Introduction to n-dimensional dynamical systems

Consider the following n-dimensional system of nonlinear ODEs

$$\begin{cases} \dfrac{d\boldsymbol{x}}{dt} = \boldsymbol{f}(\boldsymbol{x}) \\ \boldsymbol{x}(0) = \boldsymbol{x}_0 \end{cases} \tag{1}$$

where $\boldsymbol{x}(t) = [x_1(t) \cdots x_n(t)]^T$ is a vector of phase variables, $\boldsymbol{f} : D \to \mathbb{R}^n$, and D is a subset of $\mathbb{R}^n$. In an expanded notation the system (1) can be written as

$$\begin{cases} \dfrac{dx_1}{dt} = f_1(x_1, \dots, x_n) \\ \dfrac{dx_2}{dt} = f_2(x_1, \dots, x_n) \\ \quad \vdots \\ \dfrac{dx_n}{dt} = f_n(x_1, \dots, x_n) \\ x_1(0) = x_{10} \\ x_2(0) = x_{20} \\ \quad \vdots \\ x_n(0) = x_{n0} \end{cases} \tag{2}$$

Dynamical systems of the form (1) can model phenomena in classical mechanics (e.g., the pendulum equations), health and medicine (e.g., cancer models), weather patterns, material science, and quantum physics. They can also be used to approximate the dynamics of partial differential equations (PDEs). Let us provide two simple examples.

Example: Consider the following sketch of a pendulum (point mass), subject to gravity and friction with friction

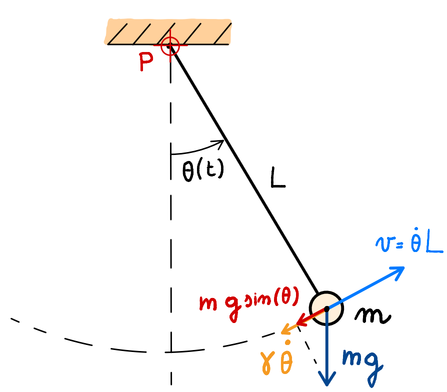

As is well-known from physics, the rate of change (time derivative) of the angular momentum of the point mass m with respect to the point P equals the momentum of the external forces acting on the point mass. The external forces in this case are gravity and friction. Setting up the balance of momenta yields

$$mL^2 \frac{d^2\theta}{dt^2} = -mgL\sin(\theta) - \gamma L \frac{d\theta}{dt}, \tag{3}$$

i.e.

$$\frac{d^2\theta}{dt^2} = -\frac{g}{L}\sin(\theta) - \frac{\gamma}{Lm}\frac{d\theta}{dt}. \tag{4}$$

This is a second-order ordinary nonlinear differential equation in $\theta(t)$. The equation can be easily transformed into a system two first-order nonlinear ODEs by defining the new variables

$$x_1(t) = \theta(t), \qquad x_2(t) = \frac{d\theta(t)}{dt}. \tag{5}$$

Clearly, based on the definition of $x_1(t)$ and $x_2(t)$ we have

$$\frac{dx_1}{dt} = x_2. \tag{6}$$

Moreover, by differentiating $x_2(t)$ with respect to time and using equation (4) we obtain

$$\frac{dx_2}{dt} = -\frac{g}{L}\sin(x_1) - \frac{\gamma}{Lm}x_2. \tag{7}$$

Hence, equation (4) is equivalent to the two-dimensional nonlinear dynamical system

$$\begin{cases} \dfrac{dx_1}{dt} = x_2, \\ \dfrac{dx_2}{dt} = -\dfrac{g}{L}\sin(x_1) - \dfrac{\gamma}{Lm}x_2. \end{cases} \tag{8}$$

Of course, in order to solve the system of ODEs (8), we need an initial condition for the position of the pendulum $x_1(0)$, and an initial condition for the velocity of the pendulum $x_2(0)$.

Example: Consider the following initial-boundary value problem for the heat equation in the periodic spatial domain $[0, 2\pi]$

$$\begin{cases} \dfrac{\partial u(t,y)}{\partial t} = \alpha \dfrac{\partial^2 u(y,t)}{\partial y^2} & \text{heat equation} \\ u(0,y) = u_0(y) & \text{initial condition} \\ u(t,0) = u(t,2\pi) & \text{periodic boundary conditions} \end{cases} \tag{9}$$

A finite-difference approximation of the PDE (9) on the evenly-spaced grid with n points

$$y_k = (k-1)\Delta y \qquad k = 1, \dots, n, \qquad \Delta y = \frac{2\pi}{n} \tag{10}$$

yields the n-dimensional linear dynamical system

$$\begin{cases} \dfrac{dx_1(t)}{dt} &= & \dfrac{\alpha}{\Delta y^2}\left[x_2(t) - 2x_1(t) + x_n(t)\right] \\ \dfrac{dx_2(t)}{dt} &= & \dfrac{\alpha}{\Delta y^2}\left[x_3(t) - 2x_2(t) + x_1(t)\right] \\ & \vdots & \\ \dfrac{dx_n(t)}{dt} &= & \dfrac{\alpha}{\Delta y^2}\left[x_1(t) - 2x_n(t) + x_{n-1}(t)\right] \end{cases} \tag{11}$$

where we defined $x_k(t) = u(y_k, t)$. Note that $x_k(t)$ represents an approximation of the solution to the partial differential equation (9) at the grid point $y = y_k$. Hence, by computing the solution to (11), we are computing an approximation of the solution to the PDE (9) at the grid points $(y_1, \dots, y_n)$.

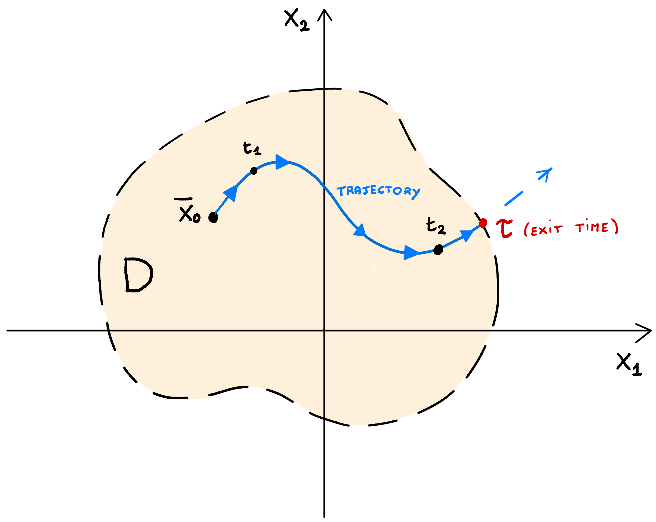

Figure 1: Illustration of the meaning of Theorem 1 in two-dimensions. Shown are the open set $D \subset \mathbb{R}^2$ in which $\boldsymbol{f}(\boldsymbol{x})$ is Lipschitz continuous, the trajectory corresponding to a particular $\boldsymbol{x}_0 \in D$ and the exit time τ for such trajectory.

Well-posedness of the initial value problem. Let us recall theorem that guarantees existence and uniqueness of the solution to the system of first order ODEs (1).

Theorem 1 (Existence and uniqueness of the solution to (1))**.** Let $D \subset \mathbb{R}^n$ be an open set, $\boldsymbol{x}_0 \in D$. If $\boldsymbol{f} : D \to \mathbb{R}^n$ is Lipschitz continuous in D then there exists a unique solution to the initial value problem (1) within the time interval $[0, \tau[$, where τ is the time instant at which the trajectory $\boldsymbol{x}(t)$ exits[1] the domain D. The solution $\boldsymbol{x}(t)$ is continuously differentiable in $[0, \tau[$.

How do we define Lipschitz continuity for a vector-valued function $\boldsymbol{f}(\boldsymbol{y})$ defined on subset of $\mathbb{R}^n$? By a simple generalization of the definition we gave for one-dimensional functions.

Definition 1. Let D be a subset of $\mathbb{R}^n$, $\boldsymbol{f} : D \to \mathbb{R}^n$. We say that $\boldsymbol{f}$ is Lipschitz continuous in D if there exists a constant $0 \leq L < \infty$ such that

$$\|\boldsymbol{f}(\boldsymbol{x}_1) - \boldsymbol{f}(\boldsymbol{x}_2)\| \leq L \|\boldsymbol{x}_1 - \boldsymbol{x}_2\| \quad \text{for all} \quad \boldsymbol{x}_1, \boldsymbol{x}_2 \in D, \tag{12}$$

where $\|\cdot\|$ is any norm defined in $\mathbb{R}^n$ (see Appendix B). Recall, in fact that all norms defined in a finite-dimensional space such as $\mathbb{R}^n$ are *equivalent*.

Similarly to what we have seen for one-dimensional dynamical systems, there conditions that simpler to verify than Lipschitz continuity.

Lemma 1. If $\boldsymbol{f}(\boldsymbol{x})$ is of class C^1 in a compact convex domain $D \subset \mathbb{R}^n$, then $\boldsymbol{f}(\boldsymbol{x})$ is Lipschitz continuous in D.

The proof of this lemma is provided in Appendix B.

Lemma 2. Let $\boldsymbol{f}(\boldsymbol{y}, t)$ be of class C^1 (continuously differentiable) in $D \subseteq \mathbb{R}^n$. If $\boldsymbol{f}(\boldsymbol{y}, t)$ has bounded derivatives $\partial f_i / \partial y_j$ then $\boldsymbol{f}(\boldsymbol{y}, t)$ is Lipschitz continuous in D.

[1] As shown in Figure 1, the "exit time" τ depends on D, $\boldsymbol{f}(x)$ and $\boldsymbol{x}_0$.

(two-dimensional system) (three-dimensional system)

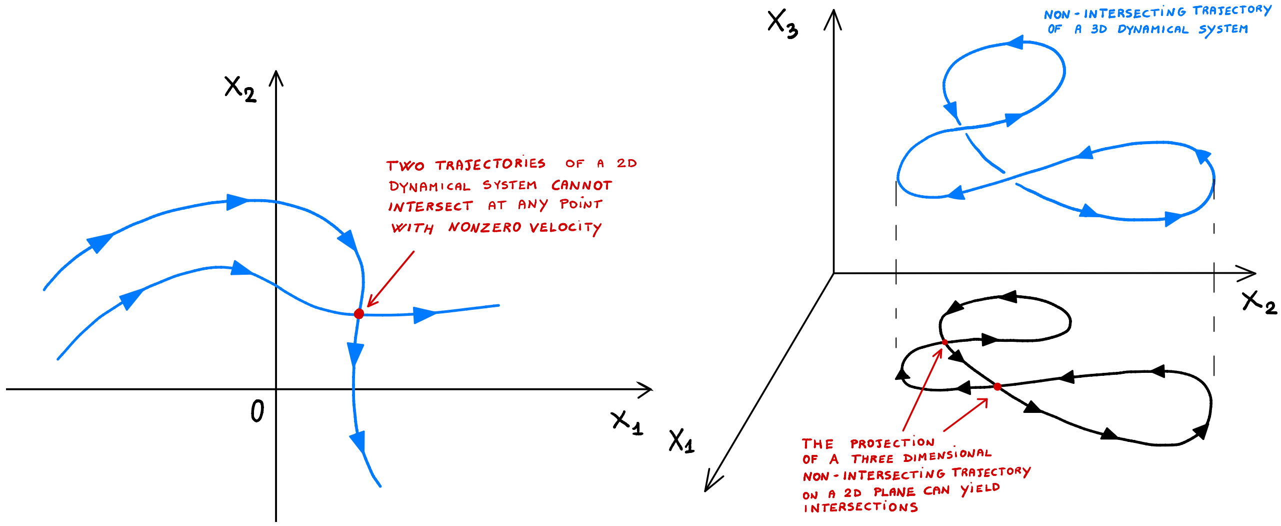

Figure 2: It is impossible for two trajectories to intersect with nonzero velocity any point in phase space. This would make $\boldsymbol{f}(\boldsymbol{x})$ non-unique at such point and also violate the existence and uniqueness Theorem 1.

Global solutions. If $\boldsymbol{f}(\boldsymbol{x})$ is Lipschitz continuous on the entire space $\mathbb{R}^n$ then the solution to the initial value problem (1) is *global*. This means that the solution exists and is unique for all $t \geq 0$. In fact, $x(t)$ never exits the domain in which $\boldsymbol{f}(\boldsymbol{x})$ is Lipschitz continuous, and therefore we can extend τ in Theorem 1 to infinity. It is important to emphasize that existence and uniqueness of the solution to (1) has nothing to do with the smoothness of $\boldsymbol{f}(\boldsymbol{x})$ but rather with the rate at which $\boldsymbol{f}(\boldsymbol{x})$ grows or decays.

Example: The solution to both dynamical systems (8) and (11) is global in time, meaning that exists and is unique for all $t \geq 0$. In fact the right hand side of such systems if globally Lipschitz in $\mathbb{R}^2$ and $\mathbb{R}^n$, respectively.

Flow generated by nonlinear dynamical systems. The solution of initial value problem (1) depends on $\boldsymbol{f}(\boldsymbol{x})$ and $\boldsymbol{x}_0$. As before, we will denote the dependence of the solution $\boldsymbol{x}(t)$ on $\boldsymbol{x}_0$ as $\boldsymbol{X}(t, \boldsymbol{x}_0)$, i.e.,

$$\boldsymbol{x}(t) = \boldsymbol{X}(t, \boldsymbol{x}_0). \tag{13}$$

Similarly to what we have seen for one-dimensional systems, it is not possible for two solutions corresponding to two different initial conditions to intersect at any finite time t. Otherwise we could use such intersection point as initial condition for (1) and conclude that there are two orbits emanating from such point (see Figure 2), hence violating the existence and uniqueness Theorem 1. This implies that $\boldsymbol{X}(t, \boldsymbol{x}_0)$ is invertible at each finite time[2] (see below), i.e., we can always identify which "particle" $\boldsymbol{x}_0$ sits at location $\boldsymbol{x}(t) = \boldsymbol{X}(t, \boldsymbol{x}_0)$ at time t. Moreover, it is impossible for two "particles" to collide at any finite time, or for one particle to split into two or more particles (Figure 2).

Theorem 2 (Regularity of the flow with respect to $\boldsymbol{x}_0$)**.** Let $D \subset \mathbb{R}^n$ be an open set, $\boldsymbol{x}_0 \in D$. If $\boldsymbol{f} : D \to \mathbb{R}^n$ is Lipschitz continuous in D then the flow $\boldsymbol{X}(t, \boldsymbol{x}_0)$ generated by the initial value problem (1), is continuous in $\boldsymbol{x}_0$. Moreover, If $\boldsymbol{f}(\boldsymbol{x})$ is of class $C^k(D)$ (continuously differentiable k-times in D) in D then $\boldsymbol{X}(t, \boldsymbol{x}_0)$ is of class $C^k(D)$ relative to $\boldsymbol{x}_0$ (continuously differentiable k-times with respect to $\boldsymbol{x}_0$).

[2] Solutions corresponding to different initial conditions can, however, intersect at $t = \infty$, e.g., when there exist an attracting set.

In summary, Theorem 2 states that the smoother $\boldsymbol{f}(\boldsymbol{x})$, the smoother the dependency of the flow $\boldsymbol{X}(t, \boldsymbol{x}_0)$ on $\boldsymbol{x}_0$. The n-dimensional function $\boldsymbol{X}(t, \boldsymbol{x}_0)$ is called *flow generated by the dynamical system* (1), and it represents the full set of solutions to (1) for each initial condition $\boldsymbol{x}_0$.

Theorem 3 (Regularity of the flow in time)**.** Let $D \subset \mathbb{R}^n$ be an open set, $\boldsymbol{x}_0 \in D$. if $\boldsymbol{f}(\boldsymbol{x})$ is of class C^k in D (continuously differentiable k-times in D with continuous derivative), then $\boldsymbol{X}(t, \boldsymbol{x}_0)$ is of class C^{k+1} in time for all $t \in [0, \tau[$, where τ is the time at which $X(t, \boldsymbol{x}_0)$ exits the domain D.

Properties of the flow: The flow $\boldsymbol{X}(t, \boldsymbol{x}_0)$ satisfies the following properties

- $\boldsymbol{X}(0, \boldsymbol{x}_0) = \boldsymbol{x}_0$. This means that at $t = 0$ the mapping $\boldsymbol{X}(0, \boldsymbol{x}_0) = \boldsymbol{x}_0$ is the identity.
- $\boldsymbol{X}(t, \boldsymbol{x}_0)$ is invertible for all t for which the solution to the initial value problem (1) exists and is unique.
- $\boldsymbol{X}(t, \boldsymbol{x}_0)$ satisfies the composition rule $\boldsymbol{X}(t+s, \boldsymbol{x}_0) = \boldsymbol{X}(t, \boldsymbol{X}(s, \boldsymbol{x}_0)) = \boldsymbol{X}(s, \boldsymbol{X}(t, \boldsymbol{x}_0))$. This property is called "semi-group property" of the flow and it follows from the fact that we can restart integration of the ODE (1) at time t (or time s) from the new initial condition $\boldsymbol{X}(t, \boldsymbol{x}_0)$ (or $\boldsymbol{X}(s, \boldsymbol{x}_0)$) to get to the final integration time $s + t$. Again, this property holds because of the existence and uniqueness theorem 1.
- The flow $\boldsymbol{X}(t, \boldsymbol{x}_0)$ satisfies the system of first-order PDEs
$$\begin{cases} \dfrac{\partial \boldsymbol{X}(t, \boldsymbol{x}_0)}{\partial t} - \boldsymbol{f}(\boldsymbol{x}_0) \cdot \nabla_{\boldsymbol{x}_0} \boldsymbol{X}(t, \boldsymbol{x}_0) = 0 \\ \boldsymbol{X}(0, x_0) = \boldsymbol{x}_0 \end{cases} \tag{14}$$
- The inverse flow $\boldsymbol{X}_0(t, \boldsymbol{x})$ satisfies the system of first-order PDEs
$$\begin{cases} \dfrac{\partial \boldsymbol{X}_0(t, \boldsymbol{x})}{\partial t} + \boldsymbol{f}(\boldsymbol{x}) \cdot \nabla_{\boldsymbol{x}} \boldsymbol{X}_0(t, \boldsymbol{x}) = 0 \\ \boldsymbol{X}_0(0, x) = \boldsymbol{x} \end{cases} \tag{15}$$

Geometric approach. The flow $\boldsymbol{X}(t, \boldsymbol{x}_0)$ generated by the ODE system (1) maps any initial condition $\boldsymbol{x}_0$ to the solution of the ODE at time t. If we think as $\boldsymbol{x}_0$ as the initial position of a particle in $\mathbb{R}^n$, then from elementary mechanics we know that $d\boldsymbol{X}(0, \boldsymbol{x}_0)/dt = \boldsymbol{f}(\boldsymbol{x}_0)$ represents the velocity of such particle. Hence, given $\boldsymbol{f}(\boldsymbol{x})$ we can immediately sketch the *vector field*[3] associated with the dynamical system, which represents where a particle sitting at any particular location in phase space is heading to To clarify this idea, consider the following two-dimensional dynamical system

$$\begin{cases} \dfrac{dx_1}{dt} = f_1(x_1, x_2) \\ \dfrac{dx_2}{dt} = f_2(x_1, x_2) \end{cases} \tag{16}$$

In Figure 3 we plot the velocity vector

$$\boldsymbol{f}(\boldsymbol{x}) = (f_1(\boldsymbol{x}), f_2(\boldsymbol{x})) \tag{17}$$

[3] A vector field is a vector that is continuously indexed by one or more variables. For one-dimensional dynamical systems the vector field $f(x)$ is indexed by coordinate x, and it is represented by a vector sitting on a line. For two-dimensional system the vector $\boldsymbol{f}(\boldsymbol{x}) = (f_1(\boldsymbol{x}), \boldsymbol{f}_2(\boldsymbol{x}))$ is a vector with two components: one along x_1 and the other along x_2. For three-dimensional systems the vector field has three components, and so on so forth.

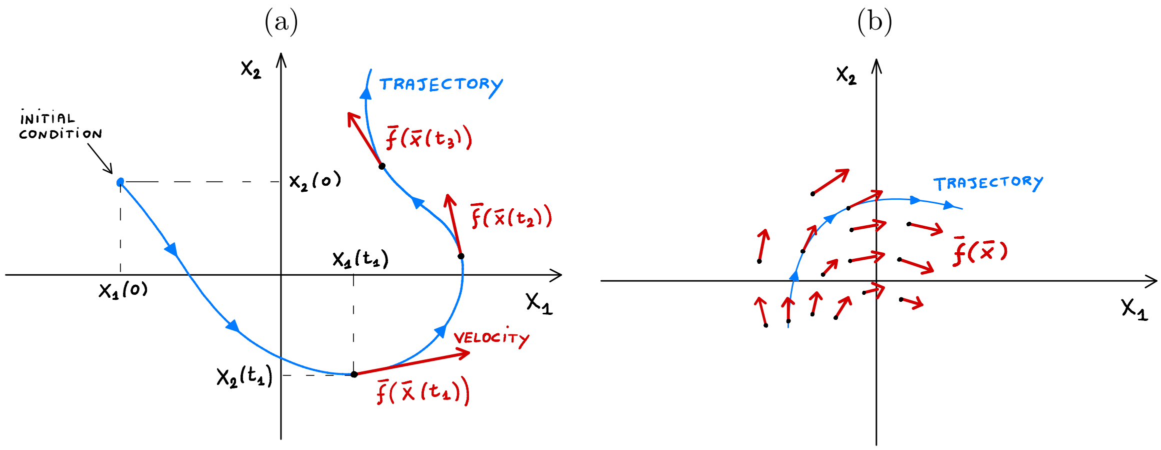

Figure 3: Geometric approach in 2D. (a) Sketch of a trajectory in the phase plane and associated velocity vectors $\boldsymbol{f}(\boldsymbol{x}(t))$. This process can be reversed in the sense that we can also guess how the trajectory $\boldsymbol{x}(t)$ looks like by plotting a bunch of velocity vectors $\boldsymbol{f}(\boldsymbol{x})$ evaluated at different points $\boldsymbol{x}$ in the plane (b).

along a trajectory $\boldsymbol{x}(t)$. This process can be reversed in the sense that we can guess how the trajectory $\boldsymbol{x}(t)$ looks like by plotting a sufficiently large number of velocity vectors $\boldsymbol{f}(\boldsymbol{x})$ evaluated at points nearby a point of interest in the phase space (see Figure 4(b)).

Example: In Figure 4 we plot the vector fields and corresponding trajectories defined by following two-dimensional dynamical systems

$$\begin{cases} \dot{x}_1 = 2x_1x_2 - 1 \\ \dot{x}_2 = -x_1^2 - x_2^2 + 10 \end{cases} \tag{18}$$

and

$$\begin{cases} \dot{x}_1 = x_2 \\ \dot{x}_2 = -\sin(x_1) - \dfrac{1}{10}x_2 \end{cases} \qquad \text{(pendulum with friction)} \tag{19}$$

As we shall see hereafter, the curves $\dot{x}_1 = 0$ and $\dot{x}_2 = 0$ are called *nullclines*. Fixed points are at the intersection of nullclines.

Fixed points. If the velocity vector $\boldsymbol{f}(\boldsymbol{x})$ is equal to zero at some point $\boldsymbol{x}^* \in \mathbb{R}^n$ then any particle placed at that point won't move at all as time increases. These points are called *fixed points* (or *equilibria*) of the dynamical system (1). Mathematically, we can define a fixed point $\boldsymbol{x}^* \in \mathbb{R}^n$ as

$$\boldsymbol{X}(t, \boldsymbol{x}^*) = \boldsymbol{x}^* \qquad \text{for all } t \geq 0. \tag{20}$$

By differentiating this previous equation with respect to time we obtain

$$\frac{\partial \boldsymbol{X}(t, \boldsymbol{x}^*)}{\partial t} = \boldsymbol{f}(\boldsymbol{X}(t, \boldsymbol{x}^*)) = \boldsymbol{f}(\boldsymbol{x}^*) = 0. \tag{21}$$

Therefore, the fixed points of the system (1) are solutions to the nonlinear system of equations

$$\boldsymbol{f}(\boldsymbol{x}) = \boldsymbol{0}. \tag{22}$$

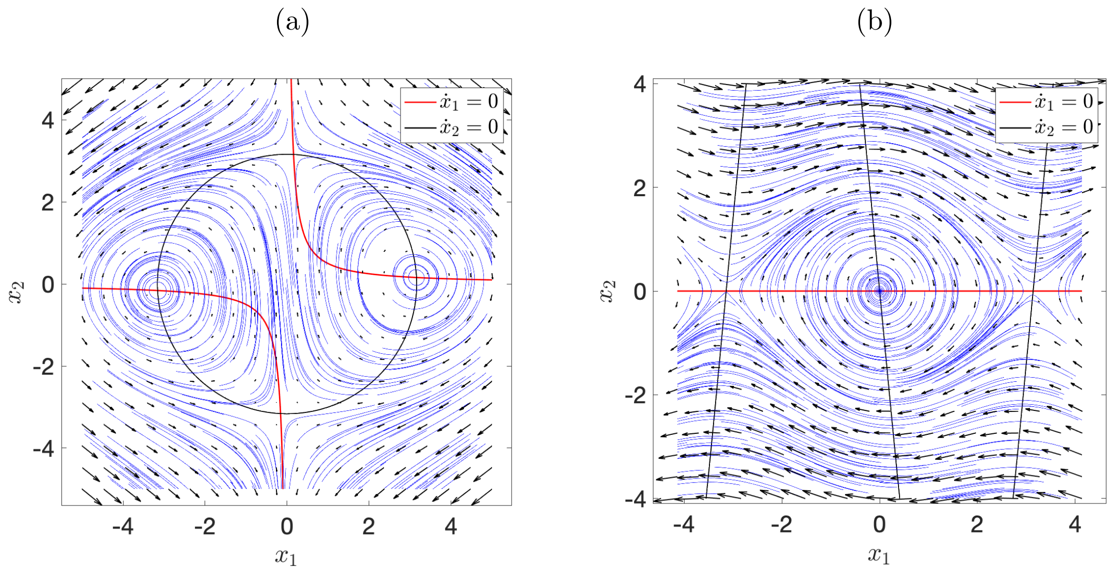

Figure 4: Vector field and trajectories generated by the two-dimensional nonlinear dynamical system (18) (Figure (a)), and (19) (Figure (b)). Shown are also the nullclines for both systems.

This system can be written as

$$\begin{cases} f_1(x_1, \ldots, x_n) = 0 \\ f_2(x_1, \ldots, x_n) = 0 \\ \quad \vdots \\ f_n(x_1, \ldots, x_n) = 0 \end{cases} \tag{23}$$

In this form, it is clear that the fixed points (if any) of the system (1) lie at the intersection of zero level sets[4] of n functions f_j each one of which is n-dimensional. In two-dimensions such zero level sets are identified by the intersection of two surfaces $f_1(x_1, x_2)$ and $f_1(x_1, x_2)$ with the (x_1, x_2) plane

$$\begin{cases} f_1(x_1, x_2) = 0 \\ f_2(x_1, x_2) = 0 \end{cases} \tag{24}$$

In two dimensions, the zero level sets of $f_1(x_1, x_2)$ and $f_1(x_1, x_2)$ are called *nullclines*. The reason for the definition is that the vector field $\boldsymbol{f}(x_1, x_2)$ is vertical at all points sitting on the nullcline $f_1(x_1, x_2) = 0$, and horizontal at all point sitting on the nullcline $f_2(x_1, x_2) = 0$. Correspondingly, the trajectories intersect the nullclines $f_1(x_1, x_2) = 0$ and $f_2(x_1, x_2) = 0$ vertically and horizontally, respectively (see Figure 4).

Example: Let us calculate the nullclines of the dynamical system (18). The first nullcline is

$$f_1(x_1, x_2) = 2x_1x_2 - 1 = 0 \qquad \Leftrightarrow \qquad x_2 = \frac{1}{2x_1} \qquad (\text{nullcline } \dot{x}_1 = 0), \tag{25}$$

i.e., the hyperbola depicted in red in Figure 4(a). The second nullcline is a circle with radius $\sqrt{10}$ centered at the origin

$$f_2(x_1, x_2) = -x_1^2 - x_2^2 + 10 \qquad \Leftrightarrow \qquad x_1^2 + x_2^2 = 10 \qquad (\text{nullcline } \dot{x}_2 = 0). \tag{26}$$

[4]The calculation of the fixed points can be done analytically only for prototype dynamical systems. In general, computing the fixed points requires a root-finding numerical algorithm for nonlinear systems of algebraic equations, e.g., the Newton's method.

This is depicted in black in Figure 4(a). We also see that the trajectories of the system intersect the nullcline $\dot{x}_2 = 0$ (black curve) horizontally. In fact, such nullcline represents the set of points in the phase plane where the velocity has zero vertical component. Similarly the trajectories intersect the nullcline $\dot{x}_1 = 0$ (red curve) vertically. The four fixed points of the system are at the intersection of the nullclines and can be computed analytically (left as exercise).

Example: Let us calculate the nullclines and the fixed points of the dynamical system (19). The first nullcline is

$$f_1(x_1, x_2) = x_2 = 0 \qquad \Leftrightarrow \qquad x_2 = 0 \qquad (\text{nullcline } \dot{x}_1 = 0), \tag{27}$$

Such nullcline is plotted in red in Figure 4(b). The second nullcline is

$$f_2(x_1, x_2) = \sin(x_1) - x_2/10 \qquad \Leftrightarrow \qquad x_2 = 10\sin(x_1) \qquad (\text{nullcline } \dot{x}_2 = 0), \tag{28}$$

and it is plotted in black in Figure 4(b).. The fixed points are at the intersection of the nullclines. In this case we obtain two physically different fixed points:

$$x_1^* = (0, 0) \qquad x_2^* = (\pi, 0) \tag{29}$$

corresponding to a pendulum in a vertical position, i.e., $x_1 = 0$ or $x_1 = \pi$ with zero velocity $x_2 = 0$.

Analysis of fixed points. A quick look at the phase portraits in Figure 4 suggests that the dynamics in a neighborhood of a fixed point can be quite different. Such dynamics can often be computed via a linearization process that is similar to the process we used in one-dimensional dynamical systems. The idea is "zoom-in" on a fixed point $\boldsymbol{x}^*$ and compute the orbits of the dynamical systems in a small neighborhood of $\boldsymbol{x}^*$ by solving a linearized version of the system (1). To this end, consider an initial condition $\boldsymbol{x}_0$ that is very close $\boldsymbol{x}^*$, and define the perturbation

$$\boldsymbol{\eta}(t, \boldsymbol{x}_0) = \boldsymbol{X}(t, \boldsymbol{x}_0) - \boldsymbol{x}^*. \tag{30}$$

By expanding $\boldsymbol{f}(\boldsymbol{X}(t, \boldsymbol{x}_0)) = \boldsymbol{f}(\boldsymbol{x}^* + \boldsymbol{\eta}(t, \boldsymbol{x}_0))$ in a neighborhood of $\boldsymbol{x}^*$, i.e., for small $\boldsymbol{\eta}(t, \boldsymbol{x}_0)$ we obtain

$$\boldsymbol{f}(\boldsymbol{x}^* + \boldsymbol{\eta}(t, \boldsymbol{x}_0)) = \underbrace{\boldsymbol{f}(\boldsymbol{x}^*)}_{=\boldsymbol{0}} + \boldsymbol{J}_{\boldsymbol{f}}(\boldsymbol{x}^*)\boldsymbol{\eta}(t, \boldsymbol{x}_0) + \cdots, \tag{31}$$

where

$$\boldsymbol{J}_{\boldsymbol{f}}(\boldsymbol{x}^*) = \begin{bmatrix} \dfrac{\partial f_1(\boldsymbol{x}^*)}{\partial x_1} & \cdots & \dfrac{\partial f_1(\boldsymbol{x}^*)}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \dfrac{\partial f_n(\boldsymbol{x}^*)}{\partial x_1} & \cdots & \dfrac{\partial f_n(\boldsymbol{x}^*)}{\partial x_n} \end{bmatrix} \tag{32}$$

is the Jacobian[5] of $\boldsymbol{f}(\boldsymbol{x})$ evaluated at the fixed point $\boldsymbol{x}^*$. Hence, the first-order approximation of the nonlinear dynamical system (1) at $\boldsymbol{x}^*$ can be written as

$$\begin{cases} \dfrac{d\boldsymbol{\eta}}{dt} = \boldsymbol{J}_{\boldsymbol{f}}(\boldsymbol{x}^*)\boldsymbol{\eta} \\ \boldsymbol{\eta}(0, \boldsymbol{x}_0) = \boldsymbol{x}_0 - \boldsymbol{x}^* \end{cases} \tag{33}$$

[5]The Jacobian of $\boldsymbol{f}(\boldsymbol{x})$ is a matrix-valued function that takes in a function $\boldsymbol{f}(\boldsymbol{x})$ and it returns a $n \times n$ matrix-valued function. The entries of such Jacobian matrix are functions. Of course, if we evaluate the Jacobian of $\boldsymbol{f}(\boldsymbol{x})$ at a specific point $\boldsymbol{x}^*$ then we obtain a matrix with real entries (provided $\boldsymbol{f}$ is real).

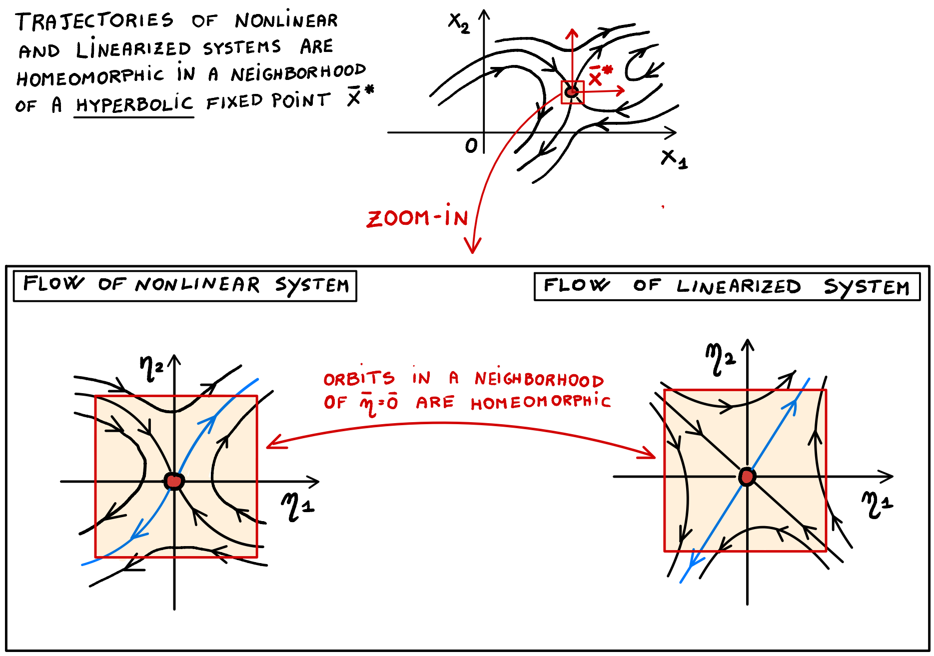

Figure 5: Geometric meaning of the Hartmman-Grobman Theorem 4. The trajectories of a nonlinear system in a neighborhood of any hyperbolic fixed point are homeomorphic to the trajectories of the linearized system at $\boldsymbol{x}^*$. This means that the trajectories of the nonlinear and linearized system are not exactly the same in the aforementioned neighborhood of $\boldsymbol{x}^*$, but they can be mapped to each other by a continuous transformation that has a continuous inverse.

Theorem 4 (Hartman-Grobman)**.** Let $\boldsymbol{x}^* \in \mathbb{R}^n$ be a fixed point of the dynamical system (1). If the Jacobian (32) has no eigenvalue with zero real part then there exists a homeomorphism (i.e., continuous invertible mapping with continuous inverse) defined on some neighborhood of $\boldsymbol{x}^*$ that takes orbits of the system (1) and maps them into orbits of the linearized system (30)-(33). The mapping preserves the orientation of the orbits.

Remark: Theorem (4) is saying that if $\boldsymbol{x}^*$ is a hyperbolic[6] fixed point then the flow of the nonlinear dynamical system (1) nearby $\boldsymbol{x}^*$ is "homemorphic" (i.e., it can be mapped back and forth by a continuous transformation) to the flow of the linearized system (30)-(33).

At this point it is natural to ask the following questions:

1. To study the flow nearby a hyperbolic fixed point of nonlinear system we need to compute the flow of the linear system (33). Is there a general method to compute such flow? Note that flows of linear systems are very important on their own as there are many system that are actually are linear (e.g., the discretized PDE system (11)).

2. What happens if the fixed point is non-hyperbolic? As we will see, in this case we need to use a generalization of the Hartman-Grobman theorem known as *center manifold* theorem.

[6] A fixed point $\boldsymbol{x}^*$ is called *hyperbolic* if the Jacobian of $\boldsymbol{J_f}(\boldsymbol{x}^*)$ has no eigenvalue with zero real part.

Appendix A: Elementary numerical methods for systems of ODEs

As before, we can re-write the Cauchy problem (1) as an integral equation

$$\boldsymbol{x}(t) = \boldsymbol{x}(0) + \int_0^t \boldsymbol{f}(\boldsymbol{x}(s))ds. \tag{34}$$

This form is quite handy to derive numerical methods to solve (1) based on quadrature rules applied to the one-dimensional integral at the right hand side. For instance, consider a partition of the $[0, T]$ into an evenly-spaced grid points such that $t_{i+1} = t_i + \Delta t$, and write (34) within each time interval

$$\boldsymbol{x}(t_{i+1}) = \boldsymbol{x}(t_i) + \int_{t_i}^{t_{i+1}} \boldsymbol{f}(\boldsymbol{x}(s))ds. \tag{35}$$

Explicit midpoint method. By approximating the integral at the right hand side of (35), e.g., using the midpoint rule yields

$$\int_{t_i}^{t_{i+1}} \boldsymbol{f}(\boldsymbol{x}(s))ds \simeq \Delta t \boldsymbol{f}\left(\boldsymbol{x}\left(t_i + \frac{\Delta t}{2}\right)\right) \tag{36}$$

At this point, we can approximate $\boldsymbol{x}(t_i + \Delta t/2)$ using the Euler forward method

$$\boldsymbol{x}\left(t_i + \frac{\Delta t}{2}\right) \simeq \boldsymbol{x}(t_i) + \frac{\Delta t}{2}\boldsymbol{f}(\boldsymbol{x}(t_i)) \tag{37}$$

to obtain the *explicit midpoint method*

$$\boldsymbol{x}(t_{i+1}) = \boldsymbol{x}(t_i) + \Delta t \boldsymbol{f}\left(\boldsymbol{x}(t_i) + \frac{\Delta t}{2}\boldsymbol{f}(\boldsymbol{x}(t_i))\right). \tag{38}$$

The explicit midpoint method is a one-step method that belongs to the larger class of Runge-Kutta methods.

Appendix B: Equivalent norms in $\mathbb{R}^n$

As is well known, all norms defined in a finite-dimensional vector space such as $\mathbb{R}^n$ are *equivalent*. This means that if we pick two arbitrary norms in $\mathbb{R}^n$, say $\|\cdot\|_a$ and $\|\cdot\|_b$, then there exist two numbers C_1 and C_2 such that

$$C_1 \|\boldsymbol{x}\|_a \leq \|\boldsymbol{x}\|_b \leq C_2 \|\boldsymbol{x}\|_a \quad \text{for all} \quad \boldsymbol{x} \in \mathbb{R}^n. \tag{39}$$

The most common norms in $\mathbb{R}^n$ are

$$\|\boldsymbol{x}\|_\infty = \max_{k=1,..,n} |x_k|, \tag{40}$$

$$\|\boldsymbol{x}\|_1 = \sum_{k=1}^n |x_k|, \tag{41}$$

$$\|\boldsymbol{x}\|_2 = \left(\sum_{k=1}^n |x_k|^2\right)^{1/2}, \tag{42}$$

$$\vdots \tag{43}$$

$$\|\boldsymbol{x}\|_p = \left(\sum_{k=1}^n |y_k|^p\right)^{1/p} \qquad p \in \mathbb{N} \setminus \{\infty\}. \tag{44}$$

Based on these definitions it can be shown that, e.g., that

$$
\begin{aligned}
\|\boldsymbol{x}\|_\infty &\le \|\boldsymbol{x}\|_1 \le n\,\|\boldsymbol{x}\|_\infty\,, &(45)\\
\|\boldsymbol{x}\|_2 &\le \|\boldsymbol{x}\|_1 \le \sqrt{n}\,\|\boldsymbol{x}\|_2\,, &(46)\\
\|\boldsymbol{x}\|_\infty &\le \|\boldsymbol{x}\|_2 \le \sqrt{n}\,\|\boldsymbol{x}\|_\infty\,. &(47)
\end{aligned}
$$

Therefore if the $\boldsymbol{f}(\boldsymbol{x})$ is Lipschitz continuous in D with respect to the 1-norm, i.e.,

$$\|\boldsymbol{f}(\boldsymbol{x}_1)-\boldsymbol{f}(\boldsymbol{x}_2)\|_1 \le L_1\,\|\boldsymbol{x}_1-\boldsymbol{x}_2\|_1 \quad \text{for all} \quad \boldsymbol{x}_1,\boldsymbol{x}_2\in D \tag{48}$$

then it is also Lipschitz continuous in D with respect to the uniform norm. In fact, by using (45) we have

$$\|\boldsymbol{f}(\boldsymbol{x}_1)-\boldsymbol{f}(\boldsymbol{x}_2)\|_\infty \le \underbrace{L_1 n}_{L_\infty}\,\|\boldsymbol{x}_1-\boldsymbol{x}_2\|_\infty\,. \tag{49}$$

Similarly, by using (46), we see that if $\boldsymbol{f}$ is Lipschitz in D with respect to the 1-norm then $\boldsymbol{f}$ is Lipschitz in D with respect to the 2-norm.

Appendix B: Proof of Lemma 1

Let $D \subseteq \mathbb{R}^n$ be a compact convex domain and let

$$M = \max_{\boldsymbol{x}\in D}\left|\frac{\partial f_j(\boldsymbol{x})}{\partial x_i}\right|. \tag{50}$$

Clearly M exists and is finite because we assumed that D is compact and that $\boldsymbol{f}$ is of class C^1 in D[7]. Consider two points $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ in D, and the line that connects $\boldsymbol{x}_1$ to $\boldsymbol{x}_2$, i.e.,

$$\boldsymbol{z}(s) = (1-s)\boldsymbol{x}_1 + s\boldsymbol{x}_2 \qquad s\in[0,1]. \tag{51}$$

Since D is convex, we have that the line $\boldsymbol{z}(s)$ lies entirely within D. Therefore we can use the mean value theorem applied to the one-dimensional function $f_i(\boldsymbol{z}(s))$ ($s\in[0,1]$) to obtain

$$f_i(\boldsymbol{x}_2) - f_i(\boldsymbol{x}_1) = \nabla f_i(\boldsymbol{z}(s^*),t)\cdot(\boldsymbol{x}_2-\boldsymbol{x}_1) \qquad \text{for some } s^*\in[0,1]. \tag{52}$$

By taking the absolute value and using the Cauchy-Schwartz inequality we obtain

$$
\begin{aligned}
|f_i(\boldsymbol{x}_1)-f_i(\boldsymbol{x}_1)|^2 =& \left|\sum_{j=1}^n \frac{\partial f_i(\boldsymbol{z}(s^*))}{\partial x_j}(x_{2j}-x_{1j})\right|^2\\
\le& \left|\sum_{j=1}^n \frac{\partial f_i(\boldsymbol{z}(s^*))}{\partial x_j}\right|^2 \left|\sum_{j=1}^n (x_{2j}-x_{1j})\right|^2\\
\le& nM^2\,\|\boldsymbol{x}_2-\boldsymbol{x}_1\|_2^2\,.
\end{aligned} \tag{53}
$$

This implies that

$$\|\boldsymbol{f}(\boldsymbol{x}_2)-\boldsymbol{f}(\boldsymbol{x}_1)\|_2 \le \underbrace{nM}_{L_2}\,\|\boldsymbol{x}_2-\boldsymbol{x}_1\|_2\,. \tag{54}$$

i.e., $\boldsymbol{f}(\boldsymbol{y},t)$ is Lipschitz continuous in the 2-norm, or any other norm that is equivalent to the 2-norm. In particular, by using the inequalities (45)-(47) we have that $\boldsymbol{f}(\boldsymbol{x})$ is Lipschitz continuous relative to the 1-norm.

[7] A compact domain is by definition bounded and closed. The minimum and maximum of a continuous function in defined on a compact domain is attained at some points within the domain or on its boundary. Note that this is not true if the domain is not compact. For example, the function $f(y) = 1/y$ is continuously differentiable on $]0,1]$ (bounded domain by not compact), but the function is unbounded on $]0,1]$.