Claremont Colleges Scholarship @ Claremont

HMC Senior Theses

HMC Student Scholarship

2021

Towards Tropical Psi Classes

Jawahar Madan

Follow this and additional works at: https://scholarship.claremont.edu/hmc_theses

Part of the Algebraic Geometry Commons

Recommended Citation

Madan, Jawahar, "Towards Tropical Psi Classes" (2021). *HMC Senior Theses*. 252. https://scholarship.claremont.edu/hmc_theses/252

This Open Access Senior Thesis is brought to you for free and open access by the HMC Student Scholarship at Scholarship @ Claremont. It has been accepted for inclusion in HMC Senior Theses by an authorized administrator of Scholarship @ Claremont. For more information, please contact scholarship@cuc.claremont.edu.

Towards Tropical ψ Classes

Jawahar Madan

Dagan Karp, Advisor

Renzo Cavalieri, Reader



Department of Mathematics

July, 2021

Copyright © 2021 Jawahar Madan.

The author grants Harvey Mudd College and the Claremont Colleges Library the nonexclusive right to make this work available for noncommercial, educational purposes, provided that this copyright statement appears on the reproduced materials and notice is given that the copying is by permission of the author. To disseminate otherwise or to republish requires written permission from the author.

Abstract

To help the interested reader get their initial bearings, I present a survey of prerequisite topics for understanding the budding field of tropical Gromov-Witten theory. These include the language and methods of enumerative geometry, an introduction to tropical geometry and its relation to classical geometry, an exposition of toric varieties and their correspondence to polyhedral fans, an intuitive picture of bundles and Euler classes, and finally an introduction to the moduli spaces of n-pointed stable rational curves and their tropical counterparts.

Contents

Ał	Abstract				
Ac	cknowledgments	ix			
1	Enumerative Geometry	1			
	1.1 What is enumerative geometry?	1			
	1.2 Moduli spaces	6			
	1.3 Rational curves and Gromov-Witten theory	7			
2	Tropical Geometry	11			
	2.1 Into the (semi)ring	11			
	2.2 Tropical polynomials	12			
	2.3 Tropical curves	14			
	2.4 Tropicalizing classical objects	15			
3	Toric Varieties	19			
	3.1 Introduction	19			
	3.2 Polyhedral fans	20			
4	Bundles and the Euler class	25			
	4.1 Introduction to Bundles	25			
	4.2 The Euler class	26			
5	Moduli spaces of <i>n</i> -pointed curves	29			
	5.1 The moduli space of n distinct points	29			
	5.2 Moduli spaces of stable <i>n</i> -pointed curves	31			
	5.3 ψ classes	33			
6	Tropical $\overline{M}_{0,n}$	35			
	6.1 Tropical stable <i>n</i> -pointed curves	35			

7	Conclusion	39
Bi	bliography	41

List of Figures

1.1	Lines on a smooth cubic	2
1.2	Topology of the projective plane	3
2.1	Graph of a tropical cubic	13
2.2	A tropical conic	16
2.3	Illustration of tropical Bézout's theorem	17
2.4	Tropicalizing a variety	18
3.1	Polyhedral fan for \mathbb{P}^2	22
5.1	Compactification of $M_{0,n}$	33
6.1	Tropical vs classical stable curve	36
6.2	Moduli cone for tropical curves	37
6.3	$\overline{M}_{0,5}^{\text{trop}}$ cone complex	37

Acknowledgments

I'd like to thank everyone who contributed their time, energy, and flexibility to this project–notably Renzo Cavalieri for offering an initial research problem and agreeing to be reader, Jon Jacobsen for running remote thesis and being gracious with the deadlines, and my friends for being there when I wanted to gripe about it.

Most of all I'd like to thank my advisor, Dagan Karp, for all the time and energy he has invested in my mathematical career–but more importantly, for reminding us that a mathematical career, especially during a pandemic, is about our health first and our output second. This paper, in light of my initial expectations for it, felt like the end of my mathematical career–a final confirmation that mathematics research was not for me. Ironically, the self-knowledge it ultimately sparked for me may end up being the biggest rebuttal to that idea.

Chapter 1

Enumerative Geometry

The material in this chapter is largely my retelling of the first two chapters of *Enumerative Geometry and String Theory* by Sheldon Katz, which I'd like to acknowledge as an excellent introduction to the topic and source of examples.

1.1 What is enumerative geometry?

In short, enumerative geometry asks the question "how many geometric structures of a given type satisfy a given collection of geometric conditions?" For example, we could ask "How many points in the plane lie on each of two given lines?" (which we know to be 1) or "How many twisted cubics are contained in a general 3-dimensional degree-5 (quintic) hypersurface?" (the answer turns out to be 317,206,375). And of course there are some more down-to-earth (but far from easily proven!) results, such as the identification of precisely 27 straight lines on a 2-dimensional cubic hypersurface, as shown in figure 1.1.

However, care must be taken that the questions we ask are well-posed: they should have well-defined answers, and ideally meaningful ones. Specifically, we would like our answers to be finite numbers that are fully determined by the constraints we impose. This means paying close attention to dimensionality: questions such as "How many times do two degree-*k* hypersurfaces in \mathbb{R}^n intersect?" only have finite answers when when the dimension of the intersection is 0 (a collection of points); this means the codimensions of the intersecting objects must add to *n*. In the aforementioned case, the objects are two hypersurfaces (codimension 1), so the question is only well-posed if n = 2. But there is more subtlety to well-posedness.

2 Enumerative Geometry



Figure 1.1 The 27 straight lines on a smooth cubic surface (from madore.org/~david)

Consider the question "Given two lines in \mathbb{R}^2 , how many points lie on both of them?" This amounts to counting the solutions to the system

$$\begin{cases} a_0 + a_1 x_1 + a_2 x_2 = 0\\ b_0 + b_1 x_1 + b_2 x_2 = 0 \end{cases}$$

where the a_i and b_i are real-number constants defining our lines and x_1, x_2 are coordinates on \mathbb{R} . But as we know, the answer depends! While the number of solutions is 'usually' 1, if the lines are parallel, it is 0, and if the lines are the same, it is ∞ . To an enumerative geometer, this signifies that we should try to improve the question. First, we can simply require that the lines are distinct. The answer is then 0 or 1. We could require that the lines are not parallel, but this is a more stringent condition than not being coincident, and there is a more elegant solution: we can work in **projective space**.

Definition 1.1.1 (Real projective plane). The real projective plane, \mathbb{RP}^2 , is the set of nonzero 3-tuples of real numbers $\mathbb{R}^3 - \{0\}$ modulo the scaling relation $(x_0 : x_1 : x_2) \sim (\lambda x_0 : \lambda x_1 : \lambda x_2)$ for $\lambda \in \mathbb{R}$.

To intuitively understand this construction, we can draw a correspondence between $(x_0 : x_1 : x_2) \in \mathbb{RP}^2$ and $(\frac{x_1}{x_0}, \frac{x_2}{x_0}) \in \mathbb{R}^2$, though this correspondence is neither unique nor defined everywhere. It is surjective, however, suggesting that \mathbb{RP}^2 looks like \mathbb{R}^2 with 'extra stuff.' This 'extra stuff' is the subset of the projective plane where $x_0 = 0$, and intuitively corresponds to 'points at infinity': $\{(0:1:a) \in \mathbb{RP}^2 | a \in \mathbb{R}\} \cup (0:0:1) \in \mathbb{RP}^2$. As $\frac{x_2}{x_1}$ is the 'ratio between the *y* and *x* values' in our map to \mathbb{R}^2 , each of these points at infinity corresponds to a different 'direction' of infinity. The topology of \mathbb{RP}^2 is visualized in Figure 1.2, where the arrows with corresponding labels are identified. As we can see, \mathbb{RP}^2 is a closed shape (albeit not one we can



Figure 1.2 Topology of the real projective plane

physically realize in 3 dimensions). This is because it is a **compactification** of the plane, a way to prevent the points we care about (such as the intersection of two lines) from 'running off' to infinity (as it does when the lines are parallel). But back to our pair of lines. Going forward, we will refer to non-projective space as **affine** space.

To work with algebraic curves in projective space, we must **homogenize** their defining polynomials. If a degree *d* polynomial p(x, y) in affine space is to be written in projective coordinates x_0, x_1, x_2 , it must be done in such a way that $(x_0 : x_1 : x_2) \in Z(p)$ implies $(\lambda x_0 : \lambda x_1 : \lambda x_2) \in Z(p)$ for all λ . This means each constituent monomial must be of the same degree; the only way to do this consistent with our projective-affine correspondence is with the mapping

$$p(x, y) \mapsto x_0^d p\left(\frac{x_1}{x_0}, \frac{x_2}{x_0}\right).$$

Finally, we return to our earlier problem, which in projective space is that of

counting the solutions to

$$\begin{cases} a_0 x_0 + a_1 x_1 + a_2 x_2 = 0\\ b_0 x_0 + b_1 x_1 + b_2 x_2 = 0 \end{cases}$$

where $(a_0, a_1, a_2) \neq (\lambda b_0, \lambda b_1, \lambda b_2) \in \mathbb{R}^3$ for any $\lambda \in \mathbb{R}$. Parallel lines are those that have $\frac{a_2}{a_1} = \frac{b_2}{b_1}$, allowing us to cancel the latter two terms on the LHS of the equations. Whereas before, this would leave us with $a_0 - \frac{a_1}{b_1}b_0 = 0$, no solutions, we now have $(a_0 - \frac{a_1}{b_1}b_0)x_0 = 0 \implies x_0 = 0$, that is, the lines simply intersect once at infinity. So the answer to our question is now an unequivocal "1"!

We now pick up the pace and start generalizing these methods, constructions, and results. For our next question: how many times does a given degree-2 curve intersect a given line in \mathbb{RP}^2 ? Again, we count solutions, this time for a more complicated system; for convenience, we choose a simple equation for the line, as the isotropy of \mathbb{RP}^2 ensures there is no loss of generality:

$$\begin{cases} x_2 = 0\\ b_0 x_0^2 + b_1 x_0 x_1 + b_2 x_0 x_2 + b_3 x_1^2 + b_4 x_1 x_2 + b_5 x_2^2 = 0\\ \implies b_0 x_0^2 + b_1 x_0 x_1 + b_3 x_1^2 = 0 \end{cases}$$

Where we have substituted the first equation into the second. To count the number of solutions to this equation, we first assume that $x_0 \neq 0$. Then we can define $x \equiv \frac{x_1}{x_0}$ and de-homogenize the equation to obtain

$$b_0 + b_1 x + b_3 x^2 = 0,$$

and from the quadratic formula we have our answer:

$$\begin{cases} 0 & \Delta < 0 \\ 2 & \Delta > 0 , \quad \Delta = b_1^2 - 4b_0 b_3 \\ 1 & \Delta = 0 \end{cases}$$

this time, the answer is ill-defined due to the varying number of roots of a quadratic. But we know exactly how to fix that: complexification! At this point it stands to introduce the most common general setting for enumerative geometry: complex projective space, which is what people usually mean when they say projective space, and is exactly what it sounds like.

Definition 1.1.2 (Projective space). Projective space, \mathbb{P}^n (sometimes \mathbb{CP}^n), is the set of nonzero (n + 1)-tuples of complex numbers $\mathbb{C}^{n+1} - \{0\}$ modulo scaling, that is, $(x_0 : x_1 : \cdots : x_n) \sim (\lambda x_0 : \lambda x_1 : \cdots : \lambda x_n)$ for $\lambda \in \mathbb{C}$.

Like real projective space, complex projective space is compact, but it carries the major additional benefit of being algebraically closed. Despite the 2n real dimensions, geometers refer to \mathbb{P}^2 as a plane, 2-real-dimensional hypersurfaces within it as embeddings of the "complex line," and so on.

With this technology, we find that a degree-2 curve intersects a line precisely twice—with one caveat. The caveat is that even over the complex numbers, a quadratic may have only a single root, if that root has multiplicity 2. So we must count intersections "with multiplicity."

We are finally equipped to understand a foundational result in enumerative geometry–perhaps the main reason we work in projective space at all:

Theorem (Bézout's theorem). Let $C \subset \mathbb{P}^2$ be a plane curve of degree c, and $D \subset \mathbb{P}^2$ be a plane curve of degree d. If $|C \cap D|$ is finite, then

$$\sum_{p \in C \cap D} m_p(C \cdot D) = cd$$

where m_p denotes the multiplicity of the intersection at p. To demystify this, first note that the number of intersection points $|C \cap D|$ is finite if and only if *C* and *D* have no shared factors; this is analogous to our requirement of distinct lines (this non-coincidence condition is often referred to as the curves being in **general position**). In lieu of proving the theorem, here is a brief sketch: in projective space, we can smoothly vary the coefficients of the polynomials defining our curves without changing the number of intersections; this is because we have both algebraic closure (ensuring that the equations we obtain in solving for intersections will not be 'missing' any roots) and analytic compactness (solutions cannot 'escape to infinity' as they do in the case of parallel lines in affine space). Given this fact, we can smoothly vary the coefficients of the polynomial defining C until it factors into *c* linear factors, that is, until *C* consists of *c* distinct lines; similarly, we can vary D until it consists of d lines, distinct from each other and from those comprising C. Since every line in C intersects precisely once, with multiplicity 1 with each line of *D*, the number of intersections is *cd*, and must have been *cd* before we deformed the curves as well.

In fact, Bézout's theorem generalizes: if $X_1, X_2, ..., X_n$ are *n* hypersurfaces in \mathbb{P}^n of degrees $d_1, d_2, ..., d_n$ (note the number of surfaces must match

the dimension of projective space), then if there is no factor shared by all of them (so they have finite intersection), we have

$$\sum_{p \in X_1 \cap \dots \cap X_n} m_p(X_1 \cdot X_2 \cdots X_n) = d_1 \cdot d_2 \cdots d_n$$

1.2 Moduli spaces

Bézout's theorem is neat, but its utility may seem limited; all it seems to do is count points. However, this is where the methods of enumerative geometry really start to shine: no matter what we are trying to count, the objects in question will be points at the intersection of various constraint surfaces in an appropriate **moduli space** of all such objects.

Starting with an easy example, suppose we want to determine how many lines in the plane \mathbb{P}^2 contain both of two given points: $P = (p_0 : p_1 : p_2)$ and $Q = (q_0 : q_1 : q_2)$. The enumerative method is to consider the space of all lines, and impose our constraints geometrically on that space. A line is defined by the equation $a_0x_0 + a_1x_1 + a_2x_2 = 0$, for some nonzero $(a_0, a_1, a_2) \in \mathbb{C}^3$; moreover, the tuple $(\lambda a_0, \lambda a_1, \lambda a_2)$ for any $\lambda \in \mathbb{C}$ defines the same line. But this is precisely the definition of \mathbb{P}^2 ; that is, the moduli space of lines is isomorphic to the projective plane. In fact, it is often called the **dual plane**, $(\mathbb{P}^2)^*$. We now impose constraints. Lines passing through *P* satisfy $a_0p_0 + a_1p_1 + a_2p_2 = 0$; this defines a line L_P in $(\mathbb{P}^2)^*$! A line passes through *P* and *Q* if and only if its tuple of coefficients lies on both L_P and L_Q in $(\mathbb{P}^2)^*$. So by Bézout's theorem, we know there is precisely one such line!

Similarly, a degree *d* plane curve is defined by a nonzero tuple of coefficients modulo scaling. The number of coefficients is equal to the number of monomial terms, which is equal to the number of partitions of *d* into three 'buckets.' The 'stars and bars' method tells us that this is equal to $\binom{d+2}{2} = \frac{(d+2)(d+1)}{2}$, so the moduli space of degree *d* plane curves is isomorphic to $\mathbb{P}^{\frac{(d+2)(d+1)}{2}-1} = \mathbb{P}^{\frac{d(d+3)}{2}}$. Since polynomials are linear in their coefficients, the requirement that a degree *d* curve passes through a given point $P \in \mathbb{P}^2$ is a linear constraint on the aforementioned moduli space; that is, it defines a hyperplane in $\mathbb{P}^{\frac{d(d+3)}{2}}$. By Bézout's theorem, $\frac{d(d+3)}{2}$ hyperplanes in this space will intersect in precisely $1 \cdot 1 \cdots 1 = 1$ point, so we can declare that there is precisely 1 degree *d* plane curve through a given set of $\frac{d(d+3)}{2}$ points.

Bézout's theorem is extra useful when we have nonlinear constraints. Recall that a conic (degree 2 curve), given by $(b_0 : b_1 : b_2 : b_3 : b_4 : b_5) \in \mathbb{P}^5$, intersects the line $x_2 = 0$ twice, except if $b_1^2 - 4b_0b_3 = 0$, in which case it intersects once with multiplicity 2. It seems that tangency–the case when there is a multiplicity 2 intersection–imposes a constraint quadratic in the coefficients of the conic. And indeed, the set of conics tangent to a given line form a degree-2 hypersurface (that is, a conic) in \mathbb{P}^5 . Applying Bézout, we may conclude that there are $2^n \cdot 1^{5-n}$ conics tangent to *n* given lines and passing through 5 - n given points. However, this is not quite true.

The reason is that we fail to filter out 'degenerate' conics: our moduli space includes *reducible* conics, which factor into pairs of lines. If those lines coincide, such a conic is a **double line**, and intersects any other line with multiplicity 2. So we have an issue: our moduli space includes objects exhibiting forms of 'tangency' we may not be interested in. In this case, since every double line is 'tangent' to every other line, the entire space of double lines is contained within all of our tangency-constraint conic hypersurfaces. Problems like these certainly give mathematicians a headache, but there exist various tools to overcome them; in this case, we can use **excess intersection theory**: we identify the moduli space of double lines (isomorphic to the moduli space of lines) and count the appropriate constraint intersections to determine how many 'extra' solutions we originally had.

While excess intersection theory works in a pinch, a more sophisticated approach is often to work with moduli spaces more finely tailored to the problem. The primary downside to this is that our standard machinery, such as Bézout's theorem, may not function in these spaces.

1.3 Rational curves and Gromov-Witten theory

While algebraic and enumerative geometry concern themselves with algebraic curves in general, Gromov-Witten theory focuses on a particular family: **rational curves**.

Definition 1.3.1 (Rational curve). A rational curve is an algebraic curve that can be parameterized by \mathbb{P}^1 . That is, $C \in \mathbb{P}^n$ is rational if

 $C = \{(g_0(x_0:x_1), g_1(x_0:x_1), \dots, g_n(x_0:x_1)) \mid (x_0:x_1) \in \mathbb{P}^1\}$

for some polynomials g_0, \ldots, g_n .

As it turns out, all lines are rational, as well as general conics. However, degree d > 2 curves are rational only in special cases. In fact, it turns out *C* is rational if and only if *C* has genus zero.

d	3d - 1	N_d	
1	2	1	
2	5	1	
3	8	12	
4	11	620	19th century (Zeuthen)
5	14	87304	early 1990s (Ran & Vainsencher)
6	17	26312976	1994 (Kontsevich)
7	20	14616808192	1994 (Kontsevich)

Table 1.1 Number of rational curves passing through 3d - 1 points

To work with rational curves, geometers set up some heavy-duty moduli spaces (the simplest of which I explain later in this paper). The space of degree *d* rational curves in \mathbb{P}^n is denoted $M(\mathbb{P}^n, d)$, and is 3d - 1 dimensional. How many degree *d* rational curves pass through 3d - 1 points? The question may seem to lend itself to a simple Bézout's theorem style argument–intersect 3d - 1 hypersurfaces, multiply their degrees, or something–but Bézout's theorem does not hold in $M(\mathbb{P}^n, d)$, and as *d* increases, the question becomes incredibly difficult, stumping mathematicians for hundreds of years at each increment. However, we now have the answer for all values of *d* (Table 1.1) thanks to the work of Kontsevich, who discovered the following recursive formula:

$$N_{d} = \sum_{d_{A}+d_{B}=d} N_{d_{A}} N_{d_{B}} d_{A}^{2} d_{B} \left(d_{B} \begin{pmatrix} 3d-4\\ 3d_{A}-2 \end{pmatrix} - d_{A} \begin{pmatrix} 3d-4\\ 3d_{A}-1 \end{pmatrix} \right)$$

I'm not going to attempt to explain here how the formula was derived, but rather present it as an example of the extraordinary success of Gromov-Witten theory.

But what is Gromov-Witten theory? In short, it is a framework for computing enumerative quantities of interest in the form of **topological invariants**. The framework is as follows. We compactify the space $M(X, \beta)$ of rational curves of class β in X to form the moduli space of **stable maps** of genus g and n marked points $\overline{M}_{g,n}(X, \beta)$. Using mappings from $\overline{M}_{g,n}(X, \beta)$ to X or other moduli spaces, we can pass (with pullbacks and pushforwards) algebraic-topological structures (homology and cohomology classes) from one moduli space to another. The topological invariants of interest are then encoded in these classes (for example, as the winding number of a certain homology class around a hole in some manifold). And just like that, we

obtain some of the most historically sought-after enumerative results, which even have applications to string theory.

An exciting recent development in Gromov-Witten theory is the transposition of many constructions and results from the classical context into tropical geometry, such as by Cavlieri, Gross, and Markwig in their upcoming paper "Tropical ψ Classes." Setting up the background material for this subfield is the goal of this thesis; we must next get acquainted with tropical geometry.

Chapter 2

Tropical Geometry

This chapter closely follows the article "Tropical Mathematics" by Speyer and Sturmfels, a wonderfully concise first foray into the subject.

2.1 Into the (semi)ring

In recent years, the tropical approach to algebraic geometry has matured from a novel subject of curiosity into an active and productive field. Tropical geometry studies the **tropical semiring**, which has a few isomorphic definitions–we will use the "min-plus" convention, which defines the semiring as follows:

Definition 2.1.1 (Tropical semiring). The tropical semiring $(\mathbb{R} \cup \{\infty\}, \oplus, \odot)$ is the set of real numbers along with 'positive infinity,' with addition defined as $x \oplus y \equiv \min(x, y)$ and multiplication defined as $x \odot y \equiv x + y$.

For example, $3 \oplus 5 = 3$, while $3 \odot 5 = 8$. As we can see, tropical addition and tropical multiplication are commutative, and the distributive law continues to hold:

 $x \odot (y \oplus z) = x + \min(y, z) = \min(x + y, x + z) = x \odot y \oplus x \odot z$

where in the last step we assume the conventional multiplication-thenaddition order of operations.

There are some features of the tropical semiring worth noting at the outset: in this semiring, 0 and 1 are not the additive and multiplicative identities; rather, ∞ is the additive identity (min(x, ∞) = x for all x) and 0 is the multiplicative identity (since tropical multiplication is classical

12 Tropical Geometry

addition). Note also that while elements have multiplicative inverses (their negations) they do not have additive inverses (we would need $\infty = x \oplus -x = \min(x, -x) \le x < \infty$). This is why the tropical semiring fails to be a ring.

Let us now do some algebra over this semiring. We quickly encounter some big deviations from classical results. For example, in the tropical setting,

$$(x \oplus y)^{n} \equiv \bigcirc_{k=1}^{n} (x \oplus y)$$

= $\bigoplus_{k=0}^{n} \bigoplus_{\ell=1}^{\binom{n}{k}} x^{k} y^{n-k}$
= $\bigoplus_{k=0}^{n} \min(x^{k} y^{n-k}, x^{k} y^{n-k}, \dots, x^{k} y^{n-k})$
= $\bigoplus_{k=0}^{n} x^{k} y^{n-k}$
= $\min(y^{n}, x y^{n-1}, \dots, x^{n})$
= $\min(x^{n}, y^{n})$
= $x^{n} \oplus y^{n}$

So the "freshman's dream" is unequivocally true in tropical arithmetic. Note that we are careful not to denote repeated addition as multiplication, since in the tropical context it is not.

We next look at **tropical polynomials**, since (as in classical algebraic geometry) they are the objects we are most interested in.

2.2 Tropical polynomials

Let x_1, \ldots, x_n be variables representing elements in the tropical semiring. A tropical monomial is a product of these variables, $x_1^{i_1}x_2^{i_2}\cdots x_n^{i_n}$, where $i_1, \ldots, i_n \in \mathbb{Z}$ (negative powers allowed since tropical elements have multiplicative inverses). Written in terms of classical operations, this is

$$i_1x_1+i_2x_2+\cdots+i_nx_n,$$

a linear function with integer coefficients. Different tropical monomials encode different linear combinations of variables, and any integer-coefficient linear combination corresponds to the tropical monomial with appropriate powers of each variable. So there is a one-to-one correspondence between tropical monomials and integer-coefficient linear functions. We are now equipped to understand tropical polynomials. **Definition 2.2.1** (Tropical polynomial). A tropical polynomial *p* is a finite linear combination of tropical monomials:

$$p(x_1,\ldots,x_n) = a \odot x_1^{i_1} \cdots x_n^{i_n} \oplus b \odot x_1^{j_1} \cdots x_n^{j_n} \oplus \cdots$$
$$= \min(a + i_1 x_1 + \cdots + i_n x_n, b + j_1 x_1 + \cdots + j_n x_n, \ldots).$$

So tropical polynomials are functions returning the minimum of a certain collection of linear functions. Some thought leads us to observe that these polynomials are continuous and piecewise-linear, as well as the more esoteric property that they are concave-down; that is,

$$p\left(\frac{\mathbf{x}+\mathbf{y}}{2}\right) \ge \frac{1}{2}(p(\mathbf{x})+p(\mathbf{y}))$$

for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$. Essentially, this says is that if we draw a line between two points ($\mathbf{x}, p(\mathbf{x})$ and ($\mathbf{y}, p(\mathbf{y})$, on a graph of p, points on the graph in between cannot be below the line. All this is easier to see with a visual example:



Figure 2.1 The graph of a cubic polynomial, with roots labeled (from "Tropical Mathematics")

Figure 2.1 is the graph of a cubic polynomial in one variable, $p(x) = ax^3 \oplus bx^2 \oplus cx \oplus d$. As we can see, it is the lower envelope of the four linear functions d, c + x, b + 2x, and a + 3x. All four lines contribute to this lower envelope if and only if the values at which these lines intersect with the next-steepest one occur in the right order: $b - a \le c - b \le d - c$. These points

are very important: in fact, if the aforementioned inequalities hold, the cubic can be factored into

$$p(x) = a \odot (x \oplus (b-a)) \odot (x \oplus (c-b)) \odot (x \oplus (d-c)).$$

To see this, consider each region separately: when x < b - a, all the tropical sums evaluate to x and we recover the cubic term, which is the linear function matching p in this interval. When b - a < x < c - b, the first sum evaluates to b - a, producing a coefficient of b and the quadratic term that contributes to the envelope in this interval. And so on. It may seem that we can read this factorization directly off the graph, and this is true: every tropical polynomial function can be written uniquely as such a tropical product of tropical linear functions. This is akin to the Fundamental Theorem of Algebra, and suggests that we consider the constants in these linear factors to be the **roots** of the tropical polynomial. In other words,

Definition 2.2.2 (Root of a tropical polynomial). $\mathbf{x} \in \mathbb{R}^n$ is the root of an *n*-variable tropical polynomial *p* if and only if *p* fails to be linear at \mathbf{x} .

These 'corner' points of a tropical graph may seem wildly different from the classical understanding of roots as zeroes, but we will see that they satisfy many classical properties. Before we move on, we must make the important note that these factorizations are based on the actual function encoded by a tropical polynomial, not the tropical polynomial itself. The distinction lies in the fact that if we add to a polynomial a monomial that is nowhere the minimum monomial, it does not contribute to the lower envelope, does not affect the roots, and therefore cannot affect the factorization of the tropical function. So these factorizations do not allow recovery of the original polynomials, only the functions they represent.

2.3 Tropical curves

We now arrive at the subject of **tropical curves**; that is, curves in the algebraic geometry sense, defined as the set of roots of a polynomial. As we will see, there is nothing much curvy about them.

Definition 2.3.1 (Tropical hypersurface). Given a tropical polynomial p : $\mathbb{R}^n \to \mathbb{R}$, its hypersurface $\mathcal{H}(p)$ is the set of points $\mathbf{x} \in \mathbb{R}^n$ such that \mathbf{x} is a root of p.

This is analogous to the classical concept of the hypersurface in \mathbb{R}^n defined by the zero locus of a function of *n* variables. What do these tropical hypersurfaces look like? In short, they look like lattices of polyhedra; the roots of a tropical polynomial are precisely the boundaries of the polyhedral domain regions where each of the constituent monomials is the minimum one.

To understand the geometry of these objects, the best starting point is to look at the two-dimensional case, where our polynomial is a function of two variables. If the polynomial is linear, that is, $p(x, y) = a \odot x \oplus b \odot y \oplus c$, we have the tropical line (degree 1 hypersurface) $\mathcal{H}(p) = \{(x, y) \in \mathbb{R}^2 \mid \min(a + x, b + y, c) \text{ is not linear}\} = \{(x, y) \in \mathbb{R}^2 \mid a + x = b + y < c, a + x = c < b + y, \text{ or } b + y = c < a + x\}$, which is a set of three half-rays: y = x + a - bextending in the (-1, -1) direction, y = c - b in the (1, 0) direction, and x = c - a in the (0, 1) direction. The rays meet at the point (c - a, c - b). This three-ray shape does not look like our intuitive notion of a line, but it has linear properties nonetheless: for example, since all tropical lines are just translations of this one, we can convince ourselves that two tropical lines intersect in precisely one point.

Next, consider a general tropical quadratic:

 $p(x, y) = a \odot x^2 \oplus b \odot xy \oplus c \odot y^2 \oplus d \odot y \oplus e \oplus f \odot x$

where a, \ldots, f are constant coefficients. What does its hypersurface (a 'tropical conic') look like? Again, it depends on the roots being 'in the right order,' but projecting the lower envelope of p onto the plane gives us a lattice of polyhedra (Figure 2.2), with a region for each monomial; the divisions between them have slopes $(0, \infty \text{ or } 1)$ corresponding to the ratio of y and x coefficients in the differences of the monomials. Moreover, slicing this curve along the x or y directions (fixing the value of x or y) produces a tropical quadratic surface in one variable, that is, a set of two points.

Again, despite the curve looking somewhat alien, it retains classical properties of conic sections. Specifically, a tropical line will intersect it precisely twice, and another tropical conic will intersect it four times (Figure 2.3). This brings us to a key fact: tropical hypersurfaces satisfy Bézout's theorem!

2.4 Tropicalizing classical objects

In the last section, we saw that the enumerative geometry of tropical varieties matches that of classical varieties, at least in some key ways. This



Figure 2.2 The graph and curve defined by a tropical quadratic (from Maclagan and Sturmfels' *Introduction to Tropical Geometry*)

means we can meaningfully 'tropicalize' classical enumerative problems, attacking them using skeletonized, combinatorial structures rather than the complexities of classical geometry.

Tropicalizing classical objects is a surprisingly confusing process: because not all classical structure can be imported to the tropical context, tropicalization is less a matter of using a one-size-fits-all method, and more of a matter of understanding a classical object deeply enough to determine the most natural way to skeletonize it. Even the concrete, general procedures that exist are presented very differently from different perspectives. Nevertheless, it is worth understanding what general ideas we can.

Let us first consider polynomials. Tropicalizing a polynomial is *almost* as simple as replacing the classical operations in its definition with tropical operations. We must also replace the constants with their counterparts in the tropical semiring, which are called **valuations** of those constants. Valuations are the images of ring elements under homomorphisms that satisfy certain properties; for our purposes we will only worry about the "log-limit" valuation $v : \mathbb{C} \to \mathbb{R} = \lim_{t\to 0} v_t : \mathbb{C} \to \mathbb{R}$ where $v_t(z) = \log_t |z|$. While it seems like this sends everything to 0, observe that this is only the case for t = 0; as t approaches 0 the map is still injective. This construction



Figure 2.3 A fun illustration of Bézout's theorem for tropical conics (from the cover of Mathematics Magazine, June 2009)

maps the classical operations to tropical ones: the logarithms give us the property

$$v_t(z_1 z_2) = \log_t |z_1 z_2| = \log_t |z_1| |z_2| = \log_t |z_1| + \log_t |z_2| = v_t(z_1) + v_t(z_2)$$

and taking the limit gives us the property $v(z_1 + z_2) = \min(v(z_1), v(z_2))$. To see the latter, take $w_1, w_2 \in \mathbb{R}$ with $w_1 \le w_2$. Then for any $z_1, z_2 \in \mathbb{C}$ with $v(z_1) = w_1, v(z_2) = w_2$, there are $\theta_1, \theta_2 \in [0, 2\pi)$ such that $z_1 = t^{w_1} e^{i\theta_1}, z_2 = t^{w_2} e^{i\theta_2}$. We then have

$$\begin{aligned} v(z_1 + z_2) &= \lim_{t \to 0} \left(\log_t |t^{w_1} e^{i\theta_1} + t^{w_2} e^{i\theta_2}| \right) \\ &= \lim_{t \to 0} \left(\log_t |t^{w_1} e^{i\theta_1}| + \log_t |1 + t^{w_2 - w_1} e^{i(\theta_2 - \theta_1)}| \right) \\ &= \lim_{t \to 0} \left(w_1 + \frac{\ln |1 + t^{w_2 - w_1} e^{i(\theta_2 - \theta_1)}|}{\ln t} \right) \\ &= w_1 \\ &= \min(w_1, w_2) = \min(v(z_1), v(z_2)) \end{aligned}$$

where the second term disappears because $\lim t^{w_2-w_1} = 0$, or 1 if $w_2 = w_1$, so the numerator is zero or order 1, while the denominator goes to $-\infty$. Note that while this process seems to produce infinitesimally small constants, the idea is to make sure the ratios between different constants and interactions between constants in a polynomial are properly ported to the tropical setting; we can always rescale the polynomial as a whole without changing its associated curve. We can now tropicalize polynomials!

Of course, in algebraic geometry, the form of tropicalization of most interest to us is the tropicalization of varieties (the curves defined by polynomials): given a classical variety defined by some polynomial, we can transform it to the **tropical variety** defined by the tropicalization of the same polynomial.

For example, the tropical quadratic $2 \odot x \odot y \oplus x \oplus y \oplus 1 = \min(2 + x + y, x, y, 1)$, defining the tropical variety shown in Figure 2.4, is the result of tropicalizing the conic section (in this case a hyperbola over the reals) defined by the classical polynomial $t^2xy + x + y - t$. Here *t* is the base of the valuation map–which we are supposedly taking to 0. What's going on? The fact is that there is some sleight of hand scaling done in making this correspondence–the limit *will* shrink a tropical variety to a fan around the origin for a given classical variety, and a tropical variety with finite nonzero constants will correspond to a degenerate limit of the associated classical variety. For this and other reasons, a more modern approach to tropicalization uses what are called Puiseux series, but the intuition for it is somewhat more involved.



Figure 2.4 A tropical variety and its classical counterpart (over the reals)

Chapter 3

Toric Varieties

3.1 Introduction

The goal of this project is to set up tropical Gromov-Witten theory; specifically on **toric varieties**, so it is worth understanding what they are.

Definition 3.1.1 (Toric variety). A toric variety *X* is a complex algebraic variety containing the algebraic torus $T = (\mathbb{C}^*)^r$ as a dense open set, equipped with an action of *T* on *X* whose restriction to $T \subset X$ is the usual multiplication on *T*.

By 'the usual multiplication on *T*' we mean the group product on $(\mathbb{C}^*)^r$, coordinate-wise multiplication.

For example, \mathbb{P}^2 is a toric variety. To see this, note first that the subset $\{(x_0, x_1, x_2) \in \mathbb{P}^2 | x_0 \neq 0, x_1 \neq 0, x_2 \neq 0\}$ is isomorphic to $(\mathbb{C}^*)^2$ under the mapping $(x_0, x_1, x_2) \mapsto (\frac{x_1}{x_0}, \frac{x_2}{x_0})$. There is an action of $T_2 = (\mathbb{C}^*)^2$ on \mathbb{P}^2 given by $(t_1, t_2) \cdot (x_0, x_1, x_2) = (x_0, t_1x_1, t_2x_2)$, and one can check that this action restricts to the usual algebraic torus product on $T_2 \subset \mathbb{P}^2$.

Importantly, this action is transitive on $T_2 \subset \mathbb{P}^2$, but not on \mathbb{P}^2 as a whole. Instead, there are various invariant subsets of \mathbb{P}^2 on which the action is transitive: specifically, the subsets $\{(x_0, x_1, 0) \mid x_0, x_1 \in \mathbb{C}^*\}$, $\{(x_0, 0, x_2) \mid x_0, x_2 \in \mathbb{C}^*\}$, and $\{(0, x_1, x_2) \mid x_1, x_2 \in \mathbb{C}^*\}$ along with the points (1, 0, 0), (0, 1, 0), and (0, 0, 1). To see this, note that the torus action lets us scale the ratios between our homogenous coordinates at will, but we cannot make them zero or ∞ .

3.2 Polyhedral fans

This section follows Chapter 7 of *Mirror Symmetry*, which is a pretty imposing book; I've done my best to translate it into words I find intuitive.

Polyhedral fans are a powerful tool for constructing and encoding toric varieties. We will define them in a moment, but first need the notion of a **rational polyhedral cone**:

Definition 3.2.1 (Strongly convex rational polyhedral cone). Let *N* be a lattice of rank *r*, and let $N_{\mathbb{R}} = N \otimes \mathbb{R}$ (\mathbb{R} -linear combinations of points of *N*, which just gives us the vector space spanned by elements of *N*). A strongly convex rational polyhedral cone $\sigma \subset N_{\mathbb{R}}$ is a set

 $\sigma = \{a_1v_1 + a_2v_2 + \dots + a_kv_k \mid a_i \ge 0\}$

generated by a finite set of vectors v_1, \ldots, v_k in N such that $\sigma \cup (-\sigma) = \{0\}$.

Intuitively, we are considering the cone of all positive linear combinations of a certain collection of lattice vectors such that said cone does not nontrivially intersect with the cone in the opposite direction (all negative linear combinations of the same vectors). We will refer to these objects as simply 'cones.' We then have the following definition:

Definition 3.2.2 (Polyhedral fan). A polyhedral fan is a collection Σ of cones in $N_{\mathbb{R}}$ such that each face of a cone in Σ is also a cone in Σ , and the intersection of two cones in Σ is a face of each.

Intuitively, this means that the cones of Σ do not 'overlap,' and their faces, and faces' faces, and so on, are also cones in Σ . We now describe how fans encode toric varieties.

To construct the toric variety X_{Σ} associated to a fan Σ , we will construct a subspace $Z(\Sigma) \in bC^n$ and a group *G* from Σ ; we will then have $X_{\Sigma} = (\mathbb{C}^n - Z(\Sigma))/G$.

First, $Z(\Sigma)$. For a fan Σ in $N_{\mathbb{R}}$ let $\Sigma(1)$ be the set of one-dimensional cones of Σ . For each $\rho \in \Sigma(1)$, there is a certain $v_{\rho} \in N_{\mathbb{R}}$ which generates $\rho \cap N$ (intuitively, the shortest lattice vector in ρ). Let $n = |\Sigma(1)|$. Now, associate a coordinate x_{ρ} to each 1-cone ρ . If *S* is a subset of $\Sigma(1)$ that does not span a cone of Σ , let $V(S) \subset \mathbb{C}^n$ be the subspace of \mathbb{C}^n defined by $x_{\rho} = 0$ for all $\rho \in S$. $Z(\Sigma)$ is then the union of V(S) for all *S*.

Next, we construct *G*. This takes several steps.

First, let *M* be the dual lattice of *N*, that is, the subset of linear functionals on the vector space $N_{\mathbb{R}}$ that take integer values on all elements of *N*. Note

that $M_{\mathbb{R}} = (N_{\mathbb{R}})^*$. Next, consider $\operatorname{Hom}(M, \mathbb{C}^*)$. A map in this space can be identified with a point of $(\mathbb{C}^*)^r$, since M is a rank-r lattice. We also need $\operatorname{Hom}(\Sigma(1), \mathbb{C}^*)$, which we can identify with $(\mathbb{C}^*)^n$, as we are treating $\Sigma(1)$ as just a set, so a map in this space simply specifies the image of each of the n 1-cones. We now define the map $\phi : \operatorname{Hom}(\Sigma(1), \mathbb{C}^*) \to \operatorname{Hom}(M, \mathbb{C}^*)$, which given $f \in \operatorname{Hom}(\Sigma(1), \mathbb{C}^*)$ produces $g \in \operatorname{Hom}(M, \mathbb{C}^*)$ defined by $g(m) = \prod_{\rho \in \Sigma(1)} f(\rho)^{m(v_{\rho})}$. In coordinates, we have $f = (t_1, \ldots, t_n) \in (\mathbb{C}^*)^n$ (where $t_i = f(\rho_i)$), and

$$\phi(f) = \phi(t_1, \dots, t_n) = \left(\prod_{i=1}^n t_i^{v_{\rho_i}^1}, \dots, \prod_{i=1}^n t_i^{v_{\rho_i}^r}\right)$$

where $v_{\rho_i}^j$ is the j^{th} component of v_{ρ_i} .

There is a lot to unpack here. ϕ essentially takes a given weighting of the 1-cones and assigns values to the elements of *m* based on this weighting. The key insight is that *M* is an additive lattice, while \mathbb{C}^* is a multiplicative group, so we need $\phi(f)(m + \ell) = \phi(f)(m)\phi(f)(\ell)$. The components of *m* and ℓ , that is, $m(v_{\rho_1}), m(v_{\rho_2}), \ldots, m(v_{\rho_n})$ and likewise for ℓ must therefore be exponentiated in the formulas for $\phi(f)(m)$ and $\phi(f)(\ell)$ so that they add together in the product of said expressions. We could exponentiate all the components together, with the same base, as in $e^{\sum_{i=1}^n m(v_{\rho_i})}$, but this destroys our ability to distinguish different elements of *M* which have the same sum of their components. Instead, we pick a set of *n* different bases, one for each component, which are the t_i . Picking a set of t_i is equivalent to picking an $f \in \text{Hom}(\Sigma(1), \mathbb{C}^*)$, which leads us to create ϕ .

Finally, we can define $G = \text{Ker}(\phi)$. G is the set of weightings $\{t_i\}$ of the $\{\rho_i\}$ that satisfy $\prod_{i=1}^{n} t_i^{v_{\rho_i}^j} = 1$ for all j, equipped with the operation of pointwise multiplication. To see that this forms a group, note that the choice $t_i = 1$ for all i has this property, for any set $\{t_i\}$ with this property $\{\frac{1}{t_i}\}$ also has the property (all the factors in the product are simply inverted), and for $\{t_i\}, \{s_i\}$ with this property, $\{t_is_i\}$ works too, because

$$\prod_{i=1}^{n} (t_i s_i)^{v_{\rho_i}^j} = \prod_{i=1}^{n} t_i^{v_{\rho_i}^j} s_i^{v_{\rho_i}^j} = \left(\prod_{i=1}^{n} t_i^{v_{\rho_i}^j}\right) \left(\prod_{i=1}^{n} s_i^{v_{\rho_i}^j}\right) = 1 \cdot 1 = 1.$$

At this point it is convenient to identify $\text{Hom}(\Sigma(1), \mathbb{C}^*)$ with $(\mathbb{C}^*)^n$, so that we can view *G* as a subgroup of $(\mathbb{C}^*)^n$. We now recall $\mathbb{C}^n - Z(\Sigma)$, which in intuitive language is \mathbb{C}^n under the condition that coordinates can only

be simultaneously zero if their corresponding v_{ρ_i} span a cone of Σ . Finally, we quotient this space by G to obtain $X_{\Sigma} = (\mathbb{C}^n - Z(\Sigma))/G$. Note that this works because the action of an element of G on \mathbb{C}^n (the natural extension of its action on $(\mathbb{C}^*)^n$) cannot set a nonzero coordinate of any point to zero and thus cannot map a point in $\mathbb{C}^n - Z(\Sigma)$ into $Z(\Sigma)$.

There are a number of interesting facts about X_{Σ} I'm still trying to wrap my mind around. The first thing to note is that $(\mathbb{C}^*)^n/G \subset X_{\Sigma}$ is an algebraic torus dense in X_{Σ} and able to act on it by the usual coordinatewise multiplication, confirming that X_{Σ} is a toric variety. Second, it turns out that X_{Σ} has the same rank as N. Third, X_{Σ} is compact if and only if the union of the cones of Σ equals the whole of $N_{\mathbb{R}}$. This has to do with the fact that each 'missing piece' contributes a closed set to $Z(\Sigma)$ which is then excised from \mathbb{C}^n , which introduces 'additional openness' in ways that break compactness even after quotienting by G.

We now come to a neat fact about X_{Σ} . All of its torus-action-invariant subvarieties are of the form $\{x \in X_{\Sigma} | x_{i_1} = x_{i_2} = \cdots = x_{i_k} = 0\}$ for some i_1, \ldots, i_k (recall the invariant subsets of \mathbb{P}^2 at the start of the chapter, which took the same form). For any given i_1, \ldots, i_k , we have the 1-cones $\rho_{i_1}, \ldots, \rho_{i_k}$. It turns out that any such invariant subvariety will correspond to a set of 1-cones that form a valid convex cone, and since we have preemptively removed from X_{Σ} such subvarieties where the cone is not included in Σ , we have a clean **order-reversing correspondence** between the cones of Σ and the torus-invariant subvarieties, which are themselves toric varieties: simply quotient *N* by the sublattice $\sigma \cap N$ (to 'forget about' the coordinates which are fixed at zero), and our new fan is the projection of Σ onto this quotient lattice.



Figure 3.1 The fan for \mathbb{P}^2

We conclude the chapter with an example: we will construct \mathbb{P}^2 from a fan. The fan in question, shown in Figure 3.1, sits in a two-dimensional lattice, so will produce a two-dimensional variety. The set of cones, not explicitly listed in the diagram, are {0}, the three 1-cones spanned by each of the edges, and the three 2-cones spanned by pairs of adjacent 1-cones. This means the only set of 1-cones not spanning a cone is the set of all three, which in turn means $Z(\Sigma) = \{(0,0,0)\} \subset \mathbb{C}^3$. We next identify *G*. Ordering the edges (-1, -1), (1, 0), (0, 1) gives us the components to explicitly write down $\phi(f)$ for $f \in \text{Hom}(\Sigma(1), \mathbb{C}^*)$. We then have

$$G = \{ f \in \operatorname{Hom}(\Sigma(1), \mathbb{C}^*) \mid \phi(f) = 1 \in \operatorname{Hom}(M, \mathbb{C}^*) \}$$

= $\{ (t_1, t_2, t_3) \in (\mathbb{C}^*)^3 \mid \phi(t_1, t_2, t_3) = (1, 1) \in (\mathbb{C}^*)^2 \}$
= $\{ (t_1, t_2, t_3) \in (\mathbb{C}^*)^3 \mid (t_1^{-1} t_2 t_3^0, t_1^{-1} t_2^0 t_3) = (1, 1) \}$
= $\{ (t_1, t_2, t_3) \in (\mathbb{C}^*)^3 \mid t_1 = t_2 = t_3 \}$
= $\{ (t, t, t) \in (\mathbb{C}^*)^3 \}$

which means

$$X_{\Sigma} = (\mathbb{C}^3 - Z(\Sigma))/G = (\mathbb{C}^3 - \{(0,0,0)\})/((x_0, x_1, x_2) \sim (tx_0, tx_1, tx_2)) \cong \mathbb{P}^2$$

as promised.

Chapter 4

Bundles and the Euler class

4.1 Introduction to Bundles

A great many geometric objects of interest in algebraic geometry can be understood as **fiber bundles**, topological spaces that are "locally a Cartesian product." More precisely,

Definition 4.1.1 (Fiber bundle). A fiber bundle (E, B, π, F) , often written synecdochically as simply $\pi : E \to B$, or even just $E \to B$, is a structure comprising topological spaces E and B and a continuous, surjective mapping $\pi : E \to B$ such that for every $b \in B$, there is a neighborhood U of b and homeomorphism $\varphi : \pi^{-1}(U) \to U \times F$ such that the projection map from $\varphi(\pi^{-1}(U))$ to U agrees with the projection map $\pi : \pi^{-1}(U) \to U$.

Intuitively, we have a space with all the local properties of the Cartesian product $B \times F$ (its dimension is the sum of the base space and fiber space dimensions, it contains copies of the base space and the fiber space, etc.) but its global topology can differ from that of the product space. Every space is trivially a bundle with itself as the base space and a one-point space as the fiber (or vice versa), and in fact any way of expressing a space as a product trivially expresses it as a bundle.

However, the need for bundles becomes apparent when we recognize that the direct product is only one way to combine two spaces: $S^1 \times [-1, 1]$ is a cylinder, so what is a Möbius strip? The answer is a **nontrivial** bundle with base space S^1 and fiber space [-1, 1]. The topological nontriviality of bundles such as the Möbius strip crucially depends on the topological nontriviality of the base space–nontrivial bundles occur when we "attach the fibers nontrivially" around an (*n*-dimensional) "hole" in the base space.

We often use this to learn about a space by investigating certain bundles over it.

Note that like any topological space, the topology of a bundle does not depend on its embedding, and in fact the embedding can be misleading–for example, a [-1, 1]-bundle over S^1 that has *two* twists is actually homeomorphic to the trivial bundle (cylinder), which is apparent if we draw its gluing diagram. Conversely, the Möbius strip is not nontrivial simply because it cannot be deformed to a cylinder in 3D; it possesses fundamentally different topology, such as having only one edge. Since we won't be able to directly visualize higher-dimensional bundles anyway, we clearly need more formal ways of differentiating and classifying them–the topic of the next section.

4.2 The Euler class

A key algebraic tool for classifying bundles, though not all-powerful, is the machinery of characteristic classes. These are algebro-topological objects, consisting of certain equivalence classes of curves in the base space of the bundle, determined by the bundle's topology.

Before we can understand what this means, we need to have a language in which to talk about curves in the base space; as always in algebraic geometry, curves will be defined by the loci of points that solve certain (usually polynomial) equations. But as usual, these solution loci can also be thought of as the zero sets of (usually polynomial) functions on the space. What will be the codomain of these functions? Since our goal is to study curves in the base space of bundles, we set the codomain to the bundle itself, in a way that respects its structure.

Definition 4.2.1 (Section of a fiber bundle). A section σ of a fiber bundle $\pi : E \to B$ is a continuous, injective function $\sigma : B \to E$ such that for all $b \in B$, $\pi(\sigma(b)) = b$.

In other words, a section is a "graph" over the base space, with the fiber space as the "vertical axis." It turns out that many types of functions on a space can be understood as sections of a relevant bundle, but for now, the important thing to focus on is the "zero loci" of these sections (understood as graphs) or equivalently the fixed points { $b \in B | \sigma(b) = b$ } of the sections (understood as mappings) or equivalently the intersections of the **zero section** *B* ('original copy of the base space *B* to which we attached the origins of the fibers, understood as a curve in *E*') with the curve $\sigma(B)$.

Again, these vanishing loci, denoted $Z(\sigma)$, are curves in the base space. The key insight we are concerned with is that *the topology of a bundle determines the topology of the vanishing loci of sections*.

For example, consider an infinite cylinder, understood as an \mathbb{R} -bundle over S^1 . Then a section is just some loop around the cylinder, which (as we can translate it arbitrarily far up or down) can be made not to intersect the zero section at all. On the other hand, a section on the infinitely wide Möbius strip, another \mathbb{R} -bundle over S^1 , must cross through the base space at least once to come back to itself as we loop around the strip (try it on the gluing diagram!). This hints that we have an invariant which distinguishes these bundles, but we need to make it more precise-sections on the infinite cylinder can have any even number of zeros, or an odd number if some are tangencies to rather than transverse crossings of the base space. On the Möbius strip, they have an odd number if the crossings are transverse, but again the number may be even if higher-order zeros are involved. We account for this variability by requiring that the sections intersect S^1 transversely, and counting points with orientation. With only transverse intersections, we know there will be an even number for the cylinder and an odd number of the Möbius strip, and incorporating orientation, defined as the "direction of the crossing," means the number of points is invariant under deformations of the section through S^1 . Finally we have a consistent answer of [0 points] for the cylinder and [1 point] for the Möbius strip. These are called the Euler classes of these bundles, and the construction generalizes to higher dimensions.

Definition 4.2.2 (Euler class). Let $\pi : E \to B$ be a fiber bundle. Let $\sigma : B \to E$ be a generic section that transversely intersects with *B*. The Euler class e(E) of *E* is the Poincaré dual of the homology class $[Z(\sigma)]$. For any topological space *B*, e(B) is defined as the Euler class of the tangent bundle of *B*.

Note that e(X) means different things depending on whether X is viewed as a bundle or just any old space (in which case we take the Euler class of the bundle over that space which most naturally captures its structure, the tangent bundle). The notion of 'homology class' just refers to the equivalence class of the curve up to deformations and cancellations of components with opposite orientations, and taking the Poincaré dual is essentially just a formality that means we turn equivalence classes of curves into ring elements that can be "added" (union) and "multiplied" (intersection), which is generally the most productive setting in which to do algebro-topological computations. In short, just think about the Euler class as a curve in the base space.

This machinery can be used to prove various topological results involving bundles. For example, the hairy ball theorem, which states that the number of zeros (with multiplicity and orientation) of any vector field on the 2-sphere is 2, is simply the statement that the Euler class $e(S^2)$ of the bundle TS^2 is 2·[a point].

A particularly useful way of thinking about the Euler class is as the The ψ classes that feature in the recent developments in tropical Gromov-Witten theory are a special case of Euler classes.

Chapter 5

Moduli spaces of *n*-pointed curves

This and the next section closely follow parts of Renzo Cavalieri's very accessible "Curve counting" lecture notes, which develop the theory further and connect it to the interesting topic of Hurwitz theory.

5.1 The moduli space of *n* distinct points

Before we can understand the interesting spaces of *n*-pointed curves at the heart of this subject, we have to understand the simpler spaces by which they are inspired:

Definition 5.1.1 (Moduli space of isomorphism classes of *n* distinct ordered points on \mathbb{P}^1). The moduli space $M_{0,n}$ is the space $\{(p_1, \ldots, p_n) | p_i \in \mathbb{P}^1, p_i \neq p_j \text{ for } i \neq j\}/\sim$ where the equivalence relation is $(p_1, \ldots, p_n) \sim (q_1, \ldots, q_n)$ if there is some $\varphi \in \text{Aut}(\mathbb{P}^1)$ such that $\varphi(p_i) = q_i$ for all *i*.

The unquotiented space, of tuples of distinct points is the easy part to understand: it's $\mathbb{P}^1 \times \mathbb{P}^1 \times \cdots \mathbb{P}^1$ but with all the diagonals removed. However, note that $(\mathbb{P}^1)^n \neq \mathbb{P}^n$: for some intuition about this, remember that \mathbb{P}^n is an *n*-sphere with antipodal points identified; the situation is analogous to comparing the *n*-sphere to the torus $(S^1)^n$.

We now come to the equivalence relation. As \mathbb{P}^n is obtained by taking \mathbb{C}^{n+1} , dropping **0** and modding out scalar multiplication, Aut(\mathbb{P}^n) $\simeq PGL(n+1, \mathbb{C})$, the projective linear group, is obtained by taking Aut(\mathbb{C}^{n+1}) $\simeq GL(n + 1, \mathbb{C})$ and modding by scalar multiplication. This means that if we represent

 \mathbb{P}^n with our usual homogeneous coordinates, which are inherited from the unquotiented vector space, $PGL(n+1, \mathbb{C})$ just acts by matrix multiplication. In our case, we have $Aut(\mathbb{P}^1) \simeq PGL(2, \mathbb{C}, \text{ equivalence classes of } 2x2 \text{ invertible matrices with respect to scalar multiplication.}$

If we consider the points $\{(x_0 : x_1) \in \mathbb{P}^1 | x_1 \neq 0\} \simeq \{(z : 1) \in \mathbb{P}^1\}$, an element of the projective linear group acts like

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} z \\ 1 \end{pmatrix} = \begin{pmatrix} az + b \\ cz + d \end{pmatrix},$$

so we can also understand the automorphisms of \mathbb{P}^1 as the group of Möbius transformations,

$$\operatorname{Aut}(\mathbb{P}^1) = \left\{ \varphi : \mathbb{P}^1 \to \mathbb{P}^1 \left| \phi(z:1) = \left(\frac{az+b}{cz+d} : 1 \right), ad - bc \neq 0 \right\} \right\},$$

where we understand that $(\infty : 1) = (1 : 0)$.

What do we get when we take the diagonal-free *n*-fold product of \mathbb{P}^1 and quotient by Mobius transformations? Let's start from n = 1 and work our way up.

The elements of $M_{0,1}$ are the equivalence classes of points of \mathbb{P}^1 . But for any point $(z : 1) \in \mathbb{P}^1$, the Mobius transformation with a = 1, b = -z, c =0, d = 1 maps it to $\left(\frac{1z-z}{0z+1} : 1\right) = (0 : 1)$ (for (1 : 0), we set a = 0, b = 1, c =1, d = 0)). So we see that every point is isomorphic to (0 : 1), and $M_{0,1}$ is a single point.

In fact, for any 3 distinct points $(z_1 : 1), (z_2 : 1), (z_3 : 1)$, we can solve the linear equations

$$\begin{cases} z_1a + b = 0\\ z_2a + b - z_2c - d = 0\\ z_3c + d = 0 \end{cases}$$

to obtain the class of projective transformations

$$\begin{pmatrix} z_2 - z_3 & -z_1(z_2 - z_3) \\ z_2 - z_1 & -z_3(z_2 - z_1) \end{pmatrix}$$

which simultaneously send $(z_1 : 1) \mapsto (0 : 1), (z_2 : 1) \mapsto (1 : 1), (z_3 : 1) \mapsto (1 : 0)$. As before, the procedure will still work for tuples involving the point (1 : 0): it is no worse behaved than any other point, just inconvenient for our algebra. So we see that any element of $M_{0,3}$ is isomorphic to ((0 : 1), (1 : 1), (1 : 0)), and thus $M_{0,3}$ and $M_{0,2}$ are also single-point spaces.

Beyond n = 3 we have some actual substance: due to scalar invariance, $PGL(2, \mathbb{C})$ is a 3, not 4 dimensional group, which means that we cannot demand more simultaneous constraints; we don't have the freedom to send 4 points to any 4 specific spots simultaneously. If we plug a fourth point $(z_4 : 1)$ into the automorphism we just found, we obtain the point $((z_2 - z_3)(z_4 - z_1) : (z_2 - z_1)(z_4 - z_3)) = \left(\frac{(z_2 - z_3)(z_4 - z_1)}{(z_2 - z_1)(z_4 - z_3)} : 1\right) \equiv (\lambda : 1)$. λ is called the **cross-ratio**; it can be thought of as "where the 4th point is when the first 3 are put in standard position" and the possible values of λ parametrize $M_{0,4}$, demonstrating that $M_{0,4} \simeq \mathbb{P}^1 \setminus \{(0:1), (1:1), (1:0)\}$.

For an *n*-tuple $(p_1, ..., p_n)$, we have n - 3 cross-ratios telling us where all the remaining points go when we fix the first 3. These can be anything except our 3 standard points or each other, so we end up with

 $M_{0,n} \simeq (\{(0:1)\} \times \{(1:1)\} \times \{(1:0)\} \times \mathbb{P}^{n-3}) \setminus \text{diagonals} \simeq (M_{0,4})^{n-3} \setminus \text{diagonals}.$

5.2 Moduli spaces of stable *n*-pointed curves

At this point, we finally have the machinery to construct and understand the (classical versions of the) central objects of our focus, moduli spaces of stable curves. The motivation to construct something new and confusing from the $M_{0,n}$ s, and expect something interesting to come of it, is that $M_{0,n}$ badly behaved for intersecting things on (and thereby for enumerative geometry), and we would like to fix it. $M_{0,n}$ is badly behaved because it is not compact: closed subsets have been removed from it (the diagonals) leaving an open set. Can we compactify it by stitching up its open ends? The obvious guess is that we would be forced to include *n*-tuples with coinciding points, giving us $\overline{M}_{0,n} \simeq \mathbb{P}^{n-3}$, and erasing any interesting structure that may have been present. Is there really a way to have a *compact* space that captures the notion of *distinct* points, through some aspect of its geometry rather than just cutting out degenerate loci where points coincide? Surprisingly, the answer is yes: we construct the space $\overline{M}_{0,n}$ of **stable rational** *n*-**pointed curves**.

Definition 5.2.1 (stable rational *n*-pointed curve). A stable rational *n*-pointed curve is a genus 0 connected algebraic curve C (possibly with singularities) with *n* ordered, distinct points (not coinciding with the singularities) such that there are no nontrivial automorphisms of C that preserve the marked and singular points.

When unpacking this definition, the trick is to understand \mathbb{P}^1 as both the Riemann sphere and as a compactified, complex *line*. As any real, open

curve (the zero locus of a polynomial in the plane) is the image of \mathbb{R} under some polynomial mapping, a "genus 0 connected algebraic curve" without singularities (the zero locus of a rational function in \mathbb{P}^2) is the image of \mathbb{P}^1 under some rational mapping (rational functions are the equivalent of polynomials when the codomain is assumed to be compact/containing ∞). The automorphisms of such curves should commute with the rational maps we consider natural in this setting, so they're just the images of the same projective linear transformations. The mental image for a non-singular genus 0 curve is therefore just \mathbb{P}^1 , warped by its embedding in \mathbb{P}^2 . For our purposes, nothing interesting will be going on in \mathbb{P}^2 , so we don't actually need to think about the embedding.

So for non-singular *C*, the *n*-pointed curves just correspond to the points of $M_{0,n}$. We brought in the language of curves so that we could include some singular ones: since the genus is still 0, these will just look like Riemann spheres/projective lines that are 'pinched' in some places, henceforth called the 'nodes.' This pinching is highly significant: the automorphisms of such curves are no longer as simple as those of \mathbb{P}^1 . Instead, an automorphism is an independent projective linear transformation on each \mathbb{P}^1 -like *component* of the curve, with consistency on the nodes.

Just as we need 3 points on \mathbb{P}^1 to 'keep it stable' (have no nontrivial point-fixing automorphisms), we need 3 points (including nodes) on each component of a general genus 0 curve to satisfy this property. If a component has at least 4 such points, we can smoothly vary the curve into one with an additional component, by bringing two of the marked points together and sprouting off a new stable component for them when they coincide. This is the essence of the compactification: where $M_{0,n}$ holds the geometry fixed and has holes where points coincide, $\overline{M}_{0,n}$ lets both points and geometry vary, so that the geometry has precisely enough freedom to accommodate those coincidences by transforming in a nontrivial way, filling the holes in $M_{0,n}$ without washing out the structure.

We can simply represent the different types of elements of $M_{0,n}$ with marked trees of lines: each line represents a component, with its marked points placed on it (order doesn't matter); there can be no closed polygons or lines with fewer than 3 (points + crossings). The *n*-marked trees with more than one line represent the classes of stable curves in the boundaries of $\overline{M}_{0,n}$: the stuff we tacked on to $M_{0,n}$. Each possible tree represents a certain locus of boundary points; the number of crossings in such a tree is the 'number of coincidences' that must occur to arrive at the corresponding curves, so is



Figure 5.1 'Sprouting off a new stable component' when marked points coincide

equal to the (complex) codimension of the locus of boundary points. When a tree can be formed by 'sprouting a component' from multiple different trees of one less component, it represents, in the moduli space, the intersection of the (closures of) the moduli space boundaries represented by its possible parent trees. We therefore have a 'stratification' of the boundary of $\overline{M}_{0,n}$: at codimension 1, the different boundaries correspond to 2-component trees with different partitions of the *n* marked points between those components, so we have one for each stable partition $[n] = A \cup A^c$. The closure of such a boundary is called the **boundary divisor** D(A). The boundary divisors intersect in a number of codimension 2 curves, which themselves intersect in codimension 3 curves, and so on until we run out of dimensions and don't have any 'spare marked points' to use for stably sprouting off new components. Enumerating the intersections of these boundary curves is nontrivial, since it involves the combinatorics of partitions, but yields intricate results.

5.3 ψ classes

Can we use the algebro-topological machinery of cohomology classes, and specifically Euler classes of bundles, on these moduli spaces to study rational

34 Moduli spaces of *n*-pointed curves

stable curves? The answer is yes, and leads us to the Euler classes of a particular collection of bundles over $\overline{M}_{0,n}$ called ψ classes.

Definition 5.3.1 (ψ class). Let $\mathbb{L}_i \to \overline{M}_{0,n}$ be a line bundle where the fiber over each point (C, p_1 , ..., p_n) is canonically identified with $T^*_{p_i}(C)$. \mathbb{L}_i is called the *i*th **cotangent line bundle**, and we define the ψ classes ψ_i , ..., $\psi_n \in A^1(\overline{M}_{0,n})$ as

$$\psi_i \equiv e(\mathbb{L}_i).$$

There's some new notation here, but it's mostly concepts already introduced. $A^{1}(M_{0,n})$ is the abelian group of codimension-1 cohomology classes in $M_{0,n}$ -recall that equivalence classes of curves form a ring under union and intersection; thinking in terms of representatives of these classes, the ψ classes are each particular codimension-1 curves in $M_{0,n}$. Specifically, they are the zero loci of generic sections of the \mathbb{L}_i bundles, measuring how twisted they are. But what are the \mathbb{L}_i ? Essentially, \mathbb{L}_i tracks the point p_i across the moduli space of stable *n*-marked curves; at each point in the moduli space, it encodes the position of p_i on the curve associated to that point, by attaching the tangent to C at p_i to that point, as a fiber. It seems sort of simultaneously tautological and the most complicated thing that could have been invented, but it makes sense that having phrased questions about automorphism-free projective curves in terms of arrangements/movements of points on those curves, we would like to get back to the machinery of geometry, by encoding the arrangements/movements of points on those curves as the twisting of a geometric object (\mathbb{L}_i). Note that the \mathbb{L}_i are thus canonically isomorphic via relabeling the points; they all 'look the same,' but there is still intricate structure in the ways their Euler classes intersect with different combinations of copies of themselves and copies of the classes of the other cotangent line bundles.

There is really a lot of structure here but meaningfully getting into it would require substantial discussion of pullbacks, pushforwards, normal bundles, self-intersection, and so on. Suffice to say the aforementioned interesting intersection theory of boundary divisors, which is really *the* interesting structure in the $\overline{M}_{0,n}$, can be computed with ψ classes.

Chapter 6

Tropical $M_{0,n}$

How might one tropicalize the structures introduced in the last chapter, such as the spaces $\overline{M}_{0,n}$ that we are primarily interested in? In short, we tropicalize the objects parameterized by the moduli space, construct the moduli space of tropical objects, and then try to draw connections between the classical and tropical moduli spaces. This process does not always 'work,' in that the moduli space of the tropicalized classical objects is not always the tropicalization of the classical moduli space, but in this case it does.

6.1 Tropical stable *n*-pointed curves

How to tropicalize a stable *n*-pointed curve? We already have combinatorial representations of such curves, the stable *n*-pointed trees. It turns out that these are almost already the tropicalizations of stable *n*-pointed curves; however, they require some modification: each tree represents a range of possible curves, and the dimension of that range of curves depends on the geometry of the tree. Specifically, we must distinguish the combinatorial representations of stable *n*-pointed curves with the same 'topology' but different 'values' for the positions of the singular points. As tropicalization takes complex dimensions to real dimensions (and associated complex numbers to positive real numbers), the complex-valued positions of the points in question will translate to positive real numbers describing the tropicalized structure. Since there is one of these values for each intersection between components of a stable tree. To make things more intuitive, we go ahead and work with the dual graph of the stable tree, where components of

the tree become points, their intersections become edges, and we represent the marked points as leaves connected to their associated component-vertex. The positive real numbers we have assigned to the component-intersections now have the natural interpretation as the lengths of the dual graph.

Definition 6.1.1 (Tropical rational *n*-pointed curve). A tropical rational *n*-pointed curve is a tree with *n* labeled leaves, and a function from the set of non-leaf edges $d : E \to \mathbb{R}^{>0}$. The curve is **stable** if each vertex has degree at least 3.

With this definition, we can represent curves where components are combined by allowing the length of their intersection-edge to go to 0.



Figure 6.1 A tropical stable curve and its classical partner

Returning to our original goal, we would like to construct the moduli space of these tropical curves. For a given tree topology, each possible curve is given by a different choice of edge lengths, so the curves of a given tree topology are parameterized by $(\mathbb{R}^{\geq 0})^n$, where *n* is the number of edges, as shown in Figure 6.2.

Of course, we want the moduli space for *all n*-pointed tropical curves, meaning all tree topologies with *n* leaves. The intuition here is that the different topologies each have cones $(\mathbb{R}^{\geq 0})^n$ for moduli spaces, which will 'fit together' into a cone *complex*, with the cones coinciding on *faces* where some edge lengths go to zero, resulting in the same topology. Of course, the cone complex quickly becomes unwieldy to visualize–the complex for



Figure 6.2 The conical moduli space for one topology of 5-pointed tropical stable curves

5 marked points, shown in Figure 6.3, demonstrates the sort of intricate incidence patterns that arise.



Figure 6.3 The cone complex $\overline{M}_{0,5}^{\text{trop}}$ (from "Curve Counting")

Here the 10 vertices of the graph represent 10 rays diverging from the point representing the 1-component curve. The colored lines between rays represent 15 2D cones, including the one from before, which together form the full moduli space $\overline{M}_{0,5}^{\text{trop}}$. Notice that the top-dimensional components of the cone complex correspond to the curves with the most components (since the number of degrees of freedom/edge lengths for a given topology is the number of connections between components), and the lower-dimensional faces and rays correspond to curves with fewer and fewer components. This is the opposite of the classical case, where generic points in the moduli space correspond to curves with a single component, and the lower-dimensional

38 Tropical $\overline{M}_{0,n}$

boundaries and their intersections correspond to curves with additional, branched-off components. This is the same inclusion-reversion we observed when looking at toric varieties and their associated fans!

Chapter 7

Conclusion

Though I wasn't able to reach a point of being able to properly bring everything together, from what I understand, the brief picture is this. The moduli space of tropical stable curves is the polyhedral fan of the toric variety that is the moduli space of classical stable curves, proving that it is the tropicalization of said moduli space. With the help of tropical intersection theory (which uses much of the same language as enumerative geometry and classical algebro-geometrical cohomology computations, but is nevertheless the biggest gap in my understanding/presentation of these topics), we are inspired to port our enumerative work on $\overline{M}_{0,n}$ to its newly constructed tropical counterpart; this makes heavy use of the aforementioned in the previous section. At the forefront of current research, these constructions are being extended to the more complicated case of moduli spaces of curves of nonzero genus, opening up vast new possibilities for the future of Gromov-Witten computations.

Bibliography

Cavalieri, Renzo. 2017. Curve counting in tropical and algebraic geometry (lecture notes).

Cavalieri, Renzo, Andreas Gross, and Hannah Markwig. 2020. Tropical ψ classes .

Hartshorne, Robin. 1977. *Algebraic Geometry*. Graduate Texts in Mathematics, Springer.

Hori, Kentaro, Sheldon Katz, Albrecht Klemm, Rahul Pandharipande, Richard Thomas, Cumrun Vafa, Ravi Vakil, and Eric Zaslow. 2000. *Mirror Symmetry, Clay Mathematics Monographs,* vol. 1. American Mathematical Society.

Katz, Sheldon. 2006. *Enumerative Geometry and String Theory*. IAS/Park City Mathematical Subseries, American Mathematical Society.

Kock, Joachim, and Israel Vainsencher. 1999. Kontsevich's formula for rational plane curves .

Maclagan, Diane, and Bernd Sturmfels. 2015. *Introduction to Tropical Geometry*. Graduate Studies in Mathematics, American Mathematical Society.

Madore, David. 2020. http://www.madore.org/~david/images/cubic3.jpg.

Meyer, Henning. 2011. Intersection theory on tropical toric varieties and compactifications of tropical parameter spaces .

Speyer, David, and Bernd Sturmfels. 2009. Tropical mathematics. *Mathematics Magazine* 82(3):163–173.