Geometric Complexity Theory V: Efficient algorithms for Noether Normalization[*]

Dedicated to Sri Ramakrishna

Ketan D. Mulmuley [†]
The University of Chicago
mulmuley@uchicago.edu

May 26, 2016

Abstract

We study a basic algorithmic problem in algebraic geometry, which we call NNL, of constructing a normalizing map as per Noether's Normalization Lemma. For general explicit varieties, as formally defined in this paper, we give a randomized polynomial-time Monte Carlo algorithm for this problem. For some interesting cases of explicit varieties, we give deterministic quasi-polynomial time algorithms. These may be contrasted with the standard EXPSPACE-algorithms for these problems in computational algebraic geometry.

In particular, we show that:

(1) The categorical quotient for any finite dimensional representation V of SL_m, with constant m, is explicit in characteristic zero.

(2) NNL for this categorical quotient can be solved deterministically in time quasi-polynomial in the dimension of V.

(3) The categorical quotient of the space of r-tuples of $m \times m$ matrices by the simultaneous conjugation action of SL_m is explicit in any characteristic.

(4) NNL for this categorical quotient can be solved deterministically in time quasi-polynomial in m and r in any characteristic $p \notin [2, \lfloor m/2 \rfloor]$.

(5) NNL for every explicit variety in zero or large enough characteristic can be solved deterministically in quasi-polynomial time, assuming the hardness hypothesis for the permanent in geometric complexity theory.

The last result leads to a geometric complexity theory approach to put NNL for every explicit variety in P.

1 Introduction

Noether's Normalization Lemma (NNL), proved by Hilbert [43], is the basis of a large number of foundational results in algebraic geometry, such as Hilbert's Nullstellensatz. It also lies at

[*] To appear in the Journal of the AMS.

[†] This work was supported by NSF grant CCF-1017760. This article is the full version of its FOCS 2012 extended abstract [69].

the heart of the foundational classification problem of algebraic geometry. For any projective variety $W \subseteq P(K^l)$, where K is an algebraically closed field and $P(K^l)$ is the projective space associated with K^l, the lemma says that any homogeneous, generic linear map $\psi : K^l \to K^k$, for any $k \geq \dim(W) + 1$, induces a regular (well defined) map on W (this means ψ does not vanish identically on the line through the origin in K^l corresponding to any point in W). Furthermore, for any such ψ, (1) $\psi(W) \subseteq P(K^k)$, the image of W, is closed in $P(K^k)$, and (2) the fiber $\psi^{-1}(p)$, for any point $p \in \psi(W)$, is a finite set. Accordingly, we call a homogeneous linear map $\psi : K^l \to K^k$, $k \geq \dim(W) + 1$, that induces a regular map on W a *normalizing map* for W. In the context of the main results of this paper, l here will be exponential in $\dim(W)$, and k will be polynomial in $\dim(W)$. In this case, Noether's Normalization Lemma expresses the variety W, embedded in the ambient space $P(K^l)$ of exponential dimension, as a finite cover of the variety $\psi(W)$, embedded in the ambient space $P(K^k)$ of polynomial dimension. This is its main significance from the complexity-theoretic perspective. We also refer to the problem of constructing a normalizing map ψ, with $k = \text{poly}(\dim(W))$, as NNL in short. This is the problem that is studied in this article for the varieties W that are given explicitly in a sense that will be made precise. We do not require $k = \dim(W) + 1$ here, and allow a polynomial slack, for the reasons explained in Section 1.2.

In algebraic geometry, the phrase "explicitly" is used informally. In this article, it is interpreted formally, from the complexity-theoretic perspective, to mean "using algebraic circuits that can be computed in deterministic polynomial time". Thus, we formally introduce in this article the notion of an *explicit* family $\{W_n\}$ of varieties (Definition 5.1) of poly(n) dimension that can be specified succinctly and uniformly by poly(n)-time-computable algebraic circuits (cf. Section 2.1) of poly(n) degree having a specification of poly(n) bit-length, even though the dimension l_n of the ambient space containing W_n can be exponential in n. If W_n is projective, we let l_n be one plus the dimension of the ambient space. If the family $\{W_n\}$ is explicit, we also say that the variety W_n is explicit, with the understanding that $n \to \infty$ in all complexity bounds. It turns out that (cf. Section 5.1) a large class of varieties that arise in practice are explicit in this formal complexity-theoretic sense.

The problem NNL for such explicit varieties W_n (cf. Definition 5.6) is the problem of constructing a specific kind (cf. Sections 1.2 and 5.3) of a normalizing map $\psi_n : K^{l_n} \to K^{k_n}$ for W_n, with $k_n = \text{poly}(n)$, having a succinct specification of poly(n) bit-length.

For general explicit varieties W_n, the standard algorithm for NNL (cf. Section 5.6), based on Gröbner basis theory [61], takes in the worst case work-space that is polynomial in l_n, and time that is exponential in l_n. In our context l_n, in general, is exponential in n, and hence, this work-space bound is exponential in n, i.e., $O(2^{\text{poly}(n)})$, and the time bound is double exponential in n. This shows that NNL for explicit varieties is in EXPSPACE. Assuming the Generalized Riemann Hypothesis, it can be shown to be in EXPH (cf. Section 5.6). Here EXPSPACE denotes the complexity class of problems that can be solved using exponential work-space in double exponential time, and EXPH denotes the exponential hierarchy [3]. Informally, these stand for the classes of problems that are computationally highly intractable (far more intractable than the problems in the class NP, which can be solved in exponential time and polynomial space). Thus, on the basis of the existing literature in computational algebraic geometry, it may appear that NNL for explicit varieties is highly intractable, and perhaps, even inherently so.

The algorithmic results in this article indicate that this is not the case.

First, it is shown in this article (cf. Theorem 1.2) that NNL for any explicit variety W_n can be solved by a poly(n)-time randomized Monte Carlo algorithm, whose output is correct with a high probability. This means, in practice, NNL for explicit varieties can be solved efficiently and correctly with a high probability. But this does not show that NNL for explicit varieties is in BPP $\subseteq$ PSPACE, for the reasons explained in Section 1.3. Hence, it does not affect the current EXPSPACE-status of NNL, or the EXPH status assuming the Generalized Riemann Hypothesis.

So we ask if NNL for any explicit variety can be solved *deterministically* in polynomial time, thereby bringing it down from EXPSPACE to P. We say that NNL for an explicit variety has an *explicit solution*, in the complexity-theoretic sense, if it can be solved in deterministic polynomial time. The motivation for such an explicit solution comes from the foundational classification problem of algebraic geometry (cf. Section 1.8 in [40]). Its goal [20] is to classify a given algebraic variety by transforming it regularly into some *canonical* normal form. Without any relaxation, this goal may be infeasible, since it is not even known at present if the isomorphism problem for algebraic varieties is decidable [90]. Hence, our goal is to do the best that we can from the complexity-theoretic perspective. As a first step in this direction, one would like an "explicit" normalizing map for an "explicitly" given variety. A random normalizing map is not enough in this context, since randomness is the opposite of canonicity. Solving NNL in deterministic polynomial time is this first step towards the classification problem of algebraic geometry, interpreted from the complexity-theoretic perspective. For the reasons explained in Section 11, this turns out to be far harder than solving NNL in randomized polynomial time by a Monte Carlo algorithm. We turn to this harder problem next.

It is shown in this article that, for some interesting cases of explicit varieties W_n, NNL can indeed be solved deterministically in quasi-poly(n)-time, i.e., in $O(2^{(\log n)^c})$ time, for some constant $c > 1$. (Here it is assumed that the dimension k_n of the target space of the constructed normalization map $\psi_n : K^{l_n} \to K^{k_n}$ is also quasi-polynomial in n.) Thus, for these explicit varieties, NNL can be brought down from EXPSPACE to quasi-P, the class of problems that can be solved deterministically in quasi-polynomial time.

The first such case of an explicit variety is the categorical quotient [75] V/G associated with any finite dimensional (rational) representation V of $G = SL_m$, with constant m, in characteristic zero. (By a rational representation, we mean that the entries of the representation matrix are rational functions of the coordinates of G. We will only be concerned with such representations in this article.) Here $V/G = \text{spec}(K[V]^G)$ [1] is the variety whose coordinate ring is $K[V]^G \subseteq K[V]$, where $K[V]$ denotes the coordinate ring of V, and $K[V]^G$ its subring of G-invariants. Explicitness of this variety is by itself a key result in this article. It means (cf. Theorem 1.5) that a succinct encoding, in the form of a symbolic determinant, of a set of (exponentially many) generators for this invariant ring can be constructed in poly(n) time, where n is the dimension of V.

This succinct and efficient encoding of generators is in the spirit of the encodings that were used in the so-called symbolic method of classical invariant theory (cf. Chapter 8 A in [94]). For example, the First Fundamental Theorem for the ring of vector invariants proved by Weyl [94] (cf. Theorem 2.6 A therein) implies such a polynomial-time-computable succinct encoding, in the form of a symbolic determinant, of a set of (exponentially many) generators for this invariant

[1] By abuse of notation, spec in this article really means max-spec; cf. [27] (page 54).

ring. The problem of proving similar First Fundamental Theorems for invariant rings has been studied intensively in the last century; cf. Section 9 in [79] for a survey. This classical problem is interpreted in this article (cf. Definition 5.2 (d)), from the complexity-theoretic perspective, as the problem of constructing an explicit encoding, in the form of a symbolic determinant or a circuit, of a set of generators of the invariant ring, where *explicit* means polynomial-time-computable. Classical invariant theory did not specify formally what "explicit" means.

For the variety V/G associated with any n-dimensional representation V of $G = SL_m$, with constant m, in characteristic zero, it is shown in this article that NNL can be solved deterministically in $O(n^{O(\log \log n)})$ time; cf. Theorem 1.6.

Noether's Normalization Lemma was, in fact, proved by Hilbert [43] to give an algorithm for constructing a finite set of generators for the invariant ring $K[V]^G$ in this context. Hilbert did not prove any explicit upper bound on its running time, or on the degrees of the generators. Such a bound on the degrees was proved in Popov [78] a century later, and improved significantly in Derksen [16]. This improved analysis yields an exponential-time algorithm for computing a set of generators for the ring $K[V]^G$ of invariants for any finite dimensional representation V of $G = SL_m$ and, in conjunction with Gröbner basis theory [61], an EXPSPACE-algorithm for NNL for this invariant ring. This algorithm for constructing a set of generators requires time that is exponential in the dimension of V, and the algorithm for NNL requires exponential work-space and double exponential time, even when m is constant. Hilbert's paper focused mainly on the case when m is three, since an algorithm to construct a finite set of generators was not known before even in this case; cf. Section 1.5.

Explicitness for constant m of the categorical quotient associated with this invariant ring $K[V]^G$ (Theorem 1.5) implies that the problem of computing an encoding, in the form of a symbolic determinant, of a set of generators for this invariant ring is in P. Thus there has been a rather remarkable change in the status of this fundamental problem of invariant theory over the course of a century from a problem that was not even known to be computable before Hilbert to a problem that is now in P, as shown in this article; cf. Section 1.5.

The quasi-polynomial-time deterministic algorithm in this article for NNL for this invariant ring (cf. Theorem 1.6), for constant m, brings the original instance of NNL in Hilbert's paper in this case from EXPSPACE to quasi-P. Analogous results hold for any connected, reductive, algebraic group of constant dimension (cf. Theorem 9.9).

The second case of an explicit variety that we consider is the categorical quotient V/G [75] associated with the space $V = M_m(K)^r$ of r-tuples of $m \times m$ matrices over K, with the simultaneous conjugation-action of $G = SL_m$ (without any restriction on m this time). It is shown in this article that this variety is explicit in characteristic zero, and is explicit in a relaxed sense in positive characteristic (cf. Theorem 1.3).

Furthermore (cf. Theorem 1.4), NNL for this variety can be solved deterministically in quasi-poly(m, r) time in any characteristic $p \notin [2, \lfloor m/2 \rfloor]$, thereby bringing NNL in this case too from EXPSPACE to quasi-P. This extends the same result in characteristic zero that is implied, as pointed out by Forbes and Shpilka [31], by a variant of a conditional result in the preliminary version [69] of this article, in conjunction with their earlier work [30] on arithmetic circuits; cf. Remark 1 in Section 1.7.

More generally (cf. Theorems 1.8, 5.11, and Section 10.5), NNL for any explicit variety in zero or large enough characteristic can be solved deterministically in quasi-polynomial time, thereby bringing it from EXPSPACE to quasi-P, assuming the hardness hypothesis for the permanent in geometric complexity theory. This hypothesis proposed in [71] (or rather its stronger variant) is that the permanent of $n \times n$ matrices cannot be approximated infinitesimally closely by symbolic determinants over K of $O(2^{n^\epsilon})$ size, for some constant $\epsilon > 0$, as $n \to \infty$. It is an algebraic geometric strengthening of the fundamental VP $\neq$ VNP conjecture in the work of Valiant [91].

In Bürgisser [13], some other consequences of this hardness hypothesis have been derived, which also crucially rely on the fundamental result of Kaltofen [50, 52], as in this paper. Consulting [71, 13, 72, 73, 14] (and especially Section 9.3 in [14] that explains in detail the precise relationship of the work in [13] with the earlier work in [71]) may help the reader to understand this hypothesis better.

1.1 Geometric complexity theory approach to the basic algorithmic problems in algebraic geometry and invariant theory

The results described above lead to the following geometric complexity theory approach to the basic algorithmic problems of algebraic geometry and invariant theory under consideration, namely, (1) the problem NNL, and (2) the problem of constructing a set of generators for the ring of invariants of a reductive group. Both these problems are motivated by Hilbert [43].

The goal of the approach in the context of the first problem is to show that it is in P for every explicit variety. The approach is to (1) first prove the hardness hypothesis [71] for the permanent in geometric complexity theory (or its weaker form, cf. Theorem 1.9), then (2) use the results in this article to show that NNL for every explicit variety is in quasi-P, in zero or large enough characteristic (cf. Theorem 5.11 and Section 10.5), and (3) finally, remove the quasi-prefix and the characteristic restriction by proving a stronger form of the hardness hypothesis (cf. Sections 5.5 and 10.5). An approach to prove the required hardness hypothesis in geometric complexity theory will be given in the sequel to this article (the revised version of [64]).

(N.B. The current version of [64] on the arxiv has become outdated in view of the recent result [46] that the occurrence-based obstructions in [73], based on the vanishing of the rectangular Kronecker coefficients, cannot be used to prove superpolynomial lower bounds for the permanent. The approach in the revised version of [64] will be based on the far more powerful multiplicity-based obstructions, to which this negative result does not apply.)

To put NNL in quasi-P, one does not need the hardness hypothesis in geometric complexity theory, or even its weaker form, in full strength for all explicit varieties. For explicit varieties of intermediate difficulty, such as explicit categorical quotients, one only needs weaker complexity-theoretic hypotheses for the classes of circuits depending on the varieties. Thus one can approach this goal step by step, increasing the hardness of the varieties in tandem with the strength of the circuits; cf. Sections 5 and 10.

The goal of the approach in the context of the second problem is to show that it is in P for every invariant ring $K[V]^G$, for any finite dimensional representation V of a connected reductive group G in characteristic zero, and after a relaxation, for any reductive group in any characteristic, allowing encoding of a set of generators by algebraic circuits; cf. Conjecture 5.3

and Section 9.6. The results in this article (Theorems 1.3 and 1.5) show this for some important rings of invariants. The class of encoding circuits depends on the ring of invariants. For example, when $G = SL_m(\mathbb{C})$, with constant m, one only needs depth four circuits; cf. Theorem 8.5. Thus one can again approach the goal step by step, increasing the hardness of the ring of invariants in tandem with the strength of the circuits; cf. Section 9.

If the goals for both the problems (1) and (2) are achieved in a stronger form, it would follow that NNL for every categorical quotient V/G is in P, and moreover, that the closed G-orbits in V have an explicit (polynomial-time-computable) parametrization, for any finite dimensional representation V of a reductive group G in any characteristic; cf. Sections 9.6 and 10.3.

There is a fundamental difference between this approach to the basic algorithmic problems of algebraic geometry and invariant theory and the standard approaches in computational algebraic geometry [61] and computational invariant theory [17]. The difference lies in how the basic objects of algebraic geometry–namely, the varieties–are specified in the computer. Computational algebraic geometry, based on Gröbner basis theory [61] and the theory of solving polynomial equations [56, 55], and computational invariant theory [78, 17] use the standard specification of the varieties in terms of their defining equations. If one uses this standard specification, then the basic algorithmic problems of algebraic geometry and invariant theory are inherently intractable. For example, Gröbner basis computation is EXPSPACE-hard [61], solving polynomial equations (Hilbert's Nullstellensatz) is NP-hard [36], NNL is NP-hard (cf. Section 3), the problem of constructing a finite set of generators for the ring $K[V]^G$ of invariants is inherently intractable, because the number of generators of this ring can be exponential in $\dim(V)$ even when $\dim(G)$ is constant (cf. the proof of Proposition 9.3), and hence, the standard [75] parametrization of the closed G-orbits in V by the points of the categorical quotient V/G is also inherently intractable.

But this article illustrates that a large class of algebraic varieties that arise in practice can be specified *explicitly* using circuits, the basic objects of complexity theory. If one uses instead this explicit complexity-theoretic specification of the basic geometric objects (varieties), as in the approach here, then the results in this article indicate that the basic algorithmic problems (1) and (2) in algebraic geometry and invariant theory, along with the basic problem in geometric invariant theory [75] of parametrizing closed orbits in representations of reductive groups, which are inherently intractable in the standard specification, are tractable in the explicit specification. The formal notion of explicitness introduced in this article (Definition 5.1), which is thus the fundamental difference between the geometric complexity theory approach to these basic problems and the standard approaches, is the driving theme of this article.

This article belongs to a series [71, 73, 70, 6] of articles on geometric complexity theory. See [67, 66] for an overview of the earlier articles in this series, and [14] for an overview of the mathematical issues therein. Preliminary versions of the results here were announced in [69].

We now state the main results of this article in more detail.

Notation: Till Section 10, K will henceforth denote an algebraically closed base field of characteristic zero, unless mentioned otherwise. We use the standard notation for the complexity classes, such as P (the class of problems that can be solved in polynomial time), BPP (the class of decision problems that can be solved by polynomial time Monte Carlo algorithms), NC (the class of problems that can be solved in poly-logarithmic parallel time using polynomial number of processors), DET (the class of problems LOGSPACE-reducible to computation of the determinant of integer matrices), EXP (the class of problems that can be solved in exponential time), EXPSPACE (the class of problems that can be solved in exponential work-space), PSPACE (the class of problems that can be solved in polynomial work-space), AC^0 (the class of problems that can be solved by constant depth Boolean circuits of polynomial size), PH (the polynomial hierarchy), EXPH (the exponential hierarchy), and so on. See [3, 15] for their formal definitions.

1.2 The problem NNL

We now define the problem NNL for the explicit variety associated with the determinant in the first article [71] in this series.

This explicit variety, denoted as $\Delta[\det, m]$, is defined as follows. Let $X = (x_1, \ldots, x_r)$ be a tuple of r variables. For convenience, let us assume that $r = m^2$, so that X can be thought of as an $m \times m$ variable matrix, identifying x_i's with the entries of X in any way. By a homogeneous symbolic determinant of size m over $X = (x_1, \ldots, x_r)$, we mean the determinant of a symbolic $m \times m$ matrix, whose each entry is a homogeneous linear function over K of $x_1, \ldots, x_r$. Let $\mathcal{X}$ be the vector space over K of homogeneous polynomials of degree m in the variables $x_1, \ldots, x_r$, and $P(\mathcal{X})$ the projective space associated with $\mathcal{X}$. Let $\Sigma[\det, m] \subseteq P(\mathcal{X})$ be the set of all points in $P(\mathcal{X})$ that correspond to nonzero homogeneous polynomials in $\mathcal{X}$ that can be expressed as homogeneous symbolic determinants of size m over X. Then $\Delta[\det, m] \subseteq P(\mathcal{X})$ is the Zariski-closure $\overline{\Sigma[\det, m]}$ of $\Sigma[\det, m]$. Its dimension is $\leq m^4$. Informally, $\Delta[\det, m]$ is explicit because it can be specified *succinctly* by a small circuit over $\mathbb{Q}$ of poly(m) total bit-size for computing $\det(X)$, and for a given m, the specification of such a circuit can be computed in poly(m) time [93]. For a formal proof of explicitness, see Section 5.1.1. Since the circuit [93, 60] for computing the determinant is very special (namely, weakly skew, cf. [60] and Section 2.1), we call $\Delta[\det, m]$ *strongly explicit*.

It has to be stressed here that $\Delta[\det, m]$ is *not* specified by giving its equations in the ambient space $P(\mathcal{X})$. This is not even possible using poly(m) bits, since the dimension of $P(\mathcal{X})$ is exponential in m. All complexity bounds for the results below for $\Delta[\det, m]$ are in terms of the $O(\text{poly}(m))$ bit-length of its succinct specification. Thus an EXPSPACE-algorithm means an algorithm that takes work-space that is exponential in m, a P-algorithm means an algorithm that takes time that is polynomial in m, and so on.

Let $\hat{\Delta}[\det, m] \subseteq \mathcal{X}$ denote the affine cone of $\Delta[\det, m]$. This is defined to be the union of all lines through the origin in $\mathcal{X}$ that correspond to the points of $\Delta[\det, m] \subseteq P(\mathcal{X})$. Let $R(\det, m)$ denote the homogeneous coordinate ring of $\Delta[\det, m]$. This is the same as the coordinate ring of $\hat{\Delta}[\det, m]$.

By Noether's Normalization Lemma (Lemma 3.1), there exists a homogeneous linear map $\psi : \mathcal{X} \to K^k$, for any $k > \dim(\Delta[\det, m])$, such that ψ does not vanish on any nonzero point in $\hat{\Delta}[\det, m] \subseteq \mathcal{X}$. Hence, ψ yields a regular (well-defined) map from $\Delta[\det, m]$ to $P(K^k)$, which we denote by ψ again. We call such a ψ a *normalizing map* (for $\Delta[\det, m]$). Any generic ψ for such k is a normalizing map. But deterministic construction or even verification of a normalizing map, as we shall see below, is very difficult.

Let x_i, $1 \leq i \leq k$, denote the coordinates of K^k, and given a normalizing map ψ, let

$\psi^*(x_i) : \mathcal{X} \to K$ denote the pullback of x_i via ψ. We also denote its restriction to $\hat{\Delta}[\det, m]$ by $\psi^*(x_i)$. If $k = \dim(\Delta[\det, m]) + 1$, the minimum possible value, then we call the subset $\{\psi^*(x_i) \mid 1 \leq i \leq k\} \subseteq R(\det, m)$ an *h.s.o.p. (homogeneous system of parameters)* for $\Delta[\det, m]$. Existence of such an h.s.o.p. is a classical fact that holds for any variety; cf. Section 3.

An h.s.o.p. for $\Delta[\det, m]$ can be constructed in work-space that is exponential in m, and in time that is double exponential in m (cf. Theorem 4.1), by first computing the equations of $\Delta[\det, m]$ as a subvariety of $P(\mathcal{X})$ using Gröbner basis theory [61]. This space requirement is exponential in m, because the number of variables in the equations of $\Delta[\det, m]$ as a subvariety of $P(\mathcal{X})$ is equal to the dimension of $\mathcal{X}$, which is exponential in m, and Gröbner basis computation [61] takes work-space that is at least polynomial in the number of variables. If we insist on an h.s.o.p., then this is the best that can be done at present. However, if we do not insist on the optimal $k = \dim(\Delta[\det, m]) + 1 \leq m^4$, but allow a slack, and only require that k be poly(m), then we can do much better.

Accordingly, we define *the problem NNL for* $\Delta[\det, m]$ as the problem of constructing a normalizing map ψ for $k = \text{poly}(m)$, not necessarily optimal, with a *succinct* specification of poly(m) bit-length. Thus we let go of optimality but insist on succinctness. We have to now explain what we mean by succinct. Obviously, the standard specification of ψ as a linear map from $\mathcal{X} \to K^k$ is not succinct, since the dimension of $\mathcal{X}$ is exponential in m. Hence, we confine ourselves to normalizing maps which have a succinct specification as follows.

For any $m \times m$ matrix B with rational entries, let ψ_B denote the homogeneous, linear, evaluation map on $\mathcal{X}$, which maps a polynomial $p(X) \in \mathcal{X}$ to $p(B)$. We denote its restriction to $\hat{\Delta}[\det, m]$ by ψ_B again. Given any set $\mathcal{B} = \{B_1, \ldots, B_k\}$ of $m \times m$ matrices with rational entries, let $\psi_{\mathcal{B}} : \mathcal{X} \to K^k$ denote the homogeneous linear map that maps $p = p(X) \in \mathcal{X}$ to $(\psi_{B_1}(p), \ldots, \psi_{B_k}(p))$. Let $S(\mathcal{B}) = \{\psi_{B_i} \mid 1 \leq i \leq k\} \subseteq R(\det, m)$.

We call $S(\mathcal{B})$ an *s.s.o.p. (small system of parameters)*[2] for $\Delta[\det, m]$ if (1) the total bit-length of B_i's is poly(m), and (2) the homogeneous linear map $\psi_{\mathcal{B}}$ does not vanish on any non-zero point in $\hat{\Delta}[\det, m] \subseteq \mathcal{X}$. Hence, $\psi_{\mathcal{B}}$ yields a regular (well-defined) map from $\Delta[\det, m]$ to $P(K^k)$, which we denote by $\psi_{\mathcal{B}}$ again. We specify the s.s.o.p. $S(\mathcal{B})$ *succinctly* by giving the matrices in $\mathcal{B}$. We call $\psi_{\mathcal{B}}$ the *succinct normalizing map* corresponding to this s.s.o.p.

It can be shown that an s.s.o.p. exists (Corollary 4.4). However, an s.s.o.p. with the optimal cardinality equal to $\dim(\Delta[\det, m]) + 1$ may not exist. This is why we allowed the slack above.

A poly(m)-time-constructible s.s.o.p. is called an *e.s.o.p. (explicit system of parameters)*, where explicit means poly(m)-time-constructible. Quasi-s.s.o.p. and quasi-e.s.o.p. are defined by replacing poly(m) by quasi-poly$(m) := 2^{\text{polylog}(m)}$ throughout in the definitions.

The *problem NNL for* $\Delta[\det, m]$ is to construct an s.s.o.p. for $\Delta[\det, m]$, given the succinct specification of $\Delta[\det, m]$ in the form a circuit for computing $\det(X)$. We say that NNL for $\Delta[\det, m]$ has an *explicit solution* if $\Delta[\det, m]$ has an e.s.o.p.

The current best, unconditional, deterministic algorithm for constructing an s.s.o.p. for $\Delta[\det, m]$, based on Gröbner basis theory [61], also takes work-space that is exponential in m, and time that is double exponential in m (cf. Theorem 4.10), as in the case of an h.s.o.p., again

[2] Such an s.s.o.p. is later called a strict s.s.o.p., as per the terminology in Section 5.3, wherein we introduce a more general definition of an s.s.o.p. But we shall not worry about this issue in this section.

because the dimension of the ambient space $P(\mathcal{X})$ containing $\Delta[\det, m]$ is exponential in m.

1.3 A Monte Carlo algorithm

The following result shows that, if we are satisfied with Monte Carlo algorithms, then an s.s.o.p. for $\Delta[\det, m]$ can be constructed efficiently and correctly, with a high probability.

Theorem 1.1 *An s.s.o.p. for $\Delta[\det, m]$ can be constructed by a poly(m)-time randomized Monte Carlo algorithm, whose output is correct with a high probability.*

Hilbert's original paper [43] itself gives a randomized Monte Carlo algorithm to construct a normalizing map for any variety. For $\Delta[\det, m]$, the algorithm is the following: Just choose a random, homogeneous, linear map from $P(\mathcal{X})$ to $P(K^k)$, with $k > \dim(\Delta[\det, m])$. It can be shown using Gröbner basis theory [61] (cf. the proof of Theorem 4.1) that it is a normalizing map with a high probability, if the entries of the matrix specifying this map are large enough randomly chosen integers of bit-length exponential in $\dim(\mathcal{X})$. Since $\dim(\mathcal{X})$ is exponential in m in our context, the number of random bits used by this algorithm and its running time are thus double exponential in m. In contrast, the randomized algorithm in Theorem 1.1 uses only poly(m) random bits and poly(m) time. This is possible because the normalizing map constructed by this algorithm has a succinct specification. Obviously, the usual matrix representation of a linear map from $\mathcal{X}$ to K^k is not succinct, since $\dim(\mathcal{X})$ is exponential in m.

The Monte Carlo algorithm in Theorem 1.1 is not a BPP-algorithm, since NNL is not a decision problem, but rather a construction problem, whose output is not uniquely defined. More importantly, BPP is known to be in PH $\subseteq$ PSPACE [3]. In contrast, Theorem 1.1 does *not* imply that NNL for $\Delta[\det, m]$ is in PSPACE. As already mentioned, at present we can only show unconditionally that it is in EXPSPACE. This is because the problem of verifying correctness of the output of the Monte Carlo algorithm in Theorem 1.1, a potential s.s.o.p., turns out to be very difficult. If the problem of verifying an s.s.o.p. were in PSPACE, then it would have followed from Theorem 1.1 that NNL for $\Delta[\det, m]$ is in PSPACE. But the current best algorithm for this verification requires exponential work-space (cf. Theorem 4.10).

Further results on NNL for $\Delta[\det, m]$ will be given in Section 1.6.

Theorem 1.2 *Theorem 1.1 holds with any explicit variety (cf. Definition 5.1) in place of $\Delta[\det, m]$.*

Theorems 1.1 and 1.2 hold in arbitrary characteristic; cf. Section 10.5.

We next turn to some exceptional instances of explicit varieties for which NNL can be solved deterministically in quasi-polynomial time using the existing techniques.

1.4 The ring of matrix invariants

The first such instance is the categorical quotient [75] associated with the ring of matrix invariants.

Let $M_m(K)$ be the space of $m \times m$ matrices over K. Let $V = M_m(K)^r$, the direct sum of r copies of $M_m(K)$, with the adjoint (simultaneous conjugate) action of $G = SL_m(K)$.

Let $U = (U_1, \ldots, U_r)$ be an r-tuple of variable $m \times m$ matrices. The variable entries of U_i's can be thought of as the coordinates of V, and the coordinate ring $K[V]$ of V can be identified with the ring $K[U_1, \ldots, U_r]$ generated by the variable entries of U_i's. An invariant in $K[V]$ is a polynomial $f(U_1, \ldots, U_r)$ in the variable entries of U_i's such that

$$f(U_1, \ldots, U_r) = f(P^{-1}U_1P, \ldots, P^{-1}U_rP),$$

for all $P \in G$. Let $n = \dim(V) = rm^2$. Let $K[V]^G \subseteq K[V]$ be the subring of invariants. It is finitely generated [43, 80]. Hence, by a general construction of algebraic geometry, one can associate with it the variety $V/G = \mathrm{spec}(K[V]^G)$, called the *categorical quotient* [75].

We say that V/G is *strongly explicit* if, given m and r, one can construct in poly(n) time a symbolic matrix $A(U, y)$ of poly(n) size such that: (1) each entry of $A(U, y)$ is a (possibly non-homogeneous) linear function, with rational coefficients, of the entries of U_i's and auxiliary variables $y = (y_1, \ldots, y_k)$, $k = \mathrm{poly}(n)$, and (2) the coefficients of $\det(A(U, y))$, considered as a polynomial in y, belong to and generate $K[V]^G$.

Theorem 1.3 *The categorical quotient V/G is strongly explicit. It is strongly explicit in a relaxed sense (cf. Definition 5.2) if K is an algebraically closed field of positive characteristic.*

We specify V/G and $K[V]^G$ *succinctly* by simply specifying V and G, which can be done by giving m and r in unary. This succinct specification is polynomial-time-equivalent to the succinct specification of V/G as an explicit variety in terms of the circuit for $\det(A(U, y))$, since, by Theorem 1.3, this circuit can be computed in poly(n) time, given m and r. The pair (m, r) in unary will thus be the input in all the problems for V/G and $K[V]^G$ described below. The bit-length of this succinct specification of V/G is $O(m + r) = O(n)$, even though the dimension of the ambient space containing V/G (cf. Section 6.2) is exponential in m. All space and time bounds for the algorithms below with this input will be in terms of n.

By Noether's Normalization Lemma (Lemma 3.2), there exists a set $S \subseteq K[V]^G$ of poly(n) homogeneous invariants such that $K[V]^G$ is integral over the subring generated by S.[3] (This statement of Noether's Normalization Lemma is equivalent to the one given in the beginning of this introduction.) In fact, there even exists such an S of optimal cardinality equal to $\dim(K[V]^G)$ (which is less than n). It is known that any generically chosen S of this cardinality has the required property. Such an S of optimal cardinality is called an *h.s.o.p. (homogeneous system of parameters)* of $K[V]^G$. (Existence of such an h.s.o.p. is again a classical fact, cf. Section 3, that holds for any finitely generated K-algebra.) It is shown here (cf. Theorem 7.3) that the problem of constructing an h.s.o.p. for V/G is in EXPH (the exponential hierarchy), assuming the Generalized Riemann Hypothesis. The hierarchy is exponential, and not polynomial, because the dimension of the ambient space containing V/G is exponential in m, and hence the current best PH-algorithm for Hilbert's Nullstellensatz in Koiran [55] becomes an EXPH-algorithm in our context. If we insist on an h.s.o.p., then this is the best that we can

[3] A ring R is said to be integral over its subring T if every $r \in R$ satisfies a monic polynomial equation of the form $r^l + b_{l-1}r^{l-1} + \ldots + b_1 r + b_0 = 0$, where each $b_i \in T$.

do at present. However, we can do much better if, as in Section 1.2, we relax the optimality constraint on the cardinality, but insist on succinctness of specification in exchange. We are thus led to the following notion of an s.s.o.p.

We call a set $S \subseteq K[V]^G$ an *s.s.o.p. (small system of parameters)* for $K[V]^G$ if (1) S contains poly(n) homogeneous invariants of poly(n) degree, (2) $K[V]^G$ is integral over the subring generated by S, and (3) each invariant $s = s(U_1, \ldots, U_r)$ in S can be expressed as a symbolic determinant of $O(\text{poly}(n))$ size, i.e., as the determinant of a symbolic matrix M_s of $O(\text{poly}(n))$ size, whose entries are linear (possibly non-homogeneous) functions, with rational coefficients, of the variable entries of U_i's, and (4) for each $s \in S$, the total bit-size of the specification of the symbolic matrix M_s (including the total bit-size of the constants therein) is $O(\text{poly}(n))$.

We call S an *e.s.o.p.(explicit system of parameters)* if, in addition, given m and r, the specification of S, consisting of a symbolic matrix M_s as above for each $s \in S$, can be computed in poly(n) time. This is a specialization to V/G of the general definition of an e.s.o.p. for strongly explicit varieties (Definition 5.6) given later. Quasi-s.s.o.p. and quasi-e.s.o.p. are defined by replacing poly(n) by quasi-poly(n) throughout in the definitions.

It can be shown that an s.s.o.p. exists (cf. Corollary 5.10).

By *the problem NNL for $K[V]^G$ or V/G*, we mean the problem of constructing an s.s.o.p. for $K[V]^G$, given the succinct specification of $K[V]^G$ in terms of the unary pair (m, r). This definition of NNL is a specialization and simplification of the general definition of NNL for explicit varieties (Definition 5.6) given later. We say that NNL for $K[V]^G$ has an *explicit solution* if $K[V]^G$ has an e.s.o.p.

Theorem 1.4 *(For characteristic zero, see [69], [30, 31], and Remark 1 in Section 1.7) Let $V = M_m(K)^r$, and $G = SL_m(K)$ as above. Then $K[V]^G$ has a quasi-e.s.o.p., assuming that K is an algebraically closed field of characteristic $p \notin [2, \lfloor m/2 \rfloor]$.*

A stronger form of this result (Theorem 10.1) implies quasi-explicit (i.e., quasi-polynomial-time computable) parametrization of the closed G-orbits in V for any characteristic $p \notin [2, \lfloor m/2 \rfloor]$; cf. Theorem 10.15. Analogous results hold for the invariant ring associated with any quiver; cf. Theorem 10.8.

1.5 The general ring of invariants

We now describe another exceptional explicit variety for which NNL can be solved deterministically in quasi-polynomial time with the existing techniques. This is the categorical quotient associated with any invariant ring of $SL_m(K)$, with constant m, in characteristic zero.

Let V be any finite dimensional representation of $G = SL_m(K)$, with arbitrary m for the moment. Since G is reductive [33], V can be decomposed as a direct sum of irreducible representations of G:

$$V = \sum_{\lambda} m(\lambda) V_\lambda(G). \tag{1}$$

Here $\lambda : \lambda_1 \geq \ldots \geq \lambda_l$, $l < m$, is a partition, i.e., a non-increasing sequence of non-negative integers, $V_\lambda(G)$ is the irreducible representation of G (Weyl module [33]) labelled by λ, and $m(\lambda)$

is its multiplicity. Fix the standard monomial basis [24, 57] for each $V_\lambda(G)$, and thus a standard monomial basis for V. Let $v = (v_1, \ldots, v_n)$, $n = \dim(V)$, be the coordinates of V in this basis. Let $K[V] = K[v_1, \ldots, v_n]$ be the coordinate ring of V. Let $K[V]^G$ be its subring of G-invariants. We call a polynomial $f(v) \in K[V]$ a *G-invariant* if $f(\sigma^{-1}v) = f(v)$ for all $\sigma \in G$. By Hilbert [43], $K[V]^G$ is finitely generated. Hence, one can associate with it the categorical quotient [75] $V/G = \text{spec}(K[V]^G)$. We specify V/G and $K[V]^G$ *succinctly* by just giving the specification $\langle V, G\rangle$ of V and G, consisting of n and m (in unary), and the multiplicities $m(\lambda)$'s (in unary) for all λ's that occur with nonzero multiplicity in the decomposition (1). The bit-length of this succinct specification is $O(n + m)$, though the dimension of the ambient space containing V/G is exponential in n, even when m is constant; cf. the proof of Proposition 9.3.

We call V/G *strongly explicit* if, given $\langle V, G\rangle$, one can compute in poly(n, m) time a symbolic matrix $A(v, y)$ such that: (1) each entry in $A(v, y)$ is a (possibly non-homogeneous) linear function, with rational coefficients, of the coordinates $v = (v_1, \ldots, v_n)$ and auxiliary variables $y = (y_1, \ldots, y_k)$, $k = \text{poly}(n, m)$, and (2) the coefficients of $\det(A(v, y))$, considered as a polynomial in y, belong to and generate $K[V]^G$.

Theorem 1.5 *Let V be a finite dimensional representation of $G = SL_m(K)$, with constant m, as above. Then V/G is strongly explicit.*

S.s.o.p., e.s.o.p., quasi-s.s.o.p., quasi-e.s.o.p. for strongly explicit V/G are defined just as for the ring of matrix invariants in Section 1.4, except that we use the coordinates $v_1, \ldots, v_n$ of V in place of the coordinates for $M_m(K)^r$ used earlier, and the succinct specification $\langle V, G\rangle$ of $K[V]^G$ in place of the succinct specification of the ring of matrix invariants by the pair (m, r) earlier; cf. Definition 9.1 for details.

For constant m, we define a *near-e.s.o.p.* for $K[V]^G$ by replacing poly(n, m) in the definition of an e.s.o.p. by $O(n^{O(\log \log n)})$ throughout.

By *the problem NNL for $K[V]^G$ or V/G*, we mean the problem of constructing an s.s.o.p. for $K[V]^G$, given $\langle V, G\rangle$ as above. We say that NNL for $K[V]^G$ has *an explicit solution* if $K[V]^G$ has an e.s.o.p.

Theorem 1.6 *Let V be a finite dimensional representation of $G = SL_m(K)$, with constant m, as above. Then $K[V]^G$ has a near-e.s.o.p.*

Analogues of Theorems 1.5 and 1.6 hold for any connected, reductive, algebraic group of constant dimension (cf. Theorem 9.9).

By Theorems 9.7 and 2.1, the ring $K[V]^G$ has a quasi-e.s.o.p., without any restriction on m, if (1) the permanent of an $n \times n$ variable matrix X cannot be computed by symbolic determinants over X of $O(2^{n^\epsilon})$ size[4], for some constant $\epsilon > 0$, as $n \to \infty$ (cf. Valiant [91]), and (2) V/G is explicit (cf. Conjecture 5.3 and the remark thereafter). Analogous results hold for any reductive, algebraic group (possibly disconnected) in zero or large enough characteristic (cf. Sections 9.6 and 10.5).

Classical invariant theory mainly studied the invariant ring $K[V]^G$ for constant m, as in Theorem 1.6, because of Gordan's seminal work (cf. [35] and Section 3.7 in [89]) that gave an

[4] Here the entries of the symbolic matrices are possibly non-homogeneous linear functions of X.

algorithm for constructing a finite set of generators for the ring of invariants of binary forms. In this case, V is the space of binary forms with the natural action of $SL_2(K)$, and $m = 2$. In the modern terminology, Gordan showed that the problem of constructing finitely many generators for the ring invariants of binary forms is computable, though the formal notion of computability was developed much later. It was not known before Hilbert if this holds for general m, or even for $m = 3$. It was not even known that finitely many generators exist when $m = 3$. This was shown by Hilbert in his first paper [42] for any m. But this proof was non-constructive. It was severely criticized by Gordan (cf. Section 3.7 in [89] for this story), since it did not give an algorithm for constructing a set of generators. In the modern terminology, it did not yield a proof, as Gordan sought, for computability of the problem of constructing a set of generators for $K[V]^G$. Such a proof was given by Hilbert in his second paper [43], as a response to Gordan's criticism. For these reasons, the second paper mainly focused on the case when $m = 3$. Theorem 1.5, or rather its stronger form (Theorem 8.5), implies that the problem of constructing a set of generators for $K[V]^G$ for constant m is, in fact, in DET $\subseteq$ NC $\subseteq$ P, allowing encoding of the set by a symbolic determinant. Theorem 1.6 shows that the original instance of NNL in Hilbert [43], with constant m, is in quasi-P.

1.6 Noether normalization vs. hardness

Next we ask if NNL for $\Delta[\det, m]$ can be solved deterministically in poly(m) time. For the reasons explained later in Section 11, this turns out to be a much harder problem than the analogous problems for the special cases of explicit varieties addressed in Theorems 1.4 and 1.6. At present, we only have a conditional result:

Theorem 1.7 *The variety $\Delta[\det, m]$ has a quasi-e.s.o.p., if the permanent of an $n \times n$ variable matrix X cannot be approximated infinitesimally closely by symbolic determinants over X of size $\leq 2^{n^\epsilon}$, for some constant $\epsilon > 0$, as $n \to \infty$.*

Entries of the symbolic matrices here are allowed to be non-homogeneous linear functions of the entries of X, with coefficients in K. When $K = \mathbb{C}$, by *infinitesimally close approximation* of the permanent, we mean that, for any $\delta > 0$, there exists a symbolic determinant over X of size $\leq 2^{n^\epsilon}$, such that the distance between the coefficient vectors of perm(X) and the symbolic determinant in the L_2-norm is less than δ.

The lower bound assumption for the permanent in the result above is a stronger form of the hardness hypothesis for the permanent in geometric complexity theory (cf. Conjecture 4.3 in [71]), with $\Omega(2^{n^\epsilon})$ lower bound in place of the superpolynomial lower bound.

Actually, we prove a stronger result (Theorem 4.7) that, under this lower bound assumption, NNL for $\Delta[\det, m]$ can even be solved fast in parallel.

Theorem 1.8 *Theorem 1.7 holds with any explicit variety (cf. Definition 5.1) in place of $\Delta[\det, m]$.*

By these results, NNL for any explicit variety can be brought down from EXPSPACE to quasi-P, assuming the hardness hypothesis [71] for the permanent in geometric complexity theory. The quasi-prefix can be removed under a stronger assumption (cf. Theorem 4.5).

Theorem 1.7 is a consequence of the following stronger result. It shows that solving NNL for $\Delta[\det, m]$ in deterministic polynomial time is in fact equivalent, ignoring a quasi-prefix, to proving a weaker implication of the hardness hypothesis in Theorem 1.7.

Theorem 1.9 *The variety $\Delta[\det, m]$ has an e.s.o.p. iff, ignoring a quasi-prefix, there exists a family $\{f_n(x_1, \ldots, x_n)\}$ of exponential-time-computable, integral, multi-linear polynomials such that f_n cannot be approximated infinitesimally closely by symbolic determinants over $(x_1, \ldots, x_n)$ of size $\leq 2^{n^\epsilon}$, for some constant $\epsilon > 0$, as $n \to \infty$.*

By exponential-time-computable, we mean that the polynomial can be computed, given an integral input, in time that is exponential in the total bit-length of the input.

Theorems 1.7, 1.8, 1.9, and their analogues for explicit varieties also hold in large enough positive characteristics (cf. Section 10.5). Furthermore, the largeness restriction on the characteristic can be removed assuming a slight extension of the hardness hypothesis; cf. Remark 1 at the end of Section 10.

These results establish an essential equivalence between the problem NNL for explicit varieties and the weaker form of the hardness hypothesis [71] in geometric complexity theory.

1.7 Proof technique

We now briefly explain how the main results in this article are proved.

The formal notion of explicitness introduced in this article (Definition 5.1) lies at the heart of the proofs, along with the fundamental work [43, 75, 78, 16, 80, 82, 32, 57] in algebraic geometry and geometric invariant theory, the fundamental work [91, 93, 88, 50, 51, 52, 60] in algebraic complexity theory, and the fundamental work [41, 76, 47, 49, 86, 2, 30, 29] on a derandomization problem in complexity theory, called black-box polynomial identity testing. Derandomization means converting a randomized efficient algorithm into a deterministic efficient algorithm by removing random bits. Theorems 1.4 and Theorems 1.6–1.9 are proved by derandomizing the Monte Carlo algorithm in Theorem 1.2 for the explicit varieties under consideration, unconditionally or assuming a suitable hardness hypothesis. Derandomization of this Monte Carlo algorithm for a given explicit variety amounts to bringing NNL for that variety from EXPSPACE, where it is by the general result (Theorem 5.12), to P. This EXPSPACE vs. P gap in the complexity of NNL that needs to be bridged to derandomize this Monte Carlo algorithm for a given explicit variety is the basic difference between derandomization in this article and derandomization in the earlier articles [76, 47, 49, 2, 30, 86] in complexity theory, wherein such a gap is absent. The use of geometric invariant theory [75] in derandomization, as in the proofs of Theorems 1.4 and 1.6, is another basic difference. In contrast, the earlier works [76, 47, 49, 2, 30, 86] on derandomization in complexity theory do not use any invariant theory.

The efficient Monte Carlo algorithm for NNL for explicit varieties in Theorems 1.1 and 1.2 is based on the classical results in algebraic geometry due to Hilbert and others, and the fundamental work in Heintz and Schnorr [41] on black-box polynomial identity testing; cf. Section 4.2.

Theorem 1.3 on explicitness of the categorical quotient associated with the ring of matrix invariants is proved in characteristic zero using the First and Second Fundamental Theorems for matrix invariants due to Procesi and Razmyslov [80, 82]; cf. Section 6.

The situation in positive characteristic turns out to be much harder. The analogous First Fundamental Theorem for matrix invariants in positive characteristic in Donkin [23] is too weak for the proof of Theorem 1.3 in positive characteristic, since the only known upper bound [21] for the degrees of the generators in [23] is exponential in the size m of the matrices. The crux of the proof of Theorem 1.3 in positive characteristic is the *geometric First Fundamental Theorem* (cf. Theorem 10.2) proved in this article, which provides a set of separating [17] matrix-invariants of polynomial degree in arbitrary characteristic. This is proved here using the criterion for stability in arbitrary characteristic due to Hilbert [43], Mumford et al. [75], and King [54], and the fundamental Brauer-Nesbitt theorem [8, 26] in modular representation theory.

The explicitness result in Theorem 1.3 lies at the heart of the proof of Theorem 1.4.

Theorem 1.3, in conjunction with Theorem 1.2 (which holds in arbitrary characteristic, cf. Section 10.5), implies that NNL for the ring of matrix invariants has a polynomial-time Monte Carlo algorithm in arbitrary characteristic.

Theorem 1.4 is proved by derandomizing this Monte Carlo algorithm, up to a quasi-prefix; cf. Sections 7 and 10.1. This is done in two steps.

The first crucial step (for the reasons that will become clear in Section 11) is to show that this Monte Carlo algorithm can be derandomized assuming the standard black-box derandomization hypothesis for symbolic determinant identity testing [41, 47, 49, 1], which is recalled in Section 2.3 here. This is shown using the fundamental work in Hilbert [43] and Mumford et al. [75]; cf. Theorem 5.13, Remark 3 thereafter, and Remark 2 in Section 10.5. This implies that NNL for the ring of matrix invariants has a deterministic polynomial-time algorithm assuming the standard black-box derandomization hypothesis for symbolic determinant identity testing.

Remark 1 (on the second step of derandomization): In the preliminary version [69] of this article, only this conditional result was proved in characteristic zero. Subsequently it was pointed out by Forbes and Shpilka [31] that the step in the proof of this result wherein symbolic determinant identity testing enters can be modified, as explained in Section 7.5 here, so as to use instead the polynomial identity testing for read-once oblivious algebraic branching programs (cf. Section 2.1). A quasi-polynomial-time deterministic black-box algorithm for this problem was already known from their earlier work [30]. Thus this instance of NNL can be solved deterministically in quasi-polynomial time in characteristic zero using the existing techniques. This is contrary to what was suggested in the preliminary version [69] of this article, because of the relationship of this problem with the wild ("impossible") problem [25] of classifying matrix tuples (though this instance of NNL itself is not wild).

This proof of Theorem 1.4 in characteristic zero can be extended to any characteristic $p \notin [2, \lfloor m/2 \rfloor]$, using the refined form of the Geometric First Fundamental Theorem for matrix invariants (cf. Theorem 10.2) proved in this article, in place of the First Fundamental Theorem for matrix invariants due to Procesi and Razmyslov [80, 82]; cf. Section 10.1.

Theorem 1.5 on explicitness of the categorical quotient associated with the general ring of invariants of SL_m, for constant m, is proved using the fundamental works in geometric invariant theory due to Hilbert [43], Mumford et al. [75], and Derksen and Kemper [16, 17], in conjunction with standard monomial theory [57], and the fundamental works in algebraic complexity theory due to Strassen [88], Valiant [91], Malod and Portier [60], and others; cf. Section 8.

The explicitness result in Theorem 1.5 lies at the heart of the proof of Theorem 1.6.

Theorem 1.5, in conjunction with Theorem 1.2, implies that NNL for the general ring of invariants of SL_m, for constant m, has a polynomial-time Monte Carlo algorithm in characteristic zero.

Theorem 1.6 is proved by derandomizing this Monte Carlo algorithm, up to a quasi-prefix; cf. Section 9. This is again done in two steps.

The first crucial step is, again, to show that this Monte Carlo algorithm can be derandomized assuming the standard black-box derandomization hypothesis for symbolic determinant identity testing. This can be done (just as in the case of Theorem 1.4) using the work of Hilbert [43] and Mumford et al. [75]; cf. Theorem 5.13 and Remark 3 thereafter. This implies that NNL for the general ring of invariants of SL_m, for constant m, has a deterministic polynomial-time algorithm in characteristic zero, assuming the standard black-box derandomization hypothesis (cf. Section 2.3) for symbolic determinant identity testing.

Using a refined form (Theorem 8.5) of Theorem 1.5 in the first step, it follows that NNL for the general ring of invariants of SL_m, for constant m, has a deterministic polynomial-time algorithm assuming a weaker black-box derandomization hypothesis for diagonal depth three circuits [83].

This hypothesis was already known to hold, up to a quasi-prefix, from the earlier work of Shpilka and Volkovich [86], and Agrawal, Saha, and Saxena [2]. Thus it follows that NNL for V/G as in Theorem 1.6 can be solved in quasi-polynomial time deterministically. This was the result that was stated in the preliminary version [69] of this article. The stronger $O(n^{O(\log\log n)})$-time bound stated in Theorem 1.6 follows in view of the recent result in Forbes, Saptharishi, and Shpilka [29], which gives an $O(s^{O(\log\log s)})$-time-computable black-box derandomization of polynomial identity testing for diagonal depth three circuits of size $\leq s$.

Let us now turn to Theorem 1.9, Theorem 1.7 being its corollary.

The first step is to show that the Monte Carlo polynomial-time algorithm for NNL for $\Delta[\det, m]$ in Theorem 1.1 can be derandomized assuming a strengthened form, introduced in this article (cf. Section 2.5), of the standard black-box derandomization hypothesis [41, 47, 49, 1] for symbolic determinant identity testing; cf. Section 4.3.

By Kabanets and Impagliazzo [49], the standard hypothesis holds, up to a quasi-prefix, assuming a sub-exponential symbolic determinant lower bound for some family of exponential-time-computable, integral, multi-linear polynomials.

It is similarly shown in this article (cf. Theorem 2.4 and the remark thereafter) that the strengthened hypothesis holds, up to a quasi-prefix, if there exists a family $\{f_n(x_1, \ldots, x_n)\}$ of exponential-time-computable, integral, multi-linear polynomials such that f_n cannot be approximated infinitesimally closely by symbolic determinants of size sub-exponential in n. This is proved using the fundamental work on black-box factorization of multivariate polynomials in Kaltofen and Trager [52], which lies at the heart of this proof, in conjunction with the fundamental hardness vs. randomness principle in Nisan and Wigderson [76], and Kabanets and Impagliazzo [49].

This implies the reduction from NNL to hardness stated in Theorem 1.9. The reduction in the other direction is easy (cf. Lemma 4.8 and Proposition 2.7).

Theorem 1.8 and the generalization of Theorem 1.9 for general explicit varieties (Theorem 5.14) follow by systematically extending the proofs of Theorems 1.7 and 1.9; cf. Section 5.

The proofs of these results can be extended to large enough positive characteristics using the standard techniques of algebraic geometry and algebraic complexity theory; cf. Section 10.5.

Conditional generalizations of Theorem 1.6 to explicit categorical quotients associated with representations of general reductive algebraic groups are given in Section 9.6. These can be proved by extending the proof for SL_m using the standard techniques in geometric invariant theory [75] and representation theory.

Organization of the paper

The rest of this paper is organized as follows.

Logical structure of the proofs: The proofs of the main results are presented in the following steps. (1) The variety under consideration is shown to be explicit. (2) An EXPSPACE-algorithm is given for constructing an h.s.o.p. for the variety. For the categorical quotient associated with the ring of matrix invariants, a more efficient EXPH-algorithm is given, assuming the Generalized Riemann Hypothesis. (3) An efficient Monte Carlo algorithm is given for constructing an s.s.o.p. for the variety. (4) This algorithm is derandomized using the strengthened or the standard form of the black-box derandomization hypothesis for an appropriate class of circuits. Which form is used depends on the closure properties of the variety. The class of circuits also depends on the variety. (5) If this class is sufficiently restricted, as happens for the categorical quotients associated with the ring of matrix invariants and the general ring of invariants of SL_m with constant m, then this black-box derandomization is carried out unconditionally, up to a quasi-prefix. (6) Otherwise, it is shown that the black-box derandomization hypothesis holds assuming an appropriate hardness hypothesis.

Organization of the sections: In Section 2, we introduce the strengthened form of the standard [41, 47, 49, 1] black-box derandomization hypothesis for polynomial identity testing, and prove the essential equivalence between strengthened black-box derandomization and sub-exponential algebraic circuit size lower bounds for infinitesimally close approximation. This is a key ingredient in the proofs of Theorems 1.7, 1.8, and 1.9. In Section 3, we recall Noether's Normalization Lemma, and show that the problem of constructing an h.s.o.p. for a general variety, specified in the standard fashion by its defining equations, belongs to PH, assuming the Generalized Riemann Hypothesis. In Section 4, we study the basic prototype $\Delta[\det, m]$ of an explicit variety, and prove Theorems 1.1, 1.7, and 1.9. In Section 5, we formulate the general notion of an explicit variety motivated by its basic prototype $\Delta[\det, m]$, define the problem NNL for explicit varieties, and prove Theorems 1.2, 1.8, and the generalization of Theorem 1.9 for explicit varieties. In Section 6, we prove Theorem 1.3 in characteristic zero. Theorem 1.4 in characteristic zero is proved in Section 7. In Section 8, we prove Theorem 1.5. Theorem 1.6 is proved in Section 9. Theorems 1.3 and 1.4 in arbitrary characteristic, their generalizations to quivers, and extensions of Theorems 1.7, 1.8, and 1.9 to large enough positive characteristics are proved in Section 10. It is also explained in Section 10 how Theorems 1.1 and 1.2 can be extended to arbitrary characteristics. Furthermore, implications of the results in this article to explicit parametrization of closed orbits and explicit parametrization of semi-simple representations of finitely generated algebras are also given in Section 10. Finally, in Section 11, we discuss the difficulties that need to be overcome to improve the current best bound for NNL for $\Delta[\det, m]$

in Theorem 4.10.

2 Black-box polynomial identity testing

In this section, we introduce the strengthened black-box derandomization hypothesis for polynomial identity testing (cf. Section 2.5), and prove the essential equivalence between strengthened black-box derandomization and sub-exponential lower bounds for infinitesimally close approximation (cf. Theorem 2.4 and Proposition 2.7). This is a key ingredient in the proofs of Theorems 1.7, 1.8, and 1.9.

2.1 Circuits and symbolic determinants

We begin by recalling [60, 87] the circuit classes for which we need these hypotheses.

By a *circuit* over the field K [60], we mean a directed acyclic graph with vertices of in-degree zero or two, in which each node of in-degree 2 is labelled with $*$ or $+$, and each node of in-degree 0 is labelled with a variable or a constant in K. By the polynomial computed by the circuit, we mean the polynomial computed at the root, in the obvious way. By *the size of the circuit*, we mean the total number of edges in it. If the constants in the circuit are in $\mathbb{Q}$ or its finite extension, then by *the bit-size or the bit-length of the circuit*, we mean the size of the circuit plus the total bit-size of the specification of all the constants.

We say that the circuit $C = C(x_1, \ldots, x_n)$ over the variables $x_1, \ldots, x_n$ has *low or small degree* if the degree of the polynomial computed by it is $O(s^a)$, for some fixed constant $a > 0$, where s is the size of C.

By a *weakly skew circuit*, we mean [60] a circuit whose each node v labelled with $*$ has at least one child u such that the sub-circuit rooted at u is connected to the rest of the circuit by just the edge (u, v).

By a *symbolic determinant* over $x_1, \ldots, x_n$ of size m, we mean the determinant of a symbolic $m \times m$ matrix, whose each entry is a linear combination (possibly non-homogeneous) over K of of $x_1, \ldots, x_n$. Weakly skew circuits are polynomially equivalent [60] to symbolic determinants.

Weakly skew circuits are also polynomially equivalent to algebraic branching programs [87, 30]. In this article we will only use a special class of such programs called *read-once oblivious algebraic branching programs* [30]. Such a program can be specified, for some n, l, and d, by a tuple $(M_1, M_2, \ldots, M_l)$ of $n \times n$ matrices such that every entry of M_i, $1 \leq i \leq l$, is a uni-variate polynomial of degree $\leq d$ over K in the distinct variable z_i associated with M_i. The uni-variate polynomials are specified by giving all their coefficients. The size of this program is $O(n^2 ld)$. The polynomial computed by this program is defined to be trace$(\prod_i M_i)$. Clearly, this polynomial can also be computed by a weakly skew circuit of poly(n, l, d) size. Hence, such programs can be viewed as restricted classes of weakly skew circuits, or equivalently, symbolic determinants.

A *diagonal depth three circuit* [83] C over the variables $x_1, \ldots, x_n$ is a circuit that computes a sum $\sum_{i=1}^{k} f_i^{e_i}$ of powers of linear functions, where each f_i is a possibly non-homogeneous linear function of x_i's with coefficients in K. Here k is called the *top fan-in* of the circuit, and $e = \max\{e_i\}$ its *degree*. The size of this circuit is $O(nek)$.

By a *circuit with oracle gates* for a function $f(y_1, \ldots, y_r)$, we mean a circuit in which some gates are labelled with f. These gates have in-degree r. The computation of f at any such gate is assigned unit cost.

2.2 Polynomial identity testing

Next, we recall the standard black-box derandomization hypothesis for polynomial identity testing [41, 47, 49, 1] over the algebraically closed base field K.

The *polynomial identity testing* problem over K is the problem of deciding if a given circuit $C(x)$, $x = (x_1, \ldots, x_r)$, over K computes an identically zero polynomial. By *the polynomial identity testing problem for small degree circuits*, or *the low-degree polynomial identity testing problem*, we mean the polynomial identity testing problem wherein the degree of the polynomial computed by $C(x)$ is assumed to be $O(s^a)$, for some constant $a > 0$, where s is the size of $C(x)$.

There is a simple randomized polynomial-time algorithm [45] for deciding if a given circuit with rational constants computes an identically zero polynomial: just substitute large enough random integer values for the variables, and test if the circuit evaluates to zero.

The *white-box derandomization problem* [49] for polynomial identity testing is to find a deterministic polynomial time algorithm for deciding if a given circuit with rational constants computes an identically zero polynomial.

The harder *black-box derandomization problem* for polynomial identity testing [41, 47, 49, 1] over K is to construct a *hitting set* against all circuits over K with size $\leq s$ on $r \leq s$ variables, given r and s in unary. By a hitting set, we mean a set $S_{r,s} \subseteq \mathbb{N}^r$ of test inputs such that, for every circuit C over K and $x = (x_1, \ldots, x_r)$ with size $\leq s$ computing a non-zero polynomial $C(x)$, $S_{r,s}$ contains a test input b such that $C(b) \neq 0$. The *standard black-box-derandomization hypothesis for polynomial identity testing* [41, 47, 49, 1] is that there exists a poly(s)-time-computable hitting set. We call such a hitting set *explicit*. More generally, if a hitting set is computable in $O(T(s))$ time, we say that the polynomial identity testing for circuits of size $\leq s$ has $O(T(s))$-time-computable black-box derandomization.

The standard black-box derandomization hypotheses for the restricted circuit classes in Section 2.1 are defined similarly.

2.3 Symbolic determinant identity testing

We now describe such a hypothesis for symbolic determinants (cf. Section 2.1) in more detail, since it plays a crucial role in this paper.

Let $Y = Y(x_1, \ldots, x_n)$ be any $m \times m$ symbolic matrix, whose each entry is a homogeneous linear form over K in the variables $x_1, \ldots, x_n$. By *symbolic determinant identity testing*, we mean the problem of deciding, given Y, if the symbolic determinant $\det(Y)$ is an identically zero polynomial in x_i's. We call a set $S_{n,m} \subseteq \mathbb{N}^n$ of test inputs a *hitting set* in this context if, for every non-zero symbolic determinant $\det(Y)$ over $x_1, \ldots, x_n$ of size m, $S_{n,m}$ contains a test input on which that symbolic determinant does not vanish. The *standard black-box derandomization hypothesis for symbolic determinant identity testing* [41, 47, 49, 1] is that, given n and m, one can construct such a hitting set in poly(n, m) time. This is a weaker form of the standard black-box derandomization hypothesis for polynomial identity testing described in Section 2.2.

2.4 Black-box polynomial identity testing vs. hardness

The following result is a variant of Theorem 7.7 in [49]. This is why polynomial identity testing is expected to have efficient black-box derandomization.

Theorem 2.1 (Kabanets and Impagliazzo) *(cf. Theorem 7.7 in [49]) Suppose there exists a family $\{f_m(x_1,\ldots,x_m)\}$ of exponential-time-computable, multi-linear, integral polynomials such that f_m cannot be evaluated by a circuit over K of $O(2^{m^a})$ size, for some constant $a>0$, as $m\to\infty$. Then polynomial identity testing for small degree circuits over K of size $\leq s$ has $O(2^{polylog(s)})$-time-computable black-box derandomization.*

Here by an exponential-time-computable, integral polynomial $f_m(x_1,\ldots,x_m)$, we mean a polynomial such that, given an integral input $a=(a_1,\ldots,a_m)$, $f_m(a)$ can be computed in time that is exponential in the total bit-length of a. Since f_m here is multi-linear, this is equivalent to saying that the coefficient vector of f_m can be computed in time exponential in m.

The proof of Theorem 2.1 is similar to that of Theorem 7.7 in [49] (which works in the black-box model). Since we are going to prove its stronger form (Theorem 2.4) later, we only point out here how to take care of the main difference between the setting in [49] and the one here. The difference is that in [49] the size of the circuit is defined to be the total number of edges in it plus the total bit-length of the constants, whereas here the size means the total number of edges. A key ingredient in the proof in [49] is an efficient algorithm in [50] for factoring multivariate polynomials (cf. Lemma 7.6. in [49]). In its place we use instead the following result in [50, 52] that does not depend on the bit-lengths of the constants in the circuit.

Theorem 2.2 (Kaltofen and Trager) *(cf. Corollary 6.2. in [50], Theorem 1 in [52], and Theorem 2.21 in Bürgisser [12])*

Suppose $\{g_n(x_1,\ldots,x_n)\}$ is a p-computable family [91] of polynomials over K. This means g_n is a polynomial of poly(n) degree that can be computed by a nonuniform circuit over K of poly(n) size. Then each factor of g_n in $K[x_1,\ldots,x_n]$ can also be computed by a nonuniform circuit over K of poly(n) size.

More generally, given any families $\{g_n(x_1,\ldots,x_n)\}$, $\{f_n(x_1,\ldots,x_n)\}$ of polynomials over K, with f_n dividing g_n, there exists for every n a nonuniform circuit over K of poly$(n,\deg(g_n))$ size, with oracle gates for g_n, that computes f_n.

A simpler proof of the first statement in Theorem 2.2 can be found in Section 2.3 in Bürgisser [12] (cf. Theorem 2.21 therein). Bürgisser also proves a stronger statement in [13] (cf. Theorem 1.3 therein) concerning complexity of infinitesimally close approximation. For the converse of Theorem 2.1, see [41, 1].

2.5 The strengthened black-box derandomization hypothesis

Next, we formulate the *strengthened black-box derandomization hypothesis for polynomial identity testing.*

Let $x = (x_1, \ldots, x_r)$ be a tuple of r variables. The *strengthened black-box derandomization problem for small degree circuits* is to construct in poly(s) time a *hitting set* against all nonzero polynomials $f(x) \in K[x]$ of degree $\leq d = O(s^a)$, $a > 0$ a constant, that can be approximated infinitesimally closely by circuits over K and x of size $\leq s$, given r, s, and d in unary.

When $K = \mathbb{C}$, by *infinitesimally close approximation*, we mean that, for any $\epsilon > 0$, there exists such a circuit C_ϵ over K of size $\leq s$ such that the distance $||C_\epsilon(x) - f(x)||_2$ between the coefficient vectors of $C_\epsilon(x)$ and $f(x)$ in the L_2-norm is less than ϵ; cf. [13] for the definition for general K. By a *hitting set*, we mean a set $S_{r,s} \subseteq \mathbb{N}^r$ of test inputs such that, for every nonzero $f(x)$ of degree $\leq d$ that can be approximated infinitesimally closely by circuits over K of size $\leq s$, $S_{r,s}$ contains a test input b such that $f(b) \neq 0$.

The following result implies that such a hitting set exists. For any positive integer u, let $[u] := \{1, \ldots, u\}$.

Theorem 2.3 (Heintz and Schnorr) *(cf. Theorem 4.4 in [41] and its proof) A randomly chosen subset $B \subseteq [u]^r$, $u = 2s(d+1)^2$, of size $q = 6(s+1+r)^2$ is with a high probability a hitting set against all non-zero polynomials that can be approximated infinitesimally closely by circuits over K and r variables of size $\leq s$ and degree $\leq d$. Specifically, at least $(1 - u^{-q/6})$-th fraction of the sequences $(b_1, \ldots, b_q)$, $b_i \in [u]^r$, are hitting.*

The *strengthened black-box-derandomization hypothesis* for polynomial identity testing for small degree circuits is that there exists a poly(s)-time-computable hitting set $S_{r,s}$. We call such a hitting set *explicit*. More generally, if a hitting set is computable in $O(T(s))$ time, we say that the polynomial identity testing for small degree circuits of size $\leq s$ has $O(T(s))$-time-computable strengthened black-box derandomization. The *strengthened black-box derandomization hypothesis for general polynomial identity testing without any degree restrictions* is defined similarly.

The similar *strengthened black-box derandomization hypothesis for symbolic determinant identity testing* is that, given n and m, one can construct in poly(n, m) time a hitting set against all nonzero homogeneous polynomials $h(x_1, \ldots, x_n)$'s over K of degree m that can be approximated infinitesimally closely by symbolic determinants (cf. Section 2.3) of size m over $x_1, \ldots, x_n$.

The strengthened black-box derandomization hypothesis is counter-intuitive unlike the standard hypothesis in Section 2.2. Conjecturally (cf. Section 11), there exist integral polynomials of small degree that can be approximated infinitesimally closely by small circuits over K but cannot be computed exactly by such circuits. Hence, a priori, there is no reason why there should exist easy-to-compute hitting sets against such hard-to-compute polynomials.

2.6 Equivalence between strengthened black-box derandomization and lower bounds for infinitesimally close approximation

The following strengthening of Theorem 2.1 says that one can still compute efficiently in quasi-polynomial time a hitting set against such polynomials, assuming a sub-exponential lower bound for infinitesimally close approximation for a family $\{p_m\}$ of exponential-time-computable, multilinear, integral polynomials. A good candidate for p_m is the permanent. It cannot be approximated infinitesimally closely by small algebraic circuits as per the hardness hypothesis [71] of geometric complexity theory. The result below is the main reason why the strengthened black-box derandomization hypothesis is expected to hold.

Theorem 2.4 *Suppose there exists a family $\{p_m(x_1, \ldots, x_m)\}$ of exponential-time-computable, multi-linear, integral polynomials such that p_m cannot be approximated infinitesimally closely by circuits over K of $O(2^{m^\epsilon})$ size, for some constant $\epsilon > 0$, as $m \to \infty$. Then polynomial identity testing for small degree circuits over K with size $\leq s$ and $n \leq s$ variables has $O(2^{polylog(s)})$-time-computable strengthened black-box derandomization.*

This result also holds if we use symbolic determinants instead of circuits in the lower bound hypothesis; cf. the proof of Theorem 4.6.

Proof: We extend the proof of Theorem 7.7 in Kabanets and Impagliazzo [49] using Theorem 2.2, which lies at the heart of this proof.

We want to construct in quasi-poly(s) time a hitting set for strengthened black-box derandomization of polynomial identity testing for small degree circuits with size $\leq s$ and $n \leq s$ variables.

Let $m = (\log s)^e$, for a large enough constant e to be fixed later. Construct an NW-design [76] for this n (the number of variables) with this choice of m. By the NW-design, we mean a family of sets $R_1, \ldots, R_n \subseteq [l]$, $l \leq m^2 = (\log s)^{2e}$, each of cardinality m, such that $|R_i \cap R_j| \leq \log n$, for all $i \neq j$. By [76] (cf. Lemma 2.23 in [49]), such a set system can be constructed in $\text{poly}(n, 2^l) = O(2^{\text{polylog}(s)})$ time.

This set system and the given hard multi-linear polynomial $p(x_1, \ldots, x_m)$ together yield an arithmetic NW-generator NW^p. By this we mean the function

$$NW^p: \quad x = (x_1, \ldots, x_l) \in \mathbb{N}^l \to (p(x|_{R_1}), \ldots, p(x|_{R_n})) \in \mathbb{Z}^n, \tag{2}$$

where $x|_R$ denotes the tuple of the elements in x indexed by R.

Claim 2.5 *The set $H = \{NW^p(a) \mid a \in [D]^l\}$, $D = dm + 1$, is a hitting set against every nonzero polynomial $f(y)$, $y = (y_1, \ldots, y_n)$, of degree $\leq d = O(s^t)$, $t > 0$ a constant, that can be approximated infinitesimally closely by circuits over K and $y = (y_1, \ldots, y_n)$ of size $\leq s$ (assuming that the constant e above is chosen to be large enough).*

Since p is exponential-time-computable, H is $O(2^{\text{polylog}(s)})$-time computable. So it remains to prove the claim.

Proof of the claim: Suppose to the contrary that $f(b) = 0$, for every $b \in H$, for some nonzero polynomial $f(y)$ of degree $\leq d$ that can be approximated infinitesimally closely by circuits over K of size $\leq s$.

Let $g_0(x_1, \ldots, x_l, y_1, \ldots, y_n) := f(y_1, \ldots, y_n)$. For $1 \leq i \leq n$, let $g_i(x_1, \ldots, x_l, y_{i+1}, \ldots y_n)$ be the polynomial obtained from f by replacing $y_1, \ldots, y_i$ by the polynomials $p(x|_{R_j})$, $1 \leq j \leq i$. Then $g_n = f(NW^p(x))$, and the degree of each g_i is $\leq dm < D$. Since $f(b) = 0$ for all $b \in H$, $g_n(a) = 0$ for all $a \in [D]^l$. Since $\deg(g_n) < D$, by the Schwarz-Zippel lemma [84], g_n is identically zero. But $g_0 = f$ is not identically zero. So there exists a smallest $0 \leq i < n$ such that g_i is not identically zero but g_{i+1} is identically zero. Fix this i. Since g_i is not identically zero, we can set $y_{i+2}, \ldots, y_n$ and x_j, $j \notin R_{i+1}$, to some integer values so that the restricted polynomial $\tilde{g}_i(x_{j_1}, \ldots, x_{j_m}, y_{i+1})$ remains a non-zero polynomial, where $R_{i+1} = \{x_{j_1}, \ldots, x_{j_m}\}$. Let us denote this polynomial by renaming the variables as $g(x_1, \ldots, x_m, y)$.

Then $g(x_1, \ldots, x_m, y)$ is a non-zero polynomial with degree $\leq dm$, but $g(x_1, \ldots, x_m, p(x_1, \ldots, x_m))$ is identically zero. By Gauss's Lemma, $h(x_1, \ldots, x_m, y) = p(x_1, \ldots, x_m) - y$ is a factor of $g(x_1, \ldots, x_m, y)$. By Theorem 2.2, $h(x_1, \ldots, x_m, y)$ has a circuit over K of $\text{poly}(m, \deg(g)) = \text{poly}(s)$ size, with oracles gates for g. Setting $y = 0$ in this circuit, we get a circuit for $p(x_1, \ldots, x_m)$ of poly(s) size with oracle gates for g.

But g has a circuit of size $O(n^2 \log n)$ with one oracle gate for f. This is because $|R_j \cap R_{i+1}|$, $j \leq i$, is at most $\log n$, by the property of the NW-design. Hence, after the specialization of the variables $y_{i+2}, \ldots, y_n$ and x_j, $j \notin R_{i+1}$, as above, each $p(x|_{R_j})$, $j \leq i$, gets restricted to a multi-linear polynomial in at most $\log n$ variables. This restricted polynomial can be computed brute-force by a circuit C_j of size at most $O(\log n 2^{\log n}) = O(n \log n)$ size. We get a circuit for g, as desired, by connecting the inputs $y_1, \ldots, y_i$ of the oracle for f to the outputs of $C_1, \ldots, C_i$, respectively, and specializing the variables $y_{i+2}, \ldots, y_n$ to their integer values chosen above.

It follows that $p(x_1, \ldots, x_m)$ can be computed by a circuit C over K of size $O(s^c)$ with oracle gates for f, for some constant $c > 0$ independent of e. Given any circuit D_δ of size $\leq s$ for approximating f within precision $\delta > 0$, let C_δ denote the circuit obtained from C by substituting D_δ for f. Since f can be approximated infinitesimally closely by circuits of size $\leq s$, by choosing δ small enough, C_δ can approximate $p(x_1, \ldots, x_m)$ to any precision. The size of C_δ is $O(s^{c+1})$. Choosing e large enough, the size of C_δ can be made $\leq 2^{m^\epsilon}$ for any $\epsilon > 0$. This contradicts hardness of infinitesimally close approximation of p. Q.E.D.

If the polynomial p in Theorem 2.4 is the permanent, the following stronger result holds.

Theorem 2.6 *Suppose the permanent of $k \times k$ matrices cannot be approximated infinitesimally closely by circuits over K of $O(2^{k^\epsilon})$ size, for some constant $\epsilon > 0$, as $k \to \infty$. Then a hitting set for strengthened black-box derandomization of small-degree circuits over K of size $\leq s$ can be constructed in $O(polylog(s))$ parallel time using $O(2^{polylog(s)})$ processors.*

Proof: The proof is like that of Theorem 2.4, letting $m = k^2$ and $p_m(x) = \text{perm}(x)$, and thinking of $x = (x_1, \ldots, x_m)$ as a $k \times k$ matrix. We follow the same notation as in the proof of Theorem 2.4. We only need to explain why the construction of a hitting set can be efficiently parallelized.

The arithmetic NW-generator, cf. (2), based on the permanent is the function

$$NW^{\text{perm}}: \quad x = (x_1, \ldots, x_l) \in \mathbb{N}^l \to (\text{perm}(x|_{R_1}), \ldots, \text{perm}(x|_{R_n})) \in \mathbb{Z}^n, \qquad (3)$$

where $x|_R$ denotes the tuple of the elements in x indexed by R with cardinality $m = k^2$.

Since $n \leq s$, we can compute each $\text{perm}(x_{R_j})$ in parallel. Thus, it suffices to explain why each $\text{perm}(x_{R_j})$ can be computed fast in parallel. Fix one R_j. Without of loss generality, assume that the elements in R_j are $x = (x_1, \ldots, x_m)$. Think of x as a $k \times k$ matrix. Then $perm(x) = \sum_\sigma \prod_i x_{i\sigma(i)}$, where σ ranges over all permutations of k letters. Since $m = \text{polylog}(s)$, the number of terms in this expansion is $O(2^{\text{polylog}(s)})$. So we can assign a processor to each monomial in the expansion. The processor can compute that monomial in poly(m) = polylog(s) time.

It follows that NW^{perm} can be computed in $O(\text{polylog}(s))$ parallel time using $O(2^{\text{polylog}(s)})$ processors. Q.E.D.

Remark 1: The crucial fact used in the proof of Theorem 2.6 is that the permanent of $k \times k$ matrices can be computed in $O(\text{polylog}(s))$ parallel time using $O(2^{\text{polylog}(s)})$ processors, if $k = O(\text{polylog}(s))$. This need not hold, in general, for the exponential-time-computable p_m in the statement of Theorem 2.4.

Remark 2: Theorem 2.6 also holds, with a similar proof, if we replace the permanent in its statement by any PSPACE-computable, integral, multi-linear polynomial satisfying a similar lower bound assumption.

The following result is the easy converse of Theorem 2.4, ignoring the quasi-prefix.

Proposition 2.7 *Suppose the strengthened black-box derandomization hypothesis for polynomial identity testing for small degree circuits over K holds. Then there exists a family $\{p_m(x_1, \ldots, x_m)\}$ of exponential-time-computable, multi-linear, integral polynomials such that p_m cannot be approximated infinitesimally closely by circuits over K of $O(2^{m/a})$ size, for some constant $a > 0$, as $m \to \infty$.*

Proof: The proof is similar that of Theorem 51 in [1].

Choose $s = 2^{m/a}$, where $a > 0$ is a large enough constant to be chosen later. Suppose there exists an $O(s^b)$-time-computable, integral hitting set T of size $\leq s^b$ against all nonzero multi-linear polynomials in m variables that can be approximated infinitesimally closely by circuits over K of size $\leq s$.

Let $p_m(x)$, $x = (x_1, \ldots, x_m)$, be a multi-linear polynomial such that

$$p_m(t) = 0, \quad \forall t \in T. \tag{4}$$

Each condition here is a linear constraint on 2^m coefficients of $p_m(x)$. The number of these constraints is $|T| \leq s^b = 2^{mb/a} < 2^m$ if $a > b$. Hence, as $m \to \infty$, there is a non-zero integral $p_m(x)$ satisfying these constraints. One such $p_m(x)$ can be computed in $2^{O(m)}$ time by solving the linear system (4). By (4), this exponential-time computable $p_m(x)$ cannot be approximated infinitesimally closely by circuits over K of size $\leq s$, since T is a hitting set. Q.E.D.

By Theorem 2.4 and Proposition 2.7, strengthened black-box derandomization and sub-exponential lower bounds for infinitesimally close approximation of exponential-time-computable, multi-linear, integral polynomials are essentially equivalent notions.

2.7 The EXPSPACE-bound for strengthened black-box derandomization

The following is the currently best unconditional deterministic upper bound for strengthened black-box derandomization.

Theorem 2.8 *The strengthened black-box derandomization problem for general polynomial identity testing belongs to EXPSPACE. It belongs to EXPH (the exponential hierarchy) assuming the Generalized Riemann Hypothesis.*

In contrast:

Proposition 2.9 *The standard black-box derandomization problem for polynomial identity testing over K belongs to PSPACE unconditionally, and to PH assuming the Generalized Riemann Hypothesis.*

This proposition can be proved using Theorem 2.3 and the following result.

Theorem 2.10 (cf. Koiran [55]) *The problem* Hilbert's Nullstellensatz *of deciding if a given system of multi-variate integral polynomials, specified as circuits, has a complex solution is in PSPACE unconditionally, and in $AM \subseteq RP^{NP} \subseteq \Pi_2$ assuming the Generalized Riemann Hypothesis. The same also holds for the homogeneous variant of Hilbert's Nullstellensatz, namely, the problem of deciding if a given system of homogeneous, multi-variate, integral polynomials has a nontrivial complex solution.*

Proof of Theorem 2.8: For simplicity, we only prove the result for symbolic determinant identity testing. The proof for general polynomial identity testing is similar, using the universal circuit polynomial $H(Y)$ introduced in [71] (and recalled in Section 5.1.3 here) in place of the symbolic determinant.

We want to construct a hitting set against all non-zero homogeneous polynomials of degree m that can be approximated infinitesimally closely by symbolic determinants of size m on r variables. Without loss of generality, we can assume that $r = m^2$ (by adding more variables or increasing the size of the determinant). We can identify these m^2 variables with the entries of a variable $m \times m$ matrix X. Then all such nonzero polynomials correspond to the nonzero points of the variety $\hat{\Delta}[\det, m] \subseteq \mathcal{X}$ constructed in Section 1.2, since the closure in the Zariski topology coincides with the closure in the complex topology; cf. Theorem 2.33 in [74]. We now follow the same terminology as in Section 1.2.

A symbolic determinant of size m over m^2 variables can be computed [60] by a circuit of size $s = O(\text{poly}(m))$. Hence, by Theorem 2.3, there exists a subset of $[u]^{m^2}$, $u = 2s(m+1)^2$, of size $q = 6(s+1+m^2)^2$ that is a hitting set against all non-zero polynomials that can be approximated infinitesimally closely by symbolic determinants of size m over the m^2 variable entries of X.

We can enumerate all subsets of $[u]^{m^2}$ of size q, and for each enumerated subset $B \subseteq [u]^{m^2}$ of size q, check if it is a hitting set. The enumeration can be done using poly(m) work-space.

However, checking if a given $B \subseteq [u]^{m^2}$ of polynomial size q is a hitting set turns out to be much more difficult for the reasons that will be explained in more detail in Section 11. This is the main difficulty, since finally we have to output a correct B of polynomial size. This check can be done using exponential space as follows.

As in Section 1.2, for any $m \times m$ rational matrix b, let ψ_b be the homogeneous linear evaluation function on $\mathcal{X}$, which maps $p(X) \in \mathcal{X}$ to $p(b)$. Let $H(b)$ denote the hyperplane that is the zero set of ψ_b. Then B is a hitting set iff $\hat{\Delta}[\det, m] \cap \bigcap_{b \in B} H(b) = \{0\}$. To carry out this test, we first compute the defining equations of $\hat{\Delta}[\det, m] \subseteq \mathcal{X}$. Using Gröbner basis theory (cf. Theorem 1 in [61]), this can be done in work-space that is polynomial in the dimension of $\mathcal{X}$ and exponential in the dimension of $\Delta[\det, m]$. This work-space requirement is exponential in m. The total bit-length of the specification of the resulting defining equations of $\hat{\Delta}[\det, m]$ is at most exponential in the work-space requirement, and thus, at most double exponential in m. After this, we again use Gröbner basis theory (cf. Theorem 1 in [61]) to carry out the test. This takes work-space that is polynomial in the dimension of $\mathcal{X}$, exponential in the dimension of $\Delta[\det, m]$, and poly-logarithmic in the total bit-length of the specification of the defining equations. This space requirement is again exponential in m, i.e., $O(2^{\text{poly}(m)})$. Overall, this is an EXPSPACE-algorithm.

In the preceding proof, one can use Theorem 2.10 in place of Gröbner basis theory. Specifically, given $B \subseteq [u]^{m^2}$ of polynomial size q, one can construct in exponential time, using Theorem 5.7 in [13], a system of polynomial equations over $\mathbb{Q}$ in exponentially many variables with the specification in terms of circuits of exponential total bit-length, such that B is a hitting set iff this system does not have a complex solution. Using Theorem 2.10, the latter test can be carried out by an EXPSPACE-algorithm unconditionally, and by an EXPH-algorithm, assuming the Generalized Riemann Hypothesis. This is an EXPH-algorithm and not a PH-algorithm, since the number of variables in the system is exponential.

Assuming the Generalized Riemann Hypothesis, this gives an EXPH-algorithm for the verification of a hitting set, and hence, an EXPH-algorithm for strengthened black-box derandomization. Q.E.D.

3 Noether's Normalization Lemma

In this section we recall Noether's Normalization Lemma and show that the problem of constructing an h.s.o.p. for a general variety, given by the standard specification (defined below) in terms of its defining equations, belongs to PH.

Lemma 3.1 (Noether's Normalization Lemma) *(Cf. page 36 in [74]) Let $X \subseteq P(K^k)$ be a projective variety of dimension n. Let $\psi : K^k \to K^m$, $m \geq n+1$, be a homogeneous linear map that does not vanish on any line through the origin in K^k corresponding to any point of X. This means ψ induces a regular (well defined) linear map from X to $P(K^m)$, which we denote by ψ again. Then the homogeneous coordinate ring $R(X)$ of X is integral over the subring generated by the pullbacks $\psi^*(x_i)$'s of the coordinate functions x_i's, $1 \leq i \leq m$, on K^m. This implies that (1) $\psi(X) \subseteq P(K^m)$, the image of X, is closed in $P(K^m)$, and (2) the fiber $\psi^{-1}(p)$, for any point $p \in \psi(X)$, is a finite set.*

Conversely, if $R(X)$ is integral over the subring generated by $\psi^*(x_i)$'s, then ψ is regular on X.

Any ψ chosen uniformly at random has the regularity property stated above if $m \geq n+1$.

The following graded version of Noether's Normalization Lemma is implicit in its proof.

Lemma 3.2 (Graded Noether Normalization) *(cf. Theorem 13.3. in [27], Corollary 2.29 in [74], and also the proof of Theorem 1.5.17 in [10]) Let R be any positively graded, finitely generated K-algebra. Let $f_1, \ldots, f_t$ be any non-constant, homogeneous generators of R, and $H \subseteq R$ any set of homogeneous elements such that, letting $I(H)$ denote the ideal generated by H, $f_i^{e_i} \in I(H)$ for some positive integer e_i, for every i. Then R is integral over the subring generated by H.*

Definition 3.3 *[27] Let R be any positively graded, finitely generated K-algebra. A set H of homogeneous invariants of cardinality equal to* $\dim(R)$ *such that R is integral over the subring generated by H is called an* h.s.o.p. (homogeneous system of parameters) *of R.*

Thus $\psi^*(x_i)$'s, $1 \leq i \leq m$, in Lemma 3.1 form an h.s.o.p. of the homogeneous coordinate ring $R(X)$ of X, if $m = n + 1$.

Let $Z \subseteq K^t$ be a variety consisting of the common zeroes of a set of homogeneous integral polynomials $f_1(z), f_2(z), \ldots$, $z = (z_1, \ldots, z_t)$. Assume that Z is specified by giving circuits for f_i's, and that the constants in these circuits are rational. We call such a specification of Z *standard*. Its *bit-length* is defined to be the total bit-length of the specification of the circuits for f_i's.

The following result shows that for general varieties over K, given by the standard specification as above, the problem of constructing an h.s.o.p. is in PH assuming the Generalized Riemann Hypothesis. The succinct specification of $\Delta[\det, m]$ (cf. Section 1.2) in terms of a small circuit for computing the determinant is not standard, since it does not specify defining equations for the variety. Hence the following result does *not* apply to $\Delta[\det, m]$ given in the succinct specification. The current best EXPSPACE-bound for $\Delta[\det, m]$ given in the succinct specification will be proved later (cf. Theorem 4.1).

Theorem 3.4 *The problem of constructing an h.s.o.p. for a general variety over K, given by the standard specification in terms of circuits for the defining equations, belongs to PH assuming the Generalized Riemann Hypothesis, and to PSPACE unconditionally.*

Here by PH, we really mean its functional analogue, since the problem under consideration is a construction problem, not a decision problem. The PSPACE-bound holds in arbitrary characteristic.

Proof: Let $Z \subseteq K^t$ be a variety consisting of the common zeroes of a set of homogeneous integral polynomials $f_1(z), f_2(z), \ldots$, $z = (z_1, \ldots, z_t)$, specified in the standard fashion by the circuits for f_i's with rational constants. Let N be the total bit-length of this specification.

Testing if $\dim(Z) = 0$ is the complement of the homogeneous Hilbert's Nullstellensatz problem in Theorem 2.10. Hence, by Theorem 2.10, we can test if $\dim(Z) = 0$ by a PSPACE-algorithm, and also by a Σ_2-algorithm assuming the Generalized Riemann Hypothesis. If $\dim(Z) = 0$, then h.s.o.p. for Z is empty, and we are done.

So let us assume that $\dim(Z) \geq 1$.

Let $s \leq \dim(Z)$ be any positive integer. Consider random linear forms $L_r(z) = \sum_k b_{k,r} z_k$, $1 \leq r \leq s$, where $b_{k,r}$'s are random integers of large enough poly(N) bit-length. Let $H_r \subseteq K^t$ be the hyperplane defined by $L_r(z) = 0$.

We claim that, if $s = \dim(Z)$, then $Z \cap \bigcap_r H_r = \{0\}$ with a high probability. If $s < \dim(Z)$, then clearly $Z \cap \bigcap_r H_r \neq \{0\}$, since it has non-zero dimension.

By Hilbert's Nullstellensatz and Lemma 3.2, it the follows that, if $s = \dim(Z)$, then the homogeneous coordinate ring of Z is integral over the subring generated by $L_r(z)$'s, and hence $\{L_r(Z)\}$ is an h.s.o.p. for Z.

So, let us first prove the claim. Accordingly, assume that $s = \dim(Z)$. Let $d = \max\{\deg(f_i)\}$. Clearly, $d \leq 2^M$, where M denotes the maximum number of multiplication gates in the circuit for any f_i. Since $M \leq N$, it follows that $d \leq 2^N$. By raising f_i's to appropriate powers, we can assume that all of them have the same degree $D \leq 2^{N^2}$. Consider generic linear combinations of f_i's and generic linear forms

$$\begin{array}{ll} F_j(z) = \sum_i y_{i,j} f_i(z), & 1 \leq j \leq t - \dim(Z), \\ L_r(z) = \sum_k w_{k,r} z_k, & 1 \leq r \leq s = \dim(Z), \end{array} \tag{5}$$

where $y_{i,j}$'s and $w_{k,r}$'s are indeterminates. Let R denote the multi-variate resultant of F_j's and L_r's. It is a polynomial in $y_{i,j}$'s and $w_{k,r}$'s of degree $\leq D^t$. By Noether's Normalization Lemma (cf. Lemma 3.1 and the remark thereafter), the system of equations (5) has only $\{0\}$ as its solution for some rational values for $y_{i,j}$'s and $w_{k,r}$'s. Hence R is not identically zero as a polynomial in $y_{i,j}$'s and $w_{k,r}$'s. By the Schwarz-Zippel lemma [84], we can specialize $y_{i,j}$'s randomly to some integers of $O(\log(D^t)) = \text{poly}(N)$ bit-length so that the resulting specialization R' of R is not identically zero. Then R' is a nonzero polynomial in $w_{k,r}$'s of degree $\leq D^t$. By the Schwarz-Zippel lemma again, R' does not vanish identically if we let $w_{k,r} = b_{k,r}$ for randomly chosen integers of $O(\log(D^t)) = \text{poly}(N)$ bit-length. For such $b_{k,r}$'s, $Z \cap \bigcap_r H_r = \{0\}$. This proves the claim.

Next, we show that $\dim(Z)$ and a specification of $L_r(z)$'s, $1 \leq r \leq \dim(Z)$, such that $Z \cap \bigcap_r H_r = \{0\}$ can be computed in poly(N) work-space.

We begin by letting $s = 1$, the first guess for $\dim(Z)$. With this choice of s, choose $b_{k,r}$'s as above randomly of large enough poly(N) bit-length and test if $Z \cap \bigcap_r H_r = \{0\}$. The latter test can be carried out in PSPACE unconditionally (cf. Theorem 2.10). If the test fails, we increase s by one and repeat the test. The test succeeds with a high probability when $Z \cap \bigcap_r H_r = \{0\}$ and $s = \dim(Z)$. Randomization in this algorithm can be removed, since RPSPACE = NPSPACE = PSPACE. This yields a PSPACE-algorithm for computing $\dim(Z)$ and $L_r(z)$'s, $1 \leq r \leq \dim(Z)$, such that $Z \cap \bigcap_r H_r = \{0\}$, as desired.

Assuming the Generalized Riemann Hypothesis, whether $Z \cap \bigcap_r H_r = \{0\}$ can be tested by a Σ_2-algorithm, by Theorem 2.10. This gives a Σ_2-algorithm for testing if there exist $L_r(z)$'s, for the given choice of s, such that $Z \cap \bigcap_r H_r = \{0\}$: guess $y_{i,j}$'s and $w_{k,r}$'s, and test if $Z \cap \bigcap_r H_r = \{0\}$ using the Σ_2-algorithm (Theorem 2.10).

Using this Σ_2-algorithm for testing the existence of $L_r(z)$'s in place of the PSPACE-algorithm before, we get a PH-algorithm for computing $\dim(Z)$ and $L_r(z)$'s, $1 \leq r \leq \dim(Z)$, such that $Z \cap \bigcap_r H_r = \{0\}$. Q.E.D.

The proof of Theorem 3.4 shows that the problem of constructing an h.s.o.p. for general varieties specified by their equations is Turing-reducible in randomized polynomial time to the complement of the homogeneous Hilbert's Nullstellensatz problem in Theorem 2.10. Conversely, the complement of the homogeneous Hilbert's Nullstellensatz problem can be reduced to the problem of constructing an h.s.o.p. for general varieties (since a projective variety X is empty iff its h.s.o.p. is empty). Since the Hilbert's Nullstellensatz problem in Theorem 2.10 is NP-hard [36], it follows that the problem of constructing an h.s.o.p. for general varieties is co-NP-hard. In analogy with the problem NNL for $\Delta[\det, m]$ in Section 1.2, we can define the problem NNL for general varieties X, specified in the standard fashion by their equations, as the problem of constructing a small homogeneous set $S \subseteq R(X)$ of cardinality polynomial in the dimension of X (but not necessarily of optimal cardinality equal to $\dim(X) + 1$) such that $R(X)$ is integral over the subring generated by S. Even this problem is co-NP-hard.

In contrast, we shall prove in the next two sections that the problem NNL of constructing an s.s.o.p. for $\Delta[\det, m]$, with a succinct specification, and more generally, the problem NNL for any explicit variety can be solved in quasi-polynomial time, assuming a lower bound for infinitesimally close approximation.

4 NNL for $\Delta[\det, m]$

In this section we prove Theorems 1.1, 1.7, and 1.9. We follow the same notation as in Section 1.2.

The variety $\Delta[\det, m]$, defined in Section 1.2, can alternatively be defined as follows. Let X be a variable $m \times m$ matrix. Let $\mathcal{X}$ be the vector space over K of homogeneous polynomials of degree m in the variable entries of X, and $P(\mathcal{X})$ the projective space associated with $\mathcal{X}$. Thus $g = \det(X)$ is an element of $\mathcal{X}$. Furthermore, $\mathcal{X}$ is a representation of $G = GL_{m^2}(K)$, where $\sigma \in GL_{m^2}(K)$ maps $h(X) \in \mathcal{X}$ to $h(\sigma^{-1}X)$, thinking of X as an m^2-vector. Then $\Delta[\det, m] \subseteq P(\mathcal{X})$ is the Zariski-closure of the orbit $Gg \subseteq P(\mathcal{X})$, thinking of g as also a point in $P(\mathcal{X})$. (We can also use $SL_{m^2}(K)$ here instead of $GL_{m^2}(K)$, since that does not change $\Delta[\det, m]$.) As in Section 1.2, let $\hat{\Delta}[\det, m] \subseteq \mathcal{X}$ be the affine cone of $\Delta[\det, m]$, and $R(\det, m)$ the homogeneous coordinate ring of $\Delta[\det, m]$.

We assume that $\Delta[\det, m]$ is specified succinctly as in Section 1.2. This can be done either by giving a small uniform circuit of poly(m) bit-length for computing $\det(X)$, or alternatively, by just giving m in unary (from which a circuit for the determinant can be computed in poly(m) time). The bit-length of this succinct specification is poly(m). All complexity bounds in this section will be in terms of this bit-length, or equivalently, in terms of m.

The problem NNL for $\Delta[\det, m]$, given in this succinct specification, is to construct an s.s.o.p. of the form $S(\mathcal{B})$, as defined in Section 1.2, for some set $\mathcal{B}$ of $m \times m$ rational matrices of poly(m) total bit-length.

Remark: Later (cf. Definition 5.6) we define a more general s.s.o.p., which need not be of the form $S(\mathcal{B})$. But s.s.o.p.'s of this form are most natural. They are called *strict s.s.o.p.* in Definition 5.7. In this section, we assume, as in Section 1.2, that an s.s.o.p. for $\Delta[\det, m]$ is always of the form $S(\mathcal{B})$. Thus NNL here is strict NNL as per the terminology in Section 5.3.

We call an s.s.o.p. $S(\mathcal{B})$ *separating* if, for any two distinct points $p, q \in \hat{\Delta}[\det, m]$, $\psi_{\mathcal{B}}(p) \neq \psi_{\mathcal{B}}(q)$, with $\psi_{\mathcal{B}}$ as in Section 1.2. Thus a separating s.s.o.p. denotes a *dimension-reducing map*, with a succinct specification, from $\mathcal{X}$ to K^k, $k = \text{poly}(m)$, that is injective on $\hat{\Delta}[\det, m]$. By the *strong form of NNL* for $\Delta[\det, m]$, we mean the problem of constructing a separating s.s.o.p. for $\Delta[\det, m]$. A poly(m)-time-constructible, separating s.s.o.p. is called a *separating* *e.s.o.p. (explicit system of parameters)*. Separating quasi-s.s.o.p. and quasi-e.s.o.p. are defined by replacing $\text{poly}(m)$ by $2^{\text{polylog}(m)}$.

4.1 Unconditional upper bound for the problem of constructing an h.s.o.p.

Before we turn to the construction of an s.s.o.p., we address the construction of an h.s.o.p. for $\Delta[\det, m]$ (cf. Section 1.2). By Theorem 3.4, the problem of constructing an h.s.o.p. for a general variety, given by the standard specification in terms of defining equations, is in PSPACE. This does *not* imply that the same problem for $\Delta[\det, m]$ is in PSPACE, since $\Delta[\det, m]$ is not specified in the standard fashion by its defining equations, but rather succinctly by a small uniform circuit for the determinant. The current best algorithm based on Gröbner basis theory [61] for converting the succinct specification of $\Delta[\det, m]$ to its standard specification takes space that is exponential in m. Hence, we only get the following EXPSPACE-bound for the succinct specification.

Theorem 4.1 *The problem of constructing an h.s.o.p. for* $\Delta[\det, m]$*, specified succinctly, belongs to EXPSPACE. (This means it can be solved in work-space that is exponential in m). Assuming the Generalized Riemann Hypothesis, it belongs to EXPH, if* $\Delta[\det, m] \subseteq P(\mathcal{X})$ *has defining equations that can be computed in time that is exponential in m.*

Proof: Given the succinct specification of $\Delta[\det, m]$, we first compute the equations defining it as a subvariety of $P(\mathcal{X})$, using Gröbner basis theory as in the proof of Theorem 2.8, in work-space that is exponential in m. The total degree and the bit-length of the specification of these equations is at most double-exponential in m.

We apply Gröbner basis theory (cf. Theorem 1 in [61]) again to compute an h.s.o.p. for $\Delta[\det, m]$, using these defining equations. This takes work-space that is polynomial in $\dim(\mathcal{X})$, exponential in $\dim(\Delta[\det, m])$, and poly-logarithmic in the total bit-length of the specification of the defining equations. This work-space requirement is single-exponential in m, i.e., $O(2^{\text{poly}(m)})$. The total running time as well as the bit-length of the output h.s.o.p. is double exponential in m. This gives an EXPSPACE algorithm for computing an h.s.o.p. for $\Delta[\det, m]$.

If $\Delta[\det, m]$ has defining equations that can be computed in exponential time, then we can skip the first step above of computing defining equations, and use these equations instead. After this, we can use the PH-algorithm for general varieties in Theorem 3.4 for computing an h.s.o.p., assuming the Generalized Riemann Hypothesis. Since the dimension of $\mathcal{X}$ is exponential in m, this PH-algorithm becomes an EXPH-algorithm in our context. Q.E.D.

The bit-length of the specification of the h.s.o.p. constructed in Theorem 4.1 is double-exponential in m. If we insist on an h.s.o.p. then Theorem 4.1 is the best that we can do at present. However, if we are willing to settle for an s.s.o.p. (which need not have the optimal cardinality like an h.s.o.p.), then Theorem 1.9, proved in this section (cf. Theorem 4.9), says that the double exponential time bound in Theorem 4.1 can be brought down to quasi-polynomial, assuming that there exists a family $\{f_n(x_1, \ldots, x_n)\}$ of exponential-time-computable, integral, multi-linear polynomials such that f_n cannot be approximated infinitesimally closely by symbolic determinants over K of sub-exponential size.

4.2 A Monte Carlo algorithm

We begin by proving the following stronger form of Theorem 1.1.

Theorem 4.2 *A separating s.s.o.p. for* $\Delta[\det, m]$ *can be constructed by a poly*(m)*-time Monte-Carlo algorithm that is correct with a high probability.*

Proof: Since the determinant can be specified by a circuit with rational constants, it follows from Gröbner basis theory [61] that $\Delta[\det, m]$ has defining equations with rational coefficients. For any rational $m \times m$ matrix B, let H_B denote the set of the zeroes of the associated homogeneous linear map ψ_B (cf. Section 1.2). Given any set $\mathcal{B} = \{B_1, \ldots, B_k\}$ of $m \times m$ rational matrices, the associated homogeneous linear map $\psi_{\mathcal{B}}$ (cf. Section 1.2) does not vanish on any non-zero point in $\hat{\Delta}[\det, m]$ iff the variety $\hat{\Delta}[\det, m] \cap \bigcap_{B \in \mathcal{B}} H_B$ does not have a non-trivial solution over K. Since $\Delta[\det, m]$ has defining equations with rational coefficients, this variety also has defining equations with rational coefficients. By Hilbert's Nullstellensatz, this variety does not have a non-trivial solution over any algebraically closed field of characteristic zero iff it does not have a non-trivial solution over $\mathbb{C}$. So, without loss of generality, we can assume that $K = \mathbb{C}$.

As explained in the beginning of this section, $\Delta[\det, m] \subseteq P(\mathcal{X})$ is the Zariski-closure of the $GL_{m^2}(\mathbb{C})$-orbit of $\det(X)$, where X is an $m \times m$ variable matrix, and $\mathcal{X}$ is the vector space over $\mathbb{C}$ of homogeneous polynomials of degree m in the variable entries of X.

Since the Zariski-closure coincides with the closure in the complex topology (cf. Theorem 2.33 in Mumford [74]), it follows that $\Delta[\det, m]$ is the closure of the $GL_{m^2}(\mathbb{C})$-orbit of $\det(X)$ in the complex topology on $P(\mathcal{X})$. In concrete terms, this means that every point in the affine cone $\hat{\Delta}[\det, m]$ of $\Delta[\det, m]$ can be approximated infinitesimally closely by symbolic determinants of size m over the m^2 variables entries of X.

Since the determinant has a small circuit, it now follows from Heintz and Schnorr (Theorem 2.3) that one can compute by a poly(m)-time Monte Carlo algorithm a hitting set $\mathcal{B} = \{B_1, \ldots, B_k\}$, $k = \text{poly}(m)$, of integral $m \times m$ matrices, with poly(m) total bit-length, such that, with a high probability, (1) for every non-zero polynomial $p(X) \in \hat{\Delta}[\det, m]$, there exists a matrix $B_i \in \mathcal{B}$ such that $p(B_i)$ is not zero, and (2) more generally, given any two distinct polynomials $p_1(X), p_2(X) \in \hat{\Delta}[\det, m]$, there exists a matrix $B_i \in \mathcal{B}$ such that $p_1(B_i) \neq p_2(B_i)$.

We assume that $\mathcal{B}$ constructed above is a hitting set with this property. Let $S(\mathcal{B}) = \{\psi_{B_i}\}$ be the associated subset of the homogeneous coordinate ring of $\Delta[\det, m]$, as defined in Section 1.2.

Claim 4.3 *The set* $S(\mathcal{B})$ *is a separating s.s.o.p. for* $\Delta[\det, m]$.

Theorem 4.2 follows from the claim.

Proof of the claim: First, we prove that the set $S(\mathcal{B})$ is an s.s.o.p. (as defined in Section 1.2) for $\Delta[\det, m]$.

Let $\psi_{\mathcal{B}} : \mathcal{X} \to K^k$ be the homogeneous linear map associated with $\mathcal{B}$ as in Section 1.2. The total bit-length of B_i's is poly(m). So it suffices to show that $\psi_{\mathcal{B}}$ does not vanish on any non-zero point in $\hat{\Delta}[\det, m]$. By Noether's Normalization Lemma (Lemma 3.1), it then follows that the homogeneous coordinate ring of $\Delta[\det, m]$ is integral over the subring generated by $S(\mathcal{B})$.

Suppose to the contrary that $\psi_{\mathcal{B}}$ does vanish on some non-zero polynomial $p = p(X) \in \hat{\Delta}[\det, m]$. Then $p(B_i) = 0$ for all $i \leq k$. Since $p(X)$ can be approximated infinitesimally closely by symbolic determinants over X of size m and $\mathcal{B}$ is a hitting set, this implies that $p(X)$ is identically zero; a contradiction.

It remains to show that $S(\mathcal{B})$ is separating. Consider any two distinct polynomials $p_1(X), p_2(X) \in \hat{\Delta}[\det, m]$. By our assumption about the hitting set $\mathcal{B}$, $p_1(B_i) \neq p_2(B_i)$ for some i. This means $\psi_{\mathcal{B}}(p_1) \neq \psi_{\mathcal{B}}(p_2)$. It follows that $S(\mathcal{B})$ is separating. Q.E.D.

Corollary 4.4 *A separating s.s.o.p. for* $\Delta[\det, m]$ *exists.*

Proof: This follows from Theorem 4.2. Q.E.D.

4.3 Reduction of NNL to strengthened black-box derandomization

The Monte Carlo algorithm in Theorem 4.2 can be derandomized assuming a suitable derandomization hypothesis.

Theorem 4.5 *The variety* $\Delta[\det, m]$ *has a separating e.s.o.p., assuming the strengthened black-box derandomization hypothesis for symbolic determinant identity testing.*

Proof: By the strengthened black-box derandomization hypothesis for symbolic determinant identity testing, we can compute in poly(m) time a hitting set $\mathcal{B} \subseteq M_m(\mathbb{Z})$ against all nonzero polynomials of degree m that can be approximated infinitesimally closely by symbolic determinants of size m over the entries of X, an $m \times m$ variable matrix. More generally, we can also assume that, given any two distinct polynomials $p_1(X)$ and $p_2(X)$ that can be approximated infinitesimally closely by symbolic determinants of size m over X, there exists $b \in \mathcal{B}$ such that $p_1(b) \neq p_2(b)$. It then follows, as in the proof of Theorem 4.2, that the associated set $S(\mathcal{B})$ is a separating s.s.o.p. for $\Delta[\det, m]$. Q.E.D.

4.4 Reduction of NNL to a lower bound hypothesis

The strengthened black-box derandomization hypothesis in Theorem 4.5 can be traded with a lower bound hypothesis as in the following result.

Theorem 4.6 *The variety* $\Delta[\det, m]$ *has a separating quasi-e.s.o.p., assuming that there exists a family* $\{f_n(x_1, \ldots, x_n)\}$ *of exponential-time-computable, integral, multi-linear polynomials such that* f_n *cannot be approximated infinitesimally closely by circuits over* K *of* $O(2^{n^\epsilon})$ *size, for some constant* $\epsilon > 0$*, as* $n \to \infty$*. Alternatively, we can assume that* f_n *cannot be approximated infinitesimally closely by symbolic determinants over* K *of* $O(2^{n^{\epsilon'}})$ *size, for some constant* $\epsilon' > 0$*, as* $n \to \infty$*.*

Proof: The first statement follows from the proof of Theorem 4.5 and Theorem 2.4, since symbolic determinant identity testing is a special case of low-degree polynomial identity testing. By [93], any circuit over K of degree d and size s can be simulated by a circuit over K of $O(\log d(logd +$

$\log s)$) depth, and hence [91], by a symbolic determinant over K of $O(2^{O(\log d(\log d+\log s)})$ size. The second statement follows from the first statement, in conjunction with this fact, letting $d = n$, $s = 2^{n^{\epsilon}}$, and $\epsilon' = 2\epsilon$. Q.E.D.

Assuming a lower bound for the permanent, we get the following stronger result. This proves a stronger form of Theorem 1.7.

Theorem 4.7 *A separating s.s.o.p. for $\Delta[\det, m]$ can be constructed in $O(polylog(m))$ parallel time using $O(2^{polylog(m)})$ processors, assuming that the permanent of $n \times n$ matrices cannot be approximated infinitesimally closely by symbolic determinants over K of $O(2^{n^{\epsilon}})$ size, for some constant $\epsilon > 0$, as $n \to \infty$.*

Proof: This follows from Theorem 4.5 and Theorem 2.6, since low-degree algebraic circuits of sub-exponential size are equivalent to symbolic determinants of sub-exponential size; cf. the proof of Theorem 4.6. Q.E.D.

Define the variety $\Delta[\text{perm}, n, m]$, just as we defined $\Delta[\det, m]$ at the beginning of this section, replacing $\det(X)$ by $z^{m-n}\text{perm}(Y)$, where Y is some $n \times n$ sub-matrix of X, and z is a variable entry in X outside Y. Then the lower bound assumption in Theorem 4.7 in the terminology of [71] is that $\Delta[\text{perm}, n, m] \not\subseteq \Delta[\det, m]$, if $m = O(2^{n^{\epsilon}})$, for some small enough constant $\epsilon > 0$. This is a stronger form of Conjecture 4.3 in [71]. If we assume instead (as in Conjecture 4.3 in [71]) that $\Delta[\text{perm}, n, m] \not\subseteq \Delta[\det, m]$, if $m = O(\text{poly}(n))$, then it can be proved similarly that NNL for $\Delta[\det, m]$ can be solved in $O(2^{n^{\epsilon}})$-time (after replacing poly(n) by $2^{n^{\epsilon}}$ in the definition of an s.s.o.p.), for every constant $\epsilon > 0$.

4.4.1 Equivalence

The following result implies the easy converse to Theorem 4.5.

Lemma 4.8 *Suppose $\mathcal{B} \subseteq M_m(\mathbb{Z})$ specifies an s.s.o.p. $S(\mathcal{B})$ for $\Delta[\det, m]$. Then $\mathcal{B}$ is a hitting set against all non-zero polynomials of degree m that can be approximated infinitesimally closely by symbolic determinants of size m over the m^2 entries of X, an $m \times m$ variable matrix. Hence, existence of an e.s.o.p. for $\Delta[\det, m]$ implies the strengthened black-box derandomization hypothesis for symbolic determinant identity testing.*

Proof: By the definition of $\Delta[\det, m]$, every non-zero polynomial $p(X)$ of degree m that can be approximated infinitesimally closely by symbolic determinants over X of size m corresponds to a non-zero point in $\hat{\Delta}[\det, m]$. Since $S(\mathcal{B})$ is an s.s.o.p., it follows that $\psi_{\mathcal{B}}$ (as defined in Section 1.2) does not vanish on any non-zero point in $\hat{\Delta}[\det, m]$. This implies that $\mathcal{B}$ is a hitting set against all non-zero polynomials of degree m that can be approximated infinitesimally closely by symbolic determinants of size m over the m^2 entries of X.

In symbolic determinant identity testing, we can assume, without loss of generality, that the number of variables is at most quadratic in the size of the matrix, increasing the size otherwise. Hence the last statement follows. Q.E.D.

The following result proves Theorem 1.9.

Theorem 4.9 *(a)The strengthened black box derandomization hypothesis for symbolic determinant identity testing holds iff* $\Delta[\det, m]$ *has an e.s.o.p.*

(b) A sub-exponential lower bound for a family of exponential-time-computable, integral, multilinear polynomials as in Theorem 4.6 holds iff, ignoring quasi-prefixes, $\Delta[\det, m]$ *has an e.s.o.p.*

Proof: (a) This follows from Theorem 4.5 and Lemma 4.8.

(b) This follows from (a), Theorem 2.4, and Proposition 2.7. Q.E.D.

4.5 The current best unconditional deterministic upper bound for NNL

The following result gives the current best unconditional deterministic bound for NNL for $\Delta[\det, m]$. It does not follow from Theorem 4.1, since an h.s.o.p. constructed there need not have a succinct specification of poly(m) bit-length.

Theorem 4.10 *The problem of constructing or verifying an s.s.o.p. for* $\Delta[\det, m]$ *belongs to EXPSPACE unconditionally. It belongs to EXPH assuming the Generalized Riemann Hypothesis.*

Proof: The statement for construction follows from Theorem 2.8 and Theorem 4.9 (a). The proof for verification is implicit in the proof for construction. Q.E.D.

5 Explicit algebraic varieties

In this section we formulate a general notion of an explicit algebraic variety, motivated by the concrete example of $\Delta[\det, m]$ studied in the preceding section, and define the problem NNL in this context (cf. Section 5.3). We then generalize the results for $\Delta[\det, m]$ in the preceding section systematically to general explicit varieties.

Definition 5.1 *(a) A family* $\{W_n\}$, $n \to \infty$, *of affine varieties is called* explicit *if there exist families of positive integers* $\{r_n\}$, $\{m_n\}$, *a family* $\{\psi_n\}$ *of maps* $\psi_n : K^{r_n} \to K^{m_n}$*:*

$$v = (v_1, \ldots, v_{r_n}) \to (f_1(v), \ldots, f_{m_n}(v)), \tag{6}$$

with $r_n = poly(n)$, $m_n = n^{\Omega(1)}$, $\log m_n = O(poly(n))$, *and each* f_j *a homogeneous polynomial of poly(n) degree, and there also exist homogeneous polynomials* $g_j(x)$, $x = (x_1, \ldots, x_n)$, $1 \leq j \leq m_n$, *of poly(n) degree, such that:*

1. W_n *is the Zariski-closure of the image* $Im(\psi_n)$ *of* ψ_n*. This means* $W_n \cong spec(R)$, *where* R *is the subring of* $K[v_1, \ldots, v_{r_n}]$ *generated by* $f_1(v), \ldots, f_{m_n}(v)$.

2. *The polynomial* $F_n(v, x) = \sum_j f_j(v) g_j(x)$ *is uniformly p-computable [91]. This means one can compute in poly(n) time a circuit* C_n*, with rational constants, over the variables* $v = (v_1, \ldots, v_{r_n})$ *and* $x = (x_1, \ldots, x_n)$ *of poly(n) total bit-size, including the bit-sizes of the constants, that computes* $F_n(v, x)$*, and the total degree* $\deg(F_n)$ *of* F_n *is poly(n).*

3. The polynomials $g_j(x)$'s are linearly independent.

We call ψ_n the map defining W_n, and F_n the polynomial defining W_n. We specify W_n succinctly by the circuit C_n. Alternatively, we can specify W_n by the circuits $C_{n,c}$'s, $1 \leq c \leq \deg(F_n)$, where $C_{n,c}$ computes the degree c-component in v of F_n. The total bit-length of this succinct specification of W_n is poly(n).

We say that $\{W_n\}$ is strongly explicit *if the circuit C_n is weakly skew (cf. Section 2.1).*

(b) A family of projective varieties is called explicit (strongly explicit) if the family of the affine cones of these varieties is explicit (respectively, strongly explicit).

(c) An explicit family of affine or projective varieties without degree restrictions *is defined just as in (a) and (b), but without putting any restriction on the degrees of f_j, g_j, and F_n.*

(d) Quasi-explicit families are defined by replacing poly(n) by $2^{polylog(n)}$.

We denote the coordinate ring of W_n by $K[W_n]$. If $\{W_n\}$ is explicit, by abuse of terminology, we also say that the variety W_n is explicit.

5.1 Examples

We now give a few examples of explicit varieties.

5.1.1 The orbit closure of the determinant

The orbit-closure $\Delta[\det, m] \subseteq P(\mathcal{X})$ studied in Section 4 is explicit. Specifically, following the same notation as in Section 4, the affine cone $\hat{\Delta}[\det, m]$ of $\Delta[\det, m]$ is explicit, with the defining map $\phi : M_{m^2}(K) \to \mathcal{X}$ that maps $v \in M_{m^2}(K)$ to $\det(vX)$, thinking of X as an m^2-vector. The polynomial $F = F(v, X)$ defining $\Delta[\det, m]$ is $\det(vX)$. The monomials in the entries of X of degree m play the role of g_j's in Definition 5.1, and f_j's are the coefficients of $\det(vX)$ considered as a polynomial in X.

5.1.2 Explicit varieties associated with depth three circuits

Let S_n^d be the space of homogeneous forms in n variables of degree d, and $P(S_n^d)$ the associated projective space. Let $Y(d, k, n) \subseteq P(S_n^d)$ be the projective closure of the set of polynomials that can be expressed as sum of k terms, each term a d-th power of a linear form in the n variables. It is the variety associated with the class of diagonal depth three circuits (cf. Section 2.1) on n variables with degree d and top-fan-in k, and is known in algebraic geometry as the k-th secant variety of the Veronese variety [58]. It is explicit, the defining polynomial being the polynomial computed by the generic, homogeneous, diagonal depth three circuit (with indeterminate constants) on n variables with degree d and top fan-in k. Specifically, this defining polynomial is $\sum_{i=1}^{k}(\sum_{j=1}^{n} y_{i,j}x_j)^d$, where x_j's are the variables in the circuit, and $y_{i,j}$'s are the indeterminate constants.

Let $X(d, k, n) \subseteq P(S_n^d)$ be the projective closure of the set of polynomials that can be expressed as sum of k terms, each term a product of d linear forms in the n variables. It is the variety associated with the class of depth three circuits on n variables with degree d and top-fan-in k, and is known in algebraic geometry as the k-th secant variety of the Chow variety [58]. It is explicit, the defining polynomial being the polynomial computed by the generic, homogeneous, depth three circuit (with indeterminate constants) on n variables with degree d and top fan-in k. Specifically, this defining polynomial is $\sum_{i=1}^{k}\prod_{r=1}^{d}(\sum_{j=1}^{n} y_{i,r,j}x_j)$, where x_j's are the variables in the circuit, and $y_{i,r,j}$'s are the indeterminate constants.

5.1.3 The explicit variety associated with the universal circuit

Following [71], we now define an explicit variety without any degree restrictions, which plays the same role in the study of general polynomial identity testing that $\Delta[\mathrm{det}, m]$ plays in the study of symbolic determinant identity testing.

First, we define a universal circuit over K of depth k and width m. Let S_i, $0 \leq i \leq k$, denote the set of nodes in this circuit with level i. We assume that S_0 contains just one node, called the *root*, and for all $i > 0$, $|S_i| = m$. For all levels $0 \leq i \leq k-1$, we introduce indeterminates $y^u_{v,w}$'s for each $u \in S_i$ and distinct $v, w \in S_{i+1}$. For the k-th level, we introduce indeterminates y^u's, $u \in S_k$. Let Y be the tuple of all these indeterminates together. Beginning at the level k, for each element u in S_i, we recursively define the form $h(u)$ in the indeterminates Y as follows. For $u \in S_k$, let $h(u) = y^u$. For $u \in S_i$, with $i < k$, let $h(u) = \sum_{v,w} y^u_{v,w} h(v)h(w)$, where the sum ranges over all distinct $v, w \in S_{i+1}$. The form $H(Y) = H_{k,m}(Y)$ computed by this universal circuit is the form $h(u)$, where $u \in S_0$ is the root.

Any circuit over K of size s can can be obtained by specializing this universal circuit with $k = O(s)$ and $m = O(s)$; cf. [71].

Let $\mathcal{X}$ be the space of homogeneous forms in Y of total degree $d := \deg(H(Y))$ over the field K. Let l denote the number of variables in Y. Then $\mathcal{X}$ has a natural action of $G = GL_l(K)$, similar to the action in Section 1.2. Let $\Delta[H(Y), k, m] \subseteq P(\mathcal{X})$ denote the closure of the G-orbit of $H(Y)$ in $P(\mathcal{X})$. This is an explicit variety without any degree restrictions (cf. Definition 5.1 (c)).

We also define an explicit variety (with the usual low-degree restrictions), which plays the same role in the study of low-degree polynomial identity testing that $\Delta[\mathrm{det}, m]$ plays in the study of symbolic determinant identity testing.

Given any positive integer c, let $H(Y)_c$ denote the homogeneous degree c part of $H(Y)$. If $c = \mathrm{poly}(k, m)$, then $H(Y)_c$ can be computed by a circuit of $\mathrm{poly}(k, m)$ size, and furthermore, the family $\{H(Y)_m\}$, with $k = c = m$, is VP-complete; cf. Section 5.6 in [12]. Now let $\mathcal{X}$ be the space of homogeneous forms in Y of total degree m over the field K. Define $\Delta[H(Y)_m, k, m]$ just as we defined $\Delta[H(Y), k, m]$ above, with $H(Y)_m$ in place of $H(Y)$. This is an explicit variety, with the usual low-degree restrictions (cf. Definition 5.1).

5.1.4 The categorical quotients

Let V be a representation of $G = SL_m(K)$ of dimension n. Then the invariant ring $K[V]^G$ is finitely generated [43]. So we can consider the variety $V/G = \mathrm{spec}(K[V]^G)$, called the *categorical quotient* [75]. It can be constructed concretely as follows.

Fix any set $F = \{f_1, \dots, f_t\}$ of non-constant homogeneous generators of $K[V]^G$. Consider the morphism $\pi_{V/G}$ from V to K^t given by

$$\pi_{V/G}: \quad v \to (f_1(v), \dots, f_t(v)). \tag{7}$$

Then V/G can be identified with the closure of the image of this morphism. As we shall see below, this image is already closed (cf. Theorem 5.4). Let $z = (z_1, \dots, z_t)$ be the coordinates of K^t, I the ideal of V/G under this embedding, and $K[V/G]$ its coordinate ring. Then $K[V/G] = K[z]/I$, and we have the comorphism $\pi^*_{V/G} : K[V/G] \to K[V]$ given by

$$\pi^*_{V/G}(z_i) = f_i. \tag{8}$$

Since f_i's are homogeneous, $K[V/G]$ is a graded ring, with the grading given by $\deg(z_i) = \deg(f_i)$. Furthermore, $\pi^*_{V/G}$ gives an isomorphism between $K[V/G]$ and $K[V]^G$. Thus, we have $\pi^*_{V/G}(K[V/G]) = K[V]^G$.

The general definition of explicit varieties (Definition 5.1) specializes, when applied to the map $\pi_{V/G}$ in (7), to the following definition. Let $v_1, \dots, v_n$ denote the standard monomial basis [57] of V as in Section 1.5.

Definition 5.2 *(a) The categorical quotient V/G is called* explicit *if, given the specification $\langle V, G\rangle$ of V and G as in Section 1.5, one can compute in poly(n,m) time a set of circuits $C = C[V,m,c]$'s, $1 \leq c \leq q =$ poly(n,m), over $\mathbb{Q}$ of poly(n,m) bit size, including the bit-sizes of the constants, and over the variables $x = (x_1, \dots, x_l)$, $l =$ poly(n,m), and $v = (v_1, \dots, v_n)$, such that the polynomials $C[V,m,c](x,v)$'s computed by $C[V,m,c]$'s are of poly(n,m) degree and can be expressed in the form*

$$C[V,m,c](x,v) = \sum_j f_{j,c}(v) g_{j,c}(x), \tag{9}$$

with homogeneous $f_{j,c}$'s $\in K[V]^G$ and $g_{j,c}$'s, so that $K[V]^G$ is generated by $f_{j,c}(v)$'s, and $g_{j,c}(x)$'s are linearly independent.

(b) It is called explicit without any degree restrictions *if the degree requirement on $C[V,m,c](x,v)$'s is dropped.*

(c) It is called strongly explicit *if, in addition to all the properties in (a), the circuits $C[V,m,c]$'s are weakly skew (cf. Section 2.1).*

(d) If V/G is explicit, we say that an explicit First Fundamental Theorem *holds for $K[V]^G$, with the circuits $C[V,m,c]$'s constituting an explicit (polynomial-time-computable) encoding of a set of generators for $K[V]^G$.*

If V/G is strongly explicit, we say that a strongly explicit First Fundamental Theorem holds for $K[V]^G$.

(e) The notions in (a)–(d) are defined in the relaxed sense *by requiring that $f_{j,c}(v)$'s in (9) only form a set of separating invariants [17] (cf. also Section 7 here) of $K[V]^G$, rather than a set of generators.*

We are abusing the terminology a bit here. Formally, instead of saying that V/G is explicit, we should really be saying that the family $\{W_{\langle V,G\rangle}\}$, indexed by the specification $\langle V, G\rangle$, where $W_{\langle V,G\rangle} := V/G$, is explicit.

Conjecture 5.3 *The categorical quotient V/G is explicit, without any degree restrictions in general.*

It may be conjectured that V/G is explicit (with the usual low-degree restrictions), if $K[V]^G$ has a set of generators of poly(n, m) degree.

For all the applications in this article and in geometric complexity theory (cf. Remark 2 after Theorem 9.7), a weaker form of this conjecture stipulating explicitness of V/G only in the relaxed sense (cf. Definition 5.2 (e)) suffices.

Conjecture 5.3 is proved in this article for $V = M_m(K)^r$ with the adjoint action of G (cf. Theorem 6.1), and for arbitrary V when m is constant (cf. Theorem 8.1). The relaxed form of the conjecture in positive characteristic is also proved for the ring of matrix invariants (cf. Theorem 10.7).

For constant m, we shall construct a $C[V, m, c]$ with depth four; cf. Theorem 8.1. The degrees of the generators encoded by $C[V, m, c]$ are at most exponential in its depth. Comparing this bound with the degree bound in Derksen [16] (cf. Theorem 8.2), one may expect $C[V, m, c]$ in Conjecture 5.3 to have $O(\text{poly}(m, \log n))$ depth in general.

The simplest instance of the conjecture that the reader can check is the following. Let $G = SL_m(K)$, and $V = K^m \oplus \cdots K^m$ (r times), with the action of G from the left. The coordinate ring $K[V]$ can be identified with the ring $K[U]$ generated by the entries of an $m \times r$ variable matrix U. By the First Fundamental Theorem of invariant theory [33, 94], the invariant ring $K[V]^G$ in this case is generated by the $r \times r$ minors of U. The corresponding map (7) in this case is the well-known Plücker map [33] $U \to (\ldots, m_\alpha(U), \ldots)$, where $m_\alpha(U)$ ranges over all $r \times r$ minors of U. The categorical quotient V/G in this case is the Grassmanian. It can be checked that the Grassmanian is strongly explicit, with the defining map being the Plücker map.

For explicit varieties in general, the image of the map ψ_n in (6) need not be closed. In contrast, for categorical quotients we have:

Theorem 5.4 (Mumford, Fogarty, and Kirwan [75]) *(cf. Theorem 1.1 in [75] and Theorem 4.6 and 4.7 in [79])*

(a) The image of $\pi_{V/G}$ in (7) is closed. Hence, the map $\pi_{V/G} : V \to V/G$ is surjective.

(b) For any $x \in V/G$, $\pi^{-1}_{V/G}(x)$ contains a unique closed G-orbit.

(c) For any G-invariant (closed) subvariety $W \subseteq V$, $\pi_{V/G}(W)$ is a closed subvariety of V/G.

(d) Given $v, w \in V$, the closures of the G-orbits of v and w intersect iff $r(v) = r(w)$ for all $r \in K[V]^G$.

These additional properties of V/G play a crucial role in this article; cf. Remark 3 after Theorem 5.13.

5.1.5 Explicit variety associated with p-computable polynomials

Let $\{p_n(v, x)\}$, $v = (v_1, \ldots, v_r)$, $r = r_n = \text{poly}(n)$, $x = (x_1, \ldots, x_n)$, be a uniform p-computable [91] family of polynomials, homogeneous in v, with rational constants. This means $p_n(v, x)$ has poly(n) degree, has a circuit over $\mathbb{Q}$ of poly(n) bit-size, and, given n, the specification of this circuit can be computed in poly(n) time.

Let $p_n(v,x) = \sum_\mu f_\mu(v)\mu(x)$, where μ ranges over all monomials in x of degree $\leq \deg(p_n) =$ poly(n). Let $m = m_n$ be the number of such monomials. Let $\psi = \psi_n$ be the map

$$\psi : v \in K^r \to (\ldots, f_\mu(v), \ldots) \in K^m.$$

Then $\{W_n = \overline{Im(\psi_n)}\}$ is an explicit family of varieties, with the defining map ψ_n and the defining polynomial p_n.

5.1.6 Explicit toric variety

Let $\{p_n(x)\}$, $x = (x_1, \ldots, x_n)$, be a uniform p-computable family of homogeneous polynomials over x and K. Let $p_n(x) = \sum_\mu a_\mu \mu(x)$, where $a_\mu \in K$, and μ ranges over all monomials in x of total degree $= \deg(p_n) =$ poly(n). Let $m = m_n$ be the number of such monomials. Consider the monomial map ψ_n:

$$\psi_n : v = (v_1, \ldots, v_n) \in K^n \to (\ldots, a_\mu \mu(v), \ldots) \in K^m.$$

Let $W_n = \overline{Im(\psi_n)}$, and $P(W_n)$ its projectivization. Then $\{P(W_n)\}$ is an explicit family of toric varieties, with the defining polynomial

$$F_n(v,x) = \sum_\mu a_\mu \mu(v)\mu(x).$$

This polynomial is p-computable and uniform, since a circuit for computing it can be obtained from the one for p_n by replacing each x_i with $v_i x_i$.

The main difference between the explicit toric variety here and the more general explicit variety in Section 5.1.5 is that $\mu(v)$ here is a monomial, whereas $f_\mu(v)$ in Section 5.1.5 can be any homogeneous polynomial.

5.2 Unconditional upper bound for the problem of constructing an h.s.o.p.

We now study the problem of constructing an h.s.o.p. for an explicit variety. The following generalization of Theorem 4.1 gives the currently best upper bound for this problem.

Theorem 5.5 *The problem of constructing an h.s.o.p. for an explicit variety W_n (cf. Definition 5.1) belongs to EXPSPACE. (This means it can be solved in work-space that is exponential in n.)*

Assuming the Generalized Riemann Hypothesis, it belongs to EXPH (the exponential hierarchy), if W_n has defining equations that can be computed in time that is exponential in n.

Proof: The proof is similar to that of Theorem 4.1, with W_n in place of $\Delta[\det, m]$. Q.E.D.

5.3 The problem NNL for explicit varieties

If we insist on an h.s.o.p., then Theorem 5.5 is the best that we can do at present. However, if we are willing to settle for a small homogeneous set $S \subseteq K[W_n]$ of poly(n) size, but not necessarily of the optimal size, such that $K[W_n]$ is integral over the subring generated by S, then Theorem 5.11 proved below says that we can do much better. Relaxing the optimality constraint on cardinality, but insisting on succinctness of specification in exchange, we are thus led to the following notion of an s.s.o.p. It generalizes the notion of an s.s.o.p for $\Delta[\det, m]$ (cf. Section 1.2) to arbitrary explicit varieties.

Definition 5.6 *Let $\{W_n\}$ be an explicit family of varieties as in Definition 5.1, $z_1, \ldots, z_{m_n}$ the coordinates of the ambient space K^{m_n} containing W_n, and ψ_n^* the comorphism of $\psi_n : K^{r_n} \to K^{m_n}$ in (6). Note that $K[W_n]$ is graded, with* $\deg(z_j) = \deg(f_j)$.

(a) We say that $s \in K[W_n]$ has a short specification *if $\psi_n^*(s)$ has a circuit over $\mathbb{Q}$ and $v_1, \ldots, v_{r_n}$ of O(poly(n)) bit-length (not just size), which computes the polynomial function on K^{r_n} corresponding to $\psi_n^*(s)$.*

(b) We say that a set $S \subseteq K[W_n]$ is a small system of parameters (s.s.o.p.) *for $K[W_n]$ (and W_n) if (1) each element $s \in S$ has a short specification as in (a) and is homogeneous of poly(n) degree, (2) $K[W_n]$ is integral over its subring generated by S, and (3) the size of S is poly(n).*

We say that S is an explicit system of parameters (e.s.o.p.) *if, in addition, the specification of S, consisting of a circuit for $\psi_n^*(s)$ for each $s \in S$ as in (a), can be computed in poly(n) time.*

If W_n is strongly explicit then, by convention, we assume that the short specification as in (a) for each element of $s \in S$ is a weakly skew circuit (cf. Section 2.1).

(c) S.s.o.p. and e.s.o.p. without any degree restrictions are defined by dropping the degree requirement in (b) (1). Quasi-e.s.o.p. and quasi-s.s.o.p. are defined by replacing poly(n) by $2^{polylog(n)}$.

(d) We call S separating *if, for any two distinct points $u, v \in W_n$, there exists an $s \in S$ such that $s(u) \neq s(v)$.*

Let F_n and C_n be as in Definition 5.1. For any $0 \leq c \leq \deg(F_n)$, let $C_{n,c}$ be the circuit computing the degree c-component (in v) of F_n. If $\deg(F_n)$ is poly(n), then, given C_n and c, we can compute $C_{n,c}$ in poly(n) time using the Vandermonde interpolation technique as per Strassen [88]; cf. also the survey by Shpilka and Yehudayoff [87]. Alternatively, the explicit variety W_n can be specified by giving the circuits $C_{n,c}$'s, $1 \leq c \leq \deg(F_n)$, instead of the circuit C_n. For any $b \in \mathbb{N}^n$ of poly(n) bit-length, let $C_{n,c,b}$ be the instantiation of $C_{n,c}$ at $x = b$.

Definition 5.7 *We say that $s \in K[W_n]$ is* strict *if, for some $b \in \mathbb{N}^n$ of poly(n) bit-length and $0 < c \leq \deg(F_n)$, $\psi_n^*(s)(v) = C_{n,c,b}(v)$. This means $s = \sum_j z_j g_j(b)$, where j ranges over all indices such that $\deg(f_j) = c$. Such a strict s can be specified succinctly by the triple (b, c, C_n), or the pair $(b, C_{n,c})$, or just b if $C_{n,c}$ is implicit (as in the case of the explicit variety $\Delta[\det, m]$).*

We say that an s.s.o.p. or an e.s.o.p. S is strict *if each $s \in S$ is strict. It can then be specified by the set of pairs $(b, C_{n,c})$'s, or just by the set of b's if $C_{n,c}$'s are implicit. Strict quasi-s.s.o.p. and quasi-e.s.o.p. are defined by replacing poly(n) by $2^{polylog(n)}$.*

Thus a strict s.s.o.p. has a short specification (cf. Definition 5.6 (a)) based on the circuit C_n defining the variety W_n itself. As such strictness is a natural way to ensure succinctness. We shall prove later (cf. Corollary 5.10) that a strict s.s.o.p. exists.

By *NNL for* W_n, we mean the problem of constructing an s.s.o.p. for $K[W_n]$. By *NNL in a strong form for* W_n, we mean the problem of constructing a separating s.s.o.p. for $K[W_n]$. By *NNL in a strict and strong form for* W_n, we mean the problem of constructing a strict, separating s.s.o.p. for $K[W_n]$. We say that NNL for W_n has an explicit solution, if $K[W_n]$ has an e.s.o.p.

As the reader can check, an s.s.o.p. for $\Delta[\det, m]$ defined in Section 1.2 is a specialization of the general definition of a strict s.s.o.p. given above. The variety W_n here is the variety $\Delta[\det, m]$ there, the map ψ_n here is the map $\phi : M_{m^2}(K) \to \mathcal{X}$ in Section 5.1.1, and a strict s.s.o.p. S specified by a set of b's here is $S(\mathcal{B})$ specified by a set $\mathcal{B}$ of $m \times m$ matrices in Section 1.2.

Strictness is used in the proof of Theorem 4.9 to derive a lower bound from NNL; cf. the proof of Lemma 4.8. It is open if similar lower bounds can be derived from non-strict NNLs. The s.s.o.p.'s constructed in all the main results of this article stated in Section 1 are strict.

For simplicity, in what follows, we often keep n implicit and denote W_n by W, m_n by m, r_n by r, ψ_n by ψ, and so on.

5.4 Monte Carlo algorithm

The following generalization of Theorem 4.2 proves a stronger form of Theorem 1.2.

Theorem 5.8 *Let $\{W_n\}$ be an explicit family of varieties. Then there is a poly(n)-time Monte Carlo algorithm to construct a separating, strict s.s.o.p. for $K[W_n]$, which is correct with a high probability.*

Proof: Let $W = W_n \subseteq K^m$, $m = m_n$, be an explicit variety as in Definition 5.1, and C_n the circuit computing $F_n(v, x)$, $v = (v_1, \ldots, v_r)$, $r = r_n$, and $x = (x_1, \ldots, x_n)$, as there. Let $s = \text{poly}(n)$ be its size, and $d = \text{poly}(n)$ its degree. Let $u = 2s(d+1)^2$.

Choose $T \subseteq [u]^n$ of size $6(s+1+n)^2$ randomly. By Theorem 2.3, it is a hitting set with a high probability against all nonzero polynomials $h(x)$ of degree $\leq d$ that can be approximated infinitesimally closely by circuits over K and x of size $\leq s$. More strongly, replacing s by $2s+1$, we can also assume that, given any two distinct polynomials $h_1(x)$ and $h_2(x)$ of degree $\leq d$ that can be approximated infinitesimally closely by circuits over K and x of size $\leq s$, there exists $b \in T$ such that $h_1(b) \neq h_2(b)$.

This probabilistic construction of T takes poly(s) = poly(n) time. In what follows, we assume that T is such a hitting set.

For each $b \in T$ and $0 < c \leq \deg(F_n)$, define $h_{b,c}(z) := \sum_j z_j g_j(b) \in K[W_n]$, where j ranges over all indices such that $\deg(f_j) = c$, and $z = (z_1, \ldots, z_m)$ denote the coordinates of the ambient space K^m containing $W = W_n$. Then $\deg(h_{b,c}) = c$, since $\deg(z_j) = \deg(f_j)$, as per the grading on $K[W_n]$. Let

$$S = \{h_{b,c}(z) \mid b \in T, 0 < c \leq \deg(F_n)\} \subseteq K[W_n]. \qquad (10)$$

It now follows from Lemma 5.9 (c) and (d) below that S is a separating, strict s.s.o.p. Q.E.D.

Lemma 5.9 *Suppose $W = W_n$ is an explicit variety. Let S and T be as in (10). Then:*

(a) $W \cap Z(S) = \{0\}$, where $Z(S) \subseteq K^m$ is the zero set of S, and 0 denotes the origin in K^m.

(b) The coordinate ring $K[W]$ is integral over the subring generated by S.

(c) The set S is a strict s.s.o.p.

(d) The set S is also separating.

Proof: Let $\psi = \psi_n$, f_j, g_j, and $F = F_n(v, x)$ be as in Definition 5.1.

(a) By Hilbert's Nullstellensatz, we can assume that $K = \mathbb{C}$, since F_n, and hence W, and $Z(S)$ are defined over $\mathbb{Q}$. Consider any nonzero point $w = (w_1, \ldots, w_m) \in W \subseteq K^m$. We have to show that $h_{b,c}(w) \neq 0$ for some $b \in T$ and $0 < c \leq \deg(F_n)$. Let $F_w(x) = \sum_j w_j g_j(x)$. Since $g_j(x)$'s are linearly independent, $F_w(x)$ is not identically zero as a polynomial in x. Recall that $K[W]$ is graded, with $\deg(z_j) = \deg(f_j)$. Let $F_w(x)_c = \sum_j w_j g_j(x)$, where j ranges over all indices such that $\deg(f_j) = c$. Then $F_w(b)_c = h_{b,c}(w)$. So we have to show that $F_w(b)_c \neq 0$ for some $b \in T$ and $0 < c \leq \deg(F_n)$.

Since $W = \overline{\mathrm{Im}(\psi)}$, and the closure in the Zariski topology coincides with the closure in the complex topology (cf. Theorem 2.33 in [74]), there exists, for any $\delta > 0$, $p_\delta \in K^r$, $r = r_n$, such that $||\psi(p_\delta) - w||_2 \leq \delta/(mA)$, where $A = \max\{||g_j||_2\}$ and, for any polynomial e, $||e||_2$ denotes the L_2-norm of the coefficient vector of e. Since $w \neq 0$, taking δ to be small enough, we can assume that $\psi(p_\delta) \neq 0$. Since $\psi(p_\delta) = (f_1(p_\delta), \ldots, f_m(p_\delta))$, and $g_j(x)$'s are linearly independent, $F_n(p_\delta, x) = \sum_j f_j(p_\delta) g_j(x)$ is not an identically zero polynomial in x. Let C_n be the circuit computing $F_n(v, x)$ as in Definition 5.1. Let $C_{n,\delta}$ be the circuit obtained from C_n by specializing v to p_δ. Then the size of $C_{n,\delta}$ is $s = \mathrm{size}(C_n) = \mathrm{poly}(n)$, and the degree is $d = \deg(C_n) = \mathrm{poly}(n)$. Furthermore,

$$||C_{n,\delta}(x) - F_w(x)||_2 = ||\sum_j (f_j(p_\delta) - w_j) g_j(x)||_2 \leq mA||\psi(p_\delta) - w||_2 \leq \delta.$$

Since δ can be made arbitrarily small, it follows that $F_w(x)$ can be approximated infinitesimally closely by circuits of degree $\leq d$ and size $\leq s$. Since T is a hitting set, and $F_w(x)$ is not identically zero as a polynomial in x, there exists $b \in T$ such that $F_w(b) \neq 0$. Hence $F_w(b)_c \neq 0$ for some $c \leq \deg(F_n)$. This proves (a).

(b) By (a) and Hilbert's Nullstellensatz, it follows that, given any $t \in K[W]$, t^l belongs to the ideal (S) in $K[W]$ generated by S, for some large enough positive integer l. Since $K[W]$ is graded, it now follows from the graded Noether's normalization lemma (Lemma 3.2) that $K[W]$ is integral over its subring generated by S. This proves (b).

(c) S is clearly strict by its definition. So it remains to verify the properties (1)-(3) in Definition 5.6 (b).

(1) We have to show that each $h_{b,c}(z) \in S$ has a short specification. We have $\psi^*(h_{b,c})(v) = F_n(v, b)_c$, the degree c component of $F_n(v, b)$. Since W is explicit, cf. Definition 5.1, we can compute the description of the circuit C_n over $\mathbb{Q}$ computing F_n in $\mathrm{poly}(n)$ time. Hence the total size of C_n, including the bit-lengths of the constants in it, is $\mathrm{poly}(n)$. Using Vandermonde interpolation as in [88, 87], we can construct, using C_n, in $\mathrm{poly}(n)$ time a circuit $C_{n,c}$, for every $0 < c \leq \deg(F_n)$, that computes the degree c-component (in v) of F_n. The circuit $C_{n,c,b}$ for

computing $\psi^*(h_{b,c})(v) = F_n(v,b)_c$ is obtained by instantiating the circuit $C_{n,c}$ at $x = b$. Its total size (including the bit-lengths of the constants) is poly(n), and its degree is poly(n). This shows that each $h_{b,c}(z)$ (or rather $\psi^*(h_{b,c})$) has a short specification.

(2) By (b), $K[W]$ is integral over the subring generated by S.

(3) Since the size of T is poly(s) = poly(n), and deg(F_n) is poly(n), the size of S is poly(n).

This shows that S is a strict s.s.o.p.

(d) Consider any two distinct points $w = (w_1, \ldots, w_m), w' = (w'_1, \ldots, w'_m) \in W \subseteq K^m$. Let $F_w(x) = \sum_j w_j g_j(x)$, and $F_{w'}(x) = \sum_j w'_j g_j(x)$. These are distinct polynomials, since w and w' are distinct and g_j's are linearly independent. It also follows as in the proof of (a) that $F_w(x)$ and $F_{w'}(x)$ can be approximated infinitesimally closely by circuits of degree $\leq d$ and size $\leq s$. Hence, by the stronger assumed property of the hitting set T, there exists $b \in T$ such that $F_w(b) \neq F_{w'}(b)$. Hence, $F_w(b)_c \neq F_{w'}(b)_c$, for some $c \leq \deg(F_n)$. This means $h_{b,c}(w) \neq h_{b,c}(w')$. Hence S is separating. Q.E.D.

The following is a corollary of Theorem 5.8.

Corollary 5.10 *A separating, strict s.s.o.p. exists for the coordinate ring $K[W_n]$ of any explicit variety W_n.*

5.5 Conditional derandomization

We now derandomize the Monte Carlo algorithm in Theorem 5.8 using an appropriate black-box-derandomization or hardness hypothesis.

The following result proves the analogues of Theorems 4.5 and 4.6 for any explicit variety.

Theorem 5.11 *Let $\{W_n\}$ be an explicit family of varieties as in Definition 5.1. Then:*

(a) The variety W_n has a separating, strict e.s.o.p., assuming the strengthened black-box derandomization hypothesis for polynomial identity testing for small degree circuits over K.

(b) The variety W_n has a separating, strict quasi-e.s.o.p., assuming that there exists a family $\{h_n(x_1, \ldots, x_n)\}$ of exponential-time-computable, multi-linear, integral polynomials such that h_n cannot be approximated infinitesimally closely by circuits over K of $O(2^{n^\epsilon})$ size, for some constant $\epsilon > 0$, as $n \to \infty$.

Proof:

(a) We have to show that a separating, strict s.s.o.p. for an explicit variety W_n can be constructed in poly(n) time, assuming the strengthened black-box derandomization hypothesis for polynomial identity testing for small degree circuits over K.

This follows if, instead of the randomly chosen hitting set T in the proof of Theorem 5.8, we use an explicit (poly(n)-time computable) hitting set T provided by the strengthened black-box derandomization hypothesis.

(b) This follows from (a) and Theorem 2.4. Q.E.D.

Remark 1: If the multi-linear polynomial in (b) is the permanent, then, as for $\Delta[\det, m]$ (cf. Theorem 4.7), a separating, strict s.s.o.p. for W_n can be constructed fast in parallel.

Remark 2: If in (b) we assume instead that there exists a family $\{h_n(x_1, \ldots, x_n)\}$ of exponential-time-computable, multi-linear, integral polynomials such that h_n cannot be approximated infinitesimally closely by circuits over K of $O(n^a)$ size, for any constant $a > 0$, as $n \to \infty$, then it can be proved similarly that the strict form of NNL for W_n can be solved in $O(2^{n^\epsilon})$-time (assuming that poly(n) is replaced by 2^{n^ϵ} in Definition 5.6), for any constant $\epsilon > 0$.

Remark 3: The statement (b), in conjunction with [39], implies that an explicit W_n has a separating, strict quasi-e.s.o.p., assuming that there exists a family $\{h_n(x_1, \ldots, x_n)\}$ of exponential-time-computable, multi-linear, integral polynomials such that h_n cannot be approximated infinitesimally closely by depth three circuits over K of $O(2^{n^{\frac{1}{2}+\epsilon}})$ size, for some constant $\epsilon > 0$, as $n \to \infty$. A similar result also holds for homogeneous depth four circuits. The known $\Omega(2^{n^{1/2} \log n})$ lower bounds for the restricted versions of these circuits [53] do not imply any nontrivial result for NNL for explicit varieties. Thus there is a sharp phase transition in the difficulty of the lower bound problem in this model at the exponent 1/2.

Remark 4: The statement (a) also holds for explicit varieties without any degree restrictions, with the general (strengthened) polynomial identity testing without any degree restrictions in place of the (strengthened) polynomial identity testing for small degree circuits.

Remark 5: If W_n is strongly explicit (cf. Definition 5.1), then the (strengthened) polynomial identity testing for small degree circuits in the statement (a) can be replaced with (strengthened) symbolic determinant identity testing. In this case it can also be shown that NNL for W_n belongs to DET, assuming that the strengthened black-box derandomization problem for symbolic determinant identity testing belongs to DET (as may be conjectured).

Remark 6: If W_n is strongly explicit, then it can be assumed that s.s.o.p.'s and e.s.o.p.'s in Theorems 5.8, 5.11 and Corollary 5.10 consist of weakly skew circuits (cf. Definition 5.6).

Remark 7: The derandomization hypothesis in the statement (a) is only needed for the class of circuits used in the definition of W_n (cf. Definition 5.1).

5.6 An unconditional EXPSPACE-algorithm

The following result gives the current best, unconditional, deterministic upper bound for NNL for explicit varieties.

Theorem 5.12 *Let $\{W_n\}$ be an explicit family of varieties. Then the problem of constructing or verifying a strict s.s.o.p. for $K[W_n]$ belongs to EXPSPACE. This means it can be solved in $O(2^{poly(n)})$ work-space. It belongs to EXPH, assuming the Generalized Riemann Hypothesis.*

Proof: For construction, this follows from Theorem 5.11 (a) and Theorem 2.8. The proof for verification is implicit in the proof for construction. Q.E.D.

5.7 NNL for explicit varieties with closed defining maps

Theorem 5.11 (a) can be improved as follows if the image of the defining map ψ_n in (6) is closed.

Theorem 5.13 *Let $\{W_n\}$ be an explicit family of varieties such that the image of the map ψ_n in (6) is closed. Then the coordinate ring $K[W_n]$ of W_n has a separating, strict e.s.o.p., assuming the standard (instead of the strengthened) black-box derandomization hypothesis for polynomial identity testing for small degree circuits over K.*

Proof: The only reason we needed the strengthened black-box derandomization hypothesis in the proof of Theorem 5.11 (a) is because the image of ψ_n need not be closed, in general. If it is closed, then we can use the standard black-box derandomization hypothesis instead. Q.E.D.

Remark 1: Theorem 5.13, in conjunction with the PSPACE-bound for the standard black-box derandomization (Proposition 2.9), implies that NNL for W_n belongs to PSPACE unconditionally if the image of ψ_n is closed.

Remark 2: Theorem 5.13, in conjunction with Theorem 2.1, implies that W_n has a strict quasi-e.s.o.p., assuming that there exists a family $\{h_n(x_1, \ldots, x_n)\}$ of exponential-time-computable, integral, multi-linear polynomials such that h_n cannot be computed by circuits over K of size sub-exponential in n. This improves Theorem 5.11 (b) when the image of ψ_n is closed.

Remark 3: If W is strongly explicit (cf. Definition 5.1), then the low-degree polynomial identity testing in Theorem 5.13 can be replaced by symbolic determinant identity testing. This result, in conjunction with Theorem 5.4 (a), implies that if W_n is a strongly explicit categorical quotient, then an e.s.o.p. for W_n exists, assuming the standard (instead of the strengthened) black-box derandomization hypothesis for symbolic determinant identity testing. This fact will play a crucial role in the proofs of Theorems 1.4 and 1.6.

Remark 4: We only need the derandomization hypothesis in Theorem 5.13 for the class of circuits used in the definition of W_n (cf. Definition 5.1).

5.8 Equivalence

The following is the analogue of Theorem 4.9 for general polynomial identity testing.

Theorem 5.14 *(a) The strengthened black-box derandomization hypothesis for general polynomial identity testing over K, without any degree restrictions, holds iff the orbit closure $\Delta[H(Y), k, m]$ (cf. Section 5.1.3), with $k = m$, has a strict e.s.o.p.*

(b) The strengthened black-box derandomization hypothesis for low-degree polynomial identity testing over K holds iff the orbit closure $\Delta[H(Y)_m, k, m]$ (cf. Section 5.1.3), with $k = m$, has a strict e.s.o.p.

(c) Ignoring a quasi prefix, a sub-exponential lower bound for a family of exponential-time-computable, integral, multi-linear polynomials as in Theorem 2.4 holds iff the orbit closure $\Delta[H(Y)_m, k, m]$ (cf. Section 5.1.3), with $k = m$, has a strict e.s.o.p.

Proof:

(a) The polynomial $H(Y)$ corresponds to a universal circuit (cf. Section 5.1.3), just as the determinant corresponds [60] to a universal weakly skew circuit. The polynomials corresponding to the points in $\Delta[H(Y), k, m]$ can be approximated infinitesimally closely by circuits over K of size $\text{poly}(k, m)$ and degree $= \deg(H(Y))$. Though $\deg(H(Y))$ is exponential in k, Theorem 2.3

still implies existence of a small hitting set, with $O(\text{poly}(k,m))$ bit-length of specification, against such polynomials. The rest of the proof is similar to that of Theorem 4.9 (a), with $\Delta[H(Y),k,m]$, with $k=m$, in place of $\Delta[\det,m]$.

(b) The proof is similar to that of (a), with $\Delta[H(Y)_m,k,m]$, with $k=m$, in place of $\Delta[H(Y),k,m]$.

(c) This follows from (b), Theorem 2.4, and Proposition 2.7. Q.E.D.

Remark: It can be shown similarly that the strengthened black-box derandomization hypothesis for polynomial identity testing for depth three circuits over K and n variables with degree $\leq d$ and top fan-in $\leq k$ holds iff the k-th secant variety $X(d,k,n)$ of the Chow variety (cf. Section 5.1.2) has a strict e.s.o.p.

5.9 The NNL for the orbit closure of the permanent

Analogue of Theorem 4.6 also holds for the orbit closure $\Delta[\text{perm},n,m]$ of the permanent defined in Section 4.4, though this variety is not explicit. So let us call it *weakly explicit*.

Specifically, given any set $\mathcal{B}=\{B_1,\ldots,B_k\}\subseteq M_m(\mathbb{N})$, define $\psi_{\mathcal{B}}:\mathcal{X}\to K^k$ as in Section 1.2, replacing $\det(X)$ by $z^{m-n}\text{perm}(Y)$ in that definition, where Y is any $n\times n$ sub-matrix of X, and z is any entry in X outside Y. We say that $\mathcal{B}$ specifies a *strict s.s.o.p.* for $\Delta[\text{perm},n,m]$ if (1) the total bit-length of B_i's is $\text{poly}(m)$, and (2) $\psi_{\mathcal{B}}$ does not vanish on any non-zero point in the affine cone $\hat{\Delta}[\text{perm},n,m]\subseteq\mathcal{X}$ of $\Delta[\text{perm},n,m]$. We say that $\mathcal{B}$ specifies a *strict e.s.o.p.* if, in addition, it is $\text{poly}(m)$-time-computable. A strict quasi-e.s.o.p. is defined by replacing $\text{poly}(m)$ by $2^{\text{polylog}(m)}$.

Theorem 5.15 *The variety $\Delta[perm,n,m]$ has a strict quasi-e.s.o.p., assuming that there exists a family $\{p_k(x_1,\ldots,x_k)\}$ of exponential-time-computable, multi-linear, integral polynomials such that p_k cannot be approximated infinitesimally closely by permanents of symbolic matrices over K of $O(2^{k^\epsilon})$ size, for some constant $\epsilon>0$, as $k\to\infty$.*

Proof: By inserting the oracle for the permanent in appropriate places in the proof of Theorem 2.4, it follows that strengthened polynomial identity testing for small degree circuits over K of size $\leq s$, with oracle gates for the permanent, has $O(2^{\text{polylog}(s)})$-time-computable hitting set (defined in the obvious way), assuming that there exists a family $\{p_k(x_1,\ldots,x_k)\}$ of exponential-time-computable, multi-linear, integral polynomials such that p_k cannot be approximated infinitesimally closely by circuits over K, with oracle gates for the permanent, of $O(2^{k^\epsilon})$ size, for some constant $\epsilon>0$, as $k\to\infty$. It easily follows from Valiant [91] that a low-degree circuit of size s, with oracle gates for the permanent, can be simulated by the permanent of a symbolic matrix of $\text{poly}(s)$ size. Hence, the same conclusion holds assuming instead that p_k cannot be approximated infinitesimally closely by permanents of symbolic matrices over K of $O(2^{k^\epsilon})$ size, for some constant $\epsilon>0$, as $k\to\infty$. The proof is now similar to that of Theorem 4.6, using this fact in place of Theorem 2.4. Q.E.D.

6 Explicitness of V/G for the ring of matrix invariants

In this section we prove Theorem 1.3, assuming that the base field K has characteristic zero.

Let $V = M_m(K)^r$, with the adjoint action of $G = SL_m(K)$, be as in Section 1.4. Given $\sigma \in G$, the adjoint action maps $(A_1, \ldots, A_r) \in V$ to $(\sigma A_1 \sigma^{-1}, \ldots, \sigma A_r \sigma^{-1})$. Let $n = \dim(V) = rm^2$. Let $U_1, \ldots, U_r$ be variable $m \times m$ matrices, and let $U = (U_1, \ldots, U_r)$. Identify the coordinate ring $K[V]$ of V with the ring $K[U_1, \ldots, U_r]$ generated by the variable entries of U_i's. Let $K[V]^G \subseteq K[V]$ be the ring of invariants with respect to the adjoint action of G. Let $V/G = \text{spec}(K[V]^G)$.

Call two words in $[r]^*$ *equivalent* if one can be obtained from the other by a circular rotation. Recall that (cf. Section 2.5) $[r]$ denotes $\{1, \ldots, r\}$.

The following is a restatement of Theorem 1.3 in characteristic zero for convenience.

Theorem 6.1 *The categorical quotient V/G is strongly explicit (cf. Definition 5.2), or in other words, a strongly explicit First Fundamental Theorem holds for $K[V]^G$.*

Specifically, let $X = (X_1, \ldots, X_r)$ be an r-tuple of $k \times k$ variable matrices, where $k = m^2$, and m is the dimension of U_i's as above. Then there exist poly(n)-time-computable weakly skew (Section 2.1) circuits C_l's, $l \leq m^2$, over $\mathbb{Q}$ and the variable entries of X_i's and U_i's, such that (1) the polynomial functions $C_l(X, U)$'s computed by C_l's are of poly(n) degree, homogeneous in X and U, and can be written as

$$C_l(X, U) = \sum_{[\alpha]} f_{[\alpha],l}(U) g_{[\alpha],l}(X), \tag{11}$$

where $[\alpha] = [\alpha_1 \cdots \alpha_l]$ ranges over the equivalence classes of all words of length l with each $\alpha_j \in [r]$, (2) $g_{[\alpha],l}(X)$'s are linearly independent homogeneous polynomials in the entries of X_i's, and (3) $f_{[\alpha],l}(U)$'s are homogeneous invariants that generate $K[V]^G$.

6.1 Geometric invariant theory

Before proving Theorem 6.1, we recall some results in geometric invariant theory that are needed for its proof.

Theorem 6.2 (Procesi-Razmyslov-Formanek) *[80, 82, 32] (The First Fundamental Theorem for matrix invariants; cf. Theorems 6 and 10 in [32]) The ring $K[V]^G$ is generated by the traces of the form trace$(U_{i_1} \cdots U_{i_l})$, $l \leq m^2$, $i_1, \ldots, i_l \in [r]$.*

Let $K[S_r]$ be the group algebra of the symmetric group S_r on r letters. Write any $\sigma \in S_r$ as a product of disjoint cycles:

$$\sigma = (a_1 \cdots a_{k_1})(b_1 \cdots b_{k_2})...,$$

where 1-cycles are included, so that each of the numbers $1, \ldots, r$ occurs exactly once. Define

$$T_\sigma(U_1, \ldots, U_r) = \text{trace}(U_{a_1} \cdots U_{a_{k_1}}) \text{trace}(U_{b_1} \cdots U_{b_{k_2}}) \cdots . \tag{12}$$

The following result is a consequence of the Second Fundamental Theorem for matrix invariants due to Procesi and Razmyslov [80, 82].

Theorem 6.3 *(cf. Theorem 1 in [32]) Define the K-linear map $\phi : K[S_r] \to K[V]^G$ by*

$$\phi(\sum a_\sigma \sigma) = \sum a_\sigma T_\sigma(U_1, \ldots, U_r).$$

Then $Ker(\phi) = \{0\}$ if $r \leq m$.

Let $X_1, \ldots, X_r$ be $k \times k$ variable matrices. For any word $\alpha = i_1, \ldots, i_l$, $i_j \in [r]$, let

$$T_\alpha(X) = \text{trace}(X_{i_1} \cdots X_{i_l}), \tag{13}$$

where $X = (X_1, \ldots, X_r)$. Let $T_{[\alpha]}(X) = T_\alpha(X)$, where $[\alpha]$ denotes the equivalence class of words equivalent to α under circular rotation. The choice of α in $[\alpha]$ does not matter.

Corollary 6.4 *The traces $\{T_{[\alpha]}(X)\}$, where $[\alpha]$ ranges over all equivalence classes of words of length $l \leq k$, are linearly independent.*

Proof: Suppose to the contrary that there is a linear dependence

$$\sum_{[\alpha]} b_{[\alpha]} T_{[\alpha]}(X) = 0, \quad b_{[\alpha]} \in K. \tag{14}$$

Without loss of generality, we can assume that this relation is homogeneous in every X_i. We can also assume that it is multi-linear in X_i's. Otherwise, we can multi-linearize it by (1) substituting

$$X_i = \sum_{j=1}^{d_i} t_{i,j} X_{i,j}$$

in the l.h.s. of (14), where d_i is the (homogeneous) degree of X_i in the relation, $t_{i,j}$'s are new variables, and $X_{i,j}$'s are new variable $k \times k$ matrices, and then (2) equating the coefficient of $\prod_i \prod_{j=1}^{d_i} t_{i,j}$ to zero.

So assume that the dependence (14) is multi-linear and homogeneous. Without of loss of generality, assume that the variables occurring in this dependence are $X_1, \ldots, X_l$, $l \leq k$. Then each $[\alpha]$ in (14) corresponds to a cyclic permutation $(i_1, \ldots, i_l) \in S_l$, which we denote by $\hat{\alpha}$. Hence, the l.h.s. of (14) equals $\phi(\sum_{\hat{\alpha}} b_{[\alpha]} \hat{\alpha})$, where $\phi : K[S_l] \to K[M_k(K)^l]^{SL_k(K)}$ is the map (cf. Theorem 6.3) that takes $\sum_\sigma a_\sigma \sigma \in K[S_l]$ to $\sum_\sigma a_\sigma T_\sigma(X_1, \ldots, X_l)$. Since $l \leq k$, it follows from (14) and Theorem 6.3 that all $b_{[\alpha]}$'s are zero. Q.E.D.

Remark: The proof above also shows that the monomials in $T_{[\alpha]}(X)$'s of total degree $l \leq k$ in X_i's are linearly independent.

6.2 Proof of Theorem 6.1

For any word $\alpha = i_1, \ldots, i_l$, $l \leq m^2$, $i_j \in [r]$, cf. (13), let

$$T_{[\alpha]}(U) = T_\alpha(U) = \text{trace}(U_{i_1} \cdots U_{i_l}), \tag{15}$$

where $U = (U_1, \ldots, U_r)$. Let

$$F = \{T_{[\alpha]}(U)\}, \tag{16}$$

where $[\alpha]$ ranges over the equivalence classes (for circular rotation) of all words in $1, \ldots, r$ of length $\leq m^2$. Then F generates $K[V]^G$ by Theorem 6.2.

Consider the map $\pi_{V/G}$ from $M_m(K)^r$ to K^t, $t = |F|$, defined as

$$\pi_{V/G}: \quad A = (A_1, \ldots, A_r) \to (\ldots, T_{[\alpha]}(A), \ldots), \tag{17}$$

where $A_i \in M_m(K)$ for all i. By Theorem 5.4 (a), its image is closed, and can be identified with V/G.

For any $l \leq k = m^2$, let

$$T_l(X, U) = \mathrm{trace}((X_1 \otimes U_1 + \cdots + X_r \otimes U_r)^l), \tag{18}$$

where X_i's are new $k \times k$ variable matrices, $X = (X_1, \ldots, X_r)$, $U = (U_1, \ldots, U_r)$, and $\otimes$ denotes the Kronecker product of matrices. Thus each $X_i \otimes U_i$ is an $m' \times m'$ matrix, where $m' = km = m^3$. We have

$$T_l(X, U) = \sum_\alpha T_\alpha(X) T_\alpha(U) = \sum_{[\alpha]} |[\alpha]| T_{[\alpha]}(X) T_{[\alpha]}(U), \tag{19}$$

where $[\alpha] = [\alpha_1 \cdots \alpha_l]$ ranges over the equivalence classes of all words of length l with each $\alpha_j \in [r]$, $|[\alpha]|$ denotes the cardinality of the equivalence class $[\alpha]$ of the word α, and $T_\alpha(U)$ and $T_\alpha(X)$ are as in (15) and (13).

Clearly $T_l(X, U)$, cf. (18), can be computed by an explicit (poly(n)-time computable) weakly skew circuit (Section 2.1). Fix such an explicit circuit C_l computing $T_l(X, U)$. Then $C_l(X, U) = T_l(X, U)$. Let $g_{[\alpha],l}(X) = |[\alpha]| T_{[\alpha]}(X)$, and $f_{[\alpha],l}(U) = T_{[\alpha]}(U)$. Then (11) holds by (19). Furthermore, $g_{[\alpha],l}(X)$'s are linearly independent by Corollary 6.4, and $f_{[\alpha],l}(U)$'s generate $K[V]^G$ by Theorem 6.2.

This proves Theorem 6.1.

7 NNL for the ring of matrix invariants

In this section, Theorem 1.4 is proved, assuming that the base field K has characteristic zero.

Let $V = M_m(K)^r$, $n = \dim(V) = rm^2$, $G = SL_m(K)$, $K[V]^G$, and V/G be as in Section 6. By Theorem 6.1, V/G is strongly explicit. Hence, we can specify V/G succinctly, as per the general definition of an explicit variety (cf. Definition 5.1), by the circuits $C_l(X, U)$'s in Theorem 6.1. Instead, we shall specify V/G and $K[V]^G$ succinctly by just giving the pair (m, r) in unary. This is sufficient and also equivalent, since, given (m, r), we can compute the circuits $C_l(X, U)$'s in Theorem 6.1 in poly(m, r) time.

An s.s.o.p. or an e.s.o.p. for $K[V]^G$ is defined as in Section 1.4. The symbolic determinants that were used in the definition of an s.s.o.p. in Section 1.4 are equivalent to weakly skew circuits (cf. Section 2.1). Quasi-s.s.o.p. and quasi-e.s.o.p. are defined, as before, by replacing poly(n) by $2^{\mathrm{polylog}(n)}$.

Following Derksen and Kemper [17], we call $S \subseteq K[V]^G$ *separating* if, for any two distinct $v, w \in V$ such that $r(v) \neq r(w)$ for some $r \in K[V]^G$, there exists an $s \in S$ such that $s(v) \neq s(w)$. This is a general notion that applies to any finite dimensional representation of a reductive group.

By the problem NNL for $K[V]^G$, we mean, as in Section 1.4, the problem of constructing an s.s.o.p., given (m, r) in unary. By the strong form of NNL, we mean the problem of constructing a separating s.s.o.p.

The reader should check that these definitions are specializations of the general Definition 5.6 for strongly explicit varieties. We can also define strict s.s.o.p. and e.s.o.p. for V/G (cf. Definition 5.7) using the circuits $C_l(X, U)$'s in Theorem 6.1. All s.s.o.p.'s constructed in this section are strict. However, strictness is not as important for V/G as it is for $\Delta[\det, m]$, since existence of a strict e.s.o.p. for V/G does not imply any lower bound. Hence, we shall not worry about strictness in this section.

We prove in this section the following stronger form of Theorem 1.4 in characteristic zero.

Theorem 7.1 (Cf. [69], [30, 31], and Remark 1 in Section 1.7) *The ring $K[V]^G$ has a separating e.s.o.p, assuming the standard black-box derandomization hypothesis for symbolic determinant identity testing. It has a separating quasi-e.s.o.p. unconditionally.*

The following will turn out to be a corollary of the proof of this result.

Theorem 7.2 *The problem of deciding if the G-orbit-closures of two rational points in V intersect belongs to DET $\subseteq$ NC.*

7.1 Construction of an h.s.o.p.

Before we prove Theorem 7.1, we study the problem of constructing an h.s.o.p. (homogeneous system of parameters) for $K[V]^G$ (cf. Definition 3.3). The following result gives the currently best upper bound for this problem.

Theorem 7.3 *The problem of constructing an h.s.o.p. for $K[V]^G$ belongs to EXPH, assuming the Generalized Riemann Hypothesis.*

For the proof, we need the following result.

Recall that the trace function T_σ defined in (12) satisfies [80] the fundamental trace identity

$$F(U_1, \ldots, U_{m+1}) = \Sigma_{\sigma \in S_{m+1}} \text{sign}(\sigma) T_\sigma(U_1, \ldots, U_{m+1}) = 0.$$

Theorem 7.4 (Procesi-Razmyslov) *(The Second Fundamental Theorem for matrix invariants) (cf. Theorem 4.5 in [80]) The ideal of all relations among the trace monomial generators of $K[V]^G$ given by Theorem 6.2 is generated by the elements of the form $F(M_1, \ldots, M_{m+1})$, where M_i's range over all possible monomials in U_j's so that the total length of M_i's is $\leq m^2$.*

This follows from the proof of Theorem 4.5 in [80].

Proof of Theorem 7.3: The defining equations for V/G given in Theorem 7.4 can clearly be computed in time exponential in n. Hence the result follows from Theorem 5.5. Q.E.D.

If we insist on an h.s.o.p., then Theorem 7.3 is the best that we can do at present. But if we only require a small homogeneous S of poly(n) cardinality such that $K[V]^G$ is integral over the subring generated by S, and do not insist on optimality of $|S|$, then Theorem 7.1 says that the double exponential time bound in Theorem 7.3 can be brought down to quasi-polynomial. (Theorem 7.3 only implies a double-exponential time bound for the problem of constructing an h.s.o.p., since conjecturally EXPH $\not\subseteq$ EXP.)

We now turn towards the proof of Theorem 7.1.

7.2 A Monte Carlo algorithm

The first step is an efficient Monte Carlo algorithm to construct an s.s.o.p.

Theorem 7.5 *A separating s.s.o.p. for $K[V]^G$ can be constructed by a poly(n)-time Monte Carlo algorithm that is correct with a high probability.*

In particular, a separating s.s.o.p. for $K[V]^G$ exists.

Proof: By Theorem 6.1, V/G is explicit. Hence the result follows from Theorem 5.8. Q.E.D.

7.3 Reduction of NNL to black-box symbolic determinant identity testing

The next step is to derandomize the Monte Carlo algorithm in Theorem 7.5 assuming a suitable derandomization hypothesis. The first statement in Theorem 7.1 following from the following result.

Theorem 7.6 *Assume that the standard black-box derandomization hypothesis for symbolic determinant identity testing over K holds. Then $K[V]^G$ has a separating e.s.o.p.*

Proof: Since V/G is strongly explicit (cf. Theorem 6.1), and the image of $\pi_{V/G}$ in (17) is closed by Theorem 5.4 (a), it follows from Theorem 5.13 and Remark 3 thereafter that the Monte Carlo algorithm in Theorem 7.5 can be derandomized assuming the standard black-box derandomization hypothesis for symbolic determinant identity testing. Q.E.D.

We now give a second more refined proof of this result, since it is needed for the proof of the second unconditional statement in Theorem 7.1. For this proof, we need the following result from geometric invariant theory. We state in a more general form than what is needed here, since it will be needed in such generality in Sections 8 and 9.

Theorem 7.7 (Derksen and Kemper) *(cf. Theorem 2.3.12 in [17]) Let W be a finite dimensional representation of any algebraic reductive group H over K. Let $S \subseteq K[W]^H$ be a finite separating set (cf. Section 2.3.2 in [17], and the beginning of this section) of homogeneous invariants. Then $K[W]^H$ is integral over the subring generated by S.*

In this section, we shall use this result with $W = V$ and $H = G$.

A refined proof of Theorem 7.6:

We follow the same notation as in Section 6.2.

Let $T_l(X, U)$ be as in (18). Let $U' = (U'_1, \ldots, U'_r)$ be another tuple of variable $m \times m$ matrices, in addition to U. For each $l \leq k = m^2$, define the symbolic trace difference

$$\tilde{T}_l(X, U, U') = T_l(X, U) - T_l(X, U'). \tag{20}$$

Clearly $\tilde{T}_l(X, U, U')$ has a weakly skew circuit (Section 2.1) over X, U, and U' of poly(n) size. Since symbolic determinants are polynomially equivalent to weakly skew circuits (cf. Section 2.1 and [60]), it follows that each $\tilde{T}_l(X, U, U')$ can be expressed as $\det(N_l(X, U, U'))$ for some symbolic matrix $N_l(X, U, U')$ of size $q = \text{poly}(n)$ over X, U, and U'.

By our black-box derandomization hypothesis for symbolic determinant identity testing, there exists an explicit (poly(n)-time computable) hitting set $B = B_{s,q} \subseteq \mathbb{N}^s$ for symbolic determinant identity testing for $q \times q$ matrices whose entries are linear functions of the $s = rk^2$ variable entries of X_i's with coefficients in K. It has to be stressed here that the hitting set B is against non-zero symbolic determinants of size q over X, not over X, U, and U'. The reason will become clear in a moment. Fix such an explicit B. We think of each $b \in B$ as an r-tuple $b = (b_1, \ldots, b_r)$ of $k \times k$ integral matrices.

Let

$$S = \{T_l(b, U) \mid b \in B, 1 \leq l \leq k\} \subseteq K[V]^G. \tag{21}$$

Suppose $A, A' \in V = M_m(K)^r$ are two r-tuples such that, for some invariant $h \in K[V]^G$, $h(A) \neq h(A')$. By Theorem 6.2, it follows that some generator $T_{[\alpha]}(U)$, cf. (16), assumes different values at A and A'. By (19) and Corollary 6.4, this implies that $\tilde{T}_l(X, A, A') = T_l(X, A) - T_l(X, A')$ is not identically zero, as a polynomial in X, for some $l \leq m^2$.

Since $\tilde{T}_l(X, U, U')$ can be expressed as a symbolic determinant of size $q = \text{poly}(n)$ over X, U, and U', $\tilde{T}_l(X, A, A')$ is a symbolic determinant of size $q = \text{poly}(n)$ over X. Since B is a hitting set against such symbolic determinants over X, and $\tilde{T}_l(X, A, A')$ is not identically zero as a polynomial in X, there exists $b \in B$ such that $\tilde{T}_l(b, A, A') \neq 0$, i.e., $T_l(b, A) \neq T_l(b, A')$. It follows that S is separating.

Every element of S is clearly homogeneous of poly(n) degree. By Theorem 7.7, it follows that $K[V]^G$ is integral over the subring generated by S.

Since the hitting set B is explicit, and matrix powering, Kronecker product, and trace have explicit weakly-skew circuits (cf. Section 2.1 and [60]), it follows from (18) that the specification of S consisting of a weakly skew circuit for its every element can be computed in poly(n) time. Hence S is a separating e.s.o.p.

This proves Theorem 7.6. Q.E.D.

Remark 1: The e.s.o.p. constructed in Theorem 7.6 is also strict (cf. Definition 5.7) with respect to the defining polynomials $C_l(X, U)$'s in Theorem 6.1 for V/G.

Remark 2: Assuming a stronger parallel black-box derandomization hypothesis for symbolic determinant identity testing over K, the problem of constructing a separating s.s.o.p. for $K[V]^G$

can be shown to belong to DET $\subseteq$ NC^2 $\subseteq$ P. This hypothesis is that the problem of constructing, given m in unary, a hitting set against non-zero symbolic determinants of size m over (say) m^2 variables belongs to DET.

7.4 Deciding if two orbit-closures intersect

The following is a consequence of the above proof in conjunction with the standard geometric invariant theory.

Theorem 7.8 *The problem of deciding if the closures of the G-orbits of two rational points in V intersect, and finding some invariant in $K[V]^G$ that separates the two if they do not, belongs to co-RDET $\subseteq$ co-RNC.*

The complexity class co-RDET here is the randomized version of co-DET (the complement of DET) [15].

Proof: By Theorem 5.4 (d) and the refined proof of Theorem 7.6, the closures of the G-orbits of $A, A' \in V$ intersect iff the symbolic trace difference $\tilde{T}_l(X, A, A') = T_l(X, A) - T_l(X, A')$ is identically zero for every $l \leq k = m^2$. For rational A and A', this can be tested by a co-RDET algorithm [45]: just substitute large enough random integer values for the entries of X and test if all the differences vanish. If the symbolic trace difference is not identically zero for some l, then this test returns a matrix C such that the test fails for that l when $X = C$. The symbolic trace $T_l(C, U)$ is an invariant that separates A and A' in that case. Q.E.D.

7.5 Replacing symbolic determinants by read-once oblivious algebraic branching programs

In this section we describe how the symbolic determinant identity testing in Theorem 7.6 can be replaced by polynomial identity testing for read-once oblivious algebraic branching programs (cf. Section 2.1), as pointed out by Forbes and Shpilka [31]. In conjunction with their earlier quasi-derandomization of polynomial identity testing for such programs in [30], this implies existence of a quasi-e.s.o.p. for $K[V]^G$, as stated in Theorem 7.1, unconditionally.

Lemma 7.9 (Forbes and Shpilka) *(cf. Lemmas 2.3 and 3.4 in [31]) For any positive integer l, there exists a read-once oblivious algebraic branching program $P_l(Y, U, U')$ over $\mathbb{Z}$, the variable entries of U, U' (thought as indeterminate constants), and the tuple $Y = (y_1, \ldots, y_l)$ of auxiliary variables, with the specification of poly(l, m, r) bit-size, such that*

$$P_l(Y, U, U') = \sum_{\alpha} Y_\alpha (T_\alpha(U) - T_\alpha(U')), \tag{22}$$

where $\alpha = \alpha_1 \alpha_2 \cdots$ ranges over all words of length l, with each $\alpha_j \in [r]$, and $Y_\alpha = \prod_j y_j^{\alpha_j}$.

Proof: The r.h.s. of (22) equals $(\mathrm{trace}(\prod_{j=1}^{l} (\sum_{i=1}^{r} y_j^i U_i))) - (\mathrm{trace}(\prod_{j=1}^{l} (\sum_{i=1}^{r} y_j^i U_i')))$, which can clearly be computed by a read-once oblivious algebraic branching program with the specification of poly(l, m, r) bit-size. Q.E.D.

Since the monomials Y_α's are linearly independent, we can replace $T_l(X, U, U')$ by $P_l(Y, U, U')$ in the refined proof of Theorem 7.6. This implies that Theorem 7.6 also holds after replacing the symbolic determinant identity testing in its statement by polynomial identity testing for read-once oblivious algebraic branching programs (cf. Section 2.1). The existence of a separating quasi-e.s.o.p. as in Theorem 7.1 (and even a quasi-NC algorithm for the strong form of NNL in this case) follows in view of the quasi-NC black-box algorithm for polynomial identity testing for read-once oblivious algebraic branching programs in [30]. This replacement also derandomizes the co-RDET-algorithm in Theorem 7.8 in view of the white-box (cf. Section 2.2) DET-algorithm for polynomial identity testing for read-once oblivious algebraic branching programs in Raz and Shpilka and Arvind et al. [81, 4]. This proves Theorem 7.2. (Unlike the co-RDET algorithm in Theorem 7.8, this algorithm does not return a separating invariant if the two orbit closures do not intersect.)

8 Explicitness of V/G when G has constant dimension

In this section we prove Theorem 1.5.

Let V be a rational representation of $G = SL_m(K)$ of dimension n. The following is a restatement of Theorem 1.5 for convenience.

Theorem 8.1 *The categorical quotient $V/G = spec(K[V]^G)$ is strongly explicit (Definition 5.2), i.e., a strongly explicit First Fundamental Theorem holds for $K[V]^G$, if m is constant.*

We begin by recalling some results from invariant theory and standard monomial theory [57, 24] that are needed to prove this result, and then we prove some complexity-theoretic lemmas.

8.1 A degree bound for the ring of invariants

First, we recall from Derksen [16] a degree bound for a set of generators for $K[V]^G$.

Since G is reductive [33], V can be decomposed as a direct sum of irreducibles:

$$V = \oplus_\lambda m(\lambda) V_\lambda(G), \tag{23}$$

where $\lambda : \lambda_1 \geq \cdots \lambda_r > 0$, $r < m$, is a partition, i.e., a non-increasing sequence of positive integers, $V_\lambda(G)$ is the irreducible Weyl module [33] of G labelled by λ, and $m(\lambda)$ is its multiplicity. We assume that V and G are specified by the tuple

$$\langle V, G \rangle := (n, m; (\lambda^1, m(\lambda^1)); \ldots; (\lambda^s, m(\lambda^s))), \tag{24}$$

which gives n and m in unary, and the multiplicity $m(\lambda^j)$ in unary for each Weyl module $V_{\lambda^j}(G)$ that occurs in the decomposition (23) with nonzero multiplicity. The bit-length of this specification is $O(n + m)$.

The *degree* d of V is defined to be the maximum of $|\lambda| = \sum_i \lambda_i$ over the λ's that occur in this decomposition with nonzero multiplicity. For each copy of $V_\lambda(G)$ that occurs in this decomposition, fix the standard monomial basis of $V_\lambda(G)$ as defined in [57]. It will be reviewed in Section 8.2 below. This yields a basis $B(V)$ of V, which we call the *standard monomial basis* of V. Let $v_1, \ldots, v_n$ be the coordinates of V in this basis. In what follows, we use these concrete coordinates of V throughout. So the elements of $K[V]$ are regarded as polynomials in $v_1, \ldots, v_n$.

Theorem 8.2 (Derksen) *(cf. Theorem 1.1, Proposition 1.2 and Example 2.1 in [16]) The invariant ring $K[V]^G$ is generated by homogeneous invariants of degree $\leq l = nm^2d^{2m^2}$.*

This bound is poly(n), when m is constant, since $d \leq n$ by the following result.

Lemma 8.3 *Let V be as in (23). Then (a) $\dim(V) = n \geq d$, and (b) $n = \Omega(2^{\Omega(m)})$, if $d = \Omega(m^2)$.*

This can be shown using the fact that the dimension of $V_\lambda(G)$ is equal [33] to the number of semi-standard tableau of shape λ. See the preliminary version [69] for the details. The fact (b) will be needed later for the proof of Theorem 8.10.

Theorem 8.2 allows the following concrete realization of V/G.

Let l be as in Theorem 8.2. Let $K[V]^G_l \subseteq K[V]^G$ be the subspace of homogeneous invariants of degree l, and $K[V]^G_{\leq l}$ the subspace of non-constant invariants of degree $\leq l$. The spaces $K[V]_l$ and $K[V]_{\leq l}$ are defined similarly. The dimension t of $K[V]^G_{\leq l}$ is bounded by $\dim(K[V]_{\leq l}) = \sum_{c \leq l} \binom{c+n-1}{n-1}$. This bound is exponential in n, even when m is constant. This worst case upper bound on t is not tight. But we cannot expect a significantly better bound, since the function $h(l) = \dim(K[V]^G_l)$ is a quasi-polynomial [5] of degree $\dim(V/G) \geq \dim(V) - \dim(G) = n - m^2$. This follows from [28] since the singularities of V/G are rational [7]. To prove Theorem 8.1, we have to show that some spanning set of $K[V]^G_{\leq l}$ of cardinality exponential in n can still be encoded by a small uniform circuit.

Let $F = \{f_1, \ldots, f_t\}$ be a set of non-constant homogeneous invariants that span $K[V]^G_{\leq l}$. By Theorem 8.2, F generates $K[V]^G$. Consider the morphism $\pi_{V/G}$ from V to K^t given by

$$\pi_{V/G} : \quad v \to (f_1(v), \ldots, f_t(v)). \tag{25}$$

By Theorem 5.4 (a), the image of this morphism is closed, and V/G can be identified with this closed image. Let $z = (z_1, \ldots, z_t)$ be the coordinates of K^t, I the ideal of V/G under this embedding, and $K[V/G]$ its coordinate ring. Then $K[V/G] = K[z]/I$, and we have the comorphism $\pi^*_{V/G} : K[V/G] \to K[V]$ given by

$$\pi^*_{V/G}(z_i) = f_i. \tag{26}$$

Since f_i's are homogeneous, $K[V/G]$ is a graded ring, with the grading given by $\deg(z_i) = \deg(f_i)$. Furthermore, $\pi^*_{V/G}$ gives the isomorphism between $K[V/G]$ and $K[V]^G$:

$$\pi^*_{V/G}(K[V/G]) = K[V]^G.$$

[5] This means there exist polynomials $h_1(l), \ldots, h_k(l)$ such that $h(l) = h_j(l)$ if $l = j \pmod{k}$. The *degree* of h is the maximum degree of h_j's.

8.2 The standard monomial basis of V

We now define the standard monomial basis of V mentioned above following [57], and prove some lemmas concerning its complexity-theoretic properties.

Let $\bar{G} = GL_m(K)$. Let Z be an $m \times m$ variable matrix. Let $K[Z]$ be the ring generated by the variable entries of Z. Let $K[Z]_d$ denote the degree d part of $K[Z]$. It has commuting left and right actions of $\bar{G}$, where $(\sigma, \sigma') \in \bar{G} \times \bar{G}$ maps $h(Z) \in K[Z]_d$ to $h(\sigma^t Z \sigma')$. For each partition $\lambda : \lambda_1 \geq \cdots \lambda_q > 0$, $q \leq m$, the Weyl module $V_\lambda(\bar{G})$ labelled by λ can be embedded in $K[Z]_d$, $d = |\lambda| = \sum_i \lambda_i$, as follows.

Let (A, B) be a bi-tableau of shape λ. This means both A and B are Young tableau [33] of shape λ such that (1) each box of A or B contains a number in $[m] = \{1, \ldots, m\}$, (2) all columns of A and B are strictly increasing, and (2) all rows are non-decreasing. Let A_i and B_i denote the i-th column of A and B, respectively. With any pair (A_i, B_i) of columns, we associate the minor $Z(A_i, B_i)$ of Z indexed by the row numbers occurring in A_i and the column numbers occurring in B_i. With each bi-tableau (A, B), we associate the monomial in the minors of Z defined by $Z(A, B) := Z(A_1, B_1)Z(A_2, B_2)Z(A_3, B_3) \cdots$. We call such a monomial *standard* of shape λ and degree $d = |\lambda|$. We call a monomial in the minors of Z *non-standard* if it is not standard. It is shown in Doubillet, Rota, and Stein [24] that the standard monomials of degree d form a basis of $K[Z]_d$. We denote this basis of $K[Z]_d$ by $B(Z)_d$.

A standard monomial $Z(A, B)$ is called *canonical* if the column B_i, for each i, just consists of the entries $1, 2, 3, \ldots$ in the increasing order. It is known [57] that, for each partition λ, the subspace of $K[Z]$ spanned by the canonical monomials of shape λ is a representation of G under its left action on $K[Z]$. It is also known [57] that this representation is isomorphic to the Weyl module $V_\lambda(\bar{G})$ of $\bar{G}$, and that the set of canonical monomials of shape λ form its basis. We refer to it as the *standard monomial basis* of $V_\lambda(\bar{G})$, and denote it by $B_\lambda = B_\lambda(\bar{G})$. Each Weyl module $V_\lambda(G)$ of $G = SL_m(K)$ is also a Weyl module of $\bar{G}$ in a natural way. Hence this also specifies the standard monomial basis B_λ of $V_\lambda(G)$.

Fix the standard monomial basis B_λ in each copy of $V_\lambda(G)$ in the complete decomposition of V as in (23). This yields a basis $B(V)$ of V, which we call its *standard monomial basis.* It depends on the choice of the decomposition of V (if the multiplicities are greater than one). But this choice does not matter in what follows.

Lemma 8.4 *(a) Given any nonstandard monomial μ of degree d in the minors of Z, the coefficients of μ in the basis $B(Z)_d$ can be computed in poly(d^{m^2}) time. More strongly, they can be computed by a uniform AC^0-circuit of poly(d^{m^2}) bit-size with oracle access to DET (the determinant function).*

(b) Consider $K[Z]_d$ as a left $\bar{G}$-module, where $g \in \bar{G}$ maps $h(Z)$ to $(g \cdot h)(Z) = h(g^t Z)$. Then, given the specifying label (a bi-tableau) of any basis element $b \in B(Z)_d$ and $g \in GL_m(\mathbb{Q})$, the coefficients of $g \cdot b$ in the basis $B(Z)_d$ can be computed in poly($d^{m^2}, \langle g \rangle$) time, where $\langle g \rangle$ denotes the bit-length of the specification of g. More strongly, they can be computed by a uniform AC^0-circuit of poly($d^{m^2}, \langle g \rangle$) bit-size with oracle access to DET.

(c) Let $V_\lambda(\bar{G})$ be a Weyl module of degree d, and B_λ its standard monomial basis as above. For any basis element $b \in B_\lambda$ specified by a tableau and $g \in GL_m(\mathbb{Q})$, the coefficients of $g \cdot b$ in the basis B_λ can be computed in poly($d^{m^2}, \langle g \rangle$) time. More strongly, they can be computed by a

uniform AC^0-circuit of poly$(d^{m^2}, \langle g \rangle)$ bit-size with oracle access to DET.

When m is constant, the poly(d^{m^2}) bound becomes poly$(d) = O(\text{poly}(n))$.

Proof:

(a) Let $B'(Z)_d$ denote the usual monomial basis of $K[Z]_d$ consisting of the monomials in the entries z_{ij} of Z of total degree d. The cardinality of $B'(Z)_d$ is equal to the number of monomials of degree d in the m^2 variables z_{ij}'s. This number is $\binom{d+m^2-1}{m^2-1} = O(\text{poly}(d^{m^2}))$. The cardinality of $B(Z)_d$ is the same. Let $\mathcal{A}_d$ be the matrix for the change of basis so that:

$$B(Z)_d = \mathcal{A}_d B'(Z)_d, \quad \text{and} \quad B'(Z)_d = \mathcal{A}_d^{-1} B(Z)_d. \tag{27}$$

The matrix $\mathcal{A}_d$ can be computed in poly(d^{m^2}) time. For this, observe that each row of $\mathcal{A}_d$ corresponds to the expansion of a standard monomial $b \in B(Z)_d$ in the usual monomial basis $B'(Z)_d$. Since the number of monomials of degree $\leq d$ in the m^2 variable entries of Z is $O(\text{poly}(d^{m^2}))$ and the degree of b is d, this expansion can be computed by a uniform weakly skew (Section 2.1) circuit of poly(d^{m^2}) bit-size (constructed by induction on d). It follows [60] that it can also be computed fast in parallel by a uniform AC^0-circuit of poly(d^{m^2}) bit-size with oracle access to DET. This yields the representation of b in the basis $B'(Z)_d$. Thus $\mathcal{A}_d$ can be computed by a uniform AC^0-circuit of poly(d^{m^2}) bit-size with oracle access to DET.

Once $\mathcal{A}_d$ has been computed, $\mathcal{A}_d^{-1}$ can also be computed fast in parallel [60] by a uniform AC^0-circuit of poly(d^{m^2}) bit-size with oracle access to DET.

The standard representation in the basis $B(Z)_d$ of any nonstandard monomial $\mu \in K[Z]_d$ in the minors of Z can now be computed fast in parallel as follows. Let $b(\mu)$ and $b'(\mu)$ be the row vectors of the coefficients of μ in the bases $B(Z)_d$ and $B'(Z)_d$, respectively. Clearly $b(\mu) = b'(\mu)\mathcal{A}_d^{-1}$. Expand μ fast in parallel (as we expanded b above) to get its representation $b'(\mu)$. Multiply this on the right by $\mathcal{A}_d^{-1}$ fast in parallel to get $b(\mu)$.

(b) First, we expand $g \cdot b$ fast in parallel (as above) to get its representation in the usual monomial basis $B'(Z)_d$. The representation in $B(Z)_d$ can now be computed fast in parallel by multiplication on the right by $\mathcal{A}_d^{-1}$.

(c) This follows from (b), using the concrete realization of $V_\lambda(G)$ described before, as the G-submodule of $K[Z]_d$, $d = |\lambda|$, spanned by the canonical monomials of shape λ. Q.E.D.

8.3 Encoding generators of $K[V]^G$ by a depth four circuit

We shall deduce Theorem 8.1 from a stronger result (Theorem 8.5) described below, which shows how to encode a set of generators of $K[V]^G$ by a depth four circuit. To state it we need a few definitions.

Let $v = (v_1, \ldots, v_n)$ be the coordinates of V in the standard monomial basis $B(V)$ of V as above. Let $x = (x_1, \ldots, x_n)$ be new variables. Let

$$X = \sum_i x_i v_i \in K[V; x] \tag{28}$$

be a generic affine combination of v_i's. Here $K[V;x]$ denotes the ring obtained by adjoining $x_1,\ldots,x_n$ to $K[V]=K[v_1,\ldots,v_n]$. Then, for any $c>0$,

$$X^c = \sum_{a_1,\ldots,a_n\geq 0:\sum a_i=c} \binom{c}{a_1,\ldots,a_n}(\prod_{i\geq 1} x_i^{a_i})(\prod_{i\geq 1} v_i^{a_i}). \tag{29}$$

Here $\binom{c}{a_1,\ldots,a_n}$ denotes the multinomial coefficient, and the monomials $(\prod_{i\geq 1} v_i^{a_i})$ occurring in this expression form a basis of the subspace $K[V]_c \subseteq K[V]$ of polynomials on V of degree c.

Let $R=R_G : K[V]\to K[V]^G$ denote the Reynolds' operator for G (cf. Section 2.2.1 in [17]). We denote the induced map from $K[V;x]$ to $K[V]^G[x]$ by R as well. Here $K[V]^G[x]$ denotes the ring obtained by adjoining $x_1,\ldots,x_n$ to $K[V]^G$. Now consider a generic invariant

$$R(X^c)(v,x) = \sum_{a_1,\ldots,a_n\geq 0:\sum a_i=c} \binom{c}{a_1,\ldots,a_n} R(\prod_{i\geq 1} v_i^{a_i})(\prod_{i\geq 1} x_i^{a_i}) \in K[V]^G[x]. \tag{30}$$

Since the monomials $(\prod_{i\geq 1} v_i^{a_i})$ in (29) form a basis of $K[V]_c$, it follows from the properties of the Reynold's operator that the elements $R(\prod_{i\geq 1} v_i^{a_i}) \in K[V]^G$ occurring in (30) span the subspace $K[V]_c^G \subseteq K[V]^G$ of invariants of degree c. By Theorem 8.2, the invariants of degree $\leq l = nm^2 d^{2m^2}$ generate $K[V]^G$. Hence, the set

$$F = \{R(\prod_{i\geq 1} v_i^{a_i}) \mid \sum_i a_i = c, 0<c\leq l\} \tag{31}$$

generates $K[V]^G$.

Let $\Delta_3[n,l,k]$ denote the class of diagonal depth three circuits (cf. Section 2.1) over K and the variables $x_1,\ldots,x_n$ with total degree $\leq l$ and top fan-in $\leq k$. The size of any such circuit is $O(knl)$.

Theorem 8.1 follows from the following stronger result.

Theorem 8.5 *Let $N = n^{m^2} d^{m^4}$, and let $l = nm^2 d^{2m^2}$ as in Theorem 8.2. Given n, m, $0 < c \leq l$, and the specification $\langle V, G\rangle$ of V and G as in (24), one can compute in poly(N) time the specification of a depth four circuit $C = C[V,m,c]$ over $\mathbb{Q}$ such that (1) C computes the polynomial $R(X^c)(v,x)$ in $x=(x_1,\ldots,x_n)$ and $v=(v_1,\ldots,v_n)$, and (2) for any fixed $h\in V$, the circuit C_h, obtained by specializing the variables v_i's in C to the coordinates of h in the standard monomial basis $B(V)$ of V, is a diagonal depth three circuit in the class $\Delta_3[n,c,k]$, with $k = O(poly(N))$.*

More strongly, C can be computed by a uniform AC^0-circuit of poly(N) bit-size with oracle access to DET.

Proof strategy: The proof proceeds in four steps: (1) Show that the computation of the Reynolds operator on $K[V;x]$ can be reduced to (a) the computation of the Reynolds operator on the coordinate ring $K[G]$ of G, and (b) the computation of a certain comorphism ψ^* on $K[V;x]$ (defined below) associated with the representation V (cf. Lemma 8.6). (2) Give an efficient algorithm for the computation of the Reynolds operator on $K[G]$, as needed in (1)(a), for constant m (cf. Lemma 8.7). (3) Show that the computation of $\psi^*(X^c)$ can be encoded by a small circuit of constant depth, for constant m (cf. Lemma 8.9). (4) Put (1), (2), and (3) together to construct efficiently a small circuit of depth four for computing $R(X^c)(v, x)$, for constant m.

The following lemma concerning the computation of the Reynolds operator $R = R_G$, $G = SL_m(K)$, addresses the first step in this proof strategy.

Consider the representation morphism $\psi : V \times G \to V$ given by: $(v, \sigma) \to \sigma^{-1}v$. Let $\psi^* : K[V] \to K[V \times G] \cong K[V] \otimes K[G]$ denote the corresponding comorphism. This is defined so that, for any $f \in K[V]$ and $t \in V \times G$,

$$\psi^*(f)(t) = f(\psi(t)). \tag{32}$$

By extending the base from K to $K[x] = K[x_1, \ldots, x_n]$, we get the morphism ψ^* from $K[V; x]$ to $K[V; x] \otimes K[G]$. Given $f \in K[V; x]$, let $\psi^*(f) = \sum_i g_i \otimes h_i$, where $g_i \in K[V; x]$ and $h_i \in K[G]$.

Lemma 8.6 *(cf. Proposition 4.5.9 and Remark 4.5.29 in [17])*

$$R_G(f) = \sum_i g_i R_G(h_i).$$

This reduces the computation of R_G on $K[V; x]$ to (a) the computation of R_G on $K[G]$, and (b) the computation of ψ^*.

Now we address the step (a) above. Since G, as an affine variety, has just one G-orbit (with respect to the left-action of G on itself), it follows that $K[G]^G = K$. Hence, R_G maps $K[G]$ to $K[G]^G = K$. Let $\mathbb{Q}[G]$ denote the coordinate ring of G over $\mathbb{Q}$. Then R_G similarly maps $\mathbb{Q}[G]$ to $\mathbb{Q}[G]^G = \mathbb{Q}$. For the proof of Theorem 8.5, we only need to compute R_G on $\mathbb{Q}[G]$.

Let Z be an $m \times m$ variable matrix. Then $K[G] = K[Z]/J$, where J is the principal ideal generated by $\det(Z) - 1$. Furthermore, by the First Fundamental Theorem of invariant theory [33], $K[Z]^G = K[\det(Z)]$, where $K[Z]$ is considered as a left G-module as in Section 8.2.

The following lemma addresses the second step in the proof strategy, namely, the computation of R_G on $K[G]$.

Lemma 8.7 *Given $g \in \mathbb{Q}[G] \subseteq K[G]$, represented as a polynomial $f \in \mathbb{Q}[Z]$, $R_G(g) \in \mathbb{Q}$ can be computed in poly*$(\deg(f)^{m^2}, \langle f \rangle)$ *time, where $\langle f \rangle$ denotes the total bit-length of the coefficients of f.*

More strongly, $R_G(g)$ can be computed by a uniform AC^0-circuit of poly$(\deg(f)^{m^2}, \langle f \rangle)$ *bit-size with oracle access to DET.*

The computation of R_G on $K[G]$ can be reduced to the computation of R_G on $K[Z]$, considered as a left G-module as in Section 8.2, where R_G maps $K[Z]$ to $K[Z]^G = K[\det(Z)]$. Indeed, if $g \in K[G]$ is represented by $f \in K[Z]$, then $R_G(g) = R_G(f) \pmod{J}$.

Hence, to prove Lemma 8.7, it suffices to show how $R_G(f)$, $f \in \mathbb{Q}[Z]$, can be computed in the stated running time.

Towards this end, we first recall how R_G on $K[Z]$ can be computed using Cayley's Ω process [43]. Here Ω is a differential operator on $K[Z]$ defined as follows. Let $z_{i,j}$'s denote the variable entries of Z. Then, for any $h(Z) \in K[Z]$,

$$\Omega(h(Z)) := \sum_{\pi \in S_m} \operatorname{sign}(\pi) \frac{\partial^m h}{\partial z_{1,\pi_1} \partial z_{2,\pi_2} \cdots \partial z_{m,\pi_m}},$$

where S_m is the symmetric group on m letters.

Lemma 8.8 *(cf. Proposition 4.5.27 in [17]) Suppose $f \in K[Z]$ is homogeneous. If the degree of f is mr, then*

$$R_G(f) = \det(Z)^r \frac{\Omega^r f}{c_{r,m}},$$

where $c_{r,m} = \Omega^r(\det(Z)^r) \in \mathbb{Z}$ (it is non-zero). If the degree of f is not divisible by m, then $R_G(f) = 0$.

If $g \in K[G]$ is represented by $f \in K[Z]$, then $R_G(g) = \frac{\Omega^r f}{c_{r,m}}$, if the degree of f is mr, and $R_G(g) = 0$, if the degree of f is not divisible by m.

Proof of Lemma 8.7: By Lemma 8.8, it suffices to compute $\Omega^r(f)$ and $c_{r,m} = \Omega^r(\det(Z)^r)$ within the stated running time, when $\deg(f) = mr$.

Write $\det(Z)^r = \sum_\alpha a_\alpha \alpha(z_{1,1}, \ldots, z_{m,m})$, where α ranges over the monomials in $z_{i,j}$'s of degree mr, and $a_\alpha \in \mathbb{Z}$. Then

$$\Omega^r = \sum_\alpha a_\alpha \alpha(\frac{\partial}{\partial z_{1,1}}, \ldots, \frac{\partial}{\partial z_{m,m}}).$$

The number of α's here is $\binom{mr+m^2-1}{m^2-1} = O(\operatorname{poly}(\deg(f)^{m^2}))$, when $\deg(f) = mr$, and the bit-length of each a_α is $\operatorname{poly}(m, r) = \operatorname{poly}(\deg(f))$. Hence $\frac{\Omega^r f}{c_{r,m}} \in \mathbb{Q}$, for $f \in \mathbb{Q}[Z]$ of degree mr, can be computed in $\operatorname{poly}(\deg(f)^{m^2}, \langle f \rangle)$ time.

The coefficients a_α's can also be computed fast in parallel by a uniform AC^0-circuit of $\operatorname{poly}(\deg(f)^{m^2})$ bit-size with oracle access to DET, using multi-variate Vandermonde interpolation [88, 87]; cf. also the proof of Lemma 8.9 below for the use of this technique. Hence $\frac{\Omega^r f}{c_{r,m}}$ can also be computed fast in parallel. Q.E.D.

Next we address the third step in the proof strategy, namely, the computation of $\psi^*(X^c)$.

Let $\bar{G} = GL_m(K)$. Then V as in (23) is also a polynomial $\bar{G}$-representation in a natural way so that, as a $\bar{G}$-module:

$$V = \oplus_\lambda m(\lambda) V_\lambda(\bar{G}). \tag{33}$$

Let $u \in \bar{G}$ be a generic (variable) matrix. Let $0 < c \leq l = O(\operatorname{poly}(n, d^{m^2}))$ and $N = n^{m^2} d^{m^4}$ be as in Theorem 8.5. Let $u^{-1} = \mathrm{A}dj(u)/\det(u)$, where $\operatorname{Adj}(u)$ denotes the adjoint of u. Let $u_{i,j}$ denote the (i, j)-th entry of u.

For any $f \in K[V; x]$, let $u \cdot f \in K[V; x]$ denote the result of applying $u \in \bar{G}$ to f, thinking of $K[V; x]$ as a $\bar{G}$-module in the natural way. Formally,

$$(u \cdot f)(w) = f(u^{-1} \cdot w), \tag{34}$$

for all $w \in V$, thinking of f as a polynomial function on V with coefficients in $K(x)$. Here $u^{-1} \cdot w$ denotes the result of applying u^{-1} to w. Let

$$(u \diamond f)(w) := f(Adj(u) \cdot w), \tag{35}$$

for all $w \in V$. If u were a generic matrix of G, instead of $\bar{G}$, then $u \diamond f$ and $u \cdot f$ would coincide.

For $X \in K[V; x]$ as in (28), $u \diamond X$ can be expressed as:

$$u \diamond X = \sum_i x_i (u \diamond v_i) = \sum_i e_i(x, u) v_i, \tag{36}$$

where $e_i \in \mathbb{Q}[x, u]$ is a polynomial in x_j's and the entries $u_{i,j}$'s of u, which is determined by the action of $\bar{G}$ on V. It is linear in x_j's, and has total degree $\leq d(m-1) \leq dm$ in the entries of u. The latter fact follows from (35), since V is a representation of $\bar{G}$ of degree d, and $Adj(u)$ has degree $m-1$ in $u_{i,j}$'s.

By (35), $(u \diamond X^c)(w) = X^c(Adj(u) \cdot w)$, for all $w \in V$. Hence, it follows from (36) that

$$(u \diamond X^c) = (u \diamond X)^c = \sum_\mu \mu \beta_\mu(v, x), \tag{37}$$

where μ ranges over the monomials in $u_{i,j}$'s of total degree at most $dmc \leq dml = O(\text{poly}(n, d^{m^2}))$, and $\beta_\mu(v, x)$ is a polynomial of degree c in $v = (v_1, \ldots, v_n)$ as well as $x = (x_1, \ldots, x_n)$. The number of μ's here is $\leq \binom{dmc+m^2-1}{m^2-1} = O(\text{poly}(N))$.

Thinking of μ's as elements of $K[G]$, $\psi^*(X^c)$, by the definition of ψ^*, cf. (32), equals the r.h.s. of (37). This is because $u \cdot X^c$ and $u \diamond X^c$ coincide if we think of u as a generic element of G (rather than $\bar{G}$), whence $\det(u) = 1$.

Thus:

$$\psi^*(X^c) = \sum_\mu \mu \beta_\mu(v, x), \tag{38}$$

where μ and β_μ are as in (37).

Hence, to encode $\psi^*(X^c)$ efficiently by a circuit, it suffices to encode $\beta_\mu(v, x)$'s efficiently by a circuit. This is done in the following result.

Lemma 8.9 *Let $N = n^{m^2} d^{m^4}$. Then, given n, m, d, c as above, and the specification $\langle V, G \rangle$ of V and G as in (24), one can compute in poly(N) time, and more strongly, by a uniform AC^0-circuit of poly(N) bit-size with oracle access to DET, the specification of a circuit C' over $\mathbb{Q}$ of poly(N) bit-size on the input variables $v_1, \ldots, v_n$ and $x_1, \ldots, x_n$, and with multiple outputs that compute the polynomials $\beta_\mu(v, x)$'s in (37). The top (output) gates of C' are all addition gates. Furthermore, for any fixed $h \in V$, the circuit C'_h obtained from C' by specializing the variables v_i's to the coordinates of h (in the standard monomial basis of V) is a diagonal depth three circuit with multiple outputs in the class $\Delta_3[n, c, e]$, $e =$ poly(N). By this, we mean that the sub-circuit of C' below each output gate is in $\Delta_3[n, c, e]$.*

Proof: We cannot compute $\beta_\mu(v,x)$ in (37) by expanding $(u \diamond X)^c$ as a polynomial in x, u, and v, since the number of terms in this expansion is exponential in n. But we can compute it by a constant depth circuit, by evaluating $(u \diamond X)^c$ at several values of u and then performing multivariate Vandermonde interpolation in the spirit of Strassen [88, 87], as follows.

First, we show how to construct, for any fixed $g \in M_m(\mathbb{Q})$, a constant depth circuit A_g that computes the polynomial in v and x given by

$$g \diamond X^c = (g \diamond X)^c = (\sum_i e_i(x,g)v_i)^c, \tag{39}$$

where $e_i(x,g)$ is a linear form in x that is obtained by evaluating $e_i(x,u)$ in (36) at $u = g$. By the definition of $\diamond$, cf. (35), this is well defined at any element of $M_m(\mathbb{Q})$ (not just $GL_m(\mathbb{Q})$).

Towards this end, we first construct a depth two circuit A'_g, with an addition gate at the top, that computes the quadratic polynomial in v and x

$$g \diamond X = \sum_i e_i(x,g)v_i, \tag{40}$$

obtained by instantiating (36) at $u = g$. Recall that $v_1, \ldots, v_n$ are the coordinates of V corresponding to the standard monomial basis $B(V)$ of V compatible with the decomposition (33). Hence, using Lemma 8.4 (c), the coefficients of the linear form $e_i(x,g)$, for given $g \in M_m(\mathbb{Q})$, can be computed in poly$(n, d^{m^2}, \langle g \rangle)$ time, and more strongly, by a uniform AC^0-circuit of poly$(n, d^{m^2}, \langle g \rangle)$ bit-size with oracle access to DET. After this, the specification of A'_g can be computed in poly$(n, d^{m^2}, \langle g \rangle)$ time, and more strongly, by a uniform AC^0-circuit of poly$(n, d^{m^2}, \langle g \rangle)$ bit-size with oracle access to DET.

Next, we construct A_g, with a single multiplication (powering) gate of fan-in c at its top, that computes the c-th power of $g \diamond X$ computed by the output node of A'_g. The polynomial $A_g(v,x)$ computed by A_g is $(g \diamond X^c)(v,x)$. Furthermore, for any fixed $h \in V$, the circuit obtained by instantiating A_g at $v = h$ is a depth two circuit with a multiplication (powering) gate at the top.

Next, we show how to efficiently construct a circuit C' for computing the polynomials β_μ's, using A_g's for several g's of poly(N) bit-length.

Let $u_{i,j}$ denote the (i,j)-th entry of u as before. Let e be the number of monomials μ's in $u_{i,j}$'s with the degree in each $u_{i,j}$ at most $d' := dmc$. Then $e = O((dmc)^{m^2}) = O(\text{poly}(N))$, since $c \leq l = \text{poly}(n, d^{m^2})$. Order these monomials lexicographically. For $r \leq e$, let μ_r denote the r-th monomial in this order. Choose $m \times m$ non-negative integer matrices $g_1, \cdots, g_e$ such that (1) the $e \times e$ matrix $B = [\mu_r(g_s)]$, whose (s,r)-th entry, for $s, r \leq e$, is $\mu_r(g_s)$, is non-singular, and (2) every entry of each g_s is $\leq d'$. We can choose such g_s's explicitly so that B is a multivariate Vandermonde matrix as described in Section 3.9 in [63]. Specifically, let $E = \{0, \ldots, d'\}^{m^2}$ be the set of e integral points in $\mathbb{Z}^{m^2}$. Order E lexicographically. Let g_s be the s-th point in E, interpreted as an $m \times m$ matrix. Then B is a non-singular multivariate Vandermonde matrix (cf. Sections 3.9 and 3.11 in [63]). It can be computed in poly(N) time, and more strongly, by a uniform AC^0-circuit of poly(N) bit-size. Its inverse B^{-1} can be computed by a uniform AC^0-circuit of poly(N) bit-size with oracle access to DET.

Let $\bar{\beta}$ denote the column-vector of length e whose r-th entry, for $r \leq e$, is $\beta_{\mu_r}(v,x)$ (which we define to be zero if the total degree of μ_r exceeds $d' = dmc$). Let $\bar{A}$ denote the column vector of length e whose s-the entry, for $s \leq e$, is $A_{g_s}(v,x) = (g_s \diamond X^c)(v,x)$. Then, by (37),

$$\bar{A} = B\bar{\beta}, \quad \text{and} \quad \bar{\beta} = B^{-1}\bar{A}.$$

Using the second equation here, we can construct a constant depth circuit C' (with multiple outputs) for computing the entries of $\bar{\beta}$, using the constant depth circuits A_{g_s}'s constructed above. Each output gate of C' is an addition gate with fan-in $e = \text{poly}(N)$. Each gate at the second level from the top is a powering gate with fan-in c, because the top gate of each A_{g_s} is the powering gate with fan-in c. For a fixed $h \in V$, the circuit C'_h obtained by instantiating C' at $v = h$ is thus a diagonal depth three circuit with multiple outputs in the class $\Delta_3[n,c,e]$.

Since A_{g_s}, for every $g_s \in E$, and B^{-1} can be constructed in poly(N) time, the construction of C' takes poly(N) time. More strongly, it can be computed by a uniform AC^0-circuit of poly(N) bit-size with oracle access to DET. Q.E.D.

In the final step, we put everything together to construct the circuit $C = C[V,m,c]$ for computing $R(X^c)$, as required in Theorem 8.5, given n, d, m, c, and the specification $\langle V, G \rangle$ of V and G as in (24).

By Lemma 8.6 and (38),

$$R(X^c)(v,x) = \sum_{\mu} R_G(\mu)\beta_\mu(v,x).$$

Here $R_G(\mu)$ is a rational number that can be computed in poly(N) time using Lemma 8.7, since the degree of μ is poly(n, d^{m^2}). Let C' be the circuit for computing β_μ's as in Lemma 8.9. The circuit C is obtained by adding a single addition gate that performs linear combinations of the various output nodes of C' computing β_μ's, the coefficients in the linear combination being the poly(N)-time-computable rational numbers $R_G(\mu)$'s. Since the top gates of C' are addition gates with fan-in e, we can ensure, by merging the addition gates in the top two levels, that the depth of C is the same as that of C'. The top gate of C after this merge is an addition gate with fan-in $k = e^2 = O(\text{poly}(N))$.

Given n, d, m, c, and $\langle V, G \rangle$ as in (24), the specification of C' can be computed in poly(N) time by Lemma 8.9. After this, the specification of the circuit C as above can also be computed in poly(N) time. More strongly, it can be computed by a uniform AC^0-circuit of poly(N) bit-size with oracle access to DET.

For any fixed $h \in V$, the circuit C_h, obtained by specializing the variables v_i's in C to the coordinates of h, is a diagonal depth three circuit in the class $\Delta_3[n,c,k]$, with $k = e^2 = O(\text{poly}(N))$. This is because, by Lemma 8.9, C'_h is a diagonal depth three circuit with multiple outputs in the class $\Delta_3[n,c,e]$, $e = O(\text{poly}(N))$.

This completes the proof of Theorem 8.5.

More generally:

Theorem 8.10 *The categorical quotient V/G is quasi-explicit (cf. Definition 5.1 (d)) when $m = O(\sqrt{d})$.*

Proof: By Lemma 8.3 (b), l and N in Theorems 8.2 and 8.5 are $O(2^{\text{polylog}(n)})$, if $m = O(\sqrt{d})$. The result follows from Theorem 8.5, in conjunction with this fact.Q.E.D.

9 NNL for the general ring of invariants

In this section we prove Theorem 1.6.

Let V be a finite dimensional representation of $G = SL_m(K)$. Let $K[V]^G \subseteq K[V]$ be the ring of invariants, and $V/G := spec(K[V]^G)$, the categorical quotient [75]. We assume that V/G and $K[V]^G$ are specified succinctly by the tuple $\langle V, G\rangle$ in (24). The bit-length of this succinct specification is $O(n + m)$.

Since V/G is explicit when m is constant (cf. Theorem 8.1), we can also specify it succinctly in this case, as per the general definition of an explicit variety (Definition 5.1), by the circuits $C[v, m, c]$'s in Theorem 8.5. This specification is equivalent when m is constant, because, given $\langle V, G\rangle$, one can compute the circuits $C[V, m, c]$'s in Theorem 8.5 in poly(n, m) time.

If V/G is explicit (cf. Conjecture 5.3), the general definition of an s.s.o.p. for explicit varieties (Definition 5.6) specializes to the following concrete definition.

Definition 9.1 *(a) A set $S \subseteq K[V]^G$ is an* s.s.o.p. (small system of parameters) *for $K[V]^G$ if (1) $K[V]^G$ is integral over the subring generated by S, (2) the cardinality of S is poly(n, m), (3) every invariant in S is homogeneous of poly(n, m) degree, and (4) every $s \in S$ has a small specification in the form of a circuit (Section 2.1) of poly(n, m) bit-length over $\mathbb{Q}$ and the coordinates $v_1, \ldots, v_n$ of V in the standard monomial basis.*

(b) A set $S \subseteq K[V]^G$ is an e.s.o.p. (explicit system of parameters) *for $K[V]^G$ if (1) S is an s.s.o.p. for $K[V]^G$, and (2) the specification of S, consisting of a circuit as above for each $s \in S$, can be computed in poly(n, m) time, given the specification $\langle V, G\rangle$ as in (24).*

If V/G is strongly explicit then, by convention, we assume that a small specification as in (a) (4) for each element of $s \in S$ is a weakly skew circuit (cf. Section 2.1).

(c) Quasi-s.s.o.p. and quasi-e.s.o.p. are defined by replacing poly(n, m) by $2^{polylog(n,m)}$.

(d) S.s.o.p., e.s.o.p., and the related notions without degree restrictions are defined by dropping the degree requirement in (a) (3).

(e) We call an s.s.o.p. or an e.s.o.p. separating if S in (a) is separating [17] (cf. Section 7).

By *the problem NNL for $K[V]^G$*, we mean the problem of constructing an s.s.o.p. for $K[V]^G$, given $\langle V, G\rangle$. By *the strong form of NNL*, we mean the problem of constructing a separating s.s.o.p.

For constant m, define a *separating near-e.s.o.p.* for $K[V]^G$ by replacing poly(n, m) in the definition of a separating e.s.o.p. above by $O(n^{O(\log \log n)})$.

We prove the following stronger form of Theorem 1.6 in this section.

Theorem 9.2 *There exists a separating near-e.s.o.p. for $K[V]^G$, if m is constant, and a separating quasi-e.s.o.p., if $m = O(\sqrt{d})$.*

9.1 An EXPSPACE-algorithm for constructing an h.s.o.p.

Before we turn to this goal, we begin with the following result for the construction of an h.s.o.p.

Proposition 9.3 *The problem of constructing an h.s.o.p. (cf. Definition 3.3) for $K[V]^G$ belongs to EXPSPACE for any m, not necessarily constant.*

When m is constant, this result follows from Theorem 5.5, since V/G is then explicit (Theorem 8.1). For general m, it cannot be deduced from Theorem 5.5, since V/G in general is not yet known be explicit, though it is conjectured to be so; cf. Conjecture 5.3.

Proof: Let $F = \{f_1, \ldots, f_t\}$ be the set of generators of $K[V]^G$ as in (25), and $\pi_{V/G}$ the morphism from V to K^t based on F as there. Here t is the dimension of $K[V]^G_{\leq l}$, with l as in Theorem 8.2. This can be exponential in n even when m is constant; cf. the discussion before (25).

Using this embedding $\pi_{V/G}$ of V/G and Gröbner basis theory, we can compute the equations of $V/G \subseteq K^t$ in work-space that is exponential in $\dim(V/G) \leq n$, polynomial in the dimension t of the ambient space, and poly-logarithmic in the maximum degree of the elements in F; cf. Theorem 1 in [61]. This work-space requirement is single exponential in n and m (since $d \leq n$ by Lemma 8.3).

Applying Gröbner basis theory [61] again to these equations of V/G, we compute an h.s.o.p. for $K[V]^G$. The work-space requirement of this algorithm is also exponential in n and m; cf. the proof of Theorem 4.1. Q.E.D.

If we insist on an h.s.o.p., then Proposition 9.3 is the best that we can do at present. But if we are willing to settle for an s.s.o.p. (which need not have the optimal cardinality) instead of an h.s.o.p., then Theorem 9.2 shows that a near-s.s.o.p. for $K[V]^G$ can be constructed in near-poly(n) time if m is constant, and more generally, a quasi-s.s.o.p. for $K[V]^G$ can be constructed in quasi-poly(n) time if $m = O(\sqrt{d})$.

We now turn to the proof of Theorem 9.2.

9.2 A Monte Carlo algorithm

We begin with the following result, which gives an efficient Monte Carlo algorithm for constructing a separating s.s.o.p., when m is constant.

Theorem 9.4 *Suppose m is constant. Then a separating s.s.o.p. for $K[V]^G$ can be constructed by a poly(n)-time Monte Carlo algorithm that is correct with a high probability. In particular, a separating s.s.o.p. for $K[V]^G$ exists.*

Proof: By Theorem 8.1, V/G is explicit when m is constant. Hence the result follows from Theorem 5.8. Q.E.D.

9.3 Reduction of NNL to the standard black-box identity testing for diagonal depth three circuits

The goal now is to derandomize this algorithm.

Since V/G is strongly explicit (cf. Theorem 8.1) when m is constant, and the image of the map $\pi_{V/G}$ in (25) is closed (cf. Theorem 5.4 (a)), it follows from Theorem 5.13 and Remark 3 thereafter that the algorithm in Theorem 9.4 can be derandomized, assuming the standard black-box derandomization hypothesis for symbolic determinant identity testing. The following result shows that this derandomization is, in fact, possible assuming a much weaker hypothesis, namely, the standard black-box derandomization hypothesis for polynomial identity testing for diagonal depth three circuits (cf. Section 2.1). This hypothesis is that a hitting set against diagonal depth three circuits on n variables with degree $\leq e$ and top fan-in $\leq k$ can be computed in poly(s) time, where $s = O(nek)$ is the size of such circuits. The parallel black-box derandomization hypothesis in this context is that such a hitting set can be computed by a uniform AC^0-circuit of poly(s) bit-size with oracle access to DET. It is known that such a hitting set can be computed by a uniform AC^0-circuit of quasi-poly(s) bit-size [2].

Theorem 9.5 *Suppose the standard black-box derandomization hypothesis for polynomial identity testing for diagonal depth three circuits over K holds. Then $K[V]^G$ has a separating e.s.o.p. if m is constant.*

Specifically, there then exists a set $S \subseteq K[V]^G$ of poly(N) homogeneous invariants, $N = n^{m^2} d^{m^4}$, such that (1) S is separating, and hence (cf. Theorem 7.7) $K[V]^G$ is integral over its subring generated by S, (2) every invariant in S has poly(N) degree, (3) every $s \in S$ has a weakly skew (Section 2.1) circuit over $\mathbb{Q}$ and the coordinates $v_1, \ldots, v_n$ of V of poly(N) bit-length, and (4) the specification of S, consisting of such a weakly-skew circuit for every invariant in S, can be computed in poly(N) time.

Assuming the parallel black-box derandomization hypothesis for polynomial identity testing for diagonal depth three circuits, the specification of S can be computed by a uniform AC^0-circuit of poly(N) bit-size with oracle access to DET.

Proof: Let N be as above. Let $k = O(\text{poly}(N))$ and l be as in Theorem 8.5. Consider the class $\Delta_3[n, l, 2k]$ (cf. Section 8.3) of diagonal depth three circuits over n variables, with total degree $\leq l$, and top fan-in $\leq 2k$.

By our black-box derandomization hypothesis for diagonal depth three circuits over K, there exists a hitting set T against $\Delta_3[n, l, 2k]$ that can be computed in $\text{poly}(n, k, l) = \text{poly}(N)$ time. Assuming the parallel black-box derandomization hypothesis, T can be computed by a uniform AC^0-circuit of poly(N) bit-size.

Fix such a T. By the definition of a hitting set, for any circuit $D \in \Delta_3[n, l, 2k]$ such that $D(x)$, $x = (x_1, \ldots, x_n)$, is not an identically zero polynomial, there exists $b \in T$ such that $D(b) \neq 0$.

For any $b \in T$ and $0 < c \leq l$, define the invariant

$$r_{b,c} := R(X^c)(v, b) \in K[V]^G,$$

where $R(X^c)$ is as in (30). Let

$$S = \{r_{b,c} \mid b \in T, 0 < c \leq l\} \subseteq K[V]^G. \qquad (41)$$

The elements of S are homogeneous polynomials in v of degree $\leq l$, which is poly(n) if m is constant.

Claim 9.6 *The set S is separating.*

Proof: Let $w_1, \ldots, w_n$ be auxiliary variables. For every $c \leq l$, define the symbolic difference

$$\tilde{R}^c(x, v, w) = R(X^c)(v, x) - R(X^c)(w, x),$$

where $R(X^c)(w, x)$ is defined just like $R(X^c)(v, x)$, substituting w for v. Suppose $e, f \in V$ are two points such that $r(e) \neq r(f)$ for some $r \in K[V]^G$. It follows that some generator in the set F in (31) assumes different values at e and f. From (30), it follows that, for some $c \leq l$, $\tilde{R}^c(x, e, f)$ is not an identically zero polynomial in x. By Theorem 8.5, $R(X^c)(e, x)$ is computed by a diagonal depth three circuit in the class $\Delta_3[n, l, k]$. Hence $\tilde{R}^c(x, e, f)$ is computed by a diagonal depth three circuit in the class $\Delta_3[n, l, 2k]$. Since T is a hitting set against such circuits, it follows that, for some $b \in T$, $\tilde{R}^c(b, e, f) \neq 0$. That is,

$$r_{b,c}(e) = R(X^c)(e, b) \neq R(X^c)(f, b) = r_{b,c}(f).$$

Thus S is separating. This proves the claim.

It follows from the claim and Theorem 7.7 that $K[V]^G$ is integral over the subring generated by S.

For any $b \in T$ and $0 < c \leq l$, let $D_{b,c}$ be the circuit obtained by specializing the circuit $C[V, m, c]$ in Theorem 8.5 at $x = b$. Then $D_{b,c}$ computes $r_{b,c} = R(X^c)(v, b)$ as a polynomial in v. We specify S by giving, for every invariant $r_{b,c} \in S$, the specification of $D_{b,c}$. By Theorem 8.5, the circuit $D_{b,c}$ has constant depth and poly(N) bit-size. Hence, it can also be specified by a weakly skew circuit of poly(N) bit-size.

By our black-box derandomization hypothesis, the specification of T can be computed in poly(N) time. Once T is computed, using Theorem 8.5, we can compute in poly(N) time, for each $b \in T$ and $c \leq l$, the specification of the circuit $D_{b,c}$ computing the invariant $r_{b,c} \in S$. Thus the specification of S in the form of a circuit $D_{b,c}$ for each $r_{b,c}$, or the corresponding weakly skew circuit, can be computed in poly(N) time. Hence, S is a separating e.s.o.p.

Assuming the parallel black-box derandomization hypothesis, T, and hence S, can be computed by a uniform AC^0-circuit of poly(N) bit-size with oracle access to DET. Q.E.D.

9.4 Proof of Theorem 9.2

Proof: By Forbes, Saptharishi, and Shpilka [29], a hitting set against diagonal depth three circuits of size $\leq s$ can be computed in $O(s^{O(\log \log s)})$ time. The result follows from the proof of Theorem 9.5 in conjunction with this fact; we also need Lemma 8.3, if $m = O(\sqrt{d})$. Q.E.D.

9.5 General m

The following is the current best result for general m.

Theorem 9.7 *Let V be a finite dimensional representation of $G = SL_m(K)$. Suppose V/G is explicit (cf. Definition 5.2). Then $K[V]^G$ has a separating e.s.o.p., assuming the standard black-box derandomization hypothesis for low-degree polynomial identity testing over K*

Proof: The proof is similar to that of Theorem 9.5, using the assumed explicitness of V/G in place of Theorem 8.5, and the black-box derandomization hypothesis for low-degree polynomial identity testing in place of the black-box derandomization hypothesis for diagonal depth three circuits. Q.E.D.

Remark 1: If V/G is strongly explicit (cf. Definition 5.2), then it follows similarly that $K[V]^G$ has a separating e.s.o.p., assuming the standard black-box derandomization hypothesis for symbolic determinant identity testing. If V/G is explicit without any degree restrictions, then one has to assume instead the black-box derandomization hypothesis for polynomial identity testing without any degree restrictions.

Remark 2: All these results (and Theorems 9.8 (a), and 9.9 (b) below) also hold assuming explicitness of V/G in the relaxed sense (cf. Definition 5.2 (e)).

Remark 3: The derandomization hypothesis in Theorem 9.7 can be traded, up to a quasi-prefix, with the hardness hypothesis in Theorem 2.1.

We also note down a consequence of the proof of Theorem 9.7.

Theorem 9.8 *Let V be a finite dimensional representation of $G = SL_m(K)$. Then:*

(a) The problem of deciding if the closures of the G-orbits of two rational points in V intersect, and finding an invariant separating the two if they do not, belongs to co-RNC, if V/G is explicit. It is in NC assuming, in addition, the white-box derandomization hypothesis [49] for low degree arithmetic circuits over $\mathbb{Q}$ [6].

(b) The problem belongs to P, if m is constant.

(c) It belongs to DET $\subseteq$ NC, for constant m, if we do not ask for a separating invariant if the closures do not intersect.

Proof: (a): The proof of the first statement is similar to that of Theorem 7.8, using the assumed explicitness of V/G in place of Theorem 6.1. The second statement is implicit in the proof of the first statement.

(b) and (c): Suppose m is constant. Using Theorem 8.5 in place of Theorem 6.1 in the proof of Theorem 7.8, we get a co-RNC-algorithm for the problem. This algorithm only uses white-box (cf. Section 2.2) polynomial identity testing for diagonal depth three circuits. It can be derandomized using the DET-algorithm for this test, which follows from Raz and Shpilka [81], Arvind et al. [4], and Saxena [83]. This yields a DET-algorithm as stated in (c) that, however, does not return a separating invariant if the orbit-closures do not intersect. To get a separating invariant if the closures intersect, as needed in (b), we use instead a polynomial time algorithm for white-box polynomial identity testing for diagonal depth three circuits that returns a witness input if the polynomial computed by the circuit is not identically zero. Such an algorithm can be obtained by combining [81, 4, 83] with a proof technique in [2] (cf. Appendix A therein). Using this witness input, a separating invariant can be constructed if the orbit-closures do not intersect; cf. the proof of Theorem 7.8. Q.E.D.

[6] This hypothesis is that, given a low-degree arithmetic circuit over $\mathbb{Q}$, one can decide in polynomial time if the circuit evaluates a non-zero polynomial, and if so, construct in the same time an input on which the evaluation is non-zero.

Remark 3: Theorem 9.8 (c) can also be proved without representation theory as follows. The orbit closures of two points $v, w \in V$ intersect iff $\forall \epsilon \in \mathbb{R}\ \exists g, h \in SL_m(\mathbb{C}) : ||g \cdot v - h \cdot w||_2 \leq \epsilon^2$, where $||\ ||_2$ denotes the L_2 norm on V. If m is constant, this can be checked in polynomial time by the algorithm in [5] for quantifier elimination in the theory of reals (which can also be parallelized). This proof does not return an invariant separating the two orbit closures if they do not intersect, as in Theorem 9.8 (a) and (b). This is a serious limitation in the context of invariant theory. Most importantly, this proof is based on an inherently white-box technique, which does not work in the black-box setting of Theorem 1.6.

Remark 4: Theorem 9.8 (a) implies that the problem of deciding if a given rational point in V belongs to the null cone [75] of the G-action belongs to co-RNC, if V/G is explicit. It belongs to NC assuming, in addition, the white-box derandomization hypothesis for low degree arithmetic circuits over $\mathbb{Q}$.

9.6 Generalization to reductive algebraic groups

The preceding results for SL_m can be generalized to other reductive algebraic groups as follows.

Let K be algebraically closed field of characteristic zero. Let G be a connected, reductive, algebraic group over K, specified by its root datum [44, 62]. Let V be a finite dimensional rational representation of G. Given any highest weight λ of G, let $V_\lambda(G)$ denote the associated irreducible representation of G [33]. We specify V, as in (24), by giving (in unary) $n = \dim(V)$ and the multiplicities of $V_\lambda(G)$'s that occur with non-zero multiplicities in V.

Theorem 9.9 *Let V and G be as above. Then:*

(a) Analogues of Theorems 9.2, 9.5, and 9.8 (b) hold, when $\dim(G)$ *is constant*

(b) Analogues of Theorems 9.7 and 9.8 (a) hold, without any restriction on $\dim(G)$*. In particular, if V/G is explicit, $K[V]^G$ has a separating e.s.o.p., assuming the standard black-box derandomization hypothesis for low-degree polynomial identity testing over K.*

This can be proved by extending the proof for SL_m; cf. the preliminary version [69] for details.

It may be conjectured that, for any finite dimensional representation V of a connected, reductive, algebraic group G in characteristic zero, with the specification of V and G as above, V/G is explicit, if $K[V]^G$ has a set of generators of poly$(n, \dim(G))$ degree, and is explicit without any degree restrictions, in general (cf. Definition 5.2).

More generally, let V be a finite dimensional representation of any reductive, possibly disconnected, algebraic group G over an algebraically closed field of any characteristic [7]. Then it

[7] Here, we assume that V and G are specified as follows, since the direct-sum decomposition as in (23) need not hold in positive characteristic. Let $G_0 \subseteq G$ be the connected component of the identity, specified by its root datum [44, 62], and T its maximal torus. We specify V by giving, in its fixed basis, the representation matrices for: (1) one-parameter multiplicative subgroups generating T, (2) one-parameter additive subgroups U_α's associated with the roots α's of G_0 with respect T (cf. Section 26.3 in [44]), and (3) a set of elements $a_1, \ldots, a_l \in G \setminus G_0$, which together with T and U_α's generate G. The entries of the representation matrix of a one-parameter subgroup are assumed to be rational functions of the parameter with coefficients in a finite extension of $\mathbb{Q}$, if the characteristic is zero, or F_p, if the characteristic is $p > 0$.

may be conjectured that V/G is explicit in the relaxed sense, without any degree restrictions, in general; cf. Definition 5.2 (e). It may be conjectured to be explicit in the relaxed sense, with the usual low degree restrictions, if there is an upper bound on the degrees of generators or separating invariants for $K[V]^G$ that is polynomial in the bit-length of the succinct specification of V and G. This is the case, for example, when G is finite and V is its permutation representation [34]. The conjecture is proved in the next section in any characteristic for $G = SL_m(K)$ and $V = M_m(K)^r$, with the adjoint action of G (cf. Theorem 10.7). Analogues of Theorems 9.7 and 9.8 (a) hold for any finite dimensional representation V of any reductive group G in any characteristic, assuming that V/G is explicit in the relaxed sense. If V/G is explicit in the relaxed sense without any degree restrictions (as conjectured in general), then analogues of Theorems 9.7 and 9.8 (a) still hold, replacing low degree polynomial identity testing with general polynomial identity testing without any degree restrictions. These results can be proved by replacing Theorem 5.4 with its generalization in [75] for arbitrary reductive groups in any characteristic.

In view of the results and arguments above, the strong form of NNL for $K[V]^G$, for any finite dimensional representation V of any reductive group G in any characteristic, may be conjectured to be in P, along with the G-orbit-closures-intersection and the null cone membership problems (cf. Theorem 9.8 and Remark 4 thereafter).

10 Extensions

In this section we extend the results in Sections 6 and 7 to arbitrary characteristics (cf. Section 10.1), and to quivers (cf. Section 10.2). We then deduce their implications for parametrization of closed orbits in representations of reductive groups (cf. Section 10.3), and for parametrization of semi-simple representations of finitely generated algebras (cf. Section 10.4). We also extend the results in Sections 4 and 5 to large enough positive characteristics (cf. Section 10.5). Henceforth, K will denote an algebraically closed field of any characteristic p.

10.1 Matrix invariants in arbitrary characteristic

First, we prove Theorem 1.4 in arbitrary characteristic. It follows from the following stronger result.

Let $V = M_m(K)^r$, with the adjoint action of $G = SL_m(K)$. Separating s.s.o.p. and e.s.o.p. for $K[V]^G$, and the black-box derandomization hypothesis for low-degree circuits are defined as in characteristic zero (cf. Sections 7 and 2.2). If the characteristic p is positive, the constants in the circuits specifying the s.s.o.p. and the entries of the elements of the hitting set against low-degree circuits are assumed to be in F_{p^l}, the finite field with p^l elements, with $l = O(\log(m))$.

Theorem 10.1 *Let V and G be as above.*

(a) Suppose $p \notin [2, \lfloor m/2 \rfloor]$. Then a separating e.s.o.p. exists for $K[V]^G$, assuming the standard black-box derandomization hypothesis for polynomial identity testing for read-once oblivious algebraic branching programs (cf. Section 2.1). A separating quasi-e.s.o.p. exists unconditionally.

(b) A separating e.s.o.p. exists for $K[V]^G$ for any p, assuming the standard black-box derandomization hypothesis for symbolic determinant identity testing over K.

(c) The problem of deciding if the closures of the G-orbits of two rational points in V intersect belongs to co-RDET $\subseteq$ co-RNC for any p. It belongs to NC if $p \notin [2, \lfloor m/2 \rfloor]$. By a rational point in V, when $p > 0$, we mean a point whose coefficients belong to a finite extension of F_p.

The known upper bound [21] on the degrees of the generators in the First Fundamental Theorem for matrix invariants in positive characteristic in Donkin [23] is exponential in m, unlike the polynomial bound in the First Fundamental Theorem for matrix invariants in characteristic zero (cf. Theorem 6.2). Hence the proof of Theorem 7.6 cannot be extended to arbitrary characteristic using Donkin [23] in place of Procesi and Razmyslov [80, 82]. But, as we shall see below, the proof can be extended using the following geometric alternative to Theorem 6.2 in arbitrary characteristic.

Let $U = (U_1, \ldots, U_r)$ denote an r-tuple of variable $m \times m$ matrices as in Section 6. Identify $K[V]$ with the ring $K[U] = K[U_1, \ldots, U_r]$ generated by the variable entries of U_i's. Given any word $\alpha \in [r]^*$, define $T_\alpha(U) \in K[V]^G$ as in (15). For any $m \times m$ matrix X, let $c_i(X)$ denote the i-th coefficient of its characteristic polynomial, so that

$$\det(\lambda I - X) = \lambda^m - c_1(X)\lambda^{m-1} + \cdots + (-1)^m c_m(X) I.$$

Define a separating set $S \subseteq K[V]^G$ in arbitrary characteristic p just as it was defined in characteristic zero in Section 7.

Theorem 10.2 (Geometric First Fundamental Theorem in arbitrary characteristic) *(a) The set $\{c_i(U_j)\} \cup \{T_\alpha(U)\} \subseteq K[V]^G$, where $\lfloor m/2 \rfloor < i \leq m$, $1 \leq j \leq r$, and $\alpha \in [r]^*$ ranges over all words of length $\leq m^3$, is separating, if $p \notin [2, \lfloor m/2 \rfloor]$.*

(b) The set $\{c_{i,\alpha}(U) \mid 0 \leq i \leq m\} \subseteq K[V]^G$, where $\alpha = i_1 i_2 \cdots \in [r]^*$ ranges over all words of length $\leq m^2$, and $c_{i,\alpha}(U) = c_i(U_{i_1} U_{i_2} \cdots)$, is separating for any p.

In characteristic zero, this result follows from Theorem 6.2, letting α range over the words of length $\leq m^2$. For a proof in arbitrary characteristic, we need the following two results.

Let $\hat{R} = K\langle U_1, \ldots, U_r \rangle$ be the free non-commutative algebra over K generated by the r matrix-variables $U_1, \ldots, U_r$ (not the rm^2 variable entries of U_i's). Given any $A = (A_1, \ldots, A_r) \in M_m(K)^r$, let $\rho_A : \hat{R} \to M_m(K)$ denote the m-dimensional representation of $\hat{R}$ given by $U_i \to A_i$. Clearly, two tuples $A, B \in M_m(K)^r$ belong to the same G-orbit iff ρ_A and ρ_B are isomorphic representations. We say that $A \in M_m(K)^r$ is semi-simple if ρ_A is a semi-simple representation of $\hat{R}$.

Theorem 10.3 (cf. Theorem 4.1 in King [54]) *The G-orbit of $A \in M_m(K)^r$ is closed iff A is semi-simple.*

Let R be any finite-dimensional algebra over K. Let $\rho : R \to M_m(K)$ be an m-dimensional representation of R. For any $r \in R$, let $\chi_\rho(r)$ denote the characteristic polynomial of $\rho(r)$. Let $Q \subseteq R$ be any subset that spans R over K.

Theorem 10.4 (Brauer and Nesbitt) *(cf. [8], Theorem 5.7 in [26] and its proof) Two finite dimensional semi-simple representations ρ and ρ' of R are isomorphic iff $\chi_\rho(q) = \chi_{\rho'}(q)$ for all $q \in Q$. If R is not a finite dimensional algebra, then the same statement also holds if $\rho(Q)$ and $\rho'(Q)$ span $\rho(R)$ and $\rho'(R)$, respectively.*

This follows from the proof of Theorem 5.7 in [26].

Proof of Theorem 10.2:

(a) By the generalization of Theorem 5.4 (d) to arbitrary characteristic (cf. Theorem 1.1 in [75]), it suffices to show that the set $\{c_i(U_j)\} \cup \{T_\alpha(U)\}$ of invariants in (a) separates closed G-orbits in $M_m(K)^r$, i.e., given two distinct closed G-orbits, there exists an invariant in the set that assumes different values on the orbits.

By Theorem 10.3, $A \in M_m(K)^r$ has a closed G-orbit iff A is semi-simple. By definition, this is so iff the m-dimensional representation ρ_A of $\hat{R} = K\langle U_1, \ldots, U_r \rangle$ given by $U_i \to A_i$ is semi-simple.

Furthermore, two semi-simple tuples $A, B \in M_m(K)^r$ are in the same (closed) G-orbit iff the two representations ρ_A and ρ_B of $\hat{R}$ are isomorphic. Let $S \subseteq [r]^*$ be the subset of words of length $\leq m^3$. It suffices to show that, given any two semi-simple $A, B \in M_m(K)^r$ with $\rho_A \not\cong \rho_B$, there exists an $\alpha \in S$ such that $T_\alpha(A) \neq T_\alpha(B)$, or there exist an i, with $\lfloor m/2 \rfloor < i \leq m$, and $j \leq r$ such that $c_i(A_j) \neq c_i(B_j)$, where A_j denotes the j-th matrix in A.

Let $K[A]$ denote the subalgebra of $M_m(K)$ generated by A_i's, the subalgebra $K[B]$ being similar. Clearly $\rho_A(\hat{R}) = K[A]$, and $\rho_B(\hat{R}) = K[B]$. Since $\dim(K[A]) \leq \dim(M_m(K)) = m^2$, it follows (cf. Pappacena [77]) that the words in A_i's of length $\leq m^2$ span $K[A]$. Similarly, the words in B_i's of length $\leq m^2$ span $K[B]$. Let Q be the set of words in U_i's of length $\leq m^2$. It follows that $\rho_A(Q)$ and $\rho_B(Q)$ span $\rho_A(\hat{R}) = K[A]$ and $\rho_B(\hat{R}) = K[B]$, respectively.

Suppose to the contrary that $T_\alpha(A) = T_\alpha(B)$, for all $\alpha \in S$, and $c_i(A_j) = c_i(B_j)$, for all $\lfloor m/2 \rfloor < i \leq m$ and $j \leq r$. Then we will show that $\chi_{\rho_A}(q) = \chi_{\rho_B}(q)$, for all $q \in Q$. For this, we have to show that $c_{k,\alpha}(A) = c_{k,\alpha}(B)$, for every $\alpha \in [r]^*$ of length $\leq m^2$ and $1 \leq k \leq m$, where $c_{k,\alpha}(A) = c_k(A_{i_1} A_{i_2} \cdots)$ if $\alpha = i_1 i_2 \cdots$.

Fix any word $\alpha = i_1 \cdots i_l$ of length $l \leq m^2$. It follows that $\alpha^j = \alpha \cdots \alpha$ (j times), for any $j \leq m$, belongs to S. Hence, by our assumption, it follows that $T_{\alpha^j}(A) = T_{\alpha^j}(B)$ for all $j \leq m$, and $c_i(A_j) = c_i(B_j)$ for all $\lfloor m/2 \rfloor < i \leq m$ and $j \leq r$. By Lemma 2 in Domokos [21], $c_{k,\alpha}(A)$, $1 \leq k \leq m$, is a polynomial in $c_{t,\alpha}(A)$'s, $t \leq \lfloor m/2 \rfloor$, and $c_i(A_j)$'s, $\lfloor m/2 \rfloor < i \leq m$ and $j \leq r$. Since $p \notin [2, \lfloor m/2 \rfloor]$, by Newton's identities, $c_{t,\alpha}(A)$, for $t \leq \lfloor m/2 \rfloor$, can be expressed as:

$$c_{t,\alpha}(A) = \frac{1}{t} \sum_{t'=1}^{t} (-1)^{t'-1} c_{t-t',\alpha}(A) T_{\alpha^{t'}}(A).$$

This shows that $c_{t,\alpha}(A)$, for $t \leq \lfloor m/2 \rfloor$, is a polynomial in $T_{\alpha^j}(A)$'s, $j \leq \lfloor m/2 \rfloor$. The story for B is similar.

It follows that $c_{k,\alpha}(A) = c_{k,\alpha}(B)$, for every $\alpha \in [r]^*$ of length $\leq m^2$ and $1 \leq k \leq m$. That is, $\chi_{\rho_A}(q) = \chi_{\rho_B}(q)$ for all $q \in Q$.

The representations ρ_A and ρ_B are semi-simple, since A and B are semi-simple. Furthermore,

$\rho_A(Q)$ and $\rho_B(Q)$ span $\rho_A(\hat{R}) = K[A]$ and $\rho_B(\hat{R}) = K[B]$, respectively. Hence, it follows from Theorem 10.4, applied to $\hat{R}$, that $\rho_A \cong \rho_B$; a contradiction.

(b) The proof is similar to that of (a). It holds in arbitrary characteristic, since we do not need to use Newton's identities now. Q.E.D.

10.1.1 Proof of Theorem 10.1

(a): Fix any $p \notin [2, \lfloor m/2 \rfloor]$. Let $Y = (y_1, \ldots, y_{m^2})$ be a tuple of auxiliary variables. For any $l \leq m^2$, let $P_l(Y, U) := \sum_\alpha Y_\alpha T_\alpha(U)$, where $\alpha = \alpha_1\alpha_2\cdots$ ranges over all words of length l with each $\alpha_j \in [r]$, and $Y_\alpha = \prod_{j=1}^{l} y_j^{\alpha_j}$.

By Theorem 10.2 (a), the coefficients $c_i(U_j)$'s of $\det(zI - U_j)$'s (considered as polynomials in z with coefficients in $K[U]$), $1 \leq j \leq r$, and the coefficients of $P_l(Y, U)$'s (considered as polynomials in Y with coefficients in $K[U]$), $l \leq m^2$, form a separating set of invariants in $K[V]^G$.

Furthermore (cf. the proof of Lemma 7.9), each $P_l(Y, U)$ can be computed by a read-once oblivious algebraic branching program over Y and U of poly(l, m, r) size, thinking of the entries of U_i's as indeterminate constants.

Let $U' = (U'_1, \ldots, U'_r)$ be another tuple of $m \times m$ variable matrices. Let $\tilde{P}_l(Y, U, U') := P_l(Y, U) - P_l(Y, U')$. It follows that $\tilde{P}_l(Y, U, U')$ can also be computed by a read-once oblivious algebraic branching program over Y, U, and U' of size $q = O(\text{poly}(l, m, r))$, thinking of the entries of U_i's and U'_i's as indeterminate constants. By our assumption, there exists an explicit poly(l, m, r)-time-computable hitting set B for polynomial identity testing for read-once oblivious algebraic branching programs of size q over Y. Fix such a B. In the definition of B, we are considering programs over Y, and not over Y, U, and U', for the reasons that will become clear soon. Fix also $m + 1$ distinct elements $a_0, \ldots, a_m \in F_{p^k}$, $k = O(\log m)$.

Claim 10.5 *The set*

$$S := \{P_l(b, U) \mid b \in B, 1 \leq l \leq m^2\} \cup \{\det(a_i I - U_j) \mid 0 \leq i \leq m, 1 \leq j \leq r\} \tag{42}$$

is a separating set of invariants in $K[V]^G$, if $p \notin [2, \lfloor m/2 \rfloor]$.

Proof of the claim: Let $A = (A_1, \ldots, A_r)$ and $A' = (A'_1, \ldots, A'_r)$ be any two r-tuples in $V = M_m(K)^r$ such that, for some invariant $h \in K[V]^G$, $h(A) \neq h(A')$. We have to show that some element in S assumes distinct values at A and A'.

By Theorem 10.2 (a), either (1) some coefficient $c_i(U_j)$ of $\det(zI - U_j)$ (considered as a polynomial in z), for some $j \leq r$, or (2) some coefficient of $P_l(Y, U)$ (considered as a polynomial in Y), for some $l \leq m^2$, assumes different values at A and A'.

In the first case, since $\det(zI - U_j)$, as a polynomial in z, has degree m, it follows that $\det(a_i I - A_j) \neq \det(a_i I - A_j)$ for some $0 \leq i \leq m$. Hence $\det(a_i I - U_j) \in S$ assumes distinct values at A and A' in this case.

In the second case, $\tilde{P}_l(Y, A, A') = P_l(Y, A) - P_l(Y, A')$ is not identically zero as a polynomial in Y. Since $\tilde{P}_l(Y, U, U')$ has a read-once oblivious algebraic branching program of size $q =$

$O(\text{poly}(m,r))$ over Y, U, and U', thinking of the entries of U_i's and U_i''s as indeterminate constants, $\tilde{P}_l(Y,A,A')$ has a read-once oblivious algebraic branching program of size q over Y. Since B is a hitting set against such programs over Y, and $\tilde{P}_l(Y,A,A')$ is not identically zero as a polynomial in Y, there exists $b \in B$ such that $\tilde{P}_l(b,A,A') \neq 0$, i.e., $P_l(b,A) \neq P_l(b,A')$. Hence, $P_l(b,U) \in S$ assumes distinct values at A and A' in this case.

It follows that S is a separating set of invariants in $K[V]^G$. This proves the claim.

Every element of S is clearly homogeneous of $\text{poly}(m,r)$ degree. By the generalization of Theorem 7.7 to arbitrary characteristic (cf. Theorem 2.3.12 in [17]), it follows that $K[V]^G$ is integral over the subring generated by S.

The size of S is $\text{poly}(m,r)$. Since the hitting set B is explicit, and $P_l(Y,U_j)$'s and $\det(a_iI - U_j)$'s have explicit weakly skew circuits, it follows that the specification of S, consisting of a weakly skew circuit for its every element, can be computed in $\text{poly}(m,r)$ time. Hence S is a separating e.s.o.p. of $K[V]^G$.

This proves the first statement in Theorem 10.1 (a).

The second statement follows from this proof of the first statement, inserting quasi-prefixes in appropriate places, in conjunction with the black-box quasi-derandomization of polynomial identity testing for read-once oblivious algebraic branching programs in Forbes and Shpilka [30], which holds in arbitrary characteristic.

(b) By Theorem 10.2 (b), the set $\{c_{i,\alpha}(U) \mid 0 \leq i \leq m\} \subseteq K[V]^G$, where $\alpha = i_1 i_2 \cdots \in [r]^*$ ranges over all words of length $\leq m^2$ and $c_{i,\alpha}(U) = c_i(U_{i_1}U_{i_2}\cdots)$, is a separating set of invariants in $K[V]^G$, for any p.

Introduce new variables y and $z_{j,s}$, $1 \leq j \leq m^2$, $0 \leq s \leq r$. Let $z = (..,z_{j,s},..)$ denote the tuple of $z_{j,s}$'s. Let

$$p(U,y,z) := \det(yI - \prod_{j=1}^{m^2}(z_{j,0}I + \sum_{s=1}^{r} z_{j,s}U_s)), \tag{43}$$

where I denotes the $m \times m$ identity matrix. This polynomial remains invariant under the adjoint action of G on the tuple $U = (U_1, \ldots, U_r)$. Hence the coefficients of $p(U,y,z)$, considered as a polynomial in y and z with coefficients in $K[U]$, belong to $K[U]^G = K[V]^G$.

For any $\alpha = i_1 i_2 \cdots \in [r]^*$ of length $\leq m^2$, we can set each $z_{j,s}$ to either zero or one so that $p(U,y,z)$ specializes to the characteristic polynomial of $U_{i_1}U_{i_2}\cdots$. It follows that the coefficients of $p(U,y,z)$, considered as a polynomial in y and z with coefficients in $K[U]^G$, form a separating set of invariants in $K[V]^G$.

The polynomial $p(U,y,z)$ in (43) has a weakly skew circuit (cf. Section 2.1) of $O(\text{poly}(m,r))$ size over y,z, and U. By the polynomial equivalence between weakly skew circuits and symbolic determinants [60], $p(U,y,z)$ can also be expressed as a symbolic determinant of size $q = O(\text{poly}(m,r))$ over y,z, and U. By our assumption, there exists an explicit $\text{poly}(m,r)$-time-computable hitting set B against all symbolic determinants over y and z of size q. Note that B is defined by considering symbolic determinants over y and z, not over y,z, and U. Fix such as a B.

Claim 10.6 *The set $S = \{p(U,b_1,b_2) \mid (b_1,b_2) \in B\}$ is a set of separating invariants in $K[V]^G$, for any p.*

The proof is similar to that of Claim 10.5, with Theorem 10.2 (b) in place of Theorem 10.2 (a). The rest of the proof of Theorem 10.1 (b) is similar to that of the first statement in Theorem 10.1 (a).

(c): Given two rational points $A, A' \in V = M_m(K)^r$, we want to decide if the closures of the G-orbits of A and A' intersect. By the generalization of Theorem 5.4 (d) to arbitrary characteristic (cf. Theorem 1 in [75]), this is so iff every invariant in $K[V]^G$ assumes the same value at A and A', or equivalently, if every invariant in any separating set of invariants in $K[V]^G$ assumes the same value at A and A'.

As noted in the proof of (b), the coefficients of the symbolic determinant $p(U, y, z)$ in (43), considered as a polynomial in y and z with coefficients in $K[U]^G$, form a separating set of invariants in $K[V]^G$. It follows that the closures of the G-orbits of A and A' intersect iff the polynomial $\tilde{p}(A, A', y, z) := p(A, y, z) - p(A', y, z)$ is identically zero as a polynomial in y and z. This can be tested by a co-RDET algorithm [45]: Just substitute random values for y and z, using a large enough extension of F_p, if $p > 0$, and test if the resulting specialization of $\tilde{p}(A, A', y, z)$ is zero.

If $p \notin [2, \lfloor m/2 \rfloor]$, then we can give a deterministic NC-algorithm for the problem as follows.

By Theorem 10.2 (a), the coefficients of $\det(zI - U_j)$'s, $1 \leq j \leq r$, considered as polynomials in z with coefficients in $K[U]$, and the coefficients of $P_l(Y, U)$'s, considered as polynomials in Y with coefficients in $K[U]$, form a separating set of invariants in $K[V]^G$ in this case. Hence it follows from the generalization of Theorem 5.4 (d) to arbitrary characteristic [75] that the closures of the G-orbits of A and A' intersect iff $f_j(z, A, A') := \det(zI - A_j) - \det(zI - A'_j)$, for every $j \leq r$, is identically zero, and $\tilde{P}_l(Y, A, A') := P_l(Y, A) - P_l(Y, A')$, for every $l \leq m^2$, is identically zero. The problem of testing whether $f_j(z, A, A')$ is identically zero belongs to DET [15], since $f_j(z, A, A')$ is the difference of two symbolic determinants in just one variable. The polynomial $\tilde{P}_l(Y, A, A')$ has a read-once oblivious algebraic branching program over Y of poly(m, r) size. Hence, whether it is identically zero can be tested by an NC-algorithm, using a straightforward parallelization of the white-box algorithm for polynomial identity testing for read-once oblivious algebraic branching programs in Raz and Shpilka [81]. (The DET-algorithm for white box polynomial identity testing for read-once oblivious algebraic branching programs in [4] works only in characteristic zero.) This proves Theorem 10.1 (c). Q.E.D.

We also note down the following consequence of the proof of Theorem 10.1 (b). Define *strong explicitness of V/G in the relaxed sense* by extending Definition 5.2 (e) to positive characteristic in the obvious way.

Theorem 10.7 *The categorical quotient V/G is strongly explicit in the relaxed sense in any characteristic, with $p(U, y, z)$ in (43) as the defining polynomial.*

Remark (on matrix semi-invariants): We can also let $G = SL_m(K) \times SL_m(K)$, and $V = M_m(K)^r$, with the left-right action of G, which maps $(C_1, \ldots, C_r) \in V$, given $(A, B) \in G$, to $(AC_1B^{-1}, \ldots, AC_rB^{-1})$. In characteristic zero, the recent polynomial degree bound in [18, 48], in conjunction with [19, 22], implies that the categorical quotient V/G is then strongly explicit. Hence, by Theorem 9.9 (b) and Remark 1 after Theorem 9.7, $K[V]^G$ has a separating e.s.o.p. in this case, assuming the black-box derandomization hypothesis for symbolic determinant identity testing. It would be interesting to make this result unconditional.

10.2 Generalization to quivers

Theorem 10.1 can be generalized to arbitrary quivers as follows.

Let Q be a quiver (a directed graph allowing loops and multiple arrows) [11, 19], i.e., a four-tuple (Q_0, Q_1, t, h), where $Q_0 = \{1, \ldots, l\}$ is a set of vertices, Q_1 is a finite set of arrows among these vertices, and the two maps $t, h : Q_1 \to Q_0$ assign to each arrow $\phi \in Q_1$ its tail $t(\phi)$ and head $h(\phi)$. A *representation* W of the quiver Q over the field K is a family $\{W(i) : i \in Q_0\}$ of finite dimensional vector spaces over K together with a family of linear maps $W(\phi) : W(t(\phi)) \to W(h(\phi))$, $\phi \in Q_1$. The l-tuple $(\dim(W(1)), \ldots, \dim(W(l)))$ of integers is called *the dimension vector* of W. For a fixed dimension vector $m = (m(1), \ldots, m(l)) \in \mathbb{N}^l$, the *representation space* $V = V(Q, m)$ of the quiver Q is the set of all representations of Q with the dimension vector m. Clearly,

$$V = V(Q, m) = \oplus_{\phi \in Q_1} \mathrm{Hom}_K(K^{m(t(\phi))}, K^{m(h(\phi))}) = \oplus_{\phi \in Q_1} M_\phi(K), \tag{44}$$

where $M_\phi(K)$ denotes the space of $m(h(\phi)) \times m(t(\phi))$ matrices with entries in K. There is a canonical action of

$$G = \prod_{i=1}^{l} GL_{m(i)}(K) \tag{45}$$

on V, defined by

$$(g \cdot W)(\phi) = g(h(\phi)) W(\phi) g(t(\phi))^{-1},$$

for any $g = (g(1), \ldots, g(l)) \in G$ and $W \in V(Q, m)$.

Let $U = (\ldots, U_\phi, \ldots)$, $\phi \in Q_1$, be a tuple of variable matrices, where U_ϕ is an $m(h(\phi)) \times m(t(\phi))$ variable matrix. Then the coordinate ring $K[V]$ of V can be identified with the ring $K[U]$ over the variable entries of U_ϕ's. Let $K[V]^G \subseteq K[V]$ be the subring of G-invariants. When Q consists of a single vertex with r self-loops, $K[V]^G$ for the dimension vector (m), $m \in \mathbb{N}$, coincides with the invariant ring in Theorem 10.1. If Q has no directed cycles then $K[V]^G = K$. So we are mainly interested in the case when Q has directed cycles.

We specify V and G by giving the graph of the quiver Q, and the dimension vector $m = (m(1), \ldots, m(l)) \in \mathbb{N}^l$ (in unary). Let $|m| := \sum_i m(i)$. The definitions of s.s.o.p. and e.s.o.p. for the ring of matrix invariants extend to this setting in a natural way.

Theorem 10.8 *The analogue of Theorem 10.1, after replacing m there with $|m|$ here, holds for V and G as above.*

For the proof, we recall some results concerning the path algebra of a quiver.

Let R_Q be the path algebra (cf. Section 1.2 in [9]) of Q. This is the associative algebra generated by the variables e_i, $i \in Q_0$, and e_ϕ, $\phi \in Q_1$, subject to the relations:

$$e_i^2 = e_i, \quad e_i e_j = 0 \ (i \neq j), \quad e_{h(\phi)} e_\phi = e_\phi e_{t(\phi)} = e_\phi. \tag{46}$$

Given two representations W_1 and W_2 of Q, a morphism $f : W_1 \to W_2$ between these two representations is a family of linear morphisms $\{f(i) : W_1(i) \to W_2(i) \mid i \in Q_0\}$ such that, for all

$\phi \in Q_1$, $W_2(\phi) \circ f(t(\phi)) = f(h(\phi)) \circ W_1(\phi)$. Thus the set of representations of Q is a category, and two representations of Q are isomorphic iff they are in the same G-orbit.

Proposition 10.9 *(cf. Proposition 1.2.2 in [9]) The category of representations of Q is equivalent to the category of left R_Q-modules.*

Let $W = (\{W(i)\}, \{W_\phi\})$ be any representation of Q with the dimension vector $m = (m(1), m(2), \ldots)$. For any $i \in Q_0$, let M_i^W denote the $|Q_0| \times |Q_0|$-block matrix whose (1) (i', j')-th block, for $1 \leq i', j' \leq |Q_0|$ with $i' \neq j'$ or $i' = j' \neq i$, is the $m(i') \times m(j')$ zero-matrix, and (2) the (i,i)-th block is the $m(i) \times m(i)$ identity matrix. For any $\phi \in Q_1$, let M_ϕ^W denote the $|Q_0| \times |Q_0|$-block matrix defined similarly, whose $(h(\phi), t(\phi))$-th block is W_ϕ, and all other blocks are zero.

It can be checked that the representation W of Q defines the left R_Q-module $\hat{W} := \bigoplus_i W(i)$, on which the action of e_i, $i \in Q_i$, is given by the matrix M_i^W, and the action of e_ϕ is given by M_ϕ^W. The representation $\hat{W}$ is completely specified by the matrix tuple $M^W := (\cdots, M_i^W, \cdots, M_\phi^W, \cdots)$, $i \in Q_0$, $\phi \in Q_1$, of $|m| \times |m|$ matrices. We think of this tuple as an element of $\hat{V} := M_{|m|}(K)^{|Q_0|+|Q_1|}$.

Theorem 10.10 *(cf. Theorem 4.1 in King [54]) The G-orbit of a representation W of Q is closed iff the R_Q-module $\hat{W}$ is semi-simple.*

Proof of Theorem 10.8: We only show how the analogue of Theorem 10.1 (a) for quivers can be deduced from Theorem 10.1 (a) for matrix invariants. The story for the analogues of Theorem 10.1 (b) and (c) is similar.

So assume that the characteristic $p \notin [2, \lfloor |m|/2 \rfloor]$, and that the black-box derandomization hypothesis for polynomial identity testing for read-once oblivious algebraic branching programs holds.

Let V be the representation space of Q associated with the dimension vector m. Consider the adjoint action of $\hat{G} = SL_{|m|}(K)$ on $\hat{V} = M_{|m|}(K)^{|Q_0|+|Q_1|}$. By Theorem 10.1 (a) and our black-box derandomization hypothesis, the invariant ring $K[\hat{V}]^{\hat{G}}$ has a separating e.s.o.p. $\hat{S}$ that can be computed in poly$(|m|, |Q_0|, |Q_1|)$ time. Fix such an $\hat{S}$.

For the tuple $U = (\ldots, U_\phi, \ldots)$, $\phi \in Q_1$, of variable matrices as before, define the matrices M_i^U, $i \in Q_0$, and M_ϕ^U, $\phi \in Q_1$, just as we defined M_i^W and M_ϕ^W, replacing W_ϕ's by U_ϕ's in the definition.

This defines a generic representation of R_Q on $\oplus_i K^{m(i)}$ specified by the matrix tuple $M^U = (\cdots, M_i^U, \cdots, M_\phi^U, \cdots)$, $i \in Q_0$, $\phi \in Q_1$, of $|m| \times |m|$ matrices. This tuple can be thought of as a generic point in $\hat{V}$, and we can evaluate each invariant in $\hat{S}$ at M^U. It is easy to see that, for each $\hat{s} \in \hat{S}$, $\hat{s}(M^U) \in K[V]^G$.

Let $S = \{\hat{s}(M^U) \mid \hat{s} \in \hat{S}\}$.

Claim 10.11 *The set S is a separating set of invariants in $K[V]^G$.*

Proof of the claim: By the generalization of Theorem 5.4 (d) to arbitrary characteristic (cf. Theorem 1.1 in [75]), it suffices to show that S separates the closed G-orbits in V.

So suppose $A, A' \in V$ are two representations of Q whose G-orbits are closed and distinct. We want to show that some invariant in S assumes distinct values on these orbits.

Since the G-orbits of A and A' are distinct, it follows that A and A' are not isomorphic representations of Q. Hence, by Proposition 10.9, the R_Q-modules $\hat{A}$ and $\hat{A}'$ are not isomorphic.

Since the G-orbits of A and A' are closed, it follows from Theorem 10.10 that the R_Q-modules $\hat{A}$ and $\hat{A}'$ are semi-simple. Hence, by Theorem 10.3, the $\hat{G}$-orbits of the matrix tuples $M^{\hat{A}}$ and $M^{\hat{A}'}$ are closed. Since $\hat{A}$ and $\hat{A}'$ are not isomorphic, it follows that the $\hat{G}$-orbits of $M^{\hat{A}}$ and $M^{\hat{A}'}$ are distinct and closed. Since $\hat{S}$ is a separating set of invariants in $K[\hat{V}]^{\hat{G}}$, it follows, by the generalization of Theorem 5.4 (d) to arbitrary characteristic [75], that there exists an invariant $\hat{s} \in \hat{S}$ that assumes distinct values at $M^{\hat{A}}$ and $M^{\hat{A}'}$.

But, $\hat{s}(M^{\hat{A}}) = \hat{s}(M^U)(A)$, and similarly, $\hat{s}(M^{\hat{A}'}) = \hat{s}(M^U)(A')$. It follows that the element $\hat{s}(M^U) \in S$ assumes distinct values at A and A'. This proves the claim.

Since $\hat{S}$ is a separating e.s.o.p., a specification of $\hat{S}$, in the form of a weakly skew circuit for its every element, can computed in poly($|m|, |Q_0|, |Q_1|$) time. It follows that the specification of S, in the form of a weakly skew circuit for its every element, can also computed in poly($|m|, |Q_0|, |Q_1|$) time. The size of S is the same as the size of $\hat{S}$, which is poly($|m|, |Q_0|, |Q_1|$). Furthermore, each element of S is homogeneous, since each element of $\hat{S}$ is homogeneous. By Claim 10.11, S is separating. Hence, by the generalization of Theorem 7.7 to arbitrary characteristic (cf. Theorem 2.3.12 in [17]), $K[V]^G$ is integral over the subring generated by S. It follows that S is a separating e.s.o.p. of $K[V]^G$. Q.E.D.

Theorem 10.8 has a simpler proof in characteristic zero. This can be obtained by extending the proof of Theorem 7.6, using Proposition 10.9, and replacing Theorem 6.2 by its generalization for quivers (cf. Theorem 1 in [11]). This generalization states that $K[V]^G$ is generated by the trace-monomials associated with the oriented cycles in Q of length $\leq |m|^2$.

10.3 Explicit parametrization of closed orbits

Now, let V be a finite dimensional representation of any reductive [62] algebraic group G over K. The set of G-orbits in V, in general, cannot be given the structure of an algebraic variety. The fundamental insight in [75] is that the set of closed G-orbits in V can be given the structure of an algebraic variety. Indeed, by the generalization of Theorem 5.4 to arbitrary characteristic [75], the points of V/G are in one-to-one correspondence with the closed G-orbits in V. But this algebraic structure is not efficient from the complexity-theoretic perspective, since typically a set of generators for $K[V]^G$, such as the one in Theorem 6.2, has exponential cardinality. So we ask if the set of closed G-orbits in V can be given the structure of a variety that is efficient from the complexity-theoretic perspective.

For simplicity, we confine ourselves to the case when $V = M_m(K)^r$, with the adjoint action of $G = SL_m(K)$, as in Section 10.1, and we assume that V and G are specified by giving m and r in unary. But the analogue of Theorem 10.13 below holds for any finite dimensional representation of any reductive algebraic group.

Given a set $S = \{s_1, \ldots, s_k\} \subseteq K[V]^G$, let $\psi_S : V \to K^k$ denote the map $v \to (s_1(v), \ldots, s_k(v))$. Let $n = \dim(V)$.

Definition 10.12 *We say that the closed G-orbits in V have an* explicit parametrization *if there is exists a subset $S = \{s_1, \ldots, s_k\} \subseteq K[V]^G$, $k = O(poly(n))$, of homogeneous invariants of poly(n) degree such that (1) the image $\psi_S(V)$ of ψ_S is closed, (2) for any $x \in \psi_S(V)$, $\psi_S^{-1}(x)$ contains a unique closed G-orbit in V, and (3) given the specification of V and G as above, the specification of S, consisting of a circuit of poly(n) bit-size for every element in it, can be computed in poly(n) time. The constants in these circuits are rational, if the characteristic p of K is zero. Otherwise, they are in a finite extension field F_{p^l}, with $l = O(\log(m))$.*

In this case, the points of the variety $\psi_S(V)$ are in one-to-one correspondence with the closed G-orbits in V, and given any $v \in V$, $\psi_S(V)$ can be computed in poly(n) arithmetic operations over K. If the circuits specifying S are weakly skew, $\psi_S(V)$, for a rational v (cf. Theorem 10.1), can be computed in time polynomial in n and the bit-length of v.

By the generalization of Theorem 5.4 to arbitrary characteristic [75], explicit parametrization of closed G-orbits in V also yields explicit parametrization of the equivalence classes of G-orbits in V, where two G-orbits are considered equivalent iff their closures intersect.

Theorem 10.13 *The closed G-orbits in V have an explicit parametrization if $K[V]^G$ has a separating e.s.o.p.*

For the proof, we need the following lemma.

Lemma 10.14 *Let $S = \{s_1, \ldots, s_k\} \subseteq K[V]^G$ be a separating set of homogeneous invariants. Then the image $\psi_S(V)$ of ψ_S is a closed subvariety of K^k. Furthermore, for any $x \in \psi_S(V)$, $\psi_S^{-1}(x)$ contains a unique closed G-orbit in V.*

Proof: The map ψ_S can be factored as:

$$V \overset{\pi_{V/G}}{\longrightarrow} V/G \overset{\psi'_S}{\longrightarrow} K^k, \tag{47}$$

where $\pi_{V/G}$ is defined as in (7). By the generalization of Theorem 5.4 (a) to arbitrary characteristic [75], the first map is surjective. Hence the image of ψ_S coincides with the image of ψ'_S. Since S is separating, by the generalization of Theorem 7.7 to arbitrary characteristic [17], the coordinate ring $K[V]^G$ of V/G is integral over the subring generated by S. This means the map ψ'_S is finite (cf. Section 5.3 in [85]), and hence its image is a closed subvariety of K^k. Thus the image of ψ_S is closed.

The map ψ'_S is also one-to-one, since S is separating. Hence, by the generalization of Theorem 5.4 (b) to arbitrary characteristic [75], for any $x \in \psi_S(V)$, $\psi_S^{-1}(x)$ contains a unique closed G-orbit in V, Q.E.D.

Proof of Theorem 10.13: Suppose $K[V]^G$ has a separating e.s.o.p. S. The properties (1) and (2) in Definition 10.12 follow from Lemma 10.14. The property (3) follows because S is an e.s.o.p. Q.E.D.

Theorem 10.1 (a), in conjunction with the proof of Theorem 10.13, implies:

Theorem 10.15 *The closed G-orbits in V have a quasi-explicit parametrization if $p \notin [2, \lfloor m/2 \rfloor]$.*

10.4 Explicit parametrization of semi-simple representations of algebras

Next, we show (cf. Theorem 10.16) that the existence of a separating e.s.o.p. for $K[V]^G$, with $V = M_m(K)^r$ and $G = SL_m(K)$ as above, implies explicit parametrization of semi-simple representations of any finitely generated algebra.

Let R be a finitely generated associative algebra over K, specified by its generators $f_1, \ldots, f_r$, and relations among them. We assume that the coefficients in the relations are in a finite extension of $\mathbb{Q}$, if the characteristic is zero, or F_p, if the characteristic is p. Let $\rho : R \to M_m(K)$ be an m-dimensional representation of R. It can be identified with the r-tuple $A = (A_1, \ldots, A_r) \in V = M_m(K)^r$ of $m \times m$ matrices, where $A_i = \rho(f_i)$. The set $W_m = W_m(R)$ of the r-tuples corresponding to m-dimensional representations of R is a closed G-subvariety of V. Two representations of R are isomorphic iff they lie in the same G-orbit, where $G = SL_m(K)$ acts on V by the adjoint action as before. By Theorem 10.3, a representation is semi-simple iff its G-orbit is closed. Thus the isomorphism classes of m-dimensional semi-simple representations of R can be identified with the closed G-orbits in W_m. Let $n = rm^2$.

We say that semi-simple representations of R of dimension m have an *explicit parametrization* if there exists a set S of poly(n) homogeneous invariants of poly(n) degree in $K[V]^G$ such that (1) the image $\psi_S(W_m)$ of W_m under the map ψ_S, defined in Section 10.3, is closed, (2) for any $x \in \psi_S(W_m)$, $\psi_S^{-1}(x)$ contains a unique closed G-orbit in W_m, and (3) given m and the specification of R, the specification of S, consisting of a circuit of poly(n) bit-size for every element in it, can be computed in time polynomial in n and the bit-length of the specification of R. The constants in these circuits are rational if the characteristic p of K is zero. Otherwise, they are in a finite extension field F_{p^l} with $l = O(\log(m))$.

In this case, the points of the variety $\psi_S(W_m)$ are in one-to-one correspondence with the isomorphism classes of m-dimensional semi-simple representations of R, and given any r-tuple $A \in V$ of matrices specifying an m-dimensional representation of R, $\psi_S(A)$ can be computed in poly(n) arithmetic operations over K.

Theorem 10.16 *For any m, the m-dimensional semi-simple representations of R over K have an explicit parametrization if $K[V]^G$ has a separating e.s.o.p.*

This result follows from Theorem 10.13 and the following result.

Proposition 10.17 *Let S be as in Lemma 10.14, with V and G as above. Then $\psi_S(W_m)$ is a closed subvariety of K^k.*

Proof: By the generalization of Theorem 5.4 (c) to arbitrary characteristic [75], $Y = \pi_{V/G}(W_m)$ is a closed subvariety of V/G. As shown in the proof of Lemma 10.14, ψ'_S in (47) is a finite morphism. Since the image of a closed variety under a finite morphism is closed (cf. Section 5.3. in [85]), the image $\psi'_S(Y) = \psi_S(W_m)$ is closed. Q.E.D.

Remark: The set S giving the explicit parametrization in Theorem 10.16 depends only on m and r, the number of generators of R, but not on the relations among the generators of R.

Theorem 10.16, in conjunction with Theorem 10.1 (a), implies:

Theorem 10.18 *For any m, the m-dimensional semi-simple representations of R over K have a quasi-explicit parametrization, if $p \not\in [2, \lfloor m/2 \rfloor]$.*

10.5 Other extensions in positive characteristic

Next, we briefly explain how the results in Sections 4 and 5 can be extended to positive characteristics.

The strengthened black-box derandomization problem for low-degree polynomial identity testing over an algebraically closed field K of positive characteristic p is defined just as in characteristic zero (cf. Section 2.5). The hitting set against low-degree circuits over K of size $\leq s$ is assumed to be a subset of $F_{p^l}^n$, n the number of variables, for a large enough $l = O(\log s)$. The phrase "infinitesimally close" is interpreted in the Zariski topology. The definition of NNL is extended from characteristic zero to positive characteristics similarly in a straightforward way.

The following result extends Theorem 4.9 to positive characteristics.

Theorem 10.19 *(a) The variety $\Delta[\det, m]$ has a strict e.s.o.p. in any characteristic iff the strengthened black-box derandomization hypothesis for symbolic determinant identity testing holds.*

(b) The variety $\Delta[\det, m]$ has a strict e.s.o.p. over an algebraically closed field of $\Omega(2^{(\log m)^a})$ characteristic, for a large enough positive constant a, iff, ignoring a quasi-prefix, there exists a family $\{f_n(x_1, \ldots, x_n)\}$ of exponential-time-computable (cf. the remark after Theorem 2.1), multi-linear, integral polynomials such that f_n cannot be approximated infinitesimally closely over an algebraically closed field of $\Omega(2^{n^\delta})$ characteristic by circuits of $O(2^{n^\epsilon})$ size, for some constants $\delta, \epsilon > 0$, as $n \to \infty$.

Analogous result holds for the explicit variety $\Delta[H(Y)_m, k, m]$ associated with the low-degree universal circuit in Section 5.1.3. Similar extensions of Theorems 5.11 (a), 5.14 (a), and 5.14 (b) to arbitrary characteristics, and of Theorems 5.11 (b) and 5.14 (c) to large enough characteristics also hold.

For the proof, we need the following results.

Theorem 10.20 (Kaltofen and Lecerf) *(cf. [51]) Suppose K is an algebraically closed field of positive characteristic p. Then:*

(a) Given any polynomial $g \in K[x_1, \ldots, x_n]$ and a polynomial $f \in K[x_1, \ldots, x_n]$ dividing g, there exists a nonuniform circuit over K, with oracle gates for g, of $O((n \deg(g))^a)$ size, for some absolute positive constant a not depending on n or p, that computes the highest power of f of the form f^{p^l}, $l \geq 0$, that divides g.

(b) In particular, given any polynomial $g \in K[x_1, \ldots, x_n]$, with $\deg(g) < p$, and a polynomial $f \in K[x_1, \ldots, x_n]$ dividing g, there exists a nonuniform circuit over K, with oracle gates for g, of $O((n \deg(g))^a)$ size, for some absolute positive constant a not depending on n or p, that computes f.

The following is the analogue of Theorem 2.4 in this setting.

Theorem 10.21 *Suppose there exists a family $\{p_m(x_1, \ldots, x_m)\}$ of exponential-time-computable, multi-linear, integral polynomials such that p_m cannot be approximated infinitesimally closely over an algebraically closed field of $\Omega(2^{m^\delta})$ characteristic by circuits of $O(2^{m^\epsilon})$ size, for some positive constants δ and ϵ, as $m \to \infty$.*

Then polynomial identity testing for low-degree circuits of size $\leq s$ over an algebraically closed field of $\Omega(2^{(\log s)^a})$ characteristic, for a large enough positive constant a, has $O(2^{polylog(s)})$-time-computable strengthened black-box derandomization.

This is proved like Theorem 2.4, using Theorem 10.20 (b) in place of Theorem 2.2. It can be checked that this replacement is possible, by choosing the constant e in the proof of Theorem 2.4 large enough, depending upon ϵ and δ. The analogue of this result also holds for exact computation in place of infinitesimally close approximation, with a similar proof.

Proof of Theorem 10.19: All results, other than Theorem 2.4, used in the proof of Theorem 4.9, namely, Theorem 2.3, Noether's Normalization Lemma (Lemma 3.1), Hilbert's Nullstellensatz, and other standard facts from algebraic geometry hold in arbitrary characteristic.

Hence, the proof of (a) is similar to that of Theorem 4.9 (a). The proof of (b) is similar to that of Theorem 4.9 (b), using Theorem 10.21 in place of Theorem 2.4. Q.E.D.

Remark 1: The restrictions on the characteristics in Theorem 10.19 (b) (and its generalizations to arbitrary explicit varieties; cf. the remark after Theorem 10.19) can be dropped, and we can let the base field be an algebraically closed field K of any fixed characteristic p, if we assume for the f_n therein that $f_n^{p^i}$, for any nonnegative $i = O(\text{poly}(n))$, cannot be approximated infinitesimally closely by circuits over K of $O(2^{n^\epsilon})$ size, for some constant $\epsilon > 0$, as $n \to \infty$.

Remark 2: Theorem 5.8 similarly holds in arbitrary characteristic. Theorem 5.13 also holds in arbitrary characteristic, since Theorem 5.4 holds in arbitrary characteristic [75].

11 Discussion

Finally, we discuss the difficulties that need to be overcome to improve the current best bound for NNL for $\Delta[\det, m]$ in Theorem 4.10.

Let K now be an algebraically closed field of characteristic zero. If every polynomial in $\Delta[\det, m]$ had a small circuit over K of poly(m) size, then the strengthened black-box derandomization problem for symbolic determinant identity testing would be in PSPACE unconditionally, like the standard black-box derandomization problem (cf. Proposition 2.9), with essentially the same proof. By (the proof of) Theorem 4.5, NNL for $\Delta[\det, m]$ would then be in PSPACE unconditionally.

However, it may be conjectured that the boundary of the orbit of the determinant in $\Delta[\det, m]$ contains points which do not have small circuits over K; cf. Section 4.2 in [71] for a preliminary investigation in this direction, and [68, 38] for further investigation. Formally, we conjecture that $\overline{\text{VP}_{ws}} \not\subseteq \text{VP}$. Here VP [91] is the class of families of polynomials of small degree having circuits of polynomial size, VP_{ws} is the class of families of polynomials that can be computed by symbolic determinants of polynomial size, and $\overline{\mathrm{VP}_{ws}}$ [13, 14] is the class of families of polynomials that can be approximated infinitesimally closely by symbolic determinants of polynomial size.

This conjecture is counter-intuitive, since one would have expected the complexity of infinitesimally close approximation of multi-linear polynomials by symbolic determinants to be polynomially related to that of exact computation. As pointed out in Bürgisser [13] (cf. Lemma 5.6 (3) and Theorem 5.7 therein), this would be the case if every point in the boundary of the orbit of the determinant could be approached by a one-parameter deformation of the determinant of polynomial order. We conjecture that this is not the case. However, for the VNP-complete polynomials such as the permanent, the complexity of infinitesimally close approximation can be conjectured to be polynomially related to that of exact computation. At present, it is not even known if $\overline{\mathrm{VP}_{ws}} \subseteq \mathrm{VNP}$, where VNP [91] is the class of p-definable families of polynomials.

The conjectural points with large circuit complexity in $\Delta[\det, m]$ constitute the main obstacle to putting NNL for $\Delta[\det, m]$ in PSPACE, or even EXP, unconditionally with the existing techniques. (This obstacle is absent for explicit categorical quotients, as in Theorems 1.3 and 1.5, by Theorem 5.4 (a).) In contrast, the Generalized Riemann Hypothesis assumption in the current EXPH-bound in Theorem 4.10 may be removed in the foreseeable future (though, this by itself is a nontrivial problem).

Thus, bringing NNL for $\Delta[\det, m]$ from EXPH, where it is currently assuming the Generalized Riemann Hypothesis, to even EXP *unconditionally* seems difficult with the existing techniques. Theorem 1.7 says that a sub-exponential algebraic circuit-size lower bound for infinitesimally close approximation of the permanent would put NNL for $\Delta[\det, m]$ in quasi-P. Theorem 1.7 (and Remark 2 after Theorem 5.11) may thus explain why the hardness hypothesis of geometric complexity theory in [71] has turned out to be so difficult. (In the terminology above, this hypothesis is that $\mathrm{VNP} \not\subseteq \overline{\mathrm{VP}_{ws}}$; cf. Proposition 9.3.2 in [14].) It is a reasonable thesis that any realistic approach to the $\mathrm{VNP} \not\subseteq \mathrm{VP}_{ws}$ conjecture in Valiant [91] would also prove this hypothesis. Indeed, all known lower bounds for the exact computation of the permanent also hold for infinitesimally close approximation; eg. see [59, 37]. Hence, Theorem 1.7 may also explain why the $\mathrm{VNP} \not\subseteq \mathrm{VP}_{ws}$ conjecture in [91] has turned out to be so difficult.

Theorem 1.7 and the equivalence results in this article (Theorems 1.9 and 5.14) thus reveal that the fundamental problems of geometry (NNL) and complexity theory (hardness) share a common root difficulty, namely, the problem of overcoming the existing EXPH vs. P gap (assuming the Generalized Riemann Hypothesis) in the complexity of NNL for general explicit varieties, or rather, the EXPH vs. NC gap; cf. Remark 1 after Theorem 5.11. We call this gap the *geometric complexity theory (GCT) chasm*. It may be viewed as the common *cause and measure* of the difficulty of these problems in geometry and complexity theory.

The superpolynomial lower bound in [65] for additive approximation of the maxflow in the PRAM model without bit-operations, which initiated geometric complexity theory (cf. the introduction of [71]), assumes special significance in view of this chasm. First, this lower bound is the main reason why the hardness hypothesis in [71] is expected to hold, despite the conjectural non-containment of $\overline{\mathrm{VP}_{ws}}$ in VP. This is because a lower bound akin to that in [65] for additive approximation of the permanent of integral matrices (instead of the maxflow) implies the hardness hypothesis in [71] for infinitesimally close approximation. Such a lower bound can be expected since, in view $\#P$-completeness [92] of the permanent, the approximation of the permanent is expected to be harder than the approximation of the maxflow. Second, the lower bound in [65] is the only known arithmetic version of a foundational conjecture in complexity theory (in this case, the P $\neq$ NC conjecture) that holds unconditionally in a natural and realistic model of computation. It is now likely to remain the only such lower bound in complexity theory, until the GCT chasm is crossed.

We conjecture that the strong form of NNL for every explicit variety is in P, and hence, the GCT chasm can be crossed, as suggested by Theorem 5.11. By geometric complexity theory, we mean henceforth any approach to cross the GCT chasm using a synthesis of geometry and complexity theory. One such approach will be described in the sequel [64].

Acknowledgement: The author is grateful to Michael Forbes and Amir Shpilka for bringing [2, 86] to his attention and for pointing out [31] an error in the preliminary version [69] of this paper (cf. Section 1.7), and to Jonah Blasiak, Peter Bürgisser, Josh Grochow, Joseph Landsberg, Nitin Saxena, Jimmy Qiao, and the referees for helpful comments.

References

[1] M. Agrawal. Proving lower bounds via pseudo-random generators. In *Proc. FSTTCS*, pages 92–105, 2005.

[2] M. Agrawal, C. Saha, and N. Saxena. Quasi-polynomial hitting-set for set-depth-delta formulas. *ArXiv:1209.2333 [cs.CC]*, 2012.

[3] S. Arora and B. Barak. *Computational complexity: a modern approach.* Cambridge University Press, 2009.

[4] V. Arvind, P. Joglekar, and S. Srinivasan. Arithmetic circuits and Hadamard product of polynomials. In *Proc. FSTTCS*, pages 25–36, 2009.

[5] S. Basu. New results on quantifier elimination in real closed fields and applications to constraint data bases. *JACM*, 46(4):537–555, 1991.

[6] J. Blasiak, K. Mulmuley, and M. Sohoni. Geometric complexity theory IV: nonstandard quantum group for the Kronecker problem. *Memoirs of the American Mathematical Society*, 235(1109, Fourth of 5 numbers), 2015.

[7] J. Boutot. Singularit'es rationelles et quotients par les groupes r'eductifs. *Invent. Math.*, 88:65–88, 1987.

[8] R. Brauer and C. Nesbitt. On the modular representations of groups of finite order I. *In Richard Brauer: Collected papers (Vol. I)*, pages 336–354, 1980.

[9] M. Brion. Representations of quivers. *Preprint, Institute Fourier*, 2012.

[10] W. Bruns and J. Herzog. *Cohen-Macauley rings.* Cambridge University Press, 1993.

[11] L. Bruyn and C. Procesi. Semisimple representations of quivers. *Transactions of the American Mathematical Society*, 317(2):585–598, 1990.

[12] P. Bürgisser. *Completeness and reduction in algebraic complexity theory.* Algorithms and Computation in Mathematics, vol. 7, Springer, 1998.

[13] P. Bürgisser. The complexity of factors of multivariate polynomials. *Found. Comput. Math.*, pages 369–396, 2004. Conference version in the proceedings of FOCS, 2001.

[14] P. Bürgisser, J. Landsberg, L. Manivel, and J. Weyman. An overview of mathematical issues arising in the geometric complexity theory approach to $VP \neq VNP$. *SIAM J. Comput.*, 40(4):1179–1209, 2011.

[15] S. Cook. A taxonomy of problems with fast parallel algorithms. *Journal of Information and Control*, 64(1-3):2–22, 1985.

[16] H. Derksen. Polynomial bounds for rings of invariants. *Proc. Amer. Math. Soc.*, 129:955–963, 2001.

[17] H. Derksen and G. Kemper. *Computational invariant theory.* Encyclopaedia of mathematical sciences, Springer, 2000.

[18] H. Derksen and V. Makam. Polynomial degree bounds for matrix semi-invariants. *ArXiv:1512.03393*, 2015.

[19] H. Derksen and J. Weyman. Semi-invariants of quivers and saturation for Littlewood-Richardson coefficients. *JAMS*, 13(3):467–479, 2000.

[20] J. Dieudonne. The historical development of algebraic geometry. *The American Mathematical Monthly*, 79(8):827–866, 1972.

[21] M. Domokos. Finite generating system of matrix invariants. *Mathematica Pannonica*, 13(2):175–181, 2002.

[22] M. Domokos and A. Zubkov. Semi-invariants of quivers as determinants. *Transformation Groups*, 6(1):9–24, 2001.

[23] S. Donkin. Invariants of several matrices. *Inv. Math.*, 110:389–401, 1992.

[24] P. Doubillet, G. Rota, and J. Stein. On the foundations of combinatorial theory, IX. *Combinatorial methods in invariant theory, Studies in Appl. Math.*, 53:185–216, 1974.

[25] J. Drozd. Tame and wild matrix problems. *Amer. Math. Soc. Transl.*, 128(2):31–55, 1986.

[26] R. Eggermont. Generalizations of a theorem of Brauer and Nesbitt. *Master Thesis, Mathematisch Instituut, Universiteit Leiden*, 2011.

[27] D. Eisenbud. *Commutative algebra with a view toward algebraic geometry.* Springer-Verlag, 1995.

[28] H. Flenner. Rationale quasi-homogene singularitäten. *Arch. Math.*, 36:35–44, 1981.

[29] M. Forbes, R. Saptharishi, and A. Shpilka. Hitting sets for multilinear read-once algebraic branching programs, in any order. In *Proc. STOC*, pages 867–875, 2014.

[30] M. Forbes and A. Shpilka. Quasi-polynomial-time identity testing of non-commutative and read-once oblivious algebraic branching programs. *ECCC, 19:115, Version 1, September*, 2012.

[31] M. Forbes and A. Shpilka. Explicit Noether normalization for simultaneous conjugation via polynomial identity testing. *ArXiv:1303.0084v2 [cs.CC]*, 2013.

[32] E. Formanek. Generating the ring of matrix invariants. *Lecture Notes in Math., Springer-Verlag*, 1195:73–82, 1986.

[33] W. Fulton and J. Harris. *Representation theory, A first course.* Springer, 1991.

[34] M. Goebel. Computing bases for rings of permutation-invariant polynomials. *J. Symb. Comp.*, 19:285–291, 1995.

[35] P. Gordan. Beweis, dass jede Covariante und Invariante einer binären Form eine ganze Funktion mit numerischen Coefficienten einer endlichen Anzahl solcher Formen ist. *J. Reine Angew. Math.*, 69:323–354, 1868.

[36] B. Grenet, P. Koiran, and N. Portier. The multivariate resultant is NP-hard in any characteristic. *ArXiv:0912.2607v3*, 2012.

[37] J. Grochow. Unifying known lower bounds via geometric complexity theory. *Computational Complexity*, 24:393–475, 2015.

[38] J. Grochow, Y. Qiao, and K. Mulmuley. Boundaries of VP and VNP. *To appear in the proceedings of ICALP, 2016.*

[39] A. Gupta, P. Kamath, N. Kayal, and R. Saptharishi. Arithmetic circuits: a chasm at depth three. *ECCC, Report No. 26*, 2013.

[40] R. Hartshorne. *Algebraic geometry.* Springer-Verlag, 1977.

[41] J. Heintz and C. Schnorr. Testing polynomials that are easy to compute. In *Proc. STOC*, pages 262–272, 1980.

[42] D. Hilbert. Über die Theorie der algebraischen Formen. *Math. Ann.*, 36:473–534, 1890.

[43] D. Hilbert. Über die vollen Invariantensysteme. *Math. Ann.*, 42:313–370, 1893.

[44] J. Humphreys. *Linear algebraic groups.* Springer Verlag, 1981.

[45] O. Ibarra and S. Moran. Probabilistic algorithms for deciding equivalence of straight-line programs. *JACM*, 30(1):217–228, 1983.

[46] C. Ikenmeyer and G. Panova. Rectangular Kronecker coefficients and plethysms in geometric complexity theory. *ArXiv:1512.03798*, 2015.

[47] R. Impagliazzo and A. Wigderson. $P = BPP$ unless E has sub-exponential circuits: Derandomizing the XOR lemma. In *Proc. STOC*, pages 220–229, 1997.

[48] G. Ivanyos, Y. Qiao, and K. Subrahmanyam. Constructive noncommutative rank computation in deterministic polynomial time over fields of arbitrary characteristic. *ArXiv:1512.03531*, 2016.

[49] V. Kabanets and R. Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Computational Complexity*, 13(1-2):1–46, 2004.

[50] E. Kaltofen. Factorization of polynomials given by straight-line programs. *Randomness and Computation*, 5:375–412, 1989.

[51] E. Kaltofen and G. Lecerf. Factorization of multivariate polynomials. *Chapter 11.5, in Handbook of Finite Fields, Discrete Mathematics and Its Applications, CRC press*, 2013.

[52] E. Kaltofen and B. Trager. Computing with polynomials given by black boxes for their evaluations: greatest common divisors, factorization, separation of numerators and denominators. *J. Symbolic Comput.*, 9(3):301–320, 1990.

[53] N. Kayal and R. Saptharishi. A selection of lower bounds for arithmetic circuits. *Progress in Computer Science and Applied Logic*, 26:77–115, 2014.

[54] A. King. Moduli of representations of finite-dimensional algebras. *Quart. J. Math. Oxford Ser. (2)*, 45(180):515–530, 1994.

[55] P. Koiran. Hilbert's Nullstellensatz is in the polynomial hierarchy. *Journal of complexity*, 12:273–286, 1996.

[56] J. Kollár. Sharp effective Nullstellensatz. *J. Amer. Math. Soc.*, 1:963–975, 1988.

[57] V. Lakshmibai and K. Raghavan. *Standard monomial theory.* Encyclopaedia of mathematical sciences, Springer, 2008.

[58] J. Landsberg. *Tensors: geometry and applications, graduate studies in mathematics*, volume 128. AMS, 2012.

[59] J. Landsberg and N. Ressayre. Hypersurfaces with degenerate duals and the geometric complexity theory program. *ArXiv:1004.4802*, 2010.

[60] G. Malod and N. Portier. Characterizing Valiant's algebraic complexity classes. *Journal of Complexity*, 24(1):16–38, 2008.

[61] E. Mayr and S. Ritscher. Space efficient Gröbner basis computation without degree bounds. In *Proc. ISAAC*, pages 257–264, 2011.

[62] J. Milne. Reductive groups. *http://www.jmilne.org/math/CourseNotes/RG.pdf*, 2012.

[63] B. Mourrain and V. Pan. Multivariate polynomials, duality, and structured matrices. *Journal of complexity*, 16:110–180, 2000.

[64] K. Mulmuley. Geometric complexity theory VI. *Revised version under preparation.* The current version of this article at ArXiv:0704.0229 is outdated; cf. the third paragraph in Section 1.1.

[65] K. Mulmuley. Lower bounds in a parallel model without bit operations. *SIAM J. Comput.*, 28:1460–1509, 1999.

[66] K. Mulmuley. On P vs. NP and Geometric complexity theory. *JACM*, 58(2, Article No. 5), 2011.

[67] K. Mulmuley. The GCT program toward the P vs. NP problem. *CACM*, 55(6):98–107, 2012.

[68] K. Mulmuley. The GCT chasm I. *Tutorial in the workshop on Geometric Complexity Theory, Simons Institute, Berkeley, http://simons.berkeley.edu/workshops/schedule/430*, 2014.

[69] K. Mulmuley. Geometric complexity theory V: equivalence between black-box derandomization of polynomial identity testing and derandomization of Noether's Normalization Lemma. In *Proc. FOCS*, pages 629–638, October, 2012. See also ArXiv:1209.5993, version 3, 2012.

[70] K. Mulmuley, H. Narayanan, and M. Sohoni. Geometric complexity theory III: on deciding nonvanishing of a Littlewood-Richardson coefficient. *Journal of Algebraic Combinatorics*, 36(1):103–110, 2012.

[71] K. Mulmuley and M. Sohoni. Geometric complexity theory I: an approach to the P vs. NP and related problems. *SIAM J. Comput.*, 31(2):496–526, 2001. This article is the journal version of the work that was presented in the workshop on complexity lower bounds, Fields Institute, Toronto, 1998, http://www.fields.utoronto.ca/programs/scientific/97-98/complexity/bounds/.

[72] K. Mulmuley and M. Sohoni. Geometric complexity theory, P vs. NP, and Explicit Obstructions. In *Advances in Algebra and Geometry, Edited by C. Musili, Proc. International Conference on Algebra and Geometry, Hyderabad*, pages 239–261, 2001.

[73] K. Mulmuley and M. Sohoni. Geometric complexity theory II: towards explicit obstructions for embeddings among class varieties. *SIAM J. Comput.*, 38(3):1175–1206, 2008.

[74] D. Mumford. *Algebraic geometry I: complex projective varieties.* Springer, 1976.

[75] D. Mumford, J. Fogarty, and F. Kirwan. *Geometric invariant theory.* Springer-Verlag, 1994.

[76] N. Nisan and A. Wigderson. Hardness vs. randomness. *JCSS*, 49(2):149–167, 1994.

[77] C. Pappacena. An upper bound on the length of a finite dimensional algebra. *J. algebra*, 197:535–545, 1997.

[78] V. Popov. The constructive theory of invariants. *Math. USSR Izvest.*, 10:359–376, 1982.

[79] V. Popov and E. Vinberg. *Invariant theory, in Encyclopaedia of mathematical sciences*, volume 55. Springer-Verlag, 1989.

[80] C. Procesi. The invariant theory of $n \times n$ matrices. *Adv. in Math.*, 19:306–381, 1976.

[81] R. Raz and A. Shpilka. Deterministic polynomial identity testing in non-commutative models. *Computational Complexity*, 14(1):1–29, 2005.

[82] Y. Razmyslov. Trace identities of full matrix algebras over a field of characteristic zero. *Math. USSR Izv.*, 8:727–760, 1974.

[83] N. Saxena. Diagonal circuit identity testing and lower bounds. *Lecture Notes in Computer Science*, 5125:60–71, 2008.

[84] J. Schwarz. Fast probabilistic algorithms for verification of polynomial identities. *JACM*, 27:701–717, 1980.

[85] I. Shafarevich. *Basic algebraic geometry I.* Sringer Verlag, 1977.

[86] A. Shpilka and I. Volkovich. Read-once polynomial identity testing. In *Proc. APPROX-RANDOM*, pages 700–713, 2009.

[87] A. Shpilka and A. Yehudayoff. Arithmetic circuits: a survey of recent results and open questions. *Foundations and Trends in Theoretical Computer Science*, 5(3-4):207–388, 2010.

[88] V. Strassen. Die Berechnungskomplexiät von elementarsymmetrischen Funktionen und von Interpolationskoeffizienten. *Numerische Mathematik*, 20:238–251, 1973.

[89] B. Sturmfels. *Algorithms in invariant theory.* Springer-Verlag, 1993.

[90] B. Totaro. *private communication.*

[91] L. Valiant. Completeness classes in algebra. In *Proc. STOC*, pages 249–261, 1979.

[92] L. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8:189–201, 1979.

[93] L. Valiant, S. Skyum, S. Berkowitz, and C. Rackoff. Fast parallel computation of polynomials using few processors. *SIAM J. Computing*, 12(4):641–644, 1983.

[94] H. Weyl. *The classical groups. Their invariants and representations.* Princeton University Press, 1939.