Classification in biological networks with hypergraphlet kernels

Jose Lugo-Martinez[1], Daniel Zeiberg[2], Thomas Gaudelet[3], Noël Malod-Dognin[4],
Nataša Pržulj[4], and Predrag Radivojac[2]

[1]Computational Biology Department, Carnegie Mellon University, Pittsburgh, Pennsylvania, U.S.A.
[2]Khoury College of Computer Sciences, Northeastern University, Boston, Massachusetts, U.S.A.
[3]Department of Computer Science, University College London, London WC1E 6BT, U.K.
[4]Barcelona Supercomputing Center, Barcelona, 08034 Spain
Contact: predrag@northeastern.edu

Abstract

Biological and cellular systems are often modeled as graphs in which vertices represent objects of interest (genes, proteins, drugs) and edges represent relational ties between these objects (binds-to, interacts-with, regulates). This approach has been highly successful owing to the theory, methodology and software that support analysis and learning on graphs. Graphs, however, suffer from information loss when modeling physical systems due to their inability to accurately represent multi-object relationships. Hypergraphs, a generalization of graphs, provide a framework to mitigate information loss and unify disparate graph-based methodologies. We present a hypergraph-based approach for modeling biological systems and formulate vertex classification, edge classification and link prediction problems on (hyper)graphs as instances of vertex classification on (extended, dual) hypergraphs. We then introduce a novel kernel method on vertex- and edge-labeled (colored) hypergraphs for analysis and learning. The method is based on exact and inexact (via hypergraph edit distances) enumeration of hypergraphlets; i.e., small hypergraphs rooted at a vertex of interest. We empirically evaluate this method on fifteen biological networks and show its potential use in a positive-unlabeled setting to estimate the interactome sizes in various species.

Availability: https://github.com/jlugomar/hypergraphlet-kernels/

1 Introduction

Graphs provide a mathematical structure for describing relationships between objects in a system. Owing to their intuitive representation, well-understood theoretical properties, the wealth of methodology and available code base, graphs have also become a major framework for modeling biological systems. Protein-protein interaction networks, protein 3D structure graphs, drug-target interaction networks, metabolic networks and gene regulatory networks are some of the major representations of biological systems. Unfortunately, molecular and cellular systems are only partially observable and may contain significant amount of noise due to their inherent stochastic nature as well as the limitations of experimental techniques. This highlights the need for the development and application of computational approaches for predictive modeling (e.g., inferring novel interactions) and identifying interesting patterns in such data.

Learning on graphs can be generally seen as supervised or unsupervised. Under a supervised setting, typical tasks involve *graph classification*; i.e., the assignment of class labels to entire graphs,

vertex or edge classification; i.e., the assignment of class labels to vertices or edges, or *link prediction*; i.e., the prediction of the existence of edges in graphs. Alternatively, frequent subgraph mining, motif finding, and clustering are traditional unsupervised approaches. Regardless of the category, the development of techniques that capture network structure, measure graph similarity and incorporate domain-specific knowledge in a principled manner lie at the core of all these problems.

The focus of this study is on classification problems across various biological networks. A straightforward approach to this problem is the use of topological and other descriptors (e.g., vertex degree, clustering coefficient, betweenness centrality) that summarize graphs and graph neighborhoods. These descriptors, much like embedding techniques [1, 2], readily form vector-space representations, after which standard machine learning algorithms can be applied to learn target functions [3]. Another strategy involves the use of kernel functions on graphs [4]. Kernels are symmetric positive semi-definite mappings of pairs of objects from an input space $\mathcal{X}$ to $\mathbb{R}$, that lead to efficient learning. Finally, classification on graphs can be seen as inference over Markov networks [5] and can be approached using related label-propagation [6] or flow-based [7] methods. These inference strategies are often well adjusted to learning smooth functions over neighboring nodes.

Despite the success and wide adoption of these methods in computational biology, it is well-understood that graph representations suffer from information loss since every edge can only encode pairwise relationships [8]. A protein complex, for instance, cannot be distinguished from a set of proteins that interact only pairwise. Such disambiguation, however, is important in order to understand the biological activity of these molecules [9]. Hypergraphs, a generalization of graphs, naturally capture these higher-order relationships [10]. As we show later, they also provide a representation that can be used to unify several conventional classification problems on (hyper)graphs as an instance of vertex classification on hypergraphs.

In this paper, we present and evaluate a kernel-based framework for the problems of vertex classification, edge classification and link prediction in graphs and hypergraphs. We first use the concepts of hypergraph duality to demonstrate that all such classification problems can be unified through the use of hypergraphs. We then describe the development of edit-distance hypergraphlet kernels; i.e., similarity functions between local vertex neighborhoods based on flexibly enumerating small labeled hypergraphs rooted at vertices of interest. These similarity functions were subsequently incorporated into a semi-supervised methodology for predicting class labels on vertices. Finally, we use fifteen biological networks to provide evidence that the proposed approaches present an attractive option in this setting.

2 Background

2.1 Graphs and hypergraphs

Graphs. A graph G is a pair (V, E), where V is a set of vertices (nodes) and $E \subseteq V \times V$ is a set of edges. In a vertex-labeled graph, a labeling function f is defined as $f : V \to \Sigma$, where Σ is a finite alphabet. Similarly, in an edge-labeled graph, another labeling function g is given as $g : E \to \Xi$, where Ξ is also a finite set. We will focus on undirected graphs (E is symmetric), without self-loops and weights associated with edges. Generalization of our approach to directed and weighted graphs is relatively straightforward.

A rooted graph G is a graph together with one distinguished vertex called the root. We denote such graphs as $G = (V, v, E)$, where $v \in V$ is the root. A walk w of length k in a graph G is a

sequence of nodes $v_0, v_1 \cdots, v_k$ such that $(v_i, v_{i+1}) \in E$, for $0 \leq i < k$. A connected graph is a graph where there exists a walk between any two nodes.

Hypergraphs. A hypergraph G is a pair (V, E), where V again is the vertex set and E is a family of non-empty subsets of V, referred to as hyperedges. A hyperedge e is said to be incident with a vertex v if $v \in e$. Two vertices are called adjacent if there is an edge that contains both vertices and two hyperedges are said to be adjacent if their intersection is non-empty. The neighbors of a vertex v in a hypergraph are the vertices adjacent to v. Finally, the degree $d(v)$ of a vertex v in a hypergraph is given by $d(v) = |\{e \in E \,|\, v \in e\}|$, whereas the degree of a hyperedge e is defined as $\delta(e) = |e|$.

As before, one can define a vertex-labeled, edge-labeled and rooted hypergraphs. When the multiplicity of each hyperedge is one, the hypergraph is said to be simple. A walk w of length k in a hypergraph is a sequence of vertices and hyperedges $v_0, e_0, v_1, \cdots, e_{k-1}, v_k$ such that $(v_i, v_{i+1}) \in e_i$ for each $0 \leq i < k$ and $e_i \in E$. A connected hypergraph is a hypergraph where there exists a walk between any two nodes.

Since we are interested in counting small predefined hypergraphs in large hypergraphs, we define two distinct ways in which these substructures can be counted, as *section hypergraphs* or *sub-hypergraphs*. Given a hypergraph $G = (V, E)$, a section hypergraph of G is a hypergraph $G' = (V', E')$ where $V' \subseteq V$ and $E' = \{e \,|\, e \in E \land e \subseteq V'\}$. A sub-hypergraph of G is a hypergraph $G' = (V', E')$ where $V' \subseteq V$ and $E' = \{e \cap V' \,|\, e \in E \land e \cap V' \neq \emptyset\}$. In other words, when nodes $V \setminus V'$ are removed from V, section hypergraphs retain only those edges that were subsets of the remaining nodes V'. In sub-hypergraphs, on the other hand, the edges that originally contained nodes from both removed $(V \setminus V')$ and unremoved (V') nodes are retained as subsets of the original edges that now include nodes only from V'. Edges from the original graph that only included nodes from V' can now have increased multiplicity.

Isomorphism. Consider two graphs, $G = (V, E)$ and $H = (W, F)$. We say that G and H are isomorphic, denoted as $G \cong H$, if there exists a bijection $f : V \to W$ such that $(u, v) \in E$ if and only if $(f(u), f(v)) \in F$ for all $u, v \in V$. If G and H are hypergraphs, an isomorphism is defined as interrelated bijections $f : V \to W$ and $g : E \to F$ such that $e = \{v_1, \cdots, v_{\delta(e)}\} \in E$ if and only if $g(e) = \{f(v_1), \cdots, f(v_{\delta(e)})\} \in F$ for all hyperedges $e \in E$. Isomorphic (hyper)graphs are structurally identical. An automorphism is an isomorphism of a (hyper)graph to itself.

Edit distance. Let G and H be two vertex- and hyperedge-labeled hypergraphs. The edit distance between these hypergraphs corresponds to the minimum number of edit operations necessary to transform G into H, where edit operations are often defined as insertion/deletion of vertices/hyperedges and substitutions of vertex and hyperedge labels. Any sequence of edit operations that transforms G into H is called an edit path; hence, the hypergraph edit distance between G and H is the length of the shortest edit path between them. This concept can be generalized to the case where each edit operation is assigned a cost. Hypergraph edit distance then corresponds to the edit path of minimum cost.

2.2 Hypergraph duality

Let $G = (V, E)$ be a hypergraph, where $V = \{v_1, \ldots, v_n\}$ and $E = \{e_1, \ldots, e_m\}$. The dual hypergraph of G, denoted as $G^* = (V^*, E^*)$, is obtained by constructing the set of vertices as $V^* = \{e_1, \ldots, e_m\}$ and the set of hyperedges as $E^* = \{\epsilon_1, \ldots, \epsilon_n\}$ such that $\epsilon_i = \{e_j \,|\, v_i \in e_j\}$. Thus, the vertices of the dual hypergraph G^* are hyperedges of the original hypergraph G, whereas the hyperedges of G^* are constructed using the hyperedges of G that are incident with the respective vertices. Figure 1A-B shows two examples of a hypergraph G and its dual hypergraph G^*.

Figure 1: Examples of hypergraph duality. Panel (A) shows a hypergraph $G = (V, E)$, where $V = \{v_1, v_2, v_3, v_4\}$ and $E = \{e_1, e_2, e_3, e_4, e_5\}$ with its dual hypergraph $G^* = (V^*, E^*)$, where $V^* = \{e_1, e_2, e_3, e_4, e_5\}$ and $E^* = \{\epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4\}$ such that $\epsilon_1 = \{e_1, e_3, e_5\}$, $\epsilon_2 = \{e_1, e_2\}$, $\epsilon_3 = \{e_2, e_3, e_4\}$ and $\epsilon_4 = \{e_4, e_5\}$. Panel (B) shows an example of graph G with two degree-one vertices that lead to the dual hypergraph G^* with self-loops; ϵ_2 and ϵ_4. Panel (C) shows an extended dual hypergraph that is proposed to formulate link prediction as an instance of vertex classification in hypergraphs. To make a prediction regarding the existence of edge $\bar{e}$, shown as a dashed line on the left side, an extended dual hypergraph is created in which $\bar{e}$ is added to the set of vertices V^*. Updates are made to hyperedges ϵ_1 and ϵ_4 (dashed) that correspond to those vertices in G that are incident with the edge $\bar{e}$.

2.3 Classification on hypergraphs

We are interested in binary classification on (possibly disconnected) hypergraphs. The following paragraphs briefly introduce three classification problems on hypergraphs, formulated here so as to naturally lead to the methodology proposed in the next section.

Vertex classification. Given a hypergraph $G = (V, E)$ and a training set $\{(v_i, t_i)\}_{i=1}^m$, where $t_i \in \{-1, +1\}$ is the class label of vertex v_i and $m < |V|$, the goal is to predict class labels of unlabeled vertices. A number of classical problems in computational biology map straightforwardly to vertex classification; e.g., network-based protein function prediction, disease gene prioritization, etc.

Hyperedge classification. Given a hypergraph $G = (V, E)$ and a training set $\{(e_i, t_i)\}_{i=1}^m$, where $t_i \in \{-1, +1\}$ is the class label of hyperedge e_i and $m < |E|$, the goal is to predict class labels of unlabeled hyperedges. An example of hyperedge classification is the prediction of functional annotations for protein complexes.

Link prediction. Let $G = (V, E)$ be a hypergraph with some missing hyperedges and let $\bar{E}$ be all non-existent hyperedges in G; i.e., $\bar{E} = \mathcal{U} - E$, where $\mathcal{U}$ represents all possible hyperedges over V. The goal is to learn a target function $t : \mathcal{U} \to \{-1, +1\}$ and infer the existence of all missing hyperedges. Examples of link prediction include predicting protein-protein interactions, predicting drug-target interactions, and so on.

2.4 Positive-unlabeled learning

A number of prediction problems in computational biology can be considered within a semi-supervised framework, where a set of labeled and a set of unlabeled examples are used to construct classifiers that discriminate between positive and negative examples. A special category of semi-supervised learning occurs when labeled data contain only positive examples; i.e., where the negative examples are either unavailable or ignored; say, if the set of available negatives is small or biased. Such problems are generally referred to as learning from positive and unlabeled data or positive-unlabeled learning [11]. Many problems in molecular biology that are often referred to as the *open world* problems lend themselves naturally to the positive-unlabeled setting.

Research in machine learning has recently established tight connections between traditional supervised learning and (non-traditional) positive-unlabeled learning. Under mild conditions, a classifier that optimizes the ranking performance; e.g., area under the ROC curve [12], in the non-traditional setting has been shown to also optimize the performance in the traditional setting [13]. Similar relationships have been established in approximating posterior distributions [14, 15] as well as in recovering the true performance accuracy in the traditional setting for a classifier evaluated in a non-traditional setting [16–18].

3 Methods

3.1 Problem formulation

We consider binary classification problems on graphs and hypergraphs and propose to unify all such learning problems through semi-supervised vertex classification on hypergraphs. First, vertex classification falls trivially into this framework. Second, the problems of edge classification in graphs and hyperedge classification in hypergraphs are equivalent to the problem of vertex classification on dual hypergraphs. As discussed in Section 2.2, both graphs and hypergraphs give rise to dual hypergraph representations and, thus, (hyper)edge classification on a graph G straightforwardly translates into vertex classification on its dual hypergraph G^*. We note here that vertices with the degree of one in G give rise to self-loops in the dual hypergraph G^*. To account for them, we add one dummy node per self-loop with the same vertex label as the original vertex and connect them with an appropriately labeled edge. Third, one can similarly see link prediction as vertex classification on dual hypergraphs, where the set of existing links is treated as positive data, the set of known non-existing links is treated as negative data, and the remaining set of missing links is treated as unlabeled data. This formulation further requires an extension of dual hypergraph representations as follows. Consider a particular negative or missing link $\bar{e} \in \bar{E}$ in the original graph G with its dual hypergraph G^* (Fig. 1C). To make a prediction on this edge $\bar{e}$, we must first introduce a new vertex $\bar{e}$ in the dual hypergraph as well as modify those hyperedges in G^* that correspond to the vertices $v \in \bar{e}$ in G (Fig. 1C). We denote this extended hypergraph as $G^*_{\bar{e}}$. It now easily follows that the sets of negative and unlabeled examples can be created by considering a collection of extended graphs $G^*_{\bar{e}}$, one at a time, for select vertices $\bar{e} \in \bar{E}$.

Since most biological networks lack large sets of representative negative examples, we approach vertex classification, (hyper)edge classification and link prediction as instances of vertex classification on (extended, dual) hypergraphs in a positive-unlabeled setting. We believe this is a novel and useful attempt at generalizing three distinct graph classification problems in a common semi-supervised setting. The following sections introduce hypergraphlet kernels that are the next step of our classification approach.

3.2 Hypergraphlets

Inspired by graphlets [19, 20], we define a *hypergraphlet* as a small, simple, connected, rooted hypergraph, without self-loops, where the root of the hypergraph is its automorphism orbit [9]. A hypergraphlet with n vertices is called an n-hypergraphlet; and the i-th hypergraphlet of order n is denoted as n_i. We consider hypergraphlets up to isomorphism and will refer to these isomorphisms as root- and label-preserving isomorphisms when hypergraphs are rooted and labeled. Figure 2 displays all non-isomorphic unlabeled n-hypergraphlets with up to three vertices. There is only one hypergraphlet of order 1 (1_1; Fig. 2A), one hypergraphlet of order 2 (2_1; Fig. 2B), nine hypergraphlets of order 3 ($3_1, \ldots, 3_9$; Fig. 2C) and 461 hypergraphlets of order 4 (Supplementary Materials). We refer to all these hypergraphlets as *base hypergraphlets* since they correspond to the case when $|\Sigma| = |\Xi| = 1$.

Figure 2: Undirected base hypergraphlets. Undirected base hypergraphlets with (A) 1, (B) 2, and (C) 3 vertices. The root node of each hypergraphlet is inscribed in a square. Hypergraphlets are presented in a compressed notation; e.g., structures 3_2 and 3_3 are shown in one drawing.

Consider now a vertex- and hyperedge-labeled (or *fully labeled* for short) hypergraphlet with n vertices and m hyperedges, where Σ and Ξ denote the vertex-label and hyperedge-label alphabets, respectively. If $|\Sigma| > 1$ and/or $|\Xi| > 1$, automorphic structures corresponding to the same base hypergraphlet may exist; hence, the number of fully labeled hypergraphlets per base structure is generally smaller than $|\Sigma|^n \cdot |\Xi|^m$. For example, if one only considers vertex-labeled 3-hypergraphlets, then there are $|\Sigma|^3$ vertex-labeled hypergraphlets corresponding to the asymmetric base hypergraphlets 3_2, 3_4 and 3_7 but only $\frac{1}{2}(|\Sigma|^3 + |\Sigma|^2)$ corresponding to the base hypergraphlets 3_1, 3_3, 3_5, 3_6, 3_8, 3_9. This is a result of symmetries in the base hypergraphlets that give rise to automorphisms among vertex-labeled structures. Similarly, if $|\Xi| > 1$, new symmetries may form with respect to the base hypergraphlets that give rise to different automorphisms among hyperedge-labeled structures.

These symmetries are important for counting vertex- and hyperedge-labeled hypergraphlets within larger graphs. The enumeration steps described above also determine the dimensionality of the Hilbert space in which the prediction is carried out. Detailed results on hypergraph enumeration are given in Supplementary Materials.

3.3 Hypergraphlet kernels

Motivated by the case for graphs [21–23], we introduce *hypergraphlet kernels*. Let $G = (V, E, f, g, \Sigma, \Xi)$ be a fully labeled hypergraph where f is a vertex-labeling function $f : V \to \Sigma$, g is a hyperedge-labeling function $g : E \to \Xi$, and $|\Sigma|, |\Xi| \geq 1$. The vertex- and hyperedge-labeled n-hypergraphlet

count vector for any $v \in V$ is defined as

$$\phi_n(v) = (\varphi_{n_1}(v), \varphi_{n_2}(v), \ldots, \varphi_{n_{\kappa(n,\Sigma,\Xi)}}(v)), \tag{1}$$

where $\varphi_{n_i}(v)$ is the count of the i-th fully labeled n-hypergraphlet rooted at v and $\kappa(n, \Sigma, \Xi)$ is the total number of vertex- and hyperedge-labeled n-hypergraphlets. Observe that, $\varphi_{n_i}(v)$ must be defined over a section hypergraph or sub-hypergraph of G which will lead to distinct vectors of counts $\phi_n(v)$. A kernel between the n-hypergraphlet counts for vertices u and v is defined as an inner product between $\phi_n(u)$ and $\phi_n(v)$; i.e.,

$$k_n(u, v) = \langle \phi_n(u), \phi_n(v) \rangle. \tag{2}$$

The hypergraphlet kernel function incorporating all hypergraphlets up to the size N is given by

$$k(u, v) = \sum_{n=1}^{N} k_n(u, v), \tag{3}$$

where N is a small integer. In this work we use $N = 4$ due to the exponential growth of the number of base hypergraphlets.

3.4 Edit-distance hypergraphlet kernels

Consider a fully labeled hypergraph $G = (V, E, f, g, \Sigma, \Xi)$. Given a vertex $v \in V$, we define the vector of counts for a τ-generalized edit-distance hypergraphlet representation as

$$\phi_{(n,\tau)}(v) = (\psi_{(n_1,\tau)}(v), \psi_{(n_2,\tau)}(v), \ldots, \psi_{(n_{\kappa(n,\Sigma,\Xi)},\tau)}(v)), \tag{4}$$

where

$$\psi_{(n_i,\tau)}(v) = \sum_{n_j \in E(n_i,\tau)} c(n_i, n_j) \cdot \varphi_{n_j}(v). \tag{5}$$

Here, $E(n_i, \tau)$ is the set of all n-hypergraphlets such that for each $n_j \in E(n_i, \tau)$ there exists an edit path of total cost at most τ that transforms n_i into n_j and $c(n_i, n_j) \geq 0$ is a user-defined parameter. In other words, the counts for each hypergraphlet n_i are updated by also counting all other hypergraphlets n_j that are in the τ vicinity of n_i. The parameter c can be used to adjust the weights of these pseudocounts or learned from data to add sophistication. We set $c(n_i, n_j) = 1$ for all i and j and the cost of all edit operations was also set to 1. This restricts τ to nonnegative integers.

The length-τ edit-distance n-hypergraphlet kernel $k_{(n,\tau)}(u, v)$ between vertices u and v can be computed as an inner product between the respective count vectors $\phi_{(n,\tau)}(u)$ and $\phi_{(n,\tau)}(v)$; i.e.,

$$k_{(n,\tau)}(u, v) = \left\langle \phi_{(n,\tau)}(u), \phi_{(n,\tau)}(v) \right\rangle. \tag{6}$$

The length-τ edit-distance hypergraphlet kernel function is given as

$$k_\tau(u, v) = \sum_{n=1}^{N} k_{(n,\tau)}(u, v). \tag{7}$$

The edit operations considered here incorporate substitutions of vertex labels, substitutions of hyperedge labels, and insertions/deletions (indels) of hyperedges (see example in Figure S1). Given these edit operations, we also define three subclasses of edit-distance hypergraphlet kernels referred

7

to as vertex label-substitution $k_\tau^{vl}(u,v)$, hyperedge label-substitution $k_\tau^{hl}(u,v)$ and hyperedge-indel kernels $k_\tau^{hi}(u,v)$. Although the functions from Equations (2) and (6) are defined as inner products, other formulations such as radial basis functions can be similarly considered [24]. We also note that the combined kernels from Equations (3) and (7) can be generalized beyond linear combinations. For the simplicity of this work, however, we only explore equal-weight linear combinations and normalize the functions from Equations (3) and (7) using the cosine transformation. Computational complexity and implementation details are described in Supplementary Materials.

3.5 Data sets

Protein-protein interaction data. The protein-protein interaction (PPI) data was used for both edge classification and link prediction. In the context of edge classification, we are given a PPI network where each interaction is annotated as either direct physical interaction or a co-membership in a complex. The objective is to predict the type of each interacting protein pair as physical vs. complex (PC). For this task, we used the budding yeast *S. cerevisiae* PPI network assembled by Ben-Hur & Noble [25].

Another important task in PPI networks is discovering whether two proteins interact. Despite the existence of high-throughput experimental methods for determining interactions between proteins, the PPI network data of all organisms is incomplete [26]. Furthermore, high-throughput PPI data contains a potentially large fraction of false positive interactions [27]. Therefore, there is a continued need for computational methods to help guide experiments for identifying novel interactions. Under this scenario, there are two classes of link prediction algorithms: (1) prediction of direct physical interactions [25, 28] and (2) prediction of co-membership in a protein complex [29, 30]. We focused on the former task and assembled six species-specific data sets comprised solely of direct protein-protein interaction data derived from public databases (BIND, BioGRID, DIP, HPRD, and IntAct) as of January 2017. We considered only one protein isoform per gene and used experimental evidence types described by Lewis et al. [26]. Specifically, we constructed link prediction tasks for: (1) *E. coli* (EC), (2) *S. pombe* (SP), (3) *R. norvegicus* (RN), (4) *M. musculus* (MM), (5) *C. elegans* (CE), and (6) *A. thaliana* (AT).

Drug-target interaction data. Identification of interactions between drugs and target proteins is an area of growing interest in drug design and therapy [31, 32]. In a drug-target interaction (DTI) network, nodes correspond to either drugs or proteins and edges indicate that a protein is a known target of the drug. Here we used DTI data for both edge classification and link prediction. In the context of edge labeling, we are given a DTI network where each interaction is annotated as direct (binding) or indirect, as well as assigned modes of action as activating or inhibiting. The objective is to predict the type of each interaction between proteins and drug compounds. For this task, we derived two data sets: (1) indirect vs. direct (ID) binding derived from MATADOR, and (2) activation vs. inhibition (AI) assembled from STITCH. Under link prediction setting, the learning task is to predict drug-target protein interactions. We focused on four drug-target classes: (1) enzymes (EZ), (2) ion channels (IC), (3) G protein-coupled receptors (GR), and (4) nuclear receptors (NR); originally assembled by Yamanishi et al. [31].

Hyperedge classification data. Hypergraphs provide a natural way to encode protein complexes. Here, we used Corum 3.0 [33] which provides manually annotated protein complexes from mammalian organisms. Under this setting, the learning objective is to predict the functional annotation of each protein complex. For this task, we defined two data sets: (1) protein binding vs. cell cycle (BC), and (2) protein modification vs. DNA processing (MP). Table 1 summarizes all data sets used in this work.

Table 1: Summary of binary classification tasks and data sets. For each learning problem, we show the number of vertices (V) and edges (E) in the full hypergraph, as well as the largest connected component (V^{lcc}, E^{lcc}). We also show the number of positive (n_+), negative (n_-), or unlabeled (n_u) data points.

Type	Dataset				
		Edge classification			
		$\|V\|$	$\|E\|$	n_+	n_-
PPI	PC	$4,761$	$22,988$	$10,517$	$12,471$
DTI	ID	544 drugs $2,261$ targets	$10,436$	$4,284$	$6,152$
	AI	378 drugs 267 targets	$1,039$	249	790
		Hyperedge classification			
		$\|V\|$	$\|E\|$	n_+	n_-
PPI	BC	$3,436$	$2,357$	145	161
	MP	$3,436$	$2,357$	175	200
		Link prediction			
		$\|V\|$	$\|E\|$	$\|V^{\mathrm{lcc}}\|$	$\|E^{\mathrm{lcc}}\|^a$
PPI	EC	393	391	100	153
	CE	$3,026$	$5,163$	$2,779$	$5,014$
	AT	$5,391$	$12,825$	$5,063$	$12,631$
	SP	853	$1,197$	685	$1,092$
	RN	526	532	301	388
	MM	$2,065$	$2,833$	$1,590$	$2,522$
DTI	EZ	445 drugs 664 targets	$2,926$	809	$2,556$
	IC	210 drugs 204 targets	$1,476$	409	$1,473$
	GR	223 drugs 95 targets	635	240	570
	NR	54 drugs 26 targets	90	42	50

a The size of n_+ and n_u is given by $\|E^{\mathrm{lcc}}\|$.

3.6 Integrating domain knowledge via a vertex alphabet

To incorporate domain knowledge into the PPI networks, we exploited the fact that each vertex (protein) in the graph is associated with its amino acid sequence. In particular, we used protein sequences to predict their Gene Ontology (GO) terms using the FANN-GO algorithm [34]. Hierarchical clustering was subsequently used on the predicted term scores to group proteins into $|\Sigma_{\text{GO}}|$ broad functional categories. In the case of DTI data, target proteins were annotated in a similar manner. For labeling drug compounds, we used the chemical structure similarity matrix computed from SIMCOMP [35], transformed it into a dissimilarity matrix and then applied hierarchical clustering to group compounds into $|\Sigma_{\text{SS}}|$ structural categories.

3.7 Evaluation methodology

We evaluated all hypergraphlet kernels by comparing them to two in-house implementations of random-walk-with-restarts kernels. Given a hypergraph G and two vertices u and v, simultaneous random walks w_u and w_v were generated from u and v for a fixed number of times (N_{walks}). In each step during a walk, one must pick hyperedges $e_{u'}$ and $e_{v'}$ incident with current vertices u' and v', respectively, and then pick next vertices $u'' \in e_{u'}$ and $v'' \in e_{v'}$. After a transition is made to the next pair of nodes u'' and v'', the walk is terminated with some probability $0 < p < 1$, or continued. In the conventional random walk implementation, a walk is scored as 1 if the sequences of all vertex and hyperedge labels between w_u and w_v are identical; otherwise, a walk is scored as 0. After $N_{\text{walks}} = 10{,}000$ walks are completed, the scores over all walks are summed to produce a kernel value between the starting vertices u and v. In order to construct a random walk similar to the hypergraphlet edit-distance approach, a cumulative random walk kernel was also implemented. Here, any match between the labels of vertices u' and v', or hyperedges $e_{u'}$ and $e_{v'}$ of each walk is scored as 1, while a mismatch is scored as 0. Thus, a walk of length ℓ can contribute between 0 and $2\ell - 1$ to the total count. In each of the random walks, the probability of restart p was selected from a set $\{0.1, 0.2, \ldots, 0.5\}$. On the link prediction data sets we also evaluated the performance of the preferential attachment method [36] and the L3 framework [37]. Furthermore, we evaluated pairwise spectrum kernels [25] on the PPI data sets (excluding the protein complex data). The k-mer size for pairwise spectrum kernels was varied from $k \in \{3, 4, 5\}$. Finally, in the case of the edit-distance kernels, we computed the set of normalized hypergraphlet kernel matrices $\mathcal{K}$ using $k_\tau^{vl}(x_i, x_j)$, $k_\tau^{hl}(x_i, x_j)$, $k_\tau^{hi}(x_i, x_j)$, and $k_\tau(x_i, x_j)$ for all pairs (x_i, x_j) obtained from a grid search over $\tau = \{0, 1\}$, $|\Sigma| = \{2, 4, 8, 16\}$ and $N = \{3, 4\}$.

The performance of each method was evaluated through a 10-fold cross-validation. In each iteration, 10% of nodes in the dual network were selected for the test set, whereas the remaining 90% were used for training. When evaluating link prediction methods, the negative examples were sampled with probabilities proportional to the product of degrees of the two nodes. Support vector machine (SVM) classifiers were used to construct all predictors and perform comparative evaluation. We used SVM^{light} with the default value for the capacity parameter [38]. Once each predictor was trained, we used Platt's correction to adjust the outputs of the predictor to the 0-1 range [39]. Finally, we estimated the area under the ROC curve (AUC), which plots the true positive rate (sensitivity, sn) as a function of false positive rate (1 - specificity, $1 - sp$). Area under the ROC curve is an easily interpretable quantity that has advantages over precision-recall curves in the positive-unlabeled setting because the class priors are difficult to estimate [17, 18].

4 Results

4.1 Performance on edge and hyperedge classification

We first evaluated the performance of hypergraphlet kernels in the task of predicting the types of interactions between pairs of proteins in a PPI network, as well as interaction types and modes of action between proteins and chemicals in DTI data. As described in Section 3.1 we first converted each input hypergraph to its dual hypergraph and then used the dual hypergraph for vertex classification. Table 2 lists AUC estimates for each method and each data set. Figure 3 shows ROC curves for one representative data set from each classification task and network type. Observe that hypergraphlet kernels achieved the highest AUCs on the three data sets, thus, outperforming all other methods. Therefore, these results provide evidence of the feasibility of this alternative approach to edge classification via exploiting hypergraph duality. Next, we evaluated the performance of the hypergraphlet kernels for predicting functional annotation of protein complexes in Corum. As shown in Table 2, traditional hypergraphlet kernel ($\tau = 0$) also performed favorably over the edit-distance kernels. In general, hypergraphlet kernels achieved the highest AUCs on both hyperedge classificatiton data sets over random walk kernels.
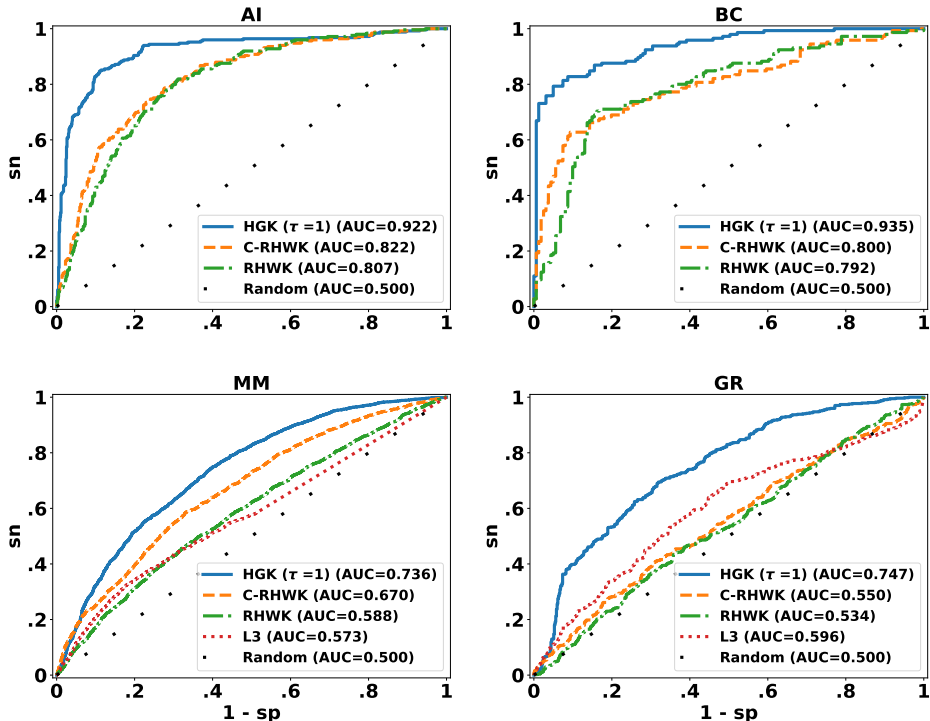
Figure 3: ROC curves for different kernel methods for four representative data sets: AI, BC, MM, and GR, as described in Table 1. Methods: RHWK = random hyperwalk kernel, C-RHWK = cumulative RHWK, HGK = hypergraphlet kernel, L3 = L3 framework.

4.2 Performance on link prediction

The performance of hypergraphlet kernels was further evaluated on the problem of link prediction on multiple PPI and DTI network data sets. Table 2 and Supplementary Table S2 show the performance accuracies for each hypergraph-based method across all link prediction data sets.

Table 2: Area under the ROC curve estimates for PPI/DTI data sets using 10-fold cross-validation. The highest performance for each data set is shown in boldface.

Method/Dataset	(Hyper)edge classification						Link prediction								
	PC	ID	AI	BC	MP	EC	SP	RN	MM	CE	AT	EZ	IC	GR	NR
Without domain information, $\lvert\Sigma\rvert = 1$															
Section hypergraphlet kernel ($\tau = 0$)	.647	.555	.529	.546	.585	.724	.639	.671	.628	.762	.578	.537	.542	.552	.783
Sub-hypergraphlet kernel ($\tau = 0$)	.838	.633	.675	.830	.787	.692	.691	.651	.612	**.873**	.681	.788	.711	.579	.732
Section hypergraphlet kernel ($\tau = 1$)	.649	.551	.521	.570	.599	.708	.641	.678	.628	.772	.581	.535	.550	.542	**.791**
Sub-hypergraphlet kernel ($\tau = 1$)	.834	.632	.672	.827	.767	.681	.686	.649	.611	.871	.683	.784	.702	.580	.736
L3 framework	-	-	-	-	-	.717	.577	.634	.573	.511	.501	.629	.642	.596	.505
Preferential attachment	-	-	-	-	-	.814	.492	.552	.498	.660	.449	.534	.561	.545	.497
With domain information, $\lvert\Sigma\rvert = \{4, 8, 16\}$ (classification) — $\Sigma = \Sigma_{\mathrm{GO}}$ (link prediction)															
	$\Sigma = \Sigma_{\mathrm{GO}}$	$\Sigma = \Sigma_{\mathrm{GO}} \bigcup \Sigma_{\mathrm{SS}}$		$\Sigma = \Sigma_{\mathrm{GO}}$					$\Sigma = \Sigma_{\mathrm{GO}}$				$\Sigma = \Sigma_{\mathrm{GO}} \bigcup \Sigma_{\mathrm{SS}}$		
Random hyperwalk	.741	.626	.807	.792	.936	.815	.610	.599	.588	.597	.608	.523	.543	.534	.548
Cumulative random hyperwalk	.711	.837	.822	.800	.937	.848	.681	.647	.670	.569	.637	.844	.617	.550	.518
Pairwise spectrum kernel ($k = \{3, 4, 5\}$)	.780	-	-	-	-	.816	.584	.716	.689	.754	.629	-	-	-	-
Section hypergraphlet kernel ($\tau = 0$)	.651	.634	.681	.663	.784	.861	.677	.698	.655	.804	.559	.633	.583	.625	.775
Sub-hypergraphlet kernel ($\tau = 0$)	**.940**	.936	.918	**.935**	**.952**	.796	.740	.676	.708	.846	.678	.936	.879	.724	.757
Section hypergraphlet kernel ($\tau = 1$)	.650	.635	.693	.674	.795	**.875**	.697	.706	.667	.786	.574	.637	.585	.619	.777
Sub-hypergraphlet kernel ($\tau = 1$)	**.940**	**.939**	**.922**	.930	**.952**	.812	**.750**	**.719**	**.736**	.860	**.728**	**.944**	**.884**	**.747**	.764

These results demonstrate good performance of our methods, with edit-distance kernels generally having the best performance. Interestingly, the *bona fide* hypergraphlet approaches displayed the highest accuracy on most PPI data sets (excluding CE), with a minor variation regarding the best method between section hypergraphlet and sub-hypergraphlet approaches. Both approaches have outperformed the state-of-the-art L3 link prediction framework as well as preferential attachment method and pairwise spectrum kernels. On the other hand, dual graphlet kernels (Supplementary Materials) showed the best performance on drug-target data sets suggesting that at this time, the increased resolution of modeling does not lead to increased performance on these networks (Supplementary Table S3). Note that pairwise spectrum kernels could not be applied to DTI data sets because they expect two pairs of objects of the same type as input, which further strengthens the appeal of our approach. Further breakdown of results based on the categories of difficulty identified by Park & Marcotte [40] is shown in Supplementary Tables S2 and S3.

4.3 Estimating interactome sizes

A positive-unlabeled formulation for link prediction presents an opportunity to estimate interactome sizes in biological networks. As a proof of concept, here we used the AlphaMax algorithm [14] for estimating class priors in positive-unlabeled learning to estimate both the number of missing links and false positive interactions in different PPI networks. To do this, we used BIND, BioGRID, DIP, HPRD, and IntAct to construct *Homo sapiens* and *Saccharomyces cerevisiae* PPI networks. The human network contained 10,841 nodes and 45,386 edges (10,729 and 45,327 in the largest connected component), whereas the yeast network contained 4,690 nodes and 26,165 edges (4,674 and 26,156 in the largest connected component). The negative examples were sampled uniformly randomly from the set of all possible edges, which was necessary in order to correctly estimate class priors and posteriors [14].

Assuming a tissue and cellular component agnostic model (i.e., any two proteins can interact), we obtained that the number of missing interactions on the largest component of the human PPI network is about 5% (i.e., approximately 2.5 million interactions), while the number of misannotated interactions is close to 11% which translates to about 4,985 false interactions. In the case of yeast, we computed that less than 1% of the potential protein interactions are missing which is close to 95,000. At the same time, the number of misannotated interactions is close to 13%, which is about 3,400 misannotated protein pairs. Some of these numbers fall within previous studies that suggest that the size of the yeast interactome is between 13,500 [41] and 137,000 [42]; however, the size of the human interactome is estimated to be within 130,000 [43] and 650,000 [41] interactions. A more recent paper by Lewis et al. [26] presents a scenario where yeast and human interactome size could reach 400,000 and over two million interactions, respectively. Although these results serve as a validation of our problem formulation and approach, additional tests and experiments, potentially involving exhaustive classifier and parameter optimization, will be necessary for more accurate and reliable estimates, especially for understanding the influence of potential biases within the PPI network data.

5 Related Work

The literature on the similarity-based measures for learning on hypergraphs is relatively scarce. Most studies revolve around the use of random walks for clustering that were first used in the field of circuit design [44]. Historically, typical hypergraph-based learning approaches can be divided into (1) tensor-based approaches, which extend traditional matrix (spectral) methods on graphs to higher-order relations for hypergraph clustering [44, 45], and (2) approximation-based

approaches that convert hypergraphs into standard weighted graphs and then exploit conventional graph clustering and semi-supervised learning [46, 47]. The methods from the first category provide a direct and mathematically rigorous treatment of hypergraph learning, although most tensor problems are NP-hard. As a consequence, this line of research remains largely unexplored despite a renewed interest in tensor decomposition approaches [48, 49]. Regarding the second category, there are two commonly used transformations for graph-based hypergraph approximation, reviewed and compared in Agarwal et al. [50].

Under a supervised learning framework, Wachman & Khardon [51] proposed walk-based hypergraph kernels on ordered hypergraphs, while Sun et al. [52] presented a hypergraph spectral learning formulation for multi-label classification. More recently, Bai et al. [53] introduced a hypergraph kernel that transforms a hypergraph into a directed line graph and computes a Weisfeiler-Lehman isomorphism test between directed graphs. A major drawback of most such approaches is that no graph representation fully captures the hypergraph structure.

6   Conslusions

This paper presents a learning framework for the problems of vertex classification, (hyper)edge classification, and link prediction in graphs and hypergraphs. The key ideas in our approach were (i) the use of hypergraph duality in order to cast each classification problem as an instance of vertex classification, and (ii) the use of a new family of kernels defined directly on labeled hypergraphs. Using the terminology of Bleakey et al. [54], our method belongs to the category of "local" learners. That is, it captures the structure of local neighborhoods, rooted at the vertex of interest, and should be distinguished from "global" models such as Markov Random Fields or diffusion kernels [55]. The body of literature on graph learning is vast; see Supplementary Materials for more details. We therefore selected to perform extensive comparisons against a limited set of methods that are most relevant to ours.

The development of hypergraphlet kernels derives from the graph reconstruction conjecture, an idea of using small graphs to probe large graphs [56, 57]. Hypergraphlet kernels prioritize accuracy over run time and are not an optimal choice for huge dense graphs where real-time performance is critical. However, biological networks (for now) are sparse and moderate in size, making accurate prediction to prioritize biological experiments an appealing choice. We additionally believe that unification of disparate prediction tasks on biological networks via hypergraph duality reduces the need for custom method development. Overall, our work provides evidence that hypergraph-based inference and hypergraphlet kernels are competitive with other approaches and readily deployable in practice.

References

[1] A. Grover and J. Leskovec. node2vec: Scalable feature learning for networks. In *Proc. 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, pages 855–864, New York, NY, USA, 2016. ACM.

[2] P. Goyal and E. Ferrara. Graph embedding techniques, applications, and performance: a survey. *Knowl-Based Syst*, 151:78–94, 2018.

[3] J. Xu and Y. Li. Discovering disease-genes by topological features in human protein-protein interaction network. *Bioinformatics*, 22(22):2800–2805, 2006.

[4] S. V. N. Vishwanathan et al. Graph kernels. *J Mach Learn Res*, 11:1201–1242, 2010.

[5] M. Deng et al. Prediction of protein function using protein-protein interaction data. *J Comput Biol*, 10(6):947–960, 2003.

[6] X. Zhu and Z. Ghahramani. Learning from labeled and unlabeled data with label propagation. In *Technical Report CMU-CALD-02-107*, Carnegie Mellon University, 2002.

[7] E. Nabieva et al. Whole-proteome prediction of protein function via graph-theoretic analysis of interaction maps. *Bioinformatics*, 21(Suppl 1):i302–i310, 2005.

[8] S. Klamt et al. Hypergraphs and cellular networks. *PLoS Comput Biol*, 5(5):1–6, 2009.

[9] T. Gaudelet et al. Higher-order molecular organization as a source of biological function. *Bioinformatics*, 34(17):i944–i953, 2018.

[10] C. Berge. Graphs and Hypergraphs. In C. Berge, editor, *North-Holland Mathematical Library*, volume 6. Elsevier, Amsterdam, Netherlands, 1973.

[11] F. Denis et al. Learning from positive and unlabeled examples. *Theor Comput Sci*, 348(16):70–83, 2005.

[12] T. Fawcett. An introduction to ROC analysis. *Pattern Recogn Lett*, 27:861–874, 2006.

[13] M. D. Reid and R. C. Williamson. Composite binary losses. *J Mach Learn Res*, 11:2387–2422, 2010.

[14] S. Jain et al. Estimating the class prior and posterior from noisy positives and unlabeled data. In *Proc. 30th Advances in Neural Information Processing Systems*, NIPS '16, pages 2693–2701, Red Hook, NY, USA, 2016. Curran Associates Inc.

[15] S. Jain et al. Nonparametric semi-supervised learning of class proportions. *arXiv preprint arXiv:1601.01944*, 2016.

[16] A. K. Menon et al. Learning from corrupted binary labels via class-probability estimation. In F. Bach and D. Blei, editors, *Proc. 32nd International Conference on Machine Learning*, ICML '15, pages 125–134, Lille, France, 2015. PMLR.

[17] S. Jain et al. Recovering true classifier performance in positive-unlabeled learning. In *Proc. 31st AAAI Conference on Artificial Intelligence*, AAAI '17, pages 2066–2072, Cambridge, MA, USA, 2017. AAAI Press.

[18] R. Ramola et al. Estimating classification accuracy in positive-unlabeled learning: characterization and correction strategies. *Pac Symp Biocomput*, 24:124–135, 2019.

[19] N. Przulj et al. Modeling interactome: scale-free or geometric? *Bioinformatics*, 20(18):3508–3515, 2004.

[20] N. Przulj. Biological network comparison using graphlet degree distribution. *Bioinformatics*, 23(2):e177–e183, 2007.

[21] N. Shervashidze et al. Efficient graphlet kernels for large graph comparison. In D. van Dyk and M. Welling, editors, *Proc. 12th International Conference on Artificial Intelligence and Statistics*, AISTATS '09, pages 488–495, Clearwater Beach, FL USA, 2009. PMLR.

[22] V. Vacic et al. Graphlet kernels for prediction of functional residues in protein structures. *J Comput Biol*, 17(1):55–72, 2010.

[23] J. Lugo-Martinez and P. Radivojac. Generalized graphlet kernels for probabilistic inference in sparse graphs. *Network Science*, 2(2):254–276, 2014.

[24] J. Shawe-Taylor and N. Cristianini. *Kernel methods for pattern analysis*. Cambridge University Press, Cambridge, MA, USA, 2001.

[25] A. Ben-Hur and W. S. Noble. Kernel methods for predicting protein-protein interactions. *Bioinformatics*, 21(Suppl 1):i38–i46, 2005.

[26] A. C. F. Lewis et al. What evidence is there for the homology of protein-protein interactions? *PLoS Comput Biol*, 8:1–14, 9 2012.

[27] C. von Mering et al. Comparative assessment of large-scale data sets of protein-protein interactions. *Nature*, 417(6887):399–403, 2002.

[28] S. M. Gomez et al. Learning to predict protein-protein interactions from protein sequences. *Bioinformatics*, 19(15):1875–1881, 2003.

[29] L. V. Zhang et al. Predicting co-complexed protein pairs using genomic and proteomic data integration. *BMC Bioinformatics*, 5(1):38, 2004.

[30] C-Y. Ma et al. Identification of protein complexes by integrating multiple alignment of protein interaction networks. *Bioinformatics*, 33(11):1681–1688, 2017.

[31] Y. Yamanishi et al. Prediction of drug-target interaction networks from the integration of chemical and genomic spaces. *Bioinformatics*, 24(13):i232, 2008.

[32] Y. Wang and J. Zeng. Predicting drug-target interactions using restricted boltzmann machines. *Bioinformatics*, 29(13):i126, 2013.

[33] M. Giurgiu et al. CORUM: the comprehensive resource of mammalian protein complexes. *Nucleic Acids Res*, 47(D1):D559–D563, 2019.

[34] W. T. Clark and P. Radivojac. Analysis of protein function and its prediction from amino acid sequence. *Proteins*, 79(7):2086–2096, 2011.

[35] M. Hattori et al. Development of a chemical structure comparison method for integrated analysis of chemical and genomic information in the metabolic pathways. *JACS*, 125(39):11853–11865, 2003.

[36] A. L. Barabási et al. Evolution of the social network of scientific collaborations. *Physica A*, 311(3):590–614, 2002.

[37] I. A. Kovács et al. Network-based prediction of protein interactions. *Nat Commun*, 10:1240–1247, 2019.

[38] T. Joachims. *Learning to classify text using support vector machines: methods, theory, and algorithms*. Kluwer Academic Publishers, Norwell, MA, USA, 2002.

[39] J. C. Platt. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. In A. Smola, P. Bartlett, B. Schölkopf, and D. Schuurmans, editors, *Advances in Large Margin Classifiers*, pages 61–74, Cambridge, MA, USA, 2000. MIT Press.

[40] Y. Park and E. M. Marcotte. Flaws in evaluation schemes for pair-input computational predictions. *Nat Methods*, 9(12):1134–1136, 2012.

[41] M. P. H. Stumpf et al. Estimating the size of the human interactome. *Proc Natl Acad Sci USA*, 105(19):6959–6964, 2008.

[42] H. Huang et al. Where have all the interactions gone? Estimating the coverage of two-hybrid protein interaction maps. *PLoS Comput Biol*, 3(11):1–20, 2007.

[43] K. Venkatesan et al. An empirical framework for binary interactome mapping. *Nat Methods*, 6:83–90, 2009.

[44] J. Cong et al. Random walks for circuit clustering. In *Proc. 4th International ASIC Conference*, ASIC '91, pages P14–2.1–P14–2.4, Piscataway, NJ, USA, 1991. IEEE.

[45] M. Leordeanu and C. Sminchisescu. Efficient hypergraph clustering. In N.D. Lawrence and M. Girolami, editors, *Proc. 15th International Conference on Artificial Intelligence and Statistics*, volume 22 of *AISTATS '12*, pages 676–684, La Palma, Canary Islands, 2012. PMLR.

[46] S. Agarwal et al. Beyond pairwise clustering. In *Proc. 18th Conference on Computer Vision and Pattern Recognition*, CVPR '05, pages 838–845, Piscataway, NJ, USA, 2005. IEEE.

[47] D. Zhou et al. Learning with hypergraphs: clustering, classification, and embedding. In *Proc. 19th Advances in Neural Information Processing Systems*, NIPS '06, pages 1601–1608, Cambridge, MA, USA, 2006. MIT Press.

[48] M. Hein et al. The total variation on hypergraphs - learning on hypergraphs revisited. In *Proc. 26th Advances in Neural Information Processing Systems*, NIPS '13, pages 2427–2435, Red Hook, NY, USA, 2013. Curran Associates Inc.

[49] P. Purkait et al. Clustering with hypergraphs: the case for large hyperedges. In *Proc. 13th European Conference on Computer Vision*, ECCV '14, pages 672–687, Piscataway, NJ, USA, 2014. IEEE.

[50] S. Agarwal et al. Higher order learning with graphs. In *Proc. 23rd International Conference on Machine Learning*, ICML '06, pages 17–24, New York, NY, USA, 2006. ACM.

[51] G. Wachman and R. Khardon. Learning from interpretations: a rooted kernel for ordered hypergraphs. In *Proc. 24th International Conference on Machine Learning*, ICML '07, pages 943–950, New York, NY, USA, 2007. ACM.

[52] L. Sun et al. Hypergraph spectral learning for multi-label classification. In *Proc. 14th International Conference on Knowledge Discovery and Data Mining*, KDD '08, pages 668–676, New York, NY, USA, 2008. ACM.

[53] L. Bai et al. A hypergraph kernel from isomorphism tests. In *Proc. 22nd International Conference on Pattern Recognition*, ICPR '14, pages 3880–3885, Piscataway, NJ, USA, 2014. IEEE.

[54] K. Bleakley et al. Supervised reconstruction of biological networks with local models. *Bioinformatics*, 23(13):i57–i65, 2007.

[55] R. I. Kondor and J. D. Lafferty. Diffusion kernels on graphs and other discrete structures. In *Proc. 19th International Conference on Machine Learning*, ICML '02, pages 315–322, Burlington, MA, USA, 2002. Morgan Kaufmann.

[56] J. A. Bondy and R. L. Hemminger. Graph reconstruction-a survey. *J Graph Theory*, 1(3):227–268, 1977.

[57] C. Borgs et al. Counting graph homomorphisms. In M. Klazar, J. Kratochvíl, M. Loebl, J. Matoušek, P. Valtr, and R. Thomas, editors, *Topics Discrete Math*, Algorithms and Combinatorics, pages 315–371, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.