J. Eur. Math. Soc. 17, 711–738

DOI 10.4171/JEMS/515

Jan Draisma · Rob H. Eggermont

Finiteness results for Abelian tree models

Received July 5, 2012 and in revised form August 13, 2013

Abstract. Equivariant tree models are statistical models used in the reconstruction of phylogenetic trees from genetic data. Here *equivariant* refers to a symmetry group imposed on the root distribution and on the transition matrices in the model. We prove that if that symmetry group is Abelian, then the Zariski closures of these models are defined by polynomial equations of bounded degree, independent of the tree. Moreover, we show that there exists a polynomial-time membership test for that Zariski closure. This generalises earlier results on tensors of bounded rank, which correspond to the case where the group is trivial and the tree is a star, and implies a qualitative variant of a quantitative conjecture by Sturmfels and Sullivant in the case where the group and the alphabet coincide. Our proofs exploit the symmetries of an infinite-dimensional projective limit of Abelian star models.

Keywords. Phylogenetic tree models, tensor rank, noetherianity up to symmetry, applied algebraic geometry

1. Introduction

Tree models are families of probability distributions used in modelling the evolution of a number of extant species from a common ancestor. Here *species* can refer to actual biological species, but tree models have also been applied to other forms of evolution, e.g. of languages. The hypothesis underlying tree models is that DNA-sequences of those extant species, arranged and suitably aligned in a table with one row for each species, can be meaningfully read off column-wise. Indeed, these columns (or *sites*) are assumed to be independent draws from one and the same probability distribution belonging to the model.

To describe that model, one fixes a finite rooted tree T whose leaves correspond to the species and whose root r corresponds to the common ancestor. One also fixes a finite alphabet B. The case where $B = \{A, C, G, T\}$ is the alphabet of nucleotides is of most interest in biology, but the theory developed here works for arbitrary finite B. Associated to each vertex of the tree is a copy of B. To r one attaches a probability distribution π

J. Draisma: Department of Mathematics and Computer Science,
Technische Universiteit Eindhoven, P.O. Box 513, 5600 MB Eindhoven, The Netherlands,
and VU and CWI, Amsterdam, The Netherlands; e-mail: j.draisma@tue.nl

R. H. Eggermont: Department of Mathematics and Computer Science, Technische Universiteit Eindhoven, P.O. Box 513, 5600 MB Eindhoven, The Netherlands; e-mail: r.h.eggermont@tue.nl

Mathematics Subject Classification (2010): Primary 13E05; Secondary 14M12, 15A69, 62P10

on B, and to each edge $q \to q'$, directed away from r, one attaches a $B \times B$-matrix $A_{qq'}$ of real nonnegative numbers whose row sums equal 1. Its entry $A_{qq'}(b, b')$ at position (b, b') records the probability that the letter b at vertex q mutates into the letter b' at vertex q'. The random process modelling evolution of the nucleotide at a single position consists of drawing a letter $b \in B$ from the distribution π and mutating it along the edges with the probabilities given by the matrices $A_{qq'}$. The probability that this leads to a given word $\mathbf{b} \in B^{\mathrm{leaf}(T)}$ equals

$$P(\mathbf{b}) = \sum_{\mathbf{b}' \in B^{\mathrm{vert}(T)} \text{ extending } \mathbf{b}} \pi(b'_r) \cdot \prod_{q \to q' \in \mathrm{edge}(T)} A_{qq'}(b'_q, b'_{q'}).$$

Now as the root distribution π and the transition matrices $A_{qq'}$ vary, the set of all probability distributions $P \in \mathbb{R}^{(B^{\mathrm{leaf}(T)})}$ thus obtained is called the *model*. The fact that the entries of P are polynomial functions of the parameters has led to an extensive study of the *algebraic variety* swept out by this parameterisation, by which we mean the Zariski closure in $\mathbb{R}^{(B^{\mathrm{leaf}(T)})}$ (or even $\mathbb{C}^{(B^{\mathrm{leaf}(T)})}$) of the model [PS05, Chapter 4]; see also the expository paper [Cip07]. The present paper also concerns that Zariski closure.

The model without further restrictions on the root distributions π or the transition matrices $A_{qq'}$ is known as the *general Markov model* for the tree T and the alphabet B. In applications the number of parameters is often reduced by imposing further symmetry, reflecting additional biological (or, say, linguistic) structure. This is often[1] done by choosing a finite group G acting by permutations on the set B, requiring that π be a G-invariant distribution (which when G acts transitively means that it is the uniform distribution), and requiring that each transition matrix $A_{qq'}$ satisfies $A_{qq'}(gb, gb') = A_{qq'}(b, b')$ for all letters $b, b' \in B$. The resulting model, which is a subset of $\mathbb{R}^{(B^{\mathrm{leaf}(T)})}$ contained in the general Markov model, has been dubbed the *equivariant tree model* for the triple (T, B, G) [DK09]; here we implicitly mean that the action of G on B is also fixed. The special case where G is Abelian and $B = G$ with the left action of G on itself is called a *group-based model*. Our first two main theorems concern the class of equivariant tree models for which G is Abelian, but does not necessarily act transitively on B. This class includes the general Markov model (with $G = \{1\}$) as well as group-based models.

Theorem 1.1 (Main Theorem I). *For any action of an Abelian group G on a finite alphabet B, there exists a uniform bound $D = D(B, G)$ such that for any finite tree T the Zariski closure of the equivariant tree model for (T, B, G) is defined by polynomial equations of degree at most D.*

In fact, we will prove the stronger statement that *finitely many types of equations* suffice to define the Zariski closures of the equivariant tree models for all T. For the general Markov model, this result first appeared in [DK14]. For group-based models, where the Zariski closure of (the cone over) the tree model for (T, B, G) is a toric variety, a much

[1] But not always! Most notably, the general *time-reversible Markov model*, where the only restriction on the transition matrices is that they be symmetric, is not of this form for $|B| > 2$. We have not tried to generalise our results to this case.

stronger conjecture was put forward in [SS05], namely, that for any tree T the ideal of that toric variety is generated by binomials of degree at most $|G|$. This would imply that $D(B, G) = |G|$ suffices when G acts transitively on B. Our result is weaker in that we do *not* prove the existence of a degree bound for polynomials generating the ideal—our result is *set-theoretic* rather than *ideal-theoretic*—and that we do not find an explicit bound. Nevertheless, Main Theorem I is the first general finiteness result even for the restricted class of group-based models, though for group-based models more recent work by Michałek [Mic13] gives finiteness results at the level of projective schemes, which are somewhere between set-theoretic and ideal-theoretic results.

Theorem 1.2 (Main Theorem II). *For any action of a finite group G on a finite alphabet B, there exists a polynomial-time algorithm that, on input a tree T and a probability distribution P on $B^{\mathrm{leaf}(T)}$, determines if P lies in the Zariski closure of the equivariant tree model for* (T, B, G).

We hasten to say that our proofs are nonconstructive. In particular, they do not yield an explicit bound $D(B, G)$ and they do not give an explicit algorithm—though the overall structure of that algorithm is clear (see Section 6). This situation is reminiscent of Robertson–Seymour's nonconstructive proof that any minor-closed property of finite graphs can be tested in polynomial time [RS95, RS04]. In Main Theorem II, the notion of *polynomial-time algorithm* depends on the (machine) representation of the entries of P. If they are rational numbers, then we mean polynomial-time in the bit-size of P (in a nonsparse representation, i.e., zero entries count). If they are abstract real numbers, then we mean a Blum–Shub–Smale machine [BSS89] whose number of arithmetic operations on real numbers is bounded by some polynomial in $|B|^{|\mathrm{leaf}(T)|}$.

Our Main Theorems I and II do *not* require that the trees T be trivalent. Indeed, for the class of trivalent trees, or indeed for the class of trees with any fixed upper bound on the valency of internal vertices, Main Theorems I and II are relatively easy consequences of known results from [AR08, CS05, SS05, DK09], which express the ideal of equations of an equivariant tree model in terms of ideals of equivariant tree models of *star trees*. Bounding the degree of polynomial equations for large star models and the complexity of testing membership of their Zariski closures is the real challenge in this paper. We stress that this leaves open the question of actually finding (practical) algorithms for testing membership of (Zariski closures of) tree models. Our results should be interpreted as a theoretical contribution to the algebraic statistics of tree models.

However, we do believe that some of the techniques that go into the proofs of our Main Theorems I and II can be of practical use. In particular, one crucial observation in our proofs is the following. Consider the equivariant star model for the triple (T, G, B), where T is a star and where G need not be Abelian. Label the leaves of T with $0, \ldots, m-1$, so that $B^{\mathrm{leaf}(T)}$ can be identified with B^m. Fix a natural number $n_0 \leq m$ and any probability distribution Q on B^{n_0} that is invariant with respect to the diagonal G-action on B^{n_0}. Then for any probability distribution P on B^m we can define a probability distribution P_Q on B^{m-n_0} by

$$P_Q(\mathbf{b}) = \frac{\sum_{\mathbf{b}' \in B^{n_0}} P(\mathbf{b}, \mathbf{b}') Q(\mathbf{b}')}{Z}$$

where $P(\mathbf{b}, \mathbf{b}')$ is the probability of observing $\mathbf{b}$ at positions $0, \ldots, m - n_0 - 1$ and $\mathbf{b}'$ at positions $m - n_0, \ldots, m - 1$. Here Z is a normalising factor, and a condition for this to be well-defined is that Z is nonzero. Let T' be the tree obtained from T by deleting the last n_0 leaves. Our elementary but useful observation is that, for any fixed G-invariant Q, the (partially defined) map $P \mapsto P_Q$ maps the equivariant model for (T, G, B) into the equivariant model for (T', G, B). As a consequence, equations for the latter model pull back to equations for the former model, and a necessary condition for P to be in (the Zariski closure of) the former model is that for all G-invariant Q the distribution P_Q lies in the latter model.

In the course of proving Main Theorems I and II we show that for some suitable n_0, chosen after fixing G and its action on B, and for some suitably chosen set of G-invariant probability distributions Q on B^{n_0}, the converse also holds: if a probability distribution P on B^m with $m \gg n_0$ has the property that P_Q lies in the star model with $m - n_0$ leaves for all chosen Q on all cardinality-n_0 subsets of the leaves, then P lies in the star model with m leaves. We do this by constructing an infinite-dimensional limit of all m-star models for the pair (G, B)—or rather n_0 of these limits, one for each congruence class of m modulo n_0—and showing that this limit lies in some infinite-dimensional *flattening variety* that is Noetherian up to its natural symmetries. This is also the technique followed in [DK14] for the case where $G = \{1\}$; there n_0 can be taken 1. We simplify some of the arguments from that paper, but our present, more general results are more subtle since they really require the use of jumps by some carefully chosen $n_0 > 1$.

This paper is organised as follows. In Section 2 we briefly recall the well-known tensorification of the set-up above (see, e.g., [AR08, DK09]) and state two theorems for this setting. Then in Section 3 we give some properties of tensors in finite-dimensional G-representations that will motivate the use of flattenings and our choice for n_0.

In Section 4, after fixing any value for n_0, we introduce an infinite-dimensional ambient space (again, n_0 of these, one for each congruence class modulo n_0), containing an infinite-dimensional limit of the equivariant models for finite stars; we dub this the *infinite star model*. In this section we define the *flattening variety* as well, a variety containing the infinite star model. This variety is defined by determinantal equations of bounded degree, roughly corresponding to the coarser star models where the leaves of a tree are partitioned into two subsets. We prove that the flattening variety is defined by finitely many orbits of determinantal equations under the natural symmetry group of the infinite tree model. Then in Section 5 we prove that the flattening variety is Noetherian under this symmetry group. Finally, our main theorems are derived from this in Section 6, and it is only here that we need the infinite star model mentioned before.

We conclude this introduction with a list recording values of our uniform bound $D(B, G)$ that are known to us.

Binary general Markov model: Here $G = \{1\}$ and B has cardinality two, and results from [LM04] imply that $D(B, G)$ can be taken equal to 3; apart from linear equations expressing that probabilities sum up to 1, the degree-3 equations are the determinantal equations defining the flattening variety (see Section 4). The paper [Rai12] proves the stronger statement, previously known as the GSS-conjecture [GSS05], that these equations generate the ideal of (the cone over) the general Markov model.

Binary Jukes–Cantor model: This is the group-based model with $G = B = \mathbb{Z}/2\mathbb{Z}$, and results from [SS05] show that $D(B, G)$ can be taken equal to 2. The nonlinear, quadratic equations are determinantal equations defining the finer flattening variety $Y_{[m]}^{\leq (k_\chi)_\chi}$ from Remark 3.7(v), and these generate the ideal of the cone over the model. The algebra and geometry of this model for varying trees is further studied in [BW07].

Kimura 3-parameter model: This is the group-based model with $G = B = \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$, and results from [Mic13] show that $D(B, G)$ can be taken equal to 4. The degree-4 equations were known from [SS05], where it was conjectured that they generate the ideal. The result of [Mic13] is slightly weaker than that but stronger than the purely set-theoretic statements that we are after. The geometry of this model is also studied in [CFS08, Mic14].

If one restricts oneself to trivalent trees, then more is known for other models as well, such as the strand-symmetric model [CS05] or the all-important 4-state general Markov model [AR08, FG12, BO11] or further group-based models with small groups G [SS05].

One might wonder where the restriction to Abelian G comes from; after all, tree models for which G is not Abelian are used in practice. At this point, before going through the proofs, all we can say is that they break down at the point where we prove that the infinite-dimensional flattening variety is defined by finitely many orbits of equations; see also Remark 5.9.

Finally, a word of self-criticism is in order here: it is unclear whether the degree bound and the algorithm from our main theorems will be useful in phylogenetic practice, even if they are made explicit. In phylogenetic reconstruction, certain determinantal equations coming from edges often suffice to distinguish the model for one tree from the model for another tree (with G and B fixed) [CFS11]. On the other hand, our characterisation of general Abelian tree models using contractions and flattenings gives more insight into the geometry of these models, and our infinite-dimensional methods will likely apply to other models from algebraic statistics.

2. Tensor formulation of the main results

Before we recall the tensorification of the model mentioned in the introduction, we introduce notation that will be used throughout this article. Let G be a finite Abelian group. For us, a G-representation over a field K will be assumed to be finite-dimensional, unless explicitly mentioned otherwise. Let K be an infinite field such that every G-representation over K splits into a direct sum of one-dimensional representations. For this it suffices, for instance, that K is algebraically closed and that char K does not divide $|G|$. For $m \in \mathbb{Z}_{\geq 0}$, set $[m] := \{0, \ldots, m-1\}$. If V_i is a G-representation over K for each $i \in [m]$ and if $I \subseteq [m]$, then we write $V_I := \bigotimes_{i \in I} V_i$ for the tensor product of the V_i with $i \in I$. The *rank* of a tensor ω in V_I is the minimal number of terms in any expression of ω as a sum of pure tensors $\bigotimes_{i \in I} v_i$ with $v_i \in V_i$. A tensor ω has *border rank* at most k if it lies in the Zariski closure of the set of tensors of rank at most k.

Given an m-tuple of linear maps $\phi_i : V_i \to U_i$, where U_i is also a vector space over K for each $i \in [m]$, we write $\phi_{[m]} := \bigotimes_{i \in [m]} \phi_i$ for the linear map $V_{[m]} \to U_{[m]}$ determined by $\bigotimes_{i\in[m]} v_i \mapsto \bigotimes_{i\in[m]} \phi_i(v_i)$. Clearly $\operatorname{rk}\phi_{[m]}\omega \leq \operatorname{rk}\omega$ for any $\omega \in V_{[m]}$, and this inequality carries over to the border rank.

If $I \subseteq [m]$ and $\xi \in \bigotimes_{i\in I} V_i^*$, then the tensor ξ induces a linear map $V_{[m]} \to V_{[m]-I}$. We call this map the *contraction* along the tensor ξ; except for a normalising factor, it is the tensorial analogue of the map $P \mapsto P_Q$ from the introduction. This map is G-equivariant if and only if ξ is G-invariant; moreover, it does not increase the rank or the border rank of any element of $V_{[m]}$. We can now state our third main theorem.

Theorem 2.1 (Main Theorem III). *For all $k \in \mathbb{Z}_{\geq 0}$ there exists M such that for all $m > M$ and for all G-modules V_i over K with $i \in [m]$, a tensor $\omega \in V_{[m]}$ has border rank at most k if and only if for all $\mu \leq M$, all its contractions in $m - \mu$ factors along G-invariant tensors have border rank at most k.*

The novelty in this theorem, compared to the results in [DK14], is that it suffices to contract along G-invariant tensors rather than general tensors, at the cost of increasing the dimension of those tensors to be contracted with. While not strictly necessary for our other main results, Main Theorem III illustrates the general approach taken in this paper, which is to replace “baby steps” for $G = 1$ with “giant steps” for general Abelian G. Our fourth main theorem, which generalises our first main theorem, requires a bit more work to formulate.

Definition 2.2. A *G-spaced tree* is a tree T together with for each vertex q a G-module V_q, a distinguished basis B_q of V_q such that G acts on B_q and a nondegenerate symmetric bilinear form $(\cdot|\cdot)_q$ defined by the property that B_q is an orthonormal basis with respect to $(\cdot|\cdot)$. For vertices q, q', we say $q \sim q'$ if and only if (q, q') is an edge of T. We denote by $\operatorname{vert}(T)$, $\operatorname{int}(T)$, respectively $\operatorname{leaf}(T)$, the set of vertices, internal vertices, respectively leaves, of T. We define

$$L(T) := \bigotimes_{q\in\operatorname{leaf}(T)} V_q \quad \text{and} \quad R(T) := \bigotimes_{q\in\operatorname{vert}(T)} V_q^{\otimes\{q'\sim q\}}.$$

Let T be a G-spaced tree. A *G-representation* of T is a collection $(A_{q'q})_{q'\sim q}$ of G-invariant elements of $V_{q'} \otimes V_q$ such that for any $q' \sim q$, the tensor $A_{q'q}$ maps to $A_{qq'}$ via the natural isomorphism $V_{q'} \otimes V_q \to V_q \otimes V_{q'}$. The set of G-representations of T is denoted $\operatorname{rep}_G(T)$.

Note that in the set-up of the introduction, each vertex of the tree has the same space attached; in other words, there is some G-representation V with some fixed basis B, some fixed symmetric bilinear form $(\cdot|\cdot)$ (and some fixed action of G) such that $V_q = V$, $B_q = B$ and $(\cdot|\cdot)_q = (\cdot|\cdot)$ for any vertex q of the tree. In this setting, we can view a probability distribution $P \in \mathbb{C}^{B^{\operatorname{leaf}(T)}}$ as an element of $L(T)$; namely, we can identify P with

$$\sum_{\mathbf{b}=(b_q)_{q\in\operatorname{leaf}(T)}\in B^{\operatorname{leaf}(T)}} P(\mathbf{b}) \cdot \bigotimes_{q\in\operatorname{leaf}(T)} b_q.$$

This is the tensorification of the set-up of the introduction. For our purposes, we will need to use the more flexible setting of Definition 2.2, as we will want to apply theorems proved in [DK09]. Usually however, it will suffice to consider trees for which each vertex has the same space attached; see for example Lemma 6.6.

There is a canonical isomorphism $\mathrm{rep}_G(T) \to R(T)$, defined by the embedding of elements in the tensor product of the $V_{q'} \otimes V_q$ ranging over the unordered pairs of edges $\{q' \sim q, q \sim q'\}$ into $R(T)$. We denote by Ψ (or sometimes Ψ_T to indicate which tree we are talking about) the composition of this map with the contraction $R(T) \to L(T)$ along the (G-invariant) tensor $\bigotimes_{q' \in \mathrm{int}(T)} \sum_{b \in B_{q'}} (b| \cdot)^{\otimes \{q \sim q'\}}$.

Definition 2.3. The *equivarant model* $\mathrm{CV}(T)$ associated to a tree T is the Zariski closure $\overline{\Psi(\mathrm{rep}_G(T))}$ of the image of Ψ.

Note the slight discrepancy with the introduction, where the term "equivariant model" was used for the image of Ψ on stochastically meaningful parameters. But the present definition is the one used in [DK09], from which we will use some results. While there the group G was allowed to be arbitrary, we stress once again that in the present paper we only consider *Abelian* G. We can now state our fourth main theorem.

Theorem 2.4 (Main Theorem IV). *If K is algebraically closed and of characteristic zero, then for all $k \in \mathbb{Z}_{\geq 0}$, there exists a $D \in \mathbb{Z}_{\geq 0}$ such that for each G-spaced tree T such that $|B_q| \leq k$ for each $q \in \mathrm{int}(T)$, the variety $\mathrm{CV}(T)$ is defined by the vanishing of a number of polynomials of degree at most D.*

The bound D will certainly have to depend on k. For instance, if G is the trivial group, and T is a star tree, then the variety $\mathrm{CV}(T)$ is the variety of tensors of rank at most k, and no polynomials of degree less than $k+1$ vanish on this variety. Main Theorem I is a direct corollary of this theorem; the details for passing from the case of unrooted trees without the restriction that the row sums of transition matrices are 1 to the case of rooted trees with that additional restriction can be found in Section 3 of [DK09].

3. Tensors and flattening

In the proofs of our main theorems, in addition to contractions, we will use a second operation on tensors, namely, *flattening*. Suppose that I, J form a partition of $[m]$ into two parts. Then there is a natural isomorphism $\flat = \flat_{I,J} : V_{[m]} \to V_I \otimes V_J$. The image $\flat\omega$ is a 2-tensor called a flattening of ω. Its rank (as a 2-tensor) is a lower bound on the border rank of ω. The first step in our proof below is a reduction to the case where all V_i are isomorphic as G-representations. Here, i can be viewed either as an element of $[m]$ (in Main Theorem III) or as an element of $\mathrm{leaf}(T)$ (in Main Theorem IV).

We have the following lemma, in which $K[G]$ stands for the regular representation of G.

Lemma 3.1. *Let m, k, n be natural numbers with $n \geq k+1$, and let $V_0, \ldots, V_{m-1}$ be G-representations over K. Then a tensor $\omega \in V_{[m]}$ has rank (respectively, border rank) at most k if and only if for all m-tuples of G-linear maps $\phi_i : V_i \to K[G]^n$ the tensor $\phi_{[m]}(\omega)$ has rank (respectively, border rank) at most k.*

Moreover, if $\omega \in V_{[m]}$ has border rank at most k, then there exist G-linear maps $\phi_i : V_i \to K[G]^k$ and $\psi_i : K[G]^k \to V_i$ $(i = 1, \ldots, m)$ such that $\psi_{[m]}(\phi_{[m]}(\omega)) = \omega$.

This lemma holds at the scheme-theoretical level, but we will not need that. For G the trivial group, the lemma reduces to [AR08, Theorem 11].

Proof. The "only if" part follows from the fact that $\phi_{[m]}$ does not increase rank or border rank. For the "if" part assume that ω has rank strictly larger than k, and we argue that there exist $\phi_0, \ldots, \phi_{m-1}$ such that $\phi_{[m]}(\omega)$ still has rank larger than k. It suffices to show how to find ϕ_0; the remaining ϕ_i are found in the same manner. Let U_0 be the image of ω regarded as a linear map from the dual space $V^*_{[m]-\{0\}}$ to V_0. Set

$$U_0' := K[G]U_0 = \left\{\sum_{g\in G} c_g g u : c_g \in K,\ u \in U_0\right\}.$$

For each irreducible G-representation χ, let k_χ be the multiplicity of χ in U_0'. If k_χ is at most n for each χ, then by elementary linear algebra and the fact that $K[G]$ is the sum of all irreducible representations of G there exist G-linear maps $\phi_0 : V_0 \to K[G]^n$ and $\psi_0 : K[G]^n \to V_0$ such that $\psi_0 \circ \phi_0$ is the identity map on U_0', and hence on U_0. Set $\omega' := (\phi_0 \otimes \bigotimes_{i>0} \mathrm{id}_{V_i})(\omega)$, so that by construction ω itself equals $(\psi_0 \otimes \bigotimes_{i>0} \mathrm{id}_{V_i})(\omega')$. By the discussion above, we have the inequalities $\mathrm{rk}\,\omega \geq \mathrm{rk}\,\omega' \geq \mathrm{rk}\,\omega$, so that both ranks are equal and larger than k, and we are done. If, on the other hand, there is χ such that $k_\chi > n$, then let $\phi_0 : V_0 \to K[G]^n$ be any G-linear map that maps the χ-component of U_0' surjectively onto the χ-component of $K[G]^n$ for each χ with $k_\chi > n$. Then the image of U_0 must have rank at least n. Defining ω' as before, we find that the image of ω' regarded as a linear map $V^*_{[m]-\{0\}} \to K[G]^n$ has rank at least n. In other words, the flattening $\flat_{\{0\},[m]-\{0\}}\omega'$ has rank at least $n > k$. This implies that ω' itself has rank larger than k. A similar argument applies to border rank.

For the second part, suppose ω has border rank at most k. Note that ω viewed as a linear map from $V^*_{[m]-\{0\}}$ to V_0 has rank at most k (since this is a closed condition that is satisfied by all tensors of rank at most k). Then as above, one finds there are ϕ_0, ψ_0 such that ω equals $(\psi_0 \otimes \bigotimes_{i>0} \mathrm{id}_{V_i})((\phi_0 \otimes \bigotimes_{i>0} \mathrm{id}_{V_m})(\omega))$; the second part follows by repeatedly applying this. □

Remark 3.2. Note that if ω is G-invariant, then all U_i will be G-stable and hence $U_i = U_i'$.

Moreover, note that we can refine Lemma 3.1 in the following way: an element ω of $V_{[m]}$ has (border) rank at most k if and only if there are m-tuples of G-linear maps $\phi_i : V_i \to K[G]^k$ and $\psi_i : K[G]^k \to V_i$ such that $\psi_{[m]}(\phi_{[m]}(\omega)) = \omega$ and such that $\phi_{[m]}(\omega)$ has (border) rank at most k.

Observe that finding m-tuples of G-linear maps as required (or finding that such m-tuples do not exist) is easily done by linear algebra. In essence, this means that the problem of finding whether the (border) rank of a tensor in some tensor product exceeds k can be reduced to the problem of finding whether the (border) rank of a tensor in the m-fold tensor product of the space $V = K[G]^k$ exceeds k.

Example 3.3. Consider the group $G = \mathbb{Z}/2\mathbb{Z} = \{e, g\}$ and the 8-dimensional G-module $V_0 = V_1 = K[G]^{\otimes[3]}$. Use shorthand notation such as $[eeg] := e \otimes e \otimes g \in V_0$. The tensor

$$\omega := [eee] \otimes [eee] + [ggg] \otimes [eeg] \in V_0 \otimes V_1$$

has rank 2. It can be regarded as a linear map from V_1^* to V_0, and as such it has image $U_0 := \langle [eee], [ggg] \rangle$. This subspace is already G-stable, so that

$$U_0' = K[G]U_0 = \langle [eee] + [ggg], [eee] - [ggg] \rangle,$$

where the last two vectors correspond to the two different characters of G. Define $\phi_0 : V_0 \to K[G]^2$ by $[eee] \mapsto (e, 0)$, $[ggg] \mapsto (g, 0)$ and by sending all other three-letter words over G to zero. This map is G-equivariant. Conversely, define $\psi_0 : K[G]^2 \to V_0$ by $\psi_0(e, 0) = [eee]$, $\psi_0(g, 0) = [ggg]$ and $\psi_0(0, K[G]) = \{0\}$. This ψ_0 is G-equivariant. We have used only one copy of $K[G]$ as both characters have multiplicity one in U_0'.

Next, consider ω as a linear map from V_0^* to V_1, and let $U_1 = \langle [eee], [eeg] \rangle$ be the image of that linear map. We find

$$U_1' = K[G]U_1 = \langle [eee] + [ggg], [eee] - [ggg], [eeg] + [gge], [eeg] - [gge] \rangle.$$

Each character has multiplicity two in U_1', and we will need the second factor $K[G]$. Define $\phi_1 : V_1 \to K[G]^2$ by

$$[eee] \mapsto (e, 0), \quad [ggg] \mapsto (g, 0), \quad [eeg] \mapsto (0, e), \quad [gge] \mapsto (0, g)$$

and by mapping all other words to zero. This map is G-equivariant and surjective. Let $\psi_1 : K[G]^2 \to V_1$ be the unique map such that $\psi_1 \circ \phi_1$ restricts to the identity on U_1'. Now we find that

$$\psi_{[2]}(\phi_{[2]}\omega) = (\psi_0 \otimes \psi_1)(\phi_0 \otimes \phi_1)\omega = \omega$$

as stated in the lemma.

Let V be a G-representation. Let $y_0, \dots, y_{d-1}$ be a basis of V^*. Let $m \in \mathbb{Z}_{\geq 0}$ and denote by $\mathcal{O}_m$ the coordinate ring of the affine space $V^{\otimes [m]}$. Let $u = (u_0, \dots, u_{m-1})$ be an element of $[d]^m$, i.e., a word of length m over the alphabet $[d]$ of length m. Then $\mathcal{O}_m$ can be viewed as the polynomial ring in the coordinates $\xi_u = \bigotimes_{i \in [m]} y_{u_i}$.

Several groups act naturally on $V^{\otimes [m]}$ in a G-equivariant way. First of all, denoting by $\mathrm{GL}_G(V)$ the group of invertible G-equivariant automorphisms of V, observe that $\mathrm{GL}_G(V)^m$ acts linearly on $V^{\otimes [m]}$ by

$$(\phi_0, \dots, \phi_{m-1})(v_0 \otimes \cdots \otimes v_{m-1}) = \phi_0 v_0 \otimes \cdots \otimes \phi_{m-1} v_{m-1},$$

and this action gives a right action on $(V^*)^{\otimes [m]}$ by

$$(z_0 \otimes \cdots \otimes z_{m-1})(\phi_0, \dots, \phi_{m-1}) = (z_0 \circ \phi_0) \otimes \cdots \otimes (z_{m-1} \circ \phi_{m-1}).$$

Second, the group S_m of permutations of $[m]$ acts by

$$\pi(v_0 \otimes \cdots \otimes v_{m-1}) = v_{\pi^{-1}(0)} \otimes \cdots \otimes v_{\pi^{-1}(m-1)}.$$

This leads to the contragredient action of S_m on the dual space $(V^*)^{\otimes [m]}$ by

$$\pi(z_0 \otimes \cdots \otimes z_{m-1}) = z_{\pi^{-1}(0)} \otimes \cdots \otimes z_{\pi^{-1}(m-1)}.$$

Both of these extend to an action on all of $\mathcal{O}_m$ by means of algebra automorphisms. Denote by H_m the group generated by S_m and $\mathrm{GL}_G(V)^m$ in their representations on $V^{\otimes[m]}$.

Let $k \in \mathbb{Z}_{\geq 0}$. Given any partition of $[m]$ into I, J we have the flattening $V^{\otimes[m]} \to V^{\otimes I} \otimes V^{\otimes J}$. Composing this flattening with a $(k+1) \times (k+1)$-subdeterminant of the resulting two-tensor gives a degree-$(k+1)$ polynomial in $\mathcal{O}_m$. The linear span of all these equations for all possible partitions I, J is an H_m-submodule of $\mathcal{O}_m$. Let $Y_{[m]}^{\leq k}$ (or more generally $Y_{I'}^{\leq k}$ for a finite set I') denote the subvariety of $V^{\otimes[m]}$ (or more generally $V^{\otimes I'}$) defined by this submodule. This is an H_m-stable variety, which will be very useful later on. Note that any contraction from $V^{\otimes[m]} \to V^{\otimes[m]-I}$ maps $Y_{[m]}^{\leq k}$ to $Y_{[m]-I}^{\leq k}$.

The following convention will be used in the remainder of this paper. Let $m \in \mathbb{Z}_{\geq 0}$ and let $n \in [m]$. If $\xi \in (V^*)^{\otimes n}$, then when we speak of the contraction from $V^{\otimes[m]} \to V^{\otimes[m-n]}$ along ξ, we mean the contraction along the tensor ξ viewed as an element of $(V^*)^{\otimes[m]-[m-n]}$ in the natural way; abusing notation, we will usually denote this contraction by ξ. We can now state the following crucial lemma.

Lemma 3.4. *Let V be a G-representation. Then there exists an $n_0 \in \mathbb{Z}_{>0}$ and a G-invariant tensor $\xi_0 \in (V^*)^{\otimes n_0}$ such that for all $k \in \mathbb{Z}_{\geq 0}$ and $m \gg k$, a tensor $\omega \in V^{\otimes[m]}$ lies in $Y_{[m]}^{\leq k}$ if (and only if) $\xi(\sigma(\omega))$ lies in $Y_{[m-n_0]}^{\leq k}$ for all $\sigma \in S_m$ and all G-equivariant contractions $V^{\otimes[m]} \to V^{\otimes[m-n_0]}$ along a tensor ξ of the form $\phi(\xi_0)$ with $\phi \in \mathrm{GL}_G(V)^{n_0}$.*

In this lemma, $m \gg k$ means that $m > M$ for some function $M = M(k)$ of k, which we will determine below. The lemma follows from the following lemma about contractions of subspaces of tensor powers.

Lemma 3.5. *Let V be a G-representation and set $n_1 := |G|$. There exists a G-invariant tensor $\xi_0 \in (V^*)^{\otimes n_1}$ such that for all $k \in \mathbb{Z}_{\geq 0}$ and all $m \gg k$ and all subspaces $W \subseteq V^{\otimes[m]}$ the following holds: if the dimension of $\xi(\sigma(W))$ is at most k for all $\sigma \in S_m$ and for all tensors $\xi \in (V^*)^{\otimes n_1}$ with $\xi = \phi(\xi_0)$ for some $\phi \in \mathrm{GL}_G(V)^{n_1}$, then* $\dim W$ *itself is at most k.*

Again, $m \gg k$ means that $m > M_1$ for some function $M_1 = M_1(k)$ of k, which we will determine below. To prove this lemma, we will make use of the following combinatorial lemma concerning words over a finite alphabet.

Lemma 3.6. *Let $k, l \in \mathbb{Z}_{\geq 0}$ and let A be a finite alphabet. Let $w_0, \ldots, w_k \in A^{[l]}$ be words of length l over A, written down as a $[k+1] \times [l]$-array of letters from A. For $\mathbf{a} \in A^{[k]}$ write*

$$J_{\mathbf{a}} := \{j \in [l] : \forall i \in [k+1] : (w_i)_j = a_i\}$$

for the set of positions j where the array has column $\mathbf{a}$, and for $J \subseteq [l]$ write $(w_i)_J \in A^J$ for the restriction of the word w_i to the positions in J. Then:

(i) *There exists an $\mathbf{a} \in A^{[k+1]}$ for which $|J_{\mathbf{a}}| \geq \lceil l/|A|^{k+1} \rceil$.*
(ii) *If $w_0, \ldots, w_k$ are pairwise distinct, then there exists a subset $J \subseteq [l]$ of cardinality at most k such that $(w_0)_J, \ldots, (w_k)_J$ are pairwise distinct.*

Proof. The first statement follows immediately from $\sum_{\mathbf{a}\in A^{[k+1]}} |J_{\mathbf{a}}| = l$. The second statement is proved by induction. It is clearly true for $k = 0$, with $J = \emptyset$. Suppose $k > 0$. By induction, we may assume that there is $J' \subseteq [l]$ of cardinality at most $k-1$ such that $(w_0)_{J'}, \dots, (w_{k-1})_{J'}$ are pairwise distinct. In particular, $(w_k)_{J'}$ can be equal to at most one $(w_i)_{J'}$ with $i < k$. If it is not equal to any of these, then take $J = J'$. If it is equal to some $(w_i)_{J'}$, $i < k$, then take $j \in [l]$ such that $(w_k)_j \neq (w_i)_j$ for this i and take $J := J' \cup \{j\}$. □

Proof of Lemma 3.5. Let $\widehat{G}$ be the group of characters of G. Note that $V = \bigoplus_{\chi\in\widehat{G}} V_\chi$ where $V_\chi = \{v \in V : \forall g \in G : gv = \chi(v)v\}$. Fix a basis of V of common G-eigenvectors, say $e_0, \dots, e_{d-1}$, and let $x_0, \dots, x_{d-1}$ be the dual basis. Such a basis exists since V splits into irreducible G-representations of dimension 1. Observe that each e_i is an element of V_χ for some character χ. Similarly, each x_i is an element of some V_χ^*. For each character χ, let $x_\chi = \sum_{\{i\in[d]:\, x_i\in V_\chi^*\}} x_i$. Note that x_χ can in principle be any nonzero element of V_χ^*, provided $V_\chi^* \neq \{0\}$. Indeed, we only choose a basis for technical reasons. Observe that $x_{\chi_0} \otimes \cdots \otimes x_{\chi_{\mu-1}}$ is G-invariant if the product of the corresponding characters is the trivial character. Since $\widehat{G}$ has cardinality $|G| = n_1$, the n_1-fold product of any element of $\widehat{G}$ is the trivial character, and therefore

$$\xi_0 := \sum_{\chi\in\widehat{G}} x_\chi^{\otimes n_1} \in (V^*)^{\otimes n_1}$$

is a G-invariant tensor. Let $k \in \mathbb{Z}_{\geq 0}$. We will show that $M_1 = k + |G|^{k+2} - |G|^{k+1}$ works for this ξ_0.

Let $\mathrm{Gr}(f, V^{\otimes[m]})$ denote the Grassmannian of f-dimensional subspaces of $V^{\otimes[m]}$, which is a projective algebraic variety over K. Set

$$Z(f,k) := \{W \in \mathrm{Gr}(f, V^{\otimes[m]}) : \dim \xi(\sigma(W)) \leq k \text{ for all } \xi = \phi(\xi_0),\ \phi \in \mathrm{GL}_G(V)^{n_1},\ \sigma \in S_m\},$$

a closed subvariety of $\mathrm{Gr}(f, V^{\otimes[m]})$. The assertion of the lemma is equivalent to the statement that the set of K-points of $Z(f,k)$ is empty if $f > k$ and $m > M_1$. So suppose the set of K-points of $Z(f,k)$ is nonempty for some $f > k$, $m > M_1$. We will use the fact that it is stable under $\mathrm{GL}_G(V)^m \subseteq H_m$.

Let $D \subseteq \mathrm{GL}_G(V)$ denote the subset of diagonal matrices with respect to the basis $e_0, \dots, e_{d-1}$. Then D^m is a connected, solvable algebraic group and hence by Borel's Fixed Point Theorem [Bor91, Theorem 15.2], D^m must have a fixed point W on the projective algebraic variety $Z(f,k)$. Then also $\sigma(W)$ is a fixed point of D^m for any $\sigma \in S_m$, so we can rearrange factors if necessary. Any D^m-stable subspace is spanned by common eigenvectors for D^m (any algebraic representation of D^m is diagonalisable). Now $\omega \in V^{\otimes[m]}$ is a D^m-eigenvector if and only if $\omega = e_{i_0} \otimes e_{i_1} \otimes \cdots \otimes e_{i_{m-1}}$ (up to a nonzero scalar) for some $i_0, \dots, i_{m-1}$ with $i_j \in [d]$ for each $j \in [m]$. Say $\omega_0, \dots, \omega_{f-1}$ form a basis of W of common D^m-eigenvectors and say $\omega_j = e_{j,0} \otimes e_{j,1} \otimes \cdots \otimes e_{j,m-1}$ (with each $e_{j,i}$ equal to some e_l). For a contradiction, it suffices to show that there exists a tensor ξ in the $\mathrm{GL}_G(V)^{n_1}$-orbit of ξ_0 and an element $\sigma \in S_m$ as above such that $\xi(\sigma(\omega_0)), \dots, \xi(\sigma(\omega_k))$ are linearly independent. Thus we will no longer need $\omega_{k+1}, \dots, \omega_{f-1}$.

By Lemma 3.6(ii) there exists a subset $J \subseteq [m]$ of cardinality at most k such that the tensors $\omega_{j,J} := \bigotimes_{l \in J} e_{j,l}$ for $j \in [k+1]$ are pairwise distinct (and hence linearly independent). Rearranging factors we may assume that $J \subseteq [k]$. We will contract the ω_i in n_1 positions that all lie beyond the first k positions. If those contractions are nonzero, then they are automatically linearly independent since their parts in the first k positions are.

We now set out to find those n_1 positions. For each $j \in [k+1]$, consider the word $w_j \in \widehat{G}^{[m]-[k]}$ of length $m-k$ with letter χ at position i if $e_{j,i} \in V_\chi$ (so we basically consider ω_j with the first k factors $e_{j,i}$ removed, and map the remaining factors to their corresponding characters). By Lemma 3.6(i), there exists a $\chi = (\chi_j)_{j \in [k+1]} \in \widehat{G}^{[k+1]}$ such that $J_\chi \subseteq [m]-[k]$ as in the lemma has cardinality at least $\lceil (m-k)/|\widehat{G}|^{k+1} \rceil$. The latter expression is at least $|G|$ by choice of M_1.

Now, pick a single such χ and take $I \subseteq J_\chi$ of cardinality $|G| = n_1$ as above; by applying some σ_2 if necessary, we may assume $I = [m]-[m-n_1]$. Note that $I \cap [k] = \emptyset$ as promised. For each $j \in [k+1]$ and $l \in I$, we have $e_{j,l} \in V_{\chi_j}$ and we observe that $\xi_0(\omega_{j,I}) = 1$ for all j. One easily verifies that $\xi(\omega_j) = \omega_{j,[m-n_1]}$ for each $j \in [k+1]$, and these $k+1$ tensors are linearly independent since $J \subseteq [k]$. This concludes the proof. □

Proof of Lemma 3.4. Let $n_0 = n_1 = |G|$ and let ξ_0 be as in the proof of the previous lemma. Let $k \in \mathbb{Z}_{\geq 0}$ and let $M = 2M_1 = 2(k + |G|^{k+2} - |G|^{k+1})$. Let $m > M$ and let $\omega \in V^{\otimes [m]}$ be such that for all $\sigma \in S_m$, the image of $\sigma(\omega)$ under any G-equivariant contraction $V^{\otimes [m]} \to V^{\otimes [m-n_0]}$ along a tensor $\xi = \phi(\xi_0)$ for some $\phi \in \mathrm{GL}_G(V)^{n_0}$ is an element of $Y^{\leq k}_{[m-n_0]}$.

Let $[m] = I \cup J$ be any partition and consider the corresponding flattening

$$\flat \colon V^{\otimes [m]} \to V^{\otimes I} \otimes V^{\otimes J}.$$

Replacing ω by $\sigma_1(\omega)$ for some $\sigma_1 \in S_m$ if necessary, we may assume $J = [m]-[\mu]$ and $I = [\mu]$ for some μ such that $m - \mu > M_1$. The statement that all $(k+1) \times (k+1)$-subdeterminants on $\flat\omega$ are zero is equivalent to $\flat\omega$ having rank at most k when regarded as a linear map from $(V^*)^{\otimes I}$ to $V^{\otimes J}$, or, in other words, to the image $W \subseteq V^{\otimes [m]-[\mu]}$ of this map having dimension at most k. Identify $V^{\otimes [m]-[\mu]}$ with $V^{\otimes [m-\mu]}$ in the natural way.

Since $|J| = m' > M_1$ we may apply Lemma 3.5 to W. Indeed, all contractions along tensors $\xi \in (V^*)^{\otimes n_0}$ of the form $\phi(\xi_0)$ for some $\phi \in \mathrm{GL}_G(V)^{n_0}$ map $\sigma'(W)$ to subspaces of $V^{\otimes [(m-\mu)-n_0]}$ of dimension at most k for all $\sigma' \in S_{m-\mu}$. This follows from the fact that this subspace is equal to the image W' of the map $(V^{\otimes I})^* \to V^{\otimes [m-\mu-n_0]}$ obtained by first applying $\flat\omega$ and then contracting along ξ. This, on the other hand, is nothing but the map $\flat'(\omega')$ where ω' is the image of ω under the same contraction but applied to $V^{\otimes [m]}$, and $\flat'$ is the flattening of $[m-n_0]$ along $[\mu], [m-\mu-n_0]$. Since ω' gives rise to a map of rank at most k by assumption, $\dim W' \leq k$ as claimed. Now this holds for all contractions and all factors and we may conclude that, indeed, $\dim W \leq k$, and $\flat\omega$ has rank at most k. □

Note that in both lemmas, we do not need to compute $\xi(\sigma(\omega))$ for all $\sigma \in S_m$: it suffices to use one σ for each subset of $[m]$ of cardinality n_0 to ensure the right factors are being contracted.

Remark 3.7. (i) Since G is Abelian, there is a natural bijection between $\widehat{G}$ and the set of isomorphism classes of irreducible G-representations. For this reason, we use the letter χ both for irreducible G-representations and for elements of $\widehat{G}$.

(ii) It is easily seen that the rank of ξ_0 as in Lemma 3.5 is bounded above by the number N of distinct characters that are represented by common G-eigenvectors in V^*; in particular, the rank can generally be bounded above by $|G|$. Moreover, observe that for any $n \in \mathbb{Z}_{>0}$, the elements $x_\chi^{\otimes n}$ (with χ ranging over those characters with $V_\chi^* \neq \{0\}$) are linearly independent. Hence, clearly, any flattening (other than the flattenings $I = \emptyset$, $J = [|G|]$ and $I = [|G|]$, $J = \emptyset$) of the ξ_0 we constructed has rank equal to N. Therefore, ξ_0 has rank N as well.

(iii) Potentially, one may do better than $n_0 = |G|$: one may take for n_0 the least common multiple of all orders of elements in G (i.e. the exponent of G), and reduce M correspondingly. For example, for the Klein 4-group, one may take $n_0 = 2$ and $M_1 = k + 4^{k+1}$ instead of $n_0 = 4$ and $M_1 = k + 4^{k+2} - 4^{k+1}$.

(iv) If G is nontrivial, then we can also take $M_1 = |G|^{k+2} - |G|^{k+1}$ instead of $k + |G|^{k+2} - |G|^{k+1}$, or even take n_0 to be the exponent $\exp G$ of G and $M_1 = (n_0 - 1)|G|^{k+1}$.

(v) If we restrict ourselves to G-stable subspaces W of $V^{\otimes[m]}$, then instead of considering merely the dimension of W, we can consider the $|G|$-tuple of multiplicities of the characters that are represented by a common G-eigenvector in W. Using the same n_1 and ξ_0 as in Lemma 3.5, for each tuple $(k_\chi)_{\chi\in\widehat{G}}$ there is an M_1 such that if for all contractions as in the lemma the multiplicity of χ in $\xi(\sigma(W))$ is at most k_χ for each χ, then the multiplicity of χ in W is at most k_χ. In this case, we can take $M_1 = (n_0 - 1)|G|^{\max_\chi(k_\chi)+1}$. Denoting by $Y_{[m]}^{\leq(k_\chi)_\chi}$ the set of G-invariant tensors ω in $V^{\otimes[m]}$ such that for each flattening, the multiplicity of χ in the image of ω is at most k_χ for each χ, we can prove an analogue of Lemma 3.4 for $Y_{[m]}^{\leq(k_\chi)_\chi}$ as well. This will be particularly useful in the case of the G-equivariant tree model later on.

(vi) In general, there may be many possible choices for ξ_0 (in fact, nearly all G-invariant tensors can be used, as the set of tensors such that the lemma is not satisfied is a closed set that is not equal to the set of G-invariant elements of $(V^*)^{\otimes n_1}$). For example, we could have taken $\xi_0 = \sum_{\chi_1,\ldots,\chi_{n_1}\in\widehat{G}:\, \chi_1\cdot\ldots\cdot\chi_{n_1}=1} \bigotimes_{j=1}^{n_1} x_{\chi_j} \in (V^*)^{\otimes n_1}$. In the specific case $V = K[G]$, this yields $\xi_0 = \sum_{g\in G} x_g^{\otimes n_1}$ (for a proper choice of a basis of G-eigenvectors of V), where $\{x_g\}$ is a basis dual to the basis $\{g\}$ of $K[G]$. In this case, our original choice would give $\xi_0 = \sum_{g_1,\ldots,g_n:\, g_1+\cdots+g_n=0} x_{g_1} \otimes \cdots \otimes x_{g_n}$.

(vii) In Lemma 3.1, if we restrict ourselves to G-invariant tensors, then we can formulate the following refinement. Let $(k_\chi)_{\chi\in\widehat{G}} \in \mathbb{Z}_{\geq 0}{}^{\widehat{G}}$ and let $k = \max_\chi(k_\chi)$.
Let m, n be natural numbers with $n \geq k + 1$, and let $V_0, \ldots, V_{m-1}$ be G-representations over K. Let $\omega \in V_{[m]}$ be G-invariant. Then the multiplicity of χ in the image of $\flat\omega$ is at most k_χ for each χ and each flattening $\flat$ if and only if there are m-tuples of G-linear maps $\phi_i : V_i \to K[G]^k$ and $\psi_i : K[G]^k \to V_i$ such that $\psi_{[m]}(\phi_{[m]}(\omega)) = \omega$ and $\phi_{[m]}(\omega) \in Y_{[m]}^{\leq(k_\chi)_\chi}$.

(viii) In this lemma, we explicitly make use of the fact that G is Abelian. Indeed, if G is non-Abelian, then the lemma is false. Suppose namely that G is non-Abelian, and let V be an irreducible G-representation of dimension $d > 1$; observe that $\mathrm{GL}_G(V) \cong K^*$. For $n \in \mathbb{Z}_{>0}$, let $\xi \in (V^*)^{\otimes n}$ be a G-invariant tensor. Let $m \in \mathbb{Z}_{>0}$ with $m \geq n$ and consider the set of S_m-invariant tensors in $V^{\otimes[m]}$. This is an $\binom{m+d-1}{d-1}$-dimensional subspace of $V^{\otimes[m]}$. The elements in this space that contract to 0 along ξ are the elements of the (nontrivial) kernel W of a set of $\binom{m-n+d-1}{d-1}$ linear equations. The actions of $\mathrm{GL}_G(V)^n$ and S_m do not give any additional linearly independent equations, so we have $\phi(\xi)(\sigma(W)) = \{0\}$ for any $\phi \in \mathrm{GL}_G(V)^n$ and $\sigma \in S_m$, while $W \neq \{0\}$. So Lemma 3.5 does not hold in this case. Likewise, Lemma 3.4 does not hold if G is non-Abelian.

Example 3.8. For $G = \mathbb{Z}/2\mathbb{Z}$, $k = 2$ and $V = K[G]^2$, the proof of the lemma combined with the remark shows we may use $n_0 = n_1 = 2$, $M_1 = 16 - 8 = 8$ and $M = 16$; taking basis $(e+g, 0), (0, e+g), (e-g, 0), (0, e-g)$ of V, with dual basis x_0, x_1, x_2, x_3 we could take $\xi_0 = (x_0 + x_1) \otimes (x_0 + x_1) + (x_2 + x_3) \otimes (x_2 + x_3)$.

In the case $V = K[G]$ and $(k_1, k_{-1}) = (1, 1)$, where we write $\widehat{G} = \{1, -1\}$, we may use $M_1 = 4$ and $M = 8$.

4. Infinite-dimensional tensors and the flattening variety

From now on, fix $k \in \mathbb{Z}_{\geq 0}$, and let V be a G-representation. Let $d = \dim V$ be the dimension of V. For each character χ with $V_\chi^* \neq \{0\}$, fix $x_\chi \in V_\chi^* - \{0\}$ and let $x_\chi = 0$ for all other characters. Let $n_0 = |G|$ and define

$$\xi_0 = \sum_{\chi \in \widehat{G}} x_\chi^{\otimes n_0} \in (V^*)^{\otimes n_0}$$

as in Section 3. For $m \in \mathbb{Z}_{\geq 0}$, we denote by ξ_0 the contraction from $V^{\otimes[m+n_0]} \to V^{\otimes[m]}$ along the tensor ξ_0. More specifically, we have

$$\xi_0(v_0 \otimes \cdots \otimes v_{m-1} \otimes v_m \otimes \cdots \otimes v_{m+n_0-1}) = \xi_0(v_m \otimes \cdots \otimes v_{m+n_0-1}) \cdot v_0 \otimes \cdots \otimes v_{m-1}.$$

Dually, this surjective map gives rise to the injective linear map

$$(V^*)^{\otimes[m]} \to (V^*)^{\otimes[m+n_0]}, \quad \xi \mapsto \xi \otimes \xi_0.$$

Let $\mathcal{O}_m$ be the coordinate ring of $V^{\otimes[m]}$. We identify $\mathcal{O}_m$ with the symmetric algebra $S((V^*)^{\otimes[m]})$ generated by the space $(V^*)^{\otimes[m]}$, and embed $\mathcal{O}_m$ into $\mathcal{O}_{m+n_0}$ by means of the linear inclusion $(V^*)^{\otimes[m]} \to (V^*)^{\otimes[m+n_0]}$ above.

From now on, fix $m_0 \in [n_0]$ and define the projective limit

$$A_\infty := \varprojlim_{m \in m_0 + n_0\mathbb{Z}_{\geq 0}} V^{\otimes[m]}$$

along the surjective linear contraction maps ξ_0. This is, in the first place, an uncountable-dimensional G-representation over K (unless $d = 1$, in which case it is one-dimensional).

But it is also the dual of the countable-dimensional *direct* limit of the $(V^*)^{\otimes[m]}$ along the inclusion maps. As a consequence, A_∞ is canonically isomorphic to the set of K-algebra homomorphisms $\mathcal{O}_\infty \to K$, where $\mathcal{O}_\infty := \bigcup_{m \in m_0 + n_0\mathbb{Z}_{\geq 0}} \mathcal{O}_m$. This gives A_∞ a Zariski topology, with closed sets given by the vanishing of subsets of $\mathcal{O}_\infty$. Since we are only concerned with set-theoretic statements, we do not need to worry about points of $\mathcal{O}_\infty$ over K-algebras other than K; the topological space A_∞ suffices for our purposes. The same applies to closed subsets (subvarieties) of A_∞ featuring below.

At a crucial step in our arguments we will use the following more concrete description of $\mathcal{O}_\infty$. Extend ξ_0 to a basis $\xi_0, \xi_1, \ldots, \xi_{d^{n_0}-1}$ of $(V^*)^{\otimes n_0}$ of G-eigenvectors. Moreover, let $y_0, \ldots, y_{d-1}$ be any basis of V^* (not necessarily consisting of G-eigenvectors). Let m be an element of $m_0 + n_0\mathbb{Z}_{\geq 0}$. Then for any $p \in \mathbb{Z}_{\geq 0}$, $(V^*)^{\otimes[m+pn_0]}$ has a basis in bijection with the pairs (u, w) with u a word in $[d]^m$ and $w = (i_0, \ldots, i_{p-1})$ a word of length p over the alphabet $[d^{n_0}]$, namely,

$$\zeta_{m,u,w} := y_{u_0} \otimes y_{u_1} \otimes \cdots \otimes y_{u_{m-1}} \otimes \xi_{i_0} \otimes \cdots \otimes \xi_{i_{p-1}}.$$

The algebra $\mathcal{O}_{m+pn_0}$ is the polynomial algebra in the variables $\zeta_{m,u,w}$ with w running over all words of length p, and u running over all words in $[d]^m$. In $\mathcal{O}_\infty$, the coordinate $\zeta_{m,u,w}$ is identified with the variable $\zeta_{m,u,w'}$ where w' is obtained from w by appending an infinite string of zeros at the end of w. If $w = 0$, then we also write $\zeta_{m,u} = \zeta_{m,u,w}$.

We conclude that $\mathcal{O}_\infty$ is a polynomial ring in countably many variables that are (for fixed $m \in m_0 + n_0\mathbb{Z}_{\geq 0}$) in bijective correspondence with triples $(m, u, (i_0, i_1, \ldots))$ in which all but finitely many i_j are 0. The finite set of positions j with $i_j \neq 0$ is called the *support* of the word $(i_1, i_2, \ldots)$; likewise, the set of positions j with $u_j \neq 0$ is called the *support* of u. Note that this gives a different set of variables for each $m \in m_0 + n_0\mathbb{Z}_{\geq 0}$; we will generally use the set of variables that is most convenient for our purposes.

Observe that for each m in $\mathbb{Z}_{\geq 0}$ we have natural embeddings $\mathrm{GL}_G(V)^m \to \mathrm{GL}_G(V)^{m+n_0}$, which render the contraction maps $V^{\otimes[m+n_0]} \to V^{\otimes[m]}$ equivariant with respect to $\mathrm{GL}_G(V)^m$. Therefore the union of $\mathrm{GL}_G(V)^m$ for all $m \in m_0 + n_0\mathbb{Z}_{\geq 0}$ acts on A_∞ and $\mathcal{O}_\infty$ by passing to the limit.

Let $S_\infty := \bigcup_{m \in m_0 + n_0\mathbb{Z}_{\geq 0}} S_m$, where S_m is embedded in S_{m+n_0} as the subgroup fixing $\{m, \ldots, m + n_0 - 1\}$. Then S_∞ is the group of all bijections $\pi : \mathbb{Z}_{\geq 0} \to \mathbb{Z}_{\geq 0}$ whose set of fixed points has a finite complement. This group acts on A_∞ and on $\mathcal{O}_\infty$ by passing to the limit.

The action of S_∞ on $\mathcal{O}_\infty$ has the following fundamental property: for each $f \in \mathcal{O}_\infty$ there exists an $m \in m_0 + n_0\mathbb{Z}_{\geq 0}$ such that whenever $\pi, \sigma \in S_\infty$ agree on the initial segment $[m]$ we have $\pi f = \sigma f$. Indeed, we may take m equal to $m_0 + (n_0$ times (1 plus the maximum of the union of the supports of words w for which $\zeta_{m_0,u,w}$ appears in f)). In this situation, there is a natural left action of the *increasing monoid* $\mathrm{Inc}(\mathbb{Z}_{\geq 0}) = \{\pi : \mathbb{Z}_{\geq 0} \to \mathbb{Z}_{\geq 0} : \pi(0) < \pi(1) < \cdots\}$ by means of injective algebra endomorphisms on $\mathcal{O}_\infty$ (see [HS12, Section 5]). The action is defined as follows: for $f \in \mathcal{O}_\infty$, let m be as above. Then to define πf for $\pi \in \mathrm{Inc}(\mathbb{Z}_{\geq 0})$ take any $\sigma \in S_\infty$ that agrees with π on the interval $[m]$ (such a σ exists) and set $\pi f := \sigma f$.

By construction, the $\mathrm{Inc}(\mathbb{Z}_{\geq 0})$-orbit of any $f \in \mathcal{O}_\infty$ is contained in the S_∞-orbit of f. Note that the left action of $\mathrm{Inc}(\mathbb{Z}_{\geq 0})$ on $\mathcal{O}_\infty$ gives rise to a *right* action of $\mathrm{Inc}(\mathbb{Z}_{\geq 0})$ by means of surjective linear maps $A_\infty \to A_\infty$. A crucial argument in Section 5 uses a map that is not equivariant with respect to S_∞ but is equivariant relative to $\mathrm{Inc}(\mathbb{Z}_{\geq 0})$.

Recall that H_m is the group generated by S_m and $\mathrm{GL}_G(V)^m$. We can now define

$$H_\infty := \bigcup_{m \in m_0 + n_0\mathbb{Z}_{\geq 0}} H_m.$$

This group acts on A_∞ and $\mathcal{O}_\infty$ by passing to the limit.

Now we get back to flattenings. Recall that $\xi_0 : V^{\otimes[m+n_0]} \to V^{\otimes[m]}$ maps $Y^{\leq k}_{[m+n_0]}$ to $Y^{\leq k}_{[m]}$; this means we can define a variety

$$Y^{\leq k}_\infty := \varprojlim_{m \in m_0 + n_0\mathbb{Z}_{\geq 0}} Y^{\leq k}_{[m]} \subseteq A_\infty.$$

We describe the determinants of flattenings in more concrete terms in the coordinates $\zeta_{m,u,w}$. Let $\mathbf{u} = (u_0, \ldots, u_k)$ be a $(k+1)$-tuple of pairwise distinct words in $[d]^m$. Let $\mathbf{u}' := (u'_0, \ldots, u'_k)$ be another such $(k+1)$-tuple. Suppose that the support of each u_i is disjoint from that of each u'_j. In this case, it makes sense to speak of $u_i + u'_j$, which is again a word in $[d]^m$. We let $\zeta[\mathbf{u}; \mathbf{u}'] = \zeta_m[\mathbf{u}; \mathbf{u}']$ be the $(k+1) \times (k+1)$ matrix with (i, j)-entry $\zeta_{m,u_i+u'_j}$. For each $m \in m_0 + n_0\mathbb{Z}_{\geq 0}$, the variety $Y^{\leq k}_{[m]}$ is defined by the determinants of all matrices $\zeta[\mathbf{u}; \mathbf{u}']$. Then the variety $Y^{\leq k}_\infty$ is defined by the determinants of all matrices $\zeta[\mathbf{u}; \mathbf{u}']$ (viewed as elements of $\mathcal{O}_\infty$) with $m \in m_0 + n_0\mathbb{Z}_{\geq 0}$.

Moreover, if $\mathbf{w} = (w_0, \ldots, w_k)$ and $\mathbf{w}' = (w'_0, \ldots, w'_k)$ are $k+1$-tuples of pairwise distinct infinite words with letters in $[d^{n_0}]$ with finite support, if $\mathbf{u}$ and $\mathbf{u}'$ are as above, and if the support of each w_i is disjoint from that of each w'_j for all $i, j \in [k]$, then we can define a $(k+1) \times (k+1)$-matrix $\zeta[\mathbf{u}, \mathbf{w}; \mathbf{u}', \mathbf{w}']$ in a way analogous to the above.

We now have the following important proposition.

Proposition 4.1. *The flattening variety $Y^{\leq k}_\infty$ is the common zero set of finitely many H_∞-orbits of $(k+1) \times (k+1)$-determinants $\det \zeta[\mathbf{u}; \mathbf{u}']$ with $\mathbf{u}, \mathbf{u}'$ as above.*

Proof. Let M be an integer such that the conclusion of Lemma 3.4 holds for the triple (M, n_0, ξ_0). Let $f_0, f_1, \ldots, f_{N-1} \in \mathcal{O}_\mu$ be finitely many $(k+1) \times (k+1)$-determinants that define $Y^{\leq k}_{[\mu]}$, where μ is the largest element of $m_0 + n_0\mathbb{Z}_{\geq 0}$ that satisfies $\mu \leq M$. Of course, in the inclusion $\mathcal{O}_\mu \subset \mathcal{O}_\infty$, each f_i may be assumed to be one of the $\det \zeta[\mathbf{u}; \mathbf{u}']$ for $\mathbf{u}, \mathbf{u}'$ each lists of $k+1$ words supported in $[\mu]$.

We will now show that $\omega \in A_\infty$ is an element of $Y^{\leq k}_\infty$ if and only if $f_i(h(\omega)) = 0$ for all i and all $h \in H_\infty$. Note that $f_i(h(\omega))$ is equal to $f_i((h(\omega))_\mu)$ where $(h(\omega))_\mu$ is the image of $h(\omega)$ in $V^{\otimes[\mu]}$ under the canonical projection $A_\infty \to V^{\otimes[\mu]}$. Now if $\omega \in Y^{\leq k}_\infty$, then obviously so is $h(\omega)$ for each $h \in H_\infty$, and hence $(h(\omega))_\mu$ is an element of $Y^{\leq k}_{[\mu]}$. This shows the "only if" part.

For the converse, suppose that $f_i(h(\omega)) = 0$ for all i and all $h \in H_\infty$. We need to show that $\omega \in Y^{\leq k}_\infty$. Equivalently, we need to show that for all $m \geq \mu$ (and $m \in m_0 + n_0\mathbb{Z}_{\geq 0}$), the image $\omega_m \in V^{\otimes[m]}$ of ω lies in $Y^{\leq k}_{[m]}$. Suppose $m = \mu + pn_0$ with $p \in \mathbb{Z}_{\geq 0}$. Recall that $f_i \in \mathcal{O}_\mu$ is identified in $\mathcal{O}_{\mu+pn_0}$ with f_i precomposed with the contraction ξ of the last pn_0 factors V along $\xi = \xi_0^{\otimes p}$. This means $f_i(\xi((h\omega)_m)) = 0$ for all $i \in [N]$ and all $h \in H_\infty$ and hence $\xi((h\omega)_m) \in Y_{[\mu]}^{\leq k}$ for all $h \in H_\infty$. Hence in particular, $\xi((h\omega)_m) \in Y_{[\mu]}^{\leq k}$ for all $h \in H_m$ of the form $\phi \circ \sigma$ with $\phi \in \mathrm{GL}_G(V)^{[m]-[\mu]}$ and $\sigma \in S_m$.

Note that for such h, one has $(h\omega)_m = h(\omega)_m$, and moreover the element $\xi((h\omega)_m)$ can be obtained by performing consecutive contractions of $\sigma(\omega_m)$ along tensors of the form $\phi'(\xi_0)$ (and in fact, all contractions of this form can be obtained in this way using some suitable h). By repeatedly applying Lemma 3.4 this means that $\omega_m \in Y_{[m]}^{\leq k}$, and we are done. □

Remark 4.2. Again, this proof can be extended to a proof for $Y_\infty^{\leq (k_\chi)_{\chi \in \widehat{G}}}$.

Example 4.3. For $G = \mathbb{Z}/2\mathbb{Z}$, $k = 2$, $m_0 = 0$ and $V = K[G]^3$, we have $M = 16$, hence $\mu = 16$. Following the proof of the proposition, we find that $Y_\infty^{\leq 2}$ is defined by the H_∞-orbits of the equations that determine $Y_{[16]}^{\leq 2}$. Let $y_0, y_1, y_2, y_3, y_4, y_5 \in V^*$ be a basis dual to the basis $(e, 0, 0), (g, 0, 0), (0, e, 0), (0, g, 0), (0, 0, g), (0, 0, e)$ of V. For $i \in [6]$ and $I \subseteq [20]$, let $u_{i,I} \in [6]^I$ be the word each of whose letters is an i. It is now an easy exercise to show that $Y_\infty^{\leq 2}$ is defined by the H_∞-orbits of $\det(\zeta[(u_{0,[n]}, u_{2,[n]}, u_{4,[n]}); (u_{0,[16]-[n]}, u_{2,[16]-[n]}, u_{4,[16]-[n]})])$ for $n \in \{1, 2, \ldots, 8\}$. Here, we make use of the fact that the $\mathrm{GL}_G(V)$-orbit of any triple of elements $v_0, v_1, v_2 \in V$ is dense in V provided that their projections to the common G-eigenspaces of V are linearly independent as well. This holds in a somewhat larger generality as well for general k-tuples and $(k+1)$-tuples of elements in V.

For $G = \mathbb{Z}/2\mathbb{Z}$ and $(k_1, k_{-1}) = (1, 1)$, things are somewhat more subtle. Let $V = K[G]$ and $m_0 = 0$. We have $M = 8$; let $y_0, y_1 \in V^*$ be a basis dual to the basis $e + g$, $e - g$ of V. Using the proof in [SS05] that the group-based model for $G = \mathbb{Z}/2\mathbb{Z}$ is defined by linear and quadratic polynomials, we can show that $Y_\infty^{\leq (1,1)}$ is defined by the H_∞-orbits of $\zeta_{8,u}$ where the cardinality of $\{i \in [8] : u_i = 1\}$ is odd and by the H_∞-orbits of $\zeta_{8,u_0}\zeta_{8,u_1} - \zeta_{8,u_2}\zeta_{8,u_3}$ such that:

(a) For each $i \in [8]$, the multiset $\{(u_0)_i, (u_1)_i\}$ equals the multiset $\{(u_2)_i, (u_3)_i\}$.
(b) For each $j \in \{1, 2, 3, 4\}$, the cardinality of $\{i \in [8] : (u_j)_i = 1\}$ is even.

We will give some more details about this in Example 6.10.

5. Equivariantly Noetherian rings and spaces

We briefly recall the notions of equivariantly Noetherian rings and topological spaces, and proceed to prove the main result of this section, namely that $Y_\infty^{\leq k}$ is H_∞-Noetherian (Theorem 5.6).

If a monoid Π has a left action by means of endomorphisms on a commutative ring R (with 1), then we call R *equivariantly Noetherian*, or Π-Noetherian, if every chain $I_0 \subseteq I_1 \subseteq \cdots$ of Π-stable ideals stabilises. This is equivalent to the statement that every Π-stable ideal in R is generated by finitely many Π-orbits. Similarly, if Π acts on a topological space X by means of continuous maps $X \to X$, then we call X equivariantly

Noetherian, or Π-Noetherian, if every chain $X_1 \supseteq X_2 \supseteq \cdots$ of Π-stable closed subsets stabilises. If R is a K-algebra, then we can endow the set X of K-valued points of R, i.e., K-algebra homomorphisms $R \to K$ (sending 1 to 1), with the Zariski topology. An endomorphism $\Phi : R \to R$ gives a continuous map $\phi : X \to X$ by pull-back, and if R has a left Π-action making it equivariantly Noetherian, then this induces a right Π-action on X making X equivariantly Noetherian. This means, more concretely, that any Π-stable closed subset of X is defined by the vanishing of finitely many Π-orbits of elements of R. If Π happens to be a group, then we can make the right action into a left action by taking inverses. Here are some further easy lemmas; for their proofs we refer to [Dra10].

Lemma 5.1. *If X is a Π-Noetherian topological space, then any Π-stable closed subset of X is Π-Noetherian with respect to the induced topology.*

Lemma 5.2. *If X and Y are Π-Noetherian topological spaces, then the disjoint union $X \cup Y$ is also Π-Noetherian with respect to the disjoint union topology and the natural action of Π.*

Lemma 5.3. *If X is a Π-Noetherian topological space, Y is a topological space with Π-action (by means of continuous maps), and $\phi : X \to Y$ is a Π-equivariant continuous map, then $\operatorname{im}\phi$ is Π-Noetherian with respect to the topology induced from Y.*

Lemma 5.4. *If Π is a group and $\Pi' \subseteq \Pi$ a subgroup acting from the left on a topological space X', and if X' is Π'-Noetherian, then the orbit space $X := (\Pi \times X')/\Pi'$ is a left-Π-Noetherian topological space.*

In this lemma, $\Pi \times X'$ carries the direct-product topology of the discrete group Π and the topological space X', the right action of Π' on it is by $(\pi, x)\sigma = (\pi\sigma, \sigma^{-1}x)$, and the topology on the quotient is the coarsest topology that makes the projection continuous. The left action of Π on the quotient comes from left-action of Π on itself. As a consequence, closed Π-stable sets in X are in one-to-one correspondence with closed Π'-stable sets in X', whence the lemma. Next we recall a fundamental example of an equivariantly Noetherian ring, which will be crucial in what follows.

Theorem 5.5 ([Coh67, HS12]). *For any Noetherian ring Q and any $l \in \mathbb{Z}_{\geq 0}$, the ring $Q[x_{ij} : i = 0, \ldots, l-1;\ j = 0, 1, 2, 3, \ldots]$ is equivariantly Noetherian with respect to the action of $\operatorname{Inc}(\mathbb{Z}_{\geq 0})$ by $\pi x_{ij} = x_{i\pi(j)}$.*

Main Theorems III and IV will be derived from the following theorem, whose proof occupies the rest of this section.

Theorem 5.6. *For every natural number k the variety $Y_\infty^{\leq k}$ is an H_∞-Noetherian topological space.*

We will proceed by induction on k. For $k = 0$ the variety $Y_\infty^{\leq k}$ consists of a single point, the zero tensor, and the theorem trivially holds. Now assume that the theorem holds for $k - 1$. By Proposition 4.1 there exists $m \in m_0 + n_0\mathbb{Z}_{\geq 0}$ and there exist k-tuples $\mathbf{u}_0, \ldots, \mathbf{u}_{N-1}, \mathbf{u}'_0, \ldots, \mathbf{u}'_{N-1}$ of words in $[d]^m$ such that $\zeta[\mathbf{u}_a; \mathbf{u}'_a]$ is defined for all $a \in [N]$ (i.e., the supports of the words in $\mathbf{u}_a$ are disjoint from the supports of the words in $\mathbf{u}'_a$) and such that $Y_\infty^{\leq k-1}$ is the common zero set of the polynomials in $\bigcup_{a=0}^{N-1} H_\infty \det(\zeta[\mathbf{u}_a; \mathbf{u}'_a])$. For each $a \in [N]$ let Z_a denote the open subset of $Y_\infty^{\leq k}$ where not all elements of $H_\infty \det(\zeta[\mathbf{u}_a; \mathbf{u}'_a])$ vanish; hence we have

$$Y_\infty^{\leq k} = Y_\infty^{\leq k-1} \cup Z_0 \cup \cdots \cup Z_{N-1}.$$

We will show that each Z_a, $a \in [N]$, is an H_∞-Noetherian topological space, with the topology induced from the Zariski topology on A_∞. Together with the induction hypothesis and Lemmas 5.2 and 5.3, this then proves that $Y_\infty^{\leq k}$ is H_∞-Noetherian, as claimed.

To prove that $Z := Z_a$ is H_∞-Noetherian, consider $\mathbf{u} := \mathbf{u}_a = (u_0, \ldots, u_{k-1})$ and $\mathbf{u}' := \mathbf{u}'_a = (u'_0, \ldots, u'_{k-1})$ with all u_i, u'_j in $[d]^m$. Let Z' denote the open subset of $Y_\infty^{\leq k}$ where $\det(\zeta[\mathbf{u}; \mathbf{u}'])$ is nonzero. This subset is stable under the group S'_∞ of all permutations σ in S_∞ that restrict to the identity on $[m]$ and such that there is $\tau \in S_\infty$ such that $\sigma(m+pn_0+i) = m+\tau(p)n_0+i$ for any $p \in \mathbb{Z}_{\geq 0}$ and $i \in [n_0]$. Note that for such σ, one has $\sigma(\zeta_{m,u,w}) = \zeta_{m,u,\tau(w)}$ where $\tau(w)_p = w_{\tau^{-1}(p)}$. More explicitly, S'_∞ consists of all permutations in S_∞ that restrict to the identity on $[m]$ and that permute the set of blocks of the form $[m + (p + 1)n_0] - [m + pn_0]$ with $p \in \mathbb{Z}_{\geq 0}$.

Lemma 5.7. *The open subset $Z' \subseteq Y_\infty^{\leq k}$ is an S'_∞-Noetherian topological space.*

Proof. We will prove that Z' is $\mathrm{Inc}(\mathbb{Z}_{\geq 0})'$-Noetherian, where $\mathrm{Inc}(\mathbb{Z}_{\geq 0})'$ is the set of all increasing maps $\pi \in \mathrm{Inc}(\mathbb{Z}_{\geq 0})$ that restrict to the identity on $[m]$ and are such that there is $\tau \in \mathrm{Inc}(\mathbb{Z}_{\geq 0})$ such that $\pi(m + pn_0 + i) = m + \tau(p)n_0 + i$ for any $i \in [n_0]$; consult Section 4 for the action of $\mathrm{Inc}(\mathbb{Z}_{\geq 0})$. Since the $\mathrm{Inc}(\mathbb{Z}_{\geq 0})'$-orbit of an equation is contained in the corresponding S'_∞-orbit, this will imply that Z' is S'_∞-Noetherian.

We start with the polynomial ring R in the variables $\zeta_{m,u,w}$, where w runs over all infinite words over the alphabet $[d^{n_0}]$ with the property that the support of w has cardinality at most 1. Among these variables there are d^m for which $w = 0$, namely the $\zeta_{m,u}$ with $u \in [d]^m$, and the remaining variables are labelled by $[d]^m \times ([d^{n_0}] - \{0\}) \times \mathbb{Z}_{\geq 0}$, where the element of $[d^{n_0}] - \{0\}$ denotes the nonzero letter of w and the element of $\mathbb{Z}_{\geq 0}$ denotes the position at which this nonzero letter occurs. On these variables acts $\mathrm{Inc}(\mathbb{Z}_{\geq 0})'$, fixing the first d^m variables and acting only on the last (position) index of the last set of variables. By Theorem 5.5 with Q the ring in the first d^m variables and $l = d^m \times (d^{n_0} - 1)$, the ring R is $\mathrm{Inc}(\mathbb{Z}_{\geq 0})'$-Noetherian. Let $S = R[\det(\zeta[\mathbf{u}; \mathbf{u}'])^{-1}]$ be the localisation of R at the determinant $\det \zeta[\mathbf{u}, \mathbf{u}']$; again, S is $\mathrm{Inc}(\mathbb{Z}_{\geq 0})'$-Noetherian. We will construct an $\mathrm{Inc}(\mathbb{Z}_{\geq 0})'$-equivariant map ϕ from the set of K-valued points of S to A_∞ whose image contains Z'. We do this, dually, by means of an $\mathrm{Inc}(\mathbb{Z}_{\geq 0})'$-equivariant homomorphism Φ from $\mathcal{O}_\infty$ to S.

To define Φ recursively, we first fix a partition I, J of $[m]$ such that the support of each u_i is contained in I and the support of each u'_j is contained in J. Now if $\zeta_{m,u,w} \in \mathcal{O}_\infty$ is one of the variables in R, then we set $\Phi(\zeta_{m,u,w}) := \zeta_{m,u,w}$. Suppose that we have already defined Φ on variables $\zeta_{m,u,w}$ such that $\mathrm{supp}(w)$ has cardinality at most b, let w be a word for which $\mathrm{supp}(w)$ has cardinality $b + 1$ and let u be a word in $[d]^m$. We will define the image of $\zeta_{m,u,w}$. Let p be the maximum of the support of w, and write

$w = w_k + w'_k$, where the support of w'_k is $\{p\}$ and the support of w_k is contained in $[p]$. Likewise, write $u = u_k + u'_k$ where the support of u_k is contained in I and the support of u'_k is contained in J. Consider the determinant of the matrix

$$\zeta[(u_0, \ldots, u_k), (w_0, \ldots, w_k); (u'_0, \ldots, u'_k), (w'_0, \ldots, w'_k)],$$

where $w_0, \ldots, w_{k-1}$ and $w'_0, \ldots, w'_{k-1}$ are all equal to the infinite word over $[d^{n_0}]$ consisting of zeroes only. This determinant equals

$$\det(\zeta[(u_0, \ldots, u_{k-1}); (u'_0, \ldots, u'_{k-1})]) \cdot \zeta_{m,u,w} - f,$$

where $f \in \mathcal{O}_\infty$ is a polynomial in variables that are of the form $\zeta_{m,u_i+u'_j,w_i+w'_j}$ with $i, j \leq k$ but not both equal to k. All of these $w_i + w'_j$ have support of cardinality at most b (since only w_k and w'_k have nonempty support and moreover, these two words have support of cardinality at most b), so $\Phi(f)$ has already been defined. Then we set

$$\Phi(\zeta_{m,u,w}) := \det(\zeta[\mathbf{u}, \mathbf{u}'])^{-1}\Phi(f).$$

The map Φ is $\mathrm{Inc}(\mathbb{Z}_{\geq 0})'$-equivariant by construction.

The set $Z' \subseteq Y_\infty^{\leq k}$ is contained in the image of the map ϕ. Indeed, this follows directly from the fact that the determinant of the matrix

$$\zeta[(u_0, \ldots, u_k), (w_0, \ldots, w_k); (u'_0, \ldots, u'_k), (w'_0, \ldots, w'_k)].$$

vanishes on Z' while $\det(\zeta[\mathbf{u}, \mathbf{u}'])$ does not. More precisely, Z' equals the intersection of $Y_\infty^{\leq k}$ with $\mathrm{im}\,\phi$, and hence by Lemmas 5.3 and 5.1 it is $\mathrm{Inc}(\mathbb{Z}_{\geq 0})'$-Noetherian. We have already pointed out that this implies that Z' is S'_∞-Noetherian. □

Now that Z' is S'_∞-Noetherian, Lemma 5.4 implies that the H_∞-space $(H_\infty \times Z')/S_\infty$ is H_∞-Noetherian. The map from this space to A_∞ sending (g, z') to gz' is H_∞-equivariant and continuous, and its image is the open set $Z \subseteq Y_\infty^{\leq k}$. Lemma 5.3 now implies that Z is S_∞-Noetherian. We conclude that, in addition to the closed subset $Y_\infty^{\leq k-1} \subseteq Y_\infty^{\leq k}$, also the open subsets $Z_0, \ldots, Z_{N-1}$ are S_∞-Noetherian. As mentioned before, this implies that $Y_\infty^{\leq k} = Y_\infty^{\leq k-1} \cup Z_0 \cup \cdots \cup Z_{N-1}$ is S_∞-Noetherian, as claimed in Theorem 5.6.

Remark 5.8. Since $Y_\infty^{\leq (k_\chi)_\chi}$ is an H-stable closed subset of $Y_\infty^{\leq \sum_\chi k_\chi}$, it is an H_∞-Noetherian topological space as well.

Remark 5.9. A natural question regarding our Main Theorems is why we restrict to Abelian groups G. Do our results carry over to general G, so that they apply to other phylogenetic models? Frankly, we do not know. Certainly the fact that G is Abelian is used in the proof of Lemma 3.5. This is used in Proposition 4.1 to prove that $Y_\infty^{\leq k}$ is defined by finitely many polynomials up to symmetry, which in turn is used in the induction proof in this section that $Y_\infty^{\leq k}$ is Noetherian. In the non-Abelian case, we have no idea whether (a suitable variant of) $Y_\infty^{\leq k}$ is defined by finitely many orbits of equations; and (a variant of) A_∞ seems simply too large to work with directly. On the other hand, in the case where G has a normal Abelian subgroup that acts transitively on B, finiteness results are proved in [Mic14].

6. Proofs of the main theorems

Recall that in Section 4, we fixed $n_0 \in \mathbb{Z}_{>0}$, a G-representation V, a tensor ξ_0 (viewed as a contraction $\xi_0 : V^{\otimes[m+n_0]} \to V^{\otimes[m]}$ for each $m \in \mathbb{Z}_{\geq 0}$) and a $k \in \mathbb{Z}_{\geq 0}$. Moreover, for each $m_0 \in [n_0]$ we defined the flattening variety $Y_\infty^{\leq k}$ which implicitly depends on all of these. In this section, n_0 and ξ_0 are still defined as before; however, we wish to stress that some of the theorems that follow hold for *any* $k \in \mathbb{Z}_{\geq 0}$ and any $m_0 \in [n_0]$; in these cases, we explicitly mention them in the statement of the theorems. If we do not mention them, then they will be defined implicitly as above. Finally, we will sometimes use specific G-representations V in our theorems.

Here are a few theorems that follow from Theorem 5.6.

Theorem 6.1. *For any fixed natural number k, any closed H_∞-stable subset Z_∞ of $Y_\infty^{\leq k}$ is the common zero set in A_∞ of finitely many H_∞-orbits of polynomials in $\mathcal{O}_\infty$.*

Proof. As Z_∞ is a closed H_∞-stable subsets of $Y_\infty^{\leq k}$, and as $Y_\infty^{\leq k}$ is an H_∞-Noetherian topological space (Theorem 5.6), Z_∞ is cut out from $Y_\infty^{\leq k}$ by finitely many H_∞-orbits of equations. Moreover, $Y_\infty^{\leq k}$ itself is cut out from A_∞ by finitely many H_∞-orbits of Equations (Proposition 4.1), and hence the same is true for Z_∞. □

Theorem 6.2. *Let Z_∞ be the projective limit in A_∞ of certain H_m-stable closed subsets $Z_m \subseteq Y_{[m]}^{\leq k}$ for m running through $m_0 + n_0\mathbb{Z}_{\geq 0}$ that satisfy $\xi_0(Z_{m+n_0}) \subseteq Z_m$ for any $m \in m_0 + n_0\mathbb{Z}_{\geq 0}$. Suppose moreover that there exists a tensor $\epsilon_0 \in V^{\otimes[n_0]}$ such that the inclusion maps $\iota : V^{\otimes[m]} \to V^{\otimes[m+n_0]}$, $\omega \mapsto \omega \otimes \epsilon_0$, map Z_m into Z_{m+n_0} and $\xi_0 \circ \iota =$* $\mathrm{id}_{V^{\otimes[m]}}$ (*i.e.* $\xi_0(\epsilon_0) = 1$). *Then there exists $D \in \mathbb{Z}_{\geq 0}$ such that for all $m \in m_0 + n_0\mathbb{Z}_{\geq 0}$, Z_m is defined by the vanishing of a number of polynomials of degree at most D.*

Proof. By Theorem 6.1 there exists a D such that Z_∞ is defined in A_∞ by polynomials of degree at most D; we prove that the same D suffices. Indeed, suppose that all polynomials of degree at most D in the ideal of Z_m vanish on a tensor $\omega \in V^{\otimes[m]}$. Let ω_∞ be the element of A_∞ obtained from ω by successively applying ι. More precisely, ω_∞ is the element in A_∞ defined by $(\omega_\infty)_{m+pn_0} = \omega \otimes \epsilon_0^{\otimes p}$ for any $p \in \mathbb{Z}_{\geq 0}$. Here, $(\omega_\infty)_{m+pn_0}$ denotes the image of ω_∞ under the natural projection $A_\infty \to V^{\otimes[m+pn_0]}$.

We claim that ω_∞ lies in Z_∞. Indeed, otherwise some $\mathcal{O}_{m'}$ contains a polynomial f of degree at most D that vanishes on $Z_{m'}$ but not on ω_∞. Now m' cannot be smaller than m, because then f vanishes on Z_m but not on ω. But if $m' = m + pn_0 \in m + n_0\mathbb{Z}_{\geq 0}$, then $f \circ \iota^p$ is a polynomial in $\mathcal{O}_m$ of degree at most D that vanishes on Z_m but not on ω. This contradicts the assumption on ω. □

The next theorem will be rather more subtle than the previous ones, as it involves contractions along G-invariant tensors that are not necessarily of length pn_0. For this reason, we will assume the existence of closed subsets Z_m of $Y_{[m]}^{\leq k}$ for each $m \in \mathbb{Z}_{\geq 0}$, rather than just for each $m \in m_0 + n_0\mathbb{Z}_{\geq 0}$.

Theorem 6.3. *For each $m \in \mathbb{Z}_{\geq 0}$, let $Z_m \subseteq Y_{[m]}^{\leq k}$ be an H_m-stable closed subset. Suppose that all contractions $V^{\otimes[m]} \to V^{\otimes[\mu]}$ along G-invariant tensors in $(V^*)^{\otimes m-\mu}$ map Z_m*

to Z_μ. Suppose moreover that there exists a G-invariant vector $e_0 \in V$ such that the inclusion maps $\iota : V^{\otimes[m]} \to V^{\otimes[m+1]}$, $\omega \mapsto \omega \otimes e_0$, map Z_m into Z_{m+1} for each $m \in \mathbb{Z}_{\geq 0}$ and $\xi_0 \circ \iota^{n_0} = \mathrm{id}_{V^{\otimes[m]}}$ for each $m \in \mathbb{Z}_{\geq 0}$. Then there exists $M \in m_0 + n_0\mathbb{Z}_{\geq 0}$ such that for all $m \in M + n_0\mathbb{Z}_{>0}$ and for all $\omega \in V^{\otimes[m]}$ the following are equivalent:

(i) *$\xi(\sigma(\omega)) \in Z_\mu$ for all $\sigma \in S_m$, and all contractions $\xi : V^{\otimes[m]} \to V^{\otimes[\mu]}$ along G-invariant tensors in $(V^*)^{\otimes m-\mu}$ with $\mu \leq M$.*
(ii) *$\omega \in Z_m$.*

Proof. The implication (ii)⇒(i) is trivial; we will show the implication (i)⇒(ii). Let Z_∞ be the projective limit in $Y_\infty^{\leq k}$ of Z_m for $m \in m_0 + n_0\mathbb{Z}_{\geq 0}$. By Theorem 6.1, Z_∞ is defined (in A_∞) by finitely many H_∞-orbits of polynomials in $\mathcal{O}_\infty$. This implies that there exists an $M \in m_0 + n_0\mathbb{Z}_{\geq 0}$ such that the H_∞-orbits of the equations of Z_M define Z_∞. We claim that this value of M suffices.

Indeed, suppose that $\omega \in V^{\otimes[m]}$ with $m \in M + n_0\mathbb{Z}_{>0}$ has the property that (for any rearrangement of its terms) all its G-equivariant contractions along tensors to $V^{\otimes[\mu]}$ lie in Z_μ and construct $\omega_\infty \in A_\infty$ as in the proof of Theorem 6.2 (using ι^{n_0} instead of ι). We claim that ω_∞ lies in Z_∞. For this it suffices to show that for each f in the ideal of Z_M and each $h \in H_\infty$ the polynomial hf vanishes on ω_∞. Let $h \in H_\infty$ and let $m' = M + pn_0 = m + p'n_0 \in m + n_0\mathbb{Z}_{\geq 0}$ be such that $h \in H_{m'}$. By construction, $f \in \mathcal{O}_M$ is identified with the function in $\mathcal{O}_{m'}$ obtained by precomposing f with the contraction $V^{\otimes[m']} \to V^{\otimes[M]}$ along the tensor $\xi_0^{\otimes p}$ on the last $m' - M$ factors. Hence hf is the same as contraction $V^{\otimes[m']} \to V^{\otimes[M]}$ along *some* G-invariant tensor (in some of the factors), followed by $h'f$ for some $h' \in H_M$. Evaluating hf at the tensor ω_∞ is the same as evaluating it at

$$\omega \otimes e_0^{\otimes p'n_0},$$

and boils down to contracting some, say l, of the factors e_0 and $m'-M-l$ of the remaining factors V along a tensor in $(V^*)^{\otimes I}$ (with $|I| = m' - M$), and evaluating $h'f$ at the result.

But this is the same thing as first applying some $\sigma \in S_m$ to ω (to ensure the right factors of ω will be contracted), then contracting $\sigma(\omega) \otimes e_0^{\otimes l} \in V^{\otimes[m+l]}$ to an element $\omega' \in V^{\otimes[\mu]}$ along some G-invariant tensor ξ' in $(V^*)^{\otimes m+l-\mu}$ (where $m - \mu = |J|$) and evaluating $h'f$ at $\sigma'(\omega' \otimes e_0^{\otimes M-\mu})$ for some $\sigma' \in S_M$. Note that σ and σ' are merely used to reorganise the terms of ω and $\omega' \otimes e_0^{\otimes M-\mu}$ to avoid some cumbersome notation.

Viewing $e_0^{\otimes l}$ as a contraction from $(V^*)^{\otimes m+l-\mu}$ to $(V^*)^{\otimes m-\mu}$ in the natural way, we have $\tilde{\xi} := e_0^{\otimes l}(\xi') \in (V^*)^{\otimes m-\mu}$. Observe that $\omega' = \tilde{\xi}(\sigma(\omega))$ and $\tilde{\xi}$ is G-invariant since both ξ' and e_0 are.

Now by assumption ω' lies in Z_μ (since $\mu \leq M$), hence $\omega' \otimes e_0^{\otimes M-\mu}$ lies in Z_M and hence $\sigma'(\omega' \otimes e_0^{\otimes M-\mu}) \in Z_M$ as well. This proves that $h'f$ vanishes on it, so that hf vanishes on ω_∞, as claimed. Hence ω_∞ lies in Z_∞. But the projection $A_\infty \to V^{\otimes[m]}$ sends ω_∞ to ω and Z_∞ to Z_m. Hence ω lies in Z_m, as required. □

With these results, we can now prove our main theorems.

Proof of Main Theorem III. By Lemma 3.1 it suffices to show that for fixed $k \in \mathbb{Z}_{\geq 0}$ and for $V = K[G]^n$ for some fixed $n \in \mathbb{Z}_{\geq 0}$ with $n > k$, there exist M, n_0 such that a tensor in $V^{\otimes[m]}$, $m \geq M$, $m \in m_0 + n_0\mathbb{Z}_{\geq 0}$, is of border rank at most k as soon as all its G-equivariant contractions along $m - \mu$-tensors to $V^{\otimes[\mu]}$ have border rank at most k (possibly after rearranging terms).

Recall that we have defined ξ_0 using $x_\chi \in V_\chi^*$. Denoting the trivial character by 0, note that V_0^* is nontrivial since the sum of all basis elements of V is G-invariant, so $x_0 \neq 0$. Moreover, x_0 vanishes outside of V_0, hence there must be an element $e_0 \in V_0$ such that $x_0(e_0) = 1$. For such e_0, observe that $\xi_0(e_0^{\otimes n_0}) = 1$ and e_0 is G-invariant because V_0 is the set of G-invariant elements of V. Now apply Theorem 6.3. □

Our fourth Main Theorem requires a bit more work. We define a *G-spaced star* to be a G-spaced tree for which the underlying tree structure is that of a star.

Lemma 6.4. *Let T be a G-spaced star with centre r and leaves $[m]$. Let $I \subsetneq [m]$ and let T' be the G-spaced star with centre r and leaves $[m] - I$ (and the same spaces attached to each vertex it shares with T). Let ξ be a G-invariant tensor in $\bigotimes_{q\in I} V_i^*$. Then the map $\xi : L(T) \to L(T')$ defined by $\bigotimes_{q\in[m]} v_q \mapsto \xi(\bigotimes_{q\in I} v_q) \cdot \bigotimes_{q\in[m]-I} v_q$ maps $\mathrm{CV}(T)$ to $\mathrm{CV}(T')$.*

Proof. We show that $\xi(\Psi(T)) \subseteq \Psi(T')$. Assume without loss of generality that $I = \{\mu, \ldots, m-1\}$. Let $A = (A_{rq})_{q\sim r} \in \mathrm{rep}_G(T)$. Write $A_{rq} = \sum_{b\in B_r} b \otimes v_{b,q}$ for any $q \in \mathrm{leaf}(T)$. Note that $gA_{rq} = \sum_{b\in B_r} (gb)\otimes(gv_{b,q}) = \sum_{b\in B_r} b\otimes(gv_{g^{-1}b,q})$. Since A_{rq} is G-invariant, we find that $g^{-1}v_{b,q} = v_{g^{-1}b,q}$ for any $b \in B_r$, $g \in G$ and $q \in \mathrm{leaf}(T)$.

Then $\xi(\Psi_T(A)) = \sum_{b\in B_r} \xi(\bigotimes_{p\in I} v_{b,q}) \cdot \bigotimes_{q\in[\mu]} v_{b,q}$. Let $c_b := \xi(\bigotimes_{q\in I} v_{b,q})$. Observe that we now have $c_b = (g\xi)(\bigotimes_{q\in I} v_{b,q}) = \xi(g^{-1} \bigotimes_{p\in I} v_{b,q}) = \xi(\bigotimes_{q\in I} v_{g^{-1}b,q}) = c_{g^{-1}b}$ for any $g \in G$.

For $q \in [\mu-1]$, define $A'_{rq} = A_{rq}$ and define $A'_{r\mu} = \sum_{b\in B_r} b\otimes c_b v_{b,\mu}$. Observe that A'_{rq} is G-invariant for each each $q \in [\mu]$, using $g^{-1}c_b v_{b,\mu} = c_b v_{g^{-1}b,\mu} = c_{g^{-1}b} v_{g^{-1}b,\mu}$ for any $g \in G$ and $b \in B_r$. This means $A' := \mathrm{rep}_G(T')$. We now easily see that $\Psi_{T'}(A') = \xi(\Psi_T(A))$, which after taking the closure concludes the proof. □

Suppose V has a distinguished basis B such that G acts on B. It is easily seen that for a G-spaced star T with centre r, leaves $[m]$ and such that $V_q = V$ for each $q \in [m]$, $\mathrm{CV}(T)$ is H_m-stable. From now on, assume that V has a distinguished basis B such that G acts on B.

Now, for $m \in \mathbb{Z}_{\geq 0}$, let T_m be a G-spaced star with centre r with space V_r and base B_r of cardinality k, leaves $[m]$, and such that $V_q = V$ for each $q \in [m]$. Denote $\mathrm{CV}_m = \mathrm{CV}(T_m)$. Observe that CV_m consists of tensors of rank at most k, hence $\mathrm{CV}_m \subseteq Y_{[m]}^{\leq k}$. Fix $m_0 \in \mathbb{Z}_{\geq 0}$. We can now define $\mathrm{CV}_\infty \subseteq Y_\infty^{\leq k} \subseteq A_\infty$ as the projective limit of the CV_m with $m \in m_0 + n_0\mathbb{Z}_{\geq 0}$. This is the *infinite star model* alluded to in the introduction.

Proposition 6.5. *For any fixed space V_r with basis B_r, the set CV_∞ is the common zero set of finitely many H_∞-orbits of polynomials in $\mathcal{O}_\infty$.*

Proof. As CV_∞ is a closed H_∞-stable subset of $Y_\infty^{\leq k}$ (with $k = |B_r|$) one can apply Theorem 6.1. □

Now, we will see how we can reduce from a star with arbitrary spaces attached to the leaves to a star for which each leaf has the space V attached. This is the analogue of Lemma 3.1 for star models.

Lemma 6.6. *Let $m \in \mathbb{Z}_{\geq 0}$ and suppose T is a G-spaced star with centre r, with space V_r and base B_r of cardinality k, and leaves $[m]$, with spaces V_q for each $q \in [m]$. Let $V = K[G]^n$ for some $n \in \mathbb{Z}_{\geq 0}$ with $n > k$ and let $B = \{gf_i : g \in G, f_i$ is the i-th standard basis vector of $K[G]^n$ viewed as a $K[G]$-module$\}$. If CV_m is defined by polynomials of degree at most D, then so is $\mathrm{CV}(T)$.*

Proof. The set $\mathrm{CV}(T)$ is contained in $L(T) = \bigotimes_{q\in[m]} V_q$ and CV_m is contained in $L(T_m) = \bigotimes_{q\in[m]} V$. Recall that $\mathrm{CV}(T)$ is the Zariski closure of $\Psi(T)$ and CV_m is the closure of the image of $\Psi(T_m)$. A generic element of $\Psi(T)$ is of the form $\sum_{b\in B_r} \bigotimes_{q\in[m]} v_{q,b}$ with $\sum_{b\in B_r} b \otimes v_{q,b}$ a G-invariant element of $V_r \otimes V_q$ for each leaf q. From this, we can easily conclude that any element of $\mathrm{CV}(T)$ has border rank at most k. Likewise, any element of CV_m has border rank at most k.

Suppose $\omega \in L(T) - \mathrm{CV}(T)$. We show that there is an m-tuple of G-linear maps $\phi_q : V_q \to V$ such that $\phi_{[m]}(\omega) \notin \mathrm{CV}_m$. Note that such a $\phi_{[m]}$ maps $\mathrm{CV}(T)$ to CV_m. If this is the case, then we can immediately conclude that there is $f \in \mathcal{O}_{L(T_m)}$ of degree at most D that vanishes on CV_m but not on $\phi_{[m]}(\omega)$, hence $\phi^*_{[m]}(f) \in \mathcal{O}_{L(T)}$ has degree at most D, vanishes on $\mathrm{CV}(T)$ and does not vanish on ω. Hence $\mathrm{CV}(T)$ is defined by polynomials of degree at most D.

If ω has border rank at most k, then by Lemma 3.1, we can find m-tuples of G-linear maps $\phi_q : V_q \to V$ and $\psi_q : V \to V_q$ such that $\psi_{[m]}(\phi_{[m]}(\omega)) = \omega$. Since $\psi_{[m]}(\phi_{[m]}(\omega)) \notin \mathrm{CV}(T)$ by assumption (and $\psi_{[m]}(\mathrm{CV}_m) \subseteq \mathrm{CV}(T)$), we conclude that $\phi_{[m]}(\omega) \notin \mathrm{CV}_m$.

If ω has border rank exceeding k, then by Lemma 3.1, there is an m-tuple of G-linear maps $\phi_i : V_i \to V$ such that $\phi_{[m]}(\omega)$ has border rank exceeding k, which implies $\phi_{[m]}(\omega) \notin \mathrm{CV}_m$. □

Remark 6.7. (i) We may in fact assume $n = k$; in this case, we first test whether some flattening of ω has rank exceeding k; this can be done by equations of degree $k + 1$. If not, then we can find m-tuples of G-linear maps $\phi_q : V_q \to V$ and $\psi_q : V \to V_q$ such that $\psi_{[m]}(\phi_{[m]}(\omega)) = \omega$ and proceed with the proof as above.

(ii) If V_r has multiplicity k_χ for each irreducible representation χ, then we may take $V = K[G]^{\max_\chi\{k_\chi\}}$ instead of $K[G]^n$. In fact, we may use $V = V_r$, using the fact that because of the given basis of V, we have $k_\chi = k_{\chi^{-1}}$ for each χ.

Moreover, observe that $\mathrm{CV}_m \subseteq Y_{[m]}^{\leq (k_\chi)_\chi}$.

Example 6.8. If $B = G$, then $K[G] \cong \bigoplus_{\chi\in\widehat{G}} \chi$ (identifying characters and irreducible representations in the natural way), and hence if $V_r = K[G]$, then $\mathrm{CV}_m \subseteq Y_{[m]}^{\leq (1)_\chi}$.

We now show that the (Zariski closure of the) equivariant model for a G-spaced star is defined in bounded degree, given a bound on the cardinality of the basis of the centre of the star. After we show this, we can finally prove Main Theorem IV.

Theorem 6.9. *Let V_r be a G-module with basis B_r of cardinality $k \in \mathbb{Z}_{\geq 0}$. Then there exists $D \in \mathbb{Z}_{\geq 0}$ such that for each $m \in \mathbb{Z}_{\geq 0}$ and each G-spaced star T with centre r with leaves $[m]$, the set $\mathrm{CV}(T)$ is defined by the vanishing of a number of polynomials of degree at most D.*

Proof. By Lemma 6.6 it suffices to prove that for fixed $k \in \mathbb{Z}_{\geq 0}$ and $V = K[G]^n$ with $n > k$, there exists a $D \in \mathbb{Z}_{\geq 0}$ such that for all $m_0 \in [n_0]$ and for all $m \in m_0 + n_0\mathbb{Z}_{\geq 0}$ the variety CV_m is defined in $V^{\otimes[m]}$ by polynomials of degree at most D.

As in the proof of Main Theorem III, observe that there is some G-invariant element e_0 such that $\xi_0(e_0^{\otimes n_0}) = 1$. Let $\epsilon_0 = e_0^{\otimes n_0}$.

Consider the inclusion maps

$$\iota : V^{\otimes m} \to V^{\otimes m+n_0}, \quad \omega \mapsto \omega \otimes \epsilon_0.$$

Observe that $\iota(\Psi_{T_m}(A)) = \Psi_{T_{m+n_0}}(A')$ where $A'_{rq} := A_{rq}$ if $q \in [m]$ and $A'_{rq} = (\sum_{b\in B_r} b) \otimes e_0$ otherwise. Moreover, each A'_{rq} is G-invariant.

Hence this map sends CV_m into CV_{m+n_0} and we easily see that it satisfies $\xi_0 \circ \iota = \mathrm{id}_{V^{\otimes[m]}}$.

Thus we can apply Theorem 6.2. $\square$

Example 6.10. Let $B = G = \mathbb{Z}/2\mathbb{Z}$ and let T be a G-spaced tree with m leaves with space $V = K[G]$ attached to each node. Let $y_0, y_1 \in V^*$ be a basis dual to the basis $e+g$, $e-g$ of V.

Using the proof in [SS05] that the group-based model for $\mathbb{Z}/2\mathbb{Z}$ is defined by linear and quadratic polynomials, we can show that CV_m is defined by the H_∞-orbits of $\zeta_{m,u}$ where the cardinality of $\{i \in [m] : u_i = 1\}$ is odd and by the H_∞-orbits of $\zeta_{m,u_0}\zeta_{m,u_1} - \zeta_{m,u_2}\zeta_{m,u_3}$ such that:

(a) For each $i \in [m]$, the multiset $\{(u_0)_i, (u_1)_i\}$ equals the multiset $\{(u_2)_i, (u_3)_i\}$.
(b) For each $j \in \{1, 2, 3, 4\}$, the cardinality of $\{i \in [m] : (u_j)_i = 1\}$ is even.

Note the similarity with Example 4.3. Indeed, these equations all vanish on $Y_{[m]}^{\leq(1,1)}$, hence $Y_{[m]}^{\leq(1,1)} \subseteq \mathrm{CV}_m$ and therefore $Y_{[m]}^{\leq(1,1)}$ is in fact equal to CV_m in this specific case. By Example 4.3, we find that we can take $M = 8$ in Theorem 6.3. A more precise examination shows that we may take $M = 5$ in this case.

Proof of Main Theorem IV. Let T be a G-spaced tree (over an algebraically closed field of characteristic 0) satisfying the conditions of the theorem. By Theorem 1.7 in [DK09], one has $I(\mathrm{CV}(T)) = \sum_{r\in\mathrm{vert}(T)} I(\mathrm{CV}(\flat_r T))$ where $\flat_r T$ is a G-spaced star with centre r. From this, we can easily conclude that if $\mathrm{CV}(\flat_r T)$ is defined by polynomials of degree at most D for each r, then so is $\mathrm{CV}(T)$. Now apply Theorem 6.9. $\square$

Remark 6.11. The proof of this theorem, along with the previous remark, shows that to describe the equations that define the equivariant model for any G-spaced tree, it suffices to describe the equations that define the equivariant model for any G-spaced star for which all nodes have the same space attached.

Proof of Main Theorem I. For the field $K = \mathbb{C}$, by Main Theorem IV there is $D \in \mathbb{Z}_{\geq 0}$ depending on G and $k = |B|$ such that $\mathrm{CV}(T)$ is defined by polynomials of degree at most D. The tensorification of the model in the introduction is the closure of the set of tensors of the form $\Psi(A)$ with $A \in \mathrm{rep}_G(T)$ such that A satisfies an additional set of linear equalities and inequalities (certain sums must be equal to 1 and certain coefficients must be nonnegative). Since Ψ is linear, these translate to linear equalities and inequalities for $\Psi(A)$. Then, clearly, the closure of the set of tensors of the form $\Psi(A)$ with A in $\mathrm{rep}_G(T)$ such that A satisfies the linear *equalities* mentioned is defined by polynomials of degree at most $\max(D, 1)$, since linear equalities can be tested by linear polynomials. The latter however equals the closure of the set of tensors of the form $\Psi(A)$ with A in $\mathrm{rep}_G(T)$ such that A satisfies both the linear equalities and the inequalities. Hence the tensorification of the model in the introduction is defined by polynomial equations of degree at most $\max(D, 1)$. □

Proof of Main Theorem II. Let $\omega \in L(T)$. We will first test whether $\omega \in \mathrm{CV}(T)$; after that, we can verify whether ω satisfies the additional linear equalities mentioned in Main Theorem I. For each vertex r, view ω as an element of $\flat_r(T)$; say $\flat_r(T)$ has leaves $[m]$ and space V_q for each $q \in [m]$. Use the construction of Lemma 3.1 to produce $\phi_{[m]}$, $\psi_{[m]}$ such that $\psi_{[m]}(\phi_{[m]}(\omega)) = \omega$, where $\phi_q : V_q \to V = K[G]^{|B|+1}$. If some flattening of ω occurring in the construction has image of rank exceeding $k = |B|$, then conclude that $\omega \notin \mathrm{CV}(T)$.

Consider $\omega' = \phi_{[m]}(\omega)$. Take M as in Theorem 6.3. Let I be a subset of $[m]$ of cardinality pn_0 with $m - pn_0 \leq M$; the number of such subsets is polynomial in m (it is $O(m^M)$).

Take a basis $\xi_1, \ldots, \xi_N$ of G-invariant tensors in $(V^*)^{\otimes I}$; let $f_1, \ldots, f_{N'}$ be a set of polynomials that defines $\mathrm{CV}(T_{[m]-I})$. We can symbolically describe the composition of a contraction of ω' along the formal linear combination $\sum x_i \xi_i$ with some f_j as a polynomial and test whether this polynomial is identically 0. If the latter is true for all I and for all flattenings, then conclude that ω lies in $\mathrm{CV}(T)$ because of Theorem 6.3. □

The set-up of our algorithm (given M) starting from T_m is as follows. In the deterministic setting:

Precomputation: Compute, once and for all, a set E_μ of equations for CV_μ for all $\mu \leq M$.
Input: $\omega \in V^{\otimes [m]}$.
Output: True or false (the answer to the question whether $\omega \in \mathrm{CV}_m$).
Algorithm: For each $I \subseteq [m]$ with $|I| \geq m - M$, check whether the composition of the equations in $E_{m-|I|}$ with the formal contraction of ω along a general G-invariant element of $(V^*)^{\otimes I}$ is identically 0. If this is the case for all I, then output 'true', else output 'false'.

The number of scalar arithmetic operations in this algorithm is bounded by a polynomial in d^m, where the degree of that polynomial depends on the degrees of the equations found in the pre-computation step. Observe that running with I over all sufficiently large subsets of $[m]$ contributes only a factor $\mathcal{O}(m^M)$, which is poly-logarithmic in d^m. In the probabilistic setting:

Precomputation: Compute, once and for all, a set E_μ of equations for CV_μ for all $\mu \leq M$.
Input: $\omega \in V^{\otimes[m]}$.
Output: True or false (the (probable) answer to the question $\omega \in \mathrm{CV}_m$?).
Algorithm: For each $I \subseteq [m]$ with $|I| \geq m - M$, generate a random element ξ of $(V^*)^{\otimes I}$ and compute whether all equations in $E_{m-|I|}$ vanish on $\xi(\omega)$ (with ξ viewed as a contraction $V^{\otimes[m]} \to V^{\otimes[m]-I}$). If this is the case for all I, then output 'true', else output 'false'.

The number of scalar arithmetic operations in this case is linear in $d^m \cdot m^M$.

Acknowledgments. We thank Bernd Sturmfels and an anonymous referee for many helpful suggestions for improving the text.

Both authors are supported by the first author's Vidi grant from the Netherlands Organisation for Scientific Research (NWO)

References

[AR08] Allman, E. S., Rhodes, J. A.: Phylogenetic ideals and varieties for the general Markov model. Adv. Appl. Math. **40**, 127–148 (2008) Zbl 1131.92046 MR 2388607

[BO11] Bates, D. J., Oeding, L.: Toward a salmon conjecture. Experiment. Math. **20**, 358–370 (2011) Zbl 1262.14056 MR 2836258

[BSS89] Blum, L., Shub, M., Smale, S.: On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines. Bull. Amer. Math. Soc. **21**, 1–46 (1989) Zbl 0681.03020 MR 0974426

[Bor91] Borel, A.: Linear Algebraic Groups. Springer, New York (1991) Zbl 0726.20030 MR 1102012

[BW07] Buczyńska, W., Wiśniewski, J. A.: On geometry of binary symmetric models of phylogenetic trees. J. Eur. Math. Soc. **9**, 609–635 (2007) Zbl 1147.14027 MR 2314109

[CFS08] Casanellas, M., Fernández-Sánchez, J.: The geometry of the Kimura 3-parameter model. Adv. Appl. Math. **41**, 265–292 (2008) Zbl 1222.14110 MR 2449591

[CFS11] Casanellas, M., Fernández-Sánchez, J.: Relevant phylogenetic invariants of evolutionary models. J. Math. Pures Appl. (9) **96**, 207–229 (2011) Zbl 1230.14063 MR 2831580

[CS05] Casanellas, M., Sullivant, S.: The strand symmetric model. In: Algebraic Statistics for Computational Biology, Cambridge Univ. Press, Cambridge, 305–321 (2005) MR 2205881

[Cip07] Cipra, B. A.: Algebraic geometers see ideal approach to phylogenetics. SIAM News **40** no. 6 (2007)

[Coh67] Cohen, D. E.: On the laws of a metabelian variety. J. Algebra **5**, 267–273 (1967) Zbl 0157.34802 MR 0206104

[Dra10] Draisma, J.: Finiteness for the k-factor model and chirality varieties. Adv. Math. **223**, 243–256 (2010) Zbl 1211.13020 MR 2563217

[DK09] Draisma, J., Kuttler, J.: On the ideals of equivariant tree models. Math. Ann. **344**, 619–644 (2009) Zbl 05565281 MR 2501304

[DK14] Draisma, J., Kuttler, J.: Bounded-rank tensors are defined in bounded degree. Duke Math. J. **163**, 35–63 (2014) Zbl 06279924 MR 3161311

[FG12] Friedland, S., Gross, E.: A proof of the set-theoretic version of the salmon conjecture. J. Algebra **356**, 374–379 (2012) Zbl 1258.14001 MR 2891138

[GSS05] Garcia, L. D., Stillman, M., Sturmfels, B.: Algebraic geometry of Bayesian networks. J. Symbolic Comput. **39**, 331–355 (2005) Zbl 1126.68102 MR 2168286

[HS12] Hillar, C. J., Sullivant, S.: Finite Gröbner bases in infinite dimensional polynomial rings and applications. Adv. Math. **221**, 1–25 (2012) Zbl 1233.13012 MR 2854168

[LM04] Landsberg, J. M., Manivel, L.: On the ideals of secant varieties of Segre varieties. Found. Comput. Math. **4**, 397–422 (2004) Zbl 1068.14068 MR 2097214

[Mic13] Michałek, M.: Constructive degree bounds for group-based models. J. Combin. Theory Ser. A **120**, 1672–1694 (2013) MR 3092693

[Mic14] Michałek, M.: Toric geometry of the 3-Kimura model for any tree. Adv. Geom. **14**, 11–30 (2014) Zbl 1300.14051 MR 3159089

[PS05] Pachter, L., Sturmfels, B. (eds.): Algebraic Statistics for Computational Biology. Cambridge Univ. Press, Cambridge (2005) Zbl 1108.62118 MR 2205865

[Rai12] Raicu, C.: Secant varieties of Segre–Veronese varieties. Algebra Number Theory **6**, 1817–1868 (2012) Zbl 1273.14102 MR 3033528

[RS95] Robertson, N., Seymour, P. D.: Graph minors. XIII: The disjoint paths problem. J. Combin. Theory Ser. B **63**, 65–110 (1995) Zbl 0823.05038 MR 1309358

[RS04] Robertson, N., Seymour, P. D.: Graph minors. XX: Wagner's conjecture. J. Combin. Theory Ser. B **92**, 325–357 (2004) Zbl 1061.05088 MR 2099147

[SS05] Sturmfels, B., Sullivant, S.: Toric ideals of phylogenetic invariants. J. Comput. Biology **12**, 204–228 (2005)