SUBHYPERGRAPHS IN NON-UNIFORM RANDOM HYPERGRAPHS

MEGAN DEWAR[1], JOHN HEALY[1], XAVIER PÉREZ-GIMÉNEZ[2], PAWEŁ PRAŁAT[3], JOHN PROOS[1], BENJAMIN REINIGER[4], KIRILL TERNOVSKY[3]

ABSTRACT. In this paper we focus on the problem of finding (small) subhypergraphs in a (large) hypergraph. We use this problem to illustrate that reducing hypergraph problems to graph problems by working with the 2-section is not always a reasonable approach. We begin by defining a generalization of the binomial random graph model to hypergraphs and formalizing several definitions of subhypergraph. The bulk of the paper focusses on determining the expected existence of these types of subhypergraph in random hypergraphs. We also touch on the problem of determining whether a given subgraph appearing in the 2-section is likely to have been induced by a certain subhypergraph in the hypergraph. To evaluate the model in relation to real-world data, we compare model prediction to two datasets with respect to (1) the existence of certain small subhypergraphs, and (2) a clustering coefficient.

1. INTRODUCTION

Myriad problems can be described in hypergraph terms, however, the theory and tools are not sufficiently developed to allow most problems to be tackled directly within this context. In particular, we lack even the most basic of hypergraph models. In this paper we introduce a natural generalization of Erdős-Rényi (binomial) random graphs to non-uniform random hypergraphs. While such a model cannot hope to capture many features of real-world datasets, it allows us to explore several fundamental questions regarding the existence of subhypergraphs and helps us to illustrate that the common practice of reducing hypergraph problems to graph problems via the 2-section operation is not reasonable in many cases.

Despite being formally defined in the 1960s (and various realizations studied long before that) hypergraph theory is patchy and often not sufficiently general. The result is a lack of machinery for investigating hypergraphs, leading researchers and practitioners to create the 2-section graph of a hypergraph of interest and then rely upon well-established graph theoretic tools for analysis. In taking the 2-section (that is, making

[1] THE TUTTE INSTITUTE FOR MATHEMATICS AND COMPUTING, OTTAWA, ON, CANADA
[2] DEPARTMENT OF MATHEMATICS, UNIVERSITY OF NEBRASKA, LINCOLN, NE, USA
[3] DEPARTMENT OF MATHEMATICS, RYERSON UNIVERSITY, TORONTO, ON
[4] DEPARTMENT OF APPLIED MATHEMATICS, ILLINOIS INSTITUTE OF TECHNOLOGY, CHICAGO, IL, USA

E-mail address: tutte.institute+MeganDewar@gmail.com, tutte.institute+JohnHealy@gmail.com, xperez@unl.edu, pralat@ryerson.ca, tutte.institute+JohnProos@gmail.com, breiniger@iit.edu, kirill.ternovsky@ryerson.ca.

each hyperedge a clique, see Section 2 for a formal definition) we lose some information about edges of size greater than two. Sometimes losing this information does not affect our ability to answer questions of interest, but in other cases it has a profound impact.

Let us further explore the fundamental issues with working on the 2-section graph by considering a small example. Suppose we are given the coauthorship hypergraph in which vertices correspond to researchers and each hyperedge consists of the set of authors of a scientific paper. We wish to answer two questions regarding this dataset: (1) What is the Erdős number of every researcher (zero for Erdős, one for coauthors of Erdős, two for coauthors of coauthors of Erdős, etc.)? and (2) Given a subhypergraph induced by only the seminal papers in a particular field, what is the minimum set of authors who between them cover all the papers in the subhypergraph? The Erdős number of an author is the minimum distance between the author's vertex and Erdős' vertex in the hypergraph and this distance is not changed by taking the 2-section. However, in the case of finding a minimum set of vertices that are incident with every hyperedge in the subhypergraph, the 2-section of the hypergraph loses the information about the set of papers that a particular author covers. In fact, the 2-section does not even retain how many papers were used to create the hypergraph! Fundamentally, if the composition of the hyperedges of size greater than two is important in solving a problem, then solving the problem in the 2-section will be difficult, if not impossible.

Besides the information loss, there is another potential downside to working with the 2-section of a hypergraph; the 2-section can be much denser than the hypergraph since a single hyperedge of size k implies $\binom{k}{2}$ edges in the 2-section. Depending on the dataset and algorithm being executed, the increased density of the 2-section can have a significant detrimental effect on compute time.

In this paper we are interested in finding subhypergraphs in hypergraphs. We study rigorously – via theorems with proofs – occurrences of a given hypergraph as a subhypergraph of random hypergraphs generated by our model. One of the implications of our work is that two hypergraphs H_1 and H_2 that induce the same subgraph in the 2-section can have drastically different thresholds for appearance. This is concrete evidence that the research community needs to develop more algorithms that deal with hypergraphs directly.

To evaluate our model and, in particular, to illuminate features of real-world networks not captured by the model, we investigate two datasets: an email hypergraph and a coauthorship hypergraph. Not surprisingly, we confirm that subhypergraphs that are not distinguishable in the 2-section graph occur with different probabilities (as predicted by the model). However, also not surprisingly, we confirm that the distribution of such subhypergraphs in the two networks is quite different from what the model predicts. This is due in large part to the fact that the model predicts that edges occur independently. In graph theory the interdependence of edges is measured by the notion of clustering coefficient. There have been a number of proposals for generalizing clustering coefficient to hypergraphs including [17], [5], [2], [14], and [20]. We calculate the hypergraph clustering coefficient from [20] for our random hypergraph model and the two real networks we are investigating. This positions us well for further model development. An extended abstract of this paper appeared in [6].

2. Definitions and Conventions

2.1. **Random graphs and random hypergraphs.** First, let us recall a classic random graph model. A *binomial random graph* $G \in \mathscr{G}(n, p)$ is a random graph with vertex set $[n] := \{1, 2, \ldots, n\}$ in which every pair $\{i, j\} \in \binom{[n]}{2}$ appears independently as an edge with probability p. Note that $p = p(n)$ may (and usually does) tend to zero as n tends to infinity.

In this paper, we are concerned with a more general combinatorial object: hypergraphs. A *hypergraph H* is an ordered pair $H = (V, E)$, where V is a finite set (the *vertex set*) and E is a family of subsets of V (the *hyperedge set*). A hypergraph $H = (V, E)$ is *r-uniform* if all hyperedges of H are of size r. For a given $r \in \mathbb{N}$, a *random r-uniform hypergraph $H \in \mathscr{H}_r(n, p)$* has n labelled vertices from a vertex set $V = [n]$, with every subset $e \subseteq V$ of size $|e| = r$ chosen to be a hyperedge of H randomly and independently with probability p. For $r = 2$, this model reduces to the model $\mathscr{G}(n, p)$.

The binomial random graph model is well known and thoroughly studied (e.g. [4, 12, 11]). Random hypergraphs are much less well understood and most of the existing papers deal with uniform hypergraphs (e.g. Hamilton cycles (both tight ones and loose ones) were recently studied in [8, 9, 10], perfect matchings were investigated in [13] and additional examples can be found in a recent book on random graphs [11]).

In this paper, we study a natural generalization of the r-uniform random hypergraph model which produces non-uniform hypergraphs. Let $\boldsymbol{p} = (p_r)_{r \geq 1}$ be any sequence of numbers such that $0 \leq p_r = p_r(n) \leq 1$ for each $r \geq 1$. A *random hypergraph $H \in \mathscr{H}(n, \boldsymbol{p})$* has n labelled vertices from a vertex set $V = [n]$, with every subset $e \subseteq V$ of size $|e| = r$ chosen to be a hyperedge of H randomly and independently with probability p_r. In other words, $\mathscr{H}(n, \boldsymbol{p}) = \bigcup_{r \geq 1} \mathscr{H}_r(n, p_r)$ is a union of independent uniform hypergraphs.

Let us mention that there are several other natural generalizations that might be worth exploring, depending on the specific application in mind. One possible generalization would be to allow hyperedges to contain repeated vertices (multiset-hyperedge hypergraphs). Another would be to allow the hyperedges to be chosen with possible repetition, resulting in parallel hyperedges. We do not address these alternative formulations in this paper.

We require a few other definitions to aid our discussions. A vertex of a hypergraph is *isolated* if it is contained in no edge. In particular, a vertex of degree one that belongs only to an edge of size one is not isolated. The *2-section* of a hypergraph H, denoted $[H]_2$, is the graph on the same vertex set as H and an edge $\{u, v\}$ if (and only if) u and v are contained in some edge of H. In other words, the 2-section is obtained by making each hyperedge of H a clique in $[H]_2$. The *complete hypergraph* on n vertices is the hypergraph with all $2^n - 1$ possible nonempty edges.

2.2. **Notation.** All asymptotics throughout are as n goes to ∞. We emphasize that the notations $o(\cdot)$ and $O(\cdot)$ refer to functions of n, not necessarily positive, whose growth is bounded. We also use the notation $f \ll g$ for $f = o(g)$ and $f \gg g$ for $g = o(f)$. We say that an event in a probability space parametrized by n holds *asymptotically almost surely* (or *a.a.s.*) if the probability that it holds tends to 1 as $n \to \infty$. Since we aim

for results that hold a.a.s., we will always assume that n is large enough. We will often abuse notation by writing $\mathscr{G}(n,p)$ or $\mathscr{H}(n,\boldsymbol{p})$ to refer to a graph or hypergraph drawn from the distributions $\mathscr{G}(n,p)$ and $\mathscr{H}(n,\boldsymbol{p})$, respectively. For simplicity, we will write $f(n) \sim g(n)$ if $f(n)/g(n) \to 1$ as $n \to \infty$ (that is, when $f(n) = (1+o(1))g(n)$). Finally, we use $\log n$ to denote natural logarithms.

2.3. **Subhypergraphs.**

In this paper, we are concerned with occurrences of a given substructure in hypergraphs. As there are at least two natural generalizations of "subgraph" to hypergraphs, we cannot simply call these substructures "subhypergraphs".

A hypergraph $H' = (V', E')$ is a *strong subhypergraph* (called *hypersubgraph* by Bahmanian and Sajna [1] and *partial hypergraph* by Duchet [7]) of $H = (V, E)$ if $V' \subseteq V$ and $E' \subseteq E$; that is, each hyperedge of H' is also an hyperedge of H. We write $H' \subseteq_s H$ when H' is a strong subhypergraph of H. For $H = (V, E)$ and $V' \subseteq V$, the *strong subhypergraph of H induced by V'*, denoted $H_s[V']$, has vertex set V' and hyperedge set $E' = \{e \in E : e \subseteq V'\}$.

A hypergraph H' is a *weak subhypergraph* of H (called *subhypergraph* by Bahmanian and Sajna) if $V' \subseteq V$ and $E' \subseteq \{e \cap V' : e \in E\}$; that is, each hyperedge of H' can be extended to one of H by adding vertices of $V \setminus V'$ to it. For $V' \subseteq V$, the *weak subhypergraph induced by V'*, denoted $H_w[V']$, has vertex set V' and hyperedge set $E' = \{e \cap V' : e \in E\}$. Note that an induced weak subhypergraph might contain repeated edges and/or the empty edge. To simplify our analysis, we tacitly replace E' by $E' \setminus \{\emptyset\}$ and assume that weak subhypergraphs do not have multiple hyperedges (that is, E' is a set, not a multiset).

Note that when G is a (2-uniform) graph, strong subhypergraphs are the usual notion of subgraph, and weak subhypergraphs are subgraphs together with possible hyperedges of size one. Each strong subhypergraph is also a weak subhypergraph but the reverse is not true.

Given hypergraphs H_1 and H_2, a weak (resp. strong) *copy* of H_1 in H_2 is a weak (resp. strong) subhypergraph of H_2 isomorphic to H_1. Most of this paper is concerned with determining the existence of strong or weak copies of a fixed H in $\mathscr{H}(n,\boldsymbol{p})$. In a mild abuse of terminology, we will often say that a hypergraph contains H as a weak (strong) subhypergraph when we actually mean that the hypergraph contains a weak (strong) copy of H. The precise meaning will always be clear from the context.
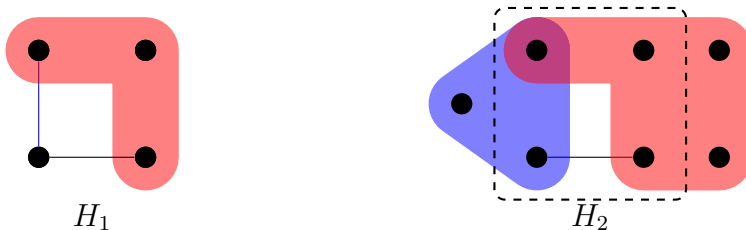

FIGURE 1. The hypergraph H_1 appears as a weak subhypergraph of H_2 (induced by the dashed vertex subset), but not as a strong subhypergraph.

Since it should not cause confusion in this paper, we will often drop the affix "hyper": we refer to hyperedges as just edges and strong/weak subhypergraphs as just strong/weak subgraphs. However, we will not drop the affix from "hypergraph."

3. SMALL SUBGRAPHS IN $\mathscr{H}(n, \boldsymbol{p})$.

We are interested in answering questions about the existence of subgraphs in $\mathscr{H}(n, \boldsymbol{p})$. This question was addressed for $\mathscr{G}(n, p)$ by Bollobás in [3]. To state his result we require two definitions. Let G be a graph. Denote by $d(G) = |E(G)|/|V(G)|$ the *density* of G, and by

$$m(G) = \max\{d(G') : G' \subseteq G\}$$

the *maximum subgraph density* of G.

Theorem 3.1 (Bollobás [3]). *For an arbitrary fixed graph G with at least one vertex,*

$$\lim_{n \to \infty} \mathbb{P}\left(G \subseteq \mathscr{G}(n, p)\right) = \begin{cases} 0 & \text{if } np^{m(G)} \to 0 \\ 1 & \text{if } np^{m(G)} \to \infty. \end{cases}$$

In other words, if $np^{m(G)} \to 0$, then a.a.s. $\mathscr{G}(n, p)$ does not contain G as a subgraph. If $np^{m(G)} \to \infty$, then a.a.s. $\mathscr{G}(n, p)$ contains G as a subgraph. The function $p_* = n^{-1/m(G)}$ (or any other function of the same asymptotic order) is called a *threshold* probability for the property that $\mathscr{G}(n, p)$ contains G as a subgraph.

Before we move to our result for random hypergraphs, let us mention why the maximum subgraph density of G, rather than simply the density of G, plays a role here. Consider the graphs G and G' depicted on Figure 2. Note that $G' \subset G$.

FIGURE 2. *Maximum subgraph density vs. density: G has density $6/5$, whilst its maximum subgraph density is $5/4$ ($> 6/5$). $G' \subset G$ has density $5/4$.*

Take any function $p = p(n)$ such that $n^{-5/6} \ll p \ll n^{-4/5}$, say $p(n) = n^{-9/11}$. Then

$$\mathbb{E}[X_n(G)] = \Theta(n^5 p^6) \to \infty$$
$$\mathbb{E}[X_n(G')] = \Theta(n^4 p^5) \to 0,$$

where $X_n(G)$ and $X_n(G')$ are random variables representing the number of copies of G and G' in $\mathscr{G}(n, p)$. Since $\mathbb{E}[X_n(G)] \to \infty$, one might expect many copies of G; however, using the first moment method we get that a.a.s. there is no copy of G' in $\mathscr{G}(n, p)$, and therefore a.a.s. there is no copy of G either.

We now generalize Theorem 3.1 to hypergraphs. In order to state our result, we need a few more definitions. Let $H = (V, E)$ be a hypergraph. Denote the number of vertices

in H by $v(H) = |V|$ and denote the number of edges by $e(H) = |E|$. For any $r \geq 1$, we will use $e_r(H) = |\{e \in E : |e| = r\}|$ to denote the number of edges of size r in H. Finally, define

$$\mu_s(H) = n^{v(H)} \prod_{r \geq 1} p_r^{e_r(H)}. \tag{1}$$

We are now ready to state our result on the appearance of strong subgraphs of $\mathscr{H}(n, \boldsymbol{p})$. We adopt the convention that $0^0 = 1$ and assume all our hypergraphs have nonempty vertex set.

Theorem 3.2. *Let H be an arbitrary fixed hypergraph. Let $\boldsymbol{p} = (p_r)_{r \geq 1}$ be any sequence such that $0 \leq p_r = p_r(n) \leq 1$ for each $r \geq 1$. Let $\mathcal{J}$ denote the family of all strong subgraphs of H.*

- (a) *If for some $H' \in \mathcal{J}$ we have $\mu_s(H') \to 0$ (as $n \to \infty$), then a.a.s. $\mathscr{H}(n, \boldsymbol{p})$ does not contain H as a strong subgraph.*
- (b) *If for all $H' \in \mathcal{J}$ we have $\mu_s(H') \to \infty$ (as $n \to \infty$), then a.a.s. $\mathscr{H}(n, \boldsymbol{p})$ contains H as a strong subgraph.*

Moreover, if there exists $\varepsilon > 0$ such that $p_r \leq 1 - \varepsilon$, for all r, then the conditions above determine whether or not H appears as an induced strong subgraph.

Let us mention that the result also holds in the multiset setting (i.e. when vertices are allowed to be repeated in each hyperedge with some multiplicity). This can be easily seen by replacing $\binom{n}{v(H)} v(H)!$ by $n^{v(H)}$ in the proof below and making several other trivial adjustments.

Proof. Denote by $X_n = X_n(H)$ the random variable that counts strong copies of H in a random hypergraph $\mathscr{H}(n, \boldsymbol{p})$. Denote by $H_1, H_2, \ldots, H_t$ all strong copies of H in the complete hypergraph on n vertices. Note that

$$t = t(H) = \binom{n}{v(H)} \frac{v(H)!}{\mathrm{aut}(H)} = \Theta\left(n^{v(H)}\right),$$

where $\mathrm{aut}(H)$ is the number of automorphisms of H. For $i \in [t]$, let

$$I_i = \begin{cases} 1 & \text{if } H_i \text{ is a strong subgraph of } \mathscr{H}(n, \boldsymbol{p}) \\ 0 & \text{otherwise,} \end{cases}$$

be an indicator random variable for the event that H_i is a strong subgraph of $\mathscr{H}(n, \boldsymbol{p})$. Then $X_n = \sum_{i=1}^{t} I_i$.

We start with part (a) of the statement. Let $H' \in \mathcal{J}$ with $\mu_s(H') \to 0$. It follows from Markov's inequality that

$$\begin{aligned} \mathbb{P}(X_n(H) > 0) \ &\leq \ \mathbb{P}(X_n(H') > 0) \ \leq \ \mathbb{E}[X_n(H')] \\ &= \ t(H') \prod_{r \geq 1} p_r^{e_r(H')} \ = \ \Theta\left(\mu_s(H')\right) = o(1). \end{aligned}$$

Hence a.a.s. $X_n(H) = 0$ and part (a) is done.

For part (b), we need to estimate the variance:

$$
\begin{aligned}
\mathbb{Var}[X_n(H)] &= \mathbb{Var}\left[\sum_{i=1}^{t} I_i\right] = \sum_{1 \leq i,j \leq t} \mathbb{Cov}(I_i, I_j) = \sum_{1 \leq i,j \leq t} \left(\mathbb{E}[I_i I_j] - (\mathbb{E}[I_i])(\mathbb{E}[I_j])\right) \\
&= \sum_{1 \leq i,j \leq t} \left(\mathbb{P}(I_i = 1, I_j = 1) - (\mathbb{P}(I_i = 1))^2\right) \\
&= \sum_{1 \leq i,j \leq t} \left(\mathbb{P}(I_i = 1, I_j = 1) - \prod_{r \geq 1} p_r^{2e_r(H)}\right).
\end{aligned}
$$

Observe that random variables I_i and I_j are independent if and only if H_i and H_j are edge-disjoint. In that case $\mathbb{P}(I_i = 1, I_j = 1) = \prod_{r \geq 1} p_r^{2e_r(H)}$, and such terms vanish from the above summation. Therefore, we may consider only the terms for which $e(H_i \cap H_j) \geq 1$. For each $H' \subseteq_s H$, there are $\Theta(n^{v(H')} n^{2(v(H) - v(H'))}) = \Theta(n^{2v(H) - v(H')})$ pairs (H_i, H_j) of copies of H in the complete hypergraph on n vertices with $H_i \cap H_j$ isomorphic to H'. Thus,

$$
\begin{aligned}
\mathbb{Var}[X_n(H)] &= \sum_{\substack{H' \subseteq_s H, \\ e(H') > 0}} \Theta(n^{2v(H) - v(H')}) \left(\prod_{r \geq 1} p_r^{2e_r(H) - e_r(H')} - \prod_{r \geq 1} p_r^{2e_r(H)}\right) \\
&= \sum_{\substack{H' \subseteq_s H, \\ e(H') > 0}} O\left(n^{2v(H) - v(H')} \prod_{r \geq 1} p_r^{2e_r(H) - e_r(H')}\right).
\end{aligned}
$$

Since $\mathbb{E}[X_n(H)] = \Theta(n^{v(H)} \prod_{r \geq 1} p_r^{e_r(H)})$, we can use the second moment method to get

$$
\mathbb{P}(X_n(H) = 0) \leq \frac{\mathbb{Var}[X_n(H)]}{(\mathbb{E}[X_n(H)])^2} = \sum_{\substack{H' \subseteq_s H, \\ e(H') > 0}} O\left(n^{-v(H')} \prod_{r \geq 1} p_r^{-e_r(H')}\right) = o(1).
$$

Note that there are a finite number of terms in the above sum and, by assumption, each term tends to zero as $n \to \infty$. Hence a.a.s. $X_n(H) \geq 1$ and part (b) is done. $\qquad\square$

In view of Theorem 3.2, we emphasize that the existence of strong copies of H in $\mathscr{H}(n, \boldsymbol{p})$ cannot be determined by translating to graphs via the 2-section. For instance, consider the three hypergraphs H_1, H_2, and H_3 in Figure 3. Each of these has H_1 as its 2-section. However, the expected number of strong copies of H_1, H_2 and H_3 in $\mathscr{H}(n, \boldsymbol{p})$ is $n^4 p_2^5$, $n^4 p_2^2 p_3$, and $n^4 p_3^2$, respectively. So if, say $p_3 = n^{-5/2}$ and $p_2 = n^{-3/4}$, then we expect many copies of H_1, a constant number of copies of H_2, and $o(1)$ copies of H_3. Moreover, by testing the conditions of Theorem 3.2 for all the strong subgraphs of H_1, H_2, and H_3, we obtain that a.a.s. $\mathscr{H}(n, \boldsymbol{p})$ contains H_1 as a strong subgraph, but not H_3 (and the theorem is inconclusive for H_2).

Next we consider the appearance of weak subgraphs of $\mathscr{H}(n, \boldsymbol{p})$. For technical reasons, we restrict ourselves to hypergraphs with bounded edge sizes. Formally, for a given $M \in \mathbb{N}$, we say that $H = (V, E)$ is an M-bounded hypergraph if $|e| \leq M$ for all
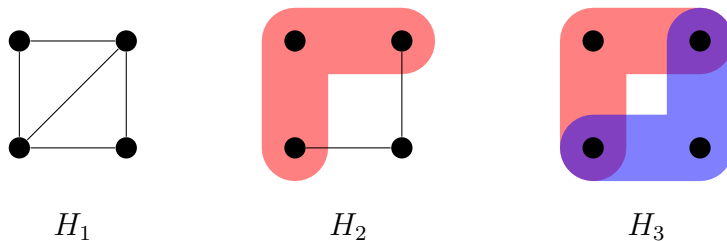
FIGURE 3. These three hypergraphs have the same 2-section, which is precisely H_1, but their behaviour as potential strong subgraphs of $\mathscr{H}(n, \boldsymbol{p})$ is different.

$e \in E$. Similarly, $\boldsymbol{p} = (p_r)_{r \geq 1}$ is an M-*bounded sequence* if $p_r = 0$ for $r > M$. We will use $\boldsymbol{p} = (p_r)_{r=1}^M$ for an M-bounded sequence instead of an infinite sequence $\boldsymbol{p} = (p_r)_{r \geq 1}$ with a bounded number of non-zero values. Clearly, if $\boldsymbol{p}$ is M-bounded, then so is $\mathscr{H}(n, \boldsymbol{p})$ (with probability 1). For $r \in [M]$, let

$$p'_r = p_r + n p_{r+1} + n^2 p_{r+2} + \cdots + n^{M-r} p_M, \tag{2}$$

and, given any fixed hypergraph H, define

$$\mu_w(H) = n^{v(H)} \prod_{r=1}^M (p'_r)^{e_r(H)}, \tag{3}$$

which will play a role analogous to that of $\mu_s(H)$.

Theorem 3.3. *Let H be an arbitrary fixed hypergraph, and let $\mathcal{J}$ be the family of all strong subgraphs of H. Let $\boldsymbol{p} = (p_r)_{r=1}^M$ be an M-bounded sequence.*

(a) *If for some $H' \in \mathcal{J}$ we have $\mu_w(H') \to 0$ (as $n \to \infty$), then a.a.s. $\mathscr{H}(n, \boldsymbol{p})$ does not contain H as a weak subgraph.*

(b) *If for all $H' \in \mathcal{J}$ we have $\mu_w(H') \to \infty$ (as $n \to \infty$), then a.a.s. $\mathscr{H}(n, \boldsymbol{p})$ contains H as a weak subgraph.*

Proof. Let $H' \in \mathcal{J}$ be a strong subgraph of H for which $n^{v(H')} \prod_{r=1}^M (p'_r)^{e_r(H')} \to 0$. As before, if there is more than one strong subgraph with this property, then choose one arbitrarily. Note that the product above still tends to 0 if all isolated vertices from H' are removed, so we can assume that H' has no isolated vertices. We will show that a.a.s. $\mathscr{H}(n, \boldsymbol{p})$ does not contain H' as a weak subgraph and so a.a.s. it does not contain H as a weak subgraph either. Let $\mathcal{F}$ be the family of all M-bounded hypergraphs (up to isomorphism) with precisely $|E(H')|$ edges, no isolated vertices, and containing H' as a weak subgraph. An important property is that each member of $\mathcal{F}$ has a bounded number of vertices (trivially $M|E(H')|$) and so $\mathcal{F}$ also has bounded size. Clearly, if $\mathscr{H}(n, \boldsymbol{p})$ contains H' as a weak subgraph, then $\mathscr{H}(n, \boldsymbol{p})$ contains some member of $\mathcal{F}$ as a strong subgraph.

Let us focus on any $F \in \mathcal{F}$. For any $e \in E(H')$, let $\bar{e} \in E(F)$ be the corresponding edge in F that e is obtained from; that is, $e = \bar{e} \cap V(H')$. Observe that

$$
\begin{aligned}
\mu_s(F) = n^{v(F)} \prod_{r=1}^{M} p_r^{e_r(F)} &= n^{v(F)} \prod_{\bar{e} \in E(F)} p_{|\bar{e}|} = n^{v(F)} \prod_{e \in E(H')} p_{|\bar{e}|} \\
&= n^{v(F) - \sum_{e \in E(H')} (|\bar{e}| - |e|)} \prod_{e \in E(H')} n^{|\bar{e}| - |e|} p_{|\bar{e}|} \\
&\leq n^{v(H')} \prod_{e \in E(H')} p'_{|e|} = \mu_w(H') \to 0.
\end{aligned}
$$

It follows from Theorem 3.2(a) that a.a.s. $\mathscr{H}(n, \boldsymbol{p})$ does not contain F as a strong subgraph. As $|\mathcal{F}|$ is bounded, a.a.s. $\mathscr{H}(n, \boldsymbol{p})$ does not contain any member of $\mathcal{F}$ as a strong subgraph and so a.a.s. $\mathscr{H}(n, \boldsymbol{p})$ does not contain H as a weak subgraph. Part (a) is finished.

Let us move to part (b). It follows immediately from the definition of p'_r that for each r there exists $i(r) \in \mathbb{Z}$, $0 \leq i(r) \leq M - r$, such that $n^{i(r)} p_{r+i(r)} \geq p'_r / (M - r + 1) \geq p'_r / M$. We construct a new hypergraph J from H as follows: for each $e \in E(H)$, add $i(|e|)$ new vertices to $V(J)$ and add them to e to form $\bar{e} \in E(J)$. See Figure 4 for an example of this construction. Our goal is to show that a.a.s. $\mathscr{H}(n, \boldsymbol{p})$ contains J as a strong subgraph, which will finish the proof as it implies that a.a.s. $\mathscr{H}(n, \boldsymbol{p})$ contains H as a weak subgraph.
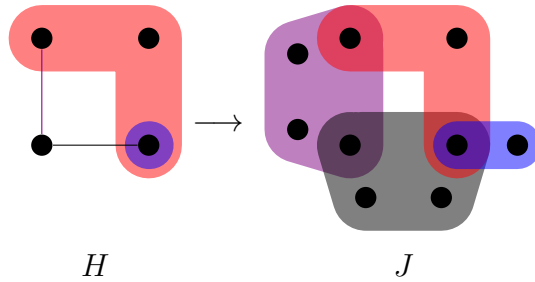

H J

FIGURE 4. The construction of J from H; here $i(1) = 1$, $i(2) = 2$, $i(3) = 0$.

Let J' be any strong subgraph of J, and let H' be obtained from J' as follows: $V(H') = V(J') \cap V(H)$ and $E(H') = \{\bar{e} \cap V(H) : \bar{e} \in E(J')\}$ (i.e., H' is the weak subgraph of J' induced by $V(H)$). Note that $E(H') \subseteq E(H)$ and so H' is a strong

subgraph of H. This time, observe that

$$
\begin{aligned}
\mu_s(J') = n^{v(J')} \prod_{r=1}^{M} p_r^{e_r(J')} &= n^{v(J')} \prod_{\bar{e} \in E(J')} p_{|\bar{e}|} = n^{v(J')} \prod_{e \in E(H')} p_{|\bar{e}|} \\
&= n^{v(J') - \sum_{e \in E(H')} i(|e|)} \prod_{e \in E(H')} n^{i(|e|)} p_{|\bar{e}|} \geq n^{v(H')} \prod_{e \in E(H')} \frac{p'_{|e|}}{M} \\
&= \frac{n^{v(H')}}{M^{e(H')}} \prod_{r=1}^{M} (p'_r)^{e_r(H')} = \Omega\big(\mu_w(H')\big) \to \infty.
\end{aligned}
$$

It follows from Theorem 3.2(b) that a.a.s. $\mathcal{H}(n, \boldsymbol{p})$ contains J as a strong subgraph. Part (b) and the proof of the theorem is finished. $\square$

We shall discuss a few details concerning Theorem 3.3. First, it is possible that a.a.s. some graph occurs as a weak subgraph but not as a strong one. For example, if

$$
p_1 = n^{-0.6}, \quad p_2 = n^{-0.9}, \quad p_3 = n^{-1.7}, \quad \text{and} \quad p_4 = n^{-3.1}, \tag{4}
$$

then a.a.s. $\mathcal{H}(n, \boldsymbol{p})$ does not contain H as a strong subgraph but a.a.s. it does contain J (both presented in Figure 4) and therefore a.a.s. it contains H as a weak subgraph. Next, observe that if we replace the collection of all strong subgraphs $\mathcal{J}$ in the statement of Theorem 3.3 by the collection $\mathcal{J}_w$ of all *weak* subgraphs of H, the theorem remains valid. This is trivially true for part (b), since $\mathcal{J}_w \supseteq \mathcal{J}$. For part (a), simply replace $H' \in \mathcal{J}$ by $H' \in \mathcal{J}_w$ in the proof, and note that the argument follows. Finally, let us comment on the definition of p'_r, and introduce related parameters p''_r and p'''_r, which will play a role later on. Our particular choice of p'_r in (3) and thus in the statement of Theorem 3.3 is the simplest function from the equivalence class of all functions of the same order. However, it is arguably more natural to replace p'_r with p''_r which is asymptotically the expected number of edges to which a given set of size r belongs. For $r \in [M]$, let

$$
p''_r = p_r + n p_{r+1} + \binom{n}{2} p_{r+2} + \cdots + \binom{n}{M-r} p_M. \tag{5}
$$

Note that p'_r and p''_r are of the same order. More precisely,

$$
(1 + o(1)) \frac{p'_r}{(M-r)!} \leq p''_r \leq p'_r.
$$

Hence, p'_r can be replaced in (3) by the more natural (but less simple) p''_r, and Theorem 3.3 remains valid. It is worth noting that both p'_r and p''_r can be greater than one or even tend to infinity as $n \to \infty$. While p''_r is not a probability, we can create a probabalistic version, p'''_r, that represents the probability that a set of size r belongs to some edge:

$$
p'''_r = 1 - (1 - p_r)(1 - p_{r+1})^{n-r} (1 - p_{r+2})^{\binom{n-r}{2}} \cdots (1 - p_M)^{\binom{n-r}{M-r}}. \tag{6}
$$

Observe that, if $p'_r = o(1)$ (or equivalently $p''_r = o(1)$), then $p''_r, p''_{r+1}, \ldots, p''_M = o(1)$, and therefore

$$
\begin{aligned}
p'''_r &= 1 - \exp\left(-(1+o(1))\left(p_r + np_{r+1} + \binom{n}{2}p_{r+2} + \cdots + \binom{n}{M-r}p_M\right)\right) \\
&= 1 - \exp\left(-(1+o(1))p''_r\right) \sim p''_r,
\end{aligned}
\tag{7}
$$

so p''_r and p'''_r asymptotically coincide.

4. Induced weak subgraphs

Let us discuss how one can use Theorem 3.3 to determine whether H appears as an *induced* weak subgraph of $\mathscr{H}(n,\boldsymbol{p})$. This seems to be a more complex question than in the case of strong subgraphs. Trivially, if H is a complete hypergraph on k vertices, then every weak copy of H in $\mathscr{H}(n,\boldsymbol{p})$ is automatically also induced. Otherwise, the non-edges of H play a crucial role in determining the existence of induced weak copies. Indeed, a weak subgraph H of $\mathscr{H}(n,\boldsymbol{p})$ is induced provided that, for every set e of vertices of H that do not form an edge, e cannot be extended to an edge of $\mathscr{H}(n,\boldsymbol{p})$ by adding vertices not in H.

First, we will give some conditions that forbid a.a.s. the existence of induced weak copies of H in $\mathscr{H}(n,\boldsymbol{p})$ (even if H *does* appear as a weak subgraph).

Proposition 4.1. *Let H be an arbitrary fixed hypergraph on k vertices with a non-edge of size r $(1 \le r \le k)$. Suppose $p''_r \ge (k+\varepsilon)\log n$ for some constant $\varepsilon > 0$, then a.a.s. H does not occur as an induced weak subgraph of $\mathscr{H}(n,\boldsymbol{p})$.*

Proof. If $\mathscr{H}(n,\boldsymbol{p})$ contains a copy H_1 of H as an induced weak subgraph, then there must be a set of r vertices in H_1 that cannot be extended to an edge of $\mathscr{H}(n,\boldsymbol{p})$ by only adding vertices outside of H_1. The expected number of such sets is

$$
\binom{n}{k}\binom{k}{r}(1-p_r)(1-p_{r+1})^{n-k}\cdots(1-p_M)^{\binom{n-k}{M-r}} \le n^k \exp(-(1+o(1))p''_r) = o(1),
$$

so a.a.s. there are none. $\qquad\square$

Proposition 4.1 implies that if H is an induced weak subgraph of $\mathscr{H}(n,\boldsymbol{p})$ of order k and $p''_r \ge (k+\varepsilon)\log n$, then H must contain all possible edges of size r. Let us return to our example hypergraph H in Figure 4 and the probabilities in (4). Note that $p''_1 \sim \binom{n}{2}p_3 \sim n^{0.3}/2$, thus, a.a.s. H does not occur as an induced weak subgraph of $\mathscr{H}(n,\boldsymbol{p})$, as not every vertex of H belongs to an edge of size one.

On the other hand, suppose that $r \ge 1$ is the size of the smallest non-edge of H and assume that

$$
\max\{p'''_r, p'''_{r+1}, \ldots, p'''_M\} \le 1 - \varepsilon
\tag{8}
$$

for some constant $\varepsilon > 0$. Then any given weak copy of H in $\mathscr{H}(n,\boldsymbol{p})$ is also induced with probability bounded away from zero. In this case, the calculations in the proof of Theorem 3.3 are still valid – with an extra $\Theta(1)$ factor – and thus the conclusions of that theorem extend to induced weak subgraphs. Since verifying condition (8) may sometimes be tedious, we give a simpler sufficient condition.

Proposition 4.2. *Let H be an arbitrary fixed hypergraph, and let r be the size of its smallest non-edge. Suppose that $p_r \leq 1 - \varepsilon$ for some constant $\varepsilon > 0$ and that $p'_r = O(1)$ (and, as a result, $p''_r = O(1)$ too). If the conditions in part (b) of Theorem 3.3 are satisfied, then a.a.s. $\mathscr{H}(n, \boldsymbol{p})$ contains H as an induced weak subgraph.*

Proof. Since $p'_{j+1} = (p'_j - p_j)/n \leq p'_j/n$, we inductively get that $p'_j \leq 1 - \varepsilon$ for all j with $r + 1 \leq j \leq M$. In particular, $p_j \leq 1 - \varepsilon$ and $p''_j = O(1)$ for $r \leq j \leq M$. In view of this, for all such j,

$$1 - p'''_j = (1 - p_j)(1 - p_{j+1})^{n-j}(1 - p_{j+2})^{\binom{n-j}{2}} \cdots (1 - p_M)^{\binom{n-j}{M-j}}$$

$$\geq \varepsilon \exp\left(-(1 + o(1))\left(np_{j+1} + \binom{n}{2}p_{j+2} + \ldots + \binom{n}{M-j}p_M\right)\right)$$

$$\geq \varepsilon \exp\left(-(1 + o(1))p''_j\right) > \varepsilon',$$

for some constant $\varepsilon' > 0$, and thus (8) holds. Hence the conclusion of Theorem 3.3 extends to weak induced subgraphs. In particular, part (b) of that theorem gives a sufficient condition for the a.a.s. existence of weak induced copies of H. $\qquad\square$

Let us return to our example from Figure 4 and the probabilities given in (4) for the last time. Since $p''_2 \sim np_3 = n^{-0.7} = o(1)$, if the "missing" edges of size 1 are added to H, then a.a.s. the resulting graph occurs as an induced weak subgraph of $\mathscr{H}(n, \boldsymbol{p})$.

5. THE 2-SECTION OF $\mathscr{H}(n, \boldsymbol{p})$

We begin by considering the question of whether a given graph G appears as a subgraph of the 2-section of $\mathscr{H}(n, \boldsymbol{p})$. Again we assume that G has no isolated vertices.

Let us start with some general observations that apply to any host hypergraph $\mathscr{H}$, not necessarily $\mathscr{H}(n, \boldsymbol{p})$. Observe that $G \subseteq [\mathscr{H}]_2$ if and only if there is a weak subgraph H of $\mathscr{H}$ such that G is a spanning subgraph of $[H]_2$. So we may test for $G \subseteq [\mathscr{H}]_2$ by finding every hypergraph H with G a spanning subgraph of $[H]_2$ and applying Theorem 3.3 to each. We can reduce the number of hypergraphs that need to be tested: if H_1 is a weak subgraph of H_2 and H_2 is a weak subgraph of $\mathscr{H}$, then H_1 is also a weak subgraph of $\mathscr{H}$. Note also that a spanning weak subgraph is actually a strong subgraph; it suffices to check only the minimal hypergraphs H (with respect to the (strong) subgraph relation) that have G as a spanning subgraph of their 2-section.

In $\mathscr{H}(n, \boldsymbol{p})$ one can reduce the number of hypergraphs H to be tested even further. Given a hypergraph H, we construct a new hypergraph H' on the same vertices and form hyperedges by taking a subset of each hyperedge of H. Any strong subgraph of H' is called a *subedge system of H*. Note that if H_1 is a subedge system of H_2 and H_2 is a weak subgraph of H, it is not necessarily true that H_1 is a weak subgraph of H, but it is true a.a.s. for $H = \mathscr{H}(n, \boldsymbol{p})$.

Proposition 5.1. *Let H_1 and H_2 be fixed hypergraphs with H_1 a spanning subedge system of H_2, and let $\boldsymbol{p}$ be M-bounded. Let $\mathcal{J}_1$ and $\mathcal{J}_2$ denote the set of all strong subgraphs of H_1 and H_2, respectively. If every $H'_2 \in \mathcal{J}_2$ satisfies $\mu_w(H'_2) \to \infty$, then every $H'_1 \in \mathcal{J}_1$ also satisfies $\mu_w(H'_1) \to \infty$.*

Proof. Let $H_1' \in \mathcal{J}_1$. Since edges of H_1' are subsets of edges of H_2, we can turn H_1' into a strong subgraph H_2' of H_2 by appropriately extending some of its edges. For each edge e_1 of H_1' of size r that is extended to an edge e_2 of size $r + i$, one factor of p_r' in $\mu_w(H_1')$ is replaced by one corresponding factor of p_{r+i}' in $\mu_w(H_2')$. We have that $p_r' \geq n^i p_{r+i}' \geq p_{r+i}'$, so $\mu_w(H_1') \geq \mu_w(H_2') \to \infty$. $\qquad\square$

Corollary 5.2. *Fix a graph G without isolated vertices. Let $\mathcal{F}$ denote the family of minimal (with respect to the subedge system relation) hypergraphs containing G in their 2-section. Let $\boldsymbol{p}$ be M-bounded.*

(a) *If for every $H \in \mathcal{F}$ there is some strong subgraph $H' \subseteq_s H$ with $\mu_w(H') \to 0$, then a.a.s. G is not a subgraph of $[\mathscr{H}(n, \boldsymbol{p})]_2$.*

(b) *If for some $H \in \mathcal{F}$ every strong subgraph $H' \subseteq_s H$ satisfies $\mu_w(H') \to \infty$, then a.a.s. G is a subgraph of $[\mathscr{H}(n, \boldsymbol{p})]_2$.*

Next we consider the following question: suppose a copy of G is found in $[\mathscr{H}(n, \boldsymbol{p})]_2$. What is the probability that this copy comes from a given weak subgraph of $\mathscr{H}(n, \boldsymbol{p})$?

Let G be a fixed graph with no isolated vertices. Let $\mathcal{F}$ denote the family of hypergraphs H on the same vertex set as G such that $G \simeq [H]_2$. Then, G appears as an induced subgraph of $[\mathscr{H}(n, \boldsymbol{p})]_2$ if and only if some $H \in \mathcal{F}$ appears as an induced weak subgraph of $\mathscr{H}(n, \boldsymbol{p})$. More precisely, for every set of vertices S inducing a copy of G in $[\mathscr{H}(n, \boldsymbol{p})]_2$, there is exactly one $H \in \mathcal{F}$ such that S induces a weak copy of H in $\mathscr{H}(n, \boldsymbol{p})$. In this case we say that the hypergraph H *originates* that particular copy of G. As a result we have the following corollary.

Proposition 5.3. *Let $\boldsymbol{p} = (p_r)_{r=1}^M$ be an M-bounded sequence. For $r \in [M]$, let p_r''' be defined as in (6). Given a copy of G in $[\mathscr{H}(n, \boldsymbol{p})]_2$, the probability that it originates from a given $H \in \mathcal{F}$ is*

$$(1 + o(1)) \frac{\mathrm{aut}(H) \prod_{r=1}^M (p_r''')^{e_r(H)} (1 - p_r''')^{\binom{v(G)}{r} - e_r(H)}}{\sum_{H' \in \mathcal{F}} \mathrm{aut}(H') \prod_{r=1}^M (p_r''')^{e_r(H')} (1 - p_r''')^{\binom{v(G)}{r} - e_r(H')}}.$$

Define the *signature* of $H \in \mathcal{F}$ as the vector $\boldsymbol{e}(H) = (e_1(H), e_2(H), \ldots, e_k(H))$, where $k = v(G)$ (and hence also $k = v(H)$). Let $\boldsymbol{e}(\mathcal{F}) = \{\boldsymbol{e}(H) : H \in \mathcal{F}\}$. For a given signature $\boldsymbol{e} \in \boldsymbol{e}(\mathcal{F})$, let $\mathcal{F}_{\boldsymbol{e}} \subseteq \mathcal{F}$ be the family of hypergraphs in $\mathcal{F}$ with signature $\boldsymbol{e}$. Notice that $\{\mathcal{F}_{\boldsymbol{e}} : \boldsymbol{e} \in \boldsymbol{e}(\mathcal{F})\}$ is a partition of $\mathcal{F}$. We can state the following corollary to Proposition 5.3 which we will make use of when comparing model predictions to real-world networks (see Section 6).

Corollary 5.4. *Let $\boldsymbol{p} = (p_r)_{r=1}^M$ be an M-bounded sequence. For $r \in [M]$, let p_r''' be defined as in (6). Then, given a copy of G in $[\mathscr{H}(n, \boldsymbol{p})]_2$, the probability that it originates from a hypergraph with a given signature $\boldsymbol{e} = (m_1, m_2, \ldots, m_k) \in \boldsymbol{e}(\mathcal{F})$ is*

$$(1 + o(1)) \frac{\sum_{H \in \mathcal{F}_{\boldsymbol{e}}} \mathrm{aut}(H) \prod_{r=1}^k (p_r''')^{m_r} (1 - p_r''')^{\binom{v(G)}{r} - m_r}}{\sum_{H' \in \mathcal{F}} \mathrm{aut}(H') \prod_{r=1}^k (p_r''')^{e_r(H')} (1 - p_r''')^{\binom{v(G)}{r} - e_r(H')}}.$$

We are particularly interested in the subgraphs induced by the action of 2-sectioning, i.e. complete graphs. Of course, these include as subgraphs all sparser graphs on the

same or fewer vertices. Let us briefly explore the implications of Corollary 5.4 when G is a complete graph. Let K_k denote the complete graph on k vertices, $k \geq 2$, and suppose that $\boldsymbol{p}$ is an M-bounded sequence satisfying $\binom{n}{j}p_j = O(n)$ for all $j \in [M]$. The latter condition is equivalent to assuming that the expected number of edges of each given size is at most linear in the number of vertices, which is a fairly reasonable assumption for many hypergraph networks. Additionally, suppose that for some r with $k \leq r \leq M$ we also have $\binom{n}{r}p_r = \Omega(n)$. From (7), we obtain that $p_j''' = O(1/n^{j-1})$ for every $j \in [M]$ and $p_k''' = \Theta(1/n^{k-1})$. Consider the signature $\widehat{\boldsymbol{e}} = (0,\ldots,0,1)$ corresponding to the hypergraph $\widehat{H}$ on k vertices with a single edge of size k. A straightforward inductive argument reveals that, for any signature $\boldsymbol{e} = (m_1, m_2, \ldots, m_k) \in \boldsymbol{e}(\mathcal{F})$,

$$\prod_{r=1}^{k}(p_r''')^{m_r}(1 - p_r''')^{\binom{k}{r}-m_r} = \begin{cases} (1 + o(1))p_k''' = \Theta(1/n^{k-1}) & \text{if } \boldsymbol{e} = \widehat{\boldsymbol{e}} \\ o(1/n^{k-1}) & \text{if } \boldsymbol{e} \neq \widehat{\boldsymbol{e}}. \end{cases}$$

As a result, applying Corollary 5.4 to all signatures different from $\widehat{\boldsymbol{e}}$, we conclude that, for a given copy of K_k in $[\mathscr{H}(n, \boldsymbol{p})]_2$, a.a.s. it must originate from $\widehat{H}$.

6. Comparing the model with reality

In this section, we look at two real-world datasets that are naturally represented as hypergraph networks. We consider how well our model captures certain features of these datasets by looking at the appearance of select weak subhypergraphs and by comparing a measure of clustering coefficient.

6.1. **Real-world datasets.** We examine two real-world datasets; a coauthorship hypergraph and an email hypergraph.

The **email hypergraph** was constructed from the Enron dataset; a version of which can be obtained from Carnegie Mellon University [16]. This dataset consists of 30,100 email messages from 151 Enron employees. We use the messages sent by these individuals to build an undirected hypergraph for analysis. From each message we extract the *from*, *to*, *cc* and *bcc* fields. The fields are merged (removing repeated addresses) and the resulting set is treated as an undirected hyperedge. We recognize that this data might be better represented as directed hyperedges, but that is outside the scope of this paper and may be considered in future work. For the purpose of this paper we also ignore the effects of multiple identical hyperedges, leaving us with 11,407 unique undirected hyperedges. The distributions of degrees and edge sizes can be seen in Figure 5(a) and Figure 5(b). Note that the degree of a vertex is defined to be the number of hyperedges the vertex is contained in, while the edge size is defined to be the number of vertices in the hyperedge. These two definitions are unambiguous in this case as edges do not contain repeated vertices.

The **coauthorship hypergraph** was generated from the ArnetMiner (AMiner) dataset. ArnetMiner [18] is a project which attempts to extract researcher social networks from the World Wide Web. Of particular interest to us is that it integrates a number of existing digital libraries. The dataset built with this tool is available for research purposes at [19]. It consists of 2,092,356 research papers from 1,712,433 authors.

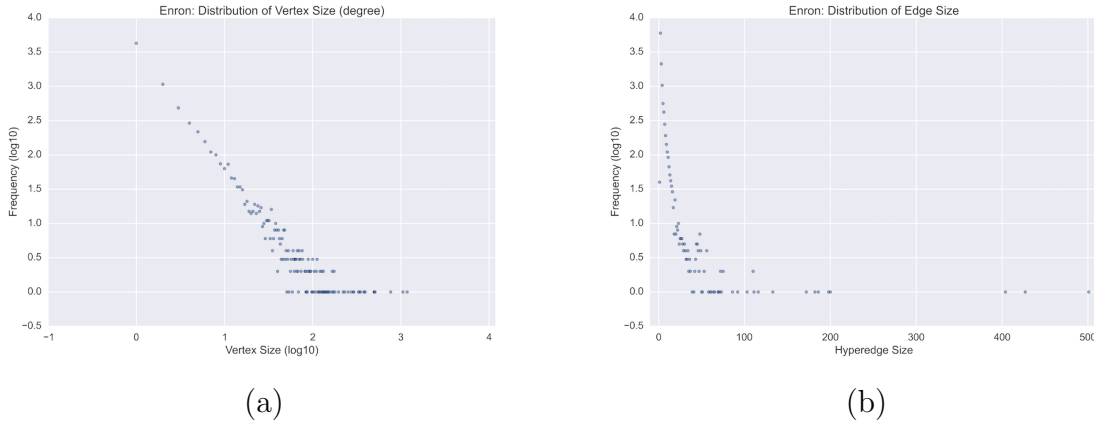(a)   (b)

FIGURE 5. Degree distribution (a) and edge-size distribution (b) of the email hypergraph from the Enron dataset.

From this we construct a coauthorship hypergraph with vertices representing authors and with one hyperedge for each paper in the database. Hyperedges are undirected and consist of the set of authors listed on each paper. Again, we eliminate duplicate edges resulting in a hypergraph containing 1,499,404 unique hyperedges. The distributions of degrees and edge sizes can be seen in Figure 6(a) and Figure 6(b).

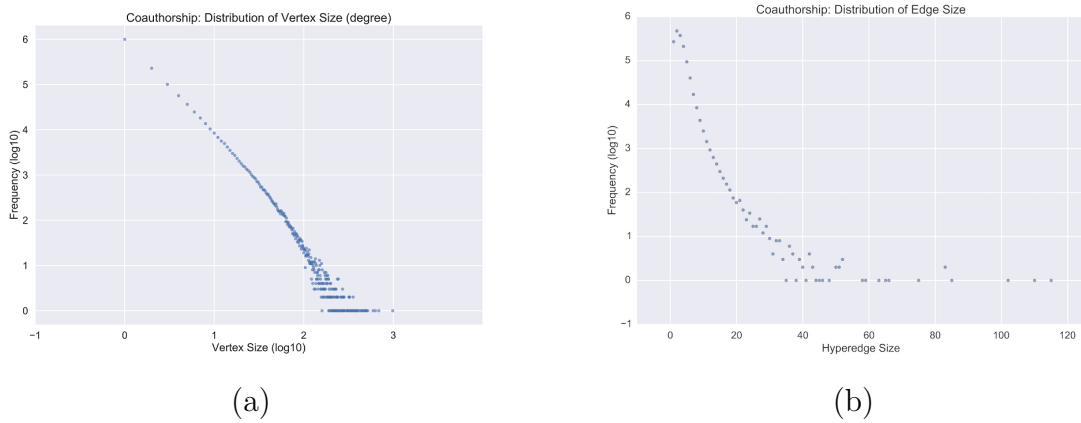(a)   (b)

FIGURE 6. Degree distribution (a) and edge-size distribution (b) of the coauthorship hypergraph from the AMiner dataset.

In our analysis we will consider truncated hypergraphs of those described above. The main reason is that large hyperedges have a significant effect on computation time and can drown out the signal of smaller, more interesting, effects. In particular, large hyperedges can have a significant impact on the appearance of small complete graphs in the 2-section. For example, each edge of size 100 introduces $\binom{100}{k}$ copies of K_k. There are sporadic hyperedges of such size in both networks we investigate.

Recall that m_i denotes the number of hyperedges of size i. In the email hypergraph we have $m_2 = 5{,}975$, $m_3 = 2{,}128$, $m_4 = 1{,}034$, and $m_5 = 561$. After removing hyperedges of size greater than or equal to 6, we were left with a hypergraph, $\mathscr{E}_5$, on $n_5 = 5{,}044$ vertices and having $m = \sum_{i=2}^{5} m_i = 9{,}698$ edges. After removing hyperedges of size greater than or equal to 5, we were left with a hypergraph, $\mathscr{E}_4$, on $n_4 = 4{,}919$ vertices and having $m = \sum_{i=2}^{4} m_i = 9{,}137$ edges. In the coauthorship hypergraph we have $m_2 = 473{,}560$, $m_3 = 373{,}262$, $m_4 = 211{,}011$, and $m_5 = 94{,}168$. After removing hyperedges of size greater than or equal to 6, we were left with a hypergraph, $\mathscr{D}_5$, on $n_5 = 1{,}264{,}602$ vertices having $m = \sum_{i=2}^{5} m_i = 1{,}152{,}001$ edges. After removing hyperedges of size greater than or equal to 5, we were left with the hypergraph $\mathscr{D}_4$ on $n_4 = 1{,}146{,}130$ vertices having $m = \sum_{i=2}^{4} m_i = 1{,}057{,}833$ edges.

Note that all isolated vertices were also removed from the networks.

6.2. **Creating the model hypergraphs.** For each real-world hypergraph we analyse, we create model hypergraphs having the same expected edge counts. That is, if m_i is the number of edges of size i in the hypergraph we wish to model, then we create a random hypergraph $\mathscr{H}(n, \boldsymbol{p})$ on n vertices where each i-set forms a hyperedge of size i with probability $p_i = m_i / \binom{n}{i}$. Note that the model hypergraph created has (an expected number of) m_i hyperedges of size i randomly distributed throughout the graph. For our comparisons we generated 1048 model hypergraphs for each of the four real-world datasets.

In Figure 7 we compare the degree distribution of the email hypergraph with edges of maximum size 5, $\mathscr{E}_5$ (a), and the coauthorship hypergraph with edges of maximum size 5, $\mathscr{D}_5$ (b), to their corresponding model hypergraphs.

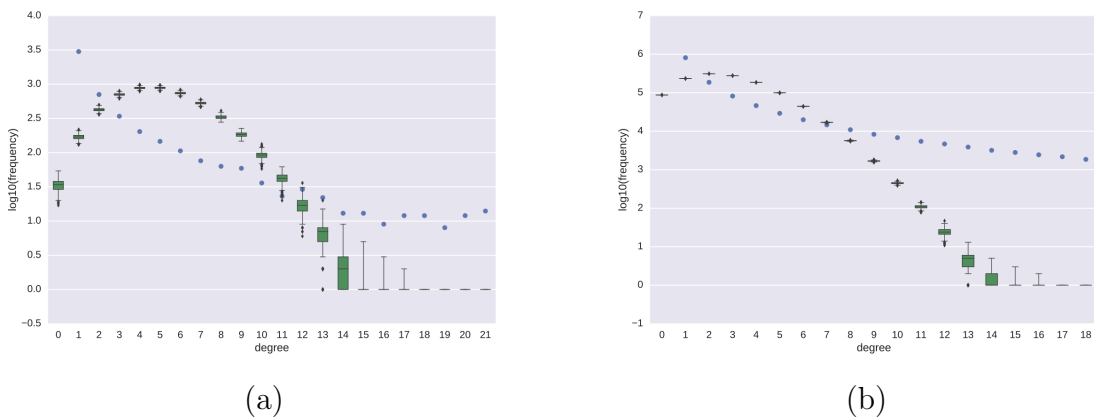(a)   (b)

FIGURE 7. Comparing degree distribution of datasets (blue dots) and generated model hypergraphs (boxplot to indicate variance): $\mathscr{E}_5$ (a) and $\mathscr{D}_5$ (b).

While it is clear from these plots that the model produces hypergraphs significantly different from reality, sections 6.3 and 6.4 further investigate these differences.

6.3. **Comparing the appearance of signatures in theory and practice.** Recall that the signature of a hypergraph H is defined as $\boldsymbol{e}(H) = (e_1(H), e_2(H), \ldots, e_k(H))$, where $k = v(H)$. Given a subset of vertices from a hypergraph, we will consider the non-$e(1)$ portion of the signature of the weak subhypergraph induced by these vertices. For our analysis, we fix a small graph G. We look for this graph G in the 2-section, $[H]_2$, of the network and then look at H to determine the signature of the originating weak subhypergraph.

For practical purposes, we focus on the complete graphs K_4 and K_5. Smaller complete graphs occur too frequently in the 2-section, whereas larger complete graphs do not occur frequently enough for us to claim anything statistically reliable. Of course, other subgraphs appear in $[H]_2$ but we are particularly interested in those specifically induced by the 2-section action. We use Corollary 5.4 to estimate the probability that a copy of K_4 or K_5 taken uniformly at random from the 2-section of $\mathscr{H}(n, \boldsymbol{p})$ originates from a weak subhypergraph with a given signature in $\mathscr{H}(n, \boldsymbol{p})$.

There are 70 signatures of hypergraphs on 4 vertices. Of these there are 60 that have at least one hypergraph inducing K_4 in the 2-section; call these *feasible*. For example, there are $36 = \binom{6}{1}\binom{4}{2}$ labelled hypergraphs with one 2-edge and two 3-edges, but only 6 of these induce K_4 in the 2-section. The number of signatures of hypergraphs on 5 vertices that can induce K_5 in the 2-section is 1,422. In order to apply Corollary 5.4 we generated all feasible signatures and all realizations of the subhypergraphs that induce a complete graph of the appropriate size in the 2-section [15].

As the results are similar for all four real-world hypergraphs described above, we choose $\mathscr{E}_5$ for illustration purposes. In Table 1 we list the most popular signatures from the theoretical point of view. Note that theoretical estimations for probabilities are decreasing very fast as we go down. As a result, it makes more sense to compare how signatures rank in both models instead of comparing their corresponding probabilities.

We ranked the 179 (of a possible 1,422) observed signatures based on their number of occurrences in $\mathscr{E}_5$; we denote this $R_{observed}$. We also ranked the 179 observed signatures based on their theoretical probabilities; we denote this $R_{theory/observed}$. Some theoretically popular signatures do not occur at all in practice. For example, signatures $(0,1,0,1)$ and $(2,1,1,0)$ were the 4th and the 7th most popular signatures from theoretical predictions; we call these ranks R_{theory}.

To understand these results at a glance, we took the 179 observed signatures and plotted the observed ranking ($R_{observed}$) against their theoretical ranking ($R_{theory/observed}$) in Figure 8. Figure 8 also contains a similar plot for the $\mathscr{E}_4$ dataset. It is clear from these two plots that the theoretical and observed ranks are highly uncorrelated.

In Figure 9 we present the same comparison of the observed versus theoretical observed ranks of signatures for the coauthorship hypergraph.

These rankings confirm our expectation that the simple random hypergraph model we have introduced does not capture features exhibited in real-world datasets. In particular, the model assumes edges occur independently, which is generally not the case in real-world networks and certainly not the case in the two datasets we considered. In order to better understand the interdependence of hyperedges in real-world networks, some notion of clustering coefficient must be investigated in the hypergraph setting.

TABLE 1. $\mathscr{E}_5$ network: popular signatures in theory

Signature	Probability (theory)	R_{theory}	$R_{theory/observed}$	Probability (observed)	$R_{observed}$
0,0,0,1	9.8118419955e-01	1	1	0.004911591	56
1,0,0,1	1.8649018608e-02	2	2	0.014734774	18
2,0,0,1	1.5950486447e-04	3	3	0.015717092	16
0,1,0,1	5.4464266334e-06	4	-		
3,0,0,1	8.0844057265e-07	5	4	0.010805501	31
4,0,1,0	4.4141198260e-07	6	5	0.000982318	148
2,1,1,0	4.0695426832e-07	7	-		
1,1,0,1	1.0351829114e-07	8	6	0.004911591	57
0,2,1,0	3.1265534058e-08	9	-		
1,0,2,0	1.8314270051e-08	10	-		
3,1,1,0	6.1878678650e-09	11	7	0.000982318	141
5,0,1,0	5.0338562002e-09	12	8	0.002946955	89
4,2,0,0	2.8645459166e-09	13	-		
4,0,0,1	2.6890048533e-09	14	9	0.013752456	21
2,1,0,1	8.8539088013e-10	15	10	0.011787819	27

TABLE 2. $\mathscr{E}_5$ network: popular signatures in practice

Signature	Probability (theory)	$R_{theory/observed}$	Probability (observed)	$R_{observed}$
4,1,0,1	1.4926318276e-14	24	0.033398821	1
4,5,1,0	9.9459608645e-34	119	0.031434185	2
4,4,1,0	1.1916310154e-27	84	0.026522593	3
3,2,0,1	1.1209412171e-17	36	0.02259332	4
4,2,0,1	3.7284328311e-20	48	0.02259332	5
4,4,2,0	4.7383972663e-37	133	0.02259332	6
9,1,0,0	2.5600700779e-16	31	0.02259332	7
10,1,0,0	1.6219446713e-19	45	0.021611002	8
4,6,0,0	6.3964285475e-31	101	0.021611002	9
7,3,0,0	3.2971609825e-21	49	0.019646365	10
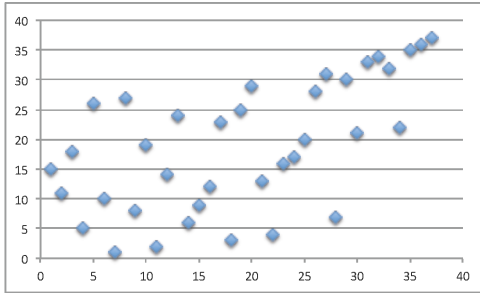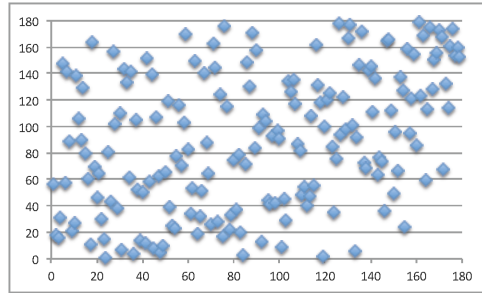
(a) $\mathscr{E}_4$ network   (b) $\mathscr{E}_5$ network

FIGURE 8. Theoretical rank, $R_{theory/observed}$ vs. Observed rank, $R_{observed}$ (x-axis vs. y-axis)

6.4. **Comparing clustering coefficient in theory and practice.** Let us recall the definition of the clustering coefficient for graphs. The clustering coefficient is an attempt to measure the degree to which vertices in a graph $G = (V, E)$ tend to cluster together by focusing on connections between neighbours of a vertex. The *local clustering coefficient* $c(v)$ for a vertex v is given by the proportion of links between the vertices within its
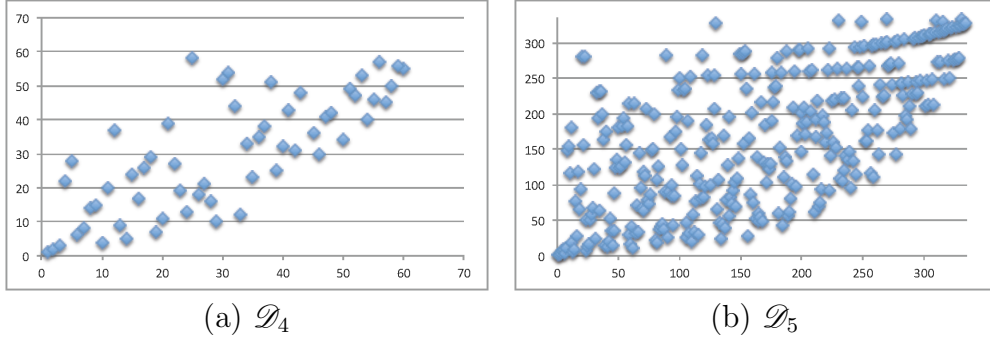
(a) $\mathscr{D}_4$ (b) $\mathscr{D}_5$

FIGURE 9. Theoretical rank, $R_{theory/observed}$ vs. Observed rank, $R_{observed}$ (x-axis vs. y-axis)

(open) neighbourhood (i.e. not including v) divided by the number of links that could possibly exist between them; that is, for vertices of degree at least 2

$$c(v) \;=\; \frac{|\{xy \in E : x, y \in N(v)\}|}{\binom{|N(v)|}{2}}$$

$$=\; \frac{\text{number of triangles involving } v}{\text{number of pairs of edges adjacent at v}},$$

while for vertices of degree less than 2, $c(v) = 0$. The *clustering coefficient* of a graph, with at least one vertex of degree at least 2, can be defined as the average of the local clustering coefficients of all the vertices; that is, when $|V_{\deg \geq 2}| \geq 1$,

$$C(G) = \frac{1}{|V_{\deg \geq 2}|} \sum_{v \in V} c(v),$$

where $|V_{\deg \geq 2}|$ is the number of vertices of degree at least 2. Note that this clustering coefficient weights the clustering coefficient of each vertex equally even though the number of possible connections of the neighbours of a vertex v is $\binom{\deg(v)}{2}$. This can lead to the clustering coefficient of the graph being skewed by the clustering coefficients of the smaller degree vertices.

Alternatively, to avoid the potential skew from smaller degree vertices, the clustering coefficient can be defined to weight every potential connection of neighbours equally. This leads to the *global clustering coefficient* for graphs having at least one vertex of degree 2 being defined as

$$C'(G) \;=\; \frac{\sum_{v \in V, \deg(v) \geq 2} |\{xy \in E : x, y \in N(v)\}|}{\sum_{v \in V, \deg(v) \geq 2} \binom{|N(v)|}{2}} \tag{9}$$

$$=\; \frac{3 \times \text{number of triangles}}{\text{number of pairs of adjacent edges}}. \tag{10}$$

This clustering coefficient is often thought of as the tendency of the graph to "close" triangles.

Suppose that G is a graph with at least one vertex of degree at least 2. Then $C(G)$ and $C'(G)$ are equal to 1 exactly when G consists of the disjoint union of complete

graphs. Both clustering coefficients are 0 exactly when G is bipartite; this is a bit of a shortcoming of the definitions since the complete bipartite graph $K_{n,n}$ has slightly more than half the edges of the complete graph, but has a clustering coefficent of 0.

A graph is considered *small-world* if its average local clustering coefficient $C(G)$ is significantly higher than a random graph constructed on the same vertex set, and if the graph has approximately the same mean shortest path length as its corresponding random graph.

There is no canonical way to generalize the idea of a graph clustering coefficient to hypergraphs, however, there have been a number of proposals: [17], [5], [2], [14], and [20]. For this paper we will focus on the following definitions of Zhou and Nakhleh [20] which generalize the idea of comparing the number of triangles to pairs of adjacent edges. Let $H = (V, E)$ be a hypergraph. If $v \in V$ and $e_i, e_j \in E$, let $M(v)$ denote the set of edges that contain v, let $N(v)$ denote the neighbours of v, and let $D_{ij} = e_i \setminus e_j = e_i \setminus (e_i \cap e_j)$. Local and global clustering coefficients on H are then defined as

$$HC_{local}(v) = \begin{cases} \frac{1}{\binom{|M(v)|}{2}} \sum_{e_i, e_j \in M(v)} EO(e_i, e_j) & \text{if } |M(v)| \geq 2 \\ 0 & \text{if } |M(v)| \leq 1, \end{cases}$$

and

$$HC_{global}(H) = \begin{cases} \frac{1}{|\mathfrak{I}|} \sum_{\{e_i, e_j\} \in \mathfrak{I}} EO(e_i, e_j) & \text{if } \mathfrak{I} \neq \emptyset \\ 0 & \text{if } \mathfrak{I} = \emptyset, \end{cases}$$

where $\mathfrak{I} = \{\{e_i, e_j\} : e_i, e_j \in E, e_i \neq e_j, \text{ and } e_i \cap e_j \neq \emptyset\}$ are the pairs of intersecting edges, and $EO(e_i, e_j)$, the *extra overlap* of a pair of edges, is defined as

$$EO(e_i, e_j) = \frac{|N(D_{ij}) \cap D_{ji}| + |N(D_{ji}) \cap D_{ij}|}{|D_{ij}| + |D_{ji}|}$$

(the proportion of the vertices in exactly one of the edges that are neighbours of vertices in only the other edge).

Note that if H is a graph then $HC_{global}(H) = C'(H)$ and $HC_{local}(v) = C(v)$. This follows from the fact that if $e_i = \{u, v\}$ and $e_j = \{u, w\}$ then the extra overlap $EO(e_i, e_j) = 1$ if $\{v, w\} \in E$ and 0 otherwise.

We calculated local and global clustering coefficients for the four hypergraphs derived from real-world datasets and over a sample of random hypergraphs generated with the same edge distribution. Figure 10 presents the distribution of local clustering coefficients for $\mathscr{E}_5$ and its corresponding model. This is illustrative of the results in general, and so we do not include other figures. In fact, in the case of the coauthorship network, the results were even more glaringly dissimilar with the models having more than half the vertices with a local clustering coefficient of zero, while most of the vertices in the real-world networks were non-zero. Specifically, the number of non-zero local clustering coefficients in $\mathscr{E}_5$ is 126844. This is 7836 standard deviations above the average for our generative model with maximum edge size 5.

Table 3 presents the global clustering coefficents of the four real-world hypergraphs, and the mean over a sample of random hypergraphs generated with the same edge distributions.
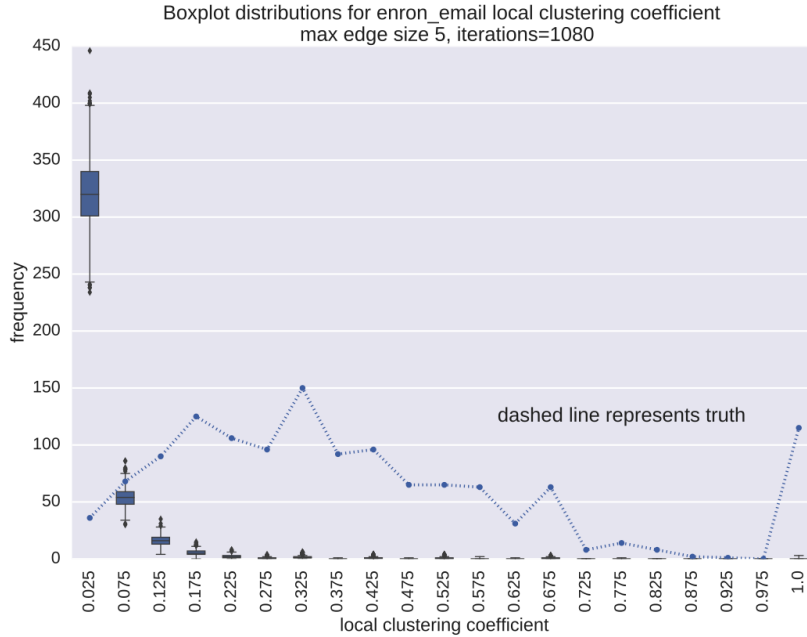
FIGURE 10. Local clustering coefficient of $\mathscr{E}_5$ and corresponding model hypergraphs

	$\mathscr{D}_4$	$\mathscr{D}_5$	$\mathscr{E}_4$	$\mathscr{E}_5$
networks we examined	0.0550	0.0687	0.1048	0.1382
random hypergraphs	$8.01 \cdot 10^{-6}$	$1.17 \cdot 10^{-5}$	0.00224	0.00308

TABLE 3. Global clustering coefficient for networks we examined and corresponding model hypergraphs

Note that the true global clustering coefficient for $\mathscr{E}_5$ is approximately 41653 standard deviations from the mean of our model with max edge size 5. Not so glaring, but still significant, the true global clustering coefficient for $\mathscr{D}_5$ is approximately 474 standard deviations above the mean of our model with max edge size of 5.

Both the local and global clustering coefficient results indicate what we expect; the model assumption of edge independence is not reflective of real-world networks. We all know researchers do not select coauthors at random!

7. Conclusions and future work

The ultimate goal of this work is to develop a reasonable model for complex networks using hypergraphs. While there are many models using graphs – including classic ones such as the binomial random graph ($\mathscr{G}(n,p)$), random d-regular graphs, and the preferential attachment model, as well as spatial ones such as random geometric graphs and the spatial preferential attachment model – there are very few using hypergraphs. The model we proposed is a generalization of $\mathscr{G}(n,p)$ and thus, as we observed, does not capture several important features of many real-world networks. However, it does

allow us to identify some structures inherent in many of these networks, in particular, the non-independence of hyperedges. It also allowed us to illustrate that some questions posed about hypergraphs cannot be addressed by looking at the 2-section.

References

[1] Mohammad A. Bahmanian and Mateja Šajna. "Connection and separation in hypergraphs". In: *Theory and Applications of Graphs* 2.2 (2015).

[2] Stevens Le Blond, Jean-Loup Guillaume, and Matthieu Latapy. "Clustering in P2P Exchanges and Consequences on Performances". In: *Lecture Notes in Computer Science, vol. 3640: Peer-to-Peer Systems IV (IPTPS)*. Ed. by M. Castro and R. van Renesse. Springer, 2005, pp. 193–204.

[3] Bela Bollobás. "Random Graphs". In: *LMS: 52 Combinatorics (London Mathematical Society Lecture Note Series)*. Ed. by H. N. V. Temperley. Cambridge: Cambridge University Press, 1981, pp. 80–102.

[4] Bela Bollobás. *Random Graphs*. Cambridge University Press, 2001.

[5] Stephen P. Borgatti and Martin G. Everett. "Network analysis of 2-mode data". In: *Social Networks* 19.3 (1997), pp. 243–269.

[6] Megan Dewar et al. "Subgraphs in non-uniform random hypergraphs". In: *Lecture Notes in Computer Science, vol. 10088: Proceedings of the 13th International Workshop on Algorithms and Models for the Web Graph*. Ed. by Anthony Bonato Fan Chung Graham and Paweł Prałat. Springer, 2016, pp. 140–151.

[7] Pierre Duchet. "Hypergraphs". In: *Handbook of Combinatorics*. Ed. by R. L. Graham, M. Grötschel, and L. Lovász. Amsterdam: Elsevier, 1995.

[8] Andrzej Dudek and Alan Frieze. "Loose Hamilton Cycles in Random Uniform Hypergraphs". In: *Electronic Journal of Combinatorics* 18.1 (2011), 14 pages.

[9] Andrzej Dudek and Alan Frieze. "Tight Hamilton cycles in random uniform hypergraphs". In: *Random Structures and Algorithms* 42.3 (2012), pp. 374–385.

[10] Asaf Ferber. "Closing Gaps in Problems related to Hamilton Cycles in Random Graphs and Hypergraphs". In: *Electronic Journal of Combinatorics* 22.1 (2015), 7 pages.

[11] Alan M. Frieze and Michał Karoński. *Introduction to Random Graphs*. Cambridge University Press, 2015.

[12] Svante Janson, Tomasz Łuczak, and Andrzej Ruciński. *Random Graphs*. New York: Wiley, 2000.

[13] Anders Johansson, Jeff Kahn, and Van Vu. "Factors in random graphs". In: *Random Structures and Algorithms* 33.1 (2008), pp. 1–28.

[14] Matthieu Latapy, Clémence Magnien, and Nathalie Del Vecchio. "Basic notions for the analysis of large two-mode networks". In: *Social Networks* 30.1 (2008), pp. 31–48.

[15] Paweł Prałat. *Feasible signatures and subhypergraphs that realize them*. 2016. URL: http://www.math.ryerson.ca/~pralat/ (visited on 01/21/2017).

[16] CALO Project. *Enron Email Dataset*. 2015. URL: https://www.cs.cmu.edu/~. /enron/ (visited on 01/21/2017).

[17] Garry Robins and Malcolm Alexander. "Small Worlds Among Interlocking Directors: Network Structures and Distance in Bipartite Graphs". In: *Computational & Mathematical Organization Theory* 10.1 (2004), pp. 69–94.

[18] Jie Tang et al. "ArnetMiner: extraction and mining of academic social networks". In: *KDD '08: Proceedings of the 14th ACM SIGKDD International Conference on Knowledge discovery and data mining*. New York: ACM, 2008, pp. 9990–998.

[19] Jie Tang et al. *Extraction and Mining of Academic Social Networks*. 2008. URL: https://aminer.org/billboard/AMinerNetwork (visited on 01/21/2017).

[20] Wanding Zhou and Luay Nakhleh. "Properties of metabolic graphs: biological organization in representation artifacts?" In: *BMC Bioinformatics* 12.132 (2011).