Notes on the Topology of Vector Fields and Flows

Daniel Asimov

Visualization ′93
San José, California
October 1993

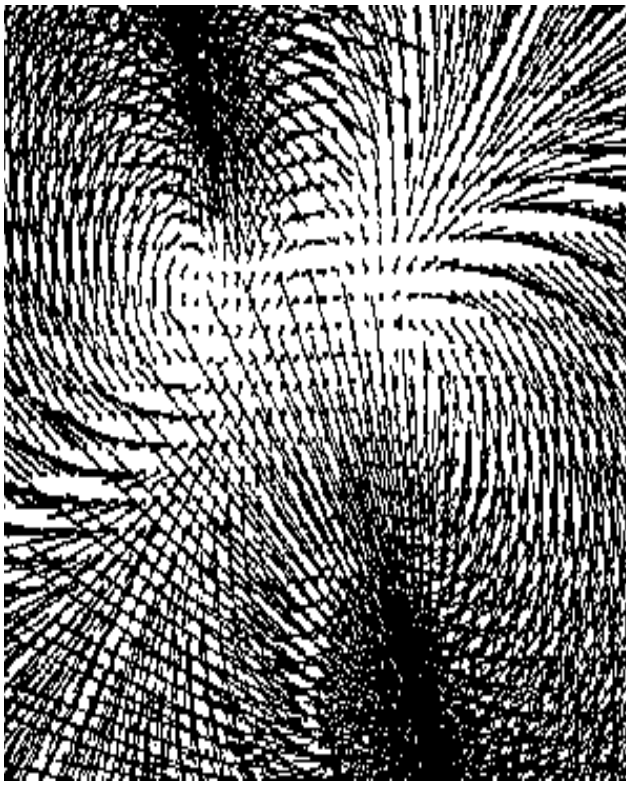
Vector fields are ubiquitous in science and mathematics, and perhaps nowhere as much as in computational fluid dynamics (CFD). In order to make sense of the vast amount of information that a vector field can carry, we present in these notes some important topological features and concepts that can be used to get a glimpse of "what a vector field is doing."

1 Vector Fields

A vector field arises in a situation where, for some reason, there is a direction and magnitude assigned to each point of space.

The classic example of a vector field in the real world is the velocity of a steady wind. We will draw a vector field as having its point of origin at the point x to which it is assigned. In this way a vector field resembles a hairdo, with all the hair being perfectly straight

Figure 1: Vector Field (unnormalized)



A **vector field** is, mathematically, the choice of a vector for each point of a region of space. In general, let U denote an open set of Euclidean space R^n. (We will usu-

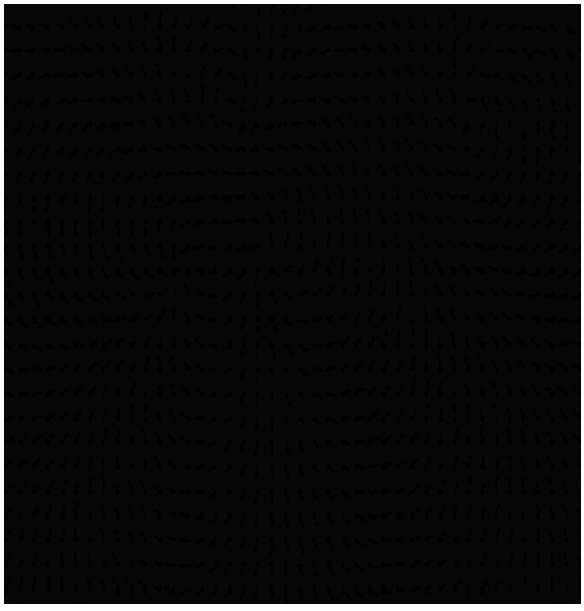
ally be interested in the cases $n = 2$ and $n = 3$.) Then a vector field on U is given by a function $f : U \to R^n$. We will assume that a vector field f is at least once continuously differentiable, denoted C^1.

More generally, one can also consider a phenomenon such as real wind, which can be represented by a vector field which changes from moment to moment. The mathematical name for this is **time-dependent vector field**. Mathematically, a time-dependent vector field on an open set U in R^n is a function of the form $f : U \times J \to R^n$, where J is some time interval $a \le t \le b$.

Think of a time-dependent vector field as a movie, each frame of which is an ordinary vector field.

In these notes we will give more emphasis to steady (non-time-dependent) vector fields.

Figure 2: Normalized Vector Field



2 Differential Equations and Solution Curves

A **differential equation** (or more precisely, an autonomous, ordinary differential equation) defined on an open set U of Euclidean space R^n is an equation of the form

$$\frac{dx}{dt} = f(x)$$

where $x : I \to U$ denotes an unknown curve parametrized by some (as yet unknown) interval I containing 0, and $f : U \to R^n$ is C^1.

Given any point $x_0 \in U$, we may assume an "initial condition" of the form $x(0) = x_0$. With this condition, the Existence Theorem for solutions of autonomous O.D.E.'s states the following: There exists a number $c > 0$ and a solution $x : (-c, c) \to U$. In other words, $x(0) = x_0$ and, for each t in the interval $(-c, c)$, we have $x'(t) = f(x(t))$.

This solution is unique in that any curve $y : (-c, c) \to U$ which also satisfies the initial condition $y(0) = x_0$ and the equation $y'(t) = f(y(t))$, must in fact be the curve x.

Think of f(x) as representing the velocity of a steady wind at the point x. Then the solution curve x(t) represents the hypothetical trajectory of a massless particle released at time = 0 at the point x_0.

One may also consider **non-autonomous differential equations**. These are equations of the form

$$\frac{dx}{dt} = f(x, t)$$

where $x : I \to U$ once again denotes an unknown curve parametrized by some interval I. However, in this case, $f : U \times J \to R^n$ is a C^1 function taking not one but two arguments. (Here J represents some interval of time.)

To say that x is a solution of a non-autonomous O.D.E. like this means that x must satisfy the initial condition, $x(t_0) = x_0$. And there must exist some number $c > 0$ such that for each t lying in the interval $(t_0 - c, t_0 + c)$, we have $x'(t) = f(x(t), t)$. (Note that unlike the interval $(-c, c)$ in the autonomous case, which must contain 0, this interval may omit 0 but must contain the value t_0.)

The classic example of such an equation is again wind: real wind, which varies with time. The function f(x, t) may be thought of as representing the velocity of the wind at location x and time t. And the solution curve x represents the trajectory of a particle that is released in the wind at time $= t_0$ at the point x_0.

3 Phase Space

In addition to concrete geometrical applications, one of the most important uses of differential equations is to understand the evolution of the state of an abstract physical system as time progresses. For example, consider the hinged roof-shaped lid atop some public trashcans. Its physical state can be described by two coordinates: the angular position of the lid, and the angular velocity with which it is rotating about its hinge. The abstract representation of all states is actually a cylinder, because the angular position repeats after 2π. This cylinder, then is the **phase space** for this problem. The differential equation governing the physics of the trashcan lid provides a vector field on this phase space.

Trajectories of this vector field decribe the change in state that the lid would experience form any initial position. Physically, this is the classic "pendulum with friction."

Figure 3: Pendulum with Friction



In this figure the horizontal coordinate represents the angular position (which has been unrolled to lay the cylinder out flat; hence the "phase portrait" above repeats with horizontal period 2π). The vertical coordinate represents the angular velocity, of the trashcan lid. For any initial condition—i.e., a point of this phase space— the unique trajectory though that point shows the time evolution of that initial condition. It is apparent from the figure that, for almost all initial conditions, the trajectory eventually spirals into a stationary point of the vector field. This

corresponds to going through some finite number of full revolutions in angular position, and then swinging back and forth, approaching the stationary equilibrium state in the limit. Anyone who has scientifically experimented with this kind of trashcan knows that this is exaactly what happens.

4 Vector Fields ↔ O.D.E.'s

In a mathematical sense, vector fields and differential equations may be considered to be the same thing. More precisely, time-independent vector fields are the same thing as autonomous ordinary differential equations.

This is simply because each one is determined by a function $f : U \to R^n$. Given a vector field defined by a function $f : U \to R^n$, one may write the differential equation given by $dx/dt = f(x)$. And conversely, given a differential equation $dx/dt = f(x)$, one may extract the vector field $f(x)$. Thus both vector fields and autonomous ordinary differential equations carry the same information.

And similarly, time-dependent vector fields may be thought of as the same thing as non-autonomous differential equations: each one is determined by a function $f : U \times J \to R^n$.

5 Flows

Vector fields and differential equations give rise to families of transformations of space called **flows**.

Let us introduce new notation for the trajectory of a point p: the solution curve which takes the value p at time = 0 will be denoted by α_p.

These trajectory functions α_p must, where defined, satisfy the fundamental **consistency condition**: Let $p \in U$ and let s and t denote any two lengths of time. Then the result of following the trajectory of p for time s, and then following the trajectory of that result for time t, is exactly the same as the result of following the trajectory of p for time s + t.

In mathematical language: Let q denote the point $\alpha_p(s)$, the result of following the trajectory of p for time s. Then we must have $\alpha_q(t) = \alpha_p(s + t)$.

It is useful to unify the trajectory functions α_p for all p into one single function of two variables. For any $p \in U$ and $t \in I_p$ define $\phi(p, t)$ to be $\alpha_p(t)$. Now $\phi(p, t)$ is just the result of following the trajectory of p for time t.

Then the consistency condition becomes just

$$\phi(\phi(p, s), t) = \phi(p, s + t)$$

wherever it is defined. (Note that this holds for negative as well as positive values of s and t.) We also have $\phi(p, 0) = p$ for all p in U.

By the way, the existence and uniqueness theorem for O.D.E.'s also tells us that the flow function ϕ of the two variables p and t must be continuously differentiable, or C^1.

Now we can see that for each fixed value of t, there is a mapping of U to itself which takes p to $\phi(p, t)$. This mapping is often denoted by ϕ_t. Thus, by definition, we

have $\phi_t(p) = \phi(p, t)$ wherever this is defined.

Finally, we can define what is meant by **the flow associated with a differential equation**: This is the family of all the transformations $\phi_t : U \to U$. The consistency condition for this flow can now be stated as $\phi_s \circ \phi_t \equiv \phi_{s+t}$ for all s and t for which this is defined. Due to the equivalence of differential equations and vector fields, the flow of the differential equation dp/dt= f(p) is also called the **flow of the vector field** f(p).

No convergent trajectories of a C^1 vector field

The uniqueness of solutions to O.D.E.'s tells us important information about the configuration of trajectories of a C^1 vector field: Each point p of the domain U lies on one and only one trajectory.

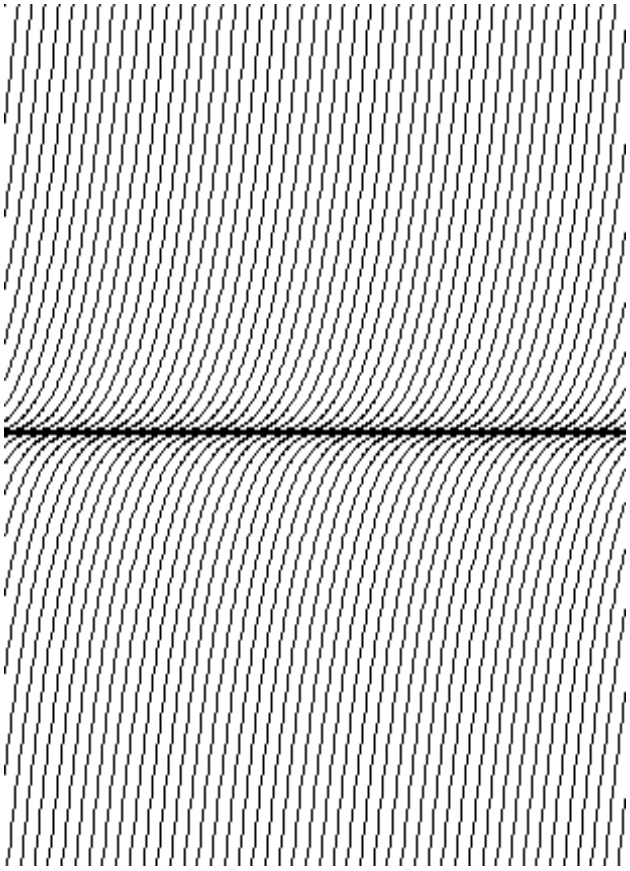
In particular, this tells us that trajectories can never truly converge, cross each other, or branch. (Two trajectories may **appear** to converge, for example, if as time $\to \infty$ both trajectories approach the same point. But this can only occur when that point is in fact its own trajectory: a stationary point.)

It is enlightening to see what can happen when we drop the condition that the vector field be continuously differentiable. An interesting example (Fig. 4) can be constructed if a vector field is defined on all of R^2 via $V(x, y) = (1, 3y^{2/3})$. (The derivative $\partial(3y^{2/3})/\partial y$ is $2y^{-1/3}$, which does not exist when y = 0, so this vector field is not continuously differentiable.)

Let a be any number. Then the curve given by $t \to (t, (t - a)^3)$ is a trajectory of this vector field. But since V(x, 0) = (1, 0) for all x, it is clear that the x-axis itself is also a trajectory, via the curve given by $t \to (t, 0)$. These two trajectories cross each other.

At the point where they cross, two distinct trajectories start from the same initial conditions.This violates the uniqueness of solutions which must hold for C^1 vector fields.

Figure 4: Vector Field $V(x, y) = (1, 3y^{2/3})$



The flow of a linear vector field

In the case where a vector field on R^n is defined by a matrix, then there is a simple explicit formula for the flow of V.

Suppose V(p)= Lp at each point p, where L is some $n \times n$ matrix. Then the flow of V is given by a very simple and beautiful formula: $\phi(p, t) \equiv e^{tL}p$. Here the exponential e^M of an $n \times n$ matrix is the $n \times n$ matrix defined as

$$e^M = \sum_{k=0}^{\infty} \frac{M^k}{k!}$$

6 Classification of Trajectories

We are now in a position to make an elementary classification of the kinds of trajectories that can occur. We may think of a trajectory as a mapping $x: R \to R^n$ from the real numbers to Euclidean space[1]. **Then one of three conditions must hold:**

1) The mapping is one-to one. In this case the trajectory is a curve that never returns to where it has been. This is called a **regular trajectory**. (More about these later.)

2) The mapping has a least period > 0, say t_0, after which it repeats exactly where it has been before: $x(t + t_0) = x(t)$ for all t. The trajectory must form a simple closed curve in Euclidean space. Such a trajectory is called a **closed orbit**. (More about these later.)

3) The trajectory stays put. The entire trajectory is a single point: x(t) = p for all t. This happens at locations where the vector field is **0**, and nowhere else. This kind of trajectory is called a **stationary point**. (More about these later.)

Terminology: Other terms for a stationary point include **equilibrium**, **singularity**, **fixed point**, and **zero**. (It is also sometimes called a "critical point," but this term is perhaps best reserved for its original definition: where the derivative of some mapping to R^1 vanishes.) A vector field without stationary points is called **non-singular**.

7 Some Essential Concepts of Topology

Some fundamental concepts of topology in Euclidean space R^n that will be used or assumed here include the following:

An **ε-neighborhood** of a point p for a number, denoted $N_\varepsilon(p)$, is the set of all points at a distance $< \varepsilon$ from p.

A point p is a **limit point** (or **accumulation point**) of a set X if for every $\varepsilon > 0$, $N_\varepsilon(p) - \{p\}$ contains at least one point of X.

An **open set** is any set that is an arbitrary union of sets of the form $N_\varepsilon(p)$. Note that the empty set is open, as is the entire space R^n.

1. In these notes, we will often not distinguish trajectories which arise from different initial conditions but constitute the same set of points.

A **closed** set is any set X whose complement $R^n - X$ is open. Equivalently, a closed set is any set which contains all its limit points.

A **neighborhood** of a point p is any open set containing p.

A set in R^n is **compact** if it is closed and bounded.

A set X is **disconnected** if there exist two disjoint open sets, each of whose intersection with X is non-empty. A set is **connected** if it is not disconnected. Intuitively, X is connected if it is all in one "piece."

A sequence $p_1, p_2,...,p_i,...$ **approaches the limit** p if for every $\varepsilon > 0$, there exists an integer $M = M(\varepsilon) > 0$, such that $i \geq M$ implies that p_i lies in $N_\varepsilon(p)$. This situation is denoted by $\lim p_i = p$. (If you are not familiar with this definition, it is worth taking the time to understand it.) Note that a sequence can approach at most one limit, although a set of points $\{p_i\}$ may have many *limit points*.

A mapping $f : U \rightarrow V$ between two open sets U and V is **continuous** if whenever there are points $\{p_i\}$ and p in U for which $\lim p_i = p$, then $\lim f(p_i) = f(p)$. (In brief: a continuous function is one which preserves limits.)

A **homeomorphism** or **topological equivalence** between two sets X and Y is a mapping $h : X \rightarrow Y$ that is continuous, one-to-one, onto, and has continuous inverse h^{-1}. Topology is concerned with properties of shapes that remain unchanged under homeomorphism.

Two vector fields or flows are called **topologically equivalent** (or homeomorphic) if there is a homeomorphism between their domains which carries trajectories to trajectories, and preserves the direction of increasing time.

Fundamental theorem of O.D.E. s

In a sense, closed orbits and stationary points are much rarer than regular trajectories, at least for most vector fields. If p is a point on a regular trajectory, we can state exactly what nearby trajectories look like in some neighborhood of p.

Theorem: If p is a point on a regular trajectory, then there is a neighborhood U of p and a homeomorphism $h : U \rightarrow R^n$ which carries each piece of a trajectory lying in U onto a straight line in R^n parallel to the x-axis.

In plain English, this theorem says that **near a point of a regular trajectory, trajectories fill up space in the same way that parallel lines do**, topologically speaking.

8 Alpha- and Omega-limit Sets

We wish to define how a regular trajectory x(t) behaves as t approaches infinity. So we define the **omega-limit set of the trajectory** x to be all those points of the domain which are limits of a sequence of the form x(t) for values of t approaching infinity. We denote this set by $\omega(x)$.

Example: Consider a vector field in which one trajectory x approaches a closed orbit

by spiraling around it closer and closer. Then the entire closed orbit is the omega-limit set of x.

Example: Consider the vector field $V(x,y) = (-x, y)$ in the plane. The trajectory x(t) satisfying the initial condition $x(0) = (x_0, y_0)$ is given by the curve $x(t) = (x_0e^{-t}, y_0e^{t})$. Thus each trajectory lies on a hyperbola of the form $xy = k$ for some value of k (and this includes the degenerate hyperbola $xy = 0$ as well).

Clearly, the only trajectories having a non-empty omega-limit set are those lying on the x-axis.

Similarly, the alpha limit set, of a trajectory x(t), denoted by $\alpha(x)$, is defined as the omega-limit set of the reversed trajectory $x(-t)$. In the above example, only trajectories lying on the y-axis have a non-empty alpha-limit set.

Suppose there is a number t_0 for which the trajectory x(t) remains inside a compact subset of the domain for all $t \geq t_0$. Then the following must be true concerning the omega-limit set of x:
1) $\omega(x)$ is non-empty
2) $\omega(x)$ is closed
3) $\omega(x)$ is invariant by the flow (i.e., it is a union of trajectories), and
4) $\omega(x)$ is connected.

9 Limit Cycles

The limit cycle is an important topological feature of some vector fields.

Definition: A **limit cycle** of a vector field V is a closed orbit C of V that is contained in $\alpha(x)$ or $\omega(x)$ for of some trajectory $x \neq C$.

Here is an important result about limit cycles in the plane:

Theorem (Poincaré-Bendixson): Suppose V is a vector field in R^2, and let x be a trajectory of V. Suppose that $\omega(x)$ is non-empty, compact, and contains no stationary point. Then $\omega(x)$ must be a closed orbit.

What makes the limit cycle particularly important is that it will persist even after a small perturbation of the vector field.

Theorem: Let V be a vector field having a limit cycle C. Let V_1 be a sufficiently small (in the C^0 sense) perturbation of V. Then V_1 will also have a limit cycle.

10 Classification of Isolated Stationary points

A stationary point is called **isolated** if is not a limit point of other stationary points. If the point p is an isolated stationary point of a vector field
$V(x_1, x_2) = (V_1(x_1,x_2),V_2(x_1,x_2))$, or
$V(x_1, x_2, x_3) = (V_1(x_1,x_2,x_3),V_2(x_1,x_2,x_3),V_3(x_1,x_2,x_3))$,
then we can further investigate the structure of the trajectories nearby by examining the so-called Jacobian matrix of partial derivatives, $JV = (\partial V_i/\partial x_j)$. This matrix will of course be 2×2 in the 2-dimensional case and 3×3 in the 3-dimensional case.

When we evaluate the Jacobian matrix at the stationary point p, we obtain a matrix whose values are numbers. As with any

such matrix, we may calculate its eigenvalues and eigenvectors.

Definition: A stationary point is called **hyperbolic** if the real parts of the eigenvalues of its Jacobian are all non-zero .

A hyperbolic stationary point must necessarily be isolated.

Theorem: If p is a hyperbolic stationary point, then the trajectories near p are determined up to topological equivalence by the number of positive and the number of negative real parts that the various eigenvalues have.

Just how often does this "non-zero real part" condition hold? Fortunately for us, it is the rule and not the exception. In a sense that can be made precise, this condition holds with probability = 1.

This theorem is a remarkable application of topology to the local classification of stationary points.

According to this theorem, we can categorize "all but a set of measure 0" of the isolated stationary points in 2 dimensions in terms of just 3 kinds: both eigenvalues with positive real part; one positive and one negative, and both negative. Similarly, in 3 dimensions there are 4 distinct types (except for that set of measure 0).

Definition: A **sink** is an isolated stationary point p each of whose nearby trajectories x satisfies $\omega(x) = \{p\}$. (Alternatively, for each nearby trajectory x, the limit of x(t) as $t \to \infty$ is p.)

Definition: A **source** of a vector field V is an isolated stationary point p which is a sink of the reversed field, −V.

11 Geometric Classification of Hyperbolic Stationary points

Let V denote a vector field with a hyperbolic stationary point p, at which the Jacobian matrix is $JV = (\partial V_i/\partial x_j)$.

The 2-dimensional case:

The eigenvalues of JV are the roots of the quadratic polynomial $P(\lambda)$ defined by $P(\lambda) = \det(J - \lambda I)$. Since this is a quadratic polynomial with real coefficients, the roots are either both real or else they are complex conjugates of each other (say K + Li and K − Li, where K and L are real). We consider the cases:

1. Both roots real:
   a. Both positive:
      source
   b. Both negative:
      sink
   c. Opposite signs:
      saddle
2. Complex conjugate roots:
   a. K positive:
      spiral source
   b. K negative:
      spiral sink

Note: The sense of the spiral in case 2. can be determined from the sign of curl(V) at the point p. If curl(V) (thought of as the scalar quantity $\partial V_1/\partial x_2 - \partial V_2/\partial x_1$) is positive, then the swirl will be counterclockwise, and vice versa.

The 3-dimensional case:

The eigenvalues of JV are the roots of the polynomial $P(\lambda)$ defined by $P(\lambda) = \det(J - \lambda I)$. Since this is a cubic polynomial with real coefficients, the theory of equations tells us that there are either 3 real roots, or 1 real root and a pair of complex conjugate roots. Case by case:

1. All 3 roots real:
   a. All positive:
      source
   b. 2 positive, 1 negative:
      saddle (2 dims. out, 1 in)
   c. 1 positive, 2 negative:
      saddle (1 dim. out, 2 in)
   d. All negative:
      sink

2. 1 real, 1 complex conjugate pair (again, say $K + Li$ and $K - Li$, where K and L are real):
   I) Real root positive:
      a. K positive:
         spiral source
      b. K negative:
         spiral saddle (1 out, 2 dims. in)
   II) Real root negative:
      a. K positive:
         spiral saddle (2 dims. out, 1 in)
      b. K negative:
         spiral sink

12 What is So Important About Hyperbolic Stationary Points?

If V is a vector field with a hyperbolic stationary point at the point p, then this is stable in the following sense: Suppose V_1 is another vector field which is sufficiently "C^1 close" to V in a neighborhood of p. (This means that V_1 and its first derivatives are close to the corresponding values of V near p.) Conclusion: the vector field V_1 must also have its own stationary point at some point p_1 near p, and furthermore **the trajectories of V_1 near p_1 must be topologically equivalent to the trajectories of V near p.** In brief: a local perturbation of a hyperbolic stationary point does not change the topology.

This kind of stability is called **local structural stability**.

This is a **very** important property of hyperbolic stationary points. It means that they will show up independent of small errors of measurement as invariably occur in the real world.

Conversely, any isolated stationary point possessing local structural stability must be hyperbolic.

13 Area- or Volume-Preserving Vector Fields

In many applications—for example, hydrodynamics—the vector fields encountered will often be volume-preserving. (In this context we use the term "volume" to mean both ordinary volume in 3 dimensions and area when we are referring to 2 dimensions.)

By definition, a vector field is **volume-preserving** when its flow ϕ_t carries any open set S in its domain to a set $\phi_t(S)$ of the same volume, for all times t.

The condition for a vector field V to preserve volume is that $\mathrm{div}(V) = 0$. (In 2D this means $\partial V_1/\partial x_1 + \partial V_2/\partial x_2 \equiv 0$; in 3D it

means $\partial V_1/\partial x_1 + \partial V_2/\partial x_2 + \partial V_3/\partial x_3 \equiv 0$.). For this reason, the term **divergence-free** is used as a synonym for volume-preserving.

A fundamental topological consequence of being volume-preserving is that **no open set of finite volume may ever flow to a proper subset of itself** by the flow, whether in positive or negative time.

This strong condition immediately rules out both sources and sinks from occurring in volume-preserving vector fields. But it is easy to see that any type of saddle is possible in a volume-preserving field.

Centers

In 2 dimensions, there is another important kind of isolated stationary point, a center. A **center** is a stationary point in 2 dimensions for which **all nearby trajectories are closed orbits.** The simplest example of a center is the point (0, 0) in the vector field given by $V(x, y) = (-y, x)$ (uniform counterclockwise rotation).

The eigenvalues of the Jacobian matrix at a center must be **pure imaginary**.

The converse is **not** true in general: If an **arbitrary** 2-dimensional vector field has an isolated stationary point with pure imaginary eigenvalues of the Jacobian, it is not necessarily a center. It could also be a source or a sink (due to non-linear terms of its Taylor series).

However, if an **area-preserving** vector field has an isolated stationary point with pure imaginary eigenvalues of the Jacobian, then it must be a center (since the only alternatives—source or sink—are excluded because they cannot preserve area).

A center is important because it possesses **constrained local structural stability**: Any *area-preserving* vector field that is sufficiently close (in the C^1 sense) to one having a center, must also have a center.

Thus if the only vector fields that can arise in some 2-dimensional situation must be area-preserving, then the only isolated stationary points that can happen are 1) saddles and 2) centers.

14 The Poincaré Map

By the use of the Poincaré map, we may reduce the topology of surrounding trajectories to a question of understanding a map of a smaller disk to a larger disk.

We will consider a vector field in 3 dimensions. The 2-dimensional case, being simpler, will follow from this.

Choose any point p belonging to C, and consider any open disk D whose intersection with the circle C is the point p, and which is not tangent to the circle C at p. (In the 2-dimensional case we would use an arc instead of a disk.) We may assume that D is chosen so that the vector field is nowhere tangent to D. Such a disk (or arc, if in 2 dimensions) is called a **local section** of the flow.

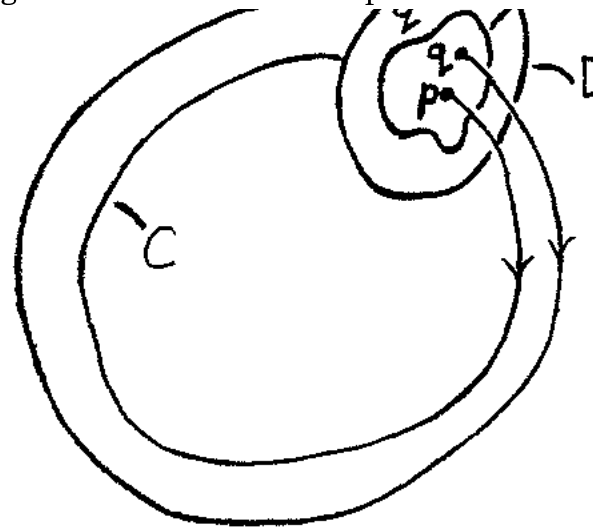
By the continuity of solutions to C^1 differential equations, there must exist a smaller disk D_0, contained in D and containing p, with the following property:

For any point q in D_0, there is a smallest value of $t = t(q) > 0$ such that following the trajectory of p for time t(q) results in a point q' that is once again in the larger disk D (but not necessarily in D_0).

The **Poincaré map** (or **first-return map**) of the closed orbit C is the mapping $f : D_0 \rightarrow D$ that takes any point q of D_0 to the point q' described above. If V is a C^1 vector field (as we have usually been assuming), then the Poincaré map f is also continuously differentiable.

This Poincaré map f must carry the original point p to itself. This is because f(p) must lie on D, and it also lies on the closed orbit C through p; the intersection of C and D is {p} by assumption.

Figure 5: The Poincaré Map



Important fact about Poincaré maps: Regardless of the choices involved in the definition (the point p, the disks D and D_0), any two resulting Poincaré maps of the same closed orbit will be topologically equivalent (in fact, differentiably equivalent).

The Poincaré map is important because it contains all the information about the topology of trajectories close to the closed orbit C. (In principle, given the Poincaré map of a closed orbit, one could reconstruct the topological structure of trajectoriesnear C.) But it is simpler to think about a mapping between two disks than all that spaghetti.

Since f has a fixed point at p, we can again look at the Jacobian matrix, not of a vector field, but of the mapping f. The Jacobian of a mapping f is again defined as the square matrix of partial derivatives, in this case $Jf = (\partial f_i/\partial x_j)$. (In order for this to make sense, we need to pretend that we have given 2-dimensional coordinates to the disk D.)

Important fact about eigenvalues: Regardless of the choice of coordinates used for the disk D, (which will affect the Jacobian matrix Jf), the **eigenvalues** of the matrix Jf will always be the same for a given closed orbit.

And so once again we can play the eigenvalue game. Analogous to the terminology for stationary points, we call a closed orbit **hyperbolic** if the eigenvalues of Jf lie **off the unit circle** in the complex plane.

As in the case of hyperbolic stationary points, **hyperbolic closed orbits are locally structurally stable**: a small perturbation of the vector field near a hyperbolic closed orbit will result in a topologically equivalent vector field.

And conversely: if any closed orbit of an arbitrary vector field is locally structurally stable, then it must be hyperbolic.

(Note how this differs from hyperbolicity for stationary points, where the eigenvalues of the appropriate Jacobian matrix JV need to lie **off the imaginary axis**. The unit circle is the image of the imaginary axis under the complex exponential mapping, and this reflects the exponential relationship between vector fields and mappings.)

Once again, the eigenvalues are the roots of a polynomial: in this case the polynomial is det(Jf - λI). Since Jf is a 2×2 matrix, this polynomial is just a quadratic polynomial, with real coefficients.

Hence the roots are either both real, or else complex conjugates. In addition, the product of the eigenvalues must be positive, since Euclidean space is orientable. (This is the 3-dimensional version of the fact that a Möbius band cannot be a subset of the plane.) So if both eigenvalues are real, they must have the same sign. We consider the possibilities :

Hyperbolic Closed orbits in 3D
($Root_1$ and $Root_2$ refer to the two eigenvalues; the numbering is of no significance)

1. Both roots real:
   a. $|Root_1| > 1$, $|Root_2| > 1$:
      source closed orbit
   b. $|Root_1| < 1$, $|Root_2| < 1$:
      sink closed orbit
   c. $|Root_1| > 1$, $|Root_2| < 1$:
      i) both roots positive**:**
      saddle closed orbit
      ii) both roots negative:
      twisted saddle closed orbit[1]
2. Complex conjugate roots:
   a. Both outside the unit circle:
      spiral source closed orbit
   b. Both inside the unit circle:
      spiral sink closed orbit

And for completeness, we also mention the rather simple classification of hyperbolic closed orbits in 2 dimensions, where the Jacobian matrix is only 1×1, and that single element **is** the lone, necessarily real, eigenvalue of Jf:

Hyperbolic Closed orbits in 2D

1. $|Root| > 1$:
   source closed orbit
2. $|Root| < 1$:
   sink closed orbit

Intrinsic vs. Extrinsic Topology:

Topology is often called upon to describe a situation where one space X is sitting inside ("embedded in") another space Y. In this case, a distinction is made between properties that depend only on X, and those that come about from **the way in which** X **is positioned inside** Y.

Those properties depending only on X are called **intrinsic**, and those depending on

1. An additional complexity is possible in the case of saddle closed orbits which cannot occur for saddle stationary points: twisting . This is detectable by the sign of the eigenvalues (which must be the same sign by a previous comment). This tells us that the eigendirections are taken by the Poincaré map to their opposite directions. The only way this can be accomplished by trajectories near a closed orbit is by twisting.

the manner in which X sits inside Y are called **extrinsic**.

Extrinsic Topology of Closed Orbits:

A. Twisting due to the embedding:

In the case of a closed orbit, we may think of X as representing a small neighborhood N(C) of the closed orbit C, and Y as representing Euclidean space R^3.

Whether or not C is a twisted closed orbit, the neighborhood N(C) may or may not be embedded in R^3 with twisting due solely to the choice of embedding. This **extrinsic** twisting can occur in any integer multiple of 360° twists.

B. Knotting:

Any simple closed curve in R^3 (which any closed orbit must be!) is intrinsically topologically equivalent to any other simple closed curve, purely as a topological space.

From the extrinsic point of view, however, there are many ways in which a simple closed curve can be embedded in R^3. Two simple closed curves C_1 and C_2 in R^3 are considered to have topologically equivalent embeddings if there is a homeomorphism from R^3 to itself that carries C_1 onto C_2. (Technically, the situation is slightly more complicated than this, but not enough to be worth mentioning.)

The resulting equivalence classes are called **knot types**. The knot type known as the **unknot** is the equivalence class of the unit circle in the xy-plane. Any simple closed curve in the knot type of the unknot **unknotted**.

The problem of classifying knots is a very difficult one; much progress has been made, but much remains to be understood. As an example of the difficulty of the subject, we cite one result that was conjectured for a long time, but proved only a couple of years ago:
Theorem: Two simple closed curves in R^3 are of the same knot type if and only if their complementary regions in R^3 are toplogically equivalent spaces.

Anything that a simple closed curve can do in R^3, a closed orbit of a vector field can potentially do as well. So any knot type can occur as the knot type of a closed orbit in some vector field.

C. Linking:

Any collection of two or more disjoint simple closed curves in R^3 may exhibit linking, another extrinsic property. Linking is really just the generalization of knotting appropriate to a finite union of simple closed curves.

Analogous to knotting, we define the collection $\{C_1,...,C_k\}$ of disjoint simple closed curves in R^3 to be **link equivalent** to another such collection $\{D_1,...,D_k\}$ if there is a homeomorphism of R^3 to itself that carries the union $C_1 \cup ... \cup C_k$ onto the union $D_1 \cup ... \cup D_k$. (This implies that, in fact, each C_i is carried onto one of the D_j.) As with knotting, an equivalence class of links is called a **link type**.

Any collection of closed orbits of a vector field is a collection of disjoint simple closed curves in R^3, and as such represents a link type. Much is known, but much remains unknown, about what link types can constitute the set of source closed orbits, saddle closed orbits, and sink closed orbits of a vector field defined on R^3 or a subset of R^3.

As an example of the intriguing state of ignorance of this situation, we mention a famous unsolved problem known as the Seifert conjecture, posed over 40 years ago:

The Seifert conjecture:

Suppose V is a non-singular smooth (i.e., C^∞) vector field defined on a solid torus (anchor ring) shaped region S in R3. Further suppose that V is pointing inward normal on the torus boundary of S.
Question: Must V necessarily have any closed orbits at all?

The first examples of such a vector field that one comes up with usually have a closed orbit that goes around the hole in the solid torus S. Surprisingly enough, there are examples of C^∞ vector fields satisfying the premises of Seifert's Conjecture, but whose closed orbits do **not** go around the hole in S. In addition, the answer is known to be false if the differentiability condition is relaxed to only C^1 or C^2.

Addendum: In late 1993 the C^∞ Seifert conjecture was found to be false when Krystyna Kuperberg of Auburn University announced a counterexample, which as of March 1994 appears correct.

15 Stable Manifolds

If H represents a hyperbolic stationary point or closed orbit of a vector field, there are important topological features lurking nearby: the so-called stable and unstable manifolds of H.

For convenience we make the

Definition: A **critical element** of a vector field V is any stationary point or closed orbit of V.

The stable manifold is the set of all points flowing into H in positive time; the unstable manifold is the set of all points flowing into H in negative time.

Of particular interest is how the stable manifold of one critical element intersects the unstable manifold of another, or possibly of itself. (It will be left to the reader to understand why two different stable manifolds cannot intersect each other.)

Stable manifolds of hyperbolic stationary points:

We now consider another important topological features of any hyperbolic critical element: the stable and unstable manifolds. Assume there is a vector field V defined on some open set U in R^2 or R^3.

Definition: Let p be a hyperbolic stationary point of the vector field V. The **stable manifold** of the point p, denoted by $W^s(p)$, is the set of points whose trajectories approach p as $t \to \infty$. The **unstable manifold** of the point p, denoted by

$W^u(p)$, is defined as the stable manifold of p for the reversed flow. (In other words, it is the set of points x whose trajectories approach p as $t \to -\infty$.)

In all cases the stable manifold of a hyperbolic stationary point is topologically equivalent to a Euclidean space. This copy of R^k is mapped into the domain of the vector field C^1 (as long as V is) and one-to-one.

In 2 dimensions, a source p has a 0-dimensional stable manifold: the point p itself. If p is a saddle, then p has a 1-dimensional stable manifold, consisting of p and the two orbits along the negative eigendirection. If p is a sink, the stable manifold is 2-dimensional, consisting at least of all the points sufficiently close to p, and all the points that eventually flow into them.

In 3 dimensions, a source p has a 0-dimensional stable manifold. A saddle of type (2 out, 1 in) has a 1-dimensional stable manifold: again it will be the point p and the precisely two orbits along the negative eigendirection. A saddle of type (1 out, 2 in) has a 2-dimensional stable manifold (i.e., it is a surface), consisting of the point p and the trajectories which get arbitrarily close to the negative real part eigenplane as $t \to \infty$. Finally, a sink p has a 3-dimensional stable manifold.

The term **n-manifold** in topology means a space each point of which has a neighborhood that is topologically equivalent to an open neighborhood in Euclidean space R^n. In brief, a manifold is a space that is "locally Euclidean" of a fixed dimension n.

A point is a 0-manifold. A curve is a 1-manifold. A surface is a 2-manifold. And a 3-manifold, if it is a subset of R^3, is any open subset of R^3.

As the terminology implies, any stable manifold is, in fact, an n-manifold for some n. The n in question, as the above examples suggest, is the number of eigenvalues having negative real part.

It follows from the definition that **a stable or unstable manifold is invariant under the action of the flow**. For this reason they are sometimes collectively called "invariant manifolds." However, since there are many other manifolds that are invariant under the flow, we prefer not to use this terminology.

Stable manifolds of hyperbolic closed orbits:

Definitions: Let C be a hyperbolic closed orbit of the vector field V. The **stable manifold** of C, denoted by $W^s(C)$, is the set of points x whose trajectories approach C as $t \to \infty$. (Here the word "approach" indicates simply that the distance of $\phi_t(x)$ to C gets arbitrarily small as t gets large.) The **unstable manifold** of C, denoted by $W^u(C)$, is defined as the stable manifold of C for the reversed flow. (In other words, it is the set of points x whose trajectories approach C as $t \to -\infty$.)

The stable manifold of a hyperbolic closed orbit is topologically equivalent to either the ordinary or twisted product of the circle and a Euclidean space. If the stable manifold is 2-dimensional, for example,

these are the cylinder and the Möbius band, respectively. As in the case of a stationary point, these stable manifolds are also mapped C^1 and one-to-one into the domain of the vector field.

In 3 dimensions, a source closed orbit C has a 1-dimensional stable manifold: just itself. A saddle closed orbit has a 2-dimensional stable manifold. A sink closed orbit has a 3-dimensional stable manifold.

Stable manifolds that are 1- or 2-dimensional look like ordinary curves or surfaces, respectively, near their associated closed orbit or stationary point. However, as they get farther away they get wound up in the dynamics of other parts of the vector field, such as other closed orbits or stationary points. This can cause them to be curly beyond belief.

16 Transversality

Suppose that a line lies in a plane in R^3. An arbitrarily small perturbation of the line or the plane or both can (and usually will) change this situation: after the perturbation, the line will almost certainly intersect the plane in one point.

Now suppose instead that a line intersects a plane in just one point. In this case a sufficiently small perturbation of the line and/or plane **cannot** alter this state of affairs.

Suppose two manifolds M_1 and M_2 in Euclidean space R^n intersect at a point p. Let T_i denote the set of all vectors tangent to M_i at p, for i = 1,2. **Now suppose that every vector in R^n is the sum of a vector in T_1 and a vector in T_2.** Then M_1 and M_2 are said to **intersect transversely at the point** p.

If for **every** point p in $M_1 \cap M_2$, it is the case that M_1 and M_2 intersect transversely at p, then M_1 and M_2 are said simply to **intersect transversely.** Note that if $M_1 \cap M_2$ is empty, then M_1 and M_2 necessarily intersect transversely.

When two manifolds of dimensions d1 and d2 intersect transversely in R^n, their intersection N is again a manifold. If d1 + d2 < n, then in fact N must be empty. Otherwise, N will be a manifold of dimension equal to d1 + d2 - n.

When two manifolds intersect transversely, then a sufficiently small perturbation of each of them will not change this fact. Like hyperbolicity, transversality is a property which will not be washed away by small errors of measurement.

And as in the case of hyperbolicity, if two manifolds do not intersect transversely, then they can be made to do so after an arbitrarily small perturbation.

17 The Non-Wandering Set

For an arbitrary vector field V, we can define the set of points that "recirculate," in a sense that can be made precise.

Definition: We say that a point is **wandering** if every neighborhood of the point, after flowing for a sufficiently large amount of time, never intersects itself again.

Precisely: p is wandering point for V if there exists a neighborhood N of p and a $T > 0$ such that, for all t with $|t| \geq T$, we have $\phi_t(N) \cap N = \varnothing$.

Definition: The **non-wandering set** $\Omega(V)$ of a vector field V consists of all points that are not wandering.

The non-wandering set must be a closed set invariant under the action of the flow.

Any critical element of a vector field is necessarily contained in its non-wandering set.

Example: For any $r > 0$ consider the constant vector field V_r on the unit square $0 \leq x \leq 2\pi$, $0 \leq y \leq 2\pi$, given by $V_r(x, y) \equiv (1, r^2)$. If we identify opposite sides of this square via $(x,0) \sim (x,2\pi)$ and $(0,y) \sim (2\pi,y)$, we get a torus with a vector field on it. Finally we can fill up a solid anchor ring A in R^3 by nesting an infinite family of such tori of revolution—one for each r satisfying $0 < r < 1/2$—about a core circle. We inscribe the vector field V_r on the torus of small radius r. Then the non-wandering set of this vector field on A is all of A.

Theorem: Suppose the vector field V is **volume-preserving on a bounded domain** U and has **no trajectories exiting its domain.** Then $\Omega(V)$ must be all of U.

Sketch of proof: Suppose, to the contrary, that there is some point p of U and some neighborhood N of p such that for all $|t| \geq T > 0$, we have $N \cap \phi_t(N) = \varnothing$.

Now, if there were two numbers s1 and s2 with $s2 \geq s1 + T$, such that $\phi_{s1}(N) \cap \phi_{s2}(N) \neq \varnothing$, then by applying ϕ_{-s1} to both sides of this expression we would obtain $N \cap \phi_{s2-s1}(N) \neq \varnothing$—contradiction.

Thus whenever t increases by at least T, $\phi_t(N)$ must occupy a new region of U with equal positive volume, which eventually will exhaust the necessarily finite volume of U—contradiction. ■

18 Structural Stability

If we are ever going to be able to accurately describe the topology of a real-world vector field, then the necessarily inaccurate numerical definition we have of the vector field must not affect the topology. We are thus interested in vector fields which keep their global topology even after a sufficiently small perturbation.

This is exactly the consideration which led to our giving prominence to hyperbolicity, and transversality, except this is not just on a local scale: it is on a global scale.

Definition: A vector field V on U is called **structurally stable** if for any new vector field V_1 on U that is sufficiently near to V (in the C^1 sense), there exists a homeomorphism $H : U \rightarrow U$ which carries the trajectories of V to those of V_1 and preserves their sense.

In brief, a small enough perturbation of V results in a new vector field that is topologically equivalent to the old one.

Note: It might be nice to be able to say that the old and new vector fields were

smoothly equivalent (i.e., topologically equivalent via a homeomorphism that is smooth and has smooth inverse). However, smooth equivalence preserves eigenvalues at stationary points, but eigenvalues can be easily changed via a C^1 perturbation.

Example: Consider the constant vector field $V(x, y) = (a, b)$ on R^2. This is structurally stable, as you might guess.

Example: Consider the vector field given by
$V(x,y) = (-y + x(x^2+y^2-1), x + y(x^2+y^2-1))$.
This vector field has one sink stationary point, and one source closed orbit. It, too, is structurally stable.

19 Morse-Smale Vector Fields

Morse-Smale vector fields are important because they are the simplest large class of structurally stable vector fields.

Definition: A vector field V is called **Morse-Smale** if the following hold:
1) V has a finite number of stationary points and closed orbits, all of which are hyperbolic.
2) The non-wandering set $\Omega(V)$ is the union of the closed orbits and the stationary points.
3) If H_1 and H_2 each represent a closed orbit or stationary point, then the stable manifold $W^s(H_1)$ of H_1 has transverse intersection with the unstable manifold $W^u(H_2)$ of H_2.

Theorem: Morse-Smale vector fields are structurally stable.

Theorem: In 2 dimensions**, every** structurally stable vector field is Morse-Smale.

It would be convenient if this were also the case in higher dimensions, but alas, this is not the case. There exist structurally stable vector fields in dimensions ≥ 3 whose non-wandering set is the entire domain (the so-called **Anosov** flows).

Decomposing Morse-Smale vector fields into handles

It would be rather complicated to describe exactly how all the stationary points and closed orbits of a Morse-Smale vector field in n dimensions (think of n = 2 or 3) fit together. But it is a remarkable fact that Morse-Smale vector fields can be decomposed into simple pieces, each of which corresponds to a critical element.

Let D^j denote the unit ball in R^j: all points at a distance of ≤ 1 from the origin. Let S^1 denote the unit circle in R^2: all points at a distance of exactly 1 from the origin.

For each stationary point with unstable manifold of dimension k, we will use one building block of the form $D^k \times D^{n-k}$ (a "handle"). For each closed orbit with unstable manifold of dimension k, we will use one building block of the form $S^1 \times D^{k-1} \times D^{n-k}$ (a "round handle"). By taking the union of these building blocks, which are allowed to intersect only along specified parts of their boundaries, we can reconstruct—in a topological sense—the original Morse-Smale vector field.

The partial ordering

Consider a Morse-Smale vector field V.

Observation 1: Let H_1, H_2, and H_3 be critical elements of V. Suppose that $W^u(H_1) \cap W^s(H_2) \neq \emptyset$, and also $W^u(H_2) \cap W^s(H_3) \neq \emptyset$. Then it must also be true that $W^u(H_1) \cap W^s(H_3) \neq \emptyset$. So there is a kind of transitivity operating here.

Observation 2: Suppose instead that we have critical elements H_1 and H_2 of V with $W^u(H_1) \cap W^s(H_2) \neq \emptyset$, and also $W^u(H_2) \cap W^s(H_1) \neq \emptyset$. Then it must follow that $H_1 = H_2$.

Now let H_1 and H_2 be any two critical elements of a Morse-Smale vector field V. Then we shall use the notation $H_1 \leq H_2$ to mean that $W^u(H_1) \cap W^s(H_2) \neq \emptyset$.

From the two observations above, we see that the condition that $\leq$ is a bona fide partial ordering on the set of critical elements of V.

Describing a Morse-Smale vector field by
A. labeling the critical elements according to
1) **stationary point or closed orbit**
2) **dimension of the stable manifold**
3) **if closed orbit, whether or not stable manifold is twisted, and**

B. specifiying the pairs H1 and H2 for which H1 ≤ H2

goes a long way towards conveying a very good description of "what the vector field V is doing topologically."

(Note, however, that there is still additional topological information—**how** the stable and unstable manifolds intersect each other—that is not covered in the above description.)

20 Beyond Morse-Smale Flows

Morse-Smale flows are important to understand, but they represent one extreme of a spectrum. The non-wandering set of a Morse-Smale vector field is concentrated in a finite number of 0- and 1-dimensional sets: the stationary points and closed orbits.

At the opposite extreme are the **Anosov flows** mentioned in Section 19. The non-wandering set of an Anosov flow is the entire domain of the flow. And there is a kind of hyperbolicity which applies to this entire non-wandering set: each point has a set of stable directions emanating from it, on which the flow is "contracting," and a set of unstable directions on which the flow is "expanding."

And the concept of hyperbolicity can be extended to include a much wider variety of sets than manifolds. When this is done properly, the resulting flows are, like Morse-Smale flows, structurally stable. Much progress has been made, but as of late 1992 such so-called **basic sets** are far from being completely classified.

21 Further Areas of Study

One fiction we have assumed to avoid complications in these notes, is that the domain of a vector field must be an open set in Euclidean space. But in fact, the

natural domain for a vector field is an arbitrary n-dimensional manifold.

On the one hand, the mathematics is usually most elegant if the domain of a vector field is a compact manifold without boundary, and if the vector field is time-independent (like the ones we have been considering for most of these notes).

But in reality, domains for real-world vector fields are most commonly manifolds with boundary, and vector fields found in the real world are frequently time-dependent.

Much in the same way that the Poincaré map can be used to understand the orbit structure around a closed orbit, the dynamics of mappings in general sheds a great deal of light on the dynamics of flows, and is well worthy of further study.

Very little is known about what kind of dynamics real-world vector fields tend to actually have, for example in aerodynamics. Valuable progress could be made by creating software to detect topological features of real-world vector fields, and reporting on the kinds of fields encountered.

Much progress could also be made by devising new, rough descriptions of how a vector field flows around, computable with a modern computer, and which are robust enough to be accurate despite numerical imprecision.

We shall see.