Lectures on Periodic Orbits

11 February 2009

Most of the contents of these notes can found in any typical text on dynamical systems, most notably Strogatz [1994], Perko [2001] and Verhulst [1996]. Complete proofs have been omitted and wherever possible, references to the literature have been given instead. As these notes are being written during the term, they will be frequently updated. Week-by-week changes are detailed on the course web page.

1 Ruling Out Periodic Orbits

Gradient Systems. A **gradient system** is a dynamical system of the form

$$\dot{x} = -\nabla V(x) \tag{1.1}$$

for a given function $V(x)$ in $\mathbb{R}^n$.

Theorem 1.1. *Gradient systems cannot have periodic orbits.*

Proof. Suppose to the contrary that $\gamma : t \mapsto x(t)$ is a periodic orbit of the gradient system (1.1) with period T. Then $V(x(T)) - V(x(0)) = 0$, but on the other hand

$$V(x(T)) - V(x(0)) = \int_0^T \frac{dV}{dt}\, dt = \int_0^T \nabla V \cdot \dot{x}\, dt = -\int_0^T \|\dot{x}\|^2\, dt < 0,$$

a contradiction. ■

Note that this result is valid for gradient systems in arbitrary dimensions, in contrast to the subsequent results, which are specific for the plane.

Dulac's Criterion. Recall that a region R of the plane is said to be **simply connected** if every closed loop within R can be shrunk to a point without leaving R (intuitively, R contains no holes).

Dulac's criterion gives sufficient conditions for the non-existence of periodic orbits of dynamical systems in simply connected regions of the plane. The downside of this method is that it depends on the choice of an appropriate multiplier, which might be hard to find.

Theorem 1.2. *Let R be a simply connected region in $\mathbb{R}^2$ and consider a planar dynamical system in R given by*

$$\dot{x} = f(x, y) \quad \textit{and} \quad \dot{y} = g(x, y), \tag{1.2}$$

where f, g are C^1 functions in R. Suppose that there exists a C^1 function $h(x, y)$ in R so that

$$\nabla \cdot h(fe_x + ge_y)$$

has a definite sign in R. Then the dynamical system (1.2) cannot have any periodic orbits in R.

Proof. Assume that (1.2) has a periodic orbit γ contained in R and let A be the area enclosed by γ. Since R is simply connected, A lies entirely in R.

By Green's theorem, we have

$$\iint_A \nabla \cdot h(fe_x + ge_y)\, dxdy = \oint_\gamma h(fe_x + ge_y) \cdot n dl,$$

where n is the outward normal to γ.

Now, note that the left-hand side of this expression is different from zero because of the sign-definiteness. However, the right-hand side is zero (since $fe_x + ge_y$ is tangent to γ), a contradiction. ■

A special case of this theorem, where the multiplier $h = 1$, is known as **Bendixson's criterion**. When asked to verify whether a vector field can have periodic orbits, the first thing to check is whether Bendixson's criterion is applicable: checking that the divergence of the vector field is sign definite can be a quick and straightforward way to rule out periodic orbits. When this fails, Dulac's criterion or other techniques can be used.

Example 1. Show that the following dynamical system does not have any closed orbits in the region $x, y > 0$:

$$\dot{x} = x(2 - x - y) \quad \text{and} \quad \dot{y} = y(4x - x^2 - 3).$$

By using the multiplier $h = \frac{1}{xy}$, we find that the divergence is given by $-1/y$, which is sign-definite in the region considered.

Lyapunov-like Functions. If a dynamical system $\dot{x} = X(x)$ has a Lyapunov-like function $V(x)$, *i.e.* $V(x)$ is monotonically decreasing along trajectories, then it is easy to see that there cannot be any periodic orbits. Indeed, assume that there is a periodic orbit with period T, then

$$0 = V(x(T)) - V(x(0)) = \int_0^T \frac{dV}{dt} dt,$$

but the right-hand side is by assumption strictly smaller than zero, a contradiction.

Like always, there is no explicit prescription for finding a Lyapunov function and in some examples, "Lyapunov-like" functions may be useful; see Strogatz [1994], example 7.2.2.

2 Index Theory

Refer to §7.3 of Strogatz [1994].

We covered the intuitive definition of the **index of a curve** as the number of times the vector field rotates in counterclockwise fashion when traversing the curve once. The index satisfies a number of elementary topological properties:

1. When C and C' are two closed curves that can be continuously deformed into each other (without moving the curve over one of the fixed points of the vector field), then $I_C = I_{C'}$.
2. If C does not enclose any fixed points, then $I_C = 0$.
3. The index is unchanged when the direction of the vector field is reversed.
4. If C is a *trajectory* of the system rather than an arbitrary closed curve, then $I_C = +1$.

As a result, we define the **index $I_{\bar{x}}$ of a fixed point** $\bar{x}$ by taking an arbitrary closed curve C around $\bar{x}$ such that C does not enclose any other fixed points, and we put $I_{\bar{x}} = I_C$. By property 1 above, $I_{\bar{x}}$ is independent of the choice of curve. Note that the index of a fixed point does not distinguish the stability types of a fixed point — sources and sinks both have index +1.

Lemma 2.1 (Thm. 6.8.1 in Strogatz [1994])**.** *If a closed curve C encloses n fixed points $\bar{x}_1, \ldots, \bar{x}_n$, then*

$$I_C = I_{\bar{x}_1} + \cdots + I_{\bar{x}_n}.$$

Some immediately useful consequences are

- Any closed orbit in the plane has to enclose fixed points whose total index sums to +1.
- In particular, if the vector field has no fixed points, there cannot be any periodic orbits.
- If the system has only one fixed point, it cannot be hyperbolic (since the index of a hyperbolic fixed point is minus one).

These conclusions can be used to rule out the existence of closed orbits, or to provide qualitative information about fixed points and closed orbits. See examples 6.8.5 and 6.8.6 in Strogatz [1994].

3 The Poincaré-Bendixson Theorem

The results in the preceding sections are all concerned with the non-existence of periodic orbits, or (in the case of index theory) with the consequences of having a given periodic orbit. The Poincaré-Bendixson theorem on the other hand gives sufficient conditions for the *existence* of a periodic orbit. There are many versions of this theorem and the proof (which we omit from these lectures) is reasonably complicated.

3.1 Introductory Definitions

Although the following concepts are strictly speaking not needed for the formulation of the Poincaré-Bendixson theorem, they are sufficiently fundamental for us to list them here anyway.

Let $\dot{x} = X(x)$ be a dynamical system and consider a trajectory $\gamma : t \mapsto x(t)$. A **positive limit point** is a point x for which there exists a sequence $t_1, t_2, \ldots \to +\infty$ such that

$$\lim_{n\to\infty} x(t_n) = x.$$

The concept of a **negative limit point** is defined similarly.

The ω**-limit set** of a trajectory γ, denoted by $\omega(\gamma)$, is defined as the set of all positive limit points of that orbit. Similarly, the α-limit set is the set of all negative limit points.

Example 2. Consider the system $\dot{x} = -x$, $\dot{y} = -2y$. The origin is a positive limit point for all orbits and there are no other positive limit points. The ω-limit set is hence the origin for all orbits. The α-limit set is empty for all orbits except the origin.

Example 3. Consider the harmonic oscillator equations $\dot{x} = y$, $\dot{y} = -x$. Every point on an orbit is simultaneously a positive and a negative limit point for that orbit.

The importance of the ω-limit set lies in the fact that trajectories in a bounded region of the plane will spiral inward to ω-limit set. Essentially, the Poincaré-Bendixson theorem tells us that $\omega(\gamma)$ will either contain a fixed point or it will be a closed orbit of the flow. The following lemma gives us a hint as to why this should be so (this is far from a complete proof, which is very complicated).

Theorem 3.1 (thm. 4.2 in Verhulst [1996])**.** *The sets $\alpha(\gamma)$ and $\omega(\gamma)$ are closed and invariant. If γ^+ is bounded, then $\omega(\gamma)$ is compact, connected and non-empty. Moreover $d(x(t), \omega(\gamma))$ goes to zero as $t \to +\infty$.*

Using the concept of the ω- and α-limit set, we can now define the concept of limit cycles rigorously.

Definition 3.2 (def. 3.2 in Perko [2001])**.** *A **limit cycle** γ of a dynamical system in the plane is a periodic orbit which is the α or ω-limit set of a trajectory γ' other than γ. If a limit cycle γ is the ω-limit set of every other trajectory in a neighborhood of γ, γ is said to be an ω-**limit cycle** or **stable limit cycle**. Likewise, if γ is the α-limit set of neighboring trajectories, γ is said to be an α-**limit cycle** or **unstable limit cycle**.*

Limit cycles are hence a truly nonlinear feature. Strogatz [1994] defines limit cycles informally as closed orbits such that nearby trajectories are either attracted to or repelled by the limit cycle: by virtue of the theorem above, we see that this definition corresponds exactly to either an ω- or an α-limit cycle.

Note that the condition that a limit cycle is the ω-limit set of a *different* trajectory is crucial. Otherwise, the trajectories of the simple oscillator (example 3) would be limit cycles, which does not capture the idea that a limit cycle is informally “where trajectories end up”.

Example 4. The dynamical system given in polar coordinates by $\dot{r} = r(1 - r^2)$, $\dot{\theta} = 1$ has exactly one stable limit cycle, given by $r = 1$. All trajectories, except for the fixed point $r = 0$, are attracted to this limit cycle. See figure 3.1.

3.2 Statement of the Theorem

Theorem 3.3 (Poincaré-Bendixson)**.** *Let R be a region of the plane which is closed and bounded. Consider a dynamical system $\dot{x} = X(x)$ in R where the vector field X is at least C^1. Assume that R contains no fixed points of X. Assume furthermore that there exists a trajectory γ of X starting in R which stays in R for all future times.*

Then, either γ is a closed orbit or γ asymptotically approaches a closed orbit; in other words, there exists a limit cycle in R.

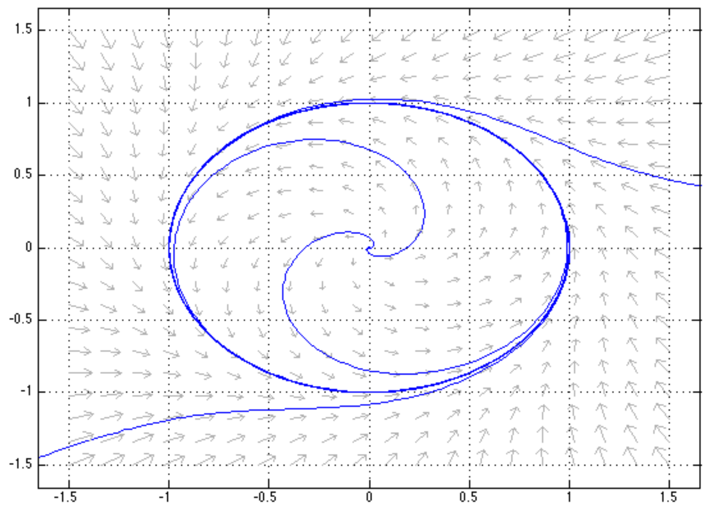

Figure 3.1: A system with a stable limit cycle.

More formally, let γ^+ be the forward part of γ. Then if γ^+ is contained in a compact subset of the plane which contains no fixed points, $\omega(\gamma)$ will be the desired periodic orbit. This is how the Poincaré-Bendixson theorem is stated in [Perko, 2001, Sec. 3.7].

The hard part of the Poincaré-Bendixson theorem consists of finding a suitable trajectory γ. However, there are a number of special cases in which this becomes easier in practice: define a **trapping region** to be any region R of the plane which is positively invariant under the flow ϕ_t of X: $\phi_t(R) \subset R$ for all $t > 0$. It follows that if R is a trapping region, then every trajectory of X starting in R stays in R for all future times. To check that R is a trapping region it is sufficient to verify that on the boundary of R, X is everywhere pointing inward.

Practically speaking, this is how one proceeds to apply the Poincaré-Bendixson theorem: by constructing an (annular) region R in the plane so that on the boundary of R, the vector field points into R.

Note that the region R will always be annular. In other words, R will contain at least one hole (topologically speaking, R is not simply connected). This can easily be seen through index theory. Denote the outer boundary of R by C: since the vector field is pointing inward on C, the index I_C will be equal to $+1$. Hence C has to enclose a given number of fixed points with total index $+1$. These fixed points have to be excised from R in order for the Poincaré-Bendixson theorem to be applicable, leaving us with a non-simply connected domain.

Remark 1. Note that there exists a "time-reversed" version of the Poincaré-Bendixson theorem, in the sense that if the backward trajectory γ^-, defined as

$$\gamma^- = \{(x(t), y(t)) : t \leq 0\}$$

is contained in a compact subset R of the plane which contains no fixed points, then there exists a limit cycle in R. More specifically, the desired limit cycle will be $\alpha(\gamma)$. This follows from the usual Poincaré-Bendixson by reversing time in the original differential equation. The result is that γ^- turns into γ^+ and $\alpha(\gamma)$ into $\omega(\gamma)$.

If one can find a **repelling region**,[1] *i.e.* an annular region such that the vector field points *outwards* of this region, then this version of the Poincaré-Bendixson theorem may be applied.

4 Liénard Equations

So far we have discussed existence results for limit cycles only. Deciding on the number of limit cycle and their relative location, however, is a wholly different matter. Indeed, this is exactly Hilbert's 16th problem, which is unresolved until this day.

Neverthelesss, for certain classes of differential equations, sharper results than just existence can be obtained. One such class are the so-called Liénard equations, for which exactly one stable limit cycle exists. Liénard equations arose in the context of circuit theory and radio-transmitters, where having a stable regime of oscillations is very important.

We refer to any differential equation of the form

$$\ddot{x} + f(x)\dot{x} + g(x) = 0 \tag{4.1}$$

as a **Liénard equation**. Here, f, g are unknown functions. For there to exist exactly one stable limit cycle, f and g have to satisfy a number of conditions.

Theorem 4.1. *Suppose the following conditions hold:*

1. f, g *are* C^1 *functions;*
2. $g(x)$ *is an odd function of* x*;*
3. $g(x) > 0$ *for* $x > 0$*;*
4. *the primitive* $F(x) = \int_0^x f(u)du$ *satisfies*
   (a) $F(x)$ *has exactly one positive root at* $x = a$*;*
   (b) $F(x) < 0$ *for* $0 < x < a$*;*
   (c) $F(x)$ *is positive and non-decreasing for* $x > a$*;*
   (d) $F(x) \to +\infty$ *for* $x \to +\infty$.

Then the Liénard equation (4.1) has exactly one stable limit cycle.

In class, we proved this theorem for the **van der Pol equation**, where $f(x) = \mu(x^2 - 1)$ with $\mu > 0$, and $g(x) = x$. Equations with similar characteristics arise throughout physics, chemistry and biology; one example is the **Fitzhugh-Nagumo system** describing the spiking of neurons. The existence proof for stable limit cycles of the van der Pol equation can be found in Hirsch et al. [2004], sec. 12.3.

Note that if $g(x) = x$ as in the case of the van der Pol equation, the system can be rewritten as a second-order equation. Kaushik raised an interesting question of whether a second-order system can ever have more than one limit cycle. While the question of determining or even bounding the number of limit cycles of a given set of ODEs is highly intricate, it turns out that by taking $g(x) = x$ and $f(x)$ to be a polynomial of sufficiently high degree, multiple limit cycles may occur. An example is due to Zhang Perko [2001]: take

$$F(x) = \frac{8}{25}x^5 - \frac{4}{3}x^3 + \frac{4}{5}x.$$

[1] This terminology is not standard.

The Liénard equation associated to this system has two limit cycles; see figure 4.1. Obviously this equation does not satify the hypotheses of theorem 4.1 as $F(x)$ has multiple positive roots.

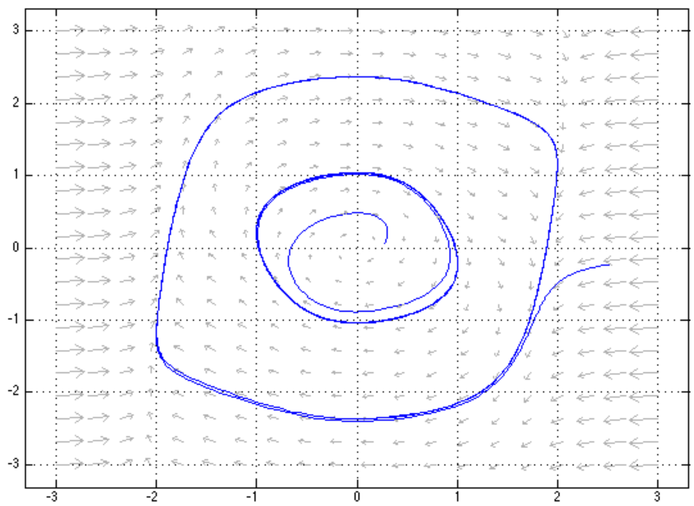

Figure 4.1: Zhang's example of a system with two limit cycles. The inner limit cycle is unstable, which can inferred from the fact that the origin is a source. The outer limit cycle is stable.

5 Quantitative Methods of Finding Periodic Orbits

Often the conclusions of the Poincaré-Bendixson, that there exists a periodic orbit in a given region, are somewhat lacking and we need an explicit expression for the periodic orbit. Series expansions provide one way of obtaining such expressions, but one should take care in interpreting them. In this section, we will look at periodic orbits of dynamical systems in the plane of the form

$$\ddot{x} + x + \epsilon h(x, \dot{x}) = 0, \tag{5.1}$$

where $\epsilon \geq 0$ is a small parameter. For $\epsilon = 0$, this is just the harmonic oscillator. What happens to the periodic orbits of the harmonic oscillator when $\epsilon > 0$?

Naive Taylor Expansions. One way of approaching this question is by expanding the solution $x_\epsilon(t)$ of (5.1) in powers of ϵ:

$$x_\epsilon(t) = x_0(t) + \epsilon x_1(t) + \epsilon^2 x_2(t) + \cdots .$$

This approach can and does run into difficulties, however: Strogatz considers the example where $h(x, \dot{x}) = 2\dot{x}$ and initial conditions are given by $x(0) = 0$, $\dot{x}(0) = 1$. This example is explicitly solvable; see Strogatz [1994] for details. One special feature of the solution is that it goes to zero exponentially fast, since the added term represents a small damping term.

From a perturbation point of view, $x_0(t)$ is easily seen to be equal to $\sin t$, and the equation for $x_1(t)$ is given by $\ddot{x}_1 + x_1 = -2\cos t$. This is an equation for a harmonic oscillator with resonant forcing: the right-hand side pumps energy into the system. The solution is given by $x_1(t) = -t \sin t$: this is a disaster since $x_1(t)$ goes off to infinity, in contrast to the exact solution, which stays bounded. Terms in the perturbation expansion like $x_1(t)$, which grow without bound, are termed **secular**.

What is the problem with all this? Mathematically speaking, $x_0(t) + \epsilon x_1(t)$ represents, for *fixed* t, the first order Taylor expansion to $x_\epsilon(t)$, and by calculating higher order contributions, we get progressively better approximations. From a dynamical systems point of view, we have no interest in fixing t of course. However, because of the secular nature of the second term, the expansion is then only accurate for times $t < \epsilon^{-1}$. See figure 7.6.2 in Strogatz [1994].

This example sets the tone for the rest of this section: even by introducing just a small friction term, we already trigger qualitatively new phenomena.

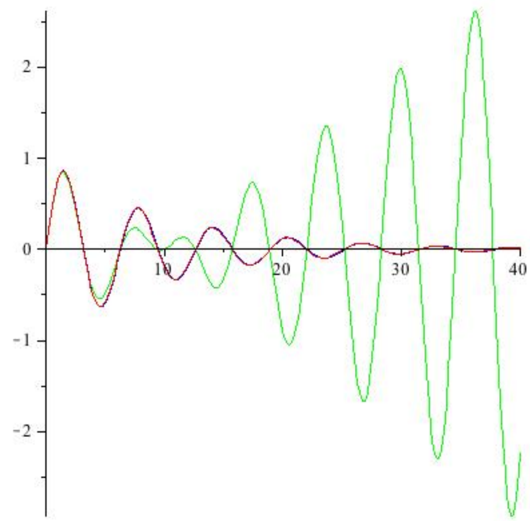

Figure 5.1: The true solution (damped) compared with its first-order Taylor expansion for $\epsilon = 0.1$. Observe that both are in good agreement for $t \ll \epsilon^{-1} = 10$.

Averaging. Another way of dealing with resonant terms is through the method of averaging, also referred to as Lagrange's method. A very thorough outline of this method can be obtained from Strogatz [1994], section 7.6. For periodic orbits, the first order approximation obtained with this method can be shown to stay within an $O(\epsilon)$ range of the true solution *for all times.*

Poincaré-Lindstedt Theory. What is needed is a way to eliminate these secular terms. There are several related methods to do this: here we use the Poincaré-Lindstedt method. This method is only applicable to approximating periodic orbits and consists of proposing a perturbation expansion for *both* the periodic orbit $x_\epsilon(t)$ and its frequency ω_ϵ.

1. Introduce a new time scale $\tau = \omega t$ so that the new period becomes 2π.
2. Substitute series expansions for

$$x_\epsilon(\tau) = x_0(\tau) + \epsilon x_1(\tau) + \cdots \quad \text{and} \quad \omega_\epsilon = \omega_0 + \epsilon\omega_1 + \cdots$$

into the equation. Note that $\omega_0 = 1$ since the solution has period 2π when $\epsilon = 0$. Substitute the same expansions into the initial conditions and find the resulting initial conditions for $x_i(t)$.

3. Collect terms of the same order in ϵ and solve the resulting equations for $x_i(t)$. Use the freedom in choosing the coefficients ω_i to kill off any secular terms.

In this way, we find an expansion to arbitrary high order to the desired periodic orbit of the system. Strogatz (exercise 7.6.19) illustrates this on an example: the Duffing oscillator $\ddot{x} + x + \epsilon x^3 = 0$ with initial conditions $x(0) = a$ and $\dot{x}(0) = 0$.

1. The new equation becomes $\omega^2 x'' + x + \epsilon x^3 = 0$, where the prime denotes derivation with respect to τ.

2. The zeroth- and first-order equations become

$$x_0'' + x_0 = 0 \quad \text{and} \quad x_1'' + x_1 = -2\omega_1 x_0'' - x_0^3$$

with initial conditions $x_0(0) = a$ and $x_1(0) = 0$, as well as $\dot{x}_0(0) = \dot{x}_1(0) = 0$. So $x_0(\tau) = a\cos\tau$ and the equation for x_1 becomes

$$x_1'' + x_1 = (2\omega_1 - \frac{3}{4}a^2)a\cos\tau - \frac{1}{4}a^3\cos 3\tau.$$

3. To eliminate secular terms we must choose $\omega_1 = \frac{3}{8}a^2$. The resulting equation for x_1 can now be solved. In this way we can proceed order by order. As in the case of the averaging method, we will usually not be interested in these higher-order contributions. Just the fact that we've eliminated secular terms in the equation for x_1 already yields important information: the first approximation x_0 is given by

$$x_0(t) = a\cos\omega t = a\cos(1 + \epsilon\frac{3}{8}a^2)t.$$

In other words, we get a circular periodic orbit with a nearly-constant frequency. Notice that this periodic orbit is not a limit cycle.

6 Limitations to the Poincaré-Bendixson Theorem

The Poincaré-Bendixson theorem essentially rules out chaos in the plane. This turns out to be a highly non-generic result, however, which does not seem to hold for other configuration spaces or other types of dynamical systems. In this section we explore a few of these possibilities.

Two-dimensional Configuration Spaces. Dynamical systems on two-dimensional manifolds other than the plane may well violate the Poincaré-Bendixson theorem. Consider for instance the following vector field on the torus, which we identify with the unit square in the plane with opposite sides identified:

$$\dot{x} = 1 \quad \text{and} \quad \dot{y} = \pi. \tag{6.1}$$

There is nothing special about the choice of π: any other irrational number would work just as well. Even though the torus is compact and the vector field (6.1) does not have any zeros, the orbits of (6.1) are not periodic: one can check that these orbits densely fill up the torus. This is referred to as **quasi-periodic motion**.

Many systems have a torus (possibly in higher dimensions) as their configuration space. The so-called Control Moment Gyroscopes (CMGs), a control mechanism for satellite reorientation maneuvers, are defined on a two-dimensional torus. Integrable systems, a special subclass of mechanical system including the Euler equations for a free rigid body, can be written in a distinguished coordinate system as the motion of a constant vector field on an n-dimensional torus, where n is the number of degrees of freedom.

Higher Dimensions. In dimension three or higher, orbits may approach a very complicated limit set known as a **strange attractor**, which is characterized by a non-integer dimension and the fact that the dynamics on it are sensitive to initial conditions. In other words, **chaos** occurs. A celebrated example of a strange attractor is the Lorentz attractor.

Maps. There exists nothing even close to the Poincaré-Bendixson theorem for discrete-time dynamical systems, *i.e.* maps. The celebrated Baker's map is a map from the unit square into itself with uncountably many chaotic orbits.

References

Hirsch, M. W., S. Smale, and R. L. Devaney [2004], *Differential equations, dynamical systems, and an introduction to chaos*, volume 60 of *Pure and Applied Mathematics (Amsterdam)*. Elsevier/Academic Press, Amsterdam, second edition.

Perko, L. [2001], *Differential equations and dynamical systems*, volume 7 of *Texts in Applied Mathematics*. Springer-Verlag, New York, third edition.

Strogatz, S. [1994], *Nonlinear Dynamics and Chaos.* Addison-Wesley.

Verhulst, F. [1996], *Nonlinear differential equations and dynamical systems.* Universitext. Springer-Verlag, Berlin, second edition.