

Subspace arrangements, graph rigidity and derandomization through submodular optimization*

Orit E. Raz[†]

Avi Wigderson[‡]

*Dedicated with admiration to László Lovász,
on the occasion of his 70th birthday.*

Abstract

This paper presents a deterministic, strongly polynomial time algorithm for computing the matrix rank for a class of symbolic matrices (whose entries are polynomials over a field). This class was introduced, in a different language, by Lovász [19] in his study of flats in matroids, and proved a duality theorem putting this problem in $NP \cap coNP$. As such, our result is another demonstration where “good characterization” in the sense of Edmonds leads to an efficient algorithm. In a different paper Lovász [16] proved that all such symbolic rank problems have efficient probabilistic algorithms, namely are in BPP . As such, our algorithm may be interpreted as a derandomization result, in the long sequence special cases of the PIT (Polynomial Identity Testing) problem. Finally, Lovász and Yemini [20] showed how the same problem generalizes the *graph rigidity* problem in two dimensions. As such, our algorithm may be seen as a generalization of the well-known deterministic algorithm for the latter problem.

There are two somewhat unusual technical features in this paper. The first is the translation of Lovász’ flats problem into a symbolic rank one. The second is the use of submodular optimization for derandomization. We hope that the tools developed for both will be useful for related problems, in particular for better understanding of graph rigidity in higher dimensions.

1 Introduction

In this paper we provide a new *deterministic*, strongly polynomial time algorithm which can be viewed in two ways. The first is as solving a derandomization problem, providing a deterministic algorithm to a new special case of the PIT (Polynomial Identity Testing) problem. The second is as computing the dimension of the span a collection of subspaces in high dimensional space. Motivating and connecting the two is the problem of testing *graph rigidity*, to which an efficient deterministic algorithm is known only in the plane, and is open for higher dimensions. Accordingly, we will divide the introduction to explain these three problems.

*The first author was partially supported from NSF grant DMS-1128155. The second author was partially supported from NSF grant CCF-1412958

[†]Department of Mathematics, University of British Columbia, Vancouver, Canada. oritraz@math.ubc.ca

[‡]School of Mathematics, Institute for Advanced Study, Princeton NJ 08540, U.S.A. avi@ias.edu

1.1 Polynomial Identity Testing (PIT)

Let \mathbb{K} be a field. Let $\mathbf{x} = (x_1, \dots, x_d)$ be a d -tuple of independent variables. The PIT problem is to determine, given a multivariate polynomial $p \in \mathbb{K}[\mathbf{x}]$, if $p \equiv 0$ (as a polynomial). Of course, the description of p as an input to this problem is central to its complexity, and many variants of this problem were considered. The most common formulation is when p is given by an arithmetic formula or circuit¹.

The original version of this question was posed by Edmonds [5]. In his formulation, p is the determinant of a matrix whose entries are linear forms in \mathbf{x} (we will refer such a matrix as a *symbolic* matrix). Lovász [16] proved that this problem is in *BPP* namely has a fast probabilistic algorithm (for fields \mathbb{K} larger than the degree of p): indeed, the algorithm simply picks random elements from \mathbb{K} and evaluates p (note that evaluating p is efficient in all three formulations above, and indeed in all formulations considered). This left open the problem of finding an efficient deterministic algorithm, namely derandomizing Lovász's algorithm for PIT.

Open Problem 1.1. *Is PIT $\in P$?*

The importance of this seemingly specific open problem was revealed in an important result of Kabanets and Impagliazzo [13]. They showed that if the answer is positive (as everyone expects), this will imply non-trivial lower bounds on either arithmetic or Boolean circuits, well beyond current techniques.

The progress towards resolving this open problem has been by providing deterministic polynomial time algorithms for a large variety of special cases of it, with the idea of building up techniques. By far, in most of these results the special cases are defined by restricting the input polynomial to lie in some complexity class. In these cases, progress in derandomization followed closely progress on lower bounds for the appropriate class (as is the case in the Boolean setting as well). There are literally dozens of such papers: many are mentioned and explained in the surveys [22, 24] and e.g. the recent paper [1].

In parallel, with motivation from algebra, geometry and other areas, a different collection of special cases of PIT was studied, of a structural nature. Here one works with Edmond's formulation, and develops an understanding (and often a polynomial time algorithm) for cases where the symbolic matrix has restricted structure. This includes for example the works [3, 4, 6, 9, 11, 21].

This paper contributes to the second line of research, providing new families of symbolic matrices for which PIT can be solved in deterministic polynomial time. To explain this structure we introduce some notation. We will work in a slightly more general setting, in two ways, as the results generalize to both. First, we will allow our symbolic matrices to have polynomial entries. In such cases, these polynomials will have simple formulas describing them. Second, we will be interested in computing the *rank* of the input symbolic matrix, not just whether its determinant vanishes. While seemingly a more general problem, this turns out to be equivalent to PIT (see e.g. [8, Appendix A]²).

¹When the input is a circuit, the degree of p is always assumed to be polynomial in the circuit's size, and in all cases considered in this paper this will be evident.

²The proof in [8] is given for *non-commutative* rank, but the exact same proof works verbatim for our usual notion of rank over $\mathbb{K}(\mathbf{x})$.

Let R be a family of polynomial maps $R = \{r : \mathbb{K}^d \rightarrow \mathbb{K}^n\}$. In all cases we assume the degree of all polynomials in all maps is at most n , and the number of variables d is at most polynomial in n , so we will think of n as the input size to the problem.

A family of maps R prescribes a family of symbolic matrices, so that each row is an image of the d -vector of variables \mathbf{x} under some map in R . More formally, define $\text{PIT}(R)$ to be the set of all symbolic matrices M (with n columns, and $\text{poly}(n)$ rows) in which every row of the matrix is of the form $r(\mathbf{x})$, for some map $r \in R$. We will be interested in families R for which the ranks of matrices in $\text{PIT}(R)$ can be computed in polynomial time³.

We first demonstrate the convenience of this notation. Call R *complete*, if a deterministic polynomial-time algorithm for $\text{PIT}(R)$ implies a deterministic polynomial-time algorithm for PIT. Very simple maps are complete! It follows from Valiant’s [28] hardness of the determinant for the class⁴ VP that

Theorem 1.2 ([28]). *The class R_{affine} of affine linear maps is complete.*

Indeed, Valiant’s original proof (see more detail here [15]) implies a stronger theorem. Even restricting the support of each row to have at most a single variable in some coordinate, is general enough to be complete.

Theorem 1.3. *The class R_{sparse} of affine linear maps, such that each map is non-constant in at most a single variable from $\{x_1, \dots, x_d\}$, is complete.*

We now turn to define the polynomial maps we will be interested in, and for which we will be able to provide efficient deterministic algorithms. Some motivation for interest in these maps will be given in the next two subsections.

Consider the following class R_2 . Here $d = n$. Every $p \in R_2$ is of the form $\mathbf{x} \mapsto (A - A^T)\mathbf{x}$, where A is a rank-1 matrix. While this family may look very special, we note that the problem of graph rigidity in \mathbb{R}^2 (for which a polynomial time algorithm is known but far from trivial) is a very special case of $\text{PIT}(R_2)$.⁵

Theorem 1.4. *$\text{PIT}(R_2)$ can be solved in deterministic polynomial time, over a field \mathbb{K} with sufficiently large characteristic (more precisely, when $\text{char}(\mathbb{K})$ is larger than the number of rows of the input matrix or $\text{char}(\mathbb{K}) = 0$).*

This construction can be generalized as follows. Here we will generate PIT instances whose entries are *polynomials*, rather than linear functions of the variables. For a k -dimensional tensor A of size n , denote by \hat{A} its “anti-symmetric” version, namely where for every entry (i_1, \dots, i_k) we have $\hat{A}(i_1, \dots, i_k) = \sum_{\sigma \in S_k} \text{sgn}(\sigma) A(i_{\sigma(1)}, \dots, i_{\sigma(k)})$. Note that for $k = 2$ we have $\hat{A} = A - A^T$.

We now extend R_2 , in which a matrix (namely a 2-dimensional tensor) acts on one vector of variables, to R_k , in which a k -dimensional tensor acts on $k - 1$ vectors of variables. Let

³We identify the set of matrices and the computational problem of determining their ranks.

⁴The arithmetic analog of the Boolean class P .

⁵Moreover, the same family of rank-2, skew symmetric matrices is featured in a very different PIT problem: determining the maximum rank of a subspace generated by given such matrices. A deterministic polynomial time solution for this problem is given by Lovasz’ celebrated matroid parity algorithm [17] (see also [18], Theorem 11.1.2).

R_k denote the following class of (degree $k - 1$) maps. Let $\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^{k-1}$ be n -vectors of independent variables, so altogether $\mathbf{x} = (\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^{k-1})$ is a vector of $(k - 1)n$ variables. A k -tensor of size n in each dimension acts on \mathbf{x} simply with the i 'th dimension acting on \mathbf{x}^i for $i \in [k - 1]$. The output of this action is a vector (along dimension k) of length n of polynomials of degree $k - 1$, each linear in \mathbf{x}^i for all i . Define R_k to be all maps defined by \hat{A} for any rank-1 tensor A . Note that with this notation R_2 is precisely the class defined above.

Generalizing the above theorem we prove:

Theorem 1.5. *For every $k < n$, $PIT(R_k)$ can be solved in deterministic polynomial time, over a field \mathbb{K} with sufficiently large characteristic (more precisely, when $\text{char}(\mathbb{K})$ is larger than the number of rows of the input matrix or $\text{char}(\mathbb{K}) = 0$).*

1.2 Graph Rigidity

The problem of graph rigidity arises from several motivations, originally, mechanical engineering (see [14]). Rigidity theory is a fast-growing area, and we refer the interested reader to [25] for more background and recent approaches. Graph rigidity has several versions, we describe perhaps the most common one, *generic* rigidity. It is supposed to capture the structural rigidity of a “bars and joints” framework described by a graph. We will not be formal here as precise definitions can be found e.g. in [2]. Here the relevant field for the geometric/physical interpretation is the Real numbers \mathbb{R} , and we use it in this subsection as in other papers on this problem (although the algebraic formulation is meaningful for every field \mathbb{K}).

Let $G(V, E)$ be an undirected graph on n vertices and m edges. An *embedding* of G in \mathbb{R}^t is a map $\phi : V \rightarrow \mathbb{R}^t$. An embedding of G is called *rigid* if there is no perturbation of the vertex positions which preserves all edge lengths, other than the rigid motions of \mathbb{R}^t . The graph G is called *rigid* if every *generic* embedding of G is rigid (equivalently, if there exists an embedding of G which is rigid, see [2]). The main question is to determine if a given graph G is rigid (and more generally, compute the dimension of the non-rigid motions of a generic embedding, in case G is not rigid).

An extremely convenient formulation of the problem (as a PIT) is the following. Let $x_{v,j}$ be a set of variables indexed by $v \in V$ and $j \in [t]$. The intuition is that $(x_{v,1}, \dots, x_{v,t})$ are the coordinates of a generic embedding of the vertex v in \mathbb{R}^t . Given G , construct a symbolic matrix $M_{G,t}$ of dimensions $m \times nt$, which may be viewed as a concatenation of t matrices, one for each dimension $j \in [t]$. Every row corresponds to an edge $\{u, v\} \in E$, and for each j , the column u, j contains the entry $x_{u,j} - x_{v,j}$, whereas the column v, j contains the the negation $x_{v,j} - x_{u,j}$.

It is not hard to prove that the rank (as usual, over $\mathbb{R}(x)$) of $M_{G,t}$ determines if G is rigid, and indeed the dimension of non-rigid motions (see [2] for the details). It is easy to see that for every graph G , the matrix $M_{G,2}$ is in the class $PIT(R_2)$ above. Indeed, let e_1, \dots, e_{2n} denote the standard basis vectors in \mathbb{R}^{2n} . For some $u < v \in [n]$, put $a = e_u - e_v$ and $b = e_{n+u} - e_{n+v}$. Consider the matrix $A = A_{u,v} := a^t b$. Then $(A - A^t)\mathbf{x}$, where $\mathbf{x} = (x_{21}, \dots, x_{2n}, x_{11}, \dots, x_{1n})$ is the $\{u, v\}$ row of $M_{G,2}$. Thus Theorem 1.4 yields as a corollary a polynomial time algorithm to determine whether a given graph G is rigid in \mathbb{R}^2 . Such algorithms for rigidity in \mathbb{R}^2 are known (see [10, Section 2.2] and references therein).

Note that the matrices $M_{G,t}$ make sense over any field \mathbb{K} , instead of \mathbb{R} , and Theorem 1.4 in fact provides a deterministic polynomial time algorithm to compute the rank of these matrices over any field \mathbb{K} with large enough characteristic.

The symbolic matrix representation above shows that for every t , the problem of testing graph rigidity in \mathbb{R}^t is in *BPP*, and it is a decades-old problem to whether it is also in *P*, even for the case $t = 3$.

Lovász and Yemini [20] have developed an alternative approach for studying graph rigidity in the plane, which obtains a somewhat finer characterization of rigidity than Laman's. What is even more interesting is their method. They show that the matrices $M_{G,2}$ can actually be obtained in the following way. First, with every edge $\{u, v\}$ associate a certain 2-dimensional subspace $f_{u,v} \subset \mathbb{R}^{2n}$. The intersection of this subspace $f_{u,v}$ with a *generic* hyperplane through the origin (of which the normal can be viewed essentially as the $2n$ -vector of variables $x_{v,j}$) yields the $\{u, v\}$ row of $M_{G,2}$. In more detail, identify the vertices of G with the set $V = [n]$, and let e_1, \dots, e_{2n} denote the standard basis in \mathbb{R}^{2n} . Define $f_{u,v}$ to be the subspace of \mathbb{R}^{2n} spanned by the pair of vectors $e_u - e_v$ and $e_{n+u} - e_{n+v}$ (note that the definition of $f_{u,v}$ is symmetric in u, v). Let $h(\mathbf{x})$ denote the subspace of \mathbb{R}^{2n} orthogonal to the vector $\mathbf{x} = (y_1, \dots, y_n, -x_1, \dots, -x_n)$. It is not hard to verify (see [20] for the details) that $h(\mathbf{x}) \cap f_{u,v}$ is spanned by the $\{u, v\}$ row of $M_{G,2}$. Thus, for a generic \mathbf{x} , we have

$$\text{rank} M_{G,2} = \dim \text{span}\{h(\mathbf{x}) \cap f_{u,v} \mid \{u, v\} \in E\}.$$

Thus, the question of computing the rank of $M_{G,2}$ becomes the question of computing the dimension of the span of the resulting intersections (which here are simply lines) with a *generic* hyperplane. To analyze this, Lovász and Yemini use a theory developed by Lovász [19] which studies a similar problem for an arbitrary family of subspaces. The relevant part of Lovász's theory is introduced in the next subsection.

The idea of [20] can be applied also to rigidity in higher dimensions. For simplicity of the presentation, let us consider only the case $t = 3$. In this case we associate with each edge $\{u, v\} \in E$ a 3-dimensional subspace $g_{u,v}$ of \mathbb{R}^{3n} . Namely, the subspace spanned by the vectors $e_u - e_v$, $e_{n+u} - e_{n+v}$, $e_{2n+u} - e_{2n+v}$, where here e_1, \dots, e_{3n} stand for the standard basis of \mathbb{R}^{3n} . Let $\mathbf{x} = (x_1, \dots, x_n, y_1, \dots, y_n, z_1, \dots, z_n)$ and define $\tilde{h}(\mathbf{x})$ to be the (codim 2) subspace of \mathbb{R}^{3n} orthogonal to the pair of vectors

$$\begin{aligned} &(y_1, \dots, y_n, -x_1, \dots, -x_n, 0, \dots, 0) \\ &(z_1, \dots, z_n, 0, \dots, 0, -x_1, \dots, -x_n). \end{aligned}$$

It is not hard to verify that $\tilde{h}(\mathbf{x}) \cap g_{u,v}$ is one dimensional and spanned by the $\{u, v\}$ row of $M_{G,3}$. Thus, for a generic choice of \mathbf{x} , we have

$$\text{rank} M_{G,3} = \dim \text{span}\{\tilde{h}(\mathbf{x}) \cap g_{u,v} \mid \{u, v\} \in E\}.$$

A crucial difference from the case $t = 2$ is that here a generic choice of \mathbf{x} does not yield a generic codim 2 subspace $\tilde{h}(\mathbf{x})$ of \mathbb{R}^{3n} . From the perspective of this method and of our paper, this is "the reason" why rigidity in higher dimensions is more challenging.

1.3 Subspaces and generic hyperplanes

Let F be a collection of subspaces in \mathbb{K}^d . Let h be a generic hyperplane in \mathbb{K}^d , which without loss of generality can be taken to be all vectors perpendicular to $\mathbf{x} = (x_1, \dots, x_d)$.

For each subspace $f \in F$, let $f' = f \cap h$. Now consider the space spanned by the subspaces in $F' := \{f' \mid f \in F\}$ (note that the flats in F' are functions of \mathbf{x}). The question is, what is the dimension of $\text{span}(F')$?

One of the major results of Lovász' paper [19] is a formula, called $\rho(F)$ (which we redefine in Section 2), that determines this dimension for every family of subspaces, and for \mathbf{x} satisfying a certain “general position” condition (see Definition 5.1). To show that a *generic* \mathbf{x} satisfies Lovász's general position condition over any field (with large enough characteristic) is one main result of our paper (see Section 7). Note that this fact is mentioned (over the field \mathbb{R}) in [19] with no proof. This fact is again mentioned⁶ and applied, again with no proof, in Tanigawa [26]. We see our paper as contributing to the completeness of these results.

When the subspaces F are derived from a graph in the manner described above to generate the rigidity matrix, Lovász and Yemini [20] write the explicit special case of the formula $\rho(F)$, which yields an elegant characterization. For the general case of an arbitrary family of subspaces F , the formula is given as the minimum, over all possible partitions of the family, of a certain easily computable function. As the number of partitions is exponential, there is no obvious efficient way of computing ρ . We have recently learned that the problem of computing ρ is a special case of minimizing, over all partitions of a set S , the *Dilworth truncation* of a given submodular function f defined over S ; a strongly polynomial algorithm for this problem is given in Frank and Tardos [7, Chapters II.1 and IV.3]. In our paper we introduce an alternative⁷ strongly polynomial algorithm for computing ρ , by reducing the original problem to a minimization problem of a certain submodular function. In fact, we prove our result to a more general quantity $\rho_c(F)$, introduced in Section 2. (Note that $\rho(F) = \rho_1(F)$ is the quantity from [19].)

Theorem 1.6. *There is a deterministic, strongly polynomial time algorithm to compute ρ_c for every real number c .*

Closing this circle, we will also prove that the problem of computing ρ_1 is *equivalent* to $\text{PIT}(R_2)$. This will yield Theorem 1.4 as a corollary to Theorem 1.6.

1.4 Related works and applications

We see our result as a step towards better understanding of the algorithmic aspects of the notions and formulas introduced in Lovász [19] and their applications.

Let us mention one related concept studied in Lovász [19] and discuss follow-up work by Tanigawa [26], which is related to Theorem 5.2 proved in this paper. It would be interesting to find efficient algorithms for the natural computational problem at hand. The reader may skip this subsection at first reading.

Let F be a finite family of subspace in \mathbb{K}^d (where \mathbb{K} is a field of characteristic 0). Let $X = \{x_f \mid f \in F\}$ be a collection of points in \mathbb{K}^d such that $x_f \in f$ for each $f \in F$. The set X is said to be in *general position* with respect to F if, for every $f \in F$ fixed, the following

⁶In Tanigawa [26] an alternative general position condition is suggested, to supposedly correct a mistake in Lovász's paper. However we find the counter example in [26, footnote on p. 1416] false. We provide a full and detailed proof of Lovász's formula in Section 5.

⁷Our algorithm seems different than the one in [7], as it does not use duality.

holds: Any subspace spanned by members of F and points of $X \setminus \{x_f\}$ containing x_f must contain the whole flat f . Lovász shows that there exists a choice of a set X in general position with respect to any given family F . He then proves the following formula:

Theorem 1.7 (Lovász [19]). *Let F be a finite family of subspace in \mathbb{K}^d , and let $X = \{x_f \mid f \in F\}$ be in general position with respect to F . Then*

$$\text{rank}(\text{span}X) = \min_{G \subseteq F} \left\{ \text{rank}(\text{span} \bigcup G) + |F \setminus G| \right\}$$

An interesting application of Theorem 1.7 to the body-rod-bar rigidity problem is obtained by Tanigawa [26]. A *body-rod-bar framework* in \mathbb{R}^d is defined as a structure consisting of d -dimensional subspaces (bodies) and $(d-2)$ -dimensional flats (rods) mutually linked by one-dimensional lines (bars). (The term “rod” is appropriate for $d = 3$.) More formally, a d -dimensional body-rod-bar-framework is a triple (G, q, r) , where $G = (V = B \cup R, E)$ is a graph, $r : R \rightarrow \text{Gr}(d-1, \mathbb{R}^{d+1}) \subset \mathbb{P}(\wedge^{d-1}(\mathbb{R}^{d+1}))$ is the *rod-configuration* mapping a vertex $v \in R$ to a $(d-1)$ -dimensional subspace r_v of \mathbb{R}^{d+1} , and $q : E \rightarrow \text{Gr}(2, \mathbb{R}^{d+1}) \subset \mathbb{P}(\wedge^2(\mathbb{R}^{d+1}))$ is the *bar-configuration* mapping an edge $e \in E$ to a 2-dimensional subspace q_e in \mathbb{R}^{d+1} , such that

q_e and r_v have a nonzero intersection, whenever $v \in R$ is a vertex of e ;

equivalently,

$$q_e \cdot r_v = 0, \text{ whenever } v \in R \text{ is a vertex of } e,$$

where here the dot product should be interpreted appropriately (see [26] for the details). Assume also that $r(u) \neq r(v)$ for every $u \neq v \in R$.

An *infinitesimal motion* of (G, q, r) is a mapping $m : B \cup R \rightarrow \wedge^{d-1}(\mathbb{R}^{d+1})$ such that

$$q_e \cdot (m(u) - m(v)) = 0, \text{ for every } e = \{u, v\} \in E. \quad (1)$$

An infinitesimal motion m is called *trivial* if either $m(u) = m(v)$ for all $u, v \in V$, or if, for some fixed $v_0 \in V$ we have $m(v_0) = r_{v_0}$ and $m(v) = 0$ for every $v \in V \setminus \{v_0\}$. Finally, a framework (G, q, r) is called *infinitesimally rigid* if every infinitesimal motion is trivial.

The body-rod-bar problem gives rise to a matroid $\text{BR}(G, q, r)$ defined on the edge set E whose rank is the maximum size of independent linear equations in (1) (for unknown m). From the definition, (G, q, r) is infinitesimally rigid if and only if the rank of $\text{BR}(G, q, r)$ is $\binom{d+1}{2}|V| - \left(\binom{d+1}{2} + |R|\right)$.

Theorem 1.8 (Tanigawa [26, Corollary 4.13]). *Let $G = (B \cup R, E)$ and suppose $d \geq 3$. Then, for almost all bar-configurations q and almost all rod-configurations r we have*

$$\text{rank}(E) = \min_{\Pi = \{F_0, \dots, F_k\}} \left\{ |F_0| + \sum_{i=1}^k \left(\binom{d+1}{2}(V(F_i) - \binom{d+1}{2}) - R(F_i) \right) \right\},$$

where the minimum is taken over all partitions Π of E .

Tanigawa’s proof is a nice combination of Theorem 1.7 with the other result of Lovász mentioned in the introduction, cited below as Theorem 5.2. Briefly, the first (simpler) step in the proof is to reduce the problem to the form of Theorem 1.7. That is, a family of

flats F is introduced, and the question becomes to find the rank of a generic set of points $X = \{x_f \mid f \in F\}$. The family F resulted from the reduction can be described as follow: Each edge $e = \{u, v\}$ of G is associated with some fixed subspace f_e in $\left(\mathbb{P}(\wedge^2(\mathbb{R}^{d+1}))\right)^{|V|}$. Then $F = \{f_e \cap h(u) \cap h(v) \mid e = \{u, v\} \in E\}$, where $h_r(u), h_r(v)$ are subspaces depending on the choice of rod configuration r . Since r is taken generically, this imposes some genericity on the subspaces $h_r(v)$, but they are not exactly generic. The proof is then complete by proving a relaxed version of Theorem 5.2, and adding the subspaces $h_r(v)$ one after the other.

For more recent applications of [19, 20] see Tanigawa [26, 27].

1.5 Organization of this paper

In Section 2 we introduce the function $\rho_c(F)$, which is the main object of this study. The rest of the paper has two separate parts. The first, in Sections 3 and 4, describes the algorithm to compute ρ_c . In Section 3, we present and prove properties of the function ρ_c . Using these properties we describe, in Section 4, a deterministic strongly polynomial time algorithm that computes ρ_c over every field via submodular optimization. Note that, as there is an alternative algorithm [7] in the literature to efficiently compute functions like ρ_c , this part can be skipped.

The second part, in Sections 5, 6, and 7, describes the genericity proof of ρ . In Section 5, we state (and reprove) the result of Lovász [19] above, relating ρ_1 to the intersection of F with a hyperplane in “general position”. A similar relation is obtained for ρ_c , for an integer $c > 0$ (see Theorem 5.5). In Section 6, we develop an explicit representation of a basis of the family F' resulting from this intersection, which give rise to the symbolic matrices $\text{PIT}(R_2)$ (and $\text{PIT}(R_k)$). Using this, we prove in Section 7 that most hyperplanes (and more generally, subspaces) satisfy the “general position” definition of Lovász, thus expressing the rank of a these symbolic matrices as appropriate $\rho(F)$. Using the algorithm above we can now compute these ranks deterministically and efficiently. This last section is the only one in which the size of the field \mathbb{K} is important.

2 Subspaces, partitions, and the function ρ_c

We introduce the main objects of this study: Families of subspaces, their partitions, and the optimization problem we solve in this paper. We consider linear subspaces f of \mathbb{K}^d . Let $d(f)$ denote the dimension of a subspace f . For a family F of subspaces, we write $\text{span}F := \text{span}\bigcup_{f \in F} f$ and

$$d(F) := d(\text{span}F).$$

A *partition* of F is a set $\Pi = \{P_1, \dots, P_t\}$ of nonempty, pairwise disjoint subfamilies of F , such that $F = \bigcup_{i=1}^t P_i$. For a partition Π of F and a family of subspaces G , we define the *restriction* of Π to G by

$$\Pi \cap G := \{P \cap G \mid P \in \Pi, P \cap G \neq \emptyset\}. \quad (2)$$

If $G \subset F$, then $\Pi \cap G$ forms a partition of G .

Lovász [19] defined the following key function ρ of a family of subspaces, whose meaning will be revealed in Section 5. We actually generalize his definition to a family of functions ρ_c , for every $c > 0$ (his ρ is our ρ_1 for $c = 1$). Computing $\rho_c(F)$ in deterministic polynomial time given F , in Section 4, will be the key to our derandomization results.

Fix a constant $c > 0$. Let F be a finite family of subspaces in \mathbb{K}^d . For a partition Π of F , we define

$$\rho_c(F, \Pi) := \sum_{P \in \Pi} (d(P) - c).$$

$$\rho_c(F) := \min_{\Pi} \rho_c(F, \Pi), \tag{3}$$

where the minimum is taken over all partitions Π of F .

Definition 2.1. We say that Π is a *minimal* partition of F , with respect to the constant $c > 0$, if Π attains $\rho_c(F)$ and has the smallest possible number of parts.

Remark. In Corollary 3.2 we prove that, fixing $c > 0$, a minimal partition Π of a family F with respect to c is unique.

Notation. We will use small letters f, g, h to denote subspaces in \mathbb{K}^d , capital letters F, G, P, Q to denote families of subspaces, and Π to denote partitions of a certain family F of subspaces. Note that the elements of a partition Π are themselves families of subspaces.

3 Properties of minimal partitions

In this and the next section we develop our algorithm in a fully self-contained manner. As mentioned in the introduction, the reader may skip these sections and apply the algorithm of [7] as a black box. In this section, we introduce some properties of minimal partitions, to be used in our algorithm. We find these properties interesting in their own right, but some may be known, indeed in more generality, for submodular functions.

3.1 Main technical lemma

We start with the following main technical lemma of this section.

Lemma 3.1. *Let F, G be families of subspaces in \mathbb{K}^d with minimal partitions Π_F, Π_G , respectively. Assume that $Q \in \Pi_G$ and $Q \subset F$. Then Q is contained in one of the parts of Π_F .*

For the proof, the idea is to show that if, when considering a minimal partition for F , it “pays off” to put the elements of Q together, then it still “pays off” (or at least, harmless) to put these elements together, when this time considering a minimal partition for G .

Proof. Consider the restriction $\Pi' := \Pi_F \cap Q$ of Π_F to Q (as defined in (2)). By assumption, $Q \subset F$, and thus Π' forms a partition of Q .

Our assumption that $Q \in \Pi_G$, and recalling that Π_G forms a minimal partition of G , implies that

$$\sum_{P \in \Pi'} (d(P) - c) \geq d(Q) - c. \quad (4)$$

Fixing some arbitrary order on the elements of Π' , we write

$$\Pi' = (P'_1, \dots, P'_t),$$

where $P'_i := P_i \cap Q$ is non-empty and $P_1, \dots, P_t \in \Pi_F$ are distinct. Set $V'_0 := \{0\}$. For each $1 \leq i \leq t$, define

$$V'_i := \text{span} \left(\bigcup_{j=1}^i P'_j \right)$$

and put $r'_i := d(V'_i) - d(V'_{i-1})$ and $s'_i := d(P'_i) - r'_i$. Note that

$$d(Q) = \sum_{i=1}^t r'_i$$

and that

$$s'_i = d((\text{span} P'_i) \cap V'_{i-1}). \quad (5)$$

With this notation, (4) can be rewritten as

$$\sum_{i=1}^t (r'_i + s'_i) - tc \geq \sum_{i=1}^t r'_i - c$$

which implies

$$\sum_{i=1}^t s'_i \geq c(t-1). \quad (6)$$

Next, we define

$$V_i := \text{span} \left(\bigcup_{j=1}^i P_j \right)$$

and put $r_i := d(V_i) - d(V_{i-1})$ and $s_i := d(P_i) - r_i$. Similar to above, we have

$$d \left(\bigcup_{i=1}^t P_i \right) = \sum_{i=1}^t r_i$$

and

$$s_i = d((\text{span} P_i) \cap V_{i-1}). \quad (7)$$

We claim that

$$\sum_{i=1}^t (d(P_i) - c) \geq d \left(\bigcup_{i=1}^t P_i \right) - c. \quad (8)$$

Indeed, the inequality (8) holds if and only if

$$\sum_{i=1}^t (r_i + s_i) - tc \geq \sum_{i=1}^t r_i - c$$

which holds if and only if

$$\sum_{i=1}^t s_i \geq c(t-1). \quad (9)$$

To prove the last inequality, notice that $V'_i \subset V_i$ and $\text{span}P'_i \subset \text{span}P_i$, for every i . Thus

$$d((\text{span}P'_i) \cap V'_{i-1}) \leq d((\text{span}P_i) \cap V_{i-1}).$$

Hence, by (5) and (7), we get $s'_i \leq s_i$. This fact combined with the inequality (6) implies (9) and hence also (8). Since Π_F is assumed to be minimal for F , we conclude that $t = 1$ and $Q \subset P_1$. This completes the proof. \square

3.2 Uniqueness of minimal partitions

We prove uniqueness of minimal partitions.

Corollary 3.2 (Uniqueness). *Let F be a family of subspaces in \mathbb{K}^d and let Π_1, Π_2 be minimal partitions of F . Then $\Pi_1 = \Pi_2$.*

Proof. Let \sim_1, \sim_2 denote the equivalence relations on F induced by the partitions Π_1, Π_2 , respectively. Let $f, g \in F$ and assume that $f \sim_1 g$. That is $f, g \in Q$, for some $Q \in \Pi_1$. Applying Lemma 3.1 (with $F, G := F$, and Q), we get that Q is contained in one of the parts in Π_2 . Thus $f \sim_2 g$. By symmetry, we conclude that $f \sim_1 g$ if and only if $f \sim_2 g$. Thus $\Pi_1 = \Pi_2$, as claimed. \square

Definition 3.3. Fix $c > 0$. Define $\Pi^*(F)$ to be *the* minimal partition of a family of subspaces F (with respect to c).

3.3 Monotonicity properties

We prove the following ‘‘monotonicity’’ property of minimal partitions.

Corollary 3.4 (Monotonicity). *Let F, G be families of subspaces in \mathbb{K}^d and assume that $G \subset F$. Then $\Pi^*(G)$ is a refinement of $\Pi^*(F) \cap G$.*

Proof. Apply Lemma 3.1 to the families F and G . \square

The following is another type of monotonicity property.

Lemma 3.5. *Let $F = \{f_1, \dots, f_n\}$ be a family of n subspaces in \mathbb{K}^d . Let $f_i \subset f'_i$, for every $i = 1, \dots, n$, and consider $F' := \{f'_1, \dots, f'_n\}$. For a partition Π of F , let Π' denote the partition of F' induced by Π , replacing each f_i by the corresponding f'_i . Then $(\Pi^*(F))'$ is a refinement of $\Pi^*(F')$.*

Proof. Let $P \in \Pi^*(F)$ and assume without loss of generality that $P = \{f_1, \dots, f_m\}$, for some $m \leq n$. It is easy to see, applying Lemma 3.1, that $\Pi^*(P) = \{P\}$.

Put $P' := \{f'_1, \dots, f'_m\}$. We claim that $\Pi^*(P') = \{P'\}$. First note that it suffices to prove the claim for the special case where $f_1 \subset f'_1$ and $f_i = f'_i$, for $i = 2, \dots, m$, and then

apply the same argument repeatedly to each i . To prove the claim for the special case, consider the family $Q = \{f_1, f'_1\}$. It is easy to see, by definition, that $\Pi^*(Q) = \{Q\}$. By Lemma 3.1, Q is contained in a part of $\Pi^*(G)$, for every family of subspaces G that contains Q . Moreover, since $f_1 \cup f'_1 \subset f'_1$, we have

$$\rho_c(G) = \rho_c(G \setminus \{f_1\}) \quad \text{and} \quad \Pi^*(G \setminus \{f_1\}) = \Pi^*(G) \cap (G \setminus \{f_1\})$$

for every such G (this follows directly from the definition of ρ_c and of Π^*).

Define $G := \{f_1, f'_1, f_2, \dots, f_m\}$. By what has just been argued, we have

$$\Pi^*(P') = \Pi^*(G) \cap P'. \tag{10}$$

Since $P, Q \subset G$, and applying Lemma 3.1, we get that each of P and Q is contained in a part of $\Pi^*(G)$. But $P \cap Q \neq \emptyset$, thus the set $P \cup Q$ must be contained in a part of $\Pi^*(G)$. Noting that $P \cup Q = G$, this implies that $\Pi^*(G) = \{G\}$. Combined with (10), this proves $\Pi^*(P') = P'$, as claimed.

Applying Lemma 3.1 to the families F' , P' , and with $P' \in \Pi^*(P')$, we conclude that P' is contained in one of the parts of $\Pi^*(F')$. Since this is true for every $P \in \Pi^*(F)$, the lemma follows. \square

3.4 The family \hat{F}

Let F be a family of subspaces in \mathbb{K}^d . We show that, in some sense, F can be replaced by a simpler family \hat{F} defined next. With each $P \in \Pi^*(F)$ associate the subspace $f_P := \text{span}P$. Then define the family

$$\hat{F} := \{f_P \mid P \in \Pi^*(F)\}.$$

Note that for $P \neq P'$ we have $f_P \neq f_{P'}$; otherwise, taking $P \cup P'$ yields a partition of F with strictly less parts and with smaller or equal value of ρ_c , contradicting the minimality of $\Pi^*(F)$.

The family F can be replaced by \hat{F} in the sense of Lemma 3.6, and \hat{F} is simpler in the sense of Lemma 3.7.

Lemma 3.6. *Let F, G be families of subspaces in \mathbb{K}^d . Then*

$$\rho_c(F \cup G) = \rho_c(\hat{F} \cup G) \quad \text{and} \quad \Pi^*(F \cup G) \simeq \Pi^*(\hat{F} \cup G).$$

By the sign \simeq we mean that the identity holds after identifying the partition $\Pi^*(\hat{F} \cup G)$ of $\hat{F} \cup G$ with the partition of $F \cup G$ naturally induced by it. Concretely, the lemma asserts that

$$\Pi^*(F \cup G) = \left\{ \left(\bigcup_{f_P \in \hat{Q}} P \right) \cup (G \cap \hat{Q}) \mid \hat{Q} \in \Pi^*(\hat{F} \cup G) \right\}.$$

Proof. In the proof we often abuse notation and regard a partition of $\hat{F} \cup G$ as a one of $F \cup G$, as explained after the statement of the lemma. Let Π^* be the partition of $F \cup G$ induced by $\Pi^*(\hat{F} \cup G)$, given by

$$\Pi^* = \left\{ \left(\bigcup_{f_P \in \hat{Q}} P \right) \cup (G \cap \hat{Q}) \mid \hat{Q} \in \Pi^*(\hat{F} \cup G) \right\}.$$

We have $|\Pi^*| = |\Pi^*(\hat{F} \cup G)|$ and

$$\rho_c(F \cup G, \Pi^*) = \rho_c(\hat{F} \cup G, \Pi^*(\hat{F} \cup G)).$$

Thus

$$\rho_c(F \cup G) \leq \rho_c(\hat{F} \cup G).$$

To prove the inverse inequality, apply Lemma 3.1 to the families F and $F \cup G$. It follows that, for every $P \in \Pi^*(F)$, there exists $Q \in \Pi^*(F \cup G)$ such that $P \subset Q$. This means that $\Pi^*(F \cup G)$ induces a well-defined partition $\hat{\Pi}^*$ of $\hat{F} \cup G$ with $|\Pi^*(F \cup G)| = |\hat{\Pi}^*|$ and

$$\rho_c(F \cup G, \Pi^*(F \cup G)) = \rho_c(\hat{F} \cup G, \hat{\Pi}^*). \quad (11)$$

Concretely, $\hat{\Pi}^*$ is given by

$$\hat{\Pi}^* := \{\hat{Q} \mid Q \in \Pi^*(F \cup G)\},$$

where

$$\hat{Q} := \{f_P \mid P \subset Q, P \in \Pi^*(F)\} \cup (Q \cap G).$$

We have

$$\begin{aligned} \rho_c(F \cup G) &= \rho_c(F \cup G, \Pi^*(F \cup G)) \\ &= \rho_c(\hat{F} \cup G, \hat{\Pi}^*) \\ &\geq \rho_c(\hat{F} \cup G). \end{aligned}$$

This proves that $\rho_c(F \cup G) = \rho_c(\hat{F} \cup G)$.

Next, we claim that $|\Pi^*(F \cup G)| = |\Pi^*(\hat{F} \cup G)|$. Indeed, by our argument above, the partition $\hat{\Pi}^*$ of $\hat{F} \cup G$ satisfies

$$\rho_c(\hat{F} \cup G, \hat{\Pi}^*) = \rho_c(\hat{F} \cup G) \quad \text{and} \quad |\hat{\Pi}^*| = |\Pi^*(F \cup G)|.$$

Since $\Pi^*(\hat{F} \cup G)$ is taken to be the smallest that attains $\rho_c(\hat{F} \cup G)$, we get

$$|\Pi^*(\hat{F} \cup G)| \leq |\Pi^*(F \cup G)|.$$

Similarly, by our argument above, the partition Π^* of $F \cup G$ satisfies

$$\rho_c(F \cup G, \Pi^*) = \rho_c(F \cup G) \quad \text{and} \quad |\Pi^*| = |\Pi^*(\hat{F} \cup G)|.$$

Thus,

$$|\Pi^*(F \cup G)| \leq |\Pi^*(\hat{F} \cup G)|.$$

This proves the claim.

By the uniqueness of minimal partition (see Corollary 3.2), we conclude that

$$\Pi^*(\hat{F} \cup G) = \hat{\Pi}^* \quad \text{and} \quad \Pi^*(F \cup G) = \Pi^*.$$

This completes the proof of the lemma. □

Lemma 3.7. *Let F be a family of subspaces in \mathbb{K}^d . Then*

$$\Pi^*(\hat{F}) = \{\{\hat{f}\} \mid \hat{f} \in \hat{F}\}.$$

Proof. Apply Lemma 3.6 with $G = \emptyset$. □

We introduce one more simple property that we need.

Lemma 3.8. $\widehat{F \cup G} = \widehat{\hat{F} \cup G}$.

Proof. By Lemma 3.6, $\Pi^*(F \cup G) = \Pi^*(\hat{F} \cup G)$. The assertion then easily follows. □

4 An algorithm for computing $\rho_c(F)$

In this section we prove Theorem 1.6. That is, we introduce an algorithm to compute $\rho_c(F)$, for any number c and a given family F of n subspaces in \mathbb{K}^d , with polynomial running time in n (and in d). While we designed our algorithm for the class of functions ρ_c , it clearly works for a wider class of submodular functions. As it is different than the one in [7], we feel it would be interesting to explore its generality. Note that the problem is trivial for $c \leq 0$, which is why we consider only $c > 0$.

As mentioned in the introduction, the problem of computing ρ_c turns out to be an instance of a more general problem to which a strongly polynomial time algorithm is already known [7]. In more detail, the *Dilworth truncation* of a set function $b' : 2^S \rightarrow \mathbb{R} \cup \{\infty\}$ is defined as the function

$$b(X) = \min_{\Pi} \sum_{P \in \Pi} b'(P),$$

where the minimum is taken over all partitions Π of X .

Theorem 4.1 (Frank and Tardos [7, IV.3]). *Let $b' : 2^S \rightarrow \mathbb{R} \cup \{\infty\}$ be a submodular set function. Suppose that a minimizing oracle for b' is available. Then $b(S)$ can be computed in a strongly polynomial time. The algorithm also constructs a partition Π of S for which $b(S) = \sum_{P \in \Pi} b'(P)$.*

Remark. In [7], a more general result is proved.

4.1 High-level description of the algorithm for ρ_c

The input to the algorithm is a number c and a family of subspaces $F = \{f_1, \dots, f_n\}$ in \mathbb{K}^d . Write $F_i := \{f_1, \dots, f_i\}$. The high-level scheme of the algorithm is the following:

1. $\hat{F}_1 \leftarrow \{f_1\}$.
2. **For** $i \leftarrow 2$ **to** n
 - 2.1. $\Pi \leftarrow$ **Compute** $\Pi^*(\hat{F}_{i-1} \cup \{f_i\})$
 - 2.2. $\hat{F}_i \leftarrow \{\text{span}(P) \mid P \in \Pi\}$
3. **Return** $\sum_{f \in \hat{F}_n} (d(f) - c)$

The heart of the algorithm is of course the missing description of Step 2.1, which computes, in the i th iteration, the minimal partition of the family $\hat{F}_{i-1} \cup \{f_i\}$ with respect to ρ .

Lemma 4.2. *The computation in Step 2.1 can be done in strongly-polynomial time.*

Recall that the minimal partition of \widehat{F}_{i-1} is the partition into singletons, by Lemma 3.7. So in this step we compute the effect on this partition of inserting one new subspace. We explain how to do so efficiently and prove Lemma 4.2 in Section 4.3 below. To describe and analyze step 2.1, we first need to recall submodular functions and optimization, which we do in Section 4.2. The proof of the lemma is then given in Section 4.3.

We are now ready to prove Theorem 1.6, assuming that Lemma 4.2 is true.

Proof of Theorem 1.6. Correctness of the algorithm. By Lemma 3.8, we have

$$\widehat{F}_i = \widehat{F}_{i-1} \widehat{\cup} \{f_i\}.$$

Thus the computation of \widehat{F}_i in Step 2.2 is correct. In view of Lemmas 3.6 and 3.7, the algorithm's output is $\rho_c(F)$, as needed.

Running time of the algorithm. We represent a k -dimensional subspace f in \mathbb{K}^d by a $k \times d$ matrix whose rows form a basis for f . The dimension $d(f)$ of a subspace f is just the number of rows in the matrix representing the subspace, and hence can be computed in a constant time. Let P be a family of subspaces in \mathbb{K}^d . To compute $\text{span}(P)$, we take the union of the rows of the matrices in P (representing subspaces) and apply Gauss elimination (using row operations only). If P has n subspaces, we will need to apply Gauss elimination to a matrix of dimensions at most $(nd) \times d$. The nonzero rows in the matrix received by this process will form a basis for $\text{span}(P)$.

Now let F be a family of n subspaces in \mathbb{K}^d . Clearly, each line in the above description of the algorithm, when applied to F , is called at most n times. In each step, excluding Step 2.1, we are required to compute at most n times one of the operations just described (finding dimension or span) or simple operations such as addition. In view of Lemma 4.2, the proof is complete. \square

4.2 A submodular set function

Recall that a function s defined on the collection of subsets of a finite set A is called *submodular* if

$$s(X) + s(Y) \geq s(X \cup Y) + s(X \cap Y)$$

for all $X, Y \subset A$.

The following is proved by Schrijver in [23].

Theorem 4.3 (Schrijver [23]). *There exists a strongly polynomial-time algorithm minimizing a submodular function s , where s is given by an oracle. The number of oracle calls is bounded by a polynomial in the size of the underlying set. The algorithm also finds a minimizer X^* of s .*

In this section we consider a set function defined as follows. Let F be a family of subspaces in \mathbb{K}^d and let $g \subset \mathbb{K}^d$ be a subspace not in F . Fix $c > 0$. Define $r_{F,g,c} : 2^F \rightarrow \mathbb{K}$ by

$$r_{F,g,c}(X) := d(X \cup \{g\}) - c + \sum_{f \in \overline{X}} (d(f) - c),$$

where $\overline{X} := F \setminus X$. We then put

$$r_{F,g,c}^* := \min_{X \subset F} r_{F,g,c}(X)$$

and we let $X_{F,g,c}^*$ denote a subset $X \subset F$ that attains $r_{F,g,c}^*$.

We show that $r_{F,g,c}$ is submodular.

Lemma 4.4. *Let F and g and c be as above. Then $r_{F,g,c}$ is submodular.*

Proof. To simplify the notation, and as F, g, c are fixed, we write for short $r = r_{F,g,c}$. Let $X, Y \subset F$. We need to show

$$r(X) + r(Y) \geq r(X \cup Y) + r(X \cap Y).$$

Put $f_X := \text{span}(X \cup \{g\})$. By definition, we have

$$\begin{aligned} r(X) + r(Y) &= d(X \cup \{g\}) + d(Y \cup \{g\}) + \sum_{f \in \overline{X}} d(f) + \sum_{f \in \overline{Y}} d(f) - c|\overline{X}| - c|\overline{Y}| - 2c \\ &= d(f_X) + d(f_Y) + \sum_{f \in \overline{X}} d(f) + \sum_{f \in \overline{Y}} d(f) - c|\overline{X}| - c|\overline{Y}| - 2c. \end{aligned}$$

By basic linear algebra, we have the identity

$$d(f_X) + d(f_Y) = d(\text{span}(f_X \cup f_Y)) + d(f_X \cap f_Y).$$

Thus the last equality, after some rearranging, is

$$\begin{aligned} r(X) + r(Y) &= \\ &= (d(\text{span}(f_X \cup f_Y)) - c + \sum_{f \in \overline{X \cap Y}} d(f) - c|\overline{X \cap Y}|) + (d(f_X \cap f_Y) - c + \sum_{f \in \overline{X \cup Y}} d(f) - c|\overline{X \cup Y}|) \end{aligned}$$

Noting that $\text{span}(f_X \cup f_Y) = \text{span}(f_{X \cup Y})$ and that $\text{span}(f_X \cap f_Y) \supset \text{span}(f_{X \cap Y})$, we get

$$\begin{aligned} r(X) + r(Y) &\geq \left(d(f_{X \cup Y}) - c + \sum_{f \in \overline{X \cup Y}} d(f) - c|\overline{X \cup Y}| \right) + \left(d(f_{X \cap Y}) - c + \sum_{f \in \overline{X \cap Y}} d(f) - c|\overline{X \cap Y}| \right) \\ &= r(X \cup Y) + r(X \cap Y). \end{aligned}$$

This proves the lemma. □

4.3 Inserting one subspace

We are now ready to describe in detail Step 2.1 which computes \widehat{F}_i given \widehat{F}_{i-1} and f_i . More precisely, we describe a subroutine that receives as an input a family F with $F = \widehat{F}$ and a subspace g , and outputs $\Pi^*(F \cup \{g\})$.

We will need the following observation.

Lemma 4.5. *Let $G = F \cup \{g\}$ be a family of subspaces in \mathbb{K}^d . Let $Q_g \in \Pi^*(G)$ be the part that contains the subspace g . Then*

$$\Pi^*(G) \setminus \{Q_g\} \subset \Pi^*(F).$$

Proof. For every $Q \in \Pi^*(G) \setminus \{Q_g\}$, we have $Q \subset F$. By Lemma 3.1, there exists $P \in \Pi^*(F)$ such that $Q \subset P$. Clearly, we also have $P \subset G$. Applying Lemma 3.1 once again, we get that also $P \subset Q$. Thus, $P = Q$ which means that $Q \in \Pi^*(F)$. \square

Corollary 4.6. *Let F be a family of n subspaces in \mathbb{K}^d with $\hat{F} = F$ and let g be another subspace in \mathbb{K}^d . Then $\rho_c(F \cup \{g\}) = r_{F,g,c}^*$ and*

$$\Pi^*(F \cup \{g\}) = \{X_{F,g,c}^* \cup \{g\}\} \cup \{\{f\} \mid f \in F \setminus X_{F,g,c}^*\},$$

where $X_{F,g,c}^*$ and $r_{F,g,c}^*$ are as defined in Section 4.2.

Proof. This follows from the definitions of ρ_c and $r_{F,g,c}^*$, combined with Lemma 4.5. \square

Proof of Lemma 4.2. Combinig Corollary 4.6 with Theorem 4.3, we get that the computation in Step 2.1 can be done in strongly-polynomial time. \square

5 Intersecting subspaces with a hyperplane

In this section we state (and reprove) a result of Lovász [19], which explains the source of the function ρ (more precisely, taking ρ_c with $c = 1$) as the dimension of the intersections of a family of subspaces with a hyperplane in “general position”. This connection has been used by Lovász to study certain questions about matroids in [19], and by Lovász and Yemini in [20] to study rigid structures in \mathbb{R}^2 . We extend Lovász’ treatment to arbitrary fields \mathbb{K} .

In Theorem 5.5 below, we further extend Lovász’s result, in a straightforward manner, to apply to the intersection of a family of subspaces with an arbitrary subspace (of any co-dimension) in “general position”, instead of only a (co-dimension 1) hyperplane.

Lovász [19] uses a very specific notion of genericity, which he calls *general position* defined below, and shows that ρ correctly computes the dimension of the intersection when the hyperplane is in general position with respect to the given family of subspaces. In Theorem 7.1 we will prove that indeed “general position” is a generic property, namely holds for almost all hyperplanes. This will complete the connection with the PIT problem solved in this paper.

A *hyperplane* in \mathbb{K}^d is a subspace (subspace of \mathbb{K}^d) of codimension 1. Let F be a family of (nonzero) subspaces in \mathbb{K}^d and let $h \subset \mathbb{K}^d$ be a hyperplane in \mathbb{K}^d . We denote by $F \cap h$ the family $\{f \cap h \mid f \in F\}$. Following Lovász, we have the following definition:

Definition 5.1 (General Position). We say that h is in *general position* with respect to F if, for every $A, B, C \subset F$, with A nonempty, we have:

- (i) If $\text{span}(A) \subset h$, then $\text{span}(A) = \{0\}$.
- (ii) If⁸

$$\text{span}((A \cap h) \cup B) \cap \text{span}((A \cap h) \cup C) \subset h,$$

then

$$\text{span}((A \cap h) \cup B) \cap \text{span}((A \cap h) \cup C) = \text{span}(A \cap h).$$

⁸Note that here one can take any of A, B, C to be the empty set, and we interpret $\text{span}(\emptyset) = \{0\}$.

Remark. In Section 6, we prove (in Theorem 7.1) that being in general position with respect to a given family F is a generic property; this fact is mentioned in [19] without a proof.

Theorem 5.2 (Lovász [19, Theorem 2.3]). *Let F be a family of subspaces in \mathbb{K}^d . Let h be a hyperplane in \mathbb{K}^d in general position with respect to F . Then*

$$\rho_1(F) = d(F \cap h)$$

For completeness, we introduce a slightly more detailed proof, based on the line of argument from [19].

Proof of Theorem 5.2. Fix F and h as in the statement. Let $F' := F \cap h$. We need to show that $\rho_1(F) = d(F')$.

We first prove that $d(F') \leq \rho_1(F)$. That is, equivalently, we show that $d(F') \leq \rho_1(F, \Pi)$, for every partition Π of the family F . Let Π be a partition of F . For $P \in \Pi$, let $P' := P \cap h$. Then

$$\text{span}(F') = \text{span}\left(\bigcup_{P \in \Pi} \text{span}(P')\right)$$

and hence

$$d(F') \leq \sum_{P \in \Pi} d(P').$$

Note also that, for every $P \in \Pi$, we have $\text{span}(P') \subset \text{span}(P) \cap h$ and hence

$$d(P') \leq d(\text{span}(P) \cap h) = d(P) - 1,$$

where here we used property (i) of the general position assumption on h , namely, we used the fact that $\text{span}(P)$ is not contained in h . We conclude that

$$d(F') \leq \sum_{P \in \Pi} (d(P) - 1), \tag{12}$$

for every partition Π of F . This implies $d(F') \leq \rho_1(F)$.

To prove the reverse inequality, we show that, for a certain partition Π^* of F , the inequality (12) is in fact tight. We will construct Π^* explicitly subsequently refining a given partition. We describe the first step, which is indeed the general step (the proof will allow us to proceed recursively).

Define an equivalence relation on F as follows: For $f_1, f_2 \in F$, $f_1 \sim f_2$ if and only if

$$\text{span}(F' \cup \{f_1\}) = \text{span}(F' \cup \{f_2\}).$$

Let $\{P_1, \dots, P_m\}$ be the partition (equivalence classes) of F induced by the relation \sim .

The main idea is to prove that after intersection with h , the spans of the parts P'_i become a direct sum decomposition of $\text{span}(F')$. As we will see below, Π^* will be achieved by refining the partition $\{P_1, \dots, P_m\}$ inductively.

Lemma 5.3. *We have*

$$\text{span}(F') = \bigoplus_{i=1}^m \text{span}(P'_i). \tag{13}$$

Before we prove Lemma 5.3, we establish some preliminary claims. Let g_1, \dots, g_m be the (distinct) subspaces $g_i := \text{span}(F' \cup \{f\})$ for some $f \in P_i$ (note that by construction g_i is independent of the specific element $f \in P_i$ that we take).

We observe that, for every $1 \leq i \leq m$,

$$d(g_i) = d(F') + 1. \quad (14)$$

Indeed, by property (i) of general position, f is not contained in h and $\dim(f \cap h) = \dim(f) - 1$, for every $f \in F$. Hence, for every $f \in F$, one can choose a basis for f with all elements of the basis in h except for exactly one element b_f which is not in h . Thus, fixing any $f \in P_i$, we have

$$g_i = \text{span}(F' \cup \{f\}) = \text{span}(F' \cup \{b_f\}) = \text{span}(F') \oplus \text{span}\{b_f\}.$$

Thus, $d(g_i) = d(F') + 1$, as needed.

Next, we observe that, for $i \neq j$, we have

$$g_i \cap g_j = \text{span}(F') \subset h. \quad (15)$$

Indeed, by construction $g_i \neq g_j$, and in particular $g_i \cap g_j \subsetneq g_i$. Combining this with (14), we get $d(g_i \cap g_j) \leq d(g_i) - 1 = d(F')$. By the definition of g_i, g_j , we also have $\text{span}(F') \subset g_i \cap g_j$. Hence $g_i \cap g_j = \text{span}(F')$ and (15) follows.

Proof of Lemma 5.3. Here property (ii) of the general position definition will be crucial for the induction step. If $m = 1$ then (13) clearly holds. For $m \geq 2$, it suffices to show that, for every $2 \leq k \leq m$ and every distinct indices $1 \leq i_1, \dots, i_k \leq m$, one has

$$\text{span}(P'_{i_1} \cup \dots \cup P'_{i_{k-1}}) \cap \text{span}(P'_{i_k}) = \{0\}. \quad (16)$$

We prove (16) by induction on k . For $k = 2$, we need to show that $\text{span}(P'_{i_1}) \cap \text{span}(P'_{i_2}) = \{0\}$, for every distinct $1 \leq i_1, i_2 \leq m$. By the definition of the subspaces g_{i_1}, g_{i_2} and applying (15), we have

$$\text{span}(P_{i_1}) \cap \text{span}(P_{i_2}) \subset g_{i_1} \cap g_{i_2} \subset h.$$

Since h is in general position, using property (ii), this implies that $\text{span}(P_{i_1}) \cap \text{span}(P_{i_2}) = \{0\}$. This proves the induction base case $k = 2$.

Assume next that (16) holds for some $2 \leq k \leq m - 1$ fixed and for every distinct indices $1 \leq i_1, \dots, i_k \leq m$. Let $1 \leq i_1, \dots, i_{k+1} \leq m$ be some distinct indices. To establish the induction step we need to prove

$$\text{span}(P'_{i_1} \cup \dots \cup P'_{i_k}) \cap \text{span}(P'_{i_{k+1}}) = \{0\}. \quad (17)$$

Observe that in order to prove (17) it suffices to show that

$$\text{span}(P'_{i_1} \cup \dots \cup P'_{i_k}) \cap \text{span}(P'_{i_2} \cup \dots \cup P'_{i_{k+1}}) \subset \text{span}(P'_{i_2} \cup \dots \cup P'_{i_k}). \quad (18)$$

Indeed, assume that (18) holds. Then

$$\begin{aligned} \text{span}(P'_{i_1} \cup \dots \cup P'_{i_k}) \cap \text{span}(P'_{i_{k+1}}) &= \text{span}(P'_{i_1} \cup \dots \cup P'_{i_k}) \cap \text{span}(P'_{i_2} \cup \dots \cup P'_{i_{k+1}}) \cap \text{span}(P'_{i_{k+1}}) \\ &\subset \text{span}(P'_{i_2} \cup \dots \cup P'_{i_k}) \cap \text{span}(P'_{i_{k+1}}), \end{aligned}$$

where the first line uses the trivial fact that $\text{span}(P'_{i_{k+1}}) \subset \text{span}(P'_{i_2} \cup \dots \cup P'_{i_{k+1}})$ and the second line is due to (18). By the induction hypothesis, we have

$$\text{span}(P'_{i_2} \cup \dots \cup P'_{i_k}) \cap \text{span}(P'_{i_{k+1}}) = \{0\}.$$

Thus, assuming that (18) is true, (17) follows.

Finally, we now prove (18). Note that, by the definition of the subspaces g_i and using (15), we have

$$\text{span}(P_{i_1} \cup (P'_{i_2} \cup \dots \cup P'_{i_k})) \cap \text{span}((P'_{i_2} \cup \dots \cup P'_{i_k}) \cup P_{i_{k+1}}) \subset g_{i_1} \cap g_{i_{k+1}} \subset h.$$

Hence, our assumption that h is in general position with respect to F implies that in fact

$$\text{span}(P_{i_1} \cup (P'_{i_2} \cup \dots \cup P'_{i_k})) \cap \text{span}((P'_{i_2} \cup \dots \cup P'_{i_k}) \cup P_{i_{k+1}}) \subset \text{span}(P'_{i_2} \cup \dots \cup P'_{i_k}).$$

This clearly implies (18). Thus we have established the inductive step and this completes the proof of Lemma 5.3. \square

Recall that our goal is to show that (12) is tight for some partition Π^* of F . In view of Lemma 5.3, for the partition $\{P_1, \dots, P_m\}$ defined above, one has

$$d(F') = \sum_{i=1}^m d(P'_i). \quad (19)$$

That is, we expressed the quantity $d(F')$ as the sum of the quantities $d(P'_i)$ for certain subfamilies P_1, \dots, P_m of F . This allows to prove the existence of Π^* using induction on the size of F .

If $|F| = 1$, the unique partition on F clearly attains (12). For $|F| \geq 1$, let $\{P_1, \dots, P_m\}$ be the partition of F given by Lemma 5.3, satisfying (19). If $m = 1$, the identity (19), combined with (14), gives

$$d(F') = d(P_1) - 1.$$

This means that (12) is tight, and thus $\Pi^* = \{P_1\}$. If $m > 1$, then each subfamily P_i has fewer elements than F . Applying the induction hypothesis, there exist subpartitions $\Pi_i^* = \{P_{i1}, \dots, P_{im_i}\}$ of P_i , for each $1 \leq i \leq m$, satisfying

$$d(P_i) = \sum_{j=1}^{m_i} (d(P_{ij}) - 1).$$

Combined with (19), we get

$$d(F') = \sum_{i=1}^m \sum_{j=1}^{m_i} (d(P_{ij}) - 1).$$

So $\Pi^* := \bigcup_{i=1}^m \Pi_i^*$ forms a partition of F that attains (12). This completes the proof of the theorem. \square

Remark 5.4. Note that in the inductive proof of Lemma 5.3, it was sufficient to consider not all k -subsets of the P_i in the given partition, but rather simply on intervals P_2, P_3, \dots, P_k . The same induction on k works without change. Thus even after refinement, in the proof of this theorem we never need to apply the “general position” condition more than $|F|$ times. This will help us later bound the show that $\rho_1(F)$ correctly computes $\dim(F \cap h)$ for most (or generic) hyperplanes h even when \mathbb{K} is finite and not too large.

We now generalize the theorem above to intersecting a family of subspaces with an arbitrary subspace. For this we need to extend the definition of “general position”.

Let F be a family of subspaces in \mathbb{K}^d . Let $\{\mathbf{x}_1, \dots, \mathbf{x}_k\}$ be a set of vectors, and define that the subspaces $h_i = \{\mathbf{x}_1, \dots, \mathbf{x}_i\}^\perp$. Note that h_i is of codimension i in \mathbb{K}^d , and that $h'_i := h_i \cap h_{i-1}$ is a hyperplane in h_{i-1} , for $i = 1, \dots, k$. We say that the subspace $h = h_k$ is in *general position* with respect to F if for all $i \in [k]$ we have that the hyperplane h'_i is in general position with respect to the family $F_i = F \cap h_{i-1}$.

Theorem 5.5. Let F be a family of subspaces in \mathbb{K}^d . Let h be a subspace in \mathbb{K}^d of codimension k in general position with respect to F . Then

$$\rho_k(F) = d(F \cap h)$$

Proof. We prove by induction on the codimension k . The case $k = 1$ is Theorem 5.2.

Let $\mathbf{x}_1, \dots, \mathbf{x}_k \in \mathbb{K}^d$ be vectors such that $h = \{\mathbf{x}_1, \dots, \mathbf{x}_k\}^\perp$ is in general position with respect to h . We know that h'_k is in general position with respect to the family $F_k := F \cap h_{k-1}$. By Theorem 5.2 again, we have

$$\begin{aligned} d(F \cap h) &= d(F_k \cap h'_k) = \rho_1(F_k) \\ &= \min_{\Pi_k} \sum_{P' \in \Pi_k} (d(P') - 1), \end{aligned}$$

where the minimum ranges over all partitions Π_k of F_k . Note that Π_k induces a partition Π on F , in the obvious way. Moreover, for every $P' \in \Pi_k$ there exists $P \subset F$ such that $P' = P \cap h_{k-1}$. By induction, we get

$$d(P') = d(P \cap h_{k-1}) = \rho_{k-1}(P).$$

Thus,

$$\begin{aligned} d(F \cap h) &= \min_{\Pi} \sum_{P \in \Pi} (\rho_{k-1}(P) - 1) \\ &= \min_{\Pi} \sum_{P \in \Pi} \left(\left(\min_{\Pi_P} \sum_{Q \in \Pi_P} (d(Q) - k + 1) \right) - 1 \right), \end{aligned}$$

where the first minimum (the outer one) in this expression is taken over all partitions Π of F , and, fixing Π and given $P \in \Pi$, the inner minimum is taken over all partitions Π_P of the family P .

Note that, for any partition Π of F , the partitions $\{\Pi_P \mid P \in \Pi\}$ induce a new partition Π' which is a refinement of Π . Namely, $\Pi' := \bigcup_{P \in \Pi} \Pi_P$. Note that taking $\Pi_P = \{P\}$ for

each $P \in \Pi$, we get

$$\begin{aligned}
d(F \cap h) &\leq \min_{\Pi} \sum_{P \in \Pi} \left(\left(\sum_{Q \in \{P\}} (d(Q) - k + 1) \right) - 1 \right) \\
&= \min_{\Pi} \sum_{P \in \Pi} (d(P) - k) \\
&= \rho_k(F).
\end{aligned} \tag{20}$$

We now prove the inverse inequality. Fix a partition Π of F , and, for $P \in \Pi$, let Π_P^* be a partition of P that attains the minimum in

$$\min_{\Pi_P} \sum_{Q \in \Pi_P} (d(Q) - k + 1).$$

That is, the partitions Π_P^* satisfy

$$\sum_{P \in \Pi} \left(\left(\min_{\Pi_P} \sum_{Q \in \Pi_P} (d(Q) - k + 1) \right) - 1 \right) = \sum_{P \in \Pi} \left(\left(\sum_{Q \in \Pi_P^*} (d(Q) - k + 1) \right) - 1 \right)$$

Let $(\Pi')^*$ be the partition of F induced by $\bigcup \{\Pi_P^* \mid P \in \Pi\}$. Observe that

$$\begin{aligned}
d(F \cap h) &= \min_{\Pi} \sum_{P \in \Pi} \left(\left(\sum_{Q \in \Pi_P^*} (d(Q) - k + 1) \right) - 1 \right) \\
&\geq \min_{\Pi} \sum_{P \in \Pi} \sum_{Q \in \Pi_P^*} ((d(Q) - k + 1) - 1) \\
&= \min_{\Pi} \sum_{Q \in (\Pi')^*} (d(Q) - k) \\
&= \min_{(\Pi')^*} \sum_{Q \in (\Pi')^*} (d(Q) - k) \\
&\geq \min_{\Pi} \sum_{Q \in \Pi} (d(Q) - k) \\
&= \rho_k(F).
\end{aligned} \tag{21}$$

Combining the inequalities (20) and (21), we get $d(F \cap h) = \rho_k(F)$. This completes the induction step, and therefore proves the theorem. \square

6 Rank of symbolic matrices

In this section we show that the quantity $\rho_c(F)$ can be interpreted as the generic rank, defined as the rank over $\mathbb{K}(\mathbf{x})$, of a certain symbolic matrix associated with F . More concretely, for $\mathbf{x} \in \mathbb{K}^d$ let

$$h(\mathbf{x}) := (\text{span}\{\mathbf{x}\})^\perp.$$

We prove that $\rho_c(F)$ equals to the generic rank of a symbolic matrix whose entries are linear combinations of the coordinates of \mathbf{x} .

Our main result for the section is the following (note that this is Theorem 1.4 in the introduction).

Theorem 6.1. *Let $u_1, \dots, u_n, v_1, \dots, v_n \in \mathbb{K}^d$ be row vectors. Consider the symbolic matrix $A(\mathbf{x})$, with unknowns $\mathbf{x} = (x_1, \dots, x_d)$, whose i th row is*

$$(v_i^t u_i - u_i^t v_i) \mathbf{x}$$

Then the (generic) rank of $A(\mathbf{x})$ can be computed in polynomial time.

To prove the theorem we use the property established in Theorem 5.2, interpreting the quantity $\rho_1(F)$ as the dimension of the space spanned by

$$F \cap h = \{f \cap h \mid f \in F\},$$

for any hyperplane h in general position with respect to F (see Definition 5.1). Taking $h = h(\mathbf{x})$ we prove, in Lemma 6.2, that the intersection $f \cap h(\mathbf{x})$ is the span of vectors with entries that are linear combinations of the coordinates of \mathbf{x} . We then prove, in Theorem 7.1, that, given a family F , $h(\mathbf{x})$ is in general position with respect to F , for every generic \mathbf{x} (namely, for almost every $\mathbf{x} \in \mathbb{K}^d$). Finally, we use the algorithm for computing ρ_1 from Section 4.

Lemma 6.2. *Let f be an m -dimensional subspace in \mathbb{K}^d and let v_1, \dots, v_m be a basis of f . Let $\mathbf{x} \in \mathbb{K}^d$ and assume that $f \not\subseteq h(\mathbf{x})$. Then $h(\mathbf{x}) \cap f$ is spanned by vectors of the form*

$$w_{ij} := (v_j \cdot \mathbf{x})v_i - (v_i \cdot \mathbf{x})v_j,$$

with $i \neq j$.

Moreover, if $(w \log) \mathbf{x} \cdot v_1 \neq 0$, then the set $\{w_{12}, \dots, w_{1m}\}$ forms a basis of $f \cap h_{\mathbf{x}}$.

Proof. We first observe that $w_{ij} \in f \cap h(\mathbf{x})$. Indeed, by definition, each w_{ij} is a linear combination of basis vectors for f , and thus $w_{ij} \in f$. We also have

$$\begin{aligned} w_{ij} \cdot \mathbf{x} &= ((v_j \cdot \mathbf{x})v_i - (v_i \cdot \mathbf{x})v_j) \cdot \mathbf{x} \\ &= (v_j \cdot \mathbf{x})(v_i \cdot \mathbf{x}) - (v_i \cdot \mathbf{x})(v_j \cdot \mathbf{x}) = 0. \end{aligned}$$

Thus $w_{ij} \in f \cap h(\mathbf{x})$.

We now show that w_{ij} also span $f \cap h(\mathbf{x})$. Indeed, we prove the stronger “moreover” statement.

Let $w \in f \cap h(\mathbf{x})$. Since $w \in f$ we may write $w = \sum_{i=1}^m a_i v_i$. Since $w \in h(\mathbf{x})$, we have $w \cdot \mathbf{x} = 0$ or

$$0 = \sum_{i=1}^m a_i v_i \cdot \mathbf{x}. \quad (22)$$

If $v_i \cdot \mathbf{x} = 0$ for every i , then $f \subseteq h(\mathbf{x})$, contradicting our assumption. We may therefore assume, without loss of generality, that $v_1 \cdot \mathbf{x} \neq 0$. In this case (22) can be rewritten as

$$a_1 = - \sum_{i=2}^m \frac{a_i v_i \cdot \mathbf{x}}{v_1 \cdot \mathbf{x}}.$$

We conclude that

$$\begin{aligned}
w &= \sum_{i=1}^m a_i v_i \\
&= - \left(\sum_{i=2}^m \frac{a_i v_i \cdot \mathbf{x}}{v_1 \cdot \mathbf{x}} \right) v_1 + \sum_{i=2}^m a_i v_i \\
&= \sum_{i=2}^m \frac{-a_i}{v_1 \cdot \mathbf{x}} ((v_i \cdot \mathbf{x})v_1 - (v_1 \cdot \mathbf{x})v_i) \\
&= \sum_{i=2}^m \frac{-a_i}{v_1 \cdot \mathbf{x}} w_{1i}.
\end{aligned}$$

This completes the proof of the lemma. \square

We observe an interesting consequence of Lemma 6.2, asserting that computing $\rho_1(F)$ for a family F can be reduced to computing $\rho_1(G)$, for a certain family G consisting only of planes (two-dimensional subspaces).

Corollary 6.3. *Let $F = \{f_1, \dots, f_n\}$ be a family of subspaces in \mathbb{K}^d and let $\{v_{i1}, \dots, v_{im_i}\}$ be a basis of f_i , for $i = 1, \dots, n$. Consider the family of two-dimensional subspaces*

$$G = \bigcup_{i=1}^n \{g_{ijk} \mid 1 \leq j \neq k \leq m_i\},$$

where $g_{ijk} = \text{span}\{v_{ij}, v_{ik}\}$. Then $\rho_1(F) = \rho_1(G)$.

Proof. It follows easily from Theorem 7.1 that $h(\mathbf{x})$ is in general position with respect to both families F and G , for every generic $\mathbf{x} \in \mathbb{K}^d$. Fixing such $\mathbf{x} \in \mathbb{K}^d$ and applying Lemma 6.2, we see that $\text{span}(F \cap h(\mathbf{x})) = \text{span}(G \cap h(\mathbf{x}))$. By Theorem 5.2 this means that $\rho_1(F) = \rho_1(G)$, as needed. \square

The following lemma is a natural extension of Lemma 6.2 to a similar description of the intersection of a given subspace with a generic one, where the latter is not necessarily of co-dimension 1. If the co-dimension is k , the basis elements of the intersection will be homogeneous polynomials of degree k in the entries of the generic vectors. This connection, together with our algorithm for computing ρ_k , will prove Theorem 1.5 from the introduction.

Lemma 6.4. *Let $k < m \leq d$ be integers. Let f be an m -dimensional subspace in \mathbb{K}^d and let v_1, \dots, v_m be a basis of f . Let $\mathbf{x}_1, \dots, \mathbf{x}_k$ be vectors in \mathbb{K}^d and define the subspace*

$$h := (\text{span}\{\mathbf{x}_1, \dots, \mathbf{x}_k\})^\perp.$$

Assume that $\dim(f \cap h) = m - k$ (this extends the assumption $f \not\subseteq h(\mathbf{x})$ of the lemma above). Let X be the $k \times d$ matrix with \mathbf{x}_i as its i th row. Let V denote the $d \times m$ matrix with v_j as its j th column. Put $M := XV$. So M is a $k \times m$ matrix with (i, j) entry being $\mathbf{x}_i \cdot v_j$. For every $I \subset [m]$ of cardinality k , let M_I denote the $k \times k$ matrix received by restricting to the columns of M with indices in I . Then $f \cap h$ is the span of vectors of the form

$$w_S := \sum_{j=1}^{k+1} (-1)^j \det(M_{I_j}) v_{s_j},$$

where $S = \{s_1 < \dots < s_{k+1}\} \subset [m]$ is of cardinality $k + 1$ and $I_j := S \setminus \{s_j\}$.

Moreover, if (wlog, given our assumption above), assuming that the last k columns of M are linearly independent, $f \cap h$ is spanned by the $m - k$ vectors w_S with S containing the last k columns.

Proof. We first show that $w_S \in f \cap h$, for every $S \subset [m]$ of cardinality $k + 1$. For S fixed, we need to verify that w_S is orthogonal to each of $\mathbf{x}_1, \dots, \mathbf{x}_k$. For every $1 \leq i \leq k$ we have

$$w_S \cdot \mathbf{x}_i = \sum_{j=1}^{k+1} (-1)^j \det(M_{I_j}) v_{s_j} \cdot \mathbf{x}_i.$$

Observe that the right-hand side is exactly the determinant of the matrix received by duplicating the i th row of M . Since the latter matrix is evidently singular, we conclude that $w_S \cdot \mathbf{x}_i = 0$, for every $i = 1, \dots, k$. Thus $w_S \in h$. Clearly, we also have $w_S \in f$. Thus $w_S \in f \cap h$, as needed.

We now turn to prove that the vectors w_S generate $f \cap h$. Indeed we prove the stronger “moreover” statement that already the $m - k$ vectors w_S with S of size $k + 1$ that contain the last k columns span $f \cap h$. Recall that the last k columns of M are independent.

It will be convenient to add one more piece of (slightly informal) notation. Let M' be the matrix extending M with one more (say, 0'th) row, that contains in the j th coordinate the vector v_j . Note that, up to a sign, the determinant of any $k + 1$ minor of M' on columns S is precisely w_S .

Note also that column operations on M' , and replacing w_S by the $k + 1$ minors of the resulting matrix, do not change the span of the vectors w_S . Moreover, note that column operations on the last k columns of M' do not change the vectors w_S , restricting to sets $S \subset I$ of size $k + 1$ that contain the indices of the last k columns. We may therefore assume, by performing such column operations, that the last k columns of M form the $k \times k$ identity matrix.

We will prove the lemma by induction on k . We already know that this statement holds for $k = 1$ (and any m) by Lemma 6.2. Assume it holds for $k - 1$ (and $m - 1$, this is all we need), and we will infer the statement for k . Consider the subspace h' orthogonal to the vectors $\mathbf{x}_1, \dots, \mathbf{x}_{k-1}$, and the subspace f' spanned by the vectors v_1, \dots, v_{m-1} , and form the associated $(k - 1) \times (m - 1)$ matrix, say N . Add to the matrix N the 0'th row to create N' . By induction, we know that the k -minors containing the last $k - 1$ columns of N' are vectors which span the $f' \cap h'$. For $i \in [m - k]$, let w'_i denote the basis vector that corresponds to the columns $\{i, m - k + 1, \dots, m - 1\}$. Note that

$$f \cap h = \text{span}((f' \cap h') \cup \{v_m\}) \cap \{\mathbf{x}_k\}^\perp.$$

Now add to N' a last column for v_m and a last row for x_k to form M' . Fix $i \in [m - k]$, and write $w_i := w_{S_i}$, where $S_i = \{i, m - k + 1, \dots, m\}$. Due to the last k columns of M being the identity matrix, we have

$$w_i = (\mathbf{x}_k \cdot v_i) v_m - w'_i.$$

Moreover, one can check that in fact

$$\begin{aligned} \mathbf{x}_k \cdot v_i &= \mathbf{x}_k \cdot w'_i & \text{and} \\ w'_i &= (\mathbf{x}_k \cdot v_m)w'_i. \end{aligned}$$

That is, $w_i = (\mathbf{x}_k \cdot w'_i)v_m - (\mathbf{x}_k \cdot v_m)w'_i$. Applying Lemma 6.2, we get that the vectors w_i , for $i \in [m - k]$, form a basis for $f \cap h$, as needed. \square

7 Generic vs. General Position

This section completes the cycle of connections, proving that most (namely, generic) hyperplanes, and indeed most subspaces, are in general position (in the Lovász sense of Section 5) with respect to any given family of subspaces. The proof will make use the explicit description we established in the previous section for a basis to the intersection of a family of subspaces and a hyperplane. Thus, computing the ranks of the symbolic matrices in Theorems 1.4 and 1.5 are equivalent to computing the functions ρ_1 and ρ_k respectively, which we can do efficiently by the algorithm of Section 4.

Theorem 7.1. *Let F be a family of subspaces in \mathbb{K}^d , and assume that either $\text{char}(\mathbb{K}) > |F|$ or $\text{char}(\mathbb{K}) = 0$. Then the hyperplane $h(\mathbf{x})$ is in general position (see Definition 5.1) with respect to F for almost every $\mathbf{x} \in \mathbb{K}^d$. More precisely, over finite fields all but $|F|/|\mathbb{K}|$ -fraction of hyperplanes are not in general position, and for infinite fields they have measure zero.*

The proof of this theorem turns out to be more intricate than we imagined. We will give below a linear-algebraic proof that is valid for all fields \mathbb{K} . In the appendix we give an alternative, geometric proof which is valid for the field \mathbb{R} of Real numbers.

Proof. Fix subsets $A, B, C \subset F$. Our goal is to show that for

$$S := \text{span}_{\mathbb{K}}((A \cap h(\mathbf{x})) \cup B) \cap \text{span}_{\mathbb{K}}((A \cap h(\mathbf{x})) \cup C)$$

either $S \not\subset h(\mathbf{x})$ generically, or $S \subset A \cap h(\mathbf{x})$ generically. Indeed, we will prove that one of these alternative holds for every \mathbf{x} , except for those \mathbf{x} that vanish on a certain nontrivial linear equation. Thus, if \mathbb{K} is finite, the fraction of such exceptional values of \mathbf{x} is $1/|\mathbb{K}|$. Since the number of choices of A, B, C is finite, we see that if \mathbb{K} is large enough this probability remains negligible. Being a bit more careful, (see Remark 5.4 at the end of the proof of Theorem 5.2), there are at most $|F|$ applications of the “general position” definition, and so the fraction of “bad” \mathbf{x} is at most $|F|/|\mathbb{K}|$ as stated.

It is easy to see that replacing B by $\text{span}B$ and C by $\text{span}C$ does not affect the subspace S . We may therefore assume that each of the families B, C contains a single subspace of \mathbb{K}^d .

Suppose that $B \cap C \neq \{0\}$, that is, that there exists $v \in B \cap C$, with $v \neq 0$. Clearly, we have $v \in S$ and the linear form $v \cdot \mathbf{x}$ not identically zero. Thus, for almost every \mathbf{x} , S is not contained in $h(\mathbf{x})$ and there is nothing to prove in this case. We may therefore assume that $B \cap C = \{0\}$. In this case, after a change of basis of \mathbb{K}^d , we may assume that

$B = \text{span}\{e_1, \dots, e_k\}$ and $C = \{e_{k+1}, \dots, e_{k+m}\}$, where $1 \leq k < k+m \leq d$ and e_1, \dots, e_d stand for the standard basis vectors in \mathbb{K}^d .

From now on we will regard \mathbf{x} as a vector of variables, and work in the field of fractions $\mathbb{K}(\mathbf{x})$. In particular this makes all subspaces under consideration, $A, B, C, A \cap h(\mathbf{x})$ and of course $S = S(\mathbf{x})$ now subspaces of $\mathbb{K}(\mathbf{x})^d$ (by taking the span of their bases in $\mathbb{K}(\mathbf{x})^d$).

With this, our task becomes proving the following about these subspaces:

Claim 7.2. *Either $S \not\subseteq h(\mathbf{x})$, or $S \subset A \cap h(\mathbf{x})$.*

We will break this task to two. Clearly, it will suffice to prove the claim for any spanning set S' replacing S . So first we will prove that we can take S' to be the affine functions (of \mathbf{x}) in S , and then we will prove the claim for S' .

Lemma 7.3. *S is spanned by its elements which are affine functions of \mathbf{x} .*

Proof of Lemma 7.3. Recall that we showed, in Lemma 6.2, that $\text{span}_{\mathbb{K}}(A \cap h)$ has a basis consisting of elements of the form $(u^t v - v^t u)\mathbf{x}$, for some $u, v \in \mathbb{K}^d$. Write $\{\mathbf{a}_1(\mathbf{x}), \dots, \mathbf{a}_n(\mathbf{x})\}$ for a basis of $\text{span}_{\mathbb{K}}(A \cap h)$ of this form.

Having bases for B, C and $A \cap h(\mathbf{x})$ we can express all elements of S as linear combinations of these bases. Thus, elements in S are described by solutions $\alpha, \alpha' \in \mathbb{K}^n, \beta \in \mathbb{K}^k, \gamma \in \mathbb{K}^m$ to the following system of linear equations.

$$\sum_{i=1}^n \alpha_i \mathbf{a}_i(\mathbf{x}) + \sum_{i=1}^k \beta_i e_i = \sum_{i=1}^n \alpha'_i \mathbf{a}_i(\mathbf{x}) + \sum_{i=1}^m \gamma_i e_{k+i} \quad (23)$$

where $\alpha_i \in \mathbb{K}$ (resp., $\alpha'_i, \beta_i, \gamma_i \in \mathbb{K}$) is the i th entry of α (resp., α', β, γ).

By basic theory of linear algebra, there exists a set of solutions, each of the form

$$w = w(\mathbf{x}) = \sum_{i=1}^n \alpha_i(\mathbf{x}) \mathbf{a}_i(\mathbf{x}) + \sum_{i=1}^k \beta_i(\mathbf{x}) e_i = \sum_{i=1}^n \alpha'_i(\mathbf{x}) \mathbf{a}_i(\mathbf{x}) + \sum_{i=1}^m \gamma_i(\mathbf{x}) e_{k+i}, \quad (24)$$

where $\alpha_i(\mathbf{x}), \alpha'_i(\mathbf{x}), \beta_i(\mathbf{x}), \gamma_i(\mathbf{x})$ are rational functions in the entries of \mathbf{x} , that together span the subspace S . Moreover, these rational functions are of degree at most $|F|$.

We will now strive to find a simpler spanning set S' for S , and then use it to prove Claim 7.2.

The first simplification is realizing (via common denominators) that without loss of generality we can assume that all $\alpha_i(\mathbf{x}), \alpha'_i(\mathbf{x}), \beta_i(\mathbf{x}), \gamma_i(\mathbf{x})$ are in fact *polynomials* in the entries of \mathbf{x} . These elements of S span the rest, after dividing by some fixed polynomial.

The next simplification (separating out homogeneous terms) shows that without loss of generality we can take all the polynomials in each of $\alpha, \alpha', \beta, \gamma$ to be homogeneous of the same degree, which we may respectively call $\deg(\alpha), \deg(\alpha'), \deg(\beta), \deg(\gamma)$. These homogeneous solutions certainly span S , and now we refine their structure further.

Indeed, inspecting the system of equations we know more: since each entry of $\mathbf{a}_i(\mathbf{x})$, for every i is of degree one, we know that for some fixed integer $r \geq 0$, they must satisfy $\deg(\alpha) = \deg(\alpha') = r$ and $\deg(\beta) = \deg(\gamma) = r + 1$. We use this to stratify solutions w by degree, and say that the associated w has degree r . Let S_r be all solutions of degree r (note

that each S_r is a subspace over \mathbb{K} , though we will not use this fact). We call solutions w of degree 0 *linear*. Our main simplification will come from showing that linear elements S_0 span S , which in this notation is a restatement of the lemma we are proving.

Claim 7.4. $\text{span}S_0 = S$

We will prove this claim by induction on r , using our stratifications S_r of members of S . It is clearly true for $r = 0$. So assume S_0 spans S_r , and we need to prove that S_0 spans S_{r+1} . By induction, it suffices to prove that S_r spans S_{r+1} . The plan for this will be as follows. We will assume we have some $w \in S_{r+1}$. We will take all partial derivatives of its constituent polynomials with respect to each variable x_t , $t \in [d]$. From each of these we will generate an element $w_t \in S_r$, as the degree decreased by 1. Finally, we will show that w is a linear combination, indeed a very simple one, of the form : $(r + 1)w = \sum_{t=1}^d x_t w_t$. We now elaborate.

Fix $t \in [d]$. Let us take a derivative with respect to the variable x_t of \mathbf{x} , of both sides of the identity (24). We get

$$\begin{aligned} & \sum_{i=1}^n \left(\frac{\partial \alpha_i(\mathbf{x})}{\partial x_t} \mathbf{a}_i(\mathbf{x}) + \alpha_i(\mathbf{x}) \frac{\partial \mathbf{a}_i(\mathbf{x})}{\partial x_t} \right) + \sum_{i=1}^k \frac{\partial \beta_i(\mathbf{x})}{\partial x_t} e_i = \\ & \sum_{i=1}^n \left(\frac{\partial \alpha'_i(\mathbf{x})}{\partial x_t} \mathbf{a}_i(\mathbf{x}) + \alpha'_i(\mathbf{x}) \frac{\partial \mathbf{a}_i(\mathbf{x})}{\partial x_t} \right) + \sum_{i=1}^m \frac{\partial \gamma_i(\mathbf{x})}{\partial x_t} e_{k+i} \end{aligned}$$

To define w_t we first define $\alpha(t), \alpha'(t), \beta(t), \gamma(t)$ by appropriately collecting homogeneous terms, and making sure that $\alpha(t), \alpha'(t) \in A \cap h$ are of degree r , and that $\beta(t) \in B$ and $\gamma(t) \in C$ are of degree $r + 1$:

- $\alpha(t)_i = \frac{\partial \alpha_i(\mathbf{x})}{\partial x_t}$
- $\alpha'(t)_i = \frac{\partial \alpha'_i(\mathbf{x})}{\partial x_t}$,
- For $i \in [k]$, $\beta(t)_i(\mathbf{x})$ is

$$\left[\sum_{s=1}^n (\alpha_s(\mathbf{x}) - \alpha'_s(\mathbf{x})) \frac{\partial \mathbf{a}_s(\mathbf{x})}{\partial x_t} \right]_i + \frac{\partial \beta_i(\mathbf{x})}{\partial x_t}$$

- For $i \in [m]$, $\gamma(t)_i(\mathbf{x})$ is

$$\left[\sum_{s=1}^n (\alpha'_s(\mathbf{x}) - \alpha_s(\mathbf{x})) \frac{\partial \mathbf{a}_s(\mathbf{x})}{\partial x_t} \right]_{k+i} + \frac{\partial \gamma_i(\mathbf{x})}{\partial x_t};$$

here we used $[v]_j$ to denote the j th entry of a vector v . Now we can formally define $w_t \in S_r$ as follows. We first observe that

$$\sum_{i=1}^n \alpha(t)_i(\mathbf{x}) \mathbf{a}_i(\mathbf{x}) + \sum_{i=1}^k \beta(t)_i(\mathbf{x}) e_i = \sum_{i=1}^n \alpha'(t)_i(\mathbf{x}) \mathbf{a}_i(\mathbf{x}) + \sum_{i=1}^m \gamma(t)_i(\mathbf{x}) e_{k+i}. \quad (25)$$

Indeed, note that (24), restricted to the j th component of the equation, implies that for every, $k + m < j \leq n$, we have

$$\left[\sum_{i=1}^n (\alpha'_i(\mathbf{x}) - \alpha_i(\mathbf{x})) \mathbf{a}_i(\mathbf{x}) \right]_j = 0.$$

From this it is straightforward to verify that the identity (25) indeed holds. Thus, letting

$$w_t := \sum_{i=1}^n \alpha(t)_i(\mathbf{x}) \mathbf{a}_i(\mathbf{x}) + \sum_{i=1}^k \beta(t)_i(\mathbf{x}) e_i,$$

for each t , the identity (25) implies that w_t is in S . Moreover, by our definition, w_t is of degree $r - 1$.

It remains to prove that w is spanned by the vectors w_t . For this, one basic fact we will need is that if $p(\mathbf{x})$ is any homogeneous polynomial of degree m , it satisfies

$$\sum_t x_t \cdot \frac{\partial p(\mathbf{x})}{\partial x_t} = mp(\mathbf{x}).$$

The second fact we will need follows from identity (24), when restricted to the j th component of the equation. For every $j \in [k]$,

$$\left[\sum_{i=1}^n (\alpha'_i(\mathbf{x}) - \alpha_i(\mathbf{x})) \mathbf{a}_i(\mathbf{x}) \right]_j = \beta_j.$$

Combining these two properties, we get

- $\sum_t x_t \alpha(t) = r\alpha$
- $\sum_t x_t \beta(t) = r\beta$

and this implies that

$$rw = \sum_t x_t w_t.$$

Note that $r \neq 0$; indeed, for \mathbb{K} with non-zero characteristic, we have $r < \text{char}(\mathbb{K})$. Thus the vectors w_t span w . This completes the induction step, and hence the proof of Lemma 7.3. \square

To complete the proof of the theorem we now prove

Lemma 7.5. *Either S_0 is not contained in $h(\mathbf{x})$, or it is contained in $A \cap h(\mathbf{x})$.*

As the elements in S_0 are affine functions of \mathbf{x} , a violation of the first possibility will imply that \mathbf{x} satisfy a linear equation, so the fraction of such vectors is at most $1/|\mathbb{K}|$ as requested.

Proof of Lemma 7.5. We first introduce some notation. Let $v(\mathbf{x})$ be a vector in $\mathbb{K}(\mathbf{x})^d$, such that each entry of $v(\mathbf{x})$ is some linear combination of x_1, \dots, x_d , the coordinates of \mathbf{x} . Then $v(\mathbf{x})$ can be represented by a matrix $M \in \text{Mat}_{d \times d}(\mathbb{K})$, with constant entries, such that

$M\mathbf{x} = v(\mathbf{x})$. Note that if M is skew-symmetric, this means that $(M\mathbf{x}) \cdot \mathbf{x} = (M^t\mathbf{x}) \cdot \mathbf{x} = -(M\mathbf{x}) \cdot \mathbf{x}$ or $2(M\mathbf{x}) \cdot \mathbf{x} = 0$, which means that $(M\mathbf{x}) \cdot \mathbf{x} = 0$, unless the characteristic of the field is 2. Conversely, if $M\mathbf{x} \cdot \mathbf{x} = 0$ for every $\mathbf{x} \in \mathbb{K}^d$ and so $M\mathbf{x} \cdot \mathbf{x}$ is the zero polynomial (in d variables), which implies that M is skew-symmetric.

Consider k such matrices M_1, \dots, M_k , representing vectors $v_1(\mathbf{x}), \dots, v_k(\mathbf{x})$, respectively. Then a linear combination $\sum_{i=1}^k \alpha_i M_i$ is a matrix that corresponds to a vector which is a linear combination of $v_1(\mathbf{x}), \dots, v_k(\mathbf{x})$, namely, $v(x) = \sum_i \alpha_i v_i(\mathbf{x})$. Thus $v(\mathbf{x})$ lies in the span of the vectors $v_i(\mathbf{x})$.

Assume first that $k + m = d$. We regard a $(k + m) \times (k + m)$ matrix M as a block matrix with $TL(M)$ (resp., $TR(M)$, $BL(M)$, $BR(M)$) denoting the top-left (resp., top-right, bottom-left, bottom-right) blocks. More precisely, $TL(M)$ (resp., $TR(M)$, $BL(M)$, $BR(M)$) stands for the submatrix induced by taking the first k (resp., first k , last m , last m) rows and first k (resp., last m , first k , last m) columns of M .

With some abuse of notation, we write $M \in Y$, for a subspace Y of $\mathbb{K}(\mathbf{x})^d$, if $M\mathbf{x} \in Y$. Recall that M is in h if and only if M is skew-symmetric. In particular, $TR(M) = -BL(M)^t$, for every $M \in \text{span}(A \cap h)$. Assume that for some $M \in \text{span}(A \cap h)$, we have $TR(M) \neq 0$ (and thus also $BL(M) \neq 0$). We claim that in this case there exists a matrix $\widetilde{M} \in S \setminus h$. To see this it is sufficient to show that there exist matrices $b \in B$ and $c \in C$ such that $M + b = c$ which is not skew-symmetric (and therefore not in h). Indeed, let b be defined by $TL(b) = -TL(M)$, $TR(b) = -TR(M)$, and $BL(b) = BR(b) = 0$. We define the matrix c by $TL(c) = TR(c) = 0$, $BL(c) = BL(M)$, $BR(c) = BR(M)$. Clearly, $b \in B$, $c \in C$ and $M + b = c$. If c is skew-symmetric, then we must have $BL(c) = BL(M) = 0$, contradicting our assumption on M . Thus $c = M + b$ is in $A \cap h$ but not in S . We conclude that in this case the general position requirement holds generically.

Assume next that for every $M \in \text{span}(A \cap h)$, we have $TR(M) = BL(M) = 0$. Recall that $\text{span}(A \cap h)$ is spanned by matrices of the form $v^t u - u^t v$ for some $u, v \in \mathbb{K}^d$. Assume that $TR(v^t u - u^t v) = BL(v^t u - u^t v) = 0$ for such a matrix. We claim that in this case at least one of $TL(v^t u - u^t v)$ or $BR(v^t u - u^t v)$ is the zero matrix. Indeed, put $M = v^t u - u^t v$, and assume that $TL(M) \neq 0$. Then for some $1 \leq i_0 \neq j_0 \leq k$ we have $u_{i_0} v_{j_0} \neq u_{j_0} v_{i_0}$. In particular, not both $u_{i_0} v_{j_0}$ and $u_{j_0} v_{i_0}$ are zero. Assume, without loss of generality, that $u_{i_0} v_{j_0} \neq 0$. That is, $u_{i_0}, v_{j_0} \neq 0$. Suppose that $u_\ell = 0$ for every $\ell > k$. In this case it is clear that $BR(M) = 0$ and the claim is proved. Therefore, we may assume that for some $\ell > k$ we have $u_\ell \neq 0$. Since we $BL(M) = 0$, we have in particular $u_\ell v_j = u_j v_\ell$, for every $j = 1, \dots, k$. In particular, $u_\ell v_{j_0} = u_{j_0} v_\ell$. Note that since $v_{j_0} \neq 0$ and $u_\ell \neq 0$, we must have that also $v_\ell, u_{j_0} \neq 0$. Thus, we get $\frac{v_{i_0}}{u_{i_0}} = \frac{v_\ell}{u_\ell}$ and $\frac{v_{j_0}}{u_{j_0}} = \frac{v_\ell}{u_\ell}$. Combining these equalities, we get that $u_{i_0} v_{j_0} = u_{j_0} v_{i_0}$, contradicting our assumption. This proves the claim.

This implies that $\text{span}(A \cap h)$ is a direct sum $U \oplus V$ of matrices with entries supported only on $TL(M)$ for $M \in U$ and matrices supported by $BR(M)$ for $M \in V$.

Now let $w \in S$. By the definition of S , w can be written as $w = a + b = a' + c$ for some $a, a' \in \text{span}(A \cap h)$, $b \in B$, $c \in C$. Write $a = a_U + a_V$, where $a_U \in U$ and $a_V \in V$. Similarly, write $a' = a'_U + a'_V$. Then $a_U + a_V + b = a'_U + a'_V + c$, or $a_U - a'_U + b = a'_V - a_V + c$. But then, we must have $b = a'_U - a_U$ and $c = a_V - a'_V$, which in particular implies that $b, c \in \text{span}(A \cap h)$.

Since $a_U - a'_U \in U$ and $a'_V - a_V \in V$, this implies that, without loss of generality, we

may assume $a \in U$ and $a' \in V$. Thus also $w = a + b = a' + c \in \text{span}(A \cap h)$. We conclude that $w \in \text{span}(A \cap h)$ for every $w \in S$. Thus the general position requirement holds in this case.

We now prove the remaining case where $k + m < d$, by reducing it to the case $k + m = d$ just discussed. Write $k + m = d - z$, for some $z > 0$. Repeat the above argument ignoring the last z rows and last z columns of every matrix used along the proof. Note that for $a \in A \cap h$, a is skew-symmetric, and adding a matrix $b \in B$ or $c \in C$ will result with a matrix which is either in h or not in h , independent of the last z rows and columns of a . Indeed, for $b \in B$ and $c \in C$ these rows and columns are zero, and therefore they cannot affect the skew-symmetry of $a + b$ or $a' + c$. \square

This completes the proof of Theorem 7.1. \square

Having established the connection between genericity and general position, we can now complete the proof of Theorem 6.1.

Proof of Theorem 6.1. Consider the family of subspaces $F = \{f_1, \dots, f_n\}$, where $f_i := \text{span}\{u_i, v_i\}$, for each $i = 1, \dots, n$. Let $\mathbf{x} = (x_1, \dots, x_d)$ and consider $h := (\text{span}\{\mathbf{x}\})^\perp$. In view of Lemma 6.2, we have

$$\text{rank}A(\mathbf{x}) = d(\{f \cap h \mid f \in F\}).$$

On the other hand, by Theorem 5.2, we have $d(\{f \cap h \mid f \in F\}) = \rho_1(F)$. Thus there exists a deterministic strongly-polynomial time algorithm to compute $\text{rank}A(\mathbf{x})$. \square

We note that in the exact same way, our ability to efficiently compute ρ_k for every integer k by Theorem 1.6, and the characterization above, completes the proof of Theorem 1.5 from the introduction.

Acknowledgements We would like to thank Ze'ev Dvir for many illuminating discussions. We thank Amir Shpilka and Roy Meshulam for useful comments on an earlier version of the paper. We also thank Jan Vondrak for telling us about Dilworth truncation.

References

- [1] M. Agrawal, C Saha, R. Saptharishi, and N. Saxena, Jacobian hits circuits: Hitting sets, lower bounds for depth-d occur-k formulas and depth-3 transcendence degree-k circuits, *SIAM J. Comput.* 45.4 (2016), 1533–1562.
- [2] L. Asimow and B. Roth, The rigidity of graphs, *Trans. Amer. Math. Soc.* 245 (1978), 279–289.
- [3] P. M. Brooksbank, and E. M. Luks, Testing isomorphism of modules, *J. Algebra* 320.11 (2008), 4020–4029.

- [4] A. Chistov, G. Ivanyos, and M. Karpinski, Polynomial time algorithms for modules over finite dimensional algebras, *Proceedings of the 1997 ACM International Symposium on Symbolic and Algebraic Computation (ISSAC)* (1997), 68–74.
- [5] J. Edmonds, Systems of distinct representatives and linear algebra, *J. Res. Natl. Bur. Stand.* 71 (1967), 241–245.
- [6] S. Fenner, R. Gurjar, and T. Thierauf, Bipartite perfect matching is in quasi-NC. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing (STOC)* (2016), 754–763.
- [7] A. Frank and É. Tardos, Generalized polymatroids and submodular flows, *Mathematical Programming* 42 (1988), 489–563.
- [8] A. Garg, L. Gurvits, R. Oliveira, and A. Wigderson, Operator scaling: theory and applications, in [arXiv:1511.03730v3](https://arxiv.org/abs/1511.03730v3).
- [9] J. F. Geelen, Maximum rank matrix completion, *Linear Algebra Appl.* 288 (1999), 211–217.
- [10] B. Hendrickson, Conditions for unique graph realizations, *SIAM J. Comput.* 21 (1992), 65–84.
- [11] G. Ivanyos, M. Karpinski, and N. Saxena, Deterministic polynomial time algorithms for matrix completion problems, *SIAM J. Comput.* 39.8 (2010), 3736–3751.
- [12] T. Jordan, Combinatorial rigidity: Graphs and matroids in the theory of rigid frameworks, *MSJ Memoirs* 34 (2016), 33–112.
- [13] V. Kabanets and R. Impagliazzo, Derandomizing polynomial identity tests means proving circuit lower bounds, *Computational Complexity* 13 (2004), 1–46.
- [14] G. Laman, On graphs and rigidity of plane skeletal structures, *J. Engrg. Math.* 4 (1970), 333–338.
- [15] H. Liu, and K. Regan, Improved construction for universality of determinant and permanent, *Inf. Process. Lett.* 100.6 (2006), 233–237.
- [16] L. Lovász, On determinants, matchings, and random algorithms, in *International Symposium on Fundamentals of Computation Theory*, (1979), 565–574.
- [17] L. Lovász, Matroid matching and some applications, *J. Combin. Theory Ser. B* 28.2 (1980), 208–236.
- [18] L. Lovász and M. D. Plummer, *Matching Theory*, American Mathematical Soc., Providence, Rhode Island, 2009.
- [19] L. Lovász, Flats in matroids and geometric graphs, in *Combinatorial Surveys*, Proc. 6th British Combinatorial Conf., P. Cameron, Ed., Academic Press, New York, 1977, pp. 45–89.
- [20] L. Lovász and Y. Yemini, On generic rigidity in the plane, *SIAM J. Alg. Disc. Math.* 3 (1982), 91–98.

- [21] R. Meshulam, On the maximal rank in a subspace of matrices, *Q. J. Math.* 36.2 (1985), 225–229.
- [22] N. Saxena, Progress on polynomial identity testing-II, *Perspectives in Computational Complexity*, Springer International Publishing, 2014, pp. 131–146.
- [23] A. Schrijver, A combinatorial algorithm minimizing submodular functions in strongly polynomial time, in *J. Combinat. Theory, Ser. B*, 80.2 (2000), 346–355.
- [24] A. Shpilka and A. Yehudayoff, Arithmetic circuits: A survey of recent results and open questions, *Found. Trends Theor. Comput. Sci.* 5 (2010), 207–388.
- [25] M. Sitharam, A. St. John, and J. Sidman (editors), *Handbook of Geometric Constraint Systems Principles*, CRC press, Taylor& Francis group (2019).
- [26] S. Tanigawa, Generic Rigidity Matroids with Dilworth Truncations, *SIAM J. Discrete Math.*, 26.3 (2012), 1412–1439.
- [27] S. Tanigawa, Matroids of gain graphs in applied discrete geometry, *Trans. Amer. Math. Soc.* 367 (2015), 8597–8641.
- [28] L. Valiant, The complexity of computing the permanent, *Theor. Comput. Sci.* 8 (1979), 189–201.

Appendix: Proof of Theorem 7.1 over \mathbb{R}

Here we provide an alternative proof of Theorem 7.1 which works over the field of Real numbers. One advantage of working over \mathbb{R} is that we have the notions of a manifold and of the dimension of a manifold available. In the proof below, we use the fact that the set of linear subspaces of \mathbb{R}^d can be viewed as a manifold. Then, to show that a certain set has measure zero, it is sufficient to show that this set has lower dimension. This allows us to obtain a more straightforward proof for the case $\mathbb{K} = \mathbb{R}$.

Proof over \mathbb{R} : We first prove that property (i) in Definition 5.1 is a generic property. Fix $A \subset F$ and put $g = \text{span}(A)$. For $\mathbf{x} \in \mathbb{S}^{d-1}$ with $g \subset h(\mathbf{x})$, we have $\mathbf{x} \in \mathbb{S}^{d-1} \cap g^\perp$. If $d(g) \geq 1$, this means that \mathbf{x} lies in a lower-dimensional sphere, which is a measure-zero subset of \mathbb{S}^{d-1} . Since F is finite (and so the number of different sub-families A is finite), we conclude that for every $\mathbf{x} \in \mathbb{S}^{d-1}$, excluding a finite union of certain lower-dimensional sub-spheres of \mathbb{S}^{d-1} , $h(\mathbf{x})$ satisfies property (i) in Definition 5.1.

We now prove that property (ii) in Definition 5.1 is a generic property. Fix some subfamilies $A, B, C \subset F$. We first handle certain degenerate cases. Note that if

$$\text{span}(A \cap h(\mathbf{x})) = \text{span}((A \cap h(\mathbf{x})) \cup B) \cap \text{span}((A \cap h(\mathbf{x})) \cup C), \quad (26)$$

for some $\mathbf{x} \in \mathbb{S}^{d-1}$, then $h(\mathbf{x})$ clearly satisfies property (ii). Using Lemma 6.2, condition (26) defines an algebraic subvariety of \mathbb{S}^{d-1} . In particular, (26) either holds for every $\mathbf{x} \in \mathbb{S}^{d-1}$ or holds only for \mathbf{x} taken from a subset of \mathbb{S}^{d-1} of measure zero. In the former case this means that, with respect to the subfamilies A, B, C , property (ii) in Definition 5.1 holds for

$h(\mathbf{x})$ for every $\mathbf{x} \in \mathbb{S}^{d-1}$ and there is nothing to prove. Therefore we can assume that we are in the complementary case. Namely, we assume that for almost every $\mathbf{x} \in \mathbb{S}^{d-1}$ we have

$$\text{span}(A \cap h(\mathbf{x})) \subsetneq \text{span}((A \cap h(\mathbf{x})) \cup B) \cap \text{span}((A \cap h(\mathbf{x})) \cup C). \quad (27)$$

Our next step is to identify the set of subspaces g of the form $g = \text{span}(A \cap h(\mathbf{x}))$, for some $\mathbf{x} \in \mathbb{S}^{d-1}$, and determine its dimension as a subset of the Grassmannian.

We need the following observation. Let

$$r := \max_{\mathbf{x} \in \mathbb{S}^{d-1}} d(A \cap h(\mathbf{x}))$$

We claim that $d(A \cap h(\mathbf{x})) = r$, for almost every $\mathbf{x} \in \mathbb{S}^{d-1}$. Indeed, by Lemma 6.2, one can write a basis for $\text{span}(A \cap h(\mathbf{x}))$ with entries that are linear combinations in the coordinates of \mathbf{x} . In particular, $d(A \cap h(\mathbf{x}))$ can be expressed as the rank of a certain symbolic matrix, with entries depending linearly in the coordinates of \mathbf{x} . This implies that $d(A \cap h(\mathbf{x})) = r$ for every $\mathbf{x} \in \mathbb{S}^{d-1}$, excluding some subset of \mathbb{S}^{d-1} of measure zero, which proves our claim. (Here we used the fact that the maximal rank of a given symbolic matrix is the same as the *generic* rank of the matrix.)

Let S_0 denote the subset of $\mathbf{x} \in \mathbb{S}^{d-1}$ such that either $d(A \cap h(\mathbf{x})) < r$ or (26) holds for $h(\mathbf{x})$. As argued above $S_0 \subset \mathbb{S}^{d-1}$ has measure zero. Let $\text{Gr}(r, d)$ denote the Grassmannian of r -dimensional subspaces of \mathbb{R}^d , regarded as an affine variety. We define a map $\phi : \mathbb{S}^{d-1} \setminus S_0 \rightarrow \text{Gr}(r, d)$ by

$$\mathbf{x} \mapsto \text{span}(A \cap h(\mathbf{x})).$$

We claim that the image of ϕ is r -dimensional. Indeed, let $g \in \text{Im}(\phi)$ and let $\mathbf{x} \in \phi^{-1}(g)$. By definition of the domain of ϕ , we have $\mathbf{x} \notin S_0$ and thus $d(g) = r$. This means g has maximal dimension. Observe that this guarantees that, for every $\mathbf{x} \in g^\perp$, we have $\text{span}(A \cap h(\mathbf{x})) = g$. (Indeed, $\mathbf{x} \in g^\perp$ certainly implies that $g \subset \text{span}(A \cap h(\mathbf{x}))$ and since $d(A \cap h(\mathbf{x})) \leq r = d(g)$, we have equality.) That is, $\phi^{-1}(g) = (\mathbb{S}^{d-1} \setminus S_0) \cap g^\perp$ and, in particular,

$$\dim(\phi^{-1}(g)) = d - 1 - r$$

(dimension here is as a manifold). We conclude that

$$\dim \text{Im}(\phi) = d - 1 - (d - 1 - r) = r, \quad (28)$$

as claimed.

Next, define

$$S'_1 = \{\mathbf{x} \in \mathbb{S}^{d-1} \mid \text{span}((A \cap h(\mathbf{x})) \cup B) \cap \text{span}((A \cap h(\mathbf{x})) \cup C) \subset h(\mathbf{x})\}.$$

Our goal is to show that S'_1 has measure zero, as a subset of the sphere. For this, it suffices to show that $S_1 := S'_1 \setminus S_0$ has measure zero (since S_0 is of measure zero). Consider the restriction of ϕ to S_1 . Let $g \in \text{Im}(\phi|_{S_1})$ and let $\mathbf{x} \in \phi|_{S_1}^{-1}(g)$. Set

$$g' := \text{span}((A \cap h(\mathbf{x})) \cup B) \cap \text{span}((A \cap h(\mathbf{x})) \cup C).$$

Since $\mathbf{x} \notin S_0$, we have (27) which means

$$d(g') \geq r + 1.$$

Since we assume also that $\mathbf{x} \in S'_1$, we have $\mathbf{x} \in (g')^\perp$. So

$$\dim(\phi|_{S'_1}^{-1}(g)) \leq d(g')^\perp \leq d - 1 - (r + 1) = d - r - 2. \quad (29)$$

Clearly we also have $\text{Im}(\phi|_{S_1}) \subset \text{Im}(\phi)$, and thus, using (28),

$$\dim(\text{Im}(\phi|_{S_1})) \leq r. \quad (30)$$

Combining (29) and (30), we get that

$$\dim S_1 = \dim(\text{Im}(\phi|_{S_1})) + \dim(\phi|_{S'_1}^{-1}(g)) \leq d - 2.$$

This completes the proof of the lemma. □