

Subgraph Covers: An Information-Theoretic Approach to Motif Analysis in Networks

Anatol E. Wegner*

Max Planck Institute for Mathematics in the Sciences, Inselstrasse 22, 04103 Leipzig, Germany
(Received 20 June 2014; published 10 November 2014)

Many real-world networks contain a statistically surprising number of certain subgraphs, called network motifs. In the prevalent approach to motif analysis, network motifs are detected by comparing subgraph frequencies in the original network with a statistical null model. In this paper, we propose an alternative approach to motif analysis where network motifs are defined to be connectivity patterns that occur in a subgraph cover that represents the network using minimal total information. A subgraph cover is defined to be a set of subgraphs such that every edge of the graph is contained in at least one of the subgraphs in the cover. Some recently introduced random graph models that can incorporate significant densities of motifs have natural formulations in terms of subgraph covers, and the presented approach can be used to match networks with such models. To prove the practical value of our approach, we also present a heuristic for the resulting NP hard optimization problem and give results for several real-world networks.

DOI: [10.1103/PhysRevX.4.041026](https://doi.org/10.1103/PhysRevX.4.041026)

Subject Areas: Complex Systems,
Interdisciplinary Physics

I. INTRODUCTION

Many complex systems can be modeled as networks where vertices represent interacting elements and edges interactions between them. A large number of real-world networks has been found to contain a statistically surprising number of certain small connectivity patterns called network motifs [1]. Network motifs, which are also commonly referred to as basic building blocks of networks, are thought to play an important role in the structural and functional organization of complex networks. For instance, in biological and technological networks, motifs are thought to contribute to the overall functioning of networks by performing modular tasks such as information processing [2]. Hence, methods for identifying such characteristic connectivity patterns are of great importance for a better understanding of complex networks.

The prevalent approach to motif analysis is due to Milo *et al.* [1] and is based on comparing the subgraph frequencies of the original network with a statistical null model that preserves some features of the original network. Part of the analysis consists of generating a representative sample of the null model which is used to determine empirical values for the mean and variance-of-motif counts in the null model. Motifs for which the frequencies significantly deviate from the null model are then classified as network motifs. In their original paper, Milo *et al.*

suggest that when detecting motifs of size n , the null model should conserve the degree distribution of the original network as well as the motif counts of size $n - 1$. For generating such networks, they propose a simulated annealing approach. However, it is not clear whether the simulated annealing approach uniformly samples such null models. Moreover, in most applications, it is computationally not feasible to preserve lower-order motif counts for motifs larger than four vertices. Consequently, in most practical applications [1,3,4], lower-order motif counts are not conserved and the configuration model [5] with the same degree distribution as the original network is used as a null model. Not conserving lower order motif counts has the unwanted consequence that most subgraphs that contain an overrepresented submotif are classified as network motifs.

In this paper, we introduce an alternative approach to motif analysis that is based on using subgraph covers as representations of graphs. A subgraph cover can be seen as a decomposition of the network into smaller building blocks. Given any network, there are many ways of decomposing it into a subgraph cover. Consequently, one needs a way of comparing subgraph covers. In this article we follow the total information approach by Gell-Mann and Lloyd [6] and look at motifs as regularities of a network which can be used to obtain a more concise representation of the network. In our approach, network motifs are defined as subgraph patterns that appear in a subgraph cover that represents the network using minimal total information. Note that this definition of network motifs is fundamentally different from the definition of Milo *et al.* [1].

Another aim of this paper is to establish a connection between motif analysis and random graph models. In contrast to most real-world networks, commonly used

*wegner@mis.mpg.de

Published by the American Physical Society under the terms of the *Creative Commons Attribution 3.0 License*. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

network models are locally treelike. Developing random graph models that can incorporate high densities of triangles and other motifs has been a long-standing problem. Recently, two random graph models that can incorporate significant densities of motifs have been proposed [7,8]. However, it remains unclear how one should select the set of motifs to be used in such models, given a specific network. As we shall see later, these models can be formulated as ensembles of subgraph covers and total information-optimal subgraph covers can be used to match networks with specific instances of these models.

The article is organized as follows: In Sec. II, we present the theory underlying our approach. Then, in Sec. III, we examine the resulting optimization problem and propose a heuristic for it. In Sec. IV, we present empirical results for several real-world networks and also test the heuristic on some synthetic networks with a predefined motif structure. Finally, in Sec. V, we summarize our results and discuss directions for future research.

II. THEORY

In this section, we first introduce necessary graph and information-theoretical concepts. We then define the total information for subgraph covers, and following the approach by Gell-Mann and Lloyd [6], we use the smallness of the total information as a criterion for selecting a subgraph cover that is an optimal representation of a given network. Finally, we discuss the relation between total information-optimal subgraph covers and model selection for random graphs.

A. Subgraph covers

A graph $G = (V(G), E(G))$ is a ordered pair of sets where $V(G)$ [$|V(G)| = N$] is the set of points called vertices and $E(G)$ is a set of links called edges that connect pairs of vertices. Depending on the kind of network, edges might be directed or undirected, although, in this article, we will not make an explicit distinction between directed and undirected graphs since the arguments and definitions apply to both equally well. In general, we will assume that G is sparse, i.e., $|E(G)| = O(N)$. Most real-world networks are sparse [9].

In graph theory, motifs correspond to isomorphism classes. Two graphs G and H are said to be isomorphic ($G \simeq H$) whenever there exists a bijection $\phi: V(G) \rightarrow V(H)$ such that $(x, y) \in E(G) \Leftrightarrow (\phi(x), \phi(y)) \in E(H)$ for all $x, y \in V(G)$. Such a map ϕ is called an isomorphism. Being isomorphic is an equivalence relation, and the corresponding equivalence classes are called isomorphism classes.

A graph H is called a subgraph of G whenever $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$. A set of subgraphs C is said to be a subgraph cover of G whenever $\bigcup_{H \in C} E(H) = E(G)$. Subgraph covers are representations of graphs, meaning

that given a subgraph cover, the corresponding graph can be recovered fully from the cover. Trivial examples of subgraph covers are the set of all edges of G and G itself. Other examples are the maximal clique and star covers, which are the sets of all cliques or stars that are not subcliques or substars. (An n clique consists of n vertices that are all mutually connected, an n star consists of a single central vertex that is connected to n peripheral vertices, and a subclique or substar is a clique or star that is a subgraph of some larger clique or star.) While the maximal star cover is closely related to the adjacency-list representation of the graph, clique covers are closely related to bipartite representations [5,10].

Given a cover C , its motif set $M(C)$ is the set of the isomorphism classes of the subgraphs in C . Similarly, given a set of isomorphism classes M , an M cover C_M is a subgraph cover of which every element belongs to some class in M .

B. The total information approach

The total information framework [6] is based on the idea that, given an entity, one can use Shannon information or entropy to describe its random or nonregular aspects and algorithmic information content to describe its regularities or rule-based features. In this approach, identifying certain regularities of an entity is equivalent to embedding it into an ensemble of objects that share these regularities while they differ in other aspects.

The first information measure of interest for the total information approach is the entropy, also known as the Shannon information. For an ensemble $E(R, p_r)$, where R is the set of possible outcomes and p_r is the respective probability of an element $r \in R$, the entropy measures the uncertainty regarding the outcome of E and is given by

$$S = -K \sum_r p_r \log p_r, \quad (1)$$

where K is a constant. When $K = 1$ and the logarithm is base 2, the entropy is measured in bits.

Another information measure that is needed in order to define the total information is the effective complexity. The effective complexity $\epsilon(E)$ of an entity that is embedded into an ensemble E as a typical member is given by the algorithmic information content (AIC) of the ensemble. The algorithmic information content of such an ensemble E with respect to a universal computer U is the length of the shortest program that instructs U to output a description of E and then halt [i.e., $\epsilon(E) = \text{AIC}_U(E)$]. In general, the effective complexity is not computable and computer dependent; therefore, in practice, one is restricted to work with approximations in the form of upper bounds. The issue of how to define a practical effective complexity for subgraph covers is dealt with in Sec. III.

The sum of the effective complexity and the entropy is the total information required to describe both the random features and regularities of an entity using a certain model:

$$\Sigma(E) = \epsilon(E) + S(E). \quad (2)$$

For a given entity, there might be a multitude of ensembles into which the entity can be embedded as a typical member and it may not always be clear which set of regularities or model provides the best description of the entity. The total information provides a basis for comparing models that describe the same entity. When comparing models, the better model is the one that minimizes the total information and then, subject to this constraint, minimizes the effective complexity. Together with additional constraints on computation time, the framework provides a method for identifying regularities or models that “most” effectively describe a given entity that in many regards is independent of the observer [6]. The minimization of the total information is closely related to the minimum description-length [11] and minimum messaging-length approaches [12].

C. Uniform subgraph covers

Following the above definitions, we define the total information of subgraph covers by embedding them into uniform subgraph covers. The uniform subgraph cover $E(m, n_m)$ is defined to be the uniform ensemble of all subgraph covers having motif set M and motif counts n_m . In order to compute the entropy, we need to compute the number of different ways a motif m can appear on N vertices, which depends on the automorphism group of the motif. An automorphism of a graph is a permutation of its vertex labels that preserves the edges of the graph. The number of all such permutations gives us the number of equivalent vertex labelings of the graph. To specify an instance of m on N vertices, one needs to specify the set of vertices m appears on and how it is embedded into this set. From the definition of the automorphism group, it follows that there are $|m|!/|\text{aut}(m)|$ different ways a motif m can appear on a set of $|m|$ vertices. Thus, a motif m with automorphism group $\text{aut}(m)$ can appear on N vertices in $N!/(N - |m|)!|\text{aut}(m)|$ different ways. Consequently, the entropy of a set of n_m distinct instances of m is given by

$$S(m, n_m) = \log \left(\frac{N!}{(N - |m|)!|\text{aut}(m)| n_m} \right), \quad (3)$$

which, given m , is the information required to specify n_m instances of m on N vertices.

Generalizing the above expression, the entropy of a cover C with motif set $M(C)$ and motif counts n_m ($m \in M$) is defined as the entropy of the uniform ensemble of all covers with motif counts n_m :

$$S(C) = \sum_{m \in M(C)} S(m, n_m). \quad (4)$$

When needed, the entropy terms can easily be approximated using Stirling's formula. For instance, when n_m and N are large enough and $|m| > 2$,

$$S(m, n_m) = n_m [|m| \log N - \log |\text{aut}(m)| - \log n_m + \log(e)] + O(\log N). \quad (5)$$

As in the case of the entropy, we define the effective complexity of a cover using uniform covers with the same motif counts: $\epsilon(C) = \text{AIC}_U[E(M(C), n_m)]$. Consequently, the total information of a cover can be defined as

$$\Sigma(C) = \epsilon(C) + S(C). \quad (6)$$

Following the total information approach, a cover is an optimal representation of the network if it minimizes the total information. As a result, we can define the network motifs of G to be the motif set of a Σ -optimal subgraph cover of G : $M[C_\Sigma(G)]$. The Σ -optimal subgraph cover also gives a decomposition of the network in terms of these motifs. In general, there might be more than one subgraph cover that minimizes Σ . If the Σ -optimal subgraph cover is not unique, additional criteria such as the minimization of the effective complexity [6] have to be considered in order to pick one of the solutions over the others.

The quantity $\epsilon[C_\Sigma(G)]$ can be interpreted as a measure of the complexity of G 's subgraph structure, which actually is in correspondence with other measures that are frequently used as indicators of a network's complexity, such as the broadness of the degree distribution and/or clustering. While the broadness of the degree distribution gives the variety of the star-shaped subgraphs that occur in the network, high clustering indicates that the network has a local structure that involves subgraphs other than trees.

Another quantity of interest is the amount of compression the optimal cover provides with respect to the edge cover of G , $C_e(G)$:

$$\Delta_\Sigma(G) = \Sigma[C_e(G)] - \Sigma[C_\Sigma(G)]. \quad (7)$$

$\Delta_\Sigma(G)$ gives us a measure of how much G deviates from an Erdős-Rényi-type random graph with respect to its subgraph structure. Motifs can also be assigned a quantitative significance similar to the z score used by Milo *et al.* [1] based on the compression they provide. One such measure is the c score c_m :

$$c_m(G) = \frac{\Sigma[C_\Sigma(G) - m]}{\Sigma[C_\Sigma(G)]} - 1, \quad (8)$$

where $[C_\Sigma(G) - m]$ is the cover obtained by replacing the m subgraphs in $C_\Sigma(G)$ with the single-edge subgraphs they cover. By definition, c_m is always non-negative and zero for the single-edge motif and motifs that are not contained in $C_\Sigma(G)$.

D. Subgraph covers and model selection

In this section, we will consider two models that are closely related to subgraph covers: the model introduced by

Bollobás, Janson, and Riordan [7] and the model introduced by Karrer and Newman [8]. Although these models can account for large densities of nontrivial subgraphs, it is not clear how one should select the set of motifs to be used in such models when matching these with a given network. In their article, Karrer and Newman [8] mention the fitting of such model to networks as an important open problem. In the following, we formulate these models in terms of subgraph covers and discuss how Σ -optimal subgraph covers can be used to associate networks with such models.

Random graphs with clustering.—In Ref. [7], Bollobás, Janson, and Riordan introduced a very general class of random graph models that is based on adding copies of certain motifs onto the vertices of a graph. For the sake of simplicity, we will only consider the homogeneous models, i.e., the case where all vertices have the same type. For the nonhomogeneous version of the model as well as various analytical results concerning the properties of the model, we refer the reader to the original paper [7]. In the homogeneous case, the model can be defined as follows: Let M be a set of motifs, each given by a labeled representative, and for every $m \in M$, let k_m be a positive constant that corresponds to the density of the motif in the model. Then, for each $m \in M$ and $|m|$ -tuple $(v_1, v_2, \dots, v_{|m|})$ of vertices, one adds a copy of m to G , such that the i th vertex of m is mapped onto v_i , with probability

$$p_m = \frac{k_m}{N^{|m|-1}}. \quad (9)$$

Since we are mainly interested in simple graphs, we will assume that any parallel edges that are formed in this process are replaced with single edges in the network. The normalization factor $1/N^{|m|-1}$ ensures that the model has $O(N)$ edges. Depending on the symmetry of m , the models formulated above allow for more multiple copies of the same m -subgraph to be added to the graph. Although the probability of such events is small, the model can be slightly modified to avoid such redundancies by considering $|m|$ -subset instead of $|m|$ -tuples. Then, for each such subset, every distinct m subgraph is added with probability

$$p_m = \frac{k_m |\text{aut}(m)|}{N^{|m|-1}}, \quad (10)$$

where the factor $|\text{aut}(m)|$ ensures that both models contain the same number of copies of m on average.

With the above modification, the model defines a multinomial distribution $P_{(M, \mathbf{k})}(\cdot)$ over the space of M covers. The corresponding distribution over graphs results from the projection of subgraph cover configurations onto edge configurations. Thus, the probability of a graph G in this model is given by

$$P_{(M, \mathbf{k})}(G) = \sum_{C \in \mathcal{C}_M(G)} P_{(M, \mathbf{k})}(C), \quad (11)$$

where $\mathcal{C}_M(G)$ is the set of all M covers of G .

Uniform subgraph covers are essentially microcanonical versions of these models. Consequently, the presented approach can be seen as a way of inferring the subgraph-cover state of such models. The Σ -optimal cover can be further used as a basis for associating the network with nonhomogeneous models of this type that also include correlations between subgraphs.

Generalized configuration models.—Another random graph model that is closely related to subgraph covers is the generalized configuration model proposed by Karrer and Newman [8]. This model is defined on the basis of a motif set M and a corresponding role sequence \mathbf{r} . Here, the role sequence specifies the number of different motifs attached to each vertex and how these motifs are attached to the vertex. The way in which a certain motif is attached to the vertex is given by the orbit of the automorphism group of the motifs the vertex belongs to. (The orbit of a vertex is the set of vertices it can be mapped onto by the automorphism group.) In order to generate a graph corresponding to a role sequence \mathbf{r} , every vertex is assigned a number of subgraph stubs corresponding to its role index. A graph is then generated by matching stubs corresponding to the same type of subgraph in appropriate combinations at random and connecting them in order to form the corresponding motif subgraph. However, in this form, the model allows for two or more stubs of the same vertex to be matched together, which results in a subgraph that is a vertex contraction of the original motif. When such problematic cases are excluded from the model, every matching of the stubs actually corresponds to an M cover. Consequently, the generalized configuration models can be formulated in terms of subgraph covers: The model corresponding to a role sequence \mathbf{r} is the uniform ensemble of all subgraph covers with role sequence \mathbf{r} .

Determining a role sequence for a network is essentially equivalent to choosing a subgraph cover for the network, since every subgraph cover produces a specific role sequence for the network. The Σ -optimal cover can be considered as a viable candidate for assigning a role sequence to a network. On the other hand, an important property of the generalized configuration models is that biconnected subgraph counts are essentially determined by the motif set while singly connected subgraphs can be mostly accounted for by the role sequence. Consequently, one can also consider restricting the analysis to biconnected subgraphs when determining a role sequence for the network. Restricting motifs to biconnected motifs also significantly reduces the number of subgraphs that have to be considered in the analysis since the majority of connected subgraphs of sparse networks is only singly connected.

The models described above suggest that, in principle, one could also consider more general or nonuniform ensembles of subgraph covers to define the total information. For instance, in the case of the generalized configuration model, one could use the ensemble of all subgraph covers that result in the same role sequence. However, there is no known simple way of calculating the entropy of such ensembles, even if only single-edge subgraphs are considered, which would be equivalent to the classical configuration model. In addition, such ensembles have high effective complexity.

E. The relation to the method of Milo *et al.*

From a compression point of view, the motifs of the Σ -optimal cover correspond to the motifs with respect to which the graph maximally deviates from a random graph. Although the Erdős-Rényi random graph is not, in general, the null model of choice in applications of the method of Milo *et al.*, both methods can be seen as trying to find motifs with respect to which the network differs from a random graph. In this sense, the subgraph-cover approach and the method of Milo *et al.* share similar goals. Using generalized configuration models to define the total information would allow for a more direct comparison of the two approaches, but as explained in the previous section, there are additional difficulties associated with using these models.

A general problem one faces when using the method of Milo *et al.* is that subgraph frequencies are, in general, highly interdependent. The most obvious of such dependencies are motif-submotif-type dependencies. That is, the presence of a motif implies the presence of larger motifs containing it as a submotif and its own submotifs. Motif-submotif-type dependencies further imply that counts of motifs that have a submotif in common are correlated. In order to avoid larger motifs to be classified as network motifs only because they contain some smaller overrepresented motif, Milo *et al.* propose using a null model that conserves lower-order motifs. However, conserving lower order motif counts does not cover all dependencies. For instance, a certain overrepresented motif might occur almost exclusively as a submotif of one or more larger network motifs. The Σ -optimal cover naturally avoids such interdependencies by simultaneously considering motifs of all sizes and effectively penalizing the sharing of edges between subgraphs. Despite their differences, one would expect both methods to find similar motifs, provided that lower-order motifs are conserved by the null model.

III. THE Σ -OPTIMAL SUBGRAPH-COVER PROBLEM

In general, when finding optimal subgraph covers, one would like to consider the most general set of potential

motifs. However, in practice, there are several technical limitations, the first being the graph isomorphism problem. That is, there exists no known polynomial time algorithm for resolving the problem of whether two finite graphs are isomorphic. The same holds for finding the automorphism group of a graph. Fortunately, there are several software packages that can efficiently compute the automorphism group of small graphs [13]. Second, the problem of finding whether a graph G contains a certain motif as a subgraph is NP complete. Thus, finding subgraph instances can be computationally expensive, especially for large motifs. Third, the number of connected motifs grows faster than exponentially with size. For instance, there are over a million different directed motifs of size 6. Therefore, the set of candidate motifs of which the subgraph instances are to be included in the analysis has to be restricted so that the analysis can be completed in reasonable time. Restricting the set of candidate motifs to all connected motifs up to a certain size seems to be an obvious choice. On the other hand, one can also include special classes of motifs of arbitrary size into the set of candidate motifs.

If one wants to include special classes of motifs into the set of candidate motifs, any prior knowledge of the structure of the network can be used to make an educated guess about which motifs are more likely to produce covers with small total information. For instance, when examining the network representing an electronic circuit, the motifs corresponding to various known subcomponents of the circuit should be included in the set of candidate motifs. Also, if the network at hand is known to have a broad degree distribution, star-shaped motifs can be included. Similarly, if some motifs are known to favor a certain type of dynamical behavior that is thought to be relevant to the network performing certain tasks, these patterns and their generalizations can be included into the analysis. As previously mentioned, if one intends to use the subgraph cover in order to determine a role degree sequence for the network, the set of motifs can be restricted to biconnected motifs. Disconnected motifs can be excluded from the analysis since the cover that independently contains the connected components of such subgraphs always has lower total information.

Another issue that has to be addressed in practice is that the algorithmic information content is not computable and in addition is computer dependent. In practice these difficulties can be circumvented by substituting the algorithmic information content of the ensemble with the code length of a reasonable encoding of it. Another simplification we make is to assume that motifs are independently encoded, which results in an effective complexity term that is additive in the motifs. One obvious way of encoding motifs is to use edge lists. In this case, we have

$$\epsilon(m) = \log^*[|V(m)|] + \log^*[|E(m)|] + S[|V(m)|, |E(m)|],$$

where $S[|V(m)|, |E(m)|]$ is the entropy of the ensemble of all graphs, with the same vertex and edge counts as m , and \log^* is the iterated logarithm. On the other hand, one can also use a predefined or fixed encoding or catalog of the candidate motifs to define their effective complexities.

After the simplifications above, the total information reduces to

$$\Sigma(C) = \log^* N + \sum_{m \in M(C)} [S(m, n_m) + \epsilon(m) + \log^* n_m]. \quad (12)$$

The choice of code used to define the effective complexity depends on the set of candidate motifs. The edge-list encoding has the advantage of being independent of the set of candidate motifs and therefore is a natural choice when considering all motifs up to a certain size. On the other hand, given a specific set of candidate motifs, the catalog approach, in general, results in shorter code lengths compared to the edge-list encoding. Therefore, the catalog approach is suitable when the set of candidate motifs contains special classes such as cliques, stars, cycles, etc., since these have obvious shorter encodings than their edge lists.

Although the edge-list encoding seems to be a natural candidate when considering general motifs, it might be argued that the choice of code used to approximate the effective complexity term is subjective. However, in practice, almost all comparable model-selection approaches require similar subjective choices [11,12,14]. From a Bayesian viewpoint, where the effective complexity term can be associated with a prior, the subjectivity of this choice might be less of a problem. The main goal of this article is not to advocate a specific approach to model selection but rather to show that subgraph covers can be used as a basis for motif analysis. The reason for us choosing the total information approach is that it accounts for the fact that the parameters of the ensembles are graphs in a rather intuitive way.

Another, more practical, way of looking at the effective complexity is as a safeguard against overfitting. From this point of view, the effective complexity of a motif corresponds to the minimal entropy gain it has to provide in order to be included in the optimal cover. Therefore, the effective complexity of a motif can be seen as setting a frequency threshold for the motif. Thus, the problem can also be formulated in terms of entropy minimization and frequency constraints. If overfitting is less of a concern, one can even set the effective complexity to 0, which in our case is equivalent to the maximum-likelihood approach. Using maximum-likelihood or similar reduced effective complexity terms or frequency thresholds might be useful if one wants to find a maximal number of potentially relevant motifs or when the network is small ($N < 200$) since in

such networks the maximal potential entropy gains of motifs are limited. The algorithms we shall present in the next section can be modified in a straightforward manner to incorporate frequency thresholds.

A. The greedy algorithm

Even with the candidate motifs restricted, finding a Σ -optimal subgraph cover is a nontrivial optimization problem. As formulated above, the problem of finding a Σ -optimal subgraph cover is a nonlinear set-covering problem, where the set to be covered is the edge set of the graph and the subsets are the edge sets of the subgraph instances of the candidate motifs. Set-covering problems are known to be NP hard, even in the linear case [15]. Consequently, in most practical applications, exact solutions are elusive and a heuristic has to be used.

The greedy algorithm we propose might be seen as a generalization of the greedy heuristic introduced by Chvatal [16] for the linear set-covering problem. The algorithm is based on the stepwise construction of a subgraph cover. At each step, the algorithm finds the motif that covers not-yet-covered edges of G most efficiently in terms of total information per edge. Given a partial cover C , the efficiency of a set S_m of m subgraphs is defined as

$$\sigma(S_m, C) = \frac{\Sigma(S_m)}{|E(S_m) - E(C)|}, \quad (13)$$

where $E(C)$ and $E(S_m)$ are the set of edges covered by C and S_m , respectively, and $\Sigma(S_m)$ is the total information corresponding to S_m . More precisely, $\Sigma(S_m) = S(m, |S_m|) + \epsilon(m) + \log^*(|S_m|)$. Following this definition, an optimal instance set of m is defined as a set of m subgraphs that minimizes σ . At each step, the algorithm determines the efficiency of all motifs in the candidate motif set by determining an optimal instance set for each of them. In the next step, the algorithm checks for each motif whether including its optimal instance set into the cover increases the overall total information of the cover. Then, from the motifs of which the optimal instance set does not increase the total information, the most efficient one is selected. Having found the most efficient motif, the corresponding optimal instance set is added to the cover and the set of covered edges is updated. The process is repeated until all edges of the graph are covered. To ensure that the algorithm terminates, we require that the single-edge motif is always included in the set of candidate motifs. The total information of partial covers is calculated by adding to them the single-edge subgraphs corresponding to the uncovered edges. Here, one should note that motifs cannot be selected based solely on their efficiency because, in general, adding the optimal instance set of a motif to the cover decreases the efficiency of other motifs which, in certain cases, might lead to an increase of the overall total information.

Algorithm 1 GreedyOptimalCover [$G(E, V), MS$]

```

CoveredEdges =  $\emptyset$ , Cover =  $\emptyset$ , Motifs =  $\emptyset$ 
While  $|\text{CoveredEdges}| < |E|$  do
   $C, m = \text{FINDMOTIF}(G, MS, \text{CoveredEdges})$ 
   $\text{CoveredEdges} \leftarrow \text{CoveredEdges} \cup_{i \in C} e(i)$ 
   $\text{Cover} \leftarrow \text{Cover} \cup C$ 
   $\text{Motifs} \leftarrow \text{Motifs} \cup \{m\}$ 
end While
return Cover, Motifs
function FINDMOTIF ( $G, MS, \text{CoveredEdges}$ )
  for  $m \in MS$  do
     $C(m) = \text{OptimalInstanceSet}[m, \text{CoveredEdges}, G(E, V)]$ 
  end for
   $M = \arg \min_{m \in MS} \{\sigma[C(m), \text{CoveredEdges}] | \Sigma[\text{Cover} \cup C(m)] \leq \Sigma(\text{Cover})\}$ 
  return  $C(M), M$ 
end function

```

Here, OptimalInstanceSet is a function that computes an optimal instance set, given a motif and a set of covered edges, Σ is the total information, and MS is the set of candidate motifs.

Given a motif m and a set of covered edges, finding an optimal instance set is a nontrivial optimization problem on its own. When subgraphs in the cover are not allowed to share edges, finding an optimal instance set is equivalent to finding a maximum independent set of m subgraphs, that is, a set of m subgraphs of maximum cardinality such that no two of the subgraphs in the set have an edge in common. This problem is equivalent to the maximum independent vertex-set problem and is NP complete [15]. As a result, some type of heuristic has to be employed. The descriptions of two such heuristics can be found in the Supplemental Material [17]. Depending on the heuristic, finding an optimal instance set requires some or all of the subgraph instances of m to be computed. There exist several well-known algorithms that can be used for this purpose [18,19].

IV. EMPIRICAL RESULTS

In the following, we apply the greedy heuristic to several real-world networks from different fields. We also consider some synthetic networks that are realizations of uniform subgraph covers with predetermined motif frequencies in order to test the heuristic.

Because of the computational resources available, the size of the subgraphs used in the analysis is limited to five in the directed and six in the undirected cases. We also consider biconnected subgraph covers in relation with generalized configuration models. All results are obtained using the maximal independent-set heuristic (for details, see the Supplemental Material [17]) for finding optimal instance sets and edge lists for encoding motifs. In the following tables, N and E stand for the number of vertices and edges, respectively. In addition to the total information of the obtained cover Σ , the tables also show the total

information of the corresponding edge cover/Erdős Renyi ensemble ERI as a benchmark. Both these quantities are rounded to the closest integer and are given in bits.

Because of the random choices involved in finding optimal instance sets, the algorithm might find different covers for the same network on different runs. The covers shown in the figures are the best solutions obtained over multiple runs. A more detailed discussion on the variability of the heuristic can be found in the Supplemental Material [17]. In Table I, the ranges of motif counts obtained over ten runs are also shown. Here, we should stress that the proposed heuristics are primarily aimed at demonstrating the feasibility of using Σ -optimal subgraph covers as a basis of motif analysis and other heuristics might be devised for the resulting covering problem.

A. Real-world networks

Table I shows the motifs found for the network representing the Western State Power Grid of the United States [20]. The motif structure indicates that, among other motifs, cycles and cliques play an important role in the organization of this network.

In Tables II and III, the motifs found for the transcription networks of the bacteria *E. coli* [21] and the yeast *S. cerevisiae* [1] are shown. For Table II, only biconnected motifs up to size 5 were considered, while Table III shows the results obtained using all connected subgraphs up to size 5. Including singly connected motifs in the candidate motif set has almost no effect on the biconnected motifs and mostly results in star-shaped motifs or motifs that consist of one vertex intersection of previously found biconnected motifs. The networks share three out of four motifs in the case of biconnected motifs.

In Table IV, we see the results for two networks representing electronic circuits that are digital fractional multipliers [1]. For these networks, the algorithm not only finds the same motifs for both networks but the motif counts in the optimal cover also scale almost exactly with network size.

TABLE I. The network motifs of the network representing the Western States Power Grid of the United States found using connected subgraphs up to size 6. The ranges of the motif counts of the covers found over ten runs are also shown in parentheses.













Network	N	E	ERI	Σ			
Power grid	4941	6594	81 084	77 109			
							
4109	141	112	44	31	17	47	
(4109–4129)	(138–145)	(111–122)	(44–45)	(30–31)	(15–17)	(45–47)	
							
11	68	2	15	42	2		
(10–11)	(67–68)	(2–2)	(15–16)	(41–43)	(2–2)		

TABLE II. The network motifs of the transcription networks of *E. coli* and *S. cerevisiae* obtained using all biconnected motifs up to size 5.
















Motifs (ID)						
						
Network	n_0	n_1	n_2	n_3	n_4	n_5
<i>E. coli</i>	323	9	14	0	13	5
<i>S. cerevisiae</i>	547	0	23	4	60	8
Network	N	E	ERI	Σ		
<i>E. coli</i>	423	519	5124	4810		
<i>S. cerevisiae</i>	688	1079	11 024	9811		

TABLE III. The motifs of the transcription networks of *E. coli* and *S. cerevisiae* obtained using all connected motifs up to size 5.

Motifs (ID)									
									
(0)	(1)	(2)	(3)	(4)	(5)				
									
(6)	(7)	(8)							
Network	n_0	n_1	n_2	n_3	n_4	n_5	n_6	n_7	n_8
<i>E.coli</i>	130	0	0	12	51	0	13	5	4
<i>S.cerevisiae</i>	59	26	16	23	94	5	61	8	0
Network	N	E	ERI	Σ					
<i>E.coli</i>	423	519	5124	4637					
<i>S.cerevisiae</i>	688	1079	11 024	9309					

In Table V, the network motifs found for the metabolic networks [22] of several species from different domains of life are shown. We find almost the same motifs in all of these networks, and most motif counts also scale approximately with network size.

Table VI shows the network motifs found in networks representing the Internet at the level of autonomous systems [23]. The counts of common motifs also scale approximately with network size.

In commonly analyzed networks, we find three- and four-node motifs that are almost identical to those found by Milo *et al.* [1] using a null model that conserves lower-order motifs. In the transcription networks, we find all the motifs found by Milo *et al.*, although in the case of the *S. cerevisiae* network, the feed-forward loop only appears as a submotif of the larger motif consisting of three feed-forward loops sharing an edge. Also, for the electronic circuit networks, we find the same three- and four-node motifs, although the three- and four-cycles appear only as submotifs. A closer analysis shows that in these networks, three- and four-cycles occur almost exclusively as

TABLE IV. The motifs of electronic circuits (digital fractional multipliers) obtained using all connected motifs up to size 5.


















Motifs (ID)					
					
(0)	(1)	(2)	(3)	(4)	
Network	n_0	n_1	n_2	n_3	n_4
$s420$	220	7	4	13	11
$s838$	456	15	8	25	23
Network	N	E	ERI	Σ	
$s420$	252	399	3491	3404	
$s838$	512	819	7995	7652	
















TABLE V. The motifs of various metabolic networks obtained using biconnected subgraphs up to size 5. AA = Aquifex aeolicus (bacteria), AB = Actinobacillus actinomycetemcomitans (bacteria), EC = Escherichia coli (bacteria), CE = Caenorhabditis elegans (eukaryote), AG = Archaeoglobus fulgidus (archaea), and AP = Aeropyrum pernix (archaea).

Motifs (ID)												
												
												
Network	n_0	n_1	n_2	n_3	n_4	n_5	n_6	n_7	n_8	n_9	n_{10}	n_{11}
AA	423	16	0	6	16	130	147	97	0	0	0	0
AB	408	22	0	4	23	128	131	82	2	0	0	0
EC	935	117	3	5	40	264	345	202	0	3	5	0
CE	478	13	2	3	31	137	178	100	0	0	0	0
AG	509	23	2	6	26	140	168	120	0	0	0	4
AP	195	11	0	0	12	55	67	46	0	0	0	0
Network	N	E	ERI	Σ								
AA	1057	2527	25 844	21 255								
AB	993	2368	24 012	19 882								
EC	2275	5763	64 842	52 590								
CE	1173	2864	29 634	24 380								
AG	1268	3011	31 616	25 960								
AP	490	1163	10 610	8856								

subgraphs of larger network motifs (three-cycles: *s420*-19/20, *s822*-39/40; four-cycles: *s420*-11/11, *s838*-23/23).

The analysis of various networks shows that networks having similar functions also have similar motif structures. The results further support that motifs play an important role in the structural organization of complex networks. Furthermore, for networks of the same type, the motif counts also scale approximately with network size. The results also show that subgraph covers can be used to obtain representations that are up to 20% shorter compared to edge-list representations.

TABLE VI. The motifs of various autonomous system networks obtained using all connected motifs up to size 5.

Motifs (ID)															
															
															
Network	n_0	n_1	n_2	n_3	n_4	n_5	n_6	n_7	n_8	n_9	n_{10}	n_{11}	n_{12}	n_{13}	n_{14}
AS-1	710	31	0	55	7	495	227	34	7	41	0	0	17	5	9
AS-2	695	18	0	73	9	505	226	35	10	40	0	0	13	9	8
AS-3	791	18	0	85	14	507	243	39	10	44	10	0	15	7	9
AS-4	1016	28	0	124	5	650	356	55	17	48	13	21	24	7	11
AS-5	980	31	0	127	10	615	368	59	17	54	12	30	38	9	18
AS-6	1227	30	0	183	14	826	569	83	23	83	13	23	35	12	13
AS-7	3162	27	0	127	8	0	299	45	13	71	0	27	34	11	14
AS-8	913	38	144	189	18	981	682	87	30	98	14	30	47	18	14
Network	N	E	ERI		Σ										
AS-1	3015	5156	57 877		52 769										
AS-2	3042	5232	58 755		53 557										
AS-3	3213	5624	63 458		57 784										
AS-4	4235	7674	89 266		80 224										
AS-5	3962	7931	90 351		80 411										
AS-6	5599	10 728	128 250		113 240										
AS-7	3570	7033	79 225		71 227										
AS-8	6474	12 572	152 686		134 176										

As previously mentioned, Σ -optimal subgraph covers can be used as a basis for associating networks with generalized configuration models, which can be used to make various predictions about the properties of these networks. The method can be further tested by comparing properties of the analyzed networks with these models. However, such comparisons are beyond the scope of this article and will be treated separately in later articles.

B. Network classification

The method can also be used to obtain a classification of networks that is similar to the classification given by Milo *et al.* in Ref. [24]. Following the approach of Ref. [24], we use motif-significance profiles that are given by the normalized c score \tilde{c}_m :

$$\tilde{c}_m = \frac{c_m}{\sqrt{\sum_{m' \in M(C_\Sigma)} c_{m'}^2}}. \quad (14)$$

The significance profiles of various network types are given in Fig. 1. We find that networks of the same type also have very similar significance profiles. The figures only show the regions of the significance profiles corresponding to the motifs that have a nonzero c score. Compared to the full significance profile, these regions are relatively small since there are 9578 connected and 7585 biconnected motifs up to size 5 in the directed case and 30 connected

motifs up to size 5 in the undirected case. Thus, compared to the subgraph profiles used in Ref. [24], the method provides a much finer-grained classification.

C. Synthetic networks

Finally, we also consider some synthetic networks that are realizations of uniform subgraph covers with predetermined motif counts in order to test whether the heuristic is able to recover the underlying motif set or subgraph cover in such cases. As shown in Table VII, for all random networks, the algorithm is able to recover the motif set. For network 1, the algorithm exactly recovers the underlying subgraph cover. Network 2 is generated to mimic the motif structure found for an electronic circuit (s838; see Table IV), and the algorithm is able to recover the original subgraph cover with only one extra subgraph. On the other hand, for networks 3 and 4, the motif counts differ significantly from the counts of the uniform subgraph covers used to generate the networks, especially with respect to the five-star counts. This deviation is probably caused by the fact that these networks contain a large number of five-stars, of which only some are explicitly contained in the underlying cover. Consequently, finding an optimal instance set of five-stars becomes more difficult. This effect is more pronounced in network 3 because it is denser compared to network 4.

As is the case with any heuristic, the quality of the solution depends on the structure of the network. One can

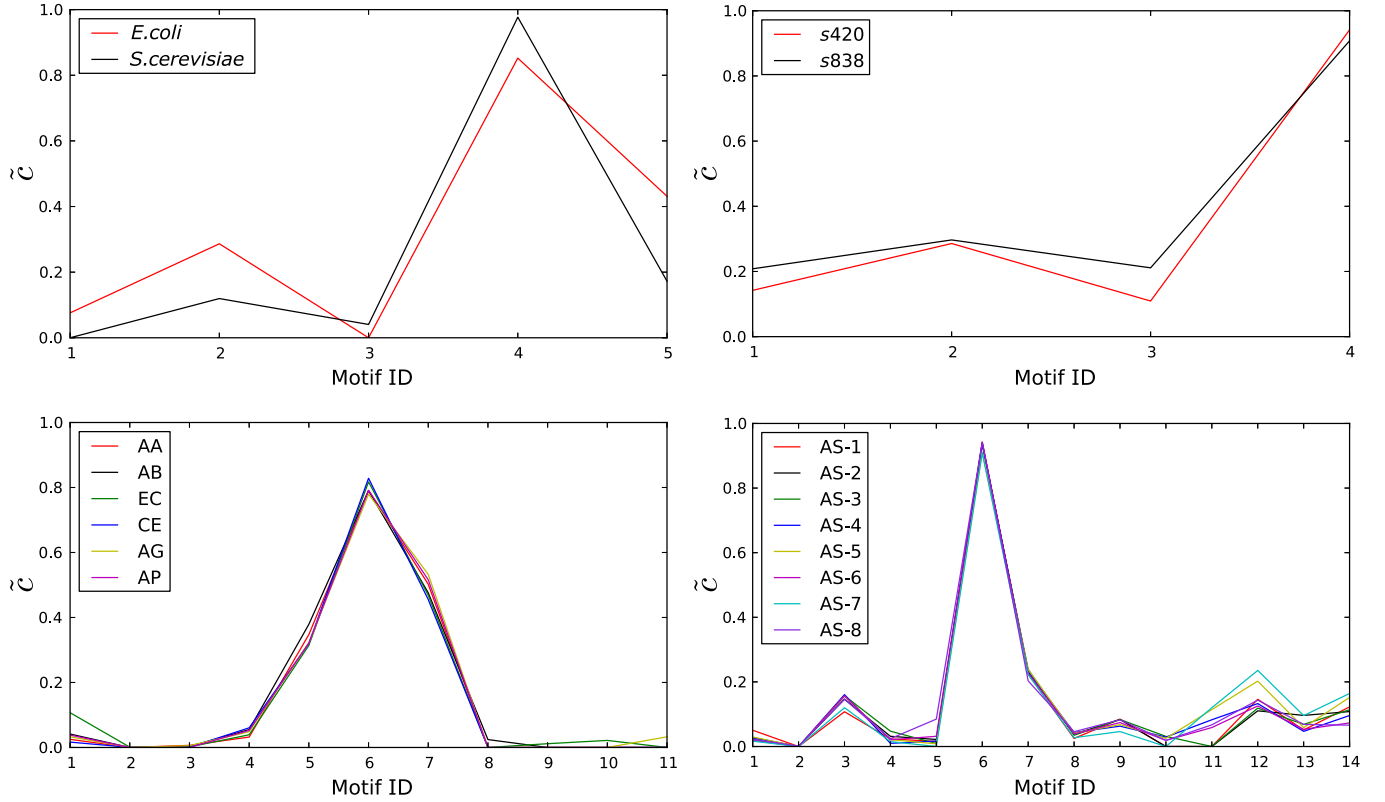


FIG. 1. The motif-significance profiles of various networks corresponding to the covers given in Tables II, IV, V, and VI.

construct examples where the greedy heuristic fails to recover all the motifs used to generate the network. In general, the greedy heuristic favors patterns that are dense, symmetric, and occur in large numbers in the network. Thus, if the graph contains only a few copies of a motif that is not very dense, the algorithm might not be able to recover that motif. Also, if a motif contains a submotif that is more dense and symmetric compared to the entire motif, the greedy algorithm might pick the submotif over the motif itself since the submotif covers edges more efficiently.
























V. DISCUSSION

In this article, we introduced an alternative approach to motif analysis in networks that is based on finding a subgraph cover of the network that represents it using minimal total information. We proposed a heuristic for the resulting NP hard optimization problem. The subgraph covers obtained for various networks show that the algorithm finds nearly identical motifs for networks with similar functions. Moreover, by considering subgraphs of various sizes simultaneously and with respect to a single global measure, the method is able to detect even large motifs consistently.

Another advantage of the method is that it provides an explicit decomposition of the network into motif subgraphs which allows motifs to be studied within the context of the rest of the network rather than in isolation.

We also showed that total information-optimal subgraph covers can be used to match networks with random graph models that incorporate the obtained motif structure which allows more accurate modeling of networks in general.

TABLE VII. The motifs obtained for networks corresponding to realizations of uniform subgraph covers. The quantities corresponding to the ensembles used to generate the networks are given in parentheses.

Network	N	E	ERI	Σ	
Network 1	1500	3115	34 069	20 566(20 566)	
					
150(150)	80(80)	125(125)	75(75)	125(125)	75(75)
Network 2	512	819	7795	7646(7652)	
					
452(456)	16(15)	8(8)	25(25)	23(23)	
Network 3	750	2065	17 594	16 296(16 089)	
					
273(150)	124(125)	119(125)	45(45)	47(45)	78(100)
Network 4	1500	2065	21 746	18 413(18 277)	
					
192(150)	123(125)	120(125)	44(45)	47(45)	95(100)

Subgraph covers can readily be generalized to graphs with labeled or colored vertices and edges as well as graphs with parallel and self-edges. Such labels might be chosen so that they correspond to known functional roles of vertices or the community structure of the network. On the other hand, the obtained subgraph covers could also be used as a starting point for detecting communities in networks or for inferring functional roles of vertices. Communities that differ with respect to their internal organization can also be expected to differ with respect to their motif structure. Similarly, one would expect the functional role of a vertex to be strongly correlated with the motifs it is a part of.

The total information approach can also be extended to ensembles more general than uniform subgraph covers. Moreover, model-selection approaches other than the total information approach can also be used. Such alternative formulations essentially correspond to using a different cost function in the covering problem.

The presented analysis strongly suggests that subgraph covers can be used to compress network data. In such applications, the total information might be replaced by the expected code length of the subgraph cover.

Finally, there is also room for improvement on the side of the heuristics. We consider the development of alternative algorithms to be an important topic for further research. While the greedy algorithm can be improved, other widely used approximation schemes such as simulated annealing or genetic algorithms can also be applied to the problem.

ACKNOWLEDGMENTS

The geng and directg tools that are a part of the NAUTY [13] package by Brendan D. McKay were used in generating the various isomorphism classes used in our analysis and the graph tool [25] PYTHON package developed by Tiago de Paula Peixoto for finding subgraphs and manipulating graphs in general. The author thanks Professor Jürgen Jost for constructive discussion as well as Güven Demirel and Murat Saglam for the critical reading of the manuscript. The author acknowledges the financial support of the International Max-Planck Research School Mathematics in the Sciences.

-
- [1] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon, *Network Motifs: Simple Building Blocks of Complex Networks*, *Science* **298**, 824 (2002).
 - [2] U. Alon, *Network Motifs: Theory and Experimental Approaches*, *Nat. Rev. Genet.* **8**, 450 (2007).
 - [3] S. Wernicke and F. Rasche, *Fanmod: A Tool for Fast Network Motif Detection*, *Bioinformatics* **22**, 1152 (2006).

- [4] F. Schreiber and H. Schwöbbermeyer, *Mavisto: A Tool for the Exploration of Network Motifs*, *Bioinformatics* **21**, 3572 (2005).
- [5] M. E. J. Newman, *Handbook of Graphs and Networks: From the Genome to the Internet* (Wiley-Vch, Weinheim, 2003).
- [6] M. Gell-Mann and S. Lloyd, *Information Measures, Effective Complexity, and Total Information*, *Complexity* **2**, 44 (1996).
- [7] B. Bollobás, S. Janson, and O. Riordan, *Sparse Random Graphs with Clustering*, *Random Struct. Algorithms* **38**, 269 (2011).
- [8] B. Karrer and M. E. J. Newman, *Random Graphs Containing Arbitrary Distributions of Subgraphs*, *Phys. Rev. E* **82**, 066118 (2010).
- [9] M. E. J. Newman, *The Structure and Function of Complex Networks*, *SIAM Rev.* **45**, 167 (2003).
- [10] J.-L. Guillaume and M. Latapy, *Bipartite Structure of All Complex Networks*, *Inf. Proc. Lett.* **90**, 215 (2004).
- [11] J. Rissanen, *Modeling by Shortest Data Description*, *Automatica* **14**, 465 (1978).
- [12] C. S. Wallace, *Statistical and Inductive Inference by Minimum Message Length* (Springer, New York, 2005).
- [13] B. D. McKay and A. Piperno, *Practical graph isomorphism, II*, *J. Symb. Comput.* **60**, 94 (2014).
- [14] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms* (Cambridge university press, Cambridge, 2003), Vol. 7.
- [15] M. R. Garey and D. S. Johnson, *Computers and Intractability* (Freeman, San Francisco, 1979), Vol. 174.
- [16] V. Chvatal, *A Greedy Heuristic for the Set-Covering Problem*, *Math. Oper. Res.* **4**, 233 (1979).
- [17] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevX.4.041026> for the description of two heuristics, of varying computational complexity, for computing optimal instance sets of motifs.
- [18] J. R. Ullmann, *An Algorithm for Subgraph Isomorphism*, *J. Assoc. Comput. Mach.* **23**, 31 (1976).
- [19] L. P. Cordella, P. Foggia, C. Sansone, and M. Vento, *A (Sub) Graph Isomorphism Algorithm for Matching Large Graphs*, *IEEE Trans. Pattern Anal. Mach. Intell.* **26**, 1367 (2004).
- [20] D. J. Watts and S. H. Strogatz, *Collective Dynamics of Small-World Networks*, *Nature (London)* **393**, 440 (1998).
- [21] S. S. Shen-Orr, R. Milo, S. Mangan, and U. Alon, *Network Motifs in the Transcriptional Regulation Network of Escherichia Coli*, *Nat. Genet.* **31**, 64 (2002).
- [22] H. Jeong, B. Tombor, R. Albert, Z. N. Oltvai, and A.-L. Barabási, *The Large-Scale Organization of Metabolic Networks*, *Nature (London)* **407**, 651 (2000).
- [23] J. Leskovec, J. Kleinberg, and C. Faloutsos, in *Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery in Data Mining* (ACM, Chicago, IL, 2005), p. 177–187.
- [24] R. Milo, S. Itzkovitz, N. Kashtan, R. Levitt, S. Shen-Orr, I. Ayzenshtat, M. Sheffer, and U. Alon, *Superfamilies of Evolved and Designed Networks*, *Science* **303**, 1538 (2004).
- [25] <http://graph-tool.skewed.de>.