A Tourist's Guide to
The General Topology of Dynamical Systems

Graduate Studies in Mathematics, V. 1,
American Mathematical Society
Providence, R.I.

Ethan Akin
Mathematics Department
The City College
137 Street and Convent Ave.
New York, N. Y. 10031
ethanakin@earthlink.net

1993

With *The General Topology of Dynamical Systems* I am attempting to provide a broad foundation for the dynamics of maps and flows on compact metric spaces. My intent is to unify concepts of chain recurrence, Lyapunov functions, attractor/repellor theory, topological perturbation theory and topological hyperbolicity in an attractive and useful way. However, some readers familiar with the subject find my approach and notation to be idiosyncratic. With this guide I hope to seduce prospective readers, enlarge the congregation of believers and generally sell the product. It provides a sketch of each of the eleven chapters of the book with emphasis on the earlier ones where the notation is framed out.

1. Closed Relations and Their Dynamic Extensions.

In order to study the dynamics of a continuous map on a compact metric space it is useful to consider more general relations on such spaces. In pursuing this strategy through the early chapters we mimic function notation to allow the reader to use intuition which has been adapted to functions. A function $g : X_1 \to X_2$ is the subset $\{(x_1, g(x_1)) : x_1 \in X_1\}$ of $X_1 \times X_2$ (usually called "the graph of g"). For example, the identity map 1_X is the diagonal subset of $X \times X$. In general, any subset g of $X_1 \times X_2$ can be regarded as a *relation* from X_1 to X_2, written $g : X_1 \to X_2$. Thus, for $x \in X_1$ $g(x)$ is the, possibly empty, subset $\{y : (x, y) \in g\}$ of X_2. For $A \subset X_1$ the *image* $g(A) = \cup\{g(x) : x \in A\} = \{y : (x, y) \in g$ for some $x \in A\}$. As the notation suggests composition of functions extends to relations. For $g : X_1 \to X_2$ and $h : X_2 \to X_3$, $h \circ g : X_1 \to X_3$ is the relation $\{(x, z) : (x, y) \in g$ and $(y, z) \in h$ for some $y \in X_2\}$ so that $(h \circ g)(A) = h(g(A))$, a special case of the associative law, which holds for relation composition. If f is a relation on X, i.e. $f : X \to X$, we define f^n as the n-fold composition for $n = 1, 2, \ldots$ and set $f^0 = 1_X$, as usual. $g : X_1 \to X_2$ has a natural inverse relation $g^{-1} : X_2 \to X_1$, $\{(y, x) : (x, y) \in g\}$. Thus, for $B \subset X_2$, $g^{-1}(B) = \{x : g(x) \cap B \neq \emptyset\}$. In particular, $g^{-1}(X_2)$ is the *domain* of g, $\mathrm{Dom}(g)$, namely $\{x : g(x) \neq \emptyset\}$. The only point requiring care here is that neither $g \circ g^{-1}$ nor $g^{-1} \circ g$ need be the identity map. So if we define f^{-n} to be $(f^{-1})^n = (f^n)^{-1}$, the equation $f^{n_1} \circ f^{n_2} = f^{n_1+n_2}$ is true when both n_1 and n_2 have the same sign, but need not hold otherwise. The precedent for this notation occurs in the theory of uniform spaces. With $\epsilon \geq 0$ and d the metric on X we define V_ϵ (or $\bar{V}_\epsilon$) to be the set of pairs (x_1, x_2) such that $d(x_1, x_2) < \epsilon$ (resp. $d(x_1, x_2) \leq \epsilon$). So $V_\epsilon^{-1} = V_\epsilon$ expresses symmetry of the metric and $V_{\epsilon_1} \circ V_{\epsilon_2} \subset V_{\epsilon_1+\epsilon_2}$ is the triangle inequality.

Because of our standing assumption of compactness, *closed relations*, i.e. closed subsets of the product, satisfy a number of special properties. The composition and inverse of closed relations are closed, and g and A closed imply $g(A)$ is closed (Proposition 1.1). In particular, the domain of g is closed. Furthermore, a closed relation $g : X_1 \to X_2$ satisfies the continuity property. For every $\epsilon > 0$ and $x \in X$, there exists $\delta > 0$ such that $f(V_\delta(x)) \subset V_\epsilon(f(x))$ (Corollary 1.2).

This relationspeak has a productive purpose beyond the delights of onanistic formal manipulation. We now introduce operations which associate to a relation various larger relations. Even for maps much of our theory is conveniently developed using these enlargements (prolongations). Fundamental is the *orbit relation* for a relation f on X:

$$\mathcal{O}f = \bigcup_{n=1}^{\infty} f^n,$$

i.e. $y \in \mathcal{O}f(x)$ if $y \in f^n(x)$ for some $n = 1, 2, \ldots$. Observe that we begin with $n = 1$ rather than $n = 0$ so that it need not be true that $x \in \mathcal{O}f(x)$. $\mathcal{O}f$ is the smallest transitive relation containing f but is usually not closed.

For a sequence of *closed* sets $\{C_n\}$ the lim sup, C, is $\cap_n \overline{\cup_{m \geq n} C_n}$ and $\overline{\cup_n C_n} = (\cup_n C_n) \cup C$ (Exercise 1.5). So for $x \in X$ we define:

$$\omega f(x) = \limsup\{f^n(x)\}$$

$$\mathcal{R}f(x) = \overline{\mathcal{O}f(x)} = \mathcal{O}f(x) \cup \omega f(x).$$

While each of these sets is closed, the *limit point relation* ωf and the *orbit closure relation* $\mathcal{R}f$ are not usually closed in the product. To get closed relations, we define:

$$\Omega f = \limsup\{f^n\}.$$

$$\mathcal{N}f = \overline{\mathcal{O}f} = \mathcal{O}f \cup \Omega f.$$

By taking the closure of a transitive relation, we may lose transitivity. Because we want both properties we define:

$$\mathcal{G}f = \text{the smallest closed, transitive relation containing } f.$$

A relation f is both closed and transitive iff $f = \mathcal{N}f$. Because the operations are defined for relations we can iterate to reach $\mathcal{G}f$ by a transfinite induction. Start with $F_0 = f$ and define $F_{\alpha+1} = \mathcal{N}F_\alpha$ and $F_\alpha = \overline{\cup_{\beta<\alpha} F_\beta}$ for a limit ordinal α. This process stabilizes at some countable ordinal to yield $\mathcal{G}f$ (Exercise 1.18).

A somewhat larger relation, easier to describe directly is Conley's chain relation:

$$\mathcal{C}f = \cap_{\epsilon>0} \mathcal{O}(\bar{V}_\epsilon \circ f).$$

As the intersection of transitive relations $\mathcal{C}f$ is transitive. It is closed because for f a closed relation:

$$\mathcal{C}f = \cap_{\epsilon>0} \mathcal{N}(\bar{V}_\epsilon \circ f \circ \bar{V}_\epsilon).$$

(Proposition 1.8).

For a map f these are familiar definitions which I hope the unfamiliar notation does not conceal. $\mathcal{O}f(x)$ consists of the points of the positive orbit sequence of x beginning with time $n = 1$ and $\omega f(x)$ is the set of limit points of the sequence. $y \in \mathcal{N}f(x)$ iff for every $\epsilon > 0$ there is a finite orbit sequence beginning ϵ close to x and ending ϵ close to y, i.e. for some x_1 with $d(x_1, x) < \epsilon$ we have $d(f^n(x_1), y) < \epsilon$ for some $n \geq 1$. $y \in \mathcal{C}f(x)$, x *chains to* y, iff for every $\epsilon > 0$ there is a sequence $\{x_0, x_1, \ldots, x_n\}$ with $n \geq 1$ such that $x_0 = x$, $x_n = y$ and $d(f(x_{i-1}), x_i) \leq \epsilon$ for $i = 1, \ldots, n$.

To contrast $\mathcal{G}$ and $\mathcal{C}$ observe that $\mathcal{G}1_X = 1_X$ but $\mathcal{C}1_X = \cup\{C \times C : C$ is a component of $X\}$ (Exercise 1.9).

For any relation F on X the *cyclic set* of F is defined by:

$$|F| = \mathrm{Dom}(F \cap 1_X) = \{x : (x, x) \in F\} = \{x : x \in F(x)\}.$$

The cyclic sets of the various enlargements of f define the following concepts of recurrence:

$x \in |f|$: x is a fixed point of f.

$x \in |\mathcal{O}f|$: x is a periodic point of f.

$x \in |\mathcal{R}f|$: x is a recurrent point of f.

$x \in |\mathcal{N}f|$: x is a non-wandering point of f.

$x \in |\mathcal{G}f|$: x is a generalized non-wandering point of f.

$x \in |\mathcal{C}f|$: x is a chain recurrent point of f.

(This is why we did not include $f^0 = 1_X$ in the enlargements of f.)

At this point one proceeds to investigate the formal properties which follow from these definitions (Proposition 1.11), e.g. $\mathcal{A}(f^{-1}) = (\mathcal{A}f)^{-1}$ for $\mathcal{A} = \mathcal{O}$, Ω, $\mathcal{N}$, $\mathcal{G}$ and $\mathcal{C}$ but not usually for $\mathcal{A} = \omega$ or $\mathcal{R}$. Certain additional properties hold when f is a map (Proposition 1.12).

A continuous map $h : X_1 \to X_2$ maps a relation f_1 on X_1 to f_2 on X_2 when $h \times h$ takes the subset f_1 into the subset f_2, or, equivalently, when $h \circ f_1 \subset f_2 \circ h$ (Definition 1.15). This is the ordinary notion of semi-conjugacy when f_1 and f_2 are continuous maps. If h maps f_1 to f_2 then it maps $\mathcal{A}f_1$ to $\mathcal{A}f_2$ for the above operations $\mathcal{A} = \mathcal{O}$, ω, etc. (Proposition 1.17).

2. Invariant Sets and Lyapunov Functions.

For a relation f on X a subset $A \subset X$ is called $+$ *invariant* if $f(A) \subset A$ and *invariant* if $f(A) = A$. If f and A are closed and A is $+$ invariant then $\{f^n(A)\}$ is a decreasing sequence whose intersection is the largest invariant subset contained in A (Proposition 2.4). A relation F is transitive iff $F^2 \subset F$ in which case $\{F^n\}$ is a decreasing sequence of relations. If F is closed as well as transitive then

$$\Omega F = \omega F = \cap_n \{F^n\}$$

is closed and satisfies $\Omega F \circ \Omega F = \Omega F$. For a closed relation f we apply this construction to $F = \mathcal{C}f$ to get the *limit chain relation* $\Omega\mathcal{C}f$. It satisfies

$$\mathcal{C}f = \mathcal{O}f \cup \Omega\mathcal{C}f.$$

(Proposition 2.4).

If A is closed and $+$ invariant for f then it is $+$ invariant for $\mathcal{R}f$. $\mathcal{N}f$ $+$ invariance is a stronger condition. In fact, A is $\mathcal{N}f$ $+$ invariant iff it is *stable* for f meaning that the f $+$ invariant subsets include a base for the neighborhood system of A (Proposition 2.7). From this result one obtains a Urysohn Lemma construction to build Lyapunov functions (Lemma 2.10).

A continuous real valued function L on X is called a *Lyapunov function* for f if $y \in f(x)$ implies $L(y) \geq L(x)$. Thus, $y_1 \in f^{-1}(x)$ and $y_2 \in f(x)$ imply $L(y_1) \leq L(x) \leq L(y_2)$. x is called a *regular point* for the Lyapunov function if these inequalities are strict for all such y_1 and y_2. Otherwise, x is called a *critical point* for the Lyapunov function L. We denote by $|L|$ the set of critical points. The subset $L(|L|)$ of $\mathbf{R}$ is the set of *critical values.* The rest of $\mathbf{R}$ are the *regular values.*

If L is a Lyapunov function for f then it is automatically a Lyapunov function for $\mathcal{G}f$ with the same notion of critical point. Hence $|\mathcal{G}f| \subset |L|$ (Proposition 2.9). Furthermore, there exists a Lyapunov function L with $|\mathcal{G}f| = |L|$ (Corollary 2.13).

A Lyapunov function for f need not be a Lyapunov function for $\mathcal{C}f$. There is a homeomorphism f on the circle X such that $|f| = |\mathcal{G}f|$ is a Cantor set but $\mathcal{C}f = X \times X$. The only $\mathcal{C}f$ Lyapunov functions are constants but f admits a *strict Lyapunov* function, L, i.e. L increases on all nonequilibrium orbits (page 35). However, if L is a Lyapunov function for a relation f and the set of critical values is nowhere dense then L is a $\mathcal{C}f$ Lyapunov function (Exercise 3.16).

As an application one can show that a closed total order (a *preference*) on X is determined by a utility function (Exercise 2.19).

3. Attractors and Basic Sets.

For a closed relation f on X a closed subset U of X is called *inward* if $f(U)$ is contained in the interior, Int U. By compactness, U is inward iff it is $V_\epsilon \circ f$ + invariant for some positive ϵ. So an inward set is $\mathcal{C}f$ + invariant. An *attractor* for f is a closed invariant subset A such that for some sequence of closed sets $\{U_n\}$ $f(U_n) \subset \text{Int } U_{n+1} \subset U_{n+1} \subset \text{Int } U_n$ and $\cap_n U_n = A$ (Exercise 3.4). This is a strong form of asymptotic stability. An attractor can be characterized by a number of conditions of different apparent strength.

A closed + invariant subset A is called a *preattractor* if it satisfies the following equivalent conditions:

(a) There is a closed, + invariant neighborhood U of A such that $\cap_n f^n(U) \subset A$.

(b) There is an inward set U containing A such that $\cap_n f^n(U) \subset A$.

(c) A is $\mathcal{C}f$ + invariant and $A \cap |\mathcal{C}f|$ is open, as well as closed, in $|\mathcal{C}f|$.

(d) $\{x : \Omega\mathcal{C}f(x) \subset A\}$ is a neighborhood of A.

An attractor is precisely an f invariant preattractor. If B is a preattractor then $\cap_n f^n(B)$ is an attractor. Finally, a closed set is $\mathcal{C}f$ + invariant iff it is the intersection of preattractors (Theorem 3.3). There are additional recognition conditions in the map case (Theorem 3.6).

An attractor A is determined by its *trace* on the chain recurrent set, $A \cap |\mathcal{C}f|$. In fact, $A = \mathcal{C}f(A \cap |\mathcal{C}f|)$ (Proposition 3.8).

A *repellor* is an attractor for f^{-1}. An attractor- repellor pair A_+, A_- are an attractor and a repellor such that $A_+ \cap A_- = \emptyset$ and $|\mathcal{C}f| \subset A_+ \cup A_-$. For every attractor A_+ there is a unique complementary repellor $A_- = \mathcal{C}f^{-1}(|\mathcal{C}f| - A_+)$. If $x \notin A_+ \cup A_-$ then $\Omega\mathcal{C}f(x) \subset A_+$ and $\Omega\mathcal{C}f^{-1}(x) \subset A_-$ (Proposition 3.9).

The chain relation can be recovered from the attractor structure. If x is a chain recurrent point then $y = x$ or $y \in \mathcal{C}f(x)$ iff for every attractor A, $x \in A \Rightarrow y \in A$ (Proposition 3.11). Using this we obtain complete $\mathcal{C}f$ Lyapunov functions.

On the closed set of chain recurrent points, $|\mathcal{C}f|$, $(\mathcal{C}f) \cap (\mathcal{C}f)^{-1}$ is an equivalence relation. The equivalence classes are called the *basic sets* for f. So x_1 and x_2 are chain recurrent points in the same basic set iff $x_1 \in \mathcal{C}f(x_2)$

and $x_2 \in \mathcal{C}f(x_1)$. A $\mathcal{C}f$ Lyapunov function is constant on each basic set.

For a closed relation f, there exists a $\mathcal{C}f$ Lyapunov function L such that $|L| = |\mathcal{C}f|$ and L takes distinct basic sets to distinct values. Furthermore, the set of critical values is closed and nowhere dense in $\mathbf{R}$. Such a map L is called a *complete $\mathcal{C}f$ Lyapunov function* (Theorem 3.12).

4. Mappings–Invariant Sets and Transitivity Concepts.

A mapping f has a rich collection of constructions leading to closed invariant sets. For example, $\omega f(x)$ is invariant for a map f but need not be + invariant for a closed relation. For any relation f and any closed subset B the restriction f_B is the relation $f \cap (B \times B)$ on B. If f is a map and B is + invariant then f_B is a map on B.

A *decomposition* of a closed, + invariant subset A is a disjoint pair of closed, + invariant subsets with union A (page 63). A is *indecomposable* if the only decomposition for A is the trivial one $\{A, \emptyset\}$.

A relation F on X is *dynamically transitive* or acts transitively on X if $F(x) = X$ for all x, i.e. $F = X \times X$. For a map f, f is dynamically transitive iff X consists of a single fixed point and $\mathcal{O}f$ is dynamically transitive iff X consists of a single periodic orbit. For $\mathcal{C}f$, $\mathcal{N}f$ and $\mathcal{R}f$ the notions are less trivial.

A map f on X is called *chain transitive* when it satisfies the equivalent conditions: (1) $\mathcal{C}f = X \times X$, (2) For some $x \in X$, $\mathcal{C}f(x) = X = \mathcal{C}f^{-1}(x)$, (3) X is the only nonempty attractor, (4) X is indecomposable and $|\mathcal{C}f| = X$. f is called *topologically transitive* when it satisfies the equivalent conditions: (1) $\mathcal{N}f = X \times X$, (2) For some $x \in X$, $\mathcal{R}f(x) = X$, (3) X is the only closed + invariant subset with a nonempty interior, (4) $\{x : \omega f(x) = X\}$ is a residual subset of X. f is called *minimal* when it satisfies the equivalent conditions: (1) $\mathcal{R}f = X \times X$, (2) For some $x \in X$, $\mathcal{R}f(x) = X = (\mathcal{R}f)^{-1}(x)$, (3) X is the only closed, nonempty + invariant subset, (4) $\{x : \omega f(x) = X\} = X$ (Theorem 4.12).

A closed, + invariant subset B of X is called a chain transitive/topologically transitive/minimal subset if the restriction f_B satisfies the corresponding property. The limit point sets $\omega f(x)$ and the basic sets, $\mathcal{C}f(x) \cap \mathcal{C}f^{-1}(x)$ (for $x \in |\mathcal{C}f|$), are chain transitive subsets (Corollary 4.13 and Proposition 4.14). In particular, they are indecomposable. If x is recurrent, i.e. $x \in |\omega f|$, then $\omega f(x)$ is a topologically transitive subset. Let $m[f]$ denote the union of all the minimal subsets for f.

$$|f| \subset |\mathcal{O}f| \subset m[f] \subset |\omega f| \subset \omega f(X) \subset |\Omega f| \subset |\mathcal{G}f| \subset |\mathcal{C}f|.$$

We call the closure of $m[f]$ the *min-center* of f. The closure of $|\omega f|$ is the Birkhoff *center* of f. The closure of $\omega f(X)$ is the positive limit point set for f, denoted $l_+[f]$. For a homeomorphism f we let αf denote $\omega(f^{-1})$ and $l_-[f]$ denote the closure of $\alpha f(X)$. $l[f] = l_+[f] \cup l_-[f]$ is the limit point set for a homeomorphism f.

5. Computation of the Chain Recurrent Set.

For a homeomorphism f on X there are two procedures for estimating the chain recurrent set.

Let $\mathcal{U}$ be a finite collection of subsets of X whose interiors cover X. Regard $\mathcal{U}$ as a discrete metric space, a finite "approximation" of X. The *mesh* of $\mathcal{U}$ is the maximum of the diameters of the elements of $\mathcal{U}$. Associated to f is the closed relation $\mathcal{U}f$ on $\mathcal{U}$.

$$\mathcal{U}f = \{(U_1, U_2) \in \mathcal{U} \times \mathcal{U} : (U_1 \times U_2) \cap f \neq \emptyset\},$$

$$\text{i.e. } U_2 \in \mathcal{U}f(U_1) \Leftrightarrow U_2 \cap f(U_1) \neq \emptyset.$$

Since the metric space $\mathcal{U}$ is discrete, $V_\epsilon \circ \mathcal{U}f = \mathcal{U}f$ for some $\epsilon > 0$. So $\mathcal{C}\mathcal{U}f = \mathcal{O}\mathcal{U}f$. Furthermore, $\mathcal{U}f$ chains and f chains can be related in a natural way (Lemma 5.1). $|\mathcal{O}\mathcal{U}f|$ consists of the elements of $\mathcal{U}$ which are cyclic for $\mathcal{O}\mathcal{U}f$, i.e. which are periodic for $\mathcal{U}f$. Define the subset of X, $|\mathcal{U}|_f = \cup|\mathcal{O}\mathcal{U}f|$. So $x \in |\mathcal{U}|_f$ iff $x \in U$ for some cyclic element U of $\mathcal{U}$. The chain recurrent set $|\mathcal{C}f|$ is contained in $|\mathcal{U}|_f$ and the latter sets close down upon $|\mathcal{C}f|$ as the mesh of $\mathcal{U}$ tends to zero (Theorem 5.2). One can similarly estimate the individual basic sets as well (Theorem 5.8).

Instead, we can approach $|\mathcal{C}f|$ from the limit point set $l[f]$, the closure of $\omega f(X) \cup \alpha f(X)$. A closed, invariant set F is called $l[f]$ *separating* if $F \cap l[f]$ is open, as well as closed, in $l[f]$, i.e. $\{F \cap l[f],\, l[f] - (F \cap l[f])\}$ is a decomposition for $l[f]$. Because $\omega f(x)$ is an indecomposable subset of $l[f]$, $\omega f(x)$ is contained in an $l[f]$ separating set F if $\omega f(x) \cap F \neq \emptyset$. The collection of all such x is the *inset*, $W^+(F)$. Similarly, the outset $W^-(F)$ consists of those points x such that $\alpha f(x)$ meets F. If F is $l[f]$ separating and $\omega f(x) \cap F = \emptyset$, i.e. $x \notin W^+(F)$, but $\Omega\mathcal{C}f(x) \cap F \neq \emptyset$, then $\Omega\mathcal{C}f(x)$ meets

$W^+(F) - W^-(F)$. This technical result, the Shub-Nitecki Lemma, requires a rather elaborate proof (Lemma 5.13).

An *invariant decomposition* $\mathcal{F}$ for f is a finite, pairwise disjoint family of closed invariant sets whose union contains $l[f]$ (page 86). Regarding $\mathcal{F}$ as a discrete metric space we associate to f a number of different relations: $(F_1, F_2) \in \mathcal{F}_1 f$ if $W^-(F_1) \cap W^+(F_2) \neq \emptyset$ i.e. for some $x \in X$, $\alpha f(x) \subset F_1$ and $\omega f(x) \subset F_2$. $(F_1, F_2) \in \mathcal{F}_5 f$ if $\mathcal{C}f(F_1) \cap F_2 \neq \emptyset$ (page 89). Clearly, $\mathcal{F}_1 f \subset \mathcal{F}_5 f$. Using the Shub-Nitechi Lemma one shows that their transitive extensions agree, $\mathcal{OF}_1 f = \mathcal{OF}_5 f$ and we denote this common extension $\mathcal{OF}f$ (Corollary 5.14). If $\mathcal{F}_+$ is an $\mathcal{OF}f$ invariant subset of $\mathcal{F}$ and $\mathcal{F}_-$ the complementary $(\mathcal{F}f)^{-1}$ invariant set then $A_+(\mathcal{F}_+) \equiv \cup\{W^-(F) : F \in \mathcal{F}_+\}$ and $A_-(\mathcal{F}_-) \equiv \cup\{W^+(F) : F \in \mathcal{F}_-\}$ are an attractor- repellor pair. For any $F \in \mathcal{F}$ the associated *$\mathcal{F}$-basic set* $B(F)$ is $A_+((\mathcal{OF}f)(F)) \cap A_-((\mathcal{OF}f^{-1})(F))$. The $\mathcal{F}$ basic sets form an invariant decomposition and every basic set is contained in a unique $\mathcal{F}$ basic set (Theorem 5.15). An invariant decomposition $\mathcal{F}$ is called a *fine decomposition* if each $F \in \mathcal{F}$ meets a unique basic set. Then the $\mathcal{F}$ basic sets are the basic sets and every attractor-repellor pair for f comes from a pair $\mathcal{F}_+, \mathcal{F}_-$ as above. f admits a fine decomposition iff there are only finitely many basic sets (or, equivalently, there are only finitely many attractors for f) (Proposition 5.17).

An equivalent way of presenting the invariant decomposition results uses *filtrations* (Exercise 5.23).

6. Chain Recurrence and Lyapunov Functions for Flows.

For a flow φ on a compact metric space X the structures of the previous chapters extend and are closely related to those of the associated time one map f. Thus, $\omega\varphi(x)$ is the smallest closed φ invariant set containing $\omega f(x)$ (Proposition 6.3) and $\Omega\mathcal{C}\varphi = \Omega\mathcal{C}f$ (Proposition 6.5). Chain recurrence, basic sets and attractors are concepts which agree for φ and f (Propositions 6.7 and 6.9).

L is called a *Lyapunov function for the flow* φ if the function $L(\varphi(x,t))$ is differentiable in t for every $x \in X$ and $\varphi \cdot L(x) \equiv \frac{d}{ds}(L(\varphi(x,s))|_{s=0}$ is continuous and nonnegative. The *critical point set* for the flow Lyapunov function, denoted $|L|_\varphi$, is $\{x : \varphi \cdot L(x) = 0\}$. L is then a Lyapunov function for f and $|L| \subset |L|_\varphi$. Conversely, if L is a Lyapunov function for f then $\bar{L}(x) = \int_0^1 L(\varphi(x,s))ds$ is a Lyapunov function for the flow with $|\bar{L}|_\varphi \subset |L|$

(Lemma 3.10).

If X is a smooth manifold and φ is the solution flow associated with C^1 vectorfield ξ then L is a *Lyapunov function for the vectorfield* ξ if L is a C^1 function and for all x, either $d_x L(\xi(x)) > 0$ or $d_x L = 0$. L is then a Lyapunov function for the flow with $\varphi \cdot L(x) = d_x L(\xi(x))$. $|L|_\varphi$ is then the critical point set in the usual sense, $\{x : d_x L = 0\}$. If ξ is a C^r vectorfield ($1 \leq r \leq \infty$) then ξ admits a C^r Lyapunov function for the vectorfield which is a complete $\mathcal{C}f$ Lyapunov function in the sense of Chapter 3 (Theorem 6.12). Using such Lyapunov functions, inward sets can be constructed which are manifolds on the boundary of which the vectorfield always points inward (Exercise 6.18). Finally, on a smooth manifold the chain relation $\Omega\mathcal{C}\varphi$ for the solution flow of ξ can be described using piecewise smooth continuous paths whose velocity vectors are close to ξ (Theorem 6.14).

7. Topologically Robust Properties of Dynamical Systems.

For a compact metric space X, the space $C(X)$ is the compact metric space of closed subsets of X with the Hausdorff metric. A relation $g : X_1 \to X_2$ is called *pointwise closed* if $g(x)$ is closed for every x. For example, with f a closed relation on X, $\mathcal{R}f$ and ωf are pointwise closed relations which are not closed. Such a relation can be regarded as a map from X_1 to $C(X_2)$. We get the notions of upper semicontinuous (usc) and lower semicontinuous (*lsc*) relations by considering such maps (Proposition 7.11). For example, a usc relation is just a closed relation. For a *usc* or *lsc* relation $g : X_1 \to X_2$ the points of continuity of the map from X_1 to $C(X_2)$ is a residual subset (Theorem 7.19). It follows, for example, that for a continuous map f on X the set $\{x : \omega f(x) = \Omega f(x)\}$ is residual (Proposition 7.22).

These ideas yield topological perturbation results. $C(X;X)$, the set of continuous functions on X with the sup metric, is a complete metric space and the map from $C(X;X) \to C(X \times X)$ regarding each such map as a closed relation is a homeomorphism onto its image (Proposition 7.20). The maps $\mathcal{C}$, $\Omega\mathcal{C} : C(X \times X) \to C(X \times X)$ and $|\ |$, $|\ | \circ \mathcal{C} : C(X \times X) \to C(X)$ associating to a closed relation f the chain relation $\mathcal{C}f$, the limit chain relation $\Omega\mathcal{C}f$, the fixed point set $|f|$ and the chain recurrent set $|\mathcal{C}f|$ are *usc* (Proposition 7.16 and Theorem 7.23). So they are continuous at points of a residual subset of $C(X \times X)$ as are the restrictions to $C(X;X)$. One can then recover Takens' partial results toward Zeeman's so-called *tolerance stability* *conjecture* (Corollary 7.30).

Miscellaneous applications include: a "Fubini Theorem" for residual sets (Exercise 7.32), and the observation that on a residual subset of $C(X;X)$ $\mathcal{C}f = \mathcal{N}f$ for certain spaces X, *generalized homogeneous spaces*, which include manifolds of dimension at least 2 and Cantor spaces (Exercise 7.40 and Exercise 9.16).

8. Invariant Measures for Mappings.

For a compact metric space X, the space $P(X)$ is the compact metric space of Borel probability measures on X with the Hutchinson metric, yielding the topology of weak convergence (Exercise 8.16). With f a continuous map on X associate to each μ in $P(X)$ the set $M(\mu)$ of limit points of the sequence of Cesaro averages $\{\frac{1}{n}\sum_{i=0}^{n-1} f_*^i\mu\}$. Each $M(\mu)$ is a closed, connected, nonempty subset of $P(X)$ consisting of f invariant measures (Proposition 8.3). Let $M(x)$ denote $M(\delta_x)$ where δ_x is the point mass at x. Call x a *convergence point* for f if $M(x)$ contains a single measure μ_x, i.e. the Cesaro averages of $\{\delta_{f^n x}\}$ converge to μ_x. The set $\mathrm{Con}(f)$ of convergence points is a Borel set of *full measure*, i.e. $\mu(\mathrm{Con}(f)) = 1$ for every f invariant measure μ (Proposition 8.4). For a topologically transitive map f there is a closed connected, nonempty set I^* of invariant measures such that $M^{-1}(I^*) = \{x : M(x) = I^*\}$ is residual. If I^* contains more than one point (the usual situation) then $M^{-1}(I^*)$ is disjoint from $\mathrm{Con}(f)$ and so has measure 0 with respect to every invariant measure (Theorem 8.11).

Ergodicity, weak mixing and mixing as well as their topological analogues are described. A chain transitive map automatically satisfies the chain analogue of mixing unless the map projects onto a periodic orbit (Exercise 8.22).

9. Examples–Circles, Simplex and Symbols.

From a constant vectorfield one obtains the classical irrational flow on the torus whose time one map is the product of rotations. Multiply by a C^1 nonnegative function α on the torus which vanishes at a single point e. The resulting flow has a single fixed point at e. By choosing α properly one can obtain flows that do and flows that do not have an invariant measure λ equivalent to Lebesque measure. In either case the flow is topologically mixing and I^* is the entire set of invariant measures. In the former case this set is the segment in $P(X)$ connecting λ with δ_e. In the latter case, δ_e is the only invariant measure (Theorem 9.2).

On the simplex of nonnegative vectors in $\mathbf{R}^n$ which sum to 1, an $n \times n$ matrix can be used to define a differential equation modeling the evolutionary dynamics of a game. Certain limit classes of measures can be used as topological invariants to distinguish similar appearing dynamics. The example which is computed uses the paper-rock-scissors game dynamics (Theorem 9.5).

With $\mathbf{Z}_2 = \{0, 1\}$, let $\mathbf{Z}_2^\infty$ be the space of sequences in $\mathbf{Z}_2$ with the product topology. On $\mathbf{Z}_2^\infty$ the shift map s is topologically mixing with the set of periodic points, $|\mathcal{O}s|$, dense. The large set of invariant measures includes the Bernoulli measures, those measures where the separate coordinate maps to $\mathbf{Z}_2$ are independent and identically distributed. Furthermore, for x in a residual subset $(M^{-1}(I^*))$ $M(x)$ is the entire set of invariant measures (Theorem 9.11). Mapping $\mathbf{Z}_2^\infty$ to the interval $I = [0, 1]$ by using base 2 expansions we get that disjoint from the set of *normal numbers* in I, $\{x : M(x) = \lambda\}$ where λ is Lebesgue measure, is a residual set of Lebesgue measure 0 whose time averages under the shift accumulate upon every invariant measure.

We can also regard $\mathbf{Z}_2^\infty$ as the ring of 2-adic integers. Then the shift map s can be interpreted by analogy with the map extending the $3x + 1$ problem of Collatz to the 2- adics (Exercise 9.17).

10. Fixed Points.

We suspend, for a chapter, our compactness assumption to consider Lipschitz maps on Banach Spaces and their linear approximations. The central result is the stable-unstable manifold theorem for a Lipschitz perturbation of a hyperbolic linear map. This is proved by applying to a generalization of the theorem, due to Conley, the elegant fixed-point proof due to Irwin (Theorems 10.13 and 10.14). Barnsley's derivation of certain fractals as the fixed points of set maps is also considered (Exercise 10.21).

11. Hyperbolic Sets and Axiom A Homeomorphisms.

For f a homeomorphism on X and K a closed invariant subset, K is *isolated* if for some $\gamma > 0$ $\cap_{n=-\infty}^{+\infty} f^n(V_\gamma(K)) = K$, i.e. if $\{f^n(x) : n \in \mathbf{Z}\}$ lies entirely in $V_\gamma(K)$ then $x \in K$. K is *isolated rel Per* (f) if $\{f^n(x) : n \in \mathbf{Z}\} \subset V_\gamma(K)$ implies $x \in K$ for any periodic point x. K is an *expansive subset* if for some $\gamma > 0$, $\cap_{n=-\infty}^{\infty}(f \times f)^n(\bar{V}_\gamma \cap [\bar{V}_\gamma(K) \times \bar{V}_\gamma(K)])$ is contained in the diagonal

1_X, i.e. $d(f^n(x_1), K) \leq \gamma$, $d(f^n(x_2) < K) \leq \gamma$ and $d(f^n(x_1), f^n(x_2)) \leq \gamma$ for all $n \in \mathbf{Z}$ imply $x_1 = x_2$. K satisfies the *Shadowing Property* if for every $\epsilon > 0$ there exists $\delta > 0$ such that for every δ chain in K, i.e. sequence $\{x_0, \dots, x_n\}$ in K such that $d(f(x_{i-1}), x_i) \leq \delta$ for $i = 1, \dots, n$, is ϵ shadowed by some 0 chain, i.e. for some $x \in X$, $d(f^i(x), x_i) \leq \epsilon$ for $i = 0, 1, \dots, n$. K is called a *topologically hyperbolic subset* if it is an expansive subset and satisfies the Shadowing Property. A hyperbolic subset for a diffeomorphism is a topologically hyperbolic subset (Theorem 11.29 and Corollary 11.30).

If a topologically hyperbolic invariant subset K is isolated rel Per f then the restriction f_K has finitely many basic sets. Each of these is an isolated, topologically transitive subset with periodic points dense (Theorem 11.13). f is called an Axiom A homeomorphism if the chain recurrent set $|\mathcal{C}f|$ is topologically hyperbolic. For such a homeomorphism there are finitely many basic sets on each of which f is topologically transitive. Furthermore, $|\mathcal{C}f|$ is an isolated invariant set in which the periodic points are dense. For every $\epsilon > 0$ there exists $\delta > 0$ so that every homeomorphism g δ-close to f in $C(X;X)$ satisfies $|\mathcal{C}g| \subset V_\epsilon(|\mathcal{C}f|)$ and there exists a continuous map $k_g : |\mathcal{C}g| \to |\mathcal{C}f|$ mapping the restriction of g to the restriction of f (Theorem 11.19). f is called *expansive* if the whole space X is an expansive subset and f is called an *Anosov homeomorphism* if X is a topologically hyperbolic subset. For an expansive homeomorphism the space X is naturally subdivided by the stable and unstable "foliations". When f is Anosov these yield a local product structure (Exercise 11.38).