

Variety Evasive Subspace Families

Zeyu Guo*

Abstract

We introduce the problem of constructing explicit *variety evasive subspace families*. Given a family \mathcal{F} of subvarieties of a projective or affine space, a collection \mathcal{H} of projective or affine k -subspaces is (\mathcal{F}, ϵ) -*evasive* if for every $\mathcal{V} \in \mathcal{F}$, all but at most ϵ -fraction of $W \in \mathcal{H}$ intersect every irreducible component of \mathcal{V} with (at most) the expected dimension. The problem of constructing such an explicit subspace family generalizes both deterministic black-box polynomial identity testing (PIT) and the problem of constructing explicit (weak) lossless rank condensers.

Using Chow forms, we construct explicit k -subspace families of polynomial size that are evasive for all varieties of bounded degree in a projective or affine n -space. As one application, we obtain a complete derandomization of Noether's normalization lemma for varieties of bounded degree in a projective or affine n -space. In another application, we obtain a simple polynomial-time black-box PIT algorithm for depth-4 arithmetic circuits with bounded top fan-in and bottom fan-in that are not in the Sylvester–Gallai configuration, improving and simplifying a result of Gupta (ECCC TR 14-130).

As a complement of our explicit construction, we prove a lower bound for the size of k -subspace families that are evasive for degree- d varieties in a projective n -space. When $n - k = \Omega(n)$, the lower bound is superpolynomial unless d is bounded. The proof uses a dimension-counting argument on Chow varieties that parametrize projective subvarieties.

1 Introduction

Polynomial identity testing (PIT) is a fundamental problem in the areas of derandomization and algebraic complexity theory. The problem asks whether a multivariate polynomial, computed by an arithmetic circuit, formula, or other algebraic computational models, is identically zero. For example, the polynomial $(X + Y)(X - Y) - X^2 - Y^2$ is identically zero while $(X + Y)^2 - X^2$ is not.

It is easy to solve PIT in randomized polynomial time, as we may simply evaluate the input polynomial at a random point and check if the evaluation is zero. On the other hand, finding a deterministic polynomial-time PIT algorithm for general arithmetic circuits is a long-standing open problem. Such algorithms are known for some special cases, and we refer the readers to the surveys [Sax09, Sax13, SY10] for details.

Black-box PIT algorithms are a special kind of PIT algorithm. A (deterministic) black-box PIT algorithm tests if a polynomial in a family \mathcal{F} is zero by constructing a *hitting set* for \mathcal{F} , which is a finite collection \mathcal{H} of evaluation points with the following property: for any nonzero $Q \in \mathcal{F}$, there exists $p \in \mathcal{H}$ such that the evaluation of Q at p is nonzero. After constructing such a hitting set \mathcal{H} , the algorithm simply checks if the evaluation of the given polynomial at every point in \mathcal{H} is zero. The problem of designing a deterministic black-box PIT algorithm is thus equivalent to constructing a hitting set. To make the algorithm efficient, such a hitting set should be small and efficiently computable.

From a geometric perspective, an n -variate nonzero polynomial Q over an algebraically-closed field \mathbb{F} defines a *hypersurface* $\mathcal{V}(Q) := \{\alpha \in \mathbb{F}^n : Q(\alpha) = 0\}$ of \mathbb{F}^n . A hitting set \mathcal{H} for \mathcal{F} has the property that for every nonzero $Q \in \mathcal{F}$, there exists a point $p \in \mathcal{H}$ that is disjoint from the hypersurface $\mathcal{V}(Q)$, or we say p

*Department of Computer Science, University of Haifa, Haifa 3498838, Israel. Email: zguotcs@gmail.com

evades $\mathcal{V}(Q)$. It is natural to consider the generalization of this property to higher dimensions/codimensions. Namely, we want to construct a finite collection \mathcal{H} of *affine k -subspaces* (i.e. affine subspaces of dimension k) such that for every *variety* $\mathcal{V} \subseteq \mathbb{F}^n$ (i.e., solution set of a set of polynomial equations) from a certain family, some (or most) $W \in \mathcal{H}$ *evade* \mathcal{V} , in the sense that the dimension of the intersection $\mathcal{V} \cap W$ is bounded by the expected dimension achieved by W in general position. A similar property can be defined for projective k -spaces, to be defined below. We call such a collection \mathcal{H} of projective or affine k -subspaces a *variety evasive subspace family*. The formal definition is given below.

1.1 Variety Evasive Subspace Families

Let \mathbb{F} be an algebraically closed field. An *affine n -space* \mathbb{A}^n , as a set, is simply defined to be the vector space \mathbb{F}^n . We also need the notion of a *projective n -space*, denoted by \mathbb{P}^n , which is (intuitively) the set of lines passing through the origin $\mathbf{0}$ of \mathbb{A}^{n+1} . Formally, it is defined to be the quotient set $(\mathbb{A}^{n+1} \setminus \{\mathbf{0}\}) / \sim$, where \sim is the equivalence relation defined by scaling, i.e., $u \sim v$ if $u = cv$ for some nonzero scalar $c \in \mathbb{F}$.

An (*affine*) *subvariety* $\mathcal{V} \subseteq \mathbb{A}^n$ is the set of common zeros of a set of n -variate polynomials over \mathbb{F} . Similarly, a (*projective*) *subvariety* $\mathcal{V} \subseteq \mathbb{P}^n$ is the set of common zeros of a set of *homogeneous* $(n+1)$ -variate polynomials over \mathbb{F} , where we represent each element of \mathbb{P}^n as an $(n+1)$ -tuple in \mathbb{A}^{n+1} . In this paper, a *variety* refers to a subvariety of a projective or affine space, and is said to be *irreducible* if it cannot be written as a union of finitely many proper subvarieties.¹

The *dimension* of a variety \mathcal{V} , denoted by $\dim(\mathcal{V})$, is intuitively the “degree of freedom” of picking a point in the variety. See Subsection 2.3 for its formal definition. For a linear subspace $\mathcal{V} \subseteq \mathbb{A}^n$, the linear-algebraic dimension of \mathcal{V} is the same as its dimension as a variety.

For two irreducible subvarieties \mathcal{V}_1 and \mathcal{V}_2 of \mathbb{P}^n or \mathbb{A}^n in *general position*, we expect the dimension of $\mathcal{V}_1 \cap \mathcal{V}_2$ to be $\dim(\mathcal{V}_1) + \dim(\mathcal{V}_2) - n$ (unless $\dim(\mathcal{V}_1) + \dim(\mathcal{V}_2) < n$, in which case we expect $\mathcal{V}_1 \cap \mathcal{V}_2 = \emptyset$). The following definition captures the condition that $\dim(\mathcal{V}_1 \cap \mathcal{V}_2)$ is bounded by the expected dimension.

Definition 1.1 (Evading). *Let \mathcal{V}_1 and \mathcal{V}_2 be irreducible subvarieties of \mathbb{P}^n or \mathbb{A}^n . We say \mathcal{V}_1 evades \mathcal{V}_2 if*

$$\dim(\mathcal{V}_1 \cap \mathcal{V}_2) \leq \dim(\mathcal{V}_1) + \dim(\mathcal{V}_2) - n,$$

where the dimension of an empty set is assumed to be $-\infty$. In particular, if $\dim(\mathcal{V}_1) + \dim(\mathcal{V}_2) < n$, then \mathcal{V}_1 evades \mathcal{V}_2 iff $\mathcal{V}_1 \cap \mathcal{V}_2 = \emptyset$.

More generally, suppose \mathcal{V}_1 is irreducible but \mathcal{V}_2 is possibly reducible. We say \mathcal{V}_1 evades \mathcal{V}_2 if it evades every irreducible component of \mathcal{V}_2 .

Next, we define *subspace families* and *variety evasive subspace families*.

Definition 1.2 (Subspace family). *For $0 \leq k \leq n$, a finite collection² of k -subspaces of \mathbb{P}^n is called a (projective) k -subspace family on \mathbb{P}^n . Similarly, a finite collection of affine k -subspaces of \mathbb{A}^n is called an affine k -subspace family on \mathbb{A}^n .*

Definition 1.3 (Variety evasive subspace family). *Let \mathcal{F} be a family of subvarieties of \mathbb{P}^n (resp. \mathbb{A}^n). Let \mathcal{H} be a k -subspace family on \mathbb{P}^n (resp. affine k -subspace family on \mathbb{A}^n) where $0 \leq k \leq n$. Then:*

- We say \mathcal{H} is \mathcal{F} -evasive if for every $\mathcal{V} \in \mathcal{F}$, there exists $W \in \mathcal{H}$ that evades \mathcal{V} .
- We say \mathcal{H} is (\mathcal{F}, ϵ) -evasive if for every $\mathcal{V} \in \mathcal{F}$, a random element $W \in \mathcal{H}$ evades \mathcal{V} with probability at least $1 - \epsilon$.

¹Varieties in this paper are not necessarily irreducible and are often called *algebraic sets* in literature.

²In this paper, a *collection* is a multiset, i.e., its elements are allowed to appear more than once.

Connection with hitting sets. Definition 1.3 naturally generalizes the notions of hitting sets in the context of PIT. For example, a collection of points in \mathbb{P}^n is a hitting set for a family \mathcal{F} of homogeneous polynomials in $\mathbb{F}[X_1, \dots, X_{n+1}]$ iff it is an \mathcal{F}' -evasive 0-subspace family, where $\mathcal{F}' = \{\mathcal{V}(P) : P \in \mathcal{F}\}$ is the family of hypersurfaces defined by the polynomials in \mathcal{F} . In other words, hitting sets may be viewed as 0-subspace families that are evasive for varieties of codimension one.

Connection with lossless rank condensers. Other than the case of codimension one, we may also consider the special case of degree one, and this leads to another important family of pseudorandom objects, called (*weak*) *lossless rank condensers* [GR08, FS12, FSS14, FG15]. These objects were used by Gabizon and Raz [GR08] to construct affine extractors. They also play a crucial role in polynomial identity testing [KS11, SS12, FS12, FSS14].

A lossless rank condenser is defined as follows: Let $r \leq t \leq n$ be positive integers. A finite collection \mathcal{H} of matrices $E \in \mathbb{F}^{t \times n}$ is called an (r, L) -*lossless rank condenser* if for every matrix $M \in \mathbb{F}^{n \times r}$ of rank r , the number of $E \in \mathcal{H}$ satisfying $\text{rank}(EM) < r$ is at most L .

The connection between lossless rank condensers and variety evasive subspace families can be seen as follows: Let us assume every matrix $E \in \mathcal{H}$ has full rank t . Such a matrix E corresponds a linear t -subspace W of \mathbb{F}^n . On the other hand, a matrix $M \in \mathbb{F}^{n \times r}$ of rank r corresponds to a linear $(n-r)$ -subspaces of \mathbb{F}^n via $M \mapsto \ker(M)$, where $\ker(M) = \{u \in \mathbb{F}^n : uM = 0\}$ denotes the left kernel of M . It is easy to see that the condition $\text{rank}(EM) = r$ is equivalent to $\dim(W \cap \ker(M)) = t - r$. Passing from \mathbb{F}^n to \mathbb{P}^{n-1} by taking the quotient modulo scalars, this condition is also equivalent to the condition that the two projective subspaces $\mathbb{P}W$ and $\mathbb{P}(\ker(M))$ evade each other.

Every projective $(n-r-1)$ -subspace of \mathbb{P}^{n-1} can be realized as $\mathbb{P}(\ker(M))$ for some rank- r matrix M . Therefore, \mathcal{H} is an (r, L) -lossless rank condenser iff it is an (\mathcal{F}, ϵ) -evasive $(t-1)$ -subspace family on \mathbb{P}^{n-1} , where $\epsilon = L/|\mathcal{H}|$ and \mathcal{F} is the family of all $(n-r-1)$ -subspaces of \mathbb{P}^{n-1} .

Rank condensers are central objects in the theory of “linear-algebraic pseudorandomness” coined by Guruswami and Forbes [FG15]. Our study of variety evasive subspace families may be seen as one step of extending the theory to a nonlinear setting.

Explicit lossless rank condensers were used to construct explicit (deterministic) *affine extractors* [GR08] and more generally, *extractors for varieties* [Dvi12]. Similar ideas were used to construct explicit *deterministic extractors* (and *rank extractors*) for *polynomial sources* [DGW09], which also generalize affine extractors. It is an interesting question to us whether explicit variety-evasive subspace families and the related derandomized Noether’s normalization lemma (see below) can be similarly useful in this area.

1.2 Our Results

We have seen that variety evasive subspace families generalize some important and well-studied pseudorandom objects. This leads to the following natural question: For which interesting families \mathcal{F} of subvarieties can we construct explicit \mathcal{F} -evasive or (\mathcal{F}, ϵ) -evasive subspace families?

In this paper, we focus on the families of subvarieties of *bounded degree*. First, we recall the definition of the *degree* of a variety.

Definition 1.4 (degree). *The degree of an irreducible variety \mathcal{V} in \mathbb{P}^n (resp. \mathbb{A}^n) is the number of intersections of \mathcal{V} with a general projective (resp. affine) subspace of codimension $\dim(\mathcal{V})$. Following [Hei83], we define the degree of a (possibly reducible) variety to be the sum of the degrees of its irreducible components.*

For convenience, we introduce the following definition.

Definition 1.5. *We say a projective (resp. affine) k -subspace family \mathcal{H} on \mathbb{P}^n (resp. \mathbb{A}^n) is (n, d) -evasive if it is \mathcal{F} -evasive, where \mathcal{F} is chosen to be the family of all subvarieties of \mathbb{P}^n (resp. \mathbb{A}^n) of degree at most d . Similarly, we say \mathcal{H} is (n, d, ϵ) -evasive if it is (\mathcal{F}, ϵ) -evasive.*

Remark. In Definition 1.5, we do not make any assumption about the dimension of the varieties in \mathcal{F} or their irreducible components. We will see in Subsection 3.1 that in fact, it suffices to consider the subfamily of equidimensional varieties or even irreducible varieties of dimension $n - k - 1$ when constructing variety evasive k -subspace families.

For $n, d \in \mathbb{N}^+$ and $k \in \{0, 1, \dots, n\}$, define $N(k, d, n)$ by

$$N(k, d, n) := \min \left\{ \binom{(k+1)(n+1+d)}{(k+1)d}, \binom{(n-k)(n+1+d)}{(n-k)d}, \binom{(d-1)(n+1+d)}{(d-1)d} \right\}.$$

Our main theorem then states as follows.

Theorem 1.6 (Main Theorem). *For $n, d \in \mathbb{N}^+$, $k \in \{0, 1, \dots, n\}$, and $\epsilon \in (0, 1)$, there exists an (n, d, ϵ) -evasive k -subspace family (resp. affine k -subspace family) \mathcal{H} on \mathbb{P}^n (resp. \mathbb{A}^n) of size $\text{poly}(N(k, d, n), n, 1/\epsilon)$, which is $\text{poly}(n^{\min\{k+1, n-k, d\}d}, 1/\epsilon)$ when $d = o(n)$. Moreover, the total time complexity of computing the linear equations defining the projective or affine subspaces in \mathcal{H} is polynomial in $|\mathcal{H}|$ (and $\log p$, if the characteristic of the base field \mathbb{F} is $p > 0$). In particular, \mathcal{H} can be constructed in polynomial time when d is bounded.*

Remark (Boundedness of coefficients). For simplicity, the base field \mathbb{F} in this paper is assumed to be an algebraically closed field. Nevertheless, we choose the coefficients of the linear equations defining the subspaces in \mathcal{H} so that they live in either \mathbb{Q} (if $\text{char}(\mathbb{F}) = 0$) or a finite extension of \mathbb{F}_p (if $\text{char}(\mathbb{F}) = p > 0$). Moreover, when $\text{char}(\mathbb{F}) = 0$, the bit-length of the numerators and denominators of these coefficients are bounded by $|\mathcal{H}|^{O(1)}$. And when $\text{char}(\mathbb{F}) = p > 0$, the finite field that contains these coefficients has size $\max\{|\mathcal{H}|^{O(1)}, p\}$. This can be readily checked from our construction. Similar properties hold for all constructions presented in this paper.

Lower bound. As a complement of the above result, we establish the following lower bound for projective k -subspace families. It implies that when $n - k = \Omega(n)$, the assumption of d being bounded is necessary for a projective (n, d) -evasive k -subspace family to have polynomial size.

Theorem 1.7. *Let $n, d \in \mathbb{N}^+$ and $k \in \{0, 1, \dots, n-1\}$. Let \mathcal{F} be the family of equidimensional projective subvarieties of \mathbb{P}^n of dimension $n - k - 1$ and degree at most d . Suppose \mathcal{H} is an \mathcal{F} -evasive k -subspace family on \mathbb{P}^n . Then*

$$|\mathcal{H}| \geq \begin{cases} (n-k)(k+1) + 1 & \text{if } d = 1, \\ \max \left\{ d(n-k)(k+1) + 1, \binom{d+n-k}{d} + (n-k+1)k \right\} & \text{if } d > 1. \end{cases}$$

In particular, $|\mathcal{H}|$ is superpolynomial in n when $n - k = \Omega(n)$ and $d = \omega(1)$.

When $d = 1$, the lower bound $|\mathcal{H}| \geq (n-k)(k+1) + 1$ in Theorem 1.7 is achieved by known explicit lossless rank condensers [FS12, FSS14, For14] (see Subsection 2.2). For general d , the lower bound in Theorem 1.7 is also tight and matched by non-explicit constructions. See Section 4 for a discussion.

Next, we list two applications of our Main Theorem (Theorem 1.6): derandomizing Noether's normalization lemma for varieties of bounded degree, and polynomial identity testing for a special family of depth-4 arithmetic circuits.

1.2.1 Derandomizing Noether’s Normalization Lemma

Noether’s normalization lemma, introduced by Noether [Noe26], is an important result in commutative algebra and algebraic geometry with many applications. For example, it is used in the development of dimension theory and can be used to prove Grothendieck’s generic freeness lemma [Eis13]. It also has applications in computational algebraic geometry, e.g., computing the dimension of a projective variety [GH93, GHL⁺00].

The usual geometric formulation of Noether’s normalization lemma states that for any affine variety $\mathcal{V} \subseteq \mathbb{A}^n$ of dimension r , there exists a surjective *finite morphism* $\pi : \mathcal{V} \rightarrow \mathbb{A}^r$. (See Subsection 2.3 for the definition of finite morphisms.) Moreover, π may be chosen to be the restriction of a linear map $\mathbb{A}^n \rightarrow \mathbb{A}^r$.³ There is also a related projective or graded version of the lemma, which states that for any projective variety \mathcal{V} of dimension r , there exists a surjective finite morphism $\pi : \mathcal{V} \rightarrow \mathbb{P}^r$. A special form of this lemma goes back to Hilbert [Hil90].

In these versions of Noether’s normalization lemma, it can be shown that with high probability, a random linear map yields a valid finite morphism π , where “random” means the coefficients of the linear map are chosen randomly from a sufficiently large finite set $S \subseteq \mathbb{F}$. It is thus a natural question to derandomize the lemma.

Mulmuley [Mul17] studied a form of Noether’s normalization lemma and proved that derandomizing it is equivalent to a strengthened form of the black-box derandomization of PIT. There, the ambient projective space has exponential dimension and the problem is constructing a finite morphism $\pi : \mathcal{V} \rightarrow \mathbb{P}^k$ with a *succinct* specification in deterministic polynomial time, where $k = \text{poly}(\dim(\mathcal{V}))$ and \mathcal{V} is an *explicit variety* [Mul17]. This problem was later shown to be in PSPACE [FS18, GSS19]. The special case for the ring of matrix invariants under simultaneous conjugation was solved in quasipolynomial time by Forbes and Shpilka [FS13].

We consider Noether’s normalization lemma in its original context and completely derandomize it for projective/affine varieties of bounded degree. The following two theorems summarize our results.

Theorem 1.8. *Let $n, d \in \mathbb{N}^+$, $r \in \{0, 1, \dots, n\}$, and $\epsilon \in (0, 1)$. There exists an explicit collection \mathcal{L} of linear maps $\mathbb{A}^{n+1} \rightarrow \mathbb{A}^{r+1}$ of size $\text{poly}(N(k, d, n), n, 1/\epsilon)$ such that for every subvariety $\mathcal{V} \subseteq \mathbb{P}^n$ of dimension r and degree at most d , all but at most ϵ -fraction of $\pi \in \mathcal{L}$ induce a surjective finite morphism from \mathcal{V} to \mathbb{P}^r . Moreover, \mathcal{L} can be computed in time polynomial in $|\mathcal{L}|$ (and $\log p$, if $\text{char}(\mathbb{F}) = p > 0$).*

Theorem 1.9. *Let $n, d \in \mathbb{N}^+$ and $r \in \{0, 1, \dots, n\}$, and $\epsilon \in (0, 1)$. There exists an explicit collection \mathcal{L} of linear maps $\mathbb{A}^n \rightarrow \mathbb{A}^r$ of size $\text{poly}(N(k, d, n), n, 1/\epsilon)$ such that for every subvariety $\mathcal{V} \subseteq \mathbb{A}^n$ of dimension r and degree at most d , all but at most ϵ -fraction of $\pi \in \mathcal{L}$ restrict to a surjective finite morphism from \mathcal{V} to \mathbb{A}^r . Moreover, \mathcal{L} can be computed in time polynomial in $|\mathcal{L}|$ (and $\log p$, if $\text{char}(\mathbb{F}) = p > 0$).*

Theorem 1.8 is proved by derandomizing a standard proof of Noether’s normalization lemma that has a geometric flavor [Sha13]. Namely, we consider a projection $\pi : \mathbb{P}^n \setminus W \rightarrow \mathbb{P}^r$ sending \mathbf{x} to $(\ell_1(\mathbf{x}), \dots, \ell_{r+1}(\mathbf{x}))$, where ℓ_1, \dots, ℓ_r are linear forms and W is the $(n - r - 1)$ -subspace where these linear forms simultaneously vanish. It is known that π restricts to a finite morphism $\mathcal{V} \rightarrow \mathbb{P}^r$ iff $W \cap \mathcal{V} = \emptyset$. So the problem reduces to choosing a family of $(n - r - 1)$ -subspaces of \mathbb{P}^n such that most of them are disjoint from \mathcal{V} . This is exactly the property satisfied by our explicit variety evasive subspace families.

³For simplicity, we assume the base field is algebraically closed and hence infinite. But the lemma and our derandomization are valid as long as the field is large enough, depending on the variety \mathcal{V} . Nagata [Nag62] proved a version of the normalization lemma that is deterministic and does not require the base field to be sufficiently large, but the morphism he used is highly nonlinear. Due to the inductive nature of Nagata’s argument, it only yields a multiply exponential degree bound for the polynomials that define the morphism. Bruce and Erman [BE19] proved an effective Noether normalization result over finite fields, which states that with high probability, a random tuple of degree- d polynomials over a finite field induces a valid finite morphism for large enough d satisfying a certain effective bound. We leave it as an open problem to derandomize their version of the normalization lemma.

Theorem 1.9 is proved similarly. Here \mathbb{A}^n is viewed as an open subset of \mathbb{P}^n whose complement is the “hyperplane at infinity” H_∞ . Then we first construct a projection $\pi : \mathbb{P}^n \setminus W \rightarrow \mathbb{P}^r$ such that W is a subspace of H_∞ and is disjoint from the *projective closure* of \mathcal{V} . Then restrict π to \mathbb{A}^n . By carefully choosing π , we can make sure that the restriction is a linear map $\mathbb{A}^n \rightarrow \mathbb{A}^r$ and is a surjective finite morphism.

Dimension-preserving morphisms vs. finite morphisms. Our construction of finite linear morphisms preserve the dimension of a variety of low degree while reducing the dimension of the ambient space. This generalizes the property of lossless rank condensers. However, for the dimension-preserving property, better constructions are known. For example, it follows implicitly from the proof in [Dvi12] that most of the linear maps $\mathbb{A}^n \rightarrow \mathbb{A}^t$ from a lossless rank condenser $\mathcal{H} \subseteq \mathbb{F}^{t \times n}$ already preserve the dimension of a variety $\mathcal{V} \subseteq \mathbb{A}^n$.⁴ This was used by Dvir [Dvi12] in his explicit constructions of *extractors for varieties*, which generalize *affine extractors* [GR08].

On the other hand, the morphisms we construct are *finite morphisms*, which are strictly stronger than morphisms that are dimension-preserving. In particular, a finite morphism π always maps a closed set onto a closed set in the Zariski topology. Moreover, the preimage $\pi^{-1}(p)$ of *every* point p in the image of π is a finite set. Neither of these two properties is necessarily satisfied by morphisms that are only dimension-preserving.

These properties of finite morphisms may be useful in extractor theory or other areas. For example, in Theorem 1.9, the cardinality of $\pi^{-1}(p)$ is bounded by the degree of \mathcal{V} for every $p \in \pi(\mathcal{V})$, which translates into a lower bound for the min-entropy of the output of π when the input random source is distributed over the variety \mathcal{V} .

1.2.2 Depth-4 Polynomial Identity Testing

Depth-4 arithmetic circuits, also known as $\Sigma\Pi\Sigma\Pi$ circuits, play a very important role in polynomial identity testing. In a surprising result, Agrawal and Vinay [AV08] proved that a complete derandomization of black-box PIT for depth-4 circuits implies an $n^{O(\log n)}$ -time derandomization of PIT for general circuits of $\text{poly}(n)$ degree.

Dvir and Shpilka [DS07] initialized the approach of applying *Sylvester–Gallai* type theorems in geometry to PIT for depth-3 ($\Sigma\Pi\Sigma$) circuits. Extending this approach, Gupta [Gup14] formulated a conjecture of Sylvester–Gallai type and proved that his conjecture implies a complete derandomization of black-box PIT for depth-4 circuits with bounded top fan-in and bottom fan-in (also called $\Sigma\Pi\Sigma\Pi(k, r)$ circuits, where $k, r = O(1)$). In a recent breakthrough (built on [Shp19, PS20a]), Peleg and Shpilka [PS20b] proved that this conjecture holds for $k = 3$ and $r = 2$, and used it to give a polynomial-time black-box PIT algorithm for $\Sigma\Pi\Sigma\Pi(3, 2)$ circuits.

In [Gup14], Gupta divided $\Sigma\Pi\Sigma\Pi(k, r)$ into two families: those in a certain Sylvester–Gallai configuration and those that are not. His conjecture states that the circuits in the first family always have bounded *transcendence degree*, depending only on k and r . If the conjecture is true, then the results in [BMS13, ASSS16] imply a complete derandomization of the black-box PIT for this family. For the second family of circuits, which we call *non-SG* circuits, he proved that the black-box PIT can also be derandomized completely.

Theorem 1.10 ([Gup14]). *There exists a deterministic black-box PIT algorithm with time complexity $(dnk)^{\text{poly}(r \cdot k^2 + k)}$ for non-SG $\Sigma\Pi\Sigma\Pi(k, r)$ circuits of degree at most d in X_1, \dots, X_n over \mathbb{C} . In particular, the algorithm runs in polynomial time when k and r are bounded.*

⁴The intuition here is that \mathcal{V} can be locally approximated at a nonsingular point $p \in \mathcal{V}$ by its *tangent space* at p . So any linear map that preserves the dimension of this tangent space also preserves the dimension of \mathcal{V} .

Gupta’s proof of Theorem 1.10 is quite complex and used tools from computational algebraic geometry, including an effective version of Bertini irreducibility theorem [HS81] and radical membership testing (which in turn depends on *effective Nullstellensatz* [Kol88, Dub93]).

We observe that what is needed here is simply an explicit construction of subspaces intersecting certain varieties with (at most) the expected dimension. Plugging in our explicit construction of variety evasive subspace families, we obtain an improved black-box PIT algorithm with a simple proof.

Theorem 1.11. *There exists a deterministic black-box PIT algorithm with time complexity polynomial in $d \cdot \binom{k(n+1+r^k)}{kr^k} \cdot \binom{k-1+d}{k-1} \leq \text{poly}(d^k, n^{r^k}, r^{k^2 r^k})$ (and $\log p$, if $\text{char}(\mathbb{F}) = p > 0$) for non-SG $\Sigma\Pi\Sigma\Pi(k, r)$ circuits of degree at most d in X_1, \dots, X_n over an algebraically closed field \mathbb{F} .*

In particular, Theorem 1.11 improves the exponent of n in the time complexity from $\text{poly}(r^{k^2} + k)$ to $O(r^k)$, and the exponent of d from $\text{poly}(r^{k^2} + k)$ to $O(k)$. Moreover, our proof is more direct and conceptually simpler than the proof in [Gup14].

Remark. In [Muk16], Mukhopadhyay gave a deterministic polynomial-time black-box PIT algorithm for $\Sigma\Pi\Sigma\Pi(k, r)$ circuits satisfying a variant of the non-SG assumption. (Its time complexity is similar to the time complexity in Theorem 1.10.) It appears to us that his assumption in fact implies the non-SG assumption. The main tool used there is the *multivariate resultant*, which may be related to our approach based on Chow forms (see Subsection 1.3). Indeed, it is known that a multivariate resultant is the Chow form of a Veronese variety [GKZ94, Chapter 3, Example 2.4].

1.3 Proof Overview

We present an overview of our proof of Theorem 1.6 and that of Theorem 1.7.

Overview of the proof of Theorem 1.6. In the proof of Theorem 1.6, we focus on constructing a k -subspace family on \mathbb{P}^n . The case of \mathbb{A}^n can be easily derived from it by viewing \mathbb{A}^n as an open subset of \mathbb{P}^n and restricting to this subset.

Consider a variety $\mathcal{V} \subseteq \mathbb{P}^n$ of degree at most d . We want to construct a k -subspace family \mathcal{H} on \mathbb{P}^n , independent of \mathcal{V} , such that all but at most ϵ -fraction of $W \in \mathcal{H}$ evade \mathcal{V} . Our key ideas can be summarized as follows.

Reducing to the equidimensional/irreducible case of dimension $n - k - 1$. As a first step, we reduce the problem to the special case that \mathcal{V} is an *equidimensional* (or even irreducible) variety of \mathbb{P}^n of dimension $n - k - 1$, which means every irreducible component of \mathcal{V} has dimension exactly $n - k - 1$. This step is explained in Subsection 3.1.

Hitting the Chow form of \mathcal{V} . Denote by $\mathbb{G}(k, n)$ the Grassmannian consisting of all k -subspaces of \mathbb{P}^n . As $\text{codim}(\mathcal{V}) = n - (n - k - 1) > k$, a *general* k -subspace $W \in \mathbb{G}(k, n)$ is disjoint from \mathcal{V} , but we want to find such W explicitly.

One remarkable fact in algebraic geometry is that there is a single polynomial $\tilde{R}_{\mathcal{V}}$ on the Grassmannian $\mathbb{G}(k, n)$ that defines precisely the subset of k -subspaces that intersect \mathcal{V} . This polynomial $\tilde{R}_{\mathcal{V}}$ is called the *Chow form* of \mathcal{V} (in *Stiefel coordinates*). Chow forms are also known as *Cayley forms* or *Cayley–van der Waerden–Chow forms* in literature. They were introduced by Cayley [Cay60] to represent curves in \mathbb{P}^3 and later generalized by Chow and van der Waerden [CvdW37]. See [DS95] for an introduction to Chow forms and [GKZ94] for an exposition in the context of elimination theory.

To be more specific, for a k -subspace $W \in \mathbb{G}(k, n)$, we choose a $(k + 1) \times (n + 1)$ matrix A that represents W . The Chow form $\tilde{R}_{\mathcal{V}}$ is a polynomial of degree $(k + 1) \deg(\mathcal{V})$ in $(k + 1)(n + 1)$ variables

with the following property: $\tilde{R}_{\mathcal{V}}$ vanishes at the matrix A (viewed as a list of $(k+1)(n+1)$ coordinates) if and only if $\mathcal{V} \cap W \neq \emptyset$. Thus, $\tilde{R}_{\mathcal{V}}$ defines precisely the subset of “bad” k -subspaces that we want to avoid.

Therefore, the problem becomes finding a collection of $(k+1) \times (n+1)$ matrices of full rank that “hit” the polynomial $\tilde{R}_{\mathcal{V}}$ of degree $(k+1) \deg(\mathcal{V}) \leq (k+1)d$. Using black-box PIT for low degree polynomials (see Subsection 2.1), we are able to construct an (n, d, ϵ) -evasive k -subspace family of size polynomial in $\binom{(k+1)(n+1+d)}{(k+1)d}$ and $1/\epsilon$, which is $\text{poly}(n, 1/\epsilon)$ when k and d are both bounded. A similar “dual” construction yields a k -subspace family of size polynomial in $\binom{(n-k)(n+1+d)}{(n-k)d}$ and $1/\epsilon$, which is $\text{poly}(n, 1/\epsilon)$ when both $n-k$ and d are bounded. For applications where d is small and either k or $n-k$ is small (e.g., Theorem 1.11), these constructions are good enough. However, when k and $n-k$ are both linear in n , the resulting k -subspace families have exponential size in n , even if d is bounded.

A two-step construction. To obtain a good construction for *arbitrary* dimension k , we use a standard fact from algebraic geometry, which states that the codimension of an irreducible subvariety $\mathcal{V} \subseteq \mathbb{P}^n$ in $\text{span}(\mathcal{V})$ is at most $\deg(\mathcal{V}) - 1$, where $\text{span}(\mathcal{V})$ denotes the smallest projective subspace containing \mathcal{V} (see Lemma 3.10). Therefore, for irreducible \mathcal{V} of degree at most d , there exists a projective subspace Λ of dimension (at most) $\dim(\mathcal{V}) + d - 1$ that contains \mathcal{V} .

Our idea is to use a two-step construction. Namely, we first construct subspaces of dimension $n - \dim \Lambda - 1$ that evade Λ , and then extend these subspaces to k -subspaces that evade \mathcal{V} . The first step is just the problem of constructing lossless rank condensers, which has an optimal solution [FS12, FSS14] (see Subsection 2.2). The second step is equivalent to extending a $((k+1) - (d-1)) \times (n+1)$ matrix B to a $(k+1) \times (n+1)$ matrix $\begin{pmatrix} A \\ B \end{pmatrix}$ such that the polynomial $\tilde{R}_{\mathcal{V}}$ does not vanish at $\begin{pmatrix} A \\ B \end{pmatrix}$. The polynomial $\tilde{R}_{\mathcal{V}}(\begin{pmatrix} \cdot \\ B \end{pmatrix})$ has degree $(d-1) \deg(\mathcal{V}) \leq (d-1)d$, as there are only $d-1$ rows of free variables. Using black-box PIT for low degree polynomials, we obtain a construction of size polynomial in $\binom{(d-1)(n+1+d)}{(d-1)d}$ and $1/\epsilon$, which is $\text{poly}(n, 1/\epsilon)$ for any bounded d .

Overview of the proof of Theorem 1.7. Our lower bound (Theorem 1.7) follows from a dimension counting argument. Let $C(r, d, n)$ be the set of all varieties $\mathcal{V} \subseteq \mathbb{P}^n$ of dimension $r := n - k - 1$ and degree d , which is the space of varieties that we want to evade.

Roughly speaking, the idea is to show that (1) $C(r, d, n)$ itself can be realized as a subvariety of some projective space \mathbb{P}^N , and (2) for every k -subspace W , the subset of $\mathcal{V} \in C(r, d, n)$ that W fails to evade is the intersection of $C(r, d, n)$ with some hyperplane H_W of \mathbb{P}^N .

To see how (1) and (2) above lead to a lower bound, suppose \mathcal{H} is a $C(r, d, n)$ -evasive k -subspace family, i.e., for any $\mathcal{V} \in C(r, d, n)$, there exists $W \in \mathcal{H}$ that is disjoint from \mathcal{V} . Then the intersection $C(r, d, n) \cap \bigcap_{W \in \mathcal{H}} H_W$ must be empty. On the other hand, taking the intersection with each hyperplane H_W reduces the dimension of a projective variety by at most one. So we have a lower bound $|\mathcal{H}| \geq \dim(C(r, d, n)) + 1$.

How do we realize $C(r, d, n)$ as a subvariety of \mathbb{P}^N ? It turns out that this is a classical problem in the study of moduli spaces and a solution was given by Cayley [Cay60] and Chow–van der Waerden [CvdW37] using the *Chow embedding*: The Chow embedding $C(r, d, n) \rightarrow \mathbb{P}^N$ simply sends a variety \mathcal{V} to its Chow form $\tilde{R}_{\mathcal{V}}$, where $\tilde{R}_{\mathcal{V}}$ is viewed as a point in the projective space \mathbb{P}^N whose homogeneous coordinates are given by the coefficients of $\tilde{R}_{\mathcal{V}}$.⁵

A technical issue here is that the image of $C(r, d, n)$ under the Chow embedding is generally not closed in the Zariski topology. To fix this issue, the definition of $C(r, d, n)$ needs to be modified so that it contains not only subvarieties of \mathbb{P}^n , but also (*effective*) *algebraic cycles* on \mathbb{P}^n , which are a generalization of subvarieties.

⁵The actual Chow embedding we use has a slightly different form, which is essentially equivalent to the one described here.

A theorem of Chow and van der Waerden [CvdW37] then states that the Chow embedding does embed $C(r, d, n)$ in a projective subspace \mathbb{P}^N as a subvariety, known as a *Chow variety*.

Finally, we also need a lower bound for the dimension of the Chow variety $C(r, d, n)$. In fact, the exact value of $\dim(C(r, d, n))$ was determined by Azcue [Azc92] and independently by Lehmann [Leh17]. Plugging in the value of $\dim(C(r, d, n))$ proves Theorem 1.7.

1.4 Other Related Work

In [DKL14], Dvir, and Kollár, and Lovett constructed explicit *variety evasive sets*, which are large subsets of \mathbb{F}_q^n over a finite field \mathbb{F}_q that have small intersection with affine varieties of fixed dimension and bounded degree. It generalizes an earlier construction of *subspace evasive sets* of Dvir and Lovett [DL12]. The definition of evasiveness there is different from ours, but they are related, since a key step in the proofs of [DL12, DKL14] is proving the intersection of two varieties has dimension zero. We also note that a subspace/variety evasive set is a single set, defined in a highly nonlinear way, whereas we define a variety evasive subspace family to be a collection of projective or affine subspaces. Finally, the results in [DL12, DKL14] hold only for affine subspaces/subvarieties, whereas we give our construction first in the projective setting and then derive the affine counterpart from it.

Guruswami and Xing in [GX13] introduced a related notion called *subspace designs*. A subspace design is a collection \mathcal{H} of large subspaces of \mathbb{F}^n such that for any small subspace $V \subseteq \mathbb{F}^n$, the number of $W \in \mathcal{H}$ satisfying $\dim(W \cap V) > 0$ is small (or even the sum $\sum_{W \in \mathcal{H}} \dim(W \cap V)$ is small). An equivalence between subspace designs and lossless rank condensers was proved in [FG15]. Explicit subspace designs were constructed by Guruswami and Kopparty [GK16] and also by Guruswami, Xing, and Yuan [GXY18]. They have applications to constructing explicit list-decodable codes with small list size [GX13, GWX16, KRZSW18, GRZ20] and explicit dimension expanders [FG15, GRX21]. Subspace designs were also used to prove lower bounds in communication complexity [CGS20].

Jeronimo, Krick, Sabia, and Sombra [JKSS04] gave a randomized algorithm, in the Blum-Shub-Smale model over fields of characteristic zero, that computes the Chow forms of varieties defined by input polynomials. The (expected) time complexity of their algorithm is polynomial in the sizes of the arithmetic circuits encoding the input polynomials and the *geometric degree* of the polynomial system. See also the survey by Krick [Kri02].

Chow varieties of effective zero-cycles and their higher secant varieties are related to lower bounds for depth-3 arithmetic circuits. They have received a considerable amount of attention in Geometric Complexity Theory [Lan12, Lan15].

Organization of the paper. Preliminaries and notations are given in Section 2. We prove the Main Theorem (Theorem 1.6) in Section 3. The lower bound (Theorem 1.7) is proved in Section 4. The applications to derandomization of Noether’s normalization lemma and PIT for depth-4 circuits are explained in Section 5. Finally, we list some open problems and future directions in Section 6.

2 Preliminaries and Notations

Define $\mathbb{N} := \{0, 1, 2, \dots\}$ and $\mathbb{N}^+ := \{1, 2, \dots\}$. Let $[n] := \{1, 2, \dots, n\}$ for $n \in \mathbb{N}$. For a set S and $k \in \mathbb{N}$, denote by $\binom{S}{k}$ the set of all subsets of S of cardinality k .

Denote by \mathbb{F} an algebraically closed field throughout this paper. We use notations like $\mathbb{F}[X_{i,j} : i \in [n], j \in [m]]$ to denote the polynomial ring over \mathbb{F} in a finite set of variables (in this case, in the set of variables $\{X_{i,j} : i \in [n], j \in [m]\}$). The vector space of $n \times m$ matrices over \mathbb{F} is denoted by $\mathbb{F}^{n \times m}$.

For an $n \times m$ matrix A and subsets $S \subseteq [n]$, $T \subseteq [m]$, denote by $A_{S,T}$ the submatrix of A whose rows and columns are selected by S and T respectively, where the orderings of rows and columns are preserved.

2.1 Black-Box PIT for Low Degree Polynomials

For convenience, we strengthen the definition of hitting sets as follows.

Definition 2.1 (ϵ -hitting set). *Let \mathcal{F} be a family of polynomials in $\mathbb{F}[X_1, \dots, X_n]$ and $\epsilon \in (0, 1)$. We say a finite collection of points $\mathcal{H} \subseteq \mathbb{F}^n$ is an ϵ -hitting set for \mathcal{F} if for any nonzero $Q \in \mathcal{F}$, the evaluation $Q(\alpha)$ is nonzero for all but at most ϵ -fraction of $\alpha \in \mathcal{H}$.*

We need an explicit construction of ϵ -hitting sets for low degree polynomials. This problem has been well studied [DL78, Zip79, Sch80, KS01, Bog05, Lu12, CTS13, Bsh14, BP20]. For completeness, we present a construction based on sparse polynomial identity testing.

Recall that a polynomial is s -sparse if it has at most s monomials. We need the following lemma from [AGKS15].

Lemma 2.2 ([AGKS15, Lemma 4, restated]). *For $n, s, d \in \mathbb{N}^+$ and $\epsilon_0 \in (0, 1)$, there exist maps $w_1, w_2, \dots, w_N : [n] \rightarrow [N \log N]$, where $N = \text{poly}(n, s, \log d, \epsilon_0^{-1})$, such that for any nonzero s -sparse polynomial $f \in \mathbb{F}[X_1, \dots, X_n]$ of individual degree at most d , all but at most ϵ_0 -fraction of w_i among w_1, w_2, \dots, w_N satisfies $f(Y^{w_i(1)}, \dots, Y^{w_i(n)}) \neq 0$. Moreover, the time complexity of computing w_1, w_2, \dots, w_N is polynomial in N .*

Given $n, d \in \mathbb{N}^+$ and $\epsilon \in (0, 1)$, we construct an ϵ -hitting set for n -variate polynomials of degree at most d as follows:

1. Let $s = \binom{n+d}{d}$, $\epsilon_0 = \epsilon/2$, and $M = \lceil \epsilon_0^{-1} dN \log N \rceil$, where N is as in Lemma 2.2.
2. Let w_1, \dots, w_N be as in Lemma 2.2, which can be computed in time $\text{poly}(N)$.
3. If $\text{char}(\mathbb{F}) = 0$, let $S = [M] \subseteq \mathbb{Z} \subseteq \mathbb{F}$. If $\text{char}(\mathbb{F}) = p > 0$, choose a finite extension \mathbb{F}_q of \mathbb{F}_p such that $M \leq q = \text{poly}(M, p)$, and choose S to be a subset of $\mathbb{F}_q \subseteq \mathbb{F}$ of cardinality M .
4. Finally, construct the following collection of points in \mathbb{F}^n of size MN

$$T = \{(\alpha^{w_i(1)}, \dots, \alpha^{w_i(n)}) : \alpha \in S, i \in [N]\} \subseteq \mathbb{F}^n.$$

Lemma 2.3. *For any nonzero polynomial $f \in \mathbb{F}[X_1, \dots, X_n]$ of degree at most d , we have $f(u) \neq 0$ for all but at most ϵ -fraction of $u \in T$. The collection T has cardinality $\text{poly}\left(\binom{n+d}{d}, 1/\epsilon\right)$ and can be computed in time $\text{poly}(|T|)$.*

Proof. Let $f \in \mathbb{F}[X_1, \dots, X_n]$ be a nonzero polynomial of degree at most d . Note that f is trivially s -sparse, where $s = \binom{n+d}{d}$. So by Lemma 2.2, for all but at most ϵ_0 -fraction of $i \in [N]$, we have $\tilde{f}_i := f(Y^{w_i(1)}, \dots, Y^{w_i(n)}) \neq 0$. Consider $i \in [N]$ such that $\tilde{f}_i \neq 0$. Note that \tilde{f}_i is a univariate polynomial of degree at most $dN \log N$. So it has at most $dN \log N \leq \epsilon_0 M$ zeros. Therefore, by the choice of M , we have $f(\alpha^{w_i(1)}, \dots, \alpha^{w_i(n)}) = \tilde{f}_i(\alpha) \neq 0$ for all but at most ϵ_0 -fraction of $\alpha \in S$. It follows that $f(u) \neq 0$ holds for all but at most ϵ -fraction of $u \in T$, as claimed. The rest of the lemma follows easily from the construction. \square

Note that the seed length required to choose a random element in T is $\log |T| = O(\log \binom{n+d}{d} + \log(1/\epsilon))$, which is optimal up to a constant factor. We have made no effort to optimize the constant hidden in $O(\cdot)$. Interested readers may find the state-of-the-art result in [BP20], which achieves the optimal constant, at least for $d = o(n)$.

2.2 Explicit Lossless Rank Condensers

We need the following lemma in the context of *lossless rank condensers*. The construction in the lemma was given by Forbes and Shpilka [FS12] and the lemma itself follows implicitly from the analysis of Forbes, Saptharishi, and Shpilka in [FSS14]. It was also stated explicitly in [For14, Theorem 5.4.3].

Lemma 2.4 ([FSS14, For14]). *Let $n \in \mathbb{N}^+$ and $r \in [n]$. Let $\omega \in \mathbb{F}^\times$ such that the multiplicative order of ω is at least n . Define the $r \times n$ matrix $W = (w_{i,j})_{i \in [r], j \in [n]}$ over $\mathbb{F}[X]$ by*

$$w_{i,j} = (\omega^{i-1} X)^{j-1}.$$

Then for every $n \times r$ matrix M over \mathbb{F} of rank r , the polynomial $\det(WM) \in \mathbb{F}[X]$ is nonzero and has degree at most $r(n-r)$ after dividing out powers of X .

Corollary 2.5. *Let n, r, W be as in Lemma 2.4 and $\epsilon \in (0, 1)$. Let $S \subseteq \mathbb{F}^\times$ be a finite set of cardinality at least $r(n-r)/\epsilon$. For every $n \times r$ matrix M over \mathbb{F} of rank r , we have $\text{rank}(W(\alpha)M) = r$ for all but at most ϵ -fraction of $\alpha \in S$, where $W(\alpha)$ denotes the matrix $(w_{i,j}(\alpha))_{i \in [r], j \in [n]}$ over \mathbb{F} .*

Corollary 2.5 states that the collection $\{W(\alpha) : \alpha \in S\}$ of matrices is a (weak) $(r, \epsilon|S|)$ -*lossless rank condenser*, as defined in [FG15]. Note that for each $\alpha \in S$, we have $\text{rank}(W(\alpha)) = r$ and hence $W(\alpha)$ correspond to an $(r-1)$ -subspace $U_{W(\alpha)}$ of \mathbb{P}^{n-1} . As explained in the introduction, the collection $\mathcal{H} = \{U_{W(\alpha)} : \alpha \in S\}$ is an (\mathcal{F}, ϵ) -*evasive* $(r-1)$ -subspace family on \mathbb{P}^{n-1} , where \mathcal{F} is the family of $(n-r-1)$ -subspaces of \mathbb{P}^{n-1} . Choosing S of size $r(n-r) + 1$ and $\epsilon = 1 - \frac{1}{r(n-r)+1}$ shows that the lower bound in Theorem 1.7 is achieved when $d = 1$.

2.3 Preliminaries on Algebraic Geometry

We list basic preliminaries and notations on algebraic geometry used in this paper. One can also refer to a standard text, e.g., [Sha13, Har13].

Affine and projective spaces. For $n \in \mathbb{N}$, write \mathbb{A}^n for the *affine n -space* over \mathbb{F} . It is defined to be the set \mathbb{F}^n equipped with the *Zariski topology*, defined as follows: A subset $S \subseteq \mathbb{A}^n$ is (*Zariski*-)closed if it is the set of common zeros of a set of polynomials in $\mathbb{F}[X_1, \dots, X_n]$. The complement of a closed set is an *open* set. The origin of an affine space is denoted by $\mathbf{0}$.

Write \mathbb{P}^n for the (*projective*) n -space over \mathbb{F} , defined to be the quotient set $(\mathbb{A}^{n+1} \setminus \{\mathbf{0}\}) / \sim$, where \sim is the equivalence relation defined by scaling, i.e., $u \sim v$ if $u = cv$ for some $c \in \mathbb{F}^\times$. The set \mathbb{P}^n is again equipped with the *Zariski topology*, where a subset is closed if it is the set of common zeros of a set of *homogeneous* polynomials in $\mathbb{F}[X_1, \dots, X_{n+1}]$. We use $(n+1)$ -tuples (x_1, \dots, x_{n+1}) to represent points in \mathbb{P}^n , called *homogeneous coordinates*.

For a vector space V over \mathbb{F} of dimension $n+1$, where $n \in \mathbb{N}$, define the projective space $\mathbb{P}V = (V \setminus \{\mathbf{0}\}) / \sim$, where \sim is again the equivalence relation defined by scaling. By fixing a coordinate system of V and identifying it with \mathbb{A}^{n+1} , we may identify $\mathbb{P}V$ with \mathbb{P}^n .

Varieties. *Varieties* in this paper refer to either projective or affine varieties. A *projective (resp. affine) variety* is simply a closed subset of a projective (resp. affine) subspace. If \mathcal{V}_1 and \mathcal{V}_2 are closed subsets of a projective or affine space and $\mathcal{V}_1 \subseteq \mathcal{V}_2$, we say \mathcal{V}_1 is a *subvariety* of \mathcal{V}_2 .

A variety is *reducible* if it is the union of finitely many proper subvarieties, and otherwise *irreducible*. Affine and projective spaces are irreducible. A variety \mathcal{V} can be uniquely written as the union of finitely many irreducible varieties, which are called the *irreducible components* of \mathcal{V} .

A projective or affine variety is called a *hypersurface* (resp. *hyperplane*) if it is definable by a single polynomial (resp. single linear polynomial).

Hilbert's Nullstellensatz. An ideal I of a commutative ring R is *radical* if $a^m \in I$ implies $a \in I$ for every $a \in R$ and $m \in \mathbb{N}^+$. For an ideal I of $\mathbb{F}[X_1, \dots, X_n]$, denote by $\mathcal{V}(I)$ the subvariety of \mathbb{A}^n defined by the polynomial in I . Define $\mathcal{V}(f_1, \dots, f_k) = \mathcal{V}(\langle f_1, \dots, f_k \rangle)$ for $f_1, \dots, f_k \in \mathbb{F}[X_1, \dots, X_n]$. For a subvariety \mathcal{V} of \mathbb{A}^n , denote by $I(\mathcal{V})$ the ideal of $\mathbb{F}[X_1, \dots, X_n]$ consisting of all the polynomials vanishing on \mathcal{V} . *Hilbert's Nullstellensatz* states that the map $\mathcal{V} \mapsto I(\mathcal{V})$ is an inclusion-reversing one-to-one correspondence between the subvarieties of \mathbb{A}^n and the radical ideals of $\mathbb{F}[X_1, \dots, X_n]$, with the inverse map $I \mapsto \mathcal{V}(I)$.

For a subvariety \mathcal{V} of \mathbb{A}^n , define $\mathbb{F}[\mathcal{V}] := \mathbb{F}[X_1, \dots, X_n]/I(\mathcal{V})$, called the *coordinate ring* of \mathcal{V} .

Projective Nullstellensatz. Consider the polynomial ring $R = \mathbb{F}[X_1, \dots, X_{n+1}]$. It can be written as a direct sum $R = \bigoplus_{d=0}^{\infty} R_d$ where each R_d denotes the space of degree- d homogeneous polynomials, called the *homogeneous part of degree d* of R or simply the *degree- d part* of R . For an ideal I of R and $d \in \mathbb{N}$, let $I_d := I \cap R_d$, called the *degree- d part* of I . We say I is a *homogeneous ideal* if $I = \bigoplus_{d=0}^{\infty} I_d$. For a homogeneous ideal I of R , we have $R/I = \bigoplus_{d=0}^{\infty} (R/I)_d$ where $(R/I)_d := R_d/I_d$.

For a homogeneous ideal I of R , denote by $\mathcal{V}(I)$ the subvariety of \mathbb{P}^n defined by the homogeneous polynomials in I . Define $\mathcal{V}(f_1, \dots, f_k) = \mathcal{V}(\langle f_1, \dots, f_k \rangle)$ for homogeneous polynomials $f_1, \dots, f_k \in R$. For a subvariety \mathcal{V} of \mathbb{P}^n , denote by $I(\mathcal{V})$ the ideal generated by the homogeneous polynomials vanishing on \mathcal{V} , which is a homogeneous ideal. The *projective Nullstellensatz* states that the map $\mathcal{V} \mapsto I(\mathcal{V})$ is an inclusion-reversing one-to-one correspondence between the nonempty subvarieties of \mathbb{P}^n and the radical homogeneous ideals of R properly contained in $\langle X_1, \dots, X_{n+1} \rangle$, with the inverse map $I \mapsto \mathcal{V}(I)$.

For a subvariety $\mathcal{V} \subseteq \mathbb{P}^n$ and the corresponding homogeneous ideal $I = I(\mathcal{V})$, we say R/I is the *homogeneous coordinate ring* of \mathcal{V} .

Morphisms. Let $\mathcal{V}_1 \subseteq \mathbb{A}^n$ and $\mathcal{V}_2 \subseteq \mathbb{A}^m$ be affine varieties. A *morphism* from \mathcal{V}_1 to \mathcal{V}_2 is a map $f : \mathcal{V}_1 \rightarrow \mathcal{V}_2$ that is a restriction of a polynomial map $\mathbb{A}^n \rightarrow \mathbb{A}^m$. Such a morphism f is associated with a ring homomorphism $f^\# : \mathbb{F}[\mathcal{V}_2] \rightarrow \mathbb{F}[\mathcal{V}_1]$, making $\mathbb{F}[\mathcal{V}_1]$ an algebra over $\mathbb{F}[\mathcal{V}_2]$. We say f is *finite* if $\mathbb{F}[\mathcal{V}_1]$ is finitely generated as an $\mathbb{F}[\mathcal{V}_2]$ -module.

Let $f : \mathcal{V}_1 \rightarrow \mathcal{V}_2$ be a map between projective varieties \mathcal{V}_1 and \mathcal{V}_2 . We say f is a morphism from \mathcal{V}_1 to \mathcal{V}_2 if there exists a collection of open subsets $\{U_i\}_{i \in I}$ of \mathcal{V}_2 such that $\mathcal{V}_2 = \bigcup_{i \in I} U_i$ (i.e., $\{U_i\}_{i \in I}$ is an open cover of \mathcal{V}_2) and for each $i \in I$, the restriction $f|_{f^{-1}(U_i)} : f^{-1}(U_i) \rightarrow U_i$ is a morphism between affine varieties. Furthermore, if each $f|_{f^{-1}(U_i)}$ is finite, then we say f is finite. Finiteness does not depend on the choice of the affine open cover. Namely, if $f : \mathcal{V}_1 \rightarrow \mathcal{V}_2$ is a finite morphism between projective varieties \mathcal{V}_1 and \mathcal{V}_2 , and U is an open subset of \mathcal{V}_2 such that $f|_{f^{-1}(U)} : f^{-1}(U) \rightarrow U$ is a morphism between affine varieties, then $f|_{f^{-1}(U)}$ is also finite.

The image of a morphism $f : \mathcal{V}_1 \rightarrow \mathcal{V}_2$ is denoted by $\text{Im}(f)$ or $f(\mathcal{V}_1)$. The image of a closed set under a finite morphism is still closed.

Dimension. The *dimension* of an irreducible variety \mathcal{V} , denoted by $\dim(\mathcal{V})$, is the largest integer m such that there exists a chain of irreducible varieties $\emptyset \subsetneq \mathcal{V}_0 \subsetneq \mathcal{V}_1 \subsetneq \dots \subsetneq \mathcal{V}_m = \mathcal{V}$. More generally, the dimension of a nonempty variety is the maximal dimension of its irreducible components. We define the dimension of an empty set to be $-\infty$. A variety is *equidimensional* if its irreducible components have the same dimension.

If $\pi : \mathcal{V} \rightarrow \mathcal{V}'$ is a finite morphism, then $\dim(\mathcal{V}) = \dim(\pi(\mathcal{V}))$.

Degree. The *degree* of an irreducible subvariety \mathcal{V} of \mathbb{P}^n (resp. \mathbb{A}^n), denoted by $\deg(\mathcal{V})$, is the number of intersections of \mathcal{V} with a projective (resp. affine) subspace of codimension $\dim(\mathcal{V})$ in general position. More generally, we define the degree of a subvariety of \mathbb{P}^n or \mathbb{A}^n to be the sum of the degrees of its irreducible components.

Projective closure. The affine n -space \mathbb{A}^n may be regarded as an open subset of \mathbb{P}^n via the map $(x_1, \dots, x_n) \mapsto (x_1, \dots, x_n, 1)$. The complement $H_\infty := \mathbb{P}^n \setminus \mathbb{A}^n$ is a hyperplane of \mathbb{P}^n defined by $X_{n+1} = 0$, called the *hyperplane at infinity*. For an affine subvariety \mathcal{V} of $\mathbb{A}^n \subseteq \mathbb{P}^n$, the smallest projective subvariety of \mathbb{P}^n containing \mathcal{V} is the *projective closure* of \mathcal{V} , which we denote by \mathcal{V}_{cl} . It is known that $\mathcal{V}_{\text{cl}} \cap \mathbb{A}^n = \mathcal{V}$, $\dim(\mathcal{V}_{\text{cl}}) = \dim(\mathcal{V})$, and $\deg(\mathcal{V}_{\text{cl}}) = \deg(\mathcal{V})$.

Joins of disjoint projective varieties. For two distinct points $p, q \in \mathbb{P}^n$, denote by \overline{pq} the unique projective line passing through them. For two *disjoint* projective subvarieties $\mathcal{V}_1, \mathcal{V}_2 \subseteq \mathbb{P}^n$, define the *join* $J(\mathcal{V}_1, \mathcal{V}_2)$ of \mathcal{V}_1 and \mathcal{V}_2 as

$$J(\mathcal{V}_1, \mathcal{V}_2) := \bigcup_{p \in \mathcal{V}_1, q \in \mathcal{V}_2} \overline{pq}.$$

Lemma 2.6 ([Har13, Examples 6.17, 11.36, and 18.17]). *$J(\mathcal{V}_1, \mathcal{V}_2)$ is a subvariety of \mathbb{P}^n of dimension $\dim(\mathcal{V}_1) + \dim(\mathcal{V}_2) + 1$ and degree at most $\deg(\mathcal{V}_1) \cdot \deg(\mathcal{V}_2)$.*

We also need the following lemmas.

Lemma 2.7. *Let \mathcal{V} be a nonempty subvariety of \mathbb{P}^n of dimension $r < n$. Let $W \subseteq \mathbb{P}^n$ be a k -subspace disjoint from \mathcal{V} . Let k' be an integer satisfying $k \leq k' \leq n - r - 1$. Then there exists a k' -subspace $W' \subseteq \mathbb{P}^n$ such that $W \subseteq W'$ and W' is disjoint from \mathcal{V} . In particular, choosing W to be a point not in \mathcal{V} shows that there exists an $(n - r - 1)$ -subspace disjoint from \mathcal{V} .*

Proof. We prove the lemma for the special case $k' = k + 1 \leq n - r - 1$ and the general case follows from iteration. By Lemma 2.6, $J(\mathcal{V}, W)$ has dimension $r + k + 1 \leq n - 1$. Pick a point $p \in \mathbb{P}^n \setminus J(\mathcal{V}, W)$ and let $W' = J(p, W)$. Then W' is a $(k + 1)$ -subspace and $W \subseteq W'$. To prove W' is disjoint from \mathcal{V} , assume to the contrary that there exists a point $q \in W' \cap \mathcal{V}$. By definition, $q \in \overline{pq'}$ for some $q' \in W$. As W is disjoint from \mathcal{V} , we have $q' \neq q$. Then $p \in \overline{pq'} = \overline{qq'} \in J(\mathcal{V}, W)$, contradicting the choice of p . \square

Lemma 2.8 ([Har13, Exercise 11.6 and Corollary 18.5]). *Let \mathcal{V} be a nonempty equidimensional subvariety of \mathbb{P}^n and H a hypersurface of \mathbb{P}^n not containing an irreducible component of \mathcal{V} . Then $\mathcal{V} \cap H$ is an equidimensional subvariety of dimension $\dim(\mathcal{V}) - 1$ and degree at most $\deg(\mathcal{V}) \cdot \deg(H)$ (or an empty set if $\dim(\mathcal{V}) = 0$).*

Lemma 2.9 ([Sha13, Section I.6.2, Theorem 6]). *Suppose \mathcal{V}_1 and \mathcal{V}_2 are subvarieties of \mathbb{P}^n and $\dim(\mathcal{V}_1) + \dim(\mathcal{V}_2) \geq n$. Then $\mathcal{V}_1 \cap \mathcal{V}_2 \neq \emptyset$ and $\dim(\mathcal{V}_1 \cap \mathcal{V}_2) \geq \dim(\mathcal{V}_1) + \dim(\mathcal{V}_2) - n$.*

3 Proof of the Main Theorem

We prove the Main Theorem (Theorem 1.6) in this section. In Subsection 3.1, we show that it suffices to consider equidimensional or irreducible subvarieties of dimension $n - k - 1$. Subsection 3.2 contains an introduction to Chow forms. Finally, in Subsection 3.3, we present the explicit constructions and complete the proof of Theorem 1.6.

3.1 Reducing to the Case of Equidimensional or Irreducible Varieties

The following lemma states that to construct k -subspace families that are evasive for subvarieties of \mathbb{P}^n , it suffices to consider equidimensional subvarieties of dimension $n - k - 1$ (i.e., codimension $k + 1$).

Lemma 3.1. *Let $n, d \in \mathbb{N}^+$ and $k \in \{0, 1, \dots, n - 1\}$. Let \mathcal{F} be the family of all equidimensional subvarieties of \mathbb{P}^n of dimension $n - k - 1$ and degree at most d . Then an (\mathcal{F}, ϵ) -evasive k -subspace family is also (n, d, ϵ) -evasive.*

The proof of Lemma 3.1 is based on the following claim.

Claim 3.2. *Let \mathcal{V} be an irreducible subvariety of \mathbb{P}^n . There exists a subvariety $\tilde{\mathcal{V}} \subseteq \mathbb{P}^n$ of dimension $n - k - 1$ and degree at most $\deg(\mathcal{V})$ such that any k -subspace of \mathbb{P}^n that evades $\tilde{\mathcal{V}}$ also evades \mathcal{V} .*

Proof. If $\dim(\mathcal{V}) = n - k - 1$, then just let $\tilde{\mathcal{V}} = \mathcal{V}$.

Now assume $\dim(\mathcal{V}) < n - k - 1$. Let $t = (n - k - 1) - \dim(\mathcal{V}) - 1$ and let $\tilde{\mathcal{V}}$ be the join of \mathcal{V} and a t -subspace disjoint from \mathcal{V} (which exists by Lemma 2.7). Then $\tilde{\mathcal{V}}$ is a projective subvariety of dimension $n - k - 1$ and degree at most $\deg(\mathcal{V})$ by Lemma 2.6. Suppose W is a k -subspace that evades $\tilde{\mathcal{V}}$. Then W is disjoint from $\tilde{\mathcal{V}} \supseteq \mathcal{V}$. So W also evades \mathcal{V} .

Finally, assume $\dim(\mathcal{V}) > n - k - 1$. Let $t = \dim(\mathcal{V}) - (n - k - 1)$. By Lemma 2.8, there exist t hyperplanes H_1, \dots, H_t of \mathbb{P}^n such that $\mathcal{V} \cap \bigcap_{i=1}^t H_i$ is equidimensional of dimension $n - k - 1$ and degree at most $\deg(\mathcal{V})$. Let $\tilde{\mathcal{V}} = \mathcal{V} \cap \bigcap_{i=1}^t H_i$. Suppose W is a k -subspace that evades $\tilde{\mathcal{V}}$. Then $W \cap \tilde{\mathcal{V}} = (W \cap \mathcal{V}) \cap \bigcap_{i=1}^t H_i = \emptyset$. Again by Lemma 2.8, we have $\dim(W \cap \mathcal{V}) \leq t - 1 = \dim(\mathcal{V}) + \dim(W) - n$. So W also evades \mathcal{V} . \square

Proof of Lemma 3.1. Consider a projective subvariety $\mathcal{V} \subseteq \mathbb{P}^n$ of degree at most d . Let $\mathcal{V}_1, \dots, \mathcal{V}_s$ be the irreducible components of \mathcal{V} . For each $i \in [s]$, use Claim 3.2 to choose a projective subvariety $\tilde{\mathcal{V}}_i \subseteq \mathbb{P}^n$ of dimension $n - k - 1$ and degree at most $\deg(\mathcal{V}_i)$ such that any k -subspace that evades $\tilde{\mathcal{V}}_i$ also evades \mathcal{V}_i . Let $\tilde{\mathcal{V}} = \bigcup_{i=1}^s \tilde{\mathcal{V}}_i$. Then $\tilde{\mathcal{V}} \in \mathcal{F}$. By construction, any k -subspace that evades $\tilde{\mathcal{V}}$ also evades \mathcal{V} . It follows that an (\mathcal{F}, ϵ) -evasive k -subspace family is also (n, d, ϵ) -evasive. \square

We further reduce to the case of irreducible varieties at the cost of blowing up the parameter ϵ by a factor of d . This is useful as we need irreducibility later in Lemma 3.10.

Lemma 3.3. *Let $n, d \in \mathbb{N}^+$ and $k \in \{0, 1, \dots, n - 1\}$. Let \mathcal{F}' be the family of all irreducible subvarieties of \mathbb{P}^n of dimension $n - k - 1$ and degree at most d . Then an (\mathcal{F}', ϵ) -evasive k -subspace family is also an $(n, d, d\epsilon)$ -evasive k -subspace family.*

Proof. Let \mathcal{F} be as in Lemma 3.1. Each $\mathcal{V} \in \mathcal{F}$ has at most d irreducible components, which are all in \mathcal{F}' since their degrees are bounded by d . By definition and the union bound, if a k -subspace family \mathcal{H} is (\mathcal{F}', ϵ) -evasive, then it is also $(\mathcal{F}, d\epsilon)$ -evasive. Combining this with Lemma 3.1 proves the lemma. \square

3.2 Chow Forms

By Lemma 3.1 and Lemma 3.3, we only need to evade equidimensional or irreducible projective subvarieties of codimension $k + 1$. The “bad” k -subspaces that intersect such a variety \mathcal{V} form a hypersurface of the Grassmannian defined by a single form called the *Chow form* of \mathcal{V} . We now explain the basic theory of Chow forms.

Grassmannians. Let $n \in \mathbb{N}$ and $k \in \{0, 1, \dots, n - 1\}$. The *Grassmannian* $G(k + 1, n + 1)$ is the set of all $(k + 1)$ -dimensional linear subspaces of \mathbb{A}^{n+1} . By taking the quotient modulo scalars, it may also be identified with the set of all k -subspaces of \mathbb{P}^n , which we denote by $G(k, n)$.

The Plücker embedding and Plücker coordinates. Consider a linear subspace $W \in G(k + 1, n + 1)$. The simplest way of representing W is using a $(k + 1) \times (n + 1)$ matrix A over \mathbb{F} such that W equals the row space of A . We call such a matrix A a *generating matrix* of W . For convenience, we also say A is a generating matrix of $\mathbb{P}W \in G(k, n)$.

The entries of A are called the (*primal*) *Stiefel coordinates* of W . However, note that A is not uniquely determined by W since for any $(k + 1) \times (k + 1)$ invertible matrix M over \mathbb{F} , the matrix MA is also a generating matrix of W .

Another way of representing W is using the vector $(\det A_{[k+1],S})_{S \in \binom{[n+1]}{k+1}}$ of maximal minors of a generating matrix A of W . For a $(k + 1) \times (k + 1)$ invertible matrix M over \mathbb{F} , replacing A by MA corresponds to multiplying all the maximal minors $\det A_{[k+1],S}$ by $\det M \in \mathbb{F}^\times$. To remove ambiguity, we could view $(\det A_{[k+1],S})_{S \in \binom{[n+1]}{k+1}}$ as a point in the projective space $\mathbb{P}^{\binom{[n+1]}{k+1}-1}$, which is then uniquely determined by W . This leads to the definition of the *Plücker embedding*.

Definition 3.4 (Plücker embedding). *Define $\phi : G(k + 1, n + 1) \rightarrow \mathbb{P}^{\binom{[n+1]}{k+1}-1}$ by*

$$\phi(W) = (\det A_{[k+1],S})_{S \in \binom{[n+1]}{k+1}}$$

where A is a generating matrix of W .

The Plücker embedding embeds the Grassmannian $G(k + 1, n + 1)$ in $\mathbb{P}^{\binom{[n+1]}{k+1}-1}$ as an irreducible projective subvariety, as stated by the following theorem. See, e.g., [Har13, Ful97] for proofs.

Theorem 3.5. *The Plücker embedding ϕ is a well-defined injective map whose image is an irreducible projective subvariety of $\mathbb{P}^{\binom{[n+1]}{k+1}-1}$.*

The homogeneous coordinates $(\det A_{[k+1],S})_{S \in \binom{[n+1]}{k+1}}$ of $\phi(W)$ are called the (*primal*) *Plücker coordinates* of W .

Denote by $R := \mathbb{F}[X_S : S \in \binom{[n+1]}{k+1}]$ the homogeneous coordinate ring of $\mathbb{P}^{\binom{[n+1]}{k+1}-1}$. The irreducible projective subvariety $\phi(G(k + 1, n + 1))$ is defined by a homogeneous prime ideal of R , which we denoted by I . Then R/I is the homogeneous coordinate ring of $\phi(G(k + 1, n + 1))$. The ideal I contains precisely the polynomial relations that the Plücker coordinates need to satisfy. It is also known that I is generated by certain quadratic forms, known as the *Plücker relations*. See [Har13, Ful97] for details.

Dual Plücker coordinates. Alternatively, we could represent a linear subspace $W \in G(k + 1, n + 1)$ by an $(n - k) \times (n + 1)$ matrix B over \mathbb{F} whose rows specify the linear equations defining W . We call such a matrix B a *parity check matrix* of W . For convenience, we also say B is a parity check matrix of $\mathbb{P}W \in \mathbb{G}(k, n)$.

The entries of B are called the *dual Stiefel coordinates* of W . This gives another embedding $\phi^\vee : G(k + 1, n + 1) \rightarrow \mathbb{P}^{\binom{[n+1]}{n-k}-1} = \mathbb{P}^{\binom{[n+1]}{k+1}-1}$, defined by

$$\phi^\vee(W) = (\det B_{[n-k],S})_{S \in \binom{[n+1]}{n-k}}. \quad (1)$$

The homogeneous coordinates $(\det B_{[n-k],S})_{S \in \binom{[n+1]}{n-k}}$ of $\phi^\vee(W)$ are called the *dual Plücker coordinates* of W .⁶ In fact, it is known that dual Plücker coordinates are equivalent to primal Plücker coordinates. Namely, if $W \in G(k + 1, n + 1)$ has primal Plücker coordinates $(c_S)_{S \in \binom{[n+1]}{k+1}}$, then it has dual Plücker coordinates $(c'_S)_{S \in \binom{[n+1]}{n-k}}$ with $c'_S = (-1)^{\sum_{i \in S} i - \sum_{i \in [k+1]} i} \cdot c_{[n+1] \setminus S}$ (see, e.g., [JT13]).

⁶Some authors use “primal” and “dual” in the opposite way (e.g., [DS95]).

Chow forms. Recall that we denote by $\mathbb{G}(k, n)$ the set of all k -subspaces of \mathbb{P}^n . By identifying $\mathbb{G}(k + 1, n + 1)$ with $\mathbb{G}(k, n)$ via $W \mapsto \mathbb{P}W$, we regard ϕ and ϕ^\vee as maps from $\mathbb{G}(k, n)$ to $\mathbb{P}^{\binom{n+1}{k+1}-1}$.

We also need the notion of *associated hypersurfaces*.

Definition 3.6 (Associated hypersurface [GKZ94]). *For an irreducible subvariety $\mathcal{V} \subseteq \mathbb{P}^n$ of dimension $n - k - 1$, define the associated hypersurfaces $\mathcal{Z}_\mathcal{V}$ of \mathcal{V} to be the set of k -subspaces intersecting \mathcal{V} , i.e.,*

$$\mathcal{Z}_\mathcal{V} := \{W \in \mathbb{G}(k, n) : \mathcal{V} \cap W \neq \emptyset\}.$$

The term ‘‘associated hypersurface’’ is justified by the following theorem.

Theorem 3.7. *Let $\mathcal{V} \subseteq \mathbb{P}^n$ be an irreducible projective subvariety of dimension $n - k - 1$ and degree $d \in \mathbb{N}^+$. Then there exists a nonzero homogeneous polynomial $P_\mathcal{V} \in R = \mathbb{F}[X_S : S \in \binom{[n+1]}{k+1}]$ of degree d such that $\phi(\mathcal{Z}_\mathcal{V})$ is defined by $P_\mathcal{V}$ as a subvariety of $\phi(\mathbb{G}(k, n))$. That is,*

$$\phi(\mathcal{Z}_\mathcal{V}) = \phi(\mathbb{G}(k, n)) \cap \mathcal{V}(P_\mathcal{V}).$$

Moreover, $\mathcal{R}_\mathcal{V} := P_\mathcal{V} + I \in (R/I)_d$ is uniquely determined by \mathcal{V} up to scalars.

Theorem 3.7 is explicitly stated as [DS95, Theorem 1.1 and Corollary 2.1]. A proof can be found in [GKZ94, Section 3.2]. We briefly explain how to find a polynomial $P_\mathcal{V}$ satisfying Theorem 3.7: Firstly, it can be shown using the trick of dimension counting via incidence varieties that $\phi(\mathcal{Z}_\mathcal{V})$ is an irreducible projective subvariety of the Grassmannian $\phi(\mathbb{G}(k, n))$ of codimension one [GKZ94, Section 3.2, Proposition 2.2]. Secondly, the homogeneous coordinate ring R/I of the Grassmannian is known to be a *unique factorization domain* [Ful97, Chapter 9]. These two facts imply that the homogeneous ideal of R/I defining $\phi(\mathcal{Z}_\mathcal{V})$ is a *principal ideal*. Choose $\mathcal{R}_\mathcal{V}$ to be a generator of this principal ideal, which is unique up to scalars. Then lift $\mathcal{R}_\mathcal{V} \in R/I$ to $P_\mathcal{V} \in R$.

Now we are ready to define the Chow form of projective subvarieties.

Definition 3.8 (Chow form). *Let $\mathcal{V} \subseteq \mathbb{P}^n$ be an irreducible subvariety of dimension $n - k - 1$ and degree $d \in \mathbb{N}^+$. Define the Chow form of \mathcal{V} in Plücker coordinates, or simply the Chow form of \mathcal{V} , to be $\mathcal{R}_\mathcal{V} \in (R/I)_d$ as in Theorem 3.7.*

More generally, for an equidimensional subvariety $\mathcal{V} = \bigcup_{i=1}^s \mathcal{V}_i \subseteq \mathbb{P}^n$ of dimension $n - k - 1$ and degree d , where $\mathcal{V}_1, \dots, \mathcal{V}_s$ are the irreducible components of \mathcal{V} , the Chow form of \mathcal{V} is $\mathcal{R}_\mathcal{V} := \prod_{i=1}^s \mathcal{R}_{\mathcal{V}_i} \in (R/I)_d$. It is uniquely determined by \mathcal{V} up to scalars.

As a k -subspace intersects $\mathcal{V} = \bigcup_{i=1}^s \mathcal{V}_i$ iff it intersects some \mathcal{V}_i , we see from Theorem 3.7 that the Chow form $\mathcal{R}_\mathcal{V}$ of an equidimensional projective subvariety \mathcal{V} of dimension $n - k - 1$ vanishes precisely at the set of k -subspaces that intersect \mathcal{V} .

Example 1. Let $k = 0$. Let $\mathcal{V} \subseteq \mathbb{P}^n$ be a hypersurface defined by a nonzero homogeneous polynomial $P \in \mathbb{F}[X_1, \dots, X_{n+1}] = R$. The ideal I of R is zero in this case. And the Chow form $\mathcal{R}_\mathcal{V}$ of \mathcal{V} is simply P (up to a scalar).

Example 2. Let $V \in \mathbb{G}(n - k, n + 1)$ and $W \in \mathbb{G}(k + 1, n + 1)$. Choose matrices $A, B \in \mathbb{F}^{(k+1) \times (n+1)}$ such that A is a generating matrix of W and B is a parity check matrix of V . Then $\mathbb{P}V \cap \mathbb{P}W \neq \emptyset$ iff $\dim(V \cap W) > 0$, which holds iff $\det(AB^T) = 0$. On the other hand, we have

$$\det(AB^T) = \sum_{S \in \binom{[n+1]}{k+1}} \det(A_{[k+1], S}) \cdot \det((B^T)_{S, [k+1]}) = \sum_{S \in \binom{[n+1]}{k+1}} \det(A_{[k+1], S}) \cdot \det(B_{[k+1], S}),$$

where the first equation is known as the *Cauchy–Binet formula* (see, e.g., [FSS14]). So $P_{\mathbb{P}V} \in R_1$ is a linear polynomial whose coefficients are given by the dual Plücker coordinates $(\det B_{[k+1], S})_{S \in \binom{[n+1]}{k+1}}$ of V (up to a scalar). The degree-one part I_1 of I is zero as I is generated by quadratic forms. So the Chow form $\mathcal{R}_{\mathbb{P}V} \in (R/I)_1 = R_1$ is simply $P_{\mathbb{P}V}$.

Chow forms in Stiefel coordinates. We may also express the Chow form in Stiefel coordinates, i.e., in the entries of a generating matrix of a linear subspace. This expression has the advantage that it is an actual polynomial rather than a member of the abstract vector space $(R/I)_d$.

Formally, let A^* be a $(k+1) \times (n+1)$ variable matrix whose (i, j) -th entry is a variable $Y_{i,j}$. Define the ring homomorphism

$$\phi^\sharp : R = \mathbb{F} \left[X_S : S \in \binom{[n+1]}{k+1} \right] \rightarrow \mathbb{F}[Y_{i,j} : i \in [k+1], j \in [n+1]]$$

that sends each variable X_S to $\det(A^*_{[k+1],S})$. Define the *Chow form of \mathcal{V} in Stiefel coordinates* to be

$$\tilde{\mathcal{R}}_{\mathcal{V}} := \phi^\sharp(P_{\mathcal{V}}) \in \mathbb{F}[Y_{i,j} : i \in [k+1], j \in [n+1]]$$

where $P_{\mathcal{V}} \in R_d$ is a lift of $\mathcal{R}_{\mathcal{V}} \in (R/I)_d$. Note that I is precisely the kernel of ϕ^\sharp . So $\tilde{\mathcal{R}}_{\mathcal{V}}$ is uniquely determined by \mathcal{V} up to scalars. By construction, for any $W \in \mathbb{G}(k+1, n+1)$ and generating matrix $A = (a_{i,j})_{i \in [k+1], j \in [n+1]}$ of W , we have $P_{\mathcal{V}}(\phi(W)) = \tilde{\mathcal{R}}_{\mathcal{V}}(A) := \tilde{\mathcal{R}}_{\mathcal{V}}(a_{1,1}, \dots, a_{k+1, n+1})$. So $\tilde{\mathcal{R}}_{\mathcal{V}}$ vanishes at A iff $\mathbb{P}W \in \mathbb{G}(k, n)$ intersects \mathcal{V} .

Chow forms in dual Stiefel coordinates. Similarly, we may express the Chow form in dual Stiefel coordinates, i.e., in the entries of a parity check matrix of a linear subspace.

More specifically, choose a homogeneous polynomial $Q_{\mathcal{V}} \in \mathbb{F} \left[X_S : S \in \binom{[n+1]}{n-k} \right]$ that defines the set of k -subspaces intersecting \mathcal{V} in terms of dual Plücker coordinates. As primal and dual Plücker coordinates are equivalent, $Q_{\mathcal{V}}$ can be obtained from the polynomial $P_{\mathcal{V}}$ above by simply negating and renaming variables. Next, compose $Q_{\mathcal{V}}$ with a ring homomorphism that substitutes dual Plücker coordinates with dual Stiefel coordinates. The resulting polynomial, which we denote by $\tilde{\mathcal{R}}_{\mathcal{V}}^\vee \in \mathbb{F}[Y_{i,j} : i \in [n-k], j \in [n+1]]$, is called the *Chow form of \mathcal{V} in dual Stiefel coordinates*.

We note that the Chow form $\tilde{\mathcal{R}}_{\mathcal{V}}$ in primal Stiefel coordinates is a homogeneous polynomial of degree $(k+1)d$ in $(k+1)(n+1)$ variables, whereas the Chow form $\tilde{\mathcal{R}}_{\mathcal{V}}^\vee$ in dual Stiefel coordinates is a homogeneous polynomial of degree $(n-k)d$ in $(n-k)(n+1)$ variables. This suggests that it is more convenient to use the Chow form in primal (resp. dual) Stiefel coordinates when k is small (resp. $n-k$ is small).⁷

3.3 Explicit Constructions of Variety Evasive Subspace Families

Let $n, d \in \mathbb{N}^+$, $k \in \{0, 1, \dots, n\}$, and $\epsilon \in (0, 1)$. In this subsection, we prove the Main Theorem (Theorem 1.6) by constructing explicit projective or affine k -subspace families that are (n, d, ϵ) -evasive. The problem is trivial when $k = n$, as we just need to choose the singleton $\{\mathbb{P}^n\}$ or $\{\mathbb{A}^n\}$. So assume $k < n$.

We first prove Theorem 1.6 in the projective case, and then derive the affine case from it by viewing \mathbb{A}^n as an open subset of \mathbb{P}^n . For the projective case, we present two constructions. The first one is simple and only uses ϵ -hitting sets for low degree polynomials (Lemma 2.3). But the size of the resulting subspace family is polynomial only when both d and k (or $n-k$) are bounded. Next, we give a more sophisticated construction, which yields subspace families of polynomial size as long as d is bounded.

3.3.1 Simple Construction

We first present a simple construction of (n, d, ϵ) -evasive k -subspace families on \mathbb{P}^n .

First assume $k+1 \leq n-k$. In this case, construct a k -subspace family \mathcal{H} on \mathbb{P}^n as follows:

⁷While both $\mathcal{R}_{\mathcal{V}}$ and $\mathcal{R}_{\mathcal{V}}^\vee$ may be viewed as elements of $(R/I)_d$, the two (injective) maps $\mathcal{R}_{\mathcal{V}} \mapsto \tilde{\mathcal{R}}_{\mathcal{V}}$ and $\mathcal{R}_{\mathcal{V}}^\vee \mapsto \tilde{\mathcal{R}}_{\mathcal{V}}^\vee$ come from different linear embeddings of $(R/I)_d$ in vector spaces of polynomials. As a result, the representation of \mathcal{V} by the polynomial $\tilde{\mathcal{R}}_{\mathcal{V}}$ and the representation by $\tilde{\mathcal{R}}_{\mathcal{V}}^\vee$ are not equally succinct in general.

1. Use Lemma 2.3 to compute an ϵ -hitting set T for the family of polynomials $f \in \mathbb{F}[Y_{i,j} : i \in [k+1], j \in [n+1]]$ of degree at most $(k+1)d$ such that $|T| = \text{poly}\left(\binom{(k+1)(n+1+d)}{(k+1)d}, 1/\epsilon\right)$. Think of T as a collection of $(k+1) \times (n+1)$ matrices over \mathbb{F} .
2. Initialize $\mathcal{H} = \emptyset$. For each matrix $A \in T$, if A has full row rank $k+1$, add to \mathcal{H} the k -subspace $W \in \mathbb{G}(k, n)$ with the generating matrix A .

Next, assume $k+1 > n-k$. In this case, construct \mathcal{H} in a similar way, but use parity check matrices instead of generating matrices. Namely, compute an ϵ -hitting set T for the family of polynomials $f \in \mathbb{F}[Y_{i,j} : i \in [n-k], j \in [n+1]]$ of degree at most $(n-k)d$ such that $|T| = \text{poly}\left(\binom{(n-k)(n+1+d)}{(n-k)d}, 1/\epsilon\right)$. Think of T as a collection of $(n-k) \times (n+1)$ matrices over \mathbb{F} . For each matrix $A \in T$, add to \mathcal{H} the k -subspace $W \in \mathbb{G}(k, n)$ with the parity check matrix A .

This construction does give an (n, d, ϵ) -evasive k -subspace family, as stated by the following lemma.

Lemma 3.9. *The k -subspace family \mathcal{H} constructed above is (n, d, ϵ) -evasive and has size polynomial in $\min\left\{\binom{(k+1)(n+1+d)}{(k+1)d}, \binom{(n-k)(n+1+d)}{(n-k)d}\right\}$ and $1/\epsilon$. Moreover, the total time complexity of computing the linear equations defining the k -subspaces in \mathcal{H} is polynomial in $|\mathcal{H}|$ (and $\log p$, if $\text{char}(\mathbb{F}) = p > 0$).*

Proof. We only show that \mathcal{H} is (n, d, ϵ) -evasive since the rest of the lemma is obvious from the construction. Let \mathcal{F} be the family of all equidimensional subvarieties of \mathbb{P}^n of dimension $n-k-1$ and degree at most d . By Lemma 3.1, it suffices to prove that \mathcal{H} is (\mathcal{F}, ϵ) -evasive. Consider any $\mathcal{V} \in \mathcal{F}$. We want to show that $\mathcal{V} \cap W = \emptyset$ for all but at most ϵ -fraction of $W \in \mathcal{H}$.

First assume $k+1 \leq n-k$. The Chow form $\tilde{\mathcal{R}}_{\mathcal{V}}$ of \mathcal{V} in Stiefel coordinates is a nonzero homogeneous polynomial in $\mathbb{F}[Y_{i,j} : i \in [k+1], j \in [n+1]]$ of degree $(k+1) \deg(\mathcal{V}) \leq (k+1)d$. By the choice of T , for all but at most ϵ -fraction of $A \in T$, we have $\tilde{\mathcal{R}}_{\mathcal{V}}(A) \neq 0$, which implies $\mathcal{V} \cap W = \emptyset$, where A is a generating matrix of W .

By construction, \mathcal{H} is the collection of k -subspaces corresponding to the matrices $A \in T$ of full row rank. So we have ignored the matrices that do not have full row rank. But this does not increase the fraction of “bad” $W \in \mathcal{H}$ since if A does not have full row rank, then the maximal minors of A are all zero, and $\tilde{\mathcal{R}}_{\mathcal{V}}(A)$ must be zero. It follows that $\mathcal{V} \cap W = \emptyset$ for all but at most ϵ -fraction of $W \in \mathcal{H}$, as desired.

Now assume $k+1 > n-k$. The proof in this case is similar and we omit the details. The only difference is that we use the Chow form $\tilde{\mathcal{R}}_{\mathcal{V}}^{\vee}$ in dual Stiefel coordinates instead of $\tilde{\mathcal{R}}_{\mathcal{V}}$. \square

3.3.2 Improved Construction

For a subvariety $\mathcal{V} \subseteq \mathbb{P}^n$, denote by $\text{span}(\mathcal{V})$ the smallest projective subspace that contains \mathcal{V} . We say \mathcal{V} is *nondegenerate* if it is not contained in a hyperplane of \mathbb{P}^n , or equivalently, $\text{span}(\mathcal{V}) = \mathbb{P}^n$.

We need the following fact from algebraic geometry (see, e.g., [EH87, Proposition 0] or [Har13, Corollary 18.12]).

Lemma 3.10. *The codimension of a nondegenerate irreducible subvariety \mathcal{V} of \mathbb{P}^n is at most $\deg(\mathcal{V}) - 1$.*

We now give an improved construction of (n, d, ϵ) -evasive k -subspace families on \mathbb{P}^n as follows.

1. If $\min\{k+1, n-k\} \leq d-1$, just use the previous simple construction. So assume $\min\{k+1, n-k\} > d-1$. Let $t = k-d+2$ and $\epsilon_0 = \epsilon/(2d)$.
2. Use Lemma 2.3 to construct an ϵ_0 -hitting set $T \subseteq \mathbb{F}^{(d-1)(n+1)}$ for the family of polynomials $f \in \mathbb{F}[Y_{i,j} : i \in [d-1], j \in [n+1]]$ of degree at most $(d-1)d$ such that $|T| = \text{poly}\left(\binom{(d-1)(n+1+d)}{(d-1)d}, d/\epsilon\right)$. Think of T as a collection of $(d-1) \times (n+1)$ matrices over \mathbb{F} .⁸

⁸When $d = 1$, just let T be the singleton $\mathbb{F}_q^{(d-1)(n+1)} = \mathbb{F}_q^0$, which consists of an “empty matrix”.

3. Use Corollary 2.5 to construct a collection U of $t \times (n+1)$ matrix over \mathbb{F} such that $|U| = \text{poly}(n, d/\epsilon)$ and for every $(n+1) \times t$ matrix M over \mathbb{F} of rank t , all but at most ϵ_0 -fraction of $B \in U$ satisfies $\text{rank}(BM) = t$.
4. Initialize $\mathcal{H} = \emptyset$. For each $(A, B) \in T \times U$, if the $(k+1) \times (n+1)$ matrix $\begin{pmatrix} A \\ B \end{pmatrix}$ has full row rank, add to \mathcal{H} the k -subspace $W \in \mathbb{G}(k, n)$ with the generating matrix $\begin{pmatrix} A \\ B \end{pmatrix}$.

See below for an illustration of a matrix $\begin{pmatrix} A \\ B \end{pmatrix}$, where $(A, B) \in T \times U$.

$$\begin{array}{c}
 \overbrace{\hspace{10em}}^{n+1} \\
 d-1 \left\{ \left(\begin{array}{c} A \\ \text{-----} \\ B \end{array} \right) \right\} \\
 t \left\{ \hspace{10em} \right\} k+1
 \end{array}$$

We use the construction above to prove the Main Theorem (Theorem 1.6) in the projective case. For convenience, we restate it in the following form.

Theorem 3.11 (Main Theorem in the projective case). *The k -subspace family \mathcal{H} constructed above is (n, d, ϵ) -evasive and has size $\text{poly}(N(k, d, n), n, 1/\epsilon)$. Moreover, the total time complexity of computing the linear equations defining the k -subspaces in \mathcal{H} is polynomial in $|\mathcal{H}|$ (and $\log p$, if $\text{char}(\mathbb{F}) = p > 0$).*

Proof. The theorem follows from Lemma 3.9 if $\min\{k+1, n-k\} \leq d-1$. So assume $\min\{k+1, n-k\} > d-1$ and hence $t \geq 1$. We only show that \mathcal{H} is (n, d, ϵ) -evasive since the rest of the theorem is obvious from the construction.

Let \mathcal{F} be the family of all irreducible subvarieties of \mathbb{P}^n of dimension $n-k-1$ and degree at most d . By Lemma 3.3, it suffices to prove that \mathcal{H} is $(\mathcal{F}, 2\epsilon_0)$ -evasive. Consider any $\mathcal{V} \in \mathcal{F}$. We want to show that $\mathcal{V} \cap W = \emptyset$ for all but at most $(2\epsilon_0)$ -fraction of $W \in \mathcal{H}$.

By definition, \mathcal{V} is a nondegenerate irreducible subvariety of $\text{span}(\mathcal{V})$. By Lemma 3.10, the codimension of \mathcal{V} in $\text{span}(\mathcal{V})$ is at most $d-1$. Therefore,

$$\dim(\text{span}(\mathcal{V})) \leq \dim(\mathcal{V}) + d - 1 = (n - k - 1) + (d - 1) = n - t.$$

Let $\Lambda \subseteq \mathbb{P}^n$ be an $(n-t)$ -subspace that contains $\text{span}(\mathcal{V})$. Let $M \in \mathbb{F}^{t \times (n+1)}$ be a parity check matrix of Λ . By the choice of U , all but at most ϵ_0 -fraction of $B \in U$ satisfies $\text{rank}(BM) = t$. Fix $B \in U$ such that $\text{rank}(BM) = t$. Let $W_0 \in \mathbb{G}(t-1, n)$ such that B is a generating matrix of W_0 . The condition $\text{rank}(BM) = t$ is equivalent to $W_0 \cap \Lambda = \emptyset$.

We make the following claim.

Claim 3.12. *For all but ϵ_0 -fraction of $A \in T$, the matrix $\begin{pmatrix} A \\ B \end{pmatrix}$ is a generating matrix of a k -subspace $W \in \mathbb{G}(k, n)$ that is disjoint from \mathcal{V} .*

Note that Claim 3.12 implies that $\mathcal{V} \cap W = \emptyset$ holds for all but at most $(2\epsilon_0)$ -fraction of $W \in \mathcal{H}$. So it remains to prove this claim.

A matrix $\begin{pmatrix} A \\ B \end{pmatrix}$ is a generating matrix of a k -subspace disjoint from \mathcal{V} as long as $\tilde{\mathcal{R}}_{\mathcal{V}}\left(\begin{pmatrix} A \\ B \end{pmatrix}\right) \neq 0$, where $\tilde{\mathcal{R}}_{\mathcal{V}} \in \mathbb{F}[Y_{i,j} : i \in [k+1], j \in [n+1]]$ is the Chow form of \mathcal{V} in Stiefel coordinates. Consider the polynomial

$$P = \tilde{\mathcal{R}}_{\mathcal{V}}\left(\begin{pmatrix} \cdot \\ B \end{pmatrix}\right) \in \mathbb{F}[Y_{i,j} : i \in [d-1], j \in [n+1]]$$

which is obtained from $\tilde{\mathcal{R}}_{\mathcal{V}}$ by assigning the $t \times (n+1)$ entries of B to the variables $Y_{d,1}, \dots, Y_{k+1,n+1}$ on the bottom t rows, with the top $d-1$ rows of variables $Y_{1,1}, \dots, Y_{d-1,n+1}$ left free.

As $W_0 \cap \Lambda = \emptyset$ and $\text{span}(\mathcal{V}) \subseteq \Lambda$, we know W_0 is disjoint from \mathcal{V} . By Lemma 2.7, W_0 extends to a k -subspace that is disjoint from \mathcal{V} . So the generating matrix B of W_0 extends to a matrix $\begin{pmatrix} A \\ B \end{pmatrix}$ such that $\tilde{\mathcal{R}}_{\mathcal{V}}\left(\begin{pmatrix} A \\ B \end{pmatrix}\right) \neq 0$. In particular, the polynomial P is not identically zero. Also note $\deg(P) = (d-1) \deg(\mathcal{V}) \leq (d-1)d$. By the choice of T , for all but ϵ_0 -fraction of $A \in T$, we have $\tilde{\mathcal{R}}_{\mathcal{V}}\left(\begin{pmatrix} A \\ B \end{pmatrix}\right) = P(A) \neq 0$, and hence $\begin{pmatrix} A \\ B \end{pmatrix}$ is a generating matrix of a k -subspace that is disjoint from \mathcal{V} . This proves Claim 3.12 and completes the proof of the theorem. \square

3.3.3 The Affine Case

In this subsection, we prove Theorem 1.6 in the affine case. Recall that we may view \mathbb{A}^n as an open subset of \mathbb{P}^n via the map $(x_1, \dots, x_n) \mapsto (x_1, \dots, x_n, 1)$. In this way, \mathbb{P}^n becomes the disjoint union of \mathbb{A}^n and the *hyperplane at infinity* H_∞ defined by $X_{n+1} = 0$.

We use the following lemma to reduce the affine case to the projective case.

Lemma 3.13. *Let $n, d \in \mathbb{N}^+$, $k \in \{0, 1, \dots, n-1\}$, and $\epsilon' \in (0, 1/2)$. Suppose \mathcal{H} is an (n, d, ϵ') -evasive k -subspace family on \mathbb{P}^n . Then*

$$\mathcal{H}' = \{W \cap \mathbb{A}^n : W \in \mathcal{H}, W \not\subseteq H_\infty\}$$

is an (n, d, ϵ) -evasive affine k -subspace family on \mathbb{A}^n , where $\epsilon = \epsilon'/(1 - \epsilon') \leq 2\epsilon'$. Moreover,

$$\mathcal{H}'' = \{W \in \mathcal{H} : W \not\subseteq H_\infty\} = \{W_{\text{cl}} : W \in \mathcal{H}'\}$$

is an (n, d, ϵ) -evasive k -subspace family on \mathbb{P}^n .

Proof. By (n, d, ϵ') -evasiveness of \mathcal{H} , at most ϵ' -fraction of $W \in \mathcal{H}$ are fully contained in H_∞ . Throwing away those k -subspaces fully contained in H_∞ increases the error parameter ϵ' by at most a factor of $1/(1 - \epsilon')$. Therefore, $\mathcal{H}'' = \{W \in \mathcal{H} : W \not\subseteq H_\infty\}$ is (n, d, ϵ) -evasive. We want to prove that $\mathcal{H}' = \{W \cap \mathbb{A}^n : W \in \mathcal{H}''\}$ is also (n, d, ϵ) -evasive.

Consider a subvariety $\mathcal{V} \subseteq \mathbb{A}^n$ of degree at most d . Let $\mathcal{V}_1, \dots, \mathcal{V}_s$ be the irreducible components of \mathcal{V} . The projective closure \mathcal{V}_{cl} of \mathcal{V} has the irreducible components $(\mathcal{V}_1)_{\text{cl}}, \dots, (\mathcal{V}_s)_{\text{cl}}$. Consider a k -subspace $W \in \mathcal{H}''$ that evades \mathcal{V}_{cl} . We just need to prove that $W \cap \mathbb{A}^n$ evades \mathcal{V} . This is true since for each $i \in [s]$,

$$\dim((W \cap \mathbb{A}^n) \cap \mathcal{V}_i) \leq \dim(W \cap (\mathcal{V}_i)_{\text{cl}}) \leq \dim(W) + \dim((\mathcal{V}_i)_{\text{cl}}) - n = \dim(W \cap \mathbb{A}^n) + \dim(\mathcal{V}_i) - n$$

where the second inequality holds since W evades \mathcal{V}_{cl} and the last equality uses the fact $W \not\subseteq H_\infty$. \square

The affine case of Theorem 1.6 now follows easily.

Proof of Theorem 1.6 in the affine case. If $k = n$, just choose $\mathcal{H} = \mathbb{A}^n$. Now assume $k < n$. Construct an $(n, d, \epsilon/2)$ -evasive k -subspace family \mathcal{H} on \mathbb{P}^n using Theorem 3.11. Then

$$\mathcal{H}' := \{W \cap \mathbb{A}^n : W \in \mathcal{H}, W \not\subseteq H_\infty\}$$

is an (n, d, ϵ) -evasive affine k -subspace family on \mathbb{A}^n by Lemma 3.13. The nonhomogeneous linear equations defining $W \cap \mathbb{A}^n \in \mathcal{H}'$ can be easily computed from the homogeneous linear equations defining $W \in \mathcal{H}$ by letting $X_{n+1} = 1$. \square

The proof of Theorem 1.6 is now complete.

Strengthening Theorem 1.6 in the affine case. For projective subvarieties $\mathcal{V}_1, \mathcal{V}_2 \subseteq \mathbb{P}^n$ such that $\dim(\mathcal{V}_1) + \dim(\mathcal{V}_2) \geq n$, the minimum possible dimension of $\mathcal{V}_1 \cap \mathcal{V}_2$ is $\dim(\mathcal{V}_1) + \dim(\mathcal{V}_2) - n$, as stated by Lemma 2.9. Nevertheless, for two affine subvarieties $\mathcal{V}_1, \mathcal{V}_2 \subseteq \mathbb{A}^n$, it is possible that the intersection of \mathcal{V}_1 and \mathcal{V}_2 is empty even if its expected dimension $\dim(\mathcal{V}_1) + \dim(\mathcal{V}_2) - n$ is nonnegative. For example, the intersection of two distinct and parallel affine hyperplanes $\mathcal{V}_1, \mathcal{V}_2 \subseteq \mathbb{A}^n$ is always empty even if $n \geq 2$. The reason this happens is that, while the dimension of $(\mathcal{V}_1)_{\text{cl}} \cap (\mathcal{V}_2)_{\text{cl}}$ is $n - 2$ (as expected), this intersection is fully contained in the hyperplane H_∞ , which is excluded from \mathbb{A}^n .

One may strengthen the definition of evading (Definition 1.1) by requiring the intersection of \mathcal{V}_1 with every irreducible component of \mathcal{V}_2 to have *exactly* the expected dimension. It is possible to construct explicit affine k -subspace families satisfying Theorem 1.6 even under this stronger definition of evading. We sketch the ideas as follows but omit the details.

First construct an $(n - 1, d, \epsilon')$ -evasive $(k - 1)$ -subspace family \mathcal{H}' on $H_\infty \cong \mathbb{P}^{n-1}$ for some sufficiently small ϵ' depending on ϵ . Then extend each $W \in \mathcal{H}'$ to a collection of k -subspaces by picking $p \in \mathbb{A}^n$ and taking the k -subspace $J(W, p)$, where the coordinates of p are chosen from an ϵ' -hitting set for polynomials of degree at most d given by Lemma 3.9. Call the resulting k -subspace family \mathcal{H} . It is easy to prove that \mathcal{H} is $(n, d, O(\epsilon'))$ -evasive.

Furthermore, the affine k -subspace family $\{W \cap \mathbb{A}^n : W \in \mathcal{H}\}$ is (n, d, ϵ) -evasive even under the stronger definition of evading. To see this, consider an affine subvariety $\mathcal{V} \subseteq \mathbb{A}^n$ of degree at most d . For most $W \in \mathcal{H}$, we have:

- For each irreducible component \mathcal{V}_i of \mathcal{V} , the dimension of $(\mathcal{V}_i)_{\text{cl}} \cap W$ is as expected by $(n, d, O(\epsilon'))$ -evasiveness of \mathcal{H} and Lemma 2.9. Call this dimension d_i , which is $-\infty$ if $(\mathcal{V}_i)_{\text{cl}} \cap W = \emptyset$.
- Moreover, the dimension of $((\mathcal{V}_i)_{\text{cl}} \cap H_\infty) \cap (W \cap H_\infty)$ is at most $d_i - 1$ by $(n - 1, d, \epsilon')$ -evasiveness of \mathcal{H}' .
- Therefore, $\mathcal{V}_i \cap (W \cap \mathbb{A}^n)$ has the expected dimension d_i for each irreducible component \mathcal{V}_i of \mathcal{V} .

4 Lower Bound

We prove Theorem 1.7 in this section. The main tool is the notion of *Chow varieties*, which parameterize projective subvarieties. More precisely, they parametrize a generalization of projective subvarieties, called (*effective*) *algebraic cycles* on a projective space.

Algebraic cycles. An *algebraic r -cycle* (or simply *r -cycle*) on \mathbb{P}^n is a formal linear combination $\sum c_i \mathcal{V}_i$ of finitely many irreducible subvarieties $\mathcal{V}_i \subseteq \mathbb{P}^n$ of dimension r , where the coefficients c_i are integers. The *degree* of $\sum c_i \mathcal{V}_i$ is $\sum c_i \deg(\mathcal{V}_i)$. An r -cycle is *effective* if all its coefficients are nonnegative. Denote by $C(r, d, n)$ the set of all effective r -cycles of degree d on \mathbb{P}^n .

Chow varieties. Let $k \in \{0, 1, \dots, n - 1\}$ and $r = n - k - 1$. The definition of Chow forms naturally extends to effective r -cycles. Namely, for an effective r -cycle $D = \sum_{i=1}^r c_i \mathcal{V}_i$ of degree d on \mathbb{P}^n , define the Chow form of D to be $\mathcal{R}_D := \prod_{i=1}^r \mathcal{R}_{\mathcal{V}_i}^{c_i}$.

Note that \mathcal{R}_D is a vector in $(R/I)_d$ and is uniquely determined by D up to scalars. Write $[\mathcal{R}_D]$ for the point in $\mathbb{P}(R/I)_d$ represented by \mathcal{R}_D . Then we have map $\psi : C(r, d, n) \rightarrow \mathbb{P}(R/I)_d$, given by

$$\psi : D \mapsto [\mathcal{R}_D],$$

called the *Chow embedding* of $C(r, d, n)$. Indeed, it embeds $C(r, d, n)$ in $\mathbb{P}(R/I)_d$ as a projective subvariety, as stated by the following theorem of Chow and van der Waerden [CvdW37].

Theorem 4.1 ([CvdW37]). *The map ψ is injective and its image is Zariski-closed.*

A proof can also be found in [GKZ94, Chapter 4]. We identify $C(r, d, n)$ with its image under ψ and view it as a projective variety. This variety is called the *Chow variety* of effective r -cycles of degree d on \mathbb{P}^n .

Example 3. Let V be the subspace of homogeneous polynomials in $\mathbb{F}[X_1, \dots, X_{n+1}]$ of degree d . Then $C(n-1, d, n)$ is simply the projective space $\mathbb{P}V$ (see Example 1).

Example 4. $C(r, 1, n)$ is the Grassmannian $G(r+1, n+1)$ (or $\mathbb{G}(r, n)$) embedded in $\mathbb{P}^{\binom{n+1}{r+1}-1} = \mathbb{P}^{\binom{n+1}{k+1}-1}$ via ϕ^\vee (see Example 2).

The dimension of Chow varieties. When $d = 1$, the Chow variety $C(r, d, n)$ is just the Grassmannian $G(r+1, n+1)$ (see Example 4) and its dimension is well known to be $(r+1)(n-r)$ [Har13]. When $d > 1$, the dimension of $C(r, d, n)$ was determined by Azcue in his Ph.D. thesis [Azc92] and independently by Lehmann [Leh17]. We state their result as follows.

Theorem 4.2 ([Azc92, Leh17]). *For $d > 1$ and $0 \leq r < n$, the dimension of $C(r, d, n)$ is*

$$\max \left\{ d(r+1)(n-r), \binom{d+r+1}{r+1} - 1 + (r+2)(n-r-1) \right\}.$$

This theorem was previously proved by Eisenbud and Harris [EH92] for the special case $r = 1$.

Remark. To prove Theorem 1.7, we only need a lower bound for the dimension of the Chow variety, which is much easier to prove than Theorem 4.2. Indeed, it is not difficult to see that $d(r+1)(n-r)$ is the dimension of the space of unions of d r -subspaces of \mathbb{P}^n , and $\binom{d+r+1}{r+1} - 1 + (r+2)(n-r-1)$ is the dimension of the space of degree- d hypersurfaces in $(r+1)$ -subspaces of \mathbb{P}^n .

Lower bound via dimension counting. We now restate Theorem 1.7 and prove it using a dimension counting argument.

Theorem 1.7. *Let $n, d \in \mathbb{N}^+$ and $k \in \{0, 1, \dots, n-1\}$. Let \mathcal{F} be the family of equidimensional projective subvarieties of \mathbb{P}^n of dimension $n-k-1$ and degree at most d . Suppose \mathcal{H} is an \mathcal{F} -evasive k -subspace family on \mathbb{P}^n . Then*

$$|\mathcal{H}| \geq \begin{cases} (n-k)(k+1) + 1 & \text{if } d = 1, \\ \max \left\{ d(n-k)(k+1) + 1, \binom{d+n-k}{d} + (n-k+1)k \right\} & \text{if } d > 1. \end{cases}$$

In particular, $|\mathcal{H}|$ is superpolynomial in n when $n-k = \Omega(n)$ and $d = \omega(1)$.

Proof. Consider an arbitrary k -subspace $W \in \mathcal{H}$. We may think of each point in $\mathbb{P}(R/I)_d$ as a homogeneous polynomial of degree d in Plücker coordinates modulo scalars and the ideal I of Plücker relations. We know Plücker coordinates always satisfy the Plücker relations. So it makes sense to talk about if a point in $\mathbb{P}(R/I)_d$ vanishes at $\phi(W)$ or not, as it does not depend on the choice of the homogeneous polynomial representing this point. Note that the constraint of $p \in \mathbb{P}(R/I)_d$ vanishing at $\phi(W)$ is a linear equation in the homogeneous coordinates of p . So the set of points in $\mathbb{P}(R/I)_d$ vanishing at $\phi(W)$ is a hyperplane of $\mathbb{P}(R/I)_d$, which we denote by H_W .

Let $r = n - k - 1$. Assume $|\mathcal{H}| \leq \dim(C(r, d, n))$. Then we have

$$\psi(C(r, d, n)) \cap \bigcap_{W \in \mathcal{H}} H_W \neq \emptyset$$

since taking the intersection with a hyperplane reduces the dimension of a projective subvariety by at most one (Lemma 2.8 or Lemma 2.9). So there exists an effective r -cycle $D \in C(r, d, n)$ such that $\psi(D) = [\mathcal{R}_D]$ vanishes at $\phi(W)$ for all $W \in \mathcal{H}$. Suppose $D = \sum_{i=1}^s c_i \mathcal{V}_i$ where $c_i \in \mathbb{N}^+$ for $i \in [s]$ and $\mathcal{V}_1, \dots, \mathcal{V}_s$ are distinct irreducible varieties.

Let $\mathcal{V} = \bigcup_{i=1}^s \mathcal{V}_i$. Note $\mathcal{V} \in \mathcal{F}$ since $\deg(\mathcal{V}) = \sum_{i=1}^s \deg(\mathcal{V}_i) \leq \sum_{i=1}^s c_i \deg(\mathcal{V}_i) = d$. For all $W \in \mathcal{H}$, we know $\mathcal{R}_D = \prod_{i=1}^s \mathcal{R}_{\mathcal{V}_i}^{c_i}$ vanishes at $\phi(W)$, or equivalently, $\mathcal{R}_{\mathcal{V}} = \prod_{i=1}^s \mathcal{R}_{\mathcal{V}_i}$ vanishes at $\phi(W)$. This implies $\mathcal{V} \cap W \neq \emptyset$ for all $W \in \mathcal{H}$. As $\mathcal{V} \in \mathcal{F}$, this contradicts our assumption about \mathcal{H} . We conclude

$$|\mathcal{H}| \geq \dim(C(r, d, n)) + 1.$$

The dimension of $C(r, d, n)$ is $(r+1)(n-r)$ when $d=1$ and is given by Theorem 4.2 when $d>1$. Plugging in $r=n-k-1$ proves the theorem. \square

Remark. It is easy to show that the lower bound in Theorem 1.7 is optimal by reversing its proof. Namely, we add random k -subspaces $W \in \mathbb{G}(k, n)$ to \mathcal{H} one by one, such that each time the dimension of $\psi(C(r, d, n)) \cap \bigcap_{W \in \mathcal{H}} H_W$ is reduced by one with high probability. It is easy to see that at each step, a general k -subspace W does reduce the dimension by one. However, it requires more work to prove a reasonable bound for the coefficients defining such a k -subspace W . This is because we need to apply a union bound over the irreducible components of $\psi(C(r, d, n)) \cap \bigcap_{W \in \mathcal{H}} H_W$. An upper bound for the number of these irreducible components can be shown by following [Kol13, Exercise 3.28]. We postpone the details to the full version of this paper.

5 Applications

In this section, we use the explicit constructions of variety-evasive subspace families in Section 3 to derandomize Noether's Normalization Lemma (Theorem 1.8 and Theorem 1.9) and black-box PIT for special depth-4 circuits (Theorem 1.11). The proof of Theorem 1.11 only uses the simple construction of variety-evasive subspace families (Lemma 3.9).

5.1 Derandomization of Noether's Normalization Lemma

Suppose W is a k -subspace of \mathbb{P}^n , and $\ell_1, \dots, \ell_{n-k} \in \mathbb{F}[X_1, \dots, X_{n+1}]$ are $n-k$ homogeneous linear polynomials such that $W = \mathcal{V}(\ell_1, \dots, \ell_{n-k})$. Then we have a map $\pi_{\ell_1, \dots, \ell_{n-k}} : \mathbb{P}^n \setminus W \rightarrow \mathbb{P}^{n-k-1}$ defined by

$$\pi_{\ell_1, \dots, \ell_{n-k}} : \mathbf{x} \mapsto (\ell_1(\mathbf{x}), \dots, \ell_{n-k}(\mathbf{x}))$$

which is well-defined since $\ell_1, \dots, \ell_{n-k}$ never simultaneously vanish on $\mathbb{P}^n \setminus W$. We say $\pi_{\ell_1, \dots, \ell_{n-k}}$ is a *projection* from $\mathbb{P}^n \setminus W$ to \mathbb{P}^{n-k-1} and W is its *center*.

The following lemma is crucial. Its proof can be found in [Sha13].

Lemma 5.1 ([Sha13, Section I.5.3, Theorem 7]). *Suppose $\pi : \mathbb{P}^n \setminus W \rightarrow \mathbb{P}^m$ is a projection with center W and \mathcal{V} is a subvariety of \mathbb{P}^n disjoint from W . Then π restricts to a finite morphism from \mathcal{V} to \mathbb{P}^m .*

We are now ready to prove Theorem 1.8 and Theorem 1.9, which we restate below for convenience.

Theorem 1.8. *Let $n, d \in \mathbb{N}^+$, $r \in \{0, 1, \dots, n\}$, and $\epsilon \in (0, 1)$. There exists an explicit collection \mathcal{L} of linear maps $\mathbb{A}^{n+1} \rightarrow \mathbb{A}^{r+1}$ of size $\text{poly}(N(k, d, n), n, 1/\epsilon)$ such that for every subvariety $\mathcal{V} \subseteq \mathbb{P}^n$ of dimension r and degree at most d , all but at most ϵ -fraction of $\pi \in \mathcal{L}$ induce a surjective finite morphism from \mathcal{V} to \mathbb{P}^r . Moreover, \mathcal{L} can be computed in time polynomial in $|\mathcal{L}|$ (and $\log p$, if $\text{char}(\mathbb{F}) = p > 0$).*

Proof. If $r = n$, we have $\mathcal{V} = \mathbb{P}^n$. Then just use the identity map $\mathbb{A}^{n+1} \rightarrow \mathbb{A}^{n+1}$. So assume $r < n$.

Let $k = n - r - 1$. Construct an (n, d, ϵ) -evasive k -subspace family \mathcal{H} on \mathbb{P}^n using Theorem 1.6. Consider $W \in \mathcal{H}$. Pick $n - k = r + 1$ homogeneous linear polynomials $\ell_1, \dots, \ell_{r+1} \in \mathbb{F}[X_1, \dots, X_{n+1}]$ such that $W = \mathcal{V}(\ell_1, \dots, \ell_{r+1})$. These $r + 1$ linear polynomials determine a linear map $\tilde{\pi}_{\ell_1, \dots, \ell_{r+1}} : \mathbb{A}^{n+1} \rightarrow \mathbb{A}^{r+1}$ sending $\mathbf{x} \in \mathbb{A}^{n+1}$ to $(\ell_1(\mathbf{x}), \dots, \ell_{r+1}(\mathbf{x}))$, and the latter induces the projection $\pi_{\ell_1, \dots, \ell_{r+1}} : \mathbb{P}^n \setminus W \rightarrow \mathbb{P}^r$. Let \mathcal{L} be the collection of all these linear maps $\tilde{\pi}_{\ell_1, \dots, \ell_{r+1}}$, one from each $W \in \mathcal{H}$.

Let \mathcal{V} be a subvariety of \mathbb{P}^n of dimension r and degree at most d . We know all but at most ϵ -fraction of $W \in \mathcal{H}$ are disjoint from \mathcal{V} . So we just need to prove that for every $W \in \mathcal{H}$ disjoint from \mathcal{V} , the corresponding projection $\pi := \pi_{\ell_1, \dots, \ell_{r+1}} : \mathbb{P}^n \setminus W \rightarrow \mathbb{P}^r$ restricts to a surjective finite morphism from \mathcal{V} to \mathbb{P}^r . The restriction $\pi|_{\mathcal{V}} : \mathcal{V} \rightarrow \mathbb{P}^r$ is indeed finite by Lemma 5.1. So its image $\pi(\mathcal{V})$ is closed and has dimension $\dim(\mathcal{V}) = r$. The only r -dimensional closed subset of \mathbb{P}^r is \mathbb{P}^r itself. So π is surjective. \square

Theorem 1.9. *Let $n, d \in \mathbb{N}^+$ and $r \in \{0, 1, \dots, n\}$, and $\epsilon \in (0, 1)$. There exists an explicit collection \mathcal{L} of linear maps $\mathbb{A}^n \rightarrow \mathbb{A}^r$ of size $\text{poly}(N(k, d, n), n, 1/\epsilon)$ such that for every subvariety $\mathcal{V} \subseteq \mathbb{A}^n$ of dimension r and degree at most d , all but at most ϵ -fraction of $\pi \in \mathcal{L}$ restrict to a surjective finite morphism from \mathcal{V} to \mathbb{A}^r . Moreover, \mathcal{L} can be computed in time polynomial in $|\mathcal{L}|$ (and $\log p$, if $\text{char}(\mathbb{F}) = p > 0$).*

Proof. If $r = n$, we have $\mathcal{V} = \mathbb{A}^n$. Then just use the identity map $\mathbb{A}^n \rightarrow \mathbb{A}^n$. If $r = 0$, use the only map $\mathbb{A}^n \rightarrow \mathbb{A}^0$. So assume $0 < r < n$. Regard \mathbb{A}^n as an open subset of \mathbb{P}^n via $(x_1, \dots, x_n) \mapsto (x_1, \dots, x_n, 1)$. Similarly, regard \mathbb{A}^r as an open subset of \mathbb{P}^r via $(x_1, \dots, x_r) \mapsto (x_1, \dots, x_r, 1)$.

Let $k = n - r - 1$. Construct an $(n - 1, d, \epsilon)$ -evasive k -subspace family \mathcal{H} on $H_\infty = \mathbb{P}^n \setminus \mathbb{A}^n \cong \mathbb{P}^{n-1}$ using Theorem 1.6. Consider $W \in \mathcal{H}$. Pick $n - k = r + 1$ homogeneous linear polynomials $\ell_1, \dots, \ell_{r+1} \in \mathbb{F}[X_1, \dots, X_{n+1}]$ such that $\ell_{r+1} = X_{n+1}$, $\ell_1, \dots, \ell_r \in \mathbb{F}[X_1, \dots, X_n]$, and $W = \mathcal{V}(\ell_1, \dots, \ell_{r+1})$. This is possible as $W \subseteq H_\infty = \mathcal{V}(X_{n+1})$. These $r + 1$ linear polynomials determine the projection $\pi_{\ell_1, \dots, \ell_{r+1}} : \mathbb{P}^n \setminus W \rightarrow \mathbb{P}^r$, defined by

$$\mathbf{x} = (x_1, \dots, x_{n+1}) \mapsto (\ell_1(\mathbf{x}), \dots, \ell_{r+1}(\mathbf{x})) = (\ell_1(\mathbf{x}), \dots, \ell_r(\mathbf{x}), x_{n+1}).$$

As $x_{n+1} = 1$ for $\mathbf{x} \in \mathbb{A}^n$, we have $\pi_{\ell_1, \dots, \ell_{r+1}}(\mathbb{A}^n) \subseteq \mathbb{A}^r$. Restricting $\pi_{\ell_1, \dots, \ell_{r+1}}$ on \mathbb{A}^n yields a map $\pi_{\ell_1, \dots, \ell_{r+1}}|_{\mathbb{A}^n} : \mathbb{A}^n \rightarrow \mathbb{A}^r$, which is a linear map as ℓ_1, \dots, ℓ_r are homogeneous linear polynomials in $\mathbb{F}[X_1, \dots, X_n]$. Let \mathcal{L} be the collection of all these linear maps $\pi_{\ell_1, \dots, \ell_{r+1}}|_{\mathbb{A}^n}$, one from each $W \in \mathcal{H}$.

Let \mathcal{V} be a subvariety of \mathbb{A}^n of dimension r and degree at most d . Its projective closure \mathcal{V}_{cl} has dimension $\dim(\mathcal{V}) = r$ and degree $\deg(\mathcal{V}) \leq d$. By the definition of \mathcal{V}_{cl} , none of the irreducible components of \mathcal{V}_{cl} is fully contained in H_∞ . So by Lemma 2.8, the projective subvariety $\mathcal{V}_{\text{cl}} \cap H_\infty$ has dimension $r - 1$ and degree at most d .

By the choice of \mathcal{H} , all but at most ϵ -fraction of $W \in \mathcal{H}$ are disjoint from $\mathcal{V}_{\text{cl}} \cap H_\infty$ and hence from \mathcal{V}_{cl} . So we just need to prove that for every $W \in \mathcal{H}$ disjoint from \mathcal{V}_{cl} and the corresponding projection $\pi := \pi_{\ell_1, \dots, \ell_{r+1}}$, the map $\pi|_{\mathcal{V}} : \mathcal{V} \rightarrow \mathbb{A}^r$ is a surjective finite morphism. We have already seen from the proof of Theorem 1.8 that, as the center W of π is disjoint from \mathcal{V}_{cl} , the projection π restricts to a surjective finite morphism $\pi|_{\mathcal{V}_{\text{cl}}} : \mathcal{V}_{\text{cl}} \rightarrow \mathbb{P}^r$. As $\mathcal{V} = \mathcal{V}_{\text{cl}} \cap \mathbb{A}^n = \mathcal{V}_{\text{cl}} \cap \pi^{-1}(\mathbb{A}^r)$, the map $\pi|_{\mathcal{V}}$ is precisely the restriction of $\pi|_{\mathcal{V}_{\text{cl}}}$ to $(\pi|_{\mathcal{V}_{\text{cl}}})^{-1}(\mathbb{A}^r)$. As $\pi|_{\mathcal{V}_{\text{cl}}}$ is a surjective finite morphism, so is $\pi|_{\mathcal{V}}$. \square

Remark. For simplicity, we have restricted to the category of varieties over an algebraically closed field \mathbb{F} when stating Theorem 1.8 and Theorem 1.9. We now mention some generalizations without proofs, which lead to the usual algebraic formulation of Noether's normalization lemma and its derandomization:

- As mentioned in the remark after Theorem 1.6, the coefficients of the linear maps that we use live in a non-algebraically closed field $\mathbb{K}_0 \subseteq \mathbb{F}$, which is either \mathbb{Q} or a finite extension of \mathbb{F}_p . For any field $\mathbb{K} \supseteq \mathbb{K}_0$, we have actually constructed explicit families of linear maps that are defined over \mathbb{K} . Theorem 1.8 and Theorem 1.9 then hold for projective/affine varieties over \mathbb{K} (which we have not defined) as well.

- Furthermore, Theorem 1.8 and Theorem 1.9 hold for *closed subschemes* of projective/affine spaces over \mathbb{K} as well. In fact, it suffices to consider the variety $\mathcal{V}_{\text{red}} := \mathcal{V}(\sqrt{I(\mathcal{V})})$ in place of a closed subscheme \mathcal{V} when checking if a linear map gives a valid surjective finite morphism. This is because the evading property that we need is set-theoretic.
- A generalization of Theorem 1.9 then translates into the following derandomization of Noether's normalization lemma: Let \mathbb{K} be a field containing all the coefficients of the linear maps in \mathcal{L} , where \mathcal{L} is as constructed in Theorem 1.9. Let $A \neq 0$ be a finitely generated commutative \mathbb{K} -algebra with generators b_1, \dots, b_n such that the *Krull dimension* of A is r and the variety $\mathcal{V} \subseteq \mathbb{A}_{\mathbb{K}}^n$ has degree at most d , where $\mathcal{V} = \mathcal{V}(\sqrt{I})$ and $I \subseteq \mathbb{K}[X_1, \dots, X_n]$ is the set of polynomial relations that b_1, \dots, b_n satisfy. For a linear map $\pi \in \mathcal{L}$ defined by $(x_1, \dots, x_n) \mapsto (\sum_{i=1}^n c_{i,1}x_i, \dots, \sum_{i=1}^n c_{i,r}x_i)$, let $y_j^\pi = \sum_{i=1}^n c_{i,j}b_i$ for $j \in [r]$. Then for all but at most ϵ -fraction of $\pi \in \mathcal{L}$, the corresponding y_1^π, \dots, y_r^π are algebraically independent and A is a finitely-generated module over $\mathbb{K}[y_1^\pi, \dots, y_r^\pi]$. The existence of such y_1^π, \dots, y_r^π is the content of the usual algebraic formulation of Noether's normalization lemma [AM69, Chapter 5, Exercise 16].

5.2 Black-Box PIT for Non-SG Depth-4 Circuits

We first define $\Sigma\Pi\Sigma\Pi(k, r)$ circuits and non-SG $\Sigma\Pi\Sigma\Pi(k, r)$ circuits.

Definition 5.2 ($\Sigma\Pi\Sigma\Pi(k, r)$ circuit). *An algebraic circuit C over \mathbb{F} is a $\Sigma\Pi\Sigma\Pi(k, r)$ circuit if it has the form*

$$C(X_1, \dots, X_n) = \sum_{i=1}^{k'} F_i = \sum_{i=1}^{k'} \prod_{j=1}^{d_i} Q_{i,j} \quad (2)$$

where $k' \leq k$, $d_1, \dots, d_{k'} \in \mathbb{N}^+$, $F_i = \prod_{j=1}^{d_i} Q_{i,j}$ for $j \in [k']$, and each $Q_{i,j}$ is a polynomial in X_1, \dots, X_n of degree at most r over \mathbb{F} . The degree of the circuit C is defined to be $\max\{\deg(F_i) : i \in [k']\}$. In addition:

- C is minimal if $\sum_{i \in I} F_i \neq 0$ for all nonempty proper subset $I \subseteq [k']$.
- C is homogeneous if all the polynomials F_i are homogeneous of the same degree.
- Let $\gcd(C) := \gcd(F_1, \dots, F_{k'})$. We say C is simple if $\gcd(C) = 1$. In general, we have $C = \gcd(C) \cdot \text{sim}(C)$ where $\text{sim}(C)$ is a simple $\Sigma\Pi\Sigma\Pi(k, r)$ circuit, called the simple part of C . Note the simple part of a minimal $\Sigma\Pi\Sigma\Pi(k, r)$ circuit is still minimal.

The polynomial computed by C is again denoted by C by an abuse of notation.

Definition 5.3 (Non-SG circuit). *A minimal, simple, and homogeneous $\Sigma\Pi\Sigma\Pi(k, r)$ circuit $C(X_1, \dots, X_n) = \sum_{i=1}^{k'} F_i$ as in (2) is non-SG if there exists $i \in [k']$ such that*

$$\bigcap_{j \in [k'] \setminus i} \mathcal{V}(F_j) \not\subseteq \mathcal{V}(F_i)$$

where $\mathcal{V}(F)$ denotes the subvariety of \mathbb{P}^n defined by F . More generally, a minimal and simple $\Sigma\Pi\Sigma\Pi(k, r)$ circuit $C(X_1, \dots, X_n) = \sum_{i=1}^{k'} F_i$ of degree d is non-SG if its homogenization

$$\tilde{C}(X_1, \dots, X_{n+1}) = \sum_{i=1}^{k'} F_i(X_1/X_{n+1}, \dots, X_n/X_{n+1}) \cdot X_{n+1}^d = \sum_{i=1}^{k'} \prod_{j=1}^{d_i} \tilde{Q}_{i,j}$$

is non-SG, where each $\tilde{Q}_{i,j}$ is either the homogenization of $Q_{i,j}$ or X_{n+1} . A minimal $\Sigma\Pi\Sigma\Pi(k, r)$ circuit C is non-SG if $\text{sim}(C)$ is non-SG. Finally, a $\Sigma\Pi\Sigma\Pi(k, r)$ circuit is non-SG if it has an equivalent minimal non-SG $\Sigma\Pi\Sigma\Pi(k, r)$ circuit.

We restate our result (Theorem 1.11) and then give a proof.

Theorem 1.11. *There exists a deterministic black-box PIT algorithm with time complexity polynomial in $d \cdot \binom{k(n+1+r^k)}{kr^k} \cdot \binom{k-1+d}{k-1} \leq \text{poly}(d^k, n^{r^k}, r^{k^2 r^k})$ (and $\log p$, if $\text{char}(\mathbb{F}) = p > 0$) for non-SG $\Sigma\Pi\Sigma\Pi(k, r)$ circuits of degree at most d in X_1, \dots, X_n over an algebraically closed field \mathbb{F} .*

Proof. If $n \leq k - 1$, we may simply use Lemma 2.3 to construct a $\frac{1}{2}$ -hitting set of size polynomial in $\binom{n+d}{n} \leq \binom{k-1+d}{k-1}$ for n -variate polynomials of degree at most d , and then run the corresponding black-box PIT algorithm. So assume $n > k - 1$.

Consider a nonzero non-SG $\Sigma\Pi\Sigma\Pi(k, r)$ circuit C of degree at most d . We want to design a black-box PIT algorithm for C . By replacing C with an equivalent minimal non-SG circuit, we may assume C is minimal. Let $D = \text{gcd}(C)$ and $E = \text{sim}(C)$. Let \tilde{C} , \tilde{D} , and \tilde{E} be the homogenization of C , D , and E respectively. Then $\tilde{D} = \text{gcd}(\tilde{C})$, $\tilde{E} = \text{sim}(\tilde{C})$, and $\tilde{C} = \tilde{D} \cdot \tilde{E}$.

Let \mathcal{H} be an affine $(k - 1)$ -subspace family on \mathbb{A}^n of size $\text{poly}(\binom{k(n+1+r^k)}{kr^k}, d)$ such that $\mathcal{H}' := \{W_{\text{cl}} : W \in \mathcal{H}\}$ is an $(n, r^k, \frac{1}{4d})$ -evasive $(k - 1)$ -subspace family on \mathbb{P}^n . Such a family \mathcal{H} can be computed using Lemma 3.13 and Lemma 3.9. We claim

1. $\tilde{D}|_W \neq 0$ for all but at most $\frac{1}{4}$ -fraction of $W \in \mathcal{H}'$, and
2. $\tilde{E}|_W \neq 0$ for all but at most $\frac{1}{4}$ -fraction of $W \in \mathcal{H}'$.

Assume these two claims hold. Then for at least half of $W \in \mathcal{H}$, we have $\tilde{C}|_{W_{\text{cl}}} \neq 0$ and hence $C|_W = \tilde{C}|_{W_{\text{cl}} \cap \mathbb{A}^n} \neq 0$, where we use the facts that $\tilde{C}(X_1, \dots, X_n, 1)$ equals $C(X_1, \dots, X_n)$ and $W_{\text{cl}} \cap \mathbb{A}^n$ is dense in W_{cl} . The restriction of C to each $W \cong \mathbb{A}^{k-1}$ is a $(k - 1)$ -variate polynomial of degree at most d . So to test if $C|_W$ is zero, we just need to use Lemma 2.3 to construct a hitting set in W of size $\text{poly}(\binom{k-1+d}{k-1})$ for $(k - 1)$ -variate polynomials of degree at most d . Take the union of these hitting sets to obtain a hitting set of size $\text{poly}(\binom{k(n+1+r^k)}{kr^k}, d, \binom{k-1+d}{k-1})$ and we are done.

So it remains to prove the two claims. Note \tilde{D} is the product of at most d factors whose degrees are bounded by r . The first claim then follows from the $(n, r^k, \frac{1}{4d})$ -evasiveness of \mathcal{H}' and the union bound.

Now we prove the second claim. By definition, \tilde{E} is a non-SG $\Sigma\Pi\Sigma\Pi(k, r)$ circuit. Suppose it has the form

$$\tilde{E} = \sum_{i=1}^{k'} F_i = \sum_{i=1}^{k'} \prod_{j=1}^{d_i} Q_{i,j} \quad (3)$$

where each $Q_{i,j}$ is a homogeneous polynomial of degree at most r . As \tilde{E} is non-SG, there exists $i_0 \in [k']$ such that

$$\bigcap_{i \in [k'] \setminus i_0} \mathcal{V}(F_i) \not\subseteq \mathcal{V}(F_{i_0})$$

Without loss of generality, we may assume $i_0 = k'$. Note $\mathcal{V}(F_i) = \bigcup_{j=1}^{d_i} \mathcal{V}(Q_{i,j})$ for $i \in [k']$. So there exists $(j_1, \dots, j_{k'-1}) \in [d_1] \times \dots \times [d_{k'-1}]$ such that

$$\bigcap_{i=1}^{k'-1} \mathcal{V}(Q_{i,j_i}) \not\subseteq \mathcal{V}(F_{k'}).$$

Let \mathcal{V}_0 be an irreducible component of $\bigcap_{i=1}^{k'-1} \mathcal{V}(Q_{i,j_i})$ such that $\mathcal{V}_0 \not\subseteq \mathcal{V}(F_{k'})$. Let $d_0 = \dim(\mathcal{V}_0) \geq 0$. By Lemma 2.8, we have $d_0 \geq n - k' + 1$ and the variety $\mathcal{V}_0 \cap \mathcal{V}(F_{k'}) = \bigcup_{j=1}^{d_{k'}} (\mathcal{V}_0 \cap \mathcal{V}(Q_{k',j}))$ has dimension at most $d_0 - 1$. For each $j \in [d_{k'}]$, the degree of $\mathcal{V}_0 \cap \mathcal{V}(Q_{k',j})$ is at most r^k by Lemma 2.8 (or by Bézout's

inequality [Hei83]). By $(n, r^k, \frac{1}{4d})$ -evasiveness of \mathcal{H}' and the union bound, all but at most $\frac{1}{4}$ -fraction of $W \in \mathcal{H}'$ evade $\mathcal{V}_0 \cap \mathcal{V}(Q_{k',j})$ for $j = 1, 2, \dots, d_{k'}$.

Consider any $W \in \mathcal{H}'$ that evades $\mathcal{V}_0 \cap \mathcal{V}(Q_{k',j})$ for $j = 1, 2, \dots, d_{k'}$. We just need to prove $\tilde{E}|_W \neq 0$, or equivalently, $W \not\subseteq \mathcal{V}(\tilde{E})$. Assume to the contrary that $W \subseteq \mathcal{V}(\tilde{E})$. Then $W \cap \mathcal{V}_0 \subseteq \mathcal{V}(\tilde{E})$. So

$$W \cap \mathcal{V}_0 = W \cap \mathcal{V}_0 \cap \mathcal{V}(\tilde{E}) = W \cap \mathcal{V}_0 \cap \mathcal{V}\left(\prod_{j=1}^{d_{k'}} Q_{k',j}\right) = \bigcup_{j=1}^{d_{k'}} (W \cap \mathcal{V}_0 \cap \mathcal{V}(Q_{k',j})) \quad (4)$$

where the second equality holds since $\tilde{E} \equiv \prod_{j=1}^{d_{k'}} Q_{k',j}$ modulo the ideal $I_0 := \langle Q_{1,j_1}, \dots, Q_{k'-1,j_{k'-1}} \rangle$ by (3) and $\mathcal{V}_0 \subseteq \bigcap_{i=1}^{k'-1} \mathcal{V}(Q_{i,j_i}) = \mathcal{V}(I_0)$. We know the dimension of $\bigcup_{j=1}^{d_{k'}} (\mathcal{V}_0 \cap \mathcal{V}(Q_{k',j}))$ is at most $d_0 - 1$. So by the choice of W , the dimension of $\bigcup_{j=1}^{d_{k'}} (W \cap \mathcal{V}_0 \cap \mathcal{V}(Q_{k',j}))$ is at most $(k-1) + (d_0 - 1) - n$. However, by Lemma 2.9, the dimension of $W \cap \mathcal{V}_0$ is at least $(k-1) + d_0 - n \geq 0$, where we use the fact $d_0 \geq n - k' + 1 \geq n - k + 1$. This contradicts (4). So $\tilde{E}|_W \neq 0$. \square

6 Open Problems and Future Directions

We have seen that constructing explicit variety evasive subspace families is a natural problem that generalizes important problems in algebraic pseudorandomness and algebraic complexity theory, including deterministic black-box polynomial identity testing (evading varieties of codimension one) and constructing explicit lossless rank condensers (evading varieties of degree one). It is closely connected with advanced topics in algebraic geometry such as Chow forms and Chow varieties, and has applications to derandomizing PIT and non-explicit results in algebraic geometry like Noether's normalization lemma.

There are many interesting open problems and potential future directions. We list some of them here.

1. Theorem 1.6 focuses on subvarieties of bounded degree in a projective or affine space. Are there other interesting families of varieties for which we could construct explicit variety evasive subspace families? Families that are defined computation-theoretically may be particularly interesting, as many results of this kind are already known for polynomial identity testing.
2. Can explicit variety evasive subspace families be used to derandomize other non-explicit results in algebraic geometry?
3. Can our explicit construction in Theorem 1.6 be improved? In the case $k = 0$ and the case $d = 1$, there are optimal or essentially optimal constructions, and our construction indeed degenerates into these constructions. In general, however, there is a significant gap between the upper bound in Theorem 1.6 and the lower bound in Theorem 1.7.
4. Extending the notion of *strong* lossless rank condensers [FG15], one could strengthen the definition of (\mathcal{F}, ϵ) -evasive subspace families in Definition 1.3 by bounding the total deviation of the dimension instead of the number of bad subspaces. At the same time, one could consider the setting where there is gap between $\dim(\mathcal{V}_1)$ and $\text{codim}(\mathcal{V}_2)$, as in typical applications of *subspace designs* [GX13, GRZ20, GRX21]. Alternatively, one could relax the definition by allowing $\dim(\mathcal{V}_1 \cap \mathcal{V}_2)$ to be slightly greater than $\dim(\mathcal{V}_1) + \dim(\mathcal{V}_2) - n$, which is related to the notion of *lossy* rank condensers in [FG15]. It is natural to study explicit constructions of these variants and their applications, which can be seen as extensions of the theory of “linear-algebraic pseudorandomness” [FG15] to a nonlinear setting.
5. Could our lower bound (Theorem 1.7) be extended to the affine case or to a “lossy” relaxation of the problem?

6. When $n - k = O(1)$, our lower bound (Theorem 1.7) is only polynomial in n and d . So one question is if there are explicit constructions of polynomial size when $n - k = O(1)$.

As a concrete special case, consider the problem of constructing an explicit affine $(n - 2)$ -subspace family \mathcal{H} on \mathbb{A}^n such that \mathcal{H} is evasive for degree- d curves that are images of morphisms $\mathbb{A}^1 \rightarrow \mathbb{A}^n$. Note that for $\varphi : \mathbb{A}^1 \rightarrow \mathbb{A}^n$ corresponding to a ring homomorphism $\varphi^\# : \mathbb{F}[X_1, \dots, X_n] \rightarrow \mathbb{F}[Y]$, an affine $(n - 2)$ -subspace defined by affine linear polynomials ℓ_1 and ℓ_2 evades the curve $\text{Im}(\varphi)$ iff $\varphi^\#(\ell_1)$ and $\varphi^\#(\ell_2)$ have no common root. Using *resultants*, we could reduce this problem to black-box PIT for symbolic determinants. We are not aware of any *unconditional* derandomization whose time complexity is subexponential in $\min\{n, d\}$, however.

Acknowledgements

We thank Nitin Saxena, Noga Ron-Zewi, Amit Sinhababu, and Suryajith Chillara for helpful discussions. We also thank the anonymous reviewers of CCC'21 for their careful reading of our manuscript and their insightful suggestions that helped improve this paper.

References

- [AGKS15] Manindra Agrawal, Rohit Gurjar, Arpita Korwar, and Nitin Saxena. Hitting-sets for ROABP and sum of set-multilinear circuits. *SIAM Journal on Computing*, 44(3):669–697, 2015.
- [AM69] Michael F. Atiyah and I. G. MacDonald. *Introduction to Commutative Algebra*. Addison-Wesley-Longman, 1969.
- [ASSS16] Manindra Agrawal, Chandan Saha, Ramprasad Saptharishi, and Nitin Saxena. Jacobian hits circuits: Hitting sets, lower bounds for depth- d occur- k formulas and depth-3 transcendence degree- k circuits. *SIAM Journal on Computing*, 45(4):1533–1562, 2016.
- [AV08] Manindra Agrawal and V Vinay. Arithmetic circuits: A chasm at depth four. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 67–75, 2008.
- [Azc92] Pablo Azcue. *On the dimension of the Chow varieties*. PhD thesis, Harvard University, 1992.
- [BE19] Juliette Bruce and Daniel Erman. A probabilistic approach to systems of parameters and Noether normalization. *Algebra & Number Theory*, 13(9):2081–2102, 2019.
- [BMS13] Malte Beecken, Johannes Mittmann, and Nitin Saxena. Algebraic independence and blackbox identity testing. *Information and Computation*, 222:2–19, 2013.
- [Bog05] Andrej Bogdanov. Pseudorandom generators for low degree polynomials. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 21–30, 2005.
- [BP20] Markus Bläser and Anurag Pandey. Polynomial identity testing for low degree polynomials with optimal randomness. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, pages 8:1–8:13, 2020.
- [Bsh14] Nader Bshouty. Testers and their applications. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, pages 327–352, 2014.
- [Cay60] Arthur Cayley. On a new analytical representation of curves in space. *The Quarterly Journal of Pure and Applied Mathematics*, 3:225–236, 1860.

- [CGS20] Arkadev Chattopadhyay, Ankit Garg, and Suhail Sherif. Towards stronger counterexamples to the log-approximate-rank conjecture. *arXiv preprint arXiv:2009.02717*, 2020.
- [CTS13] Gil Cohen and Amnon Ta-Shma. Pseudorandom generators for low degree polynomials from algebraic geometry codes. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 20, page 155, 2013.
- [CvdW37] Wei-Liang Chow and B.L. van der Waerden. Zur algebraischen Geometrie. IX. *Mathematische Annalen*, 113(1):692–704, 1937.
- [DGW09] Zeev Dvir, Ariel Gabizon, and Avi Wigderson. Extractors and rank extractors for polynomial sources. *Computational Complexity*, 18(1):1–58, 2009.
- [DKL14] Zeev Dvir, János Kollár, and Shachar Lovett. Variety evasive sets. *Computational Complexity*, 23(4):509–529, 2014.
- [DL78] Richard A. Demillo and Richard J. Lipton. A probabilistic remark on algebraic program testing. *Information Processing Letters*, 7(4):193–195, 1978.
- [DL12] Zeev Dvir and Shachar Lovett. Subspace evasive sets. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing*, pages 351–358, 2012.
- [DS95] John Dalbec and Bernd Sturmfels. Introduction to Chow forms. *Invariant Methods in Discrete and Computational Geometry*, pages 37–58, 1995.
- [DS07] Zeev Dvir and Amir Shpilka. Locally decodable codes with two queries and polynomial identity testing for depth 3 circuits. *SIAM Journal on Computing*, 36(5):1404–1434, 2007.
- [Dub93] Thomas W Dubé. A combinatorial proof of the effective Nullstellensatz. *Journal of Symbolic Computation*, 15(3):277–296, 1993.
- [Dvi12] Zeev Dvir. Extractors for varieties. *Computational Complexity*, 21(4):515–572, 2012.
- [EH87] David Eisenbud and Joe Harris. On varieties of minimal degree. In *Proceedings of Symposia in Pure Mathematics*, volume 46, pages 3–13, 1987.
- [EH92] David Eisenbud and Joe Harris. The dimension of the Chow variety of curves. *Compositio Mathematica*, 83(3):291–310, 1992.
- [Eis13] David Eisenbud. *Commutative Algebra: with a View Toward Algebraic Geometry*, volume 150. Springer Science & Business Media, 2013.
- [FG15] Michael A. Forbes and Venkatesan Guruswami. Dimension expanders via rank condensers. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, 2015.
- [For14] Michael A. Forbes. *Polynomial identity testing of read-once oblivious algebraic branching programs*. PhD thesis, Massachusetts Institute of Technology, 2014.
- [FS12] Michael A. Forbes and Amir Shpilka. On identity testing of tensors, low-rank recovery and compressed sensing. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing*, pages 163–172, 2012.

- [FS13] Michael A. Forbes and Amir Shpilka. Explicit Noether normalization for simultaneous conjugation via polynomial identity testing. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, pages 527–542, 2013.
- [FS18] Michael A. Forbes and Amir Shpilka. A PSPACE construction of a hitting set for the closure of small algebraic circuits. In *Proceedings of the 50th Annual ACM Symposium on Theory of Computing*, pages 1180–1192, 2018.
- [FSS14] Michael A. Forbes, Ramprasad Saptharishi, and Amir Shpilka. Hitting sets for multilinear read-once algebraic branching programs, in any order. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 867–875, 2014.
- [Ful97] William Fulton. *Young Tableaux: With Applications to Representation Theory and Geometry*, volume 35. Cambridge University Press, 1997.
- [GH93] Marc Giusti and Joos Heintz. La détermination des points isolés et de la dimension d’une variété algébrique peut se faire en temps polynomial. *Computational Algebraic Geometry and Commutative Algebra (Cortona, 1991)*, pages 216–256, 1993.
- [GHL⁺00] Marc Giusti, Klemens Hägele, Grégoire Lecerf, Joël Marchand, and Bruno Salvy. The projective Noether Maple package: computing the dimension of a projective variety. *Journal of Symbolic Computation*, 30(3):291–307, 2000.
- [GK16] Venkatesan Guruswami and Swastik Kopparty. Explicit subspace designs. *Combinatorica*, 36(2):161–185, 2016.
- [GKZ94] Israel M. Gelfand, Mikhail M. Kapranov, and Andrei V. Zelevinsky. *Discriminants, Resultants and Multidimensional Determinants*. Birkhäuser, 1994.
- [GR08] Ariel Gabizon and Ran Raz. Deterministic extractors for affine sources over large fields. *Combinatorica*, 28(4):415–440, 2008.
- [GRX21] Venkatesan Guruswami, Nicolas Resch, and Chaoping Xing. Lossless dimension expanders via linearized polynomials and subspace designs. *Combinatorica*, pages 1–35, 2021.
- [GRZ20] Zeyu Guo and Noga Ron-Zewi. Efficient list-decoding with constant alphabet and list sizes. *arXiv preprint arXiv:2011.05884*, 2020. To appear in STOC 2021.
- [GSS19] Zeyu Guo, Nitin Saxena, and Amit Sinhababu. Algebraic dependencies and PSPACE algorithms in approximative complexity over any field. *Theory of Computing*, 15(16):1–30, 2019.
- [Gup14] Ankit Gupta. Algebraic geometric techniques for depth-4 PIT & Sylvester-Gallai conjectures for varieties. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 21, page 130, 2014.
- [GWX16] Venkatesan Guruswami, Carol Wang, and Chaoping Xing. Explicit list-decodable rank-metric and subspace codes via subspace designs. *IEEE Transactions on Information Theory*, 62(5):2707–2718, 2016.
- [GX13] Venkatesan Guruswami and Chaoping Xing. List decoding Reed-Solomon, Algebraic-Geometric, and Gabidulin subcodes up to the Singleton bound. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*, pages 843–852, 2013.

- [GXY18] Venkatesan Guruswami, Chaoping Xing, and Chen Yuan. Subspace designs based on algebraic function fields. *Transactions of the American Mathematical Society*, 370(12):8757–8775, 2018.
- [Har13] Joe Harris. *Algebraic Geometry: A First Course*, volume 133. Springer Science & Business Media, 2013.
- [Hei83] Joos Heintz. Definability and fast quantifier elimination in algebraically closed fields. *Theoretical Computer Science*, 24(3):239–277, 1983.
- [Hi190] David Hilbert. Ueber die Theorie der algebraischen Formen. *Mathematische Annalen*, 36(4):473–534, 1890.
- [HS81] Joos Heintz and Malte Sieveking. Absolute primality of polynomials is decidable in random polynomial time in the number of variables. In *International Colloquium on Automata, Languages, and Programming*, pages 16–28, 1981.
- [JKSS04] Gabriela Jeronimo, Teresa Krick, Juan Sabia, and Martín Sombra. The computational complexity of the Chow form. *Foundations of Computational Mathematics*, 4(1):41–117, 2004.
- [JT13] Michael Joswig and Thorsten Theobald. *Polyhedral and Algebraic Methods in Computational Geometry*. Springer Science & Business Media, 2013.
- [Kol88] János Kollár. Sharp effective Nullstellensatz. *Journal of the American Mathematical Society*, pages 963–975, 1988.
- [Kol13] János Kollár. *Rational Curves on Algebraic Varieties*, volume 32. Springer Science & Business Media, 2013.
- [Kri02] Teresa Krick. Straight-line programs in polynomial equation solving. *Foundations of Computational Mathematics*, 312:96–136, 2002.
- [KRZSW18] Swastik Kopparty, Noga Ron-Zewi, Shubhangi Saraf, and Mary Wootters. Improved decoding of folded Reed-Solomon and multiplicity codes. In *Proceedings of the 59th Annual IEEE Symposium on Foundations of Computer Science*, pages 212–223, 2018.
- [KS01] Adam R. Klivans and Daniel Spielman. Randomness efficient identity testing of multivariate polynomials. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pages 216–223, 2001.
- [KS11] Zohar S. Karnin and Amir Shpilka. Black box polynomial identity testing of generalized depth-3 arithmetic circuits with bounded top fan-in. *Combinatorica*, 31(3):333, 2011.
- [Lan12] Joseph M. Landsberg. Tensors: geometry and applications. *Representation theory*, 381(402):3, 2012.
- [Lan15] Joseph M. Landsberg. Geometric complexity theory: an introduction for geometers. *Annali dell’universita’ di Ferrara*, 61(1):65–117, 2015.
- [Leh17] Brian Lehmann. Asymptotic behavior of the dimension of the Chow variety. *Advances in Mathematics*, 308:815–835, 2017.
- [Lu12] Chi-Jen Lu. Hitting set generators for sparse polynomials over any finite fields. In *Proceedings of the 27th Conference on Computational Complexity*, pages 280–286, 2012.

- [Muk16] Partha Mukhopadhyay. Depth-4 identity testing and Noether’s normalization lemma. In *Proceedings of the 11th International Computer Science Symposium in Russia*, pages 309–323, 2016.
- [Mul17] Ketan Mulmuley. Geometric complexity theory V: Efficient algorithms for Noether normalization. *Journal of the American Mathematical Society*, 30(1):225–309, 2017.
- [Nag62] Masayoshi Nagata. Local rings. *Interscience Tracts in Pure and Applied Mathematics*, 1962.
- [Noe26] Emmy Noether. Der Endlichkeitssatz der Invarianten endlicher linearer Gruppen der Charakteristik p . *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 1926:28–35, 1926.
- [PS20a] Shir Peleg and Amir Shpilka. A generalized Sylvester-Gallai type theorem for quadratic polynomials. In *Proceedings of the 35th Computational Complexity Conference*, 2020.
- [PS20b] Shir Peleg and Amir Shpilka. Polynomial time deterministic identity testing algorithm for $\Sigma^{[3]}\Pi\Sigma\Pi^{[2]}$ circuits via Edelstein-Kelly type theorem for quadratic polynomials. *arXiv preprint arXiv:2006.08263*, 2020.
- [Sax09] Nitin Saxena. Progress on polynomial identity testing. *Bulletin of the EATCS*, 99:49–79, 2009.
- [Sax13] Nitin Saxena. Progress on polynomial identity testing - II. *Electronic Colloquium on Computational Complexity (ECCC)*, 20:186, 2013.
- [Sch80] Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM*, 27(4):701–717, 1980.
- [Sha13] Igor R. Shafarevich. *Basic Algebraic Geometry 1: Varieties in Projective Space*. Springer Science & Business Media, 2013.
- [Shp19] Amir Shpilka. Sylvester-Gallai type theorems for quadratic polynomials. In *Proceedings of the 51st Annual ACM Symposium on Theory of Computing*, pages 1203–1214, 2019.
- [SS12] Nitin Saxena and Comandur Seshadhri. Blackbox identity testing for bounded top-fanin depth-3 circuits: The field doesn’t matter. *SIAM Journal on Computing*, 41(5):1285–1298, 2012.
- [SY10] Amir Shpilka and Amir Yehudayoff. *Arithmetic Circuits: A Survey of Recent Results and Open Questions*. Now Publishers Inc, 2010.
- [Zip79] Richard Zippel. Probabilistic algorithms for sparse polynomials. In *International Symposium on Symbolic and Algebraic Manipulation*, pages 216–226, 1979.