A Polynomial Degree Bound on Equations for Non-rigid Matrices and Small Linear Circuits

BEN LEE VOLK, Efi Arazi School of Computer Science, Reichman University, Israel
MRINAL KUMAR, Department of Computer Science & Engineering, IIT Bombay, India

We show that there is an equation of degree at most poly(n) for the (Zariski closure of the) set of the non-rigid matrices: That is, we show that for every large enough field $\mathbb{F}$, there is a non-zero n^2-variate polynomial $P \in \mathbb{F}[x_{1,1}, \ldots, x_{n,n}]$ of degree at most poly(n) such that every matrix M that can be written as a sum of a matrix of rank at most $n/100$ and a matrix of sparsity at most $n^2/100$ satisfies $P(M) = 0$. This confirms a conjecture of Gesmundo, Hauenstein, Ikenmeyer, and Landsberg [9] and improves the best upper bound known for this problem down from $\exp(n^2)$ [9, 12] to poly(n).

We also show a similar polynomial degree bound for the (Zariski closure of the) set of all matrices M such that the linear transformation represented by M can be computed by an algebraic circuit with at most $n^2/200$ edges (without any restriction on the depth). As far as we are aware, no such bound was known prior to this work when the depth of the circuits is unbounded.

Our methods are elementary and short and rely on a polynomial map of Shpilka and Volkovich [21] to construct low-degree "universal" maps for non-rigid matrices and small linear circuits. Combining this construction with a simple dimension counting argument to show that any such polynomial map has a low-degree annihilating polynomial completes the proof.

As a corollary, we show that any derandomization of the polynomial identity testing problem will imply new circuit lower bounds. A similar (but incomparable) theorem was proved by Kabanets and Impagliazzo [11].

CCS Concepts: • **Theory of computation → Algebraic complexity theory**; **Circuit complexity;**

Additional Key Words and Phrases: Algebraic complexity theory, rigid matrices, linear circuits, tensor rank, algebraic geometry

ACM Reference format:
Ben Lee Volk and Mrinal Kumar. 2022. A Polynomial Degree Bound on Equations for Non-rigid Matrices and Small Linear Circuits. *ACM Trans. Comput. Theory* 14, 2, Article 6 (September 2022), 14 pages.
https://doi.org/10.1145/3543685

Ben Lee Volk parts of this work was done while at Center for the Mathematics of Information, California Institute of Technology, and at the Department of Computer Science, University of Texas at Austin, supported by NSF Grant CCF-1705028.

Authors' addresses: B. L. Volk, Efi Arazi School of Computer Science, Reichman University, 8 Ha'Universita st. Herzliya 4610101, Israel; email: benleevolk@gmail.com; M. Kumar, Department of Computer Science and Engineering Indian Institute of Technology Bombay, Powai, Mumbai 400076, India; email: mrinal@cse.iitb.ac.in.

1942-3454/2022/09-ART6 $15.00
https://doi.org/10.1145/3543685

1 INTRODUCTION

1.1 Equations for Varieties in Algebraic Complexity Theory

A set $V \subseteq \mathbb{F}^n$ is called an *affine variety* if it is the common zero set of a set of polynomials in $\mathbb{F}[x_1, \ldots, x_n]$. Let $V \subseteq \mathbb{F}^n$ be an affine variety and let $\mathbf{I}(V)$ denote its ideal, i.e., the set of polynomials vanishing on V.[1] A non-zero polynomial $P \in \mathbf{I}(V)$ is called an *equation* for V. An equation for V may serve as a "proof" that a point $\mathbf{x} \in \mathbb{F}^n$ is *not* in V, by showing that $P(\mathbf{x}) \neq 0$.

A fundamental idea that dates back to Strassen [23] is that many important circuit lower bounds problems in algebraic complexity theory fit naturally into the setting of showing that a point $\mathbf{x}$ lies outside a variety V. It was vastly extended in the Geometric Complexity Theory program [17], and was recently central to the algebraic natural proofs paradigm [5, 7, 8, 10]. In this formulation, one considers V to be the closure of a class of polynomials of low complexity, and $\mathbf{x}$ is the coefficient vector of the candidate hard polynomial.

Let $\Delta(V) := \min_{0 \neq P \in \mathbf{I}(V)}\{\deg(P)\}$. The quantity $\Delta(V)$ can be thought of as a measure of complexity for the geometry of the variety V. The quantity $\Delta(V)$ is a very coarse complexity measure. The recent line of work regarding algebraic natural proofs [8, 10] suggests to study the arithmetic circuit complexity of equations for varieties V that correspond to polynomials with small circuit complexity. Having $\Delta(V)$ growing like a polynomial in n is a necessary (but not a sufficient) condition for a variety V to have an algebraic natural proof for non-containment if one insists that the "proof" (that is, the equation P) belongs to the class VP.

The usefulness of equations was also noticed by Raz [18] in his *elusive functions* approach for proving circuit lower bounds. Briefly, one component of Raz's method is the observation that it is possible to consider not only a single point $\mathbf{x}$ describing the coefficient vector of a single polynomial, but also a larger set of polynomials that is the image of an explicit polynomial map $f(\mathbf{y})$ in a small number of variables. The image of f is thought of as a family of polynomials indexed by the variables $\mathbf{y}$, or as a polynomial in the original set of variables and additional auxiliary variables $\mathbf{y}$ (but, since their number is so small, this has little effect on the total circuit complexity). If V is a variety and P is an equation for P, then if $P(f(\mathbf{y}))$ is a non-zero polynomial in $\mathbf{y}$, the map f is called "elusive" and its image contains a point that is not in V, which implies a circuit lower bound for the explicit polynomial (in the original set of variables and additional auxiliary variables $\mathbf{y}$) described by f.

1.2 Rigid Matrices

A matrix M is (r, s)-rigid if M cannot be written as a sum $R+S$ where $\mathrm{rank}(R) \leq r$ and S contains at most s non-zero entries. Valiant [24] proved that if A is $(\varepsilon n, n^{1+\delta})$-rigid for some constants $\varepsilon, \delta > 0$, then A cannot be computed by arithmetic circuits of size $O(n)$ and depth $O(\log n)$, and posed the problem of *explicitly* constructing rigid matrices with these parameters, which is still open. It is easy to prove that most matrices have much stronger rigidity parameters: Over algebraically closed fields a generic matrix is $(r, (n-r)^2)$-rigid for any target rank r.

Let $\mathbb{F}$ be an algebraically closed field. Let $A_{r,s} \subseteq \mathbb{F}^{n \times n}$ denote the set of matrices that are not (r, s)-rigid. Let $V_{r,s} = \overline{A_{r,s}}$ denote the Zariski closure of $A_{r,s}$. A geometric study of $V_{r,s}$ was initiated by Kumar, Lokam, Patankar, and Sarma [12]. Among other results, they prove that for every $s < (n-r)^2$, $\Delta(V_{r,s}) \leq n^{4n^2}$. A slightly improved (but still exponential) upper bound was obtained by Gesmundo, Hauenstein, Ikenmeyer, and Landsberg [9], who also conjectured that for some $\varepsilon, \delta > 0$, $\Delta(V_{\varepsilon n, n^{1+\delta}})$ grows like a polynomial function in n. The following theorem, which we prove in this article, confirms this conjecture:

[1] For completeness, we provide the formal (standard) definitions for these notions in Section 2.1.

Theorem 1.1. *Let $\varepsilon < 1/25$, and let $\mathbb{F}$ be a field of size at least n^2. For every large enough n, there exists a non-zero polynomial $Q \in \mathbb{F}[x_{1,1}, \ldots, x_{n,n}]$, of degree at most n^3, which is a non-trivial equation for matrices that are* not *$(\varepsilon n, \varepsilon n^2)$-rigid. That is, for every such matrix M, $Q(M) = 0$.*

In fact, the conjecture of Reference [9] was slightly weaker: They conjectured that $\Delta(U)$ is polynomial in n for every irreducible component U of $V_{\varepsilon n, n^{1+\delta}}$. As shown by Reference [12], the irreducible components are in one-to-one correspondence with subsets of $[n] \times [n]$ of size $n^{1+\delta}$ corresponding to possible supports of the sparse matrix S.

As we observe in Remark 3.3, it is somewhat simpler to show that each of these irreducible components has an equation with a polynomial degree bound. However, since the number of such irreducible components is exponentially large, it is not clear if there is a single equation for the whole variety that is of polynomially bounded degree. We do manage to reverse the order of quantifiers and prove such an upper bound in Theorem 1.1. This suggests that the set of non-rigid matrices is much less complex than what one may suspect given the results of References [9, 12].

1.3 Circuits for Linear Transformations

The original motivation for defining rigidity was in the context of proving lower bounds for algebraic circuits [24]. If $A \in \mathbb{F}^{n \times n}$ is an $(\varepsilon n, n^{1+\delta})$-rigid matrix, for any $\varepsilon, \delta > 0$, then the linear transformation represented by A cannot be computed by an algebraic circuit of depth $O(\log n)$ and size $O(n)$.

Every algebraic circuit computing a linear transformation is without loss of generality a *linear* circuit. A linear circuit is a directed acyclic graph that has n inputs labeled $X_1, \ldots, X_n$ and n output nodes. Each edge is labeled by a scalar $\alpha \in \mathbb{F}$. Each node computes a linear function in $X_1, \ldots, X_n$ defined inductively. An internal node u with children, $v_1, \ldots, v_k$, connected to it by edges labeled $\alpha_1, \ldots, \alpha_k$, computes the linear function $\ell_v = \sum_i \alpha_i \ell_{v_i}$, where ℓ_{v_i} is the linear function computed by v_i, $1 \leq i \leq k$. The size of the circuit is the number of edges in the circuit.

It is possible to use similar techniques to those used in the proof of Theorem 1.1 to prove a polynomial upper bound on an equation for a variety containing all matrices $A \in \mathbb{F}^{n \times n}$ whose corresponding linear transformation can be computed by an algebraic circuit of size at most $n^2/200$ (even without restriction on the depth). Note that this is nearly optimal as any such linear transformation can be computed by a circuit of size n^2. More formally, we show the following:

Theorem 1.2. *Let $\mathbb{F}$ be a field of size at least n^2. For every large enough n, there exists a non-zero polynomial $Q \in \mathbb{F}[x_{1,1}, \ldots, x_{n,n}]$, of degree at most n^3, which is a non-trivial equation for matrices that are computed by algebraic circuit of size at most $n^2/200$.*

Let PIT denote the set of strings that describe arithmetic circuits (say, over $\mathbb{C}$) that compute the zero polynomial. It is well known that PIT $\in$ coRP. Kabanets and Impagliazzo [11] proved that certain circuit lower bounds follow from the assumption that PIT $\in$ P. As a corollary to Theorem 1.2, we are able to prove theorem of a similar kind.

Corollary 1.3. *Suppose* PIT $\in$ P. *Then at least one of the following is true:*

(1) *There exists a family of n-variate polynomials of degree* poly(n) *over $\mathbb{C}$, which can be computed (as its list of coefficients, given the input 1^n) in* PSPACE, *which does not have polynomial size constant free arithmetic circuits.*
(2) *there exists a family of matrices, constructible in polynomial time with an* NP *oracle (given the input 1^n), which requires linear circuits of size $\Omega(n^2)$.*

A *constant free arithmetic circuit* is an arithmetic circuit that is only allowed to use the constants $\{0, \pm 1\}$.

A different way to interpret Corollary 1.3 is by saying that at least one of the following three lower bound results hold: Either PIT $\notin$ P, or (at least) one of the two circuit lower bounds stated in the corollary. We emphasize that the result holds under *any* derandomization of PIT, that is, even under *white box* derandomization of PIT, which is a weaker assumption then black-box derandomization of PIT.

Our statement is similar to, but incomparable with the result of Kabanets and Impagliazzo [11], who proved that if PIT $\in$ P, then either the permanent does not have polynomial size constant free arithmetic circuits, or NEXP $\not\subseteq$ P/poly.

Since $(\varepsilon n, \varepsilon n^2)$-rigid matrices have linear circuit of size $3\varepsilon n^2$, the last item of Corollary 1.3 in particular implies a conditional construction of $(\Omega(n), \Omega(n^2))$-rigid matrices (it is also possible to directly use Theorem 1.1 instead of Theorem 1.2 to deduce this result). Unconditional constructions of rigid matrices in polynomial time with an NP oracle were recently given in References [2, 3]. However, the rigidity parameters in these papers are not strong enough to imply circuit lower bounds (furthermore, even optimal rigidity parameters do not imply $\Omega(n^2)$ lower bounds for general linear circuits).

Since it is widely believed that PIT $\in$ P, the answer to which of the last two items of Corollary 1.3 holds boils down to the question of whether there exists an equation for non-rigid matrices of degree $\text{poly}(n)$ and circuit size $\text{poly}(n)$. If determining if a matrix is rigid is coNP-hard (as is known for some restricted ranges of parameters [15]), then it is tempting to also believe that the equations should not be easily computable, as they provide "proof" for rigidity that can be verified in randomized polynomial time. However, it could still be the case that those equations that have polynomial size circuits only prove the rigidity of "easy" instances; and it could also be the case that the *number* of equations defining the set of non-rigid matrices is very large, so even if each one of them was easily computable, one still would not get an efficient algorithm for deciding rigidity.

As another application of our our techniques used in the proofs of Theorems 1.1 and 1.2, we can also prove the following upper bound on the degree of equations for low rank tensors:

Theorem 1.4. *For every field $\mathbb{F}$ and for all $n, d \in \mathbb{N}$, there exists a non-zero polynomial Q on n^d variables and degree at most n^{2d}, which is a non-trivial equation for d-dimensional tensors $\tau : [n]^{\otimes d} \to \mathbb{F}$ of rank at most $n^{d-1}/100d$.*

Raz [19] proved that when d is a super-constant but slowly growing function of n (e.g., $d = O(\log n/\log\log n)$), strong enough lower bounds for tensor rank would imply super-polynomial lower bounds for general arithmetic formulas. While Theorem 1.4 holds for such super-constant values of d (and a rank lower bound of $n^{d-1}/100d$ would suffice for Raz's approach), the equations we get for such values of d are of slightly super-polynomial degree.

1.4 Proof Techniques

Our proof of Theorems 1.1 and 1.2 are short, elementary, and based on a two-step argument.

In the first step, we show that there is a polynomial map of *low degree* and on a *small number of variables* such that any *non-rigid* matrix (respectively, a matrix with a small algebraic circuit) lies in its image. The parameters of the polynomial map, e.g., its degree and arity depend on the upper bound on the rigidity of the matrices we are working with (and the size of the algebraic circuit for Theorem 1.2). Once we have such a map, an immediate observation is that any non-trivial annihilating polynomial of this polynomial map gives us an equation satisfied by everything in the image of the maps. To complete our proof, we use the bounds on the degree and arity of these maps to show that they do indeed have low-degree annihilating polynomials. This second step is a very simple linear algebraic argument based on dimension counting. For our construction of the

polynomial map, we rely on an application of a polynomial map due to Shpilka and Volkovich [21]. Shpilka and Volkovich constructed this map in the context of designing deterministic algorithms for polynomial identity testing for various subclasses of algebraic circuits.

This technique shares some similarities with Raz's elusive function approach for circuit lower bounds [18]. In both cases, the main underlying observation is that the set of low complexity objects is contained in an image of a low-degree polynomial map in a small number of variables.

One downside of the fact that our proofs are based on dimension counting arguments is that they are non-constructive and do not give explicit equations for the relevant varieties. It thus remains a very interesting open problem to provide explicit low-degree equations for any of the varieties considered in this article. Here, "explicit" means a polynomial that has arithmetic circuits of size $\text{poly}(n)$, although one can be even more permissive and ask for polynomials in the class VNP.

The question of whether such equations exists has a win-win flavor: If they do, then this can aid in explicit constructions of rigid matrices; if, however, all equations are hard, then we have identified a family of polynomials that requires super-polynomial arithmetic circuits. Assuming the existence of a polynomial time algorithm for polynomial identity testing, we are able to make this connection formal to prove Corollary 1.3.

2 PRELIMINARIES

2.1 Some Basic Notions in Algebraic Geometry

For completeness, in this section, we define some basic notions in algebraic geometry. A reader who is familiar with this topic may skip to the next section.

Let $\mathbb{F}$ be an algebraically closed field. A set $V \subseteq \mathbb{F}^n$ is called an *affine variety* if there exist polynomials $f_1, \ldots, f_t \in \mathbb{F}[x_1, \ldots, x_n]$ such that $V = \{\mathbf{x} : f_1(\mathbf{x}) = f_2(\mathbf{x}) = \cdots = f_t(\mathbf{x}) = 0\}$. For convenience, in this article, we often refer to affine varieties simply as varieties.

For each variety V there is a corresponding ideal $\mathbf{I}(V) \subseteq \mathbb{F}[x_1, \ldots, x_n]$ that is defined as

$$\mathbf{I}(V) := \{f \in \mathbb{F}[x_1, \ldots, x_n] : f(\mathbf{x}) = 0 \text{ for all } \mathbf{x} \in V\}.$$

Conversely, for an ideal $I \subseteq \mathbb{F}[x_1, \ldots, x_n]$, we may define the variety

$$\mathbf{V}(I) = \{\mathbf{x} : f(\mathbf{x}) = 0 \text{ for all } f \in I\}.$$

Given a set $A \subseteq \mathbb{F}^n$, we may similarly define the ideal $\mathbf{I}(A)$. The (Zariski) *closure* of a set A, denoted $\overline{A}$, is the set $\mathbf{V}(\mathbf{I}(A))$. In words, the closure of A is the set of common zeros of all the polynomials that vanish on A. It is also the smallest variety with respect to inclusion that contains A. By construction, $\overline{A}$ is a variety, and a polynomial that vanishes everywhere on A is also vanishes on $\overline{A}$.

Over $\mathbb{C}$, it is instructive to think of the Zariski closure of A as the usual Euclidean closure. In fact, for the various sets A we consider in this article (which correspond to sets of "low complexity" objects, e.g., non-rigid matrices or matrices that can be computed with a small circuit), it can be shown that these two notions of closure coincide (see, e.g., Section 4.2 of Reference [4]).

A variety V is called *irreducible* if it cannot be written as a union $V = V_1 \cup V_2$ of varieties V_1, V_2 that are properly contained in V. Every variety can be uniquely written as a union $V = V_1 \cup V_2 \cup \cdots \cup V_m$ of irreducible varieties. The varieties $V_1, \ldots, V_m$ are then called the *irreducible components* of V.

2.2 A Low-degree Equation for Images of Polynomial Maps

A key ingredient in our proofs is the following elementary lemma, which shows that images of low-degree polynomial maps in a small number of variables have a low-degree annihilator:

Lemma 2.1. *Let $\mathbb{F}$ be a field and let $P : \mathbb{F}^K \to \mathbb{F}^N$ be a polynomial map of degree at most D. Suppose Δ is such that*

$$\binom{N+\Delta}{N} > \binom{K+D\Delta}{K}.$$

Then there's a non-zero polynomial $Q \in \mathbb{F}[y_1, \ldots, y_N]$ of degree at most Δ such that for any $\boldsymbol{\alpha}$ in the image of P, $Q(\boldsymbol{\alpha}) = 0$.

Proof. Let V_1 denote the subspace of polynomials over $\mathbb{F}$ in N variables of degree at most Δ. Let V_2 denote the subspace of polynomials over $\mathbb{F}$ in K variables of degree at most $D\Delta$. Consider the linear transformation $T : V_1 \to V_2$ given by $Q \mapsto Q \circ P$, where $Q \circ P$ denotes the composition of the polynomial Q with the map P, i.e., $(Q \circ P)(\mathbf{x}) = Q(P(\mathbf{x}))$(indeed, observe that, since $\deg(Q) \leq \Delta$ and $\deg(P) \leq D$, it follows that $\deg Q \circ P \leq D\Delta$).

We have that $\dim(V_1) = \binom{N+\Delta}{N}$, whereas $\dim(V_2) = \binom{K+D\Delta}{K} < \dim(V_1)$ by assumption. This implies that T has a non-trivial kernel, that is, there exists $0 \neq Q_0 \in V_1$ such that $Q_0 \circ P \equiv 0$.

Suppose $\boldsymbol{\alpha}$ is in the image of P. That this, there exists $\boldsymbol{\beta} \in \mathbb{F}^K$ such that $P(\boldsymbol{\beta}) = \alpha$. Then

$$Q_0(\boldsymbol{\alpha}) = Q_0(P(\boldsymbol{\beta})) = Q_0 \circ P(\boldsymbol{\beta}) = 0,$$

as $Q_0 \circ P \equiv 0$. □

3 DEGREE UPPER BOUND FOR NON-RIGID MATRICES

In this section, we prove Theorem 1.1. A key component of the proof is the use of the following construction, due to Shpilka and Volkovich, which provides an explicit low-degree polynomial map on a small number of variables, which contains all sparse matrices in its image. For completeness, we provide the construction and prove its basic property.

Lemma 3.1 ([21]). *Let $\mathbb{F}$ be a field such that $|\mathbb{F}| > n$. Then for all $k \in \mathbb{N}$, there exists an explicit polynomial map $\mathrm{SV}_{n,k}(\mathbf{x}, \mathbf{y}) : \mathbb{F}^{2k} \to \mathbb{F}^n$ of degree at most n such that for any subset $T = \{i_1, \ldots, i_k\} \subseteq [n]$ of size k, there exists a setting $\mathbf{y} = \boldsymbol{\alpha}$ such that $\mathrm{SV}(\mathbf{x}, \boldsymbol{\alpha})$ is identically zero on every coordinate $j \notin T$, and equals x_j in coordinate i_j for all $j \in [k]$.*

Proof. Arbitrarily pick distinct $\alpha_1, \ldots \alpha_n \in \mathbb{F}$, and let $u_1, \ldots, u_n$ be their corresponding Lagrange's interpolation polynomials, i.e., polynomials of degree at most $n-1$ such that $u_i(\alpha_j) = 1$ if $j = i$ and 0 otherwise (more explicitly, $u_i(z) = \frac{\prod_{j \neq i}(z-\alpha_j)}{\prod_{j \neq i}(\alpha_i - \alpha_j)}$).

Let $P_i(x_1, \ldots, x_k, y_1, \ldots, y_k) = \sum_{j=1}^{k} u_i(y_j) \cdot x_j$, and finally let

$$\mathrm{SV}_{n,k}(\mathbf{x}, \mathbf{y}) = (P_1(\mathbf{x}, \mathbf{y}), \ldots, P_n(\mathbf{x}, \mathbf{y})).$$

It readily follows that given $T = \{i_1, \ldots, i_k\}$ as in the statement of the lemma, we can set $y_j = \alpha_{i_j}$ for $j \in [k]$ to derive the desired conclusion. The upper bound on the degree follows by inspection. □

As a step toward the proof of Theorem 1.1, we show there is a polynomial map on much fewer than n^2 variables with degree polynomially bounded in n such that its image contains every non-rigid matrix. In the next step, we show that the image of every such polynomial map has an equation of degree $\mathrm{poly}(n)$.

Lemma 3.2. *For every $r \leq n$ and $s \leq n^2$, there exists an explicit polynomial map $P : \mathbb{F}^{2rn+2s} \to \mathbb{F}^{n \times n}$, of degree at most n^2, such that every matrix M that is not (r, s) rigid lies in its image.*

Proof. Let $\mathbf{u}, \mathbf{v}$ be disjoint sets of rn variables each, and $\mathbf{x}, \mathbf{y}$ be disjoint sets of s variables each.

Let U be a symbolic $n \times r$ matrix whose entries are labeled by the variables $\mathbf{u}$, and similarly let V be a symbolic $r \times n$ matrix labeled by $\mathbf{v}$. Let $\mathrm{UV}(\mathbf{u}, \mathbf{v}) : \mathbb{F}^{2rn} \to \mathbb{F}^{n \times n}$ be the degree 2 polynomial map defined by the matrix multiplication UV.

Finally, let $P : \mathbb{F}^{2rn+2s} \to \mathbb{F}^{n\times n}$ be defined as

$$P(\mathbf{u},\mathbf{v},\mathbf{x},\mathbf{y}) = \mathrm{UV}(\mathbf{u},\mathbf{v}) + \mathrm{SV}_{n^2,s}(\mathbf{x},\mathbf{y}),$$

where $\mathrm{SV}_{n^2,s}$ is as defined in Lemma 3.1.

Suppose now M is a not an (r,s) rigid matrix, i.e., $M = R + S$ for R of rank at most r and S that is s-sparse. Decompose $R = U_0V_0$ for $n \times r$ matrix U_0 and $r \times n$ matrix V_0. Let T denote the support of S. For convenience, we may assume $|T| = s$ (otherwise, pad with zeros arbitrarily). Let $\boldsymbol{\alpha} \in \mathbb{F}^s$ denote the setting for $\mathbf{y}$ in $\mathrm{SV}_{n^2,s}$ that maps $x_1,\ldots,x_s$ to T, and let $\boldsymbol{\beta} = (\beta_1,\ldots,\beta_s)$ denote the non-zero entries of S. Then

$$P(U_0,V_0,\boldsymbol{\beta},\boldsymbol{\alpha}) = U_0V_0 + S = R + S = M. \qquad \square$$

To complete the proof of Theorem 1.1, we need to show that for the relevant range of parameters it is possible to combine Lemma 3.2 with Lemma 2.1.

Proof of Theorem 1.1. We apply Lemma 2.1 with the map P from Lemma 3.2 for $r = \varepsilon n$ and $s = \varepsilon n^2$ so $2rn + 2s = 4\varepsilon n^2$.

It remains to be shown that the assumption of Lemma 2.1 indeed holds if we pick $\Delta = n^3$. Indeed, note that $N = n^2$, $D = n^2$ and $K = 4\varepsilon n^2$, so $\binom{n^2+n^3}{n^2} \geq n^{n^2}$ whereas $\binom{4\varepsilon n^2+n^5}{4\varepsilon n^2} \leq (2n^5)^{4\varepsilon n^2}$, and thus by the choice of ε, this is smaller than n^{n^2} for every large enough n.

As the map P from Lemma 3.2 contains in its image all matrices that are not $(\varepsilon n, \varepsilon n^2)$-rigid, the statement follows: $\square$

Remark 3.3. If the support of the sparse matrix is fixed *a priori* to some set $S \subseteq [n] \times [n]$ of cardinality at most ϵn^2, then it is easier to come up with a universal map $\tilde{P} : \mathbb{F}^{3\epsilon n^2} \mapsto \mathbb{F}^{n\times n}$ such that every matrix M whose rank can be reduced to at most ϵn by changing entries in the set S is contained in the image of $\tilde{P}$. Just consider $\tilde{P}(\mathbf{w},\mathbf{x},\mathbf{y}) = \mathrm{UV}(\mathbf{u},\mathbf{v}) + W$, where W is a matrix such that for all $(i,j) \in [n]\times[n]$, if $(i,j) \in S$, then $W(i,j) = w_{i,j}$ and $W(i,j)$ is zero otherwise. Here, each $w_{i,j}$ is a distinct formal variable. Combined with the dimension comparison argument of Lemma 2.1, it can be seen that there is a non-zero low-degree polynomial $\tilde{Q}$ such that $\tilde{Q} \circ \tilde{P} \equiv 0$. This argument provides a (different) equation of polynomial degree for each irreducible component of the variety of non-rigid matrices. Note that the degree of $\tilde{P}$ is 2 (rather than n^2 as in the construction of P in Lemma 3.2) but other methods are currently unable to leverage that to obtain a meaningful improvement in the parameters.

Remark 3.4. It is possible to use the equation given in Theorem 1.1, and using the methods of Reference [12], to construct "semi-explicit" $(\varepsilon n, \varepsilon n^2)$-rigid matrices. These are matrices whose entries are algebraic numbers (over $\mathbb{Q}$) with short description, which are non-explicit from the computational complexity point of view. However, such constructions are also known using different methods (see Section 2.4 of Reference [14]).

4 DEGREE UPPER BOUND FOR MATRICES WITH A SMALL CIRCUIT

In this section, we prove Theorem 1.2. Our strategy, as before, is to observe that all matrices with a small circuit lie in the image of a polynomial map P on a small number of variables and small degree. Circuits of size s can have many different topologies and thus, we first construct a "universal" linear circuit, of size $s' \leq s^4$, that contains as subcircuits all linear circuits of size s. Importantly, s' will affect the degree of P but not its number of variables. We note that constructions of such universal maps whose image contain all polynomials with small circuits have appeared before in the literature. In the context of circuits computing univariate polynomials, similar maps were constructed by Strassen [23] and Lipton [13], who used those maps to prove the existence of certain

such polynomials that are hard to compute. Our approach is much more reminiscent of that of Raz [18], who constructed universal maps for circuits computing multivariate polynomials. Raz's approach first constructs a universal circuit graph, which is a graph of size poly(s) that contains a subgraph every circuit of size s. The universal circuit graph is then used to construct the universal map by associating a variable with every edge of the universal graph. This poly(s) size blow up, which translates to a polynomial map with poly(s) variables, is fine when one considers multivariate polynomials of large degree and wishes to find polynomials with exponential circuit complexity. In our context, however, this increase in the number of variables is unacceptable, since every $n \times n$ matrix defines a linear transformation that can be computed by circuits of size at most $O(n^2)$, and we wish to separate between those who can be computed by linear size circuits and those who require super-linear size. Thus, a naive use of ideas in Reference [18] is insufficient, and we compose the universal circuit construction with the Shpilka-Volkovich map to get around this difficulty.

4.1 A Construction of Universal Map for Small Linear Circuits

We now define a map $U(\mathbf{x}, \mathbf{y})$ that is "universal" for size s linear circuits, i.e., it contains in its image all $n \times n$ matrices A whose corresponding linear transformation can be computed by a linear circuit of size at most s.

Let $s \geq n$. We first define a universal graph G for size s. G has a set V_0 of n input nodes labeled $X_1, \ldots X_n$ and a set V_{s+1} of n designated output nodes. In addition, G is composed of s disjoint sets of vertices $V_1, \ldots, V_s$, each contains s vertices.

Each vertex $v \in V_i$, for $0 \leq i \leq s+1$, has as its children all vertices $u \in V_j$ for all $0 \leq j < i$. It is clear than every directed acyclic graph with s edges (and hence at most s vertices, and depth at most s) can be (perhaps non-uniquely) embedded in G as a subgraph.

We now describe the edge labeling. Let $s' \leq s^4$ be the number of edges in G and let e_i denote the ith edge, $1 \leq i \leq s'$. The edge e_i is labeled by the ith coordinate of the map $\mathrm{SV}_{s',s}(\mathbf{x}, \mathbf{y})$ given in Lemma 3.1.

Thus, the graph G with this labeling computes a linear transformation (over the field $\mathbb{F}(\mathbf{x}, \mathbf{y})$) in the variables $X_1, \ldots, X_n$. More explicitly, the (i, j)th entry of the matrix $U(\mathbf{x}, \mathbf{y})$ representing this linear transformation is given by the sum, over all paths from X_i to the jth output node, of the product of the edge labels on that path. This entry is a polynomial in $\mathbf{x}, \mathbf{y}$, so we can think of U as a polynomial map from $\mathbb{F}^{2s}$ to $\mathbb{F}^{n^2}$.

Lemma 4.1. *The map $U(\mathbf{x}, \mathbf{y}) : \mathbb{F}^{2s} \to \mathbb{F}^{n^2}$ defined above contains in its image all $n \times n$ matrices A whose corresponding linear transformation can be computed by a linear circuit of size at most s. The degree of U is at most $s' \cdot (s+1)$.*

Proof. Let A be a matrix whose linear transformation is computed by a size s circuit C. The graph of C can be embedded as a subgraph in the graph G constructed above (if the embedding is not unique, pick one arbitrarily). Let $e_{i_1}, \ldots, e_{i_s}$ be the edges of this subgraph, and let $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_s)$ be their corresponding labels in C. By the properties of the map $\mathrm{SV}_{s',s}(\mathbf{x}, \mathbf{y})$ given in Lemma 3.1, it is possible to set the tuple of variables $\mathbf{y}$ to field elements $\alpha_1, \ldots, \alpha_s$ such that the jth coordinate of $\mathrm{SV}(\boldsymbol{\beta}, \boldsymbol{\alpha})$ equals β_i if $j = i_k$ for some $1 \leq k \leq s$ the 0 otherwise. Observe that under this labeling of the edges, the circuit G computes the same transformation as the circuit C. Hence, $U(\boldsymbol{\beta}, \boldsymbol{\alpha}) = A$.

To upper bound the degree of U, note that each edge label in G is a polynomial of degree s', and each path is of length at most $s+1$. □

4.2 Low-degree Equations for Small Linear Circuits

In a similar manner to the proof of Theorem 1.1, we now combine Lemma 2.1 with the map $U(\mathbf{x}, \mathbf{y})$ to show that its image has a equation of degree at most n^3. This would complete the proof of Theorem 1.2.

Proof of Theorem 1.2. Let $U : \mathbb{F}^{2s} \to \mathbb{F}^{n^2}$ be the map given by Lemma 4.1 for $s = n^2/200$ so $s' \leq n^8$, and the degree of U is at most $s'(s+1) \leq n^{10}$.

We apply Lemma 2.1 with the map $U(\mathbf{x}, \mathbf{y})$ so now $N = n^2$, $D \leq n^{10}$ and $K = 2s$. Setting $\Delta = n^3$, we once again get that $\binom{n^3+n^2}{n^2} \geq n^{n^2}$, whereas $\binom{n^2/100+n^{13}}{n^2/100} \leq (2n^{13})^{n^2/100} < n^{n^2}$ for every large enough n, so we can get a non-zero polynomial Q of degree at most n^3 such that $Q(A) = 0$ for every matrix in the image of U. By Lemma 4.1, if A has a circuit of size $n^2/200$, then it is in the image of U. □

5 DEGREE UPPER BOUND FOR THREE-DIMENSIONAL TENSORS

In this section, we prove Theorem 1.4. We start by quickly recalling the definitions of a tensors and that of tensor rank.

Definition 5.1. Let $d, n_1, n_2, \ldots, n_d$ be natural numbers and let $\mathbb{F}$ be a field. Then, a tensor of dimension d and size $(n_1, n_2, \ldots, n_d)$ over $\mathbb{F}$ is a function $\tau : [n_1] \times [n_2] \times \cdots \times [n_d] \to \mathbb{F}$.

For tensors considered in this article, $n_1, \ldots, n_d$ are all equal and denoted by the parameter n. Thus, such tensors simply functions from $[n]^d$ to $\mathbb{F}$. Tensors are alternatively defined in terms of multilinear maps and in algebraic complexity, often as set-multilinear polynomials. For the discussion in this article, we work with Definition 5.1 and refer to the survey by Saptharishi [20] for further discussions on these alternative definitions and their equivalence. We now define the notion of a rank one tensor and use it to define tensor rank, which is the main property of interest in this section.

Definition 5.2. Let $d, n \in \mathbb{N}$, $\mathbb{F}$ be a field and $\tau : [n]^d \to \mathbb{F}$ be a tensor. Then, τ is said to be of *rank* one if there exist vectors $\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_d \in \mathbb{F}^n$ such that for every $(i_1, i_2, \ldots, i_d) \in [n]^d$,

$$\tau(i_1, i_2, \ldots, i_d) = (\mathbf{u}_1)_{i_1} \cdot (\mathbf{u}_2)_{i_2} \cdots (\mathbf{u}_d)_{i_d},$$

where, $(\mathbf{u}_j)_{i_j}$ denotes the i_jth coordinate of $\mathbf{u}_i \in \mathbb{F}^n$.

In other words, τ is a rank one tensor if there exist vectors $\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_d \in \mathbb{F}^n$ such that τ is an outer product of $\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_d$, i.e.,

$$\tau = \mathbf{u}_1 \otimes \mathbf{u}_2 \otimes \cdots \otimes \mathbf{u}_d .$$

We are now ready to define the rank of a tensor.

Definition 5.3. Let $d, n \in \mathbb{N}$, $\mathbb{F}$ be a field and $\tau : [n]^d \to \mathbb{F}$ be a tensor. Then, the rank of τ is the smallest $r \in \mathbb{N}$ such that τ can be written as a sum of r rank one tensors.

A fairly standard dimension-based argument shows for all $n, d \in \mathbb{N}$ and every field $\mathbb{F}$, there exist tensors $\tau : [n]^d \to \mathbb{F}$ of rank at least n^{d-1}/d. However, constructing an *explicit* family of tensors of such high rank (or in fact even rank $n^{d(1/2+\epsilon)}$ for any $\epsilon > 0$) continues to be a challenging open problem. In addition to being a clean and natural question on its own, one of the reasons of interest in this question of tensor rank lower bounds is its connection to arithmetic circuit lower bounds. For instance, it is known that a three-dimensional tensor of rank at least r implies a lower bound of $\Omega(r)$ on an arithmetic circuit computing the bi-linear function associated with the tensor. As mentioned earlier, Raz [19] proved that strong enough lower bounds for d-dimensional tensors (for

$d = O(\log n / \log\log n)$) would even imply super-polynomial lower bounds for general arithmetic formulas.

In this section, we observe that our techniques also provide a polynomial degree upper bounds for the set of tensors of (border) rank at most $n^2/300$. The ideas in this proof then generalize immediately to d-dimensional tensors and give Theorem 1.4.

Lemma 5.4. *Let $\mathbb{F}$ be any field. There is a polynomial map $P : \mathbb{F}^{3nr} \to \mathbb{F}^{n^3}$ of degree 3 such that for every 3-dimensional tensor $\tau : [n]^3 \to \mathbb{F}$ of rank at most r lies in its image.*

Proof. The map follows naturally from the definition of rank.

Let $\mathbf{u}_1, \ldots, \mathbf{u}_r, \mathbf{v}_1, \ldots, \mathbf{v}_r, \mathbf{w}_1, \ldots, \mathbf{w}_r$ be disjoint n-tuples of variables. Let U be a tensor over the ring $\mathbb{F}[\mathbf{u}_1, \ldots, \mathbf{u}_r, \mathbf{v}_1, \ldots, \mathbf{v}_r, \mathbf{w}_1, \ldots, \mathbf{w}_r]$ defined as follows:

$$U(\mathbf{u}, \mathbf{v}, \mathbf{w}) = \sum_{i=1}^{r} \mathbf{u}_i \otimes \mathbf{v}_i \otimes \mathbf{w}_i \,.$$

From the definition of U, it can be readily observed that for every tensor $\tau : \mathbb{F}^{[n]^3} \to \mathbb{F}$ of rank at most r, there is a setting $\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}$ of the variables in $\mathbf{u}, \mathbf{v}, \mathbf{w}$, respectively, such that $U(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}) = \tau$. Moreover, each of the coordinates of U is a polynomial of degree equal to three in the variables in $\mathbf{u}, \mathbf{v}, \mathbf{w}$. Let P be the degree three polynomial map that maps the variables $\mathbf{u}_1, \ldots, \mathbf{u}_r, \mathbf{v}_1, \ldots, \mathbf{v}_r$ and $\mathbf{w}_1, \ldots, \mathbf{w}_r$ to the coordinates of U. □

We now argue that for $r = n^2/300$, the image of the polynomial map P given by Lemma 5.4 has an equation of not too large degree.

Theorem 5.5. *Let $\mathbb{F}$ be any field. There exists a non-zero polynomial $Q \in \mathbb{F}[x_{1,1,1}, \ldots, x_{n,n,n}]$, of degree at most n^4, which is a non-trivial equation for three-dimensional tensors $\tau : [n] \times [n] \times [n] \mapsto \mathbb{F}$ of rank at most $n^2/300$.*

Proof. We once again instantiate Lemma 2.1, with $r = n^2/300$. This time $N = n^3$, $D = 3$ and $K = 3nr = n^2/100$. Picking $\Delta = n^4$, we get that $\binom{n^4+n^3}{n^3} \geq n^{n^3}$, whereas $\binom{n^3/100+3n^4}{n^3/100} \leq (2n^4)^{n^3/100}$, which implies that existence of the desired equation. □

The arguments here also generalize to tensors in higher dimensions. In particular, the following analog of Lemma 5.4 is true:

Lemma 5.6. *Let $\mathbb{F}$ be any field. Then, for all $n, d \in \mathbb{N}$, there is a polynomial map $P : \mathbb{F}^{dnr} \to \mathbb{F}^{n^d}$ of degree at most d such that for every d-dimensional tensor $\tau : [n]^{\otimes d} \to \mathbb{F}$ of rank at most r lies in its image.*

Setting $r = n^{d-1}/(100d)$ and combining this lemma with a dimension comparison argument analogous to that in the proof of Theorem 5.5 gives Theorem 1.4. We skip the details of the proof.

We remark that similar methods can be used to prove the existence of an equation of degree poly(n) for three-dimensional tensors of *slice rank* (see, e.g., Reference [5]) at most, say, $n/1{,}000$. The existence of such an equations was proved (using different techniques) in Reference [5].

6 APPLICATIONS TO CIRCUIT LOWER BOUNDS

In this section, we prove Corollary 1.3. The strategy of the proof is simple: The proof of Theorem 1.2 implies a PSPACE algorithm that produces a sequence of polynomials that are equations for the set of matrices with small linear circuits. If those equations require large circuits, then we are done, and if not, then there exists an equation with small circuits that (assuming PIT $\in$ P) can be found using an NP-oracle. Using, once again, the assumption that PIT $\in$ P, we can also find

deterministically a matrix on which the equation evaluates to non-zero, which implies the matrix requires large linear circuits.

There are some technical difficulties involved with this plan that we now describe. The first problem is that even arithmetic circuits of small size can have large description as bit strings, due to the field constants appearing in the circuits. To prevent this issue, we only consider *constant free* arithmetic circuits, which are only allowed inputs labeled by $\{0, \pm 1\}$ (but can still compute other constants in the circuit using arithmetic operations).

The second problem is that, to be able to find a non-zero of the equation in the last step of the algorithm (using the mere assumption that $\mathsf{PIT} \in \mathsf{P}$), we need not only the size of the circuit but also its *degree* to be bounded by $\mathrm{poly}(n)$. Of course, by Theorem 1.2, there exists such a circuit, but we need to be able to prevent a malicious prover from providing us with a $\mathrm{poly}(n)$ size circuit of exponential degree, and it is not known how to compute the degree of a circuit in deterministic polynomial time, even assuming $\mathsf{PIT} \in \mathsf{P}$. To solve this issue, we use an idea of Malod and Portier ([16], Theorem 1), who showed that any polynomial with circuit of size $\mathrm{poly}(n)$ and degree d also has a ***multiplicatively disjoint* (MD)** circuit of size $\mathrm{poly}(n, d)$. An MD circuit is a circuit in which every multiplication gate multiplies two disjoint subcircuits. This is a syntactic notion that is easy to verify efficiently and deterministically, and an MD circuit of size s is guaranteed to compute a polynomial of degree at most s.

A final technical issue is that the notion of MD circuits does not fit perfectly within the framework of constant free circuits. Therefore, we use the notion of "almost MD" circuits, where the inputs to a multiplication gates are not disjoint, but at least one of these inputs is the root of a subcircuit in which only constants appear.

Definition 6.1. We say a gate v in a circuit is ***constant producing* (CP)** if in the subcircuit rooted at v, all input nodes are field constants.

An *almost-MD circuit* is a circuit where every multiplication gate either multiplies two disjoint subcircuits, or at least one of its children is constant producing.

Lemma 6.2. *Suppose f is an n-variate polynomial of degree* $\mathrm{poly}(n)$ *that has a constant free arithmetic circuit of degree* $\mathrm{poly}(n)$. *Then f has a constant free almost-MD circuit of size* $\mathrm{poly}(n)$.

Proof. Let C_0 be a constant free arithmetic circuit for f. We first homogenize the circuit C_0 to obtain a circuit C_1 (a homogeneous circuit is a circuit in which every gate computes a homogeneous polynomial; see, e.g., Reference [22]). Since C_1 is homogeneous, all the gates that compute non-zero field constants are CP gates. We then eliminate all gates that compute constants by allowing the edges entering sum gates to be labeled by field scalars and interpreting a sum gate as computing a linear combination whose coefficients are given by the edge labels. We call this circuit C_2. Note that we are now in a more general model in which the edges of the circuit are allowed to be labeled by field constants, and this step does not maintain constant-freeness. However, the new circuit still computes the same polynomial, and every label appearing on the edges of C_2 was computed in C_1, so it can be computed by a constant-free arithmetic circuit of polynomial size.

We now do the transformation detailed in Lemma 2 of Reference [16] to C_2 to obtain an MD circuit C_3, which has labels on the edges (we emphasize that Lemma 2 of Reference [16] indeed applies for this more general model of circuits with labels on the edges). This step does not produce new constants. Finally, we convert C_3 to an almost-MD constant free circuit C_4 by re-computing every label appearing on the edge using a fresh subcircuit for each label and rewiring the circuit. accordingly (this step will convert the circuit from an MD circuit to an almost MD circuit). These subcircuits are guaranteed to have polynomial size constant free circuits, since these constant were all computed in C_0, which keeps the total size $\mathrm{poly}(n)$. □

For circuits that compute low-degree polynomials, the mere existence of an algorithm for the decision version of PIT allows one to construct an algorithm for the search version.

Lemma 6.3. *Suppose* PIT $\in$ P. *Then there is a polynomial time algorithm that given a non-zero almost-MD arithmetic circuit C of size s computing an n-variate polynomial, finds in time* $\mathrm{poly}(n, s)$ *an element* $\mathbf{a} \in \{0, 1, \ldots, s\}^n$ *such that* $C(\mathbf{a}) \neq 0$.

Proof. We abuse notation by denoting by C also the polynomial computed by the circuit C. Note that, since C is almost-MD, the degree of C is at most s. Thus, there exists $a_1 \in \{0, 1, \ldots, s\}$ such that $C(a_1, x_2, \ldots, x_n)$ is a non-zero polynomial in $x_2, \ldots, x_n$. By iterating over those $s + 1$ values from 0 to s and using the assumption that PIT $\in$ P, we can find such a value for a_1 in time $\mathrm{poly}(n, s)$. We then continue in the same manner with the rest of the variables. □

As we noted above, the assumption that C is almost-MD was used in Lemma 6.3 to bound the degree of the circuit. It is also useful, because it is easy to decide in deterministic polynomial time whether a circuit is almost-MD. We now complete the proof of Corollary 1.3.

Proof of Corollary 1.3. For every n, the proof of Theorem 1.2 provides an equation Q_n for the set of $n \times n$ matrices with small linear circuits. This polynomial can be found by solving a linear system of equations in a linear space whose dimension is $\exp(\mathrm{poly}(n))$. Using standard, small space algorithm for linear algebra [1, 6], this implies that there exists a PSPACE algorithm that, on input 1^n, outputs the list of coefficients of the polynomial Q_n.

Consider now the family $\{Q_n\}_{n \in \mathbb{N}}$. If for any constant $k \in \mathbb{N}$ there exist infinitely many $n \in \mathbb{N}$ such that Q_n requires circuits of size at least n^k, then it follows (by definition) that the PSPACE algorithm above outputs a family of polynomials with super-polynomial constant-free arithmetic circuits.

We are thus left to consider the case that there exists a constant $k \in \mathbb{N}$ such that for all large enough $n \in \mathbb{N}$, Q_n can be computed by circuits of size n^k. By Lemma 6.2, we may assume without loss of generality that these circuits are almost-MD circuits. Further suppose PIT $\in$ P. We will show how to construct a matrix in polynomial time using an NP oracle (for a language L that we define) that requires large linear circuits.

Consider the language L of pairs $(1^n, x)$ such that there exists a string y of length at most n^k such that xy describes an almost-MD circuit C such that C is non-zero, and $C \circ U \equiv 0$, where U is the polynomial map given in the proof of Theorem 1.2.

Assuming PIT $\in$ P, the language L is in NP, and by assumption for every large enough n there exists such a circuit. Thus, we can use the NP oracle to construct such a circuit C bit-by-bit. Finally, using Lemma 6.3, we can output a matrix M such that $C(M) \neq 0$.

By the properties of the circuit C and the map U, M does not have linear circuits of size less than $n^2/200$. □

Many variations of Corollary 1.3 can be proved as well, with virtually the same proof. By slightly modifying the language L used in the proof, it is possible to prove the same result even under the assumption PIT $\in$ NP (recall that PIT $\in$ coRP). A similar statement also holds over finite fields of size $\mathrm{poly}(n)$, in which case the proof is simpler, since there are no issues related to the bit complexity of the field constants. Finally, we record the fact that an analog of Corollary 1.3 also holds for tensor rank, using an identical proof that uses Theorem 5.5 instead of Theorem 1.2.

Corollary 6.4. *Assuming* PIT $\in$ P*, at least one of the following is true:*

(1) *There exists a family of n-variate polynomials of degree* $\mathrm{poly}(n)$ *over* $\mathbb{C}$*, which can be computed (as its list of coefficients, given the input* 1^n*) in* PSPACE*, which does not have polynomial size constant free arithmetic circuits.*

(2) *There is an efficient construction with an* NP *oracle of a three-dimensional tensor of rank* $\Omega(n^2)$.

We remark that for tensors of large rank there are no analogs of References [2, 3], i.e., there do not exist even constructions with an NP oracle of tensors with slightly super-linear rank.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their careful reading and thoughtful comments on a preliminary version of this article.

REFERENCES

[1] Eric Allender, Robert Beals, and Mitsunori Ogihara. 1999. The complexity of matrix rank and feasible systems of linear equations. *Comput. Complex.* 8, 2 (1999), 99–126. DOI: https://doi.org/10.1007/s000370050023

[2] Josh Alman and Lijie Chen. 2019. Efficient construction of rigid matrices using an NP oracle. In *Proceedings of the 60th Annual IEEE Symposium on Foundations of Computer Science (FOCS'19)*. IEEE Computer Society, 1034–1055. DOI: https://doi.org/10.1109/FOCS.2019.00067

[3] Amey Bhangale, Prahladh Harsha, Orr Paradise, and Avishay Tal. 2020. Rigid matrices from rectangular PCPs or: Hard claims have complex proofs. In *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS'20)*. IEEE, 858–869. DOI: https://doi.org/10.1109/FOCS46700.2020.00084

[4] Markus Bläser and Christian Ikenmeyer. 2017. Introduction to geometric complexity theory. http://pcwww.liv.ac.uk/~iken/teaching_sb/summer17/introtogct/gct.pdf.

[5] Markus Bläser, Christian Ikenmeyer, Vladimir Lysikov, Anurag Pandey, and Frank-Olaf Schreyer. 2019. Variety membership testing, algebraic natural proofs, and geometric complexity theory. *CoRR* abs/1911.02534 (2019).

[6] Allan Borodin, Joachim von zur Gathen, and John E. Hopcroft. 1982. Fast parallel matrix and GCD computations. *Inf. Control.* 52, 3 (1982), 241–256. DOI: https://doi.org/10.1016/S0019-9958(82)90766-5

[7] Prerona Chatterjee, Mrinal Kumar, C. Ramya, Ramprasad Saptharishi, and Anamay Tengse. 2020. On the existence of algebraically natural proofs. In *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS'20)*. https://arxiv.org/abs/2004.14147.

[8] Michael A. Forbes, Amir Shpilka, and Ben Lee Volk. 2018. Succinct hitting sets and barriers to proving lower bounds for algebraic circuits. *Theor. Comput.* 14, 1 (2018), 1–45. DOI: https://doi.org/10.4086/toc.2018.v014a018

[9] Fulvio Gesmundo, Jonathan D. Hauenstein, Christian Ikenmeyer, and J. M. Landsberg. 2016. Complexity of linear circuits and geometry. *Found. Computat. Math.* 16, 3 (2016), 599–635. DOI: https://doi.org/10.1007/s10208-015-9258-8

[10] Joshua A. Grochow, Mrinal Kumar, Michael E. Saks, and Shubhangi Saraf. 2017. Towards an algebraic natural proofs barrier via polynomial identity testing. *CoRR* abs/1701.01717 (2017).

[11] Valentine Kabanets and Russell Impagliazzo. 2004. Derandomizing polynomial identity tests means proving circuit lower bounds. *Computat. Complex.* 13, 1–2 (2004), 1–46. DOI: https://doi.org/10.1007/s00037-004-0182-6

[12] Abhinav Kumar, Satyanarayana V. Lokam, Vijay M. Patankar, and Jayalal Sarma. 2014. Using elimination theory to construct rigid matrices. *Computat. Complex.* 23, 4 (2014), 531–563. DOI: https://doi.org/10.1007/s00037-013-0061-0

[13] Richard J. Lipton. 1978. Polynomials with 0-1 coefficients that are hard to evaluate. *SIAM J. Comput.* 7, 1 (1978), 61–69. DOI: https://doi.org/10.1137/0207004

[14] Satyanarayana V. Lokam. 2009. Complexity lower bounds using linear algebra. *Found. Trends Theoret. Comput. Sci.* 4, 1-2 (2009), 1–155. DOI: https://doi.org/10.1561/0400000011

[15] Meena Mahajan and Jayalal Sarma. 2010. On the complexity of matrix rank and rigidity. *Theor. Comput. Syst.* 46, 1 (2010), 9–26. DOI: https://doi.org/10.1007/s00224-008-9136-8

[16] Guillaume Malod and Natacha Portier. 2008. Characterizing valiant's algebraic complexity classes. *J. Complex.* 24, 1 (2008), 16–38. DOI: https://doi.org/10.1016/j.jco.2006.09.006

[17] Ketan Mulmuley and Milind A. Sohoni. 2001. Geometric complexity theory I: An approach to the P vs. NP and related problems. *SIAM J. Comput.* 31, 2 (2001), 496–526. DOI: https://doi.org/10.1137/S009753970038715X

[18] Ran Raz. 2010. Elusive functions and lower bounds for arithmetic circuits. *Theor. Comput.* 6, 7 (2010), 135–177. DOI: https://doi.org/10.4086/toc.2010.v006a007

[19] Ran Raz. 2013. Tensor-rank and lower bounds for arithmetic formulas. *J. ACM* 60, 6 (2013), 40:1–40:15. DOI: https://doi.org/10.1145/2535928

[20] Ramprasad Saptharishi. 2015. A survey of lower bounds in arithmetic circuit complexity. (2015). Retrieved from https://github.com/dasarpmar/lowerbounds-survey/releases/.

[21] Amir Shpilka and Ilya Volkovich. 2015. Read-once polynomial identity testing. *Computat. Complex.* 24, 3 (2015), 477–532. DOI: https://doi.org/10.1007/s00037-015-0105-8

[22] Amir Shpilka and Amir Yehudayoff. 2010. Arithmetic circuits: A survey of recent results and open questions. *Found. Trends Theoret. Comput. Sci.* 5 (Mar. 2010), 207–388. DOI: https://doi.org/10.1561/0400000039

[23] Volker Strassen. 1974. Polynomials with rational coefficients which are hard to compute. *SIAM J. Comput.* 3, 2 (1974), 128–149. DOI: https://doi.org/10.1137/0203010

[24] Leslie G. Valiant. 1977. Graph-theoretic arguments in low-level complexity. In *Proceedings of the 2nd International Symposium on the Mathematical Foundations of Computer Science (MFCS 1977) (Lecture Notes in Computer Science)*, Jozef Gruska (Ed.), Vol. 53. Springer, 162–176. DOI: https://doi.org/10.1007/3-540-08353-7_135

Received May 2021; revised April 2022; accepted June 2022