Hyperspherical varieties

March 27, 2024

The main geometric objects of interest in [1] are hyperspherical varieties, which are supposed to be (at least a version of) the "right" geometric setting for the relative Langlands program; the prototypical example is $M = T^*(H\backslash G)$ for a suitable subgroup $H \subset G$. This in many ways carries more structure than the quotient alone; in particular hyperspherical varieties conjecturally have a very good notion of duality, pairing (G, M) with the "relative Langlands dual" $(\check{G}, \check{M})$.

Our goal for today is to build up to the definition of hyperspherical varieties, which naturally arise as the unification of various examples of interest; in particular, we can view them as unifying the cases of spherical varieties of interest to typical relative Langlands-type questions with the symplectic vector spaces which we need to incorporate things like the theta correspondence. We'll see how these simultaneously generalize, and give a structure theorem for spherical varieties; this, together with some varia on polarizations and rational forms, takes us through §1, which is the bulk of the new material for the day. In §2, we'll see how the duality we developed for spherical varieties implies a duality for *polarized* hyperspherical varieties, which arise in a certain way from spherical ones; and we'll speculate in §3 about general hyperspherical duality. Today's talk will be light on proofs, but I'll try to communicate which results are known and which are conjectural.

Except where stated otherwise, we will always work over an algebraically closed field $\mathbb{F}$.

1. Hyperspherical varieties

1.1 Motivating examples

In [2], the initial main area of interest for relative Langlands is spherical subgroups $H \hookrightarrow G$, where we study $X = H\backslash G$ as a G-variety. We generalize this to the situation of spherical G-varieties X, which are nice G-varieties satisfying certain nice properties about the orbits.

It turns out that many of the nice properties of spherical varieties X can be interpreted usefully, for our purposes, by instead studying the cotangent spaces T^*X. These carry the structure of *Hamiltonian G-spaces*: in other words it is a smooth symplectic G-variety M equipped with a moment map $\mu : M \to \mathfrak{g}^*$.

We will often want our Hamiltonian G-spaces to be graded, i.e. equipped with a $\mathbb{G}_{\mathrm{m}}$-action commuting with the G-action; following [1], we'll write $\mathbb{G}_{\mathrm{gr}}$ for this $\mathbb{G}_{\mathrm{m}}$ to avoid confusion with groups related to G. The moment map should be $\mathbb{G}_{\mathrm{gr}}$-equivariant for the squaring action on $\mathfrak{g}^*$, with action on the symplectic form by squaring. In the case $M = T^*X$ above, the $\mathbb{G}_{\mathrm{gr}}$-action is by squaring on the fibers.

A related case is the "Whittaker-type case," which can be viewed as twisting $T^*(U\backslash G)$ by a character of the unipotent radical $\psi : U \to \mathbb{G}_{\mathrm{a}}$. Explicitly, we take the fiber of $T^*G \to \mathfrak{u}^*$ over $d\psi \in \mathfrak{u}^*$, and then take the quotient by U. (The $\mathbb{G}_{\mathrm{gr}}$-action is by squaring on fibers, precomposed with left translation by $\lambda^{2\check{\rho}}$.) We'll say a lot more about this kind of construction in the following sections; it is related to Whittaker induction, which we can use more generally to construct *all* hyperspherical varieties (Theorem 1).

This can be reinterpreted as a "ψ-twisted" version of the cotangent complex example as follows. Let $U_0 = \ker \psi$, and set $\Psi = U_0 \backslash G$. This is naturally a $\mathbb{G}_a$-torsor over $U \backslash G$, and is equipped with a moment map $T^*\Psi \to \mathfrak{g}_a^*$; roughly the $\mathbb{G}_a$-quotient of the fiber at $1 \in \mathfrak{g}_a^*$ is defined to be the twisted cotangent bundle $T^*_\Psi(U\backslash G)$, which is the same as the twisted construction of the previous paragraph: $T^*_\psi(U\backslash G) = T^*_\Psi(U\backslash G)$. When Ψ is the trivial $\mathbb{G}_a$-torsor we recover the usual $T^*(U\backslash G)$.

Both of these cases are under the umbrella of [2]. However, we also want to take into account theta-type cases: take M to be a vector space, equipped with a symplectic form ω, and $G \subset \mathrm{Sp}_M$, so the moment map $\mu : M \to \mathfrak{g}^*$ factors through $M \to \mathfrak{sp}_M^*$, sending $m \in M$ to $\mathfrak{sp}_M \ni X \mapsto \frac{1}{2}\langle Xm, m\rangle$, with linear scaling action of $\mathbb{G}_{\mathrm{gr}}$.

In particular, we see that the framework of (graded) Hamiltonian G-spaces incorporates both spherical varieties and symplectic vector spaces. In the next two sections, we'll build up general machinery ("Whittaker induction") which produces examples of the second ("Whittaker") type above, and then define hyperspherical varieties as certain Hamiltonian G-spaces which give a good setting for relative Langlands which includes (almost) all of our examples of interest. We'll then see that all such varieties arise via Whittaker induction.

1.2 Hamiltonian reduction and induction

For spherical G-varieties X, we can produce the Hamiltonian G-space T^*X. If we want to produce something that preserves the symplectic structure but in some sense "quotients out" the G-action, there is a natural construction: rather than directly taking the quotient on T^*X (which would destroy a lot of the structure), we form X/G (assuming it exists in a suitable category) and then take the cotangent space $T^*(X/G)$.

We would like to generalize this to Hamiltonian G-spaces M not necessarily of the form T^*X. One way of rephrasing the above process is to take the fiber of the moment map over 0: this is a subbundle of T^*X and carries a G-action, and one can show that in fact $\mu^{-1}(0)/G \simeq T^*(X/G)$. Thus we can give a formula for this reduction which applies to all Hamiltonian G-spaces M: the *Hamiltonian reduction* of M by G is

$$M /\!\!/\!\!/ G := \mu_M^{-1}(0)/G.$$

More generally, for any $f \in \mathfrak{g}^*$ with G-orbit $\mathcal{O}_f$, we can define

$$M /\!\!/\!\!/_f G := M \times^G_{\mathfrak{g}^*} \mathcal{O}_f.$$

In general, this is a derived symplectic stack, but in cases of interest the action is free and so it is a symplectic variety. For example, we can now define the twisted cotangent complex more precisely: $T^*_\Psi(U\backslash G) = T^*\Psi /\!\!/\!\!/_1 \mathbb{G}_a$.

In the other direction, suppose we have an inclusion (or indeed any morphism) $H \hookrightarrow G$ of algebraic groups and a Hamiltonian H-space M. If $M = T^*X$ then we define the Hamiltonian induction of M to G to be simply $T^*(X \times^H G)$, where G is considered to be a right H-space via $g \cdot h = h^{-1}g$. Again, we'd like to have a version of this that works for arbitrary M, i.e. that depends only on T^*X.

The above construction gives us an idea for how to do this: we have

$$(T^*X \times T^*G) /\!\!/\!\!/ H = T^*(X \times G) /\!\!/\!\!/ H = \mu^{-1}_{T^*(X\times G)}(0)/H \simeq T^*(X \times^H G).$$

This suggests the following construction: the *Hamiltonian induction* of M from H to G is

$$\text{h-Ind}_H^G(M) := (M \times T^*G)/\!\!/\!\!/ H.$$

If we instead took the left action of H on G, this works out to be

$$\text{h-Ind}_H^G(M) = M \times_{\mathfrak{h}^*}^H T^*G \simeq (M \times_{\mathfrak{h}^*} \mathfrak{g}^*) \times^H G.$$

This description makes clear that the Hamiltonian induction is equipped with a projection to $* \times^H G = H\backslash G$, making $\text{h-Ind}_H^G(M)$ a fiber bundle over the homogeneous space $G\backslash H$. If M is graded, then so is $\text{h-Ind}_H^G(M)$ via the diagonal action of $\mathbb{G}_{\text{gr}}$ on $M \times T^*G$, which commutes with H.

One last thing worth mentioning here is the symplectic normal bundle to a G-orbit $\mathcal{O}$ in a symplectic manifold M: this is a vector bundle S over $\mathcal{O}$ with fiber at x given by

$$S = T_x\mathcal{O}^\perp/(T_x\mathcal{O}^\perp \cap T_x\mathcal{O}),$$

which carries an action of the stabilizer $H = G_x$ of x and thus has a moment map $S \to \mathfrak{h}^*$. In particular, if S is a symplectic H-representation, applying this construction to the Hamiltonian induction recovers S: if $H = G_x$, then S is the fiber at x of the symplectic normal bundle to the G-orbit of x in $\text{h-Ind}_H^G(S)$.

1.3 Whittaker induction

Hamiltonian reduction is already reminiscent of the second class of examples we introduced at the beginning. To get our full generality, though, we want to construct a more powerful version of this machinery, Whittaker induction, which depends not only on a subgroup $H \hookrightarrow G$ but in general on a morphism $H \times \mathrm{SL}_2 \to G$.

We'll think of the $\mathrm{SL}_2 \to G$ part on Lie algebras. Fixing a basis for $\mathfrak{sl}_2$ and an (invariant) identification $\mathfrak{g} \simeq \mathfrak{g}^*$, we can think of this as a triple (h, e, f) of elements of $\mathfrak{g}$, given by the image of the basis of $\mathfrak{sl}_2$.

More canonically, we can fix $f \in \mathfrak{g}^*$ and a cocharacter $\varpi : \mathbb{G}_{\text{m}} \to [G, G] \to G$, which induces $h = d\varpi(1) \in \mathfrak{g}$, such that after any invariant identification $\mathfrak{g} \simeq \mathfrak{g}^*$ we have (h, f) elements of an $\mathfrak{sl}_2$-triple (h, e, f) in $\mathfrak{g}$. We call (ϖ, f) an $\mathfrak{sl}_2$-pair; in particular the centralizer of any associated triple is independent of the identification $\mathfrak{g} \simeq \mathfrak{g}^*$, and depends only on (ϖ, f), so it makes sense to require H to be a subgroup of this centralizer to make the H and SL_2-actions commute. We'll sometimes refer to the triple (h, e, f), implicitly fixing an identification $\mathfrak{g} \simeq \mathfrak{g}^*$, but in the end everything will only depend on (ϖ, f).

If $\mathfrak{j}$ is the centralizer of $\mathfrak{sl}_2$ in $\mathfrak{g}$, we have a decomposition

$$\mathfrak{g} = \mathfrak{j} \oplus \overline{\mathfrak{u}} \oplus \mathfrak{u}^0 \oplus \mathfrak{u}$$

where $\overline{\mathfrak{u}} \oplus \mathfrak{u}^0 \oplus \mathfrak{u}$ is the sum of all the irreducible $\mathfrak{sl}_2$-subrepresentations in $\mathfrak{g}$, decomposed into weight spaces for the adjoint action of $h \in \mathfrak{g}$; so for our triple (h, e, f), we have $f \in \overline{\mathfrak{u}}$, or $f \in \mathfrak{u}^*$. Let $\overline{U}$ and U be the associated unipotent subgroups. Since the $\mathbb{G}_{\text{gr}}$-action via ϖ normalizes U, we can consider $\mathfrak{u}$ as a graded Lie algebra.

We'll assume for simplicity that all the weights of h on $\mathfrak{u}$ are even; this isn't really necessary, but simplifies some things. (In general, we let $\mathfrak{u}_+$ be the subspace with weights ≥ 2, which under this assumption is equal to $\mathfrak{u}$; without this assumption, we at various points have to replace the point $*$ with $\mathfrak{u}/\mathfrak{u}_+$.)

Write $*_f$ for the point $* = \operatorname{Spec} \mathbb{F}$ viewed as a trivial Hamiltonian HU-space with moment map sending $* \mapsto f \in \mathfrak{u}^*$. For any Hamiltonian H-space M, we can define a Hamiltonian HU-space $\tilde{M} = M \times *_f$ with trivial U-action on M; then we can define the Whittaker induction of M from H to G to be $\text{h-Ind}_{HU}^G(\tilde{M})$, or explicitly

$$(M \times *_f)^{HU}_{(\mathfrak{h}+\mathfrak{u})^*} T^*G.$$

Note that this is equipped with a natural base point if M is, given by the base point of M on the left and $(f, 1) \in T^*G \simeq \mathfrak{g}^* \times G$. Similar to Hamiltonian induction, one can lift a grading on M to its Whittaker induction, but the f-shift complicates things; we'll discuss this in terms of shearing in a moment.

First, let's give an example. Suppose M and H are trivial, so its Whittaker induction is just the Hamiltonian induction of $*_f$ from U to G. The element $f \in \mathfrak{u}^*$ defines an additive character $U \to \mathbb{G}_a$, which we call ψ. Then by definition the Hamiltonian induction in this case is the quotient of the fiber of T^*G over $f = d\psi$ by U. This is precisely the Whittaker-type example from §1.1! For f trivial, we'd recover the definition of $T^*G /\!\!/\!\!/ U$, which we know should be the same as $T^*(U\backslash G)$, so these can indeed be viewed as twists of the cotangent bundle of $U\backslash G$.

We now turn to the grading, which arises most naturally via shearing. Suppose more generally that we have an action $\varpi : \mathbb{G}_{\text{gr}} \to \operatorname{Aut}(G)$ of $\mathbb{G}_{\text{gr}}$ on G (say on the right), so G is a "graded group." We define a *sheared* Hamiltonian G-space M to be a Hamiltonian G-space with $\mathbb{G}_{\text{gr}}$-action compatible with that on G and $\mathfrak{g}^*$, in the sense that for $x \in M$, $g \in G$, and $\lambda \in \mathbb{G}_{\text{gr}}$, we have

$$x \cdot g \cdot \lambda = x \cdot \lambda \cdot \varpi(\lambda)(g), \qquad \mu(x \cdot \lambda) = \lambda^2 \varpi(\lambda)(\mu(x))$$

where we denote the action of $\mathbb{G}_{\text{gr}}$ on $\mathfrak{g}^*$ induced by ϖ again by ϖ by an abuse of notation. In particular for the trivial action of $\mathbb{G}_{\text{gr}}$ on G, a sheared Hamiltonian space is just a graded Hamiltonian space, i.e. the actions of G and $\mathbb{G}_{\text{gr}}$ on M commute and the moment map μ is equivariant for the squaring action on $\mathfrak{g}^*$.

If M is a graded Hamiltonian space and $\varpi : \mathbb{G}_{\text{gr}} \to G$ is a cocharacter, composing the $\mathbb{G}_{\text{gr}}$-action on M with ϖ gives a new $\mathbb{G}_{\text{gr}}$-action on M, such that M is now sheared with respect to the $\mathbb{G}_{\text{gr}}$-action on G induced by the (right) inner action via ϖ. Indeed if the $\mathbb{G}_{\text{gr}}$-action on G arises via some ϖ in this way, then every sheared Hamiltonian space arises from a graded one by twisting through ϖ.

In our situation above, the cocharacter ϖ induces a $\mathbb{G}_{\text{gr}}$-action on U, which for the squaring action on $\mathfrak{u}^*$ makes $*_f$ a sheared Hamiltonian U-space; putting the trivial $\mathbb{G}_{\text{gr}}$-action on H, this extends to the structure of a sheared Hamiltonian HU-space, and similarly for any graded Hamiltonian H-space M the product $\tilde{M} = M \times *_f$ is then a sheared Hamiltonian HU-space. Thus the generalization of graded spaces to sheared spaces fixes the issue with f-shifting. Thus, for a fixed $\mathfrak{sl}_2$-pair (ϖ, f), we can view Whittaker induction as the following functor from graded Hamiltonian H-spaces to graded Hamiltonian G-spaces:

- given a Hamiltonian H-space M, form the sheared Hamiltonian HU-space $\tilde{M} = M \times *_f$ (which can be thought of as twisting the grading by ϖ and shifting by f);
- via Hamiltonian induction, form the sheared Hamiltonian G-space $\text{h-Ind}_{HU}^G(\tilde{M})$;
- untwist by ϖ to form the *graded* Hamiltonian G-space $\text{h-Ind}_{HU}^G(\tilde{M})$ (which is the same as above as a Hamiltonian G-space, but with a different $\mathbb{G}_{\text{gr}}$-action).

Generalizing our example above where M is trivial, if M is a vector space, i.e. a symplectic H-representation with the scaling action of $\mathbb{G}_{\text{gr}}$, then one can show that its Whittaker induction is a vector bundle over $H\backslash G$. Explicitly, the Whittaker induction in this case works out to be isomorphic to

$$V \times^H G,$$

where

$$V = M \oplus (\mathfrak{h}^\perp \cap \mathfrak{g}_e),$$

compatibly with the various $\mathbb{G}_{\text{gr}}$-actions. Here $\mathfrak{g}_e$ is the kernel of the action of a principal nilpotent e on $\mathfrak{g}^*$.

You may recall from my last talk that for a spherical G-variety X, we were looking to construct a dual group $\check{G}_X$, together with the data of a map $\iota_X : \check{G}_X \times \text{SL}_2 \to \check{G}$ and a $\check{G}_X$-representation V_X about which we were especially vague except that it is somehow constructed from a smaller representation S_X. Although we're not yet in quite the right situation (we haven't yet said anything about duality), we can begin to see where this will come from: if we have a duality for Hamiltonian G-spaces, $M = T^*X$, and $\check{M}$ was also Whittaker induced from a datum $\check{G}_X \times \text{SL}_2 \to \check{G}$ and a $\check{G}_X$-representation S_X, then the formula for V_X from S_X is precisely the fiber of the Whittaker induction.

1.4 Hypersphericality

To have a good theory of relative Langlands duality, the generality of Hamiltonian G-spaces is actually too great; we need to impose some conditions, reflecting our examples of greatest interest from §1.1.

We say that a graded irreducible (smooth) Hamiltonian G-space M is a *hyperspherical variety* if it satisfies the following five conditions:

(1) M is affine;

(2) the field of G-invariant rational functions $\mathbb{F}(M)^G$ is commutative with respect to the Poisson bracket (M is "coisotropic");

(3) the image of the moment map $\mu : M \to \mathfrak{g}^*$ has nonempty intersection with the nilcone of $\mathfrak{g}^*$;

(4) the stabilizer in G of a generic point of M is connected;

(5) the $\mathbb{G}_{\text{gr}}$-action is "neutral."

The rest of the section will consist of (unproven) remarks on these conditions.

The neutrality condition is technical, but what it essentially means is the following: the first four conditions guarantee that there is a unique closed $G \times \mathbb{G}_{gr}$-orbit M_0 in M, whose image under the moment map is a nilpotent orbit $\mathcal{O}_f$, which can be associated to $\mathfrak{sl}_2$-triples (h, e, f), each of which as above produces another $\mathbb{G}_{gr}$-action; the neutrality condition means that on a neighborhood of M_0, the $\mathbb{G}_{gr}$-action agrees with that associated to some $\mathfrak{sl}_2$-triple of $\mathcal{O}_f$. We'll return to it after discussing some of the other conditions.

When M is the Whittaker induction of a symplectic H-representation as in the previous section, it will automatically satisfy conditions (1), (3), and (5). The "coisotropicity condition" (2) should be thought of as the "sphericality" condition; the final condition (4) is more auxiliary. In particular, if X is a smooth affine spherical variety, then T^*X is hyperspherical if and only if it satisfies condition (4), or equivalently if and only if the B-stabilizers of points in the unique open B-orbit in X are connected. (Such spaces will turn out to be the "polarized" hyperspherical varieties.)

In particular let's say a little bit more about condition (2): it is equivalent to either of the following:

- the generic G-orbit on M is coisotropic;
- the generic fiber of $\tilde{\mu}_G : M \to \mathfrak{g}^* \twoheadrightarrow \mathfrak{g}^* /\!\!/ G$ contains an open G-orbit.

In particular the latter condition in the case $M = T^*X$ is reminiscent of the sphericality condition on X, perhaps making some of the claims above more believable. Moreover under this condition it follows that the GIT quotient $M /\!\!/ G$ is isomorphic to the image of $\tilde{\mu}_G$. Then condition (3) implies that 0 is in the image of $\tilde{\mu}_G$, which together with (2) implies that $\tilde{\mu}_G$ is surjective. Assuming all conditions (1) through (4), one can show as claimed that there is a unique $G \times \mathbb{G}_{gr}$-orbit M_0 with nilpotent image under μ.

We can now return to the neutrality condition. Choose x in the closed $G \times \mathbb{G}_{gr}$-orbit $M_0 \subset M$ with image $\mu(x) = f \in \mathfrak{g}^*$, and let $H \subset G$ be the stabilizer of x; all constructions will be independent of the choice of x. The $\mathbb{G}_{gr}$-action on $M_0 \simeq H \backslash G$ commutes with G and so is given by left multiplication by a cocharacter $\varpi : \mathbb{G}_{gr} \to N(H)/H$, which acts on $f \in \mathfrak{g}^*$ by squaring. The neutrality condition (5) is then that (ϖ, f) forms an $\mathfrak{sl}_2$-pair, together with certain more technical conditions. In particular, to a hyperspherical variety M (satisfying all five conditions) we can associate a subgroup $H \subset G$ as well as a commuting $\mathfrak{sl}_2$-pair, which was exactly the data we needed to produce Whittaker induction from H-spaces to G-spaces.

1.5 Structure theorem

We are now nearly ready to state the structure theorem for hyperspherical varieties. Recall that when $H = G_x$, the symplectic normal bundle construction gave a way to recover a symplectic H-representation from its Hamiltonian induction. The following theorem can be viewed as saying that, under the conditions defining hyperspherical varieties, a similar procedure works for Whittaker induction.

Theorem 1. *Let M be a hyperspherical variety, with related notation as in the previous section, and let S be the fiber of the symplectic normal bundle to M_0 at x. Then there* *is a unique $G \times \mathbb{G}_{\mathrm{gr}}$-equivariant isomorphism of Hamiltonian G-spaces between M and the Whittaker induction of S from H to G with respect to the fixed $\mathfrak{sl}_2$-tuple (f, ϖ), sending x to the base point of the Whittaker induction.*

In particular, every hyperspherical variety is Whittaker induced from a symplectic representation. This can be viewed as combining examples of types 2 (Whittaker induction of the trivial space) and 3 (symplectic vector spaces); keeping in mind that type 2 is a twisted version of type 1, this means that (under minor restrictions) all our examples are subsumed by spherical varieties, and in fact in a certain sense generate them.

1.6 Polarization

We maintain the notation as above; in particular we keep S as the symplectic normal bundle to M_0, so M is induced from S. We say that M admits a distinguished polarization if the weight 1 component $\mathfrak{u}_1 \subset \mathfrak{u}$ vanishes (implied by our even weights assumption) and there is a Lagrangian H-stable decomposition

$$S = S^+ \oplus S^-,$$

and we call such a choice a distinguished polarization.

In this situation, as above we can identify $f \in \mathfrak{u}^*$ with $d\psi$ for a character $\psi : U \to \mathbb{G}_{\mathrm{a}}$. Then, letting $X = S^+ \times^{HU} G$, M is the ψ-twisted cotangent bundle of X. Via the conditions (1) - (5) on hyperspherical varieties, one can show that in this case X must be a spherical variety satisfying the analogue of condition (4), i.e. the B-stabilizers of points in the open B-orbit are connected; and the twisted cotangent bundles of any affine smooth spherical X satisfying this condition is hyperspherical. In particular, the data of a distinguished polarization of a hyperspherical variety M can be viewed as an identification $M \simeq T^*_\psi(X)$ for a suitable spherical X. In §2 we'll develop the duality theory for polarized hyperspherical varieties, and this goes a long way towards showing why this is easier than without the polarization data: with it, we can essentially reduce to the duality theory for spherical varieties which we've seen before. However, the same failings of that theory apply: we can't hope for a completely symmetric duality theory for polarized hyperspherical varieties, because we don't have one for spherical varieties! So in §3 we'll try to say something about the general case, where we don't have this structure to rely on.

1.7 Rationality

Before we move on to the duality theory, let's make some brief remarks about the general situation when we don't require everything to be over an algebraically closed field. It is possible to give a (not entirely satisfactory) definition of hyperspherical data over a ring, by essentially turning the structure theorem into a definition. There are some interesting issues of which is the right form to use when not over an algebraically closed field, which we skip over for the sake of time; the main point is that the most straightforward form is not always the right one from the point of view of the local Langlands conjectures. (Speculation: perhaps this is related to the observation in (geometrization of) local Langlands that it's better to simultaneously study all inner forms? Not sure but curious if the experts have thoughts!)

2.1 Duality for polarized hyperspherical varieties

Recall from my last talk: for a spherical G-variety X, we want to associate to it a "Langlands dual group" $\check{G}_X$, together with a morphism $\check{G}_X \times \mathrm{SL}_2 \to \check{G}$. This generalizes the homogeneous situation $X = H\backslash G$, where $\check{G}_X = \check{H}$. We also mentioned the further data of a representation V_X of $\check{G}_X$, which was related to another representation S_X which arises more naturally in the relative setting, but is less natural for classical applications. (The constructions in the spherical case require a lot of careful geometry and Lie algebra work, which truthfully speaking we did not go through in too great detail last time; however we will pretend we did, and treat the spherical case as completely understood.)

Our goal in the next two sections is to generalize this picture to the hyperspherical case: for a hyperspherical G-variety M, we want to associate to it a dual hyperspherical $\check{G}$-variety $\check{M}$. In the case $M = T^*X$, this should reduce in some sense to the spherical case.

With all our machinery built up, the construction is not actually too hard (although the result will be only conjecturally hyperspherical). The key idea is to use the structure theorem: each hyperspherical G-variety M has commuting $H \subset G$, $\mathfrak{sl}_2 \subset \mathfrak{g}$, and a symplectic representation S of H, and in turn arises from such data. In the case $M = T^*X$, we can associate to X the data of commuting $\check{G}_X \subset \check{G}$, $\mathfrak{sl}_2 \subset \check{\mathfrak{g}}$, and the self-dual representation S_X of $\check{G}_X$, which conjecturally should admit a $\check{G}_X$-invariant symplectic form. Conditional on this conjecture, this is exactly the data of a hyperspherical $\check{G}$-variety $\check{M}$, which should be the dual of M: explicitly $\check{M}$ is the Whittaker induction of S_X from $\check{G}_X$ to $\check{G}$ for the fixed $(H, \mathfrak{sl}_2)$-data. (As previously, the role of $\check{G}_X$ is dual to that of H, but we avoid the notation $\check{H}$ since this collides with the Langlands dual of H, which may not always be the same as $\check{G}_X$ outside the homogeneous case.)

For polarized hyperspherical varieties, it would suffice to understand not just the case of $M = T^*X$ but the slightly more general twisted cotangent bundle case $M = T^*_\psi X$. The construction of $\check{G}_X$ and the commuting $\mathfrak{sl}_2 \subset \check{G}$ is very similar; under certain conditions (the "wavefront property") the Weyl group of $\check{G}_{X,\Psi}$ should be the Weyl group associated to the $\mathbb{G}_a$-bundle Ψ viewed as a G-space (recall the construction actually does not require sphericality!). The construction of S_X is just as in the spherical case (which recall is piecemeal via defining certain weights from the geometry of X and taking S_X to be a sum of representations with highest/lowest weights coming from these), with the twist already built in to $\check{G}_X$.

In general, the $\check{M}$ thus constructed is only conjecturally hyperspherical: one would need to know that S_X was symplectic and to check various conditions on X in order to apply the structure theorem out-of-the-box. However we can prove the following interesting result towards condition (4):

Proposition 2. *Let X be a smooth affine spherical variety, $M = T^*X$, and assume that S_X is symplectic (and satisfies certain technical conditions). Then the image of the dual moment map $\check{M} \to \check{\mathfrak{g}}^*$ contains a regular nilpotent element.*

Indeed, this is the element f corresponding to the $\mathfrak{sl}_2 \subset \mathfrak{g}$ arising in the construction of $\check{M}$. This will be useful in our discussion of rationality below.

2.2 Rationality and varia

Once we've introduced duality, there are now two fields involved: the base field $\mathbb{F}$, which we've assumed algebraically closed, and the coefficient field k, also assumed algebraically closed, which the dual side $\check{M}$ and its associated data live over. (Classically, this is the situation of G over a field $\mathbb{F}$, which is often taken to be $\mathbb{C}$ in geometric Langlands but could be a number field or function field for arithmetic applications, and $\check{G}$ over the complex numbers (though for some arithmetic applications we prefer to take ℓ-adic coefficients).) Thus there are two natural questions we could ask, given by allowing either $\mathbb{F}$ or k to fail to be algebraically closed: if M (together with its polarization, if present) is defined over a field $\mathbb{F}$ which is not algebraically closed, is there an induced action of $\mathrm{Gal}(\overline{\mathbb{F}}/\mathbb{F})$ on $\check{M}$ (as there would be on $\check{G}$ in the classical case, to form ${}^L G$)? If k is not algebraically closed, is there a "distinguished" form of $\check{M}$ over k?

Recall from Proposition 2 that we have a regular nilpotent element $f \in \check{\mathfrak{g}}^*$ in the image of the moment map from $\check{M}$; we take this to correspond to the pinning of $\check{G}$. The general principle, related to both questions above, is that at least in the case $M = T^*X$, there should be an element $m \in \check{M}$, which we call a pinning of the hyperspherical variety $\check{M}$, such that $\mu(m) = f$; and the Galois action for $\mathbb{F}$ should preserve m, and similarly there should be a distinguished rational form of $\check{M}$ such that m is a k-rational point.

Naively, for X over a non-algebraically closed field $\mathbb{F}$ one gets a Galois action on the roots and so could try and construct the Galois action on $\check{G}_X$, extending it to an L-group ${}^L G_X$; but one can find examples where this demonstrably gives the wrong thing! The heuristic is that one should think of the Galois action as really on $\check{M}$, not on $\check{G}_X$. Indeed, one can carefully write down Galois actions on each piece of the hyperspherical datum defined over $\mathbb{F}$ (provided 2 is invertible in $\mathbb{F}$), and then the assembled action on the Whittaker induction $\check{M}$ (conjecturally) gives the right thing; one can confirm that it does preserve a pinning $m \in \mu^{-1}(f)$ as per the philosophy above.

3. General hyperspherical duality: speculation

We can now write down the dual hyperspherical variety for any *polarized* hyperspherical variety M; however since the polarization forces M to arise from a spherical variety, one might complain that we haven't done much new (though what we have done is synthesize these duals into the language of hyperspherical varieties, at least conjecturally). We would like to generalize this duality to all hyperspherical varieties, without resorting to polarizations.

In this (speculative) section, we'll develop a notion of "anomaly" which obstructs quantization; then we'll work out what the ideal conjecture for hyperspherical duality would be, and finally give a rational and even integral version.

3.1 Anomaly

The TQFT philosophy is that to each hyperspherical variety (G, M), we should associate an "automorphic quantization" and "spectral quantization" which give rise to the TQFTs $\mathcal{A}_{(G,M)}$ and $\mathcal{B}_{(G,M)}$; the expected overall relative Langlands duality is that the automorphic quantization of (G, M) should be Langlands dual to the spectral quantization of the dual hyperspherical variety $(\check{G}, \check{M})$. Under certain circumstances, there is an obstruction to the automorphic quantization: one instead has to pass to the metaplectic cover, which is not algebraic and so (since all our objects are algebraic) we'd prefer to avoid. (An example is $G = \mathrm{Sp}_{2n}$, $M = \mathbb{A}^{2n}$.)

In certain circumstances, we can detect the splitting of the metaplectic cover algebraically:

Proposition 3. *Let F be a nonarchimedean local field with residue characteristic different from 2, V a symplectic F-vector space, and $H \subseteq \mathrm{Sp}(V)$ an algebraic F-subgroup. If there exists a character $\theta : H \to \mathbb{G}_m$ with*

$$c_2(V) = c_1(\theta)^2$$

in the (absolute!) étale cohomology $H^4_{\text{ét}}(BH, \mathbb{Z}/2)$, then the metaplectic cover of $\mathrm{Sp}(V)$ splits.

We make no direct use of this proposition; but it motivates the following definition. Let G be a reductive group over $\mathbb{C}$, and M any symplectic G-variety. We set

$$c_2 \in H^4_G(M, \mathbb{Z}) \otimes_{\mathbb{Z}} \mathbb{Z}/2$$

to be the G-equivariant second Chern class of (the tangent bundle of) M, considered modulo 2. Say that M is strong anomaly free if $c_2 = 0$ (i.e. the equivariant Chern class vanishes modulo 2), and anomaly free if there exists an *integral* cohomology class $\beta \in H^2_G(M, \mathbb{Z})$ such that $\beta^2 \equiv c_2 \pmod 2$.

The expectation is that if M is anomaly free, it admits an automorphic quantization. (In fact the anomaly is also related to the spectral quantization, and we have the same expectation there.) There are also some physics-based motivations.

When M is a hyperspherical variety (with associated notation as above, e.g. H, S, etc.), we can rephrase this definition in a way more strongly reminiscent of Proposition 3: let $c_2(S)$ be the second Chern class of S as a vector bundle on BH in $H^4(BH, \mathbb{Z})$. Then M is strong anomaly free if and only if $c_2(S) \equiv 0 \pmod 2$, and is anomaly free if and only if there exists a character $\theta : H \to \mathbb{G}_{\mathrm{m}}$ with $c_2(S) \equiv c_1(\theta)^2 \pmod 2$. (This is under our usual evenness assumption; without it we need to replace S by $V = S \oplus \mathfrak{u}/\mathfrak{u}_+$.)

Spelling out $c_2(S)$ in terms of weights, one can check that if M admits a distinguished polarization, then it is anomaly free. If $M = V$ is a symplectic vector space and $G = \mathrm{Sp}(V)$, then M is *not* anomaly free; but it is possible that taking G instead to be certain subgroups of $\mathrm{Sp}(V)$ may make M anomaly free as a G-variety. For example, if $V = W \otimes W'$ where W is an orthogonal vector space of dimension $2n$ and W' is a symplectic vector space of dimension $2m$, so V is symplectic of dimension $4nm$, then the induced subgroup $G = \mathrm{SO}(W) \times \mathrm{Sp}(W') \hookrightarrow \mathrm{Sp}(V)$ makes V an anomaly free (hyperspherical) G-variety, even though it is not anomaly free as a $\mathrm{Sp}(V)$-variety. This recovers the fact that so long as $\dim W$ and $\dim W'$ are even we don't need to pass to the metaplectic cover for the Weil representation.

3.2 Conjectural duality

In this subsection we write down the "ideal" hyperspherical duality conjecture, taking the anomaly into account: this is not quite a conjecture but an expectation, up to minor modifications of the definitions of "hyperspherical" and "anomaly."

Expectation 4. *There exists a duality*

$$(G, M) \longleftrightarrow (\check{G}, \check{M})$$

between anomaly free hyperspherical (G, M) *over* $\mathbb{C}$ *and hyperspherical* $(\check{G}, \check{M})$ *over* $\mathbb{C}$*, such that if* M *admits a distinguished polarization then* $\check{M}$ *arises via the process described in §2.1.*

Some consequences of this expectation are as follows:

- The construction of $\check{M}$ when M admits a principal polarization is independent of the polarization.
- When M admits a distinguished polarization, $(\check{G}, \check{M})$ is anomaly free.
- If M and $\check{M}$ both admit principal polarizations, then M is constructed from $\check{M}$ via the same process.

Some sort of anomaly vanishing condition is certainly necessary, as is true even in the group case, but we may be able to extend to anomalous cases by modifying the conjectures. There is currently (as of [1]) no known example of a dual pair where neither (G, M) nor $(\check{G}, \check{M})$ admits a distinguished polarization, but there is no obvious reason why one couldn't exist.

3.3 Rationality and varia

We've discussed before the question of distinguished models of hyperspherical varieties over non-algebraically closed fields. Our philosophy is that at least in the anomaly free case, there is a universal such model of (G, M) over $\mathbb{Z}$, called the "split form." This should be compatible with our previous rationality remarks; for example, at least away from a finite set of bad primes it should be associated to the distinguished hyperspherical datum for the base change to the algebraic closure. If M and $\check{M}$ are a dual pair with $\check{M} = T^*X$, then we expect that $\mathbb{Z}[\check{M}]$ is the local Plancherel algebra. The split form expectation implies strong Galois invariance properties, and the Plancherel algebra formula would imply that the flatness of $\check{M}$ over $\mathbb{Z}$ is equivalent to certain cohomology groups (I believe $H^*(BH, k)$) being torsion-free.

For automorphic forms, a better construction would be a duality as in Expectation 3.2 taking into account rational structures on both sides (i.e. allowing both $\mathbb{F}$ (for M) and k (for $\check{M}$) to be non-algebraically closed). Following the discussion in §2.2, this should be a certain mutual Galois equivariance respecting certain pinnings (giving rise to L-groups and perhaps some sort of "L-duals" M and ${}^L M$); this picture has not yet been completely worked out. One can write down a definition of split forms under certain conditions, but this is sufficiently preliminary that the authors suggest instead hoping that the split form is obvious or easily determined by its expected properties.

References

[1] David Ben-Zvi, Yiannis Sakellaridis, and Akshay Venkatesh. Relative Langlands duality. *preprint*, 2023.

[2] Yiannis Sakellaridis, Akshay Venkatesh, and Akshay Venkatesh. *Periods and harmonic analysis on spherical varieties*. Société mathématique de France, 2017.