2. The Borsuk–Ulam Theorem

The Borsuk–Ulam theorem is one of the most useful tools offered by elementary algebraic topology to the outside world. Here are four reasons why this is such a great theorem: There are

(1) several different equivalent versions,
(2) many different proofs,
(3) a host of extensions and generalizations, and
(4) numerous interesting applications.

As for (1), Borsuk's original paper [Bor33] already gives three variants. Below we state six different but equivalent versions, all of them very useful, and several more are given in the exercises.

As for (2), there are several proofs of the Borsuk–Ulam theorem that can be labeled as completely elementary, requiring only undergraduate mathematics and no algebraic topology. On the other hand, most of the textbooks on algebraic topology, even the friendliest ones, usually place a proof of the Borsuk–Ulam theorem well beyond page 100. Some of them use just basic homology theory, others rely on properties of the cohomology ring, but in any case, significant apparatus has to be mastered for really understanding such proofs. From a "higher" point of view, it can be argued that these proofs are more conceptual and go to the heart of the matter, and thus they are preferable to the "ad hoc" elementary proofs. But this point of view can be appreciated only by someone for whom the necessary machinery is as natural as breathing.[1] Since not everyone, especially in combinatorics and computer science, belongs to this lucky group, we present some "old-fashioned" elementary proofs. The one in Section 2.2, called a *homotopy extension argument*, is geometric and very intuitive. In Section 2.3 we introduce Tucker's lemma, a combinatorial statement equivalent to the Borsuk–Ulam theorem, and we give a purely combinatorial proof. (This resembles the well-known proof of Brouwer's theorem via the Sperner lemma, but Tucker's lemma is

[1] Borsuk's footnote from [Bor33]: "Mr. H. Hopf, whom I informed about Theorem I, noted for me in a letter three other shorter proofs of this theorem. But since these proofs are founded on deep results in the theory of the mapping degree and my proof is in essence completely elementary, I think that its publication is not superfluous. [...]"

more demanding.) Next, in Section 2.4, we prove Tucker's lemma differently, introducing some of the most elementary notions of simplicial homology.

As for (3), we will examine various generalizations and strengthenings later; much more can be found in Steinlein's surveys [Ste85], [Ste93] and in the sources he quotes.

Finally, as for applications (4), just wait and see.

2.1 The Borsuk–Ulam Theorem in Various Guises

One of the versions of the Borsuk–Ulam theorem, the one that is perhaps the easiest to remember, states that *for every continuous mapping* $f\colon S^n \to \mathbb{R}^n$, *there exists a point* $\boldsymbol{x} \in S^n$ *such that* $f(\boldsymbol{x}) = f(-\boldsymbol{x})$. Here is an illustration for $n = 2$. Take a rubber ball, deflate and crumble it, and lay it flat:

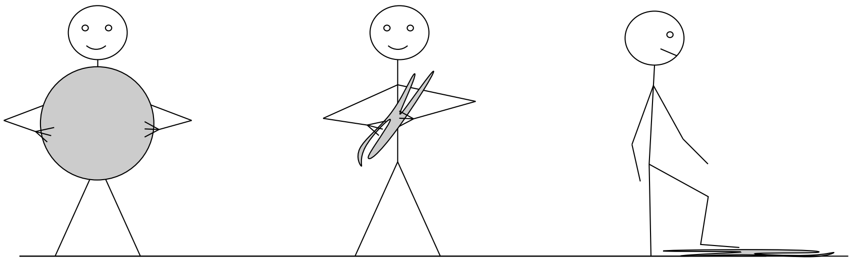

Then there are two points on the surface of the ball that were diametrically opposite (antipodal) and now are lying on top of one another!

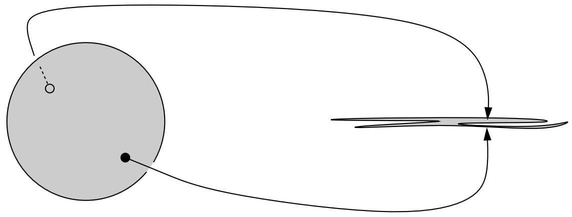

Another popular interpretation, found in almost every textbook, says that at any given time there are two antipodal places on Earth that have the same temperature and, at the same time, identical air pressure (here $n = 2$).[2]

It is instructive to compare this with the Brouwer fixed point theorem, which says that every continuous mapping $f\colon B^n \to B^n$ has a fixed point: $f(\boldsymbol{x}) = \boldsymbol{x}$ for some $\boldsymbol{x} \in B^n$. The statement of the Borsuk–Ulam theorem sounds similar (and actually, it easily implies the Brouwer theorem; see below). But it involves an extra ingredient besides the topology of the considered

[2] Although anyone who has ever touched a griddle-hot stove knows that the temperature need not be continuous.

spaces: a certain *symmetry* of these spaces, namely, the symmetry given by the mapping $\boldsymbol{x} \mapsto -\boldsymbol{x}$ (which is often called the *antipodality* on S^n and on $\mathbb{R}^n$).

Here are Borsuk's original formulations of the Borsuk–Ulam theorem:

Der Zweck dieser Arbeit ist, folgende drei Sätze zu beweisen:

Satz I [6]). *Jede antipodentreue Abbildung von S_n ist wesentlich.*

Satz II [7]). *Ist $f \in R^{n^{S_n}}$ (d. h. bildet f die Sphäre S_n auf einen Teil von R^n ab), so gibt es einen derartigen Punkt $p \in S_n$, dass $f(p) = f(p^*)$ ist.*

Satz III. *Sind $A_0, A_1, \ldots, A_n$ in sich kompakte Mengen von denen keine zwei antipodische Punkte der Sphäre S_n enthält, so enthält die Summe $\sum_{i=0}^{n} A_i$ die Sphäre S_n nicht.*

Here are the promised many equivalent versions, in English.

2.1.1 Theorem (Borsuk–Ulam theorem). *For every $n \geq 0$, the following statements are equivalent, and true:*

(BU1a) (Borsuk [Bor33, Satz II][3]) *For every continuous mapping $f\colon S^n \to \mathbb{R}^n$ there exists a point $\boldsymbol{x} \in S^n$ with $f(\boldsymbol{x}) = f(-\boldsymbol{x})$.*

(BU1b) *For every* **antipodal mapping** *$f\colon S^n \to \mathbb{R}^n$ (that is, f is continuous and $f(-\boldsymbol{x}) = -f(\boldsymbol{x})$ for all $\boldsymbol{x} \in S^n$) there exists a point $\boldsymbol{x} \in S^n$ satisfying $f(\boldsymbol{x}) = \boldsymbol{0}$.*

(BU2a) *There is no antipodal mapping $f\colon S^n \to S^{n-1}$.*

(BU2b) *There is no continuous mapping $f\colon B^n \to S^{n-1}$ that is antipodal on the boundary, i.e., satisfies $f(-\boldsymbol{x}) = -f(\boldsymbol{x})$ for all $\boldsymbol{x} \in S^{n-1} = \partial B^n$.*

(LS-c) (Lyusternik and Shnirel'man [LS30], Borsuk [Bor33, Satz III]) *For any cover $F_1, \ldots, F_{n+1}$ of the sphere S^n by $n+1$ closed sets, there is at least one set containing a pair of antipodal points (that is, $F_i \cap (-F_i) \neq \emptyset$).*

(LS-o) *For any cover $U_1, \ldots, U_{n+1}$ of the sphere S^n by $n+1$ open sets, there is at least one set containing a pair of antipodal points.*

While proving any of the versions of the Borsuk–Ulam theorem is not easy, at least without some technical apparatus, checking the equivalence of all the statements is not so hard. Deriving at least some of the equivalences before reading further is a very good way of getting a feeling for the theorem.

[3] Borsuk's footnote at this theorem reads: "This theorem was posed as a conjecture by St. Ulam."

Equivalence of (BU1a), (BU1b), and (BU2a).
(BU1a) $\Longrightarrow$ (BU1b) is clear.

(BU1b) $\Longrightarrow$ (BU1a) We apply (BU1b) to the antipodal mapping given by $g(\boldsymbol{x}) := f(\boldsymbol{x}) - f(-\boldsymbol{x})$.

(BU1b) $\Longrightarrow$ (BU2a) An antipodal mapping $S^n \to S^{n-1}$ is also a nowhere zero antipodal mapping $S^n \to \mathbb{R}^n$.

(BU2a) $\Longrightarrow$ (BU1b) Assume that $f: S^n \to \mathbb{R}^n$ is a continuous nowhere zero antipodal mapping. Then the antipodal mapping $g: S^n \to S^{n-1}$ given by $g(\boldsymbol{x}) := f(\boldsymbol{x})/\|f(\boldsymbol{x})\|$ contradicts (BU2a).

Equivalence of (BU2a) with (BU2b). This is easy once we observe that the projection $\pi: (x_1, \ldots, x_{n+1}) \mapsto (x_1, \ldots, x_n)$ is a homeomorphism of the upper hemisphere U of S^n with B^n:

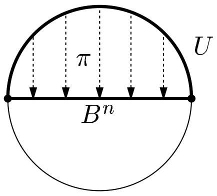

An antipodal mapping $f: S^n \to S^{n-1}$ as in (BU2a) would yield a mapping $g: B^n \to S^{n-1}$ antipodal on ∂B^n by $g(\boldsymbol{x}) = f(\pi^{-1}(\boldsymbol{x}))$.

Conversely, for $g: B^n \to S^{n-1}$ as in (BU2b) we can define $f(\boldsymbol{x}) = g(\pi(\boldsymbol{x}))$ and $f(-\boldsymbol{x}) = -g(\pi(\boldsymbol{x}))$ for $\boldsymbol{x} \in U$. This specifies f on the whole of S^n; it is consistent because g is antipodal on the equator of S^n; and the resulting f is continuous, since it is continuous on both of the closed hemispheres (see Exercise 1.1.2).

Equivalence with (LS-c), (LS-o).
(BU1a) $\Longrightarrow$ (LS-c) For a closed cover $F_1, \ldots, F_{n+1}$ we define a continuous mapping $f: S^n \to \mathbb{R}^n$ by $f(\boldsymbol{x}) := (\operatorname{dist}(\boldsymbol{x}, F_1), \ldots, \operatorname{dist}(\boldsymbol{x}, F_n))$, and we consider a point $\boldsymbol{x} \in S^n$ with $f(\boldsymbol{x}) = f(-\boldsymbol{x}) = \boldsymbol{y}$, which exists by (BU1a). If the ith coordinate of the point $\boldsymbol{y}$ is 0, then both $\boldsymbol{x}$ and $-\boldsymbol{x}$ are in F_i. If all coordinates of $\boldsymbol{y}$ are nonzero, then both $\boldsymbol{x}$ and $-\boldsymbol{x}$ lie in F_{n+1}.

(LS-c) $\Longrightarrow$ (BU2a) We need an *auxiliary result*: There exists a covering of S^{n-1} by closed sets $F_1, \ldots, F_{n+1}$ such that no F_i contains a pair of antipodal points (to see this, we consider an n-simplex in $\mathbb{R}^n$ containing $\mathbf{0}$ in its interior, and we project the facets centrally from $\mathbf{0}$ on S^{n-1}). Then if a continuous antipodal mapping $f: S^n \to S^{n-1}$ existed, the sets $f^{-1}(F_1), \ldots, f^{-1}(F_{n+1})$ would contradict (LS-c).

(LS-c) $\Longrightarrow$ (LS-o) follows from the fact that for every open cover $U_1, \ldots, U_{n+1}$ there exists a closed cover $F_1, \ldots, F_{n+1}$ satisfying $F_i \subset U_i$ for $i = 1, \ldots, n+1$:

For each point $\boldsymbol{x}$ of the sphere we choose an open neighborhood $V_{\boldsymbol{x}}$ whose closure is contained in some U_i, and apply the compactness of the sphere.

(LS-o) $\Longrightarrow$ (LS-c) Given a closed cover $F_1, \ldots, F_{n+1}$, we wrap each F_i in the open set $U_i^{\varepsilon} := \{\boldsymbol{x} \in S^n : \operatorname{dist}(\boldsymbol{x}, F_i) < \varepsilon\}$. We let $\varepsilon \to 0$ and we use the compactness of the sphere. We first obtain an infinite sequence of points $\boldsymbol{x}^0, \boldsymbol{x}^1, \boldsymbol{x}^2, \ldots$ in S^n with $\lim_{j\to\infty} \operatorname{dist}(\boldsymbol{x}^j, F_i) = \lim_{j\to\infty} \operatorname{dist}(-\boldsymbol{x}^j, F_i) = 0$ for some *fixed* i. Then we select a *convergent* subsequence of the $\boldsymbol{x}^j$. The limit of this sequence is in F_i, since F_i is closed, and it provides the required antipodal pair in F_i.

Proof of the Brouwer fixed point theorem from (BU2b). Suppose that $f: B^n \to B^n$ is continuous and has no fixed point. By a well-known construction, we show the existence of a continuous map $g: B^n \to S^{n-1}$ whose restriction to S^{n-1} is the identity map (such a g is called a *retraction* of B^n to S^{n-1}). We define $g(\boldsymbol{x})$ as the point in which the ray originating in $f(\boldsymbol{x})$ and going through $\boldsymbol{x}$ intersects S^{n-1}. This g contradicts (BU2b).

Notes. The earliest reference for what is now commonly called the Borsuk–Ulam theorem is probably Lyusternik and Shnirel'man [LS30] from 1930 (the covering version (LS-c)). Borsuk's paper [Bor33] is from 1933. The only written reference concerning Ulam's role in the matter seems to be Borsuk's footnote quoted above. Since then, hundreds of papers with various new proofs, variations of old proofs, generalizations, and applications have appeared; the most comprehensive survey known to me, Steinlein [Ste85] from 1985, lists nearly 500 items in the bibliography.

Types of proofs. In the numerous published proofs of the Borsuk–Ulam theorem, one can distinguish several basic approaches (as is done in [Ste85]). Some of these types will be treated in this book; for the others, we outline the main ideas here and give references, mostly to recent textbooks.

Degree-theoretic proofs are discussed in Section 2.4, and another such proof is outlined in the notes to Section 6.2. A related method uses the *Lefschetz number*; such a proof of a result generalizing the Borsuk–Ulam theorem is given in Section 6.2. A proof using rudimentary *Smith theory* can be found in [Bre93, Section 20].

A *proof using the cohomology ring* considers the map $g: \mathbb{R}\mathrm{P}^n \to \mathbb{R}\mathrm{P}^m$ induced by an antipodal $f: S^n \to S^m$, and shows that the corresponding homomorphism $g^*: H^*(\mathbb{R}\mathrm{P}^m, \mathbb{Z}_2) \to H^*(\mathbb{R}\mathrm{P}^n, \mathbb{Z}_2)$ of the cohomology rings carries a generator α of $H^1(\mathbb{R}\mathrm{P}^m, \mathbb{Z}_2)$ to a generator β of $H^1(\mathbb{R}\mathrm{P}^n, \mathbb{Z}_2)$. This is impossible if $m+1 \leq n$, since then α^{m+1} is trivial, while β^n is nontrivial. See, for example, [Mun84, p. 403] or [Bre93, p. 362].

A proof by a *homotopy extension argument* will be discussed in Section 2.2, and a representative of the family of *combinatorial proofs* in Section 2.3. An *algebraic proof* in [Kne82] establishes the theorem for polynomial mappings, and the general form follows by an approximation argument (for another algebraic proof see [AP83]).

The fact that the Borsuk–Ulam theorem implies Brouwer's fixed point theorem seems to be folklore; also see Su [Su97] for an alternative proof.

As for *applications* of the Borsuk–Ulam theorem, we will cover some in the subsequent sections. For a multitude of others, we refer to the surveys [Ste85], [Ste93]. The papers [Bár93] and [Alo88] give nice overviews of combinatorial applications; most of these are included in this book.

Many applications appear in existence results for solutions of nonlinear *partial differential equations* and *integral equations*; we will neglect this broad field entirely (see [KZ75], [Ste85], [Ste93]). Borsuk–Ulam-type results also play an important role in *functional analysis* and in the *geometry of Banach spaces*. A neat *algebraic application* will be outlined in the notes to Section 5.3.

A beautiful combinatorial application of the Borsuk–Ulam theorem, which we will not discuss in detail and whose original account is very nicely readable, concerns *linkless embeddings* of graphs in $\mathbb{R}^3$. Any finite graph G, regarded as a 1-dimensional finite simplicial complex, can be realized in $\mathbb{R}^3$. Such a realization is called *linkless* if any two vertex-disjoint circuits in G form two unlinked closed curves in the realization. Here two curves $\alpha, \beta \subset \mathbb{R}^3$ (each homeomorphic to S^1) are *unlinked* if they are equivalent to two isometric copies α', β' of S^1 in $\mathbb{R}^3$ lying far from one another, and the equivalence means that there is a homeomorphism $\varphi\colon \mathbb{R}^3 \to \mathbb{R}^3$ such that $\varphi(\alpha \cup \beta) = \alpha' \cup \beta'$ (these are notions from knot theory; see, e.g., Rolfsen [Rol90] for more information).

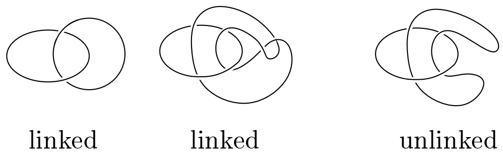

Lovász and Schrijver [LS98], building on previous work by Robertson, Seymour, and Thomas, proved that graphs possessing a linkless embedding into $\mathbb{R}^3$ are exactly those for which a numerical parameter μ, called the *Colin de Verdière number*, is at most 4. The definition of this parameter, using spectra of certain matrices, is not very intuitive at first sight (and we do not reproduce it; see [LS98] or other sources). The graph-theoretic significance of the Colin de Verdière number looks almost miraculous: Besides the incredible result about linkless embeddings, it is also known that $\mu(G) \leq 1$ iff G is a disjoint union of paths, $\mu(G) \leq 2$ iff G is outerplanar, and $\mu(G) \leq 3$ iff G is planar. In the Lovász–Schrijver proof, the Borsuk–Ulam theorem is used for establishing the following: Given any "generic" embedding of the 1-skeleton of a 5-dimensional convex polytope P into $\mathbb{R}^3$, there are two antipodal 2-dimensional faces F_1, F_2 of P (here "antipodal" means $F_1 = P \cap h_1$ and $F_2 = P \cap h_2$ for some parallel hyperplanes h_1, h_2) such that the images of the boundaries of F_1 and F_2 are linked (in fact, they have a nonzero linking number, which is stronger than being linked; the curves in the left picture above satisfy this, while those in the middle picture do not). Thus, for example, the complete graph K_6 is not linklessly embeddable. (More generally, a generic embedding of the $(d-1)$-skeleton of a $(2d+1)$-polytope into $\mathbb{R}^d$ links the boundaries of two antipodal d-faces.)

The paper [Bor33] containing the Borsuk–Ulam theorem also states the so-called *Borsuk's conjecture.* The Lyusternik–Shnirel'man theorem (about covering S^n by $n+1$ closed sets) can be restated as follows: For every closed cover of S^{n-1} by at most n sets, one of the sets has diameter 2, i.e., the same as the diameter of S^{n-1} itself. On the other hand, there are $n+1$ sets of diameter < 2 covering S^{n-1}. Borsuk asked whether any bounded set $X \subset \mathbb{R}^n$ can be split into $n+1$ parts, each having diameter strictly smaller than X. This was resolved in the negative by Kahn and Kalai [KK93]. Their spectacular combinatorial proof has made Borsuk's conjecture quite popular in recent years ([Nil94] is a two-page exposition, and the proof has been reproduced in several books, such as [AZ00]). On the other hand, Borsuk's conjecture holds for all *smooth* convex bodies, as was proved by Hadwiger [Had45], [Had46].

Kakutani-type theorems. Kakutani [Kak43] proved that for any compact convex set in $\mathbb{R}^3$ there exists a cube circumscribed about it and touching it with all 6 facets. This is an easy consequence of the following: For any continuous $f\colon S^2 \to \mathbb{R}$, there are 3 mutually perpendicular vectors $\boldsymbol{x}_1, \boldsymbol{x}_2, \boldsymbol{x}_3 \in S^2$ with $f(\boldsymbol{x}_1) = f(\boldsymbol{x}_2) = f(\boldsymbol{x}_3)$. This was generalized to dimension n (with $n+1$ mutually orthogonal vectors) by Yamabe and Yujobô [YY50], and rederived by Yang [Yan54] (in a greater generality, with a suitable abstract notion of "orthogonality"). Yang [Yan54] and Bourgin [Bou63] proved that for any continuous $f\colon S^n \to \mathbb{R}$, there are n mutually orthogonal $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n \in S^n$ with $f(\boldsymbol{x}_1) = f(-\boldsymbol{x}_1) = f(\boldsymbol{x}_2) = \cdots = f(-\boldsymbol{x}_n)$, generalizing such a result for S^2 due to Dyson [Dys51]. Here is another nice result of Yang of this type: If $f\colon S^{mn+m+n} \to \mathbb{R}^m$ is continuous, then there exists an antipodally symmetric subset of S^{mn+m+n} of dimension at least n on which f is constant. Numerous results about circumscribed geometric shapes and similar problems can be found in works of Makeev, such as [Mak96].

In this connection, we should also mention a conjecture of Knaster [Kna47], stating that for any continuous $f: S^n \to \mathbb{R}^m$ and any configuration $K \subset S^n$ of $n-m+2$ points, there exists a rotation ρ of S^n such that $f(\rho(K))$ is a single point. Although this was proved for some special configurations (for example, Hopf proved the case $m = n$ in 1944, which motivated Knaster's conjecture from 1947), the general conjecture does not hold. It was first refuted by Makeev [Mak84], stronger counterexamples were given by Babenko and Bogatyĭ [BB89], and then Chen [Che98] showed that Knaster's conjecture fails for every $n > m > 2$. Just before this book went to print, Kashin and Szarek announced a counterexample to an interesting special case of Knaster's conjecture, with $m = 1$, n sufficiently large, and K consisting of $n+1$ *linearly independent* unit vectors in $\mathbb{R}^{n+1}$. (All the previous counterexamples used configurations with linear dependencies; also note that if K is the standard orthonormal basis in $\mathbb{R}^{n+1}$, then the conjecture holds by the Yamabe–Yujobô theorem cited above).

A few of the numerous *generalizations of the Borsuk–Ulam theorem* will be discussed later. Here we mention a couple of others, which seem potentially useful for combinatorial and geometric problems.

Bourgin–Yang-type theorems are generalizations of the Borsuk–Ulam theorem of the following sort. For any continuous map $f: S^n \to \mathbb{R}^m$, the *coincidence set* $\{\boldsymbol{x} \in S^n : f(\boldsymbol{x}) = f(-\boldsymbol{x})\}$ has to be not only nonempty (as Borsuk–Ulam asserts), but even "large" if $m < n$. For example, it has dimension at least $n-m$; see [Yan54], [Bou55].

Zero sections of vector bundles. This kind of generalization is technically beyond our scope, but here we at least state a particular case (appearing in Dol'nikov [Dol'92] and, implicitly, Živaljević and Vrećica [ŽV90]; also see Fadell and Husseini [FH88]). Let $G_k(\mathbb{R}^n)$ denote the space of all k-dimensional linear subspaces of $\mathbb{R}^n$ (the *Grassmann manifold*). The natural topology on $G_k(\mathbb{R}^n)$ can be defined using a metric, for example, by saying that two k-dimensional subspaces L and L' have distance at most ε if they possess orthonormal bases $v_1, v_2, \ldots, v_k$ and $v_1', v_2', \ldots, v_k'$, respectively, such that $\|v_i - v_i'\| \leq \varepsilon$ for all $i = 1, 2, \ldots, k$. The theorem asserts that if $f_1, f_2, \ldots, f_{n-k}: G_k(\mathbb{R}^n) \to \mathbb{R}^n$ are continuous maps with $f_i(L) \in L$ for all $L \in G_k(\mathbb{R}^n)$ and all $i = 1, 2, \ldots, n-k$ (in other words, the f_i are sections of the tautological vector bundle over $G_k(\mathbb{R}^n)$), then there is a k-dimensional subspace $L \in G_k(\mathbb{R}^n)$ with $f_1(L) = f_2(L) = \cdots = f_{n-k}(L) = \mathbf{0}$.

Exercises

1. Show that the antipodality assumption in (BU2a) can be replaced by "$f(-\boldsymbol{x}) \neq f(\boldsymbol{x})$ for all $\boldsymbol{x} \in S^n$."

2. Show that the following statement is equivalent to the Borsuk–Ulam theorem: *Let* $f: B^n \to \mathbb{R}^n$ *be a continuous mapping that satisfies* $f(-\boldsymbol{x}) = -f(\boldsymbol{x})$ *for all* $\boldsymbol{x} \in S^{n-1}$*; that is, it is antipodal on the boundary. Then there is a point* $\boldsymbol{x} \in B^n$ *with* $f(\boldsymbol{x}) = \mathbf{0}$.
3.* (A "homotopy" version of the Borsuk–Ulam theorem)
(a) Derive the statement in Exercise 2 (and thus the Borsuk–Ulam theorem) from the following statement ([Bor33, Satz I]): *An antipodal mapping* $f: S^n \to S^n$ *cannot be nullhomotopic.*
(b) Show that the statement in (a) is also implied by the Borsuk–Ulam theorem.
4. (Another "homotopy" version of the Borsuk–Ulam theorem) Prove that the following statement is equivalent to the statement in Exercise 3(a): *If* $f: S^n \to S^n$ *is antipodal, then every mapping* $g: S^n \to S^n$ *that is homotopic to* f *is surjective (i.e., onto).*
5.* Prove that the validity of (any of) the statements in the Theorem 2.1.1 for n implies the validity of all the statements for $n-1$.
6.* (Generalized Lyusternik–Shnirel'man theorem [Gre02]) Derive the following common generalization of (LS-c) and (LS-o): Whenever S^n is covered by $n+1$ sets $A_1, A_2, \ldots, A_{n+1}$, each A_i open or closed, there is an i such that $A_i \cap (-A_i) \neq \emptyset$.
7. Does the Lyusternik–Shnirel'man theorem remain valid for coverings of S^n by $n+1$ sets, each of which can be obtained from open sets by finitely many set-theoretic operations (union, intersection, difference)?
8. Describe a surjective nullhomotopic map $S^n \to S^1$ (at least for $n = 1$ and $n = 2$).
9. (Borsuk graph) For a positive real number $\alpha < 2$, let $B(n+1, \alpha)$ be the (infinite) *Borsuk graph* with S^n as the vertex set and with two points connected by an edge iff their distance is at least α. Prove that the Borsuk–Ulam theorem is equivalent to the following statement: *For every* $\alpha < 2$*, we have* $\chi(B(n+1, \alpha)) \geq n+2$ (here χ denotes the usual chromatic number).
10. Let the torus be represented as $T = S^1 \times S^1$.
(a) Show that an analogue of (BU1a) for maps $T \to \mathbb{R}^2$ (formulate it!) is false.
(b) Show that it works for maps $T \to \mathbb{R}^1$.
11.* (a) Let $A_1, A_2, \ldots, A_n$ be closed subsets of S^n with $A_i \cap (-A_i) = \emptyset$. Prove, using the Borsuk–Ulam theorem, that $\bigcup_{i=1}^n (A_i \cup (-A_i)) \neq S^n$.
(b) Derive the Borsuk–Ulam theorem from the statement in (a).
12.* Consider the Borsuk–Ulam-type theorem for Grassmann manifolds stated at the end of the notes of this section.
(a) Show that the case $k = 1$ (with $n-1$ continuous maps, each assigning to each line through the origin in $\mathbb{R}^n$ a point on that line) is equivalent to the Borsuk–Ulam theorem.

(b) Prove that the case $k = n-1$ (a continuous map assigning to each hyperplane through the origin a point in that hyperplane) is equivalent to the Borsuk–Ulam theorem as well.

2.2 A Geometric Proof

We prove the version (BU1b) of the Borsuk–Ulam theorem. Let $f: S^n \to \mathbb{R}^n$ be a continuous antipodal map. We want to prove that it has a zero. First we explain the idea of the proof, assuming that f is "sufficiently generic," without making the meaning of this quite precise. Then we supply a rigorous argument, involving a suitable perturbation of f.

The intuition. Let $g: S^n \to \mathbb{R}^n$ denote the "north–south projection" map; if $S^n = \{\boldsymbol{x} \in \mathbb{R}^{n+1} : x_1^2 + \cdots + x_{n+1}^2 = 1\}$, then g is given by $g(\boldsymbol{x}) = (x_1, x_2, \ldots, x_n)$. This g has exactly two zeros, namely, the north pole and the south pole: $\boldsymbol{n} = (0, 0, \ldots, 0, 1)$, $\boldsymbol{s} = (0, 0, \ldots, 0, -1)$. (The important feature of g is that, obviously, it has a finite number of zeros; more precisely, the number of zeros is twice an odd number.)

We consider the $(n+1)$-dimensional space $X := S^n \times [0, 1]$ (a "hollow cylinder") and the mapping $F: X \to \mathbb{R}^n$ given by $F(\boldsymbol{x}, t) := (1-t)g(\boldsymbol{x}) + tf(\boldsymbol{x})$. Geometrically, we take two copies of S^n (we can think of them as placed in $\mathbb{R}^{n+2}$), one of them with the mapping g and the other one with f. We connect the corresponding points of these two spheres by segments, and the mapping F is defined on each segment by linear interpolation. For $n = 1$, we get a cylinder as in the picture:

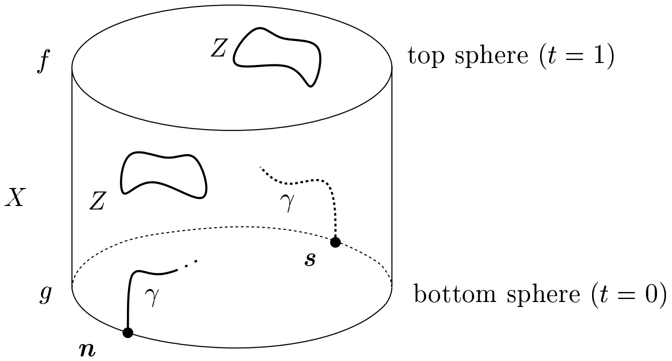

The antipodality $\boldsymbol{x} \mapsto -\boldsymbol{x}$ on S^n is extended to the map ν on X by $\nu: (\boldsymbol{x}, t) \mapsto (-\boldsymbol{x}, t)$ (note that t is unchanged). We will call ν the *antipodality on* X.

We note that F is antipodal with respect to ν; that is, $F(\nu(\boldsymbol{x}, t)) = -F(\boldsymbol{x}, t)$.

For contradiction, let us suppose that f has no zeros. We investigate the zero set $Z := F^{-1}(\boldsymbol{0})$. If f is sufficiently generic, then Z is a 1-dimensional compact manifold, and therefore, its components are cycles and paths (this is the part to be made precise later). Moreover, the endpoints of the paths lie on the bottom or top copy of S^n ($t = 0$ or $t = 1$) and are zeros of f or g, while the cycles do not reach into the top and bottom spheres.[4]

Assuming that f has no zeros and knowing that g has only the two zeros at the poles, we see that there must be a single path γ connecting $\boldsymbol{n}$ to $\boldsymbol{s}$. But at the same time, the set Z is invariant under ν. If we follow γ from $\boldsymbol{n}$ on, the other part starting from $\boldsymbol{s}$ must behave symmetrically. But then it is easy to see that the two ends cannot meet: A symmetric path from $\boldsymbol{n}$ to $\boldsymbol{s}$ does not exist in X. We have reached a contradiction.

Note that the argument actually shows that the number of zeros of a "generic" antipodal map is twice an odd number. Indeed, the zeros of f on the top sphere are paired up by paths in Z, except for two that are connected to the zeros of g on the bottom sphere.

The real thing. A rigorous proof follows the same ideas but uses a suitable small perturbation of f. Recall that the ℓ_1-norm of a point $\boldsymbol{x} \in \mathbb{R}^n$ is $\|\boldsymbol{x}\|_1 = \sum_{i=1}^n |x_i|$. Let $\hat{S}^n = \{\boldsymbol{x} \in \mathbb{R}^{n+1} : \|\boldsymbol{x}\|_1 = 1\}$ denote the unit sphere of the ℓ_1-norm. This is the boundary of a crosspolytope (Definition 1.4.1); for example, $\hat{S}^2$ is the surface of a regular octahedron. This $\hat{S}^n$ is homeomorphic to S^n, and we will consider $\hat{S}^n$ instead of S^n in the rest of the proof. The space $X := \hat{S}^n \times [0,1]$ is a union of finitely many convex polytopes (simplicial prisms). Let us call $\hat{S}^n \times \{0\}$ the *bottom sphere* and $\hat{S}^n \times \{1\}$ the *top sphere* in X.

We choose a sufficiently fine finite triangulation T of X (just how fine will be specified later) that respects the symmetry of X given by ν, in the following sense: Each simplex $\sigma \in \mathsf{T}$ is mapped bijectively onto the "opposite" simplex $\nu(\sigma) \in \mathsf{T}$, and $\sigma \cap \nu(\sigma) = \varnothing$. Moreover, the triangulation T contains triangulations T_t and T_b of the top and bottom spheres, respectively, as subcomplexes, and T_t and T_b each refine the natural triangulation of $\hat{S}^n$.

We let the mapping g be an orthogonal projection of $\hat{S}^n$ into $\mathbb{R}^n$, but not in a coordinate direction, but rather in a "generic" direction, such that the two zeros $\boldsymbol{n}$ and $\boldsymbol{s}$ of g lie in the interior of n-dimensional simplices of the triangulation T_b, as is indicated in the drawing (where $n = 2$):

[4] To gain some intuition as to why this is the case, one may think of the case $n = 1$, and unroll X to obtain a rectangle R in the plane. Then F is a real function on R, its graph is a "terrain" over R, and Z is the "zero contour." As people familiar with topographic maps will know, a typical contour on a smooth terrain is a smooth curve consisting of disjoint cycles and curve segments with both ends on the boundary of R. Other possible cases, such as two cycles meeting at a point (saddle), are exceptional, and they disappear by an arbitrarily small perturbation.

Imagining the higher-dimensional cases is more demanding. Readers knowing the implicit function theorem from analysis may want to contemplate what that theorem gives in the considered situation.

Anyway, we will soon provide a proof using a piecewise linear approximation.

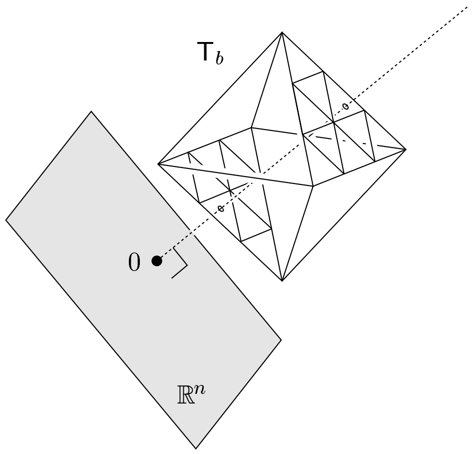

We again suppose that $f: \hat{S}^n \to \mathbb{R}^n$ has no zeros. By compactness, there is an $\varepsilon > 0$ such that $\|f(\boldsymbol{x})\| \geq \varepsilon$ for all $\boldsymbol{x} \in \hat{S}^n$. As in the informal outline, let $F(\boldsymbol{x}, t) := (1-t)g(\boldsymbol{x}) + tf(\boldsymbol{x})$, let T be a fine triangulation of X as above, and let $\bar{F}: X \to \mathbb{R}^n$ be the map that agrees with F on the vertex set $V(\mathsf{T})$ of T and is affine on each simplex of T (similar to Definition 1.5.3 of the affine extension of a simplicial map). Since F is uniformly continuous, we can assume that $\|F(\boldsymbol{y}) - \bar{F}(\boldsymbol{y})\| \leq \frac{\varepsilon}{2}$ for all $\boldsymbol{y} \in X$, provided that T is sufficiently fine. Thus,

$$\bar{F} \text{ has no zeros on the top sphere.} \tag{2.1}$$

Since our g is already affine, $\bar{F}$ coincides with g on the bottom sphere, and we have

$\bar{F}$ has exactly two zeros on the bottom sphere, lying in the interiors of n-dimensional (antipodal) simplices of T_b. (2.2)

Further, let $\tilde{F}$ be a mapping arising by a sufficiently small antipodal perturbation of $\bar{F}$. Namely, we choose a suitable map $P_0: V(\mathsf{T}) \to \mathbb{R}^n$ satisfying $P_0(\nu(\boldsymbol{v})) = -P_0(\boldsymbol{v})$ for each $\boldsymbol{v} \in V(\mathsf{T})$. Further properties required of P_0 will be specified later. We extend P_0 affinely on each simplex of T, obtaining a map $P: X \to \mathbb{R}^n$, and we set $\tilde{F} = \bar{F} + P$. We note that if all values of P_0 lie sufficiently close to $\mathbf{0}$, then the perturbed map $\tilde{F}$ still has the two properties (2.1) and (2.2).

Next, we introduce generic maps on T. We begin by noting than if $h: \mathbb{R}^{n+1} \to \mathbb{R}^n$ is an affine map, then $h^{-1}(\mathbf{0})$ either is empty, or it is an affine subspace of dimension at least 1. Now let σ be an $(n+1)$-dimensional simplex and h an affine map $\sigma \to \mathbb{R}^n$. We say that h is *generic* if $h^{-1}(\mathbf{0})$ intersects no face of σ of dimension smaller than n. In such case, $h^{-1}(\mathbf{0})$ either is empty, or it is a segment lying in the interior of σ, with endpoints lying in the interior of two (distinct) n-faces of σ:

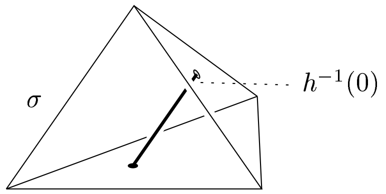

If we represent an affine map $h:\sigma \to \mathbb{R}^n$ by the $(n+2)$-tuple of values at the vertices of σ, all such maps constitute a real vector space of dimension $n(n+2)$. One can check that the set of mappings that are *not* generic is contained in a proper algebraic subvariety of this space, and so in particular, has measure zero by Sard's theorem. (Alternatively, one can check that this set is nowhere dense and use this instead of measure zero; see Exercise 1.)

Let us call a perturbed mapping $\tilde{F}: X \to \mathbb{R}^n$ *generic* if it is generic on each full-dimensional simplex of T. If T has $2N$ vertices, then the space of all possible antipodal perturbation maps P_0 on $V(\mathsf{T})$ has dimension nN (the value can be chosen freely on a set of N vertices containing no two antipodal vertices). The mappings P_0 leading to $\tilde{F}$'s that are not generic on a particular full-dimensional simplex $\sigma \in \mathsf{T}$ have measure zero in this space (here we need that $\boldsymbol{v}$ and $\nu(\boldsymbol{v})$ never lie in the same simplex of T). Therefore, arbitrarily small perturbations P_0 exist such that $\tilde{F}$ is generic.

Assuming that $\tilde{F}$ is generic and that its zeros satisfy (2.1) and (2.2), it follows that $\tilde{F}^{-1}(\mathbf{0})$ is a locally polygonal path (consisting of segments, with no branchings). This is because each n-simplex $\tau \in \mathsf{T}$ is a face of exactly two $(n+1)$-simplices $\sigma, \sigma' \in \mathsf{T}$, unless $\tau \in \mathsf{T}_t \cup \mathsf{T}_b$, in which case it is a face of exactly one $(n+1)$-simplex $\sigma \in \mathsf{T}$. Hence the components of $\tilde{F}^{-1}(\mathbf{0})$ are zero or more closed polygonal cycles (which do not intersect the top or bottom spheres) and a polygonal path γ. This γ consists of finitely many segments, and it connects $\tilde{\boldsymbol{n}}$ to $\tilde{\boldsymbol{s}}$ (these are the zeros of $\tilde{F}$ on the bottom sphere).

We choose the unit of length so that γ has length 1, and let $\gamma(z)$ denote the point of γ at distance z from $\tilde{\boldsymbol{n}}$ (measured along γ; $z \in [0,1]$). Since γ is symmetric under ν, we have $\nu(\gamma(z)) = \gamma(1-z)$, and in particular, $\nu(\gamma(\frac{1}{2})) = \gamma(\frac{1}{2})$. This is impossible, since ν has no fixed points. The Borsuk–Ulam theorem is proved.

Notes. I learned this proof from Imre Bárány, who published it, in a slightly different form, in [Bár80]. A very similar proof was given by Meyerson and Wright [MW79], and Steinlein [Ste85] has several more references for proofs of this type, all of them published between 1979 and 1981.

Exercises

1.* (a) Let $p(x_1, x_2, \ldots, x_n) = p(\boldsymbol{x})$ be a nonzero polynomial in n variables. Show that the zero set $Z(p) := \{\boldsymbol{x} \in \mathbb{R}^n : p(\boldsymbol{x}) = 0\}$ is *nowhere dense*,

meaning that any open ball B contains an open ball B' with $B' \cap Z(p) = \emptyset$.
(b) Check that a finite union of nowhere dense sets is nowhere dense.
(c) Let $\sigma := \mathrm{conv}\{\mathbf{0}, \boldsymbol{e}_1, \ldots \boldsymbol{e}_{n+1}\}$ be an $(n+1)$-dimensional simplex. Let $h: \sigma \to \mathbb{R}^n$ be an affine map (i.e., a map of the form $\boldsymbol{x} \mapsto A\boldsymbol{x}^T + \boldsymbol{b}$, where A is an $n \times (n+1)$ matrix and $\boldsymbol{b} \in \mathbb{R}^n$). If each h is represented by $(h(\mathbf{0}), h(\boldsymbol{e}_1), \ldots, h(\boldsymbol{e}_{n+1})) \in \mathbb{R}^{(n+2)n}$, show that the maps that are not generic in the sense defined in the text above form a nowhere dense set. Hint: For each possible "cause" of nongenericity, write down a determinant that becomes 0 for all maps that are nongeneric for that cause.

2.3 A Discrete Version: Tucker's Lemma

Here we derive the Borsuk–Ulam theorem from a combinatorial statement, called Tucker's lemma. It speaks about labelings of the vertices of triangulations of the n-dimensional ball. As it happens, it is also easily implied by the Borsuk–Ulam theorem: One can say that it is a "discrete version" of (BU2b).

Let T be some (finite) triangulation of the n-dimensional ball B^n. We call T *antipodally symmetric on the boundary* if the set of simplices of T contained in $S^{n-1} = \partial B^n$ is a triangulation of S^{n-1} and it is antipodally symmetric; that is, if $\sigma \subset S^{n-1}$ is a simplex of T, then $-\sigma$ is also a simplex of T.

2.3.1 Theorem (Tucker's lemma). *Let T be a triangulation of B^n that is antipodally symmetric on the boundary. Let*

$$\lambda\colon V(\mathsf{T}) \longrightarrow \{+1, -1, +2, -2, \ldots, +n, -n\}$$

be a labeling of the vertices of T that satisfies $\lambda(-\boldsymbol{v}) = -\lambda(\boldsymbol{v})$ for every vertex $\boldsymbol{v} \in \partial B^n$ (that is, λ is antipodal on the boundary). Then there exists a 1-simplex (an edge) in T that is **complementary***; i.e., its two vertices are labeled by opposite numbers.*

Here is a 2-dimensional illustration:

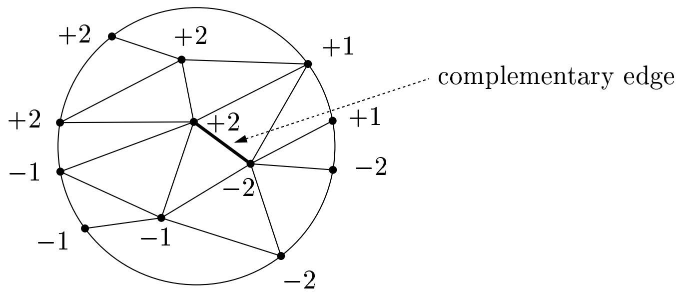

An explanation. Before we start to prove anything, we reformulate Tucker's lemma using simplicial maps into the boundary of the crosspolytope. Let $\Diamond^{n-1}$ denote the (abstract) simplicial complex with vertex set $V(\Diamond^{n-1}) = \{+1, -1, +2, -2, \ldots, +n, -n\}$, and with a subset $F \subseteq V(\Diamond^{n-1})$ forming a simplex whenever there is no $i \in [n]$ such that both $i \in F$ and $-i \in F$. By the remark below Definition 1.4.1, one can recognize $\Diamond^{n-1}$ as the boundary complex of the n-dimensional crosspolytope. The notation should suggest the case $n = 2$:

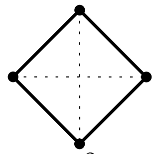

In particular, $\|\Diamond^{n-1}\| \cong S^{n-1}$. The reader is invited to check that the following statement is just a rephrasing of Theorem 2.3.1:

2.3.2 Theorem (Tucker's lemma, a reformulation). *Let* T *be a triangulation of* B^n *that is antipodally symmetric on the boundary. Then there is no map* $\lambda: V(\mathsf{T}) \to V(\Diamond^{n-1})$ *that is a simplicial map of* T *into* $\Diamond^{n-1}$ *and is antipodal on the boundary.*

Equivalence of (BU2b) with Tucker's lemma. We recall that (BU2b) claims the nonexistence of a map $B^n \to S^{n-1}$ that is antipodal on the boundary.

Deriving Tucker's lemma, in the form of Theorem 2.3.2, from (BU2b) is immediate: If there were a simplicial map λ of T into $\Diamond^{n-1}$ antipodal on the boundary, its canonical affine extension $\|\lambda\|$ would be a continuous map $B^n \to S^{n-1}$ antipodal on the boundary, and this would contradict (BU2b).

To prove the reverse implication, which is what we are actually interested in, we assume that $f: B^n \to S^{n-1}$ is a (continuous) map that is antipodal on the boundary, and we construct T and λ contradicting Theorem 2.3.2.

Here T can be chosen as any triangulation of B^n antipodal on the boundary and with simplex diameter at most δ. To specify δ, we first set $\varepsilon := \frac{1}{\sqrt{n}}$. This choice guarantees that for every $\boldsymbol{y} \in S^{n-1}$, we have $\|\boldsymbol{y}\|_\infty \geq \varepsilon$; that is, at least one of the components of $\boldsymbol{y}$ has absolute value at least ε. (If not, we would get $\sum_{i=1}^n y_i^2 < 1$.)

A continuous function on a compact set is uniformly continuous, and thus there exists a number $\delta > 0$ such that if the distance of some two points $\boldsymbol{x}, \boldsymbol{x}' \in B^n$ does not exceed δ, then $\|f(\boldsymbol{x}) - f(\boldsymbol{x}')\|_\infty < 2\varepsilon$. This is the δ bounding the diameter of the simplices of T.

Now we can define $\lambda: V(\mathsf{T}) \to \{\pm 1, \pm 2, \ldots, \pm n\}$. First we let

$$k(\boldsymbol{v}) := \min\{i : |f(\boldsymbol{v})_i| \geq \varepsilon\},$$

and then we set

$$\lambda(\boldsymbol{v}) := \begin{cases} +k(\boldsymbol{v}) & \text{if } f(\boldsymbol{v})_{k(\boldsymbol{v})} > 0, \\ -k(\boldsymbol{v}) & \text{if } f(\boldsymbol{v})_{k(\boldsymbol{v})} < 0. \end{cases}$$

Since f is antipodal on ∂B^n, we have $\lambda(-\boldsymbol{v}) = -\lambda(\boldsymbol{v})$ for each vertex $\boldsymbol{v}$ on the boundary. So Tucker's lemma applies and yields a complementary edge $\boldsymbol{v}\boldsymbol{v}'$. Let $i = \lambda(\boldsymbol{v}) = -\lambda(\boldsymbol{v}') > 0$. Then $f(\boldsymbol{v})_i \geq \varepsilon$ and $f(\boldsymbol{v}')_i \leq -\varepsilon$, and hence $\|f(\boldsymbol{v}) - f(\boldsymbol{v}')\|_\infty \geq 2\varepsilon$; a contradiction.

The definition of λ becomes more intuitive if we consider the formulation of Tucker's lemma in Theorem 2.3.2 and we think of f as going into $\|\Diamond^{n-1}\|$. Then $\lambda(\boldsymbol{v})$ is essentially the vertex of $\Diamond^{n-1}$ nearest to $f(\boldsymbol{v})$. (We have to break ties and preserve antipodality, and so the formal definition of λ above looks somewhat different.)

Special triangulations. Several combinatorial proofs of Tucker's lemma are known, but as far as I know, none establishes it in the generality stated above. One always assumes some additional properties of the triangulation T that are not necessary for the validity of the statement but that help with the proof.

Fortunately, this is no real loss of generality: For the above proof of the implication "Tucker's lemma $\Rightarrow$ Borsuk–Ulam," it is enough to know that Tucker's lemma holds for some particular sequence of triangulations with simplex diameter tending to 0. (Note that then the general form of Tucker's lemma follows from such a special case by the detour via the Borsuk–Ulam theorem.)

Two proofs of Tucker's lemma to come. In this section we present a rather direct and purely combinatorial proof. It is also constructive: It yields an algorithm for finding the complementary edge, by tracing a certain sequence of simplices.

In the next section we give another proof, completely independent of the first one (so either of them can be skipped). The second proof is perhaps more insightful, better revealing why Tucker's lemma holds. It uses some of the machinery related to simplicial homology, such as chains and the boundary operator, but in an extremely rudimentary form.

The first proof. We begin by specifying the additional requirements on the triangulation T. We first replace the Euclidean ball B^n by the crosspolytope $\hat{B}^n$, the unit ball of the ℓ_1-norm.

Let $\oplus^n$ be the natural triangulation of $\hat{B}^n$ induced by the coordinate hyperplanes. Explicitly, each simplex $\sigma \in \oplus^n$ either lies in $\Diamond^{n-1}$ (these are the simplices on the boundary), or equals $\tau \cup \{\mathbf{0}\}$ for some $\tau \in \Diamond^{n-1}$; that is, it is a cone with base σ and apex $\mathbf{0}$. The following picture shows $\oplus^2$, with some of the simplices marked by their vertex sets:

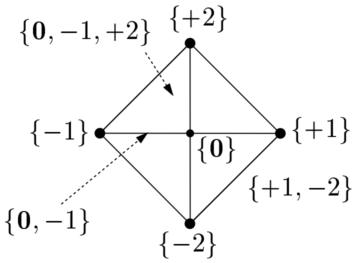

We will prove Tucker's lemma for triangulations T of $\hat{B}^n$ that are antipodally symmetric on the boundary and *refine* $\oplus^n$ (that is, for each $\sigma \in \mathsf{T}$ there is $\tau \in \oplus^n$ with $\sigma \subseteq \tau$). In other words, the second condition requires that the sign of each coordinate be constant on the relative interior of σ, for every $\sigma \in \mathsf{T}$. Let us call such a T a *special triangulation* of $\hat{B}^n$.

For $n = 2$, a special triangulation T with a labeling λ as in Tucker's lemma is shown below:

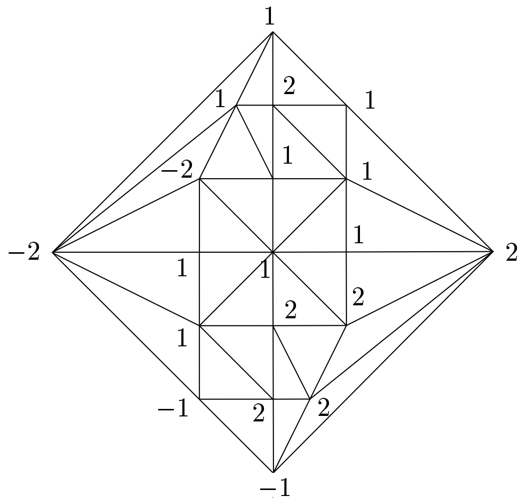

It is not hard to construct arbitrarily fine special triangulations. For example, we can start with $\oplus^n$ and repeatedly take the barycentric subdivision, until we reach a sufficiently small diameter of simplices.

We thus assume that T is a special triangulation of $\hat{B}^n$ and $\lambda\colon V(\mathsf{T}) \to \{\pm 1, \pm 2, \ldots, \pm n\}$ is a labeling antipodal on the boundary. The proof is essentially a parity argument, but not a straightforward one; we need to consider simplices of all possible dimensions. We will single out a class of simplices in T on which λ behaves in a certain way, the "happy" simplices; we will define a graph on these simplices; and we will reach a contradiction by showing that this graph has precisely one vertex of odd degree.

For a simplex $\sigma \in \mathsf{T}$, let us write $\lambda(\sigma) := \{\lambda(\boldsymbol{v}) : \boldsymbol{v} \text{ is a vertex of } \sigma\}$. We also define another set $S(\sigma)$ of labels (unrelated to the values of λ on σ).

Namely, we choose a point $\boldsymbol{x}$ in the relative interior of σ, and set

$$S(\sigma) := \{+i : x_i > 0,\, i = 1, 2, \ldots, n\} \cup \{-i : x_i < 0,\, i = 1, 2, \ldots, n\}.$$

Since T is a special triangulation, all choices of $\boldsymbol{x}$ give the same $S(\sigma)$. Geometrically speaking, $S(\sigma)$ is the vertex set of the simplex of $\diamond^{n-1}$ where σ is mapped by the central projection from $\mathbf{0}$ (and the "exceptional" simplices $\varnothing$ and $\{\mathbf{0}\}$ receive $\varnothing$).

A simplex $\sigma \in \mathsf{T}$ is called *happy* if $S(\sigma) \subseteq \lambda(\sigma)$. That is, we can regard $S(\sigma)$ as the set of "prescribed labels" for σ, and σ is happy if all of these labels actually occur on its vertices. The happy simplices are emphasized in the following picture:

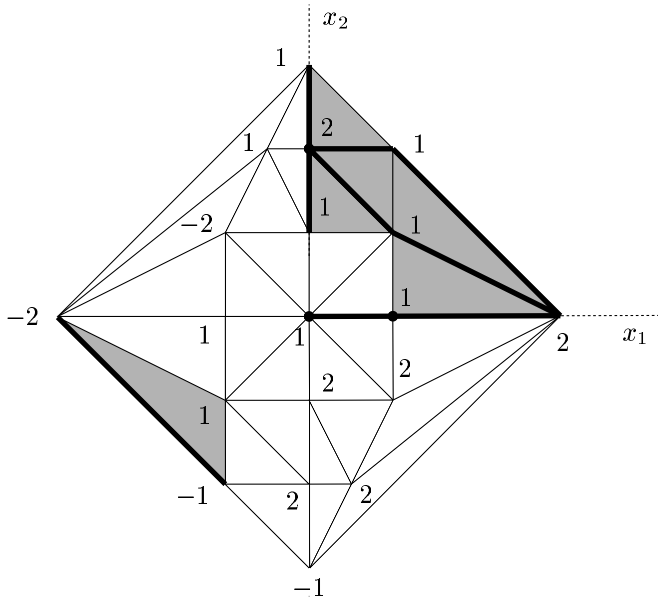

First we examine some properties of the happy simplices. Let σ be a happy simplex and let us set $k = |S(\sigma)|$. Then σ lies in the k-dimensional linear subspace L_σ spanned by the k coordinate axes x_i such that $i \in S(\sigma)$ or $-i \in S(\sigma)$. Hence $\dim \sigma \leq k$. On the other hand, $\dim \sigma \geq k-1$, since at least k vertex labels are needed to make σ happy. We call σ *tight* if $\dim \sigma = k-1$, that is, if all vertex labels are needed to make σ happy. Otherwise, if $\dim \sigma = k$, we call σ *loose*. For a loose happy simplex σ, either some vertex label occurs twice, or there is an extra label not appearing in $S(\sigma)$.

A boundary happy simplex is necessarily tight, while a nonboundary happy simplex may be tight or loose. The simplex $\{\mathbf{0}\}$ is always happy (and loose).

We define an (undirected) graph G whose vertices are all happy simplices, and in which vertices $\sigma, \tau \in \mathsf{T}$ are connected by an edge if

(a) σ and τ are antipodal boundary simplices ($\sigma = -\tau \subset \partial\hat{B}^n$); or
(b) σ is a facet of τ (i.e., a $(\dim\tau-1)$-dimensional face) with $\lambda(\sigma) = S(\tau)$; that is, the labels of σ alone already make τ happy.

The simplex $\{\mathbf{0}\}$ has degree 1 in G, since it is connected exactly to the edge of the triangulation that is made happy by the label $\lambda(\mathbf{0})$. We prove that if there is no complementary edge, then any other vertex σ of the graph G has degree 2. Since a (finite) graph cannot contain only one vertex of odd degree, this will establish Tucker's lemma.

We distinguish several cases.

1. σ is a *tight* happy simplex. Then any neighbor τ of σ either equals $-\sigma$, or has σ as a facet. We have two subcases:
   1.1. σ lies *on the boundary* $\partial\hat{B}^n$. Then $-\sigma$ is one of its neighbors. Any other neighbor τ has σ as a facet it is made happy by its labels. Thus, it has to lie in the coordinate subspace L_σ mentioned above, of dimension $k := \dim\sigma+1$. The intersection $L_\sigma \cap \hat{B}^n$ is a k-dimensional crosspolytope, and the simplices of T contained in L_σ triangulate it. If σ is a boundary $(k-1)$-dimensional simplex in a triangulation of $\hat{B}^k$, then it is a facet of precisely one k-simplex.
   1.2. σ does not lie on the boundary. Arguing in a way similar to the previous case, we see that σ is a facet of exactly two simplices made happy by its labels, and these are the two neighbors.
2. σ is a *loose* happy simplex. The subcases are:
   2.1. We have $S(\sigma) = \lambda(\sigma)$, and so one of the labels occurs twice on σ. Then σ is adjacent to exactly two of its facets (and it cannot be a facet of a happy simplex).
   2.2. There is an extra label $i \in \lambda(\sigma) \setminus S(\sigma)$. We note that $-i \notin S(\sigma)$ as well, for otherwise, we would have a complementary edge. One of the neighbors of σ is the facet of σ not containing the vertex with the extra label i. Moreover, σ is a facet of exactly one loose simplex σ' made happy by the labels of σ, namely, one with $S(\sigma') = \lambda(\sigma) = S(\sigma) \cup \{i\}$. We enter that σ' if we go from an interior point of σ in the direction of the $x_{|i|}$-axis, in the positive direction for $i > 0$ and in the negative direction for $i < 0$.

So for each possibility we have exactly two neighbors, which yields a contradiction.

Remark. The above proof proceeds by contradiction, but it can easily be turned into an algorithm for finding a complementary edge. By the above argument, a simplex σ has degree 2 in G unless $\sigma = \{\mathbf{0}\}$ or σ contains a complementary edge. So we can start at $\{\mathbf{0}\}$ and follow a path in G until we reach a simplex with a complementary edge. Such a path is indicated in the next picture:

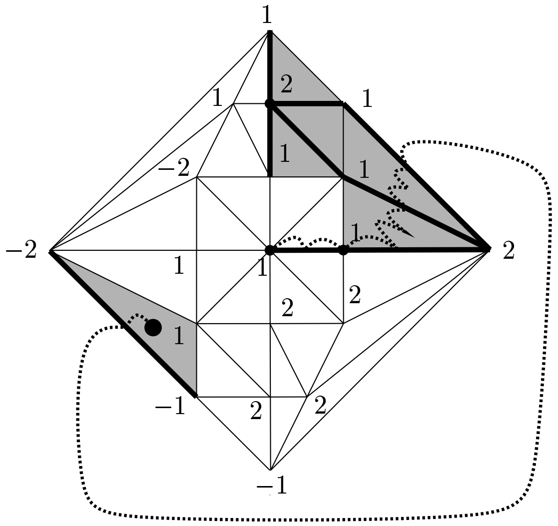

Notes. Steinlein's survey [Ste85] lists over 10 references with combinatorial proofs of the Borsuk–Ulam theorem via Tucker's lemma or some relatives of it.

Tucker's lemma is from [Tuc46]. That paper contains a 2-dimensional version, and a version for arbitrary dimension appears in the book [Lef49] (see the next section).

The proof shown above follows Freund and Todd [FT81]. They were aiming at an algorithmic proof. Such algorithms are of great interest and have actually been used for numeric computation of zeros of functions.

Exercises

1.* (A quantitative metric version of the Borsuk–Ulam theorem; Dubins and Schwarz [DS81])
(a) Let $\delta(n) = \sqrt{2(n+1)/n}$ denote the edge length of a regular simplex inscribed in the unit ball B^n. Prove that any simplex that contains $\mathbf{0}$ and has all vertices on S^{n-1} has an edge of length at least $\delta(n)$.
(b) Let T be a triangulation of the crosspolytope $\hat{B}^n$ that is antipodally symmetric on the boundary, and let $g\colon V(\mathsf{T}) \to \mathbb{R}^n$ be a mapping that satisfies $f(-\boldsymbol{v}) = -f(\boldsymbol{v}) \in S^{n-1}$ for all vertices $\boldsymbol{v} \in V(\mathsf{T})$ lying on the boundary of $\hat{B}^n$. Prove that there exist vertices $\boldsymbol{u}, \boldsymbol{v} \in V(\mathsf{T})$ with $\|g(\boldsymbol{u}) - g(\boldsymbol{v})\| \geq \delta(n)$.
(c) Derive the following theorem from (b): *Let $f\colon B^n \to S^{n-1}$ be a map that is antipodal on the boundary of B^n (continuity is not assumed). Then for every $\varepsilon > 0$ there are points $\boldsymbol{x}, \boldsymbol{y} \in B^n$ with $\|\boldsymbol{x} - \boldsymbol{y}\| \leq \varepsilon$ and $\|f(\boldsymbol{x}) - f(\boldsymbol{y})\| \geq \delta(n)$.*

This exercise is based on a simplification by Arnold Waßmer of the proof in [DS81].

2.4 Another Proof of Tucker's Lemma

Preliminaries on chains and boundaries. We introduce several simple notions, which will allow us to formulate the forthcoming proof clearly and concisely. Readers familiar with simplicial homology will recognize them immediately. But since we (implicitly) work with $\mathbb{Z}_2$ coefficients, many things become a little simpler than in the usual introductions to homology.

Let K be a simplicial complex. By a *k-chain* we mean a set C_k consisting of (some of the) k-dimensional simplices of K, $k = 0, 1, \ldots, \dim \mathsf{K}$. (The dimension will usually be shown by the subscript.) Let us emphasize that a k-chain contains *only* simplices of dimension k, and so it is not a simplicial complex.

The empty k-chain will be denoted by 0, rather than by $\varnothing$.

If C_k and D_k are k-chains, their sum $C_k + D_k$ is the k-chain that is the symmetric difference of C_k and D_k (so this addition corresponds to addition of the characteristic vectors modulo 2). In particular, $C_k + C_k = 0$.

If $F \in \mathsf{K}$ is a k-dimensional simplex, the *boundary* of F is, for the purposes of this section, the $(k-1)$-chain ∂F consisting of the facets of F (so ∂F has $k+1$ simplices). For a k-chain $C_k = \{F_1, F_2, \ldots, F_m\}$, the boundary is defined as $\partial C_k = \partial F_1 + \partial F_2 + \cdots + \partial F_m$. So it consists of the $(k-1)$-dimensional simplices that occur an odd number of times as facets of the simplices in C_k:

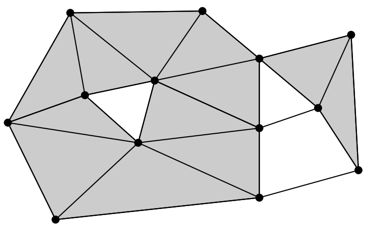

a 2-chain C_2 (gray)

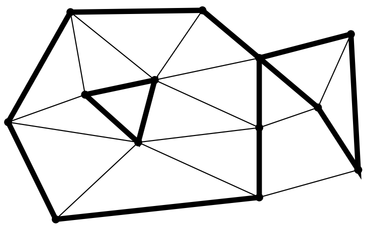

∂C_2

Important properties of the boundary operator are:

- It commutes with addition of chains: $\partial(C_k + D_k) = \partial C_k + \partial D_k$. This is obvious from the definition.
- We have $\partial\partial C_k = 0$ for any k-chain C_k. It is sufficient to verify this for C_k consisting of a single k-simplex, and this is straightforward.

A simplicial map f of a simplicial complex K into a simplicial complex L induces a mapping $f_{\#k}$ sending k-chains of K to k-chains of L. Namely, if $C_k = \{F\}$ is a k-chain consisting of a single simplex, we define $f_{\#k}(C_k)$ as $\{f(F)\}$ if $f(F)$ is a k-dimensional simplex (of L), and as 0 otherwise (so if F is "flattened" by f, it contributes nothing). Then we extend linearly to arbitrary chains: $f_{\#k}(\{F_1, F_2, \ldots, F_m\}) = f_{\#k}(\{F_1\}) + f_{\#k}(\{F_2\}) + \cdots + f_{\#k}(\{F_m\})$.

The last general fact before we take up the proof of Tucker's lemma is that these maps of chains commute with the boundary operator, in the following sense: $f_{\#k-1}(\partial C_k) = \partial f_{\#k}(C_k)$, for any k-chain C_k. It is again enough to verify this for C_k containing a single simplex.

Requirements on the triangulation. In the forthcoming proof we also need an additional condition on the triangulation T of B^n in Tucker's lemma. For $k = 0, 1, 2, \ldots, n-1$, we define

$$H_k^+ = \{\boldsymbol{x} \in S^{n-1} : x_{k+1} \geq 0, x_{k+2} = x_{k+3} = \cdots = x_n = 0\},$$

$$H_k^- = \{\boldsymbol{x} \in S^{n-1} : x_{k+1} \leq 0, x_{k+2} = x_{k+3} = \cdots = x_n = 0\}.$$

Here is a picture for $n = 3$:

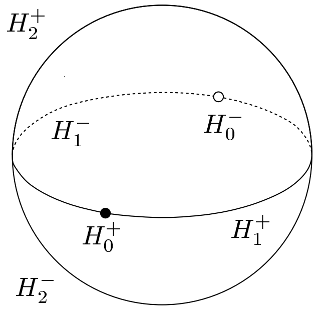

So H_{n-1}^+ and H_{n-1}^- are the "northern" and "southern" hemispheres of S^{n-1}, $H_{n-2}^+ \cup H_{n-2}^-$ is the $(n-2)$-dimensional "equator," etc., and finally, H_0^+ and H_0^- are a pair of antipodal points. We assume that T respects this structure: For each $i = 0, 1, \ldots, n-1$, there are subcomplexes that triangulate H_i^+ and H_i^-.

We prove Tucker's lemma in the version with a simplicial map into $\diamond^{n-1}$ (Theorem 2.3.2). For this proof it doesn't really matter that the mapping λ goes into $\diamond^{n-1}$; it can as well go into any antipodally symmetric triangulation L of S^{n-1}. We prove the following three claims.

2.4.1 Proposition. *Let T be a triangulation of B^n as described above, let K be the (antipodally symmetric) part of T triangulating S^{n-1}, and let L be another (finite) antipodally symmetric triangulation of S^{n-1}. Let $f\colon V(\mathsf{K}) \to V(\mathsf{L})$ be an antipodal simplicial mapping of K into L. Then we have:*

(i) *Let A_{n-1} be the $(n-1)$-chain consisting of all $(n-1)$-dimensional simplices of* K. *Then either the $(n-1)$-chain $C_{n-1} := f_{\#n-1}(A_{n-1})$ is empty, or it consists of all the $(n-1)$-dimensional simplices of* L. *In other words, either each $(n-1)$-simplex of* L *has an even number of preimages, or each has an odd number of preimages.*
In the former case (even number of preimages) we say that f has an **even degree** *and we write* $\deg_2(f) = 0$, *and in the latter case we say that f has an* **odd degree**, *writing* $\deg_2(f) = 1$.

(ii) *If $\bar{f}$ is any simplicial map of* T *into* L, *and f is the restriction of $\bar{f}$ on the boundary (i.e., on $V(\mathsf{K})$), then* $\deg_2(f) = 0$.

(iii) *If f is any antipodal simplicial map of* K *into* L, *then* $\deg_2(f) = 1$.

Hence, a simplicial map λ of T into L that is antipodal on the boundary cannot exist, since it would have an even degree by (ii) and an odd degree by (iii), which proves Tucker's lemma.

Proof of (i). This is geometrically quite intuitive, and the reader can probably invent a direct geometric proof. Here we start practicing the language of chains.

If C_{n-1} is neither empty nor everything, then there are two $(n-1)$-simplices sharing a facet such that one of them is in C_{n-1} and the other isn't. Then their common facet is in ∂C_{n-1}. At the same time, we calculate

$$\partial C_{n-1} = \partial f_{\#n-1}(A_{n-1}) = f_{\#n-2}(\partial A_{n-1}) = 0,$$

since every $(n-2)$-simplex of K is a facet of exactly two simplices of A_{n-1}. This is a contradiction.

Proof of (ii). This is again intuitive (think of an informal geometric argument) and easy. Let A_n be the n-chain consisting of all n-simplices of T. Then $A_{n-1} = \partial A_n$. At the same time, $\bar{f}_{\#n}(A_n) = 0$, simply because L has no n-simplices. Thus, $C_{n-1} = f_{\#n-1}(A_{n-1}) = \partial \bar{f}_{\#n}(A_n) = \partial 0 = 0$.

Proof of (iii). This is the challenging part. Let A_k^+ be the k-chain consisting of all k-simplices of K contained in the k-dimensional "hemisphere" H_k^+ introduced in the conditions on T, and similarly for A_k^-. We also let $A_k := A_k^+ + A_k^-$.

For $k = 1, 2, \ldots, n-1$, we have

$$\partial A_k^+ = \partial A_k^- = A_{k-1}$$

(look at the picture of the decomposition of S^{n-1} into the $H_i^{\pm}$). If we set $C_k^+ := f_{\#k}(A_k^+)$, and similarly for C_k^- and C_k, we thus obtain

$$\partial C_k^+ = \partial C_k^- = C_{k-1}.$$

Our goal is to prove $C_{n-1} \neq 0$. For contradiction, we suppose $C_{n-1} = C_{n-1}^+ + C_{n-1}^- = 0$. Then we get $C_{n-1}^+ = C_{n-1}^-$. Now the antipodality comes into play: Since A^+_{n-1} is antipodal to A^-_{n-1} and f is an antipodal map, C^+_{n-1} is antipodal to C^-_{n-1} as well, and since they are also equal, the chain $D_{n-1} := C^+_{n-1} = C^-_{n-1}$ is antipodally symmetric. Therefore, $C_{n-2} = \partial C^+_{n-1} = \partial D_{n-1}$ is the boundary of an antipodally symmetric chain.

This is a good induction hypothesis on which to proceed further. Namely, we assume for some $k > 0$ that

$$C_k = \partial D_{k+1}$$

for an antipodally symmetric chain D_{k+1}, and we infer a similar claim for C_{k-1}.

To this end, we note that the antipodally symmetric chain D_{k+1} can be partitioned into two chains, $D_{k+1} = E_{k+1} + E^{\mathrm{antip}}_{k+1}$, such that E^{antip}_{k+1} is antipodal to E_{k+1} (we divide the simplices of D_{k+1} into antipodal pairs and split each pair between E_{k+1} and E^{antip}_{k+1}). So we have $C_k = C^+_k + C^-_k = \partial(E_{k+1} + E^{\mathrm{antip}}_{k+1})$. Rearranging gives $C^+_k + \partial E_{k+1} = C^-_k + \partial E^{\mathrm{antip}}_{k+1}$. Since the left-hand side is antipodal to the right-hand side, $D_k := C^+_k + \partial E_{k+1}$ is an antipodally symmetric chain. Applying the boundary operator yields

$$\partial D_k = \partial C^+_k + \partial\partial E_{k+1} = \partial C^+_k = C_{k-1},$$

and the induction step is finished.

Proceeding all the way down to $k = 1$, we see that C_0 should be the boundary of an antipodally symmetric 1-chain. But C_0 consists of two antipodal points (0-simplices), while the boundary of any antipodally symmetric 1-chain consists of an even number of antipodal pairs (right?). This contradiction concludes the proof.

Notes. Here we have essentially reproduced Tucker's proof as presented in Lefschetz [Lef49]. Yet another degree-theoretic proof of the Borsuk–Ulam theorem is sketched in Section 6.2.

The degree of a map between spheres (or, more generally, between manifolds) is a quite useful concept. Intuitively, the degree is odd if a "generic" point in the range of the map has an odd number of preimages. We have defined rigorously the degree modulo 2 of a simplicial map between two triangulations of S^{n-1}. To extend the definition to an arbitrary continuous map f, one first defines a simplicial map $\tilde{f}$ homotopic to f (a simplicial approximation).

A similar method can be used to define the degree as an integer parameter, but one has to take the orientation of simplices into account. That is, we consider S^{n-1} as the boundary of B^n, which defines an orientation of its $(n-1)$-simplices (roughly speaking, all $(n-1)$-simplices are oriented "inwards"). To obtain the degree of f, we count the number of preimages of (any) $(n-1)$-simplex σ, where each preimage τ

such that $f(\tau)$ has the same orientation as σ is counted as $+1$, while the preimages τ with $f(\tau)$ oriented oppositely are counted as -1.

Defining the degree rigorously and establishing its basic properties (e.g., homotopy invariance) takes a nontrivial amount of work. If elementary homology theory has already been covered, a convenient definition is homological: Since the nth homology group $H_n(S^n, \mathbb{Z})$ is isomorphic to $\mathbb{Z}$, the homomorphism $f_*\colon H_n(S^n, \mathbb{Z}) \to H_n(S^n, \mathbb{Z})$ induced by f can be regarded as a homomorphism $\mathbb{Z} \to \mathbb{Z}$; thus it acts as the multiplication by some integer d, and this d is defined to be the degree of f. Dodson and Parker [DP97, Section 4.3.2] prove the Borsuk–Ulam theorem using this definition.

Another, more universal, definition of degree uses algebraic counting of the roots x of $f(x) = y$ at a "generic" image point y. The orientation of the preimages is defined using the sign of the Jacobian of the map. A proof of the Borsuk–Ulam theorem using the degree of a smooth map is sketched in [Bre93, p. 253].