysRevE.89.032809>
- [52] Céspedes P C and Marcotorchino J F 2013 Comparing different modularization criteria using relational metric *Geometric Science of Information* ed Nielsen F and Barbaresco F (Berlin, Heidelberg: Springer Berlin Heidelberg) pp 180–187 ISBN 978-3-642-40020-9
- [53] Chaudhary L and Singh B 2019 Community detection using an enhanced louvain method in complex networks *Distributed Computing and Internet Technology* ed Fahrnberger G, Gopinathan S and Parida L (Cham: Springer International Publishing) pp 243–250 ISBN 978-3-030-05366-6
- [54] Schaub M, Delvenne J, Lambiotte R and Barahona M 2019 *Physical Review E* **99** 062308–1–062308–18 URL <http://dx.doi.org/10.1103/PhysRevE.99.062308>
- [55] Leicht E A and Newman M E 2008 *Physical review letters* **100** 118703
- [56] Li L, He X and Yan G 2018 Improved louvain method for directed networks *Intelligent Information Processing IX: 10th IFIP TC 12 International Conference, IIP 2018, Nanning, China, October 19-22, 2018, Proceedings 10* (Springer) pp 192–203
- [57] Wang C, Wang F and Onega T 2021 *Transactions in GIS* **25** 1065–1081
- [58] Cazabet R and Rossetti G 2019 *Temporal network theory* 181–197
- [59] Masuda N and Lambiotte R 2016 *A guide to temporal networks* (World Scientific)
- [60] Kawadia V and Sreenivasan S 2012 *Scientific reports* **2** 794
- [61] Greene D, Doyle D and Cunningham P 2010 Tracking the evolution of communities in dynamic social networks *2010 international conference on advances in social networks analysis and mining* (IEEE) pp 176–183

- [62] Mucha P J, Richardson T, Macon K, Porter M A and Onnela J P 2010 *science* **328** 876–878
- [63] Jutla I S, Jeub L G, Mucha P J *et al.* 2011 URL <http://netwiki.amath.unc.edu/GenLouvain>
- [64] Traag V A and Bruggeman J 2009 *Physical Review E* **80** 036115
- [65] Poug e-Biyong J and Lambiotte R *In preparation*
- [66] Chunaev P 2020 *Computer Science Review* **37** 100286 ISSN 1574-0137 URL <https://www.sciencedirect.com/science/article/pii/S1574013720303865>
- [67] Dang T and Viennet E 2012 Community detection based on structural and attribute similarities *International Conference on digital society* pp 7–12
- [68] Citraro S and Rossetti G 2020 *Applied Network Science* **5** 1–20 URL <https://appliednetsci.springeropen.com/articles/10.1007/s41109-020-00302-1>
- [69] Combe D, Langeron C, G ry M and Egyed-Zsigmond E 2015 I-Louvain: An Attributed Graph Clustering Method *Intelligent Data Analysis* (Saint-Etienne, France: LaHC, University of Saint-Etienne, France) URL <https://ujm.hal.science/ujm-01219447>
- [70] Elhadi H and Agam G 2013 Structure and attributes community detection: Comparative analysis of composite, ensemble and selection methods *Proceedings of the 7th Workshop on Social Network Mining and Analysis SNAKDD '13* (New York, NY, USA: Association for Computing Machinery) ISBN 9781450323307 URL <https://doi.org/10.1145/2501025.2501034>
- [71] Bhowmick S and Srinivasan S 2013 *A template for parallelizing the louvain method for modularity maximization* (Springer Basel) pp 111–124 Modeling and Simulation in Science, Engineering and Technology funding Information: This work was completed with the support of the College of Information Science and Technology, University of Nebraska

at Omaha (UNO), and the FIRE grant from the UNO Office of Research and Creative Activity. Publisher Copyright: © 2013, Springer Science+Business Media New York.

- [72] Shi J, Dhulipala L, Eisenstat D, Lăcki J and Mirrokni V 2021 *Proc. VLDB Endow.* **14** 2305–2313 ISSN 2150-8097 URL <https://doi.org/10.14778/3476249.3476282>
- [73] Qie H, Li S, Dou Y, Xu J, Xiong Y and Gao Z 2022 *Scientific Reports* **12**
- [74] Fazlali M, Moradi E and Tabatabaee Malazi H 2017 *Microprocessors and Microsystems* **54** 26–34 ISSN 0141-9331 URL <https://www.sciencedirect.com/science/article/pii/S014193311630240X>
- [75] Tithi J J, Stasiak A, Aananthakrishnan S and Petrini F 2020 Prune the unnecessary: Parallel pull-push louvain algorithms with automatic edge pruning *Proceedings of the 49th International Conference on Parallel Processing ICPP '20* (New York, NY, USA: Association for Computing Machinery) ISBN 9781450388160 URL <https://doi.org/10.1145/3404397.3404455>
- [76] Cheong C Y, Huynh H P, Lo D and Goh R S M 2013 Hierarchical parallel algorithm for modularity-based community detection using gpus *Proceedings of the 19th International Conference on Parallel Processing Euro-Par'13* (Berlin, Heidelberg: Springer-Verlag) p 775–787 ISBN 9783642400469
- [77] Ghosh S, Halappanavar M, Tumeo A, Kalyanaraman A, Lu H, Chavarrià-Miranda D, Khan A and Gebremedhin A 2018 Distributed louvain algorithm for graph community detection *2018 IEEE International Parallel and Distributed Processing Symposium (IPDPS)* pp 885–895
- [78] Gawande N, Ghosh S, Halappanavar M, Tumeo A and Kalyanaraman A 2022 *Parallel Computing* **111** 102898 ISSN 0167-8191 URL <https://www.sciencedirect.com/science/article/pii/S0167819122000060>
- [79] Que X, Checconi F, Petrini F and Gunnels J A 2015 Scalable community detection with the louvain algorithm *2015 IEEE International Parallel and Distributed Processing Symposium* pp 28–37

- [80] Zeng J and Yu H 2015 Parallel modularity-based community detection on large-scale graphs *Proceedings of the 2015 IEEE International Conference on Cluster Computing CLUSTER '15* (USA: IEEE Computer Society) p 1–10 ISBN 9781467365987 URL <https://doi.org/10.1109/CLUSTER.2015.11>
- [81] Pearce R, Gokhale M and Amato N M 2014 Faster parallel traversal of scale free graphs at extreme scale with vertex delegates *SC '14: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* pp 549–559
- [82] Zeng J and Yu H 2018 A scalable distributed louvain algorithm for large-scale graph community detection *2018 IEEE International Conference on Cluster Computing (CLUSTER)* pp 268–278
- [83] Wickramaarachchi C, Frincu M, Small P E and Prasanna V K 2014 *2014 IEEE High Performance Extreme Computing Conference (HPEC)* 1–6 URL <https://api.semanticscholar.org/CorpusID:11970878>
- [84] Karypis G and Kumar V 1998 *SIAM Journal on Scientific Computing* **20** 359–392
- [85] Bhowmick A, Vadhiyar S and PV V 2022 *Concurrency and Computation: Practice and Experience* **34** e6987 (Preprint <https://onlinelibrary.wiley.com/doi/pdf/10.1002/cpe.6987>) URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.6987>
- [86] Lambiotte R, Rosvall M and Scholtes I 2019 *Nature physics* **15** 313–320
- [87] Battiston F, Cencetti G, Iacopini I, Latora V, Lucas M, Patania A, Young J G and Petri G 2020 *Physics Reports* **874** 1–92
- [88] Bianconi G 2021 *Higher-order networks* (Cambridge University Press)
- [89] Bick C, Gross E, Harrington H A and Schaub M T 2023 *SIAM Review* **65** 686–731
- [90] Caron F and Fox E B 2017 *Journal of the Royal Statistical Society Series B: Statistical Methodology* **79** 1295–1366
- [91] Böttcher L and Porter M A 2022 *arXiv preprint arXiv:2212.06257*

- [92] Tian Y and Lambiotte R 2023 *arXiv preprint arXiv:2307.01813*
- [93] Blondel V, Guillaume J L and Lambiotte R 2038 *In progress*