The Modularity Conjecture for Calabi-Yau Threefolds

Matt Ward

December 14, 2011

1 Introduction

A Calabi-Yau variety, X, of dimension n, is defined to be a smooth projective variety over a field k that satisfies $\omega_X := \bigwedge^n \Omega^1 \simeq \mathcal{O}_X$ and $H^j(X, \mathcal{O}_X) = 0$ for $0 < j < n$. This can be seen as a higher dimensional "cohomological" analogue of an elliptic curve. A Calabi-Yau threefold is a Calabi-Yau variety of dimension three.

The main goal of this paper is to introduce two equivalent notions of modularity for rigid Calabi-Yau threefolds defined over $\mathbb{Q}$. They are both defined via the L-series of the variety. One way is to produce a certain ℓ-adic Galois representation and the other is to define the L-series directly from Frobenius actions. There are several more equivalent notions of modularity, but we'll focus on these two.

We should probably start by noting that modularity of one-dimensional Calabi-Yau varieties was the famous Taniyama-Shimura conjecture. This took alone was very difficult to prove. Wiles and Taylor published the case of semi-stable elliptic curves in 1995, but it wasn't until 2001 that the full conjecture was published by Breuil, Conrad, Diamond, and Taylor. It may seem curious that we immediately jump to threefolds instead of focusing on two-dimensional varieties next. This has to do with the fact that threefolds can be rigid and hence produce certain small cohomology whereas in the two-dimensional case the cohomology groups have dimensions $1, 0, 22, 0$, and 1 respectively. To talk about modularity most easily we will want a two-dimensional representation. We won't dwell on this point. Instead, it should become clearer after a thorough analysis of the three-dimensional case.

2 Cohomology of Rigid Calabi-Yau Threefolds

For this section the characteristic of the field will still not matter. Define the Hodge numbers of X to be $h^{i,j} = \dim_k H^j(X, \Omega^i)$. We will assume our variety satisfies Hodge symmetry (this is always the case over a characteristic 0 field or if X arises as the reduction to a finite field from an integral model) which says that $h^{i,j} = h^{j,i}$.

By Serre duality, Hodge symmetry, and triviality of the canonical bundle we get that $h^{0,0} = h^{3,0} = h^{0,3} = h^{3,3} = 1$, $h^{1,2} = h^{2,1}$ is unknown, $h^{1,1} = h^{2,2}$ is unknown, and all other Hodge numbers are 0. This means that essentially there are only two unknowns among the Hodge numbers. A rigid Calabi-Yau threefold is one that has no non-trivial infinitesimal deformations. A standard result of deformation theory says that infinitesimal deformations are parametrized by $H^1(X, \mathcal{T}) \simeq H^1(X, \Omega^2)$, thus $h^{1,2} = 0$ for a rigid CY threefold, and we are only left with one unknown Hodge number. For our purposes that number is irrelevant, but it is still an open conjecture that it cannot be arbitrarily large. This conjecture was given in [6] and at that time the largest known Euler characteristic (alternating sum of these numbers) was 960.

Suppose now that X is a rigid CY threefold defined over $\mathbb{Q}$. Let $\overline{X}$ be the base-change of X to $\overline{\mathbb{Q}}$. Note that $\overline{X}$ is also rigid since by flat base-change $H^1(\overline{X}, \mathcal{T}) \simeq H^1(X, \mathcal{T}) \otimes \overline{\mathbb{Q}} = 0$. Since we are over an algebraically closed field of characteristic 0 it is well-known that the Hodge-de Rham spectral sequence degenerates at E_1 which implies that $H^3_{dR}(\overline{X}/\overline{\mathbb{Q}}) \simeq \bigoplus_{i+j=3} H^j(\overline{X}, \Omega^i)$. All of the Hodge numbers are known and given above, so we see that $\dim_{\overline{\mathbb{Q}}} H^3_{dR}(\overline{X}/\overline{\mathbb{Q}}) = 2$.

Similarly, since we are over an algebraically closed field of characteristic 0 we get that for any prime ℓ the middle ℓ-adic cohomology $H^3_{et}(\overline{X}, \mathbb{Q}_\ell) := \varprojlim H^3_{et}(\overline{X}, \mathbb{Z}/\ell^n) \otimes_{\mathbb{Z}_\ell} \mathbb{Q}_\ell$ is also two-dimensional.

3 The Galois Action

For the definition of the Galois action we don't need any special considerations except that our variety X is defined over $\mathbb{Q}$. We continue to use $\overline{X}$ to mean the base-change to $\overline{\mathbb{Q}}$. Consider the absolute Galois group $G := Gal(\overline{\mathbb{Q}}/\mathbb{Q})$. Any element $\sigma \in G$ is an automorphism of $\overline{\mathbb{Q}}$, so it induces an automorphism $\mathrm{Spec}(\overline{\mathbb{Q}}) \to \mathrm{Spec}(\overline{\mathbb{Q}})$. Thus we can form a commuting diagram:

$$\begin{tikzcd}
\overline{X} \arrow[r, "\sigma"] \arrow[d] & \overline{X} \arrow[d] \\
\mathrm{Spec}(\overline{\mathbb{Q}}) \arrow[r] & \mathrm{Spec}(\overline{\mathbb{Q}})
\end{tikzcd}$$

By definition the action of σ is the identity on topological spaces and the map on structure sheaves that makes the diagram commute. This means that the diagram is not a pullback (i.e. not cartesian). We will abuse notation by calling this action on schemes σ as well. Note that this is not a $\overline{\mathbb{Q}}$-map, but since $\sigma|_{\mathbb{Q}} = \mathrm{id}$ it is a map of $\mathbb{Q}$-schemes.

We can "fix" this in the standard way of just defining a new structure map to be the composition $\overline{X} \to \mathrm{Spec}\,\overline{\mathbb{Q}} \stackrel{\sigma}{\to} \mathrm{Spec}\,\overline{\mathbb{Q}}$. Now the action should be considered an isomorphism of schemes over $\overline{\mathbb{Q}}$ rather than an automorphism. Since it is an isomorphism it will induce an isomorphism on cohomology $H^3_{et}(\overline{X}, \mathbb{Q}_\ell) \stackrel{\sigma^*}{\to} H^3_{et}(\overline{X}, \mathbb{Q}_\ell)$.

This works for any element of G, so when we assume that X is a rigid Calabi-Yau threefold the action on cohomology induces a two-dimensional representation $\rho'_X : G \to \mathrm{Aut}(H^3_{et}(\overline{X}, \mathbb{Q}_\ell)) \simeq GL_2(\mathbb{Q}_\ell)$. Note that in the previous section we checked that the middle ℓ-adic cohomology was two-dimensional.

If we think to the elliptic curve case for motivation, then we define the Galois representation via its action on the Tate module $T_\ell E$. Since $T_\ell E \simeq H^1_{et}(E, \mathbb{Z}_\ell)^\vee$ the above representation will not match up. In fact, it is precisely the contragredient representation. Thus we will define the actual Galois representation ρ_X on the dual space via the contragredient representation. One notable difference is that this will invert eignevalues.

This is the convention used in [1], but since we'll only care about $\rho_X(Frob_p)$ often the convention is to not take the contragredient and instead use the geometric rather than arithmetic Frobenius. This will be explained in more detail in the next section. At this point we could just define X to be modular if this Galois representation is modular, but since that hasn't been defined we will unravel what this means in the next section.

4 Actions of the Frobenii

The notation will continue from the previous sections. We will now define modularity via defining the L-series for X and for f a modular form. A subtlety that we won't worry about is that there will be an integral model for X needed in this section and so we fix such a model for the section.

We will first define the L-series via the Galois representation. Fix p a prime not equal to ℓ and of good reduction for X (equivalently the representation is unramified at p). Standard arguments give us that ρ_X factors through some $Gal(M/\mathbb{Q})$. Here we have a conjugacy class Frob_p whose image under ρ_X has well-defined trace and determinant, since M is unramified at p.

We define

$$L(X,s) = L(H^3_{et}(\overline{X},\mathbb{Q}_\ell),s) = (*)\prod_{p\ good}\frac{1}{1-\mathrm{tr}(\rho_X(\mathrm{Frob}_p))p^{-s}+\det(\rho_X(\mathrm{Frob}_p))p^{-2s}}$$

where $(*)$ is a product of terms at the bad primes. Note that since this is a two-dimensional representation basic linear algebra tells us that the product is over the simpler expression $(\det(I-\rho_X(\mathrm{Frob}_p)p^{-s}))^{-1}$.

The real purpose of this section is to write down this L-series without reference to the Galois representation. In order to ease notation we'll denote the reduction of X at a fixed good prime p by $Y := X_{\mathbb{F}_p}$ and basechanging to the algebraic closure $\overline{Y} := X_{\overline{\mathbb{F}_p}}$. For notation let $k = \mathbb{F}_p$.

We have several natural Frobenius actions on $\overline{Y}$. The first we'll call the absolute Frobenius which we'll denote $F_{ab} : \overline{Y} \to \overline{Y}$. This is the identity on the topological space and the p-th power map on the structure sheaf. On affine patches $\mathrm{Spec}\,A \to \mathrm{Spec}\,\overline{k}$ the map is the one induced by $a \mapsto a^p$ on A. We can check that the map on topological spaces is the identity. For any prime ideal $\mathfrak{q} \in \mathrm{Spec}\,A$ the contraction $\mathfrak{q}^c = \{a \in A : a^p \in \mathfrak{q}\} = \{a \in A : a \in \mathfrak{q}\} = \mathfrak{q}$ by the property of $\mathfrak{q}$ being prime. This map translates in the language of schemes to $(id, F) : (\overline{Y}, \mathcal{O}_{\overline{Y}}) \to (\overline{Y}, \mathcal{O}_{\overline{Y}})$ where F is raising sections of the sheaf to the p-th power.

Note that the absolute Frobenius is not a map of $\overline{Y}$ over $\overline{k}$. The map is also not the pullback despite making the following diagram commute:

$$\begin{tikzcd}
\overline{Y} \arrow[r, "F_{ab}"] \arrow[d] & \overline{Y} \arrow[d] \\
\mathrm{Spec}\,\overline{k} \arrow[r, "\mathrm{Frob}_p"] & \mathrm{Spec}\,\overline{k}
\end{tikzcd}$$

Call the standard structure map $\phi : \overline{Y} \to \mathrm{Spec}\,\overline{k}$ and $\overline{Y}^{(p)}$ the pullback of the two maps in the above diagram: $\phi : \overline{Y} \to \mathrm{Spec}\,\overline{k}$ and $\mathrm{Frob}_p : \mathrm{Spec}\,\overline{k} \to \mathrm{Spec}\,\overline{k}$. Since we have the above commutative diagram we get by the universal property of a pullback diagram some map $F_r : \overline{Y} \to \overline{Y}^{(p)}$ (the relative Frobenius) and we'll call the projection on the first factor $F_{ar} : \overline{Y}^{(p)} \to \overline{Y}$. Just by definition this gives us the following diagram:

$$\begin{tikzcd}
\overline{Y} \arrow[dr, "F_r"] \arrow[drr, bend left, "F_{ab}"] \arrow[ddr, bend right, "\phi"'] & & \\
& \overline{Y}^{(p)} \arrow[r, "F_{ar}"] \arrow[d] & \overline{Y} \arrow[d, "\phi"] \\
& \mathrm{Spec}(\overline{k}) \arrow[r, "a^p"] & \mathrm{Spec}(\overline{k})
\end{tikzcd}$$

Note that there are lots of relative Frobenii. We could produce one in the same manner using any k-morphism $\overline{Y} \to Z$ by making the pullback $Y^{(p)}$ via the absolute Frobenius on Z (as opposed to $\operatorname{Spec} \overline{k}$). It might be better in other situations to notate this $F_{\overline{Y}/Z}$. This means that our $F_r = F_{\overline{Y}/\operatorname{Spec} \overline{k}}$.

Before giving an example, these can be explained in a slightly more concrete way. Consider $\overline{Y} = Y \times_{\operatorname{Spec} k} \operatorname{Spec} \overline{k}$. It is checked in the last section of [5] that these definitions give $F_r = F_{ab} \otimes id_{\overline{k}}$. The arithmetic Frobenius is $F_{ar} = id_Y \otimes \mathrm{Frob}_p$ and the inverse of this is the geometric Frobenius $F_{ge} = id_Y \otimes \mathrm{Frob}_p^{-1}$.

The above descriptions make it easier to figure out these maps for an example. Let $Y = \operatorname{Spec} k[t]$ (recall that $k = \mathbb{F}_p$). This means that $\overline{Y} = \operatorname{Spec} \overline{k}[t] = \operatorname{Spec}(k[t] \otimes_k \overline{k})$. The descriptions in terms of the ring homomorphism that induces the map on the spectra are as follows. The absolute is still just $f \mapsto f^p$. The relative is $F_{ab} \otimes id$. Since the absolute raises elements of $k[t]$ to the p, everything in k is fixed by this map, and on $\overline{k}$ it is defined to be fixed. This means that the relative only alters the t by $t \mapsto t^p$. This is sometimes referred to as "raising coordinates to the p-th power". The arithemtic Frobenius does nothing to the $k[t]$ part, but raises the $\overline{k}$ coefficients to the p, so $\sum a_n t^n \mapsto \sum a_n^p t^n$. Likewise, the geometric Frobenius takes the p-th root of the coefficients.

Straightforward (but non-trivial) computations given in [5] also give that the map that the absolute Frobenius induces on the étale site is trivial. If we look at our diagram we see that $F_{ab} = F_{ar} \circ F_r$. Since the induced map on cohomology is contravariant this gives $F_r^* \circ F_{ar}^* = id$. This means that on cohomology $F_r^* = (F_{ar}^*)^{-1} = F_{ge}^*$ by definition of the geometric Frobenius.

Now the smooth, proper base-change theorem for étale cohomology tells us that $H^3_{et}(\overline{Y}, \mathbb{Q}_\ell) \simeq H^3_{et}(\overline{X}, \mathbb{Q}_\ell)$ which is two-dimensional. Since the F_r action here is a linear operator on a vector space it makes sense to take the trace and determinant. We can define the L-series without use of the Galois representation as:

$$L(X,s) = (*) \prod_{p\ good} \frac{1}{1 - \mathrm{tr}(F_r^*)p^{-s} + \det(F_r^*)p^{-2s}}$$

where again the $(*)$ is a product of terms involving primes of bad reduction. Since there are only finitely many this is irrelevant for the definition of modularity. Of course we could have defined this without all the different Frobenius actions (we only used the relative one), but now we can get to the punchline. These two L-series are actually the same.

We just sketched above that the action of F_r and F_{ge} were the same on the étale site. But $F_{ge} = 1 \times \mathrm{Frob}_p^{-1}$ where Frob_p is the canonical generator of $Gal(\overline{k}/k)$. We have a surjection $Gal(\overline{\mathbb{Q}}/\mathbb{Q}) \to Gal(\overline{k}/k)$ and if we consider Frob_p a lift of this element by the functoriality and equivariant isomorphisms above we get that $\mathrm{tr}(\rho_X(\mathrm{Frob}_p)) = \mathrm{tr}(\rho'_X(\mathrm{Frob}_p^{-1})) = \mathrm{tr}(F_r^*)$. The determinant term turns out to always be p^3 since it can be checked to be the third power of the ℓ-adic cyclotomic character in both cases. Thus the two L-series are the same. This also tells us the representation is odd.

Note that they appear to be off by an inverse, but we actually took the contragredient representation of the one that acts on $H^3_{et}(\overline{X}, \mathbb{Q}_\ell)$, so the inverse corrects for this and they are actually the same.

5 The Modularity Conjecture

The last piece we need to state the modularity conjecture is the L-function of a modular form. Let $f \in S_4(\Gamma_0(N))$ be a weight 4 cusp form. We can write f in its q-expansion $f = \sum_{n=1}^{\infty} a_n q^n$. The Mellin transform is given by $L(f,s) = \sum_{n=1}^{\infty} \frac{a_n}{n^s}$ and this has an Euler product expansion given by $L(f,s) = \prod_{p|N}(*) \prod_{p\nmid N} \frac{1}{1 - a_p p^{-s} + p^{3-2s}}$. We say that a rigid Calabi-Yau threefold is modular if $L(X,s)$ coincides with $L(f,s)$ up to finitely many primes for some $f \in S_4(\Gamma_0(N))$.

Theorem 5.1 *Every rigid Calabi-Yau threefold defined over $\mathbb{Q}$ is modular.*

Proof The proof is given in [1] is essentially a corollary to Serre's Conjecture which was proved in [2] and [3] by Khare and Wintenberger. The earliest reference that cites this as following from Serre's Conjecture seems to be Yui in [6]. It would probably be more proper to say that the proof follows from Serre's "method". Suppose we have some fixed set of primes S and a two-dimensional, continuous, odd Galois representation ρ_λ on V_λ, a $\mathbb{Q}_\lambda$ vector space, unramified outside $S \cup \{\lambda\}$ for all rational primes λ. This method says that if some extra conditions are satisfied, then it is modular. More precisely, define $\overline{\rho}_\lambda : G \to GL_2(\mathbb{F}_\lambda)$ the reduction and semisimplification. If there is an infinite set of primes I such that the following conditions are satisfied:

1. For all $\lambda \in I$ the representation $\overline{\rho}_\lambda$ is absolutely irreducible
2. For all $\lambda \in I$ and fixing such a λ any $p \notin S \cup \{\lambda\}$ the characteristic polynomial of $\rho_\lambda(\mathrm{Frob}_p)$ acting on V_λ is independent of p
3. There is a finite universal upper bound, k_0, on the Serre weight k_λ attached to $\overline{\rho}_\lambda$ as λ ranges through I
4. There is some integer N_0 so that for all $\lambda \in I$ the Serre levels for $\overline{\rho}_\lambda$ divide N_0

then this method of Serre produces a cuspidal Hecke eigenform f of weight k_0 and new level dividing N_0 such that $\rho_{f,\lambda}$ is isomorphic to ρ_λ. All of these properties can be checked for our representation using standard geometric arguments. For instance the second condition is a consequence of the Weil conjectures. Fontaine checked that for sufficiently large λ the weight will be 4, so that is how we get the weight. A slightly more complicated argument is needed to check the level condition by checking that the ℓ-adic representations are strongly compatible. ∎

References

[1] Fernanda Gouveau and Noriko Yui. Rigid Calabi-Yau threefolds over $\mathbb{Q}$ are modular. *Expositiones Mathematicae*, 29(1):142–149, 2011.

[2] Chandrashekhar Khare and Jean-Pierre Wintenberger. Serre's modularity conjecture (i). *Inventiones Mathematicae*, 178(3):485–504, 2009.

[3] Chandrashekhar Khare and Jean-Pierre Wintenberger. Serre's modularity conjecture (ii). *Inventiones Mathematicae*, 178(3):504–586, 2009.

[4] Christian Meyer. *Modular Calabi-Yau Threefolds.* American Mathematical Society, 2005.

[5] J.S. Milne. *Étale Cohomology.* Princeton University Press, 1980.

[6] Noriko Yui. Update on the modularity of Calabi-Yau varieties. In *Calabi-Yau Varieties and Mirror Symmetry*, number 38 in Fields Inst. Commun., pages 307–362. Amer. Math. Soc., Providence, RI, 2001.