INTRODUCTION TO GEOMETRIC LANGLANDS THEORY

This is the note for Lin Chen's lecture series on geometric Langlands during the Fall 2023 semester at Tsinghua University. The original materials are available on the course webpage. The notes are taken by Wenhan Dai and the note-taker is responsible for all the mistakes and typos. This note is originally written in 2023 and was significantly revised in 2025.

Contents

Last update on June 30, 2025.

1. Introduction

This overview explains the path from global class field theory to the global, local, and quantum forms of geometric Langlands. The point is that the classical relation between Galois data and automorphic functions has a geometric lift: functions are replaced by sheaves on moduli stacks, and Frobenius eigenvalues are replaced by Hecke eigen-isomorphisms. The modern functorial formulation is developed in [GR24a, AG15].

The notes proceed as follows. Section **??** isolates Bun_G and its basic stack-theoretic and adelic descriptions. Section 2 formulates the Hecke eigenproperty and gives the GL_n minuscule Hecke argument. Section 3 develops the Whittaker model, and Section 4 constructs and analyzes Laumon's sheaf. Sections 5–9, especially the construction in Section 8, prove the main Fourier-transform, clean-extension, and vanishing statements. The categorical form is developed in Sections 10, 11, 12, and 13; the Hitchin and local inputs appear in Sections 14–15; the global proof is organized in Section 16.

1.1. **Geometric class field theory.** Let K be a global field, let $\mathbb{A}_K$ be its ring of adeles, and let $\mathrm{Gal}_{\overline{K}/K}$ be the absolute Galois group. Global class field theory gives an "almost" isomorphism; the word is intentionally imprecise in this overview, and later lectures replace this slogan by precise statements:

$$\theta_K\colon K^\times\backslash\mathbb{A}_K^\times \longrightarrow \mathrm{Gal}^{\mathrm{ab}}_{\overline{K}/K} \simeq \mathrm{Gal}_{K^{\mathrm{ab}}/K},$$

where $K^{\mathrm{ab}} \subset \overline{K}$ is the maximal abelian extension. It is characterized by local Frobenius elements: for every finite place v,

$$\theta_K(\pi_v) = \mathrm{Fr}_v.$$

Here π_v is the class of a uniformizer of K_v, well-defined modulo $\mathcal{O}_v^\times$, and Fr_v is a lift of geometric Frobenius in $\mathrm{Gal}_{\overline{\kappa}_v/\kappa_v}$, well-defined modulo the inertia group I_v.

Theorem 1.1.1 (Unramified class field theory, character form)**.** *Let $\ell \neq \mathrm{char}(K)$. There is an almost bijection*

$$\{\textit{unramified characters } \xi\colon K^\times\backslash\mathbb{A}_K^\times \to \overline{\mathbb{Q}}_\ell^\times\} \longleftrightarrow \{\textit{unramified characters } \rho\colon \mathrm{Gal}_{\overline{K}/K} \to \overline{\mathbb{Q}}_\ell^\times\},$$

characterized by

$$\xi(\pi_v) = \rho(\mathrm{Fr}_v).$$

Unramified means that ξ is trivial on $\mathcal{O}_v^\times$ and ρ is trivial on I_v, for all finite places v.

For number fields, the connected kernel of the reciprocity map is invisible to $\overline{\mathbb{Q}}_\ell^\times$-valued characters because $\overline{\mathbb{Q}}_\ell^\times$ is totally disconnected. For function fields over $\mathbb{F}_q$, the precise formulation uses the Weil group $W_{\overline{K}/K}$.

Now assume K is the function field of a smooth complete curve X over $k = \mathbb{F}_q$. Write $\mathcal{O}_K^\times \coloneqq \prod_v \mathcal{O}_v^\times$. The quotient $K^\times\backslash\mathbb{A}_K^\times/\mathcal{O}_K^\times$ is discrete, and Weil's theorem gives

$$K^\times\backslash\mathbb{A}_K^\times/\mathcal{O}_K^\times \simeq \mathrm{Pic}(X), \qquad \pi_v \longmapsto \mathcal{O}_X(v).$$

Thus unramified automorphic characters are the same as characters of the Picard group:

$$\{\text{unramified } \xi\colon K^\times\backslash\mathbb{A}_K^\times \to \overline{\mathbb{Q}}_\ell^\times\} \simeq \{\xi\colon \mathrm{Pic}(X) \to \overline{\mathbb{Q}}_\ell^\times\}.$$

On the Galois side, $\mathrm{Gal}_{\overline{K}/K} \simeq \pi_1(\mathrm{Spec}\, K, \eta)$ for a geometric point η of $\mathrm{Spec}\, K$. An unramified representation factors through the surjection

$$\pi_1(\mathrm{Spec}\, K, \eta) \longrightarrow \pi_1(X, \eta).$$

Equivalently,

$$\begin{aligned}\{\text{unramified } \rho\colon \mathrm{Gal}_{\overline{K}/K} \to \overline{\mathbb{Q}}_\ell^\times\} &\simeq \{\rho\colon \pi_1(X,\eta) \to \overline{\mathbb{Q}}_\ell^\times\}\\ &\simeq \{\text{rank-one } \ell\text{-adic local systems } \sigma \text{ on } X\}.\end{aligned}$$

Hence class field theory becomes the almost bijection

$$\{\xi\colon \mathrm{Pic}(X) \to \overline{\mathbb{Q}}_\ell^\times\} \longleftrightarrow \{\text{rank-one } \ell\text{-adic local systems } \sigma \text{ on } X\}.$$

The remaining task is to express the equality $\xi(\pi_v) = \rho(\mathrm{Fr}_v)$ geometrically. Let Y be a variety over $k = \mathbb{F}_q$ and let $\mathcal{F}$ be an ℓ-adic sheaf on Y. For $y \in Y(k)$, the stalk $\mathcal{F}_{\bar{y}}$ carries an action of $\mathrm{Gal}_{\bar{k}/k}$; define its trace function by

$$f_{\mathcal{F}}(y) \coloneqq \mathrm{Tr}(\mathrm{Fr}_y, \mathcal{F}_{\bar{y}}).$$

This is Grothendieck's sheaf-function dictionary. If ρ corresponds to a local system σ, then, at a k-point $v \in X(k)$,

$$\rho(\mathrm{Fr}_v) = f_\sigma(v) = \mathrm{Tr}(\mathrm{Fr}_v, \sigma_{\bar{v}}).$$

The same notation is used for closed points, with the corresponding geometric Frobenius.

Corollary 1.1.2 (Geometric form of the unramified character relation)**.** *There is an almost bijection*

$$\{\xi\colon \operatorname{Pic}(X) \to \overline{\mathbb{Q}}_\ell^\times\} \longleftrightarrow \{\textit{rank-one } \ell\textit{-adic local systems } \sigma \textit{ on } X\},$$

characterized by

$$\xi(\mathcal{O}_X(v)) = f_\sigma(v).$$

Replace now the Picard group by the relative Picard scheme. For every affine test scheme S/k,

$$\operatorname{Pic}_{X/k}(S) := \operatorname{Pic}(X \times_k S)/\operatorname{Pic}(S), \qquad \operatorname{Pic}_{X/k}(k) = \operatorname{Pic}(X).$$

One wants a geometric object Aut_σ on $\operatorname{Pic}_{X/k}$ whose trace function is the character attached to σ:

$$\begin{array}{ccc} \{\ell\text{-adic sheaves } \mathcal{F} \text{ on } \operatorname{Pic}_{X/k}\} & \xrightarrow{\mathcal{F}\mapsto f_{\mathcal{F}}} & \{\xi\colon \operatorname{Pic}(X) \to \overline{\mathbb{Q}}_\ell^\times\} \\ & \overset{\sigma\mapsto \mathrm{Aut}_\sigma}{\nwarrow} & \updownarrow \\ & & \{\text{rank-one } \ell\text{-adic local systems } \sigma \text{ on } X\}. \end{array}$$

The correct geometric condition on $\mathcal{F}$ is the Hecke eigenproperty.

Definition 1.1.3 (Hecke eigensheaf for GL_1)**.** Let $\mathcal{F}$ be a sheaf on $\operatorname{Pic}_{X/k}$ and let σ be a rank-one local system on X. We say that $\mathcal{F}$ is a Hecke eigensheaf with eigenvalue σ if it is equipped with an isomorphism

$$\mathrm{add}^* \mathcal{F} \simeq \sigma \boxtimes \mathcal{F},$$

satisfying the usual commutativity condition for two successive modifications. Here

$$\mathrm{add}\colon X \times_k \operatorname{Pic}_{X/k} \longrightarrow \operatorname{Pic}_{X/k}, \qquad (v, L) \longmapsto L(v).$$

Lemma 1.1.4. *Let $\mathcal{F}$ be a Hecke eigensheaf with eigenvalue σ, and assume that the restriction of $\mathcal{F}$ to the unit point $\mathcal{O}_X \in \operatorname{Pic}_{X/k}$ is the constant sheaf, so that the associated function sends the unit to* 1*. Put $\xi = f_{\mathcal{F}}$ on $\operatorname{Pic}(X) = \operatorname{Pic}_{X/k}(k)$. Then ξ is a character and, for every place v,*

$$f_{\mathcal{F}}(\mathcal{O}_X(v)) = f_\sigma(v).$$

Lemma 1.1.4 says that the Hecke eigen-isomorphism is the sheaf-level replacement of the multiplicativity and Frobenius-eigenvalue equations in class field theory.

Theorem 1.1.5 (Deligne)**.** *For every rank-one local system σ on X, there exists an essentially unique Hecke eigensheaf Aut_σ on $\operatorname{Pic}_{X/k}$ with eigenvalue σ; moreover Aut_σ is a local system. Here "essentially unique" means unique up to tensoring by an object of $\operatorname{Shv}(\operatorname{Spec} k)$, which preserves the eigenvalue.*

Theorem 1.1.5 is a purely geometric statement, and Deligne's proof is geometric. It is therefore the unramified global geometric class field theory, or equivalently the unramified global geometric Langlands correspondence for GL_1.

1.2. **Unramified global geometric Langlands correspondence.** We now pass from GL_1 to GL_n and then to split reductive groups G over k. The version for non-constant reductive group schemes over X is also important, but it is not part of these introductory lectures. From now on, X is smooth, complete, and geometrically connected over a perfect field k, and $K = K(X)$ is its function field. Fix a reasonable sheaf theory with algebraically closed coefficient field e. For a nice k-scheme or stack Y, write $\operatorname{Shv}_c(Y)^\heartsuit$ for the abelian category and $\operatorname{Shv}_c(Y)$ for the derived category. The standard choices are:

- the Betti setting: $k = \mathbb{C}$, with perverse sheaves for the complex topology;
- the ℓ-adic setting: k finite, with perverse ℓ-adic Weil sheaves;
- the de Rham setting: k algebraically closed of characteristic 0, with coherent right D-modules.

The Galois-side generalization is formal:

GL_n	rank n local systems σ on X,
G	G-local systems σ on X.

For automorphic sheaves attached to G, the actual eigenvalues are local systems for the Langlands dual group $\check{G}$ over the coefficient field e.

On the automorphic side, the abelian chart for $\mathrm{GL}_1 = \mathbb{G}_m$ must be replaced by its G-analogue, i.e.

GL_1	$K^\times\backslash\mathbb{A}_K^\times/\mathcal{O}_K^\times$	$\operatorname{Pic}(X)$	$\operatorname{Pic}_{X/k}$
G	$G(K)\backslash G(\mathbb{A}_K)/G(\mathcal{O}_K)$	?	?

Theorem 1.2.1 (Weil uniformization)**.** *Under the standard hypotheses on G and X—made precise later, with connected reductive groups over curves as the intended case—the double coset set $G(K)\backslash G(\mathbb{A}_K)/G(\mathcal{O}_K)$ is naturally identified with the set of isomorphism classes of principal G-bundles, equivalently G-torsors, on X.*

Unlike $\mathrm{Pic}_{X/k}$, the moduli object for G-bundles is usually not a scheme. Stabilizer groups vary with the bundle, so the correct object is an algebraic stack: it remembers automorphism groups of its points, extra structure that is invisible in the bare double coset set and essential for general G.

Proposition 1.2.2 (Definition of Bun_G)**.** *There is a smooth Artin stack Bun_G over k classifying families of G-torsors on X. On isomorphism classes of k-points,*

$$\mathsf{Bun}_G(k) \simeq G(K)\backslash G(\mathbb{A}_K)/G(\mathcal{O}_K).$$

A more precise notation is $\mathsf{Bun}_{G,X/k}$.

The Hecke eigenproperty also has to be generalized. For $G = \mathrm{GL}_1$, the map

$$\mathrm{add}\colon X \times_k \mathrm{Pic}_{X/k} \longrightarrow \mathrm{Pic}_{X/k}, \qquad (x, L) \longmapsto L(x)$$

modifies a line bundle at x in a canonical way. For a general G-torsor, there is no single canonical modification at x; there are many. Thus the map add is replaced by the global Hecke correspondence

$$\begin{array}{ccc} & \mathsf{Hecke}_{G,X} & \\ {\scriptstyle \overleftarrow{h}}\swarrow & & \searrow{\scriptstyle (\mathrm{supp}, \overrightarrow{h})} \\ \mathsf{Bun}_G & & X \times_k \mathsf{Bun}_G. \end{array}$$

It records a point of modification and two G-bundles identified away from that point.

Assuming the general definition of Hecke eigensheaves, the unramified global geometric Langlands correspondence can be stated as follows.

Conjecture 1.2.3 (Unramified global geometric Langlands)**.** For every geometrically irreducible $\check{G}$-local system σ on X, there exists an essentially unique Hecke eigensheaf Aut_σ on Bun_G with eigenvalue σ.

Here geometrically irreducible means that, after base change to $\bar{k}$, the associated $\check{G}$-local system has no flat reduction to a proper parabolic subgroup of $\check{G}$; for $G = \mathrm{GL}_n$ this is the usual irreducibility of the rank-n local system. If σ is reducible, one should expect Hecke eigenobjects in $\mathrm{Shv}_{\mathrm{c}}(\mathsf{Bun}_G)$ rather than eigensheaves in the heart, and essential uniqueness is no longer true.

For $G = \mathrm{GL}_n$, the existence theorem is known: Deligne proved $n = 1$, Drinfeld proved $n = 2$, and Frenkel–Gaitsgory–Vilonen proved the general case [FGKV98]. In the ℓ-adic setting, taking Frobenius traces of Aut_σ produces an automorphic function. Thus the conjecture is a geometric Galois-to-automorphic construction for function fields, and the function obtained from Aut_σ is cuspidal.

1.3. **Unramified global geometric Langlands equivalence.** Conjecture 1.2.3 assigns one automorphic sheaf to one irreducible spectral point. A categorical form asks for a functor from the whole spectral moduli object. Let $\mathsf{LocSys}^{\mathrm{irred}}_{\check{G}}$ denote the algebro-geometric object classifying families of geometrically irreducible $\check{G}$-local systems on X. One expects a functor

$$\mathrm{Coh}(\mathsf{LocSys}^{\mathrm{irred}}_{\check{G}})^{\heartsuit} \longrightarrow \mathrm{Shv}_{\mathrm{c}}(\mathsf{Bun}_G)^{\heartsuit}, \qquad \delta_\sigma \longmapsto \mathrm{Aut}_\sigma,$$

and, philosophically, an equivalence with the cuspidal part of $\mathrm{Shv}_{\mathrm{c}}(\mathsf{Bun}_G)^{\heartsuit}$.

In the de Rham setting, $\mathsf{LocSys}^{\mathrm{irred}}_{\check{G}}$ is a smooth Artin stack, and the heart-level expectation becomes the following.

Conjecture 1.3.1 (Cuspidal de Rham form)**.** There is a canonical t-exact equivalence

$$\mathbb{L}_G\colon \mathrm{DMod}_{\mathrm{coh}}(\mathsf{Bun}_G)_{\mathrm{cusp}} \simeq \mathrm{Coh}(\mathsf{LocSys}^{\mathrm{irred}}_{\check{G}}),$$

such that $\mathbb{L}_G^{-1}$ sends the skyscraper sheaf at an irreducible $\check{G}$-local system σ to the Hecke eigensheaf Aut_σ with eigenvalue σ.

If reducible local systems are included, one gives up t-exactness and works at the level of categories.

Theorem 1.3.2 (Naive slogan: full de Rham form)**.** *At the level of a first approximation, one writes a spectral-to-automorphic equivalence*

$$\mathbb{L}_G\colon \mathrm{DMod}_{\mathrm{coh}}(\mathsf{Bun}_G) \simeq \mathrm{Coh}(\mathsf{LocSys}_{\check{G}}),$$

with skyscraper sheaves at $\check{G}$-local systems mapping to Hecke eigenobjects. This display is only a guiding slogan: the precise non-cuspidal de Rham formulation replaces the coherent spectral side by the renormalized category with nilpotent singular support, such as $\mathrm{IndCoh}_{\mathrm{Nilp}}(\mathsf{LocSys}_{\check{G}})$.

This is one categorical level higher than the classical Langlands correspondence: the classical objects are functions and representations, while the geometric objects are sheaves and DG categories. More precise modern formulations replace the simplified coherent spectral side by the appropriate renormalized or nilpotent-singular-support category, such as $\mathrm{IndCoh}_{\mathrm{Nilp}}(\mathsf{LocSys}_{\check{G}})$ [AG15, GR24a], but the slogan remains: automorphic categories on Bun_G are controlled by spectral categories on $\check{G}$-local systems. The ramified global story is not treated in these lectures.

1.4. **Local geometric Langlands equivalence.** We now focus on the de Rham context. Classical local Langlands classifies representations of $G(K_v)$ for a local field K_v; these form a 1-category. Local geometric Langlands is a categorification: it seeks to classify categorical representations of $G(K_v)$, which form a 2-category.

The geometric incarnation of $G(K_v)$ is the loop group LG. Over $\mathbb{C}$, its points are $G(K_v)$-valued points. Define LG-Mod to be the $(\infty, 2)$-category of DG categories acted on by the monoidal DG category $\mathrm{DMod}(LG)$, with convolution monoidal structure. If Y is a reasonable algebro-geometric object with an LG-action, then $\mathrm{DMod}(Y)$ is an object of LG-Mod.

The basic examples are Gr_G, $\mathscr{F}\ell_G$, and $\mathsf{Bun}_G^{\infty;v}$, namely the affine Grassmannian, the affine flag ind-variety, and the moduli ind-stack of G-torsors on X of infinite level equipped with a trivialization on the formal neighborhood D_v of v. On complex points, this last object has double-coset description

$$\prod_{w\neq v} G(\mathcal{O}_w)\backslash G(\mathbb{A}_K)/G(K).$$

The category $\mathrm{DMod}(LG)$ is a bimodule over itself. Taking invariants for the right action produces DG categories that still have the left LG-action. In this way $\mathrm{DMod}(\mathsf{Gr}_G)$ and $\mathrm{DMod}(\mathscr{F}\ell_G)$ are obtained by taking, respectively, L^+G-invariants and I-invariants for the right action. Here L^+G is the arc group, the geometric avatar of $G(\mathcal{O}_v)$, and I is the Iwahori subgroup, the avatar of $G(\mathcal{O}_v)\times_{G(\kappa_v)} B(\kappa_v)$.

Another important object is the Whittaker model

$$\mathrm{Whit}(LG) \in LG\text{-Mod},$$

obtained by taking right LN-invariants against a generic character χ, where N is the unipotent radical of a Borel subgroup B. A further key object is $\mathrm{KM}_{\mathfrak{g}}$, the DG category of representations of the affine Lie algebra $\widehat{\mathfrak{g}}$ at the critical level.

On the Galois side, consider the moduli problem $\mathsf{LocSys}_{\check{G}}(D^\circ)$ of $\check{G}$-local systems on the punctured disk $D^\circ \coloneqq \operatorname{Spec} K_v$. Define

$$\mathrm{QCohCat}(\mathsf{LocSys}_{\check{G}}(D^\circ))$$

to be the $(\infty, 2)$-category of module categories for the monoidal category $\mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}(D^\circ))$, with tensor product monoidal structure. If $Z \to \mathsf{LocSys}_{\check{G}}(D^\circ)$ is a reasonable algebro-geometric object, then $\mathrm{QCoh}(Z)$ is such a module category.

The basic spectral examples are

$$\mathsf{LocSys}_{\check{G}}(D), \qquad \mathsf{LocSys}_{\check{G}}^{\mathrm{nilp}}(D^\circ), \qquad \mathsf{LocSys}_{\check{G}}(X^\circ),$$

which classify $\check{G}$-local systems on the disk $D \coloneqq \operatorname{Spec} \mathcal{O}_v$, on the punctured disk with nilpotent singularity, and on the punctured curve $X^\circ \coloneqq X \setminus \{v\}$, respectively. For a reductive group H, one has $\mathsf{LocSys}_H(D) \simeq \mathrm{pt}/H$ and $\mathsf{LocSys}_H^{\mathrm{nilp}}(D^\circ) \simeq \mathcal{N}/H$, where $\mathcal{N} \subset \mathfrak{h}$ is the nilpotent cone. One also considers $\mathsf{LocSys}_{\check{G}}(D^\circ)$ itself and the space of opers

$$\mathsf{Op}_{\check{G}}(D^\circ).$$

The local conjecture matches the preceding automorphic and spectral examples; this is the standard schematic form of local geometric Langlands [FG05, FG07].

Conjecture 1.4.1 (Local geometric Langlands, schematic form)**.** There is an almost equivalence

$$LG\text{-Mod} \simeq \mathrm{QCohCat}(\mathsf{LocSys}_{\check{G}}(D^\circ)),$$

where "almost" hides the technical mismatch between QCoh and coherent or renormalized variants: a precise formulation must choose the correct coherent parts on both sides. The equivalence has the following assignments:

$$\begin{aligned} \mathrm{DMod}(\mathsf{Gr}_G) &\longmapsto \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}(D)), \\ \mathrm{DMod}(\mathscr{F}\ell_G) &\longmapsto \mathrm{QCoh}(\mathsf{LocSys}^{\mathrm{nilp}}_{\check{G}}(D^\circ)), \\ \mathrm{DMod}(\mathsf{Bun}^{\infty;v}_G) &\longmapsto \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}(X^\circ)), \\ \mathrm{Whit}(LG) &\longmapsto \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}(D^\circ)), \\ \mathrm{KM}_{\mathfrak{g}} &\longmapsto \mathrm{QCoh}(\mathsf{Op}_{\check{G}}(D^\circ)). \end{aligned}$$

Taking Hom-categories between the listed automorphic objects, and comparing them with the Hom-categories between their spectral images, recovers many familiar equivalences. In these comparisons one replaces DMod by coherent D-modules and QCoh by the corresponding coherent spectral category.

(1) Due to Bezrukavnikov–Finkelberg, the derived geometric Satake is

$$\mathrm{End}(\mathrm{DMod}(\mathsf{Gr}_G)) \simeq \mathrm{End}(\mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}(D))).$$

(2) The Arkhipov–Bezrukavnikov–Ginzburg theorem is

$$\mathrm{Hom}(\mathrm{DMod}(\mathsf{Gr}_G), \mathrm{DMod}(\mathscr{F}\ell_G)) \simeq \mathrm{Hom}(\mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}(D)), \mathrm{QCoh}(\mathsf{LocSys}^{\mathrm{nilp}}_{\check{G}}(D^\circ))).$$

(3) By Bezrukavnikov's equivalence between two geometric realizations of the affine Hecke algebra,

$$\mathrm{End}(\mathrm{DMod}(\mathscr{F}\ell_G)) \simeq \mathrm{End}(\mathrm{QCoh}(\mathsf{LocSys}^{\mathrm{nilp}}_{\check{G}}(D^\circ))).$$

(4) We have

$$\begin{aligned} &\mathrm{Hom}(\mathrm{DMod}(\mathsf{Gr}_G), \mathrm{DMod}(\mathsf{Bun}^{\infty;v}_G)) \\ &\quad \simeq \mathrm{Hom}(\mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}(D)), \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}(X^\circ))). \end{aligned}$$

This recovers the naive global form in Theorem 1.3.2; its precise version uses the nilpotent-singular-support spectral category. Respectively replacing $\mathrm{DMod}(\mathsf{Gr}_G)$ and $\mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}(D))$ by $\mathrm{DMod}(\mathscr{F}\ell_G)$ and $\mathrm{QCoh}(\mathsf{LocSys}^{\mathrm{nilp}}_{\check{G}}(D^\circ))$ gives the corresponding tamely ramified form.

(5) Due to Frenkel–Gaitsgory–Vilonen [FGKV98], the geometric Casselman–Shalika formula is

$$\mathrm{Hom}(\mathrm{DMod}(\mathsf{Gr}_G), \mathrm{Whit}(LG)) \simeq \mathrm{Hom}(\mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}(D)), \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}(D^\circ)))$$

(6) By a theorem of Arkhipov–Bezrukavnikov,

$$\mathrm{Hom}(\mathrm{DMod}(\mathscr{F}\ell_G), \mathrm{Whit}(LG)) \simeq \mathrm{Hom}(\mathrm{QCoh}(\mathsf{LocSys}^{\mathrm{nilp}}_{\check{G}}(D^\circ))\, \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}(D^\circ))).$$

(7) Known as the Feigin–Frenkel theorem [Fre02],

$$\mathrm{Hom}(\mathrm{DMod}(\mathsf{Gr}_G), \mathrm{KM}_{\mathfrak{g}}) \simeq \mathrm{Hom}(\mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}(D)), \mathrm{QCoh}(\mathsf{Op}_{\check{G}}(D^\circ))).$$

(8) By a theorem of Raskin,

$$\mathrm{Hom}(\mathrm{Whit}(LG), \mathrm{KM}_{\mathfrak{g}}) \simeq \mathrm{Hom}(\mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}(D^\circ)), \mathrm{QCoh}(\mathsf{Op}_{\check{G}}(D^\circ))).$$

1.5. **Quantum local geometric Langlands equivalence.** Twisted D-modules deform the local conjecture in Conjecture 1.4.1. The expected quantum local equivalence is symmetric in G and its Langlands dual.

Conjecture 1.5.1 (Quantum local geometric Langlands)**.** There is an almost equivalence

$$LG\text{-Mod}_\kappa \simeq L\check{G}\text{-Mod}_{\check{\kappa}},$$

with assignments

$$\begin{aligned} \mathrm{DMod}_\kappa(\mathsf{Gr}_G) &\longmapsto \mathrm{DMod}_{\check{\kappa}}(\mathsf{Gr}_{\check{G}}), \\ \mathrm{DMod}_\kappa(\mathscr{F}\ell_G) &\longmapsto \mathrm{DMod}_{\check{\kappa}}(\mathscr{F}\ell_{\check{G}}), \\ \mathrm{DMod}_\kappa(\mathsf{Bun}^{\infty;v}_G) &\longmapsto \mathrm{DMod}_{\check{\kappa}}(\mathsf{Bun}^{\infty;v}_{\check{G}}), \\ \mathrm{Whit}_\kappa(LG) &\longmapsto \mathrm{KM}_{\check{\mathfrak{g}},\check{\kappa}}, \\ \mathrm{KM}_{\mathfrak{g},\kappa} &\longmapsto \mathrm{Whit}_{\check{\kappa}}(L\check{G}). \end{aligned}$$

The symmetry of this table is the quantum version of Langlands duality.

2. Hecke eigenproperty

The automorphic side of geometric Langlands is organized by Bun_G, the moduli stack of G-bundles on X. With Bun_G in place, the Hecke eigenproperty is formulated through correspondences that modify a G-bundle at a point of the curve.

In the following, we first recall the prestack and stack language used to define Bun_G, and record its basic finiteness properties, and relate its k-points to the usual adelic double quotient.

2.1. Prestacks and stacks.

Definition 2.1.1. A *prestack* over k is a pseudo-functor

$$y\colon \mathrm{Aff}_k^{\mathrm{op}} \longrightarrow \mathrm{Grpd}.$$

Thus every affine k-scheme S gives a groupoid $y(S)$, and every map $f\colon S_1 \to S_2$ gives a pullback functor $f^*\colon y(S_2) \to y(S_1)$, with the usual coherent compatibilities. If the target is Set, this is the usual notion of a presheaf.

Let τ be a Grothendieck topology on Aff_k, for instance the Zariski, étale, smooth, fppf, or fpqc topology.

Definition 2.1.2. A prestack y is a *τ-stack* if it satisfies descent for τ-covers. Equivalently, for every τ-cover $S^0 \to S$ with Cech nerve $S^\bullet$, the natural functor

$$y(S) \longrightarrow \lim y(S^\bullet)$$

is an equivalence of groupoids.

Remark 2.1.3 (Classifying stacks)**.** For a group scheme G/k, the τ-classifying stack $\mathbb{B}_\tau G$ is the stackification of the quotient prestack pt/G. Concretely,

$$\mathbb{B}_\tau G(S) = \{\tau\text{-locally trivial } G\text{-torsors on } S\}.$$

It may be written as the stack-theoretic colimit of the bar construction

$$\mathbb{B}_\tau G \simeq \mathrm{colim}(\mathrm{pt} \leftarrow G \leftleftarrows G \times G \cdots)$$

in Stk_k. If G is smooth over k, then torsors are locally trivial already in the étale topology, so

$$\mathbb{B}_{\mathrm{et}} G = \mathbb{B}_{\mathrm{sm}} G = \mathbb{B}_{\mathrm{fppf}} G = \mathbb{B}_{\mathrm{fpqc}} G.$$

2.2. The moduli stack of G-bundles.

Let G/k be a smooth affine algebraic group (in the geometric Langlands applications below, a connected reductive group), and let $p\colon X \to S$ be projective, flat, and of finite presentation. Define

$$\mathsf{Bun}_{G,X/S} \coloneqq \mathrm{Maps}_S(X, \mathbb{B}_{\mathrm{et}} G).$$

For any $S' \to S$,

$$\mathsf{Bun}_{G,X/S}(S') = \mathrm{Hom}(X \times_S S', \mathbb{B}_{\mathrm{et}} G) = \{\text{étale } G\text{-torsors on } X \times_S S'\}.$$

When $S = \mathrm{Spec}\, k$ we write simply Bun_G.

Theorem 2.2.1. *The stack $\mathsf{Bun}_{G,X/S}$ is an Artin stack, locally of finite presentation over S, with affine diagonal. More explicitly,*

$$\Delta_{\mathsf{Bun}_{G,X/S}}\colon \mathsf{Bun}_{G,X/S} \longrightarrow \mathsf{Bun}_{G,X/S} \times_S \mathsf{Bun}_{G,X/S}$$

is affine, and $\mathsf{Bun}_{G,X/S}$ admits an open exhaustion $\mathsf{Bun}_{G,X/S} = \bigcup_r U_r$ such that each $U_r \to \mathsf{Bun}_{G,X/S}$ is an open embedding and each U_r has a smooth atlas $V_r \to U_r$ with $V_r \to S$ of finite presentation.

Remark 2.2.2. If G/k is smooth and $X \to S$ is a relative curve, then $\mathsf{Bun}_{G,X/S} \to S$ is smooth. In the above exhaustion one can choose the atlases $V_r \to S$ smooth.

Remark 2.2.3 ($G = \mathrm{GL}_n$)**.** Fix an ample line bundle $\mathcal{O}(1)$ on X. For $r \gg 0$, let $U_r \subset \mathsf{Bun}_n$ be the open substack of rank n vector bundles $\mathcal{F}$ such that

$$R^i p_*(\mathcal{F}(r)) = 0 \quad (i > 0), \qquad p^* p_*(\mathcal{F}(r)) \twoheadrightarrow \mathcal{F}(r).$$

The open substacks U_r cover Bun_n. Nevertheless, Bun_G is usually not quasi-compact; for $G = \mathrm{GL}_n$ it is not connected, and even a fixed component Bun_G^λ need not be quasi-compact.

The following adelic description is the function-theoretic shadow of Bun_G. Let $K = k(X)$, let K_x be the completed local field at $x \in X$, set $\mathbb{A}_X = \prod'_x K_x$ and $\mathcal{O}_X = \prod_x \mathcal{O}_x$.

Theorem 2.2.4 (Adelic uniformization on k-points)**.** *Assume every G-torsor on $\operatorname{Spec} K$, on each $\operatorname{Spec} K_x$, and on $\operatorname{Spec} k'$ for each finite extension k'/k is trivial. Then the natural gluing construction identifies isomorphism classes of k-points of Bun_G with the double quotient*

$$G(\mathcal{O}_X)\backslash G(\mathbb{A}_X)/G(K),$$

using the right-torsor convention. With left torsors the same statement is written in the standard opposite order $G(K)\backslash G(\mathbb{A}_X)/G(\mathcal{O}_X)$.

Remark 2.2.5. For finite k, the triviality hypothesis holds in particular for simply connected semisimple groups, such as SL_2, and for GL_n. For $k = \bar{k}$ of characteristic 0 it holds for connected linear algebraic groups G; for $k = \bar{k}$ of characteristic p it holds for connected reductive G.

The rest of this section introduces the global and local Hecke stacks, recalls the Satake sheaves defining Hecke functors, and then specializes to GL_n to obtain the minuscule eigenproperties.

2.3. **The global Hecke stack.** Assume from now on that X is a smooth curve over k. The global Hecke stack records one modification of a G-bundle at a moving point of X.

Definition 2.3.1. The stack $\mathsf{Hecke}_{G,X}$ is defined by

$$\mathsf{Hecke}_{G,X}(S) = \{(x, \mathcal{P}, \mathcal{P}', \beta) \mid S \xrightarrow{x} X,\ \mathcal{P}, \mathcal{P}' \in \mathsf{Bun}_G(S),\ \beta\colon \mathcal{P}|_{X\times S-\Gamma_x} \xrightarrow{\sim} \mathcal{P}'|_{X\times S-\Gamma_x}\},$$

where $\Gamma_x \subset X \times S$ is the graph of x. It has the correspondence diagram

$$\mathsf{Bun}_G \xleftarrow{\overleftarrow{h}} \mathsf{Hecke}_{G,X} \xrightarrow{(\mathrm{supp}, \overrightarrow{h})} X \times \mathsf{Bun}_G.$$

Here $\overleftarrow{h}$ remembers $\mathcal{P}$, $\overrightarrow{h}$ remembers $\mathcal{P}'$, and supp remembers the point x.

2.4. **Affine Grassmannians and the local model.** The Beilinson–Drinfeld affine Grassmannian over X is the local model for the Hecke stack:

$$\mathsf{Gr}_{G,X}(S) = \{(x, \mathcal{P}, \beta) \mid x\colon S \to X,\ \mathcal{P} \text{ a } G\text{-torsor on } X \times S,\ \beta\colon \mathcal{P}|_{X\times S-\Gamma_x} \xrightarrow{\sim} \mathcal{P}^0|_{X\times S-\Gamma_x}\},$$

where $\mathcal{P}^0$ is the trivial G-torsor.

Theorem 2.4.1. *$\mathsf{Gr}_{G,X}$ is an ind-scheme ind-locally of finite type. If G is reductive, then $\mathsf{Gr}_{G,X}$ is ind-projective over X.*

Fix $x\colon S \to X$. Let $D_x = (\widehat{X \times S})_{\Gamma_x}$ be the formal disc along the graph and let $D_x^\circ = D_x \setminus \Gamma_x$ be the punctured formal disc. Define the positive and full loop prestacks by

$$L^+G_X(S) = \{(x, \gamma) \mid \gamma\colon D_x \to G\}, \qquad LG_X(S) = \{(x, \gamma) \mid \gamma\colon D_x^\circ \to G\}.$$

If $S = \operatorname{Spec} k$ and x has a parameter t, then $D_x = \operatorname{Spec} k[\![t]\!]$ and $D_x^\circ = \operatorname{Spec} k(\!(t)\!)$. The classifying stacks $\mathbb{B}L^+G_X$ and $\mathbb{B}LG_X$ classify G-torsors on D_x and D_x°, respectively. The affine Grassmannian is the quotient

$$\mathsf{Gr}_{G,X} \simeq LG_X/L^+G_X.$$

To pass from this local model to global modifications, introduce the infinite-level bundle

$$\mathsf{Bun}_G^\infty(S) = \{(x, \mathcal{P}, \alpha) \mid x\colon S \to X,\ \mathcal{P} \in \mathsf{Bun}_G(S),\ \alpha\colon \mathcal{P}|_{D_x} \xrightarrow{\sim} \mathcal{P}^0|_{D_x}\}.$$

It is an L^+G_X-torsor over $X \times \mathsf{Bun}_G$. Hence, morally and in the precise quotient sense,

$$\mathsf{Hecke}_{G,X} \simeq \mathsf{Bun}_G^\infty \times^{L^+G_X} \mathsf{Gr}_{G,X} \simeq \mathsf{Gr}_{G,X} \mathbin{\widetilde{\times}} \mathsf{Bun}_G.$$

Consequently $\mathsf{Hecke}_{G,X}$ is an ind-Artin stack, ind-locally of finite type.

2.5. **Local Hecke stack and geometric Satake.** For a fixed point $x \in X$, write $L^+G_x = G(\mathcal{O}_x)$ and $LG_x = G(K_x)$. The local Hecke stack is

$$\mathsf{Hecke}^{\mathrm{loc}}_{G,x} := [L^+G_x \backslash LG_x / L^+G_x] \simeq [L^+G_x \backslash \mathsf{Gr}_{G,x}].$$

There is a relative-position map from the global Hecke stack to the local one,

$$m_x \colon \mathsf{Hecke}_{G,x} \longrightarrow \mathsf{Hecke}^{\mathrm{loc}}_{G,x},$$

obtained by completing a modification at x.

Geometric Satake identifies the spherical sheaf category on the local Hecke stack with representations of the Langlands dual group:

$$\mathrm{Perv}_{L^+G_x}(\mathsf{Gr}_{G,x}) \simeq \mathrm{Perv}(\mathsf{Hecke}^{\mathrm{loc}}_{G,x}) \simeq \mathrm{Rep}(\check{G}).$$

For $V \in \mathrm{Rep}(\check{G})$, denote by Sat_V the corresponding perverse sheaf on $\mathsf{Hecke}^{\mathrm{loc}}_{G,x}$; in the global family over X we use the same notation for the corresponding relative Satake sheaf.

2.6. **Hecke functors and eigensheaves.** For $V \in \mathrm{Rep}(\check{G})$, the global Hecke functor is

$$H_V \colon \mathcal{D}(\mathsf{Bun}_G) \longrightarrow \mathcal{D}(X \times \mathsf{Bun}_G), \qquad H_V(\mathcal{F}) := (\mathrm{supp}, \overrightarrow{h})_!(m^* \mathrm{Sat}_V \otimes \overleftarrow{h}^* \mathcal{F}).$$

At a fixed point $x \in X$ this restricts to

$$H_{V,x} \colon \mathcal{D}(\mathsf{Bun}_G) \longrightarrow \mathcal{D}(\mathsf{Bun}_G), \qquad H_{V,x}(\mathcal{F}) := \overrightarrow{h}_{x,!}(m_x^* \mathrm{Sat}_V \otimes \overleftarrow{h}_x^* \mathcal{F}).$$

The perverse normalization of Sat_V includes the usual shifts and twists.

Let $E_{\check{G}}$ be a $\check{G}$-local system on X. For $V \in \mathrm{Rep}(\check{G})$, let V_E be the associated local system on X.

Definition 2.6.1. An object $\mathcal{F} \in \mathcal{D}(\mathsf{Bun}_G)$ is a *Hecke eigensheaf with eigenvalue* $E_{\check{G}}$ if, for every $V \in \mathrm{Rep}(\check{G})$, there are functorial isomorphisms

$$H_V(\mathcal{F}) \simeq V_E \boxtimes \mathcal{F}$$

in $\mathcal{D}(X \times \mathsf{Bun}_G)$, compatible with tensor products in $\mathrm{Rep}(\check{G})$. In particular, fiberwise at $x \in X$,

$$H_{V,x}(\mathcal{F}) \simeq (V_E)_x \otimes \mathcal{F}.$$

This is the sheaf-theoretic replacement of the classical condition that an automorphic function be an eigenvector for all spherical Hecke operators.

2.7. **The minuscule Hecke operators for** GL_n**.** For $G = \mathrm{GL}_n$ we have $\check{G} = \mathrm{GL}_n$. Let Std be the standard representation of $\check{G}$. The exterior powers $\wedge^i \mathrm{Std}$, $0 \leqslant i \leqslant n$, correspond under geometric Satake to the minuscule Schubert orbits $\mathsf{Gr}^i_G \subset \mathsf{Gr}_G$. These are the closed spherical orbits relevant for the standard Hecke operators. These give the $n+1$ standard operators indexed by $0 \leqslant i \leqslant n$. If one keeps central degree shifts for GL_n, they are translated by the central coweight; modulo this central direction ω_0 and ω_n are identified, so there are n classes.

Let Hecke^i_n be the corresponding global Hecke correspondence, and let $\mathsf{Gr}^i_{n,X}$ denote the corresponding relative Schubert stratum in $\mathsf{Gr}_{\mathrm{GL}_n,X}$. An S-point is a modification of vector bundles at $x \colon S \to X$ whose local relative position is ω_i; locally this is equivalent to a quotient of length i supported on Γ_x. The local model is the ordinary Grassmannian $\mathrm{Gr}(i,n)$, and

$$\mathsf{Hecke}^i_n \simeq \mathsf{Gr}^i_{n,X} \,\widetilde{\times}\, \mathsf{Bun}_n.$$

The corresponding Hecke functor is

$$H^i_n \colon \mathcal{D}(\mathsf{Bun}_n) \longrightarrow \mathcal{D}(X \times \mathsf{Bun}_n), \qquad H^i_n(\mathcal{F}) = (\mathrm{supp}, \overrightarrow{h}^i)_! \, \overleftarrow{h}^{i,*} \mathcal{F}[\mathrm{shift}],$$

where the shift is the Satake/perverse normalization; for the minuscule orbit one may take the shift by $\dim \mathrm{Gr}(i,n) = i(n-i)$, with the corresponding Tate twist in the ℓ-adic normalization.

Thus a Hecke eigensheaf $\mathcal{F}$ with rank n eigenvalue E satisfies

$$H^i_n(\mathcal{F}) \simeq \wedge^i E \boxtimes \mathcal{F}, \qquad 0 \leqslant i \leqslant n. \tag{2.1}$$

The remaining subsections prove this formula for the automorphic sheaf constructed from E.

We now give the detailed GL_n argument. The point is that the first Hecke eigenproperty, plus the usual symmetry under exchanging distinct points, implies all minuscule eigenproperties in (2.1). This is the sheaf-theoretic form of the elementary symmetric functions of Frobenius eigenvalues.

2.8. **The minuscule Hecke correspondences.** Let $\mathsf{Bun}_n := \mathsf{Bun}_{\mathrm{GL}_n}$. For $0 \leqslant i \leqslant n$, set $\omega_i = (1^i, 0^{n-i})$ and thus $t^{\omega_i} = \mathrm{diag}(tI_i, I_{n-i})$, and let

$$\mathsf{Gr}_n^i := L^+\mathrm{GL}_n \cdot t^{\omega_i} \cdot L^+\mathrm{GL}_n / L^+\mathrm{GL}_n.$$

Let $\mathsf{Gr}_{n,X}^i$ be the corresponding relative Schubert stratum over X. The corresponding global Hecke stack is

$$\mathsf{Hecke}_n^i \simeq \mathsf{Gr}_{n,X}^i \mathbin{\widetilde{\times}} \mathsf{Bun}_n,$$

with maps

$$\begin{array}{ccccc} & & \mathsf{Hecke}_n^i & & \\ & \overset{\leftarrow}{h} \swarrow & \downarrow{\scriptstyle \mathrm{supp}} & \searrow \vec{h} & \\ \mathsf{Bun}_n & & X & & \mathsf{Bun}_n. \end{array}$$

An S-point classifies a point $x\colon S \to X$ and an inclusion of vector bundles $\mathcal{V}' \hookrightarrow \mathcal{V}$ such that $\mathcal{V}/\mathcal{V}'$ is a length-i torsion sheaf supported on Γ_x with relative position ω_i. Over an algebraically closed point this is the condition $\mathcal{V}/\mathcal{V}' \simeq k(x)^{\oplus i}$, equivalently $\mathcal{V}' \subset \mathcal{V} \subset \mathcal{V}'(x)$ and $\mathrm{length}(\mathcal{V}/\mathcal{V}') = i$.

Put $d_i = i(n-i)$. This is the relative dimension of $\mathsf{Hecke}_{n,x}^i \to \mathsf{Bun}_n$. The normalized Hecke functor is

$$H_n^i(\mathcal{F}) := (\mathrm{supp}, \vec{h})_! \, \overset{\leftarrow}{h}{}^* \mathcal{F}[d_i](d_i/2) \in \mathcal{D}(X \times \mathsf{Bun}_n).$$

The ith Hecke eigenproperty with eigenvalue E is

$$H_n^i(\mathcal{F}) \simeq \wedge^i E \boxtimes \mathcal{F}.$$

2.9. **Orbit notation and closure order.** The global Hecke stack is obtained by twisting the affine Grassmannian with the infinite-level L^+G-torsor over $X \times \mathsf{Bun}_G$:

$$\mathsf{Hecke}_{G,X} \simeq \mathsf{Gr}_{G,X} \mathbin{\widetilde{\times}} \mathsf{Bun}_G.$$

Thus every L^+G-orbit $\mathsf{Gr}_G^\lambda \subset \mathsf{Gr}_G$ gives a corresponding relative stratum $\mathsf{Gr}_{G,X}^\lambda \subset \mathsf{Gr}_{G,X}$ and

$$\mathsf{Hecke}_G^\lambda \simeq \mathsf{Gr}_{G,X}^\lambda \mathbin{\widetilde{\times}} \mathsf{Bun}_G.$$

For $G = \mathrm{GL}_n$, dominant coweights are tuples $\lambda = (\lambda_1 \geqslant \cdots \geqslant \lambda_n)$, and

$$\mathsf{Gr}_n^\lambda := L^+\mathrm{GL}_n \cdot t^\lambda \cdot L^+\mathrm{GL}_n / L^+\mathrm{GL}_n, \qquad t^\lambda = \mathrm{diag}(t^{\lambda_1}, \ldots, t^{\lambda_n}).$$

Inside a fixed connected component, the closure order is the dominance order: $\mathsf{Gr}_n^\lambda \subset \overline{\mathsf{Gr}}_n^\mu$ if and only if $\sum_{j=1}^r \lambda_j \leqslant \sum_{j=1}^r \mu_j$ for all $1 \leqslant r < n$. The minuscule coweight $\omega_i = (1^i, 0^{n-i})$ is the one corresponding to $\wedge^i \mathrm{Std}$ under geometric Satake.

2.10. **Classical automorphic normalization.** Let

$$G_n^i = L^+\mathrm{GL}_n \cdot \mathrm{diag}(\underbrace{t, \ldots, t}_{i}, 1, \ldots, 1) \cdot L^+\mathrm{GL}_n / L^+\mathrm{GL}_n.$$

For a spherical function $f \in \mathcal{C}_c^\infty(L^+\mathrm{GL}_n \backslash L\mathrm{GL}_n / L^+\mathrm{GL}_n)$, the normalized local Hecke operator is

$$T_x^i(f)(y) = q_x^{-d_i/2} \int_{G_n^i} f(yg)\, dg.$$

If the local system E corresponds to a Weil representation $\rho\colon W_F \to \mathrm{GL}_n(\mathbb{Q}_\ell)$, then the expected eigenvalue is $\mathrm{tr}(\mathrm{Frob}_x; \wedge^i \rho)$. Equivalently,

$$\det(1 - \rho(\mathrm{Frob}_x)s) = \sum_{i=0}^n (-1)^i \,\mathrm{tr}(\mathrm{Frob}_x; \wedge^i \rho) s^i.$$

This is the function-theoretic shadow of $H_n^i(\mathcal{F}) \simeq \wedge^i E \boxtimes \mathcal{F}$.

2.11. **Reduction to the first Hecke operator.** Assume that $\mathcal{F} \in \mathrm{Perv}(\mathsf{Bun}_n)$ is equipped with

$$\alpha\colon H_n^1(\mathcal{F}) \xrightarrow{\sim} E \boxtimes \mathcal{F}.$$

For two distinct points $x \neq x'$, the stack of two elementary modifications has an $\mathfrak{S}_2$-action: consider $\mathcal{V}'' \subset \mathcal{V}' \subset \mathcal{V}$ with $\mathrm{length}(\mathcal{V}/\mathcal{V}') = \mathrm{length}(\mathcal{V}'/\mathcal{V}'') = 1$ such that $\mathrm{supp}(\mathcal{V}/\mathcal{V}') = x$ and $\mathrm{supp}(\mathcal{V}'/\mathcal{V}'') = x'$, exchanging x and x' reverses the order of the two elementary modifications. The compatibility required of an eigensheaf is that

$$H_n^1 H_n^1(\mathcal{F})|_{X^2-\Delta} \xrightarrow{\sim} E \boxtimes E \boxtimes \mathcal{F}|_{X^2-\Delta}$$

is $\mathfrak{S}_2$-equivariant. Iterating this gives, over $X^i - \Delta$, an $\mathfrak{S}_i$-equivariant isomorphism

$$(H_n^1)^i(\mathcal{F})|_{X^i-\Delta} \simeq E^{\boxtimes i} \boxtimes \mathcal{F}|_{X^i-\Delta}.$$

Proposition 2.11.1. *Under the above symmetry assumption, for every $0 \leqslant i \leqslant n$,*

$$H_n^i(\mathcal{F}) \simeq \wedge^i E \boxtimes \mathcal{F}.$$

2.12. **Geometric Satake input.** Recall

$$\mathrm{Perv}_{L^+G}(\mathsf{Gr}_G) \simeq \mathrm{Perv}(L^+G\backslash LG/L^+G) \simeq \mathrm{Rep}(\check{G}),$$

along which for all λ we have

$$\mathrm{IC}_{\overline{\mathsf{Gr}}_G^\lambda} \longleftrightarrow V^\lambda.$$

For $G = \mathrm{GL}_n$ this gives

$$\mathrm{Std} \longleftrightarrow (1, 0, \dots, 0) = (1, 0^{n-1}), \qquad \wedge^i \mathrm{Std} \longleftrightarrow (1^i, 0^{n-i}).$$

Hence $\mathrm{IC}^{\star i}_{\mathsf{Gr}_n^1}$ corresponds to $\mathrm{Std}^{\otimes i}$, and the alternating projector for the natural $\mathfrak{S}_i$-action cuts out $\wedge^i \mathrm{Std}$. The global proof realizes this projector by a Springer-type resolution of the length-i modification stack.

2.13. **The compactified stack of length-i modifications.** Let

$$\mathsf{Mod}_n^{-i} := \{(\mathcal{V}' \hookrightarrow \mathcal{V}) \mid \mathcal{V}/\mathcal{V}' \in \mathsf{Tor}^i\}.$$

It has the correspondence

$$\begin{array}{ccccc} & & \mathsf{Mod}_n^{-i} & & \\ & \overset{\leftarrow}{h}\swarrow & \downarrow{\scriptstyle \mathrm{supp}} & \searrow\vec{h} & \\ \mathsf{Bun}_n & & X^{(i)} & & \mathsf{Bun}_n, \end{array}$$

where $\mathrm{supp}(\mathcal{V}/\mathcal{V}')$ is the associated effective divisor. Let $\Delta_i\colon X \to X^{(i)}$ be $x \mapsto i \cdot x$, and put

$$\mathsf{H}_n^{i,+} := X \times_{X^{(i)}} \mathsf{Mod}_n^{-i}.$$

The stratum where $\mathcal{V}/\mathcal{V}' \simeq k(x)^{\oplus i}$ is precisely Hecke_n^i. Thus $\mathsf{H}_n^{i,+}$ compactifies the minuscule correspondence by allowing non-semisimple length-i quotients supported at x.

For $I = \{1, \dots, i\}$, the Beilinson–Drinfeld Grassmannian gives

$$\mathsf{Gr}_{G,X}^I \mathbin{\widetilde{\times}} \mathsf{Bun}_G = \mathsf{Hecke}_{G,X}^I.$$

For $G = \mathrm{GL}_n$ there is a proper map from the positive global Schubert variety to the compactified modification stack:

$$\mathsf{Gr}_{G,X}^{+,\leq(1,\dots,1)} \mathbin{\widetilde{\times}} \mathsf{Bun}_n = \mathsf{Hecke}_{n,X}^{+,\leq(1,\dots,1)} \longrightarrow \mathsf{Mod}_n^{-i}.$$

Over $X^i - \Delta$ it factorizes as

$$\mathsf{Gr}_{G,X}^{+,\leq(1,\dots,1)}|_{X^i-\Delta} \simeq (\mathsf{Gr}_{G,x_1}^1 \times \cdots \times \mathsf{Gr}_{G,x_i}^1)|_{X^i-\Delta}.$$

Now define the ordered, or Springer, resolution

$$\widetilde{\mathsf{Mod}}_n^{-i} := \{\mathcal{V}_0 \hookrightarrow \mathcal{V}_1 \hookrightarrow \cdots \hookrightarrow \mathcal{V}_i \mid \mathrm{length}(\mathcal{V}_j/\mathcal{V}_{j-1}) = 1 \text{ for all } j\}.$$

There is a proper map

$$\varphi\colon \widetilde{\mathsf{Mod}}_n^{-i} \longrightarrow \mathsf{Mod}_n^{-i}, \qquad (\mathcal{V}_0 \subset \cdots \subset \mathcal{V}_i) \longmapsto (\mathcal{V}_0 \subset \mathcal{V}_i),$$

and an ordered support map $r\colon \widetilde{\mathsf{Mod}}_n^{-i} \to X^i$. They fit into

$$\begin{CD} \widetilde{\mathsf{Mod}}_n^{-i} @>{\varphi}>> \mathsf{Mod}_n^{-i} \\ @V{r}VV @VV{\mathrm{supp}}V \\ X^i @>{\mathrm{sym}}>> X^{(i)}. \end{CD}$$

Over the disjoint locus, φ is the ordered-support cover and carries the natural $\mathfrak{S}_i$-action by permuting the order of the length-one quotients.

Claim 2.13.1. *The Springer sheaf*

$$\mathrm{Spr}_i := \varphi_! \mathrm{IC}_{\widetilde{\mathsf{Mod}}_n^{-i}}.$$

satisfies the following properties.

(1) Spr_i *is perverse.*
(2) Spr_i *is the middle extension from the disjoint locus; there it is the local system associated with the ordered-support cover and the regular representation of* $\mathfrak{S}_i$.
(3) *The alternating summand is the minuscule kernel:*

$$\mathrm{Hom}_{\mathfrak{S}_i}(\mathrm{sgn}, \mathrm{Spr}_i)|_{\mathsf{Hecke}_n^i} \simeq \mathbb{Q}_\ell|_{\mathsf{Hecke}_n^i}[d_i](d_i/2), \qquad d_i = i(n-i).$$

This is the same mechanism as for the usual Springer sheaf: the resolution orders the elementary quotients, and the sign representation extracts the exterior-power summand.

2.14. **Projection formula calculation.** The singularity-resolving maps to Bun_n are

$$\begin{array}{ccccc} & & \widetilde{\mathsf{Mod}}_n^{-i} & & \\ & \overset{\widetilde{\overleftarrow{h}}}{\swarrow} & \downarrow{\scriptstyle\varphi} & \overset{\widetilde{\overrightarrow{h}}}{\searrow} & \\ \mathsf{Bun}_n & \xleftarrow[\overleftarrow{h}]{} & \mathsf{Mod}_n^{-i} & \xrightarrow[\overrightarrow{h}]{} & \mathsf{Bun}_n. \end{array}$$

and the support maps form the support square

$$\begin{CD} \widetilde{\mathsf{Mod}}_n^{-i} @>{\varphi}>> \mathsf{Mod}_n^{-i} \\ @V{r}VV @VV{\mathrm{supp}}V \\ X^i @>{\mathrm{sym}}>> X^{(i)}. \end{CD}$$

By projection formula,

$$(\mathrm{supp}, \overrightarrow{h})_!(\overleftarrow{h}^* \mathcal{F} \otimes \mathrm{Spr}_i) \simeq (\mathrm{sym} \times \mathrm{id}_{\mathsf{Bun}_n})_!((H_n^1)^i(\mathcal{F})).$$

Here the normalized shifts and Tate twists are already contained in the IC-normalization of Spr_i and in the functors H_n^1. On $X^i - \Delta$, the right side is the symmetrized iterated first Hecke functor. Using the $\mathfrak{S}_i$-equivariant eigen-isomorphism and the middle-extension property gives

$$(\mathrm{sym} \times \mathrm{id}_{\mathsf{Bun}_n})_!((H_n^1)^i(\mathcal{F})) \simeq \mathrm{sym}_!(E^{\boxtimes i}) \boxtimes \mathcal{F}.$$

Taking the sign-isotypic summand and then restricting along $\Delta_i\colon X \to X^{(i)}$ yields

$$\Delta_i^* \mathrm{Hom}_{\mathfrak{S}_i}(\mathrm{sgn}, \mathrm{sym}_! E^{\boxtimes i}) \simeq \wedge^i E.$$

Combining this identity with the sign-Springer description of the kernel gives

$$H_n^i(\mathcal{F}) \simeq \wedge^i E \boxtimes \mathcal{F}.$$

This proves the proposition.

3. Whittaker

This section changes viewpoint from eigensheaves on Bun_n to Whittaker models. The function-theoretic statement is the uniqueness of a normalized Whittaker vector with prescribed spherical Hecke eigenvalues. Its geometric shadow is the Whittaker model of the Satake category, and for GL_n it recovers the cuspidal automorphic sheaf constructed from a local system [FGKV98, Gai18].

3.1. **Reminder: Hecke eigensheaves and cuspidality.** Let $\mathsf{Bun}_n := \mathsf{Bun}_{\mathrm{GL}_n}$. For $0 \leqslant d \leqslant n$, the dth minuscule Hecke correspondence is

$$\begin{array}{ccccc} & & \mathcal{H}_n^d & & \\ & \overset{\overleftarrow{h}}{\swarrow} & \big\downarrow \scriptstyle{\text{supp}} & \overset{\overrightarrow{h}}{\searrow} & \\ \mathsf{Bun}_n & & X & & \mathsf{Bun}_n. \end{array}$$

An S-point is a tuple

$$(x, M, M', M' \subset M \subset M'(x)),$$

where M and M' are rank n vector bundles and M/M' is finite locally free of rank d over the graph of x. The normalized Hecke functor is

$$H_n^d(K) := (\mathrm{supp}, \overrightarrow{h})_! \, \overleftarrow{h}^* K \left(\frac{d(n-d)}{2}\right) [d(n-d)] \in \mathcal{D}(X \times \mathsf{Bun}_n).$$

For a local system E on X, we say that K satisfies the dth Hecke eigenproperty with eigenvalue E if

$$H_n^d(K) \simeq \wedge^d E \boxtimes K.$$

As in Section 2, if K is perverse and satisfies the first eigenproperty, together with the usual symmetry compatibility for iterated modifications, then it satisfies all the minuscule properties.

The main existence theorem used here is the Frenkel–Gaitsgory–Vilonen theorem [FGKV98]: if E is geometrically irreducible, then there exists a cuspidal Hecke eigensheaf Aut_E on Bun_n, and its restriction to each connected component Bun_n^d is irreducible.

We also recall the geometric definition of cuspidality. Let G be reductive, let $P \subsetneq G$ be a proper parabolic with unipotent radical U and Levi quotient $M = P/U$. Then

$$\mathsf{Bun}_G \xleftarrow{\;p\;} \mathsf{Bun}_P \xrightarrow{\;q\;} \mathsf{Bun}_M$$

defines the constant-term functor

$$\mathrm{CT}_! := q_! p^*.$$

A sheaf $K \in \mathcal{D}(\mathsf{Bun}_G)$ is *cuspidal* if $\mathrm{CT}_!(K) = 0$ for every proper P.

3.2. **Classical Hecke eigenfunctions.** In the function-theoretic paragraphs below assume that k is finite. Now let $F = k(X)$, let $\mathbb{A}$ be the adele ring of F, and put $G = \mathrm{GL}_n$. Let

$$\sigma \colon \mathrm{Gal}(\overline{F}/F) \longrightarrow \mathrm{GL}_n(\mathbb{Q}_\ell)$$

be everywhere unramified, equivalently a rank n local system on X. A spherical automorphic function is a function on

$$\mathrm{GL}_n(F)\backslash \mathrm{GL}_n(\mathbb{A})/\mathrm{GL}_n(\mathcal{O}).$$

For a closed point $x \in |X|$, write K_x for the local field, $\mathcal{O}_x$ for its ring of integers, $q_x = \#k(x)$, and choose a uniformizer t_x. Set

$$M_n^d(\mathcal{O}_x) := \mathrm{GL}_n(\mathcal{O}_x)\, t_x^{(1^d, 0^{n-d})}\, \mathrm{GL}_n(\mathcal{O}_x) \subset \mathrm{GL}_n(K_x) \subset \mathrm{GL}_n(\mathbb{A}).$$

The unnormalized local Hecke operator is

$$T_x^d f(g) = \int_{M_n^d(\mathcal{O}_x)} f(gh)\, dh.$$

We say that f satisfies the dth Hecke eigenproperty for σ if, for every $x \in |X|$,

$$T_x^d f = q_x^{d(n-d)/2}\, \mathrm{Tr}(\wedge^d \sigma(\mathrm{Fr}_x)) f.$$

Equivalently, let γ_x be the normalized Satake parameter characterized by

$$\mathrm{Tr}(\wedge^d \gamma_x) = q_x^{d(n-1)/2}\, \mathrm{Tr}(\wedge^d \sigma(\mathrm{Fr}_x)), \qquad 0 \leqslant d \leqslant n.$$

In the Tate-twist convention this is denoted $\gamma_x = \sigma((n-1)/2)(\mathrm{Fr}_x)$. Then the same eigenvalue is

$$q_x^{-d(d-1)/2}\, \mathrm{Tr}(\wedge^d \gamma_x).$$

For a parabolic $P \subsetneq \mathrm{GL}_n$ with unipotent radical U, a spherical automorphic function is *cuspidal* if

$$\int_{U(F)\backslash U(\mathbb{A})} f(ug)\, du = 0, \qquad g \in \mathrm{GL}_n(\mathbb{A}).$$

Lafforgue's theorem, with the usual finite-determinant normalization, attaches to such an irreducible σ a cuspidal automorphic representation with the above unramified eigenvalues [Laf18]. The spherical vector is unique up to scalar; the Whittaker normalization fixes the scalar.

3.3. **The Whittaker character and the ρ-shift.** Let $N \subset \mathrm{GL}_n$ be the upper triangular unipotent subgroup. Fix a nontrivial additive character $\psi\colon k \to \mathbb{Q}_\ell^\times$. At x, the standard nondegenerate Whittaker character is obtained from the simple-root coordinates by

$$\begin{aligned}\dot N(K_x) \longrightarrow \dot N/[\dot N, \dot N](K_x) &\xrightarrow{\sim} \textstyle\bigoplus_{\alpha\in\Delta} K_x \otimes \Omega_x \\ &\xrightarrow{\mathrm{Res}} \textstyle\bigoplus_{\alpha\in\Delta} k(x) \xrightarrow{\mathrm{Tr}_{k(x)/k}\circ\Sigma} k \xrightarrow{\psi} \mathbb{Q}_\ell^\times.\end{aligned}$$

Here Ω_x is the completed canonical line at x. Denote the resulting character by $\Psi_x\colon \dot N(K_x) \to \mathbb{Q}_\ell^\times$.

The dot denotes the ρ-shift. For $G = \mathrm{GL}_n$ it may be written locally as

$$\dot G_n(K_x) = \{(a_{ij}) \mid a_{ij} \in K_x \otimes \Omega_x^{\otimes(j-i)}\}.$$

Thus the first superdiagonal entries lie in $K_x\Omega_x$, so residues can be applied to their simple-root coordinates. A Whittaker function on $\dot G(K_x)$ is a function W satisfying

$$W(ug) = \Psi_x(u)W(g), \qquad u \in \dot N(K_x).$$

Usually one also imposes right $\dot G(\mathcal{O}_x)$-invariance.

3.4. **The local Whittaker eigenfunction.** Let $\gamma \in \mathrm{GL}_n(\mathbb{Q}_\ell)$, for example the normalized Satake parameter γ_x above. There is a unique normalized Hecke-eigen Whittaker function $W_{\gamma,x}$ on $\dot G(K_x)/\dot G(\mathcal{O}_x)$ satisfying

(1) $W_{\gamma,x}(1) = 1$, where 1 denotes the base point of the affine Grassmannian;
(2) $W_{\gamma,x}(gh) = W_{\gamma,x}(g)$ for $h \in \dot G(\mathcal{O}_x)$;
(3) $W_{\gamma,x}(ug) = \Psi_x(u)W_{\gamma,x}(g)$ for $u \in \dot N(K_x)$;
(4) $T_x^d W_{\gamma,x} = q_x^{-d(d-1)/2}\,\mathrm{Tr}(\wedge^d\gamma)W_{\gamma,x}$ for every $0 \leqslant d \leqslant n$.

After choosing the local coordinate used to write t_x^λ, the Whittaker double quotient is indexed by the coweight lattice:

$$\dot N(K_x)\backslash \dot G(K_x)/\dot G(\mathcal{O}_x) \simeq \Lambda, \qquad \dot N(K_x)t_x^\lambda \dot G(\mathcal{O}_x) \longleftrightarrow \lambda.$$

The Casselman–Shalika formula gives the values explicitly; its geometric form is the Whittaker realization of Satake [FGKV98]. For a coweight $\lambda = (\lambda_1, \dots, \lambda_n)$,

$$W_{\gamma,x}(t_x^\lambda) = \begin{cases} q_x^{-\sum_i (i-1)\lambda_i}\,\mathrm{Tr}(\gamma, V^\lambda), & \lambda \text{ dominant},\\ 0, & \text{otherwise},\end{cases}$$

where, for dominant λ, V^λ denotes the irreducible $\check G$-representation of highest weight λ. Categorically this is the Whittaker form of geometric Satake:

$$\mathrm{Perv}^{\mathrm{Wh}}(\dot{\mathsf{Gr}}_{G,x}) \simeq \mathrm{Rep}(\check G),$$

or, at the derived level, $D^{\mathrm{Wh}}(\dot{\mathsf{Gr}}_{G,x}) \simeq \mathcal{D}^{\mathrm{b}}(\mathrm{Rep}(\check G))$.

3.5. **Global Whittaker functions and the automorphic function.** The local characters multiply to a global character

$$\Psi = \prod_x \Psi_x\colon \dot N(\mathbb{A}) \longrightarrow \mathbb{Q}_\ell^\times,$$

which is trivial on $\dot N(F)$ by the residue theorem. The preceding local theorem gives a restricted product

$$W_\sigma \coloneqq {\prod_x}' W_{\gamma_x,x}$$

on $(\dot N(\mathbb{A}), \Psi)\backslash \dot G(\mathbb{A})/\dot G(\mathcal{O})$. There is a unique normalized unramified Hecke-eigen Whittaker function on $\dot G(\mathbb{A})$ characterized by $W(1) = 1$ and

(1) $W(gh) = W(g)$ for $h \in \dot G(\mathcal{O})$;
(2) $W(ug) = \Psi(u)W(g)$ for $u \in \dot N(\mathbb{A})$;
(3) $T_x^d W = q_x^{d(n-d)/2}\,\mathrm{Tr}(\wedge^d\sigma(\mathrm{Fr}_x))W$.

For GL_n, Whittaker expansion first gives a cuspidal function on the mirabolic quotient. Let

$$Q_1 = \begin{pmatrix} \mathrm{GL}_{n-1} & V_{n-1} \\ 0 & 1 \end{pmatrix}$$

be the mirabolic subgroup. The canonical map is

$$\begin{aligned} \phi\colon \mathcal{C}^\infty((\dot N(\mathbb{A}),\Psi)\backslash \dot G(\mathbb{A})) &\xrightarrow{\ \sim\ } \mathcal{C}^\infty(\dot Q_1(F)\backslash \dot G(\mathbb{A}))_{\mathrm{cusp}} \\ W &\longmapsto \big(g \mapsto \textstyle\sum_{\gamma\in \dot N(F)\backslash \dot Q_1(F)} W(\gamma g)\big). \end{aligned}$$

Equivalently, the sum is the integral over the corresponding discrete quotient. For the normalized Hecke-eigen function W_σ, the Piatetski-Shapiro–Shalika automorphy theorem says that $\phi(W_\sigma)$ is in fact left $\dot G(F)$-invariant; this is the cuspidal Hecke eigenfunction f_σ.

3.6. **Geometric spoiler: Whittaker categories.** The local statement above has the following geometric form [Gai18]:

$$\mathrm{Wh}(\dot{\mathsf{Gr}}_{G,x}) \simeq \mathrm{Rep}(\check G).$$

After chiral integration over the curve, one expects

$$\int_X^{\mathrm{ch}} \mathrm{Wh}(\dot{\mathsf{Gr}}_{G,x}) \simeq \int_X^{\mathrm{ch}} \mathrm{Rep}(\check G).$$

The left side is the sheaf-theoretic replacement of the adelic Whittaker quotient, while the right side maps to quasi-coherent sheaves on the stack of $\check G$-local systems. The picture is

$$\begin{array}{ccc} \int_X^{\mathrm{ch}} \mathrm{Wh}(\dot{\mathsf{Gr}}_{G,x}) & \xrightarrow{\ \sim\ } & \int_X^{\mathrm{ch}} \mathrm{Rep}(\check G) \\ \downarrow{\scriptstyle \text{glob.}} & & \downarrow \\ \mathrm{Shv}((\dot N(\mathbb{A}),\Psi)\backslash \dot G(\mathbb{A})/\dot G(\mathcal{O})) & & \mathrm{QCoh}(\mathsf{LocSys}_{\check G}) \\ \downarrow{\scriptstyle \mathrm{Av}_!} & & \uparrow \\ \mathrm{Shv}(\dot Q_1(F)\backslash \dot G(\mathbb{A})/\dot G(\mathcal{O}))_{\mathrm{cusp}} & & \mathrm{QCoh}(\mathsf{LocSys}_{\check G}^{\mathrm{irred}}) \\ \uparrow & & \| \\ \mathrm{Shv}(\mathsf{Bun}_G)_{\mathrm{cusp}} & \xrightarrow{\ \text{conj.}\ } & \mathrm{QCoh}(\mathsf{LocSys}_{\check G}^{\mathrm{irred}}) \end{array}$$

Thus Whittaker models are the local-to-global bridge between Hecke eigensheaves on Bun_G and sheaves on $\mathsf{LocSys}_{\check G}$.

For the global comparison, write $K = k(X)$ for the global function field and $\mathbb{A}$ for its adeles. In the function-theoretic statements k is finite; geometrically ψ is replaced by the Artin–Schreier sheaf. Let $G = \mathrm{GL}_n$, let $N \subset B \subset G$ be the standard maximal unipotent and Borel, and fix a nontrivial additive character $\psi\colon k \to \mathbb{Q}_\ell^\times$. To reduce notation, the dot is suppressed here: whenever the Whittaker character Ψ is present, the adelic groups are the ρ-twisted forms described above. Let

$$Q_1 = \begin{pmatrix} \mathrm{GL}_{n-1} & V \\ 0 & 1 \end{pmatrix}$$

be the mirabolic subgroup. The classical goal, geometrized in the Whittaker formalism of [FGKV98, Gai18], is the Whittaker expansion map

$$\mathrm{Av}_!\colon \mathcal{C}^\infty((N(\mathbb{A}),\Psi)\backslash G(\mathbb{A})) \xrightarrow{\ \sim\ } \mathcal{C}^\infty(Q_1(K)\backslash G(\mathbb{A}))_{\mathrm{cusp}}.$$

It sends the global Whittaker function $W_\sigma = \prod'_x W_{\gamma_x,x}$ to a cuspidal function f_σ on the mirabolic quotient. Here γ_x is the normalized Satake parameter introduced above, and the local functions $W_{\gamma_x,x} \in \mathcal{C}^\infty((N(K_x),\Psi_x)\backslash G(K_x)/G(\mathcal{O}_x))$ are the normalized local Whittaker functions from the Casselman–Shalika theorem. The remaining automorphy statement is that the resulting f_σ is not only $Q_1(K)$-invariant, but actually $G(K)$-invariant.

Equivalently, since $N \subset Q_1$, we have

$$(\mathrm{Av}_!\, W)(g) = \sum_{\gamma\in N(K)\backslash Q_1(K)} W(\gamma g).$$

Cuspidality makes this summation convergent in the function-theoretic setting. Geometrically it is replaced by a !-pushforward along the corresponding map of moduli spaces of generic reductions.

3.7. **Generic reductions and the Whittaker character.** Let $\mathsf{Bun}_G^{B\text{-gen}}$ and $\mathsf{Bun}_G^{N\text{-gen}}$ denote the prestacks of G-bundles equipped with a generically defined B-, respectively N-reduction. There is a commutative diagram

$$\begin{array}{ccc} \mathsf{Bun}_G^{N\text{-gen}} & \longrightarrow & \mathsf{Bun}_G^{B\text{-gen}} \\ \uparrow{\scriptstyle \mathrm{str}} & & \uparrow{\scriptstyle \mathrm{str}} \\ \mathsf{Bun}_B^{N\text{-gen}} & \xrightarrow{L} & \mathsf{Bun}_B. \end{array}$$

Here str remembers the underlying G-bundle, and L forgets the generic N-structure. A generic N-reduction of a B-bundle is equivalently a generic trivialization of its induced T-bundle, so

$$\mathsf{Bun}_B^{N\text{-gen}} = \mathsf{Bun}_B \times_{\mathsf{Bun}_T} \mathsf{Bun}_T^{\text{rat triv}} \simeq \mathsf{Bun}_B \times_{\mathsf{Bun}_T} \mathrm{Div}_\Lambda,$$

where Div_Λ is the colored-divisor space for the coweight monoid. This is the geometric replacement for allowing a rational reduction and recording its defect divisor.

For $G = \mathrm{GL}_n$ we write Bun_N here for the ρ-twisted Whittaker stack with fixed T-bundle:

$$\mathsf{Bun}_N = \left\{0 = \overline{F}_0 \subset \overline{F}_1 \subset \cdots \subset \overline{F}_n, \quad \overline{F}_i/\overline{F}_{i-1} \simeq \Omega_X^{n-i}\right\}.$$

The simple-root extension classes give a map

$$\ell\colon \mathsf{Bun}_N \longrightarrow \prod_{i=1}^{n-1} \mathbb{G}_{\mathrm{a}} \xrightarrow{\Sigma} \mathbb{G}_{\mathrm{a}}.$$

Indeed, for $1 \leqslant i \leqslant n-1$, the ith adjacent extension

$$0 \to \Omega_X^{n-i} \to \overline{F}_{i+1}/\overline{F}_{i-1} \to \Omega_X^{n-i-1} \to 0$$

defines an element of $\mathrm{Ext}^1(\Omega_X^{n-i-1}, \Omega_X^{n-i}) \simeq H^1(X, \Omega_X) \simeq k$, and ℓ is the sum of these classes. Pulling back the Artin–Schreier sheaf $\mathcal{L}_\psi$ by ℓ gives the global Whittaker sheaf. We write $\mathcal{D}(\mathsf{Bun}_G^{N\text{-gen}})^{\mathrm{Wh}}$ for the category with this twisted $N(\mathbb{A})$-equivariance.

For $G = \mathrm{GL}_n$, the geometric form of the mirabolic Whittaker expansion is therefore

$$\mathrm{Av}_!\colon \mathcal{D}(\mathsf{Bun}_G^{N\text{-gen}})^{\mathrm{Wh}} \xrightarrow{\sim} \mathcal{D}(\mathsf{Bun}_G^{Q_1\text{-gen}})_{\mathrm{cusp}}.$$

It should be read together with the local Whittaker–Satake equivalence and the spectral side:

$$\begin{array}{ccc} \mathcal{D}(\mathsf{Bun}_G^{N\text{-gen}})^{\mathrm{Wh}} & \xrightarrow[\sim]{\mathrm{Av}_!} & \mathcal{D}(\mathsf{Bun}_G^{Q_1\text{-gen}})_{\mathrm{cusp}} \\ \downarrow{\scriptstyle \simeq} & & \downarrow{\scriptstyle \mathrm{spectral}} \\ \int_X^{\mathrm{ch}} \mathrm{Wh}(\dot{\mathsf{Gr}}_{G,x}) & \xrightarrow{\sim} & \int_X^{\mathrm{ch}} \mathrm{Rep}(\check{G}) \\ \downarrow & & \downarrow \\ \mathrm{Aut}_E \in \mathcal{D}(\mathsf{Bun}_G)_{\mathrm{cusp}} & \overset{\mathrm{conj.}}{\dashrightarrow} & \delta_E \in \mathrm{QCoh}((\mathsf{LocSys}_{\check{G}})_{\mathrm{irred}}). \end{array}$$

Here δ_E is the skyscraper at the local system E, and the chiral integral on the left is the globalization of the local Whittaker category.

3.8. **Degenerate characters and parabolics.** The characters of $N(\mathbb{A})$ that are trivial on $N(K)$ and factor through the simple-root quotient are controlled by the simple-root coordinates:

$$N(K)^\perp := \{\chi\colon N(\mathbb{A}) \to \overline{\mathbb{Q}}_\ell^\times \mid \chi|_{N(K)} = 1\}_{\mathrm{sr}} \simeq K^{\oplus(n-1)}.$$

Here "sr" means simple-root quotient. The torus $T(K)$ acts through the simple roots, and the open subset with all coordinates nonzero is identified with $T_{\mathrm{ad}}(K)$. Thus the $T(K)$-orbits are indexed by which simple-root coordinates are nonzero. Equivalently,

$$\{T(K)\text{-orbits in } N(K)^\perp\} \longleftrightarrow \{J \subset \mathrm{Dyn}(G)\} \longleftrightarrow \{P \mid B \subset P \subset G\},$$

so there are 2^{n-1} orbits. For a parabolic P with Levi M, let $J(P)$ be the set of simple roots of M, and denote by $\Psi_P\colon N(\mathbb{A}) \to \mathbb{Q}_\ell^\times$ a representative character whose simple-root coordinates are nonzero exactly on $J(P)$. The generic character is $\Psi = \Psi_G$, while the zero character is the orbit attached to $P = B$.

The extended Piatetski-Shapiro–Shalika theorem uses all these Fourier coefficients; this is the function-theoretic input behind the Whittaker comparison in [FGKV98]. Let now

$$Q = \begin{pmatrix} \mathrm{GL}_{n-1} & V \\ 0 & \mathbb{G}_\mathrm{m} \end{pmatrix}$$

be the maximal parabolic of type $(n-1, 1)$. There is a Fourier-coefficient map

$$\mathcal{C}^\infty(Q(K)\backslash G(\mathbb{A})) \longrightarrow \prod_{P \supset B} \mathcal{C}^\infty((Z_M(K) \ltimes N(\mathbb{A}), \Psi_P)\backslash G(\mathbb{A})),$$

where Z_M is the center of the Levi of P and stabilizes Ψ_P; the notation means that $Z_M(K)$ acts trivially on the character. Its restriction to the cuspidal part has all degenerate components zero and is identified with the $P = G$ factor, i.e. the nondegenerate Whittaker coefficient. This recovers the mirabolic theorem above after passing from Q to Q_1.

3.9. **Iterated Fourier transform along the mirabolic.** For the iterative proof use the mirabolic $Q_1 = L_1 \ltimes V$, with $L_1 \simeq \mathrm{GL}_{n-1}$, rather than the full maximal parabolic. Fourier transform on $V(K)\backslash V(\mathbb{A})$ uses

$$V(K)^\perp = \mathrm{Ch}(V(\mathbb{A})/V(K)).$$

The $L_1(K)$-action on $V(K)^\perp$ has two orbits:

(1) the open orbit through a nonzero character η_{n-1}, whose stabilizer is the smaller mirabolic subgroup $Q_1'(K) \subset \mathrm{GL}_{n-1}(K)$;
(2) the closed orbit $\{0\}$, whose stabilizer is $L_1(K)$.

Consequently Fourier transform gives the decomposition

$$\mathcal{C}^\infty(Q_1(K)\backslash G(\mathbb{A})) \simeq \mathcal{C}^\infty((Q_1'(K) \ltimes V(\mathbb{A}), \eta_{n-1})\backslash G(\mathbb{A})) \times \mathcal{C}^\infty(L_1(K) \ltimes V(\mathbb{A})\backslash G(\mathbb{A})).$$

If $n \geqslant 3$, then $Q_1' = L_1' \ltimes V'$ is the mirabolic subgroup inside GL_{n-1}. The first factor is again a Whittaker-type space for rank $n-1$, because $Q_1'(\mathbb{A})$ acts on $\mathcal{C}^\infty((V(\mathbb{A}), \eta_{n-1})\backslash G(\mathbb{A}))$. The second factor is contained in $\mathcal{C}^\infty(Q_1'(K) \ltimes V(\mathbb{A})\backslash G(\mathbb{A}))$ and is killed by cuspidality. Iterating this rank-lowering Fourier transform leaves only the generic Whittaker coefficient.

For the maximal parabolic Q, the categorical analogue is expressed as a fully faithful functor into a glued extended Whittaker category:

$$\mathcal{D}(\mathsf{Bun}_G^{Q\text{-gen}}) \hookrightarrow \sum_{P \supset B} \mathcal{D}(\mathsf{Bun}_G^{Z_M\text{-gen}, \mathrm{Wh}_P}).$$

The piece indexed by P geometrizes the degenerate coefficient with character Ψ_P and the $Z_M(K)$-reduction. For $G = \mathrm{GL}_n$ and P of type $(n-1, 1)$, this is the geometric shadow of the classical iterative Fourier-transform proof.

3.10. **A compactified variant using coherent sheaves.** With the right-action convention for which Q_1 stabilizes the chosen vector, a Q_1-generic reduction for $G = \mathrm{GL}_n$ may be written as a nonzero rational vector

$$\mathsf{Bun}_G^{Q_1\text{-gen}} = \left\{(M, \Omega_X^{n-1} \dashrightarrow M), \quad \text{not identically } 0\right\}.$$

A regular version is

$$\mathsf{Bun}_n' = \{(M, s\colon \Omega_X^{n-1} \to M),\ s \neq 0 \text{ generically},\ M \in \mathsf{Bun}_n\},$$

with the forgetful map $\mathsf{Bun}_n' \to \mathsf{Bun}_n$. Since M is locally free and X is a smooth curve, a map s that is nonzero at the generic point is injective as a map of coherent sheaves, and it gives

$$0 \to \Omega_X^{n-1} \to M \to M' \to 0.$$

Here M' is a coherent sheaf of rank $n-1$; it need not be locally free because the image of s need not be saturated. The extension class lies in $\mathrm{Ext}^1(M', \Omega_X^{n-1})$, and Serre duality gives

$$\mathrm{Ext}^1(M', \Omega_X^{n-1})^* \simeq \mathrm{Hom}(\Omega_X^{n-2}, M').$$

This is exactly the dual vector space on which the next Fourier transform lives.

We introduce the compactified stacks

$$\mathsf{Coh}'_m = \{(F,\, \Omega_X^{m-1} \xrightarrow{s} F),\ F \in \mathsf{Coh},\ \operatorname{rank} F = m\}, \qquad {}^{\neq 0}\mathsf{Coh}'_m \subset \mathsf{Coh}'_m$$

where the open substack is defined by $s \neq 0$ at the generic point. More precisely, for $m = n, n-1, \dots, 2$ let $\mathcal{E}_m \to \mathsf{Coh}_{m-1}$ be the relative vector stack whose fiber over M' is $\operatorname{Ext}^1(M', \Omega_X^{m-1})$. Serre duality identifies the dual relative vector stack with $\mathcal{E}_m^\vee \simeq \mathsf{Coh}'_{m-1}$, whose fiber over M' is $\operatorname{Hom}(\Omega_X^{m-2}, M')$. The map $\mathcal{E}_m \to \mathsf{Coh}'_m$ remembers the middle term with its tautological section; set ${}^{\neq 0}\mathcal{E}_m := \mathcal{E}_m \times_{\mathsf{Coh}'_m} {}^{\neq 0}\mathsf{Coh}'_m$. The successive Fourier transforms are schematically

$${}^{\neq 0}\mathsf{Coh}'_m \xleftarrow{p_m} {}^{\neq 0}\mathcal{E}_m \xrightarrow{\mathbb{F}} \mathcal{E}_m^\vee \simeq \mathsf{Coh}'_{m-1}, \qquad m = n, n-1, \dots, 2.$$

The first transform defines a Fourier–Deligne functor, in the normalization of [KL85],

$$\mathbb{F}_n \colon \mathcal{D}({}^{\neq 0}\mathsf{Coh}'_n) \longrightarrow \mathcal{D}(\mathsf{Coh}'_{n-1}).$$

Repeating this construction lowers the rank one step at a time. Hence, starting from an E-Hecke eigensheaf on ${}^{\neq 0}\mathsf{Coh}'_n$, one obtains a sheaf on ${}^{\neq 0}\mathsf{Coh}'_1$. The rank-one endpoint is forced to be the Whittaker sheaf; this is the compactified form of the Whittaker normalization used in Laumon's construction.

4. Laumon's sheaf

Let X be a smooth projective connected curve over k, and let E be a rank n local system on X. The local descriptions are étale-local, so the examples on $\mathbb{A}^1$ below are only local models. The goal is to construct a perverse sheaf $\mathcal{L}_E$ on the stack Tor of torsion sheaves on X. This is the sheaf-theoretic input that replaces the Whittaker function in Laumon's construction of automorphic sheaves for GL_n [FGKV98]; parabolic and ramified variants of this construction are discussed in [Hei04].

4.1. **Motivation from Whittaker and Satake.** The favorable picture from the Whittaker lectures (see Section 3) is

$$\begin{array}{ccc}
\displaystyle\int_X^{\mathrm{ch}} \mathrm{Wh}(\dot{\mathsf{Gr}}_{G,x}) & \xrightarrow{\quad\sim\quad} & \displaystyle\int_X^{\mathrm{ch}} \mathrm{Rep}(\check{G}) \\
\uparrow & & \uparrow \\
W_E \in \mathrm{Wh}(\mathsf{Bun}_G^{N\text{-gen}}) & & \delta_E \in \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}) \\
\big\downarrow{\scriptstyle \mathrm{Av}_!} & & \\
\mathrm{Aut}_E \in \mathcal{D}(\mathsf{Bun}_G) & &
\end{array}$$

Here δ_E is the skyscraper at the local system E. One should think of the desired Whittaker object as

$$W_E \simeq \delta_E * W_0,$$

where W_0 is the vacuum Whittaker sheaf. Applying a representation V of $\check{G}$ to the whole $\check{G}$-local system E gives a local system V_E on X, with fiber $V_{E,x}$ at x. For $G = \mathrm{GL}_n$ this is the usual operation on the rank-n local system E. Laumon's sheaf is the concrete incarnation of this principle on the positive Hecke locus for $G = \mathrm{GL}_n$. The dot in $\dot{\mathsf{Gr}}_{G,x}$ is the same ρ-twist used for Whittaker sheaves; the spherical Satake input below is written on the ordinary affine Grassmannian.

The local input is geometric Satake:

$$\mathrm{Perv}_{G(\mathcal{O}_x)}(\mathsf{Gr}_{G,x}) \simeq \mathrm{Rep}(\check{G}), \qquad \mathrm{IC}_{\overline{\mathsf{Gr}}^\lambda_{G,x}} \longleftrightarrow V^\lambda, \quad \lambda \in \Lambda^+.$$

We use $\mathsf{Gr}_{G,x} \simeq G(K_x)/G(\mathcal{O}_x)$ and $\mathsf{Gr}^\lambda_{G,x} = G(\mathcal{O}_x)\, t_x^\lambda\, G(\mathcal{O}_x)/G(\mathcal{O}_x)$.

There is also the geometric Casselman–Shalika form

$$\mathrm{Wh}(\dot{\mathsf{Gr}}_{G,x}) \simeq \mathrm{Rep}(\check{G}), \qquad W_\lambda \longleftrightarrow V^\lambda,$$

compatible with the Satake action. In the ρ-twisted notation of Sections 3, the Whittaker orbits are

$$S^\lambda := \dot{N}(K_x) t_x^\lambda \dot{G}(\mathcal{O}_x)/\dot{G}(\mathcal{O}_x), \qquad \lambda \in \Lambda.$$

There exists a $(\dot{N}(K_x), \Psi_x)$-equivariant sheaf on S^λ iff λ is dominant: indeed

$$\mathrm{Stab}_{\dot{N}(K_x)}(t_x^\lambda) \subset \ker(\Psi_x) \iff \lambda \in \Lambda^+.$$

For such λ, the corresponding twisted local system on S^λ has clean extension to $\dot{\mathsf{Gr}}_{G,x}$. The one-dimensional model is the Artin–Schreier sheaf: for $j\colon \mathbb{A}^1 \hookrightarrow \mathbb{P}^1$ and $i_\infty\colon \{\infty\} \hookrightarrow \mathbb{P}^1$,

$$i_\infty^* j_* \mathcal{L}_\psi = 0, \qquad \text{hence} \qquad j_! \mathcal{L}_\psi \simeq j_* \mathcal{L}_\psi.$$

The point is the vanishing of the boundary stalk of $j_*\mathcal{L}_\psi$ at ∞; $i_\infty^! j_!$ would vanish tautologically for an extension by zero. This is the one-dimensional model for clean extension of a nondegenerate Whittaker local system.

4.2. **From Satake kernels to torsion sheaves.** For $V \in \mathrm{Rep}(\check{G})$, Satake gives a perverse sheaf $\mathrm{Sat}(V)$ on the local spherical affine Grassmannian. The regular representation $\mathcal{O}(\check{G})$ is an ind-object of $\mathrm{Rep}(\check{G})$; applying Satake to it formally gives an ind-spherical Hecke kernel, not a single finite-type perverse sheaf. For $G = \mathrm{GL}_n$, this kernel is a local model for the global Hecke stack of triples $(M, M', \beta\colon M \simeq M'$ generically). The adelic double quotient $\mathrm{GL}_n(\mathcal{O})\backslash\mathrm{GL}_n(\mathbb{A})/\mathrm{GL}_n(\mathcal{O})$ records only the relative adelic position after choosing generic trivializations, so it should not be identified with the full global Hecke stack without that extra choice.

For $G = \mathrm{GL}_n$ we isolate the positive locus

$$\mathsf{Gr}^+_{n,x} = \bigcup_{\lambda_1 \geqslant \cdots \geqslant \lambda_n \geqslant 0} \mathsf{Gr}^\lambda_{n,x} = (\mathrm{Mat}_{n\times n}(\mathcal{O}_x) \cap \mathrm{GL}_n(K_x))/\mathrm{GL}_n(\mathcal{O}_x).$$

It consists of modifications defined on the unit disc with torsion cokernel. The full affine Grassmannian is obtained from this positive part after allowing arbitrary central shifts. The particular shift

$$(\mathcal{O}_x^n \xrightarrow{\alpha} \mathcal{E}) \longmapsto (\mathcal{O}_x(-x)^n \hookrightarrow \mathcal{O}_x^n \xrightarrow{\alpha} \mathcal{E}),$$

adds $(1, \ldots, 1)$ to the coweight. On the global positive Hecke stack one has a map

$$p\colon \mathsf{Hecke}^+_{\mathrm{GL}_n} \longrightarrow \mathsf{Tor}, \qquad (M \hookrightarrow M') \longmapsto M'/M,$$

where $\mathsf{Hecke}^+_{\mathrm{GL}_n}$ classifies inclusions of rank-n bundles with torsion quotient. Thus the universal positive Hecke kernel should be pulled back from a sheaf on Tor. This sheaf is $\mathcal{L}_E$.

4.3. **The stack of torsion sheaves.** Define

$$\mathsf{Tor}(S) \coloneqq \{\mathcal{M} \in \mathrm{Coh}(X \times S) \mid \mathcal{M} \text{ is } S\text{-flat and every geometric fiber is torsion}\}.$$

It decomposes by length:

$$\mathsf{Tor} = \bigsqcup_{d \geqslant 0} \mathsf{Tor}^d.$$

Each Tor^d is an algebraic stack. When X is proper and $H^0(X, \mathcal{O}_X) = k$, one convenient atlas is obtained from the open part of the Quot scheme

$$\mathrm{Quot}^{\circ,d}_{\mathcal{O}_X^{\oplus d}} = \{\mathcal{O}_X^{\oplus d} \twoheadrightarrow \mathcal{M},\ \mathcal{M} \in \mathsf{Tor}^d,\ H^0(\mathcal{O}_X^{\oplus d}) \xrightarrow{\sim} H^0(\mathcal{M})\}.$$

Changing the chosen basis of $H^0(\mathcal{M})$ gives the action of GL_d, and the quotient stack is

$$[\mathrm{Quot}^{\circ,d}_{\mathcal{O}_X^{\oplus d}}/\mathrm{GL}_d] \simeq \mathsf{Tor}^d.$$

For example, if $X = \mathbb{A}^1$, a length-d torsion sheaf is the same as a d-dimensional vector space with one endomorphism, hence

$$\mathsf{Tor}^d \simeq \mathfrak{gl}_d/\mathrm{GL}_d.$$

4.4. **Global stratification by partitions.** Let

$$\mathrm{Part}^d = \{(d_1 \geqslant d_2 \geqslant \cdots \geqslant d_s > 0) \mid d_1 + \cdots + d_s = d\}.$$

The stack Tor^d is stratified by locally closed substacks $\mathsf{Tor}^{(d_1,\ldots,d_s)}$. Put $d_{s+1} = 0$. The stratum indexed by $(d_1, \ldots, d_s)$ is the image of

$$X^{(d_1-d_2)} \times X^{(d_2-d_3)} \times \cdots \times X^{(d_s)} \longrightarrow \mathsf{Tor}^d,$$

where

$$(D_1, \ldots, D_s) \longmapsto \mathcal{O}_{D_1+\cdots+D_s} \oplus \mathcal{O}_{D_2+\cdots+D_s} \oplus \cdots \oplus \mathcal{O}_{D_s}.$$

The stratum $\mathsf{Tor}^{(d)}$ is open; it is the image of

$$X^{(d)} \longrightarrow \mathsf{Tor}^d, \qquad D \longmapsto \mathcal{O}_D.$$

For $X = \mathbb{A}^1$ this says

$$\mathsf{Tor}^{(d)} \simeq \mathfrak{gl}_d^{\mathrm{reg}}/\mathrm{GL}_d, \qquad \mathsf{Tor}^{(d)}|_{X^{(d)}-\Delta} \simeq \mathfrak{gl}_d^{\mathrm{rss}}/\mathrm{GL}_d, \qquad X^{(d)} \simeq \mathfrak{t}_d/W.$$

Thus the torsion stack is the curve-wise analogue of the adjoint quotient picture for $\mathfrak{gl}_d$.

4.5. **The Springer resolution of Tor^d.** Let $\widetilde{\mathsf{Tor}}^d$ be the stack of complete flags in a length-d torsion sheaf:

$$\widetilde{\mathsf{Tor}}^d = \{0 = M_0 \subset M_1 \subset \cdots \subset M_d = M \mid \mathrm{length}(M_i/M_{i-1}) = 1\}.$$

There are natural maps $\pi\colon \widetilde{\mathsf{Tor}}^d \to \mathsf{Tor}^d$ and $o_d\colon \widetilde{\mathsf{Tor}}^d \to X^d$, where π forgets the flag and o_d records the ordered supports of the quotients M_i/M_{i-1}. Write $s_d\colon \mathsf{Tor}^d \to X^{(d)}$ and $r_d\colon X^d \to X^{(d)}$ for the support map and the symmetrization map. If $q\colon \mathsf{Tor}^{(d)} \to X^{(d)}$ is the support map on the cyclic stratum, the diagram is

$$\begin{array}{ccc}
\widetilde{\mathsf{Tor}}^d & \xrightarrow{\ o_d\ } & X^d \\
\downarrow{\scriptstyle \pi} & & \downarrow{\scriptstyle r_d} \\
\mathsf{Tor}^d & \xrightarrow[\ s_d\ ]{} & X^{(d)} \\
\uparrow{\scriptstyle i^{(d)}} & \nearrow{\scriptstyle q} & \Vert \\
\mathsf{Tor}^{(d)} & \xleftarrow{\ D\mapsto \mathcal{O}_D\ } & X^{(d)}
\end{array}$$

On the regular semisimple locus $\mathsf{Tor}^{d,\mathrm{rss}}$ of multiplicity-free support, s_d exhibits $\mathsf{Tor}^{d,\mathrm{rss}}$ as the scalar-automorphism gerbe over $X^{(d),\mathrm{rss}}$, and $\widetilde{\mathsf{Tor}}^{d,\mathrm{rss}}$ is the ordered scalar-automorphism gerbe over $X^{d,\mathrm{rss}}$. After rigidifying these scalar automorphisms, π becomes the usual ordered S_d-cover $X^{d,\mathrm{rss}} \to X^{(d),\mathrm{rss}}$. This is parallel to the Grothendieck–Springer diagram [GPR20, AHJR14]

$$\mathfrak{c}_d \xrightarrow{\ s\ } \mathfrak{gl}_d^{\mathrm{reg}}/\mathrm{GL}_d \longrightarrow \mathfrak{c}_d, \qquad \widetilde{\mathfrak{gl}}_d/\mathrm{GL}_d \longrightarrow \mathfrak{gl}_d/\mathrm{GL}_d \longrightarrow \mathfrak{c}_d, \qquad \mathfrak{t}_d \longrightarrow \mathfrak{c}_d = \mathfrak{t}_d/W,$$

where s is the Kostant section. Locally on X, this analogy is literal: a torsion sheaf at a point is a nilpotent endomorphism, and a flag of torsion subsheaves is a Springer flag.

4.6. **The Springer–Laumon sheaf and Laumon's summand.** Define the Springer–Laumon sheaf in degree d by

$$\mathrm{Spr}_E^d := \pi_! o_d^*(E^{\boxtimes d})[d] \in \mathcal{D}_c^{\mathrm{b}}(\mathsf{Tor}^d).$$

The shift $[d]$ is the normalization used here for the equivariant-perverse convention. The S_d-action is the Springer monodromy action [AHJR14]: over the multiplicity-free support locus, π is the ordered-support S_d-cover after rigidification, and the resulting action extends to Spr_E^d by middle extension, equivalently by the usual simple-reflection correspondences. It is not a literal action of S_d on the whole flag stack by arbitrarily permuting a nested flag.

Theorem 4.6.1 (Laumon's extension theorem [KL85])**.** *Let $j_{\mathrm{rss}}\colon \mathsf{Tor}^{d,\mathrm{rss}} \hookrightarrow \mathsf{Tor}^d$ be the regular semisimple locus, let $s_{\mathrm{rss}}\colon \mathsf{Tor}^{d,\mathrm{rss}} \to X^{(d),\mathrm{rss}}$ be the support map, and let $r_{\mathrm{rss}}\colon X^{d,\mathrm{rss}} \to X^{(d),\mathrm{rss}}$ be the ordered cover. The sheaf Spr_E^d is perverse and is the middle extension from the regular semisimple locus:*

$$\mathrm{Spr}_E^d \simeq j_{\mathrm{rss},!*}(s_{\mathrm{rss}}^* r_{\mathrm{rss},*}(E^{\boxtimes d})[d]).$$

The larger cyclic regular locus $\mathsf{Tor}^{(d)}$ also determines the same object; its restriction there is obtained by intermediate extension across the diagonals inside $X^{(d)}$.

Laumon's sheaf is the trivial-isotypic summand for this S_d-action:

$$\mathcal{L}_E^d := \mathrm{Hom}_{S_d}(\mathrm{triv}, \mathrm{Spr}_E^d) = (\mathrm{Spr}_E^d)^{S_d} \in \mathrm{Perv}(\mathsf{Tor}^d).$$

Equivalently, if $E^{(d)} := (r_{\mathrm{rss},*} E^{\boxtimes d})^{S_d}$ on $X^{(d),\mathrm{rss}}$, then

$$\mathcal{L}_E^d \simeq j_{\mathrm{rss},!*}(s_{\mathrm{rss}}^* E^{(d)}[d]).$$

This local system has rank n^d on the multiplicity-free locus and should not be confused with $\mathrm{Sym}^d E$.

Lemma 4.6.2. *If E is geometrically irreducible, then $\mathcal{L}_E^d \in \mathrm{Perv}(\mathsf{Tor}^d)$ is geometrically irreducible.*

Sketch. After base change to $\bar{k}$, Schur's lemma gives $\mathrm{End}_{\pi_1(X_{\bar{k}})}(E_{\bar{x}}) = e$, where e is the coefficient field. The local system $E^{(d)}$ on the multiplicity-free configuration locus corresponds to $E_{\bar{x}}^{\otimes d}$ as a representation of the wreath-product group $\pi_1(X_{\bar{k}})^d \rtimes S_d$. Its endomorphism algebra is

$$\mathrm{End}_{\pi_1(X_{\bar{k}})^d \rtimes S_d}(E_{\bar{x}}^{\otimes d}) = (e^{\otimes d})^{S_d} = e,$$

so it is irreducible. Pullback to the scalar-automorphism gerbe preserves irreducibility, and middle extension preserves irreducibility of the underlying perverse sheaf. □

Finally, for $\mathcal{L}_E \in \mathrm{Perv}(\mathsf{Tor})$ define

$$\mathcal{L}_E^d := \mathcal{L}_E|_{\mathsf{Tor}^d}.$$

The following subsections analyze this object locally: after restricting to torsion sheaves supported at a single point, ordinary Springer theory decomposes the local perverse object $\mathcal{L}_{E,x}^d[-d]$ into IC-sheaves of nilpotent strata with coefficients the highest-weight representations of E_x.

4.7. **Reduction to local torsion at one point.** For the local calculation, write x for a geometric point of X. Let $i_x \colon \mathsf{Tor}_x^d \hookrightarrow \mathsf{Tor}^d$ be the fiber of s_d over $d \cdot x$. Thus Tor_x^d classifies length-d torsion sheaves supported at x. Put

$$\mathcal{L}_{E,x}^d := i_x^* \mathcal{L}_E^d.$$

Because this fiber has codimension d in the symmetric-power direction, the local perverse object is $\mathcal{L}_{E,x}^d[-d]$.

Lemma 4.7.1 (Disjoint support factorization)**.** *If $x \neq y$, then direct sum gives an isomorphism of local torsion stacks*

$$\mathsf{Tor}_x^a \times \mathsf{Tor}_y^b \xrightarrow{\sim} \mathsf{Tor}_{ax+by}^{a+b}, \qquad \mathcal{L}_{E,x}^a \boxtimes \mathcal{L}_{E,y}^b \xrightarrow{\sim} \mathcal{L}_{E,ax+by}^{a+b}.$$

More generally, for a divisor $D = \sum_x d_x x$ with distinct supports,

$$\prod_x \mathsf{Tor}_x^{d_x} \xrightarrow{\sim} \mathsf{Tor}_D^d, \qquad \boxtimes_x \mathcal{L}_{E,x}^{d_x} \simeq \mathcal{L}_{E,D}^d.$$

Hence the structure of $\mathcal{L}_E^d$ is reduced to the calculation of $\mathcal{L}_{E,x}^d[-d] \in \mathrm{Perv}(\mathsf{Tor}_x^d)$.

4.8. **Local stratification by partitions.** Let $\lambda = (\lambda_1 \geqslant \lambda_2 \geqslant \cdots \geqslant \lambda_s > 0)$ be a partition of d, and put $\lambda_{s+1} = 0$. We write $|\lambda| = d$ and $\ell(\lambda) = s$. For an effective divisor D, write $\mathcal{O}_D := \mathcal{O}_X/\mathcal{O}_X(-D)$. The global stratum $\mathsf{Tor}^\lambda \subset \mathsf{Tor}^d$ is described by the atlas

$$X^{(\lambda_1-\lambda_2)} \times X^{(\lambda_2-\lambda_3)} \times \cdots \times X^{(\lambda_s)} \longrightarrow \mathsf{Tor}^\lambda,$$

which sends $(D_1, \ldots, D_s)$ to

$$\mathcal{O}_{D_1+\cdots+D_s} \oplus \mathcal{O}_{D_2+\cdots+D_s} \oplus \cdots \oplus \mathcal{O}_{D_s}.$$

The stratum Tor^λ maps locally closedly into Tor^d.

The local stratum at x is the residual piece $\mathsf{Tor}_x^\lambda \subset \mathsf{Tor}_x^d$ corresponding to the torsion sheaf $T_{\lambda,x} := \bigoplus_{i=1}^s \mathcal{O}_X/\mathcal{O}_X(-\lambda_i x)$. Equivalently, after passing to the completed local ring at the geometric point x and choosing a uniformizer t, this is the $k[\![t]\!]$-module $\bigoplus_{i=1}^s k[\![t]\!]/(t^{\lambda_i})$. The stack Tor_x^d has dimension $-d$; under the nilpotent-cone model below, its strata are quotient stacks of nilpotent orbits.

Definition 4.8.1. For a partition $\lambda \vdash d$, define

$$B_{\lambda,x} := \mathrm{IC}_{\overline{\mathsf{Tor}_x^\lambda}}.$$

4.9. **Local decomposition of Laumon's sheaf.** Let $V = E_x$ be the fiber of E at x. For a partition $\lambda = (\lambda_1 \geqslant \cdots \geqslant \lambda_s > 0)$ with $s \leqslant n$, extend it to a dominant weight of $\mathrm{GL}(V) \simeq \mathrm{GL}_n$ by adding zeros:

$$\lambda = (\lambda_1, \ldots, \lambda_s, 0, \ldots, 0).$$

Let V^λ be the irreducible highest-weight representation of $\mathrm{GL}(V)$ with highest weight λ.

Equivalently, if $\chi^\lambda \in \mathrm{Irr}(S_d)$ is the Specht representation corresponding to λ, then Schur–Weyl duality gives

$$V^\lambda \simeq (V^{\otimes d} \otimes \chi^\lambda)^{S_d}.$$

If $s > n$, this vector space is zero.

Proposition 4.9.1 (Local formula)**.** *For every geometric point $x \in X$ and integer $d \geqslant 0$, put $n(\lambda) := \sum_{i\geqslant 1}(i-1)\lambda_i$. Via the usual Casselman–Shalika normalization view $n(\lambda) = \langle \lambda, (0, 1, \dots, n-1)\rangle$. Then*

$$\mathcal{L}^d_{E,x}[-d] \simeq \bigoplus_{\substack{\lambda \vdash d \\ \ell(\lambda) \leqslant n}} B_{\lambda,x} \otimes E^\lambda_x(-n(\lambda)).$$

4.10. **Relation with the affine Grassmannian.** Let $\mathsf{Gr}_{n,x}$ be the affine Grassmannian of GL_n at x. Let $\mathsf{Gr}^{d,\mathrm{mod}}_{n,x}$ be the degree-d positive modification locus, i.e. the locus whose points are quotients of $\mathcal{O}^n_x$ of length d; this notation should not be confused with the regular nilpotent locus. Let

$$f\colon \mathsf{Gr}^{d,\mathrm{mod}}_{n,x} \longrightarrow \mathsf{Tor}^d_x$$

be the map taking a modification to its quotient torsion sheaf. Its image is the open substack of torsion sheaves generated by at most n elements, and over this image f is smooth of relative dimension nd. For $|\lambda| = d$ with $\ell(\lambda) \leqslant n$, let $\mathsf{Gr}^\lambda_{n,x}$ be the Schubert stratum of type λ. The diagram is

$$\begin{array}{ccc} \mathsf{Gr}^\lambda_{n,x} & \hookrightarrow & \mathsf{Gr}^{d,\mathrm{mod}}_{n,x} \\ \downarrow & & \downarrow{\scriptstyle f}\ {\scriptstyle \text{smooth}} \\ \mathsf{Tor}^\lambda_x & \hookrightarrow & \mathsf{Tor}^d_x. \end{array}$$

Put

$$A_\lambda := \mathrm{IC}_{\overline{\mathsf{Gr}^\lambda_{n,x}}}.$$

With the convention that IC includes the perverse cohomological shift but no Tate twist, smooth pullback gives

$$f^* B_{\lambda,x} \simeq A_\lambda[-dn].$$

Equivalently, $f^*B_{\lambda,x}[dn] \simeq A_\lambda$. The Tate factor $-n(\lambda)$ in the local Laumon formula is part of the Whittaker/Casselman–Shalika normalization and is not part of the smooth-pullback comparison of IC sheaves.

4.11. **Proof of the local formula by Springer theory.** We now explain the proof of the local decomposition, using the classical Springer correspondence [GPR20, AHJR14]. Work étale-locally at x, or replace X by the formal disk at x. After choosing a uniformizer, a length-d torsion sheaf supported at x is a d-dimensional vector space with a nilpotent endomorphism. Therefore

$$\mathsf{Tor}^d_x \simeq \mathcal{N}_d/\mathrm{GL}_d,$$

where $\mathcal{N}_d \subset \mathfrak{gl}_d$ is the nilpotent cone. The non-local statement $\mathsf{Tor}^d \simeq \mathfrak{gl}_d/\mathrm{GL}_d$ is literal for $X = \mathbb{A}^1$; for a general curve it is the corresponding étale local model near a divisor of degree d. Under this local identification, the map $\widetilde{\mathsf{Tor}}^d \to \mathsf{Tor}^d$ becomes the Grothendieck–Springer resolution, and its restriction over Tor^d_x becomes the Springer resolution of the nilpotent cone.

Base change to a geometric point $\bar{x}$ and suppress Tate twists temporarily. The local perverse Springer sheaf decomposes as

$$\mathrm{Spr}^d_{\bar{x}}[-d] \simeq \bigoplus_{\chi \in \mathrm{Irr}(S_d)} \chi \otimes \mathrm{IC}_{\overline{\mathcal{O}_\chi}}.$$

Here $\mathrm{Irr}(S_d)$ is identified with partitions of d, and also with nilpotent orbits in $\mathfrak{gl}_d$. We use the following convention:

$$\lambda \longleftrightarrow \mathcal{O}_\lambda, \qquad \text{trivial representation} \longleftrightarrow \text{regular nilpotent orbit } (d).$$

This fixes the no-transpose convention.

For Laumon's sheaf, the fiber $E^{\otimes d}_x$ is carried along the Springer sheaf, and the S_d-action on this tensor factor is by permutation of the d factors. Hence

$$\mathrm{Spr}^d_{E,x}[-d] \simeq \bigoplus_{\chi \in \mathrm{Irr}(S_d)} \chi \otimes E^{\otimes d}_x \otimes \mathrm{IC}_{\overline{\mathcal{O}_\chi}}.$$

Taking the S_d-invariant summand gives

$$(\mathrm{Spr}^d_{E,x}[-d])^{S_d} \simeq \bigoplus_{\chi \in \mathrm{Irr}(S_d)} (\chi \otimes E^{\otimes d}_x)^{S_d} \otimes \mathrm{IC}_{\overline{\mathcal{O}_\chi}}.$$

Now write $\chi = \chi^\lambda$, where $d = \lambda_1 + \cdots + \lambda_s$ with $\lambda_1 \geqslant \cdots \geqslant \lambda_s > 0$. Extend λ by zeros to length n. By Schur–Weyl duality,

$$(\chi^\lambda \otimes E_x^{\otimes d})^{S_d} \simeq E_x^\lambda,$$

and this vanishes if $s > \dim(E_x) = n$. Finally,

$$\mathrm{IC}_{\overline{\mathcal{O}_\lambda}} = B_{\lambda,x}$$

under $\mathsf{Tor}_x^d \simeq \mathcal{N}_d/\mathrm{GL}_d$. Reinstating the Tate twist gives the formula of the proposition for $\mathcal{L}_{E,x}^d[-d]$.

5. Fourier transform

5.1. **Reminder: from Laumon's sheaf to a preliminary automorphic sheaf.** Keep $G = \mathrm{GL}_n$ and let E be a rank n local system on X. Section 4 constructed $\mathcal{L}_E \in \mathrm{Perv}(\mathsf{Tor})$. The aim here is to use it to build an automorphic sheaf Aut_E on Bun_n.

Let Bun_n' be the stack of pairs (M, s) where M is a rank n vector bundle and $s\colon \Omega_X^{n-1} \hookrightarrow M$ is injective. Let $\mathcal{Q}_n$ be the positive global Hecke–Whittaker stack

$$\mathcal{Q}_n = \left\{ \begin{array}{c} 0 = M_0 \subset M_1 \subset \cdots \subset M_n, \quad M_i/M_{i-1} \simeq \Omega_X^{n-i}, \\ M_n \hookrightarrow M \in C_n \text{ with torsion cokernel} \end{array} \right\}.$$

Equivalently, this is the compactified global version of the positive Hecke correspondence fibered with the ρ-twisted Whittaker stack; it is not a literal product of a local affine Grassmannian with Bun_N. It has three evident maps: $\pi\colon \mathcal{Q}_n \to \mathsf{Bun}_n'$ remembers $\Omega_X^{n-1} = M_1 \hookrightarrow M$, the map $p_1\colon \mathcal{Q}_n \to \mathsf{Tor}$ records the torsion quotient M/M_n, and $p_2\colon \mathcal{Q}_n \to \mathsf{Bun}_N \to \mathbb{G}_\mathrm{a}$ is the Whittaker character, i.e. the sum of the adjacent extension classes of the flag.

Let $\mathcal{L}_\psi$ be the Artin–Schreier sheaf on $\mathbb{G}_\mathrm{a}$. On the component where $\mathrm{length}(M/M_n) = d$, the preliminary object is

$$\mathrm{Aut}_E'^{,d} \coloneqq \pi_!(p_1^*\mathcal{L}_E^d \otimes p_2^*\mathcal{L}_\psi)[\dim_\mathrm{rel}](\dim_\mathrm{rel}/2).$$

The bracket denotes the relative shift and Tate twist used in the Fourier normalization. The exclamation mark is compactly supported pushforward; the map π is not proper in general.

The desired properties are as follows.

1. If E is irreducible, then $\mathrm{Aut}_E'^{,d}$ is perverse and irreducible for $d \gg 0$.
2. For $d \gg 0$, $\mathrm{Aut}_E'^{,d}$ descends along $\mathsf{Bun}_n' \to \mathsf{Bun}_n^d$, giving $\mathrm{Aut}_E^d \in \mathrm{Perv}(\mathsf{Bun}_n^d)$.
3. Aut_E satisfies the truncated top Hecke relation. With the Hecke-correspondence convention of Section 2,

$$t_x^*\mathrm{Aut}_E^d \simeq \det(E_x) \otimes \mathrm{Aut}_E^{d+n},$$

with $t_x\colon \mathsf{Bun}_n^{d+n} \xrightarrow{\sim} \mathsf{Bun}_n^d$, $M \mapsto M(-x)$. Equivalently, $(t_x)_!\mathrm{Aut}_E^{d+n} \simeq \det(E_x)^{-1} \otimes \mathrm{Aut}_E^d$.

The difficulty is that there is no formal reason for the compactly supported pushforward $\pi_!$ to preserve perversity or irreducibility. The point of the section is that, after rewriting the construction as Fourier–Deligne transform [KL85], this becomes accessible.

5.2. **Fourier transform for vector bundles.** Let $V \to S$ be a vector bundle of rank r, $V^\vee \to S$ its dual, and $\langle -, - \rangle\colon V \times_S V^\vee \to \mathbb{A}^1$ the pairing. The Fourier–Deligne transform [KL85] is

$$\mathbb{F}_{V/S}(K) \coloneqq p_!^\vee(p^*K \otimes \langle -, - \rangle^*\mathcal{L}_\psi)[r](r/2), \qquad K \in \mathcal{D}_\mathrm{c}^\mathrm{b}(V).$$

With this normalization it is t-exact for the perverse t-structure and preserves irreducibility on simple perverse sheaves. Thus the problem is to recognize the sheaf defining Aut_E' as the Fourier transform of a clean extension.

5.3. **The case $G = \mathrm{GL}_2$.** For $G = \mathrm{GL}_2$ write

$$\overline{\mathsf{Bun}}_2' = \{0 \to \Omega_X \to M \to \mathcal{L} \to 0, \quad \mathcal{L} \in \mathsf{Coh}_1\}$$

for the coherent extension stack. The stack Bun_2' used for automorphic sheaves is the open substack where the middle term M is locally free. Let Coh_1 be the stack of rank-one coherent sheaves, and let

$$\mathsf{Coh}_1^+ \coloneqq \{(\mathcal{L}, \mathcal{O}_X \xrightarrow{s} \mathcal{L})\}, \qquad \mathsf{Coh}_1' \subset \mathsf{Coh}_1^+$$

be the open substack where s is generically injective. On Coh_1' the cokernel of s is a torsion sheaf; let $q\colon \mathsf{Coh}_1' \to \mathsf{Tor}$ be this quotient map. Thus the sheaf used below is $q^*\mathcal{L}_E$, not a literal restriction of $\mathcal{L}_E$ to Coh_1'.

There is a Cartesian square

$$\begin{array}{ccc} \square = \left\{ \begin{matrix} 0 \to \Omega_X \to M' \to \mathcal{O}_X \to 0, \\ M' \to M \end{matrix} \right\} & \longrightarrow & \overline{\mathsf{Bun}}_2' \\ \big\downarrow & & \big\downarrow \\ \mathsf{Coh}_1' & \longrightarrow & \mathsf{Coh}_1, \end{array}$$

where, given an extension $0 \to \Omega_X \to M \to \mathcal{L} \to 0$ and $\mathcal{O}_X \to \mathcal{L}$, the bundle M' is the pullback. Over the open of Coh_1 where $H^1(X, \mathcal{L}) = \mathrm{Ext}^1(\mathcal{O}_X, \mathcal{L}) = 0$, the map $\mathsf{Coh}_1^+ \to \mathsf{Coh}_1$ is a vector bundle with fiber $\mathrm{Hom}(\mathcal{O}_X, \mathcal{L})$. The extension stack with fiber $\mathrm{Ext}^1(\mathcal{L}, \Omega_X)$ is the dual vector bundle, since Serre duality gives

$$\mathrm{Ext}^1(\mathcal{L}, \Omega_X) \simeq \mathrm{Hom}(\mathcal{O}_X, \mathcal{L})^\vee.$$

The open substack where the middle term is locally free is the corresponding part of Bun_2'. Thus one must work on this open substack of Coh_1, and one must use Coh_1^+, not Coh_1', because the zero section is part of the vector bundle.

The Whittaker map $\square \to \mathsf{Bun}_N \to \mathbb{G}_a$ is exactly the pairing

$$\mathrm{Hom}(\mathcal{O}_X, \mathcal{L}) \times \mathrm{Ext}^1(\mathcal{L}, \Omega_X) \longrightarrow \mathrm{Ext}^1(\mathcal{O}_X, \Omega_X) \simeq k.$$

Therefore, over this open in Coh_1,

$$\mathrm{Aut}_E' \simeq \mathbb{F}(j_! q^* \mathcal{L}_E), \qquad j \colon \mathsf{Coh}_1' \hookrightarrow \mathsf{Coh}_1^+.$$

If $j_! = j_{!*} = j_*$, then $j_! q^* \mathcal{L}_E$ is a simple perverse sheaf, hence its Fourier transform is perverse and irreducible. This proves the desired statement at least on the open of Bun_2' where the vector-bundle description holds.

5.4. **Drinfeld's clean-extension theorem for** GL_2**.** Restrict now to line bundles $\mathsf{Bun}_1 \subset \mathsf{Coh}_1$. Since a map between line bundles is either zero or injective, over Bun_1 one has

$$\mathsf{Bun}_1' \xrightarrow{\ j\ } \mathsf{Bun}_1^+ \xleftarrow{\ i\ } \mathsf{Bun}_1,$$

where i is the zero section and j is its open complement.

Theorem 5.4.1 (Drinfeld, clean extension; see [FGKV98, KL85])**.** *Assume E is geometrically irreducible of rank 2. The sheaf $j_! q^* \mathcal{L}_E^d$ is clean over $\mathsf{Bun}_1^d \subset \mathsf{Coh}_1$, provided $d > 4(g-1)$.*

Let $F = q^* \mathcal{L}_E^d$ on Bun_1'. Cleanness of $j_! F$ means that the natural map $j_! F \to j_* F$ is an isomorphism. Since the cone is supported on the zero section, this is equivalent to

$$i^* j_* F = 0.$$

By Verdier duality it suffices to prove the compactly supported boundary vanishing for F and for its dual. For the dilation action of $\mathbb{G}_m$ on the vector bundle $\mathsf{Bun}_1^+ \to \mathsf{Bun}_1$, the contraction principle gives the compactly supported form

$$i^! j_! F \simeq \pi_! F,$$

where $\pi \colon \mathsf{Bun}_1' \to \mathsf{Bun}_1$ is the projection. Modulo the stacky $\mathbb{G}_m$-factor, π is the Abel–Jacobi map

$$\mathrm{AJ}^d \colon X^{(d)} \longrightarrow \mathrm{Pic}^d.$$

Thus the boundary term vanishes by Deligne's theorem: for E geometrically irreducible of rank $n > 1$ (or nontrivial of rank 1),

$$d > n(2g-2) \implies \mathrm{AJ}^d_!(E^{(d)}) = 0.$$

For $n = 2$ this condition is precisely $d > 4(g-1)$.

5.5. **The naive iteration for** GL_n**.** For general n one wants to iterate the preceding Fourier transform. Let Coh_k be the stack of rank-k coherent sheaves. Define

$$\mathsf{Coh}_k^+ := \{(M_k, s\colon \Omega_X^{k-1} \to M_k)\}, \qquad \mathsf{Coh}_k' \subset \mathsf{Coh}_k^+$$

to be the open locus where s is generically injective. If s is generically injective, then $\mathrm{coker}(s)$ has rank $k-1$, so Coh_k' maps to Coh_{k-1}. The recursion starts with the pullback along the quotient map $q_1\colon \mathsf{Coh}_1' \to \mathsf{Tor}$,

$$\mathcal{F}_1 := q_1^* \mathcal{L}_E$$

and naively sets

$$\mathcal{F}_{k+1} := \mathbb{F}(j_{k,!}\mathcal{F}_k)|_{\mathsf{Coh}_{k+1}'}, \qquad j_k\colon \mathsf{Coh}_k' \hookrightarrow \mathsf{Coh}_k^+,$$

where Coh_{k+1}' is viewed as the open substack of the Serre-dual extension bundle over Coh_k. At the end, $\mathcal{F}_n$ is the sheaf whose restriction to Bun_n' is the Fourier description of Aut_E'.

One cannot simply restrict all intermediate Coh_k to Bun_k. If $C_k \subset \mathsf{Bun}_k$ for $k < n$, then the construction sees only the part of the compactified Whittaker stack where the intermediate quotients are vector bundles. This misses defect strata in which some $M_k \to M$ has torsion cokernel. Hence the correct opens $C_k \subset \mathsf{Coh}_k$ must be larger than Bun_k.

5.6. **The opens** C_k **and the two required conditions.** The construction postpones the final definition of C_k. The two conditions used here are the following. Let $M_k \in C_k$, and write M_k^{tf} for its torsion-free quotient.

First, the vector-bundle condition is

$$\mathrm{Ext}^1(\Omega_X^{k-1}, M_k) \simeq \mathrm{Ext}^1(\Omega_X^{k-1}, M_k^{\mathrm{tf}}) = 0.$$

It implies that $\mathsf{Coh}_k^+ \to \mathsf{Coh}_k$ is a vector bundle with fiber $\mathrm{Hom}(\Omega_X^{k-1}, M_k)$. By the same Serre-duality calculation it also gives $\mathrm{Hom}(M_k, \Omega_X^k) = 0$, so the dual extension stack has no automorphism directions. Its linear dual is the honest vector bundle of extension classes

$$\mathcal{E}_{k+1} \to C_k, \qquad 0 \to \Omega_X^k \to M_{k+1} \to M_k \to 0.$$

The open $C_{k+1}' \subset \mathcal{E}_{k+1}$ is the locus where the middle term lies in the chosen open C_{k+1} and the induced section $\Omega_X^k \to M_{k+1}$ is generically injective. The duality follows from Serre duality:

$$\mathrm{Ext}^1(M_k, \Omega_X^k) \simeq \mathrm{Hom}(\Omega_X^{k-1}, M_k)^\vee.$$

Second, the normalized degree must be large:

$$\mathring{\deg}\,(M_k^{\mathrm{tf}}) := \deg(M_k^{\mathrm{tf}}) - \deg(\mathcal{O}_X \oplus \Omega_X \oplus \cdots \oplus \Omega_X^{k-1}) > nk(2g-2).$$

From this point on, all degrees are meant to be normalized degrees $\mathring{\deg}$, even when the dot is omitted.

5.7. **Cleanness and averaging vanishing.** The expected clean-extension statement, at the stages where the next Fourier transform is taken, is:

Theorem 5.7.1 (Cleanness)**.** *For $1 \leqslant k < n$, the extension $j_{k,!}\mathcal{F}_k$ is clean at least over*

$$\mathsf{Bun}_k \cap C_k \subset \mathsf{Coh}_k.$$

If $\mathcal{F}_k$ is simple perverse on C_k', then this clean extension is again simple. With additional conditions in the definition of C_k, the same cleanness will hold over C_k.

The input for cleanness is Gaitsgory's vanishing theorem for averaging functors [Gai04, Gai16b]. Let Mod_k^d be the Hecke stack

$$\mathsf{Mod}_k^d = \{M \subset M',\ \mathrm{length}(M'/M) = d\},$$

with maps

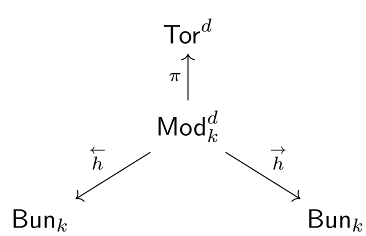

where $\pi(M \subset M') = M'/M$. The averaging functor is

$$\mathrm{Av}_E^d(-) := \overrightarrow{h}_!(\overleftarrow{h}^*(-) \otimes \pi^*\mathcal{L}_E^d)[dk](dk/2).$$

Theorem 5.7.2 (Vanishing)**.** *If E is geometrically irreducible of rank n, with $k < n$, and $d > nk(2g-2)$, then*

$$\mathrm{Av}_E^d = 0.$$

More generally, the vanishing theorem applies to averaging on Bun_k when $\mathrm{rank}(E) > k$, with bound $d > (2g-2)k \cdot \mathrm{rank}(E)$.

For $k = 1$, this is Deligne's vanishing for Abel–Jacobi. Section 6 uses Theorem 5.7.2 to prove the cleanness theorem.

6. Cleanness

6.1. **Fourier tower over the opens C_k.** This section develops the Fourier-transform recursion. For $1 \leqslant k \leqslant n$, let

$$\mathsf{Coh}_k := \{M_k \in \mathrm{Coh}(X) : \mathrm{rank}(M_k) = k\}.$$

There are two section stacks:

$$\mathsf{Coh}'_k := \{\Omega_X^{k-1} \to M_k \text{ generically injective}\}, \qquad \overline{\mathsf{Coh}}'_k := \{\Omega_X^{k-1} \to M_k\}.$$

The bar allows the section to be non-injective or zero. Fix opens $C_k \subset \mathsf{Coh}_k$ and write

$$C'_k := \mathsf{Coh}'_k \times_{\mathsf{Coh}_k} C_k, \qquad \overline{C}'_k := \overline{\mathsf{Coh}}'_k \times_{\mathsf{Coh}_k} C_k.$$

The opens are required to satisfy the following conditions.

(1) For $M_k \in C_k$, one has $\mathrm{Ext}^1(\Omega_X^{k-1}, M_k) = 0$. Hence $\overline{C}'_k \to C_k$ is a vector bundle with fiber $\mathrm{Hom}(\Omega_X^{k-1}, M_k)$.
(2) If M_k^{tf} is the torsion-free quotient, then

$$\mathring{\deg}\,(M_k^{\mathrm{tf}}) > nk(2g-2).$$

The proof uses this normalized-degree form.
(3) If $M_k \in C_k$ and $\Omega_X^{k-1} \hookrightarrow M_k$ is generically injective, then $M_k/\Omega_X^{k-1} \in C_{k-1}$.

By Serre duality, the linear dual to $\overline{C}'_k \to C_k$ has fiber

$$\mathrm{Ext}^1(M_k, \Omega_X^k).$$

Denote this dual vector bundle by $\mathcal{E}_{k+1} \to C_k$; its points are extension classes

$$0 \to \Omega_X^k \to M_{k+1} \to M_k \to 0.$$

This is why the next Fourier transform moves from level k to level $k+1$.

Define $\mathcal{F}_k$ by induction. First

$$\mathcal{F}_1 := q_1^*\mathcal{L}_E[\text{shift}](\text{twist}),$$

where $q_1 \colon C'_1 \to \mathsf{Tor}$ takes the quotient of $\mathcal{O}_X = \Omega_X^0 \hookrightarrow M_1$. If $j_k \colon C'_k \hookrightarrow \overline{C}'_k$ is the open embedding, then

$$\mathcal{F}_{k+1} := \mathbb{F}_{\overline{C}'_k/C_k}(j_{k,!}\mathcal{F}_k)|_{\mathcal{E}_{k+1}^{\neq 0}}.$$

Here $\mathcal{E}_{k+1}^{\neq 0}$ is the open where the middle term lies in C_{k+1} and the tautological section $\Omega_X^k \to M_{k+1}$ is generically injective. When no confusion is possible, we regard this restriction as a sheaf on C'_{k+1} via the forgetful map from extension classes to middle terms. The functor $\mathbb{F}_{\overline{C}'_k/C_k}$ is the Fourier–Deligne transform between the dual vector bundles over C_k.

On $\mathsf{Bun}'_n \cap C'_n$, the result is the preliminary automorphic sheaf:

$$\mathcal{F}_n|_{\mathsf{Bun}'_n \cap C'_n} = \mathrm{Aut}'_E|_{\mathsf{Bun}'_n \cap C'_n}.$$

More generally, $\mathcal{F}_n$ is computed from the compactified Whittaker stack

$$\mathcal{Q}_n = \left\{ \begin{array}{c} 0 = M_0 \subset M_1 \subset \cdots \subset M_n, \quad M_i/M_{i-1} \simeq \Omega_X^{n-i}, \\ M_n \hookrightarrow M \in C_n \text{ with torsion cokernel} \end{array} \right\},$$

which replaces the non-compactified Whittaker stack. The target is allowed to lie in Coh_n, not only in Bun_n, because torsion defects must be retained.

6.2. **The cleanness theorem.** The goal, for the levels needed to form the next Fourier transform, is:

Theorem 6.2.1 (Cleanness)**.** *For $1 \leqslant k < n$, the extension of $\mathcal{F}_k$ across the non-injective locus is clean, i.e. on $\overline{C}'_k$ we have*

$$j_{k,!}\mathcal{F}_k \simeq j_{k,*}\mathcal{F}_k.$$

Thus the Fourier induction preserves the intended irreducible perverse sheaves on the chosen opens.

Filter by torsion length. Put

$$\mathsf{Coh}_k^{\mathrm{tor},\leqslant \ell} := \{M_k \in \mathsf{Coh}_k : \mathrm{length}(\mathrm{Tor}(M_k)) \leqslant \ell\}, \qquad C_k^{\mathrm{tor},\leqslant \ell} := C_k \cap \mathsf{Coh}_k^{\mathrm{tor},\leqslant \ell}.$$

Then $\mathsf{Coh}_k^{\mathrm{tor}=0} = \mathsf{Bun}_k$ and $C_k = \bigcup_{\ell \geqslant 0} C_k^{\mathrm{tor},\leqslant \ell}$. We prove cleanness over $\overline{C}'^{\,\mathrm{tor},\leqslant \ell}_k$ by induction on ℓ.

6.3. **Base case: torsion-free sheaves.** Let $\ell = 0$. A map from a line bundle to a vector bundle is either zero or injective. Hence the complement of $C'^{\,\mathrm{tor}=0}_k$ inside $\overline{C}'^{\,\mathrm{tor}=0}_k$ is the zero section

$$i\colon C_k^{\mathrm{tor}=0} \hookrightarrow \overline{C}'^{\,\mathrm{tor}=0}_k.$$

Let $q\colon \overline{C}'^{\,\mathrm{tor}=0}_k \to C_k^{\mathrm{tor}=0}$ be the vector-bundle projection. The sheaf is monodromic for the dilation action on the fibers. As in the GL_2 case, cleanness is equivalent to the boundary vanishing $i^* j_{k,*}\mathcal{F}_k = 0$; by Verdier duality it suffices to prove the compactly supported version for $\mathcal{F}_k$ and for its dual. The contraction principle gives

$$i^! j_{k,!}\mathcal{F}_k \simeq q_! j_{k,!}\mathcal{F}_k.$$

Thus it remains to show $q_! j_{k,!}\mathcal{F}_k = 0$; the same argument applied to the dual gives the $*$-boundary vanishing.

Using the Whittaker formula for $\mathcal{F}_k$, this object is

$$q_! j_{k,!}\mathcal{F}_k \simeq (r \circ \pi)_!(p_1^*\mathcal{L}_E \otimes p_2^*\mathcal{L}_\psi)[\mathrm{shift}](\mathrm{twist})|_{C_k^{\mathrm{tor}=0}},$$

where p_1 records the torsion quotient and p_2 is the Whittaker character. Over the torsion-free zero-section boundary, the compactified Whittaker stack restricts to the ordinary Whittaker stack

$$\mathsf{Bun}_{N_k} \xrightarrow{\ \pi\ } \mathsf{Bun}'_k \xrightarrow{\ r\ } \mathsf{Bun}_k.$$

Let W_k be the base Whittaker sheaf, i.e. the direct image of $\exp_\psi$ from Bun_{N_k} to Bun_k with the same shift and twist. On the normalized degree-d component,

$$q_! j_{k,!}\mathcal{F}_k \simeq \mathrm{Av}_E^d(W_k)|_{C_k^{\mathrm{tor}=0}}.$$

The normalized degree convention places Bun_{N_k} over the degree-zero component of Bun_k. Condition (2) gives

$$C_k^{\mathrm{tor}=0} \subset \bigcup_{d > nk(2g-2)} \mathsf{Bun}_k^d.$$

Since $k < n$ in the induction levels under consideration, Theorem 5.7.2 says $\mathrm{Av}_E^d = 0$ for $d > nk(2g-2)$. Therefore $q_! j_{k,!}\mathcal{F}_k = 0$, proving the base case.

6.4. **Induction step and the base change.** Assume cleanness is known over $C_k^{\mathrm{tor},\leqslant \ell-1}$. To avoid confusing the target sheaf with the endpoint of a Whittaker flag, write the target of the section as M. The new boundary stratum is

$$Z := \left\{ \Omega_X^{k-1} \to M \;\middle|\; \text{not injective, } \mathrm{length}(\mathrm{Tor}(M)) = \ell \right\}.$$

Equivalently, in $\overline{C}'^{\,\mathrm{tor},\leqslant \ell}_k$ the open already handled is

$$U := C'^{\,\mathrm{tor},\leqslant \ell}_k \cup \overline{C}'^{\,\mathrm{tor},\leqslant \ell-1}_k,$$

and Z is the remaining closed piece.

For $\ell > 0$ there is no canonical projection from the locus of torsion $\leqslant \ell$ to the locus of torsion ℓ. Put $Y_\ell := \overline{C}'^{\,\mathrm{tor},\leqslant \ell}_k$ and add such a presentation by base change. Define

$$\widetilde{Y} := \left\{ \begin{array}{c} \Omega_X^{k-1} \to M, \quad 0 \to \mathring{M} \to M \to \mathcal{T} \to 0, \\ \mathring{M} \in \mathsf{Bun}_k,\ \mathcal{T} \in \mathsf{Tor}^\ell,\ M \in C_k,\ \mathring{M} \oplus \mathcal{T} \in C_k \end{array} \right\}.$$

The map $\widetilde{Y} \to Y_\ell$ forgetting the presentation is smooth over the torsion-ℓ stratum. Without the section, this is the usual presentation of the torsion-ℓ stratum of Coh_k by modifications of vector bundles; the relative tangent is $\mathrm{Hom}(\mathring{M}, \mathcal{T})$, so the relative dimension is locally constant.

The fixed locus is

$$\widetilde{Z} := \{\Omega_X^{k-1} \to \mathcal{T},\ \mathring{M} \in \mathsf{Bun}_k,\ \mathring{M} \oplus \mathcal{T} \in C_k\}.$$

It maps to $\widetilde{Y}$ by the split extension $M = \mathring{M} \oplus \mathcal{T}$. Notice that $\widetilde{Z} \neq Z$: in $\widetilde{Z}$ the splitting/presentation is part of the data. There is a projection

$$\tilde{f} \colon \widetilde{Y} \longrightarrow \widetilde{Z}, \qquad (\Omega_X^{k-1} \to M,\ 0 \to \mathring{M} \to M \to \mathcal{T} \to 0) \longmapsto (\Omega_X^{k-1} \to M \to \mathcal{T}, \mathring{M}).$$

Scalar multiplication on the extension class in $\operatorname{Ext}^1(\mathcal{T}, \mathring{M})$, together with scalar multiplication on the $\mathring{M}$-component of the lifted section, defines a $\mathbb{G}_{\mathrm{m}}$-action on $\widetilde{Y}$ and extends to an $\mathbb{A}^1$-contraction onto $\widetilde{Z}$. Scaling only the extension class would leave an unwanted vector-bundle component of the section at the fixed point. Therefore the contraction principle reduces the induction step to

$$\tilde{f}_! \tilde{j}_! (\mathcal{F}_k|_{\widetilde{U}}) = 0, \qquad \widetilde{U} := U \times_{\overline{C}'_k} \widetilde{Y}.$$

6.5. **Fiber calculation and the surjective case.** Use the compactified Whittaker formula

$$\mathcal{F}_k \simeq \pi_!(p_1^* \mathcal{L}_E \otimes p_2^* \mathcal{L}_\psi)[\text{shift}](\text{twist})|_{C'_k},$$

where p_1 records the torsion quotient and p_2 is the Whittaker character. After base change, fix a point $(s_T \colon \Omega_X^{k-1} \to \mathcal{T}, \mathring{M}) \in \widetilde{Z}$. The relevant fiber consists of Whittaker flags and a map $M_k \to M$ over the fixed extension

$$0 \to \mathring{M} \to M \to \mathcal{T} \to 0,$$

such that the composite $\Omega_X^{k-1} = M_1 \to M_k \to M \to \mathcal{T}$ is the fixed map s_T. The induced map $M_k \to \mathcal{T}$ need not be surjective.

If $M_k \to \mathcal{T}$ is surjective, then, with $\mathring{M}_k := \ker(M_k \to \mathcal{T})$, one has the exact diagram

$$\begin{array}{ccccccccc}
0 & \longrightarrow & \mathring{M}_k & \longrightarrow & M_k & \twoheadrightarrow & \mathcal{T} & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \| & & \\
0 & \longrightarrow & \mathring{M} & \longrightarrow & M & \twoheadrightarrow & \mathcal{T} & \longrightarrow & 0,
\end{array}$$

and the fiber is equivalent to the ordinary modification stack

$$\{\mathring{M}_k \to \mathring{M} \mid \mathring{M}/\mathring{M}_k \in \mathsf{Tor}^d\}.$$

The required !-push is therefore a stalk of Av_E^d. In normalized degrees, the endpoint M_k of the Whittaker flag has $\mathring{\deg}\, M_k = 0$, hence $\mathring{\deg}\, \mathring{M}_k = -\ell$. Since $\mathring{M} \oplus \mathcal{T} \in C_k$, the condition (2) gives $\mathring{\deg}\, \mathring{M} > nk(2g-2)$. Thus

$$d = \mathring{\deg}\, \mathring{M} - \mathring{\deg}\, \mathring{M}_k > nk(2g-2) + \ell.$$

So $d > nk(2g-2)$, and the vanishing theorem gives $\mathrm{Av}_E^d = 0$. This handles the surjective case.

6.6. **Reduction of the non-surjective case.** Suppose $M_k \to \mathcal{T}$ is not surjective, and let $\mathcal{T}' \subsetneq \mathcal{T}$ be its image. Put $\mathring{M}_k := \ker(M_k \to \mathcal{T}')$. Let $M' \subset M$ be the inverse image of $\mathcal{T}'$ under $M \to \mathcal{T}$. Then the data factor through a diagram

$$\begin{array}{ccccccccc}
0 & \longrightarrow & \mathring{M}_k & \longrightarrow & M_k & \twoheadrightarrow & \mathcal{T}' & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \| & & \\
0 & \longrightarrow & \mathring{M} & \longrightarrow & M' & \twoheadrightarrow & \mathcal{T}' & \longrightarrow & 0 \\
 & & \| & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & \mathring{M} & \longrightarrow & M & \twoheadrightarrow & \mathcal{T} & \longrightarrow & 0.
\end{array}$$

Consequently

$$0 \to M'/M_k \to M/M_k \to \mathcal{T}/\mathcal{T}' \to 0.$$

Let $H_{\mathcal{T}}$ and $H_{\mathcal{T}'}$ be the corresponding modification fibers. The map $H_{\mathcal{T}} \to H_{\mathcal{T}'}$ is compatible with the torsion diagram

$$\begin{array}{ccccc}
 & & \mathsf{Tor}^{\mathrm{ses}} & & \\
 & \overset{m}{\swarrow} & & \overset{(s,q)}{\searrow} & \\
\mathsf{Tor} & & & & \mathsf{Tor} \times \mathsf{Tor},
\end{array}$$

where $\mathsf{Tor}^{\mathrm{ses}}$ classifies short exact sequences of torsion sheaves; m records the middle term, while (s,q) records subobject and quotient. In the present situation, these are M/M_k, M'/M_k, and $\mathcal{T}/\mathcal{T}'$.

After pulling back to this diagram, $H_{\mathcal{T}} \to H_{\mathcal{T}'}$ is a fibration in affine spaces. The affine directions are controlled by extension data such as $\mathrm{Ext}^1(\mathcal{T}/\mathcal{T}', \mathring{M}_k)$ together with choices of the compatible lift $M_k \to M'$ over the quotient $\mathcal{T}'$. Thus, up to the usual affine-bundle shift and Tate twist, the remaining term is reduced by base change to a pushforward along the residual projection pr,

$$\mathrm{pr}_!((s,q)_! m^* \mathcal{L}_E|_{H_{\mathcal{T}'}}) = 0.$$

The needed proposition is the factorization identity, in the quotient-stack normalization used for Tor,

$$(s,q)_! m^* \mathcal{L}_E \simeq \mathcal{L}_E \boxtimes \mathcal{L}_E,$$

which is proved in Section 7. It splits the torsion quotient into the part M'/M_k and the residual part $\mathcal{T}/\mathcal{T}'$, reducing the non-surjective case to the surjective case. This completes the proof of Cleanness modulo that proposition.

7. Laumon's factorization

7.1. **The remaining lemma from cleanness.** This section completes the last step of the proof that vanishing implies cleanness. Fix $M_k \in \mathsf{Bun}_k^0$, $J \in \mathsf{Tor}^\ell$, $\mathring{M} \in \mathsf{Bun}_k$, $a\colon M_k \to J$, where Bun_k^0 denotes the normalized degree-zero component occurring as the endpoint of the Whittaker flag. Assume that $\mathring{M} \oplus J \in C_k$. Consider the stack

$$H := \left\{ \begin{array}{c} 0 \to \mathring{M} \to M \to J \to 0, \quad M \in C_k, \\ u\colon M_k \hookrightarrow M, \quad (M_k \xrightarrow{u} M \twoheadrightarrow J) = a \end{array} \right\}.$$

It has a map $q_H\colon H \to \mathsf{Tor}$, $(M,u) \mapsto M/M_k$. The lemma needed in the cleanness proof is the compactly supported vanishing

$$R\Gamma_{\mathrm{c}}(H, q_H^* \mathcal{L}_E) \simeq 0.$$

There is one technical strengthening of the good open C_k. We may require that, for every extension $0 \to \mathring{M} \to M \to J \to 0$,

$$\mathring{M} \oplus J \in C_k \quad \Longrightarrow \quad M \in C_k.$$

One way to force this implication is to choose a line bundle L^{est} and impose the additional open condition

$$\mathrm{Hom}(M, L^{\mathrm{est}}) = 0.$$

Indeed $\mathrm{Hom}(J, L^{\mathrm{est}}) = 0$ for torsion J and torsion-free target L^{est}. Hence $\mathring{M} \oplus J \in C_k$ is equivalent to $\mathring{M} \in C_k$, and applying $\mathrm{Hom}(-, L^{\mathrm{est}})$ to the extension gives $M \in C_k$. Thus the condition $M \in C_k$ can be ignored in the fiber calculation.

7.2. **The surjective case.** First suppose that $a\colon M_k \twoheadrightarrow J$. Put $\mathring{M}_k := \ker(a)$. Then an object of H is the same as a diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \mathring{M}_k & \longrightarrow & M_k & \longrightarrow\!\!\!\!\rightarrow & J & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \| & & \\ 0 & \longrightarrow & \mathring{M} & \longrightarrow & M & \longrightarrow\!\!\!\!\rightarrow & J & \longrightarrow & 0. \end{array}$$

Equivalently, it is a modification $\mathring{M}_k \hookrightarrow \mathring{M}$. Thus

$$H \simeq \{\mathring{M}_k\} \times_{\mathsf{Bun}_k} \mathsf{Mod}_k^d \times_{\mathsf{Bun}_k} \{\mathring{M}\}, \qquad \mathsf{Bun}_k \leftarrow \mathsf{Mod}_k^d \to \mathsf{Bun}_k.$$

Under this identification q_H records $\mathring{M}/\mathring{M}_k$. Up to the harmless normalizing one-dimensional line appearing in the definition of Av_E,

$$R\Gamma_{\mathrm{c}}(H, q_H^* \mathcal{L}_E) \simeq (\mathrm{Av}_E^d(\delta_{\mathring{M}_k}))_{\mathring{M}} \otimes (\text{line}), \qquad d = \mathring{\deg}\, \mathring{M} - \mathring{\deg}\, \mathring{M}_k.$$

In the normalized component used in the cleanness proof, $\mathring{\deg}\, M_k = 0$, so $\mathring{\deg}\, \mathring{M}_k = -\ell$. The condition $\mathring{M} \oplus J \in C_k$ gives $\mathring{\deg}\, \mathring{M} > nk(2g-2)$. Hence

$$d = \mathring{\deg}\, \mathring{M} - \mathring{\deg}\, \mathring{M}_k > nk(2g-2) + \ell > nk(2g-2).$$

The vanishing theorem for Av_E^d therefore gives $R\Gamma_{\mathrm{c}}(H, q_H^* \mathcal{L}_E) = 0$.

7.3. **Reduction to the image of $M_k \to J$.** For a general map $a\colon M_k \to J$, let $J' \subset J$ be the image and put $J'' \coloneqq J/J'$. Let $\mathring{M}_k = \ker(M_k \twoheadrightarrow J')$. For an object of H, set $M' \coloneqq M \times_J J'$. Then we have

$$\begin{array}{ccccccccc}
0 & \longrightarrow & \mathring{M}_k & \longrightarrow & M_k & \twoheadrightarrow & J' & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \| & & \\
0 & \longrightarrow & \mathring{M} & \longrightarrow & M' & \twoheadrightarrow & J' & \longrightarrow & 0 \\
 & & \| & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & \mathring{M} & \longrightarrow & M & \twoheadrightarrow & J & \longrightarrow & 0.
\end{array}$$

Let H' be the stack of the first two rows. It is the surjective version with target J'.

Let $\mathsf{Tor}^{\mathrm{ses}}$ be the stack of short exact sequences of torsion sheaves

$$0 \to J_s \to J_m \to J_q \to 0.$$

Write

$$m(J_\bullet) = J_m, \qquad s(J_\bullet) = J_s, \qquad q(J_\bullet) = J_q.$$

The quotient sequence

$$0 \to M'/M_k \to M/M_k \to M/M' \to 0$$

defines a map $H \to \mathsf{Tor}^{\mathrm{ses}}$. Here

$$M'/M_k \simeq \mathring{M}/\mathring{M}_k, \qquad M/M' \simeq J/J'.$$

The induced diagram

$$\begin{array}{ccc}
H & \longrightarrow & H' \\
\downarrow & & \downarrow \\
\mathsf{Tor}^{\mathrm{ses}} & \xrightarrow{(s,q)} & \mathsf{Tor} \times \mathsf{Tor} \\
{\scriptstyle m}\downarrow & & \\
\mathsf{Tor} & &
\end{array}$$

is not Cartesian. Nevertheless base-change holds for the sheaves under consideration. The reason is that

$$H \longrightarrow \mathsf{Tor}^{\mathrm{ses}} \times_{\mathsf{Tor}\times\mathsf{Tor}} H'$$

is an affine bundle. At a point

$$0 \to J_s \to J_m \to J_q \to 0, \qquad 0 \to M_k \to M' \to J_s \to 0,$$

the fiber is the affine space

$$\mathrm{Ext}^1(J_m, M_k) \times_{\mathrm{Ext}^1(J_s, M_k)} \{[0 \to M_k \to M' \to J_s \to 0]\}.$$

Its dimension is constant on each connected component; indeed it is modeled on $\ker(\mathrm{Ext}^1(J_m, M_k) \to \mathrm{Ext}^1(J_s, M_k)) \simeq \mathrm{Ext}^1(J_q, M_k)$, whose dimension is k length(J_q). Hence, up to the constant affine-bundle shift and Tate twist, it remains to prove

$$R\Gamma_c((s,q)_! m^* \mathcal{L}_E|_{H'}) = 0.$$

By the surjective case, this will follow from the factorization identity below.

7.4. **Factorization of Laumon's sheaf through short exact sequences.**

Proposition 7.4.1 (Laumon factorization [KL85])**.** *For any local system E for which the Laumon complexes are defined, in particular for the geometrically irreducible E used in the cleanness argument, one has on* $\mathsf{Tor} \times \mathsf{Tor}$

$$(s,q)_! m^* \mathcal{L}_E \simeq \mathcal{L}_E \boxtimes \mathcal{L}_E.$$

Equivalently, on the component of $\mathsf{Tor}^{\mathrm{ses}}$ *where the subobject has length d_1 and the quotient has length d_2, over* $\mathsf{Tor}^{d_1} \times \mathsf{Tor}^{d_2}$,

$$(s,q)_! m^* \mathcal{L}_E^{d_1+d_2} \simeq \mathcal{L}_E^{d_1} \boxtimes \mathcal{L}_E^{d_2},$$

with shifts and Tate twists already included in the normalization of $\mathcal{L}_E$.

Remark 7.4.2. The surrounding cleanness argument assumes geometric irreducibility of E, but the factorization identity itself is a structural property of the Springer–Laumon construction and does not use this irreducibility.

Proof. Recall the regular semisimple locus

$$\mathsf{Tor}^{d,\mathrm{rss}} = \left\{ \bigoplus_{i=1}^{d} k(x_i) \;\middle|\; x_i \neq x_j \text{ for } i \neq j \right\} \subset \mathsf{Tor}^d,$$

which is a gerbe over $X^{(d),\mathrm{rss}} = X^{(d)} - \Delta$ whose stabilizer at a geometric point is $(\mathbb{G}_\mathrm{m})^d$. If $r_\mathrm{rss}\colon X^{d,\mathrm{rss}} \to X^{(d),\mathrm{rss}}$ is the ordered-support cover and $s_\mathrm{rss}\colon \mathsf{Tor}^{d,\mathrm{rss}} \to X^{(d),\mathrm{rss}}$ is the support map, then

$$E^{(d)} \coloneqq (r_{\mathrm{rss},*} E^{\boxtimes d})^{S_d}, \qquad \mathcal{L}^d_E \simeq j_{\mathrm{rss},!*}\left(s^*_\mathrm{rss} E^{(d)}[d]\right).$$

Let

$$U_{d_1,d_2} \coloneqq (\mathsf{Tor}^{d_1,\mathrm{rss}} \times \mathsf{Tor}^{d_2,\mathrm{rss}})_\mathrm{disj}$$

be the open locus where the two supports are disjoint, and let $s_i\colon \mathsf{Tor}^{d_i,\mathrm{rss}} \to X^{(d_i),\mathrm{rss}}$ be the two support maps. Over this locus every short exact sequence of the two torsion sheaves splits uniquely, and the middle-term map m is identified with addition of divisors. Hence

$$(s,q)_! m^* \mathcal{L}^d_E|_{U_{d_1,d_2}} \simeq \mathrm{add}^*\, s^*_\mathrm{rss} E^{(d)}[d]|_{U_{d_1,d_2}},$$

while

$$(\mathcal{L}^{d_1}_E \boxtimes \mathcal{L}^{d_2}_E)|_{U_{d_1,d_2}} \simeq (s_1^* E^{(d_1)}[d_1] \boxtimes s_2^* E^{(d_2)}[d_2])|_{U_{d_1,d_2}}.$$

These two local systems are canonically identified by symmetric descent. Moreover the right-hand side is the IC-extension of its restriction to U_{d_1,d_2}: the product local system already extends smoothly across the cross-diagonals where the two supports meet. Thus it is enough to show that

$$(s,q)_! m^* \mathcal{L}^d_E \simeq \mathrm{IC}((s,q)_! m^* \mathcal{L}^d_E|_{U_{d_1,d_2}}).$$

Since $\mathcal{L}^d_E = (\mathrm{Spr}^d_E)^{S_d}$ is a direct summand of the Springer–Laumon sheaf, it suffices to prove the corresponding statement for Spr^d_E, equivariantly for $S_{d_1} \times S_{d_2} \subset S_d$, and then take the invariant summand. The assertion is étale-local on X. After such a localization the local system E is constant with fiber V; the coefficient factor $V^{\otimes d}$ is carried through the construction. Thus the geometric input is the constant-coefficient Springer statement.

For $X = \mathbb{A}^1$, a length-d torsion sheaf is a d-dimensional vector space with an endomorphism, hence

$$\mathsf{Tor}^d \simeq \mathfrak{gl}_d/\mathrm{GL}_d.$$

For $d = d_1 + d_2$, the stack $\mathsf{Tor}^{\mathrm{ses},d_1,d_2}$ is the parabolic correspondence

$$\mathfrak{g}/G \longleftarrow \mathfrak{p}/P \longrightarrow \mathfrak{m}/M,$$

where $G = \mathrm{GL}_d$, $P = P_{d_1,d_2} = UM$, $M \simeq \mathrm{GL}_{d_1} \times \mathrm{GL}_{d_2}$, and $\mathfrak{p}, \mathfrak{m}$ are the corresponding Lie algebras. The required geometric input is the standard parabolic-restriction theorem for the Grothendieck–Springer sheaf [AHJR14, GPR20]: the complex

$$(\mathfrak{p}/P \to \mathfrak{m}/M)_!(\mathfrak{p}/P \to \mathfrak{g}/G)^* \mathrm{Spr}_\mathfrak{g}$$

is the IC-extension from the locus in $\mathfrak{m}/M$ where the two characteristic polynomials are regular and disjoint, with the Weyl local system obtained by restricting the S_d-monodromy to $S_{d_1} \times S_{d_2}$.

For reference, a Bruhat proof stratifies

$$\widetilde{\mathfrak{p}} \coloneqq \{(x, gB) \mid x \in \mathfrak{p} \cap \mathrm{Ad}_g\, \mathfrak{b}\}$$

by the double-coset pieces PwB/B, with $w \in W_M \backslash W_G$. For a minimal coset representative w, the map to the Grothendieck–Springer resolution for $\mathfrak{m}$ is an affine fibration whose fiber is

$$\mathfrak{u}_w \coloneqq \mathfrak{u} \cap \mathrm{Ad}_w\, \mathfrak{b},$$

not the whole $\mathfrak{u}$; the dimension depends on w. With the usual perverse normalization these strata assemble to the parabolic-restriction IC-extension above, and no summand supported on the complement of the regular-disjoint locus appears.

Taking the S_d-invariant summand with coefficients $V^{\otimes d}$ identifies this IC-extension with $\mathcal{L}^{d_1}_E \boxtimes \mathcal{L}^{d_2}_E$, because over the regular-disjoint locus it is exactly the symmetric-descent identification already described. This proves the Laumon factorization proposition. ∎

7.5. **Completion of the lemma.** Apply the proposition to the residual expression obtained above:

$$(s,q)_! m^* \mathcal{L}_E|_{H'} \simeq (\mathcal{L}_E \boxtimes \mathcal{L}_E)|_{H'}.$$

The first factor is exactly the surjective case for $M_k \twoheadrightarrow J'$, and the second factor is the fixed residual quotient J/J'. The !-push along H' therefore vanishes by the surjective case. By the affine-bundle base-change above,

$$R\Gamma_c(H, q_H^* \mathcal{L}_E) = 0.$$

This proves the remaining lemma and completes the missing step in the proof of cleanness.

8. Construction of Aut_E

8.1. **What has already been constructed.** Keep $G = \mathrm{GL}_n$, and let E be a geometrically irreducible rank-n local system on X. Sections 5–7 constructed a preliminary Whittaker–Laumon sheaf

$$\mathrm{Aut}'_E \in \mathcal{D}(C'_n)$$

on the good open $C'_n \subset \mathsf{Coh}'_n$. It is obtained from the compactified Whittaker stack

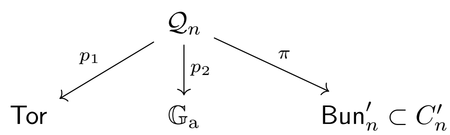

by the formula

$$\mathrm{Aut}'_E = \pi_!(p_1^* \mathcal{L}_E \otimes p_2^* \mathcal{L}_\psi)[\mathrm{shift}](\mathrm{twist}).$$

The already proved theorem is that, on C'_n, this object is perverse and irreducible.

Let

$$C_n^\circ := C_n \cap \mathsf{Bun}_n, \qquad C_n^{\prime,\circ} := C'_n \cap \mathsf{Bun}'_n,$$

and let

$$p_C : C_n^{\prime,\circ} \longrightarrow C_n^\circ$$

be the forgetful map from a bundle with a generically nonzero section $\Omega_X^{n-1} \to M$ to the bundle M. The descent statement to be proved is: if $C_n \cap \mathsf{Bun}_n^d \neq \varnothing$, then

$$\mathrm{Aut}'_E|_{C_n^{\prime,\circ} \cap (\mathsf{Bun}_n^d)'}$$

descends along p_C to a sheaf

$$\mathrm{Aut}_E|_{C_n \cap \mathsf{Bun}_n^d} \in \mathrm{Perv}(C_n \cap \mathsf{Bun}_n^d).$$

Since p_C is smooth on the chosen good open, a nonzero descended sheaf is automatically perverse and irreducible. The final goal is to extend these local pieces uniquely to a Hecke eigensheaf Aut_E on all of Bun_n.

8.2. **Descent along the projective-bundle fibers.** Work on a smooth open chart $V \subset C_n^\circ \cap \mathsf{Bun}_n^d$. After shrinking V if necessary, the space of sections $\Omega_X^{n-1} \to M$ is a vector bundle over V, and after quotienting by the scalar action on the section one gets a projective-bundle-type map

$$p_V : V' \longrightarrow V.$$

The descent picture is that $\mathrm{Aut}'_E|_{V'}$ is first written as the IC-extension of a local system on a dense open of V', and then one proves that this IC-extension is actually a local system on all of V'.

Two conditions are used. First, for every geometric point $m \in V$, the intersection of the good open with the fiber is dense:

$$V' \cap p_V^{-1}(m) \subset p_V^{-1}(m) \quad \text{is open and dense.}$$

This can be included as part of the definition of the good open. Second, $\mathrm{Aut}'_E|_{V'}$ must be a local system, rather than merely an IC-sheaf. Once this is known, the local system descends uniquely along p_V, because the geometric fibers are projective spaces and

$$\pi_1(\mathbb{P}^m_{\bar{k}}) = 1.$$

Equivalently, the pullback functor from local systems on V to local systems on V' is an equivalence. This gives the desired descent from $C_n^{\prime,\circ}$ to C_n°.

The criterion used to prove that an irreducible perverse sheaf is a local system is a finite-field trace test in the pure IC situation at hand. Let Y be smooth and let $K \in \mathrm{Perv}(Y)$ be irreducible. Write

$$\chi_K(y) := \sum_i (-1)^i \dim H^i(K_{\bar{y}})$$

for the Euler characteristic of the stalk. For arbitrary perverse sheaves, this function alone is not a local-system criterion. Here K is an irreducible IC-extension produced by the construction; in this setting, constancy and nonvanishing of χ_K rule out boundary contribution, hence force K to be a shifted local system on Y. Therefore it is enough to prove:

$$\chi(\mathrm{Aut}'_E) \text{ is constant on the fibers of } \mathsf{Bun}'_n \to \mathsf{Bun}_n, \qquad \chi(\mathrm{Aut}'_E) \neq 0 \text{ over } C_n^\circ.$$

8.3. **Constancy of the trace along fibers.** The constancy of $\chi(\mathrm{Aut}'_E)$ is reduced to a case where an automorphic function is already known to exist. The black-box input is Deligne's proper-pushforward invariance: if $f\colon Y_1 \to Y_2$ is proper and $K_1, K_2 \in \mathcal{D}(Y_1)$ are étale-locally isomorphic, then the Euler characteristic functions of $f_! K_1$ and $f_! K_2$ agree.

For any two rank-n local systems E_1, E_2, even reducible ones, the local systems are étale locally isomorphic on X. Hence the corresponding Laumon sheaves are étale locally isomorphic,

$$\mathcal{L}_{E_1} \simeq_{\mathrm{et,loc}} \mathcal{L}_{E_2}.$$

If the map defining Aut'_E were proper, Deligne's theorem would immediately give

$$\chi(\mathrm{Aut}'_{E_1}) = \chi(\mathrm{Aut}'_{E_2}).$$

The map is not literally proper, but the compactified Whittaker construction admits a Drinfeld compactification, so the same argument applies after replacing the correspondence by its compactified version.

Thus, to prove the constancy proposition for a given E, it is enough to prove it for a convenient rank-n local system E'. Choose E' so that there exists a cuspidal Hecke eigenfunction with Hecke eigenvalue E'. By the comparison theorem of Frenkel–Gaitsgory–Kazhdan–Vilonen, its restriction to $\mathsf{Bun}'_n(\mathbb{F}_q)$ is the trace function

$$M' \longmapsto \mathrm{Tr}(\mathrm{Frob}_{M'}, \mathrm{Aut}'_{E'}).$$

Since this function is pulled back from $\mathsf{Bun}_n(\mathbb{F}_q)$, it is constant on each fiber of $\mathsf{Bun}'_n \to \mathsf{Bun}_n$. By the preceding reduction, the same constancy holds for Aut'_E. This proves the descent theorem.

8.4. **The Hecke–Laumon correspondence.** The usual Hecke property is deduced from a stronger-looking property on coherent sheaves. For $k \geqslant 0$ and $d \geqslant 0$, define the Hecke–Laumon stack

$$\mathsf{HL}_k^d := \{0 \to M' \to M \to T \to 0 \mid M, M' \in \mathsf{Coh}_k, T \in \mathsf{Tor}^d\}.$$

It has maps

$$\mathsf{Coh}_k \xleftarrow{\overleftarrow{h}} \mathsf{HL}_k^d \xrightarrow{\overrightarrow{h}} \mathsf{Tor}^d \times \mathsf{Coh}_k,$$

where $\overleftarrow{h}$ remembers M and $\overrightarrow{h}$ remembers (T, M'). The corresponding functor is

$$\mathsf{HL}_k^d(K) := \overrightarrow{h}_! \overleftarrow{h}^* K[dk](dk/2), \qquad K \in \mathcal{D}(\mathsf{Coh}_k).$$

A perverse sheaf $K \in \mathrm{Perv}(\mathsf{Coh}_k)$ is called a Hecke–Laumon eigensheaf with eigenvalue E if there are isomorphisms

$$\alpha_d\colon \mathsf{HL}_k^d(K) \xrightarrow{\sim} \mathcal{L}_E^d \boxtimes K \qquad (d \geqslant 0),$$

with the evident normalization for $d = 0$, and these isomorphisms are compatible with iterated modifications. More explicitly, for $d = d_1 + d_2$, one compares the two ways of passing from M to M'' through

$$0 \to M' \to M \to T' \to 0, \qquad 0 \to M'' \to M' \to T'' \to 0.$$

The stack of such two-step modifications maps both to HL_k^d and to $\mathsf{Tor}^{d_1} \times \mathsf{HL}_k^{d_2}$. The required compatibility says that the diagram obtained from

$$\mathsf{HL}_k^{d_2} \circ \mathsf{HL}_k^{d_1} \simeq ((s,q)_! m^* \boxtimes \mathrm{Id}) \circ \mathsf{HL}_k^d$$

commutes with the factorization isomorphism

$$(s,q)_! m^* \mathcal{L}_E^d \simeq \mathcal{L}_E^{d_1} \boxtimes \mathcal{L}_E^{d_2}$$

from Section 7. Here m, s, q are the maps from the short-exact-sequence stack of torsion sheaves

$$0 \to T'' \to T \to T' \to 0.$$

For $k = 0$, this is precisely the Laumon factorization theorem: $\mathcal{L}_E \in \mathrm{Perv}(\mathsf{Tor})$ is a Hecke–Laumon eigensheaf. The Fourier induction from the preceding sections then gives, by induction on k, that

$$\mathcal{F}_k \in \mathrm{Perv}(C'_k)$$

is a Hecke–Laumon eigensheaf. In particular, $\mathrm{Aut}'_E|_{C_n'^{,\circ}}$ satisfies the Hecke–Laumon eigenproperty. After descent, the sheaf $\mathrm{Aut}_E|_{C_n^\circ}$ inherits the same property.

8.5. **From Hecke–Laumon to usual Hecke.** The remaining question is why the Hecke–Laumon property implies the usual Hecke eigenproperty on Bun_n. The statement used here is:

Proposition 8.5.1. *Let $S \in \mathrm{Perv}(C_n^\circ)$ be a Hecke–Laumon eigensheaf with eigenvalue E. Suppose also that its Verdier dual $\mathbb{D}S$ is a Hecke–Laumon eigensheaf with eigenvalue E^*. Then S satisfies the usual Hecke eigenproperty with eigenvalue E.*

It is enough to prove the first Hecke relation. Consider HL_n^1. Removing the split short exact sequences gives the open substack

$$\mathsf{HL}_n^{1,\circ} := \{0 \to M' \to M \to T \to 0 \mid \mathrm{length}(T) = 1,\ \text{the sequence is not split}\}.$$

After base change along $X \to \mathsf{Tor}^1$, which sends x to the length-one torsion sheaf supported at x, this stack is the usual first Hecke correspondence:

$$\mathsf{HL}_n^{1,\circ} \times_{\mathsf{Tor}^1} X = \{(x, 0 \to M' \to M) \mid x \in X, \mathrm{Supp}(M/M') = x, \mathrm{length}(M/M') = 1\}.$$

Equivalently,

$$M' \subset M \subset M'(x), \qquad \mathrm{length}(M/M') = 1.$$

The residual scalar automorphism of the length-one torsion quotient is a $\mathbb{G}_\mathrm{m}$-factor; after quotienting by it one obtains the usual Hecke stack Hecke_n^1.

The perverse estimate needed for this quotient is supplied by the following elementary lemma.

Lemma 8.5.2. *Let $\mathcal{E} \xrightarrow{\pi} B$ be a vector bundle, and let $\mathbb{P}(\mathcal{E}) \xrightarrow{\bar{\pi}} B$ be its projectivization. Suppose $K \in \mathrm{Perv}(\mathcal{E})$ is $\mathbb{G}_\mathrm{m}$-equivariant and that $\pi_! K$, $\pi_* K \in \mathrm{Perv}(B)$. Let $\bar{K}$ be the descent of $K|_{\mathcal{E}\setminus 0}$ to $\mathbb{P}(\mathcal{E})$. Then $\bar{\pi}_! \bar{K}$ is perverse.*

Apply Lemma 8.5.2 fiberwise to the quotient from $\mathsf{HL}_n^{1,\circ}$ to the usual Hecke stack. The Hecke–Laumon eigenproperty from Subsection 8.4 gives the required perversity for $\pi_! K$, and the same property for $\mathbb{D}S$ gives it for $\pi_* K$. Since

$$\mathsf{Tor}^1 \simeq X \times \mathbb{B}\mathbb{G}_\mathrm{m}$$

and $\mathcal{L}_E^1|_X = E$, the first usual Hecke transform of S is perverse and has eigenvalue E. The higher minuscule Hecke relations follow from the compatibility of the first Hecke correspondence under convolution, with eigenvalues $\wedge^i E$. Thus S is a usual Hecke eigensheaf.

For $S = \mathrm{Aut}_E|_{C_n^\circ}$, the hypotheses hold by the Hecke–Laumon induction and Verdier duality: $\mathbb{D}(\mathrm{Aut}_E)$ has eigenvalue E^*. Hence $\mathrm{Aut}_E|_{C_n^\circ}$ satisfies the usual Hecke eigenproperty. In particular, with the Hecke-correspondence convention of Section 2,

$$(t_x^{d+n})^* \mathrm{Aut}_E^d \simeq \det(E_x) \otimes \mathrm{Aut}_E^{d+n}$$

where the isomorphism $t_x^{d+n} \colon \mathsf{Bun}_n^{d+n} \to \mathsf{Bun}_n^d,\ M \mapsto M(-x)$. Equivalently, for $\tau_x^d \colon \mathsf{Bun}_n^d \to \mathsf{Bun}_n^{d+n}$, $M \mapsto M(x)$, one has

$$(\tau_x^d)^* \mathrm{Aut}_E^{d+n} \simeq \det(E_x)^{-1} \otimes \mathrm{Aut}_E^d.$$

8.6. **Extension to all of Bun_n and cuspidality.** Let $j_d\colon C_n \cap \mathsf{Bun}_n^d \hookrightarrow \mathsf{Bun}_n^d$ be the good open in the degree-d component. The construction first extends the descended sheaf by middle extension

$$\mathrm{Aut}_E^d \coloneqq j_{d,!*}(\mathrm{Aut}_E|_{C_n \cap \mathsf{Bun}_n^d}).$$

In the clean range this is the same object as $j_{d,!}$ and $j_{d,*}$. The top Hecke relation above then gives the unique extension between degree components, because $M \mapsto M(-x)$ changes the degree by $-n$. Compatibility between Hecke_n^1 and the top Hecke correspondence propagates the first Hecke relation from the good open to the extended sheaf. This gives a Hecke eigensheaf

$$\mathrm{Aut}_E \in \mathcal{D}(\mathsf{Bun}_n)$$

with eigenvalue E.

Finally, Aut_E is cuspidal. The reason is the same vanishing theorem used in the clean-extension argument: every proper constant-term functor for GL_n factors through an averaging operation on some Bun_k with $k < n$, and the averaging theorem for the geometrically irreducible rank-n local system E forces that contribution to vanish. Thus all proper constant terms of Aut_E vanish.

Over a finite field, the sheaf-function comparison with FGKV says the following. If a cuspidal Hecke eigenfunction on $\mathsf{Bun}_n(\mathbb{F}_q)$ with eigenvalue E is given, then its restriction to $\mathsf{Bun}_n'(\mathbb{F}_q)$ is

$$M' \longmapsto \mathrm{Tr}(\mathrm{Frob}_{M'}, \mathrm{Aut}_E').$$

Thus the sheaf Aut_E constructed above geometrizes the expected automorphic eigenfunction, matching the comparison with Whittaker functions in [FGKV98].

9. Vanishing conjectures

This section explains the averaging vanishing theorem used in the clean-extension argument of Section 6 and the cuspidality argument of Subsection 8.6 [Gai04, Gai16b]. The notation is the same as before: X is a smooth projective curve of genus g, and $\mathcal{D}(\boldsymbol{-})$ denotes the chosen sheaf theory. In the de Rham discussion below, take $k = \mathbb{C}$ and $\mathcal{D}(\boldsymbol{-}) = D\text{-mod}(\boldsymbol{-})$.

9.1. **The averaging functor.** Let E be an irreducible local system of rank r and let $m < r$. For $d = d_2 - d_1 \geqslant 0$ consider upper modifications of rank-m bundles,

$$\begin{array}{ccccc} \mathsf{Bun}_m^{d_1} & \xleftarrow{\overleftarrow{h}} & \mathrm{Mod}_m^{d_1,d_2} & \xrightarrow{\overrightarrow{h}} & \mathsf{Bun}_m^{d_2} \\ & & \big\downarrow{\scriptstyle \pi} & & \\ & & \mathsf{Tor}^d & & \end{array}$$

where a point of $\mathrm{Mod}_m^{d_1,d_2}$ is $M \subset M'$ with $\mathrm{length}(M'/M) = d$, and $\pi(M \subset M') = M'/M$. Let $\mathcal{L}_E^d \in \mathcal{D}(\mathsf{Tor}^d)$ be Laumon's sheaf: on the preimage of the multiplicity-free locus $X^{(d),\circ} \subset X^{(d)}$ under the support map $\mathsf{Tor}^d \to X^{(d)}$ it is the symmetric descent $E^{(d)}$ of $E^{\boxtimes d}$, and globally it is its Goresky–MacPherson extension to Tor^d. Up to the usual normalizing shift and Tate twist,

$$\mathrm{Av}_E^d(\mathcal{F}) \coloneqq \overrightarrow{h}_!(\overleftarrow{h}^*\mathcal{F} \otimes \pi^*\mathcal{L}_E^d) \in \mathcal{D}(\mathsf{Bun}_m^{d_2}).$$

Equivalently, after summing over connected components, this is an endofunctor of $\mathcal{D}(\mathsf{Bun}_m)$. We call this "upper modification by $\mathcal{L}_E^d$"; the same averaging functor is reviewed from the spectral side in [Gai16b].

Theorem 9.1.1 (Vanishing theorem)**.** *Assume E is irreducible,* $\mathrm{rank}(E) = r > m$*, and, for simplicity, $g > 1$. Then, whenever $d > rm(2g-2)$, we have*

$$\mathrm{Av}_E^d = 0.$$

The same formulation works in the usual sheaf-theoretic settings, with the standard normalizations of shifts and twists.

Theorem 9.1.1 is the theorem that was used earlier with $r = n$ and $m = k < n$. The point here is not the construction of $\mathcal{L}_E^d$, but why the functor should vanish from the spectral side of geometric Langlands.

9.2. **Spectral action heuristic.** In the de Rham form of geometric Langlands one expects, schematically, the spectral action and nilpotent-support formulation of [GR24a, AG15]:

$$D\text{-mod}(\mathsf{Bun}_G)^{\text{temp}} \simeq \text{QCoh}(\text{LocSys}_{\check{G}}),$$
$$D\text{-mod}(\mathsf{Bun}_G) \simeq \text{IndCoh}_{\text{Nilp}}(\text{LocSys}_{\check{G}}).$$

Restricting the right hand side to compact objects, this is often heuristically reflected by coherent sheaves on the spectral stack. Thus $\text{QCoh}(\text{LocSys}_{\check{G}})$ should act on $D\text{-mod}(\mathsf{Bun}_G)$. This action is not arbitrary. For $E \in \text{LocSys}_{\check{G}}$, the fiber functor

$$(-)_E\colon\ \text{QCoh}(\text{LocSys}_{\check{G}}) \to \mathsf{Vect}$$

should cut out the E-Hecke category by base change:

$$D\text{-mod}(\mathsf{Bun}_G) \otimes_{\text{QCoh}(\text{LocSys}_{\check{G}})} \mathsf{Vect} \simeq D\text{-mod}(\mathsf{Bun}_G)_{E\text{-Hecke}}.$$

Here the right side consists of objects whose Hecke transforms are obtained by tensoring with the local system E.

At a fixed point $x \in X$, geometric Satake gives a functor

$$\text{Rep}(\check{G}) \otimes \mathcal{D}(\mathsf{Bun}_G) \longrightarrow \mathcal{D}(\mathsf{Bun}_G), \qquad (V, \mathcal{K}) \longmapsto \mathcal{H}_{V,x}(\mathcal{K}),$$

where $\text{Rep}(\check{G})$ is the Satake heart inside the spherical category $\mathcal{D}(G(\mathcal{O}_x)\backslash G(\mathcal{K}_x)/G(\mathcal{O}_x))$. If the point moves, one gets

$$\text{Rep}(\check{G}) \otimes \mathcal{D}(\mathsf{Bun}_G) \longrightarrow \mathcal{D}(\mathsf{Bun}_G \times X).$$

An E-Hecke eigensheaf is an object $\mathcal{K}$ equipped with isomorphisms

$$\mathcal{H}_V(\mathcal{K}) \simeq \mathcal{K} \boxtimes V_E$$

for all $V \in \text{Rep}(\check{G})$, plus the higher compatibilities over X^I.

The naive functor $\text{Rep}(\check{G}) \otimes \mathcal{D}(X) \to \text{End}(\mathcal{D}(\mathsf{Bun}_G))$ is a functor but not a monoidal functor in the required sense: when two points collide, convolution is governed by the monoidal structure on geometric Satake, not by the external tensor product on X^2. For example, one must have

- $\mathcal{H}_{V,x} \circ \mathcal{H}_{W,y} = \mathcal{H}_{V,W;x,y}$ for $x \neq y$, and
- $\mathcal{H}_{V,W;x,x} = \mathcal{H}_{V\otimes W,x}$.

This is exactly a factorization condition, so the right parameter space must be the Ran space.

9.3. **The Ran action and local-to-global.** Let

$$\text{Ran}(X) = \{I \subset X \mid 0 < |I| < \infty\}.$$

It is understood as the colimit of the powers X^I along surjections. The factorization category $\text{Rep}(\check{G})_{\text{Ran}}$ has fiber $\bigotimes_{x\in I} \text{Rep}(\check{G})$ over a finite set $I \subset X$, with fusion when points collide. The Hecke functors assemble to a monoidal action

$$\text{Rep}(\check{G})_{\text{Ran}} \longrightarrow \text{End}(D\text{-mod}(\mathsf{Bun}_G)).$$

There is also a monoidal "localization" functor

$$\text{Loc}_{\check{G}}\colon\ \text{Rep}(\check{G})_{\text{Ran}} \longrightarrow \text{QCoh}(\text{LocSys}_{\check{G}}),$$

which sends a collection $(V_x)_{x\in I}$ to the evaluation vector bundle $\bigotimes_{x\in I} V_{\mathcal{E},x}$ on $\text{LocSys}_{\check{G}}$.

The following result refers to [Gai10].

Theorem 9.3.1 (Gaitsgory–Lurie, local-to-global)**.** *The functor* $\text{Loc}_{\check{G}}$ *realizes* $\text{QCoh}(\text{LocSys}_{\check{G}})$ *as a Verdier quotient of* $\text{Rep}(\check{G})_{\text{Ran}}$. *Equivalently, the right adjoint of* $\text{Loc}_{\check{G}}$ *is fully faithful.*

The proof is formal from the description

$$\text{LocSys}_{\check{G}} = \{\text{horizontal sections of } \mathbb{B}\check{G} \times X \to X\},$$

and the same argument works for any stack $Z \to X$ equipped with a connection relative to X. Thus the spectral action has no further choices once the Ran Hecke action is known.

Theorem 9.3.2 (Generalized vanishing theorem [Gai10, Gai04])**.** *The monoidal action*

$$\mathrm{Rep}(\check{G})_{\mathrm{Ran}} \longrightarrow \mathrm{End}(D\text{-mod}(\mathsf{Bun}_G))$$

uniquely factors through $\mathrm{QCoh}(\mathrm{LocSys}_{\check{G}})$*:*

$$\begin{array}{ccc} \mathrm{Rep}(\check{G})_{\mathrm{Ran}} & \longrightarrow & \mathrm{End}(D\text{-mod}(\mathsf{Bun}_G)) \\ & \searrow^{\mathrm{Loc}_{\check{G}}} & \uparrow \\ & & \mathrm{QCoh}(\mathrm{LocSys}_{\check{G}}) \end{array}$$

This is the categorical form of the generalized vanishing conjecture. Its proof uses Beilinson–Drinfeld localization for Kac–Moody representations.

9.4. **Recovering the functor** Av_E^d**.** Take $G = \mathrm{GL}_m$, so $\check{G} = \mathrm{GL}_m$, and let St_m be the standard representation. The object of $\mathrm{Rep}(\mathrm{GL}_m)_{\mathrm{Ran}}$ corresponding to the averaging kernel is

$$\mathcal{A}_{E,d} := (\mathrm{Sym}^d(\mathrm{St}_m \otimes E))_{X^{(d)}}.$$

It is supported on the natural map $X^{(d)} \to \mathrm{Ran}(X)$. Up to the same normalization as in the geometric Satake dictionary, the Ran Hecke action sends $\mathcal{A}_{E,d}$ to Av_E^d.

The reason is visible on the local Satake fiber. Over a point $x \in X$, the Cauchy decomposition gives

$$\mathrm{Sym}^d(\mathrm{St}_m \otimes E_x) \simeq \bigoplus_{\substack{\mu \vdash d \\ \ell(\mu) \leqslant \min(m,r)}} \mathbb{S}_\mu(\mathrm{St}_m) \otimes \mathbb{S}_\mu(E_x).$$

Under geometric Satake, $\mathbb{S}_\mu(\mathrm{St}_m)$ corresponds to the IC sheaf of the Schubert stratum indexed by μ, while $\mathbb{S}_\mu(E_x) = E_x^\mu$. Equivalently, the right-hand side is the invariant space $(\mathrm{St}_m^{\otimes d} \otimes E_x^{\otimes d})^{S_d}$, but not an additional S_μ-invariant of each summand. This is the Satake meaning of the Laumon sheaf $\mathcal{L}_E^d$ appearing in the modification kernel.

By Theorem 9.3.2, Av_E^d is obtained from $\mathrm{Loc}_{\mathrm{GL}_m}(\mathcal{A}_{E,d})$. If L is a local system on X, write $L^{(d)}$ for its symmetric descent to $X^{(d)}$. The fiber of this quasi-coherent sheaf at a rank-m local system $E' \in \mathrm{LocSys}_{\mathrm{GL}_m}$ is, up to dual conventions, $R\Gamma(X^{(d)}, (E' \otimes E)^{(d)})$. Hence it is enough to prove the following elementary vanishing.

Lemma 9.4.1. *Let E be irreducible of rank $r > m$, and let E' be any rank-m local system. If $d > rm(2g-2)$, then*

$$R\Gamma(X^{(d)}, (E' \otimes E)^{(d)}) = 0.$$

Proof. Set $L = E' \otimes E$. Since E is irreducible and $\mathrm{rank}(E) > \mathrm{rank}(E')$, any map $(E')^\vee \to E$ is zero; hence $H^0(X, L) = 0$. Applying the same argument to $E^\vee$ gives $H^0(X, L^\vee) = 0$, and Verdier duality gives $H^2(X, L) = 0$. Therefore $R\Gamma(X, L)$ is concentrated in degree 1, and

$$\dim H^1(X, L) = (2g-2)\,\mathrm{rank}(L) = rm(2g-2).$$

And Künneth gives

$$R\Gamma(X^{(d)}, L^{(d)}) \simeq \mathrm{Sym}^d\, R\Gamma(X, L) \simeq \wedge^d H^1(X, L)[-d].$$

The last exterior power vanishes for $d > \dim H^1(X, L)$. □

Thus the spectral action sends $\mathcal{A}_{E,d}$ to zero, and this gives the vanishing theorem for Av_E^d in the de Rham setting.

9.5. **Other sheaf theories and general reductive groups.** For sheaf theories other than de Rham D-modules, the spectral stack in the preceding argument is replaced by a restricted stack of local systems [AGK+20a], denoted informally by $\mathrm{LocSys}_{\check{G}}^{\mathrm{restr}}$. One can describe it by the functor of points

$$\mathrm{LocSys}_{\check{G}}^{\mathrm{restr}}(S) = \left\{ \mathrm{Rep}(\check{G}) \xrightarrow{\otimes} \mathrm{QCoh}(S) \otimes \mathrm{QLisse}(X) \right\},$$

while in the de Rham case one has the parallel expression with $D\text{-mod}(X)$ in place of $\mathrm{QLisse}(X)$. The known factorization statement in the Betti/constructible setting lands in a nilpotent-singular-support category, schematically

$$\mathrm{Rep}(\check{G})_{\mathrm{Ran}}^{\mathrm{QLisse}} \longrightarrow \mathrm{QCoh}(\mathrm{LocSys}_{\check{G}}^{\mathrm{restr}}) \longrightarrow \mathrm{End}(D_{\mathrm{Nilp}}(\mathsf{Bun}_G)),$$

so it does not by itself imply the full vanishing theorem on all of $\mathcal{D}(\mathsf{Bun}_G)$.

The expected general reductive statement is the same. Let $V \in \mathrm{Rep}(\check{G})$ and let E be an irreducible vector local system with $\mathrm{rank}(E) > \dim V$. The object $\mathrm{Sym}^d(V \otimes E)$ in the Ran representation category should act by zero on $\mathcal{D}(\mathsf{Bun}_G)$ for

$$d > (2g-2)\dim(V \otimes E).$$

In de Rham theory this follows from the generalized vanishing theorem by the same cohomological calculation as above. In the ℓ-adic setting the corresponding statement is still treated as a conjectural input in this formulation.

9.6. **Gaitsgory's original proof of vanishing.** This section also sketches Gaitsgory's original proof of the vanishing theorem [Gai04]. The argument is geometric, but it mirrors the analytic statement that the relevant Rankin–Selberg L-function is a polynomial of bounded degree.

First, define lower modifications and the functor $\mathrm{Av}_{E^*}^{-d}$. Up to the standard shifts, $\mathrm{Av}_{E^*}^{-d}$ is both the left and the right adjoint of Av_E^d, and Verdier duality gives

$$\mathbb{D} \circ \mathrm{Av}_{E^*}^{-d} \simeq \mathrm{Av}_E^d \circ \mathbb{D}.$$

To prove $\mathrm{Av}_E^d = 0$, it is enough to prove

$$\mathrm{Av}_E^d \circ \mathrm{Av}_{E^*}^{-d} = 0.$$

Indeed, if $G := \mathrm{Av}_{E^*}^{-d}$ is a left adjoint of $F := \mathrm{Av}_E^d$ and $FG = 0$, then $\mathrm{Hom}(G\mathcal{K}, G\mathcal{K}) \simeq \mathrm{Hom}(\mathcal{K}, FG\mathcal{K}) = 0$, so $G = 0$; since G is also a right adjoint of F, this forces $F = 0$. For d in the vanishing range, an induction lemma says that $\mathrm{Av}_E^d(\mathcal{F})$ and $\mathrm{Av}_{E^*}^{-d}(\mathcal{F})$ are cuspidal. Hence the desired composition vanishes once Av_E^d kills cuspidal sheaves.

Second, one proves the exactness needed for the Euler-characteristic reduction. More precisely, Av_E^d is first shown to be exact after passing to the auxiliary Verdier quotient described below; since its image is cuspidal and the quotient is conservative on cuspidal objects, this gives the required exactness for the vanishing argument. Then a standard Euler-characteristic argument reduces vanishing to

$$\chi(\mathrm{Av}_E^d(\mathcal{F})) = 0.$$

For Euler characteristics, the local system E can be replaced by another local system of the same rank; one reduces to the trivial local system, although the exactness input itself uses the irreducibility of E.

For the trivial rank-r system $\mathbb{1}^{\oplus r}$, Braverman's calculation gives

$$\mathrm{Av}_{\mathbb{1}^{\oplus r}}^d \simeq \bigoplus_{d_1+\cdots+d_r=d} \mathrm{Av}_{\mathbb{1}}^{d_1} \circ \cdots \circ \mathrm{Av}_{\mathbb{1}}^{d_r}.$$

Therefore it suffices to show

$$\mathrm{Av}_{\mathbb{1}}^d(\mathcal{F}) = 0 \qquad \text{for cuspidal } \mathcal{F} \text{ and } d > m(2g-2).$$

This last step is a direct calculation. On functions it is the Jacquet–Godement statement that the standard L-function of a cuspidal automorphic representation of GL_m is a polynomial of degree at most $m(2g-2)$; Braverman's calculation geometrizes this.

9.7. **Why the averaging functor has cuspidal image.** Let $m = m_1 + m_2$, and let

$$\mathrm{CT}\colon \mathcal{D}(\mathsf{Bun}_m) \longrightarrow \mathcal{D}(\mathsf{Bun}_{m_1} \times \mathsf{Bun}_{m_2})$$

be the constant-term functor for the standard parabolic with Levi $\mathrm{GL}_{m_1} \times \mathrm{GL}_{m_2}$. There is a filtration on $\mathrm{CT} \circ \mathrm{Av}_E^d$ whose associated graded pieces are

$$\bigoplus_{d_1+d_2=d} (\mathrm{Av}_E^{d_1} \boxtimes \mathrm{Av}_E^{d_2}) \circ \mathrm{CT}.$$

This is a purely geometric filtration of modification stacks. Since $\mathrm{rank}(E) > m_i$ for both Levi factors, the vanishing theorem for smaller ranks kills every graded piece by induction whenever $d > rm(2g-2)$: for each $d = d_1 + d_2$, at least one d_i exceeds $rm_i(2g-2)$. Hence

$$\mathrm{CT} \circ \mathrm{Av}_E^d = 0,$$

so the image of Av_E^d is cuspidal.

This also gives a quick proof of a fact used above. If $\mathcal{F} \in \mathcal{D}(\mathsf{Bun}_m)$ is an E-Hecke eigensheaf, then

$$\mathrm{Av}_E^d(\mathcal{F}) \simeq \mathcal{F} \otimes R\Gamma(X^{(d)}, (E^* \otimes E)^{(d)}),$$

as predicted by the spectral action. Applying CT gives

$$\mathrm{CT}(\mathrm{Av}_E^d(\mathcal{F})) \simeq \mathrm{CT}(\mathcal{F}) \otimes R\Gamma(X^{(d)}, (E^* \otimes E)^{(d)}).$$

On the other hand, the filtration above expresses the same object as a sum of proper-rank averaging functors applied to $\mathrm{CT}(\mathcal{F})$, hence it vanishes for $d \gg 0$. Because $E^* \otimes E$ contains the trivial local system, the displayed cohomology is nonzero for suitable large d. Therefore $\mathrm{CT}(\mathcal{F}) = 0$, i.e. $\mathcal{F}$ is cuspidal.

9.8. **The quotient category behind t-exactness.** The difficult point in the original proof is t-exactness. The tempting statement "Av_E^1 is t-exact" is false. Instead, one constructs a Verdier quotient

$$\mathcal{D}(\mathsf{Bun}_m) \twoheadrightarrow \widetilde{\mathcal{D}}(\mathsf{Bun}_m)$$

with three properties:

$$\begin{array}{ccc} \mathcal{D}(\mathsf{Bun}_m) & \twoheadrightarrow & \widetilde{\mathcal{D}}(\mathsf{Bun}_m) \\ \downarrow{\scriptstyle \mathrm{Av}_E^1} & & \downarrow{\scriptstyle \widetilde{\mathrm{Av}}_E^1} \\ \mathcal{D}(\mathsf{Bun}_m) & \twoheadrightarrow & \widetilde{\mathcal{D}}(\mathsf{Bun}_m), \end{array}$$

where $\widetilde{\mathrm{Av}}_E^1$ is t-exact, the quotient inherits a perverse t-structure, and the composite

$$\mathcal{D}(\mathsf{Bun}_m)^{\mathrm{cusp}} \longrightarrow \mathcal{D}(\mathsf{Bun}_m) \longrightarrow \widetilde{\mathcal{D}}(\mathsf{Bun}_m)$$

is conservative. This is a Springer-type replacement for the false exactness of Av_E^1. Gaitsgory's original construction of $\widetilde{D}$ is technically involved.

The modern interpretation is that one should take the quotient to be the tempered part:

$$\widetilde{\mathcal{D}}(\mathsf{Bun}_G) = \mathcal{D}(\mathsf{Bun}_G)^{\mathrm{temp}}.$$

The geometric Ramanujan theorem says that cuspidal objects are tempered [Ber21], and the Eisenstein/constant term formalism supplies the remaining non-cuspidal pieces. This gives a cleaner conceptual replacement for the old quotient category.

10. Categorical conjectures

We now pass from individual Hecke eigensheaves to the categorical form of geometric Langlands. In this section $k = \mathbb{C}$ and the coefficient field is also $\mathbb{C}$. We work in the de Rham setting unless stated otherwise. Write $\mathsf{LocSys}_{\check{G}}$ for the derived stack of de Rham $\check{G}$-local systems on X. Ignoring the customary half-twist, the expected equivalence is the Arinkin–Gaitsgory/Gaitsgory–Raskin formulation [AG15, GR24a]:

$$\mathrm{DMod}(\mathsf{Bun}_G) \simeq \mathrm{IndCoh}_{\mathrm{Nilp}}(\mathsf{LocSys}_{\check{G}}).$$

The point of the right side is that $\mathsf{LocSys}_{\check{G}}$ is usually singular as a derived stack. Thus $\mathrm{QCoh}(\mathsf{LocSys}_{\check{G}})$, $\mathrm{IndCoh}(\mathsf{LocSys}_{\check{G}})$, and the intermediate singular-support category $\mathrm{IndCoh}_{\mathrm{Nilp}}(\mathsf{LocSys}_{\check{G}})$ are genuinely different.

10.1. **Why ind-coherent sheaves appear.** For a singular derived scheme or stack Y, perfect and coherent objects need not coincide; in general one only has an inclusion

$$\mathrm{Perf}(Y) \subseteq \mathrm{Coh}(Y),$$

which is often strict. For example, take $A = k[\varepsilon]$ with $\deg(\varepsilon) = -1$, and $Y = \operatorname{Spec} A$. Then $k \in \mathrm{Coh}(Y)$, but $k \notin \mathrm{Perf}(Y)$; one sees this from the infinite periodic resolution

$$\cdots \xrightarrow{\varepsilon} A \xrightarrow{\varepsilon} A \xrightarrow{\varepsilon} A \longrightarrow k.$$

After ind-completion, for reasonable Y,

$$\mathrm{QCoh}(Y) = \mathrm{Ind}(\mathrm{Perf}(Y)) \subseteq \mathrm{IndCoh}(Y) = \mathrm{Ind}(\mathrm{Coh}(Y)),$$

with strictness in the singular examples above. Thus IndCoh retains coherent objects that QCoh can miss.

There is, however, a controlled family of categories between Perf and Coh. Following Arinkin–Gaitsgory [AG15], if Y is quasi-smooth, define its scheme of singularities by

$$\mathrm{Sing}(Y) := \mathrm{Spec}_{Y^{\mathrm{cl}}} \mathrm{Sym}_{\mathcal{O}_{Y^{\mathrm{cl}}}}(H^1(T_Y)).$$

Equivalently, this is the classical stack underlying the (-1)-shifted cotangent stack. If Y is smooth, $\mathrm{Sing}(Y)$ is the zero vector bundle over Y, i.e. the zero section. For $\mathcal{F} \in \mathrm{Coh}(Y)$, one can attach a conical support $\mathrm{supp}(\mathcal{F}) \subset \mathrm{Sing}(Y)$. For a conical closed subset $\mathcal{N} \subset \mathrm{Sing}(Y)$, set

$$\mathrm{Coh}_{\mathcal{N}}(Y) := \{\mathcal{F} \in \mathrm{Coh}(Y) \mid \mathrm{supp}(\mathcal{F}) \subset \mathcal{N}\}, \qquad \mathrm{IndCoh}_{\mathcal{N}}(Y) := \mathrm{Ind}(\mathrm{Coh}_{\mathcal{N}}(Y)).$$

The zero-support category recovers perfect objects:

$$\mathrm{Coh}_0(Y) \simeq \mathrm{Perf}(Y), \qquad \mathrm{IndCoh}_0(Y) \simeq \mathrm{QCoh}(Y).$$

A useful local model is the derived intersection $Y = \mathrm{pt} \times_V \mathrm{pt}$, where both maps to the vector space V are the zero map. Then $H^1(T_Y) = V$, so $\mathrm{Sing}(Y) = V^*$. Koszul duality gives

$$\mathrm{IndCoh}(\mathrm{pt} \times_V \mathrm{pt}) \simeq \mathrm{Sym}(V[-2])\text{-Mod}.$$

Thus $H^\bullet(\mathcal{F})$ is a graded $\mathrm{Sym}(V)$-module, equivalently a quasi-coherent sheaf on V^*, and $\mathrm{supp}(\mathcal{F})$ is the usual conical support of this module. For a general scheme, the formal neighborhood $\widehat{Y}_y$ is controlled to first order by the derived self-intersection over T_yY; this produces the same support. For stacks, the construction is glued by descent.

10.2. **The nilpotent singular support on the spectral side.** Let $(\mathcal{E}, \nabla) \in \mathsf{LocSys}_{\check{G}}$. The adjoint bundle $\check{\mathfrak{g}}_{\mathcal{E}}$ carries the induced connection. The tangent complex is

$$T_{(\mathcal{E},\nabla)}\mathsf{LocSys}_{\check{G}} \simeq R\Gamma_{\mathrm{dR}}(X, \check{\mathfrak{g}}_{\mathcal{E}})[1] \simeq R\Gamma(X, \check{\mathfrak{g}}_{\mathcal{E}} \xrightarrow{\mathrm{ad}(\nabla)} \check{\mathfrak{g}}_{\mathcal{E}} \otimes \Omega_X)[1].$$

The fiber of $\mathrm{Sing}(\mathsf{LocSys}_{\check{G}})$ over $(\mathcal{E}, \nabla)$ is $H^1(T_{(\mathcal{E},\nabla)}\mathsf{LocSys}_{\check{G}})^\vee$. Since $H^1(T) = H^2_{\mathrm{dR}}(X, \check{\mathfrak{g}}_{\mathcal{E}})$, Poincare duality identifies this dual space with $H^0_{\mathrm{dR}}(X, \check{\mathfrak{g}}^*_{\mathcal{E}})$. Consequently the classical singularity stack has points

$$\mathrm{Sing}(\mathsf{LocSys}_{\check{G}}) = \{(\mathcal{E}, \nabla, A) \mid A \in H^0_{\mathrm{dR}}(X, \check{\mathfrak{g}}^*_{\mathcal{E}})\}.$$

Equivalently, A is a horizontal section of the coadjoint local system. Using an invariant form $\check{\mathfrak{g}}^* \simeq \check{\mathfrak{g}}$, define

$$\mathrm{Nilp} := \{(\mathcal{E}, \nabla, A) \mid A_x \in \check{\mathfrak{g}} \text{ is nilpotent for every } x \in X\} \subset \mathrm{Sing}(\mathsf{LocSys}_{\check{G}}).$$

Thus $\mathrm{IndCoh}_{\mathrm{Nilp}}(\mathsf{LocSys}_{\check{G}})$ consists of ind-coherent sheaves whose singular directions are globally nilpotent horizontal fields.

Even the simplest curve shows why the derived structure matters. For $X = \mathbb{P}^1$, the underived stack of de Rham local systems is essentially $\mathrm{pt}/\check{G}$, but the derived stack is the derived zero fiber

$$\mathsf{LocSys}_{\check{G}}(\mathbb{P}^1) \simeq \mathrm{pt}/\check{G} \times_{\check{\mathfrak{g}}/\check{G}} \mathrm{pt}/\check{G}.$$

If one kept only $\mathrm{QCoh}(\mathrm{pt}/\check{G}) \simeq \mathrm{Rep}(\check{G})$, one would lose the derived Ext-data visible on the automorphic side, for instance in the spherical category $\mathrm{DMod}(G(\mathcal{O})\backslash G(K)/G(\mathcal{O}))$. The category $\mathrm{IndCoh}_{\mathrm{Nilp}}$ on the derived self-intersection is the correct refinement.

10.3. **Compatibility test I: derived Satake.** The local spherical category is, compatibly with the Arinkin–Gaitsgory singular-support formalism [AG15],

$$\mathrm{Sph}_{G,x} := \mathrm{DMod}(G(\mathcal{O}_x)\backslash G(K_x)/G(\mathcal{O}_x)).$$

The usual geometric Satake equivalence gives the heart $\mathrm{Rep}(\check{G})$. At the derived categorical level, the correct spectral object is the spectral Hecke stack

$$\mathsf{Hecke}^{\mathrm{spec}}_{\check{G}} := \mathrm{pt}/\check{G} \times_{\check{\mathfrak{g}}/\check{G}} \mathrm{pt}/\check{G} \simeq (\mathrm{pt} \times_{\check{\mathfrak{g}}} \mathrm{pt})/\check{G}.$$

The derived Satake theorem identifies

$$\mathrm{Sph}_{G,x} \simeq \mathrm{IndCoh}_{\mathrm{Nilp}(\check{\mathfrak{g}}^*)/\check{G}}(\mathsf{Hecke}^{\mathrm{spec}}_{\check{G}}).$$

Here $\mathrm{Sing}(\mathsf{Hecke}^{\mathrm{spec}}_{\check{G}}) \simeq \check{\mathfrak{g}}^*/\check{G}$, and the nilpotent cone $\mathrm{Nilp}(\check{\mathfrak{g}}^*)/\check{G}$ is exactly the allowed singular support. This is the local model for the global category $\mathrm{IndCoh}_{\mathrm{Nilp}}(\mathsf{LocSys}_{\check{G}})$.

10.4. **Compatibility test II: Eisenstein series.** Let $P \subset G$ be a parabolic with Levi quotient M. On the automorphic side one has

$$\begin{array}{ccc} & \mathsf{Bun}_P & \\ {}^{p}\swarrow & & \searrow^{q} \\ \mathsf{Bun}_G & & \mathsf{Bun}_M, \end{array}$$

and the Eisenstein functor

$$\mathrm{Eis}_{P,!} := p_! q^* \colon \mathrm{DMod}(\mathsf{Bun}_M) \to \mathrm{DMod}(\mathsf{Bun}_G).$$

Its continuous right adjoint is the constant-term functor

$$\mathrm{CT}_{P,*} := q_* p^! \colon \mathrm{DMod}(\mathsf{Bun}_G) \longrightarrow \mathrm{DMod}(\mathsf{Bun}_M).$$

Thus the automorphic formalism is stable under adjunction.

On the other hand, the spectral diagram is

$$\begin{array}{ccc} & \mathsf{LocSys}_{\check{P}} & \\ {}^{\check{p}}\swarrow & & \searrow^{\check{q}} \\ \mathsf{LocSys}_{\check{G}} & & \mathsf{LocSys}_{\check{M}}. \end{array}$$

A naive functor

$$\check{p}_* \check{q}^* \colon \mathrm{QCoh}(\mathsf{LocSys}_{\check{M}}) \longrightarrow \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}})$$

does not have the correct continuous right adjoint. The obstruction is that $\check{p}$ is proper enough to preserve coherent sheaves, but it does not preserve perfect complexes. The repair is precisely the nilpotent ind-coherent category. Schematically, the spectral Eisenstein functor is

$$\mathrm{Eis}^{\mathrm{spec}}_{\check{P}} := (\check{p})^{\mathrm{IndCoh}}_* \circ \check{q}^!,$$

and it sends perfect objects on the Levi side into nilpotent singular support:

$$\mathrm{Eis}^{\mathrm{spec}}_{\check{P}}(\mathrm{Perf}(\mathsf{LocSys}_{\check{M}})) \subset \mathrm{Coh}_{\mathrm{Nilp}}(\mathsf{LocSys}_{\check{G}}).$$

Moreover, $\mathrm{IndCoh}_{\mathrm{Nilp}}(\mathsf{LocSys}_{\check{G}})$ is generated by the spectral Eisenstein objects $\mathrm{Eis}^{\mathrm{spec}}_{\check{P}}(\mathrm{Perf}(\mathsf{LocSys}_{\check{M}}))$ for all Levi subgroups $\check{M}$, including $\check{M} = \check{G}$. This explains why the nilpotent condition is not optional: it is forced by Eisenstein compatibility.

10.5. **Tempered objects and Whittaker detection.** There is also a smaller expected equivalence

$$\mathrm{DMod}(\mathsf{Bun}_G)^{\mathrm{temp}} \simeq \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}).$$

Ideologically, the tempered part is the portion of $\mathrm{DMod}(\mathsf{Bun}_G)$ detected by Whittaker coefficients. One has a Whittaker category for generically defined N-reductions and a natural !-pushforward

$$\mathrm{Wh}(\mathsf{Bun}_G^{N\text{-gen}}) \longrightarrow \mathrm{DMod}(\mathsf{Bun}_G).$$

The local input is that the chiral homology of the Whittaker category of the affine Grassmannian matches the chiral homology of $\mathrm{Rep}(\check{G})$. This gives generators, but it does not by itself identify the Verdier quotient defining the tempered subcategory. In practice, the modern definition uses the spherical Hecke action.

For every $x \in X$, the category $\mathrm{Sph}_{G,x}$ acts monoidally on $\mathrm{DMod}(\mathsf{Bun}_G)$. Under derived Satake,

$$\mathrm{Sph}_{G,x} \simeq \mathrm{IndCoh}_{\mathrm{Nilp}}(\mathrm{pt}/\check{G} \times_{\check{\mathfrak{g}}/\check{G}} \mathrm{pt}/\check{G}),$$

while the tempered, or left-completed, spectral Hecke category is

$$\mathrm{QCoh}(\mathrm{pt}/\check{G} \times_{\check{\mathfrak{g}}/\check{G}} \mathrm{pt}/\check{G}).$$

One defines $\mathcal{F} \in \mathrm{DMod}(\mathsf{Bun}_G)$ to be *tempered* if the $\mathrm{Sph}_{G,x}$-module generated by $\mathcal{F}$ factors through this left-completed category. A theorem of Færgeman–Raskin says that the resulting subcategory is independent of the chosen point x [FR23].

The full automorphic category is generated by Eisenstein series from tempered compact objects:

$$\mathrm{DMod}(\mathsf{Bun}_G) = \left\langle \mathrm{Eis}_P(\mathrm{DMod}(\mathsf{Bun}_M)^{\mathrm{temp,c}}) \mid P \subset G \right\rangle.$$

Therefore the tempered conjecture, together with the Eisenstein compatibilities above, is equivalent to the full categorical conjecture.

10.6. **Betti and other sheaf theories.** For other sheaf theories, especially Betti or ℓ-adic sheaves, the automorphic category itself must also be restricted by singular support [BZN16, NY19, AGK+20a]. In the Betti setting, $\mathsf{LocSys}_{\check{G}}$ is the derived character stack of representations of $\pi_1(X)$ into $\check{G}$; in the most precise sheaf-theoretic formulations one may replace it by the corresponding restricted variant. The conjectural form is

$$\mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G) \simeq \mathrm{IndCoh}_{\mathrm{Nilp}}(\mathsf{LocSys}_{\check{G}}).$$

Here $\mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G)$ is defined using microlocal singular support. For a sheaf $\mathcal{F} \in \mathrm{Shv}(Y)$, its singular support is a closed conical subset $\mathrm{SS}(\mathcal{F}) \subset T^*Y$, defined by microlocal analysis, equivalently by the vanishing of nearby cycles in non-characteristic directions. A sheaf is a local system precisely when $\mathrm{SS}(\mathcal{F})$ is the zero section.

At a G-bundle $\mathcal{P}$,

$$T_{\mathcal{P}}\mathsf{Bun}_G \simeq R\Gamma(X, \mathfrak{g}_{\mathcal{P}})[1], \qquad T^*_{\mathcal{P}}\mathsf{Bun}_G \simeq R\Gamma(X, \mathfrak{g}^*_{\mathcal{P}} \otimes \Omega_X).$$

The classical cotangent stack is the Higgs stack

$$T^*\mathsf{Bun}_G = \mathsf{Higgs}_G, \qquad \mathsf{Higgs}_G(\mathbb{C}) = \{(\mathcal{P}, \phi) \mid \phi \in H^0(X, \mathfrak{g}^*_{\mathcal{P}} \otimes \Omega_X)\}.$$

The global nilpotent cone is

$$\mathrm{Nilp}_{\mathrm{glob}} := \{(\mathcal{P}, \phi) \in \mathsf{Higgs}_G \mid \phi_x \in \mathcal{N}_{\mathfrak{g}^*} \text{ for every } x \in X\}.$$

Then

$$\mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G) := \{\mathcal{F} \in \mathrm{Shv}(\mathsf{Bun}_G) \mid \mathrm{SS}(\mathcal{F}) \subset \mathrm{Nilp}_{\mathrm{glob}}\}.$$

This is the Betti analogue of imposing nilpotent singular support on the spectral ind-coherent side.

11. The category $\mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G)$

Section 10 stated the de Rham categorical conjecture [AG15, GR24a]

$$\mathrm{DMod}(\mathsf{Bun}_G) \simeq \mathrm{IndCoh}_{\mathrm{Nilp}}(\mathsf{LocSys}^{\mathrm{dR}}_{\check{G}}).$$

This section asks for the analogue in other sheaf theories, especially the Betti formulation of [BZN16, NY19]. In the Betti setting one works with ordinary sheaves on the complex analytic stack $\mathsf{Bun}_G(\mathbb{C})$. The correct automorphic category is not all sheaves, but the full subcategory with nilpotent microlocal singular support.

11.1. **All sheaves on an algebraic stack.** For a finite type affine scheme S over $\mathbb{C}$, set

$$\mathrm{Shv}^{\mathrm{all}}(S) := \mathrm{Shv}^{\mathrm{all}}(S(\mathbb{C})),$$

with no constructibility condition. If Y is an algebraic stack, define sheaves on Y by smooth descent:

$$\mathrm{Shv}^{\mathrm{all}}(Y) := \lim_{(S \to Y) \in (\mathrm{Aff}^{\mathrm{sm}}_{/Y})^{\mathrm{op}}} \mathrm{Shv}^{\mathrm{all}}(S),$$

where the transition functors are pullbacks. Equivalently, by adjunction, this may be written as a colimit

$$\mathrm{Shv}^{\mathrm{all}}(Y) \simeq \operatorname*{colim}_{\substack{S \to Y, \\ S \text{ smooth}}} \mathrm{Shv}^{\mathrm{all}}(S),$$

where the transition functors are the corresponding left adjoints. The general categorical principle is the following: if $\mathrm{L}_{ij} : C_i \rightleftarrows C_j : \mathrm{R}_{ij}$ are adjoint functors between presentable categories, then

$$\operatorname*{colim}_{\mathrm{L}_{ij}} C_i \simeq \lim_{\mathrm{R}_{ij}} C_i.$$

Thus $\mathrm{Shv}^{\mathrm{all}}(\mathsf{Bun}_G)$ is a large presentable category built from analytic charts of Bun_G.

11.2. **Betti versus de Rham local systems.** In the de Rham setting, spectral parameters are de Rham local systems:

$$\mathsf{LocSys}^{\mathrm{dR}}_{\check{G}}(X) = \mathrm{Conn}_{\check{G}}(X) := \{(\mathcal{E}_{\check{G}}, \nabla) \mid \mathcal{E}_{\check{G}} \text{ is a } \check{G}\text{-torsor and } \nabla \text{ is a connection}\}.$$

In the Betti setting, spectral parameters are Betti local systems:

$$\mathsf{LocSys}^{\mathrm{Betti}}_{\check{G}}(X) := \mathrm{Map}(X^{\mathrm{top}}, \mathbb{B}\check{G}).$$

After choosing a base point $x \in X(\mathbb{C})$, the classical truncation of the rigidified Betti moduli is

$$\mathsf{LocSys}^{\mathrm{Betti,rig,cl}}_{\check{G},x}(X) \simeq \mathrm{Hom}(\pi_1(X(\mathbb{C}), x), \check{G}).$$

The unrigidified classical stack is obtained by quotienting by conjugation. This description is only a first approximation: the Betti moduli used in categorical Langlands is derived, with tangent complex governed by cochains of $X(\mathbb{C})$ with coefficients in the adjoint local system.

The two moduli stacks are not algebraically the same. The de Rham stack depends on the complex algebraic curve X, while the Betti stack depends only on the topology of $X(\mathbb{C})$. For example, let $X = E$ be an elliptic curve and choose $e \in E$. Then

$$\mathsf{LocSys}^{\mathrm{dR,rig}}_{\mathbb{G}_{\mathrm{m}},e}(E) \longrightarrow \mathrm{Pic}^0_E \simeq E^\vee$$

is the universal vector extension, hence a $\mathbb{G}_{\mathrm{a}}$-torsor, equivalently an $\mathbb{A}^1$-bundle, over $E^\vee$. By contrast,

$$\mathsf{LocSys}^{\mathrm{Betti,rig}}_{\mathbb{G}_{\mathrm{m}},e}(E) \simeq \mathrm{Hom}(\mathbb{Z}^2, \mathbb{G}_{\mathrm{m}}) \simeq \mathbb{G}_{\mathrm{m}}^2.$$

The Riemann–Hilbert correspondence identifies the complex points analytically, but it uses the exponential map and does not identify the algebraic structures.

11.3. **Why the naive Betti conjecture is too large.** A first guess would be

$$\mathrm{Shv}^{\mathrm{all}}(\mathsf{Bun}_G) \overset{?}{\simeq} \mathrm{IndCoh}_{\mathrm{Nilp}}(\mathsf{LocSys}^{\mathrm{Betti}}_{\check{G}}).$$

This is false already for $G = \mathrm{GL}_1$ and $X = E$ an elliptic curve. On the automorphic side,

$$\mathsf{Bun}_{\mathbb{G}_{\mathrm{m}}} \simeq \coprod_{d\in\mathbb{Z}} \mathrm{Pic}^d_E \times \mathbb{B}\mathbb{G}_{\mathrm{m}} \simeq E \times \mathbb{Z} \times \mathbb{B}\mathbb{G}_{\mathrm{m}}.$$

Topological Fourier–Mellin transforms identify the finitary or constructible normalization with

$$\mathrm{Shv}(\mathbb{Z})_{\mathrm{fin}} \simeq \mathrm{QCoh}(\mathbb{B}\mathbb{G}_{\mathrm{m}}), \qquad \mathrm{Shv}(\mathbb{B}\mathbb{G}_{\mathrm{m}})_{\mathrm{fin}} \simeq \mathrm{QCoh}(\mathrm{pt} \times_{\mathbb{A}^1} \mathrm{pt}).$$

For literal all sheaves these categories are completed; this completion issue is not the source of the mismatch in the example. The derived Betti stack for rank-one local systems has the corresponding factors

$$\mathsf{LocSys}^{\mathrm{Betti}}_{\mathbb{G}_{\mathrm{m}}}(E) \simeq \mathbb{G}_{\mathrm{m}}^2 \times \mathbb{B}\mathbb{G}_{\mathrm{m}} \times (\mathrm{pt} \times_{\mathbb{A}^1} \mathrm{pt}).$$

Thus the only mismatch is the elliptic-curve factor:

$$\mathrm{Shv}^{\mathrm{all}}(E) \neq \mathrm{QCoh}(\mathbb{G}_{\mathrm{m}}^2).$$

The correct replacement of $\mathrm{Shv}^{\mathrm{all}}(E)$ is the full subcategory of lisse sheaves

$$\mathrm{QLisse}(E) := \{\mathcal{F} \in \mathrm{Shv}^{\mathrm{all}}(E) \mid H^i(\mathcal{F}) \text{ ind-object generated by finite-dimensional local systems}\}.$$

Equivalently,

$$\mathrm{QLisse}(E) = \mathrm{Shv}^{\mathrm{all}}_0(E) := \{\mathcal{F} \mid \mathrm{SS}(\mathcal{F}) \subset T^*_E E\}.$$

Here $T^*_E E$ is the zero section. Thus the error in the naive conjecture is exactly the presence of sheaves with nonzero microlocal support.

11.4. **Microlocal singular support for sheaves.** Let S be a smooth complex scheme. For constructible $\mathcal{F}$, microlocal analysis attaches a closed conical subset, in the sense recalled in [Bei16],

$$\mathrm{SS}(\mathcal{F}) \subset T^*S,$$

called the singular support. It measures the failure of $\mathcal{F}$ to be locally constant. More precisely, if a tangent direction $v \in T_sS$ is annihilated by all covectors in $\mathrm{SS}(\mathcal{F})_s$, then $\mathcal{F}$ is locally acyclic, or lisse, along that direction. For a closed conical subset $\mathcal{N} \subset T^*S$, define the all-sheaf version by ind-completion:

$$\mathrm{Shv}^{\mathrm{all}}_{\mathcal{N}}(S) := \mathrm{Ind}\{\mathcal{F} \in \mathrm{Shv}_{\mathrm{cons}}(S) \mid \mathrm{SS}(\mathcal{F}) \subset \mathcal{N}\}.$$

Equivalently, one may view it as the full cocomplete subcategory of $\mathrm{Shv}^{\mathrm{all}}(S)$ generated by constructible objects with singular support in $\mathcal{N}$. The zero-support category is the lisse, or quasi-lisse, category:

$$\mathrm{Shv}^{\mathrm{all}}_0(S) \simeq \mathrm{QLisse}(S).$$

In the AGKRRV normalization this means the left completion of the ind-category of finite-rank local systems.

For comparison with the de Rham side, if M is a regular holonomic D-module and $\mathrm{Sol}(M)$ is its solution sheaf, then

$$\mathrm{SS}(\mathrm{Sol}(M)) = \mathrm{Char}(M).$$

Thus the microlocal support of Betti sheaves is the Betti counterpart of the characteristic variety of holonomic D-modules.

11.5. **The global nilpotent cone.** For a G-bundle $\mathcal{P}$, write $\mathfrak{g}_{\mathcal{P}} = \mathcal{P} \times^G \mathfrak{g}$. The classical cotangent stack of Bun_G is the Higgs stack

$$T^*\mathsf{Bun}_G \simeq \mathsf{Higgs}_G, \qquad \mathsf{Higgs}_G(\mathbb{C}) = \{(\mathcal{P}, \phi) \mid \phi \in H^0(X, \mathfrak{g}^*_{\mathcal{P}} \otimes \Omega_X)\}.$$

The global nilpotent cone is

$$\mathrm{Nilp}_{\mathrm{glob}} := \{(\mathcal{P}, \phi) \in T^*\mathsf{Bun}_G \mid \phi_x \in \mathcal{N}_{\mathfrak{g}^*} \text{ for every } x \in X\}.$$

It is the global analogue of the ordinary nilpotent cone $\mathcal{N}_{\mathfrak{g}^*} \subset \mathfrak{g}^*$.

Equivalently, $\mathrm{Nilp}_{\mathrm{glob}}$ is the zero fiber of the Hitchin fibration. If $p_i \in \mathbb{C}[\mathfrak{g}^*]^G$ are homogeneous generators of degrees d_i, set

$$A_G := \bigoplus_i H^0(X, \Omega_X^{d_i}).$$

The Hitchin map is

$$h\colon T^*\mathsf{Bun}_G \longrightarrow A_G, \qquad (\mathcal{P}, \phi) \longmapsto (p_i(\phi))_i.$$

Then

$$\mathrm{Nilp}_{\mathrm{glob}} = h^{-1}(0).$$

In invariant notation this is the map

$$\Gamma(X, \mathfrak{g}^*_{\mathcal{P}} \otimes \Omega_X) \longrightarrow \Gamma(X, (\mathfrak{g}^* /\!\!/ G)_{\Omega_X}),$$

where $(-)_{\Omega_X}$ denotes the $\mathbb{G}_{\mathrm{m}}$-twist coming from dilation on $\mathfrak{g}^*$. With homogeneous generators p_i of degrees d_i, this target is $\bigoplus_i H^0(X, \Omega_X^{d_i})$.

11.6. **The Betti categorical conjecture with nilpotent support.** The corrected Betti conjecture is the nilpotent-support version of Betti geometric Langlands [BZN16, NY19]:

$$\mathrm{Shv}^{\mathrm{all}}_{\mathrm{Nilp}}(\mathsf{Bun}_G) \simeq \mathrm{IndCoh}_{\mathrm{Nilp}}(\mathsf{LocSys}^{\mathrm{Betti}}_{\check{G}}),$$

where

$$\mathrm{Shv}^{\mathrm{all}}_{\mathrm{Nilp}}(\mathsf{Bun}_G) := \{\mathcal{F} \in \mathrm{Shv}^{\mathrm{all}}(\mathsf{Bun}_G) \mid \mathrm{SS}(\mathcal{F}) \subset \mathrm{Nilp}_{\mathrm{glob}}\}.$$

In the formulation valid for general sheaf theories, especially ℓ-adic sheaves, the spectral stack is replaced by the restricted-variation version $\mathsf{LocSys}^{\mathrm{restr}}_{\check{G}}$, but the automorphic category remains the nilpotent-support subcategory of sheaves on Bun_G.

This condition has three basic checks. First, it is true for tori: since the nilpotent cone in $\mathfrak{t}^*$ is the zero section, $\mathrm{Shv}^{\mathrm{all}}_{\mathrm{Nilp}}(\mathsf{Bun}_T)$ is the lisse part, and the Mellin transform identifies it with quasi-coherent sheaves on the Betti dual torus local-system stack. Second, it is not too small: a Hecke eigensheaf should have nilpotent singular support. Indeed, Nadler–Yun prove local constancy in the point of X for Hecke functors on sheaves with nilpotent singular support [NY19]; combined with the converse/extension used in AGKRRV [AGK+20a], Hecke-lisse behavior is equivalent to nilpotent singular support. In particular, eigensheaves belong to $\mathrm{Shv}^{\mathrm{all}}_{\mathrm{Nilp}}(\mathsf{Bun}_G)$. Third, it is not too large: $\mathrm{Shv}^{\mathrm{all}}_{\mathrm{Nilp}}(\mathsf{Bun}_G)$ is compactly generated, whereas $\mathrm{Shv}^{\mathrm{all}}(\mathsf{Bun}_G)$ is generally not.

More generally, for a smooth scheme S and conical subset $\mathcal{N} \subset T^*S$, the category $\mathrm{Shv}^{\mathrm{all}}_{\mathcal{N}}(S)$ is controlled by a stratification adapted to $\mathcal{N}$: objects with singular support in $\mathcal{N}$ are locally constant along the strata. This is the microlocal reason that nilpotent support restores finiteness properties.

12. Spectral decomposition and $\mathsf{LocSys}^{\mathrm{restr}}_{\check{G}}$

The preceding section introduced the Betti form of categorical Langlands [BZN16, NY19]:

$$\mathrm{Shv}^{\mathrm{all}}_{\mathrm{Nilp}}(\mathsf{Bun}_G) \simeq \mathrm{IndCoh}_{\mathrm{Nilp}}(\mathsf{LocSys}^{\mathrm{Betti}}_{\check{G}}).$$

In particular the compact spectral category $\mathrm{Perf}(\mathsf{LocSys}^{\mathrm{Betti}}_{\check{G}})$ acts on $\mathrm{Shv}^{\mathrm{all}}_{\mathrm{Nilp}}(\mathsf{Bun}_G)$; in the presentable form this is often written as a $\mathrm{QCoh}(\mathsf{LocSys}^{\mathrm{Betti}}_{\check{G}})$-action. The point of this section is that, in the Betti setting, this action can be obtained cleanly from local Hecke actions, as in [NY19, AGK+20a]. The same mechanism also produces the restricted spectral stack $\mathsf{LocSys}^{\mathrm{restr}}_{\check{G}}$, which is the correct spectral object for constructible and ℓ-adic sheaf theories.

12.1. **Local systems on a homotopy type.** Let $Y \in \mathsf{Spc}$ be a homotopy type; for example, Y may be the homotopy type of a topological space. Define

$$\mathsf{LocSys}(Y) := \{\text{local systems on } Y\} \simeq \mathrm{Fun}(Y, \mathsf{Vect}).$$

Equivalently, viewing Y as an ∞-groupoid and using the constant diagram with value Vect,

$$\mathsf{LocSys}(Y) \simeq \lim_{Y} \mathsf{Vect} \simeq \operatorname*{colim}_{Y} \mathsf{Vect}.$$

The last identification uses the self-duality of Vect as a DG category. Caution that $\mathsf{LocSys}(Y)$ depends on the homotopy type of Y, not only on $\pi_1(Y)$, even when Y is connected.

For a map $f\colon Y_1 \to Y_2$, pullback gives

$$f^*\colon \mathsf{LocSys}(Y_2) \longrightarrow \mathsf{LocSys}(Y_1).$$

It has adjoints in the presentable setting, denoted f_* and $f_!$ when they are defined by the usual Kan-extension formulas. The functor $f_!$ is already subtle topologically: for instance, for $\mathrm{pt} \to S^1$, the result corresponds to a local system with an infinite Jordan block. Thus one should not expect all pushforwards to stay inside a small constructible or finite-rank category.

12.2. **The Betti local-to-global theorem.** For a homotopy type Y, consider the Betti spectral stack

$$\mathrm{Maps}(Y, \mathbb{B}\check{G}).$$

It is characterized by the following functor-of-points description. For a test affine scheme S,

$$\mathrm{Maps}(S, \mathrm{Maps}(Y, \mathbb{B}\check{G})) \simeq \mathrm{Maps}(S \times Y, \mathbb{B}\check{G}).$$

Equivalently, an S-point is a symmetric monoidal functor

$$\mathrm{Rep}(\check{G}) \longrightarrow \mathrm{QCoh}(S) \otimes \mathsf{LocSys}(Y),$$

with the usual right-exact or right t-exact condition when the chosen sheaf theory carries such a structure. For $Y = X(\mathbb{C})$, this is $\mathsf{LocSys}^{\mathrm{Betti}}_{\check{G}}(X)$.

The Betti local-to-global theorem of [NY19, AGK+20a] says that an action of the spectral monoidal category $\mathrm{Perf}(\mathrm{Maps}(Y, \mathbb{B}\check{G}))$, or of its presentable QCoh-completion when such completion is allowed, on a dualizable DG category $\mathcal{C}$ is encoded by the following factorization data. For every finite set I, one has a monoidal functor

$$\mathrm{Rep}(\check{G})^{\otimes I} \longrightarrow \mathrm{End}(\mathcal{C}) \otimes \mathsf{LocSys}(Y^I),$$

compatible with all maps of finite sets. If $I \to J$ is such a map, the associated diagonal $\Delta\colon Y^J \to Y^I$ gives the compatibility square

$$\begin{array}{ccc}
\mathrm{Rep}(\check{G})^{\otimes I} & \longrightarrow & \mathrm{End}(\mathcal{C}) \otimes \mathsf{LocSys}(Y^I) \\
\downarrow & & \downarrow{\scriptstyle \mathrm{Id} \otimes \Delta^*} \\
\mathrm{Rep}(\check{G})^{\otimes J} & \longrightarrow & \mathrm{End}(\mathcal{C}) \otimes \mathsf{LocSys}(Y^J).
\end{array}$$

The point is that the single spectral category $\mathrm{Perf}(\mathrm{Maps}(Y, \mathbb{B}\check{G}))$, or its QCoh-completion, packages all finite-set Hecke compatibilities.

12.3. **Comparison with the de Rham Ran construction.** In the de Rham setting one replaces $Y = X(\mathbb{C})$ and $\mathsf{LocSys}(Y^I)$ by the de Rham sheaf theory $\mathrm{DMod}(X^I)$. The local Hecke categories are organized into a Ran-category expression

$$\mathrm{Rep}(\check{G})_{\mathrm{Ran}} \simeq \operatorname*{colim}_I (\mathrm{Rep}(\check{G})^{\otimes I} \otimes \mathrm{DMod}(X^I)),$$

with transition maps built from diagonals. There is a natural monoidal functor

$$\mathrm{Rep}(\check{G})_{\mathrm{Ran}} \longrightarrow \mathrm{QCoh}(\mathsf{LocSys}^{\mathrm{dR}}_{\check{G}}).$$

Its kernel is large and hard to describe; this is the source of the vanishing conjectures from the preceding sections.

The Betti situation is cleaner but not completely formal. For diagonals $\Delta\colon Y^J \to Y^I$, the left Kan extension

$$\Delta_!\colon \mathsf{LocSys}(Y^J) \longrightarrow \mathsf{LocSys}(Y^I)$$

needed for a Ran-style colimit is not monoidal in general; already maps such as $\mathrm{pt} \to S^1$ can create infinite Jordan-block local systems. Hence one cannot naively form a single monoidal category by taking colimits of the $\mathsf{LocSys}(Y^I)$-valued Hecke data. The local-to-global theorem, Theorem 9.3.1, says that the correct monoidal object is nevertheless $\mathrm{Perf}(\mathrm{Maps}(Y, \mathbb{B}\check{G}))$, or its QCoh-completion.

12.4. **Hecke-lisse sheaves and nilpotent singular support.** Let X be the curve and let $\mathcal{F} \in \mathrm{Shv}^{\mathrm{all}}(\mathsf{Bun}_G)$. The global Hecke correspondence gives, for $V \in \mathrm{Rep}(\check{G})$, a functor

$$\mathcal{H}(V, -)\colon \mathrm{Shv}^{\mathrm{all}}(\mathsf{Bun}_G) \longrightarrow \mathrm{Shv}^{\mathrm{all}}(\mathsf{Bun}_G \times X).$$

We say that $\mathcal{F}$ is *Hecke-lisse* if for every finite set I and every $V_I \in \mathrm{Rep}(\check{G})^{\otimes I}$, the iterated Hecke transform lies in the lisse-in-the-curve subcategory:

$$\mathcal{H}_I(V_I, \mathcal{F}) \in \mathrm{Shv}^{\mathrm{all}}(\mathsf{Bun}_G) \otimes \mathsf{LocSys}(X(\mathbb{C})^I) \subset \mathrm{Shv}^{\mathrm{all}}(\mathsf{Bun}_G \times X^I).$$

Equivalently, all Hecke transforms are locally constant along the X^I-directions.

The Betti local-to-global theorem in Subsection 12.2 shows the nilpotent-support condition implies the Hecke-lisse property [NY19], and the form used in AGKRRV upgrades this to the equivalence [AGK+20a]

$$\mathrm{Shv}^{\mathrm{all}}_{\mathrm{Nilp}}(\mathsf{Bun}_G) \simeq \mathrm{Shv}^{\mathrm{all}}(\mathsf{Bun}_G)^{\text{Hecke-lisse}}.$$

Therefore the Hecke action on $\mathrm{Shv}^{\mathrm{all}}_{\mathrm{Nilp}}(\mathsf{Bun}_G)$ satisfies the hypotheses of the Betti local-to-global theorem, and one obtains the spectral action

$$\mathrm{Perf}(\mathsf{LocSys}^{\mathrm{Betti}}_{\check{G}}) \curvearrowright \mathrm{Shv}^{\mathrm{all}}_{\mathrm{Nilp}}(\mathsf{Bun}_G),$$

or, after presentable completion, the corresponding QCoh-action.

12.5. **Other sheaf theories and the restricted stack.** The same theorem has versions in several sheaf-theoretic contexts:

- full Betti sheaves over $\mathbb{C}$, with $\mathrm{Shv}^{\mathrm{all}}$;
- constructible Betti sheaves, usually after passing to $\mathrm{Ind}(\mathrm{Shv}_{\mathrm{cons}})$;
- étale or ℓ-adic sheaves with coefficients such as $\mathbb{Z}_\ell$, $\mathbb{Q}_\ell$, or $\overline{\mathbb{Q}}_\ell$;
- de Rham sheaves, including full D-modules and the holonomic or regular-holonomic ind-completions.

In a constructible context the category $\mathsf{LocSys}(X(\mathbb{C}))$ is too large. The correct replacement is

$$\mathrm{QLisse}(X) \coloneqq \mathrm{Shv}_0(X),$$

where the subscript means zero singular support. In the AGKRRV construction [AGK+20a] this is more precisely the left completion of $\mathrm{Ind}(\mathrm{Lisse}(X))$, where $\mathrm{Lisse}(X)$ denotes finite-rank local systems in the chosen sheaf theory. Thus the restriction is not "finite rank only"; it keeps the quasi-lisse part and removes sheaves with nonzero microlocal support. The notation r. t. ex. means right t-exact; this is the constructible analogue of the right-exact tensor-functor condition in the ordinary Betti stack.

Thus one defines the restricted spectral stack by the functor-of-points rule

$$\mathsf{LocSys}^{\mathrm{restr}}_{\check{G}}(S) \coloneqq \mathrm{Fun}^{\otimes,\mathrm{r.t.ex.}}(\mathrm{Rep}(\check{G}), \mathrm{QLisse}(S \times X)).$$

For affine test schemes this is often written heuristically as the target $\mathrm{QCoh}(S)\otimes\mathrm{QLisse}(X)$; the point is that the family is allowed to vary only in the zero-singular-support, or lisse, direction along X. The local-to-global theorem then gives an action

$$\mathrm{QCoh}(\mathsf{LocSys}^{\mathrm{restr}}_{\check{G}}) \curvearrowright \mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G)$$

in any constructible or ℓ-adic context where $\mathrm{Shv}_{\mathrm{Nilp}}$ is defined.

Over $\mathbb{C}$, the restricted stack is best viewed as a formal refinement of the Betti stack, with an analytic comparison to the de Rham picture through Riemann–Hilbert. There is a natural map

$$\mathsf{LocSys}^{\mathrm{restr}}_{\check{G}} \longrightarrow \mathsf{LocSys}^{\mathrm{Betti}}_{\check{G}}.$$

Roughly, if $\pi\colon \mathsf{LocSys}^{\mathrm{Betti}}_{\check{G}} \to \mathsf{LocSys}^{\mathrm{coarse}}_{\check{G}}$ denotes the map to the GIT coarse moduli of semisimple local systems, then locally over a semisimplification $\bar{\sigma}$ the restricted stack is modeled by the formal completion $(\pi^{-1}(\bar{\sigma}))^{\wedge}$. This is a formal-local description, not an ordinary disjoint union of all such completions. In this description $\pi(\sigma_1) = \pi(\sigma_2)$ exactly when σ_1 and σ_2 have the same semisimplification. The de Rham analogue is obtained only after passing through the analytic Riemann–Hilbert comparison, so it should not be read as a literal algebraic map $\mathsf{LocSys}^{\mathrm{restr}} \to \mathsf{LocSys}^{\mathrm{dR}}$.

12.6. **Restricted categorical Langlands.** The resulting restricted form of the Betti/constructible conjecture is

$$\mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G) \overset{?}{\simeq} \mathrm{IndCoh}_{\mathrm{Nilp}}(\mathsf{LocSys}^{\mathrm{restr}}_{\check{G}}).$$

This section records a caveat: in some constructible formulations this naive equivalence may be too large or may require restricting to a direct summand, for instance to connected components selected by the coefficient theory. What is robust, and what is used for spectral decomposition, is the action

$$\mathrm{QCoh}(\mathsf{LocSys}^{\mathrm{restr}}_{\check{G}}) \curvearrowright \mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G).$$

This is the sheaf-theoretic replacement for decomposing automorphic objects by Langlands parameters.

12.7. **Categorical trace and functions over finite fields.** Now take $k = \mathbb{F}_q$ and work with ℓ-adic sheaves. Frobenius acts on the automorphic category:

$$\mathrm{Fr}\colon \mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G) \longrightarrow \mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G).$$

The expected sheaf-function statement is the categorical trace identity

$$\mathrm{Tr}(\mathrm{Fr}; \mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G)) \simeq \mathcal{C}^{\infty}_{\mathrm{stk}}(\mathsf{Bun}_G(\mathbb{F}_q)),$$

up to the usual choices of completions, finiteness conditions, and stacky automorphism weights on automorphic functions.

Recall the definition of categorical trace. For a finite-dimensional vector space V and an endomorphism $f\colon V \to V$, the ordinary trace is the scalar obtained from

$$\big[\, k \xrightarrow{\mathrm{coev}} V\otimes V^* \xrightarrow{f\otimes\mathrm{Id}} V\otimes V^* \xrightarrow{\mathrm{ev}} k \,\big].$$

For a dualizable DG category $\mathcal{C}$ and an endofunctor $F\colon \mathcal{C}\to\mathcal{C}$, the same formula gives an object of Vect:

$$\mathrm{Tr}(F;\mathcal{C}) := \big[\, \mathsf{Vect} \xrightarrow{\mathrm{coev}} \mathcal{C}\otimes\mathcal{C}^{\vee} \xrightarrow{F\otimes\mathrm{Id}} \mathcal{C}\otimes\mathcal{C}^{\vee} \xrightarrow{\mathrm{ev}} \mathsf{Vect} \,\big].$$

The category $\mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G)$ is expected, and in the restricted theory known in the required sense, to be dualizable [AGK+20b], so this trace makes sense.

More generally, for $V_I \in \mathrm{Rep}(\check{G})^{\otimes I}$, form the endofunctor with parameters in X^I

$$\mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G) \xrightarrow{\mathrm{Fr}} \mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G) \xrightarrow{\mathcal{H}_I(V_I,-)} \mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G)\otimes\mathrm{QLisse}(X^I).$$

Its categorical trace is an object

$$\mathrm{Tr}(\mathcal{H}_I(V_I,-)\circ\mathrm{Fr}; \mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G)) \in \mathrm{QLisse}(X^I).$$

The conjectural identification is that this object is the cohomology of the corresponding shtuka moduli space. Thus the categorical trace formalism geometrizes the passage from geometric to function-theoretic Langlands and V. Lafforgue's excursion operators [Gai16a, Laf18].

Finally, the action of $\mathrm{QCoh}(\mathsf{LocSys}^{\mathrm{restr}}_{\check{G}})$ on $\mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G)$ upgrades the trace $\mathrm{Tr}(\mathrm{Fr};\mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G))$ to a quasi-coherent sheaf on the Frobenius-twisted fixed stack

$$\mathsf{LocSys}^{\mathrm{restr,Fr}}_{\check{G}} := \mathsf{LocSys}^{\mathrm{restr}}_{\check{G}} \times_{\mathsf{LocSys}^{\mathrm{restr}}_{\check{G}}\times\mathsf{LocSys}^{\mathrm{restr}}_{\check{G}},\Delta,\,(\mathrm{Id},\mathrm{Fr})} \mathsf{LocSys}^{\mathrm{restr}}_{\check{G}},$$

whose points are restricted Weil $\check{G}$-local systems. Taking global sections recovers the automorphic trace space above. This is the spectral decomposition statement behind the function-theoretic Langlands parametrization.

13. Conservativity of Whittaker Coefficients

The aim is to explain why Whittaker coefficients detect the tempered part of the automorphic category [FR25, Gai18]. In the de Rham form of categorical Langlands one expects

$$\mathrm{DMod}(\mathsf{Bun}_G) \simeq \mathrm{IndCoh}_{\mathrm{Nilp}}(\mathsf{LocSys}_{\check{G}}),$$

and

$$\mathrm{DMod}(\mathsf{Bun}_G)^{\mathrm{temp}} \simeq \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}).$$

In a Betti or constructible theory the corresponding spectral stack is $\mathsf{LocSys}_{\check{G}}^{\mathrm{restr}}$, and one works with $\mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G)$, the nilpotent-singular-support part of the sheaf category. The theorem recorded in this section is the conservativity of the renormalized Whittaker coefficient functors

$$\mathrm{DMod}(\mathsf{Bun}_G)^{\mathrm{temp}} \xrightarrow{\mathrm{Coeff}^{\mathrm{ren}}} \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}),$$

$$\mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G)^{\mathrm{temp}} \xrightarrow{\mathrm{Coeff}^{\mathrm{ren}}_{\mathrm{restr}}} \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}^{\mathrm{restr}}).$$

Here "conservative" means that if the coefficient is zero, then the object itself is zero. In the de Rham incarnation, the analogous constructible category is the ind-completion of regular holonomic D-modules whose characteristic variety is contained in the global nilpotent cone. This is not the whole $\mathrm{DMod}(\mathsf{Bun}_G)$, but the regular-holonomic/nilpotent-support part.

13.1. **The ordinary and renormalized coefficient functors.** Fix a Borel $B = TN$ and a non-degenerate additive character $\psi\colon N \to \mathbb{G}_{\mathrm{a}}$. Let Bun_N^{Ω} be the usual Ω-twisted stack of N-bundles, with projection $p\colon \mathsf{Bun}_N^{\Omega} \to \mathsf{Bun}_G$ and Whittaker character $\psi\colon \mathsf{Bun}_N^{\Omega} \to \mathbb{A}^1$. Up to the standard cohomological normalization, the first Whittaker coefficient is the compactly supported de Rham cohomology

$$\mathrm{coeff}(\mathcal{F}) \coloneqq C_{\mathrm{dR,c}}(\mathsf{Bun}_N^{\Omega}, p^!(\mathcal{F}) \otimes \psi^!(\exp)) \in \mathsf{Vect}.$$

Equivalently, it is the !-pushforward to a point of the pullback of $\mathcal{F}$ to the Whittaker stack, twisted by the exponential sheaf. The shift is chosen compatibly with Verdier duality and with the microlocal index formula below.

The spectral action gives

$$\mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}) \otimes \mathrm{DMod}(\mathsf{Bun}_G) \longrightarrow \mathrm{DMod}(\mathsf{Bun}_G).$$

By duality, a scalar functional $c\colon \mathrm{DMod}(\mathsf{Bun}_G) \to \mathsf{Vect}$ produces a spectral-valued functional

$$\mathrm{DMod}(\mathsf{Bun}_G) \xrightarrow{\mathrm{coact}} \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}) \otimes \mathrm{DMod}(\mathsf{Bun}_G) \xrightarrow{\mathrm{Id} \otimes c} \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}).$$

For $c = \mathrm{coeff}$ this is $\mathrm{Coeff}^{\mathrm{ren}}$. Its counit, or global-sections, shadow recovers the scalar coefficient:

$$\Gamma(\mathsf{LocSys}_{\check{G}}, \mathrm{Coeff}^{\mathrm{ren}}(\mathcal{F})) \simeq \mathrm{coeff}(\mathcal{F}).$$

This recovery does not by itself characterize the $\mathrm{QCoh}(\mathsf{LocSys}_{\check{G}})$-valued object, since $\mathsf{LocSys}_{\check{G}}$ is not affine in general; the definition is the coaction formula above. The same remarks apply in the restricted Betti/constructible setting. Under the usual rigidity hypotheses on the spectral category, giving a $\mathrm{QCoh}(\mathsf{LocSys}_{\check{G}})$-linear functor to $\mathrm{QCoh}(\mathsf{LocSys}_{\check{G}})$ is equivalent to giving the underlying scalar functor to Vect. This is the formal reason why a scalar Whittaker functional can define a spectral-valued one.

The proof uses two basic facts. First,

$$\mathrm{coeff}\colon \mathrm{DMod}(\mathsf{Bun}_G) \longrightarrow \mathsf{Vect}$$

factors through the projection to the tempered quotient. Second, the action of $\mathrm{QCoh}(\mathsf{LocSys}_{\check{G}})$ on $\mathrm{DMod}(\mathsf{Bun}_G)^{\mathrm{temp}}$ is stable under all Satake/Hecke operators, hence $\mathrm{Coeff}^{\mathrm{ren}}$ also factors through the tempered part.

13.2. **Local Whittaker categories and temperedness.** Let $K = \mathbb{C}((t))$ and $\mathcal{O} = \mathbb{C}[\![t]\!]$. If a DG category $\mathcal{C}$ carries a $G(K)$-action, its non-degenerate Whittaker category is

$$\mathrm{Wh}(\mathcal{C}) := \mathcal{C}^{N(K),\psi}.$$

The local geometric Langlands philosophy predicts, in the Whittaker formalism of [Gai18, FR23],

$$\mathrm{DMod}(G(K))\text{-}\mathrm{Mod}^{\mathrm{temp}} \simeq \mathrm{ShvCat}(\mathsf{LocSys}_{\check{G}}(D^\circ)),$$

where $D^\circ = \operatorname{Spec} K$. Under this picture the bi-Whittaker category

$$\mathrm{BiWhit}(G) := \mathrm{DMod}((N(K),\psi)\backslash G(K)/(N(K),\psi))$$

should identify with

$$\mathrm{BiWhit}(G) \simeq \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}(D^\circ)).$$

For a $G(K)$-category $\mathcal{C}$, the expected reconstruction map is

$$\mathrm{Wh}(\mathrm{DMod}(G(K))) \otimes_{\mathrm{BiWhit}(G)} \mathrm{Wh}(\mathcal{C}) \longrightarrow \mathcal{C}.$$

The reconstruction map is fully faithful onto the $\mathrm{DMod}(G(K))$-submodule generated by $\mathrm{Wh}(\mathcal{C})$. Thus $\mathcal{C}$ is tempered precisely when this image is all of $\mathcal{C}$.

The unramified case is more explicit. Let

$$\mathrm{Sph}_{G,x} := \mathrm{DMod}(G(\mathcal{O})\backslash G(K)/G(\mathcal{O}))$$

be the spherical Hecke category at $x \in X$. A $G(K)$-category $\mathcal{C}$ with a strong $G(\mathcal{O})$-action is controlled by its spherical invariants $\mathcal{C}^{G(\mathcal{O})}$, an $\mathrm{Sph}_{G,x}$-module. Derived Satake gives the adjunction/reconstruction picture

$$\mathrm{DMod}(G(K)/G(\mathcal{O})) \otimes_{\mathrm{Sph}_{G,x}} \mathcal{C}^{G(\mathcal{O})} \rightleftarrows \mathcal{C},$$

and the geometric Casselman–Shalika equivalence identifies the spherical-to-Whittaker comparison with the usual passage from $\mathrm{Sph}_{G,x}$ to $\mathrm{Rep}(\check{G})$. Thus, in the unramified situation, $\mathcal{C}$ is tempered exactly when the functor from spherical invariants to Whittaker invariants is conservative on the relevant generated subcategory.

For Bun_G, every $x \in X$ gives an action of $\mathrm{Sph}_{G,x}$ by Hecke modifications. The Arinkin–Gaitsgory definition of $\mathrm{DMod}(\mathsf{Bun}_G)^{\mathrm{temp}}$ is obtained from this action by passing to the tempered spectral Hecke category. A theorem of Færgeman–Raskin proves that this definition is independent of x [FR23]. There are also left- and right-tempered variants; locally they are expected to coincide, and globally this is reflected by the fact that the Whittaker coefficient does not depend on the chosen point.

13.3. **Function-theoretic shadow.** The picture is parallel to the classical automorphic one. Automorphic functions live on

$$G(F)\backslash G(\mathbb{A})/G(\mathbb{O}),$$

while Whittaker coefficients integrate against the character ψ along

$$(N(\mathbb{A}),\psi)\backslash G(\mathbb{A})/G(\mathbb{O}).$$

Geometrically, the lower quotient is replaced by the Whittaker stack of generically defined N-reductions, and the integration is replaced by the functor coeff. The conservativity theorem described in Section 13 says that, after restricting to the tempered category, this geometric Whittaker functional loses no information.

13.4. **Reduction to nilpotent singular support.** The proof begins with the comparison square supplied by the restricted/nilpotent form of categorical Langlands. Schematically, in the direction used here, it is

$$\begin{array}{ccc} \mathrm{DMod}(\mathsf{Bun}_G)^{\mathrm{temp}} & \xrightarrow{\mathrm{Coeff}^{\mathrm{ren}}} & \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}) \\ \Big\downarrow{\scriptstyle \mathrm{restr}} & & \Big\downarrow{\scriptstyle \mathrm{formal/restriction}} \\ \mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G)^{\mathrm{temp}} & \xrightarrow{\mathrm{Coeff}^{\mathrm{ren}}_{\mathrm{restr}}} & \mathrm{QCoh}(\mathsf{LocSys}^{\mathrm{restr}}_{\check{G}}). \end{array}$$

The vertical arrows should be read as the comparison or projection functors available in the chosen formalism, not as literal inclusions in one common category. Thus it is enough to understand the conservativity mechanism inside $\mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G)$, where microlocal singular support and characteristic cycles are available.

Recall

$$T^*\mathsf{Bun}_G = \mathsf{Higgs}_G = \left\{(\mathcal{P}_G, \phi) \mid \phi \in H^0(X, \mathfrak{g}^*_{\mathcal{P}_G} \otimes \Omega_X)\right\}.$$

The Hitchin map is

$$h\colon T^*\mathsf{Bun}_G \longrightarrow \mathcal{A}_G, \qquad \mathcal{A}_G := \Gamma(X, (\mathfrak{g}^* /\!\!/ G)_{\Omega_X}),$$

and the global nilpotent cone is $\mathrm{Nilp} := h^{-1}(0)$. Let $\mathrm{Nilp}_{\mathrm{irreg}} \subset \mathrm{Nilp}$ be the closed complement of the locus where the nilpotent Higgs field is generically regular. Thus the open complement is the locus of generically regular nilpotent Higgs bundles; "irregular" here means not generically regular, not irregular at every point of X.

A key microlocal theorem in the nilpotent sheaf setting [FR25, NT25] says

$$\mathrm{Shv}_{\mathrm{Nilp}_{\mathrm{irreg}}}(\mathsf{Bun}_G) = \mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G)^{\text{anti-temp}}$$

as full subcategories. It says that if the singular support never meets the generically regular nilpotent locus, then the object is invisible to the tempered quotient, and conversely.

13.5. **The Kostant component and the first coefficient.** The global Kostant section is the global analogue of a Kostant section $\mathfrak{g}^* /\!\!/ G \to \mathfrak{g}^{*,\mathrm{reg}}$ of the Chevalley map, viewed after quotienting by G. After the usual global twisting choices, its value over the zero point of the Hitchin base gives a distinguished regular nilpotent locus, denoted $f^{\mathrm{glob}} \subset \mathrm{Nilp}$. Let $\mathrm{Nilp}_{\mathrm{Kos}}$ be the unique irreducible component of Nilp containing this locus. Microlocally, $\mathrm{Nilp}_{\mathrm{Kos}}$ is the shadow of the exponential Whittaker kernel, as in the shifted-microstalk interpretation of [NT25]: composing the Lagrangian graph $d\psi \subset T^*\mathsf{Bun}_N^{\Omega}$ with the Lagrangian correspondence from Bun_N^{Ω} to Bun_G gives the global Kostant slice.

The characteristic-cycle calculation is first stated for compact constructible objects; the presentable-category conservativity statement follows by continuity and compact generation.

For constructible $\mathcal{F} \in \mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G)$, write its characteristic cycle as

$$\mathrm{CC}(\mathcal{F}) = \sum_{Z \subset \mathrm{Nilp}} c_Z(\mathcal{F})[Z], \qquad c_Z(\mathcal{F}) \in \mathbb{Z},$$

where the sum runs over irreducible components of Nilp. The Færgeman–Raskin microlocal index theorem gives [FR25]

$$\chi(\mathrm{coeff}(\mathcal{F})) = (-1)^{\dim \mathsf{Bun}_G} c_{\mathrm{Nilp}_{\mathrm{Kos}}}(\mathcal{F}).$$

Thus, if $\mathrm{Nilp}_{\mathrm{Kos}}$ occurs as an irreducible component of $\mathrm{SS}(\mathcal{F})$, equivalently with non-zero characteristic-cycle multiplicity, then $\mathrm{coeff}(\mathcal{F}) \neq 0$. The proof uses the smoothness of the regular nilpotent locus and the fact that the global Kostant section meets the zero Hitchin fiber along the distinguished locus f^{glob}.

13.6. **Other Whittaker coefficients and Hecke translates.** Let $D = \sum_i \check{\lambda}_i x_i$ be an effective $\check{\Lambda}^+$-valued divisor on X. It has an associated object $V^D \in \mathrm{Rep}(\check{G})_{\mathrm{Ran}}$. Over the locus of distinct points this is the tensor product of the representations $V_{\check{\lambda}_i}$ placed at the points x_i; on the Ran space it is the corresponding factorization extension. There is a corresponding D-shifted Whittaker stack, denoted here by $\mathsf{Bun}_N^{\Omega,D}$ to avoid committing to a sign convention, and a coefficient coeff_D. The geometric Casselman–Shalika formula says

$$\mathrm{coeff}_D(\mathcal{F}) \simeq \mathrm{coeff}(V^D \star \mathcal{F}).$$

Hence, if $\mathrm{Nilp}_{\mathrm{Kos}}$ occurs as a component of $\mathrm{SS}(V^D \star \mathcal{F})$, then $\mathrm{coeff}_D(\mathcal{F}) \neq 0$. As $\mathrm{Coeff}^{\mathrm{ren}}$ is compatible with the Hecke/spectral action, this non-vanishing implies $\mathrm{Coeff}^{\mathrm{ren}}(\mathcal{F}) \neq 0$.

The remaining geometric input is the Hecke-moving theorem of [FR25]:

$$\mathrm{SS}(\mathcal{F}) \not\subset \mathrm{Nilp}_{\mathrm{irreg}} \quad \Longrightarrow \quad \exists D \text{ such that } \mathrm{Nilp}_{\mathrm{Kos}} \text{ occurs as a component of } \mathrm{SS}(V^D \star \mathcal{F}).$$

This is proved by a local calculation with affine Springer fibers and a global spread-out argument. Informally, Hecke modifications can move any generically regular nilpotent singular direction to the Kostant component.

Now suppose $\mathcal{F} \in \mathrm{Shv}_{\mathrm{Nilp}}(\mathsf{Bun}_G)^{\mathrm{temp}}$ and assume the appropriate renormalized coefficient, namely $\mathrm{Coeff}^{\mathrm{ren}}_{\mathrm{restr}}(\mathcal{F})$ in the restricted theory vanishes. Hecke compatibility then implies that every D-coefficient vanishes. By the Hecke-moving theorem, $\mathrm{SS}(\mathcal{F}) \subset \mathrm{Nilp}_{\mathrm{irreg}}$. Hence $\mathcal{F}$ is anti-tempered. Since the intersection of the tempered and anti-tempered subcategories is zero, $\mathcal{F} = 0$. This proves conservativity. The same argument gives the de Rham statement via the comparison square above.

14. Quantization of Hitchin systems

The goal is to explain the Beilinson–Drinfeld quantization of Hitchin Hamiltonians and the resulting Hecke eigensheaves [BD91, BD05]. We keep the automorphic group G and its Langlands dual $\check{G}$, and assume that $\check{G}$ is of adjoint type. Write $\check{\mathfrak{g}} = \mathrm{Lie}(\check{G})$, fix a Borel $\check{B} \subset \check{G}$, and put

$$\mathfrak{c}_{\check{G}} := \check{\mathfrak{g}} /\!\!/ \check{G} = \operatorname{Spec} k[\check{\mathfrak{g}}]^{\check{G}}, \qquad \mathfrak{c}_{\check{G},\Omega} := (\mathfrak{c}_{\check{G}})_{\Omega_X}.$$

If $d_1, \ldots, d_r$ are the exponents of $\check{\mathfrak{g}}$, then $\mathfrak{c}_{\check{G},\Omega} \simeq \bigoplus_i \Omega_X^{d_i+1}$, and the Hitchin base is

$$\mathsf{Hitch}_{\check{G}}(X) := \Gamma(X, \mathfrak{c}_{\check{G},\Omega}).$$

The point of the section is that the same affine space appears twice: classically as the Hitchin base, and quantum mechanically as the associated graded of the center at critical level.

14.1. **Opers as a torsor over the Hitchin base.** Let $\pi\colon Y \to X$ be a D_X-scheme. A $\check{G}$-oper on Y, relative to X, in the sense of Beilinson–Drinfeld [BD05], is a pair $(\mathcal{F}_{\check{B}}, \nabla)$ consisting of a $\check{B}$-torsor on Y and a connection ∇ on the induced $\check{G}$-torsor $\mathcal{F}_{\check{G}} := \check{G} \times^{\check{B}} \mathcal{F}_{\check{B}}$, relative to π, such that the second fundamental form of the $\check{B}$-reduction satisfies the oper condition:

$$c(\nabla) \in \Gamma(Y, (\check{\mathfrak{g}}/\check{\mathfrak{b}})_{\mathcal{F}_{\check{B}}} \otimes \pi^*\Omega_X)$$

lies in the open $\check{B}$-orbit, equivalently its simple negative-root components are nowhere zero. The functor $Y \mapsto \mathsf{Op}_{\check{G}}(Y)$ is represented by a D_X-scheme, denoted $\mathsf{Op}_{\check{G}}^D \to X$. Its global sections are the ordinary opers $\mathsf{Op}_{\check{G}}(X)$.

Proposition 14.1.1. *The prestack $\mathsf{Op}_{\check{G}}(X)$ is a torsor under the vector space $\mathsf{Hitch}_{\check{G}}(X)$. Consequently, if*

$$A := \Gamma(\mathsf{Op}_{\check{G}}(X), \mathcal{O}), \qquad A^{\mathrm{cl}} := \Gamma(\mathsf{Hitch}_{\check{G}}(X), \mathcal{O}),$$

then the natural filtration on A has associated graded $\operatorname{gr} A \simeq A^{\mathrm{cl}}$.

The proof of Proposition 14.1.1 is reduced to Lemmas 14.1.2, 14.1.3, and 14.1.4.

Lemma 14.1.2. *$\mathsf{Op}_{\check{G}}(X)$ is non-empty.*

Lemma 14.1.3. *Let W be the right adjoint to the forgetful functor $\mathrm{Sch}_{/X_{\mathrm{dR}}} \to \mathrm{Sch}_{/X}$. For a vector bundle E on X, one may think of $W(E)$ as the D_X-scheme of jets, so that*

$$W(E)(D_x) = \mathrm{Jets}_x(E).$$

Then $\mathsf{Op}_{\check{G}}^D$ is a torsor under the D_X-scheme $W(\mathfrak{c}_{\check{G},\Omega})$.

Lemma 14.1.4. *Choose a principal $\mathfrak{sl}_2$-triple (e, h, f) in $\check{\mathfrak{g}}$. There is a canonical $\check{B}$-torsor $\mathcal{F}_{\check{B}}^{\mathrm{op}}$ on X, depending on this choice, such that for every $\check{G}$-oper on any $Y \to X$, the underlying $\check{B}$-torsor is canonically the pullback of $\mathcal{F}_{\check{B}}^{\mathrm{op}}$. The corresponding point of $\mathsf{Bun}_{\check{G}}$ is independent of the auxiliary choice up to connected component.*

In particular, forgetting the connection gives a constant map to $\mathsf{Bun}_{\check{G}}$:

$$\begin{array}{ccc}
\mathsf{Op}_{\check{G}}(X) & \longrightarrow & \mathsf{LocSys}_{\check{G}}(X) \\
\downarrow & & \downarrow \\
\mathrm{pt} & \xrightarrow{\ \mathcal{F}_{\check{G}}^{\mathrm{op}}\ } & \mathsf{Bun}_{\check{G}}(X).
\end{array}$$

14.2. **The local normal form.** Work locally on X, choose a coordinate t, and first suppose that the canonical $\check{B}$-torsor is trivial. In such a trivialization an oper connection has the form

$$\nabla = d + (f + \rho(t))\, dt, \qquad \rho(t) \in \check{\mathfrak{b}}(\!(t)\!),$$

and $\rho(t) \in \check{\mathfrak{b}}[\![t]\!]$ in the regular case. This form is not unique. The Drinfeld–Sokolov gauge statement for opers [BD05, Fre02] is that each $\check{N}(\!(t)\!)$-gauge class has a unique representative

$$\nabla = d + (f + v(t))\, dt, \qquad v(t) \in \check{\mathfrak{g}}^e(\!(t)\!),$$

with $v(t) \in \check{\mathfrak{g}}^e[\![t]\!]$ for regular opers. Here $\check{\mathfrak{g}}^e$ is the centralizer of e. This is Kostant's slice: the true parameters are the Kostant coordinates in $\check{\mathfrak{g}}^e$, not all of $\check{\mathfrak{b}}$.

For $\check{G} = \mathrm{PGL}_2$, the oper condition trivializes $(\check{\mathfrak{g}}/\check{\mathfrak{b}})_{\mathcal{F}_{\check{B}}} \otimes \pi^*\Omega_X$. Hence $\mathsf{Op}_{\mathrm{PGL}_2}(Y)$ is a possibly empty torsor over $\pi^*\Omega_X^{\otimes 2}$. It is non-empty for $g > 1$ because $H^1(X, \Omega_X^{\otimes 2}) = 0$, and for $g = 0, 1$ by direct construction. Therefore the three lemmas hold for PGL_2. In this case the usual statement is that projective connections form a torsor over quadratic differentials.

For general adjoint $\check{G}$, induce a PGL_2-oper along the principal homomorphism $\mathrm{PGL}_2 \to \check{G}$. This gives a $\check{G}$-oper and hence proves non-emptiness. Kostant's isomorphism gives

$$(\check{\mathfrak{g}}^e)_{\mathcal{F}^{\mathrm{op}}_{\check{B}}} \otimes \Omega_X \simeq \mathfrak{c}_{\check{G},\Omega},$$

so the same argument identifies the freedom in a general oper with the Hitchin vector bundle. This proves that $\mathsf{Op}^D_{\check{G}}$ is a $W(\mathfrak{c}_{\check{G},\Omega})$-torsor and hence that $\mathsf{Op}_{\check{G}}(X)$ is a $\mathsf{Hitch}_{\check{G}}(X)$-torsor.

14.3. **The Feigin–Frenkel center.** Let $x \in X$, $D_x = \operatorname{Spec} k[\![t]\!]$, and $D_x^\times = \operatorname{Spec} k(\!(t)\!)$. Denote by $\hat{\mathfrak{g}}_{\mathrm{crit},x}$ the affine Kac–Moody algebra at critical level and by

$$V_{\mathrm{crit},x} \coloneqq \operatorname{Ind}^{\hat{\mathfrak{g}}_{\mathrm{crit},x}}_{\mathfrak{g}[\![t]\!] \oplus k\mathbb{1}} k$$

the vacuum module. There are two related centers:

$$\mathfrak{Z}_{\mathrm{crit},x} \coloneqq Z(\widehat{\mathcal{U}}_{\mathrm{crit},x}), \qquad \mathfrak{z}_{\mathrm{crit}}(D_x) \coloneqq \operatorname{End}_{\hat{\mathfrak{g}}_{\mathrm{crit},x}}(V_{\mathrm{crit},x}).$$

The vacuum modules vary with x and form the critical chiral algebra V_{crit}. Its center $\mathfrak{z}_{\mathrm{crit}} \coloneqq Z(V_{\mathrm{crit}})$ is a commutative chiral algebra; equivalently, it is a commutative D_X-algebra in the factorization sense, so $\operatorname{Spec} \mathfrak{z}_{\mathrm{crit}}$ is a D_X-scheme.

Theorem 14.3.1 (Feigin–Frenkel, see [Fre02])**.** *There are canonical identifications*

$$\operatorname{Spf} \mathfrak{Z}_{\mathrm{crit},x} \simeq \mathsf{Op}_{\check{G}}(D_x^\times), \qquad \operatorname{Spec} \mathfrak{z}_{\mathrm{crit}}(D_x) \simeq \mathsf{Op}_{\check{G}}(D_x),$$

and globally

$$\operatorname{Spec} \mathfrak{z}_{\mathrm{crit}} \simeq \mathsf{Op}^D_{\check{G}}$$

as D_X-schemes. Under the PBW filtration, the classical limit is

$$\operatorname{Spec} \operatorname{gr} \mathfrak{z}_{\mathrm{crit}} \simeq W(\mathfrak{c}_{\check{G},\Omega}).$$

The proof idea is as follows. First, the chiral-algebra nature of V_{crit} makes $\mathfrak{z}_{\mathrm{crit}}$ factorizable; locally this produces a map from the formal spectrum of the ordinary critical center to the punctured-disc center. The Theorem 14.3.1 says this map is an isomorphism. Thus it remains to identify $\operatorname{Spec} \mathfrak{z}_{\mathrm{crit}}(D)$ with regular opers on D, compatibly with $\operatorname{Aut}(D)$.

The regular map $\operatorname{Spec} \mathfrak{z}_{\mathrm{crit}}(D) \to \mathsf{Op}_{\check{G}}(D)$ is constructed Tannakianly. One needs a $\check{G}$-torsor on $\operatorname{Spec} \mathfrak{z}_{\mathrm{crit}}(D) \times D$, a connection, and a $\check{B}$-reduction satisfying the oper condition, all $\operatorname{Aut}(D)$-equivariant. Geometric Satake and critical global sections give a symmetric monoidal functor

$$F_D \colon \operatorname{Rep}(\check{G}) \longrightarrow \mathfrak{z}_{\mathrm{crit}}(D)\text{-Mod}$$

characterized by

$$\Gamma(\mathsf{Gr}_G, \operatorname{Sat}(V))_{\mathrm{crit}} \simeq F_D(V) \otimes_{\mathfrak{z}_{\mathrm{crit}}(D)} V_{\mathrm{crit}}.$$

By Tannaka duality, F_D determines a $\check{G}$-torsor over $\operatorname{Spec} \mathfrak{z}_{\mathrm{crit}}(D)$. The $\check{B}$-reduction is obtained from lowest rotation weights. For the irreducible $\check{G}$-representation $V^{\check{\lambda}}$, let $\mathrm{IC}_{\check{\lambda}} = \operatorname{Sat}(V^{\check{\lambda}})$. If $L_0 = -t\partial_t \in \operatorname{Lie} \operatorname{Aut}^0(D)$, then the lowest L_0-eigenvalue of $\Gamma(\mathsf{Gr}_G, \mathrm{IC}_{\check{\lambda}})_{\mathrm{crit}}$ is $-\langle \check{\lambda}, \rho \rangle$; the corresponding eigenspace is one-dimensional and $\check{B}$-invariant. These lines satisfy the Plucker relations, hence give the desired $\check{B}$-reduction. The last step is that the resulting map to $\mathsf{Op}_{\check{G}}(D)$ is an isomorphism.

14.4. **Global differential operators and the quantum Hitchin map.** Let $\mathcal{D}_{\mathrm{crit}}$ be the sheaf of critically twisted differential operators on Bun_G. The local critical center acts on global differential operators by localization. Locally, one uses the standard construction: if a pair $(\mathfrak{h}, K)$ acts on a smooth stack Y and $y \in Y$ is a K-fixed point, then invariant differential operators on the homogeneous model map to differential operators on Y:

$$(U(\mathfrak{h})/U(\mathfrak{h})\mathfrak{k})^K \longrightarrow \Gamma(Y, \mathcal{D}_Y).$$

For $Y = \mathsf{Bun}_G$, $\mathfrak{h} = \mathfrak{g}(\!(t)\!)$, and $K = G[\![t]\!]$, this gives

$$\mathfrak{z}_{\mathrm{crit}}(D_x) = \mathrm{End}(V_{\mathrm{crit},x}) \longrightarrow \Gamma(\mathsf{Bun}_G, \mathcal{D}_{\mathrm{crit}}).$$

These local maps are compatible with the connection in the variable x, hence pass to horizontal sections:

$$\mathfrak{z}_{\mathrm{crit}}(X) \longrightarrow \Gamma(\mathsf{Bun}_G, \mathcal{D}_{\mathrm{crit}}).$$

Taking associated graded gives the classical Hitchin map

$$\mathrm{gr}\, \mathfrak{z}_{\mathrm{crit}}(X) \longrightarrow \Gamma(T^*\mathsf{Bun}_G, \mathcal{O}).$$

Using the Feigin–Frenkel identification and the torsor statement above, set

$$A := \Gamma(\mathsf{Op}_{\check{G}}(X), \mathcal{O}) = \Gamma(X_{\mathrm{dR}}, \mathfrak{z}_{\mathrm{crit}}),$$
$$A^{\mathrm{cl}} := \Gamma(\mathsf{Hitch}_{\check{G}}(X), \mathcal{O}) \simeq \mathrm{gr}\, A.$$

Theorem 14.5.1 says that the preceding maps are isomorphisms [BD91]; for a positive-characteristic approach in the GL_n direction see [BCTZ16]:

$$A \xrightarrow{\sim} \Gamma(\mathsf{Bun}_G, \mathcal{D}_{\mathrm{crit}}), \qquad A^{\mathrm{cl}} \xrightarrow{\sim} \Gamma(T^*\mathsf{Bun}_G, \mathcal{O}).$$

Thus the commutative algebra of quantum Hitchin Hamiltonians is the algebra of functions on the space of $\check{G}$-opers. The second isomorphism is the classical statement that global functions on $T^*\mathsf{Bun}_G$ come from the Hitchin base.

14.5. **The automorphic module attached to an oper.** Let $\sigma \in \mathsf{Op}_{\check{G}}(X)$. Via the inclusion $\mathsf{Op}_{\check{G}}(X) \subset \mathsf{LocSys}_{\check{G}}(X)$, it has an underlying $\check{G}$-local system, still denoted σ. Let k_σ be the corresponding character of $A = \Gamma(\mathsf{Op}_{\check{G}}(X), \mathcal{O})$. Define the critically twisted D-module

$$\mathrm{Aut}_\sigma := \mathcal{D}_{\mathrm{crit}} \otimes_{A,\sigma} k_\sigma \in \mathrm{DMod}(\mathsf{Bun}_G).$$

Theorem 14.5.1 (Beilinson–Drinfeld [BD91])**.** *Aut_σ is a Hecke eigensheaf with eigenvalue σ. Equivalently, for every $x \in X$ and every $V \in \mathrm{Rep}(\check{G})$,*

$$\mathcal{H}_{V,x}(\mathrm{Aut}_\sigma) \simeq V_{\sigma,x} \otimes \mathrm{Aut}_\sigma,$$

where $\mathcal{H}_{V,x}$ is the Hecke functor corresponding to V, and $V_{\sigma,x}$ is the fiber at x of the local system associated to V and σ.

The proof uses Beilinson–Drinfeld localization

$$\mathrm{Loc}_x \colon \hat{\mathfrak{g}}_{\mathrm{crit},x}\text{-}\mathrm{Mod}^{G[\![t]\!]} \longrightarrow \mathrm{DMod}_{\mathrm{crit}}(\mathsf{Bun}_G)$$

with the property

$$\mathrm{Loc}_x(V_{\mathrm{crit},x}) \simeq \mathcal{D}_{\mathrm{crit}}.$$

For $V \in \mathrm{Rep}(\check{G})$, the Hecke action is compatible with localization:

$$\begin{aligned}\mathrm{Sat}_x(V) \star \mathcal{D}_{\mathrm{crit}} &\simeq \mathrm{Loc}_x(\mathrm{Sat}_x(V) \star V_{\mathrm{crit},x}) \\ &\simeq \mathrm{Loc}_x\left(F_{D_x}(V) \otimes_{\mathfrak{z}_{\mathrm{crit}}(D_x)} V_{\mathrm{crit},x}\right).\end{aligned}$$

Let $\mathcal{F}$ be the universal $\check{G}$-torsor, or equivalently the universal oper local system, on $\mathsf{Op}_{\check{G}}(X)$. The functor F_{D_x} is the restriction of this universal torsor to the formal disk, so the last display becomes

$$\mathrm{Sat}_x(V) \star \mathcal{D}_{\mathrm{crit}} \simeq V_{\mathcal{F},x} \otimes_A \mathcal{D}_{\mathrm{crit}}.$$

Tensoring over A with k_σ gives

$$\mathrm{Sat}_x(V) \star \mathrm{Aut}_\sigma \simeq V_{\sigma,x} \otimes \mathrm{Aut}_\sigma.$$

This is exactly the Hecke eigenproperty.

14.6. **Vacuum generation and the next local picture.** The construction in Subsection 14.5 only uses the part of the critical Kac–Moody category generated by the vacuum module. Schematically, one has

$$\begin{array}{ccc} \mathrm{KL}^{V\text{-gen}}_{\mathrm{crit},x} & \xrightarrow{\mathrm{Loc}_x} & \mathrm{DMod}_{\mathrm{crit}}(\mathsf{Bun}_G) \\ \uparrow & & \uparrow \\ \mathrm{QCoh}(\mathsf{Op}_{\check{G}}(D_x)) & \longrightarrow & \mathrm{QCoh}(\mathsf{Op}_{\check{G}}(X)), \end{array}$$

where V-generated means generated by the vacuum object. The next step is to replace this vacuum-generated subcategory by the full critical category. The expected local form is the fundamental local equivalence. Put

$$\mathsf{Op}^{\mathrm{unr}}_{\check{G},x} := \mathsf{Op}_{\check{G}}(D_x^{\times}) \times_{\mathsf{LocSys}_{\check{G}}(D_x^{\times})} \mathsf{LocSys}_{\check{G}}(D_x).$$

Then, up to the usual $*$- versus !-IndCoh convention,

$$\mathrm{KL}_{G,\mathrm{crit},x} \simeq \mathrm{IndCoh}(\mathsf{Op}^{\mathrm{unr}}_{\check{G},x}).$$

Localization should enhance to a relative construction of the form

$$\mathrm{KL}_{G,\mathrm{crit},x} \otimes_{\mathrm{QCoh}(\mathsf{Op}^{\mathrm{unr}}_{\check{G},x})} \mathrm{QCoh}(\mathsf{Op}_{\check{G}}(X \setminus \{x\})^{\mathrm{unr}}) \longrightarrow \mathrm{DMod}_{\mathrm{crit}}(\mathsf{Bun}_G),$$

where we define

$$\mathsf{Op}_{\check{G}}(X \setminus \{x\})^{\mathrm{unr}} := \mathsf{Op}_{\check{G}}(X \setminus \{x\}) \times_{\mathsf{LocSys}_{\check{G}}(X \setminus \{x\})} \mathsf{LocSys}_{\check{G}}(X).$$

This is the bridge from the Beilinson–Drinfeld eigensheaves for opers to the broader critical-level local theory of [FG05, FG07].

15. Fundamental local equivalence

This section explains the local input behind the global localization arguments [FG05, FG07, ABC+24a]. The slogan is that critical-level spherical Kac–Moody representations are the same as sheaves on the space of unramified, or monodromy-free, opers. This upgrades the vacuum-generated picture from Section 14 to the full Kazhdan–Lusztig category.

15.1. **From Beilinson–Drinfeld to a local equivalence.** Recall the two inputs used in the proof of the Hecke eigenproperty for Aut_σ. First, for a point $x \in X$, the Beilinson–Drinfeld localization functor gives

$$\mathrm{Loc}_x \colon \hat{\mathfrak{g}}_{\mathrm{crit},x}\text{-Mod}^{\mathfrak{L}^+G_x} \longrightarrow \mathrm{DMod}_{\mathrm{crit}}(\mathsf{Bun}_G),$$

and the center acts by pulling functions from local opers to global opers:

$$\begin{array}{ccc} \hat{\mathfrak{g}}_{\mathrm{crit},x}\text{-Mod}^{V\text{-gen}} & \xrightarrow{\mathrm{Loc}_x} & \mathrm{DMod}_{\mathrm{crit}}(\mathsf{Bun}_G) \\ \uparrow & & \uparrow \\ \mathrm{QCoh}(\mathsf{Op}_{\check{G}}(D_x)) & \longrightarrow & \mathrm{QCoh}(\mathsf{Op}_{\check{G}}(X)). \end{array}$$

Here the superscript V-generated means the subcategory generated by the vacuum module. Second, for $V \in \mathrm{Rep}(\check{G})$, geometric Satake and critical global sections give

$$\mathrm{Sat}_x(V) \star \mathbb{V}_{\mathrm{crit},x} \simeq F_x(V) \otimes_{\mathfrak{z}_{\mathrm{crit}}(D_x)} \mathbb{V}_{\mathrm{crit},x},$$

where $F_x(V)$ is the vector bundle on $\mathsf{Op}_{\check{G}}(D_x) \simeq \mathrm{Spec}\, \mathfrak{z}_{\mathrm{crit}}(D_x)$ attached to the universal $\check{G}$-local system. Fundamental local equivalence replaces this vacuum-generated statement by an equivalence for the whole spherical critical category.

15.2. **Unramified opers.** Let $D = \operatorname{Spec} k[\![t]\!]$, $D^\times = \operatorname{Spec} k(\!(t)\!)$, $\mathfrak{L}^+G = G[\![t]\!]$, and $\mathfrak{L}G = G(\!(t)\!)$. Write

$$\mathrm{KL}_{G,\mathrm{crit}} \coloneqq \hat{\mathfrak{g}}_{\mathrm{crit}}\text{-}\mathrm{Mod}^{\mathfrak{L}^+G}$$

for the critical Kazhdan–Lusztig category. The spectral space in the FLE is not the space of regular opers on D, but the larger fiber product

$$\mathsf{Op}^{\mathrm{unr}}_{\check{G}} \coloneqq \mathsf{Op}_{\check{G}}(D^\times) \times_{\mathsf{LocSys}_{\check{G}}(D^\times)} \mathsf{LocSys}_{\check{G}}(D).$$

Thus a point is an oper on the punctured disk whose underlying $\check{G}$-local system extends over D. The map $\mathsf{Op}_{\check{G}}(D) \to \mathsf{Op}^{\mathrm{unr}}_{\check{G}}$ is not an isomorphism: extending the local system does not force the oper reduction to extend regularly.

Concretely, after choosing a trivialization, an oper on $D^\times$ has the form $\nabla = d+\omega$, with $\omega \in \check{\mathfrak{g}}(\!(t)\!)dt$. Being unramified means that after some $\check{G}(\!(t)\!)$-gauge transformation it becomes $d+\eta$, with $\eta \in \check{\mathfrak{g}}[\![t]\!]dt$. It is not necessary that the same gauge makes $\eta \bmod \check{\mathfrak{b}}[\![t]\!]$ generic at $t = 0$. The defect is measured by a dominant weight $\check{\lambda} \in X^*(\check{T})^+$, equivalently by a dominant coweight of G. On the corresponding regular stratum one can write

$$\eta \bmod \check{\mathfrak{b}}[\![t]\!] = \sum_{i\in I} t^{\langle \check{\lambda}, \check{\alpha}_i\rangle} \psi_i(t) f_i \, dt,$$

with $\psi_i \in k[\![t]\!]$ such that $\psi_i(0) \neq 0$, where α_i are the simple roots of $\check{\mathfrak{g}}$ and $f_i \in \check{\mathfrak{g}}_{-\alpha_i}$ are simple negative-root vectors. Denote this locus by $\mathsf{Op}^{\mathrm{reg},\check{\lambda}}_{\check{G}}$. Set-theoretically,

$$\mathsf{Op}^{\mathrm{unr}}_{\check{G}}(k) = \bigcup_{\check{\lambda}\in\check{\Lambda}^+} \mathsf{Op}^{\mathrm{reg},\check{\lambda}}_{\check{G}}(k).$$

More precisely, $\mathsf{Op}^{\mathrm{unr}}_{\check{G}}$ has components, or formal pieces, $\mathsf{Op}^{\mathrm{unr},\check{\lambda}}_{\check{G}}$, and

$$\mathsf{Op}^{\mathrm{reg},\check{\lambda}}_{\check{G}} \subset (\mathsf{Op}^{\mathrm{unr},\check{\lambda}}_{\check{G}})_{\mathrm{red}}.$$

Theorem 15.2.1 (critical FLE, pointwise form [FG07, ABC+24a])**.** *There is an equivalence of DG categories*

$$\mathrm{FLE}_{G,\mathrm{crit}} \colon \hat{\mathfrak{g}}_{\mathrm{crit}}\text{-}\mathrm{Mod}^{\mathfrak{L}^+G} \xrightarrow{\sim} \mathrm{IndCoh}(\mathsf{Op}^{\mathrm{unr}}_{\check{G}}),$$

t-exact for the standard Frenkel–Gaitsgory abelian t-structures, with the conventional $*$- or $!$-IndCoh choice suppressed. The equivalence is compatible with the geometric Satake action of the spherical category.

This compatibility is visible on Weyl objects. Let $W^{\check{\lambda}}_{\mathrm{crit}}$ be the critical Weyl object corresponding to the Schubert label $\check{\lambda}$, and let $i_{\check{\lambda}} \colon \mathsf{Op}^{\mathrm{reg},\check{\lambda}}_{\check{G}} \hookrightarrow \mathsf{Op}^{\mathrm{unr}}_{\check{G}}$. Drinfeld–Sokolov reduction gives

$$\Psi(W^{\check{\lambda}}_{\mathrm{crit}}) \simeq \mathfrak{z}^{\mathrm{reg},\check{\lambda}} \simeq \mathrm{End}(W^{\check{\lambda}}_{\mathrm{crit}}), \qquad \operatorname{Spec} \mathfrak{z}^{\mathrm{reg},\check{\lambda}} \simeq \mathsf{Op}^{\mathrm{reg},\check{\lambda}}_{\check{G}}.$$

Under FLE this says, up to the IndCoh convention, that $W^{\check{\lambda}}_{\mathrm{crit}}$ corresponds to $(i_{\check{\lambda}})^{\mathrm{IndCoh}}_* \mathcal{O}_{\mathsf{Op}^{\mathrm{reg},\check{\lambda}}_{\check{G}}}$. Thus the support of the spherical critical category is exactly the unramified oper locus; unramified on the automorphic side corresponds to monodromy-free on the spectral side.

15.3. **Drinfeld–Sokolov reduction and the Whittaker form.** The functor underlying the FLE is Drinfeld–Sokolov reduction [FG05, Fre02]. Choose the standard non-degenerate additive character $\chi \colon \mathfrak{n}(\!(t)\!) \to k$. For a critical module M, put

$$\Psi(M) \coloneqq C^{\frac{\infty}{2}}(\mathfrak{n}(\!(t)\!), M \otimes k_\chi).$$

This is the semi-infinite cohomology, or BRST reduction, of M. At critical level it has an enhanced version landing in sheaves over opers on the punctured disk. The corresponding Whittaker statement is read as

$$\hat{\mathfrak{g}}_{\mathrm{crit}}\text{-}\mathrm{Mod}^{\mathrm{Whit}} \simeq \mathrm{IndCoh}(\mathsf{Op}_{\check{G}}(D^\times)).$$

Passing from the Whittaker form to the spherical/Kazhdan–Lusztig subcategory amounts spectrally to imposing that the underlying local system extend over D, hence to replacing $\mathsf{Op}(D^\times)$ by $\mathsf{Op}^{\mathrm{unr}}$.

The Satake action has the parallel spectral description. The critical spherical Hecke category is identified with ind-coherent sheaves on the local-system Steinberg stack

$$\mathrm{DMod}(\mathfrak{L}^+G\backslash\mathfrak{L}G/\mathfrak{L}^+G)_{\mathrm{crit}} \simeq \mathrm{IndCoh}(\mathsf{LocSys}_{\check{G}}(D) \times_{\mathsf{LocSys}_{\check{G}}(D^\times)} \mathsf{LocSys}_{\check{G}}(D)).$$

Under FLE, convolution by this category becomes the natural correspondence action on $\mathrm{IndCoh}(\mathsf{Op}^{\mathrm{unr}}_{\check{G}})$. In the original Frenkel–Gaitsgory proof this appears at the level of abelian categories and abelian geometric Satake; the DG/factorization formulation is the version needed for global Langlands.

15.4. **Relation with quantum local Langlands.** The critical FLE is a boundary value of quantum local geometric Langlands [FG05]. Away from the boundary, dual levels are related by

$$\kappa = \kappa_{G,\mathrm{crit}} + c\,\kappa_{\mathrm{Kil},G}, \qquad \check{\kappa} = \kappa_{\check{G},\mathrm{crit}} + c^{-1}\kappa_{\mathrm{Kil},\check{G}}.$$

For $c \neq 0$, the expected local equivalence has the form

$$\hat{\mathfrak{g}}_\kappa\text{-Mod}^{\mathfrak{L}^+G} \simeq \mathrm{Whit}(D_{\check{\kappa}}(\mathsf{Gr}_{\check{G}})).$$

Its boundary shadows are the standard Satake and Casselman–Shalika equivalences. In the large-level limit the spherical and Whittaker Grassmannian categories degenerate to

$$\mathcal{D}(\mathsf{Gr}_G)^{\mathfrak{L}^+G}_{\mathrm{sph}} \simeq \mathrm{Rep}(\check{G}), \qquad \mathrm{Whit}(\mathcal{D}(\mathsf{Gr}_{\check{G}})) \simeq \mathrm{Rep}(G).$$

As $c \to 0$, the G-side becomes critical and the $\check{G}$-Whittaker side becomes the classical spectral geometry of opers; the limiting equivalence is the FLE

$$\hat{\mathfrak{g}}_{\mathrm{crit}}\text{-Mod}^{\mathfrak{L}^+G} \simeq \mathrm{IndCoh}(\mathsf{Op}^{\mathrm{unr}}_{\check{G}}).$$

This is one row in the broader quantum local dictionary:

$$\begin{array}{ccc}
\mathfrak{L}G\text{-Mod}_\kappa & \longleftrightarrow & \mathfrak{L}\check{G}\text{-Mod}_{\check{\kappa}}, \\
\hat{\mathfrak{g}}_\kappa\text{-Mod} & \longleftrightarrow & \mathrm{Whit}(D_{\check{\kappa}}(\mathfrak{L}\check{G})), \\
\mathrm{Whit}(D_\kappa(\mathfrak{L}G)) & \longleftrightarrow & \widehat{\check{\mathfrak{g}}}_{\check{\kappa}}\text{-Mod}, \\
D_\kappa(\mathsf{Gr}_G) & \longleftrightarrow & D_{\check{\kappa}}(\mathsf{Gr}_{\check{G}}), \\
D_\kappa(\mathrm{Fl}_G) & \longleftrightarrow & D_{\check{\kappa}}(\mathrm{Fl}_{\check{G}}).
\end{array}$$

At $\kappa = \kappa_{\mathrm{crit}}$, the full local $\mathfrak{L}G$-module theory is expected to be a sheaf of categories on $\mathsf{LocSys}_{\check{G}}(D^\times)$:

$$\mathfrak{L}G\text{-Mod}_{\mathrm{crit}} \rightsquigarrow \mathrm{ShvCat}(\mathsf{LocSys}_{\check{G}}(D^\times)).$$

The Kac–Moody category is therefore a category over the oper locus: its critical center is functions on $\mathsf{Op}_{\check{G}}(D^\times)$. After Whittaker/Drinfeld–Sokolov reduction, this category over opers is identified with

$$\mathrm{IndCoh}(\mathsf{Op}_{\check{G}}(D^\times)),$$

which is consistent with the Feigin–Frenkel center $\mathrm{Spec}\,\mathfrak{Z}_{\mathrm{crit}} \simeq \mathsf{Op}_{\check{G}}(D^\times)$.

15.5. **Factorization FLE and compatibility with global localization.** The pointwise equivalence in Theorem 15.2.1 has a factorization version over the Ran space [ABC+24a]:

$$\mathrm{FLE}_{G,\mathrm{crit},\mathrm{Ran}} \colon \mathrm{KL}_{G,\mathrm{crit},\mathrm{Ran}} \xrightarrow{\sim} \mathrm{IndCoh}(\mathsf{Op}^{\mathrm{unr}}_{\check{G}})_{\mathrm{Ran}}.$$

It is compatible with the localization functor to Bun_G and with the global spectral functor. The basic diagram is

$$\begin{array}{ccc}
\mathrm{KL}_{G,\mathrm{crit},\mathrm{Ran}} & \xrightarrow[\sim]{\mathrm{FLE}_{\mathrm{crit}}} & \mathrm{IndCoh}(\mathsf{Op}^{\mathrm{unr}}_{\check{G}})_{\mathrm{Ran}} \\
\Big\downarrow{\scriptstyle \mathrm{Loc}} & & \Big\downarrow{\scriptstyle \mathrm{Poinc}^{\mathrm{spec}}} \\
\mathrm{DMod}_{\mathrm{crit}}(\mathsf{Bun}_G) & \xrightarrow{\mathbb{L}} & \mathrm{IndCoh}(\mathsf{LocSys}_{\check{G}}(X)).
\end{array}$$

Here $\mathbb{L}$ denotes the global Langlands functor. The functor $\mathrm{Poinc}^{\mathrm{spec}}$ is the push–pull transform along the correspondence which compares a global local system with its formal-disk restrictions. For a point x, write

$$\mathsf{Op}_{\check{G}}(X \setminus \{x\})^{\mathrm{unr}} := \mathsf{Op}_{\check{G}}(X \setminus \{x\}) \times_{\mathsf{LocSys}_{\check{G}}(X\setminus\{x\})} \mathsf{LocSys}_{\check{G}}(X)$$

and the relevant spectral correspondence is

$$\mathsf{Op}_{\check{G}}(X \setminus \{x\})^{\mathrm{unr}} \longrightarrow \mathsf{Op}_{\check{G}}(D_x^\times) \times_{\mathsf{LocSys}_{\check{G}}(D_x^\times)} \mathsf{LocSys}_{\check{G}}(D_x).$$

Later, the proof that $\mathbb{L}$ is an equivalence uses this diagram together with the coefficient functor and the infinite-level FLE:

$$\begin{array}{ccc} \mathrm{DMod}_{\mathrm{crit}}(\mathsf{Bun}_G) & \xrightarrow{\mathbb{L}} & \mathrm{IndCoh}(\mathsf{LocSys}_{\check{G}}(X)) \\ \downarrow{\scriptstyle \mathrm{coeff}} & & \downarrow{\scriptstyle \Gamma^{\mathrm{spec}}} \\ \mathrm{Whit}(\mathsf{Gr}_G)_{\mathrm{crit,Ran}} & \xrightarrow[\sim]{\mathrm{FLE}_\infty} & \mathrm{Rep}(\check{G})_{\mathrm{Ran}}. \end{array}$$

15.6. **Many points and generation.** For a finite set of points $\underline{x} = \{x_1, \ldots, x_n\}$, factorization gives

$$\mathrm{KL}_{G,\mathrm{crit},\underline{x}} \simeq \bigotimes_{i=1}^{n} \hat{\mathfrak{g}}_{\mathrm{crit},x_i}\text{-}\mathrm{Mod}^{\mathfrak{L}^+G_{x_i}}.$$

The usual localization functor

$$\mathrm{Loc}_{\underline{x}} \colon \mathrm{KL}_{G,\mathrm{crit},\underline{x}} \longrightarrow \mathrm{DMod}_{\mathrm{crit}}(\mathsf{Bun}_G)$$

can be enhanced by adding global spectral parameters on the punctured curve. More precisely, writing $\mathsf{Op}^{\mathrm{unr}}_{\check{G},\underline{x}} := \prod_i \mathsf{Op}^{\mathrm{unr}}_{\check{G},x_i}$, the enhanced functor has the schematic form

$$\mathrm{KL}_{G,\mathrm{crit},\underline{x}} \otimes_{\mathrm{QCoh}(\mathsf{Op}^{\mathrm{unr}}_{\check{G},\underline{x}})} \mathrm{QCoh}(\mathsf{Op}_{\check{G}}(X \setminus \underline{x})^{\mathrm{unr}}) \longrightarrow \mathrm{DMod}_{\mathrm{crit}}(\mathsf{Bun}_G),$$

where

$$\mathsf{Op}_{\check{G}}(X \setminus \underline{x})^{\mathrm{unr}} := \mathsf{Op}_{\check{G}}(X \setminus \underline{x}) \times_{\mathsf{LocSys}_{\check{G}}(X \setminus \underline{x})} \mathsf{LocSys}_{\check{G}}(X).$$

Equivalently, the action of $\mathrm{Rep}(\check{G})_{\underline{x}}$ on the image of $\mathrm{Loc}_{\underline{x}}$ factors through $\mathsf{LocSys}_{\check{G}}(X)$.

As $\underline{x}$ becomes large, these local categories generate the cuspidal part. A useful schematic form is: for every quasi-compact open $U \subset \mathsf{Bun}_G$, the image of

$$\mathrm{KL}_{G,\mathrm{crit,Ran}} \longrightarrow \mathrm{DMod}_{\mathrm{crit}}(\mathsf{Bun}_G) \longrightarrow \mathrm{DMod}_{\mathrm{crit}}(U)$$

generates the cuspidal contribution after restriction to U. The remaining part of the global argument is Eisenstein: one handles non-cuspidal objects by induction and by compatibility of FLE with Eisenstein series and constant-term functors.

16. Proof of the global conjecture

This section proves the global categorical Langlands theorem after the local input of Section 15, following the modern parabolic-induction input [CCF+24]. We suppress the usual half-density twist and write $\mathrm{DMod}(\mathsf{Bun}_G)$ for the automorphic category used above. The global theorem asserts that the Langlands functor is an equivalence

$$\mathbb{L}_G \colon \mathrm{DMod}(\mathsf{Bun}_G) \xrightarrow{\sim} \mathrm{IndCoh}_{\mathrm{Nilp}}(\mathsf{LocSys}_{\check{G}})$$

compatible with derived Satake, Eisenstein series, constant term functors, and duality. The proof first treats the parabolic part (i.e. Eisenstein part) and then proves the remaining cuspidal equivalence.

16.1. **The coarse Langlands functor.** The first construction produces the *coarse* shadow, via the construction of the Langlands functor in [GR24a],

$$\mathbb{L}^0_G \colon \mathrm{DMod}(\mathsf{Bun}_G) \longrightarrow \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}).$$

It is characterized by the generalized vanishing theorem: the local Hecke action of $\mathrm{Rep}(\check{G})_{\mathrm{Ran}}$ on $\mathrm{DMod}(\mathsf{Bun}_G)$ factors through the spectral action of $\mathrm{QCoh}(\mathsf{LocSys}_{\check{G}})$. In the notation of Sections 13–15, the Whittaker coefficient functor and the spectral action give the diagram

$$\begin{array}{ccc} \mathrm{DMod}(\mathsf{Bun}_G) & \xrightarrow{\mathbb{L}^0_G} & \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}) \\ \downarrow{\scriptstyle \mathrm{coeff}_G} & & \downarrow{\scriptstyle \Gamma^{\mathrm{spec}}} \\ \mathrm{Whit}(G)_{\mathrm{Ran}} & \xrightarrow{\mathrm{CS}} & \mathrm{Rep}(\check{G})_{\mathrm{Ran}}. \end{array}$$

Here CS denotes the geometric Casselman–Shalika equivalence. The full Langlands functor is a lift of $\mathbb{L}^0_G$ through the forgetful functor

$$\Psi \colon \mathrm{IndCoh}_{\mathrm{Nilp}}(\mathsf{LocSys}_{\check{G}}) \longrightarrow \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}),$$

i.e. $\mathbb{L}_G^0 = \Psi \circ \mathbb{L}_G$. On compact objects one first lands in a bounded-below coherent subcategory; the presentable functor $\mathbb{L}_G$ is obtained by ind-completion. It is convenient to write $\mathrm{QCoh}^!(\mathsf{LocSys}_{\check{G}})$ for the corresponding renormalized spectral target. Thus the shorthand is

$$\mathrm{DMod}(\mathsf{Bun}_G) \xrightarrow{\mathbb{L}_G} \mathrm{QCoh}^!(\mathsf{LocSys}_{\check{G}}).$$

Compatibility with derived Satake is built into the construction. The new work is compatibility with Eisenstein and constant term functors [CCF+24].

16.2. **Eisenstein and constant term compatibility.** Let $P \subset G$ be a parabolic with Levi quotient M, and let $\check{P} \subset \check{G}$ be the dual parabolic. The automorphic Eisenstein and constant term functors are, up to the standard shifts and twists,

$$\mathrm{Eis}_P\colon \mathrm{DMod}(\mathsf{Bun}_M) \longrightarrow \mathrm{DMod}(\mathsf{Bun}_G),$$
$$\mathrm{CT}_P\colon \mathrm{DMod}(\mathsf{Bun}_G) \longrightarrow \mathrm{DMod}(\mathsf{Bun}_M).$$

On the spectral side they are defined by the correspondence

$$\mathsf{LocSys}_{\check{M}} \xleftarrow{q} \mathsf{LocSys}_{\check{P}} \xrightarrow{p} \mathsf{LocSys}_{\check{G}}.$$

The target statement is the commutativity of

$$\begin{array}{ccc} \mathrm{DMod}(\mathsf{Bun}_M) & \xrightarrow{\mathbb{L}_M} & \mathrm{QCoh}^!(\mathsf{LocSys}_{\check{M}}) \\ \downarrow{\scriptstyle \mathrm{Eis}_P} & & \downarrow{\scriptstyle \mathrm{Eis}^{\mathrm{spec}}_{\check{P}}} \\ \mathrm{DMod}(\mathsf{Bun}_G) & \xrightarrow{\mathbb{L}_G} & \mathrm{QCoh}^!(\mathsf{LocSys}_{\check{G}}) \end{array} \tag{16.1}$$

and

$$\begin{array}{ccc} \mathrm{DMod}(\mathsf{Bun}_G) & \xrightarrow{\mathbb{L}_G} & \mathrm{QCoh}^!(\mathsf{LocSys}_{\check{G}}) \\ \downarrow{\scriptstyle \mathrm{CT}_P} & & \downarrow{\scriptstyle \mathrm{CT}^{\mathrm{spec}}_{\check{P}}} \\ \mathrm{DMod}(\mathsf{Bun}_M) & \xrightarrow{\mathbb{L}_M} & \mathrm{QCoh}^!(\mathsf{LocSys}_{\check{M}}). \end{array} \tag{16.2}$$

The proof of (16.1), in the form of [CCF+24], is reduced to two squares: a geometric square involving Whittaker coefficients and a spectral square involving representations. Under geometric Casselman–Shalika, $\mathrm{Whit}(H)_x \simeq \mathrm{Rep}(\check{H})_x$, so the comparison becomes a compatibility between parabolic restriction of Whittaker coefficients and the corresponding Chevalley functor on representations.

16.3. **The spectral square: push versus localization.** The basic problem is the following. For a homomorphism $H_1 \to H_2$, one has $\mathsf{LocSys}_{H_1} \to \mathsf{LocSys}_{H_2}$, and for a point $x \in X$ one also has local localization functors

$$\mathrm{Loc}_x\colon \mathrm{Rep}(H_i)_x \longrightarrow \mathrm{QCoh}^!(\mathsf{LocSys}_{H_i}).$$

There is no naive square comparing pullback along $\mathsf{LocSys}_{H_1} \to \mathsf{LocSys}_{H_2}$ with restriction of representations. The correct square uses *pushforward* and Ran insertions.

Lemma 16.3.1. *Assume $H_1 \to H_2$ has unipotent kernel. Let* $\mathrm{Ran}_{x,\mathrm{c}} := \{I \in \mathrm{Ran}(X) \mid x \in I\}$. *Then the following diagram commutes.*

$$\begin{array}{ccc} \mathrm{QCoh}^!(\mathsf{LocSys}_{H_1}) & \xleftarrow{\mathrm{Loc}_x} & \mathrm{Rep}(H_1)_x \\ \downarrow{\scriptstyle \mathrm{push}} & & \downarrow{\scriptstyle \mathrm{inv}^+} \\ \mathrm{QCoh}^!(\mathsf{LocSys}_{H_2}) & \xleftarrow{\mathrm{Loc}_{\mathrm{Ran}_{x,\mathrm{c}}}} & \mathrm{Rep}(H_2)_{\mathrm{Ran}_{x,\mathrm{c}}} \end{array} \tag{16.3}$$

Here inv^+ first inserts the object at the marked point and then takes derived invariants, equivalently Chevalley cochains, with respect to the unipotent kernel. Explicitly, if $V_x \in \mathrm{Rep}(H_1)_x$, then at a point $(x, y_1, \ldots, y_m) \in \mathrm{Ran}_{x,\mathrm{c}}$ the insertion is

$$\mathrm{ins}(V_x)_{(x,y_1,\ldots,y_m)} = V_x \otimes \mathbb{1}_{y_1} \otimes \cdots \otimes \mathbb{1}_{y_m}.$$

The proof is a Ran-space form of the following elementary principle. If $\pi\colon Y \to X$ is a D-scheme and Y is quasi-affine, coaffine, or satisfies the same descent condition, then

$$\mathrm{QCoh}(Y) \simeq \pi_* \mathcal{O}_Y\text{-}\mathrm{Mod}(\mathrm{QCoh}(X)),$$

where $\pi_*\mathcal{O}_Y$ is generally a complex. Examples such as $\mathbb{A}^2 \setminus \{0\} \to \mathrm{pt}$ and $\mathbb{B}\mathbb{G}_{\mathrm{a}} \to \mathrm{pt}$ illustrate why one must keep this complex, not only its degree-zero functions. Applying this to the formal disk and using the compatibility of insertions with the Ran convolution gives (16.3).

For the parabolic $\check{P} \to \check{M}$, whose kernel is the unipotent radical $\check{U}$, the lemma gives

$$\begin{array}{ccc} \mathrm{QCoh}^!(\mathsf{LocSys}_{\check{P}}) & \xleftarrow{\ \mathrm{Loc}_x\ } & \mathrm{Rep}(\check{P})_x \\ {\scriptstyle \mathrm{push}}\downarrow & & \downarrow{\scriptstyle \mathrm{inv}^+} \\ \mathrm{QCoh}^!(\mathsf{LocSys}_{\check{M}}) & \xleftarrow{\ \mathrm{Loc}_{\mathrm{Ran}_{x,\mathrm{c}}}\ } & \mathrm{Rep}(\check{M})_{\mathrm{Ran}_{x,\mathrm{c}}}. \end{array}$$

After composing with $\mathsf{LocSys}_{\check{P}} \to \mathsf{LocSys}_{\check{G}}$, this gives the constant-term-oriented spectral square

$$\begin{array}{ccc} \mathrm{QCoh}^!(\mathsf{LocSys}_{\check{G}}) & \xrightarrow{\ \Gamma^{\mathrm{spec}}\ } & \mathrm{Rep}(\check{G})_x \\ {\scriptstyle \mathrm{CT}^{\mathrm{spec}}_{\check{P}}}\downarrow & & \downarrow{\scriptstyle \mathrm{Chev}^+} \\ \mathrm{QCoh}^!(\mathsf{LocSys}_{\check{M}}) & \xrightarrow{\ \Gamma^{\mathrm{spec}}\ } & \mathrm{Rep}(\check{M})_{\mathrm{Ran}_{x,\mathrm{c}}}. \end{array}$$

Here Chev^+ is the Chevalley functor. If $V_x \in \mathrm{Rep}(\check{G})_x$, then schematically

$$\mathrm{Chev}^+(V_x)_{(x,y_1,\dots,y_m)} = C^\bullet(\check{\mathfrak{u}}, V_x)_x \otimes C^\bullet(\check{\mathfrak{u}})_{y_1} \otimes \cdots \otimes C^\bullet(\check{\mathfrak{u}})_{y_m},$$

where $\check{\mathfrak{u}} = \mathrm{Lie}(\check{U})$ is the Lie algebra of the unipotent radical of $\check{P}$. The opposite radical appears after passing to the dual, or equivalently to left adjoints on the representation side, gives the Eisenstein-oriented square

$$\begin{array}{ccc} \mathrm{QCoh}^!(\mathsf{LocSys}_{\check{G}}) & \xrightarrow{\ \Gamma^{\mathrm{spec}}\ } & \mathrm{Rep}(\check{G})_x \\ {\scriptstyle \mathrm{Eis}^{\mathrm{spec}}_{\check{P}}}\uparrow & & \uparrow{\scriptstyle (\mathrm{Chev}^+)^{\mathrm{L}}} \\ \mathrm{QCoh}^!(\mathsf{LocSys}_{\check{M}}) & \xrightarrow{\ \Gamma^{\mathrm{spec}}\ } & \mathrm{Rep}(\check{M})_{\mathrm{Ran}_{x,\mathrm{c}}}. \end{array}$$

16.4. **The geometric square and the semi-infinite category.** The geometric side is the square below:

$$\begin{array}{ccc} \mathrm{DMod}(\mathsf{Bun}_G) & \xrightarrow{\ \mathrm{coeff}_G\ } & \mathrm{Whit}(G)_x \\ {\scriptstyle \mathrm{Eis}_P}\uparrow & & \uparrow{\scriptstyle J^-} \\ \mathrm{DMod}(\mathsf{Bun}_M) & \xrightarrow{\ \mathrm{coeff}_M\ } & \mathrm{Whit}(M)_{\mathrm{Ran}_{x,\mathrm{c}}}. \end{array}$$

Here $J^-\colon \mathrm{Whit}(M)_{\mathrm{Ran}_{x,\mathrm{c}}} \to \mathrm{Whit}(G)_x$ is the induction-side functor in the Eisenstein square. The same semi-infinite kernel first gives a local restriction functor $r_P^-\colon \mathrm{Whit}(G) \to \mathrm{Whit}(M)$; J^- is its adjoint form, as predicted by second adjointness. Locally, identify

$$\mathrm{Whit}(G) = \mathrm{DMod}(\mathsf{Gr}_G)^{\mathfrak{L}N,\chi_G}, \qquad \mathrm{Whit}(M) = \mathrm{DMod}(\mathsf{Gr}_M)^{\mathfrak{L}N_M,\chi_M}.$$

The kernel is constructed from the negative semi-infinite category

$$\begin{aligned} \mathrm{SI}_P^- &:= \mathrm{DMod}(\mathsf{Gr}_G)^{\mathfrak{L}N^-\cdot\mathfrak{L}^+M} \\ &\simeq (\mathrm{DMod}(\mathsf{Gr}_G) \otimes \mathrm{DMod}(\mathsf{Gr}_M))^{\mathfrak{L}P^-} \\ &\simeq \mathrm{Fun}_{\mathfrak{L}P^-}(\mathrm{DMod}(\mathsf{Gr}_G), \mathrm{DMod}(\mathsf{Gr}_M)). \end{aligned}$$

Thus any $\mathfrak{L}P^-$-linear functor $F\colon \mathrm{DMod}(\mathsf{Gr}_G) \to \mathrm{DMod}(\mathsf{Gr}_M)$ gives a restriction-type functor

$$r_F\colon \mathrm{DMod}(\mathsf{Gr}_G)^{\mathfrak{L}N,\chi_G} \longrightarrow \mathrm{DMod}(\mathsf{Gr}_M)^{\mathfrak{L}N_M,\chi_M}.$$

The functor J^- used above is the corresponding adjoint for the unit object of SI_P^-, namely the standard object on the orbit

$$\mathfrak{L}N^- \cdot 1 \cdot \mathfrak{L}^+G/\mathfrak{L}^+G \subset \mathsf{Gr}_G.$$

The Whittaker character χ_G is then degenerated to a character χ_P trivial on the unipotent radical of P and inducing χ_M on $\mathfrak{L}N_M$. This gives the local pattern

$$\mathrm{DMod}(\mathsf{Gr}_G)^{\mathfrak{L}N,\chi_G} \xrightarrow{\text{nearby cycles}} \mathrm{DMod}(\mathsf{Gr}_G)^{\mathfrak{L}N,\chi_P} \longrightarrow \mathrm{DMod}(\mathsf{Gr}_M)^{\mathfrak{L}N_M,\chi_M}.$$

The use of P^- in the semi-infinite correspondence and P in the degeneration of characters is a manifestation of second adjointness.

Let $\mathrm{Chev}^- := (\mathrm{Chev}^+)^{\mathrm{L}}$. Under Casselman–Shalika the desired top square is as follows:

$$\begin{array}{ccc} \mathrm{Whit}(G)_x & \xrightarrow{\ \sim\ } & \mathrm{Rep}(\check{G})_x \\ {\scriptstyle J^-}\uparrow & & \uparrow{\scriptstyle \mathrm{Chev}^-} \\ \mathrm{Whit}(M)_{\mathrm{Ran}_{x,\mathrm{c}}} & \xrightarrow[\sim]{} & \mathrm{Rep}(\check{M})_{\mathrm{Ran}_{x,\mathrm{c}}}. \end{array}$$

The proof uses the spectral description of SI_P^-:

$$\mathrm{SI}_P^- \simeq \mathrm{QCoh}^! \left(\mathsf{LocSys}_{\check{G}}(D) \times_{\mathsf{LocSys}_{\check{G}}(D^\circ)} \mathsf{LocSys}_{\check{P}}(D^\circ) \times_{\mathsf{LocSys}_{\check{M}}(D^\circ)} \mathsf{LocSys}_{\check{M}}(D)\right). \tag{16.4}$$

When $P = G$, this is exactly the spectral side of derived Satake. More generally, it is the parabolic semi-infinite analogue of derived Satake.

The local-Langlands intuition is that an $\mathfrak{L}G$-module category should correspond to a sheaf of categories over $\mathsf{LocSys}_{\check{G}}(D^\circ)$. Restricting the loop group from G to M corresponds to pulling along $\mathsf{LocSys}_{\check{M}}(D^\circ) \to \mathsf{LocSys}_{\check{G}}(D^\circ)$. Thus one expects

$$\begin{aligned} \mathrm{DMod}(\mathsf{Gr}_G) &\rightsquigarrow \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}(D)), \\ \mathrm{DMod}(\mathsf{Gr}_M) &\rightsquigarrow \mathrm{QCoh}(\mathsf{LocSys}_{\check{M}}(D)), \\ \mathrm{DMod}(\mathsf{Gr}_G)^{\mathfrak{L}M} &\rightsquigarrow \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}(D) \times_{\mathsf{LocSys}_{\check{G}}(D^\circ)} \mathsf{LocSys}_{\check{M}}(D^\circ)). \end{aligned}$$

Computing $\mathrm{Fun}(-,-)$ on the two sides gives (16.4), and hence proves the Eisenstein compatibility (16.1).

16.5. **Why constant term is harder.** The method proving (16.1) does not directly prove constant-term compatibility (16.2). In the Eisenstein diagram all arrows point in the direction needed for pushforward. For CT_P one would have to reverse those arrows, and the simple geometric argument no longer applies.

The replacement is Kac–Moody–Whittaker duality. At non-critical level one compares

$$\begin{array}{ccc} \mathrm{Whit}_\kappa(M) & \longrightarrow & \mathrm{Whit}_\kappa(G) \\ {\scriptstyle \mathrm{coeff}_M}\uparrow & & \uparrow{\scriptstyle \mathrm{coeff}_G} \\ \mathrm{DMod}_\kappa(\mathsf{Bun}_M) & \xrightarrow{\mathrm{Eis}_P} & \mathrm{DMod}_\kappa(\mathsf{Bun}_G). \end{array}$$

Specializing to the critical level, or equivalently to the boundary value of the quantum parameter used in the FLE degeneration, the local Whittaker categories are replaced by spherical critical Kac–Moody categories. In the IndCoh convention, FLE identifies them with unramified opers:

$$\mathrm{KL}_{G,\mathrm{crit}} \simeq \mathrm{IndCoh}(\mathsf{Op}_{\check{G}}^{\mathrm{unr}}), \qquad \mathrm{KL}_{M,\mathrm{crit}} \simeq \mathrm{IndCoh}(\mathsf{Op}_{\check{M}}^{\mathrm{unr}}).$$

The relevant local functor is

$$r_P^{\mathrm{KM}} \colon \mathrm{KL}_{G,\mathrm{crit}} \longrightarrow \mathrm{KL}_{M,\mathrm{crit}}, \qquad \mathcal{M} \longmapsto C^{\infty/2}(\mathfrak{u}(\!(t)\!), \mathcal{M}),$$

where $\mathfrak{u}$ is the Lie algebra of the unipotent radical of P. Its spectral counterpart is the functor

$$r_{\check{P}}^{\mathrm{op}} \colon \mathrm{IndCoh}(\mathsf{Op}_{\check{G}}^{\mathrm{unr}}) \longrightarrow \mathrm{IndCoh}(\mathsf{Op}_{\check{M}}^{\mathrm{unr}})$$

given by the unit object in $\mathrm{SI}_{\check{P}^-,\mathrm{spec}}$. These functors prove the top and side faces of the constant-term cube [CCF+24]. The remaining front and back faces are not consequences of the semi-infinite argument alone; they are proved by a chiral method. This completes the proof of the constant-term compatibility (16.2).

16.6. **Adjoints and the Eisenstein-generated part.** The last part of the Eisenstein argument explains how the compatibilities (16.1) and (16.2) enter the global induction. Let $\mathrm{DMod}(\mathsf{Bun}_G)_{\mathrm{Eis}}$ be the subcategory generated by $\mathrm{Eis}_P(\mathrm{DMod}(\mathsf{Bun}_M))$ for proper parabolics P, and let $\mathsf{LocSys}_{\check{G}}^{\mathrm{red}} \subset \mathsf{LocSys}_{\check{G}}$ be the reducible locus, i.e. the union of the images of $\mathsf{LocSys}_{\check{P}} \to \mathsf{LocSys}_{\check{G}}$ over proper parabolics $\check{P}$. Equivalently, these are the local systems admitting a proper parabolic reduction. Its complement is the irreducible locus $\mathsf{LocSys}_{\check{G}}^{\mathrm{irr}}$.

One needs that the coarse functor $\mathbb{L}_G^0$ has a left adjoint. The point is that it suffices to construct the left adjoint on generators. Spectrally, generators come from $\mathrm{Eis}_{\check{P}}^{\mathrm{spec}}(\mathrm{QCoh}(\mathsf{LocSys}_{\check{M}}))$, and the Eisenstein compatibility gives the required automorphic functor by transporting adjoints from M. Equivalently, one uses the fact that the coarse functor already has a left adjoint after passing to the Ran-action description: $\mathrm{Rep}(\check{M})_{\mathrm{Ran}}$ acts on the automorphic side, and the quotient $\mathrm{Rep}(\check{M})_{\mathrm{Ran}} \to \mathrm{QCoh}(\mathsf{LocSys}_{\check{M}})$ identifies the spectral action.

Consequently, assuming the Langlands equivalence for every proper Levi M, the functor $\mathbb{L}_G$ induces an equivalence on the Eisenstein summands [CCF+24]:

$$\mathbb{L}_G\colon\ \mathrm{DMod}(\mathsf{Bun}_G)_{\mathrm{Eis}} \xrightarrow{\sim} \mathrm{QCoh}^!(\mathsf{LocSys}_{\check{G}})_{\mathrm{red}}. \tag{16.5}$$

Thus the global conjecture is reduced to the cuspidal part. The remaining statement is

$$\mathbb{L}_{G,\mathrm{cusp}}\colon\ \mathrm{DMod}(\mathsf{Bun}_G)_{\mathrm{cusp}} \xrightarrow{\sim} \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}^{\mathrm{irr}}).$$

This reduction supplies the input for the cuspidal argument below.

We now finish the proof after the Eisenstein reduction above, following the multiplicity-one part of the proof series [GR24b]. We keep the notation

$$\mathbb{L}_G\colon\ \mathrm{DMod}(\mathsf{Bun}_G) \longrightarrow \mathrm{QCoh}^!(\mathsf{LocSys}_{\check{G}})$$

for the renormalized Langlands functor and write $\mathbb{L}_G^{\mathrm{L}}$ and $\mathbb{L}_G^{\mathrm{R}}$ for its left and right adjoints when restricted to the relevant summands. The preceding subsections constructed $\mathbb{L}_G$, constructed $\mathbb{L}_G^{\mathrm{L}}$, and proved the Eisenstein equivalence

$$\mathbb{L}_G\colon\ \mathrm{DMod}(\mathsf{Bun}_G)_{\mathrm{Eis}} \xrightarrow{\sim} \mathrm{QCoh}^!(\mathsf{LocSys}_{\check{G}})_{\mathrm{red}}.$$

It remains to prove the cuspidal equivalence

$$\mathbb{L}_{G,\mathrm{cusp}}\colon\ \mathrm{DMod}(\mathsf{Bun}_G)_{\mathrm{cusp}} \xrightarrow{\sim} \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}^{\mathrm{irr}}).$$

Here $\mathsf{LocSys}_{\check{G}}^{\mathrm{irr}}$ is the open substack of irreducible local systems, i.e. those admitting no reduction to a proper parabolic $\check{P} \subsetneq \check{G}$.

16.7. **Known inputs on the cuspidal functor.** The first input is conservativity.

Theorem 16.7.1. *The functor* $\mathbb{L}_{G,\mathrm{cusp}}$ *is conservative.*

Proof Idea. The proof roughly uses Whittaker coefficients and the geometric Ramanujan input [FR25, Ber21, GR24b]. The automorphic coefficient functor gives a square of the form

$$\begin{array}{ccc}
\mathrm{DMod}(\mathsf{Bun}_G)_{\mathrm{cusp}} & \xrightarrow{\mathbb{L}_{G,\mathrm{cusp}}} & \mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}^{\mathrm{irr}}) \\
{\scriptstyle \mathrm{coeff}_G}\big\downarrow & & \big\downarrow \\
\mathrm{Whit}(G)_{\mathrm{Ran}} & \xrightarrow{\mathrm{CS}} & \mathrm{Rep}(\check{G})_{\mathrm{Ran}}.
\end{array}$$

The lower horizontal functor is the Casselman–Shalika identification. By the geometric Ramanujan input, cuspidal objects are tempered, and the Whittaker coefficient functor is conservative on the tempered subcategory. Hence $\mathbb{L}_{G,\mathrm{cusp}}$ is conservative. ☐

The second input is the existence of the adjoint already constructed in Subsection 16.6. Thus, to prove that $\mathbb{L}_{G,\mathrm{cusp}}$ is an equivalence, it suffices to identify the monad

$$\mathbb{L}_{G,\mathrm{cusp}} \circ \mathbb{L}_{G,\mathrm{cusp}}^{\mathrm{L}} \in \mathrm{End}(\mathrm{QCoh}(\mathsf{LocSys}_{\check{G}}^{\mathrm{irr}})).$$

The spectral category is $\mathrm{QCoh}(\mathsf{LocSys}^{\mathrm{irr}}_{\check{G}})$-linear. Hence the monad is tensoring with an associative algebra object

$$A_{G,\mathrm{irr}} \coloneqq (\mathbb{L}_{G,\mathrm{cusp}} \circ \mathbb{L}^{\mathrm{L}}_{G,\mathrm{cusp}})(\mathcal{O}_{\mathsf{LocSys}^{\mathrm{irr}}_{\check{G}}}) \in \mathrm{AssocAlg}(\mathrm{QCoh}(\mathsf{LocSys}^{\mathrm{irr}}_{\check{G}})).$$

Equivalently, if $A_G = \mathbb{L}_G \mathbb{L}^{\mathrm{L}}_G(\mathcal{O})$ on the whole spectral side, then $A_{G,\mathrm{irr}} \simeq A_G|_{\mathsf{LocSys}^{\mathrm{irr}}_{\check{G}}}$. This restriction is not merely formal from the notation $\mathrm{QCoh}^!$, but follows from the compatibility of the renormalized spectral category with the ordinary QCoh-module structure on the irreducible open. The only remaining task is therefore

$$\mathcal{O}_{\mathsf{LocSys}^{\mathrm{irr}}_{\check{G}}} \xrightarrow{\sim} A_{G,\mathrm{irr}}.$$

The reducible part of this statement was proved above by Eisenstein induction.

16.8. **Ambidexterity and the self-dual algebra.**

Theorem 16.8.1 (Ambidexterity [ABC+24b])**.** *There is a canonical equivalence of adjoints*

$$\mathbb{L}^{\mathrm{L}}_{G,\mathrm{cusp}} \simeq \mathbb{L}^{\mathrm{R}}_{G,\mathrm{cusp}}.$$

Sketch. Both sides are compared by the localization and coefficient diagrams used above. The automorphic cuspidal category and $\mathrm{QCoh}(\mathsf{LocSys}^{\mathrm{irr}}_{\check{G}})$ carry their natural dualities. The compatibility of $\mathbb{L}_G$ with FLE, spectral Poincare functors, and Whittaker coefficients identifies the functor dual to $\mathbb{L}_{G,\mathrm{cusp}}$ with the same functor. Thus the left and right adjoints coincide. □

The consequence is that the algebra $A_{G,\mathrm{irr}}$ is perfect as a complex on $\mathsf{LocSys}^{\mathrm{irr}}_{\check{G}}$ and is canonically self-dual. Indeed, the right adjoint gives the same kernel as the left adjoint but with the opposite adjunction convention. Therefore the object dual to $A_{G,\mathrm{irr}}$ is the same object, now viewed as a coalgebra. Denote this coalgebra by $B_{G,\mathrm{irr}}$. In formulas,

$$B_{G,\mathrm{irr}} \coloneqq A^{\vee}_{G,\mathrm{irr}}, \qquad A^{\vee}_{G,\mathrm{irr}} \simeq A_{G,\mathrm{irr}}.$$

This is the point where ambidexterity converts a monadic problem into a geometric statement about a pseudo-proper Ran pushforward on the spectral side.

16.9. **Generic opers and the coalgebra.** Let

$$\pi\colon \overline{\mathsf{Op}}^{\mathrm{gen}}_{\check{G},\mathrm{irr}} \longrightarrow \mathsf{LocSys}^{\mathrm{irr}}_{\check{G}}$$

be the stack of generic, more precisely extended generic, $\check{G}$-oper structures on an irreducible local system. Informally, the fiber over $\sigma \in \mathsf{LocSys}^{\mathrm{irr}}_{\check{G}}$ is the space of generic reductions of σ to $\check{B}$ satisfying the oper transversality condition away from finitely many points of X. The FLE/localization formalism identifies the coalgebra above with the de Rham direct image [ABC+24a, GR24b]

$$B_{G,\mathrm{irr}} \simeq \pi_{\mathrm{dR},!}\omega_{\overline{\mathsf{Op}}^{\mathrm{gen}}_{\check{G},\mathrm{irr}}}$$

as a cocommutative coalgebra object of $\mathrm{QCoh}(\mathsf{LocSys}^{\mathrm{irr}}_{\check{G}})$. In the proof this object is constructed from the vacuum Poincare kernel: local FLE says that the Kac–Moody vacuum is sent to the unramified-oper object, and the Ran factorization glues the local generic reductions into the above global generic-oper correspondence.

The following elementary consequences are used. First, $B_{G,\mathrm{irr}}$ is cocommutative and lives in cohomological degrees $\leqslant 0$, with a flat connection along $\mathsf{LocSys}^{\mathrm{irr}}_{\check{G}}$. Second, by self-duality, $A_{G,\mathrm{irr}}$ is commutative and lives in degrees $\geqslant 0$. Since duality carries this perfect object back to itself, the opposite bound also holds; hence $A_{G,\mathrm{irr}}$ lies in the heart and is a vector bundle with flat connection. Third, $\pi_{\mathrm{dR},!}\omega$ has finite monodromy. Hence:

Theorem 16.9.1. *The algebra $A_{G,\mathrm{irr}}$ is a finite étale commutative $\mathcal{O}_{\mathsf{LocSys}^{\mathrm{irr}}_{\check{G}}}$-algebra. Equivalently, it defines a finite étale cover of $\mathsf{LocSys}^{\mathrm{irr}}_{\check{G}}$.*

Remark 16.9.2. For classical groups this finite-monodromy input is strengthened by the theorem of Beraldo–Kazhdan–Schlank: the relevant generic-oper fibers are homologically contractible [BKS18]. Thus in types A, B, C, D the finite étale cover is already forced to be trivial. The proof uses this as the model case and then reduces the general case to the same finite-étale argument.

16.10. **Reduction to a numerical rank.** We now restrict to the standard good case used in the proof: G is of adjoint type and $g_X > 1$, with the small exceptional case $g_X = 2$ and $\check{G} = \mathrm{SL}_2$ handled by a separate reduction. In this case $\mathsf{LocSys}^{\mathrm{irr}}_{\check{G}}$ has no nontrivial finite étale covers in the sense needed here. Therefore any finite étale algebra on it is constant:

$$A_{G,\mathrm{irr}} \simeq \mathcal{O}^{\oplus n}_{\mathsf{LocSys}^{\mathrm{irr}}_{\check{G}}}$$

for some integer $n \geqslant 1$. To prove $n = 1$, it is enough to show that the map on global sections

$$\Gamma(\mathsf{LocSys}^{\mathrm{irr}}_{\check{G}}, \mathcal{O}) \longrightarrow \Gamma(\mathsf{LocSys}^{\mathrm{irr}}_{\check{G}}, A_{G,\mathrm{irr}})$$

is an isomorphism. The complement of the irreducible locus has codimension at least two in the good case, and the stack of local systems is Cohen–Macaulay; hence the same question may be checked after extending $A_{G,\mathrm{irr}}$ to the global algebra A_G on all of $\mathsf{LocSys}_{\check{G}}$. It is enough to prove

$$\Gamma(\mathsf{LocSys}_{\check{G}}, \mathcal{O}) \xrightarrow{\sim} \Gamma(\mathsf{LocSys}_{\check{G}}, A_G).$$

The global-functions input gives $\Gamma(\mathsf{LocSys}_{\check{G}}, \mathcal{O}) = k$. By definition of the global Poincare-vacuum kernel, the right-hand side is the endomorphism algebra of the basic Whittaker object W_i:

$$\Gamma(\mathsf{LocSys}_{\check{G}}, A_G) \simeq \mathrm{End}(W_i).$$

Thus the final computation is $\mathrm{End}(W_i) = k$.

The Whittaker chart used in the calculation is the open stratum in the usual Drinfeld compactification. For sufficiently negative degree data i there is a locally closed diagram

$$\begin{array}{ccccc} \mathsf{Bun}^i_N & \longrightarrow & \mathsf{Bun}^i_N/T & \xhookrightarrow[\text{closed}]{\text{locally}} & \mathsf{Bun}_G \\ \Big\downarrow{\scriptstyle \chi_{\mathrm{glob}}} & & \Big\downarrow & & \\ \mathbb{G}_{\mathrm{m}}^{r(G)} & \longrightarrow & \mathbb{G}_{\mathrm{m}}^{r(G)}/T, & & \end{array}$$

where $r(G) = \mathrm{rank}_{\mathrm{ss}}(G)$ and χ_{glob} is the product of the simple-root Whittaker characters. The relevant minimal stratum has the form

$$\mathring{\mathcal{A}}_{\min} \simeq \mathbb{G}_{\mathrm{m}}^{r(G)} \times \mathbb{B}U_i,$$

with U_i unipotent. For the clean Whittaker object the $\mathbb{B}U_i$-factor contributes only the unit; thus the relevant endomorphism calculation reduces to

$$\mathrm{End}_{\mathrm{DMod}(\mathring{\mathcal{A}}_{\min})}(W_i|_{\mathring{\mathcal{A}}_{\min}}) \simeq \mathrm{End}_{\mathrm{DMod}(\mathbb{G}_{\mathrm{m}}^{r(G)})}(\mathcal{E}^{\boxtimes r(G)}).$$

The object W_i restricts to the clean exponential object $\chi^!_{\mathrm{glob}}\mathcal{E}$ on this chart. Therefore the calculation is the elementary one for $\mathcal{E}^{\boxtimes r(G)}$ on $\mathbb{G}_{\mathrm{m}}^{r(G)}$, and

$$\mathrm{End}(W_i) = k.$$

Consequently $n = 1$, so $A_{G,\mathrm{irr}} \simeq \mathcal{O}$. Hence the monad $\mathbb{L}_{G,\mathrm{cusp}}\mathbb{L}^{\mathrm{L}}_{G,\mathrm{cusp}}$ is the identity. Since $\mathbb{L}_{G,\mathrm{cusp}}$ is conservative and has the required adjoint, it is an equivalence:

$$\mathbb{L}_{G,\mathrm{cusp}}\colon \mathrm{DMod}(\mathsf{Bun}_G)_{\mathrm{cusp}} \xrightarrow{\sim} \mathrm{QCoh}(\mathsf{LocSys}^{\mathrm{irr}}_{\check{G}}).$$

Combining this with the (reducible) Eisenstein equivalence (16.5) proves the global geometric Langlands equivalence [CCF+24, ABC+24b, GR24b].

References

[ABC+24a] Dima Arinkin, Dario Beraldo, Justin Campbell, Lin Chen, Joakim Færgeman, Dennis Gaitsgory, Kevin Lin, Sam Raskin, and Nick Rozenblyum. Proof of the geometric Langlands conjecture II: Kac–Moody localization and the FLE, 2024. Preprint. arXiv:2405.03648. ↑56, 57, 58, 64

[ABC+24b] Dima Arinkin, Dario Beraldo, Lin Chen, Joakim Færgeman, Dennis Gaitsgory, Kevin Lin, Sam Raskin, and Nick Rozenblyum. Proof of the geometric Langlands conjecture IV: Ambidexterity, 2024. Preprint. arXiv:2409.08670. ↑64, 65

[AG15] Dima Arinkin and Dennis Gaitsgory. Singular support of coherent sheaves, and the geometric Langlands conjecture. *Selecta Mathematica. New Series*, 21(1):1–199, 2015. doi:10.1007/s00029-014-0167-5. ↑3, 7, 38, 41, 42, 44

[AGK+20a] Dima Arinkin, Dennis Gaitsgory, David Kazhdan, Sam Raskin, Nick Rozenblyum, and Yakov Varshavsky. The stack of local systems with restricted variation and geometric Langlands theory with nilpotent singular support, 2020. Preprint. arXiv:2010.01906. ↑39, 44, 47, 48

[AGK+20b] Dima Arinkin, Dennis Gaitsgory, David Kazhdan, Sam Raskin, Nick Rozenblyum, and Yakov Varshavsky. Duality for automorphic sheaves with nilpotent singular support, 2020. Preprint. arXiv:2012.07665. ↑49

[AHJR14] Pramod N. Achar, Anthony Henderson, Daniel Juteau, and Simon Riche. Weyl group actions on the Springer sheaf. *Proceedings of the London Mathematical Society*, 108(6):1501–1528, 2014. doi:10.1112/plms/pdt055. ↑22, 24, 33

[BCTZ16] Roman Bezrukavnikov, Tsao-Hsien Chen, Roman Travkin, and Xinwen Zhu. Quantization of Hitchin integrable system via positive characteristic, 2016. Preprint. arXiv:1603.01327. ↑55

[BD91] Alexander Beilinson and Vladimir Drinfeld. Quantization of Hitchin's integrable system and Hecke eigensheaves, 1991. Unpublished manuscript. URL. ↑53, 55

[BD05] Alexander Beilinson and Vladimir Drinfeld. Opers, 2005. Preprint. arXiv:math/0501398. ↑53, 54

[Bei16] Alexander Beilinson. Constructible sheaves are holonomic. *Selecta Mathematica. New Series*, 22:1797–1819, 2016. arXiv:1505.06768. ↑46

[Ber21] Dario Beraldo. On the geometric Ramanujan conjecture, 2021. Preprint. arXiv:2103.17211. ↑41, 63

[BKS18] Dario Beraldo, David Kazhdan, and Tomer M. Schlank. Contractibility of the space of generic opers for classical groups, 2018. Preprint. arXiv:1801.00655. ↑64

[BZN16] David Ben-Zvi and David Nadler. Betti geometric Langlands, 2016. Preprint. arXiv:1606.08523. ↑44, 46, 47

[CCF+24] Justin Campbell, Lin Chen, Joakim Færgeman, Dennis Gaitsgory, Kevin Lin, Sam Raskin, and Nick Rozenblyum. Proof of the geometric Langlands conjecture III: Compatibility with parabolic induction, 2024. Preprint. arXiv:2409.07051. ↑59, 60, 62, 63, 65

[FG05] Edward Frenkel and Dennis Gaitsgory. Local geometric Langlands correspondence and affine Kac–Moody algebras, 2005. Preprint. arXiv:math/0508382. ↑7, 56, 57, 58

[FG07] Edward Frenkel and Dennis Gaitsgory. Local geometric Langlands correspondence: The spherical case, 2007. Preprint. arXiv:0711.1132. ↑7, 56, 57

[FGKV98] Edward Frenkel, Dennis Gaitsgory, David Kazhdan, and Kari Vilonen. Geometric realization of Whittaker functions and the Langlands conjecture. *Journal of the American Mathematical Society*, 11(2):451–484, 1998. arXiv:alg-geom/9703022. ↑6, 8, 14, 15, 16, 17, 19, 20, 26, 37

[FR23] Joakim Færgeman and Sam Raskin. The Arinkin–Gaitsgory temperedness conjecture. *Bulletin of the London Mathematical Society*, 55(3):1419–1446, 2023. arXiv:2108.02719. ↑43, 51

[FR25] Joakim Færgeman and Sam Raskin. Non-vanishing of geometric Whittaker coefficients for reductive groups. *Journal of the American Mathematical Society*, 38(4):919–995, 2025. arXiv:2207.02955. ↑50, 52, 63

[Fre02] Edward Frenkel. Lectures on Wakimoto modules, opers and the center at the critical level, 2002. Preprint. arXiv:math/0210029. ↑8, 54, 57

[Gai04] Dennis Gaitsgory. On a vanishing conjecture appearing in the geometric Langlands correspondence. *Annals of Mathematics*, 160(2):617–682, 2004. arXiv:math/0204081. ↑27, 37, 39, 40

[Gai10] Dennis Gaitsgory. A generalized vanishing conjecture, 2010. Manuscript. URL. ↑38, 39

[Gai16a] Dennis Gaitsgory. From geometric to function-theoretic Langlands (or how to invent shtukas), 2016. Informal note. arXiv:1606.09608. ↑49

[Gai16b] Dennis Gaitsgory. Recent progress in geometric Langlands theory, 2016. Bourbaki exposé. arXiv:1606.09462. ↑27, 37

[Gai18] Dennis Gaitsgory. The local and global versions of the Whittaker category, 2018. Preprint. arXiv:1811.02468. ↑14, 17, 50, 51

[GPR20] William Graham, Martha Precup, and Amber Russell. A new approach to the generalized Springer correspondence, 2020. Preprint. arXiv:2002.12480. ↑22, 24, 33

[GR24a] Dennis Gaitsgory and Sam Raskin. Proof of the geometric Langlands conjecture I: Construction of the functor, 2024. Preprint. arXiv:2405.03599. ↑3, 7, 38, 41, 44, 59

[GR24b] Dennis Gaitsgory and Sam Raskin. Proof of the geometric Langlands conjecture V: The multiplicity one theorem, 2024. Preprint. arXiv:2409.09856. ↑63, 64, 65

[Hei04] Jochen Heinloth. Coherent sheaves with parabolic structure and construction of Hecke eigensheaves for some ramified local systems. *Annales de l'Institut Fourier*, 54(7):2235–2325, 2004. doi:10.5802/aif.2080. ↑20

[KL85] Nicholas M. Katz and Gérard Laumon. Transformation de Fourier et majoration de sommes exponentielles. *Publications Mathématiques de l'IHÉS*, 62:145–202, 1985. doi:10.1007/BF02698808. ↑20, 22, 25, 26, 32

[Laf18] Vincent Lafforgue. Chtoucas pour les groupes réductifs et paramétrisation de Langlands globale. *Journal of the American Mathematical Society*, 31(3):719–891, 2018. doi:10.1090/jams/897. ↑16, 49

[NT25] David Nadler and Jeremy Taylor. The Whittaker functional is a shifted microstalk. *Transformation Groups*, 30:1425–1450, 2025. doi:10.1007/s00031-023-09836-x. ↑52

[NY19] David Nadler and Zhiwei Yun. Spectral action in Betti geometric Langlands. *Israel Journal of Mathematics*, 232:537–582, 2019. doi:10.1007/s11856-019-1871-9. ↑44, 46, 47, 48

Department of Mathematics, National University of Singapore, Block S17, 10 Lower Kent Ridge Road, Singapore 119076

Email address: daiwenhan@u.nus.edu