Connections and known cases

Avi Zeff

Last time, we stated the p-adic Langlands conjectures, in terms of the existence and properties of a certain fully faithful functor $\mathfrak{A} : \operatorname{Ind} D^b_{\text{f.p.}}(\text{sm.}\, G) \hookrightarrow \operatorname{Ind}\operatorname{Coh}(\mathcal{X}_d)$, and made a series of remarks giving further properties of it. Today's talk could be viewed as a continuation of this series of remarks, but they will be somewhat more substantial ones: first, we'll see how this conjecture implies the geometric Breuil–Mézard conjecture. Then we'll look at some known cases and some recent progress: the GL_1 case is reasonably simple, but already the $\mathrm{GL}_2(\mathbb{Q}_p)$ case, though (more or less) known, is complicated, and we'll only vaguely sketch the theory. For higher extensions of $\mathbb{Q}_p$, the theory is even more complicated; we'll mention a bit of how some of this goes, although the general theory is still unknown. Finally, we'll look at the right adjoint $\mathfrak{B}$ of $\mathfrak{A}$, giving approximately the other direction of the correspondence.

1. The geometric Breuil–Mézard conjecture

This conjecture stems from the following fact: let k_F be the residue field of our p-adic field F, and $k = k_E = \mathcal{O}/\pi$ (we maintain the same notation as last time, so $\mathcal{O} = \mathcal{O}_E$). Then the irreducible $\bar{k} = \overline{\mathbb{F}_p}$-representations of $\mathrm{GL}_d(k_F)$ have a simple classification: via highest weight vectors they are in bijection with Serre weights, i.e. tuples $\underline{m} = (m_{\bar{\sigma},i})$ for $\bar{\sigma} : k_F \hookrightarrow \bar{k}$ and $1 \leq i \leq d$ satisfying certain restrictions, unimportant for us. If $\underline{m}$ is a Serre weight, we write $F_{\underline{m}}$ for the corresponding $\bar{k}$-representation of $\mathrm{GL}_d(k_F)$.

Recall that to a regular Hodge type $\underline{\lambda}$ and an inertial type τ we can associate representations $V_{\underline{\lambda}}$ and $\sigma^{\text{crys},\circ}(\tau)$ of $K = \mathrm{GL}_d(\mathcal{O}_F)$ over $\mathcal{O}$. To the pair $(\underline{\lambda}, \tau)$, it is then natural to associate their tensor product; because of the above observation, we're interested in k-representations, not $\mathcal{O}$-representations, so we simply reduce modulo π_E and call (the semisimplification of) the resulting representation $\overline{\sigma}^{\text{crys}}(\underline{\lambda}, \tau) = (V_{\underline{\lambda}} \otimes_{\mathcal{O}} \sigma^{\text{crys},\circ}(\tau) \otimes_{\mathcal{O}} k)^{\text{ss}}$, which is now a representation of $\mathrm{GL}_d(k_F)$. Therefore we can write it as a direct sum of irreducibles, which are classified by Serre weights. In other words, there exist integer $n^{\text{crys}}_{\underline{m}}(\underline{\lambda}, \tau)$ such that

$$\overline{\sigma}^{\text{crys}}(\underline{\lambda}, \tau) = \bigoplus_{\underline{m}} F_{\underline{m}}^{\bigoplus n^{\text{crys}}_{\underline{m}}(\underline{\lambda},\tau)}.$$

Now, we expect the data of $\underline{\lambda}$ and τ to pick out a component of $\mathcal{X}_d$, namely $\mathcal{X}_d^{\text{crys},\underline{\lambda},\tau}$. In particular if we take the underlying reduced stack (which now is entirely in characteristic p) we get a cycle $Z_{\text{crys},\underline{\lambda},\tau}$ on $\mathcal{X}_{d,\text{red}}$, i.e. an element of the free abelian group on its irreducible components. For purposes such as automorphy lifting, we would like to be able to geometrize the above statement, i.e. interpret the $F_{\underline{m}}$ similarly as cycles on $\mathcal{X}_{d,\text{red}}$ to form the following.

Conjecture (Geometric Breuil–Mézard conjecture)**.** *There exist cycles $Z_{\underline{m}}$ on $\mathcal{X}_{d,\text{red}}$ such that for every regular Hodge type $\underline{\lambda}$ and inertial type τ we have*

$$\overline{\sigma}^{crys}(\underline{\lambda}, \tau) = \sum_{\underline{m}} n^{crys}_{\underline{m}}(\underline{\lambda}, \tau) \cdot Z_{\underline{m}}.$$

Now, assume our main conjecture from last time. We can take the representation $V_{\underline{\lambda}} \otimes_{\mathcal{O}} \sigma^{\mathrm{crys},\circ}(\tau)$ of $K = \mathrm{GL}_d(\mathcal{O}_F)$ and take compact induction up to $G = \mathrm{GL}_d(F)$ to get a smooth G-representation; applying $\mathfrak{A}$ gives a derived ind-coherent sheaf $\mathcal{F}_{\underline{\lambda},\tau}$ on $\mathcal{X}_d$, which we saw last time is actually a coherent sheaf concentrated in degree 0. Similarly from $F_{\underline{m}}$, the action of $\mathrm{GL}_d(k_F)$ gives an action of $\mathrm{GL}_d(\mathcal{O}_F)$ and so we can similarly induce up to G and apply $\mathfrak{A}$ to get a complex $\mathcal{F}_{\underline{m}} = \mathfrak{A}(c\text{-}\mathrm{Ind}_K^G F_{\underline{m}})$, which again turns out to be coherent concentrated in degree 0 without too much more work, supported on the special fiber of $\mathcal{X}_d$.

Now, by the definitions and the functoriality in $\mathfrak{A}$ we have in $K_0(\mathrm{Coh}(\mathcal{X}_d))$ the equality

$$[\mathcal{F}_{\underline{\lambda},\tau} \otimes_{\mathcal{O}} k] = \sum_{\underline{m}} n_{\underline{m}}^{\mathrm{crys}}(\underline{\lambda}, \tau) \cdot [\mathcal{F}_{\underline{m}}].$$

Taking supports and using the maximality condition we skipped lightly over last time gives cycles $Z_{\underline{m}}$ as the support of $\mathcal{F}_{\underline{m}}$ satisfying the hypotheses of the conjecture, so the geometric Breuil–Mézard conjecture is a consequence of our p-adic Langlands conjecture.

2. GL_1

We expect the local Langlands program for GL_1 to be essentially local class field theory, and indeed this will be true in this case. Here, we can even describe the image explicitly to get an equivalence of derived categories, which comes from an equivalence of abelian categories and so the statement will be true even without deriving.

By local class field theory, we have an isomorphism $W_F^{\mathrm{ab}} \xrightarrow{\sim} F^\times = \mathrm{GL}_1(F)$, and one can work out that in this case $\mathcal{X}_1$ is genuinely the stack of one-dimensional continuous representations of W_F, or equivalently of $W_F^{\mathrm{ab}} \xrightarrow{\sim} \mathrm{GL}_1(F)$. In particular one can describe $\mathcal{X}_1$ explicitly as

$$\mathrm{Spf}\, \mathcal{O}[[F^\times]]/\widehat{\mathbb{G}_{\mathrm{m}}}$$

where $\widehat{\mathbb{G}_{\mathrm{m}}}$ is the p-completed multiplicative group acting trivially on $\mathrm{Spf}\, \mathcal{O}[[F^\times]]$. Therefore sheaves on $\mathcal{X}_1$ with trivial $\widehat{\mathbb{G}_{\mathrm{m}}}$-action are equivalent to $\mathcal{O}[[F^\times]]$-modules. Recalling that the latter are (a generalization of) an incarnation of smooth $\mathrm{GL}_1(F)$-representations, this gives an explicit version of $\mathfrak{A}$, sending a smooth $\mathrm{GL}_1(F)$-representation to the corresponding sheaf on $\mathcal{X}_1$.

Another way to formulate this is to look at the structure sheaf $\mathcal{O}_{\mathcal{X}_1}$: via this description, it naturally carries an $\mathcal{O}[[F^\times]]$-action, so given another $\mathcal{O}[[F^\times]]$-module π our functor is just given by

$$\pi \mapsto \mathcal{O}_{\mathcal{X}_1} \otimes_{\mathcal{O}[[F^\times]]} \pi,$$

i.e. $L_\infty = \mathcal{O}_{\mathcal{X}_1}$.

3. $\mathrm{GL}_2(\mathbb{Q}_p)$

There is work in progress of Dotto, Emerton, and Gee to prove the conjecture for $\mathrm{GL}_2(\mathbb{Q}_p)$, for which a less categorical formulation of p-adic Langlands is already known (and is essentially the only known case beyond GL_1). The strategy is to explicitly construct L_∞.

In this case, $\mathcal{X}_2$ classifies rank 2 (φ, Γ)-modules, and in particular id : $\mathcal{X}_2 \to \mathcal{X}_2$ gives a universal rank 2 (φ, Γ)-module D over $\mathcal{X}_2$. For any morphism $\operatorname{Spec} A \to \mathcal{X}_2$, we write D_A for the pullback of D to A (twisted by the cyclotomic character for technical reasons). Then (recall David's talk) D_A is a rank 2 projective module over

$$\mathbb{A}_A = A((T))$$

with commuting A-linear semilinear actions of φ and Γ by $\varphi(1+T) = (1+T)^p$ and $\gamma(1+T) = (1+T)^{\epsilon(\gamma)}$ for $\gamma \in \Gamma = \operatorname{Gal}(\mathbb{Q}_p(\zeta_{p^\infty})/\mathbb{Q}_p)$ and ϵ the cyclotomic character. The action of φ is injective and so has a left inverse ψ, which one can work out is an A-linear surjection commuting with the Γ-action and satisfying

$$\psi(\varphi(a)m) = a\psi(m), \qquad \psi(a\varphi(m)) = \psi(a)m$$

for all $a \in \mathbb{A}_A$ and $m \in D_A$.

Using the actions of φ, ψ, and Γ on D_A, we will explicitly build a $\mathrm{GL}_2(\mathbb{Q}_p)$-representation. Our first attempt will be a representation called $D \boxtimes \mathbb{P}^1$; this will not be quite right, so we build a subrepresentation which is the right thing for L_∞.

To get $D \boxtimes \mathbb{P}^1$, we'll build $D \boxtimes \mathbb{Z}_p$, viewing $\mathbb{Z}_p$ as the affine line, and then glue two copies along $D \boxtimes \mathbb{Z}_p^\times$ and $z \mapsto z^{-1}$. We start by giving an action of the monoid

$$P^+ = \begin{pmatrix} \mathbb{Z}_p \setminus \{0\} & \mathbb{Z}_p \\ 0 & 1 \end{pmatrix}$$

on D_A by

$$\begin{pmatrix} p^m a & b \\ 0 & 1 \end{pmatrix} = (1+T)^b \varphi^m \circ a$$

for $a \in \mathbb{Z}_p^\times$, $m \geq 0$, $b \in \mathbb{Z}_p$ where we view a as an element of Γ via the standard isomorphism. There is also a natural action of P^+ on $\mathbb{Z}_p$ by

$$\begin{pmatrix} \mathbb{Z}_p \setminus \{0\} & \mathbb{Z}_p \\ 0 & 1 \end{pmatrix} \cdot z = p^m a z + b,$$

and both actions extend to all of $\mathrm{GL}_2(\mathbb{Z}_p)$ (they're somewhat vague about this, but I guess it's via Bruhat decomposition or similar?) to get a representation $D_A \boxtimes \mathbb{Z}_p$ of $\mathrm{GL}_2(\mathbb{Z}_p)$. One can check that the subrepresentation $D_A \boxtimes \mathbb{Z}_p^\times$ is the locus on which $\psi = 0$. Gluing along this locus and the map $z \mapsto z^{-1}$ allows us to extend to a $\mathrm{GL}_2(\mathbb{Z}_p)$-action on the resulting representation $D_A \boxtimes \mathbb{P}^1$. This is functorial in A and so lifts to a representation $D \boxtimes \mathbb{P}^1$.

However, this turns out not to be the right choice for L_∞: it realizes not the local Langlands correspondence but some extension by a dual. Therefore we want to pick out some subrepresentation by choosing a lattice: set $\mathbb{A}_A^+ = A[[T]]$, which again has actions of φ, ψ, and Γ. Then viewing D_A as an $\mathbb{A}_A^+$-module, it turns out to contain a minimal ψ-stable $\mathbb{A}_A^+$-lattice $D_A^\natural$ (i.e. it is finitely generated over $\mathbb{A}_A^+$ and spans D_A over $\mathbb{A}_A$) compatibly with flat base change. Taking the product with $\mathbb{P}^1$ as above turns out (very non-obviously, using facts from the classical p-adic local Langlands correspondence for $\mathrm{GL}_2(\mathbb{Q}_p)$) to give a $\mathrm{GL}_2(\mathbb{Q}_p)$-stable lattice $D_A^\natural \boxtimes \mathbb{P}^1$. Again this gives a universal lattice $D^\natural \boxtimes \mathbb{P}^1$ on $\mathcal{X}_2$, which turns out to be the right candidate for L_∞.

It is also possible to describe the geometry of $\mathcal{X}_2$ fairly explicitly: in particular its reduced closed points look like a chain of copies of $\mathbb{P}^1$. We will not get into this.

4. $\mathrm{GL}_2(\mathbb{Q}_{p^r})$

We next consider the case where we replace $\mathbb{Q}_p$ by an unramified extension $\mathbb{Q}_{p^r}$, focusing on the simplest case $\mathbb{Q}_{p^2}$. Already this is much more complicated. One example demonstrating this is as follows. Recall that for a Serre weight $\underline{m}$ we got an irreducible representation $F_{\underline{m}}$ of $\mathrm{GL}_2(\mathbb{F}_{p^2})$, and thus of $K = \mathrm{GL}_2(\mathbb{Z}_{p^2})$. When we take compact induction up to $G = \mathrm{GL}_2(\mathbb{Q}_{p^2})$ and apply $\mathfrak{A}$, we get an explicit sheaf $\mathcal{F}_{\underline{m}}$. In the $\mathbb{Q}_p$ case, we saw that this sheaf was concentrated in degree 0; further, its support is a point, i.e. it is a skyscraper sheaf at the corresponding L-parameter. This is the geometric incarnation of a particularly nice case of the local Langlands program, just with p-adic coefficients.

In the $\mathbb{Q}_{p^2}$ case, things are much more complicated: we still have $\mathcal{F}_{\underline{m}}$, but now the support is one-dimensional, and so representation has infinite length. In particular, unlike the $\mathbb{Q}_p$ case *or* the case $\ell \neq p$, we can't think of a smooth G-representation as an actual sheaf; the derived setup is very much necessary.

Nevertheless, under some restrictions on the G-representations there is a result in progress (by many people, including Michael Harris) giving a functor which is expected to be the restriction of $\mathfrak{A}$ to suitably nice representations. This is again done by an explicit construction which is too complicated for this talk.

5. The adjoint functor $\mathfrak{B}$

The last thing I want to talk about is the other direction of the correspondence. Since $\mathfrak{A} : \mathrm{Ind}\, D^b_{\text{f.p.}}(\mathrm{sm.}\, G) \to \mathrm{Ind}\,\mathrm{Coh}(\mathcal{X}_d)$ is a cocontinuous functor between suitably nice ∞-categories, by the adjoint functor theorem it admits a right adjoint $\mathfrak{B} : \mathrm{Ind}\,\mathrm{Coh}(\mathcal{X}_d) \to \mathrm{Ind}\, D^b_{\text{f.p.}}(\mathrm{sm.}\, G)$. If we expect that our conjecture should have an upgrading to an equivalence of categories, replacing $\mathrm{Ind}\, D^b_{\text{f.p.}}(\mathrm{sm.}\, G)$ by some sort of derived category $D(\mathrm{Bun}_G)$ (suitably decorated), with $i : [*/G(F)] \hookrightarrow \mathrm{Bun}_G$ giving the relevant stratum, then we expect that our functor $\mathfrak{A}$ should be the composition of this equivalence of categories with $i_! : D(\mathrm{sm.}\, G) \to D(\mathrm{Bun}_G)$, and $\mathfrak{B}$ should be the composition of the inverse of the equivalence with $i^! : D(\mathrm{Bun}_G) \to D(\mathrm{sm.}\, G)$.

We can try to understand this functor using what we know about $\mathfrak{A}$, and in particular we hope to be able to use it to understand the reverse direction of the local Langlands correspondence: given a suitable Galois representation $\mathrm{Gal}_F \to \mathrm{GL}_d(\mathcal{O})$, we can assign it to a skyscraper sheaf on $\mathcal{X}_d$ and would like to be able to say what representation of $\mathrm{GL}_d(F)$ this corresponds to.

Let L_∞ be the sheaf corresponding to $\mathfrak{A}$ by the conjecture. We have $\mathfrak{A}(\pi) = L_\infty \otimes^{\mathbb{L}}_{\mathcal{O}[[G]]} \pi$, so we can give an equally explicit description of $\mathfrak{B}$: if $\mathcal{F}$ is a derived ind-coherent sheaf on $\mathcal{X}_d$, then

$$\mathfrak{B}(\mathcal{F}) = R\,\mathrm{Hom}_{\mathcal{O}_{\mathcal{X}_d}}(L_\infty, \mathcal{F}).$$

Recall the connection to Taylor–Wiles patching: we expect that if R_∞ is the power series ring in the Taylor–Wiles variables over the universal deformation ring of a point $x \in \mathcal{X}_d(\overline{\mathbb{F}_p})$ corresponding to $\overline{\rho} : \mathrm{Gal}_F \to \mathrm{GL}_d(\overline{\mathbb{F}_p})$, then the universal lift corresponds to a morphism $f : \mathrm{Spf}\, R_\infty \to \mathcal{X}_d$, and we expect that the patched module M_∞ is given by $f^* L_\infty$. Generically, we expect that R_∞ is formally smooth, and M_∞ is pro-flat (even pro-free) over R_∞. Thus we should be able to compute $\mathfrak{B}$ as

$$\begin{aligned}\mathfrak{B}(\mathcal{F}) &= R\operatorname{Hom}_{\mathcal{O}_{\mathcal{X}_d}}(L_\infty, \mathcal{F}) \\ &= R\Gamma(R\mathscr{H}\!om_{\mathcal{O}_X}(L_\infty, \mathcal{F})) \\ &= R\Gamma(\mathcal{X}_d, \mathscr{H}\!om_{\mathcal{O}_{\mathcal{X}_d}}(L_\infty, \mathcal{F})).\end{aligned}$$

If $x : \operatorname{Spec} \overline{\mathbb{F}_p} \to \mathcal{X}_d$ is a closed point corresponding to $\overline{\rho}$ semisimple and non-scalar and $G_x = \operatorname{Aut}_{\mathrm{Gal}_F}(\overline{\rho})$, then x gives rise to a closed embedding $i : \operatorname{Spec} \overline{\mathbb{F}_p}/G_x \hookrightarrow \mathcal{X}_d$, and we can form the skyscraper sheaf $\delta_x = i_* \mathcal{O}_{\operatorname{Spec} \overline{\mathbb{F}_p}}$. If $f : \operatorname{Spf} R_\infty \to \mathcal{X}_d$ is the versal morphism at x coming from Taylor–Wiles patching, then

$$\mathscr{H}\!om_{\mathcal{O}_{\mathcal{X}_d}}(L_\infty, \delta_x) = Ri_* \mathscr{H}\!om_{\mathcal{O}_{\operatorname{Spec} \overline{\mathbb{F}_p}/G_x}}(i^* L_\infty, \mathcal{O}_{\operatorname{Spec} \overline{\mathbb{F}_p}}),$$

which is just the $\mathfrak{m}$-torsion in $M_\infty^\vee$ (as dual to $M_\infty/\mathfrak{m}_\infty$), i.e. $M_\infty^\vee[\mathfrak{m}]$. Thus $\mathfrak{B}(\delta_x) = M_\infty^\vee[\mathfrak{m}]^{G_x}$.

We can now compute some examples. Take $d = 2$ and $F = \mathbb{Q}_p$, and suppose $\overline{\rho} = \overline{\chi}_1 \oplus \overline{\chi}_2$ for $\overline{\chi}_1 \neq \overline{\chi}_2$. Then $G_x = \mathbb{G}_{\mathrm{m}}$ and $M_\infty^\vee[\mathfrak{m}]$ is the direct sum $\pi_1 \oplus \pi_2$ of two irreducible principal series representations, with

$$\mathscr{H}\!om_{\mathcal{O}_{\mathcal{X}_d}}(L_\infty, \delta_x) = \pi_1(1) \oplus \pi_2(-1)$$

(in some order) with the twists denoting the weights of the $\mathbb{G}_{\mathrm{m}}$-action. Thus $\mathfrak{B}(\delta_x(1)) = \pi_2$ and $\mathfrak{B}(\delta_x(-1)) = \pi_1$.

If $\overline{\rho}$ is irreducible, then $G_x = \{\pm 1\}$ and $M_\infty^\vee[\mathfrak{m}]$ is an irreducible supersingular representation, with

$$\mathscr{H}\!om_{\mathcal{O}_{\mathcal{X}_d}}(L_\infty, \delta_x) = \pi(1)$$

with the twist denoting a nontrivial action of $\{\pm 1\}$. Therefore $\mathfrak{B}(\delta_x(1)) = \mathfrak{B}(\delta_x(-1)) = \pi$. One can also do similar examples for $\mathbb{Q}_{p^2}$.

More generally, we can think of the representation π associated to $\overline{\rho}$ (in suitably nice contexts) as the unique irreducible representation π such that the support of $H^0(\mathfrak{A}(\pi))$ contains the point $x \in \mathcal{X}_d$ corresponding to $\overline{\rho}$. (For more general groups G, we would no longer expect there to be a unique such representation.)

References

[1] Matthew Emerton, Toby Gee, and Eugen Hellmann. An introduction to the categorical p-adic Langlands program. *arXiv preprint arXiv:2210.01404*, 2022.