

Hypergraph Representation via Axis-Aligned Point-Subspace Cover

Oksana Firman¹

Joachim Spoerhase²

¹ Universität Würzburg, Würzburg, Germany

² University of Liverpool, Liverpool, United Kingdom

revisions 31st July 2023, 16th Jan. 2025; accepted 8th Feb. 2025.

A k -hypergraph is a k -partite k -uniform hypergraph, that is, a hypergraph with a partition of vertices into k parts such that each hyperedge contains exactly one vertex of each part. We propose a new geometric representation of k -hypergraphs. Namely, given positive integers ℓ, d , and k with $\ell \leq d - 1$ and $k = \binom{d}{\ell}$, any finite set P of points in \mathbb{R}^d represents a k -hypergraph G_P as follows. Each point in P is covered by k many axis-aligned affine ℓ -dimensional subspaces of \mathbb{R}^d , which we call ℓ -subspaces for brevity and which form the vertex set of G_P . We interpret each point in P as a hyperedge of G_P that contains each of the covering ℓ -subspaces as a vertex. The class of (d, ℓ) -hypergraphs is the class of k -hypergraphs that can be represented in this way. The resulting classes of hypergraphs are fairly rich, since every k -hypergraph is a $(k, k - 1)$ -hypergraph. On the other hand, for $\ell < d - 1$, there exists a k -hypergraph which is not a (d, ℓ) -hypergraph.

In this paper we give a natural structural characterization of (d, ℓ) -hypergraphs based on vertex cuts. This characterization leads to a polynomial-time recognition algorithm that decides for a given k -hypergraph G whether or not G is a (d, ℓ) -hypergraph and that computes a representation of G if one exists. Here we assume that the dimension d is constant and that the partition of the vertex set of G is prescribed.

Keywords: hypergraph, point-line cover, graph representation

1 Introduction

Geometric representations of graphs or hypergraphs provide a wide and intensively studied field of research in the intersection of combinatorial geometry and graph theory. Well-known examples are geometric intersection or incidence graphs, with a large amount of literature, such as McKee and McMorris (1999), Coxeter (1950), and Pisanski and Randić (2001). The benefits of adopting these two perspectives—the geometric and the graph-theoretic one—are as follows. On the one hand, knowing that a given graph can be represented geometrically may give new insights about its structure, because the geometric perspective is often rather intuitive. On the other hand, giving a graphical characterization for certain types of geometric objects may help pin down the essential combinatorial properties that can be exploited in the geometric setting.

One example of this interplay is the study of geometric set cover and hitting set problems, as explored by Brönnimann and Goodrich (1995), Varadarajan (2010), and Chan et al. (2012). In this important branch of

geometric optimization, incidence relations of two types of geometric objects are studied where one object type is represented by vertices of a hypergraph whose hyperedges are, in turn, represented by the other object type. In this representation a vertex is contained in a hyperedge if and only if the corresponding geometric objects have a certain geometric relation such as containment or intersection. The objective is to find the minimum number of vertices hitting all hyperedges⁽ⁱ⁾. In this line of research, the goal is to exploit geometric properties in order to improve upon the state of the art for general hypergraphs. This is known to be surprisingly challenging even in many seemingly elementary settings.

For example, in the well-studied point-line cover problem, as studied by Hassin and Megiddo (1991) and Kratsch et al. (2014), we are given a set of points in the plane and a set of lines. The goal is to identify a smallest subset of the lines to cover all the points. This problem can be cast as a hypergraph vertex cover problem. Points can be viewed as hyperedges containing the incident lines as vertices. The objective is to cover all the hyperedges by the smallest number of vertices.

It seems quite clear that point-line cover instances form a heavily restricted subclass of general hypergraph vertex cover. For example, they have the natural intersection property that two lines can intersect in at most one point. However, somewhat surprisingly, in terms of approximation algorithms, no worst-case result improving the ratios for general hypergraph vertex cover is known (Aharoni et al., 1996; Guruswami and Saket, 2010; Chvátal, 1979). In fact, it has been shown that merely exploiting the above intersection property in the hypergraph vertex cover is not sufficient to give improved approximations, as showed by Kumar et al. (2000). Giving a simple combinatorial characterization of the point-line cover instances seems to be challenging.

In this paper, we study a representation of hypergraphs that arises from a natural variant of point-line cover where we want to cover a given set P of points in \mathbb{R}^d by *axis-parallel* lines, see Fig. 1 for an illustration. While the axis-parallel case of point-line cover has been considered before by Gaur and Bhat-tacharya (2007), the known algorithms do not improve upon the general case of hypergraph vertex cover, as studied by Aharoni et al. (1996); Chvátal (1979). More generally, we investigate the generalization where we are additionally given a parameter $\ell \leq d - 1$ and we would like to cover P by axis-aligned affine ℓ -dimensional subspaces of \mathbb{R}^d , which we call ℓ -subspaces. The resulting class of hypergraphs is fairly rich as any k -partite k -uniform hypergraph (i.e., a hypergraph with a partition of vertices into k parts such that each hyperedge contains exactly one vertex of each type) can be represented by a set of points in \mathbb{R}^k to be covered by $(k - 1)$ -subspaces. On the other hand, for $\ell < d - 1$, there exists a k -partite k -uniform hypergraph which is not a (d, ℓ) -hypergraph, i.e., can not be represented by a set of points in \mathbb{R}^d to be covered by ℓ -dimensional subspaces of \mathbb{R}^d .

Note that for any sufficiently large set X , a similar representation can be defined in the power X^d , where hyperedges correspond to elements of X^d and vertices correspond to subsets of X^d with $d - \ell$ coordinates fixed. We keep $X = \mathbb{R}$ to provide more intuitive representations.

Related work. The question of representing (hyper-)graphs geometrically is related to the area of *graph drawing*. For example, Evans et al. (2019) study drawing hypergraphs in \mathbb{R}^3 by representing vertices as points and hyperedges as convex polygons while preserving incidence relations. Note that in our work we study the “dual” representation, where hyperedges are represented as points and vertices are represented as axis-parallel lines or affine subspaces. Another related problem in graph drawing has been introduced by Chaplick et al. (2016). They study drawings of graphs in finite-dimensional Euclidean spaces where

⁽ⁱ⁾ For the sake of presentation, we use here the representation as hitting set problem rather than the equivalent and maybe more common geometric set cover interpretation.



Fig. 1: A graph (a) and a hypergraph (b) and their representations in \mathbb{R}^2 and \mathbb{R}^3 , respectively.

the vertices (and the edges, in some variants) of the graph can be covered by a minimum number of lines or hyperplanes.

Our contribution and outlook. We initiate the study of representing k -hypergraphs, i.e., k -partite k -uniform hypergraphs, via axis-aligned point-subspace cover instances in this generality. Our main insight is that the axis-aligned case of point-subspace cover allows for a natural, combinatorial characterization contrasting what is known for the non-aligned case (see discussion above). The characterization is based on vertex cuts and can be leveraged to obtain a polynomial time recognition algorithm for such hypergraphs assuming the dimension d is a constant and that we are given the partition of the vertices, which is NP-hard to compute for $k \geq 3$ (Ilie et al., 2005).

We believe that it is an interesting future research direction to exploit these combinatorial properties in order to obtain improved results for various optimization problems in hypergraphs such as hypergraph vertex cover or hypergraph matching. We also hope that our combinatorial characterization may help make progress on geometric problems. We conclude our paper by an outlook containing related open questions and some first motivating results in this direction.

Notation and problem definition. We use d and ℓ to denote the dimensions of the space and the axis-aligned subspaces, respectively. We use $k = \binom{d}{\ell}$ to denote the number of parts in the corresponding hypergraph. For the special case of point-line cover that we consider in Section 2 we have $k = d$. Let $[k] = \{1, \dots, k\}$ for any $k \in \mathbb{N}$.

Let P be a finite set of points in \mathbb{R}^d . We define the k -hypergraph G_P as follows. The vertex set of G_P is the set of axis-parallel affine subspaces (lines in the special case) containing at least one point in P . The hyperedges in G_P correspond to the points in P . The hyperedge corresponding to a point $p \in P$ contains the k axis-parallel subspaces (lines) incident to p as vertices. Note that G_P is k -partite and k -uniform where the k parts of the partition correspond to $\binom{d}{\ell}$ (d in the special case of point-line cover) sets of axis-aligned subspaces.

Our main task is to decide, for a given hypergraph $G = (V, E)$ with the partition $V = V_1 \cup V_2 \cup \dots \cup V_k$ of the vertex set V of G , whether there is an axis-aligned point-subspace cover (or point-line cover in the special case) instance P such that G and G_P are isomorphic. In this case we say that G is *represented* by P and, thus, *representable*. A representable hypergraph we also call (d, ℓ) -hypergraph. We assume that the partition of G into k parts is given.

We remark that every bipartite graph is representable in \mathbb{R}^2 by point-line cover because we can derive a grid-like representation as shown in Fig. 1(a) directly from the adjacency matrix of this graph. However,



Fig. 2: Two hypergraphs that cannot be represented by a set of points in \mathbb{R}^3 to be covered by lines. In (a), any two lines corresponding to the shared vertices already define a unique point; in (b), one of the points representing a hyperedge would have to lie on two skew lines, which is impossible.

for $k \geq 3$ the class of $(k, 1)$ -hypergraphs forms a non-trivial subclass of all k -hypergraphs. For example, no pair of hyperedges can share more than one vertex in any representable hypergraph, see Fig. 2(a). But this constraint alone is not enough to characterize representable hypergraphs, see Fig. 2(b).

2 Characterization of Hypergraphs Representable by Point-Line Cover

For the sake of an easier presentation, we first describe the result for the special case of point-line covers, that is, for $(d, 1)$ -hypergraphs. We later describe how the result generalizes to higher-dimensional axis-aligned affine subspaces.

Let $G = (V, E)$ be a k -hypergraph with the partition $V = V_1 \cup \dots \cup V_k$. We aim to compute a point-line cover instance P (if such exists) such that each $e = (v_1, \dots, v_k) \in E$ corresponds to some $p^e = (x_1^e, \dots, x_k^e) \in P$, where these correspondences are induced by the isomorphism between G and G_P . In addition, for each $i \in [k]$, each vertex $v_i \in V_i$ corresponds to the line ℓ^{v_i} that is parallel to the i -th coordinate axis and contains p^e , that is, for all $j \neq i$, we fix the j -th coordinate, whereas the i -th coordinate is *free*. Refer to Fig. 1 for illustrative examples.

Definition 1 For any vertices $s, t \in V$, an s - t path is a sequence of vertices $s = v_1, \dots, v_r = t$ such that v_i and v_{i+1} are both contained in some hyperedge $e \in E$ for all $i \in [r-1]$. Similarly, for any hyperedges $e, e' \in E$, an e - e' path is a v - v' path such that $v \in e$ and $v' \in e'$.

The following two separability conditions based on vertex cuts are key for our characterization.

Definition 2 (Vertex separability) For a given k -hypergraph G we say that distinct vertices $v, v' \in V_i$ with $i \in [k]$ are separable if there exists $j \in [k]$ with $j \neq i$ such that every v - v' path contains a vertex in V_j . (Informally, removing V_j from the vertex set and from the edges of G separates v and v' .) A k -hypergraph G is called vertex-separable if every two distinct vertices from the same vertex part of G are separable.

Definition 3 (Edge separability) For a given k -hypergraph G we say that its distinct hyperedges $e, e' \in E$ are separable if there exists $j \in [k]$ such that every e - e' path contains a vertex in V_j . A k -hypergraph G is called edge-separable if every two distinct hyperedges of G are separable.

Note that any pair of hyperedges sharing two or more vertices are not separable. Therefore, edge-separable k -hypergraphs do not contain such hyperedge pairs.

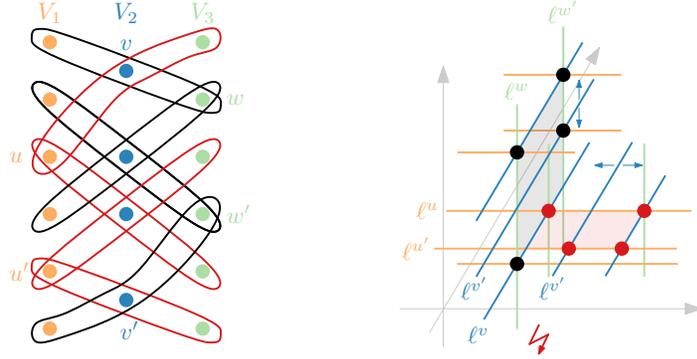


Fig. 3: A hypergraph G (on the left) that is edge-separable, but not vertex-separable (the vertices v and v' from V_2 are not separable). In a hypothetical representation of G (on the right), the line $\ell^{v'}$ must simultaneously intersect $\ell^{u'}$ and $\ell^{w'}$ and therefore must be equal to ℓ^v , which is impossible

Lemma 1 Any vertex-separable k -hypergraph is edge-separable.

Proof: Let G be any vertex-separable k -hypergraph, and let e and e' be any distinct hyperedges of G . There are an integer $i \in [k]$ and distinct vertices $v, v' \in V_i$ such that $v \in e$ and $v' \in e'$. Since G is vertex-separable, there exists $j \in [k]$ with $j \neq i$ such that every v - v' path contains a vertex in V_j . Let π be any e - e' path. Then v, π, v' is a v - v' path, so it contains a vertex in V_j , and so does π . Thus G is edge-separable. \square

The converse of Lemma 1 is not true, see Fig. 3. In the instance depicted, the two red edges containing v, u and u', v' or the two black edges containing v, w and w', v' , for example, are separated by removing the blue vertex part V_2 (which we cannot do to show that vertices v and v' are separable).

Definition 4 Let $G = (V, E)$ be a k -hypergraph with the partition $V = V_1 \cup V_2 \cup \dots \cup V_k$. For each $i \in [k]$, let $G_i = (E, E_i)$ be the graph whose distinct vertices e and e' are adjacent if and only if they have a common vertex in a part V_j with $j \neq i$.

In the following theorem we state our characterization of k -hypergraphs representable by point-line covers via vertex separability.

Theorem 1 Let k be any positive integer. A k -hypergraph G is representable if and only if G is vertex-separable.

Proof: We construct for each hyperedge e a point $p^e \in \mathbb{R}^k$ and for each vertex $v_i \in V_i$ with $i \in [k]$ a line $\ell^{v_i} \subseteq \mathbb{R}^k$ that is parallel to the x_i -axis. We do this as follows. For G we construct the graphs G_i , $i \in [k]$ defined in Definition 4. For each graph G_i , $i \in [k]$ we consider the connected components of G_i and assign to each connected component a unique (integer) value in $\{1, 2, \dots, |V(G_i)|\}$.

Now, if p_i^e is the value of the connected component in G_i that contains e then we let the point $p^e = (p_1^e, \dots, p_k^e)$ represent the hyperedge e , see Fig. 4 for an example.

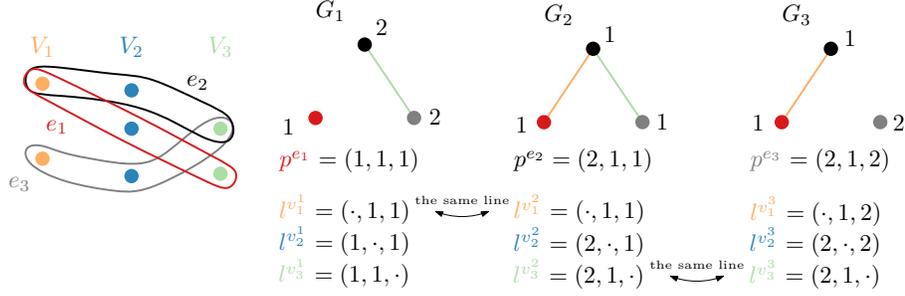


Fig. 4: A hypergraph G , the graphs G_1, G_2, G_3 , and the coordinates of the points and lines corresponding to the hyperedges and vertices. The dots placed instead of coordinates mean that those coordinates are free.

Recall that any line parallel to the x_i -axis can be defined by fixing its x_j -coordinate for all $j \neq i$, while leaving x_i free. Now, if the hyperedge $e = \{v_1, \dots, v_k\}$ is represented by $p^e = (p_1^e, \dots, p_k^e)$ then for each $i \in [k]$, the line l^{v_i} that represents the vertex v_i is defined by coordinates $p_j^e, j \neq i$ while leaving the x_i -coordinate free, see Fig. 4. It is important to note, that the representation l^{v_i} is well-defined although v_i may be contained in multiple hyperedges in G . This follows from the fact that all the hyperedges containing v_i belong to the same connected component in $G_j, j \in [k], j \neq i$ because each pair of them is joined by some edge in G_j corresponding to v_i and in particular these hyperedges form a clique. Therefore, there is no disagreement in the x_j -coordinate where $j \neq i$. Hence, we uniquely define the coordinates that determine a line.

(\Leftarrow) Assume that G is vertex-separable. By the construction of the point-line cover instance we have:

- every point p^e is in fact covered by the lines l^{v_1}, \dots, l^{v_k} where $e = \{v_1, \dots, v_k\}$, because by construction every line l^{v_i} and point p^e have the same x_j -coordinate with $j \neq i$.
- $\forall v \neq v' \in V$ it holds that $l^v \neq l^{v'}$. This is obviously true if vertices belong to different parts, because then the free coordinate of v is fixed for v' and vice versa. If $v, v' \in V_i$ for some $i \in [k]$ then, by vertex separability, there exists $j \neq i$ such that v and v' are not connected in graph G_j and get different x_j -coordinates. So they represent distinct lines.
- $\forall e \neq e' \in E$ it holds that $p^e \neq p^{e'}$. Indeed, by Lemma 1, G is edge-separable and by definition of edge separability (see Definition 3) distinct hyperedges are not connected in at least one graph G_i and get different x_i -coordinates. So they represent distinct points.

By the above construction, for every incident vertex-hyperedge pair $v \in V, e \in E$, that is, $v \in e$, the corresponding geometric objects l^v and p^e are incident as well. We claim that if v and e are not incident, that is, $v \notin e$ then l^v and p^e are not incident as well. This is because every point p^e is already incident to precisely k lines l^v by construction, because the lines l^v are pairwise distinct, and because p^e cannot be incident on more than k axis-parallel lines. Thus, we have constructed a point-line cover instance that represents the hypergraph G and this means that G is representable.

(\Rightarrow) Assume that G is not vertex-separable but that it has a point-line cover representation. This means that it contains at least two distinct vertices v and v' from the same part V_i that are not separable. Then for each part V_j with $j \neq i$, there exists a v - v' path $v = v_1, \dots, v_r = v'$ such that $v_t \notin V_j$ for each

$t \in [r]$. All lines ℓ^{v_t} with $t \in [r]$ that represent the vertices v_1, \dots, v_r lie on the same hyperplane H_j perpendicular to the x_j -axis. This is because successive line pairs are joined by a common point (representing the hyperedge containing both) and since none of these lines is parallel to the x_j -axis and so the x_j -coordinate stays fixed. Since this holds for all $j \in [k]$, $j \neq i$, the lines ℓ^v and $\ell^{v'}$ lie in the intersection $\bigcap_{j \neq i} H_j$. But the intersection of such hyperplanes is a single line. This contradicts that v and v' correspond to the distinct lines. \square

Note that our characterization in Theorem 1 directly gives rise to an efficient recognition algorithm to check whether or not a given k -hypergraph is representable. We check for every vertex pair from the same part whether they are separable.

Below we give an algorithm to actually compute a representation for a given k -hypergraph (or output that a representation does not exist). The algorithm is an algorithmic implementation of the construction used in the proof of Theorem 1 to recover a representation from a given (representable) k -hypergraph.

Computing a representation. For each $i \in [k]$ we compute a graph G_i as in Definition 4. We use the construction from the proof of Theorem 1 to assign the coordinates for all points and lines (using connected components of G_i). In particular, we uniquely number the connected components of each graph G_i , $i \in [k]$. Then, for any hyperedge e , we construct the point p^e whose i -th coordinate is the number of the connected components containing e in G_i . Based on the representation of the hyperedges we can easily obtain the representation of the vertices as lines as well. To verify if the resulting candidate representation is in fact valid, we have to check if any two points $p^e, p^{e'}$ representing distinct hyperedges e, e' are in fact distinct, that is, differ in at least one coordinate. Similarly, we have to check distinctness of the lines. If this is the case then the hypergraph G is representable and we get the point-line cover instance corresponding to G .

The construction of a point-line cover instance for a given hypergraph can be done in $\mathcal{O}(k \cdot m^2)$, where m is the number of hyperedges. Checking distinctness of the vertex and edge representations takes $\mathcal{O}(k \cdot (m \log m + n \log n))$ using lexicographic sorting, where n is the number of vertices. In total, the runtime of the algorithm is $\mathcal{O}(k \cdot (m^2 + n \log n))$.

3 Generalization to Point-Subspace Cover

In this section, we generalize the previous result for $(d, 1)$ -hypergraphs to the general case of (d, ℓ) -hypergraphs. That is, we want to characterize k -hypergraphs (where $k = \binom{d}{\ell}$) that are representable as sets of points in \mathbb{R}^d covered by ℓ -dimensional axis-aligned subspaces where $1 \leq \ell \leq d - 1$ and $d \geq 2$. Such an axis-aligned affine subspace L in \mathbb{R}^d is defined by fixing $d - \ell$ coordinates (called *fixed coordinates*) to specific values and allowing ℓ coordinates (called *free coordinates*) to vary freely. Formally, there exists a point $(x_1, \dots, x_d) \in \mathbb{R}^d$ and an index set $I \subseteq [d]$ with $|I| = \ell$ such that

$$L = \{(y_1, \dots, y_d) \in \mathbb{R}^d \mid y_i = x_i \text{ for all } i \notin I\}.$$

In this formulation, the coordinates in I define the ℓ -dimensional degrees of freedom, while the remaining coordinates are fixed to the values of x_i .

Note that in the case $\ell = d - 1$ every k -hypergraph can be represented analogously to the case of bipartite graphs, each of which is representable in two-dimensional space. This follows immediately from our characterization below but can also be seen in a direct way analogously to the two-dimensional setting.

For each positive integer $\ell \leq d-1$, each (d, ℓ) -hypergraph is a k -hypergraph, for $k = \binom{d}{\ell}$. Every part of the vertex partition in the hypergraph represents a set of ℓ -subspaces that have the same free coordinates. In the following we assume that the *labeling* of the parts of the hypergraph is given. That is, for each part of the hypergraph, we know which ℓ -subset I of free coordinates this part corresponds; we write V_I to denote this correspondence. If the labeling is not given, we guess the correct labeling in a brute-force fashion by checking all $\mathcal{O}(d^{\ell d})$ many labelings whether it satisfies the characterization for fixed labeling as described below.

The following definitions and proofs are natural generalizations of these for $(d, 1)$ -hypergraphs. In the case of $\ell = 1$ the idea of separability is related to cutting the part V_j of the graph corresponding to some coordinate $j \in [k]$. The key idea in the generalized setting is to cut instead all parts V_I that contain a specific coordinate j . Below, we give the full definitions and proofs, which are along the lines of $\ell = 1$ but rather use the generalized concept of cuts and separability.

Generalization of the vertex separability property. Let ℓ and d with $\ell \leq d-1$ be any positive integers and let $\binom{[d]}{\ell}$ be the family of subsets of $[d]$ of cardinality ℓ . Let $k = \left| \binom{[d]}{\ell} \right| = \binom{d}{\ell}$. Let $G = (V, E)$ be a k -hypergraph and let $V = \bigcup_{I \in \binom{[d]}{\ell}} V_I$ be the partition of V . For each $i \in [d]$, let $\widehat{V}_i = \bigcup \{V_I \mid I \in \binom{[d]}{\ell}, I \ni i\}$, i.e., a set of parts of V that share the free coordinate i .

Now we generalize the vertex separability property using the definition of a path (see Definition 1 from Section 2).

Definition 5 (Vertex separability) *Let $I \in \binom{[d]}{\ell}$ be any set. For a given $\binom{[d]}{\ell}$ -hypergraph G , two distinct vertices v and $v' \in V_I$ are separable if there exists $j \in [d] \setminus I$ such that every v - v' path contains a vertex in \widehat{V}_j . (Informally, removing \widehat{V}_j from the vertex set and from the hyperedges of G separates v and v' .) A hypergraph G is called vertex-separable if every two distinct vertices from the same vertex part of G are separable.*

In an analogous way we can generalize edge separability and Lemma 1 that vertex separability implies edge separability.

Definition 6 (Edge separability) *For a given $\binom{[d]}{\ell}$ -hypergraph G , two distinct hyperedges e and e' of G are separable if there exists $j \in [d]$ such that every e - e' path contains a vertex in \widehat{V}_j . A hypergraph G is called edge-separable if every two distinct hyperedges of G are separable.*

Lemma 2 *Any vertex-separable $\binom{[d]}{\ell}$ -hypergraph is edge-separable.*

Proof: Assume that a given $\binom{[d]}{\ell}$ -hypergraph G is not edge-separable. This means that there are two distinct hyperedges e and e' that are not separable. Then $\forall j \in [d]$ there is an e - e' path that does not contain a vertex from \widehat{V}_j . Because e and e' are distinct, there are distinct vertices v and v' with $v \in e$ and $v' \in e'$ from the same part V_I for some $I \in \binom{[d]}{\ell}$. Now, for each $j \in [d] \setminus I$, there exists an e - e' path π_j that does not contain a vertex from \widehat{V}_j . But then v, π_j, v' forms a v - v' path that does not contain a vertex from \widehat{V}_j . This means that G is not vertex-separable. \square

As was already mentioned in the previous section, the converse is not true, see Fig. 3.

Definition 7 Let $G = (V, E)$ be a $\binom{d}{\ell}$ -hypergraph with the partition $V = \bigcup_{I \in \binom{[d]}{\ell}} V_I$. For each $i \in [d]$, let $G_i = (E, E_i)$ be the graph whose distinct vertices e and e' are adjacent if and only if they have a common vertex in $V \setminus \widehat{V}_i$.

Theorem 2 Let ℓ and d with $\ell \leq d - 1$ be any positive integers. A $\binom{d}{\ell}$ -hypergraph G is representable if and only if G is vertex-separable.

Proof: The proof is a natural generalization of the proof of Theorem 1.

We construct for each hyperedge e a point $p^e \in \mathbb{R}^d$ and for each vertex $v_I \in V_I$ with $I \in \binom{[d]}{\ell}$ a subspace $\ell^{v_I} \in \mathbb{R}^d$ that is parallel to all x_i -axis where $i \in I$. We do this as follows. For G , we construct the graphs G_i with $i \in [d]$ defined in Definition 7. For each graph G_i we consider the connected components of the graph and assign to each of them a unique (integer) value.

Now, if p_i^e is the value of the connected component in G_i that contains e then we let the point $p^e = (p_1^e, \dots, p_d^e)$ represent the hyperedge e , see Fig. 4 for an example.

Recall that any subspace parallel to all x_i -axis where $i \in I$ can be defined by fixing its x_j -coordinate for all $j \neq i$, while leaving x_i -coordinates free. Now, if the hyperedge $e = \{v_I\}$ for all $I \in \binom{[d]}{\ell}$ is represented by $p^e = (p_1^e, \dots, p_d^e)$ then for each $I \in \binom{[d]}{\ell}$, the subspace ℓ^{v_I} that represents the vertex v_I is defined by coordinates p_j^e , $j \notin I$ while leaving the x_i -coordinates free. It is important to note, that the representation ℓ^{v_I} is well-defined although v_I may be contained in multiple hyperedges in G . This follows from the fact that all the hyperedges containing v_I belong to the same connected component in G_j with $j \in [d]$ because each pair of them is joined by some edge in G_j corresponding to v_I and in particular these hyperedges form a clique. Therefore, there is no disagreement in the x_j -coordinate where $j \neq i$ for all $i \in I$. Hence, we uniquely define the coordinates that determine a subspace.

(\Leftarrow) Assume that G is vertex-separable. By the construction of the point-subspace cover instance we have:

- every point p^e is in fact covered by the lines ℓ^{v_I} where $e = \{v_I\}$ for all $I \in \binom{[d]}{\ell}$. This is because by construction every line ℓ^{v_I} and point p^e have the same x_j -coordinate with $j \notin I$.
- $\forall v \neq v' \in V$ it holds that $\ell^v \neq \ell^{v'}$. This is obviously true if vertices belong to different parts, because then there exists at least one free coordinate of v that is fixed for v' and vice versa. If $v, v' \in V_I$ for some $I \in \binom{[d]}{\ell}$ then, by vertex separability, there exists $j \notin I$ such that v and v' are not connected in graph G_j and get different x_j -coordinates. So they represent distinct subspaces.
- $\forall e \neq e' \in E$ it holds that $p^e \neq p^{e'}$. Indeed, by Lemma 2, G is edge-separable and by the definition of edge separability (see Definition 6) distinct hyperedges are not connected in at least one graph G_i and get different x_i -coordinates. So they represent distinct points.

By the above construction, for every incident vertex-hyperedge pair $v \in V, e \in E$, that is, $v \in e$, the corresponding geometric objects ℓ^v and p^e are incident as well. We claim that if v and e are not incident, that is, $v \notin e$ then ℓ^v and p^e are not incident as well. This is because every point p^e is already incident to precisely $\binom{d}{\ell}$ subspaces ℓ^v by construction, because the subspaces ℓ^v are pairwise distinct, and because p^e cannot be incident to more than $\binom{d}{\ell}$ axis-parallel subspaces. Thus, we have constructed a point-subspace cover instance that represents the hypergraph G and this means that G is representable.

(\Rightarrow) Assume that G is not vertex-separable but that it has a point-subspace cover representation. This means that it contains at least two distinct vertices v and v' from the same part V_I that are not separable. Then, for each coordinate $j \in [d] \setminus I$, there exists a v - v' path such that none of the vertices on this path has j as a free coordinate. All subspaces that represent the vertices from the path lie in the same hyperplane H_j perpendicular to the x_j -axis, i.e., x_j -coordinate is fixed and all other coordinates are free. This is because successive subspace pairs are joined by a common point (representing the hyperedge containing both successive vertices from the path). Since none of these subspaces is parallel to the x_j -axis, the x_j -coordinate stays fixed. Since this holds for all $j \in [d] \setminus I$, the subspaces ℓ^v and $\ell^{v'}$ lie in the intersection $\bigcap_{j \notin I} H_j$. But the intersection of such hyperplanes is a single subspace that has all coordinates $i \in I$ free. This contradicts that v and v' correspond to the distinct subspaces. \square

Analogously to the special case of $(d, 1)$ -hypergraphs the above proof leads directly to an algorithm computing a representation (or notifying about non-existence). A straightforward implementation of this algorithm gives a running time of $\mathcal{O}(d^{\ell d^{\ell}+2}(m^2 + n \log n))$ where the exponential factor in d comes from guessing the labeling of the parts. It is a very interesting question if the exponential dependence on d can be removed to get a polynomial-time algorithm also for non-constant d .

4 Conclusion and Outlook

There is a large amount of literature in algorithms and graph theory on hypergraph problems. This motivates various future research directions.

1. Can the structure of (d, ℓ) -hypergraphs be leveraged in (optimization) problems for hypergraphs such as matching or vertex cover?
2. What is the relation of the class of (d, ℓ) -hypergraphs to other classes of hypergraphs (for example geometrically representable hypergraphs)?

For example, maximum matching in k -hypergraphs is a very well-studied problem still exhibiting large gaps in our current understanding (Cygan, 2013; Hazan et al., 2006). Another example is hypergraph vertex cover, for which tight approximability results are known for general k -hypergraphs, as shown by Guruswami and Saket (2010). While the problem has been considered for $(k, 1)$ -hypergraphs before by Gaur and Bhattacharya (2007) (from the geometric perspective of point-line covering), no improvement upon the general case is known. We hope that our structural characterization can help to obtain such improvements. In the following, we state some first motivating results opening up these lines of research. Namely, we show that matching on $(3, 1)$ -hypergraphs is NP-hard and consider the relation of $(d, 1)$ -hypergraphs to point-line cover representations on the plane where we drop the requirement of axis-alignment. We hope that our structural characterization helps make progress on these questions.

4.1 Complexity of Matching on $(3, 1)$ -Hypergraphs

Finding maximum matchings in hypergraphs and in k -hypergraphs in particular is a classic optimization problem. The latter problem is also known as k -dimensional matching. This is a classic NP-hard problem, and thus its approximation algorithms have been studied. The currently best approximation algorithm for k -dimensional matching belongs to Cygan (2013) and has the ratio $(k + 1 + \epsilon)/3$ for an arbitrarily small

constant $\epsilon > 0$. Despite the problem is very well-studied, there is still a big gap to the best known lower bound of $\Omega(k/\log k)$ on the approximability by Hazan et al. (2006).

It is an interesting question if improved upper bounds can be obtained for (d, ℓ) -hypergraphs. As a first motivating result, we show in the following that indeed matching on $(k, 1)$ -hypergraphs is NP-hard for any $k \geq 3$. (For $k = 2$ the problem is matching on bipartite graphs, and thus its optimum can be found in polynomial time.)

Below we consider the matching problem for $(3, 1)$ -hypergraphs. Specifically, we are given a 3-uniform 3-partite vertex-separable hypergraph, i.e., a 3-hypergraph that can be represented as a point-line cover instance in \mathbb{R}^3 . The aim is to find a largest *3-dimensional matching*, that is, a set of hyperedges of the hypergraph such that no two hyperedges share a common vertex. We call the problem MATCHING ON $(3, 1)$ -HYPERGRAPHS.

The problem can be translated to point-line cover terms as follows. Given a finite set of points in \mathbb{R}^3 , the aim is to find a largest subset of points, such that no two lie on the same axis-parallel line.

In the following, we show NP-hardness of this problem by reduction from the decision version of MAXIMUM INDEPENDENT SET, which we describe next. Given a graph $G = (V, E)$ with $|V| = n$, the task is to decide whether there exists an *independent set* of size at least r , that is, a subset of vertices $I \subseteq V$ with $|I| \geq r$, such that no two vertices from I are adjacent. We denote by $\alpha(G)$ the size of a largest independent set of the graph G .

Theorem 3 *The decision version of MATCHING ON $(3, 1)$ -HYPERGRAPHS is NP-hard.*

Proof: Given a graph $G = (V, E)$ with n vertices and m edges, we construct a $(3, 1)$ -hypergraph $H(G)$ in polynomial time such that, for any integer $r \geq 0$, $H(G)$ has a matching of size at least $n^2 - 2n + r + 2m$ if and only if G has an independent set of size at least r .

Below, we create such an instance where we represent vertices by axis-parallel lines and hyperedges by points. First we introduce a vertex gadget. For every vertex $v \in V$, we create a set of $2(n-1)$ axis-parallel lines $\ell_1(v), \dots, \ell_{2n-2}(v)$ that lie on the same plane perpendicular to the z -axis. We require that, for each $i \in [2n-2]$, the line $\ell_i(v)$ is parallel to the x -axis, if i is odd, and to the y -axis, if i is even, see Fig. 5(a). Moreover, for each odd (even) $i \in [2n-3]$, the line $\ell_i(v)$ has smaller fixed y -coordinate (x -coordinate) than the line $\ell_{i+2}(v)$. For each $i \in [2n-3]$, we introduce a point at the intersection of the lines $\ell_i(v)$ and $\ell_{i+1}(v)$. We call the set of these $2n-3$ points the *vertex path* of v and denote it by $P(v)$. We say that the point of the vertex path that lies at the intersection of lines $\ell_i(v)$ and $\ell_{i+1}(v)$, for $i \in [2n-3]$, is of *type A* if i is odd, and of *type B*, otherwise. Let $A(v) \subseteq P(v)$ be the set of points of type A (see the blue disks at Fig. 5(a)) and let $B(v) \subseteq P(v)$ be the set of points of type B (see the green squares at Fig. 5(a)). Note that the points of $A(v)$ and the points of $B(v)$ alternate in $P(v)$, $|A(v)| = n-1$, and $|B(v)| = n-2$. Hence, both $A(v)$ and $B(v)$ are matchings in the set $P(v) = A(v) \cup B(v)$, and since $|A(v)| > |B(v)|$ it holds that $A(v)$ is the unique maximum matching in $P(v)$.

Now we introduce an edge gadget. For each edge $e = (u, v) \in E$ we create a set of five axis-parallel lines $\ell_1(e), \dots, \ell_5(e)$ connecting a point in $A(u)$ with a point in $A(v)$, such that $\ell_1(e)$ and $\ell_5(e)$ are parallel to the z -axis, $\ell_2(e)$ and $\ell_4(e)$ are parallel to the x -axis, and $\ell_3(e)$ is parallel to the y -axis, see Fig. 5(b). Moreover, two consecutive lines cross and the first and the last lines contain a point of $A(u)$ and a point of $A(v)$, respectively. For each $i \in [4]$, we introduce a point at the intersection of the lines $\ell_i(e)$ and $\ell_{i+1}(e)$. These four points (which we call *inner points*) together with the two points from $A(u)$ and $A(v)$ (which we call *endpoints*) create an *edge path* $P(e)$ of e . For each edge path we pick different endpoints,

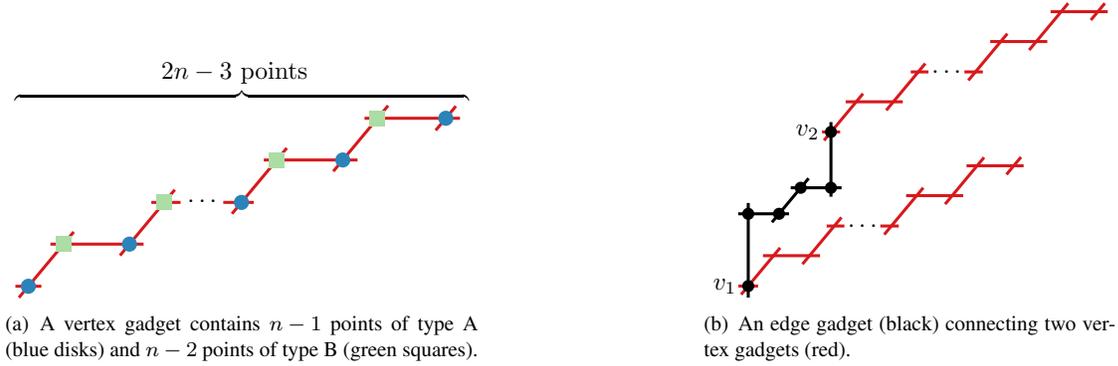


Fig. 5: A vertex (a) and an edge (b) gadgets.

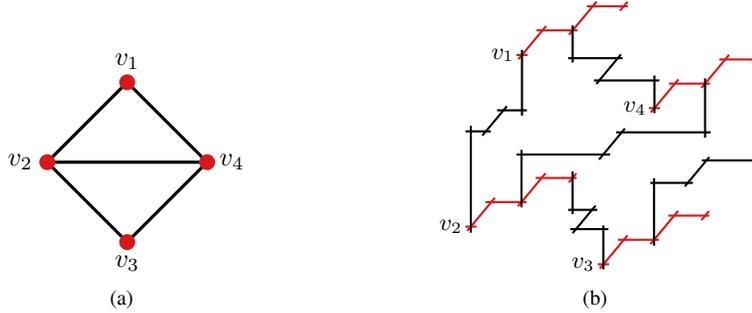


Fig. 6: Illustration of the reduction: a graph (on the left) and a corresponding point-line cover instance (on the right). The input points lie at the intersections of the lines drawn at the figure.

so that the edge paths are pairwise disjoint. This is possible because every vertex v has $|A(v)| = n - 1$ points of type A on its vertex path and is incident to at most $n - 1$ edges. Note that any matching of an edge gadget contains at most two inner points. Moreover, a matching can contain two inner points only if at most one of the endpoints is in this matching.

For an example of the construction of a matching instance from a given graph, see Fig. 6. We construct the instance in such a way that no two lines introduced above coincide and so that every vertex path lies on its own plane perpendicular to the z -axis.

It remains to prove that the instance $H(G)$ has a matching M of size at least $t = n^2 - 2n + r + 2m$ if and only if G has an independent set I of size at least r , for any integer $r \geq 0$.

Assume that there is an independent set I of G of size $r' \geq r$. We construct a matching M of size at least t as follows. For each vertex v of G , if $v \in I$, we add the set $A(v)$ to M ; otherwise, we add the set $B(v)$ to M . This ensures that, for any edge path corresponding to an edge (u, v) of G , at most one of its end points will be included in M , because at most one of the vertices u, v is in I . This allows us to add



Fig. 7: A not vertex-separable hypergraph and its representation on the plane with three directions.

two inner points of each edge path to M . Thus,

$$|M| = r' \cdot (n - 1) + (n - r') \cdot (n - 2) + 2m = n^2 - 2n + r' + 2m \geq t.$$

Now, assume conversely that $H(G)$ has a matching M of size at least t . We first modify M such that it remains a valid matching of size at least t but has a structure corresponding to an independent set in G as described above. Consider vertex paths of $H(G)$. Observe that if, for a vertex v of G , the matching M contains points both of $A(v)$ and $B(v)$, then $M \cap P(v)$ has size at most $n - 2$. In this case, we change the matching on $P(v)$ to $B(v)$ without decreasing its size. Thus, assume that $M \cap P(v) = B(v)$. Using this assumption, we can ensure that every vertex path has either only points of type A or only points of type B in M . Assume that there is an edge path $P((u, v))$ corresponding to some edge $(u, v) \in E$ where both endpoints are in M . This implies $M \cap P(u) = A(u)$ and $M \cap P(v) = A(v)$. Then $P((u, v))$ contains only *one* inner point that is in M . In this case, we change the matching M so that $M \cap P(u) = B(u)$ and so that $P((u, v))$ contains *two* inner points in M without decreasing the size of M overall. Hence, we can assume that, for any edge $(u, v) \in E$, we do not have $M \cap P(u) = A(u)$ and $M \cap P(v) = A(v)$ simultaneously. We can also assume that any edge path contains two inner points in M .

Let $I(M)$ be the set of vertices $v \in V$ where $P(v) \cap M = A(v)$. Note that $I(M)$ is an independent set by our assumption above. Since any edge path contains two inner points in M , the total size of M is $|I(M)|(n - 1) + (n - 2)(n - |I(M)|) + 2m$. This quantity is at least t only if $|I(M)| \geq r$, which completes the proof. \square

4.2 Plane Many-Directions Representations

To address the second research direction mentioned above, we give a first result relating the class of $(d, 1)$ -hypergraphs to other geometrically represented hypergraphs. Here we consider in particular point-line cover representations on the plane where we drop the requirement of axis-alignment.

Obtaining a many-directions representation for a $(d, 1)$ -hypergraph. Assume we are given an integer $d \geq 0$ and a $(d, 1)$ -hypergraph with an axis-aligned point-line cover representation P in \mathbb{R}^d . We create an equivalent point-line cover representation on the plane with d distinct directions. To this end, we create a random $2 \times d$ matrix M , where each entry is picked uniformly at random from the interval $[0, 1]$. Note that, for any two distinct vectors $x, y \in \mathbb{R}^d$, the probability that $Mx = My$ is 0. Hence, with probability 1, the set $P' = \{Mp \mid p \in P\} \subseteq \mathbb{R}^2$ has cardinality $|P|$. By the same reasoning, with probability 1, any two non-parallel lines in \mathbb{R}^d stay non-parallel under the mapping induced by M . Thus, the axis-parallel lines in \mathbb{R}^d covering the points in P are transformed into a set of distinct lines in the plane using precisely d distinct directions.

The converse does not hold. For an example of a hypergraph that is not vertex-separable (thus, does not have a representation in \mathbb{R}^3), but is representable on the plane with three directions, see Fig. 7.

Observe that the plane many-directions case is sensitive to the directions we allow. In the configuration in Fig. 8 there are four points and six directions of lines, namely, vertical, horizontal, and with the slopes 2, -2 , $1/2$, and $-1/2$. However, it is not hard to verify that it is not possible to represent an isomorphic hypergraph with the directions vertical, horizontal, and with the slopes 2, -2 , $1/2$, and $1/3$.

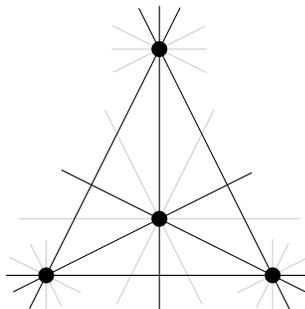


Fig. 8: An example of a many-directions representation with four points and six directions: vertical, horizontal, and slopes 2, -2 , $1/2$, $-1/2$.

References

- Ron Aharoni, Ron Holzman, and Michael Krivelevich. On a theorem of Lovász on covers in τ -partite hypergraphs. *Combinatorica*, 16(2):149–174, 1996. URL <https://doi.org/10.1007/BF01844843>.
- Hervé Brönnimann and Michael T. Goodrich. Almost optimal set covers in finite VC-dimension. 14(4): 463–479, 1995. URL <https://doi.org/10.1007/BF02570718>.
- Timothy M. Chan, Elyot Grant, Jochen Köneemann, and Malcolm Sharpe. Weighted capacitated, priority, and geometric set cover via improved quasi-uniform sampling. In Yuval Rabani, editor, *Proc. 23rd Annu. Symp. on Discrete Algorithms (SODA'12)*, pages 1576–1585. SIAM, 2012. URL <https://doi.org/10.1137/1.9781611973099.125>.
- Steven Chaplick, Krzysztof Fleszar, Fabian Lipp, Alexander Ravsky, Oleg Verbitsky, and Alexander Wolff. Drawing graphs on few lines and few planes. In Yifan Hu and Martin Nöllenburg, editors, *Proc. 24th Int. Symp. Graph Drawing & Network Vis. (GD'16)*, volume 9801 of *LNCIS*, pages 166–180. Springer, 2016. URL https://doi.org/10.1007/978-3-319-50106-2_14.
- Vasek Chvátal. A greedy heuristic for the set-covering problem. *Math. Oper. Res.*, 4(3):233–235, 1979. URL <https://doi.org/10.1287/moor.4.3.233>.
- Harold Scott MacDonald Coxeter. Self-dual configurations and regular graphs. *Bull. Amer. Math. Soc.*, 56(5):413–455, 09 1950. URL <https://doi.org/10.1090/S0002-9904-1950-08645-X>.
- Marek Cygan. Improved approximation for 3-dimensional matching via bounded pathwidth local search. In *Proc. 54th Annu. Symp. on Found. of Comp. Sci. (FOCS'13)*, pages 509–518. IEEE Comp. Soc., 2013. URL <https://doi.org/10.1109/FOCS.2013.61>.

- William S. Evans, Pawel Rzazewski, Noushin Saeedi, Chan-Su Shin, and Alexander Wolff. Representing graphs and hypergraphs by touching polygons in 3D. In Daniel Archambault and Csaba D. Tóth, editors, *Proc. 27th Int. Symp. Graph Drawing & Network Vis. (GD'19)*, volume 11904 of *LNCS*, pages 18–32. Springer, 2019. URL https://doi.org/10.1007/978-3-030-35802-0_2.
- Daya Ram Gaur and Binay Bhattacharya. Covering points by axis parallel lines. In *Proc. 23rd Europ. Workshop on Comput. Geom. (EuroCG'07)*, pages 42–45. Citeseer, 2007.
- Venkatesan Guruswami and Rishi Saket. On the inapproximability of vertex cover on k -partite k -uniform hypergraphs. In Samson Abramsky, Cyril Gavoille, Claude Kirchner, Friedhelm Meyer auf der Heide, and Paul G. Spirakis, editors, *Proc. 37th Int. Colloq. on Automata, Languages and Programming (ICALP'10)*, volume 6198 of *LNCS*, pages 360–371. Springer, 2010. URL https://doi.org/10.1007/978-3-642-14165-2_31.
- Refael Hassin and Nimrod Megiddo. Approximation algorithms for hitting objects with straight lines. *Discrete Appl. Math.*, 30(1):29–42, 1991. URL [https://doi.org/10.1016/0166-218X\(91\)90011-K](https://doi.org/10.1016/0166-218X(91)90011-K).
- Elad Hazan, Shmuel Safra, and Oded Schwartz. On the complexity of approximating k -set packing. *Comput. Complex.*, 15(1):20–39, 2006. URL <https://doi.org/10.1007/s00037-006-0205-6>.
- Lucian Ilie, Roberto Solis-Oba, and Sheng Yu. Reducing the size of NFAs by using equivalences and pre-orders. In Alberto Apostolico, Maxime Crochemore, and Kunsoo Park, editors, *Combinatorial pattern matching*, pages 310–321. Springer, 2005. URL https://doi.org/10.1007/11496656_27.
- Stefan Kratsch, Geevarghese Philip, and Saurabh Ray. Point line cover: the easy kernel is essentially tight. In Chandra Chekuri, editor, *Proc. 25th Annu. Symp. on Discrete Algorithms (SODA'14)*, pages 1596–1606. SIAM, 2014. URL <https://doi.org/10.1137/1.9781611973402.116>.
- V. S. Anil Kumar, Sunil Arya, and H. Ramesh. Hardness of set cover with intersection 1. In Ugo Montanari, José D. P. Rolim, and Emo Welzl, editors, *Proc. Int. Colloq. on Automata, Languages and Programming (ICALP'00)*, volume 1853 of *LNCS*, pages 624–635. Springer, 2000. URL https://doi.org/10.1007/3-540-45022-x_53.
- Terry A. McKee and F. R. McMorris. *Topics in intersection graph theory*. SIAM, 1999.
- Tomaž Pisanski and Milan Randić. Bridges between geometry and graph theory. In Catherine A. Gorini, editor, *Geometry at Work: Papers in Applied Geometry*, pages 174–194. Math. Assoc. of Amer., 2001.
- Kasturi R. Varadarajan. Weighted geometric set cover via quasi-uniform sampling. In Leonard J. Schulman, editor, *Proc. 42nd Symp. on Theory of Computing (STOC'10)*, pages 641–648. ACM, 2010. URL <https://doi.org/10.1145/1806689.1806777>.