COMPOSITIO MATHEMATICA

Canonical bases and higher representation theory

Ben Webster

Compositio Math. **151** (2015), 121–166.

doi:10.1112/S0010437X1400760X

Compositio Math. 151 (2015) 121–166
doi:10.1112/S0010437X1400760X

Canonical bases and higher representation theory

Ben Webster

Abstract

This paper develops a general theory of canonical bases and how they arise naturally in the context of categorification. As an application, we show that Lusztig's canonical basis in the whole quantized universal enveloping algebra is given by the classes of the indecomposable 1-morphisms in a categorification when the associated Lie algebra is of finite type and simply laced. We also introduce natural categories whose Grothendieck groups correspond to the tensor products of lowest- and highest-weight integrable representations. This generalizes past work of the author's in the highest-weight case.

Contents

1 **Pre-canonical structures and canonical bases** 125
2 **Dual canonical bases** 133
3 **The 2-category $\mathcal{U}$** 136
4 **The 2-category $\mathcal{T}$** 141
5 **Tensor product algebras** 147
6 **Representation categories and standard modules** 153
7 **Orthodox bases in higher representation theory** 157
8 **Canonical bases in higher representation theory** 161
Acknowledgements 164
References 164

One of the consistent motivations for the construction of categorifications has been the accompanying appearance of canonical bases in the original object under consideration. At its core, this is a consequence of a very simple principle: the indecomposable objects of any Krull–Schmidt category give a basis of its split Grothendieck group. Furthermore, any map between Grothendieck groups which lifts to a functor must have positive integer coefficients in this basis.

While this positivity is an appealing consequence, on its own, it has trouble making up for the difficulty of computing this basis in many situations. For example, irreducible characters give a basis of class functions on a finite group where multiplication has positive integral structure coefficients, but finding irreducible characters is still very hard in general.

This computational problem eases if this basis has the additional property of being *canonical*; we will make precise in § 1 what this means. Defining a canonical basis requires a choice of *pre-canonical structure*, which consists of a bar-involution, a sesquilinear pairing and a

Received 12 January 2013, accepted in final form 29 April 2014, published online 7 October 2014.
2010 Mathematics Subject Classification 17B35, 17B37, 18D05 (primary).
Keywords: canonical bases, categorification.
This work was supported by the NSF under grant DMS-1151473 and by the NSA under grant H98230-10-1-0199.

'first approximation' to our basis. Once this data is chosen, there can be at most one canonical basis, though it may be that none exists. Many readers will be familiar with how these elements induce a canonical basis from work of Kazhdan, Lusztig and others [KL79, Lus93], but we will give a general account making this definition precise. If we can prove that the basis coming from our categorification is canonical (using categorical properties), then we can reduce its construction to the Gram–Schmidt algorithm.

It is worth noting that the term 'canonical basis' had not (to the author's knowledge) had a precise meaning in the past, but was applied to a variety of cases with properties in common. We shall give a formal definition of a *canonical basis* (Definition 1.7) which recovers most examples that we are aware of; most importantly, it will recover the canonical bases defined by Lusztig on the modified quantum universal enveloping algebra $\dot{U}$ and tensor products in the cases where they are defined in [Lus93]. To avoid confusion, in this paper we will refer to the bases defined in [Lus93] as *Lusztig's bases* and use the term 'canonical' to mean only bases satisfying Definition 1.7. See Theorem 7.4 for the precise connection between these types of bases.

In the first two sections, we develop the theory of *humorous categories*, which are categories well-suited to a connection with a canonical basis. These categories always have an *orthodox*[1] basis arising from their indecomposable objects. This basis will be canonical when the category satisfies an additional condition on positivity of gradings, which we call *mixedness* by analogy with mixed structures in geometry. In particular, we will show (in Lemmata 1.17 and 1.18) how information about mixedness and canonical bases can pass back and forth between categories and certain special quotients, and give a useful general principle for extracting mixed humorous categories from t-structures on dg-categories (Lemma 1.20). All of these lemmata are key to understanding the canonical basis of $\dot{U}$ in categorical terms.

The aim of the rest of the paper is to apply this theory to give an account of the canonical bases arising in quantized universal enveloping algebras and their representations. Our primary tool will be higher representation theory, as developed by Rouquier, Khovanov, Lauda and others. We shall give a brief reminder about the categorification $\dot{\mathcal{U}}$ of the universal enveloping algebra itself, and then define a categorification $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ of a tensor product of a sequence of highest and lowest integrable representations. However, this sequence cannot be in an arbitrary order; in effect, the categorification forces us to put lowest-weight representations on the left and highest-weight representations on the right. The significance of this condition is not clear at present, but it matches the underlying algebra and combinatorics of these representations. This can be seen in [BW14], where the existence of a canonical basis in precisely these tensor products is proved. These latter categories are generalizations of the categorifications of highest-weight representations defined by the author in [Web13a].

We build on very important results of Vasserot and Varagnolo [VV11, Theorem 4.5] to show the following.

Theorem A (Theorems 8.8 and 8.11).

(a) *If $\mathfrak{g}$ is an arbitrary Kac–Moody algebra with symmetric Cartan matrix, the canonical basis of a tensor product of highest-weight integrable representations coincides with the classes of indecomposable objects in the categorification $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ defined in § 5.*

(b) *If $\mathfrak{g}$ is finite-dimensional and simply laced (that is, of ADE type), then the canonical basis of the modified quantized universal enveloping algebra $\dot{U}$ coincides with the classes of indecomposable objects in the categorification $\dot{\mathcal{U}}$ defined in § 3.*

[1] The word 'orthodox' comes from the Greek ὀρθός ('correct') + δόξα ('belief'); it is a basis we can believe in.

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

Let us reiterate: here 'canonical basis' refers to Definition 1.7, but these bases agree with Lusztig's in all cases where the latter are defined by Theorem 7.4.

Here is a brief outline of the proof of Theorem A. The general theory of § 1 allows us to reduce to the question of whether the accompanying categorifications $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ and $\mathcal{U}$ are mixed. First, using results relating $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ to the geometry of perverse sheaves on certain quiver varieties from [Web12b] (building on older results of Varagnolo and Vasserot [VV11]), we prove that $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ is mixed using Lemma 1.17. This shows Theorem A(a).

In the finite-type case, highest- and lowest-weight modules coincide, and a similar property holds for categorifications (Proposition 5.7). This allows us to show that $\mathcal{X}^{w_0\lambda,\mu}$ is mixed when both λ and μ are dominant. Lemma 1.18 allows us to conclude that $\mathcal{U}$ is mixed as well, since $\mathcal{X}^{w_0\lambda,\mu}$ forms a series of quotients that jointly capture all the structure in $\mathcal{U}$. This establishes Theorem A(b).

As indicated earlier in the introduction, this sort of result is particularly interesting because of its positivity consequences.

Corollary B. *If $\mathfrak{g}$ is finite-dimensional and simply laced, then the structure coefficients of multiplication of the canonical basis of $\dot{U}$ are all Laurent polynomials with positive integer coefficients. The same holds for matrix coefficients in the canonical basis of its action on any tensor product of finite-dimensional modules.*

Let us indicate why the hypotheses in Theorem A are necessary. If the Cartan matrix is not symmetric, then Theorem A simply will not hold: the categorifications we use exist, but their indecomposables do not give a canonical basis, since we know that the analogue of Corollary B fails.

If the Cartan matrix is symmetric but of infinite type, Theorem A(b) applies, but at the moment we know no proof of the analogue of Theorem A(a). It seems likely that this categorical interpretation of the canonical basis will hold for both $\dot{U}$ and arbitrary tensor products of a group of lowest-weight representations with a group of highest-weight representations. Later in this paper, we will define pre-canonical structures on these spaces. However, the techniques in the proof of Theorem A do not suffice to prove that a canonical basis exists in this case, let alone that such a basis arises from a categorification. The proof of Theorem A(b) uses that highest- and lowest-weight modules of $\dot{U}$ are the same in a very strong way, so it cannot proceed here.

A proof of Theorem A(b) in infinite symmetric type will require very different techniques. One promising approach would be to apply Lemma 1.18 to the categorical actions on quantizations of quiver varieties described in [Web12a]. However, in order to do so, we must prove that certain functors are full, and this fullness is equivalent to Kirwan surjectivity for quiver varieties. This is a long-standing open problem, so until it finds a solution, we cannot use this approach. See [JKK09] for a more extensive discussion of the Kirwan surjectivity problem for a general hyperkähler quotient.

Even if this approach is successful, it can only be applied to the canonical basis of $1_\lambda\dot{U}1_{\lambda'}$ where λ and λ' both lie in the positive or negative Tits cone $\mathfrak{g}$ (the union of the Weyl group orbits of the dominant or anti-dominant Weyl chamber). For example, for affine Lie algebras, it would exclude the case where λ or λ' has level 0. Another approach which could potentially apply in all cases is to use the techniques of [EW12] in this context, but it remains to be seen if such an approach can be successful.

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

Notation. We let $\mathfrak{g}$ be a symmetrizable Kac–Moody algebra. Consider the weight lattice $Y(\mathfrak{g})$ and root lattice $X(\mathfrak{g})$, and the simple roots α_i and coroots $\alpha_i^\vee$. We let Y^+ denote the set of dominant weights and Y^- the anti-dominant weights. Let $c_{ij} = \alpha_i^\vee(\alpha_j)$ be the entries of the Cartan matrix.

We let $\langle -, - \rangle$ denote the symmetrized inner product on $Y(\mathfrak{g})$, fixed by the fact that the shortest root has length $\sqrt{2}$ and

$$2\frac{\langle \alpha_i, \lambda \rangle}{\langle \alpha_i, \alpha_i \rangle} = \alpha_i^\vee(\lambda).$$

As usual, we let $2d_i = \langle \alpha_i, \alpha_i \rangle$, and for $\lambda \in Y(\mathfrak{g})$ we let

$$\lambda^i = \alpha_i^\vee(\lambda) = \langle \alpha_i, \lambda \rangle / d_i.$$

Throughout the paper, we will use $\underline{\boldsymbol{\lambda}} = (\lambda_1, \dots, \lambda_\ell)$ to denote an ordered ℓ-tuple of dominant or anti-dominant weights, and we will always use the notation $\lambda = \sum_i \lambda_i$.

We let $U_q(\mathfrak{g})$ denote the deformed universal enveloping algebra of $\mathfrak{g}$, that is, the associative $\mathbb{C}(q)$-algebra given by generators E_i and F_i for α_i a simple root and K_μ with $\mu \in Y(\mathfrak{g})$, subject to the following relations:

(i) $K_0 = 1$ and $K_\mu K_{\mu'} = K_{\mu+\mu'}$ for all $\mu, \mu' \in Y(\mathfrak{g})$;
(ii) $K_\mu E_i = q^{\alpha_i^\vee(\mu)} E_i K_\mu$ for all $\mu \in Y(\mathfrak{g})$;
(iii) $K_\mu F_i = q^{-\alpha_i^\vee(\mu)} F_i K_\mu$ for all $\mu \in Y(\mathfrak{g})$;
(iv) $E_i F_j - F_j E_i = \delta_{ij}(\tilde{K}_i - \tilde{K}_{-i})/(q^{d_i} - q^{-d_i})$, where $\tilde{K}_{\pm i} = K_{\pm d_i \alpha_i}$;
(v) for all $i \neq j$,

$$\sum_{a+b=-c_{ij}+1} (-1)^a E_i^{(a)} E_j E_i^{(b)} = 0 \quad \text{and} \quad \sum_{a+b=-c_{ij}+1} (-1)^a F_i^{(a)} F_j F_i^{(b)} = 0,$$

where $E_i^{(a)} = E_i^a/[a]_q!$ and $[a]_q! = [a]_q[a-1]_q \cdots$ for $[a]_q = (q^a - q^{-a})/(q - q^{-1})$.

This is a Hopf algebra with coproduct on Chevalley generators given by

$$\Delta(E_i) = E_i \otimes 1 + \tilde{K}_i \otimes E_i, \quad \Delta(F_i) = F_i \otimes \tilde{K}_{-i} + 1 \otimes F_i$$

and with the antipode on these generators defined by $S(E_i) = -\tilde{K}_{-i}E_i, S(F_i) = -F_i\tilde{K}_i$. We shall also need to use the opposite coproduct

$$\Delta^{\mathrm{op}}(E_i) = E_i \otimes \tilde{K}_i + 1 \otimes E_i, \quad \Delta(F_i) = F_i \otimes 1 + \tilde{K}_{-i} \otimes F_i.$$

Consider the Cartan involution

$$\omega(E_i) = F_i, \quad \omega(F_i) = E_i, \quad \omega(K_i^{\pm 1}) = K_i^{\mp 1}$$

on $\dot{U}$, and note that this involution intertwines the usual and opposite coproducts, $(\omega \otimes \omega) \circ \Delta \circ \omega = \Delta^{\mathrm{op}}$.

We let $U_q^{\mathbb{Z}}(\mathfrak{g})$ denote the Lusztig (divided powers) integral form generated over $\mathbb{Z}[q, q^{-1}]$ by $E_i^{(n)}$, $F_i^{(n)}$ and K_μ for $n \geqslant 1$. The integral form of the representation of highest (respectively, lowest) weight λ if λ is dominant (respectively, anti-dominant) will be denoted by $V_\lambda^{\mathbb{Z}}$. It will be natural for us to use a slightly unusual convention for our tensor products. When we tensor on the right with a highest-weight representation, we use the usual coproduct; but when we tensor with a lowest-weight representation, then more precisely let $\boldsymbol{\epsilon}$ be the sign vector such that

$\epsilon_k = -1$ if λ_k is dominant and $\epsilon_k = 1$ if it is anti-dominant, and let $\otimes^{-1}$ be the tensor product with the usual coproduct and $\otimes^1$ the tensor product using Δ^{op}. Then

$$V^{\mathbb{Z}}_{\underline{\boldsymbol{\lambda}}} = (\cdots((V^{\mathbb{Z}}_{\lambda_1} \otimes^{\epsilon_2}_{\mathbb{Z}[q,q^{-1}]} V^{\mathbb{Z}}_{\lambda_2}) \otimes^{\epsilon_3}_{\mathbb{Z}[q,q^{-1}]} V^{\mathbb{Z}}_{\lambda_3}) \otimes^{\epsilon_4}_{\mathbb{Z}[q,q^{-1}]} \cdots) \otimes^{\epsilon_\ell}_{\mathbb{Z}[q,q^{-1}]} V^{\mathbb{Z}}_{\lambda_\ell}.$$

If we let $i_1, \dots, i_p$ be the indices with $\epsilon_{i_k} = 1$ and $j_1, \dots, j_q$ the indices with $\epsilon_{j_k} = -1$, then this is isomorphic to the tensor product

$$V^{\mathbb{Z}}_{\underline{\boldsymbol{\lambda}}} \cong V^{\mathbb{Z}}_{\lambda_{i_p}} \otimes_{\mathbb{Z}[q,q^{-1}]} \cdots \otimes_{\mathbb{Z}[q,q^{-1}]} V^{\mathbb{Z}}_{\lambda_{i_1}} \otimes_{\mathbb{Z}[q,q^{-1}]} V^{\mathbb{Z}}_{\lambda_{j_1}} \otimes_{\mathbb{Z}[q,q^{-1}]} \cdots \otimes_{\mathbb{Z}[q,q^{-1}]} V^{\mathbb{Z}}_{\lambda_{j_q}}$$

using only the usual coproduct. We let $\bar{V}^{\mathbb{Z}}_{\underline{\boldsymbol{\lambda}}}$ denote the reduction of $V^{\mathbb{Z}}_{\underline{\boldsymbol{\lambda}}}$ at $q = 1$. Note that since $\Delta \equiv \Delta^{\mathrm{op}} \pmod{q-1}$, these ordering issues are unimportant after this specialization.

1. Pre-canonical structures and canonical bases

1.1 Humorous categories

We let V be a free $\mathbb{Z}[q, q^{-1}]$-module.

Definition 1.1. A *pre-canonical structure* on V is a triple consisting of:

- a 'bar-involution' $\psi : V \to V$ which is $\mathbb{Z}[q, q^{-1}]$-anti-linear;
- a sesquilinear inner product $\langle -, - \rangle : V \times V \to \mathbb{Z}((q^{-1}))$, for which ψ is flip-unitary, i.e.

$$\langle u, v \rangle = \langle \psi(v), \psi(u) \rangle;$$

- a 'standard basis' a_c with partially ordered index set $(C, <)$ such that

$$\psi(a_c) \in a_c + \sum_{c' < c} \mathbb{Z}[q, q^{-1}] \cdot a_{c'}.$$

Pre-canonical structures arise as shadows of categorical structures. One of our tasks will be to describe the sorts of categories that will interest us. Fix a field $\Bbbk$.

Definition 1.2. A *humorous category* is an additive $\Bbbk$-linear Krull–Schmidt idempotent complete category $\mathcal{C}$ equipped with:

- an invertible grading shift functor (1) such that the induced $\mathbb{Z}$-action on the set of indecomposables is free with finitely many orbits (we denote the set of such orbits by C); and
- a duality functor $M \mapsto M^{\circledast}$ which satisfies $(M(1))^{\circledast} \cong M^{\circledast}(-1)$ and has a unique fixed point $P_c \cong P_c^{\circledast}$ on each orbit $c \in C$.

These must satisfy the following conditions.

- $\mathrm{Hom}_{\mathcal{C}}(M, N(i)) = 0$ for $i \ll 0$, and $\dim \mathrm{Hom}_{\mathcal{C}}(M, N) < \infty$ for all objects M and N.
- The local ring $\mathrm{Hom}(P_c, P_c)$ has residue field $\Bbbk$ for all c (as opposed to a general division algebra), i.e. P_c is absolutely indecomposable over $\Bbbk$; it remains indecomposable under any finite-degree field extension of $\Bbbk$.

We can call a morphism $M \to N(i)$ a morphism of degree i. As in [KL09], we let HOM $(M, N) \cong \bigoplus_{i \in \mathbb{Z}} \mathrm{Hom}(M, N(i))$ and call elements of this space *degraded morphisms*. The reader might prefer to think of the *degraded category* $\mathcal{C}_{\mathbb{Z}}$ with the same objects as $\mathcal{C}$ and morphisms given by $\mathrm{HOM}(M, N)$ as graded $\Bbbk$-vector spaces, but it is usually more convenient to work with $\mathcal{C}$.

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

We say that a graded $\Bbbk$-algebra A equipped with anti-automorphism $\gamma : A \to A$ is *humorous* if:

- each graded piece of A is finite-dimensional, and the grades that appear in A are bounded below;
- every indecomposable graded A-module is absolutely indecomposable;
- the anti-automorphism $\gamma : A \to A$ preserves at least one primitive idempotent in each isomorphism class.

Lemma 1.3. *A category (with additional data) is humorous if and only if it is equivalent to the category of finitely generated projective graded left modules over a humorous $\Bbbk$-algebra A with morphisms given by homogeneous maps of degree* 0.

In this case, the humorous structure is given by:

- *the grading shift $M(i)$, which is the same module as M but with all gradings decreased by i;*
- *the duality $\circledast$ such that $M^{\circledast} = \mathrm{HOM}_A(M, A)$, with the anti-automorphism γ used to switch this right module to a left module.*

Proof. Most of the proof consists of simply applying definitions, so we will only give a sketch. First, let $\mathcal{C}$ be a humorous category; this contains a self-dual object O in which every indecomposable module appears as a summand with multiplicity 1. The algebra A is the opposite algebra of graded endomorphisms of O, that is, $A^{\mathrm{op}} \cong \mathrm{HOM}_{\mathcal{C}}(O, O)$, with the equivalence given by $M \mapsto \mathrm{HOM}_{\mathcal{C}}(O, M)$. Commutation with grading shift is clear. Of course,

$$\mathrm{HOM}_{\mathcal{C}}(O, M^{\circledast}) \cong \mathrm{HOM}_{\mathcal{C}}(M, O) \cong \mathrm{HOM}_A(\mathrm{HOM}_{\mathcal{C}}(O, M), A),$$

so this shows the commutation with duality. The anti-involution γ is induced by any isomorphism $O \cong O^{\circledast}$; this must preserve each primitive idempotent since there is only one in each isomorphism class.

Now, assume that A is a $\Bbbk$-algebra satisfying our conditions. Its category of representations is Krull–Schmidt and has finitely many indecomposables (up to shift), since A_0 is finite-dimensional. Since every projective module is a summand of $A^{\oplus n}$, we have that the finiteness of Hom spaces between projectives follows from the finiteness of the grades of A.

Finally, if we define $M^{\circledast} = \mathrm{HOM}_A(M, A)$ with this right module turned into a left module using γ, then this obviously commutes with grading shift and sends $(Ae)^{\circledast} \cong A\gamma(e)$. Thus we can define P_c to be the image of any γ-invariant primitive idempotent in the corresponding isomorphism class. □

The reader may rightly wonder about the full category of finite-dimensional graded modules; this appears as the category of representations of $\mathcal{C}$.

Definition 1.4. Let $\mathrm{Rep}(\mathcal{C})$ be the category of additive functors from $\mathcal{C}$ to the category of finite-dimensional $\Bbbk$-vector spaces. One can easily see that if $\mathcal{C}$ consists of the projective modules over A, then $\mathrm{Rep}(\mathcal{C})$ is the category of finite-dimensional modules over A.

There is a duality $\star$ on the abelian category $\mathrm{Rep}(\mathcal{C})$, defined by the property that

$$\mathrm{Hom}(P^{\circledast}, M) = \mathrm{Hom}(P, M^{\star})^{*}. \tag{1.1}$$

If $\mathrm{Rep}(\mathcal{C})$ is the category of finite-dimensional modules over A, then this is simply vector space duality, with γ used to switch left and right modules as before.

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

A humorous structure on a category induces certain natural structures on the Grothendieck group $K^0(\mathcal{C})$. For example:

- the grading shift functor induces an invertible endomorphism of $K^0(\mathcal{C})$ which we call q, making $K^0(\mathcal{C})$ a $\mathbb{Z}[q,q^{-1}]$-module, i.e. $q[M]=[M(1)]$;
- the decategorification of the duality functor gives an anti-linear involution ψ;
- the graded Euler pairing

$$\langle [M],[N]\rangle = \sum_{i\in\mathbb{Z}} q^{-i}\dim \operatorname{Hom}(M,N(i))$$

gives a sesquilinear pairing, since

$$\begin{aligned}\langle q^a[M],q^b[N]\rangle &= \sum_{i\in\mathbb{Z}} q^{-i}\dim \operatorname{Hom}(M(a),N(i+b)) = \sum_{i\in\mathbb{Z}} q^{-i}\dim \operatorname{Hom}(M(a),N(i+b))\\ &= \sum_{j\in\mathbb{Z}} q^{-j-a+b}\dim \operatorname{Hom}(M,N(j)) = q^{b-a}\langle [M],[N]\rangle.\end{aligned}$$

DEFINITION 1.5. A pre-canonical structure is said to be *balanced* if $\langle \psi(a_c),a_d\rangle=\delta_{c,d}$ and *almost balanced* if $\langle \psi(a_c),a_d\rangle\in \delta_{c,d}+q^{-1}\mathbb{Z}[q^{-1}]$.

LEMMA 1.6. *Assume that $\mathcal{C}$ is a graded humorous category with a partial order $(C,<)$, and that either:*

(1) *we have a collection of self-dual objects $M_c\cong M_c^{\circledast}$ in $\mathcal{C}$ such that $M_c\cong P_c\oplus(\bigoplus_{c'<c}P_{c'}^{\oplus m_{c'}})$ and take $a_c=[M_c]$; or*

(2) *the category of $\operatorname{Rep}(\mathcal{C})$ is of highest weight*[2] *for the order $(C,<)$, with standard modules Δ_c (graded so that Δ_c is a quotient of P_c), and we take $a_c=[\Delta_c]=[M_c^0]-[M_c^1]+\cdots$ where $\cdots\to M_c^1\to M_c^0\to\Delta_c$ is a projective resolution.*

Then, the basis $\{a_c\}$, the involution ψ and the graded Euler pairing together define a pre-canonical structure; in the second case, the pre-canonical structure is balanced.

Proof. We need to check that ψ is flip-unitary, which follows from the isomorphism

$$\operatorname{Hom}(P,Q(i))=\operatorname{Hom}(Q^{\circledast},P^{\circledast}(i)).$$

We also need to confirm that

$$\psi(a_c)\in a_c+\sum_{c'<c}\mathbb{Z}[q,q^{-1}]\cdot a_{c'}.$$

This is clear in the first case, since $\psi(a_c)=a_c$. In the second, since $[\Delta_c]=[P_c]+\sum_{c'<c}m_{c'}[P_{c'}]$, we have that

$$\psi([\Delta_c])-[\Delta_c]=\sum_{c'<c}(\bar{m}_{c'}-m_{c'})[P_{c'}].$$

Since $[P_{c'}]$ is in the span of $[\Delta_d]$ for $d\leqslant c'$, this completes the proof.

Finally, we must prove the balanced structure in the highest-weight case. This involves looking at $\operatorname{Ext}^\bullet(\Delta_c^{\circledast},\Delta_{c'})$. Thus, we have that

$$\operatorname{Ext}^\bullet(\Delta_c^{\circledast},\Delta_{c'})\cong \operatorname{Ext}^\bullet(\Delta_c,\Delta_{c'}^{\star})^*\cong\begin{cases}\Bbbk & \text{if } c=c',\\ 0 & \text{if } c\neq c'.\end{cases}$$ □

[2] To fix conventions, the indecomposable projective module P_c has a filtration by the modules $\Delta_{c'}$ with $c'\leqslant c$; this is the opposite of the most common convention for highest-weight categories.

1.2 Canonical bases

As suggested in the introduction, in some cases, vector spaces with a pre-canonical structure will have a special basis called a canonical basis.

Definition 1.7. A basis $\{b_c\}$ of V is said to be *canonical* if:

(I) each vector b_c in the basis is invariant under ψ;

(II) each vector b_c in the basis is in the set $a_c + \sum_{c'<c} \mathbb{Z}[q, q^{-1}] \cdot a_{c'}$;

(III) the vectors b_c are *almost orthonormal* in the sense that

$$\langle b_c, b_{c'} \rangle \in \delta_{c,c'} + q^{-1}\mathbb{Z}[[q^{-1}]].$$

Throughout the paper, we will use (I), (II) and (III) to refer to the statements above. The reader might have expected (II) and (III) to instead read:

(II′) the transition matrix from b_c to a_c is the identity modulo q^{-1}; that is, b_c is in the set $a_c + \sum_{c'<c} q^{-1}\mathbb{Z}[q^{-1}] \cdot a_{c'}$.

In many important cases this is an equivalent condition, but the first definition will prove more flexible.

Lemma 1.8. *If the standard basis a_c is almost balanced, then conditions* (II) *and* (III) *are equivalent to condition* (II′).

Proof. Assume that (II) and (III) hold. Then $b_c = a_c + \sum_{d<c} m_d a_d$ and assume that c' is minimal among elements such that $m_{c'} \notin q^{-1}\mathbb{Z}[[q^{-1}]]$. Let $b_{c'} = a_{c'} + \sum_{d<c'} n_d a_d$. By (III) we have that $\langle b_c, b_{c'} \rangle \in q^{-1}\mathbb{Z}[[q^{-1}]]$, so

$$\langle b_c, b_{c'} \rangle = \langle \psi(b_c), b_{c'} \rangle = \left\langle \psi(a_c) + \sum_{d<c} \bar{m}_d \psi(a_d),\ a_{c'} + \sum_{d'<c'} n_{d'} a_{d'} \right\rangle \in q^{-1}\mathbb{Z}[[q^{-1}]].$$

By the minimality of c', we have that $m_d n_{d'} \langle \psi(a_d), a_{d'} \rangle \in q^{-1}\mathbb{Z}[[q^{-1}]]$ unless $d = d' = c'$, so $m_{c'} \langle \psi(a_{c'}), a_{c'} \rangle \in q^{-1}\mathbb{Z}[[q^{-1}]]$ as well. This is possible only if $m_{c'} \in q^{-1}\mathbb{Z}[[q^{-1}]]$ as well, which is a contradiction, so (II′) holds.

Assume that (II′) holds, so that $b_c = a_c + \sum_{d<c} m_d a_d$ and $b_{c'} = a_{c'} + \sum_{d<c'} n_d a_d$ for every c and c'. As calculated above,

$$\langle b_c, b_{c'} \rangle = \langle \psi(b_c), b_{c'} \rangle = \left\langle \psi(a_c) + \sum_{d<c} \bar{m}_d \psi(a_d),\ a_{c'} + \sum_{d<c'} n_d a_d \right\rangle.$$

If $c \neq c'$, each term $m_d n_{d'} \langle \psi(a_d), a_{d'} \rangle \in q^{-1}\mathbb{Z}[[q^{-1}]]$, so the same is true of $\langle b_c, b_{c'} \rangle$. On the other hand, when the basis vectors coincide,

$$\langle b_c, b_c \rangle \equiv \langle \psi(a_c), a_c \rangle \equiv 1 \pmod{q^{-1}}. \qquad \square$$

In fact, the dependence on the standard basis is quite weak.

Theorem 1.9. *Assume that V is a finitely generated $\mathbb{Z}[q, q^{-1}]$-module, equipped with a pre-canonical structure and a canonical basis. If $v \in V$ is any vector in V such that $\bar{v} = v$ and $\langle v, v \rangle \in 1 + q^{-1}\mathbb{Z}[q^{-1}]$, then either v or $-v$ is a canonical basis vector. In particular, the canonical bases of any two pre-canonical structures with the same bar-involution and form coincide up to sign.*

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

Proof. We can write $v = \sum f_c(q) b_c$, with $\bar{f}_c = f_c \in \mathbb{Z}[q, q^{-1}]$. Thus, $\sum_c f_c^2 \in 1 + q^{-1}\mathbb{Z}[q^{-1}]$. This is possible only if $f_c = \pm 1$ for some c and 0 otherwise. □

Canonical bases have the distinct advantage of being computable using a Gram–Schmidt algorithm. Embedded in this algorithm is a well-trodden argument showing the uniqueness of this basis.

Proposition 1.10. *A pre-canonical structure has at most one canonical basis. In fact, if bases $\{a_c\}$ and $\{a'_c\}$ define pre-canonical structures with the same bar-involution and bilinear form, and $a_c = a'_c + \sum_{d<c} p_d a'_d$, then any canonical basis for $\{a_c\}$ coincides with any for $\{a'_c\}$.*

Proof. Assume that $\{b_c\}$ and $\{b'_c\}$ are canonical bases for $\{a_c\}$ and $\{a'_c\}$ which are not identical. Assume that c is minimal with $b_c \neq b'_c$. By assumption, $b_c - b'_c = \sum_{d<c} m_d b_d$ for some $m_d \in \mathbb{Z}((q^{-1}))$, since the span of a_d for $d < c$, the span of a'_d for $d < c$, and the span of b_d for $d < c$ all coincide. By ψ-invariance of canonical bases, m_d must be a palindromic Laurent polynomial in q. On the other hand, by almost orthogonality, $\langle b_c - b'_c, b_d \rangle \in q^{-1}\mathbb{Z}[[q^{-1}]]$, so we must also have $m_d \in q^{-1}\mathbb{Z}[[q^{-1}]]$. This is a contradiction, so $\{b_c\}$ must be unique. □

1.3 Mixed humorous categories

However, showing existence is generally quite difficult, unless the pre-canonical structure comes from a categorification. In this case, we have an easy restatement of the canonical property which is more 'categorical' in nature.

Definition 1.11. Following Beilinson *et al.* [BGS96] and Achar and Stroppel [AS13], we define a humorous category to be *mixed* if there is a weight function wt from indecomposable objects to $\mathbb{Z}$ satisfying

$$\mathrm{wt}(M^{\circledast}) = -\mathrm{wt}(M), \quad \mathrm{wt}(M(1)) = \mathrm{wt}(M) + 1$$

such that $\mathrm{Hom}(M, N) = 0$ whenever $\mathrm{wt}(N) < \mathrm{wt}(M)$ or when $M \ncong N$ and $\mathrm{wt}(N) = \mathrm{wt}(M)$, and $\mathrm{Hom}(M, M) \cong \Bbbk$ when M is indecomposable.

Note that a humorous category can have at most one such weight function, uniquely determined by $\mathrm{wt}(P_c(a)) = a$. For an algebra A satisfying the hypotheses of Lemma 1.3, the category of projectives over A is mixed if and only if A is positively graded.

Definition 1.12. We let the *orthodox basis* of $K^0(\mathcal{C})$ be the basis defined by the classes $[P_c]$ of self-dual indecomposable modules in $\mathcal{C}$.

Note that while a canonical basis depends only on the pre-canonical structure, orthodox bases exist only in Grothendieck groups and depend explicitly on the category. Later we shall see examples of categories with canonically isomorphic Grothendieck groups that give different orthodox bases.

Let k be any ring. For any k-linear idempotent complete category $\mathcal{C}$ and a k-algebra k', we can consider $\mathcal{C} \otimes_k k'$, which is the idempotent completion of the category with the same objects as $\mathcal{C}$ but with

$$\mathrm{Hom}_{\mathcal{C}\otimes_k k'}(M, N) := \mathrm{Hom}_{\mathcal{C}}(M, N) \otimes_k k'.$$

In particular, if $\mathcal{C}$ is a humorous category over $\Bbbk$ and $\mathbb{K} \supset \Bbbk$ is an overfield, then $\mathcal{C} \otimes_{\Bbbk} \mathbb{K}$ is clearly again humorous. Note that since we have assumed splitness from the start, the orthodox basis will be unchanged by field extensions, in opposition to situations such as the representation theory of finite groups.

Proposition 1.13. *The orthodox basis of $\mathcal{C} \otimes_{\Bbbk} \mathbb{K}$ coincides with that for $\mathcal{C}$.*

Proof. By assumption, each indecomposable object in $\mathcal{C}$ is absolutely indecomposable; thus it remains indecomposable on base extension to $\mathbb{K}$, so we have the same orthodox basis. □

On the other hand, if we consider an additive category $\mathcal{C}_{\mathbb{Z}}$ with a grading shift and anti-automorphism that induces a humorous structure on $\mathcal{C}_{\mathbb{Z}} \otimes \Bbbk$ for any field, then we can compare the orthodox bases of $\mathcal{C}_{\mathbb{Z}} \otimes \Bbbk$ for different $\Bbbk$ using familiar methods from modular representation theory.

In particular, every indecomposable object in $\mathcal{C}_{\mathbb{Z}} \otimes \mathbb{F}_p$ can be lifted canonically to an object in $\mathcal{C}_{\mathbb{Z}} \otimes \mathbb{Z}_p$ and then base-changed to an (not necessarily indecomposable) object in $\mathcal{C}_{\mathbb{Z}} \otimes \mathbb{Q}_p$. This defines an *extension map* $E : K^0(\mathcal{C}_{\mathbb{Z}} \otimes \mathbb{F}_p) \to K^0(\mathcal{C}_{\mathbb{Z}} \otimes \mathbb{Q}_p) \cong K^0(\mathcal{C}_{\mathbb{Z}} \otimes \mathbb{Q})$. By construction, we have the following property.

Proposition 1.14. *The matrix coefficients of E in the orthodox bases are non-negative integers.*

The map E is often, but not always, an isomorphism. It will be an isomorphism in all the examples of importance in this paper (categorifications of universal enveloping algebras and tensor products), but it fails to be so if $\mathcal{C}_{\mathbb{Z}} = \mathbb{Z}[G]$-mod for G a finite group of order divisible by p. Examples where this map is an isomorphism and extension does not preserve indecomposability (we will discuss one in Example 7.8) show that the same space can be naturally endowed with different orthodox bases.

Lemma 1.15. *If $\mathcal{C}$ is a category satisfying the hypotheses of Lemma 1.6, including one of the conditions (1) or (2), then the orthodox basis is canonical for the pre-canonical structure if and only if the category $\mathcal{C}$ is mixed.*

Proof. In the cases (1) and (2), the basis $[P_c]$ satisfies the conditions (I) and (II) of a canonical basis; thus we need only check that almost orthogonality is equivalent to mixedness. Since the bilinear form is the graded Euler form in this case, almost orthogonality is exactly the statement that $\mathrm{Hom}(P_c, P_{c'}(i))_0 = 0$ for $i \leqslant 0$ unless $i = 0$ and $c = c'$, in which case $\mathrm{Hom}(P_c, P_c) \cong \Bbbk$, which is precisely the same as mixedness if $\mathrm{wt}(P_c) = 0$. The compatibility between weight and duality shows that this is the only possible weight function for $\mathcal{C}$, which completes the proof. □

In fact, when a canonical basis and an orthodox basis coincide, there are stronger positivity properties than implied by just the canonical property. In particular, we have the following result.

Corollary 1.16. *If a basis $\{b_c\}$ is simultaneously orthodox for some humorous category and canonical for the induced pre-canonical structure, then*

$$\langle b_c, b_{c'} \rangle \in \delta_{c,c'} + q^{-1}\mathbb{Z}_{\geqslant 0}[[q^{-1}]]. \tag{1.2}$$

Let us make an essentially trivial observation about the compatibility of mixed structures, which will nevertheless prove quite useful. Assume that $\mathcal{C}$ and $\mathcal{C}'$ are humorous categories, and that there is a full and essentially surjective functor $a : \mathcal{C} \to \mathcal{C}'$ which commutes with grading shifts and duality.

Lemma 1.17. *If $\mathcal{C}$ is mixed, then so is $\mathcal{C}'$. In this case, each orthodox basis vector $K^0(\mathcal{C}')$ is the image of a unique orthodox basis vector in $K^0(\mathcal{C})$ under the induced map $[a] : K^0(\mathcal{C}) \to K^0(\mathcal{C}')$; every other orthodox basis vector in $K^0(\mathcal{C})$ is sent to 0. Put differently, the image of the orthodox basis of $K^0(\mathcal{C}')$ (plus a suitable number of zeros) is the image of the orthodox basis of $K^0(\mathcal{C})$.*

Proof. Since the functor a is full, it sends indecomposable objects to indecomposable objects. By essential surjectivity, each indecomposable M is the image of an object N. Every idempotent in $\mathrm{End}(N)$ is sent to 0 or 1 in $\mathrm{End}(M)$; by the finite dimensionality of degree-0 endomorphisms, there exists an idempotent e whose image in $\mathrm{End}(M)$ is 1 and which cannot be written as the sum of two commuting idempotents. The image eN must be indecomposable, and we have $a(eN) = M$; thus we may as well assume that N is indecomposable. If N' is another indecomposable object such that $M = a(N')$, by fullness we must have that $\mathrm{Hom}(N, N') \neq 0$ and $\mathrm{Hom}(N', N) \neq 0$, since the identity of M must be the image of some morphism. By mixedness of $\mathcal{C}$, we have $N' \cong N$. Thus, we can define a weight function by giving M the same weight as N.

Similarly, it follows that $\mathcal{C}'$ is mixed for this weight; if $\mathrm{wt}(M) > \mathrm{wt}(M')$ or $M \ncong M'$ with $\mathrm{wt}(M) = \mathrm{wt}(M')$, then we indeed have $\mathrm{Hom}(M, M') = 0$ by fullness, since these objects are the images of indecomposables with the same vanishing.

We have already shown above that the class of each indecomposable object in $\mathcal{C}'$ is the image of one from $\mathcal{C}$, and that no non-isomorphic pair of objects can be sent to the same class. This shows the desired statement on canonical bases. □

The converse of the above lemma is false, but we can make an 'if and only if' statement if we strengthen the conclusion, as follows.

Lemma 1.18. *Suppose that we have a sequence of full and essentially surjective functors $a_i : \mathcal{C} \to \mathcal{C}_i$ commuting with grading shifts and duality such that for every object M in $\mathcal{C}$, there is some i such that the natural map $\mathrm{End}(M) \to \mathrm{End}(a_i(M))$ is an isomorphism. The category $\mathcal{C}$ is mixed if and only if all $\mathcal{C}_i$ are. A vector in $K^0(\mathcal{C})$ lies in the orthodox basis if and only if its image under $[a_i]$ lies in the orthodox basis for $\mathcal{C}_i$ for every i.*

Proof. The 'only if' direction is Lemma 1.17. If $\mathcal{C}$ is not mixed, then either there is a map between modules with the wrong weights, or there exists some indecomposable module M such that $\mathrm{End}(M)$ is a local ring with non-trivial Jacobson radical. In either case, we can choose a_i so that it does not kill this bad morphism, contradicting the assumption that $\mathcal{C}_i$ is mixed. Similarly, if a class in $K^0(\mathcal{C})$ is not an orthodox vector, then it must be a linear combination of orthodox vectors $[P_{c_1}], \dots, [P_{c_m}]$. Let $P \cong P_{c_1} \oplus \cdots \oplus P_{c_m}$. By assumption, there is some i such that $\mathrm{End}(P) \cong \mathrm{End}(a_i(P))$. Thus, $[a_i]$ does not kill any of these classes. Since the image of our class is orthodox under $[a_i]$, this gives a contradiction. □

Mixedness appears naturally in geometry and category theory. Let $\mathcal{J}$ be a compactly generated $\Bbbk$-linear dg-category such that $\dim \mathrm{Ext}^i_{\mathcal{C}}(M, N) < \infty$ for all objects M and N. Assume that $\mathcal{J}$ is endowed with a t-structure $(\mathcal{J}^{\geqslant 0}, \mathcal{J}^{\leqslant 0})$ and a duality functor $\circledast$ which are compatible in the sense that $\mathcal{J}^{\geqslant 0}$ and $\mathcal{J}^{\leqslant 0}$ are both invariant under duality. Assume further that every simple object in the heart $\mathcal{H} = \mathcal{J}^{\geqslant 0} \cap \mathcal{J}^{\leqslant 0}$ is self-dual and absolutely irreducible.

Definition 1.19. Let $\mathcal{J}$ be the additive subcategory of $\mathcal{J}$ generated by shifts of simple objects of $\mathcal{H}$. We think of this as a category in the usual sense by taking the morphisms between two objects to be the zeroth cohomology of the morphism complex in the dg-category $\mathcal{J}$.

Lemma 1.20. *The category $\mathcal{J}$ is a mixed humorous category, with* (1) *given by homological shift in $\mathcal{J}$ and duality induced by $\circledast$.*

Proof. That $\mathcal{J}$ is additive over $\Bbbk$ is automatic from the definition. The Krull–Schmidt property follows from the assumption that $\dim \mathrm{Ext}^0_{\mathcal{C}}(M, M) < \infty$ for every M. Since we have a t-structure, any module in $\mathcal{H}$ has no negative self-Exts. In particular, no shift of a module in $\mathcal{H}$ can lie

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

in $\mathcal{H}$, so the action of grading shift on indecomposable objects is free. The finiteness of the number of orbits follows from the fact that $\mathcal{H}$ has finitely many simples. This follows from the compact generation of the category by an object M with $\dim \operatorname{Ext}^i_{\mathcal{C}}(M, M) < \infty$. Every orbit has a fixed point given by a representative in $\mathcal{H}$.

Any duality on a compactly generated dg-category must satisfy $(M[1])^{\circledast} \cong M^{\circledast}[-1]$, so that property is automatic. Finally, the mixedness follows from the fact that two simple objects $M, N \in \mathcal{H}$ have trivial negative Exts and

$$\operatorname{Ext}^0_{\mathcal{C}}(M, N) \cong \begin{cases} \Bbbk & \text{if } M \cong N, \\ 0 & \text{if } M \ncong N, \end{cases}$$

by absolute irreducibility. □

The example of greatest interest to us is when $\mathcal{J}$ is a full subcategory of the category of constructible complexes of sheaves on an algebraic variety or Artin stack generated by IC sheaves for trivial local systems. In this case, the t-structure we will want to take is the perverse one and the duality $\circledast$ will be Verdier duality.

1.4 Examples

While we know of no works in the literature where most of the definitions above are made in such generality, there are many examples. In each case, we will leave the details of the pre-canonical structure to the references.

- Kazhdan and Lusztig [KL79] showed that the Hecke algebra of a Weyl group has a canonical basis, now usually called the *Kazhdan–Lusztig basis*.
- Lusztig showed that the simple integrable representations of quantized universal enveloping algebras of Kac–Moody algebras, as well as a small modification $\dot{U}$ of the algebras themselves, have canonical bases [Lus93]. These have also appeared in the work of Kashiwara as *global crystal bases*.
- In finite type, the tensor product of simple representations also carries a natural canonical basis [Lus93, § 27.3]; in the special case of a tensor product of highest- and lowest-weight representations, this works for infinite-type Kac–Moody algebras as well [Lus92].
- Lascoux *et al.* [LLT96] showed that a level-1 Fock space representation of $\widehat{\mathfrak{sl}}_n$ carries a canonical basis. This was extended to higher-level twisted Fock spaces by Uglov [Ugl00]. Brundan and Kleshchev [BK09a] showed that tensor products of level-1 Fock spaces also have a canonical basis arising as a 'limit' of Uglov's.

All of the bases and pre-canonical structures listed above have close ties to humorous categories as in Lemma 1.6.

- The Kazhdan–Lusztig basis arises from the categorification of the Hecke algebra by $B \times B$-equivariant mixed sheaves on G, the associated algebraic group [Spr82]; alternatively, there is an equivalent approach using indecomposable Soergel bimodules [Soe92]. Recently, Elias and Williamson established in [EW12] that the category of Soergel bimodules is mixed for an arbitrary reflection group over $\mathbb{R}$.
- For $\mathfrak{sl}_n$, the canonical basis of a tensor product of fundamental representations corresponds to the projective (or tilting, depending on convention) objects in a parabolic category $\mathcal{O}$, equipped with its Koszul grading [Sus08, Theorem 6].

- For $\widehat{\mathfrak{sl}}_n$, the canonical basis of a level-1 Fock space arises from a graded version of the q-Schur algebra (for q being an nth root of unity) [Ari09]. For higher-level Fock spaces this canonical basis comes from category $\mathcal{O}$ for the Cherednik algebra of the complex reflection group $G(r, 1, \ell)$, by recent work of several authors [RSVV13, Los13, Web13b].
- For $\mathfrak{sl}_2$, the indecomposable objects of $\mathcal{U}$ match Lusztig's canonical basis, by the work of Lauda [Lau10, Corollary 9.12], and similarly for $\mathfrak{sl}_3$, by the work of Stošić [Sto11].

Our aim in this paper is to complete the connection of categorifications to each of the canonical bases listed above, covering the examples of $\dot{U}$ itself and tensor products of highest-weight modules.

2. Dual canonical bases

2.1 Duality of pre-canonical structures

For any $\mathbb{Z}[q, q^{-1}]$-module V, let $\bar{V}$ be the same underlying abelian group with the action of $\mathbb{Z}[q, q^{-1}]$ twisted by the bar-involution.

Throughout § 2, we assume that V is a free $\mathbb{Z}[q, q^{-1}]$-module with a pre-canonical structure $\{\langle -, -\rangle, \psi, \{a_c\}\}$, where $\langle -, -\rangle$ is valued in $\mathbb{Z}[q, q^{-1}]$ (rather than $\mathbb{Z}((q^{-1})))$. In this case, we can consider the dual space V^* of $\mathbb{Z}[q, q^{-1}]$-linear functionals that kill all but finitely many a_c and the anti-linear evaluation map ϵ sending v to $\langle v, -\rangle$. We will furthermore assume that this map is an isomorphism. Versions of the results in this section are possible in more general situations, but we have aimed for the setup that will give us the cleanest statements of results.

Lemma 2.1. *If V is the Grothendieck group of a humorous category $\mathcal{C}$, then $\langle -, -\rangle$ is valued in $\mathbb{Z}[q, q^{-1}]$ if and only if for every pair of objects M and N in $\mathcal{C}$, we have that $\mathrm{Hom}_{\mathcal{C}}(M, N(i)) = 0$ for all but finitely many $i \in \mathbb{Z}$.*

We can naturally identify V^* with the Grothendieck group of $\mathrm{Rep}(\mathcal{C})$ using the pairing $\{[P], [R]\} = \sum_i q^{-i} \dim R(P(i))$. With this identification, we can think of the map $\epsilon : \bar{V} \to V^*$ as categorifying the Yoneda embedding sending $P \mapsto \mathrm{Hom}(P, -)$.

Lemma 2.2. *If the category $\mathrm{Rep}(\mathcal{C})$ has finite projective dimension, then $\epsilon : \bar{V} \to V^*$ is an isomorphism.*

Proof. The space V^* is spanned by the dual basis to the classes of indecomposables in V. These are the classes of the simple representations L_c, namely those killing all but one indecomposable object, with the remaining indecomposable to $\mathbb{k}$. Thus, we need only prove that these objects are in the image of ϵ. Since $\mathrm{Rep}(\mathcal{C})$ has finite global dimension, we have a finite projective resolution $L_c \leftarrow P_0 \leftarrow P_1 \leftarrow \cdots$. Thus, $[L_c] = \sum_{i=0}^{\infty} (-1)^i [P_i]$. □

Throughout this section, we take $\mathcal{C}$ to be a humorous category satisfying the hypotheses of Lemmata 2.1 and 2.2.

In this case, the same underlying abelian group of V has a second natural pre-canonical structure. Let $\psi^* : V \to V$ be the involution defined by

$$\langle \psi u, v\rangle = \overline{\langle u, \psi^* v\rangle}. \tag{2.3}$$

We call this the *dual bar-involution*. If we identify the Grothendieck groups of $\mathcal{C}$ and $\mathrm{Rep}(\mathcal{C})$ as discussed above, comparing (1.1) and (2.3) shows that ψ^* categorifies the duality $\star$ defined in (1.1).

Let $\{a_c^*\}$ denote the right dual basis of $\{a_c\}$, equipped with the opposite partial order, i.e. it is the unique basis satisfying $\langle a_c, a_{c'}^*\rangle = \delta_{c,c'}$. In certain infinite-dimensional situations, the

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

Gram–Schmidt construction of this basis will not converge; it is enough to assume that V is a sum of finite-dimensional orthogonal spaces with bases given by sets of a_c. This will hold for a tensor product of highest-weight modules (where we break into weight spaces) but not for $U_q(\mathfrak{g})$.

Note that in this notation, a pre-canonical structure is balanced if $\psi^*(a_c) = a_c^*$.

Proposition 2.3. *The (Koszul) dual pre-canonical structure on $\bar{V}$ is defined by the triple $\{\overline{\langle -,-\rangle}, \psi^*, \{a_c^*\}\}$, where we endow C with the opposite partial order. If the primal structure $\{\langle -,-\rangle, \psi, \{a_c\}\}$ has a canonical basis $\{b_c\}$, then the right dual basis $\{b_c^*\}$ is canonical for the dual pre-canonical structure.*

Thus, $\{b_c^*\}$ doubly merits the title *dual canonical basis*: it is both the dual basis to a canonical one and canonical for the dual pre-canonical structure.

Proof. Since ψ is flip-unitary, its conjugate is flip-unitary as well:

$$\langle \psi^* u, \psi^* v\rangle = \overline{\langle \psi\psi^* u, v\rangle} = \overline{\langle \psi v, \psi^* u\rangle} = \langle v, u\rangle.$$

Furthermore, the upper-triangularity of ψ on the basis $\{a_c\}$ exactly translates into lower-triangularity for the transpose. Thus, for the reversed order, we have that ψ^* is upper-triangular.

Obviously, the dual basis to any ψ-invariant basis will consist of ψ^*-invariant elements. Similarly, if a basis is triangular with respect to $\{a_c\}$, its dual basis will be obtained from $\{a_c^*\}$ by the transposed basis-change matrix and thus also be triangular.

Note that by duality we have $b_c = \sum_{c'} \langle b_{c'}, b_c\rangle b_{c'}^*$; thus,

$$\delta_{cc''} = \langle b_c, b_{c''}^*\rangle = \sum_{c'} \langle b_{c'}, b_c\rangle \overline{\langle b_{c'}^*, b_{c''}^*\rangle}.$$

That is, the matrix of $\langle -,-\rangle$ for the basis $\{b_c\}$ is inverse to that for $\overline{\langle -,-\rangle}$ in $\{b_c^*\}$.

Since $\langle b_{c'}, b_c\rangle \in \delta_{c'c} + q^{-1}\mathbb{Z}[[q^{-1}]]$, it must also hold that $\overline{\langle b_{c'}^*, b_{c''}^*\rangle} \in \delta_{c'c''} + q^{-1}\mathbb{Z}[[q^{-1}]]$. □

In this case we have a second way of thinking about the dual canonical basis.

Corollary 2.4. *Assume that $\mathcal{C}$ is a humorous category such that its orthodox basis defines a canonical basis for the induced pre-canonical structure. The dual canonical basis of this pre-canonical structure is defined by the classes of the simple modules in* $\mathrm{Rep}(\mathcal{C})$.

The significance of using the barred Euler form in the dual pre-canonical structure is that the Ext space between two simple modules in a mixed category is *negatively* graded, as a simple calculation with free resolutions shows, so this reversal is necessary if we are to have any hope of canonicity holding.

When the pre-canonical structure is balanced, we have yet a third pre-canonical structure, given by $(\langle -,-\rangle, \psi^*, \{a_c\})$. We call this the *Ringel dual* pre-canonical structure, since on the categorical level it matches with Ringel duality.

2.2 Balanced positivity

One particularly common phenomenon is that if a_c is carefully chosen, the basis $\{b_c\}$ will have good positivity properties.

Definition 2.5. We say a pre-canonical structure with canonical basis $b_c = \sum_{c'} m_{cc'}(q) a_{c'}$ is *balanced positive* if it is balanced, $m_{cc'}(q) \in q^{-1}\mathbb{Z}_{\geqslant 0}[[q^{-1}]]$, and

$$a_c = \sum_{c'} n_{cc'}(-q) b_{c'}$$

for $n_{cc'}(q) \in q^{-1}\mathbb{Z}_{\geqslant 0}[[q^{-1}]]$.

Such bases arise naturally in representation theory through categorifications. Assume that $\mathcal{C}$ is a mixed humorous category such that $\mathrm{Rep}(\mathcal{C})$ is of highest weight and is standard Koszul (see [ÁDL03] for definitions). As before, we let C be the set of self-dual indecomposable modules in $\mathcal{C}$ and let $\Delta(c)$ be the standard modules with cosocle concentrated in weight 0.

Theorem 2.6. *With $\mathcal{C}$ as discussed above, the Grothendieck group $K^0(\mathcal{C})$ has a balanced positive pre-canonical structure such that*

$$\psi = [\circledast], \quad \psi^* = [\star], \quad \langle [M],[N]\rangle = \sum_{i,j}(-1)^i q^{-j} \dim \mathrm{Ext}^i_A(M, N(j)), \quad a_c = [\Delta(c)].$$

The canonical basis is given by the classes of the $\circledast$-self-dual indecomposable projective modules and the dual canonical basis by the classes of the $\star$-self-dual simple modules. The canonical basis of the Ringel dual pre-canonical structure is given by the classes of $\star$-self-dual indecomposable tilting modules.

This setup seems quite specialized (and indeed it is), but it has already made several appearances in the literature. Categories that satisfy the hypotheses of the theorem include:

- the category $\mathcal{O}$ of a semi-simple Lie algebra, by [ÁDL03, Corollary 3.8], so the Kazhdan–Lusztig basis of the Hecke algebra is balanced positive;
- the truncated parabolic category $\mathcal{O}$ of Shan and Vasserot, by [SVV11, Theorem 3.12], so Uglov's canonical basis of twisted Fock spaces is balanced positive; as a special case, the same holds for tensor products of wedge powers of the natural representation of $\mathfrak{sl}_n$;
- the hypertoric category $\mathcal{O}$ defined by the author jointly with Braden, Licata and Proudfoot [BLPW10, BLPW12].

Proof. That we have a pre-canonical structure follows from Lemma 1.6. The compatibility of these two involutions is precisely the decategorification of (1.1). By Lemma 1.15, the basis $b_c = [P_c]$ is canonical, so by Proposition 2.3 the right dual basis $b^*_c = [L_c]$ consists of the classes of the simple modules. The classes of the indecomposable tiltings are invariant under ψ^*, and almost orthogonality follows from the positivity of the grading on the Ringel dual; the upper-triangularity with respect to classes of the standards is a standard property of tiltings.

Finally, we wish to show balanced positivity. The positivity of $[P_c]$ in terms of $[\Delta_c]$ follows from the fact that P_c has a standard filtration; the polynomials $m_{**}(q)$ are just the graded multiplicities of this filtration. On the other hand, the (twisted) positivity of $[\Delta_c]$ follows from the fact that $[\Delta_c]$ has a *linear* resolution by projectives (this is the definition of standard Koszulity). Thus $n_{**}(q)$ consists of the graded multiplicities of this resolution, where q could measure either grading shift or homological shift, which coincide by the linearity of the resolution. □

The above theorem shows that balanced positivity is a natural condition from the categorical perspective. Now we give a combinatorial consequence of this definition.

Theorem 2.7. *If $\{\langle -,-\rangle, \psi, \{a_c\}\}$ is a balanced positive pre-canonical structure, then its dual structure will be balanced positive for the variable $p = -q^{-1}$.*

Proof. The essential point is that the polynomials m_{**} and n_{**} will switch roles. The proof of this fact is a combinatorial version of BGG reciprocity. Note that

$$\langle b_c, a^*_{c'}\rangle = \overline{m_{cc'}(q)} = m_{cc'}(-p), \quad \langle a_c, b^*_{c'}\rangle = \overline{n_{cc'}(-q)} = n_{cc'}(p).$$

By the definition of dual bases, we have

$$a_c^* = \sum_{c'} \langle b_{c'}, a_c^* \rangle b_{c'}^* = \sum_{c'} m_{c'c}(-p) b_{c'}^*,$$
$$b_c^* = \sum \langle a_{c'}, b_c^* \rangle a_{c'}^* = \sum_{c'} n_{c'c}(p) a_{c'}^*,$$

and this exactly shows balanced positivity. □

Remark 2.8. When we are considering a basis which comes from a category $\mathcal{C} = A$-mod for some standard Koszul algebra A satisfying the the hypotheses of Theorem 2.6, this fact has a categorical proof. As usual, the algebra A has a Koszul dual $A^! \cong \mathrm{Ext}_A^\bullet(A_0, A_0)$, which satisfies the same conditions. The Koszul duality functor K defined in [BGS96, § 2.12] induces an isomorphism of graded Grothendieck groups $K^0(A\text{-mod}) \cong K^0(A^!\text{-mod})$ sending q to $-q^{-1}$, which sends the dual pre-canonical structure to the primal pre-canonical structure of Theorem 2.6; thus, if one has the desired positivity, the other does as well.

3. The 2-category $\mathcal{U}$

In this paper, our notation builds on that of Khovanov and Lauda, who gave a graphical version of the 2-quantum group, which we denote by $\mathcal{U}$ (leaving $\mathfrak{g}$ understood). These constructions could also be rephrased in terms of Rouquier's description, and we have striven to make the paper readable following either [KL10] or [Rou08]; however, it is most sensible for us to use the 2-category defined by Cautis and Lauda [CL11], which is a variation on both of these. See the introduction of [CL11] for more detail on the connections between these different approaches.

The object of interest in this section is a strict 2-category. As described, for example, in [Lau10], one natural way of discussing strict 2-categories is via planar diagrammatics. The 2-category $\mathcal{U}$ is thus most clearly described in this language.

Definition 3.1. A *blank KL diagram* is a collection of finitely many oriented curves in $\mathbb{R} \times [0, 1]$ which has no triple points, decorated with finitely many dots. Every strand is labeled with an element of Γ, and any open end must meet one of the lines $y = 0$ or $y = 1$ at a distinct point from all other ends.

A *KL diagram* is a blank KL diagram together with a labeling of regions between strands (the components of its complement), with weights obeying the following rule.

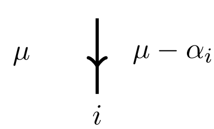

We identify two KL diagrams if they are isotopic via an isotopy which does not cancel any critical points of the height function or move critical points through crossings or dots. We will deal with isotopies that do have these features later. In the interest of simplifying diagrams, we will often write a dot with a number beside it to indicate a group of that number of dots.

We call the lines $y = 0$ and $y = 1$ the *bottom* and *top* of the diagram, respectively. Reading across the bottom and the top from left to right, we obtain a sequence of elements of Γ, which we wish to record in order from left to right. Since orientations are quite important, we let $\pm\Gamma$ denote $\Gamma \times \{\pm 1\}$, and associate i to a strand labeled with i which is oriented upward and

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

associate $-i$ to a strand oriented downward. For example, below is a blank KL diagram

$m =$ 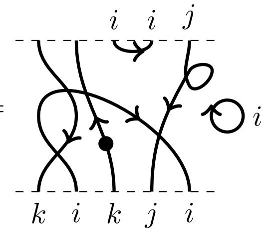

with top given by $(-k, k, -i, i-j)$ and bottom given by $(-k, i, k, -j, -i)$.

We also wish to record the labeling on regions; since fixing the label on one region determines all the others, we will typically record only $\mathcal{L}$, the weight of the region at the far left, and $\mathcal{R}$, the weight at the far right. In addition, in the interest of simplifying pictures we will typically not draw the weights on all regions. We call the pair consisting of a sequence $\mathbf{i} \in (\pm\Gamma)^n$ and the weight $\mathcal{L}$ a *KL pair*; let $\mathcal{R} := \mathcal{L} + \sum_{j=1}^{n} \alpha_{i_j}$, where we let $\alpha_{-i} = -\alpha_i$.

Definition 3.2. Given KL diagrams a and b, their *(vertical) composition ab* is obtained by stacking a on top of b and attempting to join the bottom of a and the top of b. If the sequences from the bottom of a and the top of b do not match or if $\mathcal{L}_a \neq \mathcal{L}_b$, then the composition is not defined and by convention is 0, which is not a KL diagram, just a formal symbol.

The *horizontal composition* $a \circ b$ of KL diagrams is the diagram which pastes together the strips where a and b live with a to the *right* of b. The only compatibility we require is that $\mathcal{L}_a = \mathcal{R}_b$, so that the regions of the new diagram can be labeled consistently. If $\mathcal{L}_a \neq \mathcal{R}_b$, the horizontal composition is 0 as well.

Implicit in this definition is a rule for *horizontal composition* of KL pairs in $\pm\Gamma$, which is the reverse of the concatenation $(i_1, \ldots, i_m) \circ (j_1, \ldots, j_n) = (j_1, \ldots, j_n, i_1, \ldots, i_m)$ and gives 0 unless $\mathcal{L}_{\mathbf{i}} = \mathcal{R}_{\mathbf{j}}$.

We should warn the reader that this convention requires us to read our diagrams differently from the conventions of [Lau10, KL10, CL11]; in our diagrammatic calculus, 1-morphisms point from left to right, not from right to left as indicated in [Lau10, § 4]. The practical implication will be that our relations are the reflection through a vertical line of those of Cautis and Lauda.

We can define a *degree* function on KL diagrams. The degrees are given on elementary diagrams by

deg 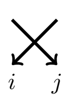 $= \langle \alpha_i, \alpha_j \rangle$, deg  $= -\langle \alpha_i, \alpha_i \rangle$, deg 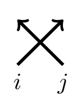 $= \langle \alpha_i, \alpha_j \rangle$, deg  $= -\langle \alpha_i, \alpha_i \rangle$,

deg 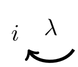 $= \langle \lambda, \alpha_i \rangle - d_i$, deg 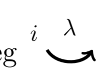 $= -\langle \lambda, \alpha_i \rangle - d_i$,

deg 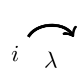 $= \langle \lambda, \alpha_i \rangle - d_i$, deg 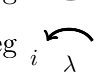 $= -\langle \lambda, \alpha_i \rangle - d_i$.

For a general diagram, we sum together the degrees of the elementary diagrams it is constructed from.

Now, we wish to assemble these into a linear category over a ring $\Bbbk$; we shall be most interested in the case where $\Bbbk$ is a field, but it will be convenient for us to let $\Bbbk$ be a commutative complete local ring as well. Once and for all, fix a matrix of polynomials $Q_{ij}(u, v) = \sum_{k,m} Q_{ij}^{(k,m)} u^k v^m$ valued in $\Bbbk$ and indexed by $i \neq j \in \Gamma$; by convention $Q_{ii} = 0$. We assume that each polynomial is

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

homogeneous of degree $-\langle \alpha_i, \alpha_j \rangle = -2d_i c_{ij} = -2d_j c_{ji}$ when u is given degree $2d_i$ and v degree $2d_j$. We will always assume that the leading order of Q_{ij} in u is $-c_{ij}$ and that $Q_{ij}(u,v) = Q_{ji}(v,u)$. We let $t_{ij} = Q_{ij}^{(-c_{ij},0)} = Q_{ij}(1,0)$; by convention $t_{ii} = 1$. In [CL11], the coefficients of this polynomial are denoted by

$$Q_{ij}(u,v) = t_{ij}u^{-c_{ij}} + t_{ji}v^{-c_{ji}} + \sum_{pd_i+qd_j=d_ic_{ij}} s_{ij}^{pq}u^p v^q.$$

Khovanov and Lauda's original category uses the choice $Q_{ij} = u^{-c_{ij}} + v^{-c_{ji}}$.

Definition 3.3. Let $\mathcal{U}$ be the 2-category whose objects are the weights $Y(\mathfrak{g})$, whose 1-morphisms $\mu \to \nu$ are grading shifts of KL pairs with $\mathcal{L} = \mu$ and $\mathcal{R} = \nu$, and whose degraded 2-morphisms are a quotient of the formal span over $\Bbbk$ of KL diagrams. Before giving these, we note that each such diagram has a degree, which we adjust by the grading shift of its source and target to arrive at the degree of the 2-morphism; by our conventions, these will be 'honest' 2-morphisms only if they have degree 0. As before, we denote the space of 2-morphisms of degree 0 between two 1-morphisms u and v by $\operatorname{Hom}(u,v)$, and we let $\operatorname{HOM}(u,v) := \bigoplus_i \operatorname{Hom}(u, v(i))$ denote the space of degraded morphisms.

The relations that we impose on degraded 2-morphisms are the following.

- The cups and caps are the units and counits of a biadjunction. The dot morphism is cyclic. The cyclicity for crossings can be derived from the pitchfork relation as follows.

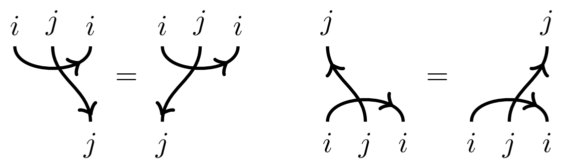

(3.4a)

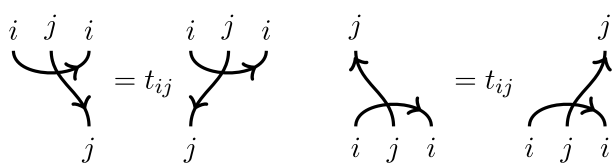

(3.4b)

  The mirror images of these relations through a vertical axis also hold.
- Recall that a *bubble* is a morphism given by a closed circle, endowed with some number of dots. Any bubble of negative degree is 0, and any bubble of degree 0 is equal to 1. We must add formal symbols, called 'fake bubbles', which are bubbles labeled with a negative number of dots (these are explained in [KL10, § 3.1.1]); given these, we have the following inversion formula for bubbles:

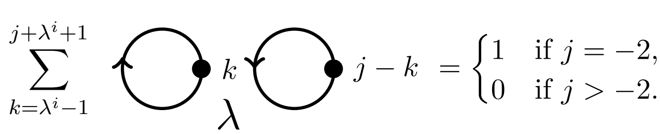

(3.5)

- Two relations connecting the crossing with cups and caps are as shown in (3.6a)–(3.6d).

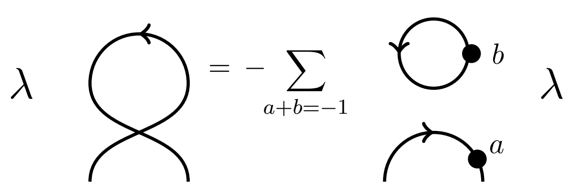

(3.6a)

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

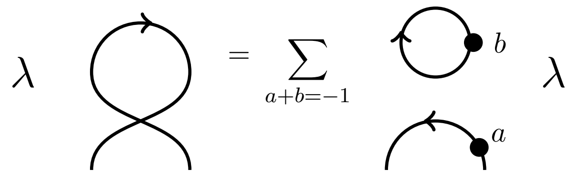

(3.6b)

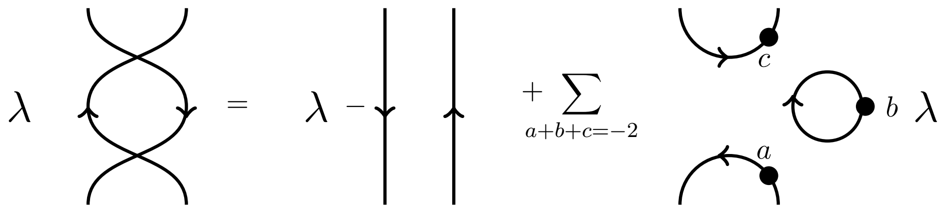

(3.6c)

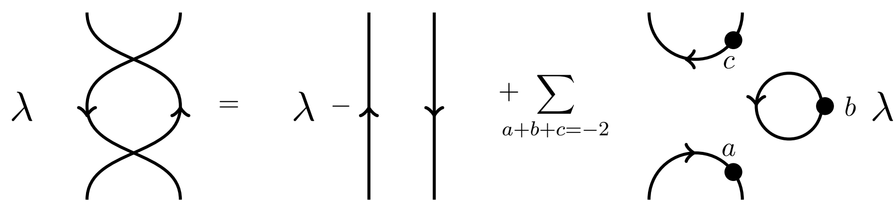

(3.6d)

- Oppositely oriented crossings of differently colored strands simply cancel with a scalar.

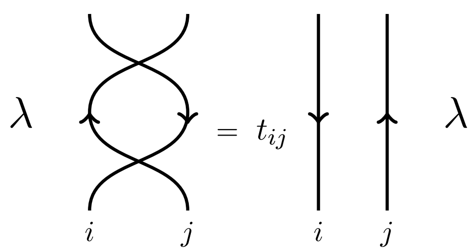

(3.7a)

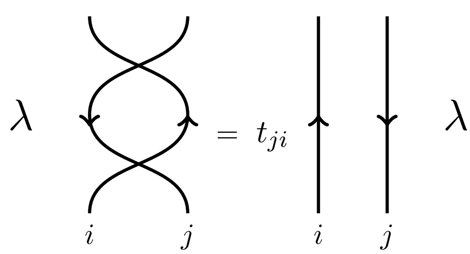

(3.7b)

- The endomorphisms of words only using only $-\Gamma$ (or, by duality, only $+\Gamma$) satisfy the following relations (3.8a)–(3.8g) of the *quiver Hecke algebra* R.

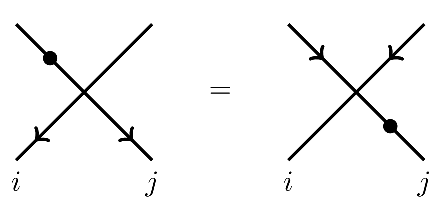

unless $i = j$ (3.8a)

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

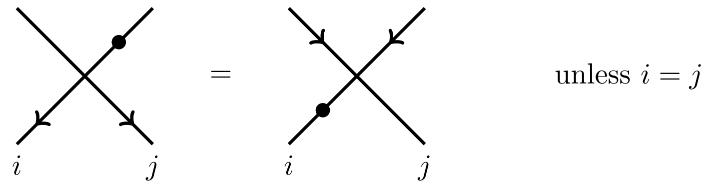

(3.8b)

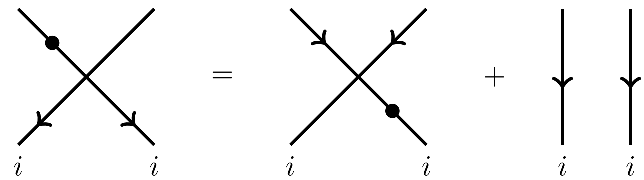

(3.8c)

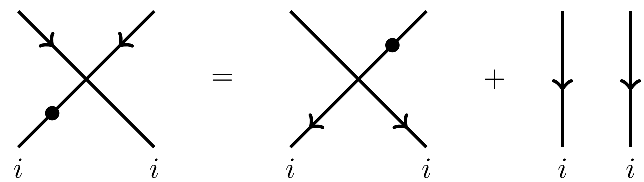

(3.8d)

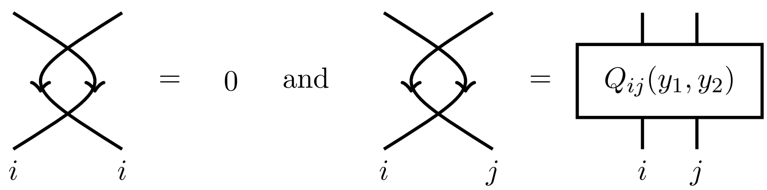

(3.8e)

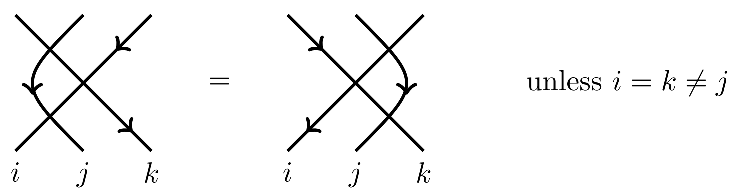

(3.8f)

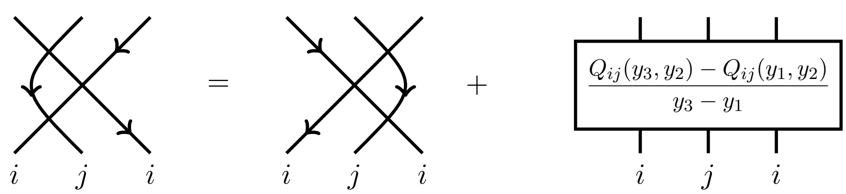

(3.8g)

As in [KL10], we let $\dot{\mathcal{U}}$ denote the strict 2-category where every Hom-category is replaced by its idempotent completion. We note that since every object in $\mathcal{U}$ has a finite-dimensional degree-0 part of its endomorphism algebra, every Hom-category in $\dot{\mathcal{U}}$ satisfies the Krull–Schmidt property.

This 2-category is a categorification of the universal enveloping algebra in the following sense.

Theorem 3.4 [Web13a, Theorem B.2]. *The Grothendieck group of $\dot{\mathcal{U}}$ is isomorphic to $\dot{U}$, and its graded Euler form is given by Lusztig's inner product $(-,-)$ on $\dot{U}$.*

This theorem was first conjectured by Khovanov and Lauda [KL10] and then proved by them in the special case of $\mathfrak{sl}_n$. While not explicitly stated in their paper, in the finite-type case this also follows easily from [CL11, Corollary 7.2], which was proved independently of the work above, relying on the paper [KK12] of Kang and Kashiwara in its stead.

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

We recall from [KL10, § 3.3.2] that we have an involution

$$\tilde{\psi} : \mathrm{HOM}(\mathcal{E}_{i_1} \cdots \mathcal{E}_{i_m}\lambda, \mathcal{E}_{j_1} \cdots \mathcal{E}_{j_n}\lambda) \to \mathrm{HOM}(\mathcal{E}_{j_1} \cdots \mathcal{E}_{j_n}\lambda, \mathcal{E}_{i_1} \cdots \mathcal{E}_{i_m}\lambda)$$

reflecting the diagrams of two morphisms through a horizontal line and reversing orientation. This extends to a 2-functor $\dot{\mathcal{U}} \to \dot{\mathcal{U}}$ which is covariant on 1-morphisms and contravariant on 2-morphisms, sending $\mathcal{E}_i(k) \mapsto \mathcal{E}_i(-k)$ and $\mathcal{F}_i(k) \mapsto \mathcal{F}_i(-k)$.

Proposition 3.5 [KL10, Proposition 3.28]. *The 2-functor $\tilde{\psi}$ categorifies the bar-involution of $\dot{U}_q(\mathfrak{g})$ (denoted by ψ in [KL10]).*

This inner product and involution are part of the pre-canonical structure used by Lusztig to define the canonical basis of $\dot{U}$; the role of the standard basis can be played by a number of different bases of $\dot{U}$. We will use one defined using string parametrizations of crystal elements. This is perhaps less elegant on the level of the quantum groups than the PBW basis defined via the braid action used by Lusztig in [Lus90] (in particular, it is not balanced as in Definition 1.5), but it is easier to handle in the categorification.

4. The 2-category $\mathcal{T}$

4.1 Tricolore diagrams

In the next three sections, we will present a construction of a categorification of tensor products of highest- and lowest-weight representations. Almost all of the results which appear here have equivalents in the author's earlier paper [Web13a], and in most cases the nature of the proofs is quite similar. First, we present an auxiliary category which generalizes that presented in [Web13a, § 4.5].

Definition 4.1. A *blank tricolore diagram*[3] is a collection of finitely many oriented curves in $\mathbb{R} \times [0, 1]$. Each curve is either:

- thick, smooth and labeled with a dominant weight of $\mathfrak{g}$ (when color is available, we color these red); or
- thick, wavy and labeled with an anti-dominant weight of $\mathfrak{g}$ (when color is available, we color these blue); or
- thin, labeled with $i \in \Gamma$ and decorated with finitely many dots.

The smooth strands are constrained to be oriented downwards and the wavy strands to be oriented upwards; we will generally not draw the orientation on these strands. Furthermore, smooth and wavy strands are forbidden to intersect with any other smooth or wavy strand. The thin strands are allowed to close into circles, self-intersect, intersect smooth and wavy strands, and so on.

Blank tricolore diagrams divide their complement in $\mathbb{R}^2 \times [0, 1]$ into finitely many connected components, and we define a *tricolore diagram* to be a blank tricolore diagram together with a labeling of these regions by weights consistent with the following rules.

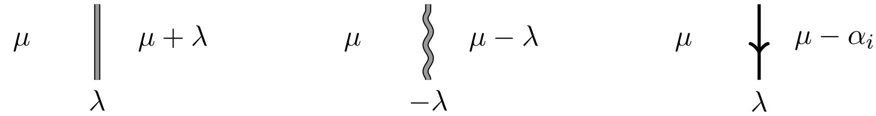

[3] When drawn on a blackboard, such a diagram involves red, white and blue colors. Citizens of Australia, Cambodia, Chile, the Cook Islands, Costa Rica, Croatia, Cuba, the Czech Republic, the Dominican Republic, the Faroe Islands, France, Haiti, Iceland, North Korea, Laos, Liberia, Luxembourg, the Netherlands, Norway, Panama, Paraguay, Russia, Samoa, Serbia, Sint Maarten, Slovakia, Slovenia, Taiwan, Thailand, the United Kingdom and the United States are all free to regard this patriotically according to their preferences.

Since this labeling is fixed as soon as one region is labeled, in the interest of simplifying pictures we will typically not draw in the weights in all regions.

For example,

$a' =$ 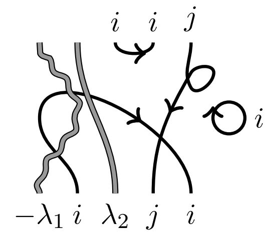

is a blank tricolore diagram. Both the notion of *KL diagrams* and that of *double Stendhal diagrams* from [Web13a] are special cases of tricolore diagrams: a KL diagram is a tricolore diagram with no smooth and wavy strands, while a double Stendhal diagram is a tricolore diagram with no wavy strands.

As usual, we will want to record the horizontal slices at $y = 0$ and $y = 1$, the *bottom* and *top* of the diagram. These will be encoded as a *tricolore quadruple*, consisting of:

- the sequence $\mathbf{i} \in (\pm\Gamma)^n$ of simple roots and their negatives on thin strands, read from the left;
- a sequence $\underline{\boldsymbol{\lambda}} \in (Y^{\pm})^{\ell}$ of dominant or anti-dominant weights on smooth and wavy strands, read from the left;
- the weakly increasing function $\kappa : [1, \ell] \to [0, n]$ such that $\kappa(m)$ is the number of thin strands to the left of the mth smooth or wavy strand (both counted from the left). By convention, we write $\kappa(i) = 0$ if the ith smooth or wavy strand is to the left of all thin strands.
- the weights $\mathcal{L}$ and $\mathcal{R}$ at the far left and far right of the diagram; these are related by

$$\mathcal{L} + \sum_{k=1}^{\ell} \lambda_k + \sum_{m=1}^{n} \alpha_{i_m} = \mathcal{R}.$$

We shall often condense the first three items above into a single sequence whose entries are elements of both $\pm\Gamma$ and $Y^{\pm}$; we will typically write such sequences with upper-case sans-serif letters, such as I.

Tricolore diagrams are endowed with horizontal and vertical composition operations, just like KL and double Stendhal diagrams; similarly, tricolore quadruples are endowed with a horizontal composition. As in [Web13a], we maintain the dyslexic convention that the horizontal composition $a \circ b$ places a to the *right* of b, and we read diagrams from bottom to top.

Definition 4.2. We let $\tilde{\tilde{\mathcal{T}}}$ be the strict 2-category where:

- objects are weights in $X(\mathfrak{g})$;
- 1-morphisms $\mu \to \nu$ are tricolore quadruples with $\mathcal{L} = \mu$ and $\mathcal{R} = \nu$, and composition is given by horizontal composition as above;
- degraded 2-morphisms $h \to h'$ between tricolore quadruples are $\Bbbk$-linear combinations of tricolore diagrams with h as bottom and h' as top, and vertical and horizontal composition of 2-morphisms is as defined above. As with $\mathcal{U}$, we should only take elements of degree 0 as 'honest' 2-morphisms.

We can grade the 2-morphism spaces of this 2-category by endowing each tricolore diagram with a degree. KL diagrams are graded by the degrees given in § 3. The smooth/thin or wavy/thin

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

crossings have the following degrees, which are invariant under reflection through a vertical line:

deg 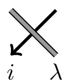 $= \langle \alpha_i, \lambda \rangle$, deg 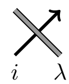 $= 0$, deg 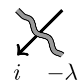 $= 0$, deg 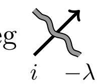 $= \langle \alpha_i, \lambda \rangle$.

Definition 4.3. Let $\mathcal{T}$ be[4] the quotient of $\tilde{\tilde{\mathcal{T}}}$ by the following relations on 2-morphisms.

- All the relations of $\mathcal{U}$ given in (3.4a)–(3.8g) hold on thin strands.
- Oppositely oriented crossings of differently labeled strands simply cancel, as shown below in (4.10). This includes crossings of smooth/wavy strands with thin ones.

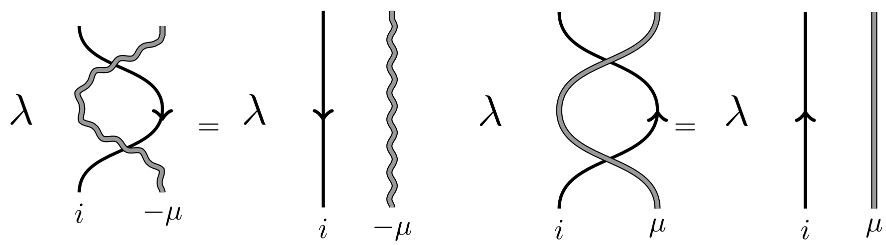

(4.10)

- All thin crossings and dots can pass through smooth or wavy lines, with a correction term similar to Khovanov and Lauda's, as shown below. (For the relations of (4.11c)–(4.11d), we also include their mirror images through a vertical line.)

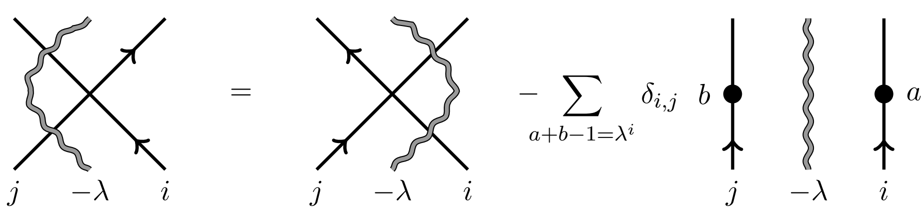

(4.11a)

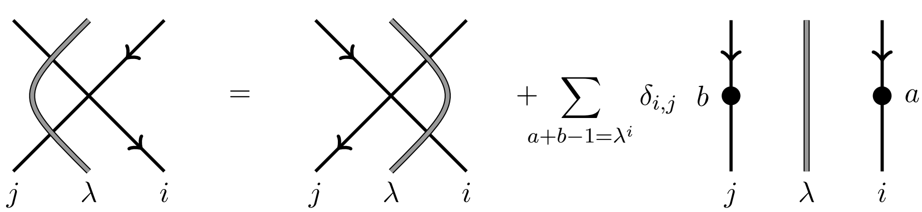

(4.11b)

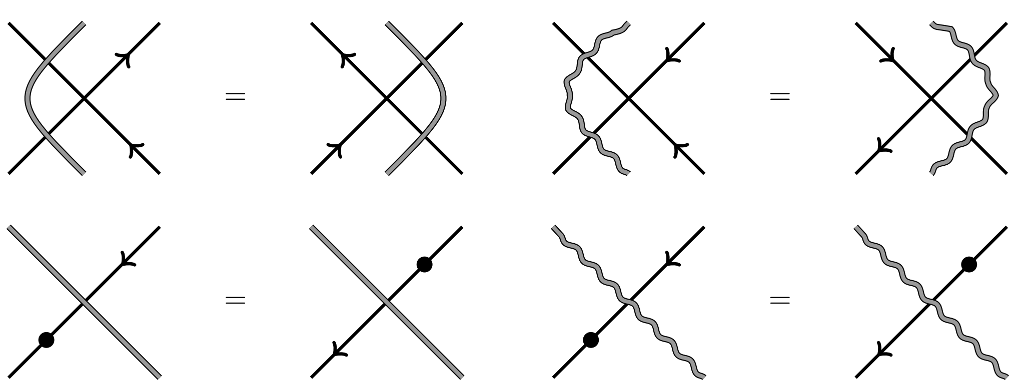

(4.11c)

[4] The author recognizes that the same symbol is used in [Web13a] to denote the subcategory of this one where no wavy strands are allowed. From context, we do not think that confusion is likely.

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

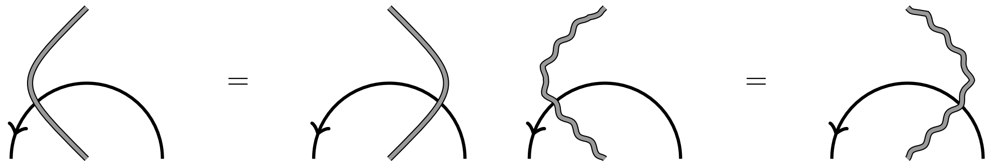

(4.11d)

- The ‘cost’ of separating similarly oriented smooth/wavy and thin lines is adding $\lambda^i = \alpha_i^\vee(\lambda)$ dots to the thin strand as shown below in (4.12).

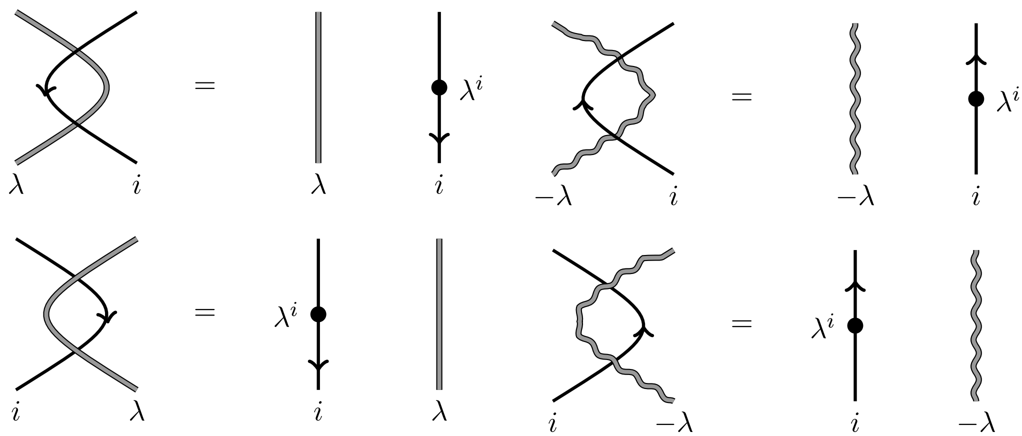

(4.12)

The category $\mathcal{T}$ has a subcategory, which we will call the *Polish*[5] *subcategory*, given by diagrams with no wavy lines, only smooth and thin ones. We have a complementary *Scottish*[6] *subcategory* of diagrams with no smooth strands, only wavy and thin ones.

As with $\mathcal{U}$ and $\dot{\mathcal{U}}$, we let $\dot{\mathcal{T}}$ be the idempotent completion of the Hom-categories of $\mathcal{T}$. We have an obvious 2-functor $\mathcal{U} \to \mathcal{T}$ (and thus $\dot{\mathcal{U}} \to \dot{\mathcal{T}}$), thinking of a KL diagram as a tricolore diagram. The 2-category $\mathcal{T}$ thus carries an action of $\mathcal{U}$ by horizontal composition on the right and on the left.

4.2 Non-degeneracy

Unfortunately, as is usual with a presentation by generators and relations, it is far from obvious that the category even has any non-zero objects.

We shall prove this by showing that $\mathcal{T}$ has a ‘polynomial representation’ much like a KLR algebra, but not on a polynomial ring. Instead, we let $\mathcal{Q}$ be any finite set of KL pairs and consider the representation $\mathcal{T}$ where the quadruple triple $(\underline{\boldsymbol{\lambda}}, \mathbf{i}, \kappa)$ with $\mathcal{R} = \mu$ is sent to $\bigoplus_{q\in\mathcal{Q}} \mathrm{HOM}_{\mathcal{U}'}(\mathbf{i}, q)$, where $\mathbf{i}$ represents the KL pair with this sequence and $\mathcal{R} = \mu$, and $\mathcal{U}'$ is the localization of the 2-category by adjoining a formal inverse to $y : \mathcal{F}_i \to \mathcal{F}_i$ and requiring the relation (3.5) to hold with honest bubbles, as discussed in [Web13a, § 3.6].

The basic 2-morphisms in $\mathcal{T}$ act as follows.

- The crossing ψ acts by simply removing the smooth and wavy strands and acting by the resulting KL diagram, as does the dot y.
- The clockwise cup ι between the kth and $(k+1)$st smooth/wavy strands acts by $(1 \otimes y^{-\lambda_+^i})\iota$ where λ_+ is the sum of the dominant weights in $\{\lambda_{k+1}, \ldots, \lambda_\ell\}$, that is, the labels on the

[5] No slight to the residents of Austria, Bahrain etc. intended.

[6] Similar apologies to Salvadoreños, Finns etc.

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

smooth strands right of the cup. Similarly, the counterclockwise cap ϵ acts by $\epsilon(1 \otimes y^{\lambda_+^i})$.

- The counterclockwise cup ι' acts by $(1 \otimes y^{\lambda_-^i})\iota'$, where λ_- is the sum of the labels on wavy strands right of the cup. Similarly, the clockwise cap ϵ' acts by $\epsilon'(1 \otimes y^{-\lambda_-^i})$.

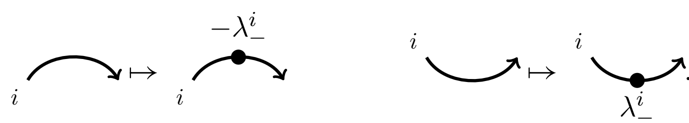

- A smooth/thin (respectively, wavy/thin) crossing where the thin strand runs NW/SE or points from SW to NE (respectively, from NE to SW) acts by the identity.
- a smooth/thin (respectively, wavy/thin) crossing where the thin strand points from NE to SW (respectively, from SW to NE) acts by adding λ^i (respectively, $-\lambda^i$) dots on the thin strand.

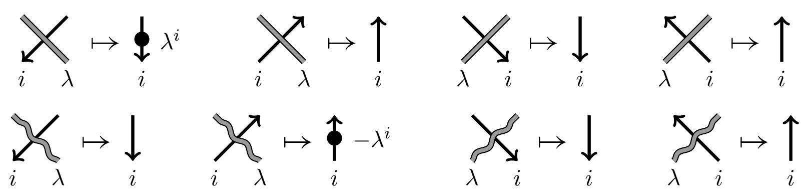

Diagrammatically, we can keep track of this action by letting θ^κ be the unique tricolore diagram having top $(\underline{\boldsymbol{\lambda}}, \mathbf{i}, \kappa)$ and bottom $(\underline{\boldsymbol{\lambda}}, \mathbf{i}, 0)$, with no dots or thin/thin crossings. When a tricolore diagram acts on the top of θ_κ, this is equivalent via our relations to acting on the bottom by a KL diagram, if we allow ourselves to use negative powers of dots. The representation we have described is exactly sending a tricolore diagram to the corresponding KL diagram.

Theorem 4.4. *For any $\mathcal{Q}$, the above defines an action of $\mathcal{T}$.*

Proof. The relations (3.4a)–(3.8g) follow since the map we have given is that of [Web13a, Proposition 3.30] with $a = \lambda_+^i$ and $b = -\lambda_-^i$.

The relations (4.11a)–(4.11b) follow from the same calculation as the polynomial representation of $\tilde{T}$ given in [Web13a, Lemma 4.12]; that is, (4.11a) follows from the calculation

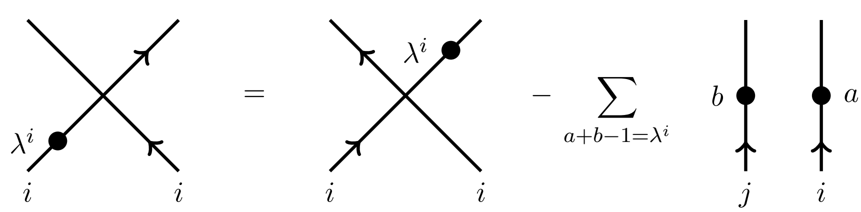

and (4.11b) from its rotation. The relations (4.11c) are clear. Now consider the smooth case of (4.11d), where the strand we can see is labeled with λ_j. On the right-hand side, the action is by λ_+^i dots on a cap, where λ_+ is calculated from the far left of the picture. On the left-hand side, we add $(\lambda_+ - \lambda_j)^i$ dots to the cap (since now we only count smooth lines to the right of the picture we can see) but must add λ_j^i dots to the crossing. Thus, the two sides are the same. The wavy case of (4.11d) follows similarly. Finally, (4.10) and (4.12) are clear. □

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

We say that a tricolore diagram is *minimal* if no two strands (of any color) cross twice in it. Consider a matching between the elements of the sequences $\mathbf{i}$ and $\mathbf{i}'$ which matches appearances of $i \in \pm\Gamma$ in $\mathbf{i}$ with ones in $\mathbf{i}'$, or matches i with $-i$ within $\mathbf{i}$ or within $\mathbf{i}'$. These matchings are precisely like those in [Web13a, Definition 2.7]. For each such matching, there is a minimal diagram with no dots and with top $(\underline{\boldsymbol{\lambda}}, \mathbf{i}, \kappa)$ and bottom $(\underline{\boldsymbol{\lambda}}, \mathbf{i}', \kappa')$ whose strands attach matched elements. Any two such diagrams differ by isotopy and a finite number of switches through triple points. For each such matching and fixed $(\underline{\boldsymbol{\lambda}}, \mathbf{i}, \kappa)$ and $(\underline{\boldsymbol{\lambda}}, \mathbf{i}', \kappa')$, choose a minimal diagram d_π which traces out this matching and fix a point on each strand. Let D be the collection of the diagrams obtained from d_π by ranging over all placings of possible numbers of dots at the chosen points (and nowhere else in the diagram).

Note that the elements of D are in canonical bijection with the basis $B_{\mathbf{i},\mathbf{i}'}$ of Khovanov and Lauda, defined in [KL10, § 2.2].

Proposition 4.5. *The set D is a basis for the morphism space* $\mathrm{HOM}_{\mathcal{T}}((\underline{\boldsymbol{\lambda}}, \mathbf{i}, \kappa), (\underline{\boldsymbol{\lambda}}, \mathbf{i}', \kappa'))$.

Proof. The proof that the elements of D form a spanning set is essentially equivalent to the proof of [KL10, Proposition 3.11]. First, note that any two minimal diagrams for the same matching are equivalent modulo those with fewer crossings (using the relations (3.8f), (3.8g), (4.11a) and (4.11b)). Similarly, moving dots to the chosen positions only introduces diagrams with fewer crossings.

Thus, we only need to show that all minimal diagrams span. Of course, if a diagram is non-minimal, then it can be rewritten in terms of diagrams with fewer crossings, by using the relations to clear all strands out from a bigon, and then making use of (3.6c), (3.6d), (3.7a), (3.7b), (3.8e) and (4.12) to remove it. Thus, by induction, this process must terminate at an expression in terms of minimal diagrams.

The proof that these vectors form a basis proceeds by reducing to the case of $\mathcal{U}$. So suppose not; let $\sum_{d \in D} a_c c = 0$ be a non-trivial linear combination of elements of D, and let d be an element for which $a_d \neq 0$ and the number of crossings and dots in d is maximal among elements with this property.

We will use 0 to denote the constant function that is valued 0 on $[1, \ell]$, and we consider the elements of the Hom-space $\mathrm{HOM}_{\mathcal{T}}((\underline{\boldsymbol{\lambda}}, \mathbf{i}, 0), (\underline{\boldsymbol{\lambda}}, \mathbf{i}', 0))$ obtained from multiplying by $\dot{\theta}_{\kappa'}$ and θ_κ, the elements that sweep thin strands to the right and smooth or wavy strands to the left, exactly as in the proof of [Web13a, Theorem 4.16].

For any basis vector $d \in D$, we can consider $\dot{\theta}_{\kappa'} c \theta_\kappa$. The thin strands still trace out the same element of Khovanov and Lauda's basis, but we have introduced a number of bigons, one for each pair of a thin strand and a smooth or wavy strand such that at least one end of the thin strand is to the left of the smooth/wavy one. Moving the thin strands to the right, we will need to use relations (4.12), (4.11a) and (4.11b). The latter two allow an isotopy through a triple point, modulo terms with fewer crossings. Using the first relation to remove the bigons just adds a fixed number of dots to each strand of a diagram, depending only on its shape (not the number of dots on it originally). This shows that we have obtained a non-trivial linear combination of elements of D in $\mathrm{HOM}_{\mathcal{T}}((\underline{\boldsymbol{\lambda}}, \mathbf{i}, 0), (\underline{\boldsymbol{\lambda}}, \mathbf{i}', 0))$. Thus, it suffices to show that D supplies a basis in this Hom-space, i.e. that the natural map from $\mathrm{HOM}_{\mathcal{U}}(\mathbf{i}, \mathbf{i}')$ is injective.

Now, we use the representation constructed from $\mathcal{Q}$. The space $\mathrm{HOM}_{\mathcal{T}}((\underline{\boldsymbol{\lambda}}, \mathbf{i}, 0), (\underline{\boldsymbol{\lambda}}, \mathbf{i}', 0))$ acts on $\bigoplus_{q \in \mathcal{Q}} \mathrm{HOM}_{\mathcal{U}'}(\mathbf{i}', q)$ by simply forgetting smooth and thin strands and the usual composition of KL diagrams. Thus, the composition of the map $\mathrm{HOM}_{\mathcal{U}}(\mathbf{i}, \mathbf{i}') \to \mathrm{HOM}_{\mathcal{T}}((\underline{\boldsymbol{\lambda}}, \mathbf{i}, 0), (\underline{\boldsymbol{\lambda}}, \mathbf{i}', 0))$ with this action is the usual vertical compostion in $\mathcal{U}$. As long as $\mathcal{Q}$ contains $\mathbf{i}'$, any element of $\mathrm{HOM}_{\mathcal{U}}(\mathbf{i}, \mathbf{i}')$ will act non-trivially by [Web13a, Proposition 3.31], showing the desired injectivity.

Thus, the vectors of D in $\mathrm{HOM}_{\mathcal{T}}((\underline{\boldsymbol{\lambda}}, \mathbf{i}, 0), (\underline{\boldsymbol{\lambda}}, \mathbf{i}', 0))$ are linearly independent in this case, which completes the proof. □

Just as on $\mathcal{U}$, the 2-category $\mathcal{T}$ has an autofunctor flipping diagram which is covariant on 1-morphisms and contravariant on 2-morphisms. Abusing notation, we will denote this also by $\tilde{\psi}$.

5. Tensor product algebras

The category $\mathcal{T}$ is quite auxiliary from our perspective. The fundamental object of this paper is an induced module category over this 2-category. We wish to consider representations of $\mathcal{U}$; for our purposes, this means strict 2-functors $\mathcal{U} \to \mathsf{Cat}$ to the strict 2-category of categories, functors and natural transformations.

Recall that the category $\mathcal{U}$ has a 'trivial' representation on $\mathsf{Vect}_{\Bbbk}$. Every 1-morphism corresponding to a KL pair with $\mathbf{i} \neq \emptyset$ acts by the zero functor, as does the identity 1-morphism of any non-zero weight, while $\mathrm{id}_0 \cdot V \cong V$ for all vector spaces.

Definition 5.1. We let $\mathcal{X}$ denote the 'induction' of this representation to $\mathcal{T}$; that is, an object of $\mathcal{X}$ is a sum of 1-morphisms of $\mathcal{T}$ formally applied to objects of $\mathsf{Vect}_{\Bbbk}$. In addition to the morphisms given by tensor products, we also add a natural isomorphism

$$tu \cdot V \cong t \cdot uV \quad \text{for } t \in \mathrm{Hom}_{\mathcal{T}}(\lambda, \mu),\ u \in \mathrm{Hom}_{\mathcal{U}}(\mu, \nu),\ V \in \mathsf{Ob}(\mathsf{Vect}_{\Bbbk}).$$

Remember that our convention for switching between formulas and diagrams is 'dyslexic': it switches left and right. In essence, then, $\mathcal{X}$ is the quotient of all diagrams in $\mathcal{T}$ (which we view as objects in $\mathcal{X}$ by tensoring with $\Bbbk$ itself) with an $\mathcal{F}_i$ or $\mathcal{E}_i$ or a weight other than 0 at the far *left*, since we can move these over to act (trivially) on the vector space $\Bbbk$. The reader is free to imagine the object $\Bbbk$ as a horde of zombies at the far left of the plane which hungrily eats any thin strand or non-zero weight it can reach, but which is unable to pass through smooth or wavy lines.

Obviously, the only tricolore quadruples that will survive are those where $\mathcal{L} = 0$. Thus, we define a *tricolore triple* to be a tricolore quadruple with the weight $\mathcal{L}$ left out and understood to be 0.

The category $\mathcal{X}$ still carries a $\mathcal{U}$-action by horizontal composition on the right; note that $\mathcal{U}$ is unable to change the labeling or ordering of the smooth and wavy strands.

Definition 5.2. Let $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ denote the subcategory consisting of all 1-morphisms (now thought of as objects of $\mathcal{X}$) where the sequence of labels on smooth and wavy lines is exactly $\underline{\boldsymbol{\lambda}}$. This inherits an action of $\mathcal{U}$ from $\mathcal{X}$. Let $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}_{\mu}$ be the subcategory of $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ where the weight $\mathcal{R}$ is μ.

Recall that the author already defined in [Web13a, § 4.2] a categorification of the tensor product of highest-weight representations, based on certain algebras $T^{\underline{\boldsymbol{\lambda}}}$; these are, in fact, a special case of the categorifications discussed in this paper.

Theorem 5.3. *If $\underline{\boldsymbol{\lambda}}$ consists only of dominant weights, then $\mathcal{X}^{\underline{\boldsymbol{\lambda}}} \cong T^{\underline{\boldsymbol{\lambda}}}$-$\mathsf{pmod}$.*

Proof. If one sums over all tricolore triples, then the resulting object has endomorphism algebra given by the algebra $DT^{\underline{\boldsymbol{\lambda}}}$ defined in [Web13a, Definition 4.27]. This is Morita-equivalent to $T^{\underline{\boldsymbol{\lambda}}}$ by [Web13a, Theorem 4.29]. □

For future results, we must have a precise notion of equivariant morphisms between representations of $\mathcal{U}$. Let $\aleph_1, \aleph_2 : \mathcal{U} \to \mathsf{Cat}$ be two strict 2-functors.

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

DEFINITION 5.4. A *strongly equivariant* functor β is a collection of functors $\beta(\lambda) : \aleph_1(\lambda) \to \aleph_2(\lambda)$ together with natural isomorphisms of functors $c_u : \beta \circ \aleph_1(u) \cong \aleph_2(u) \circ \beta$ for every 1-morphism $u \in \mathcal{U}$ such that

$$c_v \circ (\mathrm{id}_\beta \otimes \aleph_1(\alpha)) = (\aleph_2(\alpha) \otimes \mathrm{id}_\beta) \circ c_u$$

for every 2-morphism $\alpha : u \to v$ in $\mathcal{U}$. (Here we use $\otimes$ for horizontal composition and $\circ$ for vertical composition of 2-morphisms.)

As usual, we let $\mathcal{X}^\lambda = \mathcal{X}^{(\lambda)}$.

THEOREM 5.5. *If $\mathfrak{g}$ is finite-dimensional, we have a strongly $\mathcal{U}$-equivariant equivalence $\mathcal{X}^\lambda \cong \mathcal{X}^{w_0\lambda}$.*

Proof. By symmetry, it suffices to assume that λ is dominant.

Consider the $\mathcal{U}$-module $\mathcal{X}^{w_0\lambda}$. Since the Grothendieck group of this category is an irreducible module, the Jordan–Hölder filtration introduced by Rouquier in [Rou08, § 5] must have a single step; that is, by [Rou08, Theorem 5.8] we have that the category $\mathcal{X}^{w_0\lambda}$ is strongly equivariantly equivalent to a base-change category, which Rouquier denotes by $\mathcal{L}(\lambda) \otimes_{\mathrm{End}(\bar{1}_\lambda)} \mathcal{X}^{w_0\lambda}_\lambda$. By [Web13a, Corollary 3.27], we can unpack this a bit more explicitly. The category $\mathcal{X}^{w_0\lambda}$ is equivalent to the category of projective modules over an algebra $\check{R}^\lambda \otimes_{\check{R}^\lambda_\lambda} A$, where:

- $\check{R}^\lambda$ is a free deformation of T^λ defined and shown to be free in [Web13a, Corollary 3.26], and the deformation base $\check{R}^\lambda_\lambda$ can be identified with a polynomial ring freely generated by the fake bubble endomorphisms of id_λ in $\mathcal{U}$;
- A is an Artinian algebra such that the weight space category $\mathcal{X}^{w_0\lambda}_\lambda$ is equivalent to A-pmod; this inherits a map $\check{R}^\lambda_\lambda \to A$ from the action of the endomorphisms of id_λ in $\mathcal{U}$.

Since $\mathcal{X}^{w_0\lambda}$ has a unique simple module, A is in fact local. By the freeness of $\check{R}^\lambda$ over $\check{R}^\lambda_\lambda$, the base change $\check{R}^\lambda \otimes_{\check{R}^\lambda_\lambda} A$ is free over A, and thus so is any projective over $\check{R}^\lambda \otimes_{\check{R}^\lambda_\lambda} A$. Therefore, the morphism space between any two objects in $\mathcal{X}^{w_0\lambda}$ must be a free A-module.

Since the homomorphism space between the tricolore triples with $\mathbf{i} = \emptyset$ is one-dimensional, this is only possible if $A = \Bbbk$. Thus, $\mathcal{X}^{w_0\lambda}$ is strongly equivariantly equivalent to $\mathcal{X}^\lambda$. □

One very interesting special case is where there is one smooth and one wavy line; assume that $-\lambda$ and μ are dominant and consider $\mathcal{X}^{(\lambda,\mu)}$.

PROPOSITION 5.6. *Every object in $\mathcal{X}^{(\lambda,\mu)}$ is a summand of $(\lambda, \mu, \mathbf{i})$ for some $\mathbf{i}$. The morphism space $(\lambda, \mu, \mathbf{i}) \to (\lambda, \mu, \mathbf{j})$ is the quotient of the morphisms $\mathrm{HOM}_{\mathcal{U}}(\mathbf{i}, \mathbf{j})$ by the relations shown below in (5.13).*

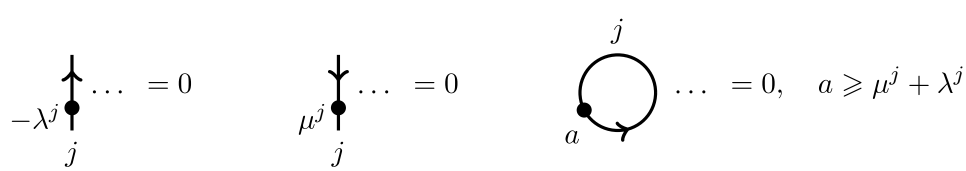

(5.13)

Proof. By definition, any object is a summand of some sequence $(\lambda, \mathbf{i}', \mu, \mathbf{i}'')$. Now let us prove the first statement by induction on the number of pairs in $\mathbf{i}'$ of entries where $i_k \in -\Gamma, i_{k'} \in \Gamma$ and $k < k'$.

If this number is zero, then we can move all elements in $-\Gamma$ left past the wavy strand labeled λ by the Scottish relation in (4.10); similarly, we can move all elements in Γ right past the strand labeled μ by the Polish version of the same relation.

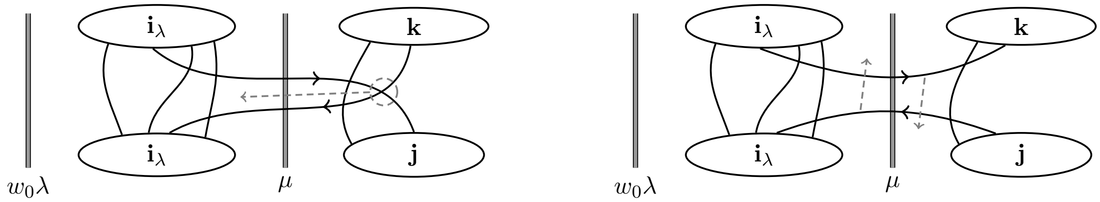

Figure 1. The argument for fullness in Proposition 5.7.

On the other hand, if the number is not 0, there is a pair where k and k' are consecutive. We can apply the relation (3.6c) or (3.7a) to rewrite $e_{(\lambda,\mathbf{i}',\mu,\mathbf{i}'')}$ as factoring through quadruples with a lower number of such pairs.

The relations (5.13) follow immediately from (4.10) and (4.12). The reduction to these relations is essentially the same as in the proof of [Web13a, Theorem 4.20]. The kernel is spanned by diagrams with a strand to the left of both the smooth and wavy strands. If there are multiple such strands, at least one being attached to the top or the bottom, we can use the relations much as in the proof of [Web13a, Theorem 4.20] to remove these until we are left with a single strand to the left of this point, which is thus a consequence of one of the first two relations in (5.13). Otherwise, we can reduce to the case where a single bubble of positive degree is left of all smooth and blus strands, which is thus a consequence of the last relation in (5.13). □

Proposition 5.7. *If $\mathfrak{g}$ is finite-dimensional, we have a strongly $\mathcal{U}$-equivariant equivalence* $\mathcal{X}^{(\lambda,\mu)} \cong \mathcal{X}^{(w_0\lambda,\mu)}$.

Proof. We may as well assume that $-\lambda$ and μ are dominant, since all other cases follow from this one by symmetry.

We have proven in Theorem 5.5 above that $\mathcal{X}^{w_0\lambda}_{\lambda}$ is equivalent to $\mathsf{Vect}_{\Bbbk}$. Let $(w_0\lambda, \mathbf{i}_\lambda)$ be the unique indecomposable object in the λ-weight space of $\mathcal{X}^{w_0\lambda}$.

The equivalence must send the tricolore triple (λ,μ) to $(w_0\lambda,\mathbf{i}_\lambda,\mu)$. Such a functor exists since $(w_0\lambda,\mathbf{i}_\lambda,\mu,i)$ is killed by the $-\lambda^i$th power of the dot on the last $\mathcal{E}_i$ and, similarly, $(w_0\lambda,\mathbf{i}_\lambda,\mu,-i)$ is killed by the μ^ith power of the dot on the last $\mathcal{F}_i$; this confirms (5.13), so a functor exists by Proposition 5.6. Since the ungraded Euler forms of the two categories coincide by Theorem 5.13, we need only prove that this functor is full.

Now, consider a morphism between $(w_0\lambda,\mathbf{i}_\lambda,\mu,\mathbf{j})$ and $(w_0\lambda,\mathbf{i}_\lambda,\mu,\mathbf{k})$, where $\mathbf{j}$ and $\mathbf{k}$ are arbitrary KL pairs with $\mathcal{L} = \mu + \lambda$. We wish to show that this is induced by a 2-morphism in $\mathcal{U}$ from $\mathbf{j} \to \mathbf{k}$.

When we draw the diagram of such a morphism, a terminal in $\mathbf{i}_\lambda$ at the bottom may connect to one in $\mathbf{j}$ or $\mathbf{k}$ (corresponding to either a $\mathcal{E}_i$ or a $\mathcal{F}_i$, respectively). We must show that we can write a diagram of either type in terms of diagrams where strands do not cross the second smooth strand. A diagram where we have a connect to $\mathbf{k}$ must have a crossing between the strand passing from $\mathbf{i}_\lambda$ to $\mathbf{k}$ and one passing from $\mathbf{j}$ to the copy of $\mathbf{i}_\lambda$ at the top. By the freedom we have to choose the spanning set D, this crossing can be assumed to occur to the left of the smooth line for μ by Theorem 4.5. Thus this diagram factors through a tricolore triple of the form $(w_0\lambda,\mathbf{i}_\lambda,-\alpha_i,\dots)$ for some α_i; but $(w_0\lambda,\mathbf{i}_\lambda,-\alpha_i) \cong 0$ since the $\lambda-\lambda_i$ weight space of $V^{\mathbb{Z}}_{w_0\lambda}$ is trivial. Thus, this diagram is 0, and by induction we can write our diagram with no strands from $\mathbf{i}_\lambda$ connecting to $\mathbf{k}$. This argument is represented schematically in the left-hand picture of Figure 1.

Thus we can assume that there is a strand from $\mathbf{i}_\lambda$ at the bottom connecting to $\mathbf{j}$. In this case, there must be at least one strand opposite it which arcs from the top copy of $\mathbf{i}_\lambda$ to $\mathbf{k}$. We can push these strands together using relation (3.6c) left of the strand labeled μ. There are two terms: the correction terms have fewer strands passing from $\mathbf{i}_\lambda$ to $\mathbf{j}$, and the resulting diagram with a bigon factors through $(w_0\lambda, \mathbf{i}_\lambda, -\alpha_i, \dots)$; therefore we can use the argument from above to see that this diagram is 0. This argument is represented schematically in the right-hand picture of Figure 1. This shows the fullness of the functor and completes the proof. □

As in [Web13a, Definition 5.16], we fix an infinite list $\mathbf{p} = \{p_1, p_2, \dots\} \in \Gamma$ of simple roots such that each element of Γ appears infinitely often. For any element v of a highest-weight crystal $\mathcal{B}^\lambda$, there are unique integers $\{a_1, \dots\}$ such that $\cdots \tilde{e}_{p_2}^{a_2}\tilde{e}_{p_1}^{a_1} v = v_{\text{high}}$ and $\tilde{e}_k^{a_k+1}\cdots\tilde{e}_{p_1}^{a_1} v = 0$. The parametrization of the elements of the crystal by this tuple is called the *string parametrization.* We can associate this to a sequence with multiplicities $(\dots, p_2^{(a_2)}, p_1^{(a_1)})$. While this sequence is *a priori* infinite, $a_j = 0$ for all but finitely many j, so upon deleting entries with multiplicity 0 we obtain a finite sequence, which we call the *string parametrization* of the corresponding crystal element. By convention, we write $|\mathbf{a}| = \sum a_i < \infty$.

Let $\boldsymbol{\epsilon}$ be the sign vector such that $\epsilon_k = -1$ if λ_k is dominant and $\epsilon_k = 1$ if it is anti-dominant. For an ℓ-tuple of words $(\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(\ell)})$, we let $\mathsf{l}(\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(\ell)})$ be the tricolore triple such that the thin block after the jth smooth or wavy strand is the sequence associated to the word $a^{(j)}$, with upward strands if λ_j is anti-dominant or downward strands if it is dominant. More formally, this is the tricolore triple with our chosen $\underline{\boldsymbol{\lambda}}$,

$$\mathbf{i} = (\dots, \epsilon_1 p_2^{(a_2^{(1)})}, \epsilon_1 p_1^{(a_1^{(1)})}, \dots, \epsilon_2 p_2^{(a_2^{(2)})}, \epsilon_2 p_1^{(a_1^{(2)})}, \dots, \epsilon_\ell p_2^{(a_2^{(\ell)})}, \epsilon_\ell p_1^{(a_1^{(\ell)})}),$$

and $\kappa(j) = |\mathbf{a}^{(1)}| + \cdots + |\mathbf{a}^{(j-1)}|$.

DEFINITION 5.8. We define an ordering on compositions of length ℓ, called the *reverse dominance order*, by $\nu \geqslant \nu'$ if and only if $\sum_{k=j}^{\ell} \nu'_k \geqslant \sum_{k=j}^{\ell} \nu_k$ for all $j \in [1, \ell]$. If $|\nu| = |\nu'|$, then this coincides with the usual dominance order.

We shall order ℓ-tuples of words by reverse dominance order on the composition $|\mathbf{a}^\bullet|$ given by taking sums of each word, with ties broken by lexicographic order on $\mathbf{a}^{(\ell)}$, then lexicographic order on $\mathbf{a}^{(\ell-1)}$, and so on.

We define a *stringy tricolore triple* for the sequence $\mathbf{p}$ to be $\mathsf{l}(\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(\ell)})$, where $\mathbf{a}^{(k)}$ is the string parametrization for an element of the crystal of highest or lowest weight λ.

Since we will use this fact many times, let us remind the reader that in a graded category where the degree-0 part of the endomorphisms of any object are finite-dimensional (a condition satisfied by $\mathcal{U}, \mathcal{T}$ and $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$), an object is indecomposable if and only if its endomorphism algebra is graded local, i.e. has a unique maximal homogeneous ideal.

LEMMA 5.9. *Every indecomposable object of $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ is isomorphic to a summand of $\mathsf{l}(\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(\ell)})$ for some ℓ-tuple $(\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(\ell)})$.*

Proof. Since every element of Γ occurs infinitely often, every sequence in $\epsilon_p\Gamma$ occurs attached to some word $\mathbf{a}^{(p)}$. Thus we need to show that every indecomposable is a summand of an object of a specific form, where the thin block to the right of a smooth strand is all downward and that after a wavy strand is all upward. This is essentially the proof of Theorem 5.3 with a slightly more delicate induction. We induct on the statement that every indecomposable is a summand of an tricolore triple where the m rightmost thin blocks have the desired form. We can start with a tricolore triple for which this is true of the $m-1$ rightmost thin blocks and concentrate on the mth from the right. As in the proof of Theorem 5.3, the relations (3.6c), (3.6d) and (4.10) allow us to push badly oriented strands further left until they are out of the mth thin block. Thus, we are done. □

Lemma 5.10. *The object* $\mathsf{I}(\mathbf{a}^{(1)}, \ldots, \mathbf{a}^{(\ell)})$ *has at most one summand that is not a summand of the tricolore triple for a greater word, and has no such summand unless this triple is stringy.*

Proof. First, we claim that it is sufficient to show this for $\Bbbk$ being a field. We let $\mathbb{r} = \Bbbk/\mathfrak{m}$ be the residue field of $\Bbbk$ and assume that the theorem holds in this case. We have a natural functor $\mathcal{X}^{\boldsymbol{\lambda}}_{\Bbbk} \to \mathcal{X}^{\boldsymbol{\lambda}}_{\mathbb{r}}$ given by simply killing $\mathfrak{m}$. By Hensel's lemma, this functor induces a bijection between indecomposable projectives and between summands of a given tricolore sequence. This reduces us to the case where $\Bbbk$ is a field.

Now, fix an ℓ-tuple of words $(\mathbf{a}^{(1)}, \ldots, \mathbf{a}^{(\ell)})$. Let I be the two-sided ideal in $\mathrm{End}(\mathsf{I}(\mathbf{a}^{\bullet}))$ generated by elements factoring through $\mathsf{I}(\mathbf{b}^{(1)}, \ldots, \mathbf{b}^{(\ell)})$ such that $\mathbf{b}^{\bullet} > \mathbf{a}^{\bullet}$. We wish to show by induction that the quotient $\mathrm{End}(\mathsf{I}(\mathbf{a}^{\bullet}))/I$ is graded local.

We can assume that $\mathbf{a}^{(\ell)} \neq 0$, since if $\mathbf{a}^{(\ell)} = 0$, we can simply remove the rightmost smooth or wavy strand without changing $\mathrm{End}(\mathsf{I}(\mathbf{a}^{\bullet}))$ or I, and reduce to the case where we have $\ell - 1$ smooth and wavy strands. Let q be minimal such that $a = a_q^{(\ell)} \neq 0$. Thus, the rightmost part of the tricolore triple $\mathsf{I}(\mathbf{a}^{\bullet})$ consists of a strands labeled p_q. Let $\mathbf{c}^{\bullet}$ be the sequence of words which coincides with $\mathbf{a}^{\bullet}$, except that $\mathbf{c}_q^{(\ell)} = 0$. By induction, we can assume that $\mathrm{End}(\mathsf{I}(\mathbf{c}^{\bullet}))/I_{\mathbf{c}}$ is graded local. Let R_a be the ring of symmetric functions in a variables, which acts by natural endomorphisms on the functor $\mathcal{F}_{i_q}^{a}$ as symmetric polynomials in the dots; this is also graded local. Thus, we have a natural map $\phi : \mathrm{End}(\mathsf{I}(\mathbf{c}^{\bullet}))/I_{\mathbf{c}} \otimes R_a \to \mathrm{End}(\mathsf{I}(\mathbf{a}^{\bullet}))/I$ where the first term acts on all but the rightmost a strands and the second acts on the last a strands. If we prove that ϕ is surjective, then we will know that $\mathrm{End}(\mathsf{I}(\mathbf{a}^{\bullet}))/I$ is necessarily graded local.

We can divide the matchings with top and bottom $\mathsf{I}(\mathbf{a}^{\bullet})$ whose associated diagrams are not in the image of ϕ into the following two categories:

(1) those where one of the last a strands is connected by a cap to the bottom;
(2) those where there is no such cap, but one of the last a strands is connected to another node at the top.

If we choose our basis vectors D carefully, we can assume that the diagrams associated to these matchings lie in I. In case (1), this is because any diagram having a single cup connecting one of the last a strands with the bottom of the diagram and no other crossings has bottom smaller than the top in reverse dominance order. For any matching of type (1), we can choose the basis vector so that the bottom portion of the diagram is of this form. In case (2), we can push every crossing between strands from the rightmost a at the top and the rightmost a at the bottom to the far right; after doing this, there will be a slice through the middle of the diagram which corresponds to a tricolore triple for a word with $a_q^{(\ell)} > a$, and this is higher in our order. Thus, ϕ is indeed surjective, and we find that we have at most one new projective.

Now, we need only show that if $\mathsf{I}(\mathbf{a}^{\bullet})$ is not stringy, then this quotient is trivial. If $\mathbf{a}^{(k)}$ is not a string parametrization, the identity of $\mathsf{I}(\mathbf{a}^{\bullet})$ can be rewritten as factoring through triples where $\mathbf{a}^{(k)}$ is replaced by higher words in lexicographic order, or where one of the strands is pulled left through the kth smooth or wavy strand. Both these are higher in the order, so this triple has no 'new' summands. □

We say that an object of a $\Bbbk$-linear category is *absolutely indecomposable* if it remains indecomposable after extension by any local ring homomorphism $\Bbbk \to \Bbbk'$ between complete local rings. If $\Bbbk$ is a field, this just means that the object remains indecomposable under field extensions.

Corollary 5.11. *Every indecomposable object of $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ is absolutely indecomposable.*

Proof. Ring extension sends the object $\mathsf{I}(\mathbf{a}^{\bullet})$ to the same object over the same ring. Fix a local homomorphism $\Bbbk \to \Bbbk'$ and let $\mathbf{a}^{\bullet}$ be the greatest stringy sequence whose associated indecomposable M splits after extension to $\Bbbk'$. None of the summands of $M \otimes_{\Bbbk} \Bbbk'$ appears in a greater stringy sequence, so $\mathsf{I}(\mathbf{a}^{\bullet})$ has either two new indecomposable summands or a new summand with multiplicity greater than 1. Both of these cases are impossible by Lemma 5.10, so we have arrived at a contradiction. □

This result can easily be extended to the 2-category $\mathcal{U}$. We say that a 1-morphism $(\mathbf{i}, -\mathbf{j})$ in $\dot{\mathcal{U}}$ is *stringy* if $\mathbf{i}$ and $\mathbf{j}$ are both positive string parametrizations for elements of the crystal $B(-\infty)$. We can endow these with a similar variant of lexicographic order, i.e. provided that we fix the weight of the 1-morphism, then $(\mathbf{i}, -\mathbf{j}) > (\mathbf{i}', -\mathbf{j}')$ if:

- $|\mathbf{i}| < |\mathbf{i}'|$; or
- $|\mathbf{i}| = |\mathbf{i}'|$ and $\mathbf{j} > \mathbf{j}'$ in lexicographic order on the corresponding words; or
- $\mathbf{j} = \mathbf{j}'$ and $\mathbf{i} > \mathbf{i}'$ in lexicographic order on the corresponding words.

The proof can be extended to show the following result.

Proposition 5.12. *The 1-morphism $(\mathbf{i}, -\mathbf{j})$ in $\dot{\mathcal{U}}$ has at most one summand that is not a summand of any higher stringy sequence, unless $(\mathbf{i}, -\mathbf{j})$ is itself stringy. This summand is absolutely indecomposable and every indecomposable appears in this way for a unique stringy sequence.*

As in [Web13a, § 4.7], we can define vectors $v_{\mathbf{i}}^{\kappa} = v_{\mathsf{I}}$ in $V_{\underline{\boldsymbol{\lambda}}}^{\mathbb{Z}}$ inductively as follows.

- If $\kappa(\ell) = n$, then $v_{\mathbf{i}}^{\kappa} = v_{\mathbf{i}}^{\kappa^-} \otimes v_{\ell}$, where v_{ℓ} is the highest-weight (respectively, lowest-weight) vector of V_{λ_ℓ} if λ_ℓ is dominant (respectively, anti-dominant), and κ^- is the restriction to $[1, \ell - 1]$.
- If $\kappa(\ell) \neq n$, then $v_{\mathbf{i}}^{\kappa} = E_{i_n} v_{\mathbf{i}^-}^{\kappa}$, where $\mathbf{i}^- = (i_1, \ldots, i_{n-1})$, using the convention that $F_i = E_{-i}$.

Theorem 5.13. *The ungraded Grothendieck group $K^0(\mathcal{X}^{\underline{\boldsymbol{\lambda}}})$ is isomorphic to $\bar{V}_{\underline{\boldsymbol{\lambda}}}^{\mathbb{Z}}$ via the map sending $[(\underline{\boldsymbol{\lambda}}, \mathbf{i}, \kappa)]$ to $v_{\mathbf{i}}^{\kappa}$. This map intertwines the Euler form with the factorwise Shapovalov form $\langle -, - \rangle_s$ on the tensor product.*

Remark 5.14. Note that the comparable theorem for $\underline{\boldsymbol{\lambda}}$ all of highest weight [Web13a, Theorem 4.38] does not require setting $q = 1$. This is because we already knew the form on the quantum tensor product we expected to match the Euler form with, whereas here we do not know *a priori* that a suitable bilinear form exists on a tensor product. We will establish later that this result holds with q not specialized.

Proof. Once we have proved the equality of dimensions

$$\sum \dim \mathrm{HOM}(\mathsf{I}, \mathsf{I}') = \langle v_{\mathsf{I}}, v_{\mathsf{I}'} \rangle_s, \tag{5.14}$$

the proof is precisely the same as that of [Web13a, Theorem 4.38] with its associated lemmata, which we leave as an exercise for the reader.

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

The proof of (5.14) is also quite similar, but requires small changes. We need only prove that (5.14) holds when the tricolore triples are of the form $\mathsf{l}(\mathbf{a}^{\bullet})$. As in [Web13a, Theorem 4.38], the induction is easier to swing if we allow one extra strand which points in the 'wrong' direction.

We induct on the reverse dominance order on words. If l and l' both have $\mathbf{a}^{(\ell)} = 0$, then (5.14) will hold after we remove the rightmost strand, which is equivalent to the desired case of (5.14). If one of l and l' has $\mathbf{a}^{(\ell)} \neq 0$, then we can apply the adjunction to pull a strand from one side to the other, and then slide it left using the relations (3.6c) and (4.10). □

This in particular shows that the classes of stringy sequences are linearly independent, and so the 'new' summand of each stringy sequence must be non-zero.

Corollary 5.15. *Each stringy sequence in $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ or $\mathcal{U}$ has exactly one indecomposable non-zero summand which is not isomorphic to any summand of a larger sequence.*

The autofunctor $\tilde{\psi}$ obviously preserves violating morphisms and thus descends to an involution on $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$, which we denote by $\tilde{\psi}^{\underline{\boldsymbol{\lambda}}}$.

This functor defines an involution on $V^{\underline{\boldsymbol{\lambda}}}$ for any $\underline{\boldsymbol{\lambda}}$. We will denote this involution by $\psi^{\underline{\boldsymbol{\lambda}}}$.

Proposition 5.16. *For each indecomposable projective P in $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ (respectively, $\dot{\mathcal{U}}$), there is a unique grading shift $P(n)$ such that $\tilde{\psi}^{\underline{\boldsymbol{\lambda}}}(P(n)) \cong P(n)$ (respectively, $\tilde{\psi}(P(n)) \cong P(n)$).*

Proof. Such a shift is obviously unique, so we need only prove that it exists. There is a unique n such that $P(n)$ is a summand of the corresponding stringy sequence. Since the latter module is self-dual, $\tilde{\psi}^{\underline{\boldsymbol{\lambda}}}(P(n))$ is a summand of it, and by the uniqueness of Proposition 5.12 we must have $\tilde{\psi}^{\underline{\boldsymbol{\lambda}}}(P(n)) \cong P(n)$. □

These results show the following fact.

Theorem 5.17. *The categories $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ and $\mathcal{U}$ are humorous categories with the obvious grading shift and with dualities given by $\tilde{\psi}^{\underline{\boldsymbol{\lambda}}}$ and $\tilde{\psi}$.*

6. Representation categories and standard modules

As in [Web13a], it will be useful to deal with an abelian category, not just an additive one. In particular (as far as the author is aware), this is necessary for checking that the Grothendieck group of $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ is the tensor product representation as a representation of $U_q(\mathfrak{g})$; we have thus far only checked that this holds at $q = 1$.

Definition 6.1. Let $\mathfrak{V}^{\underline{\boldsymbol{\lambda}}} := \mathrm{Rep}(\mathcal{X}^{\underline{\boldsymbol{\lambda}}})$. Let $\mathscr{Y} : \mathcal{X}^{\underline{\boldsymbol{\lambda}}} \to \mathfrak{V}^{\underline{\boldsymbol{\lambda}}}$ be the Yoneda embedding $\mathsf{l} \mapsto \mathrm{Hom}(\mathsf{l}, -)$.

Note that we do *not* require an object in $\mathfrak{V}^{\underline{\boldsymbol{\lambda}}}$ to be finitely generated.

Definition 6.2. Let $H^{\underline{\boldsymbol{\lambda}}}$ be the subring of the opposite endomorphism ring $\mathrm{End}_{\mathcal{X}^{\underline{\boldsymbol{\lambda}}}}(\bigoplus_{\mathsf{l}} \mathsf{l})^{\mathrm{op}}$ which kills all but finitely many summands.

Note that this is a non-unital ring; by a $H^{\underline{\boldsymbol{\lambda}}}$-module M, we always mean one which is the direct sum of the images of the idempotents $e_{\mathbf{i},\underline{\boldsymbol{\lambda}},\kappa}$. Since $H^{\underline{\boldsymbol{\lambda}}}$ is locally unital for the system of idempotents $e_{\mathbf{i},\underline{\boldsymbol{\lambda}},\kappa}$, this is the natural generalization of the condition that the identity of a ring must act by the identity on a module.

We can interpret an object in $\mathfrak{V}^{\underline{\boldsymbol{\lambda}}}$ as a module over $H^{\underline{\boldsymbol{\lambda}}}$ using the obvious functor $\mathcal{X}^{\underline{\boldsymbol{\lambda}}} \to H^{\underline{\boldsymbol{\lambda}}}\text{-}\mathsf{pmod}$ given by the morphism space $X \mapsto \mathrm{HOM}_{\mathcal{X}}(\bigoplus_{\mathsf{l}} \mathsf{l}, -)$. Of course, $\mathfrak{V}^{\underline{\boldsymbol{\lambda}}}$ is an abelian

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

category, and since $\mathscr{Y}(P)$ is projective for any $P \in \mathsf{Ob}(\mathcal{X}^{\underline{\boldsymbol{\lambda}}})$, $\mathfrak{V}^{\underline{\boldsymbol{\lambda}}}$ has enough projectives. However, it is not clear that if M is finitely generated and P is a finitely generated projective with a surjection $P \to M$, then the kernel of this map is finitely generated.

We can define an action of $\mathcal{U}$ on $\mathfrak{V}^{\underline{\boldsymbol{\lambda}}}$ by exact functors using the biadjunction between $\mathcal{E}_i$ and $\mathcal{F}_i$ as a definition, i.e.

$$\mathcal{F}_i \cdot M(\mathsf{I}) := M(\mathcal{E}_i \mathsf{I}(\langle -\alpha_i, \mu\rangle + 1)),$$
$$\mathcal{E}_i \cdot M(\mathsf{I}) := M(\mathcal{F}_i \mathsf{I}(\langle -\alpha_i, \mu\rangle + 1)).$$

Note the switch of $\mathcal{E}_i$ and $\mathcal{F}_i$ above; this is what is required so that the Yoneda embedding intertwines the categorification functors, since $\mathcal{E}_i$ and $\mathcal{F}_i$ are biadjoint on $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$.

Theorem 6.3. *The Grothendieck group $K^0(\mathfrak{V}^{\underline{\boldsymbol{\lambda}}})$ is isomorphic to the lattice dual to $\bar{V}^{\mathbb{Z}}_{\underline{\boldsymbol{\lambda}}}$, with the map induced by $\mathscr{Y}$ given by the Shapovalov form.*

Since we have twisted our action by the Cartan involution, the Shapovalov pairing defines a map of representations. If $\mathfrak{g}$ is infinite-dimensional, then we take the full dual; that is, as abstract abelian groups, $\bar{V}^{\mathbb{Z}}_{\underline{\boldsymbol{\lambda}}}$ is a direct *sum* of copies of $\mathbb{Z}$ while $K^0(\mathfrak{V}^{\underline{\boldsymbol{\lambda}}})$ is a direct *product*. We let $\widehat{V}_{\underline{\boldsymbol{\lambda}}} = K^0(\mathfrak{V}^{\underline{\boldsymbol{\lambda}}}) \otimes_{\mathbb{Z}} \mathbb{C}$; this is a $\mathfrak{g}$ representation defined by taking the direct product of the weight spaces of a highest-weight representation, rather than the direct sum.

Even in finite type, the Shapovalov form is not always unimodular over the integers, so it will not usually coincide with $K^0(\mathcal{X}^{\underline{\boldsymbol{\lambda}}})$; this will only happen if all entries in $\underline{\boldsymbol{\lambda}}$ are minuscule and $\mathfrak{g}$ is finite-dimensional. In particular, this shows that $\mathfrak{V}^{\underline{\boldsymbol{\lambda}}}$ extremely rarely has finite global dimension, since that is only possible when these lattices coincide.

For a tricolore triple $\mathsf{I} = (\mathbf{i}, \underline{\boldsymbol{\lambda}}, \kappa)$, the number of thin strands in each thin block defines a composition, which we denote by ν_{I}. We also have a function $\boldsymbol{\alpha}_{\mathsf{I}} : [1, \ell] \to X$ given by the sum of the roots labeling each thin block. For each such function $\boldsymbol{\alpha}$, we have a map $\wp : R_{\alpha(1)} \otimes \cdots \otimes R_{\alpha(\ell)} \to H^{\underline{\boldsymbol{\lambda}}}$ sending an ℓ-tuple of KL diagrams to their horizontal composition with smooth and wavy lines added, as in [Web13a, (4.3)].

Definition 6.4. The standard representation S_{I} is the maximal quotient of $\mathscr{Y}(\mathsf{I})$ such that $S_{\mathsf{I}}(\mathsf{I}') = 0$ if $\nu_{\mathsf{I}'} > \nu_{\mathsf{I}}$ in the reverse dominance order on compositions.

More generally, we let the *standardization* S_M of an object M in $\mathcal{X}^{\lambda_1}_{\lambda_1 - \alpha(1)} \times \cdots \times \mathcal{X}^{\lambda_\ell}_{\lambda_\ell - \alpha(\ell)}$ for some fixed $\boldsymbol{\alpha}$ be the initial object in $\mathfrak{V}^{\underline{\boldsymbol{\lambda}}}$ among those such that $S_M(\bigoplus_{\boldsymbol{\alpha}_{\mathsf{I}} = \boldsymbol{\alpha}} \mathsf{I}) \cong M$ as $R_{\alpha(1)} \otimes \cdots \otimes R_{\alpha(\ell)}$-modules (via $\wp$) and $S_M(\mathsf{I}') = 0$ if $\nu_{\mathsf{I}'} > \nu_{\mathsf{I}}$.

We think of the relation induced from reverse dominance order on compositions as a pre-order of the set of sequences I. In terms of $H^{\underline{\boldsymbol{\lambda}}}$-modules, this module has a presentation much like that of the standard modules of [Web13a]: one can define S_{I} as a quotient of $\mathscr{Y}(\mathsf{I})$ by the image of every map from $\mathscr{Y}(\mathsf{I}')$ with $\mathsf{I}' > \mathsf{I}$. In terms of Stendhal diagrams, this means that we quotient $H^{\underline{\boldsymbol{\lambda}}} e_{\mathsf{I}}$ by all diagrams which are *standardly violated*,[7] that is, diagrams whose bottom is I and where some horizontal slice $y = a$ is I' with $\mathsf{I}' > \mathsf{I}$.

As in [Web13a], we let ${}_k\mathcal{E}_i$ and ${}_k\mathcal{F}_i$ denote the categorification functors on the kth factor of $\mathcal{X}^{\lambda_1} \times \cdots \times \mathcal{X}^{\lambda_\ell}$.

Lemma 6.5. *The module $\mathcal{E}_i S_M$ has a natural filtration $Q_1 \supset Q_2 \supset \cdots$ such that*

$$Q_k / Q_{k+1} \cong S_{{}_k\mathcal{E}_i M}.$$

[7] This is the reflection of the definition of standardly violated from [Web13a], since we are looking at left standard modules rather than right modules.

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

The module $\mathcal{F}_i S_M$ has a natural filtration $O_m \supset O_{m-1} \supset \cdots$ such that

$$O_k/O_{k-1} \cong S_{{}_k\mathcal{F}_i M}.$$

Proof. Since this proof is quite close to that of [Web13a, Proposition 5.5], we will only give a short sketch covering the points to be changed (which themselves are almost the same as the changes made for the proof of Theorem 5.13). The construction of the filtrations and the surjective maps from standardizations is exactly as in [Web13a, Proposition 5.5]; thus we need only show that the successive quotients have the correct dimensions. That is, we need to prove that for any I,

$$\dim \mathcal{E}_i S_M(\mathsf{I}) = \sum \dim S_{{}_k\mathcal{E}_i M}(\mathsf{I}), \tag{6.15}$$

$$\dim \mathcal{F}_i S_M(\mathsf{I}) = \sum \dim S_{{}_k\mathcal{F}_i M}(\mathsf{I}). \tag{6.16}$$

This must be shown via an induction over I in reverse dominance order (rather than weight). Note that both I and S_M have associated compositions, which do not coincide. Our induction step is to prove (6.15) and (6.16) where the composition for either I or S_M coincides with ν, assuming that the equations hold when either of the compositions is greater than ν in reverse dominance order.

As usual, we can assume that if λ_ℓ is dominant (respectively, anti-dominant), then all strands after the last smooth or wavy strand are downwards (respectively, upwards). If the leftmost strand in I is smooth or wavy, the equations (6.15) and (6.16) follow from the case where this strand is removed, so we can assume that the leftmost strand is thin.

If λ_ℓ is dominant, then we can apply the argument in the proof of [Web13a, Proposition 5.5]. We note that

$$\dim \mathcal{F}_i S_M(\mathsf{I}) = \dim \mathcal{E}_j \mathcal{F}_i S_M(\mathsf{I}') \leqslant \sum \dim S_{{}_p\mathcal{E}_{j\,k}\mathcal{F}_i M}(\mathsf{I}'), \tag{6.17}$$

which implies that

$$\dim \mathcal{F}_i \mathcal{E}_j S_M(\mathsf{I}') \leqslant \sum \dim S_{{}_k\mathcal{F}_{i\,p}\mathcal{E}_j M}(\mathsf{I}'), \tag{6.18}$$

with equality holding in both (6.17) and (6.18) or neither. It holds in (6.18) by induction, since $\mathcal{E}_i$ applied to a tricolore triple with λ_ℓ dominant can always be rewritten as a summand of triples higher in reverse dominance order. Thus, we also have equality in (6.17), which is only possible if (6.15) and (6.16) hold as well (since we already know they are inequalities).

If λ_ℓ is an anti-dominant map, we just apply the same argument with $\mathcal{F}_i$ and $\mathcal{E}_i$ reversed. This is possible since now $\mathcal{F}_i$ applied to this triple will be a summand of triples higher in reverse dominance order. □

If we write I as the concatenation of tricolore triples I_i consisting of one smooth or wavy strand and then the thin block, we let

$$s_{\mathsf{I}} = v_{\mathsf{I}_1} \otimes \cdots \otimes v_{\mathsf{I}_\ell}.$$

However, if we calculate the grading shifts of these filtrations, a new wrinkle appears. If the new line we add is wavy, the grading shifts will behave as though we have tensored the representations in the opposite order. In other words, this can be stated as follows.

Proposition 6.6. *As representations of $U_q(\mathfrak{g})$, we have an isomorphism*

$$K_q^0(\mathcal{X}^{(\lambda_1,\dots,\lambda_\ell)}) \cong \begin{cases} K_q^0(\mathcal{X}^{(\lambda_1,\dots,\lambda_{\ell-1})}) \otimes K_q^0(\mathcal{X}^{(\lambda_\ell)}) & \text{for } \lambda_\ell \text{ dominant,} \\ K_q^0(\mathcal{X}^{(\lambda_1,\dots,\lambda_{\ell-1})}) \otimes^{\mathrm{op}} K_q^0(\mathcal{X}^{(\lambda_\ell)}) & \text{for } -\lambda_\ell \text{ dominant,} \end{cases}$$

induced by the functor

$$\mathbb{S} : \mathcal{X}^{(\lambda_1,\dots,\lambda_{\ell-1})} \times \mathcal{X}^{(\lambda_\ell)} \to \mathcal{X}^{\underline{\boldsymbol{\lambda}}}.$$

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

Proof. We obtain a basis of $K^0(\mathcal{X}^{\underline{\boldsymbol{\lambda}}})$ by taking the stringy sequences. Each of these has a standard quotient, and the matrix giving the multiplicities of standards in the stringy basis is upper-triangular, with ones on the diagonal. Thus, it has an inverse with the same property, which we can use to define classes $[S_{\mathsf{i}}]$ in $K^0(\mathcal{X}^{\underline{\boldsymbol{\lambda}}})$ which form a basis. Since $\mathbb{S}$ sends pairs of these standard quotients to standard quotients, it sends a basis to a basis and thus defines an isomorphism. Therefore, we need only check how E_i and F_i act.

The proof for λ_ℓ dominant is essentially the same as in [Web13a]: one can consider the action on a standardization $\mathbb{S}(M_1, M_2)$ from $\mathcal{X}^{(\lambda_1,\dots,\lambda_{\ell-1})} \times \mathcal{X}^{(\lambda_\ell)}$. The module $\mathcal{F}_i\mathbb{S}(M_1, M_2)$ has a submodule N generated by

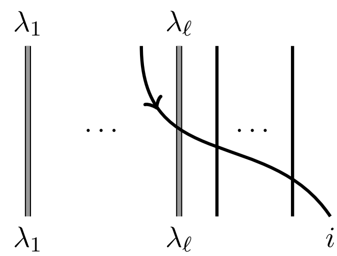

(6.19)

and this submodule N is isomorphic to the standardization $\mathbb{S}(\mathcal{F}_iM_1, M_2)$ but with a shift by the degree of the element (6.19), which is $\mu_\ell^i = \alpha_i^\vee(\mu_\ell)$ where μ_ℓ is the weight of M_2. The quotient $\mathcal{F}_i\mathbb{S}(M_1, M_2)/N$ is isomorphic to $\mathbb{S}(M_1, \mathcal{F}_iM_2)$; that is, we have an equality

$$\begin{aligned}[\mathcal{F}_i\mathbb{S}(M_1, M_2)] &= q^{\mu_\ell^i}[\mathbb{S}(\mathcal{F}_iM_1, M_2)] + [\mathbb{S}(M_1, \mathcal{F}_iM_2)] \\ &= q^{\mu_\ell^i}(F_i[M_1]) \otimes [M_2] + [M_1] \otimes (F_i[M_2]) = \Delta(F_i)([M_1] \otimes [M_2]).\end{aligned}$$

On the other hand, the module $\mathcal{E}_i\mathbb{S}(M_1, M_2)$ has a similar submodule generated by

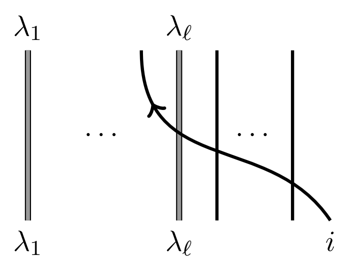

(6.20)

Since the element (6.19) has degree 0, the submodule N it generates is isomorphic to $\mathbb{S}(\mathcal{E}_iM_1, M_2)$ with no grading shift, whereas the quotient is isomorphic to $\mathcal{E}_i\mathbb{S}(M_1, M_2)/N \cong \mathbb{S}(M_1, \mathcal{E}_iM_2)(\mu^i - \mu_\ell^i)$, since every diagram in the quotient has a cap to the right of the leftmost smooth strand, and the degree of this cap decreases by $\mu^i - \mu_\ell^i$ when the label on the surrounding region changes from μ to μ_ℓ. Thus, we have that

$$\begin{aligned}[\mathcal{E}_i\mathbb{S}(M_1, M_2)] &= [\mathbb{S}(\mathcal{E}_iM_1, M_2)] + q^{\mu_\ell^i - \mu^i}[\mathbb{S}(M_1, \mathcal{E}_iM_2)] \\ &= (E_i[M_1]) \otimes [M_2] + q^{\mu_\ell^i - \mu^i}[M_1] \otimes (E_i[M_2]) = \Delta(E_i)([M_1] \otimes [M_2]).\end{aligned}$$

This shows the result for λ dominant. The result in the case where λ is anti-dominant follows from the same argument with the places of $\mathcal{E}_i$ and $\mathcal{F}_i$ switched. In fact, there is an equivalence of categories, $\mathcal{X}^{\underline{\boldsymbol{\lambda}}} \cong \mathcal{X}^{-\underline{\boldsymbol{\lambda}}}$, which reverses the orientation on every thin strand and switches smooth and thin strands (one must multiply all thin/thin crossings by -1). This functor interchanges the action of $\mathcal{E}_i$ and $\mathcal{F}_i$, and thus is compatible with the Cartan involution ω. Since $(\omega \otimes \omega) \circ \Delta \circ \omega = \Delta^{\mathrm{op}}$, this shows that the action when λ is anti-dominant is given by the opposite coproduct. □

Thus, applying this formula inductively, we obtain the following result.

Proposition 6.7. *There is a unique isomorphism of $U_q(\mathfrak{g})$-modules $K^0(\mathcal{X}^{\underline{\boldsymbol{\lambda}}}) \cong V^{\mathbb{Z}}_{\underline{\boldsymbol{\lambda}}}$ sending $[\mathsf{I}] \mapsto v_{\mathsf{I}}$ and $[S_{\mathsf{I}}] \mapsto s_{\mathsf{I}}$.*

This result also shows that the category of objects filtered by standards is invariant under the action of $\mathcal{F}_i, \mathcal{E}_i$ and the addition of smooth and wavy lines. In particular, this shows that the objects $\mathscr{Y}(\mathsf{I})$ all have standard filtrations.

We refer to [DPS98, 1.2.4] for the definition of a *strict stratifying system*. While in that paper they only consider the case of a finite quasi-poset, their definition makes sense even for an countably infinite quasi-poset, such as compositions endowed with reverse dominance order. Similarly, for us a *standard stratification* will be allowed to be indexed by a countably infinite quasi-poset such that the interval between any two elements is finite; in terms of [CPS96], this means allowing a stratification of an algebra to be an infinite chain of ideals $\cdots \subset J_i \subset J_{i+1} \subset \cdots$ with $i \in \mathbb{Z}$ such that $\bigcup J_i = A$ and $\bigcap J_i = \{0\}$, rather than a finite chain of this type.

Corollary 6.8. *The objects S_{I} with the induced pre-order define a strict stratifying system of $\mathfrak{V}_{\underline{\boldsymbol{\lambda}}}$; thus, they define a standard stratification of the algebra $H^{\underline{\boldsymbol{\lambda}}}$. In particular, every indecomposable projective in $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ has a filtration by standardizations of projective indecomposables in $\mathcal{X}^{\lambda_1}_{\lambda_1-\alpha(1)} \times \cdots \times \mathcal{X}^{\lambda_\ell}_{\lambda_\ell-\alpha(\ell)}$.*

We note that outside finite type, the standards will typically be infinite-dimensional (assuming that both smooth and wavy strands are used). However, there are only finitely many compositions of any size larger than a fixed one in reverse dominance order, and there are only finitely many sequences with a given composition. Thus, only finitely many standards will occur in the stratification of $\mathscr{Y}(\mathsf{I})$.

7. Orthodox bases in higher representation theory

In this paper, we have developed the theory of categorifications of $\dot{U}$ and its representations with a particular application in mind: constructing bases for these representations. We remind the reader that we have fixed a complete local ring $\Bbbk$ and polynomials Q_{ij}. We should note that the bases we consider depend in an essential way on the ring and polynomials chosen.

Definition 7.1. Let C denote the set of indecomposable $\tilde{\psi}$-invariant 1-morphisms (up to shift) in $\dot{\mathcal{U}}$ and let $C_{\underline{\boldsymbol{\lambda}}}$ be the set of $\tilde{\psi}^{\underline{\boldsymbol{\lambda}}}$-invariant objects of $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$.

Let the *orthodox basis* $\{o_P = [P]\}_{P\in C}$ of $\dot{U}$ be defined by classes of $\tilde{\psi}$-invariant indecomposables under the isomorphism $K^0(\dot{\mathcal{U}}) \cong \dot{U}$. Similarly, the orthodox basis $\{o_P = [P]\}_{P\in C_{\underline{\boldsymbol{\lambda}}}}$ of $V^{\mathbb{Z}}_{\underline{\boldsymbol{\lambda}}}$ is defined by $\tilde{\psi}^{\underline{\boldsymbol{\lambda}}}$-invariant indecomposable classes of $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$.

The orthodox bases of $\dot{U}$ and its representations carry over a surprising amount of structure that occurs for canonical bases.

We have already shown in Theorem 5.17 that $\dot{\mathcal{U}}$ and $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ are humorous categories. Thus, by Lemma 1.6, we have a pre-canonical structure on the Grothendieck groups of these categories.

Definition 7.2. The *orthodox pre-canonical structure* on the vector spaces $\dot{U}$ and $V^{\mathbb{Z}}_{\underline{\boldsymbol{\lambda}}}$ is defined as in Lemma 1.6 as follows.

- The bar-involution is given by ψ (respectively, $\psi^{\underline{\boldsymbol{\lambda}}}$).
- The inner product is given by the Euler form $\langle [M], [N]\rangle = \sum_i q^{-i} \dim \operatorname{Hom}(M, N)$.
- The standard basis a_c is given by the classes of tricolore triples attached to stringy sequences.

In several cases, the bar-involutions and inner products of these pre-canonical strctures include previously defined structures. Lusztig has defined bar-involutions and inner products on the following structures.

(1) The modified quantized enveloping algebra $\dot{U}$: the bar-involution $^-$ is defined in [Lus93, 23.1.8], and in [Lus93, 26.1.2] he defines a bilinear form $(-,-)$; we wish to consider the induced sesquilinear form $\langle u,v\rangle := (\bar{u},v)$.
(2) The tensor product $V^{\mathbb{Z}}_{-\lambda,\mu}$ for λ and μ both dominant: the bar-involution Ψ is defined in [Lus93, 24.3.2] and the bilinear form $(-,-)_{\lambda,\mu}$ in [Lus93, 26.2.1]; as above, we take the induced sesquilinear form $\langle u,v\rangle = (\Psi(u),v)_{\lambda,\mu}$.
(3) The tensor product $V^{\mathbb{Z}}_{\underline{\boldsymbol{\lambda}}}$ if all λ_i are dominant and $\mathfrak{g}$ is finite-dimensional: we can use the bar-involution Ψ defined in [Lus93, 27.3.1] for the tensor product of these as based modules; we define a bilinear form on these spaces as the tensor product of the bilinear forms $(-,-)_{0,\lambda_i}$, and then turn this into a sesquilinear form $\langle u,v\rangle = (\Psi(u),v)_{\lambda_*,0}$.

Proposition 7.3. *The orthodox bar-involution and inner product agree with Lusztig's in the cases* (1)–(3) *above.*

Proof.

(1) The involution $^-$ on $\dot{U}$ is distinguished by fixing monomials in E_i and F_i. The same is true of that induced by $\tilde{\psi}$. The agreement of forms follows from Theorem 3.4.
(2) The involution Ψ on $V_{-\lambda}\otimes V_\mu$ is the unique involution which satisfies

$$\Psi(u\cdot(v_{-\lambda}\otimes v_\mu)) = \bar{u}\cdot(v_{-\lambda}\otimes v_\mu).$$

Thus we need only show that $\tilde{\psi}^{-\lambda,\mu}$ satisfies this property. Flipping over a diagram commutes with acting on the right side of a wavy and smooth strand $(-\lambda,\mu)$ with it, so $\tilde{\psi}^{-\lambda,\mu}$ and $\tilde{\psi}$ (on $\mathcal{U}$) are compatible.
For both forms, we have $\langle v_{-\lambda}\otimes v_\mu, v_{-\lambda}\otimes v_\mu\rangle = 1$ and $\langle uv,w\rangle = \langle v,\tau(u)w\rangle$ where τ is the q-anti-linear anti-automorphism defined by

$$\tau(E_i) = q_i^{-1}\tilde{K}_{-i}F_i, \quad \tau(F_i) = q_i^{-1}\tilde{K}_iE_i, \quad \tau(K_\mu) = K_{-\mu}.$$

These properties characterize the form uniquely, so they must agree.
(3) Lusztig's bar-involution Ψ on a tensor product is compatible with the action of $U_q(\mathfrak{g})$ on the tensor product, and if $v\in M'$ is a highest-weight vector, then it also commutes with the inclusion $M\to M\otimes M'$ sending $m\mapsto m\otimes v$. Furthermore, if M' is irreducible, Ψ is uniquely characterized by these properties, since $M\otimes\{v\}$ generates $M\otimes M'$ over $U_q(\mathfrak{g})$. Note that, in particular, the vectors v_{l} defined above are all Ψ-invariant and span the space $V_{\underline{\boldsymbol{\lambda}}}$; thus any other anti-linear map that fixes these vectors must be Ψ.
Thus, we need only check that $\tilde{\psi}^{\underline{\boldsymbol{\lambda}}}$ categorifies a bar-involution that fixes the same vectors v_{l}. The tricolore triples l are obviously $\tilde{\psi}^{\underline{\boldsymbol{\lambda}}}$-invariant (essentially by definition), and so the result follows from Theorem 5.13. The match of the Euler form and Lusztig's form in this case is precisely [Web13a, 3.22]. □

Our standard bases, however, are not the same as those typically used by Lusztig; luckily, as noted before, the dependence of canonical bases on standard bases is very weak, so we can still show that Lusztig's bases are canonical bases in our sense.

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

Theorem 7.4. *The bases defined by Lusztig in cases* (1)–(3) *above are canonical bases for the orthodox pre-canonical structure, in the sense of Definition 1.7.*

It is worth noting that, unlike Theorem A, Theorem 7.4 does not require a symmetric Cartan matrix. It holds for any symmetrizable Kac–Moody algebra, in particular for finite-dimensional Lie algebras of type BCDFG.

Proof. In cases (1)–(3), the canonical basis of Lusztig satisfies the almost orthogonality conditions. This is proved for cases (1) and (2) in [Lus93, 26.3.1], and the same argument extends easily to case (3). Therefore, as argued in Lemma 1.9, any canonical basis vector for the orthodox pre-canonical structure must either lie in Lusztig's basis, or its negative must.

As in [Lus93, 14.4.2], the property that distinguishes Lusztig's basis from its more easily found signed version is compatibility with the action of E_i and F_i. First, as a base case, this basis contains (in the respective cases):

(1) the vector 1_ν;
(2) the vector $v_{-\lambda} \otimes v_\mu$;
(3) all vectors of the form $B_d \otimes v_{\lambda_\ell}$ for the canonical basis of $V^{\mathbb{Z}}_{\lambda_1} \otimes \cdots \otimes V^{\mathbb{Z}}_{\lambda_{\ell-1}}$.

There is also an inductive piece of the definition: the indexing set of the canonical basis in cases (1)–(3) can be identified with the corresponding set of stringy sequences. If such a sequence does not end with $\pm$ a simple root (that is, with a thin strand at its far right), then it belongs to one of the base cases above. Otherwise, its last term is of the form $i^{(n)}$ for some $n \geqslant 0$ and $i \in \pm\Gamma$. If c is an element of this indexing set, let c' be the object indexed by the stringy sequence with this last appearance of $i^{(n)}$ deleted.

Using the convention that $F_i = E_{-i}$, the positivity condition for the corresponding Lusztig basis vector B_c is that

$$B_c \in E_i^{(n)} B_{c'} + \sum_{d<c} \mathbb{Z}[q, q^{-1}] B_d.$$

This is explicit in the case of a basis vector in U^+ by [Lus93, 14.3.2(c) and 14.4.2], and easily carries over to the more general cases (1)–(3), which use the basis of U^+ in their definition.

The base case is easily established for the canonical bases of the orthodox pre-canonical structure. The only case that is not tautological is (3), where we must work by induction and assume that we have proven the theorem for tensor products with $\ell - 1$ factors. Once this is assumed, we need only note that $- \otimes v_\lambda$ commutes with Ψ, sends stringy sequences to stringy sequences and preserves the inner product, and thus sends canonical basis vectors to canonical basis vectors.

Now consider a minimal Lusztig canonical basis vector such that $B_c \notin v_{\mathsf{l}_c} + \sum_{d<c} \mathbb{Z}[q, q^{-1}] \cdot v_{\mathsf{l}_d}$, where we let l_d denote the stringy sequence for a crystal element d. By minimality, we must have $B_{c'} \in v_{\mathsf{l}_{c'}} + \sum_{d'<c'} \mathbb{Z}[q, q^{-1}] \cdot v_{\mathsf{l}'_d}$. Then

$$E_i^{(n)} B_{c'} \in E_i^{(n)} v_{\mathsf{l}_{c'}} + \sum_{d'<c'} \mathbb{Z}[q, q^{-1}] \cdot E_i^{(n)} v_{\mathsf{l}'_d}.$$

Now, by definition, $E_i^{(n)} v_{\mathsf{l}_{c'}} = v_{\mathsf{l}_c}$, and use of the lexicographic ordering implies that

$$\sum_{d'<c'} \mathbb{Z}[q, q^{-1}] \cdot E_i^{(n)} v_{\mathsf{l}'_d} \subset \sum_{d<c} \mathbb{Z}[q, q^{-1}] \cdot v_{\mathsf{l}_d}.$$

Thus we must have that

$$B_c - v_{\mathsf{l}_c} = (B_c - E_i^{(n)} B_{c'}) + (E_i^{(n)} B_{c'} - v_{\mathsf{l}_c}) \in \sum_{d<c} \mathbb{Z}[q, q^{-1}] \cdot v_{\mathsf{l}_d},$$

which is a contradiction. This shows that Lusztig's basis is also canonical in our sense. □

Orthodox bases and canonical bases sometimes coincide and sometimes do not. For future comparisons between them, we note that the only issue is the condition (III).

Proposition 7.5. *The bases* o_P *satisfy conditions* (I) *and* (II) *of Definition 1.7 for the orthodox pre-canonical structure.*

Proof. Condition (I) is clear from the definition. Condition (II) follows immediately from Lemma 5.10 and Proposition 5.12. □

Furthermore, there is at least one property in which orthodox bases are an improvement over canonical bases.

Proposition 7.6. *For any* $\Bbbk$ *and* Q_{ij}*, the structure coefficients of multiplication in* $\dot{U}$ *and the matrix coefficients for the action on* $V_{\underline{\boldsymbol{\lambda}}}^{\mathbb{Z}}$ *where the orthodox basis is used are in* $\mathbb{Z}_{\geqslant 0}[q, q^{-1}]$.

Similarly, by Corollary 6.8, we have the following.

Proposition 7.7. *For any* $\Bbbk$ *and* Q_{ij}*, the coefficients of any orthodox basis vector for* $V_{\underline{\boldsymbol{\lambda}}}^{\mathbb{Z}}$ *in terms of pure tensor products of orthodox basis vectors in the factors are in* $\mathbb{Z}_{\geqslant 0}[q, q^{-1}]$.

Very loosely, the difference between orthodox and canonical bases lies in the trading-off of positivity in coefficients for positivity in exponents of q; Lusztig's basis is defined in a way that depends strongly on the latter, at the cost of positivity of coefficients in the non-symmetric case.

In fact, the dependence of this basis on the base ring $\Bbbk$ is quite crude; the corresponding basis depends only on the characteristic of the residue field, by Corollary 1.13. Thus, we can assume that $\Bbbk$ is generated by the coefficients of Q_{ij} over a prime field. If these coefficients are integers, then we need only consider the prime fields. As shown in Proposition 1.14, the orthodox bases over different fields may not coincide, but those for positive characteristic are positive linear combinations of the bases for characteristic 0.

Example 7.8. Perhaps the easiest example where these bases differ is when $\mathfrak{g} = \widehat{\mathfrak{sl}}_2$; in this case, we have chosen the polynomial $Q_{01}(u, v) = u^2 - 2uv + v^2$. Consider the object in $\mathcal{U}$ given by $\mathcal{F}_1\mathcal{F}_0\mathcal{F}_1\mathcal{F}_0$. This has only a two-dimensional space of degree-0 endomorphisms, spanned by the following two elements.

$1 =$ 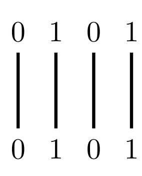

$\psi_2\psi_3\psi_1\psi_2 =$ 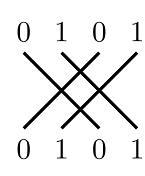

One can easily calculate that $(\psi_2\psi_3\psi_1\psi_2)^2 = 2\psi_2\psi_3\psi_1\psi_2$.

Thus, if 2 is a unit in the ring $\Bbbk$, we have that $\frac{1}{2}(\psi_2\psi_3\psi_1\psi_2)$ is a primitive idempotent and that $\mathrm{End}_0(\mathcal{F}_1\mathcal{F}_0\mathcal{F}_1\mathcal{F}_0) \cong \Bbbk \oplus \Bbbk$ (so this object is the sum of two distinct summands). On the other hand, if $\Bbbk$ has characteristic 2, then the same calculation shows that $\psi_2\psi_3\psi_1\psi_2$ is nilpotent. We have an isomorphism $\mathrm{End}_0(\mathcal{F}_1\mathcal{F}_0\mathcal{F}_1\mathcal{F}_0) \cong \Bbbk[t]/(t^2)$, so this object is indecomposable.

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

This example equally shows the dependence of the orthodox basis on the choice of Q_{ij}: for any ring $\Bbbk$, if $Q_{01}(u,v)=u^2+v^2$, then $\mathcal{F}_1\mathcal{F}_0\mathcal{F}_1\mathcal{F}_0$ is indecomposable, as in the characteristic-2 case. Note that in this case, the diagram $\psi_3\psi_1\psi_2$ gives a map of degree -2 from $\mathcal{F}_1\mathcal{F}_0\mathcal{F}_1\mathcal{F}_0$ to $\mathcal{F}_1^{(2)}\mathcal{F}_0^{(2)}$, which is indecomposable for any choice of Q_{01}. This shows that there is a negative-degree map between indecomposable $\tilde{\psi}$-invariant projectives.

The same example is treated by Tingley and the author in [TW12, § 3.5] and by Kashiwara in [Kas12, Example 3.3] from the dual perspective (in terms of simples rather than projectives). A variant displaying similar behavior was considered by Khovanov and Lauda in [KL09, 3.25].

This categorification framework provides us with a wealth of bases, which are actually quite difficult to study in general. For $U_q^+(\widehat{\mathfrak{sl}}_n)$, the orthodox bases of the basic representation V_{ω_0} (for the choice of Q_{ij} fixed in [Web13a, § 9]) over $\Bbbk=\mathbb{F}_p$ were defined by Grojnowski as 'p-canonical bases' [Gro99, § 14.1]. The equivalence of our approach and that of Grojnowski is shown by the 'Main Theorem' of [BK09b]. This shows that the simple modules over the symmetric groups over a field of characteristic p give the dual orthodox basis of the basic representation of $\mathfrak{sl}_p$ over the field $\mathbb{F}_p$; the determination of these classes is one of the most important questions in modular representation theory.

Similarly, in finite type, examples where the canonical and orthodox bases do not coincide were recently described by Williamson in [Wil12]; in fact, for any prime p, the orthodox basis of $U(\mathfrak{sl}_{8p-1})$ from characteristic p differs from the canonical basis.

Of course, if we are given any representation which seems to have a natural choice of categorification, we can use this to define an orthodox basis. At the moment, the most obvious example is the case where $\widehat{\mathfrak{sl}}_e$ and the associated representation is a higher-level Fock space. As shown in [Sha11], the categories $\mathcal{O}$ of symplectic reflection algebras provide one such categorification. By recent independent work of the author [Web13b], Rouquier *et al.* [RSVV13] and Losev [Los13], these have a graded lift where the classes of projectives give a canonical basis; in [Web13b], we also give a diagrammatic description of these categories that fits more with the philosophy of this paper.

8. Canonical bases in higher representation theory

In this section, we consider the question of when the orthodox and canonical bases coincide.

Definition 8.1. An orthodox basis of $\dot{U}$ or $V_{\underline{\boldsymbol{\lambda}}}^{\mathbb{Z}}$ is said to be *canonical* if its elements are *almost orthonormal*, i.e. $\langle o_P, o_{P'}\rangle \in \delta_{P,P'}+q^{-1}\mathbb{Z}_{\geqslant 0}[[q^{-1}]]$ for all $P,P'\in C$ or $C_{\underline{\boldsymbol{\lambda}}}$, that is, when the orthodox basis is canonical in the sense of § 1.

Note that this is *a priori* stronger than (III); but in fact, for any orthodox basis, we have $\langle o_P, o_{P'}\rangle \in \mathbb{Z}_{\geqslant 0}((q))$, so it is an equivalent condition. By Theorem 7.4, this is also the same as requiring the orthodox bases to match Lusztig's basis in the cases (1)–(3) discussed there.

Remark 8.2. We point out that the methods of Kashiwara [Kas12, 3.1-2] can be easily extended to show that if the orthodox basis of a $\dot{U}$ or $V_{\underline{\boldsymbol{\lambda}}}^{\mathbb{Z}}$ is canonical for some value of $Q_{*,*}$ over $\Bbbk$, then the same holds for generic $Q_{*,*}$, that is, when we take the coefficients of $Q_{*,*}$ to be formal variables on a space $\mathfrak{Q}$ of possible choices and replace $\Bbbk$ by $\Bbbk(\mathfrak{Q})$.

Proposition 8.3. *The orthodox basis of $V_{-\lambda}^{\mathbb{Z}}$, $V_{\mu}^{\mathbb{Z}}$ or $V_{-\lambda,\mu}^{\mathbb{Z}}$ is a crystal basis if and only if it is canonical.*

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

Proof. Any of these modules has at most one bar-invariant crystal basis, which in these cases coincides with Lusztig's basis. Since the orthodox basis is bar-invariant by Proposition 7.5, the result follows. □

The next result follows immediately from Lemma 1.15.

Proposition 8.4. *The orthodox basis of $\dot{U}$ or $V^{\mathbb{Z}}_{\underline{\boldsymbol{\lambda}}}$ is canonical if and only if the categorification $\dot{\mathcal{U}}$ or $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ is mixed in the sense of Definition 1.11.*

Thus, the question of when orthodox bases are canonical reduces to computing when categorifications are mixed. As suggested by the name, typically this is proved using relations to geometry: one shows that there is a functor with nice properties sending $\tilde{\psi}$-invariant objects in one's categorification to perverse sheaves on some space, and deduces positivity of the grading from the fact that perverse sheaves are the heart of a t-structure. This connection with geometry holds only in certain situations. In Example 7.8 we showed that if $\mathfrak{g} = \widehat{\mathfrak{sl}}_2$ and $Q_{01}(u, v) = u^2 + v^2$, then the category $\mathcal{U}$ cannot be mixed.

For the rest of the paper, we will assume that $\mathfrak{g}$ has symmetric Cartan matrix (so that we may use quiver varieties) and $\Bbbk$ is a field of characteristic 0 (so that we may use the decomposition theorem), and we fix a particular choice of $Q_{*,*}$, which coincides with the choice used in [VV11, § 3.3] and [Rou08, § 3.2.4]. This choice is forced upon us by geometry and is of the following nature: we choose an orientation Ω on our Dynkin diagram, let ϵ_{ij} denote the number of edges oriented from i to j, and fix

$$Q_{ij}(u, v) = (-1)^{\epsilon_{ij}}(u - v)^{c_{ij}}. \tag{8.21}$$

Note that these hypotheses include $\mathfrak{g} = \widehat{\mathfrak{sl}}_2$ but with $\pm Q_{01}(u, v) = u^2 - 2uv + v^2$.

Assuming these hypotheses, the result [VV11, 4.5] says the following (in different language).

Theorem 8.5. *The orthodox basis of $V^{\mathbb{Z}}_{\lambda}$ is canonical.*

The proof of this fact can be rephrased in terms of Lemma 1.17. We apply Lemma 1.20 to the category $\mathcal{J}$ of constructible sheaves on a moduli space of quiver representations *à la* Lusztig generated by certain pushforwards. Vasserot and Varagnolo showed that the associated mixed category $\mathcal{J}$ is equivalent to the projective modules over the KLR algebra [VV11, 3.6]. Thus, the functor of dividing by the cyclotomic ideal defines a functor $\mathcal{J} \to \mathcal{X}^{\lambda}$ intertwining Verdier duality with the duality $\tilde{\psi}$, to which we can apply Lemma 1.17.

Remark 8.6. Very similar arguments were also used by the author and Stroppel in their study of quiver Schur algebras [SW11], in order to show that the indecomposable projectives over these algebras correspond to canonical bases of higher-level Fock spaces; in fact, Proposition 8.4 and Lemma 1.17 seem to be applicable in essentially any categorical $\mathfrak{g}$-module yet dreamed up. The difficult part is to understand the relevant pre-canonical structure in terms of previously understood representation theory.

We can extend this proof to the tensor product of highest-weight representations using an extension of Vasserot and Varagnolo's geometric techniques. This relies on a more general result on certain generalizations of KLR algebras called *weighted KLR algebras*.

Proposition 8.7 [Web12b, 4.9]. *Assume that $\mathfrak{g}$ has symmetric Cartan matrix, $\Bbbk$ is a field of characteristic 0, and $Q_{*,*}$ is as in (8.21). Then, the algebra $\tilde{T}^{\boldsymbol{\lambda}}_{\mu}$ defined in [Web13a] is isomorphic to the Ext-algebra of an object Y in the constructible derived category of a moduli space of quiver representations, denoted by $E_{\lambda-\mu}/G'_{\lambda-\mu}$, which is a sum of shifts of semi-simple perverse* *sheaves. This isomorphism intertwines the duality $\tilde{\psi}$ for $\tilde{T}^{\underline{\boldsymbol{\lambda}}}_{\mu}$-modules and Verdier duality on the constructible derived category.*

Theorem 8.8. *Assume that $\mathfrak{g}$ has symmetric Cartan matrix, all the λ_i are dominant, $\Bbbk$ is a field of characteristic 0, and $Q_{*,*}$ is as in (8.21). Then the orthodox basis of $V^{\mathbb{Z}}_{\underline{\boldsymbol{\lambda}}}$ is canonical.*

Proof. We apply Lemma 1.17 with $\mathcal{C}$ being the sums of shifts of the summands of Y, morphisms given by Exts in the constructible derived category and the grading given by homological grading. This category is a mixed humorous category by Lemma 1.20.

Since this category is equivalent to the graded projective modules over $\tilde{T}^{\underline{\boldsymbol{\lambda}}}_{\mu}$, dividing by the violating ideal defines a full and essentially surjective functor. Pre-composed with the inverse of the equivalence of Proposition 8.7, this gives a functor $\mathcal{C} \to \mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ which satisfies the assumptions of Lemma 1.17 and thus shows that $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ is mixed. Lemma 1.15 shows that the orthodox basis for $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$ is canonical. □

The reader familiar with Lusztig's construction of this basis for tensor products might wonder where his standard basis, the pure tensor products of canonical basis vectors in the factors, has disappeared to. Of course, by Theorem 8.5, these are the same as the pure tensor products of orthodox basis vectors, and thus are given by the classes of standardizations of indecomposable projectives by Proposition 6.7; in the context of standardly stratified categories, these would usually just be called the 'standards'. We can also define a pre-canonical structure with the same bar and form as the orthodox pre-canonical structure, but using these pure tensor products as a standard basis. This pre-canonical structure will be almost orthonormal under the hypotheses of Theorem 8.8.

Thus, by Lemma 1.8, for *this* standard basis we could use condition (II′) from § 1 as the definition of the canonical basis, as Lusztig does. In terms of representation theory, the coefficients of the canonical basis in terms of the pure tensors are multiplicities of the standard filtration on indecomposable tensor products. These are obviously positive integer Laurent polynomials, and condition (II′) corresponds to the fact that only positive grading shifts of standards occur.

Combining this observation with Proposition 7.7, we see that the following holds.

Corollary 8.9. *The canonical basis of $V^{\mathbb{Z}}_{\underline{\boldsymbol{\lambda}}}$ or $V^{\mathbb{Z}}_{-\lambda,\mu}$ is a linear combination of pure tensors of canonical basis elements with coefficients in $\mathbb{Z}_{\geqslant 0}[q^{-1}]$.*

Up to this point, we have had to be very careful about choosing our polynomials $Q_{*,*}$. Example 7.8 and the example from [KL09, § 3.3] show that this choice is very important in the affine case. However, in type ADE, it is much ado about nothing.

Lemma 8.10. *If $\mathfrak{g}$ is finite-dimensional and simply laced, then all choices of $Q_{*,*}$ result in equivalent categories $\mathcal{U}$ and $\mathcal{X}^{\underline{\boldsymbol{\lambda}}}$.*

Proof. The argument is precisely the same as that given in [KL11, p. 17] for KLR algebras. Since $\mathfrak{g}$ is simply laced, the polynomial $Q_{*,*}$ is determined uniquely by t_{ij}. We simply note that if the products $t_{ij}t_{ji}^{-1}$ coincide with those for another choice t'_{ij}, then these algebras are isomorphic by a rescaling of the crossings between differently colored strands. Furthermore, if we multiply $t_{ij}t_{ji}^{-1}$ by a coboundary in $\Bbbk^*$, then we get an algebra isomorphic by rescaling like colored crossings and dots. Since all 1-cocycles on a tree are coboundary, we are done. □

Corollary 8.11. *Assume $\mathfrak{g}$ is finite-dimensional and simply laced, $\Bbbk$ is a field of characteristic 0 and $Q_{*,*}$ is arbitrary. Then, the orthodox basis of $\dot{U}$ or $V^{\mathbb{Z}}_{\underline{\boldsymbol{\lambda}}}$ is canonical.*

Proof. First, we can replace $\Bbbk$ by its algebraic closure $\bar{\Bbbk}$, by Theorem 1.13. By Lemma 8.10, we can reduce to the case where $Q_{*,*}$ is as in (8.21); Theorem 8.8 thus establishes the case of $V^{\mathbb{Z}}_{\underline{\boldsymbol{\lambda}}}$.

Now, we consider $\dot{U}$. By Proposition 5.7, the orthodox bases of $V^{\mathbb{Z}}_{\lambda,\mu}$ and $V^{\mathbb{Z}}_{w_0\lambda,\mu}$ coincide. By Theorem 8.8, the former coincides with the canonical basis as well, so the same is true of $V^{\mathbb{Z}}_{w_0\lambda,\mu}$. Since the canonical basis of $\dot{U}$ is uniquely determined by the fact that it lands on the canonical basis of $V^{\mathbb{Z}}_{w_0\lambda,\mu}$ under the map $u \mapsto u \cdot (v_{-\lambda} \otimes v_\mu)$, Lemma 1.18 shows that the categorification $\dot{\mathcal{U}}$ is mixed and the orthodox basis agrees with the canonical basis. □

Acknowledgements

I thank George Lusztig; without his questioning, this paper might never have happened. Ben Elias was also helpful far beyond the call of duty in giving comments and suggesting revisions. I also thank Jon Brundan, Yiqiang Li, Aaron Lauda, Ivan Losev, Marco Mackaay, Catharina Stroppel, Éric Vasserot and Weiqiang Wang for useful discussions.

References

AS13 P. N. Achar and C. Stroppel, *Completions of Grothendieck groups*, Bull. Lond. Math. Soc. (N.S.) **45** (2013), 200–212.

ÁDL03 I. Ágoston, V. Dlab and E. Lukács, *Quasi-hereditary extension algebras*, Algebr. Represent. Theory **6** (2003), 97–117; MR 1960515 (2004c:16010).

Ari09 S. Ariki, *Graded q-Schur algebras*, Preprint (2009), arXiv:0903.3453.

BW14 H. Bao and W. Wang, *Canonical bases in tensor products revisited*, Preprint (2014), arXiv:1403.0039.

BGS96 A. Beilinson, V. Ginzburg and W. Soergel, *Koszul duality patterns in representation theory*, J. Amer. Math. Soc. **9** (1996), 473–527.

BLPW10 T. Braden, A. Licata, N. Proudfoot and B. Webster, *Gale duality and Koszul duality*, Adv. Math. **225** (2010), 2002–2049.

BLPW12 T. Braden, A. Licata, N. Proudfoot and B. Webster, *Hypertoric category $\mathcal{O}$*, Adv. Math. **231** (2012), 1487–1545.

BK09a J. Brundan and A. Kleshchev, *Graded decomposition numbers for cyclotomic Hecke algebras*, Adv. Math. **222** (2009), 1883–1942.

BK09b J. Brundan and A. Kleshchev, *Blocks of cyclotomic Hecke algebras and Khovanov–Lauda algebras*, Invent. Math. **178** (2009), 451–484.

CL11 S. Cautis and A. Lauda, *Implicit structure in 2-representations of quantum groups*, Preprint (2011), arXiv:1111.1431.

CPS96 E. Cline, B. Parshall and L. Scott, *Stratifying endomorphism algebras*, Mem. Amer. Math. Soc. **124** (1996), no. 591.

DPS98 J. Du, B. Parshall and L. Scott, *Stratifying endomorphism algebras associated to Hecke algebras*, J. Algebra **203** (1998), 169–210; MR 1620729 (99e:20006).

EW12 B. Elias and G. Williamson, *The Hodge theory of Soergel bimodules*, Preprint (2012), arXiv:1212.0791.

Gro99 I. Grojnowski, *Affine $\mathfrak{sl}_p$ controls the representation theory of the symmetric group and related Hecke algebras*, Preprint (1999), arXiv:math/9907129.

JKK09 L. Jeffrey, Y.-H. Kiem and F. Kirwan, *On the cohomology of hyperKähler quotients*, Transform. Groups **14** (2009), 801–823.

KK12 S.-J. Kang and M. Kashiwara, *Categorification of highest weight modules via Khovanov–Lauda–Rouquier algebras*, Invent. Math. **190** (2012), 699–742.

Kas12 M. Kashiwara, *Notes on parameters of quiver Hecke algebras*, Preprint (2012), arXiv:1204.1776.

KL79 D. Kazhdan and G. Lusztig, *Representations of Coxeter groups and Hecke algebras*, Invent. Math. **53** (1979), 165–184.

KL09 M. Khovanov and A. D. Lauda, *A diagrammatic approach to categorification of quantum groups. I*, Represent. Theory **13** (2009), 309–347.

KL10 M. Khovanov and A. D. Lauda, *A categorification of quantum* sl(n), Quantum Topol. **1** (2010), 1–92; MR 2628852 (2011g:17028).

KL11 M. Khovanov and A. D. Lauda, *A diagrammatic approach to categorification of quantum groups II*, Trans. Amer. Math. Soc. **363** (2011), 2685–2700; MR 2763732 (2012a:17021).

LLT96 A. Lascoux, B. Leclerc and J.-Y. Thibon, *Hecke algebras at roots of unity and crystal bases of quantum affine algebras*, Comm. Math. Phys. **181** (1996), 205–263.

Lau10 A. D. Lauda, *A categorification of quantum* sl(2), Adv. Math. **225** (2010), 3327–3424; MR 2729010 (2012b:17036).

Los13 I. Losev, *Proof of Varagnolo-Vasserot conjecture on cyclotomic categories* $\mathcal{O}$, Preprint (2013), arXiv:1305.4894.

Lus90 G. Lusztig, *Canonical bases arising from quantized enveloping algebras*, J. Amer. Math. Soc. **3** (1990), 447–498.

Lus92 G. Lusztig, *Canonical bases in tensor products*, Proc. Natl. Acad. Sci. USA **89** (1992), 8177–8179.

Lus93 G. Lusztig, *Introduction to quantum groups*, Progress in Mathematics, vol. 110 (Birkhäuser, Boston, 1993).

Rou08 R. Rouquier, *2-Kac–Moody algebras*, Preprint (2008), arXiv:0812.5023.

RSVV13 R. Rouquier, P. Shan, M. Varagnolo and E. Vasserot, *Categorifications and cyclotomic rational double affine Hecke algebras*, Preprint (2013), arXiv:1305.4456.

Sha11 P. Shan, *Crystals of Fock spaces and cyclotomic rational double affine Hecke algebras*, Ann. Sci. Éc. Norm. Supér. (4) **44** (2011), 147–182; MR 2760196 (2012c:20009).

SVV11 P. Shan, M. Varagnolo and E. Vasserot, *Koszul duality of affine Kac–Moody algebras and cyclotomic rational DAHA*, Preprint (2011), arXiv:1107.0146.

Soe92 W. Soergel, *The combinatorics of Harish-Chandra bimodules*, J. Reine Angew. Math. **429** (1992), 49–74; MR 1173115 (94b:17011).

Spr82 T. A. Springer, *Quelques applications de la cohomologie d'intersection*, in *Bourbaki Seminar, Vol. 1981/1982*, Astérisque, vol. 92 (Société Mathématique de France, Paris, 1982), 249–273; MR 689533 (85i:32016b).

Sto11 M. Stošić, *Indecomposable 1-morphisms of* $\dot{\mathcal{U}}_3^+$ *and the canonical basis of* $U_q^+(\mathfrak{sl}_3)$, Preprint (2011), arXiv:1105.4458.

SW11 C. Stroppel and B. Webster, *Quiver Schur algebras and* q*-Fock space*, Preprint (2011), arXiv:1110.1115.

Sus08 J. Sussan, *Positivity and the canonical basis of tensor products of finite-dimensional irreducible representations of quantum* $\mathfrak{sl}(k)$, Preprint (2008), arXiv:0804.2034.

TW12 P. Tingley and B. Webster, *Mirković–Vilonen polytopes and Khovanov–Lauda–Rouquier algebras*, Preprint (2012), arXiv:1210.6921.

Ugl00 D. Uglov, *Canonical bases of higher-level* q*-deformed Fock spaces and Kazhdan–Lusztig polynomials*, in *Physical combinatorics (Kyoto, 1999)*, Progress in Mathematics, vol. 191 (Birkhäuser, Boston, 2000), 249–299.

VV11 M. Varagnolo and E. Vasserot, *Canonical bases and KLR-algebras*, J. Reine Angew. Math. **659** (2011), 67–100.

Web12a B. Webster, *A categorical action on quantized quiver varieties*, Preprint (2012), arXiv:1208.5957.

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press

Web12b B. Webster, *Weighted Khovanov–Lauda–Rouquier algebras*, Preprint (2012), arXiv:1209.2463.

Web13a B. Webster, *Knot invariants and higher representation theory*, Preprint (2013), arXiv:1309.3796.

Web13b B. Webster, *Rouquier's conjecture and diagrammatic algebra*, Preprint (2013), arXiv:1306.0074.

Wil12 G. Williamson, *On an analogue of the James conjecture*, Preprint (2012), arXiv:1212.0794.

Ben Webster bwebster@virginia.edu
Department of Mathematics, University of Virginia, Charlottesville, VA 22908, USA

https://doi.org/10.1112/S0010437X1400760X Published online by Cambridge University Press