Automorphic forms (informal course, May-July 2003.)

Kevin Buzzard

April 26, 2012

These are notes from a course I once gave. I wrote the notes as I gave the course and I learnt the theory as I gave the course too, so there may be some confusion in some places as the theory sank in. Read at your own risk!

Last modified 1st June 2004.

Introduction

I wanted to cover both local and global stuff in this course but I only had time to do non-archimedian local stuff; this is basically because I was having to learn it as I went along. My main motivation was initially that whenever I go to a talk about automorphic forms I am lost within minutes because I don't know what a generic tempered automorphic representation of a reductive group over a global field which has no Whittaker model but which is in a certain L-packet blah blah blah, means. But now I have another motivation: it seems that at the time of writing (June 2003) a p-adic Local Langlands conjecture is being formulated (although it has not yet been formulated in precise terms, there is clearly something going on), and one should surely know something about the classical story before trying to understand this new forthcoming one.

I thought that one motivation for the course would be to give a precise statement of the Local Langlands conjecture for a general reductive group—I didn't get this far, although I scraped the surface in the case of $\mathrm{GL}_n(F)$, F non-archimedian local. What I did end up doing was giving lots of basic definitions, and discussing some explicit constructions of supercuspidals. I learnt a lot from the survey article [1] of Bernstein and Zelevinsky, and this definitely influenced the direction that the lectures took.

Perhaps in another series of lectures I could finish the non-archimedian local stuff, do some analysis, and then do some global stuff: I would like one day to understand the definition of an automorphic representation over a global field, show how modular eigenforms give examples of such things, show how the representation-theoretic point of view actually clarifies the theory of modular forms, show (if this is true) that any automorphic form on GL_2 over $\mathbf{Q}$ is either a modular form, a Maass form, or an Eisenstein series (or show why it's not true if it's not true, or show the state of the art if no-one knows whether it's true or not). I would like to say something about multiplicity one and its failure for GSp_4 and its success for GL_n. I would like to talk about global Whittaker models. I would like to see a host of definitions: generic representations, tempered representations, cuspidal representations, and so on. I would like to see the definition of the L-function of an automorphic form and what is known about analytic continuation of these things.

I discovered from reading Chapter 4 of Bump that there is a strong analogy between the local case and the case of a finite field, so I will start with:

1 Representations of GL_2 of a finite field.

I will assume all the representation theory I was taught as an undergraduate. Let G be a finite group and let V be a finite-dimensional vector space over C, an algebraically closed field of characteristic zero. A representation of G on V is a group homomorphism $G \to \mathrm{GL}(V) = \mathrm{Aut}_C(V)$ and we abuse notation by suppressing the name of the map, which is typically ρ. So $\rho(g)v = gv$. Recall that to give a representation of G is just to give a left $C[G]$-module, where $C[G]$ is the group ring. The ring $C[G]$ is actually a sum of matrix rings (useful to know but I don't think we'll use it). A representation of G is irreducible if it's non-zero and doesn't contain any non-trivial sub-$C[G]$-modules. Any finite-dimensional representation of G is a direct sum of irreducible ones and there are only finitely many isomorphism classes of irreducible representations. If G is abelian then a finite-dimensional C-representation is irreducible iff it's 1-dimensional and we call a 1-dimensional representation of a group a character.

If V, W are two representations of G then $\mathrm{Hom}_G(V, W)$ is the C-linear maps $V \to W$ which commute with the G-action, $f(gv) = gf(v)$, so it's a sub-C-space of $\mathrm{Hom}_C(V, W)$. If V is an irreducible $C[G]$-module then $\mathrm{Hom}_G(V, V) = \mathrm{End}_G(V)$ is a C-vector space with a ring structure. If ϕ is any element of $\mathrm{End}_G(V)$ then let λ be an eigenvalue of ϕ; then $\phi - \lambda$ isn't invertible so it's zero. So the natural map $C \to \mathrm{End}_G(V)$ is an isomorphism. Next if V, W are non-isomorphic and irreducible then $\mathrm{Hom}_C(V, W) = 0$ as any non-zero map is an isomorphism. We conclude that if $V = \oplus e_i V_i$ and $W = \oplus f_i V_i$ then the dimension of $\mathrm{Hom}_C(V, W)$ is $\sum_i e_i f_i$. This is just the orthogonality relations.

Recall restriction and induction. If H is a subgroup of G then any G-representation is an H-representation by restriction. On the other hand if V is a $C[H]$-module then $V^G := C[G] \otimes_{C[H]} V$ is a $C[G]$-module and hence a representation of G. Here's a concrete way of realising this module. If V is a representation of H then consider the space X of functions

$$\{f : G \to V | f(hg) = hf(g) \forall g \in G, h \in H\}.$$

This is naturally a C-vector space and it's also naturally a G-module by $gf(\gamma) =$

$f(\gamma g)$. I claim that X is isomorphic to V^G and this is easily checked; if $G = \coprod_i g_i H$ then sending $f \in X$ to $\sum_i g_i \otimes f(g_i^{-1})$ is independent of choice of g_i and works. Bump uses this as his definition of induction; again it's helpful, I think, to keep both definitions in mind. Exercise: check the map is G-linear.

Recall Frobenius reciprocity: if V is an H-module and W is a G-module, for H a subgroup of G, then $\hom_H(V, W_H) = \hom_G(V^G, W)$ and $\hom_H(W_H, V) = \hom_G(W, V^G)$. We're actually spoilt here: when we go onto infinite groups and put continuity constraints on our representations we'll find that only one of these is true, the other fails.

Mackey theory. Again there's not much here, it's just a case of unravelling the definitions. However I didn't see this as an undergraduate so I'll make it more explicit. Let G be finite and let H_1, H_2 be two subgroups of G. Let V_i be a finite-dimensional C-representation of H_i for $1 \leq i \leq 2$. We're interested in $\mathrm{Hom}_G(V_1^G, V_2^G)$. Here is a very concrete way of working it out: this result isn't deep but the proof is messy.

Lemma 1. *There's a canonical C-linear bijection between the following two C-vector spaces:* $\mathrm{Hom}_G(V_1^G, V_2^G)$ *and* $\{\Delta : G \to \mathrm{Hom}_C(V_1, V_2) : \Delta(h_2 g h_1) = h_2 \circ \Delta(g) \circ h_1\}$.

Corollary 2. *The dimensions of the two spaces are the same.*

Proof.

$$\begin{aligned}\mathrm{Hom}_G(V_1^G, V_2^G) &= \mathrm{Hom}_{H_2}(V_1^G, V_2)\\ &= \mathrm{Hom}_{H_2}(C[G] \otimes_{C[H_1]} V_1, V_2)\end{aligned}$$

Now note that there's a canonical surjection $C[G] \otimes_C V_1 \to C[G] \otimes_{C[H_1]} V_1$ sending $g \otimes v$ to $g \otimes v$ and hence a canonical injection $\mathrm{Hom}_{H_2}(C[G] \otimes_{C[H_1]} V_1, V_2) \to \mathrm{Hom}_{H_2}(C[G] \otimes_C V_1, V_2)$. Moreover there's a canonical injection from the latter to $\mathrm{Hom}_C(C[G] \otimes_C V_1, V_2)$ which is just $\mathrm{Hom}_C(C[G], \mathrm{Hom}_C(V_1, V_2))$ which is just $\{\Delta : G \to \mathrm{Hom}_C(V_1, V_2)\}$. So the question is: which Δ correspond to the elements of $\mathrm{Hom}_{H_2}(C[G] \otimes_{C[H_1]} V_1, V_2)$? We see that Δ just has to satisfy $\Delta(h_2 g h_1) = h_2 \Delta(g) h_1 : V_1 \to V_2$. So we're done. For a gory proof involving an explicit construction, see Bump, Proposition 4.1.2. $\square$

Now let k be a finite field of cardinality $q = p^n$ and let G be the group $\mathrm{GL}_2(k)$. What are some examples of irreducible representations of G? There are some easy 1-dimensional examples; $k^\times$ is a finite cyclic group of order $q-1$ and the map $G \to k^\times$ (determinant) gives us $q-1$ 1-dimensional representations. It turns out that the other representations are $q-1$, q and $q+1$-dimensional. How do we get to these? Let T denote the subgroup of diagonal matrices in G and let B denote the upper triangular matrices in G. If χ_1, χ_2 are characters of $k^\times$ then let $B(\chi_1, \chi_2)$ be the induction from B to G of the 1-dimensional representation $\chi : \left(\begin{smallmatrix} a & b \\ 0 & d \end{smallmatrix}\right) \mapsto \chi_1(a)\chi_2(d)$. Note that the index of B in G is $q+1$ because B is the stabiliser of $(1:0)$ under the natural left action of G on $\mathbf{P}^1(k)$, which is transitive. Note also that the order matters (a priori at least). Is $B(\chi_1, \chi_2)$ irreducible? Mackey gives a very nice way of working this out.

Lemma 3 (Bruhat decomposition for GL_2 and SL_2)**.** *If F is any field, G is the group $\mathrm{GL}_2(F)$ or $\mathrm{SL}_2(F)$, and B is the upper triangular matrices in G, then $G = B \coprod BwB$ where $w = \left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right)$.*

Remark. This statement can be massively generalised to reductive groups.

Proof. A dull explicit calculation works in this case. If $g \in BwB$ then one checks that the bottom left entry of g is non-zero so the union is certainly disjoint. All that's left to do is to check that the union is the whole thing and this can be done, for example, by checking that if $g = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ and $D = \det(g)$ and $c \neq 0$ then

$$g = \begin{pmatrix} 1 & a/c \\ 0 & 1 \end{pmatrix} w \begin{pmatrix} -c & -d \\ 0 & -D/c \end{pmatrix}.$$

□

Corollary 4. *The group $\mathrm{SL}_2(F)$ is generated by the diagonal matrices, the upper triangular unipotent matrices, and w.*

Proof. Indeed the diagonal and upper triangular matrices trivially generate B. □

Proposition 5. *Let $\chi_1, \chi_2, \mu_1, \mu_2$ be characters of $k^\times$. Define $e_1 = 1$ if $\chi_1 = \mu_1$ and $\chi_2 = \mu_2$ and $e_1 = 0$ otherwise. Define $e_2 = 1$ if $\chi_1 = \mu_2$ and $\chi_2 = \mu_1$ and $e_2 = 0$ otherwise. Then the dimension of $\mathrm{Hom}_G(B(\chi_1, \chi_2), B(\mu_1, \mu_2))$ is $e_1 + e_2$.*

Proof. Let χ denote the χ_i-representation of B and similarly μ. By the lemma and Mackey, the dimension of the space is equal to the dimension of the space X of maps $\Delta : G \to \mathbf{C}$ such that $\Delta(b_2 g b_1) = \mu(b_2)\Delta(g)\chi(b_1)$ for all $b_1, b_2 \in B$. Such a Δ is determined by $\Delta(1)$ and $\Delta(w)$ so the dimension of this space is clearly at most 2. If $e_1 = 0$ then I claim $\Delta(1) = 0$ and this is because $e_1 = 0$ iff there exists $b \in B$ such that $\chi(b) \neq \mu(b)$ and then $\Delta(1) = \Delta(b^{-1}b) = \mu(b)^{-1}\Delta(1)\chi(b) = c\Delta(1)$ with $c \neq 1$, so $\Delta(1) = 0$. Similarly if $e_2 = 0$ then set $b_1 = \left(\begin{smallmatrix} s & 0 \\ 0 & t \end{smallmatrix}\right)$ and $b_2 = \left(\begin{smallmatrix} t & 0 \\ 0 & s \end{smallmatrix}\right)$ with $s, t \in k^\times$ and such that $\chi(b_1) \neq \mu(b_2)$; then use the fact that $b_1 w b_2^{-1} = w$ to deduce $\Delta(w) = 0$. Finally if $e_1 = 1$ then the function $\Delta : G \to C$ which is zero outside B and which satisfies $\Delta(b) = \chi(b)$ is non-zero and in X; and if $e_2 = 1$ then the function supported on BwB and defined by $\Delta(b_1 w b_2) = \chi(b_1)\mu(b_2)$ is a well-defined non-zero element of X. □

Corollary 6. *$B(\chi_1, \chi_2)$ is irreducible if $\chi_1 \neq \chi_2$ and is the direct sum of two irreducible representations if $\chi_1 = \chi_2$. If $\chi_1 \neq \chi_2$ and $\mu_1 \neq \mu_2$ and $e_1 = e_2 = 0$ in the above notation then $B(\chi_1, \chi_2)$ and $B(\mu_1, \mu_2)$ aren't isomorphic.*

Proof. $\sum_i e_i^2 = 1, 2$ has unique solutions and if both e_i are zero then there are no non-zero maps between the two spaces $B(\chi_1, \chi_2)$ and $B(\mu_1, \mu_2)$. □

In fact one can work out what the two irreducible representations are if $\chi_1 = \chi_2$ because using the description of the induced representation as functions $G \to \mathbf{C}$ such that $f(bg) = \chi(b)f(g)$ we spot that the function f defined by

$f(g) = \chi_1(\det(g))$ is in, and generates a 1-dimensional space which is clearly G-invariant. So what's left is a q-dimensional space and this must be irreducible. Again, although we don't need it, these are all the irreducible q and $q+1$-dimensional representations of G.

I bumped into Martin Liebeck and asked him how to construct the others, namely the ones of dimension $q-1$. He said "aah, these are the cuspidal ones—it's hard to construct them". Here is an amazing construction.

Proposition 7. *Let F be any field. Let S be the following abstract group; it has generators $t(y)$ for any $y \in F^\times$, $n(z)$ for any $z \in F$, and w, and the relations are:*

- $t(y_1)t(y_2) = t(y_1y_2)$
- $n(z_1)n(z_2) = n(z_1 + z_2)$
- $t(y)n(z)t(y^{-1}) = n(y^2z)$
- $wt(y)w = t(-y^{-1})$
- $wn(y)w = t(-y^{-1})n(-y)wn(-y^{-1})$

Then the natural map $S \to \mathrm{SL}_2(F)$ sending $t(y)$ to $\left(\begin{smallmatrix} y & 0 \\ 0 & y^{-1}\end{smallmatrix}\right)$, $n(z)$ to $\left(\begin{smallmatrix} 1 & z \\ 0 & 1\end{smallmatrix}\right)$, w to $\left(\begin{smallmatrix} 0 & 1 \\ -1 & 0\end{smallmatrix}\right)$ is an isomorphism.

Proof. The natural map α is well-defined (easy check) and surjective, because of Bruhat (t and n generate the Borel). We write down a map the other way! We send $\left(\begin{smallmatrix} a & b \\ c & d\end{smallmatrix}\right)$ to $n(a/c)t(-c^{-1})wn(d/c)$ if $c \neq 0$ and to $t(a)n(b/a)$ if $c = 0$ and call this map β. It's certainly a map of sets! To check that this is a well-defined group homomorphism we have to do a mammoth computation. We have to check that if $gh = k$ then $\beta(g)\beta(h) = \beta(k)$. This is five very unpleasant games because β depends on whether we're in the Borel or not. Bump writes down the details of the goriest case, when none of them are in the Borel, although I didn't check it. If g is but the others aren't then we get

$$\begin{pmatrix} a & b \\ 0 & d \end{pmatrix}\begin{pmatrix} A & B \\ C & D \end{pmatrix} = \begin{pmatrix} aA + bC & aB + bD \\ dC & dD \end{pmatrix}$$

and we have to check that in S we have $t(a)n(b/a)n(A/C)t(-C^{-1})wn(D/C) = n(a^2A/C + ba)t(-aC^{-1})wn(D/C)$ that is $t(a)n(b/a)n(A/C) = n(a^2A/C + ba)t(a)$ that is $t(a)n(b/a + A/C) = t(a)n(A/C + b/a)$ and we are there. Yowzer! Finally, because the composite $\alpha\beta$ is the identity on S this proves that the canonical map $S \to \mathrm{SL}_2(F)$ is injective as well. Amazing. $\square$

Now we describe the Weil representation of $\mathrm{SL}_2(k)$, with k a finite field. This representation depends on two things. It depends crucially on the choice of a k-algebra E which is either $k \oplus k$ (with k embedded diagonally, of course) (this is called the "split" case) or the quadratic field extension of k. It depends much less crucially on the choice of a non-trivial additive character $\psi : k \to C^\times$ and so let's fix one once and for all and suppress it from our notation. Let's remark though that there are of course $q-1$ such characters, and if ψ is one of them then for any $a \in k$ the map $\psi_a : k \to C^\times$ defined by $\psi_a(\lambda) = \psi(a\lambda)$ is also an additive character which is non-trivial iff $a \in k^\times$; furthermore if $a \neq b$ then $\psi_a \neq \psi_b$ so the map $a \mapsto \psi_a$ defines a bijection between k and its character group.

While we're here, let's define the subgroup N of G to be the upper triangular unipotent matrices, that is, all those of the form $\left(\begin{smallmatrix} 1 & x \\ 0 & 1 \end{smallmatrix}\right)$, and if V is a representation of N and $a \in k$ then define $V(a)$ to be the subspace of V where N acts via ψ_a; then V is the direct sum of the $V(a)$.

We have fixed such a ψ so we have "Fourier transforms". Define the conjugation map $E \to E$, $e \mapsto \overline{e}$, in the obvious way if E is a field, and as $\overline{(k_1, k_2)} = (k_2, k_1)$ in the split case. Define trace and norm $E \to k$ as sum and product of conjugates as usual. Let W denote the q^2-dimensional C-vector space of all C-valued functions on E. If $f \in W$ then define $\widehat{f} \in W$ by

$$\widehat{f}(x) = \epsilon q^{-1} \sum_{y \in E} f(y)\psi(\mathrm{tr}(\overline{x}y))$$

where ϵ is $+1$ in the split case and -1 in the non-split case. It's W for Weil. The result we need is

Theorem 8 (Weil representation for $\mathrm{SL}_2(k)$)**.** *There's a unique representation $\mathrm{SL}_2(F) \to \mathrm{Aut}_C(W)$ with the following properties:*

- $(\left(\begin{smallmatrix} a & 0 \\ 0 & a^{-1} \end{smallmatrix}\right) f)(x) = f(ax)$ *for all* $a \in F^\times, x \in E$
- $(\left(\begin{smallmatrix} 1 & z \\ 0 & 1 \end{smallmatrix}\right) f)(x) = \psi(zN(x))f(x)$ *for all* $z \in F, x \in E$
- $(wf)(x) = \widehat{f}(x)$

Proof. Clearly there is at most one such representation, because the matrices given will generate $\mathrm{SL}_2(k)$ by Corollary 4. All we have to do is to verify that there is one, that is, that the given properties don't give a contradiction. So we have to check that our five explicit relations are satisfied. The last one is really messy: see, for example, Bump Proposition 4.1.3; the key point is that $\sum_{\alpha \in E} \psi(N(\alpha)) = \epsilon q$ which you check by counting, for example, the number of solutions to $N(\alpha) = x$. What about the other four relations? The first two relations are obviously fine. The third just follows from the fact that if $y \in k$ and $x \in E$ then $N(yx) = y^2 N(x)$. The fourth is interesting and I'll give a proof. I need to check that if $f : E \to C$ and h is the function $x \mapsto \widehat{f}(yx)$ then

$\widehat{h}(x) = f(-y^{-1}x)$ (Fourier inversion formula when $y = 1$). Well

$$\begin{aligned}
\widehat{h}(x) &= \epsilon q^{-1} \sum_{\alpha \in E} h(\alpha)\psi(\mathrm{tr}(\overline{x}\alpha)) \\
&= \epsilon q^{-1} \sum_{\alpha} \widehat{f}(y\alpha)\psi(\mathrm{tr}(\overline{x}\alpha)) \\
&= q^{-2} \sum_{\alpha,\beta \in E} f(\beta)\psi(\mathrm{tr}(\overline{y\alpha}\beta))\psi(\mathrm{tr}(\overline{x}\alpha)) \\
&= q^{-2} \sum_{\beta} f(\beta) \sum_{\alpha} \psi(y\alpha\overline{\beta} + \overline{x}\alpha)
\end{aligned}$$

and if $e \in E^\times$ then $\sum_{\alpha \in E} \psi(\mathrm{tr}(e\alpha)) = \sum_{\beta \in E} \psi(\beta) = 0$ because trace is a surjective linear map $E \to k$ as E/k is separable. So the sum over α vanishes unless β is $-x/y$ in which case it's of course q^2, which is what we need. $\square$

So now we have this big q^2-dimensional representation W of $\mathrm{SL}_2(k)$. Later on we will bump it up to a representation of $\mathrm{GL}_2(k)$; it will turn out that if E is split then W will be composed mostly of the $q+1$-dimensional representations of $\mathrm{GL}_2(k)$ that we have seen already—but if E is anisotropic (that is, not split) then we will get some new ones. Here's the idea. Let E_1 denote the kernel of $N : E^\times \to k^\times$. Let χ be a character of $E^\times$ that doesn't factor through the norm (that is, such that χ isn't trivial on E_1). Now define $W(\chi)$ to be

$$\{f \in W : f(yx) = \chi(y)^{-1} f(x)\, \forall y \in E_1^\times\}.$$

Easy check: if E is anisotropic then $f \in W(\chi)$ has $f(0) = 0$ and E_1 of order $q+1$ is acting transitively on $E^\times$ so the dimension of $W(\chi)$ is $q-1$. In the split case then E_1 has order $q-1$ and the dimensions are all $q+1$ (note that if $q = 2$ then no χ exists in the split case).

One checks easily that the generators of $\mathrm{SL}_2(k)$ send $W(\chi)$ to itself (checking w is a bit of fun, use the fact that E_1 has order $q \pm 1$), so $W(\chi)$ has an action of $\mathrm{SL}_2(k)$. How do we extend to $\mathrm{GL}_2(k)$? Let's try to do this by, for $d \in k^\times$, letting $g = \left(\begin{smallmatrix} d & 0 \\ 0 & 1 \end{smallmatrix}\right)$ act via $(gf)(x) = \chi(b)f(bx)$ where b is any element in E whose norm is d. This is a well-defined element of $W(\chi)$. Now if $g \in \mathrm{GL}_2(k)$ and $d = \det(g)$ then $g = \gamma\left(\begin{smallmatrix} d & 0 \\ 0 & 1 \end{smallmatrix}\right)$ where γ has determinant 1 and so we define the action of g in the obvious way. We have to check it's an action and this boils down to checking that if $\gamma \in \mathrm{SL}_2(k)$ and $d \in k^\times$ then setting $\delta = \left(\begin{smallmatrix} d & 0 \\ 0 & 1 \end{smallmatrix}\right)$, the two actions of $\delta^{-1}\gamma\delta$ that we have defined, coincide. So for fixed δ let X be the subset of $\mathrm{SL}_2(k)$ such that $\delta^{-1}x\delta$ and $\delta^{-1} \circ x \circ \delta$ give the same map. It's easy to check that X is a subgroup of $\mathrm{SL}_2(k)$ and that it contains all the generators of Corollary 4 so we're done. We have representations.

I will not prove that in the split case we have seen these representations before! It's easy though (Bump Proposition 4.1.4) (note that a character of $E^\times$ that doesn't factor through the norm is just two non-equal characters of $k^\times$). In the anisotropic case we have some $q-1$-dimensional representations.

Lemma 9. *If V is any representation of G (or even just the upper triangular matrices in G) and we restrict V to a representation of N and write, for $a \in E$, $V(a)$ for $\{v \in V : \left(\begin{smallmatrix}1 & x\\ 0 & 1\end{smallmatrix}\right)v = \psi(ax)v\}$, then all the $V(a)$ have the same dimension for $a \neq 0$.*

Proof. If $v \in V(a)$ and $d \in k^\times$ then consider $v' = \left(\begin{smallmatrix}d^{-1} & 0\\ 0 & 1\end{smallmatrix}\right)v$; we see that $\left(\begin{smallmatrix}1 & z\\ 0 & 1\end{smallmatrix}\right)v' = \left(\begin{smallmatrix}d^{-1} & 0\\ 0 & 1\end{smallmatrix}\right)\left(\begin{smallmatrix}1 & zd\\ 0 & 1\end{smallmatrix}\right)v = \psi(azd)v'$ so $v' \in V(ad)$ and so $\left(\begin{smallmatrix}d^{-1} & 0\\ 0 & 1\end{smallmatrix}\right)$ gives an isomorphism $V(a) \to V(ad)$. $\square$

Corollary 10. *In the anisotropic case, the $W(\chi)$ are irreducible.*

Proof. Firstly note that there is no element of $W(\chi)$ which is invariant under $N := \left(\begin{smallmatrix}1 & *\\ 0 & 1\end{smallmatrix}\right)$ because if f is N-invariant then $f(0) = 0$ as χ is non-trivial on E_1, and for any $x \neq 0$ we can choose $z \in k^\times$ such that $\psi(zN(x)) \neq 1$ and then unravel the definition of the action of $\left(\begin{smallmatrix}1 & z\\ 0 & 1\end{smallmatrix}\right)$ to prove that $f(x) = 0$. Now consider V an irreducible G-submodule of $W(\chi)$; there exists $a \in k$ such that $V(a)$ is non-zero, and we've just shown $a \neq 0$; by the previous lemma $V(a)$ is non-zero for all $a \in k^\times$; so the dimension of V is at least $q-1$. This is enough. $\square$

Remark. It turns out that there's more to this trick than meets the eye. The $q-1$-dimensional Vs had no N-invariant fixed vector and hence their duals don't have an N-invariant fixed vector either hence there's no non-zero linear map $l : V \to C$ such that $l(nv) = l(v)$ for all $n \in N$. Compare this with the $B(\chi_1, \chi_2)$ which certainly contain N-invariant vectors ($B(\chi_1, \chi_2)$ as a B-module certainly contains the trivial representation of N because it contains the original 1-dimensional representation of B by Frobenius reciprocity). Thus the representations fall into two types—those with a non-zero N-invariant functional and those without. I think that in the local case this is how we detect supercuspidality but I haven't read that far in Bump yet!

Definition. *If V is any finite-dimensional representation of G, let's define $A(V)$ to be the dimension of the N-invariant subspace of V, that is, the dimension of $V(0)$, and $B(V)$ to be the dimension of $V(a)$ for any $a \in k^\times$; by Lemma 9 this is independent of the choice of a.*

We have just seen in corollary 10 that if $V = W(\chi)$ then $A(V) = 0$ and $B(V) = 1$. In general, Lemma 9 implies that

$$A(V) + (q-1)B(V) = \dim(V).$$

Exercise: show in the anisotropic case that $W(\chi)$ and $W(\chi^q)$ are isomorphic, via the map sending f to the function $\overline{f}$ defined by $\overline{f}(x) = f(x^q)$.

Bump claims in his book that all the "missing" $q-1$-dimensional representations of G are of the form $W(\chi)$ but his proof is only complete if one can show that for χ_1, χ_2 as above, if $W(\chi_1)$ and $W(\chi_2)$ are isomorphic then $\chi_1 \in \{\chi_2, \chi_2^q\}$. To my annoyance, I don't see any proof of this at all at the minute, although it is doubtless true. Let us assume it.

Assuming it, we have constructed $\big((q^2-1)-(q-1)\big)/2$ new representations, all of which are $q-1$-dimensional, and now the formula

$$(q-1)+(q^2-q)/2(q-1)^2+(q-1)q^2+(q-1)(q-2)/2(q+1)^2=(q^2-1)(q^2-q)$$

implies that we've found all the representations of G.

Let's look more closely at the irreducible representations that we have seen constructed, as representations of N. In particular, let's compute $A(V)$ and $B(V)$ for all the irreducible representations of G that we have computed. For a given V, we can compute $B(V)$ from $A(V)$ and to compute $A(V)$ all we have to do is to count N-fixed vectors, that is, compute the dimension of $\mathrm{Hom}_N(V,1)$. This is easy if V is of the form $\chi\circ\det$; $A(V)=1$ and $B(V)=0$. We have done the case of $W(\chi)$ earlier. We can do all other cases using Frobenius reciprocity and Mackey theory: it suffices to deal with the case $V=B(\chi_1,\chi_2)$ and $\mathrm{Hom}_N(V,1)=\mathrm{Hom}_G(V,1^G)$ by Frobenius reciprocity.

Lemma 11. *$B(\chi_1,\chi_2)$ contains the trivial representation of N with multiplicity 2. In other words, $A(B(\chi_1,\chi_2))=2$.*

Proof. We do this by observing that $G=B\coprod NwB$ by Lemma 3 and now by Mackey all we have to do is to compute the space of functions $f:G\to C$ such that $f(ngb)=\chi(b)f(g)$. The same arguments as before show that this space has dimension at most 2, but in fact it does always have dimension 2 because one can assign $f(1)$ and $f(w)$ arbitrarily and then bootstrap up. To check this one has to check that if $nb=1$ then $\chi(b)=1$ but this is clear because b must be unipotent, and also one has to check that if $nwb=w$ then b is unipotent, but this follows because $n=wb^{-1}w^{-1}$ so n and b^{-1} have the same eigenvalues. □

We summarise what we have:

Proposition 12. *The irreducible representations of G are as follows:*

- *The 1-dimensional representations $\chi\circ\det$, with $A(V)=1$ and $B(V)=0$*
- *The $(q^2-q)/2$ irreducible $q-1$-dimensional representations of the form $V=W(\chi)$ have $A(V)=0$ and $B(V)=1$.*
- *The $q-1$ Steinberg q-dimensional representations have $A(V)=B(V)=1$.*
- *The $(q-1)(q-2)/2$ irreducible representations of dimension $q+1$ have $A(V)=2$ and $B(V)=1$.*

Note that every representation V not of the form $\chi\circ\det$ has $B(V)=1$. So

Corollary 13. *If $\psi:N\to C^\times$ is defined by $\psi\left(\begin{smallmatrix}1&z\\0&1\end{smallmatrix}\right)=\psi(z)$ then the induced representation $\mathrm{Ind}_N^G(\psi)$ is just the direct sum of all the representations not of the form $\chi\circ\det$.*

Proof. Frobenius reciprocity. □

As a check, one sees that we should have an equality

$$(q-1)(q^2-q)/2 + q(q-1) + (q+1)(q-1)(q-2)/2 = (q^2-1)(q^2-q)/q$$

which indeed we do.

Corollary 14. *If V is a representation of G which isn't of the form $\chi\circ\det$ then there is a unique set of functions $X : \{f : G \to C\}$ with the following properties:*

- *X is a C-vector space and if $f \in X$ and $g \in G$ then the function (gf) defined by $gf(\gamma) = f(\gamma g)$ is also in X (so G acts on X on the left)*
- *If $f \in X$ then $f(ng) = \psi(n)f(g)$*
- *X is isomorphic to V as a G-module.*

This is called a Whittaker model of V. We have proved that for V irreducible, such a model exists iff V is not of the form $\chi\circ\det$, and we have proved uniqueness in all cases. This will be a key ingredient in "multiplicity one" arguments later. I think. Closely related to Whittaker models are Whittaker functionals:

Corollary 15. *If V is a representation of G not of the form $\chi\circ\det$ then there's a non-zero linear map $l : V \to C$ such that $l(nv) = \psi(n)l(v)$ and l is unique up to non-zero scalar. If V is of the form $\chi\circ\det$ then there are no such l.*

Remark. I think that such things might be very handy indeed when we come to study L-functions?

2 Smooth and admissible representations: basic definitions.

Here I follow Bernstein and Zelevinsky [1], and Cartier [3]. We will consider topological groups G with the following property:

($*$) G is Hausdorff, and any open neighbourhood of the identity contains a compact open subgroup.

This is called an l-group by Bernstein and Zelevinsky, and a "group of *td* type" by Cartier, and probably more things by other people. Let's go with Bernstein and say that G is an l-group if it satisfies ($*$). Note that an l-group is locally compact and furthermore if S is any subset of G with more than one point then S cannot be connected. So G is totally disconnected too. From now on, all our groups G will be l-groups. Note that any closed subgroup of an l-group is an l-group.

We can write down lots of examples of these things if we know something about local fields. A local field, for me, is the field of fractions K of a complete DVR A with finite residue field. One proves (see, for example, most of Chapter 2 of Serre's "Local Fields") that any such field is either isomorphic to the field of fractions of $A = k[[T]]$ where k is finite, or K is isomorphic to a finite extension of the p-adic numbers $\mathbf{Q}_p$; then A is the integral closure of $\mathbf{Z}_p$ in K.

If K is a local field, then it contains its integers A and A is a local ring so has a natural filtration $A \supset \mathfrak{m} \supset \mathfrak{m}^2 \supset \ldots$, where $\mathfrak{m}$ is the maximal ideal of A. Let G be the group $\mathrm{GL}_n(K)$. Then G is an l-group because G contains $\mathrm{GL}_n(A)$ which is compact, being the projective limit of the finite groups $\mathrm{GL}_n(A/\mathfrak{m}^r)$, and $\mathrm{GL}_n(A)$ contains the kernels of the maps $\mathrm{GL}_n(A) \to \mathrm{GL}_n(A/\mathfrak{m}^r)$ for all r and these are open compact subgroups and a neighbourhood basis of the identity. So there's an example, indeed it will be our key example. Other examples, that we will get onto later, are:

- The K-valued points of any affine algebraic group—as they give a closed subgroup of some $\mathrm{GL}_n(K)$.
- Products of finitely many such things.
- Finite groups with the discrete topology!
- Infinite "restricted" products (for example, $G(\mathbf{A}_f)$ where G is an affine algebraic group and $\mathbf{A}_f$ is the finite adeles (I will explain all this later)).

We will restrict to some explicit examples later, but for a while I'll mostly do generalities which work for all l-groups.

Now let G be any l-group, and let V be a vector space over C, possibly infinite-dimensional, and let $\pi : G \to \mathrm{Aut}_C(V)$ be a representation, that is, a group homomorphism.

Definition.

- *We say that π is* irreducible *if $V \neq 0$ and if there are no G-invariant subspaces other than $0, V$ (so this is just the same as before).*
- *We say that π is* smooth *(some people say* algebraic*) if for any $v \in V$ the stabiliser of v is open. Easy check: this is just the statement that the induced map $G \times V \to V$ is continuous if V has the discrete topology. If V is any representation then set $V_s := \{v \in V : \mathrm{stab}(v)$ is open$\}$; then V_s is a G-invariant subspace (easy check) called the* smooth part *of V.*
- *We say that π is* admissible *if it's smooth and also if, for any $U \subseteq G$ open subgroup, the space V^U is finite-dimensional.*

Note that if V is finite-dimensional and G acts in some ridiculous non-continuous way then V^U is finite-dimensional for all U, but V is not smooth. Note also that if V is infinite-dimensional and G acts trivially then V is smooth but not admissible. However, in many cases smooth and irreducible implies admissible. Local Langlands conjectures are about a canonical 1-1 correspondence between admissible representations of $\mathrm{GL}_n(K)$, K local, and certain n-dimensional representations of the Weil-Deligne group of K, a group closely related to the absolute Galois group of K. It's easy to write down representations of groups. But I have seen very few examples of admissible representations, and one of my many local motivations is to see some more, and also to associate L and ϵ factors to such things, or at least discover that this isn't easy. My understanding is that the Bernstein-Zelevinsky paper I'm following reduces the

Local Langlands conjecture to matching up supercuspidal representations with irreducible representations of the Weil-Deligne group (I'll define these later). Let's see some examples of these things.

Example. Say G is an l-group, and let $\chi : G \to C^\times$ be any character (that is, group homomorphism) with open kernel. Then χ is irreducible, smooth and admissible. Conversely if V is a 1-dimensional representation of an l-group and V is smooth then clearly it is of the above form. For example if $G = \mathrm{GL}_1(F)$ with F local, then the kernel could be something like $\mathcal{O}_F^\times$ or $1 + m^r$ for any $r \geq 1$. The quotient is then the product of $\mathbf{Z}$ and a finite group, with the discrete topology, and one takes any representation at all.

Coming up are some versions of Schur's lemma, which will prove that in fact the only irreducible smooth representations of an abelian l-group with certain finiteness conditions are of this form:

Definition. *We say that an l-group G is* countable at infinity *if G is a countable union of compact sets.*

Remark.

- Note that if G has this property then so does any closed subgroup, and if F is a local field then $\mathrm{GL}_n(F)$ has this property because if we consider it as a closed subset of F^{n^2+1} in the usual way then we can use $(m^{-r}O)^{n^2+1}$ intersect $\mathrm{GL}_n(F)$ for $r = 0, 1, 2, \ldots$. Hence the F-points of any affine algebraic group will be countable at infinity.
- Note that if G is countable at infinity then any compact open K will have at most countable index in G: it suffices to prove that if L is any compact subset of G then L is covered by finitely many translates of K, and this is immediate because G is a disjoint union of translates of K, all of which are open, and now just intersect with L.

Here are two versions of Schur's Lemma. Let us assume from now on that C is uncountable (this is OK, I want to include the case that $C = \overline{\mathbf{Q}}_p$ but this is uncountable).

Lemma 16. *(a) If G is an l-group which is countable at infinity, and if V is an irreducible smooth representation of G then the only G-invariant linear maps $V \to V$ are the scalars.*

(b) If G is any l-group and V is an irreducible admissible representation of G then the only G-invariant linear maps $V \to V$ are the scalars.

Proof. (a) Let v be a non-zero element of V. Let K be its stabilizer. Then V must be the space spanned by Gv and Gv is a countable set because G/K is; hence V has at most countable dimension. Now if there is $A : V \to V$ that isn't in C then $A - \lambda$ is invertible in $\mathrm{End}_G(V)$ for all $\lambda \in C$, and if R_λ is its inverse and $v_\lambda = R_\lambda v$ then it will be a contradiction if all the uncountably many v_λ are linearly independent. However they are indeed linearly independent: if $\sum c_i v_{\lambda_i} = 0$ and none of the (finitely many) c_i are zero then we can write this as $F(A) \prod_i R_i v = 0$ and $F(A)$ is $c \prod (A - \mu_i)$ and now we have a contradiction as everything is invertible so can't kill v.

(b) Choose $0 \neq v \in V$ and let K be its stabilizer. Then V^K is non-zero but finite-dimensional, and any $A : V \to V$ which is G-invariant sends V^K to itself, so it must have an eigenvalue λ on V^K. Now $A - \lambda$ is not injective and its kernel is G-invariant so it's zero. □

Corollary 17. *Every smooth irreducible representation of* $\mathrm{GL}_1(K)$ *with* K *local (or more generally any abelian* l*-group which is countable at infinity) is 1-dimensional and hence automatically admissible.*

Proof. $g \in G$ gives a non-zero G-invariant map so it must be multiplication by an element of $C^\times$ and hence all subspaces are invariant. □

This is a bit misleading—the moment we move away from the abelian case, the interesting representations will usually be infinite-dimensional. Note however that if G is finite with the discrete topology then there could be non-trivial finite-dimensional representations of G, so it's not a general theorem about l-groups that all finite-dimensional admissible representations are 1-dimensional.

Our goal is to write down lots of representations of $\mathrm{GL}_n(K)$ for K local (in fact I'd like to write down all of them in some cases but I'm not sure how hard this is). I cannot do this at present but I can do the easy cases because they're easy. I gave you two Schur's Lemmas; I'll now give you two ways of doing induction! Before we start I'll mention the following: If $f : X \to V$ is a function from a topological space to a vector space then we define the *support* of f to be the closure of $\{x \in X | f(x) \neq 0\}$.

> **Definition.** *Let* G *be an* l*-group and let* H *be a closed subgroup. Let* $\rho : H \to \mathrm{Aut}_C(V)$ *be a smooth representation of* H.

- *We define* $\mathrm{Ind}_H^G(\rho)$ *to be the space of functions* $f : G \to V$ *such that*

  (i) $f(hg) = \rho(h)f(g)$ *for all* $h \in H, g \in G$,

  (ii) There exists an open compact subgroup $N = N(f)$ *of* G *such that* $f(gn) = f(g)$ *for all* $g \in G, n \in N$.

  We define a G*-action by* $gf(\gamma) = f(\gamma g)$.

- *We define* $\mathrm{ind}_H^G(\rho)$ *to be the subspace of* $\mathrm{Ind}_H^G(\rho)$ *consisting of functions* f *such that the support of* f *is compact mod* H*, which means that there exists some compact subset* $K = K(f)$ *of* G *such that the support of* f *is contained in* $H.K := \{hk : h \in H, k \in K\}$*. The* G*-action is the induced one.*

Note that condition (ii) above guarantees that the induced representations are smooth. Note also that if the quotient topological space $H\backslash G$ is compact then the two definitions coincide, because for any $g \in G$ we can choose a compact open neighbourhood of g and the pushforward of these will be open, so finitely many will cover and so on. I think this notation is terrible, by the way. As far as I know, we will only be interested in the case where $H\backslash G$ is compact, so don't worry about the distinction and just use the easier of the definitions. Here's another useful fact:

Lemma 18. *If $H\backslash G$ is compact and $\rho : H \to \mathrm{GL}_C(V)$ is admissible then $W := \mathrm{Ind}_H^G(\rho) = \mathrm{ind}_H^G(\rho)$ is admissible.*

Proof. Let N be an open subgroup of G; we must show that W^N is finite-dimensional. Now $H\backslash G/N$ is finite because $H\backslash G$ is compact; let Ω denote a subset of G representing this double coset space; then $f \in W^N$ is determined by its values on Ω. Of course this doesn't do it yet because f takes values in V which is typically infinite-dimensional. But if $\omega \in \Omega$ and N' denotes the open compact subgroup $H \cap \omega N \omega^{-1}$ then I claim that $f(\omega) \in V^{N'}$ and of course this does do it because $V^{N'}$ is finite-dimensional by admissibility of V. To check $f(\omega) \in V^{N'}$ is a formality: if $f \in W^N$ and $n' \in N'$ then then $\rho(n')(f(\omega)) = f(n'\omega) = f(\omega n)$ for some $n \in N$ so this equals $f(\omega)$ because $f \in W^N$. □

Proposition 19 (Frobenius reciprocity)**.** *If H is a closed subgroup of the l-group G, if V is a smooth representation of G and W is a smooth representation of H then $\mathrm{Hom}_G(V, \mathrm{Ind}_H^G(W)) = \mathrm{Hom}_H(V|H, W)$ (canonically blah blah blah).*

Proof. (sketch) Here's the dictionary: if $a : V \to \mathrm{Ind}_H^G(W)$ then define a map $V|H \to W$ by $v \mapsto (a(v))(1)$. Conversely if $b : V \to W$ then define a map $V \to \mathrm{Ind}_H^G(W)$ by sending v to the function $g \mapsto b(gv)$. It's now two pages of trivial algebra to check that these maps are well-defined maps between the two spaces that are inverse to one another, so done. □

Now let's talk about Haar measure. If G is an l-group then it's locally compact (every point has a compact neighbourhood) so there will be a (left-invariant) Haar measure on G. In fact in this case it's very easy to write it down explicitly, it's much more combinatorial than something like Lebesgue measure. Here's what we're looking for: define a *Schwartz function* on G to be a function $f : G \to C$ which is locally constant and has compact support (the support of f is the closure of $\{g \in G | f(g) \neq 0\}$). Let $S(G)$ denote the space of Schwartz functions on G. Here is a helpful lemma that

Lemma 20. *If $f \in S(G)$ then there exists a compact open subgroup K, an integer $n \geq 0$, elements $g_1, g_2, \ldots, g_n \in G$ and elements $c_1, c_2, \ldots, c_n \in C$ such that the union $\cup_i K g_i K$ is disjoint and $f = \sum_i c_i \chi_{K g_i K}$, where χ_X denotes the characteristic function of the set X.*

Proof. Thanks to Richard Hill for pointing out the following proof, which is simpler than my original one. By definition, f is locally constant with compact support. Because f is locally constant it's a continuous map $G \to C$ if we give C the discrete topology. Hence $X := \{g \in G | f(g) \neq 0\}$ is closed and open, and hence this is the support of f. For every $x \in X$ we now construct an open subset U_x of X as follows: consider the map $\phi_x : G \times G \to C$ defined by $\phi_x(g_1, g_2) = f(g_1 x g_2)$. Then ϕ_x is continuous if C has the discrete topology, and so there is a compact open subgroup K_x of G such that ϕ_x is constant on $K_X \times K_x$. Define $U_x = K_x x K_x$. Then U_x is compact and open, and f is constant on U_x. Now X is covered by the U_x and hence by finitely many of them, and distinct double cosets are disjoint and we can let K be the intersection of the finitely many K_x that occur and we're home. □

A Haar measure on G is a non-zero linear map $\mu : S(G) \to C$ which is left-invariant, that is, if $f \in S(G)$ and $g \in G$ and we define a new function $h \in S(G)$ by $h(\gamma) = f(g\gamma)$, then $\mu(h) = \mu(f)$. If μ is a Haar measure and K is an open compact subset of G, we define $\mu(K)$, the measure of K, to be $\mu(\chi_K)$, where χ_K is the characteristic function of K.

Proposition 21. *A Haar measure exists, and is unique up to non-zero scalar.*

Proof. (sketch) Firstly let's construct a measure. Choose any compact open subgroup K_1 of G and let's decree that $\mu(K_1) = 1$. Now let L be any compact open subgroup of G. Then $M := K_1 \cap L$ is also compact and open, and both K_1 and L are finite unions of cosets of M, say $K_1 = \coprod_{i=1}^{k} g_i M$ and $L = \coprod_{j=1}^{l} h_i M$ and so we are forced to define $\mu(M) = 1/k$ and then $\mu(L) = l/k$. This works! Finally if X is any compact open set then one can write X as a union of subsets of the form xK where K is a compact open subgroup; finitely many of these will cover X and the intersection of all these finitely many $K's$ is a compact open K_0 such that X is a finite disjoint union of sets of the form xK_0. Say there are n sets in the union: then define $\mu(X) = n\mu(K_0)$.

Now if $f \in S(G)$ we can write f as $\sum_{i=1}^{n} a_i \chi_{X_i}$ with X_i disjoint compact open sets; we of course define $\mu(f) = \sum_{i=1}^{n} a_i \mu(X_i)$. This works; conversely any measure μ' which is left invariant and such that $\mu'(K_1) = 1$ must be equal to μ. ∎

Notation:

$$\mu(f) = \int_G f(x)\mathbf{d}\mu(x).$$

Note that in this 1-dimensional space of left-invariant measures there is a canonical $\mathbf{Q}$-structure: a measure μ is *rational* if $\mu(K)$ is a rational number for one, and hence all, compact open subgroups of G.

It's certainly not true in general that Haar measure μ is also right-invariant; if $g \in G$ then the map $S(G) \to C$ defined by sending f to $\int_G f(xg)\mathbf{d}\mu(x)$ is however a constant multiple of Haar measure, so it's $\Delta(g)\mu$ for some $\Delta(g) \in C^\times$; one can do better in fact: $\Delta(g) = \mu(K_1 g)/\mu(K_1)$ if K_1 is open and compact, and hence $\Delta(g)$ is a positive rational. One easily checks that $\Delta(g) = 1$ if $g \in K_1$ (or indeed if g is contained in any compact subgroup of G) and that $\Delta : G \to C^\times$ is an admissible representation. Note that Δ is identically 1 if G is compact or abelian. It's also trivial on $\mathrm{GL}_n(K)$, where K is local, or more generally if G is the K-points of some reductive group over K, although I don't know why (Richard Hill tells me it's not hard, and suggested the following proof sketch: clearly the centre of the group is in the kernel of Δ by definition, as is the derived subgroup of G, but these should generate a subgroup of finite index (certainly this is true on the level of algebraic groups) and the positive rationals have no non-trivial finite subgroups.

One example where Δ isn't identically 1 is the case of a parabolic subgroup of $\mathrm{GL}_n(F)$, F local: I will compute Δ in the case of the upper triangular matrices in $\mathrm{GL}_2(F)$ later on (Lemma 32) when I need it.

2.1 Duality.

Lemma 22. *If K is a compact l-group then all irreducible smooth representations of K are finite-dimensional, and the action of K factors through a finite quotient.*

Proof. Let V be such a representation, and choose $0 \neq v \in V$. Then the stabiliser of v is some compact open subgroup L of K and compactness of K implies that L has finite index, and hence only finitely many conjugates. Let M be the intersection of these conjugates. Then M is compact and open, and normal in K, and $v \in V^M$ which is G-invariant and hence by irreducibility $V = V^M$. Finally G acts via the finite group G/M and an irreducible representation of a finite group is finite-dimensional. $\square$

Remark. Hence compact groups also have the property that irreducible and smooth implies admissible.

Now let G be any l-group and choose a compact open K in G. Let V be any representation of G (not necessarily smooth or admissible). Let θ denote any irreducible admissible representation of K and define $V(\theta)$ to be the C-subspace of V generated by the images of all the elements of $\mathrm{Hom}_K(\theta, V)$. That is, $V(\theta)$ is the maximal subspace where K acts as θ. Now let θ run over all the irreducible admissible representations of K. One sees easily that the sum of the $V(\theta)$ is direct, because if a finite linear combination of v_θ in distinct $V(\theta)$ were zero then consider the intersection of the kernels of the differing θ; quotient K out by this compact open subgroup and now we get a finite group G and $C[G]$ is semi-simple. Hence $\oplus_\theta V(\theta)$ is a subspace of V.

Lemma 23. *(i) $V_s = \oplus_\theta V(\theta)$. In particular V is smooth iff $V = \oplus_\theta V(\theta)$.*
(ii) V is admissible iff $V = \oplus_\theta V(\theta)$ and all the $V(\theta)$ are finite-dimensional.

Proof. (i) if $v \in V_s$ then it's fixed by some L which WLOG is a compact open in K, normal in K, and V^L is a representation of the finite group K/L which hence decomposes as a sum of the parts corresponding to the finitely many irreducible representations of K/L. The converse is immediate.

(ii) One just has to note that if L is a compact open in K, normal in K, then V^L is just the direct sum of the $V(\theta)$ as θ ranges through the representations with L in the kernel. $\square$

Remark. Although I haven't said anything about representations of real and complex groups (which aren't in general l-groups), this previous lemma is a very very big clue as to the correct definition of admissibility in the archimedean case. See Wallach's article in the Corvallis proceedings [4].

If V is a smooth representation of an l-group G then we define the dual $\tilde{V}$ to be the smooth vectors in the algebraic dual of V. This is also a smooth representation of V.

Lemma 24. *Let V be a smooth representation of an l-group G.*
(i) If V is admissible then so is $\tilde{V}$

(ii) The canonical map $V \to \tilde{\tilde{V}}$ is always an injection, and is an isomorphism iff V is admissible.

(iii) V is admissible iff $\tilde{V}$ is admissible.

Proof. (i) We know $V = \oplus_\theta V(\theta)$ and hence $V^* = \prod_\theta V(\theta)^*$. One now checks easily that the smooth vectors in this space are $\oplus_\theta V(\theta)^*$, because any open subgroup of G contains a compact open subgroup L of K which is normal in K, and $(V^*)^L$ is the direct sum of the finitely many $V^(\theta)^*$ which factor through L. Now it's clear that the θ-eigenspace of $\oplus_\theta V(\theta)^*$ is just $V(\theta^*)^*$ where θ^* is the usual dual representation in the theory of finite groups. Everything now follows from the previous lemma.

(ii) We see that $\tilde{\tilde{V}} = \oplus_\theta V(\theta)^{**}$ and the double dual of a vector space is itself iff the vector space is finite-dimensional.

(iii) Immediate from (i) and (ii). $\square$

One reason I mentioned everything I've mentioned today is the relationship between duals and induction. It has a twist that you might not expect: a twist by Δ.

Lemma 25. *Let H be a closed subgroup of the l-group G and let W be a smooth representation of H. Then the dual of $\mathrm{ind}_H^G(W)$ is isomorphic to $\mathrm{Ind}_H^G(\tilde{W} \otimes_C \chi)$ where $\chi : H \to C^\times$ is the map $\chi(h) = \Delta_G(h)/\Delta_H(h)$.*

Proof. Elementary once one knows that this is what you're trying to prove. A bit messy though; I'll omit it. It's Bernstein-Zelevinsky 2.25(c). $\square$

Remark. One conclusion that we can draw from this is that it might have been a better idea to define the induction from H to G of W to be the maps $f : G \to W$ such that $f(hg) = (\Delta_G(h)/\Delta_H(h))^{1/2} hf(g)$ and that f is constant on left cosets of some $N = N(f)$. Let's call this *normalised induction* and write it as NInd (and Nind for the corresponding compact support induction). I am quite pleased to discover this—this is a definition that I had read before for GL_2 and I never had any idea about why the fudge factor was there. If we do this then I *think* this messes up Frobenius reciprocity. But it makes duals much better: with notation as above, we see

$$\mathrm{Nind}_H^G(W) = \mathrm{ind}_H^G(W \otimes \chi^{1/2})$$

and similarly for NInd, and then that

$$\begin{aligned}
\widetilde{\mathrm{Nind}_H^G}(W) &= \mathrm{ind}_H^G\widetilde{(W \otimes} \chi^{1/2}) \\
&= \mathrm{Ind}_H^G(\tilde{W} \otimes \chi^{-1/2} \otimes \chi) \\
&= \mathrm{Ind}_H^G(\tilde{W} \otimes \chi^{1/2}) \\
&= \mathrm{NInd}_H^G(\tilde{W}).
\end{aligned}$$

We will also see later, once I have defined unitary representations, that this "normalised" induction takes a unitary representation to a unitary representation, so that's another reason why it's a good idea.

The last general thing I want to talk about is Hecke algebras. Let G be any l-group and K any compact open subgroup. Fix a positive left-invariant Haar measure on G. Define $H(G,K)$ to be the subset of $S(G)$ consisting of functions $f : G \to C$ which have compact support and are bi-K-invariant, that is $f(k_1gk_2) = f(g)$ for $g \in G$ and $k_1, k_2 \in K$. By lemma 20 $S(G)$ is the union of the $H(G,K)$ as K runs through all compact opens. Define a multiplication on $H(G,K)$ (depending on the choice of Haar measure) by $f_1 * f_2(g) = \int_G f_1(x)f_2(x^{-1}g)\mathbf{d}\mu(x)$. This integral makes sense because for fixed g the integrand, as a function of x, is compactly supported and locally constant. Furthermore the resulting function is in $H(G,K)$ again. This gives $H(G,K)$ an associative multiplication, and it has an obvious addition. It also has an identity: one checks easily that $e_K := 1/\mu(K)\chi_K$ works. As a C-vector space $H(G,K)$ has a basis consisting of functions which are 1 on one double coset KxK and zero everywhere else.

If $L \subset K$ is a strictly smaller compact open then there's a natural inclusion $H(G,K) \to H(G,L)$ but unfortunately it doesn't send the identity to the identity. The union over K of the $H(G,K)$ is $S(G)$, which we write $H(G)$ if we're considering it as an associative algebra, and this is a ring which is commutative iff if G is, and has a 1 iff G is discrete. This ring $H(G)$ is the analogue of $C[G]$ in the finite-dimensional case, for smooth representations, as we shall now see. One thing I should mention before we go on is that one can extract $H(G,K)$ from $H(G)$ thus: $H(G,K) = e_K H(G) e_K$ (easy check).

A *representation* of $H(G)$ is, of course, a C-vector space V and a map $H(G) \to \mathrm{End}_C(V)$ which is C-linear and commutes with multiplication. So it's just a module for this ring without a one. For rings without a one there's a notion of "non-degeneracy" which one doesn't see for rings with ones. A module for $H(G)$ is *non-degenerate* if for all $v \in V$ there is $h_i \in H$ and $v_i \in V$, $1 \leq i \leq n$, such that $v = \sum_i h_i v_i$.

Let V be a smooth representation of G, and write $\tilde{V}$ be the smooth dual of V. Choose $v \in V$ and $f \in \tilde{V}$. The *matrix coefficient* $\pi_{v,f}$ associated to v and f is the canonical map $G \to C$ defined by $\pi_{v,f}(g) = f(gv)$.

Now fix a Haar measure on G.

Lemma 26. *Let V be a smooth representation of G, and say $h \in H(G)$. Then there's a unique C-linear map $h : V \to V$ with the property that for all $v \in V$ and $f \in \tilde{V}$ we have*

$$f(hv) = \int_G h(g)\pi_{v,f}(g)\mathbf{d}\mu(g).$$

This makes V into a non-degenerate $H(G)$-module. Furthermore this construction establishes an equivalence of categories between the category of smooth G-modules and the category of non-degenerate $H(G)$-modules.

Proof. We could just define

$$hv = \int_G h(g)gv\mathbf{d}\mu(g)$$

if people are happy with this: then everything will work. The problem is that the integrand isn't a C-valued function, it's a V-valued one, so let's just make clear what we mean here: for $v \in V$ choose K compact open and so small that K stabilises v and that h is constant on cosets of the form gK; then $h = \sum_{i=1}^{n} c_i \chi_{g_i K}$ and we define the integral to be $\mu(K) \sum_{i=1}^{n} c_i g_i v$. This is independent of the choice of K and everything works; the check that the action commutes with multiplication is an elementary exercise. Note also that if $v \in V$ and K is in the stabiliser of v then $e_K v = v$ (recall e_K was the unit in $H(G,K)$) and hence V is non-degenerate. In fact one checks easily that the action above makes V^K into an $H(G,K)$-module and $H(G,K)$ has a one so we're back to sanity if you think in this way.

One checks easily that sub-G-modules and sub-$H(G)$-modules coincide. To get the equivalence we have to construct a smooth G-module from a non-degenerate $H(G)$-module (and then do lots of elementary checks); if we have a non-degenerate $H(G)$-module then for $v \in V$ and $g \in G$ we must define gv. By non-degeneracy, $v = \sum h_i v_i$ and each h_i is in $H(G, K_i)$ for some compact open K_i so they're all in $H(G,K)$ for some compact open K; then $e_K v = v$ and we can define $gv = \frac{1}{\mu(K)} \chi_{gK} v$. One checks that this is well-defined if one shrinks K and then all the other things follow. Note that the reason one gets a smooth representation is that K acts trivially on v. **I should check that these functors are quasi-inverse of one another because something seems a bit fishy to me** □

We say that a smooth representation of G is *finitely-generated* if the associated $H(G)$-module is finitely-generated. Here's a nice way of detecting irreducibility:

Lemma 27. *Let V be a non-degenerate representation of $H(G)$. Then V is irreducible iff $V \neq 0$ and for each compact open subgroup K of G, either $V^K = 0$ or V^K is an irreducible $H(G,K)$-module.*

Proof. If V is a non-degenerate representation of $H(G)$ and there's K such that V^K contains an $H(G,K)$-invariant subspace U not equal to 0 or V^K then let W be the G-submodule of V generated by U; then $W^K = e_K W$ certainly contains U and is in fact equal to U because it's spanned by elements of the form $e_K h u$ with $h \in H(G)$ and $u \in U$, and $e_K h u = e_K h e_K u \in U$ as $e_K h e_K \in H(G,K)$. Hence W must be a non-trivial subspace of V. Conversely, if W is a non-trivial subspace of V then choose $0 \neq w \in W$ and $v \in V$ with $v \notin W$ and then for small enough K we'll have $v, w \in V^K$ so we have proper inclusions $0 \subset W^K \subset V^K$. □

How does one extract the admissible representations in this context?

Definition. *A representation of $H(G)$ is* admissible *if it's non-degenerate and every element in $H(G)$ has finite-dimensional image.*

Lemma 28. *If V is a smooth representation of G then it's admissible in the old sense iff the corresponding representation of $H(G)$ is admissible in this sense.*

Proof. If e_K has finite rank then $e_K V = V^K$ is finite-dimensional, so one way is very easy; the other is easy too, if all the V^K are finite-dimensional then for $h \in H(G)$ we have $h \in H(G, K)$ for some K and then $hV = he_K V = h(V^K)$ must be finite-dimensional. □

One thing that comes from this is that admissible representations have traces! Let V be an admissible representation; then define a linear map $H(G) \to C$ by sending h to the trace of h acting on V. This map, a linear functional on $H(G)$ and hence sometimes called a *distribution* on G, is called the character of the admissible representation V. The standard fact from representation theory of finite groups, that traces of non-isomorphic representations are linearly independent, also is true in this setting, but I won't prove it because I don't need it; one can even prove some kind of orthogonality relations—see the first few Theorems of Cartier's Corvallis article [3].

Finally, here are two important notions, one perhaps rather unfashionable nowadays, and one absolutely essential. People once used to consider continuous, or perhaps unitary, representations of groups like $\mathrm{GL}_n(\mathbf{Q}_p)$ on complex Hilbert spaces; given such a thing one can take the smooth vectors and get a smooth representation, but I don't think that all smooth representations arise in this way. Furthermore two non-isomorphic topological representations can, I think, give rise to isomorphic smooth representations (Jacquet and Langlands don't say this but they sometimes talk about an even larger class of representations, on separable complete locally convex spaces, and here this phenomenon can occur, they say (on p30)). If V is an admissible representation and $C = \mathbf{C}$, one says that V is *unitary* if there exists a positive definite Hermitian sesquilinear form $(,)$ on V such that $(gv, gw) = (v, w)$ for all $v, w \in V$ and $g \in G$. This implies that $\tilde{V}$ is isomorphic to the conjugate of V. The reason I found it hard to extract precise statements of the form "if G is (blah) then smooth and irreducible implies admissible" is that sometimes people prove it only for unitary representations. One passes from irreducible unitary Hilbert space representations to smooth representations of the kind we've been considering by taking smooth vectors (surprise surprise) and then one has to work hard to check that the result is admissible (this was an open question for a while, but Bernstein and Zelevinsky state in section 4.21 of [1] that it's true for $\mathrm{GL}_n(K)$, K local, and I think it's true for reductive groups over local fields). Conversely of course, and one goes from unitary smooth representations to Hilbert space representations by completing with respect to the norm. This gives a bijection between the topologically irreducible unitary Hilbert space representations of G and the admissible irreducible unitary representations of G. One other reason that one should perhaps mention unitary representations in this algebraic setting is that if $G = \mathrm{GL}_n(K)$ and H is a parabolic subgroup of G then it's a theorem that if W is a unitary representation of H then $NInd_H^G(W)$ is a unitary representation of G **reference?**.

The other notion I wanted to mention was the definition of supercuspidal representations. Let V be a smooth representation of G. Then V is called *supercuspidal* if each matrix coefficient $\pi_{v,f} : G \to C$ is compact mod centre, that is, if for all v, f as above there exists a compact subset $X = X(v, f)$ of G such that $\pi_{v,f}$ vanishes outside the subset XZ of G, where where $Z = Z(G)$ is the centre of G. Note that V is admissible so by Schur's lemma the centre of G acts via a character, and so certainly the support of G is a union of cosets of Z; this is why we don't just say "compact", because Z might not be compact. I guess that when people just thought about representations of semisimple groups they didn't have this problem. Note for example that if G is compact (or even if G/Z is compact, for example if G is abelian) then all representations are supercuspidal.

Here's one nice thing about supercuspidal representations. Assume that C is uncountable and G is countable at infinity. Then Schur's lemma works, and using it we can see

Lemma 29. *If V is an irreducible smooth supercuspidal representation of G then V is admissible.*

Remark. Hence we re-prove that smooth and irreducible implies admissible for compact or abelian groups.

Proof. Suppose V is irreducible and smooth and not admissible. We produce a matrix coefficient that is not compact mod centre. Firstly choose K compact and open such that V^K is infinite-dimensional. Then choose $0 \neq v \in V^K$. Then V is irreducible so is spanned by $\{gv : g \in G\}$, and $V^K = e_K V$ is hence spanned by $\{e_K gv : g \in G\}$. Choose a countably infinite subset $\{g_1, g_2, \ldots\}$ of G such that the $e_K g_i v$ are linearly independent (note that this implies that the $g_i v$ are all linearly independent too). If 1 denotes the trivial representation of K then $V^K = V(1)$ and we showed when discussing duals that $\tilde{V}^K = V^*(1)$, so there's an element $f \in \tilde{V}^K \subseteq \tilde{V}$ such that $f(e_K g_i v) = i$ for all i. The corresponding matrix coefficient certainly contains all the g_i in its support, as $\pi_{v,f}(g_i) = f(g_i v)$ and $e_K f = f$ so this is $f(e_K g_i v) = i$. Finally we have to show that the g_i don't all lie in a subset that is compact mod centre, and this is because if $X \subseteq G$ is compact and all the g_i were in XZ, then cover X by finitely many translates gK of K and then note that the C-span of the set XZv is (by Schur) the same as the C-span of Xv, which is finite-dimensional, yet the space spanned by the $g_i v$ is infinite-dimensional, contradiction. □

So when we want to prove that all irreducible smooth representations of some group are admissible, we only have to deal with the non-supercuspidal ones. This is of great use when we restrict to groups like $\mathrm{GL}_n(F)$ later, F local.

3 Representations of $\mathrm{GL}_2(F)$, F local: Jacquet-Langlands' theory.

Having seen these generalities, we now turn to the special case of $G = \mathrm{GL}_2(F)$. We extensively rely on the non-archimedean local part of Godement's notes [5], who is in turn summarising the book by Jacquet and Langlands [6], but we also talk about §1 of [6], which Godement omits.

For this section let F be a non-archimedean local field and let G denote the group $\mathrm{GL}_2(F)$. Fix a Haar measure μ on G, and for convenience let's choose the one such that $\mu(\mathrm{GL}_2(\mathcal{O}_F)) = 1$. We must also fix a non-trivial continuous map $U \to C^\times$ where U is the matrices in G of the form $\left\{\left(\begin{smallmatrix} 1 & x \\ 0 & 1 \end{smallmatrix}\right)\right\}$. In other words, we have to fix $\tau : F \to C^\times$ a non-trivial locally constant group homomorphism. Here's one example for $K = \mathbf{Q}_p$ and $C = \mathbf{C}$: if $t \in \mathbf{Q}_p$ we can write $t = r + s$ where $r \in \mathbf{Z}_p$ and $s = a/p^n \in \mathbf{Q}, 0 \leq s < 1$. Send t to $e^{2\pi i s}$. This is well-defined. Now here's an example for F a finite extension of $\mathbf{Q}_p$: take the trace down from F to $\mathbf{Q}_p$ and then use the map before. Actually, this might not be such a good choice: a locally constant $\tau : F \to C^\times$ is called *unramified* if $\mathcal{O}_F$ is the largest fractional ideal of F on which τ is trivial; but if τ is as above one can check that there's $f \in F$ such that the map $x \mapsto \tau(fx)$ is unramified.

Here's an example of an unramified additive character if $F = k((t))$: if $f \in F$ then take the residue (that is, the coefficient of t^{-1}; this is in k; now take the trace down to $\mathbf{F}_p$; choose a pth root of unity ζ in C and send 1 to ζ.

We also need to fix once and for all a Haar measure on F; we can for example take the one such that the measure of $\mathcal{O}_F$ is 1. This works out rather nicely: if τ is unramified and we use this choice of μ then we have the following compatibility:

Lemma 30. *If τ is unramified and $\mu(\mathcal{O}_F) = 1$ then for all $g \in S(F)$, if we define*

$$\hat{g}(y) = \int_F g(x)\tau(xy)\mathbf{d}\mu(x)$$

then

$$g(z) = \int_F \hat{g}(y)\overline{\tau(yz)}\mathbf{d}\mu(y).$$

Proof. Note that $\overline{\tau(x)} = \tau(-x)$ firstly; next note that everything is finitely additive and so it suffices to prove the lemma for g the characteristic function of $(\pi)^n$ for all $n \in \mathbf{Z}$ in which case it's an explicit computation. □

The big question: what are all the irreducible admissible representations of G? We have to do two things to work this out: firstly we have to find some methods of constructing them, and secondly we have to prove that we have constructed them all. Our construction will exactly mimic the case of $\mathrm{GL}_2(k)$ for k finite; we have one-dimensional representations coming via the determinant map, we have representations induced from the subgroup of upper triangular matrices, and we have supercuspidal ones as constructed by Weil by a more elaborate construction involving the explicit presentation of $\mathrm{SL}_2(F)$.

Here we go then. Let F be local. If $\chi : F^\times \to C^\times$ is any continuous (that is, locally constant) character then $\chi \circ \det$ is a map $\mathrm{GL}_2(F) \to C^\times$ and this gives us a 1-dimensional smooth representation. On the other hand, it's easy to check that these are all the finite-dimensional ones.

Lemma 31. *If V is a finite-dimensional smooth irreducible representation of G then V is of the above form.*

Proof. The kernel of the representation is the intersection of the stabilisers of a basis, so it's open and compact, and normal in G. I claim that any such subgroup contains $\mathrm{SL}_2(F)$ and of course to check this all I have to do is to check that it contains all upper triangular unipotent matrices, all diagonal matrices with determinant one, and w. But any open and compact subgroup of G will contain an open neighbourhood of the identity, and hence a matrix of the form $\left(\begin{smallmatrix}1 & t\\ 0 & 1\end{smallmatrix}\right)$ for some $t \neq 0$. If this subgroup is normal in G then it must hence contain $\left(\begin{smallmatrix}1 & \lambda\\ 0 & 1\end{smallmatrix}\right)$ for any $\lambda \in K$ (conjugate by a diagonal matrix). Similarly all lower triangular unipotent matrices are in too. So $\left(\begin{smallmatrix}1 & 1\\ 0 & 1\end{smallmatrix}\right)\left(\begin{smallmatrix}1 & 0\\ -1 & 1\end{smallmatrix}\right)\left(\begin{smallmatrix}1 & 1\\ 0 & 1\end{smallmatrix}\right)$ is in and this is w; its square is -1 so this is in too; finally I need to get $\left(\begin{smallmatrix}\lambda & 0\\ 0 & \lambda^{-1}\end{smallmatrix}\right)$ in for $\lambda \in F^\times$ and WLOG $\lambda \neq \lambda^{-1}$; now note that $\left(\begin{smallmatrix}1 & x\\ 0 & 1\end{smallmatrix}\right)\left(\begin{smallmatrix}1 & 0\\ 1 & 1\end{smallmatrix}\right)$ is in but if $x = \lambda + \lambda^{-1} - 2$ then this matrix has eigenvalues λ and λ^{-1} so must be conjugate to what we want. Now by Schur's Lemma (which applies because the representation is admissible) G acts via scalars and hence V is one-dimensional and the induced character χ is locally constant (easy check). $\square$

It would be nice to use normalised induction when inducing up from upper triangular matrices, so let's work out the character Δ for the upper triangular matrices P in $\mathrm{GL}_2(F)$. Let's fix a left-invariant Haar measure on P such that the measure of the compact open K consisting of $\left(\begin{smallmatrix}a & b\\ 0 & d\end{smallmatrix}\right)$ with $a, b \in \mathcal{O}_F^\times$ and $b \in \mathcal{O}_F$ has measure 1. By the same argument as above, Δ will contain the upper triangular unipotent matrices, so it suffices to evaluate Δ on $\left(\begin{smallmatrix}\pi & 0\\ 0 & 1\end{smallmatrix}\right)$ and $\left(\begin{smallmatrix}1 & 0\\ 0 & \pi\end{smallmatrix}\right)$. The next thing to observe is that if we choose elements $f_1, \ldots, f_q$ in $\mathcal{O}_F$ lifting the elements of $\mathcal{O}_F/(\pi)$ (letting q be the size of the residue field) and let g_i denote the matrix $\left(\begin{smallmatrix}1 & f_i/\pi\\ 0 & 1\end{smallmatrix}\right)$ then the union $g_i K$ is disjoint and equals the matrices $\left(\begin{smallmatrix}a & b\\ 0 & d\end{smallmatrix}\right)$ with $a, b \in \mathcal{O}_F^\times$ and $b \in \pi^{-1}\mathcal{O}_F$. So this set must have measure q. But it's $\left(\begin{smallmatrix}\pi^{-1} & 0\\ 0 & 1\end{smallmatrix}\right)K\left(\begin{smallmatrix}\pi & 0\\ 0 & 1\end{smallmatrix}\right)$ so $\Delta(\left(\begin{smallmatrix}\pi & 0\\ 0 & 1\end{smallmatrix}\right)) = q$. Because Δ is trivial on the centre we have proved

Lemma 32. $\Delta(\left(\begin{smallmatrix}a & b\\ 0 & d\end{smallmatrix}\right)) = |d|/|a|$ *where* $|.| : F^\times \to C^\times$ *is trivial on units and sends* π *to* $1/q$, *where* q *is the size of the residue field.*

In particular, P is not unimodular. Now let χ_1, χ_2 be two continuous characters $F^\times \to C^\times$; they induce a character of P sending $\left(\begin{smallmatrix}a & b\\ 0 & d\end{smallmatrix}\right)$ to $\chi_1(a)\chi_2(d)$. This is clearly an admissible representation of P so representation $B(\chi_1, \chi_2)$, the normalised induction of this P-representation to G, that is the functions $G \to C$ such that $f(\left(\begin{smallmatrix}a & b\\ 0 & d\end{smallmatrix}\right)g) = \chi_1(a)\chi_2(d)(|a|/|d|)^{1/2}f(g)$ and such that there exists open compact N with $f(gn) = f(g)$, is an admissible representation of G. It's unfortunately not always irreducible. For example, let χ be a continuous character of $F^\times$ and consider the function $f = \chi \circ \det$. This will be in $B(\chi_1, \chi_2)$ iff $\chi(ad) = \chi_1(a)\chi_2(d)(|a|/|d|)^{1/2}$, that is $\chi_1|.|^{1/2} = \chi = \chi_2|.|^{-1/2}$, which can happen iff $\chi_1|.| = \chi_2$. But $B(\chi_1, \chi_2)$ usually is irreducible. Unfortunately this lies a bit deeper in the theory; we can't mimic the methods we used in the finite field case because of problems with non-semisimplicity. To work with these representations we need to know about Whittaker and Kirillov models for representations of $\mathrm{GL}_2(F)$: these are ways of getting our hands on representations of $\mathrm{GL}_2(F)$ (and presumably of other groups too, but I haven't got that far).

3.1 Kirillov models.

Let V be a representation of $\mathrm{GL}_2(F)$. A *Kirillov model* of V is a C-vector subspace $\mathcal{K}$ of the space of C-valued functions on $F^\times$, and an action of $\mathrm{GL}_2(F)$ on $\mathcal{K}$ with the property that $\left(\left(\begin{smallmatrix} a & b \\ 0 & 1 \end{smallmatrix}\right)k\right)(x) = \tau(bx)k(ax)$ for all $a, x \in F^\times$, $b \in F$, and $k \in K$, such that the representations V and $\mathcal{K}$ are isomorphic. The main result, which will be of great use to us, is:

Theorem 33. *If V is an irreducible admissible infinite-dimensional representation of $\mathrm{GL}_2(F)$ then V has a unique Kirillov model.*

Remark. Here uniqueness is meant in a very strong sense: if there are two subspaces $\mathcal{K}_1$ and $\mathcal{K}_2$ and two G-actions both satisfying everything above then $\mathcal{K}_1 = \mathcal{K}_2$ (not just an isomorphism, an equality) and the G-actions are also exactly the same, not just isomorphic.

Proof. (sketch of the idea). If V has a Kirillov model $\mathcal{K}$ then the subspace $\mathcal{K}_0$ of K consisting of k such that $k(1) = 0$ has codimension at most 1, and in fact has codimension exactly 1 because if $k(1) = 0$ for all $k \in \mathcal{K}$ then $k(a) = \left(\left(\begin{smallmatrix} a & 0 \\ 0 & 1 \end{smallmatrix}\right)k\right)(1) = 0$ for all $a \in F^\times$ so $\mathcal{K} = 0$, contradiction. Now K_0 corresponds to a subspace V_0 of V and this subspace, assuming the theorem, has an intrinsic definition: one checks easily that $v \in V_0$ iff

$$\int_{(\pi)^{-n}} \overline{\tau(x)}\left(\begin{smallmatrix} 1 & x \\ 0 & 1 \end{smallmatrix}\right)v\mathbf{d}x = 0$$

for all sufficiently large n. So let's now forget that V has a Kirillov model, let's start with an arbitrary infinite-dimensional irreducible admissible V and define V_0 to be the subspace of v for which the integrals above vanish for n sufficiently large. The main technical result, which is elementary but long, and whose proof occupies seven pages of Godement's notes (sections I.2 and I.3), is that V_0 has codimension 1 in V. The strategy is, to my mind, rather nice: one writes down lots and lots of linear endomorphisms of V/V_0 and then proves firstly that any endomorphism of V/V_0 which commutes with all of these endomorphisms must be a scalar, and then one proves that all the endomorphisms we've written down commute with one another! This is enough. Now one just chooses a C-isomorphism α of V/V_0 with C and for $v \in V$ one defines the map $k_v : F^\times \to C$ by $k_v(x) = \alpha\left(\left(\begin{smallmatrix} x & 0 \\ 0 & 1 \end{smallmatrix}\right)v\right)$; this does existence, and uniqueness follows very quickly from the fact that V/V_0 is 1-dimensional. $\square$

In fact, one can say much more about $\mathcal{K}$:

Theorem 34. *Let V be an infinite-dimensional irreducible admissible representation. Let $\mathcal{K}$ be its Kirillov model. Then every $k \in \mathcal{K}$ is a locally constant* *function on $F^\times$ and furthermore for every $k \in \mathcal{K}$ there's a compact subset C_k of F (*note: not of $F^\times$*) such that k vanishes outside C_k. Moreover, the space $S(F^\times)$ of locally constant functions on $F^\times$ with compact support is a subset of $\mathcal{K}$, and the quotient is finite-dimensional.*

The proof is another four pages in Godement's notes (all of section I.4). Again it's entirely self-contained but rather long. The idea is to check firstly that one non-zero element of $S(F^\times)$ must be in $\mathcal{K}$ and then to check that the upper triangular matrices acting on one element will generate all of $S(F^\times)$. That the codimension is finite is a messy computation that I don't really understand (and haven't read properly).

3.2 Whittaker models.

Let V be a representation of $G = \mathrm{GL}_2(F)$. A *Whittaker model* for V is a sub-C-vector space $\mathcal{W}$ of the space of all locally constant C-valued functions on G satisfying $w\left(\left(\begin{smallmatrix}1 & x\\ 0 & 1\end{smallmatrix}\right)g\right) = \tau(x)w(g)$ for all $g \in G, x \in F$, and with the G-action defined by $(gw)(\gamma) = w(\gamma g)$.

Theorem 35. *Let V be an irreducible admissible infinite-dimensional representation of G. Then V has a unique Whittaker model.*

Proof. This is easy from existence and uniqueness of the Kirillov model: given the Kirillov model $\mathcal{K}$ of V define $\mathcal{W}$ thus: for $k \in \mathcal{K}$ define w by $w(g) = (gk)(1)$. Conversely given a Whittaker model $\mathcal{W}$ define $\mathcal{K}$ by, for $w \in \mathcal{W}$, setting $k(x) = w\left(\left(\begin{smallmatrix}x & 0\\ 0 & 1\end{smallmatrix}\right)\right)$. $\square$

A *Whittaker functional* on a G-representation V is a non-zero linear map $L : V \to C$ such that $L\left(\left(\begin{smallmatrix}1 & x\\ 0 & 1\end{smallmatrix}\right)v\right) = \tau(x)L(v)$ for all $v \in V, x \in F$. One checks that to give a Whittaker functional (up to scaling) is to give a Whittaker model and vice-versa: again all we have to supply is the dictionary, which is this: given $\mathcal{W}$ define L by $L(w) = w(1)$ and given L define $\mathcal{W}$ by, for $v \in V$, defining $w(g) = L(gv)$.

What is the use of these things? Well, here is one use.

Lemma 36. *If V is an irreducible admissible representation of G then $\tilde{V}$ is isomorphic to $V \otimes \omega^{-1}$, where ω is the central character of V.*

Proof. This is trivial if V is 1-dimensional, so WLOG V is infinite-dimensional and hence has a Kirillov model $\mathcal{K}$. Let $\tilde{\mathcal{K}}$ denote the set of functions of the form $x \mapsto \omega\left(\left(\begin{smallmatrix}x & 0\\ 0 & x\end{smallmatrix}\right)\right)^{-1} k(x)$ as k runs through $\mathcal{K}$. Then $\tilde{K}$ is the Kirillov model for the twist of V. One can now construct without too much difficulty a pairing $\mathcal{K} \times \tilde{K} \to C$ proving that $\tilde{K}$ is also a Kirillov model for $\tilde{V}$, and we're done. See §I.6 of Godement for more details. $\square$

Here's another application. Let V be an infinite-dimensional irreducible admissible representation of G. Let $\mathcal{K}$ denote its Kirillov model. We have seen already that $S(F^\times)$ is a subspace of $\mathcal{K}$ with finite codimension. We have

Theorem 37. *V is supercuspidal iff $S(F^\times) = \mathcal{K}$.*

Proof. (sketch) It's an easy check that if $k \in \mathcal{K}$ then $k \in S(F^\times)$ iff k vanishes in a (punctured) neighbourhood of 0 iff the function

$$\int_{(\pi)^{-n}} \left(\begin{smallmatrix} 1 & x \\ 0 & 1 \end{smallmatrix}\right) k \mathbf{d}x = 0$$

for all n sufficiently large (the integral denoting a function $F^\times \to C$, because it's a finite sum of elements of K).

Firstly let's assume that $S(F^\times) = \mathcal{K}$ and prove that $\mathcal{K}$ is supercuspidal. We know the Kirillov model for $\mathcal{K}$ and $\tilde{\mathcal{K}}$, and because $\tilde{\mathcal{K}}$ is just a twist of $\mathcal{K}$ we can deduce that $S(F^\times) = \tilde{\mathcal{K}}$ as well. Now it's an easy calculation: a matrix coefficient is gotten by choosing an element of $\mathcal{K}$ and an element of $\tilde{\mathcal{K}}$ and both of these elements are in $S(F^\times)$; of course one can only control the action of the upper triangular matrices on the Kirillov models but this is enough: one checks that if π is the corresponding matrix coefficient then $\pi\left(\left(\begin{smallmatrix} x & 0 \\ 0 & 1 \end{smallmatrix}\right)\right)$ is in $S(F^\times)$ and this is enough because G is generated by these matrices, $\mathrm{GL}_2(\mathcal{O}_F)$ (which is compact), and the centre.

Finally, assume $\mathcal{K}$ is supercuspidal, and choose $k \in \mathcal{K}$; we know that any element of $S(F^\times)$ is in the Kirillov model of $\tilde{\mathcal{K}}$ and now writing down the statement that the resulting matrix coefficient is compact mod centre, the fact that $k \in S(F^\times)$ easily follows. $\square$

So we now have three equivalent formulations of supercuspidality for infinite-dimensional irreducible admissible representations of $\mathrm{GL}_2(F)$: firstly the definition, secondly some statement about the Kirillov model only having functions in it which vanish in a neighbourhood of 0, and thirdly the statement that for all $v \in V$ there's N such that for all $n \geq N$ we have

$$\int_{(\pi)^{-n}} \left(\begin{smallmatrix} 1 & x \\ 1 & < \end{smallmatrix}\right) 01 v \mathbf{d}x = 0.$$

Note that the equivalence of the first and third conditions went via the Kirillov model, but neither of them mentions a Kirillov model.

Here's an amazing consequence of what we have done so far:

Proposition 38. *Let V be an infinite-dimensional irreducible admissible representation of $\mathrm{GL}_2(F)$; if V is not supercuspidal, then it's a subrepresentation of some $B(\chi_1, \chi_2)$ for some locally constant characters χ_i of $F^\times$.*

Proof. If V is not supercuspidal then $S(F^\times)$ is a proper subset of the Kirillov model $\mathcal{K}$ of V, and the quotient is finite-dimensional. But $S(F^\times)$ is a P-invariant subspace of $\mathcal{K}$ where P denotes the upper triangular matrices in G; so the quotient $\mathcal{K}/S(F^\times)$ is a non-zero smooth representation X of P. Now if $k \in \mathcal{K}$ one checks easily that $\left(\begin{smallmatrix} 1 & x \\ 0 & 1 \end{smallmatrix}\right) k - k \in S(F^\times)$ because $\tau(f) = 1$ for f sufficiently small. So the unipotent upper triangular matrices act trivially on X and now it follows easily that X has a 1-dimensional quotient on which P acts via the character

$\left(\begin{smallmatrix} a & b \\ 0 & d \end{smallmatrix}\right) \mapsto \chi_1(a)\chi_2(d)|a/d|^{1/2}$. This proves that the space $\mathrm{Hom}_P(V, C(\chi_1,\chi_2))$ is non-zero and hence by Frobenius reciprocity the space $\mathrm{Hom}_G(V, B(\chi_1,\chi_2))$ is non-zero. But V is irreducible so any non-zero map is an injection. $\square$

It turns out that it's quite easy to write down a natural space of functions on $F^\times$ and a map from $B(\chi_1,\chi_2)$ to this space, and using it and what we know about Kirillov models it's not so hard to analyse the $B(\chi_1,\chi_2)$ explicitly. We get

Theorem 39 (Theorem 3.3 of [6]; Theorem 6 of [5])**.** *(i) If χ_1/χ_2 is neither $|.|$ nor $|.|^{-1}$ then $B(\chi_1,\chi_2)$ is irreducible. In this case the representations $B(\chi_1,\chi_2)$ and $B(\chi_2,\chi_1)$ are isomorphic. Call the corresponding representation the* principal series representation *corresponding to the pair χ_1,χ_2.*

(ii) If $\chi_1|.| = \chi_2$ then the 1-dimensional subspace $\chi\circ\det$ is the unique non-trivial G-invariant subspace of $B(\chi_1,\chi_2)$ and the quotient is an infinite-dimensional irreducible admissible representation $B_s(\chi_1,\chi_2)$ of G. Call this representation $B_s(\chi_1,\chi_2)$ the special *representation associated to χ_1,χ_2. Finally if $\chi_1 = \chi_2|.|$ then $B(\chi_1,\chi_2)$ contains a unique non-trivial G-invariant subspace, isomorphic to $B_s(\chi_2,\chi_1)$, and the quotient is isomorphic to $\chi\circ\det$.*

All we have to do now is to construct some supercuspidal representations of G. Godement doesn't do this, but one very powerful method, essentially due to Weil, is in section 1 of Jacquet-Langlands. Here's the sketch. Again note the similarity between this and the finite-dimensional case.

Let K be one of the following F-algebras: either $F\oplus F$, a separable quadratic extension of F, a non-split quaternion algebra over F (there's a unique such thing: if L is the unique unramified quadratic extension of F, π is a uniformiser of F and σ is the non-trivial F-automorphism of L then it's the subring

$$\left\{\left(\begin{smallmatrix} a & b \\ \pi\sigma(b) & \sigma(a) \end{smallmatrix}\right) : a,b\in L\right\}$$

of $M_2(L)$), or $M_2(F)$, the two by two matrix algebra over F.

For each of these F-algebras K there's an obvious involution (an antiautomorphism) ι and we define $t,N : K\to F$ by $t(k) = k+\iota(k)$ and $N(k) = k.\iota(k)$. We define τ_K by $\tau_K = \tau\circ N$ and choose a left-invariant Haar measure on K normalised so that something we do below works. We choose a constant γ in a case-by-case way: $\gamma = 1$ if $K = F\oplus F$ or $M_2(F)$, $\gamma = -1$ if K is the non-split quaternion algebra, and γ is some messy Gauss sum in the quadratic case. We define $\omega : F^\times \to \pm 1$ to be the trivial character if K isn't a separable quadratic extension of F, and ω is the quadratic character associated to K in the separable quadratic case. Let $S(K)$ denote the locally constant functions on K with compact support.

Proposition 40. *There's a unique representation of $\mathrm{SL}_2(F)$ on $S(K)$ with the following properties:*

(i) $\left(\left(\begin{smallmatrix} a & 0 \\ 0 & a^{-1} \end{smallmatrix}\right)\phi\right)(x) = \omega(a)|a|_K^{1/2}\phi(ax)$

(ii) $\left(\left(\begin{smallmatrix} 1 & z \\ 0 & 1 \end{smallmatrix}\right)\phi\right)(x) = \tau_F(zN(x))\phi(x)$

(iii) $(w\phi)(x) = \gamma\hat{\phi}(\iota(x))$
where $\hat{\phi}(x) = \int_K \phi(y)\tau_K(xy)\mathbf{d}y$.

Proof. Exactly the same as in the case when F was finite. Note that I have slightly renormalised my Fourier transforms, blame Bump for this. □

Now let U be a finite-dimensional C-vector space and define $S(K, U)$ to be $S(K) \otimes_C U$. Now let Ω be a map $K^\times \to \mathrm{Aut}_C(U)$ making U into a smooth irreducible representation of $K^\times$ (note that this implies that U is 1-dimensional unless K is the non-split quaternion algebra, which has irreducible representations of dimension greater than 1), and define $S(K, \Omega)$ to be the subspace of $S(K, U)$ consisting of functions ϕ such that $\phi(xh) = \omega(h)^{-1}\phi(x)$ for all $x \in K$ and $h \in K^1$, where K^1 is the subset of $K^\times$ consisting of all elements of norm 1. One checks that this is still a representation of $\mathrm{SL}_2(F)$ (by checking it's stable under generators) and now we'll beef this up to a representation of $\mathrm{GL}_2(F)$ as follows. Define G_+ to be the subgroup of $\mathrm{GL}_2(F)$ consisting of all g with determinant in $N(K^\times)$. So $G_+ = \mathrm{GL}_2(F)$ if $K = F \oplus F$ or $K = M_2(F)$, and one checks that $G_+ = \mathrm{GL}_2(F)$ if K is the non-split quaternion algebra as well. If K is a separable quadratic extension of F however, then by local class field theory, or an explicit computation, $N(K^\times)$ has index 2 in $F^\times$ and so G_+ has index 2 in $\mathrm{GL}_2(F)$.

Extend the action of $\mathrm{SL}_2(F)$ on $S(K, \Omega)$ to an action of G_+ by decreeing that

$$\left(\left(\begin{smallmatrix} a & 0 \\ 0 & 1 \end{smallmatrix}\right)\phi\right)(x) = |h|_K^{1/2}\Omega(h)\phi(xh)$$

for $x \in K$, $a \in N(K^\times)$, $a = N(h)$, $h \in K^\times$. Here $|h|_K := |N(h)|_F$ and $|.|_F$ is the canonical norm on a local field (coming from the additive Haar measure). Here are some theorems, proved in [6]: again the proofs are long but essentially elementary, it seems to me. Firstly let K be the non-split quaternion algebra over F.

Theorem 41 (Theorem 4.2 of [6])**.** *The representation of* $\mathrm{GL}_2(F)$ *that we have just constructed is admissible. If* $d = \dim(U)$ *then it's isomorphic to the direct sum of* d *copies of an irreducible representation called* $JL(\Omega)$ *of* $\mathrm{GL}_2(F)$*. If* $d = 1$ *then* Ω *is of the form* $\chi \circ N$ *where* χ *is a character of* $F^\times$ *and in this case* $JL(\Omega)$ *is isomorphic to the special representation associated to* $\chi|.|_F^{1/2}$*. If* $d > 1$ *then* $JL(\Omega)$ *is supercuspidal.*

The proof of course goes by constructing the Kirillov realisation of the representation: one shows that for $d > 1$ the space can be identified with $S(F^\times) \otimes_C U$, for example.

The following theorem might be deeper, it's certainly deeper into [6]. But it might not be, because in the introduction Jacquet and Langlands say that they discovered the results of chapters 15 and 16 whilst writing the first 14.

Theorem 42 (Theorem 15.1 of[6])**.** *The construction taking a representation of* $K^\times$ *and giving a representation of* $\mathrm{GL}_2(F)$ *is injective (i.e. sends non-isomorphic representations to non-isomorphic ones) and its image is the supercuspidals and the specials.*

Remark. Let π be an irreducible admissible representation of an l-group G, and let $C = \mathbf{C}$. We say that π is *square-integrable* if for any $u_1, u_2 \in \pi$ and any $\tilde{u}_1, \tilde{u}_2$ in $\tilde{\pi}$, the integral

$$\int_{Z\backslash G} \tilde{u}_1(gu_1)\tilde{u}_2(g^{-1}u_2)\mathbf{d}g$$

converges absolutely: note that the integrand is constant on cosets of the centre, by Schur's Lemma, and the Haar measure used is not the Haar measure on G but the Haar measure on G/Z. This is an analytic notion: the integral isn't a finite sum in general, it's an infinite sum. It's obvious that supercuspidals are square-integrable because all the matrix coefficients are compact mod centre so the integral in this case is over a compact set. An explicit check (Lemma 15.2 of [6]) shows that special representations are square-integrable too; the integrals are infinite sums but they're bounded above by geometric progressions. Although I can't find it in [6] I think it's the case that the principal series representations aren't square integrable; so this gives an analytic characterization of the image of the JL map.

The idea of the proof is to look at traces of Hecke operators on the representations, and prove some orthogonality relations which are enough to show the injectivity of the map. For surjectivity they seem to use more analytic methods. I think that perhaps to read a proof all you might have to do is to read bits of sections 1,4,7 and 15 of [6] rather than the whole book, and I have read most of section 1 to you already and told you most of the interesting theorems in section 4. But I might be wrong about all this.

I haven't seen the paper "Représentations des algèbres centrales simple p-adiques" by Deligne, Kazhdan and Vignéras, but it's nearly 100 pages long and in it they construct a bijection between the square-integrable representations of $\mathrm{GL}_n(F)$ and $\mathrm{GL}_m(D)$ where D is an F-central division algebra of degree d^2 over F and $dm = n$, although apparently for $m = 1$ Rogowski did the general case; the Steinberg representation of $\mathrm{GL}_n(F)$ corresponds to the reduced norm of D if you're interested in one special case of the dictionary. The proof of Deligne et al is global and uses the trace formula. I don't have a clue what this actually means, to be honest. For the proof to work (I am taking all these statements from the Math Reviews review of the article, by the way) it's necessary to know some growth conditions on the characters of the representations, and these conditions are (or were, in 1984) only known in characteristic 0.

Now let K be a separable quadratic extension of F. We have constructed a representation of G_+ and this has index 2; we now simply induce up to get a representation of $\mathrm{GL}_2(F)$. Note that in this case Ω is just a character of $K^\times$, and U is 1-dimensional.

Theorem 43 (Theorem 4.6 of [6])**.** *This induced representation is admissible and irreducible. If $\Omega = \chi \circ N$ for some character χ of $F^\times$ then the representation is isomorphic to the principal series representation associated to the pair $\chi, \chi\omega$ where ω is the character of $F^\times$ associated to K/F by local Class Field theory. If Ω doesn't factor through N then the induced representation is supercuspidal.*

I think it's true that in residue characteristic greater than 2 we again have constructed all the supercuspidals.

Here are some remarks that hint at the Local Langlands correspondence. We have constructed irreducible admissible representations of $\mathrm{GL}_2(F)$ from the following data: principal series and special, we used two characters of $F^\times$. Supercuspidal, we used a character of a quadratic extension of F. On the other hand, how would one construct two-dimensional, not necessarily irreducible, representations of $\mathrm{Gal}(\overline{F}/F)$? Recall that by class field theory, the abelianisation of $\mathrm{Gal}(\overline{F}/F)$ is isomorphic to some kind of completion of $F^\times$. One could form the direct sum of two 1-dimensional representations of $\mathrm{Gal}(\overline{F}/F)$, which would essentially be two characters of $F^\times$, or to get some irreducible ones (indeed, to get all the irreducible ones if the residue characteristic is greater than 2) one could induce characters of a separable quadratic extension of K. This leads us to wondering whether there is a connection between irreducible admissible representations of $\mathrm{GL}_2(F)$ and 2-dimensional representations of $\mathrm{Gal}(\overline{F}/F)$ or some related group. We have to explain those special representations though, maybe they correspond to non-semi-simple representations or something. I will explain all of this in the next lecture.

4 Bernstein and Zelevinsky's work on $\mathrm{GL}_n(F)$.

Here is a far-reaching generalisation of some of the stuff we just did for $\mathrm{GL}_2(F)$. In fact it's much more convenient to work with groups of the form $\mathrm{GL}_{n_1}(F) \times \mathrm{GL}_{n_2}(F) \times \ldots \times \mathrm{GL}_{n_r}(F)$, which will always be thought of as embedded in $\mathrm{GL}_n(F)$, $n = n_1 + n_2 + \ldots + n_r$, in the obvious way.

Let's start with some definitions. Firstly a general situation. If G is an l-group and H is a closed subgroup and V is a smooth representation of G then let $V(H)$ denote the subspace of V spanned by $hv - v : v \in V, h \in H$. This is a sub-$H$-representation but in general not a sub-G-representation. In fact we can do a little better: if $N_G(H) = \{g \in G | ghg^{-1} \in H \text{ for all } h \in H\}$ then $V(H)$ is $N_G(H)$-invariant. We're actually interested in the quotient $V_H := V/V(H)$, which is also $N_G(H)$-invariant.

The case which we're particularly interested in is the following. For the time being α will denote an (ordered) finite sequence $n_1, n_2, \ldots, n_r$ and G_α will denote $\mathrm{GL}_{n_1}(F) \times \ldots \times \mathrm{GL}_{n_r}(F)$, embedded in $GL_n(F)$ via diagonal matrices, $n = n_1 + n_2 + \ldots + n_r$. If β is another finite sequence $m_1, m_2, \ldots, m_s$ and $m_1 + m_2 + \ldots + m_s = n$ as well, then we say $\beta \leq \alpha$ if $G_\beta \subseteq G_\alpha$, that is, if for all $1 \leq i \leq r$ there is $1 \leq j \leq s$ such that $m_1 + \ldots + m_j = n_1 + \ldots + n_i$ (picture).

More notation: P_α is the matrices in $\mathrm{GL}_n(F)$ which are "upper triangular with respect to α", that is, perhaps not literally upper triangular, but block upper triangular (picture). There's a canonical surjection $P_\alpha \to G_\alpha$; let U_α denote the kernel (picture). Finally, if $\beta \leq \alpha$ then define $P_\beta(\alpha) = P_\beta \cap G_\alpha$ (picture) and $U_\beta(\alpha) = U_\beta \cap G_\alpha$ (picture). The observation we'll use later on is that G_β is in $N_{G_\alpha}(U_\beta(\alpha))$.

Say $\beta \leq \alpha$. Here are the Jacquet functors between representations of G_β and

G_α. If W is a representation of G_β then we can consider it as a representation of $P_\beta(\alpha)$ and induce up: define $i_{\alpha,\beta}(W) = \mathrm{ind}_{P_\beta(\alpha)}^{G_\alpha}(W)$. This takes smooth representations to smooth representations and we also saw earlier (Lemma 18) that it takes admissible representations to admissible representations. Here's the new thing: if V is a representation of G_α then define $r_{\beta,\alpha}(V) = V_{U_\beta(\alpha)}$, the coinvariants. It's trivial that this takes smooth representations to smooth representations.

Exercise. If V is an admissible representation of G_α then $r_{\beta,\alpha}(V)$ is too. Hint: it's Theorem 3.14 of [1]. Another hint: if K is an open subgroup of G_α and $K^0 = K \cap G_\beta$ then the natural projection $V^K \to (r_{\alpha,\beta}(V))^{K^0}$ is an isomorphism and this is enough.

Lemma 44. *$i_{\alpha,\beta}$ and $r_{\beta,\alpha}$ are adjoints on the category of smooth representations.*

Proof. This follows formally from Frobenius reciprocity. If V is a smooth representation of G_α and W is a smooth representation of G_β then we know

$$\mathrm{Hom}_{G_\beta}(V, i_{\alpha,\beta}(W)) = \mathrm{Hom}_{P_\beta(\alpha)}(V, W)$$

by Frobenius reciprocity, and because $U_\beta(\alpha)$ acts trivially on W by definition the right hand side is just $\mathrm{Hom}_{G_\beta}(r_{\beta,\alpha}(V), W)$. $\square$

The quite-simple-to-prove but powerful theorem, whose "philosophical" significance will be apparent later, is:

Theorem 45. *If V is a smooth representation of G_α then V is supercuspidal iff $r_{\alpha,\beta}(V) = 0$ for all $\beta < \alpha$.*

Remark. Of course taking $\alpha = \{n\}$ we get theorems about representations of $\mathrm{GL}_n(F)$.

Look at the consequences of this theorem!

Corollary 46. *If V is an irreducible smooth representation of G_α then there's some $\beta \leq \alpha$ and some irreducible supercuspidal representation W of G_β such that V is a subrepresentation of $i_{\alpha,\beta}(W)$.*

Proof. (of corollary) If V is supercuspidal take $\beta = \alpha$. If not then choose $\beta < \alpha$ such that $r_{\beta,\alpha}(V) \neq 0$ but $r_{\beta',\alpha} = 0$ for all $\beta' < \beta$. Set $W = r_{\beta,\alpha}(V)$. It's easy to check that V is finitely-generated as a $P_\beta(\alpha)$-module because one checks that $G_\alpha = P_\beta(\alpha)\Gamma$ where $\Gamma = \prod_i \mathrm{GL}_{n_i}(\mathcal{O}_F)$, and hence that if $0 \neq v \in V$ then Γv is a finite set which generates V as a $P_\beta(\alpha)$-module. Hence W is finitely-generated as a G_β-module and hence has an irreducible quotient W'. Then W' is supercuspidal by the theorem, and $\mathrm{Hom}_{P_\beta(\alpha)}(V, W') \neq 0$ and so we're done by Frobenius reciprocity. $\square$

Corollary 47. *Irreducible smooth representations of $\mathrm{GL}_n(F)$ are admissible.*

Proof. (of corollary) Supercuspidals are admissible for any group, and inducing an admissible gives an admissible. □

Remark. In fact the proof that we give will generalise to a proof that if G is the F-points of any connected reductive group over F then irreducible smooth representations of G are admissible.

Remark. We proved the first corollary for $G = \mathrm{GL}_2(F)$ recently (Proposition 38) using the Kirillov model. I wouldn't have done this if I'd realised that there was a Kirillov-model-free proof.

Proof. (of theorem) It's 3.21 in BZ (note that they don't use the word "supercuspidal" and they *define* quasi-cuspidal to mean $r_{\beta,\alpha}(V) = 0$ for all $\beta < \alpha$ so don't get confused. BZ attribute the theorem to Harish-Chandra. The key point is to prove that both statements are equivalent to the following, for V a smooth representation of G_α: for any $v \in V$ and any congruence subgroup $N \subseteq G$ the set

$$K_{v,N} := \{g \in G | e_N g^{-1} v \neq 0\}$$

is compact mod centre. It's just now a long exercise. The idea is that if V is smooth then you can classify $V(U_\beta(\alpha))$ as the $v \in V$ such that for any congruence subgroup N there is a number $t = t(v, N)$ such that $e_N g^{-1} v = 0$ for all g diagonal and such that the absolute value of log of the ratio of the gs entries which are in the same α but different β segments, are all at least t. It's elementary but messy, it's again only half a page though. □

4.1 Whittaker and Kirillov models.

These do exist for $\mathrm{GL}_n(F)$, and they are unique (multiplicity one). Recall the theorem for $\mathrm{GL}_2(F)$: a representation had a Whittaker model iff it was infinite-dimensional. It's not so easy for $\mathrm{GL}_n(F)$. A representation has a Whittaker model iff it's *non-degenerate* (this is a definition) and there are infinite-dimensional degenerate representations. This muddles me a bit because I thought that you needed Whittaker models to define L-functions. This is all I'm going to say about this.

4.2 Local Langlands for $\mathrm{GL}_n(F)$.

We have proved that the irreducible smooth representations of $\mathrm{GL}_1(F)$ are just the characters of $F^\times$. I will just assume some Galois theory of local fields now. Let G_F denote the absolute Galois group of F. It fits in an exact sequence

$$0 \to I_F \to G_F \to \hat{\mathbf{Z}} \to 0$$

Define the Weil group to be the subgroup of G_F consisting of elements which map to $\mathbf{Z}$, topologised so that I_F is open. Then the abelianisation of W_F is isomorphic to $F^\times$. We define a *Weil-Deligne representation* of W_F to be a pair consisting of a continuous (that is, locally constant) representation $W_K \to$

$\mathrm{GL}_n(C)$, and a nilpotent matrix $N \in M_n(C)$ such that $\sigma N \sigma^{-1} = |\sigma|^{-1} N$ where $|\sigma|$ is the map $W_F \to \mathbf{Z} \to C^\times$ sending 1 to the reciprocal of the size of the residue field. A Weil-Deligne representation is *semi-simple* if the underlying representation of W_F is semi-simple, that is, a direct sum of irreducibles. It's a bit late in the day to explain why these things might be the right thing to look at, but one good reason is

Proposition 48. *(Grothendieck) Let C be the algebraic closure of $\mathbf{Q}_l$ for some prime l not equal to the residue characteristic of F. Then there's a natural bijection between n-dimensional l-adic representations, that is, (isomorphism classes of) continuous (with respect to the l-adic topology) Galois representations $G_F \to \mathrm{GL}_n(C)$ and Weil-Deligne representations (ρ_0, N) such that the eigenvalues of $\rho_0(\phi)$ (ϕ any lift of $1 \in \mathbf{Z}$ to W_F) are l-adic units.*

The proof is an explicit construction which doesn't really use any new ideas; see Tate's Corvalis article. But the proposition links Weil-Deligne representations with things naturally arising in number theory (e.g. etale cohomology, Tate modules, and so on).

Much much much deeper than this is (back to any algebraically closed C of characteristic 0, I think):

Theorem 49 (Local Langlands conjecture)**.** *(Weil, Jacquet-Langlands, Drinfeld, Tunnell, Kutzko, Henniart, Laumon-Rapoport-Stuhler, Harris, Harris-Taylor, others as well probably) There is a canonical bijection between irreducible smooth representations of $\mathrm{GL}_n(F)$ and n-dimensional F-semi-simple Weil-Deligne representations of W_F.*

It's hard to make "canonical" meaningful—this was one of Henniart's contributions. Say π and σ are in bijection. It turns out that one can define conductors of both sides, and they should be equal. For $n = 1$ the canonical bijection should be that of Local Class Field Theory. For general n oOne can go from GL_n to GL_1 on both sides: take the determinant of the Galois representation, and take the central character of the smooth representation; these should match up. One can define L-functions (I *think*: does one need a Whittaker model?) and ϵ-factors of both sides and these should match up. Unfortunately one can prove that if there is one bijection where all these match up then there is more than one! So one asks for more (ϵ factors for pairs etc). It's a bit messy.

I will make some comments about the case $n = 2$. Here we have pretty much written everything down. Let's try and classify 2-dimensional Weil-Deligne representations. Firstly let's try $N = 0$. Then we have irreducible ones and reducible ones. The reducible ones are just two characters, so by GL_1 case we get two characters on the automorphic side, and we can put them together to get a principal series, which is the associated smooth representation. Let's forget about the irreducible ones for a minute and consider the case $N \neq 0$. Then there's only one choice for N and if I had time to unravel the definitions I'd show you that it's easily shown that the only possibility for the representation of the Weil group is diagonal with two characters whose ratio is the norm; this corresponds to the special case, which was exactly when we had a "spare" representation after inducing. Finally we have the irreducible Galois representations and the supercuspidal smooth representations; these should match up. We saw a very elaborate construction of supercuspidals: start with a separable quadratic extension K of F and a character of K. By Local Class Field Theory this gives a Weil-Deligne representation of W_K. If the character didn't factor through the norm then the resulting smooth representation of $\mathrm{GL}_2(F)$ was supercuspidal: the corresponding statement on the Galois side is that the Weil-Deligne representation of W_K, when induced to W_F, is irreducible. When $p = 2$ there are representations of W_F that aren't induced (just write them down) and conversely Weil wrote down some smooth representations that weren't obtainable via his elaborate construction, adding more weight to the speculation when $n = 2$. So most of those theorems for GL_2 on the representation theory side were very elaborate analogues of very basic constructions on the Galois side.

Note finally that in the GL_2 case the supercuspidals matched up with the irreducible Galois representations (where necessarily $N = 0$). Imagine that one could "only" prove a bijection between supercuspidals and irreducibles for all n. Then in fact one would be done because of the result I told you about earlier: the statement that a non-supercuspidal is a subrepresentation of an induced would be enough, it's the analogue of saying that if a representation isn't irreducible, it's reducible and hence comes from a smaller G_α.

The proofs of Local Langlands, in both the function field case (Laumon, Rapoport, Stuhler, 1994) and the p-adic field case (Harris-Taylor, with Harris first when $p > n$ I think) are all global. Basically I am saying that I haven't explained anywhere near enough tools to show you how they did it. Colin Bushnell tries to prove the conjecture using local methods by mimicing the classical approaches: one simply tries to write down all the elements of both sides and match them up. This works for GL_2 if $p > 2$ and I think Henniart did it for GL_3 too. Henniart's talk a few months ago seemed to indicate to me that they were getting close to doing this if $p \nmid n$.

There is a Local Langlands conjecture for general reductive groups, I think, but it seems to me that people haven't worked out what "canonical" means and furthermore I think people don't think that it should be a bijection. Even for GSp_4 I think there are things called "L-packets" that should all correspond to the same thing on the Galois side. I don't understand this stuff at all.

5 Other things in the local theory that I didn't have time to cover.

Non-degenerate (is this the same as generic?) representations. Whittaker and Kirillov models for $\mathrm{GL}_n(F)$. Structure of certain local Hecke algebras—the Satake isomorphism. L-functions and ϵ-factors. Precise statement of the Local Langlands conjecture for $\mathrm{GL}_n(F)$ and perhaps for arbitary reductive groups. L-groups. Definitions of essentially square-integrable, tempered, essentially tempered, discrete series, complementary series.

References

[1] Bernstein and Zelevinsky, *Representations of the group* $\mathrm{GL}(n, F)$*, where* F *is a non-archimedian local field*, Russian Math Surveys 31:3 (1976), pp1–68.

[2] D. Bump, "Automorphic forms and representations", Cambridge studies in advanced mathematics 55, CUP.

[3] P. Cartier, *Representations of* p*-adic groups*, in [4].

[4] "Automorphic forms, representations, and L-functions", Proc Sympos Pure Math Vol XXXIII Parts 1 and 2, ed. A. Borel and W. Casselman.

[5] R. Godement, *Notes on Jacquet-Langlands Theory*, mimeographed lecture notes.

[6] Jacquet-Langlands *Automorphic forms on* GL_2, Springer LNM 114.