THE MODULI SPACE OF SPATIAL POLYGONS

ALEX WRIGHT

January 14, 2019

Abstract. These are rough notes to accompany a series of three expository lectures in the RTG Seminar on Geometry, Dynamics and Topology at the University of Michigan in fall 2018. Corrections appreciated.

Contents

1. Introduction

Let $r = (r_1, \ldots, r_n)$ be a tuple of positive numbers. Let $S_{r_i} \subset \mathbb{R}^3$ denote the sphere of radius r_i. Let

$$\mu : S_{r_1} \times \cdots \times S_{r_n} \to \mathbb{R}^3$$

denote the addition map, so $\mu(e_1, \ldots, e_n) = e_1 + \cdots + e_n$.

Lemma 1.1. *μ is a submersion away from the locus where all r_i are collinear.*

Proof. The image of the derivative is the sum of the tangent spaces of the spheres. These two-dimensional tangent spaces span $\mathbb{R}^3$ as long as they are not all equal. $\square$

Corollary 1.2. 0 *is a regular value of* μ *as long as there is no solution to* $\sum \varepsilon_i r_i = 0$ *with* $\varepsilon_i \in \{\pm 1\}$.

Proof. If $\sum v_i = 0$ and all the v_i are collinear, then the signed sum of their lengths must be 0. More formally, take the inner product with a unit vector that is collinear to all the v_i to get the result. There are two choices of unit vector, which correspond to the two possible sign conventions. ☐

Standing assumption: From now on, we will assume there is no solution to $\sum \varepsilon_i r_i = 0$ with $\varepsilon_i \in \{\pm 1\}$. Note that this assumption is satisfied by most choices of r.

Corollary 1.3. $\mu^{-1}(0)$ *is a submanifold of* $\prod S_{r_i}$ *of dimension* $2n-3$

Remark 1.4. Without our standing assumption, this would not be true in general; $\mu^{-1}(0)$ might have singularities. The singularities are mild and understandable, but we prefer to stick to the non-singular situation.

Lemma 1.5. $\mu^{-1}(0)$ *is invariant under* $SO(3)$, *and the quotient*

$$\mathcal{M}(r) = \mu^{-1}(0)/SO(3)$$

is a manifold of dimension $2n-6$.

Proof. Note that because of the genericity assumption, the $SO(3)$ action on $\mu^{-1}(0)$ is free.

That the quotient is a manifold follows from the fact that the $SO(3)$ actions is free and proper and smooth, see for example [**Lee13**, Theorem 7.10]. The dimension of the quotient is the dimension of $\mu^{-1}(0)$ minus the dimension of $SO(3)$. ☐

This manifold $\mathcal{M}(r)$ is called the moduli space of spatial polygons, or polygons in 3-space $\mathbb{R}^3$. The polygon should be through of as the edge vectors $(e_1, \ldots, e_n) \in \prod S_{r_i}$, placed in three space from tip to tail, up to orientation preserving rigid motions of $\mathbb{R}^3$. Note that the polygons may intersect themselves.

The standing assumption says that $\mathcal{M}(r)$ does not contain any polygon contained in a line.

General references on $\mathcal{M}(r)$ include the two papers [**KM96, Kly94**] that independently developed some of the general theory, and the more recent thesis [**Man**], whose Section 1.1 corresponds to the material above. We will give additional references as we proceed, but we have not made any attempt to be comprehensive in our bibliography. The author is not an expert on this subject, and hence may have inadvertently left out some important references and attributions.

2. The case of $n = 4$

If $n = 4$ then the manifold is of dimension 2, i.e. it is a smooth surface, and the standing assumption isn't required. (It is required for bigger n for the moduli space to be a smooth manifold.)

Any spacial polygon can be obtained by gluing together two triangles. The side lengths of one of the triangles are r_1, r_2 and ℓ, and the sides lengths of the other triangles are r_3, r_4 and ℓ. Here there is an interval of possibilities for ℓ. When it achieves its maximum and minimum values, one of the polygons will be degenerate, in that all three edges will be collinear.

The spacial polygon is obtained by gluing the two triangles together along some angle θ. Different θ give different polygons, except when one of the triangles is a line. We think of ℓ as a map $\mathcal{M}(r) \to \mathbb{R}$, and we see that

- the image is an interval,
- the preimage of an interior point is a circle, and
- with the standing assumption, the preimage of an endpoint is a point.

Without the standing assumption, the preimage of an endpoint isn't necessarily a point. (Think for example of the case when all r_i are equal. Then the preimage of $\ell = 0$ is an interval.) Nonetheless, it isn't too hard to conclude the following.

Proposition 2.1. *When $n = 4$, the moduli space $\mathcal{M}(r)$ is a sphere whenever it is non-empty and not a point.*

It is empty when $r_1 + r_2 + r_3 < r_4$, or when the same inequality is true after permuting the r_i. It is a point when equality holds.

Remark 2.2. For $n > 4$ the topology is more complicated, and depends on r.

3. $\mathfrak{so}(3)$

The purpose of this section and the next section is to show that μ is a moment mapping, and hence that some results concerning μ that can be understood very explicitly can also be viewed as instances of general phenomena.

The Lie algebra of $SO(3)$ is the space $\mathfrak{so}(3)$ of trace 0 anti-symmetric matrices. Note that $\mathfrak{so}(3)$ is a three dimensional algebra, where the algebra product is given by Lie bracket. Another three dimensional algebra is $\mathbb{R}^3$ endowed with cross product. In fact, the Jacobi identity

$$a \times (b \times c) + b \times (c \times a) + c \times (a \times b) = 0$$

for cross product shows that $(\mathbb{R}^3, \times)$ is a Lie algebra. Our goal is now to show this Lie algebra is isomorphic to $\mathfrak{so}(3)$.

For $a \in \mathbb{R}^3$, define $T_a \in \mathrm{End}(\mathbb{R}^3)$ by $T_a(b) = a \times b$.

Lemma 3.1. *$T_a \in \mathfrak{so}(3)$.*

Proof. Recall the formula for the triple scalar product:

$$\langle a \times b, c\rangle = \det(a, b, c).$$

The fact that T_a is anti-symmetric corresponds to the fact that determinant changes signs when you swap a pair of columns.

That T_a is trace zero is geometrically clear: T_a annihilates a, and acts as a rotation followed by a scaling on its orthogonal complement. □

Remark 3.2. T_a is the derivative of a on parameter group of rotations that fixes a. This can be used to give an alternative proof of the lemma.

Lemma 3.3. *$T_{a\times b} = [T_a, T_b]$.*

Proof. Using the Jacobi identity and the anti-commutativity of cross product, we calculate

$$\begin{aligned} T_{a\times b}(c) &= (a \times b) \times c \\ &= -c \times (a \times b) \\ &= +b \times (c \times a) + a \times (b \times c) \\ &= -b \times (a \times c) + a \times (b \times c) \\ &= -T_b(T_a(c) + T_a(T_b(c)), \end{aligned}$$

which is the desired result. □

Lemma 3.4. *If $B \in SO(3)$, then $T_{Ba} = BT_aB^{-1}$.*

Proof. We compute

$$T_{Ba}(c) = (Ba) \times c = B(a \times (B^{-1}c)) = (BT_aB^{-1})(c),$$

using the fact $(Bv) \times (Bw) = B(v \times w)$. □

The three lemmas show the following.

Proposition 3.5. *The map $a \mapsto T_a$ is an isomorphism between the Lie algebras $(\mathbb{R}^3, \times)$ and $\mathfrak{so}(3)$, conjugating the usual action of $SO(3)$ on $\mathbb{R}^3$ with the Adjoint action of $SO(3)$ on $\mathfrak{so}(3)$.*

4. Moment mappings for $SO(3)$

We now consider the sphere S_r of radius r in $\mathbb{R}^3$. We deviate from our typical notational conventions in this section and allow r to denote a single positive real number rather than a tuple, and similarly for other notation.

We endow this sphere with the symplectic form ω that is $1/r$ times the usual area form. That is, if $v, w \in T_e(S_r)$,

$$\omega(v, w) = \frac{\det(v, w, e)}{r^2}.$$

Note that e/r is the unit normal vector to the sphere. We will presently see why the extra factor of r is included; it is so the lemma below is true. The factor of r is more important that you might think, since we will be using different values of r simultaneously.

Suppose that a Lie group, say $SO(3)$, acts on a symplectic manifold, say S_r, via symplectomorphisms. This gives a collection of flows, via the one parameter subgroups $\exp(t\xi)$, for $\xi \in \mathfrak{so}(3)$. We say that

$$\mu : S_r \to \mathfrak{so}(3)^*$$

is a moment map if, for each $\xi \in \mathfrak{so}(3)$, we have that $\langle \mu, \xi \rangle$ is a Hamiltonian function for the flow $\exp(t\xi)$. Here we can in fact consider μ as a map to $\mathfrak{so}(3)$ rather than its dual, using the standard $SO(3)$ invariant inner product $\langle \cdot, \cdot \rangle$ to identity $\mathbb{R}^3 \simeq \mathfrak{so}(3)$ and $(\mathbb{R}^3)^* \simeq \mathfrak{so}(3)^*$.

Concretely, that $\langle \mu, \xi \rangle$ is a Hamiltonian function for the flow means that, for $v \in T_e S_r$,

$$(d\langle \mu, \xi \rangle)(v) = \omega \left(\frac{d}{dt}\bigg|_{t=0} \exp(t\xi)e, v \right).$$

In other words, the exterior derivative of the function $\langle \mu, \xi \rangle$ gives a 1-form, i.e. a co-vector at each point of the manifold. The symplectic form, being a non-degenerate bi-linear form on each tangent space, allows us to identity the tangent and co-tangent spaces, and transform this co-vector into a tangent vector at each point. In this way we obtain a vector field, which generates the flow.

Lemma 4.1. *The inclusion $S_r \subset \mathbb{R}^3 = \mathfrak{so}(3)$ is the moment mapping for the $SO(3)$ action on S_r.*

Proof. We compute, using μ to denote the identity map,

$$\begin{aligned}(d\langle \mu, \xi\rangle)(v) &= \frac{d}{dt}\langle \mu(e+tv), \xi\rangle\Big|_{t=0} \\ &= \frac{d}{dt}\langle e+tv, \xi\rangle\Big|_{t=0} \\ &= \langle v, \xi\rangle\end{aligned}$$

and also

$$\begin{aligned}\omega\left(\frac{d}{dt}\Big|_{t=0}\exp(t\xi)e, v\right) &= \omega\left(\xi\times e, v\right) \\ &= \frac{\det(\xi\times e, v, e)}{\|e\|^2}.\end{aligned}$$

We wish to show that this is the above expression for $(d\langle \mu, \xi\rangle)(v)$. By $SO(3)$ invariance, homogeneity, and the fact that v is orthogonal to e, it suffices to check this for $e = (0,0,1)$ and $v = (0,1,0)$. In this case, the determinant is the x component of $\xi \times e = \xi \times (0,0,1)$, which is the y component $\langle \xi, v\rangle$ of ξ as desired. □

Remark 4.2. We have observed that

$$\frac{\det(\xi\times e, v, e)}{\|e\|^2} = \langle v, \xi\rangle$$

whenever v is orthogonal to e.

In general, if a group G acts on two manifolds, it acts diagonally on their product. In this context, if each action has a moment map, then the sum of these maps is easily seen to be a moment map for the diagonal action. Hence we get the following.

Corollary 4.3. *The moment map for the $SO(3)$ action on a product $S_{r_1} \times \cdots \times S_{r_n}$ of spheres is given by the addition map μ.*

We close by commenting that we can view S_r as contained in $\mathfrak{so}(3)$ via the isomorphism to $\mathbb{R}^3$. In general, any orbit for the action of a Lie group on the (dual of) its Lie algebra has a *Kirillov-Kostant* symplectic form. (Here we can use an invariant inner product to identify $\mathfrak{so}(3)$ and its dual.) The symplectic pairing of two tangent vectors $[X, \lambda]$ and $[Y, \lambda]$ at the point λ in the adjoint orbit is

$$-\langle \lambda, [X, Y]\rangle.$$

Note that the tangent space to the orbit at λ is $\{[\lambda, X]\}$, where $[\cdot,\cdot]$ is the Lie bracket. Returning to our previous notation, we have

$$e = \lambda, v = [X, \lambda], w = [Y, \lambda].$$

Since cross product by e/r acts by rotation by $\pi/2$ on T_e, we can set $X = -e \times v/r^2$ and $Y = -e \times w/r^2$. This gives

$$\begin{aligned} \omega(v,w) &= \langle e, (-e \times v/r^2) \times (-e \times w/r^2) \langle \\ &= \det(e, (e/r) \times (v/r), (e/r) \times (w/r)) \\ &= \frac{1}{r^2} \det(e, v, w). \end{aligned}$$

So we see that our normalization of the symplectic form exactly recovers the Kirillov-Kostant form.

It is a general fact that the inclusion of a (co)adjoint orbit is a moment map.

5. The symplectic structure on $\mathcal{M}(r)$

General principles in symplectic topology guarantee that $\mathcal{M}(r)$ has a natural symplectic structure, because it was constructed using a symplectic space $(\prod S_{r_i}, \omega)$, and a moment map μ. Sometimes this is called symplectic reduction, or the Marsden-Weinstein Theorem. Natural means that, to compute the symplectic pairing of of two tangent vectors to $\mathcal{M}(r)$ we lift them to $\prod S_{r_i}$ and take the pairing there. There are many possible lifts, but it turns out that they all give the same pairing. This is for the following reason.

Lemma 5.1. *The tangent vectors in the kernel of the derivative of the map $\mu^{-1}(0) \to \mathcal{M}(r)$ have zero symplectic pairing with any tangent vector to $\mu^{-1}(0)$.*

There are easy proofs of this in most symplectic topology books. Here we also provide an algebraic proof in our situation.

Proof. We take $e = (e_1, \ldots, e_k) \in \mu^{-1}(0)$. The kernel of the derivative is all vectors of the form $(\xi \times e_1, \ldots, \xi \times e_n)$ for some $\xi \in \mathbb{R}^3$. Indeed, these are the derivatives for the action of one parameter subgroups of $SO(3)$, so generate the tangent space to the $SO(3)$ orbit.

Now consider a general vector $(v_1, \ldots v_n)$ in $T_e\mu^{-1}(0)$. So $\langle v_i, e_i \rangle = 0$ for each i and $\sum v_i = 0$. We need to show that

$$0 = \sum \frac{\det(\xi \times e_i, v_i, e_i)}{\|e_i\|^2}.$$

By Remark 4.2, the right hand side is equal to

$$\sum \langle \xi, v_i \rangle$$

which is indeed zero since $\sum v_i = 0$. □

6. Action-angle coordinates

The reference for this section is [**KM96**, Sections 3, 4].

Consider an oriented polygon in $\mathbb{R}^3$ given by $(e_1, \dots, e_n) \in \prod S_{r_i}$. We define the length of the i-th diagonal as

$$\ell_i = \|e_1 + \cdots + e_{i+1}\|.$$

Let $\mathcal{M}'(r) \subset \mathcal{M}(r)$ denote the locus where all ℓ_i are non-zero and each $e_0 + \cdots + e_i$ is not collinear to e_i. On this locus, we can define triangles T_i that span the initial vertex and the $(i+1)$-st edge. The triangles T_i, T_{i+1} share an edge which is exactly the i-the diagonal. We define θ_i to be the dihedral angle from T_i to T_{i+1}. So θ_i is valued in $\mathbb{R}/2\pi i\mathbb{Z}$. We observe the following.

Lemma 6.1. *The $\ell_i, \theta_i, i = 1, \dots, n-3$ are smooth coordinates for $\mathcal{M}'(r)$.*

These coordinates are analogous to Fenchel-Nielsen coordinates.

Remark 6.2. For each $\sigma \in S_n$, there is an isomorphism $\sigma : \mathcal{M}(r) \to \mathcal{M}(\sigma(r))$. Hence one obtains (a slight generalization of) action angle coordinates on $\sigma^{-1}\mathcal{M}(\sigma(r))$. These coordinates are very natural when σ is a power of the cyclic permutation $(123 \dots n)$. In this case it simply amounts to picking a different base vertex of the spatial polygon. Other permutations give rise to less geometric coordinate charts.

One can see that the union of the geometric $\sigma^{-1}\mathcal{M}(\sigma(r))$ do not always cover $\mathcal{M}(r)$. For example, consider the case when $n = 9$ and all r_i are equal. Then $\mathcal{M}(r)$ contains a polygon that goes three times around a triangle. More generally, when $n = 9$ and $r_1 + r_2 + r_3 = r_4 + r_5 + r_6 = r_7 + r_8 + r_9$ then $\mathcal{M}(r)$ contains polygons that are triangles where each edge is actually composed of three collinear edges. These polygons are not in any of the geometric coordinate charts arising from cyclic permutations. (Thanks to Gabe Khan for these examples.)

However, the union of all $\sigma^{-1}\mathcal{M}(\sigma(r))$ do always cover $\mathcal{M}(r)$. To see this we need to take a polygon, and produce a permutation. To do this, start with any edge. Then repeatedly pick the "next" edge by picking any of the remaining edges that are not collinear to the sum of the previously chosen edges. This algorithm might get stuck, but if it does, then all the remaining edges are collinear. If this is the case, remove the previous edge e_{i_0}, add all the remaining collinear edges, and then add e_{i_0} as the final edge. The order of the edges will produce the desired permutation.

For each i, there is a bending flow which twists the first part of the polygon around the i-th diagonal. In coordinates, it fixes all coordinates except θ_i, and increases that coordinates linearly.

Lifting to $\mu^{-1}(0)$, the flow rotates $e_1, \ldots, e_{i+1}$ about the direction of $e_1 + \cdots + e_{i+1}$. We will see that, up to re-parametrization, this is the Hamiltonian flow for the function ℓ_i. Since the flow fixes the coordinates $e_j, i > 1$, it suffices to prove the following result.

Lemma 6.3. *The Hamiltonian flow on* $\prod S_{r_i}$ *generated by the function* $H = \frac{1}{2}\|\sum e_i\|^2$ *is a rotation about the direction of* $\sum e_i$ *with speed* $K = \|\sum e_i\|$. *If one instead uses* K *as the Hamiltonian, the the flow is the rotation with speed one.*

Speed refers to how quickly angles change. So speed one means that the period is 2π, speed $1/2$ means that the period is π, etc.

Proof. We begin by observing that

$$dH(v) = \langle \sum v_i, \sum e_i \rangle.$$

Recall the definition of the symplectic form:

$$\omega(v, w) = \sum \frac{\det(v_i, w_i, e_i)}{r_i^2}.$$

We plug in $w = ((\sum e_j) \times e_i)$, which we wish to show is the Hamiltonian vector field, to get

$$\sum \frac{\det(v_i, (\sum e_j) \times e_i, e_i)}{r_i^2} = \sum \langle \sum e_j, v_i \rangle$$

as desired.

Now, note that $H = K^2/2$ so $dH = KdK$, so $dK = (dH)/K$ and the speed of the flow is changed to 1. $\square$

Corollary 6.4. *The Hamiltonian flow for the function* ℓ_i *preserves* $\mathcal{M}'(r)$ *and fixes all coordinates except* θ_i*, on which it acts by* $\theta_i \mapsto \theta_i + t$. *In other words,*

$$d\ell_i(\cdot) = \omega\left(\frac{\partial}{\partial \theta_i}, \cdot\right).$$

This Corollary was discovered in **[KM96]**, where they are called bending flows. (We may also call them twisting flows, because they are analogous to Fenchel-Nielsen twists.) The Hamiltonian flows were also discussed in **[Kly94]**, but without giving a geometric interpretation.

Remark 6.5. A symplectic manifold with a faithful torus action of half the dimension is called a toric manifold. $\mathcal{M}(r)$ is almost toric, but is not actually toric. This is because the i-th bending flow cannot be extended to the locus with $\ell_i = 0$. Intuitively, one could change the polygon so ℓ_i is infinitesimal, but with the i-th diagonal pointing in a different direction, and all these different directions correspond to different, incompatible bends. Formally, at these points $d\ell_i$ is zero, so the Hamiltonian flow is not well defined. Nonetheless, for $n < 7$ and r generic, $\mathcal{M}(r)$ is toric [**KM96**, Remark 3.13].

In [**KM96**] the above, plus a trick, was used to understand the symplectic structure on $\mathcal{M}'(r)$.

Theorem 6.6 (Kapovich-Millson)**.** *The symplectic form on* $\mathcal{M}'(r)$ *is*

$$\omega = \sum d\theta_i d\ell_i.$$

Proof. Write

$$\omega = \sum (f_{ij} d\theta_i d\ell_j + g_{ij} d\theta_i d\theta_j + h_{ij} d\ell_i d\ell_j).$$

The corollary, and the facts that $d\ell_i(\ell_j) = \delta_{ij}$ and $d\ell_i(\theta_j) = 0$ give that $f_{ij} = \delta_{ij}$ and $g_{ij} = 0$. So it remains only to show that the coefficients of $d\ell_i d\ell_j$ is zero.

For this, we use that any spacial polygon can be moved using the bending flows to be planar. Indeed, being planar just means that all $\theta_i = 0$. This doesn't change the lengths. Since ω is closed and these flows are Hamiltonian, this doesn't change the coefficients of ω. (This is basically the fact that Hamiltonian flows are symplectomorphisms.)

There is an isometry on $\mathcal{M}'(r)$ given by applying any orientation reversing isometry of $\mathbb{R}^3$ to the polygon. It doesn't matter which one, since we've already modded out by orientation preserving isometries. The planar polygons are fixed by involution. All the lengths are obviously also fixed, but the angles are negated.

The involution negates the symplectic form. For example, if you lift it to $\mu^{-1}(0)$, then it is just negation, and there are three sign changes (the two tangent vectors, and e), and $(-1)^3 = -1$.

Since the symplectic form is negated, and $d\ell_i$ is preserved that means that the coefficient of $d\ell_i d\ell_j$ is zero. This gives the result. □

Remark 6.7. If one only wishes to know the associated volume form, the proof is simpler, in that it does not require the symmetry argument at the end, because

$$\sum (d\theta_i d\ell_j + h_{ij} d\ell_i d\ell_j)^{n-3} = \prod d\theta_i d\ell_j.$$

Remark 6.8. This proof follows the same outline as one of the proofs of Wolpert's formula for the Weil-Petersson symplectic form.

Corollary 6.9. *The volume of $\mathcal{M}(r)$ is $(2\pi)^{n-3}$ times the usual Euclidean volume of the image the map $(\ell_1, \ldots, \ell_{n-3})$.*

In particular, it should be possible to use this to see that the volume is piecewise polynomial. There is a beautiful closed form formula for the volume of $\mathcal{M}(r)$ [**Tak01, Kho05, Man14**]. It would be great to find an elementary proof of this formula directly from the corollary. One would think that at least an elementary inductive proof would be possible. (I haven't tried to write one down. It would be preferable to find a proof that encapsulated more geometric meaning while still being elementary.)

It's nice to make the convention that $\ell_0 = r_1$ and $\ell_{n-2} = r_n$. (The convention is that ℓ_i is a side length of a triangle whose other side lengths are ℓ_{i-1} and r_{i+1}.) Then the image of the map in the corollary is the polytope described by the three inequalities indicating that there is a triangle of side lengths $(\ell_{i-1}, \ell_i, r_i)$. It is a general feature of moment maps that the image is a polytope.

In particular, this makes it easy to see that for $r = (r_1, r_2, r_3, r_4)$ with $\sum r_i$ constant, the greatest volume is realized for $r_1 = r_2 = r_3 = r_4$. In other words, the regular quadrilateral is the "most flexible" quadrilateral, in that the volume of its deformation space is greatest. In fact, this was generalized in [**Kho05**] to arbitrary n. We will give a short, elementary argument now.

Theorem 6.10. *The maximum volume of $\mathcal{M}(r)$ among r with $\sum r_i$ a constant is uniquely achieved when all the lengths are equal.*

Proof. We first give the proof without worrying about uniqueness.

Suppose that not all r_i are equal. Then we claim that, without loss of generality we can assume that $r_1 > r_2$.

We compare the polytope of allowable ℓ_i values to the polytope obtained by replacing both r_1, r_2 with their average $a = (r_1 + r_2)/2$.

Before making this change, the allowable ℓ_1 values range from $r_1 + r_2 = 2a$ to $r_1 - r_2 > 0$. After the change, the allowable ℓ_1 values range from $2a$ to 0.

This shows that the max is achieved when all the lengths are equal. But it doesn't show uniqueness. The problem is that there are other constraints on ℓ_1. The upper bound that $\ell_1 \leq r_3 + \ldots + r_n$ is of no concern to us, because we see the improvement of the range of ℓ_1 on the lower bound. But it may be that $r_1 - r_2$ is the minimum value of the length of a sum of vectors of lengths $r_3, \ldots, r_n$, and then we would see no improvement. In all other cases we see improvement.

One can see uniqueness in the following way, suggested in outline by Benjamin Krakoff. One can iteratively average the largest and smallest

r_i's. If the process ends after finitely many iterations with all r_i's being equal, then the last averaging is easily seen to increase the volume; the remaining edges all have equal length, so they can be arranged to sum to zero. Otherwise, eventually the r_i are all approximately the same. In this case, the remaining edges for the next averaging can also arranged to sum to zero, so the next averaging increases the volume. □

7. The Kähler structure on $\mathcal{M}(r)$

A nice online general reference for Kähler quotients is [**Zhu**]. The other references for this section are [**KM96, Kly94**] as usual.

Recall that
$$T_e(SO(3)e) \subset T_e\mu^{-1}(0)$$
has zero symplectic pairing with all of $T_e\mu^{-1}(0)$. Hence $T_e\mu^{-1}(0)$ is the symplectic perp of $T_p(SO(3)e)$. Hence it is the Riemannian perp of $IT_e(SO(3)e)$, where I is the complex structure. Hence, the subspace of $T_e\mu^{-1}(0)$ that is orthogonal to the $SO(3)$ orbit is
$$(IT_e(SO(3)e))^\perp \cap (T_e(SO(3)e))^\perp.$$
This subspace is isomorphic to $T\mathcal{M}(r)$. Since it is I invariant, we get that $T\mathcal{M}(r)$ has an almost complex structure. This complex structure together with ω gives $\mathcal{M}(r)$ a Kähler structure.

One can write I concretely as follows, for $(v_1, \ldots, v_n) \in T_e \prod S_{r_i}$:
$$I(v_1, \ldots, v_n) = \left(v_1 \times \frac{e_1}{r_1}, \ldots, v_n \times \frac{e_n}{r_n} \right).$$

The subspace of $T_e\mu^{-1}(0)$ that is identified with $T\mathcal{M}(r)$ consists of those $(v_1, \ldots, v_n) \in \mu^{-1}(0)$ that are orthogonal to the $SO(3)$ orbit, i.e., are orthogonal to all vectors of the form $(\xi \times e_1, \ldots, \xi \times e_n)$ for any $\xi \in \mathbb{R}^3$. This is equivalent to
$$\sum \frac{v_i \times e_i}{r_i} = 0.$$
Note that the r_i^{-1} comes from the definition of the Kähler inner product, which like the symplectic form, is scaled from what is typical in $\mathbb{R}^3$. This equivalence is easily seen using that scalar triple product is determinant, to cyclically permute the role of the vectors.

In the previous section we discussed an involution on $\mathcal{M}(r)$ that we will now denote R. This R applies an orientation reversing isometry of $\mathbb{R}^3$ to the polygon. It is an isometry, so hence its fixed point set is totally geodesic.

When $n = 5$, the fixed point set has real dimension 2, and it is known that it can be a surface of positive, zero, or negative Euler characteristic, see for example [**Gor**, Table 1]. (Note that the fixed point set of R is not the same thing as the $SO(2)$ symplectic quotient of circles; rather it is the $O(2)$ symplectic quotient.) It follows that, in general, $\mathcal{M}(r)$ does not always have non-negative sectional curvature.

I don't know the answers to any of the following. (Other people besides me may know the answer.)

- Does $\mathcal{M}(r)$ have positive scalar curvature? Are there nice formulas for its curvatures?
- Can the isometry group of $\mathcal{M}(r)$ ever be larger than the stabilizer of r in S_n? (Possibly at least when $\dim \mathcal{M}(r) > 2$, in analogy with the fact that Royden's Theorem fails for two dimensional Teichmüller spaces?)
- Can geodesics be computed? Is there any sense in which they represent the most efficient way to deform one polygon into another, in some physical sense?
- Does the geodesic flow have dense orbits?

8. Moduli of points on $\mathbb{P}^1$

Consider $\mathbb{H}^3$ in the unit ball model, so its boundary is the unit sphere $S^2 \subset \mathbb{R}^3$.

Douady and Earle defined, for each probability measure ν on S^2 that has no atoms of mass at least $\frac{1}{2}$, a conformal center of mass $C(\nu) \in \mathbb{H}^3$ [**DE86**]. It has the property that for any $g \in PSL(2, \mathbb{C})$ then

$$C(g_* \nu) = gC(\nu).$$

Since $PSL(2, \mathbb{C})$ acts on $\mathbb{H}^3$ transitively, for any such ν there is a $g_\nu \in PSL(2, \mathbb{C})$ such that $g_\nu C(\nu) = 0$. This g_ν is well defined up to $SO(3) \subset PSL(2, \mathbb{C})$, where $SO(3)$ is the stabilizer of $0 \in \mathbb{H}^3$.

Douady and Earle also showed that $C(\nu) = 0$ if and only if the usual Euclidean center of mass of ν is 0. Hence, a consequence of their work is that there is a $g_\nu \in PSL(2, \mathbb{C})$ that is well defined up to $SO(3)$ such that $(g_\nu)_* \mu$ has Euclidean center of mass zero.

Remark 8.1. In Teichmüller theory, the Douady-Earle result gives a nice way to extend a quasi-symmetric map of the circle to a quasi-conformal map of the disc. Each point in the disc $\mathbb{H}^2$ gives a visual measure on the circle S^1 at infinity. This measure is pushed forward via the quasi-symmetric map, and the point is mapped to the conformal center of mass.

Fix $r = (r_1, \dots, r_n)$ with our usual genericity assumption, which means the r_i can't be divided into two subsets of equal total sum.

We now defined a space $\overline{\mathcal{M}}_{0,r}$ which is a compactification of $\mathcal{M}_{0,n}$, the moduli space of n distinct ordered points on $\mathbb{P}^1$ up to $PSL(2,\mathbb{C})$. Of course, $\mathcal{M}_{0,n}$ can be constructed as the subset of n distinct points in $\mathbb{P}^1$, quotiented by the diagonal $PSL(2,\mathbb{C})$ action. Similarly, we can consider the subset of n points in $\mathbb{P}^1$ such that the sum of the masses at any given point is strictly less than $\frac{1}{2}$. Here we give the i-th point mass r_i. The quotient of this space by $PSL(2,\mathbb{C})$ is defined to be $\overline{\mathcal{M}}_{0,r}$.

Remark 8.2. In $\overline{\mathcal{M}}_{0,r}$, there must always be at least three distinct points, since the total weight at any point is less than $\frac{1}{2}$. There should always be a map from the Deligne-Mumford compactification to $\overline{\mathcal{M}}_{0,r}$. This map can be defined as follows. Start on a leaf of the nodal Riemann surface. If the sum of the marked points on that leaf is less than $\frac{1}{2}$, contract it to a point, and give that point the sum of the weights. (There will always be a leaf with this property, unless we're done.) Repeat to eventually get a point in $\overline{\mathcal{M}}_{0,r}$.

It is intuitive that $\overline{\mathcal{M}}_{0,r}$ is compact. For example, if some points collide in such a way that the total mass would be more than $\frac{1}{2}$, then we can use a Möbius transformation to zoom in on the collision. This changes the point of view, so that now the remaining points, which have mass less than $\frac{1}{2}$ are colliding. But compactness will also follow from the following.

Theorem 8.3. *$\overline{\mathcal{M}}_{0,r}$ is isomorphic to $\mathcal{M}(r)$.*

Proof. We can consider any point of $\overline{\mathcal{M}}_{0,r}$ as a measure on S^2, up to $PSL(2,\mathbb{C})$. Apply a Möbius transformation to this measure so that the (conformal) center of mass is zero. The measure is now a collection of points $(q_1,\dots,q_n)\in S^2$, defined up to $SO(3)$, such that $\sum r_i q_i = 0$. This corresponds to the polygon with edge vectors $e_i = r_i q_i$. □

One can also see that the complex structures agree.

Proposition 8.4. *$\mathcal{M}(r)$ is simply connected.*

Proof. $\pi_1(\mathcal{M}_{0,n}) \to \pi_1(\mathcal{M}(r))$ is surjective; since the boundary is real codimension at least 2, any loop can be isotoped to be disjoint from the boundary.

$\pi_1(\mathcal{M}_{0,n})$ is generated by loops that move one point around another. If the sum of the weights is less than one half, the points can be collided, showing the that loop is trivial. If the sum is more than a half, all remaining points can be collided, showing that the loop is trivial. (Note that $\mathcal{M}_{0,3}$ is a point, so you in fact can't have such a loop when there are only three points.) □

It would be nice to come up with a direct proof not using the connection to moduli of points. I imagine this is possible, but haven't done it. A proof might attempt to show that $\pi_1(\mathcal{M}(r))$ is generated by the loops given by bending deformations, which are all obviously trivial.

Example 8.5. We consider the case $n = 5$, and show that the topology of $\mathcal{M}(r)$ depends on r.

In general, for any tuple of four r_i such that any three of them add up to at least 1/2, we get a map $\mathcal{M}(r) \to \overline{\mathcal{M}}_{0,4}$.

In particular, we get a map $\mathcal{M}(.2, .2, .2, .2, .2) \to \mathbb{P}^1 \times \mathbb{P}^1$, by consider the first 4 points and the last 4 points. It seems that there are three contracted curves, for the three pairs of two of the central three points that can overlap. The three contracted curves are disjoint. So $\mathcal{M}(.2, .2, .2, .2, .2)$ is $\mathbb{P}^1 \times \mathbb{P}^1$ blown up at 3 points. (Compare to [**Kly94**, Example 1.4.4], which claims that it is $\mathbb{P}^2$ blown up at 4 points.)

On the other hand, we also get a map $\mathcal{M}(.03, .31, .31, .31, .04) \to \mathbb{P}^1 \times \mathbb{P}^1$, again by consider the first four and the last four points. In this case this map is an isomorphism, since no pair of the middle three points can collide.

Example 8.6. We consider the example

$$\mathcal{M}\left(\frac{1}{2} - \varepsilon, \frac{\frac{1}{2} + \varepsilon}{n-1}, \dots, \frac{\frac{1}{2} + \varepsilon}{n-1}\right).$$

The first point cannot collide with any other. We can move that point to ∞ on $\mathbb{P}^1$, and we get that we have $n-1$ points on $\mathbb{C}$ up to affine transformation, any $n-2$ of which can collide. We can normalize the second point so it is at the origin, and then the remaining points are defined up to scaling, and we have the condition that they can't all be 0. Hence the moduli space is $\mathbb{P}^{n-3}$.

Example 8.7. We consider the example $\mathcal{M}(1, 1, 1, \varepsilon, \dots, \varepsilon)$. We have a bunch of maps to $\mathbb{P}^1$ by taking cross ratios of four tuples including the first three points. Since the first three points can't collide, this gives an isomorphism to $(\mathbb{P}^1)^{n-3}$.

These moduli spaces were used by Deligne-Mostow to build lattices in $SU(k, 1)$ for some values of k. In short, they looked at monodromy of families of cyclic covers of $\mathbb{P}^1$ branched over k points, and had to allow some collisions of points. Thurston gave a different point of view, related to the moduli of points point of view. He used the fact that, giving a collection of points on $\mathbb{P}^1$ and collection of r_i so that $\sum r_i = 1$, then one can find a cone metric on the sphere with cone angles $2\pi(2-r_i)$ at the given points. So $\mathcal{M}(r)$ is also a compactification of the space of cone metrics on S^2 with given cone angles!

9. Chamber structure

$\mathcal{M}(r)$ is non-empty if and only if all permutations of the inequality

$$r_1 \leq r_2 + \ldots + r_n$$

hold. Thus, the space of valid r is a polyhedron in $\mathbb{R}^n$. Since there are n-inequalities, the polyhedron is actually a cone on a simplex. For example, if we choose to normalize so $\sum r_i$ is a constant, then the region of r_i is a simplex.

This simplex is cut by a number of hyperplanes corresponding to the r not satisfying our standing assumption. These hyperplanes cut the region of valid r in to sub-regions called chambers.

In each chamber, the topology of $\mathcal{M}(r)$ is constant. One nice way to see this is that the corresponding moduli of points are isomorphic. One can just change the weights from one r to another r' in the same chamber, and this will never result in a point with total weight $\frac{1}{2}$ or more. Indeed, one can linearly or smoothly interpolate between r and r'. If the weights at some point reached $\frac{1}{2}$ at a point, this would prove that this intermediate value of r' did not lie in the same chamber.

10. Circle bundles, Euler classes and intersection numbers

For each $1 \leq i \leq n$, there is a circle bundle $\mathcal{C}_i \to \mathcal{M}(r)$. A point in $\mathcal{C}_i$ is an element of $\mu^{-1}(0)$ whose i-th edge begins at the origin and goes in the direction of the x-axis. The group $SO(2)$ of rotations that fix the x-axis acts on these configurations, and the quotient is $\mathcal{M}(r)$.

Remark 10.1. Recording the length of a diagonal gives a map $\ell_1 : \mathcal{M}(r) \to \mathbb{R}$, whose image is an interval. The preimage of a point in the interior of isomorphic to a circle bundle $\mathcal{C}_i$ over some $\mathcal{M}(r')$ of smaller dimension.

Let $\mathcal{L}_i$ denote the normal bundle of the i-the edge. To be pedantic, it is easiest to start by defining the normal bundle over $\mu^{-1}(0)$. Here it is given by the subspace of $\mathbb{R}^3$ that is normal to the i-th vector. Then we mod out by the $SO(3)$ action to get the total space of $\mathcal{L}_i \to \mathcal{M}(r)$.

Let $\mathcal{U}_i \subset \mathcal{L}_i$ denote the unit normal bundle. Note that $\mathcal{U}_i$ is isomorphic to $\mathcal{C}_i$. Indeed, a map is given by taking a polygon and a unit normal vector and placing the polygon in $\mathbb{R}^3$ so the i-th edge starts at the origin and goes in the x-direction, and the unit vector maps to the y direction. The inverse of this map takes a point in $\mathcal{C}_i$ and choose the unit normal vector in the y direction.

Hence $\mathcal{L}_i$ is the $\mathbb{R}^2$ bundle associated to the circle bundle $\mathcal{C}_i$. (This also suggests that $\mathcal{C}_i$ may be viewed as the relative tangent bundle at the i-th point, from the moduli of points point of view.)

There is a cohomology class

$$\mathrm{eu}(\mathcal{C}_1) \in H^2(\mathcal{M}(r), \mathbb{Z})$$

called the Euler class of the circle bundle. If the Euler class is non-zero, then the circle bundle doesn't have a section. We will use the following to compute the Euler class. For us it can serve as a definition; we'll skip more standard definitions of the Euler class.

Lemma 10.2. *Suppose that $\mathcal{C}$ is a circle bundle over a space $\mathcal{M}$. Let $\mathcal{L}$ be the associated $\mathbb{R}^2$ bundle. Let s be a section of $\mathcal{L} \to \mathcal{M}$ that intersects the zero section transversely. Then* $\mathrm{eu}(\mathcal{C})$ *is Poincare dual to the zero locus of s.*

Note that this lemma in theory applies to any circle bundle, since the zero section of $\mathcal{L}$ can always be perturbed to be transverse.

In our situation, we can easily define a section s of $\mathcal{L}_1$ as follows. Just take the second vector and project it to the normal direction. The zero section is the locus where the first and second edges are parallel, which has one component isomorphic to

$$D^+_{1,2} = \mathcal{M}(r_1 + r_2, r_3, \ldots, r_n)$$

and one component isomorphic to

$$D^-_{1,2} = \mathcal{M}(|r_1 - r_2|, r_3, \ldots, r_n).$$

The orientations are given by [**AG09**, Section 4] and could presumably be worked out from first principles fairly easily. We get

$$\mathrm{eu}(\mathcal{C}_1) = D^+_{1,2} + \mathrm{sign}(r_1 - r_2) D^-_{1,2}.$$

We might have made a sign error, if our conventions differ from other sources, and we are assuming $r_1 \neq r_2$ to get a transverse intersection. Note that the $D^\pm$ have natural orientations, since they are embedded copies of smaller polygon spaces, and these polygons spaces cary a natural complex structure. But D^- is not a complex submanifold of $\mathcal{M}(r)$. That is, the embedding is not holomorphic. One can already see a hint at this in the fact that the negation map is not holomorphic on the sphere, and hence the graph of this map is not a complex submanifold of a product of spheres. Note also that sometimes one or even both of these two loci could be empty. This happens sometimes even in the case of $n = 4$.

We can obtain integers by taking a product of $n - 3$ of the $\mathrm{eu}(\mathcal{L}_i)$ (for different i's). These integers count the signed number of transverse intersections of the Poincare dual classes. Note that even if one wants, for example, to consider the product of $\mathrm{eu}(\mathcal{L}_1)$ with itself in the case of $n-3$, this can often be done by considering two different sections, which might give transverse Poincare duals.

The information of these intersection numbers can be packaged into a generating function. They satisfy recursive relations analogous to those provided by Witten's conjectures for the case of moduli of curves, and they can also be computed explicitly, see for example [**AG09**]. There are a number of other papers that have been written about the intersection theory and topology of $\mathcal{M}(r)$. For example, the $\mathrm{eu}(\mathcal{L}_i)$ generate the cohomology ring, and an explicit presentation of this ring is given in [**HK98, Man14**], which also give results on intersection numbers.

References

[AG09] José Agapito and Leonor Godinho, *Intersection numbers of polygon spaces*, Trans. Amer. Math. Soc. **361** (2009), no. 9, 4969–4997.

[DE86] Adrien Douady and Clifford J. Earle, *Conformally natural extension of homeomorphisms of the circle*, Acta Math. **157** (1986), no. 1-2, 23–48.

[Gor] Irina Gorodetskaya, *Moduli spaces of planar pentagonal linkages: Combinatorial description*, arXiv:1305.6756.

[HK98] J.-C. Hausmann and A. Knutson, *The cohomology ring of polygon spaces*, Ann. Inst. Fourier (Grenoble) **48** (1998), no. 1, 281–321.

[Kho05] Vu The Khoi, *On the symplectic volume of the moduli space of spherical and Euclidean polygons*, Kodai Math. J. **28** (2005), no. 1, 199–208.

[Kly94] Alexander A. Klyachko, *Spatial polygons and stable configurations of points in the projective line*, Algebraic geometry and its applications, Aspects Math., E25, Friedr. Vieweg, Braunschweig, 1994, pp. 67–84. MR 1282021

[KM96] Michael Kapovich and John J. Millson, *The symplectic geometry of polygons in Euclidean space*, J. Differential Geom. **44** (1996), no. 3, 479–513.

[Lee13] John M. Lee, *Introduction to smooth manifolds*, second ed., Graduate Texts in Mathematics, vol. 218, Springer, New York, 2013.

[Man] Alessia Mandini, *The geometry of the moduli space of polygons in the euclidean space.*

[Man14] ______, *The Duistermaat-Heckman formula and the cohomology of moduli spaces of polygons*, J. Symplectic Geom. **12** (2014), no. 1, 171–213.

[Tak01] Tatsuru Takakura, *Intersection theory on symplectic quotients of products of spheres*, Internat. J. Math. **12** (2001), no. 1, 97–111.

[Zhu] Chenchang Zhu, *Marsden-Weinstein Reductions for Kähler, Hyperkähler and Quaternionic Kähler Manifolds.*