

Introduction to Control Theory

Mathematical Foundations of Control and Machine Learning

Enrique Zuazua, Daniel López Montero

Winter 2026

FAU Erlangen-Nürnberg

Plan of the lecture

Part I. Dynamical systems, control and the three questions

1. A short history
2. Dynamical systems
3. Acting on a system: the ingredients of a control problem
4. The linear setting and the harmonic oscillator
5. Optimal control, controllability, feedback;

Part II. Controllability of linear systems

6. The reachable set
7. The Kalman rank condition
8. Combining motions: parking
9. The Gramian and an explicit control; the cost of control
10. Control and learning; outlook

Notebooks **01** and **02** of the course complement these slides with proofs, worked examples, code and exercises.


Part I

Dynamical systems, control and the three questions

A short history


Control Theory: a short history

Third century BCE. Ktesibios' water clock keeps its water supply steady with a float valve: the first known feedback device. Antiquity already saw automation as a source of freedom, in a passage of Aristotle's *Politics*.



Ktesibios' water clock.

18th century. In 1788 James Watt fits a *centrifugal governor* to the steam engine. As the engine speeds up, the two rotating balls move outward and close the steam valve; as it slows down, they fall and open it.



Watt's governor and steam valve.

Maxwell, 1868: *On Governors*

Maxwell studied why a governor can fail to settle at a steady speed and instead oscillate [2].

- For small deviations, use a **linear differential equation**.
- Disturbances decay if **every root has negative real part**. Complex roots with positive real part make oscillations grow.

Routh (1877) and Hurwitz (1895) later developed a general stability test.

276

Mr. J. C. Maxwell *on Governors*.

[Mar. 5,

The equation of motion of the machine itself is

$$M \frac{d^2 x}{dt^2} = P - R - F \left(\frac{dx}{dt} - V_1 \right) - Gy. \quad \dots \quad (10)$$

This must be combined with equation (7) to determine the motion of the whole apparatus. The solution is of the form

$$x = A_1 e^{n_1 t} + A_2 e^{n_2 t} + A_3 e^{n_3 t} + V_1 t, \quad \dots \quad (11)$$

where n_1, n_2, n_3 are the roots of the cubic equation

$$MBn^3 + (MY + FB)n^2 + FYN + FG = 0. \quad \dots \quad (12)$$

If n be a pair of roots of this equation of the form $a \pm \sqrt{-1}b$, then the part of x corresponding to these roots will be of the form

$$e^{at} \cos (bt + \beta).$$

If a is a negative quantity, this will indicate an oscillation the amplitude of which continually decreases. If a is zero, the amplitude will remain constant, and if a is positive, the amplitude will continually increase.


One root of the equation (12) is evidently a real negative quantity. The condition that the real part of the other roots should be negative is

$$\left(\frac{F}{M} + \frac{Y}{B} \right) \frac{Y}{B} - \frac{G}{B} = \text{a positive quantity.}$$

On Governors (1868), p. 276.

The original equation, exponential modes and stability condition.

A timeline



Dynamical systems

What is a dynamical system?

A **dynamical system** is a state space X together with a rule for how the state changes over time.

In **discrete time**, one map gives one step:


$$F : X \rightarrow X, \quad x_{k+1} = F(x_k), \quad k = 0, 1, 2, \dots$$

- The **state** records what we need to know now to predict what happens later.
- In **continuous time**, the rule is usually a differential equation, $x' = f(x)$.
- A map F gives a **deterministic** update. A **stochastic** system uses transition probabilities instead.

Discrete time, finite state set

X is finite. Each step follows a **deterministic rule** or **transition probabilities**.

A traffic light




The state is the active light: green, yellow or red.

$$X = \{G, Y, R\}.$$

In a simplified cycle, F advances to the next light:

$$G \rightarrow Y \rightarrow R \rightarrow G.$$

A magnet: Ising spins




N spins, each pointing up or down:

$$X = \{-1, +1\}^N.$$

One step attempts a spin flip; its acceptance probability depends on the energy change.

A Markov chain



The next state is random, with probabilities depending only on the current state.


$$p_{k+1} = P^T p_k.$$

p_k is the probability vector over X ; P is the transition matrix.

Discrete time in $\mathbb{R}^n$

The state is a **vector**; each step applies an update map, $x_{k+1} = f_k(x_k)$.

The logistic map




The state is a population with limited resources:

$$x_{k+1} = r x_k (1 - x_k).$$

The two initial states shown differ by 10^{-4} .

A residual network




The state is a data point's features. Each layer adds a learned increment:

$$x_{k+1} = x_k + h g_k(x_k),$$

where g_k is a neural network and $h > 0$ is a step size.

Gradient descent



The state is the parameter vector; F is the loss:


$$w_{k+1} = w_k - \eta \nabla F(w_k).$$

The step size is $\eta > 0$.

Continuous time in $\mathbb{R}^n$

The state evolves continuously, following a **deterministic** or **stochastic** rule.


Newton's law



Position and velocity form the state $x = (q, v)$:

$$q' = v, \quad m v' = F(q, v).$$

The pendulum




The state is angle and angular velocity,

$$x = (\theta, \theta'):$$

$$\theta'' = -\sin \theta.$$

Brownian motion




A particle's position is random, $X_t \in \mathbb{R}^2$:

$$dX_t = dW_t, \quad X_0 = 0.$$

Infinite dimension: partial differential equations

The state is a **function** of space; its evolution follows a PDE.

Heat




The state is the temperature profile $y(t, \cdot) \in L^2(0, 1)$:

$$y_t = y_{xx},$$

$$y(t, 0) = y(t, 1) = 0.$$

Diffusion smooths the profile.

Waves




Displacement and velocity form the state (y, y_t) :

$$y_{tt} = y_{xx}.$$

A pulse splits into two waves travelling in opposite directions.

Transport



The state is a moving profile $y(t, \cdot)$:

$$y_t + c y_x = 0,$$

$$y(t, x) = y_0(x - ct).$$

The shape moves at speed c .

Dynamical systems: summary

Time, state space	Rule	Examples
discrete, X finite	$x_{k+1} = F(x_k)$	traffic light, board game; a Markov chain on the probabilities
discrete, $X = \mathbb{R}^n$	$x_{k+1} = f(x_k)$	logistic map; Euler's method; a residual network; gradient descent
continuous, $X = \mathbb{R}^n$	$x' = f(x)$ or $dX = dW$	Newton's law; pendulum; Brownian motion
continuous, X infinite dimensional	$y_t = \Delta y, \quad y_{tt} = \Delta y,$ $y_t + c y_x = 0$	heat, waves, transport

Dynamical systems with control

Dynamical systems: adding a control

A control is an **input u that we choose** to influence the evolution.

Without control

$$x'(t) = f(x(t)).$$


With control

$$x'(t) = f(x(t), u(t)).$$

Control systems: three examples

Three of the systems from before, now with an input u that we choose.

The pendulum




A motor at the pivot applies a torque $u(t)$:

$$\theta'' = -\sin \theta + u.$$

A feedback $u(\theta, \theta')$ can hold the pendulum upright, at $\theta = \pi$.

A residual network




Each layer has weights and a bias; σ is a fixed activation:

$$x_{k+1} = x_k + h \sigma(W_k x_k + b_k).$$

Training chooses $u_k = (W_k, b_k)$; the same control acts on every data point.

Heat



A heater acts only on a subregion $\omega \subset (0, 1)$:

$$y_t = y_{xx} + 1_\omega u,$$

$$y(t, 0) = y(t, 1) = 0.$$

Here u is constant and $y(0, \cdot) = 0$: heat enters through ω , then spreads.

The ingredients of a control problem

1. A **state equation**, which describes how the state evolves and where we can act. For most of the course,

$$x'(t) = f(x(t), u(t)), \quad t \in (0, T), \quad x(0) = x^0,$$

with the **state** $x(t) \in \mathbb{R}^n$ and the **control** $u(t) \in \mathbb{R}^m$.

2. A **set of admissible controls** $\mathcal{U}_{\text{ad}}$ the control is taken from. Typical choices: $\mathcal{U}_{\text{ad}} = L^2(0, T; \mathbb{R}^m)$, finite energy; or controls with bounds.
3. An **objective**: a target to reach, a cost to minimize, or an equilibrium to keep.

From nonlinear to linear: linearization

Let $(\bar{x}, \bar{u})$ be an equilibrium, $f(\bar{x}, \bar{u}) = 0$, and write $x = \bar{x} + \xi$, $u = \bar{u} + v$. Taylor's formula gives

$$\xi' = f(\bar{x} + \xi, \bar{u} + v) = \underbrace{D_x f(\bar{x}, \bar{u})}_A \xi + \underbrace{D_u f(\bar{x}, \bar{u})}_B v + O(|\xi|^2 + |v|^2).$$

To first order the system is **linear and time-invariant**: $\xi' = A\xi + Bv$. This is what Maxwell did with the governor, and what every engineer does first.

Linear systems and the harmonic oscillator

The linear time-invariant setting

Let $n, m \geq 1$ and $T > 0$. We study

$$x'(t) = Ax(t) + Bu(t), \quad t \in (0, T), \quad x(0) = x^0,$$

with $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{n \times m}$. The state has n components and the control m .


Variation of constants

$$x(t) = e^{At}x^0 + \int_0^t e^{A(t-s)}Bu(s) ds.$$

Proof. Multiply the equation by e^{-At} :

$(e^{-At}x(t))' = e^{-At}(x'(t) - Ax(t)) = e^{-At}Bu(t)$. Integrate from 0 to t and multiply by e^{At} . □

Example: the harmonic oscillator



A mass m_0 attached to a spring of stiffness k , pushed by a force $u(t)$.
Newton's law:

$$m_0 q'' = -k q + u.$$

With the state $x = (x_1, x_2) = (q, q')$, position and velocity, one second-order equation becomes two first-order ones:

$$x' = \underbrace{\begin{pmatrix} 0 & 1 \\ -k/m_0 & 0 \end{pmatrix}}_A x + \underbrace{\begin{pmatrix} 0 \\ 1/m_0 \end{pmatrix}}_B u.$$

Three questions

Three different questions


Several kinds of control problems fit the same framework, depending on how the objective is formulated.

Controllability

Given x^0 , x^1 and T : is there $u \in \mathcal{U}_{\text{ad}}$ such that

$$x(T) = x^1?$$

A dynamical notion: can the target be *reached*?




Optimal control

Find which $u \in \mathcal{U}_{\text{ad}}$ minimizes a cost, for instance

$$J(u) = |x(T) - x^1|^2 + \int_0^T |u|^2 dt$$

It need not reach the target: it makes the best compromise allowed by the cost.




Stabilization

Is there a *rule* $u = F(x)$ such that every solution of


$$x' = f(x, F(x))$$

tends to an equilibrium?




Open loop and closed loop

open loop — plan once, hope



closed loop — measure and react



The last question introduces a distinction that runs through the whole course. A control can be

1. **open loop**, $u = u(t)$: a function of time, computed in advance from the model and then applied blindly;
2. **closed loop** or **feedback**, $u = F(x(t))$: computed in real time from the current state, so it can react to disturbances the model does not know about.

Why feedback? A thermostat

Temperature deviation x , heating u , setpoint 1, and an *unknown* constant disturbance d (an open window):


$$x' = -x + u + d, \quad x(0) = 0.$$

Open loop. The model says $u \equiv 1$ brings x to 1. With the disturbance, $x(t) \rightarrow 1 + d$: the error is the full disturbance.

Feedback. $u = 1 + L(1 - x)$ with gain $L > 0$. Then $x' = -(1 + L)x + 1 + L + d$ and

$$x(t) \rightarrow 1 + \frac{d}{1 + L} \quad \text{as } t \rightarrow \infty.$$

The error shrinks like $1/(1 + L)$, *without knowing d* , and the response is faster.



Feedback everywhere

The concept of feedback comes from the capacity of biological systems to regulate themselves, and was incorporated into control engineering in the 1920s.

Engineering

- thermostats, cruise control, autopilots
- the governor, and the hunting it may cause
- active noise cancelling: measure the sound, emit the opposite
- robots, rockets, chemical reactors, power grids

Nature and learning

- homeostasis: body temperature, glucose, blood pressure
- a cyclist balancing, a hand reaching for a cup
- reinforcement learning: an action depends on the observed state; the policy is *learned* from rewards instead of *designed* from a model
- gradient descent with momentum

Part I: key takeaways

1. A dynamical system is a state space and an evolution rule:
 $x_{k+1} = F(x_k)$ in discrete time, or a differential equation in continuous time.
2. A control problem consists of a state equation, a set of admissible controls and an objective.
3. Controllability: *can* the target be reached? Optimal control: which admissible control is *best*? Stabilization: is there a *rule* $u = F(x)$ that makes the equilibrium attracting?
4. Open loop: $u = u(t)$, computed in advance. Closed loop: $u = F(x(t))$, computed in real time from the state, so it can correct for unknown disturbances.

Part II

Controllability of linear systems

Controllability of linear systems

Where are we?

When can a linear system be steered between arbitrary states?

Definition (controllability)

The system $x' = Ax + Bu$ is **controllable in time T** if for every $x^0, x^1 \in \mathbb{R}^n$ there is $u \in L^2(0, T; \mathbb{R}^m)$ such that the solution satisfies $x(T) = x^1$.

The control process drives the solution from x^0 to x^1 in time T , by acting on the system through u .

- The answer will be the **Kalman rank condition**, a finite computation on (A, B) , and an **explicit** steering control.
- A weaker notion, **approximate controllability**, only asks that the reachable states be dense: $|x(T) - x^1| \leq \varepsilon$ for every $\varepsilon > 0$. In $\mathbb{R}^n$ the two notions coincide, but for PDEs the two notions differ.

Example 1: an uncontrollable system

With $A = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ and $B = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ the system reads


$$x_1' = x_1 + u, \quad x_2' = x_2 \implies x_2(t) = x_2^0 e^t.$$

The control does not act on the second component, which is completely determined by its initial value, so the system **is not controllable**.

What *can* be reached? From x^0 the reachable states at time T form the horizontal line

$$\{x_2 = x_2^0 e^T\} = e^{AT} x^0 + \text{span}\{e_1\}.$$

This picture turns out to be general.



Example 2: the harmonic oscillator is controllable


With $A = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$ and $B = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ the system reads

$$\begin{cases} x_1' = x_2, \\ x_2' = -x_1 + u \end{cases} \iff x_1'' + x_1 = u.$$

To stop the oscillator, from $x^0 = (1, 0)$ to $x^1 = (0, 0)$, we can use a cubic polynomial

$$u(t) = 1 - \frac{6}{T^2}t + \frac{12}{T^3}t^2 - 3\left(\frac{t}{T}\right)^2 + 2\left(\frac{t}{T}\right)^3,$$

with $\|u\|_{L^2(0,T)} = 1.21$ for $T = 5$.



The reachable set

Definition (reachable set)

The states that can be reached from x^0 in time T form

$$\mathcal{R}(T, x^0) = \{x(T) : x' = Ax + Bu, x(0) = x^0, u \in L^2(0, T; \mathbb{R}^m)\}.$$

Controllability in time T means $\mathcal{R}(T, x^0) = \mathbb{R}^n$ for every x^0 .

By variation of constants, $x(T) = e^{AT}x^0 + L_T u$ with the **input-to-state map**

$$L_T : L^2(0, T; \mathbb{R}^m) \rightarrow \mathbb{R}^n, \quad L_T u = \int_0^T e^{A(T-s)} Bu(s) ds,$$

which is linear, so $\mathcal{R}(T, x^0) = e^{AT}x^0 + \text{Range } L_T$ is an **affine subspace**: the free evolution of x^0 , plus everything the control can add.

Which directions can the control add? Expand $e^{A(T-s)}B = \sum_{k \geq 0} \frac{(T-s)^k}{k!} A^k B$: by **Cayley–Hamilton**, $A^k \in \text{span}\{I, A, \dots, A^{n-1}\}$ for $k \geq n$, so only the columns of $B, AB, \dots, A^{n-1}B$ ever appear.

The Kalman rank condition

Definition (Kalman matrix)

$$\mathcal{K}(A, B) = [B, AB, A^2B, \dots, A^{n-1}B] \in \mathbb{R}^{n \times nm}.$$

Theorem 1 (Kalman, 1960)

The system $x' = Ax + Bu$ is controllable if and only if

$$\text{rank } \mathcal{K}(A, B) = n.$$

- A **finite computation**: $n - 1$ matrix products and one rank. No differential equation needs to be solved.
- **Time plays no role**: controllable in a second if controllable in an hour, at the price of a larger control.

Example 1. $B = e_1$, $AB = e_1$: $\mathcal{K} = \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix}$ has rank 1: not controllable, as found by hand.

Example 2. $B = e_2$, $AB = e_1$: $\mathcal{K} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ has rank 2: controllable, in every time.

Proof of Theorem 1

Necessity. Since $\mathcal{R}(T, x^0) = e^{AT}x^0 + \text{Range } L_T$, controllability in time T means $\text{Range } L_T = \mathbb{R}^n$. The Cayley–Hamilton expansion of $e^{A(T-s)}B$ on the reachable-set slide shows that every $L_T u$ is a combination of columns of $B, AB, \dots, A^{n-1}B$, so $\text{Range } L_T \subseteq \text{Range } \mathcal{K}$. Hence controllability forces $\text{rank } \mathcal{K} = n$.

Sufficiency. Assume $\text{rank } \mathcal{K} = n$. Both ranges are subspaces of $\mathbb{R}^n$, so it suffices to show that every v orthogonal to $\text{Range } L_T$ is orthogonal to $\text{Range } \mathcal{K}$. Let $v \perp \text{Range } L_T$. For every control,

$$0 = \langle v, L_T u \rangle = \int_0^T \langle B^* e^{A^*(T-s)} v, u(s) \rangle ds.$$

Choose the control well. Take $u(s) = B^* e^{A^*(T-s)} v$. Then $\int_0^T |B^* e^{A^*(T-s)} v|^2 ds = 0$. The integrand is continuous, so $g(s) := B^* e^{A^*(T-s)} v$ vanishes for all $s \in [0, T]$.

All derivatives vanish. The function g is smooth and vanishes identically, so all its derivatives vanish. At $s = T$,

$$g^{(j)}(T) = (-1)^j B^* (A^*)^j v = 0 \quad \text{for all } j \geq 0.$$

In other words, v is orthogonal to every column of $A^j B$, hence to $\text{Range } \mathcal{K} = \mathbb{R}^n$. So $v = 0$, and $\text{Range } L_T = \mathbb{R}^n$: the system is controllable. $\square$

Computing the rank

```
def kalman(A, B): # [B | AB | ... | A
                  ^{n-1}B]
    cols = [B]
    for _ in range(A.shape[0] - 1):
        cols.append(A @ cols[-1])
    return torch.cat(cols, dim=1)

def num_rank(M):
    s = torch.linalg.svdvals(M)
    tol = max(M.shape) * eps * s[0]
    return int((s > tol).sum()), s
```

The near-uncontrollable family.

$$A = \begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & -1 \end{pmatrix}, \quad B_\varepsilon = \begin{pmatrix} 0 \\ 1 \\ \varepsilon \end{pmatrix} :$$


For $\varepsilon = 0$ the actuator acts only on the oscillator and $\text{Range } \mathcal{K} = \{x_3 = 0\}$.

Numerical rank of $[B]$, $[B, AB]$, $[B, AB, A^2B]$, and the singular values of $\mathcal{K}$:

ε	$[B]$	$[B, AB]$	$[B, AB, A^2B]$	singular values of $\mathcal{K}$
0	1	2	2	1.414, 1.000, 0.000
0.1	1	2	3	1.414, 1.005, 0.141
0.3	1	2	3	1.414, 1.052, 0.403

For $\varepsilon = 0.1$ the pair is controllable, but it is *nearly* uncontrollable.

Combining motions: parking



A car has two controls, **drive** and **steer**, but a four-dimensional state: position, heading and the angle of the wheels. No control moves it *sideways*, which is what parallel parking requires. Six moves do it:

- | | | | | | |
|---|------------------------------|---|------------------------------|---|--------------|
| 1 | steer right | 2 | reverse (the rear swings in) | 3 | steer left |
| 4 | reverse (the nose swings in) | 5 | straighten the wheels | 6 | pull forward |

New directions are generated by *combining* motions that are available. This is the mechanism behind the Kalman matrix.

Nelson's car and Lie brackets

State (x, y, θ, φ) : position of the rear axle, heading, wheel angle. Controls: speed v and steering rate ω (Nelson's car [3]):

$$\begin{pmatrix} x \\ y \\ \theta \\ \varphi \end{pmatrix}' = v \underbrace{\begin{pmatrix} \cos \theta \\ \sin \theta \\ \frac{1}{L} \tan \varphi \\ 0 \end{pmatrix}}_{g_1: \text{drive}} + \omega \underbrace{\begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}}_{g_2: \text{steer}}.$$

Going along g_2 , then g_1 , then back along g_2 , then back along g_1 does *not* return to the start: the net displacement is the **Lie bracket**

$$[g_2, g_1] = Dg_1 g_2 - Dg_2 g_1 = \left(0, 0, \frac{1}{L \cos^2 \varphi}, 0\right) : \text{wriggle},$$

and $[[g_2, g_1], g_1]$ at $\varphi = 0$ is $\frac{1}{L}(-\sin \theta, \cos \theta, 0, 0)$: *slide*, sideways. Drive, steer, wriggle and slide span $\mathbb{R}^4$: the car is controllable [4, 1].

The linear case

For $x' = Ax + bu$ the drift is $f(x) = Ax$ and the control field is the constant b : $[f, b] = Dbf - Df b = -Ab$. Brackets with the drift produce $Ab, A^2b, \dots$: the Kalman matrix is the linear instance of a rank condition on Lie brackets.

The Kalman rank, block by block

Consider the dynamical system

$$x^{(n)} = u$$

which is the simplest controllable system.


Every block is needed: For the state

$$(x, x', \dots, x^{(n-1)}),$$

$$A = \begin{pmatrix} 0 & 1 & & \\ & 0 & \ddots & \\ & & \ddots & 1 \\ & & & 0 \end{pmatrix}, \quad B = e_n,$$

so

$$B = e_n, \quad AB = e_{n-1}, \quad \dots, \quad A^{n-1}B = e_1.$$



$x''' = u$ with u a unit pulse on $[0, 1]$: the pulse enters x'' at once, x' after one integration and x after two. Each block of the Kalman matrix is one integration.

Most systems are controllable

Controllable pairs are open and dense. Controllability fails if and only if all $n \times n$ minors of $\mathcal{K}(A, B)$ vanish. Each minor is a polynomial in the entries of (A, B) . The set where finitely many polynomials vanish, not all of them identically zero, is closed and has empty interior. So controllability is

- **robust:** it survives small perturbations of A and B ;
- **generic:** most systems are controllable, and an arbitrarily small perturbation makes an uncontrollable pair controllable.

Rank is not a measure of “how controllable”. A pair can be controllable *and* nearly uncontrollable, like the family at $\varepsilon = 0.1$. The singular values tell more than the rank.

Cost estimate. When controllability holds there is a control with

$$\|u\|_{L^2(0,T)} \leq C |e^{AT}x^0 - x^1|,$$

where C depends on (A, B, T) . The size of C is the quantitative side of controllability. The Gramian, below, makes it computable.

The Gramian and the cost of control

The controllability Gramian

The proof of Theorem 1 used the control $u(s) = B^* e^{A^*(T-s)} v$. Applying L_T to controls of this form produces a matrix.

Definition (controllability Gramian)

$$\mathcal{G}_T = \int_0^T e^{A(T-s)} B B^* e^{A^*(T-s)} ds \in \mathbb{R}^{n \times n}.$$

Proposition 2 (properties of the Gramian)

1. $\mathcal{G}_T$ is symmetric and positive semidefinite:
 $v^* \mathcal{G}_T v = \int_0^T |B^* e^{A^*(T-s)} v|^2 ds \geq 0.$
2. $\ker \mathcal{G}_T = (\text{Range } \mathcal{K})^\perp.$
3. $\mathcal{G}_T$ is invertible if and only if $\text{rank } \mathcal{K} = n$, that is, if and only if the system is controllable.
4. $L_T(B^* e^{A^*(T-\cdot)} c) = \mathcal{G}_T c$ for every $c \in \mathbb{R}^n.$

One $n \times n$ matrix records how strongly the control acts in each direction of the state space over $[0, T]$.

Proof of Proposition 2

(1) Symmetry is clear from the formula, since $(e^{A\tau} B B^* e^{A^*\tau})^*$ is the same matrix. The identity follows from $v^* e^{A(T-s)} B B^* e^{A^*(T-s)} v = |B^* e^{A^*(T-s)} v|^2$, integrated over s .

(2) If $\mathcal{G}_T v = 0$ then $v^* \mathcal{G}_T v = 0$, hence $g(s) = B^* e^{A^*(T-s)} v \equiv 0$ by (1), and the derivative argument in the proof of Theorem 1 gives $v \perp \text{Range } \mathcal{K}$.

Conversely, if $v \perp \text{Range } \mathcal{K}$, then v is orthogonal to every column of $e^{A(T-s)} B$ by the Cayley–Hamilton step in the proof of Theorem 1. So $g \equiv 0$, and $\mathcal{G}_T v = \int_0^T e^{A(T-s)} B g(s) ds = 0$.

(3) By (2), $\mathcal{G}_T$ is invertible iff $(\text{Range } \mathcal{K})^\perp = \{0\}$ iff $\text{Range } \mathcal{K} = \mathbb{R}^n$ iff $\text{rank } \mathcal{K} = n$, which is controllability by Theorem 1.

(4) By the definition of L_T applied to $u(s) = B^* e^{A^*(T-s)} c$,

$$L_T u = \int_0^T e^{A(T-s)} B B^* e^{A^*(T-s)} c ds = \mathcal{G}_T c. \quad \square$$

An explicit steering control

If the system is controllable, property (4) tells us how to reach any target. We want $L_T u = x^1 - e^{AT} x^0$; try $u = B^* e^{A^*(T-\cdot)} c$ and solve $\mathcal{G}_T c = x^1 - e^{AT} x^0$:

The Gramian control


$$u(t) = B^* e^{A^*(T-t)} \mathcal{G}_T^{-1} (x^1 - e^{AT} x^0). \quad (\text{G})$$

Verification. By variation of constants and the definition of $\mathcal{G}_T$,

$$\begin{aligned} x(T) &= e^{AT} x^0 + \int_0^T e^{A(T-s)} B B^* e^{A^*(T-s)} ds \mathcal{G}_T^{-1} (x^1 - e^{AT} x^0) \\ &= e^{AT} x^0 + \mathcal{G}_T \mathcal{G}_T^{-1} (x^1 - e^{AT} x^0) = x^1. \end{aligned} \quad \square$$

What we do not have yet. Why this particular form $u = B^* e^{A^*(T-\cdot)} c$? Is it the cheapest control? The next lecture answers both: it is the control of *minimal* L^2 norm, and it comes from a dual, *adjoint* problem.

The control (G) stops the oscillator



For $A = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$, $B = e_2$ and $T = 5$: $\mathcal{K} = [B, AB] = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ has rank 2,
and

$$\mathcal{G}_T = \begin{pmatrix} 2.636 & 0.460 \\ 0.460 & 2.364 \end{pmatrix}, \quad \lambda(\mathcal{G}_T) = \{2.02, 2.98\}.$$

The cost of control

The effort of (G). With $d = x^1 - e^{AT}x^0$ and $c = \mathcal{G}_T^{-1}d$,

$$\int_0^T |u(t)|^2 dt = \int_0^T c^* e^{A(T-t)} B B^* e^{A^*(T-t)} c dt = c^* \mathcal{G}_T c = d^* \mathcal{G}_T^{-1} d.$$

And the constant in the cost estimate is $\lambda_{\min}(\mathcal{G}_T)^{-1/2}$: the smallest eigenvalue of the Gramian measures the hardest direction.

Longer horizons are cheaper. For $T' > T$

$$\mathcal{G}_{T'} - \mathcal{G}_T = \int_T^{T'} e^{A\tau} B B^* e^{A^*\tau} d\tau \succeq 0.$$

Thus $\lambda_{\min}(\mathcal{G}_T)$ is non-decreasing in T , and the same displacement d costs no more effort on a longer horizon.

How hard is it to control? The near-uncontrollable family

The Kalman condition is a yes/no answer, and a “yes” can hide very different behaviours. Take again


$$A = \begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & -1 \end{pmatrix}, \quad B_\varepsilon = \begin{pmatrix} 0 \\ 1 \\ \varepsilon \end{pmatrix}.$$

A direct computation gives $\det \mathcal{K}(A, B_\varepsilon) = -2\varepsilon$: controllable for every $\varepsilon \neq 0$.

Steering $x^0 = 0$ to the weak mode e_3 in time $T = 5$ with (G) costs

$$\int_0^T |u|^2 dt = d^* \mathcal{G}_T^{-1} d \sim \frac{C}{\varepsilon^2} \quad (\varepsilon \rightarrow 0),$$

at $\varepsilon = 10^{-3}$ it is 10^6 times its value at $\varepsilon = 1$.



The energy of the control (G) against the coupling ε , on logarithmic axes.

Outlook

Control $\leftrightarrow$ machine learning

- **Feedback $\leftrightarrow$ reinforcement learning.** Reinforcement learning is “learning based on feedback”: there are no correct input–output pairs, but there is a reward. In both, an action depends on the observed state. A feedback law is *designed* from a known model; a policy is *learned* from rewards.
- **Damping $\leftrightarrow$ momentum.** The heavy-ball picture of momentum in gradient descent is a damped second-order dynamics. Its damping parameter can be too small, giving oscillations, or too large, giving slow creeping, as in the overdamping paradox of Part I.

The overdamping paradox

The model. Without friction, the linearized pendulum $x_1'' + x_1 = 0$ conserves its energy

$$e(t) = \frac{1}{2} (x_1(t)^2 + x_1'(t)^2).$$

With friction proportional to the velocity,

$$x_1'' + x_1 = -k x_1', \quad k > 0.$$


The control is $u = -kx_2$: a feedback law where the control opposes the measured velocity. It is the simplest instance of Maxwell's governor analysis.

Energy decreases. Along solutions,

$$e'(t) = x_1'(x_1 + x_1'') = -k (x_1')^2 \leq 0.$$

So more damping removes energy faster *at each instant*, for the same velocity.

Does it follow that a larger k always brings the system to rest faster?



Beyond linear, finite-dimensional, unconstrained

- **Nonlinear systems.** If the linearization at an equilibrium is controllable, the nonlinear system is controllable *near* it. Far from equilibria, Lie brackets replace the Kalman blocks (parking). Controllability can be global, local, or fail altogether.
- **Bounded controls.** With $|u(t)| \leq 1$ the reachable set is no longer a subspace but a convex set that grows with T ; time-optimal controls switch between the extreme values: *bang-bang*.
- **Stochastic systems.** Controls act on distributions; the natural questions are about expectations and risks, and feedback becomes indispensable.

Observation: we rarely see the whole state

In practice we only observe an output

$$x'(t) = f(x(t), u(t)), \quad y(t) = g(x(t)) \in \mathbb{R}^p, \quad \text{usually } p < n.$$

In the linear case, $x' = Ax + Bu$ and $y = Cx$ with $C \in \mathbb{R}^{p \times n}$.

System	State x	Observation y
oscillator, pendulum	position and velocity	position only: $C = \begin{pmatrix} 1 & 0 \end{pmatrix}$
heat in a domain Ω	temperature field	thermometers in a subregion $\omega \subset \Omega$
a spacecraft	position, velocity, attitude	radar ranges, star sightings
reinforcement learning	the state of the environment	a camera image of it

Knowing u and y on $(0, T)$, can we recover $x(t)$? Can we still stabilize?

Part II: key takeaways

1. Kalman: controllable $\iff \text{rank}[B, AB, \dots, A^{n-1}B] = n$.
2. New directions come from combining motions: AB for linear systems, Lie brackets for the car.
3. The Gramian measures the effort needed to steer x^0 to x^1 .
4. The rank is algebraic; singular values and Gramian eigenvalues measure how hard control is, direction by direction.

- [1] Jean-Michel Coron. *Control and Nonlinearity*, volume 136 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2007. doi: 10.1090/surv/136.
- [2] James Clerk Maxwell. On governors. *Proceedings of the Royal Society of London*, 16: 270–283, 1868. doi: 10.1098/rsp1.1867.0055.
- [3] Edward Nelson. *Tensor Analysis*. Princeton University Press, Princeton, NJ, 1967.
- [4] Eduardo D. Sontag. *Mathematical Control Theory: Deterministic Finite Dimensional Systems*. Springer, New York, 2 edition, 1998. doi: 10.1007/978-1-4612-0577-7.