p-Adic families of automorphic forms

2026.09.04 Y-Rant VIII

$\S\ 1$. Introduction and Motivation

Last year at Y-RANT VII in Nottingham, I spoke about the local geometry of the Coleman--Mazur eigencurve. Today I want to step back and discuss one of the key inputs in the eigenvariety machine -- p-adic families of automorphic forms. More precisely, I will use the following two papers as our guide:

- [Pil13] Pilloni, Vincent. *Overconvergent modular forms*. Ann. Inst. Fourier (Grenoble) 63 (2013), no. 1, 219–239.
- [AIP15] Andreatta, Fabrizio; Iovita, Adrian; Pilloni, Vincent. _p-Adic families of Siegel modular cuspforms._ Ann. of Math. (2) 181 (2015), no. 2, 623–697.

If time permits, I will then talk about the associated eigenvariety.

Fear not if you are not familiar with these terms - the main motivation/take away is the following:

- we would like to interpolate automorphic forms p-adically into families as their weights vary (because, for instance, the constant term of Eisenstein families give rise to an immediate construction of certain p-adic L-functions; Hida and Coleman families have found important applications in arithmetic geometry and the Langlands program);
- once we have constructed well-behaved spaces of p-adic families automorphic forms (together with a compact U_p operator), the eigenvariety machine then produces for us an associated "eigenvariety" - a rigid analytic space which one may heuristically regard as a "moduli space" of p-adic systems of Hecke eigenvalues.

Before getting started, please bare in mind that I am far from an expert in this fascinating topic. If you would like to learn more about it, I highly recommend you to talk to Alice Pozzi or Charlotte Clare-Hunt.

$\S\ 2$. Toy Case: Elliptic Modular Forms

Let me begin with a slightly different way of thinking of modular forms. Katz (1970's) taught us that a modular form of weight k is an element $f \in \mathrm{H}^0\left(X, \omega^k\right)$ where

- X denotes the modular curve (that is, the moduli space of elliptic curves + certain level structures), and
- ω is the automorphic line bundle on X

or equivalently, it is a rule that assigns to a pair $(E/S, \eta)$ consisting of

- an elliptic curve E over a scheme S, and
- nowhere vanishing invariant differential η on E

a value $f(E/S, \eta) \in \mathcal{O}_S(S)$, such that

- $f(E/S, \lambda\eta) = \lambda^{-k} f(E/S, \eta)$ for all $\lambda \in \mathcal{O}_S(S)^\times$, and that
- the association is functorial in S.

Over $S = \mathbb{C}$, $E = \mathbb{C}/\Lambda$ is a complex torus and there is a canonical invariant differential $\eta = dz$ on E. Some computations recover our familiar classical notion of modular forms.

N.B. I am being vague about the base scheme and level because they are not super relevant for us.

Question: How to p-adically interpolate the space of modular forms of weight k, that is, how to allow κ to be an arbitrary p-adic character $\mathbb{Z}_p^\times \to \mathbb{C}_p^\times$?

The first step towards p-adic interpolation is to switch to a more presentation-theoretic point of view: Why don't we look at the moduli space of pairs $(E/S, \eta)$ instead? This would lead us to the frame bundles over X,

- $T := \operatorname{Hom}(\mathcal{O}_X, \omega) \to X$;
- $T^\times := \operatorname{Isom}(\mathcal{O}_X, \omega) \xrightarrow[\pi]{\mathbb{G}_{\mathrm{m}}} X$.

The S-points of $\mathcal{T}^\times$ are exactly given by (equivalent classes of) all the pairs $(E/S, \eta)$ we saw above.

By doing so, f can be thought of as a function

$$T^\times(S) \to \mathcal{O}_S(S) = \Gamma\left(S, \mathcal{O}_S\right) = \mathbb{A}^1(S)$$

which is functorial in S, or equivalently

$$f \in \mathrm{H}^0\left(T^\times, \mathcal{O}_{T^\times}\right) = \mathrm{H}^0\left(X, \pi_*\mathcal{O}_{T^\times}\right).$$

The homogeneity property of weight k becomes κ'-equivariance under the $\mathbb{G}_\mathrm{m}$-action, namely

$$\mathrm{H}^0\left(X, \omega^k\right) = \mathrm{H}^0\left(X, \pi_*\mathcal{O}_{T^\times}\left[\kappa'\right]\right),$$

where $\kappa\colon x \mapsto x^k$ is an algebraic character of $\mathbb{G}_\mathrm{m}$ and $\kappa'\colon x \mapsto x^{-k}$ its inverse.

Now, for κ any p-adic character, it is tempting to simply apply the same formula on the rigid-analytic frame bundle. Namely, take the rigid analytic spaces $\mathcal{X}$ (resp. $\mathcal{T}^\times := T_{\mathrm{rig}} \setminus \{0\}$) associated to X (resp. $T^\times$) and put

$$\omega^\kappa := \mathcal{O}_{\mathcal{T}^\times}\left[\kappa'\right].$$

N.B. One may ask why we didn't take $\left(T^\times\right)_{\mathrm{rig}}$. In fact, for classical weights it doesn't really matter, but $T_{\mathrm{rig}} \setminus \{0\}$ will be more convenient for us later.

However, this naïve ω^κ admits no global sections unless κ is algebraic. Indeed, the fibers are given by

$$\{\varphi\colon B(0,1) \setminus \{0\} \to \mathbb{C}_p \text{ analytic} : \varphi(\lambda z) = \kappa^{-1}(\lambda)\varphi(z), \forall\, \lambda \in \mathbb{Z}_p^\times\}.$$

Writing $\varphi(z) = \sum\limits_{n\in\mathbb{Z}} a_n z^n$, the homogeneity condition then forces $\kappa(\lambda) = \lambda^k$ for some $k \in \mathbb{Z}$. Also, κ is only locally analytic, namely, it extends analytically on to sufficiently small neighborhoods of $\mathbb{Z}_p^\times$.
The solution therefore is to not force κ onto every frame, but retain only those frames which are close to a distinguished p-adic family of frames.

Slogan: The Hodge--Tate map selects the right frames

- $\mathcal{X}(v) := \{E \in \mathcal{X}\colon \mathrm{Hdg}(E) \leq v\}$, $0 \leq v \leq 1$, strict neighborhood of the ordinary locus
  - the condition $\mathrm{Hdg}(E) \leq v$ means that the elliptic curves E is "not too supersingular"

    - $\mathcal{X}(0) =$ the ordinary locus; $\mathcal{X}(1) =$ the entire modular curve
- (Fargues, ...) there exists $H_n \subset E[p^n]$ "the canonical subgroup of order n" if v is small relative to n, whose Cartier dual $H_n^\vee$ is "étale-like"
- $\mathrm{HT}\colon H_n^\vee \to \omega_{H_n}$ the Hodge--Tate map, defined as $x \mapsto x^*(dt/t)$

We thus have a diagram

$$\left(H_n^\vee\right)^\times \xrightarrow{\mathrm{HT}} \omega_{H_n} \xleftarrow{\mathrm{res}} \omega_E.$$

And Pilloni defined the desired smaller frame bundle $\mathcal{F}_n^\times$ by pulling back along these two arrows:

$$\mathcal{F}_n^\times = \left\{(E, x, \omega)\colon E \in \mathcal{X}(v), x \in \left(H_n^\vee\right)^\times, \eta \in \omega_E \text{ s.t. } \mathrm{HT}(x) = \omega|_{\omega_{H_n}}\right\}.$$

The equation $\mathrm{HT}(x) = \omega|_{\omega_{H_n}}$ roughly means that ω is congruent to $\mathrm{HT}(x)$ modulo p^w for some w depending on v and n. Hence, as x runs through generators of $H_n^\vee$, one sees that $\mathcal{F}_n^\times$ is a union of small discs inside the frame bundle $\mathcal{T}^\times$

Remark: Taking $v = 0$ and passing to $n \to \infty$, one recovers the ordinary Igusa tower. As a result, $\mathcal{F}_n^\times$ can be viewed as a partial extension of the Igusa tower to overconvergent loci.

For n sufficiently large, we define the overconvergent modular sheaf of weight κ to be

$$\omega^{\dagger,\kappa} := \pi_* \mathcal{O}_{\mathcal{F}_n^\times}\left[\kappa'\right],$$

and the space of overconvergent modular forms of weight κ to be

$$M_{\kappa'}^\dagger := \lim_{v \to 0^+} \mathrm{H}^0\left(X(v), \omega^{\dagger,\kappa}\right).$$

Note that $\omega^{\dagger,,\kappa}$ is independent of the auxiliary level n. Indeed, whenever two consecutive levels are admissible, restriction from $\mathcal{F}_{n-1}^\times$ to $\mathcal{F}_n^\times$ is injective by analytic continuation and surjective because every κ'-homogeneous analytic function extends uniquely to the larger discs.

It turns out that [Pil13, Prop. 3.6] for any classical weight $\kappa\colon x \mapsto x^k$, $\omega^k|_{\mathcal{X}(v)} \xrightarrow{\sim} \omega^{\dagger,\kappa}$.

$\S\,3$. General Case: $G = \mathrm{GSp}_g$ for $g \geq 2$

Rough idea:

- the automorphic bundle ω has rank g
- fibers of the classical modular sheaf ω^κ of weight $\kappa \in X^+(T)$ (T denotes the torus in G here) is given by the algebraic induction $V_\kappa = \mathrm{Ind}_B^G(\kappa')$ of κ' from the Borel B to G
- the canonical subgroup H_n of level n has rank p^{ng}
    - its dual $H_n^\vee$ is locally trivialized by $(\mathbb{Z}/p^n\mathbb{Z})^g$ and so admits a standard basis & filtration/flag of length g
    - this standard flag + basis is sent by the Hodge--Tate map HT to a flag + basis in ω_{H_n}
- AIP considered the space of $\mathcal{IW}$ over $\mathcal{X}(v)$ consisting of flags + bases which coincide with the standard flag + basis under HT
    - this again gives a congruence condition modulo p^w for some w depending on v and n

Write $\pi\colon \mathcal{IW} \to \mathcal{X}(v)$, and define

$$\omega^{\dagger,\,\kappa} := \pi_* \mathcal{O}_{\mathcal{IW}}\left[\kappa'\right].$$

One can show that fibers of this sheaf can be described as the w-analytic induction $V_{\kappa'}$ of κ'. To summarize, for any $\kappa \in X^+(T)$ there is a canonical restriction map

$$\omega^\kappa|_{\mathcal{X}(v)} \hookrightarrow \omega^{\dagger,\,\kappa},$$

which is (étale) locally isomorphic to the inclusion

$$V_{\kappa'} \hookrightarrow V_{\kappa'}^{w-\mathrm{an}}$$

of the algebraic induction into the analytic induction.

Warning: For $g \geq 2$, $\omega^{\dagger,\,\kappa}$ is only a Banach sheaf because of the analytic induction process. This degenerates when $g = 1$ simply because $G = B = T = \mathrm{GL}_1$ in that case.

$\S\,4$. Families and the Cuspidal Eigenvariety

- $\mathcal{W} = \mathrm{Hom}\left(T\left(\mathbb{Z}_p\right), \mathbb{C}_p^\times\right)$ the weight space, whose connected components are g-dimensional open polydiscs
- $\mathcal{U} \subset \mathcal{W}$ affinoid open
- $\kappa^{\mathrm{univ}}: T\left(\mathbb{Z}_p\right) \to \mathcal{O}_U(\mathcal{U})^\times$ the universal character

Since the universal character κ^{univ} is locally analytic, one may substitute it into the preceding construction and obtain a Banach sheaf

$$\omega^{\dagger,\,\kappa^{\mathrm{univ}}}$$

over $\mathcal{X}(v) \times \mathcal{U}$, whose specialization at each fiber $\kappa \in \mathcal{U}$ recovers $\omega_w^{\dagger,\kappa}$, and thus define the space of families of overconvergent forms (resp. cuspforms) as

- $M_{\mathcal{U}}^{\dagger} := \mathrm{H}^0\left(\mathcal{X}_{\mathrm{Iw}(p)}(v) \times \mathcal{U}, \omega_w^{\dagger,\,\kappa^{\mathrm{univ}}}\right)$, and
- $M_{\mathcal{U}}^{\dagger,\,\mathrm{cusp}} := \mathrm{H}^0\left(\mathcal{X}_{\mathrm{Iw}(p)}(v) \times \mathcal{U}, \omega_w^{\dagger,\,\kappa^{\mathrm{univ}}}(-D)\right)$, where D denotes the boundary of $\mathcal{X}(v)$

We restrict ourselves to cuspforms because

- the ordinary locus in modular curves is affinoid, whereas when $g \geq 2$ this is not the case;
- for modular curves, the classical modular sheaves are interpolated by coherent sheaves, whereas when $g \geq 2$, the sheaves $\omega_w^{\dagger,\kappa}$ are only Banach sheaves.

Given these facts, it is highly non-trivial to prove that

1. $M_{\mathcal{U}}^{\dagger,\,\mathrm{cusp}}$ is a projective Banach $\mathcal{O}(\mathcal{U})$-module;
2. for every $\kappa \in \mathcal{U}$, the specialization map

   $$M_{\mathcal{U}}^{\dagger,\,\mathrm{cusp}} \to \mathrm{H}^0\left(\mathcal{X}(v), \omega^{\dagger,\,\kappa}(-D)\right)$$

   is surjective.

And cuspidality assumption is expected to be necessary when $g \geq 2$. 1) & 2) guarantee that $M_{\mathcal{U}}^{\dagger,\,\mathrm{cusp}}$ are nice enough to be fed into the eigenvariety machine. Denote the resulting eigenvarieity by $w: \mathcal{E} \to \mathcal{W}$.

What has the geometric construction bought us? AIP were able to prove

- $\mathcal{E}$ is equidimensional of dimension g -- which is not known for Emerton's approach to eigenvariety via completed cohomology
- overconvergent cuspidal eigenforms of finite slope can be deformed over all the g weight directions
- $\cdots$