

Let \mathbb{F} be a field (usually in this talk you lose nothing from assuming $\mathbb{F} = \mathbb{R}$ or $\mathbb{F} = \mathbb{C}$). Let $\mathbb{F}^{m \times n}$ denote the vector space of $m \times n$ matrices with entries in \mathbb{F} . In this talk, we will be studying subspaces of the vector space $\mathbb{F}^{m \times n}$; in other words, we will be studying *matrix spaces*, i.e. linear spaces of matrices.

At first, it's not clear if there's anything to study here. Indeed, we know that all vector spaces of the same dimension over the same field are isomorphic, so it's not clear why a d -dimensional subspace of $\mathbb{F}^{m \times n}$ should be any different from \mathbb{F}^d . However, we gain additional structure from the fact that we view the elements as matrices: namely, for every property of matrices, we can ask about matrix spaces where all elements have that property. Here are a few examples we will discuss later.

- What can we say about a matrix space $S \subseteq \mathbb{F}^{m \times n}$ if every $M \in S$ has rank at most r ?
- What can we say about a matrix space $S \subseteq \mathbb{F}^{n \times n}$ if every $M \in S$ is singular?
- What can we say about a matrix space $S \subseteq \mathbb{F}^{n \times n}$ if every *non-zero* $M \in S$ is *non-singular*?
- What can we say about a matrix space $S \subseteq \mathbb{F}^{n \times n}$ if every $M \in S$ is nilpotent?

In general, answering these questions can be very difficult. For example, finding a polynomial-time algorithm to determine whether a matrix space $S \subseteq \mathbb{F}^{n \times n}$ is *singular* (meaning that every matrix in it is singular) is a famously difficult problem in theoretical computer science. In fact, by combining results of Valiant and Impagliazzo–Kabanets, there is a remarkable fact: the existence of a polynomial-time algorithm for this problem implies a result that is very akin to $P \neq NP$. More precisely, if such a polynomial-time algorithm exists, then one obtains unconditional lower bounds on either arithmetic or Boolean circuits: either $VNP \not\subseteq VP$ or $NEXP \not\subseteq P/\text{poly}$. I won't explain what this means, nor how to prove it, but it shows that there is a surprising depth to these innocent-looking problems.

1 Thirteen¹ ways of looking at a matrix space

We can always think of a matrix space $S \subseteq \mathbb{F}^{m \times n}$ as just that—a vector subspace of $\mathbb{F}^{m \times n}$. However, there is another very useful perspective on matrix spaces, which also reveals the close connection between this topic and abstract algebra.

Given a matrix space $S \subseteq \mathbb{F}^{m \times n}$, we may pick a basis A_1, \dots, A_d of S , so that each A_i is an $m \times n$ matrix. Then every element of S can be written as $x_1 A_1 + \dots + x_d A_d$ for some $x_1, \dots, x_d \in \mathbb{F}$. If we treat x_1, \dots, x_d as variables (or just abstract symbols), we can add all of this and find that a generic element of S is a matrix whose entries are homogeneous linear polynomials in x_1, \dots, x_d . Namely, the (i, j) entry of some generic element of S is simply $(A_1)_{i,j} x_1 + \dots + (A_d)_{i,j} x_d$, where here $(A_1)_{i,j}, \dots, (A_d)_{i,j}$ are scalars in \mathbb{F} .

Definition 1.1. A *symbolic matrix* over \mathbb{F} is a matrix whose entries are homogeneous linear polynomials in variables x_1, \dots, x_d , where the coefficients of the polynomials are from \mathbb{F} . In other words, a symbolic matrix over \mathbb{F} is simply a matrix over some polynomial ring $\mathbb{F}[x_1, \dots, x_d]$, where we insist that all the entries are linear forms.

¹Here, $13 = 2$.

The discussion above shows that given a matrix space S with a fixed basis A_1, \dots, A_d , we may create a symbolic matrix which we will call $X(S)$ (note that we don't record the dependence on the basis, since it won't end up mattering to us). Conversely, given a symbolic matrix X , we may form a matrix space S by simply taking all possible evaluations of the variables in \mathbb{F} . Since we insist on the entries of our symbolic matrices being linear forms, it is clear that doing this will yield a matrix space, and it is also clear that doing this to $X(S)$ recovers the space S .

In particular, when discussing matrix spaces, it is often easiest to represent them as symbolic matrices. For example, rather than speaking of the matrix space spanned by $\begin{pmatrix} 3 & 2 \\ -1 & 0 \end{pmatrix}$ and $\begin{pmatrix} 1 & 0 \\ -2 & 2 \end{pmatrix}$, we can simply discuss the symbolic matrix $\begin{pmatrix} 3x+y & 2x \\ -x-2y & 2y \end{pmatrix}$.

The key fact we will need is that this conversion between matrix spaces and symbolic matrices allows us to translate properties from one world to another: every matrix in S has some property if and only if the *single* matrix $X(S)$ has that same property, when viewed as a matrix over a polynomial ring. Two specific instances of this fact are in the following lemma, and it is usually easy to generalize this to any other property you might be interested in.

Lemma 1.2. *Let $S \subseteq \mathbb{F}^{n \times n}$ be a matrix space, and let $X(S) \in \mathbb{F}[x_1, \dots, x_d]^{n \times n}$ be the corresponding symbolic matrix. The following hold assuming $|\mathbb{F}|$ is sufficiently large ($|\mathbb{F}| > n$ suffices).*

(a) *Every $M \in S$ is singular if and only if $X(S)$ is singular.*

(b) *Every $M \in S$ is nilpotent if and only if $X(S)$ is nilpotent.*

Proof. To prove (a), first suppose that $X(S)$ is singular. This implies that $\det(X(S))$, which is an element of $\mathbb{F}[x_1, \dots, x_d]$, is the zero polynomial. That shows that for every evaluation of x_1, \dots, x_d , the corresponding scalar matrix will have zero determinant, because we are simply evaluating the zero polynomial at some point $(x_1, \dots, x_d) = (a_1, \dots, a_d)$. Since every $M \in S$ can be obtained from $X(S)$ by some such evaluation, we conclude that $\det(M) = 0$ for all $M \in S$. For the converse, suppose that $\det(X(S))$ is not the zero polynomial. Since $|\mathbb{F}|$ is sufficiently large², there is some evaluation $(x_1, \dots, x_d) = (a_1, \dots, a_d)$ such that $\det(X(S))(a_1, \dots, a_d)$ is not zero. But this evaluation yields some matrix $M \in S$, whose determinant is non-zero.

For (b), the proof is nearly identical. A single $n \times n$ matrix M is nilpotent if and only if $M^n = 0$. Every entry of $X(S)^n$ is some polynomial in the entries of $X(S)$, and thus some polynomial in x_1, \dots, x_d . If $X(S)$ is nilpotent, then all these polynomials are identically zero, so every evaluation of them is zero, so every $M \in S$ is nilpotent. Conversely, if one of these polynomials is non-zero, then there is some non-zero evaluation, and thus some $M \in S$ which is not nilpotent. \square

Remark. The assumption that $|\mathbb{F}|$ is large is actually necessary. For example, suppose

²By being careful here, one can obtain the bound $|\mathbb{F}| > n$: we see that $\det(X(S))$ has degree n , and the Schwartz-Zippel lemma implies that no polynomial of degree n can vanish everywhere on a field of size larger than n .

$\mathbb{F} = \mathbb{F}_2$. Then consider the matrix space given by

$$X(S) = \begin{pmatrix} 0 & x & 0 \\ 0 & 0 & y \\ x + y & 0 & 0 \end{pmatrix}$$

It is easy to see that for every choice of $x, y \in \mathbb{F}_2$, the corresponding matrix is singular: it will have at least one all-zero row. However, $\det(X(S)) = x^2y + xy^2$, which is not the zero polynomial.

Similarly, we see that $X(S)$ is not nilpotent (it certainly can't be nilpotent if it's non-singular). However, for every evaluation of $x, y \in \mathbb{F}_2$, the matrix one gets is similar to a strictly upper-triangular matrix, and thus is nilpotent.

In particular, when dealing with infinite fields like \mathbb{C} , we see that we lose nothing by studying properties of the symbolic matrix $X(S)$ when we care about properties of matrices in S . We also discover a connection to algebraic geometry. For example, every element of a matrix space S is singular if and only if a *single* polynomial whose coefficients depend on S —obtained as $\det(X(S))$ —is the zero polynomial.

One other consequence is that this seems to give us a nice algorithm for determining if a matrix space $S \subseteq \mathbb{F}^{n \times n}$ is singular. Namely, compute the symbolic matrix $X(S)$, then compute the polynomial $\det(X(S))$ (which can be done in polynomial time by Gaussian elimination), and then just check if this is the zero polynomial. The unfortunate issue is that $\det(X(S))$ is a polynomial of degree n in d variables. Such a polynomial will, in general, have roughly d^n monomials, and this is exponential in n . So even *writing down* this polynomial, when expanded into monomials, will take exponential time. However, all is not lost, and this technique does give a *randomized* polynomial-time algorithm for checking if S is symbolic. Namely, pick a random evaluation for each of the variables in $X(S)$, compute the determinant of that, and see if it's 0. If $\det(X(S))$ is not the zero polynomial, then for a random assignment of variables, we will get a non-zero value with high probability. If $\det(X(S))$ is the zero polynomial, then we will certainly get 0. In other words, if there *is* some non-singular linear combination of a collection of matrices, then a *random* linear combination will be non-singular with high probability.

This simple algorithm remains the best thing we know how to do for this problem, and in fact, the main point of the Impagliazzo–Kabanets result mentioned earlier—that finding a deterministic algorithm would yield computational lower bounds—is really a statement about derandomization. If we can derandomize this probabilistic algorithm, we can prove hardness results.

2 Matrix spaces of bounded rank

Given that fully understanding singular matrix spaces seems hard, it makes sense to start with easier problems. For example, can we bound the dimension of a singular space of matrices?

We can certainly say something: for example, if $S \subseteq \mathbb{F}^{n \times n}$ has dimension n^2 , then it must contain every $n \times n$ matrix, and thus cannot be singular. So if S is singular, then

$\dim S \leq n^2 - 1$. On the other hand, it is not too hard to come up with a singular matrix space of dimension $n^2 - n$: simply take the space of all matrices whose last row is 0. Where does the truth lie between $n^2 - n$ and $n^2 - 1$?

This question was first studied by Dieudonné in 1948, who solved it in order to address an issue in invariant theory. Namely, Dieudonné was interested in characterizing the group of symmetries of the space of singular matrices, i.e. the invariants of the algebraic variety in \mathbb{F}^{n^2} cut out by the single polynomial equation $\det(X) = 0$. These symmetries were actually already characterized by Frobenius in 1897 (apparently Dieudonné did not know this when he studied this problem), using a different technique. Dieudonné's technique, which is algebro-geometric in nature, required classifying the maximal-dimensional singular matrix spaces. In fact, there is a close connection between classifying maximal singular spaces and classifying certain torsion-free sheaves on projective spaces, but I don't really understand what this means.

Anyway, Dieudonné proved that the lower bound is correct, namely that an $n \times n$ singular matrix space has dimension at most $n^2 - n$. There are two natural ways to generalize Dieudonné's theorem. The first, considered by Flanders, is to more generally study a space of $n \times n$ matrices all of which have rank at most r , for some $r < n$. In this case, a natural lower bound is nr , obtained by taking the space of all matrices whose last $n - r$ rows are all zero. Flanders proved that this lower bound is tight, at least when $|\mathbb{F}|$ is large enough. The second natural generalization, considered (sort of implicitly) by Meshulam, is to move beyond square matrices, and to consider $m \times n$ matrix spaces. In this case, we again have a natural lower bound of $\max\{mr, nr\}$, obtained by taking either $m - r$ zero rows or $n - r$ zero columns. Again, this lower bound turns out to be tight (as shown by Meshulam), without any restriction on $|\mathbb{F}|$.

Theorem 2.1 (Dieudonné, Flanders, Meshulam). *Let $S \subseteq \mathbb{F}^{m \times n}$ have the property that every matrix in S has rank at most r . Then $\dim S \leq \max\{mr, nr\}$.*

I will present Meshulam's proof, which reveals a surprising connection between this topic and graph theory. His proof needs a little bit of setup. First, let \prec denote the lexicographic ordering on $[m] \times [n]$, i.e. the ordering where $(i, j) \prec (i', j')$ if and only if $i < i'$ or $i = i'$ and $j < j'$. For an $m \times n$ matrix A , let $p(A)$ denote the *pivot* of A , which is defined to be the lexicographically first position in which A has a non-zero entry. Given a collection $\mathcal{A} = \{A_1, \dots, A_d\}$ of $m \times n$ matrices, let $G(\mathcal{A})$ be the bipartite graph whose two parts are $[m]$ and $[n]$, and whose edges are the pairs $p(A_1), \dots, p(A_d)$. In other words, we connect $i \in [m]$ to $j \in [n]$ by an edge if (i, j) is the pivot of one of the matrices in \mathcal{A} .

Recall that a *matching* in a graph is a collection of vertex-disjoint edges, and its *size* is the number of such edges. The *matching number* $\nu(G)$ of a graph G is the size of the largest matching it contains. Meshulam's key lemma, which will easily imply Theorem 2.1, is the following.

Lemma 2.2. *Let $\mathcal{A} = \{A_1, \dots, A_d\}$ be any collection of $m \times n$ matrices over any field \mathbb{F} . Then $\text{span}(A_1, \dots, A_d)$ contains a matrix of rank at least $\nu(G(\mathcal{A}))$.*

Proof sketch. Suppose first, for simplicity, that each A_k is an elementary matrix, i.e. consists of all zeroes except for a single one at position $p(A_k)$. Suppose without loss of generality that

a maximum matching in $G(\mathcal{A})$ is given by the edges $p(A_1), \dots, p(A_\nu)$. This means that in the matrix $B = A_1 + \dots + A_\nu$, we have ν ones and all other entries are zero, and these ν ones appear in ν distinct rows and ν distinct columns. Indeed, this is exactly what a matching in $G(\mathcal{A})$ is: pivots of the matrices A_k which appear in distinct rows and columns. But this shows that $\text{rank } B = \nu$, and clearly $B \in \text{span}(A_1, \dots, A_d)$, which proves the lemma.

The main remaining point is that, although the special case above was extremely special, the general case behaves essentially the same. Suppose again that $p(A_1), \dots, p(A_\nu)$ form a maximum matching in $G(\mathcal{A})$. It is no longer the case that $B = A_1 + \dots + A_\nu$ has a $\nu \times \nu$ submatrix with exactly one 1 in each row and column, but it is “almost” like that. Indeed, since we declared a pivot to be the lexicographically first non-zero entry of a matrix, we see that the structure of B is something like that of an upper-triangular matrix. By exploiting this structure, one can find a different matrix B' , still a linear combination of A_1, \dots, A_ν , whose rank is at least ν . \square

Proof of Theorem 2.1. Let $S \subseteq \mathbb{F}^{m \times n}$ be a matrix space, and suppose that every matrix in S has rank at most r . By performing Gaussian elimination, we can pick a basis A_1, \dots, A_d of S such that $p(A_1), \dots, p(A_d)$ are all distinct: this is a special case of the fact that every subspace of \mathbb{F}^N has a basis where each basis vector has a distinct first non-zero coordinate. Letting $\mathcal{A} = \{A_1, \dots, A_d\}$, we see that the graph $G(\mathcal{A})$ has $d = \dim(S)$ edges. Moreover, the matching number of $G(\mathcal{A})$ is at most r , by Lemma 2.2.

Now, we need to recall König’s theorem, a fundamental result in graph theory. It says that in any bipartite graph G , the matching number is equal to the *vertex cover number*, namely the minimum number of vertices of G such that every edge is incident to one of them. Recall that the vertices of $G(\mathcal{A})$ are $[m] \cup [n]$, namely correspond to rows and columns of the matrices. Because of this, we see that a vertex cover of $G(\mathcal{A})$ is a collection of rows and columns whose union contains $p(A_1), \dots, p(A_d)$.

Since every row contains at most n pivots, and every column contains at most m pivots, we see that we need at least $d/\max\{m, n\}$ rows and columns to cover $p(A_1), \dots, p(A_d)$. But by König’s theorem, we can find a collection of $\nu(G(\mathcal{A})) \leq r$ rows and columns that cover all these pivots. We conclude that

$$r \geq \nu(G(\mathcal{A})) \geq \frac{d}{\max\{m, n\}},$$

which is what we wanted to prove. \square

In addition to proving the bound on the dimension of a space all of whose elements have rank at most r , this proof also shows us that the structure we came up with—namely a matrix space consisting of r non-zero rows or columns—is really fundamental. Indeed, the proof shows that in any matrix space with elements of rank at most r , the pivots of an appropriately chosen basis must indeed lie in only r rows or columns. In fact, by using the same ideas and a little more work, one can completely classify the spaces of dimension $\max\{rm, rn\}$ with rank at most r : they can all be made into the form of having only r non-zero rows or columns by applying an appropriate change of basis on the left or the right.

However, such a classification is much harder to find when the dimension is substantially smaller than $\max\{rm, rn\}$. Indeed, as we discussed above, it seems difficult to determine if

an $n \times n$ matrix space is singular, i.e. if all its entries have rank at most $n - 1$. If we had a good classification of such spaces, then one could expect an efficient algorithm to follow. But there are a number of “exotic” singular spaces. A nice example is the space of $n \times n$ anti-symmetric matrices, where n is odd.

3 Nilpotent spaces

Let $S \subseteq \mathbb{F}^{n \times n}$ be a matrix space all of whose entries are nilpotent. What can we say about S ? This question was first studied by Gerstenhaber in the 1950s, motivated by some questions in algebra: he was interested in developing the theory of non-associative algebras, and questions about nilpotent matrix spaces arose naturally for much the same reason as the nilradical arises naturally in the study of commutative rings.

As above, the most natural question is how “big” can such a nilpotent space be. Namely, what is the maximum dimension of $S \subseteq \mathbb{F}^{n \times n}$ if every matrix in S is nilpotent?

Since every nilpotent matrix is singular, we automatically get an upper bound of $n^2 - n$ from Theorem 2.1. For the lower bound, a natural construction is the space of all strictly upper-triangular matrices, i.e. the space of all matrices with zeroes on and below the main diagonal. Every such matrix is certainly nilpotent, and this space has dimension $\binom{n}{2} = \frac{1}{2}(n^2 - n)$, since there are $\binom{n}{2}$ positions above the main diagonal. Thus, in this case, we have a factor of two between our upper and lower bounds. Which is closer to the truth?

As with the case of matrix spaces of bounded rank, it turns out that this natural lower bound is actually correct. This was proven for sufficiently large fields by Gerstenhaber, and for all fields by Serezhkin in the 1980s.

Theorem 3.1 (Gerstenhaber, Serezhkin). *Let $S \subseteq \mathbb{F}^{n \times n}$ be a matrix space. If every matrix in S is nilpotent, then $\dim S \leq \binom{n}{2}$.*

There are now a number of proofs known for this theorem, all beautiful and using a variety of different ideas. My personal favorite proof is due to de Seguins Pazzis, who proves Theorem 3.1 via a beautiful and surprising inductive argument. It is not at all obvious from his proof as written, but it in fact turns out that this proof is also graph-theoretic at heart. Indeed, much as Meshulam’s proof exploited a connection between the ranks of matrices and matchings in bipartite graphs, de Seguins Pazzis’s proof is fundamentally about the connection between nilpotent matrices and acyclic directed graphs. More generally, there turns out to be a rich theory connecting between graph-theoretic properties and matrix-theoretic ones, and many results about matrix spaces can be viewed as “quantum” versions of graph-theoretic results.

4 Nowhere singular matrix spaces

For the last topic, I want to go in the opposite direction from the Dieudonné–Flanders–Meshulam theorem. Namely, let us say that $S \subseteq \mathbb{F}^{n \times n}$ is *nowhere singular* if every non-zero matrix in S is invertible. This is basically the opposite extreme: rather than demanding that every matrix in S be singular, we are demanding that every matrix in S be non-singular, with the obvious exception of the zero matrix.

It is very easy to come up with examples of one-dimensional nowhere singular spaces. Namely, we can just pick any invertible matrix $A \in \mathbb{F}^{n \times n}$, and let S be the space spanned by A . But it is not at all obvious how to construct even two-dimensional nowhere singular spaces. In fact, in many instances, such spaces simply do not exist.

Proposition 4.1. *Let $S \subseteq \mathbb{F}^{n \times n}$ be nowhere singular. If \mathbb{F} is algebraically closed, then $\dim S \leq 1$.*

Proof. Suppose for contradiction that $\dim S \geq 2$, and let $A, B \in S$ be linearly independent. Consider the polynomial $p(\lambda) = \det(A + \lambda B)$, whose coefficients are in \mathbb{F} . Since \mathbb{F} is algebraically closed, we conclude that p has a root, namely there exists some $\lambda \in \mathbb{F}$ so that $p(\lambda) = 0$. But this implies that $A + \lambda B$ is a singular matrix, and this is not the zero matrix because A and B are linearly independent. This is a contradiction. \square

Because of this, we can't hope to find interesting nowhere singular spaces over algebraically closed fields, such as \mathbb{C} . What about over non-algebraically closed fields, like \mathbb{Q} or \mathbb{R} ? The same proof shows that even in these cases, we might need to temper our expectations.

Proposition 4.2. *Let $S \subseteq \mathbb{R}^{n \times n}$ be nowhere singular, and let n be odd. Then $\dim S \leq 1$.*

Proof. As above, suppose for contradiction that $A, B \in S$ are linearly independent, and let $p(\lambda) = \det(A + \lambda B)$. This is a polynomial with real coefficients whose degree, n , is odd. Since every odd-degree real polynomial has a root, we can find some $\lambda \in \mathbb{R}$ such that $p(\lambda) = 0$, which yields the same contradiction as above. \square

Nonetheless, when n is even, we *can* find nowhere singular spaces of real matrices with dimension greater than 1. The simplest example is the 2×2 space defined by the symbolic matrix

$$X = \begin{pmatrix} a & b \\ -b & a \end{pmatrix}.$$

We can directly compute that $\det(X) = a^2 + b^2$, which is never zero for non-zero real a, b . Thus, we have found a two-dimensional nowhere singular matrix space.

Going up to $n = 4$, we can find an even more interesting example, whose dimension is four:

$$X = \begin{pmatrix} a & b & c & d \\ -b & a & -d & c \\ -c & d & a & -b \\ -d & -c & b & a \end{pmatrix}$$

One can directly check that

$$\det(X) = (a^2 + b^2 + c^2 + d^2)^2,$$

which is never zero for non-zero real a, b, c, d . This yields a 4-dimensional nowhere singular real space for $n = 4$.

It is also not too hard to see that these two examples are best possible, and that a nowhere singular $n \times n$ matrix space has dimension at most n .

Proposition 4.3. *If $S \subseteq \mathbb{F}^{n \times n}$ is nowhere singular, then $\dim S \leq n$.*

Proof. Let A_1, \dots, A_d be a basis for S , and fix some non-zero vector $v \in \mathbb{F}^n$. Consider the d vectors A_1v, \dots, A_dv . If $d > n$, then these must be linearly independent, meaning that we can find $\lambda_1, \dots, \lambda_d \in \mathbb{F}$, not all zero, so that $\lambda_1 A_1 v + \dots + \lambda_d A_d v = 0$. But this implies that $\lambda_1 A_1 + \dots + \lambda_d A_d \in S$ has a non-trivial kernel, a contradiction. \square

Thus, we see that for odd n , the largest dimension of a nowhere singular space is 1, while for $n \in \{2, 4\}$, the largest such dimension is n . It is not clear what pattern generalizes this, and one can come up with several natural guesses. But I would certainly never guess the following pattern.

Definition 4.4. For a positive integer n , we write $n = (2a+1)2^{b+c}$ for non-negative integers a, b, c with $c < 4$; in other words, we consider the highest power of 2 dividing n , and then further write that exponent in terms of its mod-4 behavior. Then the *Hurwitz–Radon number* is defined as

$$\text{HR}(n) = 2^c + 8b.$$

Theorem 4.5 (Hurwitz, Radon, Adams). *The maximum dimension of an $n \times n$ nowhere singular space of real matrices is $\text{HR}(n)$.*

The lower bound was proved by Hurwitz and Radon in the 1920s, by explicitly constructing nowhere singular spaces in $\mathbb{R}^{n \times n}$ of dimension $\text{HR}(n)$. To understand where these constructions come from, we return to our examples above: the first one screams out “complex numbers” while the latter screams out “quaternions”. This is, of course, not a coincidence, and we make the following definition.

Definition 4.6. A family A_1, \dots, A_t of $n \times n$ real orthogonal matrices is called a *Hurwitz–Radon family* if $A_i^2 = -I$ for all i and $A_i A_j = -A_j A_i$ for all $i \neq j$.

Proposition 4.7. *Let $A_1, \dots, A_t \in \mathbb{R}^{n \times n}$ be a Hurwitz–Radon family. Then the space $S = \text{span}(A_1, \dots, A_t, I)$ is a nowhere singular matrix space of dimension $t + 1$.*

Proof. Note first that since each A_i is orthogonal and satisfies $A_i^2 = -I$, we conclude that $A_i^T = -A_i$ for all i . Suppose for contradiction that there exist $\lambda_1, \dots, \lambda_t, \lambda \in \mathbb{R}$ so that $M := \lambda_1 A_1 + \dots + \lambda_t A_t + \lambda I$ is singular. Note that

$$MM^T = (\lambda_1 A_1 + \dots + \lambda_t A_t + \lambda I)(-\lambda_1 A_1 - \dots - \lambda_t A_t + \lambda I) = (\lambda_1^2 + \dots + \lambda_t^2 + \lambda^2)I.$$

The left-hand side is singular, so we must have $\lambda_1^2 + \dots + \lambda_t^2 + \lambda^2 = 0$, which proves that S is nowhere singular. \square

The relations we require of a Hurwitz–Radon family are essentially the same as those we require of the generators of \mathbb{C} and \mathbb{H} . More generally, a *Clifford algebra* is an algebraic structure whose generators satisfy such relations, and \mathbb{C} and \mathbb{H} are the two smallest non-trivial Clifford algebras. In order to turn a Clifford algebra into a Hurwitz–Radon family, we simply need a map that embeds the Clifford algebra inside the matrix algebra $\mathbb{R}^{n \times n}$, namely an orthogonal representation of the Clifford algebra.

As it turns out, Clifford algebras and their representations are fairly easy to understand, and they can be fully classified. Using this classification, one can prove the following theorem.

Theorem 4.8 (Hurwitz, Radon). *The maximum size of a Hurwitz–Radon family in $\mathbb{R}^{n \times n}$ is $\text{HR}(n) - 1$.*

This immediately implies the lower bound in Theorem 4.5, thanks to Proposition 4.7. The upper bound is substantially harder, and was proved by Adams in the 1960s, in the process of studying a different problem which arises naturally in topology. Namely, Adams was studying families of linearly independent vector fields on spheres. Recall that a vector field V on S^{n-1} is an assignment of a vector $V_x \in \mathbb{R}^n$ to every point x of S^{n-1} , with the property that V_x is orthogonal to x , and so that V_x varies continuously in x . A collection of vector fields $V^{(1)}, \dots, V^{(t)}$ is called *linearly independent* if $V_x^{(1)}, \dots, V_x^{(t)}$ are linearly independent vectors in \mathbb{R}^n for all $x \in S^{n-1}$.

The hairy ball theorem states that if n is odd, then every vector field V on S^{n-1} must vanish at some point, i.e. $V_x = 0$ for some $x \in S^{n-1}$. In particular, if n is odd, then there does not exist any family of linearly independent vector fields on S^{n-1} . However, when n is even, then one can always find at least one non-vanishing vector field on S^{n-1} . As a natural generalization of the hairy ball theorem, one can ask for the maximum number of linearly independent vector fields on S^{n-1} . The connection between this problem and that of nowhere singular matrix spaces is given by the following.

Proposition 4.9. *Let $S \subseteq \mathbb{R}^{n \times n}$ be a nowhere singular space of dimension $t + 1$. Then there exist t linearly independent vector fields on S^{n-1} .*

Proof. Let a basis of S be given by $A_1, \dots, A_{t+1} \in \mathbb{R}^{n \times n}$. By multiplying every matrix in S by A_{t+1}^{-1} , we may assume that $A_{t+1} = I$. For every $x \in S^{n-1}$ and every $1 \leq i \leq t$, we define

$$V_x^{(i)} = A_i x - (x \cdot A_i x)x.$$

Note that $V_x^{(i)}$ is orthogonal to x by design, since we have simply performed the Gram-Schmidt orthogonalization process to ensure this. Additionally, the definition of $V_x^{(i)}$ is clearly continuous in x , so each $V^{(i)}$ is a vector field on S^{n-1} . To see that they are linearly independent, suppose for contradiction that there exist $x \in S^{n-1}$ and $\lambda_1, \dots, \lambda_t \in \mathbb{R}$ so that $\lambda_1 V_x^{(1)} + \dots + \lambda_t V_x^{(t)} = 0$. This implies that

$$\lambda_1 A_1 x + \dots + \lambda_t A_t x - \lambda x = 0,$$

where $\lambda = \sum_{i=1}^t \lambda_i (x \cdot A_i x)$. But this in turn means that $\lambda_1 A_1 + \dots + \lambda_t A_t - \lambda I \in S$ is a matrix with non-trivial kernel, contradicting the assumption that S is nowhere singular. So we conclude that $\lambda_1 = \dots = \lambda_t = 0$, as desired. \square

Thanks to this proposition, we see that we can find at least $\text{HR}(n) - 1$ linearly independent vector fields on S^{n-1} . Adams's theorem is that this is actually best possible.

Theorem 4.10 (Adams). *There do not exist $\text{HR}(n)$ linearly independent vector fields on S^{n-1} .*

In order to prove this theorem, Adams had to define the so-called *Adams operations* in topological K -theory. This allowed him to compute the topological K -theory of real projective space, which turns out to be more or less equivalent to determining the number of linearly independent vector fields on spheres. These operations are now central tools in K -theory.