COMPOSITIO MATHEMATICA

A tensor restriction theorem over finite fields

Andreas Blatter, Jan Draisma and Filip Rupniewski

Compositio Math. **161** (2025), 1021–1053.

doi: 10.1112/S0010437X25007067

FOUNDATION COMPOSITIO MATHEMATICA

LONDON MATHEMATICAL SOCIETY
EST. 1865

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

Compositio Math. 161 (2025) 1021–1053
doi:10.1112/S0010437X25007067

A tensor restriction theorem over finite fields

Andreas Blatter, Jan Draisma and Filip Rupniewski

Abstract

Restriction is a natural quasi-order on d-way tensors. We establish a remarkable aspect of this quasi-order in the case of tensors over a fixed finite field; namely, that it is a well-quasi-order: it admits no infinite antichains and no infinite strictly decreasing sequences. This result, reminiscent of the graph minor theorem, has important consequences for an arbitrary restriction-closed tensor property X. For instance, X admits a characterisation by finitely many forbidden restrictions and can be tested by looking at subtensors of a fixed size. Our proof involves an induction over polynomial generic representations, establishes a generalisation of the tensor restriction theorem to other such representations (e.g., homogeneous polynomials of a fixed degree), and also describes the coarse structure of any restriction-closed property.

1. Introduction and results

1.1 Tensor restriction

Let K be a finite field and let d be a natural number. This paper concerns properties of d-way tensors that are preserved under taking linear maps. For a vector space V over K we denote by $V^{\otimes d}$ the d-fold tensor product $V \otimes V \otimes \cdots \otimes V$ over K, and for a linear map $\varphi : V \to W$ we denote by $\varphi^{\otimes d} : V^{\otimes d} \to W^{\otimes d}$ the linear map determined by

$$\varphi^{\otimes d}(v_1 \otimes \cdots \otimes v_d) := \varphi(v_1) \otimes \cdots \otimes \varphi(v_d).$$

Definition 1.1.1. Let V, W be finite-dimensional vector spaces over K and let $S \in V^{\otimes d}$ and $T \in W^{\otimes d}$. We call T a *restriction* of S if there exists a linear map $\varphi : V \to W$ such that $\varphi^{\otimes d} S = T$. We then write $S \succeq T$.

The rationale for this terminology is that S can be thought of as a multilinear map $(V^*)^d \to K$, and composing this map with $(\varphi^*)^d : (W^*)^d \to (V^*)^d$ gives the multilinear map T. In particular, if φ^* is injective, so that we can use it to identify W^* with a subspace of V^*, then we can think of T as the restriction of S to the subspace $(W^*)^d$.

Remark 1.1.2. Much of the literature on tensors considers tensor products $V_1 \otimes \cdots \otimes V_d$ of different vector spaces V_i, and for restriction allows the application of distinct linear maps $\varphi_i : V_i \to W_i$ to the individual factors. The theorems that we will prove imply the corresponding theorems for this setting (see Remark 1.6.3).

Received 27 April 2023, accepted in final form 8 November 2024.
2020 Mathematics Subject Classification 15A69, 14N07 (primary), 20G05, 13E05 (secondary).
Keywords: tensors; tensor restriction; well-quasi-orders.

1.2 The tensor restriction theorem over finite fields

The relation $\succeq$ is reflexive and transitive, so it is a quasi-order on tensors over K. We will prove that this quasi-order is a well-quasi-order.

Theorem 1.2.1 (Tensor restriction theorem). *Fix a natural number d. For every $i \in \mathbb{N}$ let V_i be a finite-dimensional vector space over the finite field K and let $T_i \in V_i^{\otimes d}$. Then there exist $i < j$ such that $T_j \succeq T_i$.*

As the following example shows, the requirement that K be finite is essential.

Example 1.2.2. If $|K| = \infty$, then the statement of the theorem fails already for $d = 2$. Indeed, if $\operatorname{char} K \neq 2$, then consider the matrices

$$M_a := \begin{bmatrix} 1 & a \\ -a & 1 \end{bmatrix} \in (K^2)^{\otimes 2}$$

for a ranging through K. If $M_a \succeq M_b$, then there exists a 2×2 matrix g over K such that $gM_a g^T = M_b$. Looking at the symmetric parts of M_a and M_b, we find that $gIg^T = I$, so g is an orthogonal matrix and M_a, M_b have the same characteristic polynomial. But the characteristic polynomial of M_a equals $(t-1)^2 + a^2$, so $M_a \succeq M_b$ holds (if and) only if $a^2 = b^2$. Since $|K| = \infty$, we have found infinitely many two-way tensors that are incomparable with respect to $\succeq$. A similar construction works when $\operatorname{char} K = 2$. It is easy to see that the failure for $d = 2$ implies the failure for all larger d.

1.3 Consequences of the tensor restriction theorem

The tensor restriction theorem is reminiscent of the celebrated graph minor theorem [RS04], which says that finite graphs are well-quasi-ordered by the minor order. We are not aware of any logical dependence between these theorems, but the tensor restriction theorem has similar far-reaching consequences for tensors to those the graph minor theorem has for graphs. These consequences are best formulated using the following notion.

Definition 1.3.1. A *restriction-closed property* of d-way tensors is a property such that if a tensor S has it, and $S \succeq T$ holds, then T also has it. We can identify such a property with the data of a subset $X(V) \subseteq V^{\otimes d}$ for every finite-dimensional vector space V over K, such that if $\varphi : V \to W$ is a linear map, then $\varphi^{\otimes d}$ maps $X(V)$ into $X(W)$.

Example 1.3.2. Let $T \in W^{\otimes d}$. Then the property of *not* having T as a restriction is restriction-closed. We denote this property by $X_{\not\succeq T}$.

If X is a restriction-closed property, and T is a tensor that does not satisfy it, then we call T a *forbidden restriction* for X.

Corollary 1.3.3. *For d-way tensors over the fixed finite field K the following statements hold.*

1. *Restriction-closed properties satisfy the descending chain condition: any chain $X_1 \supseteq X_2 \supseteq \cdots$ of such properties stabilises.*
2. *Every restriction-closed property X is characterised by finitely many forbidden restrictions, i.e., we have $X = \bigcap_{i=1}^{k} X_{\not\succeq T_i}$ for some k and some tensors $T_i \in V_i^{\otimes d}$.*
3. *For every restriction-closed property X there exists a finite-dimensional vector space U such that for any V and any $T \in V^{\otimes d}$, we have $T \in X(V)$ if (and only if) $\varphi^{\otimes d} T \in X(U)$ for all linear maps $\varphi : V \to U$.*

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

(4) *For every restriction-closed property X there exists a number n_0 such that a tensor $T \in (K^n)^{\otimes d}$ satisfies X if and only if for every subset $S \subseteq [n] := \{1, \dots, n\}$ of size n_0 the subtensor of T in $(K^S)^{\otimes d}$ satisfies X.*

(5) *For every restriction-closed property X there exists a polynomial-time deterministic algorithm that on input n and a $T \in (K^n)^{\otimes d}$ decides whether T satisfies X.*

The proofs of (1), (2), and (3) are straightforward from the tensor restriction theorem, and conversely the tensor restriction theorem follows from each of these.

Proofs of (1), (2), and (3) from the tensor restriction theorem and vice versa. Assuming the tensor restriction theorem, we prove (1): whenever X_i and X_{i+1} are not the same property, there exists a tensor T_i that satisfies X_i but not X_{i+1}. For $i < j$ we then have $T_j \not\geq T_i$, and hence $X_i \neq X_{i+1}$ holds only finitely many times.

Next we prove (1) $\Rightarrow$ (2). Start with $k = 0$. While X is strictly contained in $X_k := \bigcap_{i=1}^{k} X_{\not\geq T_i}$, there is a tensor T_{k+1} that does not satisfy X but does not have any of the tensors $T_1, \dots, T_k$ as a restriction. This yields a strictly descending chain $X_0 \supsetneq X_1 \supsetneq \cdots$, which by (1) must terminate, so that X is equal to X_k for some k.

For (2) $\Rightarrow$ (3), we take for U any space of dimension at least that of all of the spaces V_i, $i = 1, \dots, k$, where $T_i \in V_i^{\otimes d}$. If $T \in V^{\otimes d}$ does not lie in $X(V)$, then it has a restriction equal to some T_i, so that $\psi^{\otimes d} T = T_i$ for some linear map $\psi : V \to V_i$. Now ψ factors via a linear map $\varphi : V \to U$, and it follows that $\varphi^{\otimes d} T \notin X(U)$.

Finally, (3) implies the tensor restriction theorem: let $T_i \in V_i^{\otimes d}$, $i = 1, 2, \dots$, and define $X_n := \bigcap_{i=1}^{n} X_{\not\geq T_i}$ and $X := \bigcap_{i=1}^{\infty} X_{\not\geq T_i}$. Let U be as in (3) for X. Then, since $X_0(U)$ is a finite set, the chain of subsets

$$X_0(U) \supseteq X_1(U) \supseteq \cdots$$

stabilises after finitely many steps: $X_n(U) = X(U)$. Then in particular T_{n+1}, which is not in $X(V_{n+1})$, is not in $X_n(V_{n+1})$, which means that it must have some T_i with $i \leq n$ as a restriction. □

Example 1.3.4. Let $d = 2$, assume $\operatorname{char} K > 2$, and let $X(V) \subseteq V \otimes V$ be the set of tensors $\{c \cdot v \otimes v \mid v \in V, c \in K\}$ (symmetric matrices of rank 1). Clearly, X is a restriction-closed property. If $T \in V \otimes V$ is not in $X(V)$, then write $T = T_1 + T_2$ with T_1 symmetric and T_2 skew-symmetric. If $T_2 \neq 0$, then T_2 has rank at least 2 and there exists a linear map $\varphi : V \to K^2$ with $\varphi^{\otimes 2} T_2 = e_1 \otimes e_2 - e_2 \otimes e_1$. It follows that $\varphi^{\otimes 2} T \notin X(K^2)$. Similarly, if T_1 has rank at least 2 , then there exists a linear map $\varphi : V \to K^2$ such that $\varphi^{\otimes 2} T_1$ has rank at least 2 , and again $\varphi^{\otimes 2} T \notin X(K^2)$. So we obtain a (finite and minimal) set of forbidden minors characterising X by picking a representative from each $\mathrm{GL}(K^2)$-orbit on elements in $(K^2 \otimes K^2) \setminus X(K^2)$. This minimal set of forbidden minors increases with $|K|$.

Note that the difference between (3) and (4) is that in (4) we only consider coordinate projections $K^n \to K^I$. The proofs of (4) and (5) are slightly more involved and deferred to § 5.1.

Remark 1.3.5. Versions of (1), (3), (4), and (5) also hold for restriction-closed tensor properties over an *infinite* field, provided that the tensor property can be expressed by polynomial equations in the tensor entries. For (1) and (3) this follows from [Dra19]. For (4) and (5) this follows from (1) and (3) and the technique in § 5.1 below.

1.4 Restriction-monotone functions

Tensor restriction plays an important role in theoretical computer science, in particular through many notions of tensor rank, of which we briefly discuss two here.

Definition 1.4.1. The *rank* $\mathrm{rk}(S)$ of $S \in V^{\otimes d}$ is the minimal r such that S can be written as

$$S = \sum_{i=1}^{r} v_{i,1} \otimes \cdots \otimes v_{i,d}$$

for suitable vectors $v_{i,1}, \ldots, v_{i,d}$. The *vertical tensor product* of $S \in V^{\otimes d} =: V_1 \otimes \cdots \otimes V_d$ and $T \in W^{\otimes d} = W_1 \otimes \cdots \otimes W_d$ is the d-way tensor $S \boxtimes T$ obtained by regarding $S \otimes T$ as a tensor in

$$(V_1 \otimes W_1) \otimes \cdots \otimes (V_d \otimes W_d) = (V \otimes W)^{\otimes d},$$

where we 'forget' the tensor product structure in each $V_i \otimes W_i$. One has $\mathrm{rk}(S \boxtimes T) \leq \mathrm{rk}(S) \cdot \mathrm{rk}(T)$.

The *asymptotic rank* of S is the limit

$$\lim_{t \to \infty} \sqrt[t]{\mathrm{rk}(S^{\boxtimes t})},$$

where $S^{\boxtimes t} \in (V^{\otimes t})^{\otimes d}$ is the t-fold tensor power of S with itself.

If $S \succeq T$, then $\mathrm{rk}(S) \geq \mathrm{rk}(T)$ and $S^{\boxtimes t} \succeq T^{\boxtimes t}$ for every natural number t, so that also the asymptotic rank of S is at least that of T. This shows that rank and asymptotic rank are both monotone in the following sense.

Definition 1.4.2. A function f that assigns to any d-way tensor a real number is called *restriction-monotone* if $S \succeq T$ implies that $f(S) \geq f(T)$.

Corollary 1.4.3. *Let f be any restriction-monotone function on d-way tensors over the finite field K. Then the set of values of f in $\mathbb{R}$ is a well-ordered set.*

Proof. If not, then there exist tensors $T_1, T_2, \ldots$ on which f takes values $a_1 > a_2 > \cdots$. Let $X_{\leq a_i}$ be the tensor property of having f-value at most a_i. Since f is restriction-monotone, this property is restriction-closed. Furthermore, since $T_i \in X_{\leq a_i} \setminus X_{\leq a_{i+1}}$, we have

$$X_{\leq a_1} \supsetneq X_{\leq a_2} \supsetneq \cdots .$$

But this contradicts Corollary 1.3.3, part (1). □

In particular, the set of asymptotic ranks of d-way tensors over a fixed finite field is well-ordered.

Example 1.4.4. Take $d = 3$. By Corollary 1.4.3, the set $S \subseteq \mathbb{R}_{\geq 0}$ of asymptotic ranks of three-way tensors is well-ordered. This means that $S \setminus [0, 4]$ contains a minimal element $4 + \epsilon$ with $\epsilon > 0$. Hence in particular, the asymptotic rank of 2×2 matrix multiplication, a tensor in $K^4 \otimes K^4 \otimes K^4$, is either 4 (which is equivalent to the well-known conjecture that the exponent of matrix multiplication over K is 2; see [CGL+21]) or at least $4 + \epsilon$. We point out, though, that we do not know whether asymptotic ranks of tensors over an *infinite* field are well-ordered, because the property of having asymptotic rank at most some real number is not (evidently, at least) a Zariski-closed property (see Remark 1.3.5).

Example 1.4.5. Another restriction-monotone function on tensors is the *asymptotic subrank*. In [CGZ23], it is proved that a tensor in $V^{\otimes d}$ has asymptotic subrank 0 (in which case it is 0), 1 (in which case it is, up to a permutation of the tensor factors, of the form $v \otimes T$ with $v \in V$ and

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

$T \in V^{\otimes d-1}$), or at least $2^{h(1/d)}$, where $h:(0,1)\to\mathbb{R}$ is the binary entropy function defined by

$$h(p) := -p\log_2(p) - (1-p)\log_2(1-p);$$

in the latter case, it has a specific tensor W_d as a restriction in the broader sense of Remark 1.1.2, and W_d is known to have the asymptotic subrank above. Interestingly, this result holds over *any* field. Our Corollary 1.4.3 only guarantees such gaps over a fixed finite field. It would be interesting to see to what extent Corollary 1.4.3 extends to other fields.

Example 1.4.6. Another notion of rank that is restriction-monotone is *analytic rank* [Lov19]. Fix a nontrivial character (group homomorphism) $\chi:(K,+)\to(\mathbb{C}^*,\cdot)$. Thinking of $T\in V^{\otimes d}$ as a multilinear form $(V^*)^d \to K$, the analytic rank of T equals

$$-\log_{|K|}\mathbb{E}(\chi(T(x_1,\dots,x_d))),$$

where $\mathbb{E}$ stands for expectation in the probabilistic model where $(x_1,\dots,x_d)$ is picked uniformly at random in $(V^*)^d$. The analytic rank is restriction-monotone by Lemma 1.4.7 below. Hence, by Corollary 1.4.3, the set of analytic ranks of d-linear forms over K is a well-ordered subset of the real numbers.

The following is well known to experts, but we did not find a proof in the published literature, so we provide one here.

Lemma 1.4.7. *The analytic rank is restriction-monotone.*

Proof. It is convenient to see this in the more general setting of Remark 1.1.2, where we have different vector spaces $V_1,\dots,V_d$ and $T\in V_1\otimes\cdots\otimes V_d$ is a regarded as a multilinear function $V_1^*\times\cdots\times V_d^*\to K$.

Consider a linear map $\varphi:V_d\to W$ and define $T' := \mathrm{id}_{V_1}\otimes\cdots\otimes\mathrm{id}_{V_{d-1}}\otimes\varphi$. For fixed $(x_1,\dots,x_{d-1})\in V_1^*\times\cdots\times V_{d-1}^*$, the linear form $T(x_1,\dots,x_{d-1},\cdot)$ is either zero, in which case $\chi(T(x_1,\dots,x_{d-1},x_d))=1$ for all $|V_d|$ choices of x_d, or it is nonzero, in which case, as x_d varies through V, $T(x_1,\dots,x_d)$ takes all values equally often, and therefore the values of χ cancel out. We conclude that the expectation in the analytic rank of T equals

$$a|V_d|/(|V_1|\cdots|V_d|)$$

where a is the number of tuples $(x_1,\dots,x_{d-1})$ for which the linear form is zero. By the same reasoning, the expectation in the analytic rank of T' equals

$$a'|W|/(|V_1|\cdots|V_{d-1}|\cdot|W|)$$

where now a' is the number of tuples $(x_1,\dots,x_{d-1},\cdot)$ for which $T(x_1,\dots,x_{d-1})$ is zero *on the image of* $\varphi^*:W^*\to V^*$. Now $a'\geq a$ and therefore the expression for T' is at least that for T. Taking $-\log_{|K|}$ on both sides, and repeating this argument with linear maps in the other $d-1$ tensor factors, we are done. □

1.5 Generic representations

Let **Vec** be the category of finite-dimensional vector spaces over the finite field K.

Definition 1.5.1. A *generic representation* is a functor $F:\mathbf{Vec}\to\mathbf{Vec}$.

The terminology is explained by the observation that if F is a generic representation, then for each n, $F(K^n)$ is a representation of the finite group $\mathrm{GL}_n(K)$ and of the finite monoid $\mathrm{End}(K^n)$ of $n\times n$ matrices. Generic representations can therefore be thought of as sequences of representations of $\mathrm{End}(K^n)$, one for each n, that depend in a suitably generic manner

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

on n. Generic representations form an abelian category in which the morphisms are natural transformations.

Example 1.5.2. Here are two rather different examples of generic representations:

(1) the functor T^d that sends V to $V^{\otimes d}$ and $\varphi: V \to W$ to $\varphi^{\otimes d}$; and
(2) the functor that sends V to the K-vector space KV with basis $\{e_v \mid v \in V\}$ and $\varphi: V \to W$ to the unique linear map $KV \to KW$ that sends the basis vector $e_v \in V$ to the basis vector $e_{\varphi(v)} \in W$.

The following beautiful theorem characterises a particularly nice class of generic representations.

Theorem 1.5.3 [Kuh94a, Theorem 4.14]. *For a generic representation $F: \mathbf{Vec} \to \mathbf{Vec}$ the following properties are equivalent:*

(1) *F has a finite composition series in the abelian category of generic representations;*
(2) *the function $d_F: \mathbb{Z}_{\geq 0} \to \mathbb{Z}_{\geq 0}$ defined by $d_F(n) := \dim F(K^n)$ is (bounded above by) a polynomial in n; and*
(3) *F is a subquotient of a finite direct sum $T^{d_1} \oplus \cdots \oplus T^{d_n}$ for suitable $d_1, \ldots, d_n \in \mathbb{Z}_{\geq 0}$.*

We call a generic representation satisfying any of the equivalent properties above *polynomial*. Often, we will drop the adjective generic and just speak of *polynomial representations*.

Example 1.5.4. The generic representation $V \mapsto V^{\otimes d}$ is polynomial, and so is the generic representation $V \mapsto S^d V$. The generic representation $V \mapsto KV$ is not polynomial, because $\dim KV = |V| = |K|^{\dim V}$ is exponential in $\dim V$.

1.6 The restriction theorem for polynomial representations

The tensor restriction theorem generalises as follows.

Theorem 1.6.1 (The restriction theorem in polynomial representations). *Let P be a polynomial generic representation over the finite field K, and for $i \in \mathbb{N}$ let $T_i \in P(V_i)$. Then there exist $i < j$ and a linear map $\varphi: V_j \to V_i$ such that $T_i = P(\varphi) T_j$.*

We will use the term 'restriction' also in this more general context, i.e., the conclusion of the theorem says that T_i is a restriction of T_j.

Remark 1.6.2. The condition that P be polynomial cannot be dropped. For instance, let P be the functor that sends V to KV. For each $n \geq 3$ let $T_n \in P(K^{n-1})$ be the formal sum

$$T_n := e_{v_1} + \cdots + e_{v_n} \in P(K^{n-1}),$$

where $\{v_1, \ldots, v_n\} \in K^{n-1}$ is a *circuit*: any $n-1$ of the v_i are a basis of K^{n-1}. We claim that no T_n is a restriction of any T_m with $m \neq n$. Indeed, if it were, then writing $T_m = e_{v'_1} + \cdots + e_{v'_m}$, there would be a linear map $K^{m-1} \to K^{n-1}$ that maps the circuit $\{v'_1, \ldots, v'_m\}$ to the circuit $\{v_1, \ldots, v_n\}$. By basic linear algebra, such linear maps do not exist.

Corollary 1.3.3 generalises verbatim to polynomial representations, and so does Corollary 1.4.3.

Remark 1.6.3. Versions of Theorem 1.6.1 and its corollaries also hold for *multivariate polynomial representations*, defined as functors $P: \mathbf{Vec}^k \to \mathbf{Vec}$ for which $\dim_K P(K^{n_1}, \ldots, K^{n_k})$ is a polynomial in $n_1, \ldots, n_k$. Indeed, given elements $T_i \in P(V_i^{(1)}, \ldots, V_i^{(k)})$ for $i = 1, 2, \ldots$, we can

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

choose linear injections $\iota_i^{(j)}$ from $V_i^{(j)}$ into a $U_i \in \mathbf{Vec}$ (which depends only on i), and linear surjections $\pi_i^{(j)} : U_i \to V_i^{(j)}$ with $\pi_i^{(j)} \circ \iota_i^{(j)} = \mathrm{id}_{V_i^{(j)}}$. Then define

$$T_i' := P(\iota_i^{(1)}, \ldots, \iota_i^{(k)}) T_i \in P(U_i, \ldots, U_i) =: Q(U_i)$$

where Q is now a univariate polynomial generic representation. Theorem 1.6.1 applied to Q says that there exist $i < j$ and a linear map $\psi : U_j \to U_i$ such that

$$Q(\psi) T_j' = P(\psi, \ldots, \psi) T_j' = T_i'.$$

We then have

$$P(\pi_i^{(1)} \circ \psi \circ \iota_j^{(1)}, \ldots, \pi_i^{(k)} \circ \psi \circ \iota_j^{(1)}) T_j = T_i,$$

as desired.

1.7 Proof strategy: the polynomial method

Rather than proving the restriction theorem for polynomial representations directly, we will prove Noetherianity, corresponding to (1) in Corollary 1.3.3: if P is a polynomial representation, and $X_1 \supseteq X_2 \supseteq \cdots$ are restriction-closed properties, then $X_n = X_{n+1}$ for all sufficiently large n.

To establish Noetherianity, we adapt the proof method of [Dra19] for polynomial functors over *infinite fields* to our current setting. This is far from straightforward. For instance, a polynomial functor over an infinite field and its coordinate ring both have a $\mathbb{Z}_{\geq 0}$-grading, whereas a polynomial representation over the finite field K and its coordinate ring only have a grading by $\{0, 1, \ldots, |K| - 1\}$. Nevertheless, after introducing the *degree* d of the polynomial representation P, we show that P has a unique minimal subrepresentation $P_{>d-1}$, the quotient by which has degree at most $d-1$. We think of $P_{>d-1}$ as the *top-degree part* of P. We then take an irreducible subrepresentation R in $P_{>d-1}$, and assume that the Noetherianity statement holds for P/R and various other polynomial representations that have the same top-degree part as P/R and are therefore in a lexicographic sense smaller than P. This means that if $X_1 \supseteq X_2 \supseteq \cdots$ is a chain of restriction-closed properties in P, then their projections $X_1' \supseteq X_2' \supseteq \cdots$ in P/R stabilise. Therefore, it suffices to prove Noetherianity for properties $X \subseteq P$ that have a fixed projection $X' \subseteq P/R$. Then, to prove that any property $X \subseteq P$ with projection X' is Noetherian, we think of each $X(V)$ as a *Zariski-closed subset* of $P(V)$, i.e., as given by polynomial equations in the finite vector space $P(V)$. We do induction on the minimal degree of an equation that vanishes identically on X but not on X'. Using *spreading operators*, we show that from such an equation we can construct many equations of the same degree that are affine-linear in the R-direction. This allows us to embed a certain subset of X into a strictly smaller polynomial functor, while on the complement of that subset a polynomial of strictly smaller degree vanishes. Both subsets can therefore be handled by induction.

We stress that this proof never actually looks at concrete tensors or elements of $P(V)$; all reasoning uses polynomial equations, and exploits the fact that every subset of a finite vector space is given by polynomial equations. In this sense, the proof can be regarded an instance of the polynomial method.

We remark that polynomial generic representations are not the same thing as strict polynomial functors in the sense of Friedlander and Suslin [FS97]. Roughly speaking, while former deal with sequences of representations of the finite groups $\mathrm{GL}_n(K)$, the latter deal with sequences of algebraic representations of the group schemes GL_n. Topological Noetherianity of strict polynomial functors, over arbitrary rings with Noetherian spectrum and hence certainly over finite

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

fields, was established in [BDD23], using the techniques from [Dra19]. However, strict polynomial functors have a scheme structure built in and are therefore much more amenable to the techniques of [Dra19] than the polynomial generic representations that we study here. Furthermore, even if one is interested only in polynomial generic representations that come from strict polynomial functors by forgetting some of the data (such as the functor $V \mapsto V^{\otimes d}$) our proof, in which we mod out an irreducible subrepresentation R, requires that one leaves the realm of these special representations. This explains the need for the new ideas developed in this paper.

1.8 Further relations to the literature

Restriction-closed properties of tensors are a rapidly expanding research area. Here is a very small selection of recent research related to our work.

In [Kar22] it is proved, for various notions of rank including ordinary tensor rank, slice rank, and partition rank, that a large tensor of rank r has a subtensor whose size depends only on r and whose rank is at least some function of r. For finite fields, this result also follows from Corollary 1.3.3, item (3); in fact, by that item, a subtensor of fixed size can be found of rank equal to r. However, Karam also finds an explicit formula for the size, while our theorem does not give such a bound. It would be very interesting to see whether *the proof* of our theorem could shed further light on such bounds.

In [CM23] it is proved that over sufficiently large fields, partition rank is bounded by a linear function of the analytic rank of a tensor; and in [MZ22] the condition on the field size is removed at the cost of a polylogarithmic factor. This is the culmination of many years of research by many authors into the relation between bias and rank of tensors, starting with [GT09] via polynomial bounds in [Mil19] and linear bounds for trilinear forms in [AKZ21]. Using the proof of our tensor restriction theorem and techniques from [BDE19], it is easy to recover the result that partition rank is bounded from above by at least *some* function of the analytic rank. However, again, our techniques do not yield bounds that can compete with the state of the art.

In [CGL$^+$21], motivated by Strassen's asymptotic rank conjecture that says that any *tight* tensor has the minimal possible asymptotic rank, the authors study the geometry of various varieties of tensors, such as the (closure of) the set of tight tensors. It would be interesting to study these varieties from the perspective of this paper (over finite fields) and from the perspective of [BDES23] (over infinite fields). In both cases, after a shift and a localisation, these varieties become of the form a fixed finite-dimensional variety times an affine space that depends on the size of the tensor. Over infinite fields this follows from the *shift theorem* in [BDES23], and over finite fields it follows from the weak shift theorem in this paper.

In [PS17] and [SS17], the long-standing *Lannes–Schwartz Artinian conjecture* was resolved, which says that any finitely generated (not necessarily polynomial) generic representation $F\colon \mathbf{Vec} \to \mathbf{Vec}$ is Noetherian in the module sense: it satisfies the ascending chain condition on subrepresentations. Dually, this means that any descending chain of subrepresentations of $F^*\colon V^* \mapsto F(V)^*$ stabilises. Interpreting the elements of $F(V)$ as linear functions on $F(V)^*$, one may interpret this as Noetherianity for *linear* functorial subsets of F^*. However, already for $F\colon V \mapsto K \cdot V$, one can show that F^* does not satisfy the descending chain condition on *nonlinear* subsets. So topological Noetherianity as we prove it seems restricted only to polynomial generic representations. It would be nice to know a precise statement to this effect. For instance, is it true that the *only* generic representations for which the restriction theorem holds are the polynomial representations?

In [Sno21], Snowden extends many results about GL-algebras from [BDES23] to modules over GL-algebras equipped with a compatible GL-action. Along the way, he also gives a proof of the shift theorem that differs slightly from the proof in [BDES23], and which uses the search for an element of weight $(1, \ldots, 1)$ in a suitable GL-representation. This inspired the development of weight theory in our current, different context in § 3 and the idea that a suitably spread-out element in the vanishing ideal of a tensor property would have weight $(1, \ldots, 1)$, which is a key insight in the proof of the embedding theorem in § 4.6.

1.9 Organisation of this paper

In § 2 we discuss the theory of generic polynomial representations, including the definition of top-degree parts and shift functors. In particular, we will see that the ring of functions on a polynomial representation is itself a countable union of polynomial representations.

In § 3 we develop a partial analogue of the classical weight theory for representations of group schemes GL_n. This includes the spreading operators alluded to above. Since functions on a polynomial representation themselves live in a polynomial representation, these spreading operators also act on functions.

In § 4 we prove Noetherianity for polynomial representations over the fixed finite field K, which implies the restriction theorems for tensors and polynomial representations and items (1)–(3) of Corollary 1.3.3, both for tensors and for polynomial representations, and also implies the existence of a unique decomposition of a restriction-closed tensor property into irreducible such properties (see Theorem 4.2.2). We do so by first deriving Noetherianity from an auxiliary result that we call the embedding theorem, since it is the finite-field analogue of the embedding theorem in [BDES23]. The proof of the embedding theorem, then, is the heart of the paper. We also derive from it a version of the shift theorem in [BDES23].

Finally, in § 5 we use the theory of finitely generated FI-modules to prove item (4) from Corollary 1.3.3; as we have seen, (5) is then a direct consequence.

2. Polynomial generic representations

Throughout the paper, K is a fixed finite field, with q elements. All linear and multilinear algebra will be over K. We denote by **Vec** the category of finite-dimensional K-vector spaces, and for $V \in \mathbf{Vec}$ we denote the dual space by V^*.

2.1 Functions as polynomials

We introduce the ring of functions on a vector space; we will also call this the coordinate ring.

Definition 2.1.1. Given $V \in \mathbf{Vec}$, we write $K[V]$ for the K-algebra of functions $V \to K$. This has a natural algebra filtration

$$\{0\} = K[V]_{\leq -1} \subseteq K[V]_{\leq 0} \subseteq K[V]_{\leq 1} \subseteq K[V]_{\leq 2} \subseteq \cdots$$

where $K[V]_{\leq d}$ is the set of functions $f : V \to K$ for which there exists an element of $\bigoplus_{e=0}^{d} S^e V^*$ that defines the function f.

We stress that $K[V]$ is an algebra of functions, not of polynomials. More precisely, $K[V]$ is the quotient of the symmetric algebra SV^* by the ideal generated by the polynomials $x^q - x$ as x runs through (a basis of) V^*. Since these polynomials are not homogeneous, $K[V]$ has no natural *grading*; however, as seen above, it does have a natural *filtration*.

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

Note further that $K[V]$ is a finite-dimensional K-vector space, of dimension $q^{\dim(V)}$, the number of elements of V.

Definition 2.1.2. Given a basis $x_1, \dots, x_n$ of V^*, every element f of $K[V]$ has a unique representative polynomial in which all exponents of all variables are less than or equal to $q-1$; we will call this representative, which depends on the choice of basis, the *reduced* polynomial representation for f relative to the choice of coordinates.

The following lemma is immediate; the natural isomorphisms in it will be interpreted as equalities throughout the paper.

Lemma 2.1.3. *For $V, W \in \mathbf{Vec}$ we have $K[V \times W] \cong K[V] \otimes K[W]$ via the K-linear map from right to left that sends $f \otimes g$ to the function $(v,w) \mapsto f(v)g(w)$; this is a K-algebra isomorphism.*

Similarly, the set of arbitrary maps $V \to W$ is canonically isomorphic to $K[V] \otimes W$ via the K-linear map from right to left that sends $f \otimes w$ to the function $v \mapsto f(v) \cdot w$.

Furthermore, we write $K[V]_0 = K$ for the sub-K-algebra of constant functions, and $K[V]_{>0}$ for the K-vector space spanned by all functions that vanish at zero.

2.2 Polynomial generic representations over K

Recall Theorem 1.5.3, which characterises polynomial representations among all generic representations. We will use the following alternative characterisation instead.

Definition 2.2.1. A generic representation $P: \mathbf{Vec} \to \mathbf{Vec}$ is called *polynomial* if there exists a d such that for all $U, V \in \mathbf{Vec}$ the map $P: \mathrm{Hom}(U,V) \to \mathrm{Hom}(P(U), P(V))$ lies in $K[\mathrm{Hom}(U,V)]_{\leq d} \otimes \mathrm{Hom}(P(U), P(V))$. The minimal such $d \in \mathbb{Z}_{\geq -1}$ is called the *degree* of P and denoted $\deg(P)$.

Polynomial representations form an abelian category, in which the morphisms are natural transformations.

Any polynomial representation in the sense of Theorem 1.5.3 is a subquotient of a direct sum $T^{d_1} \oplus \cdots \oplus T^{d_k}$, and this implies that it is polynomial in the sense of the definition above, of degree at most the maximum of the d_i. In Remark 2.4.2, we will see that, conversely, any generic representation that is polynomial in the sense of the definition above is polynomial in the sense of Theorem 1.5.3.

Every finite-degree *strict* polynomial functor $\mathbf{Vec} \to \mathbf{Vec}$ in the sense of Friedlander-Suslin [FS97] gives rise to a polynomial representation. But this forgetful functor is not an equivalence of abelian categories. For instance, if Q is a strict polynomial functor of degree d over K, then its q-Frobenius twist is a polynomial functor of degree dq over K and hence not isomorphic to Q. However, Q and its q-Frobenius twist give rise to the same generic representation.

2.3 Schur algebras over K

In spite of the discrepancy between strict polynomial functors and polynomial generic representations, a version of the theorem by Friedlander and Suslin that relates polynomial functors to representations of the Schur algebra, does hold.

Fix a natural number d and a $U \in \mathbf{Vec}$. The composition map $\mathrm{End}(U) \times \mathrm{End}(U) \to \mathrm{End}(U)$ gives rise, via pullback of functions, to a K-linear map

$$K[\mathrm{End}(U)]_{\leq d} \to K[\mathrm{End}(U)]_{\leq d} \otimes K[\mathrm{End}(U)]_{\leq d}.$$

We write $A_{\leq d}(U) := K[\mathrm{End}(U)]^*_{\leq d}$. Dualising the map above, we obtain a K-bilinear map

$$A_{\leq d}(U) \times A_{\leq d}(U) \to A_{\leq d}(U).$$

A straightforward computation, using the associativity of composition of linear maps, shows that this turns $A_{\leq d}(U)$ into a unital, associative algebra, with unit element $f \mapsto f(\mathrm{id}_U)$.

DEFINITION 2.3.1. The unital, associative algebra $A_{\leq d}(U)$ with the multiplication above is called the *Schur algebra* over K.

We remark that this is in fact a subalgebra of the Schur algebra in [FS97], which is the dual space to the space of *polynomials* of degree at most d on $\mathrm{End}(U)$. Our Schur algebra consists of only those linear functions that vanish on the ideal of polynomials that define the zero function on $\mathrm{End}(U)$.

The Schur algebra comes with a homomorphism of monoids (not of K-algebras) $\mathrm{End}(U) \to A_{\leq d}(U)$ defined by $\varphi \mapsto (f \mapsto f(\varphi))$. This homomorphism is an embedding if $d \geq 1$.

2.4 A finite-field analogue of the Friedlander–Suslin lemma

Given a polynomial generic representation P of degree at most d and a vector space U, we turn $P(U)$ into an $A_{\leq d}(U)$-module by a construction very similar to the construction of $A_{\leq d}(U)$: first, the map

$$\mathrm{End}(U) \times P(U) \to P(U), (\varphi, p) \mapsto P(\varphi)(p)$$

gives rise, via pullback, to a K-linear map

$$P(U)^* \to K[\mathrm{End}(U)]_{\leq d} \otimes P(U)^*.$$

Dualising, we obtain a K-bilinear map

$$A_{\leq d}(U) \times P(U) \to P(U)$$

that turns $P(U)$ into a (unital) $A_{\leq d}(U)$-module.

The following proposition is proved exactly as Friedlander and Suslin's corresponding theorem, and it is also almost equivalent to [Kuh94a, Proposition 4.10].

PROPOSITION 2.4.1. *Fix a natural number d and a $U \in \mathbf{Vec}$ of dimension at least d. Then $P \mapsto P(U)$ is an equivalence of abelian categories from the category of polynomial generic representations $\mathbf{Vec} \to \mathbf{Vec}$ of degree at most d to the category of $A_{\leq d}(U)$-representations.*

Remark 2.4.2. It follows from this proposition that every polynomial representation of degree at most d in the sense of Definition 2.2.1 has finite length. Therefore, it is also polynomial in the sense of Theorem 1.5.3.

Remark 2.4.3. If P is an irreducible polynomial generic representation, then for each $U \in \mathbf{Vec}$, $P(U)$ is (zero or) an irreducible $\mathrm{End}(U)$-module. Indeed, if M were a nonzero proper submodule, then, for varying V,

$$Q(V) := \{p \in P(V) \mid \forall \varphi \in \mathrm{Hom}_{\mathbf{Vec}}(V, U) \ P(\varphi)p \in M\}$$

would define a nonzero proper subrepresentation.

2.5 Filtering a polynomial representation by degree

A strict polynomial functor in the sense of Friedlander and Suslin has a *grading* by degree. In contrast, we will see that a polynomial generic representation only has a *filtration* by degree. One notable exception is the degree-zero part of a polynomial representation.

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

Definition 2.5.1. Let $P: \mathbf{Vec} \to \mathbf{Vec}$ be a polynomial generic representation. Then we define $P_0: \mathbf{Vec} \to \mathbf{Vec}$ by $P_0(V) := P(0) =: U$ for all $V \in \mathbf{Vec}$ and $P_0(\varphi) := \mathrm{id}_U$ for all $\varphi \in \mathrm{Hom}_{\mathbf{Vec}}(V, W)$. This is a direct summand of P in the abelian category of polynomial generic representations, called the *degree-zero part* or *constant* part of P; we will also informally say that U is the constant part of P. The constant part P_0 has a unique complement in P, namely,

$$P_{>0}(V) := \{p \in P(V) \mid P(0 \cdot \mathrm{id}_V)p = 0\}.$$

A polynomial representation P of degree at most some $e \leq d$ is in particular a polynomial representation of degree at most d. On the Schur algebra side, this inclusion of abelian categories is made explicit as follows. Take a vector space U of dimension at least d. Then the inclusion $K[\mathrm{End}(U)]_{\leq e} \to K[\mathrm{End}(U)]_{\leq d}$ dualises to a linear surjection $A_{\leq d}(U) \to A_{\leq e}(U)$. This surjection is an algebra homomorphism, and hence if M is a module over the latter algebra, then it is also naturally a module over the former algebra.

This interpretation also shows *which* $A_{\leq d}(U)$-modules are also $A_{\leq e}(U)$-modules, namely, those for which the kernel I of the surjection acts as zero. Furthermore, if M is an $A_{\leq d}(U)$-module, and N is an $A_{\leq d}(U)$-submodule of M, then M/N is an $A_{\leq e}(U)$-module if and only if $I \cdot (M/N) = 0$, i.e., if and only if N contains the $A_{\leq d}(U)$-submodule $I \cdot M$. We conclude that there is a unique inclusionwise minimal $A_{\leq d}(U)$-submodule N of M such that M/N is an $A_{\leq e}(U)$-module, namely, $N = I \cdot M$.

By Proposition 2.4.1 we may translate this back to polynomial representations.

Proposition 2.5.2. *For any polynomial representation P and any $e \in \mathbb{Z}_{\geq -1}$, there is a unique inclusionwise minimal subrepresentation Q such that P/Q is a polynomial representation of degree at most e.*

Definition 2.5.3. Let P be a polynomial representation and let $e \in \mathbb{Z}_{\geq -1}$. The unique inclusionwise minimal subrepresentation Q of P such that P/Q has degree at most e is denoted by $P_{>e}$.

Example 2.5.4. Suppose that $\mathrm{char}\, K = 2$. Consider the polynomial representation $P: V \mapsto S^2 V$ and the polynomial representation Q that sends V to the space of symmetric tensors in $V \otimes V$. Then P has as a subrepresentation the representation R that maps V to the space of squares of elements of V, and this is the only nontrivial subrepresentation unequal to P itself. The quotient P/R has degree 2, so $P_{>1} = P$. On the other hand, Q has the subrepresentation T that assigns to V the set of skew-symmetric tensors in $V \otimes V$ (i.e., those in the linear span of tensors of the form $u \otimes v - v \otimes u$ as u, v range through V) and the quotient Q/T is isomorphic to R. Now if $K = \mathbb{F}_2$, then R has degree 1, so that $Q_{>1} = T$; while if $K \neq \mathbb{F}_2$, then R has degree 2, and therefore $Q_{>1} = Q$.

We clearly have

$$P = P_{>-1} \supseteq P_{>0} \supseteq \cdots \supseteq P_{>d} = \{0\}$$

where $d = \deg(P)$; and a straightforward check shows that $P_{>0}$ in this definition agrees with the direct complement $P_{>0}$ of P_0 in Definition 2.5.1.

Lemma 2.5.5. *If $\alpha: P \to Q$ is a morphism in the abelian category of polynomial representations, then for each e we have $\alpha(P_{>e}) \subseteq Q_{>e}$.*

Proof. We compute

$$P/\alpha^{-1}(Q_{>e}) \cong \mathrm{im}(\alpha)/(\mathrm{im}(\alpha) \cap Q_{>e}) \subseteq Q/Q_{>e}.$$

Since the latter representation has degree at most e, so does the first. The defining property of $P_{>e}$ then implies that $P_{>e} \subseteq \alpha^{-1}(Q_{>e})$. This is equivalent to the statement in the lemma. □

Lemma 2.5.6. *Let P be a polynomial representation, let $e \in \mathbb{Z}_{\geq -1}$, and let R be a subobject of $P_{>e}$, and hence of P. Then $(P/R)_{>e} \cong P_{>e}/R$.*

Proof. By Lemma 2.5.5, the morphism $P \to P/R$ maps $P_{>e}$ into $(P/R)_{>e}$, and its kernel on $P_{>e}$ is R, so that $P_{>e}/R$ maps injectively into $(P/R)_{>e}$. To see that it also maps surjectively, we note that

$$(P/R)/(P_{>e}/R) \cong P/P_{>e}$$

has degree at most e. Hence $P_{>e}/R$ contains $(P/R)_{>e}$ by definition of the latter object. □

2.6 Shifting

Just as a univariate polynomial can be shifted over a constant, and then its leading term does not change, a polynomial representation can be shifted over a constant vector space, and we will see that its top-degree part does not change.

Definition 2.6.1. Given a $U \in \mathbf{Vec}$ and a representation $P : \mathbf{Vec} \to \mathbf{Vec}$, we define the representation $\mathrm{Sh}_U P$ by $(\mathrm{Sh}_U P)(V) := P(U \oplus V)$ and $(\mathrm{Sh}_U P)(\varphi) := P(\mathrm{id}_U \oplus \varphi)$ for $\varphi \in \mathrm{Hom}(V, W)$. We call $\mathrm{Sh}_U P$ the *shift of P by U*.

If P is polynomial of degree at most d, then $\mathrm{Sh}_U P$ is also polynomial of degree at most d; below we will prove a more precise statement.

We have a morphism $\alpha : P \to \mathrm{Sh}_U P$ in the abelian category of polynomial generic representations defined by $\alpha_V = P(\iota_V) : P(V) \to P(U \oplus V)$, where $\iota_V : V \to U \oplus V$ is the inclusion $v \mapsto 0 + v$. Indeed, that $(\alpha_V)_V$ is a morphism follows from the commutativity of the following diagram, for any $\varphi \in \mathrm{Hom}_{\mathbf{Vec}}(V, W)$.

$$\begin{array}{ccc} P(V) & \xrightarrow{P(\iota_V)} & P(U \oplus V) \\ {\scriptstyle P(\varphi)}\downarrow & & \downarrow{\scriptstyle P(\mathrm{id}_U \oplus \varphi)} \\ P(W) & \xrightarrow[P(\iota_W)]{} & P(U \oplus W) \end{array}$$

This in turn follows from the fact that P is a representation and that

$$(\mathrm{id}_U \oplus \varphi) \circ \iota_V = \iota_W \circ \varphi.$$

Similarly, we have a morphism $\beta : \mathrm{Sh}_U P \to P$ defined by $\beta_V = P(\pi_V) : P(U \oplus V) \to P(V)$, where $\pi : U \oplus V \to V$ is the projection $u + v \mapsto v$. The relation $\pi \circ \iota = \mathrm{id}_V$ translates to $\beta \circ \alpha = \mathrm{id}_P$. This implies that $\mathrm{Sh}_U P$ is the direct sum of $\mathrm{im}(\alpha) \cong P$ and the polynomial representation $Q := \ker(\beta)$.

The following lemma says, informally, that the top-degree part of a polynomial representation is invariant under shifting.

Lemma 2.6.2. *Assume that $\deg(P) = d \geq 0$. Then $(\mathrm{Sh}_U P)_{>d-1} \cong P_{>d-1}$.*

Proof. Using the notation α and β from above, we have $\alpha(P_{>d-1}) \subseteq (\mathrm{Sh}_U P)_{>d-1}$ and $\beta((\mathrm{Sh}_U P)_{>d-1}) \subseteq P_{>d-1}$ by Lemma 2.5.5. Combining these facts shows that α maps $P_{>d-1}$ injectively into $(\mathrm{Sh}_U P)_{>d-1}$. To argue that it also maps surjectively there, it suffices to show that $(\mathrm{Sh}_U P)/\alpha(P_{>d-1})$ has degree at most $d-1$.

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

To see this, we recall that $\mathrm{Sh}_U P = \mathrm{im}(\alpha) \oplus Q$, where $Q = \ker(\beta)$. Accordingly,

$$(\mathrm{Sh}_U P)/\alpha(P_{>d-1}) \cong (\alpha(P)/\alpha(P_{>d-1})) \oplus Q.$$

Here the first summand on the right is isomorphic to $P/P_{>d-1}$, hence of degree at most $d-1$. So it suffices to show that Q has degree at most $d-1$ as well. Consider a vector $q \in Q(V)$ and a linear map $\varphi \in \mathrm{Hom}(V,W)$. Then we have

$$\begin{aligned} Q(\varphi)(q) &= P(\mathrm{id}_U \oplus \varphi)(q) = P(\mathrm{id}_U \oplus \varphi)(q) - P(\mathrm{id}_U \oplus \varphi)(\alpha_V(\beta_V(q))) \\ &= (P(\mathrm{id}_U \oplus \varphi) - P(0_U \oplus \varphi))(q) \end{aligned}$$

where the second equality follows from $\beta_V(q) = 0$ and the last equality follows from the definition of α and β. Now, for ψ running through $\mathrm{Hom}(U \oplus V, U \oplus W)$, $P(\psi)$ can be described by a polynomial map of degree at most d. If in this map we substitute for ψ the maps $\mathrm{id}_U \oplus \varphi$ and $0_U \oplus \varphi$ respectively, we obtain the same degree-d parts in φ. Hence the map $\varphi \mapsto P(\mathrm{id}_U \oplus \varphi) - P(0_U \oplus \varphi)$ is given by a polynomial map of degree at most $d-1$ in the entries of φ. This shows that Q has degree at most $d-1$, as desired. □

2.7 A well-founded order on polynomial representations

Definition 2.7.1. Given polynomial representations $Q, P : \mathbf{Vec} \to \mathbf{Vec}$, we write $Q \preceq P$ if $Q \cong P$ or else for the largest e such that $Q_{>e} \ncong P_{>e}$ the former is a quotient of the latter. We write $Q \prec P$ to mean $Q \preceq P$ and $Q \ncong P$.

Lemma 2.7.2. *The relation $\preceq$ is a well-founded pre-order on polynomial representations.*

Proof. Reflexivity is immediate. To see transitivity, assume $R \preceq Q \preceq P$. If one of the inequalities is an isomorphism, it follows immediately that $R \preceq P$. Suppose that they are both not isomorphisms. Let e be maximal such that $Q_{>e} \ncong P_{>e}$ and let e' be maximal such that $R_{>e'} \ncong Q_{>e'}$. If $e' \geq e$, then e' is maximal such that $R_{>e'} \ncong P_{>e'}$, and the former is a quotient of the latter. If $e' < e$, then e is maximal such that $(Q_{>e} \cong) R_{>e} \ncong P_{>e}$, and the former is a quotient of the latter. In both cases, we find $R \prec P$, as desired.

To see that $\preceq$ is well-founded, suppose we had an infinite chain

$$P_1 \succ P_2 \succ \cdots .$$

To each P_i we associate a length sequence $\ell(P_i) \in \mathbb{Z}_{\geq 0}^{\{-1,0,1,2,\ldots\}}$, where $\ell(P_i)(e)$ is the length of any composition chain of $(P_i)_{>e}$ in the abelian category of polynomial representations; by Proposition 2.4.1 this length is finite.

Note that $\ell(P_i)(e) = 0$ for $e \geq \deg(P_i)$, i.e., $\ell(P_i)$ has finite support. Now $P_i \succ P_{i+1}$ implies that $\ell(P_{i+1})$ is lexicographically strictly smaller than $\ell(P_i)$. Since the lexicographic order on sequences with finite support is a well-order, we arrive at a contradiction. □

The following construction will play a crucial role in our proof.

Lemma 2.7.3. *Let $P \neq 0$ be a polynomial representation of degree $d \geq 0$, and let R be an irreducible subobject of $P_{>d-1}$. Let $U \in \mathbf{Vec}$ and set $Q := \mathrm{Sh}_U P$. Then $Q/R \prec P$.*

Proof. By Lemma 2.6.2, R is also naturally a subobject of $Q_{>d-1}$, which in turn is a subobject of Q. This explains the notation Q/R. By Lemma 2.5.6, we have $(Q/R)_{>d-1} \cong Q_{>d-1}/R$, which in turn is isomorphic to $P_{>d-1}/R$, a quotient of $P_{>d-1}$. Since $P_{>e} = Q_{>e} = 0$ for $e \geq d$, we conclude that $Q/R \prec P$. □

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

2.8 The coordinate ring of a polynomial representation

DEFINITION 2.8.1. Let $P : \mathbf{Vec} \to \mathbf{Vec}$ be a polynomial representation. We define $K[P]$ as the contravariant functor from **Vec** to K-algebras that assigns to V the ring $K[P(V)]$ and to a linear map $\varphi : V \to W$ the pullback $P(\varphi)^{\#} : K[P(W)] \to K[P(V)]$. We call $K[P]$ the *coordinate ring of* P.

Note that $P(\varphi)^{\#}$ is an algebra homomorphism; this is going to be of crucial importance in § 4.6. The coordinate ring comes with a natural ring filtration,

$$\{0\} = K[P]_{\le -1} \subseteq K[P]_{\le 0} \subseteq K[P]_{\le 1} \subseteq K[P]_{\le 2} \subseteq \cdots ,$$

where $K[P]_{\le e}$ assigns to V the space $K[P(V)]_{\le e}$.

LEMMA 2.8.2. *If P is a polynomial representation of degree at most d, then $V^* \mapsto K[P(V)]_{\le e}$ is a polynomial representation of degree at most $d.e$.*

Proof. This representation assigns to a linear map $\varphi : V^* \to W^*$ the restriction of the pullback $P(\varphi^*)^{\#} : K[P(V)] \to K[P(W)]$ to $K[P(V)]_{\le e}$. Since $P(\varphi^*)$ is a linear map, this pullback does indeed map $K[P(V)]_{\le e}$ into $K[P(W)]_{\le e}$, and it does so via a linear map that is polynomial of degree at most e in $P(\varphi^*)$, hence of degree at most $d.e$ in φ^*, which in turn depends linearly on φ. □

Example 2.8.3. Let $P = S^2$ and assume $|K| > 2$. Take $V = K^n$ with basis $e_1, \ldots, e_n$, so that $P(V)$ has basis $e_i e_j$ with $i \le j$. For $k > l$, let $g_{kl}(s) \in \mathrm{End}(V)$ be the matrix with 1s on the diagonal, an s in position (k, l), and 0s elsewhere. We have

$$\begin{aligned} P(g_{kl}(s)) \sum_{i \le j} a_{ij} e_i e_j &= \sum_{i \le j} a_{ij} (g_{kl}(s) e_i)(g_{kl}(s) e_j) \\ &= \sum_{i \le j} a_{ij} (e_i + \delta_{il} s e_k)(e_j + \delta_{jl} s e_k) \\ &= \sum_{i \le j} a_{ij} (e_i e_j + s(\delta_{jl} e_i e_k + \delta_{il} e_k e_j) + s^2 \delta_{il} \delta_{jl} e_k^2) \\ &= \left(\sum_{i \le j} a_{ij} e_i e_j \right) + s \left(\sum_{i \le l} a_{il} e_i e_k + \sum_{j \ge l} a_{lj} e_k e_j \right) + s^2 a_{ll} e_k^2. \end{aligned}$$

Observe that by acting with g_{kl} on (linear combinations of) the basis vectors $e_i e_j$, indices l either remain the same or turn into indices k.

We now look at the dual. Let $\{x_{ij} \mid i \le j\}$ be the basis of $P(V)^*$ dual to the given basis of $P(V)$. Then, for instance, for $l < i < k$ we have

$$P(g_{kl}(s))^{\#} x_{ik} = x_{ik} + s x_{li},$$

as can be seen by taking the coefficient of $e_i e_k$ in the expression above. We observe here that indices k either remain the same or turn into indices l. We can also write the above as

$$P(g_{lk}(s)^T)^{\#} x_{ik} = x_{ik} + s x_{li}.$$

Note that $P(g)^{\#}$ is contravariant in g, and hence $P(g^T)^{\#}$ is again covariant. This explains the V^* in Lemma 2.8.2.

3. Weight theory

In the representation theory of the group scheme GL_n, the weight space decomposition of a representation, i.e., its decomposition as a module over the subgroup of diagonal matrices, is of

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

crucial importance. For the finite group $G = \mathrm{GL}_n(K)$ with $K = \mathbb{F}_q$, it was already observed in [Ste16, p. 129, Remark] that weights alone do not suffice to distinguish the roles played by various vectors in a representation. The example given there is that the highest weight vector 1 in the trivial $\mathrm{GL}_n(K)$-representation K^1 and the highest weight vector e_1^{q-1} in the $(q-1)$th symmetric power $S^{q-1}K^n$ of the standard representation both have weight $(0,\dots,0)$. It is explained there how to act with elements of the group algebra of G to distinguish the two.

In this particular example, we *can* distinguish these vectors by extending the action of (diagonal matrices in) $\mathrm{GL}_n(K)$ to (diagonal matrices in) $\mathrm{End}(K^n)$, as we do in the context of generic polynomial representations: the first highest weight vector then has weight $(0,\dots,0)$, while the latter has weight $(q-1,0,\dots,0)$; see the definitions below. However, the vector $e_1^{\otimes q}$ in the qth tensor power cannot be distinguished from the vector e_1 in the standard representation via diagonal matrices. Therefore we, too, will act with suitable elements of the monoid algebra of $\mathrm{End}(K^n)$ to get a better grasp on weight vectors. However, our focus will not be on *highest* weight vector; rather, we will look for *middle* weight vectors, i.e., weight vectors whose weight is maximally spread out in a sense that we will make precise below.

We are by no means the first to study weights in this context. For instance, they also feature as *reduced weights* in [Kuh94b]. However, the procedure of maximally spreading out weight that we introduce below does seem to be new.

3.1 Multiplicative monoid homomorphisms $K \to K$

A monoid homomorphism $(K,\cdot) \to (K,\cdot)$ is a map $\varphi : K \to K$ with $\varphi(1) = 1$ and $\varphi(ab) = \varphi(a)\varphi(b)$ for all $a,b \in K$. In particular, φ restricts to a group homomorphism from the multiplicative group $K^\times := K \setminus \{0\}$ to itself. Since $K^\times$ is cyclic, say with generator g, the monoid homomorphism φ is uniquely determined by its values on g and on 0. Write $\varphi(g) = g^e$ for a unique exponent $e \in \{1,\dots,q-1\}$. If $e \neq q-1$, so that $\varphi(g) \neq 1$, then $\varphi(0)$ is forced to be 0, since otherwise $\varphi(g)\cdot\varphi(0)$ does not equal $\varphi(g\cdot 0) = \varphi(0)$. If $e = q-1$, then there are two possibilities for $\varphi(0)$, namely, $\varphi(0) = 1$ and $\varphi(0) = 0$. In the first case we will denote φ by $c \mapsto c^0$, and in the second case we denote by $c \mapsto c^{q-1}$. The following lemma is now straightforward.

Lemma 3.1.1. *The monoid of monoid homomorphisms $K \to K$ is isomorphic to the monoid $\{0,\dots,q-1\}$ with operation $i \oplus j$ defined by $i \oplus j = i+j$ if $i+j \le q-1$ and $i \oplus j = i+j-(q-1)$ otherwise.*

Note that this monoid is not cancellative, since $0 \oplus j = (q-1) \oplus j$ for all $j \in \{1,\dots,q-1\}$. Nevertheless, it will be convenient to have a notation for subtracting elements in the following sense: for $i \in \{1,\dots,q-1\}$ and $j \in \{0,\dots,q-1\}$ we write $i \ominus j$ for the unique element in $\{1,\dots,q-1\}$ that equals $i-j$ modulo $q-1$.

3.2 Acting with diagonal matrices

Let $P : \mathbf{Vec} \to \mathbf{Vec}$ be a polynomial representation and set $V := K^n$, so that we may identify $\mathrm{End}(V)$ with the space of $n \times n$ matrices. Then the monoid $\mathrm{End}(V)$ acts linearly on $P(V)$ via the map $\mathrm{End}(V) \to \mathrm{End}(P(V)), \varphi \mapsto P(\varphi)$, and hence so does its submonoid $D_n \subseteq \mathrm{End}(V)$ of diagonal matrices.

Lemma 3.2.1. *We have*

$$P(V) = \bigoplus_{\chi : D_n \to K} P(V)_\chi$$

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

where χ runs over all monoid homomorphisms $(D_n,\cdot)\to(K,\cdot)$ and where

$$P(V)_\chi := \{p\in P(V)\mid \forall\varphi\in D_n : P(\varphi)p=\chi(\varphi)p\}.$$

Proof. Each element $\varphi\in D_n$ satisfies $\varphi^q=\varphi$, and therefore also $P(\varphi)^q=P(\varphi^q)=P(\varphi)$. Consequently, $P(\varphi)$ is a root of the polynomial $h=T\cdot(T^{q-1}-1)\in K[T]$. This polynomial is square-free, so that $P(\varphi)$ is diagonalisable over a separable closure of K. But also the eigenvalues of $P(\varphi)$ are roots of h, i.e., elements of K, so $P(\varphi)$ is diagonalisable over K. Moreover, all elements of D_n commute, and therefore so do all elements of $P(D_n)$. Hence the latter are all simultaneously diagonalisable. We therefore have

$$P(V)=\bigoplus_{\chi:D_n\to K} P(V)_\chi$$

where, *a priori*, χ runs through all maps $D_n\to K$.

Now if $P(V)_\chi\neq 0$, then it follows that $\chi(\mathrm{diag}(1,\dots,1))=1$ and $\chi(\varphi\psi)=\chi(\varphi)\chi(\psi)$, i.e., χ is a monoid homomorphism $D_n\to K$. □

Note that monoid homomorphisms $D_n\to K$ can be naturally identified with n-tuples of monoid homomorphisms $K\to K$, and hence, by Lemma 3.1.1, with elements of $\{0,\dots,q-1\}^n$. Explicitly, χ is identified with the tuple $(a_1,\dots,a_n)$ if $\chi(\mathrm{diag}(t_1,\dots,t_n))=t_1^{a_1}\cdots t_n^{a_n}$ for all $(t_1,\dots,t_n)\in K^n$.

In analogy with the theory of representations of algebraic groups, we will use the word *weight* for monoid homomorphisms $\chi: D_n\to K$, and we call a vector in $P(V)_\chi$ a *weight vector* of weight χ.

We use the notation $\oplus$ also in this context: if $\chi,\mu\in\{0,\dots,q-1\}^n$ are weights, then $\chi\oplus\mu$ is their componentwise sum with respect to $\oplus$. Note that

$$(\chi\oplus\mu)(\mathrm{diag}(t_1,\dots,t_n))=\chi(\mathrm{diag}(t_1,\dots,t_n))\cdot\mu(\mathrm{diag}(t_1,\dots,t_n)).$$

Example 3.2.2. If U is the subspace of V spanned by the first k basis vectors, then $P(U)$, regarded as a subspace of $P(V)$, is the direct sum of all $P(V)_\chi$ where χ runs over the characters in $\{0,\dots,q-1\}^k\times\{0\}^{n-k}$. In particular, the constant part of P is $P(0)=P(V)_{(0,\dots,0)}$.

LEMMA 3.2.3. *Let $\chi=(a_1,\dots,a_n)\in\{0,\dots,q-1\}^n$ be a weight such that $P(K^n)_\chi$ is nonzero. Then $\sum_i a_i$ is at most $\deg(P)$.*

Proof. Choose a nonzero $p\in P(K^n)_\chi$. Then $P(\mathrm{diag}(t_1,\dots,t_n))p=t_1^{a_1}\cdots t_n^{a_n}p$, and we note that $t_1^{a_1}\cdots t_n^{a_n}$ is a reduced polynomial in $t_1,\dots,t_n$. On the other hand, $P(\mathrm{diag}(t_1,\dots,t_n))$ can be expressed as a reduced polynomial of degree at most $\deg(P)$ in $t_1,\dots,t_n$ with coefficients that are linear maps $P(K^n)\to P(K^n)$. Evaluating this at p yields a reduced polynomial of degree at most $\deg(P)$ in $t_1,\dots,t_n$ whose coefficients are elements of $P(K^n)$. But we already know which polynomial that is, namely $t_1^{a_1}\cdots t_n^{a_n}p$. Hence $\sum_i a_i\le\deg(P)$. □

3.3 Acting with additive one-parameter subgroups

Let $P:\mathbf{Vec}\to\mathbf{Vec}$ be a polynomial representation, $n\in\mathbb{Z}_{\geq 2}$, and $i,j\in[n]$ distinct. Then we have a one-parameter subgroup

$$g_{ij}:(K,+)\to\mathrm{GL}_n(K),\quad g_{ij}(s):=I+sE_{ij},$$

where E_{ij} is the matrix with zeros everywhere except for a 1 in position (i,j). For $b=0,\dots,q-1$ we define the linear map $F_{ij}[b]:P(K^n)\to P(K^n)$ by

$$F_{ij}[b]p=\text{the coefficient of } s^b \text{ in } P(g_{ij}(s))p,$$

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

where we write $P(g_{ij}(s))p$ as a reduced polynomial in s with coefficients in $P(K^n)$.

Lemma 3.3.1. *For any subrepresentation Q of P, the linear space $Q(K^n)$ is stable under $F_{ij}[b]$.*

Proof. Let $p \in Q(K^n)$. Then for all $s \in K$ the element

$$P(g_{ij}(s))p = F_{ij}[0]p + sF_{ij}[1]p + \cdots + s^{q-1}F_{ij}[q-1]p$$

lies in $Q(K^n)$. The Vandermonde matrix $(s^e)_{s\in K, e\in\{0,\ldots,q-1\}}$ is invertible, and this implies that each of the $F_{ij}[e]p$ above are linear combinations of the $P(g_{ij}(s))p$, and therefore in $Q(K^n)$. □

Lemma 3.3.2. *Let $p \in P(K^n)$ be a weight vector of weight $a = (a_1, \ldots, a_n)$, let $b \in \{0, \ldots, q-1\}$, and set $\tilde{p} := F_{ij}[b]p$. Then the following statements hold.*

(1) *We have $\tilde{p} = p$ for $b = 0$.*
(2) *If $a_j = 0$, then $\tilde{p} = 0$ for $b \neq 0$.*
(3) *If $0 < a_j \neq b$, then $\tilde{p}$ is a weight vector of weight $a \ominus (be_j) \oplus (be_i)$.*
(4) *If $0 < a_j = b$, then $\tilde{p}$ is a sum of a weight vector of weight*

$$a \ominus be_j \oplus be_i = (a_1, \ldots, a_i \oplus b, \ldots, q-1, \ldots, a_n)$$

and a weight vector of weight

$$a - be_j \oplus be_i = (a_1, \ldots, a_i \oplus b, \ldots, 0, \ldots, a_n).$$

Proof. We write

$$P(g_{ij}(s))p = p_0 + sp_1 + \cdots + s^{q-1}p_{q-1},$$

where the $p_b \in P(K^n)$ are uniquely determined by the fact that the identity above holds for all $s \in K$.

By setting s equal to zero we obtain $P(g_{ij}(0))p = P(\mathrm{id}_{K^n})p = p$ on the left-hand side, and p_0 on the right-hand side. This proves the first item.

If $a_j = 0$, then

$$P(\mathrm{diag}(1, \ldots, 1, 0, 1 \ldots, 1))p = p$$

where the 0 is in position j. Therefore

$$P(g_{ij}(s))p = P(g_{ij}(s)\,\mathrm{diag}(1, \ldots, 1, 0, 1, \ldots, 1))p = P(\mathrm{diag}(1, \ldots, 1, 0, 1, \ldots, 1))p$$

does not depend on s and hence $F_{ij}[b]p = 0$ for $b \neq 0$.

We now assume $a_j > 0$. We have $F_{ij}[b]p = p_b$. To determine the weight(s) appearing in p_b, we act on p_b with diagonal matrices. For $t = (t_1, \ldots, t_n) \in K^n$ and $t_j \neq 0$ we have

$$\mathrm{diag}(t_1, \ldots, t_n) \cdot g_{ij}(s) = g_{ij}(t_i s t_j^{-1}) \cdot \mathrm{diag}(t_1, \ldots, t_n)$$

and therefore

$$\begin{aligned}\sum_{d=0}^{q-1} s^d P(\mathrm{diag}(t_1, \ldots, t_n))p_d &= P(\mathrm{diag}(t_1, \ldots, t_n) g_{ij}(s))p \\ &= P(g_{ij}(t_i s t_j^{-1})\,\mathrm{diag}(t_1, \ldots, t_n))p \\ &= t_1^{a_1} \cdots t_n^{a_n} \cdot P(g_{ij}(t_i s t_j^{-1}))p \\ &= t_1^{a_1} \cdots t_n^{a_n} \cdot \sum_{d=0}^{q-1} (t_i s t_j^{-1})^d p_d.\end{aligned}$$

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

Comparing coefficients of s^b, we find

$$P(\mathrm{diag}(t))p_b = t^{a-be_j+be_i}p_b$$

for all $t \in K^{j-1} \times K^\times \times K^{n-j} =: D$. Hence p_b is a linear combination of weight vectors with weights that on D agree with the weight $a \ominus be_j \oplus be_i$. If $a_j \neq b$, then there is only one such weight, namely, $a \ominus be_j \oplus be_i$. If $a_j = b$, then there are two such weights, namely, $a \ominus be_j \oplus be_i$ and $a - be_j \oplus be_i$. □

3.4 Spreading out weight

Retaining the notation from § 3.3, suppose we are given a nonzero weight vector $p \in P(K^n)$ of weight $(a_1, \ldots, a_n) \in \{0, \ldots, q-1\}^n$ and a $j \in [n]$ with $a_j > 0$. We construct vectors $\tilde{p} \in P(K^{n+1})$ by identifying p with $P(\iota)p$, where $\iota : K^n \to K^{n+1}$ is the embedding adding a 0 in the last position. Then p is a vector of weight $a = (a_1, \ldots, a_n, 0)$ in $P(K^{n+1})$, and we compute

$$\tilde{p} := F_{n+1,j}[b]p$$

for various b. The vector $\tilde{p}$ is guaranteed to be nonzero for at least two values of b, namely, for $b = 0$ (in which case $\tilde{p} = p$), and, as we will now see, for $b = a_j$. Indeed, in the latter case, by Lemma 3.3.2, $\tilde{p}$ is the sum of a weight vector $\tilde{p}_0$ of weight $a - a_j e_j + a_j e_{n+1}$ and a weight vector $\tilde{p}_1$ of weight $a + (q-1-a_j)e_j + a_j e_{n+1}$.

LEMMA 3.4.1. *In the case where $b = a_j$, we have $\tilde{p}_0 = P((j, n+1))p$, where $(j, n+1)$ is shorthand for the permutation matrix corresponding to the transposition $(j, n+1)$.*

Proof. The vector $\tilde{p}_0$ is obtained by applying $P(\pi_j)$ to $\tilde{p}$, where π_j is the projection $K^{n+1} \to K^{n+1}$ that sets the jth coordinate to zero. Furthermore, we have $P(\pi_{n+1})p = p$, where π_{n+1} sets the $(n+1)$th coordinate to zero. We can then compute $\tilde{p}_0$ as the coefficient of s^{a_j} in

$$\begin{aligned} P(\pi_j)P(g_{n+1,j}(s))p &= P(\pi_j g_{n+1,j}(s)\pi_{n+1})p \\ &= P((j, n+1))P(\mathrm{diag}(1, \ldots, 1, s, 1, \ldots, 1, 0))p \\ &= P((j, n+1))s^{a_j}p. \end{aligned}$$

□

If $F_{n+1,j}[b]p \neq 0$ for some $b \neq 0, a_j$ or if $F_{n+1,j}[a_j]p \neq P((j, n+1))p$, then we find a new vector p' in the subrepresentation of P generated by p whose weight has strictly more nonzero entries; we have *spread out the weight* of p.

DEFINITION 3.4.2. A nonzero weight vector $p \in P(K^n) \subseteq P(K^{n+1})$ of weight $a \in \{0, \ldots, q-1\}^n$ is called *maximally spread out* if for all $j \in [n]$ with $a_j > 0$ we have

$$P(g_{n+1,j}(s))p = p + s^{a_j}P((j, n+1))p.$$

PROPOSITION 3.4.3. *For any nonzero polynomial representation P, there exist an n and a nonzero weight vector $p \in P(K^n)$ that is maximally spread out.*

Proof. Let $p \in P(K^m)$ be a nonzero weight vector. As long as p is not maximally spread out, by the above discussion we can replace p by a nonzero weight vector in $P(K^{m+1})$ whose weight has strictly more nonzero entries. But by Lemma 3.2.3, the number of nonzero entries is bounded from above by $\deg(P)$. Hence this process must terminate, with a maximally spread-out vector. □

Example 3.4.4. It is not true that every polynomial representation is generated by its maximally spread-out vectors. Consider, for instance, K of characteristic 2 and the representation Q that sends V to the space of symmetric tensors in $V \otimes V$. The weight vectors in $Q(K^n)$ are of the forms $e_i \otimes e_i$ and $e_i \otimes e_j + e_j \otimes e_i \in Q(K^n)$ with $i \neq j$. Only the latter are maximally spread out. But they generate the subrepresentation of Q consisting of all skew-symmetric tensors in $V \otimes V$.

3.5 The prime field case

In this section we assume that q is a prime, so that K is a prime field. We retain the notation from above.

Definition 3.5.1. Let $\iota : K^n \to K^{n+1}$ be the standard embedding and $F_{n+1,j} := F_{n+1,j}[1] : P(K^n) \to P(K^{n+1})$ be the operator that sends p to the coefficient of s^1 in $P(g_{n+1,j}(s) \circ \iota)(p)$.

Lemma 3.5.2. *Assume that K is a prime field. Then the operator $F_{n+1,j} : P(K^n) \to P(K^{n+1})$ is injective on the direct sum of all weight spaces corresponding to weights $\chi = (a_1, \ldots, a_n)$ with $a_j > 0$, and it is zero on the weight spaces corresponding to weights with $a_j = 0$.*

Proof. The last part follows immediately from Lemma 3.3.2; we now prove the first part. The operator $F_{n+1,j}$ maps the weight space of χ into that of $\chi \ominus e_j + e_{n+1}$ if $a_j > 1$ and into the sum of the weight spaces with weights $\chi - e_j + e_{n+1}$ and $\chi \ominus e_j + e_{n+1}$ if $a_j = 1$. Since these weights are distinct for distinct χ, it suffices to show that $F_{n+1,j}$ is injective on a single weight space, corresponding to the weight $(a_1, \ldots, a_n)$, where $a_j > 0$. Let p be a nonzero vector in this weight space.

Define $\varphi : K^{n+1} \to K^n$ by

$$\varphi(c_1, \ldots, c_{n+1}) := (c_1, \ldots, c_j + c_{n+1}, \ldots, c_n).$$

We then have

$$\varphi \circ g_{n+1,j}(s) \circ \iota = \mathrm{diag}(1, \ldots, 1 + s, \ldots, 1)$$

and therefore

$$P(\varphi) P(g_{n+1,j}(s)) P(\iota) p = (1 + s)^{a_j} \cdot p.$$

The coefficient of s^1 in the latter expression is $a_j \cdot p$, which is nonzero since $a_j < q$ and q is prime. That coefficient is also equal to $P(\varphi)\tilde{p}$, where $\tilde{p} := F_{n+1,j} p$. Hence $\tilde{p} \neq 0$. □

By Lemma 3.5.2, if $\chi = (a_1, \ldots, a_n)$ with $a_j > 1$, then $F_{n+1,j}$ maps $P(K^n)_\chi$ injectively into $P(K^{n+1})_{\chi'}$, where $\chi' = \chi - e_j + e_{n+1}$. On the other hand, if $a_j = 1$, then by Lemma 3.4.1, $F_{n+1,j}$ followed by the projection to the weight space of $\chi' = (a_1, \ldots, 0, \ldots, a_n, 1)$ agrees on $P(K^n)_\chi$ with the map $P((n+1, j))$, which of course we already knew is injective.

Example 3.5.3. We note that Lemma 3.5.2 is false for nonprime fields. Indeed, take $K = \mathbb{F}_4$ and $P = S^2$. Consider the element $p := e_1^2 \in P(K^1)$, of weight (2). Now $P(g_{21}(s))p = (e_1 + s e_2)^2 = e_1^2 + s^2 e_2^2$, and hence $F_{21} p = 0$. On the other hand, if $K = \mathbb{F}_2$, then $s^2 = s$, and $F_{21} p = e_2^2$.

Remark 3.5.4. Note that, as a consequence of the lemma, if a weight vector p of weight $(a_1, \ldots, a_n)$ is maximally spread out, then $a_j \in \{0, 1\}$ for all j, and moreover $F_{n+1,j} p = P((n+1, j))p$ for all j with $a_j = 1$. Indeed, if $a_j > 1$, then $F_{n+1,j} p$ is a weight vector of weight $(a_1, \ldots, a_j - 1, \ldots, a_n, 1)$, and if $a_j = 1$ but $F_{n+1,j} p \neq P((n+1, j))p$, then the left-hand side has a component of weight $(a_1, \ldots, q - 1, \ldots, a_n, 1)$. In either case, p was not maximally spread out.

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

4. Noetherianity for polynomial representations

4.1 The main result

We recall that the tensor restriction theorem, its generalisation to polynomial representations, and Corollary 1.3.3 concern restriction-closed properties. We will simply use the term *subset* for such a property.

DEFINITION 4.1.1. Let P be a polynomial representation. A *subset* of P is the data of a subset $X(V)$ of $P(V)$ for every $V \in \mathbf{Vec}$, subject to the condition that for all $V, W, \varphi \in \mathrm{Hom}_{\mathbf{Vec}}(V, W)$, $P(\varphi)X(V) \subseteq X(W)$.

THEOREM 4.1.2 (Noetherianity). *Let P be a polynomial representation over the finite field K. Then any descending chain*

$$P \supseteq X_1 \supseteq X_2 \supseteq \cdots$$

of subsets stabilises.

Example 4.1.3. For polynomial functors of the form $P = S^{d_1} \oplus \cdots \oplus S^{d_k}$ with all $d_i < \mathrm{char}\, K$, Theorem 4.1.2 can be derived from [KZ20, Theorem 1.4 and Remark 1.8] as follows. Assume that $d_1 \leq \cdots \leq d_k$. Let $X = X_1$ be a proper subset of P, and consider a descending chain $X_1 \supseteq X_2 \supseteq \cdots$ of subsets. Then by the results in [KZ20], the Schmidt rank of tuples $(p_1, \ldots, p_k)$ in X is uniformly bounded. This means that, for some $i \in \{1, \ldots, k\}$, some l, and some positive integers $e_1, \ldots, e_l < d_i$, X is contained in the image of the map

$$P' := S^{e_1} \oplus S^{d_i - e_1} \oplus \cdots \oplus S^{e_l} \oplus S^{d_i - e_l} \oplus \bigoplus_{j \neq i} S^{d_j} \to S^{d_1} \oplus S^{d_2} \oplus \cdots \oplus S^{d_k},$$

$$(g_1, h_1, \ldots, g_l, h_l, p_1, \ldots, p_{i-1}, p_{i+1}, \ldots, p_k) \mapsto \left(p_1, \ldots, p_{i-1}, \sum_{j=1}^{l} g_j h_j, p_{i+1}, \ldots, p_k\right).$$

Let $X_i' \subseteq P'$ be the pre-image of X_i. Since the tuple of degrees in P' is lexicographically smaller than that in P, we may assume by induction that the chain $X_1' \supseteq X_2' \supseteq \cdots$ stabilises. And since the map $X_i' \to X_i$ is surjective, so does the chain $(X_i)_i$.

The proof for general P uses a similar induction along a well-founded order introduced in § 2.7, but it is considerably more subtle. Moreover, rather than *parameterising* subsets of P by smaller P', we *embed* subsets of P into smaller P'.

4.2 Irreducible decomposition of restriction-closed tensor properties

Before proceeding with the proof of Noetherianity, we deduce from it the fact that any subset of P admits a unique decomposition into irreducible subsets.

DEFINITION 4.2.1. Let P be a polynomial generic representation over the finite field K and let X be a subset of P. We call X *irreducible* if $X(0) \neq \emptyset$ and if whenever X_1, X_2 are subsets of P such that $X(V) = X_1(V) \cup X_2(V)$ holds for all $V \in \mathbf{Vec}$, it follows that $X = X_1$ or $X = X_2$.

THEOREM 4.2.2. *For any subset X of a polynomial representation P over the finite field K, there is a unique decomposition*

$$X = X_1 \cup \cdots \cup X_k$$

where all X_i are irreducible and none is contained in any other.

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

Proof. This is an immediate consequence of Noetherianity (Theorem 4.1.2), and the proof is identical to the proof that any Noetherian topological space admits a unique decomposition into irreducible closed subspaces. □

For an instructive example, we need the following lemma.

Lemma 4.2.3. *Suppose that $P(0)=0$. Then the subset $X:=P$ is irreducible in the sense above.*

We note that the requirement that $P(0)$ be zero is necessary for irreducibility; otherwise, one can take any partition $S_1 \sqcup S_2$ of the finite set $P(0)$ into two nonempty parts, define $X_i(V)$ to be the set of elements in $P(V)$ that map into the S_i, and note that $P=X_1 \cup X_2$.

Proof. Suppose that $X=X_1 \cup X_2$ where $X_i \subsetneq X$ for $i=1,2$. Then X_1 has at least one forbidden restriction $T_1 \in P(V_1)$, and X_2 has at least one forbidden restriction $T_2 \in P(V_2)$. Let $\iota_i : V_i \to V_1 \oplus V_2$ be the canonical inclusion, and write $T := P(\iota_1)(T_1)+P(\iota_2)(T_2)$. Let $\pi_i : V_1 \oplus V_2 \to V_i$ be the projection. Then $P(\pi_1 \circ \iota_2)=P(0_{V_2 \to V_1})$, which is the zero map since it factors via $P(0)=0$; and similarly $P(\pi_2 \circ \iota_1)=0$. We conclude that

$$T_1 = P(\pi_1 \circ \iota_1)T_1 = P(\pi_1)(P(\iota_1)T_1 + P(\iota_2)T_2) = P(\pi_1)T,$$

so T_1 is a restriction of T. Similarly, T_2 is a restriction of T. It follows that T lies neither in $X_1(V_1 \oplus V_2)$ nor in $X_2(V_1 \oplus V_2)$, a contradiction. Hence X is irreducible as claimed. □

Example 4.2.4. For any $r \in \mathbb{R}_{\geq 0}$, let X_r be the locus in T^d where the *partition rank* is at most r, and let Y_r be the locus in T^d where the *analytic rank* is at most r.

For X_r (with r an integer) it is easy to write down a decomposition into irreducible subsets: for any of the $2^{d-1}-1$ unordered partitions $\{I,J\}$ of $[d]$ into two nonempty sets, choose a number $r_{\{I,J\}}$ such that these numbers add up to r. This choice gives a natural map

$$\prod_{\{I,J\}} (T^{|I|} \times T^{|J|})^{r_{\{I,J\}}} \to T^d \quad (f_{I,k}, g_{J,k})_{\{I,J\},k} \mapsto \sum_{\{I,J\}} \sum_{k=1}^{r_{\{I,J\}}} f_{I,k} \otimes g_{J,k}$$

where $f_{I,k}$ and $g_{J,k}$ are tensors in the copies of V labelled by I and J, respectively. This parameterises the locus in X_r of tensors with a partition rank at most r decomposition of a fixed type; this image is irreducible by virtue of Lemma 4.2.3 applied to the left-hand side above. The total number of components of X_r that we find is the number of ways of partitioning r into $2^{d-1}-1$ nonnegative integers, which is polynomial in r.

Given the (almost) linear relation between partition rank and analytic rank [CM23, MZ22], it is natural to ask whether the number of components of Y_r, too, is polynomial in r.

4.3 The vanishing ideal of a subset

We will prove Noetherianity by looking at functions that vanish identically on a subset.

Definition 4.3.1. Given a subset $X \subseteq P$, we denote by $I_X(V) \subseteq K[P(V)]$ the ideal of all functions $P(V) \to K$ that vanish identically on $X(V)$.

We stress that conversely, since K is finite, $X(V)$ is also the set of all common zeros of $I_X(V)$ in $P(V)$.

4.4 Shifting and localising

Definition 2.6.1 can be extended to subsets of polynomial representations.

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

DEFINITION 4.4.1. Given a subset X of a polynomial representation P and a $U \in \mathbf{Vec}$, the shift $\mathrm{Sh}_U X$ is the subset of $\mathrm{Sh}_U P$ defined by $(\mathrm{Sh}_U X)(V) := X(U \oplus V)$.

DEFINITION 4.4.2. Given a subset X of P and a function $h \in K[P(0)]$, we can think of h as a function on any $P(V)$ via pullback along the linear map $P(V) \to P(0)$, and hence also as a function on $X(V)$. We define $X[1/h]$ as the functor

$$V \mapsto X[1/h](V) := \{p \in X(V) \mid h(p) \neq 0\}.$$

Clearly, $X[1/h]$ is a subset of P.

We will often combine a shift and a localisation: given a function $h \in K[P(U)]$, we can think of h as a function on $K[(\mathrm{Sh}_U P)(0)]$, and hence localise.

Example 4.4.3. Let $P: V \to V \otimes V$ and let $X(V)$ be the set of tensors (matrices) of rank at most n in $P(V)$. Let $h \in K[P(K^n)]$ be the $n \times n$ determinant and set $U := K^n$. Then $(\mathrm{Sh}_U X)[1/h]$ is isomorphic to the functor that sends V to $B \times V^{2n}$, where $B := X(U)[1/h]$ is the set of invertible $n \times n$ matrices, and the isomorphism $(\mathrm{Sh}_U X)[1/h](V) \to B \times V^{2n}$ comes from observing that

$$(\mathrm{Sh}_U X)[1/h](V) \subseteq P(U \oplus V) = (K^n \otimes K^n) \times (K^n \otimes V) \times (V \otimes K^n) \times (V \otimes V),$$

and realising that the $V \otimes V$ component of a matrix of rank at most n is completely determined by its remaining three components, provided that the $K^n \otimes K^n$ component has nonzero determinant. This phenomenon, that X becomes an affine space up to shifting and localising, holds in greater generality, at least at a counting level (see Corollary 4.9.1).

4.5 Reduction to the prime field case

PROPOSITION 4.5.1. *Suppose that Theorem 4.1.2 holds when K is a prime field. Then it also holds when K is an arbitrary finite field.*

Proof. Let F be the prime field of K and set $e := \dim_F K$. For an n-dimensional K-vector space U, we write U_F for the $e \cdot n$-dimensional F-vector space obtained by restricting the scalar multiplication on U from $K \times U \to U$ to $F \times U \to U$.

Now let P be a polynomial representation over K. Define a generic representation P_F over F by setting, for a finite-dimensional F-vector space U, $P_F(U) := (P(K \otimes_F U))_F$, and sending an F-linear map $\varphi: U \to V$ to the map $P_F(\varphi) := P(\mathrm{id}_K \otimes \varphi)$, which is K-linear and therefore also F-linear. It is easy to see from the definitions that P_F is polynomial of the same degree as P.

For a subset X of P, we define a subset of P_F via $X_F(U) := X(K \otimes_F U)$. If $X_1 \supseteq X_2 \supseteq \cdots$ is a chain of subsets in P, then $(X_1)_F \supseteq (X_2)_F \supseteq \cdots$ is a chain of subsets in P_F. By assumption, the latter stabilises, say at $(X_{n_0})_F$. Then it follows that, for any $n \geq n_0$ and any m,

$$X_n(K^m) = X_n(K \otimes_F F^m) = (X_n)_F(F^m) = (X_{n_0})_F(F^m) = X_{n_0}(K^m),$$

and this suffices to conclude that $X_n = X_{n_0}$. □

In view of Proposition 4.5.1, from now on we assume that K is a prime field. An important reason for this assumption is that we can then use Lemma 3.5.2. We believe that the proof below can be adapted to arbitrary finite fields, and this might actually give more general results. In particular, in the proof below we will act with the operators $F_{n+1,j} = F_{n+1,j}[1]$; and in the general case we would have to work with the operators $F_{n+1,j}[b]$ for $b \in \{1, \ldots, q-1\}$. But the reasoning below is already rather subtle, and we prefer not to make it more opaque by the additional technicalities coming from nonprime fields.

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

4.6 The embedding theorem

We will prove Theorem 4.1.2 via an auxiliary result of independent interest. Let P be a polynomial representation of positive degree d and let R an irreducible subobject of $P_{>d-1}$. Let $\pi: P \to P/R =: P'$ be the projection. Dually, this gives rise to an embedding $K[P/R] \subseteq K[P]$. For a fixed $V \in \mathbf{Vec}$, if we choose elements $y_1, \ldots, y_n \in P(V)^*$ that map to a basis of $R(V)^*$, then we can write elements of $K[P(V)]$ as reduced polynomials in $y_1, \ldots, y_n$ with coefficients that are elements of $K[P'(V)]$. We note, however, that R is typically not a direct summand of P. This implies, for instance, that when acting with $\mathrm{End}(V)$ on y_i, we typically do not stay within the linear span of the $y_1, \ldots, y_n$ but also get terms that are linear functions in $K[P'(V)]$.

Let X be a subset of P, and let X' be the image of X in P/R, i.e., $X'(V) := \pi(X(V))$ (to simplify notation, we write π instead of π_V).

Now there are two possibilities:

(1) $X = \pi^{-1}(X')$, i.e., $X(V) = \pi^{-1}(X'(V))$ for all V; in this case, I_X is generated by $I_{X'} \subseteq K[P'] \subseteq K[P]$;

(2) There exist a space V and an element $f \in I_X(V)$ such that f does not lie in $K[P] \cdot I_{X'}$.

Theorem 4.6.1 (Embedding theorem). *Assume, as above, that K is a prime field. From any $f \in I_X(V) \setminus K[P(V)] \cdot I_{X'}(V)$, we can construct a $U \in \mathbf{Vec}$ and a polynomial h in $K[P(U)]$ of degree strictly smaller than that of f, such that also h does not vanish identically on $\pi^{-1}(X'(U))$ and such that the projection $\mathrm{Sh}_U P \to (\mathrm{Sh}_U P)/R$ restricts to an injective map on $(\mathrm{Sh}_U X)[1/h]$.*

Here $(\mathrm{Sh}_U X)(V) := X(U \oplus V) \subseteq P(U \oplus V) = (\mathrm{Sh}_U P)(V)$ and $(\mathrm{Sh}_U X)[1/h]$ is the subset of $\mathrm{Sh}_U P$ consisting of points p where $h(p) \neq 0$. A warning here is that h may actually vanish identically on $X(U) \subseteq \pi^{-1}(X'(U))$, in which case the conclusion is trivial because $(\mathrm{Sh}_U X)[1/h]$ is empty. But in our application to the Noetherianity theorem, this will be irrelevant.

We will now first prove Theorem 4.1.2 using the embedding theorem. The proof of the embedding theorem itself is given in § 4.8.

4.7 Proof of Noetherianity from the embedding theorem

Proceeding by induction on P along the partial order from § 2.7, we may assume that Noetherianity holds for every representation $Q \prec P$; we call this the *outer* induction hypothesis.

Let d be the degree of P. If $d = 0$, then $P(V)$ is a fixed finite set independent of V, and clearly any chain of subsets of this set stabilises. So we may assume that $d > 0$.

Let R be an irreducible subrepresentation in the subrepresentation $P_{>d-1}$ of P. Given a subset X of P, we write X' for its projection in $P' := P/R$.

We define $\delta_X \in \{1, 2, \ldots, \infty\}$ as the minimal degree of a polynomial in $I_X \setminus K[P] \cdot I_{X'}$; this is ∞ if $I_X = K[P] \cdot I_{X'}$.

For X, Y subsets of P we write $X > Y$ if we either have a strict inequality $X' \supsetneq Y'$ among the projections of X, Y in P', or else $X' = Y'$ but $\delta_X > \delta_Y$. Since, by the outer induction assumption, P' is Noetherian, this is a well-founded partial order on subsets of P. To prove that a given subset $X \subseteq P$ is Noetherian, we may therefore assume that all subsets $Y \subseteq P$ with $Y < X$ are Noetherian; this is the *inner* induction hypothesis.

Now if $\delta_X = \infty$, then any proper subset Y of X satisfies $Y < X$, so we are done. We are therefore left with the case where $\delta_X \in \mathbb{Z}_{\geq 1}$.

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

Then let $V \in \mathbf{Vec}$ and $f \in I_X(V) \setminus (K[P(V)] \cdot I_{X'}(V))$ be an element of degree δ_X. By the embedding theorem, there exists an element $h \in K[P(U)] \setminus I_X(U)$ of degree $< \delta_X$ such that $(\mathrm{Sh}_U X)[1/h] \to (\mathrm{Sh}_U P)/R$ is an injective map. Since $(\mathrm{Sh}_U P)/R \prec P$ by Lemma 2.7.3, $(\mathrm{Sh}_U X)[1/h]$ is Noetherian by the outer induction hypothesis.

Define Y as the subset of X defined by the vanishing of h. Explicitly,

$$Y(V) := \{p \in X(V) \mid \forall \varphi \in \mathrm{Hom}_{\mathbf{Vec}}(V,U) : h(P(\varphi)p) = 0\}.$$

Let $Y' \subseteq X'$ be the projection of Y in P/R. If $Y' \subsetneq X'$, then $Y < X$ and hence Y is Noetherian by the inner induction hypothesis. If $Y' = X'$, then $h \in I_Y(U) \setminus (K[P(U)] \cdot I_{Y'}(U))$ and hence $\delta_Y \leq \deg(h) < \delta_X$. So then, too, $Y < X$, and Y is Noetherian by the inner induction hypothesis.

Now consider a chain

$$X \supseteq X_1 \supseteq X_2 \supseteq \cdots$$

of subsets. By the above two paragraphs, from some point on both the chain $(X_i \cap Y)_i$ and the chain $((\mathrm{Sh}_U X_i)[1/h])_i$ have stabilised. We claim that then the chain $(X_i)_i$ has also stabilised.

Indeed, take $p \in X_i(W)$. If $p \in X_i(W) \cap Y(W)$, then also $p \in X_{i+1}(W) \cap Y(W)$ by the first chain, and we are done. If not, then let $\varphi : W \to U$ be a linear map such that $h(P(\varphi)p) \neq 0$. Let $\iota : W \to U \oplus W$ be the embedding $w \mapsto (\varphi(w), w)$. Then we find that

$$\begin{aligned} P(\iota)p \in X_i(U \oplus W)[1/h] &= (\mathrm{Sh}_U X_i)(W)[1/h] \\ &= (\mathrm{Sh}_U X_{i+1})(W)[1/h] \subseteq X_{i+1}(U \oplus W). \end{aligned}$$

Now if $\rho : U \oplus W \to W$ is the projection, then we find that $p = P(\rho)P(\iota)p \in P(\rho)X_{i+1}(U \oplus W) = X_{i+1}(W)$, as desired.

4.8 Proof of the embedding theorem

Recall that P has degree $d > 0$, $X \subseteq P$ is a subset, R an irreducible subrepresentation of $P_{>d-1}$, $\pi : P \to P/R$ is the projection, $X' = \pi(X)$, $X \neq \pi^{-1}(X')$, and $f \in I_X(V) \setminus (K[P(V)] \cdot I_{X'}(V))$. Assume that f has degree δ. Recall from Lemma 2.8.2 that $V^* \mapsto K[P(V)]_{\leq \delta}$ is a polynomial representation. Furthermore, this has subrepresentations $V^* \mapsto I_X(V)_{\leq \delta}$ and $V^* \mapsto (K[P(V)] \cdot I_{X'}(V))_{\leq \delta}$.

We may assume that $V = K^n$ and, without loss of generality, f is a weight vector. We will act on f with elements $g_{n+1,j}(s)^T$ (see Example 2.8.3 for an explanation of the transpose). The part of degree b in s is then captured by the operator $F_{n+1,j}[b]$.

After acting repeatedly with operators $F_{n+1,j}[b]$ (for increasing values of n and possibly j and observing that this does not increase the degree of f), we may assume that the image of $f \in K[P(K^n)]$ in the quotient representation

$$I_X(K^n)_{\leq \delta} / (K[P(K^n)] \cdot I_{X'}(K^n))_{\leq \delta}$$

is maximally spread out (see Proposition 3.4.3). After passing to a coordinate subspace, by Remark 3.5.4, this implies that the weight of f is $(1, \ldots, 1)$. Moreover, it implies that if we split, for any $j \in \{1, \ldots, n\}$, $\tilde{f} := F_{n+1,j} f$ as $\tilde{f}_0 + \tilde{f}_1$ where $\tilde{f}_0$ has weight $(1, \ldots, 0, \ldots, 1, 1)$ and $\tilde{f}_1$ has weight $(1, \ldots, q-1, \ldots, 1, 1)$, then $\tilde{f}_1$ vanishes identically on $X'(K^{n+1})$; indeed, otherwise $\tilde{f}_1$ would be a more spread-out polynomial that vanishes identically on X but not on X'.

Choose a basis $\mathbf{x}$ of $P'(K^n)^* \subseteq P(K^n)^*$ consisting of weight vectors, and extend this to a basis $\mathbf{x}, \mathbf{y}$ of $P(K^n)^*$ of weight vectors. This means that $\mathbf{y}$ maps to a weight basis of $R(K^n)^*$. Relative to these choices, we can write f as a reduced polynomial

$$f = \sum_\alpha f_\alpha(\mathbf{x}) \mathbf{y}^\alpha \tag{1}$$

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

for suitable exponent vectors α and nonzero functions $f_\alpha \in K[P'(K^n)]$. We choose this expression *minimal* relative to $I_{X'}(K^n)$ in the following sense: no nonempty subset of the terms of any f_α add up to a polynomial in $I_{X'}(K^n)$. This implies that no f_α is in the ideal of $I_{X'}(K^n)$, but the requirement is a bit stronger than that.

Let y_0 be one of the elements in $\mathbf{y}$ that appears in f; we further choose y_0 such that the support in $\{1,\dots,n\}$ of its weight is inclusionwise minimal. Consider the expression (coarser than (1))

$$f = f_0(\mathbf{x}, \mathbf{y} \setminus \{y_0\})y_0^0 + \cdots + f_c(\mathbf{x}, \mathbf{y} \setminus \{y_0\})y_0^c$$

where every f_e is a reduced polynomial in $\mathbf{x}$ and the variables in $\mathbf{y}$ except for y_0, and where $f_c \neq 0$ and $c \in \{1,\dots,q-1\}$. Note that f_0 is a weight vector of the same weight as f. *A priori*, the coefficients f_e with $e > 0$ need not be weight vectors, since the weight monoid $(\{0,\dots,q-1\}^n, \oplus)$ is not cancellative. However, all terms in f_e have the same weight up to identifying 0 and $q-1$, and upon adding e times the weight of y_0 to any of the weights of a term in f_e (using the operation $\oplus$), one obtains the weight $(1,\dots,1)$ of f.

Lemma 4.8.1. *We have $c = 1$, f_1 is a weight vector, and after a permutation f_1 has weight $(1^m, 0^{n-m})$ and y_0 has weight $(0^m, 1^{n-m})$ for some m.*

Proof. To prove the claim, let $j \in [n]$ be such that the weight $\chi = (a_1,\dots,a_n)$ of y_0 has $a_j > 0$.

We partition the variables $\mathbf{y}$ into three subsets: those whose weight has an entry 0 in position j are collected in the tuple $\mathbf{y}_0$; those with a 1 in position j in the tuple $\mathbf{y}_1$; and those with an entry greater than 1 there in the tuple $\mathbf{y}_{>1}$.

We construct a weight basis of $P(K^{n+1})^*$ consisting of:

- $\mathbf{x}, \mathbf{y}_0, \mathbf{y}_1$, and $\mathbf{y}_{>1}$;
- the tuple $(n+1,j)\mathbf{y}_1$ obtained by applying $(n+1,j)$ to each variable in $\mathbf{y}_1$ (we suppress in this notation the polynomial generic representation in which these variables live);
- the tuple $F_{n+1,j}\mathbf{y}_{>1}$ obtained by applying $F_{n+1,j}$ to each variable in the tuple $\mathbf{y}_{>1}$;
- weight elements that together with $\mathbf{x}$ form a basis of $P'(K^{n+1})^*$; and
- weight elements that along with $\mathbf{y}_0, \mathbf{y}_1, \mathbf{y}_{>1}, (n+1,j)\mathbf{y}_1, F_{n+1,j}\mathbf{y}_{>1}$ project to a weight basis of $R(K^{n+1})^*$.

The only nonobvious thing here is that the elements in $F_{n+1,j}\mathbf{y}_{>1}$ can be chosen as part of a set mapping to a basis of $R(K^{n+1})^*$, and this follows from Lemma 3.5.2. Note that none of the variables in the tuples $(n+1,j)\mathbf{y}_1$ and $F_{n+1,j}\mathbf{y}_{>1}$ has a $q-1$ in position j of its weight.

Either y_0 belongs to $\mathbf{y}_1$ or to $\mathbf{y}_{>1}$. In the first case we define $y_1 := (n+1,j)y_0$, and in the second case we define $y_1 := F_{n+1,j}y_0$. In both cases, y_1 is the (nonzero) weight-$(a_1,\dots,a_j-1,\dots,a_n,1)$ component of $F_{n+1,j}y_0$ and one of the chosen variables. (In the first case, this uses Lemma 3.4.1.)

Consider

$$g_{n+1,j}(s)^T f = \sum_{e=0}^{c} (f_e(g_{n+1,j}(s)^T \mathbf{x}, g_{n+1,j}(s)^T(\mathbf{y} \setminus \{y_0\})))(g_{n+1,j}(s)^T y_0)^e, \tag{2}$$

where on both sides we have suppressed the polynomial generic representations where f and the coordinate tuples live (see Lemma 2.8.2 and Example 2.8.3). From the term with $e = c$ we get a contribution

$$c \cdot (f_c + F_{n+1,j}[q-1]f_c) \cdot s \cdot y_0^{c-1} \cdot y_1. \tag{3}$$

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

There is no cancellation in the $+$ in (3): the weights of monomials in f_c are distinct from the weights of monomials in $F_{n+1,j}[q-1]f_c$, because the latter have a positive entry on position $n+1$. Furthermore, we have $0<c<q$ and q is prime, so the term $c\cdot f_c\cdot s\cdot y_0^{c-1}\cdot y_1$ is nonzero.

Rewriting (2) as a reduced polynomial in s,y_0,y_1 with coefficients that are reduced polynomials in the remaining chosen variables in $P(K^{n+1})^*$, we claim that the coefficient of $s\cdot y_0^{c-1}\cdot y_1$ is precisely that in (3). Indeed, y_0,y_1 only appear in the terms $y_0=F_{n+1,j}[0]y_0$ and $F_{n+1,j}y_0$ from $g_{n+1,j}(s)^Ty_0$ and nowhere in $f_e(g_{n+1,j}(s)^T\mathbf{x}, g_{n+1,j}(s)^T(\mathbf{y}\setminus\{y_0\}))$ or in $F_{n+1,j}[b]y_0$ with $b>0$ because:

- $g_{n+1,j}(s)^T$ maps the coordinates $\mathbf{x}$ into linear combinations of $\mathbf{x}$ and the further chosen variables in $P'(K^{n+1})^*$;
- y_0,y_1 do not appear in $F_{n+1,j}[b]\mathbf{y}$ for $b>1$ for weight reasons; expressing the elements in the latter tuple on the basis of the chosen variables, all variables have weights with a $b>1$ at position $n+1$, while y_0,y_1 have a 0 and 1 there, respectively;
- y_0 does not appear in $F_{n+1,j}y$ for any variable y in $\mathbf{y}$, again by comparing the weights in position $n+1$;
- y_1 is different from all variables $F_{n+1,j}y$ where y ranges over the variables in $\mathbf{y}_{>1}$ (other than y_0, if y_0 is in $\mathbf{y}_{>1}$);
- y_1 is different from all variables $(n+1,j)y$ where y ranges over the variables in $\mathbf{y}_1$ (other than y_0, if y_0 is in $\mathbf{y}_1$); and indeed,
- y_1 does not appear in the weight component $y'=F_{n+1,j}y-(n+1,j)y$ of any variable y in $\mathbf{y}_1$. Indeed, if $(a_1',\ldots,1,\ldots,a_n')$ is the weight of y, then y' has weight $(a_1',\ldots,q-1,\ldots,a_n',1)$. But, as remarked earlier, the variable y_1 constructed from y_0 does not have a $q-1$ on position j in its weight.

This proves the claim. We conclude that, when writing $\tilde{f}=F_{n+1,j}f$ as a polynomial in all of the chosen variables, the terms divisible by $y_0^{c-1}y_1$ are precisely those in

$$c\cdot(f_c+F_{n+1,j}[q-1]f_c)\cdot y_0^{c-1}y_1.$$

Now in f_c, expanded as a reduced polynomial in $\mathbf{y}\setminus\{y_0\}$ with coefficients that are reduced polynomials in $\mathbf{x}$, consider any nonzero term $\ell(\mathbf{x})\cdot(\mathbf{y}\setminus\{y_0\})^\alpha$.

Group the terms in $\ell(\mathbf{x})$ into two parts: $\ell(\mathbf{x})=\ell_0(\mathbf{x})+\ell_1(\mathbf{x})$, in such a manner that $\ell_0(\mathbf{x})\cdot(\mathbf{y}\setminus\{y_0\})^\alpha y_0^{c-1}y_1$ is the part of $\ell(\mathbf{x})\cdot(\mathbf{y}\setminus\{y_0\})^\alpha y_0^{c-1}y_1$ that has weight $(1,\ldots,0,\ldots,1,1)$ and hence is part of $(1/c)\tilde{f}_0$; and $\ell_1(\mathbf{x})\cdot(\mathbf{y}\setminus\{y_0\})^\alpha y_0^{c-1}y_1$ has weight $(1,\ldots,q-1,\ldots,1,1)$ and hence is part of $(1/c)\tilde{f}_1$.

For the same exponent vector α, $F_{n+1,j}[q-1]f_c$ may also contain a term $\tilde{\ell}\cdot(\mathbf{y}\setminus\{y_0\})^\alpha$, where $\tilde{\ell}$ is a polynomial in $\mathbf{x}$ and the remaining chosen coordinates on $P'(K^{n+1})$. We can similarly decompose $\tilde{\ell}=\tilde{\ell}_0+\tilde{\ell}_1$, where $\tilde{\ell}_0\cdot(\mathbf{y}\setminus\{y_0\})^\alpha y_0^{c-1}y_1$ is part of $(1/c)\tilde{f}_0$ and $\tilde{\ell}_1\cdot(\mathbf{y}\setminus\{0\})^\alpha y_0^{c-1}y_1$ is part of $(1/c)\tilde{f}_1$.

Since $\tilde{f}_1$ vanishes identically on $\pi^{-1}(X'(K^{n+1}))$ (this was the point of choosing f such that its image in a suitable representation is maximally spread out) we find that $\ell_1(\mathbf{x})+\tilde{\ell}_1$ vanishes identically on $X'(K^{n+1})$. Furthermore, since the weight of every monomial in $F_{n+1,q}[q-1]f_c$ has a positive entry in position $n+1$ and the weight of $(\mathbf{y}\setminus\{y_0\})^\alpha$ has a zero in that position, the weight of every term in $\tilde{\ell}_1$ has a positive entry in position $n+1$. This implies that $\tilde{\ell}_1$ vanishes identically on the linear subspace $P'(K^n)\subseteq P'(K^{n+1})$. We conclude that $\ell_1(\mathbf{x})$ vanishes identically on $X'(K^n)$. However, by minimality of the expression for f relative to $I_{X'}(K^n)$, no nonempty subset of the terms of $\ell(\mathbf{x})$ add up to a polynomial that vanishes identically on $X'(K^n)$. So we have $\ell_1(\mathbf{x})=0$ and $\ell(\mathbf{x})=\ell_0(\mathbf{x})$. It follows that $\ell(\mathbf{x})(\mathbf{y}\setminus\{y_0\})^\alpha y_0^{c-1}y_1$ is a weight vector

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

of weight $(1,\dots,0,\dots,1,1)$. Since the term $\ell(\mathbf{x})(\mathbf{y}\setminus\{y_0\})^\alpha$ in $f_c(\mathbf{x},\mathbf{y}\setminus\{y_0\})$ was arbitrary, we find that $f_c y_0^{c-1} y_1$ is a weight vector of weight $(1,\dots,0,\dots,1,1)$. As the weight of y_0 has a positive entry in position j, we find that $c-1=0$ and all weights appearing in f_c have a 0 in position j.

Now j was arbitrary in the support of the weight of y_0, so the weights appearing in f_c all have disjoint support from that of y_0. But the only way, in the weight monoid $\{0,1,\dots,q-1\}^n$, to obtain the weight $(1,\dots,1)$ as a $\oplus$-sum of two weights with disjoint supports is if, after a permutation, one weight is $(1^m,0^{n-m})$ and the other weight is $(0^m,1^{n-m})$. Hence f_c is a weight vector that, after that permutation, has the former weight, and then y_0 has the latter.

Now we have found that

$$f = f_0 + f_1\cdot y_0$$

where f_1 does not vanish identically on $\pi^{-1}(X'(K^n))$; f_1 has weight $(1^m,0^{n-m})$, y_0 has weight $\chi=(0^m,1^{n-m})$, and f_0 does not involve y_0. It might be, though, that f_0 still contains other variables y in $\mathbf{y}$ of the same weight $\chi=(0^m,1^{n-m})$. Therefore, among the $\mathbf{y}$-variables, let $y_0=\hat{y}_1,\hat{y}_2,\dots,\hat{y}_N$ be those that have weight equal to χ; so N is the multiplicity of χ in $R(K^n)^*$. Then the above implies that

$$f = \hat{f}_1\hat{y}_1 + \cdots + \hat{f}_N\hat{y}_N + r \tag{4}$$

where each $\hat{f}_i$ has weight $(1^m,0^{n-m})$ and where the $\mathbf{y}$-variables that appear in r have weights with at least one nonzero entry in the first m positions (here we use that y_0 had a weight vector of minimal support). Note that $\hat{f}_1$ equals f_1, and does not vanish identically on $\pi^{-1}(X'(K^n))$.

Now set $U:=K^m$, $W:=K^{n-m}$, and $h:=\hat{f}_1$. Note that $h\in K[P(U)]$, since its weight is $(1^m,0^{n-m})$. Also, h has lower degree than f, as desired, and does not vanish identically on $\pi^{-1}(X'(U))$. In fact, all $\hat{f}_i$ are polynomials in $K[P(U)]$, the $\hat{y}_i$ map to coordinates on $R(W)$, and r is a polynomial in $K[(\mathrm{Sh}_U P)(W)/R(W)]$ because every $\mathbf{y}$-variable in r has at least one nonzero entry among the first m entries of its weight.

We claim that $(\mathrm{Sh}_U X)[1/h]\to(\mathrm{Sh}_U P)/R$ is injective. We first show that this is the case when evaluating at $W=K^{n-m}$. Consider two points $p,p'\in(\mathrm{Sh}_U X)[1/h](W)$ with the same projection in $(\mathrm{Sh}_U P)(W)/R(W)$, so that $p-p'\in R(W)$. Then f vanishes at both p and p' and, in (4), we have $\hat{f}_i(p)=\hat{f}_i(p')=:c_i\in K$ for all i, as well as $r(p)=r(p')$. Then (4) shows that

$$c_1\hat{y}_1(p)+\cdots+c_N\hat{y}_N(p)=c_1\hat{y}_1(p')+\cdots+c_N\hat{y}_N(p').$$

This can be expressed as $L(p-p')=0$ for a linear form $L\in R(W)^*$ which is nonzero because $c_1=h(p)=h(p')\neq 0$. Now act with an element $\psi\in\mathrm{End}(K^{n-m})$ on (4), and then substitute p and p'. This yields the identity $L(R(\psi)(p-p'))=0$. Hence we obtain a nonzero $\mathrm{End}(K^{n-m})$-submodule of linear forms in $R(K^{n-m})^*$ that are zero on $p-p'$. But since $R(K^{n-m})$, and hence $R(K^{n-m})^*$, are irreducible $\mathrm{End}(K^{n-m})$-modules by Remark 2.4.3, this means that $p-p'=0$. The same argument applies when W is replaced by K^s for any s. This completes the proof of the embedding theorem. □

4.9 The weak shift theorem

The embedding theorem can be used to show that the behaviour of Example 4.4.3 is typical.

Corollary 4.9.1 (Weak shift theorem). *Suppose, as in the embedding theorem, that K is a prime field of cardinality q. For any nonempty subset X of some polynomial generic representation P, there exist a $U\in\mathbf{Vec}$, a nonzero function h on $X(U)$, and a polynomial $A(n)\in\mathbb{Q}[n]$ such that for all $n\in\mathbb{N}$ the cardinality of $(\mathrm{Sh}_U X)[1/h](K^n)$ equals $q^{A(n)}$.*

In other words, at least in a counting sense, $(\mathrm{Sh}_U X)[1/h](K^n)$ is an affine space of dimension $A(n)$ over K. We expect there to be a stronger version of this theorem, similar to the shift theorem in [BDES23], which says that this affine space is functorial in V. But we do not yet know the precise statement of this stronger theorem. Note that, by applying the weak shift theorem to the subset Y of X defined by the vanishing of h, and so on, we obtain a kind of stratification of X by finitely many affine spaces. A stronger version of the weak shift theorem would therefore give deeper insight into the geometric structure of general restriction-closed properties of tensors.

Proof of the weak shift theorem from the embedding theorem. If P has degree 0, then $X = X(0)$ is a finite set, and we can choose $U = 0$ and h to vanish on all but one point of $X(0)$, so that $(\mathrm{Sh}_U X)[1/h]$ is that remaining point.

Now assume that P has degree $d > 0$ and that the result holds for all polynomial representations $Q \prec P$. Let R be an irreducible subobject of $P_{>d-1}$ and let X' be the projection of X in $P' := P/R$.

There are two cases. First assume that X is the pre-image of X'. Since $P' \prec P$, by the induction assumption there exist a U and an $h \in P'(U)$ that does not vanish on $X'(U)$ such that $|(\mathrm{Sh}_U X')[1/h](K^n)| = q^{A(n)}$ for some polynomial $A(n)$. Now $(\mathrm{Sh}_U X)[1/h](K^n)$ is the pre-image of $(\mathrm{Sh}_U X')[1/h](K^n)$, with fibres $(\mathrm{Sh}_U R)(K^n)$. The fibre is a finite-dimensional vector space over K whose dimension is a polynomial $B(n)$. Hence $|(\mathrm{Sh}_U X)[1/h'](K^n)| = q^{A(n)+B(n)}$, as desired.

If X is not the pre-image of X', then we have seen that there exist a $U_1 \in \mathbf{Vec}$, a polynomial $h_1 \in K[P(U_1)]$ that does not vanish on X, and an injection

$$(\mathrm{Sh}_{U_1} X)[1/h_1] \to ((\mathrm{Sh}_{U_1} P)/R) =: Q.$$

Let Y be the image of this injection. Since $Q \prec P$, there exist $U_2 \in \mathbf{Vec}$ and $h_2 \in K[Q(U_2)]$ such that $|(\mathrm{Sh}_{U_2} Y)[1/h_2](K^n)| = q^{A(n)}$ for some polynomial n. Now set $U := U_1 \oplus U_2$ and $h := h_1 \cdot h_2$ and we find that

$$|(\mathrm{Sh}_U X)[1/h](K^n)| = q^{A(n)},$$

as desired. □

5. Relations to FI and algorithms

5.1 FI and testing properties via subtensors

We recall from [CEF15] that **FI** is the category of finite sets with injections and that an **FI**-module over K is a functor from **FI** to **Vec**. The central result that we will use is the following theorem.

Theorem 5.1.1 [CEF15]. *For any field K, any finitely generated **FI**-module M over K is Noetherian in the sense that every **FI**-submodule is finitely generated.*

We now use this to establish items (4) and (5) in Corollary 1.3.3. Their generalisation to arbitrary polynomial generic representations is as follows.

Theorem 5.1.2. *Let P be a polynomial generic representation over the finite field K, and let $X \subseteq P$ be a subset. Then there exists an n_0 such that for any $n \in \mathbb{Z}_{\geq 0}$, an element $p \in P(K^n)$ lies in $X(K^n)$ if and only if, for every subset S of $[n]$ of size n_0, the image of p in $P(K^S)$ under the linear map corresponding to the coordinate projection $K^n \to K^S$ lies in $X(K^S)$.*

Consequently, there exists a polynomial-time algorithm that on input an $n \in \mathbb{Z}_{\geq 0}$ and a $T \in P(K^n)$ decides whether T lies in $X(K^n)$.

Proof. Noetherianity for subsets of P (Theorem 4.1.2) implies that the ideal I_X is finitely generated. In particular, I_X is generated by $(I_X)_{\leq e}$ for some finite degree e. Now consider the functor F from **FI** to **Vec** that assigns to any finite set S the space $K[P(K^S)]_{\leq e}$ and to every injection $\iota : S \to T$ the embedding $K[P(K^S)]_{\leq e} \to K[P(K^T)]_{\leq e}$ coming from the pullback along the linear map $P(K^T) \to P(K^S)$ associated to ι. Since weights in $K[P(K^S)]_{\leq e}$ have at most de nonzero entries, where $d = \deg(P)$ (see Lemma 2.8.2), F is generated by $F([de])$, hence a finitely generated **FI**-module. By Theorem 5.1.1, the **FI**-submodule $S \mapsto I_X(K^S)$ is also finitely generated, say by $I_X(K^{n_0})$. This n_0 has the desired property.

Now for the last statement. By the above, testing whether T lies in $X(K^n)$ boils down to testing whether the image of T in $P(K^S)$ lies in $X(K^S)$ for each subset $S \subseteq [n]$ of size n_0. Each of these tests takes constant time, and there are $\binom{n}{n_0}$ of these subsets, which is a degree-n_0 polynomial in n. □

Remark 5.1.3. In the latter theorem and proof, we have not specified how elements of $P(K^n)$ are represented on a computer, and it can depend on this representation how fast the image of T in $P(K^S)$ is computed. But since P is a subquotient of a generic representation of the form $V \mapsto V^{\otimes d_1} \oplus \cdots \oplus V^{\otimes d_k}$ (Theorem 1.5.3), a natural such representation is as a k-tuple of arrays. Then computing the image boils down to looking up the entries of T in positions all of whose indices are in S. In this representation, the algorithm explained in the last paragraph is polynomial-time even in the setting where each tensor is given as a *sparse* array, i.e., as a list of pairs consisting of a position $(i_1, \ldots, i_{d_l})$ and an entry in K.

5.2 Infinite tensors

We conclude with a theorem about infinite tensors. Let $P : \mathbf{Vec} \to \mathbf{Vec}$ be a polynomial generic representation over the finite field K. Define

$$P_\infty := \lim_{\leftarrow n} P(K^n)$$

where the limit is along the projections $P(K^{n+1}) \to P(K^n)$ coming from the projections $K^{n+1} \to K^n$ forgetting the last entry. In the case where $P(V) = V^{\otimes d}$, P_∞ can be thought of as the space of $\mathbb{N} \times \cdots \times \mathbb{N}$ tensors (with d factors $\mathbb{N}$). The space P_∞ carries the inverse limit of discrete topologies, or, equivalently, the Zariski topology in which closed subsets are defined by the vanishing of (possibly infinitely many) functions in the ring

$$K[P_\infty] := \lim_{n \to \infty} K[P(K^n)].$$

The monoid Π of matrices that differ from the identity matrix only in finitely many positions acts on P_∞ and $K[P_\infty]$. Let $X \subseteq P$ be a subset in the sense of Definition 4.1.1. Then $X_\infty := \lim_{\leftarrow n} X(K^n)$ is a subset of P_∞.

Proposition 5.2.1. *The correspondence that sends $X \subseteq P$ to $X_\infty \subseteq P_\infty$ is a bijection between subsets of P and closed, Π-stable subsets of P_∞.*

Proof. The subset X_∞ is clearly closed and Π-stable. Conversely, let $Y \subseteq P_\infty$ be closed and Π-stable. Define $Y_n \subseteq P(K^n)$ as the image of Y under the projection $P_\infty \to P_n$, and for any $V \in \mathbf{Vec}$ define $X(V)$ to be the image of Y_n under $P(\varphi)$ for any linear isomorphism $K^n \to V$; this is independent of the choice of φ.

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

We claim that X is a subset of P. Indeed, if $\psi : V \to W$ is any linear map and $\varphi : K^n \to V$ a bijection, then we have to show that $P(\psi)P(\varphi)Y_n$ is contained in $P(\varphi')Y_m$ where $\varphi' : K^{n'} \to W$ is a bijection. Now the map $(\varphi')^{-1}\psi\varphi : K^n \to K^{n'}$ extends to a linear map $\alpha : K^m \to K^m$, where $m := \max\{n, n'\}$ and we regard K^n and $K^{n'}$ as the subspaces of K^m where the last $m-n$ and $m-n'$ entries, respectively, are zero. The map α, in turn, can be regarded an element of Π. Consequently, Y is invariant under α, and therefore so is Y_m. This shows that the map $P((\varphi')^{-1}\psi\varphi)$ maps Y_n into $Y_{n'}$, so that $P(\psi)$ maps $X(V)$ into $X(W)$, as desired.

Next we claim that $X_\infty = Y$. Indeed, the fact that Y is closed means precisely that to test whether $(y_0, y_1, y_2, \ldots)$ lies in Y, it suffices to check whether y_n lies in Y_n for all y. And on the other hand, this is precisely the definition of X_∞. □

Theorem 4.1.2 now implies the following corollary.

Corollary 5.2.2. *Closed, Π-stable subsets of P_∞ satisfy the descending chain condition. Dually, Π-stable ideals of $K[P_\infty]$ satisfy the ascending chain condition.*

The latter, ring-theoretic Noetherianity is not known in the context of infinite fields [Dra19], except in characteristic 0 for a few polynomial functors of degree 2 [NSS16, SS22]. However, in the current context, all ideals are radical, since in $K[P_\infty]$ every element f satisfies the identity $f^q = f$. This implies that the two statements in the corollary are equivalent.

Acknowledgements

We thank Nate Harman and Andrew Snowden for useful discussions and for references to relevant literature.

Conflicts of interest

None.

Financial support

AB and FR are supported by Swiss National Science Foundation (SNSF) project grant 200021_191981, and JD is partially supported by that grant and partially supported by Vici grant 639.033.514 from the Netherlands Organisation for Scientific Research (NWO). We thank the referees for numerous suggestions for improvement. JD thanks the Institute for Advanced Study for the excellent working conditions, under which part of this project was carried out.

Journal information

Compositio Mathematica is owned by the Foundation Compositio Mathematica and published by the London Mathematical Society in partnership with Cambridge University Press. All surplus income from the publication of *Compositio Mathematica* is returned to mathematics and higher education through the charitable activities of the Foundation, the London Mathematical Society and Cambridge University Press.

References

AKZ21 K. Adiprasito, D. Kazhdan and T. Ziegler. *On the Schmidt and analytic ranks for trilinear forms*. Preprint (2021), arXiv:2102.03659.

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

BDD23 A. Bik, A. Danelon and J. Draisma. *Topological Noetherianity of polynomial functors II: Base rings with Noetherian spectrum.* Math. Ann. **385** (2023), 1879–1921.

BDE19 A. Bik, J. Draisma and R. H. Eggermont. *Polynomials and tensors of bounded strength.* Commun. Contemp. Math. **21** (2019), 1850062.

BDES23 A. Bik, J. Draisma, R. H. Eggermont and A. Snowden. *The geometry of polynomial representations.* Int. Math. Res. Not. **2023** (2023), 14131–14195.

CEF15 T. Church, J. S. Ellenberg and B. Farb. *FI-modules and stability for representations of symmetric groups.* Duke Math. J. **164** (2015), 1833–1910.

CGL$^+$21 A. Conner, F. Gesmundo, J. M. Landsberg, E. Ventura and Y. Wang. *Towards a geometric approach to Strassen's asymptotic rank conjecture.* Collect. Math. **72** (2021), 63–86.

CGZ23 M. Christandl, F. Gesmundo and J. Zuiddam. *A gap in the subrank of tensors.* SIAM J. Appl. Algebra Geom. **7** (2023), 742–767.

CM23 A. Cohen and G. Moshkovitz. *Partition and analytic rank are equivalent over large fields.* Duke Math. J. **172** (2023), 2433–2470.

Dra19 J. Draisma. *Topological Noetherianity of polynomial functors.* J. Amer. Math. Soc. **32** (2019), 691–707.

FS97 E. M. Friedlander and A. Suslin. *Cohomology of finite group schemes over a field.* Invent. Math. **127** (1997), 209–270.

GT09 B. Green and T. Tao. *The distribution of polynomials over finite fields, with applications to the Gowers norms.* Contrib. Discrete Math. **4** (2009), 1–36.

Kar22 T. Karam. *High-rank subtensors of high-rank tensors*, Preprint (2022), arXiv:2207.08030.

Kuh94a N. J. Kuhn. *Generic representations of the finite general linear groups and the Steenrod algebra. I.* Amer. J. Math. **116** (1994), 327–360.

Kuh94b N. J. Kuhn. *Generic representations of the finite general linear groups and the Steenrod algebra. II.* K-Theory **8** (1994), 395–428.

KZ20 D. Kazhdan and T. Ziegler. *Properties of high rank subvarieties of affine spaces.* Geom. Funct. Anal. **30** (2020), 1063–1096.

Lov19 S. Lovett. *The analytic rank of tensors and its applications.* Discrete Anal. **2019** (2019), 10.

Mil19 L. Milićević. *Polynomial bound for partition rank in terms of analytic rank.* Geom. Funct. Anal. **29** (2019), 1503–1530.

MZ22 G. Moshkovitz and D. G. Zhu. *Quasi-linear relation between partition and analytic rank.* Preprint (2022), arXiv:2211.05780.

NSS16 R. Nagpal, S. V. Sam and A. Snowden. *Noetherianity of some degree two twisted commutative algebras.* Selecta Math. (N.S.) **22** (2016), 913–937.

PS17 A. Putman and S. V. Sam. *Representation stability and finite linear groups.* Duke Math. J. **166** (2017), 2521–2598.

RS04 N. Robertson and P. D. Seymour. *Graph minors. XX: Wagner's conjecture.* J. Comb. Theory, Ser. B **92** (2004), 325–357.

SS17 S. V. Sam and A. Snowden. *Gröbner methods for representations of combinatorial categories.* J. Amer. Math. Soc. **30** (2017), 159–203.

SS22 S. V. Sam and A. Snowden. Sp-*equivariant modules over polynomial rings in infinitely many variables.* Trans. Amer. Math. Soc. **375** (2022), 1671–1701.

Sno21 A. Snowden. *Stable representation theory: beyond the classical groups*, Preprint (2021), arXiv:2109.11702.

https://doi.org/10.1112/S0010437X25007067 Published online by Cambridge University Press

Ste16 R. Steinberg. *Lectures on Chevalley groups*, University Lecture Series, vol. 66 (American Mathematical Society, Providence, RI, 2016).

Andreas Blatter andreas.blatter@unibe.ch
Mathematical Institute, University of Bern, Alpeneggstrasse 22, 3012 Bern, Switzerland

Jan Draisma jan.draisma@unibe.ch
Mathematical Institute, University of Bern, Sidlerstrasse 5, 3012 Bern, Switzerland; and Department of Mathematics and Computer Science, P.O. Box 513, 5600 MB, Eindhoven, the Netherlands

Filip Rupniewski filip.rupniewski@unibe.ch
Mathematical Institute, University of Bern, Alpeneggstrasse 22, 3012 Bern, Switzerland