

Duality for matrix space questions

Yuval Wigderson*

Abstract

We present a new proof of a classical theorem of Dieudonné: if a linear space of $n \times n$ matrices consists entirely of singular matrices, then its dimension is at most $n^2 - n$. Our proof is based on a surprising “duality” argument: we prove this universal upper bound by exhibiting a *single* matrix space that serves as a *lower bound* for a related problem. Interestingly, this approach only works for certain fields, but we use model-theoretic arguments to obtain the same result for all fields. We hope that this approach can be generalized to provide new duality-based proofs of other classical theorems on matrix spaces, and give some preliminary results in this direction.

1 Introduction.

Let F be a field, and let $F^{n \times n}$ denote the space of $n \times n$ matrices over F . $F^{n \times n}$ is a vector space over F , and a subspace¹ $W \leq F^{n \times n}$ is called a *matrix space* or *space of $n \times n$ matrices over F* . The study of matrix spaces has rich and surprising connections to many areas of mathematics, including algebraic geometry [11], algebraic topology [2], graph theory [16], group theory [6], invariant theory [8], matroid theory [17], quantum information theory [27], and theoretical computer science [10]. In most of these connections, one studies matrix spaces $W \leq F^{n \times n}$ such that every matrix in W satisfies some natural linear-algebraic property.

One of the most interesting and well-studied cases concerns *singular matrix spaces*, namely a space $S \leq F^{n \times n}$ such that every matrix $M \in S$ is singular (non-invertible). The study of singular matrix spaces goes back at least 75 years [8], and many of the connections mentioned above actually concern singular matrix spaces. As a striking example of the complexity of this simple definition, let us remark that it is unknown how to efficiently determine if a given matrix space is singular; moreover, finding an efficient deterministic algorithm for this problem would come close² to proving that $P \neq NP$ [15, 26] (see also the survey [25, Theorems 1.2 and 4.5]).

At the other extreme, we might ask to study spaces of matrices all of whose elements are invertible. A moment’s thought reveals that this makes no sense, since the zero matrix lies in every matrix space, and is certainly singular. However, if we exclude this trivial obstruction, we come to the definition of *totally non-singular* matrix spaces, namely a space $T \leq F^{n \times n}$ such that every non-zero matrix $0 \neq M \in T$ is invertible. The structure of totally non-singular matrix spaces turns out to be quite interesting, and yields unexpected connections to other areas of mathematics; for example, determining the maximum dimension of an $n \times n$ totally non-singular space over $F = \mathbb{R}$

*IST Austria, 3400 Klosterneuburg, Austria. Email: yuval.wigderson@ista.ac.at. Research supported by Dr. Max Rössler, the Walter Haefner Foundation, and the ETH Zürich Foundation.

¹Throughout, we use the notation $U \leq V$ to mean that U is a subspace of the vector space V .

²More precisely, it would prove much stronger circuit lower bounds than anything currently known.

is equivalent [2] to determining the maximum number of non-vanishing vector fields on the sphere S^{n-1} , which was famously resolved by Adams using K -theoretic tools [1]. This question, in turn, is closely connected to the representation theory of Clifford algebras.

Having made these definitions, one is immediately led to the following trivial observation: if $S \leq F^{n \times n}$ is singular and $T \leq F^{n \times n}$ is totally non-singular, then

$$S \cap T = \{0\}. \quad (1)$$

Indeed, any non-zero matrix in $S \cap T$ would need to simultaneously be invertible and non-invertible, hence cannot exist. Using the fact that $S, T \leq F^{n \times n}$, we conclude from (1) that

$$\dim S + \dim T \leq n^2, \quad (2)$$

since $\dim(F^{n \times n}) = n^2$. The goal of this paper is to show that some classical results about matrix spaces, whose proofs involve non-trivial algebraic machinery, can be derived from the extremely simple observations (1) and (2).

Example 1. Let us consider first the case $F = \mathbb{R}, n = 2$. It is not hard to come up with a 2-dimensional singular space of matrices; for example, we can take

$$S^* = \left\{ \begin{pmatrix} x & y \\ 0 & 0 \end{pmatrix} : x, y \in \mathbb{R} \right\},$$

which is certainly a singular matrix space as every matrix in S^* has a row of zeros. It is a bit harder, but not by much, to come up with a 2-dimensional totally non-singular space of matrices in $\mathbb{R}^{2 \times 2}$; for example, we can define

$$T^* = \left\{ \begin{pmatrix} x & y \\ -y & x \end{pmatrix} : x, y \in \mathbb{R} \right\}. \quad (3)$$

The determinant of the matrix $\begin{pmatrix} x & y \\ -y & x \end{pmatrix}$ is $x^2 + y^2$, which is strictly positive unless $x = y = 0$, hence we see that T^* is indeed totally non-singular. However, just from the *existence* of these two simple examples, we can derive that they are the largest possible examples.

Proposition 2. *The maximum dimension of a singular matrix space in $\mathbb{R}^{2 \times 2}$ is 2. Similarly, the maximum dimension of a totally non-singular matrix space in $\mathbb{R}^{2 \times 2}$ is 2.*

Proof. Since T^* is totally non-singular and satisfies $\dim T^* = 2$, we find from (2) that $\dim S \leq 4 - \dim T^* = 2$ for any singular matrix space $S \leq \mathbb{R}^{2 \times 2}$. Similarly, for any totally non-singular space T , (2) again shows $\dim T \leq 4 - \dim S^* = 2$. $\square$

Let us pause and reflect on how remarkable this is: we have proven an upper bound on the dimension of *any* singular matrix space S , simply by exhibiting *a single* totally non-singular matrix space T^* (and vice versa). That is, we have proved a “for all” statement by proving a “there exists” statement, which is (usually) a much simpler task.

We can encapsulate the preceding discussion in the following more general statement, which is again a simple corollary of (2).

Proposition 3. Suppose that $S^* \leq F^{n \times n}$ is singular and $T^* \leq F^{n \times n}$ is totally non-singular, and suppose that $\dim S^* + \dim T^* = n^2$. Then

$$\max_{\text{singular } S \leq F^{n \times n}} \dim S = \dim S^* \quad \text{and} \quad \max_{\text{totally non-singular } T \leq F^{n \times n}} \dim T = \dim T^*.$$

Proof. By (2), we have $\dim S \leq n^2 - \dim T^* = \dim S^*$ for any singular $S \leq F^{n \times n}$. Similarly, $\dim T \leq n^2 - \dim S^* = \dim T^*$ for any totally non-singular T . $\square$

Again, the point of this is that the existence of a *single* pair (S^*, T^*) as in the proposition statement implies a sharp upper bound on the dimension of *any* singular or totally non-singular space.

Before proceeding, let us remark on the similarity between Proposition 3 and the notion of *duality*, which is one of the most important concepts in mathematical optimization (and in particular in linear, semidefinite, and convex programming). In such areas, our goal is to solve an optimization problem of the form $\max_{x \in \mathcal{X}} f(x)$, where f is some “nice” (e.g. concave) function, and $\mathcal{X}$ is some “nice” (e.g. convex) domain. Under certain hypotheses, the theory of duality yields a *dual* optimization problem, of the form $\min_{y \in \mathcal{Y}} g(y)$, for some other nice function g and nice domain $\mathcal{Y}$, with the property that $f(x) \leq g(y)$ for every $x \in \mathcal{X}, y \in \mathcal{Y}$. In particular, this implies that the optima of the two optimization problems satisfy

$$\max_{x \in \mathcal{X}} f(x) \leq \min_{y \in \mathcal{Y}} g(y). \quad (4)$$

As a consequence, if we can find a single pair $(x^*, y^*) \in \mathcal{X} \times \mathcal{Y}$ satisfying $f(x^*) = g(y^*) = m$, then we conclude that

$$\max_{x \in \mathcal{X}} f(x) \geq f(x^*) = m = g(y^*) \geq \min_{y \in \mathcal{Y}} g(y).$$

Combined with (4), we see that all these inequalities must actually be equalities, hence we learn that this m is the optimum value of both optimization problems: we must have $\max_{x \in \mathcal{X}} f(x) = f(x^*) = m = g(y^*) = \min_{y \in \mathcal{Y}} g(y)$. In other words, proving the “for all” statement $f(x) \leq m$ for all $x \in \mathcal{X}$ boils down to proving the existence of the single pair (x^*, y^*) . Moreover, in the setting of convex optimization, we actually have a much stronger fact: such a pair (x^*, y^*) is guaranteed to exist (under mild hypotheses). That is, such an optimization problem can *always* be solved by exhibiting a single pair of solutions (x^*, y^*) .

At this level of abstraction, we see that Proposition 3 fits into exactly the same framework. Namely, let us define $\mathcal{X}$ to be the set of singular matrix spaces $S \leq F^{n \times n}$, and $\mathcal{Y}$ to be the set of totally non-singular matrix spaces $T \leq F^{n \times n}$. If we further set $f(S) = \dim S$ and $g(T) = n^2 - \dim T$, then (2) precisely states that $f(S) \leq g(T)$ for all $S \in \mathcal{X}, T \in \mathcal{Y}$. Additionally, Proposition 3 says nothing more than if we can find some $S^* \in \mathcal{X}, T^* \in \mathcal{Y}$ with $f(S^*) = g(T^*)$, then we automatically solve both optimization problems.

With that in mind, let us state the main theorem that we will prove using this duality approach.

Theorem 4. Let F be a field and n a positive integer. The maximum dimension of a singular space $S \leq F^{n \times n}$ is $n^2 - n$.

Theorem 4 was first proved by Dieudonné [8] under the assumption that $|F| > n$, an assumption that was finally lifted by Meshulam [19], who found a proof that works for all fields. Both proofs use somewhat involved algebraic techniques and properties of high-degree polynomials; our main goal

in this paper is to demonstrate that Theorem 4 can be proved much more simply, as a consequence of the duality inequality (2). Let us now see how this works.

First, the existence of a singular space $S^* \leq F^{n \times n}$ of dimension $\dim S^* = n^2 - n$ is not too hard; for example, we can take S^* to consist of all matrices whose last row is 0, as we did in Example 1. So, following the framework discussed above, all that remains to do is to find a totally non-singular space $T^* \leq F^{n \times n}$ of dimension $\dim T^* = n$, which will then imply Theorem 4 by Proposition 3.

Unfortunately, this last step is not always possible. For example, one can show that if F is algebraically closed (e.g. if $F = \mathbb{C}$), then any totally non-singular space in $F^{n \times n}$ has dimension at most 1. In order to circumvent this (serious) issue, we actually prove Theorem 4 by proving the following two lemmas.

Lemma 5. *If F is a finite field, then there is a totally non-singular $T^* \leq F^{n \times n}$ with $\dim T^* = n$. In particular, Theorem 4 holds whenever F is a finite field.*

Lemma 6. *If Theorem 4 holds whenever F is a finite field, then it holds for all fields.*

The proof of Lemma 5 is fairly short and straightforward, and actually works for a larger class of fields which we term *algebraically open fields*, which includes some infinite fields such as $\mathbb{Q}$. The construction of T^* for finite fields is fairly well-known, and the fact that it can be used to provide a proof of Theorem 4 for finite fields was also observed in [20, Corollary 4].

Thus, in some sense, our main contribution is the proof of Lemma 6, which is somewhat involved: we first use an elementary reduction to prove that Theorem 4 holds for the algebraically closed fields $\overline{\mathbb{F}_p}$ for all primes p , then use model-theoretic techniques (related to the proof of the Ax–Grothendieck theorem [5, 14]) to deduce that Theorem 4 holds for all algebraically closed fields, and finally use another elementary reduction to deduce the statement for all fields.

The reader might justly complain that they were promised a *simple* proof of Theorem 4, and this multi-step process looks rather convoluted. This is a fair criticism, but we believe that our proof of Theorem 4 has some advantages over the earlier proofs of Dieudonné and Meshulam. First, for many fields that we are interested in (e.g. the rationals and all finite fields), only Lemma 5 is necessary, and hence for such fields the proof really is quite simple and follows the duality framework laid out above. Note, in particular, that small finite fields are precisely the fields for which Dieudonné’s original proof fails, and we obtain a very simple proof for such fields.

Second, while the proof of Lemma 6 is somewhat involved, it is also extremely general, and uses almost nothing about the fact that we are working with singular spaces. In fact, we prove a much more general result (see Theorem 14 below), which states that for a wide variety of natural properties of matrix spaces, to upper-bound the maximum dimension of a space whose matrices all have this property, it suffices to solve this problem over all finite fields.

In particular, we are hopeful that the duality technique can be used to prove other extremal results in matrix spaces. As a proof of concept, we prove the following generalization of Theorem 4 (which is precisely the case $m = n, r = n - 1$ of the following result).

Theorem 7. *Let F be a field and let $r \leq m \leq n$ be positive integers. The maximum dimension of a space $R \leq F^{m \times n}$ all of whose elements have rank at most r is rn .*

The lower bound is again easy: the space R^* consisting of those matrices whose last $m - r$ rows are zero consists only of matrices of rank at most r , and has dimension rn . The upper bound in Theorem 7 was first proved by Flanders [12] under the assumption that $|F| > r$, and the general case was again proved by Meshulam [19]. Again, we provide a new duality-based proof of Theorem 7:

we first show that for all finite fields, we can construct a space $Q^* \leq F^{m \times n}$ with $\dim Q^* = (m-r)n$ such that all non-zero matrices in Q^* have rank strictly larger than r . We then have

$$\dim R^* + \dim Q^* = rn + (m-r)n = mn = \dim F^{m \times n},$$

which implies Theorem 7 for all finite fields, via the same argument as in Proposition 3. Finally, we then use (the general form of) Lemma 6 to deduce that the same result holds for all fields.

There are many other examples of extremal results on matrix spaces, and we hope that some of them can be proved along the same lines. A notable example is a theorem of Gerstenhaber and Serezhkin [13, 23], which states that if $N \leq F^{n \times n}$ consists entirely of nilpotent matrices, then $\dim N \leq \binom{n}{2}$. It would again be very interesting to find a proof of this theorem using duality, i.e. by exhibiting a space of dimension $n^2 - \binom{n}{2}$ none of whose non-zero elements is nilpotent. This is straightforward over $\mathbb{R}$, as the space of symmetric matrices does the job, but it is not clear to us how to do it (or whether it is even true) over finite fields.

Conjecture 8. *For every finite field F and every integer $n \geq 1$, there is a matrix space $W \leq F^{n \times n}$ with $\dim W = n^2 - \binom{n}{2}$ such that no non-zero matrix in W is nilpotent.*

If one could find such a construction, our general form of Lemma 6 would again yield the Gerstenhaber–Serezhkin theorem over all fields. A more ambitious goal would be to use this approach to prove a theorem of Atkinson [4], which simultaneously generalizes Theorem 4 and the Gerstenhaber–Serezhkin theorem; Atkinson’s proof is again restricted only to large fields, and to the best of our knowledge, there is currently no known proof of Atkinson’s theorem that works over all fields.

2 Algebraically open fields.

In this section, we prove Lemma 5. As indicated above, we actually prove it for a more general class of fields than just finite fields.

Before making a formal definition, let us think back to (3). There, we considered the space of all matrices of the form $\begin{pmatrix} x & y \\ -y & x \end{pmatrix}$ over $\mathbb{R}$, and observed that this space is totally non-singular. This matrix space is in fact an algebra (the product of two such matrices is another matrix of the same form). Moreover, a bit of thought will reveal that this algebra is isomorphic to $\mathbb{C}$: the matrix $\begin{pmatrix} x & y \\ -y & x \end{pmatrix}$ corresponds to the complex number $x + iy$. The fact that this matrix space is totally non-singular boils down to the fact that $\mathbb{C}$ is a field: every non-zero element in $\mathbb{C}$ is invertible.

With this understanding of the construction, it is fairly straightforward to generalize it as follows. We recall that a *field extension* $F \subset K$ is just an inclusion of one field in another; in any such situation, the field K is a vector space over F , and the *degree* of the extension is the dimension of this vector space.

Proposition 9. *Suppose that $F \subset K$ is a field extension of degree n . Then there is a totally non-singular matrix space $T^* \leq F^{n \times n}$ with $\dim T^* = n$.*

Proof. For every $\alpha \in K$, consider the multiplication map $m_\alpha : K \rightarrow K$ defined by $m_\alpha(x) = \alpha x$. This is an F -linear map, hence $m_\alpha \in \text{End}_F(K) \cong F^{n \times n}$. Moreover, we clearly have that $m_{\alpha+\beta} = m_\alpha + m_\beta$ and that $m_{\gamma\alpha} = \gamma m_\alpha$ for $\gamma \in F$, hence the set of all m_α defines a vector subspace of $F^{n \times n}$, and we define $T^* = \{m_\alpha : \alpha \in K\} \leq F^{n \times n}$ to be this subspace.

Moreover, as each non-zero α is invertible in K , we see that m_α is an invertible linear map. This implies that every non-zero element of T^* is invertible, hence T^* is totally non-singular. Moreover, it implies the F -linear map $\alpha \mapsto m_\alpha$ is injective, as no non-zero α can be in its kernel. As the domain of this map is K , which is n -dimensional, we conclude that $\dim T^* = n$. $\square$

This proposition is exactly what we wanted: it gives a general setting under which we can find a totally non-singular space of dimension n in $F^{n \times n}$. All we need to apply this is access to a field extension $K \supset F$ of degree n . This motivates the following definition³.

Definition 10. A field F is called *algebraically open* if, for any positive integer n , there is a field extension $K \supset F$ of degree n .

Thanks to Proposition 9, we see that if F is algebraically open, then there is an n -dimensional totally non-singular space in $F^{n \times n}$ for all n . Thus, in order to complete the proof of Lemma 5, it suffices to observe the following.

Lemma 11. *All finite fields are algebraically open.*

Proof. This is an immediate and well-known corollary of the classification of finite fields: for every prime power q and every integer n , there is a field $\mathbb{F}_{q^n}$ which is a degree- n extension of the field $\mathbb{F}_q$, implying that $\mathbb{F}_q$ is algebraically open. As every finite field is of the form $\mathbb{F}_q$ for some prime power q , we have the desired conclusion. $\square$

Combining Proposition 9 and Lemma 11, we obtain Lemma 5. Given that we only need Lemma 5 for finite fields, there is arguably not much advantage to defining algebraically open fields and proving Proposition 9 in this greater generality. However, the reason we did so is that many other fields besides finite fields are algebraically open. For example, it is a well-known consequence of properties of cyclotomic fields that $\mathbb{Q}$ is algebraically open (see e.g. [9, Corollary 14.28]). With more effort, one can prove that in fact, every number field is algebraically open⁴. Thus, Proposition 9 gives an extremely simple proof of Theorem 4 for all such fields.

The other reason we work in the greater generality of algebraically open fields is that, more generally, the existence of totally non-singular spaces over F is closely related to the set of field extensions of F . We have already mentioned that algebraically closed fields do not admit any non-trivial totally non-singular spaces. Over $\mathbb{R}$, the largest $n \times n$ totally non-singular spaces arise from the action of Clifford algebras, and this construction can be viewed as a generalization of Proposition 9 where we allow K to be a non-commutative ring. For a more precise statement of this connection, see [7, Section 4].

3 The reduction to general fields.

In this section, we state and prove a generalization of Lemma 6, which implies that for a wide array of extremal problems on matrix spaces, a bound that holds over finite fields immediately implies that the same bound holds over all fields. We begin by defining the class of questions that this

³We have not been able to find this notion in the literature, so do not think it has a standard name. The name “algebraically open” is meant to convey that such a field lies at the opposite extreme from an algebraically closed field, which has no non-trivial finite-degree extensions whatsoever.

⁴This also has an elementary proof using cyclotomic fields, or can alternatively be obtained as a very special case of the Grunwald–Wang theorem, a deep result in class field theory (see e.g. [3, Theorem X.6]).

reduction theorem applies to. Roughly speaking, this is the class of matrix properties that can be defined by polynomials with integer coefficients.

Before giving the formal definition, we recall that for every field F , there is a unique ring homomorphism $\mathbb{Z} \rightarrow F$. This map naturally extends to a map of polynomial rings $\mathbb{Z}[x_1, \dots, x_\ell] \rightarrow F[x_1, \dots, x_\ell]$, and this allows us to view any integer polynomial as a polynomial defined over F . We will frequently abuse notation and use the same symbol to refer both to the polynomial with coefficients in $\mathbb{Z}$ and the polynomial with coefficients in F .

Definition 12. Let $p_1, \dots, p_k \in \mathbb{Z}[x_{11}, \dots, x_{mn}]$ be polynomials with integer coefficients in the variables $\{x_{ij} : 1 \leq i \leq m, 1 \leq j \leq n\}$. Let $P = \{p_1, \dots, p_k\}$, and let $Z(P) \subseteq F^{m \times n}$ be the set of $m \times n$ matrices M satisfying $p_1(M) = \dots = p_k(M) = 0$; here, we view $p_1, \dots, p_k$ as polynomials over F , and we plug in the mn entries of M as the variables of the polynomials. If $M \in Z(P)$, we say that M *satisfies* P .

Finally, we say that a matrix space $W \subseteq F^{m \times n}$ is a P -space if $W \subseteq Z(P)$. That is, a P -space is a matrix space all of whose elements satisfy P .

Example 13. The determinant of an $n \times n$ matrix is an integer polynomial in the entries of the matrix. Hence, if we let $P_{\text{sing}} = \{\det\}$, then a matrix satisfies P_{sing} if and only if its determinant is 0, i.e. if and only if it is singular. Therefore, a P_{sing} -space is the same as a singular space.

Similarly, the property of having rank at most r can be encoded by polynomials. Namely, let $P_{\leq r}$ denote the set of all determinants of all $(r+1) \times (r+1)$ submatrices of the $m \times n$ matrix $\begin{pmatrix} x_{11} & \dots & x_{1n} \\ \vdots & \ddots & \vdots \\ x_{m1} & \dots & x_{mn} \end{pmatrix}$. Then a matrix satisfies $P_{\leq r}$ if and only if each of its $(r+1) \times (r+1)$ submatrices are singular, which happens if and only if its rank is at most r . Hence, a $P_{\leq r}$ -space is the same as a matrix space all of whose elements have rank at most r .

Another example is the property of nilpotency, which can be encoded by polynomials in at least two different ways. First, an $n \times n$ matrix M is nilpotent if and only if $M^n = 0$. The n^2 entries of M^n are each integer polynomials of the entries in M , so we can encode nilpotency by insisting that these n^2 polynomials are all zero. Second, a matrix is nilpotent if and only if its eigenvalues are all zero, which happens if and only if all coefficients of its characteristic polynomial, apart from the leading coefficient, are zero. The coefficients of the characteristic polynomial are themselves polynomials in the entries of the matrix, so we could instead demand that these n polynomials in the entries are all zero.⁵

Finally, totally non-singular spaces are a non-example: there is no collection of polynomials P such that a matrix space is totally non-singular if and only if it is a P -space.

We now introduce one final piece of notation: for a set $P \subseteq \mathbb{Z}[x_{11}, \dots, x_{mn}]$ of integer polynomials and a field F , let us denote by $\text{maxdim}_F(P)$ the maximum dimension of a P -space $W \subseteq F^{m \times n}$. With these definitions in hand, we can now state our general reduction theorem.

Theorem 14. Let $P \subseteq \mathbb{Z}[x_{11}, \dots, x_{mn}]$ be a finite set of integer polynomials. Suppose that

$$\text{maxdim}_F(P) \leq d \tag{5}$$

for every finite field F . Then the same holds for any field.

⁵As pointed out by an anonymous referee, the same idea can be used to obtain a different encoding of the property of rank at most r , at least for square matrices: an $n \times n$ matrix has rank at most r if and only if it has at least $n - r$ zero eigenvalues, which happens if and only if the $n - r$ lowest coefficients of its characteristic polynomial are zero.

In particular, this theorem immediately implies Lemma 6, which corresponds to the case $m = n$, $d = n^2 - n$, and $P = P_{\text{sing}} = \{\det\}$.

The proof of Theorem 14 breaks into three steps, which are accomplished in each of the following three subsections. For the remainder of this section, we treat P, m, n, d as fixed.

3.1 From finite fields to $\overline{\mathbb{F}_p}$.

We begin by showing that if (5) holds for all finite fields, then it holds for their algebraic closures $\overline{\mathbb{F}_p}$, for all primes p .

Thus, let us fix some P -space $W \leq \overline{\mathbb{F}_p}^{m \times n}$; our goal is to show that $\dim W \leq d$. We fix a basis $M_1, \dots, M_t$ for W . Recall that $\overline{\mathbb{F}_p}$ is the union of the finite fields $\mathbb{F}_{p^k}$ for all k ; this means that every element of $\overline{\mathbb{F}_p}$ is an element of $\mathbb{F}_{p^k}$ for some k . In particular, we can find some integer k such that each of the tmn entries of $M_1, \dots, M_t$ all lie in the field $F = \mathbb{F}_{p^k}$.

Let $W_F \leq F^{m \times n}$ be the matrix space over F spanned by these matrices. As $M_1, \dots, M_t$ are linearly independent over $\overline{\mathbb{F}_p}$, they are also linearly independent over the subfield F , hence $\dim W_F = \dim W$. Moreover, every matrix in W_F is also a matrix in W : it is an F -linear combination of $M_1, \dots, M_t$, hence it is also a $\overline{\mathbb{F}_p}$ -linear combination of $M_1, \dots, M_t$, hence it lies in W . In particular, every matrix in W_F satisfies P , since W is a P -space. As a consequence,

$$\dim W = \dim W_F \leq \max \dim_F(P) \leq d,$$

which is what we wanted to prove.

3.2 From $\overline{\mathbb{F}_p}$ to all algebraically closed fields via model theory.

We now turn to the crux of the argument, where we use model-theoretic tools to convert our knowledge of the bound (5) from the special algebraically closed fields $\overline{\mathbb{F}_p}$ to all algebraically closed fields. Such arguments are now fairly commonplace, and to our knowledge they originate with the classical proofs of the Ax–Grothendieck theorem [5, 14]. This is a non-trivial statement about polynomial maps $\mathbb{C}^n \rightarrow \mathbb{C}^n$, but the analogous statement for $\overline{\mathbb{F}_p}^n$ is rather simple; the observation of Ax and Grothendieck is that one can deduce the statement for $\mathbb{C}$ from the statement for $\overline{\mathbb{F}_p}$.

To state the main model-theoretic result we need, we recall that a *first-order sentence in the language of rings* is any statement that can be formed with the quantifiers $\forall$ and $\exists$, the logical operations $\neg, \wedge, \vee$, variables for ring elements, and the symbols $0, 1, +, -, \times$ corresponding to the additive and multiplicative identities, and to the basic operations in a ring.

A given first-order sentence ϕ can be true in some rings but not in others. For example, the sentence $1 + 1 = 0$ is true in rings of characteristic 2 but in no others. The sentence $\forall x : (x = 0) \vee (\exists y : x \times y = 1)$ states that every non-zero element has a multiplicative inverse, and thus is true in all fields but false in any ring that is not a field. The sentence $\exists x : x \times x + 1 = 0$ is true in any algebraically closed field, but not true in fields such as $\mathbb{R}$ in which -1 is not a perfect square.

Most importantly for our purposes, the statement “ $\max \dim_F(P) \leq d$ ” can be encoded as a first-order sentence in the language of rings, such that this sentence is true in F if and only if we actually have $\max \dim_F(P) \leq d$. Indeed, the statement $\max \dim_F(P) \leq d$ holds if and only if for all linearly independent $M_1, \dots, M_{d+1} \in F^{m \times n}$, there is some matrix in their span which does not satisfy P . The fact that $M_1, \dots, M_{d+1}$ are linearly independent can be expressed as a first-order sentence (for all $\lambda_1, \dots, \lambda_{d+1}$, either $\lambda_1 = \dots = \lambda_{d+1} = 0$, or the linear combination $\sum \lambda_i M_i$ is

non-zero), and the fact that there exists a matrix in the span not satisfying P is also a first-order statement (there exist $\lambda_1, \dots, \lambda_{d+1}$ such that for $M = \sum \lambda_i M_i$, either $p_1(M) \neq 0$, or $p_2(M) \neq 0$, or $\dots$). We remark that this is the only place in the proof where we use that the polynomials in P have integer coefficients: this fact means that every polynomial in P can be expressed using variables and the symbols $0, 1, +, -, \times$.

In particular, the next step of our reduction follows immediately from the following theorem.

Theorem 15. *Let ϕ be a first-order sentence in the language of rings. Then the following are equivalent.*

- (i) ϕ holds in the fields $\overline{\mathbb{F}_p}$ for all primes p .
- (ii) ϕ holds in all algebraically closed fields.

In particular, given that we now know that $\text{maxdim}_F(P) \leq d$ holds for $F = \overline{\mathbb{F}_p}$ for all p , we conclude that it holds for all algebraically closed fields.

As Theorem 15 may (and should!) appear completely magical to those unfamiliar with the model-theoretic background, we include a proof sketch of Theorem 15 in Section 4, and a full proof can be found in [18, Corollaries 2.2.9 and 2.2.10]. For further applications of this result and its variants, we recommend the survey [24]. However, for the moment, we continue with our proof of Theorem 14.

3.3 From algebraically closed fields to all fields.

Before we begin the final step, we record the following simple lemma.

Lemma 16. *Let F, K be fields with $F \subseteq K$. Let $v_1, \dots, v_t \in F^k$ be linearly independent over F . Then $v_1, \dots, v_t$ are also linearly independent over K , when viewed as vectors in the vector space K^k .*

Proof. Let V be the $k \times t$ matrix whose columns are $v_1, \dots, v_t$. As $v_1, \dots, v_t$ are linearly independent over F , we see that V has rank t , when viewed as a matrix over F . Therefore, there is some collection of t rows of V such that the corresponding $t \times t$ submatrix V' also has rank t over F . In particular, this implies that $\det V'$ is some non-zero element of F .

However, if we now turn to compute $\det V'$ over the field K , we will get the same answer, as we are working with the same matrix and the determinant is computed identically over all fields. Hence $\det V'$ is also a non-zero element of K , meaning that V' is invertible over K . Thus V' has rank t over K , and hence so does V . As a consequence, the columns of V , which are the vectors $v_1, \dots, v_t$, are linearly independent over K . $\square$

Let now F be a field, and let $\overline{F}$ be its algebraic closure. Thanks to Theorem 15, we know that $\text{maxdim}_{\overline{F}}(P) \leq d$. Our goal now is to prove that $\text{maxdim}_F(P) \leq d$ as well.

As the assumption of Theorem 14 is that (5) holds for all finite fields, we may assume that F is infinite. Suppose for contradiction that $\text{maxdim}_F(P) > d$, namely that there exists some P -space $W \in F^{m \times n}$ with $\dim W > d$. Fix a basis $M_1, \dots, M_t$ for W , and note that by Lemma 16, $M_1, \dots, M_t$ are linearly independent over $\overline{F}$. Therefore, if we let $\overline{W} \in \overline{F}^{m \times n}$ be the $\overline{F}$ -span of $M_1, \dots, M_t$ we have that $\dim \overline{W} = \dim W > d \geq \text{maxdim}_{\overline{F}}(P)$. Therefore, $\overline{W}$ must not be a P -space, hence there exists some $M \in \overline{W}$ not satisfying P . In particular, there is some fixed polynomial $p \in P$ such that $p(M) \neq 0$.

Let us define the polynomial $q \in F[y_1, \dots, y_t]$ by

$$q(y_1, \dots, y_t) = p(y_1 M_1 + \dots + y_t M_t).$$

Since the entries of $M_1, \dots, M_t$ are in F , we see that q is indeed a polynomial with coefficients in F . Since $F \subseteq \overline{F}$, we can also view it as a polynomial over $\overline{F}$.

Since $M \in \overline{W}$, there exist some $\lambda_1, \dots, \lambda_t \in \overline{F}$ such that $M = \lambda_1 M_1 + \dots + \lambda_t M_t$. And since $p(M) \neq 0$, we conclude that $q(\lambda_1, \dots, \lambda_t) \neq 0$. In particular, we see that q is not the zero polynomial; if it were, we would have $q(\lambda_1, \dots, \lambda_t) = 0$ for all $(\lambda_1, \dots, \lambda_t) \in \overline{F}^t$.

To conclude, we will use the following simple and well-known lemma about polynomials.

Lemma 17. *If F is an infinite field and $q \in F[y_1, \dots, y_t]$ is not the zero polynomial, then there exists some $(\mu_1, \dots, \mu_t) \in F^t$ such that $q(\mu_1, \dots, \mu_t) \neq 0$.*

Lemma 17 is essentially (a weak form of) the Schwartz–Zippel lemma, and we include a short proof for completeness momentarily. But before we do, let us conclude the proof of Theorem 14. Applying Lemma 17 to the polynomial q defined above, we obtain $(\mu_1, \dots, \mu_t) \in F^t$ with $q(\mu_1, \dots, \mu_t) \neq 0$. Let $M' = \mu_1 M_1 + \dots + \mu_t M_t$, and note that $M' \in W$, as M' is an F -linear combination of the matrices $M_1, \dots, M_t$. Moreover, since $q(\mu_1, \dots, \mu_t) \neq 0$, we conclude that $p(M') \neq 0$. But this contradicts our assumption that W is a P -space, completing the proof.

As promised, we now present the proof of Lemma 17.

Proof of Lemma 17. We proceed by induction on t . The base case $t = 1$ is immediate, since a non-zero polynomial in one variable has at most as many roots as its degree, and in particular has only finitely many roots; since F is infinite it must have some non-root in F .

For the inductive step, let us write

$$q(y_1, \dots, y_t) = \sum_{i=0}^s y_t^i q_i(y_1, \dots, y_{t-1})$$

for some polynomials $q_0, \dots, q_s \in F[y_1, \dots, y_{t-1}]$, where s is the degree of y_t in q . Note that we may assume that $s > 0$, since otherwise y_t does not appear at all in q , and we are already done by the inductive hypothesis. This in particular implies that q_s is not the zero polynomial. By the inductive hypothesis, we can find some $\mu_1, \dots, \mu_{t-1} \in F$ with $q_s(\mu_1, \dots, \mu_{t-1}) \neq 0$. But this implies that $q(\mu_1, \dots, \mu_{t-1}, y_t)$, viewed as a one-variable polynomial, is not the zero polynomial. So by again applying the base case we can find some μ_t such that $q(\mu_1, \dots, \mu_t) \neq 0$, as claimed. $\square$

4 Proof sketch of Theorem 15.

In this section, we give a brief sketch of the model-theoretic background needed to prove Theorem 15. We stress that while this sketch conveys all of the key ideas, it is by no means a full proof, and we refer to a standard introductory text in model theory such as [18] for more formal definitions and proofs.

For our purposes, a *theory* is simply a collection of axioms, each of which is a first-order sentence in the language of rings⁶. For example, the theory of rings consists of all of the usual ring axioms,

⁶Of course, there is no reason to restrict oneself to the language of rings, and indeed the real definition of a theory works with an arbitrary language. But our only goal is Theorem 15, so we restrict ourselves to this less general setting. We also assume throughout that our rings are commutative and have a multiplicative identity.

e.g. the axiom “ $\forall x : x + 0 = x$ ” that 0 is the additive identity, the axiom “ $\forall x \forall y : x + y = y + x$ ” that addition is commutative, and the axiom “ $\forall x \forall y \forall z : x \times (y + z) = (x \times y) + (x \times z)$ ” that multiplication distributes over addition. If we add to these axioms the axiom “ $\forall x : (x = 0) \vee (\exists y : x \times y = 1)$ ” that every non-zero element is invertible, we obtain the theory of fields. We can then add further axioms to get the theory of algebraically closed fields; this involves adding, for every degree d , an axiom stating that every monic polynomial of degree d has a root.

A closely related concept is that of a *model*. Loosely speaking, a model $\mathcal{M}$ of a theory $\mathcal{T}$ is an object in which all the axioms of $\mathcal{T}$ hold. Thus, for instance, $\mathbb{Z}$ is a model of the theory of rings, but is not a model of the theory of fields; this is saying nothing more but that $\mathbb{Z}$ is a ring and is not a field.

Given a theory $\mathcal{T}$ and a sentence ϕ , we say that $\mathcal{T}$ *proves* ϕ if ϕ is true in every model $\mathcal{M}$ of $\mathcal{T}$. For example, an elementary exercise in any first course on rings is that additive inverses are unique; this means that the theory of rings proves the sentence $\forall x \forall y \forall z : (y = z) \vee (x + y \neq 0) \vee (x + z \neq 0)$. The reader might complain that this is not what “prove” should mean. Instead, it should mean that one can perform a series of logical deductions to derive ϕ from the axioms of $\mathcal{T}$. Indeed, when one solves the elementary ring theory exercise above, one does not actually verify that this property holds by inspecting all rings one by one, but instead derives this property from the axioms of rings. Luckily, the famous *completeness theorem* of Gödel says that these two notions of proof are equivalent: one can logically deduce ϕ from the axioms in $\mathcal{T}$ if and only if ϕ is true in every model of $\mathcal{T}$.

A theory is called *complete* if it can determine the veracity of any statement, i.e. for every sentence ϕ , it can either prove ϕ or $\neg\phi$. This is a very strong condition: it implies that for every first-order sentence ϕ , either ϕ is true in all models or it is false in all models. Thus, for example, we can easily see that the theory of rings is incomplete: the fact that some rings are fields and some are not implies that the theory of rings cannot prove or disprove the axiom “all non-zero elements have a multiplicative inverse”, and hence the theory of rings is incomplete. Similarly, the theory of fields is incomplete, since some fields are algebraically closed and others are not. And again, the theory of algebraically closed fields is incomplete, because it can neither prove nor disprove the sentence “ $1 + 1 = 0$ ”; this is because there are algebraically closed fields of characteristic 2, and also algebraically closed fields of other characteristics.

Rather remarkably, however, this is essentially the only obstruction to completeness. More precisely, let ACF denote the theory of algebraically closed fields. For every prime p , let χ_p denote the sentence $1 + \dots + 1 = 0$, where there are p ones on the left-hand side. Then let $\text{ACF}_p = \text{ACF} \cup \{\chi_p\}$ denote the theory of algebraically closed fields of characteristic p , and let $\text{ACF}_0 = \text{ACF} \cup \{\neg\chi_2, \neg\chi_3, \neg\chi_5, \dots\}$ denote the theory⁷ of algebraically closed fields of characteristic 0. Then we have the following remarkable theorem.

Theorem 18. *The theories ACF_0 and ACF_p , for every prime p , are complete.*

The proof of Theorem 18 makes it seem a bit less remarkable. Indeed, unpacking what it means, it states that for every sentence ϕ , either ϕ or $\neg\phi$ is true in every algebraically closed field of a given characteristic. But the reason for this is just that there are not so many different algebraically closed fields of a given characteristic: a standard result in field theory is that algebraically closed fields are

⁷We are implicitly using here the fact that every algebraically closed field either has characteristic 0 or characteristic p for some prime p , hence if we add in the negations of all the axioms χ_p , we get the theory of algebraically closed fields of characteristic 0.

uniquely determined by their characteristic and their transcendence degree (see e.g. [21, Theorem 9.16]). To prove Theorem 18, suppose for contradiction that F, F' are two algebraically closed fields of the same characteristic, such that ϕ is true in F but false in F' . By adding uncountably many “dummy elements” (see [18, Proposition 2.2.2] for details), we can moreover assume that F and F' have the same uncountable cardinality. But an uncountable field has transcendence degree equal to its cardinality, hence F and F' must be isomorphic, a contradiction.

Given Theorem 18, it is quite easy to prove Theorem 15.

Proof of Theorem 15. Clearly, (ii) implies (i), so we only need to prove the converse. So suppose that ϕ is true in $\overline{\mathbb{F}}_p$ for all primes p . This means that ϕ is true in some model of ACF_p (namely $\overline{\mathbb{F}}_p$), hence it must be true in all models of ACF_p , by the completeness of ACF_p . That is, ϕ is true in every algebraically closed field of positive characteristic.

So all that remains is to handle the case of characteristic 0. Suppose for contradiction that ACF_0 does not prove ϕ . By the completeness of ACF_0 , this means that ACF_0 proves $\neg\phi$. Recall that, by the completeness theorem, this means that one can deduce $\neg\phi$ from the axioms of ACF_0 via some finite sequence of logical deductions. In this finite sequence of deductions, we must only use finitely many axioms of ACF_0 , and in particular only use finitely many of the axioms $\neg\chi_p$. But if p_0 is some other prime not in this finite list, then the same proof is also valid in the theory ACF_{p_0} . Unpacking what all this means, we have found that if ACF_0 proves $\neg\phi$, then ACF_{p_0} proves $\neg\phi$ as well, for all but finitely many primes p_0 . But this is impossible, as we just saw that ACF_{p_0} proves ϕ for all p_0 . $\square$

5 Proof of Theorem 7.

In this section, we use the same duality machinery we have developed to prove Theorem 7. Recall the statement: for all integers $r \leq m \leq n$, if every matrix in $R \leq F^{m \times n}$ has rank at most r , then $\dim R \leq rn$. As the property of having rank at most r can be described by integer polynomials (Example 13), Theorem 14 implies that it suffices to prove this statement for finite fields. And over finite fields, we have the following elegant result of Roth [22].

Theorem 19. *If F is a finite field and $s \leq n$ are non-negative integers, then there is a matrix space $Q \leq F^{n \times n}$ of dimension $\dim Q = (n - s)n$, all of whose non-zero elements have rank greater than s .*

Note that the case $s = n - 1$ is precisely Lemma 5. Although Theorem 19 is stated only for $n \times n$ matrices, it gives as a simple corollary an analogous result for rectangular matrices.

Corollary 20. *If F is a finite field and $r \leq m \leq n$ are integers, there is a matrix space $Q^* \leq F^{m \times n}$ of dimension $\dim Q^* = (m - r)n$, all of whose non-zero elements have rank greater than r .*

Proof. Let $Q \leq F^{n \times n}$ be the space given by Theorem 19, applied with parameter $s = n - m + r$, so that $\dim Q = (n - s)n = (m - r)n$. Let Q^* be obtained from Q by deleting the last $n - m$ rows of each matrix, to obtain a matrix space in $F^{m \times n}$. More formally, letting $\pi : F^{n \times n} \rightarrow F^{m \times n}$ be the projection sending each matrix to the submatrix given by the first m rows, we define $Q^* = \pi(Q)$. By the rank-nullity theorem, we thus have

$$\dim Q^* = \dim Q - \dim(Q \cap \ker \pi).$$

We claim that $Q \cap \ker \pi = \{0\}$. Indeed, if M is a non-zero matrix in $Q \cap \ker \pi$, then M must have rank greater than $s \geq n - m$ (by the definition of Q), and must also be supported entirely in the last $n - m$ rows (by the definition of π), and these two properties are incompatible. Therefore $\dim(Q \cap \ker \pi) = 0$, and we find that $\dim Q^* = \dim Q = (m - r)n$.

Finally, we claim that every non-zero matrix in Q^* has rank greater than r . Indeed, fix some non-zero matrix $A \in Q^*$, and fix some $M \in Q$ with $\pi(M) = A$. As $M \in Q$, we know that $\text{rank } M > s$, hence there is a collection of more than s linearly independent rows in M . Of these, at most $n - m$ are in the last $n - m$ rows, hence A must have a collection of more than $s - (n - m) = r$ linearly independent rows, implying that $\text{rank } A > r$, as claimed. $\square$

Given these ingredients, the proof of Theorem 7 follows immediately.

Proof of Theorem 7. By Theorem 14, it suffices to prove the statement for finite fields F . We already saw that the space $R^* \leq F^{m \times n}$ of all matrices supported in the first r rows has dimension rn and all of its elements have rank at most r . By Corollary 20, we also have some $Q^* \leq F^{m \times n}$ of dimension $(m - r)n$, all of whose non-zero elements have rank greater than r . Since $\dim R^* + \dim Q^* = mn$, we conclude that both of these spaces must be of maximum dimension under these constraints, as desired. $\square$

For completeness, let us also see a proof of Theorem 19, due to Meshulam [20].

Proof of Theorem 19. Let $F = \mathbb{F}_q$ be a finite field, and let $k = n - s$. If $s = n$ we may take $Q = \{0\}$, so we assume henceforth that $k > 0$. Consider the set Q of all polynomials of the form

$$f(x) = \sum_{i=0}^{k-1} a_i x^{q^i},$$

where $a_i \in \mathbb{F}_{q^n}$, and where we view the variable x as also living in the extension field $\mathbb{F}_{q^n}$. The first observation is that each such $f \in Q$ is an $\mathbb{F}_q$ -linear map $\mathbb{F}_{q^n} \rightarrow \mathbb{F}_{q^n}$, since $(x + y)^q = x^q + y^q$ for every $x, y \in \mathbb{F}_{q^n}$. Hence we may view Q as a subset of $\text{End}_{\mathbb{F}_q}(\mathbb{F}_{q^n}) \cong \mathbb{F}_q^{n \times n}$. The second observation is that every non-zero f in Q is a polynomial of degree at most q^{k-1} , hence has at most q^{k-1} roots in $\mathbb{F}_{q^n}$. This means that, viewed as a linear function, its kernel has dimension at most $k - 1 < k$. By the rank-nullity theorem, this means that every non-zero element of Q has rank greater than $n - k = s$. It also implies that the linear map sending a coefficient vector $(a_0, \dots, a_{k-1})$ to the corresponding element $\sum a_i x^{q^i}$ of Q is injective, and thus that distinct coefficient vectors $(a_0, \dots, a_{k-1})$ yield distinct elements of Q . The final observation is that any $\mathbb{F}_q$ -linear combination of two $f, f' \in Q$ is another element of Q , hence Q is actually a subspace of $\mathbb{F}_q^{n \times n}$. We have q^n choices for each of the k coefficients a_i , hence $|Q| = (q^n)^k$, which in turn implies that $\dim Q = kn = (n - s)n$, as desired. $\square$

Just as we proved Lemma 5 for all algebraically open fields, and not only for finite fields, we could also extend this proof to cover a wider class of fields. Namely, let us call a field F *cyclically open* if for all positive integers n , there is a cyclic field extension $K \supset F$ of degree n . Then one can check that (essentially) the same proof works to show that if F is cyclically open, then there is $Q \leq F^{n \times n}$ of dimension $\dim Q = (n - s)n$, all of whose non-zero elements have rank greater than s . Rather than working with the polynomials f as defined above, instead one works with (formal) polynomials in the variable σ , where σ is a generator of the Galois group $\text{Gal}(K/F)$, which is assumed to be cyclic of order n . Each such polynomial yields an F -linear map $K \rightarrow K$, and one

can verify that the kernel of this map has dimension strictly less than the degree of the polynomial, which was the key property we used in the proof of Theorem 19.

As with algebraically open fields, there are numerous examples of cyclically open fields that are not finite, such as $\mathbb{Q}$ and all number fields [3, Theorem X.6]. In fact, we are not aware of any example of an algebraically open field that is not cyclically open (although we believe that such examples must surely exist).

Acknowledgments: I am indebted to Vivian Kuperberg, Youming Qiao, and Avi Wigderson for enlightening discussions on this topic, and to the anonymous referees for helpful comments on an earlier version of this article.

AI disclosure statement ChatGPT 5.5 was used during the final revision of this manuscript to check for typos. All of the ideas and writing in this paper are due entirely to the author.

References

- [1] J. F. Adams, Vector fields on spheres, *Ann. of Math. (2)* **75** (1962), 603–632.
- [2] J. F. Adams, P. D. Lax, and R. S. Phillips, On matrices whose real linear combinations are non-singular, *Proc. Amer. Math. Soc.* **16** (1965), 318–322.
- [3] E. Artin and J. Tate, *Class field theory*, AMS Chelsea Publishing, Providence, RI, 2009. Reprinted with corrections from the 1967 original.
- [4] M. D. Atkinson, Spaces of matrices with several zero eigenvalues, *Bull. London Math. Soc.* **12** (1980), 89–95.
- [5] J. Ax, The elementary theory of finite fields, *Ann. of Math. (2)* **88** (1968), 239–271.
- [6] R. Baer, Groups with abelian central quotient group, *Trans. Amer. Math. Soc.* **44** (1938), 357–386.
- [7] C. de Seguins Pazzis, The singular linear preservers of non-singular matrices, *Linear Algebra Appl.* **433** (2010), 483–490.
- [8] J. Dieudonné, Sur une généralisation du groupe orthogonal à quatre variables, *Arch. Math.* **1** (1949), 282–287.
- [9] D. S. Dummit and R. M. Foote, *Abstract algebra*, third ed., John Wiley & Sons, Inc., Hoboken, NJ, 2004.
- [10] J. Edmonds, Systems of distinct representatives and linear algebra, *J. Res. Nat. Bur. Standards Sect. B* **71B** (1967), 241–245.
- [11] D. Eisenbud and J. Harris, Vector spaces of matrices of low rank, *Adv. Math.* **70** (1988), 135–155.
- [12] H. Flanders, On spaces of linear transformations with bounded rank, *J. London Math. Soc.* **37** (1962), 10–16.

- [13] M. Gerstenhaber, On nilalgebras and linear varieties of nilpotent matrices. I, *Amer. J. Math.* **80** (1958), 614–622.
- [14] A. Grothendieck, Éléments de géométrie algébrique. IV. Étude locale des schémas et des morphismes de schémas. III, *Inst. Hautes Études Sci. Publ. Math.* **28** (1966), 5–255.
- [15] V. Kabanets and R. Impagliazzo, Derandomizing polynomial identity tests means proving circuit lower bounds, *Comput. Complexity* **13** (2004), 1–46.
- [16] Y. Li, Y. Qiao, A. Wigderson, Y. Wigderson, and C. Zhang, Connections between graphs and matrix spaces, *Israel J. Math.* **256** (2023), 513–580.
- [17] L. Lovász, Singular spaces of matrices and their application in combinatorics, *Bol. Soc. Brasil. Mat. (N.S.)* **20** (1989), 87–99.
- [18] D. Marker, *Model theory: an introduction*, *Graduate Texts in Mathematics*, vol. 217, Springer-Verlag, New York, 2002.
- [19] R. Meshulam, On the maximal rank in a subspace of matrices, *Quart. J. Math. Oxford Ser. (2)* **36** (1985), 225–229.
- [20] R. Meshulam, Spaces of Hankel matrices over finite fields, *Linear Algebra Appl.* **218** (1995), 73–76.
- [21] J. S. Milne, *Fields and Galois theory*, Kea Books, Ann Arbor, MI, 2022.
- [22] R. M. Roth, Maximum-rank array codes and their application to crisscross error correction, *IEEE Trans. Inform. Theory* **37** (1991), 328–336.
- [23] V. N. Serezhkin, Linear transformations preserving nilpotency, *Vestsī Akad. Navuk BSSR Ser. Fiz.-Mat. Navuk* (1985), 46–50, 125.
- [24] J.-P. Serre, How to use finite fields for problems concerning infinite fields, in *Arithmetic, geometry, cryptography and coding theory*, *Contemp. Math.*, vol. 487, Amer. Math. Soc., Providence, RI, 2009, 183–193.
- [25] A. Shpilka and A. Yehudayoff, Arithmetic circuits: a survey of recent results and open questions, *Found. Trends Theor. Comput. Sci.* **5** (2009), 207–388.
- [26] L. G. Valiant, Completeness classes in algebra, in *Conference Record of the Eleventh Annual ACM Symposium on Theory of Computing (Atlanta, Ga., 1979)*, ACM, New York, 1979, 249–261.
- [27] N. Weaver, Quantum graphs as quantum relations, *J. Geom. Anal.* **31** (2021), 9090–9112.