

University of Chicago - Autumn 2026

Applied dynamical systems

Vincent X. Wang

October 18, 2026

These notes are primarily based off of lecture notes taken from CAAM/STAT 31410 at University of Chicago, taught by Nisha Chandramoorthy during Autumn 2026. The course follows [\[Mei17\]](#) loosely.

Contents

<i>What is a dynamical system? (and examples)</i>	2
<i>Linear systems</i>	3
<i>Review of linear algebra</i>	5
<i>More on linear systems</i>	7
<i>Complex eigenvalues</i>	8
<i>Intro to stability analysis</i>	10
<i>Non-linear dynamics</i>	11
<i>An aside on geometry</i>	11
<i>Contraction maps</i>	11
<i>Existence and uniqueness of ODEs</i>	12

What is a dynamical system? (and examples)

“A dynamical system is a set of variables and their evolution in time.”

More formally, let M be a compact phase space. We call $\varphi^t : M \rightarrow M$ the *dynamical system*. We may enforce $t \in \mathbb{Z}$ (wherein this is called a map) or $t \in \mathbb{R}$ (wherein this is called a flow¹). We say a map φ^t is an *autonomous* map if

$$\varphi^t = \underbrace{\varphi \circ \dots \circ \varphi}_{t \text{ times}}$$

We say φ^t is a *non-autonomous map*² if

$$\varphi^t = \varphi_t \circ \dots \circ \varphi_1.$$

The main moral of all of this distinction is make clear that *dynamics are not differential equations and/or group theory* (in the same sense that what is happening at a point in time is not always fixed, as those subjects can be thought of as autonomous dynamical systems). In general, this definition of a dynamical system is quite loose and thus dynamical systems are a very interdisciplinary subject that draws from many areas of mathematics (e.g., perturbation theory is closely tied to differential topology, probability theory when dealing with randomness).

Example 1. Let $d = 1$ and $M = [0, 1]$, and $\varphi(x) = x$. This is kind of uninteresting because nothing happens.

Non-example 2. Let $d = 1$ and $M = [0, 1]$. Let $\varphi(x) = 2x$. This fails to be a dynamical system because φ does not map M to M .

Example 3. Let

$$\frac{dx}{dt} = v(t, x) \quad x \in \mathbb{R}^d.$$

This is implicitly a non-autonomous dynamical system. This can be rewritten:

$$\frac{d\varphi^t(x)}{dt} = v(t, \varphi^t(x)) \quad x \in \mathbb{R}^d.$$

We think of x as an “initial condition” rather than a position (given by $\varphi^t(x)$).

We call $\{\varphi^t(x)\}_t$ the *orbit* (or trajectory) of x ; dynamical systems is the study of all possible orbits. People are also very interested in perturbations of dynamical systems, but there are many ways and perspectives to look at this.

Definition 4. A *linear* dynamical system can be written as

$$\frac{dx}{dt} = Ax \quad x \in \mathbb{R}^d.$$

Example 5. Let $\phi(x) = 2x \% 1$, i.e.,

$$\phi(x) = \begin{cases} 2x & x < 1/2 \\ 2x - 1 & x \geq 1/2 \end{cases}.$$

We call this a *linear hyperbolic map*, but this is an easy example of a nonlinear dynamical system.

¹ Note that we also allow invertible flows/maps here, which may not necessarily always be the case.

² with abuse of notation for φ^t

Example 6. We can think of *next-token prediction* as a (random) dynamical system too. Let X_t be the tokens. Then,

$$\varphi^t(x) \sim P(X_t | X_1, \dots, X_{t-1}).$$

In this example, M is the discrete space of strings. This is a “non-autonomous” dynamical system.

Example 7. The Navier-Stokes equations (which have been in the news) are an autonomous dynamical system. But the solution of the PDE lives in a function space, so the phase space is a function space (i.e. L^2). Here, our ‘ x ’ is the initial velocity (which is a function). We will not study this because M is not compact and is also infinite-dimensional.

Linear systems

Definition 8. A *linear system* is a dynamical system given by

$$\frac{\partial \varphi^t}{\partial t}(x) = A\varphi^t(x)$$

where A is a matrix.

For now, let us assume the solutions exist and are unique. Linear systems are important because they are just about the only systems that have closed-form solutions. In practice, though, nearly everything we write down does not have a closed-form solution and needs to be approximated. Recall that the solution is given by

$$\varphi^t(x) = e^{tA}x,$$

which is further discussed in the book.

We have two motivations for studying linear systems:

- They arise in linear stability analysis / perturbation theory.
- They are kind of basic for analysis.

In contemporary research, we don’t really think about continuous-time dynamics anymore – we just take the discretization as the “ground-truth”. This is just because the systems we look at now are so complicated we can’t hope to garner much about an analytical solution.


Let us just look at maps instead of ODEs. This case is very easy: the solution to

$$x_{t+1} = Bx_t$$

is given by

$$\varphi^t(x) = B^t x.$$


Example 9. Let $B = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda \end{pmatrix}$. What is $\varphi^t(x)$? In the case of $\lambda \neq 1$,



. We call the origin a *node* in this case. It is *stable* if $0 < \lambda < 1$ and *unstable* if $\lambda > 1$.


We should think of coordinates $x^{(1)}$ as a “projection” of x onto the standard basis.

Example 10. Let $B = \begin{pmatrix} \lambda & 0 \\ 0 & \mu \end{pmatrix}$. What if $\lambda > \mu > 0$?



Example 11. Let $B = \lambda \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}$. We call the origin a *focus*.

Example 12 (Hyperbolic system). Let $B = \begin{pmatrix} \lambda & 0 \\ 0 & \mu \end{pmatrix}$.



We call the origin a *saddle point*. THIS IS WRONG!

These are called phase portraits, but these only make sense for flows (and not maps)! As an exercise, one can prove for each of the exercises that the curve is invariant under the map; i.e., let a curve be written as $x^{(2)} = f(x^{(1)})$. Show that

$$x^{(2)}(\varphi^t(x)) = f(x^{(1)}(\varphi^t(x))).$$

What is the relationship between the map and the flow? Let

$$\begin{aligned}\dot{x} &= Ax \\ x_{t+1} &= Bx_t.\end{aligned}$$

We have discussed the latter map, but what is the continuous time flow when

$$B = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda \end{pmatrix},$$

then using the solution to the continuous ODE, we can write

$$A = \begin{pmatrix} \log \lambda & 0 \\ 0 & \log \lambda \end{pmatrix}.$$

Thus, we can't study the phase portraits with for *maps* with negative λ (but we certainly can for *flows*!).

Review of linear algebra

In this section, we consider **matrices with entries over the reals**, and briefly review some linear algebra.

When is a (square) matrix A diagonalizable (i.e., $A = P^{-1}\Lambda P$ where Λ is diagonal)? Such a diagonalization exists when the column space spans all of $\mathbb{R}^d$ (or when the eigenvectors span $\mathbb{R}^d$). If A is diagonalizable, it is easy to compute $A^{-1} = P\Lambda^{-1}P^{-1}$. Thus, a matrix and its inverse have the same eigenvectors.

We introduce two notions of multiplicity:

Definition 13. The **algebraic multiplicity** of an eigenvalue λ is the multiplicity of λ with respect to the characteristic polynomial $\det(A - \lambda I) = 0$. An eigenvalue is *simple* if its (algebraic) multiplicity is one.

Definition 14. The **geometric multiplicity** of λ is equal to $\dim(\text{null}(A - \lambda I))$.

If the algebraic multiplicity is equal to the geometric multiplicity for all eigenvalues, then the matrix is diagonalizable. However, in general, the following fact is true:

Fact 15. *The algebraic multiplicity is always greater than or equal to the geometric multiplicity.*

In some ways, the algebraic multiplicity is also geometric (here, we diverge from the book). Let us introduce the **Jordan decomposition**, which is the most general decomposition of a matrix that always exists. It is of the form

$$P^{-1}AP = J,$$

where J has the Jordan block structure

$$J_\lambda = \begin{pmatrix} \lambda & 1 & & \\ & \lambda & 1 & \\ & & \ddots & \ddots \\ & & & \lambda & 1 \\ & & & & \lambda \end{pmatrix}$$

and

$$J = \begin{pmatrix} J_{\lambda_1} & & \\ & J_{\lambda_2} & \\ & & \ddots \\ & & & J_{\lambda_k} \end{pmatrix},$$

where J_λ is a Jordan block.

To interpret the Jordan decomposition, we introduce generalized eigenvectors of A .

Definition 16. We say v is a **generalized eigenvectors** of A if, for some $k > 0$, $(A - \lambda I)^k v = 0$. Thus, the generalized eigenvalue satisfies $\det((A - \lambda I)^k) = 0$.

As an example, let us consider $k = 2$: If

$$(A - \lambda I)(A - \lambda I)v_2 = 0$$

and

$$(A - \lambda I)v_1 = 0,$$

then (check this)

$$(A - \lambda I)v_2 = v_1,$$

so we can relate the $k = 2$ generalized eigenvectors to the eigenvectors of A . We call v_1, v_2 a *chain*. It is a deep theorem that the span of the generalized eigenvalues $\{v_1, \dots, v_d\}$ span $\mathbb{R}^d$, which is equivalent to the fact that a Jordan decomposition exists. Call E^λ the *generalized eigenspace* associated with λ .

Fact 17. E^λ is invariant under the map and the flow (i.e., for flows, if $x \in E^\lambda$, then $e^{tA}x \in E^\lambda$).

This is the reason we discuss the Jordan form. Some fast facts:

Fact 18. The number of Jordan blocks J_λ for that eigenvalue is equal to the geometric multiplicity of that eigenvalue.

Fact 19. The size of the Jordan block is how many eigenvectors we can chain. For given λ , the sum of the sizes of the Jordan blocks J_λ is equal to the algebraic multiplicity.³

Now, when A is diagonalizable, let us reconsider the general solution to

$$\frac{\partial \varphi^t(x)}{\partial t} = A\varphi^t(x)$$

³ There is some discrepancy in that the book introduces algebraic multiplicity as something like the size of the largest Jordan block

is

$$\varphi^t(x) = \sum_{i=1}^d c_i v_i e^{\lambda_i t}.$$

It is simple to verify that this is a solution just by taking the derivative.

More on linear systems

Let us rediscuss the issue of invertibility. It is easy to tell when a map $x_{t+1} = Ax_t$ is invertible just by looking at the matrix A . It is more complicated for a flow. Recall again the general solution of a linear flow $\varphi^t(x) = e^{tA}x$, and the following facts:

Definition 20.

$$e^{tA} = I + tA + \frac{t^2 A^2}{2!} + \dots$$

Fact 21. Let $\|\cdot\|$ denote the operator norm. If $\|A\| < \infty$, then $\|e^{tA}\| < \infty$.

Proof. In the book. □

Fact 22.

$$e^{tA}A = Ae^{tA}.$$

Proof. This is easy to see because the powers of a matrix commute (i.e., $AA^n = A^nA$). □

One may now verify

$$\frac{\partial \varphi^t(x)}{\partial t} = e^{tA}Ax = Ae^{tA}x = A\varphi^t(x).$$

We have not assumed anything about A (other than it is bounded). This solution is unique via the following short proof. If $y(t)$ is a solution, then one can compute

$$\frac{\partial}{\partial t} (e^{-tA}y(t)) = 0,$$

so $y(t) = e^{tA}c$, and $y(0) = c$ implies $x = c$. Thus, $y(t) = e^{tA}x$.

The last point that is addressed is computing e^{tA} , and we follow a different approach than the book. Suppose, in the easy case, that A is diagonalizable with $A = P\Lambda P^{-1}$. One can check easily that

$$e^{tA} = Pe^{\Lambda t}P^{-1}.$$

Now note that $e^{\Lambda t}$ is very easy to compute, as it is just $\text{diag}(e^{\lambda_1 t}, \dots, e^{\lambda_n t})$. Observe that after making explicit the choice of basis and picking constants c_i , the solution is equivalent to $\varphi^t(x) = \sum_{i=1}^d c_i v_i e^{\lambda_i t}$.

More generally, let A be any matrix, not necessarily diagonalizable. Then, we can write $e^{tA} = Pe^{Jt}P^{-1}$, where the columns of P are generalized eigenvectors. What is e^{Jt} ? Suppose we have a Jordan block of size 2 with generalized eigenvectors v_1, v_2 , and corresponding

$$J_\lambda = \begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix}.$$

The action of A on some v_1 is the action of J_λ on $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$, which evaluates to $\lambda \begin{pmatrix} 1 \\ 0 \end{pmatrix}$. This does not really say anything interesting because the first generalized eigenvector v_1 of the chain is just the regular eigenvector. So now, we consider

$$J_\lambda \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ \lambda \end{pmatrix}.$$

In Euclidean coordinates, this is

$$J_\lambda v_2 = v_1 + \lambda v_2,$$

which is exactly what we did earlier. The point of the above computations is really to show that the action of A on the generalized eigenvectors gets us a “full-dimensional” subspace.

As an exercise, compute⁴:

⁴ Splitting J_λ into diagonal and nilpotent parts may be useful!

Fact 23.

$$e^{tJ_\lambda} = \begin{pmatrix} e^{t\lambda} & te^{t\lambda} \\ 0 & e^{t\lambda} \end{pmatrix}.$$

Using this, we get that the part of the solution corresponding to E_λ is $c_1 e^{t\lambda} v_1 + c_2 (te^{t\lambda} v_1 + e^{t\lambda} v_2)$. Can we extend this to $d \times d$ matrices?

Complex eigenvalues

We consider A with real entries, where $\|A\| < \infty$.⁵ In the context of dynamical systems, we want e^{tA} to be real so that the solution sends the phase space to itself. Recall

⁵ We do not need to specify what kind of norm we are working with here because of norm equivalence. As an exercise, one can show that a norm is continuous.

$$e^{tA} = \sum_n \frac{t^n A^n}{n!}$$

Suppose A has a complex eigenvalue $\lambda = a + bi$ and A can be diagonalizable $P^{-1}AP = \Lambda$ on $\mathbb{C}$. Let $v = u + iw$. Recall the full solution is

$$\sum c_i e^{t\lambda} v_i.$$

Let's just look at a term which could become complex:

$$\begin{aligned} e^{t\lambda} v &= e^{t(a+ib)}(u + iw) \\ &= e^{at} e^{ibt}(u + iw) \\ &= e^{at} (\cos bt + i \sin bt)(u + iw) \\ &= e^{at} (u \cos bt - w \sin bt) + i e^{at} (w \cos bt + u \sin bt) \end{aligned}$$

There is no “catch” here because we can think of restricting the complex dynamical system just to the real part. If we just put

$$\begin{bmatrix} u & w \end{bmatrix}^{-1} A \begin{bmatrix} u & w \end{bmatrix} = \begin{bmatrix} a & -b \\ b & a \end{bmatrix},$$

we get a sort of diagonal matrix. [why?] This is the *normal form* for a complex matrix.

More generally,

$$J_\lambda = \begin{bmatrix} \lambda & 1 & & & \\ & \lambda & 1 & & \\ & & \ddots & \ddots & \\ & & & \lambda & 1 \\ & & & & \lambda \end{bmatrix}$$

with $P_\lambda = [v_1, \dots, v_{\lambda_n}]$ generalized eigenvectors corresponding to λ . The book talks about something that *looks* like a Jordan form, which is a decomposition into $A = S + N$ where S is semi-simple (i.e., diagonalizable over $\mathbb{C}$) and N is nilpotent (there is some k for which $N^k = 0$ and $N^{k-1} \neq 0$). This is called the *additive Jordan decomposition*. We like to introduce the Jordan normal form instead because it still allows us to think of the operator as acting on a basis, and explicitly write out the change of bases via the generalized eigenvectors. They are a good basis to study the invariance of a linear dynamical system. But in the book's decomposition, it is unclear what's going on even though they are spiritually the same.

Let us see why: consider a Jordan block of A

$$J_\lambda = \begin{pmatrix} 2 & 1 \\ 0 & 2 \end{pmatrix} = \begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix} + \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}.$$

Consider the action $J_\lambda \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 2 \\ 0 \end{pmatrix}$, which represents J_λ on v_1 really in Euclidean space. This is exactly what S tells us how to do. The generalized eigenvector v_2 .

$$(A - 2I)^2 v_2 = (A - 2I)v_1 = 0,$$

so $N = A - 2I$ is nilpotent. We can also observe the more general fact that the order of nilpotency of N is precisely the maximum algebraic multiplicity.

The whole moral of what the book does is that e^{tA} is "easy" even when A isn't diagonalizable. As the book does it:

$$e^{tS} = P^{-1} e^{t\Lambda} P$$

$$e^{tN} = \sum_{n=0}^{\infty} \frac{(tN)^n}{n!} = \sum_{n=k-1}^{\infty} \frac{(tN)^n}{n!},$$

so we can literally just compute it. It turns out that this is also useful for analysis later.

Fact 24. Let $A = B + C$. Then.

$$e^A = e^B e^C$$

if and only if $BC - CB = 0$ (this is the *commutator* $\begin{bmatrix} B & C \end{bmatrix}$).

You can check the above and see what the problem is. We exploit

Fact 25. Let $A = S + N$ be the additive Jordan decomposition. Then,

$$\begin{bmatrix} S & N \end{bmatrix} = 0,$$

which directly implies

$$e^{tA} = PA^{t\Lambda}P^{-1} \sum_{n=k-1}^{\infty} \frac{(tN)^n}{n!}.$$

Intro to stability analysis

We introduce linear stability analysis. We can write

$$\mathbb{R}^d = \bigoplus_{\lambda} E_{\lambda} = E_u \oplus E_s \oplus E_c,$$

where

$$E_u = \{v \in E_{\lambda} : \operatorname{Re}(\lambda) > 0\}$$

$$E_s = \{v \in E_{\lambda} : \operatorname{Re}(\lambda) < 0\}$$

$$E_c = \{v \in E_{\lambda} : \operatorname{Re}(\lambda) = 0\}.$$


These are called the **unstable, stable, and center/neutral** subspaces⁶.

There are basically the parts of the system that go to 0 exponentially fast. Similarly, for $v \in E_u$,

$$e^{tA}v = e^{t\lambda}v,$$

so $\|e^{t\lambda}v\|$ runs off to infinity. Let's ignore E_c for now.

- No one told us E_u and E_s are necessarily orthogonal.
- The vectors in E_s get infinitely small and can be defined as so. But the same is not true for E_u , as any vector in E_s is in this set.
- E_u , however, can be defined as the set of vectors that go to 0 *backwards* in time!



More formally, E_s is the subspace that is all vectors v satisfying the condition that there exists $\alpha > 0, C > 0$ such that

$$\|e^{tA}v\| \leq Ce^{-\alpha t}\|v\|$$

⁶ Note why being real matters: what if $\lambda = 2i$?

for all $t > 0$. Furthermore, E_u is the subspace that is all vectors v satisfying the condition that there exists $\alpha > 0, C > 0$ such that

$$\|e^{tA}v\| \leq Ce^{-\alpha|t|}\|v\|$$

for all $t < 0$.

The center subspace E_c corresponds to rotations in the phase space.

Exercise: For the case of maps (which are not in the book, what are the characterizations of these subspaces)? Answer:

$$E_u = \{v \in E_\lambda : |\lambda| > 1\}$$

$$E_s = \{v \in E_\lambda : |\lambda| < 1\}$$

$$E_c = \{v \in E_\lambda : |\lambda| = 1\}.$$

Do the characterizations in terms of exponential blow-up change if ϕ is discrete? It turns out, not really (other than restricting t to the integers).

Non-linear dynamics

An aside on geometry

Consider an arbitrary dynamical system

$$\frac{\partial \phi^t(x)}{\partial t} = f(t, \phi^t(x)).$$

The differential geometric way to think about this is that the RHS is a vector field $f : M \rightarrow TM$. It so happens that in this class we will let $M = \mathbb{R}^d$, but this is good to keep in mind in general.

Contraction maps

Let us consider a very simple (but important!) non-linear dynamical system, the contraction maps.

Definition 26. Let M be Banach. A function $F : M \rightarrow M$ is a **contraction map** if $\|F(x) - F(y)\| \leq \lambda \|x - y\|$ for all $x, y \in M$ for some $\lambda \in (0, 1)$.

We study the dynamical system

$$x_{t+1} = F(x_t), \quad t \in \mathbb{Z}^+.$$

Via a simple proof, one can show that the orbit of any $x \in M$ all have the same limit (i.e., everything gets squished to a point). The sketch is that we should use Cauchy to show that the orbit converges to some x^* , then show that x^* is a fixed point.

Lemma 27 (Contraction mapping principle). *Let M be a complete normed⁷ space. Let $F : M \rightarrow M$ be a contraction map. Then, F has a unique fixed point and the orbit of any $x \in M$ converges exponentially fast to the fixed point.*

⁷ Metric is OK too.

Let $x \in M$ and

$$\frac{d\phi^t(x)}{dt} = f(\phi^t(x)).$$

Q: For what x does the solution exist in linear systems? A: Any x (as we were able to write out the solution, and this also implies uniqueness). For a non-linear system, because we cannot do such a thing, we need other ways to prove existence and uniqueness. We will do this via contraction maps in the next section.

Existence and uniqueness of ODEs

Let

$$\frac{\partial \phi^t(x)}{\partial t} = f(t, \phi^t(x)).$$

We will typically assume that functions are Lipschitz, which come equipped with some nice properties:

Theorem 28 (Rademacher's theorem). *A Lipschitz function is differentiable almost everywhere.*

Lemma 29. *Lipschitz implies uniform continuity (which implies continuity).*

Theorem 30 (Picard-Lindelöf). *Let⁸ $f : J \times B_b(x_0) \rightarrow \mathbb{R}^n$ be Lipschitz continuous in space with constant k and continuous in time. Then,*

$$\phi^t(x) = x + \int_{t_0}^t f(u, \phi^u(x)) du$$

exists and is unique for all $x \in B_b(x_0)$ in the time interval⁹ $J = [t_0 - a, t_0 + a]$ whenever $a = b/M$ where $M = \max_{(t,x) \in J \times B_b(x_0)} |f(t, x)|$.

Proof. There are some detailed proofs in [Mei17], but there are some main insights:

- For $u \in C^0(J, B_b(x_0))$, define the operator

$$T(u) := x_0 + \int_{t_0}^t f(\tau, u) d\tau.$$

The solution to the ODE is equivalent to finding a fixed point of T .

- The main idea is that we want T to be a contraction. This is done through the combination of two tools: first, we need to assume f is Lipschitz; second, we need to work in the Bielecki norm, defined by

$$\|f\|_L = \sup_{t \in J} e^{-L|t-t_0|} |f(t)|$$

for some L . See the following.

□

Let us elaborate a little on the norm described above. Recall the elementary fact from analysis:

⁸ We let $B_b(x_0)$ denote the closed ball of radius b about x_0 .

⁹ We are a bit sloppy in how we defined J since we assumed J to already exist earlier. This can be quickly fixed by just taking a slightly larger interval to be the input to f in time.

Fact 31. $C^0(J; B_b(x_0))$ is complete under the sup norm $\|\cdot\|_\infty$.

Lemma 32. The norms $\|\cdot\|_\infty$ and $\|\cdot\|_L$ are equivalent. More specifically,

$$\|f\|_\infty e^{-La} \leq \|f\|_L \leq e^{La} \|f\|_\infty.$$

Proof. Should be clear by looking at it since J is closed. □

Putting the two facts together directly proves:

Lemma 33. The function space $C^0(J; B_b(x_0))$ is complete under $\|\cdot\|_L$.

References

- [Mei17] James D. Meiss. *Differential Dynamical Systems*, volume 22 of *Mathematical Modeling and Computation*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, revised edition, 2017.