



Annals of Representation Theory

BEN WEBSTER 

Gelfand–Tsetlin modules in the Coulomb context

Volume 1, issue 3 (2024), p. 393–437

<https://doi.org/10.5802/art.14>

Communicated by Andrea Solotar.

© The author, 2024



This article is licensed under the
CREATIVE COMMONS ATTRIBUTION (CC-BY) 4.0 LICENSE.
<http://creativecommons.org/licenses/by/4.0/>



Annals of Representation Theory is published by the
Norwegian University of Science and Technology
and a member of the
Centre Mersenne for Open Scientific Publishing

e-ISSN: 2704-2081



Gelfand–Tsetlin modules in the Coulomb context

Ben Webster 

ABSTRACT. This paper gives a new perspective on the theory of principal Galois orders, as developed by Futorny, Ovsienko, Hartwig, and others. Every principal Galois order can be written as eFe for any idempotent e in an algebra F , which we call a flag Galois order; and in most important cases we can assume that these algebras are Morita equivalent. These algebras have the property that the completed algebra controlling the fiber over a maximal ideal has the same form as a subalgebra in a skew group ring, which gives a new perspective to a number of results about these algebras.

We also discuss how this approach relates to the study of Coulomb branches in the sense of Braverman–Finkelberg–Nakajima, which are particularly beautiful examples of principal Galois orders. These include most of the interesting examples of principal Galois orders, such as $U(\mathfrak{gl}_n)$. In this case, all the objects discussed have a geometric interpretation, which endows the category of Gelfand–Tsetlin modules with a graded lift and allows us to interpret the classes of simple Gelfand–Tsetlin modules in terms of dual canonical bases for the Grothendieck group. In particular, we classify the Gelfand–Tsetlin modules over $U(\mathfrak{gl}_n)$ and relate their characters to a generalization of Leclerc’s shuffle expansion for dual canonical basis vectors.

Finally, as an application, we disprove a conjecture of Mazorchuk, showing that the fiber over a maximal ideal of the Gelfand–Tsetlin subalgebra appearing in a finite-dimensional representation has an infinite-dimensional module in its fiber for $n \geq 6$.

1. INTRODUCTION

Let Λ be a Noetherian commutative ring and $\widehat{W}$ a monoid acting faithfully on Λ ; let $L = \text{Frac}(\Lambda)$ be the fraction field of Λ . Assume that $\widehat{W}$ is the semi-direct product of a finite subgroup W and a submonoid $\mathcal{M}$ and that $\#W$ is invertible in Λ . For simplicity, we assume throughout the introduction that $\mathcal{M}$ has finite stabilizers in its action on $\text{MaxSpec}(\Lambda)$.

Manuscript received 2023-02-28, revised 2024-06-26 and accepted 2024-08-11.

Keywords. Gelfand–Tsetlin module, weight module, Coulomb branch.

2020 *Mathematics Subject Classification.* 14M15, 16G99, 17B10.

The author was supported by NSERC through a Discovery Grant. This research was supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Research, Innovation and Science.

A *principal Galois order* (Definition 2.1) U is a subalgebra of invariants $(L\#\mathcal{M})^W$ of the skew group ring $L\#\mathcal{M}$ equipped with (among other structures) an inclusion of $\Gamma = \Lambda^W$ as a subalgebra (usually called the Gelfand–Tsetlin subalgebra) and a faithful action on Γ .

We call a finitely generated U -module *Gelfand–Tsetlin* if it is locally finite under the action of Γ , and thus decomposes as a direct sum of generalized weight spaces. An important motivating question for a great deal of work in recent years has been the question:

Question. Given a principal Galois order U , classify the simple Gelfand–Tsetlin modules and describe the dimensions of their generalized weight spaces for the different maximal ideals of Γ .

1.1. Generalities on Galois orders. Work of Drozd–Futorny–Ovsienko [9, Th. 18] shows that the simple GT modules containing a maximal ideal $\mathfrak{m}_\gamma$ of Γ in their support (the *fiber* over γ) are controlled by a profinite length algebra $\widehat{U}_\gamma$, which naturally acts on the corresponding generalized weight space for any U -module. The simple discrete modules over $\widehat{U}_\gamma$ are the non-zero γ -generalized weight spaces of the different simple Gelfand–Tsetlin modules. Thus, we can rephrase the above question as how to understand these algebras in specific special cases.

One perspective shift we want to strongly emphasize is that taking invariants for a group action conceals a great deal of structure—we can see this structure more clearly if we instead consider subalgebras F in the skew group ring $L\#\widehat{W}$ with coefficients in L of the semi-direct product $\widehat{W} = W \ltimes \mathcal{M}$, which we call *principal flag orders* (Definition 2.2). These are simply the principal Galois orders containing the smash product $\Lambda\#W$ where we take $W' = \{1\}$ and $\mathcal{M}' = \widehat{W}$.

If we let $e \in \mathbb{Z}[\frac{1}{\#W}][W]$ be the symmetrization idempotent, then for any principal flag order F , the centralizer $U = eFe$ is a principal Galois order for our original data, and every principal Galois order appears this way (Lemma 2.5). One theme we will use throughout the paper is the interplay between a maximal ideal $\mathfrak{m}_\lambda \subset \Lambda$, and the maximal ideal $\mathfrak{m}_\gamma = \mathfrak{m}_\lambda \cap \Gamma$ lying under it in Γ .

Applying the results of [9] to F in this situation, we have an algebra $\widehat{F}_\lambda$ which plays the same role as $\widehat{U}_\gamma$, controlling the GT modules in this fiber. One of the advantages of this approach is that the algebra $\widehat{F}_\lambda$ has the same flavor as F itself, but with the group $\widehat{W}$ replaced by the stabilizer of λ in this group. Let $\widehat{W}_\lambda \subset \widehat{W}$ be the stabilizer of $\lambda \in \text{MaxSpec}(\Lambda)$ and let W_λ be the stabilizer of λ in W . Let $\widehat{\Lambda}_\lambda$ be the completion of Λ with respect to this maximal ideal and $\widehat{L}_\lambda$ the fraction field of this completion. Consider the symmetrizing idempotent e_λ in $\mathbb{Z}[\frac{1}{\#W}][W_\lambda]$.

Theorem 1.1 (Propositions 2.11, 2.12 & Lemma 2.13). *The algebra $\widehat{F}_\lambda$ is a principal flag order for the ring $\widehat{\Lambda}$ and the group $\widehat{W}_\lambda$, that is, it is a subalgebra of the skew group ring $\widehat{L}_\lambda\#\widehat{W}_\lambda$ such that $\widehat{F}_\lambda \otimes_{\widehat{\Lambda}_\lambda} \widehat{L}_\lambda \cong \widehat{L}_\lambda\#\widehat{W}_\lambda$, with an induced action on $\widehat{\Lambda}$.*

Furthermore, we have a natural isomorphism

$$\widehat{U}_\gamma = e_\lambda \widehat{F}_\lambda e_\lambda.$$

In particular, if $W_\lambda = \{1\}$, these algebras are isomorphic.

By [17, Theorem 4.1 (4)], the center of $\widehat{F}_\lambda$ is the subalgebra of invariants $\widehat{\Lambda}_\lambda = \widehat{\Lambda}^{\widehat{W}_\lambda}$ and any simple module over $\widehat{F}_\lambda$ will factor through the quotient $F_\lambda^{(1)}$ by the unique maximal

ideal of the center. Thus, this gives a canonical way to choose a finite-dimensional quotient of $\widehat{F}_\lambda$ through which all simples factor.

1.2. The reflection case. The situation will be simpler if we work in the context (studied in [22, § 4.1] and [16]) where we assume that:

- ($\star$) The algebra Λ is the symmetric algebra on a vector space V , the group W is a complex reflection group acting on V , $\mathcal{M}$ is a subgroup of translations, and F is free as a left Λ -module.

In this case, we can always choose F so that U and F are Morita equivalent via the bimodules eF and Fe , and the dimension of $F_\lambda^{(1)}$ is easy to calculate: it is just $(\#\widehat{W}_\lambda)^2$. Furthermore, the quotient by the maximal ideal $\mathfrak{m}_\lambda$ has dimension $\#\widehat{W}_\lambda$ and has every simple module as a quotient. In particular, the sum of the dimensions of the λ -generalized weight space for all simple Gelfand–Tsetlin-modules is $\leq \#\widehat{W}_\lambda$.

If we consider how the results apply to $\widehat{U}_\gamma$, then they are almost unchanged, except that we replace the order of the group $\widehat{W}_\lambda$ with the number of cosets

$$C(\gamma) = \frac{\#\widehat{W}_\lambda}{\#W_\lambda}$$

for any maximal ideal $\mathfrak{m}_\lambda$ lying over $\mathfrak{m}_\gamma$ in Λ . With the assumptions ($\star$), the algebra $U_\gamma^{(1)}$ is $C(\gamma)^2$ -dimensional, and the sum of the dimensions of the γ -generalized weight space for all simple Gelfand–Tsetlin-modules is $\leq S(\gamma)$. This seems to be implicit in the results of [18] such as Theorem 4.12(c), but some of these results are left unstated there¹.

1.3. Coulomb branches. The results of the previous sections are fairly abstract and give no indication of how to actually compute the algebras $U_\gamma^{(1)}$ and understand their representation theory. In this section, we discuss the source of many of the most interesting examples of principal Galois orders: the *Coulomb branches* of Braverman, Finkelberg, and Nakajima [3, 37]. These include the primary motivating example, the orthogonal Gelfand–Zetlin² algebras of Mazorchuk [33] (including $U(\mathfrak{gl}_n)$), and a number of examples that seem to have escaped the notice of experts, such as the spherical Cherednik algebras of the groups $G(\ell, 1, n)$ and hypertoric enveloping algebras.

The Coulomb branch is an algebra constructed from the data of a gauge group G and matter representation N . For example:

- In the case where G is abelian and N arbitrary, the Coulomb branch is a hypertoric enveloping algebra as defined in [2]; the isomorphism of this with a Coulomb branch (defined at a “physical level of rigor”) is proven in [7, § 6.6.2]; it was confirmed this matches the BFN definition of the Coulomb branch in [3, § 4(vii)].
- In the case where $G = GL_n$ and $N = \mathfrak{gl}_n \oplus (\mathbb{C}^n)^{\oplus \ell}$, the Coulomb branch is a spherical Cherednik algebra of the group $G(\ell, 1, n)$ by [29]. Recent work of the author and LePage confirms that the spherical Cherednik algebra for $G(\ell, p, n)$ is also a principal Galois order [31, Proposition 3.16].

¹Note that the published and arXived versions of [18] have different section numbering. We follow the numbering of the published version; in the arXiv version, this is Theorem 5.2(3).

²As any savvy observer knows, there is no universally agreed-upon spelling of Гельфанд–Цетлин in the Latin alphabet; in fact, it’s not even spelled consistently in Russian, since some authors write Цейтлин, a different transliteration of the same name. We will write “Tsetlin” as this is the spelling that will elicit the most correct pronunciation from an English-speaker. However, since “OGZ” is well-established as an acronym, we will not change the spelling of the name of these algebras.

- In the case where

$$G = GL_{v_1} \times \cdots \times GL_{v_{n-1}} \quad (1.1a)$$

$$N = M_{v_n, v_{n-1}}(\mathbb{C}) \oplus M_{v_{n-1}, v_{n-2}}(\mathbb{C}) \oplus \cdots \oplus M_{v_2, v_1}(\mathbb{C}), \quad (1.1b)$$

the Coulomb branch is an orthogonal Gelfand–Zetlin algebra associated to the dimension vector $\mathbf{v} = (v_1, \dots, v_n)$ as shown in [51, § 3.5]. In particular, $U(\mathfrak{gl}_n)$ arises from the vector $\mathbf{v} = (1, 2, 3, \dots, n)$.

In this case, the algebras $U_\gamma^{(1)}$ also have a geometric interpretation in terms of convolution in homology:

Theorem 1.2 (Theorem 4.4). *The Coulomb branch for any group G and representation N is a principal Galois order with $\Lambda = \text{Sym}^\bullet(\mathfrak{t})[\hbar]$, the symmetric algebra on the Cartan of G with an extra loop parameter $\hbar$ and $\widehat{W}$ given by the affine Weyl group of G acting naturally on this space.*

For each maximal ideal $\mathfrak{m}_\gamma$ of Γ , there is a Levi subgroup $G_\gamma \subset G$, with parabolic P_γ and a P_γ -submodule $N_\gamma^- \subset N$ such that

$$\begin{aligned} U_\gamma^{(1)} &\cong H_*^{BM} \left(\left\{ (gP_\gamma, g'P_\gamma, n) \in \frac{G_\gamma}{P_\gamma} \times \frac{G_\gamma}{P_\gamma} \times N \mid n \in gN_\gamma^- \cap g'N_\gamma^- \right\} \right) \\ U_{\mathcal{S}}^{(1)} &\cong \bigoplus_{\gamma, \gamma' \in \mathcal{S}} H_*^{BM} \left(\left\{ (gP_\gamma, g'P_{\gamma'}, n) \in \frac{G_\gamma}{P_\gamma} \times \frac{G_{\gamma'}}{P_{\gamma'}} \times N \mid n \in gN_\gamma^- \cap g'N_{\gamma'}^- \right\} \right) \end{aligned} \quad (1.2)$$

for any set $\mathcal{S}$ contained in a single $\widehat{W}$ -orbit, where the right hand side is endowed with the usual convolution multiplication (as in [8, (2.7.9)]).

The algebra $U_{\mathcal{S}}^{(1)}$ is a *Steinberg algebra* in the sense of Sauter [40]. One notable point to consider is that this algebra is naturally graded. Thus, for any choice of (G, N) and $\widehat{W}$ -orbit $\mathcal{S}$, this gives a graded lift $\widehat{\Gamma\Pi}(\mathcal{S})$ of the category of Gelfand–Zetlin modules supported on this orbit. It's a consequence of the Decomposition Theorem that the classes of simple modules form a dual canonical basis of the Grothendieck group $K_0(\widehat{\Gamma\Pi}(\mathcal{S}))$ (see Theorem 4.11).

Let us now focus on the case of orthogonal Gelfand–Zetlin algebras, so G and N are of the form (1.1a)–(1.1b). The convolution algebras of (1.2) have appeared in numerous places in the literature: they are very closely related to the *KLRW algebras*³ $\widetilde{T}$ as defined in [46, Definition 4.5] corresponding to the Lie algebra $\mathfrak{sl}_n$, with its Dynkin diagram identified as usual with the set $\{1, \dots, n-1\}$. These algebras correspond to a list of highest weights, which we will take to be v_n copies of the $(n-1)^{\text{st}}$ fundamental weight ω_{n-1} ; the dimension vector $(v_1, \dots, v_{n-1})$ determines the number of times that each Dynkin node appears as a label on a black strand. Readers unfamiliar with these algebras can also refer to [23, § 3.1]. The author has proven in [48, Corollary 4.9] that there is a set $\mathcal{S}$ such that:

$$\widetilde{T} \cong \bigoplus_{\gamma, \gamma' \in \mathcal{S}} H_*^{BM, G} \left(\left\{ (gP_\gamma, g'P_{\gamma'}, n) \in \frac{G_\gamma}{P_\gamma} \times \frac{G_{\gamma'}}{P_{\gamma'}} \times N \mid n \in gN_\gamma^- \cap g'N_{\gamma'}^- \right\} \right), \quad (1.3)$$

³Called “Stendhal algebras” or “Webster algebras” in some other sources, and a special case of “reduced weighted KLR algebras” by [48, Theorem 3.5].

n	3	4	5	6	7	8	9
# of simples	20	259	6005	235,546	14,981,789	1.494×10^9	2.275×10^{11}

That is, $\tilde{T}$ is an equivariant Steinberg algebra for the space appearing in (1.2). The set $\mathcal{S}$ appearing here is finite, but if we change it to the set of all integral elements, the RHS of (1.3) gives an algebra Morita equivalent to $\tilde{T}$.

The algebra $\tilde{T}$ is a cousin of the KLR algebras [25], but instead of categorifying the universal enveloping algebra $U(\mathfrak{n})$ of the strictly lower triangular matrices in $\mathfrak{gl}_n$, they categorify the tensor product of $U(\mathfrak{n})$ with the v_n th tensor power of the defining representation of $\mathfrak{gl}_n$; this is proven in [46, Prop. 4.39]. The classes of simple modules over this algebra match the dual canonical basis in this space (which is proven in the course of the proof of [45, Th. 8.7]).

Of course, the difference between the RHS of equations (1.2) and (1.3) is between non-equivariant and equivariant homology. We can account for the difference between these on the LHS by taking an appropriate quotient. That is:

Corollary 1.3. *For $\mathcal{S}$ the set of integral elements of $\text{MaxSpec}(\Gamma)$, the algebra $U_{\mathcal{S}}^{(1)}$ is Morita equivalent to the algebra $\tilde{T}'$, the quotient of $\tilde{T}$ by all positive degree central elements.*

This gives a new way of interpreting the results of [23, § 6]; in particular, Corollary 1.3 is effectively equivalent to Theorem 6.4 of loc. cit. This gives us a criterion in terms of which weight spaces are not zero that classifies the different simple Gelfand–Tsetlin modules with integral weights for an orthogonal Gelfand–Zetlin algebra (Theorem 5.9).

In joint work with Silverthorne [41], we develop the consequences of this connection further and present computer calculations that completely answer Question 1 for Gelfand–Tsetlin modules of $\mathfrak{sl}_3$ and $\mathfrak{sl}_4$ (higher values of n proved to be too much for our computer). This matches the results of Futorny–Grantcharov–Ramirez [14]. We can answer at least one basic question for much higher values of n : the number of simple integral Gelfand–Tsetlin modules in the principal block of $\mathfrak{sl}_n$. These are given by a modified version of the Kostant partition function (for which we know no closed form). In small ranks, these are given by:

This shows the difficulty of answering this question in an *ad hoc* case-by-case manner once $n > 3$ (and especially $n > 4$).

However, this does not preclude systematic study of these questions. As an illustration, we use these results to resolve a question of Mazorchuk [32].

Standard calculations (for example, [35, Theorem 2.20]) compute the spectrum of Γ on every finite-dimensional representation. These correspond precisely to *Gelfand–Tsetlin patterns*, and in the fiber over such a pattern, there is a unique finite-dimensional module. It is natural to ask if there are any infinite-dimensional modules in the fiber over a Gelfand–Tsetlin pattern. One can easily confirm by hand that there are no such infinite-dimensional modules for $\mathfrak{gl}_2$, and the explicit calculations of [14, 41] show that there are no such modules for $n \leq 4$. However, low-rank cases like this can often be deceptive:

Theorem 1.4 (Theorem 6.2). *Let $U = U(\mathfrak{gl}_n)$ and Γ its usual Gelfand–Tsetlin subalgebra. For the maximal ideal $\mathfrak{m}_\gamma \subset \Gamma$ corresponding to a Gelfand–Tsetlin pattern:*

- (1) *If $n \leq 5$, the fiber over γ is a single finite-dimensional irreducible representation.*
- (2) *If $n \geq 6$, the fiber over γ contains an infinite-dimensional irreducible representation.*

We establish this by a hands-on computation using KLR algebras (though locating the relevant example required some help from SageMath).

2. GENERALITIES ON GALOIS ORDERS

Following the notation of [22], let Λ be a noetherian integrally closed domain and L its fraction field. Note that this implies Hartwig's condition (A3), and we lose no generality in assuming this by [22, Lemma 2.1]. Let W be a finite group⁴ acting faithfully on Λ and $\Gamma = \Lambda^W, K = L^W$. Let $\mathcal{M}$ be a submonoid of $\text{Aut}(\Lambda)$ which is normalized by W , and let $\widehat{W} = W \ltimes \mathcal{M}$, which we also assume acts faithfully (this implies Hartwig's (A1) and (A2)). Let $\mathcal{L}$ be the smash product $L \# \mathcal{M}$, $\mathcal{F} = \mathcal{L} \# W$, and $\mathcal{K} = \mathcal{L}^W$. Note that L is an $\mathcal{L}$ module in the obvious way and thus K is a $\mathcal{K}$ -module.

The more general notion of Galois orders was introduced by Futorny and Ovsienko [17], but we will only be interested in a special class of these considered by Hartwig in [22], which makes these properties easy to check.

Definition 2.1 ([22, Definitions 2.22 & 2.24]). The *standard order* (or “universal ring” in the terminology of [34, 43]) is the subalgebra

$$\mathcal{K}_\Gamma = \{X \in \mathcal{K} \mid X(\Gamma) \subset \Gamma\}.$$

A subalgebra $A \subset \mathcal{K}_\Gamma$ containing Γ is a *principal Galois order* if $KA = \mathcal{K}$.

It is a well-known principle in the analysis of quotient singularities that taking the smash product of an algebra with a group acting on it is a much better-behaved operation than taking invariants. Similarly, in the world of Galois orders, there is a larger algebra that considerably simplifies the analysis of these algebras.

Definition 2.2. The *standard flag order* is the subalgebra

$$\mathcal{F}_\Lambda = \{X \in \mathcal{F} \mid X(\Lambda) \subset \Lambda\}.$$

A subalgebra $F \subset \mathcal{F}_\Lambda$ containing Λ is called a *principal flag order* if $LF = \mathcal{F}$ and $W \subset F$.

It is an easy check, via the same proofs, that the analogues of [22, Propositions 2.5, 2.14 & Theorem 2.21] hold here: that is F is a Galois order inside $\mathcal{F}$ with Λ maximal commutative; in order to match the notation of [17], we must take $G = \{1\}$ and $\mathcal{M} = W \ltimes \mathcal{M}$.

Let $e = \frac{1}{\#W} \sum_{w \in W} w \in \mathcal{F}_\Lambda$. Note that $\mathcal{K} \subset \mathcal{F}$ through the obvious inclusion. Given $k \in \mathcal{K}$, the element $eke \in \mathcal{F}$ acts on Γ by the same operator as k . Thus, $k \mapsto eke$ is an algebra isomorphism $\mathcal{K} \cong e\mathcal{F}e$.

Lemma 2.3. *The isomorphism above induces an isomorphism $\mathcal{K}_\Gamma \cong e\mathcal{F}_\Lambda e$.*

Proof. If $a \in \mathcal{F}_\Lambda$, then $ea\Gamma = ea\Gamma \subset e\Lambda = \Gamma$, so $ea\Gamma \subset \Gamma$, so $ea\Gamma \subset \Gamma$, so $ea\Gamma \subset \Gamma$, so $ea\Gamma \subset \Gamma$, so $ea\Gamma \subset \Gamma$. On the other hand, the subalgebra $e\mathcal{K}_\Gamma e$ acts trivially on the elements of Λ that transform by any non-trivial irrep of W , and sends Λ to Λ . This shows that we also have the opposite inclusion $e\mathcal{K}_\Gamma e \subset e\mathcal{F}_\Lambda e$. $\square$

Thus, we have that for any flag order F , the centralizer algebra $U = eFe$ is a principal Galois order. As usual with the centralizer algebra of an idempotent:

⁴Note that this is a departure from the notation of [22], where this group is denoted by G . We will be most interested in the case where W is the Weyl group of a semisimple Lie group acting on the Cartan, so we prefer to save G for the name of the Lie group.

Lemma 2.4. *The category of U -modules is a quotient of the category of F -modules through the functor $M \mapsto eM$; that is, this functor is exact and has right and left adjoints $N \mapsto Fe \otimes_U N$ and $N \mapsto \operatorname{Hom}_U(eF, N)$ that split the quotient functor.*

Furthermore, every principal Galois order appears this way. Consider the smash product $\Lambda \# W \subset \operatorname{End}_{\Lambda^W}(\Lambda)$, and let D be a subalgebra that satisfies $\Lambda \# W \subset D \subset \operatorname{End}_{\Gamma}(\Lambda) \subset L \# W$. Note that in this case, $eDe = \Gamma$, since this is true when $D = \Lambda \# W$ or $D = \operatorname{End}_{\Gamma}(\Lambda)$. Let F_D be the subalgebra generated by $D \subset \mathcal{F}_{\Lambda}$ and by eUe , which is the image of $U \subset \mathcal{K}_{\Gamma}$ under its isomorphism to $e\mathcal{F}_{\Lambda}e$. Although we have a canonical isomorphism $U \cong eUe$, it is helpful to distinguish these in the notation, since their natural actions on Λ are different. Indeed, every element of eUe acts by 0 on $(1 - e)L$, which is a complementary K -subspace to K inside L . While eUe is not a unital subalgebra of $\mathcal{F}_{\Lambda}$, the subalgebra D is, so F_D is a unital subalgebra.

Lemma 2.5. *For any principal Galois order U , and any D as above, the subalgebra F_D is a principal flag order such that $U \cong eF_De$.*

Proof. $[F_D \subset \mathcal{F}_{\Lambda}]$: As discussed above, we have that $eUe \cdot \Lambda = eU\Gamma$. Thus, by the principal Galois order property of U , we have

$$eUe \cdot \Lambda = eU \cdot \Gamma \subset \Gamma \subset \Lambda. \quad (2.1)$$

The subalgebra D preserves Λ by definition, so combined with (2.1), we have $F_D \subset \mathcal{F}_{\Lambda}$.

$[LF_D = \mathcal{F} \text{ and } W \subset F_D]$: Furthermore, we have $LF_D \supset \mathcal{L}\Lambda W = \mathcal{L}W = \mathcal{F}$ and by construction, the algebra F_D contains the smash product $\Lambda \# W$.

Thus, the paragraphs above show that F_D satisfies the conditions of a principal flag order.

$[eF_De = U]$: The algebra F_D is spanned by elements of the form $f = d_0 e u_1 e d_1 e \cdots e u_n e d_n$ for $d_i \in D$ and $u_i \in U$. Thus, we have

$$efe = (ed_0e)u_1(ed_1e) \cdots u_n(ed_ne).$$

Since $ed_ie \in \Gamma$, this product lies in U . □

2.1. Gelfand–Tsetlin modules. Now, fix a principal flag order $F \subset \mathcal{F}_{\Lambda}$. We wish to understand the representation theory of this algebra. Consider the weight functors

$$\mathcal{W}_{\lambda}(M) = \left\{ m \in M \mid \mathfrak{m}_{\lambda}^N m = 0 \text{ for some } N \gg 0 \right\}$$

for $\lambda \in \operatorname{MaxSpec}(\Lambda)$. The reader might reasonably be concerned about the fact that this is a *generalized* eigenspace; in this paper, we will always want to consider these, and thus will omit “generalized” before instances of “weight.”

Definition 2.6. We call a finitely generated F -module M a *weight module* or *Gelfand–Tsetlin module* if $M = \bigoplus_{\lambda \in \operatorname{MaxSpec}(\Lambda)} \mathcal{W}_{\lambda}(M)$.

Remark 2.7. One subtlety here is that we have not assumed that $\mathcal{W}_{\lambda}(M)$ is finite-dimensional. We’ll see below that this holds automatically if the stabilizer of λ in $\widehat{W}$ is finite.

Since many readers will be more interested in the Galois order $U = eFe$, let us compare the weight spaces of a module M with those of the U -module eM . Recall that W_{λ} is the stabilizer of λ in W , and let $e_{\lambda} \in \mathbb{Z}[\frac{1}{\#W}][W_{\lambda}]$ be the symmetrizing idempotent. Of course,

in U , we only have an action of Γ . Let $\gamma \in \text{MaxSpec}(\Gamma)$ be the image of λ under the obvious map, $\mathfrak{m}_\gamma \subset \Gamma$ the corresponding maximal ideal and

$$\mathcal{W}_\gamma(M) = \left\{ m \in eM \mid \mathfrak{m}_\gamma^N m = 0 \ \forall N \gg 0 \right\}.$$

Lemma 2.8. *If M is a Gelfand–Tsetlin F -module, then eM is a Gelfand–Tsetlin U -module with*

$$\mathcal{W}_\gamma(eM) \cong e_\lambda \mathcal{W}_\lambda(M).$$

Proof. Let $\mathfrak{m}_\gamma = \Gamma \cap \mathfrak{m}_\lambda$. By standard commutative algebra, the other maximal ideals lying over $\mathfrak{m}_\gamma$ are those in the orbit $W \cdot \lambda$. Thus, we have that

$$\mathcal{W}_\gamma(eM) = e \cdot \left(\bigoplus_{\lambda' \in W\lambda} \mathcal{W}_{\lambda'}(M) \right).$$

This space $\bigoplus_{\lambda' \in W\lambda} \mathcal{W}_{\lambda'}(M)$ has a W -action induced by the inclusion $W \subset F$, and is isomorphic to the induced representation $\text{Ind}_{W_\lambda}^W \mathcal{W}_\lambda(M)$ since it is a sum of subspaces which it permutes like the cosets of this subgroup. Thus, its invariants are canonically isomorphic to the invariants for W_λ on $\mathcal{W}_\lambda(M)$. $\square$

2.2. The fiber for a flag order.

Definition 2.9. Fix an integer N . The *universal Gelfand–Tsetlin module* of weight λ and length N is the quotient $F/F\mathfrak{m}_\lambda^N$.

This is indeed a Gelfand–Tsetlin -module by [22, Lem. 3.2]. Since a homomorphism $F/F\mathfrak{m}_\lambda^N \rightarrow M$ for any module M is determined by the image of $\bar{1} \in F/F\mathfrak{m}_\lambda^N$, the module $F/F\mathfrak{m}_\lambda^N$ represents the functor of taking generalized weight vectors killed by $\mathfrak{m}_\lambda^N$:

$$\text{Hom}_F(F/F\mathfrak{m}_\lambda^N, M) = \left\{ m \in M \mid \mathfrak{m}_\lambda^N m = 0 \right\}.$$

In particular, every simple Gelfand–Tsetlin-module with $\mathcal{W}_\lambda(S) \neq 0$ is a quotient of $F/F\mathfrak{m}_\lambda$, since it must have a vector killed by $\mathfrak{m}_\lambda$. Taking the inverse limit $\varprojlim F/F\mathfrak{m}_\lambda^N$, we obtain a universal (topological) Gelfand–Tsetlin module of arbitrary length. Consider the algebra

$$\widehat{F}_\lambda = \varprojlim F / (F\mathfrak{m}_\lambda^N + \mathfrak{m}_\lambda^N F)$$

As noted in [9, Th. 18], this algebra controls the λ weight spaces of all modules, and in particular simple modules, in the sense that for every simple Gelfand–Tsetlin-module with $\mathcal{W}_\lambda(S) \neq 0$, the $\widehat{F}_\lambda$ -module $\mathcal{W}_\lambda(S)$ is simple, and every simple $\widehat{F}_\lambda$ appears this way for a unique simple Gelfand–Tsetlin-module.

Let $\widehat{W}_\lambda$ be the subgroup of $\widehat{W} = W \ltimes \mathcal{M}$ which fixes λ . For the remainder of this section, we assume that $\widehat{W}_\lambda$ is finite. This implies that Λ is finitely generated over $\Lambda_\lambda = \Lambda^{\widehat{W}_\lambda}$.

Definition 2.10. Let F_λ be the intersection $F \cap L \cdot \widehat{W}_\lambda \subset \mathcal{F} = L\widehat{W}$ with the L -span of $\widehat{W}_\lambda$. Since F_λ is the intersection of two subalgebras, it is itself a subalgebra.

Let ${}_\lambda \widehat{W}_\mu$ be the set of elements of $\widehat{W}$ such that $w \cdot \mu = \lambda$. Let ${}_\lambda F_\mu = F \cap K \cdot {}_\lambda \widehat{W}_\mu$ be the elements of F which are in the K -span of ${}_\lambda \widehat{W}_\mu$. This is an F_λ - F_μ bimodule

This has an obvious left and right module structure over Λ but Λ is not central. In the notation of [18, (3)], this would be $F(\widehat{W}_\lambda)$. Let $\widehat{\Lambda}$ be the completion of Λ in the $\mathfrak{m}_\lambda$ -adic topology, and let $\widehat{L}$ be the fraction field of $\widehat{\Lambda}$.

Proposition 2.11.

- (1) The bimodule ${}_{\lambda}F_{\mu}$ is finitely generated as a left module and as a right module over Λ and satisfies ${}_{\lambda}F_{\mu}L = L{}_{\lambda}F_{\mu} = L \cdot {}_{\lambda}\widehat{W}_{\mu}$.
- (2) In fact, F_{λ} is a Galois order for the group $\mathcal{M} = \widehat{W}_{\lambda}$ and the commutative ring Λ , using the notation of [17].
- (3) The image of ${}_{\lambda}F_{\mu}$ spans $F/(F\mathfrak{m}_{\mu}^N + \mathfrak{m}_{\lambda}^N F)$ for all N .
- (4) The bimodule ${}_{\lambda}\widehat{F}_{\mu}$ is the completion of ${}_{\lambda}F_{\mu}$ with respect to the topology induced by the basis of neighborhoods of the identity $F_{\lambda}\mathfrak{m}_{\mu}^N + \mathfrak{m}_{\lambda}^N F_{\lambda}$. In particular, the $\widehat{L}$ -vector space $\widehat{L} \otimes_{\Lambda} {}_{\lambda}\widehat{F}_{\mu}$ has basis ${}_{\lambda}\widehat{W}_{\mu}$.

This shows in particular that Λ is *big at λ* in the terminology of [9].

Proof.

(1): Finite generation is an immediate consequence of the fact that F is an order.

Thus, it only remains to show that ${}_{\lambda}\widehat{F}_{\mu}L = L{}_{\lambda}\widehat{F}_{\mu} = L \cdot {}_{\lambda}\widehat{W}_{\mu}$. The inclusions ${}_{\lambda}\widehat{F}_{\mu}L \subset L \cdot {}_{\lambda}\widehat{W}_{\mu} \supset L{}_{\lambda}\widehat{F}_{\mu}$ are obvious by definition, so we only need to prove the opposite inclusions. Since $\mathcal{F} = LF$, for any $w \in {}_{\lambda}\widehat{W}_{\mu}$, we have $w = \sum k_i f_i$ for $k_i \in K$, and $f_i \in F$. Let $T \subset \widehat{W}$ be the support of the f_i 's. If $T \subset {}_{\lambda}\widehat{W}_{\mu}$, then we are done, so let us prove this by induction on the number of elements $t \in T \setminus {}_{\lambda}\widehat{W}_{\mu}$. Fix such a t . We have a polynomial p vanishing at λ , but not at $t^{-1} \cdot \lambda$. Note that for w as above, we have $w = \frac{1}{p^t - p^w} (p^t w - wp)$, with the $p^t - p^w$ being non-zero in K since it does not vanish at λ . Substituting into our formula for w , we have

$$w = \sum \frac{k_i}{p^t - p^w} (p^t f_i - f_i p).$$

The element $p^t f_i - f_i p \in F$ has support on $T \setminus \{t\}$. Thus, we can inductively reduce the size of T until $T \subset {}_{\lambda}\widehat{W}_{\mu}$. That is, we can assume that $f_i \in {}_{\lambda}\widehat{F}_{\mu}$. This completes the proof that $L \cdot {}_{\lambda}\widehat{W}_{\mu} = L{}_{\lambda}\widehat{F}_{\mu}$; the proof for $L \cdot {}_{\lambda}\widehat{W}_{\mu} = {}_{\lambda}\widehat{F}_{\mu}L$ is identical.

(2): The property $F_{\lambda}L = LF_{\lambda} = L \cdot \widehat{W}_{\lambda}$ which we have already verified shows that F_{λ} is a Galois ring.

The ring F_{λ} inherits the order property, i.e. its intersection with any finite-dimensional L -subspace Z for the left/right action of $\mathcal{F}_{\lambda}$ is finitely generated for the left/right action of Λ , from the order F . After all, Z is a finite-dimensional subspace of $\mathcal{F}$, so the order property of F implies that $Z \cap F_{\lambda} = Z \cap F$ is finitely generated over Λ .

(3): This is [18, Theorem 4.7] in the case where $S = {}_{\lambda}\widehat{W}_{\mu}$, $\mathfrak{m} = \mathfrak{m}_{\lambda}$ and $\mathfrak{n} = \mathfrak{m}_{\mu}$.

(4): By point (3), ${}_{\lambda}\widehat{F}_{\mu}$ is the completion of ${}_{\lambda}F_{\mu}$ with respect to the subspace topology, that is, the topology with a basis of neighborhoods given by ${}_{\lambda}F_{\mu} \cap (F\mathfrak{m}_{\mu}^N + \mathfrak{m}_{\lambda}^N F)$. Since

$${}_{\lambda}F_{\mu}\mathfrak{m}_{\mu}^N + \mathfrak{m}_{\lambda}^N {}_{\lambda}F_{\mu} \subset {}_{\lambda}F_{\mu} \cap (F\mathfrak{m}_{\mu}^N + \mathfrak{m}_{\lambda}^N F),$$

we will have the desired equivalence of topologies if we prove that the inclusion above is an equality. Consider the quotient Λ - Λ -bimodule

$$Q_N = \frac{{}_{\lambda}F_{\mu} \cap (F\mathfrak{m}_{\mu}^N + \mathfrak{m}_{\lambda}^N F)}{F_{\lambda}\mathfrak{m}_{\mu}^N + \mathfrak{m}_{\lambda}^N F_{\lambda}}.$$

Consider the ideal $\mathfrak{M}_{\lambda} = \mathfrak{m}_{\lambda} \otimes \Lambda + \Lambda \otimes \mathfrak{m}_{\mu}$. Note that

$$\mathfrak{M}_{\lambda}^{2N} \cdot Q_N \subset Q_N \mathfrak{m}_{\mu}^N + \mathfrak{m}_{\lambda}^N Q_N = 0. \quad (2.2)$$

Assume that $f = \sum a_i g_i + h_i b_i$ for $a_i, b_i \in F$ and $g_i \in \mathfrak{m}_\lambda^N, h_i \in \mathfrak{m}_\mu^N$. By [18, Theorem 4.7], we can choose $a'_i, b'_i \in \mathfrak{M}_\lambda \cdot F$ such that $a_i - a'_i, b_i - b'_i \in {}_\lambda F_\mu$. The image of f in Q_N is the same as that of $f' = \sum a'_i g_i + h_i b'_i \in \mathfrak{M}_\lambda Q_N$. This implies that $Q_N = \mathfrak{M}_\lambda Q_N$. Combining this with (2.2), Nakayama's lemma shows that $Q_N = 0$, completing the proof. $\square$

Alternatively, we can think about this topology by noting that F_λ is finitely generated over $\Lambda_\lambda = \Lambda^{\widehat{W}_\lambda}$. Furthermore, Λ_λ is central in F_λ , since it commutes with $L \cdot \widehat{W}_\lambda$; in fact, by Lemma 2.11 above and [17, Theorem 4.1(4)], it is the full center of this algebra. Let $\mathfrak{n}_\lambda = \mathfrak{m}_\lambda \cap \Lambda_\lambda$. Since λ is fixed by $\widehat{W}_\lambda$ (by definition), the ideal $\mathfrak{n}_\lambda \Lambda$ still only vanishes at λ , that is, $\mathfrak{n}_\lambda \Lambda \supset \mathfrak{m}_\lambda^k$ for some k .

Let $\widehat{\Lambda}_\lambda$ be the completion of Λ_λ in the $\mathfrak{n}_\lambda$ -adic topology.

Proposition 2.12. *We have an isomorphism of topological bimodules*

$${}_\lambda \widehat{F}_\mu \cong {}_\lambda F_\mu \otimes_{\Lambda_\mu} \widehat{\Lambda}_\mu \cong \widehat{\Lambda}_\lambda \otimes_{\Lambda_\lambda} {}_\lambda F_\mu$$

In particular, the ring $\widehat{F}_\lambda$ is a Galois order for $\mathcal{M} = \widehat{W}_\lambda$ and the ring $\widehat{\Lambda}$.

Proof. The tensor product $F_\lambda \otimes_{\Lambda_\lambda} \widehat{\Lambda}_\lambda$ is the completion of F_λ with respect to the topology with basis of 0 given by the 2-sided ideals $F_\lambda \mathfrak{n}_\lambda^N$. Since $\Lambda \mathfrak{n}_\lambda \supset \mathfrak{m}_\lambda^k$ for some k , we have that

$$F_\lambda \mathfrak{m}_\lambda^{kN} + \mathfrak{m}_\lambda^{kN} F_\lambda \subset F_\lambda \mathfrak{n}_\lambda^N \subset F_\lambda \mathfrak{m}_\lambda^N + \mathfrak{m}_\lambda^N F_\lambda$$

which shows the equivalence of the topologies, and thus the isomorphism of completions. Faithful base change by a central subalgebra preserves the properties of being a Galois order, so this follows from Lemma 2.11. $\square$

We can also use these results to understand the fiber for U as well for any principal Galois order. By Lemma 2.5, we can choose a principal flag order with $U = eFe$. The algebra F_λ contains the stabilizer W_λ and its symmetrizing idempotent e_λ . As before, let γ be the image of λ in $\text{MaxSpec}(\Gamma)$ and $\mathfrak{m}_\gamma = \Gamma \cap \mathfrak{m}_\lambda$.

It is worth noting that there is no obvious analogue of $\widehat{W}_\lambda$ and F_λ in the context of U . The closest analogue is the set $S(\mathfrak{m}_\gamma, \mathfrak{m}_\gamma)$ defined in [18, § 4.1]. This is the subset of $\mathcal{M} = \widehat{W}/W$ such that $m \cdot w'\lambda = w\lambda$ for some $w, w' \in W$. In this case, we have that $w^{-1}mw \in \widehat{W}_\lambda$. Put differently, $S(\mathfrak{m}_\gamma, \mathfrak{m}_\gamma)$ is given by the W -saturation of the image of $\widehat{W}_\lambda$, i.e. the union of all W - W double cosets WwW/W for $w \in \widehat{W}_\lambda$. We have a surjective map $\widehat{W}_\lambda/W_\lambda \twoheadrightarrow S(\mathfrak{m}_\gamma, \mathfrak{m}_\gamma)$, but this is not necessarily injective: The image contains an element of each W -orbit, but is not W -invariant.

However, there is a close relationship between F_λ and the algebra $\widehat{U}_\gamma$. Note that the ideal $\Lambda \mathfrak{m}_\gamma$ has vanishing set given by the orbit $W \cdot \lambda$. Thus, if we consider the completion of Λ at this ideal, the result is $\bigoplus_{\lambda' \in W \cdot \lambda} \widehat{\Lambda}_{\lambda'}$. Since passing invariants for W is exact, it commutes with inverse limits. Thus, the completion $\widehat{\Gamma}_\gamma$ of Γ at the maximal ideal $\mathfrak{m}_\gamma$ is isomorphic to

$$\widehat{\Gamma}_\gamma \cong e_\lambda \widehat{\Lambda}_\lambda = \widehat{\Lambda}_\lambda^{W_\lambda} \cong \left(\bigoplus_{\lambda' \in W \cdot \lambda} \widehat{\Lambda}_{\lambda'} \right)^W.$$

Lemma 2.13. *The above isomorphism induces an isomorphism $\widehat{U}_\gamma \cong e_\lambda \widehat{F}_\lambda e_\lambda$.*

Proof. Consider the completion $\widehat{F}_\gamma = F/(F\mathfrak{m}_\gamma^N + \mathfrak{m}_\gamma^N F)$; as discussed above, since $\mathfrak{m}_\gamma$ is an ideal defining the orbit $W \cdot \lambda$, this decomposes as the sum $\bigoplus_{\lambda', \lambda'' \in W \cdot \lambda} F/(F\mathfrak{m}_{\lambda'}^N + \mathfrak{m}_{\lambda''}^N F)$ by the Chinese Remainder theorem. In particular, this means that as a $W \times W$ -representation $\widehat{F}_\gamma = \text{Ind}_{W_\lambda \times W_\lambda}^{W \times W} \widehat{F}_\lambda$.

The exactness of taking invariants shows that $\widehat{U}_\gamma$ is $e\widehat{F}_\gamma e$, or put differently, the invariants of F_γ under the action of W by left and right multiplication. Thus, we obtain the desired isomorphism. $\square$

2.3. Universal modules. While this is largely redundant with [9], it will be helpful to explain how we construct simple Gelfand–Tsetlin modules.

Definition 2.14. Fix an integer N . The *central universal Gelfand–Tsetlin module* of weight λ and length N is the quotient $P_\lambda^{(N)} = F/F\mathfrak{n}_\lambda^N$.

Consider the quotient algebra $F_\lambda^{(N)} := F_\lambda/F_\lambda\mathfrak{n}_\lambda^N$.

Theorem 2.15. The module $P_\lambda^{(N)}$ is a Gelfand–Tsetlin module such that

$$\mathcal{W}_\lambda(P_\lambda^{(N)}) \cong \text{End}(P_\lambda^{(N)}) \cong F_\lambda^{(N)}.$$

More generally, we have that

$$\text{Hom}_F(P_\lambda^{(N)}, M) = \{m \in M \mid \mathfrak{n}_\lambda^N m = 0\}. \quad (2.3)$$

Proof. Equation (2.3) is a basic property of left ideals. This is a Gelfand–Tsetlin module by [22, Lem. 3.2].

Note that the map $F_\lambda \rightarrow \mathcal{W}_\lambda(P_\lambda^{(N)})$ is surjective by construction. Of course, the kernel of this map is $F_\lambda \cap F\mathfrak{n}_\lambda^N = F_\lambda\mathfrak{n}_\lambda^N$. This shows that $\mathcal{W}_\lambda(P_\lambda^{(N)}) \cong F_\lambda/F_\lambda\mathfrak{n}_\lambda^N$. Since $\mathfrak{n}_\lambda^N$ is central in F_λ , it acts trivially on this weight space, and the identification with $\text{End}(P_\lambda^{(N)})$ follows from (2.3). $\square$

Note that “length N ” refers to the maximal length of a Jordan block of an element of $\mathfrak{n}_\lambda$, not of $\mathfrak{m}_\lambda$. Since $\mathfrak{n}_\lambda$ is central in F_λ , the ideal $\mathfrak{n}_\lambda^N$ acts trivially on $P_\lambda^{(N)}$. More generally:

Lemma 2.16. The ideal $\mathfrak{n}_\mu^N$ acts trivially on the weight space $\mathcal{W}_\mu(P_\lambda^{(N)})$.

Proof. Note that translation by $\mu - \lambda$ induces an automorphism $\sigma_{\mu, \lambda}: \Lambda \rightarrow \Lambda$, and that $\sigma_{\mu, \lambda}(\mathfrak{n}_\lambda^N) = \mathfrak{n}_\mu^N$. We have $\sigma_{\mu, \lambda}(g)w = wg$ for any $w \in {}_\mu\widehat{W}_\lambda, g \in \Lambda^W$, so by linearity, $\sigma_{\mu, \lambda}(g)f = fg$ for any $f \in {}_\mu F_\lambda, g \in \Lambda^W$. Since $\mathfrak{n}_\lambda^N$ is generated by W -invariant elements, this implies that $\mathfrak{n}_\mu^N f = f\mathfrak{n}_\lambda^N$ for any $f \in {}_\mu F_\lambda$.

Let $\bar{1}$ denote that image of $1 \in F$ in $P_\lambda^{(N)}$. The elements of $\mathcal{W}_\mu(P_\lambda^{(N)})$ are precisely those of the form $f \cdot \bar{1}$ for $f \in {}_\mu F_\lambda$. Thus, the commutation above shows that $\mathfrak{n}_\mu^N f \cdot \bar{1} = f\mathfrak{n}_\lambda^N \cdot \bar{1} = 0$, showing the desired vanishing. $\square$

On the other hand the nilpotent length of the action of $\mathfrak{m}_\mu$ on $\mathcal{W}_\mu(P_\lambda^{(N)})$ is typically more than N ; the argument above fails because the generators of $\mathfrak{m}_\lambda$ aren’t W -invariant.

It follows immediately from [9, Theorem 18] that:

Theorem 2.17. The map sending $S \mapsto \mathcal{W}_\lambda(S)$ is a bijection between the isoclasses of simple Gelfand–Tsetlin F -modules in the fiber over λ and simple $F_\lambda^{(1)}$ -modules.

Similarly, we can define a U -module $Q_\gamma^{(N)} = eP_\lambda^{(N)}e_\lambda = e(P_\lambda^{(N)})^{W_\lambda}$ such that

$$\mathcal{W}_\gamma(Q_\gamma^{(N)}) \cong \text{End}(Q_\gamma^{(N)}) \cong U_\lambda^{(N)} = e_\lambda F_\lambda^{(N)} e_\lambda.$$

More generally, we have that

$$\text{Hom}_F(Q_\lambda^{(N)}, M) = \{n \in N \mid (\Lambda \mathfrak{n}_\lambda^N \cap \Gamma) m = 0\}. \quad (2.4)$$

Applying [9, Theorem 18] again shows that the map sending $S \mapsto \mathcal{W}_\gamma(S)$ is a bijection between the isoclasses of simple Gelfand–Tsetlin U -modules in the fiber over γ and simple $U_\gamma^{(1)}$ -modules.

2.4. Dimension bounds. In [18, Theorem 4.12], bounds are given on the number and dimensions of the irreducible representations in the fiber over a given maximal ideal γ . In this section, we explain how related bounds can be recovered in our framework.

Let $\sigma(\lambda, \lambda)$ be the minimum number of generators of $\widehat{F}_\lambda$ as a right $\widehat{\Lambda}_\lambda$ -module. Note that the minimum number of generators of F_λ is an upper bound on $\sigma(\lambda, \lambda)$. Since $\widehat{F}_\lambda \otimes_{\widehat{\Lambda}_\lambda} \widehat{L}_\lambda$ is $\#\widehat{W}_\lambda$ dimensional over L , we have that $\sigma(\lambda, \lambda) \geq \#\widehat{W}_\lambda$. We will have equality if and only if $\widehat{F}_\lambda$ is free over $\widehat{\Lambda}_\lambda$, which will follow if F_λ is free over Λ (in particular if F is free over Λ).

Let us note how these statistics compare with those in [18, § 4]. In [18, 4.1(c)], the set $\widehat{\mathcal{M}}_\lambda = \{m \in \mathcal{M} \mid m\lambda \in W \cdot \lambda\}$ is considered. If we write an element of $\widehat{W}_\lambda$ as $w^{-1}m$ for $w \in W, m \in \mathcal{M}$, then we will have $m \in \widehat{\mathcal{M}}_\lambda$. The induced map $\widehat{W}_\lambda \rightarrow \widehat{\mathcal{M}}_\lambda$ has fiber given by the choices of $w \in W$ such that $w\lambda = m\lambda$; these form a single coset in W/W_λ . This shows that $\#\widehat{\mathcal{M}}_\lambda = \frac{\#\widehat{W}_\lambda}{\#W_\lambda}$.

In any simple $F_\lambda^{(1)}$ -module, there is a vector where $\mathfrak{m}_\lambda$ acts trivially. As discussed before, this means that:

Proposition 2.18. *Any simple $F_\lambda^{(1)}$ -module appears as a quotient of $F_\lambda/F_\lambda\mathfrak{m}_\lambda$. If $\widehat{F}_\lambda$ is a free module over $\widehat{\Lambda}$ (necessarily of rank $\#\widehat{W}_\lambda$) then $\dim_{\Lambda/\mathfrak{m}_\lambda} F_\lambda/F_\lambda\mathfrak{m}_\lambda = \#\widehat{W}_\lambda$.*

Proof. The algebra $\Lambda_\lambda^{(1)} = \Lambda_\lambda/\mathfrak{n}_\lambda$ is a local commutative subalgebra of the finite-length algebra $F_\lambda^{(1)}$. Thus, for any $F_\lambda^{(1)}$ -module M , some power of the maximal ideal $\mathfrak{m}_\lambda$ kills M . Let n be maximal such that $\mathfrak{m}_\lambda^n M \neq 0$. In this case, $\mathfrak{m}_\lambda^n M$ is a nonzero subspace of M killed by $\mathfrak{m}_\lambda$, so any non-zero element of this space induces a non-zero map $F_\lambda/F_\lambda\mathfrak{m}_\lambda \rightarrow M$, which is surjective if M is simple. $\square$

Combining this with Theorem 2.17 above, we have that:

Corollary 2.19. *The dimensions of the λ -weight spaces in the simples over F in the fiber over λ have sum $\leq \sigma(\lambda, \lambda)$, and thus $\leq \#\widehat{W}_\lambda$ if F_λ is a free right module over Λ .*

The dimensions of the γ -weight spaces in the simple U -modules in the fiber over γ have sum $\leq \frac{\sigma(\lambda, \lambda)}{\#W_\lambda}$, and thus $\leq \frac{\#\widehat{W}_\lambda}{\#W_\lambda}$ if F_λ is a free right module over Λ .

We can generalize these results to be closer to [18]. Let $\sigma(\mu, \lambda)$ be the minimal number of generators of ${}_\lambda F_\mu$ as a Λ -module.

Corollary 2.20. *The dimensions of the λ -weight spaces in the simples over F in the fiber over μ have sum $\leq \sigma(\mu, \lambda)$, and thus $\leq \#{}_\lambda \widehat{W}_\mu$ if ${}_\lambda F_\mu$ is a free right module over Λ .*

If γ' lies under μ , then the dimensions of the γ -weight spaces in the simple U -modules in the fiber over γ' have sum $\leq \frac{\sigma(\mu, \lambda)}{\#W_\lambda}$, and thus $\leq \frac{\#_\lambda \widehat{W}_\mu}{\#W_\lambda}$ if ${}_\lambda F_\mu$ is a free module over Λ .

Remark 2.21. In [18, Theorem 4.12(c)], Futorny and Ovsienko show similar bounds but using a slightly different looking statistic $\#(W \backslash S(\mathfrak{m}_\mu, \mathfrak{m}_\lambda))$. In [18, Theorem 4.1(c)], they show that this is less than or equal to the size of the set $\{m \in \mathcal{M} \mid m\mu \in W \cdot \lambda\}$. As mentioned above, sending $wm \mapsto m$ is a $\#W_\lambda$ -to-1 map from ${}_\lambda \widehat{W}_\mu$ to the set $\{m \in \mathcal{M} \mid m\mu \in W \cdot \lambda\}$, so [18, Theorem 4.1(c)] can be rewritten as

$$\#(W \backslash S(\mathfrak{m}_\mu, \mathfrak{m}_\lambda)) \leq \frac{\# \widehat{W}_\lambda}{\#W_\lambda}. \quad (2.5)$$

While we have not found an example where this inequality is strict, it seems likely that they exist. The explanation for this difference between these bounds is that there could potentially be F -modules M cyclically generated by a vector of weight μ such that eM is not cyclically generated over U . Such a module exists for $F = \widehat{W} \# \Lambda$ if and only if the inequality (2.5) is strict for some λ .

2.5. Weightification and canonical modules. There is another natural way to try to construct Gelfand–Tsetlin modules. Consider any F -module M , and fix an $\widehat{W}$ -invariant subset $\mathcal{S} \subset \text{MaxSpec}(\Lambda)$.

Definition 2.22. Consider the sums

$$M^{\mathcal{S}} = \bigoplus_{\lambda \in \mathcal{S}} \{m \in M \mid \mathfrak{n}_\lambda m = 0\} \quad M_{\mathcal{S}} = \bigoplus_{\lambda \in \mathcal{S}} M/\mathfrak{n}_\lambda M$$

We can define actions of F on these sums as follows: given $f \in F$, the Λ -module $Q_{f, \lambda} = \Lambda f \Lambda / \Lambda f \Lambda \mathfrak{n}_\lambda$ is finite length, and thus the sum of finitely many weight spaces

$$Q_{f, \lambda} = \bigoplus_{i=1}^m \mathcal{W}_{\mu_i}(Q_{f, \lambda}).$$

By [18, Theorem 4.7], we can write $f = f_{(1)} + f'$ where $f_{(1)} \in {}_{\mu_1} F_\lambda$, $f' \in \mathfrak{n}_{\mu_1} f + f \mathfrak{n}_\lambda$; applying this inductively, we can write

$$f = f_{(1)} + f_{(2)} + \cdots + f_{(k)} + f_0 \quad \text{where} \quad f_{(i)} \in {}_{\mu_i} F_\lambda, \quad f_0 \in F \mathfrak{n}_\lambda.$$

The elements $f_{(i)}$ are unique up to the addition of an element of ${}_{\mu_i} F_\lambda \cap (F \mathfrak{n}_\lambda + \mathfrak{n}_{\mu_1} F) = {}_{\mu_i} F_\lambda \mathfrak{n}_\lambda$. Acting by $f_{(i)}$ gives natural maps

$$\{m \in M \mid \mathfrak{n}_\lambda m = 0\} \rightarrow \{m \in M \mid \mathfrak{n}_{\mu_i} m = 0\} \quad M/\mathfrak{n}_\lambda M \rightarrow M/\mathfrak{n}_{\mu_i} M;$$

these maps are independent of the choice of $f_{(i)}$ since ${}_{\mu_i} F_\lambda \mathfrak{n}_\lambda$ acts trivially in both cases.

Theorem 2.23. The ring F acts on $M^{\mathcal{S}}$ and $M_{\mathcal{S}}$ by the formula $f \cdot m = \sum_{i=1}^k f_{(i)} m$ and this module structure is Gelfand–Tsetlin.

Note that even if M is a finitely generated module, the modules $M^{\mathcal{S}}$ and $M_{\mathcal{S}}$ may not be finitely generated, though the individual weight spaces

$$\mathcal{W}_\lambda(M^{\mathcal{S}}) = \{m \in M \mid \mathfrak{n}_\lambda m = 0\} \quad \mathcal{W}_\lambda(M_{\mathcal{S}}) = M/\mathfrak{n}_\lambda M$$

will be finitely generated over $\Lambda_\lambda^{(1)} = \Lambda/\Lambda \mathfrak{n}_\lambda$.

Proof. The module action is well-defined: We need to check that $f(gm) = (fg)m$ for all $f, g \in F$, and $m \in M^{\mathcal{S}}$ or $M_{\mathcal{S}}$. Without loss of generality, we can assume that f, g both have $m = 1$, i.e., that gm is a weight vector, as is $f(gm)$. That is $f = f_{(1)} + f_0, g = g_{(1)} + g_0$ where

$$f_{(1)} \in {}_{\nu}F_{\mu} \quad g_{(1)} \in {}_{\mu}F_{\lambda} \quad f_0 \in F\mathfrak{n}_{\mu} \quad g_0 \in F\mathfrak{n}_{\lambda}.$$

for some ν, μ, λ , and that $\mathfrak{n}_{\lambda}m = 0$ or $m \in M/\mathfrak{n}_{\lambda}M$. In either case, $fg = f_{(1)}g_{(1)} + f_{(1)}g_0 + f_0g_{(1)} + f_0g_0$. The first term lies in ${}_{\nu}F_{\lambda}$, and the last three all lie in $F\mathfrak{n}_{\lambda}$. Thus, we have:

$$f(gm) = f(g_{(1)}m) = f_{(1)}g_{(1)}m = (fg)m$$

completing the proof that the module action is well-defined.

The module is Gelfand–Tsetlin: By construction, $\mathcal{W}_{\lambda}(M^{\mathcal{S}}) = \{m \in M \mid \mathfrak{n}_{\lambda}m = 0\}$ and $\mathcal{W}_{\lambda}(M_{\mathcal{S}}) = M/\mathfrak{n}_{\lambda}M$, and these modules are the direct sums of these spaces by construction. $\square$

We could similarly consider “thicker” versions of these modules where we replace $\mathfrak{n}_{\lambda}$ with powers of this ideal, and direct/inverse limits of the resulting modules. Since we have no application in mind for these modules, we will leave discussion of them to another time.

One particularly interesting case is $M = \Lambda$ itself. In this case, $\Lambda_{\mathcal{S}}$ is a Gelfand–Tsetlin module such that $\mathcal{W}_{\lambda}(\Lambda_{\mathcal{S}}) = \Lambda_{\lambda}^{(1)}$ for all $\lambda \in \mathcal{S}$. The same module has been constructed by Mazorchuk and Vishnyakova [34, Theorem 4]. The dual version of this construction given by taking the vector space dual $\Lambda^* = \text{Hom}_{\mathbb{k}}(\Lambda, \mathbb{k})$ for some subfield $\mathbb{k}$ and considering $(\Lambda^*)^{\mathcal{S}}$ has been studied by several authors, including Early–Mazorchuk–Vishnyakova [10], Hartwig [22] and Futorny–Grantcharov–Ramirez–Zadunaisky [16]; in particular, it appears to the author that $e(\Lambda^*)^{\mathcal{S}}$ is precisely the $U = eFe$ module $V(\Omega, T(v))$ defined in [16, Definition 7.3] when $\mathcal{S} = \widehat{W} \cdot v$ and Ω is a base of the group $\widehat{W}_{\lambda}$ for any $\lambda \in \mathcal{S}$.

Based on the structure of this module, we can construct a “canonical” module as in [10, 22]; the author is not especially fond of this name as the embedding of F in $\mathcal{F}$ is not itself canonical if the algebra F is the object of interest. For example, $U(\mathfrak{gl}_n)$ has an embedding into $\mathcal{F}$ for each orientation of the linear quiver, each with its own notion of “canonical module.”

For every $\lambda \in \mathcal{S}$, we can consider the submodule C'_{λ} of $\Lambda_{\mathcal{S}}$ generated by $\mathcal{W}_{\lambda}(\Lambda_{\mathcal{S}})$ which is clearly finitely (in fact, cyclically) generated.

Lemma 2.24. *The submodule C'_{λ} has a unique simple quotient C_{λ} , and corresponds to the unique simple quotient of $\Lambda_{\lambda}^{(1)}$ as a $F_{\lambda}^{(1)}$ -module under Theorem 2.17.*

Proof. Given any proper submodule $M \subset C'_{\lambda}$, consider $M \cap \mathcal{W}_{\lambda}(\Lambda_{\mathcal{S}}) \subset \Lambda_{\lambda}^{(1)}$. This must be a proper submodule, because $\mathcal{W}_{\lambda}(\Lambda_{\mathcal{S}})$ generates C'_{λ} . As a $\Lambda_{\lambda}^{(1)}$ -module, $\Lambda_{\lambda}^{(1)}$ has a unique maximal submodule, the ideal $\mathfrak{m}_{\lambda}/\mathfrak{n}_{\lambda}$, which thus contains $M \cap \mathcal{W}_{\lambda}(\Lambda_{\mathcal{S}})$. Thus, the sum of two proper submodules has the same property and is again proper. This shows that there is a unique maximal proper submodule, and thus a unique simple quotient. $\square$

In the terminology of [22], the canonical module is the right module C_{λ}^* obtained by dualizing this construction with respect to a subfield $\mathbb{k}$. Since we avoid dualizing, our result here is both a bit stronger and a bit weaker than [22, Theorem 3.3]. That result does not depend on the finiteness of $\widehat{W}_{\lambda}$, though as a result, one pays the price of not knowing whether $\mathcal{W}_{\lambda}$ is finite-dimensional. However, our construction applies when Λ is arbitrary, making no assumption on characteristic or linearity over a field.

A natural question, first posed to us by Mazorchuk, is which simple GT modules appear as canonical modules. In particular, one could hope that each simple module is the canonical module of some maximal ideal; jointly with Silverthorne, we have shown that this is the case for OGZ algebras with their usual Galois order structure [41, Theorem A(1)]. This heavily uses the combinatorics of that special case and does not readily generalize to other cases.

2.6. Interaction between weight spaces. In this section, we continue to assume that every weight considered has finite stabilizer in $\widehat{W}$.

Whereas in the previous 3 subsections, we focused attention on a single weight, in this section we study how we can understand the classification of modules by considering how different weights interact. For now, fix two different weights $\lambda, \mu \in \text{MaxSpec}(\Lambda)$. Recall that ${}_{\lambda}\widehat{W}_{\mu}$ is the set of elements of $\widehat{W}$ such that $w \cdot \mu = \lambda$ and ${}_{\lambda}F_{\mu} = F \cap K \cdot {}_{\lambda}\widehat{W}_{\mu}$. This is a F_{λ} - F_{μ} -bimodule, and we have a multiplication ${}_{\lambda}F_{\mu} \otimes_{F_{\mu}} F_{\nu} \rightarrow {}_{\lambda}F_{\nu}$. Thus, we can define a matrix algebra:

$$F(\lambda_1, \dots, \lambda_k) = \begin{bmatrix} F_{\lambda_1} & {}_{\lambda_1}F_{\lambda_2} & \cdots & {}_{\lambda_1}F_{\lambda_k} \\ {}_{\lambda_2}F_{\lambda_1} & F_{\lambda_2} & \cdots & {}_{\lambda_2}F_{\lambda_k} \\ \vdots & \vdots & \ddots & \vdots \\ {}_{\lambda_k}F_{\lambda_1} & {}_{\lambda_k}F_{\lambda_2} & \cdots & F_{\lambda_k} \end{bmatrix} \quad (2.6)$$

More generally, for any subset $S \subset \text{MaxSpec}(\Lambda)$, we let $F(S)$ be the direct limit of this matrix algebra over all finite subsets. Note that if S is not finite, this is not a unital algebra, but is locally unital. This acts by natural transformations on the functor $\bigoplus_{\lambda \in S} \mathcal{W}_{\lambda}$.

Note that if λ and μ are not in the same orbit of $\widehat{W}$, then ${}_{\lambda}F_{\mu} = 0$, so $F(S)$ naturally breaks up as a direct sum over the different $\widehat{W}$ -orbits that these weights lie in.

If λ and μ are in the same orbit, then we have a canonical isomorphism $\Lambda_{\lambda} \cong \Lambda_{\mu}$ induced by any element of ${}_{\lambda}\widehat{W}_{\mu}$, which identifies the ideals $\mathfrak{n}_{\lambda}$ and $\mathfrak{n}_{\mu}$. For $\mathcal{S}$ a single $\widehat{W}$ -orbit, we can identify these with a single algebra $\mathcal{Z}(\mathcal{S}) \supset \mathfrak{n}$.

Proposition 2.25. *If $S \subset \mathcal{S}$, then $\mathcal{Z}(\mathcal{S})$ is the center of $F(S)$.*

Proof. As discussed before, we have an isomorphism $F_{\lambda} \otimes_{\Gamma} K \cong L \# \widehat{W}_{\lambda}$, and ${}_{\lambda_1}F_{\lambda_2} \otimes_{\Gamma} K$ is just the bimodule induced by an isomorphism between these algebras. Thus $F(S) \otimes K$ is Morita equivalent to $L \# \widehat{W}_{\lambda}$, and its center is the subfield $L^{\widehat{W}_{\lambda}} \subset L$. We have that $Z(F(\lambda_1, \dots, \lambda_k)) = F(S) \cap Z(L \# \widehat{W}_{\lambda}) = \mathcal{Z}(\mathcal{S})$. $\square$

Let

$$\begin{aligned} F^{(N)}(S) &= F(S)/\mathfrak{n}^N F(S) \\ \widehat{F}(S) &= F(S) \otimes_{\Lambda, \mathcal{S}} \widehat{\Lambda}_{\mathcal{S}}. \end{aligned}$$

As a consequence of [9, Theorem 17], we can easily extend Theorem 2.17 to incorporate any number of weight spaces. Since S might be infinite, the module $\bigoplus_{\lambda \in S} \mathcal{W}_{\lambda}(S)$ might not be finite-length as a module over Λ . We call a module M over $\widehat{F}(S)$ locally finite-length if for each idempotent $1_{\lambda} \in \widehat{F}_{\lambda}$, then image $1_{\lambda}M$ is finite-length.

Theorem 2.26. *The simple Gelfand–Tsetlin F -modules S such that $\mathcal{W}_{\lambda}(S) \neq 0$ for some $\lambda \in S$ are in bijection with locally finite-length simple modules over $F^{(1)}(S)$, sending $S \mapsto \bigoplus_{\lambda \in S} \mathcal{W}_{\lambda}(S)$.*

We can also extend this to an equivalence of categories: let $\Gamma\mathbb{I}(\mathbf{S})$ be the category of all Gelfand–Tsetlin modules modulo the subcategory of modules such that $\mathcal{W}_\lambda(M) = 0$ for all $\lambda \in \mathbf{S}$.

Theorem 2.27. *The functor $S \mapsto \bigoplus_{i=1}^k \mathcal{W}_{\lambda_i}(S)$ gives an equivalence between $\Gamma\mathbb{I}(\mathbf{S})$ and locally finite-length modules over the completion $\widehat{F}(\mathbf{S})$ which are continuous with respect to the discrete topology.*

By [12, Theorem 4.7], if we remove the assumption that stabilizers are finite (or consider modules which are not finitely generated), some care is needed about topologies; representations which are continuous in the discrete topology will correspond to *strong* Gelfand–Tsetlin modules, i.e. those where $\mathfrak{m}_\lambda^N \mathcal{W}_\lambda(M) = 0$ for $N \gg 0$. With our assumptions, Gelfand–Tsetlin modules are automatically strong.

As before, let $\mathcal{S}$ be a $\widehat{W}$ -orbit in $\text{MaxSpec}(\Lambda)$ and let $\Gamma\mathbb{I}(\mathcal{S})$ the category of Gelfand–Tsetlin modules where if $\lambda \notin \mathcal{S}$, we have $\mathcal{W}_\lambda(M) = 0$.

Definition 2.28. We call a set of weights $\mathbf{S} \subset \mathcal{S}$ *complete* for the orbit $\mathcal{S}$ if $\Gamma\mathbb{I}(\mathbf{S}) = \Gamma\mathbb{I}(\mathcal{S})$, that is, if any module M with $\mathcal{W}_{\lambda_i}(M) = 0$ for all i satisfies $\mathcal{W}_\lambda(M) = 0$ for all $\lambda \in \mathcal{S}$.

A finite set $\mathbf{S}$ is complete for the orbit $\mathcal{S}$, if and only if $\Gamma\mathbb{I}(\mathcal{S}) \cong \widehat{F}(\mathbf{S})\text{-fdmod}$.

Of course, many readers will be more interested in understanding modules of the original principal Galois order. For simplicity, assume that $\mathbf{S}$ only contains at most one element of each W -orbit. We can derive the weight spaces of U from those of F by taking invariants under the stabilizer W_λ . Let e_λ be the idempotent in $\widehat{F}_\lambda$ which projects to the invariants of W_λ , and $e_\lambda \in \widehat{F}(\mathbf{S})$ the matrix with these as diagonal entries for the different $\lambda \in \mathbf{S}$. Let $U^{(1)}(\mathbf{S}) = e_\lambda F^{(1)}(\mathbf{S}) e_\lambda$.

Theorem 2.29. *The simple Gelfand–Tsetlin U -modules S such that $\mathcal{W}_\gamma(S) \neq 0$ for γ in the image of $\mathbf{S}$ are in bijection with simple modules over $U^{(1)}(\mathbf{S})$, sending $S \mapsto \bigoplus_{\lambda \in \mathbf{S}} e_\lambda \mathcal{W}_\lambda(S)$.*

3. THE REFLECTION CASE

In Section 2, we worked in the same generality as in [22]. In this section, we wish to specialize to a much simpler case. Let V be a $\mathbb{C}$ -vector space with an action of a complex reflection group W , and $\mathcal{M}$ a finitely generated (over $\mathbb{Z}$) subgroup of V^* . We assume from now on that $\Lambda = \text{Sym}^\bullet(V)$ is the symmetric algebra on this vector space, with the obvious induced $\mathcal{M}$ -action. Note that the stabilizer $\widehat{W}_\lambda$ for any $\lambda \in V^*$ is finite, and in fact a subgroup of W via the usual quotient map $\widehat{W} \rightarrow W$. It is generated by the $\mathcal{M}$ -translates of root hyperplanes containing λ , and thus is again a complex reflection group, acting by the translation of a linear action.

This simplifies matters in one key way: the module Λ is a free Frobenius extension over Λ_λ and over Γ . Recall that we call a ring extension $A \subset B$ *free Frobenius* if B is a free A -module, and $\text{Hom}_A(B, A)$ is a free B module of rank 1 for its induced left B -action or right B -action; a *Frobenius trace* is a generator of $\text{Hom}_A(B, A)$ as a B -module (again, as a left module or a right module).

The fact that Λ is free Frobenius over Γ is well-known, and easily derived from results in [5]: following the notation of loc. cit., we have a map $\Lambda \rightarrow \Gamma$ defined by $D(J^*)$, which is the desired trace. In slightly more down-to-earth terms, we have a unique element $J \in \Lambda$ of minimal degree that transforms under the determinant character of the action

on V^* ; this is obtained by taking a suitable power of the linear form defining each root hyperplane. There is a unique homogeneous Frobenius trace up to scalar multiplication, which is characterized by sending this element to $1 \in \Gamma$ and killing all other isotypic components for the action of W .

In particular, this means that $D = \text{End}_\Gamma(\Lambda)$, the *nilHecke algebra* of W , is Morita equivalent to Γ ; see for example [20, Lemma 7.1.5].

Definition 3.1. We call a flag order F *Morita* if the symmetrization idempotent gives a Morita equivalence between $U = eFe$ and F ; that is, if $F = FeF$.

Recall that for a fixed principal Galois order U , we have an associated flag Galois order F_D . Since $D = DeD$ when $D = \text{End}_\Gamma(\Lambda)$ in the complex reflection case, we find that the flag order F_D is Morita for any principal Galois order in this case.

Thus, for any principal Galois order, we can study the representation theory of the corresponding flag order instead. This approach is implicit in much recent work on this subject, which uses the nilHecke algebra, such as [13, 15, 38], but many issues are considerably simplified if we think of the flag order as the basic object.

It is easy to see how Gelfand–Tsetlin modules behave under this equivalence. We can strengthen Lemma 2.8 to:

Lemma 3.2. *If F is Morita, then $\mathcal{W}_\lambda(M)$ is free as a $\mathbb{C}W_\lambda$ -module and we have isomorphisms*

$$\mathcal{W}_\gamma(eM) \cong \mathcal{W}_\lambda(M)^{W_\lambda} \quad \mathcal{W}_\lambda(M) \cong (\mathcal{W}_\gamma(eM))^{\oplus \#W_\lambda}.$$

The reflection hypothesis also allows us to define a dual version of the canonical module C_λ . We can consider the quotient $\tilde{C}'_\lambda$ of the module $\Lambda_{\mathcal{S}}$ by all submodules having trivial intersection with $\mathcal{W}_\lambda(\Lambda_{\mathcal{S}})$.

The algebra $\Lambda_\lambda^{(1)}$ is a Frobenius algebra, so its socle as a $\Lambda_\lambda^{(1)}$ -module is 1-dimensional, and every non-zero submodule of $\tilde{C}'_\lambda$ has nontrivial intersection with $\mathcal{W}_\lambda(\Lambda_{\mathcal{S}})$, and thus contains this socle. This shows that the intersection of all non-zero submodules is non-trivial, giving a simple socle $\tilde{C}_\lambda \subset \tilde{C}'_\lambda$. This will sometimes be isomorphic to C_λ , and sometimes not.

3.1. Special cases of interest.

Definition 3.3. We call a weight λ *non-singular* if $\widehat{W}_\lambda = \{1\}$ and more generally *p-singular* if $\widehat{W}_\lambda$ has a minimal generating set of p reflections.

In this case, we have an equality $\widehat{F}_\lambda = \widehat{\Lambda}_\lambda$, which is a complete local ring, and thus has a single simple module $\widehat{\Lambda}_\lambda/\mathfrak{m}_\lambda$. Theorem 2.17 shows that:

Corollary 3.4. *If λ is non-singular, there is a unique simple Gelfand–Tsetlin module S with $\mathcal{W}_\lambda(S) \cong \mathbb{C}$ and for all other simples S' we have $\mathcal{W}_\lambda(S') = 0$.*

A natural question to consider is when two non-singular weights λ, μ have the same simple, and when they do not. Of course, they can only give the same simple if $\mu = w \cdot \lambda$ for some $w \in \widehat{W}$.

Corollary 3.5. *Given non-singular weights λ and μ as above, we have a simple Gelfand–Tsetlin module S with $\mathcal{W}_\lambda(S) \cong \mathcal{W}_\mu(S) \cong \mathbb{C}$ if and only if ${}_\lambda F_\mu \cdot {}_\mu F_\lambda \notin \mathfrak{m}_\lambda$.*

Now assume λ is 1-singular and F_λ is a free module over Λ . In this case, $\widehat{W}_\lambda \cong S_2$, so $F_\lambda^{(1)}$ is 4-dimensional. Thus, there are 3 possibilities for the behavior of such a weight:

Corollary 3.6. *Exactly 1 of the following holds:*

- (1) $F_\lambda^{(1)} \cong M_2(\mathbb{C})$ and there is a unique simple Gelfand–Tsetlin module S with $\mathcal{W}_\lambda(S) \cong \mathbb{C}^2$ and for all other simples it is 0.
- (2) The Jacobson radical of $F_\lambda^{(1)}$ is 2-dimensional and there are two simple Gelfand–Tsetlin modules S_1, S_2 with $\mathcal{W}_\lambda(S_i) \cong \mathbb{C}$ and for all other simples it is 0.
- (3) The Jacobson radical of $F_\lambda^{(1)}$ is 3-dimensional and there is a unique simple Gelfand–Tsetlin module S with $\mathcal{W}_\lambda(S) \cong \mathbb{C}$ and for all other simples it is 0.

4. COULOMB BRANCHES

Throughout this section, we fix a field $\mathbb{k}$, and all (co)homology will be calculated with coefficients in this field. For now, $\mathbb{k}$ can have any characteristic not dividing $\#W$, but for most of the sequel, we will assume that $\mathbb{k}$ is characteristic 0.

4.1. Coulomb branches and principal orders. One extremely interesting collection of examples of principal Galois orders are the Coulomb branches defined by Braverman, Finkelberg, and Nakajima [3]. These algebras have attracted considerable interest in recent years, and subsume most examples of interesting principal Galois orders known to the author.

There is a Coulomb branch attached to each connected reductive complex group⁵ G and representation N . Let $G[[t]]$ be the Taylor series points of the group G , and $G((t))$ its Laurent series points. Let

$$\mathcal{Y} = (G((t)) \times N[[t]])/G[[t]],$$

equipped with its obvious map $\pi: \mathcal{Y} \rightarrow N((t))$; we can think of this as a vector bundle over the affine Grassmannian $G((t))/G[[t]]$. Readers who prefer moduli theoretic interpretations can think of this as the moduli space of principal bundles on a formal disk with choice of section of the associated bundle for N and of trivialization away from the origin.

Let $H = N_{GL(N)}(G)^\circ$ be the connected component of the identity in the normalizer of G , and let Q be a group equipped with an inclusion $G \hookrightarrow Q$ with Q/G a torus, and a compatible map $Q \rightarrow H$. The choice we will want to make most often is to assume that this map induces an isomorphism of Q/G to a maximal torus of H/G , but it can be useful to have the freedom to make a different choice. Given a maximal torus T_Q of Q , its intersection with G gives a maximal torus T of G . Note that $\mathcal{Y}$ has a Q -action via $q \cdot (g(t), n(t)) = (qg(t)q^{-1}, qn(t))$. It also carries a canonical principal Q -bundle $\mathcal{Y}_Q$ given by the quotient of $G((t)) \times Q \times N[[t]]$ via the action $g(t) \cdot (g'(t), q, n(t)) = (g'(t)g^{-1}(t), qg^{-1}(0), g(t)n(t))$. We can extend this to an action of $Q \times \mathbb{C}^*$ where the factor of $\mathbb{C}^*$ acts by the loop scaling, and let $\hbar$ denote the equivariant parameter of the loop scaling.

Definition 4.1. The (quantum) Coulomb branch is the convolution algebra

$$\mathcal{A} = H_*^{Q \times \mathbb{C}^*} \left(\pi^{-1}(N[[t]]) \right),$$

It might not be readily apparent what the algebra structure on this space is. However, it is uniquely determined by the fact that it acts on $H_*^{Q \times \mathbb{C}^*}(N[[t]]) = H_{Q \times \mathbb{C}^*}^*(*)$ by

$$a \star b = \pi_*(a \cap \iota(b)) \tag{4.1}$$

⁵Note that in most previous work on Galois orders such as [16, 17, 22], G has denoted the finite group which we denote W ; since in all cases of interest to us, W is the Weyl group of a reductive group, and as discussed below, this is the context where we find it, we feel this switch in notation is justified.

where ι is the inclusion of this algebra into $\mathcal{A}$ as the Chern classes of the principal bundle $\mathcal{Y}_Q$ and the obvious inclusion of $\mathbb{C}[\hbar] \cong H_*^{\mathbb{C}^*}(N[[t]])$. Obviously, there are a lot of technical issues that are being swept under the rug here; a reader concerned about this point should refer to [3] for more details.

Let $J = Q/G$, and $\mathfrak{j}$ the Lie algebra of this group. The subalgebra $H_{J \times \mathbb{C}^*}^*(*) = \text{Sym}(\mathfrak{j}^*)[\hbar] \subset \mathcal{A}$ induced by the $Q \times \mathbb{C}^*$ -action is central; borrowing terminology from physics, we call these *flavor parameters*. Thus, we can consider the quotient of $\mathcal{A}$ by a maximal ideal in this ring. This quotient is what is called the “Coulomb branch” in [3, Definition 3.13] and our Definition 4.1 matches the deformation constructed in [3, § 3 (viii)].

We let W be the Weyl group of G (which is also the Weyl group of Q), let $V = \mathfrak{t}_Q^* \oplus \mathbb{C} \cdot \hbar$ where $\mathfrak{t}_Q$ is the (abstract) Cartan Lie algebra of Q and let $\mathcal{M}$ be the cocharacter lattice of T_G , acting by the $\hbar$ -scaled translations

$$\chi \cdot (\nu + k\hbar) = \nu + k\langle \chi, \nu \rangle + k\hbar.$$

Note that the action has finite stabilizers on any point where $\hbar \neq 0$ if $\mathbb{k}$ has characteristic 0, but any point with $\hbar = 0$ will have infinite stabilizer. We’ll ultimately only be interested in modules over the specialization $\hbar = 1$, so this will not cause an issue for the moment. Note that

$$\Lambda \cong H_{T_Q \times \mathbb{C}^*}^*(*) = \text{Sym}^\bullet(\mathfrak{t}_Q)[\hbar] \quad \Gamma \cong H_{Q \times \mathbb{C}^*}^*(*) = \text{Sym}^\bullet(\mathfrak{t}_Q)^W[\hbar],$$

and $W \ltimes \mathcal{M}$ is the extended affine Weyl group of G . Localization in equivariant cohomology shows that the action of (4.1) induces an inclusion $\mathcal{A} \hookrightarrow \mathcal{K}_\Gamma$ for the data above; see [3, (5.18) & Proposition 5.19]. Thus, it immediately follows that:

Proposition 4.2. *The Coulomb branch is a principal Galois order for these data.*

If we fix the flavor parameters, the result will also be a principal Galois order for an appropriate quotient of Λ .

The flag order attached to these data also has an interpretation as the flag BFN algebra from [47, Definition 3.2]. Let $\mathcal{X} = (G((t)) \times N[[t]])/I$, where I is the standard Iwahori, $\pi_{\mathcal{X}}: \mathcal{X} \rightarrow N((t))$ the obvious map and ${}_0\mathcal{X}_0 = \pi_{\mathcal{X}}^{-1}(N[[t]])$.

Definition 4.3. The *Iwahori Coulomb branch* is the convolution algebra

$$F = H_*^{T_Q \times \mathbb{C}^*}({}_0\mathcal{X}_0).$$

This is the Morita flag order F_D associated to $\mathcal{A}$ with $D = \text{End}_\Gamma(\Lambda)$ the nilHecke algebra of W , as is shown in [47, Theorem 3.3].

As mentioned before, we wish to consider the specializations of these algebras where $\hbar = 1$. These are again principal/flag Galois orders in their own right, but are harder to interpret geometrically. Note that by homogeneity, the specializations of this algebra at all different nonzero values of $\hbar$ are isomorphic. The specialization $\hbar = 0$ is quite different in nature, since in this case, the action of $\mathcal{M}$ is trivial.

4.2. Representations of Coulomb branches. From now on, we assume that $\mathbb{k}$ has characteristic 0. For a Coulomb branch, the algebra $F_\lambda^{(1)}$ has a geometric interpretation. Since we assume that $\hbar = 1$, when we interpret λ as an element of the Lie algebra $\mathfrak{t}_Q \oplus \mathbb{C}$, the second component is 1. Let G_λ (resp. Q_λ) be the Levi subgroup of G (resp. Q) which only contains the roots which are integral at λ , and N_λ the span of the weight spaces for weights integral on λ . Let B_λ be the Borel in G_λ such that $\text{Lie}(B_\lambda)$ is generated by the roots α such that $\langle \lambda, \alpha \rangle$ is negative and those in the fixed Borel $\mathfrak{b}_G$ such that $\langle \lambda, \alpha \rangle = 0$.

The element λ integrates to a character acting on N_λ . Let N_λ^- be the subspace of N_λ which is non-positive for the cocharacter corresponding to λ ; this subspace is preserved by the action of B_λ . Consider the associated vector bundle $X_\lambda = (G_\lambda \times N_\lambda^-)/B_\lambda$ and p_λ the associated map $p: X_\lambda \rightarrow N_\lambda$. If $W_\lambda \neq \{1\}$, then there is also a parabolic version of these spaces. Let $P_\lambda \subset G_\lambda$ be the parabolic corresponding to W_λ , and let $Y_\lambda = (G_\lambda \times N_\lambda^-)/P_\lambda$.

As usual, we have associated Steinberg varieties:

$$\begin{aligned}\mathbb{X}_\lambda &= X_\lambda \times_{N_\lambda} X_\lambda = \left\{ (g_1 B_\lambda, g_2 B_\lambda, n) \mid n \in g_1 N_\lambda^- \cap g_2 N_\lambda^- \right\} \\ {}_\lambda \mathbb{X}_\mu &= X_\lambda \times_{N_\lambda} X_\mu = \left\{ (g_1 B_\lambda, g_2 B_\mu, n) \mid n \in g_1 N_\lambda^- \cap g_2 N_\mu^- \right\} \\ \mathbb{Y}_\lambda &= Y_\lambda \times_{N_\lambda} Y_\lambda = \left\{ (g_1 P_\lambda, g_2 P_\lambda, n) \mid n \in g_1 N_\lambda^- \cap g_2 N_\lambda^- \right\} \\ {}_\lambda \mathbb{Y}_\mu &= Y_\lambda \times_{N_\lambda} Y_\mu = \left\{ (g_1 P_\lambda, g_2 P_\mu, n) \mid n \in g_1 N_\lambda^- \cap g_2 N_\mu^- \right\}\end{aligned}$$

Recall that the *Borel–Moore homology* of an algebraic variety X over $\mathbb{C}$ is the hypercohomology of the dualizing sheaf $\mathbb{D}k_X$ indexed backwards; as usual, this pushforward needs to be computed for the classical topological space X_{an} rather than in the Zariski topology. We use the same convention for equivariant Borel–Moore homology:

$$H_i^{BM}(X) = \mathbb{H}^{-i}(X_{\text{an}}; \mathbb{D}k_X) \quad H_i^{BM,G}(X) = \mathbb{H}_G^{-i}(X_{\text{an}}; \mathbb{D}k_X). \quad (4.2)$$

Note that this convention makes $H_*^{BM,G}(X)$ into a module over $H_G^*(X)$ which is homogeneous when this ring is given the negative of its usual homological grading; similarly, the group $H_i^{BM,G}(X)$ must be 0 if $i > \dim_{\mathbb{R}} X$, but this can be non-zero in infinitely many negative degrees. We let $\hat{H}_*^{BM,G_\lambda}(X)$ denote the completion of G_λ -equivariant Borel–Moore homology with respect to its grading, with all elements of degree $\leq k$ being a neighborhood of the identity for all k .

The Borel–Moore homology $H_*^{BM}(\mathbb{X}_\lambda)$ has a convolution algebra structure and $H_*^{BM}({}_\lambda \mathbb{X}_\mu)$ a bimodule structure defined by [8, (2.7.9)].

Theorem 4.4. *Keeping the assumption that $\mathbb{k}$ has characteristic 0, we have isomorphisms of algebras and bimodules*

$$F_\lambda^{(1)} \cong H_*^{BM}(\mathbb{X}_\lambda) \quad {}_\lambda F_\mu^{(1)} \cong H_*^{BM}({}_\lambda \mathbb{X}_\mu) \quad (4.3)$$

$$\hat{F}_\lambda \cong \hat{H}_*^{BM,Q_\lambda}(\mathbb{X}_\lambda) \quad {}_\lambda \hat{F}_\mu \cong \hat{H}_*^{BM,Q_\lambda}({}_\lambda \mathbb{X}_\mu) \quad (4.4)$$

$$U_\lambda^{(1)} \cong H_*^{BM}(\mathbb{Y}_\lambda) \quad {}_\lambda U_\mu^{(1)} \cong H_*^{BM}({}_\lambda \mathbb{Y}_\mu) \quad (4.5)$$

$$\hat{U}_\lambda \cong \hat{H}_*^{BM,Q_\lambda}(\mathbb{Y}_\lambda) \quad {}_\lambda \hat{U}_\mu \cong \hat{H}_*^{BM,Q_\lambda}({}_\lambda \mathbb{Y}_\mu) \quad (4.6)$$

If we specialize F by fixing the flavor parameters, then the same result holds with Q_λ , replaced by G_λ .

This theorem is a consequence of [47, Theorem 4.3], which is proven purely algebraically. H. Nakajima has also communicated a more direct geometric proof to the author, based on the earlier work of Varagnolo–Vasserot [42, § 2]. We will include a sketch of that argument here, but there are some slightly subtle points about infinite-dimensional topology which we will skip over.

Proof (sketch). Note first how the left and right actions of Λ on F operate. The left action is simply induced by the equivariant cohomology of a point, whereas the right action is induced by the Chern classes of tautological bundles on $G((t))/I$.

Consider the 1-parameter subgroup $\mathbb{T}$ of $G \times \mathbb{C}^*$ obtained by exponentiating λ . By the localization theorem in equivariant cohomology, the completion $\varinjlim F/\mathfrak{n}_\lambda^k F$ is isomorphic to the completion of the T_Q -equivariant Borel–Moore homology of ${}_0\mathcal{X}_0^\mathbb{T}$ with respect to the usual grading. This is easily seen from [21, (6.2)(1)]: the T_Q -equivariant Borel–Moore homology of the complement of the fixed points is a torsion module whose support avoids λ , since the action of $\mathbb{T}$ is locally free. Thus, after completion, the long exact sequence in Borel–Moore homology gives the desired result. Note that here we also use the fact that since the action of $\mathbb{T}$ on the fixed points is trivial, the completion at any point in $\mathfrak{t}$ gives the same result.

First, note that the fixed points $N[[t]]^\mathbb{T}$ are isomorphic to N_λ^- via the map $\tau_\lambda: N_\lambda \rightarrow N((t))$ sending an element n of weight $-a$ in N_λ to $t^a n$.

We can also apply this to the adjoint representation and find the fixed points of the 1-parameter subgroup on $\mathfrak{g}((t))$; this is a copy of $\mathfrak{g}_\lambda$, embedded according to the description above. Accordingly, the centralizer of this 1-parameter subgroup in $G((t))$ is a copy of G_λ generated by the roots SL_2 ’s of the roots $t^{-\langle \lambda, \alpha \rangle} \alpha$. The Borel B_λ is the intersection of this copy of G_λ with the Iwahori I .

Now consider the fixed points of $\mathbb{T}$ in $G((t))/I$. Each component of this space is a G_λ -orbit and these components are in bijection with elements of the orbit $\widehat{W} \cdot \lambda$; that is, wI and $w'I$ are in the same orbit if and only if $w \cdot \lambda = w' \cdot \lambda$. If w is of minimal length with $\mu = w \cdot \lambda$, the stabilizer of wI under the action of G_λ is the Borel B_μ . Considering the vector bundles induced by the tautological bundles shows that elements of $\mathfrak{n}_\mu$ act by elements with trivial degree 0 term, i.e. that the homology of this component is ${}_\lambda \widehat{F}_\mu$.

Thus, the fixed points $\mathcal{X}^\mathbb{T}$ break into components corresponding to these orbits as well, with the fiber over gwI for $g \in G_\lambda$ and w as defined above is given by $g\mathfrak{n}_\mu^-$, via the map $g \cdot \tau_\mu$. The map $\pi_\mathcal{X}$ maps this to $N((t))$ via the map $\tau_\lambda \circ \tau_\mu^{-1} \circ g^{-1}$, so its intersection with the preimage of $N[[t]]$ is $N_\lambda^- \cap g\mathfrak{n}_\mu^-$.

The relevant T_Q -equivariant homology group is thus

$$H_*^{T_Q} \left(\left\{ (gB_\mu, x) \mid g \in G_\mu, x \in N_\lambda^- \cap g\mathfrak{n}_\mu^- \right\} \right) \cong H_*^{Q_\lambda}({}_\lambda \mathbb{X}_\mu).$$

Taking quotient by $\mathfrak{n}_\lambda$, we obtain the non-equivariant Borel–Moore homology of this variety as desired. This shows that we have a vector space isomorphism in (4.3).

The row of isomorphisms (4.5) follows from the same argument applied to $\pi^{-1}(N[[[t]]])$ and the affine Grassmannian.

Note that we have not checked that the resulting isomorphism is compatible with multiplication, and doing so is somewhat subtle. For a finite-dimensional manifold X , we have two isomorphisms between $H_*^\mathbb{T}(X)$ and $H_*^\mathbb{T}(X^\mathbb{T})$ after completion at any non-zero point in $\mathfrak{t}$: pullback (defined using Poincaré duality) and pushforward, which differ by the (invertible) Euler class of the normal bundle by the adjunction formula. To obtain an isomorphism $H_*^\mathbb{T}(X \times X)$ and $H_*^\mathbb{T}(X^\mathbb{T} \times X^\mathbb{T})$ that commutes with convolution, one must take the middle road between these, using pullback times the inverse of the Euler class of the normal bundle along the first factor, which is the same as the inverse of pushforward times the Euler class of the normal bundle along the second factor (effectively, we use the pushforward isomorphism in the first factor and the pullback in the second factor). Due to the infinite dimensionality of the factors $\mathcal{X}$ and $\mathcal{Y}$, and the nature of the cycles we use, neither the pushforward nor the pullback isomorphisms make sense, but this intermediate isomorphism does.

As we said above, we will not give a detailed account of this isomorphism, since we have already constructed a ring isomorphism using the algebraic arguments of [47]. That “half” of the Euler class we need to invert should match with [47, (4.3a)]. $\square$

Remark 4.5. This theorem can be modified to work in characteristic p , but with a rather different variety than $\mathbb{X}_\lambda$. Since the stabilizer of λ is the extended affine Weyl group of a Levi subgroup in this case, the algebra $\widehat{F}_\lambda$ is again a principal flag order for an affine Coxeter group, and is actually either a Coulomb branch itself or a close relative. We develop this theory in [44].

The stabilizer $\widehat{W}_\lambda$ is always isomorphic to a parabolic subgroup of the original Weyl group W .

Definition 4.6. We call an orbit *integral* if $\widehat{W}_\lambda \cong W$ and $N = N_\lambda$.

One especially satisfying consequence of Theorem 4.4 is that the category of modules with weights in the non-integral orbit is equivalent to the same category for an integral orbit but of the Coulomb branch for the corresponding Levi subgroup G_λ and subrepresentation N_λ .

More precisely, fix an orbit $\mathcal{S}$ of $\widehat{W}$, and let $G' = G_\lambda$ and $N' = N_\lambda$ for arbitrary $\lambda \in \mathcal{S}$. Let $\mathcal{S}' \subset \mathcal{S}$ be an orbit of the subgroup $\widehat{W}' \subset \widehat{W}$ generated by the Weyl group of G' and the subgroup $\mathcal{M}$. Let $\Gamma\Pi'(\mathcal{S}')$ be the category of weight modules with all weights concentrated in the set $\mathcal{S}'$ for the Coulomb branch of (G', N') . Note that since all the different orbits $\mathcal{S}' \subset \mathcal{S}$ are conjugate under the action of W , this category only depends on $\mathcal{S}$. Of course, for this smaller group, $\mathcal{S}'$ is an integral orbit. By Theorem 4.4, we have that:

Corollary 4.7. *We have an equivalence of categories $\Gamma\Pi(\mathcal{S}) \cong \Gamma\Pi'(\mathcal{S}')$.*

This equivalence does not change the underlying vector space and its weight space decomposition; it simply multiplies the action of elements of F by elements of the appropriate completion of Γ to adjust the relations. This can be proven in the spirit of Theorem 4.4 by presenting the Coulomb branch of (G_λ, N_λ) as the homology of the fixed points of the torus action, and noting that the Euler class of the normal bundle acts invertibly on all the modules in the relevant subcategory.

4.3. Gradings. In this section, we’ll assume for simplicity that we are in the integral case. This is a particularly nice description since the convolution algebras in question are graded, and a simple geometric argument shows that they are graded free over the subalgebra $\Lambda_\lambda^{(1)}$, with the degrees of the generators read off from the dimensions of the preimages of the orbits in $\mathbb{X}_\lambda$. For reasons of Poincaré duality, we grade $H_*^{BM}(\mathbb{X}_\lambda)$ so that a cycle of dimension d has degree $\dim X_\lambda - d$, and $H_*^{BM}(\lambda\mathbb{X}_\mu)$ so that a cycle of dimension d has degree $\frac{\dim X_\lambda + \dim X_\mu}{2} - d$. This is homogeneous by [8, (2.7.9)].

Note that since we have reversed the homological grading again, we’ve effectively gotten rid of the minus sign in (4.2), and now cohomology will act homogeneously with its usual grading rather than its negative. In particular, $H_*^{BM, Q_\lambda}(\lambda\mathbb{X}_\mu)$ will be a homogeneous module over $H_{Q_\lambda}^*(pt)$ in the usual grading.

Proposition 4.8. $F_\lambda^{(1)}$ has a set of free generators with degrees given by $\dim(N_\lambda^-) - \dim(wN_\lambda^- \cap N_\lambda^-) - \ell(w)$ ranging over $w \in \widehat{W}_\lambda$, identified with the Weyl group of G_λ .

Proof. The product $(G_\lambda/B_\lambda)^2$ breaks up into finitely many G_λ -orbits, each one of which contains the pair of cosets (B_λ, wB_λ) for a unique $w \in \widehat{W}_\lambda$. This orbit is isomorphic to an affine bundle over G_λ/B_λ with fiber $B_\lambda/(B_\lambda \cap wB_\lambda w^{-1})$, which is an affine space of dimension $\ell(w)$. Furthermore, the preimage of this orbit in $\mathbb{X}_\lambda$ is a vector bundle of dimension $\dim(wN_\lambda^- \cap N_\lambda^-)$. This means that, under the usual grading on the convolution algebra, the fundamental class has degree equal to $\dim X_\lambda$ minus the dimension of this orbit. These fundamental classes give free generators over $\Lambda_\lambda^{(1)}$, since the homology of each of these vector bundles is free of rank 1. $\square$

In particular, if these degrees are always non-negative, then all elements of positive degree are in the Jacobson radical.

Corollary 4.9. *If $\dim(N_\lambda^-) - \dim(wN_\lambda^- \cap N_\lambda^-) - \ell(w) \geq 0$ for all $w \in \widehat{W}_\lambda$, then the sum of $(\dim \mathcal{W}_\lambda(S))^2$ over all simple Gelfand–Tsetlin modules is*

$$\leq \# \left\{ w \in \widehat{W}_\lambda \mid \dim(N_\lambda^-) - \dim(wN_\lambda^- \cap N_\lambda^-) = \ell(w) \right\}.$$

Note that the fact that the algebra $F^{(1)}(\mathbf{S})$ is graded allows us to define a graded lift $\widehat{\Gamma\mathbb{I}}$ of the category of Gelfand–Tsetlin modules by considering graded modules over $F^{(1)}(\lambda_1, \dots, \lambda_k)$.

Following Ginzburg and Chriss [8, 8.6.7], we can restate Theorem 4.4 as

$$\begin{aligned} F_\lambda^{(1)} &\cong \operatorname{Ext}^\bullet((p_\lambda)_* \mathbb{k}_{X_\lambda}, (p_\lambda)_* \mathbb{k}_{X_\lambda}) \\ F^{(1)}(\mathbf{S}) &\cong \operatorname{Ext}^\bullet\left(\bigoplus_{i=1}^k (p_{\lambda_i})_* \mathbb{k}_{X_{\lambda_i}}, \bigoplus_{i=1}^k (p_{\lambda_i})_* \mathbb{k}_{X_{\lambda_i}}\right) \end{aligned} \quad (4.7)$$

The geometric description of (4.7) has an important combinatorial consequence when combined with the Decomposition Theorem of Beilinson–Bernstein–Deligne–Gabber [8, Theorem 8.4.8]:

Theorem 4.10. *The simple Gelfand–Tsetlin modules S such that $\mathcal{W}_{\lambda_i}(S) \neq 0$ for some i are in bijection with simple perverse sheaves $\operatorname{IC}(Y, \chi)$ appearing as summands up to shift of $\oplus_i (p_{\lambda_i})_* \mathbb{k}_{X_{\lambda_i}}$, with the dimension of $\mathcal{W}_{\lambda_i}(S)$ being the multiplicity of all shifts of $\operatorname{IC}(Y, \chi)$ in $(p_{\lambda_i})_* \mathbb{k}_{X_{\lambda_i}}$.*

Note that this result is implicit in [8, § 8.7] and [40, p. 9] but unfortunately is not stated clearly in either source.

Proof. By the Decomposition Theorem, $(p_\lambda)_* \mathbb{k}_{X_\lambda}$ is a direct sum of shifts of simple perverse sheaves. In the notation of [8, Theorem 8.4.8], we have

$$(p_\lambda)_* \mathbb{k}_{X_\lambda} \cong \bigoplus_{(i, Y, \chi)} L_{Y, \chi}(i, \lambda) \otimes \operatorname{IC}(Y, \chi)[i].$$

Let $L_{Y, \chi} \cong \oplus_{i, \lambda_j} L_{Y, \chi}(i, \lambda_j)$ be the $\mathbb{Z}$ -graded vector space obtained by summing the multiplicity spaces. Let

$$A = \operatorname{Ext}^\bullet\left(\bigoplus_{L_{Y, \chi} \neq 0} \operatorname{IC}(Y, \chi)\right) \quad B = \operatorname{Ext}^\bullet\left(\bigoplus_j (p_{\lambda_j})_* \mathbb{k}_{X_{\lambda_j}}, \bigoplus_{L_{Y, \chi} \neq 0} \operatorname{IC}(Y, \chi)\right).$$

By [8, Corollary 8.4.4], this algebra A is a positively graded basic algebra with irreps indexed by pairs (Y, χ) such that $L_{Y, \chi} \neq 0$, and B is an A - $F^{(1)}(\mathbf{S})$. The bimodule B induces

a graded Morita equivalence between A and $F^{(1)}(\mathbf{S})$. Thus, the simple representations of $F^{(1)}(\mathbf{S})$ are the images of these 1-dimensional irreps under the Morita equivalence, that is, the multiplicity spaces $L_{Y,\chi}$, with the dimension of the different weight spaces given by $\dim L_{Y,\chi}(*, \lambda)$, the multiplicity of all shifts of $\mathrm{IC}(Y, \chi)$ in $(p_\lambda)_* \mathbb{k}_{X_\lambda}$. $\square$

The additive category of perverse sheaves given by sums of shifts of summands of $(p_{\lambda_i})_* \mathbb{k}_{X_\lambda}$ satisfies the hypotheses of [45, Lemma 1.18], and so by [45, Lemma 1.13 & Corollary 2.4], we have that (as proven in [47, Def. 4.7]):

Theorem 4.11. *The classes of the simple Gelfand–Tsetlin modules form a dual canonical basis (in the sense of [45, § 2]) in the Grothendieck group of $\widetilde{\Gamma\Pi}$.*

We only truly need the Decomposition theorem to prove a single purely algebraic, but extremely non-trivial fact:

Corollary 4.12. *The graded algebra $F^{(1)}(\mathbf{S})$ is graded Morita equivalent to an algebra which is non-negatively graded and semi-simple in degree 0.*

This property is called “mixedness” in [1, 45]; the celebrated recent work of Elias and Williamson [11] gives an algebraic proof of this fact in some related contexts and could possibly be applied here as well.

4.4. Applications. As before, this description is particularly useful in the 1-singular case. In this case, we must have $G_\lambda/B_\lambda \cong \mathbb{P}^1$.

Corollary 4.13. *For a 1-singular weight, the different possibilities of Corollary 3.6 hold when:*

- (1) *There is exactly one simple Gelfand–Tsetlin module S with $\mathcal{W}_\lambda(S) \neq 0$ and this space is 2-dimensional iff $N_\lambda^- = sN_\lambda^-$.*
- (2) *There are exactly two simple Gelfand–Tsetlin modules S_1, S_2 with $\mathcal{W}_\lambda(S_i) \neq 0$ and for both it is 1-dimensional iff $N_\lambda^- \cap sN_\lambda^-$ is codimension 1 in N_λ^- .*
- (3) *there is exactly one simple Gelfand–Tsetlin module S with $\mathcal{W}_\lambda(S) \neq 0$ and this space is 1-dimensional in all other cases.*

Geometrically, these correspond to the situations where the map $X_\lambda \rightarrow G_\lambda \cdot N_\lambda^-$ is (1) the projection $X_\lambda = \mathbb{P}^1 \times N_\lambda^- \rightarrow N_\lambda^-$, (2) strictly semismall, or (3) small.

Of course, in the non-singular case, there is no difficulty in classifying the simple modules where a given weight appears: There is always a unique one. However, it is still an interesting question when these simples are the same for 2 different weights. Note that if λ, μ are in the same orbit of $\widehat{W}$, then $N_\lambda = \mathfrak{n}_\mu$, but the positive subspaces are not necessarily equal.

Corollary 4.14. *Assume that λ, μ are non-singular and in the same $\widehat{W}$ -orbit. Then there is a simple Gelfand–Tsetlin module with $\mathcal{W}_\lambda(S)$ and $\mathcal{W}_\mu(S)$ both non-zero if and only if $N_\lambda^- = N_\mu^-$.*

Outside the nonsingular case, we can still usefully compare weights. We can define an equivalence relation on weights such that λ, μ are equivalent if: For some $\mu' = w\mu$ with $w \in W$, we have $\lambda - \mu' \in \mathfrak{t}_\mathbb{Z}$, and for some $g \in G$, we have $B_\lambda = gB_{\mu'}g^{-1}$ and $N_\lambda^- = gN_{\mu'}^-$. We call the equivalence classes of this relation *clans*.

Lemma 4.15. *If λ, μ are in the same clan, then the weight spaces $\mathcal{W}_\lambda(M)$ and $\mathcal{W}_\mu(M)$ are canonically isomorphic for all modules M .*

Proof. The graph of the element g defines the desired isomorphism. □

Since whether a given weight space lies in N_λ^- only depends on which side of a hyperplane λ lives on, the points in a given coset of $\mathfrak{t}_\mathbb{Z}$ such that $N = N_\lambda^-$ for any given N are precisely those in the intersection of a finite number of half-spaces, i.e. a polyhedron. Thus, the corresponding clan is the W -orbit of these points.

Since only finitely many subspaces may appear as N_λ^- as λ ranges over an orbit of $\widehat{W}$:

Corollary 4.16. *Every $\widehat{W}$ -orbit is a union of finitely many clans, each defined by the W -orbit of the intersection of a $\mathfrak{t}_\mathbb{Z}$ -coset with a polyhedron, and thus has a finite complete set in the sense of Definition 2.28.*

Note that this result is not true for a general principal Galois order.

A *seed* is a weight $\gamma \in \text{MaxSpec}(\Gamma)$ which is the image of $\lambda \in \text{MaxSpec}(\Lambda)$ such that $P_\lambda = G_\lambda$.

Theorem 4.17. *If λ is a seed, there is a unique simple Gelfand–Tsetlin U -module S with $\mathcal{W}_\gamma(S) \cong \mathbb{k}$, and for all other simples S' we have $\mathcal{W}_\gamma(S') = 0$. The weight spaces of S satisfy $\dim \mathcal{W}_{\gamma'}(S) \leq \#(W/W_{\lambda'})$, and this bound is sharp if $N_\lambda^- = N_{\lambda'}^-$.*

Proof. First, we note that $U_\lambda^{(1)} \cong \mathbb{k}$, so this shows the desired uniqueness. The module $eP_\lambda^{(1)}$ is a weight module with S as cosocle satisfying $\dim \mathcal{W}_{\gamma'}(eP_\lambda^{(1)}) \leq \#W_\lambda/W_{\lambda'}$ whenever $\lambda' \in \widehat{W} \cdot \lambda$. This shows that desired upper bound.

We have that $\dim \mathcal{W}_{\gamma'}(S) = \#(W/W_{\lambda'})$ if and only if S is also the only Gelfand–Tsetlin module such that this weight space is non-zero, i.e. if and only if ${}_{\lambda}U_{\lambda'}^{(1)}$ is a Morita equivalence. This is clear if $N_{\lambda}^{-} = N_{\lambda'}^{-}$, since in this case $F_{\lambda}^{(1)} = F_{\lambda'}^{(1)}$ with ${}_{\lambda}F_{\lambda'}^{(1)}$ giving the obvious Morita equivalence. $\square$

Note that this shows that the module S discussed above has all the properties proven for the socle of the tableau module in [15, Theorem 1.1]. Using the numbering of that paper,

- (i) The weight γ itself lies in the essential support.
- (ii) This follows from Corollary 2.19.
- (iii) This follows from Theorem 4.17.
- (iv) For any parabolic subgroup $W' \subset W$, we can find a λ' such that $N_{\lambda'} = N_{\lambda}$, and $W' = W_{\lambda}$. The result then follows from Corollary 2.19.

4.5. Gelfand–Kirillov dimension. It will be useful for future applications to know some general facts about the Gelfand–Kirillov dimensions of Gelfand–Tsetlin modules⁶. Consider a field $\mathbb{k}$ and a $\mathbb{k}$ -algebra A which is generated by a finite-dimensional subspace A_0 , and a left A -module M which is finitely generated by a finite dimensional subspace M_0 . In this context, the *Gelfand–Kirillov dimension* $\mathrm{GKdim}_A(M)$ is defined by:

$$\mathrm{GKdim}_A(M) = \limsup_{n \rightarrow \infty} \log_n \dim_{\mathbb{k}}(A_0^n M_0) \quad (4.8)$$

It's a standard result that this number is independent of choice of A_0 and M_0 , and only depends on the structure of M as an A -module.

Let F be an Iwahori–Coulomb branch as in Definition 4.3. Let M be a GT module over F with support $\mathrm{supp}(L) \subset \mathfrak{t}_Q^* = \mathrm{Spec} \Lambda$. An important statistic that measures the “size” or “growth” of L is the dimension of the Zariski closure $m = \dim \overline{\mathrm{supp}(L)}$. Consider the quotient $A = F/\mathrm{ann}(M)$ by the annihilator of M .

In [36], Musson and van der Bergh prove two fundamental results about the Gelfand–Kirillov dimension of Gelfand–Tsetlin modules over Coulomb branches in the case where G is a torus (though they don't use this terminology):

Proposition 4.18 ([36, Corollary 8.2.5]). *If G is a torus, then*

$$\mathrm{GKdim}(A) = 2 \mathrm{GKdim}(M) = 2m.$$

For use in the future, we'd like to prove that this result holds for a general Coulomb branch.

Theorem 4.19. *For any connected reductive group G , representation N , and any Gelfand–Tsetlin module M over F , we have an equality:*

$$\mathrm{GKdim}(A) = 2 \mathrm{GKdim}(M) = 2d.$$

Proof. The algebra F contains as a subalgebra the Coulomb branch algebra F_{ab} associated to the maximal torus $T \subset G$ with the same matter representation. Let A_{ab} be the image of F_{ab} in F . Thus, we can restrict M to be an A_{ab} module and apply Proposition 4.18. This shows that:

$$\mathrm{GKdim}(A) \geq \mathrm{GKdim}(A_{\mathrm{ab}}) = 2m \quad \mathrm{GKdim}_A(M) \geq \mathrm{GKdim}_{A_{\mathrm{ab}}}(M) = m \quad (4.9)$$

To complete the proof, we need to show the reverse inequalities.

⁶Both of these notions are named after Israel Gelfand, but otherwise are unrelated.

$\boxed{\text{GKdim}(A) \leq 2m}$: In order to prove the reverse of the first equality of (4.9), we need to calculate some estimates on dimensions. For $w \in \widehat{W}$, let $F(\leq w) = F \cap K(\leq w)$ be the Λ -submodule given by the K -span of $w' \leq w$ in Bruhat order on $\widehat{W}$; this is a Λ - Λ -subbimodule. Using the geometric model for this algebra (following the notation of [24, Definition 2.2]), this is the homology of $R_{G,N}^B(\leq w)$, the preimage of the Schubert variety $\overline{IwI}/I$ in $R_{G,N}^B$.

Consider $F(\leq w)/F(< w)$. This is a free module of rank 1 over Γ as a left module or as a right module, this is spanned by a single element of minimal degree, which we denote $\mathfrak{r}_w$. The left and right actions differ by the action of w by [47, (3.6c) & (3.9d)]. For $n \geq 0$, let $F(\leq n)$ be the span of $F(\leq w)$ for all w of length $\ell(w) \leq n$.

Taking the corresponding quotient $A(\leq w)/A(< w)$, we thus obtain a Γ - Γ -bimodule whose support as a left and a right module must be in $\overline{\text{supp}(M')}$. Since these actions differ by w , the support as a left Γ module must lie in $\overline{\text{supp}(M')} \cap w \cdot \overline{\text{supp}(M')}$. The affine Weyl group elements where this intersection is $\geq k$ dimensional have translation parts that lie in a $2d - k$ dimensional variety, since all the components of $\overline{\text{supp}(M')}$ are affine subspaces which are $\leq d$ dimensional. This shows that:

- (i) The number of $w \in \widehat{W}$ of length $\leq \ell$ such that $\dim \overline{\text{supp}(M')} \cap w \cdot \overline{\text{supp}(M')} = k$ is bounded above by $D\ell^{2d-k}$ for some constant D .

Now, consider the span A_0 of

- (1) the degree 1 elements $\mathfrak{t}^* \subset \Gamma$ and
- (2) generators of $F(\leq n)$ as a left Γ -module for a fixed n .

If we choose n sufficiently large, this subspace will be a set of generators of F as an algebra. The q th power A_0^q lies in $F(\leq nq)$. Furthermore, if we let $d(w)$ be the degree of the unique generator of $F(\leq w)/F(< w)$, then this depends at worst linearly on $\ell(w)$: we have $|d(w)| \leq C'\ell(w)$ for some constant $C' > 1$. This shows that elements of A_0^q have degree no more than $C'nq$.

If $\dim \overline{\text{supp}(M')} \cap w \cdot \overline{\text{supp}(M')} \leq k$, then we must have that the dimension of the span of the elements of degree $\leq p$ in Γ times the cyclic generator in $A(\leq w)/A(< w)$ must be bounded by $C''p^k$ for some constant C'' ; since this intersection is a union of affine spaces, whose number of components is bounded by the number of pairs of components $\overline{\text{supp}(M')}$, we can choose one C'' which works for all w .

- (ii) The dimension of $(A(\leq w) \cap A_0^q)/(A(< w) \cap A_0^q)$ is bounded above by $C''(C'nq)^k$ if $\ell(w) \leq nq$. Note that if we choose $C = C''(C'n)^{2d}$, then this dimension is bounded above by $\leq Cq^k$.

Combining observations (i) and (ii) and summing over $k = 1, \dots, 2d$, we have $\dim A_0^q \leq 2dCDn^{2d}q^{2d}$. Thus, we have

$$\log_q(\dim A_0^q) \leq 2d + \frac{\log(2dCD) + 2d \log(n)}{\log q}$$

so taking the limit, we have $\text{GKdim}(A) \leq 2d$. Combining with (4.9), we find that $\text{GKdim}(A) = 2d$.

$\boxed{\text{GKdim}_A(M) \leq m}$: Now we turn to showing the reverse of the second inequality in (4.9). For finite dimensional subset $A_0 \subset A$ and any metric on $\tilde{\mathfrak{t}}$, there is a point $x \in \tilde{\mathfrak{t}}$ and a real number $\epsilon > 0$ such that the ball $B_t(x)$ of radius t around x satisfies

$$A_0 \cdot \bigoplus_{\lambda \in B_t(x)} \mathcal{W}_\lambda(M) \subset \bigoplus_{\lambda \in B_{t+\epsilon}(x)} \mathcal{W}_\lambda(M).$$

For $t \gg 0$, the sum $\bigoplus_{\lambda \in B_t(x)} \mathcal{W}_\lambda(M)$ generates M as an A module. Thus, the Gelfand–Kirillov dimension satisfies:

$$\mathrm{GKdim}(M) \leq \lim_{t \rightarrow \infty} \frac{\log \dim \bigoplus_{\lambda \in B_t(x)} \mathcal{W}_\lambda(M)}{\log t}$$

Since the Zariski closure $\overline{\mathrm{supp}(M)}$ is unchanged by considering M as an A_{ab} -module, by [36, Proposition 7.2.4], the closure $\overline{\mathrm{supp}(M)}$ is the union of finitely many affine spaces. Of course, d is the maximum of these dimensions of affine spaces, and the support of M is the intersection of a lattice with this union of affine spaces. This shows that

$$\lim_{t \rightarrow \infty} \frac{\log \dim \bigoplus_{\lambda \in B_t(x)} \mathcal{W}_\lambda(M^B)}{\log t} = d$$

which completes the proof that $\mathrm{GKdim}_A(M) = m$. $\square$

Since it will be useful at other times, let us note that $\mathrm{supp}(M)$ is a union of finitely many clans, and $\overline{\mathrm{supp}(M)}$ is the union of the Zariski closure of these clans. Thus, we have:

Lemma 4.20. *The Gelfand–Tsetlin dimension of M is $\geq d$ if and only if M has non-zero multiplicity on a clan whose Zariski closure is $\geq d$ -dimensional.*

5. THE CASE OF ORTHOGONAL GELFAND–TSETLIN ALGEBRAS

We’ll continue to assume that $\mathbb{k}$ has characteristic 0. This is not strictly necessary for Theorem 5.1, but will be needed for all later results in this section.

5.1. Orthogonal Gelfand–Tsetlin algebras as Coulomb branches. Let us now briefly describe how one can interpret the results of this paper for orthogonal Gelfand–Tsetlin algebras [33] over $\mathbb{k}$ in terms of [23]. As in the introduction, choose a dimension vector $\mathbf{v} = (v_1, \dots, v_n)$ and fix scalars $(\lambda_{n,1}, \dots, \lambda_{n,v_n}) \in \mathbb{k}^{v_n}$. Let

$$\Omega = \{(i, r) \mid 1 \leq i \leq n, 1 \leq r \leq v_i\}.$$

Let $U = U_{\mathbf{v}}$ be the associated orthogonal Gelfand–Zetlin algebra modulo the ideal generated by specializing $x_{n,r} = \lambda_{n,r}$. This is a principal Galois order with the data:

- The ring Λ given by the polynomial ring generated by $x_{i,j}$ with $(i, j) \in \Omega$ and $i < n$. Note that we have not included the variables $x_{n,1}, \dots, x_{n,v_n}$, since these are already specialized to scalars.
- The monoid $\mathcal{M}$ given by the subgroup of $\mathrm{Aut}(\Lambda)$ generated by $\varphi_{i,j}$, the translation satisfying

$$\varphi_{i,j}(x_{k,\ell}) = (x_{k,\ell} + \delta_{ik}\delta_{j\ell})\varphi_{i,j}$$

- The group $W = S_{v_1} \times \dots \times S_{v_{n-1}}$, acting by permuting each alphabet of variables.

By definition, U is the subalgebra of $\mathcal{K}$ generated by $\Gamma = \Lambda^W$ and the elements

$$X_i^\pm = \mp \sum_{j=1}^{v_i} \frac{\prod_{k=1}^{v_{i+1}} (x_{i,j} - x_{i+1,k})}{\prod_{k \neq j} (x_{i,j} - x_{i,k})} \varphi_{i,j}^\pm$$

Let $F = F_D$ be the corresponding Morita flag order. This is the subalgebra of $\mathcal{F}$ generated by U embedded in $e\mathcal{F}e \cong \mathcal{K}$ and the nilHecke algebra $D = \text{End}_\Gamma(\Lambda)$.

As mentioned in the introduction, it is proven in [51] that:

Theorem 5.1 ([51, Corollary 3.16 & Theorem A]). *We have an isomorphism between the OGZ algebra attached to the dimension vector $\mathbf{v}$ and the Coulomb branch at $\hbar = 1$ for the (G, N)*

$$G = GL_{v_1} \times \cdots \times GL_{v_{n-1}}$$

$$N = M_{v_n, v_{n-1}}(\mathbb{C}) \oplus M_{v_{n-1}, v_{n-2}}(\mathbb{C}) \oplus \cdots \oplus M_{v_2, v_1}(\mathbb{C}),$$

where Q is given by the product of G with the diagonal matrices in GL_{v_n} and the variables $x_{n,1}, \dots, x_{n,v_n}$ are given by the equivariant parameters for $Q/G \cong (\mathbb{C}^\times)^{v_n}$.

If we assume that

$$v_1 \leq v_2 - v_1 \leq v_3 - v_2 \leq \cdots \leq v_n - v_{n-1},$$

then U is isomorphic to the quotient of a finite W -algebra of $U(\mathfrak{gl}_{v_n})$ for a nilpotent matrix of Jordan type $(v_1, v_2 - v_1, \dots, v_{n-1})$, modulo a maximal ideal of the center fixed by the scalars $\lambda_{n,*}$. In particular, if $\mathbf{v} = (1, 2, \dots, n)$, then U is the universal enveloping algebra $U(\mathfrak{gl}_n)$ itself modulo this maximal ideal.

Note that here we use the realization of W -algebras as quotients of shifted Yangians proven in [6, Theorem 10.1] (refer to [50, Theorem 4.3(a)] for a version of this more compatible with Weekes' notation). If you would prefer not to mod out by this maximal ideal, we can leave $x_{n,*}$ as variables, and take invariants of S_{v_n} permuting these variables to obtain the full W -algebra.

Proof. Using the substitutions above, the element denoted by X_k^+ in [22, (4.6)] is the image of $\pm E_i^{(1)}$ under the map of [4, Theorem B.15], and similarly X_k^- is the image of the z^{-1} -coefficient of $\pm F_i^{(1)}$. Thus, the orthogonal Gelfand–Tsetlin algebra is contained in the image of this map. On the other hand, the relation [4, (B.5)] shows that these elements and the polynomials in $w_{i,r}$'s generate the image, so the image is precisely the OGZ algebra. This map is surjective by [4, Theorem B.28], which induces the desired isomorphism. $\square$

Thus, we can apply the results of Section 4 to OGZ algebras. An element $\lambda \in \text{MaxSpec}(\Lambda)$ is exactly choosing a numerical value $x_{i,r} = \lambda_{i,r}$ for all $(i, r) \in \Omega$, and the corresponding $\gamma \in \text{MaxSpec}(\Gamma)$ only remembers these values up to permutation of the second index. A choice of λ partitions the set Ω according to which coset of $\mathbb{Z}$ the value $\lambda_{i,r}$ lies in. Given a coset $[a] \in \mathbb{k}/\mathbb{Z}$, let

$$\Omega_{[a]} = \{(i, r) \in \Omega \mid \lambda_{i,r} \equiv a \pmod{\mathbb{Z}}\}.$$

The maximal ideal λ has an integral orbit if there is one coset such that $\Omega = \Omega_{[a]}$.

Note that the representation N is spanned by the dual basis to the matrix coefficients of the maps $\mathbb{C}^{v_k} \rightarrow \mathbb{C}^{v_{k+1}}$, which we denote $h_{r,s}^{(k)}$ for $1 \leq r \leq v_k$ and $1 \leq s \leq v_{k+1}$.

Proposition 5.2. *Given $\lambda \in \text{MaxSpec}(\Lambda)$, we have that N_λ is the span of the elements $h_{r,s}^{(k)}$ such that $\lambda_{k,r} - \lambda_{k+1,s} \in \mathbb{Z}$, and N_λ^- is the span of these elements with $\lambda_{k,r} - \lambda_{k+1,s} \in \mathbb{Z}_{\geq 0}$.*

Remark 5.3. Note that equivalence classes of weights in a $\widehat{W}$ -orbit with N_{λ}^{-} fixed also appear in the discussion of generic regular modules in [10, § 3.3]. That is, the subspace N_{λ}^{-} changes precisely when the numerator of one of the Gelfand–Tsetlin formulae vanishes.

We can encapsulate this with an order on the set Ω which is the coarsest such that $(i, r) \prec (i + 1, s)$ if $\lambda_{i,r} - \lambda_{i+1,s} \in \mathbb{Z}_{<0}$ and $(i, r) \succ (i + 1, s)$ if $\lambda_{i,r} - \lambda_{i+1,s} \in \mathbb{Z}_{\geq 0}$. Lemma 4.15 then shows that:

Proposition 5.4. *The weights λ and λ' are in the same clan if and only if for all pairs (i, r) and $r \in [1, v_i]$, we have $\lambda_{i,r} - \lambda'_{i,r} \in \mathbb{Z}$, and the induced order on the set Ω is the same.*

While interesting, these observations are not a large advance over what was known in the literature. To get a more detailed answer, we must use Theorem 4.4 more carefully. As we’ve discussed, this depends sensitively on the integrality conditions of $\mathcal{S}$. If $\mathcal{S}$ is not integral, then by Corollary 4.7, the category $\Gamma\Pi(\mathcal{S})$ is equivalent to the category of Gelfand–Tsetlin modules supported on the same orbit for a tensor product $\bigotimes_{[a] \in \mathbb{k}/\mathbb{Z}} U_{[a]}$ where $U_{[a]}$ is the OGZ algebra attached to the set $\Omega_{[a]}$, that is, to the dimension vector $\mathbf{v}^{(a)}$ given by the number of indices k such that $\lambda_{i,k} \equiv a \pmod{\mathbb{Z}}$. Since the simple Gelfand–Tsetlin modules over this tensor product are just outer tensor products of the simple Gelfand–Tsetlin modules over the individual factors (and in fact, the category $\Gamma\Pi(\mathcal{S})$ is a Deligne tensor product of the corresponding category for the factors), let us focus our attention on the integral case.

5.2. The integral case. Let $\mathcal{S}_{\mathbb{Z}}$ be the $\widehat{W}$ -orbit where $\lambda_{i,r} \in \mathbb{Z}$ for all $(i, r) \in \Omega$, and we fix integral values $\lambda_{n,1} \leq \dots \leq \lambda_{n,v_n}$. All integral orbits differ from this one by a uniform shift, and all these orbits are equivalent via the functor of tensor product with a one-dimensional representation where $\mathfrak{gl}_n$ acts by a multiple of the trace.

In this case, we are effectively rephrasing [23, Theorem 5.2] in a slightly different language and in the notation of this paper. Identify $I = \{1, \dots, n-1\}$ with the Dynkin diagram of $\mathfrak{sl}_n$ as usual. Let $\widehat{T}_{\mathbf{v}}$ be the block of the KLRW algebra as discussed in [23, § 3.1], attached to the sequence $(\omega_{n-1}, \dots, \omega_{n-1})$ with this fundamental weight appearing v_n times and where v_i black strands have the label i for all $i \in I$. Note that this algebra contains a central copy of the algebra

$$\mathcal{Z}(\mathcal{S}_{\mathbb{Z}}) = \bigotimes_{i=1}^{n-1} \mathbb{k}[x_{i,1}, \dots, x_{i,v_i}]^{S_{v_i}},$$

given by the polynomials in the dots which are symmetric under permutation of all strands.

Fix a very small real number $0 < \epsilon \ll 1$. Given a weight λ , we define a map

$$x: \Omega \rightarrow \mathbb{R} \quad x(i, s) = \lambda_{i,s} - i\epsilon - s\epsilon^2.$$

Note that under this map, the partial order $\prec$ is compatible with the usual order on $\mathbb{R}$; this map thus gives a canonical way to refine $\prec$ and the order on Ω induced by the usual partial order on $\lambda_{i,s}$ to a total order on Ω . The ϵ term is very important for assuring the compatibility with $\prec$, whereas the ϵ^2 term is essentially arbitrary and is only there to avoid issues when two strands go to the same place.

Let $w(x)$ be the word in $[1, n]$ given by ordering the elements of Ω according to the function x , and then projecting to the first index.

Now, consider the idempotent $e(\lambda)$ in $\tilde{T}_{\mathbf{v}}$ where we place a red strand with label ω_{n-1} at $x(n, r)$ for all $r = 1, \dots, v_n$, and a black strand with label i at $x(i, s)$ for all $i \in I$ and $s = 1, \dots, v_i$.

Definition 5.5. Let $w(\lambda)$ be the word in $[1, n]$ given by ordering the elements of Ω according to the function x described above for a given λ .

The labels of strands read left to right are just the word $w(\lambda)$. The isomorphism type of this idempotent only depends on the partial order $\prec$, and it would be the same for any map x that preserves this order. For example, we would match [23] more closely if we used $x(i, s) = 2\lambda_{i,s} - i$ (again with a perturbation to ensure that all elements have distinct images), which works equally well. This choice matches better with the parameterization of Γ by the variables $w_{i,k}$ used in [4].

Let $S \subset \mathcal{S}_{\mathbb{Z}}$ be a finite set. For simplicity, we assume that this set does not have pairs of weights that correspond as in Proposition 5.4, up to the action of W . Of course, this set will be complete if every possible partial order $\prec$ that appears in the orbit $\mathcal{S}_{\mathbb{Z}}$ is realized. Let e_S be the sum of these idempotents in $\tilde{T}_{\mathbf{v}}$.

Theorem 5.6. *The algebra $\hat{F}(S)$ is isomorphic to the completion with respect to its grading of $e_S \tilde{T}_{\mathbf{v}} e_S$, and $F^{(1)}(S)$ is isomorphic to $e_S \tilde{T}_{\mathbf{v}} e_S$ modulo all positive degree elements of $\mathcal{X}(\mathcal{S}_{\mathbb{Z}})$.*

This is truly a restatement of [23, Theorem 5.2], but can also be derived from Theorem 4.4, using the convolution description of $\tilde{T}_{\mathbf{v}}$ as a convolution algebra based on [48, Theorems 4.5 & 3.5]. If you prefer to keep $x_{n,r}$ as variables rather than specializing them, then the resulting algebra is the deformation $\tilde{\mathbb{T}}_{\mathbf{v}}$ of $\tilde{T}_{\mathbf{v}}$ defined in [41, Definition 2.7]; geometrically, this is reflected by whether we keep equivariance for the group $J = Q/G$. In [41, Proposition 3.4], we give a more algebraic proof of this result, which incorporates the variables $x_{n,r}$ and thus accounts for modules over the OGZ algebra where the action of $x_{n,r}$ is not nilpotent; [49, Lemma 4.11] provides a useful summary of how other properties of $U(\mathfrak{gl}_n)$ -modules transfer.

This reduces the question of understanding Gelfand–Tsetlin modules to studying the simple representations of these algebras. The usual theory of translation functors shows that the structure of this category only depends on the stabilizer under the action of S_{v_n} on the element $(\lambda_{n,1}, \dots, \lambda_{n,v_n})$. This is a Young subgroup of the form $S_{\mathbf{h}} = S_{h_1} \times \dots \times S_{h_\ell}$; of course, a regular block will have all $h_k = 1$. Consider the sequence of dominant weights $\mathbf{h} = (h_1 \omega_{n-1}, \dots, h_\ell \omega_{n-1})$. This corresponds to the tensor product $\text{Sym}^{h_1}(Y) \otimes \text{Sym}^{h_2}(Y) \otimes \dots \otimes \text{Sym}^{h_\ell}(Y)$, where Y is the dual of the vector representation of $\mathfrak{sl}_n$. Thus, by [23, Proposition 3.1], we have that: $K^0(\tilde{T}_{\mathbf{v}}^{\mathbf{h}}) \cong U(\mathbf{h})$ where $\mathbf{n}_-$ is the algebra of $n \times n$ strictly lower triangular matrices and

$$U(\mathbf{h}) := U(\mathbf{n}_-) \otimes \text{Sym}^{h_1}(Y) \otimes \text{Sym}^{h_2}(Y) \otimes \dots \otimes \text{Sym}^{h_\ell}(Y).$$

While we have a general theorem connecting simples over $\tilde{T}_{\mathbf{v}}^{\mathbf{h}}$ to the dual canonical basis of $U(\mathbf{h})$, because we are looking at a particularly simple special case, this combinatorics simplifies.

5.3. Goodly combinatorics. Following the work of Leclerc [30] and the relation of this work to KLR algebras discussed in [28], we can give a simple indexing set of this dual canonical basis. Consider a simple Gelfand–Tsetlin module S , and the set $\mathbb{L}(S)$ of words

$w(\lambda)$ for $\lambda \in \mathcal{S}_{\mathbb{Z}}$ such that $\mathcal{W}_{\lambda}(S) \neq 0$. We order words in the set $[1, n]$ lexicographically, with the rule that $(i_1, \dots, i_{k-1}) > (i_1, \dots, i_k)$.

Definition 5.7. We call a word *red-good* if it is minimal in lexicographic order amongst $\mathbb{L}(S)$ for some simple S . Since $\mathbb{L}(S)$ is finite, every simple has a unique good word.

Let $\mathcal{GL}$ be the set of words of the form $(k, k-1, \dots, k-p)$ for $k \leq n-1$, and $0 \leq p < k$, and $\mathcal{GL}'$ be the set of words of the form $(n, n-1, \dots, n-p)$ for $0 \leq p < n$; as noted in [30, § 6.6], these together form the good Lyndon words of the A_n root system in the obvious order on the nodes in the Dynkin diagram (which we identify with $[1, n]$).

Definition 5.8. We say a word $\mathbf{i}$ is *goodly* if it is the concatenation $\mathbf{i} = a_1 \cdots a_p b_1 \cdots b_{v_n}$ of words for $a_k \in \mathcal{GL}$, and $b_k \in \mathcal{GL}'$ that satisfies $a_1 \leq a_2 \leq \cdots \leq a_p$ in lexicographic order.

For simplicity, assume that the central character $(\lambda_{n,1}, \dots, \lambda_{n,v_n})$ is regular, that is, $S_{\mathbf{h}} = \{1\}$. In this case, a goodly word can always be realized as $w(\lambda^{(\mathbf{i})})$ for a weight $\lambda^{(\mathbf{i})}$ chosen as follows: pick integers $\mu_1, \dots, \mu_p$ so that $\mu_1 < \cdots < \mu_p < \lambda_{n,1} < \cdots < \lambda_{n,v_n}$. Now, choose the set $\lambda_{i,*}^{(\mathbf{i})}$ so that μ_k appears (always with multiplicity 1) if and only if i appears as a letter in a_k , and $\lambda_{n,q}$ if and only if i appears as a letter in b_q . This weight depends on the choice of μ_* , but all these choices are equivalent via Lemma 4.15.

Theorem 5.9. *The map sending a simple Gelfand–Tsetlin module to its red-good word is a bijection, and a word is red-good if and only if it is goodly.*

Note that implicit in the theorem above is that we consider the set of all red-good words for all different $\mathbf{v}$'s, but $\mathbf{v}$ is easily reconstructed from the word, just letting v_i be the number of times i appears.

Proof. Note that the words in $\mathcal{GL}$ index cuspidal representations of the KLR algebra of $\mathfrak{sl}_n$ in the sense of Kleshchev–Ram [28]; thus, concatenations of these words in increasing lexicographic order give the good words for $\mathfrak{sl}_n$, and the lex maximal word in the different simple representations of the KLR algebra of $\mathfrak{sl}_n$ by [28, Theorem 7.2].

On the other hand, the words $\mathcal{GL}'$ give the idempotents corresponding to the different simples over the cyclotomic quotient $T^{\omega_{n-1}}$, which are all 1-dimensional.

Thus, given a red-good word $\mathbf{i} = a_1 \cdots a_p b_1 \cdots b_{v_n}$, there is a unique simple L_0 over $\tilde{T}^0$ corresponding to $a_1 \cdots a_p$ and v_n simple modules $L_1, \dots, L_n$ over $T^{\omega_{n-1}}$ corresponding to $b_1, \dots, b_n$. By [46, Corollary 5.23], the standardization $M(\mathbf{i})$ over these simples has a unique simple quotient $L(\mathbf{i})$, and every simple appears this way for a unique goodly word. Note that the standardization $M(\mathbf{i})$ has the property that if $e(\mathbf{j})M(\mathbf{i}) \neq 0$, then $\mathbf{j}$ is a shuffle of words whose idempotents have non-zero image on $L_0, \dots, L_n$. Since n is first letter of such a word for L_i with $i > 0$, and n does not appear in any word for L_0 , any such shuffle which is non-trivial will be lex-greater than the trivial shuffle of the same words. In particular, the lex-minimal word $\mathbf{j}$ such that $e(\mathbf{j})M(\mathbf{i}) \neq 0$ must be the concatenation of the corresponding lex-minimal words for $L_0, \dots, L_n$. This is precisely the goodly word $\mathbf{i}$.

The image $e_S L$ gives a simple module over $F_S^{(1)}$ for any set S containing the weight $\lambda^{(\mathbf{i})}$ and thus a simple Gelfand–Tsetlin -module S by Theorem 2.27. We claim that $\mathbf{i}$ is the red-good word for this simple.

For any other word that appears as $w(\lambda) < \mathbf{i}$, we can add λ to S , and by the discussion above, we have $\mathcal{W}_{\lambda}(S) = e(\lambda)L = 0$, showing that $\mathbf{i}$ is the red-good word of S . This shows that the map from representations to red-good words is surjective.

Consider any other simple S' . By the discussion above, this comes from a simple $\tilde{T}_{\mathbf{v}}$ representation L' , which is the quotient of the standardization of a different goodly word $\mathbf{i}'$. As we've already argued, this means that $\mathbf{i}' \neq \mathbf{i}$ is its red-good word. This shows that the map on red-good words is injective and completes the proof. $\square$

Example 5.10. For example, the case of integral Gelfand–Tsetlin modules of $\mathfrak{sl}_3$ corresponds to $\mathbf{v} = (1, 2, 3)$. Thus, the red-good words are of the form:

$$\begin{array}{ccc}
 (1|2|2|3|3|3) & (2, 1|2|3|3|3) & \\
 (1|2|3, 2|3|3) & (1|2|3|3, 2|3) & (1|2|3|3|3, 2) \\
 (2, 1|3, 2|3|3) & (2, 1|3|3, 2|3) & (2, 1|3|3|3, 2) \\
 (2|3, 2, 1|3|3) & (2|3|3, 2, 1|3) & (2|3|3|3, 2, 1) \\
 (1|3, 2|3, 2|3) & (1|3|3, 2|3, 2) & (1|3, 2|3|3, 2) \\
 (3, 2, 1|3, 2|3) & (3|3, 2, 1|3, 2) & (3, 2, 1|3|3, 2) \\
 (3, 2|3, 2, 1|3) & (3|3, 2|3, 2, 1) & (3, 2|3|3, 2, 1)
 \end{array}$$

We've included vertical bars $|$ between the Lyndon factors of each word.

In order to construct the actual weights appearing, we choose

$$\mu_1 = -2 < \mu_2 = -1 < \mu_3 = 0 < \lambda_{3,1} = 1 < \lambda_{3,2} = 2 < \lambda_{3,3} = 3.$$

We'll represent maximal ideals of the Gelfand–Tsetlin subalgebra using tableaux, where the entries of the k^{th} row from the bottom are the roots of $\prod_{j=1}^k (u - x_{k,j}) \in \Lambda[u]$ reduced modulo the maximal ideal. Accordingly, these entries come as an unordered k -tuple, which we write below in decreasing order.

Using this notation, the corresponding weight spaces $\lambda^{(\mathbf{i})}$ for the words above are shown in Figure 5.1.

Thus, each generic integral block for $\mathfrak{gl}_3$ has 20 simple Gelfand–Tsetlin modules. We discuss the structure of these modules and extend this calculation to other low-rank cases in joint work with Silverthorne [41]. As mentioned in the introduction, we have done computer computation of the dimensions of the weight spaces of simples through $\mathfrak{sl}_4$, and of the number of simples in the principal block up through $\mathfrak{sl}_9$.

5.4. The singular case. This theorem is a little more awkward to state for the singular case where $S_{\mathbf{h}} \neq \{1\}$. To understand this case, it will help to recall a few facts about the cyclotomic quotient corresponding to the highest weight $h\omega_{n-1}$.

Lemma 5.11.

- (1) The algebra $T_{\mathbf{v}}^{h\omega_{n-1}}$ is non-zero if and only if $h \geq v_{n-1} \geq \cdots \geq v_1$.
- (2) The algebra $T_{\mathbf{v}}^{h\omega_{n-1}}$ is Morita equivalent to the cohomology ring of the variety partial flags in $\mathbb{C}^h$ with subspaces of size v_* , and thus has a unique simple module $M_{\mathbf{v}}$.
- (3) The image $e(\mathbf{i})M$ is non-zero if and only if $\mathbf{i}$ is a shuffle of words from $\mathcal{GL}'$ with their initial n removed.




FIGURE 5.1. The tableaux corresponding to the weights $\lambda^{(i)}$ for the red-good words for the principal block of $\mathfrak{sl}_3$.

Proof.

- (1) By [46, Proposition 3.21], the Grothendieck group of the category of $T_{\mathbf{v}}^{h\omega_{n-1}-}$ modules is the $h\omega_{n-1} - v_{n-1}\alpha_{n-1} - \cdots - v_1\alpha_1$ weight space of an integral form of the representation $\text{Sym}^h(Y)$. In the usual description of the integral weights of $\mathfrak{sl}_n$ as n -tuples of integers modulo the span of $(1, \dots, 1)$, we have the following.

$$\begin{aligned} h\omega_{n-1} - v_{n-1}\alpha_{n-1} - \cdots - v_1\alpha_1 &= (0, \dots, 0, -h) - (0, \dots, v_{n-1}, -v_{n-1}) - \cdots - (v_1, -v_1) \\ &= (-v_1, -v_2 + v - 1, \dots, -h + v_{n-1}). \end{aligned}$$

The condition that $h \geq v_{n-1} \geq \cdots \geq v_1$ is equivalent to all the entries of this vector being negative, which indeed describes exactly the weights of $\mathrm{Sym}^h(Y)$.

- (2) By [49, Theorem 3.18], the deformed cyclotomic quotient $T_{\mathbf{v}}^{h\omega_{n-1}}$ is Morita equivalent to the GL_n equivariant cohomology ring of this partial flag variety, since the constant sheaf generates the derived category of GL_n -equivariant constructible sheaves on this flag variety. Passing to the undeformed quotient $T_{\mathbf{v}}^{h\omega_{n-1}}$ kills the equivariant parameters, giving the result by the equivariant formality of the partial flag variety.
- (3) The image $e(\mathbf{i})M$ is nonzero if and only if the idempotent $e(\mathbf{i})$ itself is. Using [49, Theorem 3.18] again, we see that this is the case if and only if the sheaf $\mathrm{Res}(\mathcal{F}_{\mathbf{i}})$ (in the notation of [49, Theorem 3.18]) is nonzero. This is the pushforward to the partial flag variety of a particular quiver flag variety. It's the pushforward of the set of $\sum_i v_i$ -tuples of flags, where for each k , we choose a flag $V_1^{(k)} \subseteq V_2^{(k)} \subset \cdots$ such that $V_i^{(k)} \subset V_i^{(k+1)}$ for all i and k , and the dimension of $V_i^{(k)}$ is the number of times i appears in the first k letters of $\mathbf{i}$. We leave it to the reader to check that $\mathbf{i}$ being a shuffle of the desired form is equivalent to the existence of such a flag for simple dimension reasons. $\square$

Note that this implies that there is a bijection between simple modules over $T_{\mathbf{v}}^{h\omega_{n-1}}$ and unordered h -tuples of words from $\mathcal{GL}'$.

For slightly silly reasons, the red-good words as we have defined them depend on the choice of $\lambda_{n,*}$, but we can still consider goodly words $\mathbf{i} = a_1 \cdots a_p b_1 \cdots b_{v_n}$ and the associated weight $\lambda^{(\mathbf{i})}$. Note that this now only depends on the choice of $b_1, \dots, b_{v_n}$ up to permutations under $S_{\mathbf{h}}$.

Proposition 5.12.

- (1) For each goodly word $\mathbf{i} = a_1 \cdots a_p b_1 \cdots b_{v_n}$ which is lex maximal in its $S_{\mathbf{h}}$ -orbit, there is a unique simple Gelfand–Tsetlin module S such that $\mathcal{W}_{\lambda^{(\mathbf{i})}}(S) \neq 0$, and $\mathcal{W}_{\lambda^{(\mathbf{i}')}}(S) = 0$ for all $\mathbf{i}'$ of the same form with $\mathbf{i}' < \mathbf{i}$, and this gives a complete irredundant list of simple modules in $\Pi(\mathcal{S}_{\mathbb{Z}})$ for the corresponding central character.
- (2) If S is a complete set, then $\widehat{F}_S$ is Morita equivalent to the completion of $\widetilde{T}_{\mathbf{v}}^{\mathbf{h}}$ with respect to its grading for $\mathbf{h} = (h_1\omega_{n-1}, \dots, h_{\ell}\omega_{n-1})$, and $F_S^{(1)}$ to the quotient of this algebra by positive degree elements of $\Lambda_{\mathcal{S}_{\mathbb{Z}}}$.

Proof. We'll actually prove part (2) first. By Theorem 5.6, it's enough to show that the ring $e_S \widetilde{T}_{\mathbf{v}} e_S$ has the desired Morita equivalence.

Given $\lambda \in \mathrm{Spec}(\Lambda)$, since we will never have a black strand between the red strands that correspond to $\lambda_{n,k} = \lambda_{n,k+1}$, we have that $e(\lambda) \in \widetilde{T}_{\mathbf{v}}^{\mathbf{h}}$ embedded as in [46, Proposition 4.21] by “zipping” the red strands. Thus, e_S and $e_S \widetilde{T}_{\mathbf{v}} e_S$ will lie in this subalgebra.

By standard results of Morita theory, it's enough to check that no simple module over $\widetilde{T}_{\mathbf{v}}^{\mathbf{h}}$ is killed by e_S . By [46, Corollary 5.23], every such simple is obtained by standardization of a module L_0 over the usual KLR algebra $\widetilde{T}^{\emptyset}$, and then of modules L_i over the cyclotomic quotient corresponding to $h_i\omega_{n-1}$. Of course, by Lemma 5.11, L_i is uniquely determined by an unordered h -tuple of words from $\mathcal{GL}'$; we can uniquely construct a word from these by taking the lex-maximal element of the set of such concatenations. Construct a goodly word $\mathbf{i}_L$ by concatenating the good word corresponding to L_0 , with the words just attached to $L_1, L_2, \dots, L_{\ell}$.

This word has a corresponding weight $\lambda_{\mathbf{i}_L}$; note that turning this back into a word via the usual rule, we don't get $\mathbf{i}$ back, but instead, for each h_i -tuple, we get the word sorted in descending order $(n, \dots, n, n-1, \dots, n-1, \dots, 1, \dots, 1)$, since all the variables $\lambda_{*,*}$ assigned to these black strands have the same longitude. This is the red-good word for the simple, following the definition precisely. Since this is a shuffle of the h_i -tuple of words in $\mathcal{GL}'$, the corresponding idempotent has nonzero image on the simple L_i . This means that $e(\lambda_{\mathbf{i}_L})$ has non-zero image on the standard module, and its image contains a pure tensor of non-zero vectors in the simples $L_0, \dots, L_\ell$, and thus generates the standard module. This shows that $e(\lambda_{\mathbf{i}_L})$ also has non-zero image in L . This shows the Morita equivalence, since no simple is killed by e_S ; thus (2) holds.

Now, let us show (1). Let S be the GT module corresponding to L . We have already noted that $\mathcal{W}_{\lambda^{(\mathbf{i})}}(S) \neq 0$. We wish to show that $\mathcal{W}_{\lambda^{(\mathbf{i}')}}(S) = 0$ for all $\mathbf{i}'$ of the same form with $\mathbf{i}' < \mathbf{i}$. By construction, the word $\mathbf{i}'$ must be a shuffle of the words $a_1, \dots, a_p, b_1, \dots, b_q$ without crossing any red strands. Consider the first letter in $\mathbf{i}'$ which is different from the corresponding letter in $\mathbf{i}$. This must be the first letter of one of the words a_i or b_i . If it comes from one of the words $b_i \in \mathcal{GL}'$, then it is n in $\mathbf{i}'$, so we must have $\mathbf{i}' > \mathbf{i}$. If it is from one of the words a_j , then simply deleting the letter from b_* gives a shuffle of the words $a_1, \dots, a_p$ which is lex-lower. This is impossible by [30, Lemma 15]. $\square$

6. ON A CONJECTURE OF MAZORCHUK

We say that a maximal ideal $\Gamma \subset U(\mathfrak{gl}_n)$ is a *Gelfand–Tsetlin pattern* if all $\lambda_{i,k}$ lie in the same coset of $\mathbb{Z}$ in $\mathbb{C}$, and the order $\prec$ satisfies $(i, s) \prec (i-1, s) \prec (i, s+1)$ for $i = 2, \dots, n$ and $s = 1, \dots, i-1$. As discussed previously, if a representation of $\mathfrak{gl}_n$ is finite dimensional, then its spectrum consists of precisely the Gelfand–Tsetlin patterns with fixed $\lambda_{n,*}$. This result is implicit in the original work of Gelfand and Tsetlin [19] and was developed further by Zhelobenko [52, Theorem 13.5]; see [35, Theorem 2.20] for a more modern treatment.

Mazorchuk communicated to us a conjecture which would be a strong converse to this result:

Conjecture 6.1. If S is a simple $U(\mathfrak{gl}_n)$ module, and $\mathcal{W}_\gamma(S) \neq 0$ for γ a Gelfand–Tsetlin pattern, then S is finite-dimensional. That is, for any $\gamma \in \text{MaxSpec}(\Gamma)$, then either:

- (1) $\mathcal{W}_\gamma(S) = 0$ for all infinite-dimensional simple modules S and $\mathcal{W}_\gamma(S') \neq 0$ for some finite-dimensional simple module S' (i.e. γ is a Gelfand–Tsetlin pattern) or
- (2) $\mathcal{W}_\gamma(S') = 0$ for all finite-dimensional simple modules S and $\mathcal{W}_\gamma(S) \neq 0$ for some infinite-dimensional simple module S (i.e. γ is not a Gelfand–Tsetlin pattern).


Embarrassingly, we at one point claimed to have a proof of this fact. Unfortunately, this proof was incorrect and a more careful computer search showed that:

Theorem 6.2. *Conjecture 6.1 holds for $n \leq 5$ and is false for $n \geq 6$. That is, a Gelfand–Tsetlin pattern has an infinite-dimensional module in its fiber if and only if $n \geq 6$.*


The key to this proof is studying the algebra $U_\gamma^{(1)} = F_\lambda^{(1)}$ for γ a Gelfand–Tsetlin pattern, which we can write as a quotient of $e(\lambda)T^{n\omega_n-1}e(\lambda)$ by Theorem 5.6. Since simple modules in the fiber are in bijection with simple $U_\gamma^{(1)}$ -modules, and exactly one of these modules is finite-dimensional, the conjecture above holds if and only if $U_\gamma^{(1)}$ has only one simple module.

Proof. $[n \leq 5]$ In this case, we wish to prove that there are no other simple modules with γ in their support, that is, that $U_\gamma^{(1)}$ has a unique simple module. This will hold if the algebra has a non-negative grading, with degree 0 piece spanned by the scalars. By Proposition 4.8, as a module over the positively graded coinvariant algebra, the algebra $U_\gamma^{(1)}$ has a set of free generators indexed by the elements of W .


Thus, we need only confirm that any nontrivial element $w \in W = S_{n-1} \times \cdots \times S_2 \times S_1$ gives a generator of positive degree. In the cases where $n \leq 4$, this is easy to do by hand. For example, if $n = 4$, there are 12 elements of $S_3 \times S_2 \times S_1$. The resulting diagrams have degree 0, 2 or 4. The only diagram with degree 0 is:



The diagrams of degree 2 are:



The diagrams of degree 4 are:



For $n = 5$, there are 288 elements of $S_4 \times S_3 \times S_2 \times S_1$, so this is impractical to check by hand. We have checked this by computer calculation using SageMath; the code we

used can be found on the public GitHub repository for this paper. The Jupyter notebook `ComputeDegrees.ipynb` will guide you through the required computations. The identity is the only element of this group giving a generator of degree 1, while there are 29 of degree 2 and degree 8, 114 of degree 4 and of degree 6, and 1 of degree 10.


This confirms the conjecture for $n \leq 5$. In contrast, for $n = 6$, we find that the number of elements of degree 0, 2, 4, 6, 8, 10, 12, 14, 16, 18 is 2, 222, 2406, 8598, 12418, 8122, 2434, 338, and 20. The additional element of degree 0 in this case is the first hint that the conjecture will fail for higher values of n .

$[n \geq 6]$ To see that this conjecture fails in the case where $n \geq 6$, we need only find one γ where it fails for each n . Of course, since all Gelfand–Tsetlin patterns are in the same clan, and all choices of central character where Gelfand–Tsetlin patterns exist are equivalent by translation functors, the answer will be the same for all Gelfand–Tsetlin patterns for a fixed n .

First, consider the case $n = 6$. For concreteness, we choose one where the corresponding word is


$$\mathbf{j} = (6, 5, 4, 3, 2, 1, 6, 5, 4, 3, 2, 6, 5, 4, 3, 6, 5, 4, 6, 5, 6).$$

Consider the element $\tau = ((24), (12)(34), (13), 1, 1) \in S_5 \times S_4 \times S_3 \times S_2 \times S_1$. The corresponding diagram D is shown below:



The degree of D is 0: there are 8 crossings of strands with the same label, and 4 crossings of strands with adjacent labels for each of the pairs 6/5, 5/4, 4/3, and 3/2. We have verified by computer that it is the unique non-trivial diagram with degree 0 in the case $n = 6$.

In order to understand this case, the key calculation is to find D^2 , that is:



Of course, this is a complex calculation. It will be easier if we consider the action of D on the polynomial representation $\mathcal{P}$ of the KLRW algebra. It is simpler to use slightly different notation from earlier appearances of this representation, such as [26, 39, 41], so let us introduce this faithful module over $\tilde{T}_{\mathbf{v}}^{n\omega_{n-1}}$. We are following the conventions of [41, § 2.2.1] with changed notation—instead of having a single alphabet of variables, we separate them according to the labels on the corresponding strands. Let

$$S = \mathbb{k}[z_{i,j}]_{(i,j) \in \Omega} \quad \Omega = \{(i,j) \mid 1 \leq j \leq i \leq 6\}.$$

For a permutation σ , we let $\sigma^{(k)}$ denote the action of this permutation on the variables $z_{k,*}$. We define an action of the KLRW algebra $\tilde{T}_{\mathbf{v}}^{n\omega_{n-1}}$ on the sum $\mathcal{P} = \bigoplus_{\mathbf{i}} S \cdot 1_{\mathbf{i}}$ where

- $e(\mathbf{i})$ acts by projection to the corresponding summand $S \cdot 1_{\mathbf{i}}$,
- a dot on the k^{th} strand from the left with label i acts by multiplication by $z_{i,k}$,
- a crossing of the k^{th} and $k+1^{\text{st}}$ strands with $\mathbf{i}$ at the bottom and $\mathbf{i}'$ at top acts by
 - If $i_k = i_{k+1}$ and these are the r^{th} and $(r+1)^{\text{st}}$ strands with this label, the divided difference operator

$$f1_{\mathbf{i}} \mapsto \frac{f - (r, r+1)^{(i_k)} \cdot f}{z_{i_k, r} - z_{i_k, r+1}} 1_{\mathbf{i}'}$$

- If $i_k + 1 = i_{k+1}$ and these are the r^{th} and s^{th} strands with these labels, the multiplication

$$f1_{\mathbf{i}} \mapsto (z_{i_k, r} - z_{i_{k+1}, s}) f1_{\mathbf{i}'}$$

- Otherwise, the identity map

$$f1_{\mathbf{i}} \mapsto f1_{\mathbf{i}'}$$

By [41, Proposition 2.20], this representation is faithful. The action of the diagram D in this representation is given by $p\partial_{\tau}q$ where

$$\begin{aligned} p &= (z_{2,1} - z_{3,1})(z_{4,2} - z_{5,2})(z_{3,3} - z_{4,3})(z_{5,4} - z_{6,4}) \\ q &= (z_{5,2} - z_{6,3})(z_{3,1} - z_{4,2})(z_{4,3} - z_{5,4})(z_{2,2} - z_{3,3}) \\ \partial_{\tau} &= \partial_{(24)}^{(5)} \partial_{(12)(34)}^{(4)} \partial_{(13)}^{(3)} \end{aligned}$$

where $\partial_{\sigma}^{(j)}$ is the divided difference operator for this element of the symmetric group S_j . Since we are only considering the longest elements in different parabolic subgroups, we have

$$\begin{aligned} \partial_{(24)}^{(5)} &= \frac{1 - (23)^{(5)} - (34)^{(5)} - (24)^{(5)} + (234)^{(5)} + (243)^{(5)}}{\prod_{2 \leq r < s \leq 4} (z_{j,r} - z_{j,s})} \\ \partial_{(12)(34)}^{(4)} &= \frac{1 - (12)^{(4)} - (34)^{(4)} + (12)(34)^{(4)}}{(z_{4,1} - z_{4,2})(z_{4,3} - z_{4,4})} \\ \partial_{(13)}^{(3)} &= \frac{1 - (12)^{(3)} - (23)^{(3)} - (13)^{(3)} + (123)^{(3)} + (132)^{(3)}}{\prod_{1 \leq r < s \leq 3} (z_{j,r} - z_{j,s})} \end{aligned}$$

Thus, D^2 acts by $p\partial_{\tau}qp\partial_{\tau}q$. Since τ is the longest element of a parabolic subgroup, we find that

$$p\partial_{\tau}qp\partial_{\tau}q = p\partial_{\tau}q \cdot \partial_{\tau}(qp)$$

Thus, we have that $D^2 = \partial_{\tau}(qp)D$. Before calculating $\partial_{\tau}(qp)$, let us note that it is not too hard to rephrase this calculation in terms of the diagram above: The polynomial qp will be obtained by resolving all the bigons involving different colors on the center of the diagram (in terms of the dots), and $\partial_{\tau}(qp)$ will be obtained by resolving the bigons of the same color using [27, (2.20)] on the strands with label 3, those with label 5, and on the two crossing groups of the strands with label 4.

First, note that any term that includes $z_{2,*}$ or $z_{6,*}$ will be killed by this divided difference operator, so we can simply set these variables to 0. After this substitution, we have

$$\begin{aligned} \partial_{(12)(34)}^{(4)}(pq) &= z_{3,1}z_{5,2}z_{5,4}z_{3,3} \\ &\cdot \partial_{(12)(34)}^{(4)} \left(\left(z_{4,2}^2 - (z_{5,2} + z_{3,1})z_{4,2} + z_{5,2}z_{3,1} \right) \left(z_{4,3}^2 - (z_{5,4} + z_{3,3})z_{4,3} + z_{5,4}z_{3,3} \right) \right) \end{aligned}$$

Furthermore, note that only terms with degree 3 in the variables $z_{5,*}$ and $z_{3,*}$ will have nonzero image, so we have

$$\begin{aligned}\partial_\tau(qp) &= -\partial_{(24)}^{(5)} \partial_{(13)}^{(3)} (z_{3,1} z_{5,2} z_{5,4} z_{3,3} (z_{5,2} + z_{3,1}) (z_{5,4} + z_{3,3})) \\ &= -\partial_{(24)}^{(5)} (z_{5,2}^2 z_{5,4}) \partial_{(13)}^{(3)} (z_{3,1} z_{3,3}^2) - \partial_{(24)}^{(5)} (z_{5,2} z_{5,4}^2) \partial_{(13)}^{(3)} (z_{3,1}^2 z_{3,3}) \\ &= 2\end{aligned}$$

Thus, we have that $D^2 = 2D$. In particular, $D/2$ and $1 - D/2$ are orthogonal idempotents.

The same calculation addresses any value of $n > 6$ by considering the diagram where the strands with labels ≤ 6 trace out D , and all the others are straight vertical. Let us abuse notation and also denote this diagram D .

Let L_0 be the unique finite-dimensional $U(\mathfrak{gl}_n)$ -module such that $\ell_0 = \mathcal{W}_\gamma(L_0) \neq 0$. The image ℓ_0 is a one-dimensional module over $U_\gamma^{(1)}$ killed by all elements of non-zero degree. Of course, we have that $D\mathcal{W}_\gamma(L_0) = 0$ since D factors through weight spaces that don't correspond to Gelfand–Tsetlin patterns, which thus have trivial weight spaces for L_0 . However, since $D/2$ is idempotent, we must have a simple $U_\gamma^{(1)}$ -module ℓ_1 such that $D\ell_1 \neq 0$; in fact, any simple quotient of the projective $U_\gamma^{(1)}D$ will work. There is a unique corresponding simple $U(\mathfrak{gl}_n)$ -module L_1 such that $\mathcal{W}_\gamma(L_1) = \ell_1$. The module L_1 is necessarily infinite-dimensional, since L_0 is the unique finite-dimensional module with this infinitesimal character. $\square$

In the case $n = 6$, the algebra $U_\gamma^{(1)}$ has no elements of negative degree, so every positive degree element lies in the Jacobson radical. The quotient $\bar{U}_\gamma^{(1)}$ by the ideal of positive degree elements is two-dimensional and is spanned by the orthogonal idempotents $D/2$ and $1 - D/2$. That is, $\bar{U}_\gamma^{(1)} \cong \mathbb{C} \oplus \mathbb{C}$. This shows that $U_\gamma^{(1)}$ has exactly two simple modules, which are distinguished by whether D acts by zero. The simple ℓ_0 on which D acts trivially corresponds to a finite-dimensional $U(\mathfrak{gl}_6)$ -module L_0 and the simple ℓ_1 on which $1 - D/2$ acts trivially to an infinite-dimensional module L_1 . Recall that for a word $\mathbf{i}$, the canonical module $C(\mathbf{i})$ is the unique simple quotient of the submodule of the polynomial representation of the KLRW algebra generated by the image of $e(\mathbf{i})$; the canonical module for $U(\mathfrak{gl}_n)$ is the corresponding simple Gelfand–Tsetlin module. It's easy to check that L_0 is the canonical module of the word $w(\lambda)$, but we can also use this language to describe L_1

Lemma 6.3. *The module L_1 is the canonical module for the words*

$$\begin{aligned}\mathbf{i} &= (6, 5, 2, 4, 4, 1, 6, 3, 3, 3, 2, 6, 5, 5, 5, 4, 4, 6, 6, 5, 6) \\ \mathbf{i}' &= (2, 1, 3, 4, 5, 4, 3, 5, 6, 5, 4, 3, 2, 6, 6, 5, 4, 6, 6, 5, 6)\end{aligned}$$

and under the bijection of Theorem 5.9, the word $\mathbf{i}'$ is the corresponding red-good word.

Proof. Divide the diagram D in half by cutting at the line $y = \frac{1}{2} + \epsilon$ in diagrams D_1 above this line and D_2 below, so $D = D_1 D_2$. Note that the word $\mathbf{i}$ is obtained exactly by reading left to right on this horizontal line. Consider $D_1 1_{\mathbf{i}} \in \mathcal{P}$. This is non-zero since the proof that $D^2 = D$ also showed that $D_2 D_1 1_{\mathbf{i}} = \partial_\tau(qp) \cdot 1_{\mathbf{i}} = 2 \cdot 1_{\mathbf{i}}$. This shows that the canonical module $C(\mathbf{i})$ is infinite-dimensional (since $\mathbf{i}$ is not a Gelfand–Tsetlin pattern) and $\mathcal{W}_\gamma(C(\mathbf{i})) \neq 0$. Thus, this canonical module must be L_1 . We can find other words with the same canonical module by applying the rules of [41, Lemma 2.24]:

$$\begin{aligned} & C(6, 5, 2, 4, 4, 1, 6, 3, 3, 3, 2, 6, 5, 5, 5, 4, 4, 6, 6, 5, 6) \\ & \cong C(2, 1, 6, 5, 4, 4, 3, 3, 3, 2, 6, 6, 5, 5, 5, 4, 4, 6, 6, 5, 6) \\ & \cong C(2, 1, 4, 6, 5, 4, 3, 3, 3, 2, 6, 6, 5, 5, 5, 4, 4, 6, 6, 5, 6) \\ & \cong C(2, 1, 3, 4, 3, 6, 5, 4, 3, 2, 6, 6, 5, 5, 5, 4, 4, 6, 6, 5, 6) \\ & \cong C(2, 1, 3, 5, 4, 3, 5, 6, 5, 4, 3, 2, 6, 6, 5, 4, 4, 6, 6, 5, 6) \\ & \cong C(2, 1, 3, 4, 5, 4, 3, 5, 6, 5, 4, 3, 2, 6, 6, 5, 4, 6, 6, 5, 6) \end{aligned}$$

This last word is red-good, so by [41, Theorem 2.23], it is the red-good word of this module as desired. □

INDEX OF NOTATION

This index of notation gives a brief description of the main notation used in the paper, together with the section and page where the notation is defined.

§	Symbol	Description	Page
1.0	Λ	A Noetherian commutative ring.	393
	L	The fraction field of L .	393
	$\widehat{W}$	A monoid with a faithful action on Λ , which is the semi-direct product of a finite subgroup W and a submonoid $\mathcal{M}$.	393
	Γ	The invariants Λ^W .	393
	K	The fraction field of Γ .	393
	U	A principal Galois order.	394
1.1	$\widehat{U}_\gamma$	The algebra controlling the Gelfand–Tsetlin modules with non-zero γ weight space.	394
	F	A principal flag order.	394
	$\mathfrak{m}_\lambda$	A maximal ideal in Λ .	394
	$\widehat{F}_\lambda$	The endomorphism algebra of the weight functor $\mathcal{W}_\lambda$.	394
	$\widehat{W}_\lambda$	The stabilizer of λ in $\widehat{W}$.	394
	W_λ	The stabilizer of λ in W .	394
	$\widehat{\Lambda}_\lambda$	The invariants of $\widehat{W}_\lambda$ acting on $\widehat{\Lambda}$	394
	$F_\lambda^{(1)}$	The quotient of F_λ by the maximal ideal of $\widehat{\Lambda}$.	394
1.3	G	A reductive connected group.	395
	N	A representation of N .	395
	$\widetilde{T}_\mathbf{v}$	The KLRW algebra with v_i black strands.	396
2.0	$\mathcal{K}_\Gamma$	The standard order $\{X \in \mathcal{K} \mid X(\Gamma) = \Gamma\}$.	398
	$\mathcal{F}_\Lambda$	The standard flag order $\{X \in \mathcal{F} \mid X(\Lambda) = \Lambda\}$.	398
	F_D	A flag Galois order canonically constructed from U and D by considering $De \otimes_\Gamma U \otimes_\Gamma eD$.	399
2.1	$\mathcal{W}_\lambda$	The weight functor $\mathcal{W}_\lambda(M) = \{m \in M \mid \mathfrak{m}_\lambda^N m = 0 \text{ for some } N \gg 0\}$.	399
2.2	Λ_λ	The invariants $\Lambda^{\widehat{W}_\lambda}$.	400

§	Symbol	Description	Page
	${}_{\lambda}\widehat{W}_{\mu}$	The set of elements of $\widehat{W}$ such that $w \cdot \mu = \lambda$.	400
	${}_{\lambda}F_{\mu}$	The elements of F which are in the K -span of ${}_{\lambda}\widehat{W}_{\mu}$.	400
	$\mathfrak{n}_{\lambda}$	The maximal ideal $\mathfrak{n}_{\lambda} = \mathfrak{m}_{\lambda} \cap \Lambda_{\lambda}$	402
2.6	$F(\mathbf{S})$	The matrix algebra of (2.6).	407
	$\Gamma\mathbf{U}(\mathbf{S})$	The category of all Gelfand–Tsetlin modules modulo the subcategory of modules such that $\mathcal{W}_{\lambda}(M) = 0$ for all $\lambda \in \mathbf{S}$.	408
3.0	V	A vector space such that $\Lambda = \mathrm{Sym}^{\bullet}(V)$.	408
4.0	$\mathbb{k}$	A field, assumed to be of characteristic 0 through most of the paper.	410
4.1	G	A complex reductive group, often called the “gauge group.”	410
	N	A representation of G , often called the “matter representation.”	410
	$\mathcal{Y}$	The moduli space $\mathcal{Y} = (G((t)) \times N[[t]])/G[[t]]$ of principal bundles on a formal disk with choice of section of the associated bundle for N and of trivialization away from the origin.	410
	Q	A group acting on N such that $G \subset Q$ and Q/G is a torus.	410
	$\mathcal{A}$	The quantum Coulomb branch attached to the data G, N and Q (Definition 4.1)	410
	W	The common Weyl group of G and Q .	411
	I	The standard Iwahori $I \subset G[[t]]$	411
4.2	G_{λ}	The Levi subgroup of G which only contains the roots which are integral at λ .	411
	N_{λ}	The span in N of the weight spaces for weights integral on λ .	411
	B_{λ}	The Borel in G_{λ} generated by the roots α such that $\langle \lambda, \alpha \rangle$ is negative.	411
	N_{λ}^{-}	The subspace of N_{λ} which is non-positive for the cocharacter corresponding to λ .	411
	X_{λ}	The associated vector bundle $(G_{\lambda} \times N_{\lambda}^{-})/B_{\lambda}$.	411
	${}_{\lambda}\mathbb{X}_{\mu}$	The generalized Steinberg variety $X_{\lambda} \times_{N_{\lambda}} X_{\mu}$.	411
	$\mathbb{T}$	The 1-parameter subgroup of $G \times \mathbb{C}^*$ obtained by exponentiating $\lambda \in \mathrm{MaxSpec}(\Lambda)$.	412
4.5	GKdim	The Gelfand–Kirillov dimension $\limsup_{n \rightarrow \infty} \log_n \dim_{\mathbb{k}}(A_0^n M_0)$.	418
5.1	$\mathbf{v}$	The dimension vector defining the OGZ algebra.	420
	Ω	The set $\{(i, r) \mid 1 \leq i \leq n, 1 \leq r \leq v_i\}$.	420
5.2	$w(\lambda)$	The word in $[1, n]$ given by ordering the elements of Ω according to the weight λ and then projecting to the first index. (Definition 5.5)	423
5.3	$\mathbb{L}(S)$	The set of words $w(\lambda)$ for $\lambda \in \mathcal{S}_{\mathbb{Z}}$ such that $\mathcal{W}_{\lambda}(S) \neq 0$.	423
	$\mathcal{GL}$	The set of words of the form $(k, k - 1, \dots, k - p)$ for $k \leq n - 1$, and $0 \leq p < k$.	424
	$\mathcal{GL}'$	The set of words of the form $(n, n - 1, \dots, n - p)$ for $0 \leq p < n$.	424

6.0	D	The KLRW diagram in $U_\gamma^{(1)}$ induced by the permutation $\tau = ((24), (12)(34), (13), 1, 1)$.	430
-----	-----	---	-----

ACKNOWLEDGMENTS

A great number of people deserve acknowledgment in the creation of this paper: my collaborators Joel Kamnitzer, Peter Tingley, Alex Weekes, and Oded Yacobi, since this project grew out of our joint work; Volodymyr Mazorchuk, who first suggested to me that a connection existed between our previous work and the study of Gelfand–Tsetlin modules, and who, along with Elizaveta Vishnyakova, Jonas Hartwig, Slava Futorny, Dimitar Grantcharov, and Pablo Zadunaisky, suggested many references and improvements; Turner Silverthorne, who collaborated with me on the computer program used in some of the calculations in this paper; Hiraku Nakajima who, amongst many other things, pointed out to me the argument used in the proof of Theorem 4.4; and several anonymous referees. In particular, prodding from a referee helped me to find the mistake in the “proof” of Conjecture 6.1.

REFERENCES

- [1] Alexander Beilinson, Victor Ginzburg, and Wolfgang Soergel, *Koszul duality patterns in representation theory*, J. Am. Math. Soc. **9** (1996), no. 2, 473–527.
- [2] Tom Braden, Anthony Licata, Nicholas Proudfoot, and Ben Webster, *Hypertoric category O* , Adv. Math. **231** (2012), no. 3–4, 1487–1545.
- [3] Alexander Braverman, Michael Finkelberg, and Hiraku Nakajima, *Towards a mathematical definition of Coulomb branches of 3-dimensional $\mathcal{N} = 4$ gauge theories. II*, Adv. Theor. Math. Phys. **22** (2018), no. 5, 1071–1147.
- [4] ———, *Coulomb branches of 3d $\mathcal{N} = 4$ quiver gauge theories and slices in the affine Grassmannian*, Adv. Theor. Math. Phys. **23** (2019), no. 1, 75–166.
- [5] Michel Broué, *Introduction to complex reflection groups and their braid groups*, Lecture Notes in Mathematics, vol. 1988, Springer, 2010.
- [6] Jonathan Brundan and Alexander Kleshchev, *Shifted Yangians and finite W -algebras*, Adv. Math. **200** (2006), no. 1, 136–195.
- [7] Mathew Bullimore, Tudor Dimofte, Davide Gaiotto, and Justin Hilburn, *Boundaries, mirror symmetry, and symplectic duality in 3d $\mathcal{N} = 4$ gauge theory*, J. High Energy Phys. **10** (2016), Paper no. 108.
- [8] Neil Chriss and Victor Ginzburg, *Representation theory and complex geometry*, Birkhäuser, 1997.
- [9] Yu. A. Drozd, Vyacheslav M. Futorny, and Serge A. Ovsienko, *Harish–Chandra subalgebras and Gelfand–Zetlin modules*, in Finite-dimensional algebras and related topics (Ottawa, ON, 1992), NATO ASI Series. Series C. Mathematical and Physical Sciences, vol. 424, Kluwer Academic Publishers, 1994, pp. 79–93.
- [10] Nick Early, Volodymyr Mazorchuk, and Elizaveta Vishnyakova, *Canonical Gelfand–Zeitlin Modules over Orthogonal Gelfand–Zeitlin Algebras*, Int. Math. Res. Not. **2020** (2020), no. 20, 6947–6966.
- [11] Ben Elias and Geordie Williamson, *The Hodge theory of Soergel bimodules*, Ann. Math. **180** (2014), no. 3, 1089–1136.
- [12] Dylan Fillmore, *On the Category of Harish–Chandra Block Modules*, 2023, <https://arxiv.org/abs/2304.06448>.
- [13] Vyacheslav Futorny, Dimitar Grantcharov, and Luis Enrique Ramirez, *Singular Gelfand–Tsetlin modules of $gl(n)$* , Adv. Math. **290** (2016), 453–482.
- [14] ———, *Gelfand–Tsetlin modules of $sl(3)$ in the principal block*, in Representations of Lie algebras, quantum groups and related topics, Contemporary Mathematics, vol. 713, American Mathematical Society, 2018, pp. 121–134.
- [15] Vyacheslav Futorny, Dimitar Grantcharov, Luis Enrique Ramirez, and Pablo Zadunaisky, *Bounds of Gelfand–Tsetlin multiplicities and tableaux realizations of Verma modules*, J. Algebra **556** (2020), 412–436.

- [16] ———, *Gelfand–Tsetlin theory for rational Galois algebras*, Isr. J. Math. **239** (2020), no. 1, 99–128.
- [17] Vyacheslav Futorny and Serge A. Ovsienko, *Galois orders in skew monoid rings*, J. Algebra **324** (2010), no. 4, 598–630.
- [18] ———, *Fibers of characters in Gelfand–Tsetlin categories*, Trans. Am. Math. Soc. **366** (2014), no. 8, 4173–4208.
- [19] I. M. Gelfand and M. L. Cetlin, *Finite-dimensional representations of the group of unimodular matrices*, Dokl. Akad. Nauk SSSR, n. Ser. **71** (1950), 825–828.
- [20] Victor Ginzburg, *Nil–Hecke algebras and Whittaker $\mathfrak{D}$ -Modules*, in Lie Groups, Geometry, and Representation Theory, Springer, 2018, pp. 137–184.
- [21] Mark Goresky, Robert Kottwitz, and Robert MacPherson, *Equivariant cohomology, Koszul duality, and the localization theorem*, Invent. Math. **131** (1998), no. 1, 25–83.
- [22] Jonas T. Hartwig, *Principal Galois orders and Gelfand–Zeitlin modules*, Adv. Math. **359** (2020), Paper no. 106806.
- [23] Joel Kamnitzer, Peter Tingley, Ben Webster, Alex Weekes, and Oded Yacobi, *On category O for affine Grassmannian slices and categorified tensor products*, Proc. Lond. Math. Soc. **119** (2019), no. 5, 1179–1233.
- [24] Joel Kamnitzer, Ben Webster, Alex Weekes, and Oded Yacobi, *Lie algebra actions on module categories for truncated shifted yangians*, Forum Math. Sigma **12** (2024), Paper no. e18.
- [25] Mikhail Khovanov and Aaron D. Lauda, *A diagrammatic approach to categorification of quantum groups. I*, Represent. Theory **13** (2009), 309–347.
- [26] ———, *A diagrammatic approach to categorification of quantum groups. II*, Trans. Am. Math. Soc. **363** (2011), no. 5, 2685–2700.
- [27] Mikhail Khovanov, Aaron D. Lauda, Marco Mackaay, and Marko Stošić, *Extended graphical calculus for categorified quantum $sl(2)$* , Memoirs of the American Mathematical Society, vol. 1029, American Mathematical Society, 2012.
- [28] Alexander Kleshchev and Arun Ram, *Representations of Khovanov–Lauda–Rouquier algebras and combinatorics of Lyndon words*, Math. Ann. **349** (2011), no. 4, 943–975.
- [29] Ryosuke Kodera and Hiraku Nakajima, *Quantized Coulomb branches of Jordan quiver gauge theories and cyclotomic rational Cherednik algebras*, in String-Math 2016, Proceedings of Symposia in Pure Mathematics, vol. 98, American Mathematical Society; International Press, 2018, pp. 49–78.
- [30] Bernard Leclerc, *Dual canonical bases, quantum shuffles and q -characters*, Math. Z. **246** (2004), no. 4, 691–732.
- [31] Elise LePage and Ben Webster, *Rational Cherednik algebras of $G(\ell, p, n)$ from the Coulomb perspective*, Adv. Math. **433** (2023), Paper no. 109295.
- [32] Volodymyr Mazorchuk, *Personal communication*.
- [33] ———, *Orthogonal Gelfand–Zeitlin algebras. I*, Beitr. Algebra Geom. **40** (1999), no. 2, 399–415.
- [34] Volodymyr Mazorchuk and Elizaveta Vishnyakova, *Harish–Chandra modules over invariant subalgebras in a skew-group ring*, Asian J. Math. **25** (2021), no. 3, 431–454.
- [35] Alexander I. Molev, *Gelfand–Tsetlin Bases for Classical Lie Algebras*, in Handbook of Algebra, vol. 4, Elsevier, 2006, pp. 109–170.
- [36] Ian M. Musson and Michel Van den Bergh, *Invariants under tori of rings of differential operators and related topics*, Memoirs of the American Mathematical Society, vol. 650, American Mathematical Society, 1998.
- [37] Hiraku Nakajima, *Towards a mathematical definition of Coulomb branches of 3-dimensional $N = 4$ gauge theories. I*, Adv. Theor. Math. Phys. **20** (2016), no. 3, 595–669.
- [38] Luis E. Ramirez and Pablo Zadunaisky, *Gelfand–Tsetlin modules over $\mathfrak{gl}(n)$ with arbitrary characters*, J. Algebra **502** (2018), 328–346.
- [39] Raphael Rouquier, *2-Kac–Moody algebras*, 2008, <https://arxiv.org/abs/0812.5023>.
- [40] Julia Sauter, *A Survey on Springer Theory*, 2013, <https://arxiv.org/abs/1307.0973>.
- [41] Turner Silverthorne and Ben Webster, *Gelfand–Tsetlin Modules: Canonicity and Calculations*, Algebr. Represent. Theory **27** (2024), no. 2, 1405–1455.
- [42] Michela Varagnolo and Eric Vasserot, *Double affine Hecke algebras and affine flag manifolds, I*, in Affine flag manifolds and principal bundles, Trends in Mathematics, Springer, 2010, pp. 233–289.
- [43] Elizaveta Vishnyakova, *A geometric approach to 1-singular Gelfand–Tsetlin $\mathfrak{gl}_n$ -modules*, Differ. Geom. Appl. **56** (2018), 155–160.

- [44] Ben Webster, *Coherent sheaves and quantum Coulomb branches I: tilting bundles from integrable systems*, <https://arxiv.org/abs/1905.04623>.
- [45] ———, *Canonical bases and higher representation theory*, Compos. Math. **151** (2015), no. 1, 121–166.
- [46] ———, *Knot invariants and higher representation theory*, Memoirs of the American Mathematical Society, vol. 1191, American Mathematical Society, 2017.
- [47] ———, *Koszul duality between Higgs and Coulomb categories $\mathcal{O}$* , 2019, <https://arxiv.org/abs/1611.06541>.
- [48] ———, *Weighted Khovanov–Lauda–Rouquier algebras*, Doc. Math. **24** (2019), 209–250.
- [49] ———, *Three perspectives on categorical symmetric Howe duality*, 2020, <https://arxiv.org/abs/2001.07584>.
- [50] Ben Webster, Alex Weekes, and Oded Yacobi, *A quantum Mirković–Vybornov isomorphism*, Represent. Theory **24** (2020), no. 2, 38–84.
- [51] Alex Weekes, *Generators for Coulomb branches of quiver gauge theories*, 2019, <https://arxiv.org/abs/1903.07734>.
- [52] Dmitry P. Zhelobenko, *The classical groups: spectral analysis of their finite-dimensional representations*, Russ. Math. Surv. **17** (1962), no. 1, 1–94.

— BEN WEBSTER —

URL: <https://uwaterloo.ca/scholar/b2webste>

DEPARTMENT OF PURE MATHEMATICS, UNIVERSITY OF WATERLOO & PERIMETER INSTITUTE FOR THEORETICAL PHYSICS, CANADA

E-mail address: ben.webster@uwaterloo.ca