Representation Theory in Intermediate Characteristic

Jacob Lurie
Notes by Tony Feng

1 Introduction

We are going to work p-locally, i.e. fix a prime p and work over a "field" in which all other primes are invertible. In algebra, this means that we are working in a characteristic 0 field or a field extension of $\mathbb{F}_p$.

Last time we introduced the notion of a cohomological field. Now the notion of work p-locally becomes a little more complicated. There are Morava K-theories which can be thought of as "intermediate fields" interpolating between characteristic 0 and characteristic p.

Question. What happens to the representation theory of finite groups over intermediate fields?

1.1 (Classical) representation Theory of finite groups

Let G be a finite group and k a field. What does the representation theory of G look like over k?

Characteristic zero. We have *complete reducibility*, which can be described by the equivalent statements:

- Every representation is a direct sum of irreducibles.
- Every exact sequence of representations splits.

In particular, for any V the subspace of invariants V^G is a direct summand of V. Although this is a special case of completely reducibility, it's not hard to show that it also implies complete reducibility.

Proof. We introduce also the *co-invariants* $V_G = V/(gv - v)$, which is the maximal quotient of V with trivial G-action. So we have

$$V^G \hookrightarrow V \twoheadrightarrow V_G.$$

In particular, the composition is a map from V^G to V_G. If we want to split off V^G, it would suffice to show that this composite is an isomorphism.

You can show this by writing down a map in the other direction, involving averaging:

$$v \mapsto \sum_{g \in G} g \cdot v.$$

A priori this is a map $V \to V$, but it is easily shown to factor through V^G. Dually, it factors through V_G:

$$V \twoheadrightarrow V_G \xrightarrow{N_G} V^G \hookrightarrow V.$$

The composition in either direction is multiplication by $|G| \neq 0$, which is where we use characteristic 0. □

We highlight:

Phenomenon: In characteristic 0, the coinvariants and invariants are isomorphic.

Characteristic p. If $|G|$ is coprime to p, then everything is the same. However, the argument falls apart if $|G|$ is divisible by p and complete reducibilty fails.

In fact let's consider an extreme: suppose $|G|$ is a *power* of p. Then complete reducibility fails badly, but we get the consolation prize of *unipotence*.

- Every irreducible representation of G is trivial.
- $V \neq 0 \implies V^G \neq 0$.

Therefore, every representation can be formed by extensions of trivial representations.

So if G is a p-group, we get a dichotomy:

Characteristic 0	Characteristic p
Complete reducibilty	Unipotence
Irreducibles interesting	Irreducibles trivial
Extensions trivial	Extensions interesting

2 Local systems

First I want to recast the notion of a representation. Fix a field k.

Definition 2.1. A *local system* $\mathcal{L}$ of k-vector space on X assigns:

- To each $x \in X$ a k-vector space $\mathcal{L}_x$.
- For each path $p\colon [0,1] \to X$ from $x = p(0)$ to $y = p(1)$ an isomorphism $\mathcal{L}_p\colon \mathcal{L}_x \cong \mathcal{L}_y$ (thought of as "parallel transport along p")

- For every two-simplix $\Delta^2 \to X$

  $$\begin{array}{ccc} & y & \\ {}^{p}\nearrow & & \searrow^{q} \\ x & \xrightarrow{r} & z \end{array}$$

  we require $\mathcal{L}_r = \mathcal{L}_q \circ \mathcal{L}_p$.

  This says two things simultaneously. First, that $\mathcal{L}_p$ depends only on the homotopy class of p. Second, that the isomorphism associated to a composite of paths is obtained by composing the isomorphisms of the individual paths (some kind of functoriality).

If X is connected, then *a local system on X is the same as a representation of* $\pi_1 X$.

2.1 Local systems of $K(n)$-modules

Now I want to go to homotopy theory and talk about local systems of $K(n)$-modules on a space X, i.e. spectra (or cohomology theories) equipped with an action of $K(n)$.

Definition 2.2. A *local system of* $K(n)$*-modules* consists of:

- To each $x \in X$ a $K(n)$-module spectrum $\mathcal{L}_x$.
- To each path $p\colon [0,1] \to X$ from $x = p(0)$ to $y = p(1)$ an isomorphism $\mathcal{L}_p\colon \mathcal{L}_x \cong \mathcal{L}_y$ (thought of as "parallel transport along p")
- For every two-simplex $\Delta^2 \to X$

  $$\begin{array}{ccc} & y & \\ {}^{p}\nearrow & & \searrow^{q} \\ x & \xrightarrow{r} & z \end{array}$$

  a homotopy $\mathcal{L}_r \cong \mathcal{L}_q \circ \mathcal{L}_p$.

  In homotopy theory, it's not natural to demand equality. Instead, we demand a homotopy, and this should be supplied as part of the data of a local system of $K(n)$-modules.
- Analogous data for simplices of all dimensions.

Warning: this does not depend only on $\pi_1 X$. This notion potentially sees the entire homotopy type of X. The topological space X serves as a replacement for the finite group G.

Classically, local systems on X are representations of $\pi_1(X)$. We need some finiteness conditions, which strengthen the notion of finite fundamental group.

Definition 2.3. A space X is π*-finite* if:

- the set $\pi_0(X)$ is finite
- each homotopy group $\pi_n(X, x)$ is finite
- the groups $\pi_n(X, x)$ vanish for $n \gg 0$ (so there are only finitely many non-zero such).

We say that X is *p-finite* if it satisfies all of these and also all the non-zero homotopy groups are p-groups.

Example 2.4. If G is finite, then BG is π-finite.

If G is a finite p-group, then BG is p-finite.

Local systems on BG are the same as "representations of G." This is literally true for local systems of vector spaces, but for local systems of $K(n)$ modules it is a definition.

We're going to talk about local systems on general π (or p) finite groups. There are at least two reasons for this level of generality. First, the general story works out well so we may as well tell it. Second, some features would be obscured if we restricted ourselves to BG.

Representations	Local Systems
Finite group G	Space X
Representations V	Local system $\mathcal{L}$
Invariants V^G	$C^*(X;\mathcal{L}) := \varprojlim_{x\in X} \mathcal{L}_x$ (i.e. global sections)
Coinvariants	$C_*(X;\mathcal{L}) := \varinjlim_{x\in X} \mathcal{L}_x$

This limit here is over all points and paths.

3 Norm Isomorphisms

Theorem 3.1 (Hopkins-Lurie)**.** *Let X be a π-finite space and $\mathcal{L}$ a local system of $K(n)$-modules on X. There is a canonical norm isomorphism*

$$N_X : C_*(X;\mathcal{L}) \xrightarrow{\sim} C^*(X;\mathcal{L}).$$

This should surprise you. This was a feature of characteristic 0, which worked out because you could divide by the order of the group. However, over $K(n)$ of a point, $p = 0$ so you can't "divide by the order of the group." But this is telling you that in some sense $K(n)$ is like working in characteristic 0.

This is pretty abstract, so let's spell out what it tells you in some special cases.

Example 3.2. If $\mathcal{L}$ is the trivial local system, then $K(n)_*X$ is like the homology of X with coefficients in $K(n)$ and $K(n)^*X$ is the cohomology, so the isomorphism

$$K(n)_*X \xrightarrow{\sim} K(n)^*X$$

is like Poincaré duality. In particular, $K(n)^0X$ is *canonically self-dual*.

Example 3.3. $K(1) = K/p$. For BG, this is related to the representation theory of G. Let G be a finite p-group and $\operatorname{Rep}(G)$ its representation ring. Then there is a non-degenerate bilinear form

$$b\colon \operatorname{Rep}(G) \otimes \operatorname{Rep}(G) \to \mathbb{Z}$$

given by

$$(V, W) \mapsto \dim_{\mathbb{C}} \operatorname{Hom}(V, W) = \frac{1}{|G|} \sum_{g \in G} \overline{\chi_V(g)} \chi_W(g). \tag{1}$$

The Atiyah-Segal completion theorem gives an isomorphism

$$K(1)^0 B \cong \operatorname{Rep}(G) \otimes \mathbb{F}_p.$$

The form b gives an identification of $K(1)^0 BG$ with its dual (over $\mathbb{F}_p$).

The Theorem 3.1 is the driving force for all the results I'll discuss.

Sketch of proof. You build the norm map in such a way that it has nice naturality properties. That lets you run a devissage to a few elemental cases.

One is the case where X is an Eilenberg-MacLane space and $\mathcal{L}$ is trivial. In this case, you can think of the norm map as a bilinear form on the homology. You need to show that this is non-degenerate. The way that goes is sort of mimicking the known argument by "lifting to characteristic 0," at which point you can write down a formula. It looks like the rightmost expression in (1).

It's clear that it's non-degenerate rationally but not integrally, and in fact it's not even clearly well-defined integrally. You have to find an interpretation like the middle expression of (1). □

Let X be any space. The construction $\mathcal{L} \mapsto C^*(X; \mathcal{L})$ commutes with inverse lmits because it is an inverse limit. But it doesn't usually commute with direct limits, since inverse limits don't need to play well with direct limits.

Corollary 3.4. *If X is π-finite, then the two constructions are the same so they both commute with both types of limits.*

Definition 3.5. A local system $\mathcal{L}$ of $K(n)$-modules on a space X is *unipotent* if it can be built from constant local systems using direct limits.

This looks different from the usual notion, because you usually build unipotent things with extensions. But local systems are a triangulated category, and extensions are replaced with cones, so this subsumes extensions.

Corollary 3.6. *If X is π-finite then every lcoal system can be written as an extension*

$$\mathcal{L}^{unip} \to \mathcal{L} \to \mathcal{T}$$

where $C^*(X; \mathcal{T}) = 0$.

Sketch. Take $\mathcal{L}^{\mathrm{unip}}$ to be the direct limit of all constant local systems wiht a map to $\mathcal{L}$. You need to show that $\mathcal{T}$ has no global sections, and for that you need to know if global sections commutes with limits. □

4 Phenomena

4.1 Unipotence

Theorem 4.1 (Hopkins-Lurie). *Let X be a p-finite space and assume that $\pi_m X \cong 0$ for $m > n$. Then every local system of $K(n)$-modules on X is unipotent.*

Example 4.2. In particular, if G is a finite p-group then every representation of G (on a $K(n)$-module) is unipotent.

This is saying that for p-finite spaces, working over Morava k-theories is like working in characateristic p.

4.2 Complete reducibility

What if the homotopy groups are not concentrated in degree $\leq n$? Then you get a kind of "complete reducibility. "

Theorem 4.3. *Let X be a π-finite space. Assume that $|\pi_m X|$ is not divisible by p for $m \leq n$. Then*

- *Every unipotent local system on X is constant.*
- *For every local system $\mathcal{L}$ on X, the extension*

$$\mathcal{L}^{unip} \to \mathcal{L} \to \mathcal{K}$$

  splits (breaks off into invariant piece and a piece with no invariatns).

Remark 4.4. If $\pi_m X \cong 0$ for $m \leq n + 1$, then any local system of $K(n)$-modules on X is constant. This says that really the only thing that matters in degree $n + 1$, and everything above is not that relevant.

4.3 Summary

What is this saying? For local systems of $K(n)$-modules on a p-finite space X:

- There is a "unipotent regime" if the homotopy groups are concentrated in degree up to n,
- There is a "complete reducibility regime" if the homotopy groups are concentrated in degree $n + 1$,
- If the homotopy groups are concentrated above degree $n + 1$ then everything is trivial.

So there is a sort of "phase transition." It depends on where the homotopy groups are concentrated. Notice that this is why it was so important not merely to study spaces BG.

The larger n is, the more unipotence (characteristic p) we see, and the "closer" we are to being in characteristic p. The above result is a quantitative articulation of this philosophy.