5-1-2021

Representation Stability of the Cohomology of Springer Varieties and Some Combinatorial Consequences

Aba Mbirika
University of Wisconsin-Eau Claire

Julianna Tymoczko
Smith College, jtymoczko@smith.edu

Recommended Citation

REPRESENTATION STABILITY OF THE COHOMOLOGY OF SPRINGER VARIETIES AND SOME COMBINATORIAL CONSEQUENCES

ABA MBIRIKA AND JULIANNA TYMOCZKO

ABSTRACT. A sequence of S_n-representations $\{V_n\}$ is said to be uniformly representation stable if the decomposition of $V_n = \bigoplus_\mu c_{\mu,n} V(\mu)_n$ into irreducible representations is independent of n for each μ—that is, the multiplicities $c_{\mu,n}$ are eventually independent of n for each μ. Church-Ellenberg-Farb proved that the cohomology of flag varieties (the so-called diagonal coinvariant algebra) is uniformly representation stable. We generalize their result from flag varieties to all Springer fibers. More precisely, we show that for any increasing subsequence of Young diagrams, the corresponding sequence of Springer representations form a graded co-FI-module of finite type (in the sense of Church-Ellenberg-Farb). We also explore some combinatorial consequences of this stability.

1. INTRODUCTION

Homological stability is a topological property of certain sequences of topological spaces: given a sequence $\{X_n\}_{n=1}^\infty$ of topological spaces with maps $\phi_n : X_n \to X_{n+1}$ for each n, then $\{X_n\}_{n=1}^\infty$ is homologically stable if there exists a positive integer N so that the maps $(\phi_n)_* : H_i(X_n) \to H_i(X_{n+1})$ is an isomorphism whenever $n \geq N$. In other words the homology groups stabilize after a certain point in this sequence; however the topological spaces change later in the sequence, the changes do not affect the i^{th} homology group.

Church and Farb [9], and later Church, Ellenberg, and Farb [6] defined *representation stability* to mimic the topological definition. Informally, the sequence of S_n-representations $V_1 \xrightarrow{f_1} V_2 \xrightarrow{f_2} V_3 \xrightarrow{f_3} \cdots$ is representation-stable if there are S_n-equivariant linear injections $f_n : V_n \to V_{n+1}$ and if there exists an N so that the multiplicities $c_{\mu,n}$ in the decomposition into irreducibles

$$V_n = \bigoplus_\mu c_{\mu,n} V(\mu)_n$$

are independent of n for all $n \geq N$. (Section 6 defines representation stability precisely.)

One important problem is to find families of representations that are representation stable; these families often arise from geometric considerations. In their original work, Church, Ellenberg, and Farb identify a number of representation stable families arising from geometry/topology and classical representation theory [6]. Indeed their example of the diagonal coinvariant algebra [6, Section 5] is the sequence of cohomology rings of flag varieties $\{H^*(GL_n/B_n)\}_{n=1}^\infty$ which in some sense is the springboard of this paper. Since then, others have demonstrated that representation stability arises naturally in many contexts, including

arrangements associated to root systems [2], linear subspace arrangements [16], configuration spaces in $\mathbb{R}^d$ [18], filtrations of Torelli groups [25], moduli spaces of Riemann surfaces of genus g with n labeled marked points [27], and others [12, 28].

The central goal of this paper is to prove that an important family of S_n-representations called *Springer representations* is representation stable. The Springer representation is the archetypal geometric representation: in its most basic form, it arises when the symmetric group S_n acts on the cohomology of a family of subvarieties of the flag variety called *Springer fibers*. The flag variety can be described as the set of nested vector subspaces

$$V_1 \subseteq V_2 \subseteq \cdots \subseteq V_{n-1} \subseteq \mathbb{C}^n$$

where each V_i is i-dimensional. Given a nilpotent matrix X, the Springer fiber consists of the flags that are fixed by X in the sense that $XV_i \subseteq V_i$ for all i. Every nilpotent matrix is conjugate to one in Jordan form, which is determined by a partition of n into Jordan blocks. Since Springer fibers associated to conjugate matrices X and gXg^{-1} are homeomorphic, the Springer fibers are parametrized by partitions of n.

Springer first constructed a representation of the symmetric group S_n on the cohomology of Springer fibers [29]. Since then, the representation has been recreated in many different ways [10, 30, 4, 23]. The geometry of Springer fibers encodes key data about representations of the symmetric group. For instance the top-dimensional cohomology is an irreducible S_n-representation [29]; and the ungraded representation on $H^*(Spr_\lambda)$ is Young's representation associated to the partition λ [17, Introduction].

Our analysis of Springer representations uses the **co-FI** category, which Church, Ellenberg, and Farb defined to concisely describe the compatibility conditions needed for representation stability. Theorem 3 proves that for any increasing sequence of Young diagrams, the corresponding Springer representations form a graded co-FI-module of finite type. The main consequence from our point of view is that sequences of Springer representations are representation stable (see Corollary 3 for a precise statement).

From this many other properties follow: (1) for each fixed degree and n large enough, the character is given by a polynomial that is independent of n, and in particular (2) the dimension of the Springer representation is eventually polynomial in n. It is this second consequence that we explore in Section 7.

Remark 1. Curiously, the literature does not make a clear distinction between the Springer representation and its dual. Hotta first observed this and classified existing constructions of the Springer representation up to that point [19]. However, the ambiguity persisted with subsequent constructions of the Springer representation. In this paper, we treat "the" Springer representation interchangeably with its dual. We prove that the Garsia-Procesi construction of the Springer representation is a co-FI-module (graded, of finite type) and its *dual* is representation stable. Properties like dimension and decomposition into irreducibles are well-behaved with respect to duality and apply to the original Garsia-Procesi construction, too.

Kim proves a different kind of stability of Springer representations, giving conditions under which the lower-graded parts of the Springer representation for λ coincide with those of λ' for partitions λ, λ' of the same n [21]. Kim recovers Theorem 6 using his notion of stability [21, Corollary 4.3].

This paper is organized as follows. In Section 2, we define the conditions that guarantee the stability of a sequence of S_n-representations. In Section 3, we describe Garsia-Procesi's combinatorial description of the Springer representation, and in particular the so-called *Tanisaki ideal*. In Section 4, we prove there is no FI-module structure on sequences of Springer representations, except for the trivial representation. On the other hand, in Section 5 we prove that the sequence of Tanisaki ideals forms a co-FI-ideal and deduce that a co-FI-module structure exists on all sequences of Springer representations. In Section 6, we conclude with our main result on the stability of the Springer representations. We give some concrete combinatorial consequences of this stability in Section 7, and a collection of open questions in Section 8 that probe the new combinatorial ideas raised by representation stability.

2. **FI** AND **co-FI**

In this section we describe FI-modules and co-FI-modules, defined by Church, Ellenberg, and Farb [6] to streamline and extend the essential features of Church and Farb's earlier notion of representation stability [9]. We begin with the categories **FI** and **co-FI**, which carry actions of the permutation groups S_n and are constructed to be compatible with inclusions. The key example of FI- and co-FI-modules for this manuscript is the sequence of polynomial rings $\{k[x_1, \ldots, x_n]\}$. We then list the properties about FI- and co-FI-modules that we will need to establish that sequences of Springer representations are graded co-FI-modules.

This section provides only what is needed in this paper. The interested reader is referred to Church, Ellenberg, and Farb's work for many other interesting results [6].

Remark 2. We assume k is a field of characteristic zero. Parts of Church-Ellenberg-Farb's theory extends to other fields as well [6].

Definition 1 (FI-module, FI-algebra graded FI-algebra)**.** Let $\mathbf{n}$ denote the set $\{1, \ldots, n\}$. **FI** is the category whose objects are finite sets and whose morphisms are injections. This is equivalent to the category whose objects are sets $\mathbf{n} = \{1, 2, \ldots, n\}$ and whose morphisms are injections $\mathbf{m} \to \mathbf{n}$. An *FI-module* (*FI-algebra, graded FI-algebra*) over a commutative ring k is a functor V from **FI** to the category of modules over k (k-algebras, graded k-algebras). We usually denote the k-module (respectively algebra) $V(\mathbf{n})$ by V_n.

Modules that carry permutation actions compatible with the permutation action on integers provide a rich source of examples of FI-modules. The next example is the most important for our purposes.

Example 1 ([6, Example 4.1.2 and Remark 4.1.3])**.** Define a functor R by:
- R sends the object $\mathbf{n}$ to the polynomial ring $k[x_1, \ldots, x_n]$ and
- R sends the morphism $f : \mathbf{m} \to \mathbf{n}$ to the homomorphism

$$f_* : k[x_1, \ldots, x_m] \to k[x_1, \ldots, x_n]$$

  induced by the condition that $f_*(x_i) = x_{f(i)}$ for all $i \in \mathbf{m}$.

Then R is a graded FI-algebra. In particular, we have that R_n is the polynomial ring $k[x_1, \ldots, x_n]$.

The category **co-FI** is opposite to **FI**. The polynomial algebras also form a graded co-FI-algebra, as described below.

Definition 2 (co-FI-module, co-FI-algebra, graded co-FI-algebra). Denote the opposite category of **FI** by **co-FI**. In particular the objects in **co-FI** are finite sets, without loss of generality the sets **n**, and the morphisms in **co-FI** from **n** to **m** are the morphisms in **FI** from **m** to **n**. A *co-FI-module* (*co-FI-algebra, graded co-FI-algebra*) over a commutative ring k is a functor from **co-FI** to the category of k-modules (k-algebras, graded k-algebras).

For instance if V is an FI-module over k then the dual $V^\vee$ forms a co-FI-module, and vice versa.

Example 2 ([6, Example 4.1.2 and Remark 4.1.3]). Define a functor R as follows:
- R sends the object **n** to the polynomial ring $k[x_1, \ldots, x_n]$ and
- R sends the FI-morphism $f : \mathbf{m} \to \mathbf{n}$ to the homomorphism

$$f^* : k[x_1, \ldots, x_n] \to k[x_1, \ldots, x_m]$$

induced by the condition that

$$f^*(x_i) = \begin{cases} x_{f^{-1}(i)} & \text{if } i \in \text{Im}(f) \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

Then R is a graded co-FI-algebra.

The module categories **FI** and **co-FI** are abelian and so admit many of the algebraic constructions that modules and algebras do. In particular we can consider co-FI-quotients of a co-FI-module and FI-submodules of an FI-module.

Example 3 ([6, **Classical coinvariant algebra** in Section 5]). For each n let I_n denote the ideal of symmetric polynomials with no constant term. The sequence of ideals I is not an FI-submodule of the FI-module R because the image of an S_m-symmetric polynomial under the inclusion map $\iota : \{1, \ldots, m\} \to \{1, \ldots, n\}$ is not symmetric under the larger group S_n. However the ideals I do form a co-FI-submodule of the co-FI-module R.

The next lemma proves that the image of a set of S_m-invariant polynomials under an arbitrary injection is the same as the image under the inclusion ι that is the identity on the integers $\{1, \ldots, m\}$. We use it to simplify later calculations.

Lemma 1. *Suppose that $m \leq n$ and that $f : \{1, \ldots, m\} \to \{1, \ldots, n\}$ is an injection. Let ι be the inclusion $\iota : \{1, \ldots, m\} \to \{1, \ldots, n\}$ that sends $\iota(i) = i$ for each $1 \leq i \leq m$.*

Consider the action of S_n on R_n under which for each $w \in S_n$ and $x_i \in R_n$ we have $w(x_i) = x_{w(i)}$. Let $\mathcal{J} \subseteq R_n$ be any set of polynomials that are preserved under the S_n-action, in the sense that $w\mathcal{J} \subseteq \mathcal{J}$ for all permutations $w \in S_n$. Then

$$f^*(\mathcal{J}) = \iota^*(\mathcal{J}).$$

Proof. Our hypothesis means that the polynomial $p \in \mathcal{J}$ if and only if the image $w(p) \in \mathcal{J}$ for each $w \in S_n$. The map $f : \{1, \ldots, m\} \hookrightarrow \{1, \ldots, n\}$ is an injection so we can define a permutation $w_f \in S_n$ by

$$w_f = \begin{cases} f(i) \mapsto i & \text{for all } i \text{ with } 1 \leq i \leq m \\ j_i \mapsto m + i & \text{for } j_i \text{ such that both } j_1 < j_2 < \cdots < j_{n-m} \text{ and} \\ & \quad \{j_1, j_2, \cdots, j_{n-m}\} \cup \text{Im}(f) = \{1, 2, \ldots, n\}. \end{cases}$$

We know that for each i

$$\iota^*(w_f x_i) = \iota^*(x_{w_f(i)})$$

by definition of the S_n-action on permutations. By construction of ι^* we have

$$\iota^*(x_{w_f(i)}) = \begin{cases} 0 & \text{if } i \in \{j_1, j_2, \cdots, j_{n-m}\} \\ x_{f^{-1}(i)} & \text{otherwise.} \end{cases}$$

This is exactly $f^*(x_i)$. Thus $f^*(p) = \iota^*(w_f p)$ for all polynomials $p \in \mathcal{J}$ and hence as desired $f^*(\mathcal{J}) = \iota^*(\mathcal{J})$. □

It follows that if I is a sequence of symmetric homogeneous ideals in the graded co-FI-algebra R then we can prove I is a co-FI-submodule simply by considering inclusions—or even just a subset of inclusions.

Corollary 1. *Let R be a graded co-FI-algebra and I be a sequence of ideals $\{I_1, I_2, \ldots\}$ with each I_n an S_n-invariant homogeneous ideal in R_n. The following are equivalent:*

(1) *I is a co-FI-submodule of R.*
(2) *For each $m < n$ and inclusion $\iota : \{1, \ldots, m\} \to \{1, \ldots, n\}$ defined by $\iota(i) = i$ for all i, the induced map satisfies $\iota^*(I_n) \subseteq I_m$.*
(3) *For each n and inclusion $\iota_n : \{1, \ldots, n-1\} \to \{1, \ldots, n\}$, the induced map satisfies $\iota_n^*(I_n) \subseteq I_{n-1}$.*

Proof. Part (1) is equivalent to Part (2) by Lemma 1. Part (2) implies Part (3) by definition. The composition of inclusions $\iota_n \circ \iota_{n-1} \circ \cdots \circ \iota_{m+1}$ is the inclusion $\iota : \{1, \ldots, m\} \hookrightarrow \{1, \ldots, n\}$. By functoriality if $\iota_i^*(I_i) \subseteq I_{i-1}$ for each $i \geq 2$ then $\iota^*(I_n) \subseteq I_m$ for each pair $n \geq m \geq 1$. So Part (3) implies Part (2). □

The definition of finitely-generated FI-modules is crucial to representation stability. It differs importantly from the corresponding definition for modules or rings because it incorporates the underlying S_n-action. As we see in Example 4, this implies that FI-modules often have fewer generators than we might expect.

In fact, though this does not appear explicitly in the literature, the category of *finitely generated* **FI** modules over Noetherian rings is also abelian (by the Noetherian property, proven over $\mathbb{C}$ in [6] and over other Noetherian rings in [8]) [5]. This is the thrust of the arguments that we cite in this paper.

Definition 3 (Finite generation, finite type)**.** An FI-module V is *finitely generated* if there is a finite set S of elements in $\coprod_i V_i$ so that no proper sub-FI-module of V contains S. A graded FI-module V has *finite type* if the i^{th} graded part V^i is finitely generated for each i. A graded co-FI-module W is of finite type if its dual $W^\vee$ is a graded FI-module of finite type [6, Co-FI-algebras in Section 4.2]

Example 4 ([6, **Graded FI-modules of finite type** in Section 4.2], [13, Example 4.2], [31, Example 1.4])**.** The sequence $R = \{k[x_1, \ldots, x_n]\}$ of polynomial rings are not finitely generated as an FI-module; indeed no ring $k[x_1, \ldots, x_n]$ is a finite dimensional k-vector space. However when graded by polynomial degree, each graded part of the sequence R is finitely generated as an FI-module. For instance the graded part of degree 3 is generated by x_1^3, $x_1^2 x_2$, and $x_1 x_2 x_3$ since every monomial of degree three is obtained by permuting indices of one of these three. More generally the graded part of degree d is generated by all monomials of the form $x_1^{d_1} x_2^{d_2} \cdots x_n^{d_n}$ over partitions $d_1 \geq d_2 \geq \cdots \geq d_n \geq 0$ of d.

For completeness, we define the quotient of an FI-algebra or co-FI-algebra.

Definition 4 ([7, Definition 2.76])**.** Let R be a graded FI-algebra. If I is a graded FI-submodule (respectively co-FI-submodule) for which each object I_n is a homogeneous ideal in R_n then I is called an *FI-ideal* (respectively *co-FI-ideal*). The quotient FI-module R/I is defined so that $(R/I)_n = R_n/I_n$ for each n (respectively co-FI).

The following proposition is our main tool, and sketches the main points of [8, Theorem F]. This result implies that we only need to prove each sequence of Tanisaki ideals forms a co-FI-submodule; after that, a straightforward algebraic argument allows us allow us to conclude that sequences of Springer representations are representation-stable.

Proposition 1. *Let R be a graded co-FI-module over a Noetherian ring and let I be a co-FI-ideal. Then the dual $(R/I)^\vee$ is a graded FI-module. If $R^\vee$ has finite type then so does $(R/I)^\vee$.*

Proof. The dual of a graded co-FI-module is a graded FI-module by purely formal properties. Suppose further that $R^\vee$ has finite type and consider the j^{th} graded parts $(R_n)_j^\vee$ and $(R_n/I_n)_j$ for each n. The dual $(R_n/I_n)_j^\vee$ is the FI-submodule of $(R_n)_j^\vee$ consisting of those functionals that vanish on the j^{th} graded parts $(I_n)_j = I_n \cap (R_n)_j$ by definition of graded quotients. Each FI-submodule of an FI-module of finite type over a Noetherian ring is also of finite type [8, Theorem A], thus proving the claim.   □

In Section 7 and 8 of this manuscript we analyze properties of FI-modules in the case of Springer representations.

3. Springer theory

This section summarizes two key combinatorial tools: Biagioli-Faridi-Rosas's description of generators for each Tanisaki ideal and Garsia-Procesi's description of a basis for the Springer representation.

3.1. **The Tanisaki ideal and its generators.** We use Garsia-Procesi's presentation of the cohomology of the Springer fiber, which describes the cohomology as a quotient of a polynomial ring analogous to the Borel construction of the cohomology of the flag variety [3]. Let R_n be the polynomial ring $\mathbb{C}[x_1, \ldots, x_n]$. For each partition λ of n Tanisaki defined an ideal $I_\lambda \subseteq R_n$ that is now called the *Tanisaki ideal*. To describe the Tanisaki ideal, we define certain sets of elementary symmetric functions.

Definition 5. Given a subset $S \subseteq \{x_1, \ldots, x_n\}$ the elementary symmetric function $e_i(S)$ is the polynomial

$$e_i(S) = \sum_{\substack{T \text{ such that} \\ T \subseteq S \text{ and } |T| = i}} \prod_{x_j \in T} x_j.$$

The set $E_{i,j}^n$ is the set of elementary symmetric functions $e_i(S)$ over all possible subsets $S \subseteq \{x_1, \ldots, x_n\}$ of cardinality j.

Example 5. Let $S = \{x_1, x_3, x_4\} \subseteq \{x_1, x_2, \ldots, x_5\}$. Then $e_2(S) = x_1 x_3 + x_1 x_4 + x_3 x_4$. Letting S run over all $\binom{5}{3} = 10$ size three subsets of $\{x_1, x_2, \ldots, x_5\}$ gives the 10 elementary symmetric polynomials $e_2(S)$ comprising the set $E_{2,3}^5$.

We follow Biagioli, Faridi and Rosas's construction of the Tanisaki ideal [1, Definition 3.4].

Definition 6. Let $\lambda = (\lambda_1, \ldots, \lambda_k)$ be a partition of n. The *BFR-filling* of λ is constructed as follows. From the leftmost column of λ to the rightmost column, place the numbers $1, 2, \ldots, n - \lambda_1$ bottom to top skipping the top row. Finally fill the top row from right to left with the remaining numbers $n - \lambda_1 + 1, \ldots, n$. The *BFR-generators* are polynomials in the set $\mathcal{G}(\lambda)$ defined as the following union: If the box filled with i has j in the top row of its column, then include the elements of $E_{i,j}^n$ in the set $\mathcal{G}(\lambda)$.

Example 6. Consider the partition $\lambda = (4, 2, 2, 1)$. Then the BFR-filling of λ is

$$
\begin{array}{|c|c|c|c|}
\hline
9 & 8 & 7 & 6 \\
\hline
3 & 5 \\
\cline{1-2}
2 & 4 \\
\cline{1-2}
1 \\
\cline{1-1}
\end{array}
$$

and hence the set $\mathcal{G}(\lambda)$ is

$$
\mathcal{G}(\lambda) = \bigcup_{i \in \{1,2,3,9\}} E_{i,9}^9 \;\cup\; \bigcup_{i \in \{4,5,8\}} E_{i,8}^9 \;\cup\; E_{7,7}^9 \;\cup\; E_{6,6}^9.
$$

The first union of sets arises from the first column, the second from the second, and so on. Each set of the form $E_{i,9}^9$ contains the single elementary symmetric function $e_i(x_1, x_2, \ldots, x_9)$.

Theorem 1 (Biagioli-Faridi-Rosas [1, Corollary 3.11]). *Given a partition λ of n the Tanisaki ideal I_λ is generated by the set $\mathcal{G}(\lambda)$.*

3.2. **Combinatorial presentation for the Springer representation and the Garsia-Procesi basis.** It turns out that the natural S_n-action on R_n given by

$$
w \cdot p(x_1, \ldots, x_n) = p(x_{w(1)}, \ldots, x_{w(n)})
$$

restricts to the Tanisaki ideal I_λ. Thus the S_n-action on R_n induces an S_n-action on R_n/I_λ. The key result of Garsia-Procesi's work (with others [22, 10, 30]) is that this quotient is the cohomology of the Springer variety.

Proposition 2 (Garsia-Procesi [17]). *Let $R_n = \mathbb{C}[x_1, \ldots, x_n]$ and let I_λ be the Tanisaki ideal corresponding to the partition λ. The quotient R_n/I_λ is isomorphic to the cohomology of the Springer variety $H^*(Spr_\lambda)$ as a graded S_n-representation.*

Moreover Garsia-Procesi describe an algorithm to compute a nice basis $\mathcal{B}(\lambda)$ of monomials for R_n/I_λ. Our exposition owes much to the presentation in the first author's work [24, Definition 2.3.2].

Definition 7 (Garsia-Procesi [17, Section 1]). If $\lambda = (1)$ is the unique partition of 1 then $\mathcal{B}(\lambda) = \{1\}$. If $n > 1$ and λ is a partition of n with k parts then:

- Number the rightmost box in the i^{th} row of λ with i for each $i \in \{1, \ldots, k\}$.
- For each $i \in \{1, \ldots, k\}$ construct the partition λ_i of $n-1$ by erasing the box labeled i and rearranging rows if needed to obtain a Young diagram once again.
- Recursively define $\mathcal{B}(\lambda)$ as

$$
\mathcal{B}(\lambda) = \bigcup_{i=1}^{k} x_n^{i-1} \, \mathcal{B}(\lambda_i).
$$

From the *GP-algorithm* above we construct the *GP-tree* as follows. Let λ sit alone at Level n in a rooted tree directed down. Since λ has k parts, create k edges labeled $x_n^0, x_n^1, \ldots, x_n^{k-1}$ left-to-right to the k subdiagrams λ_i for each $i \in \{1, \ldots, k\}$ in Level $n - 1$. For each of these subdiagrams and their descendants, recursively repeat this process. The process ends at Level 1, whose diagrams all contain one single box. Multiplying the edge labels on any downward path gives a unique *GP-monomial* in the *GP-basis*.

Example 7. We illustrate the first two steps of the GP-algorithm on the partition of 5 given by $\lambda = (2, 2, 1)$. Number the far-right boxes and branch down from Level 5 to Level 4 of the recursion as follows:
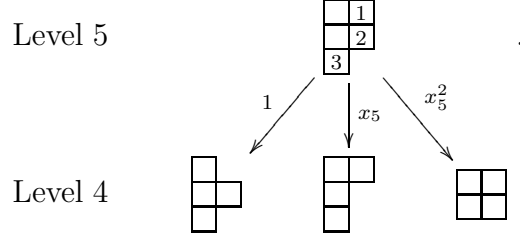
After rearranging rows to obtain a Young diagram, we begin the recursion again on each of the three partitions to produce Level 3 of the tree.
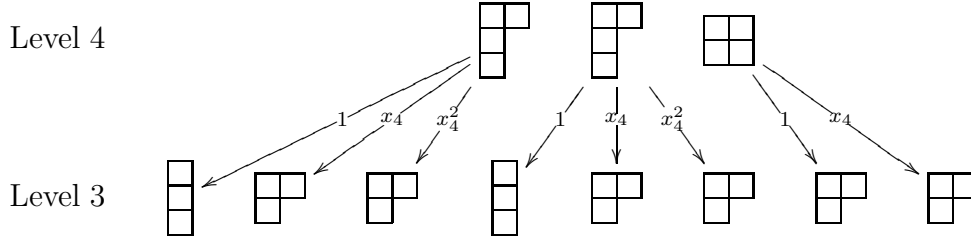
By Level 1, there will be 30 diagrams, each equal to the partition (1). Recovering the Garsia-Procesi basis from this tree is equivalent to multiplying the edge labels of the 30 paths. The reader can verify that we obtain the following basis $\mathcal{B}(\lambda)$:

degree	#	monomials in $\mathcal{B}(\lambda)$
0	1	1
1	4	x_i for $2 \leq i \leq 5$
2	9	x_i^2 for $3 \leq i \leq 5$ and $x_i x_j$ for $2 \leq i < j \leq 5$
3	11	$x_i x_j x_5$ for $2 \leq i < j \leq 4$, $x_i x_j^2$ for $2 \leq i < j \leq 5$, and $x_i^2 x_5$ for $3 \leq i \leq 4$
4	5	$x_i x_j^2 x_5$ for $2 \leq i < j \leq 4$ and $x_i x_4 x_5^2$ for $2 \leq i \leq 3$

Garsia-Procesi bases have several nice containment properties that we use when analyzing the FI- and co-FI-structure of Springer representations. The first describes the relationship between dominance order and the Garsia-Procesi bases.

Proposition 3 (Garsia-Procesi [17, Proposition 4.1]). *Let λ and λ' be partitions of n and suppose that $\lambda \trianglelefteq \lambda'$ in dominance order, namely that we have $\lambda_1 + \lambda_2 + \cdots + \lambda_i \leq \lambda_1' + \lambda_2' + \cdots + \lambda_i'$ for all $i \geq 1$. Then $\mathcal{B}(\lambda') \subseteq \mathcal{B}(\lambda)$.*

The second describes the relationship between containment of Young diagrams and the Garsia-Procesi bases.

Lemma 2. *If $\lambda \subseteq \lambda'$ then $\mathcal{B}(\lambda) \subseteq \mathcal{B}(\lambda')$.*

Proof. We prove the claim assuming that λ' has exactly one more box than λ. Repeating the argument gives the desired result.

Consider the subdiagrams $\lambda_1, \lambda_2, \ldots, \lambda_k$ obtained from λ' in the recursive definition of the Garsia-Procesi algorithm. The Garsia-Procesi algorithm says

$$\mathcal{B}(\lambda') = \bigcup_{i=1}^{k} x_n^{i-1} \, \mathcal{B}(\lambda_i) = \mathcal{B}(\lambda_1) \cup \bigcup_{i=2}^{k} x_n^{i-1} \, \mathcal{B}(\lambda_i)$$

so $\mathcal{B}(\lambda') \supseteq \mathcal{B}(\lambda_1)$. By construction λ_i is obtained from λ' by removing a box that is above and possibly to the right of the box removed for λ_{i+1} for each $i \in \{1, \ldots, k-1\}$. In particular $\lambda_i \trianglelefteq \lambda_{i+1}$ in dominance order. Proposition 3 implies that $\mathcal{B}(\lambda_{i+1}) \subseteq \mathcal{B}(\lambda_i)$ for each i and so $\mathcal{B}(\lambda') \supseteq \mathcal{B}(\lambda_i)$ for each i. Since λ is obtained from λ' by removing a single box, we know $\lambda = \lambda_i$ for some i. The claim follows. $\qquad\square$

4. The Springer representations with FI-module structure

Recall that the sequence of polynomial rings carries an FI-module structure and that both the Tanisaki ideal and the quotients R_n/I_λ carry an S_n-action. The question in this section is: do these fit together to give an FI-module structure on Springer representations? Lemma 2 suggests that the answer could be yes, since it proved that if $\lambda \subseteq \lambda'$ then the Garsia-Procesi basis for R_n/I_λ is contained in the Garsia-Procesi basis for $R_{n'}/I_{\lambda'}$.

We prove that this is misleading: the inclusion $R_n/I_\lambda \hookrightarrow R_{n'}/I_{\lambda'}$ in no way preserves the S_n action (and is not what Church and Farb call a consistent sequence [9, pg.6]). In particular we prove that there is *no* FI-module structure on sequences of Springer representations, except for the trivial representation. Church, Ellenberg, and Farb observed that the sequence of ideals of symmetric functions with no constant term is not an FI-ideal (see Example 3); in our language, they study the special case of the sequence of Tanisaki ideals $\{I_{\lambda_1}, I_{\lambda_2}, \ldots\}$ where each λ_i is a column with i boxes.

Theorem 2. *For each positive integer n, let λ_n denote a Young diagram with n boxes. The only sequence of Young diagrams $\lambda_1 \subseteq \lambda_2 \subseteq \lambda_3 \subseteq \cdots$ for which the sequence of Tanisaki ideals $\{I_{\lambda_1}, I_{\lambda_2}, I_{\lambda_3}, \ldots\}$ forms an FI-ideal is* $\square \to \square\square \to \square\square\square \to \cdots$.

Proof. Suppose that λ is a partition of $n - 1$ and λ' is a partition of n with $\lambda \subseteq \lambda'$. We prove that if $f : \{1, \ldots, n-1\} \to \{1, \ldots, n\}$ is an injection for which $f_*(I_\lambda) \subseteq I_{\lambda'}$ then $I_{\lambda'} = \langle x_1, x_2, \ldots, x_n \rangle$. By definition of FI-modules this suffices to prove our claim.

We first show that $I_{\lambda'}$ contains one of the variables $x_1, \ldots, x_n$. Construct the BFR-fillings of both λ and λ' and compare the boxes labeled 1. A Young diagram with at least two rows has label 1 in the bottom box of the first column. A Young diagram with only one row has label 1 in the rightmost box of its row. If either λ or λ' has just one row then the generating set for its Tanisaki ideal contains $e_1(x_1)$. Hence $I_{\lambda'} \supseteq I_\lambda$ contains the variable x_1. Otherwise the top-left boxes of the BFR-fillings for λ and λ' have the labels $n - 1$ and n respectively, so $\mathcal{G}(\lambda)$ contains $e_1(x_1, \ldots, x_{n-1})$ and $\mathcal{G}(\lambda')$ contains $e_1(x_1, \ldots, x_n)$. In this case

$$e_1(x_1, \ldots, x_n) - f_*(e_1(x_1, \ldots, x_{n-1})) = x_j$$

lies in $I_{\lambda'}$ for the unique $j \in \{1, 2, \ldots, n\} - \mathrm{Im}(f)$.

Thus $I_{\lambda'}$ contains at least one of the variables $x_1, \ldots, x_n$. The Tanisaki ideal is invariant under the action of S_n that permutes the variables $x_1, \ldots, x_n$ so in fact $I_{\lambda'} = \langle x_1, \ldots, x_n \rangle$.

We conclude that
$$R_n/I_{\lambda'} = R_n/\langle x_1, \ldots, x_n \rangle \cong \mathbb{C}$$
is trivial. Since λ' gives the trivial representation, it consists of a single row, as desired. $\square$

5. The Springer representation with the co-FI-module structure

In this section we prove that Springer representations admit a co-FI-module structure. Recall from Section 2 that the sequence of polynomial rings forms a co-FI-module. This section shows that the natural restriction maps from $k[x_1, \ldots, x_n]$ to $k[x_1, \ldots, x_{n-1}]$ are defined on Tanisaki ideals, too. Our proof mimics a similar proof for the cohomology of flag varieties in the arXiv version of a paper by Church, Ellenberg, and Farb [7, Theorem 3.4]. A surprising feature of our result is that this co-FI-module structure exists for Springer representations corresponding to *every possible* sequence $\lambda_1 \subseteq \lambda_2 \subseteq \lambda_3 \subseteq \cdots$ of Young diagrams.

The main features of the proof were outlined in previous sections, especially Section 2, which collected steps that reduce the proof that a sequence of Springer representations form a co-FI-module to proving that particular inclusions $\iota_n^* : R_n \to R_{n-1}$ preserve Tanisaki ideals. We prove the main theorem first and then prove the key lemma.

Theorem 3. *Suppose $\lambda_1 \subseteq \lambda_2 \subseteq \cdots$ is a sequence of Young diagrams for which λ_n has n boxes for each n. Then the sequence $\{R_n/I_{\lambda_n}\}$ forms a graded co-FI-module of finite type.*

Proof. Definition 3 states that to show $\{R_n/I_{\lambda_n}\}$ is a graded co-FI-module of finite type, we must show 1) it is a graded co-FI-module and 2) its dual is a graded FI-module of finite type.

Assuming a field k of characteristic zero, Proposition 1 states that if R is a graded co-FI-module and I is a co-FI-ideal then R/I is a graded co-FI-module. The sequence of polynomial rings is a graded co-FI-algebra under the co-FI-algebra structure that sends the map $f : \{1, \ldots, m\} \to \{1, \ldots, n\}$ to the map with $f^*(x_i)$ equal to $x_{f^{-1}(i)}$ if $i \leq m$ and zero otherwise (see Example 4). Corollary 1 proves that if $\{\lambda_1 \subseteq \lambda_2 \subseteq \cdots\}$ is a sequence of Young diagrams for which $\iota_n^*(I_{\lambda_n}) \subseteq I_{\lambda_{n-1}}$ for every $n \geq 2$ then the sequence of Tanisaki ideals $\{I_{\lambda_1}, I_{\lambda_2}, \ldots\}$ forms a co-FI-ideal. Lemma 5 proves that $\iota_n^*(I_{\lambda_n}) \subseteq I_{\lambda_{n-1}}$ for every $n \geq 2$ and for every sequence of Young diagrams $\{\lambda_1 \subseteq \lambda_2 \subseteq \cdots\}$. Thus $\{R_n/I_{\lambda_n}\}$ is a graded co-FI module.

It is known that the dual $(R_n)^\vee$ is a graded FI-module of finite type. The proof uses the facts that 0) the variables $x_1, \ldots, x_n \in R_n$ form a basis for the dual to k^n, 1) the j^{th} graded part $(R_n)_j^\vee$ is isomorphic to the j^{th} symmetric power of k^n, and 2) the k-modules $k, k^2, k^3, \ldots, k^n, \ldots$ are the parts of a natural FI-module $M(1)$ whose algebraic properties are well understood; for details, see [8, proof of Theorem F].

Proposition 1 then implies that $\{(R_n/I_{\lambda_n})^\vee\}$ is a graded FI-module of finite type, proving the claim. $\square$

5.1. **Sequences of Tanisaki ideals form a co-FI-ideal.** We now analyze the BFR-generators of the Tanisaki ideals to show that $\iota_n^*(I_{\lambda'}) \subseteq I_\lambda$ when $\lambda \subseteq \lambda'$. In an interesting twist, this co-FI-module structure exists for any path in the poset of Young diagrams. Our proof that these sequences of Tanisaki ideals $\{I_{\lambda_i}\}$ form a co-FI-ideal occurs over the course of the following lemmas. Each result actually analyzes a subset of homogeneous polynomials of the same degree inside I_λ so our proofs preserve grading (though we don't use this fact).

We note the following well-known relation (e.g., [26, pg. 21], [14, Equation (3.1)]).

Proposition 4. *Let $S \subseteq \{x_1, \ldots, x_n\}$. If $x_n \in S$ then $e_i(S)$ decomposes as*

$$e_i(S) = x_n \cdot e_{i-1}(S - \{x_n\}) + e_i(S - \{x_n\}).$$

This immediately implies the following. (Recall that i is the degree of each function in the set $E_{i,j}^n$ while j is the cardinality $|S|$ of the variable subset $S \subseteq \{x_1, x_2, \ldots, x_n\}$.)

Lemma 3.

$$\iota_n^* \left(E_{i,j+1}^n \right) = E_{i,j+1}^{n-1} \cup E_{i,j}^{n-1}$$

where $E_{i,k}^{n-1}$ is empty if $k > n - 1$ or if $i > k$ or if $n - 1 < 1$.

The next lemma is our main tool: if an ideal contains $E_{i,j}^n$ then it contains $E_{i,j+1}^n$ as well.

Lemma 4. *The $\mathbb{Q}$-linear span of $E_{i,j}^n$ contains $E_{i,j+1}^n$.*

Proof. Let $S' \subseteq \{x_1, \ldots, x_n\}$ be an arbitrary subset of cardinality $j + 1$. We construct $e_i(S')$ explicitly in the $\mathbb{Q}$-linear span of $E_{i,j}^n$. Consider the polynomial

$$p = \sum_{\substack{S \text{ such that} \\ S \subseteq S' \text{ and } |S| = j}} e_i(S)$$

which is in $\langle E_{i,j}^n \rangle_{\mathbb{Q}}$ by definition. Let T be a subset of S' of cardinality i and consider the coefficient of the monomial $\prod_{x_k \in T} x_k$ in p. Whenever $S \subseteq S'$ contains T the monomial $\prod_{x_k \in T} x_k$ appears with coefficient 1 in $e_i(S)$. There are $j + 1$ subsets of S' with cardinality j and all but i of them contain T. Hence the coefficient of $\prod_{x_k \in T} x_k$ in p is exactly $j + 1 - i$. Thus the polynomial $\frac{1}{j+1-i} p = e_i(S')$ as desired. $\square$

Example 8. Consider the sets $E_{2,3}^5$ and $E_{2,4}^5$ and let $S' = \{x_1, x_2, x_4, x_5\}$. The polynomial p from Lemma 4 is

$$\begin{aligned} p = &(x_1 x_2 + x_1 x_4 + x_2 x_4) + (x_1 x_2 + x_1 x_5 + x_2 x_5) \\ &+ (x_1 x_4 + x_1 x_5 + x_4 x_5) + (x_2 x_4 + x_2 x_5 + x_4 x_5) \end{aligned}$$

which simplifies as desired to

$$p = 2(x_1 x_2 + x_1 x_4 + x_1 x_5 + x_2 x_4 + x_2 x_5 + x_4 x_5) = (4 - 2) e_2(S').$$

Applying Lemma 4 repeatedly gives the following.

Corollary 2. *If an ideal $I \subseteq \mathbb{Q}[x_1, \ldots, x_n]$ contains the subset $E_{i,j}^n$ then it contains $E_{i,j+k}^n$ for all $k \in \{1, 2, \ldots, n - j\}$.*

We now use these combinatorial properties of symmetric functions together with the BFR-generators of the Tanisaki ideal to prove the main lemma of this section

Lemma 5. *Suppose that λ is a Young diagram with $n - 1$ boxes. If λ' is obtained from λ by adding one box then $\iota_n^*(I_{\lambda'}) \subseteq I_\lambda$.*

Proof. Below we give schematics for the relative configurations of λ, λ', and the deleted box (shown in light grey).
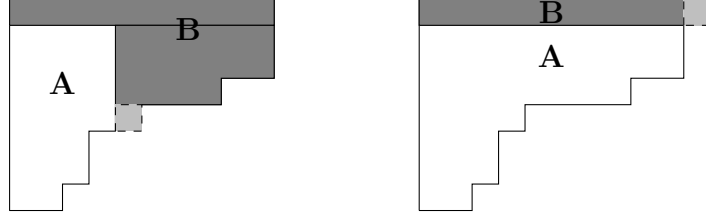
We will compare the BFR-generators for $I_{\lambda'}$ and I_λ for region A, region B, and the two different cases of grey boxes.

Each box in region A corresponds to BFR-generators $E_{i,j}^{n-1}$ in I_λ and $E_{i,j+1}^n$ in $I_{\lambda'}$. In each of these cases the image

$$\iota_n^*(E_{i,j+1}^n) = E_{i,j+1}^{n-1} \cup E_{i,j}^{n-1}$$

by Lemma 3. By construction I_λ contains $E_{i,j}^{n-1}$ so by Corollary 2 we know I_λ contains $E_{i,j+1}^{n-1}$ as well. It follows that I_λ contains $\iota_n^*(E_{i,j+1}^n)$ for every box in region A.

If the box in $\lambda' - \lambda$ is not on the first row, then it is labeled $E_{i,j+1}^n$ in $I_{\lambda'}$ and the box above it is labeled $E_{i,j}^{n-1}$ in I_λ. This is the case of the previous paragraph.

If the box in $\lambda' - \lambda$ is on the first row, then $I_{\lambda'}$ contains $E_{i,i}^n$ while I_λ contains $E_{i,i}^{n-1}$. Lemma 3 shows that $\iota_n^*(E_{i,i}^n) = E_{i,i}^{n-1} \cup E_{i,i-1}^{n-1}$. However $E_{i,i-1}^{n-1}$ is empty by definition. So I_λ contains $\iota^*(E_{i,i}^n)$ in this case too.

Finally, each box in region B corresponds to BFR-generators $E_{i,j}^{n-1}$ in I_λ and $E_{i+1,j+1}^n$ in $I_{\lambda'}$. As above we know

$$\iota_n^*(E_{i+1,j+1}^n) = E_{i+1,j+1}^{n-1} \cup E_{i+1,j}^{n-1}$$

by Lemma 3. The box labeled $i + 1$ in λ is either above but in the same column as in λ' or in a column to the right of $i + 1$ in λ'. In either case $E_{i+1,j-k}^{n-1}$ is a BFR-generator for λ for some nonnegative integer k. Corollary 2 implies that both $E_{i+1,j+1}^{n-1}$ and $E_{i+1,j}^{n-1}$ are in I_λ. Hence $\iota_n^*(E_{i+1,j+1}^n)$ is in I_λ for every box in region B, proving the claim.   □

6. Representation stability of Springer representations

We want to say that a sequence of representations $\{V_n\}$ is representation stable if, when decomposed into irreducible representations, the sequence of multiplicities of each irreducible representation eventually becomes constant. This doesn't quite make sense because the irreducible representations of the symmetric group S_n depend on n, and in fact correspond to the partitions of n. The next definition describes a particular family of irreducible representations whose multiplicities we use to define representation stability.

Definition 8 (Irreducible S_n-representation $V(\mu)_n$ [9, Section 2.1]). Let $\mu = (\mu_1, \ldots, \mu_l)$ be a partition of k. For any $n \geq k + \mu_1$ define the *padded partition* to be

$$\mu[n] := (n - k, \mu_1, \ldots, \mu_l).$$

Define $V(\mu)_n$ to be the irreducible S_n-representation

$$V(\mu)_n := V_{\mu[n]}.$$

Note that every partition of n can be written as $\mu[n]$ for a unique partition μ and hence every irreducible S_n-representation is of the form $V(\mu)_n$ for a unique partition μ.

We can now define representation stability precisely.

Definition 9. A sequence of S_n-representations $V_1 \xrightarrow{f_1} V_2 \xrightarrow{f_2} V_3 \xrightarrow{f_3} \cdots$ is *representation stable* if

(1) the linear maps $f_n : V_n \to V_{n+1}$ are S_n-equivariant, in the sense that for each $w \in S_n$ we have $f_n \circ w = w(n+1) \circ f_n$ where $w(n+1) \in S_{n+1}$ is the permutation that sends $n+1 \mapsto n+1$ while otherwise acting as w;

(2) the maps $f_n : V_n \to V_{n+1}$ are injective;

(3) the span of the S_{n+1}-orbit of the image $f_n(V_n)$ is all of V_{n+1}; and

(4) if V_n decomposes into irreducible representations as

$$V_n = \bigoplus_\mu c_{\mu,n} V(\mu)_n$$

then there exists N so that the multiplicities $c_{\mu,n}$ are independent of n for all $n \geq N$.

If N is independent of μ then $V_1 \xrightarrow{f_1} V_2 \xrightarrow{f_2} V_3 \xrightarrow{f_3} \cdots$ is *uniformly* representation stable.

FI-module structure corresponds to a sequence of representations being uniformly representation stable. Moreover the stability of the multiplicities $c_{\mu,n}$ corresponds to a kind of stabilization of the characters of the representations, and thus the dimensions of the representations.

Since $\{R_n/I_{\lambda_n}\}$ forms a co-FI-module, we will actually show that the dual $\{(R_n/I_{\lambda_n})^\vee\}$ is representation stable. The term *Springer representation* is applied interchangeably to these two dual representations (as described in the Introduction–see Remark 1). Moreover, the key implications of representation stability (including dimensions and characters) apply to $\{R_n/I_{\lambda_n}\}$ by virtue of applying to its dual.

More precisely we have the following.

Corollary 3. *Suppose $\lambda_1 \subseteq \lambda_2 \subseteq \lambda_3 \subseteq \cdots$ is a sequence of Young diagrams for which λ_n has n boxes for each n. Then the following are all true:*

(1) *Each graded part of the sequence $\{(R_n/I_{\lambda_n})^\vee\}$ is uniformly representation stable.*

(2) *For each k and each n let $\chi_{k,n}$ be the character in the k^{th} graded part of R_n/I_{λ_n}. The sequence $(\chi_{k,1}, \chi_{k,2}, \chi_{k,3}, ...)$ is eventually polynomial in the sense of [6].*

(3) *For each k and n let $d_{k,n}$ be the dimension of the k^{th} graded piece of R_n/I_{λ_n} as a complex vector space. Then the sequence $(d_{k,1}, d_{k,2}, d_{k,3}, \ldots)$ is eventually polynomial, in the sense that there is an integer s_k and polynomial $p_k(n)$ in n so that for all $n \geq s_k$ the dimensions $d_{k,n} = p_k(n)$.*

Proof. This follows immediately from corresponding results of Church, Ellenberg, and Farb. Part (1) follows from Theorem 3 together with the definition of representation stability (or, e.g., [6, Theorem 1.13]). Part (2) follows for $(R_n/I_{\lambda_n})^\vee$ from [6, Theorem 1.5] and for the dual representation R_n/I_{λ_n} because the characters of complex representations of S_n respect the operation of taking duals. Part (3) is an easy corollary of Part (2) for R_n/I_{λ_n} by [6, Theorem 1.5] and for its dual because they have the same dimension as complex representations. $\square$

7. Combinatorial consequences

The results of Corollary 3 open up new combinatorial questions about Springer fibers. Representation stability guarantees that for any sequence of Young diagrams $\{\lambda_n\}$, the k^{th} degree of the cohomology of the corresponding Springer fibers stabilizes as a polynomial $p_k(n)$. But what is this polynomial? For instance, given the sequence $\lambda_1 \subseteq \lambda_2 \subseteq \lambda_3 \subseteq \cdots$ of Young diagrams:

- Is there an explicit formula for the dimensions $p_k(n)$ of the k^{th} graded part of R_n/I_{λ_n}? What if the sequence contains a particular family of Young diagrams (e.g. hooks, two-row, two-column, etc.)?
- What is the minimal integer s_k at which the dimension of R_n/I_{λ_n} becomes polynomial? Given k and a sequence of Young diagrams, can we find some s_k (not necessarily minimal) after which the dimension of R_n/I_{λ_n} is polynomial?
- Can we show the Springer representations have polynomial dimension via the monomial bases $\mathcal{B}(\lambda_n)$ from Section 3?

To the best of our knowledge, these are entirely new questions about Springer representations; they arise only because of the consequences of representation stability. Moreover, our preliminary answers to these questions suggest additional combinatorial structures that may undergird Springer representations. We give details and some concrete results in this section, followed in the next by (more) open questions.

7.1. **Minimal monomials.**
Our first results suggest that the exponents that appear in the monomials within each $\mathcal{B}(\lambda)$ are more rigid than previously thought.

Definition 10 (Monomial type). Let x^α be a monomial with exponent vector $(\alpha_1, \ldots, \alpha_n)$. In other words, the exponent of the variable x_i in x^α is α_i. Let $(\alpha_{i_1}, \ldots, \alpha_{i_r})$ be the nonzero exponents in the order in which they appear in x^α. This r-tuple is called the *monomial type* of x^α.

For instance both $x_2 x_3^2 x_4$ and $x_3 x_9^2 x_{11}$ have monomial type $(1, 2, 1)$.

The following result says that if x_I^α is a monomial of a fixed monomial type in the GP-basis $\mathcal{B}(\lambda)$ then increasing the indices of the variables lexicographically while preserving the monomial type produces another monomial in the GP-basis $\mathcal{B}(\lambda)$.

Theorem 4. *Fix a composition $\alpha = (\alpha_1, \ldots, \alpha_k)$. For each ordered sequence $I = (i_1, \ldots, i_k)$ let x_I^α denote the monomial $\prod_{j=1}^k x_{i_j}^{\alpha_j}$. If $x_I^\alpha \in \mathcal{B}(\lambda)$ and I' satisfies*

$$I \leq I' \leq (n-k+1, \ldots, n-1, n)$$

in lexicographic order then $x_{I'}^\alpha \in \mathcal{B}(\lambda)$.

Proof. First note that we may simply consider the case when I and I' differ by exactly one in exactly one entry. Indeed, for each j with $1 \leq j \leq k+1$, define the sequences $I_j = (i_1, i_2, \ldots, i_{j-1}, i'_j, i'_{j+1}, \ldots, i'_k)$. Note that $I = I_{k+1}$ and $I' = I_1$. Thus it suffices to prove that if $x_{I_{j+1}}^\alpha \in \mathcal{B}(\lambda)$ then $x_{I_j}^\alpha \in \mathcal{B}(\lambda)$. Furthermore we may assume that $i'_j = i_j + 1$ since repeating the argument would successively increment the index and imply the result for more general i'_j. Thus we prove the claim for I_j and I_{j+1} assuming that $i'_j = i_j + 1$.

The GP-algorithm proceeds the same for both monomials $x_{I_j}^\alpha$ and $x_{I_{j+1}}^\alpha$ in the variables $x_n, x_{n-1}, \ldots, x_{i_j+2}$. Let $\lambda^{(j+2)}$ be the Young diagram left after those steps. Implementing the GP-algorithm for $x_{I_{j+1}}^\alpha$ removes a box from the first row and then from the $\alpha_j + 1^{\text{th}}$ row of

$\lambda^{(j+2)}$, whereas for $x_{I_j}^{\alpha}$ it removes a box from the $\alpha_j + 1^{\text{th}}$ row and then from the first row. Call the resulting Young diagrams $\lambda^{(j+1)}$ and $\lambda^{(j)}$ respectively.

We now show that $\lambda^{(j+1)} \trianglerighteq \lambda^{(j)}$ in dominance order. If the first row of $\lambda^{(j+2)}$ is strictly larger than the $\alpha_j + 1^{\text{th}}$ row of $\lambda^{(j+2)}$ then the result is clear, since then $\lambda^{(j+1)} = \lambda^{(j)}$. If not, then removing a box from the first row of $\lambda^{(j+2)}$ leaves a Young diagram whose shape is constrained: at least the first α_j rows have the same length and there is at least one row immediately after of length one less. Thus $\lambda^{(j+1)} \trianglerighteq \lambda^{(j)}$ with equality in the case that exactly the first $\alpha_j + 1$ rows of $\lambda^{(j+2)}$ are the same length, and any subsequent rows are smaller.

By Proposition 3, we have $\mathcal{B}(\lambda^{(j+1)}) \subseteq \mathcal{B}(\lambda^{(j)})$. Since $x_{i_1}^{\alpha_1} x_{i_2}^{\alpha_2} \cdots x_{i_{j-1}}^{\alpha_{j-1}}$ is a GP-monomial in $\mathcal{B}(\lambda^{(j+1)})$ it is also a GP-monomial in $\mathcal{B}(\lambda^{(j)})$ so the monomial $x_{I_j}^{\alpha}$ can be obtained from the GP-algorithm starting with the shape $\lambda^{(j+2)}$. Thus $x_{I_j}^{\alpha}$ is in the GP-basis for $\mathcal{B}(\lambda)$ as desired, and the claim follows. $\qquad \square$

In fact we conjecture that there is a unique minimal monomial of each monomial type, in the following lexicographic sense.

Definition 11 (Minimal monomials)**.** Fix a monomial type α and a partition λ of n. The monomial $x_I^{\alpha} \in \mathcal{B}(\lambda)$ is a *minimal monomial* of type α if for any other $x_{I'}^{\alpha} \in \mathcal{B}(\lambda)$ of the same monomial type, the index sets satisfy $I \leq I'$ in lexicographic order.

The next result proves that when λ has exactly two rows then there is a unique minimal monomial of each monomial type in $\mathcal{B}(\lambda)$. The previous result then says that this minimal monomial produces all other monomials of that fixed monomial type within $\mathcal{B}(\lambda)$.

Theorem 5. *Fix a two-row Young diagram $\lambda = (n - k, k)$. Then the basis $\mathcal{B}(\lambda)$ of R_n/I_λ has the unique minimal monomial set $\left\{ \prod_{j=1}^{i} x_{2j} \right\}_{i=1}^{k}$.*

Proof. Since λ has two rows, the maximum exponent that appears in a monomial for λ is 1. Since λ has k boxes in the second row, the maximum degree in a monomial for λ is k. So every monomial type that appears in $\mathcal{B}(\lambda)$ is the partition given by j copies of 1, where j is any integer with $0 \leq j \leq k$.

We claim that the minimal monomial of degree j is $x_2 x_4 \cdots x_{2j}$. To see this, let $x_I \in \mathcal{B}(\lambda)$ have degree j and index set $I = (i_1, \ldots, i_j)$. It suffices to show that $(2, 4, \ldots, 2j) \leq I$ in lexicographic order, namely that $i_r \geq 2r$ for each $r \in \{1, \ldots, j\}$. Recall that in the GP-algorithm, exactly one rightmost box in a row is removed when going down each level of the GP-tree. Going from Level i to Level $i - 1$, the box removed is in the first row of the Level i diagram if $i \notin \{i_1, \ldots, i_j\}$ and the second row otherwise. After a box is removed from the first row, the rows of the resulting Level $i - 1$ diagram are switched if necessary to maintain the shape of a Young diagram.

Now suppose $r \in \{1, \ldots, j\}$. The path in the GP-tree from Level i_r to Level $i_r - 1$ removes a box from the second row of the Level i_r diagram $\lambda^{(i_r)}$. The diagram $\lambda^{(i_r)}$ must have at least r boxes in the second row since the $r - 1$ factors $x_{i_{r-1}}, x_{i_{r-2}}, \ldots, x_{i_1}$ remaining in x_I correspond to $r - 1$ more second-row boxes removed in the remaining $i_r - 2$ levels of the GP-tree below Level $i_r - 1$. Since all Level i_r diagrams have i_r boxes and $\lambda^{(i_r)}$ has at least r boxes in its second row, we conclude $i_r \geq 2r$ as desired. $\qquad \square$

We conjecture that for each $\mathcal{B}(\lambda)$ and each monomial type that appears in $\mathcal{B}(\lambda)$, there is a unique minimal monomial of that monomial type (see Section 8.1).

7.2. **Polynomial dimension when the number of rows exceed a certain minimum.**
We can explicitly compute the polynomial dimensions described in Corollary 3 in some cases,
including when the Young diagram has "enough" rows. In this case, the dimensions in low
degree coincide with the corresponding dimensions for the diagonal coinvariant algebra, as
described in the next theorem. Kim's version of stability coincides with ours in this case [21,
Remark in Section 4].

Theorem 6. *Let* $\dim_i(R_n/I_{\lambda_n})$ *denote the dimension of the* i^{th}-*degree part of* R_n/I_{λ_n}. *Let*
$\{\lambda_n\}$ *be a nested sequence of Young diagrams such that* $|\lambda_n| = n$ *for all* n. *If for some* N *the*
diagram λ_N *contains at least* $k+1$ *rows, then* $\dim_i\left(R_n\big/I_{\lambda_n}\right)$ *agrees with the dimension of*
the i^{th}-*degree part of the diagonal coinvariant algebra* $R_n\big/\langle e_j\rangle_{j=1}^n$ *for all* $i \leq k$ *and* $n \geq N$.

  In particular for all $i \leq k$ *and* $n \geq N$ *we have* $\dim_i\left(R_n\big/I_{\lambda_n}\right) = p_i(n)$ *where* $p_i(n)$ *is*
the polynomial that gives the number of permutations of the set $\{1, 2, \ldots, n\}$ *with exactly* i
inversions.

Proof. Let $\text{col}_n = (1, \ldots, 1)$ denote the column partition with exactly n parts. The diagonal
coinvariant algebra is R_n/I_{col_n}. There is a unique monomial $x_2 x_3^2 x_4^3 \cdots x_n^{n-1} \in \mathcal{B}(\text{col}_n)$ of
maximal degree; it is obtained by choosing the far-right edge at each level in the GP-tree.
Garsia-Procesi proved that their basis elements form a lower-order ideal with respect to
division [17, Proposition 4.2]. It follows that the other elements of $\mathcal{B}(\text{col}_n)$ are precisely the
divisors of this maximal monomial, so

$$\mathcal{B}(\text{col}_n) = \{x_2^{\alpha_2} x_3^{\alpha_3} x_4^{\alpha_4} \cdots x_n^{\alpha_n} \mid 0 \leq \alpha_j \leq j-1\}. \tag{1}$$

  Now suppose that λ_n has at least $k+1$ rows. Then $\lambda_n \supseteq \text{col}_{k+1}$ and so $\mathcal{B}(\lambda_n) \supseteq \mathcal{B}(\text{col}_{k+1})$
by Lemma 2. By Theorem 4 we can lexicographically increase the subscripts of each mono-
mial in $\mathcal{B}(\text{col}_{k+1})$ to produce monomials x_I^α for each k-element subset $I \subseteq \{x_2, \ldots, x_n\}$ and
exponents α which satisfy $0 \leq \alpha_j \leq j-1$ for each $j = 1, 2, \ldots, k+1$. This proves that
$\mathcal{B}(\lambda_n)$ also contains all of those possible elements of degrees $0, 1, \ldots, k$ in the variable set
$\{x_2, \ldots, x_n\}$. Trivially, these are the same as the monomials of degrees $0, 1, \ldots, k$ in the basis
for $\mathcal{B}(\text{col}_n)$ above, so in these degrees the basis elements for R_n/I_{λ_n} and the diagonal coin-
variant algebra R_n/I_{col_n} coincide. The claim about $p_i(n)$ follows from the similar statement
for the i^{th} graded part of R_n/I_{col_n}.   □

  The polynomial function $p_i(n)$ can be found explicitly as an alternating sum of certain
combinations. Example 10 gives closed formulas for $p_i(n)$ for the first few values of i.

Remark 3. Since the diagonal coinvariant algebra is the cohomology ring of the full flag
variety, this result implies that the Garsia-Procesi bases for the cohomology of the Springer
fibers coincide with the cohomology of the full flag variety in many degrees. More precisely
if $\mathcal{S}_{\lambda_n}$ denotes the Springer fiber corresponding to the nilpotent of Jordan type λ_n then we
have an isomorphism $H^i(\mathcal{S}_{\lambda_n}) \cong H^*(GL_n(\mathbb{C})/B)$ for all $i \leq k$ and $n \geq N$, where k and N
are as given in Theorem 6.

Example 9. Below we give the basis elements for $\mathcal{B}(\text{col}_5)$ in degrees 0 through 3 and the
cardinalities of the k^{th} degree parts for $4 \leq k \leq 10$, using the unimodality of the sequence
of dimensions for $k \geq 6$.

degree	#	monomials in $\mathcal{B}(\mathrm{col}_5)$
0	1	1
1	4	x_i for $2 \leq i \leq 5$
2	9	x_i^2 for $3 \leq i \leq 5$ and $x_i x_j$ for $2 \leq i < j \leq 5$
3	15	$x_i x_j x_k$ for $2 \leq i < j < k \leq 5$, $x_i x_j^2$ for $2 \leq i < j \leq 5$, $x_i^2 x_j$ for $3 \leq i < j \leq 5$, and x_i^3 for $4 \leq i \leq 5$
4	20	
5	22	
$6 \leq k \leq 10$	same cardinality as $(10-k)^{\text{th}}$ degree part	
$k > 10$	0	

Compare this to the basis $\mathcal{B}(\lambda)$ for the partition $\lambda = (2,2,1)$ that we computed in Example 7. The basis elements for $\mathcal{B}(\lambda)$ and $\mathcal{B}(\mathrm{col}_n)$ coincide in degrees 0, 1, and 2, as expected. However in degree 3 the set $\mathcal{B}(\mathrm{col}_5)$ contains four basis elements that are not in $\mathcal{B}(\lambda)$, namely $x_2 x_3 x_4$, $x_3^2 x_4$, x_4^3, and x_5^3.

Example 10. We give the following closed formulas for the dimensions $p_i(n)$ from Theorem 6 when $0 \leq i \leq 6$:

- When $i = 0$ we have $p_0(n) = 1$.
- When $i = 1, 2, 3, 4$ we have

$$p_i(n) = \binom{n-2+i}{i} - \binom{n-3+i}{i-2}. \tag{2}$$

- When $i = 5$ we have

$$p_5(n) = \binom{n+3}{5} - \binom{n+2}{3} + 1.$$

- When $i = 6$ we have

$$p_6(n) = \binom{n+4}{6} - \binom{n+3}{4} + n.$$

To prove these, let $i \in \{1, 2, 3, 4\}$. Our convention is that $0 \in \mathbb{N}$. Define the set

$$\widetilde{E_i} = \left\{ (d_2, \ldots, d_n) \in \mathbb{N}^{n-1} \ \middle| \ \sum_{k=2}^{n} d_k = i \text{ and } d_k \leq i \right\}$$

of nonnegative integer solutions to the equation $d_2 + \cdots + d_n = i$. We know $|\widetilde{E_i}| = \binom{(n-2)+i}{i}$. Define $E_i \subseteq \widetilde{E_i}$ to be the subset satisfying also $d_k \leq k - 1$ for each k so that $p_i(n) = |E_i|$. We now compute $|E_i| = |\widetilde{E_i}| - |E_i^c|$.

For each m, define the set $F_{m,i} \subseteq \widetilde{E_i}$ by the condition that $d_m > m - 1$. Observe that $E_i^c = \bigcup_{m=2}^{n} F_{m,i}$. Moreover the sets $F_{m,i}$ are pairwise disjoint. Indeed, suppose there were $m_1 < m_2$ with an element $(d_2, \ldots, d_n) \in F_{m_1,i} \cap F_{m_2,i}$. Then $d_{m_1} > m_1 - 1 \geq 1$ and $d_{m_2} > m_2 - 1 \geq 2$ and so $d_2 + \cdots + d_n \geq 5$. This contradicts $i \in \{1, 2, 3, 4\}$.

So $|E_i^c| = \sum_{m=2}^{n} |F_{m,i}|$. We now find $|F_{m,i}|$. Since $d_m > m - 1$ we write $d_m = m + d_m'$ for some $d_m' \geq 0$ and then count instead the solutions to $d_2 + \cdots + d_m' + \cdots + d_n = i - m$. Thus

$$|E_i^c| = \sum_{m=2}^{n} |F_{m,i}| = \sum_{m=2}^{n} \binom{(n-2) + (i-m)}{i - m}.$$

As long as $i \leq n$ this sum has the form $|E_i^c| = \cdots + \binom{r_m+2}{2} + \binom{r_m+1}{1} + \binom{r_m}{0}$ for some r_m. Using a combinatorial identity we write $|E_i^c| = \binom{(n-2)+(i-2)+1}{i-2}$. This gives the formula for $p_i(n)$ in Equation (2).

When $i = 5$ and $i = 6$ we can use essentially the previous argument, since only a few pairs of $F_{m,i}$ share elements. When $i = 5$ the sets $F_{2,5}$ and $F_{3,5}$ both contain $(2, 3, 0, 0, 0, \ldots, 0)$. Inclusion-Exclusion gives the desired formula. Similarly if $i = 6$ then $F_{2,6} \cap F_{3,6}$ contains the two elements $(3, 3, 0, 0, \ldots, 0)$ and $(2, 4, 0, 0, \ldots, 0)$ as well as the $n - 3$ elements $(2, 3, d_4, \ldots, d_n)$ where exactly one of the remaining entries d_k is nonzero (and thus is 1). Additionally $F_{2,6} \cap F_{4,6}$ contains the unique element $(2, 0, 4, 0, 0, \ldots, 0)$. Inclusion-Exclusion gives the desired formula.

This process becomes more complicated when i is larger.

Kim gives equivalent but different versions of the formulas in Example 10 [21, Example 3.3].

8. Open questions

This section describes a set of questions that arise from our analysis of representation stability, including how to characterize the polynomials into which the dimensions stabilize, and how to describe the monomials in $\mathcal{B}(\lambda)$.

8.1. **Minimal monomials and monomial types.** In Subsection 7.1 we discussed monomial types and minimal monomials. We pose several questions here.

Question 1. For each Young diagram λ and monomial type α, is there a unique monomial $x_I^\alpha \in \mathcal{B}(\lambda)$ that is minimal with respect to lexicographic ordering, in the sense that if $x_{I'}^\alpha \in \mathcal{B}(\lambda)$ then $I \leq I'$?

All examples that we have computed do in fact have minimal monomials.

There are several related questions about how the set of monomial types that appear in $\mathcal{B}(\lambda)$ are related to λ. The first essentially asks how to predict the monomial types that appear for a given λ, while the second asks which $\mathcal{B}(\lambda)$ contain a given monomial type α.

Question 2. For each Young diagram α, let $\mathcal{A}(\lambda)$ be the set of monomial types in $\mathcal{B}(\lambda)$. Is there a quick algorithm to construct $\mathcal{A}(\lambda)$ for arbitrary λ?

We know there is an algorithm to construct $\mathcal{A}(\lambda)$, namely use the GP-algorithm to find all of $\mathcal{B}(\lambda)$ and then identify which monomial types appear. The previous question is asking for a direct algorithm, or for characterizations of the monomial types that do or do not appear.

Question 3. Fix a monomial type α. For which λ is there a monomial of type α in $\mathcal{B}(\lambda)$?

For instance $\mathcal{B}(\lambda)$ contains a monomial of type (k) if and only if λ has at least $k + 1$ rows.

8.2. **Polynomial dimension for particular families of Young diagrams.** In this section, we return to questions from Subsection 7.2 to ask for concrete formulas for dimensions of the Springer representations, and especially for the polynomials $p_k(n)$ to which they stabilize.

Representation stability relies on sequences of diagrams while the literature on Springer fibers instead uses families of Young diagrams like two-row, two-column, hook, and so on. To study families of Young diagrams, we choose a sequence $\{\lambda_n\}$ that characterizes particular shapes, as Example 11 does for two-column Young diagrams.

Question 4. For what families of Young diagrams $\{\lambda_n\}$ can we find an explicit closed formula for the polynomial dimension $p_k(n)$?

For instance, Kim gave a formula for two-row Young diagrams [20], following ideas due to Fresse [15, Example 4.5].

When we can identify minimal monomials (unique or not), strictly combinatorial techniques might then give information about the polynomial dimension formula. More precisely:

Question 5. Can we find minimal monomials for other shapes, and use them together with the index-incrementation tools in Theorem 4 to describe the dimension polynomials, either partially or completely?

For instance, Theorem 5 gave the minimal monomials for the case of the two-row diagrams $\lambda = (n-k, k)$. In that case, the previous question sketches an alternate proof of Kim's result for two-row Young diagrams.
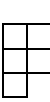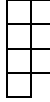
8.3. **Stable shapes.** The central question of this section is to determine the integer s_k after which the polynomial $p_k(n)$ counts the degree-k monomials in $\mathcal{B}(\lambda_n)$. More colloquially, when does the dimension stabilize?

We conjecture that there is a core shape such that as soon as λ_n contains the core shape, the dimension stabilizes. The following describes one possible choice of a core shape, though we do not believe that it is in general minimal.

Definition 12. Let $\lambda_1 \subseteq \lambda_2 \subseteq \cdots$ be a sequence of Young diagrams with $|\lambda_n| = n$ for each n. Denote by τ_{k+1} the staircase partition $(k+1, k, \ldots, 1)$. The *k-stable shape* of $\{\lambda_n\}$ is the partition $(\cup_n \lambda_n) \cap \tau_{k+1}$.

Example 11. Let $\{\lambda_n\}$ be the nested sequence of Young diagrams defined by $\lambda_1 = (1)$, $\lambda_{2n} = (\underbrace{2, 2, \ldots, 2}_{n \text{ times}})$, and $\lambda_{2n+1} = (\underbrace{2, 2, \ldots, 2}_{n \text{ times}}, 1)$. Then $\{\lambda_n\}$ is the following
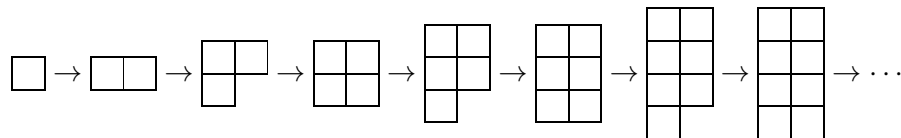
The 1-, 2- and 3-stable shapes are , and , respectively.

Conjecture 1. *Let $\{\lambda_i\}$ be a sequence of Young diagrams such that $\lambda_1 \subseteq \lambda_2 \subseteq \cdots$ with $|\lambda_i| = i$. Let N be the smallest integer for which λ_N contains the k-stable shape of $\{\lambda_i\}$. Then there exists a polynomial $p_k(n)$ that gives the dimension of the k^{th} graded part of the Springer representation R_n / I_{λ_n} for all $n \geq N$.*

If true, this conjecture would provide more tools to explicitly compute $p_k(n)$.

Acknowledgements

References

[1] R. Biagioli, S. Faridi, and M. Rosas: The defining ideals of conjugacy classes of nilpotent matrices and a conjecture of Weyman. *Int. Math. Res. Not. IMRN* (2008)

[2] C. Bibby: Representation stability for the cohomology of arrangements associated to root systems. *J. Algebraic Combin.* **48**, no. 1, 51–75 (2018)

[3] A. Borel: Sur la cohomologie des espaces fibrés principaux et des espaces homogènes de groupes de Lie compacts. *Ann. of Math.* (2) **57**, 115–207 (1953)

[4] W. Borho and R. MacPherson: Partial Resolutions of Nilpotent Varieties. *Astérisque* **101-102**, 23–74, Soc. Math. France, Paris (1983)

[5] T. Church. Personal communication, 2020.

[6] T. Church, J.S. Ellenberg, and B. Farb: FI-modules and stability for representations of symmetric groups. *Duke Math. J.* **164**, no. 9, 1833–1910 (2015)

[7] T. Church, J.S. Ellenberg, and B. Farb: FI-modules: a new approach to stability for S_n-representations. This is the 2012 version of the paper above. Available at http://arxiv.org/pdf/1204.4533v2.pdf.

[8] T. Church, J.S. Ellenberg, B. Farb, and R. Nagpal: FI-modules over Noetherian rings. *Geom. Topol.* **18**, no. 5, 2951–2984 (2014).

[9] T. Church and B. Farb: Representation theory and homological stability. *Adv. Math.* **245**, 250–314 (2013)

[10] C. De Concini and C. Procesi: Symmetric functions, conjugacy classes and the flag variety. *Invent. Math.* **64**(2), 203–219 (1981)

[11] F. De Mari, C. Procesi, and M.A. Shayman: Hessenberg varieties. *Trans. Amer. Math. Soc.* **332**(2), 529–534 (1992)

[12] J. Duque and R. Rolland: Representation stability for the pure cactus group. Available at https://arxiv.org/pdf/1501.02835.

[13] B. Farb: Representation Stability. Available at http://arxiv.org/pdf/1404.4065.pdf.

[14] S. Fomin, S. Gelfand, and A. Postnikov: Quantum Schubert polynomials. *J. Amer. Math. Soc.* **10**(3), 565–596 (1997)

[15] L. Fresse: Betti numbers of Springer fibers in type A. *J. Algebra* **322**, 2566–2579 (2009)

[16] N. Gadish: Representation stability for families of linear subspace arrangements. *Adv. Math.* **322**, 341–377 (2017)

[17] A. Garsia and C. Procesi: On certain graded S_n-modules and the q-Kostka polynomials. *Adv. Math.* **94**, 82–138 (1992)

[18] P. Hersh and V. Reiner: Representation stability for cohomology of configuration spaces in $\mathbb{R}^d$. *Int. Math. Res. Not. IMRN* **5**, 1433–1486 (2017)

[19] R. Hotta.: On Springer's representations. *J. Fac. Sci. Univ. Tokyo Sect. IA Math.* **28** (1981), no. 3, 863–876 (1982)

[20] D. Kim: Euler characteristic of Springer fibers. *Transform. Groups.* **24**, no. 2, 403–428 (2019)

[21] D. Kim: Stability of Springer representations for type A. Available at https://arxiv.org/pdf/1612.02381.

[22] H. Kraft: Conjugacy classes and Weyl group representations. In: Young tableaux and Schur functors in algebra and geometry (Toruń, 1980), *Astérisque* **87-88**, 191–205, Soc. Math. France, Paris (1981)

[23] G. Lusztig: Intersection cohomology complexes on a reductive group. *Invent. Math.* **75**, 205–272 (1984)

[24] A. Mbirika: A Hessenberg generalization of the Garsia-Procesi basis for the cohomology ring of Springer varieties. *Electron. J. Combin.* **17**(1), Research Paper 153 (2010)

[25] P. Patzt: Representation stability for filtrations of Torelli groups. *Math. Ann.* **372**, 257–298 (2018)

[26] C. Procesi: *Lie groups – An approach through invariants and representations.* Universitext. Springer, New York (2007)

[27] R. Rolland: Representation stability for the cohomology of the moduli space M_g^n. *Algebr. Geom. Topol.* **11**, 3011–3041 (2011)

[28] R. Rolland: The cohomology of $\mathcal{M}_{0,n}$ as an FI-module. *Springer INdAM Ser.* **14**, 313–323 (2016)

[29] T.A. Springer: A construction of representations of Weyl groups. *Invent. Math.* **44**(3), 279–293 (1978)

[30] T. Tanisaki: Defining ideals of the closures of the conjugacy classes and representations of the Weyl groups. *Tôhoku Math. J.* (2) **34**(4), 575–585 (1982)

[31] J. Wilson: FI$_W$-modules and stability criteria for representations of classical Weyl groups. *J. Algebra* **420**, 269–332 (2014)

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF WISCONSIN-EAU CLAIRE, EAU CLAIRE, WI, U.S.A.
E-mail address: mbirika@uwec.edu

DEPARTMENT OF MATHEMATICS AND STATISTICS, SMITH COLLEGE, NORTHAMPTON, MA, U.S.A.
E-mail address: jtymoczko@smith.edu