A circle method for algebraic geometers

Matthew Hase-Liu

Abstract

Browning and Vishe studied the moduli space of smooth genus zero curves of fixed degree on a smooth low-degree hypersurface using the circle method, a technique from analytic number theory. I'll explain how their strategy can be interpreted completely algebro-geometrically, and then use this perspective to generalize their results to the higher genus setting. Time permitting, I'll also discuss some applications to Geometric Manin's conjecture and terminal singularities of these moduli spaces, the latter of which is joint work with Jakob Glas.

1 Introduction

Let $X \hookrightarrow \mathbb{P}^n_{\mathbb{C}}$ be a smooth hypersurface of degree d, say defined by $F(x_0, \ldots, x_n) = 0$.

The "naive" moduli space of rational curves of degree e on X, denoted by $\mathrm{Mor}_e(\mathbb{P}^1, X)$, can be thought of as tuples $(f_0(u,v), \ldots, f_n(u,v))$ of homogeneous polynomials with no common roots, with each f_i of degree e, and such that $F(f_0, \ldots, f_n)$ vanishes identically.

Its expected dimension is given by the $(n+1)(e+1)$ conditions to choose coefficients for the f_i's, minus the $de+1$ conditions defined by the vanishing, minus one more for scaling:

$$(n+1)(e+1) - (de+1) - 1.$$

Two related spaces are $\mathcal{M}_{0,0}(X, e)$, which are the degree e morphisms up to automorphisms of $\mathbb{P}^1$, and its Kontsevich compactification $\overline{\mathcal{M}}_{0,0}(X, e)$. These have the expected dimension

$$(n+1)(e+1) - (de+1) - 1 - 3.$$

Coskun-Harris-Starr conjectured the following:

Conjecture 1. *Suppose X is a general hypersurface, $d \geq 3$, and $n \geq d+1$. Then, $\overline{\mathcal{M}}_{0,0}(X, e)$ (and hence $\mathcal{M}_{0,0}(X, e)$) is irreducible and has the expected dimension.*

For general X:

- Harris-Roth-Starr '04 for $d < (n+1)/2$.
- Beheshti-Kumar '13 for $d < 2n/3$ and $n \geq 20$.
- Riedl-Yang '14 for $n \geq d+2$.

In the direction of arbitrary X:

- Coskun-Starr '09 for any smooth cubic X and $n \geq 5$. But for $n = 4$, you have the right dimension but two irreducible components.

Browning-Vishe '17 showed the following result, using analytic number theory over function fields:

Theorem 2. For $n \geq 2^{d-1}(5d-4)$, $e \geq 1$, and X smooth of degree d, $\mathcal{M}_{0,0}(X, e)$ is irreducible and has the expected dimension.

To do this, they work with the space $\mathrm{Mor}_e(\mathbb{P}^1, X)$.

- Originally proposed by Ellenberg-Venkatesh.
- Pugin's thesis '11 for diagonal cubic hypersurfaces.
- Browning-Sawin '18 improved this to $n \geq 2^d(d - 1/2)$.

2 What did I do?

(i) Re-interpreted Browning-Vishe analytic strategy geometrically.

(ii) Extension to higher genus:

Theorem 3. Let C be a smooth projective curve of genus $g \geq 1$ and $X \hookrightarrow \mathbb{P}^n$ a smooth hypersurface of degree d, both defined over $\mathbb{C}$. Suppose the following holds:

$$n > 2^d(d-1) + 1, e \geq C(d, g)$$

for some explicit but large constant $C(d, g)$ depending only on d and g.

Then, the mapping space $\mathrm{Mor}_e(C, X)$ is irreducible and of the expected dimension $(n + 1)(e - g + 1) - (de - g + 1) + g$. Similarly, $\mathcal{M}_g(X, e)$ is irreducible and of the expected dimension.

Remark 4. The constant $C(d, g)$ is necessary: for e small compared to d and g, we can construct maps $C \to X$ that factor through $\mathbb{P}^1$ as hyperelliptic covers $C \to \mathbb{P}^1$ that give you many more maps than expected!

Remark 5. $\overline{\mathcal{M}}_g(X, e)$ pretty much always has larger than expected dimension for $d \geq 3$, so this strategy can't be made to work for the compactification.

(iii) Converse to Geometric Manin's conjecture:
For a fixed smooth projective variety $X/\mathbb{C}$, the "a-numbers" of a subvariety $V \subset X$ are roughly expected to control the components of $\mathrm{Mor}(\mathbb{P}^1, X)$ that have larger than expected dimension. Specifically, work of Lehman-Riedl-Sengupta-Tanimoto predicts that if $a(V) < 1$ for all proper subvarieties $V \subset X$, then each component of $\mathrm{Mor}(\mathbb{P}^1, X)$ has the expected dimension.

Corollary 6. If $X \hookrightarrow \mathbb{P}^n$ a smooth hypersurface of degree d such that $n > 2^d (d-1)$, then for any proper subvariety $V \subset X$, we have $a(V) < 1$.

(iv) Application to terminal singularities (joint work with Glas):

Theorem 7. Using the same assumptions as Theorem 3 but with $n > 2^{d-2}(d-1)(4d^2-4d+3)$, the moduli space $\mathcal{M}_g(X,e)$ has at worst terminal singularities.

For the rest of the talk, I want to discuss this geometric interpretation of Browning-Vishe's analytic strategy.

(i) Reduction to characteristic p and point-counting.

(ii) Expressing the point-counts by integrating $S(\alpha)$, ranging over $\alpha \in$"circle," where $S(\alpha)$ is an exponential sum associated to each α.

(iii) Dividing up the circle in major and minor arcs.

(iv) Major arc contribution: a direct computation using PIE.

(v) Minor arc contribution:

   (a) Relating $S(\alpha)$ to a "linearized" $N(\alpha)$ via "Weyl differencing."

   (b) Bounding $N(\alpha)$ via the "shrinking lemma."

Step 1: Reduction to characteristic p and point-counting

To go to characteristic p, the idea is to spread out and realize $\operatorname{Mor}_e(C,X)$ as the generic fiber of some family over $\operatorname{Spec}\Lambda$, where Λ is obtained by adjoining all the coefficients needed to define $\operatorname{Mor}_e(C,X)$ to $\mathbb{Z}$. Every residue field of any maximal ideal of Λ is a finite field, so proving the irreducibility + expected dimension over finite fields implies the result for the generic fiber.

This is the two-sentence explanation for the reduction to characteristic p.

How do we relate this to point-counting? Let μ be the expected dimension. Suppose that C and X are defined over $\mathbb{F}_q$ for p sufficiently large. Then, if we can show that

$$\frac{\# \operatorname{Mor}_e(C,X)\left(\mathbb{F}_{q^\ell}\right)}{q^{\ell\mu}} \to 1$$

as $\ell \to \infty$, then the Lang-Weil bounds imply that $\# \operatorname{Mor}_e(C,X)$ is both irreducible and of the expected dimension. Indeed, the Lang-Weil bounds tell us that the number of $\mathbb{F}_{q^\ell}$-points of a variety is roughly the number of geometrically irreducible components times $(q^\ell)^{\text{dim of the variety}}$.

So we've reduced to a point-counting problem, namely bounding $\# \operatorname{Mor}_e(C,X)\left(\mathbb{F}_{q^\ell}\right)$.

Step 2: Point-counts as exponential sums

Let L be a degree e line bundle on C. The "circle" is defined as the space of linear funtionals $H^0(C, L^{\otimes d})^\vee$. For $\alpha \in H^0(C, L^{\otimes d})^\vee$, we define the exponential sum associated to α as follows:

$$S(\alpha) \coloneqq \sum_{\vec{x} \in H^0(C,L)^{n+1}} \psi\left(\alpha\left(F\left(\vec{x}\right)\right)\right),$$

where ψ be a non-trivial additive character $\mathbb{F}_q \to \mathbb{C}^\times$. This expression makes sense because F can be viewed as a map $H^0\left(C, L\right)^{n+1} \to H^0\left(C, L^{\otimes d}\right)$.

Then the "igniting spark" of the circle method is the observation that

$$\begin{aligned}
\sum_{\alpha \in H^0(C,L^{\otimes d})^\vee} S(\alpha) &= \sum_{\vec{x} \in H^0(C,L)^{n+1}} \sum_{\alpha \in H^0(C,L^{\otimes d})^\vee} \psi\left(\alpha\left(F\left(\vec{x}\right)\right)\right) \\
&= \sum_{\vec{x} \in H^0(C,L)^{n+1}} \begin{cases} 0 & \text{if } F\left(\vec{x}\right) \neq 0, \\ \#H^0\left(C, L^{\otimes d}\right)^\vee & \text{else} \end{cases} \\
&= \#H^0\left(C, L^{\otimes d}\right)^\vee \underbrace{\#\left\{\left(x_0, \ldots, x_n\right) \in H^0\left(C, L\right)^{n+1} : F\left(x_0, \ldots, x_n\right) = 0\right\}}_{\text{almost } \operatorname{Mor}_e(C,X)}.
\end{aligned}$$

The term "almost $\operatorname{Mor}_e\left(C, X\right)$" is basically $\operatorname{Mor}_e\left(C, X\right)$, except that we remember the line bundle L and forget the base-point free conditions on the sections $x_0, \ldots, x_n$. But otherwise we can basically express point-counts of $\operatorname{Mor}_e\left(C, X\right)$ in terms of $\sum_{\alpha \in H^0(C,L^{\otimes d})^\vee} S(\alpha)$.

Since $\#H^0\left(C, L^{\otimes d}\right)^\vee = q^{de-g+1}$, the point-counting bound we'd like to show now looks like

$$\frac{\sum_{\alpha \in H^0(C,L^{\otimes d})^\vee} S(\alpha)}{q^{(n+1)(e-g+1)}} \to 1.$$

Step 3: Major and minor arcs

To define the major and minor arcs, we consider subschemes of C for which α factors through—this is the analog of Dirichlet approximation.

- Any α factors through a closed subscheme of degree at most $de/2 + 1$.
- Let $\deg(\alpha)$ be the smallest degree of a subscheme α factors through.
- Major arcs: α such that $\deg(\alpha) \leq e - 2g + 1$.
- Minor arcs: all other α.

A direct computation shows

$$\frac{\sum_{\deg \alpha \leq e-2g+1} S(\alpha)}{q^{(n+1)(e-g+1)}} \to 1.$$

So it suffices to prove that

$$\frac{\sum_{\deg \alpha > e-2g+1} S(\alpha)}{q^{(n+1)(e-g+1)}} \to 0,$$

for which it suffices to show each $S(\alpha)$ is small.

Step 5a: "Linearizing" $S(\alpha)$ (Weyl differencing)

Note that if $\alpha \circ F$ was linear, then

$$\sum_{\vec{x} \in H^0(C,L)^{n+1}} \psi\left(\alpha\left(F\left(\vec{x}\right)\right)\right) = \begin{cases} \#H^0\left(C, L\right)^{n+1} & \alpha \circ F = 0 \\ 0 & \text{otherwise} \end{cases}.$$

Slogan: Weyl differencing allows you to reduce to the linear situation.

There are certain multilinear forms $\Psi_j\left(\vec{x}^{(1)}, \ldots, \vec{x}^{(d-1)}\right)$ that depend on the equation F. Repeated application of Cauchy-Schwarz gives:

$$|S(\alpha)|^{2^{d-1}} \leq \left(\#H^0\left(C, L\right)^{n+1}\right)^{2^{d-1}-d+1} N(\alpha),$$

where $N(\alpha)$ is the number of tuples

$$\left(\vec{x}^{(1)}, \ldots, \vec{x}^{(d-1)}\right) \in \left(H^0\left(C, L\right)^{n+1}\right)^{d-1}$$

such that $\alpha\left(\Psi_j\left(x^{(1)}, \ldots, x^{(d-1)}\right) x\right)$ vanishes as a function of $x \in H^0\left(C, L\right)$.

This is the part of the argument (namely applying Cauchy-Schwarz a bunch of times) that gives an exponential lower bound on n.

Step 5b: "Shrinking" $N(\alpha)$

Slogan: bound $N(\alpha)$ in terms of a related quantity $N_s(\alpha)$ by interpreting both geometrically.

Define $N_s(\alpha)$ as the number of tuples

$$\left(\vec{x}^{(1)}, \ldots, \vec{x}^{(d-1)}\right) \in \left(H^0\left(C, L(-s)\right)^{n+1}\right)^{d-1}$$

such that $\alpha\left(\Psi_j\left(x^{(1)}, \ldots, x^{(d-1)}\right) x\right)$ vanishes as a function of $x \in H^0\left(C, L\left(\left(d-1\right) s\right)\right)$.

The geometry of numbers over function fields gives the following inequality:

$$\frac{N(\alpha)}{N_s(\alpha)} \leq q^{\text{something in terms of } d,n,s}$$

(Browning-Vishe) Arithmetically, this comes from analyzing a certain lattice and its successive minima, and expressing $N(\alpha)$ as the number of norm-bounded elements of the lattice.

Geometrically, in the case of genus zero, a lattice is a locally-free module over $\mathbb{A}^1$ and a choice of norm from ∞—this is equivalent to a vector bundle over $\mathbb{P}^1$ by the Beauville-Laszlo theorem. By Birkhoff-Grothendieck, these vector bundles split into $\mathcal{O}(a)$'s, which are analogous to the successive minima.

For a smooth projective curve C, Beauville-Laszlo tells you that a vector bundle on C is equivalent to choosing a vector bundle on $C - \infty$, a vector bundle on $\widehat{\mathcal{O}_\infty}$, and an isomorphism of these two on $\operatorname{Frac}\left(\widehat{\mathcal{O}_\infty}\right)$. Vector bundles don't typically split, so we use the slopes coming from the Harder-Narasimhan filtration instead.

Then, the quantities $N(\alpha)$ and $N_s(\alpha)$ can be expressed in terms of global sections of vector bundles, which can be analyzed and compared using these slopes and Riemann-Roch.

(End of talk) Final remarks

Thanks for listening/reading(?), and hopefully I've kind of convinced you that the circle method can be done with just algebraic geometry and no analytic number theory. There are also many further applications and variants: singularities, cohomology, different targets like complete intersections, and different sources like higher-dimensional things.

3 (Extra stuff) Application to Fujita invariants

Fix a smooth projective variety $X/\mathbb{C}$. For a subvariety V, let $Y \to V$ be a resolution of singularities.

Let K_X, K_Y be the canonical divisors. The Fujita invariant

$$a(V) := \min \{t \in \mathbb{R} : t\left[-K_X|_Y\right] + \left[K_Y\right] \text{ is pseudo-effective}\}.$$

Lehmann-Tanimoto showed for a smooth projective (weak) Fano variety, there is a proper closed subset of X that is the closure of all subvarieties V with $a(V)$ larger than $a(X)$. Moreover, any component of $\operatorname{Mor}(\mathbb{P}^1, X)$ parametrizing a curve not in V has the expected dimension.

Conversely, it is expected that the converse is true: a subvariety with larger a-value will contain families of rational curves with dimension higher than the expected dimension in X.

Corollary 8. Let X be a smooth hypersurface in $\mathbb{P}^n$ of degree d satisfying $n \geq 2^d(d-1)+1$. Then, if V is a proper subvariety of X, we have $a(V) < 1 = a(X)$.

The general strategy for the corollary is to show there is a smooth projective curve C such that $\dim \operatorname{Mor}_e(C, V) > \dim \operatorname{Mor}_{e'}(C, X)$, which is a contradiction; our main result gives upper bounds on the latter, and we always have lower bounds on the former.

The assumption on large Fujita invariant implies $K_Y - K_X|_Y$ is not pseudo-effective.

It is a hard theorem (Boucksom-Demailly-Păun-Peternell '04) that for a smooth projective variety $Y/\mathbb{C}$ that if a line bundle L is not pseudo-effective, then for a general point in Y, we can find a curve passing through Y such that the intersection of L and the curve is negative.

Applying this to $K_Y - K_X|_Y$ produces a curve that intersects negatively with $K_Y - K_X|_Y$, but we have little control over its degree. By passing to finite characteristic, replacing the curve with an Artin-Schreier cover to increase genus, and increasing the degree without changing genus by using the Frobenius (like in the bend-and-break lemmas), we can find a curve C that gives us our desired contradiction.