The Langlands Classification and Irreducible Characters for Real Reductive Groups

JEFFREY ADAMS
DAN BARBASCH
DAVID A. VOGAN, JR.

1. Introduction.

In [2] and [3], Arthur has formulated a number of conjectures about automorphic forms. These conjectures would have profound consequences for the unitary representation theory of the group $G(\mathbb{R})$ of real points of a connected reductive algebraic group G defined over $\mathbb{R}$. Our purpose in this book is to establish a few of these local consequences. In order to do that, we have been led to combine the ideas of Langlands and Shelstad (concerning dual groups and endoscopy) with those of Kazhdan and Lusztig (concerning the fine structure of irreducible representations).

We will recall Arthur's conjectures in detail in Chapters 22 and 26, but for the moment it is enough to understand their general shape. We begin by recalling the form of the Langlands classification. Define

$$\Pi(G(\mathbb{R})) \supset \Pi_{\mathrm{unit}}(G(\mathbb{R})) \supset \Pi_{\mathrm{temp}}(G(\mathbb{R})) \tag{1.1}$$

to be the set of equivalence classes of irreducible admissible (respectively unitary or tempered) representations of $G(\mathbb{R})$. Now define

$$\Phi(G(\mathbb{R})) \supset \Phi_{\mathrm{temp}}(G(\mathbb{R})) \tag{1.2}$$

to be the set of Langlands parameters for irreducible admissible (respectively tempered) representations of $G(\mathbb{R})$ (see [34], [10], [1], Chapter 5, and Definition 22.3). To each $\phi \in \Phi(G(\mathbb{R}))$, Langlands attaches a finite set $\Pi_\phi \subset \Pi(G(\mathbb{R}))$, called an *L-packet of representations*. The L-packets Π_ϕ partition $\Pi(G(\mathbb{R}))$. If $\phi \in \Phi_{\mathrm{temp}}(G(\mathbb{R}))$, then the representations in Π_ϕ are all tempered, and in this way one gets also a partition of $\Pi_{\mathrm{temp}}(G(\mathbb{R}))$.

Now the classification of the unitary representations of $G(\mathbb{R})$ is one of the most interesting unsolved problems in harmonic analysis. Langlands' results immediately suggest that one should look for a set between $\Phi(G(\mathbb{R}))$ and $\Phi_{\mathrm{temp}}(G(\mathbb{R}))$ parametrizing exactly the unitary representations. Unfortunately, nothing quite so complete is possible: Knapp has found examples in which some members of the set Π_ϕ are unitary and some are not.

The next most interesting possibility is to describe a set of parameters giving rise to a large (but incomplete) family of unitary representations. This is the local aim of Arthur's conjectures. A little more precisely, Arthur defines a new set

$$\Psi(G/\mathbb{R}) \tag{1.3)(a}$$

of parameters (Definition 22.4). (We write $G/\mathbb{R}$ rather than $G(\mathbb{R})$ because Arthur's parameters depend only on an inner class of real forms, and not on one particular real form.) Now assume that $G(\mathbb{R})$ is quasisplit. Then Arthur defines an inclusion

$$\Psi(G/\mathbb{R}) \hookrightarrow \Phi(G(\mathbb{R})), \quad \psi \mapsto \phi_\psi. \tag{1.3)(b}$$

Write $\Phi_{\mathrm{Arthur}}(G(\mathbb{R}))$ for the image of this inclusion. Then

$$\Phi(G(\mathbb{R})) \supset \Phi_{\mathrm{Arthur}}(G(\mathbb{R})) \supset \Phi_{\mathrm{temp}}(G(\mathbb{R})). \tag{1.3)(c}$$

Roughly speaking, Arthur proposes that $\Psi(G/\mathbb{R})$ should parametrize all the unitary representations of $G(\mathbb{R})$ that are of interest for global applications. More specifically, he proposed the following problems (still for $G(\mathbb{R})$ quasisplit at first).

Problem A. Associate to each parameter $\psi \in \Psi(G/\mathbb{R})$ a finite set $\Pi_\psi \subset \Pi(G(\mathbb{R}))$. This set (which we might call an *Arthur packet*) should contain the L-packet Π_{ϕ_ψ} (cf. (1.3)(b)) and should have other nice properties, some of which are specified below.

The Arthur packet will not in general turn out to be a union of L-packets; so we cannot hope to define it simply by attaching some additional Langlands parameters to ψ.

Associated to each Arthur parameter is a certain finite group A_ψ (Definition 21.4).

Problem B. Associate to each $\pi \in \Pi_\psi$ a non-zero finite-dimensional representation $\tau_\psi(\pi)$ of A_ψ.

Problem C. Show that the distribution on $G(\mathbb{R})$

$$\sum_{\pi \in \Pi_\psi} \left(\epsilon_\pi \dim(\tau_\psi(\pi))\right) \Theta(\pi)$$

is a stable distribution in the sense of Langlands and Shelstad ([35], [48]). Here $\epsilon_\pi = \pm 1$ is also to be defined.

Problem D. Prove analogues of Shelstad's theorems on lifting tempered characters (cf. [48]) in this setting.

Problem E. Extend the definition of Π_ψ to non-quasisplit G, in a manner consistent with appropriate generalizations of Problems B, C, and D.

Problem F. Show that every representation $\pi \in \Pi_\psi$ is unitary.

We give here complete solutions of problems A, B, C, D, and E. Our methods offer no information about Problem F. (In that direction the best results are those of [5], where Problem F is solved for complex classical groups.)

The central idea of the proofs is by now a familiar one in the representation theory of reductive groups. It is to describe the representations of $G(\mathbb{R})$ in terms of an appropriate geometry on an L-group. So let ${}^\vee G$ be the (complex reductive) dual group of G, and ${}^\vee G^\Gamma$ the (Galois form of the) L-group attached to the real form $G(\mathbb{R})$. The L-group is a complex Lie group, and we have a short exact sequence

$$1 \to {}^\vee G \to {}^\vee G^\Gamma \to \mathrm{Gal}(\mathbb{C}/\mathbb{R}) \to 1 \tag{1.4)(a}$$

(The complete definition of the L-group is recalled in Chapter 4.) We also need the Weil group $W_\mathbb{R}$ of $\mathbb{C}/\mathbb{R}$; this is a real Lie group, and there is a short exact sequence

$$1 \to \mathbb{C}^\times \to W_\mathbb{R} \to \mathrm{Gal}(\mathbb{C}/\mathbb{R}) \to 1. \tag{1.4)(b}$$

(The Weil group is not a complex Lie group because the action of the Galois group on $\mathbb{C}^\times$ is the non-trivial one, which does not preserve the complex structure.)

Definition 1.5 ([34], [10]). A *quasiadmissible homomorphism* ϕ from $W_\mathbb{R}$ to ${}^\vee G^\Gamma$ is a continuous group homomorphism satisfying

(a) ϕ respects the homomorphisms to $\mathrm{Gal}(\mathbb{C}/\mathbb{R})$ defined by (1.4); and
(b) $\phi(\mathbb{C}^\times)$ consists of semisimple elements of ${}^\vee G$.

(Langlands' notion of "admissible homomorphism" includes an additional "relevance" hypothesis on ϕ, which will not concern us. This additional hypothesis is empty if $G(\mathbb{R})$ is quasisplit.) Define

$$P({}^\vee G^\Gamma) = \{\, \phi : W_\mathbb{R} \to {}^\vee G^\Gamma \mid \phi \text{ is quasiadmissible} \,\}.$$

Clearly ${}^\vee G$ acts on $P({}^\vee G^\Gamma)$ by conjugation on the range of a homomorphism, and we define

$$\Phi(G/\mathbb{R}) = \{\, {}^\vee G \text{ orbits on } P({}^\vee G^\Gamma) \,\}.$$

(If $G(\mathbb{R})$ is quasisplit, this is precisely the set of parameters in (1.2). In general Langlands omits the "irrelevant" orbits.)

Now a homomorphism ϕ is determined by the value of its differential on a basis of the real Lie algebra of $\mathbb{C}^\times$, together with its value at a single specified element of the non-identity component of $W_\mathbb{R}$; that is, by an element of the complex manifold ${}^\vee\mathfrak{g} \times {}^\vee\mathfrak{g} \times {}^\vee G^\Gamma$. The conditions (a) and (b) of Definition 1.5 amount to requiring the first two factors to be semisimple, and the third to lie in the non-identity component. Requiring that these elements define a group homomorphism imposes a finite number of complex-analytic relations, such as commutativity of the first two factors. Pursuing this analysis, we will prove in Chapter 5

Proposition 1.6. *Suppose ${}^\vee G^\Gamma$ is an L-group. The set $P({}^\vee G^\Gamma)$ of quasiadmissible homomorphisms from $W_\mathbb{R}$ into ${}^\vee G^\Gamma$ may be identified with the set of pairs (y, λ) satisfying the following conditions:*
a) $y \in {}^\vee G^\Gamma - {}^\vee G$, and $\lambda \in {}^\vee\mathfrak{g}$ is a semisimple element;
b) $y^2 = \exp(2\pi i\lambda)$; and
c) $[\lambda, \mathrm{Ad}(y)\lambda] = 0$.

The Langlands classification described after (1.2) is thus already geometric: L-packets are parametrized by the orbits of a reductive group acting on a topological space. Subsequent work of Langlands and Shelstad supports the importance of this geometry. For example, one can interpret some of the results of [48] as saying that the L-packet Π_ϕ may be parametrized using ${}^\vee G$-equivariant local systems on the ${}^\vee G$ orbit of ϕ.

By analogy with the theory created by Kazhdan-Lusztig and Beilinson-Bernstein in [28] and [8], one might hope that information about irreducible characters is encoded by perverse sheaves on the closures of ${}^\vee G$-orbits on $P({}^\vee G^\Gamma)$. Unfortunately, it turns out that the orbits are already closed, so these perverse sheaves are nothing but the local systems mentioned above. On the other hand, one can often parametrize the orbits of several rather different group actions using the same parameters; so we sought a different space with a ${}^\vee G$ action, having the same set of orbits as $P({}^\vee G^\Gamma)$, but with a more interesting geometry.

In order to define our new space, we need some simple structure theory for reductive groups. (This will be applied in a moment to ${}^\vee G$.)

Definition 1.7. Suppose H is a complex reductive group, with Lie algebra $\mathfrak{h}$, and $\lambda \in \mathfrak{h}$ is a semisimple element. Set

$$\mathfrak{h}(\lambda)_n = \{\, \mu \in \mathfrak{h} \mid [\lambda, \mu] = n\mu \,\} \quad (n \in \mathbb{Z}) \tag{1.7)(a}$$

$$\mathfrak{n}(\lambda) = \sum_{n=1,2,\ldots} \mathfrak{h}(\lambda)_n \tag{1.7)(b}$$

$$e(\lambda) = \exp(2\pi i\lambda) \in H. \tag{1.7)(c}$$

The *canonical flat through* λ is the affine subspace

$$\mathcal{F}(\lambda) = \lambda + \mathfrak{n}(\lambda) \subset \mathfrak{h}. \tag{1.7)(d}$$

We will see in Chapter 6 that the canonical flats partition the semisimple elements of $\mathfrak{h}$ — in fact they partition each conjugacy class — and that the map e is constant on each canonical flat. If Λ is a canonical flat, we may therefore write

$$e(\Lambda) = \exp(2\pi i\lambda) \qquad (\text{any } \lambda \in \Lambda). \tag{1.7)(e}$$

Finally, write $\mathcal{F}(\mathfrak{h})$ for the set of all canonical flats in $\mathfrak{h}$.

Definition 1.8 Suppose ${}^\vee G^\Gamma$ is the L-group of a real reductive group (cf. (1.4)). The *geometric parameter space for* ${}^\vee G^\Gamma$ is the set

$$X = X({}^\vee G^\Gamma) = \{\, (y, \Lambda) \mid y \in {}^\vee G^\Gamma - {}^\vee G,\ \Lambda \in \mathcal{F}({}^\vee\mathfrak{g}),\ y^2 = e(\Lambda) \,\}.$$

This is our proposed substitute for Langlands' space $P({}^\vee G^\Gamma)$. The set $\mathcal{F}({}^\vee\mathfrak{g})$ is difficult to topologize nicely, as one can see already for $SL(2)$; this difficulty is inherited by X. To make use of geometric methods we will always restrict to the subspaces appearing in the following lemma.

Lemma 1.9 (cf. Proposition 6.16 below). *In the setting of Definition 1.8, fix a single orbit $\mathcal{O}$ of ${}^\vee G$ on the semisimple elements of ${}^\vee\mathfrak{g}$, and set*

$$X(\mathcal{O}, {}^\vee G^\Gamma) = \{\, (y, \Lambda) \in X \mid \Lambda \subset \mathcal{O} \,\}.$$

Then $X(\mathcal{O}, {}^\vee G^\Gamma)$ has in a natural way the structure of a smooth complex algebraic variety, on which ${}^\vee G$ acts with a finite number of orbits.

The variety $X(\mathcal{O}, {}^\vee G^\Gamma)$ need not be connected or equidimensional, but this will cause no difficulties. We topologize X by making the subsets $X(\mathcal{O}, {}^\vee G^\Gamma)$ open and closed. (It seems likely that a more subtle topology will be important for harmonic analysis, as soon as continuous families of representations are involved.)

Of course the first problem is to check that the original Langlands classification still holds.

Proposition 1.10. (cf. Proposition 6.17 below). *Suppose ${}^\vee G^\Gamma$ is an L-group. Then there is a natural ${}^\vee G$-equivariant map*

$$p : P({}^\vee G^\Gamma) \to X({}^\vee G^\Gamma), \quad p(y, \lambda) = (y, \mathcal{F}(\lambda))$$

inducing a bijection on the level of ${}^\vee G$-orbits. The fibers of p are principal homogeneous spaces for unipotent algebraic groups. More precisely, suppose $x = p(\phi)$. Then the isotropy group ${}^\vee G_\phi$ is a Levi subgroup of ${}^\vee G_x$.

This proposition shows that (always locally over $\mathbb{R}$!) the geometric parameter space X shares all the formal properties of $P({}^\vee G^\Gamma)$ needed for the Langlands classification. In particular, if $G(\mathbb{R})$ is quasisplit, L-packets in $\Pi(G(\mathbb{R}))$ are parametrized precisely by ${}^\vee G$-orbits on X. What has changed is that the orbits on the new space X are not closed; so the first new question to consider is the meaning of the closure relation.

Proposition 1.11. *Suppose $G(\mathbb{R})$ is quasisplit. Let ϕ, $\phi' \in \Phi(G/\mathbb{R})$ be two Langlands parameters, and S, $S' \subset X$ the corresponding ${}^\vee G$-orbits. Then the following conditions are equivalent:*

i) S is contained in the closure of S'.

ii) there are irreducible representations $\pi \in \Pi_\phi$ and $\pi' \in \Pi_{\phi'}$ with the property that π' is a composition factor of the standard representation of which π is the unique quotient.

If ϕ is a tempered parameter, then the orbit S is open in the variety $X(\mathcal{O}, {}^\vee G^\Gamma)$ containing it (cf. Lemma 1.9).

(In the interest of mathematical honesty, we should admit that this result is included only for expository purposes; we will not give a complete proof. That (ii) implies (i) (even for $G(\mathbb{R})$ not quasisplit) follows from Corollary 1.25(b) and (7.11). The other implication in the quasisplit case can be established by a subtle and not very interesting trick. The last assertion follows from Proposition 22.9(b) (applied to an Arthur parameter with trivial $SL(2)$ part).)

Proposition 1.11 suggests the possibility of a deeper relationship between irreducible representations and the geometry of orbit closures on X. To make the cleanest statements, we need to introduce some auxiliary ideas. (These have not been emphasized in the existing literature on the Langlands classification, because they reflect phenomena over $\mathbb{R}$ that are non-existent or uninteresting globally.) The reader should assume at first that G is adjoint. In that case the notion of "strong real form" introduced below amounts to the usual notion of real form, and the "algebraic universal covering" of ${}^\vee G$ is trivial.

Definition 1.12. Suppose G is a complex connected reductive algebraic group. An *extended group for* $G/\mathbb{R}$ is a pair $(G^\Gamma, \mathcal{W})$, subject to the following conditions.

(a) G^Γ is a real Lie group containing G as a subgroup of index two, and every element of $G^\Gamma - G$ acts on G (by conjugation) by antiholomorphic automorphisms.

(b) $\mathcal{W}$ is a G-conjugacy class of triples (δ, N, χ), with

   (1) The element δ belongs to $G^\Gamma - G$, and $\delta^2 \in Z(G)$ has finite order. (Write $\sigma = \sigma(\delta)$ for the conjugation action of δ on G, and $G(\mathbb{R})$ or $G(\mathbb{R}, \delta)$ for the fixed points of σ; this is a real form of G.)

   (2) $N \subset G$ is a maximal unipotent subgroup, and δ normalizes N. (Then N is defined over $\mathbb{R}$; write $N(\mathbb{R}) = N(\mathbb{R}, \delta)$ for the subgroup of real points.)

   (3) The element χ is a one-dimensional non-degenerate unitary character of $N(\mathbb{R})$. (Here "non-degenerate" means non-trivial on each simple restricted root subgroup of N.)

We will discuss this definition in more detail in Chapters 2 and 3. For now it suffices to know that each inner class of real forms of G gives rise to an extended group. The groups $G(\mathbb{R})$ appearing in the definition are quasisplit (because of (b)(2)) and the pair $(N(\mathbb{R}), \chi)$ is the set of data needed to define a Whittaker model for $G(\mathbb{R})$.

Definition 1.13. Suppose $(G^\Gamma, \mathcal{W})$ is an extended group. A *strong real form* of $(G^\Gamma, \mathcal{W})$ (briefly, of G) is an element $\delta \in G^\Gamma - G$ such that $\delta^2 \in Z(G)$ has finite order. Given such a δ, we write $\sigma = \sigma(\delta)$ for its conjugation action on G, and

$$G(\mathbb{R}) = G(\mathbb{R}, \delta)$$

for the fixed points of σ. Two strong real forms δ and δ' are called *equivalent* if they are conjugate by G; we write $\delta \sim \delta'$. (The elements δ of Definition 1.12 constitute a single equivalence class of strong real forms, but in general there will be many others.)

The usual notion of a real form can be described as an antiholomorphic involution σ of G. Two such are equivalent if they differ by the conjugation action of G. This is exactly the same as our definition if G is adjoint. The various groups $G(\mathbb{R}, \delta)$ (for δ a strong real form of $(G^\Gamma, \mathcal{W})$) constitute exactly one inner class of real forms of G.

Definition 1.14. Suppose $(G^\Gamma, \mathcal{W})$ is an extended group. A *representation of a strong real form* of $(G^\Gamma, \mathcal{W})$ (briefly, of G) is a pair (π, δ), subject to
(a) δ is a strong real form of $(G^\Gamma, \mathcal{W})$ (Definition 1.13); and
(b) π is an admissible representation of $G(\mathbb{R}, \delta)$.

Two such representations (π, δ) and (π', δ') are said to be *(infinitesimally) equivalent* if there is an element $g \in G$ such that $g\delta g^{-1} = \delta'$, and $\pi \circ \mathrm{Ad}(g^{-1})$ is (infinitesimally) equivalent to π'. (In particular, this is possible only if the strong real forms are equivalent.) Finally, define

$$\Pi(G^\Gamma, \mathcal{W}) = \Pi(G/\mathbb{R})$$

to be the set of (infinitesimal) equivalence classes of irreducible representations of strong real forms of G.

Lemma 1.15. *Suppose $(G^\Gamma, \mathcal{W})$ is an extended group for G (Definition 1.12). Choose representatives $\{\, \delta_s \mid s \in \Sigma \,\}$ for the equivalence classes of strong real forms of G (Definition 1.13). Then the natural map from left to right induces a bijection*

$$\coprod_{s \in \Sigma} \Pi(G(\mathbb{R}, \delta_s)) \simeq \Pi(G/\mathbb{R})$$

(Definition 1.14; the set on the left is a disjoint union).

This lemma is an immediate consequence of the definitions; we will give the argument in Chapter 2.

The set $\Pi(G/\mathbb{R})$ is the set of representations we wish to parametrize. To do so requires one more definition on the geometric side.

Definition 1.16. Suppose ${}^\vee G^\Gamma$ is the L-group of the inner class of real forms represented by the extended group G^Γ (cf. (1.4) and Definition 1.13). The *algebraic universal covering* ${}^\vee G^{alg}$ is the projective limit of all the finite coverings of ${}^\vee G$. This is a pro-algebraic group, of which each finite-dimensional representation factors to some finite cover of ${}^\vee G$.

With the algebraic universal covering in hand, we can define a complete set of geometric parameters for representations.

Definition 1.17. Suppose G is a connected reductive algebraic group endowed with an inner class of real forms, and ${}^\vee G^\Gamma$ is a corresponding L-group for G. A *complete geometric parameter* for G is a pair

$$\xi = (S, \mathcal{V}),$$

where
(a) S is an orbit of ${}^\vee G$ on $X({}^\vee G^\Gamma)$ (Definition 1.8); and
(b) $\mathcal{V}$ is an irreducible ${}^\vee \widetilde{G}$-equivariant local system on S, for some finite covering ${}^\vee \widetilde{G}$ of ${}^\vee G$.

We may write $(S_\xi, \mathcal{V}_\xi)$ to emphasize the dependence on ξ. In (b), it is equivalent to require $\mathcal{V}$ to be ${}^\vee G^{alg}$-equivariant. Write $\Xi(G/\mathbb{R})$ for the set of all complete geometric parameters.

A slightly different formulation of this definition is sometimes helpful. Fix a ${}^\vee G$-orbit S on X, and a point $x \in S$. Write ${}^\vee G_x^{alg}$ for the stabilizer of x in ${}^\vee G^{alg}$, and define

$$A_S^{loc,alg} = {}^\vee G_x^{alg} / \left({}^\vee G_x^{alg}\right)_0$$

for its (pro-finite) component group. We call $A_S^{loc,alg}$ the *equivariant fundamental group of* S; like a fundamental group, it is defined only up to inner automorphism (because of its dependence on x). Representations of $A_S^{loc,alg}$ classify equivariant local systems on S, so we may also define a complete geometric parameter for G as a pair

$$\xi = (S, \tau),$$

where

(a) S is an orbit of ${}^\vee G$ on $X({}^\vee G^\Gamma)$; and

(b) τ is an irreducible representation of $A_S^{loc,alg}$.

Again we may write (S_ξ, τ_ξ).

Theorem 1.18. *Suppose $(G^\Gamma, \mathcal{W})$ is an extended group for G (Definition 1.12), and ${}^\vee G^\Gamma$ is an L-group for the corresponding inner class of real forms. Then there is a natural bijection between the set $\Pi(G/\mathbb{R})$ of equivalence classes of irreducible representations of strong real forms of G (Definition 1.14), and the set $\Xi(G/\mathbb{R})$ of complete geometric parameters for G (Definition 1.17). In this parametrization, the set of representations of a fixed real form $G(\mathbb{R})$ corresponding to complete geometric parameters supported on a single orbit is precisely the L-packet for $G(\mathbb{R})$ attached to that orbit (Proposition 1.10).*

As we remarked after Proposition 1.6, one can find results of this nature in [48].

For each complete geometric parameter ξ, we define (using Theorem 1.18 and Definition 1.14)

$$(\pi(\xi), \delta(\xi)) = \text{ some irreducible representation parametrized by } \xi \tag{1.19)(a}$$

$$M(\xi) = \text{ standard representation with Langlands quotient } \pi(\xi). \tag{1.19)(b}$$

As a natural setting in which to study character theory, we will also use

$$K\Pi(G/\mathbb{R}) = \text{ free } \mathbb{Z}\text{-module with basis } \Pi(G/\mathbb{R}). \tag{1.19)(c}$$

We will sometimes call this the *lattice of virtual characters*. One can think of it as a Grothendieck group of an appropriate category of representations of strong real forms. In particular, any such representation ρ has a well-defined image

$$[\rho] \in K\Pi(G/\mathbb{R}).$$

By abuse of notation, we will usually drop the brackets, writing for example $M(\xi) \in K\Pi(G/\mathbb{R})$. (All of these definitions are discussed in somewhat more depth in Chapters 11 and 15.)

In order to write character formulas, we will also need a slight variant on the notation of (1.19)(a). Fix a strong real form δ of G, and a complete geometric parameter ξ. By the proof of Lemma 1.15, there is at most one irreducible representation π of $G(\mathbb{R}, \delta)$ so that (π, δ) is equivalent to $(\pi(\xi), \delta(\xi))$. We define

$$\pi(\xi, \delta) = \pi. \tag{1.19(d)}$$

If no such representation π exists, then we define

$$\pi(\xi, \delta) = 0. \tag{1.19)(e}$$

Similarly we can define $M(\xi, \delta)$.

Lemma 1.20 (Langlands — see [54], [56]). *The (image in $K\Pi(G/\mathbb{R})$ of the) set*

$$\{\, M(\xi) \mid \xi \in \Xi(G/\mathbb{R}) \,\}$$

is a basis for $K\Pi(G/\mathbb{R})$.

Because the standard representations of real groups are fairly well understood, it is natural to try to describe the irreducible representations in terms of the standard ones. On the level of character theory, this means relating the two bases $\{\pi(\xi)\}$ and $\{M(\xi)\}$ of $K\Pi(G/\mathbb{R})$:

$$M(\xi) = \sum_{\gamma\in\Xi} m_r(\gamma,\xi)\pi(\gamma). \tag{1.21}$$

(The subscript r stands for "representation-theoretic," and is included to distinguish this matrix from an analogous one to be introduced in Definition 1.22.) Here the *multiplicity matrix* $m_r(\gamma,\xi)$ is what we want. The Kazhdan-Lusztig conjectures (now proved) provide a way to compute the multiplicity matrix, and a geometric interpretation of it — the "Beilinson-Bernstein picture" of [8]. Unfortunately, this geometric interpretation is more complicated than one would like in the case of non-integral infinitesimal character, and it has some fairly serious technical shortcomings in the case of singular infinitesimal character. (What one has to do is compute first at nonsingular infinitesimal character, then apply the "translation principle." The translation principle can introduce substantial cancellations, which are not easy to understand in the Beilinson-Bernstein picture.) We have therefore sought a somewhat different geometric interpretation of the multiplicity matrix. Here are the ingredients. (A more detailed discussion appears in Chapter 7.)

Definition 1.22. Suppose Y is a complex algebraic variety on which the pro-algebraic group H acts with finitely many orbits. Define

$$\mathcal{C}(Y,H) = \text{ category of } H\text{-equivariant constructible sheaves on } Y. \tag{1.22)(a}$$

$$\mathcal{P}(Y,H) = \text{ category of } H\text{-equivariant perverse sheaves on } Y. \tag{1.22)(b}$$

(For the definition of perverse sheaves we refer to [9]. The definition of H-equivariant requires some care in the perverse case; see [38], section 0, or [39], (1.9.1) for the case of connected H.) Each of these categories is abelian, and every object has finite length. (One does not ordinarily expect the latter property in a category of constructible sheaves; it is a consequence of the strong assumption about the group action.) The simple objects in the two categories may be parametrized in exactly the same way: by the set of pairs

$$\xi = (S_\xi, \mathcal{V}_\xi) = (S, \mathcal{V}) \tag{1.22)(c}$$

with S an orbit of H on Y, and $\mathcal{V}$ an irreducible H-equivariant local system on S. The set of all such pairs will be written $\Xi(Y,H)$, the set of *complete geometric parameters for H acting on Y*. Just as in Definition 1.17, we may formulate this definition in terms of the equivariant fundamental group

$$A_S^{loc} = H_y/(H_y)_0 \qquad (y\in S)$$

and its representations. We write $\mu(\xi)$ for the irreducible constructible sheaf corresponding to ξ (the extension of ξ by zero), and $P(\xi)$ for the irreducible perverse sheaf (the "intermediate extension" of ξ — cf. [9], Definition 1.4.22).

The Grothendieck groups of the two categories $\mathcal{P}(Y,H)$ and $\mathcal{C}(Y,H)$ are naturally isomorphic (by the map sending a perverse sheaf to the alternating sum of its cohomology sheaves, which are constructible). Write $K(Y,H)$ for this free abelian group. The two sets $\{P(\xi) \mid \xi\in\Xi\}$ and $\{\mu(\xi)\mid \xi\in\Xi\}$ are obviously bases of their respective Grothendieck groups, but they are *not* identified by the isomorphism. Write $d(\xi)$ for the dimension of the underlying orbit S_ξ. We can write in $K(Y,H)$

$$\mu(\xi) = (-1)^{d(\xi)} \sum_{\gamma\in\Xi(Y,H)} m_g(\gamma,\xi)P(\gamma) \tag{1.22)(d}$$

with $m_g(\gamma,\xi)$ an integer. (The subscript g stands for "geometric.") In this formula, it follows easily from the definitions that

$$m_g(\xi,\xi)=1, \qquad m_g(\gamma,\xi)\neq 0 \text{ only if } S_\gamma \subset \left(\overline{S_\xi} - S_\xi\right) \quad (\gamma\neq\xi). \tag{1.22)(e}$$

The matrix $m_g(\gamma, \xi)$ is essentially the matrix relating our two bases of $K(Y, H)$. It is clearly analogous to (1.21). In each case, we have a relationship between something uncomplicated (the standard representations, or the extensions by zero) and something interesting (the irreducible representations, or the simple perverse sheaves). One can expect the matrix m to contain interesting information, and to be difficult to compute explicitly.

Definition 1.23. In the setting of Definition 1.8, define

$$\mathcal{C}(X({}^\vee G^\Gamma), {}^\vee G^{alg})$$

to be the direct sum over semisimple orbits $\mathcal{O} \subset {}^\vee\mathfrak{g}$ of the categories $\mathcal{C}(X(\mathcal{O}, {}^\vee G^\Gamma), {}^\vee G^{alg})$ of Definition 1.22. The objects of this category are called (by a slight abuse of terminology) ${}^\vee G^{alg}$*-equivariant constructible sheaves on* X. Similarly we define

$$\mathcal{P}(X({}^\vee G^\Gamma), {}^\vee G^{alg}),$$

the ${}^\vee G^{alg}$*-equivariant perverse sheaves on* X. The irreducible objects in either category are parametrized by $\Xi(G/\mathbb{R})$ (cf. Definition 1.22), and we write

$$KX({}^\vee G^\Gamma)$$

for their common Grothendieck group. We write $\mu(\xi)$ and $P(\xi)$ for the irreducible objects constructed in Definition 1.22, or their images in $KX({}^\vee G^\Gamma)$. These satisfy (1.22)(d) and (e).

Since Theorem 1.18 tells us that the two Grothendieck groups $KX({}^\vee G^\Gamma)$ and $K\Pi(G/\mathbb{R})$ have bases in natural one-to-one correspondence, it is natural to look for a functorial relationship between a category of representations of strong real forms of G, and one of the geometric categories of Definition 1.23. We do not know what form such a relationship should take, or how one might hope to establish it directly. What we are able to establish is a formal relationship on the level of Grothendieck groups. This will be sufficient for studying character theory.

Theorem 1.24 *Suppose $(G^\Gamma, \mathcal{W})$ is an extended group for G (Definition 1.12), and ${}^\vee G^\Gamma$ is an L-group for the corresponding inner class of real forms. Then there is a natural perfect pairing*

$$<,>: K\Pi(G/\mathbb{R}) \times KX({}^\vee G^\Gamma) \to \mathbb{Z}$$

between the Grothendieck group of the category of finite length representations of strong real forms of G, and that of ${}^\vee G^{alg}$-equivariant (constructible or perverse) sheaves on X (cf. (1.19) and Definition 1.23). This pairing is defined on the level of basis vectors by

$$< M(\xi), \mu(\xi') >= e(G(\mathbb{R}, \delta(\xi)))\delta_{\xi,\xi'}.$$

Here we use the notation of (1.19) and Definition 1.22. The group $G(\mathbb{R}, \delta(\xi))$ is the real form represented by $M(\xi)$; the constant $e(G(\mathbb{R})) = \pm 1$ is the one defined in [32] (see also Definition 15.8), and the last δ is a Kronecker delta. In terms of the other bases of (1.19) and Definition 1.23, we have

$$< \pi(\xi), P(\xi') >= e(G(\mathbb{R}, \delta(\xi)))(-1)^{d(\xi)}\delta_{\xi,\xi'}.$$

The content of this theorem is in the equivalence of the two possible definitions of the pairing. We will deduce it from the main result of [56]. As an indication of what the theorem says, here are three simple reformulations.

Corollary 1.25.

a) *The matrices m_r and m_g of (1.21) and Definition 1.22(d) are essentially inverse transposes of each other:*

$$\sum_\gamma (-1)^{d(\gamma)} m_r(\gamma, \xi) m_g(\gamma, \xi') = (-1)^{d(\xi)}\delta_{\xi,\xi'}.$$

b) *The multiplicity of the irreducible representation $\pi(\gamma)$ in the standard representation $M(\xi)$ is up to a sign the multiplicity of the local system $\mathcal{V}_\xi$ in the restriction to S_ξ of the Euler characteristic of the perverse sheaf $P(\gamma)$:*

$$m_r(\gamma,\xi) = (-1)^{d(\gamma)-d(\xi)}\sum_i (-1)^i(\textit{multiplicity of } \mathcal{V}_\xi \textit{ in } H^iP(\gamma)\mid_{S_\xi}).$$

c) *The coefficient of the standard representation $M(\gamma)$ in the expression of the irreducible representation $\pi(\xi)$ is equal to $(-1)^{d(\gamma)-d(\xi)}$ times the multiplicity of the perverse sheaf $P(\xi)$ in the expression of $\mu(\gamma)[-d(\gamma)]$.*

Here part (c) refers to the expansion of $\pi(\xi)$ in the Grothendieck group as a linear combination of standard representations (cf. Lemma 1.20); and similarly for $\mu(\gamma)$.

Another way to think of Theorem 1.24 is this.

Corollary 1.26. *In the setting of Theorem 1.24, write*

$$\overline{K} = \overline{K}\Pi(G/\mathbb{R})$$

for the set of (possibly infinite) integer combinations of irreducible representations of strong real forms of G. Then $\overline{K}$ may be identified with the space of $\mathbb{Z}$-linear functionals on the Grothendieck group $KX({}^\vee G^\Gamma)$:

$$\overline{K}\Pi(G/\mathbb{R}) \simeq \mathrm{Hom}_{\mathbb{Z}}(KX({}^\vee G^\Gamma),\mathbb{Z}).$$

In this identification,

a) *the standard representation $M(\xi)$ of $G(\mathbb{R},\delta(\xi))$ corresponds to $e(G(\mathbb{R},\delta(\xi)))$ times the linear functional "multiplicity of $\mathcal{V}_\xi$ in the restriction to S_ξ of the constructible sheaf C;" and*

b) *the irreducible representation $\pi(\xi)$ of $G(\mathbb{R},\delta(\xi))$ corresponds to $e(G(\mathbb{R},\delta(\xi)))(-1)^{d(\xi)}$ times the linear functional "multiplicity of $P(\xi)$ as a composition factor of the perverse sheaf Q."*

Here in (a) we are interpreting $KX({}^\vee G^\Gamma)$ as the Grothendieck group of constructible sheaves, and in (b) as the Grothendieck group of perverse sheaves.

We call elements of $\overline{K}\Pi(G/\mathbb{R})$ *formal virtual characters of strong real forms of G.*

In order to bring Langlands' notion of stability into this picture, we must first reformulate it slightly.

Definition 1.27. In the setting of Definition 1.14 and (1.19), suppose

$$\eta = \sum_{\xi\in\Xi} n(\xi)(\pi(\xi),\delta(\xi))$$

is a formal virtual character. We say that η is *locally finite* if for each strong real form δ there are only finitely many ξ with $n(\xi)\neq 0$ and $\delta(\xi)$ equivalent to δ. Suppose that η is locally finite, and that δ is a strong real form of G. There is a finite set $\pi_1,\ldots,\pi_r$ of inequivalent irreducible representations of $G(\mathbb{R},\delta)$ so that each (π_j,δ) is equivalent to some $(\pi(\xi_j),\delta(\xi_j))$ with $n(\xi_j)\neq 0$. Each of these representations has a character $\Theta(\pi_j)$, a generalized function on $G(\mathbb{R},\delta)$; and we define

$$\Theta(\eta,\delta) = \sum_j n(\xi_j)\Theta(\pi_j),$$

a generalized function on $G(\mathbb{R},\delta)$. This generalized function has well-defined values at the regular semisimple elements of $G(\mathbb{R},\delta)$, and these values determine $\Theta(\eta,\delta)$. In the notation of (1.19)(d,e), we can write

$$\Theta(\eta,\delta) = \sum_\xi n(\xi)\Theta(\pi(\xi,\delta)).$$

We say that η is *strongly stable* if it is locally finite, and the following condition is satisfied. Suppose δ and δ' are strong real forms of G, and $g \in G(\mathbb{R},\delta) \cap G(\mathbb{R},\delta')$ is a strongly regular semsimple element. Then

$$\Theta(\eta,\delta)(g) = \Theta(\eta,\delta')(g).$$

A necessary condition for η to be strongly stable is that each $\Theta(\eta,\delta)$ should be stable in Langlands' sense. Conversely, Shelstad's results in [48] imply that if Θ is a stable finite integer combination of characters on a real form $G(\mathbb{R},\delta)$, then there is a strongly stable η with $\Theta = \Theta(\eta,\delta)$.

Corollary 1.26 gives a geometric interpretation of formal virtual characters. We can now give a geometric interpretation of the notion of stability.

Definition 1.28. In the setting of Definition 1.22, fix an H-orbit $S \subset Y$, and a point $y \in S$. For a constructible sheaf C on Y, write C_y for the stalk of C at y, a finite-dimensional vector space. The map

$$\chi_S^{loc} : \mathrm{Ob}\,\mathcal{C}(Y,H) \to \mathbb{N}, \quad \chi_S^{loc}(C) = \dim(C_y)$$

is independent of the choice of y in S. It is additive for short exact sequences, and so defines a $\mathbb{Z}$-linear map

$$\chi_S^{loc} : K(Y,H) \to \mathbb{Z},$$

the *local multiplicity along* S. If we regard $K(Y,H)$ as a Grothendieck group of perverse sheaves, then the formula for χ_S^{loc} on a perverse sheaf P is

$$\chi_S^{loc}(P) = \sum (-1)^i \dim(H^i P)_y.$$

Any $\mathbb{Z}$-linear functional η on $K(Y,H)$ is called *geometrically stable* if it is in the $\mathbb{Z}$-span of the various χ_S^{loc}.

In the setting of Definition 1.23, a $\mathbb{Z}$-linear functional η on $K(Y,H)$ is called *geometrically stable* if its restriction to each summand $K(X(\mathcal{O}), {}^\vee G^\Gamma)$ is geometrically stable, and vanishes for all but finitely many $\mathcal{O}$.

Theorem 1.29. *In the identification of Corollary 1.26, the strongly stable formal virtual characters correspond precisely to the geometrically stable linear functionals.*

This is an immediate consequence of Corollary 1.26 and Shelstad's description of stable characters in [47]. (It is less easy to give a geometric description of the stable characters on a single real form of G, even a quasisplit one.)

In a sense Arthur's conjectures concern the search for interesting new stable characters. We have now formulated that problem geometrically, but the formulation alone offers little help. The only obvious geometrically stable linear functionals are the χ_S^{loc}. For S corresponding to an L-packet by Proposition 1.10, the corresponding strongly stable formal virtual character is essentially the sum of all the standard representations attached to the L-packet. This sum is stable and interesting, but not new, and not what is needed for Arthur's conjectures. To continue, we need a different construction of geometrically stable linear functionals on $K(Y,H)$.

Definition 1.30. Suppose Y is a smooth complex algebraic variety on which the pro-algebraic group H acts with finitely many orbits. To each orbit S we associate its conormal bundle

$$T_S^*(Y) \subset T^*(Y);$$

this is an H-invariant smooth Lagrangian subvariety of the cotangent bundle. Attached to every H-equivariant perverse sheaf P on Y is a *characteristic cycle*

$$\mathrm{Ch}(P) = \sum_S \chi_S^{mic}(P)\overline{T_S^*(Y)}.$$

Here the coefficients $\chi_S^{mic}(P)$ are non-negative integers, equal to zero unless S is contained in the support of P. One way to construct $\mathrm{Ch}(P)$ is through the Riemann-Hilbert correspondence ([12]): the category of H-equivariant perverse sheaves on Y is equivalent to the category of H-equivariant regular holonomic D-modules on Y, and the characteristic cycle of a D-module is fairly easy to define (see for example [16] or [24]). The functions χ_S^{mic} are additive for short exact sequences, and so define $\mathbb{Z}$-linear functionals

$$\chi_S^{mic} : K(Y,H) \to \mathbb{Z},$$

the *microlocal multiplicity along S.*

Theorem 1.31 (Kashiwara — see [23], [24], Theorem 6.3.1, or [16], Theorem 8.2.) *The linear functionals χ_S^{mic} of Definition 1.30 are geometrically stable. More precisely, for every H-orbit S' such that $\overline{S'} \supset S$ there is an integer $c(S,S')$ so that for every H-equivariant perverse sheaf on Y,*

$$\chi_S^{mic}(P) = \sum_{S'} c(S,S')\chi_{S'}^{loc}(P).$$

Here $\chi_{S'}^{loc}$ is defined in Definition 1.28.

In fact Kashiwara's interest was in an inverted form of this relationship, expressing χ_S^{loc} in terms of the various $\chi_{S'}^{mic}$. (The invertibility of the matrix $c(S,S')$ is an immediate consequence of the facts that $c(S,S) = (-1)^{\dim S}$, and that $c(S,S') \neq 0$ only if $\overline{S'} \supset S$.) We could therefore have defined geometrically stable in terms of the linear functionals χ_S^{mic}.

Perhaps the main difficulty in Theorem 1.31 is the definition of the matrix $c(S,S')$. That definition is due independently to Macpherson in [41]. Although the D-module approach to characteristic cycles is intuitively very simple, it entails some great technical problems (notably that of lifting [44]). We will therefore find it convenient to use a geometric definition of χ_S^{mic} due to MacPherson (see (24.10) and Definition 24.11 below). With this definition, Theorem 1.31 has a very simple proof due to MacPherson; we reproduce it at the end of Chapter 24.

The matrix $c(S,S')$ and its inverse have been extensively studied from several points of view (see for example the references in [16]). If $S \neq S'$ is contained in the smooth part of $\overline{S'}$, then $c(S,S') = 0$. Nevertheless (and in contrast with the multiplicity matrices of (1.22)(d)) there is no algorithm known for computing it in all the cases of interest to us.

Corollary 1.32. *Suppose $({}^\vee G, \mathcal{W})$ is an extended group for G (Definition 1.12), and ${}^\vee G^\Gamma$ is an L-group for the corresponding inner class of real forms. Fix an orbit S of ${}^\vee G$ on $X({}^\vee G^\Gamma)$ (Definition 1.8) (or, equivalently, an L-packet for the quasisplit form of G (Proposition 1.10)). Then the linear functional χ_S^{mic} on $KX({}^\vee G^\Gamma)$ (Definition 1.30) corresponds via Corollary 1.26 to a strongly stable formal virtual representation η_S^{mic}. The irreducible representations of strong real forms occurring in η_S^{mic} are those for which the corresponding perverse sheaf P has $\chi_S^{mic}(P) \neq 0$. This includes all perverse sheaves attached to the orbit S itself, and certain sheaves attached to orbits S' containing S in their closures. With notation as in (1.19) and Definition 1.27, the corresponding stable distribution on $G(\mathbb{R},\delta)$ is*

$$\Theta(\eta_S^{mic},\delta) = e(G(\mathbb{R},\delta)) \sum_{\xi' \in \Xi} (-1)^{d(\xi') - \dim S} \chi_S^{mic}(P(\xi'))\Theta(\pi(\xi',\delta)).$$

In terms of standard representations, this distribution may be expressed as

$$\Theta(\eta_S^{mic},\delta) = e(G(\mathbb{R},\delta))(-1)^{\dim S} \sum_{\xi' \in \Xi} c(S, S_{\xi'})\Theta(M(\xi',\delta)).$$

The set $\{ \eta_S^{mic} \}$ (as S varies) is a basis of the lattice of strongly stable formal virtual representations.

(Recall that the tempered representations correspond to open orbits; in that case χ_S^{mic} is equal to $(-1)^{\dim S}\chi_S^{loc}$, and we get nothing new.) $\blacksquare$

As the second formula of Corollary 1.32 shows, obtaining explicit character formulas for η_S^{mic} amounts to computing the matrix $c(S,S')$ in Theorem 1.31.

To approach Arthur's conjectures, we need an extension of some of the notation in Definitions 1.22 and 1.30.

Definition 1.33. Suppose Y is a smooth complex algebraic variety on which the pro-algebraic group H acts with finitely many orbits. Fix a point y belonging to an H-orbit $S \subset Y$, and write $T^*_{S,y}(Y)$ for the conormal bundle at y to the orbit S. (This is a subspace of the cotangent space at y, having dimension equal to the codimension of S in Y.) The isotropy group H_y acts linearly on $T^*_{S,y}(Y)$; so for any $\nu \in T^*_{S,y}(Y)$, the isotropy group $H_{y,\nu}$ is a pro-algebraic subgroup of H_y. We can therefore form the pro-finite component group $A_{y,\nu} = H_{y,\nu}/(H_{y,\nu})_0$. This family of groups will be locally constant in the variable ν over most of $T^*_{S,y}(Y)$ (Lemma 24.3 below), so we can define the *equivariant micro-fundamental group* A^{mic}_S to be $A_{y,\nu}$ for generic ν.

Attached to every H-equivariant perverse sheaf P on Y is a representation $\tau^{mic}_S(P)$ of A^{mic}_S (Theorem 24.8 and Corollary 24.9 below), of dimension equal to $\chi^{mic}_S(P)$ (Definition 1.30).

The differences between Langlands' original conjectures and those of Arthur amount geometrically to the difference between local geometry (on the orbits S) and microlocal geometry (on the union of the conormal bundles $T^*_S(Y)$); formally, to the difference between the equivariant fundamental group A^{loc}_S and the equivariant micro-fundamental group A^{mic}_S. (Here we are writing Y for the geometric parameter space $X(\mathcal{O}, {}^\vee G^\Gamma)$ containing S.) Notice that if S is open, then $T^*_{S,y}(Y)$ is zero, and A^{mic}_S coincides with the equivariant fundamental group A^{loc}_S. Because tempered representations correspond to open orbits in Theorem 1.18 (Proposition 1.11), we see why the Langlands theory is so effective for tempered representations: the local and the microlocal geometry coincide.

More generally, an L-packet is called *generic* if some irreducible representation in it admits a Whittaker model (see (3.11) and (14.11)–(14.14) below. Here we must understand L-packets as extending over all strong real forms of G.) The generic L-packets (in fact the individual generic representations) were explicitly determined in [30] and [52]. Using those results and Proposition 1.11, it is not difficult to show that an L-packet is generic if and only if the corresponding orbit S is open; that is, if and only if S coincides with the conormal bundle $T^*_S(Y)$. We believe that this collapsing of microlocal to local geometry explains why representation theory should be simpler for generic L-packets.

The existence of the representation $\tau^{mic}_S(P)$ in Definition 1.33 is well-known but quite subtle. We will construct it, following Goresky and MacPherson in [17], using a pair of small spaces $J \supset K$ that reflect the local nature of the singularity of the orbit stratification of Y along S. The representation $\tau^{mic}_S(P)$ will be the hypercohomology of the pair (J, K) with coefficients in P. By a "purity" theorem of Goresky-MacPherson and Kashiwara-Schapira ([17], [26]; see also Chapter 24 below) this hypercohomology is non-zero in only one degree.

Our approach to Arthur's conjectures is now fairly straightforward. Arthur attaches to his parameter ψ a Langlands parameter ϕ_ψ, and thus (by Proposition 1.10) an orbit S_ψ. We define Π_ψ to consist of all those representations appearing in $\eta^{mic}_{S_\psi}$; that is, representations for which the corresponding perverse sheaf has the conormal bundle of S_ψ in its characteristic cycle. We will show (Corollary 27.13) that this agrees with the previous definition of Barbasch and Vogan (in the case of "unipotent" parameters) in terms of primitive ideals. It follows from Proposition 22.9 that Arthur's group A_ψ is isomorphic to a quotient of the equivariant micro-fundamental group $A^{mic,alg}_{S_\psi}$. The difference arises only from our use of the algebraic universal covering of ${}^\vee G$, and for these local purposes our choice seems preferable. We therefore define

$$A^{alg}_\psi = A^{mic,alg}_{S_\psi} \tag{1.34)(a}$$

Definition 1.33 provides a representation $\tau^{mic}_{S_\psi}(P)$ of $A^{mic,alg}_{S_\psi}$, of dimension equal to $\chi^{mic}_{S_\psi}(P)$. Now Problems A, B, C, and E are resolved as special cases of Corollary 1.32 and the preceding definitions. In particular, the representation $\tau_\psi(\pi)$ of Problem B is defined to be

$$\tau_\psi = \tau^{mic}_{S_\psi}(P(\pi)), \tag{1.34)(b}$$

with $P(\pi)$ the irreducible perverse sheaf corresponding to π (Theorem 1.24).

Arthur's Problem C identifies one interesting linear combination of the representations in Π_ψ, using the dimensions of the representations $\tau_\psi(\pi)$. If we use instead other character values, we can immediately define several more. Fix an element $\sigma \in A_\psi^{alg}$, and consider the complex-valued linear functional on equivariant perverse sheaves given by

$$P \mapsto \mathrm{tr}\,[\tau^{mic}_{S_\psi}(P)](\sigma) \qquad (1.34)(c)$$

By Corollary 1.26, this linear functional corresponds to a complex formal virtual representation $\eta_\psi(\sigma)$; that is, to a formal sum with complex coefficients of representations of various strong real forms of G. Just as in Corollary 1.32, we can write this virtual representation on a single strong real form δ as

$$\Theta(\eta_\psi(\sigma),\delta) = e(G(\mathbb{R},\delta)) \sum_{\pi\in\Pi(G(\mathbb{R},\delta))_\psi} (-1)^{d(\pi)-\dim S_\psi}\,\mathrm{tr}\,\tau_\psi(\pi)(\sigma)\Theta(\pi). \qquad (1.34)(d)$$

(More details may be found in Definition 26.8.) Arthur's Problem D asks for a description of the (complex formal virtual) character $\eta_\psi(\sigma)$ in terms of stable characters like $\eta_\psi(1)$ on smaller groups.

To see how this might be possible, we need a long digression about Langlands' functoriality principle. This principle concerns relationships between the representations of real forms of G and those of a smaller reductive group H. The simplest way that such relationships arise is when G and H are equipped with fixed real forms, and $H(\mathbb{R}) \subset G(\mathbb{R})$. In that case we have functors of induction (carrying representations of $H(\mathbb{R})$ to representations of $G(\mathbb{R})$) and restriction (carrying representations of $G(\mathbb{R})$ to representations of $H(\mathbb{R})$). Except in a few very special cases (for example, when G/H is symmetric and we restrict attention to the trivial representation of $H(\mathbb{R})$) these functors are poorly behaved on irreducible representations, and offer little insight into their structure.

A more interesting situation arises when H is a Levi subgroup of a real parabolic subgroup $P = HN$ of G. Then we have the functor of parabolic induction, which carries irreducible representations of $H(\mathbb{R})$ to finite-length representations of $G(\mathbb{R})$. (There are also various "Jacquet functors," analogous to restriction, carrying irreducible representations of $G(\mathbb{R})$ to finite-length (sometimes only virtual) representations of $H(\mathbb{R})$.)

The parabolic induction functors provide the basic model for Langlands functoriality. They exhibit a number of important features of functoriality in general, of which we will mention two. First, they are most simply defined on the level of virtual representations. The reason is that we want to go directly from representations of $H(\mathbb{R})$ to representations of $G(\mathbb{R})$. The definition of parabolic induction requires the choice of a parabolic subgroup P with Levi subgroup H. Different choices of P lead to inequivalent representations, but to the same virtual representations. If we work with virtual representations, we may therefore suppress the dependence on P.

The second feature is actually hidden within the first. To get independence of P even on the level of virtual representations, we must normalize parabolic induction using certain "ρ-shifts." The definition of these shifts requires the extraction of a square root of a character of H (on the top exterior power of the Lie algebra of N). This character happens to be real-valued on $H(\mathbb{R})$, so the square root more or less exists on $H(\mathbb{R})$. (The problem of square roots of -1 can be swept under the rug.) Nevertheless it is clear that (linear) coverings of H are waiting in the wings.

To get correspondences from representations of a small group H to those of a larger group G that behave like parabolic induction, we will use Theorem 1.24. It is therefore natural to begin with extended groups $(G^\Gamma,\mathcal{W})$ and $(H^\Gamma,\mathcal{W}_H)$, and corresponding L-groups ${}^\vee G^\Gamma$ and ${}^\vee H^\Gamma$. Adopting the suggestion from the preceding paragraph that we should seek only a correspondence of virtual representations, we find that we want something like a $\mathbb{Z}$-linear map

$$\epsilon_* : K\Pi(H/\mathbb{R}) \to K\Pi(G/\mathbb{R}). \qquad (1.35)(a)$$

According to Theorem 1.24 (compare Corollary 1.26), such a map is more or less the same as the transpose of a $\mathbb{Z}$-linear map

$$\epsilon^* : KX({}^\vee G^\Gamma) \to KX({}^\vee H^\Gamma). \qquad (1.35)(b)$$

(The "more or less" refers only to issues of finiteness — the difference between K and $\overline{K}$ in Corollary 1.26.) Recall now that the Grothendieck groups in (1.35)(b) are built from equivariant constructible sheaves on geometric parameter spaces. Bearing in mind that H is supposed to be smaller than G, we find that a natural source for a map like (1.35)(b) is the pullback of constructible sheaves by an equivariant morphism of varieties. It is therefore natural to seek a morphism of pro-algebraic groups

$$\epsilon_\bullet : {}^\vee H^{alg} \to {}^\vee G^{alg}, \tag{1.35)(c}$$

and a compatible morphism of geometric parameter spaces

$$X(\epsilon) : X({}^\vee H^\Gamma) \to X({}^\vee G^\Gamma). \tag{1.35)(d}$$

Because the geometric parameter spaces are constructed from the L-groups, this suggests finally the definition at the heart of the functoriality principle.

Definition 1.36 (cf. Definitions 5.1 and 26.3 below). Suppose $(G^\Gamma, \mathcal{W})$ and $(H^\Gamma, \mathcal{W}_H)$ are extended groups (Definition 1.12), with corresponding L-groups ${}^\vee G^\Gamma$ and ${}^\vee H^\Gamma$. An *L-homomorphism* is a morphism

$$\epsilon : {}^\vee H^\Gamma \to {}^\vee G^\Gamma$$

respecting the homomorphisms to $\mathrm{Gal}(\mathbb{C}/\mathbb{R})$ (cf. (1.4)(a)).

Proposition 1.37 (Corollary 6.21 and Proposition 7.18 below). *Suppose $(G^\Gamma, \mathcal{W})$ and $(H^\Gamma, \mathcal{W}_H)$ are extended groups with corresponding L-groups ${}^\vee G^\Gamma$ and ${}^\vee H^\Gamma$, and $\epsilon : {}^\vee H^\Gamma \to {}^\vee G^\Gamma$ is an L-homomorphism.*

a) The restriction of ϵ to the identity component ${}^\vee H$ induces a morphism of pro-algebraic groups

$$\epsilon_\bullet : {}^\vee H^{alg} \to {}^\vee G^{alg}$$

as in (1.35)(c).

b) The map ϵ induces a morphism of geometric parameter spaces

$$X(\epsilon) : X({}^\vee H^\Gamma) \to X({}^\vee G^\Gamma)$$

as in (1.35)(d), compatible with the actions of ${}^\vee H^{alg}$ and ${}^\vee G^{alg}$ and the morphism $\epsilon_\bullet$ of (a).

c) Pullback of constructible sheaves via the morphism $X(\epsilon)$ of (b) defines a $\mathbb{Z}$-linear map

$$\epsilon^* : KX({}^\vee G^\Gamma) \to KX({}^\vee H^\Gamma)$$

as in (1.35)(b).

d) The transpose of the map in (c) is a $\mathbb{Z}$-linear map

$$\epsilon_* : K\Pi(H/\mathbb{R}) \to \overline{K}\Pi(G/\mathbb{R})$$

(cf. Corollary 1.26), which we call Langlands functoriality. It carries representations of strong real forms of H to formal virtual characters of strong real forms of G.

The most important point about this proposition is that the relationship between H and G is entirely on the dual group side; there may be no homomorphism from H to G dual to ϵ in any sense. (A typical example is provided by the split symplectic group $G = Sp(2n)$, and the split orthogonal group $H = SO(2n)$. The corresponding L-groups are ${}^\vee G^\Gamma = SO(2n+1) \times \Gamma$ and ${}^\vee H^\Gamma = SO(2n) \times \Gamma$, so there is an obvious L-homomorphism as in Definition 1.36. For n at least 3, however, any homomorphism from H to G must be trivial.) Similarly, a group homomorphism from H to G need not give rise to an L-homomorphism in general. On the other hand, a real parabolic subgroup $P = HN$ of G *does* provide an L-homomorphism. (Here "real" should be interpreted with respect to one of the special real forms defining the extended group structure on G.) It is not very difficult to show that in this case the Langlands functoriality map of Proposition 1.37 implements the parabolic induction functor discussed earlier; this is implicit in Proposition 26.4.

The primary motivations for studying Langlands functoriality are connected with automorphic representation theory and the trace formula (see for example [2]); we will not discuss them further. However, there are also purely local motivations. One is that the distribution character of the (virtual) representation

$\epsilon_*(\pi_H)$ can be expressed in terms of the distribution character of π_H. We will not explain in detail how to do this; but in Proposition 26.4(b) we will solve the closely related problem of calculating ϵ_* in the basis of standard representations.

In the setting of (1.34), we can now refine slightly our formulation of Arthur's Problem D: the goal is to find $\eta_\psi(\sigma)$ in the image of a Langlands functoriality map ϵ_*. Choose an elliptic element

$$\tilde{s} \in {}^\vee G^{alg} \tag{1.38)(b}$$

representing the class $\sigma \in A_\psi^{alg}$; this is possible by Lemma 26.20. (We may actually choose $\tilde{s}$ to have finite order in every algebraic quotient of ${}^\vee G^{alg}$.) The idea is to use the element $\tilde{s}$ to construct the data (ϵ, H) required for Proposition 1.37. To do this, we begin by defining

$${}^\vee H = \text{identity component of centralizer in } {}^\vee G \text{ of } \tilde{s}. \tag{1.38)(c}$$

We can then define H to be a complex reductive group with dual group ${}^\vee H$. The Arthur parameter ψ defines a Langlands parameter ϕ_ψ, and therefore a geometric parameter x_ψ (Proposition 1.10). As a point of a geometric parameter space, x_ψ is a pair (y_ψ, Λ_ψ) (Definition 1.8), with y_ψ an element of ${}^\vee G^\Gamma - {}^\vee G$ commuting with $\tilde{s}$. Set

$${}^\vee H^\Gamma = \text{group generated by } y_\psi \text{ and } {}^\vee H, \tag{1.38)(d}$$

and let $\epsilon : {}^\vee H^\Gamma \to {}^\vee G^\Gamma$ be the identity map. Because of our choice of $\tilde{s}$, the Arthur parameter ψ takes values in the group ${}^\vee H^\Gamma$; we may write it as ψ_H when we wish to emphasize this.

We are now nearly in the setting of Proposition 1.37. The group ${}^\vee H^\Gamma$ inherits from ${}^\vee G^\Gamma$ a short exact sequence

$$1 \to {}^\vee H \to {}^\vee H^\Gamma \to \mathrm{Gal}(\mathbb{C}/\mathbb{R}) \to 1 \tag{1.38)(e}$$

as in (1.4)(a). In particular, there is an action of $\mathrm{Gal}(\mathbb{C}/\mathbb{R})$ on the based root datum for ${}^\vee H$ and for H (see Definition 2.10). This gives rise to an extended group structure $(H^\Gamma, \mathcal{W}_H)$ (see Chapter 3). The reductive group H, together with the inner class of real forms defined by the extended group structure, is an *endoscopic group* for G (see Definition 26.15 and (26.17) below). Unfortunately, ${}^\vee H^\Gamma$ fails to be an L-group for H^Γ (Definition 4.6 below) for two reasons. First, the sequence (1.38)(e) may admit no distinguished splittings. Second, even if such splittings exist, there is no natural way to fix a ${}^\vee H$-conjugacy class of them. The second of these problems is only a minor nuisance, but the first is somewhat more serious. We will postpone discussing it for a moment in order to formulate a solution to Arthur's Problem D.

Theorem 1.39 (cf. Theorem 26.25 below). *Suppose $(G^\Gamma, \mathcal{W})$ is an extended group with L-group ${}^\vee G^\Gamma$, ψ is an Arthur parameter for G, and $\sigma \in A_\psi^{alg}$. Following (1.34), define a complex formal virtual character $\eta_\psi(\sigma)$ for strong rational forms of G. Choose an elliptic representative $\tilde{s} \in {}^\vee G^{alg}$ for σ, and define ${}^\vee H^\Gamma$, ϵ, and $(H, \mathcal{W}_H)$ as in (1.38). Finally, choose any preimage $\tilde{s}_H \in {}^\vee H^{alg}$ for $\tilde{s}$ under the map ϵ_* of (1.35)(c).*

Assume that ${}^\vee H^\Gamma$ is endowed with the structure of an L-group for $(H, \mathcal{W}_H)$. Then ψ_H may be regarded as an Arthur parameter for H, and $\tilde{s}_H$ represents a class $\sigma_H \in A_{\psi_H}^{alg}$. The complex formal virtual characters of strong rational forms of H and G defined in (1.34) are related by Langlands functoriality (Proposition 1.37) as follows:

$$\eta_\psi(\sigma) = \epsilon_*(\eta_{\psi_H}(\sigma_H)).$$

We will discuss the proof of this result in a moment; first there are some formal issues to address. At (1.34), we asked for a description of $\eta_\psi(\sigma)$ in terms of virtual representations like $\eta_\psi(1)$ on smaller groups. The right side of the formula in Theorem 1.39 involves not $\eta_{\psi_H}(1)$, but rather $\eta_{\psi_H}(\sigma_H)$. The difference is harmless, for the following reason. The element $\tilde{s}_H$ representing σ_H is central in ${}^\vee H^{alg}$ (by (1.38)(c)). Its image in ${}^\vee H$ is fixed by the action of Γ on $Z({}^\vee H)$ (by (1.38)(d)). Now Lemma 26.12 below shows that for any strong real form δ_H of H there is a non-zero complex number $c = \tau_{univ}(\delta_H)(\tilde{s}_H)$ so that

$$\Theta(\eta_{\psi_H}(\sigma_H), \delta_H) = c\Theta(\eta_{\psi_H}(1), \delta_H). \tag{1.40}$$

(Here we use the notation of Definition 1.27.) If δ_H is one of the distinguished (quasisplit) strong real forms of H defining the extended group structure, then $c = 1$.

The second formal issue is the one we postponed a moment ago: what happens when ${}^\vee H^\Gamma$ is not the L-group for H? The answer is implicit in [35]. All of the geometry we have discussed for L-groups can still be carried out on ${}^\vee H^\Gamma$. The resulting geometric parameter space is slightly different from the one constructed using the L-group of H, and so it should correspond to something slightly different from representations of real forms of H. The right objects turn out to be projective representations. The failure of ${}^\vee H^\Gamma$ to be an L-group is measured by a certain cocycle (the "second invariant of an E-group" introduced in Definition 4.6 below). This same cocycle defines a class of projective representations of real forms of H (Definition 10.3). There is a version of Theorem 1.18 (Theorem 10.4 below) relating these representations to geometry on ${}^\vee H^\Gamma$. Once all this extra formalism is assembled, Theorem 1.39 makes sense (and is true) without the hypothesis that ${}^\vee H^\Gamma$ is an L-group. It is this version that is proved in Chapter 26. (Recall also that the use of projective representations in Langlands functoriality is one of the possibilities suggested by the example of parabolic induction.)

A third formal issue is exactly how much Theorem 1.39 is telling us about distribution characters. In light of the remarks after Proposition 1.37 on the computability of Langlands functoriality (cf. Proposition 26.4(a) and (b)), Theorem 1.39 reduces the calculation of the complex formal virtual characters $\eta_\psi(\sigma)$ to the case $\sigma = 1$. This is the case considered in Corollary 1.32, where we found that it was equivalent to an interesting but (in general) unsolved geometric problem. (Another approach that is sometimes effective is described in the next paragraph.)

A fourth issue in the formulation of Theorem 1.39 is the dependence of the endoscopic group H on the choice $\tilde{s}$ of a representative of σ. Simple examples show that this dependence is very strong: different choices lead to very different endoscopic groups. For example, if the identity component of the group ${}^\vee G_\psi$ (the centralizer of the Arthur parameter ψ) is not central in ${}^\vee G$, one can choose a non-central element $\tilde{s}$ to represent $1 \in A_\psi^{alg}$. Theorem 1.39 then computes the stable character $\eta_\psi(1)$ in terms of the same kind of character on a strictly smaller group, bypassing the problem of computing the matrix $c(S, S')$. (More precisely, we are showing how to compute the matrix $c(S, S')$ in the presence of a non-trivial torus action.) The variation of H with the choices should therefore be regarded as a helpful tool, rather than as a weakness of the result.

The last formal issue is that of computing irreducible characters. It is natural to consider the identities (1.34)(d) for fixed ψ and varying σ, and to try to invert them to get formulas for the individual irreducible characters $\Theta(\pi)$ (for $\pi \in \Pi_\psi$) as linear combinations of the formal virtual representations $\eta_\psi(\sigma)$. This was done by Shelstad for tempered representations in [48], and by Barbasch-Vogan for special unipotent representations of complex groups in [7]. In both cases the result is an elementary consequence of two facts peculiar to these cases: the representation $\tau_\psi(\pi)$ of A_ψ^{alg} is irreducible, and the map $\pi \mapsto \tau_\pi$ (from Π_ψ to $\widehat{A}_\psi^{alg}$) is injective. The second fact is certainly not true for Arthur packets in general (see Theorem 27.18 and the remarks after it). It seems likely that the first fails as well; but a counterexample would have to be geometrically rather complicated, and we have not found one. In any case, the identities (1.34)(d) cannot be inverted in general. We do not know whether to expect the existence of a larger set of natural identities that could be inverted.

Here is a sketch of the proof of Theorem 1.39. After unwinding the definition in (1.34) of the virtual representations $\eta_\psi(\sigma)$, and the definition in Proposition 1.37 of the Langlands functoriality map ϵ_*, what must be proved is the following formula. Suppose C is any ${}^\vee G^{alg}$-equivariant constructible complex on the geometric parameter space $X({}^\vee G^\Gamma)$. Then

$$\sum_p (-1)^p \mathrm{tr}\,(\sigma \text{ on } H^p(J, K; C)) = \sum_q (-1)^q \mathrm{tr}\,(\sigma_H \text{ on } H^q(J_H, K_H; \epsilon^*(C)))\,. \tag{1.41}$$

Here $J \supset K$ is the pair of spaces arising in the definition of τ_ψ at (1.34); the cohomology is the hypercohomology of the pair with coefficients in the complex C. (By taking C perverse, we could arrange for this cohomology to be zero except in one degree. But even in that case $\epsilon^*(C)$ would not necessarily be perverse, so we would still need the alternating sum on the right.) The objects on the right are defined similarly for H and the Arthur parameter ψ_H. The spaces J and K may be chosen invariant under the action of $\tilde{s}$; then the left side of (1.41) is a Lefschetz number for $\tilde{s}$ acting on $J \supset K$ (an automorphism of finite order). If we recall that ${}^\vee H$ was defined as the centralizer of $\tilde{s}$, it is perhaps not surprising that $J_H \supset K_H$ turns out to be the fixed point set of $\tilde{s}$. We should be able to compute a Lefschetz number in terms of local contributions along this fixed point set; a theorem of Goresky and MacPherson allows us to do this explicitly, leading to (1.41). The answer is so simple because the automorphism $\tilde{s}$ is of finite order. The geometric details are in Chapter 25 (Theorem 25.8).

Shave and a haircut, two bits.