Hodge-Tate theory.

Kevin Buzzard

February 9, 2012

Last modified 22/05/2009. This was my attempt to read an old paper of Serre's on the notion of Hodge-Tate.

1 Basic definitions.

For simplicity let's define a *Fontaine field* to be a finite extension K of the field of fractions of $W(k)$, where k is a perfect field of characteristic p and $W(k)$ is its Witt vectors. Examples are finite extensions of $\mathbf{Q}_p$ but if we work in this slightly greater generality then it helps us (well, it helps me, at least) to see which arguments are "the right ones".

Choose an algebraic closure $\overline{K}$ of K, let $\Gamma_K = \mathrm{Gal}(\overline{K}/K)$, let C denote the completion of $\overline{K}$. Tate showed that Γ_K acts on C, and that the invariants C^{Γ_K} were just K again.

Note that k contains a copy of $\mathbf{Z}/p\mathbf{Z}$, so $W(k)$ contains a copy of $\mathbf{Z}_p$ and so K and hence C contains a copy of $\mathbf{Q}_p$. There is a canonical character $\chi : \Gamma_K \to \mathbf{Z}_p^\times$, the cyclotomic character, defined in the usual way by the action of Galois on p-power roots of unity. We can regard χ as taking values in $\mathbf{Z}_p^\times$ or $\mathbf{Q}_p^\times$ or even $C^\times$. We define $C(i)$ to be the 1-dimensional C-representation of Γ_K on which Γ_K acts via χ^i. Let B_{HT} denote the ring $C[T, T^{-1}]$, with Γ_K acting on T^i via χ^i.

If $V/\mathbf{Q}_p$ is a finite-dimensional vector space with an action of Γ_K then we can let Γ_K act on $V \otimes_{\mathbf{Q}_p} B_{HT}$ via the diagonal action; this is not C-linear but it is $\mathbf{Q}_p$-linear, and in fact it is K-linear, with K acting via its action on B_{HT}. Note that $B_{HT} = \oplus_{i\in\mathbf{Z}} C(i)$ is a Galois-stable decomposition of B_{HT} into graded pieces.

Define

$$\begin{aligned} D(V) &= (V \otimes_{\mathbf{Q}_p} B_{HT})^{\Gamma_K} \\ &= \oplus_i (V \otimes_{\mathbf{Q}_p} C(i))^{\Gamma_K} \end{aligned}$$

Then $D(V)$ is a graded K-vector space. Tate proved that $D(V)$ is finite-dimensional over K, and in fact that the natural map

$$\alpha : D(V) \otimes_K B_{HT} \to V \otimes_{\mathbf{Q}_p} B_{HT}$$

sending $d \otimes c$ to dc was an injection. Hence $\dim_K D(V) \leq \dim_{\mathbf{Q}_p} V$ (tensor up to the field of fractions of B_{HT}). Furthermore, Fontaine observed that B_{HT} was an "admissible ring", from which it follows that α is an isomorphism if and only if $\dim_K D(V) = \dim_{\mathbf{Q}_p} V$ (no doubt this was known to Tate; I'm too lazy to check if he writes this explicitly). We say V is *Hodge-Tate* if equality holds.

In fact another thing we can do is to consider the map $B_{HT} \to C$ sending T to 1; if α is an isomorphism then the base change to C of α via this map is an isomorphism, giving us that the canonical map

$$\oplus_i (V \otimes_{\mathbf{Q}_p} C(i))^{\Gamma_K} \otimes_K C \to V \otimes_{\mathbf{Q}_p} C$$

is an isomorphism; hence $V \otimes_{\mathbf{Q}_p} C$ inherits a grading, and hence we get a map $\mu : (\mathbf{G}_m)_C \to \mathrm{GL}(V)_C$—this is the "Hodge-Tate character" of V. The Hodge-Tate weights of V are the i for which $(V \otimes_{\mathbf{Q}_p} C(i))^{\Gamma_K}$ are non-zero and the multiplicity of the weight i is the K-dimension of this latter space. Hence (counted with multiplicities), an n-dimensional representation has n Hodge-Tate weights.

Example: the cyclotomic character has Hodge-Tate weight -1, which seems to be the usual convention these days. If we change the action of T to be via the inverse of the cyclotomic character then it'll have weight $+1$, which Toby says is better for us. Becuase I'll mostly be working with general representations, nothing I say below will depend on conventions.

2 Some basic properties of D that we need.

I'm just cribbing all of these straight out of periodes p-adiques.

The map D is an additive functor from p-adic representations of Γ_K (by which I mean finite-dimensional $\mathbf{Q}_p$-vector spaces plus a continuous Γ_K-action) to finite-dimensional graded K-vector spaces with grading-preserving maps (p141).

Hodge-Tate representations are stable under duality, tensor product, direct sum, and passage to submodules and quotients. Serre says this on p471 of Ouevres III (his article on algebraic groups associated to Hodge-Tate Galois representations). Moreover the category of Hodge-Tate representations is an abelian tensor category.

But better: the category of Hodge-Tate representations of Γ_K is a sub-Tannakian category of the category of all representations, and the map D is in fact an exact faithful tensor-functor from the category of Hodge-Tate representations to the category of graded vector spaces (p142 of periodes p-adiques).

This last piece of twaddle implies that there's a canonical map $D(V_1)\otimes_K D(V_2) \to D(V_1 \otimes_{\mathbf{Q}_p} V_2)$ and it's an isomorphism and furthermore it's grading-preserving, that the dual of D is D of the dual (as graded K-vector space) and that D is an exact functor.

3 More general coefficient fields.

Now say $E/\mathbf{Q}_p$ is finite of degree d, and V is an n-dimensional E-vector space with an E-linear action of Γ_K, with K as always a Fontaine field. The correct way to define Hodge-Tate is: consider V as an nd-dimensional $\mathbf{Q}_p$-vector space V_0; the E-representation V of Γ_K is Hodge-Tate iff (by definition) V_0 is. Looking at it this way we get nd Hodge-Tate weights. But we can do better: the E-action gives some structure on these weights. Define $D(V) = D(V_0)$. Then $D(V)$ is actually a module for $E \otimes_{\mathbf{Q}_p} K$ and Hodge-Tateness of V is equivalent to the statement that the K-dimension of $D(V)$ is nd. But it's better not to think about it this way: think about Hodge-Tateness as saying that the induced map $D(V) \otimes_K B_{HT} \to V \otimes_{\mathbf{Q}_p} B_{HT}$ is an isomorphism of B_{HT}-modules, because if it is then it's an isomorphism of $E \otimes_{\mathbf{Q}_p} B_{HT}$-modules. Now V is free of rank n over E, so $V \otimes_{\mathbf{Q}_p} B_{HT}$ is free of rank n over $E \otimes_{\mathbf{Q}_p} B_{HT}$, and hence Hodge-Tateness implies that $D(V) \otimes_K B_{HT}$ is free of rank n over $E \otimes_{\mathbf{Q}_p} B_{HT}$. Now $D(V)$ is an $E \otimes_{\mathbf{Q}_p} K$-module and I claim that it must be free of rank n over this ring. For $E \otimes_{\mathbf{Q}_p} K$ is isomorphic to a direct sum of finitely many finite field extensions L_i of K, and so $D(V)$ is just a vector space D_i over L_i for each i. The claim is that all these vector spaces are n-dimensional, and this must be true because they have *some* dimension, and when tensored over K to B_{HT} they become free of rank n over $L_i \otimes_K B_{HT}$, so had better have been of rank n to start with.

The conclusion is that V being Hodge-Tate is equivalent to $D(V)$ having K-dimension nd but is also equivalent to the superficially stronger assertion that $D(V)$ is free of rank n over $E \otimes_{\mathbf{Q}_p} K$.

How do Hodge-Tate numbers work in this setting? Well, $D(V)$ is a K-vector space with a grading whose ith piece is $(V \otimes_{\mathbf{Q}_p} C(i))^{\Gamma_K}$, but this piece is a module for $E \otimes_{\mathbf{Q}_p} K$. So it seems to me that the natural numbers to associate with this situation are to write $E \otimes_{\mathbf{Q}_p} K$ as a direct sum of fields L_i and for each such field we have a graded vector space and we can count the dimensions of each graded piece. The case that Toby is particularly fond of is when $K/\mathbf{Q}_p$ is finite and E contains the normal closure of $K/\mathbf{Q}_p$; then each L_i is E, and the i's run through the maps

$K \to E$. But another case would be when K contains the normal closure of $E/\mathbf{Q}_p$ (for example when $E = \mathbf{Q}_p$); in this case one gets numbers for each $E \to K$.

How does μ work in this setting? Well, if G is the algebraic group $\mathrm{Aut}(V)$ over E, and $H = \mathrm{Res}_{E/\mathbf{Q}_p} G$, then the point is that H is smaller than $\mathrm{Aut}_{\mathbf{Q}_p}(V)$ (for A a $\mathbf{Q}_p$-algebra, an A-valued point of H is an endomorphism of the A-module $V \otimes_{\mathbf{Q}_p} A$ which is both A- and E-linear; for an A-valued point of $\mathrm{Aut}_{\mathbf{Q}_p}(V)$ you don't demand the E-linearity). Now I claim that μ is really a map $\mathrm{GL}_1/C \to H_C$. One can see this directly: indeed $D(V)$ is graded; its graded pieces are also modules for $E \otimes_{\mathbf{Q}_p} K$ (but the graded pieces might not themselves be free: see example below). The isomorphism $D(V) \otimes_K C = V \otimes_{\mathbf{Q}_p} C$ means that $V \otimes_{\mathbf{Q}_p} C$ inherits a grading, but the graded pieces are modules for $E \otimes_{\mathbf{Q}_p} C$, and hence the induced map from $\mathrm{GL}_1(C)$ is taking values in the $E \otimes_{\mathbf{Q}_p} C$-linear endomorphisms of $V \otimes_{\mathbf{Q}_p} C$, and this is precisely what $H(C)$ is. More generally one works with C-algebras, get exactly the same result, and concludes that μ goes from GL_1/C to H_C. One can go a bit further here: because C contains the normal closure of $E/\mathbf{Q}_p$ we have that H_C is canonically $\prod_{\tau:E\to C} \mathrm{GL}(V)_{\tau:E\to C}$, the product of the base-changes of our original group $\mathrm{GL}(V)/E$ to C via the d embeddings. So we could think of μ as being d maps $\mu_\tau : (\mathrm{GL}_1)_C \to \mathrm{GL}(V)_C$, one for each base extension of $\mathrm{GL}(V)$ to C via $\tau : E \to C$.

Note of course that μ itself gives rise to numbers, so we are also getting n numbers for each $\tau : E \to C$, and of course one checks easily that what is happening here is that any $\tau : E \to C$ induces a map $E \otimes_{\mathbf{Q}_p} K \to C$ whose image is a field and a component of $E \otimes_{\mathbf{Q}_p} K$.

Summary: any component of $E \otimes_{\mathbf{Q}_p} K$ gives us n numbers. Any $\tau : E \to C$ gives us a component of $E \otimes_{\mathbf{Q}_p} K$, and hence numbers, but it also gives us a μ.

An easy example of a $D(V)$ which is Hodge-Tate and hence free over $E \otimes_{\mathbf{Q}_p} K$ but whose graded pieces aren't free, is when $E = K$ is the unramified quadratic extension of $\mathbf{Q}_p$. Class field theory tells us that if we choose a uniformiser in K (for example, the number p) then there's an extension of K with Galois group canonically $R^\times$, with R the integers of K, and which corresponds via class field theory to the quotient of $K^\times$ obtained by sending our chosen uniformiser to 1. If we map $R^\times$ to $E^\times$ via the canonical inclusion $R \to E$ then the resulting character is Hodge-Tate. Thought of as a 2-dimensional representation of Γ_K then its Hodge-Tate weights are 0 and -1. But in fact $D(V)$ is free of rank 1 over $E \otimes_{\mathbf{Q}_p} K$, and its graded pieces are K-vector spaces of dimension 1 in degrees 0 and -1, and clearly these can't be free over $E \otimes_{\mathbf{Q}_p} K$.

Now let's put ourselves into Toby's favourite setting—the setting we're in in our paper, when K is a finite extension of $\mathbf{Q}_p$ of degree m, and when E is a sufficiently large finite extension that E contains the normal closure of $K/\mathbf{Q}_p$. Let's see how these notions simplify in this case. In this case we have $E \otimes_{\mathbf{Q}_p} K = \oplus_{\sigma:K\to E} E$ as σ ranges over the m $\mathbf{Q}_p$-algebra maps $K \to E$, and so in this case $D(V)$ becomes the direct sum of m E-vector spaces, each of which is graded and n-dimensional. The upshot is that in this case one attaches n Hodge-Tate weights to each $\sigma : K \to E$; this is the sort of thing that Clozel-Harris-Taylor like to do. It seems to me from thinking about this sort of thing now that from a purely local point of view this seems a little unnatural—the motivation from doing this will, it seems to me, typically be coming from a global setting (i.e. we will probably be in the case where other global features of the situation are predicting Hodge-Tate weights). Furthermore the Hodge-Tate cocharacter doesn't seem to work any more simply in this setting: to get the filtration one still has to tensor up to C so the natural ring over which one is working is still $E \otimes_{\mathbf{Q}_p} C$, and so one is naturally led to consider maps $E \to \overline{K}$ rather than maps $K \to E$, which is was what I had always assumed would come out in the wash: the filtration on $D(V)$ is by $E \otimes_{\mathbf{Q}_p} K$-modules and if E is big then this breaks up as a bunch of E-vector spaces, but to get the μ *on* V you need to go up to C: you can have a μ on $D(V)$ but we want something on V because later on $\mathrm{GL}(V)$ will be replaced by a general reductive group.

4 Coefficient field $\overline{\mathbf{Q}}_p$.

This is just an extension of the previous idea, but for convenience Toby and I will stick to coefficient fields in $\overline{\mathbf{Q}}_p$ and I just wanted to check it all works out fine. Now we have V an n-dimensional vector space over $\overline{\mathbf{Q}}_p$, and an action of Γ_K on V (with K a general Fontaine field at this point).

Pick a basis temporarily. The image of Γ_K in $\mathrm{GL}_n(\overline{\mathbf{Q}}_p)$ is clearly compact and Hausdorff, so by the Baire Category Theorem it's a Baire space, which means that given countably many closed subsets of it, each with empty interior, their union will also have empty interior. But Γ_K is the countable union of $\Gamma_K \cap \mathrm{GL}_n(E)$ as E varies through the finite extensions of $\mathbf{Q}_p$, so $\Gamma_K \cap \mathrm{GL}_n(E)$ has non-empty interior for some E, and hence contains an open and hence finite index subgroup of Γ_K. Enlarging E if necessary, we may assume $\Gamma_K \subseteq \mathrm{GL}_n(E)$. So there's an n-dimensional vector space V_0 over E equipped with an action of Γ_K which tensors up to give V over $\overline{\mathbf{Q}}_p$.

We apply the above twaddle to V_0. The "numbers" version gives us this: for every component of $E \otimes_{\mathbf{Q}_p} K$ we get n numbers. The "μ" version gives us this: for each $E \to C$ an embedding of $\mathbf{Q}_p$-algebras we get $\mu_\tau : (\mathrm{GL}_1)_C \to \mathrm{GL}(V_0)_{\tau,C}$.

But now we need to think about how this data depends on E. For a start, there is no E in the question, so there shouldn't be an E in the answer. But this is OK: we can fix this thus. For the "numbers" version of Hodge-Tate weights, I think we are getting n numbers for every component of $K \otimes_{\mathbf{Q}_p} \overline{\mathbf{Q}}_p$. I am not sure I'm going to persevere with this in the general Fontaine field case, because I'm not really interested in numbers. So let's make the simplifying assumption that $K/\mathbf{Q}_p$ is finite and then see if we can get numbers out. Now a component of $K \otimes_{\mathbf{Q}_p} E$ is just a finite extension of E. Now we get one of these for every map $K \to \overline{\mathbf{Q}}_p$, the point being that E is an explicit subfield of $\overline{\mathbf{Q}}_p$ so any map $K \to \overline{\mathbf{Q}}_p$ gives a map $K \otimes_{\mathbf{Q}_p} E \to \overline{\mathbf{Q}}_p$ whose image has no zero divisors and is hence a field; this is the component we're after. So the "numbers" version gives n numbers for each $K \to \overline{\mathbf{Q}}_p$. And the "$\mu$" version (back to a general Fontaine field): we get a $\mu : \mathrm{GL}_1 \to \mathrm{GL}(V_0)_C$ for each $E \to C$, so given a $\mathbf{Q}_p$-algebra homomorphism $\overline{\mathbf{Q}}_p \to C$ we can base extend $\mathrm{GL}(V_0)$ from E to $\overline{\mathbf{Q}}_p$ and then to C, so the "μ" version of the story is that for each $\overline{\mathbf{Q}}_p \to C$ we're getting $\mu : \mathrm{GL}_1 \to \mathrm{GL}(V)_C$.

But this isn't quite the end of the story. We could have chosen a totally different basis and a totally different E. Then what? As Toby and I both observed, it suffices to check that the data we get doesn't change under field extension (because then you can go up and go down). So say we have V_0/E and a finite extension E' of E within $\overline{\mathbf{Q}}_p$. The "numbers" game (with $K/\mathbf{Q}_p$ finite) gives us a list of H-T weights for each component of $E \otimes_{\mathbf{Q}_p} K$, and if we base change to E' (note that D commutes with tensoring up to E' over E) then dimensions don't change and hence numbers don't change. So indeed the notion of Hodge-Tate weights attached to an embedding $K \to \overline{\mathbf{Q}}_p$ are well-defined.

What about μ? Back to K a general Fontaine field. Here the question is the following. We have E in E' in $\overline{\mathbf{Q}}_p$ and we have a map $\tau : \overline{\mathbf{Q}}_p \to C$. Both E and E' give us $\mu_\tau : \mathrm{GL}_1 \to \mathrm{GL}(V)_{\tau,C}$ and we need to check that they give us the same map. It's easier to think in terms of filtrations. Here's the explicit question. Over E we have a filtration on $V_0 \otimes_{\mathbf{Q}_p} C$ but we're only interested in the component corresponding to $\tau|E$ which is simply the submodule $V_0 \otimes_E C$. Now of course things are becoming clear. Over E' we're interested in the filtration on $E' \otimes_E V_0 \otimes_{\mathbf{Q}_p} C$ but we're only interested in the part of this $E' \otimes_{\mathbf{Q}_p} C$-module corresponding to the component of $E' \otimes_{\mathbf{Q}_p} C$ corresponding to $\tau|E'$ which is $V_0 \otimes_E E' \otimes_{E'} C = V_0 \otimes_E C$ again. Finally the filtration on $D(V_0) \otimes_E E'$ is the same as the filtration on $D(V_0 \otimes_E E')$ and so the filtrations on $V \otimes_\tau C$ coincide. So indeed we have a well-defined μ attached to a Hodge-Tate $V/\overline{\mathbf{Q}}_p$ and an embedding $\overline{\mathbf{Q}}_p \to C$.

5 Representations to more general groups over $\mathbf{Q}_p$.

Let's start with the story of G, a reductive group over $\mathbf{Q}_p$, and K a Fontaine field and a map $\Gamma_K \to G(\mathbf{Q}_p)$. Let's furthermore choose a faithful representation $G \to \mathrm{GL}(V)$ over $\mathbf{Q}_p$; now we have a map $\Gamma_K \to \mathrm{GL}(V)(\mathbf{Q}_p)$. If this map is Hodge-Tate then we get $\mu : (\mathrm{GL}_1)_C \to \mathrm{GL}(V)_C$. The claim is that μ factors through a map to G_C, and let me spend some time convincing myself of this because we'll need to generalise it later.

The argument in Serre's paper does *not* go "let's think about things carefully and unravel the definitions and check it"; it goes like this: "μ is a canonical and functorial thing, and hence can be regarded as some sort of map from certain representations of G to representations of $\mathbf{G}_m$, and so by the Tannakian Formalism it will give rise to a map $\mathbf{G}_m \to G$. It's this argument that we have to understand because we're going to need to generalise it.

Here's the proof then. First take an arbitrary map $G \to \mathrm{GL}(W)$ over $\mathbf{Q}_p$, with W a $\mathbf{Q}_p$-vector space of dimension N. This gives us an induced N-dimensional Galois representation. I claim that this Galois representation is Hodge-Tate. This is true because V is a faithful representation of G, and any representation of G, it says here in Proposition 3.1 of Deligne's article in SLNM900 (Deligne-Milne-Ogus-Shih), that the N-dimensional Galois representation is a subobject of $V \otimes V \otimes \ldots \otimes V \otimes V^\vee \otimes V^\vee \otimes \ldots \otimes V^\vee$ (This is only asserted for G reductive; I don't know if it's true for a general algebraic group G). But Hodge-Tate reps are stable under tensor products, duals, and subobjects, so we're done. Note that this argument also shows that there is a sensible notion of what it means for $\Gamma_K \to G(\mathbf{Q}_p)$ to be Hodge-Tate: it just means that for a faithful representation to GL_n (equivalently, all faithful reps, equivalently all reps), the induced Galois rep is (are) Hodge-Tate.

In particular, for any representation $G \to \mathrm{GL}(W)$ we get a map $\mu : (\mathrm{GL}_1)_C \to \mathrm{GL}(W_C)$.

Hence, for any representation $X : G \to \mathrm{GL}(W)$ of G and any C-algebra A, we have, for $a \in A^\times$, we have an endomorphism of $W \otimes_{\mathbf{Q}_p} A$ corresponding to a, which on the ith graded piece is multiplication by a^i. These endomorphisms satisfy some very natural properties: recall $D(W_1) \otimes_K D(W_2) = D(W_1 \otimes_{\mathbf{Q}_p} W_2)$ (this is in Fontaine Periodes p-adiques p124) and, crucially, if $\alpha : W_1 \to W_2$ is G-equivariant then $D(\alpha) : D(W_1) \to D(W_2)$ exists and is grading-preserving (D is a functor to graded vector spaces). The upshot of all of this is that, for a fixed A and a, we have an A-valued automorphism of the fibre (forgetful) functor on the category of representations of G, and hence a map $A^\times \to G(A)$ by one of the standard theorems about Tannakian categories. This is a morphism of functors and turns into a map $(\mathrm{GL}_1)_C \to G_C$ which is of course just μ again.

Of course we don't have "numbers" in this general setting, only μs. So far we have shown that for a Hodge-Tate representation $\Gamma_K \to G(\mathbf{Q}_p)$ with G reductive, we get $\mu : (\mathrm{GL}_1)_C \to G_C$.

6 More general coefficient fields.

Now say we have $\Gamma_K \to G(E)$ with G now reductive over a finite extension E of $\mathbf{Q}_p$. Of course we just follow our noses; we set $H = \mathrm{Res}_{E/\mathbf{Q}_p} G$, so now we have a representation to $H(\mathbf{Q}_p)$. A faithful E-representation of G gives a faithful $\mathbf{Q}_p$-representation of H, so there's our notion of Hodge-Tate. So we get $\mu : (\mathrm{GL}_1)_C \to H_C$ and, because C is so big, we see $H_C = \prod_{\tau:E\to C} G_{\tau,C}$, so we get a $\mu_t au : (\mathrm{GL}_1)_C \to G_{\tau,C}$ for each $\tau : E \to C$.

Let's just pause here and think what τ is again. To give τ is to give a K-algebra map $E \otimes_{\mathbf{Q}_p} K \to C$ and hence it's to give a component L_i of $E \otimes_{\mathbf{Q}_p} K$ and a K-algebra map $L_i \to C$.

7 Our ultimate goal: general G and coefficient field $\overline{\mathbf{Q}}_p$.

We now put ourselves in the position that Toby and I find ourselves in: we have a representation $\Gamma_K \to G(\overline{\mathbf{Q}}_p)$, with G reductive over $\overline{\mathbf{Q}}_p$ (and hence with finite component group; we'll have to reduce to this case but this is fine) and, for the moment at least, K a general Fontaine field. The first thing to do is to choose a totally arbitrary model G_0 for G over a field E_0; the Baire Category argument shows that we can now find E/E_0 large enough such that that Γ_K is taking values in $G_0(E)$. We now (from the previous section) have a $\mu : (\mathrm{GL}_1)_C \to (G_0)_{\tau:E\to C}$ for every $\tau : E \to C$. To make this notion intrinsic we consider instead maps $\tau : \overline{\mathbf{Q}}_p \to C$; to every such map we get an induced map $E \to \overline{\mathbf{Q}}_p \to C$, and the base change of G_0 to $\overline{\mathbf{Q}}_p$ is G again, so for every such τ we get $\mu : (\mathrm{GL}_1)_C \to G_{\tau:E\to C}$.

Finally we have to check that this is well-defined, so we have to worry about a completely different model G_0' and a completely different E_0' and so on. But both G_0 and G_0' will become isomorphic after some huge field extension containing both E and E', so it suffices to check that the recipe attaching μ to τ is unchanged if we replace E above by a bigger field E'. We have a representation to $G_0(E)$ and also to $G_0(E')$. By the recipe of the previous section, if $\tau' : E' \to C$ restricts to $\tau : E \to C$ then we'll get a mu and a mu' and we need to check they're the same. But in fact we can reduce to the vector space case! Just choose a faithful representation of G_0 defined over E; this induces a faithful representation over E' and we already checked here that the two μs coincided after this faithful representation; hence they coincide.

8 The case $K/\mathbf{Q}_p$ finite and G an L-group.

Given any Fontaine field K and a representation $\Gamma_K \to G(\overline{\mathbf{Q}}_p)$ with G reductive, there's a good notion of being Hodge-Tate. If the representation is Hodge-Tate then for each $\tau : \overline{\mathbf{Q}}_p \to C$ we get $\mu : (\mathrm{GL}_1)_C \to G_{\tau,C}$. The image of μ will land in the connected component of G_C and will hence give us a cocharacter of the dual torus (recall G is an L-group) and hence an element of $X^*(T)/W$, where T is a torus over $\overline{\mathbf{Q}}$ in our conn red group, and W is the "absolute" Weyl group.

So we're getting an element of $X^*(T)/W$ for each $\tau : \overline{\mathbf{Q}}_p \to C$, and in practice there will be a coefficient field E_0 such that this element will only depend on the restriction of τ to E_0. I do wonder whether E_0 will be the coefficient field of the automorphic form, because if there is *any* form of G defined over E_0 and which is recieving the Galois representation then μ_τ will only depend on $\tau|E_0$.

Now if $K/\mathbf{Q}_p$ is finite then τ induces an isomorphism $\overline{\mathbf{Q}}_p = \overline{K}$ and so to give τ is to give a map $\sigma : \overline{K} \to \overline{\mathbf{Q}}_p$. Furthermore if L/K is a finite extension of K containing the normal closure of $E_0/\mathbf{Q}_p$ (that is, of all L_i/K with L_i components of $E_0 \otimes_{\mathbf{Q}_p} K$) then μ_σ will only depend on $\sigma|L$. This is some sort of "continuity" (indeed—even a local constancy) of the association $\sigma \mapsto \mu_\sigma$. But to actually get the cocharacter one does need more than $K \to \overline{\mathbf{Q}}_p$; to see what one needs in practice one should take "the smallest field E such that ρ can be realised over E" and then take the Galois closure of K and of E and put it all together; that's what one has to embed into $\overline{\mathbf{Q}}_p$. If you don't embed enough then you only get "Hodge-Tate numbers".