

The computational complexity to evaluate representations of general linear groups

Extended Abstract, March 17, 1998

Peter Bürgisser

Institut für Mathematik, Universität Zürich, Winterthurerstr. 190,
CH-8057 Zürich, Switzerland. E-mail: buerg@amath.unizh.ch

Abstract. We describe a fast algorithm to evaluate irreducible matrix representations of general linear groups $GL_m(\mathbb{C})$ with respect to a symmetry adapted basis (Gelfand-Tsetlin basis). This is complemented by a lower bound, which shows that our algorithm is optimal up to a factor m^2 with regard to nonscalar complexity. Our algorithm can be used for the fast evaluation of special functions: for instance, we obtain an $O(\ell \log \ell)$ algorithm to evaluate all associated Legendre functions of degree ℓ . For studying the complexity to evaluate single entries of the representation matrix, we investigate the complexity of immanants. We obtain an algorithm to evaluate immanants which is faster than previous algorithms due to Hartmann and Barvinok. Finally, we show that the evaluation of certain immanants is inherently difficult by proving their completeness in a complexity model due to Valiant.

Résumé. Nous décrivons un algorithme rapide pour évaluer des représentations irréductibles du groupe linéaire général $GL_m(\mathbb{C})$ par rapport à une base adaptée à la symétrie (base de Gelfand et Tsetlin). Ce résultat est complété par l'obtention d'une borne inférieure qui montre que notre algorithme est optimal à un facteur m^2 près, par rapport à la complexité non scalaire. On peut l'employer pour évaluer des fonctions spéciales : par exemple, nous obtenons un algorithme $O(\ell \log \ell)$ pour évaluer toutes les fonctions de Legendre associées de degré ℓ . Afin d'étudier la complexité d'évaluer des coefficients particuliers de la matrice de représentation, nous examinons la complexité des immanants. Nous obtenons un algorithme pour l'évaluation des immanants plus rapide que les algorithmes précédents dus à Hartmann et Barvinok. Finalement, nous montrons que l'évaluation de certains immanants est intrinsèquement difficile en prouvant leur complétude dans un modèle dû à Valiant.

1 Introduction

The theory of representations of Lie groups has countless applications in mathematics and physics. In particular, it is an indispensable tool of quantum mechanics. Let $D: G \rightarrow GL_d$ be an irreducible (finite dimensional) continuous representation of the Lie group G . How fast can we compute the representation matrix $D(g)$ for

a given $g \in G$? This question contains the problem of the efficient evaluation of special functions and orthogonal polynomials, which can be interpreted as matrix entries of some suitable representation D (see Vilenkin and Klimyk [18]).

In this paper, we investigate the above question for the complex general linear group $G = \mathrm{GL}_m := \mathrm{GL}(m, \mathbb{C})$. This includes the case of the unitary group $\mathrm{U}(m)$, which is of particular importance for physics, since all the continuous irreducible representations of $\mathrm{U}(m)$ can be obtained from the rational irreducible ones of GL_m by restriction.

It is well known that the irreducible polynomial representations of the general linear group GL_m can be labelled by decreasing sequences $\lambda \in \mathbb{N}^m$ (cf. [3]). With respect to some chosen basis, such a representation is a group morphism

$$D_\lambda : \mathrm{GL}_m \longrightarrow \mathrm{GL}_{d_\lambda}, \quad A \mapsto D_\lambda(A) \quad ,$$

and the entries of D_λ are homogeneous polynomials of degree $|\lambda| := \sum_i \lambda_i$ in the entries of A . The matrix $D_\lambda(A)$ is usually called an *invariant matrix* when the entries of A are interpreted as indeterminates. Littlewood [12, 13] found an explicit construction of invariant matrices. He obtained formulas for the polynomial entries of $D_\lambda(A)$ in terms of representations of the symmetric group. Grabmeier and Kerber [9] gave a modern derivation of such formulas, and based on this, they designed an algorithm for computing invariant matrices. This algorithm in fact computes a sparse representation of the polynomial entries of the invariant matrix. However, the example of the $m \times m$ determinant already shows that this approach cannot be efficient for evaluating $D_\lambda(A)$ at specific entries $A \in \mathrm{GL}_m$ for larger m : the determinant polynomial has $m!$ terms, but it can be computed with only $O(m^3)$ arithmetic operations using Gaussian elimination.

As unitary representations of $\mathrm{U}(m)$ are important in quantum mechanics, it is not astonishing that various explicit constructions of representations have been developed by physicists. (See for instance the books by Biedenharn and Louck [2], where the entries of invariant matrices are called boson polynomials.)

Gelfand and Tsetlin [8] derived explicit expressions for $D_\lambda(A)$ for generators A of GL_m with respect to bases *adapted* to the chain of subgroups

$$\mathrm{GL}_m > \mathrm{GL}_{m-1} \times \mathbb{C}^\times > \mathrm{GL}_{m-2} \times (\mathbb{C}^\times)^2 > \dots > (\mathbb{C}^\times)^m \quad .$$

Such symmetry adapted bases are also called *Gelfand-Tsetlin bases*. (For a definition see Sect. 2.) A very detailed account of this can be found in the book by Vilenkin and Klimyk [18, Chap. 18].

Our main result is an efficient algorithm to evaluate the map D_λ with respect to a Gelfand-Tsetlin basis. The underlying idea is roughly as follows: Using Gaussian elimination, we factor the given matrix $A \in \mathrm{GL}_m$ into a product $A = A_N A_{N-1} \dots A_1 \Delta$ of a diagonal matrix Δ and of elementary matrices A_i of the form $B_{n-1,n}(t)$ or $B_{n,n-1}(t)$. For instance, the matrix $B_{n-1,n}(t)$ has the following effect ($t \in \mathbb{C}$, e_i canonical basis vector, $n \leq m$)

$$B(t)e_n = te_{n-1} + e_n, \quad B(t)e_i = e_i \quad (i \neq n) \quad .$$

We have $D_\lambda(A) = D_\lambda(A_N)D_\lambda(A_{N-1}) \cdots D_\lambda(A_1)D_\lambda(\Delta)$. For showing that $D_\lambda(A)$ can be multiplied fast with a given vector $v \in \mathbb{C}^{d_\lambda}$, it suffices to show this for $D_\lambda(\Delta)$ and for each $D_\lambda(A_i)$. Since Gelfand-Tsetlin bases consist of weight vectors, $D_\lambda(\Delta)$ is diagonal and causes no problem. Now note that $B_{n-1,n}(t)$ is in GL_n and commutes with GL_{n-2} . Therefore, by Schur's lemma and the fact that our basis is symmetry adapted, the matrix $F(t) := D_\lambda(B_{n-1,n}(t))$ has a sparse block structure. The same is true for the infinitesimal generator Γ of the one-parameter subgroup $(\mathbb{C}, +) \rightarrow \mathrm{GL}_{d_\lambda}$, $t \mapsto F(t)$. (Note that $F(t) = e^{t\Gamma}$.) The block structure turns out to be even finer by taking into account the action on the weight spaces. The point is now that Γ can be transformed efficiently by a block diagonal matrix to a direct sum of nilpotent Jordan blocks $J_1 \oplus \dots \oplus J_s$. On the other hand, the e^{tJ_σ} are Toeplitz matrices. We may now use a well-known FFT-based algorithm for multiplying $r \times r$ Toeplitz matrices with a vector with only $O(r \log r)$ arithmetic operations. These reasonings yield an algorithm for computing $D_\lambda(A)v$ for given $A \in \mathrm{GL}_m$ and $v \in \mathbb{C}^{d_\lambda}$ using only

$$O(m^2(\mathrm{mult}(\lambda) + \log |\lambda|) d_\lambda)$$

arithmetic operations, provided $\lambda \in \mathbb{N}^m$ is not constant. (We say that a partition λ is constant iff its components are all equal. For the definition of the invariant $\mathrm{mult}(\lambda)$ see Sect. 2.) The number of nonscalar operations can be bounded by $O(m^2 d_\lambda + m \lambda_1)$. Hereby, we assume exact arithmetic of complex numbers. We have not yet studied the effect of round-off errors.

We remark that our algorithm was inspired by Clausen's fast Fourier transform [6] for the symmetric group, as well as Maslen's extension to compact Lie groups, see the survey [14]. Their techniques also rely heavily on symmetry adapted bases.

In Section 4 we complement our algorithmic result by proving that d_λ nonscalar operations are indeed necessary for the computation of $D_\lambda(A)v$. This is easily obtained by combining Burnside's lemma with the dimension bound of algebraic complexity. It shows that our algorithm is optimal up to a factor of m^2 with respect to nonscalar complexity.

Our algorithm provides already in the special case of GL_2 results which are, to our best knowledge, new. We obtain a fast rational $O(\ell \log \ell)$ algorithm for computing all the associated Legendre functions $P_\ell^\mu(\cos \theta)$, $|\mu| \leq \ell$, of degree ℓ from $\cos \theta$ and $\sin \theta$. (See Sect. 5.)

For computing individual entries of the invariant matrix, the cost of our algorithm may appear to be prohibitively high, mainly because the dimension d_λ can be very large. For instance, for $\lambda = (m, 0, \dots, 0) \in \mathbb{N}^m$ we have $\mathrm{mult}(\lambda) = 1$ and $d_\lambda = \binom{2m-1}{m}$, which is exponential in m . Is this inherent to the problem, or are there faster algorithms running with a number of steps polynomially bounded in m ? For approaching this question, we do not focus on individual entries of the invariant matrix, but we study related functions having some invariant meaning. Namely, we consider the sum of the diagonal entries of invariant matrices corresponding to the weight $(1, \dots, 1)$. These turn out to be the so-called immanant polynomials. Our algorithm yields upper bounds on the computational complexity

of immanants, which improve previous upper bounds due to Hartmann [10] and Barvinok [1]. This is discussed in Section 6.

Finally, in Section 7 we describe attempts to prove the hardness of immanants by reducing them to permanents. This fits well into the framework of Valiant’s algebraic P-NP theory [17] (see also [5, Chap. 21]). Valiant’s hypothesis claims that the permanent of m by m matrices cannot be computed by arithmetic straight-line programs with a number of steps polynomially bounded in m . We generalize a result of Hartmann [10] by proving that (families of) immanants corresponding to certain hook diagrams are VNP-complete. Moreover, we succeed in proving that immanants corresponding to square diagrams (or more generally, to certain rectangular diagrams) are also VNP-complete, thereby solving an open problem posed by Strassen [16, Problem 14.2]. The meaning of these results is that these immanants cannot be computed with a polynomial number of arithmetic operations if Valiant’s hypothesis is true.

Acknowledgements. I am grateful to M. Clausen and A. Shokrollahi for pointing out to me the papers by Grabmeier and Kerber, and Barvinok, respectively.

2 Preliminaries

We collect first some facts about the representations of the general linear group $\mathrm{GL}_m := \mathrm{GL}(m, \mathbb{C})$ (see [3, 7]). Consider GL_{m-1} as the subgroup of GL_m fixing the last canonical basis vector e_m . A polynomial representation $D_\lambda: \mathrm{GL}_m \rightarrow \mathrm{GL}(V)$ with highest weight $\lambda \in \mathbb{N}^m$ restricted to GL_{m-1} splits according to the branching rule [3, V§6] into a direct sum of representations with highest weights $\mu \in \mathbb{N}^{m-1}$, which satisfy the betweenness conditions $\lambda_j \geq \mu_j \geq \lambda_{j+1}$ for $1 \leq j < m$. It is important that each representation corresponding to such μ occurs with *multiplicity one*. Thus the decomposition of V into corresponding submodules V_μ is unique. Moreover, note that this is also the decomposition of V restricted to the subgroup $\mathrm{GL}_{m-1} \times \mathbb{C}^\times$: the diagonal matrix $\mathrm{diag}(1, \dots, 1, t)$ operates on V_μ by multiplication with $t^{|\lambda| - |\mu|}$. (We write $|\lambda| := \sum_i \lambda_i$.)

If we iterate this procedure by restricting according to the chain of subgroups

$$\mathrm{GL}_m > \mathrm{GL}_{m-1} \times \mathbb{C}^\times > \mathrm{GL}_{m-2} \times (\mathbb{C}^\times)^2 > \dots > (\mathbb{C}^\times)^m, \quad (1)$$

we finally end up with a decomposition of V into one-dimensional subspaces of weight vectors. This decomposition is unique, and bases of V adapted to this decomposition are called *Gelfand-Tsetlin bases*. (Recall that a vector $v \in V$ is said to be of weight $w \in \mathbb{N}^m$ iff $D_\lambda(\mathrm{diag}(t_1, \dots, t_m))v = t_1^{w_1} \dots t_m^{w_m} v$.) The splitting behaviour can be conveniently visualized by a layered graph $G(\lambda)$, whose nodes on level n ($0 \leq n \leq m$) are the occurring irreducible representations of V restricted to $\mathrm{GL}_n \times (\mathbb{C}^\times)^{m-n}$. These nodes can thus be uniquely described by pairs (ν, w) , where $\nu \in \mathbb{N}^n$ is a partition and $w \in \mathbb{N}^{m-n}$ satisfies $|\nu| + |w| = |\lambda|$. A node on level n is connected in the graph $G(\lambda)$ with a node on level $n-1$ if the latter appears in the decomposition of the former upon restriction to $\mathrm{GL}_{n-1} \times (\mathbb{C}^\times)^{m-n+1}$ (see Fig. 1).




Fig. 1. The graph $G(\lambda)$ for $\lambda = (2, 1, 0)$, $m = 3$.

The number of paths in $G(\lambda)$ between a node $x = (\nu, w)$ at level n and a node $x' = (\nu', w')$ at level $n' < n$ is just the multiplicity with which x' occurs in x when restricted to $\mathrm{GL}_{n'} \times (\mathbb{C}^\times)^{m-n'}$. We denote this multiplicity by $\mathrm{mult}(x, x')$. We call the maximum of $\mathrm{mult}(x, x')$ taken over all pair of nodes *two levels apart* the *multiplicity* $\mathrm{mult}(\lambda)$ of λ . (In Fig. 1 we have $\mathrm{mult}(\lambda) = 2$.)

The vectors of a Gelfand-Tsetlin basis can be labelled by paths in $G(\lambda)$ going from the top node λ to a node at level zero. Such paths can be encoded as semi-standard tableaux: for instance, in Fig. 1 we have two vectors of weight $(1, 1, 1)$ corresponding to the tableaux $\begin{array}{|c|c|} \hline 1 & 2 \\ \hline 3 & \\ \hline \end{array}$ and $\begin{array}{|c|c|} \hline 1 & 3 \\ \hline 2 & \\ \hline \end{array}$. The quantity $\mathrm{mult}(x, x')$ can thus be alternatively described as the number of semistandard tableaux of the skew diagram $\nu \setminus \nu'$ in which j occurs exactly w'_j times ($n' < j \leq n$).

By proper scaling of the vectors of a Gelfand-Tsetlin basis, we can obtain a basis of V , which is *adapted* to the chain (1) of subgroups. This means that the corresponding matrix representation D of GL_m has the following properties for all n :

1. The restriction $D \downarrow \mathrm{GL}_n \times (\mathbb{C}^\times)^{m-n}$ is *equal* to a direct sum of matrix representations of this subgroup.
2. Equivalent irreducible constituents of $D \downarrow \mathrm{GL}_n \times (\mathbb{C}^\times)^{m-n}$ are *equal*.

These additional requirements are crucial for our computational purpose. For convenience, we will call such adapted bases also Gelfand-Tsetlin bases. We remark that if V is a (finite dimensional) Hilbert space and D_λ restricted to $\mathrm{U}(m)$ is unitary, then a Gelfand-Tsetlin basis can be chosen to orthonormal.

Consider the matrix $B_{i,j}(t) \in \mathrm{GL}_m$ with entries 1 in the diagonal, entry t at position (i, j) , and having entries 0 elsewhere ($i \neq j$). Let $D: \mathrm{GL}_m \rightarrow \mathrm{GL}(V)$ be a rational representation. Then $F = F_{i,j}: \mathbb{C} \rightarrow \mathrm{GL}(V)$, $t \mapsto D(B_{i,j}(t))$ is a one-

parameter subgroup: we have $F(s+t) = F(s)F(t)$ for $s, t \in \mathbb{C}$. Hence $F'(t) = F'(0)F(t)$ and therefore $F(t) = e^{tF'(0)}$. (Note that $F'(0)$ must be nilpotent.) Let $\epsilon_i \in \mathbb{N}^m$ denote the basis vector having components 0 except at position i , where the component equals 1.

Lemma 1. $F'_{i,j}(0)$ maps a vector of weight $w \in \mathbb{Z}^m$ into one of weight $w + \epsilon_i - \epsilon_j$.

3 The algorithm

Our main result is expressed in the following theorem.

Theorem 2. Let D_λ be the matrix representation of GL_m with highest weight $\lambda \in \mathbb{N}^m$ with respect to a Gelfand-Tsetlin basis. We suppose that λ is not constant. Then the map

$$\mathrm{GL}_m \times \mathbb{C}^{d_\lambda} \longrightarrow \mathbb{C}^{d_\lambda}, (A, v) \mapsto D_\lambda(A)v$$

can be computed with $O(m^2(\mathrm{mult}(\lambda) + \log |\lambda|) d_\lambda)$ arithmetic operations. The non-scalar complexity is bounded by $O(m^2 d_\lambda + m \lambda_1)$.

Remark. 1. In the formulation of the algorithm, we assume exact arithmetic of complex numbers.

2. In the above upper bound the cost for *constructing* the algorithm (preconditioning) is not taken into account. The explicit formulas in Vilenkin and Klimyk [18, Chap. 18] for invariant matrices evaluated at special generators of GL_m suggest that this can also be done very efficiently.

3. We need the assumption that λ is not constant. Otherwise, the determinant could be evaluated with $O(m^2 \log m)$ operations (take $\lambda = (1, \dots, 1)$).

Corollary 3. The invariant matrix with respect to a nonconstant partition $\lambda \in \mathbb{N}^m$ and a Gelfand-Tsetlin basis can be evaluated at a matrix $A \in \mathrm{GL}_m$ with $O(m^2(\mathrm{mult}(\lambda) + \log |\lambda|) d_\lambda^2)$ arithmetic operations.

Before giving the proof of Thm. 2 we state some auxiliary results. The first one is easily obtained by Gaussian elimination.

Lemma 4. Any matrix $A \in \mathrm{GL}_m$ can be factored as $A = A_N A_{N-1} \cdots A_1 \Delta$, where $N \leq 2m^2$, Δ is a diagonal matrix, and all A_i are elementary matrices of the form $B_{n-1,n}(t)$ or $B_{n,n-1}(t)$. Moreover, such a decomposition can be computed with $O(m^3)$ arithmetic operations.

The next result is well-known and relies on the fast Fourier transform (see for instance [5, Cor. 13.13]).

Proposition 5. Suppose $J \in \mathbb{C}^{r \times r}$ is a nilpotent Jordan block. Then e^{tJ} is a Toeplitz matrix. Thus $e^{tJ}u$ can be computed from $t \in \mathbb{C}$ and $u \in \mathbb{C}^r$ with $O(r \log r)$ arithmetic operations. For this, $O(r)$ nonscalar operations are sufficient.

Part (b) of the following lemma follows easily from its Part (a) and Prop. 5.

Lemma 6. Let $M \in \mathbb{C}^{m \times m}$ be a matrix decomposed into r^2 blocks $M_{ij} \in \mathbb{C}^{m_i \times m_j}$, $m = m_1 + \dots + m_r$. Suppose that all block entries outside the lower diagonal are zero, that is, $M_{ij} = 0$ if $i \neq j + 1$. Then the following is true:

- (a) There is an invertible block diagonal matrix $S = [S_{ij}] \in \mathbb{C}^{m \times m}$ ($S_{ij} \in \mathbb{C}^{m_i \times m_j}$, $S_{ij} = 0$ for $i \neq j$) and a permutation matrix P such that the $PSMS^{-1}P^{-1}$ is a direct sum of nilpotent Jordan blocks of size at most r .
- (b) The product $e^{tM}u$ can be computed with $O(\sum_{\rho=1}^r m_\rho^2 + m \log r)$ arithmetic operations from $t \in \mathbb{C}$ and $u \in \mathbb{C}^m$.

Proof. (of Thm. 2) We first factor the given matrix $A \in \text{GL}_m$ according to Lemma 4 as $A = A_N A_{N-1} \dots A_1 \Delta$. Note that the cost for doing this is dominated by $O(m^2 d_\lambda)$, since $d_\lambda \geq m$ for a nonconstant $\lambda \in \mathbb{N}^m$. We therefore have $D_\lambda(A) = D_\lambda(A_N) D_\lambda(A_{N-1}) \dots D_\lambda(A_1) D_\lambda(\Delta)$. For given $v \in \mathbb{C}^{d_\lambda}$ we first compute $v_0 = D_\lambda(\Delta)v$ and then successively $v_i = D_\lambda(A_i)v_{i-1}$ for $1 \leq i \leq N$. Obviously, $v_i = D_\lambda(A)v$.

As a Gelfand-Tsetlin basis consists of weight vectors, the matrix $D_\lambda(\Delta)$ is diagonal with entries $t_1^{w_1} \dots t_m^{w_m}$, where $w_i \leq \lambda_1$. Thus $D_\lambda(\Delta)v$ can be certainly computed with $O(m d_\lambda \log \lambda_1) \leq O(m d_\lambda \log |\lambda|)$ arithmetic operations. (Note that $2 \log w_i$ multiplications are sufficient to obtain $t_i^{w_i}$.)

It remains to show that we can compute each of the products $D_\lambda(A_i)v_i$ with $O((\text{mult}(\lambda) + \log |\lambda|)d_\lambda)$ arithmetic operations. We assume that $A_i = B_{n-1,n}(t)$, the case of $A_i = B_{n,n-1}(t)$ being analogous. Let us write $V = \mathbb{C}^{d_\lambda}$ and interpret V as an GL_m -module via D_λ . Recall the graph $G(\lambda)$ introduced in Sect. 2, which describes the splitting behaviour of V . We have

$$V \downarrow \text{GL}_n \times (\mathbb{C}^\times)^{m-n} = \bigoplus_x \bigoplus_{j=1}^{f(x)} V_{x,j}, \quad (2)$$

where the first sum is over all nodes x of $G(\lambda)$ at level n , $f(x)$ equals $\text{mult}(\lambda, x)$, and $V_{x,j}$ is an irreducible $\text{GL}_n \times (\mathbb{C}^\times)^{m-n}$ -module of type x . Because of the symmetry adaptation, the decomposition (2) is compatible with our Gelfand-Tsetlin basis (i.e., subsets of this basis form a basis of each $V_{x,j}$). As A_i is contained in GL_n , $D_\lambda(A_i)$ decomposes according to (2) into

$$D_\lambda(A_i) = \bigoplus_x \bigoplus_j D_{x,j}(A_i).$$

If we write $v_i = \bigoplus_x \bigoplus_j v_{x,j}$ according to (2) (this decomposition can be done for free), we have $D_\lambda(A_i)v_i = \bigoplus_x \bigoplus_j D_{x,j}(A_i)v_{x,j}$. Now it is sufficient to show that each of the products $D_{x,j}(A_i)v_{x,j}$ can be computed with $O((\text{mult}(\nu_x) + |\nu_x|)d_{\nu_x})$ arithmetic operations, where $x = (\nu_x, w)$, $\nu_x \in \mathbb{N}^n$, $d_{\nu_x} = \dim V_{x,j}$. (Indeed, then $D_\lambda(A_i)v_i$ can be computed with a number of arithmetic operations bounded by

$$\sum_x (\text{mult}(\nu_x) + |\nu_x|) f(x) d_{\nu_x} \leq (\text{mult}(\lambda) + |\lambda|) \sum_x f(x) d_{\nu_x} = (\text{mult}(\lambda) + |\lambda|) d_\lambda$$

up to a constant factor.) By symmetry adaptation, a subset of the original Gelfand-Tsetlin basis forms a Gelfand-Tsetlin basis of $V_{x,j}$ and $D_{x,j}$ is the corresponding

From Lemma 6(b) we know that a product $e^{tM^\nu} u$ can be computed from $t \in \mathbb{C}$ and $u \in \mathbb{C}^{g(\nu)}$ with a number of arithmetic operations bounded by

$$\sum_a f(\nu, a)^2 + g(\nu) \log(|\lambda| - |\nu| + 1) \leq g(\nu) \text{mult}(\lambda) + g(\nu) \log(|\lambda| + 1)$$

up to a constant factor. Therefore, a product $F(t)v$ can be computed with

$$O\left(\sum_\nu d_\nu g(\nu) (\text{mult}(\lambda) + \log(|\lambda| + 1))\right)$$

arithmetic operations, which proves the claim, as $d_\lambda = \sum_\nu d_\nu g(\nu)$.

The estimation of the nonscalar complexity is similar. □

4 Lower bounds

The subsequent lower bound result shows that our algorithm is optimal up to a factor of m^2 with respect to the number of nonscalar operations.

Theorem 7. *Let $\lambda \in \mathbb{N}^m$ be a partition with $|\lambda| > 1$ and $v \in \mathbb{C}^{d_\lambda}$ be nonzero. Let D_λ denote a matrix representation of GL_m with highest weight λ . Then any arithmetic algorithm (formally, algebraic computation tree) computing the map $\text{GL}_m \rightarrow \mathbb{C}^{d_\lambda}$, $A \mapsto D_\lambda(A)v$ requires at least d_λ nonscalar operations. The evaluation of the invariant matrix $\text{GL}_m \rightarrow \text{GL}_{d_\lambda}$, $A \mapsto D_\lambda(A)$ requires at least d_λ^2 nonscalar operations.*

Proof. The theorem of Burnside (cf. [11]) states that the linear hull of the image of D_λ equals $\mathbb{C}^{d_\lambda \times d_\lambda}$, as D_λ is irreducible. Therefore, the entries of the invariant matrix corresponding to λ are linearly independent polynomials of degree $|\lambda| > 1$. The dimension bound in [5, (4.12)] easily implies the claims. □

5 Fast evaluation of Legendre functions

The algorithm of Section 3 can be applied to evaluate many special functions and orthogonal polynomials (Legendre, Jacobi, Gegenbauer polynomials, generalized Beta functions, etc.), since all these are matrix entries of a suitable representation of GL_m (compare [18]).

We illustrate this here just by the example of the associated Legendre functions. Let D^ℓ be an irreducible representation of GL_2 with highest weight $\lambda = (2\ell, 0)$ corresponding to a Gelfand-Tsetlin basis ($\ell \in \mathbb{N}$). Consider the unitary matrix ($0 \leq \theta < \pi$)

$$A(\theta) := \begin{pmatrix} \cos \theta/2 & i \sin \theta/2 \\ i \sin \theta/2 & \cos \theta/2 \end{pmatrix}.$$

It turns out that the column of $D^\ell(A(\theta))$ corresponding to the weight (ℓ, ℓ) just contains, up to some scalar factors, the associated Legendre functions $P_\ell^\mu(\cos \theta)$, $|\mu| \leq \ell$ (see [18, §6.2]). Note that $P_\ell^\mu(\cos \theta)$ is a polynomial in $\cos \theta$ and $\sin \theta$.

Corollary 8. *All the associated Legendre functions $P_\ell^\mu(\cos \theta)$, $|\mu| \leq \ell$, can be computed from $\cos \theta$ and $\sin \theta$ by an algorithm using only rational operations with $O(\ell \log \ell)$ arithmetic operations. The nonscalar complexity is bounded by $O(\ell)$.*

We remark that the obvious algorithm using a well-known recursion formula for P_ℓ^μ needs $O(\ell^2)$ arithmetic operations.

6 Computation of immanants

The *immanant* of a matrix $A \in \mathbb{C}^{m \times m}$ corresponding to a partition λ of m is defined as (cf. [13])

$$\text{im}_\lambda(A) = \sum_{\pi \in S_m} \chi_\lambda(\pi) \prod_{i=1}^m A_{i, \pi(i)} ,$$

where χ_λ denotes the irreducible character of the symmetric group S_m belonging to λ . Note that this notion contains the permanent and determinant as special cases. The next lemma states that immanants appear naturally as the sum of the diagonal entries of invariant matrices corresponding to the weight $(1, \dots, 1)$. The proof easily follows from Thm. 2.4 in [9].

Lemma 9. *Let λ be a partition of m . Suppose $D_\lambda = [D_{i,j}]$ is an irreducible matrix representation of GL_m of type λ with respect to a basis of weight vectors. Then we have for all $A \in \text{GL}_m$ that*

$$\text{im}_\lambda(A) = \sum_i D_{i,i} ,$$

where the sum is over all i corresponding to weight vectors of weight $(1, \dots, 1)$.

Remark. If we extend the above sum over all indices, we get the character of D_λ evaluated at A . This value can be computed by a *polynomial number* of arithmetic operations (in m). This follows readily from Giambelli's formula (cf. [7]), which expresses a Schur polynomial as a determinant of a matrix of size m involving elementary symmetric functions.

By combining Thm. 2 with the above lemma we get the following result.

Theorem 10. *Let λ be a nonconstant partition of m and s_λ denote the number of standard tableaux on the diagram of λ . One can compute $\text{im}_\lambda(A)$ from $A \in \text{GL}_m$ with a number of arithmetic operations bounded by*

$$O(m^2(\text{mult}(\lambda) + \log m) s_\lambda d_\lambda) .$$

The nonscalar complexity is bounded by $O(m^2 s_\lambda d_\lambda)$.

Remark. Hartmann [10] proved the upper bound $m^{6(m-s)+4}$ for the nonscalar complexity to evaluate permanents corresponding to partitions with at most s parts. Barvinok [1] showed the upper bound $O(m^3 d_\lambda^4)$ for the total complexity. Our bound improves Barvinok's, as $s_\lambda \leq d_\lambda$ and $\text{mult}(\lambda) \leq d_\lambda$.

To compare our bound with those of Hartmann, consider hook partitions $\lambda = (k, 1, \dots, 1) \in \mathbb{N}^m$. For such λ one can show that

$$s_\lambda = \binom{m-1}{k-1}, \quad d_\lambda = \frac{m+k-1}{m} \binom{m+k-2}{m-k, k-1, k-1},$$

hence $m^2 s_\lambda d_\lambda \leq m^2 18^m$. This is considerably smaller than Hartmann's bound m^{6k} if $k \geq m/2$.

For permanents Thm. 10 yields the bound $O(m^{1.54^m} \log m)$, which is not too far away from the best known upper bound $O(m2^m)$ due to Ryser [15].

7 Lower bounds for the complexity of immanants

We discuss the computational complexity for evaluating immanants within the framework of Valiant's algebraic P-NP theory [17]. (For a detailed account see [5, Chap. 21].) In this theory, a hierarchy of complexity classes $\text{VP} \subseteq \text{VNP}$, as well as a quasi-order \leq_p on VNP (p -projection) is defined. VP consists of the families of multivariate polynomials $f = (f_n)$, which are computable with a number of arithmetic steps polynomially bounded in n . One can prove that there exist families which are maximal in VNP with respect to \leq_p ; they are said to be VNP-complete. Natural examples are provided by the families of permanent polynomials and Hamilton cycle polynomials. Valiant's hypothesis claims that VP is strictly contained in VNP. We note that VNP-complete families do not lie in VP under this hypothesis.

For giving evidence that certain polynomials cannot be computed with a polynomial number of arithmetic operations, one can try to prove that they define a VNP-complete family. That is what we have done for certain immanants.

The subsequent theorem generalizes a result by Hartmann [10].

Theorem 11. *Let (k_m) be a sequence of natural numbers satisfying $m^\epsilon \leq k_m \leq m$ for some positive ϵ . Then the sequence of immanant polynomials corresponding to the hook partitions $(k_m, 1, \dots, 1) \in \mathbb{N}^m$ is VNP-complete.*

Our next result answers a question posed by Strassen [16, Problem 14.2].

Theorem 12. *Let (s_n) be a polynomially bounded sequence of natural numbers. Then the sequence of immanant polynomials belonging to the rectangular partitions $(n, \dots, n) \in \mathbb{N}^{n \times s_n}$ is VNP-complete.*

Our central tool for proving completeness is a consequence of the Murnaghan-Nakayama rule for the characters of the symmetric group. It allows to identify certain alternating sums of immanants corresponding to smaller partitions as a projection of a given immanant. Using this strategy repeatedly, it is possible to obtain a permanent or Hamilton cycle polynomial (of large size) as a projection of the immanant under investigation. Due to lack of space we can not provide more details. Proofs of the Thms. 11 and 12 can be found in the preprint [4].

The complexity of immanants is still full of mysteries. For instance, we do not even know whether the sequence of immanant polynomials belonging to the partitions $(2, 2, \dots, 2) \in \mathbb{N}^{2n}$ is VNP-complete. We have the following conjecture.

Conjecture 13. Let $(\lambda^{(n)})$ be a sequence of partitions such that $|\lambda^{(n)}| = n$ and $\max \lambda^{(n)} \geq n^\epsilon$ for some positive ϵ . Then the corresponding sequence of immanant polynomials is VNP-complete.

References

1. A.T. Barvinok. Computational complexity of immanants and representations of the full linear group. *Functional Analysis Applications*, 24:144–145, 1990.
2. L.C. Biedenharn and J.D. Louck. *Angular Momentum in Quantum Physics: Theory and Application; The Racah-Wigner Algebra in Quantum Theory*. Number 8 and 9 in Encyclopedia of Mathematics and its Applications. Addison-Wesley, 1981.
3. H. Boerner. *Representations of groups*. Elsevier North-Holland, Amsterdam, 1970.
4. P. Bürgisser. The computational complexity of immanants. Preprint, University of Zurich, 1997.
5. P. Bürgisser, M. Clausen, and M.A. Shokrollahi. *Algebraic Complexity Theory*. Number 315 in Grundlehren der mathematischen Wissenschaften. Springer Verlag, 1996.
6. M. Clausen. Fast generalized Fourier transforms. *Theoret. Comp. Sci.*, 67:55–63, 1989.
7. W. Fulton and J. Harris. *Representation Theory*. Number 129 in GTM. Springer Verlag, 1991.
8. I.M. Gelfand and M.L. Tsetlin. Finite dimensional representations of the group of unimodular matrices. *Dokl. Akad. Nauk SSSR*, 71:825–828, 1950. (in Russian).
9. J. Grabmeier and A. Kerber. The Evaluation of Irreducible Polynomial Representations of the General Linear Groups and of the Unitary Groups over Fields of Characteristic 0. *Acta Applicandae Mathematicae*, 8:271–291, 1987.
10. W. Hartmann. On the complexity of immanants. *Linear and Multilinear Algebra*, 18:127–140, 1985.
11. S. Lang. *Algebra*. Addison-Wesley, second edition, 1984.
12. D.E. Littlewood. The construction of invariant matrices. *Proc. London Math. Soc.* (2), 43:226–240, 1937.
13. D.E. Littlewood. *The Theory of Group Characters and Matrix Representations of Groups*. Oxford University Press, 1940.
14. D.K. Maslen and D.N. Rockmore. Generalized FFTs - a survey of some recent results. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 28:183–237, 1997.
15. H. J. Ryser. *Combinatorial Mathematics*. Number 14 in Carus Math. Monographs. Math. Assoc. America, New York, 1963.
16. V. Strassen. Algebraic complexity theory. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume A, chapter 11, pages 634–672. Elsevier Science Publishers B. V., Amsterdam, 1990.
17. L.G. Valiant. Completeness classes in algebra. In *Proc. 11th ACM STOC*, pages 249–261, 1979.
18. N.J. Vilenkin and A.U. Klimyk. *Representations of Lie groups and special functions, I, II, and III*. Number 72, 74, and 75 in Mathematics and Its Applications. Kluwer Academic Publishers, Dordrecht, 1991, 1992, and 1993.

This article was processed using the L^AT_EX macro package with LLNCS style