TMA4165 Dynamical Systems

Assorted notes on dynamical systems

Harald Hanche-Olsen

https://folk.ntnu.no/hanche/

Abstract. These are supplementary notes for a course on dynamical systems. The notes were first made for the course in 2007. For 2008, those notes were worked into a single document.

The basic text for the course is D. W. Jordan and P. Smith's *Nonlinear Ordinary Differential Equations*. These notes are only intended to fill in some material that is not in the book, or to present a different exposition.

Version 2017-08-24

Chapter 1

Well-posedness for ODEs

This chapter is about well-posedness of the initial-value problem for a system of ordinary differential equations:

$$\begin{aligned} \dot{x}(t) &= f\big(x(t)\big), \\ x(0) &= a. \end{aligned} \tag{1.1}$$

Here $f\colon \Omega \to \mathbb{R}^n$ is a mapping defined on an open set $\Omega \subseteq \mathbb{R}^n$. The initial value a is supposed to belong to Ω, and the unknown function x is to be defined on an open interval containing $t = 0$.

By *well-posedness* of the problem we mean a positive answer to three questions: (1) Does a solution exist? (2) Is the solution unique? (3) Does the solution depend continuously on the data (a and the function f)?

Lipschitz continuity

The answer to the question of well-posedness is in general negative. It turns out that the natural requirement to obtain a well-posed problem is Lipschitz continuity of the righthand side f. The function f is called *Lipschitz continuous* if there exists a finite constant L so that

$$\big|f(x) - f(y)\big| \le L|x - y|, \qquad \text{for all } x, y \in \Omega.$$

The smallest such constant L is called the *Lipschitz constant* for f on Ω.

Lipschitz continuity is not uncommon. For example, assume that f is a C^1 function, by which we mean that its first order partial derivatives exist and are continuous. We write Df for the Jacobian matrix of f:

$$Df = \begin{pmatrix} \frac{\partial f_1}{\partial x_1} & \cdots & \frac{\partial f_1}{\partial x_n} \\ \vdots & \ddots & \cdots \\ \frac{\partial f_n}{\partial x_1} & \cdots & \frac{\partial f_n}{\partial x_n} \end{pmatrix}$$

Then, if the whole line segment $[x, y]$ with end points x and y lies within Ω, we can write

$$f(y) - f(x) = \int_0^1 \frac{d}{dt} f\big((1-t)x + ty\big)\, dt = \int_0^1 Df\big((1-t)x + ty\big)\, dt \cdot (y - x)$$

with the result that, if $\|Df(z)\| \le L$ for all z, (where the norm is the operator norm of the matrix, seen as an operator on $\mathbb{R}^n$), then $\big|f(x) - f(y)\big| \le L|x - y|$.

If f belongs to C^1 then f is *locally Lipschitz continuous*, which means that every point $x \in \Omega$ has a neighbourhood in which f is Lipschitz continuous. Equivalently, f is Lipschitz continuous on every compact (i.e., closed and bounded) subset of its domain.

Uniqueness

1 Theorem. *Assume that f is locally Lipschitz continuous. Then* (1.1) *has at most one solution on any given interval containing* 0.

Proof: Assume that x and y are two solutions of (1.1), and note that

$$x(t) = a + \int_0^t f(x(\tau))\, d\tau, \quad y(t) = a + \int_0^t f(y(\tau))\, d\tau.$$

Put

$$u(t) = \int_0^t \big|f(x(\tau)) - f(y(\tau))\big|\, d\tau,$$

so that $\big|x(t) - y(t)\big| \le u(t)$ for $t > 0$. Further, if f is Lipschitz continuous with Lipschitz constant L, then for $t > 0$ we find

$$\dot{u}(t) = \big|f(x(t)) - f(y(t))\big| \le L\big|x(t) - y(t)\big| \le Lu(t)$$

This implies that $e^{-Lt}u$ is non-increasing. But $u(0) = 0$, so this implies $u(t) = 0$ for $t > 0$, and therefore also $x(t) = y(t)$. (The same argument holds for negative t, by time reversal: If $x(t)$ solves (1.1) then $\tilde{x}(t) = x(-t)$ solves a similar problem with f replaced by $-f$. So if we have uniqueness forward in time, the same must hold backward in time.)

The proof as written requires f to be (globally) Lipschitz continuous, but in fact local Lipschitz continuity is enough: If x and y are two solutions and $x(t_0) = y(t_0)$ for *some* t_0 then use the above proof, with L being the Lipschitz constant of f on some neighbourhood of $x(t_0)$, to conclude $x(t) = y(t)$ of all t in some neighbourhood of t_0, and this implies that $x(t) = y(t)$ for all t. We skip the details, but provide the following hint: Assume that $x(t_1) \ne y(t_1)$ for some $t_1 > 0$. Put $t_0 = \sup\big\{t \in [0, t_1] : x(t) = y(t)\big\}$, and immediately arrive at a contradiction. ∎

Existence

2 Theorem. *Assume that f is Lipschitz continuous on the closed ball $B = \overline{B}_R(a) = \{x\colon |x-a| \le r\}$, where r is some positive number. Let M be the maximum value of $|f|$ on B, and define T by $MT = r$. Then* (1.1) *has a solution defined for $|t| \le T$.*

Proof: Note that (1.1) is equivalent with

$$x(t) = a + \int_0^t f\big(x(\tau)\big)\,d\tau. \tag{1.2}$$

We proceed to find a solution by a procedure known as *Picard iteration*:

$$x_k(t) = \begin{cases} a & k = 0, \\ a + \displaystyle\int_0^t f\big(x_{k-1}(\tau)\big)\,d\tau & k = 1,2,\dots \end{cases}$$

It is easy to prove by induction that $x_k(t) \in B$ for $t \in [-T,T]$. Working a bit harder we can show, again by induction, that

$$\big|x_k(t) - x_{k-1}(t)\big| \le \frac{ML^{k-1}|t|^k}{k!}$$

where L is the Lipschitz constant of f on B. For $k = 1$ the above inequality reduces to the obvious $|\int_0^t f(a)\,d\tau| \le M|t|$. And assuming the inequality holds for some k, compute

$$\begin{aligned} |x_{k+1}(t) - x_k(t)| &= \left|\int_0^t \big(f(x_k(\tau)) - f(x_{k-1}(\tau))\big)\,d\tau\right| \\ &\le \left|\int_0^t \big|f(x_k(\tau)) - f(x_{k-1}(\tau))\big|\,d\tau\right| \le \left|\int_0^t L\big|x_k(\tau) - x_{k-1}(\tau)\big|\,d\tau\right| \\ &\le \left|\int_0^t \frac{ML^k|\tau|^k}{k!}\,d\tau\right| = \frac{ML^k|t|^{k+1}}{(k+1)!}, \end{aligned}$$

which finishes the induction part of the proof.

We shall only need the simpler version

$$|x_k(t) - x_{k-1}(t)| \le \frac{ML^{k-1}T^k}{k!}$$

of the above inequality.[1] Assuming $m \le k_1 < k_2$ we therefore have

$$\big|x_{k_2}(t) - x_{k_1}(t)\big| = \Big|\sum_{k=k_1+1}^{k_2} x_k(t) - x_{k-1}(t)\Big| \le \sum_{k=m}^{\infty} \frac{ML^{k-1}T^k}{k!}.$$

[1] We could hardly manage to prove *that* by induction: The t dependency in the estimate was essential. But when we *use* the inequality, we prefer an estimate independent of t.

Noting that the final sum here (1) does not depend on t, and (2) goes to zero as $m \to \infty$, we conclude first that the sequence $\big(x_k(t)\big)_{k=1}^{\infty}$ is a Cauchy sequence, and hence has a limit $x(t)$, and then that this convergence is uniform, so that $x(t)$ is a continuous function.

Finally, we let $k \to \infty$ in the relation $x_k(t) = a + \int_0^t f(x_{k-1}(\tau))\,d\tau$ to deduce (1.2). The only slightly nontrivial part of this limit procedure is to show that $\int_0^t f(x_{k-1}(\tau))\,d\tau \to \int_0^t f(x(\tau))\,d\tau$. To see this, estimate the difference:

$$\Big|\int_0^t \big(f(x_k(\tau)) - f(x(\tau))\big)\,d\tau\Big| \le \int_0^t \big|f(x_k(\tau)) - f(x(\tau))\big|\,d\tau \le \int_0^t L\big|x_k(\tau) - x(\tau)\big|\,d\tau \to 0$$

because $x_k \to x$ uniformly. ∎

In fact, it is not hard to show that there exists a *maximal interval of existence*, that is an open interval I on which (1.1) has a solution, and so I contains any other open interval with a solution on it. One simply takes I to be the union of all open intervals J containing 0 so that (1.1) has a solution on J. For any $t \in I$, pick some J on which there exists a solution y, and define $x(t) = y(t)$. If K is another such interval, and z is a solution on K, then $J \cap K$ is yet another interval, so the uniqueness theorem shows that $y = z$ on $J \cap K$. Therefore our definition of $x(t)$ does not depend on the particular choice of J.

3 Theorem. *Let the maximal interval of existence be (α, β), where $-\infty \le \alpha < 0 < \beta \le \infty$. Write $\delta(x) = \min\{\operatorname{dist}(x, \partial\Omega), 1/|x|\}$.*

If $\beta < \infty$, then $\delta\big(x(t)\big) \to 0$ when $t \to \beta$. Similarly, if $\alpha > -\infty$, then $\delta\big(x(t)\big) \to 0$ when $t \to \alpha$.[2]

Proof sketch: For any $\varepsilon > 0$, write $K_\varepsilon = \big\{x \in \Omega\colon \delta(x) \ge \varepsilon\big\}$. Then K_ε is compact.

Assuming $\beta < \infty$ and that the stated conclusion does not hold. That means there exists some $\varepsilon > 0$ and a sequence $t_k \to \beta$ with $x(t_k) \in K_\varepsilon$ for all k.

We may as well assume that $\varepsilon < 2$, in which case $\operatorname{dist}(K_\varepsilon, \mathbb{R}^n \setminus K_{\varepsilon/2}) < \varepsilon/2$. Let $M = \max_{K_{\varepsilon/2}} |f|$. It is clear that $x(t)$ exists and belongs to $M = K_{\varepsilon/2}$ for any $t \in [t_k, t_k + \varepsilon/(2M)]$, since the solution does not move fast enough to leave $M = K_{\varepsilon/2}$ in time $\varepsilon/(2M)$.

As $t_k \to \beta$, this contradicts the definition of β as soon as $t_k + \varepsilon/(2M) > \beta$. This contradiction proves the first part.

The second part is proved similarly, or it follows from the first part by time reversal. ∎

[2] A more common statement of the theorem says that $x(t)$ will leave any compact set $K \subset \Omega$, in the sense that there is some $\beta' < \beta$ so that $x(t) \notin K$ when $\beta' < t < \beta$, and a similar statement at the negative end.

Continuous dependence on data

Let us consider the dependence on the initial value a. Assume that x solves (1.1), and that y solves the same system, but with initial data $y(0) = b$. If f is Lipschitz continuous with Lipschitz constant L, we have seen that $e^{-Lt}\big|x(t) - y(t)\big|$ is non-increasing, so that

$$\big|x(t) - y(t)\big| \leq e^{L|t|}\big|x(0) - y(0)\big|$$

(I added a strategic absolute value in the exponent on the righthand side, so the result can also be used for $t < 0$.) So the solution depends continuously on the initial data. (The dependence is locally Lipschitz continuous, but that takes a bit of effort to prove, so I'll skip it.)

If f is smoother, then we can even conclude that the solution depends on the initial data in a *differentiable* way:

Write now $x(t, a)$ for the solution with initial condition a, so that (1.1) can be written

$$\begin{aligned} \frac{\partial x}{\partial t} &= f\big(x(t,a)\big), \\ x(0,a) &= a \end{aligned}$$

Assuming for a moment that x is differentiable with respect to a, with continuous partial derivatives, we expect to find

$$\frac{\partial}{\partial t}\frac{\partial x}{\partial a_j} = \frac{\partial}{\partial a_j}\frac{\partial x}{\partial t} = \frac{\partial}{\partial a_j} f\big(x(t,a)\big) = Df\big(x(t,a)\big)\frac{\partial x}{\partial a_j}$$

so that $\partial x/\partial a_j$ itself satisfies a differential equation. It will also satify the initial condition $\partial x/\partial a_j(0) = e_j$, where e_j is the jth unit vector.

One can turn this argument inside out: Assuming that Df is Lipschitz continuous, the problem $\dot{z}_j = Df\big(x(t,a)\big)z_j$, $z_j(0) = e_j$ has a solution, and that solution can then be shown to be the partial derivative $\partial x/\partial a_j$. But we shall skip the details.

Odds and ends

Non-autonomous systems. The initial value problem for a non-autonomous system

$$\begin{aligned} \dot{x}(t) &= f\big(x(t), t\big), \\ x(0) &= a. \end{aligned}$$

can be reduced to the autonomous form (1.1) by writing $w(t) = \big(x(t), t\big)$ and solving the autonomous system

$$\begin{aligned} \dot{w}(t) &= \big(f(w(t)), 0\big), \\ w(0) &= (a, 0). \end{aligned}$$

This may not be the best way to study non-autonomous systems, but it does show that the well-posedness results extends to this case.

Continuous dependence on f. Assume that f depends on further parameters $b \in \mathbb{R}^m$:

$$\begin{aligned} \dot{x}(t) &= f\big(x(t), b\big), \\ x(0) &= a. \end{aligned}$$

A rather silly looking way to solve this is to write $w(t) = \big(x(t), b\big)$ and to solve

$$\begin{aligned} \dot{w}(t) &= \big(f\big(w(t)\big), 0\big), \\ w(0) &= (a, b). \end{aligned}$$

That is, we add the components b to x and add equations saying that those components of w are constants (their derivatives are zero).

Note that the b moved from f into the initial conditions. It follows that the solution depends continuously (smoothly, if f is smooth) on b.

The flow

Consider an autonomous system $\dot{x} = f(x)$. From the well-posedness result we know that the solution $x(t)$ at a time t is uniquely determined by t and the initial data $x(0) = a$. Thus we could write $x(t) = \Phi(t, a)$ for some function Φ of two variables, but I shall prefer instead a slightly different notation, and write

$$x(t) = \Phi^t\big(x(0)\big)$$

if x is a solution which is defined in an interval containing 0 and t.

Stated differently, Φ is defined by

$$\begin{aligned}\frac{\partial}{\partial t}\Phi^t(a) &= f\big(\Phi^t(a)\big),\\ \Phi^0(a) &= a.\end{aligned}$$

The map Φ is called the *flow* of the given system. According to the well-posedness result, Φ is continuous (even Lipschitz). It may not be defined everywhere, since a solution to the system may go off to infinity or leave the domain of f in finite time. However, it is not difficult to show that Φ is defined in an open subset D_Φ of $\mathbb{R} \times D_f$ (where D_f is the domain of f), and so that for any $a \in D_f$, the set $\big\{t \in \mathbb{R}\colon (t, a) \in D_\Phi\big\}$ is an interval containing 0.

The most important property of the flow is called the *semigroup* property:

$$\Phi^{t_1+t_2}(a) = \Phi^{t_2}\big(\Phi^{t_1}(a)\big)$$

whenever the righthand side is defined. To show this, just note that $t \mapsto \Phi^{t+t_2}(a)$ and $t \mapsto \Phi^{t_2}\big(\Phi^{t_1}(a)\big)$ are two solutions to the system (autonomy is needed for the first one), with the same initial data $\Phi^{t_1}(a)$ at $t = 0$, so they are the same solution.

We often write the semigroup property in the shorter form

$$\Phi^{t_1+t_2} = \Phi^{t_2} \circ \Phi^{t_1},$$

but it should be noted that $\Phi^{t_1+t_2}$ may be defined at some point where $\Phi^{t_2} \circ \Phi^{t_1}$ is not.

For a simple example, consider the scalar equation $\dot{x} = \lambda x$, where λ is a constant. The general solution is of the form $x = ae^{\lambda t}$ for a constant a, and then clearly $x(0) = a$, so the flow is given by $\Phi^t(a) = ae^{\lambda t}$, and the semigroup property corresponds to the familiar addition formula for the exponential function. (Incidentally, this is why I like to put the t in Φ^t in the exponent position.)

As a further example, consider the scalar equation $\dot{x} = x^2$. Separation of variables yields $x = 1/(C - t)$ for the general solution. Initial data $x(0) = a$ forces $C = 1/a$, so we can write the flow as

$$\Phi^t(a) = a/(1 - at), \qquad at < 1.$$

The semigroup property is less obvious but still easily verified. This time, however, $\Phi^{t_2}\big(\Phi^{t_1}(a)\big)$ can well be undefined while $\Phi^{t_1+t_2}(a)$ is defined: For a trivial example, pick $a = 1$, $t_1 = 1$, $t_2 = -\frac{1}{2}$.

The flow encodes everything about the dynamical system. For example, f itself can be recovered from the flow via the simple rule

$$f(x) = \frac{\partial}{\partial t}\Phi^t(x)\Big|_{t=0}$$

In fact, it is possible to characterize precisely those maps Φ which are flows of a dynamical system: All that is required is that Φ and $\partial\Phi/\partial t$ be (locally) Lipschitz, and that they have the semigroup property (including $\Phi^0(x) = x$).

Accordingly, some authors consider the flow to *be* the dynamical system. This is a point of view that fits better with the notion of a discrete dynamical system, which is typically defined by iterating a map $f\colon \mathbb{R}^n \to \mathbb{R}^n$. Writing f^n for the n-fold composition $f \circ f \circ \cdots \circ f$, we again recognize the semigroup property $f^{n_1+n_2} = f^{n_2} \circ f^{n_1}$, thus making the analogy between the two kinds of dynamical system clear.

Chapter 2

Equilibria for planar systems

This note is about the classification of equilibrium points of nonlinear systems in the plane. We shall deal only with those cases which can be (mostly) decided on the basis of the linearization of the systems around the equilibrium.

But first, a few generalities that apply in any dimension: Consider a system of the form

$$\dot{z} = f(z)$$

where the unknown function z is vector valued: $z(t) \in \mathbb{R}^n$.

If z_0 is an equilibrium point, i.e., if $f(z_0) = 0$, then the change of variables $w = z - z_0$ changes the system into one with an equilibrium at 0. So we lose no generality in assuming that $z_0 = 0$, and shall do so through this note.

So assume now that $f(0) = 0$, and also that f is *differentiable* at 0: Thus

$$f(z) = Az + o(r), \qquad z \to 0,$$

where A is an $n \times n$ matrix, called the *Jacobian matrix* or the *derivative* of f at 0.[1] We write $df(0)$ or even $f'(0)$ for this matrix.

In order to see what information can be gained from the Jacobian matrix, it is useful to reduce attention to a handful of *normal forms*. If we perform a linear change of variables:

$$z = Vw,$$

where V is an invertible matrix, the system is transformed into the form $V\dot{w} = f(Vw)$, or

$$\dot{w} = g(w), \qquad \text{where } g(w) = V^{-1}f(Vw).$$

Since $g(w) = V^{-1}f(Vw) = V^{-1}\big(f'(0)Vw + o(|Vw|)\big) = V^{-1}f'(0)Vw + o(|w|)$, we must have

$$g'(0) = V^{-1}f'(0)V.$$

In other words, $g'(0)$ is *similar* to $f'(0)$.

[1] Actually, I prefer to think of it as a linear map $\mathbb{R}^n \to \mathbb{R}^n$.

Normal forms of equilibria for planar systems

It is a fundamental result of linear algebra that every real 2×2 matrix is similar to one of the following *normal forms*:

Two distinct, real eigenvalues:

$$\begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix}$$

These give rise to the *nodes* ($\lambda_1\lambda_2 > 0$) and *saddle points* ($\lambda_1\lambda_2 < 0$). The degenerate cases ($\lambda_1\lambda_2 = 0$) require more detailed analysis and will be skipped here.

Two complex eigenvalues $\sigma \pm i\omega$, with $\sigma, \omega \in \mathbb{R}$ and $\omega \neq 0$:

$$\begin{pmatrix} \sigma & -\omega \\ \omega & \sigma \end{pmatrix}$$

These give rise to *foci* [singular: *focus*] when $\sigma \neq 0$. The cases where $\sigma = 0$ (the linearization is a center) require more detailed analysis

One real eigenvalue (two possibilities):

$$\begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix} \qquad \begin{pmatrix} \lambda & 0 \\ 0 & \lambda \end{pmatrix}$$

These are borderline cases between the nodes and the foci.

Nodes

We consider a planar system of the form

$$\begin{aligned} \dot{x} &= \lambda_1 x + g(x,y), & g(x,y) &= o(r), \\ \dot{y} &= \lambda_2 y + h(x,y), & h(x,y) &= o(r), \end{aligned}$$

where the "little-oh" notation refers to the limit as $r = \sqrt{x^2+y^2} \to 0$. We must also assume that g and h are sufficiently regular that the existence and uniqueness theorems hold: C^1 is the usual assumption, but Lipschitz continuity is sufficient in the first part of the analysis.

To investigate stability of the system, use polar coordinates:

$$\dot{r} = \frac{x\dot{x} + y\dot{y}}{r} = \frac{\lambda_1 x^2 + \lambda_2 y^2 + xg(x,y) + yh(x,y)}{r} = \frac{\lambda_1 x^2 + \lambda_2 y^2}{r} + o(r).$$

First, assuming $\lambda_1 \leq \lambda_2 < 0$ we conclude

$$\dot{r} < \lambda_2 r + o(r) < (\lambda_2 + \varepsilon) r$$

for any ε and r small enough.[2] Thus (ensure $\lambda_2 + \varepsilon < 0$) $r \to 0$ exponentially as $t \to \infty$, if the starting value is small enough. This is the *stable* case.

But we can say more: We also find

$$\dot\theta = \frac{x\dot y - y\dot x}{r^2} = \frac{\lambda_2 x\big(y + h(x,y)\big) - \lambda_1 y\big(x + g(x,y)\big)}{r^2} = (\lambda_2 - \lambda_1)\frac{xy}{r^2} + o(1),$$

and using $xy = r\cos\theta \cdot r\sin\theta = \frac{1}{2} r \sin 2\theta$,

$$\dot\theta = \tfrac{1}{2}(\lambda_2 - \lambda_1)\sin 2\theta + o(1).$$

If $\lambda_1 < \lambda_2 < 0$ and r is small enough, the final $o(1)$ term will be small enough so that $\dot\theta$ must have the same sign as $\sin 2\theta$, except inside four narrow sectors around the axes, as indicated in figure 2.1.

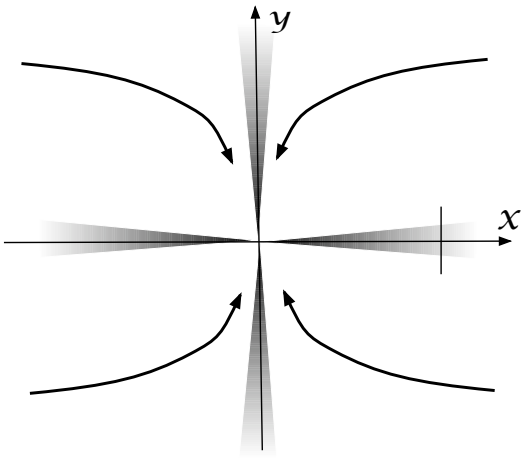

Figure 2.1: An attractive node

Outside the shaded sectors, the movement must be roughly as indicated, in the sense that $\dot r < 0$ and the sign of $\dot\theta$ must be as shown by the arrows.

In particular, any trajectory that does not stay within the horizontal sectors must end up inside the vertical ones. (For example, looking in the first quadrant but outside the shaded sectors, we have a definite lower positive bound on $\dot\theta$, so the trajectory must cross into the sector around the y axis in a finite time.)

By considering ever smaller scales, we can redraw the figure with ever narrower shaded sectors. So we conclude that the majority of trajectories will approach the

[2] The precise statement: Given any $\varepsilon > 0$ there is some $\delta > 0$ so that the inequality holds whenever $r < \delta$.

origin from the vertical direction, but it is conceivable that some trajectories will do so horizontally.

In fact, some of them must do so, as we show next.

Consider initial data (x_0, y_0) with $x_0 > 0$ fixed and small, while y_0 varies (i.e., on the thin vertical line indicated in the righthand side of Figure 2.1). For large enough (but still small) $|y_0|$, the initial point will be outside the shaded sector, and so the solution will remain outside the sector. For others, the solution will escape the sector either upwards or downwards. Now, by the continuous dependence of solutions with respect to initial data, the set of y_0 for which the solution escapes upwards will be an open set A, and the set for which it escapes downwards will be another open set B. In fact, these sets will be intervals, since solution curves cannot cross, so any solution trapped above one that escapes upwards will itself do so, and similarly for the downward escaping ones. So there must be at least one y_0 that belongs to neither A nor B, and the trajectory through this point must approach the origin while staying inside the sector (since there is nowhere else for it to go).

If we assume a bit more regularity of the righthand side of the equation, one can show that only one curve from each side will approach the origin horizontally, but we shall not prove it here.

If both eigenvalues are positive rather than negative, we get the described behaviour of solutions, but this time as $t \to -\infty$. There is no need to repeat the analysis: Just apply the results of the above reasoning to the reversed system.

As an alternative to polar coordinates, we might note that the trajectories satisfy

$$\frac{x}{y}\frac{dy}{dx} = \frac{\lambda_2 + h(x,y)/y}{\lambda_1 + g(x,y)/x} \to \frac{\lambda_2}{\lambda_1} \quad \text{as } r \to 0$$

so long as they stay out of the shaded sectors around the axes. Thus the trajectory satisfies a differential equation of the form

$$\frac{dy}{dx} = \alpha(x)\frac{y}{x}$$

with $\alpha(x) \to \gamma = \lambda_2/\lambda_1$ when $x \to 0$. Note that $0 < \gamma < 1$. That equation implies

$$\ln|y| = \int \alpha(x)\frac{dx}{x} \sim \gamma \ln|x| \quad \text{as } x \to 0$$

so that y behaves like $C|x|^\gamma$ for small x, and (x, y) does not stay out of the sector around the y axis after all.

If g and h are C^1 functions, and if $y_1(x)$ and $y_2(x)$ are two trajectories that both approach the origin along the x axis, then with a bit more effort we can show that $z = y_2 - y_1$ satisfies an equation of the form

$$\frac{dz}{dx} = \beta(x)\frac{z}{x}$$

with $\beta(x) \to \gamma$ as above (subtract the two equations for y_2 and y_1 and manipulate the result for a while), and again this leads to a contradiction. Thus only one trajectory approaches the origin from each side of the x axis.

Coinciding eigenvalues.

$$\begin{aligned} \dot{x} &= \lambda x + \varepsilon y + g(x,y), & g(x,y) &= o(r), \\ \dot{y} &= \lambda y + h(x,y), & h(x,y) &= o(r), \end{aligned}$$

with $\lambda \neq 0$. The case $\varepsilon = 0$ is the case where the eigenspace corresponding to the single eigenvalue is two-dimensional, while the case $\varepsilon \neq 0$ corresponds to a Jordan normal form with a one-dimensional eigenspace. The latter is usually specified with $\varepsilon = 1$, but replacing y by εy yields the above form.

We shall insist on having $0 \leq \varepsilon < |\lambda|$: For then

$$\dot{r} = \frac{x\dot{x} + y\dot{y}}{r} = \frac{\lambda r^2 + \varepsilon xy}{r} + o(r).$$

Noting as before that $xy = \frac{1}{2}r^2 \sin 2\theta$, we find $|\varepsilon xy| \leq \frac{1}{2}r^2 < \frac{1}{2}|\lambda|$, so that $\dot{r}$ has the same sign as λ when r is small enough, and solutions tend to 0 exponentially as $t \to \infty$ (if $\lambda < 0$) or $t \to -\infty$ (if $\lambda > 0$).

In other words, the question of stability is settled just as previously. What happens in the angular (θ) direction is far less clear-cut.

Saddles

Next, we consider the case

$$\begin{aligned} \dot{x} &= -\lambda x + g(x,y), & g(x,y) &= o(r), \\ \dot{y} &= \mu y + h(x,y), & h(x,y) &= o(r), \end{aligned}$$

where $\lambda, \mu > 0$ and the "little-oh" notation again refers to the limit as $r = \sqrt{x^2+y^2} \to 0$.

Polar coordinates are not quite as useful in this case, but we can instead note that $\dot{x}$ and x have opposite signs, so long as r is small and we stay outside the thin sectors around the y axis in Figure 2.2.

In detail: Let $\varepsilon > 0$. If r is small enough then $|g(x,y)| < \varepsilon r$. Whenever $\varepsilon r < \lambda |x|$, then $\dot{x}$ and x must indeed have opposite signs. Squaring the inequality we get $\varepsilon^2(x^2+y^2) < \lambda^2 x^2$, and assuming $\varepsilon < \lambda$, that is true if $\sqrt{\lambda^2 - \varepsilon^2}\,|x| < \varepsilon |y|$.

A similar analysis shows that $\dot{y}$ and y have the same signs outside the small sectors around the x axis.

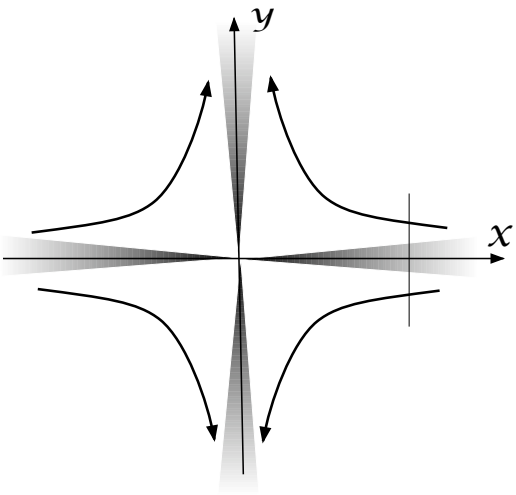

Figure 2.2: A saddle point

So outside the four sectors, all solutions must move in the general direction indicated by the arrows.

In particular, any solution that strays outside the two horizontal sectors will escape out of the neighbourhood in the vertical direction.

Repeating the argument from the node case, considering initial data on the small vertical line segment across the positive x axis, we find that some trajectory will approach the origin horizontally from the right (and similarly, one from the left).

Once more, assuming a bit more regularity of the system we can show that there is only one such trajectory from each side.

Assume the righthand side of the system is C^1. Thus g and h are C^1 functions. Since both are $o(r)$, the first order partial derivatives of each are 0 at $(0,0)$, and it follows that $\partial_y g = \partial g/\partial y = o(1)$ as $r \to 0$. Applying the mean value formula in the y direction, we conclude from this that $g(x,y_2) = g(x,y1) + o(1)(y_2 - y_1)$, and the similar formula for h.

Now assume two different trajectories approach the origin from the right half plane. Write them in the form $y_i(x)$ for $i = 1,2$, let $z = y_2 - y_1$, and subtract the two equations

$$\frac{dy_i}{dx} = \frac{\mu y_i + h(x,y_i)}{-\lambda x + g(x,y_i)}, \qquad i = 1,2,$$

and get (after multiplying by the common denominator)

$$\begin{aligned}(-\lambda x + g(x, y_1))&(-\lambda x + g(x, y_2)) \\ &= (\mu y_2 + h(x, y_2))(-\lambda x + g(x, y_1)) - (\mu y_1 + h(x, y_1))(-\lambda x + g(x, y_2)) \\ &= (-\lambda\mu x + o(r))z\end{aligned}$$

where we have substituted $g(x, y_2) = g(x, y1) + o(1)z$ and $h(x, y_2) = h(x, y1) + o(1)z$ and noticed plenty of cancellation in the last line. As we have seen, the trajectories y_i must approach the origin in a sector, so the $o(r)$ term above can be written $o(x)$, and we conclude that, if $y_1 < y_2$ then $dz/dx < 0$. Thus z increases as $x \to 0$ from above, but this is impossible if both orbits are to stay within a sector.

Taken together with the origin itself, the two curves approaching the origin from each side form what is known as the *stable curve* of the equilibrium point. This curve is tangential with the x axis. All initial data near the equilibrium point and not on the stable curve, must escape out of small neighbourhoods.

Now we can reverse time and repeat the argument. In the original system, we conclude that there is an *unstable curve* which is tangent to the y axis, along which solutions tend to the equilibrium points as $t \to -\infty$. And all other initial data produces solutions that escape the neighbourhood when time runs backwards.

Foci

Finally, we have saved the easiest case for last: systems on the form

$$\begin{aligned}\dot{x} &= \sigma x - \omega y + g(x, y), & g(x, y) &= o(r), \\ \dot{y} &= \omega x + \sigma y + h(x, y), & h(x, y) &= o(r),\end{aligned}$$

where $\omega \neq 0$. In polar coordinates we get

$$\dot{r} = \sigma r + o(r), \quad \dot{\theta} = \omega + o(1), \qquad \text{as } r \to 0.$$

The first equation shows that if $\sigma < 0$ then $r \to 0$ exponentially fast as $t \to \infty$, and the equilibrium is *stable*.

Similarly, if $\sigma > 0$ then $r \to 0$ exponentially fast as $t \to -\infty$, and the equilibrium is *unstable*.

In either case, θ grows approximately at a linear rate as $r \to 0$, so the solution spirals around the equilibrium point an infinite number of times.

This behaviour defines a *focus* in general.

Note that in the degenerate case $\sigma = 0$ the linearized system is a center, but anything might happen to the nonlinear system: It could be a center or a focus, or there could be an infinite sequence of closed trajectories around the origin, typically with spirals in between.

Summary

Via a linear change of coordinates the above analysis applies to equilibrium points of any C^1 planar system, so long as the eigenvalues of the Jacobian matrix are distinct and have nonzero real part.

If both eigenvalues are real and of the same sign, we get a *node*. By definition, a node is either stable or unstable: In the stable case, all nearby trajectories approach the equilibrium as $t \to \infty$. In the ustable case, they approach the equilibrium as $t \to -\infty$. And in either case, they all do so tangentially to a common line through the equilibrium point, with just two exceptions, which approach from opposite sides tangentially to a different line. The two lines are parallel to the eigenspaces of the Jacobian matrix.

The node is stable if the eigenvalues are negative, and unstable if they are positive. This result holds even if the eigenvalues are equal, but in the latter case behaviour may be like either a node or a focus, and more detailed analysis is needed.

If both eigenvalues are real and of opposite signs, we get a *saddle point*. Through the saddle point are two curves, the *stable curve* which is tangent to the eigenspace corresponding to the negative eigenvalue, and the *unstable curve* which is tangent to the eigenspace corresponding to the positive eigenvalue. The stable and unstable curves are each composed of two trajectories and the equilibrium point itself. The two trajectories on the stable curve approach the equilibrium point as $t \to \infty$, and those on the unstable curve approach the equilibrium point as $t \to -\infty$.

Finally, in the case of non-real eigenvalues, the two eigenvalues will be complex conjugates of each other, and we get a *focus*: Solutions will spiral towards the equilibrium either as $t \to \infty$ (the stable case) or $t \to -\infty$ (the unstable case).

Chapter 3

Linearization at equilibrium points

This chapter is about the behaviour of a nonlinear autonomous system $\dot{x} = f(x)$ (where $x(t) \in \mathbb{R}^n$) near an equilibrium point x_0 (i.e., $f(x_0) = 0$).

The *Hartman–Grobman* theorem states that the system behaves "just like" its linearization near the equilibrium point. However, this theorem requires of the linearization that no eigenvalues have real part zero: Thus it is not the appropriate tool for deciding instability where some eigenvalue has a positive real part, since other eigenvalues may have a real part equal to zero in general. The Hartman–Grobman theorem will decide stability when all eigenvalues have negative real parts, but this is sort of a sledgehammer approach where simpler tools will do the job.

We will first develop and use these simpler tools, then return to the Hartman–Grobman theorem and its more appropriate uses later.

Stability and instability of equilibrium points

In this section, we use suitable Liapunov functions to prove the standard results on stability and instability of equilibria based on the eigenvalues of the linearization. We consider estimates for the linear part first.

After a change of variables, a linear system $\dot{x} = Ax$ can be written on the form $\dot{u} = Ju$, where J is a matrix on Jordan normal form: I assume that you know what this means, but remind you of the basic Jordan building block:

$$\lambda I + N = \begin{pmatrix} \lambda & 1 & 0 & 0 & \dots & 0 \\ 0 & \lambda & 1 & 0 & \dots & 0 \\ 0 & 0 & \ddots & \ddots & & \vdots \\ \vdots & & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & \dots & 0 & \lambda & 1 \\ 0 & \dots & \dots & \dots & 0 & \lambda \end{pmatrix} \tag{3.1}$$

where N is the matrix with 1 just above the main diagonal, and zeroes elsewhere. For our purposes, it is useful to note that this is similar to a matrix $\lambda I + \varepsilon N$, where

each 1 above the diagonal is replaced by an ε. To be more precise, $D^{-1}(\lambda I + N)D = \lambda I + \varepsilon N$ where D is the diagonal matrix with $1, \varepsilon, \varepsilon^2, \ldots$ on the diagonal.

The reason this is interesting is the estimate

$$(\lambda - \varepsilon)|u|^2 \le u^{\mathrm{T}}(\lambda I + \varepsilon N)u \le (\lambda + \varepsilon)|u|^2 \tag{3.2}$$

for any vector u. It follows via a simple calculation from $|u^{\mathrm{T}}Nu| = |u_1u_2 + u_2u_3 + \cdots| \le |u|^2$, which in turn comes from the Cauchy–Schwarz inequality.

Now, I lied a bit above, for there are complex eigenvalues to be considered as well. To make a long story short, complex eigenvalues come in mutually conjugate pairs $\lambda = \sigma \pm i\omega$ where $\sigma, \omega \in \mathbb{R}$ and $\omega \neq 0$. These can give rise to Jordan blocks almost like (3.2), except each λ must be replaced by a 2×2matrix

$$\begin{pmatrix} \sigma & -\omega \\ \omega & \sigma \end{pmatrix}$$

and each 1 above the diagonal by a 2×2 identity matrix. But then we can perform the same rescaling trick as before, essentially replacing each of these identity matrices by ε times the identity, and we obtain an estimate just like (3.2), but with $\sigma = \operatorname{Re}\lambda$ replacing λ in the upper and lower bounds.

All this handwaving amounts to a proof of the following:

4 Lemma. *If A is a real, quadratic matrix, and each eigenvalue λ of A satisfies $\alpha \le \operatorname{Re}\lambda \le \beta$, then for any $\varepsilon > 0$, A is similar to a matrix $\tilde{A}$ satisfying*

$$(\alpha - \varepsilon)|u|^2 \le u^{\mathrm{T}}\tilde{A}u \le (\beta + \varepsilon)|u|^2$$

for each vector u.

We are now in a position to show the stability result:

5 Proposition. *Let x_0 be an equilibrium point of $\dot{x} = f(x)$, where f is a C^1 vector field. If $\operatorname{Re}\lambda < 0$ for each eigenvalue of the Jacobian matrix of f at x_0, the equilibrium is asymptotically stable.*

Proof: We may assume $x_0 = 0$ without loss of generality (after all, it's just a change of variables). So the system is of the form

$$\dot{x} = f(x) = Ax + o(|x|).$$

Now the Jacobian A of f at 0 has only a finite number of eigenvalues, all of which have negative real part – so there is some $\varepsilon > 0$ with $\operatorname{Re}\lambda \le -2\varepsilon$ for each eigenvalue

λ. By Lemma 4, we can perform a further linear change of variables so that the system takes the form

$$\dot{u} = \tilde{A}u + o(|u|),$$

and where $u^{\mathrm{T}}\tilde{A}u \le -\varepsilon|u|^2$ for all u.

Consider the function $V(u) = \frac{1}{2}|u|^2$. Then

$$\dot{V} = u^{\mathrm{T}}\dot{u} = u^{\mathrm{T}}\tilde{A}u + o(|u|^2) \le -\varepsilon|u|^2 + o(|u|^2) < 0$$

when $|u|$ is small enough, so V is a strong Liapunov function, and 0 is indeed asymptotically stable. ∎

We take a moment for a couple digressions, just because they are easy consequences of Lemma 4.

First, however, we state and prove a lemma. Note that we could have used this at the end of the proof above, rather than referring to Liapunov theory.

6 Lemma. ("Grönwall light") *Assume u is a differentiable function, v is continuous, and $\dot{u} \le uv$ for $t > 0$. Then*

$$u(t) \le u(0)\exp\int_0^t v(s)\,ds \qquad \text{for } t \ge 0.$$

Proof: Define $V(t) = \int_0^t v(s)\,ds$. Then $d(e^{-V}u) = e^{-V}(\dot{u} - uv) \le 0$, so $e^{-V}u$ is non-decreasing, and in particular $e^{-V(t)}u(t) \le e^{-V(0)}u(0) = u(0)$. ∎

The first digression is a proof of the following theorem, from chapter 8.11 of the book.

7 Theorem. *Consider the system*

$$\dot{x} = Ax + h(x,t)$$

in which A is an $n \times n$ matrix all of whose eigenvalues have negative real part, and where

$$h(x,t) = o(|x|) \quad \text{uniformly in } t \text{ for } t \ge 0.$$

Then the zero solution is asymptotically stable.

Proof: According to Lemma 4, by a suitable linear change of variables (and reusing the old names for the new variables) we can arrange things so that

$$x^{\mathrm{T}}Ax \le -2\varepsilon|x|^2$$

for all $x \in \mathbb{R}^n$, where $\varepsilon > 0$. Furthermore, the assumption on h means that there is some $\delta > 0$ so that

$$|h(x,t)| < \varepsilon|x| \qquad \text{whenever } |x| < \delta \text{ and } t \ge 0.$$

Whenever x is a solution to the system and $|x(t)| < \delta$, then, by the above two inequalities and the Cauchy–Schwarz inequality,

$$\frac{d}{dt}|x|^2 = 2x^{\mathrm{T}}\dot{x} = 2x^{\mathrm{T}}Ax + 2x^{\mathrm{T}}h(x,t) < -4\varepsilon|x|^2 + 2\varepsilon|x|^2 = -2\varepsilon|x|^2$$

Thus $|x|^2$ decays exponentially to zero as $t \to \infty$. ∎

8 Proposition. *Let A be a constant $n \times n$ matrix whose eigenvalues satisfy* $\operatorname{Re}\lambda < 0$, *and let B be a matrix valued function on $[0,\infty)$ so that $\int_0^\infty \|B(t)\|\,dt < \infty$. Then the zero solution of $\dot{x} = (A + B(t))x$ is asymptotically stable.*

Proof: Neither the hypotheses nor the conclusion are altered by a linear change of variables. Thus, we may assume without loss of generality that $x^{\mathrm{T}}Ax \le -\eta|x|^2$ where $\eta > 0$. Define $u(t) = |x(t)|^2$ and note that

$$\dot{u} = 2x^{\mathrm{T}}\dot{x} = 2(x^{\mathrm{T}}Ax + x^{\mathrm{T}}B(t)x) \le 2(\|B(t)\| - \eta)u.$$

By "Grönwall light", we have

$$u(t) \le u(0)\exp 2\int_0^t (\|B(s)\| - \eta)\,ds \le u(0)e^{2(M-\eta t)}, \qquad M = \int_0^\infty \|B\|\,dt$$

In particular $|x(t)| \le |x(0)|e^M$, which proves stability. Furthermore we get $|x(t)| \to 0$ as $t \to \infty$, so the stability is asymptotic. ∎

9 Lemma. *Consider an equilibrium point* 0 *for a dynamical system $\dot{u} = g(u)$.*

Let U be a C^1 function so that $U(0) = 0$, every neighbourhood of 0 *contains some u with $U(u) > 0$, and assume there is some neighbourhood of* 0 *so that whenever u belongs to this neighbourhood and $U(u) > 0$, then $\dot{U}(u) > 0$ as well. Then* 0 *is an unstable equilibrium point.*

Proof: Let $\varepsilon > 0$ be so that whenever $|u| \le \varepsilon$ and $U(u) > 0$, then $\dot{U}(u) > 0$.

Consider any $\delta > 0$. We shall prove that there exists an orbit starting within the δ-neighbourhood of 0 which must escape the ε-neighbourhood of 0.

So pick any u_0 with $|u_0| < \delta$ and $U(u_0) > 0$. Write

$$K = \{u\colon |u| \le \varepsilon \text{ and } U(u) \ge U(u_0)\}.$$

K is closed and bounded, therefore compact. Since $\dot{U} > 0$ on K, $\dot{U}$ has a positive lower bound on K, say $\dot{U}(u) \ge \gamma > 0$ whenever $u \in K$.

Now let u be the solution with initial value u_0. So long as $u(t) \in K$ then $U(u)$ will grow with a rate at least γ, so if $u(t) \in K$ for all t then $U\big(u(t)\big)$ will grow without bound, which is impossible because U is bounded on the compact set K. Therefore u must leave K, and it can only do that by getting $|u| > \varepsilon$, i.e., by escaping the ε-neighbourhood as claimed. ∎

10 Proposition. *Let x_0 be an equilibrium point of $\dot{x} = f(x)$, where f is a C^1 vector field. If $\operatorname{Re}\lambda > 0$ for some eigenvalue of the Jacobian matrix of f at x_0, the equilibrium is unstable.*

Proof: As before, assume $x_0 = 0$ without loss of generality. So the system is of the form

$$\dot{x} = f(x) = Ax + o(|x|).$$

We may as well assume we have already changed the variables so that A has Jordan normal form. We can also assume that the Jordan blocks of A appear in decreasing order of $\operatorname{Re}\lambda$. Lump together all the blocks with the largest value of $\operatorname{Re}\lambda$, and write u for the corresponding components of x. Write v for the remaining components. The system now has the form

$$\begin{aligned}\dot{u} &= Bu + o(\sqrt{|u|^2 + |v^2|}),\\ \dot{v} &= Cv + o(\sqrt{|u|^2 + |v^2|}),\end{aligned}$$

where each eigenvalue of B satisfies $\operatorname{Re}\lambda = \beta > 0$, while each eigenvalue of C satisfies $\operatorname{Re}\lambda \le \alpha < \beta$. We can certainly insist that $\alpha > 0$ as well. Let $0 < \varepsilon < \frac{1}{2}(\beta - \alpha)$.

We shall change variables yet again, separately for u and v this time, but we will reuse the old variable names for u, v, B and C. We shall use Lemma 4 so that, after the variable change, we find

$$v^{\mathrm{T}} C v < (\alpha + \varepsilon)|v|^2, \quad (\beta - \varepsilon)|u|^2 < u^{\mathrm{T}} B u.$$

Let $U(u, v) = \frac{1}{2}(|u|^2 - |v|^2)$. We claim that U satisfies the conditions of Lemma 9, which will finish the proof.

The only property of U that is nontrivial to prove is the one on the sign of $\dot{U}$. Now we find

$$\begin{aligned}\dot{U} &= u^{\mathrm{T}}\dot{u} - v^{\mathrm{T}}\dot{v}\\ &= u^{\mathrm{T}}Bu - v^{\mathrm{T}}Cv + o(|u|^2 + |v^2|)\\ &> (\beta - \varepsilon)|u|^2 - (\alpha + \varepsilon)|v|^2 + o(|u|^2 + |v^2|),\end{aligned}$$

and when $U > 0$ we have $|u| > |v|$, so we find

$$\dot{U} > (\beta - \alpha - 2\varepsilon)|u|^2 + o(|u|^2) > 0$$

when $|u|$ is small enough. ∎

The Hartman–Grobman theorem

Consider the autonomous system $\dot{x} = f(x)$ with an equilibrium point x_0. We shall assume that f is a C^1 function. The linearization of this system is $\dot{u} = Au$, where A is the Jacobian matrix of f at x_0. The general solution of the linearized system is $u = e^{tA}u_0$, while that for the nonlinear system is written $x = \Phi^t(x_0)$, where Φ is the *flow* of the system.

The proof of the following theorem is beyond the scope of this text. A relatively easy proof can be found in [4]. However, the proof is done in a Banach space setting, which might make it less accessible. The theorem was originally proved independently by Grobman [1] and Hartman [2].

11 Theorem. (Hartman–Grobman) *Under the above assumptions, and with the extra condition that every eigenvalue of A has nonzero real part, there is a homeomorphism H from a neighbourhood S of 0 to a neighbourhood R of x_0, which maps the flow of the linearized system to the flow of the original system, in the sense that $H(e^{tA}u_0) = \Phi^t\big(H(u_0)\big)$ whenever $e^{sA}u_0 \in R$ for all s between 0 and t inclusive.*

A *homeomorphism* is just a continuous map with a continuous inverse.

One weakness of the Hartman–Grobman theorem is the assumption on the eigenvalues, which cannot be avoided: When some eigenvalues have real part zero, the detailed behaviour of the system near the equilibrium cannot be derived from the linearization.

Another weakness of the theorem is that the conclusion is too weak for many applications: A homeomorphism can map a node to a focus!

For example, consider the function

$$H(u,v) = (u\cos s - v\sin s, u\sin s + v\cos s), \qquad s = -\tfrac{1}{2}\ln(u^2+v^2).$$

Assuming that (u,v) satisfy the equations $\dot{u} = -u$ and $\dot{v} = -v$ (corresponding to a stable node), we find $\dot{s} = 1$, and $(x,y) = H(u,v)$ solves the system

$$\dot{x} = -x - y, \quad \dot{y} = x - y,$$

which corresponds to a stable focus.

Of course, the function H above is not differentiable at 0. (It is continuous there, if we add $H(0,0) = (0,0)$ to the definition.) If we require that H and its inverse be differentiable, such behaviour as seen in the example above becomes impossible.

Unfortunately, we cannot guarantee differentiability of H in general. But the following result [3] helps:

12 Theorem. (Hartman) *Under the assumptions of the Hartman–Grobman theorem, if additionally f is a C^2 function and the real parts of all the eigenvalues of A have the same sign, then the homeomorphism H can in fact be chosen to be a C^1 diffeomorphism.*

And by that we mean that H is C^1, and so is its inverse.

As a consequence of this, so long as the vector field f is C^2, any equilibrium whose linearization is a node or focus is itself of the same type.

No matter how differentiable f may be, we cannot conclude any higher degree of differentiability for H. And when eigenvalues exist in both the left and right halfplanes, a homeomorphism is all we can hope for. All this makes the Hartman–Grobman theorem quite a bit less useful than it looks at first sight.

Chapter 4

Linear systems of ODEs with variable coefficients

Let A be a matrix valued function defined on some interval, with each $A(t)$ being an $n \times n$ matrix. A is supposed to be a Lipschitz continuous function of its argument.

This note is about the linear system

$$\dot{x} = Ax + b(t) \tag{4.1}$$

where $x(t)$ is a (column) n-vector for each t, and b is a vector valued function of t, assumed throughout to be continuous.

Consider the following ODE for a matrix valued function Φ, where each $\Phi(t)$ is also supposed to be an $n \times n$ matrix:

$$\dot{\Phi} = A\Phi \tag{4.2}$$

13 Proposition. *Let Φ be a matrix valued function satisfying (4.2). If $\Phi(t_0)$ is invertible for some t_0 then $\Phi(t)$ is in fact invertible for every t, and the inverse $\Psi(t) = \Phi(t)^{-1}$ satisfies the differential equation*

$$\dot{\Psi} = -\Psi A. \tag{4.3}$$

Proof: The differential equation for Ψ is easy to derive: Just differentiate the relation $\Psi\Phi = I$ to get

$$0 = \frac{d}{dt}(\Psi\Phi) = \dot{\Psi}\Phi + \Psi\dot{\Phi} = \dot{\Psi}\Phi + \Psi A\Phi,$$

which when multiplied on the right by Ψ (and using $\Phi\Psi = I$) yields (4.3).

The above proof requires of course not only that Φ is invertible for all t, but also that the inverse is differentiable.

We can make the argument more rigourous by turning inside out, *defining* Ψ to be the solution of (4.3) satisfying the initial condition $\Psi(t_0) = \Phi(t_0)^{-1}$. Then we differentiate:

$$\frac{d}{dt}(\Psi\Phi) = \dot{\Psi}\Phi + \Psi\dot{\Phi} = -\Psi A\Phi + \Psi A\Phi = 0,$$

so that $\Psi\Phi = I$ for all t, since it is so at $t = t_0$. ∎

14 Definition. A matrix valued solution of (4.2), which is invertible for all t, is called a *fundamental matrix* for (4.1).

Clearly, there are many fundamental matrices, for if Φ is one such and B is any constant invertible matrix, then ΦB is also a fundamental matrix.

However, a fundamental matrix is uniquely determined by its value at any given t_0, and if Φ_1 and Φ_2 are two fundamental matrices, we can set $B = \Phi_1^{-1}(t_0)\Phi_2(t_0)$, so that $\Phi_1 B = \Phi_2$ – at $t = t_0$, and hence for all t.

We now show how the fundamental matrix solves the general initial-value problem for (4.1).

In fact, let x be any solution of (4.1). Let Φ be a fundamental matrix, and write $x = \Phi y$. Then $\dot{x} = \dot{\Phi} y + \Phi \dot{y} = A\Phi y + \Phi \dot{y}$, so that (4.1) becomes

$$A\Phi y + \Phi \dot{y} = A\Phi y + b.$$

Two terms cancel of course, and after multiplying both sides by Φ^{-1} on the left what remains is

$$\dot{y} = \Phi^{-1} b,$$

which is trivial to solve. Given the initial condition $x(t_0) = x_0$, that translates into $y(t_0) = \Phi(t_0)^{-1} x_0$, so the solution for y is

$$y(t) = \Phi(t_0)^{-1} x_0 + \int_{t_0}^{t} \Phi(s)^{-1} b(s)\, ds.$$

Multiplying by $\Phi(t)$ on the left we finally have the solution

$$x(t) = \Phi(t)\Phi(t_0)^{-1} x_0 + \Phi(t) \int_{t_0}^{t} \Phi(s)^{-1} b(s)\, ds.$$

Chapter 5

Sectors and Bendixson's index formula

The picture below shows a phase portrait in the vicinity of an equilibrium point.

In this picture, the neighbourhood contains two *elliptic sectors*, recognizable by orbits starting and ending at the equilibrium itself (and these orbits forming ever smaller loops converging on the equilibrium point).

There are also two *hyperbolic sectors*, filled with orbits roughly resembling hyperbolas. The sectors are separated in one case, by *separatrices*, being orbits with one end at the equilibrium separating the sectors.

Finally, there are two *parabolic sectors*, filled with orbits having just one end at the equilibrium, and surrounded by other orbits of the same kind.

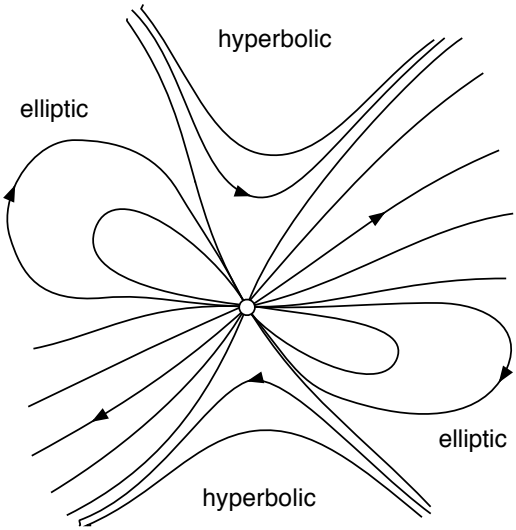

Bendixson's index formula states that the index of the equilibrium point is

$$1 + \frac{e-h}{2}$$

where e is the number of elliptic sectors and h is the number of hyperbolic sectors.

To see why this is true, recall the definition of the index: At any non-equilibrium point, let ϕ be the angle made by the underlying vector field with the x axis. As we traverse a small closed curve (small enough not to surround any other equilibrium) going once around the equilibrium point in the positive direction, keep track of ϕ as a continuously varying quantity. The index is the increase in ϕ, divided by 2π. Written symbolically, it is $\Delta\phi/(2\pi)$.

Now let (r,θ) be polar coordinates centered at the equilibrium point (move the equilibrium to the origin if you prefer). Write $\phi = \theta + \psi$, where ψ must then be the angle between the line from the equilibrium to the current point and the vector field at that point (see the figure).

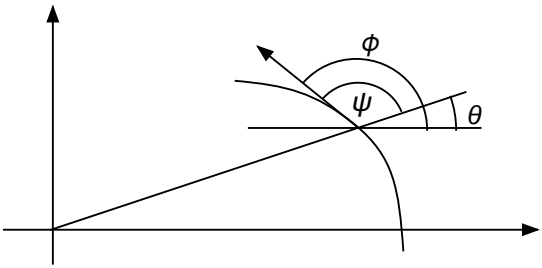

Since θ by definition increases by 2π as we go around a curve of the prescribed sort, symmbolically $\Delta\theta = 2\pi$, it follows that the index must be $1 + \Delta\psi/(2\pi)$.

Now it only remains to notice that ψ increases by π across an elliptic sector, ψ decreases by π across a hyperbolic sector, and it does not change at all across a parabolic sector.

Add these changes to ψ over all sectors, with the result $\Delta\psi = (e-h)\pi$, and Bendixson's formula is the result.

The above argument is exact if the separatrices are straight lines. In practice they need not be, but if we move in sufficiently close to the equilibrium, they can be approximated by straight lines. Since the total increase in ψ must be an integer times 2π, any small errors in estimating the angles will not matter.

Notice finally that e and h must both be odd numbers, or else they must both be even. To see this more directly, note that any separatrix or parabolic sector has a definite direction: Either towards the equilibrium point or away from it. And opposite sides of an elliptic or hyperbolic sectors have opposite directions. Thus there must be in total an even number of these sectors.

Chapter 6

The Poincaré–Bendixson theorem

The Poincaré–Bendixson theorem is often misstated in the literature. The purpose of this note is to try to set the record straight, and to provide the outline of a proof.

Throughout this note we are considering an autonomous dynamical system on the form

$$\dot{x} = f(x), \qquad x(t) \in \Omega \subseteq \mathbb{R}^2$$

where $f\colon \Omega \to \mathbb{R}^2$ is a locally Lipschitz continuous vector field on the open set Ω.

Furthermore, we are considering a solution x whose *forward half orbit* $O_+ = \{x(t)\colon\ t \geq 0\}$ is contained in a compact set $K \subset \Omega$.

An *omega point* of O_+ is a point z so that one can find $t_n \to +\infty$ with $x(t_n) \to z$. It is a consequence of the compactness of K that omega points exist. Write ω for the set of all omega points of O_+.

It should be clear that ω is a closed subset of K, and therefore compact. Also, as a consequence of the continuous dependence of initial data and the general nature of solutions of autonomous systems, ω is an invariant set (both forward and backward) of the dynamical system.

We can now state our version of the main theorem.

15 Theorem. (Poincaré–Bendixson) *Under the above assumptions, if ω does not contain any equilibrium points, then ω is a cycle. Furthermore, either the given solution x traverses the cycle ω, or it approaches ω as $t \to +\infty$.*

What happens if ω does contain an equilibrium point?

The simplest case is the case $\omega = \{x_0\}$ for an equilibrium point x_0. Then it is not hard to show that $x(t) \to x_0$ as $t \to +\infty$. (If not, there is some $\varepsilon > 0$ so that $|x(t) - x_0| \geq \varepsilon$ for arbitrarily large t, but then compactness guarantees the existence of another omega point in $\{z \in K\colon\ |z - x_0| \geq \varepsilon\}$.)

I said in the introduction that the Poincaré–Bendixson theorem is often misstated. The problem is that the above two possibilities are claimed to be the only possibilities. But a third possibility exists: ω can consist of one or more equilibrium points joined by solution paths starting and ending at these equilibrium points (i.e., heteroclinic or homoclinic orbits).

16 Example. Consider the dynamical system

$$\left.\begin{aligned}\dot{x} &= \frac{\partial H}{\partial y} + \mu H \frac{\partial H}{\partial x} \\ \dot{y} &= -\frac{\partial H}{\partial x} + \mu H \frac{\partial H}{\partial y}\end{aligned}\right\}, \qquad H(x,y) = \tfrac{1}{2}y^2 - \tfrac{1}{2}x^2 + \tfrac{1}{4}x^4.$$

Notice that if we set the parameter μ to zero, this is a Hamiltonian system. Of particular interest is the set given by $H = 0$, which consists of the equilibrium point at zero and two homoclinic paths starting and ending at this equilibrium, roughly forming an ∞ sign.

In general, an easy calculation gives

$$\dot{H} = \frac{\partial H}{\partial x}\dot{x} + \frac{\partial H}{\partial y}\dot{y} = \mu H \cdot \left[\left(\frac{\partial H}{\partial x}\right)^2 + \left(\frac{\partial H}{\partial y}\right)^2\right]$$

so that H will tend towards 0 if $\mu < 0$. In particular, any orbit starting outside the "∞ sign" will approach it from the outside, and the "∞ sign" itself will be the omega set of this orbit.

Figure 6.1 shows a phase portrait for $\mu = -0.02$.

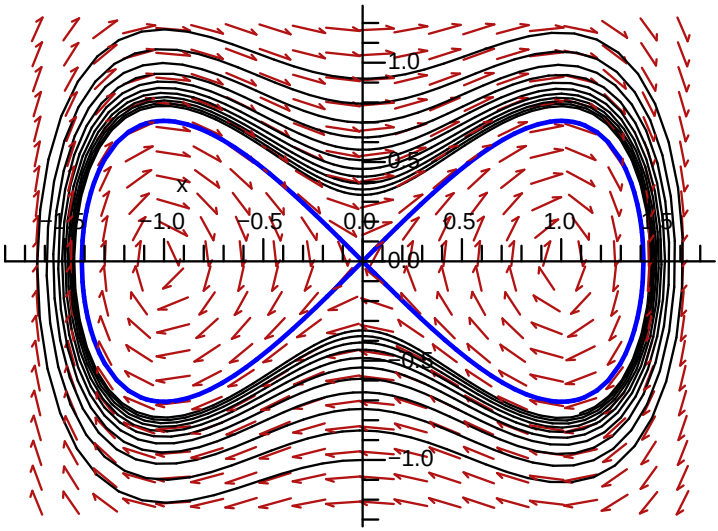

Figure 6.1: An orbit and its omega set.

We now turn to the proof of theorem 15.

By a *transverse line segment* we mean a closed line segment contained in Ω, so that f is not parallel to the line segment at any point of the segment. Thus the vector field points consistently to one side of the segment.

Clearly, any non-equilibrium point of Ω is in the interior of some transverse line segment.

17 Lemma. *If an orbit crosses a transverse line segment L in at least two different points, the orbit is not closed. Furthermore, if it crosses L several times, the crossing points are ordered along L in the same way as on the orbit itself.*

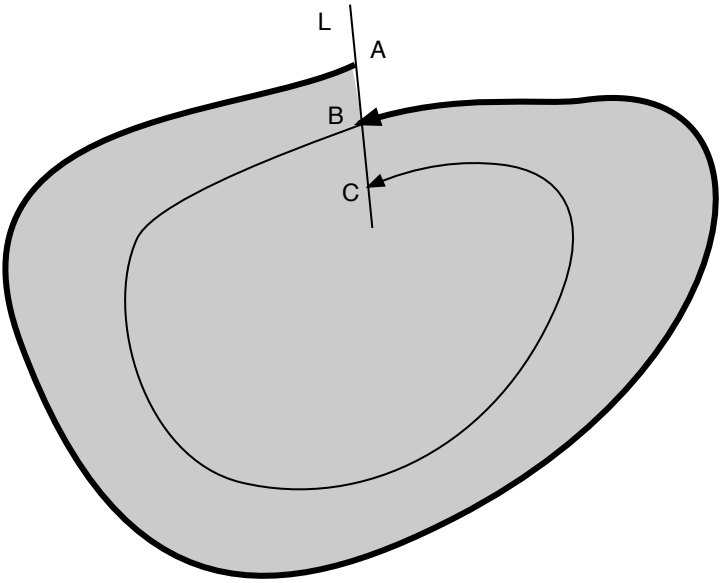

Figure 6.2: Crossings of a transverse line segment

Proof: Figure 6.2 shows a transverse line segment L and an orbit that crosses L, first at A, then at B. Note that the boundary of the shaded area consists of part of the orbit, which is of course not crossed by any other orbit, and a piece of the L, at which the flow enters the shaded region. (If B were to the other side of A, we would need to consider the outside, not the inside, of the curve.) In particular, there is no way the given orbit can ever return to A. Thus the orbit is not closed.

It cannot return to any other point on L between A and B either, so if it ever crosses L again, it will have to be further along in the same direction on L, as in

the point C indicated in the figure. (Hopefully, this clarifies the somewhat vague statement at the end of the lemma.) ∎

18 Corollary. *A point on some orbit is an omega point of that orbit if, and only if, the orbit is closed.*

Proof: The "if" part is obvious. For the "only if" part, assume that A is a point that is also an omega point of the orbit through A. If A is an equilibrium point, we have a special case of a closed orbit, and nothing more to prove. Otherwise, draw a transverse line L through A. Since A is also an omega point, some future point on the orbit through A will pass sufficiently close to A that it must cross L at some point B. If the orbit is not closed then $A \neq B$, but then any future point on the orbit is barred from entering a neighbourhood of A (consult Figure 6.2 again), which therefore cannot be an omega point after all. This contradiction concludes the proof. ∎

Outline of the proof of Theorem 15 Fix some $x_0 \in \omega$, and a transverse line segment L with x_0 in its interior.

If x_0 happens to lie on O_+ the corollary above shows that the orbit through x_0 must be closed, so ω in fact equals that orbit.

If x_0 does not lie on O_+ then O_+ is not closed. However, I claim that the orbit through x_0 is still closed. In fact, let z_0 be an omega point of the orbit through x_0, and draw a transverse line L through z_0. If the orbit through x_0 is not closed, it must pass close enough to z_0 that it must cross L, infinitely often in a sequence that approaches z_0 from one side. In particular, it crosses at least twice, say, first at A and then again at B (again, refer to Figure 6.2).

But B is an omega point of O_+, so O_+ crosses L arbitrarily close to B, and so O_+ enters the shaded area in the figure. But then it can never again get close to A. This is a contradiction, since A is also an omega point of O_+.

We have shown that x_0 lies on a closed path. This closed path must be all of ω. The solution x gets closer and closer to ω, since it crosses a transverse line segment through x_0 in a sequence of points approaching x_0, and the theorem on continuous dependence on initial data does the rest. ∎

Bibliography

[1] D.M. Grobman, Homeomorphisms of systems of differential equations (in Russian). *Dokl. Akad. Nauk SSSR* **128** (1959) 880–881.

[2] Philip Hartman, A lemma in the theory of structural stability of differential equations. *Proc. A.M.S.***11** (1960) 610–620.

[3] Philip Hartman, On local homeomorphisms of Euclidean spaces. *Bol. Soc. Math. Mexicana* **5** (1960) 220–241.

[4] Charles C. Pugh, On a theorem of P. Hartman. *Amer. J. Math.* **91** (1969) 363–367.