Twistor structures and Hecke correspondences

November 25, 2024

Last time, we looked at the stacks

$$X^\diamondsuit \to X^\diamondsuit / \mathrm{U}(1)_{\mathrm{Betti}} \leftarrow X^{\mathrm{tw}}$$

for real or complex manifolds X, and claimed that vector bundles on these stacks corresponded to variations of Hodge structures on the middle term; variations of twistor structures on the right; and "variations of vector bundles on $\mathrm{Div}^1_{\mathbb{C}}$" (or Div^1 after Galois descent) on the left. Our first goal today is to justify these claims.

1. The analytic Hodge–Tate stack

Last time, we studied the open subset $X^\diamondsuit \times_{\mathrm{Div}^1_{\mathbb{C}}} \mathbb{C}^\times_{\mathrm{Betti}}$, which we saw was isomorphic to its Betti stack $X(\mathbb{C})_{\mathrm{Betti}} \times \mathbb{C}^\times_{\mathrm{Betti}}$. One way of saying this is that over this locus of $\mathrm{Div}^1_{\mathbb{C}}$, $X^\diamondsuit$ is uniform: indeed we proved this by observing that away from 0 and ∞, every degree 1 divisor of $X_{\mathbb{C},A}$ is isomorphic to $\mathrm{AnSpec}\,\mathrm{Cont}(S,\mathbb{R})$.

By contrast, over 0 and ∞ we have gluing data involving A, and so the degree 1 divisors supported over these points can be more complicated. Once again the stories at 0 and ∞ are symmetric (and interchanged under complex conjugation), so for concreteness we study the fiber over ∞. This is the analytic Hodge–Tate stack: X^{HT} is the functor sending A to the anima of maps

$$X_{\mathbb{C},A} \times_{X_{\mathbb{C},\infty}} \mathrm{AnSpec}\,\mathbb{C}_{\mathrm{gas}} \to X.$$

Locally near ∞, $X_{\mathbb{C},A}$ is given by gluing $\mathrm{AnSpec}\,A$ to $\mathbb{A}^1_{\mathbb{C}} \times_{\mathrm{AnSpec}\,\mathbb{C}_{\mathrm{gas}}} \mathrm{AnSpec}\,\mathrm{Cont}(S,\mathbb{C})$ along $\mathrm{AnSpec}\,\mathrm{Cont}(S,\mathbb{C})$ via the point at infinity; this is the affine analytic stack given by the spectrum of

$$\mathrm{Cont}(S,\mathbb{C})[T] \times_{\mathrm{Cont}(S,\mathbb{C})} A$$

with fiber at infinity given by evaluating at $T = 0$. Naively, this is just A; but we should take the fiber in the derived sense, giving

$$\widetilde{A} = (\mathrm{Cont}(S,\mathbb{C})[T] \times_{\mathrm{Cont}(S,\mathbb{C})} A)/^{\mathbb{L}}T.$$

This maps to A, with fiber $\mathrm{Nil}^\dagger(A)[1]$; the inclusion $\mathrm{Cont}(S,\mathbb{C}) \to \mathrm{Cont}(S,\mathbb{C})[T]$ induces a section, making this a split extension.[1] We can think of this as a sort of deformation of A; when A is "perfectoid," i.e. nil-reduced so that $A \simeq \mathrm{Cont}(S,\mathbb{C})$, then in fact $\widetilde{A} \simeq A$.

This is analogous to the following phenomenon in the p-adic situation: if X is a p-adic formal scheme, an R-point of its prismatization X^{Δ} is a Cartier–Witt divisor $I \to W(R)$ together with a map $\mathrm{Spec}\,W(R)/I \to X$, and the Hodge–Tate stack $X^{\mathrm{HT}} \subset X^{\Delta}$ is the locus on which $I \to W(R)$ factors through the Verschiebung operator $V : W(R) \to W(R)$. To each R-point of X^{Δ} we can associate the "untilt" $W(R)/I$; when R is (integral) perfectoid

[1] I don't follow this very well, so don't take this too literally unless you know animated rings better than I do and can verify.

in characteristic p, i.e. a perfect $\mathbb{F}_p$-algebra, this can almost be made literal, and the Hodge–Tate "untilt" $W(R)/I$ is just the original ring R. One however can take much more general test objects R, and in general $W(R)/VW(R)$ does not agree with R. Next week, we'll briefly mention an analogue of prismatization in our setting, which presumably should fit into an analogous picture with the analytic Hodge–Tate stack.

Proposition 1. *There is a natural map $X^{\mathrm{HT}} \to X$ which is a gerbe for $T_X^\dagger$ where T_X is the tangent bundle. In fact this gerbe is split, giving an isomorphism $X^{\mathrm{HT}} \simeq BT_X^\dagger$.*

Proof. Since X is a complex manifold, it is locally isomorphic to $\mathbb{A}^{n,\mathrm{an}}$, maps from $X_{\mathbb{C},A} \times_{X_{\mathbb{C},\infty}} \operatorname{AnSpec} \mathbb{C}_{\mathrm{gas}}$ to which are equivalent to n-tuples of elements of $\widetilde{A}$, i.e. A-points of $(\mathbb{A}^{n,\mathrm{an}})^{\mathrm{HT}}$ are n-tuples of points of $\widetilde{A}$. Since $\widetilde{A}$ is an extension of A by $\mathrm{Nil}^\dagger(A)[1]$, we can locally view these as $\dagger$-torsors for the trivialized tangent complex $T_{\mathbb{A}^{n,\mathrm{an}}}$ over A, i.e. A-points of $BT^\dagger_{\mathbb{A}^{n,\mathrm{an}}}$. These local isomorphisms glue, and the fact that the gerbe is split follows from the fact that $\widetilde{A}$ is a split extension. □

This is parallel to the Hodge–Tate gerbe in the p-adic setting, though there are some differences; there $T_X^\dagger$ is replaced by $T_X\{1\}^\sharp$, the PD-hull of the Breuil–Kisin-twisted tangent bundle. I'm not sure how much of the difference reflects the analytic rather than algebraic setting and how much is due to the archimedean vs. nonarchimedean distinction.

As a consequence we get the following description:

Corollary 2. *Vector bundles on X^{HT} are equivalent to Higgs bundles on X, i.e. vector bundles E on X together with a map $\theta : E \to E \otimes \Omega^1_X$ such that $\theta \wedge \theta : E \to E \otimes \Omega^2_X$ vanishes.*

Proof. Locally, $T_X^\dagger$ is isomorphic to $(\mathbb{G}_{\mathrm{a}}^\dagger)^n$ for some n, so vector bundles on $X^{\mathrm{HT}} \simeq BT_X^\dagger$ are locally equivalent to representations of $(\mathbb{G}_{\mathrm{a}}^\dagger)^n$, which is to say vector bundles on X together with n commuting endomorphisms. This can be massaged into the data of Higgs bundles. □

2. T-connections and twistor structures

Last time, we studied the diagram

$$X^\diamondsuit \to X^\diamondsuit / \mathrm{U}(1)_{\mathrm{Betti}} \leftarrow X^{\mathrm{tw}}$$

obtained by base change from the "absolute" case

$$\mathrm{Div}^1_{\mathbb{C}} \to \mathrm{Div}^1_{\mathbb{C}} / \mathrm{U}(1)_{\mathrm{Betti}} \leftarrow \mathbb{P}^1_{\mathbb{C}},$$

as well as its real analogue, and observed that on the locus over $\mathbb{C}^\times_{\mathrm{Betti}} \subset \mathrm{Div}^1_{\mathbb{C}}$ ("away from the poles") the diamondization has a simple description: $X^\diamondsuit \times_{\mathrm{Div}^1_{\mathbb{C}}} \mathbb{C}^\times_{\mathrm{Betti}}$ is isomorphic to its Betti stack, which is $X(\mathbb{C})_{\mathrm{Betti}} \times \mathbb{C}^\times_{\mathrm{Betti}}$. Now that we understand the fiber X^{HT} of $X^\diamondsuit$ over the point at infinity, our next goal is to extend this description to the neighborhood of ∞, where we've seen that the absolute diagram is locally

$$\mathbb{A}^{1,\mathrm{an}} / \mathrm{U}(1)^\dagger \to \mathbb{A}^{1,\mathrm{an}} / \mathrm{U}(1)^{\mathrm{la}} \leftarrow \mathbb{A}^{1,\mathrm{an}}.$$

On this locus, the map $\mathrm{U}(1)^\dagger \to \mathrm{U}(1)^{\mathrm{la}}$ does induce a lift of $\mathbb{A}^{1,\mathrm{an}} \to \mathbb{A}^{1,\mathrm{an}}/\mathrm{U}(1)^{\mathrm{la}}$ to $\mathbb{A}^{1,\mathrm{an}}/\mathrm{U}(1)^\dagger$, though this fails globally.

Suppose first that $X = \mathbb{A}^{n,\mathrm{an}}$, and choose coordinates $U_1, \ldots, U_n$. The action of $\mathrm{U}(1)^{\mathrm{la}} \subset \mathbb{G}_{\mathrm{m}}^{\mathrm{an}}$ on $\mathbb{G}_{\mathrm{a}}^{\mathrm{an}}$ by conjugation preserves $\mathbb{G}_{\mathrm{a}}^\dagger$, which lets us form the semidirect product $(\mathbb{G}_{\mathrm{a}}^\dagger)^n \rtimes \mathrm{U}(1)^{\mathrm{la}}$. This acts naturally on $\mathbb{A}^{n+1,\mathrm{la}}$ via $(u_1, \ldots, u_n, t) \cdot (U_1, \ldots, U_n, T) = (U_1 + u_1 T, \ldots, U_n + u_n T, tT)$. This is analogous to the action of $\mathbb{Z}_p^n \rtimes \mathbb{Z}_p^\times$ in p-adic Hodge theory, with the additive group acting on the "geometric" variables and the multiplicative group acting on the "arithmetic" variable.

Proposition 3. *There is a Cartesian diagram*

$$\begin{array}{ccc} \mathbb{A}^{n+1,\mathrm{an}}/(\mathbb{G}_a^\dagger)^n \rtimes \mathrm{U}(1)^{\mathrm{la}} & \longrightarrow & \mathbb{A}^{1,\mathrm{an}}/\mathrm{U}(1)^{\mathrm{la}} \\ \downarrow & & \downarrow \\ (\mathbb{A}^{n,\mathrm{an}})^\diamondsuit/\mathrm{U}(1)_{\mathrm{Betti}} & \longrightarrow & \mathrm{Div}^1_{\mathbb{C}}/\mathrm{U}(1)_{\mathrm{Betti}} \end{array}.$$

In particular, pulling back to $\mathbb{P}^1_{\mathbb{C}}$ and its neighborhood near ∞, we get an isomorphism

$$(\mathbb{A}^{n,\mathrm{an}})^{\mathrm{tw}} \times_{\mathbb{P}^1_{\mathbb{C}}} \mathbb{A}^{1,\mathrm{an}} \simeq \mathbb{A}^{n+1,\mathrm{an}}/(\mathbb{G}_{\mathrm{a}}^\dagger)^n,$$

where the action of $(\mathbb{G}_{\mathrm{a}}^\dagger)^n$ on $\mathbb{A}^{n+1,\mathrm{an}}$ is by specializing the formula above at $t = 1$, i.e.

$$(u_1, \ldots, u_n) \cdot (U_1, \ldots, U_n, T) = (U_1 + u_1 T, \ldots, U_n + u_n T, T).$$

Proof. The claim is that there is a surjective map $\mathbb{A}^{n+1,\mathrm{an}} \simeq \mathbb{A}^{n,\mathrm{an}} \times \mathbb{A}^{1,\mathrm{an}} \to (\mathbb{A}^{n,\mathrm{an}})^\diamondsuit \times_{\mathrm{Div}^1_{\mathbb{C}}} \mathbb{A}^{1,\mathrm{an}}$ which, after quotienting by $\mathrm{U}(1)_{\mathrm{Betti}}$, induces the claimed equivalence relation. We first construct this map: locally near ∞, as discussed in the previous section $X_{\mathbb{C},A}$ is given by $\mathrm{AnSpec}(\mathrm{Cont}(S, \mathbb{C})[T] \times_{\mathrm{Cont}(S,\mathbb{C})} A)$[2] which via the inclusion $\mathrm{Cont}(S, \mathbb{C}) \to \mathrm{Cont}(S, \mathbb{C})[T]$ maps to $\mathrm{AnSpec}\, A$. Thus every degree 1 divisor $Z \subset X_{\mathbb{C},A}$ away from 0 is equipped with a morphism $Z \to \mathrm{AnSpec}\, A$, and so composing with this morphism gives a map from A-points of a manifold X to Z-points of X. Allowing Z to vary (with support away from 0) and specializing to $X = \mathbb{A}^{n,\mathrm{an}}$ gives a map

$$\mathbb{A}^{n,\mathrm{an}} \times \mathbb{A}^{1,\mathrm{an}}/\mathrm{U}(1)^\dagger \to (\mathbb{A}^{n,\mathrm{an}})^\diamondsuit \times_{\mathrm{Div}^1_{\mathbb{C}}} \mathbb{A}^{1,\mathrm{an}}/\mathrm{U}(1)^\dagger.$$

Since other maps $\mathrm{AnSpec}(\mathrm{Cont}(S, \mathbb{C})[T] \times_{\mathrm{Cont}(S,\mathbb{C})} A) \to \mathrm{AnSpec}\, A$ at worst factor through quotients of A, in fact every Z-point arises in this way and so this map is surjective. One can also eliminate the quotient by $\mathrm{U}(1)^\dagger$ if desired, but for our purposes it's better to further take the quotient by $\mathrm{U}(1)_{\mathrm{Betti}}$ to obtain the quotient by $\mathrm{U}(1)^{\mathrm{la}}$ as in the diagram.

It remains to understand the fiber, which amounts to understanding the fiber of each map $Z \to \mathrm{AnSpec}\, A$. Using our description of $\widetilde{A}$ from last section, if I is the ideal sheaf of $Z \subset X_{\mathbb{C},A}$ one can compute that $A \to \mathcal{O}(Z)$ has cofiber $\mathrm{Nil}^\dagger(A)[1] \otimes I/I^2$, so that as I varies this is the action of the group stack corepresented by $\mathbb{A}^{n,\mathrm{an}}(\mathrm{Nil}^\dagger(A))$, i.e. $(\mathbb{G}_{\mathrm{a}}^\dagger)^n$. Unwinding definitions gives the claimed formulas. □

[2] Scholze's notes actually say $X_{\mathbb{C},A}$ is locally a subset of this affine analytic stack; this might just be an ambiguity of "neighborhood," or maybe I'm misunderstanding something.

Gluing affine pieces, we arrive at the following description.

Corollary 4. *Let X be a complex manifold. Vector bundles on*

$$X^{\mathrm{tw}} \times_{\mathbb{P}^1_{\mathbb{C}}} \mathbb{A}^{1,\mathrm{an}}$$

are equivalent to vector bundles E on $X \times \mathbb{A}^{1,\mathrm{an}}$ together with a flat T-connection, i.e. a map

$$\Delta : E \to E \otimes_{\mathcal{O}_X} \Omega^1_X$$

such that $\Delta \wedge \Delta = 0$ and

$$\Delta(fv) = f\Delta(v) + T\Delta(f)v.$$

In particular, gluing the above at 0 and ∞ lets us describe vector bundles on X^{tw} as pairs of vector bundles with T-connections on $X \times \mathbb{A}^{1,\mathrm{an}}$ for complementary affine loci (with one copy of T given by the inverse of the other), whose restrictions agree: the restrictions to $X \times \mathbb{G}_{\mathrm{m}}^{\mathrm{an}}$ give, by analytic Riemann–Hilbert, a $\mathbb{G}_{\mathrm{m}}$-family of local systems on X, which agree after twisting one family by complex conjugation. This is, for our purposes, a variation of $\mathbb{C}$-twistor structures on X.

We note that the connection is purely in the "geometric direction," i.e. on X rather than on the "arithmetic" $\mathbb{A}^{1,\mathrm{an}}$.

Taking $\mathrm{U}(1)^{\mathrm{la}}$-equivariant objects in the above description, we can study vector bundles on $X^{\diamondsuit}/\,\mathrm{U}(1)_{\mathrm{Betti}}$. We know that over $\mathbb{C}^{\times}_{\mathrm{Betti}} \subset \mathrm{Div}^1_{\mathbb{C}}$, $X^{\diamondsuit}$ is isomorphic to its Betti stack $X(\mathbb{C})_{\mathrm{Betti}} \times \mathbb{C}^{\times}_{\mathrm{Betti}}$, so quotienting by $\mathrm{U}(1)_{\mathrm{Betti}}$ on this open piece $\mathbb{C}^{\times}_{\mathrm{Betti}}/\,\mathrm{U}(1)_{\mathrm{Betti}} \simeq \mathbb{R}^{\times}_{>0,\mathrm{Betti}}$ we find that $X^{\diamondsuit}/\,\mathrm{U}(1)_{\mathrm{Betti}}$ is isomorphic to $X(\mathbb{C})_{\mathrm{Betti}} \times \mathbb{R}_{>0,\mathrm{Betti}}$, which since $\mathbb{R}_{>0}$ is contractible is (at least for the purposes of vector bundles) just $X(\mathbb{C})_{\mathrm{Betti}}$ and so vector bundles on it are just local systems on $X(\mathbb{C})$. Extensions to the poles give filtrations on this local system as above, which now must be stable under the T-connection; this amounts to Griffiths transversality. Thus we obtain the following result:

Corollary 5. *Vector bundles on $X^{\diamondsuit}/\,\mathrm{U}(1)_{\mathrm{Betti}} \times_{\mathrm{Div}^1_{\mathbb{C}}/\,\mathrm{U}(1)_{\mathrm{Betti}}} \mathbb{A}^{1,\mathrm{an}}/\,\mathrm{U}(1)^{\mathrm{la}}$ are equivalent to local systems $\mathbb{L}$ on X together with a (separated, exhaustive) filtration $\mathrm{Fil}^{\bullet}(\mathbb{L} \otimes_{\mathbb{C}} \mathcal{O}_X)$ of the corresponding vector bundle satisfying Griffiths transversality.*

In particular, gluing the descriptions at 0 and ∞, we find that vector bundles on $X^{\diamondsuit}/\,\mathrm{U}(1)_{\mathrm{Betti}}$ are precisely variations of $\mathbb{C}$-Hodge structures. Taking Galois descent, we find that vector bundles on the real version of $X^{\diamondsuit}/\,\mathrm{U}(1)_{\mathrm{Betti}}$ for real manifolds X gives variations of real Hodge structures.

Since moduli of variations of Hodge structures are given by Hermitian symmetric domains, this suggests interpreting the latter as spaces of vector bundles on $X^{\diamondsuit}/\,\mathrm{U}(1)_{\mathrm{Betti}}$ for various manifolds X, or $\mathrm{U}(1)_{\mathrm{Betti}}$-equivariant vector bundles on $X^{\diamondsuit}$. In the case $X = *$, we recover vector bundles on Div^1, the space of which is our candidate stack of L-parameters.

It is possible to get a similar description of vector bundles on $X^{\diamondsuit}$ itself, but these no longer have a classical description; they can be thought of as combining a T-connection in the geometric direction (as for vector bundles on X^{tw}) and another in the "arithmetic" direction, i.e. along $\mathbb{A}^{1,\mathrm{an}}$. These do not commute; they satisfy conditions corresponding to the noncommutative group $(\mathbb{G}_{\mathrm{m}}^{\dagger})^{\mathrm{an}} \rtimes \mathrm{U}(1)^{\dagger}$. If we divide by T to think about these as connections with logarithmic singularities, we can be more explicit: a vector bundle on $X^{\diamondsuit}$ should be equivalent to two vector bundles on $X \times \mathbb{A}^{1,\mathrm{an}}$ (with complementary coordinates) with connections with logarithmic singularities, whose restrictions to $X \times \mathbb{G}_{\mathrm{m}}^{\mathrm{an}}$ define the same local system on $X(\mathbb{C}) \times \mathbb{C}^{\times}$.

Each of our operations $-^{\diamondsuit}$, $-^{\diamondsuit}/\,\mathrm{U}(1)_{\mathrm{Betti}}$, and $-^{\mathrm{tw}}$ are functorial, and given a morphism of manifolds pushforward of the resulting maps gives relative cohomology of variations of Hodge and twistor structures. These functors preserve cohomological smoothness and properness, and so for f proper and smooth $f^{\diamondsuit}$ and f^{tw} preserve perfect complexes. Under suitable assumptions, they should also preserve vector bundle objects in each degree, but this is not yet known.

3. Bun_G

We now turn to the stack Bun_G, in a sense the core of the geometrization program. Our first goal is to describe its stratification, analogous to the decomposition of Pic by degree; in the next section we'll introduce Hecke operators, which are how we'll access the most interesting features of Bun_G. For today we'll speak in generalities, and next week we'll look into the simplest nontrivial example, which we'll see is closely related to the modular curve.

With our machinery set up, the definition is easy: let G be a real reductive group. Then Bun_G is the totally disconnected stack sending A to the anima of G-bundles on $X_{\mathbb{R},A}$. (We will sometimes conflate it with its analytic realization, when it is more convenient to talk about analytic stacks.) The proof that this is in fact a stack is essentially the same as for $\mathrm{Pic} = \mathrm{Bun}_{\mathbb{G}_{\mathrm{m}}}$.

Since $X_{\mathbb{R},A}$ is defined by a pushout diagram, the stack of G-bundles on it is defined by a pullback diagram:

Proposition 6. *The diagram*

$$\begin{array}{ccc} \mathrm{Bun}_G & \longrightarrow & (\mathrm{Bun}_G)_{\mathrm{Betti}} \\ \downarrow & & \downarrow \\ */G_{\mathbb{C}}^{\mathrm{an}} & \longrightarrow & */G(\mathbb{C})_{\mathrm{Betti}} \end{array}$$

is Cartesian.

Here the left vertical arrow is given by taking the fiber at $\infty : \mathrm{AnSpec}\, A \to X_{\mathbb{R},A}$ and the right vertical arrow is the Betti stack functor applied to this map.

This lets us prove our first key result about Bun_G:

Proposition 7. *The natural map $*/G(\mathbb{R})^{\mathrm{la}} \to \mathrm{Bun}_G$, induced by the trivial bundle $* \to \mathrm{Bun}_G$ together with its automorphism group $G(\mathbb{R})^{\mathrm{la}}$, is an open immersion, with image the locus of G-bundles which are fiberwise trivial.*

Proof. As with the inclusion $*/\mathbb{R}^{\times,\mathrm{la}} \to \mathrm{Pic}$, injectivity is clear, and the description of the image is essentially tautological. The main thing to prove is that the map is an open immersion. By Proposition 6, it suffices to prove that it is after taking Betti stacks, i.e. that

$$*/G(\mathbb{R})_{\mathrm{Betti}} \to (\mathrm{Bun}_G)_{\mathrm{Betti}}$$

is an open immersion, since then in the diagram

$$\begin{array}{ccc} */G(\mathbb{R})^{\mathrm{la}} & \longrightarrow & */G(\mathbb{R})_{\mathrm{Betti}} \\ \downarrow & & \downarrow \\ \mathrm{Bun}_G & \longrightarrow & (\mathrm{Bun}_G)_{\mathrm{Betti}} \\ \downarrow & & \downarrow \\ */G_{\mathbb{C}}^{\mathrm{an}} & \longrightarrow & */G(\mathbb{C})_{\mathrm{Betti}} \end{array}$$

the bottom square and outer square are Cartesian so so is the upper square and so $*/G(\mathbb{R})^{\mathrm{la}} \to \mathrm{Bun}_G$ is the pullback of $*/G(\mathbb{R})_{\mathrm{Betti}} \to (\mathrm{Bun}_G)_{\mathrm{Betti}}$. Since we know this is an inclusion and we can describe its image, the remaining problem can be stated as follows: if S is a light profinite set and $A = \mathrm{Cont}(S, \mathbb{C})$, for any G-bundle E on $X_{\mathbb{R},A}$, the locus of $s \in S$ for which E_s is trivial is open in S, and if it is all of S then E is in fact trivial.

First, suppose $G = \mathrm{GL}_n$. The openness of the subset of S on which E_s is trivial follows from the semicontinuity of the Newton polygon, whose proof is similar to the p-adic case; this apparently largely amounts to the properness of projectivized Banach–Colmez spaces, which we have nearly already seen and which is easier in this setting than p-adically.

Next, we want to show that if this set is all of S, then E is trivial; so assume that E_s is trivial for all s. The cohomology $V = R\Gamma(X_{\mathbb{R},A}, E)$ is (over $\mathrm{Cont}(S, \mathbb{R})$) a perfect complex concentrated in degrees 0 and 1; since each E_s is trivial, $H^1(X_{\mathbb{R},A}, E)$ vanishes (as previously) and so this is actually a vector bundle over $\mathrm{Cont}(S, \mathbb{R})$. Pulling back to $X_{\mathbb{R},A}$, we get a (trivial) vector bundle $V \otimes_{\mathbb{R}} \mathcal{O}_{X_{\mathbb{R},A}}$ with a natural map to E, which on each fiber is an isomorphism and hence is an isomorphism globally, so E itself is trivial.

For a general reductive group G, we have an embedding $G \hookrightarrow \mathrm{GL}_n$ with affine smooth quotient GL_n / G. Viewing the G-torsor E as a GL_n-torsor via this map, via the previous paragraph we know the relevant results for GL_n, so we may as well assume that E is trivial as a GL_n-torsor; choosing a trivialization, the data of a G-torsor together on $X_{\mathbb{R},A}$ with a GL_n-trivialization is equivalent to a map $X_{\mathbb{R},A} \to \mathrm{GL}_n / G$, where the target is smooth and affine. Referring back to the definition of $X_{\mathbb{R},A}$, we deduce that this is equivalent to a map $\mathrm{Spec}\,\mathrm{Cont}(S, \mathbb{R}) \to \mathrm{GL}_n / G$, which in turn is equivalent to a continuous map $S \to (\mathrm{GL}_n / G)(\mathbb{R})$. The locus on which E is trivial corresponds to the image of $\mathrm{GL}_n(\mathbb{R}) \to (\mathrm{GL}_n / G)(\mathbb{R})$, which is a submersion of real manifolds and hence has open image. If the image of S is inside this image, we can lift it to GL_n and so the underlying G-bundle E is globally trivial since its GL_n-lift is. □

Recall that G-bundles on the absolute curve $X_{\mathbb{R}}$ are classified by Kottwitz's set $B(\mathbb{R}, G)$ (this is discussed in Chapter 5 of Jaburi's master's thesis). As for Pic, this gives a point $* \to \mathrm{Bun}_G$ for every $b \in B(\mathbb{R}, G)$; we write Bun_G^b for the image of this map. For the trivial element $b = 1$, this is the open substack $\mathrm{Bun}_G^1 \simeq */G(\mathbb{R})^{\mathrm{la}}$. When b is basic, the corresponding G-bundle E^b is semistable and its automorphism group G_b is an inner form of G.

Theorem 8. *The inclusion* $\mathrm{Bun}_G^b \subset \mathrm{Bun}_G$ *is the pullback of* $(\mathrm{Bun}_G^b)_{\mathrm{Betti}} \subset (\mathrm{Bun}_G)_{\mathrm{Betti}}$*, which is a locally closed substack. As b varies, the* $(\mathrm{Bun}_G^b)_{\mathrm{Betti}}$ *give a locally finite stratification of* $(\mathrm{Bun}_G)_{\mathrm{Betti}}$*, and therefore the* Bun_G^b *give a locally finite stratification of* Bun_G*.*

Proof. The first statement amounts to the claim that for a strongly totally disconnected $\mathbb{C}$-algebra A, $S = \mathrm{Hom}(A, \mathbb{C})$, and E a G-bundle on $X_{\mathbb{R},A}$, if E is isomorphic to E^b on $X_{\mathbb{R}} \times_{\mathrm{AnSpec}\,\mathbb{R}} \mathrm{AnSpec}\,\mathrm{Cont}(S, \mathbb{R})$ then it is isomorphic to E^b globally. This amounts to lifting an isomorphism of G-torsors along $A \to \mathrm{Cont}(S, \mathbb{C})$, which is a henselian thickening so this is formal.

Via the Harder–Narasimhan filtration by semistable G-bundles, we can reduce to understanding Bun_G^b for b basic and show that this covers the semistable locus. The argument above shows that to classify semistable G-bundles up to isomorphism, we can reduce to the case $A = \mathrm{Cont}(S, \mathbb{C})$, in which case they are just those coming from the absolute case and so are classified by the basic b; since here the automorphism group is $G_b(\mathbb{R})^{\mathrm{la}}$, we can describe Bun_G^b as $*/G_b(\mathbb{R})^{\mathrm{la}}$, and repeating the argument for Proposition 7 for G_b in place of G and identifying $\mathrm{Bun}_{G_b} \simeq \mathrm{Bun}_G$ by shuffling the inner forms (and the Kottwitz set more generally) we conclude that this is also open in Bun_G. Letting b vary gives the semistable locus. □

In the p-adic setting, each connected component of Bun_G contains a unique semistable stratum, which for some purposes lets us restrict attention to the basic case. This does not occur here: components may have either multiple semistable strata or none. In particular, while we saw that all semistable strata are open, the converse does not necessarily hold.

4. Hecke operators

The general strategy to define Hecke operators over a curve X is as follows: first, we define a Hecke stack $\mathrm{Hck}_G \to \mathrm{Bun}_G \times \mathrm{Bun}_G \times X$, parametrizing some sort of modifications of vector bundles; next, we find a version of geometric Satake associating to representations V of $\widehat{G}$ some sheaf $\mathcal{S}_V$ on Hck_G; and we use the resulting sheaf as a kernel on the correspondence defined by the Hecke stack to define the Hecke operator associated to V.

This is the strategy we will carry out, replacing X by Div^1, except that we want to avoid discussing geometric Satake (at least for now). We'll therefore restrict to (representations corresponding to) minuscule cocharacters, for which we can take the structure sheaf on the stratum Hck_μ.

It remains to understand what a modification should mean in this context. We generally think of this as an isomorphism away from a divisor: at that divisor we might have a map which fails to be an isomorphism in one direction or the other, or a correspondence or so forth, depending on the bounding cocharacter. An archetypical example is for a map of vector bundles which is injective everywhere but fails to be surjective at one point, with some finite length cokernel, say length 1; then at that point instead we have a parabolic subgroup stabilizing the partial flag preserved by the modification. We can actually think of modifications like this in general, at least for modifications of the trivial vector bundle: we keep track of flags at certain points of certain types, and these essentially determine a modification of the trivial bundle at those points. Thus the stack of modifications of the trivial vector bundle should classify divisors of the curve (of suitable degree) together with flags at those points.

We can actually make this precise in our language, though it won't be enough to pin down the full Hecke stack: for a fixed minuscule cocharacter μ, the Grassmannian Gr_μ should classify degree 1 divisors $Z \subset X_{\mathbb{R},A}$ together with a type μ flag at Z, i.e. a Z-point of the flag variety $\mathrm{Fl}_\mu = G/P_\mu$. In the language above, this is the isomorphism $\mathrm{Gr}_\mu \simeq \mathrm{Fl}_\mu^\diamondsuit$. (In general for non-minuscule μ, we would expect to have a map $\mathrm{Gr}_\mu \to \mathrm{Fl}_\mu^\diamondsuit$, but not for it to be an isomorphism; making this precise would require giving a precise definition of Gr_μ independent of this result, which we prefer to avoid.)

Thinking of the Grassmannian as classifying modifications $\mathcal{E}^1 \dashrightarrow \mathcal{E}'$ of the trivial bundle $\mathcal{E}^1$, we get a projection $\mathrm{Gr}_\mu \to \mathrm{Bun}_G$. More generally, we allow both bundles to vary: we define the Hecke stack Hck_μ to send a test ring A to the space of G-bundles $\mathcal{E}$ on $\mathcal{X}_{\mathbb{R},A}$; degree 1 divisors $Z \subset X_{\mathbb{R},A}$; and sections $Z \to \mathcal{E} \times^G \mathrm{Fl}_\mu$ of the projection $\mathcal{E} \times^G \mathrm{Fl}_\mu \to \mathcal{E} \to X_{\mathbb{R},A}$ over Z. We think of this as the space of modifications $\mathcal{E} \dashrightarrow \mathcal{E}'$, and therefore we have *two* projections $\mathrm{Hck}_\mu \to \mathrm{Bun}_G$, sending the above data to either $\mathcal{E}$ or $\mathcal{E}'$, as well as the projection to Div^1 sending it to Z. Note that the fiber of the first projection over $\mathcal{E} = \mathcal{E}^1$ is precisely $\mathrm{Fl}_\mu^\diamondsuit \simeq \mathrm{Gr}_\mu$, while the other fibers are twisted versions of the flag variety.

Thus we get a correspondence

$$\begin{array}{ccc} & \mathrm{Hck}_\mu & \\ \swarrow & & \searrow \\ \mathrm{Bun}_G & & \mathrm{Bun}_G \times \mathrm{Div}^1 \end{array}$$

with both projections cohomologically smooth and proper, and so we can define the pull-push along this correspondence to be the Hecke operator

$$T_\mu : D(\mathrm{Bun}_G) \to D(\mathrm{Bun}_G \times \mathrm{Div}^1).$$

Note that by properness and cohomological smoothness it doesn't matter (up to twist) whether we take $*$- or !-functors.

By further projecting along $\mathrm{Div}^1 \to */W_{\mathbb{R}}^{\mathrm{la}}$—which we've seen is, though not an equivalence, reasonably close to one, and we might guess induces an equivalence in this case as happens p-adically—we can replace the target by $D(\mathrm{Bun}_G \times */W_{\mathbb{R}}^{\mathrm{la}}) \simeq D(\mathrm{Bun}_G)^{BW_{\mathbb{R}}^{\mathrm{la}}}$, so Hecke operators give objects equivariant for the Weil group.

We note that Hecke operators incorporate three disparate themes of the talks so far: locally analytic G-representations, as sheaves on Bun_G; L-parameters, as vector bundles on Div^1; and variations of Hodge and twistor structures, via the fiber $\mathrm{Fl}_\mu^\diamondsuit$ and its twists.

We now have all of the necessary objects in hand to say what the categorical archimedean local Langlands conjecture should be. However, we are still missing some key points. First let's sketch what we expect to be true: let $\mathrm{LocSys}_{\widehat{G}}$ be the stack of $\widehat{G}$-local systems on Div^1. Then, parallel to the nonarchimedean case, we expect that we should have an equivalence of ∞-categories between the subcategory of compact objects in $D(\mathrm{Bun}_G)$ and the subcategory of bounded objects with suitable support conditions (e.g. quasicompact, nilpotent singular) in $D(\mathrm{LocSys}_{\widehat{G}})$. We hesitate however to write down a precise statement, not only due to the technical question of working out these conditions but because we lack a spectral action: the equivalence should depend on a choice of Whittaker data (with the structure sheaf on $\mathrm{LocSys}_{\widehat{G}}$ corresponding to the Whittaker sheaf on Bun_G) and be equivariant for the action of something like $\mathbf{Perf}(\mathrm{LocSys}_{\widehat{G}})$, which encodes the action of the Hecke operators on each side. This would necessitate having a more complete theory of geometric Satake in our setting; I don't know of any deep obstruction to this, but I don't think it's been developed thus far.

As a final remark, let's sketch the connection between Hecke operators and the action of $\mathbf{Perf}(\mathrm{LocSys}_{\widehat{G}})$, though a more complete explanation is more complicated. The Hecke operators (once we have a complete geometric Satake picture) should give a family of functors

$$\mathrm{Rep}(\widehat{G} \rtimes W_{\mathbb{R}}^{\mathrm{la}})^I \to \mathrm{End}(D(\mathrm{Bun}_G)^\omega)$$

for finite sets I which is exact, monoidal, and linear over $\mathrm{Rep}(W_{\mathbb{R}}^{\mathrm{la},I})$. On the other hand, suppose we have a family of functors $\mathrm{Rep}(\widehat{G} \rtimes W_{\mathbb{R}}^{\mathrm{la}})^I \to \mathbf{Perf}(\mathrm{LocSys}_{\widehat{G}})^{BW_{\mathbb{R}}^{\mathrm{la},I}}$ with similar properties; then an action of $\mathbf{Perf}(\mathrm{LocSys}_{\widehat{G}})$ on $D(\mathrm{Bun}_G)^\omega$ would induce the action of the Hecke operators by precomposition. Indeed we have such a family, induced via tensor products from the case $I = \{*\}$, where the functor

$$\mathrm{Rep}(\widehat{G} \rtimes W_{\mathbb{R}}^{\mathrm{la}}) \to \mathbf{Perf}(\mathrm{LocSys}_{\widehat{G}})^{BW_{\mathbb{R}}^{\mathrm{la}}}$$

is equivalent, via the Tannakian formalism, to a $W_{\mathbb{R}}^{\mathrm{la}}$-equivariant $\widehat{G} \rtimes W_{\mathbb{R}}^{\mathrm{la}}$-torsor on $\mathrm{LocSys}_{\widehat{G}}$: the latter stack parametrizes $\widehat{G}$-torsors on $S \times \mathrm{Div}^1$ for a test space S, and so there is a universal $\widehat{G}$-torsor on $\mathrm{LocSys}_{\widehat{G}} \times \mathrm{Div}^1$ which as above gives rise to a $W_{\mathbb{R}}^{\mathrm{la}}$-equivariant $\widehat{G} \rtimes W_{\mathbb{R}}^{\mathrm{la}}$-torsor on $\mathrm{LocSys}_{\widehat{G}}$.

Next time, we'll specialize to the simplest nonabelian case, corresponding to the modular curve, and see how the Hecke operators let us produce a version of the correspondence.