# Algebraic Metacomplexity and Representation Theory

**Maxim van den Berg**
Ruhr University Bochum, Germany
University of Amsterdam, The Netherlands

**Pranjal Dutta**
Nanyang Technological University, Singapore

**Fulvio Gesmundo**
University of Toulouse, France

**Christian Ikenmeyer**
University of Warwick, UK

**Vladimir Lysikov**
Ruhr University Bochum, Germany

## Abstract

In the algebraic metacomplexity framework we prove that the decomposition of metapolynomials into their isotypic components can be implemented efficiently, namely with only a quasipolynomial blowup in the circuit size. We use this to resolve an open question posed by Grochow, Kumar, Saks & Saraf (2017). Our result means that many existing algebraic complexity lower bound proofs can be efficiently converted into isotypic lower bound proofs via highest weight metapolynomials, a notion studied in geometric complexity theory. In the context of algebraic natural proofs, it means that without loss of generality algebraic natural proofs can be assumed to be isotypic. Our proof is built on the Poincaré–Birkhoff–Witt theorem for Lie algebras and on Gelfand–Tsetlin theory, for which we give the necessary comprehensive background.

**2012 ACM Subject Classification** Theory of computation → Algebraic complexity theory

**Keywords and phrases** Algebraic complexity theory, metacomplexity, representation theory, geometric complexity theory

**Digital Object Identifier** 10.4230/LIPIcs.CCC.2025.26

**Funding** *Maxim van den Berg*: supported by the European Research Council through an ERC Starting Grant (grant agreement No. 101040907, SYMOPTIC) and the Dutch National Growth Fund (NGF) as part of the Quantum Delta NL visitor programme.
*Pranjal Dutta*: The work was done while PD was at the National University of Singapore, funded by NRF Singapore under the project "Computational Hardness of Lattice Problems and Implications" [NRF-NRFI09-0005].
*Christian Ikenmeyer*: supported by EPSRC grants EP/W014882/1 and EP/W014882/2.
*Vladimir Lysikov*: supported by the European Research Council through an ERC Starting Grant (grant agreement No. 101040907, SYMOPTIC).
Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

**Acknowledgements** We are grateful to Alessandro Danelon and Nutan Limaye for helpful discussions. We thank anonymous referees for their comments, and especially for bringing our attention to Open Question 2 in [40].



40th Computational Complexity Conference (CCC 2025).
Editor: Srikanth Srinivasan; Article No. 26; pp. 26:1–26:35

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

# 1 Introduction

Strong computational complexity lower bounds are often difficult to prove, for many different complexity measures. In most cases, the inability to prove lower bounds can be explained by barrier results such as the *Relativization barrier* due to Baker, Gill and Solovay [5], the *Algebrization barrier* due to Aaronson and Wigderson [1] and the *Natural Proofs barrier* due to Razborov and Rudich [66]. Metacomplexity studies the complexity of computing various complexity measures, such as the minimum circuit complexity of a function. One eventual goal of the metacomplexity program is to understand mathematically why lower bounds are hard to prove. In this paper, we study metacomplexity in the algebraic circuit complexity setup from a representation-theoretic viewpoint.

Algebraic complexity focuses on proving complexity lower bounds for various problems related to the computation of multivariate polynomials and rational functions. Prominent examples include the discrete Fourier transform, matrix multiplication, and solving systems of linear equations; see [14]. In algebraic complexity, the main model of computation is the *algebraic circuit*, which is a labeled directed graph encoding an algorithm to evaluate a polynomial. Most polynomials require large algebraic circuits; however, exhibiting explicit ones with such property is difficult. Valiant's flagship conjecture [71], known as $\mathsf{VP} \neq \mathsf{VNP}$, states that the algebraic circuit complexity of the permanent polynomial $\mathsf{per}_d(x_{1,1}, \ldots, x_{d,d}) = \sum_{\pi \in \mathfrak{S}_d} \prod_{i=1}^{d} x_{i,\pi(i)}$ is not polynomially bounded as a function of $d$. This conjecture is an algebraic analogue of the $\mathsf{P} \neq \mathsf{NP}$ conjecture: indeed, the permanent of the adjacency matrix of a bipartite graph counts the number of perfect matchings, which is a #P-hard problem.

Many current proof techniques in algebraic complexity theory are implicitly based on *metapolynomials*. Examples of such techniques are *rank methods*, dating back to [70], and used more recently in complexity theory both for polynomials [62, 55] and for tensors [69, 54], methods based on geometric features [53, 36], and the geometric complexity program [60, 59]. We refer to Section 2.2 for a more extensive discussion. Informally, metapolynomials are polynomials whose variables are coefficients of polynomials. The concept dates back to resultants and invariants of forms [20, 70], while the term "metapolynomial" was coined in [67] and made popular in [40]. A classical example is the metapolynomial $b^2 - 4ac$ in the coefficients of a univariate degree 2 polynomial $ax^2 + bx + c$, which vanishes exactly when the two roots of the polynomial coincide. Grochow pointed out in [39] that for many border complexity measures the complexity lower bounds can without loss of generality be proved by metapolynomials that lie in separating modules, or are isotypic, or are highest weight vectors. These are all closely related concepts (and we prove analogous results for all of them), where being isotypic is the most natural one because it does not depend on choosing any bases. Being a *homogeneous* metapolynomial is the same as being isotypic with respect to the action of $1 \times 1$ matrices. Being isotypic also generalizes the property of being *invariant* under the special linear group. While invariants can have large circuit size, see [34], we prove that not all separating *isotypic* metapolynomials have large circuit size, unless *all* separating metapolynomials have large circuit size. More precisely, we prove that without loss of generality one may assume the metapolynomial to be isotypic, or to be a highest weight vector, or to lie in a separating module, with only a quasipolynomial blowup in the circuit size. We use this to answer Open Question 2 in [40] by proving that *all* metapolynomials that lie in an irreducible representation have the same circuit complexity up to quasipolynomial blowup, see Corollary 1.2.

Our result reduces the search space for lower bounds from *all* homogeneous metapolynomials to the isotypic metapolynomials, which is a technique for example used in [2, 15, 16, 43, 3, 23, 11]. In particular the large-scale cluster computations in [3, 23, 11]

would be infeasible without this search space reduction. For example, [43] reduces the search space for proving that the border complexity of the $2 \times 2$ matrix multiplication tensor is 7 to a 4-dimensional linear subspace of isotypic metapolynomials in a $\binom{64+20-1}{20} = 8,179,808,679,272,664,720$-dimensional space of homogeneous degree 20 metapolynomials on $4 \times 4 \times 4$ tensors. On the reduced search space, the problem is then easily solved by linear algebra. This specific approach was used in [8] to prove that the border support rank of the $2 \times 2$ matrix multiplication tensor is 7. To this date, this is the only proof of this fact. Similar methods are used in [63, 37], where the polynomials of interest are invariants for the action of a special linear group: this information is crucially used to reduce the search space and determine the polynomial via interpolation methods on a small space. These are finite results, but the techniques work in the asymptotic settings as well, for example for lower bounds on the rank of matrix multiplication as in [15, 16], or the best known lower bound for the border determinantal complexity of the permanent [53].

Our result also points to potential future lower bounds: Our method can be used to extract isotypic metapolynomials for example from rank-based lower bounds proofs, and the obtained metapolynomials give at least the same lower bounds, but potentially better ones. It is conceivable that these metapolynomials can be adjusted to *not* be affected by barriers for rank methods [26, 33].

Our procedure uses explicit efficient circuit constructions and hence makes a contribution to algorithmic representation theory that is of independent interest.

## Metapolynomials

We work over the field $\mathbb{C}$ of complex numbers, and the set $\{x_1, x_2, \ldots\}$ of variables. For every sequence $\boldsymbol{i}$ of nonnegative integers with $\sum_j i_j < \infty$, a monomial $x^{\boldsymbol{i}}$ is the finite product of variables $x_1^{i_1} x_2^{i_2} \cdots$, and a polynomial $f$ is a finite linear combination of monomials. For every monomial $x^{\boldsymbol{i}} = x_1^{i_1} x_2^{i_2} \cdots$, define a *metavariable* $c_{\boldsymbol{i}}$. A metavariable $c_{\boldsymbol{i}}$ has degree 1, and it has *weight* $\boldsymbol{i}$. Finite products of metavariables are called metamonomials, and finite linear combinations of metamonomials are called metapolynomials. The weight of a metamonomial is the sum of the weights of its metavariables. If every metamonomial in a metapolynomial $\Delta$ has the same weight, then $\Delta$ is called a weight metapolynomial (and the zero metapolynomial has every weight). For a fixed weight $w$ the metapolynomials of that weight form a vector space, the weight $w$ space. Every metapolynomial can be written as a unique sum of weight metapolynomials of different weights. A metapolynomial is called affine linear if all its metamonomials have degree at most 1, i.e., are just a metavariable or a constant. Every metapolynomial $\Delta$ can be evaluated at any polynomial $f$ by assigning to each metavariable the coefficient of the corresponding monomial in $f$. For example, the discriminant of a degree two homogeneous polynomial in two variables $c_{20}x_0^2 + c_{11}x_0x_1 + c_{02}x_1^2$ is the metapolynomial $\Delta = c_{11}^2 - 4c_{20}c_{02}$. It is a classical fact that $\Delta(f) = 0$ if and only if $f = \ell^2$ for some homogeneous linear polynomial $\ell = \alpha_0 x_0 + \alpha_1 x_1$. In principle, metapolynomials can involve metavariables that correspond to monomials of different degrees, but for our purposes it is sufficient to consider the case where all metavariables correspond to monomials of the same degree $d$, see Remark 1.3. Every such metapolynomial can be decomposed uniquely as a sum of its homogeneous degree $\delta$ components, and we call a homogeneous degree $\delta$ metapolynomial a metapolynomial of format $(\delta, d, k)$ if all its metavariables only involve the variables $x_1, \ldots, x_k$.

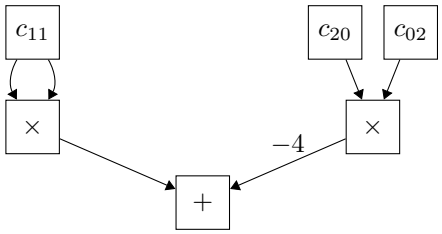


**Figure 1** An algebraic circuit computing the discriminant $\Delta = c_{11}^2 - 4c_{20}c_{02}$.

## Algebraic circuits

In this paper, we are mostly interested in the algebraic circuit size of metapolynomials. Algebraic circuits for metapolynomials are defined in the same way as for usual polynomials, as follows.

An algebraic circuit is a directed acyclic graph in which each vertex has indegree 0 or 2. Every indegree 0 vertex is called *input gate* and it is labelled by an affine linear metapolynomial; every indegree 2 vertex is called a *computation gate* and it is labeled by either "+" or "×"; every edge is labeled by a complex number, which is assumed to be 1 if this scalar is omitted; there is exactly one vertex of outdegree 0. An algebraic circuit computes a metapolynomial at every vertex by induction over the directed acyclic graph structure (the labels on the edges rescale the computed values), i.e., if the values computed at the children of a gate $v$ are $\Delta_1$ and $\Delta_2$ and the incoming edges to $v$ are labelled with $\alpha_1$ and $\alpha_2$, respectively, then the value computed at $v$ is $\alpha_1\Delta_1 + \alpha_2\Delta_2$ or $\alpha_1\alpha_2\Delta_1\Delta_2$, depending on whether $v$ is labelled with a "+" or a "×". An algebraic circuit *computes* the metapolynomial computed at its outdegree 0 vertex. The *size* of a circuit is the total number of its vertices. The algebraic circuit complexity of a metapolynomial is the *minimum* size of an algebraic circuit computing it. For example, Figure 1 shows that the algebraic circuit complexity of the discriminant is at most 6. Algebraic circuits for polynomials are defined analogously.

## Representation Theory

Let $\mathbb{N}$ denote the set of natural numbers including zero. A *partition* $\lambda = (\lambda_1, \lambda_2, \ldots)$ is a finite sequence of non-increasing natural numbers. Let $\ell(\lambda) = \max\{j \mid \lambda_j \neq 0\}$. We define $\lambda_j = 0$ for all $j > \ell(\lambda)$. Write $|\lambda| := \sum_j \lambda_j$, and say that $\lambda$ is a partition of $|\lambda|$ of length $\ell(\lambda)$. In this case we write $\lambda \vdash |\lambda|$ or $\lambda \vdash_{\ell(\lambda)} |\lambda|$ depending on whether the length is relevant. For example, $\lambda = (4, 2, 2)$ is a partition of 8 of length 3, i.e., $\lambda \vdash_3 8$. The *Young diagram* of a partition $\lambda$ is a top-left justified array of boxes with $\lambda_j$ many boxes in row $j$. For example, the Young diagram for $(4, 2, 2)$ is  . A *Young tableau* $T$ is a filling of the boxes of a Young diagram with positive integer numbers. The partition $\lambda$ is called the *shape* of $T$. For example,  is a Young tableau of shape $(4, 2, 2)$. A Young tableau is *semistandard* if the numbers within each row from left to right are non-decreasing, and the numbers within each column from top to bottom are strictly increasing, for example  . For each partition $\lambda$ there exists a unique *superstandard* tableau, that is, the tableau with only the number $j$ in row $j$, e.g.,  .

The group $\mathrm{GL}_k$ acts on the space $\mathbb{C}[x_1, \ldots, x_k]_d$ of homogeneous polynomials on $\mathbb{C}^k$ of degree $d$ by linear change of coordinates. This action induces, again by linear change of coordinates, an action on the space of metapolynomials with format $(\delta, d, k)$ for every $\delta$. For instance, the change of basis swapping the two coordinates on $\mathbb{C}^2$, defines the change

of basis on $\mathbb{C}[x_0, x_1]_2$ obtained by exchanging $x_0$ and $x_1$. This induces a linear change of coordinates on $\mathbb{C}[c_{20}, c_{11}, c_{02}]$ which exchanges $c_{20}$ and $c_{02}$ and leaves $c_{11}$ fixed. The discriminant $c_{11}^2 - 4c_{02}c_{20}$ is sent to itself by this change of coordinates. More precisely the group actions are defined as follows. For every homogeneous degree $d$ polynomial $f$ in $k$ variables and every $g \in \mathrm{GL}_k$, the polynomial $g \cdot f$ is defined by $(g \cdot f)(x) = f(g^{-1}x)$ for every $x \in \mathbb{C}^k$. Similarly, for every format $(\delta, d, k)$ metapolynomial $\Delta$, the metapolynomial $g \cdot \Delta$ is defined by $(g \cdot \Delta)(f) = \Delta(g^{-1} \cdot f)$ for every homogeneous degree $d$ polynomial $f$. The details about this action are provided in Section 4.

The vector space of metapolynomials of format $(\delta, d, k)$ is closed under this action of $\mathrm{GL}_k$, and it decomposes into a direct sum of subspaces which are also closed under the action of $\mathrm{GL}_k$: these subspaces are called *isotypic components* and they are in one to one correspondence with partitions $\lambda \vdash_k \delta d$. The summand corresponding to the partition $\lambda$ is called the $\lambda$-isotypic component. This is called the isotypic decomposition. Each $\lambda$-isotypic component decomposes even further into a direct sum of subspaces with one summand for each semistandard tableau $T$ of shape $\lambda$. We call the component for $T$ the *$T$-isotypic* component. If $V$ is $\lambda$-isotypic and *irreducible* (i.e., has no nontrivial subrepresentations), then $\lambda$ is called the *isomorphism type* of $V$, and each $T$-isotypic component of $V$ is 1-dimensional. If $T$ is a superstandard tableau of shape $\lambda$, then the corresponding $T$-isotypic component is called the *highest weight metapolynomial vector space of weight $\lambda$*. Equivalently, a highest weight vector of weight $\lambda$ is a metapolynomial $\Delta$ with $g \cdot \Delta = t_1^{\lambda_1} \cdots t_k^{\lambda_k} \Delta$, for every $g \in \mathrm{GL}_k$ upper triangular with $t_1, \ldots, t_k$ on the main diagonal. The $\lambda$-isotypic component is known to be the linear span of the union of the orbits of all its highest weight vectors.

For example (not easily verifiable by hand, see Example 5.4), the space of metapolynomials of format $(\delta, d, k) = (3, 2, 3)$ decomposes into three nonzero isotypic components, corresponding to partitions $(6, 0, 0)$, $(4, 2, 0)$ and $(2, 2, 2)$ of $\delta \cdot d = 6$. The isotypic component $(4, 2, 0)$ further decomposes into three $T$-isotypic components for the three semistandard tableaux of shape $(4, 2, 0)$: $\begin{smallmatrix}1&1&2&2\\3&3\end{smallmatrix}$, $\begin{smallmatrix}1&1&2&3\\2&3\end{smallmatrix}$, $\begin{smallmatrix}1&1&3&3\\2&2\end{smallmatrix}$. The metapolynomial $\Delta = c_{200}c_{020}c_{002}$ can be written according to this decomposition as

$$\begin{aligned}
60\,\Delta \;&=\; 2\big(c_{101}^2c_{020} + 2c_{110}c_{101}c_{011} + c_{200}c_{011}^2 + c_{110}^2c_{002} + 2c_{200}c_{020}c_{002}\big)\\
&+\; \Big[\, 2\big(-c_{020}c_{101}^2 - 2c_{011}c_{101}c_{110} + 4c_{002}c_{110}^2 - c_{011}^2c_{200} + 8c_{002}c_{020}c_{200}\big)\\
&+\; \quad 0\\
&+\; \quad 5\big(c_{020}c_{101}^2 - c_{011}c_{101}c_{110} - c_{002}c_{110}^2 + c_{011}^2c_{200} + 4c_{002}c_{020}c_{200}\big)\,\Big]\\
&+\; 5\big(-c_{101}^2c_{020} + c_{110}c_{101}c_{011} - c_{200}c_{011}^2 - c_{110}^2c_{002} + 4c_{200}c_{020}c_{002}\big),
\end{aligned}$$

where the five summands from top to bottom are $\begin{smallmatrix}1&1&2&2&3&3\end{smallmatrix}$-isotypic, $\begin{smallmatrix}1&1&2&2\\3&3\end{smallmatrix}$-isotypic, $\begin{smallmatrix}1&1&2&3\\2&3\end{smallmatrix}$-isotypic, $\begin{smallmatrix}1&1&3&3\\2&2\end{smallmatrix}$-isotypic, and $\begin{smallmatrix}1&1\\2&2\\3&3\end{smallmatrix}$-isotypic, respectively. The fifth summand is a highest weight vector, because $\begin{smallmatrix}1&1\\2&2\\3&3\end{smallmatrix}$ is superstandard. All summands have weight (2,2,2), because $\Delta$ has weight (2,2,2).

▶ **Theorem 1.1** (Main theorem). *Let $\Delta\colon \mathbb{C}[x_1, \ldots, x_k]_d \to \mathbb{C}$ be a metapolynomial of format $(\delta, d, k)$ computed by an algebraic circuit of size $s$. Then*

1. *For every weight $\mu \in \mathbb{N}^k$, $|\mu| = d\delta$, the projection of $\Delta$ onto the weight space of weight $\mu$ can be computed by a circuit of size $O(s(\delta d)^{2k^3})$.*
2. *For every partition $\lambda \vdash d\delta$ the projection of $\Delta$ onto the $\lambda$-isotypic component can be computed by a circuit of size $O(sk^{2k^2}(\delta d)^{2k^3})$.*

3. *For every partition $\lambda \vdash d\delta$ the projection of $\Delta$ onto the highest weight space of weight $\lambda$ can be computed by a circuit of size $O(s(k+1)^{2k^2}(\delta d)^{2k^3})$.*
4. *For every semistandard tableau $T$ of shape $\lambda \vdash d\delta$ the projection of $\Delta$ onto the $T$-isotypic space can be computed by a circuit of size $O(sk^{2k^2}(\delta d)^{2k^4})$.*

In all applications of Theorem 1.1, we will always have that $k$ is logarithmic in the number of metavariables, which we explain in Section 2.5. Open Question 2 in [40] asks: given a sequence of irreducible representations of metapolynomials, what is the sequence of metapolynomials in it that has the lowest circuit complexity? Theorem 1.1 answers this question in a very satisfying way: *Every* such sequence of metapolynomials has the same circuit complexity up to quasipolynomial blowup, as seen in the following corollary.

▶ **Corollary 1.2.** *Let $s$ be the smallest circuit complexity of a nonzero metapolynomial in an irreducible $\mathrm{GL}_k$-representation $V$ of type $\lambda \vdash \delta d$. Then* every *metapolynomial in $V$ has circuit complexity at most $O(sk^{2k^2}(\delta d)^{2k^4+k^2})$.*

**Proof.** If $V$ is irreducible of isomorphism type $\lambda$, then it decomposes into a direct sum of 1-dimensional $T$-isotypic components, where $T$ ranges over all semistandard tableaux of shape $\lambda$. Take a nonzero metapolynomial $\Delta$ with minimal complexity in $V$ and apply a generic basis change, i.e., apply a generic element of $\mathrm{GL}_k$. Then independently project the result onto all $T$-isotypic components. We claim that since the basis change was generic, the projections are nonzero, and therefore form a basis of $V$. In fact, we can be more precise about how we pick the basis change element $g \in \mathrm{GL}_k$. For every semistandard tableau $T$, define $X_T := \{g \in \mathrm{GL}_k \mid g\Delta \text{ has zero } T\text{-isotypic projection}\}$. Such $X_T$ is Zariski closed in $\mathrm{GL}_k$, but it is not the whole $\mathrm{GL}_k$, otherwise the linear span of $\mathrm{GL}_k\Delta$ would be a nontrivial subrepresentation of the irreducible representation $V$. Therefore the finite union $X := \bigcup_T X_T$ is also a proper Zariski closed subset of $V$, and its complement $\mathrm{GL}_k \setminus X$ is a nonempty Zariski open subset. If we pick $g \in \mathrm{GL}_k \setminus X$, then the projections of $g\Delta$ to every $T$-isotypic space are nonzero. Moreover, by Theorem 1.1, all these projections have complexity at most $O(sk^{2k^2}(\delta d)^{2k^4})$. But the dimension of $V$ equals the number of semistandard tableaux of shape $\lambda$ with entries $1, \dots, k$, which is at most $(\delta d)^{k^2}$ (each tableaux has $k$ rows, and each row has length at most $\delta d$, so can be filled in at most $(\delta d)^k$ ways). ◀

▶ Remark 1.3. In general, algebraic complexity classes are defined for possibly nonhomogeneous polynomials. In this paper, we assume that polynomials and metapolynomials are homogeneous. However, one can lift the results to the nonhomogeneous setting as follows.

A polynomial $f(x_1, \dots, x_k)$ is called homogeneous if all its monomials have the same degree. For $d \geq \deg(f)$, define the degree $d$ homogenization

$$f^{\sharp d}(x_0, x_1, \dots, x_k) := x_0^d f(x_1/x_0, \dots, x_k/x_0).$$

If $f$ has an algebraic circuit of size $s$, then $f^{\sharp d}$ has an algebraic circuit of size at most $O(sd)$: $s(d+1)$ for extracting the homogeneous parts of $f$ via interpolation, $d$ operations to compute all values $x_0, x_0^2, \dots, x_0^d$, and another $d+1$ for multiplying the homogeneous components with the correct powers of $x_0$, and a final $d$ for adding up the results. Given a metapolynomial $\Delta$ that does not involve $x_0$, define $\Delta^{\sharp d}$ by replacing every metavariable $c_{\boldsymbol{i}}$ with $c_{(d-|\boldsymbol{i}|,i_1,i_2,\dots)}$, and setting every metavariable $c_{\boldsymbol{i}}$ with $|\boldsymbol{i}| > d$ to zero. Clearly we have $\Delta^{\sharp d}(f^{\sharp d}) = \Delta(f)$, provided that $d \geq \deg(f)$. Hence, in this paper we can always assume that $f$ is homogeneous.

### Organization of the paper

In the first part of Section 2 we illustrate the motivations for studying the circuit complexity of the isotypic components and the highest weight vectors in the space of metapolynomials, in the context of algebraic complexity lower bounds and border complexity classes. In Section 2.5, based on Theorem 1.1 we prove that algebraic natural proofs can without loss of generality be assumed to be isotypic, see Theorem 2.4. Section 3 sketches the main ideas behind Theorem 1.1. In Section 4, we provide an introduction to Lie algebras, universal enveloping algebras, and their representation theory, on which the core of the proof is built. In Section 5.1, we explain the construction of the projection maps used in the proof of the main theorem; we describe the efficient circuit implementation in Section 5.2. In Section A, we prove the results of Section 2.5. In Section B, we show one part of the proof of the Poincaré–Birkhoff–Witt Theorem (Theorem 4.9); this proof is explicit and can be used to construct the projectors in Theorem 1.1.

## 2 Context and consequences of the main theorem

### 2.1 Complexity classes defined by invariant complexity measures

A p-family is a sequence of polynomials $(f_n)_{n\in\mathbb{N}}$ for which the number of variables of $f_n$ and the degree of $f_n$ are polynomially bounded functions of $n$. The class $\mathsf{VP}$ is the set of p-families whose algebraic circuit complexity is polynomially bounded. Valiant's landmark conjecture is that the permanent sequence $(\mathsf{per}_n)_{n\in\mathbb{N}}$ with $\mathsf{per}_n = \sum_{\sigma\in\mathfrak{S}_n}\prod_{i=1}^{n} x_{i,\sigma(i)}$ is not contained in $\mathsf{VP}$, i.e., $\mathsf{VP}\neq\mathsf{VNP}$; for details see [14, 57, 13]. The extended conjecture states that $(\mathsf{per}_n)\notin\mathsf{VQP}$ [12], i.e., that the algebraic circuit complexity of the permanent is not quasipolynomially bounded.

Let $\mathrm{cc}(f)$ denote the algebraic circuit complexity of the polynomial $f\in\mathbb{C}[x_1,\dots,x_k]$. It is easy to see that $\mathrm{cc}(f)=\mathrm{cc}(g\cdot f)$ for any $g\in\mathrm{GL}_k$: A circuit for $g\cdot f$ is obtained from the one for $f$ by applying $g$ at the input gates. Recall that in the definition of algebraic circuits we allow affine linear forms at the input gates. We say that the complexity measure cc is an *invariant* complexity measure. In more restrictive circuit definitions, for instance if the input gates are required to be variables, the associated circuit complexity $c'$ might change under the action of $\mathrm{GL}_k$. However, both definitions give the same class $\mathsf{VP}$.

Several other classical classes in algebraic complexity theory can be defined as the set of p-families for which some invariant complexity measure is polynomially bounded. For example, the class $\mathsf{VF}$ is the set of p-families whose algebraic formula size is polynomially bounded; this is the smallest size of an algebraic circuit whose underlying directed graph is a tree. The class $\mathsf{VBP}$ is the set of p-families whose algebraic branching program width is polynomially bounded; we do not define this notion here, but it is an easy consequence of the definition that it is an invariant complexity measure, e.g. [68, Def. 2.2]. More classical classes can be readily defined in this way, for example the class of p-families of polynomially bounded circuit size and constant depth (still allowing arbitrary linear forms as inputs), such as $\Sigma\Pi\Sigma$, or the class of p-families of sums of polynomially many powering gates $\Sigma\Lambda\Sigma$. Finally, we mention $\mathsf{VNP}$, which is the class of p-families of polynomially bounded permanental complexity: The smallest $r$ such that $f$ can be written as the permanent of an $r\times r$ matrix with affine linear entries. In all these cases, the relevant complexity measure is invariant.

Notable complexity measures that are *not* invariant are those related to the sparsity of the polynomial, in the sense of [49], or to restrictions on how often a variable can occur, such as read-once or read-$k$ models [27]. A non-invariant complexity measure $c'$ can be converted

into an invariant complexity measure $c$ by defining $c(f) := \min\{c'(gf) \mid g \in \mathrm{GL}_k\}$, i.e., by taking the minimum complexity over all base changes, but for some non-invariant measures this can change the corresponding complexity class. In this paper, we only study complexity classes defined by invariant measures.

## 2.2 Lower bounds via isotypic metapolynomials

An important problem in algebraic complexity theory concerns proving concrete lower bounds for interesting complexity measures $c$. More precisely, given a polynomial $f_{\mathrm{hard}}$ and a certain integer $r$, one aims to show $c(f_{\mathrm{hard}}) > r$ or equivalently $f_{\mathrm{hard}} \notin X_r$ where $X_r = \{f \in \mathbb{C}[x_1, \ldots, x_k]_d \mid c(f) \leq r\}$ is the set of polynomials satisfying the upper bound $r$ on the complexity measure $c(\cdot)$. A standard approach is to determine a function $\Delta$ with the property that $\Delta(f) = 0$ for every $f \in X_r$ and $\Delta(f_{\mathrm{hard}}) \neq 0$.

Several lower bound methods in complexity theory are based on this approach, and most often the function $\Delta$ is a metapolynomial. This is the case for *rank methods* such as the method of partial derivatives [70, 22], the method of shifted partial derivatives [41, 28], and other augmented flattening methods such as Koszul–Young flattenings [54]: in all these cases the complexity lower bound is obtained in terms of a lower bound of the rank of a matrix whose entries depend on the metavariables, or equivalently on the nonvanishing of a suitable set of minors, which are metapolynomials. Rank methods yield the recent breakthrough of [55] giving the first superpolynomial lower bounds for explicit polynomials to be computed by low-product-depth algebraic circuits in large characteristic, further extended to *any* field in [30]. Similarly, the current best known lower bounds for the border determinantal complexity of the permanent, as well as the one of the sum of powers, are based on a non-degeneracy condition which can be translated into the vanishing of a suitable metapolynomial [53, 52]. Finally, the algebraic branching program lower bounds discussed in [50, 36] are built on geometric conditions of the zero set of the polynomial of interest, which in turn can be translated into the vanishing of metapolynomials.

When the complexity measure $c(\cdot)$ is invariant, one can use representation theory to enhance the search for the metapolynomials yielding a separation. Our main result Theorem 1.1 shows that this enhancement is not expensive: the computational overhead is quasi-polynomial.

To illustrate the effect of our result, first consider a simpler reduction. Let $\Delta$ be a metapolynomial vanishing on $X_r$ such that $\Delta(f_{\mathrm{hard}}) \neq 0$. Write $\Delta = \sum_i \Delta^{(i)}$ for homogeneous metapolynomials $\Delta^{(i)}$ of degree $i$. If $X_r$ is invariant under rescaling, one has $\Delta^{(i)}(X_r) = 0$ for every $i$. On the other hand, there exists at least one $i$ such that $\Delta^{(i)}(f_{\mathrm{hard}}) \neq 0$. The relevant homogeneous component $\Delta^{(i)}$ can be extracted from $\Delta$ with a simple interpolation argument, and one has $\mathrm{cc}(\Delta^{(i)}) \leq (\deg(\Delta) + 1) \cdot \mathrm{cc}(\Delta)$. In summary, if the lower bound $c(f_{\mathrm{hard}}) > r$ can be proved via a metapolynomial $\Delta$ satisfying $\mathrm{cc}(\Delta) \leq s$, then the same lower bound can be proved via a *homogeneous* metapolynomial $\Delta^{(i)}$ satisfying $\mathrm{cc}(\Delta^{(i)}) \leq (\deg(\Delta) + 1)s$; and even further, the homogeneous metapolynomial $\Delta^{(i)}$ can be realized as the projection of $\Delta$ onto the $i$-th homogeneous component of the space of metapolynomials.

Theorem 1.1 shows similar results for other projections and in particular for the projection onto a specific isotypic component. Indeed, suppose $\Delta$ is a (homogeneous) metapolynomial vanishing on $X_r$ and such that $\Delta(f_{\mathrm{hard}}) \neq 0$. Write $\Delta = \sum_\lambda \Delta^{(\lambda)}$ as the sum of its isotypic components. If $c(\cdot)$ is an invariant complexity measure, then $X_r$ is invariant under the action of $\mathrm{GL}_k$. In this case, one has $\Delta^{(\lambda)}(X_r) = 0$ for every $\lambda$. On the other hand, there exists at least one $\lambda$ such that $\Delta^{(\lambda)}(f_{\mathrm{hard}}) \neq 0$. Theorem 1.1 provides an upper bound on the circuit complexity of $\Delta^{(\lambda)}$ in terms of the circuit complexity of $\Delta$, showing that if the lower

bound $c(f_{\text{hard}}) > r$ can be proved via a metapolynomial $\Delta$ satisfying $\text{cc}(\Delta) \leq s$, then the same lower bound can be proved via an *isotypic* metapolynomial with circuit complexity controlled by the parameter $s$. Theorem 1.1 further proves analogous results for more refined projections, such as the $T$-isotypic projections or the projection on the highest weight space. These are discussed in Section 2.3.

## 2.3 Highest weight vectors and representation theoretic obstructions

Theorem 1.1 can be used to convert lower bounds proved by $\Delta$ into lower bounds proved by a highest weight vector (HWV for short). To see this, note that if $\Delta$ vanishes on $X_r$ and $\Delta(f_{\text{hard}}) \neq 0$, then there exists a highest weight vector $\Delta'$ vanishing on $X_r$ and such that $\Delta'(g \cdot f_{\text{hard}}) \neq 0$ for some $g \in \text{GL}_k$; in fact, this can be achieved with a random $g \in \text{GL}_k$. Hence a HWV proves the lower bound $c(g \cdot f_{\text{hard}}) > r$; if $c(\cdot)$ is an invariant complexity measure, the same lower bound holds for $f_{\text{hard}}$. Similarly to before, Theorem 1.1 provides an upper bound on the circuit complexity $\text{cc}(\Delta')$ in terms of $\text{cc}(\Delta)$, showing only a quasipolynomial blowup in the complexity of the metapolynomial.

A coarser view on HWVs is given via representation theoretic multiplicities, which is a central idea of geometric complexity theory [60, 61, 59]. Let $Y := \overline{\text{GL}_k f_{\text{hard}}}$ denote the orbit closure, and let $X := \overline{X_s}$. It is easy to see that $f_{\text{hard}} \in X$ if and only if $Y \subseteq X$. The vector space of HWVs of weight $\lambda$ defines a vector space of polynomial functions on $X$; the dimension of such space is called the multiplicity $\mathsf{mult}_\lambda \mathbb{C}[X]$ of $\lambda$ in the coordinate ring $\mathbb{C}[X]$. Standard facts in representation theory, and in particular Schur's Lemma (see Lemma 4.4), guarantee that if $\mathsf{mult}_\lambda \mathbb{C}[X] > \mathsf{mult}_\lambda \mathbb{C}[Y]$ for some $\lambda$, then $X \not\subset Y$. Therefore, multiplicities give a method to potentially separate two varieties, and obtain a lower bound on $c(f_{\text{hard}})$. In this case, we say $\lambda$ is a *multiplicity obstruction*. If additionally $\mathsf{mult}_\lambda \mathbb{C}[X] > 0 = \mathsf{mult}_\lambda \mathbb{C}[Y]$, then $\lambda$ is called an *occurrence obstruction*. Occurrence obstructions [15, 16] and multiplicity obstructions [46] are sometimes sufficient to obtain separations. However, it is known that occurrence obstructions do not work in certain setups [47, 38, 18, 24].

Theorem 1.1 makes no statement about the viability of the method of multiplicity obstructions. Indeed, one main hope of geometric complexity theory is that information about $\mathsf{mult}_\lambda \mathbb{C}[X]$ can be obtained without ever having to write down any circuit for a HWV. Information can be gained by decomposing much easier coordinate rings of orbits, see [19, 17]. This gives upper bounds on $\mathsf{mult}_\lambda \mathbb{C}[X]$ by using classical representation theoretic branching rules. The lower bounds on $\mathsf{mult}_\lambda \mathbb{C}[\overline{\text{GL}_k f_{\text{hard}}}]$ are more difficult to obtain, see [46] for a recent implementation.

## 2.4 Border complexity

Lower bounds achieved via metapolynomials are better described in the setting of *border complexity*. Generally, for a given algebraic complexity measure $c(\cdot)$, one can define a corresponding *border* measure, $\underline{c}(\cdot)$ as follows:

$$\underline{c}(f) = \min \left\{ r : \begin{array}{l} \text{there exists a sequence } f_\varepsilon \text{ such that } f = \lim_{\varepsilon \to 0} f_\varepsilon \\ \text{and } c(f_\varepsilon) = r \text{ for all but finitely many } \varepsilon \in \mathbb{C} \end{array} \right\};$$

correspondingly, a complexity class $\mathcal{C}$ defined in terms of the growth of a given complexity measure naturally induces a *border* complexity class $\bar{\mathcal{C}}$, described by the growth of the corresponding border complexity measure. In fact, $\bar{\mathcal{C}}$ can be defined as the closure of $\mathcal{C}$ with respect to a suitable topology [48]. In the setting of algebraic circuits, the circuit size $\text{cc}(\cdot)$ has a corresponding border circuit measure $\underline{\text{cc}}(\cdot)$ which defines the complexity class $\overline{\mathsf{VP}}$.

A natural question concerns whether $\mathcal{C} = \bar{\mathcal{C}}$ and more generally how much larger $\bar{\mathcal{C}}$ is compared to $\mathcal{C}$. For example, the border class of sequences admitting (border) polynomial size $\Sigma\Lambda\Sigma$ circuits is contained in $\mathsf{VBP}$, see [29] and [9, Theorem 4.2]. In [25], the same containment was shown for the border classes of $\Sigma^{[k]}\Pi\Sigma$ circuits (for every constant $k$). In some restricted setting, for instance in the study of matrix multiplication complexity, it is known that a complexity measure and its associated border complexity are equivalent [7]. In general, the question is wide open for most complexity classes and, in particular, for $\mathsf{VP}$.

Metapolynomials characterize border complexity classes in the following sense: if a sequence of polynomials $(f_n)$ does not belong to a border complexity class $\bar{\mathcal{C}}$ then there is a sequence of metapolynomials *witnessing* the non-membership. This is made precise in Theorem 2.2.

## 2.5 Algebraic complexity and algebraic natural proofs

Razborov and Rudich [66] introduced the notion of natural proofs in Boolean circuit complexity.

▶ **Definition 2.1** (Natural property). *A subset $P \subset \{f : \{0,1\}^n \to \{0,1\}\}$ of Boolean functions is said to be a natural property useful against a class $\mathcal{C}$ of Boolean circuits if the following is true:*

- **(Usefulness).** *Any Boolean function $f : \{0,1\}^n \to \{0,1\}$ that can be computed by a Boolean circuit in $\mathcal{C}$ does not have the property $P$.*
- **(Constructivity).** *Given the truth table of a Boolean function $f : \{0,1\}^n \to \{0,1\}$, whether it has the property $P$ can be decided in time polynomial in the length of the input, i.e., in time $2^{O(n)}$.*
- **(Largeness).** *For all large enough $n$, at least a $2^{-O(n)}$ fraction of all $n$ variate Boolean functions have the property $P$.*

A proof that a certain family of Boolean functions cannot be computed by circuits in $\mathcal{C}$ is said to be a *natural lower bound proof* if the proof (perhaps implicitly) proceeds via establishing a natural property useful against $\mathcal{C}$, and showing that the candidate hard function has this property. Razborov and Rudich showed that most of the known Boolean circuit lower bound proofs, such as lower bounds for $\mathsf{AC}^0$ (constant-depth) circuits have the natural property. One can set the same framework in the algebraic setting; see [39, 40, 31, 21, 51]. The exact details of the definitions are not fully settled yet, see page 19 of [21].

Following [40], the *stretched* classes $\mathsf{metaVP}$ and $\mathsf{metaVQP}$ arise as follows. The class $\mathsf{metaVP}$ is the class of sequences $(\Delta_1, \Delta_2, \ldots)$ of metapolynomials of format $(\delta(n), n, k(n))$ with $k(n)$ polynomially bounded in $n$, and $\delta(n)$ and $\mathrm{cc}(\Delta_n)$ polynomially bounded in $N(n) = \binom{k(n)+n-1}{n}$, which is the number of degree $n$ monomials in $k(n)$ variables. The class $\mathsf{metaVQP}$ is defined analogously, but with $\mathrm{cc}(\Delta_n)$ quasipolynomially bounded in $N(n)$.

We will now argue that $k(n)$ is bounded by a polynomial in $\log(N(n))$. To see this, choose $\gamma$ such that $k(n) \leq n^\gamma$ for $n$ large enough. Take any such large enough $n$ and make a case distinction. If $k(n) \leq n$, then $N = \binom{k+n-1}{k-1} \geq (1 + \frac{n}{k-1})^{k-1}$ implies $\log N \geq (k-1)\log(1 + \frac{n}{k-1}) \geq (k-1)$, since $\frac{n}{k-1} \geq 1$; here we use the fact that $\binom{m}{r} \geq (m/r)^r$. On the other hand, if $k(n) \geq n+1$, then since $\binom{m}{r}$ is an increasing function on $m$ for a fixed $r$, we get that $N = \binom{k+n-1}{n} \geq \binom{2n}{n} \geq 2^n$, for $n \geq 1$. Therefore, $(\log N)^\gamma \geq n^\gamma \geq k$. Hence, the claim is proved in both cases.

Fix an invariant complexity measure $c$ and let $X_{d,r}$ denote the set of homogeneous degree $d$ polynomials $f$ with $c(f) \leq r$. For a format $(\delta, d, k)$ metapolynomial $\Delta : \mathbb{C}[x_1, \ldots, x_k]_d \to \mathbb{C}$ let $\Delta|_{X_{d,r}} : X_{d,r} \to \mathbb{C}$ denote the restriction of $\Delta$ to $X_{d,r}$. Let $\mathcal{C}$ denote the set of p-families with polynomially bounded $c$. The *vanishing ideal* of $\mathcal{C}$ is defined as

$$I(\mathcal{C}) := \left\{ (\Delta_n)_{n\in\mathbb{N}} \text{ of format } (*,n,*) \;:\; \begin{array}{l} \text{for every polynomially bounded } r(n), \\ \text{the sequence } \left(\Delta_1|_{X_{1,r(1)}}, \Delta_2|_{X_{2,r(2)}}, \dots\right) \\ \text{eventually vanishes identically} \end{array} \right\}$$

To emphasize how natural this definition is, we provide a corresponding Nullstellensatz.

▶ **Theorem 2.2** (Nullstellensatz for $\mathcal{C}$). *Let $(f_n)_{n\in\mathbb{N}}$ be a p-family of homogeneous polynomials such that $(\deg f_n)$ is a strictly increasing sequence. The sequence $(f_n)$ lies in $\overline{\mathcal{C}}$ if and only if for every $(\Delta_n) \in I(\mathcal{C})$ we have $\Delta_{\deg(f_n)}(f_n) = 0$ eventually.*

The proof of Theorem 2.2 is given in Section A.

▶ **Definition 2.3** (Algebraic natural proof). *We say that $\mathcal{C}$ has* algebraic natural proofs *if $I(\mathcal{C}) \cap \mathsf{metaVQP}$ contains a sequence that is not eventually zero.*

Here we loosened the definition from the literature slightly from metaVP (see [40, 21]) to metaVQP, but note that the literature does not have just one standard definition, see page 19 of [21]. The landmark question in algebraic metacomplexity is whether VP has algebraic natural proofs.

Definition 2.3 should be compared with Definition 2.1 in the Boolean setting. If $I(\mathsf{VP}) \cap \mathsf{metaVQP}$ contains an eventually nonzero sequence $\Delta$ of format $(\delta(n), n, k(n))$, then clearly there exists the largest $r(n)$ such that $\Delta_n(X_{n,r(n)}) = \{0\}$, and further $\mathrm{cc}(\Delta_n) \leq 2^{\mathrm{poly}(\log N)}$, where $N := \binom{k(n)+n-1}{n}$, and $k(n) \leq \mathrm{poly}(n)$. Therefore, $\Delta$ certifies the non-membership, a *useful natural property*. Furthermore, it trivially satisfies the *largeness* criterion, since $\Delta_n$ does not vanish on most polynomials. Finally, the quasipolynomial bound on the algebraic circuit complexity of $\Delta_n$ gives a natural algebraic analogue of the notion of *constructivity*.

Let Isotypic denote the set of sequences of isotypic metapolynomials. The following theorem follows from Theorem 1.1. It states in a concise way that only isotypic metapolynomials have to be considered for algebraic natural proofs.

▶ **Theorem 2.4** (Main theorem about algebraic natural proofs). *$\mathcal{C}$ has algebraic natural proofs if and only if $I(\mathcal{C}) \cap \mathsf{metaVQP} \cap \mathsf{Isotypic}$ contains a sequence that is not eventually zero.*

**Proof.** One direction is clear, because $I(\mathcal{C}) \cap \mathsf{metaVQP} \cap \mathsf{Isotypic} \subseteq I(\mathcal{C}) \cap \mathsf{metaVQP}$.

For the other direction, we first consider a general property of isotypic components of vanishing ideals. Let $I(X)_\delta = \{\Delta \text{ of format } (\delta, d, k) \mid \Delta|_X = 0\}$. Let $X$ be closed under the action of $\mathrm{GL}_k$. If $\Delta \in I(X)_\delta$, then $g\Delta \in I(X)_\delta$ for every $g \in \mathrm{GL}_k$. Indeed, in this case for all $x \in X$ we have $(g\Delta)(x) = (\Delta)(g^{-1}x) = 0$, since $g^{-1}x \in X$. Moreover, $I(X)_\delta$ is closed under linear combinations, hence $I(X)_\delta$ is a vector space, and in particular closed under taking limits and thus closed under the action of the Lie algebra. We have $\Delta \in I(X)_\delta$ if and only if for all $\lambda$ we have $\Delta^{(\lambda)} \in I(X)_\delta$, because the projection of $\Delta$ to an isotypic component $\Delta^{(\lambda)}$ involves only finite linear combinations of Lie algebra actions, see Section 5.

Let $\mathcal{C}$ have algebraic natural proofs, witnessed by a sequence $(\Delta_n)_{n\in\mathbb{N}} \in I(\mathcal{C}) \cap \mathsf{metaVQP}$ of format $(\delta(n), n, k(n))$ that is not eventually zero. Now, let $(\lambda^n)_{n\in\mathbb{N}}$ be a sequence of partitions, $\lambda^n \vdash_{k(n)} n \cdot \delta(n)$, such that the isotypic component $\Delta_n^{(\lambda^n)}$ of $\Delta_n$ is nonzero whenever $\Delta_n$ is nonzero. By construction, $(\Delta_n^{(\lambda_n)})_{n\in\mathbb{N}} \in \mathsf{Isotypic}$ and $(\Delta_n^{(\lambda_n)})_{n\in\mathbb{N}}$ is not eventually zero. Since each $X_{n,i}$ is invariant under the group action, we have for all $i$ and $n$: Whenever $\Delta_n|_{X_{n,i}} = 0$, then also $\Delta_n^{(\lambda^n)}|_{X_{n,i}} = 0$. For every polynomially bounded $r$ we have $\Delta_n|_{X_{n,r(n)}} = 0$ eventually, hence we have $\Delta_n^{(\lambda^n)}|_{X_{n,r(n)}} = 0$ eventually, and hence $(\Delta_n^{(\lambda_n)})_{n\in\mathbb{N}} \in$

$I(\mathcal{C})$. Let $N(n) = \binom{k(n)+n-1}{n}$. Let $s(n) = \mathrm{cc}(\Delta_n) \in 2^{\mathrm{poly}(\log(N))}$. By Theorem 1.1 with $k(n) \in \mathrm{poly}(\log(N))$, $\delta(n) \in \mathrm{poly}(N)$, and $d = n$, we have $\mathrm{cc}(\Delta_n^{(\lambda^n)}) \in 2^{\mathrm{poly}(\log(N))}$, which implies $(\Delta_n^{(\lambda_n)})_{n\in\mathbb{N}} \in \mathsf{metaVQP}$, which finishes the proof. ◀

## 3 Outline of the proof of Theorem 1.1

The proof of Theorem 1.1 is built on an efficient construction of the image of metapolynomials under projections onto the isotypic spaces.

Roughly speaking, we will construct the projectors of Theorem 1.1 as linear combinations of repeated derivations applied to metapolynomials. These derivations live naturally in a vector space called the *Lie algebra* $\mathfrak{gl}_k$ of $\mathrm{GL}_k$. They are defined as the linearization of linear change of coordinates by elements of $\mathrm{GL}_k$ applied to polynomials on which a metapolynomial is evaluated. Repeated applications of these operations is formally described using products in an associative algebra $\mathcal{U}(\mathfrak{gl}_k) \supseteq \mathfrak{gl}_k$ called the *universal enveloping algebra* of $\mathfrak{gl}_k$. More concretely, a standard basis element in $\mathfrak{gl}_k$ corresponds to an analogue of a shifted partial derivative for metapolynomials (see Claim 4.2), and an element of $\mathcal{U}(\mathfrak{gl}_k)$ captures repeated applications of such shifted partial derivatives.

Lie algebras and their universal enveloping algebras have a long history in mathematics, and their theory is rich and well-developed. We review the theory in Section 4 to the extent required to construct the projectors. The construction is completely explicit. The universal enveloping algebra $\mathcal{U}(\mathfrak{gl}_k)$ contains elements $\{C_1, \ldots, C_r\}$ which *separate* isotypic components in the following sense: each $C_i$ acts by scalar multiplication on each isotypic component, and for any two distinct isotypic components there is a $C_i$ for which the corresponding scalars are different; the precise statement is the content of Theorem 4.11 and Theorem 4.13. This property is used to construct the projectors via an analogue of standard multivariate polynomial interpolation, see Section 5.1. The elements $\{C_1, \ldots, C_r\}$ are called *Casimir operators*, and they are defined as generators of certain commutative subalgebras of $\mathcal{U}(\mathfrak{gl}_k)$. For each case of Theorem 1.1, there is a suitable algebra for which the corresponding isotypic components are the spaces of interest, see Table 1.

The proof that the projectors can be implemented efficiently is built on the following insights. The Poincaré–Birkhoff–Witt theorem (Theorem 4.9) allows one to express each projector as a linear combination in $\mathcal{U}(\mathfrak{gl}_k)$ of $m$-fold products of elements of $\mathfrak{gl}_k$ with $m$ bounded polynomially in the number of isotypic components. This number is exponential in $k$, which is polylogarithmic in $N$ as observed in Section 2.5. We use this fact to construct circuits of quasipolynomial size for the projectors in Section 5.2. In turn, this will complete the proof of Theorem 1.1.

## 4 Representations of Lie algebras

In this section, we provide an introduction to Lie algebras, universal enveloping algebras, and their representation theory, following mainly [32, 42]. Afterwards, we introduce and explain the results needed to construct the projectors in Theorem 1.1. These results are typically stated in the language of abstract Lie algebras, and therefore we take time to build up this general theory. In our case of interest, the Lie algebra arises from $\mathrm{GL}_k$ and many parts simplify. We will regularly indicate this in the text for clarity and concreteness. We establish homogeneous polynomials and metamonomials as a Lie algebra representation, which we will use as running examples.

A Lie algebra $\mathfrak{g}$ is a vector space endowed with a bilinear *bracket* operation $[-,-]$ which is skew-commutative, i.e., $[X,Y] = -[Y,X]$, and satisfies a relation called the Jacobi identity $[X,[Y,Z]] + [Z,[X,Y]] + [Y,[Z,X]] = 0$; in particular, the bracket operation is usually not associative. Every associative algebra $\mathcal{A}$ is a Lie algebra with the commutator bracket given by $[a,b] \coloneqq ab - ba$ for every $a,b \in \mathcal{A}$. In this paper, we are interested in the representation theory of the general linear group $\mathrm{GL}_k$, and we study it via the representation theory of its Lie algebra $\mathfrak{gl}_k$: this is the space of $k \times k$ matrices endowed with the commutator bracket.

A Lie algebra representation is a vector space $V$ with an associated linear map $\rho : \mathfrak{g} \to \mathrm{End}(V)$ such that $\rho([X,Y]) = [\rho(X), \rho(Y)]$ where the bracket on the right is the commutator bracket. We equivalently say that $\mathfrak{g}$ (or its elements) acts on $V$. For $v \in V$, we often write $X.v \coloneqq \rho(X)v$ as a shorthand. A simple example is $V = \mathbb{C}^k$, where $\mathfrak{gl}_k$ acts by matrix-vector multiplication. This is called the standard representation of $\mathfrak{gl}_k$. In this paper we examine the space of metapolynomials $\mathbb{C}[\mathbb{C}[x_1,\dots,x_k]_d]_\delta$ as a Lie algebra representation of $\mathfrak{gl}_k$. The Lie algebra action is explicitly given in in Claim 4.2.

Representation theory of Lie algebras is used to study the representation theory of Lie groups because it allows one to "linearize" a group action. More precisely, every Lie group $G$ has an associated Lie algebra $\mathfrak{g}$ and every Lie group representation (a non-linear map) defines naturally a Lie algebra representation (a linear map). All Lie algebra representations we will encounter arise from this correspondence. We will describe the $\mathfrak{gl}_k$ action on polynomials and metapolynomials explicitly in Claim 4.1 and Claim 4.2, which may be taken as the definition. The details of the correspondence between Lie groups and Lie algebras are not crucial for the rest of the paper. We briefly outline them in Remark 4.3.

For a vector space $V$ with basis $v_1,\dots,v_r$, let $v_1^*,\dots,v_r^* \in V^*$ be the dual basis of $V^*$. Explicitly, $v_\ell^*$ is the linear form mapping a vector $v = \sum_i c_i v_i \in V$, with $c_i \in \mathbb{C}$, to the coefficient $c_\ell$. The basis dual to the standard basis of $\mathbb{C}^k$ is the set of variables $x_1,\dots,x_k$. Let $E_{i,j} \in \mathfrak{gl}_k$ be the $k \times k$ matrix having 1 at the entry $(i,j)$ and zero elsewhere. Regarding $E_{i,j} : \mathbb{C}^k \to \mathbb{C}^k$ as a linear map, we have $x_\ell \circ E_{i,j} = 0$ if $\ell \neq i$ and $x_\ell \circ E_{i,j} = x_j$ if $\ell = i$; here $\circ$ denotes the composition of functions. Since $\{E_{i,j}\}_{i,j}$ forms a basis for $\mathfrak{gl}_k$, the action of $\mathfrak{gl}_k$ on a representation $V$ is uniquely determined by the action of the $E_{i,j}$'s.

▷ Claim 4.1 (Action on polynomials). The vector space $\mathbb{C}[x_1,\dots,x_k]_d$ of homogeneous degree $d$ polynomials is a representation of $\mathfrak{gl}_k$. It is the linearization of the action of $\mathrm{GL}_k$ acting by base change on the variables $x_1,\dots,x_k$. For $f \in \mathbb{C}[x_1,\dots,x_k]_d$, the action is defined by the *shifted partial derivative*

$$E_{i,j}.f = -\sum_{\ell=1}^{k} (x_\ell \circ E_{i,j}) \frac{\partial}{\partial x_\ell} f = -x_j \frac{\partial}{\partial x_i} f. \tag{1}$$

In particular, for a monomial $f = x_{\ell_1} \cdots x_{\ell_d}$, we have

$$E_{i,j}.f = \textstyle\sum_{p=1}^{d} x_{\ell_1} \cdots x_{\ell_{p-1}} (E_{i,j}.x_{\ell_p}) x_{\ell_{p+1}} \cdots x_{\ell_d}.$$

For a multiindex $\mu \in \mathbb{N}^k, \sum_i \mu_i = d$, let $c_\mu$ be the corresponding variable, that is, the linear map sending a polynomial $f$ to the coefficient of $x^\mu$. Let $e_1,\dots,e_k \in \mathbb{N}^k$ be the standard basis vectors: $e_i$ is the vector with 1 at the $i$-th entry and 0 elsewhere.

▷ Claim 4.2 (Action on metapolynomials). The vector space of metapolynomials of format $(\delta,d,k)$, that is $\mathbb{C}[\mathbb{C}[x_1,\dots,x_k]_d]_\delta = \mathbb{C}[c_\mu : \mu \in \mathbb{N}^k, |\mu| = d]_\delta$ is a representation of $\mathfrak{gl}_k$. It is the linearization of the action of $\mathrm{GL}_k$ acting by base change on the variables $x_1,\dots,x_k$. For a metapolynomial $\Delta \in \mathbb{C}[c_\mu : \mu \in \mathbb{N}^k, |\mu| = d]_\delta$, the action is given by

$$E_{i,j}.\Delta = \sum_{\mu} (E_{i,j}.c_\mu) \frac{\partial}{\partial c_\mu} \Delta, \tag{2}$$

where the summation is over all multi-indices $\mu \in \mathbb{N}^k$ with $\sum_i \mu_i = d$. The action on metavariables is given by

$$E_{i,j}.c_\mu = \begin{cases} (\mu_i + 1)\, c_{\mu + e_i - e_j} & \text{if } i \neq j, \\ \mu_i\, c_\mu & \text{otherwise,} \end{cases} \tag{3}$$

where we set $c_{\mu+e_i-e_j} = 0$ if $\mu_j = 0$. In particular, for a metamonomial $\Delta = c_{\mu_1} \cdots c_{\mu_\delta}$, we have

$$E_{i,j}.\Delta = \textstyle\sum_{p=1}^{\delta} c_{\mu_1} \cdots c_{\mu_{p-1}} (E_{i,j}.c_{\mu_p}) c_{\mu_{p+1}} \cdots c_{\mu_\delta}.$$

The following remark illustrates the correspondence between representations of complex algebraic Lie groups and representations of the corresponding Lie algebras. As mentioned before, it is not essential for the rest of the paper, but the proof of Claim 4.1 and Claim 4.2 are built on this correspondence.

▶ Remark 4.3. A complex algebraic Lie group is a group with the structure of a complex algebraic variety, [32, Ch. 7]. One can show that $\mathrm{GL}_k$ is such a group. Let $V$ be a finite-dimensional vector space. A representation of a complex algebraic Lie group $G$ on $V$ is a group homomorphism $\rho : G \to \mathrm{GL}(V)$ which is a regular map, in the sense that the entries of $\rho(g) \in \mathrm{GL}(V)$ regarded as a matrix in a fixed basis, are polynomial functions on $G$.

As a vector space, the Lie algebra of $G$ is, by definition, the tangent space of $G$ at its identity element $\mathfrak{g} = T_{\mathsf{id}} G$. The bracket operation is defined via the *adjoint representation* of $G$ on $\mathfrak{g}$, see [32, Ch. 8] In the case of $\mathrm{GL}_k$ one has $\mathfrak{gl}_k \simeq T_{\mathsf{id}_k} \mathrm{GL}_k$ is the space of $k \times k$ matrices, and the bracket is the standard commutator.

Every representation $\rho : G \to \mathrm{GL}(V)$ of $G$ defines a representation $\mathfrak{g}$ via its differential at the identity: $\mathrm{d}\rho : T_{\mathsf{id}} G \to T_{\mathsf{id}} \mathrm{GL}(V)$; this is a linear map and defines a Lie algebra homomorphism from $T_{\mathsf{id}} G = \mathfrak{g}$ to $T_{\mathsf{id}} \mathrm{GL}(V) = \mathrm{End}(V) = \mathfrak{gl}(V)$.

Explicitly, the value of $\mathrm{d}\rho$ at $X \in \mathfrak{g}$ is the element $\mathrm{d}\rho(X) \in \mathfrak{gl}(V) \simeq \mathrm{End}(V)$ defined by

$$\big(\mathrm{d}\rho(X)\big)v = \frac{\mathrm{d}}{\mathrm{d}s}\Big|_{s=0} \Big(\rho\big(\gamma(s)\big)v\Big) \quad \text{for every } v \in V; \tag{4}$$

here $\gamma : \mathbb{C} \to G$ is any smooth curve such that $\gamma(0) = \mathsf{id}_G$ and $\gamma'(0) = X$. It is easy to see that $\mathrm{d}\rho(X)$ does not depend on the choice of $\gamma$.

**Proof of Claim 4.1 and Claim 4.2.** We derive the Lie algebra representations from their respective group representations, as in Remark 4.3.

Every representation $\rho : G \to \mathrm{GL}(V)$ defines a *pullback* representation on the space of polynomials of degree $p$ on $V$, $\mathbb{C}[V]_p$. For $g \in G$ and $f \in \mathbb{C}[V]_p$, this is given by $g \cdot f = f \circ \rho(g^{-1})$ where $\rho(g^{-1}) : V \to V$ is regarded as a linear map and $\circ$ denotes the composition of functions.

The standard representation of $\mathrm{GL}_k$ on $\mathbb{C}^k$ induces in this way a representation on the space of polynomials $\mathbb{C}[x_1, \ldots, x_k]_d$. Similarly, the action of $\mathrm{GL}(\mathbb{C}[x_1, \ldots, x_k]_d)$ induces a representation on the space of metapolynomials $\mathbb{C}[\mathbb{C}[x_1, \ldots, x_k]_d]_\delta$; the action of $\mathrm{GL}_k$ on metapolynomials is the composition of the two representations

$$\mathrm{GL}_k \to \mathrm{GL}(\mathbb{C}[x_1, \ldots, x_k]_d) \to \mathrm{GL}(\mathbb{C}[\mathbb{C}[x_1, \ldots, x_k]_d]_\delta).$$

Now, if $\rho : G \to \mathrm{GL}(V)$ induces the Lie algebra representation $\mathrm{d}\rho : \mathfrak{g} \to \mathrm{GL}(V)$, the pullback representation of $G$ on $\mathbb{C}[V]_p$ induces the Lie algebra representation described as follows. For $X \in \mathfrak{g}$, let $\gamma(s)$ be a smooth curve in $G$ such that $\gamma(0) = \mathsf{id}_G$ and $\gamma'(0) = X$. The chain rule of derivatives implies $\frac{\mathrm{d}}{\mathrm{d}s}\big|_{s=0}\gamma(s)^{-1} = -X$. We show this in the case where $G$ is a group of matrices: in this case, one can *multiply* elements of $G$ by elements of $\mathfrak{g}$ and obtain

$$\frac{\mathrm{d}}{\mathrm{d}s}\Big|_{s=0}\gamma(s)^{-1} = -\big[\frac{\mathrm{d}}{\mathrm{d}s}\Big|_{s=0}\gamma(s)\big]\cdot\big[(\gamma(s)^{-2})\big|_{s=0}\big] = -\gamma'(0)\cdot\gamma(0)^{-2} = -X\cdot\mathsf{id}^{-2} = -X.$$

In general, the multiplication of matrices should be replaced by the action of $G$ on $\mathfrak{g}$.

Let $V^*$ be the dual space of $V$ with basis $z_1, \ldots, z_r$; in particular $\mathbb{C}[V]_p$ is the space of polynomials of degree $p$ in $z_1, \ldots, z_r$. By Equation (4), for every $F \in \mathbb{C}[V]_p$, $X.F$ is the polynomial function described as follows, on every $\zeta \in V$:

$$\begin{aligned} X.F(\zeta) =& \frac{\mathrm{d}}{\mathrm{d}s}\Big|_0(\gamma(s)\cdot F)(\zeta) = \\ & \frac{\mathrm{d}}{\mathrm{d}s}\Big|_0(F\circ\gamma(s)^{-1})(\zeta) = \qquad \text{(using chain rule)} \\ & \sum_{\ell=1}^{r}\frac{\partial}{\partial z_\ell}F(\zeta)\cdot(z_\ell(\frac{\mathrm{d}}{\mathrm{d}s}\Big|_0\gamma(s)^{-1}\zeta)) = \\ & \sum_{\ell=1}^{r}\frac{\partial}{\partial z_\ell}F(\zeta)\cdot(z_\ell(-X.\zeta)) = -\sum_{\ell=1}^{r}\frac{\partial}{\partial z_\ell}F(\zeta)\cdot z_\ell(X.\zeta). \end{aligned}$$

Specialize the calculation above to the case of polynomials and metapolynomials to obtain Equation (1) and Equation (2). It remains to show Equation (3), which follows from a direct calculation: regarding $c_\mu$ as a function of a polynomial $f$, we have

$$E_{ij}.c_\mu(f) = -c_\mu(E_{ij}.f) = c_\mu(x_j\frac{\partial}{\partial x_i}f) = \begin{cases}(\mu_i+1)\,c_{\mu+e_i-e_j}(f) & \text{if } i\neq j,\\ \mu_i\,c_\mu(f) & \text{otherwise.}\end{cases}$$ ◀

Let $V$ be a representation of a Lie algebra $\mathfrak{g}$. A subspace $V' \subseteq V$ that is itself a $\mathfrak{g}$-representation is called a subrepresentation. A representation $V$ is irreducible if it contains no subrepresentations except the zero space of $V$ and $V$ itself. A map $\phi : V \to W$ between two representations is called equivariant if it commutes with the action of $\mathfrak{g}$. Schur's Lemma characterizes equivariant maps between irreducible representations:

▶ **Lemma 4.4** (Schur's Lemma, [32, Lem. 1.7]). *Let $V, W$ be irreducible representations for a Lie algebra $\mathfrak{g}$ and let $\phi : V \to W$ be an equivariant map. Then either $\phi = 0$ or $\phi$ is an isomorphism. Moreover, if $V = W$, then $\phi = c\cdot\mathsf{id}_V$ for some $c \in \mathbb{C}$.*

A complex algebraic Lie group is called *reductive* if every representation is completely reducible, namely it splits into direct sum of irreducible subrepresentations; the general linear group $\mathrm{GL}_k$ is reductive, see e.g. [45, Ch.X]. In the following $\mathfrak{g}$ is the Lie algebra of a reductive group, and explicitly we are interested in the case where $G = \mathrm{GL}_k$ and $\mathfrak{g} = \mathfrak{gl}_k$.

## 4.1 Cartan subalgebras and weights

A subalgebra $\mathfrak{t} \subseteq \mathfrak{g}$ is called abelian if $[\mathfrak{t}, \mathfrak{t}] = 0$. A Cartan subalgebra $\mathfrak{h}$ of $\mathfrak{g}$ is a maximal abelian subalgebra with the property that if $[X, \mathfrak{h}] \subseteq \mathfrak{h}$ then $X \in \mathfrak{h}$. The subalgebra of diagonal matrices in $\mathfrak{gl}_k$ is a Cartan subalgebra of $\mathfrak{gl}_k$. It turns out that if $G$ is reductive, then all Cartan subalgebras of $\mathfrak{g}$ have the same dimension; in fact they are all equivalent

under a natural action of $G$ on $\mathfrak{g}$ [32, Sec. D3]. In the case of $\mathrm{GL}_k$, the action on $\mathfrak{gl}_k$ is given by conjugation: $g \cdot X = gXg^{-1}$ for $g \in \mathrm{GL}_k$ and $X \in \mathfrak{gl}_k$. In this case, one can show that any Cartan subalgebra is equivalent to the subspace of diagonal matrices.

If $\mathfrak{g}$ is the Lie algebra of a complex reductive Lie group $G$, then Cartan subalgebras of $\mathfrak{g}$ correspond to maximal abelian subgroups in $G$. For a fixed Cartan subalgebra $\mathfrak{h}$, let $\mathbb{T}$ be the corresponding abelian subgroup of $G$. If $\mathfrak{h}$ is the subset of diagonal matrices in $\mathfrak{gl}_k$, then $\mathfrak{h}$ is the Lie algebra of the abelian group $\mathbb{T}$ of diagonal matrices with nonzero diagonal entries. Irreducible $\mathbb{T}$-representations are 1-dimensional and they are in one-to-one correspondence with a lattice $\Lambda$ in the space $\mathfrak{h}^*$, called the *weight lattice* of $\mathfrak{g}$; the elements of $\Lambda$ are called *weights*. When $\mathbb{T}$ is the subgroup of $\mathrm{GL}_k$ of diagonal matrices, then this definition coincides with the one used in Section 1 for metapolynomials; see also Example 4.7. The construction of the weight lattice is not straightforward for arbitrary Lie algebras, but it is very explicit in the case where $\mathfrak{h}$ is the subalgebra of diagonal elements in $\mathfrak{gl}_k$:

▷ Claim 4.5. The weight lattice of $\mathfrak{gl}_k$ with respect to the Cartan subalgebra $\mathfrak{h}$ of diagonal elements is the abelian group of integer linear combinations of the diagonal elements, that is, the group of linear functions

$$\begin{aligned} \alpha : \mathfrak{h} &\to \mathbb{C} \\ \begin{pmatrix} h_1 & & \\ & \ddots & \\ & & h_k \end{pmatrix} &\mapsto \alpha_1 h_1 + \cdots + \alpha_k h_k. \end{aligned}$$

with $\alpha_1, \ldots, \alpha_k \in \mathbb{Z}$. In the following, we usually identify the weight $\alpha$ with the $k$-tuple of its integer coefficients $\alpha = (\alpha_1, \ldots, \alpha_k)$.

Proof. Let $\mathbb{T}$ be the subgroup of diagonal matrices in $\mathrm{GL}_k$. Let $(\alpha_1, \ldots, \alpha_k) \in \mathbb{Z}^k$ be an integer vector.

$$\begin{aligned} \rho : \mathbb{T} &\to \mathbb{C}^\times \simeq \mathrm{GL}_1 \\ \begin{pmatrix} t_1 & & \\ & \ddots & \\ & & t_k \end{pmatrix} &\mapsto t_1^{\alpha_1} \cdots t_k^{\alpha_k} \end{aligned}$$

is a group homomorphism, hence it defines a 1-dimensional (irreducible) representation of $\mathbb{T}$. It is not hard to prove that all irreducible representations of $\mathbb{T}$ arise in this way.

The induced $\mathfrak{h}$-representation is

$$\begin{aligned} \mathrm{d}\rho : \mathfrak{h} &\to \mathbb{C} \simeq \mathfrak{gl}_1 \\ \begin{pmatrix} h_1 & & \\ & \ddots & \\ & & h_k \end{pmatrix} &\mapsto \alpha_1 h_1 + \cdots + \alpha_k h_k. \end{aligned}$$

In other words, irreducible $\mathfrak{h}$-representations induced from $\mathbb{T}$-representations are in one-to-one correspondence with functions $\alpha \in \mathfrak{h}^*$ mapping a diagonal element $H \in \mathfrak{h}$ to an integer linear combination of its diagonal entries. ◁

Every $\mathfrak{g}$-representation $V$ restricts to an $\mathfrak{h}$-representation, hence it decomposes into the direct sum of irreducible 1-dimensional $\mathfrak{h}$-representations. Since they are uniquely determined by their weights, we have $V = \bigoplus_{\alpha \in \Lambda} V_\alpha^{\oplus m_\alpha}$; in particular $m_\alpha \neq 0$ only for finitely many $\alpha \in \Lambda$. If $m_\alpha \neq 0$, then $\alpha$ is called a weight of $V$, $m_\alpha$ is called *multiplicity* and the direct summand $V_\alpha^{\oplus m_\alpha}$ is called the *weight subspace* of $V$ of weight $\alpha$. An element $X \in \mathfrak{h}$ acts simultaneously on the $m_\alpha$ copies of $V_\alpha$ via scalar multiplication by $\alpha(X) \in \mathbb{C}$. The elements of a weight subspace are called *weight vectors*; a basis of $V$ consisting of weight vectors is called a *weight basis*.

▶ **Example 4.6.** Consider the $\mathfrak{gl}_3$-representation $V = \mathbb{C}[x_1, x_2, x_3]_2$ with the action described in Claim 4.1. Let $\mathfrak{h} \subseteq \mathfrak{gl}_3$ be the subalgebra of diagonal matrices. Consider the six monomials of $V$, which define a basis:

$$\begin{array}{lll} x_1^2, & x_2^2, & x_3^2, \\ x_2x_3, & x_1x_3, & x_1x_2. \end{array}$$

One can verify that they form a weight basis of $V$. For instance, given $H = \left(\begin{smallmatrix} t_1 & & \\ & t_2 & \\ & & t_3 \end{smallmatrix}\right) = t_1E_{11} + t_2E_{22} + t_3E_{33} \in \mathfrak{h}$, we have

$$H.(x_1^2) = \textstyle\sum_i t_iE_{ii}.(x_1^2) = -\sum_i t_i \sum_j x_j \dfrac{\partial}{\partial x_j}(x_1^2) = -(2t_1)x_1^2$$

showing that $x_1^2$ is a weight vector of weight $(-2, 0, 0)$. More generally, we have

$$\begin{array}{lll} H.(x_1^2) = -2t_1\, x_1^2, & H.(x_2^2) = -2t_2\, x_2^2, & H.(x_3^2) = -2t_3\, x_3^2, \\ H.(x_2x_3) = -(t_2+t_3)\, x_2x_3, & H.(x_1x_3) = -(t_1+t_3)\, x_1x_3, & H.(x_1x_2) = -(t_1+t_2)\, x_1x_2. \end{array}$$

This shows that $V$ decomposes into the sum of six weight spaces, each spanned by one of the weight vectors above. The corresponding weights are the six elements of $\mathfrak{h}^*$ mapping $H$ to the eigenvalues of $H$ on each weight vector. As $k$-tuple of integers, they are

$$\begin{array}{lll} (-2,0,0), & (0,-2,0), & (0,0,-2), \\ (0,-1,-1), & (-1,0,-1), & (-1,-1,0). \end{array}$$

More generally, in a space of polynomials $V = \mathbb{C}[x_1, \ldots, x_k]_d$, the monomials form a weight basis. The weight of a monomial $x^\mu = x_1^{\mu_1} \cdots x_k^{\mu_k}$, regarded as a vector of integers, is $-\mu$.

▶ **Example 4.7.** Consider the $\mathfrak{gl}_2$-representation $V = \mathbb{C}[\mathbb{C}[x_1, x_2]_3]_2$ of metapolynomials of format $(2, 3, 2)$, with the action described in Claim 4.2. One can verify that the metamonomials form a weight basis. For instance, given $H = \left(\begin{smallmatrix} t_1 & \\ & t_2 \end{smallmatrix}\right) = t_1E_{11} + t_2E_{22} \in \mathfrak{h}$, we have

$$H.(c_{30}^2) = \textstyle\sum_{i=1}^2 t_iE_{ii}.(c_{30}^2) = \sum_{i=1}^2 t_i \sum_\mu E_{ii}.c_\mu \cdot \dfrac{\partial}{\partial c_\mu}(c_{30}^2) = 3 \cdot 2 \cdot t_1 \cdot c_{30}^2 = (6t_1)c_{30}^2.$$

showing that $c_{30}^2$ is a weight vector of weight $(6, 0)$. Similarly, one can compute the weight of every metamonomial:

$$\begin{array}{lll} H.(c_{21}^2) = (4t_1+2t_2)c_{21}^2, & H.(c_{12}^2) = (2t_1+4t_2)c_{12}^2, & H.(c_{03}^2) = (6t_2)c_{03}^2, \\ H.(c_{30}c_{21}) = (5t_1+\ t_2)c_{30}c_{21}, & H.(c_{30}c_{12}) = (4t_1+2t_2)c_{30}c_{12}, & H.(c_{30}c_{03}) = (3t_1+3t_2)c_{30}c_{03}, \\ H.(c_{21}c_{12}) = (3t_1+3t_2)c_{21}c_{12}, & H.(c_{21}c_{03}) = (2t_1+4t_2)c_{21}c_{03}, & H.(c_{12}c_{03}) = (\ t_1+5t_2)c_{21}c_{12}. \end{array}$$

Therefore $V$ has seven weight spaces of weight $(p, 6-p)$ with $p = 0, \ldots, 6$. The weight spaces of weight $(4, 2), (3, 3), (2, 4)$ have dimension 2, the others have dimension 1.

More generally, the metamonomials give a weight basis for the space of metapolynomials of format $(\delta, d, k)$. A metamonomial $c_{\mu_1} \cdots c_{\mu_\delta}$ has weight $\mu_1 + \cdots + \mu_\delta$, in accordance with what we defined in the introduction.

## 4.2 Roots and highest weights

An important representation is given by the action of the Lie algebra on itself, called the adjoint representation $\mathrm{ad} : \mathfrak{g} \to \mathrm{End}(\mathfrak{g})$ and defined by $\mathrm{ad}(X) = [X, -]$. By definition, the weight space of weight 0 is $\mathfrak{h}$ itself; moreover, one can prove that all other weight spaces

are 1-dimensional [32, Ch. 14]. The nonzero weights of the adjoint representation are called the *roots* of the Lie algebra $\mathfrak{g}$ (with respect to $\mathfrak{h}$); let $\Phi_{\mathfrak{g}} \subseteq \Lambda$ denote the set of roots of $\mathfrak{g}$. It turns out that $\eta \in \mathfrak{h}^*$ is a root if and only if $-\eta$ is a root and $\pm\eta$ are the only integer multiples of $\eta$ that are roots.

We compute the roots for $\mathfrak{gl}_k$ with $\mathfrak{h}$ the algebra of diagonal matrices. Let $H = \mathrm{diag}(t_1, \ldots, t_k) \in \mathfrak{h}$. Then

$$\mathrm{ad}(H)E_{i,j} = HE_{i,j} - E_{i,j}H = t_iE_{i,j} - t_jE_{i,j} = (t_i - t_j)E_{i,j}. \tag{5}$$

So $\mathbb{C}E_{i,j}$ is the weight space corresponding to the weight $(e_i - e_j) \in \mathfrak{h}^*$. Since the matrices $E_{i,j}$ span $\mathfrak{gl}_k$, we obtain that the set of roots is $\Phi_{\mathfrak{gl}_k} = \{e_i - e_j \mid i, j \in [k], i \neq j\}$.

The roots of the Lie algebra *shift* the weights of the Lie algebra representations: more precisely, let $V$ be a representation of $\mathfrak{g}$, $v \in V$ a weight vector of weight $\lambda$ and $X \in \mathfrak{g}_\alpha$ for some root $\alpha$; then $X.v$ is a weight vector of weight $\lambda + \alpha$.

A choice of positive roots is a subset $\Phi^+_{\mathfrak{g}}$ of $\Phi_{\mathfrak{g}}$ with the following two properties: for every root $\alpha$, either $\alpha$ or $-\alpha$ belong to $\Phi^+_{\mathfrak{g}}$; for every pair of roots $\alpha_1, \alpha_2$ such that $\alpha_1 + \alpha_2$ is a root, if $\alpha_1, \alpha_2 \in \Phi^+_{\mathfrak{g}}$ then $\alpha_1 + \alpha_2 \in \Phi^+_{\mathfrak{g}}$. For $\mathfrak{gl}_k$, one such choice is $\Phi^+_{\mathfrak{gl}_k} = \{e_i - e_j \mid i < j\}$.

All choices of positive roots are equivalent under the action of a finite group $\mathcal{W}_{\mathfrak{g}}$, called the Weyl group of $\mathfrak{g}$. This group acts on $\mathfrak{h}$ and hence on $\mathfrak{h}^*$; the action of $\mathfrak{h}^*$ sends the set of roots to itself. All possible sets of positive roots are obtained one from the other via the action of $\mathcal{W}_{\mathfrak{g}}$. The Weyl group of $\mathfrak{gl}_k$ is the permutation group $\mathfrak{S}_k$ on $k$ elements. If $\mathfrak{h}$ is the Cartan subalgebra of diagonal matrices, then $\mathcal{W}_{\mathfrak{g}}$ acts by permuting the diagonal elements; the induced action on $\mathfrak{h}^*$ permutes the weights $e_i$: for $\sigma \in \mathfrak{S}_k$, $\sigma \cdot e_i$ is the weight $e_{\sigma^{-1}(i)}$.

Since for every root $\alpha$, either $\alpha$ or $-\alpha$ is positive, the weight decomposition of $\mathfrak{g}$ gives $\mathfrak{g} = \mathfrak{h} \oplus \bigoplus_{\alpha \in \Phi^+_{\mathfrak{g}}} (\mathfrak{g}_\alpha \oplus \mathfrak{g}_{-\alpha})$. In the case of $\mathfrak{gl}_k$, this decomposition is $\mathfrak{gl}_k = \mathfrak{h} \oplus \bigoplus_{i<j} (\mathbb{C}E_{ij} \oplus \mathbb{C}E_{ji})$. A *raising operator* (resp. *lowering operator*) is an element $X \in \bigoplus_{\alpha \in \Phi^+_{\mathfrak{g}}} \mathfrak{g}_\alpha$ (resp. $X \in \bigoplus_{\alpha \in \Phi^+_{\mathfrak{g}}} \mathfrak{g}_{-\alpha}$). Let $\mathfrak{n} = \bigoplus_{\alpha \in \Phi^+_{\mathfrak{g}}}$ be the space of raising operators; the condition that a root $\alpha_1 + \alpha_2$ is positive whenever $\alpha_1, \alpha_2$ are positive make $\mathfrak{n}$ into a Lie algebra.

For $\mathfrak{gl}_k$, the space $\mathfrak{n} = \mathrm{span}\{E_{i,j} : i < j\}$ of raising operators is the space of strictly upper triangular matrices and the space of lowering operators is the space of strictly lower triangular matrices.

Let $V$ be a representation of $\mathfrak{g}$. A *highest weight vector* for $V$ is a weight vector $v \in V$ with the property that $\mathfrak{n}.v = 0$. The corresponding weight is called a *highest weight* for $V$. When $\mathfrak{g} = \mathfrak{gl}_k$ and $V$ is a space of metapolynomials, it is not hard to verify that this condition is equivalent to the characterization of highest weight metapolynomials in Section 1; the characterization in terms of tableaux is explained in details in Section 4.5. A fundamental result in representation theory is that if $V$ is irreducible, then $V$ has a unique highest weight and this highest weight uniquely determines $V$ up to isomorphism [32, Prop. 14.13]. In particular, this result classifies finite-dimensional representations of $\mathfrak{g}$.

Only a subset of weights occur as highest weights of finite-dimensional irreducible representations. Such vectors are called *dominant weights* and their characterization in general requires the introduction of a suitable inner product structure on $\mathfrak{h}^*$, which is not straightforward. Let $\Lambda^+$ denote the subset of the weight lattice consisting of dominant weights. The representation associated to the highest weight $\lambda$ is denoted $V_\lambda$ and it is the unique $\mathfrak{g}$-representation having highest weight $\lambda$. In the case of $\mathfrak{gl}_k$, the set of dominant weights consist of weights $\lambda = (\lambda_1, \ldots, \lambda_k)$ such that $\lambda_j \geq \lambda_{j+1}$. If $\lambda_i \geq 0$ for every $i$, we usually identify the weight $\lambda$ with a partition of length $k$.

▶ **Example 4.8.** We determine the highest weights vectors in the representations $\mathbb{C}[x_1, x_2, x_3]_2$ and $\mathbb{C}[\mathbb{C}[x_1, x_2, x_3]_2]_3$. The space of raising operators of $\mathfrak{gl}_3$ is $\mathfrak{n} = \operatorname{span}\{E_{12}, E_{23}, E_{13}\}$ of upper triangular matrices. The highest weight vectors are weight vectors characterized by the condition $\mathfrak{n}.v = 0$. Since $E_{13} = [E_{12}, E_{23}]$, the condition $\mathfrak{n}.v = 0$ is equivalent to $E_{12}.v = E_{23}.v = 0$.

For $\mathbb{C}[x_1, x_2, x_3]_2$, one can check the unique highest weight vector is $x_3^2$ and the corresponding highest weight is $(0, 0, -2)$. Therefore $\mathbb{C}[x_1, x_2, x_3]_2$ is irreducible, and it is isomorphic to the $\mathfrak{gl}_3$-representation $V_{(0,0,-2)}$.

For $\mathbb{C}[\mathbb{C}[x, y, z]_2]_3$, the highest weight vectors are the metapolynomials

$$c_{2,0,0}^3, \qquad c_{2,0,0}c_{1,1,0}^2 - 4c_{2,0,0}^2 c_{0,2,0} \qquad \text{and}$$
$$4c_{2,0,0}c_{0,2,0}c_{0,0,2} - c_{0,0,2}c_{1,1,0}^2 + c_{1,0,1}c_{1,1,0}c_{0,1,1} - c_{0,2,0}c_{1,0,1}^2 - c_{2,0,0}c_{0,1,1}^2.$$

Their weights are $(6, 0, 0)$, $(4, 2, 0)$ and $(2, 2, 2)$. Therefore $\mathbb{C}[\mathbb{C}[x_1, x_2, x_3]_2]_3$ is a direct sum of three irreducible $\mathfrak{gl}_3$-representations

$$\mathbb{C}[\mathbb{C}[x_1, x_2, x_3]_2]_3 = V_{(6,0,0)} \oplus V_{(4,2,0)} \oplus V_{(2,2,2)}.$$

## 4.3 Universal enveloping algebra

The universal enveloping algebra of a Lie algebra $\mathfrak{g}$ is an associative algebra associated to the Lie algebra $\mathfrak{g}$. It is used to study representation theory of Lie algebras using representation theory of associative algebras.

A representation of an associative (unital) algebra $\mathcal{A}$ is a vector space $V$ equipped with a homomorphism of associative algebras $\rho\colon \mathcal{A} \to \operatorname{End}(V)$. For $v \in V$ and $U \in \mathcal{A}$, write $U.v = \rho(U)v$, as in the Lie algebra setting. Subrepresentations and equivariant maps are defined for representations of an associative algebra in the same way as for Lie algebra representations. For a subset $S \subseteq \mathcal{A}$, the subalgebra generated by $S$ is the smallest subalgebra of $\mathcal{A}$ containing $S$; explicitly, it can be realized as the set of (noncommutative) polynomials in elements of $S$.

Let $\mathfrak{g}$ be a Lie algebra. Consider a representation $\rho\colon \mathfrak{g} \to \operatorname{End}(V)$. Since $\operatorname{End}(V)$ is an associative algebra, we can consider the subalgebra of $\operatorname{End}(V)$ generated by the linear maps in the image $\rho(\mathfrak{g})$. Since $\rho$ is a representation, these generators satisfy the relations

$$\rho(X)\rho(Y) - \rho(Y)\rho(X) = \rho([X, Y]). \tag{6}$$

The *universal enveloping algebra* $\mathcal{U}(\mathfrak{g})$ is an associative algebra constructed in such a way that the relations (6) are forced for every representation $\rho$ of $\mathcal{U}(\mathfrak{g})$. It is a quotient of the *tensor algebra* $\mathcal{T}(\mathfrak{g})$ over $\mathfrak{g}$. The tensor algebra is spanned by formal products of elements of $\mathfrak{g}$,

$$\mathcal{T}(\mathfrak{g}) = \bigoplus_{k=0}^{\infty} \mathfrak{g}^{\otimes k} = \mathbb{C} \oplus \mathfrak{g} \oplus \mathfrak{g}^{\otimes 2} \oplus \mathfrak{g}^{\otimes 3} \oplus \cdots \tag{7}$$

with multiplication induced by the tensor products $\otimes\colon \mathfrak{g}^{\otimes p} \times \mathfrak{g}^{\otimes q} \to \mathfrak{g}^{\otimes (p+q)}$. In particular, the tensor algebra is graded, with $\mathfrak{g}^{\otimes p}$ being the component of degree $p$. Consider the ideal $\mathcal{I}(\mathfrak{g}) \subset \mathcal{T}(\mathfrak{g})$ generated by the elements of the form

$$X \otimes Y - Y \otimes X - [X, Y]$$

for $X, Y \in \mathfrak{g}$. The universal enveloping algebra of $\mathfrak{g}$ is defined as $\mathcal{U}(\mathfrak{g}) = \mathcal{T}(\mathfrak{g})/\mathcal{I}(\mathfrak{g})$. The multiplication in $\mathcal{U}(\mathfrak{g})$ is denoted by the usual multiplication sign rather than $\otimes$. The inclusion of $\mathfrak{g}$ into $\mathcal{T}(\mathfrak{g})$ (as a vector space) gives rise to a linear map $\mathfrak{g} \to \mathcal{U}(\mathfrak{g})$, which is also injective. We identify $\mathfrak{g}$ with the image of this injection. The elements of $\mathcal{T}(\mathfrak{g})$ and, therefore, $\mathcal{U}(\mathfrak{g})$ can be written as noncommutative polynomials in elements of $\mathfrak{g}$. When $\mathfrak{g}$ is abelian (for instance, in the case of a Cartan algebra), the ideal $\mathcal{I}(\mathfrak{g})$ is indeed homogeneous, since $[X, Y] = 0$; in this case $\mathcal{U}(\mathfrak{g})$ is the symmetric algebra of $\mathfrak{g}$, isomorphic to the polynomial ring on $\mathfrak{g}^*$.

By construction of the ideal $\mathcal{I}(\mathfrak{g})$ we have $XY - YX = [X, Y]$ in $\mathcal{U}(\mathfrak{g})$ for every $X, Y \in \mathfrak{g}$. This implies that every representation of $\mathcal{U}(\mathfrak{g})$ restricts to a Lie algebra representation of the Lie algebra $\mathfrak{g}$, via the embedding $\mathfrak{g} \subseteq \mathcal{U}(\mathfrak{g})$; conversely every representation $\rho\colon \mathfrak{g} \to \mathrm{End}(V)$ of $\mathfrak{g}$ uniquely extends to a representation of $\mathcal{U}(\mathfrak{g})$ by setting $\rho(X_1 X_2 \dots X_\ell) = \rho(X_1)\rho(X_2)\dots\rho(X_\ell)$ for all $X_i \in \mathfrak{g}$.

Note that in general the ideal $\mathcal{I}(\mathfrak{g})$ is not homogeneous, so $\mathcal{U}(\mathfrak{g})$ does not inherit the grading from the tensor algebra $\mathcal{T}(\mathfrak{g})$. However, $\mathcal{U}(\mathfrak{g})$ has a *filtration*, in the following sense: define $\mathcal{U}(\mathfrak{g})_{\le \ell} = \{X + \mathcal{I}(\mathfrak{g}) \mid X \in \bigoplus_{k=0}^{\ell} \mathfrak{g}^{\otimes k}\}$; then $\mathcal{U}(\mathfrak{g})_{\le \ell} \subseteq \mathcal{U}(\mathfrak{g})_{\le (\ell+1)}$ for every $\ell$, and if $X \in \mathcal{U}(\mathfrak{g})_{\le p}, Y \in \mathcal{U}(\mathfrak{g})_{\le q}$ then $XY \in \mathcal{U}(\mathfrak{g})_{\le (p+q)}$.

In particular, the elements of $\mathcal{U}(\mathfrak{g})_{\le \ell}$ are represented by noncommutative polynomials of degree at most $\ell$ in elements of $\mathfrak{g}$. The *length* of an element $X \in \mathcal{U}(\mathfrak{g})$ is the minimal $\ell$ such that $X \in \mathcal{U}(\mathfrak{g})_{\le \ell}$.

▶ **Theorem 4.9** (Poincaré–Birkhoff–Witt [42, Theorem 9.10])**.** *Let $X_1, \dots, X_r$ be a basis of $\mathfrak{g}$. Then the set of ordered monomials $X_{i_1} X_{i_2} \cdots X_{i_d}$, $i_1 \le i_2 \le \cdots \le i_d$ is a basis of $\mathcal{U}(\mathfrak{g})$. Moreover, the set of ordered monomials of length at most $\ell$ is a basis of the subspace $\mathcal{U}(\mathfrak{g})_{\le \ell}$ of elements with length at most $\ell$.*

In fact, we will only use that ordered monomials are a set of generators for the spaces $\mathcal{U}(\mathfrak{g})_{\le \ell}$. This is proved explicitly in Theorem B.1. The Poincaré–Birkhoff–Witt theorem (PBW theorem) shows that elements of $\mathcal{U}(\mathfrak{g})$ have unique representation as noncommutative polynomials in $X_1, \dots, X_r$ with ordered terms, and the length of the element is given by the longest monomial in this polynomial.

▶ **Example 4.10.** Consider $\mathfrak{g} = \mathfrak{gl}_2$ with a $\mathfrak{gl}_2$-representation $V$. Then $F \coloneqq E_{2,1}, E \coloneqq E_{1,2}, H_i \coloneqq E_{i,i}$ gives an ordered basis $(F, E, H_1, H_2)$ of $\mathfrak{gl}_2$. Examples of elements of $\mathcal{U}(\mathfrak{gl}_2)$ are $H_1 EF$ and $1 - \frac{1}{2}FE + \frac{1}{12}F^2E^2$. Note that these are formal products, not matrix multiplication (general Lie algebras do not have this additional algebra structure). Taking $v \in V$, the second element acts as $(1 - \frac{1}{2}FE + \frac{1}{12}F^2E^2).v = v - \frac{1}{2}F.(E.v) + \frac{1}{12}F.(F.(E.(E.v))) \in V$. The PBW theorem guarantees that these elements can be written as polynomials with ordered monomials. Indeed, we have

$$\begin{aligned} H_1EF &= H_1(FE + [E, F]) = H_1FE + H_1(H_1 - H_2) \\ &= \cdots = FEH_1 + 2FE + H_1^2 - H_1H_2. \end{aligned}$$

## 4.4 Central characters and Casimir elements

Let $\mathfrak{g}$ be a Lie algebra. An element $X \in \mathcal{U}(\mathfrak{g})$ of the universal enveloping algebra is *central* if $XY = YX$ for every $Y \in \mathcal{U}(\mathfrak{g})$ (or equivalently every $Y \in \mathfrak{g}$). The center $\mathcal{Z}(\mathfrak{g})$ of $\mathcal{U}(\mathfrak{g})$ is the set of all central elements; $\mathcal{Z}(\mathfrak{g})$ is a commutative associative subalgebra of $\mathcal{U}(\mathfrak{g})$.

A representation $\rho\colon \mathfrak{g} \to \mathrm{End}(V)$ of $\mathfrak{g}$ extends to a representation of $\mathcal{U}(\mathfrak{g})$, and if $X \in \mathcal{Z}(\mathfrak{g})$ is a central element, the map $\rho(X) : V \to V$ is $\mathcal{U}(\mathfrak{g})$-equivariant. If $V$ is irreducible, then by Schur's lemma (Lemma 4.4) $\rho(X)$ is a multiplication by some scalar $\chi_V(X)$. More precisely, $\rho$ induces a 1-dimensional representation $\chi_V : \mathcal{Z}(\mathfrak{g}) \to \mathbb{C}$. This representation is called the *central character* of $V$.

It is a crucial fact that for Lie algebras $\mathfrak{g}$ of reductive groups central characters separate non-isomorphic irreducible representations.

▶ **Theorem 4.11** ([65], see also [10, §VIII.8.5, Cor.2]). *Let $\mathfrak{g}$ be a reductive Lie algebra. If $V$ and $W$ are irreducible representations of $\mathfrak{g}$ and the central characters $\chi_V$ and $\chi_W$ coincide, then $V$ and $W$ are isomorphic.*

It follows that for reductive $\mathfrak{g}$ the isotypic components of a $\mathfrak{g}$-representation are also isotypic components of the induced $\mathcal{Z}(\mathfrak{g})$-representation. More specifically, if a $\mathfrak{g}$-representation $V$ decomposes as $\bigoplus_{\lambda\in\Lambda} V_\lambda^{\oplus m_\lambda}$ with pairwise non-isomorphic $V_\lambda$, then as a $\mathcal{Z}(\mathfrak{g})$-representation we have $V \cong \bigoplus_{\lambda\in\Lambda} \chi_\lambda^{\oplus m_\lambda \cdot \dim V_\lambda}$, where we denote by $\chi_\lambda$ the 1-dimensional representation of $\mathcal{Z}(\mathfrak{g})$ corresponding to the central character of $V_\lambda$.

▶ **Example 4.12.** Consider the $\mathfrak{gl}_3$ representation $V = \mathbb{C}[x_1, x_2, x_3]_d$ of homogeneous degree 3 polynomials in three variables (Claim 4.1). Consider the element $I \coloneqq -(E_{1,1} + E_{2,2} + E_{3,3})$. Clearly, $I \in \mathcal{Z}(\mathfrak{gl}_3)$. We determine the value of the central characters. Take $f \in V$. We find

$$I.f = x_1 \frac{\partial}{\partial x_1} f + x_2 \frac{\partial}{\partial x_2} f + x_3 \frac{\partial}{\partial x_3} f = df.$$

So $\chi_V(I) = d$. As a simple applications of Theorem 4.11, this allows us to conclude $\mathbb{C}[x_1, x_2, x_3]_d$ and $\mathbb{C}[x_1, x_2, x_3]_{d'}$ contain no isomorphic irreducible subrepresentations unless $d = d'$.

When $\mathfrak{g}$ is the Lie algebra of a reductive group, then $\mathcal{Z}(\mathfrak{g})$ can be described explicitly as the ring of invariants for the the action of the Weyl group on the Cartan subalgebra. Note that since the Cartan subalgebra $\mathfrak{h}$ is abelian, the universal enveloping algebra $\mathcal{U}(\mathfrak{h})$ is in fact commutative and can be identified with the polynomial algebra $\mathbb{C}[\mathfrak{h}^*]$. Recall that the Weyl group $\mathcal{W}_\mathfrak{g}$ acts on $\mathfrak{h}$ and, therefore, on $\mathcal{U}(\mathfrak{h}) \cong \mathbb{C}[\mathfrak{h}^*]$.

▶ **Theorem 4.13** (Harish–Chandra isomorphism, [44, Ch. 23];[6]). *Let $\mathfrak{g}$ be a reductive algebra and let $\mathfrak{h}$ be a Cartan subalgebra. The center $\mathcal{Z}(\mathfrak{g})$ of the universal enveloping algebra is isomorphic to the algebra of invariants $\mathcal{U}(\mathfrak{h})^{\mathcal{W}_\mathfrak{g}}$. Moreover, if $\dim \mathfrak{h} = k$, there exist algebraically independent $C_1, \dots, C_k \in \mathcal{U}(\mathfrak{g})$ such that $\mathcal{Z}(\mathfrak{g}) = \mathbb{C}[C_1, \dots, C_k]$. The lengths of these elements coincide with the degrees of the fundamental invariants of $\mathcal{U}(\mathfrak{h})^{\mathcal{W}_\mathfrak{g}}$.*

The generators $C_1, \dots, C_k$ of the center $\mathcal{Z}(\mathfrak{g})$ are known as *Casimir elements*. These generators are not unique.

Consider the case of $\mathfrak{gl}_k$, with $\mathfrak{h}$ the Cartan subalgebra of diagonal matrices. Then $\mathcal{W}_{\mathfrak{gl}_k}$ is the symmetric group $\mathfrak{S}_k$, and it acts on $\mathfrak{h}$ by permuting the diagonal elements. Then $\mathcal{U}(\mathfrak{h})$ is the polynomial ring in $k$ variables, and $\mathcal{U}(\mathfrak{h})^{\mathcal{W}_{\mathfrak{gl}_k}}$ is the subring of polynomials invariant under permutation of variables: this is the classical ring of symmetric polynomials in $k$ variables. It is generated, for instance, by the elementary symmetric polynomials, which have degrees $1, \dots, k$. These numbers give lengths of Casimir elements generating the center $\mathcal{Z}(\mathfrak{g})$. There are several constructions of Casimir elements, one possible choice being

$$C_p \coloneqq \sum_{i_1=1}^{k} \cdots \sum_{i_p=1}^{k} E_{i_1,i_2} E_{i_2,i_3} \cdots E_{i_p,i_1} \quad \text{with } 1 \le p \le k. \tag{8}$$

▶ **Example 4.14.** Let $\mathfrak{g} = \mathfrak{gl}_2$. Then the Casimir elements are given by

$$C_1 = E_{1,1} + E_{2,2} \qquad \text{and} \qquad C_2 = E_{1,1}E_{1,1} + E_{1,2}E_{2,1} + E_{2,1}E_{1,2} + E_{2,2}E_{2,2}.$$

We will use Claim 4.2 to evaluate the application of $C_1$ and $C_2$ to the metapolynomials

$$\Delta = c_{11}^2 + 2\,c_{02}c_{20} \qquad \text{and} \qquad \Gamma = -c_{11}^2 + 4\,c_{02}c_{20}.$$

We find that

$$C_1.\Delta = E_{1,1}.\Delta + E_{2,2}.\Delta = 2c_{11}^2 + 4c_{02}c_{20} + 2c_{11}^2 + 4c_{02}c_{20} = 4\Delta$$

and similarly $C_1.\Gamma = 4\Gamma$ (in fact, one can show $C_1$ always scales a metapolynomial of format $(\delta, d, k)$ by $d\delta$). We also compute

$$\begin{aligned} E_{2,1}.\Delta &= 2\,c_{02}^2c_{11} + 2\,c_{11}^2c_{02} + 2\cdot 0c_{20} + 2\,c_{02}c_{11} = 6\,c_{02}^2c_{11} \\ (E_{1,2}E_{2,1}).\Delta &= 6\,c_{11}c_{11} + 6\cdot 2\,c_{02}c_{20} = 6\Delta \end{aligned}$$

By symmetry of multiindices in the metapolynomial we know $(E_{2,1}E_{1,2}).\Delta = 6\Delta$ too. An analogous computation shows $E_{2,1}.\Gamma = 0$ and $E_{1,2}.\Gamma = 0$. We find

$$C_2.\Delta = \big(2^2 + 6 + 6 + 2^2\big)\Delta = 20\Delta \qquad \text{and} \qquad C_2.\Gamma = \big(2^2 + 0 + 0 + 2^2\big)\Delta = 8\Delta.$$

We obtain that $\Delta$ and $\Gamma$ live in different isotypic components with respect to $\mathcal{Z}(\mathfrak{gl}_2)$, and hence with respect to $\mathfrak{gl}_2$.

The eigenvalues of these Casimir elements on irreducible representations of $\mathfrak{gl}_k$ are computed by Perelomov and Popov [64]:

$$\chi_\lambda(C_p) = \mathsf{Tr}(A_\lambda^p E) \tag{9}$$

where $E$ is the $k \times k$ matrix consisting of all ones, and the entries of the $k \times k$ matrix $A_\lambda = (a_{ij}^{(\lambda)})$ are given by

$$a_{ij}^{(\lambda)} = \begin{cases} \lambda_i + k - i, & \text{if } i = j, \\ 0, & \text{if } i > j, \\ -1, & \text{if } i < j. \end{cases} \tag{10}$$

▶ **Example 4.15.** Observe first that $\mathsf{Tr}(A_\lambda^p E)$ is simply the sum of the entries of $A_\lambda^p$. Consider again Example 4.14. We compute the formula in Equation (9).

$$\begin{aligned} A_{(4,0)} &= \left[\begin{smallmatrix} 5 & -1 \\ 0 & 2 \end{smallmatrix}\right] \quad \text{and thus} \quad \chi_{(4,0)}(C_2) = \mathsf{Tr}\left(\left[\begin{smallmatrix} 25 & -7 \\ 0 & 4 \end{smallmatrix}\right]\left[\begin{smallmatrix} 1 & 1 \\ 1 & 1 \end{smallmatrix}\right]\right) = 20 \\ A_{(2,2)} &= \left[\begin{smallmatrix} 3 & -1 \\ 0 & 2 \end{smallmatrix}\right] \quad \text{and thus} \quad \chi_{(2,2)}(C_2) = \mathsf{Tr}\left(\left[\begin{smallmatrix} 9 & -5 \\ 0 & 4 \end{smallmatrix}\right]\left[\begin{smallmatrix} 1 & 1 \\ 1 & 1 \end{smallmatrix}\right]\right) = 8. \end{aligned}$$

Observe that this matches our results in Example 4.14. We give one more example for $\mathfrak{gl}_3$.

$$A_{(6,0,0)} = \left[\begin{smallmatrix} 8 & -1 & -1 \\ 0 & 1 & -1 \\ 0 & 0 & 0 \end{smallmatrix}\right] \quad \text{and thus} \quad \chi_{(6,0,0)}(C_2) = \mathsf{Tr}\left(\left[\begin{smallmatrix} 64 & -9 & -7 \\ 0 & 1 & -1 \\ 0 & 0 & 0 \end{smallmatrix}\right]\left[\begin{smallmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{smallmatrix}\right]\right) = 48.$$

## 4.5 Gelfand–Tsetlin theory

Gelfand–Tsetlin bases are used to explicitly construct irreducible representations of some classical Lie algebras. The explicit bases first appeared in [35] without much explanation. The construction was expanded upon by Želobenko [72] and independently by Baird and Biedenharn [4]. It can be explained in terms of restrictions of irreducible representations along a chain of Lie algebra inclusions.

Let $\mathfrak{a}$ be a Lie subalgebra of $\mathfrak{b}$. Every representation $V$ of $\mathfrak{b}$ can be regarded as a representation of $\mathfrak{a}$ using this inclusion. The resulting representation is called the *restriction of $V$ to $\mathfrak{a}$* and is denoted as $V{\downarrow}^{\mathfrak{b}}_{\mathfrak{a}}$.

Let $\mathfrak{g}_1 \subset \mathfrak{g}_2 \subset \cdots \subset \mathfrak{g}_k$ be an inclusion chain of reductive Lie algebras. We call this chain a *Gelfand–Tsetlin chain* if $\mathfrak{g}_1$ is abelian and every irreducible representation $V$ of $\mathfrak{g}_\ell$ $(1 < \ell \leq k)$, when restricted to $\mathfrak{g}_{\ell-1}$, is a multiplicity-free completely reducible representation, that is,

$$V_\lambda{\downarrow}^{\mathfrak{g}_\ell}_{\mathfrak{g}_{\ell-1}} \cong \bigoplus_{\mu \in R_\ell(\lambda)} V_\mu \tag{11}$$

for some set $R_\ell(\lambda)$ of isomorphism types of irreducible representations of $\mathfrak{g}_{i-1}$. We write $\mu \to \lambda$ if $V_\mu$ appears in the decomposition for $V_\lambda$.

Let $V_\lambda$ be a irreducible representation of $\mathfrak{g}_k$. By repeatedly restricting irreducible representations of $\mathfrak{g}_i$ we obtain a decomposition of $V_\lambda$ as a representation of $\mathfrak{g}_1$.

$$V_\lambda{\downarrow}^{\mathfrak{g}_k}_{\mathfrak{g}_1} = \bigoplus_{\lambda^{(1)} \to \lambda^{(2)} \to \cdots \to \lambda^{(k)} = \lambda} V_{\lambda^{(1)}} \tag{12}$$

Since $\mathfrak{g}_1$ is abelian, its irreducible representations are 1-dimensional, so the decomposition above distinguishes a set of 1-dimensional subspaces in $V_\lambda$ indexed by chains $\lambda^{(1)} \to \lambda^{(2)} \to \cdots \to \lambda^{(k)}$ called *Gelfand–Tsetlin patterns*. We denote the subspace of $V_{\lambda^{(k)}}$ corresponding to a Gelfand–Tsetlin pattern $T = \lambda^{(1)} \to \lambda^{(2)} \to \cdots \to \lambda^{(k)}$ by $V_T$.

The inclusions $\mathfrak{g}_\ell \subseteq \mathfrak{g}_k$ give an inclusion of universal enveloping algebras $\mathcal{U}(\mathfrak{g}_\ell) \subseteq \mathcal{U}(\mathfrak{g}_k)$, and therefore a chain $\mathcal{U}(\mathfrak{g}_1) \subset \mathcal{U}(\mathfrak{g}_2) \subset \cdots \subset \mathcal{U}(\mathfrak{g}_k)$. The *Gelfand–Tsetlin algebra* $\mathcal{GZ}(\mathfrak{g}_k)$ of $\mathfrak{g}_k$ is the subalgebra of $\mathcal{U}(\mathfrak{g})$ generated by all centers $\mathcal{Z}(\mathfrak{g}_\ell)$ with $\ell \leq k$. Since the center $\mathcal{Z}(\mathfrak{g}_\ell)$ of $\mathcal{U}(\mathfrak{g}_\ell)$ commutes with all elements of $\mathcal{U}(\mathfrak{g}_\ell)$, it commutes in particular with the centers $\mathcal{Z}(\mathfrak{g}_j)$ for $j < \ell$; as a result, $\mathcal{GZ}(\mathfrak{g}_k)$ is a commutative subalgebra.

Consider a Gelfand–Tsetlin pattern $T = \lambda^{(1)} \to \lambda^{(2)} \to \cdots \to \lambda^{(k)}$ and the corresponding 1-dimensional subspace $V_T \subset V_{\lambda^{(k)}}$. For every $\ell \leq k$ the subspace $V_T$ is contained in an irreducible representation of $\mathfrak{g}_\ell$ isomorphic to $V_{\lambda^{(\ell)}}$, and the center $\mathcal{Z}(\mathfrak{g}_\ell)$ acts on $V_T$ via the central character $\chi_{\lambda^{(\ell)}}$. It follows that $V_T$ is a 1-dimensional representation of $\mathcal{GZ}(\mathfrak{g}_k)$ given by a *Gelfand–Tsetlin character* $\chi_T \colon \mathcal{GZ}(\mathfrak{g}_k) \to \mathbb{C}$ such that $X.v = \chi_T(X)v$ for $X \in \mathcal{GZ}(\mathfrak{g}_k)$, $v \in V_T$. More specifically, on $\mathcal{Z}(\mathfrak{g}_\ell)$ we have $\chi_T(X) = \chi_{\lambda^{(\ell)}}(X)$. Theorem 4.11 implies that different Gelfand–Tsetlin patterns correspond to different Gelfand–Tsetlin characters.

By the previous discussion, every completely reducible representation of $\mathfrak{g}_k$ is also completely reducible as a representation of $\mathcal{GZ}(\mathfrak{g}_k)$. That is, if $V$ is a completely reducible representation of $\mathfrak{g}_k$ with isotypic decomposition $V = \bigoplus_{\lambda \in \Lambda} V_\lambda^{\oplus m_\lambda}$, then the action of $\mathcal{GZ}(\mathfrak{g}_k)$ on $V$ splits the decomposition further into Gelfand–Tsetlin characters as

$$V = \bigoplus_{T \colon \lambda^{(1)} \to \cdots \to \lambda^{(k)}, \lambda^{(k)} \in \Lambda} \chi_T^{\oplus m_{\lambda^{(k)}}}. \tag{13}$$

We call the components of this decomposition *GZ-isotypic* spaces of $V$.

The classical case – used to construct a basis of irreducible representations of $\mathfrak{gl}_n$ – corresponds to the chain of inclusions $\mathfrak{gl}_1 \subset \mathfrak{gl}_2 \subset \cdots \subset \mathfrak{gl}_k$ where $\mathfrak{gl}_{k-1}$ is included in $\mathfrak{gl}_k$ as the subalgebra of $\mathfrak{gl}_k$ consisting of matrices having the $k$-th row and the $k$-th column identically equal to 0. The Gelfand–Tsetlin algebra in this case is the polynomial algebra $\mathcal{GZ}(\mathfrak{g}_k) = \mathbb{C}[C_{\ell,p}]$ generated by the Casimir elements

$$C_{\ell,p} \coloneqq \sum_{i_1=1}^{\ell} \cdots \sum_{i_p=1}^{\ell} E_{i_1,i_2} E_{i_2,i_3} \cdots E_{i_p,i_1} \quad \text{with } 1 \le \ell \le k \text{ and } 1 \le p \le \ell. \tag{14}$$

The irreducible representations of $\mathfrak{gl}_k$ correspond highest weights $\lambda = (\lambda_1, \ldots, \lambda_k)$ with integer $\lambda_1 \ge \cdots \ge \lambda_k$. The restriction $V_\lambda{\downarrow}^{\mathfrak{gl}_k}_{\mathfrak{gl}_{k-1}}$ is given by the classical branching rule (attributed to Schur in [58]) which states that $\mu \to \lambda$ if and only if $\mu$ and $\lambda$ *interleave*, that is $\lambda_\ell \ge \mu_\ell \ge \lambda_{\ell+1}$ for all $\ell$, $1 \le \ell \le k-1$. Moreover, $E_{k,k}$ acts on the copy of $V_\mu$ in $V_\lambda$ as multiplication by $|\lambda| - |\mu| = \sum_{\ell=1}^{k-1} (\lambda_k - \mu_k)$.

A *Gelfand-Tsetlin pattern* is given by a sequence $\lambda^{(1)} \to \lambda^{(2)} \to \cdots \to \lambda^{(k)}$, that is, by integers $\lambda_i^{(\ell)}$ with $1 \le \ell \le k$, $1 \le i \le \ell$ such that $\lambda_i^{(\ell)} \ge \lambda_i^{(\ell-1)} \ge \lambda_{i+1}^{(\ell)}$, often arranged into a triangular shape, see Figure 2.

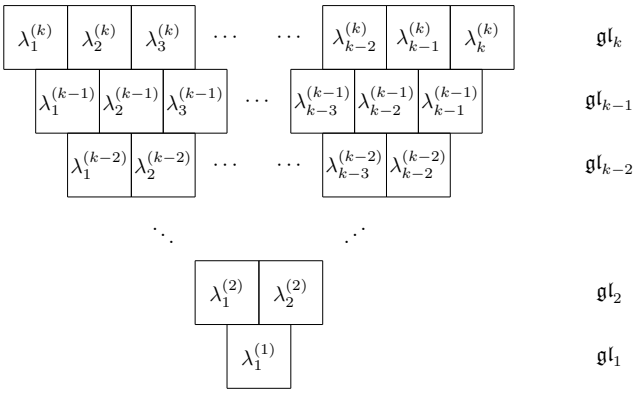


**Figure 2** Gelfand–Tsetlin pattern of $\mathfrak{gl}_k$.

The 1-dimensional subspace $V_T$ corresponding to the pattern $T = \lambda^{(1)} \to \lambda^{(2)} \to \cdots \to \lambda^{(k)}$ is spanned by a weight vector of weight $(|\lambda^{(1)}|, |\lambda^{(2)}| - |\lambda^{(1)}|, \ldots, |\lambda^{(k)}| - |\lambda^{(k-1)}|)$. The generators $C_{\ell,p}$ of $\mathcal{GZ}(\mathfrak{gl}_k)$ act on this space as multiplications by the corresponding eigenvalues $\chi_T(C_{\ell,p})$. Since $C_{\ell,p}$ is a Casimir operator for $\mathfrak{gl}_\ell$, we have $\chi_T(C_{\ell,p}) = \chi_{\lambda^{(\ell)}}(C_{\ell,p}) = \mathsf{Tr}(A^p_{\lambda^{(\ell)}} E)$ from Perelomov–Popov formula.

If $\lambda^{(k)}$ is a partition, that is, all $\lambda_i^{(k)}$ are nonnegative, then the Gelfand–Tsetlin patterns for $\mathfrak{gl}_k$ are in bijective correspondence with semistandard Young tableaux filled with numbers from $\{1, 2, \ldots, k\}$. Recall that a Young tableau is semistandard whenever the numbers are strictly increasing along the columns, and weakly increasing along the rows.

A semistandard tableau $T$ gives rise to a sequence of partitions $\lambda^{(1)}, \ldots, \lambda^{(k)}$ where $\lambda^{(\ell)}$ is the shape of the tableau $T^{(\ell)}$ obtained from $T$ be removing all boxes labeled with numbers greater than $\ell$. Semistandardness of $T$ guarantees that all $T^{(\ell)}$ are also semistandard tableaux and the interleaving property $\lambda_i^{(\ell)} \ge \lambda_i^{(\ell-1)} \ge \lambda_{i+1}^{(\ell)}$ holds (see e. g. [56, p.5]). The weight is given by the content of the tableau $T$, that is, the vector $\mu = (\mu_1, \ldots, \mu_k)$ where $\mu_i$ records the number of occurrences of the label $i$ in the tableau.

For example, the tableau $T = \begin{array}{|c|c|c|c|}\hline 1&1&2&2\\\hline 2&3\\\cline{1-2} 4&4\\\cline{1-2}\end{array}$ gives

$$T^{(1)} = \begin{array}{|c|c|}\hline 1&1\\\hline\end{array}, \quad T^{(2)} = \begin{array}{|c|c|c|c|}\hline 1&1&2&2\\\hline 2\\\cline{1-1}\end{array}, \quad T^{(3)} = \begin{array}{|c|c|c|c|}\hline 1&1&2&2\\\hline 2&3\\\cline{1-2}\end{array}, \quad T^{(4)} = T = \begin{array}{|c|c|c|c|}\hline 1&1&2&2\\\hline 2&3\\\cline{1-2} 4&4\\\cline{1-2}\end{array},$$

and the sequence of shapes is $\lambda^{(1)} = (2)$, $\lambda^{(2)} = (4,1)$, $\lambda^{(3)} = (4,2,0)$, $\lambda^{(4)} = (4,2,2,0)$. The corresponding weight is $(2,3,1,2)$, the content of the tableau.

The superstandard tableau of shape $\lambda = (\lambda_1, \ldots, \lambda_n)$ corresponds to the Gelfand-Tsetlin pattern with $\lambda^{(\ell)} = (\lambda_1, \ldots, \lambda_\ell)$. The corresponding 1-dimensional subspace of $V_\lambda$ contains the highest weight vectors. To prove this, recall that since $E_{i,j}$ with $i \neq j$ span root spaces of $\mathfrak{gl}_k$, they shift the weights of weight vectors, that is, $E_{i,j}.v$ is a weight vector of weight $\mu + e_i - e_j$. The content of the superstandard tableau of shape $\lambda$ is $\lambda$, so every nonzero vector $v$ in the corresponding 1-dimensional subspace has weight $\lambda$, and for every $E_{i,j}$ the vector $E_{i,j}.v$ has weight $\lambda + e_i - e_j$. But there are no semistandard tableaux with shape $\lambda$ and content $\lambda + e_i - e_j$ when $i < j$. Therefore, $E_{i,j}.v = 0$ if $i < j$, and it follows that $v$ is annihilated by every strictly upper triangular matrix.

## 5 Efficient construction of projectors

In this section we prove Theorem 1.1. We explain the construction of the projectors in Section 5.1, and describe the efficient circuit implementation in Section 5.2.

### 5.1 Construction of projectors

We construct the required projections as projections onto a common eigenspace of several commuting matrices. The construction is elementary and is based on the classical construction of eigenprojectors from basic linear algebra.

We present the construction in full generality in Theorem 5.2. The projectors in the cases relevant for Theorem 1.1 are obtained by applying the general construction to suitable commutative subalgebras of the universal enveloping algebra $\mathcal{U}(\mathfrak{gl}_k)$, see Table 1. If $\mathcal{A}$ is a commutative algebra generated by $U_1, \ldots, U_q$, then a representation of $\rho : \mathcal{A} \to \mathrm{End}(V)$ is uniquely determined by $\rho(U_1), \ldots, \rho(U_q)$, which are $q$ commuting endomorphisms of $V \to V$. Every irreducible representation of $\mathcal{A}$ is 1-dimensional, so it is given by an algebra homomorphism $\chi\colon \mathcal{A} \to \mathbb{C}$. In other words, an irreducible representation is spanned by a common eigenvector of the generators $U_1, \ldots, U_q$ of $\mathcal{A}$ with the corresponding eigenvalues $\chi(U_1), \ldots, \chi(U_q)$.

Recall that a representation is called *completely reducible* if it can be decomposed into a direct sum of irreducible representations. All representations in this paper have this property. For a commutative algebra this means that the generators of the algebra act not just by commuting but by simultaneously diagonalizable linear maps.

We slightly abuse notation and denote the character $\chi\colon \mathcal{A} \to \mathbb{C}$ and the corresponding irreducible representation by the same letter. Every completely reducible representation $V$ of $\mathcal{A}$ has a unique *isotypic decomposition*

$$V \cong \bigoplus_{\lambda \in \Lambda} \chi_\lambda^{\oplus m_\lambda}, \tag{15}$$

where $\Lambda$ is a set indexing irreducible representations of $\mathcal{A}$. We denote the isotypic component $\chi_\lambda^{\oplus m_\lambda}$ in $V$ by $V_{\mathcal{A},\lambda}$.

The length of the projectors as elements of $\mathcal{U}(\mathfrak{g})$ depends mainly on the number of different isotypic components of the representation.

▶ **Definition 5.1.** *Let $V$ be a completely reducible representation of an algebra $\mathcal{A}$. We call the number of isomorphism types of irreducible subrepresentations of $V$ the* diversity *of $V$.*

▶ **Theorem 5.2.** *Let $\mathcal{A}$ be a commutative algebra generated by $U_1, \dots, U_q$ and $V$ be a completely reducible representation of $\mathcal{A}$ with diversity at most $p$. For every irreducible representation $\chi_\lambda$ of $\mathcal{A}$ there exists a polynomial $f_\lambda$ of degree at most $p-1$ such that*

$$f_\lambda(U_1, \dots, U_q).x = \Pi_\lambda x \tag{16}$$

*where $\Pi_\lambda \colon V \to V$ is the unique $\mathcal{A}$-equivariant projection onto the isotypic component $V_{\mathcal{A},\lambda}$.*

**Proof.** Let $\Lambda$ be the set of isomorphism types appearing in the isotypic decomposition of $V$, so that $V \cong \bigoplus_{\mu\in\Lambda} V_{\mathcal{A},\mu}$ with $V_{\mathcal{A},\mu} \neq 0$. An element $U \in \mathcal{A}$ acts as $U.x = \chi_\mu(U)x$ on $V_{\mathcal{A},\mu}$. If $\lambda \notin \Lambda$, set $f_\lambda = 0$.

Fix $\lambda \in \Lambda$. For every $\mu \in \Lambda$ with $\mu \neq \lambda$, let $U_{i_\mu}$ be a generator such that $\chi_\mu(U_{i_\mu}) \neq \chi_\lambda(U_{i_\mu})$. Such generator exist because $\chi_\lambda$ and $\chi_\mu$ are different characters of $\mathcal{A}$. Define

$$P_\lambda \coloneqq \prod_{\mu\in\Lambda,\mu\neq\lambda} \frac{U_{i_\mu} - \chi_\mu(U_{i_\mu})}{\chi_\lambda(U_{i_\mu}) - \chi_\mu(U_{i_\mu})} \quad \in \mathcal{A}. \tag{17}$$

The factor $\dfrac{U_{i_\mu} - \chi_\mu(U_{i_\mu})}{\chi_\lambda(U_{i_\mu}) - \chi_\mu(U_{i_\mu})}$ acts on $V$ as the zero map on $V_{\mathcal{A},\mu}$ and as the identity map on $V_{\mathcal{A},\lambda}$; moreover, its action maps isotypic components to themselves. The element $P_\lambda$ acts as the composition of these factors; therefore its image in $\mathrm{End}(V)$ is the projection $\Pi_\lambda$ onto $V_{\mathcal{A},\lambda}$. Finally, since every factor is linear in $U_1, \dots, U_q$, $P_\lambda$ is a polynomial in $U_1, \dots, U_q$. Its degree is (at most) $p-1$. ◀

The projections required in Theorem 1.1 arise as special cases of this construction. In each case, the image of the projection is an isotypic component of some commutative subalgebra $\mathcal{A} \subset \mathcal{U}(\mathfrak{gl}_k)$. Theorem 5.2 then gives a construction of the projectors as elements of $\mathcal{U}(\mathfrak{gl}_k)$, expressed as polynomials in the generators of $\mathcal{A}$. The length of the projectors is easily bounded in terms of lengths of generators of $\mathcal{A}$ in $\mathcal{U}(\mathfrak{gl}_K)$ and the diversity of $\mathbb{C}[\mathbb{C}[x_1, \dots, x_k]_d]_\delta$ as a representation of $\mathcal{A}$. Indeed, if the generators $U_1, \dots, U_q$ of $\mathcal{A}$ have length at most $\ell$, and the projectors are expressed as polynomials in $U_1, \dots, U_q$ of degree at most $p-1$, then the lengths of the projectors are bounded by $\ell(p-1)$. The following table summarizes these parameters for the three cases considered.

**Table 1** Isotypic components, generators of their corresponding subalgebras, and their diversity in $\mathbb{C}[\mathbb{C}[x_1, \dots, x_k]_d]_\delta$. See Equation (8) and Equation (14) for the definitions of $C_p$ and $C_{\ell,p}$.

| Isotypic components | Weight spaces | $\lambda$-isotypic spaces | $T$-isotypic spaces |
|---|---|---|---|
| Algebra $\mathcal{A}$ | $\mathcal{U}(\mathfrak{h}) \cong \mathbb{C}[\mathfrak{h}^*]$ | $\mathcal{Z}(\mathfrak{gl}_k)$ | $\mathcal{GZ}(\mathfrak{gl}_k)$ |
| Generators $U_i$ | $\{E_{i,i} \mid i \in [k]\}$ | $\{C_p \mid p \in [k]\}$ | $\{C_{\ell,p} \mid \ell \in [k], p \in [\ell]\}$ |
| Generator lengths | 1 | $\leq k$ | $\leq k$ |
| Diversity | $\leq (\delta d)^k$ | $\leq (\delta d)^k$ | $\leq (\delta d)^{k^2}$ |

▶ Remark 5.3. In the construction of Theorem 5.2 we require knowledge of the value of $\chi_\mu$ evaluated on the generators $U_i$, for all irreducible representation types $\mu$ occurring in the representation. For the weight spaces these values are given by the weights themselves, and for the other two algebra they can be computed using Equation (9). It is then also possible to remove any choices in the construction by taking for each term instead a linear combination $U$ of all $U_i$ with algebraically independent irrational weights. Since the values $\chi_\mu$ are always integer, this ensures $\chi_\mu(U) \neq \chi_\lambda(U)$ when $\mu \neq \lambda$.

▶ **Example 5.4.** We construct and evaluate a $\lambda$-isotypic projector for metapolynomials of format $(3,2,3)$. From Example 4.8 we know $\Lambda = \{(6,0,0),(4,2,0),(2,2,2)\}$. By Equation (9) we find that the Casimir element $C_2$ scales these three irreducible representations with different scalars, namely 48, 28 and 12 respectively. Hence we can use $C_2$ to construct the projectors from Theorem 5.2. Consider first the projector $P_{(6,0,0)}$ to the $(6,0,0)$-isotypic component. We have that

$$P_{(6,0,0)} = \frac{C_2 - \chi_{(2,2,2)}(C_2)}{\chi_{(6,0,0)}(C_2) - \chi_{(2,2,2)}(C_2)} \cdot \frac{C_2 - \chi_{(4,2,0)}(C_2)}{\chi_{(6,0,0)}(C_2) - \chi_{(4,2,0)}(C_2)} = \frac{C_2 - 12}{36} \cdot \frac{C_2 - 28}{20}.$$

Consider again the example $\Delta = c_{002}c_{020}c_{200}$ from the introduction. We will compute $P_{(6,0,0)}.\Delta$. Using the equations in Claim 4.2 we find

$$C_2.\Delta = \sum_{i=1}^{3}\sum_{j=1}^{3} \big(E_{i,j}E_{j,i}\big).\Delta = 2\,c_{020}c_{101}^2 + 2\,c_{002}c_{110}^2 + 2\,c_{011}^2c_{200} + 24\,c_{002}c_{020}c_{200}.$$

It directly follows that

$$C_2.\Delta - 28\Delta = 2\,c_{020}c_{101}^2 + 2\,c_{002}c_{110}^2 + 2\,c_{011}^2c_{200} - 4\,c_{002}c_{020}c_{200}.$$

We apply $C_2$ again (at this point a computer algebra program is recommended) and find

$$C_2.(C_2.\Delta - 28\Delta) = 48\big(c_{020}c_{101}^2 + c_{011}c_{101}c_{110} + c_{002}c_{110}^2 + c_{011}^2c_{200}\big).$$

Hence,

$$\begin{aligned} 60\,P_{(6,0,0)}.\Delta &= \frac{C_2.(C_2.\Delta - 28\Delta) - 12(C_2.\Delta - 28\Delta)}{12} \\ &= 2\,c_{020}c_{101}^2 + 4\,c_{011}c_{101}c_{110} + 2\,c_{002}c_{110}^2 + 2\,c_{011}^2c_{200} + 4\,c_{002}c_{020}c_{200}. \end{aligned}$$

This is indeed the isotypic component as given in the introduction. The other two can be determined by an analogous computation.

### 5.1.1 Weight projectors

Let $V$ be a representation of a complex reductive Lie group $G$. The action of $G$ on $V$ induces the action of the corresponding Lie algebra $\mathfrak{g}$. Fix a maximal torus $T \subset G$ and let $\mathfrak{h} \subset \mathfrak{g}$ be the corresponding Cartan subalgebra. The torus $T$ is reductive, and the isotypic components of $V$ with respect to the action of $T$ are exactly the weight spaces. The action of $T$ induces the action of $\mathfrak{h}$ with the same isotypic components, so the action of $\mathfrak{h}$ is completely reducible.

Let $H_1, \ldots, H_r$ be a basis of the Cartan subalgebra $\mathfrak{h}$. Since $\mathfrak{h}$ is abelian, the universal enveloping algebra $\mathcal{U}(\mathfrak{h}) \subset \mathcal{U}(\mathfrak{g})$ is commutative. It can be identified with the polynomial algebra $\mathbb{C}[H_1, \ldots, H_r]$. We can now apply Theorem 5.2 to obtain for every weight $\mu \in \mathfrak{h}^*$ an element $H_\mu \in \mathcal{U}(\mathfrak{g})$ such that $H_\mu.v$ is the projection of $v \in V$ onto the weight space of

weight $\mu$. The projectors $H_\mu$ are obtained as polynomial expressions $H_\mu = f_\mu(H_1, \dots, H_r)$ with $\deg f_\mu \le p-1$ where $p$ is the diversity of $V$ as a representation of $\mathfrak{h}$ (or, equivalently, of $\mathcal{U}(\mathfrak{h})$), that is, the number of weights of $V$. Since $H_i$ as elements of $\mathcal{U}(\mathfrak{g})$ have length 1, the length of $H_\mu$ is also bounded by $p-1$.

We can specialize this construction to the action of $\mathfrak{gl}_k$ on the space of metapolynomials $W \coloneqq \mathbb{C}[\mathbb{C}[x_1, \dots, x_k]_d]_\delta$. The standard Cartan subalgebra consists of all diagonal matrices, and we can take $H_i \coloneqq E_{i,i}$. The weights appearing in $W$ correspond to $k$-tuples of nonnegative integers summing to $\delta d$ (see Example 4.7). The number of such tuples is $\binom{\delta d+k-1}{k-1} \le (\delta d)^k$.

▶ **Corollary 5.5.** *For each weight $\mu = (\mu_1, \dots, \mu_k)$ with $\sum_{i=1}^k \mu_i = \delta d$ there is an element $H_\mu \in \mathcal{U}(\mathfrak{gl}_k)$ of length at most $(\delta d)^k$ which acts on $\mathbb{C}[\mathbb{C}[x_1, \dots, x_k]_d]_\delta$ as the projector onto the weight space corresponding to $\mu$.*

### 5.1.2 Isotypic projectors

Let $V$ be a representation of a complex reductive Lie group $G$ with the corresponding Lie algebra $\mathfrak{g}$. Then $V$ is completely reducible as a representation of $\mathfrak{g}$, with the same isotypic components as for $G$. As discussed in Section 4.4, they are also isotypic components of $V$ as a representation of $\mathcal{Z}(\mathfrak{g})$, which is a commutative algebra generated by Casimir elements $C_1, \dots, C_r$.

Applying Theorem 5.2, we obtain for every highest weight $\lambda$ a corresponding element $Z_\lambda \in \mathcal{U}(\mathfrak{g})$ such that $Z_\lambda.v$ is the projection of $v$ onto the isotypic component of corresponding to the irreducible representation with highest weight $\lambda$. The projectors $Z_\lambda$ are obtained as polynomial expressions $Z_\lambda = f_\lambda(C_1, \dots, C_r)$ with $\deg f_\lambda$ is at most $p-1$, where $p$ is the diversity of $V$ as a representation of $\mathfrak{g}$. It follows that if the Casimir elements $C_1, \dots, C_r$ have length at most $\ell$, then the length of $Z_\lambda$ is at most $\ell(p-1)$.

For the Lie algebra $\mathfrak{gl}_k$ Casimir elements have length at most $k$. Since every highest weight is a weight, the diversity of $\mathbb{C}[\mathbb{C}[x_1, \dots, x_k]_d]_\delta$ is bounded by $(d\delta)^k$ as in the previous section.

▶ **Corollary 5.6.** *For each partition $\lambda \vdash_k \delta d$ there is an element $Z_\lambda \in \mathcal{U}(\mathfrak{gl}_k)$ of length at most $k(\delta d)^k$ which acts on $\mathbb{C}[\mathbb{C}[x_1, \dots, x_k]_d]_\delta$ as the projector onto the isotypic component corresponding to $\lambda$.*

### 5.1.3 Highest weight projectors

Since the highest weight subspace of weight $\lambda$ is exactly the weight $\lambda$ subspace of the isotypic component corresponding to the irreducible representation $V_\lambda$, the projector onto the highest weight subspace is the composition $H_\lambda Z_\lambda$ of the projectors constructed in the previous sections.

▶ **Corollary 5.7.** *For each partition $\lambda \vdash_k \delta d$ there is an element $X_\lambda = H_\lambda Z_\lambda \in \mathcal{U}(\mathfrak{gl}_k)$ of length at most $(k+1)(\delta d)^k$ which acts on $\mathbb{C}[\mathbb{C}[x_1, \dots, x_k]_d]_\delta$ as the projector onto the highest weight space of weight $\lambda$.*

### 5.1.4 GZ-isotypic projectors

Let $V$ be a representation of a complex reductive Lie group $G$ with the corresponding Lie algebra $\mathfrak{g}$ admitting a Gelfand–Tsetlin chain $\mathfrak{g}_1 \subset \mathfrak{g}_2 \subset \cdots \subset \mathfrak{g}_k = \mathfrak{g}$. As discussed in Section 4.5, under the action of the Gelfand–Tsetlin algebra $\mathcal{GZ}(\mathfrak{g})$ the representation $V$

decomposes into GZ-isotypic components indexed by Gelfand–Tsetlin patterns, and Theorem 5.2 can be used to construct for every Gelfand–Tsetlin pattern $T$ an element $Y_T \in \mathcal{U}(\mathfrak{g})$ such that $Y_T.x$ is the projection of $x$ onto the $T$-isotypic subspace of $V$. The length of $Y_T$ is bounded in terms of the lengths of Casimir elements of $\mathfrak{g}_\ell$ and the number of Gelfand–Tsetlin patterns appearing in the decomposition of $V$.

Consider the case of $\mathfrak{gl}_k$ acting on $\mathbb{C}[\mathbb{C}[x_1,\dots,x_k]_d]_\delta$. For the chain $\mathfrak{gl}_1 \subset \mathfrak{gl}_2 \subset \cdots \subset \mathfrak{gl}_k$, Casimir elements have length at most $k$. Every irreducible subrepresentation of $\mathbb{C}[\mathbb{C}[x_1,\dots,x_k]_d]_\delta$ corresponds to a partition $\lambda \vdash_k d\delta$. Gelfand–Tsetlin patterns are in bijective correspondence with semistandard Young tableaux with $\delta d$ cells and $k$ rows. Each row of a tableau is an ordered tuple of integers from $\{1,\dots,k\}$, so there are at most $\binom{d\delta+k-1}{k-1} \le (d\delta)^k$ tuples that can serve as rows of a tableau, and therefore at most $(d\delta)^{k^2}$ possible semistandard tableaux.

▶ **Corollary 5.8.** *For each semistandard tableau $T$ of shape $\lambda \vdash_k \delta d$ there is an element $Y_T \in \mathcal{U}(\mathfrak{gl}_k)$ of length at most $k(\delta d)^{k^2}$ which acts on $\mathbb{C}[\mathbb{C}[x_1,\dots,x_k]_d]_\delta$ as the projector onto the $T$-isotypic space corresponding to $T$.*

## 5.2 Circuit transformations

We now combine the bounds on the lengths of the projectors with the PBW theorem (Theorem 4.9) to prove Theorem 1.1. Pick any ordered basis $(X_1,\dots,X_{k^2})$ for $\mathfrak{gl}_k$, such as $\{E_{i,j}\}_{i,j}$ with $(i,j)$ ordered lexicographically.

We first introduce some notation. Set $K \coloneqq k^2$. For $\boldsymbol{i} = (i_1,\dots,i_K) \in \mathbb{N}^K$ we will use the shorthand notation $X^{\boldsymbol{i}} \coloneqq X_1^{i_1}\cdots X_K^{i_K} \in \mathcal{U}(\mathfrak{gl}_k)$. Let $\boldsymbol{i}' \le \boldsymbol{i}$ denote that $\boldsymbol{i}' = (i'_1,\dots,i'_K) \in \mathbb{N}^K$ with $i'_r \le i_r$ for all $r \in [K]$. Write $\boldsymbol{i} - \boldsymbol{i}' \coloneqq (i_1 - i'_1,\dots,i_K - i'_K)$ and $\binom{\boldsymbol{i}}{\boldsymbol{i}'} \coloneqq \binom{i_1}{i'_1}\cdots\binom{i_K}{i'_K}$. Lastly, define $|\boldsymbol{i}| \coloneqq \sum_r i_r$ and $\mathbb{N}^K_{\le L} \coloneqq \{\boldsymbol{i} \in \mathbb{N}^K \mid |\boldsymbol{i}| \le L\}$.

By Claim 4.2, any element $X \in \mathcal{U}(\mathfrak{gl}_k)$ acts as a derivation on the space of metapolynomials $\mathbb{C}[\mathbb{C}[x_1,\dots,x_k]_d] = \bigoplus_{\delta\ge 0}\mathbb{C}[\mathbb{C}[x_1,\dots,x_k]_d]_\delta$.

▶ **Lemma 5.9.** *Suppose $\Phi,\Gamma \in \mathbb{C}[\mathbb{C}[x_1,\dots,x_k]_d]$ are metapolynomials. Then*

1. $X^{\boldsymbol{i}}.(\Phi+\Gamma) = X^{\boldsymbol{i}}.\Phi + X^{\boldsymbol{i}}.\Gamma$
2. $X^{\boldsymbol{i}}.(\Phi\Gamma) = \displaystyle\sum_{\boldsymbol{i}' :\, \boldsymbol{i}' \le \boldsymbol{i}} \binom{\boldsymbol{i}}{\boldsymbol{i}'} \big(X^{\boldsymbol{i}'}.\Phi\big)\big(X^{\boldsymbol{i}-\boldsymbol{i}'}.\Gamma\big)$

**Proof.** Item 1 follows directly from linearity. Since $\mathfrak{gl}_k$ acts by derivations (see Claim 4.2) we have $X.(fg) = (X.f)g + f(X.g)$. For $i \in \mathbb{N}$ a straightforward induction argument shows $X^i.(fg) = \sum_{i' :\, i' \le i} \binom{i}{i'}\big(X^{i'}.f\big)\big(X^{i-i'}.g\big)$. Then Item 2 follows. ◀

▶ **Theorem 5.10.** *Let $P \in \mathcal{U}(\mathfrak{gl}_k)$ be an element of length $L$. Suppose there exists an arithmetic circuit $C$ of size $s$ computing a metapolynomial $\Delta\colon \mathbb{C}[x_1,\dots,x_k]_d \to \mathbb{C}$. Then there exists an arithmetic circuit $C'$ of size $O\big(sL^{2k^2}\big)$ computing $P.\Delta$.*

**Proof.** Again set $K \coloneqq k^2$. Construct $C'$ from $C$ using the following recursive procedure:

1. Add for every input gate $\Theta$ in $C$ the input gates $X^{\boldsymbol{i}}.\Theta$ for $\boldsymbol{i} \in \mathbb{N}^K_{\le L}$.
2. Add for every addition gate $\Theta = \Phi + \Gamma$ in $C$ the addition gates computing $X^{\boldsymbol{i}}.\Theta$ for $\boldsymbol{i} \in \mathbb{N}^K_{\le L}$, using the expression in Lemma 5.9.
3. Add for every product gate $\Theta = \Phi\Gamma$ in $C$:
   - product gates computing $\big(X^{\boldsymbol{i}'}.\Phi\big)\big(X^{\boldsymbol{i}-\boldsymbol{i}'}.\Gamma\big)$ for $\boldsymbol{i} \in \mathbb{N}^K_{\le L}$ and $\boldsymbol{i}' \le \boldsymbol{i}$,
   - addition gates computing $X^{\boldsymbol{i}}.\Theta$ for $\boldsymbol{i} \in \mathbb{N}^K_{\le L}$, using the expression in Lemma 5.9.

4. The resulting circuit computes $X^{\boldsymbol{i}}.\Delta$ for $\boldsymbol{i} \in \mathbb{N}^K_{\leq L}$. By the PBW theorem (Theorem 4.9), there exists scalars $\beta_{\boldsymbol{i}} \in \mathbb{C}$ such that

$$P = \sum_{\boldsymbol{i} \in \mathbb{N}^K_{\leq L}} \beta_{\boldsymbol{i}} X^{\boldsymbol{i}}.$$

Add to $C'$ the addition gates computing $P.\Delta$.

The largest blow-up occurs in step 3. Let $A \coloneqq \{(\boldsymbol{i}, \boldsymbol{i}') \mid \boldsymbol{i} \in \mathbb{N}^K_{\leq L}, \boldsymbol{i}' \leq \boldsymbol{i}\}$. We added for every product gate in $C$ exactly $|A|$ product gates computing $\left(X^{\boldsymbol{i}'}.f\right)\left(X^{\boldsymbol{i}-\boldsymbol{i}'}.g\right)$ for $(\boldsymbol{i}, \boldsymbol{i}') \in A$, and $|A| - 1$ addition gates computing $X^{\boldsymbol{i}}.(fg) = \sum_{(\boldsymbol{i},\boldsymbol{i}') \in A} \binom{\boldsymbol{i}}{\boldsymbol{i}'} \left(X^{\boldsymbol{i}'}.f\right)\left(X^{\boldsymbol{i}-\boldsymbol{i}'}.g\right)$.

The blow-up of the other steps is bounded above by $|A|$, and the total blow-up is therefore upper bounded by $2|A| - 1$. We now prove that $|A| = \binom{L+2K}{2K}$.

Let $(\boldsymbol{i}, \boldsymbol{i}') \in A$. By definition, $i_1 + \cdots + i_K \leq L$, and $(i_j - i'_j) \geq 0$, for all $j \in [K]$, with $i'_j \geq 0$. Let $a_j \coloneqq (i_j - i'_j)$, and $b_j \coloneqq i'_j$. We want to find the number of $(\boldsymbol{a}, \boldsymbol{b}) \in \mathbb{N}^{2K}$ such that $\sum_{j=1}^{2K}(a_j + b_j) \leq L$. So the number of pairs $(\boldsymbol{a}, \boldsymbol{b}) \in \mathbb{N}^{2K}$ such that $\sum_{j=1}^{2K}(a_j + b_j) = t$ is $\binom{t+2K-1}{2K-1}$. Therefore, by the hockey-stick identity,

$$|A| \;=\; \sum_{t=0}^{L} \binom{t+2K-1}{2K-1} \;=\; \binom{L+2K}{2K}.$$

The above equality readily gives us the upper bound of $|A| = O(L^{2K})$, which gives the desired complexity. ◀

**Proof of Theorem 1.1.** Let $\Delta$ be a metapolynomial of format $(d, \delta, n)$ admitting an algebraic circuit of size $s$. For each case of Theorem 1.1, we have a corresponding projector:

1. For every weight $\mu$, Corollary 5.5 provides an element $H_\mu \in \mathcal{U}(\mathfrak{gl}_k)$ of length at most $O((\delta d)^k)$ acting as projection on the weight space of weight $\mu$. Theorem 5.10 guarantees that $H_\mu.\Delta$ admits a circuit of size at most $O(s((\delta d)^k)^{2k^2}) = O(s(\delta d)^{2k^3})$.
2. For every partition $\lambda$, Corollary 5.6 provides an element $Z_\mu \in \mathcal{U}(\mathfrak{gl}_k)$ of length at most $O(k(\delta d)^k)$ acting as projection on the isotypic component of type $\lambda$. Theorem 5.10 guarantees that $Z_\lambda.\Delta$ admits a circuit of size at most $O(s(k(\delta d)^k)^{2k^2}) = O(sk^{2k^2}(\delta d)^{2k^3})$.
3. For every partition $\lambda$, Corollary 5.7 provides an element $X_\lambda \in \mathcal{U}(\mathfrak{gl}_k)$ of length at most $O((k+1)(\delta d)^k)$ acting as projection on the highest weight space of weight $\lambda$. Theorem 5.10 guarantees that $X_\lambda.\Delta$ admits a circuit of size at most $O(s((k+1)(\delta d)^k)^{2k^2}) = O(s(k+1)^{2k^2}(\delta d)^{2k^3})$.
4. For every semistandard tableaux $T$ of shape $\lambda \vdash d\delta$, Corollary 5.8 provides an element $Y_\lambda \in \mathcal{U}(\mathfrak{gl}_k)$ of length at most $O(k(\delta d)^{k^2})$ acting as projection on $T$-isotypic space. Theorem 5.10 guarantees that $Y_T.\Delta$ admits a circuit of size at most $O(s(k(\delta d)^{k^2})^{2k^2}) = O(sk^{2k^2}(\delta d)^{2k^4})$. ◀

## References


1 S. Aaronson and A. Wigderson. Algebrization: A new barrier in complexity theory. *ACM Transactions on Computation Theory (TOCT)*, 1(1):1–54, 2009. doi:10.1145/1490270.1490272.

2 A. Abdesselam and J. Chipalkatti. Brill–Gordan loci, transvectants and an analogue of the Foulkes conjecture. *Advances in Mathematics*, 208(2):491–520, 2007. doi:10.1016/j.aim.2006.03.003.

3 A. Abdesselam, C. Ikenmeyer, and G. Royle. 16,051 formulas for Ottaviani's invariant of cubic threefolds. *J. Algebra*, 447:649–663, 2016. doi:10.1016/j.jalgebra.2015.11.009.

4 G. E. Baird and L. C. Biedenharn. On the representations of the semisimple Lie groups. II. *J. Mathematical Phys.*, 4:1449–1466, 1963. doi:10.1063/1.1703926.

5 T. Baker, J. Gill, and R. Solovay. Relativizations of the $\mathcal{P} =? \mathcal{NP}$ question. *SIAM Journal on computing*, 4(4):431–442, 1975. doi:10.1137/0204037.

6 F. Berdjis. A criterion for completeness of Casimir operators. *J. Math. Phys.*, 22:1851–1856, 1981. doi:10.1063/1.525156.

7 D. Bini, M. Capovani, G. Lotti, and F. Romani. $O(n^{2.7799})$ complexity for $n \times n$ approximate matrix multiplication. *Inform. Process. Lett.*, 8(5):234–235, 1979. doi:10.1016/0020-0190(79)90113-3.

8 M. Bläser, M. Christandl, and J. Zuiddam. The border support rank of two-by-two matrix multiplication is seven. *Chicago Journal OF Theoretical Computer Science*, 5:1–16, 2018.

9 M. Bläser, J. Dörfler, and C. Ikenmeyer. On the Complexity of Evaluating Highest Weight Vectors. *36th Computational Complexity Conference (CCC 2021)*, 200:29:1–29:36, 2021. doi:10.4230/LIPIcs.CCC.2021.29.

10 N. Bourbaki. *Lie groups and Lie algebras. Chapters 7–9*. Elements of Mathematics (Berlin). Springer-Verlag, Berlin, 2005. Translated from the 1975 and 1982 French originals by Andrew Pressley.

11 P. Breiding, R. Hodges, C. Ikenmeyer, and M. Michałek. Equations for GL invariant families of polynomials. *Vietnam Journal of Mathematics*, 50(2):545–556, 2022. doi:10.1007/s10013-022-00549-4.

12 P. Bürgisser. *Completeness and reduction in algebraic complexity theory*, volume 7. Springer Science & Business Media, 2000. doi:10.1007/978-3-662-04179-6.

13 P. Bürgisser. Completeness classes in algebraic complexity theory. *arXiv:2406.06217*, 2024.

14 P. Bürgisser, M. Clausen, and M. A. Shokrollahi. *Algebraic complexity theory*, volume 315. Springer Science & Business Media, 2013.

15 P. Bürgisser and C. Ikenmeyer. Geometric complexity theory and tensor rank. In *Proceedings of the 43rd Annual ACM Symp. on Th. of Comp.*, pages 509–518, New York, 2011. STOC '11, ACM. doi:10.1145/1993636.1993704.

16 P. Bürgisser and C. Ikenmeyer. Explicit Lower Bounds via Geometric Complexity Theory. In *Proceedings of the 45th Annual ACM Symp. on Th. of Comp.*, pages 141–150. ACM, 2013. doi:10.1145/2488608.2488627.

17 P. Bürgisser and C. Ikenmeyer. Fundamental invariants of orbit closures. *J. Algebra*, 477:390–434, 2017. doi:10.1016/j.jalgebra.2016.12.035.

18 P. Bürgisser, C. Ikenmeyer, and G. Panova. No occurrence obstructions in geometric complexity theory. *J. Amer. Math. Soc.*, 32(1):163–193, 2019. doi:10.1090/jams/908.

19 P. Bürgisser, J. M. 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. doi:10.1137/090765328.

20 A. Cayley. On the theory of linear transformations. *Cambridge Math. J.*, iv:193–209, 1845.

21 P. Chatterjee, M. Kumar, C. Ramya, R. Saptharishi, and A. Tengse. On the Existence of Algebraically Natural Proofs. In *2020 IEEE 61st Annual Symp. on Found. of Comp. Sc. (FOCS)*, pages 870–880, 2020. Full version at arXiv:2004.14147. doi:10.1109/FOCS46700.2020.00085.

22 X. Chen, N. Kayal, and A. Wigderson. Partial derivatives in arithmetic complexity and beyond. *Found. Trends Theor. Comput. Sci.*, 6(1–2):1–138, 2011. doi:10.1561/0400000043.

23 M.-W. Cheung, C. Ikenmeyer, and S. Mkrtchyan. Symmetrizing tableaux and the 5th case of the Foulkes conjecture. *J. Symbolic Comp.*, 80:833–843, 2017. doi:10.1016/j.jsc.2016.09.002.

24 J. Dörfler, C. Ikenmeyer, and G. Panova. On geometric complexity theory: Multiplicity obstructions are stronger than occurrence obstructions. *SIAM Journal on Applied Algebra and Geometry*, 4(2):354–376, 2020. doi:10.4230/LIPIcs.ICALP.2019.51.

25 P. Dutta, P. Dwivedi, and N. Saxena. Demystifying the border of depth-3 algebraic circuits. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 92–103. IEEE, 2021. doi:10.1109/FOCS52979.2021.00018.

26 K. Efremenko, A. Garg, R. Oliveira, and A. Wigderson. Barriers for Rank Methods in Arithmetic Complexity. In *9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*, volume 94 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 1:1–1:19, Dagstuhl, Germany, 2018. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2018.1.

27 M. A. Forbes. *Polynomial identity testing of read-once oblivious algebraic branching programs*. PhD thesis, Massachusetts Institute of Technology, 2014.

28 M. A. Forbes. Deterministic divisibility testing via shifted partial derivatives. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 451–465. IEEE, 2015. doi:10.1109/FOCS.2015.35.

29 M. A. Forbes. Some concrete questions on the border complexity of polynomials. Presentation given at the Workshop on Algebraic Complexity Theory WACT 2016 in Tel Aviv, 2016. URL: https://www.youtube.com/watch?v=1HMogQIHT6Q.

30 M. A. Forbes. Low-Depth Algebraic Circuit Lower Bounds over Any Field. In *39th Computational Complexity Conference, CCC 2024, July 22-25, 2024*, volume 300 of *LIPIcs*, pages 31:1–31:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICS.CCC.2024.31.

31 M. A. Forbes, A. Shpilka, and B. L. Volk. Succinct hitting sets and barriers to proving lower bounds for algebraic circuits. *Theory of Computing*, 14(1):1–45, 2018. doi:10.4086/toc.2018.v014a018.

32 W. Fulton and J. Harris. *Representation theory: a first course*, volume 129 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1991.

33 M. Gałązka. Vector bundles give equations of cactus varieties. *Lin. Alg. Appl.*, 521:254–262, 2017. doi:10.1016/j.laa.2016.12.005.

34 A. Garg, C. Ikenmeyer, V. Makam, R. Oliveira, M. Walter, and A. Wigderson. Search problems in algebraic complexity, GCT, and hardness of generators for invariant rings. In *35th Computational Complexity Conference*, page 1, 2020.

35 I. M. Gelfand and M. L. Tsetlin. Finite-dimensional representations of the group of unimodular matrices. *Dokl. Akad. Nauk SSSR*, 71(8):825, 1950.

36 F. Gesmundo, P. Ghosal, C. Ikenmeyer, and V. Lysikov. Degree-Restricted Strength Decompositions and Algebraic Branching Programs. In *42nd IARCS Ann. Conf. Found. Soft. Tech. and TCS (FSTTCS 2022)*, volume 250 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 20:1–20:15, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.FSTTCS.2022.20.

37 F. Gesmundo, Y.-I. Han, and B. Lovitz. Linear preservers of secant varieties and other varieties of tensors. *J. Symbolic Comp.*, 131:102449, 2025. doi:10.1016/j.jsc.2025.102449.

38 F. Gesmundo, C. Ikenmeyer, and G. Panova. Geometric complexity theory and matrix powering. *Diff. Geom. Appl.*, 55:106–127, 2017. doi:10.1016/j.difgeo.2017.07.001.

39 J. A. Grochow. Unifying known lower bounds via geometric complexity theory. *computational complexity*, 24:393–475, 2015. doi:10.1007/s00037-015-0103-x.

40 J. A. Grochow, M. Kumar, M. Saks, and S. Saraf. Towards an algebraic natural proofs barrier via polynomial identity testing. *arXiv*, 2017. arXiv:1701.01717.

41 A. Gupta, P. Kamath, N. Kayal, and R. Saptharishi. Arithmetic circuits: A chasm at depth three. *Electronic Colloquium on Computational Complexity (ECCC)*, 20:26, 2013. doi:10.1109/FOCS.2013.68.

42 B. C. Hall. *Lie Groups, Lie Algebras, and Representations: An Elementary Introduction*, volume 222 of *Graduate Texts in Mathematics*. Springer, New York, 2010.

43 J. Hauenstein, C. Ikenmeyer, and J. M. Landsberg. Equations for lower bounds on border rank. *Experimental Mathematics*, 22(4):372–383, 2013. doi:10.1080/10586458.2013.825892.

44 J. E. Humphreys. *Introduction to Lie algebras and representation theory*, volume 9 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1978.

45 J. E. Humphreys. *Linear algebraic groups*, volume 21. Springer Science & Business Media, 2012.

46 C. Ikenmeyer and U. Kandasamy. Implementing geometric complexity theory: On the separation of orbit closures via symmetries. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 713–726, 2020. doi:10.1145/3357713.3384257.

47 C. Ikenmeyer and G. Panova. Rectangular Kronecker coefficients and plethysms in geometric complexity theory. *Adv. Math.*, 319:40–66, 2017. doi:10.1016/j.aim.2017.08.024.

48 C. Ikenmeyer and A. Sanyal. A note on VNP-completeness and border complexity. *Information Processing Letters*, 176:106243, 2022. doi:10.1016/j.ipl.2021.106243.

49 A. G. Khovanskiĭ. *Fewnomials*, volume 88. American Mathematical Soc., 1991.

50 M. Kumar. A quadratic lower bound for homogeneous algebraic branching programs. *computational complexity*, 28:409–435, 2019. doi:10.1007/s00037-019-00186-3.

51 M. Kumar, C. Ramya, R. Saptharishi, and A. Tengse. If VNP Is Hard, Then so Are Equations for It. In *39th Int. Symp. on Th. Asp. of Comp. Sc. (STACS 2022)*, volume 219 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 44:1–44:13, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.STACS.2022.44.

52 M. Kumar and B. L. Volk. A lower bound on determinantal complexity. *Comput. Complex.*, 31(2):12, 2022. doi:10.1007/S00037-022-00228-3.

53 J. M. Landsberg, L. Manivel, and N. Ressayre. Hypersurfaces with degenerate duals and the Geometric Complexity Theory program. *Comment. Math. Helv.*, 88(2):469–484, 2013. doi:10.4171/CMH/292.

54 J. M. Landsberg and G. Ottaviani. New lower bounds for the border rank of matrix multiplication. *Th. of Comp.*, 11(11):285–298, 2015. doi:10.4086/toc.2015.v011a011.

55 N. Limaye, Srinivasan S., and S. Tavenas. Superpolynomial lower bounds against low-depth algebraic circuits. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 804–814. IEEE, 2021. doi:10.1109/FOCS52979.2021.00083.

56 I. G. Macdonald. *Symmetric functions and Hall polynomials*. Oxford Mathematical Monographs. The Clarendon Press Oxford University Press, New York, second edition, 1995.

57 M. Mahajan. Algebraic complexity classes. *Perspectives in Computational Complexity: The Somenath Biswas Anniversary Volume*, pages 51–75, 2014. doi:10.1007/978-3-319-05446-9_4.

58 A. I. Molev. Gelfand-Tsetlin bases for classical Lie algebras. In *Handbook of algebra. Vol. 4*, volume 4 of *Handb. Algebr.*, pages 109–170. Elsevier/North-Holland, Amsterdam, 2006. doi:10.1016/S1570-7954(06)80006-9.

59 K. D. Mulmuley. The GCT program toward the P vs. NP problem. *Communications of the ACM*, 55(6):98–107, 2012. doi:10.1145/2184319.2184341.

60 K. D. 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. doi:10.1137/S009753970038715X.

61 K. D. Mulmuley and M. Sohoni. Geometric complexity theory II: towards explicit obstructions for embeddings among class varieties. *SIAM Journal on Computing*, 38(3):1175–1206, 2008. doi:10.1137/080718115.

62 N. Nisan and A. Wigderson. Lower bounds on arithmetic circuits via partial derivatives. *Computational complexity*, 6:217–234, 1996. doi:10.1007/BF01294256.

63 L. Oeding and S.V Sam. Equations for the fifth secant variety of Segre products of projective spaces. *Exp. Math.*, 25(1):94–99, 2016. doi:10.1080/10586458.2015.1037872.

64 A. M. Perelomov and V. S. Popov. Casimir Operators for Semisimple Lie Groups. *Mathematics of the USSR-Izvestiya*, 2(6):1313, December 1968. doi:10.1070/IM1968v002n06ABEH000731.

65 G. Racah. Sulla caratterizzazione delle rappresentazioni irriducibili dei gruppi semisemplici di Lie. *Atti Accad. Naz. Lincei Rend. Cl. Sci. Fis. Mat. Nat. (8)*, 8:108–112, 1950.

66 A. A. Razborov and S. Rudich. Natural proofs. In *Proc. of the 26th Ann. ACM Symp. Th. of Comp.*, pages 204–213, 1994.

67 K. W. Regan. Understanding the Mulmuley–Sohoni approach to P vs. NP. *Bulletin of the EATCS*, 78:86–99, 2002.

68 R. Saptharishi. A survey of lower bounds in arithmetic circuit complexity. `https://github.com/dasarpmar/lowerbounds-survey/`, 2021. Version 9.0.3.

69 V. Strassen. Rank and optimal computation of generic tensors. *Lin. Alg. Appl.*, 52/53:645–685, 1983. `doi:10.1016/0024-3795(83)80041-X`.

70 J. J. Sylvester. On the principles of the calculus of forms. *Cambridge and Dublin Math. J.*, 7:52–97, 1852.

71 L. G. Valiant. Completeness classes in algebra. In *Proceedings of the 11th ACM Symp. on Th. of Comp.*, STOC '79, pages 249–261, New York, 1979. ACM. `doi:10.1145/800135.804419`.

72 D. P. Žhelobenko. The classical groups. Spectral analysis of their finite-dimensional representations. *Uspehi Mat. Nauk*, 17(1(103)):27–120, 1962. `doi:10.1070/RM1962v017n01ABEH001123`.


## A Proofs on algebraic natural proofs

In this section we prove Theorem 2.2.

▶ **Definition A.1.** *We call a sequence* $r\colon \mathbb{N} \to \mathbb{N}$ strongly superpolynomial *if* $r(n)$ *dominates every polynomial sequence eventually, that is, for every polynomial* $p$ *there exists* $N$ *such that* $r(n) > p(n)$ *for all* $n \geq N$.

▶ **Lemma A.2.** $(\Delta_n) \in I(\mathcal{C})$ *if and only if there exists a strongly superpolynomial sequence* $r$ *such that* $\Delta_n(X_{n,r(n)}) = 0$.

**Proof.** Clearly, if $\Delta_n(X_{n,r(n)}) = 0$, then $\Delta_n(X_{n,p(n)})$ vanishes eventually for every polynomially bounded sequence $p$, since $p(n) \leq r(n)$ eventually.

On the other hand, if $(\Delta_n)$ is a sequence of metapolynomials in $I(\mathcal{C})$, we can define $r(n)$ as $r(n) := \max\{r \mid \Delta_n(X_{n,r}) = 0\}$. For every polynomially bounded sequence $p(n)$ we have that $\Delta_n(X_{n,p(n)})$ vanishes eventually, so $r(n) \geq p(n)$ eventually. ◀

**Proof of Theorem 2.2.** We first observe we may assume $\deg(f_n) = n$. Indeed, for every p-family $(f_n)$ with the property that $\deg(f_n)$ is strictly increasing, let $(f'_n)$ be the sequence defined by

$$f'_m = \begin{cases} f_n & \text{if } m = \deg(f_n) \text{ for some } n \\ 0 & \text{otherwise} \end{cases}$$

Since $(f_n)$ is a p-family, the sequence $\deg(f_n)$ is polynomially bounded, hence $(f'_n)$ is a p-family as well. Moreover $(f_n) \in \mathcal{C}$ if and only if $(f_n) \in \mathcal{C}$, and similarly for $\bar{\mathcal{C}}$. Finally, for a sequence of metapolynomials $(\Delta_n)$ of format $(*, n, *)$, we have $\Delta_{\deg(f_n)}(f_n) = 0$ eventually if and only if $\Delta_n(f'_n) = 0$ eventually. In particular, proving the statement for $(f'_n)$ is equivalent to proving the statement for $(f_n)$. Since $\deg(f'_n) = n$, this shows we may assume $\deg(f_n) = n$.

Suppose $(f_n) \in \overline{\mathcal{C}}$, that is, the sequence $p(n) = \underline{c}(f_n)$ is polynomially bounded. For every sequence of metapolynomials $(\Delta_n) \in I(\mathcal{C})$ we have that $\Delta_n(X_{n,p(n)})$ vanishes eventually. Since the metapolynomials $\Delta_n$ are continuous, $\Delta_n(\overline{X_{n,p(n)}})$ vanishes eventually, so $\Delta_n(f_n) = 0$ eventually.

Conversely, suppose $(f_n) \notin \overline{\mathcal{C}}$. Therefore $r(n) = \underline{c}(f_n)$ is not polynomially bounded. Let $\theta(n) = \log_n r(n)$, $\hat{\theta}(n) = \max\{\theta(m) \mid m \leq n\}$, and $R(n) = \lceil n^{\hat{\theta}(n)} \rceil$. Note that $R(n)$ is monotonically increasing.

The sequence $R(n)$ is strongly superpolynomial. Indeed, since $r(n)$ is not polynomially bounded, for every $k$ there exists $n_0$ such that $\theta(n_0) \geq k$ and, therefore, $\hat{\theta}(n) \geq k$ for all $n \geq n_0$. It follows that for every $k$ we have $R(n) \geq n^k$ eventually.

For all $\Theta$, if $n_\Theta$ is the minimal value such that $\theta(n_\Theta) \geq \Theta$, then $\hat{\theta}(n_\Theta) = \theta(n_\Theta)$. Since $\theta$ is unbounded, it follows that $\theta$ and $\hat{\theta}$ coincide on an infinite subset of $\mathbb{N}$. On this subset we also have $R(n) = r(n)$.

Consider a metapolynomial sequence $\Delta_n$ such that $\Delta_n(\overline{X_{n,R(n)-1}}) = 0$, and if $r(n) = R(n)$, then $\Delta_n$ is chosen so that $\Delta_n(f_n) \neq 0$ (this is possible because $\underline{c}(f_n) = r(n) = R(n) > R(n) - 1$). Since $R(n)$ is strongly superpolynomial, $R(n) - 1$ is strongly superpolynomial, and $(\Delta_n) \in I(\mathcal{C})$ by Lemma A.2, and since $R(n) = r(n)$ infinitely often, $\Delta_n(f_n) \neq 0$ infinitely often. ◀

## B PBW, proof that monomials span the space

In this section, we give an elementary proof of the implication in Theorem 4.9 that we use. We include this standard result in this appendix, because it is the main result that enables our algorithmic speedup compared to a naive implementation.

Recall that the universal enveloping algebra $\mathcal{U}(\mathfrak{g})$ is the quotient of the tensor algebra $\mathcal{T}(\mathfrak{g})$ by the ideal $\mathcal{I}(\mathfrak{g}) = \{X \otimes Y - Y \otimes X - [X, Y] : X, Y \in \mathfrak{g}\}$. The quotient induces a filtration $\mathcal{U}(\mathfrak{g})_{\leq m} \subseteq \mathcal{U}(\mathfrak{g})_{\leq (m+1)}$ and the elements of $\mathcal{U}(\mathfrak{g})_{\leq m}$ are represented by noncommutative polynomials of degree at most $m$ in the elements of $\mathfrak{g}$.

▶ **Theorem B.1** (One direction of the PBW Theorem). *Let $X_1, \ldots, X_r$ be a basis for a Lie algebra $\mathfrak{g}$. Then*

$$\mathcal{B}_m \coloneqq \{X_{i_1} X_{i_2} \cdots X_{i_\ell} \mid i_1 \leq i_2 \leq \cdots \leq i_\ell \text{ and } \ell \leq m\}$$

*spans $\mathcal{U}(\mathfrak{g})_{\leq m}$.*

**Proof.** It is clear that $\mathcal{T}(\mathfrak{g})_m$ is spanned by all elements $X_{i_1} \otimes \cdots \otimes X_{i_m}$; hence it suffices to prove that any such element is equivalent, modulo $\mathcal{I}(\mathfrak{g})$, to a linear combination of elements of $\mathcal{B}_m$. We prove this statement by induction on $m$.

Let $X_{j_1} \otimes \cdots \otimes X_{j_m} \in \mathcal{T}(\mathfrak{g})_m$. For every $a = 1, \ldots, m-1$, we have $X_{j_a} \otimes X_{j_{a+1}} - X_{j_{a+1}} \otimes X_{j_a} - [X_{j_a}, X_{j_{a+1}}] \in \mathcal{I}(\mathfrak{g})$; in particular $X_{j_a} \otimes X_{j_{a+1}} = X_{j_{a+1}} \otimes X_{j_a} + Y$ where $Y$ is a linear combination of elements $\mathcal{I}(\mathfrak{g})$ and of $\mathcal{T}(\mathfrak{g})_{\leq 1}$. We obtain

$$X_{j_1} \otimes \cdots \otimes X_{j_\ell} = X_{j_1} \otimes \cdots \otimes X_{j_{a-1}} \otimes X_{j_{a+1}} \otimes X_{j_a} \otimes X_{j_{a+2}} \otimes \cdots \otimes X_{j_\ell} + Y \tag{18}$$

where $Y$ is a linear combination of elements of $\mathcal{I}(\mathfrak{g})$ and of $\mathcal{T}(\mathfrak{g})_{\leq m-1}$.

Let $\pi \in \mathfrak{S}_m$ be a permutation such that $(j_{\pi(1)} \cdots j_{\pi(m)})$ is nondecreasing. Write $\pi$ as a composition of elementary transpositions $\pi = \sigma_s \cdots \sigma_1$, where $\sigma_q = (a, a+1)$ for some $a = 1, \ldots, m-1$. Applying Equation (18) $s$ times we obtain

$$X_{j_1} \otimes \cdots \otimes X_{j_m} = X_{j_{\pi(1)}} \otimes \cdots \otimes X_{j_{\pi(m)}} + U$$

where $U \in \mathcal{I}(\mathfrak{g}) + \mathcal{U}_{\leq (m-1)}$. Passing to the quotient modulo $\mathcal{I}(\mathfrak{g})$, we obtain the following identity in $\mathcal{U}(\mathfrak{g})_{\leq m}$:

$$X_{j_1} \cdots X_{j_m} = X_{j_{\pi(1)}} \cdots X_{j_{\pi(m)}} + U$$

with $U \in \mathcal{U}(\mathfrak{g})_{\leq (m-1)}$. Apply the induction hypothesis to $U$ to conclude. ◀