Representing Hypergraphs by Point-Line Incidences

Alexander Dobler[1], Stephen Kobourov[2], William J. Lenhart[3], Tamara Mchedlidze[4], Martin Nöllenburg[1], and Antonios Symvonis[5]

1 **Algorithms and Complexity Group, TU Wien, Austria**
adobler@ac.tuwien.ac.at,noellenburg@ac.tuwien.ac.at
2 **University of Arizona, Tucson, AZ, USA**
kobourov@cs.arizona.edu
3 **Williams College, Williamstown, USA**
wlenhart@williams.edu
4 **Utrecht University, Utrecht, The Netherlands**
t.mtsentlintze@uu.nl
5 **NTUA, Athens, Greece**
symvonis@math.ntua.gr

Abstract

We consider hypergraph visualization that represent vertices as points and hyperedges as lines with few bends passing through points of their incident vertices. Guided by point-line incidence theory we show several theoretical results: if every vertex is part of at most two hyperedges, then we can find such a visualization without bends. There exist hypergraphs with three vertices per hyperedge and three hyperedges incident to each vertex requiring an arbitrary number of bends. It is $\exists\mathbb{R}$-hard to decide whether an arbitrary hypergraph can be visualized without bends. This only answers some interesting questions for such visualizations and we conclude with many open research questions.

1 Introduction

Hypergraphs arise in many domains and visualizing them is a non-trivial challenge. Classical approaches, such as Venn and Euler diagrams [1, 12, 15], do not scale to large instances. Recent experimental work [20] has shown that representing hypergraphs as collections of polylines (for the hyperedges) and common intersection points (for the vertices) allows for faster and more accurate performance of hypergraph-related tasks. In particular, the MetroSets approach [11] uses the metro map metaphor by representing each hyperedge with a metro line and each vertex as an interchange station. It attempts to visualize the result in the traditional octolinear fashion; see Figure 1. The visual complexity of the result depends on the total number of bends along the metro lines in the embedding. Minimizing the visual complexity makes the representations simpler to understand and work with. The natural question is to ask which hypergraphs can be represented with just one bendless line per hyperedge. Naturally, only hypergraphs such that hyperedges share at most one vertex – called *linear hypergraphs* – can be represented in such a way. Otherwise, lines of hyperedges could coincide and could not be distinguished. Further, the *rank* of a hypergraph is the maximum cardinality of a hyperedge; the *degree* is defined equivalently as for classic graphs. With this in mind we show that:

- Maximum degree two linear hypergraphs can be represented with one line per hyperedge.
- Not all rank-three maximum degree-three linear hypergraphs can be represented with one line per hyperedge. In fact, there is a family of such hypergraphs requiring an arbitrary number of bends.
- Determining whether an arbitrary-rank linear hypergraph can be represented with one line per hyperedge is $\exists\mathbb{R}$-hard.

40th European Workshop on Computational Geometry, Ioannina, Greece, March 13–15, 2024.
This is an extended abstract of a presentation given at EuroCG'24. It has been made public for the benefit of the community and should be considered a preprint rather than a formally reviewed paper. Thus, this work is expected to appear eventually in more final form at a conference with formal proceedings and/or in a journal.

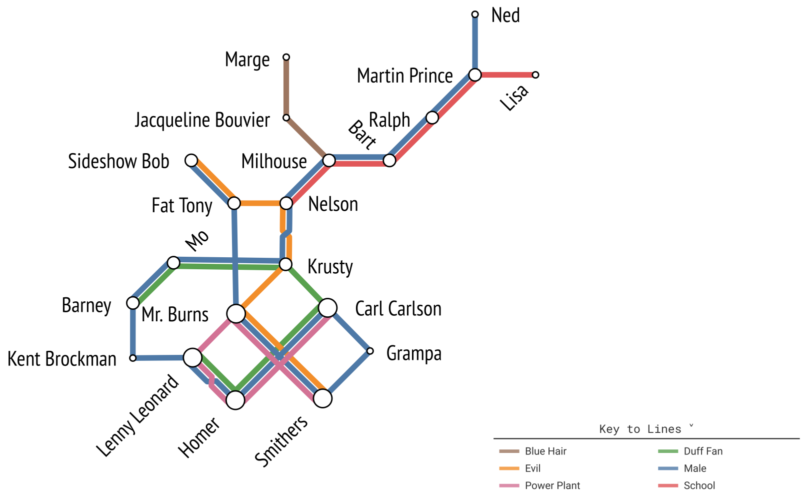

Figure 1 A visualization of a Simpsons hypergraph dataset using the MetroSets metaphor [11], the visualization is taken from `https://metrosets.ac.tuwien.ac.at/`. Hyperedges are represented by metro lines and elements are represented by stations.

Here, lines are infinite lines and not line segments.

2 Related Work

Representing hypergraphs with one line per hyperedge is related to classical geometric problems going as far back as the 19th century. In particular, it is related to the study of *configurations*, a set of points and an arrangement of lines, such that each point is incident to the same number of lines and each line is incident to the same number of points [9,10,16]. This was the main topic of the PhD thesis of Steinitz [18] and was of interest to many with notable examples including the configurations of Desargues, Pappus, and Möbius–Kantor [10]. If the representation of a hypergraph should however form a non-crossing straight-line drawing of a tree, then it can be decided in polynomial time whether such a representation exists [19]. Graph-based techniques for drawing hypergraphs via support graphs [2–4] have a different focus and do not take into account geometric straightness or bends of hyperedges.

If we ask for crossing-free representations of hypergraphs with line segments instead of lines, there is only one line of research that we are aware of [6,8]. Namely, Gonçalves [8] has shown that there are planar linear hypergraphs which cannot be represented with straight line segments (in contract to planar graphs which can always be drawn with straight lines).

The problem seems to be related to stretchability of pseudolines [17], but is different because the order of the vertices along each hyperedge is not specified. A similar $\exists\mathbb{R}$-hard problem, called matroid representability [13] will be used to show $\exists\mathbb{R}$-hardness of one of the problems studied in this paper. We will formally explain matroid representability in the corresponding Section 6.

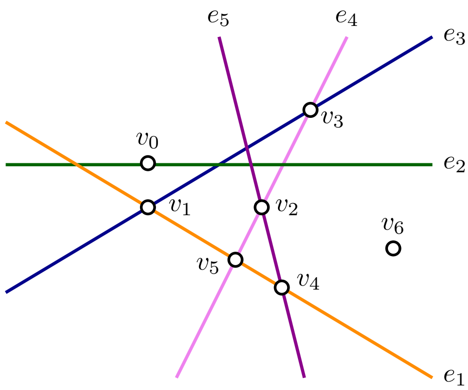

Figure 2 Illustration of a max-degree-2 hypergraph $H = (V, E)$ with $V = \{v_0, \ldots, v_6\}$ and $E = \{e_1, \ldots, e_5\}$, where $e_1 = \{v_1, v_4, v_5\}$, $e_2 = \{v_0\}$, $e_3 = \{v_1, v_3\}$, $e_4 = \{v_2, v_3, v_5\}$, $e_5 = \{v_2, v_4\}$.

3 Preliminaries

A hypergraph $H = (V, E)$ is defined by a vertex set V and a hyperedge set E, where each $e \in E$ is a subset of V. The *degree* of a vertex v is the number of hyperedges containing v. The *rank* of a hypergraph is the maximum cardinality of a hyperedge $|e|$ taken over all e in E; hence a rank-2 hypergraph is an ordinary graph. A hypergraph is k*-uniform* if very hyperedge has cardinality exactly k and it has degree k if every vertex has degree k. A hypergraph is *linear* if $|e \cap e'| \leq 1$ for every pair of distinct hyperedges $e, e' \in E$. A representation of a hypergraph consists of an injective mapping α of vertices to points in $\mathbb{R}^2$ and an injective mapping β of hyperedges to lines in $\mathbb{R}^2$ such that $v \in e$ if and only if $\alpha(v) \in \beta(e)$ for $v \in V, e \in E$.

4 Max-Degree-2 Hypergraphs

▶ **Theorem 4.1.** *There exists a representation for any max-degree-2 linear hypergraph.*

Proof. Let H be a hypergraph with n vertices $v_1, \ldots, v_n$ and m hyperedges $e_1, \ldots, e_m$. Consider m lines $\ell_1, \ldots, \ell_m$ on $\mathbb{R}^2$ in general position, such that any two of these lines cross; see Figure 2. Let v be a vertex of H. If the degree of v is exactly 2, then we place v at the intersection of ℓ_i and ℓ_j, where ℓ_i and ℓ_j are the lines corresponding to the hyperedges e_i and e_j containing v. If the degree of v is 1, then we place v at any point of the line corresponding to the hyperedge containing v that is not an intersection point of the lines $\ell_1, \ldots, \ell_m$. If the degree of v is zero, then we place v at any point of $\mathbb{R}^2$ that is not on any of the lines $\ell_1, \ldots, \ell_m$. This yields a representation of H with one line per hyperedge. ◀

5 Rank-3 Degree-3 Hypergraphs

The PhD thesis of Ernst Steinitz [18] claims that every 3-uniform degree-3 hypergraph can be represented with one line per hyperedge, except maybe of one hyperedge (which could be represented as a polyline with one bend). However, more careful consideration shows that this is indeed not true as already pointed out by Grünbaum [10]. Similar results exist, but none show the claim of Steinitz [7, 14]. We show a construction that has at least two hyperedges that must have a bend and this construction can be generalized. For this, we

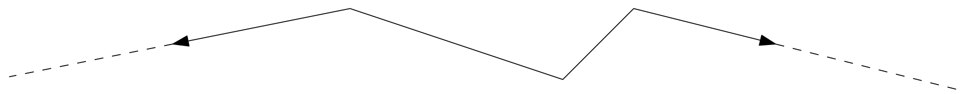

Figure 3 Illustration of an infinite polygonal chain with 3 bends that extends infinitely in the left and right direction.

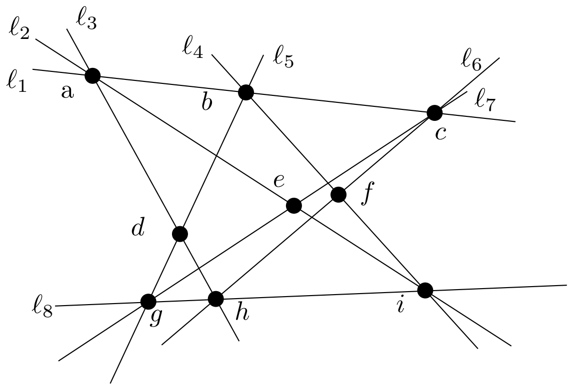

Figure 4 The Pappus configuration without the line through d, e, f.

define for $t \in \mathbb{N}_0$, t-bend representations for hypergraphs H as follows. In the original definition of representations, we replace lines by what we call *infinite polygonal chains*. An infinite polygonal chain consists of a (possibly empty) polygonal chain and two rays, one ending at the first point of the chain, one ending at the last point (see Figure 3). Essentially, an infinite polygonal chain is a polyline whose first and last segment is extended infinitely. Further, two distinct infinite polygonal chains $\beta(e), \beta(e')$ for $e, e' \in E$ must not share a line segment nor a bend point. Lastly, we require that the total number of all bends in $\beta(E)$ is exactly t. A representation in the original sense is hence a 0-bend representation.

Consider the 3-uniform degree-3 hypergraph H defined as follows. Let H_1 be the hypergraph depicted in Figure 4 with points $a, b, \ldots, i$ and hyperedges defined by the lines $\ell_1, \ell_2, \ldots, \ell_8$. Pappus's theorem [5, Chapter 3.5] says that in any representation of H_1, the points d, e, f must be collinear[1]. Now let H_2 be a copy of H_1 with the points $a', b', \ldots, i'$ and hyperedges defined equivalently. The hypergraph H is the union of H_1, H_2, and the two hyperedges $\{d, e, f'\}$ and $\{d', e', f\}$.

▶ **Lemma 5.1.** *There is no t-bend representation for H with $t < 2$.*

Proof. If every hyperedge in the subhypergraph H_1 is represented without a bend, then $\beta(\{d, e, f'\})$ must pass through f. Thus, at least one hyperedge of H_1 must be represented with at least one bend. Applying the same argument to H_2 we see that we require at least two bends. ◀

The above construction can be generalized so that there is no t-bend representation for any $t < x$ for an arbitrary $x \in \mathbb{N}$. Instead of one copy of H_1, we add $x - 1$ copies $H_2, H_3, \ldots, H_x$. Further we add x hyperedges $e_1, e_2, \ldots, e_x$ such that e_i contains d, e of H_i and f of $H_{(i \bmod x)+1}$.

[1] Formally, Pappus's theorem is as follows: let points a, b, c be on one line, and points g, h, i be on another line. Let the three points d, e, f be defined by the intersections of line ah with bg, ai with cg, and bi with ch, respectively. Then d, e, f are collinear. It is clear that our formulation is equivalent.

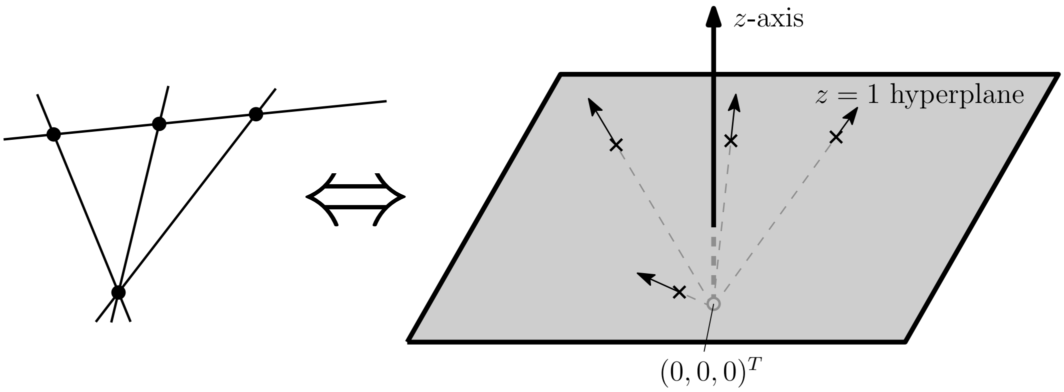

Figure 5 Illustration of the equivalence between representations of H and representations of M. Lines on the left side correspond to linear subpaces of dimension two that are common to at least two vectors on the right side.

▶ **Theorem 5.2.** *For any $x \in \mathbb{N}$ there exists a rank-3 hypergraph H such that there exists no t-bend representation for H for any $t < x$.*

Similar constructions exist (at least for $x = 2$) [10], the generalization has not been stated explicitly.

6 $\exists\mathbb{R}$-hardness

In this section we want to show that it is $\exists\mathbb{R}$-hard to decide whether there exists a representation for a given hypergraph H. We reduce from a problem that lends itself for a nice reduction, called MATROID REPRESENTABILITY [13]. For the purposes of a simple description, we give here a simplified description of a variant of that problem that is still $\exists\mathbb{R}$-hard [13]. We start with definitions. A matroid M is given as $M = (X, \mathcal{I})$ where X is the finite *ground set* and $\mathcal{I} \subseteq 2^X$ is the set of independent sets with

1. $\emptyset \in \mathcal{I}$,
2. $I' \subset I \in \mathcal{I}$ implies $I' \in \mathcal{I}$, and
3. $I_1, I_2 \in \mathcal{I}$ with $|I_1| < |I_2|$ implies that there is an $x \in I_2 \setminus I_1$ with $I_1 \cup \{x\} \in \mathcal{I}$.

A representation of M is an injective mapping $f(X) : X \to \mathbb{R}^3$ such that for any $Y \subseteq X$ we have $Y \in \mathcal{I}$ if and only if $f(Y)$ forms a set of linearly independent vectors in $\mathbb{R}^3$. The $\exists\mathbb{R}$-hard problem MATROID REPRESENTABILITY is given as input a matroid and the question is whether there is a representation f of M. For the vectors $v \in \mathbb{R}^3$ we call the first, second, and third coordinate the x, y, and z-coordinates, respectively.

We start by making some normalizations to the instance M. First, we can assume that every independent set $I \in \mathcal{I}$ has cardinality of at most 3, as there is otherwise no representation. Second, we can assume that each pair $x, x' \in X$ of distinct elements forms an independent set, i.e. $\{x, x'\} \in \mathcal{I}$. Otherwise, $f(x) = cf(x')$ for some $c \in \mathbb{R}$ must hold for any representation. We can remove x' from X and replace any occurence of x' in $\mathcal{I}$ by x, and obtain an equivalent instance w.r.t. representability.

We are ready to state our reduction. The main idea is to identify the vectors $f(X)$ by points in the plane due to a projective transformation; see Figure 5.

▶ **Theorem 6.1.** *It is $\exists\mathbb{R}$-hard to decide whether a linear hypergraph can be represented.*

Proof. We reduce from MATROID REPRESENTABILITY. We are given a matroid $M = (X, \mathcal{I})$ (we assume both above normalizations were applied already) and transform it to a hypergraph $H = (V, E)$ as follows. First, we set $V = X$. The hyperedges E are defined as follows. Let $Y \subseteq X$ be a maximal subset of X that must lie in the same linear subspace U of $\mathbb{R}^3$ (which by definition contains the origin) with $\dim U = 2$ for any representation f of M. Add Y to E. All such Y can be determined in polynomial time: first, we can determine in polynomial time all triples $x, x', x'' \in X$ that do not form an independent set. If two triples share two elements, the union of both triples forms a four-tuple that has to be part of a common subspace in any representation. Continuing this process greedily, we can find all such Y. This process terminates with the desired sets Y, as the matroid M specifies for each triple exactly if the vectors of their representation are dependent or not. Notice also that there can only exist polynomially many such Y as each pair of $x, x' \in X$ can be part of at most one such Y. Lastly, we have to add pairs $\{x, x'\}$, $x, x' \in X$, to E, that are not part of some Y defined above. The vectors $f(x), f(x')$ certainly have to span a subspace U of dimension two not containing any other vectors $f(x'')$.

We claim that M has a representation f if and only if H has a representation (α, β).

"$\Rightarrow$": let f be a representation of M. First, if any $f(x')$, $x' \in X$ has z-coordinate 0 we multiply all $f(x)$, $x \in X$, by the same rotation matrix $R \in SO(3) \subseteq \mathbb{R}^{3\times 3}$ such that no $f(x)$ has z-coordinate 0. This is possible as X is finite. Second, we scale each $f(x)$ by 1 divided by its z-coordinate so that the z-coordinates of all vectors in $f(X)$ are 1. The new f still is a representation. For $x \in X$, we set $\alpha(x)$ equal to the point defined by the first two coordinates of $f(x)$. Essentially, we applied a projective transformation. Lastly, for each hyperedge $e \in E$ we set $\beta(e)$ to the line through any two distinct points $\alpha(x)$ and $\alpha(x')$ with $x, x' \in e$ (E does not contain edges of size 1). It is now easy to verify that (α, β) is a representation of H: let $e \in E$. For each $x \in e$ the point $\alpha(x)$ must lie on the line $\beta(e)$ because $f(x)$ must be in the same linear subspace of dimension two as all $f(x'), x' \in e$. No other point can lie on that line, as e would not have been maximal for our construction of Y above.

"$\Leftarrow$": let (α, β) be a representation of H. For $x \in X$, let $\alpha(x) = (r, s)^T$. We set $f(x) = (r, s, 1)^T$ and claim that f is a representation of M. Indeed, if $Y \subseteq X$ does not form an independent set in M, then $Y \subseteq e$ for some $e \in E$. Thus, points $\alpha(Y)$ lie on the same line and vectors $f(Y)$ are in the same linear subspace of dimension two and are thus dependent as $|Y| \geq 3$. If $Y \in \mathcal{I}$ there are two cases.

- If $|Y| \leq 2$, $f(Y)$ is clearly independent as α is injective.
- If $|Y| = 3$, let $Y = \{x_1, x_2, x_3\}$. Because of the construction of E, there exists $e \in E$ such that $x_1, x_2 \in e$ and $x_3 \notin e$. Thus $\alpha(Y)$ forms a triangle and $f(Y)$ is independent. ◀

7 Conclusions and Open Problems

Motivated by a hypergraph visualization using polylines for hyperedges [11], we initiated the investigation of visualizations with hyperedges drawn with as few bends as possible. We provide results for special classes of hypergraphs. If the maximum vertex degree is 2, any linear hypergraph can always be drawn without a bend. For rank-3 linear hypergraphs arbitrarily many bends may be required. Lastly, it is even $\exists\mathbb{R}$-hard to decide whether an arbitrary linear hypergraph can be drawn without bends.

Our results are highly inconclusive and many open research directions are possible.

- We have considered representations with lines and infinite polygonal chains as we could use results from point-line incidence theory. If we replace these with line segments and polygonal chains, respectively, our results from Sections 5 and 6 do not extend. It may as well be that for such definitions each connected 3-uniform degree-3 hypergraph (besides the Fano and Möbius-Kantor configurations [10]) has a representation with line segments without bends, and it might be that it is "easy" to decide whether such a line segment representation exists.

- Our $\exists\mathbb{R}$-hardness result from Section 6 works for arbitrary rank hypergraphs. Can we state a similar hardness result for bounded-rank linear hypergraphs, e.g. rank 3?
- Are there polynomial-time algorithms to represent a 3-uniform degree-3 linear hypergraph in such a way that the number of bends in the representation is a constant factor away from optimum (minimum number of bends)? Are there families of 3-uniform degree-3 hypergraphs requiring more than a linear number of bends?

Acknowledgments.

This work started at the Bertinoro Workshop on Graph Drawing 2023. We thank Michael A. Bekos for the discussions.

References

1 Bilal Alsallakh, Luana Micallef, Wolfgang Aigner, Helwig Hauser, Silvia Miksch, and Peter J. Rodgers. The state-of-the-art of set visualization. *Comput. Graph. Forum*, 35(1):234–260, 2016. URL: `https://doi.org/10.1111/cgf.12722`, `doi:10.1111/CGF.12722`.

2 Ulrik Brandes, Sabine Cornelsen, Barbara Pampel, and Arnaud Sallaberry. Path-based supports for hypergraphs. *J. Discrete Algorithms*, 14:248–261, 2012. URL: `https://doi.org/10.1016/j.jda.2011.12.009`, `doi:10.1016/J.JDA.2011.12.009`.

3 Kevin Buchin, Marc van Kreveld, Henk Meijer, Bettina Speckmann, and Kevin Verbeek. On planar supports for hypergraphs. *J. Graph Algorithms Appl.*, 15(4):533–549, 2011. `doi:10.7155/jgaa.00237`.

4 Thom Castermans, Mereke van Garderen, Wouter Meulemans, Martin Nöllenburg, and Xiaoru Yuan. Short plane supports for spatial hypergraphs. *J. Graph Algorithms Appl.*, 23(3):463–498, 2019. `doi:10.7155/jgaa.00499`.

5 Harold Scott Macdonald Coxeter and Samuel L Greitzer. *Geometry revisited*, volume 19. Mathematical Association of America, 1967.

6 Hubert de Fraysseix and Patrice Ossona de Mendez. Representations by contact and intersection of segments. *Algorithmica*, 47(4):453–463, 2007. `doi:10.1007/S00453-006-0157-X`.

7 David G. Glynn. On the representation of configurations in projective spaces. *Journal of Statistical Planning and Inference*, 86(2):443–456, May 2000. `doi:10.1016/S0378-3758(99)00124-X`.

8 Daniel Gonçalves. A planar linear hypergraph whose edges cannot be represented as straight line segments. *Eur. J. Comb.*, 30(1):280–282, 2009. URL: `https://doi.org/10.1016/j.ejc.2007.12.004`, `doi:10.1016/J.EJC.2007.12.004`.

9 Harald Gropp. Configurations and their realization. *Discret. Math.*, 174(1-3):137–151, 1997. `doi:10.1016/S0012-365X(96)00327-5`.

10 Branko Grünbaum. *Configurations of Points and Lines.* American Mathematical Society, 2009.

11 Ben Jacobsen, Markus Wallinger, Stephen G. Kobourov, and Martin Nöllenburg. Metrosets: Visualizing sets as metro maps. *IEEE Trans. Vis. Comput. Graph.*, 27(2):1257–1267, 2021. `doi:10.1109/TVCG.2020.3030475`.

12 David S. Johnson and H. O. Pollak. Hypergraph planarity and the complexity of drawing Venn diagrams. *J. Graph Theory*, 11(3):309–325, 1987. `doi:10.1002/jgt.3190110306`.

13 Eunjung Kim, Arnaud de Mesmay, and Tillmann Miltzow. Representing matroids over the reals is $\exists\mathbb{R}$-complete. *CoRR*, abs/2301.03221, 2023. `doi:10.48550/arXiv.2301.03221`.

14 William Kocay and Ryan Szypowski. The application of determining sets to projective configurations. *Ars Combinatoria*, 53:193–208, 1999.

15 Erkki Mäkinen. How to draw a hypergraph. *Int. J. Computer Math.*, 34(3–4):177–185, 1990. doi:10.1080/00207169008803875.

16 Tomaž Pisanski and Brigitte Servatius. Configurations from a graphical viewpoint. *Ars Math. Contemp.*, 5(2), 2012. URL: https://amc-journal.eu/index.php/amc/article/view/342/193.

17 Peter W. Shor. Stretchability of pseudolines is NP-hard. In Peter Gritzmann and Bernd Sturmfels, editors, *Proc. Applied Geometry And Discrete Mathematics (DIMACS 1990)*, volume 4 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 531–554. DIMACS/AMS, 1990. doi:10.1090/DIMACS/004/41.

18 Ernst Steinitz. *Über die Construction der Configurationen n_3*. PhD thesis, Breslau, 1894.

19 R. Swaminathan and Donald K. Wagner. On the consecutive-retrieval problem. *SIAM J. Comput.*, 23(2):398–414, 1994. doi:10.1137/S0097539792235487.

20 Markus Wallinger, Ben Jacobsen, Stephen G. Kobourov, and Martin Nöllenburg. On the readability of abstract set visualizations. *IEEE Trans. Vis. Comput. Graph.*, 27(6):2821–2832, 2021. doi:10.1109/TVCG.2021.3074615.