1 Introduction to Legendrian contact homology

1.1 The problem of classification

A very general goal of mathematics is *classification*: given some definition, we want a complete list of all objects meeting that definition. Moreover, when we come across an object, we want to know which one we've got.

For a very famous example, recall that a finite simple group is a finite group whose only proper normal subgroup is the trivial subgroup. As of 2004, there is a complete list of the finite simple groups (which we won't try to reproduce here): given any finite simple group, it must be isomorphic to a group on the list, and the groups in the list are pairwise non-isomorphic.

It's interesting to think about how we might start trying to produce a classification. Of course we're taking two things as given:

1. a definition (e.g., of finite simple groups);
2. a notion of sameness (e.g., group isomorphism).

From here, essentially the only way to prove that two objects are the same is by checking the definition (i.e., by producing an isomorphism). But how can we show that two objects are *not* the same? Can you think of two finite simple groups which are not isomorphic? How do you know they're not isomorphic?

Being able to obstruct sameness is important: when we come across a new object which satisfies our definition, we want to know whether it's truly new, or is in fact the same as an object already on our list. Hopefully the obstruction we gave to sameness for finite simple groups was the *order* of a group. Because group isomorphism preserves order, a pair of groups with distinct orders cannot be isomorphic. So, for instance, each of the groups $\mathbb{F}_p$, where p is prime, appears as a distinct entry in our list of finite simple groups. Knowing that we've listed *all* objects of a particular type is a different, more difficult problem.

We call the order of a finite simple group an *invariant* because it is not changed by isomorphism. One of the major themes of modern mathematics is defining and computing invariants of various objects. Intuitively, invariants should make comparisons between objects easier: it's much easier to check the sameness of a pair of integers than that of a pair of finite simple groups.

Our goal over the next couple of weeks is to define some complicated invariants of Legendrian knots. In order to do this, we'll need to define Legendrian knots and what it means for two Legendrian knots to be the same, as mentioned above. But we'll also need to define other mathematical objects — mainly complicated algebraic gadgets — and notions of sameness for those objects. The invariants we want to extract from Legendrian knots aren't as simple as integers, but are themselves complicated objects. In fact, we'll then need to compute invariants of our invariants until we have objects for which comparisons are reasonable.

1.2 Legendrian knots

The objects we want to classify this summer are *Legendrian knots*. We'll give a more careful description of these later in the week, but for now let's think about *parametrized* Legendrian knots. So we'll assume we have a regular parametrization $\gamma\colon \mathbb{R} \to \mathbb{R}^3$, given by $\gamma(t) = (x(t), y(t), z(t))$, satisfying

$$z'(t) - y(t)x'(t) = 0 \quad \text{and} \quad \gamma(t+1) = \gamma(t), \forall t \in \mathbb{R}.$$

The first condition — the differential equation — ensures that γ is Legendrian, while the second condition ensures that we have a knot.

Our notion of sameness for Legendrian knots will be Legendrian isotopy. We'll say that Legendrian knots γ and γ' are *Legendrian isotopic* if we can find a smooth family of Legendrian knots γ_s, $s \in [0,1]$, such that

$\gamma_0 = \gamma$ and $\gamma_1 = \gamma'$.

Intuitively, isotopy is how you would naturally manipulate a knot in 3D space. If you have a closed up knot — think of plugging an extension cord into itself — then you can basically do anything you want except cut the knot open. The catch for Legendrian knots is that we insist that all of our intermediate knots satisfy the differential equation. So maybe you have two knots which are *topologically* the same — that is, you can manipulate one into the other in $\mathbb{R}^3$ — but which are not Legendrian isotopic.

Examples: Show the *Mathematica* examples.

1.3 The Thurston-Bennequin invariant and the front projection

One invariant of Legendrian knots is obvious: the topological knot type. That is, if we can't turn one knot into another, even ignoring the Legendrian condition, then we certainly can't win in the Legendrian setting. Of course the knot type is itself a complicated invariant — and thus subject of much study — but let's ignore that for now and think about distinguishing two Legendrian knots L and L' which are smoothly isotopic (such as the examples we just saw).

One way to do this uses the *Thurston-Bennequin invariant*. We compute the TB invariant of L by first considering the knot L_ϵ we get by translating L up a short distance ϵ in the z-direction. Notice that L_ϵ is also a Legendrian knot. Then

$$\mathrm{tb}(L) = lk(L, L_\epsilon).$$

That is, $\mathrm{tb}(L)$ counts the number of times L_ϵ wraps around L. Let's try to compute this for our first Legendrian unknot above.

Aside: Maybe it seems strange that this linking number can be nonzero for an unknot. How could shifting our knot up give us a knot which links with the original? Try playing around to see what causes linking.

We can compute the Thurston-Bennequin invariant — and just draw Legendrian knots in general — more simply in 2D. The *front projection* of a Legendrian knot $\gamma(t) = (x(t), y(t), z(t))$ is given by

$$\pi_F(\gamma(t)) = (x(t), z(t)) \in \mathbb{R}^2_{x,z}.$$

That is, we just ignore the y-coordinate. Notice that we can recover $y(t)$ from the Legendrian condition — we have $y(t) = z'(t)/x'(t)$. You'll also notice that the front projection is not smooth — it has a cusp whenever $x'(t) = 0$. That's okay.

Now we can compute the Thurston-Bennequin invariant in the front projection: translating up in the z-direction is easy, and the differential equation allows us to determine crossing data. Let's do this for the *Mathematica* examples.

Exercise. Find a formula for the Thurston-Bennequin invariant, computed from the front projection. Your formula should use the number of cusps and crossings in your diagram, and you'll need to distinguish between positive crossings and negative crossings.

Reidemeister moves (aside, will probably skip)

Why should you believe me that the Thurston-Bennequin invariant is actually an invariant of Legendrian knots? We need to know that applying a Legendrian isotopy won't change the invariant, but that seems like a lot to check — we have to consider all possible Legendrian isotopies.

Thankfully, we only have to check a finite number of isotopies. It is a fact I'm not going to prove that if L and L' are Legendrian isotopic, then their front diagrams can be related by *Legendrian Reidemeister moves*:

show the moves.

So if we want to know that some quantity is invariant under Legendrian isotopy, we just need to show that it doesn't change when we apply any of the Legendrian Reidemeister moves. The logic goes: if L and L' are Legendrian isotopic, then there's a chain of Legendrian Reidemeister moves connecting the front diagram of one to the front diagram of the other; since our invariant doesn't change values as we apply the moves, L an L' have the same invariant value.

Exercise. Using your formula for the Thurston-Bennequin invariant from before, check that tb is an invariant of Legendrian knots.

1.4 Chekanov-Eliashberg DGA

- Into the 1990s, there were only three invariants of Legendrian knots: the smooth knot type, the Thurston-Bennequin invariant, and the rotation number.
- The Thurston-Bennequin invariant and rotation number are integers; if we want an invariant to store lots of information, it'll probably need to be a more complicated object.
- In the late 1990s, Chekanov introduced an invariant now known as the Chekanov-Eliashberg DGA, inspired by ideas due to Eliashberg. Here DGA stands for *differential graded algebra*, a rather complicated algebraic gadget.
- Though the Chekanov-Eliashberg DGA is an invariant of Legendrian knots (up to stable tame isomorphism), it is often too complicated to work with directly. By way of analogy, think about how difficult it is to determine the isomorphism type of a group given a presentation. It is similarly difficult to determine whether or not two Legendrian knots lead to the same Chekanov-Eliashberg DGA.
- At this point, the task becomes algebraic: given two DGAs, what sort of information can we extract that will be invariant under stable tame isomorphism? One of these things is the *homology* of the DGA. We call the homology of the Chekanov-Eliashberg DGA of a Legendrian knot L the *Legendrian contact homology* of L.
- And now the algebra continues: comparing the LCH of two different Legendrian knots is still a problem in noncommutative algebra, and therefore very difficult. We'll also need to define *augmentations* of DGAs, which allow us to compute the *linearized Legendrian contact homology* of a Legendrian knot. Linearized LCH is a much more reasonable invariant, and augmentations of the Chekanov-Eliashberg DGA are interesting invariants to study in their own right.

2 Differential Forms on $\mathbb{R}^3$ I

2.1 Covectors

We start by giving an abstract definition of covectors.

Definition. A **covector** on $\mathbb{R}^3$ is a linear map $\lambda\colon \mathbb{R}^3 \to \mathbb{R}$. That is, λ is a map which satisfies

$$\lambda(\mathbf{v}+\mathbf{w}) = \lambda(\mathbf{v}) + \lambda(\mathbf{w}) \quad \text{and} \quad \lambda(c\mathbf{v}) = c\lambda(\mathbf{v}),$$

for any vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^3$ and any scalar $c \in \mathbb{R}$. We denote the collection of all covectors on $\mathbb{R}^3$ by $(\mathbb{R}^3)^*$.

Abstractly, we could develop the entire theory of covectors from the above definition, we'll take a slightly more grounded approach.

Example 2.1. We can define covectors dx, dy, and dz on $\mathbb{R}^3$ which simply give the relevant coordinate of a vector:

$$dx\left(\begin{bmatrix} a \\ b \\ c \end{bmatrix}\right) = a, \quad dy\left(\begin{bmatrix} a \\ b \\ c \end{bmatrix}\right) = b, \quad dz\left(\begin{bmatrix} a \\ b \\ c \end{bmatrix}\right) = c.$$

You should verify that each of these is a linear map $\mathbb{R}^3 \to \mathbb{R}$.

We call dx, dy, dz the **standard covectors** on $\mathbb{R}^3$. Another way to write them uses the dot product on $\mathbb{R}^3$:

$$dx(\mathbf{v}) := \mathbf{v} \cdot \mathbf{i}, \quad dy(\mathbf{v}) := \mathbf{v} \cdot \mathbf{j}, \quad dz(\mathbf{v}) := \mathbf{v} \cdot \mathbf{k}, \tag{2.1}$$

where $\mathbf{i}, \mathbf{j}, \mathbf{k}$ is the standard basis for $\mathbb{R}^3$.

Remark. The collection $(\mathbb{R}^3)^*$ of covectors on $\mathbb{R}^3$ is a 3-dimensional vector space, called the *dual* of $\mathbb{R}^3$. The standard covectors form a basis for this dual space, and the standard covectors are by definition dual to the standard basis vectors of $\mathbb{R}^3$. The fact that we can relate these bases via equation 2.1 is sort of a special feature of the standard inner product on $\mathbb{R}^3$.

Exercise 2.2. Show that the standard covectors dx, dy, dz form a basis for $(\mathbb{R}^3)^*$. That is, show that every covector λ can be written as a linear combination of dx, dy, and dz, and show that the standard covectors are linearly independent.
Hint: To prove linear independence, we'll need to show that if

$$0 = a\,dx + b\,dy + c\,dz,$$

for some real numbers $a, b, c \in \mathbb{R}$, then we must have $a = b = c = 0$.

According to Exercise 2.2, every covector on $\mathbb{R}^3$ can be uniquely written as a **linear combination** of the standard covectors. This motivates the following re-definition of covectors.

Definition. A **covector** on $\mathbb{R}^3$ is a linear combination of the standard covectors on $\mathbb{R}^3$. That is, a covector is a map $\lambda\colon \mathbb{R}^3 \to \mathbb{R}$ of the form

$$\lambda = a\,dx + b\,dy + c\,dz,$$

for some real numbers $a, b, c \in \mathbb{R}$.

This is all well and good, but what does a covector actually *do*? The answer is that a covector eats vectors, and outputs the component of the input vector along some fixed vector.

Example 2.3. Say we have a fixed vector $\mathbf{v}_0 := \langle a_0, b_0, c_0 \rangle \in \mathbb{R}^3$, and we want a map which takes some vector $\mathbf{v} \in \mathbb{R}^3$ as input and tells us how much $\mathbf{v}$ pushes along $\mathbf{v}_0$. We can define a covector $\lambda = (\mathbf{v}_0)^*$ as

$$\lambda = a_0\,dx + b_0\,dy + c_0\,dz.$$

Then we notice that

$$\lambda(\mathbf{v}) = a_0(\mathbf{v}\cdot\mathbf{i}) + b_0(\mathbf{v}\cdot\mathbf{j}) + c_0(\mathbf{v}\cdot\mathbf{k}) = \mathbf{v}\cdot\mathbf{v}_0$$

gives us the component of $\mathbf{v}$ along $\mathbf{v}_0$. So *covectors are dual to vectors* (which is how they get the name).

2.2 Differential 1-forms

Remember that a vector field on $\mathbb{R}^3$ assigns a vector to each point of $\mathbb{R}^3$. We usually write these as

$$\mathbf{X} = F_1\,\mathbf{i} + F_2\,\mathbf{j} + F_3\,\mathbf{k},$$

where F_1, F_2, F_3 are some functions $\mathbb{R}^3 \to \mathbb{R}$ which we typically assume are smooth.

Just as we can define vector fields, we can define covector fields. However, we will usually refer to covector fields as **differential 1-forms**. Actually, we'll often drop the word "differential" and just talk about **1-forms**.

Definition. A **differential 1-form** (or **1-form** or **covector field**) on $\mathbb{R}^3$ is a smooth[1]assignment of a covector on $\mathbb{R}^3$ to each point of $\mathbb{R}^3$. That is, a differential 1-form is a smooth map $\mathbb{R}^3 \to (\mathbb{R}^3)^*$.

In practice, we won't think about the above definition too much. Instead, we'll consider a 1-form to be a linear combination of the **standard 1-forms** dx, dy, dz, except that scalar multiplication actually means multiplication by scalar-valued functions. That is, a differential 1-form on $\mathbb{R}^3$ can be written as

$$\lambda = F_1\,dx + F_2\,dy + F_3\,dz,$$

for some real-valued functions $F_1, F_2, F_3 : \mathbb{R}^3 \to \mathbb{R}$.

Notation. It might be a little confusing that we used λ to denote a covector and also a covector field. Similarly, we used dx, dy, dz to denote the standard covectors, and now they denote the standard 1-forms. This is because we can think of covectors as constant covector fields: they give us the same map at every point of $\mathbb{R}^3$. Going forward, we will think pretty much exclusively about differential 1-forms, and rarely about a single, stand-alone covector.

Just as we did with covectors, we give a more down-to-earth definition of differential 1-forms which is suited to the needs of this course. This is somehow less pleasing from an abstract point of view, but will meet our needs.

Definition. A **differential 1-form** on $\mathbb{R}^3$ is a linear combination over $C^\infty(\mathbb{R}^3)$ of the standard 1-forms on $\mathbb{R}^3$. That is, a differential 1-form λ is an expression

$$\lambda = F_1\,dx + F_2\,dy + F_3\,dz,$$

where $F_1, F_2, F_3 : \mathbb{R}^3 \to \mathbb{R}$ are smooth maps. The set of all differential 1-forms on $\mathbb{R}^3$ is denoted $\Omega^1(\mathbb{R}^3)$.

Remark. Unlike $(\mathbb{R}^3)^*$, $\Omega^1(\mathbb{R}^3)$ is not a vector space. Instead, $\Omega^1(\mathbb{R}^3)$ is a module over the ring $C^\infty(\mathbb{R}^3)$. The basic idea is that, where scalar multiplication in $(\mathbb{R}^3)^*$ can be undone (provided the scalar is nonzero), the same is not true for $\Omega^1(\mathbb{R}^3)$. For instance, $\lambda = x\,dx$ is a differential 1-form on $\mathbb{R}^3$, and there does not exist any smooth function $F : \mathbb{R}^3 \to \mathbb{R}$ with the property that $F\,\lambda = dx$.

A differential 1-form on $\mathbb{R}^3$ corresponds in a natural way to a vector field, and we'll formalize this correspondence with notation:

$$\lambda = F_1\,dx + F_2\,dy + F_3\,dz \quad \Rightarrow \quad \mathbf{X}_\lambda = F_1\,\mathbf{i} + F_2\,\mathbf{j} + F_3\,\mathbf{k}.$$

[1]Making this precise would require thinking about a differential structure on the dual space $(\mathbb{R}^3)^*$. This isn't hard, but we won't do it. Instead, we'll think about 1-forms as linear combinations of standard 1-forms, where the "scalars" in our linear combinations are smooth functions $f : \mathbb{R}^3 \to \mathbb{R}$.

2.3 Integrating 1-forms

There's enough abstraction going on that it might be difficult to keep track of what we're doing. But the mantra so far is that *covectors measure (signed) lengths of projections*. That is, every covector λ corresponds to some fixed vector $\mathbf{v}_0$, and evaluating λ on some vector $\mathbf{v}$ tells us the component of $\mathbf{v}$ along $\mathbf{v}_0$.

The mantra that many people find helpful for 1-forms is to say that *1-forms are objects which we can integrate over curves*. Indeed, if a covector tells us the component of one vector along some other vector, then we can use a covector field (i.e., a 1-form) to determine the component of the unit tangent vector $\mathbf{T}$ of some curve along a vector field. Here's the defining equation:

$$\int_{\mathcal{C}} \lambda := \int_{\mathcal{C}} \lambda(\mathbf{T})\, ds.$$

On the left, we have some 1-form λ which we want to integrate over an oriented curve $\mathcal{C}$. On the right, we have a scalar line integral over $\mathcal{C}$, and the scalar-valued function we're integrating is given by $\lambda(\mathbf{T})$.

Remember that if $\mathcal{C}$ is an oriented curve in $\mathbb{R}^3$, we can integrate a function $f : \mathbb{R}^3 \to \mathbb{R}$ over $\mathcal{C}$ by choosing a regular parametrization $\mathbf{r}\colon [a,b] \to \mathbb{R}^3$ of $\mathcal{C}$ and using the equation

$$\int_{\mathcal{C}} f\, ds = \int_a^b f(\mathbf{r}(t))\, \|\mathbf{r}'(t)\|\, dt.$$

(Let's assume that $\mathbf{r}$ is injective on $[a,b]$.) Applying this to the above definition leads us to

$$\boxed{\int_{\mathcal{C}} \lambda = \int_a^b \lambda_{\mathbf{r}(t)}(\mathbf{r}'(t))\, dt,}$$

for any regular parametrization $\mathbf{r}(t)$, $a \le t \le b$, of $\mathcal{C}$. This is the equation we'll use in practice.

Example 2.4. Consider the 1-form $\lambda = x\, dy - y\, dx$, and let $\mathcal{C}$ be the circle of radius R centered at $(0,0,0)$ and contained in the xy-plane. Let's compute the integral of λ over $\mathcal{C}$. We can parametrize $\mathcal{C}$ via

$$\mathbf{r}(t) = (R\cos t, R\sin t), \quad 0 \le t \le 2\pi.$$

Then $\mathbf{r}'(t) = (-R\sin t, R\cos t)$, so

$$\lambda_{\mathbf{r}(t)}(\mathbf{r}'(t)) = x(t)y'(t) - y(t)x'(t) = (R\cos t)(R\cos t) - (R\sin t)(-R\sin t) = R^2,$$

so we have

$$\int_{\mathcal{C}} \lambda = \int_0^{2\pi} \lambda_{\mathbf{r}(t)}(\mathbf{r}'(t))\, dt = \int_0^2 \pi R^2\, dt = \boxed{2\pi R^2}.$$

2.4 2-forms

Since we integrate 1-forms over curves, it seems natural that we should integrate 2-forms over surfaces. Unlike 1-forms, we won't give the abstract definition of differential 2-forms — it's just a bit *too* abstract. Instead, we'll define the **standard 2-forms** on $\mathbb{R}^3$ and, as we did with 1-forms, consider an arbitrary 2-form to be a linear combination of these.

Recall that the 1-forms dx, dy, and dz tell us the components of a vector $\mathbf{v}$ in the x-, y-, and z-directions, respectively. We want to similarly define standard 2-forms

$$dy \wedge dz, \quad dz \wedge dx, \quad \text{and} \quad dx \wedge dy$$

on $\mathbb{R}^3$. A 2-form will accept as input a pair $\mathbf{v}$, $\mathbf{w}$ of vectors and output a real number. In each case, this real number will correspond to the signed area of a projection of the parallelogram in $\mathbb{R}^3$ spanned by $\mathbf{v}$ and $\mathbf{w}$.

Let's think about what this means for $dx \wedge dy$. Say our vectors are

$$\mathbf{v} = (v_1, v_2, v_3) \quad \text{and} \quad \mathbf{w} = (w_1, w_2, w_3).$$

Projecting these vectors to the xy-plane gives

$$\text{proj}_{xy}(\mathbf{v}) = (v_1, v_2) \quad \text{and} \quad \text{proj}_{xy}(\mathbf{w}) = (w_1, w_2),$$

and these two vectors span a parallelogram with (signed) area given by

$$\det\big(\text{proj}_{xy}(\mathbf{v}) \quad \text{proj}_{xy}(\mathbf{w})\big) = \det\begin{pmatrix} v_1 & w_1 \\ v_2 & w_2 \end{pmatrix} = v_1 w_2 - w_1 v_2.$$

Based on this computation, we declare that $dx \wedge dy$ is the bilinear map $\mathbb{R}^2 \to \mathbb{R}$ defined by

$$(dx \wedge dy)(\mathbf{v}, \mathbf{w}) = dx(\mathbf{v})dy(\mathbf{w}) - dx(\mathbf{w})dy(\mathbf{v}).$$

Analogously, we have

$$(dy \wedge dz)(\mathbf{v}, \mathbf{w}) = dy(\mathbf{v})dz(\mathbf{w}) - dy(\mathbf{w})dz(\mathbf{v}) \quad \text{and} \quad (dz \wedge dx)(\mathbf{v}, \mathbf{w}) = dz(\mathbf{v})dx(\mathbf{w}) - dz(\mathbf{w})dx(\mathbf{v}).$$

We refer to $dy \wedge dz$, $dz \wedge dx$, and $dx \wedge dy$ as the **standard 2-forms** on $\mathbb{R}^3$.

Definition. A **differential 2-form** on $\mathbb{R}^3$ is a linear combination over $C^\infty(\mathbb{R}^3)$ of the standard 2-forms on $\mathbb{R}^3$. That is, a differential 2-form η is an expression

$$\eta = F_1\, dy \wedge dz + F_2\, dz \wedge dx + F_3\, dx \wedge dy,$$

where $F_1, F_2, F_3 \colon \mathbb{R}^3 \to \mathbb{R}$ are smooth maps. The set of all differential 2-forms on $\mathbb{R}^3$ is denoted $\Omega^2(\mathbb{R}^3)$.

Example 2.5. Notice that we haven't defined the symbols $dz \wedge dy$, $dx \wedge dz$, or $dy \wedge dx$. But it's easy to do so by copying the definitions for the standard 2-forms. For instance,

$$(dy \wedge dx)(\mathbf{v}, \mathbf{w}) = dy(\mathbf{v})dx(\mathbf{w}) - dy(\mathbf{w})dx(\mathbf{v}).$$

Using this equation, how do we write $dz \wedge dy$, $dx \wedge dz$, and $dy \wedge dx$ as differential 2-forms, according to the definition? What do you think the symbols $dx \wedge dx$, $dy \wedge dy$, and $dz \wedge dz$ should mean?

As with 1-forms, we can associate a vector field to each differential 2-form on $\mathbb{R}^3$:

$$\eta = F_1\, dy \wedge dz + F_2\, dz \wedge dx + F_3\, dx \wedge dy \quad \Rightarrow \quad \mathbf{X}_\eta = F_1\,\mathbf{i} + F_2\,\mathbf{j} + F_3\,\mathbf{k}.$$

This vector field encodes the coefficients of η in terms of the standard 2-forms, and we can recover η by projection to the plane perpendicular to $\mathbf{X}_\eta$.

2.5 3-forms and beyond

So far we have **1-forms** and **2-forms**, which measure signed lengths and areas[2], respectively. Namely, a 1-form λ on $\mathbb{R}^3$ corresponds to some vector field $\mathbf{X}$ on $\mathbb{R}^3$, and the value $\lambda(\mathbf{v})$ measures the component of a vector $\mathbf{v}$ along $\mathbf{X}$. Similarly, a 2-form ω corresponds to some plane field[3] P, and $\omega(\mathbf{v}, \mathbf{w})$ measures the signed area of the parallelogram spanned by $\mathbf{v}$ and $\mathbf{w}$, when projected to P.

[2]This is only true because we're working in $\mathbb{R}^3$. In higher dimensions, a 2-form does not generally correspond to a signed area.
[3]That is, a choice of two-dimensional plane at every point in $\mathbb{R}^3$.

By analogy, a 3-form on $\mathbb{R}^3$ should measure 3-dimensional signed volumes. Just as we did for 1-forms and 2-forms we have a standard 3-form, called $dx \wedge dy \wedge dz$, and the value of

$$(dx \wedge dy \wedge dz)(\mathbf{u}, \mathbf{v}, \mathbf{w})$$

is the signed volume of the parallelepiped spanned by $\mathbf{u}, \mathbf{v}, \mathbf{w}$. By *signed* volume, we mean that this volume will be positive if $\mathbf{u}, \mathbf{v}, \mathbf{w}$ satisfies the right hand rule as a basis for $\mathbb{R}^3$, and will be negative if these vectors give a left-handed basis. Thankfully, linear algebra gives us a nice way of expressing this signed volume:

$$\boxed{(dx \wedge dy \wedge dz)(\mathbf{u}, \mathbf{v}, \mathbf{w}) := \det\begin{pmatrix}\mathbf{u} & \mathbf{v} & \mathbf{w}\end{pmatrix}.}$$

Just as with 1-forms and 2-forms, a general 3-form is then a linear combination of the single standard 3-form.

Definition. A **differential 3-form** on $\mathbb{R}^3$ is a linear combination over $C^\infty(\mathbb{R}^3)$ of the standard 3-form on $\mathbb{R}^3$. That is, a differential 3-form λ is an expression

$$\lambda = F\, dx \wedge dy \wedge dz,$$

where $F : \mathbb{R}^3 \to \mathbb{R}$ is a smooth map. The set of all differential 3-forms on $\mathbb{R}^3$ is denoted $\Omega^3(\mathbb{R}^3)$.

For the sake of completeness, we can now declare that a differential 0-form is simply a smooth function. That is, $\Omega^0(\mathbb{R}^3) = C^\infty(\mathbb{R}^3)$. We'll also declare that if $k > 3$, then 0 is the only k-form on $\mathbb{R}^3$. These statements probably seem silly for now, but they help us work out the algebra (and calculus) of differential forms.

2.6 Wedge product

Given a k-form λ and an ℓ-form η, we can produce a $(k+\ell)$-form $\lambda \wedge \eta$, essentially by naively trusting our notation. For instance, if λ and η are both 1-forms, then we can write

$$\lambda = f_1\, dx + f_2\, dy + f_3\, dz \quad \text{and} \quad \eta = g_1\, dx + g_2\, dy + g_3\, dz,$$

for some smooth functions f_i, g_j. Then

$$\begin{aligned}
\lambda \wedge \eta &= (f_1\, dx + f_2\, dy + f_3\, dz) \wedge (g_1\, dx + g_2\, dy + g_3\, dz) \\
&= f_1 g_1\, dx \wedge dx + f_2 g_1\, dy \wedge dx + f_3 g_1\, dz \wedge dx \\
&\quad f_1 g_2\, dx \wedge dy + f_2 g_2\, dy \wedge dy + f_3 g_2\, dz \wedge dy \\
&\quad f_1 g_3\, dx \wedge dz + f_2 g_3\, dy \wedge dz + f_3 g_3\, dz \wedge dz \\
&= \cdots = (f_2 g_3 - g_2 f_3)\, dy \wedge dz + (g_1 f_3 - f_1 g_3)\, dz \wedge dx + (f_1 g_2 - g_1 f_2)\, dx \wedge dy.
\end{aligned}$$

Now, just because we *can* push symbols around like this doesn't necessarily mean it's a good idea. Why is this a useful thing to do?

In dimension 3, our intuition can be quite helpful. Each of the 1-forms λ and η corresponds to a vector field, and evaluating the 1-form on a vector means asking how much that vector pushes in the direction of the corresponding vector field. If the vector fields corresponding to λ and η are linearly independent, then they span a plane in $\mathbb{R}^3$. But in $\mathbb{R}^3$, planes correspond to 2-forms! So $\lambda \wedge \eta$ should be the 2-form which corresponds to the plane built from $\mathbf{X}_\lambda$ and $\mathbf{X}_\eta$. Indeed, $\mathbf{X}_\lambda \times \mathbf{X}_\eta$ is a normal vector for the plane spanned by $\mathbf{X}_\lambda$ and $\mathbf{X}_\eta$, and the components of this vector match the coefficients of $\lambda \wedge \eta$:

$$\mathbf{X}_\lambda \times \mathbf{X}_\eta = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ f_1 & f_2 & f_3 \\ g_1 & g_2 & g_3 \end{vmatrix} = (f_2 g_3 - g_2 f_3)\mathbf{i} + (g_1 f_3 - f_1 g_3)\mathbf{j} + (f_1 g_2 - g_1 f_2)\mathbf{k} = \mathbf{X}_{\lambda \wedge \eta}.$$

The upshot is that the plane in which the area computations of $\lambda \wedge \eta$ take place is precisely the plane spanned by X_λ and X_η — the vectors along which the length computations of λ and η take place.

We can similarly work out a formula for the wedge product of a 1-form with a 2-form:

$$(f_1\,dx + f_2\,dy + f_3\,dz) \wedge (h_1\,dy \wedge dz + h_2\,dz \wedge dx + h_3\,dx \wedge dy) = \cdots = (f_1h_1 + f_2h_2 + f_3h_3)\,dx \wedge dy \wedge dz.$$

Since a 0-form is just a smooth function, the wedge product of a 0-form with a k-form is as simple as scalar multiplication: $f \wedge \lambda = f\,\lambda$. Because the wedge product of a k-form with an ℓ-form is a $(k+\ell)$-form, the remaining wedge products are uninteresting. For instance, the wedge product of a 3-form with a 2-form is a 5-form, which is necessarily 0, since we're focused on $\mathbb{R}^3$.

Finally, here are several facts about the wedge product that we might find useful. Unless otherwise stated, λ is a k-form and η is an ℓ-form.

- (forms are multi-linear) For any auxiliary vector field $\mathbf{w}$,

$$\lambda(\mathbf{v}_1, \ldots, \mathbf{v}_i, \ldots, \mathbf{v}_k) + \lambda(\mathbf{v}_1, \ldots, \mathbf{w}, \ldots, \mathbf{v}_k) = \lambda(\mathbf{v}_1, \ldots, \mathbf{v}_i + \mathbf{w}, \ldots, \mathbf{v}_k),$$

  where $\mathbf{w}$ is in the i^{th} position.

- (forms are alternating) If the i^{th} and j^{th} vectors in the expression $\lambda(\mathbf{v}_1, \ldots, \mathbf{v}_k)$ are swapped, the output is scaled by -1.

- (forms are alternating, again) If the vectors $\mathbf{v}_1, \ldots, \mathbf{v}_k$ are linearly dependent, then $\lambda(\mathbf{v}_1, \ldots, \mathbf{v}_k) = 0$. In particular, if $\mathbf{v}_i = \mathbf{v}_j$ for some $i \neq j$, then $\lambda(\mathbf{v}_1, \ldots, \mathbf{v}_k) = 0$.

- (dimension property) If $k > 3$, then $\lambda = 0$.

- (skew-commutativity of wedge product) $\eta \wedge \lambda = (-1)^{k\ell} \lambda \wedge \eta$

- (associativity of wedge product) For any $\lambda, \eta, \theta \in \Omega^\bullet(\mathbb{R}^3)$,

$$(\lambda \wedge \eta) \wedge \theta = \lambda \wedge (\eta \wedge \theta).$$

- (homogeneity of wedge product) If c is a scalar, then

$$(c\,\lambda) \wedge \eta = c\,(\lambda \wedge \eta) = \lambda \wedge (c\,\eta).$$

- (distributivity of wedge product) If η and θ have the same degree, then

$$\lambda \wedge (\eta + \theta) = \lambda \wedge \eta + \lambda \wedge \theta.$$

- (1-forms square to zero) If λ is a 1-form, then $\lambda \wedge \lambda = 0$.

- (wedge products of 1-forms) If λ and η are 1-forms, then

$$(\lambda \wedge \eta)(\mathbf{v}_1, \mathbf{v}_2) = \lambda(\mathbf{v}_1)\eta(\mathbf{v}_2) - \lambda(\mathbf{v}_2)\eta(\mathbf{v}_1).$$

You might be asked to prove some of these facts as exercises.

3 Differential Forms on $\mathbb{R}^3$ II

3.1 Exterior differentiation

We want to define a notion of differentiation for differential forms which at least moderately agrees with the philosophy that derivatives should measure how quickly a quantity is changing. The only possibly-artificial requirement we impose on our differentiation is that we want the derivative of a differential form to be another differential form.

0-forms

Let's start with the easiest case: 0-forms. A 0-form is just a smooth function $f : \mathbb{R}^3 \to \mathbb{R}$, and we're quite familiar with derivatives for smooth functions. Our preferred derivative here will be the *directional derivative*. That is, the derivative of a 0-form f is the 1-form $df : \mathbb{R}^3 \to \mathbb{R}$ defined by

$$df(\mathbf{v}) := D_{\mathbf{v}}f = \nabla f \cdot \mathbf{v} = \frac{\partial f}{\partial x}\,dx + \frac{\partial f}{\partial y}\,dy + \frac{\partial f}{\partial z}\,dz.$$

We call df the **total differential** of f.

> **Definition.** Let $f \in \Omega^0(\mathcal{R})$ be a 0-form on some open subset $\mathcal{R}$ of $\mathbb{R}^3$. The **derivative** (or **exterior derivative**) df of f is the total differential of f.

Exercise 3.1. Prove the product rule for total differentials: $d(fg) = g\,df + f\,dg$, for any smooth functions $f, g : \mathbb{R}^3 \to \mathbb{R}$.

1-forms

That wasn't so bad. We started with a 0-form f and asked, given a vector $\mathbf{v}$, how quickly f changes as we move in the direction (and at the speed) determined by $\mathbf{v}$. Maybe we can do something similar for a 1-form $\lambda \in \Omega^1(\mathcal{R})$. Let's ask the vague question, "How does λ vary as we move in the direction of $\mathbf{v}$?"

The first problem we run into is nailing down what we mean by, "How does λ vary?" In order to get a real number out of λ, we need to plug in a vector — this is in addition to the vector $\mathbf{v}$ along which we're taking a derivative. So we realize that *the derivative of a 1-form will eat two vectors and output a real number*. Since we expect derivatives of differential forms to be differential forms, we conclude that the derivative of a 1-form will be a 2-form.

Remark. One point about which we could have been more careful is that a differential k-form on $\mathcal{R}$ doesn't just eat k vectors — it eats a point p in $\mathcal{R}$, and then k vectors based at p. A typical notation for this is

$$\eta_p(\mathbf{v}_1, \ldots, \mathbf{v}_k),$$

to indicate that the vectors $\mathbf{v}_1, \ldots, \mathbf{v}_k$ are all based at p. So if $\mathbf{w}$ is a vector field on $\mathcal{R}$, then we can think of the quantity $\lambda(\mathbf{w})$ as a smooth function $\mathcal{R} \to \mathbb{R}$: for each point $p \in \mathcal{R}$, we compute $\lambda_p(\mathbf{w}(p))$. This perspective should be helpful as we define the derivative.

A reasonable first attempt at defining $d\lambda(\mathbf{v}, \mathbf{w})$ would treat $\mathbf{w}$ as a (constant) vector field and simply compute the directional derivative $D_{\mathbf{v}}(\lambda(\mathbf{w}))$. Since $\lambda(\mathbf{w})$ is a smooth function, we can take its directional derivative in the direction of $\mathbf{v}$, so this makes sense. The problem is that $D_{\mathbf{v}}(\lambda(\mathbf{w}))$ treats the vectors $\mathbf{v}$ and $\mathbf{w}$ very unequally. We want $d\lambda$ to be a 2-form, but there's no reason to expect the expression $D_{\mathbf{v}}(\lambda(\mathbf{w}))$ to be alternating in the vectors $\mathbf{v}$ and $\mathbf{w}$. We fix this by taking the *alternatization* of the expression.

> **Definition.** Let $\lambda \in \Omega^1(\mathcal{R})$ be a differential 1-form on some open subset $\mathcal{R}$ of $\mathbb{R}^3$. The **derivative** (or

exterior derivative) $d\lambda$ of λ is a 2-form on $\mathcal{R}$ defined by

$$d\lambda(\mathbf{v},\mathbf{w}) := D_{\mathbf{v}}(\lambda(\mathbf{w})) - D_{\mathbf{w}}(\lambda(\mathbf{v})),$$

for vectors $\mathbf{v},\mathbf{w}$, treated as constant vector fields on $\mathcal{R}$.

Remark. In a more abstract setting, it's not as easy to just pretend that vectors are in fact vector fields; i.e., there's not generally such a thing as a constant vector field. Writing down this definition of the derivative in such a setting is a little more work.

Exercise 3.2. Show that exterior differentiation is a linear map $d\colon \Omega^1(\mathcal{R}) \to \Omega^2(\mathcal{R})$. That is, show that if $c \in \mathbb{R}$ is a scalar and λ, η are 1-forms on $\mathcal{R}$, then

$$d(c\,\lambda) = c\,d\lambda \quad \text{and} \quad d(\lambda+\eta) = d\lambda + d\eta.$$

Note: It's important that c is a constant scalar, and not a scalar function.

Example 3.3. Let's compute the derivative of a standard 1-form on $\mathbb{R}^3$. Take $\mathbf{v},\mathbf{w} \in \mathbb{R}^3$. Then

$$d(dx)(\mathbf{v},\mathbf{w}) = D_{\mathbf{v}}(dx(\mathbf{w})) - D_{\mathbf{w}}(dx(\mathbf{v})).$$

But the function $dx(\mathbf{w})$ is constant, so $D_{\mathbf{v}}(dx(\mathbf{w})) = 0$. We similarly find that $D_{\mathbf{w}}(dx(\mathbf{v})) = 0$, so $\boxed{d(dx) = 0}$. The same argument shows that $\boxed{d(dy) = 0}$ and $\boxed{d(dz) = 0}$.

Example 3.4. Together, Exercise 3.2 and Example 3.3 might worry us. We've shown that the derivative of a standard 1-form is 0, and that differentiation is linear. Since an arbitrary 1-form is a linear combination (over $C^\infty(\mathcal{R})$) of standard 1-forms, we might be tempted to think that the derivative of a 1-form is always 0. Instead, we have a product rule. Suppose we have a 0-form f on $\mathcal{R} \subset \mathbb{R}^n$. Then for any vectors $\mathbf{v},\mathbf{w} \in \mathbb{R}^3$,

$$d(f\,dx)(\mathbf{v},\mathbf{w}) = D_{\mathbf{v}}(f\,dx(\mathbf{w})) - D_{\mathbf{w}}(f\,dx(\mathbf{v})).$$

But again, $dx(\mathbf{w})$ doesn't depend on $\mathbf{v}$, so our product rule for directional derivatives tells us that

$$\begin{aligned} d(f\,dx)(\mathbf{v},\mathbf{w}) &= D_{\mathbf{v}}(f)\,dx(\mathbf{w}) - D_{\mathbf{w}}(f)\,dx(\mathbf{v}) \\ &= df(\mathbf{v})\,dx(\mathbf{w}) - df(\mathbf{w})\,dx(\mathbf{v}) \\ &= (df \wedge dx)(\mathbf{v},\mathbf{w}). \end{aligned}$$

So we find that $\boxed{d(f\,dx) = df \wedge dx}$. We similarly see that $\boxed{d(f\,dy) = df \wedge dy}$ and $\boxed{d(f\,dz) = df \wedge dz}$.

The upshot is the following formula for the derivative of a 1-form:

$$\boxed{d(f\,dx + g\,dy + h\,dz) = df \wedge dx + dg \wedge dy + dh \wedge dz.}$$

We can unwind the right hand side using the definition of the total differential and the anti-commutativity of the wedge product.

2-forms and 3-forms

Now that we've differentiated 1-forms, we could run through the same abstraction for 2- and 3-forms. Instead, for the sake of brevity, we'll skip the abstraction and give formulas. We have

$$\boxed{d(f\,dy \wedge dz + g\,dz \wedge dx + h\,dx \wedge dy) = df \wedge dy \wedge dz + dg \wedge dz \wedge dz + dh \wedge dx \wedge dy.}$$

Once again, the right hand side can be unwound using the definition of the total differential and the anti-commutativity of the wedge product. The formula for 3-forms is even easier: since the derivative of a k-form should be a $(k+1)$-form and 0 is the only 4-form on $\mathbb{R}^3$, we have

$$\boxed{d(f\,dx \wedge dy \wedge dz) = df \wedge dx \wedge dy \wedge dz = 0.}$$

Remark. The linearity of differentiation for 2- and 3-forms is built into the definition we've given here.

Example 3.5. As was the case for 1-forms, the standard 2-forms have derivative 0:

$$d(dx \wedge dy) = d(1 \cdot dx \wedge dy) = d(1) \wedge dx \wedge dy = 0 \wedge dx \wedge dy.$$

Exercise 3.6. Verify that exterior differentiation satisfies the **Leibniz rule**:

$$d(\lambda \wedge \eta) = d\lambda \wedge \eta + (-1)^{\deg(\lambda)} \lambda \wedge d\eta,$$

for any k-form λ and any ℓ-form η. (You should only have to check a few cases.)

Exercise 3.7. Prove that if λ is a k-form, then $d(d\lambda) = 0$. This is sometimes written as $d^2 = 0$. (Since we're working in $\mathbb{R}^3$, the cases $k = 2$ and $k = 3$ should basically be free.) This fact will be *hugely* important for us.

3.2 Some special differential forms

Remember that a parametrized curve $\gamma(t) = (x(t), y(t), z(t))$ in $\mathbb{R}^3$ is said to be *Legendrian* if it satisfies the differential equation

$$z'(t) - y(t)x'(t) = 0.$$

Another way to express this condition is to say that we have $\alpha(\dot{\gamma}(t)) = 0$, where

$$\boxed{\alpha := dz - y\,dx}$$

is the **standard contact form** on $\mathbb{R}^3$. Hopefully we'll have a chance to say more about what makes this form so special a little later, but for now let's just give the definition of a contact form on $\mathbb{R}^3$, punting on the motivation.

Definition. We say that a 1-form λ on $\mathbb{R}^3$ is a (positive) **contact form** if

$$\lambda \wedge d\lambda = f\, dx \wedge dy \wedge dz,$$

for some function $f > 0$.

Exercise 3.8. Check that the differential equation $\alpha(\dot{\gamma}(t)) = 0$ is equivalent to $z'(t) - y(t)x'(t) = 0$.

Exercise 3.9. Compute $\alpha \wedge d\alpha$, where α is the standard contact form on $\mathbb{R}^3$.

Exercise 3.10. Check that $dz + r^2\, d\theta$ is a contact form on $\mathbb{R}^3$, where (r, θ, z) are cylindrical coordinates on $\mathbb{R}^3$. That is, $x = r\cos\theta$ and $y = r\sin\theta$.

We can define contact forms on any odd-dimensional Euclidean space $\mathbb{R}^{2n+1}$ (indeed, on any odd-dimensional manifold). The analogue for even-dimensional spaces is a symplectic form. Though we haven't discussed differential forms on $\mathbb{R}^4$, let's assume that we can get there by extending our knowledge of forms on $\mathbb{R}^3$.

Definition. Consider $\mathbb{R}^4$ with coordinates (x_1, y_1, x_2, y_2). A 2-form ω on $\mathbb{R}^4$ is said to be **symplectic** if

1. $d\omega = 0$;
2. $\omega \wedge \omega = f\, dx_1 \wedge dy_1 \wedge dx_2 \wedge dy_2$, for some function $f > 0$.

Exercise 3.11. Show that $\omega_0 = dx_1 \wedge dy_1 + dx_2 \wedge dy_2$ is a symplectic form on $\mathbb{R}^4$. This is called the **standard symplectic form** on $\mathbb{R}^4$.

Exercise 3.12. Let $\mathbb{R}^4$ have coordinates (s, x, y, z). Given that α is a contact form on $\mathbb{R}^3$ (with coordinates (x, y, z)), show that $\omega = d(e^s\,\alpha)$ is a symplectic form on $\mathbb{R}^4$.

3.3 Integration

When we first introduced 1-forms, we saw that they could be integrated over parametrized curves $\mathbf{r}\colon [a,b] \to \mathbb{R}^3$. Indeed, many mathematicians describe differential forms as "things to be integrated." (Personally, I think that sells forms short a bit.) We won't have a ton of reason to integrate 2- and 3-forms this summer, so we'll just give a quick-and-dirty definition for integration here.

Just as 1-forms are integrated over 1-dimensional objects (curves), 2-forms are integrated over 2-dimensional objects (surfaces). Also as in the 1-dimensional case, we'll take the easy way out by assuming that we have a parametrization of our surface, though this isn't necessary.

Definition. Let $F\colon \mathcal{D} \to \mathcal{S}$ be a smooth parametrization[a] of a surface $\mathcal{S} \subset \mathbb{R}^3_{x,y,z}$ with domain $\mathcal{D} \subset \mathbb{R}^2_{u,v}$. Assume that the vectors

$$F_u := \frac{\partial F}{\partial u} \quad \text{and} \quad F_v := \frac{\partial F}{\partial v}$$

are everywhere linearly independent. Then, for any 2-form λ on $\mathbb{R}^3$,

$$\int_{\mathcal{S}} \lambda := \iint_{\mathcal{D}} \lambda_{F(u,v)}(F_u, F_v)\, du\, dv.$$

[a] We're intentionally not defining the word 'parametrization' right now.

Notice that the integral on the right side of the defining equation is a usual double integral.

Exercise 3.13. Let $\mathcal{S} \subset \mathbb{R}^3$ be the surface parametrized by

$$F(u,v) = (\sin u \cos v, \sin u \sin v, \cos u), \quad 0 \le u \le \pi, 0 \le v \le 2\pi.$$

Show that

$$\int_{\mathcal{S}} \nu = 4\pi,$$

where

$$\nu = x\, dy \wedge dz + y\, dz \wedge dx + z\, dx \wedge dy.$$

Can you identify the surface $\mathcal{S}$?

Because we've restricted ourselves to $\mathbb{R}^3$, integrating 3-forms will be particularly easy. A 3-form ought to be integrated over a 3-dimensional object, which will simply[1] mean a region in $\mathbb{R}^3$ with nonempty interior. Because a 3-form on $\mathbb{R}^3$ is determined by a function, we can integrate the 3-form by just integrating that function.

Definition. Let $\mathcal{R} \subset \mathbb{R}^3_{x,y,z}$ be a region with nonempty interior, and let $\lambda = f\, dx \wedge dy \wedge dz$ be a 3-form on $\mathbb{R}^3$. Then

$$\int_{\mathcal{R}} \lambda = \int_{\mathcal{R}} f\, dx \wedge dy \wedge dz := \iiint_{\mathcal{R}} f\, dx\, dy\, dz,$$

where the integral on the right is a usual triple integral.

Exercise 3.14. Let $\mathcal{B}$ be the ball of radius 1 centered at the origin in $\mathbb{R}^3$, and compute

$$\int_{\mathcal{B}} d\nu,$$

where ν is the 2-form identified in Exercise 3.13.
Hint: You shouldn't need to evaluate a triple integral if you remember the volume of a ball.

[1] It's not as simple as we're pretending, but the details aren't really worth the trouble for our purposes.

Having integrated 1-, 2-, and 3-forms, we might as well have a definition for the integration of 0-forms. Of course, these are just smooth functions, and should be integrated over 0-dimensional objects — i.e., points. To integrate a function over a point, we just evaluate the function at that point.

Definition. Fix a point $p_0 = (x_0, y_0, z_0) \in \mathbb{R}^3$ and consider a 0-form f on $\mathbb{R}^3$. Then

$$\int_{p_0} f = f(x_0, y_0, z_0).$$

Remark. One caveat to the above definition has to do with signs. Since we integrate over *parametrized* curves and surfaces, we should also allow our points to somehow be "parametrized." All this really means is assigning a sign ± to the point. If we're thinking of the point as negative, then this negative sign carries over into the integral.

Now that we've defined integration over points, curves, surfaces, and regions, we can write down a single equation which encodes the various versions of the fundamental theorem of integral calculus that you've learned in your mathematical career. These may have been called the fundamental theorem of line integrals, Stokes' theorem, or the divergence theorem, but all of them have the property of exchanging a derivative for a boundary.

Theorem 3.15: Stokes' theorem

Let λ be a k-form, with $k = 0$, 1, or 2. Let $\mathcal{W}$ be:

- a curve in $\mathbb{R}^3$, if $k = 0$;
- a surface in $\mathbb{R}^3$, if $k = 1$;
- a region in $\mathbb{R}^3$, if $k = 2$.

Assuming that $\partial\mathcal{W}$ is nice[a],

$$\int_{\mathcal{W}} d\lambda = \int_{\partial\mathcal{W}} \lambda.$$

[a] Whatever that means.

Remark. It's very important to be careful about orientations when using Stokes' theorem. For a parametrized curve $\mathbf{r}(t)$, $a \le t \le b$, we must treat $\mathbf{r}(b)$ as positive and $\mathbf{r}(a)$ as negative. For a parametrized surface $F : \mathcal{D} \to \mathcal{S}$, the boundary curve $\partial\mathcal{S}$ carries the orientation of $\partial\mathcal{D}$ (mapped over by F). The boundary $\partial\mathcal{R}$ of a region $\mathcal{R}$ is oriented using an outward-pointing normal vector.

Exercise 3.16. Use Stokes' theorem, as stated above, to recover the fundamental theorem of line integrals, Stokes' theorem (as you would have seen in a calculus course), and the divergence theorem.

4 Differential Graded Algebras I

Given an open subset $\mathcal{R}$ of $\mathbb{R}^3$, we can think about $\Omega^k(\mathcal{R})$, the collection of k-forms on $\mathcal{R}$. Just like $\Omega^k(\mathbb{R}^3)$, $\Omega^k(\mathcal{R})$ is a module generated by the standard k-forms (which are unchanged); the difference is that $\Omega^k(\mathcal{R})$ is a module over $C^\infty(\mathcal{R})$, instead of $C^\infty(\mathbb{R}^3)$. On the one hand, simply restricting the domains of our scalar functions doesn't sound that interesting. But it turns out that the collection of differential forms on $\mathcal{R}$ records some deep information about $\mathcal{R}$.

Let's denote by $\Omega^\bullet(\mathcal{R})$ the collection of **formal sums** of differential forms on $\mathcal{R}$. That is, an element of $\Omega^\bullet(\mathcal{R})$ looks like

$$\lambda_0 + \lambda_1 + \lambda_2 + \lambda_3 \in \Omega^\bullet(\mathcal{R}),$$

where λ_k is a k-form, for $k = 0, 1, 2, 3$. It doesn't make much technical sense to add, say, a 1-form to a 2-form. But that's what makes these *formal* sums — they have the *form* of a sum, even if that doesn't make perfect sense.

In order to extract any information about $\mathcal{R}$ from $\Omega^\bullet(\mathcal{R})$, we'll need to think about the algebraic structure carried by $\Omega^\bullet(\mathcal{R})$ — what sort of object do we have?

- By definition, $\Omega^\bullet(\mathcal{R})$ has an addition operation with an obvious additive inverse. Moreover, we have scalar multiplication by elements of $C^\infty(\mathcal{R})$, just by using the scalar multiplication on each $\Omega^k(\mathcal{R})$. Because the addition and scalar multiplication are compatible, $\Omega^\bullet(\mathcal{R})$ is a $C^\infty(\mathcal{R})$-module.
- Now we'll start describing algebraic structures you may not have seen before, but which we'll soon define. Wedge product gives us a way of multiplying two elements of $\Omega^\bullet(\mathcal{R})$. We know how to compute the wedge product of a pair of differential forms, and we can define the wedge product of a pair of formal sums by forcing the product to be distributive. A module which has a multiplication that plays nicely with addition and scalar multiplication is called an **algebra**, so we see that $\Omega^\bullet(\mathcal{R})$ is a $C^\infty(\mathcal{R})$-algebra.
- The fact that we can write our formal sums out in their constituent pieces is important. This allows us to write
  $$\Omega^\bullet(\mathcal{R}) = \Omega^0(\mathcal{R}) \oplus \Omega^1(\mathcal{R}) \oplus \Omega^2(\mathcal{R}) \oplus \Omega^3(\mathcal{R}).$$
  We refer to honest differential forms as **homogeneous** terms, and the above decomposition tells us that we can write every element of $\Omega^\bullet(\mathcal{R})$ uniquely as a sum of homogeneous terms. Notice that the wedge product gives us maps
  $$\wedge : \Omega^i(\mathcal{R}) \times \Omega^j(\mathcal{R}) \to \Omega^{i+j}(\mathcal{R}),$$
  and that each $\Omega^k(\mathcal{R})$ is an additive group. These two facts tell us that $\Omega^\bullet(\mathcal{R})$ has the structure of a *graded* $C^\infty(\mathcal{R})$-algebra. The "grading" refers to our ability to identify elements of various degree in $\Omega^\bullet(\mathcal{R})$.
- A final bit of structure comes from some facts you've already checked. We know that $\Omega^\bullet(\mathcal{R})$ carries a map $d : \Omega^\bullet(\mathcal{R}) \to \Omega^\bullet(\mathcal{R})$, built by distributing the derivative map across sums. This map has three important properties:
  - d increases grading by 1;
  - $d \circ d = 0$;
  - $d(\lambda \wedge \eta) = (d\lambda) \wedge \eta + (-1)^{\deg \lambda} \lambda \wedge (d\eta)$.

  Recall that the last equation is called the Leibniz rule. This bit of structure might not seem worth naming, but it turns out that many interesting graded algebras carry such a map. Because of this map, we say that $\Omega^\bullet(\mathcal{R})$ is a differential graded algebra over $C^\infty(\mathcal{R})$.

So $\Omega^\bullet(\mathcal{R})$ carries a *lot* of structure. An upcoming goal of ours is to define an invariant of Legendrian knots which carries the same structure, so now we want to define more carefully the words we used above. In the second DGAs talk, we wil learn how to extract more computable information from the structure carried by a DGA.

We begin with a notion you have likely encountered in an algebra course.

Definition. Let R be a commutative ring with multiplicative unit 1_R. An R**-module** or a **module over** R is a set M together with

1. a binary operation $+$ on M so that $(M,+)$ is an abelian group;
2. a map $R \times M \to M$, denoted by rm, which satisfies:
   (a) $(r+s)m = rm + sm$, for all $r,s \in R$, $m \in M$;
   (b) $(rs)m = r(sm)$, for all $r,s \in R$, $m \in M$;
   (c) $r(m+n) = rm + rn$, for all $r \in R$, $m,n \in M$;
   (d) $1_R m = m$, for all $m \in M$.

Remark. It's not actually necessary to assume that R is commutative or unital, but then we'd have to use adjectives like *left* and *right* that we'd rather avoid for now. While our *algebras* will not typically be commutative, their underlying rings will be commutative.

Example 4.1. If R is a field, then the conditions for M to be a module are precisely the conditions for M to be a vector space. This should guide our intuition: we think of modules as vector spaces where the set of scalars is only required to be a ring, rather than a field.

Example 4.2. For any ring R with a multiplicative unit, we can define a module

$$R^n = \{(a_1, a_2, \dots, a_n) \,|\, a_k \in R, 1 \le k \le n\},$$

which we call **the free module of rank** n **over** R. Addition and scalar multiplication are defined component-wise (recall that the additive structure of a ring is always commutative). This example should reinforce our intuition that modules are to rings as vector spaces are to fields.

Exercise 4.3. Show that any abelian group G can be made into a $\mathbb{Z}$-module as follows: first, write the group operation of G as $+$. Then, for any $n \in \mathbb{Z}$ and $g \in G$, define

$$ng = \begin{cases} g + g + \cdots + g \ (n \text{ times}), & n > 0 \\ 0_G, & n = 0 \\ -g - g - \cdots - g \ (-n \text{ times}), & n < 0 \end{cases},$$

where 0_G is the additive identity of G.

Definition. Let R be a commutative, unital ring, and let A be an R-module. We say that A is an **algebra** over R if there is a binary operation $\cdot : A \times A \to A$ such that, for any $r,s \in R$ and any $x,y,z \in A$ we have:

1. $(x+y)\cdot z = x \cdot z + y \cdot z$;
2. $z \cdot (x+y) = z \cdot x + z \cdot y$;
3. $(rx)\cdot(sy) = (rs)(x \cdot y)$.

More simply, an algebra over R is an R-module with an R-bilinear binary operation.

Example 4.4. The complex numbers $\mathbb{C}$ can be treated as a vector space over $\mathbb{R}$, and thus constitute an $\mathbb{R}$-module. We also know how to multiply complex numbers with each other, and it's easy to check that this multiplication plays nicely with scalar multiplication by real numbers. So $\mathbb{C}$ is an example of an associative, commutative $\mathbb{R}$-algebra.

Example 4.5. Complex multiplication gives an $\mathbb{R}$-algebra structure to $\mathbb{R}^2$; the cross product endows $\mathbb{R}^3$ with such a structure. However, this $\mathbb{R}$-algebra is neither associative nor commutative. You will recall that the cross product is anti-commutative: $\mathbf{v} \times \mathbf{w} = -\mathbf{w} \times \mathbf{v}$. To see that associativity fails, notice that

$$(\mathbf{i} \times \mathbf{i}) \times \mathbf{j} = \mathbf{0} \times \mathbf{j} = \mathbf{0},$$

while

$$\mathbf{i} \times (\mathbf{i} \times \mathbf{j}) = \mathbf{i} \times \mathbf{k} = -\mathbf{k}.$$

Example 4.6. Not every notion of vector product will give $\mathbb{R}^3$ the structure of an algebra. For instance, we also know how to take the dot product of a pair of vectors in $\mathbb{R}^3$ (or, more generally, $\mathbb{R}^n$). But the output of this binary operation is a scalar, rather than another vector.

Example 4.7. For any commutative, unital ring R, we can consider the R-module $M_n(R)$ of $n \times n$ matrices with entries in R. Matrix multiplication gives $M_n(R)$ the structure of an R-algebra which is associative, but not commutative.

Example 4.8. For any commutative, unital ring R, the polynomial ring $R[x]$ is an associative, commutative R-algebra. Here

$$R[x] := \{r_0 + r_1 x + r_2 x^2 + \cdots + r_n x^n \mid n \in \mathbb{N} \text{ and } r_k \in R, 0 \leq k \leq n\},$$

and the algebraic operations are defined in the obvious way.

Definition. A **grading** on a ring R is a decomposition

$$R = \bigoplus_{n=-\infty}^{\infty} R_n = \cdots \oplus R_{-2} \oplus R_{-1} \oplus R_0 \oplus R_1 \oplus R_2 \oplus \cdots$$

of R as a direct sum such that $R_m R_n \subseteq R_{m+n}$. We call a ring R along with a choice of grading a **graded ring**, and we refer to a nonzero element of R_n as a **homogeneous element of degree** n.

Definition. A **graded R-algebra** is an R-algebra which is graded as a ring. That is, an R-algebra A with a decomposition

$$A = \bigoplus_{n=-\infty}^{\infty} A_n = \cdots \oplus A_{-2} \oplus A_{-1} \oplus A_0 \oplus A_1 \oplus A_2 \oplus \cdots$$

as a direct sum of R-modules such that $A_m A_n \subseteq A_{m+n}$. Elements of R are given grading 0, so that $R \subseteq A_0$.

Example 4.9. For any commutative, unital ring R, the polynomial ring $R[x]$ has a natural grading. (As usual so far, we don't strictly need R to be commutative and unital, but always make this assumption.) In particular, we can define

$$R_n[x] := \{r x^n \mid r \in R\},$$

for every integer $n \geq 0$, and $R_n[x] = 0$ for $n < 0$. We then find that

$$R[x] = \bigoplus_{n=0}^{\infty} R_n[x].$$

You may recall that the polynomial $r x^n$ is said to be *homogeneous*, and its *degree* is n. This terminology is reflected in the definition of graded rings. Finally, we see that

$$R_n[x] R_m[x] = \{(r x^n)(s x^m) \mid r, s \in R\} = \{(rs) x^{n+m} \mid r, s \in R\} = R_{n+m}[x].$$

(The last equality uses the fact that R is unital. But we don't need strict equality; subset inclusion is fine.) So $R[x]$ is indeed a graded algebra over R, with the grading given by the degree of homogeneous polynomials.

Definition. A **chain complex** $(A_\bullet, \partial_\bullet)$ consists of a doubly infinite sequence of abelian groups

$$\ldots, A_{-2}, A_{-1}, A_0, A_1, A_2, \ldots,$$

along with homomorphisms (called **boundary operators**)

$$\partial_n : A_n \to A_{n-1}$$

such that $\partial_n \circ \partial_{n+1} = 0$. Similarly, a **cochain complex** $(A_\bullet, d_\bullet)$ consists of a doubly infinite sequence of abelian groups $\ldots, A_{-2}, A_{-1}, A_0, A_1, A_2, \ldots$ along with homomorphisms (called **differentials**)

$$d_n : A_n \to A_{n+1}$$

such that $d_{n+1} \circ d_n = 0$.

Example 4.10. Here's a very important example that we won't try to make precise. For $0 \le k \le 3$, we want $\mathcal{C}_k$ to be an abelian group generated by all the connected "shapes" of dimension k in $\mathbb{R}^3$. A 0-dimensional shape is a point, a 1-dimensional shape is a curve, a 2-dimensional shape is a surface, and a 3-dimensional shape is a region. We really ought to be more careful than this, but let's just give an impression of what's going on. Also, to make sure that $\mathcal{C}_k$ is a group, we have to take formal sums of shapes, whatever that means. What we really want to observe is that we have a map

$$\partial : \mathcal{C}_k \to \mathcal{C}_{k-1},$$

obtained by taking the boundary of our shape. (For $k \neq 0, 1, 2, 3$, we'll define $\mathcal{C}_k$ to be the trivial group.) The boundary of a 3-dimensional region is a collection of surfaces, the boundary of a surface is a (possibly empty) collection of curves, and so on. So we have a chain of homomorphisms

$$\cdots \longleftarrow 0 \xleftarrow{\partial_0 = \partial} \mathcal{C}_0 \xleftarrow{\partial_1 = \partial} \mathcal{C}_1 \xleftarrow{\partial_2 = \partial} \mathcal{C}_2 \xleftarrow{\partial_3 = \partial} \mathcal{C}_3 \longleftarrow 0 \longleftarrow \cdots$$

It's not hard to convince ourselves that a curve which arises as the boundary of a surface will itself have empty boundary, and the same is true for a surface which is the boundary of a region. (This is where our lack of care about the meaning of "shape" haunts us.) Even though we haven't made this example precise, we bring it up because it's the original example which motivated the algebraic definition of chain complexes in the first place.

Example 4.11. We also have an example of a cochain complex: $(\Omega^\bullet(\mathcal{R}), d)$. This gives us a sequence

$$\cdots \longrightarrow 0 \longrightarrow \Omega^0(\mathcal{R}) \xrightarrow{d_0 = d} \Omega^1(\mathcal{R}) \xrightarrow{d_1 = d} \Omega^2(\mathcal{R}) \xrightarrow{d_2 = d} \Omega^3(\mathcal{R}) \xrightarrow{d_3 = d} 0 \longrightarrow \cdots$$

which satisfies $d_{n+1} \circ d_n = 0$. (For $k \neq 0, 1, 2, 3$, $\Omega^k(\mathcal{R})$ is the trivial group.)

Example 4.12. It might be tempting to think that the abelian groups $R_n[x]$, seen in Example 4.8, naturally form a cochain complex with the usual derivative. But the usual derivative does not square to 0, so $(R_\bullet[x], \frac{d}{dx})$ is not a cochain complex.

We've now seen that $\Omega^\bullet(\mathcal{R})$ has both the structure of a graded algebra and the structure of a cochain complex. Moreover, the cochain complex structure has a special compatibility with the multiplication coming from the algebra: it satisfies the Leibniz rule, which is a graded version of the product rule. We have a special name for this amount of structure.

Definition. Let R be a commutative, unital ring. A **differential graded algebra over** R is a graded algebra over R, along with a (co)chain complex structure whose boundary operators/differentials satisfy the Leibniz rule. That is, a graded R-algebra A along with a map $d : A \to A$ which raises (or lowers) degree by 1 and satisfies the equations $d \circ d = 0$ and

$$d(x \cdot y) = (dx) \cdot y + (-1)^{\deg(x)} x \cdot (dy)$$

for all homogeneous elements $x, y \in A$.

5 Chekanov-Eliashberg DGA I

Throughout this talk, Λ will be an oriented Legendrian knot in $(\mathbb{R}^3, \xi_{\text{std}})$. Our goal is to define the **Chekanov-Eliashberg DGA** $(\mathcal{A}_\Lambda, \partial_\Lambda)$ of Λ. Let us point out that different versions of $(\mathcal{A}_\Lambda, \partial_\Lambda)$ exist in the literature; we will define the *fully noncommutative* version here, and make some comments about its relation to other versions.

We also assume that the only singularities which exist in the Lagrangian projection $\Pi(\Lambda)$ are transverse double points. Moreover, we assume that the strands of $\Pi(\Lambda)$ intersect orthogonally at double points. This can always be accomplished via a generic Legendrian isotopy. Finally, our definition of $(\mathcal{A}_\Lambda, \partial_\Lambda)$ will also require fixing a base point $*$ on Λ which does not project to a double point in $\Pi(\Lambda)$.

As we saw in the DGAs talks, the structures of a DGA are layered: we first define the algebra, then the grading, and finally the differential.

5.1 The algebra

The algebra $\mathcal{A}_\Lambda$ is a unital $\mathbb{Z}$-algebra, generated by the double points of $\Pi(\Lambda)$, plus an auxiliary variable t and its inverse t^{-1}. If we label the double points of $\Pi(\Lambda)$ $a_1, \dots, a_n$, then we may write

$$\mathcal{A}_\Lambda = \mathbb{Z}\langle a_1, \dots, a_n, t^{\pm}\rangle.$$

Let us emphasize that $a_1, \dots, a_n, t, t^{-1}$ are the generators of $\mathcal{A}_\Lambda$ *as an algebra*, and that the only relation among these generators is $t \cdot t^{-1} = t^{-1} \cdot t = 1$. So we build elements of $\mathcal{A}_\Lambda$ by multiplying the generators together and summing the resulting products: for instance

$$a_1 t a_3 t^{-1} a_1 t a_2 + 1 + a_7 a_8 \in \mathcal{A}_\Lambda.$$

(Remember that $\mathcal{A}_\Lambda$ is unital by definition; this is what allows us to include 1 in our sum.)

While $\mathcal{A}_\Lambda$ is finitely-generated as an algebra, it has infinitely many generators *as a module*. The generators of $\mathcal{A}_\Lambda$ as a $\mathbb{Z}$-module are all of the **words** in the letters $a_1, \dots, a_n, t, t^{-1}$. The only relation among these letters is the one given above: $t \cdot t^{-1} = t^{-1} \cdot t = 1$. Otherwise, these letters are noncommuting, and do not have multiplicative inverses.

Remark. We're going to define the Chekanov-Eliashberg DGA completely combinatorially, so it's not necessary to know why the double points of $\Pi(\Lambda)$ might be involved in the definition of an invariant of Λ; but it's an interesting topic, once you learn a bit more geometry and topology. A generally productive way to study a space X is to study the critical points of a function $f : X \to \mathbb{R}$. When X is a smooth manifold and f is a sufficiently nice function, this is the purview of *Morse theory*. But if X is the space of all possible paths $\gamma : [0,1] \to \mathbb{R}^3$ with $\gamma(0)$ and $\gamma(1)$ on Λ — that is, all possible paths connecting a pair of points on Λ — then we can define a function $A : X \to \mathbb{R}$ called the *action*. The action of γ is defined by

$$A(\gamma) := \int_\gamma dz - y\,dx.$$

With a little work, we can show that γ is a critical point of the action function(al) if and only if γ is a Reeb chord. So the algebra $\mathcal{A}_\Lambda$ that we've defined combinatorially is really trying to capture some sort of infinite-dimensional Morse theory.

Example 5.1. Consider the standard Legendrian unknot Λ. (Diagrams included in the lecture.) The Lagrangian projection of Λ has a single double point, so

$$\mathcal{A}_\Lambda = \mathbb{Z}\langle a, t^{\pm}\rangle,$$

where a represents the unique double point. We emphasize once again that the elements of $\mathcal{A}_\Lambda$ are formal sums of words in the letters a, t, t^{-1}, with the only relation being $tt^{-1} = t^{-1}t = 1$.

5.2 The grading

Now we'll define the degree of each of the algebra generators of $\mathcal{A}_\Lambda$; that is, for each of the elements $a_1, \ldots, a_n, t, t^{-1}$. The degree of a word in these generators is then the sum of the degrees of the letters:

$$|a_1 t a_4| = |a_1| + |t| + |a_4|.$$

We can then write $\mathcal{A}_\Lambda^n$ for the module generated by the words of degree n, and find that

$$\mathcal{A}_\Lambda = \bigoplus_{n=-\infty}^{\infty} \mathcal{A}_\Lambda^n$$

gives a grading on $\mathcal{A}_\Lambda$.

We use the *Maslov grading* for the auxiliary variable t, simply defining

$$|t| = -2\,\mathrm{rot}(\Lambda) \quad \text{and} \quad |t^{-1}| = 2\,\mathrm{rot}(\Lambda).$$

(Because t^{-1} is the multiplicative inverse of t, the grading of t^{-1} is determined by that of t.)

Next, we need to define the grading/degree of each generator a_i. To do this, let α_i be the path in $\Pi(\Lambda)$ which runs from the overcrossing of a_i to the undercrossing of a_i and does not intersect the chosen basepoint $*$. This is called a **capping path** for a_i, and we let $\mathrm{rot}(\alpha_i) \in \mathbb{R}$ denote the number of counterclockwise rotations its tangent vector makes as we move from the overcrossing to the undercrossing. Our assumption that all crossings in $\Pi(\Lambda)$ are orthogonal ensures that $\mathrm{rot}(\alpha_i)$ will be an odd multiple of $1/4$, which allows us to define

$$|a_i| = 2\,\mathrm{rot}(\alpha_i) - 1/2.$$

Example 5.2. Let's return to the standard Legendrian unknot from before. This knot has rotation number 0, so $|t| = -2\mathrm{rot}(\Lambda) = 0$, and thus $|t^{-1}| = 0$. As we move from the overcrossing of a to the undercrossing of a — without crossing the basepoint $*$ — we see that the tangent vector to Λ makes a counterclockwise rotation from $3\pi/4$ to $\pi/4$. That is, the capping path γ associated to a has $\mathrm{rot}(\gamma) = 3/4$, and we see that

$$|a| = 2(3/4) - (1/2) = 1.$$

So $\mathcal{A}_\Lambda^k$ is the $\mathbb{Z}$-module generated by words in a, t, and t^{-1} in which a appears exactly k times. So $\mathcal{A}_\Lambda^0$ is generated (as a module) by words of the form t^m, $m \in \mathbb{Z}$, while $\mathcal{A}_\Lambda^1$ is generated (as a module) by words of the form $t^{m_0} a t^{m_1}$, with $m_0, m_1 \in \mathbb{Z}$. In general, $\mathcal{A}_\Lambda^k$ is generated (as a module) by words of the form

$$t^{m_0} a^{p_1} t^{m_1} a^{p_2} t^{m_2} \cdots a^{p_\ell} t^{m_\ell},$$

where $\ell \geq 1$, $m_0, \ldots, m_\ell \in \mathbb{Z}$, $p_1, \ldots, p_\ell \geq 0$, and $p_1 + \cdots + p_\ell = k$.

5.3 The differential

Finally, we define a differential $\partial : \mathcal{A}_\Lambda \to \mathcal{A}_\Lambda$ which decreases grading by 1. Notice that this conflicts with our definition of differentials from yesterday — when grading was decreased, we said we had a *boundary operator*. The literature isn't 100% consistent.

Before we can write down the differential, we need to determine some signs associated to each crossing.

5.3.1 Reeb signs and orientation signs

Let a be a double point of the Lagrangian projection $\Pi(\Lambda)$ in $\mathbb{R}^2_{xy}$. In a neighborhood of a, $\mathbb{R}^2_{xy}$ is divided into four quadrants, each of which has a corresponding **Reeb sign**. Each quadrant gives its boundary a counterclockwise orientation, and if this orientation takes us from an understrand of Λ to an overstrand, then the quadrant has a positive Reeb sign. If this orientation takes us from an overstrand to an understrand, the quadrant has a negative Reeb sign.

Next, each quadrant needs an **orientation sign**. If a is a negative crossing, all quadrants have positive orientation sign; if a is a positive crossing, then two quadrants have negative orientation sign:

See Figure 4 of Etnyre-Ng

5.3.2 The moduli space

The differential on $\mathcal{A}_\Lambda$ will be defined by counting the number of immersed discs satisfying certain conditions. We define those conditions now.

Definition. $(\Delta(a; b_1, \dots, b_m))$
For $m \geq 0$, let $\xi = e^{2\pi i/(n+1)} \in \mathbb{C}$ and define

$$D^2_m := D^2 - \{\xi^0, \xi^1, \xi^2, \dots, \xi^m\} \subset \mathbb{C}.$$

We refer to $\xi^0, \xi^1, \dots, \xi^m$ as the boundary punctures of D^2_m. Now let each of $a, b_1, \dots, b_m$ be a double point of $\Pi(\Lambda)$. Then we define $\mathcal{M}(a; b_1, \dots, b_m)$ to be the set of smooth maps

$$u \colon (D^2_m, \partial D^2_m) \to (\mathbb{R}^2_{xy}, \Pi(\Lambda))$$

satisfying

1. u is an *immersion*;
2. $u(\xi^0) = a$, and u sends a neighborhood of ξ^0 in D^2_m to a quadrant of a with positive Reeb sign;
3. for $1 \leq k \leq m$, $u(\xi^k) = b_k$, and u sends a neighborhood of ξ^k in D^2_m to a quadrant of b_k labeled with a negative Reeb sign.

Finally, we define an equivalence relation $\sim$ on $\mathcal{M}(a; b_1, \dots, b_m)$ by declaring $u \sim u'$ if there is a diffeomorphism $\varphi \colon (D^2_m, \partial D^2_m) \to (D^2_m, \partial D^2_m)$ such that $u' = u \circ \varphi$. Then we have

$$\Delta(a; b_1, \dots, b_m) := \mathcal{M}(a; b_1, \dots, b_m)/\sim.$$

Remark. We used the word *immersion* in the above definition, but there's a good chance you haven't seen this word before. Given any smooth map $F \colon \mathbb{R}^2 \to \mathbb{R}^2$, we can think about the matrix

$$\begin{pmatrix} \partial F_1/\partial x & \partial F_1/\partial y \\ \partial F_2/\partial x & \partial F_2/\partial y \end{pmatrix},$$

where F_1 and F_2 are the components of F. We'll say that F is an immersion if the determinant of this matrix is nowhere zero. Since u is a map between subsets of $\mathbb{R}^2$, we can apply this definition of immersion to u. (Really immersion means something more general, but this is good enough for now.) In practice, we're not going to worry about this condition too much.

Example 5.3. In our Lagrangian projection of the standard Legendrian unknot, we only get a nonempty set when $n = 0$. That is, $\Delta(a)$ contains two discs, and there are no discs with more punctures that satisfy the conditions above.

Fact. If $\Delta(a; b_1, \ldots, b_m)$ is nonempty, then we must have

$$|a| = 1 + \sum_{k=1}^{m} |b_k|.$$

5.3.3 The differential

We're almost ready to define the differential. The last preparation we need to make is to associate a word $w(u)$ and a sign $\epsilon(u)$ to each disc $u \in \Delta(a; b_1, \ldots, b_m)$.

To do this, notice that $u|_{\partial D^2_m} : \partial D^2_m \to \Pi(\Lambda)$ gives us a union of paths $\gamma_0, \gamma_1, \ldots, \gamma_m$, with γ_k starting at b_k for $1 \leq k \leq m$, and γ_0 starting at a. Let's denote by $\#\gamma_k$ the number of times the path γ_k crosses the basepoint $*$. This is a signed count according to the orientation on Λ. Then we can define the word $w(u)$:

$$w(u) := t^{\#\gamma_0} b_1 t^{\#\gamma_1} b_2 \cdots b_m t^{\#\gamma_m}.$$

We define the sign of u to be

$$\epsilon(u) := \epsilon_a(u)\epsilon_{b_1}(u) \cdots \epsilon_{b_m}(u),$$

where for each crossing c, $\epsilon_c(u)$ is the orientation sign of the quadrant of c covered by u.

At last, we can define the differential. We will define $\partial_\Lambda : \mathcal{A}_\Lambda \to \mathcal{A}_\Lambda$ by giving the definition for each generator $a_1, \ldots, a_m, t^{\pm}$ and then extending this definition to words using the Leibniz rule. By definition we will have

$$\partial_\Lambda(t) := 0 \quad \text{and} \quad \partial_\Lambda(t^{-1}) := 0.$$

For $a_1, \ldots, a_m$ we define

$$\partial_\Lambda(a_k) := \sum_{u \in \Delta(a; b_1, \ldots, b_m)} \epsilon(u) w(u),$$

where the sum allows $b_1, \ldots, b_m$ to be *any* collection of double points of $\Pi(\Lambda)$, $m \geq 0$.

Example 5.4. The standard Legendrian unknot has just one double point a, so there's only one differential to compute. Two discs contribute to this differential, and we'll find in lecture that $\partial_\Lambda(a) = 1 + t^{-1}$. (Or, depending on where we put the basepoint, $\partial_\Lambda(a) = 1 + t$.)

With a differential defined, we now have the Chekanov-Eliashberg DGA $(\mathcal{A}_\Lambda, \partial_\Lambda)$ of Λ defined. But it's really not obvious that this is a DGA. For instance, we still need to check that:

- the differential $\partial_\Lambda(a_k)$ is a finite sum, and thus well-defined;
- the differential decreases degree by 1;
- the differential squares to 0.

Moreover, the notation $(\mathcal{A}_\Lambda, \partial_\Lambda)$ suggests that this DGA depends only on Λ, and not on any other choices. We'll punt on all of these details for now, maybe checking them later.

5.4 The DGA for a Legendrian trefoil

Here's a front projection for a Legendrian trefoil, which we can resolve to a Lagrangian projection:

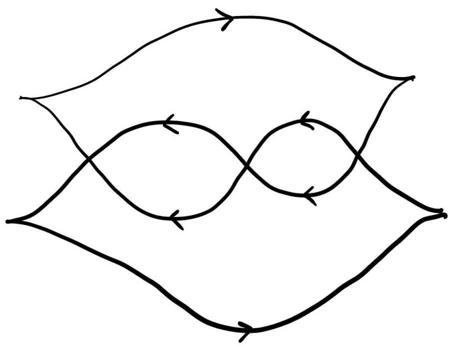

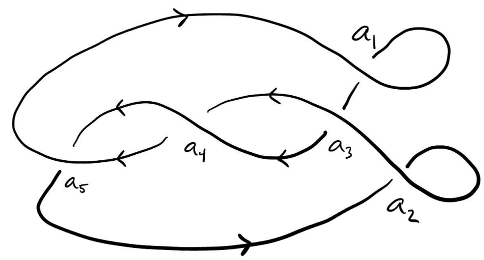

In lecture we'll find that $\partial_\Lambda(a_1) = 1 + a_3 + a_5 + a_5 a_4 a_3$. See Example 3.6 of Etnyre-Ng for details.