Consistent Sets of Lines with no Colorful Incidence

Boris Bukh[1]
Carnegie Mellon University, Department of Mathematical Sciences
Pittsburgh, PA 15213, USA.
bbukh@math.cmu.edu

Xavier Goaoc[2]
Université Paris-Est, LIGM)
UMR 8049, CNRS, ENPC, ESIEE, UPEM, F-77454, Marne-la-Vallée, France.
xavier.goaoc@u-pem.fr

Alfredo Hubard
Université Paris-Est, LIGM
UMR 8049, CNRS, ENPC, ESIEE, UPEM, F-77454, Marne-la-Vallée, France.
alfredo.hubard@u-pem.fr

Matthew Trager[3]
PSL Research University
Inria, École Normale Supérieure, and CNRS.
matthew.trager@inria.fr

Abstract

We consider incidences among colored sets of lines in $\mathbb{R}^d$ and examine whether the existence of certain concurrences between lines of k colors force the existence of at least one concurrence between lines of $k+1$ colors. This question is relevant for problems in 3D reconstruction in computer vision.

2012 ACM Subject Classification Theory of computation → Randomness, geometry and discrete structures, Computing methodologies → Scene understanding

Keywords and phrases Incidence geometry, image consistency, probabilistic construction, algebraic construction, projective configuration

Digital Object Identifier 10.4230/LIPIcs.SoCG.2018.17

Related Version A full version of this paper is available at `https://arxiv.org/abs/1803.06267`.

Funding Part of the work was done during the visit of BB to Université Paris-Est Marne-la-Vallée supported by LabEx Bézout (ANR-10-LABX-58).

Acknowledgements We thank Éric Colin de Verdière and Vojta Kalusza for discussions at an early stage of this work.

[1] Supported in part by Sloan Research Fellowship and by U.S. taxpayers through NSF CAREER grant DMS-1555149.
[2] Supported by Institut Universitaire de France.
[3] Supported in part by the ERC grant VideoWorld and the Institut Universitaire de France

34th International Symposium on Computational Geometry (SoCG 2018).
Editors: Bettina Speckmann and Csaba D. Tóth; Article No. 17; pp. 17:1–17:14
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

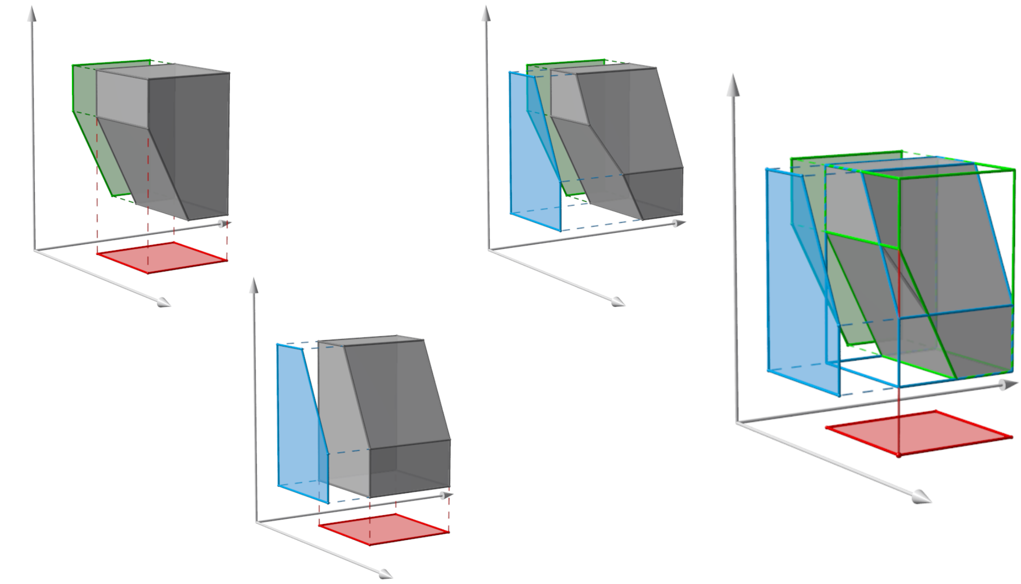

Figure 1 Three silhouettes that are 2-consistent but not globally consistent for three orthogonal projections. Each of the first three figures shows a three-dimensional set that projects onto two of the three silhouettes. The fourth figure illustrates that no set can project simultaneously onto all three silhouettes: the highlighted red image point cannot be lifted in 3D, since no point that projects onto it belongs to the pre-images of both the blue and green silhouettes.

1 Introduction

A central problem in computer vision is the reconstruction of a three-dimensional scene from multiple photographs. Trager et al. [16, Definition 1] defined a set of images as consistent if they represent the same scene from different points of view. They constructed examples (like that of Figure 1) of a set of images which is pairwise consistent while being altogether inconsistent. They also showed [16, Proposition 4] that under a certain convexity hypothesis, images that are consistent three at a time are globally consistent. In this paper we drop the convexity condition and consider these affairs from the point of view of incidence geometry.

Problem statement. An *incidence* is a set of lines that meet at a single point. Let $\mathcal{L} = \mathcal{L}_1 \cup \mathcal{L}_2 \cup \ldots \cup \mathcal{L}_m$ be a set of lines of m colors in $\mathbb{R}^d$ (where each $\mathcal{L}_i$ is a color class). Given $S \subset \{1, 2 \ldots m\}$, an *$S$-incidence* in $\mathcal{L}$ is an incidence between lines of every color in S. This paper focuses on the following notions:

▶ **Definition 1.** For $1 \leq k \leq m$, a *k-incidence* in $\mathcal{L}$ is a S-incidence where $|S| = k$. A *colorful incidence* in $\mathcal{L}$ is an incidence that contains lines of every color.

▶ **Definition 2.** The set $\mathcal{L}$ is *k-consistent* if for every k-tuple of colors $S \subset \{1, 2 \ldots m\}$, every line in $\cup_{i \in S} \mathcal{L}_i$ belongs to an S-incidence. The set $\mathcal{L}$ is *consistent* if every line belongs to (at least) one colorful incidence.

Instead of wondering if k-consistency implies consistency, we aim for a more modest goal:

▶ **Problem 3.** *Under which conditions does the k-consistency assumption imply the existence of a $(k+1)$-incidence?*

The main results of this paper are two constructions of (infinite families of) finite sets of lines which are k-consistent and have no colorful incidence. Thus, consistency does not propagate.

▶ Remark. Unless indicated otherwise, the set $\mathcal{L}$ is assumed to be finite. We also assume throughout that the lines in $\mathcal{L}$ are pairwise distinct. This has no consequence on Problem 3: repeating a line in a color class is useless, and if two lines of distinct colors coincide, then the k-consistency assumption trivially implies that this line has a $(k+1)$-incidence.

Relation to photograph consistency. Let us explain how our initial image consistency question relates to Problem 3. Firstly, we ignore color or intensity information, and treat the scene as a set of opaque objects and the images as their projections onto certain planes. In this setting, images are *consistent* if and only if there exists a subset $R \subset \mathbb{R}^3$ that projects into each of them. Assuming that light travels along straight lines, the set of 3D points that are mapped to a given image point is a ray, or more conveniently a line, in $\mathbb{R}^3$. Starting with m photographs, if we let $\mathcal{L}_i$ denote the lines that are pre-images of the projection on the ith photograph, then the photographs are consistent if and only if $\cup_{i=1}^m \mathcal{L}_i$ is consistent: R is the set of points of colorful incidences.

Setting. In the basic set-up for computer vision, all lines used to project the scene onto a given image plane pass through a "pinhole". We therefore define a color class $\mathcal{L}_i$ as *concurrent* if it consists of concurrent lines. We consider, however, the problem more generally since it is possible to build other imaging systems. For example, there are cameras that use the lines secant to two fixed skew lines; other cameras use the lines secant to an algebraic curve γ and to a line intersecting γ in $\deg \gamma - 1$ points. For a discussion of the geometry of families of lines arising in the modeling of imaging systems, see [2, 17] and the references therein.

We focus in this paper on the consistency of finite sets of lines. This restriction is technically convenient and remains relevant to the initial motivation on continuous sets of lines. On the one hand, our constructions for the finite problem turn out to readily extend to infinite families of lines (see Section 4). On the other hand, the finite problem is already relevant to 3D reconstruction, when one has to recover the camera parameters (settings or position) used in the photographs. Indeed, this recovery is typically done by identifying pixels in different images that are likely to be the projection of the same 3D element, and using the incidence structure of their inverse images to infer the position of the camera; this process is called *structure from motion* [14]. The number of lines required to determine the cameras is typically 5 to 7 per image. Although pixels are usually matched across pairs of images, there are good reasons for wanting to match them across more images, firstly for robustness to noise, but also because this avoids ambiguities in the reconstruction in the case of degenerate camera configurations (for example, pairwise matches are never sufficient to reconstruct a scene from images when all the camera pinholes are exactly aligned [12, Chapter 15.4.2]). Understanding the consistency propagation may simplify the certification of such matchings.

1.1 Results

We focus on Problem 3 for $k \geq 3$ because examples of tricolor sets of lines that are 2-consistent but without a colorful incidence are relatively easy to build:

▶ **Example 4.** Let $(\vec{x}_0, \vec{x}_1, \vec{x}_2)$ be a basis of $\mathbb{R}^3$. Let $p_0, p_2 \ldots, p_{3n}$ be a set of points where p_0 is arbitrary, $p_{i+1} \in p_i + \mathbb{R}\vec{x}_{(i \bmod 3)}$ and $p_{3n} = p_0$. For each $i \in \{0,1,2\}$ define $\mathcal{L}_i$ to be the set of lines in coordinate direction i that are incident to points p_j with $j \equiv i \pmod 3$ and $j \equiv i-1 \pmod 3$. If desired, we may apply a projective transformation that turns parallelism into concurrence.

Constructions from higher-dimensional grids. We present two constructions of arbitrary large sets of lines in $\mathbb{R}^d$ of $k+1$ colors that are k-consistent and have no colorful incidence, for every $k \geq 3$ and $k+1 \geq d \geq 2$. Both constructions are based on selecting subsets of lines from a regular grid in $\mathbb{R}^{k+1}$. In one case, the selection is probabilistic (Theorems 5), while in the other case it uses linear algebra over finite vector spaces (Theorem 6). In both constructions, every color class is concurrent. The probabilistic argument is asymptotic and proves the existence of configurations where every line is involved in many k-incidences for every choice of $k-1$ other colors. The algebraic construction is explicit and is minimal in the sense that removing any line breaks the k-consistency.

Restrictions on higher-dimensional grids. We then test the sharpness and potential of constructions from higher-dimensional grids. On the one hand, we examine the number of lines of such constructions. The algebraic selection method picks at least 2^{k^2-k-1} lines of each color (we leave aside the probabilistic selection method as its analysis is asymptotic). This construction has the property that the lines meeting at a k-incidence are not "flat", in the sense that they are not contained in a $k-1$-dimensional subspace. We show, using the polynomial method [11], that for any construction with this property, the number of lines must be at least exponential in k (Proposition 7). On the other hand, we examine the possibility of designing similar constructions for models of cameras in which the lines are not all concurrent. We observe that when every color class is secant to two fixed lines, lines from two color classes cannot form a complete bipartite intersection graph (Proposition 8).

Small configurations. We also investigate small-size configurations of lines in $\mathbb{R}^3$ with 4 colors that are 3-consistent but have no colorful incidence. The smallest example provided by our constructions has 32 lines per color, which says little for applications like structure from motion, where each color class has very few lines. Figure 2 shows two non-planar examples with 12 lines each. We prove that they are the only non-planar constructions with these parameters (Theorem 11). We also show that any configuration with these parameters and concurrent color classes must have at least 24 lines or be planar (Theorem 10).

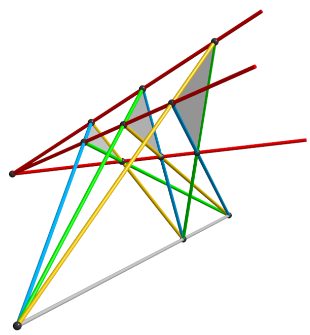

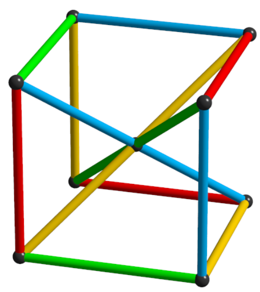

Figure 2 Two non-planar examples of 12 lines in 4 colors that are 3-consistent and have no 4-incidence. (Left) A variation around Desargues' configuration. (Right) A subset of the $(12_4 16_3)$ configuration of Reye; note that triples of parallel lines intersect at infinity.

1.2 Related work

The study of consistent families of colored lines relates most prominently to classical questions in computer vision and in discrete geometry.

In computer vision. The simplest and most extensively studied setting for consistency deals with families where each color class has a single line. The study of n-tuples of lines that are incident at a point (or "point correspondences"), is central in *multi-view geometry* [12], that is the foundation of 3D-reconstructions algorithms. In this setting, consistency propagates trivially: n lines are concurrent if and only if any three of them are (even better: n lines not all coplanar are concurrent if and only if every pair of them is). Concurrency constraints are traditionally expressed algebraically as polynomials in image coordinates (see, *e.g.* [6]).

A more systematic study of consistency for silhouettes (*i.e.*, for infinite families of lines) was proposed to design reconstruction methods based on shapes more complex than points or lines [3, 13]. Pairwise consistency for silhouettes can be encoded in a "generalized epipolar constraint", which can be viewed as an extension of the epipolar constraint for points, and expresses 2-consistency in terms of certain simple tangency conditions [1, 16]. There is no known similar characterization for k-consistency with $k > 2$. Consistency propagation is only known for convex silhouettes: 3-consistency implies consistency [16].

In the dual, consistency expresses conditions for a family of planar sets to be sections of the same 3D object [16], a question classical in geometric tomography or stereology. We are not aware of any relevant result on consistency in these directions.

Discrete geometry. As evidenced by Figure 2, our analysis of small configurations relates to the classical configurations of Reye and Desargues in projective geometry. Our problem and results for larger configurations relate to various lines of research in incidence geometry. Inspired by the Sylvester–Gallai theorem, Erdős [5] asked for the largest number of collinear k-tuples in a planar point set with no collinear $k+1$-tuple. The best construction for $k = 3$ come from irreducible cubic curves[4]. For higher k the best constructions were given by Solymosi and Stojakovíc [15] and are projections of higher-dimensional subsets of the regular grid (selected, unlike ours, by taking concentric spheres). In the plane, our problem is dual to a colorful variant of Erdős's question. An intermediate between Erdős's problem and the one treated here would ask for the existence of a set of lines $\mathcal{L}$ in which each line is involved in many (colorless) k-incidences but there are no (colorless) $k+1$-incidences. Since the Solymosi-Stojakovíc construction provides $n^{2-\frac{c}{\sqrt{n}}}$ aligned k tuples of points, it is not hard to see, using a greedy deletion argument, that this alternative problem is essentially equivalent to Erdős's original one.

In higher dimensions, the question of finding sets of lines with many k-rich points (in the terminology of [10]) is interesting even without the condition of having no $(k+1)$rich point. Much of the recent research around this question has followed the solution to the joint problem [11] and has been driven by algebraic considerations (see [10] and the references therein). Here, we also ask for many k-rich points, but our questions are driven by combinatorial considerations. Our assumptions trade the usual density requirements (we assume linearly many, rather than polynomially many, k-rich points) for structural hypotheses in the form of conditions on the colors. We can still use some of the algebraic methods; the proof of Proposition 7 is, for instance, modeled on the upper bound on the number of joints of Guth and Katz [11].

[4] This case is closely connected with the famous *orchard problem* recently solved in its asymptotic version [8]

2 Probabilistic construction

In this section we prove:

▶ **Theorem 5.** *For any $k \geq 3$, $k+1 \geq d \geq 2$, and arbitrarily large $N \in \mathbb{N}$, there exists a finite set of lines in $\mathbb{R}^d$ of $k+1$ colors that is k-consistent, has no $(k+1)$-incidence, and in which each color class consists of between N and $3N$ lines, all concurrent.*

We describe our construction in $\mathbb{R}^{k+1}$ with color classes consisting of parallel lines. We then apply an adequate projective transform (to turn parallelism into concurrence) and a generic projection to a d-dimensional space; both transformations preserve incidences and therefore the properties of the construction.

Construction. Consider the finite subset $[n]^{k+1} = \{1, 2, ..., n\}^{k+1} \subset \mathbb{R}^{k+1}$ of the integer grid. We make our construction in two stages:

- Consider the set $\mathcal{L}_i^\#$ of n^k lines that are parallel to the ith coordinate axis and contain at least one point of our grid. We pick a random subset $\mathcal{L}'_i$, where each line from $\mathcal{L}_i^\#$ is chosen to be in $\mathcal{L}'_i$ independently with probability $p \stackrel{\text{def}}{=} 2n^{-\frac{2}{2k-1}}$ (the value of p is chosen with foresight).
- We then delete from $\mathcal{L}'_i$ all lines that are concurrent with k other lines from $\cup_{j\neq i}\mathcal{L}'_j$ and denote the resulting set $\mathcal{L}_i$.

We let $\mathcal{L}$ denote the colored set of lines $\mathcal{L} = \mathcal{L}_1 \cup \mathcal{L}_2 \cup \ldots \cup \mathcal{L}_{k+1}$. The second stage of the construction ensures that $\mathcal{L}$ has no $(k+1)$-incidence.[5] To prove Theorem 5, it thus suffices to show that with positive probability, $\mathcal{L}$ is k-consistent and each $\mathcal{L}_i$ has the announced size. Let us clarify that all lines considered in the proof are in $\cup_{i=1}^{k+1}\mathcal{L}_i^\#$ unless stated otherwise.

Consistency. Let us argue that $\mathcal{L}$ is k-consistent with high probability. For a set $I \subset [k+1]$, let

$$S_I \stackrel{\text{def}}{=} \{Q \in [n]^{k+1} : \forall i \in I \text{ there is a line of } \mathcal{L}_i \text{ containing } Q\},$$
$$S'_I \stackrel{\text{def}}{=} \{Q \in [n]^{k+1} : \forall i \in I \text{ there is a line of } \mathcal{L}'_i \text{ containing } Q\}.$$

We say that $\ell \in \mathcal{L}_i^\#$ is *j-bad* (for $j \neq i$) if ℓ contains no point of $S_{[k+1]\setminus\{i,j\}}$. Note that $\mathcal{L}$ is not k-consistent precisely when some $\ell \in \mathcal{L}_i^\#$ is j-bad and ℓ ends up in $\mathcal{L}_i$.

Let $\ell \in \mathcal{L}_i^\#$ and let $L \subset \mathcal{L}_i^\#$ be any set containing ℓ. Let $j \neq i$. We shall estimate $\mathrm{P}\left[(\ell \in \mathcal{L}_i) \wedge (\ell \text{ is } j\text{-bad}) \mid \mathcal{L}'_i = L\right]$. For ease of notation, we may assume that $i = k+1$, $j = k$ and ℓ is the line $\{(1, 1, \ldots, 1, x) : x \in \mathbb{R}\}$. Call a point $Q \in [n]^{k+1}$ *regular* if $Q \notin \ell$.

The randomness in the construction comes from $(k+1)n^k$ random choices, one for each line in $\mathcal{L}_1^\# \cup \cdots \cup \mathcal{L}_{k+1}^\#$. We refer to these random choices as 'coin flips' since we can think of each as being a result of a toss of a (biased) coin.

Let $\ell_{r,x}$ denote the line $\{(1, 1, \ldots, 1, y, 1, \ldots, 1, x) : y \in \mathbb{R}\}$, where y is at position r. If a line $\ell' \notin \mathcal{L}_r^\#$ intersects $\ell_{r,x}$ in point $(1, 1, \ldots, 1, y, 1, \ldots, 1, x)$, then all points of ℓ' have y in the rth position. Note that a point $(1, 1, \ldots, 1, y, 1, \ldots, 1, x)$ is regular if $y \neq 1$. A crucial observation is that if a line $\ell' \notin \mathcal{L}_{k+1}^\#$ intersects $\ell_{r,x}$ in a regular point and a line $\ell'' \notin \mathcal{L}_{k+1}^\#$ intersects $\ell_{r',x'}$ in a regular point and $(r, x) \neq (r', x')$, then ℓ' is different from ℓ''. This implies that sets of coin flips on which the events of the form

[5] Deleting one line per concurrence of size $k+1$ would suffice, but deleting all lines as we do simplifies the analysis and suffices for our purpose.

$$\text{“there is a regular } Q \in \ell_{r,x}\ Q \in S'_{[k]\setminus\{r\}}\text{”}$$

are disjoint for distinct (r,x), apart from the flips associated to the lines in $\mathcal{L}^{\#}_{k+1}$.

For a point $Q \in [n]^{k+1}$, let $\lambda(Q)$ be the line in $\mathcal{L}^{\#}_{k+1}$ containing Q. Hence,

$$\begin{aligned}
&\mathrm{P}\left[(\ell \in \mathcal{L}'_{k+1}) \wedge (\ell \text{ is } k\text{-bad}) \mid \mathcal{L}'_{k+1} = L\right] \\
&= \mathrm{P}\left[(\ell \in \mathcal{L}'_{k+1}) \wedge \bigwedge_{x\in[n]} (1,1,\ldots,1,x) \notin S_{[k-1]} \mid \mathcal{L}'_{k+1} = L\right] \\
&= \mathrm{P}\left[(\ell \in \mathcal{L}'_{k+1}) \wedge \bigwedge_{x\in[n]} (\exists r \in [k-1]\ \ell_{r,x} \notin \mathcal{L}_r) \mid \mathcal{L}'_{k+1} = L\right] \\
&= \mathrm{P}\left[(\ell \in \mathcal{L}'_{k+1}) \wedge \bigwedge_{x\in[n]} \Big(\exists r \in [k-1]\ \ell_{r,x} \notin \mathcal{L}'_r \vee (\exists Q \in \ell_{r,x} \cap S'_{[k+1]})\Big) \mid \mathcal{L}'_{k+1} = L\right]
\end{aligned}$$

In this last formula, the point Q can be assumed to be regular because $\ell \in L$, by assumption. Now we may drop $\ell \in \mathcal{L}'_{k+1}$ to obtain that the above is

$$\le \mathrm{P}\left[\bigwedge_{x\in[n]} \Big(\exists r \in [k-1]\ \ell_{r,x} \notin \mathcal{L}'_r \vee (\exists\,\text{reg. } Q \in \ell_{r,x} \cap S'_{[k+1]})\Big) \mid \mathcal{L}'_{k+1} = L\right]$$

Observe that if $\ell_{r,x} \in \mathcal{L}'_r$ then $Q \in \ell_{r,x} \cap S'_{[k+1]}$ holds if and only if $Q \in \ell_{r,x} \cap S'_{[k]\setminus\{r\}}$ and $\lambda(Q) \in L$. By the observation above, the set of coin flips on which these latter events depend for different x are disjoint, so this probability is

$$\begin{aligned}
&= \prod_{x\in[n]} \mathrm{P}\left[\exists r \in [k-1]\ \ell_{r,x} \notin \mathcal{L}'_r \vee (\exists\,\text{reg. } Q \in \ell_{r,x} \cap S'_{[k]\setminus\{r\}} \wedge \lambda(Q) \in L) \mid \mathcal{L}'_{k+1} = L\right] \\
&= \prod_{x\in[n]} \Bigg(1 - \mathrm{P}\Bigg[\forall r \in [k-1]\ \ell_{r,x} \in \mathcal{L}'_r \\
&\qquad\qquad \wedge (\forall\,\text{reg. } Q \in \ell_{r,x} Q \notin S'_{[k]\setminus\{r\}} \vee \lambda(Q) \notin L) \mid \mathcal{L}'_{k+1} = L\Bigg]\Bigg) \\
&= \prod_{x\in[n]} \Bigg(1 - \prod_{r\in[k-1]} \Bigg(p \cdot \prod_{\substack{\text{regular } Q\in\ell_{r,x} \\ \lambda(Q)\in L}} \mathrm{P}\left[Q \notin S'_{[k]\setminus\{r\}}\right]\Bigg)\Bigg)
\end{aligned}$$

Call $L \subset \mathcal{L}^{\#}_{k+1}$ *unbiased* if for every pair $(r,x) \in [k-1] \times [n]$ the number of points $Q \in \ell_{r,x}$ such that $\lambda(Q) \in L$ is at most $2pn$. For unbiased L, we obtain that the above is at most

$$\left(1 - \left(p \cdot (1-p^{k-1})^{2pn}\right)^{k-1}\right)^n \le \left(1 - \left(\tfrac{1}{2}p\right)^{k-1}\right)^n \le e^{-n\left(\frac{1}{2}p\right)^{k-1}} = e^{-n^{\frac{1}{2k-1}}}$$

If we pick L uniformly at random, then, for every $(r,x) \in [k-1] \times [n]$, the number of points $Q \in \ell_{r,x}$ such that $\lambda(Q) \in L$ is a binomial random variable. Chernoff's bound then yields

$$\begin{aligned}
&\mathrm{P}\left[(\ell \in \mathcal{L}'_{k+1}) \wedge (\ell \text{ is } k\text{-bad})\right] \\
&\leq \mathrm{P}\left[\mathcal{L}_{k+1} \text{ is biased}\right] + \sum_{\text{unbiased } L} \mathrm{P}\left[\mathcal{L}'_{k+1} = L\right] \mathrm{P}\left[(\ell \in \mathcal{L}'_{k+1}) \wedge (\ell \text{ is } k\text{-bad}) \mid \mathcal{L}'_{k+1} = L\right] \\
&\leq \sum_{(r,x)\in[k-1]\times[n]} e^{-(pn)^2/2n} + e^{-n^{\frac{1}{2k-1}}} = e^{-cn^{\frac{1}{2k-1}}}.
\end{aligned}$$

By taking the union bound over all i, j and ℓ we obtain that

$$\begin{aligned}
\mathrm{P}\left[\mathcal{L} \text{ is not } k\text{-consistent}\right] &\leq \mathrm{P}\left[\exists i, j\ \exists \ell \in \mathcal{L}_i^{\#}\ \left((\ell \in \mathcal{L}'_i) \wedge (\ell \text{ is } j\text{-bad})\right)\right] \\
&\leq (k+1)^2 n^k e^{-cn^{\frac{1}{2k-1}}} \leq e^{-c'n^{\frac{1}{2k-1}}}.
\end{aligned}$$

Size. We now analyze the probability that $\mathcal{L}_1$ is large (the bound will hold for each $\mathcal{L}_i$). Let us write $\mathcal{L}' = \cup_{i=1}^{k+1} \mathcal{L}'_i$ and label $\ell_1, \ell_2, \ldots, \ell_{n^k}$ the lines parallel to the 1st coordinate axis that intersect our grid. Put $X_i = \mathbb{1}_{\ell_i \in \mathcal{L}_1}$ and let $X = |\mathcal{L}_1| = X_1 + X_2 + \ldots + X_{n^k}$. We have

$$\begin{aligned}
&\mathrm{E}\left[X_i\right] = \mathrm{P}\left[X_i = 1\right] = \mathrm{P}\left[\ell_i \in \mathcal{L}'\right] \mathrm{P}\left[\ell_i \in \mathcal{L} \mid \ell_i \in \mathcal{L}'\right] = p(1-p^k)^n, \text{ and} \\
&\mathrm{E}\left[X\right] = n^k p(1-p^k)^n = \left(1 - n^{-\frac{2k}{2k-1}}\right)^n n^{k-\frac{2}{2k-1}} \geq \left(1 - \frac{1}{n}\right)^n n^{k-\frac{2}{2k-1}} \geq \frac{1}{4} n^{k-\frac{2}{2k-1}}.
\end{aligned}$$

Thus $\mathrm{E}\left[X\right] = N \in [\frac{1}{4} n^{k-\frac{2}{2k-1}}, n^{k-\frac{2}{2k-1}}]$. We next use a concentration inequality to pass from $\mathrm{E}\left[X\right]$ to an estimate on the probability that X is large. The second step introduces some dependency between some of the variables X_i, so we use Chebychev's inequality:

$$\mathrm{P}\left[|X - \mathrm{E}\left[X\right]| > \frac{1}{2}\mathrm{E}\left[X\right]\right] \leq 4 \frac{\mathrm{Var}\left[X\right]}{\mathrm{E}\left[X\right]^2} \leq 64 \mathrm{Var}\left[X\right] n^{-\left(2k - \frac{4}{2k-1}\right)}.$$

Recall that

$$\mathrm{Var}\left[X\right] = \sum_{i=1}^{n^k} \mathrm{Var}\left[X_i\right] + \sum_{1 \leq i < j \leq n^k} \mathrm{Cov}\left[X_i, X_j\right].$$

Since X_i takes values in $\{0, 1\}$, the first sum in the right-hand term is bounded by n^k. Moreover, there are $O(n^{k+1})$ pairs of variables X_i and X_j with non-zero covariance, since this requires the two lines ℓ_i and ℓ_j to belong to a common axis-aligned 2-plane. Again, each non-zero covariance is at most 1. Altogether, $\mathrm{Var}\left[X\right] = O(n^{k+1})$, so

$$\mathrm{P}\left[|X - \mathrm{E}\left[X\right]| > \frac{1}{2}\mathrm{E}\left[X\right]\right] = O\left(n^{\frac{2k+3}{2k-1} - k}\right).$$

For $k \geq 3$, the probability that X is in $\left[\frac{1}{8} n^{k-\frac{2}{2k-1}}, \frac{3}{2} n^{k-\frac{2}{2k-1}}\right]$ goes to 1 as n goes to infinity.

3 Algebraic construction

In this section we prove:

▶ **Theorem 6.** *For any $k \geq 3$, $k+1 \geq d \geq 2$, and arbitrarily large N, there exists a finite set of lines in $\mathbb{R}^d$ with $k+1$ colors that is k-consistent, has no $(k+1)$-incidence, and in which each color class consists of N lines, all concurrent.*

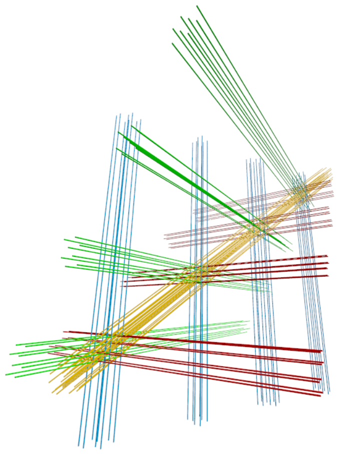

Figure 3 A projection to $\mathbb{R}^3$ of our construction for $k = 3$ and $p = 2$ (reprojected to the plane).

As in Section 2, we describe our construction in $\mathbb{R}^{k+1}$ with parallel families of lines, and obtain the desired configuration by an adequate projective transformation and a projection. We again consider the finite portion of the integer grid $[n]^{k+1} \subset \mathbb{R}^{k+1}$ and the axis-aligned lines that intersects it. Unlike in Section 2, we give an explicit way to select some of these lines to achieve the desired configuration.

Construction. We work with axis-aligned lines that intersect in points of our grid. Hence, identifying each line with the subset of the grid that it contains does not affect incidences. We fix a prime number p and parameterize $[n]$ by the vector space $\mathbb{V} = (\mathbb{Z}/p\mathbb{Z})^{k-1}$; this restricts the choice of n to certain prime powers, but still allows to make it arbitrarily large. We use this parameterization to describe the lines in our configuration as solutions of well-chosen linear equations.

Let $v_1, v_2, \ldots, v_k \in \mathbb{V}$ such that $v_1 + v_2 + \ldots + v_k = 0$ and any proper subset of them are linearly independent. Let $\cdot$ denote the inner product of the vector space $\mathbb{V}$. For $i = 1 \ldots k$, our set $\mathcal{L}_i$ consist of all the lines parallel to the ith coordinates and passing through a point with parameters $(X_1, \ldots, X_{k+1}) \in \mathbb{V}^{k+1}$ such that

$$v_{i-1} \cdot X_1 + v_{i-1} \cdot X_2 + \ldots + v_{i-1} \cdot X_{i-1} + v_i \cdot X_{i+1} + \ldots + v_i \cdot X_{k+1} = 0. \quad (1)$$

(Keep in mind that each X_i is a vector in $(\mathbb{Z}/p\mathbb{Z})^{k-1}$.) We define $\mathcal{L}_{k+1}$ similarly but replace Equation (1) by

$$v_k \cdot X_1 + v_k \cdot X_2 + \ldots + v_k \cdot X_k = 1. \quad (2)$$

No $(k+1)$-incidence. Any $(k+1)$-incidence is a point of the grid whose parameters $(X_1, \ldots, X_{k+1})$ satisfy the system:

$$\begin{cases} & v_1 \cdot X_2 + v_1 \cdot X_3 + \ldots + v_1 \cdot X_k + v_1 \cdot X_{k+1} = 0 \\ v_1 \cdot X_1 & + v_2 \cdot X_3 + \ldots + v_2 \cdot X_k + v_2 \cdot X_{k+1} = 0 \\ v_2 \cdot X_1 + v_2 \cdot X_2 & + \ldots + v_3 \cdot X_k + v_3 \cdot X_{k+1} = 0 \\ v_3 \cdot X_1 + v_3 \cdot X_2 + v_3 \cdot X_3 & + \ldots + v_4 \cdot X_k + v_4 \cdot X_{k+1} = 0 \\ \ldots & \\ v_{k-1} \cdot X_1 + v_{k-1} \cdot X_2 + v_{k-1} \cdot X_3 + \ldots & + v_k \cdot X_{k+1} = 0 \\ v_k \cdot X_1 + v_k \cdot X_2 + v_k \cdot X_3 + \ldots + v_k \cdot X_k & = 1 \end{cases}$$

Summing all these conditions yields

$$\left(\sum_{i=1}^{k} v_i\right) \cdot \left(\sum_{i=1}^{k+1} X_i\right) = 1,$$

which contradicts $v_1 + v_2 + \ldots + v_k = 0$. So there is no $(k+1)$-incidence.

k-consistency. Fix a line $\ell \in \mathcal{L}_1$. It corresponds to some solution $(X_2^*, \ldots, X_{k+1}^*)$ of Equation (1). The grid points on ℓ are precisely the points of the form $(X_1, X_2^*, X_3^*, \ldots, X_{k+1}^*)$ and are parameterized by X_1. Each equation in the system above reduces to $v_j \cdot X_1 = c_j$, where c_j is some constant vector (computed from the v_j's and the X_j^*'s). Since $X_1 \in (\mathbb{Z}/p\mathbb{Z})^{k-1}$ and any $k-1$ of the v_j are linearly independent, any choice of $k-1$ equations has a solution. This means that for any i, the line ℓ is concurrent with lines from all $\mathcal{L}_j$ with $j \in [k+1] \setminus \{i\}$. The same goes with the lines of $\mathcal{L}_2, \ldots, \mathcal{L}_{k+1}$ so the configuration is consistent.

Size. In this construction, the size of $\mathcal{L}_i$ is the number of $(X_1, \ldots, X_{i-1}, X_{i+1}, \ldots, X_{k+1})$ in $\mathbb{V}^k$ satisfying Equation (1) – or (2) if $i = k+1$. Hence $|\mathcal{L}_i| = p^{k^2-k-1}$ for every i. The smallest configuration built in this way thus has 2^{k^2-k-1} lines per set (which is 32 for $k = 3$); refer to Figure 3.

4 More on grid-like examples

Both Theorems 5 and 6 construct examples as projections of subsets of a regular grid in higher dimension. We discuss here the properties of such constructions.

Number of lines. Consider a colored set of lines $\mathcal{L} = \mathcal{L}_1 \cup \mathcal{L}_2 \cup \ldots \cup \mathcal{L}_{k+1}$ in $\mathbb{R}^d$. We say that a t-incidence of $\mathcal{L}$ is *flat* if the lines meeting there are contained in an affine subspace of dimension at most $\min(d, t) - 1$. In any grid-like construction such as those in Theorems 5 and 6, every k-incidence is non-flat.

▶ **Proposition 7.** *Let $\mathcal{L} = \mathcal{L}_1 \cup \mathcal{L}_2 \cup \ldots \cup \mathcal{L}_m$ be a k-consistent colored set of lines in $\mathbb{R}^k$ with no $(k+1)$-incidence. If no k-incidence of $\mathcal{L}$ is flat, then*

$$\sum_{i=1}^{m} |\mathcal{L}_i| \geq \frac{\binom{\binom{m-1}{k-1}+k-1}{k}}{\binom{m-1}{k-1}}.$$

The proof essentially follows the argument of Guth and Katz [11] for bounding the number of joint among n lines; the main difference is that the consistency assumption makes their initial pruning step unnecessary. We spell out the details in [4, Proposition 7]. For $m = k+1$, the bound of Proposition 7 is $\frac{1}{k}\binom{2k}{k}$, so the number of lines required grows exponentially with k.

Non-concurrent colors. Theorems 5 and 6 both use a grid in $\mathbb{R}^{k+1}$ to start with $k+1$ color classes, each of size n^k, where every line is involved in n colorful incidences. Recall that in this setup, every color class is concurrent (it consists of parallel lines). This is in fact important, perhaps essential. To see this, note that any two of our starting color classes contain arbitrarily large subsets whose intersection graph is dense. This is impossible, generically, if we try to work with families of lines that are secant to two skew lines in $\mathbb{R}^3$. [6]

[6] This choice is motivated by the design of *two-slit cameras* [17, 2].

▶ **Proposition 8.** *For $i = 1, 2$, let Γ_i denote the set of lines secant to two fixed lines s_i and s'_i in $\mathbb{R}^3$. Let A and B be two sets of n lines from Γ_1 and Γ_2, respectively. If the lines s_1, s'_1, s_2 and s'_2 are in generic position then the intersection graph of A and B has $O\left(n^{4/3}\right)$ edges.*

Proof. First, note that the intersection graph of A and B is semi-algebraic: parameterizing Γ_i by $s_i \times s'_i \simeq \mathbb{R}^2$ makes the incidence an algebraic relation, as can be deduced from the bilinearity of incidence in Plücker coordinates. Next, remark that if this graph contains a complete bipartite subgraph $K_{3,3}$, then the lines $\{s_1, s'_1, s_2, s'_2\}$ are in a special position. Indeed, in the generic case, these two triples of lines come from the two families of rulings of a quadric surface [18, §10]; the lines s_1, s'_1, s_2 and s'_2 are also rulings of that quadric, so *both* s_1 and s'_1 intersect *both* s_2 and s'_2. In the non-generic cases, the six lines must be coplanar with s_1 and s_2. Now, we apply the semi-algebraic version of the Kővári –Sós–Turán theorem [7], and obtain that the number of edges of our graph is $O\left(n^{4/3}\right)$. ◀

We see the previous result as an indication that a straightforward adaptation of our probabilistic construction to the case of two-slits is unlikely. Can the bound in Proposition 8 be improved from $O\left(n^{4/3}\right)$ to $O(n)$?

▶ Remark. Note that the genericity assumption in Proposition 8 is on the sets Γ_i, not on their subsets. The analogue for concurrent sets of lines would be to require that the centers of concurrence are in generic position; this clearly does not prevent finding arbitrarily large subsets with dense intersection graphs.

Extension to continuous sets of lines. The constructions of Theorems 5 and 6 can be turned into continuous families of lines as follows.

First, we follow either construction up to the point where we have a family $\mathcal{L}$ of lines of $k+1$ colors in $\mathbb{R}^{k+1}$ that is k-consistent, without colorful incidence, and where each color class is parallel. Consider a parameter $\epsilon > 0$, to be fixed later. For every i, we build a set $\mathcal{L}_i(\epsilon)$ by considering every line $\ell \in \mathcal{L}_i$ in turn, and adding to $\mathcal{L}_i$ every line ℓ' parallel to ℓ such that the distance between ℓ and ℓ' is most ϵ. Note that for $\epsilon < 1/2$ the family $\mathcal{L}(\epsilon)$ is k-consistent and without colorful incidence.

Now, consider a generic projection $f\colon \mathbb{R}^{k+1} \to \mathbb{R}^d$ for the desired d. For any $\epsilon > 0$ the family $\mathcal{L}(\epsilon)$ is k-consistent. We observe that for $\epsilon > 0$ small enough, it also remains without colorful incidence. Let τ denote the minimum distance, in the projection, between a k-incidence and a line (of any color) not involved in that incidence. Every k-incidence in $\mathcal{L}$ gives rise, in $\mathcal{L}(\epsilon)$, to k tubes that intersect in a bounded convex set B of size $O(\epsilon)$. Choosing $\epsilon > 0$ such that the diameter of $f(B)$ is less than $\tau/2$ ensures that the corresponding family $f(\mathcal{L}(\epsilon))$ has no colorful incidence.

For a given family of colored lines $\mathcal{L}$ define the set P_S to be the set of points incident to at least one line of each of the color classes in S; see Figure 1. Notice that in our examples, for each set S of k colors the set P_S is highly disconnected. As mentioned in the introduction, Trager et al. [16] showed that if a family of sets of lines is 3-consistent and for each S of size 3, the set P_S is convex, then the whole family is consistent. An interesting open question is whether an analogue theorem holds if instead of convexity, we assume that for every set S of size k, the set P_S is sufficiently connected.

5 Constructions with few lines

The configurations constructed in Sections 2 and 3 have at least 32 lines per color. This is considerably larger than the sets of lines involved in some of the questions around consistency that arise in computer vision. In the example of structure-from-motion mentioned in introduction, when the camera is central, every color class has only 5 to 7 lines. It turns out that for sufficiently small configurations, k-consistency does imply some colorful incidences:

▶ **Lemma 9.** *Any 3-consistent colored set of lines $\mathcal{L} = \mathcal{L}_1 \cup \mathcal{L}_2 \cup \mathcal{L}_3 \cup \mathcal{L}_4$ in $\mathbb{R}^d$ with $|\mathcal{L}_1| = |\mathcal{L}_2| = |\mathcal{L}_3| = |\mathcal{L}_4| = 2$ contains a colorful incidence.*

Proof. Let us prove the case where $d = 2$; the general case follows by projecting onto a generic 2-plane. Let P_i denote the dual of $\mathcal{L}_i$, and let $P = P_1 \cup P_2 \cup P_3 \cup P_4$. Assume, by contradiction, that $\mathcal{L}$ contains no colorful incidence, *i.e.* that no line intersects every P_i. Let $P' = P_2 \cup P_3 \cup P_4$ and let us apply a projective transform to map the points of P_1 to the horizontal and vertical directions, respectively. We call a line that contains a point of each of P_2, P_3 and P_4 a *rainbow line.*

Since $\mathcal{L}$ is 3-consistent, for any point $x \in P$ and any choice of 2 other colors, there is a line through x that contains a point of each of these colors. Since $\mathcal{L}$ has no colorful incidence, there must exist three horizontal lines and three vertical lines that intersect P', and each must contain exactly two points of P' of distinct colors. Moreover, no rainbow line can be horizontal or vertical. But this implies that out of the 9 intersections between horizontal and vertical lines, only 5 (the corners and the center) can be on a rainbow line. This contradicts $|P'| = 6$. ◀

We prove here a slightly stronger lower bound:

▶ **Theorem 10.** *Let $\mathcal{L} = \mathcal{L}_1 \cup \mathcal{L}_2 \cup \mathcal{L}_3 \cup \mathcal{L}_4$ be a 3-consistent colored set of lines in $\mathbb{R}^3$ with no colorful incidence and concurrent colors. If $|\mathcal{L}| < 24$, then $\mathcal{L}$ is contained in a 2-plane.*

Sketch of proof. We only outline our proof here and refer to [4] for the details. Assume, by contradiction, that $\mathcal{L}$ is a configuration with all required properties and $|\mathcal{L}| < 24$.

We first argue that $\mathcal{L}$ decomposes into a disjoint union of two colored sets of lines, each of which has 4 colors, is 3-consistent, has no colorful incidence, and has concurrent color classes. To do so, we consider the planes spanned by a line from the smallest color class, say $\mathcal{L}_1$, and the center of concurrence of another color class, say $\mathcal{L}_2$. The assumptions force every line of a color to intersect the lines of any other color in at least two points. This means that any such plane contains at least two lines of $\mathcal{L}_1$, so there are at most two planes. The same reason forces all lines from $\mathcal{L}_2$ to be contained in one plane or the other, and eventually the same goes for $\mathcal{L}_3$ and $\mathcal{L}_4$.

We then conclude by arguing that each of the subsets has at least 12 lines, forcing $\mathcal{L}$ to have at least 24 lines. This is straightforward if every color class has size at least 3. We argue that if a color class has size two, then all other color classes must have size at least 4. ◀

Classification. We also provide a characterization of 3-consistent, 4-colored sets of lines in $\mathbb{R}^3$ with no colorful incidence and 3 lines per color. Forgetting for a moment about colors, any such configuration must consist of 12 lines and 12 points, every point on 3 lines and every line through 3 points; in the classical tabulation of projective configurations, they are called (12_3) *configurations.* It turns out that there are 229 possible incidence structures meeting this description, and that every single one of them is realizable in $\mathbb{R}^3$ [9]. To analyze what happens when we add back the colors and the consistency assumption, we consider two special (12_3) configurations:

- A *Reye-type configuration*[7] is a configuration obtained by selecting 12 out of the 16 lines supporting the 12 edges and four long diagonals of a cube, in a way that produces a $(12)_3$ configuration.
- A *Desargues-type* configuration is defined from six planes $\Pi_1, \Pi_2, \ldots, \Pi_6$ in $\mathbb{R}^3$ where (i) each of $\{\Pi_1, \Pi_2, \ldots, \Pi_5\}$ and $\{\Pi_1, \Pi_2, \Pi_3, \Pi_6\}$ is in general position, and (ii) Π_4, Π_5 and Π_6 intersect in a line. The configuration consists of all lines that are contained in exactly two planes.

Here is our classification:

▶ **Theorem 11.** *Let $\mathcal{L} = \mathcal{L}_1 \cup \mathcal{L}_2 \cup \mathcal{L}_3 \cup \mathcal{L}_4$ be a 3-consistent colored set of lines in $\mathbb{R}^3$ with no colorful incidence. If every color class has size 3, and $\mathcal{L}$ is not contained in a 2-plane, then it is a Desargues-type or a Reye-type configuration colored as in Figure 2.*

Sketch of proof. We only outline our proof here, and refer to [4] for the details. The hypothesis imply that every line must intersect the lines of any other color in at least two points. This essentially allows us to establish that every color class consists of lines that are either pairwise skew, or concurrent and not coplanar. This geometric restriction implies, in turn, that for $i \neq j$, every line of $\mathcal{L}_i$ intersects exactly two lines of $\mathcal{L}_j$, and for any two lines of $\mathcal{L}_j$, there is exactly one line of $\mathcal{L}_i$ that intersects them both.

We use these two observations to reduce the sets of candidates for the incidence structure of the 12 lines. We fix a color class, say $A = \mathcal{L}_1$, and build a graph whose vertices are the 3-incidences involving a line of A, and where two vertices form an edge if the corresponding incidences have a line in common. This graph can be checked to be one of two candidates:

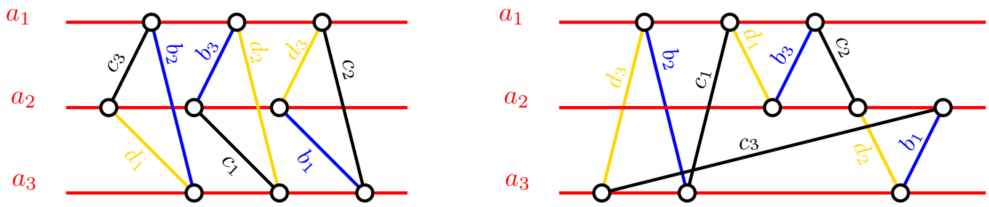

This already determines all 3-incidences that involve $\mathcal{L}_1$. The rest follows by observing that if the lines of $\mathcal{L}_1$ are pairwise skew (resp. concurrent and not coplanar) then two lines in $\mathcal{L}_2 \cup \mathcal{L}_3 \cup \mathcal{L}_4$ that intersect the same pair (resp. different pairs) of lines of $\mathcal{L}_1$ cannot intersect outside of $\mathcal{L}_1$. We end up with only two possible incidence structures (up to isomorphism, and possibly relabeling):

$$(I): \begin{matrix} a_1b_2c_3 & a_1b_3d_2 & a_1c_2d_3 & b_1c_3d_2 \\ a_2b_3c_1 & a_2b_1d_3 & a_2c_3d_1 & b_2c_1d_3 \\ a_3b_1c_2 & a_3b_2d_1 & a_3c_1d_2 & b_3c_2d_1 \end{matrix} \quad \text{or} \quad (II): \begin{matrix} a_1b_2c_3 & a_1b_3d_2 & a_1c_2d_3 & b_1c_1d_1 \\ a_2b_3c_1 & a_2b_1d_3 & a_2c_3d_1 & b_2c_2d_2 \\ a_3b_1c_2 & a_3b_2d_1 & a_3c_1d_2 & b_3c_3d_3. \end{matrix}$$

Figure 2 gives non-planar realizations of both set of incidences. We argue that these are essentially the only realizations by choosing a particular subset of points of incidence, and showing that their coordinates determines the whole geometric realization. This last step amounts, in each case, to an incidence theorem in projective geometry similar to the classic theorems of Reye or Desargues. ◀

[7] The $(12_4 16_3)$ configuration of Reye consists of 12 points and 16 lines in $\mathbb{R}^3$ such that every point is on 4 lines and every line contains 3 points; its realizations are projectively equivalent to the 16 lines supporting the 12 edges and four long diagonals of a cube, together with that cube's vertices and center and the 3 points at infinity in the directions of its edges

References

1 Kalle Åström, Roberto Cipolla, and Peter J Giblin. Generalised epipolar constraints. In *European Conference on Computer Vision*, pages 95–108. Springer, 1996.

2 Guillaume Batog, Xavier Goaoc, and Jean Ponce. Admissible linear map models of linear cameras. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 1578–1585. IEEE, 2010.

3 Edmond Boyer. On using silhouettes for camera calibration. *Computer Vision–ACCV 2006*, pages 1–10, 2006.

4 Boris Bukh, Xavier Goaoc, Alfredo Hubard, and Matthew Trager. Consistent sets of lines with no colorful incidence. *Preprint arXiv:1803.06267*, 2018.

5 Paul Erdős and George Purdy. Some extremal problems in geometry. *Discrete Math*, pages 305–315, 1974.

6 Olivier Faugeras and Bernard Mourrain. On the geometry and algebra of the point and line correspondences between n images. In *Computer Vision, 1995. Proceedings., Fifth International Conference on*, pages 951–956. IEEE, 1995.

7 Jacob Fox, János Pach, Adam Sheffer, Andrew Suk, and Joshua Zahl. A semi-algebraic version of Zarankiewicz's problem. *Journal of the European Mathematical Society*, 19:1785–1810, 2017.

8 Ben Green and Terence Tao. On sets defining few ordinary lines. *Discrete & Computational Geometry*, 50(2):409–468, 2013.

9 Harald Gropp. Configurations and their realization. *Discrete Mathematics*, 174(1-3):137–151, 1997.

10 Larry Guth. Ruled surface theory and incidence geometry. In *A Journey Through Discrete Mathematics*, pages 449–466. Springer, 2017.

11 Larry Guth and Nets Hawk Katz. Algebraic methods in discrete analogs of the Kakeya problem. *Advances in Mathematics*, 225(5):2828–2839, 2010.

12 Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.

13 Carlos Hernández, Francis Schmitt, and Roberto Cipolla. Silhouette coherence for camera calibration under circular motion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(2), 2007.

14 Onur Özyeşil, Vladislav Voroninski, Ronen Basri, and Amit Singer. A survey of structure from motion. *Acta Numerica*, 26:305–364, 2017.

15 József Solymosi and Miloš Stojaković. Many collinear k-tuples with no $k+1$ collinear points. *Discrete & Computational Geometry*, 50(3):811–820, 2013.

16 Matthew Trager, Martial Hebert, and Jean Ponce. Consistency of silhouettes and their duals. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3346–3354, 2016.

17 Matthew Trager, Bernd Sturmfels, John Canny, Martial Hebert, and Jean Ponce. General models for rational cameras and the case of two-slit projections. In *CVPR 2017 - IEEE Conference on Computer Vision and Pattern Recognition*, Honolulu, United States, 2017. URL: `https://hal.archives-ouvertes.fr/hal-01506996`.

18 Oswald Veblen and John Wesley Young. *Projective geometry*, volume 1. Ginn, 1918.