Paradoxical Decompositions: Galileo to Banach & Tarski

Joel H. Shapiro

January 6, 2015

Abstract. These notes describe some of the paradoxical decompositions arising in mathematics, beginning with Galileo's 1638 observation that there are as many squares as positive integers, and ending with the Banach-Tarski Paradox, which implies that the unit ball of $\mathbb{R}^3$ contains finitely many disjoint pieces which can rigidly shifted and then reassembled into *two* copies of the unit ball.

Contents

1 *Galileo's Paradox*

In 1638, Galileo described the problem of trying to count the points of an infinite set. He observed that within the set of natural numbers $\mathbb{N}$, the subset of squares is in one-to-one correspondence with the entire set. Thanks to Dedekind and Cantor the resolution of this apparent paradox is now learned by every undergraduate mathematics major. Nevertheless, let's take Galileo's problem a little further and observe that the subcollection of non-squares is also in one-to-one correspondence with the entire set $\mathbb{N}$. Upon denoting by S the set of squares and T the set of non-squares, we have $\mathbb{N}$ expressed as the union of disjoint sets S and T, even though both $S \sim \mathbb{N}$ and $T \sim \mathbb{N}$ (where "$\sim$" means "is in one-to-one correspondence with"). In other words, $\mathbb{N}$ is partitioned into two disjoint sets, each of which has, in the sense of cardinality, the same size as $\mathbb{N}$ itself!

Dialogues Concerning Two New Sciences, Dover, New York 1954., pp. 31–33.

This phenomenon is, in fact, present in *any* infinite set. Let's recall that a set is defined to be *infinite* provided it can be put into one-to-one correspondence with a proper subset (Dedekind 1888), and that two sets are said to have the *same cardinality* if they can be put into one-to-one correspondence with each other, i.e. if there is a bijective map taking one onto the other.

Theorem 1.1 (Paradoxical decomposition of any infinite set)**.** *Every infinite set X can be partitioned into two disjoint subsets, each having the same cardinality as X.*

Using $|\cdot|$ to denote cardinality, this theorem says: $|X| = 2|X|$ if X is infinite.

Proof. By Zorn's Lemma there is a maximal family $\mathscr{C} = \{C_i : i \in I\}$ of pairwise disjoint countable (="countably infinite") subsets of X. The complement of the union of this family of sets must be finite; if not it would have a countable subset A that could be added to the family $\mathscr{C}$, violating its maximality. Fix an index $i \in I$ and replace C_i by $C_i \cup A$, resulting in a new countable set, which we'll still designate by C_i. The new C_i is still disjoint from all the other ones, and now $X = \bigcup_{i\in I} C_i$, where on the right-hand side we have a union of pairwise disjoint countable sets. Now *à la* Galileo we can partition each set C_i into two countable sets S_i and T_i (e.g. index the elements of C_i by natural numbers and take one of the sets to be the elements with even indices and the other with odd indices). Let $S^* = \bigcup_{i\in I} S_i$ and $T^* = \bigcup_{i\in I} T_i$. Then X is the disjoint union of S^* and T^*, but $S^* \sim X \sim T^*$. □

Every chain of families of pairwise disjoint countable subsets of X (the order being "family-inclusion") has an upper bound, namely the collection of all the families in the chain (which is still a family of pairwise disjoint countable sets). Thus Zorn's Lemma applies.

I thank John Erdman for his helpful comments, correcting an earlier version of this proof.

2 *Paradoxical Decomposition of a Remarkable Group*

THE GROUP IN QUESTION is F_2, the *free group group on two generators.* Its elements are "reduced words:" formal expressions $x_1 x_2 \cdots x_n$ for $n \in \mathbb{N}$ where each x_j comes from the set of symbols $\{a, a^{-1}, b, b^{-1}\}$, subject only to the restriction that no symbol occurs next to its "inverse." Multiplication in F_2 is defined to be concatenation of words, followed by "reduction" (e.g. $aba^{-1} \cdot abba = abbba = ab^3a$) [1]. Upon allowing the "empty word" to belong to F_2 we obtain a group.

In what follows we use the symbol "$\uplus$" to denote "disjoint union" i.e. the union of pairwise disjoint sets.

Theorem 2.1. *There exist pairwise disjoint subsets $\{A_1, A_2, B_1, B_2\}$ of F_2 such that*

$$F_2 = A_1 \uplus aA_2 = B_1 \uplus bB_2 \,.$$

Thus some portion of F_2 can be broken up, and then reassembled, using only the action of F_2 on itself, into *two* copies of F_2.

Proof. For $x \in \{a, a^{-1}, b, b^{-1}\}$ let $W(x)$ denote the set of reduced words that begin with x. Let

$$A_1 = W(a), \quad A_2 = W(a^{-1}), \quad B_1 = W(b), \quad \text{and} \quad B_2 = W(b^{-1}) \,.$$

Clearly these are pairwise disjoint subsets of F_2; note that $aA_2 = aW(a^{-1})$ is the set of reduced words in F_2 that don't begin with a, hence $F_2 = A_1 \uplus aA_2$. Similarly $bB_2 = W(b^{-1})$ is the set of reduced words in F_2 that don't begin with b, hence $F_2 = B_1 \uplus bB_2$, as promised. □

SIMPLE AS IT SEEMS, this "paradoxical" property of F_2 is the foundation for the Banach-Tarski Paradox, the proof of which depends crucially upon finding, within the group of rotations of the ball, a subgroup isomorphic to F_2. The paradoxical nature of F_2 can then be transferred to provide a similar phenomenon for the ball by means of a process that we'll now explore.

[1] To render this operation "well-defined" it must be shown that the same reduced word results no matter how this reduction is performed. This is not completely trivial; see e.g. Magnus, Karrass, and Solitar, *Combinatorial Group Theory* (Dover 1976), Theorem 1.2, pp. 134-5 for the details.

Exercise: Convince yourself that F_2 is isomorphic to the fundamental group of a figure-eight.

For example, a and $ab^{-1}ab^2$ belong to $W(a)$, while b, b^{-1}, and $a^{-1}ba^2b^{-1}$ do not.

Note that $A_1 \uplus A_2 \uplus B_1 \uplus B_2$ is *not* F_2; rather, it is $F_2 \backslash \{e\}$, where e is the identity element of F_2 (the "empty word"). It's an easy exercise to modify the argument given here so that the empty word is not excluded from the union of the A's and B's. Just replace A_1 by $A_1 \uplus \{e, a^{-1}, a^{-2}, \ldots\}$ and A_2 by $A_2 \backslash \{a^{-1}, a^{-2}, \ldots\}$. We still have $F_2 = A_1 \uplus aA_2 = B_1 \uplus bB_2$, but now $A_1 \uplus A_2 \uplus B_1 \uplus B_2 = F_2$.

3 *Paradoxical decompositions of sets*

NOW SUPPOSE THAT X is a set and G is a group of self-maps of X. The group operation is composition, hence each of the maps in G is a *bijection,* i.e. a one-to-one map taking X onto itself.

Definition 3.1. To say a subset E of X is *G-paradoxical* means that there is an "initial family" of pairwise disjoint subsets $\{E_i : 1 \leq i \leq n\}$ of E, group elements $\{g_i : 1 \leq i \leq n\}$, and an integer $1 \leq m < n$ such that

$$E = \bigcup_{i=1}^{m} g_i E_i = \bigcup_{i=m+1}^{n} g_i A_i .$$

In other words, a G-paradoxical subset of X contains a subset that can be partitioned into a finite number of pieces which can be re-assembled by group actions alone into *two* copies of the original group.

Remarks 3.2. Regarding the definition of "paradoxical:"

(a) If E is paradoxical with respect to a subgroup of G then it's clearly paradoxical with respect to G itself.

(b) It's not required that either of the "final families" $\{g_i E_i\}_1^m$ or $\{g_i E_i\}_{m+1}^n$ be pairwise disjoint, but *this "final disjointness" can be always be arranged.*

Proof. Suppose more generally that $\{E_i\}_1^n$ is any pairwise disjoint family of subsets of E and $F = \cup_1^n g_i E_i$, where $\{g_i\}_1^n$ is a collection of distinct transformations in G. For $1 \leq i \leq n$ let

$$E_i' = E_i \setminus \cup_{j=i+1}^n g_i^{-1} E_j .$$

Then the new initial family $\{E_i'\}_1^n$ is pairwise disjoint with its union contained in that of the original initial family, and the new final family $\{g_i E_i'\}_1^n$ is now pairwise disjoint, and has the same union as the original one. □

(c) It's not required that the initial family $\{E_i\}_1^n$ exhaust E, but we'll see in §7 (Corollary 7.2) that this, too, can always be arranged. The proof is remarkably remarkable.

Example 3.3 (Galileo's Paradox revisited)**.** *Suppose G is the group of bijections of the set $\mathbb{Z}$ of all integers. Then the natural numbers $\mathbb{N}$ form a G-paradoxical subset of $\mathbb{Z}$.*

Proof. We have $\mathbb{N} = E_1 \uplus E_2$, where E_1 denotes the set of even (positive) integers, and E_2 the odds. Define $g_1 : E_1 \to \mathbb{N}$ by $g_1(n) = n/2$. Since $\mathbb{Z}\setminus E_1$ and $\mathbb{Z}\setminus\mathbb{N}$ are both countable sets, we can extend g_1 to a bijection (which we'll still denote by g_1) of $\mathbb{Z}$ onto itself. Thus there exists $g_1 \in G$ with $g_1 E_1 = \mathbb{N}$, and similarly there exists $g_2 \in G$ with $g_2 E_2 = \mathbb{N}$. □

Any group acts on itself by (say) left-multiplication, and so can be paradoxical with respect to this self-action. In such a case we'll just say the group is paradoxical. In this terminology, Theorem 2.1 asserts:

Theorem 2.1 revisited. *The free group F_2 on two letters is paradoxical.*

HERE IS THE KEY to all that follows: it's possible for groups can transfer paradoxicality to sets upon which they act. This will happen whenever a paradoxical group acts *freely* on the set, meaning that only the identity transformation is allowed to have fixed points.

Theorem 3.4. *Suppose G is a paradoxical group of transformations that acts freely on a set X. Then X is G-paradoxical.*

Proof. We're given the situation of Definition 3.1. Now G partitions X in two different ways:

- Into pairwise disjoint orbits $\{Gx : x \in X\}$, and
- Into pairwise disjoint sets "transverse" to orbits, i.e. sets $\{gM : g \in G\}$, where the M is the "choice set" obtained by choosing one point from each G-orbit.

Axiom of Choice warning!!

The first of these results is a standard exercise, which I leave to you. The second one is also not difficult, but requires the hypothesis that the G act freely on X; it goes like this. By the definition of M, each G-orbit has the form Gm for some $m \in M$. Fix $x \in X$; it belongs to some G-orbit, so $x = gm$ for some $g \in G$ and $m \in M$, i.e. $x \in gM$. Thus $X = \bigcup\{gM : g \in G\}$. As for the disjointness, suppose g and h belong to G, and that gM has nonvoid intersection with hM, i.e. there exists $m_1, m_2 \in M$ with $gm_1 = hm_2$. Thus $h^{-1}gm_1 = m_2$, so m_2 lies in the G-orbit of m_1. By the definition of M this demands that $m_1 = m_2$, so $h^{-1}g$ fixes this point of X. Since G acts freely on X this implies that $h^{-1}g$ is the identity map on X, i.e. that $g = h$. Thus the sets gM for different $g \in G$ must be pairwise disjoint, as desired.

With these observations in hand we transfer the paradoxicality of G to X by setting $A_i^* = A_iM = \biguplus\{g(M) : g \in A_i\}$. Then the sets $\{A_i^* : 1 \le i \le n\}$ are pairwise disjoint in X, and

$$\bigcup_{i=1}^{m} g_i(A_i^*) \;=\; \bigcup_{i=1}^{m} (g_iA_i)M \;=\; \left(\bigcup_{i=1}^{m} g_iA_i\right) M \;=\; GM \;=\; X.$$

By the same computation, $\bigcup_{i=m+1}^{n} g_i(A_i^*) = X$, which establishes the promised G-paradoxicality of X. □

Remark 3.5. If our initial family $\{A_i\}_1^n$ exhausts G (as we've already noted can be arranged for $G = F_2$) then the corresponding family

$\{A_i^*\}_1^n$ exhausts X. Similarly, if one of the two final families of g_iA_i's is pairwise disjoint, so is the corresponding family of $g_iA_i^*$'s.

4 *The Hausdorff Paradox*

Let S^2 denote the unit sphere of $\mathbb{R}^3$, i.e. the set of points of three dimensional euclidean space that lie at distance 1 from the origin. Let $\mathscr{R}_3$ denote the group of rotations of $\mathbb{R}^3$ about the origin.

The Banach-Tarski Paradox is built upon an earlier result due to Hausdorff, who showed—about a century ago—that if a certain countable subset is removed from S^2, then "two-thirds of what remains" is $\mathscr{R}_3$-paradoxical. What's usually known these days as Hausdorff's paradox is a better result:

By a "rotation" of $\mathbb{R}^3$ about the point p we'll mean an isometric bijection of $\mathbb{R}^3$ that fixes each point of a line (the "axis" of the rotation) through p. In Appendix A below we'll deal carefully with this notion, showing that $\mathscr{R}_3$ actually is a group by establishing its isomorphism with the group $SO(3)$ of 3×3 orthogonal matrices of determinant 1.

Grundzüge der Mengenlehre, Veig, Leipzig, 1914. Reprinted by Chelsea, New York 1949, 1965, 1978.

Theorem 4.1. *There is a countable subset C of S^2 such that $S^2\backslash C$ is $\mathscr{R}_3$-paradoxical.*

The key to this result is the following property of the rotation group $\mathscr{R}_3$, the proof of which we'll defer for a moment.

Theorem 4.2. *$\mathscr{R}_3$ contains a subgroup $\mathscr{F}_2$ isomorphic to F_2, the free group on two letters.*

Proof of Theorem 4.1. Since F_2 is paradoxical, so is $\mathscr{F}_2$. Since F_2 is countable, so is $\mathscr{F}_2$, hence the set C in which the axes of the rotations in $\mathscr{F}_2$ intersect S^2 is a countable set. Moreover, C is taken into itself by each of the rotations in $\mathscr{F}_2$ (if p is fixed by $\rho \in \mathscr{F}_2$, then for each $\sigma \in \mathscr{F}_2$ the point $\sigma(p)$ is fixed by $\sigma\rho\sigma^{-1} \in \mathscr{F}_2$). Thus $\mathscr{F}_2$ acts *freely* on $S^2\backslash C$, so Theorem 3.4 guarantees that $S^2\backslash C$ is paradoxical for $\mathscr{F}_2$, and hence for $\mathscr{R}_3$. □

Our work on the paradoxicality of F_2 implies, more precisely, that

$$S^2\backslash D = \biguplus_{i=1}^{4} E_i = E_1 \uplus \rho E_2 = E_3 \uplus \sigma E_4$$

for some rotations $\rho, \sigma \in \mathscr{R}_3$.

Proof of Theorem 4.2. Our goal is to find two rotations of $\mathbb{R}^3$ about the origin with the property that no reduced word in these rotations and their inverses reduces to the identity transformation. For these we'll choose rotation through $\theta = \sin^{-1}\left(\frac{4}{5}\right)$ radians about the z-axis and the x-axis, respectively. The matrices representing these maps are ρ (for the rotation about the z-axis) and σ (for the rotation about the x-axis) below.

$$\rho := \begin{pmatrix} \frac{3}{5} & -\frac{4}{5} & 0 \\ \frac{4}{5} & \frac{3}{5} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \qquad \sigma := \begin{pmatrix} 1 & 0 & 0 \\ 0 & \frac{3}{5} & -\frac{4}{5} \\ 0 & \frac{4}{5} & \frac{3}{5} \end{pmatrix}.$$

Our goal now is to show that no reduced word in these matrices and their inverses multiplies out to I, the 3×3 identity matrix.

We'll deal instead with integer matrices:

$$5\rho = \begin{pmatrix} 3 & -4 & 0 \\ 4 & 3 & 0 \\ 0 & 0 & 5 \end{pmatrix} \qquad \text{and} \qquad 5\sigma = \begin{pmatrix} 5 & 0 & 0 \\ 0 & 3 & -4 \\ 0 & 4 & 3 \end{pmatrix}.$$

Since ρ and σ are orthogonal matrices, their inverses are their transposes, so to say a reduced word of length n in these matrices and their inverses is nontrivial is the same is saying that the corresponding word in 5ρ, 5σ, and *their transposes* does not multiply out to $5^n I$. For *this* it's enough to show that no such word multiplies out to a matrix all of whose entries are divisible by 5, i.e. that *over the field $\mathbb{Z}_5$ of integers modulo 5, no such word multiplies out to the zero-matrix!*

Over the field $\mathbb{Z}_5$ our matrices 5ρ, 5σ and their transposes become

$$r = \begin{pmatrix} 3 & 1 & 0 \\ 4 & 3 & 0 \\ 0 & 0 & 0 \end{pmatrix}, r' = \begin{pmatrix} 3 & 4 & 0 \\ 1 & 3 & 0 \\ 0 & 0 & 0 \end{pmatrix}, s = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 3 & 1 \\ 0 & 4 & 3 \end{pmatrix}, s' = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 3 & 4 \\ 0 & 1 & 3 \end{pmatrix}.$$

Let's call a word in the letters r, r', s, s' *admissible* if r never stands next to r', and s never next to s'. Our job now is to show that no admissible word in these new matrices multiplies out to the zero-matrix. We'll do this by proving something more precise:

We eschew the term "reduced" now because, while in our original setup we had, e.g. $\rho\rho^{-1} = \rho^{-1}\rho = I$, now we have $rr' = r'r = 0$, and similarly for s.

The kernel of each admissible word in the letters r, r', s, s' is the kernel of its last letter.

Here "kernel" means "left null-space," e.g. $\ker r = \{x \in (\mathbb{Z}_5)^3 : rx = 0\}$.

Proof. Each of the matrices r, r', s, s' has one dimensional range (i.e. column space) and two dimensional kernel. Upon calculating these ranges and kernels explicitly we find that the ranges of the "r-matrices" intersect the kernels of "s-matrices" in $\{0\}$, and the same is true of the way the kernels of r-matrices intersect the ranges of s-matrices.

We now proceed by induction on word-length. The result is trivial for words of length one. Suppose $n \geq 1$ and that the kernel of each admissible word of length n in r, r', s, s' equals the kernel of that word's last letter. We wish to prove that the same is true of every admissible word of length $n+1$. Let w be such a word, so $w = va$ where v is an admissible word of length n and $a \in \{r, r', s, s'\}$. Then $x \in \ker w$ means that $vax = 0$, i.e. that $ax \in \ker v \cap \operatorname{ran} a$. Since w is an admissible word, the last letter of v, call it b, is not a', and by the induction hypothesis $\ker v = \ker b$. Thus $ax \in \ker b \cap \operatorname{ran} a = \{0\}$, so $x \in \ker a$. We've shown that $\ker w \subset \ker a$. The opposite inclusion is trivial, so $\ker w = \ker a$, as we wished to show. □

This completes the proof of the result we're calling the Hausdorff Paradox.

5 *Equidecomposability*

THE HAUSDORFF PARADOX tells us that $S^2\backslash C$ is paradoxical for some countable subset C of S^2. In the next section we'll show that S^2 itself is paradoxical by "absorbing" the set C via a technique reminiscent of the one described for F_2 in the side-note to the proof of Theorem 2.1. To do this efficiently we need to examine more thoroughly the definition of "paradoxical." We'll begin by restating the conclusion of Remark 3.2(b).

Throughout this section, G will denote a group of self-mappings of a set X.

Proposition 5.1 (Definition of "paradoxical," revisited)**.** *$E \subset X$ is G-paradoxical if and only if there exists a pairwise disjoint family $\{E_i\}_1^n$, mappings $\{g_i\}_1^n \subset G$, and an integer $m \in [1, n)$ such that $E = \biguplus_{i=1}^m g_i E_i = \biguplus_{i=m+1}^n g_i E_i$.*

This result asserts that, in the original definition, the "initial family" $\{E_i\}_1^n$ can be chosen so that the two "final families" $\{g_i E_i\}_1^m$ and $\{g_i E_i\}_{m+1}^n$ are each *pairwise disjoint*.

At the heart of this refinement of the notion of "paradoxical" lies an important concept.

Definition 5.2 (Equidecomposability)**.** To say that subsets E and F are *G-equidecomposable* means that there exists a partition $\{E_i\}_1^n$ of E and a partition $\{F_i\}_i^n$ of F and mappings $\{g_i\}_1^n \subset G$ such that $F_i = g_i E_i$ $(1 \leq i \leq n)$).

A *partition* of a set is a pairwise disjoint family of subsets whose union is the whole set.

We'll abbreviate "E and F are G-equidecomposable" by "$E \sim_G F$." If the group G is understood (which it usually is) we'll just say "E and F are equidecomposable," and write $E \sim F$.

This definition allows a very efficient restatement of the definition of "paradoxical:"

Proposition 5.3 ("Paradoxical" re-visited yet again)**.** *$E \subset X$ is G-paradoxical if and only if there exist disjoint subsets A and B contained in E such that $A \sim_G E \sim_G B$.*

In the next section we'll revisit the definition of "paradoxical" one last time, and show that the sets A and B of the above Proposition can always be taken to partition E. Right now let's note that it's an easy exercise to show the relation "$\sim$" (a.k.a "$\sim_G$") on $\mathscr{P}(X)$ is reflexive ($E \sim E$ for every $E \subset X$) and symmetric ($E \sim F \implies F \sim E$ for every pair of subsets E, F of X). In fact, it's also transitive ($E \sim F$ & $F \sim H \implies E \sim H$), i.e. an *equivalence relation* on $\mathscr{P}(X)$.

Theorem 5.4. *Equidecomposability is an equivalence relation.*

Proof. To prove transitivity, suppose $E \sim_G F$ and $F \sim_G H$ for subsets E, F, H of X. Thus there exists partitions $\{E_i\}_1^n$ and $\{F_i\}_1^n$ of E and F

respectively, and transformations $\{g_i\}_1^n \subset G$ such that $g_i E_i = F_i$ for $1 \leq i \leq n$. There also exist partitions $\{F'_j\}_1^m$ and $\{H_j\}_1^m$ of F and H respectively, and transformations $\{h_j\}_1^m \subset G$ such that $h_j F'_j = H_j$.

Let $E_{i,j} = E_i \cap g_i^{-1} F'_j$, and set $\gamma_{i,j} = h_j g_i$ on $E_{i,j}$. Thus each $\gamma_{i,j} \in G$, and one checks easily that $\{E_{i,j} : 1 \leq i \leq n, 1 \leq j \leq m\}$ and $\{\gamma_{i,j} E_{i,j} : 1 \leq i \leq n, 1 \leq j \leq m\}$ partition E and H respectively. Thus $E \sim_G H$, as desired. □

The notion of "same cardinality" is defined in terms of arbitrary bijections. In this vein, "equidecomposable" is a refinement of that concept, defined in terms of special bijections. More precisely:

Definition 5.5 (Puzzle map)**.** For subsets E and F of X, to say a bijection φ of E onto F is a *G-puzzle map* (or, if G is understood, just a "puzzle map") means that there is a partition $\{E_i\}_1^n$ of E and transformations $\{g_i\}_1^n \subset G$ such that $\varphi \equiv g_i$ on E_i.

"Piecewise G-map" would also be reasonable terminology.

The terminology suggests that we think of E as a jigsaw puzzle assembled from some finite collection of pieces, which the puzzle map φ reassembles into another jigsaw puzzle F. With this definition we have the following equivalent formulation of the notion of equidecomposability:

Proposition 5.6 (Equidecomposability via puzzle maps)**.** *Subsets E and F of X are G-equidecomposable if and only if there is a G-puzzle map taking E onto F.*

The fact that G-equidecomposability is an equivalence relation can be explained in terms of puzzle maps: reflexivity means that the identy map is a puzzle map, symmetry means that the inverse of a puzzle map is a puzzle map, and the just-proved symmetry means that compositions of puzzle maps are puzzle maps.

The next result shows that paradoxicality is a property, not just of subsets of X, but even of $\sim_G$-equivalence classes.

Corollary 5.7. *Suppose E and F are G-equidecomposable subsets of X. If E is G-paradoxical then so is F.*

By symmetry of "$\sim$" this is really an "if and only if" result.

Proof. By Proposition 5.3 we're given disjoint subsets A and B of E such that $A \sim E \sim B$. Since $E \sim F$ we're given a puzzle map $\varphi\colon E \to F$. Since φ is one-to-one, $A' := \varphi(A)$ and $B' := \varphi(B)$ are disjoint subsets of F and since the restriction of a puzzle map is clearly a puzzle map, we know that $A' \sim A$ and $B' \sim B$. Thus by transitivity:

$$A' \sim A \sim E \sim F \quad \text{and} \quad B' \sim B \sim E \sim F$$

hence $A' \sim F \sim B'$. □

6 *The Banach-Tarski Paradox for S^2 and $\mathbb{B}^3$*

To this point we know that if we remove a certain countable subset C from S^2, then the remainder $S^2\backslash C$ is paradoxical for the group $\mathscr{R}_3$ of rotations of $\mathbb{R}^3$ about the origin. Now, with the help of the previous section's work on equidecomposability, we can show that S^2 itself is paradoxical. For this we'll build on the "absorption" technique used in the sidenote to Theorem 2.1 to refine the paradoxicality of the free group F_2. The point now, as it was then, is to create a set that contains the exceptional set and is taken into itself by some group element.

Proposition 6.1. *Suppose C is any countable subset of S^2. Then there exists a rotation $\rho \in \mathscr{R}_3$ of infinite order such that the infinite family of sets $\{\rho^n(C)\}_1^\infty$ is pairwise disjoint.*

Proof. By "infinite order" for a rotation ρ (or for any element of any group, for that matter) we mean that $\rho^n \neq$ identity for any positive integer n. To prove the Proposition it's enough find $\rho \in \mathscr{R}_3$ such that $\rho^n(C) \cap C = \emptyset$ for each positive integer n. For then if m and n are integers with $0 \leq m < n$ we'll have

$$\rho^m(C) \cap \rho^n(C) = \rho^m\left(C \cap \rho^{n-m}(C)\right) = \rho^m(\emptyset) = \emptyset$$

as desired.

To find the desired rotation ρ, choose a line L through the origin of $\mathbb{R}^3$ that does not intersect the set C and let ρ_θ denote the mapping rotation through angle θ with axis L (choose in advance either direction of L to be the positive one). For n a positive integer and $c \in C$ let $\Theta(n, c)$ denote the set of angles $\theta \in [0, 2\pi)$ for which $\rho_\theta^n(c) \in C$. There are at most countably many such angles, hence the set $\Theta := \bigcup\{\Theta(n, c) \colon n \in \mathbb{N}, c \in C\}$, the set of angles θ for which $\rho^n(C)$ intersects C for some $n \in \mathbb{N}$, is at most countable. Thus for any $\theta \in [0, 2\pi)\backslash\Theta$, the rotation ρ_θ will do the job. □

Theorem 6.2 (Banach-Tarski for S^2)**.** *The unit sphere S^2 of $\mathbb{R}^3$ is $\mathscr{R}_3$-paradoxical.*

Proof. We know from our strong version of the Hausdorff Paradox (Theorem 4.1) that S^2 contains a countable subset C such that $S^2\backslash C$ is paradoxical. For this set C choose $\rho \in \mathscr{R}_3$ according to Proposition 6.1. Let $C_\infty = \bigcup\{\rho^n(C) \colon n \in \mathbb{N}\}$, and note that $\rho(C_\infty) = C_\infty\backslash C$. Thus we have on the one hand: $S^2 = (S^2\backslash C_\infty) \cup C_\infty$, while on the other hand: $S^2\backslash C = (S^2\backslash C_\infty) \cup \rho(C_\infty)$. Thus $S^2 \sim_{\mathscr{R}_3} S^2\backslash C$ and since $S^2\backslash C$ is paradoxical, Corollary 5.7 assures us that the same is true of S^2. □

Theorem 6.2 gives the following start on the Banach-Tarski Paradox for the unit ball of $\mathbb{R}^3$.

Corollary 6.3. $\mathbb{B}^3\backslash\{0\}$ *is* $\mathscr{R}_3$*-paradoxical.*

Proof. According to Theorem 6.2 and Corollary 5.3 we know that S^2 contains disjoint subsets A and B such that

$$A \sim_{\mathscr{R}_3} S^2 \sim_{\mathscr{R}_3} B\,. \tag{1}$$

Let $A^* := \bigcup_{a\in A}\{ra\colon 0 < r \leq 1\}$, and similarly define B^*. Thus A^* and B^* are disjoint subsets of $\mathbb{B}^3\backslash\{0\}$, and it's an easy exercise to check that the same rotations that effect the equivalences of (1) above also show that

$$A^* \sim_{\mathscr{R}_3} \mathbb{B}^3\backslash\{0\} \sim_{\mathscr{R}_3} B^*\,.$$

Thus $\mathbb{B}^3\backslash\{0\}$ is (again by Corollary 5.7) paradoxical. □

Let's recall our original statement of the Banach-Tarski Paradox:

> ... the unit ball of $\mathbb{R}^3$ contains finitely many disjoint pieces which can rigidly shifted and then reassembled into *two* copies of the unit ball.

See the Abstract on page 1.

So far "rigidly shifted" has meant "rotated about the origin." However to go further we'll have to enlist a larger transformation group. This will be $\mathscr{G}_3$, the group of all isometric transformations of $\mathbb{R}^3$. In particular, any rotation about any center belongs to $\mathscr{G}_3$.

In what follows the group $\mathscr{G}_3$ of transformations does *not* act on $\mathbb{B}^3$. Rather, it acts on $\mathbb{R}^3$, and what we'll prove is that $\mathbb{B}^3$ is $\mathscr{G}_3$-paradoxical *as a subset of* $\mathbb{R}^3$.

Theorem 6.4 (Banach-Tarski for $\mathbb{B}^3$)**.** *The three dimensional unit ball is* $\mathscr{G}_3$*-paradoxical.*

Proof. Let L be the line through the point $(0, 0, \frac{1}{2})$ parallel to the x-axis, oriented (say) in the direction of the x-axis. Let ρ be rotation about L through an angle that's an irrational multiple of π. Thus ρ is an infinite-order element of $\mathscr{G}_3$. Let C denote the orbit of the origin under ρ, i.e. $C = \{\rho^n(0)\colon n = 0, 1, 2, \ldots\}$. Thus $C \subset \mathbb{B}^3$ and $\rho(C) = C\backslash\{0\}$. So on the one hand we have $\mathbb{B}^3 = (\mathbb{B}^3\backslash C) \cup C$, while on the other hand $\mathbb{B}^3\backslash\{0\} = (\mathbb{B}^3\backslash C) \cup \rho(C)$. Consequently $\mathbb{B}^3\backslash\{0\} \sim_{\mathscr{G}_3} \mathbb{B}^3$, hence $\mathbb{B}^3$ inherits the $\mathscr{G}_3$-paradoxicality of $\mathbb{B}^3\backslash\{0\}$. □

$\mathbb{B}^3\backslash\{0\}$ is $\mathscr{R}_3$-paradoxical, and $\mathscr{R}_3 \subset \mathscr{G}_3$, hence $\mathbb{B}^3\backslash\{0\}$ is $\mathscr{G}_3$ paradoxical.

In summary: Any ball of radius one can be thought of as containing a three dimensional jigsaw puzzle that can be reassembled, using only rotations (not all of them about the ball's center), into *two* balls of radius one. This raises further questions: Is every ball $\mathscr{G}_3$-equidecomposable with any other ball? With a cube? We'll take up these matters in the next section.

A careful look at our arguments from the point of view of puzzle maps shows that this initial puzzle is, in fact, the whole ball. This is not an accident, as we'll see in the next section.

7 *Banach-Tarski beyond* $\mathbb{B}^3$

In Section 1 we gave a modern rephrasing of Galileo's 1638 discussion of the difficulties one encounters in comparing the sizes of infinite sets. Using the notation "$A \sim B$" for "there exists a bijection of set A onto set B" (i.e. "A and B have the same cardinality") we expressed Galileo's paradox as follows:

> If $\mathbb{N}$ is the set of natural numbers, S the subset of squares, and T the subset of nonsquares, then, even though $\mathbb{N}$ is the disjoint union of S and T it's nonetheless true that $S \sim \mathbb{N} \sim T$.

Theorem 5.3 (page 12) phrases the notion of paradoxicality in similar terms, but now using a more sophisticated equivalence relation on sets: that of "equidecomposability." Like the notion of "same cardinality," equidecomposability can be defined in terms of bijections, but now the bijections are not arbitrary; they are *puzzle maps* (Proposition 5.6).

See Definition 5.5, page 9.

The deepest elementary result about "same cardinality" is the Schröder-Bernstein Theorem: if set A has the same cardinality as a subset of set B, and B has the same cardinality as a subset of A, then A and B have the same cardinality. It turns out that the same is true for equidecomposability, and that the two results have a common proof! In this section we'll give this proof, and examine its astonishing consequences for the notion of paradoxicality.

We'll assume as usual that G is a group of self-maps of a set X, and we'll continue to write $A \sim_G B$ for "A and B are G-equidecomposable." We'll introduce the new notation $A \preceq_G B$ to mean "A is G-equidecomposable with a subset of B," or in terms of bijections: "There is a puzzle map taking A onto a subset of B." Thus the relation $\preceq_G$ is *reflexive* since the identity map is a puzzle map, and *transitive* since the composition of puzzle maps is a puzzle map. The main result of this section and the next is that $\preceq_G$ is also *symmetric*, hence a partial order on $\mathscr{P}(X)$. This is the content of:

Theorem 7.1 (The Banach-Schröder-Bernstein Theorem)**.** *If A and B are subsets of X with $A \preceq_G B$ and $B \preceq_G A$, then $A \sim_G B$.*

Banach and Tarski, Fundamenta Math. 6 (1924) 244–277, Theorem 8, page 251. All the results of this section, as well as Theorem 6.4, come from this remarkable paper.

This theorem implies that the relation $\preceq_G$ is a *partial order* on $\mathscr{P}(X)$. We'll prove it in the next section; for the rest of this one let's examine some consequences. First recall that a subset E if X is *G-paradoxical* if and only if there are disjoint subsets A and B of E such that $A \sim_G E \sim_G B$ (Proposition 5.3). Thanks to Theorem 7.1 we can (finally!) prove that the sets A and B can be chosen so that their union is E.

See Proposition 5.3, page 8.

Corollary 7.2. *$E \subset X$ is G-paradoxical if and only if there is a partition $\{A, B\}$ of E with $A \sim_G E \sim_G B$.*

Proof. The key to this one (besides the Banach-Schröder-Bernstein Theorem) is the trivial fact that if $F \subset E$ then $F \preceq_G E$ (Proof: The identity map on F is a G-puzzle map). Suppose E is G-paradoxical, so there exist disjoint subsets A_0 and B of E with $A_0 \sim_G E \sim_G B$. Let $A = A_0 \uplus (E \backslash (A_0 \cup B))$. Then $A_0 \subset A$ so $A_0 \preceq_G A \preceq_G E$. But $E \preceq_G A_0$, so $E \preceq_G A \preceq_G E$, hence by Theorem 7.1 $E \sim_G A$. This proves the "paradoxical implies partition" direction. For the other direction there's nothing to prove. □

For subsets A and B of X we've already noted that

$$A \subset B \implies A \preceq_G B.$$

To proceed further we'll need another simple property of the ordering $\preceq_G$:

Lemma 7.3. *Suppose $\{A_j\}_1^n$ is a pairwise is disjoint family of subsets of X, as is also $\{B_j\}_1^n$.*

(a) If $A_j \preceq_G B_j$ for each index j, then $\uplus_{j=1}^n A_j \preceq_G \uplus_{j=1}^n A_j$.

(b) If $A_j \sim_G B_j$ for each index j, then $\uplus_{j=1}^n A_j \sim_G \uplus_{j=1}^n A_j$.

Proof. (a) Our hypothesis is that for each j there is a puzzle map φ_j taking A_j into B_j. Then it's easy to check that the map φ defined by setting $\varphi = \varphi_j$ on A_j is a puzzle map taking the union of the A_j's onto the union of the B_j's.

(b) Same as (a), except now the puzzle map φ_j takes A_j *onto* B_j $(1 \leq j \leq n)$, and therefore φ takes $\uplus_1^n A_j$ *onto* $\uplus_1^n B_j$. □

Corollary 7.4. *Any two balls in $\mathbb{R}^3$ are $\mathscr{G}_3$-equidecomposable.*

Proof. Fix a ball B in $\mathbb{R}^3$. It's enough to prove that B is equidecomposable with the closed unit ball $\mathbb{B}^3$.

Suppose first that the radius of B is > 1. Cover B by balls $\{B_j\}_1^n$ of radius equal to one, and "disjointify" this collection of B_j's by setting $B'_n = Bn$ and $B'_j := B_j \backslash \cup_{k=j+1}^n B_k$ for $1 \leq j < n$. Then $B'_j \subset B_j$ for each index j, and the new collection $\{B'_j\}_1^n$ has the same union as the original one; in particular it still covers B. Now let $\{C_j\}_1^n$ be a pairwise disjoint collection of balls of radius 1 in $\mathbb{R}^3$. Then (omitting the subscript $\mathscr{G}_3$)

$$\mathbb{B}^3 \preceq B \preceq \biguplus_{j=1}^n B'_j \preceq \biguplus_{j=1}^n C_j \sim \mathbb{B}^3$$

where the first "inequality" comes from the containment of $\mathbb{B}^3$ in a translate of B, the second one from the containment of B in the union of the B'_j s, the third one from Lemma 7.3 above along with the containment of each B'_j in a translate of the corresponding C_j, and the final "equality" from iteration of the Banach-Tarski Theorem (Theorem 6.4, page 11). Thus $\mathbb{B}^3 \preceq B \preceq \mathbb{B}^3$, so $B \sim \mathbb{B}^3$ by the Banach-Schröder-Bernstein Theorem.

If the radius of B is > 1, repeat the above argument with the roles of B and $\mathbb{B}^3$ reversed. If the radius of B is equal to 1 then B, being a translate of $\mathbb{B}^3$, is trivially $\mathcal{G}_3$-equidecomposable with that set. □

Corollary 7.5 (The "ultimate" Banach-Tarski Theorem)**.** *Any two bounded subsets of* $\mathbb{R}^3$ *with nonempty interior are* $\mathcal{G}_3$*-equidecomposable.*

Proof. Let E be a bounded subset of $\mathbb{R}^3$ with nonempty interior. It's enough to show that $E \sim \mathbb{B}^3$. Since E contains a ball B we know from Corollary 7.4 that $\mathbb{B}^3 \sim B \preceq E$. Since E is bounded it is contained in a ball B', so again by Corollary 7.4, $E \preceq B' \sim \mathbb{B}^3$. Thus $\mathbb{B}^3 \preceq E \preceq \mathbb{B}^3$, hence $E \sim \mathbb{B}^3$ by Banach-Schröder-Bernstein. □

8 *Proof of the Banach-Schröder-Bernstein Theorem.*

The argument depends crucially on a mapping result due to Banach.

Theorem 8.1 (The Banach Mapping Theorem)**.** *Suppose A is a set with subset C, and that there are mappings f taking A into another set B and g taking B back into A. Then there is a subset C of A such that g takes $B\backslash f(C)$* onto *$A\backslash C$.*

S. Banach, Fundamenta Math. 24 (1924) 236-239. In the proof below, the mappings f and g are not required (as they are in the above reference) to be one-to-one.

Discussion. The figure below illustrates what is going on here.

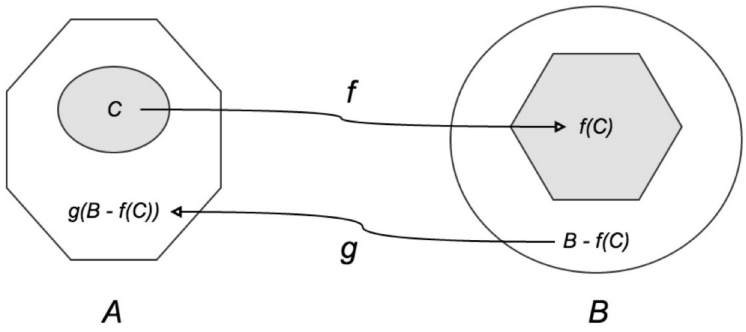

In symbols, the conclusion of Theorem 8.1 says that

$$g(B\backslash f(C)) = A\backslash C \tag{2}$$

or, upon taking the complement in A of both sides of that equation, that

$$C = A\backslash g(B\backslash f(C)) \quad \text{for some } C \in \mathscr{P}(A). \tag{3}$$

Let's define $\Phi : \mathscr{P}(A) \to \mathscr{P}(A)$ by

$$\Phi(E) = A\backslash g(B\backslash f(E)) \qquad (E \in \mathscr{P}(A)). \tag{4}$$

Then the conclusion of the Banach Mapping Theorem, in its form (3) above, says that:

The mapping Φ has a fixed point in $\mathscr{P}(A)$.

It's in this form that we'll prove the Banach Mapping Theorem. Right now let's see how it yields the Banach-Schröder-Bernstein Theorem.

Proof of the Banach-Schröder-Bernstein Theorem. The hypotheses assert that there are puzzle maps $f\colon A \to B_1 \subset B$ and $g\colon B \to A_1 \subset A$. By the Banach Mapping Theorem there is a subset C of A satisfying equation (2) above. Since g is a puzzle map, and the restriction of a puzzle map is again a puzzle map, this equation asserts that $B\backslash f(C) \sim A\backslash C$. Since f is a puzzle map we know that $f(C) \sim C$. Thus Lemma 7.3 insures that

$$B = (B\backslash f(C)) \cup f(C) \sim (A\backslash C) \cup C = A$$

as desired. □

Proof of the Banach Mapping Theorem. The mapping Φ defined by (4) is best understood as the composition of four simple set-mappings:

$$\mathscr{P}(A) \xrightarrow{f} \mathscr{P}(B) \xrightarrow{C_B} \mathscr{P}(B) \xrightarrow{g} \mathscr{P}(A) \xrightarrow{C_A} \mathscr{P}(A)$$

where C_A denotes "complement in A," and similarly for C_B. Two of these maps, f and g, preserve set-containment, while the other two reverse it. Thus the composite mapping *preserves* set-containment ($E \subset F \implies \Phi(E) \subset \Phi(F)$).

With these observations the Banach Mapping theorem, and with it the Schröder-Bernstein Theorem, follows from:

Theorem 8.2 (The Knaster-Tarski Theorem)**.** *Suppose A is a set, and $\Phi\colon \mathscr{P}(A) \to \mathscr{P}(A)$ is a mapping that preserves set-containment. Then Φ has a fixed point.*

B. Knaster, Ann. Soc. Polon. Math. 6 (1928) 133–134.

Proof. Let $\mathscr{E}$ be the collection of subsets E of A for which $E \subset \Phi(E)$. Note that $\mathscr{E}$ is non-empty; it contains the empty subset of A. Let C be the union of all the sets in the collection $\mathscr{E}$.

Claim: $\Phi(C) = C$.

Proof of CLAIM. Suppose $E \in \mathscr{E}$. Then $E \subset \Phi(E)$ by the definition of $\mathscr{E}$, and $\Phi(E) \subset \Phi(C)$ by the set-preserving nature of Φ. Thus $E \subset \Phi(C)$, so by the definition of C we have half of what we want: $C \subset \Phi(C)$. But this means that $\Phi(C) \subset \Phi(\Phi(C))$, so $\Phi(C) \in \mathscr{E}$, hence $\Phi(C) \subset C$. Thus $\Phi(C) = C$, which completes the proof of the Knaster-Tarski Theorem. □

IN THE PROOF JUST GIVEN for the Banach-Schröder-Bernstein Theorem, the only properties we used of the equidecomposability relation " $\sim$" are that it's an equivalence relation on $\mathscr{P}(X)$ that obeys the following two additional conditions:

(a) If $A, B \in \mathscr{P}(X)$ with $A \sim B$ then there exists a bijection $g : A \to B$ such that $C \sim g(C)$ for every $C \subset A$.

(b) If $A_1, A_2, B_1, B_2 \subset X$ with $A_1 \cap A_2 = \emptyset = B_1 \cap B_2$, $A_1 \sim B_2$, and $A_2 \sim B_2$, then $A_1 \cup B_1 \sim A_2 \cup B_2$.

If $\sim$ denotes equidecomposability then: for (a) one chooses g to be the puzzle map that defines the relations $A \sim B$ and observes that the restriction of a puzzle map is a puzzle map, while (b) is a special case of Lemma 7.3.

Thus, if $\sim$ is an equivalence relation on $\mathscr{P}(X)$ having properties (a) and (b) above, and if "$A \preceq B$" means "$A \sim A_0$ for some subset A_0 of B," then the relation $\preceq$ is a partial order on $\mathscr{P}(X)$.

In particular the notion of "same cardinality" is an equivalence relation on $\mathscr{P}(X)$ that satisfies (a) and (b) above, so the argument just given also proves the classical Schröder-Bernstein Theorem.

9 *Notes*

Amenability. We noted on page 6 that Hausdorff didn't quite prove the full strength of what we've called "The Hausdorff Paradox" (that S^2 with a certain countable subset is paradoxical.) Hausdorff's interest lay in the related question of whether or not one there exists a "finitely additive, rotation-invariant probability measure" on $\mathscr{P}(S^2)$. Hausdorff proved enough to be able to show that such objects cannot exist. It's an easy exercise to derive this from our Banach-Tarski Theorem (Theorem 6.2) for S^2.

The question of existence of such invariant set functions on groups and their connection with paradoxicality is of considerable interest. A group whose collection of all subsets supports a finitely additive probability measure is termed "amenable." In fact, Tarski proved that a group is paradoxical if and only if it's not amenable[2].

[2] Fundamenta Math. 31 (1938) 47–66.

Expository references. The gold standard for exposition on the Banach Tarski Paradox is Stan Wagon's eponymous book[3]. The internet is replete with expositions. Two of my favorites are Søren Knudby's undergraduate thesis[4] at the University of Copenhagen, and Terry Tao's preprint,[5] which contains the idea to use the field $\mathbb{Z}_5$ to prove that $SO(3)$ contains a free subgroup on two generators.

[3] *The Banach-Tarski Paradox,* Cambridge Univ. Press 1993.

[4] *The Banach-Tarski Paradox,* available online at `http://www.math.ku.dk/~knudby/`

[5] *The Banach-Tarski Paradox*, available online from `www.math.ucla.edu/~tao/preprints/`

Bibliography

Stefan Banach, *Un théorème sur les transformations biunivoques,* Fundamenta Mathematicae 24 (1924) 236-239.

Stefan Banach and Alfred Tarski, *Sur la décomposition des ensembles de points en parties respectivement congruent,* Fundamenta Mathematicae. 6 (1924) 244–277,

Galileo Galilie, *Dialogues Concerning Two New Sciences,* Dover, New York 1954. Translation by Crew and de Salvio of *Discorsi e Dimonstrazione Matematiche intorno à due nuove scienze,* Elzevir, Leiden 1638. This translation originally published by Macmillan, New York, 1914.

Felix Hausdorff, *Grundzüge der Mengenlehre,* Veig, Leipzig, 1914. Reprinted by Chelsea, New York 1949, 1965, 1978.

Bronisław Knaster, *Un théorème sur les fonctions d'ensembles,* Ann. Soc. Polon. Math. 6 (1928) 133–134.

Søren Knudby, *The Banach-Tarski Paradox,* available online at `http://www.math.ku.dk/~knudby/`

Terry Tao, *The Banach-Tarski Paradox*, available online at `www.math.ucla.edu/~tao/preprints/Expository/banach-tarski.pdf`.

Alfred Tarski, *Algebraische Fassung des Massproblems,* Fundamenta Math. 31 (1938) 47–66.

Stan Wagon, *The Banach-Tarski Paradox,* Cambridge Univ. Press 1993.

A Appendix: $\mathscr{R}_3$ as a matrix group

TO THIS POINT we've treated the notion of "rotation" (about the origin, in $\mathbb{R}^3$) informally, assuming that each rotation has an axis, each point of which it fixes, and that the rotations form a group under composition. This latter property is not obvious: If two rotations have different axes, why should their compositions have an axis? In this section we'll take some time to carefully define what we mean by "rotation about the origin in $\mathbb{R}^3$" and develop the properties needed for the arguments we've just given for the paradoxicality of S^2, $\mathbb{B}^3$ and beyond.

A.1 The isometries of $\mathbb{B}^n$

BY AN *isometry* of a subset of a metric space we mean a mapping of the set into itself that preserves distances. In Euclidean space, translations and rotations are isometries. For the unit ball of $\mathbb{R}^n$ the isometries are easily characterized in terms of orthogonal matrices, whose definition and basic properties we'll now review.

For $\mathbb{R}^n$ let's denote the inner product by $\langle \cdot, \cdot \rangle$, and the standard unit vector basis by $(e_j : 1 \leq j \leq n)$; e_j is the vector with 1 in the j-th coordinate and zeros elsewhere. We'll think of $\mathbb{R}^n$ as a space of column vectors. For any matrix A we'll denote its transpose by A^T.

Of crucial importance is the fundamental connection between the inner product and matrix transpose:

Proposition A.1. *For any $n \times n$ real matrix A,*

$$\langle Av, w \rangle = \langle v, A^T w \rangle \qquad (v, w \in \mathbb{R}^n).$$

Proof. It's enough to prove the result for vectors in the standard basis, so let $v = e_i$ and $w = e_j$. Then the left-hand side of the identity is just the (i, j)-element of the matrix A, while the right-hand side is, by the symmetry of the real inner product, $\langle A^T e_j, e_i \rangle$, the (j, i)-element of the transpose of A. Thus the right-hand side equals the left-hand side for these vectors hence, by the bilinearity of the inner product, for all vectors. □

Definition A.2 (Orthogonal matrices)**.** To say a square matrix with real entries is *orthogonal* means that its transpose is its inverse.

More precisely: an $n \times n$ matrix A orthogonal if and only if $AA^T = A^T A = I$ where I is the $n \times n$ identity matrix. We'll use $O(n)$ to denote the collection of all $n \times n$ orthogonal matrices.

Exercise: $O(n)$ is, for each positive integer n, a group under matrix multiplication.

Proposition A.3. *An $n \times n$ real matrix is orthogonal if and only if its columns form an orthonormal basis for $\mathbb{R}^n$.*

Proof. For an n-tuple of vectors in $\mathbb{R}^n$, orthonormality implies linear independence, and hence "basis-ness." Suppose A is an $n \times n$ real matrix. Then its j-th column is Ae_j, so by Proposition A.1 the inner product of the j-th and k-th columns is:

$$\langle Ae_j, Ae_k \rangle = \langle A^T Ae_j, e_k \rangle = \text{the } (i,j)\text{-element of } A^T A$$

Thus the n-tuple of vectors $(f_j : 1 \leq j \leq n)$ is orthonormal if and only if $A^T A = I$. Linear algebra (or the argument above, with A replaced by A^T) shows that this happens if and only if A^T and A are inverse to each other, ie. if and only if A is orthogonal. □

Proposition A.4. *If A is an $n \times n$ orthogonal matrix, then:*

(a) $\langle Av, Aw \rangle = \langle v, w \rangle$ for any pair v, w of vectors in $\mathbb{R}^n$.

(b) The linear transformation $v \to Av$ is an isometry of $\mathbb{R}^n$ that takes $\mathbb{B}^n$ onto itself.

Proof. (a) Using successively Proposition A.1 and the definition of orthogonality:

$$\langle Av, Aw \rangle = \langle A^T Av, w \rangle = \langle v, w \rangle.$$

(b) Upon setting $v = w$ in part (a) we obtain

$$\|Av\|^2 = \langle Av, Av \rangle = \langle v, v \rangle = \|v\|^2,$$

so the mapping induced on $\mathbb{R}^n$ by A is an isometry, hence takes the unit ball into itself. Since A^T is also an orthogonal matrix (its transpose is its inverse), it too induces an isometry on $\mathbb{R}^n$, which also takes $\mathbb{B}^n$ into itself. Thus for every $v \in \mathbb{B}^n$ $A^T v \in \mathbb{B}^n$ and $v = A(A^T v)$, hence the isometry induced by A takes $\mathbb{B}^n$ onto itself. □

Proposition A.4 is just half of the real story.

Theorem A.5. *A mapping $T\colon \mathbb{B}^n \to \mathbb{B}^n$ is an isometry if and only if there exists $A \in O(n)$ such that $T(v) = Av$ for every $v \in \mathbb{B}^n$. The matrix A is uniquely determined by T.*

In particular, every such isometry maps the ball *onto* itself, and extends to a linear transformation of the ambient Euclidean space.

Proof. In view of Proposition A.4 We need only show that isometries arise from orthogonal matrices. To this end, let's for the moment allow T to be an isometry taking $\mathbb{B}^n$ into $\mathbb{R}^n$.

(a) *Suppose, in addition, that $T(0) = 0$. Then $\langle T(u), T(v) \rangle = \langle u, v \rangle$ for every pair u, v of vectors in $\mathbb{B}^n$.*

This follows immediately from the relationship between norms of differences and inner products. For $u, v \in \mathbb{B}^n$:

$$\|u - v\|^2 = \langle u - v, u - v \rangle = \|u\|^2 - 2\langle u, v \rangle + \|v\|^2$$

Upon replacing u and v in the above calculation with $T(u)$ and $T(v)$ respectively (and being careful not to assume linearity for T):

$$\begin{aligned} \|T(u) - T(v)\|^2 &= \|T(u)\|^2 - 2\langle T(u), T(v) \rangle + \|T(v)\|^2 \\ &= \|u\|^2 - 2\langle T(u), T(v) \rangle + \|v\|^2, \end{aligned}$$

where the second equality arises from the fact that the distance from the vector 0 to v is the same as that from $0 = T(0)$ to Tv. Similarly the distance from u to v is the same as that from Tu to Tv, so the left-hand sides of both of the equations above are equal, hence so are the right-hand sides, and this yields the desired identity.

(b) *If T is an isometry taking $\mathbb{B}^n$ into $\mathbb{R}^n$ with $T(0) = 0$, then there exists $A \in O(n)$ for which $T(v) = Av$ for every $v \in \mathbb{B}^n$.*

Let $(e_1, e_2, \ldots e_n)$ denote the standard orthonormal basis for $\mathbb{R}^n$. Since each of these vectors belongs to $\mathbb{B}^n$ we can apply T to them, and, since, by (a) above, T preserves inner products the result is another orthonormal basis $(f_1, f_2, \ldots f_n)$ for $\mathbb{R}^n$. Let A be the matrix that has as its j-th column the coefficients of $f_j = T(e_j)$ with respect to the original basis $\{e_j\}$. Then $A \in O(n)$ by Proposition A.3, and $T(e_j) = Ae_j$, hence for every $v \in \mathbb{B}^n$:

$$T(v) = \sum_{j=1}^{n} \langle T(v), f_j \rangle f_j = \sum_{j=1}^{n} \langle T(v), T(e_j) \rangle T(e_j) = \sum_{j=1}^{n} \langle v, e_j \rangle Ae_j = Av$$

as desired.

(c) Suppose again that T is an isometry of $\mathbb{B}^n$ into $\mathbb{R}^n$. Then $T - T(0)$ is an isometry $\mathbb{B}^n \to \mathbb{R}^n$ that fixes the origin, so it follows from part (b) that for some orthogonal matrix A we have $T(v) = Av + T(0)$ for each $v \in \mathbb{R}^n$, where the transformation induced by A takes $\mathbb{B}^n$ onto itself. Thus Proposition A.4 insures that T takes $\mathbb{B}^n$ onto the closed ball of radius 1 centered at $T(0)$, so if $T(\mathbb{B}^n) = \mathbb{B}^n$ then $T(0) = 0$, hence T must be effected by multiplication by the orthogonal matrix A.

That this matrix is unique is clear, since—as we saw in part (b)—its columns are just the images of the standard basis vectors for $\mathbb{R}^n$ under the action of T. This completes our characterization of isometries of $\mathbb{B}^n$. □

BEFORE PROCEEDING let's note that:

If $n \geq 2$ then the group $O(n)$ is not commutative.

Indeed, here are two matrices in $O(2)$ that do not commute:

$$\begin{pmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} \qquad \text{and} \qquad \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

the first of which induces rotation through an angle of 45 degrees, while the second induces reflection about the horizontal axis. To get an example in $O(n)$ for $n > 2$ just put each of the above matrices in the upper right hand corner of an $n \times n$ matrix, and fill in the remaining entries with ones on the main diagonal and zeros off it.

A.2 Rotations

Proposition A.6. *If $A \in SO(2)$ (i.e. orthogonal with determinant 1) then it induces on $\mathbb{R}^2$ a rotation about the origin. If $A \in O(2)$ has determinant -1, then it induces on $\mathbb{R}^2$ a reflection in a line through the origin.*

Proof. Each $A \in O(2)$ takes the pair of unit vectors (e_1, e_2) (respectively along the horizontal and vertical axes) to an orthogonal pair (u, v) of unit vectors, where u is the rotate of e_1 through some angle θ, and v is either the rotate of e_2 through that angle, in which case the determinant of A is 1 and A is the mapping of "rotation by θ," or v is the negative of that vector. In this latter case $\det A = -1$, and A effects the mapping of reflection in the line through the origin parallel to u. □

Proposition A.7. *If $A \in SO(3)$ then the map $x \to Ax$ is a rotation of $\mathbb{R}^3$, with center at the origin.*

We're saying that for each $A \in SO(3)$ the associated linear transformation fixes each point of a line L through the origin, and in any plane P perpendicular to L acts as a rotation with center at the intersection of P and L. This isn't obvious. Clearly the product of two matrices in $SO(3)$ also belongs to $SO(3)$ (multiplicativity of the determinant), but it's not so clear that the result of composing two rotations about different axes has to fix a line through the origin.

Proof. Suppose $A \in SO(3)$. To find the axis of rotation we need to show that $Av = v$ for some unit vector $v \in \mathbb{R}^3$, i.e. that 1 is an eigenvalue of A; equivalently, $\det(A - I) = 0$. For this, note that since $AA^t = I$ we have

$$(A - I)A^t = AA^t - A^t = I - A^t = -(A - I)^t$$

hence, since $\det A = \det A^t = 1$:

$$\begin{aligned}\det(A-I) &= \det(A-I)\det(A^t) = \det[(A-I)A^t] \\ &= \det[-(A-I)^t] = (-1)^3\det(A-I)^t \\ &= -\det(A-I)\end{aligned}$$

so $\det(A-I) = 0$, as desired.

Let $v_1 \in \mathbb{R}^3$ be the unit vector promised by the last paragraph: $Av_1 = v_1$. Let (v_2, v_3) be an orthonormal basis for the subspace E of $\mathbb{R}^3$ orthogonal to v_1. Then (v_1, v_2, v_3) is an orthonormal basis for $\mathbb{R}^3$, relative to which the matrix of the transformation $x \to Ax$ has block diagonal form $\left[\begin{smallmatrix} 1 & 0 \\ 0 & B \end{smallmatrix}\right]$ where B is a 2×2 orthogonal matrix. Thus A and B have the same determinant, so $\det B = 1$, i.e. $B \in SO(2)$, so by the previous proposition B induces on E either the identity map or a rotation (about the origin). □

It's easy to see that the rotation group of the ball does *not* share the commutativity of that of the disc; take, for an example, a pair of 45° rotations about different orthogonal axes. Thus, while the matrix group $SO(2)$ is commutative, $SO(3)$ is not.

ROTATIONS ABOUT THE ORIGIN in 3-space are linear transformations, and linear transformations have matrix representations. Let $R_u(\rho)$ denote the matrix (with respect to the standard basis of $\mathbb{R}^3$) of the transformation of rotation about the origin through angle ρ with axis the unit vector $u \in \mathbb{R}^3$ (the "right-hand rule" determining the positive direction of ρ). Although somewhat complicated, this matrix factors readily as a product of simpler matrices. Here are the three "elementary" rotation matrices; the ones that represent rotations about the coordinate axes:

1. Rotation through angle ρ about the z-axis

$$R_z := \begin{bmatrix} \cos\rho & -\sin\rho & 0 \\ \sin\rho & \cos\rho & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

2. Rotation through angle ρ about the x-axis

$$R_x := \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\rho & -\sin\rho \\ 0 & \sin\rho & \cos\rho \end{bmatrix}$$

3. Rotation through angle ρ about the y-axis

$$R_y := \begin{bmatrix} \cos\rho & 0 & \sin\rho \\ 0 & 1 & 0 \\ -\sin\rho & 0 & \cos\rho \end{bmatrix}$$

Now fix the unit vector $u \in \mathbb{R}^3$ and the angle $\rho \in [-\pi, \pi)$, and let L_u denote the oriented line through the origin and u, in the direction from the origin to u. We're going to understand the transformation $R_u(\rho)$ of rotation about L_u through angle ρ by factoring it into a product of several elementary ones.

Let (φ, θ) be the spherical coordinates of u, i.e.

$$u = [\sin\varphi\cos\theta, \sin\varphi\sin\theta, \cos\varphi]^T$$

where $\varphi \in [0, \pi]$ is the angle between u and the z-axis, and $\theta \in [-\pi, \pi]$ is the angle between the x-axis and the projection of the u into the x, y-plane. Let $T = R_y(-\varphi)R_z(-\theta)$, so that T rotates u through angle $-\theta$ about the z-axis, depositing it into the x, z-plane, then in that plane (i.e. about the y-axis) rotates the resulting vector through angle ρ so that it ends up at the "north pole" $e_3 := [0, 0, 1]^T$. Thus $T^{-1}R_z(\rho)T$ fixes u and, since T belongs to $SO(3)$, and therefore preserves both distances and angles, it rotates points of $\mathbb{R}^3$ about L_u through angle ρ, i.e. it's none other than $R_u(\rho)$. Explicitly:

$$R_u(\rho) = R_z(\theta)R_y(\varphi)R_z(\rho)R_y(-\varphi)R_z(-\theta). \tag{5}$$

To FIND THE MATRIX OF $R_u(\rho)$ (with respect to the standard unit vector basis of $\mathbb{R}^3$) one "need only" multiply the elementary matrices for the five transformations on the right-hand side of (5). This is best done with your favorite computer-algebra program; the result is nevertheless quite a mess. To bring it into some kind of reasonable form it helps to invert the spherical coordinate representation of u, noting that $\cos\varphi = z$, $\sin\varphi = \sqrt{x^2 + y^2} = \sqrt{1 - z^2}$ (non-negative square root because $0 \le \varphi \le \pi$), $\cos\theta = x/\sqrt{1 - z^2}$, and $\sin\theta = y/\sqrt{1 - z^2}$. Again with the help of your computer-algebra program, most likely aided by some paper and pencil algebraic simplifications, there will result the following matrix representation of $R_u(\rho)$:

$$\begin{bmatrix} x^2 + (1 - x^2)\cos\rho & xy(1 - \cos\rho) - z\sin\rho & xz(1 - \cos\rho) + y\sin\rho \\ xy(1 - \cos\rho) + z\sin\rho & y^2 + (1 - y^2)\cos\rho & yz(1 - \cos\rho) - x\sin\rho \\ xz(1 - \cos\rho) - y\sin\rho & yz(1 - \cos\rho) + x\sin\rho & z^2 + (1 - z^2)\cos\rho \end{bmatrix}.$$

Fariborz Maseeh Department of Mathematics & Statistics
Portland State University
PO Box 751, Portland, OR 97207-0751
joels314(at)gmail.com
www.joelshapiro.org