

RESEARCH OUTPUTS / RÉSULTATS DE RECHERCHE

Flow-Based Community Detection in Hypergraphs

Eriksson, Anton; Carletti, Timoteo; Lambiotte, Renaud; Rojas, Alexis; Rosvall, Martin

Published in:
Higher-Order Systems

DOI:
https://doi.org/10.1007/978-3-030-91374-8_4

Publication date:
2022

[Link to publication](#)

Citation for pulished version (HARVARD):

Eriksson, A, Carletti, T, Lambiotte, R, Rojas, A & Rosvall, M 2022, Flow-Based Community Detection in Hypergraphs. in *Higher-Order Systems*. Understanding Complex Systems, Springer, pp. 141-161.
https://doi.org/10.1007/978-3-030-91374-8_4

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Flow-based Community Detection in Hypergraphs

Anton Eriksson, Timoteo Carletti, Renaud Lambiotte, Alexis Rojas, and Martin Rosvall

Abstract To connect structure, dynamics and function in systems with multibody interactions, network scientists model random walks on hypergraphs and identify communities that confine the walks for a long time. The two flow-based community-detection methods Markov stability and the map equation identify such communities based on different principles and search algorithms. But how similar are the resulting communities? We explain both methods' machinery applied to hypergraphs and compare them on synthetic and real-world hypergraphs using various hyperedge-size biased random walks and time scales. We find that the map equation is more sensitive to time-scale changes and that Markov stability is more sensitive to hyperedge-size biases.

With new introduction, experiments, and conclusion sections, this text reuses some text from Carletti, T., Fanelli, D., Lambiotte, R.: Random walks and community detection in hypergraphs. *J. Phys. Complex.* 2 015011 (2021) and Eriksson, A., Edler, D., Rojas, A., Rosvall, M.: Mapping flows on hypergraphs: How choosing random-walk model and network representation matters for community detection. *Comm. Phys.* (2021).

Anton Eriksson
Umeå University, Sweden, e-mail: anton.eriksson@umu.se

Timoteo Carletti
University of Namur, Belgium, e-mail: timoteo.carletti@unamur.be

Renaud Lambiotte
University of Oxford, United Kingdom, e-mail: renaud.lambiotte@maths.ox.ac.uk

Alexis Rojas
Umeå University, Sweden, e-mail: alexis.rojas-briceno@umu.se

Martin Rosvall
Umeå University, Sweden, e-mail: martin.rosvall@umu.se

1 Introduction

Researchers model and map flows on networks to identify important nodes and detect significant communities [6, 46, 24, 40, 16, 34]. From small to large system scales, random walk-based methods help to uncover the inner workings of networks [5, 23]. When standard network models with dyadic relations between pairs of nodes fail to adequately represent a system’s interactions, researchers turn to higher-order models of complex systems [32, 2], including multilayer networks [35, 14, 30, 12] for multitype interactions, memory networks [42, 45, 49] for multistep interactions, and simplicial complexes [37, 43, 25, 44] and hypergraphs [50, 10, 7, 8] for multibody interactions.

While several methods can identify flow-based communities in multilayer [35, 13, 26] and memory [42, 45, 49] networks with higher-order Markov dynamics, researchers have focused on combinatorial methods to identify communities in hypergraphs [1, 9, 33, 27, 28, 11] and only recently begun to unravel flow-based community structures associated with random walks guided by hyperedges sizes [8]. Two such methods are Markov stability and the map equation. Both algorithms exploit random walkers’ tendency to stay unexpectedly long times in assortative communities, albeit in different ways. Markov stability measures the fraction of random walks that reside in the community where they started after time t compared with stationarity. Instead, the map equation measures the shortest possible modular codelength required to describe the random walker on the network with given communities. Also their optimisation algorithms differ and it remains unclear how their detected communities in hypergraphs compare.

For Markov stability, we have previously analysed random walks on hypergraphs with a parameter controlling the bias of the dynamics towards small or large hyperedges and identified widely differing communities [8]. For the map equation and its optimisation algorithm Infomap [21], we have derived and clustered unipartite, bipartite, and multilayer network representations of hypergraph flows with different advantages [22]. Both papers highlight that the network and the research question should decide which model to use, but did not compare the two methods. For example, when modelling the flow of ideas with random walks in a co-authorship hypergraph, how does the organisation of authors in communities, at scales from research groups to research areas, change with community-detection method?

We compare Markov stability and the map equation applied to random walks on hypergraphs with hyperedge-size bias. After explaining the random-walk model and the flow-based community-detection methods using a schematic hypergraph for illustration, we consider experiments on three real-world hypergraphs: a zoo hypergraph with 101 nodes, a collaboration hypergraph with 361 nodes, and a fossil-record hypergraph with 13, 276 nodes. We find that the map equation responds faster to time-scale changes and that Markov stability is more sensitive to hyperedge-size biases.




Fig. 1 A schematic hypergraph in (a) projected into three different weighted networks with hyperedge-size bias $\sigma = -1$ in (b), 0 in (c), and 1 in (d).

2 Random walks on hypergraphs

We consider hypergraphs $\mathcal{H}(V, E)$ with n nodes $V = \{v_1, \dots, v_n\}$ and m hyperedges $E = \{E_1, \dots, E_m\}$. Each hyperedge $E_\alpha \subset V$ is an unordered collection of nodes. A hypergraph can be encoded by its incidence matrix $e_{i\alpha}$, where we use Roman indexes for nodes and Greek indexes for hyperedges:

$$e_{i\alpha} = \begin{cases} 1 & v_i \in E_\alpha \\ 0 & \text{otherwise} . \end{cases} \quad (1)$$

The $n \times n$ adjacency matrix of the hypergraph is $\mathbf{A} = \mathbf{e}\mathbf{e}^\top$, whose entry A_{ij} represents the number of hyperedges containing both nodes i and j with the diagonal elements set to zero. We also define the $m \times m$ hyperedge matrix $\mathbf{B} = \mathbf{e}^\top \mathbf{e}$, whose entry $B_{\alpha\beta}$ counts the number of nodes the two hyperedges share in the original hypergraph, $E_\alpha \cap E_\beta$. \mathbf{B} corresponds to the weighted adjacency matrix of the dual hypergraph, where hyperedges of the original hypergraph become nodes of the new structure.

A random walk process on a hypergraph can be defined by the transition probability, T_{ij} , allowing the walker to move across any pair of nodes (i, j) . The resulting continuous-time Markov process can be defined by

$$\dot{p}_i(t) = \sum_j p_j(t)T_{ji} - \sum_j p_i(t)T_{ij}, \quad (2)$$

where $p_i(t)$ is the probability of finding the walker on node i at time t . As often assumed when dealing with Markov processes, $\mathbf{p} = (p_1, \dots, p_n)$ is a row vector. With normalised transition probabilities such that $\sum_j T_{ij} = 1$, we rewrite the continuous-time Markov process as

$$\dot{p}_i = \sum_j p_j(T_{ji} - \delta_{ij}) = - \sum_j p_j L_{ji}, \quad (3)$$

where $L_{ij} := \delta_{ij} - T_{ij}$ is the random-walk Laplace operator.

Assigning weights to nodes and hyperedges biases the transition probabilities and leads to different probability flows between nodes [50, 7, 8].

Fixed weights and independent moves. [50]

Setting a weight $\omega(E_\alpha)$ to each hyperedge E_α and a weight $\gamma_{E_\alpha}(i)$ to each node $i \in E_\alpha$ allows us to decompose the random walk into three steps:

1. A random walker on node i chooses an incident hyperedge E_α proportional to its weight,

$$Q_{\alpha i} = \frac{\omega(E_\alpha) e_{i\alpha}}{\sum_{\beta} \omega(E_\beta) e_{i\beta}}. \quad (4)$$

2. The random walker picks a node $j \neq i$ of hyperedge E_α proportional to its weight,

$$R_{j\alpha} = \frac{\gamma_{E_\alpha}(j) e_{j\alpha}}{\sum_{\ell \neq i} \gamma_{E_\alpha}(\ell) e_{\ell\alpha}}. \quad (5)$$

3. The random walker moves to node j with probability

$$T_{ij} = \sum_{\alpha} R_{j\alpha} Q_{\alpha i}. \quad (6)$$

Several random-walk processes can be defined with various node and hyperedge weights. With hyperedge size as a proxy for the higher-order interactions' nature, we consider the transition probability

$$T_{ij}^{(\sigma)} = \frac{K_{ij}^{(\sigma)}}{\sum_{\ell \neq i} K_{i\ell}^{(\sigma)}} \quad \forall i \neq j \quad \text{and} \quad T_{ii}^{(\sigma)} = 0, \quad (7)$$

with

$$K_{ij}^{(\sigma)} = \sum_{\alpha} (B_{\alpha\alpha} - 1)^{\sigma} e_{i\alpha} e_{j\alpha} \quad \forall i \neq j \quad \text{and} \quad K_{ii}^{(\sigma)} = 0, \quad (8)$$

for some real σ , biases random walks to hyperedges depending on their size [7, 8]. By setting $\omega(E_\alpha) = (B_{\alpha\alpha} - 1)^{\sigma+1}$ and $\gamma_{E_\alpha}(j) = (B_{\alpha\alpha} - 1)$ for all $j \in E_\alpha$, Eq. (7) takes the form of Eq. (6). Indeed, from

$$\sum_{\ell \neq i} K_{i\ell}^{(\sigma)} = \sum_{\ell \neq i} \sum_{\beta} (B_{\beta\beta} - 1)^{\sigma} e_{i\beta} e_{\ell\beta} = \sum_{\beta} (B_{\beta\beta} - 1)^{\sigma+1} e_{i\beta},$$

where we used $\sum_{\ell \neq i} e_{\ell\beta} = B_{\beta\beta} - 1$, we can conclude

$$\begin{aligned}
T_{ij}^{(\sigma)} &= \frac{1}{\sum_{\beta} (B_{\beta\beta} - 1)^{\sigma+1} e_{i\beta}} \sum_{\alpha} (B_{\alpha\alpha} - 1)^{\sigma} e_{i\alpha} e_{j\alpha} \\
&= \frac{1}{\sum_{\beta} (B_{\beta\beta} - 1)^{\sigma+1} e_{i\beta}} \sum_{\alpha} (B_{\alpha\alpha} - 1)^{\sigma+1} e_{i\alpha} \frac{e_{j\alpha}}{B_{\alpha\alpha} - 1} \\
&= \sum_{\alpha} Q_{\alpha i} R_{j\alpha}.
\end{aligned}$$

For large positive values of σ , hyperedges with many nodes contribute more to $T_{ij}^{(\sigma)}$ and guide the random process. For large negative values of σ , the large hyperedges' contributions are negligible and only the small hyperedges drive the random walk process. In this way, σ is a hyperedge-size bias parameter.

When $\sigma = 0$, the random walk moves on the so-called clique-reduced multigraph, where each pair of nodes is connected by a number of edges equal to the number of hyperedges containing that pair in the hypergraph. The transition matrix takes the form

$$T_{ij}^{(0)} = \frac{K_{ij}^{(0)}}{\sum_{\ell \neq i} K_{i\ell}^{(0)}} = \frac{A_{ij}}{\sum_{\ell \neq i} A_{i\ell}},$$

where we used the hyperadjacency matrix $A_{ij} = \sum_{\alpha} e_{i\alpha} e_{j\alpha}$. The clique reduced multigraph is different from the projected network obtained by associating to each hyperedge a clique of the same size. The latter can be interpreted as the unweighted version of the clique-reduced multigraph.

The matrix $K_{ij}^{(\sigma)}$ given by Eq. (8) can be considered as the weighted adjacency matrix of an undirected network. With the associated Laplace operator

$$L_{ij}^{(\sigma)} = \delta_{ij} - T_{ij}^{(\sigma)} = \delta_{ij} - \frac{K_{ij}^{(\sigma)}}{\sum_{\ell \neq i} K_{i\ell}^{(\sigma)}}, \quad (9)$$

the continuous-time random walk is

$$\dot{p}_i(t) = - \sum_j p_j(t) L_{ji}^{(\sigma)}, \quad (10)$$

and the continuous-time transition matrix, which forwards the continuous-time random-walk process by time t , is

$$\mathbf{T}^{(\sigma)}(t) = e^{-t\mathbf{L}^{(\sigma)}}. \quad (11)$$

Based on this projection, we can invoke standard results about random walks and in particular prove that the stationary state ¹

¹ To lighten the notation, we do not show explicitly the dependence on σ on the asymptotic distribution $\boldsymbol{\pi}$.

$$\pi_j = \frac{d_j^{(\sigma)}}{\sum_{\ell} d_{\ell}^{(\sigma)}}, \quad (12)$$

where $d_j^{(\sigma)} = \sum_{\ell \neq j} K_{j\ell}^{(\sigma)}$ is the strength of node j in the weighted graph. This quantity is an immediate generalisation of the standard node degree that accounts for different hyperedge sizes.

The standard formulation of the map equation uses a discrete-time random-walk process. For discrete time steps k , the discrete-time version of Eq. (3) takes the form

$$p_i(k+1) = \sum_j p_j(k) T_{ji}, \quad (13)$$

with the same stationary distribution as for the continuous-time process.

3 Flow-based community detection in hypergraphs

To illustrate how the flow-based community-detection methods known as Markov stability and the map equation identify communities in hypergraphs, we explain their disparate machinery and illustrate with a schematic hypergraph.

3.1 Markov stability

The Markov stability is a quality function for partitioning a network into communities based on the persistence of random-walk flows inside a group of nodes.

Consider an ergodic random walk process in its stationary state. The Markov stability [31, 18] of a partition at time t is defined as the difference between the probability of a random walker to be in the same community at time 0 and at time t and the analogous quantity computed once the system settles in the stationary state. At a given time t , Markov stability is large when random walkers are unlikely to have escaped their initial community. Because the process is assumed to be ergodic and knowledge of the initial conditions is lost asymptotically, the second probability is equal to the probability of two randomly chosen walkers residing in the same community.

Markov stability and autocovariance of the walker signal

The Markov stability is equivalent to the autocovariance of a signal encoding the sequence of communities visited by a random walker in the stationary state.

Starting from the continuous-time random walk in Eq. (3) or Eq. (10) and with a partition encoded by an $n \times C$ indicator matrix \mathbf{C} , we assign the values X_α ($\alpha = 1, \dots, C$) to the vertices of each community. The autocovariance of the sequence of values $X(t)$ is

$$\text{cov}[X(0)X(t)] = \langle X(0)X(t) \rangle - \langle X(0) \rangle \langle X(t) \rangle = \mathbf{X}^\top \mathbf{R}(t, \mathbf{C}) \mathbf{X}, \quad (14)$$

where $\langle X(t) \rangle$ is the expectation of the random variable $X(t)$, \mathbf{X} is the $1 \times C$ column vector of labels assigned to the C communities and $\mathbf{R}(t, \mathbf{C})$ is the $C \times C$ clustered covariance matrix,

$$\mathbf{R}(t, \mathbf{C}) = \mathbf{C}^\top [\mathbf{\Pi} \exp(-t\mathbf{L}) - \boldsymbol{\pi} \boldsymbol{\pi}^\top] \mathbf{C}, \quad (15)$$

where $\mathbf{\Pi} = \text{diag}(\boldsymbol{\pi})$ is a diagonal matrix encoding the stationary distribution. This $C \times C$ matrix does not depend on our arbitrary choice for the values X_α .

Since $(\mathbf{\Pi} \exp(-t\mathbf{L}))_{ij}$ measures the flow of probability from node i to node j over time t , the first term of Eq. (15) measures the probability flow between two communities α and β over a time t . In a partition with a strong assortative community structure, the probability flows should be captured for long times within communities, and large entries in $\mathbf{R}(t, \mathbf{C})$ should be concentrated on its diagonal. Following this reasoning, the Markov stability $r(t, \mathbf{C})$ is defined as the trace of the clustered autocovariance matrix [18, 17]:

$$r(t, \mathbf{C}) = \text{Tr}[\mathbf{R}(t, \mathbf{C})]. \quad (16)$$

The Markov stability quantifies the quality of a partition at different time scales. For every t , maximising $r(t, \mathbf{C})$ gives the best partition of the network into communities, resulting in a sequence of optimal partitions at different times. Time acts as a resolution parameter that enables us to tune the typical size of the communities in the optimal partition, as longer times typically lead to fewer and larger communities. Markov stability shows connections with several concepts related to community detection. For small t , the first order of its Taylor expansion recovers, up to a constant, a parametric generalisation[38] of Newman-Girvan modularity.

3.1.1 Markov stability for hypergraphs

The random walk process above defined in Eq. (10) corresponds to a random walk on a weighted undirected network. The link weights are self-consistently defined starting from the process on the hypergraph and thus taking into account the multibody

interactions encoded by the size of the hyperedges. Exploiting this relation, we can optimise Markov Stability on the weighted network to partition the hypergraph nodes into communities. Consider again a partition of the nodes of a hypergraph into C non-overlapping communities, encoded by the $n \times C$ indicator matrix \mathbf{C} . Expanding Eq. (16), the corresponding Markov stability is explicitly given by

$$r(t, \mathbf{C}) = \text{Tr} \left[\mathbf{C}^\top \left(\mathbf{\Pi} e^{-t\mathbf{L}^{(\sigma)}} - \boldsymbol{\pi}^\top \boldsymbol{\pi} \right) \mathbf{C} \right], \quad (17)$$

where $\boldsymbol{\pi}$ is the asymptotic solution of Eq. (12), $\mathbf{\Pi}$ is the diagonal matrix containing $\boldsymbol{\pi}$ on the diagonal, $\mathbf{L}^{(\sigma)}$ is the random-walk Laplace matrix and $\boldsymbol{\pi}^\top \boldsymbol{\pi}$ is the matrix whose (i, j) entry is $\pi_i \pi_j$.

As an illustration, we compute the Markov stability for a schematic hypergraph where nodes have two features represented by letters (A , B or C) and numbers (1 or 2) (Fig. 1). The hypergraph has six nodes, $\{A_1, A_2, B_1, B_2, C_1, C_2\}$, and five hyperedges, three of size two connecting nodes with the same letter, and two of size three connecting nodes with the same number. With $\mathbf{K}^{(\sigma)}$ defined in Eq. (8), the hyperedges of size two have a weight of 1^σ to the associated link, whereas the hyperedges of size three have a weight of 2^σ . The corresponding transition probabilities in Eq. (7) are

$$\begin{aligned} T_{X_1, X_2} &= \frac{1}{2 \times 2^\sigma + 1}, \\ T_{X_1, Y_1} &= \frac{2^\sigma}{2 \times 2^\sigma + 1}, \\ T_{X_1, Y_2} &= 0, \\ T_{X_1, X_1} &= 0, \\ T_{X_2, X_2} &= 0 \end{aligned}$$

for all X, Y in $\{A, B, C\}$.

When $\sigma = 0$, the transition probabilities are $T_{X_1, X_2} = T_{X_1, Y_1} = 1/3$. Thus, a walker remains twice as likely in the same 3-hyperedge than in a 2-hyperedge. The Markov stability thus returns two modules for sufficiently large Markov times ($t > 1$). For $\sigma = 1$, $T_{X_1, X_2} = 1/5$ and $T_{X_1, Y_1} = 2/5$. Then, a walker is four times more likely to remain in the same 3-hyperedge than in a 2-hyperedge. For $\sigma = -1$, $T_{X_1, X_2} = 1/2$ and $T_{X_1, Y_1} = 1/4$, and the probability to stay in the 3-hyperedge is the same as leaving it. In the limit as $\sigma \rightarrow \infty$,

$$\lim_{\sigma \rightarrow \infty} T_{X_1, X_2} = 0, \quad \lim_{\sigma \rightarrow \infty} T_{X_1, Y_1} = 1/2 \quad \forall X, Y \in \{A, B, C\},$$

hops among nodes with the same number are strongly favoured and a walker remains for a long time in the same 3-hyperedge. In the other limit,

$$\lim_{\sigma \rightarrow -\infty} T_{X_1, X_2} = 1, \quad \lim_{\sigma \rightarrow -\infty} T_{X_1, Y_1} = 0 \quad \forall X, Y \in \{A, B, C\}$$

a walker spends longer time in the 2-hyperedges. In the first case, Markov stability favours two communities for sufficiently large Markov times, while it favours three communities in the second case (Fig. 2).

3.2 The Map Equation

Like Markov stability, the map equation is a quality function for partitioning a network into communities based on the persistence of random-walk flows inside groups of nodes. But instead of using autocovariance, the map equations casts the problem of finding flow-based communities in networks into a minimum-description-length problem [40].

The map equation measures, in bits, the optimal codelength L per step of a discrete random walk on a network for a given node partition M with C modules. When all nodes are in the same module, the map equation is simply the Shannon entropy H of the node-visit rates $\mathcal{P} = \{\pi_i\}$.

In partitions with more than one module, the map equation combines within and between module codelengths for describing flows within and between modules. For modules $\alpha = 1, \dots, C$ with

$$\begin{aligned} \text{entry flow rates } q_{\alpha\curvearrowright} &= \sum_{i \notin \alpha, j \in \alpha} \pi_i T_{ij}, \\ \text{exit flow rates } q_{\alpha\curvearrowleft} &= \sum_{i \in \alpha, j \notin \alpha} \pi_i T_{ij}, \\ \text{entry flow rate random variable } \mathcal{Q} &= \{q_{\alpha\curvearrowright}\} \\ &\text{with total flow rate } q_{\curvearrowright} = \sum_{\alpha} q_{\alpha\curvearrowright}, \\ \text{exit and node-visit rate random variables } \mathcal{P}_{\alpha} &= \{q_{\alpha\curvearrowright}, \pi_{i \in \alpha}\} \\ &\text{with total flow rate } p_{\alpha\cup} = q_{\alpha\curvearrowright} + \sum_{i \in \alpha} \pi_i, \end{aligned}$$

the map equation takes its general two-level form

$$L(M) = q_{\curvearrowright} H(\mathcal{Q}) + \sum_{\alpha} p_{\alpha\cup} H(\mathcal{P}_{\alpha}). \quad (18)$$

The first term is the codelength for between-module movements, followed by the sum of codelengths for within-module movements over all modules.

This standard formulation of the map equation encodes every transition of a random walker. For Markov times other than 1, the map equation uses a linearised continuous-time transition matrix,

$$\tilde{T}(t) = \begin{cases} (1-t)I + tT & t < 1 \\ tT & t \geq 1, \end{cases} \quad (19)$$

which captures Markov times below 1 with self-links and Markov times above 1 with transition rates proportional to the average rate of the underlying Poisson process [29]. At Markov time 1, it recovers the discrete-time process in Eq. (13). Unlike the exponential of the Laplacian, the linearisation keeps the transition matrix sparse also for Markov times larger than zero.²

The linearisation is appealing for the map equation because the node-visit rates π_i remain the same for all Markov times t – since the relative visit rates at steady state do not depend on how often the visits are sampled – and the module exit and entry rates change linearly – since the number of random walkers moving along any link between nodes during time t is directly proportional to t . Therefore, we can define

$$q_{\alpha \curvearrowright}(t) = tq_{\alpha \curvearrowright} \quad (20)$$

$$q_{\alpha \curvearrowleft}(t) = tq_{\alpha \curvearrowleft}. \quad (21)$$

With time dependence on all variables in Eq. (18) that depend on the module entry or exit flow rates, the map equation for Markov time t takes the form

$$L(M, t) = q_{\curvearrowright}(t)H(Q(t)) + \sum_{\alpha} p_{\alpha \cup}(t)H(\mathcal{P}_{\alpha}(t)). \quad (22)$$

While the standard formulation of the map equation can be applied directly to the continuous-time transition matrix for various Markov times in Eq. (11), deriving the matrix is computationally expensive for large networks. For long Markov times, it also generates dense networks with longer clustering times. In contrast, the map equation for Markov time t in Eq. (22) has no overhead compared with the standard map equation. From a flow modelling perspective, the map equation for Markov time t counts each time a random walker moves across a modular boundary during time t . Instead, the standard map equation applied to the continuous-time transition matrix counts at most one boundary-crossing since it only considers the final position after time t . For short Markov times, the two approaches are similar. But for longer Markov times, the map equation for Markov time t generally favours solutions with fewer modules [29].

To identify the optimal partition – the one that compresses the modular description the most – we use the community-detection algorithm Infomap [20] available on www.mapequation.org. Infomap is to the map equation what the Louvain [4] and the Leiden [48] methods are to the objective function modularity [36], which favours partitions with a high internal density of links compared with a statistical null model. Infomap uses a similar search algorithm as the Leiden method but tries to find the node assignment that minimises the map equation’s codelength. Infomap also finds deeper hierarchical partitions – from top-level supermodules with multiple levels of submodules down to leaf-level modules containing the nodes – if such multilevel solutions give higher modular compression [41]. The search algorithm works on many network types, including weighted, directed, bipartite [3], and multilayer [13]

² Markov stability can also apply the linearisation to speed up calculations [31], but here we use the exponential of the Laplacian for Markov stability in all examples.

networks. Infomap can also identify overlapping multilevel communities in higher-order network flows modelled with so-called memory networks [20]. Recent work introduces a Bayesian estimate of the map equation for identifying flows on sparse networks with missing links [47], but here we focus on the standard map equation and identify hard two-level partitions in weighted unipartite networks obtained with the hypergraph projection described in section 2.

3.2.1 The map equation for hypergraphs

We have previously derived unipartite, bipartite, and multilayer network representations of hypergraph flows and analysed their different advantages when mapping flows on hypergraphs with Infomap [22]. We found that when the research question does not require hyperedge assignments from using bipartite networks or overlapping modules from using multilayer networks, a unipartite network representation like the one we use here provides a good trade-off between speed and ability to reveal modular regularities.

We illustrate the map equation for hypergraphs with the schematic example in Fig. 1 for which the stationary distribution of the random walk is independent of σ such that $\pi_i^{(\sigma)} = \pi_i$ for all nodes i . The one-module codelength is

$$\begin{aligned} L(M_1, t) &= H(\mathcal{P}) \\ &= H(\pi_{A_1}, \pi_{A_2}, \pi_{B_1}, \pi_{B_2}, \pi_{C_1}, \pi_{C_2}) \\ &= 2.58 \text{ bits for all values of } \sigma \text{ and } t. \end{aligned} \tag{23}$$

When a network has modular regularities, a partition captures the modular flows when the random walker spends long times within the modules with few transitions between them. The codelength is shorter than in the one-module solution because the information required to specify a random walker's position in a module decreases with its size. But for partitions with too many modules, the information required for describing between-module movements exceeds the gain from using small modules. The optimal partition has the shortest codelength. Its node assignment best captures the modular regularities of flows on the network.

Using the three-module solution in Fig. 1(b), the codelengths for the different hyperedge-size biased random walks parametrised with σ are

$$\begin{aligned}
L(M_3, t) &= q_{\curvearrowright} H(q_{1\curvearrowright}, q_{2\curvearrowright}, q_{3\curvearrowright}) \\
&\quad + (q_{1\curvearrowright} + \pi_{A_1} + \pi_{A_2}) H(q_{1\curvearrowright}, \pi_{A_1}, \pi_{A_2}) \\
&\quad + (q_{2\curvearrowright} + \pi_{B_1} + \pi_{B_2}) H(q_{2\curvearrowright}, \pi_{B_1}, \pi_{B_2}) \\
&\quad + (q_{3\curvearrowright} + \pi_{C_1} + \pi_{C_2}) H(q_{3\curvearrowright}, \pi_{C_1}, \pi_{C_2}) \\
&= \begin{cases} \mathbf{1.00} & \mathbf{1.00} & \mathbf{1.00} & \text{bits for } \sigma \rightarrow -\infty, \\ \mathbf{2.30} & 3.17 & 4.58 & \text{bits for } \sigma = -1, \\ 2.61 & 3.67 & 5.41 & \text{bits for } \sigma = 0, \\ 2.84 & 4.05 & 6.04 & \text{bits for } \sigma = 1, \text{ and} \\ \underbrace{3.17}_{t=0.5} & \underbrace{4.58}_{t=1} & \underbrace{6.92}_{t=2} & \text{bits for } \sigma \rightarrow \infty. \end{cases}
\end{aligned} \tag{24}$$

This three-module letter-based solution gives shorter codelengths than the one-module solution for small σ -values and short Markov times, indicated by bold numbers in Eq. (24). For these values, the random walker spends long times within the three letter-based modules with few transitions between.

Using the two-module solution in Fig. 1(d), the codelengths for the different hyperedge-size biased random walks parametrised with σ are

$$\begin{aligned}
L(M_2, t) &= q_{\curvearrowright} H(q_{1\curvearrowright}, q_{2\curvearrowright}) \\
&\quad + (q_{1\curvearrowright} + \pi_{A_1} + \pi_{B_1} + \pi_{C_1}) H(q_{1\curvearrowright}, \pi_{A_1}, \pi_{B_1}, \pi_{C_1}) \\
&\quad + (q_{2\curvearrowright} + \pi_{A_2} + \pi_{B_2} + \pi_{C_2}) H(q_{2\curvearrowright}, \pi_{A_2}, \pi_{B_2}, \pi_{C_2}) \\
&= \begin{cases} 3.46 & 4.58 & 6.34 & \text{bits for } \sigma \rightarrow -\infty, \\ 2.74 & 3.46 & 4.58 & \text{bits for } \sigma = -1, \\ \mathbf{2.44} & 3.00 & 3.87 & \text{bits for } \sigma = 0, \\ \mathbf{2.17} & \mathbf{2.56} & 3.19 & \text{bits for } \sigma = 1, \text{ and} \\ \underbrace{\mathbf{1.58}}_{t=0.5} & \underbrace{\mathbf{1.58}}_{t=1} & \underbrace{\mathbf{1.58}}_{t=2} & \text{bits for } \sigma \rightarrow \infty. \end{cases}
\end{aligned} \tag{25}$$

This two-module number-based solution gives shorter codelengths than the one-module solution for large σ -values and short Markov times, again indicated by bold numbers in Eq. (25). For these values, the random walker instead spends long times within the two number-based modules so the codelength savings from using smaller modules exceed the extra cost from switching modules.

3.3 Comparing Markov stability with the map equation

While both Markov stability and the map equation identify flow-based communities in hypergraphs, they highlight different structures. Both methods prefer the solution with one node in each community for minimal Markov times when nearly all flow remains in the nodes. But the Markov time at which they highlight non-trivial solutions differs. For Markov stability this happens when more flows stay in larger

communities than in singleton communities compared with respective stationary expectation. For the map equation, it happens when describing transitions between communities costs more than gained from describing within-community movements in singleton communities. Which one comes first depends on the network structure. For the schematic network in Fig. 1, the map equation first prefers non-trivial solutions (Fig. 2). But in the larger networks analysed in section 4, Markov stability highlights non-trivial solutions for shorter Markov times than the map equation.

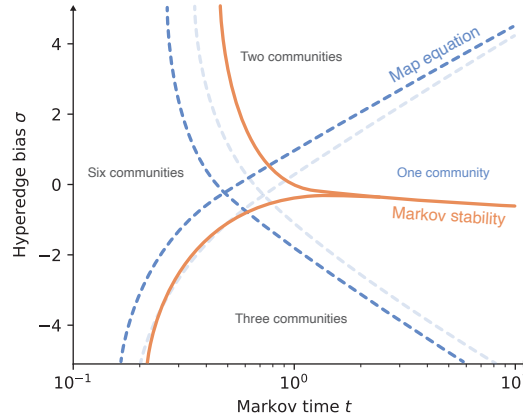


Fig. 2 Optimal solutions of the schematic hypergraph in Fig. 1 for different hyperedge biases and Markov times. Solutions from Fig. 1(b-d) and the singleton solution with six communities. Light dashed line for the map equation applied to a continuous-time process.

The map equation strictly prefers the one-community solution over a non-trivial solution for long Markov times. It prefers a two-community solution only when describing transitions between the communities costs no more than gained from describing within-community movements in smaller communities. The map equation’s modular code structure with costs in bits for both community exits and entries has a regularising effect as long as the network is connected. In contrast, Markov stability prefers any solution that traps the slightest more flows within the communities for Markov time t than expected after infinite time. For example, only the map equation prefers the one-community solution for the longer Markov times included in Fig. 2. As a result, Markov stability typically has a longer range of non-trivial solution and responds slower to changes in Markov time.

For the schematic hypergraph, both methods prefer two communities with flows biased to larger hyperedges and three communities when biased to smaller hyperedges (Fig. 2). The map equation’s continuous-time approximation slightly shifts the optimal solutions to longer Markov times in this case.

4 Experiments

To illustrate how Markov stability and the map equation differ applied to real hypergraphs, we analysed a zoo hypergraph, a collaboration hypergraph, and a hypergraph from fossil records.

4.1 A zoo hypergraph

The zoo hypergraph has been built using data from the UCI Machine Learning Depository [19]; it consists of 101 animals nodes and 20 features hyperedges such as fur, having a tail, and the number of legs. We identified communities for random-walk bias $\sigma = -4, -2, \text{ and } 2$ as a function of Markov time for Markov stability using public code [15] and the map equation using Infomap [21].

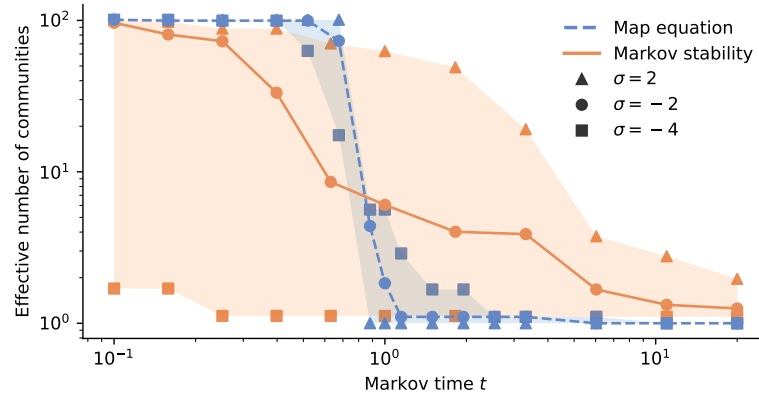


Fig. 3 Effective number of communities in the zoo hypergraph for different hyperedge-size bias σ and Markov time t . Filled areas between the highest and lowest number of effective communities for Markov stability (orange) and the map equation (blue).

The effective number of communities, measured as the perplexity of the relative community sizes \mathcal{S} , $2^{H(\mathcal{S})}$, decreases gradually for Markov stability and more abruptly for the map equation (Fig. 3). Already at Markov time 0.1, Markov stability applied to the zoo hypergraph with hyperedge-size bias $\sigma = -4$ identifies less than two effective communities. The bias toward small hyperedges fragments the network into weakly connected components. The penalty term $-\pi^\top \pi$ in Eq. (17) with global flows makes Markov stability sense this large-scale network structure even if most flows – more than 90 per cent in this case – remain at the node for the short Markov time.

In contrast, the map equation prefers the singleton solution for the three tested hyperedge-size biases for Markov time up to 0.5. A short Markov time makes encoding each node in its singleton community cheap since most flows remain there in each step. The extra cost from describing movements in larger communities is higher than the gain from describing fewer transitions between communities until the Markov time approaches 1.

Markov stability identifies the largest communities with flows biased to small hyperedges for all Markov times. The map equation applied to a continuous-time process shows similar behaviour. But for the approximation shown in Fig. 3, it favours the one-community solution with flows biased to large hyperedges for Markov time about 0.9 or higher. Encoding all nodes in the same community best compresses the interconnected flows biased to large hyperedges.

The various community structures at Markov time 1 illustrates that Markov stability and the map equation work differently. They identify different communities for similar flows and similar communities for different flows. Between $\sigma = -2$ and -4 , Markov stability and the map equation find closest solutions for differing σ . Markov stability for $\sigma = -2$ and the map equation for $\sigma = -4$ find communities that best capture animal classes (Fig. 4).




Fig. 4 Alluvial diagram of the zoo hypergraph solutions for Markov time $t = 1$ and different hyperedge biases σ . Each block represents a community with block height proportional to the number of nodes in the community. Stream fields connect blocks with shared nodes. Colours indicate animal classes.

4.2 A collaboration hypergraph

We used a collaboration hypergraph compiled from the 734 references in the review paper *Networks beyond pairwise interactions: Structure and dynamics* [2, 22]. Hyperedges represent referenced articles and nodes represent their authors. Authors with multiple articles form connections between the hyperedges. We analysed the largest connected component with 361 author nodes in 220 hyperedges.

Again, changes in Markov time have a more gradual effect on Markov stability than on the map equation. But for the collaboration network, the map equation transitions from the singleton to the all-in-one solution over three orders of magnitude and identifies non-trivial solutions for Markov times over 100 (Fig. 5). Different Markov times highlight different hierarchical levels in the hypergraph’s hierarchical community structure with smaller communities nested in larger communities. The largest communities remain clear despite long walks that multiply flows tenfold between the communities.




Fig. 5 Effective number of communities in the collaboration hypergraph for different hyperedge-size bias σ and Markov time t . Filled areas between the highest and lowest number of effective communities for Markov stability (orange) and the map equation (blue).

4.3 A fossil-record hypergraph

We analysed a hypergraph representation of marine fossil animals from Cambrian (541 MY) to Cretaceous (66 MY) [39]. Geological stages in the underlying sample-based occurrence data form hyperedges connecting all genera occurring at each stage. Genera occurring in multiple geological stages connect hyperedges. We weighted the hyperedges by dividing the number of samples where a genus occurs in a given

geological stage by the total number of samples recorded at the stage. The assembled hypergraph comprises 77 geological stage hyperedges and 13,276 fossil genera nodes [22].




Fig. 6 Effective number of communities in the fossil hypergraph for different hyperedge-size bias σ and Markov time t . Filled areas between the highest and lowest number of effective communities for Markov stability (orange) and the map equation (blue).

Once more, the effective number of identified communities decreases gradually for the Markov stability without reaching the all-in-one solution even at high Markov times (Fig. 6). In contrast, the map equation solutions quickly transition from single- to all-in-one communities around Markov time 1. In the narrow range of Markov times with non-trivial solutions, the effective number of communities is two to five. For this substantially larger hypergraph, the extensive hyperedges form a weak community structure that dissolves for long Markov times and prevents modular compression.

Non-trivial Infomap solutions reproduce the underlying temporal structure of the paleontological data with faunas organised into units of geological time (Fig. 7). Although fossil genera can occur in more than one of these large-scale temporal units, Infomap identifies successive global faunas that replace each other at their boundaries. Overall, faunas from geological periods are clustered together or combined into coarser temporal units at Markov time 1. These salient structures appear from singletons without intermediate smaller structures for shorter Markov times.

Markov stability also delineates temporal faunas limited by geological units at Markov time around 1. However, the temporal structure of the data cannot explain the numerous communities Markov stability identifies for short Markov times (Fig. 8).




Fig. 7 Alluvial diagram of the map equation’s fossil hypergraph solutions for hyperedge-size bias $\sigma = 0$. Each block represents a community with block height proportional to the number of nodes in the community. Stream fields connect blocks with shared nodes. Singleton communities aggregated in orange blocks for clarity.




Fig. 8 Alluvial diagram of Markov stability’s fossil hypergraph solutions for hyperedge-size bias $\sigma = 0$. Each block represents a community with block height proportional to the number of nodes in the community. Stream fields connect blocks with shared nodes. Markov times chosen to match the effective number of communities in Fig. 7 and small communities with fewer than five nodes aggregated in an orange block for clarity.

5 Conclusions

We have derived Markov stability and the map equation for a random-walk process on hypergraphs with hyperedge-size bias. Both methods identify communities where flows stay for long times. Still, with disparate machinery – Markov stability measures overrepresentation of random walkers in communities where they started, whereas the map equation measures the modular description length in bits – they capitalise on distinct flow-based structures. By comparing with the stationary expectation, Markov stability is more sensitive to hyperedge-size biases and gradually finds larger communities for longer Markov times. When increasing the Markov time, the map equation instead transitions more abruptly from identifying many small to few large communities. Compared with the influence from their disparate machineries, the map equation’s approximation of a continuous-time process has a negligible effect. Whereas the map equation identifies salient structures, Markov stability can identify communities of any size irrespective of the large-scale hypergraph structure. The question and hypergraph at hand decide which method identifies the flow-based communities that best elucidate the hypergraph’s studied function.

Acknowledgements A.E was supported by the Swedish Foundation for Strategic Research, Grant No. SB16-0089. A.R. and M.R. were supported by the Swedish Research Council, Grant No. 2016-00796

References

1. Angelini, M.C., Caltagirone, F., Krzakala, F., Zdeborová, L.: Spectral detection on sparse hypergraphs. In: 2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton), pp. 66–73. IEEE (2015)
2. Battiston, F., Cencetti, G., Iacopini, I., Latora, V., Lucas, M., Patania, A., Young, J.G., Petri, G.: Networks beyond pairwise interactions: structure and dynamics. *Phys. Rep.* (2020)
3. Blöcker, C., Rosvall, M.: Mapping flows on bipartite networks. *Phys. Rev. E* **102**(5), 052305 (2020)
4. Blondel, V.D., Guillaume, J.L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. *J. Stat. Mech.* **2008**(10), P10008 (2008)
5. Boccaletti, S., Latora, V., Moreno, Y., Chavez, M., Hwang, D.U.: Complex networks: Structure and dynamics. *Phys. Rep.* **424**(4-5), 175–308 (2006)
6. Brin, S., Page, L.: The anatomy of a large-scale hypertextual web search engine. *Comput. Netw.* **30**(1-7), 107–117 (1998)
7. Carletti, T., Battiston, F., Cencetti, G., Fanelli, D.: Random walks on hypergraphs. *Phys. Rev. E* **101**(2), 022308 (2020)
8. Carletti, T., Fanelli, D., Lambiotte, R.: Random walks and community detection in hypergraphs. *Journal of Physics: Complexity* **2**(1), 015011 (2021)
9. Chien, I., Lin, C.Y., Wang, I.H.: Community detection in hypergraphs: Optimal statistical limit and efficient algorithms. In: International Conference on Artificial Intelligence and Statistics, pp. 871–879. PMLR (2018)
10. Chitra, U., Raphael, B.J.: Random walks on hypergraphs with edge-dependent vertex weights. In: 36th International Conference on Machine Learning, ICML 2019, pp. 2002–2011. International Machine Learning Society (IMLS) (2019)

11. Chodrow, P.S., Veldt, N., Benson, A.R.: Hypergraph clustering: from blockmodels to modularity. arXiv:2101.09611 (2021)
12. De Domenico, M., Granell, C., Porter, M.A., Arenas, A.: The physics of spreading processes in multilayer networks. *Nat. Phys.* **12**(10), 901–906 (2016)
13. De Domenico, M., Lancichinetti, A., Arenas, A., Rosvall, M.: Identifying modular flows on multilayer networks reveals highly overlapping organization in interconnected systems. *Phys. Rev. X* **5**(1), 011027 (2015)
14. De Domenico, M., Solé-Ribalta, A., Cozzo, E., Kivelä, M., Moreno, Y., Porter, M.A., Gómez, S., Arenas, A.: Mathematical formulation of multilayer networks. *Physical Review X* **3**(4), 041022 (2013)
15. Delmotte, A., Schaub, M.: Community detection using the stability of a graph partition (2012). URL http://wwwf.imperial.ac.uk/~mpbara/Partition_Stability/
16. Delvenne, J., Yaliraki, S., Barahona, M.: Stability of graph communities across time scales. *Proc. Natl. Acad. Sci. U.S.A* **107**(29), 12755–12760 (2010)
17. Delvenne, J.C., Schaub, M., Yaliraki, S.N., Barahona, M.: The stability of a graph partition: A dynamics-based framework for community detection. In: *Dynamics On and Of Complex Networks (II)*, p. 221 (2013)
18. Delvenne, J.C., Yaliraki, S.N., Barahona, M.: Stability of graph communities across time scales. *Proceedings of the National Academy of Sciences* **107**, 12755 (2010)
19. Dua, D., Graff, C.: UCI machine learning repository (2017). URL <https://archive.ics.uci.edu/ml/datasets/Zoo>
20. Edler, D., Bohlin, L., et al.: Mapping higher-order network flows in memory and multilayer networks with infomap. *Algorithms* **10**(4), 112 (2017)
21. Edler, D., Eriksson, A., Rosvall, M.: The Infomap Software Package (2021). URL <https://www.mapequation.org>
22. Eriksson, A., Edler, D., Rojas, A., Rosvall, M.: Mapping flows on hypergraphs: How choosing random-walk model and network representation matters for community detection. arXiv:2101.00656 (2021)
23. Fortunato, S.: Community detection in graphs. *Phys. Rep.* **486**(3-5), 75–174 (2010)
24. Guimera, R., Amaral, L.A.N.: Functional cartography of complex metabolic networks. *Nature* **433**(7028), 895–900 (2005)
25. Iacopini, I., Petri, G., Barrat, A., Latora, V.: Simplicial models of social contagion. *Nat. Commun.* **10**(1), 1–9 (2019)
26. Jeub, L.G., Mahoney, M.W., Mucha, P.J., Porter, M.A., et al.: A local perspective on community structure in multilayer networks. *Netw. Sci.* **5**(2), 144–163 (2017)
27. Kamiński, B., Poulin, V., Prałat, P., Szufel, P., Théberge, F.: Clustering via hypergraph modularity. *PloS one* **14**(11), e0224307 (2019)
28. Ke, Z.T., Shi, F., Xia, D.: Community detection for hypergraph networks via regularized tensor power iteration. arXiv:1909.06503 (2019)
29. Kheirkhahzadeh, M., Lancichinetti, A., Rosvall, M.: Efficient community detection of network flows for varying markov times and bipartite networks. *Phys. Rev. E* **93**(3), 032309 (2016)
30. Kivelä, M., Arenas, A., Barthelemy, M., Gleeson, J.P., Moreno, Y., Porter, M.A.: Multilayer networks. *J. Complex Netw.* **2**(3), 203–271 (2014)
31. Lambiotte, R., Delvenne, J.C., Barahona, M.: Random walks, markov processes and the multi-scale modular organization of complex networks. *IEEE Transactions on Network Science and Engineering* **1**, 76 (2014)
32. Lambiotte, R., Rosvall, M., Scholtes, I.: From networks to optimal higher-order models of complex systems. *Nat. Phys.* **15**(4), 313–320 (2019)
33. Li, P., Milenkovic, O.: Inhomogeneous hypergraph clustering with applications. In: I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (eds.) *Advances in Neural Information Processing Systems*, vol. 30. Curran Associates, Inc. (2017). URL <https://proceedings.neurips.cc/paper/2017/file/a50abba8132a77191791390c3eb19fe7-Paper.pdf>

34. Mangioni, G., Jurman, G., De Domenico, M.: Multilayer flows in molecular networks identify biological modules in the human proteome. *IEEE Transactions on Network Science and Engineering* **7**(1), 411–420 (2018)
35. Mucha, P.J., Richardson, T., Macon, K., Porter, M.A., Onnela, J.P.: Community structure in time-dependent, multiscale, and multiplex networks. *Science* **328**(5980), 876–878 (2010)
36. Newman, M.E., Girvan, M.: Finding and evaluating community structure in networks. *Phys. Rev. E* **69**(2), 026113 (2004)
37. Parzanchevski, O., Rosenthal, R.: Simplicial complexes: spectrum, homology and random walks. *Random Struct. Algorithms* **50**(2), 225–261 (2017)
38. Reichardt, J., Bornholdt, S.: Statistical mechanics of community detection. *Physical review E* **74**(1), 016110 (2006)
39. Rojas, A., Calatayud, J., Kowalewski, M., Neuman, M., Rosvall, M.: A multiscale view of the Phanerozoic fossil record reveals the three major biotic transitions. *Communications Biology* **4**(1), 309 (2021). DOI 10.1038/s42003-021-01805-y. URL <http://www.nature.com/articles/s42003-021-01805-y>
40. Rosvall, M., Bergstrom, C.T.: Maps of random walks on complex networks reveal community structure. *Proc. Natl. Acad. Sci. U.S.A.* **105**(4), 1118–1123 (2008)
41. Rosvall, M., Bergstrom, C.T.: Multilevel compression of random walks on networks reveals hierarchical organization in large integrated systems. *PloS one* **6**(4), e18209 (2011)
42. Rosvall, M., Esquivel, A.V., Lancichinetti, A., West, J.D., Lambiotte, R.: Memory in network flows and its effects on spreading dynamics and community detection. *Nat. Commun.* **5**(1), 1–13 (2014)
43. Salmikov, V., Cassese, D., Lambiotte, R.: Simplicial complexes and complex systems. *Eur. J. Phys.* **40**(1), 014001 (2018)
44. Schaub, M.T., Benson, A.R., Horn, P., Lippner, G., Jadbabaie, A.: Random walks on simplicial complexes and the normalized hodge 1-laplacian. *SIAM Rev. Soc. Ind. Appl. Math* **62**(2), 353–391 (2020)
45. Scholtes, I., Wider, N., Pfitzner, R., Garas, A., Tessone, C.J., Schweitzer, F.: Causality-driven slow-down and speed-up of diffusion in non-markovian temporal networks. *Nat. Commun.* **5**(1), 1–9 (2014)
46. Simonsen, I., Eriksen, K.A., Maslov, S., Sneppen, K.: Diffusion on complex networks: a way to probe their large-scale topological structures. *Physica A* **336**(1-2), 163–173 (2004)
47. Smiljanić, J., Edler, D., Rosvall, M.: Mapping flows on sparse networks with missing links. *Phys. Rev. E* **102**(1), 012302 (2020)
48. Traag, V.A., Waltman, L., Van Eck, N.J.: From louvain to leiden: guaranteeing well-connected communities. *Sci. Rep.* **9**(1), 1–12 (2019)
49. Xu, J., Wickramaratne, T.L., Chawla, N.V.: Representing higher-order dependencies in networks. *Science Adv.* **2**(5), e1600028 (2016)
50. Zhou, D., Huang, J., Schölkopf, B.: Learning with hypergraphs: Clustering, classification, and embedding. In: *Advances in neural information processing systems*, pp. 1601–1608 (2007)