Part I

Volume Preserving Homeomorphisms of the Cube

1
Introduction to Parts I and II (Compact Manifolds)

1.1 Dynamics on Compact Manifolds

Two of the principal analytic structures that may be put on a set X are measure and topology. We are interested in transformations of X which preserve *both* of these structures: measure preserving homeomorphisms. In the first half of the book, Parts I and II, the topological space X will be a compact manifold, possibly with boundary. (In fact Part I specializes to the case where X is simply the unit cube I^n in some dimension $n \geq 2$.) The measure, denoted μ, will be a nonatomic Borel probability measure which assigns the manifold boundary measure zero and is positive on all nonempty open sets (a property we call *locally positive*). (In Part I, μ is simply the volume measure on the cube.) The first two parts of the book are concerned with determining typical properties of μ-preserving homeomorphisms of the (arbitrary) compact manifold X. We denote the set of all such homeomorphisms by $\mathcal{M}[X, \mu]$, which we endow with the uniform topology, with respect to which it is complete. We call a property *typical*, or *generic*, if it is possessed by a dense G_δ (or larger) subset of transformations. The purpose of this introductory chapter is to give a nontechnical presentation of the main results, and the definitions they involve, for measure preserving homeomorphisms of compact manifolds. Both the definitions and theorems mentioned in this chapter will be presented more rigorously in later chapters.

1.2 Automorphisms of a Measure Space

Given X and μ, we will often consider more general transformations called *automorphisms*, which are bimeasurable bijections of X which preserve the measure μ. In particular, automorphisms do not need to be continuous. Since the topological structure of X is ignored the remaining measure space (X, μ) is measure theoretically the same as the unit interval with Lebesgue measure. Such a measure space is called a *finite Lebesgue space* (see [71]). Consequently the space $\mathcal{G} = \mathcal{G}[X, \mu]$ consisting of all automorphisms of (X, μ) is essentially the same as the space of all Lebesgue measure preserving bijections of the unit interval. We endow the space $\mathcal{G}[X, \mu]$ with the *weak topology*, which is determined by defining the sequential convergence of a sequence of automorphisms g_i to a limit automorphism g if $\mu(g_i(A) \bigtriangleup g(A)) \to 0$ for all measurable sets A. Here the symbol $\bigtriangleup$ denotes the *symmetric difference* between sets, defined by $A \bigtriangleup B = \left(A \cap \tilde{B}\right) \cup \left(\tilde{A} \cap B\right) = (A - B) \cup (B - A)$. The space $\mathcal{G}[X, \mu]$ is complete with respect to the weak topology.

1.3 Main Results for Compact Manifolds

Historically, the question of typical properties has been studied quite separately for the two settings ($\mathcal{G}[X, \mu]$, weak topology) and ($\mathcal{M}[X, \mu]$, uniform topology) with different techniques being applied. In each case, the first property shown to be typical was *ergodicity*. (An automorphism of a measure space is called ergodic if every invariant set either has measure zero or its complement has measure zero.) Ergodicity was proved to be typical for $\mathcal{G}[X, \mu]$, that is for automorphisms of any finite Lebesgue space, by Halmos in 1944 [69]. This followed the slightly earlier (1941) and more difficult proof of Oxtoby and Ulam [88] that ergodicity is typical in $\mathcal{M}[X, \mu]$. In a second 1944 paper, Halmos further proved that *weak mixing* automorphisms are typical in $\mathcal{G}[X, \mu]$ (an automorphism f of (X, μ) is weak mixing if $f \times f$ is ergodic on $(X \times X, \mu \times \mu)$). However, it was not until 1970 that Katok and Stepin [76] proved the corresponding result for $\mathcal{M}[X, \mu]$. Other properties have also been shown to be typical in both spaces, first in $\mathcal{G}[X, \mu]$ and later in $\mathcal{M}[X, \mu]$. In the case of homeomorphisms these results are also existence results for the specified measure theoretic behavior on arbitrary compact manifolds, since it is not known how to construct examples. However, it is easy to construct automorphisms with the required behavior. The main purpose of this part of the book is to unify these two theories in the following Theorem C obtained by the first author in 1978 [11]. In the form given, it is Corollary 10.4, which follows from a symmetric version giving simultaneous typicality in both contexts (Theorem 10.3). By a 'measure theoretic property', we mean a set $\mathcal{V}$ of automorphisms which is invariant under conjugation by any automorphism (i.e., $\mathcal{V} \subset \mathcal{G}[X, \mu]$ such that

$g^{-1}\mathcal{V}g = \mathcal{V}$ for all $g \in \mathcal{G}[X, \mu]$). See also Theorem 8.2 for a version of the theorem for the cube.

Theorem C *If a measure theoretic property is typical for length preserving automorphisms of the unit interval then it is also typical for homeomorphisms of a compact manifold which preserve a given finite nonatomic locally positive measure.*

The main idea of this part of the book, used to obtain the unification mentioned above in Theorem C, is to view the space $\mathcal{M}[X, \mu]$ as a subset of $\mathcal{G}[X, \mu]$. Thus even when the questions are entirely about homeomorphisms in $\mathcal{M}[X, \mu]$, we may employ approximations which go outside that space into $\mathcal{G}[X, \mu]$ and hence do not have to be continuous.

In order to obtain Theorem C, we need two results on the embedding of $\mathcal{M}[X, \mu]$ in $\mathcal{G}[X, \mu]$. The first, Theorem A (Theorem 8.4), lets us uniformly approximate any homeomorphism in $\mathcal{M}[X, \mu]$ by an automorphism with a desired measure theoretic property (e.g., weak mixing). An automorphism is called antiperiodic if its set of periodic points has zero measure.

Theorem A (Conjugacy Approximation) *Any homeomorphism in $\mathcal{M}[X, \mu]$ may be uniformly approximated by an automorphism of the underlying measure space which is conjugate to any given antiperiodic automorphism.*

For example, if the given automorphism is taken to be ergodic, this says that any μ-preserving homeomorphism may be uniformly approximated by an ergodic automorphism. However, since the approximating automorphism is not necessarily continuous (may lie outside $\mathcal{M}[X, \mu]$), we need an additional mechanism to eventually go back into the space $\mathcal{M}[X, \mu]$ of homeomorphisms. The relevant mechanism is a type of Lusin Theorem which says that

Theorem B (Lusin Theorem for Measure Preserving Homeomorphisms) *The space $\mathcal{M}[X, \mu]$ is dense in the space $\mathcal{G}[X, \mu]$, in the weak topology.*

Actually a stronger version of this result, Theorem 6.2, is needed. These two results (Theorem A (8.4) and Theorem B (6.2)) on the embedding of $\mathcal{M}[X, \mu]$ in $\mathcal{G}[X, \mu]$ are exactly what is needed to obtain the synthesis of Theorem C (Corollary 10.4) mentioned above regarding the identity of typical measure theoretic properties in the two spaces. These three results, Theorems A, B, C, form the core of this half of the book, on compact manifolds. In addition, we make extensive use of the 'Homeomorphic Measures Theorem' of von Neumann, and Oxtoby and Ulam, which enables us to restrict ourselves to the simple case of the unit cube with Lebesgue measure for the first eight chapters (which we call Part I), and then to simply extend the theory in Chapters 9 and 10 (which we call Part II) to any finite nonatomic locally positive measure on any compact manifold. Thus the core of this half of the book is contained in Chapters 2 (definitions), 6 (Theorem B), 8 (Theorems A, C), and 9, 10 (covering the applications of the Homeomorphic Measures Theorem).

All of these theorems establish typical ergodic theoretic behavior for volume preserving homeomorphisms. Some of the techniques can be used to establish some typical topological dynamical properties for volume preserving homeomorphisms such as transitivity or chaos. Theorem 4.8 shows that every volume preserving homeomorphism of the n-cube ($n \geq 2$) can be uniformly approximated by one which is maximally chaotic (the latter notion is stronger than the usual notion of chaos in the sense of Devaney – see Chapter 4). This result can be combined with a result of Daalderop and Fokkink [55] to prove

Theorem D *Maximal chaos is typical for volume preserving homeomorphisms of the cube.*

This is a purely topological result which has no counterpart in $\mathcal{G}[I^n, \lambda]$, the space of volume preserving automorphisms.

In addition to the above core results of this half of the book, we present a number of ancillary results based on Peter Lax's idea of approximating volume preserving homeomorphisms of the cube by permutations of the cells of some dyadic decomposition. This is a very powerful and intuitive technique which often lead to the initial proofs of new results. Indeed, the first (slightly weaker) versions of Theorems A, B, C were based on this combinatorial idea. For this reason we have included a chapter on these combinatorial techniques, as well as chapters on some applications: existence of a transitive homeomorphism of the cube and of R^n, a proof of Poincaré's Last Geometric Theorem, and the typicality of ergodicity and chaos for volume preserving homeomorphisms of the cube. The results of these Chapters (3, 4, 5, 7) will not be used elsewhere, so these chapters may be considered optional. However, they will certainly increase the reader's intuitive grasp of the ideas in this book.

2

Measure Preserving Homeomorphisms

2.1 The Spaces $\mathcal{M}, \mathcal{H}, \mathcal{G}$

This book is primarily concerned with typical measure theoretic properties (such as ergodicity or weak mixing) of members of the space $\mathcal{M}[X, \mu]$ consisting of all self-homeomorphisms of a manifold X which preserve a given Borel measure μ. To a much lesser extent, we will also consider topological properties, such as transitivity or the existence of fixed points. We will only consider manifolds of dimension at least 2. The transformations we study preserve both the measure theoretic and topological structure of the underlying space. That is, they belong to both the space of self-homeomorphisms of the manifold (denoted $\mathcal{H}[X]$) and to the space of *automorphisms* of the underlying Borel measure space (X, μ) (denoted $\mathcal{G}[X, \mu]$). An automorphism $g \in \mathcal{G}[X, \mu]$ is a bijection $g : X \to X$ with both g and g^{-1} measurable and $\mu(A) = \mu(g(A)) = \mu(g^{-1}(A))$ for all measurable sets A. Automorphisms which differ on a set of measure zero will be identified. The measure theoretic properties that we are interested in, such as ergodicity and weak mixing, do not rely on the topology of the underlying space – rather they depend only on the measure theoretic structure of the space, and for the manifolds we consider these are all the same: namely the manifolds that we consider are all measure theoretically the same as the standard Lebesgue space (I, λ), the unit interval with the sigma algebra of Lebesgue measurable sets and Lebesgue measure λ (length measure). Such measure spaces (X, μ) are called Lebesgue spaces and are distinguished only by their total measure $\mu(X)$.

In Parts I and II we consider compact manifolds with probability measures and indeed for Part I (Chapters 1–8) we consider only Lebesgue measure λ (n-dimensional volume measure) on the n-cube I^n. We denote by $\mathcal{M}[I^n, \lambda]$ the space of volume preserving homeomorphisms of the unit cube I^n. In Part II, Chapters 9 and 10 we will show that all the results obtained for this special case can be easily extended to compact manifolds with finite nonatomic measures which are positive on open sets.

In the compact case we will endow the spaces $\mathcal{H}[X]$, $\mathcal{M}[X, \mu]$, and $\mathcal{G}[X, \mu]$ with the *uniform topology* defined by the metric $\|f - g\| \equiv \operatorname{ess\,sup}_{x \in X} d(f(x), g(x))$, where d is a metric on the manifold X, usually the Euclidean or maximum metrics on the cube or torus, and denoted by $|x - y|$. We will also denote $\|f\| \equiv \operatorname{ess\,sup}_{x \in X} d(f(x), x)$ as the *norm* of f, observing that $\|fg^{-1}\| = \|f - g\|$ in our notation. Of course for the spaces $\mathcal{H}[X]$, $\mathcal{M}[X, \mu]$, the essential supremum reduces to the maximum. We will use the notation $\mathcal{H}[X, Y]$ ($\mathcal{M}[X, Y, \mu]$) to denote the subspace of $\mathcal{H}[X]$ (respectively $\mathcal{M}[X, \mu]$) consisting of homeomorphisms equal to the identity on the subset Y.

The spaces $\mathcal{H}[X]$ and its closed subset $\mathcal{M}[X, \mu]$ are not complete under the uniform topology metric given above. However, they are *topologically complete*, since that metric is equivalent to the complete metric defined by $u(f, g) = \|f - g\| + \|f^{-1} - g^{-1}\|$ (see [91]). We call this the *uniform metric*. This will justify our repeated application of the Baire Category Theorem (see [91] for discussion and proof):

Theorem 2.1 *In a complete metric space the countable intersection of dense open sets is dense.*

A set which is the countable intersection of open sets is called a G_δ set. We shall call a property *typical*, or *generic*, if the set of points with this property contains a dense G_δ set. Typical properties represent sets which are large in a topological sense, and in particular, nonempty. For this reason many of the results to be presented here can be considered existence proofs. For example, the main classical result of Oxtoby and Ulam says that ergodicity is typical among measure preserving homeomorphisms of a compact manifold. We note that a set $V \subset X$ is *nowhere dense* if for every nonempty open set U there is a nonempty open set in $U - V$ (i.e., every open set U has an open subset, 'a hole', missing V). It is easy to see that V is nowhere dense if and only if the interior of the closure of V is empty. Thus a nowhere dense set V is a 'topologically small' set (V is like a piece of Swiss cheese where the 'holes' missing V are dense in every open set). Furthermore, in a complete metric space the countable union of closed nowhere dense sets is small (since by the

Category Theorem, the complement would be a dense G_δ set). Any set which is the countable union of closed nowhere dense sets in a complete metric space is said to be a *set of first (Baire) category* (or *Baire category I*) – the complement of a set of first category is called a *residual set*. For a delightful investigation of the analogies between notions of topological smallness and measure theoretic smallness (zero measure) see J. C. Oxtoby's book *Measure and Category* [91].

At this point in the exposition, the reader would probably like to see some examples of measure preserving homeomorphisms. There is always of course the identity map. On manifolds with an additive structure which leaves the measure invariant (e.g., Euclidean space or the torus), translations of the form $x \mapsto x + c$ give simple examples. Unfortunately these will be of no use to us on general manifolds, or on the important special case of the cube, because they cannot be localized. On Euclidean space, rotations form another important example. These will in fact be useful in general because they can be localized. For example, given a planar disk of radius r centered about a point p, and a continuous function $\alpha : [0, r] \to [0, \infty)$, $\alpha(r) = 0$, we may consider the transformation which rotates the circle of radius t by an angle $\alpha(t)$, for $t \leq r$. We call this a *variable rotation*. This is clearly an area preserving homeomorphism, and we shall find that most of our constructions are ultimately limits of compositions of such local variable rotations. (An exception to this is the construction in Chapter 6.)

2.2 Extending a Finite Map

A simple question one may ask about the space $\mathcal{M}[I^n, \lambda]$ of volume preserving homeomorphisms of the cube, is whether it acts transitively on the interior. That is, given any pair of interior points x, y, can one always find a transformation h in $\mathcal{M}[I^n, \lambda]$ with $h(x) = y$? Actually, the space $\mathcal{M}[I^n, \lambda]$ possesses the stronger *finite extension property*: Any embedding $\check{h} : F \to I^n$ of a finite set $F \subset \operatorname{Int} I^n$, the interior of I^n, can be extended to a homeomorphism h in $\mathcal{M}[I^n, \lambda]$ with the norm $\|h\|$ as close to that of $\check{h}$ as desired. It is this property that allows us to combine combinatorial constructions based on finite sets with continuous approximations of various sorts. The actual construction of the extension h outlined in the lemmas below is slightly more explicit than in the original proof of Oxtoby and Ulam, to allow some additional applications not given in their original paper (in particular to the Lusin theory given in Chapter 6). It uses the variable rotations discussed in the previous section.

Lemma 2.2 *Given any two points p, q in R^n $(n \geq 2)$, and any positive number δ, let $B = B(p, q; \delta)$ denote the closed Euclidean ball centered at the midpoint of p and q, and with radius $|p-q|/2+\delta$, where $|x-y|$ denotes Euclidean distance. Then there is a volume preserving homeomorphism h of R^n which equals the identity off B, satisfies $h(p) = q$, and maps some neighborhood of p rigidly onto a neighborhood of q.*

Proof First consider the case $n = 2$. Define h by rotating the disk $B(p, q; \delta/2)$ by an angle π and rotating the circle given by the boundary of $B(p, q; \delta/2 + t)$ by the angle $\pi - 2\pi t/\delta$, for $0 \leq t \leq \delta/2$. For $n > 2$, let $D = D(p, q; \delta)$ be the intersection of $B(p, q; \delta)$ with any 2-dimensional plane through p and q. Define h on D as in the 2-dimensional case and extend it to $B(p, q; \delta)$ by requiring it to be a rigid motion of every sphere concentric to $B(p, q; \delta/2)$. Finally, extend h to the rest of R^n by setting it equal to the identity off B. □

Lemma 2.3 *Let $U \subset I^n$ be an open neighborhood of an arc L from p to q. Then there is a volume preserving homeomorphism h of I^n with $h(p) = q$, which maps some neighborhood of p onto a neighborhood of q by simple translation, and equals the identity off U.*

Proof Choose $\delta > 0$ and points $p = p_0, p_1, \ldots, p_k = q$ in L sufficiently close so that $B(p_i, p_{i+1}; \delta) \subset U$, for $i = 0, \ldots, k-1$. For $i = 0, \ldots, k-1$, let h_i be the homeomorphism given by the previous lemma for the points p_i and p_{i+1}. Then the composition $h = h_k \circ h_{k-1} \circ \cdots \circ h_1 \circ h_0$ will be the required homeomorphism if the homeomorphism h_k is an appropriate rigid motion of $B(q, q; \delta)$ which equals the identity off U. □

Theorem 2.4 *Let $\{p_i\}_{i=1}^N$ and $\{q_i\}_{i=1}^N$ be two sets of N distinct interior points of I^n, with $|p_i - q_i| < \epsilon$. Then there is a volume preserving homeomorphism h, with $\|h\| < \epsilon$ and equal to the identity on the boundary of I^n, which for each i maps some neighborhood of p_i by simple translation onto a neighborhood of q_i, and p_i into q_i. The homeomorphism h can be made to equal the identity on a given finite set disjoint from the p's and q's.*

This result also holds for any manifold possessing a metric with the property that any two points at a distance less than δ can be joined *by an arc of length less than δ (the underlying metric will be denoted by $|x-y|$). Note that the maximum metric (on I^n or the torus T^n) has this property. In all cases the required homeomorphism h can be constructed as the composition of homeomorphisms with support ($\{x : h(x) \neq x\}$) in balls.*

Proof Select N arcs $L_i : [0,1] \to \operatorname{Int} I^n$, satisfying

$$d(L_i(t), L_i(t')) < \epsilon|t-t'|, \ L_i(0) = p_i, \ L_i(1) = q_i.$$

We first prove the result under the assumption that the N sets $L_i\,[0,1]$ are disjoint, and then use this special case to prove the general result.

Assuming disjoint arcs L_i, we may choose a $\delta > 0$ sufficiently small so that the sets U_i of points within distance δ of the arc L_i are disjoint open subsets of the interior of I^n, with diameter less than ϵ. Applying the previous lemma for each i, we obtain volume preserving homeomorphisms h_i with supports ($\{x : h_i(x) \neq x\}$) in U_i, which map each p_i into q_i, and are locally translations at p_i. The composition of the h_i gives the required homeomorphism h. The L_i and U_i can always be chosen to avoid the given finite set.

In the general case, where the arcs L_i are not necessarily disjoint, we may at least assume (by suitable small displacements of the L_i, if necessary) that the arcs L_i intersect in a finite subset F, and (by small reparameterization, if necessary) that for each fixed t in $[0,1]$ the points $L_i(t)$, $i = 1, \ldots, N$, are distinct. Let $t_1 < t_2 < \cdots < t_k$ be all the values of t for which $L_i(t) \in F$ for some i. Choose numbers s_j for $j = 1, \ldots, k-1$ so that

$$0 = s_0 \leq t_1 < s_1 < t_2 < \cdots < s_{k-1} < t_k < s_k = 1.$$

Next define $p_{ij} = L_i(s_j)$ and $q_{ij} = L_i(s_{j+1})$ for $i = 1, \ldots, N$ and $j = 0, \ldots, k-1$. Then for each fixed j, the sets $\{p_{ij}\}_{i=1}^N$ and $\{q_{ij}\}_{i=1}^N$ satisfy the disjoint arc assumption (with respect to the arcs obtained by restricting the L_i to the interval $[s_j, s_{j+1}]$) and the distance condition $|p_{ij} - q_{ij}| < \epsilon(s_{j+1} - s_j)$. Hence by the special case already established, we obtain for each $j = 0, \ldots, k-1$ a volume preserving homeomorphism h_j with $\|h_j\| < \epsilon(s_{j+1} - s_j)$ and $h_j(p_{ij}) = q_{ij}$. The composition $h = h_{k-1} \circ h_{k-2} \circ \cdots \circ h_0$ satisfies the requirements of the theorem.

Since all the constructions used in the above proof for I^n are local, they can be carried out on any manifold, and produce an h which is the composition of homeomorphisms supported by balls. Also note that the