Some introductory notes on the Local Langlands correspondence.

Last modified some time in the 1990s.

A lot of the introductory papers published on Langlands' stuff are written for people who know a fair amount of analysis. These notes were written principally for number theorists (in fact they were written principally for David Jones!), to temporarily fill in the gap in the literature which will one day be permanently filled by Leila Schneps' book on automorphic representations. She's writing up some notes from a conference which took place in Cambridge in 1993, and that conference is where I learnt a lot of this stuff. I also learnt a lot from various lectures and a course given by Richard Taylor in 1992. There are essentially no proofs in this note, just some recipes.

I essentially sat down and just wrote these notes over the course of 2 evenings, so there are bound to be some inaccuracies. If you spot any, however trivial, or have any ideas on how improve these notes, then feel free to email me at `buzzard@dpmms.cam.ac.uk`.

Let K be a finite extension of $\mathbb{Q}_p$ for some prime p. I should perhaps add before we start that there is a wider class of fields for which things like all this work, for example fields like $\mathbb{F}_p((t))$, and even the real numbers and the complexes, but we'll stick to finite extensions of $\mathbb{Q}_p$ because it makes life easier.

This document makes a very vague and brief attempt to deal with the following questions:

1) What is an admissible representation of $\mathrm{GL}_n(K)$?

2) What are all the admissible irreducible representations of $\mathrm{GL}_n(K)$ for $n \leq 2$?

3) What is the Local Langlands correspondence?

4) What *is* the Local Langlands correspondence, explicitly, for small values of n?

5) What is the Jacquet-Langlands theorem, vaguely speaking?

§0. Notation.

Throughout, p is a prime, l is a prime not equal to p, K is a finite extension of $\mathbb{Q}_p$, and $\mathcal{O}$ denotes the integers in K. Also, q is the size of the residue field, and π is a uniformiser (although in the literature π is frequently used to denote an admissible representation, we shall use letters like V to denote these).

§1. What is an admissible representation of $\mathrm{GL}_n(K)$?

For simplicity let's just talk about complex representations. Let K be as in §0. Then K is a topological ring, so $\mathrm{GL}_n(K)$ inherits the structure of a topological group. We want to study (usually infinite-dimensional) complex representations of this group, and there are problems with topologies here. The topology on $\mathrm{GL}_n(K)$ is rather p-adic, and any sane topology that one might one to put on the automorphisms of a complex vector space will not be. But there is a pretty neat way of expressing a useful concept of continuity, as I shall now explain. By the way, I never got my head around $\TeX$ and pagebreaks, so if this definition is over 2 pages then rotten luck.

Definition. A (complex) *admissible representation* of $\mathrm{GL}_n(K)$ is a complex vector space V equipped with an action of $\mathrm{GL}_n(K)$ (by which I of course mean a group homomorphism $\rho\colon \mathrm{GL}_n(K) \to \mathrm{Aut}_{\mathbb{C}}(V)$) such that

a) If $U \subseteq \mathrm{GL}_n(K)$ is an open subgroup, then V^U, the set of vectors $v \in V$ which are fixed by every $u \in U$, is a finite-dimensional vector space, and

b) If $v \in V$ then the stabilizer of v in $\mathrm{GL}_n(K)$ is open.

These notions above give a (perhaps rather strong) notion of continuity.

Definition. We say that an admissible representation $\mathrm{GL}_n(K) \to \mathrm{Aut}(V)$ is *irreducible* if V is non-zero but the only stable subspaces are 0 and V.

Note that if V is admissible and $W \subseteq V$ is $\mathrm{GL}_n(K)$-stable, then W is automatically admissible.

One of the aims of the game is to try and classify all irreducible admissible representations. For small n one can try this "by hand", and I'll explain vaguely what the classification is for $n = 1$ now and $n = 2$ later on in these notes.

§2 What are all the admissible irreducible representations of $\mathrm{GL}_n(K)$ for $n = 1$?

If $n = 1$ and V is a non-zero admissible representation of $\mathrm{GL}_1(K) = K^\times$ then choose $0 \neq v \in V$. The stabilizer of v is then open in $K^\times$; call it U. Because $K^\times$ is commutative we have that V^U is stable under the group action, and it contains v so it is non-zero. Hence if V is irreducible then $V = V^U$, and hence V is finite-dimensional. Now it's rather elementary to see that V must in fact be 1-dimensional (use the fact that over the complexes, every matrix has an eigenvector) and the representation is in fact a map $\chi : K^\times \to \mathbb{C}^\times$ with open kernel, that is, with kernel containing $1 + \pi^n\mathcal{O}$ for some $n > 0$.

We shall refer to 1-dimensional irreducible admissible representations of $\mathrm{GL}_n(K)$ as *admissible characters*, or even just *characters*.

§3 Weil-Deligne representations.

If K is a finite extension of $\mathbb{Q}_p$ then let's recall the definition of the Weil group W_K associated with K. First recall that K has a maximal unramified extension K^{nr}, and the Galois groups of unramified extensions are controlled by the Galois groups of the residue fields, so $\mathrm{Gal}(K^{nr}/K) = \widehat{\mathbb{Z}}$ canonically. Moreover, there is a canonical reduction (or valuation) map $v : \mathrm{Gal}(\overline{K}/K) \to \mathrm{Gal}(K^{nr}/K) = \widehat{\mathbb{Z}}$.

Definition. The *Weil group* W_K of K is the subgroup of $\mathrm{Gal}(\overline{K}/K)$ consisting of elements σ such that $v(\sigma) \in \mathbb{Z}$.

We topologise W_K, but not by giving it the subspace topology. We give W_K the topology such that the subgroup $I_K = \mathrm{Gal}(\overline{K}/K^{nr}) = \ker(v)$ of W_K is open and has its usual profinite topology. This is enough to define the topology on W_K uniquely.

Definition. A (complex) Weil-Deligne representation of K is a pair (ρ, N) where

a) $\rho : W_K \to \mathrm{GL}_n(\mathbb{C})$ is continuous with respect to the discrete topology on $\mathrm{GL}_n(\mathbb{C})$

b) N is a nilpotent complex n by n matrix, and

c) $\rho(\sigma)N\rho(\sigma)^{-1} = |\sigma|^{-1}N$, where $|\sigma| := q^{-v(\sigma)}$ (recall that q is the order of the residue field of K and v is the valuation map).

I should remark that there is in fact a Weil-Deligne group (scheme) and Weil-Deligne representations as defined above actually correspond bijectively to representations of the Weil-Deligne group.

Let's try to work out some examples of Weil-Deligne representations. If $n = 1$, we see that N is nilpotent and 1 by 1, so N has to be zero and a Weil-Deligne representation is just a continuous homomorphism $W_K \to \mathbb{C}^\times$. Note also that I_K is profinite and $\mathrm{GL}_n(\mathbb{C})$ contains no small subgroups, so the image of I_K is finite and hence ρ "isn't too bad", in the sense that the hard bit of W_K has got finite image, and the easy bit of W_K is just $\mathbb{Z}$ which is very easy.

Because the image of I_K is finite, it can be given a filtration using the standard "lower numbering", and the usual definition of the conductor of ρ makes sense. We shall think of conductors as being positive integers, as opposed to powers of π. The conductor of a Weil-Deligne representation should perhaps involve N though, so we make the following

Definition. If (ρ, N) is a Weil-Deligne representation, with ρ and N thought of as acting on the vector space $V = \mathbb{C}^n$, then the *conductor* of (ρ, N) is the sum of the conductor of ρ and $\dim(V^{I_K}/(\ker(N))^{I_K})$.

One can semisimplify Weil-Deligne representations: given a Weil-Deligne representation of W_K there is a unique unipotent $u \in \mathrm{GL}_n(\mathbb{C})$ such that if ρ_u is defined by $\rho_u(\sigma) = \rho(\sigma)u^{v(\sigma)}$ then (ρ_u, N) is a Weil-Deligne representation and ρ_u is semi-simple (that is, the direct sum of irreducibles). Of course, if ρ is already semi-simple then one can take u to be the identity. I think Tate's Corvalis paper (Proc Sympos Pure Math XXXIII) might say something about this.

If (ρ, N) is a Weil-Deligne representation and ρ is semi-simple, then (ρ, N) is said to be *F-semi-simple.*

Now let l be a prime not equal to p. The field $\overline{\mathbb{Q}}_l$ is algebraically closed of characteristic 0 and of cardinality that of the complexes, and so it's a theorem that it is isomorphic, as a field, to the complexes. Choose such an isomorphism. Then we have the pretty neat (but not too hard) theorem of Grothendieck:

Theorem (Grothendieck): There is a canonical bijection between

1) Isomorphism classes of continuous (with respect to the *usual topologies*) n-dimensional representations $\mathrm{Gal}(\overline{K}/K) \to \mathrm{GL}_n(\overline{\mathbb{Q}}_l)$, and

2) Isomorphism classes of complex Weil-Deligne representations of W_K such that the eigenvalues of some chosen lifting of Frobenius are l-adic units.

Furthermore, under this correspondence, the F-semi-simple Weil-Deligne representations match up with the Galois representations ρ with the property that $\rho(\phi)$ is diagonalisable, where ϕ is any element of $\mathrm{Gal}(\overline{K}/K)$ such that $v(\phi) = 1$.

The proof of this theorem is a long exercise, and I won't do it. As I recall it's also done in Tate's Corvalis article. The main point is that we're considering local representations and the only hard part of Galois is the p-part of I_K, and this part of the Galois group can't behave too badly because we're interested in complex or l-adic representations of this group and $l \neq p$.

I shall explain explicitly a recipe explaining this theorem for $n = 2$ later on. Note that the case $n = 1$ is not too hard.

§4 Local Langlands.

We have seen that an irreducible admissible representation of $K^\times$ is nothing more than a map $K^\times \to \mathbb{C}^\times$ which is continuous with respect to the discrete topology on $\mathbb{C}^\times$. But this is precisely what a complex 1-dimensional Weil-Deligne representation of W_K is, because the topological abelianisation of W_K is just (by class field theory) isomorphic to $K^\times$. So (if we believe local class field theory) we believe the case $n = 1$ of the following conjecture:

Conjecture (Local Langlands): If $n > 0$ is an integer, then there is a canonical bijection between irreducible admissible complex representations of $\mathrm{GL}_n(K)$ and F-semisimple n-dimensional complex Weil-Deligne representations.

I could be more specific about how canonical this bijection is, but I won't. It certainly preserves, for example, conductors, although I didn't say what the conductor of an admissible representation was. It also preserves ϵ-factors (so I'm told). There are more things that it does as well. I'll be more specific for $n = 2$ later.

The Local Langlands conjecture is a theorem for $n = 1$ and also for $n = 2$. It's also a theorem in many more cases, but I don't know exactly what is known. It might actually even be a theorem for all n by now. My understanding is that the hard cases are those when $p \leq n$, and indeed we shall see later that $p = n = 2$ is the first hard case.

Let's now restrict to $n = 2$ and try and really understand both sides of this Local Langlands conjecture.

§5 Admissible irreducible representations of $\mathrm{GL}_2(K)$.

As ever, K is a finite extension of $\mathbb{Q}_p$. Here is a list of irreducible admissible representations of $\mathrm{GL}_2(K)$.

Firstly, there are the 1-dimensional ones. Let χ be any admissible complex character of $K^\times$. Then composing with the determinant map $\mathrm{GL}_2(K) \to K^\times$ we get an irreducible admissible representation, which is of course 1-dimensional. It turns out that these are the only finite-dimensional irreducible admissible representations of $\mathrm{GL}_2(K)$.

It is *possible* to write down explicitly some infinite-dimensional representations of $\mathrm{GL}_2(K)$, but I have chosen not to do so. What I shall do is just to say that they exist and to give them names. They are not that mysterious, and one can see definitions in many of the books or in Richard Taylor's amazing notes on this stuff from Caltech, 1992. So here are some names for the infinite-dimensional ones.

Let χ_1 and χ_2 be two admissible complex characters of $K^\times$, such that χ_1/χ_2 does not equal either the usual norm on K (sending a uniformiser to $1/q$), or its inverse. Then there is an irreducible admissible representation $PS(\chi_1, \chi_2)$ associated to the (unordered) pair $\{\chi_1, \chi_2\}$, which turns out to be infinite-dimensional but not too hard to describe. A concrete description of some of these representations is given in Richard Taylor's amazing notes on this stuff from Caltech, 1992. They're just explicit infinite-dimensional spaces of functions with an explicit action of $\mathrm{GL}_2(K)$ and I don't think we'd gain much if I said what they were.

The above two classes of representations are called *principal series* representations. One can think of the admissible 1-dimensional representation as being associated to the "missing case" above, that is, the admissible 1-dimensional representation $\chi \circ \det$ is the principal series representation (if I got the normalisation right) $PS(\chi|.|^{1/2}, \chi|.|^{-1/2})$.

Next, I'll tell you almost nothing about the *special* representations. If χ is an admissible character of $K^\times$ then there is an irreducible admissible representation $S(\chi)$ associated with χ, called the special representation associated to χ. Distinct characters give distinct representations. One can also write down an explicit countably infinite-dimensional vector space with an explicit action of $\mathrm{GL}_2(K)$, which corresponds to this reprsentation, and again I'm not going to do this.

Finally, there are the *supercuspidal* ones. These are rather tricky to describe in general, but there is a trick which works in many cases. Let L/K be an arbitrary quadratic field extension, and let χ be an admissible character of $L^\times$ such that χ does not equal $\chi\tau$, where τ is the non-trivial element of $\mathrm{Gal}(L/K)$. Then one can obtain a supercuspidal representation (called a "base change") of $\mathrm{GL}_2(K)$ from χ. Let's call this representation $BC(L/K, \chi)$. It turns out that if $p \neq 2$ then *every* supercuspidal representation arises in this way. Moreover $BC(L/K, \chi)$ and $BC(L'/K, \chi')$ are isomorphic if and only if the induced representations $\mathrm{Ind}_{W_L}^{W_K}(\chi)$ and $\mathrm{Ind}_{W_{L'}}^{W_K}(\chi')$ are isomorphic (note that I am switching between thinking of χ as being a representation of $K^\times$ and of W_K, so in other words I'm using class field theory implicitly here).

So far I've just told you that there are certain representations which have names, and the whole point of giving these representations names is that I can talk about them, and even state now that if $p \neq 2$ then we have now given a name to every irreducible admissible representation of $\mathrm{GL}_2(K)$.

If $p = 2$ then there are more, which shouldn't be surprising when you consider that there are also "more Weil-Deligne representations" in the case of $p = 2$, as if $p = 2$ then the image of I_K can be more complicated than if $p \neq 2$. The extra irreducible admissible representations for $p = 2$ are called "extraordinary", and there are infinitely many of them, but they're not too hard to understand via the Langlands correspondence.

By the way, I think French for "supercuspidal" is just "cuspidale" because their cusp forms are "formes paraboliques", so for them there is no confusion. Don't think that supercuspidal representations are somehow to do with cusp forms (although they are, in a weak sense—see later).

§6. Conductors.

We shall give an ad hoc definition of conductor, for admissible representations of $\mathrm{GL}_n(K)$ with n at most 2. For $n = 1$ define a decreasing sequence of subgroups $V(t)$ of $K^\times$ for t a non-negative integer, by $V(0) = \mathcal{O}^\times$ and $V(t) = 1 + (\pi)^t$ for $t > 0$.

Definition. Let χ be an admissible character of $K^\times$. Then the *conductor* of χ is the smallest integer t such that χ is trivial on $V(t)$.

For $n = 2$ things are harder and we shall just deal with everything case by case.

1) If V is a finite-dimensional irreducible admissible representation of $\mathrm{GL}_2(K)$ then V is a 1-dimensional principal series and the representation factors through det and a character χ of $K^\times$. Define the conductor of V to be twice the conductor of χ.

If V is infinite-dimensional then it's either principal series, special, or supercuspidal.

2) If V is infinite-dimensional principal series, then $V = PS(\chi_1, \chi_2)$ and the conductor of V is the sum of the conductors of χ_1 and χ_2.

3) If V is special, isomorphic to $S(\chi)$, then the conductor of V is defined to be either 1 (if χ is unramified) or twice the conductor of χ (if χ is ramified).

4) If V is supercuspidal then things are harder. If $V = BC(\chi)$ where χ is an admissible character of a quadratic extension L/K of K, then the conductor of V is twice the conductor of χ if L/K is unramified, and the conductor of χ plus the conductor of (the quadratic character of the absolute Galois group of K corresponding to) L/K if L/K is ramified.

If $p = 2$ then there are still some representations to go, and it's hard to define conductors for them because we still haven't even given them names. But here is *another* definition of conductor which works for any irreducible admissible infinite-dimensional representation of $\mathrm{GL}_2(K)$, even if $p = 2$.

For t a non-negative integer, set $U_1(t) = \{g \in \mathrm{GL}_2(\mathcal{O}) \mid g \equiv \begin{pmatrix} * & * \\ 0 & 1 \end{pmatrix} \bmod \pi^t\}$. Let V be an admissible irreducible infinite-dimensional representation of $\mathrm{GL}_2(K)$.

Theorem. There is a unique minimal non-negative integer t such that $V^{U_1(t)}$ is non-zero.

Definition. The conductor of V is defined to be t as above.

One can check that the conductor of a supercuspidal representation of $\mathrm{GL}_2(K)$ is always at least 2.

§7 A recipe for Local Langlands if $n = 2$.

We have already seen some explicit theorems which say that certain admissible representations should biject with certain other things, but I didn't really explain how they bijected. This section will make the Local Langlands theorem explicit for $n = 2$.

For $n = 2$, the theorem says that there is a bijection between isomorphism classes of irreducible admissible complex representations of $\mathrm{GL}_2(K)$ and isomorphism classes of 2-dimensional complex F-semi-simple Weil-Deligne representations of W_K. Recall that we have already shown that for $n = 1$ this is a rather easy consequence of local class field theory, and we shall continue to identify admissible characters of $K^\times$ with characters of the Weil-Deligne group. Here's the recipe for $n = 2$.

Let $|.|^{1/2}$ denote the character which is the root of the norm character $|.|$, that is, $|\sigma|^{1/2} = q^{-v(\sigma)/2}$.

1) Let χ be an admissible character of $K^\times$, and let V denote the associated 1-dimensional admissible representation of $\mathrm{GL}_2(K)$ obtained by composing with the determinant map.

Then V is 1-dimensional principal series. The associated Weil-Deligne representation is $\chi|.|^{1/2} \oplus \chi|.|^{-1/2}$, with $N = 0$.

2) Let χ_1 and χ_2 be two admissible characters of $K^\times$ whose ratio is not the norm or its inverse. Then the Weil-Deligne representation associated to $PS(\chi_1, \chi_2)$ is just $\chi_1 \oplus \chi_2$, with $N = 0$. Now one can see that even though the two classes of principal series representations were pretty different, the associated Weil-Deligne representations are very similar-looking.

3) Let χ be an admissible character of $K^\times$. Then the Weil-Deligne representation associated to $S(\chi)$ is $\chi \oplus \chi|.|$ with $N = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$. This is the only case where N is non-zero.

4) Let χ be an admissible character of $L^\times$, where L is a quadratic extension of K. Then the Weil-Deligne representation associated to $BC(L/K, \chi)$ is the 2-dimensional representation $\mathrm{Ind}_{W_L}^{W_K}(\chi)$ induced from the 1-dimensional representation of W_L which corresponds to χ, and $N = 0$.

5) Finally, if we have an extraordinary supercuspidal admissible representation, then this corresponds to a Weil-Deligne representation with $N = 0$ and image isomorphic to something more complicated than everything before, for example, to Weil-Deligne representations such that the image of inertia in $\mathrm{PGL}_2(\mathbb{C})$ is isomorphic to A_4 or S_4. Note that if $p > 2$ then no such representations exist.

So that was an "explicit" Local Langlands correspondence. Note that principal series representations correspond to semi-simple but reducible Weil-Deligne representations with $N = 0$, special representations correspond to semi-simple reducible Weil-Deligne representations with $N \neq 0$, and supercuspidal ones correspond to irreducible Weil-Deligne representations.

§8 A recipe for Grothendieck's theorem if $n = 2$.

As ever, l is a prime distinct from p. Recall that, after fixing an isomorphism $\overline{\mathbb{Q}}_l \cong \mathbb{C}$, then for certain Weil-Deligne representations one will get, by the theorem of Grothendieck in §3, an associated l-adic local Galois representation. If $N = 0$ then it turns out that to get from the Weil-Deligne representation to the Galois representation, one just "extends by continuity" from W_K to $\mathrm{Gal}(\overline{K}/K)$. In particular, if $N = 0$ then the image of inertia in the Galois representation is always finite. Tate's paper is pretty essential reading if you want to know the nuts and bolts of this result, and how to proceed if $N \neq 0$.

In the case $N \neq 0$, that is, the case of special admissible representations, it turns out that the image of inertia in the associated Galois representation is never finite. In this case, there is a basis for which the Galois representation has image in the upper-triangular matrices, and the top right hand corner really is non-trivial. It's this top right hand corner which causes the image of inertia to be infinite. This is all obvious if you read Tate's paper!

Note that the image of inertia is infinite in the Galois representation iff $N \neq 0$ iff the Weil-Deligne representation is associated to a special representation. Note moreover that the Galois representation is a direct sum of two 1-dimensional representations iff the Weil-Deligne representation is associated to a principal series representation. Note finally that the Galois representation is irreducible iff the Weil-Deligne representation is associated to a supercuspidal representation.

§9. The automorphic representations associated to newforms.

It turns out that given an eigenform f of weight k on $S_k(\Gamma_1(N))$, then there is a canonical way of associating to f an admissible irreducible representation $V_{f,p}$ of $\mathrm{GL}_2(\mathbb{Q}_p)$ for all primes p. These admissible representations have the property that if f is a newform of level N and p^t exactly divides N, then the conductor of $V_{f,p}$ is t. The way to construct $V_{f,p}$ from f is surprisingly easy but I am a bit too lazy to explain it here, and anyway I have explained so little else. It's all done in Taylor's Caltech notes.

Given the admissible representation $V_{f,p}$, one gets, via Local Langlands, an F-semi-simple complex Weil-Deligne representation of $W_{\mathbb{Q}_p}$. Now choose $l \neq p$ a prime and choose an isomorphism $\overline{\mathbb{Q}}_l \cong \mathbb{C}$. Take a lifting ϕ of Frobenius to $W_{\mathbb{Q}_p}$ and look at its eigenvalues, considered as elements of $\overline{\mathbb{Q}}_l$. I believe that it turns out that they are l-adic units (I guess they are global integers whose product is a unit) and hence one can apply Grothendieck's theorem. Hence if $l \neq p$ is a prime, one gets an l-adic representation of $\mathrm{Gal}(\overline{\mathbb{Q}}_p/\mathbb{Q}_p)$ from $V_{f,p}$ and hence from f. Unsurprisingly, this representation is precisely the restriction to the local Galois group of the global representation constructed by Deligne. Hence we have a second way of constructing a local representation from a modular form. Furthermore, this second way is much easier than the first, although local representations are very easy to construct, whereas global ones are not.

The fact that the 2 methods of getting local representations from modular forms give the same answer was only proved in about 1989, by Carayol, as far as I can see, although for unramified primes it was done by Deligne much earlier: the unramified primes are essentially taken care of by the fact that we know the trace and determinant of the Galois representation.

§10 A recipe for getting from f to $V_{f,p}$.

Let f be a newform of level N. It's sometimes possible to work out things about $V_{f,p}$ just by looking at f "at p", and here I'll explain what I know about this. Putting these results together with an earlier recipe one gets a recipe for getting from f to the local l-adic Galois representations associated to f. Of course, one already knows this latter recipe when f has level N and p is prime to N, but this recipe works in the harder cases as well.

If p does not divide the level of f then $V_{f,p}$ is principal series, and is in fact one of the infinite-dimensional ones. I have to tell you what the two characters are. Well, they are both unramified. To give an unramified character is to give the value it takes on a uniformiser of K, and the two values in question in this case are the two roots of $X^2 - a_pX + \chi(p)p^{k-1}$ where k is the weight and χ the character of f. Note that the ratio of these characters can't be the norm, because of things we know about a_p, I guess. Note that of course, this gives us the right associated l-adic Galois representation, that is, the trace and determinant of Frobenius are correct.

If p divides the level of f exactly once, then there are two cases. If the character of f has conductor prime to p, then $V_{f,p}$ is special associated to an unramified character which might or might not send Frobenius to a_p, depending on whether or not I got all the normalisations right, but I guess it should. If however the character of f has conductor a multiple of p, then $V_{f,p}$ must be a principal series representation associated to two characters, an unramified one called χ_1 such that $\chi_1(\pi) = a_p$, and a tamely ramified one.

If p divides the level of f at least twice, then anything could be going on. If a_p is non-zero then I think $V_{f,p}$ has to be principal series associated to two characters, at least one of which is unramified. If f can be "untwisted", that is, f is a twist of a newform of conductor prime to p, then one can twist the corresponding admissible representation and this sometimes tells you something, but basically it's harder to say what's going on in this case.

One very useful corollary of this theorem is that now one can sit down and work out what the local Galois representation associated to a modular form is, even in some of the cases where p divides the level of the form.

Of course, if one happens to know things about the Galois representation, then this can also help to determine facts about the admissible representations attached to the form. One silly example of this is the case of weight 1 forms. One knows that the image of Galois is finite in the representations associated to a weight 1 form f. One can deduce that $V_{f,p}$ cannot be special for any p, because then already the local representation would have infinite image of inertia at p in these cases. In particular, if f is a weight 1 newform and p exactly divides the conductor of f then p also exactly divides the conductor of the character of f. Another example is elliptic curves, but I'll explain them more carefully in the next recipe.

§11. A recipe for elliptic curves.

Here's another recipe. If E is a modular elliptic curve coming from a modular form f, then one can ask how the possible local behaviours of E at a prime p match up with the admissible representation $V_{f,p}$ associated to f. Here's how it works. Choose a prime p.

1) $V_{f,p}$ is unramified principal series iff E has good reduction at p. This is because these two cases are the only cases where the conductor is 0, so they must match up.

2) $V_{f,p}$ is special iff E has potentially multiplicative reduction at p. This is because these are the only cases where the image of inertia in the associated Galois representation is infinite, so they must match up.

3) $V_{f,p}$ is special associated to an unramified character iff E has multiplicative reduction. This is because these are the only two subcases of case 2) where the conductor is 1.

4) $V_{f,p}$ is ramified principal series or supercuspidal iff E has bad, but potentially good, reduction. This is because these are the only cases left.

I'm not sure I understand why the two statements below are true, but someone once told me that they were:

5) V is ramified principal series iff E attains good reduction over an abelian extension of $\mathbb{Q}_p$.

6) V is supercuspidal iff E attains good reduction over a non-abelian extension of $\mathbb{Q}_p$.

Presumably 5) and 6) come out in the wash, if they're true, if you think about the associated Galois representations.

§12. Quaternion algebras.

Now the reader will soon have to take a lot on board (or read the notes on quaternion algebras that I will write up soon). The goal in these final few sections is to give some hint of what the Jacquet-Langlands theorem says.

Let K be a finite extension of $\mathbb{Q}_p$ as usual. Let's consider the following rather weird thing: there is a quaternion algebra over K which is not isomorphic to $M_2(K)$. It's not too hard to describe actually, but I won't. Let's just think of it as a weird non-commutative ring, with K in its centre, and such that it's 4-dimensional over K. Let G be the non-zero elements of this quaternion algebra. Then G is a topological group and it makes sense to talk about admissible irreducible complex representations of G. In stark contrast to the $\mathrm{GL}_2(K)$ case, *all* irreducible admissible representations of G are finite-dimensional (because $G/K^\times$ is compact, he said cryptically). One can get some of the way to describing explicitly these admissible representations. Here's at least an example. There is a norm map $\nu : G \to K^\times$. Let χ be an admissible character of $K^\times$. Then $\chi \circ \nu$ is an irreducible admissible representation of G.

§13. The local Jacquet-Langlands theorem.

This theorem is pretty amazing (but not as amazing as the global theorem, which is mind-blowing!). Let G be as in the previous section. Then

Theorem (Jacquet-Langlands): There is a canonical bijection between

1) The irreducible admissible representations of G, and

2) The irreducible admissible representations of $\mathrm{GL}_2(K)$ which are not principal series.

Again one can go via a recipe, but I didn't even tell you what all the irreducible admissible representations of G were so you will have to be content with a fragment of it: here it is. The admissible representation $\chi \circ \nu$ of G in the previous section corresponds to the special representation of $\mathrm{GL}_2(K)$ corresponding to $\chi|.|^{-1/2}$ (again if I got the normalisations right).

Another way of saying "non-principal-series" is "special or supercuspidal" of course, and yet another way is "essentiallement de carré intégrable" which I think means "essentially square-integrable". But if you want to know what we're integrating then you have to start getting your hands dirty with what these representations actually *are*, which is way beyond what I'm prepared to do in this note.

§14. The Global Jacquet-Langlands theorem.

Now let D be a quaternion algebra over $\mathbb{Q}$. Then one can tensor up D to $\mathbb{Q}_p$ for any prime p and get a quaternion algebra D_p over $\mathbb{Q}_p$. This quaternion algebra is usually $M_2(\mathbb{Q}_p)$ but for finitely many primes it isn't, and these primes are called the ramified primes. There is a notion of modular forms on quaternion algebras, and indeed one can even talk about eigenforms. For example, if $D = M_2(\mathbb{Q})$ then a modular form on D is just a classical modular form.

Given an eigenform f on D, and a prime p, one can associate an admissible irreducible representation $V_{f,p}$ of $(D_p)^\times$ to f. This is of course just a generalisation of the case where

$D = M_2(\mathbb{Q})$. Recall that whatever D is, then most of the time $(D_p)^\times$ will be $\mathrm{GL}_2(\mathbb{Q}_p)$ anyway, and we know what we're doing, but some of the time it's the units in a weird quaternion algebra and then we get an admissible representation of this weird thing.

So the most naïve thing that one can do is the following. Take an eigenform f of weight k for D. For simplicity I'll assume that $k \geq 2$ and moreover that if $k = 2$ then f is not "Eisenstein", a notion that I won't bore you with. Anyway, now go through all the primes p. For each prime p there is an admissible representation $V_{f,p}$ of $(D_p)^\times$ associated to f. If p isn't ramified in D then $(D_p)^\times$ is just $\mathrm{GL}_2(\mathbb{Q}_p)$ and let's call this representation X_p. If p *is* ramified then we've got ourselves an admissible representation of the units in a weird quaternion algebra, but via local JL we can again get ourselves an admissible representation of $\mathrm{GL}_2(\mathbb{Q}_p)$, which isn't principal series. Call this latter representation X_p. So we have X_p, an admissible irreducible automorphic representation of $\mathrm{GL}_2(\mathbb{Q}_p)$ for every prime p.

Amazing thing number 1: there is a classical modular eigenform g of weight k such that for all primes p we have $V_{g,p} = X_p$, where of course $V_{g,p}$ is the admissible irreducible representation of $\mathrm{GL}_2(\mathbb{Q}_p)$ associated to g.

Amazing thing number 2: If g is a classical modular eigenform of weight k such that for all primes p ramifying in D, the admissible representation of $\mathrm{GL}_2(\mathbb{Q}_p)$ attached to g is not principal series, then there is always an eigenform f on D such that f and g are associated via Amazing thing number 1.

The upshot of all this is that certain modular forms can be "found" on quaternion algebras, as long as they're not principal series everywhere. Of course, there are examples of modular forms which are principal series everywhere, for example any eigenform of level 1. But there are also plenty of examples of modular forms which aren't principal series somewhere. For example, consider the eigenforms in $S_2(\Gamma_0(p))$. These eigenforms are new at p and must be special, and so one can also find them as eigenforms of weight 2 on the quaternion algebra of discriminant p. Rather surprisingly (see more notes I'm in the process of typing up) the definition of an eigenform on a quaternion algebra of discriminant p isn't so difficult, and results of this form are sometimes useful.

Kevin Buzzard, May 1998; some typos corrected Feb 2007 (thanks to Atsushi Yamagami).