The noncommutative Springer resolution in type A and KLRW algebras

Ben Webster

University of Waterloo
Perimeter Institute for Theoretical Physics

January 5, 2023

UNIVERSITY OF WATERLOO
FACULTY OF MATHEMATICS
Department of Pure Mathematics

PERIMETER INSTITUT
INSTITUTE FOR THEORETICAL PHYSICS
PÉRIMÈTRE DE PHYSIQUE THÉORIQUE

Let $X = T^*\mathrm{Fl}_n$ be the cotangent bundle of the flag variety $X_0 = \mathrm{Fl}_n$ over a field $\Bbbk$ of characteristic $p \geq 0$.

Let $\mathsf{Coh}_0(X)$ denote the abelian category of coherent sheaves on X which are (set-theoretically) supported on X_0.

Non-commutative Springer resolution

Consider the algebra $A = U\mathfrak{gl}_n(\Bbbk)$. Let $\mathfrak{U}\text{-mod}_0$ be the principal block of the category of finite dimensional modules with central character.

HC

Let $X = T^*\mathrm{Fl}_n$ be the cotangent bundle of the flag variety $X_0 = \mathrm{Fl}_n$ over a field $\Bbbk$ of characteristic $p \geq 0$.

Let $\mathsf{Coh}_0(X)$ denote the abelian category of coherent sheaves on X which are (set-theoretically) supported on X_0.

Consider the algebra $A = U\mathfrak{gl}_n(\Bbbk)$. Let $\mathfrak{U}$-mod_0 be the principal block of the category of finite dimensional modules with central character.

Theorem (Bezrukavnikov-Mirkovič)

If $p \gg 0$, there is an equivalence of derived categories

$$D^b(\mathsf{Coh}_0(X)) \cong D^b(\mathcal{U}).$$

Bezrukavnikov calls this a "non-commutative counterpart of the Springer resolution."

This is a beautiful equivalence, but it's quite abstract. I want to give you a somewhat more concrete way of thinking about it.

$\mathsf{Coh}_0(X)$	$\mathfrak{U}\text{-mod}_0$
geometry	representation theory

Bezrukavnikov calls this a "non-commutative counterpart of the Springer resolution."

This is a beautiful equivalence, but it's quite abstract. I want to give you a somewhat more concrete way of thinking about it.

$\mathsf{Coh}_0(X)$	$\mathfrak{U}\text{-mod}_0$
geometry	representation theory

$\longleftrightarrow$ $\mathring{R}\text{-mod}_0$ $\longleftrightarrow$

combinatorial algebra

Definition

A (planar) KLRW diagram is a generic collection of curves in $\mathbb{R} \times [0,1]$ which are of the form $\{(\pi(t), t) \mid t \in [0,1]\}$ for $\pi\colon [0,1] \to \mathbb{R}$.

1. Each strand is labeled from $[1, n]$. If this label is n, we color the strand red, otherwise we color it black. 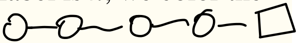
2. Red strands must be vertical at fixed, distinct x-values (for example, $x = 1/n, 2/n, \dots, 1$).
3. We place dots at a finite number of points on black strands, avoiding crossings.

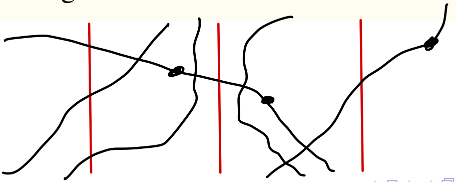

Definition

A cylindrical KLRW diagram is a generic collection of curves in $\mathbb{R}/\mathbb{Z} \times [0,1]$ which are of the form $\{(\pi(t), t) \mid t \in [0,1]\}$ for $\pi : [0,1] \to \mathbb{R}/\mathbb{Z}$.

1. Each strand is labeled from $[1, n]$. If this label is n, we color the strand red, otherwise we color it black.
2. Red strands must be vertical at fixed, distinct x-values (for example, $x = 1/n, 2/n, \ldots, 1$).
3. We place dots at a finite number of points on black strands, avoiding crossings.

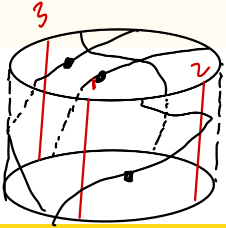

We can compose KLRW diagrams by stacking, if the labels on the bottom of one and top of the other match up to isotopy (never moving red strands).

Definition

The (planar) KLRW algebra R is the formal $\Bbbk$-span of planar KLRW diagrams modulo the local relations below.

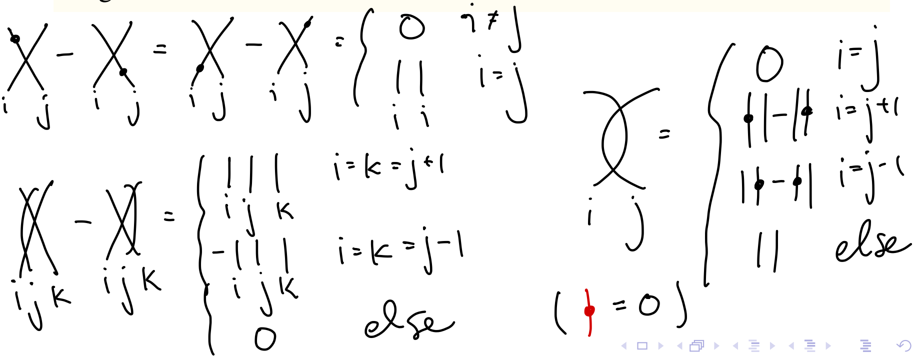

We can compose KLRW diagrams by stacking, if the labels on the bottom of one and top of the other match up to isotopy (never moving red strands).

Definition

The cylindrical KLRW algebra $\mathring{R}$ is the formal $\Bbbk$-span of cylindrical KLRW diagrams modulo the local relations below.

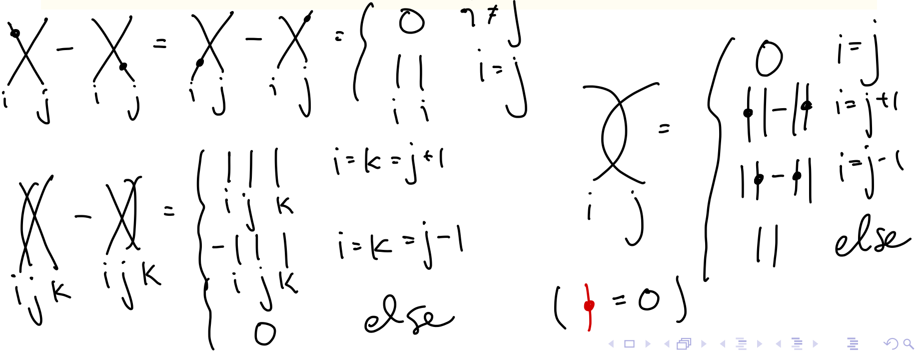

Important role is played by idempotents where all strands are vertical.

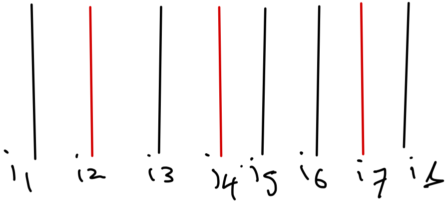

There's one of these for each possible order on strands. Can encode this in a word $\mathbf{i}$ in $\{1, \dots, n-1, n\}$. Denote by $e(\mathbf{i})$.

Definition

The (planar) KLRW category is the category whose objects are words as above, and where $\mathrm{Hom}(\mathbf{i}, \mathbf{j}) = e(\mathbf{j})\mathring{R}e(\mathbf{i})$.

Important role is played by idempotents where all strands are vertical.

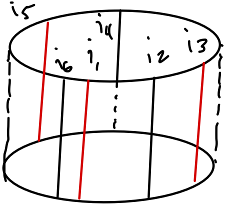

There's one of these for each possible order on strands. Can encode this in a word $\mathbf{i}$ in $\{1, \dots, n-1, n\}$. Denote by $e(\mathbf{i})$.

For $\mathring{R}$, this word is really cyclic, but can always start with red at $x = 0$.

Definition

The cylindrical KLRW category is the category whose objects are words as above, and where $\mathrm{Hom}(\mathbf{i}, \mathbf{j}) = e(\mathbf{j})\mathring{R}e(\mathbf{i})$.

The key to the connection to representation theory is the Gelfand-Tsetlin subalgebra Γ:

$$\Gamma = \langle Z_{HC}(U\mathfrak{gl}_n), Z_{HC}(U\mathfrak{gl}_{n-1}), \dots, Z_{HC}(U\mathfrak{gl}_1)\rangle$$

Theorem (Harish-Chandra)

We have an isomorphism:

$$Z_{HC}(U\mathfrak{gl}_k) = \mathbb{C}[z_{k,1}, \dots, z_{k,k}]^{S_k}$$

where $f(\mathbf{z})$ acts on the Verma module with highest weight $(a_1, \dots, a_k)$ with scalar $f(a_1, a_2 - 1, \dots, a_k - (k-1))$.

The ring Γ is a tensor product of these factors, so it's polynomials invariant under $S_n \times \cdots \times S_1$.

Thus, given any finite dimensional representation of $A = U\mathfrak{gl}_n(\mathbb{k})$, the spectrum of the representation is a subset of

$$\operatorname{Spec}\Gamma = \prod_{k=1}^{n} \mathbb{A}^k_{\mathbb{k}}/S_k$$

If $\operatorname{char}(\mathbb{k}) = 0$, then this spectrum is simple (all multiplicities 1), and determined by **Gelfand-Tsetlin patterns**.

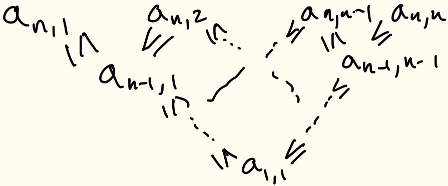

On the other hand, in characteristic $p > 0$, this is very much not the case.

Problem?

The fact that there are finitely many integers mod p, and infinitely many mod 0, makes it much easier to be finite-dimensional in characteristic p.

Better characteristic 0 generalization:

Definition

We call a finitely generated A-module M **Gelfand-Tsetlin** if it is Γ-locally-finite (i.e. $\dim(\Gamma m) < \infty$ for all $m \in M$).

~~If $p > 0$, then equivalent to finite dimensional.~~ If $p = 0$, then many infinite dimensional examples.

This looks like an innocent enough definition, but these modules have proved tricky to work with.

Theorem (Futorny-Grantcharov-Ramirez (2018))

The principal block of the category of Gelfand-Tsetlin modules for $\mathfrak{sl}(3)$ contains 20 simple modules, exactly 1 of which does not lie in category $\mathcal{O}$ for some Borel.

There's no indexing set as obvious as highest weights for category $\mathcal{O}$ so how do we analyze this category and do something like find all simples?

Use the weight functors associated to the kernel $\mathfrak{m}_{\mathbf{a}}$ of the map sending $z_{i,j} \mapsto a_{i,j}$:

$$\mathcal{W}_{\mathbf{a}}(M) = \{m \in M \mid \mathfrak{m}_{\mathbf{a}}^N m = 0 \text{ for } N \gg 0\}$$

The values $a_{n,*}$ describe how $Z_{HC}(U\mathfrak{gl}_n)$ acts, and thus play a special role; $\mathrm{Hom}_{\mathcal{C}}(\mathbf{a},\mathbf{b}) = 0$ unless in same S_n-orbit.

This is related to $A_{\boldsymbol{\alpha}}$, the quotient of A by the corresponding maximal ideal of $Z_{HC}(U\mathfrak{gl}_n)$.

Consider the topological category $\mathcal{C}_{\boldsymbol{\alpha}}$ whose:

- objects are the maximal ideals $\mathfrak{m}_{\mathbf{a}}$ for $a_{i,j} \in \mathbb{Z}/p\mathbb{Z}$ for all i,j, with $a_{n,k} = \alpha_k$.
- morphisms are given by:

$$\mathrm{Hom}_{\mathcal{C}}(\mathbf{a},\mathbf{b}) = \varprojlim A_{\boldsymbol{\alpha}}/(A_{\boldsymbol{\alpha}}\mathfrak{m}_{\mathbf{a}}^N + \mathfrak{m}_{\mathbf{b}}^N A_{\boldsymbol{\alpha}}) = \mathrm{Hom}(\mathcal{W}_{\mathbf{a}},\mathcal{W}_{\mathbf{b}}).$$

Any Gelfand-Tsetlin A_α-module M defines a representation of this category, i.e. a functor to the category $\Bbbk\text{-}\mathsf{vect}$, sending $\mathbf{a} \mapsto \mathcal{W}_\mathbf{a}(M)$.

Theorem (Drozd-Futorny-Ovsienko)

This functor defines an equivalence of categories from integral Gelfand-Tsetlin modules with central character α to (discrete continuous) representations of $\mathcal{C}_\alpha$.

Fix $\alpha \in (\mathbb{Z}/p\mathbb{Z})^n/S_n$ to be a free S_n-orbit, i.e. a regular integral central character.

Theorem

If $p = 0$, then the category $\mathcal{C}_\alpha$ is (Karoubi) equivalent to the planar KLRW category with k strands of label k, completed by adding power series in dots.

If $p > 0$, then the category $\mathcal{C}_\alpha$ is (Karoubi) equivalent to the corresponding cylindrical KLRW category, completed by adding power series in dots.

Composing these theorems:

Theorem

If $p = 0$, then the category of integral GT A_α-modules is equivalent the category of finite dimensional R-modules where dots are nilpotent.

If $p > 0$, then the category of finite dimensional A_α modules is equivalent the category of finite dimensional $\mathring{R}$-modules where dots are nilpotent.

I am sure that essentially the same proof shows that $U_q(\mathfrak{gl}_n(\mathbb{C}))$ for generic q and q a root of unity satisfy same theorem, but haven't checked carefully.

This enabled first classification of simple Gelfand-Tsetlin modules for $p = 0$ (Kamnitzer-Weekes-W.-Yacobi).

What is this equivalence?

It must send the maximal ideal $\mathfrak{m}_{\mathbf{a}}$ to a word $\mathbf{i}(\mathbf{a})$.

Let $\Omega = \{(i,j) \mid 1 \leq j \leq i\}$. If $p = 0$, then there is a unique order on Ω such that:

1. If $a_{i,j} < a_{k,\ell}$, then $(i,j) < (k,\ell)$.
2. if $a_{i,j} = a_{k,\ell}$ and $i < k$, then $(i,j) < (k,\ell)$.

$a_{2,1} = -2 \quad a_{2,2} = 3 \quad a_{1,1} = 5 \implies \underline{i} = (2,2,1)$

Theorem

The word $\mathbf{i}(\mathbf{a})$ is obtained by writing elements of Ω in order, and taking first entries. That is, for a GT module M and KLRW module $\Theta(M)$, we have $\mathcal{W}_{\mathbf{a}}(M) = e(\mathbf{i}(\mathbf{a}))\Theta(M)$.

If $p > 0$, then $\mathbb{Z}/p\mathbb{Z}$ is not ordered, but it is cyclically ordered. A version of the theorem above, but with cyclic orders, holds in this case.

$p=11$

$a_{3,1}=2 \quad a_{3,2}=3 \quad a_{3,3}=7$

$a_{2,1}=5 \quad a_{2,2}=7$

$a_{1,1}=8$

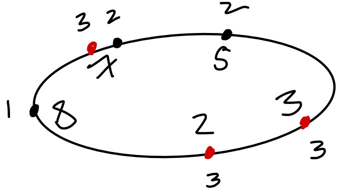

$\underline{i}=(3,3,2,2,3,1)$

This is interesting from a characteristic p perspective:

- It shows that we can match irreps for different p's and different α so that the dimensions of the GT generalized eigenspaces is independent of these choices.
- Thus, dimensions and characters of representations only depend on the number of maximal ideals which correspond to a given $\mathbf{i}$. This is the number of integral points in a polytope, and thus depends quasi-polynomially on α and p.

But this also clarifies the connection to geometry, which previously had only been possible for $p > 0$.

How do we relate this story to geometry?

Recall that we call a vector bundle T on an algebraic variety X a tilting generator if $\mathbb{R}\mathrm{Hom}(T, -)$ induces an equivalence of derived categories $D^b(\mathsf{Coh}(X)) \cong D^b(\mathrm{End}(T)^{\mathrm{op}})\text{-mod}$.

Theorem (W.)

Unless $0 < p < n$, *there is a tilting generator* T *on* $X = T^*\mathrm{Fl}_n$ *such that* $\mathrm{End}(T)^{\mathrm{op}} = \mathring{R}$.

$$D^b(\mathsf{Coh}(X)) \cong D^b(\mathring{R}\text{-mod}).$$

In particular, the ring $\mathring{R}$ *is a non-commutative crepant resolution of singularities of* $\mathcal{N}$ *which is D-equivalent to* X.

How do we relate this story to geometry?

Recall that we call a vector bundle T on an algebraic variety X a tilting generator if $\mathbb{R}\mathrm{Hom}(T,-)$ induces an equivalence of derived categories $D^b(\mathsf{Coh}(X)) \cong D^b(\mathrm{End}(T)^{\mathrm{op}})\text{-mod}$.

Conjecture (W.)

There is a tilting generator T on $X = T^*\mathrm{Fl}_n$ such that $\mathrm{End}(T)^{\mathrm{op}} = \mathring{R}$.

$$D^b(\mathsf{Coh}(X)) \cong D^b(R\text{-mod}).$$

In particular, the ring $\mathring{R}$ is a non-commutative crepant resolution of singularities of $\mathcal{N}$ which is D-equivalent to X.

What is T?

- It's the tilting generator which arises from crystalline (twisted) differential operators.
- It also has an explicit construction by modules over the projective coordinate ring.

This latter description will look very complicated if I write it down. Important point: comes from BFN description of $U\mathfrak{gl}_n(\mathbb{k})$ as a Coulomb branch (in the sense of Braverman-Finkelberg-Nakajima).

TDOs turn into an Azumaya algebra $\mathcal{A}$ of rank $p^{n(n-1)}$ whose sections are A_α which does not split for some silly characteristic p reasons, but it really *wants* to.

In the restriction $\widehat{\mathcal{A}}$ to the formal neighborhood of X_0, $f(z_{i,*}^p - z_{i,*})$ for f symmetric acts nilpotently. This implies that Γ can be (generalized) diagonalized with spectrum $\mathcal{S} = \prod_{k=1}^n \mathbb{F}_p^k/S_k$ and $1 = \sum_{\mathbf{a}\in\mathcal{S}} 1_{\mathbf{a}}$ is the sum of the projections.

Theorem

$\widehat{T} = \widehat{\mathcal{A}} \cdot 1_{\mathbf{0}}$ *is a splitting bundle for the Azumaya algebra* $\widehat{\mathcal{A}} = \widehat{T}^\vee \otimes \widehat{T}$.

It follows from our previous calculations with Γ that $\mathrm{End}(\widehat{T}) = \mathring{R}$.

Versions of all of these results apply to Coulomb branches of all minuscule ADE cases (affine Grassmannian slices) and affine type A quiver gauge theories (also affine type A quiver varieties).

All of these have tilting generators whose endomorphisms are versions of KLRW algebras.

Q: What about other classical Lie algebras? Maybe if you think very hard about previous talk.....

Thanks for listening.

KLRW algebras