

Revealing Biological Modules via Graph Summarization

SAKET NAVLAKHA,¹ MICHAEL C. SCHATZ,¹ and CARL KINGSFORD^{1,2}

ABSTRACT

The division of a protein interaction network into biologically meaningful modules can aid with automated detection of protein complexes and prediction of biological processes and can uncover the global organization of the cell. We propose the use of a graph summarization (GS) technique, based on graph compression, to cluster protein interaction graphs into biologically relevant modules. The method is motivated by defining a biological module as a set of proteins that have similar sets of interaction partners. We show this definition, put into practice by a GS algorithm, reveals modules that are more biologically enriched than those found by other methods. We also apply GS to predict complex memberships, biological processes, and co-complexed pairs and show that in most settings GS is preferable over existing methods of protein interaction graph clustering.

Key words: function prediction, graph summarization, module detection, protein interaction networks.

1. INTRODUCTION

HIGH-THROUGHPUT EXPERIMENTAL PROTOCOLS have provided a noisy, incomplete picture of the network of interactions between proteins in the cells of many organisms (Shoemaker and Panchenko, 2007; Zhu et al., 2007). These putative protein associations, determined largely by yeast two-hybrid (Fields and Song, 1989) and co-immunoprecipitation (Gavin et al., 2006; Krogan et al., 2006), need to be analyzed to uncover new biology. It is a central computational problem to derive from these networks the sets of proteins that form complexes or are involved in the same biological processes. A common approach to this task is to partition the interaction graph into modules—subsets of proteins—based on the connectivity pattern of the nodes, exploiting the fact that, despite the large amount of noise, proteins with related processes tend to be situated near one another within the network (Bu et al., 2003; Guimera and Luis, 2005; Han et al., 2004; Hartwell et al., 1999; Rives and Galitski, 2003). These modules can be mined to uncover the physical, logical, and evolutionary organization of the cell.

¹Department of Computer Science and Center for Bioinformatics and Computational Biology, University of Maryland, College Park, Maryland.

²Institute for Advanced Computer Studies, University of Maryland, College Park, Maryland

Despite intense interest (Clauset et al., 2004; Kumar et al., 1999; Radicchi et al., 2004; Rosvall and Bergstrom, 2007; Sharan et al., 2007; Spirin and Mirny, 2003; Wakita and Tsurumi, 2007; Wang et al., 2007; Ziv et al., 2005), the proper definition of a network module has remained elusive. It is likely that the appropriate definition must be motivated by the application domain on a case-by-case basis. Complexes are sets of proteins which interact stably over time, and often appear in networks as dense subgraphs. In the context of uncovering complexes and proteins involved in similar biological processes within protein interaction networks, a natural definition of a module is a set of proteins that all have a similar set of interaction partners. If two proteins interact with similar partners, they likely have related cellular roles. For example, every pair of members in an isolated clique representing a stable complex interact with nearly exactly the same set of proteins: the other members of the clique. Conversely, unrelated proteins are more likely to be widely separated in the network and to thus share few, if any, common interaction partners.

Motivated by this definition of a module, we investigate the application of a technique called *graph summarization* (GS) (Navlakha et al., 2008) to the problem of biological module detection. In general, GS seeks to produce a minimal cost representation of an input graph, compressing it according to the well-known minimal description length (MDL) (Rissanen, 1978) ideology. GS has been used to reduce the cost of large networks by 50% or more without loss of information (Navlakha et al., 2008). It achieves such savings by representing nodes with common edge patterns as *supernodes* connected by *superedges* in a summary graph, in conjunction with a *correction list* of topologically inconsistent or missing edges. Graph compression schemes have been previously suggested for community detection (Rosvall and Bergstrom, 2007; Ziv et al., 2005), although this is their first application in the context of biological networks. In our context, supernodes are composed of proteins with highly similar patterns of interactions. These supernodes are thus natural candidates for biologically meaningful modules.

Detection of modules and complexes within biological networks has received much attention recently, with several widely used algorithms having been developed. The molecular complex detection algorithm (MCODE) (Bader and Hogue, 2003) is a graph-theoretic clustering algorithm specifically designed to find complexes by identifying densely connected subgraphs in networks. The Markov clustering algorithm (MCL) (van Dongen, 1998) is based on simulating flow expansion and flow contraction on graphs, and has been applied (Enright et al., 2002) to the detection of protein families. In a recent comparison of graph clustering algorithms (Brohee and van Helden, 2006) MCL was also shown to be the most robust at recovering MIPS (Guldener et al., 2005) complexes embedded within a simulated interaction graph with noise, outperforming restricted neighborhood search clustering (King et al., 2004) and superparamagnetic clustering (Spirin and Mirny, 2003). Recently, Newman's (2006) spectral algorithm has become a popular choice for community detection, especially within social networks. It is based on modularity (Newman and Girvan, 2004), a measure of the number of edges falling within modules minus the expected number in a random network, though there is evidence that this definition prefers modules of a certain characteristic size (Fortunato and Barthelemy, 2007). Other graph clustering algorithms based on highly connected subgraphs (Przulj et al., 2004) have also been applied to biological networks. All of these algorithms use the graph topology only, an approach we follow here as well. This allows us to assess and improve the ability to extract biological information using only interaction data. Any improved network-based analysis can subsequently be incorporated into a more comprehensive, integrative system (Jansen et al., 2003; Qiu and Noble, 2008; Tanay et al., 2004; Troyanskaya et al., 2003).

We show that our GS approach can recover both stable protein complexes and broader modules enriched for biological processes from a protein interaction network of *Saccharomyces cerevisiae*. Using several evaluation metrics, we show that GS outperforms other standard complex- and module-detection methods, such as MCL, MCODE, and Newman's spectral partitioning algorithm. In particular, a higher percentage of MIPS (Guldener et al., 2005) complexes and interesting Gene Ontology annotations (Ashburner et al., 2000; Myers et al., 2006) are represented in a statistically significant manner by the modules constructed by GS than those constructed by the other methods.

The modules identified by GS are also more useful for annotating proteins with unknown complex membership and biological processes. We evaluate several different schemes for labeling unannotated proteins within a module, based on transfer of the majority, plurality, or statistically enriched annotations. With nearly every approach for both complexes and processes, GS has the highest *F*-score, which demonstrates its ability to make precise predictions covering large portions of the proteome. Further,

GS generally predicts annotations for the largest number of proteins under leave-one-out cross-validation testing as compared with existing methods.

These improvements over MCL, MCODE, and Newman's algorithm hold for both an unfiltered yeast network derived from IntAct (Kerrien et al., 2007) and generally for a smaller, higher-confidence network constructed by eliminating edges from the network that have weaker support in the literature. Further, while our GS approach works well overall, the predictions made are not a strict superset of those found by MCL or MCODE, and a large fraction of predictions are made only by one method. Thus a combination of multiple methods may be useful to maximize coverage.

Finally, in addition to identifying modules, the corrections list produced by the GS algorithm can also be analyzed to predict co-complexed pairs by identifying missing and false edges in the network. This use of GS generalizes the popular method of predicting edges to complete defective cliques (DCC) (Yu et al., 2006). Our testing shows GS is more precise than DCC over a wide range of parameters for DCC. In addition, unlike DCC which only predicts missing edges, GS accurately predicts edges to remove from the network as well.

2. METHODS

2.1. Graph summarization

The goal of graph summarization (GS) (Navlakha et al., 2008) is to produce a compressed representation of the input network with minimal cost. Given an input graph $G = (V_G, E_G)$, the summarizer produces a new graph $H = (V_H, E_H)$, an onto mapping f from V_G to V_H , and a list of *corrections*. Each node in H is called a *supernode* and is composed of one or more original nodes from V_G collapsed into a single node. Edges between supernodes u and v in E_H imply an edge between all pairs of original nodes contained in u and v . In other words, superedge $\{u, v\}$ implies all edges $f^{-1}(u) \times f^{-1}(v)$ exist in the original graph.

To account for errors that would be introduced by this superedge expansion rule, the summarizer also produces a list \mathcal{C} of corrections. This list augments the superedges with edges from the original network that must be added or subtracted to fully reconstruct the original network. Edges that are implied by a superedge but that are missing from E_G are recorded as edges to subtract (negative corrections); on the other hand, atypical edges that are not covered by any superedge are recorded as edges that must be added (positive corrections). Together the supernodes and superedges effectively summarize the original network and reveal the major topological structure, while the corrections list reveals the exceptions to that overall structure. The original graph G can then be exactly reconstructed from H by expanding each supernode into its constituent nodes, adding all edges implied by the superedges, and then applying the positive and negative corrections from the corrections list.

In accordance with the MDL principle, GS attempts to find the representation H, \mathcal{C} of G that minimizes the number of superedges plus the number of corrections, $|E_H| + |\mathcal{C}|$. For example, if G is composed of two nearly complete cliques of size m and n , each missing one edge, and connected by a single edge, then G will be compressed into H composed of two supernodes with self-edges, and a corrections list containing the two missing edges from the cliques (negative corrections) and the individual edge between the cliques (positive correction). The cost of the summary is the number of superedges (2) plus the number of edges on the correction list (3) giving a dramatic saving over the original cost $\binom{m}{2} + \binom{n}{2} - 1$. For a detailed, step-by-step example, we refer the reader to Navlakha et al. (2008).

Here, a compression of the protein interaction network is found using the greedy graph summarization method (Navlakha et al., 2008). This method iteratively merges the pair of nodes that results in the greatest cost reduction (generally the pair with the most similar set of edges), much like bottom-up hierarchical clustering (Brun et al., 2003; Rives and Galitski, 2003; Samanta and Liang, 2003). Graph summarization, however, seeks to minimize a global cost based on description length, and the greedy algorithm naturally terminates when merging any pair of nodes increases the cost of H . In the context of finding modules within protein networks, we add a self-edge to each node before compression so that every member of a clique will have exactly the same neighbors (otherwise, the neighbor sets of any two interacting proteins will be different). Once the compressed graph is found, the supernodes of proteins with similar interaction partners are taken to represent biological modules.

2.2. Interaction network

We constructed a protein interaction network called Y_{ppi} for the yeast *S. cerevisiae* from data deposited in the IntAct (Kerrien et al., 2007) database. The Y_{ppi} network includes all deposited edges and contains 5,492 proteins with 40,332 interactions. Most of these interactions were derived from yeast-two-hybrid and TAP experiments, while a smaller number were obtained through traditional low-throughput assays.

2.3. Complex and biological process annotations

We assess the biological quality of the modules found by GS by their ability to recapitulate known protein complexes and biological processes. We use known annotations for assessment and prediction only—no annotations were used when constructing the modules.

Yeast complex data was taken from the MIPS (Guldener et al., 2005) complex catalog. This set of complexes has been widely used to assess computational methods (Jansen et al., 2003; Qiu and Noble, 2008; Yu et al., 2006). We ignore complexes from the “550” section of the MIPS tree, which represent computationally inferred annotations. We are interested in making predictions which are as specific as possible, and we therefore use the set of complexes at the leaves of the tree. The leaf set contains 267 complexes, of which 266 are represented by at least one protein in the Y_{ppi} network.

We obtained known biological process annotations from the Saccharomyces Genome Database (SGD) (Cherry et al., 1997) corresponding to the biological process sub-ontology of the Gene Ontology (GO) database (Ashburner et al., 2000). We used the ancestor relationships of GO to determine which proteins participate in each biological process. To ensure that we are detecting processes at an interesting level of specificity, we focus on a comprehensive, expert-curated subset of biologically interesting annotations selected by Myers et al. (2006). All analysis of biological processes in this paper is done using this set of 295 terms, 182 of which are represented in the Y_{ppi} network by at least one protein.

2.4. Measuring enrichment of biological processes and complex membership

The goal of module-detection algorithms is to discover modules from an interaction network that are “enriched.” In other words, proteins in a module should all be part of the same complex or biological process. We measure the enrichment of a given annotation F in a given module M with the hypergeometric p -value, computed by $\sum_{i=k}^m \frac{\binom{f}{i} \binom{n-f}{m-i}}{\binom{n}{m}}$, where n is the number of nodes in the network, m is the number of nodes in M , f is the number of nodes in the network annotated with F , and k is the number of nodes in M annotated with F . The computed hypergeometric p -values are Bonferroni-corrected to account for multiple testing. An annotation with a p -value of less than 0.001 was considered enriched. We will refer to the hypergeometric with this p -value as “hypergeometric.”

We use two measures to assess the quality of the modules. The first is the percentage of complexes or biological processes that are enriched in at least one module. This measures the diversity of annotations present in the modules, with large values indicating that a wide spectrum of biology is represented. Small values, in contrast, suggest the modules recapitulate only a few biological annotations (such as the ribosome). Conversely, we measure the percentage of the reported modules enriched for at least one annotation, as used previously in Ulitsky and Shamir (2007).

2.5. Predicting new annotations

Another test of the quality of a module decomposition is how well it can be used to make new predictions for membership within complexes or biological processes. Given a decomposition of proteins into sets of modules, we employ several module-assisted prediction techniques to infer new annotations for proteins of unknown complexes or biological processes. Each technique is based on transferring an annotation that is common to many proteins in a module to every protein in the module.

The first approach, “majority,” transfers an annotation to a protein if more than 50% of the other annotated proteins in the module have that annotation. If no annotation exists on more than 50% of the proteins (or if there exists only one annotated protein in the module), no predictions are made. Relaxing the requirement for a strict majority leads to the second method of annotation transfer, here called “plurality.” Under this scheme, a protein is predicted to have the most common annotation within its module. The third approach transfers all annotations that are statistically enriched within a module to every protein in the module.

An annotation with a Bonferroni-corrected hypergeometric p -value of <0.001 was considered enriched. In all cases, modules consisting of a single protein are ignored. These schemes are applied separately for complex membership and biological process annotations.

We tested the efficacy of these schemes for each module detection algorithm using leave-one-out cross-validation. For every annotated protein p , all of its annotations were forgotten, the majority, plurality, or enriched annotations were computed for its module and then transferred to p as predicted annotations. If multiple annotations were transferred, each transferred annotation was considered one prediction. A prediction was correct if the protein is known to belong to that complex or biological process, and incorrect if it is only known to belong to other complexes or processes. Naturally, given the incomplete state of knowledge, some “incorrect” predictions may in fact represent correct, new biology.

We measure performance by the precision and recall of the predictions made. Precision is the percent of predicted annotations that are correct. Recall is the number of correct annotations made divided by the total number of possible correct annotations. There are a large number of possible annotations over all the proteins in the network, which generally leads to low recall for all methods. Even relatively low recall values, however, represent hundreds of correct annotations inferred using only interaction topology. We also report the widely used F -score to evaluate a method’s balance between precision and recall. The F -score is a number between 0 and 1 (the larger the better) and is computed as the harmonic mean of the precision and recall: $2 \cdot \text{precision} \cdot \text{recall} / (\text{precision} + \text{recall})$.

2.6. Clustering parameter variation

We compare the results of GS with those obtained by MCL, MCODE, and the spectral partitioning algorithm of Newman (which we label here as NSP). Both MCL and MCODE make use of parameter values that can affect the clustering they produce. MCL uses a single parameter I that indirectly controls the sizes of the obtained modules. Larger values of I favor smaller modules. We tried 9 different values of I between 1.8 and 4.6, and chose $I = 3.8$ because this value maximized the precision of the predictions made by MCL on Y_{ppi} using both the majority and plurality rules. We also report the predictive performance using the suggested default value of $I = 2.0$, which was considerably worse. A recent survey (Brohee and van Helden, 2006) found the value of 1.8 for this parameter most effective, but with our data set, this value produced clusters that were too large, and did not perform as well as $I = 3.8$. The performance of MCL presented below should thus be considered *trained*.

MCODE supports several parameters (degree cutoff, node score cutoff, haircut, fluff, k-core, max depth). Experiments were run using both the default parameters as well as the parameters suggested in Brohee and van Helden (2006). The latter set of parameters was used because it produced clusters with the greatest predictive precision using both the majority and plurality rules (although less than 6% of the proteins were clustered). The “node score cutoff” had the greatest effect on the modules. It influences the cluster size and was set to 0.0 as in Brohee and van Helden (2006) to favor small clusters. Again, because parameters were selected based on their performance on the test set, the results for MCODE should also be considered *trained*. While an exhaustive search of parameter space may reveal a set of parameters for which performance is improved, it is unclear in practice how these parameters should be set without fitting to a training set.

GS and NSP require no parameters to be set.

3. RESULTS AND DISCUSSION

3.1. Application of graph clustering techniques

Each of the MCODE, MCL, NSP, and GS methods were run on the unweighted Y_{ppi} network. GS, MCL, and NSP give a complete partitioning of the graph; MCODE, however, did not place all proteins in a module. In all cases, modules that contained only a single protein were discarded. All numbers presented in this section are for the tuned versions of MCL and MCODE. The module decompositions were very different amongst the four algorithms.

NSP divides the network into only 8 modules, consistent with the previous observation that the modularity statistic does not easily find small modules within larger graphs (Fortunato and Barthelemy, 2007). While these large modules produced by NSP do yield some biological information, they are generally difficult to

interpret and do not likely correspond to natural biological divisions of the graph. MCODE also produced far fewer modules than either MCL or GS, creating 80 modules covering only 308 proteins (less than 6% of the total network). MCL covered 4383 proteins with 1185 modules. GS produced the largest number of modules: 1632 modules that covered 4997 proteins. GS thus placed nearly all proteins in modules of size ≥ 2 . The average size of the modules produced by GS was 3.1 and 1043 of the modules contained only two proteins. The largest module produced by GS contained 129 proteins.

Figure 1 shows a visualization of the high-level summary structure returned by GS on the Y_{ppi} network. For clarity, we only show the summary induced by the supernodes containing at least two proteins and with at least one superedge (which may be a self-loop). Shown at the bottom are some of the supernodes GS found which represent cliques or near-cliques, indicated by a supernode with a self-loop.

3.2. Comparison of complex and biological process enrichment of the modules

A primary test of the quality of a module decomposition is its similarity to the modules induced by known biology. We use two different evaluation measures to assess the modules produced by each method (for details, see Section 2.4). A summary of the performance of these measures on both protein complexes and biological processes is shown in Figure 2.

The modules produced by GS cover the largest number of complexes and biological processes (Fig. 2a). Out of the 266 complexes appearing at least once in our network, 152 (57.1%) are enriched in at least one GS module. This indicates that GS performs well on many different biological units and is not simply highly successful for a few large complexes. MCL has 40.9% of the complexes enriched in at least one module. MCODE is limited in its ability to make predictions for many complexes because it only clusters 308 proteins and thus only covers 17.3% of the complexes. Only 16.5% are enriched in some NSP module, despite NSP clustering all nodes.

All of the 8 modules produced by NSP are enriched for at least one complex and typically enriched for several. This suggests that the graph partition produced by NSP does yield some useful information. However, because the modules are so large (average size: 687) it is not clear how to make use of them. Further, as we will see below, the percentage of modules enriched is not a good indicator of a method's ability to predict annotations.

GS also has the largest percentage (70.9%) of biological processes enriched in at least one module (Fig. 2b). MCL is again the second best-performing method with 63.7% of processes enriched. NSP and MCODE again have the largest percentage of their few modules enriched for some biological process. This, combined with the similar results for complex enrichment, suggests that GS gives the best coverage of possible biology. This is in contrast with MCODE, which produces good modules, but leaves 94% of the graph unclustered.

3.3. Improvement in module-assisted annotation prediction

Ultimately, the goal of dividing an interaction network into modules is to learn new biology. We test the utility of the various module-decomposition methods for predicting new complex and biological process annotations using three different methods for annotation transfer within a module (see Section 2.5). Figure 3 shows the performance of these prediction schemes for both complex membership and biological processes. Because MCL and MCODE both require parameter tuning, we include in Figure 3 the results of both approaches with and without tuning. The tuned versions are marked with a + symbol and show a significant improvement in performance when compared to the untuned versions. All numbers discussed below apply to the tuned versions of MCL and MCODE.

Complex membership prediction. The most conservative annotation transfer technique is to label each protein in a module with all the majority annotations within the module. Of the three annotation transfer methods, the majority approach results in fewer correctly predicted annotations, although these annotations tend to be more accurate. For complex membership, GS makes more correct predictions than MCL (305 for GS compared with 241 for MCL) and has higher precision (91.9% vs. 79.3%). MCODE makes slightly more precise predictions than GS (92.7% vs. 91.9%) but has a much lower recall (9.2% vs. 22.1%). Further, the predictions made by GS cover a larger number of proteins than MCL or MCODE (265 for GS, 219 for MCL, 104 for MCODE). NSP performs the worst in all cases due to its large module sizes.

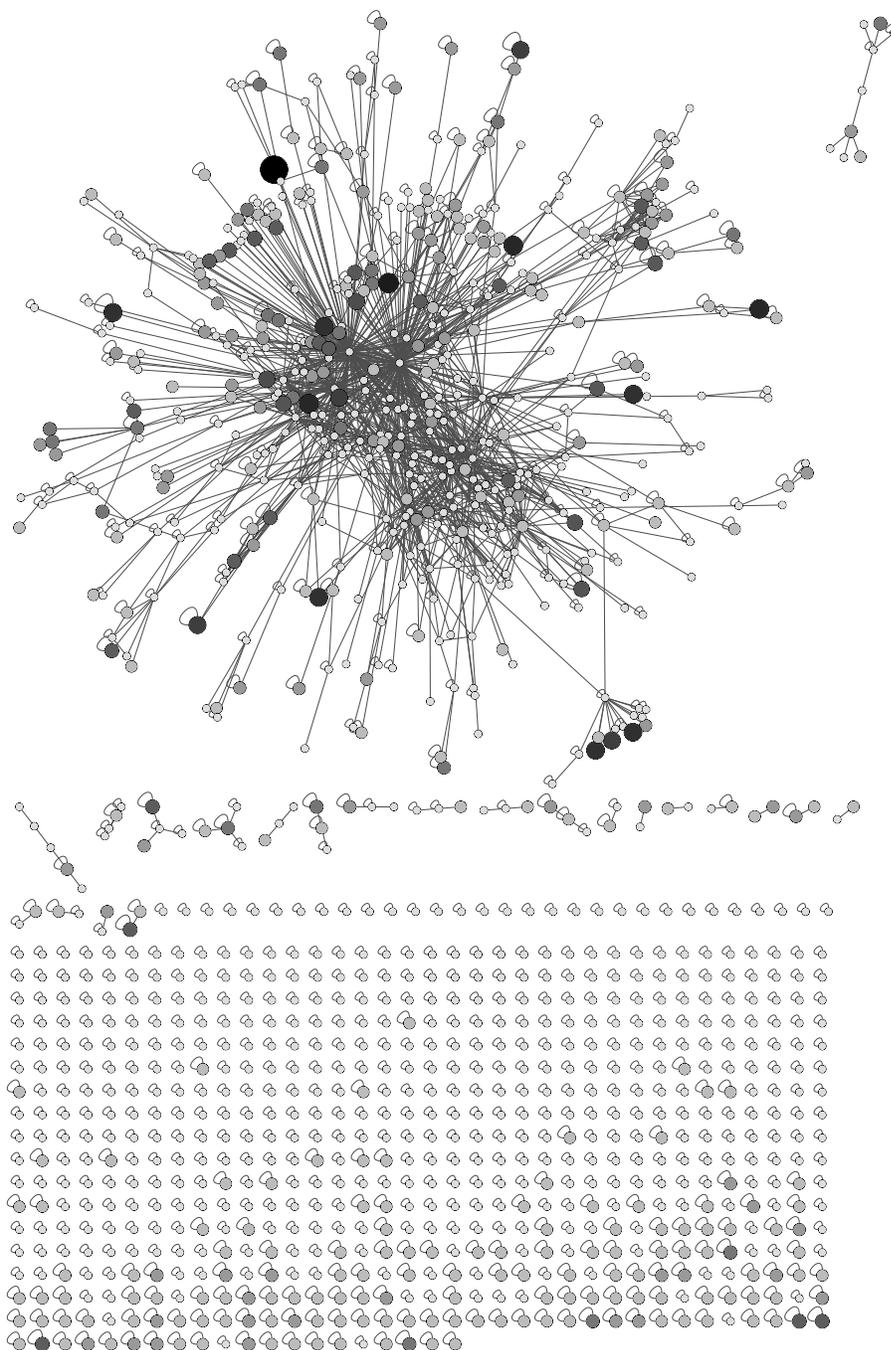


FIG. 1. A visualization, made in Cytoscape (Shannon et al., 2003), of the summary structure returned by GS on the Y_{ppi} network. Circles represent supernodes, with sizes proportional to their member proteins. Lines represent superedges.

When the transfer rule is relaxed to permit predictions based on the most common complex annotations within a module (plurality), many more correct predictions are made, though the predictions are less accurate. For complex membership, GS is again able to make more correct predictions (598) than MCL (436) and MCODE (164). Although MCODE's predictions are considerably more accurate than GS (86.8% vs. 60.5%), GS's recall is almost four times that of MCODE's (11.9% vs. 43.4%). GS also correctly predicts complexes for substantially more proteins than MCODE (530 vs. 140). This indicates that MCODE is able

make 330 correct predictions, more than with any other algorithm. It also has the highest recall (23.9% compared with MCL at 20.0%) and the highest precision (86.6% compared with MCODE at 83.6%). NSP makes correct predictions for slightly more proteins than GS (293 vs. 285) but NSP suffers from very low precision (10.6%). While the hypergeometric approach does not produce quite as many correct predictions as using the plurality rule, it makes more accurate predictions.

Figure 3 shows that the transfer schemes using the GS modules are all Pareto optimal: no other method dominates them in both precision and recall. For each transfer rule, GS also has the largest F -score compared to any other method. For example, using the majority rule, GS has a F -score of 35.7% compared to the next best at 28.7% for MCL. In general, MCL consistently does second-best. MCODE suffers from the fact that its clustering only covers 6% of the total nodes in the network.

Repeating the above experiments using more general complex annotations coming from level 3 of the MIPS complex catalog tree showed little change in the results. Hence, GS appears to work well for complexes with varying levels of specificity.

Biological processes prediction. Using the three annotation transfer schemes to predict biological process annotations for proteins reveals a similar pattern as with complexes. For the methods that produce smaller module sizes (GS and MCL), plurality makes the highest number of correct predictions, although at the expense of much lower accuracy compared with majority and hypergeometric. Interestingly, the hypergeometric rule makes fewer biological process predictions than majority, but is slightly more accurate—the opposite of what happens for complex prediction.

GS makes the largest number of correct predictions under all schemes except the hypergeometric (685 compared with 779 for MCL), but our accuracy is 26% higher than that of MCL (88.0% compared with 62.5%). The accuracy of both GS and MCL is slightly lower for predicting biological processes than for predicting complex membership.

In general, hypergeometric produces more correct predictions than majority at a slight loss in precision for complex annotations. For biological processes, the opposite is true: majority produces more correct predictions at a slight loss in precision. Although always lower in accuracy, the plurality rule makes more correct predictions and has a larger coverage than majority or hypergeometric for both biological processes and complexes. Thus, the choice of annotation transfer approach must be made with a desired level of predictions, precision, and coverage in mind.

Comparison of prediction sets. Because of their very different philosophies and approaches, GS, MCODE, and MCL complement each other well for the detection of biologically meaningful modules. GS produces a large number of correct predictions that none of the other methods make. Of the 305 correct complex predictions made by GS using the majority rule, 135 of the predictions (or 44%) are not made by either MCL or MCODE. Similarly, 39% of MCL's predictions are not made by either GS or MCODE. Of MCODE's 127 predictions, 23% are unique. The same is true using the hypergeometric rule for annotation transfer: Of the predictions made by GS, MCL, and MCODE, 44%, 38%, and 11% respectively are unique to each method. These numbers also carry over similarly to predicting biological processes: using the majority rule, 40% (550), 36% (371), and 17% (107) of the predictions made by GS, MCL, and MCODE, respectively are unique.

The large number of unique and correct predictions made by each method suggests each are able to uncover new biology that the other methods cannot. It also suggests that combining predictions from various methods may be useful, and emphasizes the point that no one approach to clustering is likely to be universally applicable.

3.4. High-confidence network

Protein interaction networks derived from high-throughput experiments are known to be noisy, with false positive rates potentially reaching 90% (Hart et al., 2006). As a result, we constructed an additional high-confidence yeast interaction network, $Y_{\text{high-conf}}$, that includes edges from IntAct that are associated with more than one PubMed identifier. (If the same identifier was listed twice by IntAct in support of an interaction, it was counted twice and thus included in $Y_{\text{high-conf}}$.) Two experiments are likely to both capture a true interaction, but are less likely to both return a false interaction. The $Y_{\text{high-conf}}$ network

TABLE 1. PREDICTING MIPS COMPLEXES AND GENE ONTOLOGY BIOLOGICAL PROCESSES IN $Y_{\text{high-conf}}$

Annotation	Method	Clusters		Majority			Plurality			Hyper		
		n	m	$prot$	R	P	$prot$	R	P	$prot$	R	P
Complexes	GS	2336	807	321	31.1	87.2	589	56.3	77.4	312	30.0	85.4
	MCL+	2323	615	386	36.4	84.0	571	55.2	71.1	395	38.6	67.7
	MCODE+	293	72	132	13.0	88.4	166	16.0	87.3	129	12.0	88.1
	NSP	2604	81	236	22.5	66.3	434	40.9	38.8	577	56.2	23.3
	MCL	2570	466	440	41.8	77.1	606	57.3	60.5	582	57.3	50.8
	MCODE	717	89	174	17.7	69.6	233	23.0	55.7	250	25.4	56.8
Processes	GS	2336	807	685	23.8	82.1	1308	41.9	51.9	345	8.9	86.3
	MCL+	2323	615	878	27.9	78.4	1328	37.7	53.4	632	18.9	82.6
	MCODE+	293	72	254	10.1	90.2	262	9.0	88.9	136	3.8	89.5
	NSP	2604	81	568	13.9	66.2	1014	17.2	43.3	1061	29.0	30.2
	MCL	2570	466	1048	31.6	75.6	1462	35.7	50.5	905	26.8	66.5
	MCODE	717	89	439	14.7	81.9	507	13.6	74.3	382	10.9	64.1

Columns list number of proteins clustered (n) and number of modules (m). For each transfer method, the number of proteins for which at least one correct prediction is made ($prot$), the recall (R), and precision (P) are given. Methods marked with + used tuned parameters.

contains 2604 proteins and 8341 interactions, fewer than half the proteins of the Y_{ppi} network. Using $Y_{\text{high-conf}}$, we reproduce the results described above to probe the effects of noise on the performance of the module-detection methods. The qualitative performance is similar to using the more comprehensive unfiltered network (Table 1). The GS approach makes more accurate complex predictions when compared with MCL and NSP. MCODE generally makes slightly more precise predictions, but at a large loss in recall and predicted annotations for different proteins.

3.5. Predicting co-complexed pairs

GS has the benefit of producing both a set of modules (supernodes) and a list of corrections to that modular structure. Analysis of the corrections list can lead to further insights. For example, we can use the corrections to predict missing edges that are indicative of co-complexed proteins. We consider negative corrections, where an edge must be removed from the summary to match the original graph, as predictions of pairs that are co-complexed. Positive corrections, where an edge must be added to restore the input graph, are predictions of noisy or erroneous edges that should be filtered. This edge prediction approach thus has the highly desirable feature of making predictions for both edges to add and edges to remove. GS applied this way can be thought of as a generalization of the popular method of completing defective cliques (DCC) (Yu et al., 2006) for predicting edges within apparent protein complexes.

We evaluated the performance of the GS co-complex predictions by comparing the edges predicted for $Y_{\text{high-conf}}$ to a gold standard set of co-complexed pairs composed of 11,014 edges between proteins annotated from the same MIPS complex (at the leaves of the hierarchy) and a negative set containing 2,705,720 edges between proteins annotated with different subcellular localizations (Yu et al., 2006). For this test we used the unmodified GS algorithm that does not add self edges before compression.

The DCC algorithm takes two parameters k and l controlling the overlap and size of defective cliques considered. We compared the results of the GS method to the results of the DCC algorithm over a range of parameters for k ($4 \leq k \leq 10$) and l ($2 \leq l \leq 5$). The precision of GS (66.7% for 224 predicted new edges) is better than DCC under all parameters for DCC (between 37.1% for 2317 predictions using $k = 4, l = 5$, and 62.5% for 39 predictions using $k = 9, l = 2$). Further, GS has very high specificity for accurately filtering incorrect edges from the network (97.4% for 3331 predicted edges). Because DCC does not predict edges to filter from the network, it is not applicable to compute its specificity. On the unfiltered Y_{ppi} network, DCC achieved higher precision for some parameters than GS, although the precision of both GS and DCC was generally poor. This is most likely because the large number of false edges in Y_{ppi} obfuscated the true complexes.

4. CONCLUSION

While the definition of a biological network “module” will likely remain unsettled for some time to come, the results presented here suggest it is biologically meaningful and informative to define a module as a set of proteins that have similar interaction partners. Such a definition generalizes cliques, defective cliques, and other types of dense subgraphs. Based on the minimum description length principle, graph summarization uses this notion of similar interaction partners to identify modules. It is well-suited to analyze protein interaction data because it can tolerate noisy edges in the network (which are placed in the corrections list). The GS algorithm also has no parameters that must be optimized, unlike both MCL and MCODE. Using GS to predict membership in protein complexes and biological processes led to increased performance compared with other approaches, even when their parameters are tuned to fit the data. The GS modules also cover a larger fraction of complexes and biological processes than other methods. Finally, GS works well for predicting co-complexed pairs. The general utility of graph summarization for extracting meaningful biological modules suggests that GS will be a useful technique for the analysis of biological networks.

ACKNOWLEDGMENTS

C.K. was partially supported by the NSF (grant IIS-0812111).

DISCLOSURE STATEMENT

No competing financial interests exist.

REFERENCES

- Ashburner, M., Ball, C.A., Blake, J.A., et al. 2000. Gene ontology: tool for the unification of biology. The Gene Ontology Consortium. *Nat. Genet.* 25, 25–29.
- Bader, G.D., and Hogue, C.W.V. 2003. An automated method for finding molecular complexes in large protein interaction networks. *BMC Bioinform.* 4, 2.
- Brohee, S., and van Helden, J. 2006. Evaluation of clustering algorithms for protein-protein interaction networks. *BMC Bioinform.* 7, 488.
- Brun, C., Chevenet, F., Martin, D., et al. 2003. Functional classification of proteins for the prediction of cellular function from a protein-protein interaction network. *Genome Biol.* 5, R6.
- Bu, D., Zhao, Y., Cai, L., et al. 2003. Topological structure analysis of the protein-protein interaction network in budding yeast. *Nucleic Acids Res.* 31, 2443–2450.
- Cherry, J.M., Ball, C., Weng, S., et al. 1997. Genetic and physical maps of *Saccharomyces cerevisiae*. *Nature* 387, 67–73.
- Clauset, A., Newman, M.E.J., and Moore, C. 2004. Finding community structure in very large networks. *Phys. Rev. E* 70, 066111.
- Enright, A.J., Van Dongen, S., and Ouzounis, C.A. 2002. An efficient algorithm for large-scale detection of protein families. *Nucleic Acids Res.* 30, 1575–1584.
- Fields, S., and Song, O. 1989. A novel genetic system to detect protein-protein interactions. *Nature* 340, 245–246.
- Fortunato, S., and Barthelemy, M. 2007. Resolution limit in community detection. *Proc. Natl. Acad. Sci. USA* 104, 36–41.
- Gavin, A.-C., Aloy, P., Grandi, P., et al. 2006. Proteome survey reveals modularity of the yeast cell machinery. *Nature* 440, 631–636.
- Guimera, R., and Luis, L.A.N. 2005. Functional cartography of complex metabolic networks. *Nature* 433, 895–900.
- Guldener, U., Munsterkötter, M., Kastenmüller, G., et al. 2005. CYGD: the comprehensive yeast genome database. *Nucleic Acids Res.* 33, D364.
- Han, J.-D.J., Bertin, N., Hao, T., et al. 2004. Evidence for dynamically organized modularity in the yeast protein-protein interaction network. *Nature* 430, 88–93.
- Hart, T.G., Ramani, A.K., and Marcotte, E.M. 2006. How complete are current yeast and human protein-interaction networks? *Genome Biol.* 7, 120.

- Hartwell, L.H., Hopfield, J.J., Leibler, S., et al. 1999. From molecular to modular cell biology. *Nature* 402, C47–C52.
- Jansen, R., Yu, H., Greenbaum, D., et al. 2003. A bayesian networks approach for predicting protein-protein interactions from genomic data. *Science* 302, 449–453.
- Kerrien, S., Alam-Faruque, Y., Aranda, B., et al. 2007. Intact–open source resource for molecular interaction data. *Nucleic Acids Res.* 35, D561–D565.
- King, A.D., Przulj, N., and Jurisica, I. 2004. Protein complex prediction via cost-based clustering. *Bioinformatics* 20, 3013–3020.
- Krogan, N.J., Cagney, G., Yu, H., et al. 2006. Global landscape of protein complexes in the yeast *Saccharomyces cerevisiae*. *Nature* 440, 637–643.
- Kumar, R., Raghavan, P., Rajagopalan, S., et al. 1999. Trawling the web for emerging cyber-communities. *Comput. Netw.* 31, 1481–1493.
- Myers, C.L., Barrett, D.R., Hibbs, M.A., et al. 2006. Finding function: evaluation methods for functional genomic data. *BMC Genomics* 7, 187.
- Navlakha, S., Rastogi, R., and Shrivastava, N. 2008. Graph summarization with bounded error. *SIGMOD 2008* 419–432.
- Newman, M.E.J. 2006. Modularity and community structure in networks. *Proc. Natl. Acad. Sci. USA* 103, 8577–8582.
- Newman, M.E.J., and Girvan, M. 2004. Finding and evaluating community structure in networks. *Phys. Rev. E* 69.
- Przulj, N., Wgle, D.A., and Jurisica, I. 2004. Functional topology in a network of protein interactions. *Bioinformatics* 20, 340–348.
- Qiu, J., and Noble, W.S. 2008. Predicting co-complexed protein pairs from heterogeneous data. *PLoS Comput. Biol.* 4.
- Radicchi, F., Castellano, C., Cecconi, F., et al. 2004. Defining and identifying communities in networks. *Proc. Natl. Acad. Sci. USA* 101, 2658–2663.
- Rissanen, J. 1978. Modeling by shortest data description. *Automatica* 14, 465–471.
- Rives, A.W., and Galitski, T. 2003. Modular organization of cellular networks. *Proc. Natl. Acad. Sci. USA* 100, 1128–1133.
- Rosvall, M., and Bergstrom, C.T. 2007. An information-theoretic framework for resolving community structure in complex networks. *Proc. Natl. Acad. Sci. USA* 104, 7327–7331.
- Samanta, M.P., and Liang, S. 2003. Predicting protein functions from redundancies in large-scale protein interaction networks. *Proc. Natl. Acad. Sci. USA* 100, 12579–12583.
- Shannon, P., Markiel, A., Ozier, O., et al. 2003. Cytoscape: a software environment for integrated models of biomolecular interaction networks. *Genome Res.* 13, 2498–2504.
- Sharan, R., Ulitsky, I., and Shamir, R. 2007. Network-based prediction of protein function. *Nat. Mol. Syst. Biol.* 3, 88.
- Shoemaker, B.A., and Panchenko, A.R. 2007. Deciphering protein-protein interactions. Part I. Experimental techniques and databases. *PLoS Comput. Biol.* 3, e42.
- Spirin, V., and Mirny, L.A. 2003. Protein complexes and functional modules in molecular networks. *Proc. Natl. Acad. Sci. USA* 100, 12123–12128.
- Tanay, A., Sharan, R., Kupiec, M., et al. 2004. Revealing modularity and organization in the yeast molecular network by integrated analysis of highly heterogeneous genomewide data. *Proc. Natl. Acad. Sci. USA* 101, 2981–2986.
- Troyanskaya, O.G., Dolinski, K., Owen, A.B., et al. 2003. A Bayesian framework for combining heterogeneous data sources for gene function prediction (in *Saccharomyces cerevisiae*). *Proc. Natl. Acad. Sci. USA* 100, 8348–8353.
- Ulitsky, I., and Shamir, R. 2007. Identification of functional modules using network topology and high-throughput data. *BMC Syst. Biol.* 1, 8.
- van Dongen, S. 1998. A new cluster algorithm for graphs, T.R. 9814, Centrum voor Wiskunde en Informatica (CWI).
- Wakita, K., and Tsurumi, T. 2007. Finding community structure in mega-scale social networks. arXiv:cs/0702048v1.
- Wang, C., Ding, C., Yang, Q., et al. 2007. Consistent dissection of the protein interaction network by combining global and local metrics. *Genome Biol.* 8, R271.
- Yu, H., Paccanaro, A., Trifonov, V., et al. 2006. Predicting interactions in protein networks by completing defective cliques. *Bioinformatics* 22, 823–829.
- Zhu, X., Gerstein, M., and Snyder, M. 2007. Getting connected: analysis and principles of biological networks. *Genes Dev.* 21, 1010–1024.
- Ziv, E., Middendorf, M., and Wiggins, C.H. 2005. Information-theoretic approach to network modularity. *Phys. Rev. E* 71, 046117.

Address reprint requests to:
Dr. Carl Kingsford
Department of Computer Science
University of Maryland
College Park, MD 20742

E-mail: carlk@cs.umd.edu