Properties of the Eigencurve

Gabor Wiese

27th January 2010

Abstract

These are notes for my talk in the seminar 'The Eigencurve' at the Universität Duisburg-Essen on 28 January 2010. The aim of the seminar is to go through the seminal paper by Coleman and Mazur first defining the eigencurve. The present notes try to cover most parts of Chapter 7. The philosophy was to completely disregard technicalities and to try to stress the underlying ideas.

The reader be warned that I learned the subject only while preparing this talk (part time during less than 2 weeks), so misunderstandings and, hence, wrong or imprecise statements are very likely.

1 Introductory words

In Johan's talk, the eigencurve was constructed. In this talk, we recover the nilreduction of the eigencurve by patching together the affinoid spaces attached to certain Hecke algebras.

In short:

(1) At the basis of the eigencurve is the theorem (Hida, Gouvêa) that to an overconvergent eigenform one can attach a pseudo representation. This allows one to see overconvergent eigenforms inside $X_p \times \mathbb{A}^1$, where X_p is the rigid analytic space attached to a universal deformation ring of pseudo-representations.

The eigencurve C_p is **by construction** the subspace of $X_p \times \mathbb{A}^1$ whose $\mathbb{C}_p$-points are the images of the overconvergent eigenforms under the Hida-Gouvêa construction.

Like this, the geometry is hidden. It is even not a priori clear that the eigencurve is a curve.

(2) Today we will construct a curve D (it is obviously a curve) by patching together rigid analytic spaces associated with certain Hecke algebras. It will turn out that D is the nilreduction of C_p.

I am speculating that the French translation 'courbe de Hecke' for eigencurve refers to this construction.

2 Collection of definitions and statements

2.1 Standing assumptions and notation

- p an odd prime number.
- $\mathbb{C}_p$ the completion of a fixed algebraic closure $\overline{\mathbb{Q}}_p$ of $\mathbb{Q}_p$.

2.2 Notions from rigid analytic geometry

Definition 2.1 *(i) $B[0,p^m]_K = \mathrm{Max}(\mathcal{O}_K\langle p^m T\rangle)$ is the closed disc of radius p^m at the origin. It is the rigid analytic subspace of $\mathbb{A}^1_K$ whose $\mathbb{C}_p$-valued points are $\{x \in \mathbb{C}_p \mid |x| \leq p^m\}$.*

(ii) If A is a complete noetherian local ring, then X_A is the attached rigid analytic space over K. The $\mathbb{C}_p$-valued points are the continuous $\mathcal{O}_K$-algebra homomorphisms $A \to \mathbb{C}_p$, i.e. $X_A(\mathbb{C}_p) = \mathrm{Hom}_{\mathcal{O}_K}(A, \mathbb{C}_p)$.

(iii) The spectral semi-norm *(or* supremum norm*):*

$$|f| := \max_{x \in \mathrm{Max}(A)} |f(x)|,$$

where $f(x) := f + \mathfrak{m}_x \in A/\mathfrak{m}_x \hookrightarrow \overline{K}$. Note that we write $\mathfrak{m}_x$ for x to remind us that we are talking about maximal ideals. Note also that $|f(x)|$ is independent of the embedding $A/\mathfrak{m}_x \hookrightarrow \overline{K}$.

(iv) If Y is a rigid analytic variety over K, then $A(Y)$ is the K-algebra of rigid analytic functions. $A^0(Y)$ is the $\mathcal{O}_K$-subalgebra of $A(Y)$ consisting of rigid analytic functions with spectral semi-norm ≤ 1.

(v) A rigid space X over a complete non-archimedean valued field K is called nested *if it admits an admissible covering by a countable increasing sequence of affinoid subdomains*

$$X_0 \subseteq X_1 \subseteq \cdots \subseteq X.$$

Definition 2.2 (Power series) *Let A be a topological ring.*

(i) $A\langle T\rangle := \{f = \sum_{n=0}^{\infty} a_n T^n \mid a_n \xrightarrow{n\to\infty} 0\}$, the Tate algebra of restricted or strictly convergent power series on the closed unit ball.

(ii) Assume A is a local ring with maximal ideal $\mathfrak{m}$. Then

$$A\{\{T\}\} := \{f = \sum_{n=0}^{\infty} a_n T^n \mid \exists c_n \in \mathbb{R} \text{ s.t. } c_n/n \xrightarrow{n\to\infty} \infty \text{ and } a_n \in \mathfrak{m}^{c_n}\ \forall\ n\}$$

are the entire power series over A. They are precisely those series having an infinite radius of convergence.

2.3 Fredholm

Definition 2.3 *(i) Let A be a local ring. A* Fredholm series over A *is an element $F \in A\{\{T\}\}$ with constant term* 1.

(ii) Let X be a rigid analytic space. A Fredholm series over X *is an entire power series over $A(X)$ with constant term* 1.

For more, see Stefan's talk.

2.4 Weight space

For simplicity, I only list the tame level 1 weight space.

Definition 2.4 *The* Iwasawa algebra Λ *is defined as*

$$\Lambda := \mathbb{Z}_p[[\mathbb{Z}_p^\times]] := \varprojlim_n \mathbb{Z}_p[(\mathbb{Z}/p^n\mathbb{Z})^\times].$$

The weight space of tame level 1 *is defined as the rigid analytic space over $\mathbb{Q}_p$ $\mathcal{W} := \mathcal{W}_1$ associated with Λ.*

The Iwasawa algebra Λ is a complete regular local ring of Krull dimension 2. Since Λ can also be equipped with the structure of a Hopf algebra, the weight space $\mathcal{W}$ is a commutative rigid analytic group object. Note

$$\mathcal{W}(\mathbb{C}_p) = \mathrm{Hom}_{\mathrm{cts}}(\mathbb{Z}_p^\times, \mathbb{C}_p) = \ \mathbb{C}_p\text{-values characters of } \mathbb{Z}_p^\times.$$

We mention that the p-cyclotomic character defines a surjection $\chi_p : \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \to \mathbb{Z}_p^\times$ and by the Kronecker-Weber theorem an isomorphism $G^{\mathrm{ab}}_{\mathbb{Q},\{p,\infty\}} \cong \mathbb{Z}_p^\times$.

For more, see Panos' talk.

2.5 Overconvergent modular forms

See Ralf's talk.

2.6 Pseudo representations

Let $\overline{\rho}_1, \ldots, \overline{\rho}_r$ be the distinct p-modular Galois representations $G_{\mathbb{Q},\{p,\infty\}} \to \mathrm{GL}_2(\overline{\mathbb{F}}_p)$ and let $R_{\overline{\rho}_i}$ be the universal deformation rings attached to the corresponding pseudo representations. Put $R_p = R_{\overline{\rho}_1} \times \cdots \times R_{\overline{\rho}_r}$ and let X_p be the associated rigid analytic space over $\mathbb{Q}_p$.

We recall the rigid analytic map $\mu_{\mathrm{wt}} : X_p \to \mathcal{W}$ defined in Juan's talk. This map comes from a continuous ring homomorphism $\mu_{\mathrm{wt}} : \Lambda \to R_p$, equipping R_p with a Λ-algebra structure. We now explain this notion. Let $G := G_{\mathbb{Q},\{p,\infty\}}$ and let $\rho : G \to \mathrm{GL}_2(R)$ be any continuous 2-dimensional (pseudo-) representation for any topological ring R. Let $\det(\rho)$ be the determinant of ρ. There is a map $\mu_{\det}$ making the following diagram commutative:

$$\begin{array}{ccc} G & \xrightarrow{\det(\rho)} & R^\times \\ \downarrow & & \uparrow{\scriptstyle \mu_{\det(\rho)}} \\ G^{\mathrm{ab}} & \xrightarrow[\chi_p]{\sim} & \mathbb{Z}_p^\times . \end{array}$$

Recall that $\det(\rho_f) = \chi_p^{k-1}$ for a classical modular form of weight k (and trivial nebentype). We find that $\mu_{\det(\rho_f)}$ sends $\gamma \in \mathbb{Z}$ to $\gamma^{k-1} \in R$. In order to recover k (instead of $k-1$), one defines μ_{wt} (for given ρ) as

$$\mu_{\mathrm{wt}} : \mathbb{Z}_p^\times \to R_p^\times, \quad \gamma \mapsto \gamma \cdot \mu_{\det(\rho)}(\gamma).$$

This continuous group homomorphism extends uniquely to the desired

$$\mu_{\mathrm{wt}} : \Lambda = \mathbb{Z}_p[[\mathbb{Z}_p^\times]] \to R_p,$$

hence, giving rise to

$$\mu_{\mathrm{wt}} : X_p \to \mathcal{W}.$$

The meaning of μ_{wt} is just that it sends a pseudo-representation $\lambda \in X_p(\mathbb{C}_p)$ to its weight character $\mu_{\mathrm{wt}}(\lambda) =: \kappa_\lambda \in \mathcal{W}(\mathbb{C}_p)$.

For more, see Juan's talk.

3 The Eigencurve

In this section, we will recall the description of the eigencurve on $\mathbb{C}_p$-points. In fact, in the whole talk we are mostly going to treat $\mathbb{C}_p$-points and we shall not so much worry about the structures as rigid analytic spaces.

Recall that

- $\mathcal{H}' := \Lambda[T_\ell : \ell \neq p \text{ prime }] = \Lambda[T_2, T_3, \ldots, \hat{T_p}, \ldots]$, the *restricted abstract Hecke algebra*.
- $\mathcal{H} := \mathcal{H}'[U_p]$, the *full abstract Hecke algebra*.
- $\iota : \mathcal{H}' \to R_p$ is the Λ-algebra homomorphism defined by sending T_ℓ to the trace of the universal pseudo-reprsentation at Frob_ℓ.
- For $\alpha \in \mathcal{H}'$ such that $\iota(\alpha) \in R_p^\times$ we have the Fredholm power series $P_{\iota(\alpha)U_p}(T) \in \Lambda\{\{T\}\}$ which is the unique power series satisfying the following assertion: For every weight (character) $\kappa : \Lambda \to \mathbb{C}_p$ (i.e. every point of the weight space $\mathcal{W}$), the power series obtained by applying κ to the coefficients is the Fredholm determinant

$$P^\kappa_{\iota(\alpha)U_p}(T) = \det(1 - (\iota(\alpha)U_p)T | M^\dagger_\kappa(\mathbb{C}_p)) = \prod_i (1 - u_i T)^{e_i} \in \mathbb{C}_p\{\{T\}\}.$$

The u_i are precisely the eigenvalues of the operator $\iota(\alpha)U_p$ on the space $M^\dagger_\kappa(\mathbb{C}_p))$ of overconvergent modular forms of weight κ (and tame level 1, as always).

- The spectral curve $\mathcal{Z}_\alpha := \mathcal{Z}_{\iota(\alpha)U_p}$ is the rigid analytic subspace of $\mathcal{W} \times \mathbb{A}^1$ cut out by $P_{\iota(\alpha)U_p}$. We have the following description on the $\mathbb{C}_p$-points:

$$\begin{aligned}\mathcal{Z}_\alpha(\mathbb{C}_p) &= \{(\kappa, u^{-1}) \in (\mathcal{W} \times \mathbb{A}^1)(\mathbb{C}_p) \mid P^\kappa_{\iota(\alpha)U_p}(u^{-1}) = 0\} \\ &= \{(\kappa, u^{-1}) \in (\mathcal{W} \times \mathbb{A}^1)(\mathbb{C}_p) \mid \exists f \in M^\dagger_\kappa(\mathbb{C}_p) : (\iota(\alpha))U_p f = uf\}.\end{aligned}$$

  As I am taking inverses, I should treat zero explicitly. I just ignore this.

- One defines a rigid analytic map $r_\alpha : X_p \times \mathbb{A}^1 \to \mathcal{W} \times \mathbb{A}^1$ which can explicitly be given on $\mathbb{C}_p$-points like this:

$$(\lambda, t) \mapsto (\kappa_\lambda, \frac{t}{\lambda(\iota(\alpha))}).$$

  Note that $\lambda(\iota(\alpha))$ is just the evaluation of the pseudo representation $\lambda : R_p \to \mathbb{C}_p$ at the image of the Hecke operator α. If $\alpha = T_\ell$, then this is just λ evaluated at the trace of Frob_ℓ under the universal pseudo-representation.

- We have furthermore:

$$\begin{aligned}r_\alpha^{-1}(\mathcal{Z}_\alpha)(\mathbb{C}_p) &= \{(\lambda, u^{-1}) \in (X_p \times \mathbb{A}^1)(\mathbb{C}_p) \mid (\kappa_\lambda, (\lambda(\iota(\alpha))u)^{-1}) \in \mathcal{Z}_\alpha(\mathbb{C}_p)\} \\ &= \{(\lambda, u^{-1}) \in (X_p \times \mathbb{A}^1)(\mathbb{C}_p) \mid \exists f \in M^\dagger_{\kappa_\lambda}(\mathbb{C}_p) : (\iota(\alpha))U_p f = \lambda(\iota(\alpha))uf\}.\end{aligned}$$

- The eigencurve C_p (of tame level 1) is by definition the rigid analytic space

$$C_p = \bigcap_{\alpha \in \mathcal{H}' \,:\, \iota(\alpha) \in R_p^\times} r_\alpha^{-1}(\mathcal{Z}_\alpha).$$

  For $(\lambda, u^{-1}) \in (X_p \times \mathbb{A}^1)(\mathbb{C}_p)$ the following statements are equivalent:

  (i) $(\lambda, u^{-1}) \in C_p(\mathbb{C}_p)$.

  (ii) For all $\alpha \in \mathcal{H}'$ such that $\iota(\alpha) \in R_p^\times$ there exists a normalised overconvergent Hecke eigenform $f \in M^\dagger_{\kappa_\lambda}(\mathbb{C}_p)$ satisfying

$$(\iota(\alpha)U_p)f = \lambda(\iota(\alpha))uf.$$

- By a theorem of Gouvêa and Hida, one can attach a pseudo representation λ_f to any normalised overconvergent Hecke eigenform f. This provides one with a map

$$\{\text{ normalised overconv. eigenforms of finite slope over } \mathbb{C}_p \} \xrightarrow{f \mapsto (\lambda_f, u_f^{-1})} (X_p \times \mathbb{A}^1)(\mathbb{C}_p),$$

  where u_f is the U_p-eigenvalue of f.

- The main theorem from Johan's talk is that this map is a bijection onto $C_p(\mathbb{C}_p)$.

  This is not a miracle at all. The construction of C_p as the intersection of the $r_\alpha^{-1}(\mathcal{Z}_\alpha)$ was made precisely to have this property.

  The idea is that due to the flexibility given by the choice of α, all eigenvalues at all T_ℓ can be encoded as eigenvalues of $\iota(\alpha)U_p$. This forces the following assertion, which is the key point in the proof presented in Johan's talk: Given $(\lambda, u^{-1}) \in C_p(\mathbb{C}_p)$. For all n, there is some overconvergent modular form f_n of weight κ_λ such that its q-expansion coincides with the formal q-expansion of λ up to n and such that its U_p-eigenvalue is u.

4 Construction of a reduced eigencurve

In this section, we shall construct a curve D which will be isomorphic to the nilreduction of C_p. Conceptually, there are the following differences with C_p:

- We do not use single weights, but we work locally over the weight space.
- The curve D will be patched together from pieces which are locally finite over the weight space. They are hence curves, so that also D is a curve.

Here is the plan for the construction of D. First take any $\alpha \in \mathcal{H}'$ such that $\iota(\alpha) \in R_p^\times$. We will construct 1-dimensional pieces $D_\alpha(V)$. This will take up most of the time. Then the curve D_α will be patched together from the $D_\alpha(V)$. Finally, we will set $D := D_1$, state its properties and relate it to C_p.

4.1 Construction of the 1-dimensional piece $D_\alpha(V)$

We start by constructing $D_\alpha(V)$ and, hence, fix α as above for the time being. Consider the morphism:

$$\pi_\alpha : \mathcal{Z}_\alpha \xrightarrow{\text{natural inclusion}} \mathcal{W} \times \mathbb{A}^1 \xrightarrow{\text{first projection}} \mathcal{W}.$$

Fact: If $V \subset \mathcal{Z}_\alpha$ is an affinoid subdomain, then $\pi_\alpha(V) \subset \mathcal{W}$ is also an affinoid subdomain.

Let

$$\mathcal{C}_\alpha := \{V \subset \mathcal{Z}_\alpha \text{ affinoid subdomain} \mid V \xrightarrow{\pi_\alpha} \pi_\alpha(V) \text{ is finite }, V \text{ admissibly closed-open in } \mathcal{Z}_\alpha\}$$

and

$$\mathcal{C}_\alpha^{\text{irr}} := \{V \in \mathcal{C}_\alpha \mid V \text{ irreducible }\}.$$

Fact: $\mathcal{C}_\alpha$ is an admissible covering of $\mathcal{Z}_\alpha$.

Fact: Fix an affinoid $Y \subset \mathcal{W}$ with ring of rigid analytic functions $A := A(Y)$. Then the set

$$\{V \in \mathcal{C}_\alpha^{\text{irr}} \mid \pi_\alpha(V) = Y\}$$

is in bijection with the set of tuples $(Q, H) \in A[T] \times A\{\{T\}\}$ such that

- the constant term of Q is 1 and the leading term a unit in A (think of as $Q = \prod_{i=1}^{d}(1 - u_i T)$),
- $\mathrm{gcd}(Q, H) = 1$ and
- $P_{\iota(\alpha)U_p} = QH$.

I am wondering whether the condition that Q is irreducible is missing. Explicitly, the bijection is such that $V = \mathrm{Sp}\big(A\langle T\rangle/(Q^*(T))\big)$, where Q^* is the polynomial Q 'the other way around', i.e. $Q^*(T) = T^{\deg(T)}Q(T^{-1})$. That Q is a polynomial is a consequence of the assumed finiteness of $V \xrightarrow{\pi_\alpha} \pi_\alpha(V)$. We shall sometimes write Q_V to indicate the relationship between V and Q.

By construction, we have the

Fact: The degree of $\pi_\alpha : V \to Y$ is equal to $d := \deg(Q)$.

We just pause for a second to explicit what the finiteness means on $\mathbb{C}_p$-points: For every $\kappa \in Y(\mathbb{C}_p)$ there are generically precisely d points $(\kappa, u_i^{-1}) \in V(\mathbb{C}_p)$, namely the occuring u_i are precisely the zeros of $Q^{\kappa,*}$. (Recall that $Q \in A[T]$. Specialising to a weight κ means applying the weight character κ to the coefficients of Q making it a polynomial $Q^\kappa \in \mathbb{C}_p[T]$.) We know by what we recalled in the previous section that there are overconvergent eigenforms of the chosen weight κ with the $\iota(\alpha)U_p$-eigenvalues u_i.

It is not this discrete point of view, i.e. single κs, that we want to assume. Instead, we use the fact that overconvergent eigenforms come in families; in particular, we will now consider the A-module $M_Y^\dagger$ (see Ralf's talk), which consists of families of overconvergent modular forms over Y.

Fact: The polynomial $Q = Q(V)$ cuts $M_Y^\dagger$ into two pieces (closed $A(Y)$-submodules)

$$M_Y^\dagger = N_\alpha(V) \oplus F_\alpha(V)$$

such that $Q^*(\iota(\alpha)U_p)$ acts invertibly on $F_\alpha(V)$ and kills $N_\alpha(V)$.

The A-module $N_\alpha(V)$ is called the space of *overconvergent finite slope modular forms over V*. An element of this space is a family of overconvergent modular eigenforms with weights in Y which we regard as eigenfunctions for $\iota(\alpha)U_p$ with eigenvalues one of the zeros of Q^*.

Fact: $N_\alpha(V)$ is (of course) Hecke stable and it is locally free of rank $d = \deg(Q)$ as an $A = A(Y)$-module.

We want a geometric object whose points are precisely the overconvergent modular forms just described. **The principal idea of the present construction is to use the rigid analytic space associated with a suitable Hecke algebra.**

We define the *overconvergent finite slope Hecke algebra over V* as the $A = A(Y)$-module

$$\mathbb{T}_\alpha(V) := \mathrm{im}\left(\mathcal{H} \otimes_\Lambda A \longrightarrow \mathrm{End}_A(N_\alpha(V))\right).$$

Let us make the important observation that the map

$$A(V) = A(Y)\langle T\rangle/(Q^*(T)) \xrightarrow{T \mapsto \iota(\alpha)U_p} \mathbb{T}_\alpha(V)$$

is a well-defined A-algebra homomorphism, since by construction the endomorphism $\iota(\alpha)U_p$ of $N_\alpha(V)$ is the zero map.

Fact: $\mathbb{T}_\alpha(V)$ and $N_\alpha(V)$ are 'almost' dual under the q-expansion pairing:

$$\langle T, f\rangle = a_1(f|T),$$

in the sense that we have an injection $\mathbb{T}_\alpha(V) \hookrightarrow \mathrm{Hom}_A(N_\alpha(V), A)$ with cokernel corresponding to the eigenforms with constant q-expansion (if they exist for the allowed weights Y).

Finally, we can define the desired piece $D_\alpha(V)$ as

$$D_\alpha(V) := \mathrm{Sp}(\mathbb{T}_\alpha(V)).$$

By the above observation, it comes equipped with maps

$$D_\alpha(V) \xrightarrow{z_\alpha} V \xrightarrow{\pi_\alpha} Y$$

and (from the 'almost' duality) we conclude the

Fact: The map $D_\alpha(V) \to Y$ is finite flat of degree $d = \deg(Q)$.

Recall that V was chosen in $\mathcal{C}_\alpha^{\mathrm{irr}}$. For general $V \in \mathcal{C}_\alpha$ such that $V = V_1 \sqcup \cdots \sqcup V_n$ with $V_i \in \mathcal{C}_\alpha^{\mathrm{irr}}$ we define

$$D_\alpha(V) := D_\alpha(V_1) \sqcup \cdots \sqcup D_\alpha(V_n).$$

It is now obvious that $D_\alpha(V)$ **is 1-dimensional**. Note that $D_\alpha(V)(\mathbb{C}_p)$ are overconvergent eigenforms with weight in $Y(\mathbb{C}_p)$ and $\iota(\alpha)U_p$-eigenvalue a zero of $Q_V^{\kappa,*} \in \mathbb{C}_p[T]$.

Moreover, we have the

Fact: If $Q_V(T)$ is squarefree, then the map $D_\alpha(V) \to V$ is generically an isomorphism and both V and $D_\alpha(V)$ are reduced.

We summarise what we did so far. We still keep $\alpha \in \mathcal{H}'$ with $\iota(\alpha) \in R_p^\times$ fixed. Also fix an affinoid subdomain $Y \in \mathcal{W}$.

$M_Y^\dagger$			$\{N_\alpha(V)\}$	$\leftrightarrow$	$\{\mathbb{T}_\alpha(V)\}$	$\leftrightarrow$	$\{D_\alpha(V)\}$
			$\updownarrow$				
$\mathcal{Z}_\alpha$	$\{V \in \mathcal{C}_\alpha \mid \pi_\alpha(V) = Y\}$	$\leftrightarrow$	$\{Q_V\}$				

The passage from the spectral curve to the overconvergent modular forms is made via the polynomials Q_V, thus retaining the information on the eigenvalues of $\iota(\alpha)U_p$.

4.2 Glueing

Also in this subsection, we shall keep $\alpha \in \mathcal{H}'$ with $\iota(\alpha) \in R_p^\times$ fixed. The aim is to obtain a curve D_α by glueing the $D_\alpha(V)$ together. We shall be short and only describe the main steps.

Fact: Let $V_1, V_2 \in \mathcal{C}_\alpha$. Then $V_1 \cap V_2 \in \mathcal{C}_\alpha$.

Fact: Suppose, moreover, that $V_1 \subseteq V_2$. Then $D_\alpha(V_1) \subseteq D_\alpha(V_2)$.

As a consequence, for $U, V \in \mathcal{C}_\alpha$, we have $D_\alpha(U \cap V) \subseteq D_\alpha(U)$. Denote the image by $D_\alpha(U, V)$ (then $D_\alpha(V, U)$, of course, is the image in $D_\alpha(V)$).

Fact: $D_\alpha(U, V) \cong D_\alpha(V, U)$.

Looking only at $\mathbb{C}_p$-points (which we explicited above), the previous two facts are obvious. More precisely, the assumption $V_1 \subseteq V_2$ translates to a surjection

$$A\langle T\rangle/(Q^*_{V_2}) \twoheadrightarrow A\langle T\rangle/(Q^*_{V_1}),$$

i.e. to the condition $Q^*_{V_1}|Q^*_{V_2}$. Due to

$$D_\alpha(V_i)(\mathbb{C}_p) = \{f \in M^\dagger_\kappa(\mathbb{C}_p) \mid \kappa \in \pi_\alpha(V_i)(\mathbb{C}_p), f \text{ eigenform}, \\ \iota(\alpha)U_p f = uf \text{ for } u \text{ such that } Q^{\kappa,*}_{V_i}(u) = 0\},$$

we find $D_\alpha(V_1)(\mathbb{C}_p) \subseteq D_\alpha(V_2)(\mathbb{C}_p)$.

The final fact for $\mathbb{C}_p$-points means that the $D_\alpha(V)(\mathbb{C}_p)$ can be glued together (as sets).

Fact: The $D_\alpha(V)$ can be glued together to give a rigid analytic curve D_α. Moreover, the maps $D_\alpha(V) \xrightarrow{z_\alpha} V \xrightarrow{\pi_\alpha} Y$ give rise to rigid analytic morphisms

$$\pi : D_\alpha \xrightarrow{z_\alpha} \mathcal{Z}_\alpha \xrightarrow{\pi_\alpha} \mathcal{W}.$$

Proposition 4.1 *(a) z_α is a finite morphism.*

(b) π is locally in-the-domain finite flat: D_α is covered by affinoid subdomains, namely the $D_\alpha(V)$, which have the property that their images $Y := \pi(D_\alpha(V)) \subset \mathcal{W}$ are affinoid subdomains of $\mathcal{W}$ and moreover $D_\alpha(V)$ is finite flat over Y.

(c) D_α is a curve (equidimensional of dimension 1).

That D_α is a curve follows from the fact that all the $D_\alpha(V)$ are 1-dimensional (or alternatively from the finiteness over the weight space, which is a curve).

4.3 The curve D

Let $\mathbb{T}_p$ be the Katz Hecke algebra of tame level 1 (see Tommaso's talk), which is defined as the completion of the image of $\mathcal{H}$ in the endomorphisms of the space $\mathbb{V}$ of all Katz p-adic modular functions of tame level 1 (with respect to the compact-open topology on the ring of $\mathbb{Z}_p$-endomorphisms of $\mathbb{V}$).

Proposition 4.2 *Let $\alpha, \beta \in \mathcal{H}'$ such that $\iota(\alpha), \iota(\beta) \in R_p^\times$. If the images of α and β in $\mathbb{T}_p$ are associate (i.e. they generate the same principal ideal), then there is a natural isomorphism $D_\alpha \cong D_\beta$.*

I imagine this proposition somehow like this. The assumption on α and β means that the slopes of the operators $\iota(\alpha)U_p$ and $\iota(\beta)U_p$ are the same. The patches from which D_α is made up are just something like translates of the patches for D_β. It is certainly possible to work this out on points, too, but I have not done so.

In particular, $D_\alpha \cong D_1$ for all α such that the image of α in $\mathbb{T}_p$ is a unit. We let

$$D := D_1.$$

Proposition 4.3 *The rigid analytic curve D is reduced, nested and each irreducible component of D maps surjectively and generic isomorphically onto a Fredholm hypersurface.*

I cannot say much about the proof of this proposition during the talk. It uses the 'flexibility' in choosing 'good' α locally. The 'nestedness' implies that D is the union of its irreducible components.

Recall the decomposition

$$\mathcal{W} = \bigsqcup_{i \in (\mathbb{Z}/p\mathbb{Z})^\times} \mathcal{W}_i,$$

corresponding to the isomorphism

$$\Lambda = \mathbb{Z}_p[[\mathbb{Z}_p^\times]] = \mathbb{Z}_p[[(\mathbb{Z}/p\mathbb{Z})^\times \times \mathbb{Z}_p]] \cong \prod_{i \in (\mathbb{Z}/p\mathbb{Z})^\times} \mathbb{Z}_p[[\mathbb{Z}_p]].$$

Fact: Every irreducible component $\mathcal{D}$ of D maps almost surjectively onto one of the $\mathcal{W}_i$, where 'almost surjectively' means that at most a finite number of points are missed.

The vague reason for this is that the Fredholm power series P_{U_p} also factors into a product $\prod_{i \in (\mathbb{Z}/p\mathbb{Z})^\times} P_{U_p,i}$.

4.4 Relationship between C_p and D

We will now present the construction of a rigid analytic morphism

$$\delta : D \to X_p \times \mathbb{A}^1 \setminus \{0\},$$

which on $\mathbb{C}_p$-points is given by

$$c \mapsto (\lambda_c, \frac{1}{u_c}),$$

where c is an overconvergent eigenform (as seen in the construction of D) with attached pseudo-reprsentation λ_c and U_p-eigenvalue u_c.

We descibe the maps $D \to X_p$ and $D \to \mathbb{A}^1 \setminus \{0\}$. Disregarding some slight complications and just taking $D_1(V)$ instead of D, the first map is just the one coming from $R_p \to \mathbb{T}_\alpha(V)$ (attaching the pseudo-representation). The second map is the composite

$$D \xrightarrow{z_1} \mathcal{Z}_1 \hookrightarrow W \times \mathbb{A}^1 \xrightarrow{\text{2nd projection}} \mathbb{A}^1.$$

The key point is that δ factors through the eigencurve C_p. On $\mathbb{C}_p$-points this seems quite clear (the points correspond to overconvergent eigenforms with finite slope). Note that we are disregarding some subtleties concerning the comparison between overconvergent and p-adic Katz modular forms.

Theorem 4.4 *The above map induces an isomorphism $D \cong C_p^{\text{red}}$.*

As D is reduced, the map factors not only through C_p, but also through C_p^{red}. We have seen that it is a bijection on $\mathbb{C}_p$-points. I cannot say anything on the two pages of proof showing that it is a rigid analytic isomorphism.

5 Summary and main theorems

In this final part, we gather and explain, as well as possible, the main statements of the paper by Coleman and Mazur appearing in Sections 1.5 and 7.6.

The *classical modular locus of tame level* 1 is the set of points

$$\mathcal{M} \subset (X_p \times \mathbb{A}^1)(\mathbb{C}_p)$$

consisting of $\overline{x}_f = (\lambda_f, 1/u_f)$ for f running through all classical modular eigenforms of finite slope on $\Gamma_1(p^n)$ for any n. As above, λ_f denotes the corresponding pseudo-representation and u_f the U_p-eigenvalue.

Theorem 5.1 *The eigencurve C_p is the Fredholm closure of $\mathcal{M}$, i.e. the smallest Fredholm variety containing $\mathcal{M}$.*

A Fredholm variety is – by definition – an arbitrary intersection of Fredholm hypersurfaces. These, in turn, are those rigid analytic subspaces (of $X_p \times \mathbb{A}^1$) cut out by a Fredholm series. As $C_p = \bigcap_\alpha r_\alpha^{-1}(\mathcal{Z}_\alpha)$, it is clear that C_p is a Fredholm variety. Moreover, it clearly contains $\mathcal{M}$. Since among the zeros of each $P_{\iota(\alpha)U_p}$ (after specialising to any integral weight character) there are some corresponding to the inverse U_p-eigenvalue of a classical modular form, it should be possible to show that C_p is indeed the Fredholm closure. But, to be honest, I cannot turn this into a correct proof right now.

From Propositions 4.1 and 4.3 and Theorem 4.4 we obtain:

Theorem 5.2 *(a) The eigencurve C_p is a curve.*

(b) The natural projection of any irreducible component of the reduced eigencurve to weight space is component-wise almost surjective in the sense that given any irreducible component of the reduced eigencurve the complement of its image in the unique irreducible component of weight space containing that image (is empty, or) consists of at most a finite number of weights.

(c) The projection of the reduced eigencurve to weight space is locally in-the-domain finite flat in the sense that C_p^{red} is covered by admissible affinoid domains $\mathfrak{U}$ such that the restriction of projection to weight space to $\mathfrak{U}$ is a finite flat mapping of $\mathfrak{U}$ onto its image in $\mathcal{W}$.

Next, recall that R_p was made in Juan's talk by taking the direct product of the universal deformation rings of the pseudo-reprsentations of the finitely many semi-simple residual p-modular Galois representations. Whenever we write $\overline{\rho}$, we now mean one of those residual representations. Hence,

$$X_p \cong \bigsqcup_{\overline{\rho}} X_{\overline{\rho}} \quad \text{and} \quad X_p \times \mathbb{A}^1 \cong \bigsqcup_{\overline{\rho}} (X_{\overline{\rho}} \times \mathbb{A}^1).$$

Define

$$C_{\overline{\rho}} := C_p \cap X_{\overline{\rho}}.$$

Theorem 5.3 *(a) We have $C_p = \bigsqcup_{\overline{\rho}} C_{\overline{\rho}}$, so that the $C_{\overline{\rho}}$ are admissibly closed open in C_p (in the sense that they are disjoint and form an admissible covering of C_p).*

(b) If two classical modular points $\overline{x}_f$ and $\overline{x}_g$ (in the notation above) lie on the same connected component of the reduced eigencurve, then the reductions of their Fourier expansions are congruent modulo $\mathcal{O}_{\mathbb{C}_p}$ (at almost all prime coefficients) and their attached semi-simple residual Galois representations are equivalent.

Theorem 5.4 *The reduced eigencurve is the rigid Zariski closure of the classical modular locus.*

This theorem, which we will not prove, needs a deep theorem of Coleman's as input: Every overconvergent eigenform of integral weight k and slope less than $k-1$ and different from $\frac{k-1}{2}$ is classical. The proof then uses that every irreducible component of the reduced eigencurve contains forms of all integral weights bigger than a certain k_0 (depending on the component) due to the almost surjectivity.

We will finish with the following theorem.

Theorem 5.5 *If $(f_n)_n$ is a sequence of normalised eigenforms (of tame level 1) with Fourier coefficients in $\mathbb{C}_p$ such that f_n has weight κ_n and the sequence of slopes $(v(a_p(f_n)))_n$ is bounded independently of n and the sequence of q-expansions in $\mathbb{C}_p[[q]](1/q)$ converges coefficient-wise to a series $f(q)$, then the sequence of weights converges to a weight κ and $f(q)$ is the Fourier expansion of an overconvergent modular eigenform of tame level 1, weight κ and finite slope.*

Since the weights can be read off from the q-expansions, it should be an almost formal consequence that the weights converge (Coleman and Mazur refer to an article by Serre, which I did not look up). The point is that due to the convergence of the weights, eventually all the f_n will correspond to $\mathbb{C}_p$-points on $D_1(V)$ for some $V \in \mathcal{C}_1$. Denote by $\eta_n : \mathbb{T}_1(V) \to \mathbb{C}_p$ the corresponding ring homomorphisms so that $f_n = \sum_m \eta_n(T_m) q^m$. The point-wise convergence of the q-expansions now means that for each m, the $\eta_n(T_m)$ converge to some $\eta(T_m)$. As the images of T_m generate all of $\mathbb{T}_1(V)$, like this one obtains a limit function $\eta : \mathbb{T}_1(V) \to \mathbb{C}_p$, i.e. the desired overconvergent modular form.