Lecture 3: Galois deformation rings

Mok
October 16, 2009
Notes by Sam Lichtenstein

Let G be a profinite group and $\overline{\rho} : G \to \mathrm{GL}_n(k)$ a representation defined over a finite field k of characteristic p. Let Λ be a complete discrete valuation ring with residue field k, e.g. $\Lambda = W(k)$. Let $\mathcal{C}_\Lambda$ be the category of artinian local Λ-algebras with residue field k, and local morphisms. Let $\widehat{\mathcal{C}}_\Lambda$ be the category of complete Noetherian local Λ-algebras with residue field k, i.e. the pro-category of $\mathcal{C}_\Lambda$.

1. Deformation functors

Define $\mathrm{Def}(\overline{\rho}) : \widehat{\mathcal{C}}_\Lambda \to \mathrm{Sets}$ by

$$\mathrm{Def}(\overline{\rho})(A) = \{(\rho, M, \iota)\}/\sim$$

where M is a free A-module of rank n, $\rho : G \to \mathrm{GL}_A(M)$ is a continuous representation, $\iota : \rho \otimes_A k \cong \overline{\rho}$ is an isomorphism, and two such triples are equivalent when the representations are isomorphic in a manner which respects the ι's. Define the **framed** deformation functor $\mathrm{Def}^\square(\overline{\rho})$ by

$$\mathrm{Def}^\square(\overline{\rho})(A) = \{(\rho, M, \iota, \beta)/\sim$$

where β is a basis for M lifting the standard basis for k^n under ι. Morally, $\mathrm{Def}^\square$ is the set of liftings of $\overline{\rho}$ into $\mathrm{GL}_n(A)$.

There is a forgetful functor $\mathrm{Def}^\square \to \mathrm{Def}$.

Equivalent definitions are

$$\mathrm{Def}^\square(\overline{\rho})(A) = \{\rho : G \to \mathrm{GL}_n A \,|\, \rho \bmod \mathfrak{m}_A = \overline{\rho}\},$$

$$\mathrm{Def}(\overline{\rho})(A) = \mathrm{Def}^\square(\overline{\rho})(A)/(\text{conjugation by } \Gamma_n(A) := \ker(\mathrm{GL}_n(A) \to \mathrm{GL}_n(k))).$$

Note: it is easy to see that $\mathrm{Def}^\square(\overline{\rho})(A) = \varprojlim_i \mathrm{Def}^\square(\overline{\rho})(A/\mathfrak{m}_A^i)$. It is also true (but requires an argument) that $\mathrm{Def}(\overline{\rho})(A) = \varprojlim_i \mathrm{Def}(\overline{\rho})(A/\mathfrak{m}_A^i)$. In other words "we can compute these functors on the level of artinian quotients", so we just need to consider them on the category $\mathcal{C}_\Lambda$.

2. p-finiteness

We cannot hope to represent $\mathrm{Def}(\overline{\rho})$ or $\mathrm{Def}^\square(\overline{\rho})$ in $\widehat{\mathcal{C}}_\Lambda$ (which only contains Noetherian rings) unless G is "not too big".

Definition. We say G **satisfies the p-finiteness condition** if for every open subgroup $H \subset G$ of finite index, there are only finitely many continuous group homomorphisms $H \to \mathbb{Z}/p\mathbb{Z}$ (i.e., only finitely many open subgroups of index p). (This holds if and only if for any such H, the maximal pro-p quotient of H is topologically finitely generated.)

We are interested in two cases.

(1) $G = G_K$ for a local field K finite over $\mathbb{Q}_\ell$ (allowing $\ell = p$!).
(2) $G = G_{K,S}$ for a number field K and S a finite set of ramified primes.

In case (1), $H = G_{K'}$ for a finite extension K'/K, and the p-finiteness condition follows from the fact that the local field K' of characteristic 0 has only finitely many extensions of any given degree (such as degree p). For (2), H corresponds to some finite extension K'/K unramified outside of S, so the index-p open subgroups of H correspond to certain degree-p extensions of K' unramified away from the places of K' over S. Thus, the p-finiteness follows from the **Hermite-Minkowski theorem**, which says that only finitely many extensions of K of bounded degree unramified outside S.

Returning to the general situation, assume G satisfies p-finiteness. By *Schlessinger's criterion*, we will eventually see that $\mathrm{Def}^\square(\overline{\rho})$ is always representable in $\widehat{\mathcal{C}}_\Lambda$, so there exists a universal framed deformation ring $R^\square_{\overline{\rho}} \in \widehat{\mathcal{C}}_\Lambda$ and a universal framed deformation $\rho^\square_{\overline{\rho}}$ satisfying the natural universality property. We will also see that $\mathrm{Def}(\overline{\rho})$ is itself representable by a universal deformation ring $(R_{\overline{\rho}}, \rho_{\mathrm{univ}})$, at least when $\mathrm{End}_G(\overline{\rho}) = k$. This will be the case if $\overline{\rho}$ is absolutely irreducible, and also if $n = 2$ and $\overline{\rho}$ is a non-split extension of distinct characters.

3. Zariski tangent space to the deformation functors

Let $k[\epsilon]$ denote the ring of dual numbers of k. The **tangent space** to a functor $F : \widehat{\mathcal{C}}_\Lambda \to Sets$ is $F(k[\epsilon]) =: t_F$. Initially this is just a set; the hypotheses of Schlessinger's criterion give it a natural structure of k-vector space (compatibly with natural transformations in F).

Let $V \in \mathrm{Def}(\overline{\rho})(k[\epsilon]) = t_{\mathrm{Def}(\overline{\rho})}$. Then by definition there is given a specified isomorphism $V/\epsilon V \cong \overline{\rho}$, so we obtain an exact sequence

$$0 \to \epsilon V \to V \to \overline{\rho} \to 0.$$

But it is easy to see that ϵV is naturally $k[G]$-isomorphic to $\overline{\rho}$ as well. Hence we see

$$t_{\mathrm{Def}(\overline{\rho})} = \mathrm{Ext}^1_{k[G]}(\overline{\rho}, \overline{\rho}) = \mathrm{H}^1(G, \mathrm{Ad}(\overline{\rho}));$$

this respects the k-linear structure on both sides.

More explicitly, given $\rho \in \mathrm{Def}^\square(\overline{\rho})(k[\epsilon])$ we can write $\rho(g) = \overline{\rho}(g) + \epsilon\Phi(g)\overline{\rho}(g)$ for $\Phi(g) \in \mathrm{Ad}(\overline{\rho})$. One can compute that the condition that ρ is a group homomorphism is the 1-cocycle condition on Φ. So $t_{\mathrm{Def}^\square(\overline{\rho})} = Z^1(G, \mathrm{Ad}(\overline{\rho}))$. Similarly one checks that two framed deformations are conjugate under $\Gamma_0(k[\epsilon]) = I_n + \epsilon M_n(k)$ if and only if their associated cocycles differ by a 1-coboundary. We conclude that $t_{\mathrm{Def}(\overline{\rho})} = \mathrm{H}^1(G, \mathrm{Ad}(\overline{\rho}))$, and

$$\dim_k B^1(G, \mathrm{Ad}(\overline{\rho})) = \dim \mathrm{Ad}(\overline{\rho}) - \dim \mathrm{H}^0(G, \mathrm{Ad}(\overline{\rho}))$$

is the **number of framed variables**. The p-finiteness hypothesis says precisely that $\dim Z^1, \dim \mathrm{H}^1 < \infty$. If moreover $\mathrm{End}_G(\overline{\rho}) = k$ then $h^0(G, \mathrm{Ad}(\overline{\rho})) = 1$, and we are in the representable situation. The forgetful functor $\mathrm{Def}^\square(\overline{\rho}) \to \mathrm{Def}(\overline{\rho})$ induces a map $R_{\overline{\rho}} \to R^\square_{\overline{\rho}}$, which turns out to be formally smooth, and thus realizes $R^\square_{\overline{\rho}}$ as a ring of formal power series (in some number d of variables) over $R_{\overline{\rho}}$. The number d is precisely the number of framed variables, which in this case is $n^2 - 1$.

Concretely, what is going on is that if $\overline{\rho}$ has only scalar endomorphism (so likewise for any lifting of $\overline{\rho}$) and we consider the universal deformation $R_{\overline{\rho}}$ then to "universally" specify a basis which residually lifts the identity is precisely to applying conjugation by a residually trivial matrix which is unique up to a unit scaling factor. And we can eliminate the unit scaling ambiguity by demanding (as we always may in a unique way) that the upper left matrix entry is not merely a unit but is equal to 1. Thus, the framing amounts to specifying a "point" of the formal $R_{\overline{\rho}}$-group of PGL_n at the identity, which thereby proves the asserted description of the universal framed deformation ring in these cases as a formal power series ring over $R_{\overline{\rho}}$ in $n^2 - 1$ variables. To be explicit, over

$$R^\square(\overline{\rho}) = R(\overline{\rho})[\![Y_{i,j}]\!]_{1 \le i,j \le n, (i,j) \ne (1,1)}$$

the universal framed deformation is the lifting $\overline{\rho}_{\mathrm{univ}}$ equipped with the basis obtained from the standard one by applying the invertible matrix $1_n + (Y_{i,j})$ where $Y_{1,1} := 0$.

It must be stressed that we will later need to work with cases in which $\overline{\rho}$ is *trivial* (of dimension 2), so $R_{\overline{\rho}}$ does not generally exist. This is why the framed deformation ring is useful.

4. References

- Mazur's articles in "Galois groups over $\mathbb{Q}$" and "Modular Forms and Fermat's Last Theorem".
- Kisin's notes from CMI summer school in Hawaii.

5. More on Zariski tangent spaces to deformation functors

From now on fix G to be either G_K for local K or $G_{K,S}$ for a number field K. Fix $\overline{\rho} : G \to \mathrm{GL}_n(k)$ and suppose the characteristic of the finite field k is p. If F is a deformation functor represented by $R \in \widehat{\mathcal{C}}_\Lambda$, recall that

$$F(A) = \mathrm{Hom}_{\Lambda\text{-alg}}(R, A), \qquad t_F = F(k[\epsilon]) = \mathrm{Hom}_{\Lambda\text{-alg}}(R, k[\epsilon]) = \mathrm{Hom}_{\Lambda\text{-alg}}(R/(\mathfrak{m}_R^2 + \mathfrak{m}_\Lambda R), k[\epsilon]).$$

The last equality is because the Λ-algebra maps are *local* morphisms, so in particular they send $\mathfrak{m}_R$ to $\epsilon k[\epsilon]$, and hence $\mathfrak{m}_R^2$ to zero. But by general nonsense we have

$$R/(\mathfrak{m}_R^2 + \mathfrak{m}_\Lambda R) = k \oplus \frac{\mathfrak{m}_R}{\mathfrak{m}_R^2 + \mathfrak{m}_\Lambda R},$$

where the second summand is square zero. Thus we see

$$t_F = \mathrm{Hom}_k\left(\frac{\mathfrak{m}_R}{\mathfrak{m}_R^2 + \mathfrak{m}_\Lambda R}, k\right) = t_R^*,$$

where for $A \in \widehat{\mathcal{C}}_\Lambda$ we define the **reduced Zariski cotangent space of** A to be

$$t_A^* = \frac{\mathfrak{m}_A}{\mathfrak{m}_A^2 + \mathfrak{m}_\Lambda R}.$$

Exercise. Fix a map $A \xrightarrow{f} B$ in $\widehat{\mathcal{C}}_\Lambda$. Then f is surjective if and only if $t_f^* : t_A^* \to t_B^*$ is surjective. [Use completeness... it's a Nakayamal's lemma sort of thing.]

A corollary of the Exercise is that if $d = \dim_k t_F = \dim_k t_R^*$ then we can pick a k-basis $x_1, \ldots, x_d$ of t_R^*, lift it to a collection $\widetilde{x}_i \in \mathfrak{m}_R \subset R$, and then the map $\Lambda[\![X_1, \ldots, X_d]\!] \to R$ sending X_i to $\widetilde{x}_i$ will be surjective. *A priori* bounds for the number of generators in the kernel (and hence on the dimension of R) can be obtained by estimating certain H^2s in the cohomology of G, which will be discussed later. These dimension bounds are sometimes useful, but usually not strong enough to give good control on R.

6. Examples

A local case. Let $K/\mathbb{Q}_\ell$ be local with $\ell \neq p$ and $G = G_K$. Let $\overline{\rho}$ be the trivial representation of dimension n. Then in particular $\mathrm{End}_G \overline{\rho} \supsetneq k$, so only the framed deformation functor is representable. In this case we can actually construct $R^\square(\overline{\rho})$ by hand. If $\rho : G \to \mathrm{GL}_n A$ is a deformation of the trivial representation $\overline{\rho}$, then G lands in the kernel $\Gamma_n A \subset \mathrm{GL}_n A$. Now $\Gamma_n A = I_n + M_n(\mathfrak{m}_A)$, explicitly, which is a pro-p group isomorphism to the additive group $M_n(A)$. In particular ρ factors throug h the maximal pro-p quotient of G.

In particular $\rho|_{I_K}$ factors through the p-part of the tame quotient $I_K^{\mathrm{tame}} = I_K / I_K^{\mathrm{wild}}$ of the inertia I_K of K. The picture to keep in mind is the tower of field extensions

$$K \hookrightarrow K^{\mathrm{unr}} \hookrightarrow K^{\mathrm{tame}} \hookrightarrow \overline{K}.$$

Now from the structure of local fields we know that the p-part of I_K^{tame} is

$$I_K^{\mathrm{tame},(p)} = \mathbb{Z}_p(1).$$

Here the twist means that if $\sigma \in I_K^{\mathrm{tame},(p)}$ then $\mathrm{Frob}_K \sigma \mathrm{Frob}_K^{-1} = \sigma^q$ where $q = \ell^r = \#(\mathcal{O}_K/\mathfrak{m}_K)$. Fix a lift $f \in G$ of Frob_K and τ a topological generator of $I_K^{\mathrm{tame},(p)}$. What we can conclude is that a lift ρ to any A is specified by the images of f and τ, subject to the relation

$$\rho(f)\rho(\tau) = \rho(\tau)^q \rho(f).$$

So we can take

$$R^\square(\overline{\rho}) = \Lambda[\![\{f_{ij}, \tau_{ij}\}_{1 \le i,j \le n}]\!]/I$$

where the ideal of relations I is generated by the ones given by the matrix equations

$$[I_n + (f_{ij})][I_n + (\tau_{ij})] = [I_n + (\tau_{ij})]^q [I_n + (f_{ij})].$$

A global case. For a global case we'll consider *characters* of $G = G_{K,S}$. Note that we have a wonderful fact in this case. The Teichmüller lift $[\cdot] : k \to W(k)$ is a multiplicative section of $W(k) \to k$. This allows us to twist any character $\overline{\rho}$ by the Teichmüller lift $[\overline{\rho}^{-1}]$ of its reciprocal. to conclude that $R(\overline{\rho}) = R(\mathbf{1})$ where $\mathbf{1} : G \to k^\times$ is the trivial character. In other words, the universal deformation of a character $\overline{\rho}$ is just a twist of the universal deformation of the trivial character (using the same coefficient ring).

Arguing just like in the local case, it follows that any lift ρ to A of the trivial mod p character $\overline{\rho}$, must factor through the maximal pro-p quotient $G_{K,S}^{\mathrm{ab},(p)}$ of the abelianization of $G_{K,S}$.

Let us specialize now to the case $K = \mathbb{Q}$ [the case of a general number field is similar, but requires class field theory]. Assume $p \in S$, since otherwise stuff is boring. By the Kronecker-Weber theorem we know that $G_{\mathbb{Q},S}^{\mathrm{ab}} = \prod_{\ell \in S} \mathbb{Z}_\ell^\times$, which implies that the maximal pro-p quotient is

$$G_{\mathbb{Q},S}^{\mathrm{ab},(p)} = \prod_{\ell \in S, \ell \equiv 1 (p)} (\mathbb{F}_\ell^\times)^{(p)} \times (1 + p\mathbb{Z}_p).$$

So we can, in this case, simply take $R = \Lambda[\![G^{\mathrm{ab},(p)}_{\mathbb{Q},S}]\!]$ to be the formal group algebra over Λ. From the description of $G^{\mathrm{ab},(p)}_{\mathbb{Q},S}$ we can be very explicit:

$$R = \frac{\Lambda[\![\{X_\ell\}_{\ell\in S,\ell\equiv 1(p)}, T]\!]}{(\{((X_\ell+1)^{p^{\mathrm{ord}_p(\ell-1)}}-1)\}_{\ell\in S,\ell\equiv 1(p)})}$$

In particular if $S=\{p,\infty\}$ then $R\cong\Lambda[\![T]\!]$.

For a general number field K this relates to the *Leopoldt conjecture* which says that $\mathrm{rk}_{\mathbb{Z}_p}(G^{\mathrm{ab},(p)}_{K,S}) = 1+r_2$, where r_2 is the number of conjugate pairs of complex embeddings of K.

7. Local and Global

We can relate the two examples from the last subsection in the following manner, which will be extremely important later in one of Kisin's key improvements of Wiles' method. Let $G=G_{K,S}$, $\overline{\rho}: G_{K,S}\to \mathrm{GL}_n(k)$ a fixed residual representations, and Σ a finite set of primes. For each $v\in\Sigma$ we have

$$\overline{\rho}|_{G_v} : G_{K_v} = G_v \hookrightarrow G_K \twoheadrightarrow G_{K,S} \xrightarrow{\overline{\rho}} \mathrm{GL}_n(k).$$

We have local framed deformation rings $R_v^\square := R^\square(\overline{\rho}|_{G_v})$. Define a variation of the global framed deformation functor by

$$\mathrm{Def}^{\square,\Sigma}(\overline{\rho})(A) = \{(\rho_A, \{\beta_v\}_{v\in\Sigma})\}/\sim;$$

here, ρ_A is a deformation of $\overline{\rho}$ and β_v is a basis for $\rho_A|_{G_v}$ which reduces to the standard basis for $\overline{\rho}$. Then in fact $\mathrm{Def}^{\square,\Sigma}(\overline{\rho})$ is also representable, by a ring $R^{\square,\Sigma}_{K,S}$. For each $v\in\Sigma$ we have a forgetful map

$$\mathrm{Def}^{\square,\Sigma}(\overline{\rho}) \to \mathrm{Def}(\overline{\rho}|_{G_v})$$

and hence on the revel of representing objects, an algebra

$$R_v^\square \to R^{\square,\Sigma}_{K,S}.$$

In concrete terms, this is saying that if we form the universal deformation of $\overline{\rho}$ equipped with a framing along Σ and then forget the framing away from v and restrict to G_v, the resulting framed deformation of $\overline{\rho}|_{G_v}$ with coefficients in $R^{\square,\Sigma}_{K,S}$ is uniquely obtained by specializing the universal framed deformation of $\overline{\rho}|_{G_v}$ along a unique local Λ-algebra homomorphism $R_v^\square \to R^{\square,\Sigma}_{K,S}$.

Hence, by the universal property of completed tensor products (to be discussed in Samit's talk rather generally) we get an important map

$$\widehat{\bigotimes}_\Lambda R_v^\square \to R^{\square,\Sigma}_{K,S}$$

in $\widehat{\mathcal{C}}_\Lambda$. (Note that we have to take the completion of the algebraic tensor product, which is not itself a complete ring. For example, $\Lambda[\![x]\!]\otimes_\Lambda\Lambda[\![y]\!]$ is a gigantic non-noetherian ring, but the corresponding completed tensor product is $\Lambda[\![x,y]\!]$.) This is a rather interesting extra algebra structure on the global framed deformation ring, much richer than its mere Λ-algebra structure; of course, this all has perfectly good analogues without the framings, assuming that $\overline{\rho}$ and its local restrictions at each $v\in\Sigma$ have only scalar endomorphisms.

This idea of viewing a global deformation ring as an algebra over a (completed) tensor product of local deformation rings is the key to Kisin's method for "patching" deformation rings in settings going far beyond the original Taylor-Wiles method (where only the Λ-algebra structure was used).