Camels and symplectic rigidity in vector spaces

Luke Volk

September 17, 2019

In these notes we explore the concept of symplectic rigidity. To do this with minimal prerequisites, we have decided to limit ourselves to the much easier linear background of symplectic *vector spaces* (rather than manifolds). As a result, all that is required is a first course in linear algebra (including vector spaces) and a curious mind. All vector spaces are assumed to be finite-dimensional.

Contents

1 Symplectic vector spaces

1.1 Area and the symplectic forms

Recall that the magnitude of the cross product of two vectors $u, v \in \mathbb{R}^3$ gives the area of the parallelogram with sides spanned by u and v. Symbolically:

$$\text{area}(u, v) = \|u \times v\|.$$

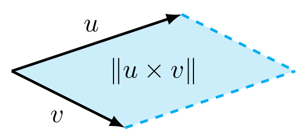

In particular, we can restrict ourselves to $u, v \in \mathbb{R}^2$ by setting the z-coordinate to be 0. Those familiar with determinants will then note that the *signed* area of the parallelogram spanned by $u, v \in \mathbb{R}^2$ is given by the determinant:

$$\operatorname{area}(u, v) = \det \begin{bmatrix} | & | \\ u & v \\ | & | \end{bmatrix},$$

where the vectors u and v are written as columns.

Exercise 1.1. Show that these formulas for area agree up to a change in sign.

From now on, our areas will always be signed, unlike in the first formula.

Still considering $u, v \in \mathbb{R}^2$, note that we may write:

$$\operatorname{area}(u, v) = u^\top \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} v,$$

where $u^\top$ denotes the transposition of u from a column to row vector. The matrix

$$\begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

is an example of a **bilinear** form on $\mathbb{R}^2$, meaning that it is a function $\mathbb{R}^2 \times \mathbb{R}^2 \to \mathbb{R}$ which is linear in both components. From either direct calculation or the properties of the cross product or determinant, one can deduce (and the reader is encouraged to verify) the following properties:

(i) area is a bilinear form.

(ii) area is skew-symmetric:

$$\operatorname{area}(u, v) = -\operatorname{area}(v, u).$$

(iii) area is non-degenerate:

$$\operatorname{area}(u, v) = 0 \text{ for all } v \in \mathbb{R}^2 \implies u = 0.$$

More generally, a non-degenerate bilinear form $B : V \times V \to \mathbb{R}$ on a vector space V which is skew-symmetric is called a **symplectic form** on V. A vector space endowed with a symplectic form is called a **symplectic vector space**.

Example 1.2. The bilinear form area is a bilinear form for $V = \mathbb{R}^2$.

This example can be modestly generalised to arbitrary even dimension by doing the following exercise:

Exercise 1.3. Show that $n \geqq 1$ the matrix:

$$\omega_0 := \begin{bmatrix} 0 & I_n \\ -I_n & 0 \end{bmatrix},$$

is a symplectic form on $\mathbb{R}^{2n}$, using the transformation rule:

$$(u, v) \mapsto \langle u, \omega_0 v \rangle = u^\top \omega_0 v.$$

The matrix ω_0 is called the **standard symplectic form** on $\mathbb{R}^{2n}$, and in a way, it gives a higher dimensional analogue to area (but not *volume*, like the determinant gives).

Given an arbitrary symplectic form ω on a vector space V, there is a standard form[1] in which we can express it as a matrix:

Proposition 1.4. *Let V be a finite dimensional vector space endowed with a non-degenerate, skew-symmetric bilinear form ω. Then there exists a basis $\{e_i, f_i \mid 1 \leqq i \leqq n\}$ of V such that:*

$$\omega(e_i, e_j) = \omega(f_i, f_j) = 0, \qquad \omega(e_i, f_j) = \delta_{i,j}.$$

Proof. Pick a non-zero vector in V and call it e_1. Note that ω is non-degenerate, and so we may find an $f_1 \in V$ such that $\omega(e_1, f_1) = 1$. Define $V_1 := \operatorname{span}\{e_1, f_1\}$ and define a respective set:

$$V_1^\omega := \{v \in V \mid \forall w \in V_1 \,.\, \omega(w, v) = 0\}.$$

Exercise 1.5. Show that $V = V_1 \oplus V_1^\omega$. That is, show both of the following:

$$V = V_1 + V_1^\omega, \quad \text{and} \quad V_1 \cap V_1^\omega = \{0\}.$$

This set acts as an "orthogonal" complement with respect to the bilinear form ω.

Repeat this process, picking a non-zero $e_2 \in V_1^\omega$, and an f_2 such that $\omega(e_2, f_2) = 1$, where we then define V_2 and so forth.

This process reduces the dimension by two each iteration and eventually halts as V is finite dimensional. Moreover, V is even-dimensional because

[1] For an elaboration on this process, ask me for my notes on the Gram-Schmidt algorithm.

otherwise we would conclude this process with a 1-dimensional subspace spanned by a vector u such that $\omega(u, \cdot) \equiv 0$. ■

Corollary 1.6. *If V admits a symplectic form, then V is necessarily even dimensional.*

We will call this basis a **symplectic basis** for ω. In particular, we see that in this basis the symplectic form ω is represented by the matrix:

$$\omega = \begin{bmatrix} 0 & I_n \\ -I_n & 0 \end{bmatrix}.$$

That is to say, we may choose a basis so that ω is represented by exactly the standard symplectic form ω_0 on $\mathbb{R}^{2n}$. Note that regardless of the basis, if ω is a bilinear form represented by the matrix Ω, then we may write:

$$\omega(u, v) = \langle u, \Omega v \rangle.$$

Thus as a result of the previous proposition, all symplectic forms "act" like Example 1.3. With regards to this notation, we will abuse notation sometimes and treat ω as equivalent to its matrix when the basis is clear, and so the identity would be stated as $\omega(u, v) = \langle u, \omega v \rangle$.

Note that if ω were degenerate and admitted a non-zero $u \in V$ such that $\omega(u, \cdot) = 0$, then we would have additional rows (and columns) of zeros in the matrix representation of ω. Indeed, if

$$U = \{u \in V \mid \omega(u, \cdot) = 0\},$$

then the dimension of this subspace U gives exactly how many rows/columns of zeros are added, the matrix resembling:

$$\begin{bmatrix} 0 & I_n & 0 & \cdots & 0 \\ -I_n & 0 & & & 0 \\ 0 & & & & \vdots \\ \vdots & & & & 0 \\ 0 & 0 & \cdots & 0 & 0 \end{bmatrix}$$

Exercise 1.7. Show that ω is non-degenerate if and only if the assignment

$$v \in V \mapsto \omega(v, \cdot),$$

is an isomorphism from V to the vector space of *linear* forms, V^*.

1.2 The orthogonal group

Consider a vector space V with an inner product $\langle \cdot, \cdot \rangle$. When we impose additional structure (such as inner products) on vector spaces, we tend to pay special attention to functions between the spaces which *preserve* this additional structure. In this way, the concept of orthogonal maps arises from considering inner product spaces. An **orthogonal map** on V is a linear map $A\colon V \to V$ such that:

$$\langle A(u), A(v) \rangle = \langle u, v \rangle, \quad \text{for all } u, v \in V.$$

Denote the set of orthogonal maps on V by $O(V)$, and in particular we denote $O(n) := O(\mathbb{R}^n)$. Orthogonal maps preserve the notions of angle and length in inner product spaces.

> **Exercise 1.8.** Show that each $A \in O(V)$ is an isomorphism—it is both injective and surjective.

From the above exercise, we know that each element of $O(V)$ is invertible (and the inverses are orthogonal themselves). Moreover, the identity map is in $O(V)$, and thus $O(V)$ forms a group.

> **Exercise 1.9.** Show that one may use the Gram-Schmidt process in order to find a basis of V that induces an isomorphism $O(V) \cong O(n)$, where $\dim V = n$.

Thus we have that only the dimension of V matters (up to ismorphism), so we can restrict our attention to $O(n)$. We call $O(n)$ the **orthogonal group** in dimension n.

In $\mathbb{R}^n$, the standard inner product of vectors $u, v \in \mathbb{R}^n$ is given by:

$$\langle u, v \rangle = u^\top v.$$

If $A \in O(n)$ is viewed as an $n \times n$ matrix in $M(n, \mathbb{R})$, then we know by definition:

$$\begin{aligned} \langle Au, Av \rangle = \langle u, v \rangle &\iff (Au)^\top (Av) = u^\top v \\ &\iff u^\top A^\top A v = u^\top v. \end{aligned}$$

> **Exercise 1.10.** Using $u^\top A^\top A v = u^\top v$ for all $u, v \in \mathbb{R}^n$, deduce that:
>
> $$A^\top A = I_n,$$
>
> the $n \times n$ identity matrix.

Note that for square matrices, a left-inverse is also a right inverse, and so this gives an alternative definition for $O(n)$:

$$O(n) = \{A \in M(n, \mathbb{R}) \mid A^\top A = AA^\top = I_n\}.$$

1.3 The symplectic group

Now instead of considering maps in inner product spaces, we will consider maps in symplectic vector spaces. Let V be a vector space with a symplectic form ω. We call a linear map $\Psi\colon V \to V$ a **symplectomorphism**[2] such that:

$$\omega(\Psi(u), \Psi(v)) = \omega(u, v), \quad \text{for all } u, v \in V.$$

Exercise 1.11. Show that a symplectomorphism is necessarily an isomorphism.

Similarly to the orthogonal group, we can justify the definition of the **symplectic group** in dimension $2n$, $\mathrm{Sp}(2n)$, namely that we only have one such space up to symplectomorphism.

Exercise 1.12. Taking hints from the calculation for the orthogonal group, show that:

$$\mathrm{Sp}(2n) = \{\Psi \in M(2n, \mathbb{R}) \mid \Psi^T \omega_0 \Psi = \omega_0\}.$$

1.4 The affine symplectic group

Recall that a **linear map** between (real) vector spaces V and W is a map $A\colon V \to W$ satisfying for each $v, v' \in V$ and $\lambda \in \mathbb{R}$:

(i) $A(v + v') = A(v) + A(v')$,

(ii) $A(\lambda v) = \lambda A(v)$.

Linear maps act as the subspace-preserving maps in the study of linear algebra. An important corollary of the above definition is that f fixes the origin:

$$A(0) = 0.$$

[2] Technically we can define a map between symplectic vector spaces (V, ω_V) and (W, ω_W) which preserves the symplectic forms:

$$\omega_W(\Psi(u), \Psi(v)) = \omega_V(u, v), \quad \text{for all } u, v \in V.$$

However, finite dimensionality of the domain and codomain implies that such a map will be an isomorphism, and so we will simply consider symplectomorphisms on a symplectic vector space V itself rather than between symplectic vector spaces.

Example 1.13. When $V = W = \mathbb{R}$, a map A is linear if and only if there exists a $\lambda \in \mathbb{R}$ such that:

$$A(v) = \lambda v.$$

Geometrically, A (more correctly: $\operatorname{im} A$) must be a line passing through the origin. This might be at odds with some terminology the reader has encountered before. In particular, a so-called linear functions of the form:

$$B(v) = \lambda v + \mu,$$

for fixed $\lambda, \mu \in \mathbb{R}$. Such a map is not "linear" as defined above, but actually satisfies a looser definition called "affine".

An **affine map** between (real) vector spaces V and W is a map $A \colon V \to W$ such that there exists a linear map $B \colon V \to W$ and $b \in W$ such that:

$$A(v) = B(v) + b.$$

The idea is that an affine mimics subspace-preserving properties of a linear function, but without the datum of a distinguished "origin" for your vector spaces. In this way, an affine map is a composition of a linear map with a translation.

Example 1.14. Contrasting with the linear case, when $V = W = \mathbb{R}$, a map A is affine if and only if there exists $\lambda, \mu \in \mathbb{R}$ such that:

$$A(v) = \lambda v + \mu.$$

So the affine maps consist of all lines passing through any point, not only those that pass through the origin.

Just as we defined the symplectic group $\mathrm{Sp}(2n)$ using linear maps, we can do similarly with affine maps. Define an **affine symplectomorphism** on $\mathbb{R}^{2n}$ to be a map $A \colon \mathbb{R}^{2n} \to \mathbb{R}^{2n}$ such that there exists a $\Psi \in \mathrm{Sp}(2n)$ and $b \in \mathbb{R}^{2n}$ such that:

$$A(v) = \Psi(v) + b.$$

Denote the group[3] of affine symplectomorphisms to be $\mathrm{ASp}(2n)$.

2 Non-squeezing

2.1 Volume preserving maps

Recall (or learn for the first time!) that the determinant of an $n \times n$ matrix A is the n-dimensional volume of the parallelotope spanned by the columns (or

[3]An additional exercise might be to justify the use of this word here.

equivalently, the rows) of A. A linear map $\Psi\colon \mathbb{R}^n \to \mathbb{R}^n$ is said to be **volume preserving** if for any $n \times n$ matrix,

$$A = \begin{bmatrix} | & | & & | \\ c_1 & c_2 & \cdots & c_n \\ | & | & & | \end{bmatrix},$$

we have:

$$|\det(\Psi A)| = \left|\det\left(\begin{bmatrix} | & | & & | \\ \Psi(c_1) & \Psi(c_2) & \cdots & \Psi(c_n) \\ | & | & & | \end{bmatrix}\right)\right| = |\det A|.$$

That is to say, Ψ is volume preserving if and only if $\det \Psi = \pm 1$, it taking parallelotopes to parallelotopes of the same volume modulo sign. From the properties of the determinant, we see that for a symplectopmorphism $\Psi \in \mathrm{Sp}(2n)$:

$$\det(\Psi^T \omega_0 \Psi) = \det \omega_0 \implies (\det \Psi)^2 = 1 \implies \det \Psi = \pm 1.$$

And thus symplectomorphisms are examples of volume preserving maps.

2.2 Camels

Consider the following biblical quote:

> And again I say unto you, It is easier for a camel to go through the eye of a needle, than for a rich man to enter into the kingdom of God.
>
> Matthew 19:24 (KJV)

Of course, where only volume preserving is concerned, a mathematician equipped would no doubt be able to *squeeze* a camel through the eye of a needle, the camel effectively deforming into a noodle to thread it through the eye.

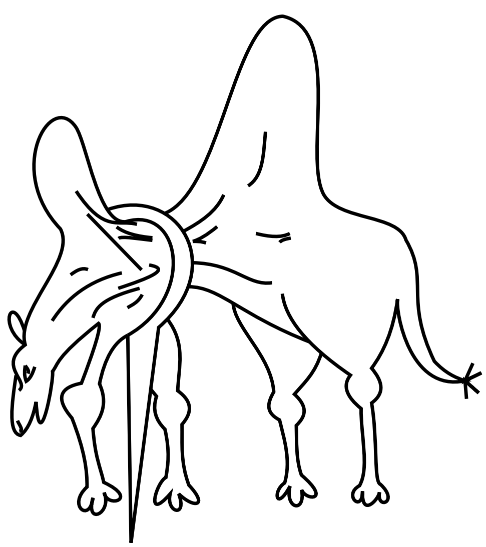

A little easier on the imagination is wondering how simple shapes such as balls deform under particular classes of transformations. Early in history it was not clear how symplectomorphisms (in the smooth, differential geometric sense of the word—not linear ones!) behaved with regards to something even as simple as a ball. As it turns out, Mikhail Gromov famously proved that the camel definitely cannot go through the eye of the needle if Jesus were assuming the camel to be *symplectic*.

2.3 Non-squeezing

We simplify Jesus's setup. Consider a (2-dimensional) ball of radius r (the "camel") and a wall with a cylindrical hole of radius R ("the needle"). The idea is that in order for the ball to pass through the hole from one side of the wall to the other (that is, "go through the eye of a needle"), it surely must be able to embed (symplectically) in a cylinder of infinite length. Consider $\mathbb{R}^{2n}$ as the set of points with coordinates:

$$(x_1, x_2, \ldots, x_n, y_1, y_2, \ldots, y_n).$$

Define the following objects:

- The **symplectic (closed) ball** of radius $r > 0$,

$$B(r) := \{(x_1, y_1, \ldots, x_n, y_n) \in \mathbb{R}^{2n} \mid \sum {x_i}^2 + {y_i}^2 \leqq r^2\}.$$

- The **symplectic cylinder** of radius $r > 0$,

$$Z(r) := \{(x_1, y_1, \ldots, x_n, y_n) \in \mathbb{R}^{2n} \mid {x_1}^2 + {y_1}^2 \leqq r^2\}.$$

Note that the cylinder is "symplectic" in the sense that the coordinates x_1, y_1 spanning the disk correspond to basis elements e_1 and f_1 so that $\omega_0(e_1, f_1) \neq 0$.

Now the theorem we will prove is that in order for a ball $B(r)$ to symplectically embed (via an affine symplectomorphism in $\mathrm{ASp}(2n)$) into the cylinder $Z(R)$, we must have:

$$r \leqq R.$$

That is, the radius of the ball is less than or equal to that of the cylinder (the hole).

Theorem 2.1 (Affine non-squeezing)**.** *Let $A\colon \mathbb{R}^{2n} \to \mathbb{R}^{2n}$ be an affine symplectomorphism such that:*

$$A(B(r)) \subseteq Z(R).$$

Then $r \leqq R$.

Proof. Suppose without loss of generality that the radius of the ball is $r = 1$. Decompose A into its $\mathrm{Sp}(2n)$ and translation parts:

$$A(z) = \Psi(z) + b.$$

Note that $A(B(r)) \subseteq Z(R)$ is equivalent to the boundary of $B(r)$ being embedded in $Z(R)$. In particular, if $\{e_1, f_1, \ldots, e_n, f_n\}$ is the standard symplectic basis then the condition $A(B(r)) \subseteq Z(R)$ becomes[a]:

$$\sup_{\|z\|=1} \left\{ \langle A(z), e_1 \rangle^2 + \langle A(z), f_1 \rangle^2 \right\} \leqq R^2.$$

Because $A(z) = \Psi(z) + b$, we may write it as:

$$\begin{aligned} R^2 &\geqq \sup_{\|z\|=1} \left\{ \langle \Psi(z) + b, e_1 \rangle^2 + \langle \Psi(z) + b, f_1 \rangle^2 \right\} \\ &= \sup_{\|z\|=1} \left\{ (\langle \Psi(z), e_1 \rangle + \langle b, e_1 \rangle)^2 + (\langle \Psi(z), f_1 \rangle + \langle b, f_1 \rangle)^2 \right\} \\ &= \sup_{\|z\|=1} \left\{ (\langle z, \Psi^\top e_1 \rangle + \langle b, e_1 \rangle)^2 + (\langle z, \Psi^\top f_1 \rangle + \langle b, f_1 \rangle)^2 \right\}. \end{aligned} \tag{$\star$}$$

Note that:

$$\begin{aligned} \omega(\Psi^\top e_1, \Psi^\top f_1) &= \langle \Psi^\top e_1, \omega \Psi^\top f_1 \rangle \\ &= \langle e_1,\ \Psi \omega \Psi^\top f_1 \rangle, \end{aligned}$$

but $\Psi \omega \Psi^\top = \omega$, because $\Psi \in \mathrm{Sp}(2n)$, hence:

$$\omega(\Psi^\top e_1, \Psi^\top f_1) = \langle e_1, \omega f_1 \rangle = \omega(e_1, f_1) = 1.$$

Thus the Cauchy-Schwarz inequality gives:

$$\|\Psi^\top e_1\| \cdot \|\Psi^\top f_1\| \geqq \omega(\Psi^\top e_1, \Psi^\top f_1) = 1,$$

implying that one of $\|\Psi^\top e_1\|$ or $\|\Psi^\top f_1\|$ is greater than or equal to 1. Assuming without loss of generality that $\|\Psi^\top e_1\| \geqq 1$, then we can set

$$z_0 := \pm \frac{\Psi^\top e_1}{\|\Psi^\top e_1\|},$$

which is on the boundary of B and gives us a lowerbound for the supremum in ($\star$):

$$\underbrace{(\langle z_0, \Psi^\top e_1 \rangle + \langle b, e_1 \rangle)^2}_{\geqq (1+\varepsilon)^2 \geqq 1} + \underbrace{(\langle z_0, \Psi^\top f_1 \rangle + \langle b, f_1 \rangle)^2}_{\geqq 0} \quad \geqq 1.$$

Hence we have that the supremum is greater or equal to 1, concluding:

$$R \geqq 1.$$

■

[a] Why does this make sense? The products $\langle A(z), e_1 \rangle$ and $\langle A(z), f_1 \rangle$ fetch the coefficients of $A(z)$ in the e_1 and f_1 coordinates, respectively. Thus what we have inside the supremum is the square of the "radius" of z (with respect to the circular face of the cylinder), which we'd expect for all z on the boundary (i.e. z such that $\|z\| = 1$) to be less than the square of the radius of the cylinder, R^2.

So Jesus's statement holds extra meaning if he was referring to a symplectic camel, and not just your regular joe volume preserving camel. In particular, there are balls of equal volume which are not symplectomorphic.

2.4 A step further

The story for affine non-squeezing does not stop there. We can define a set $B \subseteq \mathbb{R}^{2n}$ to be a **linear symplectic ball** of radius r if it is symplectomorphic to $B(r)$, and a set $Z \subseteq \mathbb{R}^{2n}$ a **linear symplectic cylinder** of radius R if it is symplectimorphic to $Z(R)$.

A $\Psi \in M(2n, \mathbb{R})$ has the **linear non-squeezing property** if for each linear symplectic ball B of radius r and linear symplectic cylinder Z of radius R, we have:

$$\Psi(B) \subseteq Z \implies r \leqq R.$$

With some effort, the following theorem may be proved:

Theorem 2.2. *If $\Psi \in M(2n, \mathbb{R})$ is an invertible matrix such that Ψ and Ψ^{-1} both have the linear non-squeezing property, then for all $u, v \in \mathbb{R}^{2n}$ either:*

$$\omega_0(\Psi u, \Psi v) = \omega_0(u, v), \quad \text{or} \quad \omega_0(\Psi u, \Psi v) = -\omega_0(u, v).$$

That is to say, either the matrix Ψ is symplectic ($\Psi \in \mathrm{Sp}(2n)$), or it is *anti-symplectic*.