Consider the collection G of isometric transformations on $\mathbb{R}^n$. Recall that a map $\psi : \mathbb{R}^n \to \mathbb{R}^n$ is called an **isometry** if for all x, y, $\|x - y\| = \|f(x) - f(y)\|$. Our first claim draws a connection between a specific class of isometries and orthogonal matrices:

Proposition 0.1

A transformation $\psi : \mathbb{R}^n \to \mathbb{R}^n$ is an origin-preserving (meaning $\psi(0) = 0$) isometry if and only if $\psi(x) = Ax$ for some orthogonal matrix A (meaning AA^T equals the identity).

Proof. Let ψ be an origin-preserving isometry. Let $\{e_i\}_{i=1}^n$ be the standard Euclidean basis of $\mathbb{R}^n$. Note that an isometry sends triangles to congruent triangles, so in particular law of cosine implies that it also preserves angles and dot products, which we call the "dot product identity." Consequently $\{\psi(e_i)\}_{i=1}^n$ forms an orthonormal basis of $\mathbb{R}^n$ like $\{e_i\}$ does. Our first observation is that ψ is linear. To see this, pick any $u, v, w \in \mathbb{R}^n$.The "dot product identity" shows

$$\psi(u + v) \cdot \psi(w) = (u + v) \cdot w$$

and

$$(\psi(u) + \psi(v)) \cdot \psi(w) = \psi(u) \cdot \psi(w) + \psi(v) \cdot \psi(w) = u \cdot w + v \cdot w = (u + v) \cdot w.$$

Keeping u, v fixed and letting $w = e_i$ shows that $\psi(u + v)$ and $\psi(u) + \psi(v)$ agree coordinate-wise, so they equal. Similarly $\psi(cu) = c\psi(u)$. Therefore ψ is a linear transformation characterized by some matrix A. But then using the dot product identity once again, for any standard basis vectors e_i, e_j,

$$\mathbf{1}[i = j] = (Ae_i) \cdot (Ae_j) = e_j^T A^T A e_i = (A^T A)_{i,j}.$$

This shows $A^T A$ is the identity (and so is AA^T).
The converse is trivial: if $\psi(x) = Ax$ then $(Au) \cdot (Av) = v^T A^T Au = u \cdot v$. Setting $u = v$ we see A maps basis to basis while also preserving length. □

With this result, we are able to fully characterize isometries of $\mathbb{R}^n$ as the composition of an origin-preserving isometry and a translation:

Theorem 0.2: TODO: iso char

Every isometry ψ of $\mathbb{R}^n$ can be characterized by an orthogonal matrix $A \in \mathrm{O}_n(\mathbb{R})$ and a "translation" vector b, such that $\psi(x) = Ax + b$.

Proof. First suppose $\psi(x) = Ax+b$ subject to the constraints above. This implies ψ is an isometry, since orthogonal matrices preserves distance and so does translation.
Conversely, let an isometry ψ be given. We first represent ψ as the composition of an origin-preserving isometry and a translation. To do so, write $\psi = (\psi - \psi(0)) + \psi(0)$. It remains to show that there exists an orthogonal matrix A such that $Ax = \psi(x) - \psi(0)$ for all x. This result now follows from the following lemma. □

It is easy to see that isometries form a group under composition. Among these transformations, **reflections** are of particular significance. These are the transformations that uses a certain **affine subspaces** as mirror, sending points across the mirror along the line perpendicular (orthogonal) to it. Formally a reflection ρ is an isometry of order 2

(i.e. $\rho = \rho^{-1}$ or $\rho^2 = I$, identity): indeed, mirroring twice and one ends up staying at the starting point. To transform u into v via a reflection, the affine mirror must lie in the "center" of the u, v. The canonical example in this case would be a $(n-1)$-dimensional hyperplane containing the "average" $(u+v)/2$ but also remaining orthogonal to the difference, $u - v$. Orthogonality gives rise to the following formula

$$H = \{x \in \mathbb{R}^n : (x - (u+v)/2) \cdot (u - v) = 0\} = \{x \in \mathbb{R}^n : x \cdot (u - v) = (u - v) \cdot (u + v)/2 = (\|u\| - \|v\|)/2\}.$$

From now on, given a reflection ρ, we use $\mathrm{aff}(\rho)$ to denote its corresponding affine subspace. In the next theorem, we show that reflections generate the entire group of isometry.

Theorem 0.3: TODO: iso ref

Every isometry ψ of $\mathbb{R}^n$ is a composition of at most $n+1$ reflections.

Proof. We may WLOG assume $\psi(0) = 0$, so that ψ is completely characterized by some $A \in \mathrm{O}_n(\mathbb{R})$, at the cost of at most one reflection. To see this, suppose $\psi(0) \neq 0$. We consider a reflection $\rho^{(0)}$ such that $\mathrm{aff}(\rho^{(0)})$ contains the midpoint $\psi(0)/2$ and is orthogonal to the line crossing 0 and $\psi(0)$. One can also verify that one such example is the reflection across the $(n-1)$-dimensional hyperplane described by $\{u \in \mathbb{R}^n : u \cdot (\psi(0) - 0) = u \cdot \psi(0) = (\|\psi(0)\| - \|0\|)/2 = \psi(0)^2/2\}$.

Now we assume $\psi(0) = 0$ and prove by induction. The case $n = 1$ is clear, since an origin-preserving isometry on $\mathbb{R}$ is either identity or negative identity.

For the inductive step assume the claim holds for $\mathbb{R}^{n-1}$. Consider an arbitrary ψ represented by $A \in \mathrm{O}_n(\mathbb{R})$, and let $\{e_i\}_{i=1}^n$ be the standard basis of $\mathbb{R}^n$. Repeating the argument in the first paragraph, we see that there exists a reflection f_n, possibly identity, such that $f_n(Ae_n) = (f_n \circ \psi)(e_n) = e_n$. Furthermore, since $A \in \mathrm{O}_n(\mathbb{R})$ we know $\|Ae_n\| = \|e_n\|$. This means $\mathrm{aff}(f_n) = \{u \in \mathbb{R}^n : u \cdot (Ae_n - e_n) = (\|Ae_n\| - \|e_n\|)/2 = 0\}$ contains the origin, and therefore $f_n \in \mathrm{O}_n(\mathbb{R})$.

By construction, $f_n \circ \psi$ equals identity on $E = \{0\}^{n-1} \times \mathbb{R}$, the subspace spanned by e_n, whereas its orthogonal complement, $H = \mathbb{R}^{n-1} \times \{0\}$, remains invariant under f_n. Our induction hypothesis states that there exist reflections $f_1, \cdots, f_{n-1}$ whose composition agrees with $f_n \circ \psi$ on H. We can easily define an extension of $f_1 \circ \cdots \circ f_{n-1}$ by setting its n^{th} component to be an identity mapping, i.e., $(f_1 \circ \cdots \circ f_{n-1})(\cdot, x_n) = ((f_1 \circ \cdots \circ f_{n-1})(\cdot), x_n)$. This extended function now coincides with $f_n \circ \psi$ on all of $\mathbb{R}^n$, so $\psi = f_n^{-1} \circ f_1 \circ \cdots \circ f_{n-1}$, and the inductive step is complete.

To sum up, we showed that every origin-preserving isometry of $\mathbb{R}^n$ is the composition of up to n reflections, and more generally, every isometry of $\mathbb{R}^n$ is the composition of up to $n+1$ reflections. □

Observe that in $\mathbb{R}^n$, any reflection ρ satisfies $\dim \mathrm{aff}(\rho) \leqslant n - 1$. For a simple example, in $\mathbb{R}^2$, reflections are characterized by 0- or 1-dimensional affine subspaces, namely, points or lines. And a simple drawing shows that if we have two line reflections across ℓ_1 and ℓ_2, then they commute if and only if ($\ell_1 = \ell_2$ or $\ell_1 \perp \ell_2$). In $\mathbb{R}^n$, the high-dimensional analogy also holds.

Proposition 0.4: TODO: commute

Let ρ_1, ρ_2 be two reflections of $\mathbb{R}^n$, with affine subspaces H_1, H_2. Then they commute if and only if one of the following **commutativity criteria** holds:

(1) One of the subspaces is contained in the other, or

(2) $H_1 \cap H_2 \neq \varnothing$, and every $u \in H_1 \backslash (H_1 \cap H_2) = H_1 \backslash H_2$ is orthogonal to every $v \in H_2 \backslash H_1$.

For convenience we call these the "reflection commutativity criterion." In either cases, $\operatorname{aff}(\rho_1 \circ \rho_2) = \operatorname{aff}(\rho_2 \circ \rho_1) = H_1 \cap H_2$.

Remark. Despite sounding like an orthogonality criterion, condition (2) cannot be replaced with "H_1 and H_2 are orthogonal." Consider for example two reflections in $\mathbb{R}^3$, one across the xy-plane: $(x, y, z) \mapsto (x, y, -z)$ and one across the yz-plane: $(x, y, z) \mapsto (-x, y, z)$. They clearly commute, but the two planes, when viewed as subspaces, are not orthogonal — they intersect at a line.

This observation suggests a certain structure of composition of commutative reflections. Note that if ρ_1 commutes with ρ_2, then their composition is yet another reflection: $(\rho_1\rho_2)(\rho_1\rho_2)^{-1} = \rho_1\rho_2\rho_2^{-1}\rho_1^{-1} = \rho_1\rho_1^{-1} = I$. This suggests we can define commutative groups of reflection, with identity transformation being the identity.

Theorem 0.5: TODO: max commute

In $\mathbb{R}^n$, a commutative group of reflections under composition has at most 2^n elements, and this upper bound can be attained.

Proof. It's trivial to construct a commutative group of reflections with 2^n elements. For $1 \leqslant i \leqslant n$, define ρ_i to be the mapping that flips the sign of the i^{th} coordinate while keeping others unchanged. It follows that the ρ_i's commute, and that they generate a group of 2^n elements — pick any subset of the n coordinates in $\mathbb{R}^n$ and flip the sign.

Conversely, let G_n be a commutative group of reflections of $\mathbb{R}^n$. For each $1 \leqslant k \leqslant n$, consider $\{\rho \in G_n : \dim \operatorname{aff}(\rho) = k\}$, the subset of reflections whose affine subspace has dimension k. They have to commute pairwise so they have to satisfy the reflection commutativity criterion, and clearly they need to satisfy (2). When their dimensions match, criterion (2) is equivalent to requiring that the normal vectors to these subspaces are orthogonal. In $\mathbb{R}^n$ there exist $\binom{n}{k}$ pairwise orthogonal vectors that can serve as normal vectors to k-dimensional affine subspaces. Therefore $|\{\rho \in G_n : \dim \operatorname{aff}(\rho) = k\}| \leqslant \binom{n}{k}$, and

$$|G_n| = \sum_{k=0}^{n} |\{\rho \in G : \dim \operatorname{aff}(\rho) = k\}| \leqslant \sum_{k=0}^{n} \binom{n}{k} = 2^n. \tag{Δ}$$

Let $|G_n| = 2^n$ be a commutative group of reflections of $\mathbb{R}^n$. There are two "special" elements in this group, namely when $k = n$ and when $k = 0$ as in (Δ). The former corresponds to a reflection whose axis is the entire space — this is the identity transformation. The latter, on the other hand, is the reflection across a 0-dimensional affine subspace, namely a point, if one assumes its existence a priori. We call it the **point reflection** associated with G_n. This point

reflection is special in the sense that it "looks the same from every perspective[1]," whereas for every other reflection $\rho \in G_n$, the effect imposed by ρ on a region depends on the location of the region and its relation to the affine subspace associated with ρ.

In other words, the point reflection is the "special" element of G_n displaying a certain invariance property. This should remind one of the notion of conjugation in abstract algebra. Let us quickly recall that if G is a group, then $g, h \in G$ are **conjugates** of each other if there exists a $f \in G$ such that $h = fgf^{-1}$. It follows by group properties that conjugacy is an equivalence relation, which gives rise to **conjugacy classes**.

Going back to the group of isometries. Let f, g be two isometric transformations. We say x is invariant under g if $x = g(x)$. An immediate result following conjugation is that x is invariant under g if and only if $f(x)$ is invariant under fgf^{-1}:

$$x = g(x) \Rightarrow (fgf^{-1})(f(x)) = (fg)(x) = f(x)$$

and conversely

$$(fgf^{-1})(f(x)) = f(x) \Rightarrow (fgf^{-1})(f(x)) = fg(x) = f(x) \Rightarrow g(x) = x$$

since isometries are injective, directly by definition: $\|f(x) - f(y)\| = 0$ if and only if $\|x - y\| = 0$.

Now we restrict our attention to point reflections — x is invariant under a reflection ρ if and only if x belongs to the affine subspace associated with ρ. Point reflections only have one fixed point, namely the "point" across which the reflection is performed. Therefore, if p is the **point reflection** associated with G_n and f is *any* other isometry (we don't require it to be another reflection), then fpf^{-1} is also a point reflection. Using the characterizations of commutative reflections, we see that these two point reflections p_1, p_2 commute if and only if they are identical — otherwise, $\mathrm{aff}(p_1) \cap \mathrm{aff}(p_2) = \varnothing$, and both condition fails. Put formally:

Definition 0.6: TODO: point ref

A point reflection p of $\mathbb{R}^n$ is an isometry satisfying the following:

(1) p is not the identity mapping, and

(2) For any other isometry f, the composition fpf^{-1} either equals p or does not commute with it.

Note that this definition itself does not invoke any notion of points, nor is it dependent on dimensionality, unlike the theorem on maximal commutative group of reflections. Having characterized point reflections, now we may work our way backwards and recover other primitives and notions under standard Euclidean geometry.

First, linearity. Just like how two points define a line, we can fix two point reflections and define a corresponding line reflection. In order to do so, we need to establish a strict partial order on reflections with respect to the dimension and containment of their associated affine subspaces. The idea builds on the following fact: if $H_1 \subset H_2$ are with dimensions $d_1 < d_2$, then there exists an $H_1' \subset H_2$ such that $\dim(H_1') = \dim(H_1)$, and they have orthogonal normal vectors. In order to formally state these, we need to (i) define subspace containment, and (ii) define orthogonal subspaces of the same dimension. Luckily, with conjugation, we have all the ingredients to cook up these notions. We use the symbol $\rho_1 \prec \rho_2$ to represent that ρ_1's subspace is strictly contained in ρ_2's.

[1] When I first learned abstract algebra, I would always visit Math SE for intuitive explanations behind concepts. Here I quoted an answer on conjugation and group actions.

Definition 0.7: TODO: prec

Given two commutative reflections ρ_1, ρ_2 of $\mathbb{R}^n$, we say $\rho_1 \prec \rho_2$ if the following holds:

(1) There exists a translation f leaving ρ_2 invariant under conjugation but not ρ_1, i.e., $\rho_1 \neq f\rho_1 f^{-1}$, but $\rho_2 = f\rho_2 f^{-1}$, and

(2) There does *not* exist a translation g leaving ρ_1 invariant but not ρ_2.

The conditions impose structural constraints on ρ_1 and ρ_2's subspaces: $\mathrm{aff}(\rho_2)$ has "extra" free dimensions whereas $\mathrm{aff}(\rho_1)$ doesn't. This ensures that if $\rho_1 \prec \rho_2$, then $\mathrm{aff}(\rho_1) \subset \mathrm{aff}(\rho_2)$ with a strictly lower dimension. By definition, $\prec$ is irreflexive and asymmetric. To show transitivity, suppose $\rho_1 \prec \rho_2 \prec \rho_3$, with f_{12} fixing ρ_2 but not ρ_1, and f_{23} fixing ρ_3 but not ρ_2. Then

$$(f_{12}f_{23})\rho_3(f_{12}f_{23})^{-1} = f_{12}(f_{23}\rho_3 f_{23}^{-1})f_{12}^{-1} = f_{12}\rho_3 f_{12}^{-1}.$$

Since f_{12} leaves ρ_2 invariant and $\rho_2 \prec \rho_3$, by definition f_{12} also leaves ρ_3 invariant, so the above equals ρ_3. The other direction is trivial — if a translation leaves ρ_1 invariant under conjugation then by definition it leaves ρ_2 invariant. Repeating this argument once again, we see it must also leave ρ_3 invariant.

While we proved $\prec$ to be a partial order, it is still a much weaker notion than dimensioanlity, since only reflections with overlapping affine subspaces are comparable. Indeed, in $\mathbb{R}^n$ we may compare the dimension of two arbitrary suitably nice regions But this will suffice in our upcoming definitions.

Definition 0.8: TODO: line ref

Given two point reflections p_1 and p_2, the **line reflection** $\ell = \ell(p_1, p_2)$ associated with p_1 and p_2 is the minimal reflection (w.r.t. $\prec$) such that $p_1 \prec \ell$ and $p_2 \prec \ell$. In other words, if $p_1, p_2 \prec \rho$ for some other reflection ρ, then $\ell \prec \rho$.

It naturally follows that a point reflection p_3 is **colinear** with p_1 and p_2 if $p_3 \prec \ell(p_1, p_2)$. More generally, we can iteratively apply this definition and recover the entire structure of the Euclidean space. Define a point reflection to be of **dimension** 0 and a line reflection of dimension 1. Given two reflections ρ_1, ρ_2 of dimensions $n-1$, we may define a corresponding reflection $\rho = \rho(\rho_1, \rho_2)$ of dimension n to be the minimal reflection w.r.t. $\prec$ satisfying $\rho_1 \prec \rho, \rho_2 \prec \rho$. Finally, given three reflections ρ_1, ρ_2, ρ_3 of the same dimension, we say (ρ_1, ρ_2) is **congruent** to (ρ_2, ρ_3) if there exists a translation τ such that $\tau \circ \rho_1 = \rho_2$ and $\tau \circ \rho_2 = \rho_3$.