Hierarchical Hidden Community Detection for Protein Complex Prediction

Chao Li

Huazhong University of Science and Technology - Main Campus: Huazhong University of Science and Technology <https://orcid.org/0000-0001-9066-1440>

Kun He (✉ brooklet60@hust.edu.cn)

Huazhong University of Science and Technology - Main Campus: Huazhong University of Science and Technology

Guang shuai Liu

Huazhong University of Science and Technology - Main Campus: Huazhong University of Science and Technology

John E. Hopcroft

Cornell University

Research article

Keywords: Protein complex, Community detection, Interaction networks, Hidden structure

Posted Date: December 3rd, 2020

DOI: <https://doi.org/10.21203/rs.3.rs-116708/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

RESEARCH

Hierarchical hidden community detection for protein complex prediction

Chao Li¹, Kun He^{1*}, Guangshuai Liu¹ and John E. Hopcroft²

Abstract

Background: Discovering functional modules in protein-protein interaction networks through optimization remains a longstanding challenge in Biology. Traditional algorithms simply consider strong protein complexes found in the original network by optimizing some metric, which may cause obstacles for discovering weak and hidden complexes that are overshadowed by strong complexes. Additionally, protein complexes have not only different densities but also various ranges of scales, making them extremely difficult to be detected. We address these issues and propose a hierarchical hidden community detection approach to predict protein complexes of various strengths and scales accurately.

Results: We propose a meta-method called HirHide (Hierarchical Hidden Community Detection). It is the first combination of hierarchical structure with hidden structure, which provides a new perspective for finding protein complexes of various strengths and scales. We compare the performance of several standard community detection methods with their HirHide versions. Experimental results show that the HirHide versions achieve better performance and sometimes even significantly outperform the baselines.

Conclusions: HirHide can adopt any standard community detection method as the base algorithm and enable it to discover hidden hierarchical communities as well as boosting the detection of strong hierarchical communities. Some biological networks are too complex for standard community detection algorithms to produce a positive performance. Most of the time, a better choice is to choose a corresponding algorithm based on the characteristics of a specific biological network. Under these circumstances, HirHide has clear advantages because of its flexibility. At the same time, according to the natural hierarchy of cells, organelle, intracellular compound *etc.*, hierarchical structure with hidden structure is in line with the characteristics of the data itself, thus helping researchers to study biological interactions more deeply.

Keywords: Protein complex; Community detection; Interaction networks; Hidden structure

Background

A protein complex is a group of proteins that interact with each other for specific biological activities [1]. The identification of protein complexes is crucial for predicting protein functions [2–5], disease genes [6, 7], phenotypic effects of genetic mutations [8], and drug-disease associations [9]. Given a protein-protein interaction (PPI) network, where nodes represent proteins and edges represent interactions, the protein complexes can be searched by detecting densely connected subgraphs in the network. Mathematically, such subgraphs are called communities, in which nodes are joined together in tightly-knit groups, and there

are only looser inter-connections among the groups [10, 11].

Community detection plays a significant role in the biological network analysis and provides insight into the underlying structure existing in networks [10, 11]. Over the last decade, numerous algorithms have been proposed to detect communities in social and biological networks. Early works focus primarily on identifying disjoint communities that partition the nodes [12–14], then researchers observed the overlapping membership among the communities and develop overlap community detection techniques [15–18]. Some partition-based community detection algorithms are also extended to address the overlapping case [19, 20]. However, all these algorithms focus primarily on the prominent community structure in networks, which is potentially problematic. Weak communities that are

* Correspondence: brooklet60@hust.edu.cn

¹School of Computer Science and Technology, Huazhong University of Science and Technology, 430074 WuHan, China

Full list of author information is available at the end of the article

shielded by some dominant communities are sometimes of high value. For example, real-world protein complexes are not always dense, and sometimes they can be very sparse [21]. As proteins typically get involved in several interactions, there exist many overlaps among the protein complexes, and thus weak protein complexes are usually hidden behind the stronger ones.

Additionally, there could be some undiscovered protein interactions [22]. Therefore we have no connection between these proteins in PPI networks. As a consequence, sparse protein complexes may be only sparse in existing incomplete PPI networks. But they are more likely to be dominant in real-world PPI networks if all protein interactions had been discovered. That is, sparse protein complexes may also be potential strong protein complexes. However, these sparse complexes are overlooked by previous technique.

Though standard community detection methods can find a portion of the sparse protein complexes by detecting weak communities, they cannot deal with the case where most nodes of the weak communities also belong to other stronger communities. Under such a case, these weak communities are defined as hidden communities [23, 24]. In PPI networks, we call these communities the hidden protein complexes. For instance, in Fig. 1 (a), we build a network with hidden community structures. Because standard algorithms focus on discovering dominant communities, the weaker community with green nodes is generally overlooked. Even if it is detected, its structure is considered as the structure of green nodes in Fig. 1 (a), which contains four smaller blocks with dense intra-connections. However, its structure is actually like the structure of green nodes in Fig. 1 (b). The community in (b) becomes detectable because the stronger communities have been weakened. And as edges belonging to stronger communities are removed, this hidden community doesn't contain four smaller blocks with dense intra-connections like (a) shows. Mathematical graph representations of many protein complexes related to sparser communities are hidden communities, we call them hidden complexes. They are partially or completely covered by stronger protein complexes, and finding such hidden complexes is very difficult. Traditional algorithms simply consider hidden protein complexes as a part of the stronger protein complexes, causing big obstacles for the hidden ones to be discovered, which can partially explain why standard methods are not working well.

To address this problem, we design HirHide, which is inspired by a meta-approach called HiCode (Hidden Community Detection) [23, 25], a novel approach that first addresses the hidden community structure. However, HiCode does not consider hierarchical community

structure, and hence it cannot handle complicated networks where communities are organized hierarchically, which exists in many real-world networks. In large-scale PPI networks, many protein complexes are organized hierarchically, indicating that protein complexes may consist of sub-complexes extending to several hierarchical levels deep. An example of such deeply embedded complex is the SAGA complex (MIPS identifier 510.190.10.20.10), a multi-functional coactivator that regulates the transcription by RNA polymerase II [26].

The main contributions of this work include:

- We propose HirHide that combines the detection of hierarchical structure with hidden structure, and could detect communities of various strengths (related to density) as well as communities of various scales (related to size).
- HirHide is designed as a general method that can be combined with standard community detection methods and enables them to discover hidden hierarchical communities and boosts the detection of dominant hierarchical communities.
- We also propose the concept of hidden complexes and try to explain and relate the biological significance of hidden complexes based on gene ontology.

For easy distinction, we use **level** to denote the stratification of hierarchical structure and **layer** to denote the stratification of hidden structure throughout the experiments. Multi-level hierarchical structure together with multi-layer hidden structure build a multi-granular characteristic of HirHide (see Fig. 1 (c)). In experiments, we set two layers for hidden structure (one is dominant, the other is hidden), two levels for hierarchical structures to avoid over-complication (a total of 4 granular divisions). Layer1-level1 represents the first level of layer 1, which corresponds to the root, strong communities; Layer1-level2 indicates the second level of layer 1, which is a more detailed division of the root, strong communities; Layer2-level1 indicates the root, hidden communities in layer 2; Layer2-level2 denotes a more detailed division of the root, hidden communities. In the following, the four community divisions are shortly recorded as L_{11} , L_{12} , L_{21} , and L_{22} respectively.

Methods

Hiddenness Value

A PPI network can be represented as a graph $G = (V, E)$, where V is the node set and E is the edge set. Suppose the network is divided into K communities, denoted by $\mathcal{C} = \{C_1, \dots, C_k, \dots, C_K\}$.

Let $\mathcal{S}_k$ be the set of communities stronger than a community C_k :

$$\mathcal{S}_k = \{C_i | F_i > F_k, C_i \in \mathcal{C}\}. \quad (1)$$

Figure 1 Example to show the hidden hierarchical structure. (a) A network with hidden hierarchical structure. The graph exists of four large dense communities consisting of 128 nodes, each with an internal subdivision of four small communities with 32 nodes. Additional, there is a sparse, hidden community structure consisted by the four small blocks in the center. (b) The network after weakening the stronger communities. The community with green nodes is the hidden community, detected by weakening the stronger communities. So it doesn't contain dense edges. (c) Schematic diagram of a hidden hierarchical structure. Level and layer are separately used to denote the stratification of hierarchical structure and hidden structure. Each level of hierarchical structure can contain several layers of hidden structure, and vice versa.

F_k represents the strength of community C_k , which can be calculated by the modularity metric [11]. The larger the value is, the stronger the corresponding community structure is. The modularity Q is defined as:

$$Q = \frac{1}{2m} \sum_{i,j} \left[A_{ij} - \frac{d_i d_j}{2m} \right] \delta(Com_i, Com_j), \quad (2)$$

where m indicates the number of edges in the graph, d_i , d_j separately indicate the degree of nodes i and j , Com_i , Com_j separately indicate the community to which nodes i and j belong. And $\delta(Com_i, Com_j)$ indicates whether nodes i, j are in the same community. If so, $\delta(Com_i, Com_j) = 1$, and otherwise 0.

He et al. raise a formula to calculate the hiddenness value of a community [23]. This definition calculates the fraction of nodes of C_k belonging to other stronger communities.

$$H(C_k) = \frac{|\bigcup_{C_i \in S_k} C_i \cap C_k|}{|C_k|}. \quad (3)$$

The larger the hiddenness value is, the higher the probability the community is hidden, and the less likely the community can be detected. Note that there is no single specific threshold between a hidden community and a dominant community because the hiddenness value indicates the fraction of nodes of C_k belonging to other stronger communities. For two communities $C_i, C_j \in \mathcal{C}$, if $H(C_i) > H(C_j)$, then community C_i is comparatively hidden compared to C_j , and community C_j is comparatively dominant compared to C_i .

Hidden communities versus hierarchical communities

Hierarchical community structure and hidden community structure are two aspects of community structure. An algorithm combining these two concepts may

be confusing. For example, a subgraph with a hidden community and one or multiple stronger communities can be wrongly detected as a hierarchical structure. Fig. 1 (a) and (b) show this case. So we first discuss the difference between hidden structure and hierarchical structure to clarify the concepts. When considering the hierarchical community structure, an algorithm gradually detects stronger and smaller communities, and ignore weaker communities. When considering the hidden community structure, an algorithm detects communities weaker than the overhead communities. Intuitively, when the density of a community is calculated by modularity, if the average density of root communities is 0.5, then the average density of the sub-hierarchical level is larger than 0.5, and the average density of the next hidden layer is less than 0.5. Note that, a hidden community can be completely covered by one stronger community or partially covered by one or multiple stronger communities. And the hidden community does not need to be smaller than the overhead communities. Due to their different characteristics, each layer of hidden structure can contain several levels of hierarchical structure, and vice versa (Fig. 1 (c)). HirHide combines the advantages of both types of algorithms, and guarantees a complete protein complex detection.

HirHide

We propose a hierarchical hidden detection approach (HirHide) for community mining tasks. Our algorithm consists of three steps. In the first step of initialization, HirHide identifies a layer of communities *Layer1* via the base algorithm (Fig. 1 (c)), which can be a standard community detection algorithm with promising performance.

The second step is called the hierarchical detection step. HirHide constructs a hierarchical structure by recursively capturing sub-communities and iterating until an appropriate number of levels are found. A crucial dimension of this step is to determine the number of levels in a network. We simply count the number of nodes in each community. When the average amount is smaller than a certain threshold, the algorithm stops capturing sub-communities. In our experiments, the default threshold is set to 9.

The third step is called the hidden detection step. HirHide weakens the structure of the previously detected layer $Layer_1$ to get a reduced graph G' . In G' , the base algorithm is used to identify new communities to form a hidden layer $Layer_2$. Because G' does not contain the strong communities of $Layer_1$, the weaker communities can be easily discovered. The hierarchical detection step can be used on $Layer_2$ in the reduced graph G' to find hierarchical structures for this layer. Then HirHide weakens $Layer_1$ and $Layer_2$ and calls the base algorithm again to detect $Layer_3$. This process iterates until no communities can be detected.

In HirHide, the step of digging out hidden structure is in a similar spirit with HiCode [23]. But the step of combining hidden structure and hierarchical structure is new and well-designed to guarantee both hierarchical communities and hidden communities to be well captured. A key issue is which level of hierarchical communities should be weakened in the detected layer. When these hierarchical levels are organized into trees with each community in the first level as root, we can weaken the communities in roots or the communities in leaves (Fig. 1 (c)). Because the edge connection and the total size of nodes in each hierarchical level are similar (considering separate nodes and too small communities are removed) and the weakening step is a global operation, which level is weakened does not make a big difference. To double-check, we analyzed the performance of HirHide framework when it separately weakens the communities in roots or in leaves in the confirmatory experiments.

In the hierarchical step, the sub-communities are captured on the original graph. They can also be caught on the reduced graph G' after weakening the structure of strong communities. Intuitively, the reduced graph has weakened the influence of other layers, so the detected sub-communities should be more precise. In our confirmatory experiments, recursively capturing sub-communities on the original graph or the weakened graph are compared.

Although the base algorithm is recursively called in both the hierarchical and hidden steps, its role is significantly different. In the hierarchical step, the base algorithm is repeatedly applied in sub-communities to

discover smaller sub-communities. In the hidden step, the base algorithm is repeatedly used in the reduced graphs to dig out hidden communities.

Results

Algorithms, data, and metrics

Algorithms

We select three state-of-the-art algorithms as the baseline methods as well as the base algorithms of HirHide, which are MOD [12], Infomap [14], and LC [15]. None of them can detect hidden communities without HirHide. After they are combined with HirHide, their HirHide versions are called HirHide-MOD, HirHide-Infomap, and HirHide-LC. We also compare HirHide with the HiCode [23] version of the three methods, separately called HiCode-MOD, HiCode-Infomap, and HiCode-LC. One of our metrics comes from ClusterONE [26], but the authors did not provide the complete source code. As a consequence, we use ClusterONE as a baseline method without the corresponding HirHide version. Overall, the algorithms involved in the comparisons are HirHide-MOD, HirHide-Infomap, HirHide-LC with HiCode-MOD, HiCode-Infomap, HiCode-LC, MOD, Infomap, LC and ClusterONE.

Except HirHide and HiCode, the above comparison involves four other algorithms. MOD [12] treats each node as a separate community in initialization and gradually optimizes the modularity value by expanding each community's size. After repeatedly iterating, MOD can get a community division with the largest modularity value. Infomap [14] is based on the principle of information theory and defines the community from the perspective of coding. To get the maximum compression ratio, Infomap uses Huffman coding and community structure secondary coding. In this way, the problem of community detection is transformed into an optimization problem: finding a community division so that the codeword length of random code walks within and between communities is the smallest. LC [15] reinvents communities as groups of links rather than nodes. This approach successfully reconciles the antagonistic organizing principles of overlapping communities and hierarchical structure. Link communities naturally incorporate overlap while revealing hierarchical organization. ClusterONE [26] outlines the concept of cohesiveness score and uses a greedy growth process to find groups that are likely to correspond to protein complexes in a PPI network.

Data

We compare these algorithms in three large scale yeast PPI networks, which consist of core experimental yeast PPI network [27], a combined computational interaction network [28] and the entire set of

physical protein-protein interactions in yeast from BioGRID [29]. These three datasets are referred to as *Krogan_core*, *YeastNet*, and *BioGRID*.

To evaluate the performance of each algorithm, we use Munich Information Center of Protein Sequences (MIPS) [30] and CYC [31] as the reference sets. The yeast protein complexes cataloged by the MIPS database have been widely used to generate protein-protein interaction reference sets. And CYC is a comprehensive catalog of manually curated 408 heteromeric protein complexes in *S. cerevisiae*. For convincing, we choose the latest version of the two reference sets. Additionally, we only consider complexes containing 3 to 100 proteins as the reference protein complexes to avoid the selection bias.

Table 1 shows the basic information of the three PPI networks. Because different PPI networks contain different nodes and edges and the two reference sets have different reference complexes, we have removed proteins that only exist in PPI networks or reference datasets.

Evaluation Metrics

It has become a standard practice to compare the performance of different methods by assessing their ability to identify the reference communities. When eval-

Table 1 Basic information of the three PPI networks.

Dataset	MIPS		CYC	
	Nodes	Edges	Nodes	Edges
BioGRID	1155	10825	1355	13098
YeastNet	1085	11271	1289	13564
Krogan_core	679	1797	910	2808

uating the performance of community detection algorithms, the recognized metric is the F1 score [32], which is the harmonic mean of precision and recall. However, the reference communities, which only contain protein complexes whose interactions can be discovered under the current experimental conditions, have incomplete nature [33]. Using the F1 score as an evaluation metric is unreasonable when the reference sets are incomplete. Because under the same condition, the more communities an algorithm detects, the smaller the precision is. Traditional algorithms are typically designed to detect fewer complexes than the reference complexes (much less than the real-world complexes) to increase the precision. Consequently, comparing the F1 score is unfair for algorithms that detect more complexes.

As the current reference sets are incomplete [33], we evaluate the algorithms' performance by two other measures. One is the maximum matching ra-

Figure 3 The performance of different algorithms with the increasing hiddenness value of communities. The HirHide versions detect more accurately on hidden communities with higher hiddenness values.

tio (MMR), which is designed specifically for protein complex detection [26]. MMR guarantees that each detected community only matches one reference community and vice versa. And it maximizes the total score of all one-to-one connections between predicted and reference complexes. This measure is inspired by the bipartite graph maximum matching problem, in which the two sets of nodes respectively represent detected complexes and reference complexes. MMR tries to find the best match for each reference complex. So even if an algorithm detects more complexes, MMR won't be reduced.

The other is the recall, which only measures the capacity of discovering the reference complexes. So additional hidden hierarchical communities do not decrease the score. The recall scores each pair composed of a predicted complex and a reference complex by their similarity. Given a set of detected communities $\mathcal{D}$ and a set of reference communities $\mathcal{G}$. Each reference community G_j has its recall:

$$R(G_j) = \max_{D_i \in \mathcal{D}} \frac{|G_j \cap D_i|}{|G_j \cup D_i|}. \quad (4)$$

The final recall is defined as the average of $R(G_j)$ over all reference communities.

Experimental results on real-world data

We compare the performance of standard algorithms with their HirHide versions and their HiCode versions in real-world networks. Fig. 2 shows the comparative performance of the ten algorithms using MIPS and CYC separately as the reference sets. The different bar clusters represent using different datasets (*Biogrid*, *Krogan_core*, and *YeastNet*) as the PPI networks. The bars on the bottom indicate the maximum matching ratio (MMR) scores, and the bars on the top represent the recall scores. Higher bars represent better performance. In each bar, the specific score is shown if it is larger than 0.2. We can see that when benchmark algorithms are combined with HirHide, they achieve better MMR, recall, and composite scores on most

PPI datasets. As mentioned before, HirHide does not change the core of a standard algorithm but enables it to detect hidden hierarchical community structures. So the better performance demonstrates that detecting hidden hierarchical community structures helps detect complexes in PPI networks. Furthermore, the HirHide versions of the three benchmark algorithms perform better than their HiCode versions in all the datasets, which demonstrates that the hierarchical structure is necessary.

Fig. 3 illustrates the performance of different algorithms on communities with increasing hiddenness values. We show the results of HirHide-MOD versus MOD and HirHide-Infomap versus Infomap on *Biogrid*. The reference dataset is MIPS. A higher hiddenness value indicates a deeper hidden degree. Because there are hundreds of communities, we smooth the results to make them more concise. As illustrated in Fig. 3, HirHide-MOD and HirHide-Infomap show increasing advantage with the higher hiddenness value, indicating that HirHide has significant benefits in detecting the hidden communities.

Synthetic data and experimental results

To some extent, the two evaluation metrics can solve the problem that the reference sets are incomplete, but they may not be persuasive enough. As a result, the comparison of synthetic networks is necessary. We build three synthetic networks, each of which contains two layers for hidden structure. The first layer consists of strong communities, and the second layer consists of relatively weak communities. In addition to the multi-layer feature, hierarchical structures are added. Specifically, in each layer, smaller and denser communities are added to make up the next hierarchical level.

The choice of edge probability follows a certain rule. According to the hidden community conception, the interactions between proteins in the hidden layer are relatively sparse, so the edge probability is lower than that of the first layer. When constructing the hierarchical structure, the edge probability of the leaf communities at the second level is slightly higher than that

Figure 4 Constructing hidden hierarchical community structure. When constructing a community, it is treated as a random block model and edges are added within the community. (a) Constructing communities in root of the first layer. (b) Constructing communities in leaves of the first layer. (c) The graph after randomly scrambling the node number. (d) Constructing communities in root of the second layer. (e) Constructing communities in leaves of the second layer.

Table 2 The performance of HirHide-MOD (H-MOD) and MOD on synthetic data.

Truth	G_{360}		G_{2000}		G_{3000}	
	H-MOD	MOD	H-MOD	MOD	H-MOD	MOD
L_{11}	1.000	1.000	0.999	0.980	1.000	0.999
L_{12}	0.997	0.857	0.686	0.681	0.859	0.690
L_{21}	1.000	0.337	1.000	0.100	1.000	0.079
L_{22}	0.606	0.215	0.474	0.109	0.721	0.080

of the root communities, to further highlight the structure of sub-communities. The first network called G_{360} contains 360 nodes. At the first layer, 360 nodes are divided into three communities, each with 120 nodes. According to the definition of the community, internal edges are added to the three communities with probability $P_{11} = 0.2$, and no edge is added outside (Fig. 4 (a)). Then the hierarchical structure is appended. The 120 nodes of each community are subdivided into two sub-communities. Six sub-communities make up the second hierarchical level, each of which contains 60 nodes. Edges are added to the inside of the six sub-communities with probability $P_{12} = 0.3$ (Fig. 4 (b)). Then the first layer is successfully built.

According to the idea of hidden structure, before building the second layer, the node numbers in the first layer are randomly scrambled so that the previous community structure is evenly distributed in the adjacency matrix (Fig. 4 (c)).

The second layer is constructed on the randomly scrambled graph. In the second layer, 360 nodes are divided into four communities, each of which has 90 nodes. Internal edges are added in these communities with a probability of $P_{21} = 0.15$ (Fig. 4 (d)). Then we add a hierarchical structure, and each community is divided into three small sub-communities, each of which has 30 nodes. Edges are added inside these sub-communities with a probability of $P_{22} = 0.25$ (Fig. 4 (e)).

Similarly, we also construct a network with 2000 nodes and a network with 3000 nodes, called G_{2000} and G_{3000} .

Because MOD has the best performance among the three benchmark algorithms, We compare HirHide-MOD with MOD on our synthetic data. Table 2 shows the experimental results. L_{11} , L_{12} represent the dominant community layer and L_{21} , L_{22} represent the hidden community layer. As illustrated in Table 2, HirHide-MOD has a slight advantage over MOD on L_{11} and L_{12} , which means HirHide doesn't reduce the performance of traditional algorithms in networks without hidden structure. Moreover, HirHide-MOD has significant advantage over MOD on L_{21} and L_{22} . The scores of MOD on L_{21} and L_{22} are extremely low while the scores of HirHide-MOD on L_{21} and L_{22} are normal and high. Consequently, we can conclude that MOD can not detect the hidden community layer. But after it is combined with the HirHide framework, HirHide-MOD can detect the hidden community layer well. These results are consistent with the results in real-world networks.

Confirmatory experiments

In the HirHide framework, the strong community structure detected by the base algorithm needs to be weakened. After combining the concept of hierarchical structure, there is a choice of weakening communities at the top level of the hierarchical structure or weakening communities at the bottom level of the hierarchical structure. To determine which level we should choose, we experiment on the synthetic network G_{3000} . Because the synthetic data is complete, we have chosen the F1 score as the evaluation metric. As illustrated in Table 3, weakening communities in roots can make HirHide have a better performance, especially in level 2.

Figure 5 Comparison of hidden complexes and dominant complexes. We sort each set of complexes in ascending order of ancestor evidence.

In HirHide, there are two options of grabbing sub-communities in a layer: to grab sub-communities on the original graph or to grab sub-communities on the reduced graph after weakening other layers. We still use the F1 score to evaluate the performance. As illustrated in Table 4, in graph G_{3000} , grabbing the sub-communities on the original graph has a slight advantage. But in graph G_{2000} , we have the opposite result. Consequently, overall which graph is chosen does not make a big difference.

Table 3 Results of weakening communities in roots or leaves. *Weakening Roots* means that weakening communities in roots and *Weakening Leaves* means that weakening communities in leaves.

Truth	Weakening Roots	Weakening Leaves
L_{11}	1.000	0.913
L_{12}	0.997	0.691
L_{21}	1.000	0.911
L_{22}	0.606	0.485

Table 4 Results of grabbing sub-communities on original graph or weakened graph.

Graph	L_{12}		L_{22}	
	Original	Weakened	Original	Weakened
G_{2000}	0.686	0.707	0.474	0.479
G_{3000}	0.859	0.826	0.721	0.717

Biological significance of hidden complexes

To explore the biological relevance of hidden communities, we need to compare two sets of complexes. One set is composed of complexes with the highest hiddenness values, and the other set contains complexes with the lowest hiddenness values. The former set contains all complexes with at least 25 percent of proteins also

belonging to other stronger complexes. The latter set includes no protein that belongs to other stronger complexes. As all algorithms perform better in CYC than in MIPS in our experiment, we filter the sets from reference set CYC. We don't divide all complexes of CYC into two sets, because there is no clear boundary between the hidden complexes and the dominant complexes. We calculate the hiddenness values of the two sets of complexes based on three graphs mentioned in subsection Data and select all complexes with the highest hiddenness values or the lowest hiddenness values from the three graphs.

We compare the information of the two sets about gene ontology in AmiGO 2 [34], and find a significant difference in the biological aspect of ancestor evidence used in manual assertion between the two sets. The hidden complexes typically have much more ancestor evidence (about 2.5 times) than the dominant complexes. Fig. 5 visualizes the difference. Approximately 82% of the hidden complexes have more than 480 ancestor evidence. For the dominant complexes, this fraction is around 40% (See details for each complex in the Additional file 1: Table S1). Although the number of evidence can be affected by other factors like the importance of a complex, such a significant difference indicates consistency between the hiddenness value and the number of ancestor evidence. Therefore, we conjecture that complexes with high hiddenness values have more ancestors or have a deeper relationship with their ancestors. Traditional methods determine hierarchy by the relationship of containing or be contained, which can not work in the hierarchy of gene ontology because of the complicated relationship among the gene

products. We find that determining hierarchy by hiddenness value is more consistent within gene ontology. So the hidden structure found by HirHide is a useful supplement to the gene ontology. In case the ancestor evidence is inconsistent with the hiddenness value for a gene product, it is worth further exploration.

Discussion

A HirHide-combined algorithm can detect additional hidden hierarchical communities that can not be detected without HirHide. Some of them fit the characteristics of protein complexes with sparse internal connections. So they can serve as predicted protein complexes. Here, the predicted protein complexes mean these protein complexes do not appear in the reference sets, and they have the potential to be protein complexes.

However, detecting these hidden hierarchical communities are solely based on density property in PPI networks. Not all of them are reliable enough, which is a main weakness of HirHide. Emerging patterns (EPs) are conjunctive patterns that contrast sharply between different classes of data, which contain more informative properties such as degree statistics, clustering coefficient, topological coefficients and eigenvalues of a sub-graph. Recently, EPs are exploited to address the complex prediction problem [21]. In EPs, a feature vector is first constructed to describe the critical properties of the reference protein complexes as well as those of random non-complexes communities. Then to discover EPs by contrasting feature vectors of reference protein complexes and random non-complexes communities. Next, the discovered EPs are used to find out potential complexes. We combine the results of the two methods to screen more reliable predicted protein complexes. For each of the complexes predicted by HirHide, if it is quite similar to complexes discovered by EPs, it is considered a more reliable predicted complex. Examples of the predicted complexes are illustrated in the Additional file 1: Figure S1. Note that the communities detected by HirHide are based on mathematical prediction. More reliable results need to be conformed by biological experiments.

Conclusions

The main contribution of this work is to propose a new method called HirHide and apply HirHide for protein complex detection and prediction. HirHide serves as a meta-method that can be combined with existing standard community detection algorithms and enable them to discover hidden hierarchical communities. We improve the definition of the necessity of detecting hidden structures and explain why detecting hidden hierarchical structures is more helpful for complex

detection in PPI networks. We propose the concept of hidden complexes. Besides, we describe the biological significance of hidden complexes. Also, we have compared standard detection algorithms with their HirHide versions in synthetic data and real-world PPI networks. Experimental results illustrate that the performance of standard algorithms is boosted when combined with HirHide. And their HirHide versions have better performance than their HiCode versions. In networks without hidden structure, HirHide doesn't reduce their performance.

How to combine the concept of hierarchical structure with hidden structure is a critical issue in our work. We verify which level of the hierarchical structure is selected when iterating in HirHide and conclude that weakening communities at the root level can achieve the best weakening effectiveness. There is almost no difference to grab sub-communities on the original graph or on the reduced graph. Finally, we select some new and reliable hidden hierarchical communities as preliminary predictions of undiscovered protein complexes.

In this work, we applied HirHide to detect protein complexes in PPI networks. HirHide can be used for most community detection questions because of the flexibility of the base algorithms. HirHide is also suitable for issues related to other large-scale biological data such as neural networks and gene regulatory networks. Some of these biological networks are too complex for standard community detection algorithms to produce a positive performance. Most of the time, a better choice is to choose a corresponding algorithm based on the characteristics of a specific biological network. Under these circumstances, HirHide has clear advantages because of its flexibility. At the same time, according to the natural hierarchy of cells, organelle, intracellular compound *etc.*, hierarchical structure with hidden structure is in line with the characteristics of the data itself, thus helping researchers to study biological interactions more deeply.

Additional Files

Additional file 1 — Supplementary materials
(Supplementary Tables S1, Supplementary Figure S1).(PDF 384 kb)

Abbreviations

HirHide: Hierarchical hidden community detection;PPI: Protein-protein interaction;HiCode: Hidden community detection;MMR: Maximum matching ratio;MIPS: Munich information center of protein sequences;EPs: Emerging patterns.

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Competing interests

The authors declare that they have no competing interests.

Funding

No Funding.

Availability of data and materials

The three benchmark datasets supporting the conclusions of this article are available online:

BioGRID: <https://thebiogrid.org/> [29],

YeastNet: <https://www.inetbio.org/yeastnet/> [28],

Krogan_core: <https://www.yeastgenome.org/> [27].

They can also be download on <https://github.com/JHL-HUST/HirHide> in a convenient manner. The data of ancestor evidence (search for all the protein complexes one by one):

http://amigo.geneontology.org/amigo/dd_browse [34].

Author's contributions

KH and CL designed the study. CL and KH wrote the manuscript. CL, GSL and KH planned, designed, and conducted the experiment to compare HirHide with other state-of-the-art algorithms and contributed to its analysis. All authors revised and approved the final manuscript.

Acknowledgements

Not applicable.

Author details

¹School of Computer Science and Technology, Huazhong University of Science and Technology, 430074 WuHan, China. ²Department of Computer Science, Cornell University, 14853 Ithaca NY, U.S.A..

References

- Fiannaca, A., La Rosa, M., Urso, A., Rizzo, R., Gaglio, S.: A knowledge-based decision support system in bioinformatics: an application to protein complex extraction. *BMC bioinformatics* **14**(1), 5 (2013)
- Schwikowski, B., Uetz, P., Fields, S.: A network of protein–protein interactions in yeast. *Nature biotechnology* **18**(12), 1257 (2000)
- King, A.D., Pržulj, N., Jurisica, I.: Protein complex prediction via cost-based clustering. *Bioinformatics* **20**(17), 3013–3020 (2004)
- Xie, Z., Kwok, C.K., Li, X.-L., Wu, M.: Construction of co-complex score matrix for protein complex prediction from ap-ms data. *Bioinformatics* **27**(13), 159–166 (2011)
- Winterhalter, C., Nicolle, R., Louis, A., To, C., Radvanyi, F., Elati, M.: Pepper: cytoscape app for protein complex expansion using protein–protein interaction networks. *Bioinformatics* **30**(23), 3419–3420 (2014)
- Lage, K., Karlberg, E.O., Størting, Z.M., Olason, P.I., Pedersen, A.G., Riggins, O., Hinsby, A.M., Tümer, Z., Pociot, F., Tommerup, N., et al.: A human phenome-interactome network of protein complexes implicated in genetic disorders. *Nature biotechnology* **25**(3), 309 (2007)
- Yang, P., Li, X., Wu, M., Kwok, C.-K., Ng, S.-K.: Inferring gene-phenotype associations via global protein complex network propagation. *PLoS one* **6**(7), 21502 (2011)
- Fraser, H.B., Plotkin, J.B.: Using protein complexes to predict phenotypic effects of gene mutation. *Genome biology* **8**(11), 252 (2007)
- Yu, L., Huang, J., Ma, Z., Zhang, J., Zou, Y., Gao, L.: Inferring drug-disease associations based on known protein complexes. *BMC medical genomics* **8**(2), 2 (2015)
- Girvan, M., Newman, M.E.: Community structure in social and biological networks. *Proceedings of the national academy of sciences* **99**(12), 7821–7826 (2002)
- Newman, M.E.: The structure and function of complex networks. *SIAM review* **45**(2), 167–256 (2003)
- Blondel, V.D., Guillaume, J.-L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment* **2008**(10), 10008 (2008)
- Pons, P., Latapy, M.: Computing communities in large networks using random walks. In: *International Symposium on Computer and Information Sciences*, pp. 284–293 (2005). Springer
- Rosvall, M., Bergstrom, C.T.: Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences* **105**(4), 1118–1123 (2008)
- Ahn, Y.-Y., Bagrow, J.P., Lehmann, S.: Link communities reveal multiscale complexity in networks. *nature* **466**(7307), 761 (2010)
- Coscia, M., Rossetti, G., Giannotti, F., Pedreschi, D.: Demon: a local-first discovery method for overlapping communities. In: *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 615–623 (2012). ACM
- Lancichinetti, A., Radicchi, F., Ramasco, J.J., Fortunato, S.: Finding statistically significant communities in networks. *PLoS one* **6**(4), 18961 (2011)
- Yang, J., Leskovec, J.: Community-affiliation graph model for overlapping network community detection. In: *2012 IEEE 12th International Conference on Data Mining*, pp. 1170–1175 (2012). IEEE
- Gregory, S.: Finding overlapping communities using disjoint community detection algorithms. In: *Complex Networks*, pp. 47–61. Springer, ??? (2009)
- Zhang, S., Wang, R.-S., Zhang, X.-S.: Identification of overlapping community structure in complex networks using fuzzy c-means clustering. *Physica A: Statistical Mechanics and its Applications* **374**(1), 483–490 (2007)
- Liu, Q., Song, J., Li, J.: Using contrast patterns between true complexes and random subgraphs in ppi networks to predict unknown protein complexes. *Scientific reports* **6**, 21223 (2016)
- Stumpf, M.P., Wiuf, C., May, R.M.: Subnets of scale-free networks are not scale-free: sampling properties of networks. *Proceedings of the National Academy of Sciences* **102**(12), 4221–4224 (2005)
- He, K., Li, Y., Soundarajan, S., Hopcroft, J.E.: Hidden community detection in social networks. *Information Sciences* **425**, 92–106 (2018)
- He, K., Soundarajan, S., Cao, X., Hopcroft, J., Huang, M.: Revealing multiple layers of hidden community structure in networks. *arXiv preprint arXiv:1501.05700* (2015)
- He, K., Soundarajan, S., Cao, X., Hopcroft, J.E., Huang, M.: Revealing multiple layers of hidden community structure in networks. *CoRR abs/1501.05700* (2015)
- Nepusz, T., Yu, H., Paccanaro, A.: Detecting overlapping protein complexes in protein–protein interaction networks. *Nature methods* **9**(5), 471 (2012)
- Krogan, N.J., Cagney, G., Yu, H., Zhong, G., Guo, X., Ignatchenko, A., Li, J., Pu, S., Datta, N., Tikuisis, A.P., et al.: Global landscape of protein complexes in the yeast *saccharomyces cerevisiae*. *Nature* **440**(7084), 637 (2006)
- Kim, H., Shin, J., Kim, E., Kim, H., Hwang, S., Shim, J.E., Lee, I.: Yeastnet v3: a public database of data-specific and integrated functional gene networks for *saccharomyces cerevisiae*. *Nucleic acids research* **42**(D1), 731–736 (2013)
- Stark, C., Breitkreutz, B.-J., Reguly, T., Boucher, L., Breitkreutz, A., Tyers, M.: Biogrid: a general repository for interaction datasets. *Nucleic acids research* **34**(suppl_1), 535–539 (2006)
- Mewes, H.-W., Amid, C., Arnold, R., Frishman, D., Güldener, U., Mannhaupt, G., Münsterkötter, M., Pagel, P., Strack, N., Stümpflen, V., et al.: Mips: analysis and annotation of proteins from whole genomes. *Nucleic acids research* **32**(suppl_1), 41–44 (2004)
- Pu, S., Wong, J., Turner, B., Cho, E., Wodak, S.J.: Up-to-date catalogues of yeast protein complexes. *Nucleic acids research* **37**(3), 825–831 (2008)
- Chase Lipton, Z., Elkan, C., Narayanaswamy, B.: Thresholding classifiers to maximize f1 score. *arXiv preprint arXiv:1402.1892* (2014)
- Li, X., Wu, M., Kwok, C.-K., Ng, S.-K.: Computational approaches for detecting protein complexes from protein interaction networks: a survey. *BMC genomics* **11**(1), 3 (2010)
- Carbon, S., Ireland, A., Mungall, C.J., Shu, S., Marshall, B., Lewis, S., Hub, A., Group, W.P.W.: Amigo: online access to ontology and annotation data. *Bioinformatics* **25**(2), 288–289 (2009)

Figures

Figure 1

Example to show the hidden hierarchical structure. (a) A network with hidden hierarchical structure. The graph exists of four large dense communities consisting of 128 nodes, each with an internal subdivision of four small communities with 32 nodes. Additional, there is a sparse, hidden community structure consisted by the four small blocks in the center. (b) The network after weakening the stronger communities. The community with green nodes is the hidden community, detected by weakening the stronger communities. So it doesn't contain dense edges. (c) Schematic diagram of a hidden hierarchical structure. Level and layer are separately used to denote the stratification of hierarchical structure and hidden structure. Each level of hierarchical structure can contain several layers of hidden structure, and vice versa.

Figure 2

Benchmark results. The different bar clusters represent results for different datasets (Biogrid, Krogan core and YeastNet) as the PPI networks. The bars on bottom indicate the scores of maximum matching ratio (MMR) and the bars on top represent the scores of the recall. The number inside each bar is the score and higher bars represent better performance.

Figure 3

The performance of different algorithms with the increasing hiddenness value of communities. The HirHide versions detect more accurately on hidden communities with higher hiddenness values.

Figure 4

Constructing hidden hierarchical community structure. When constructing a community, it is treated as a random block model and edges are added within the community. (a) Constructing communities in root of the first layer. (b) Constructing communities in leaves of the first layer. (c) The graph after randomly scrambling the node number. (d) Constructing communities in root of the second layer. (e) Constructing communities in leaves of the second layer.

Figure 5

Comparison of hidden complexes and dominant complexes. We sort each set of complexes in ascending order of ancestor evidence.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [Additionalfile1.pdf](#)