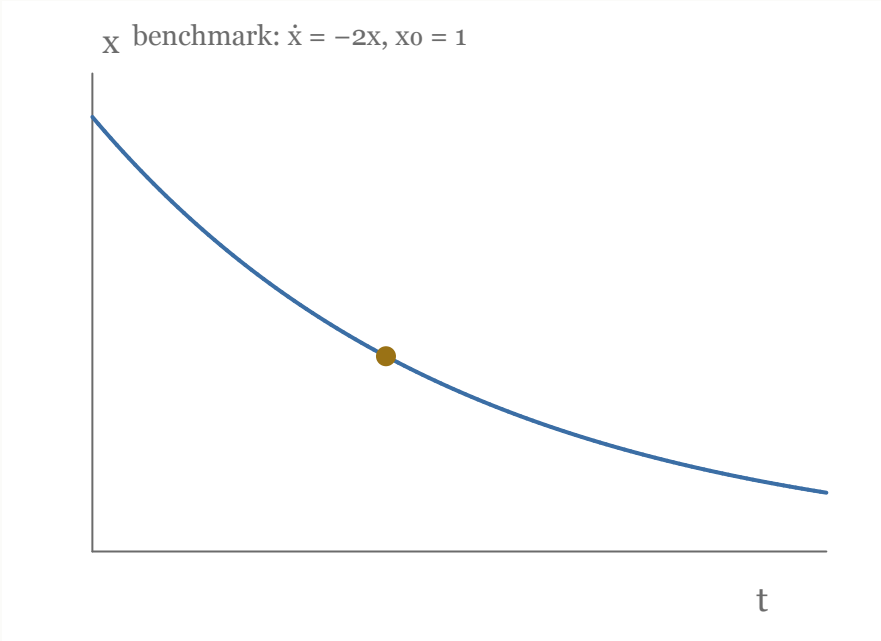
ODE and Simulation

How does a simulator advance the world?

$$\dot{x} = f(x, t), \qquad x(t_0) = x_0$$

- `x` state column
- `f` rate of change
- `t` time; to is the start
- x_0 known initial state

Ordinary differential equation (ODE): a state's rate of change with time.



OBJECTS

ODE and Simulation

A trajectory and its stored samples

$$t_k = t_0 + kh, \qquad x_k \approx x(t_k)$$

`h, k` step size and sample index

`x(t), xₖ` exact state and numerical sample

Initial state + rate model

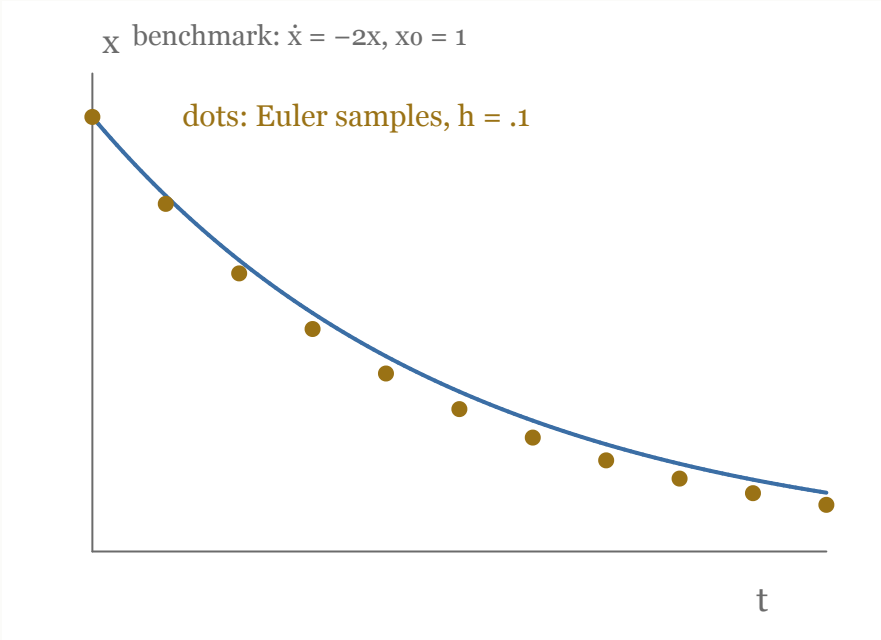
↓ exact solution / numerical integration

Trajectory at requested times

Accuracy: error · Stability: repeated steps

Solver verification: exact / refined solution

Model validation: measurements



Exact solution; numerical solution — [4] FNC §6.2 · Course map

DEFINITION

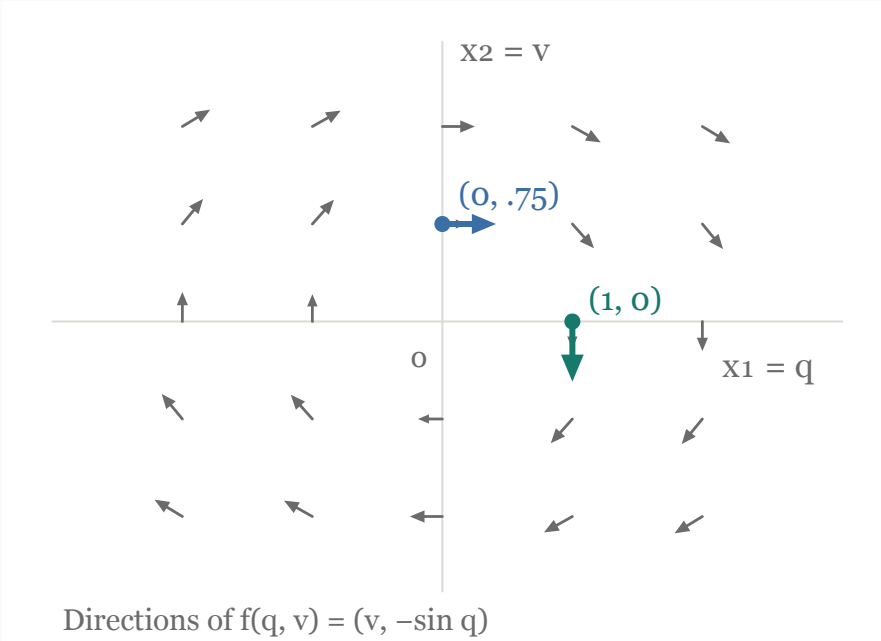
Vector Fields

An arrow gives a rate, not a destination

$$\dot{x} = f(x, t) \in \mathbb{R}^d$$

- `x` d-component state column
- `f(x,t)` d-component state rate
- `t` time held fixed in a field picture

Illustration: $x=(q,v)^T$ and $f=(v,-\sin q)^T$. Arrow lengths are scaled.



Directions of f(q, v) = (v, −sin q)

Vector field; phase portrait — [1] Tedrake, ch.2 · Course illustration

ASSUMPTIONS

Vector Fields

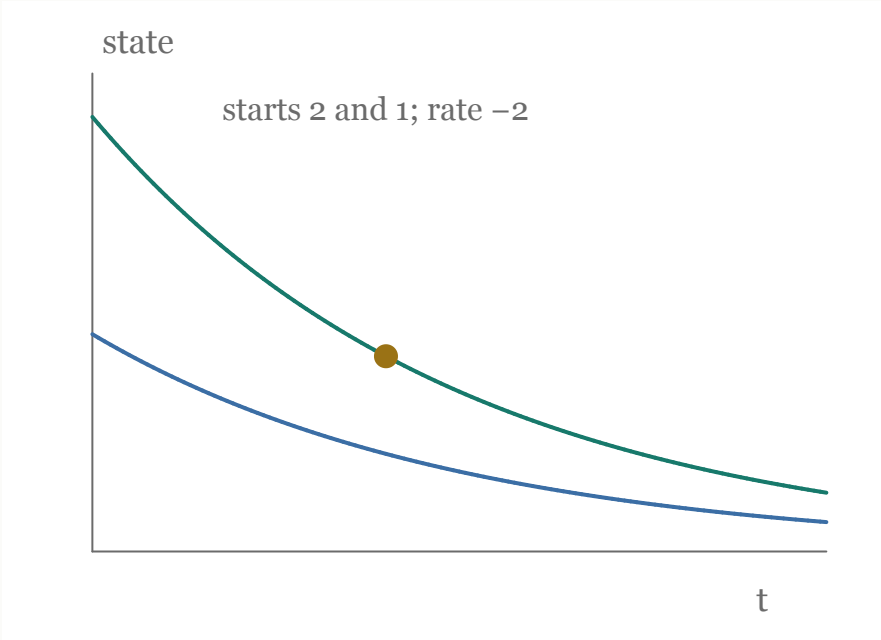
One field, different initial states

$$\dot{x} = -2x, \quad x(0) = 1 \text{ or } 2$$

`x(0)` two possible starts

`-2x` the same rate model for both

Initial-value problem (IVP): a differential equation plus a starting state.



Initial condition; solution trajectory — [4] FNC §6.1 · Course example

WORKED EXAMPLE

Vector Fields

Verify an exact decay solution

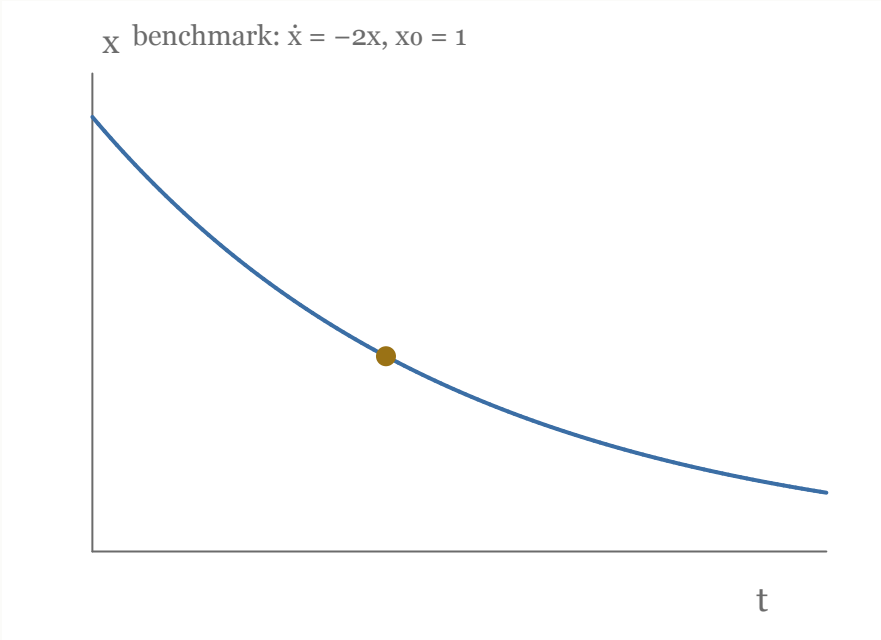
$$x(t) = e^{-2t} \quad \Longrightarrow \quad \dot{x} = -2x$$

`x(0)` one

`-2` decay rate in inverse time

`t` elapsed time

> Worked example: x(0.1) = 0.8187307531.



Exponential decay — [4] FNC §6.1 · Course example

YOUR TURN

Vector Fields

Your turn: a changed decay rate

$$\dot{x} = -3x, \qquad x(0) = 2$$

- `x` scalar state
- `-3` fixed decay rate
- `2` initial value

Which trajectory satisfies BOTH the rate and the start?

A. exp(−3t)

B. 2 exp(−3t)

C. 2 exp(3t)

Exponential decay — [4] FNC §6.1 · Course example

WHY

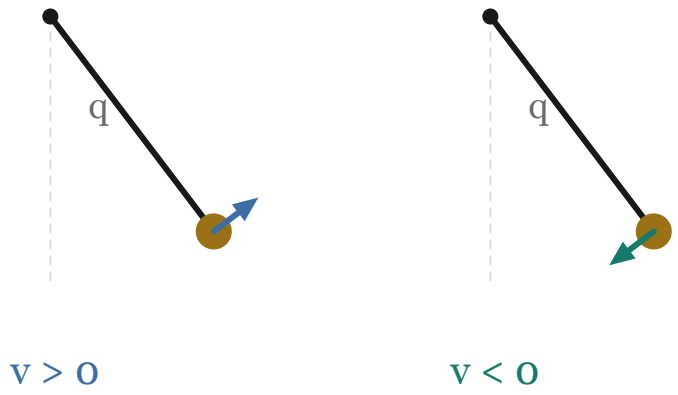
Robot State

Position alone does not predict motion

$$x = \begin{bmatrix} q \\ v \end{bmatrix}, \qquad v = \dot{q}$$

- q — n joint positions
- v — n joint velocities
- x — 2n-component state

The same angle can occur while a joint moves in either direction.



same angle; opposite velocities

First-order system; state vector — [4] FNC §6.3

MODEL

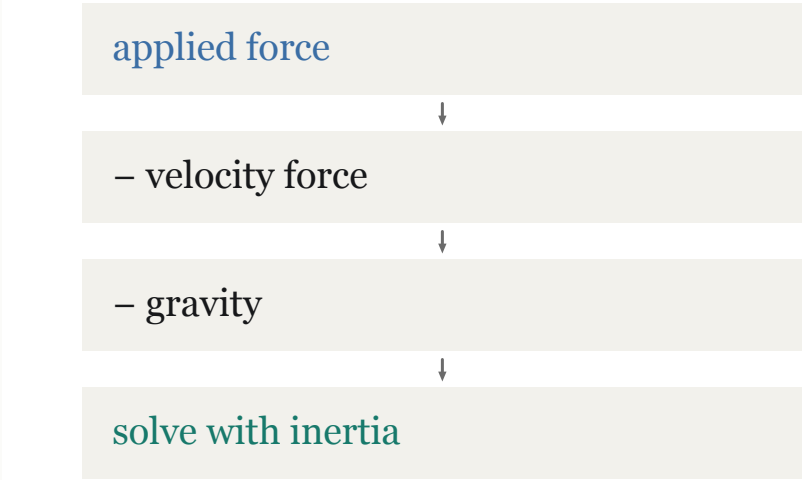
Robot State

The manipulator equation is a force balance

$$M(q)\ddot{q} + C(q,v)v + g(q) = \tau$$

- `M` n×n inertia; $z^T Mz > 0$ for every $z \neq 0$
- `Cv` velocity-dependent generalized forces
- `g, τ` gravity and applied generalized forces
- `q,v` positions and velocities

The model balances inertial, velocity-dependent and gravitational terms.



Manipulator dynamics; mass matrix — [2] Modern Robotics §8.1

DERIVATION

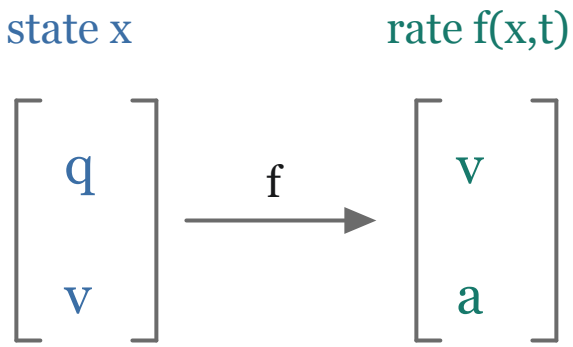
Robot State

Rewrite the mechanics as a first-order ODE

$$\dot{x} = \begin{bmatrix} v \\ M(q)^{-1}[\tau - C(q,v)v - g(q)] \end{bmatrix}$$

- x [q;v], a 2n-column
- τ applied generalized force
- M^{-1} exists for positive-definite inertia

First order: only ẋ appears. Solve M a = τ − Cv − g for acceleration.



position rate = velocity

velocity rate = acceleration

First-order system; forward dynamics — [2] Modern Robotics §8.1 · [4] FNC §6.3

WORKED EXAMPLE

Robot State

A one-joint calculation

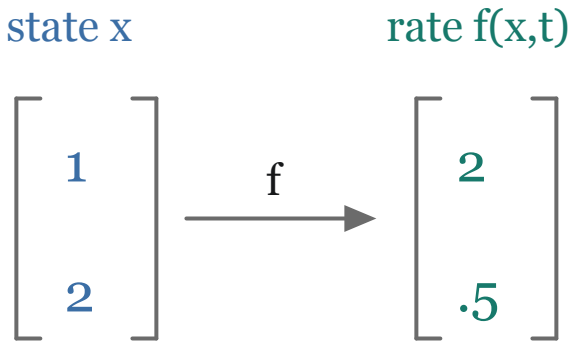
$$a = \frac{5 - 0.5(2) - 3}{2} = 0.5, \qquad \dot{x} = \begin{bmatrix} 2 \\ 0.5 \end{bmatrix}$$

`M,C` 2 and 0.5 at this state

`g,τ` 3 and 5

`q,v` 1 and 2

> Worked example: solve acceleration, then assemble the full derivative.



position rate = velocity

velocity rate = acceleration

Forward dynamics — [2] Modern Robotics §8.1 · Course example

YOUR TURN

Robot State

Your turn: reduce the applied force

$$M = 2,\ C = .5,\ g = 3,\ \tau = 1,\ v = 2$$
$$Ma = \tau - Cv - g, \quad \dot{x} = (v, a)^T$$

`a`	acceleration
`x=(q,v)ᵀ`	position then velocity

Which state rate follows from the force balance?

A. $(-1.5, 2)^T$

B. $(2, -3)^T$

C. $(2, -1.5)^T$

Forward dynamics — [2] Modern Robotics §8.1 · Course example

MODEL

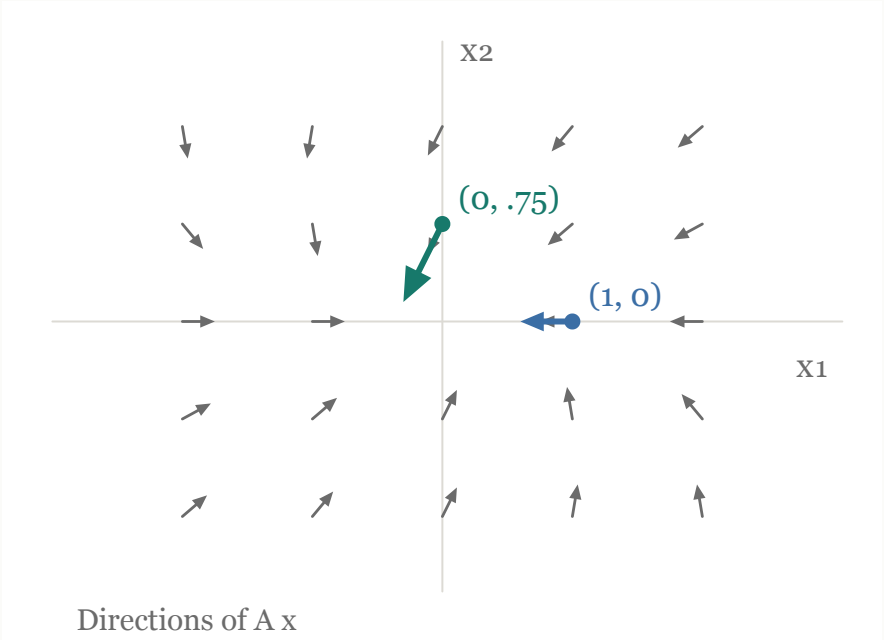
Linear Dynamics

A constant matrix gives a linear field

$$\dot{x} = Ax, \qquad A \in \mathbb{R}^{d \times d}$$

- `A` constant matrix
- `x` d-component state
- `Ax` state rate at x

Illustrated A has rows (−1,−1) and (0,−2); the field is homogeneous.



Directions of A x

Linear constant-coefficient system — [3] Strang, lecture 23 · Course example

RECALL

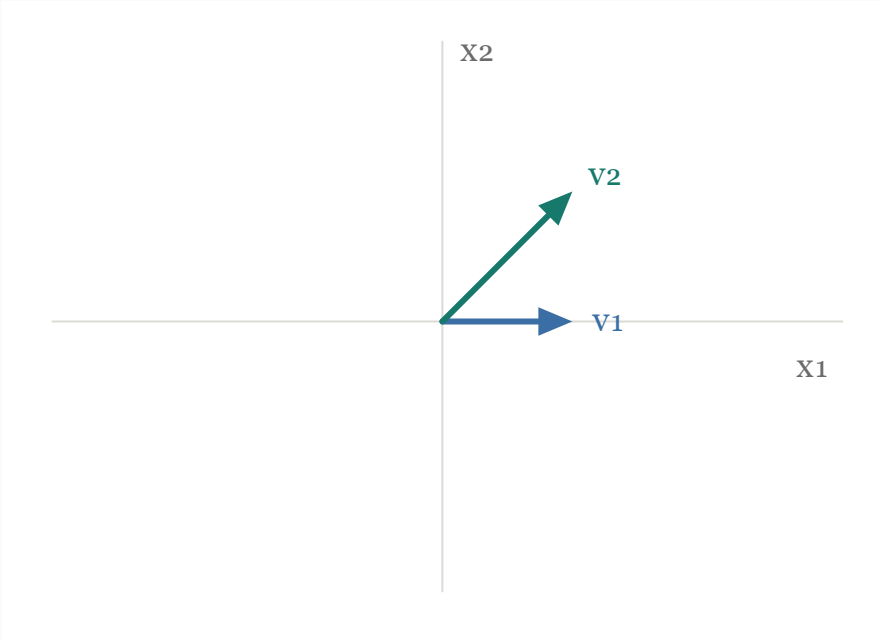
Linear Dynamics

Recall the eigenvector coordinates

$$AV = V\Lambda, \qquad x = Vc$$

- V d independent eigenvectors as columns
- Λ eigenvalues in matching order
- c coordinates of x in that basis

Eigenbasis: d independent eigenvectors spanning the state space. Here $v_1 = (1,0)^T$, $v_2 = (1,1)^T$.



Diagonalization; modal coordinates — [3] Strang, lecture 23

DERIVATION

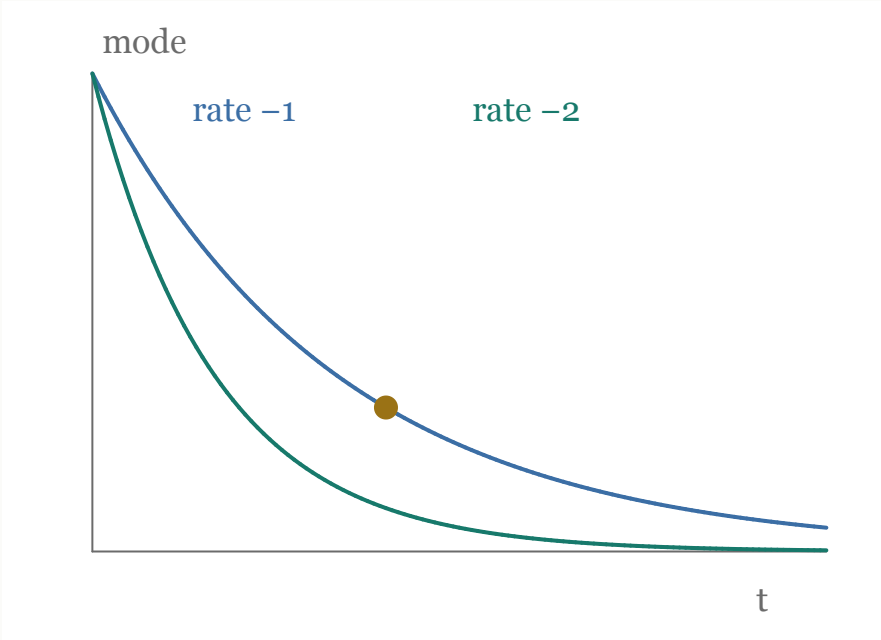
Linear Dynamics

The ODE separates into scalar modes

$$V\dot{c} = AVc = V\Lambda c \quad \Rightarrow \quad \dot{c} = \Lambda c$$

- V constant invertible basis
- c_i coefficient of eigenvector i; one mode
- Λ diagonal rate matrix

Multiply on the left by V^{-1}; each coordinate now evolves alone.



Decoupled linear ODEs — [3] Strang, lecture 23 · Course derivation

RESULT

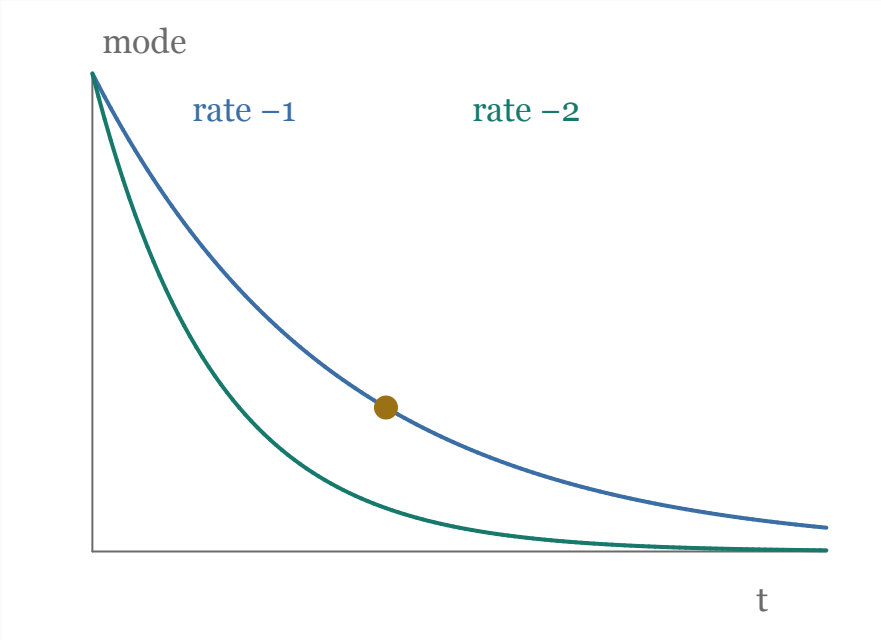
Linear Dynamics

Each eigenvalue is a continuous growth rate

$$c_i(t) = e^{\lambda_i t} c_i(0)$$

- λ_i rate of mode i
- $c_i(0)$ initial amount of that mode
- t time measured from zero

Negative rates decay; positive rates grow if the mode is present.



Eigenvalues; exponential growth and decay — [3] Strang, lecture 23 · Course example

YOUR TURN

Linear Dynamics

Return from modal coordinates

$$V = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \quad \Lambda = \mathrm{diag}(-1, -2)$$

- V eigenvector basis
- Λ matching eigenvalues
- c_0 coordinates of x_0 in the V basis

For NEW $x_0=(1,1)^T$, which modal description is correct?

A. $c_0=(0,1)^T$; only the -2 mode

B. $c_0=(1,2)^T$; both modes

C. $c_0=(1,1)^T$; both modes

Change of basis; matrix exponential — [3] Strang, lecture 23 · Course derivation

DEFINITION

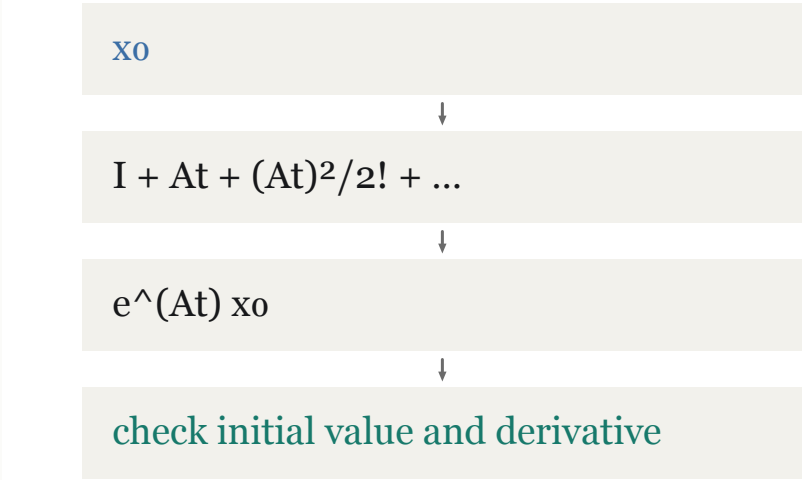
Linear Dynamics

The matrix exponential always exists

$$e^{At} = \sum_{j=0}^{\infty} \frac{(At)^j}{j!}, \qquad x(t) = e^{At} x_0$$

- A any square constant matrix
- j series index; $A^0 = I$
- x_0 initial state

Diagonalization is one way to evaluate the solution, not an existence requirement.



Matrix exponential — [3] Strang, lecture 23

YOUR TURN

Linear Dynamics

A defective system

$$A = \begin{bmatrix} -1 & 1 \\ 0 & -1 \end{bmatrix}, \quad x_0 = (0,1)^T, \quad x(t) = (te^{-t}, e^{-t})^T$$

`defective` fewer than n independent eigenvectors in $\mathbb{C}^n$

`v` eigenvector; here v lies along $(1,0)^T$

Does x(t) satisfy $\dot{x}=Ax$ and $x(0)=x_0$?

A. No: a defective matrix has no solution.

B. No: the factor t forces growth forever.

C. Yes: its derivative and initial state both match.

Defective matrix; matrix exponential — [3] Strang, lecture 23 · [7] NIST DLMF §1.2 · Course example

WORKED EXAMPLE

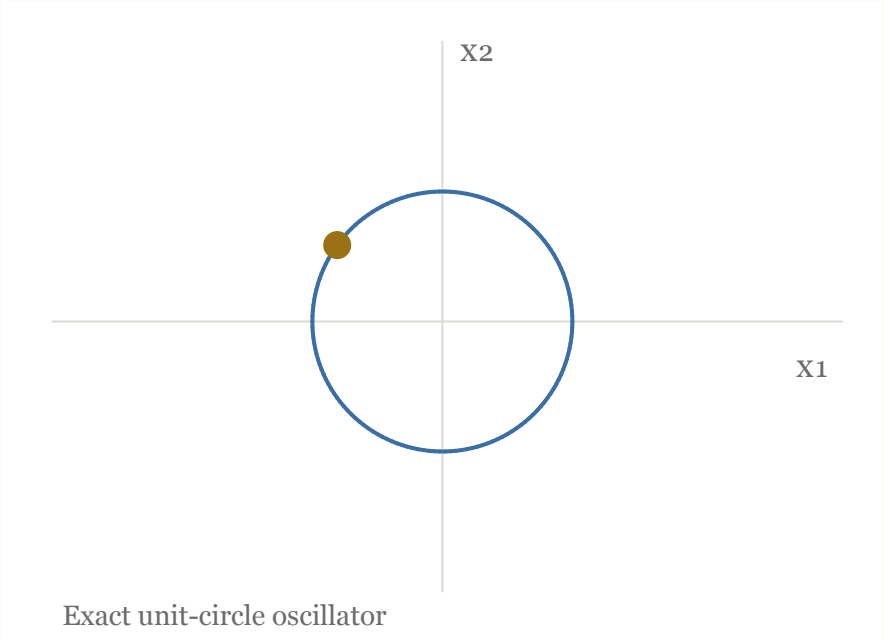
Linear Dynamics

A real rotation has complex eigenvalues

$$A = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}, \quad x(t) = (\cos t, \sin t)^T$$

- x_0 $(1,0)^T$
- A constant rotation generator
- λ $\pm i$ over the complex field

> This trajectory stays on the unit circle.



Exact unit-circle oscillator

Complex eigenvalues; rotation — [3] Strang, lecture 23 · Course example

YOUR TURN

Linear Dynamics

Zero real parts need a closer look

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \qquad x_0 = (0, 1)^T$$

- `A` both eigenvalues zero
- `x0` chosen initial state
- `x(t)` state to determine

What does this zero-eigenvalue example do?

A. $x(t)=(0,1)^T$: fixed

B. $x(t)=(t,1)^T$: unbounded

C. $x(t)=(\exp(t),1)^T$: exponential growth

Boundedness; asymptotic decay — [3] Strang, lecture 23

DERIVATION

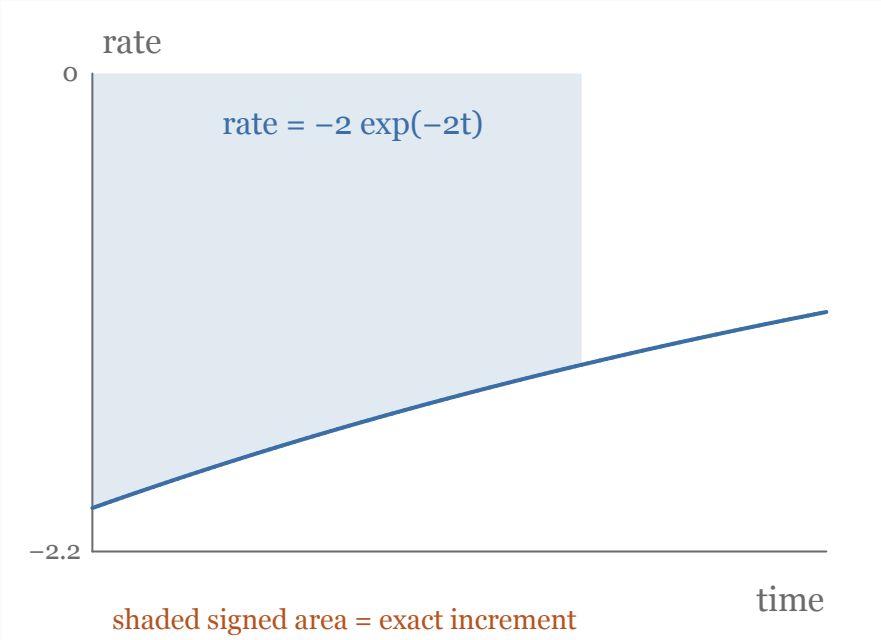
Euler's Method

An exact step is an integral

$$x(t+h) - x(t) = \int_t^{t+h} f(x(s), s)\, ds$$

- `h` positive interval length
- `s` integration variable
- `f(x(s),s)` rate along the unknown path

The fundamental theorem of calculus gives the starting relation.



Integral form of an ODE — [4] FNC §6.2 · Course derivation

RULE

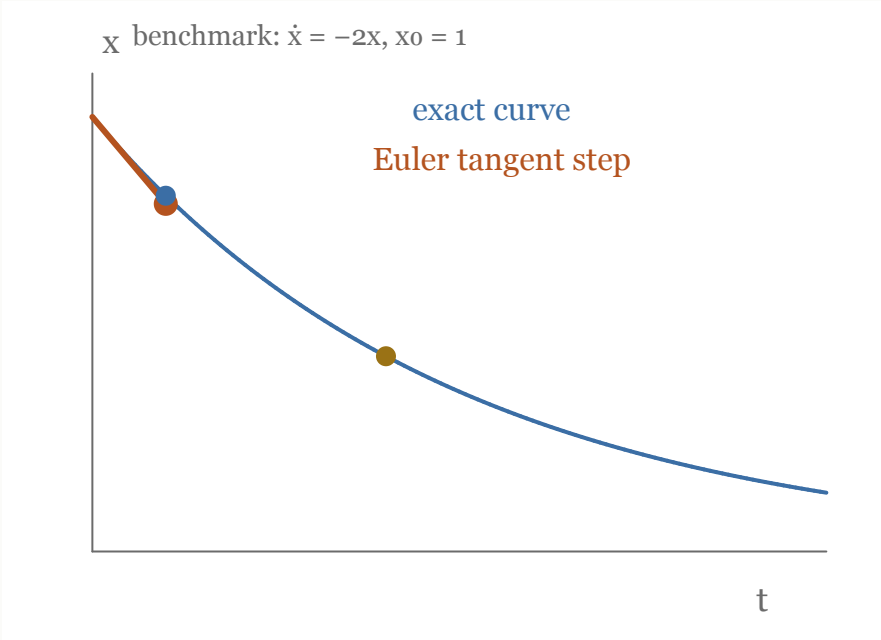
Euler's Method

Euler uses the rate at the start

$$x_{k+1} = x_k + hf(x_k, t_k)$$

- x_k current numerical state
- h positive time step
- $f(x_k, t_k)$ rate at the old state

Every component uses the same old state.



Explicit Euler method — [4] FNC §6.2

WORKED EXAMPLE

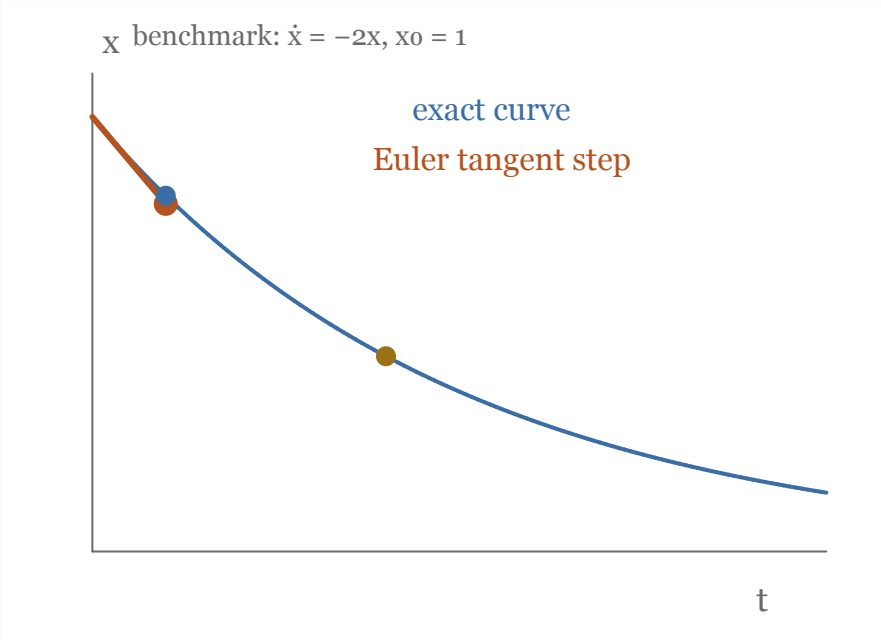
Euler's Method

One step against an exact answer

$$x_1 = 1 + 0.1(-2) = 0.8$$

x_0	one
h	0.1
$f(x)$	$-2x$

> Worked example: exact x(0.1)=0.8187307531; Euler is lower.



Explicit Euler method; one-step error — [4] FNC §6.2 · Course example

YOUR TURN

Euler's Method

Your turn: two larger steps

$$\dot{x} = -2x, \quad x_0 = 1, \quad h = 0.2$$

- h fixed time step
- x_0 initial numerical state
- x_2 approximation at t=0.4

Which Euler pair uses a fresh rate at each step?

A. x1=.6, x2=.36

B. x1=.6, x2=.2

C. x1=.8, x2=.64

Explicit Euler method — [4] FNC §6.2 · Course example

Euler's Method

Local error and accumulated error differ

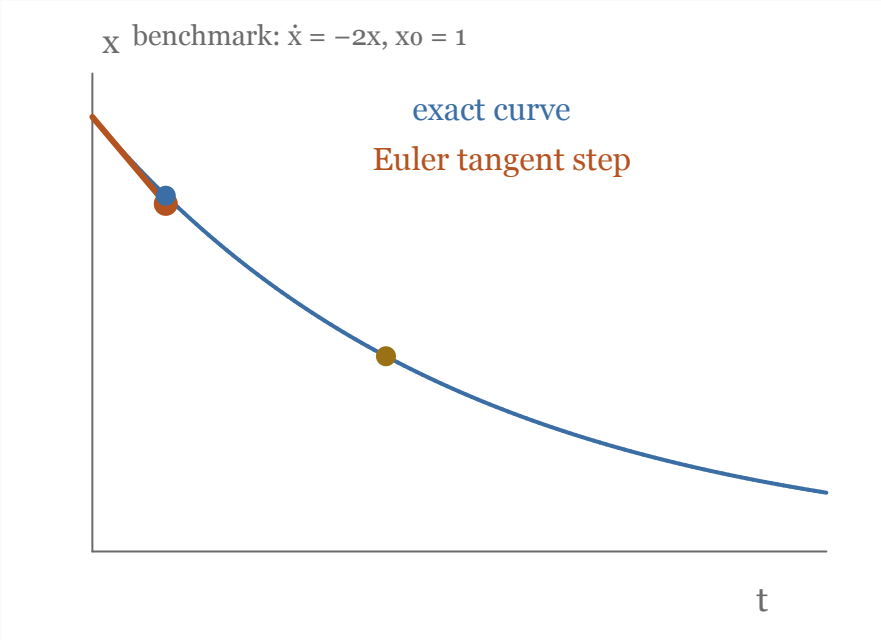
$$x(t+h) = x + h\dot{x} + \tfrac{1}{2}h^2\ddot{x} + O(h^3)$$

$$\ddot{x} = f_t + J_f f$$

`x, f` evaluated on the exact path at time t

f_t, `Jf` time derivative; $(Jf)_{ij} = \partial f_i / \partial x_j$

C^2: continuous derivatives through order 2. $O(h^p)$: size $\le C h^p$ for small h. Local state error = h × normalized local truncation error.



Local error; global error — [4] FNC §6.2

RECALL AND DERIVATION

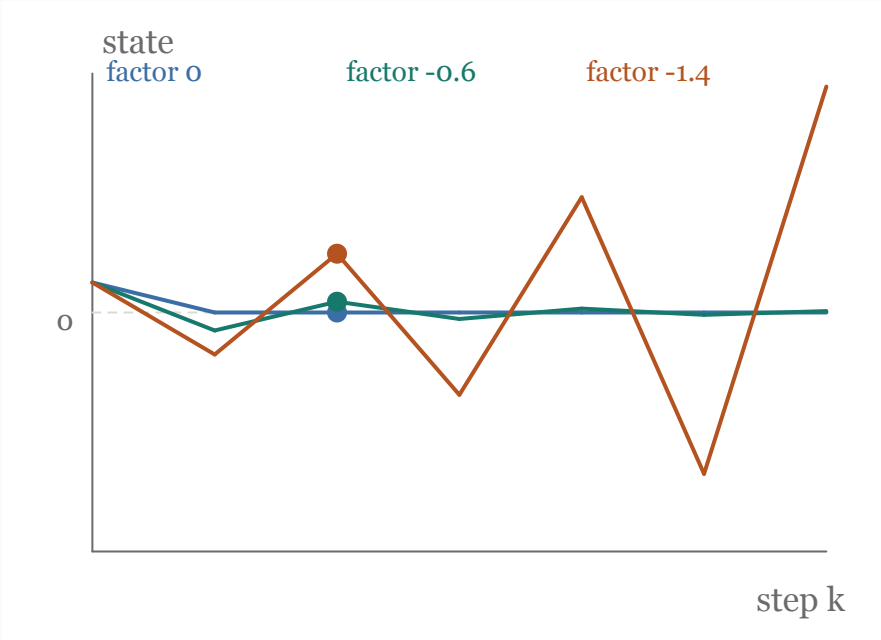
Numerical Stability

A repeated step has a multiplier

$$\dot{x} = \lambda x \quad \Rightarrow \quad x_{k+1} = (1 + h\lambda) x_k$$

- λ constant scalar rate
- h step size
- $1+h\lambda$ numerical amplification factor
- H Hessian: second derivatives of the fixed quadratic

Recall Lecture 6: A=−H turns this into the gradient-descent multiplier.



Amplification factor — [4] FNC §11.3

DERIVATION

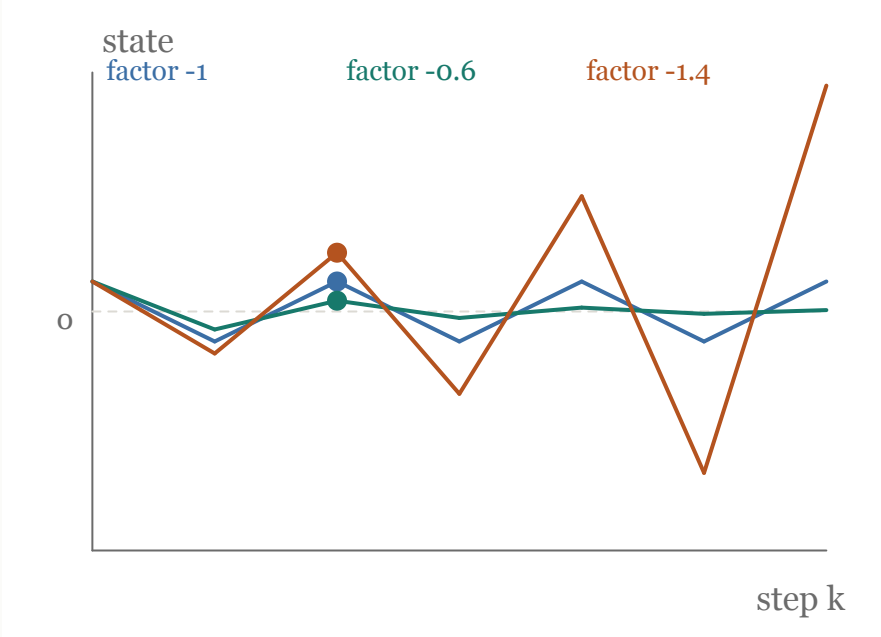
Numerical Stability

The strict inequality decides decay

$$|1 + h\lambda| < 1 \quad \Longleftrightarrow \quad 0 < h < -2/\lambda$$

- λ real negative rate
- h positive step size
- $|1+h\lambda|$ size of each numerical multiplier

FNC: bounded if $|1+h\lambda| \le 1$. Here: decay needs $|1+h\lambda| < 1$.



Absolute stability; asymptotic decay — [4] FNC §11.3

YOUR TURN

Numerical Stability

Predict three step-size regimes

$$\dot{x} = -4x, \quad h \in \{.25, .4, .6\}$$

- λ fixed decay rate
- h candidate steps
- x_0 any nonzero start

For h=.25, .4, .6, which Euler behaviors occur in order?

A. Zero; alternating growth; alternating growth

B. All decay because the ODE is stable

C. Zero; alternating decay; alternating growth

Step-size restriction; asymptotic decay — [4] FNC §11.3 · Course example

YOUR TURN

Numerical Stability

Every mode must satisfy the bound

$$\dot{x} = Ax, \quad A = \mathrm{diag}(-2, -8)$$

- `A` constant state matrix
- `h` one step used for both modes
- x_0 arbitrary initial state

For Euler, which interval gives decay from EVERY start?

A. $0<h\leq 1/4$

B. $0<h<1/4$

C. $0<h<1$

Stability of a linear system — [4] FNC §11.3 · Course example

DIAGNOSTIC

Numerical Stability

Euler adds energy to an oscillator

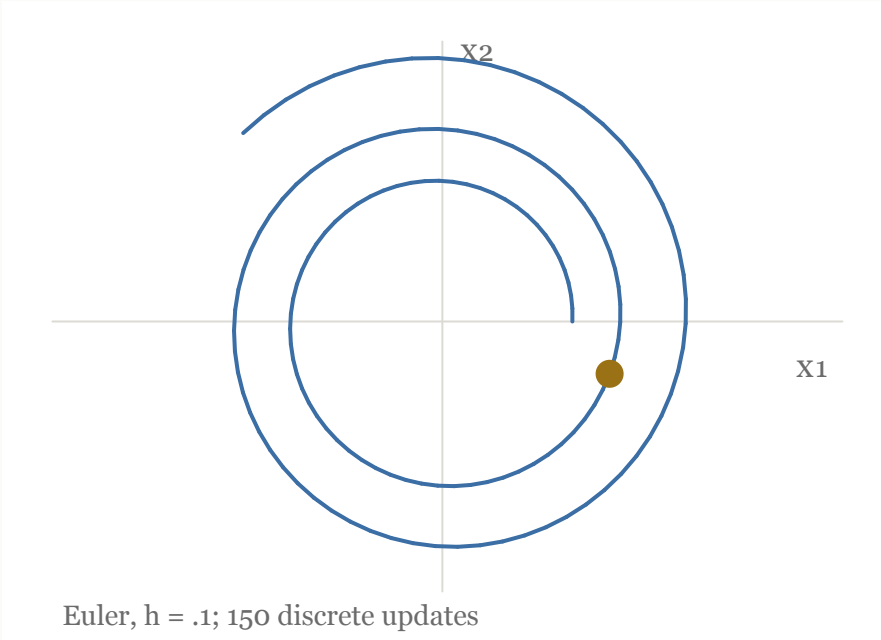
$$\dot{x} = Ax, \quad A = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$$

$$|1 + ih| = \sqrt{1 + h^2} > 1 \quad (h > 0)$$

`±i` eigenvalues of A

`h` positive Euler step

A bounded exact orbit can produce an outward numerical spiral.



Euler, h = .1; 150 discrete updates

Absolute stability; harmonic oscillator — [4] FNC §11.3

RECAP

Numerical Stability

Model validation and solver verification

$$\text{model: } e^{\lambda h} \qquad \text{Euler: } 1 + h\lambda$$

`λ`	continuous rate
`h`	step size
`two factors`	exact and numerical evolution over the same interval

Invariant: constant along an exact trajectory. Solver verification checks equations; model validation checks measurements.

Same model, two comparisons

Numerical solution

↓ compare with

Exact / refined solution

tests the solver and step size

Model prediction ↔ measurements

tests the physical model

Verification; validation — [5] NASA, V&V glossary

WHY

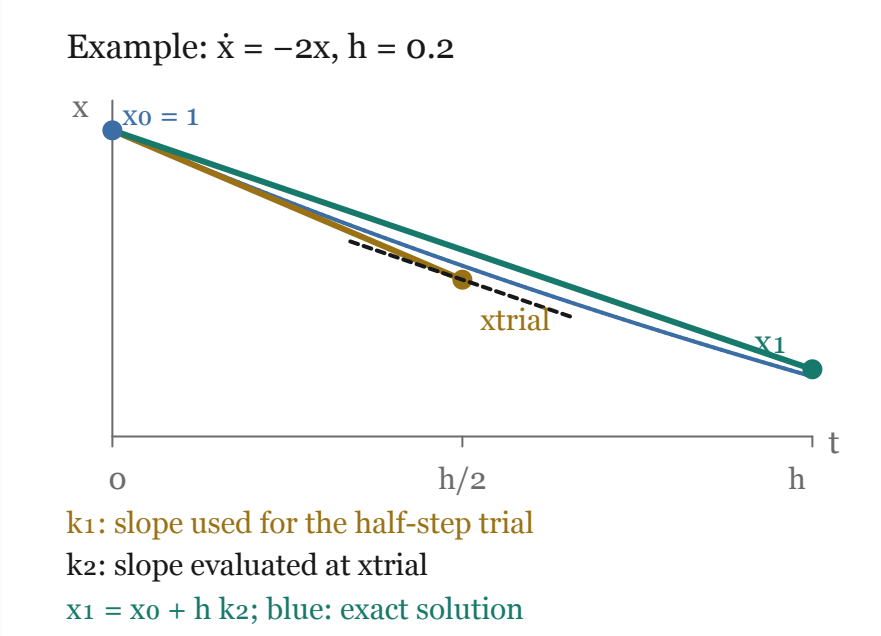
Runge–Kutta Methods

Sample the field inside the step

$$\Delta x = \int_{t_k}^{t_k+h} f(x(s), s)\, ds$$

- Δx exact increment from the current exact state
- f state-rate model along the exact trajectory
- h one integration interval

A better slope estimate can improve the same time step.



Runge–Kutta methods — [4] FNC §6.4 · Course diagram

RULE

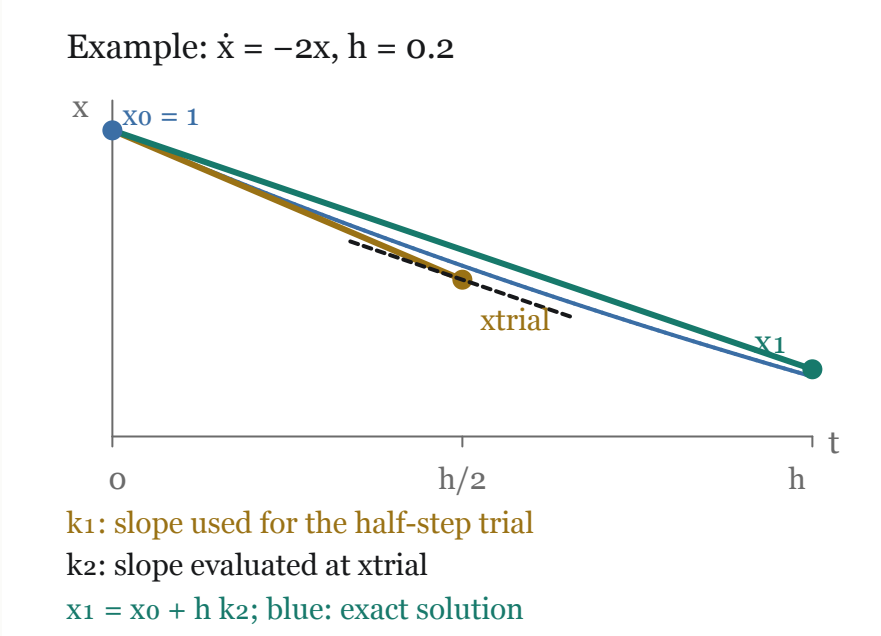
Runge–Kutta Methods

Construct a trial midpoint

$$k_1 = f(x_k, t_k), \quad x_{\text{trial}} = x_k + \frac{h}{2} k_1$$
$$k_2 = f(x_{\text{trial}}, t_k + \frac{h}{2})$$

k_1, k_2	initial and trial-midpoint rates
xtrial	temporary state at $t_k + h/2$
h k_1/2	estimated change over half a step

Explicit midpoint: second-order Runge–Kutta (RK2). FNC also calls it improved Euler (IE2).



Explicit midpoint method (RK2) — [4] FNC §6.4 · Course diagram

WORKED EXAMPLE

Runge–Kutta Methods

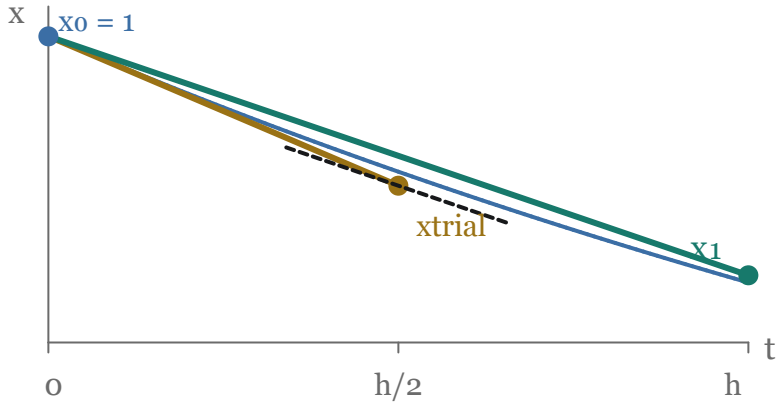
Advance with the midpoint rate

$$x_{k+1} = x_k + hk_2$$

- x_k original current state
- k_2 rate evaluated at the temporary midpoint state
- h full step length

> $\dot{x}=-2x$, $x_0=1$, $h=.2$: $k_1=-2 \rightarrow$ trial$=.8 \rightarrow k_2=-1.6 \rightarrow x_1=.68$.

Example: $\dot{x} = -2x$, $h = 0.2$



k1: slope used for the half-step trial

k2: slope evaluated at xtrial

x1 = x0 + h k2; blue: exact solution

Explicit midpoint method (RK2) — [4] FNC §6.4 · Course diagram

YOUR TURN

Runge–Kutta Methods

Your turn: halve the midpoint step

$$\dot{x} = -2x, \quad x(0) = x_0 = 1, \quad h = 0.1$$

`xtrial` $x_0+(h/2)k_1$; temporary midpoint state

`k1,k2` $f(x_0)$, $f(\text{xtrial})$: rates

`h k2` estimated change over the whole step

Explicit midpoint: which x_1 estimates $x(0+h)$?

A. .9: xtrial

B. .72: xtrial + h k2

C. .82: x0 + h k2

Explicit midpoint method (RK2) — [4] FNC §6.4 · Course diagram

RULE

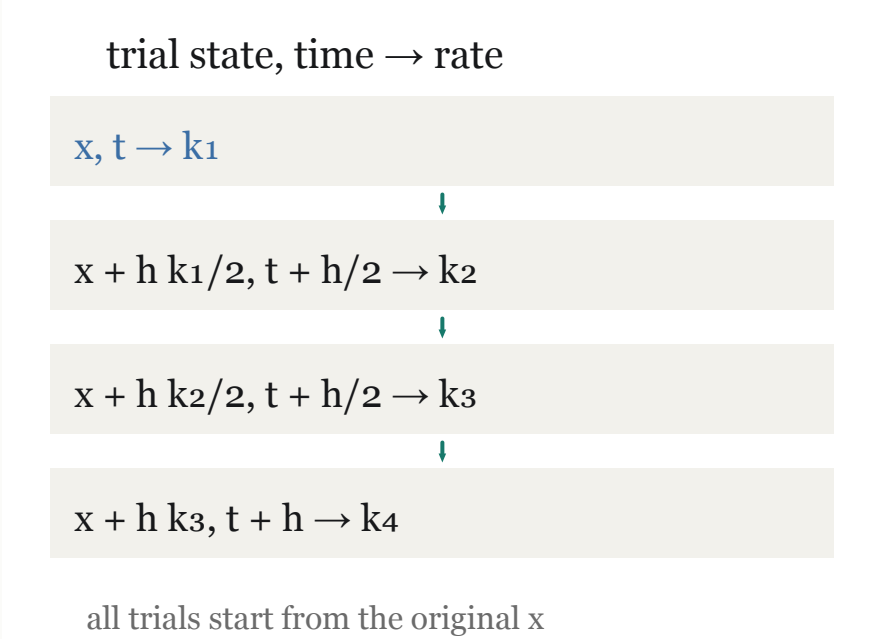
Runge–Kutta Methods

Classical RK4 evaluates four rates

$$
\begin{aligned}
k_1 &= f(x, t) \\
k_2 &= f(x + hk_1/2, t + h/2) \\
k_3 &= f(x + hk_2/2, t + h/2) \\
k_4 &= f(x + hk_3, t + h)
\end{aligned}
$$

`x,t`	current state and time
`h`	one full step
`k₁…k₄`	four rates with state/time units

RK4: classical fourth-order Runge–Kutta. Each trial starts from the same x.



Classical fourth-order Runge–Kutta (RK4) — [4] FNC §6.4 · Course diagram

WORKED EXAMPLE

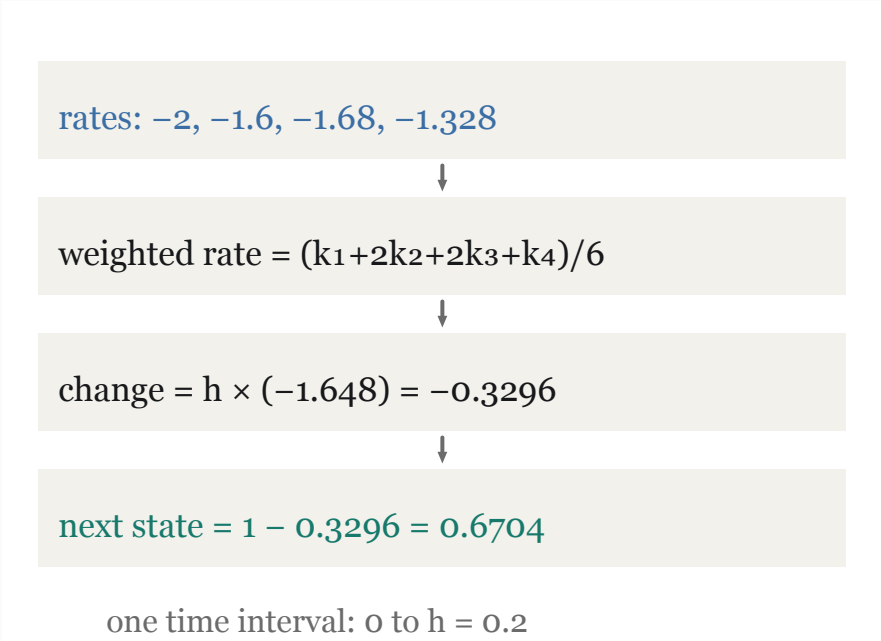
Runge–Kutta Methods

Combine the four rates

$$x_{k+1} = x_k + \frac{h}{6}(k_1 + 2k_2 + 2k_3 + k_4)$$

$k_1 \ldots k_4$	RK4 rates
1,2,2,1	weights before division by six
h	full step length

> $\dot{x}=-2x$, $x_0=1$, $h=.2$: rates $(-2,-1.6,-1.68,-1.328)$ give $x_1=.6704$.



Classical RK4; weighted stage sum — [4] FNC §6.4 · Course diagram

YOUR TURN

Runge–Kutta Methods

Your turn: compare equal model-call budgets

$$\dot{x} = -2x, \quad x_0 = 1, \quad h = 0.1$$

`h` one RK4 step interval

`cost` number of calls to f

How many Euler substeps match ONE RK4 step in model calls?

A. Four, each h/4, over the same interval

B. One, with step h

C. Four, each h, over the same interval

Function evaluations; order of accuracy — [4] FNC §6.4

COMPARISON

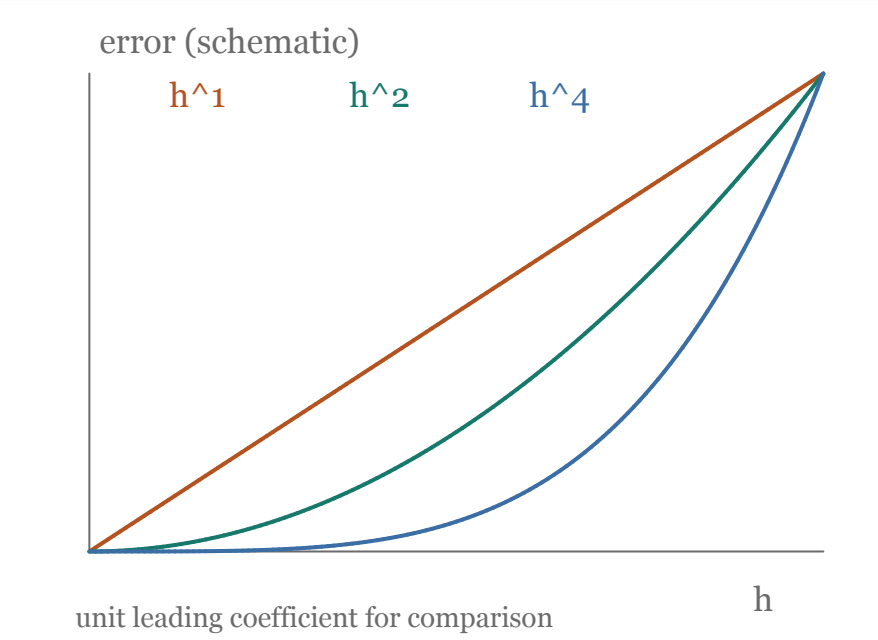
Runge–Kutta Methods

Accuracy improves at a computational price

$$\varepsilon_{\text{Euler}} = O(h), \quad \varepsilon_{\text{midpoint}} = O(h^2), \quad \varepsilon_{\text{RK4}} = O(h^4)$$

- ε — norm of state error at a fixed final time
- h — step tending to zero
- assumptions — bounded derivatives; field locally Lipschitz near solution

Absolute-stability region: hλ values with bounded repeated updates on ẋ=λx. A large RK4 step can fail.



Order of accuracy; absolute stability — [4] FNC §6.4 · [4] FNC §11.3

MODEL

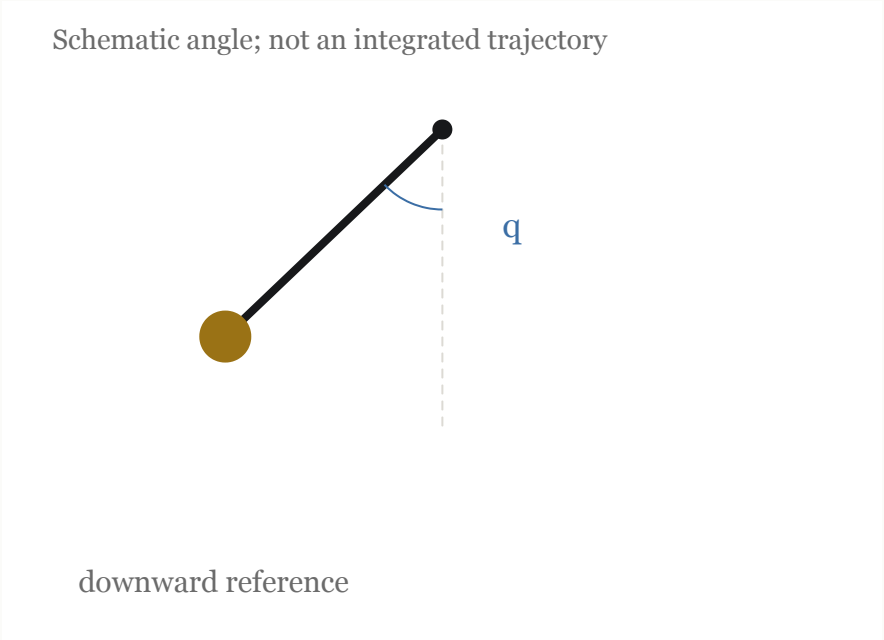
Pendulum Simulation

The restoring acceleration has a sign

$$\dot{q} = v, \qquad \dot{v} = -\frac{g_0}{\ell} \sin q$$

- `q` angle from downward vertical
- `v` angular velocity
- g_0, ℓ positive gravity acceleration and length

Unforced, undamped motion is the model for this comparison.



Simple pendulum; first-order system — [1] Tedrake, ch.2

DEFINITION

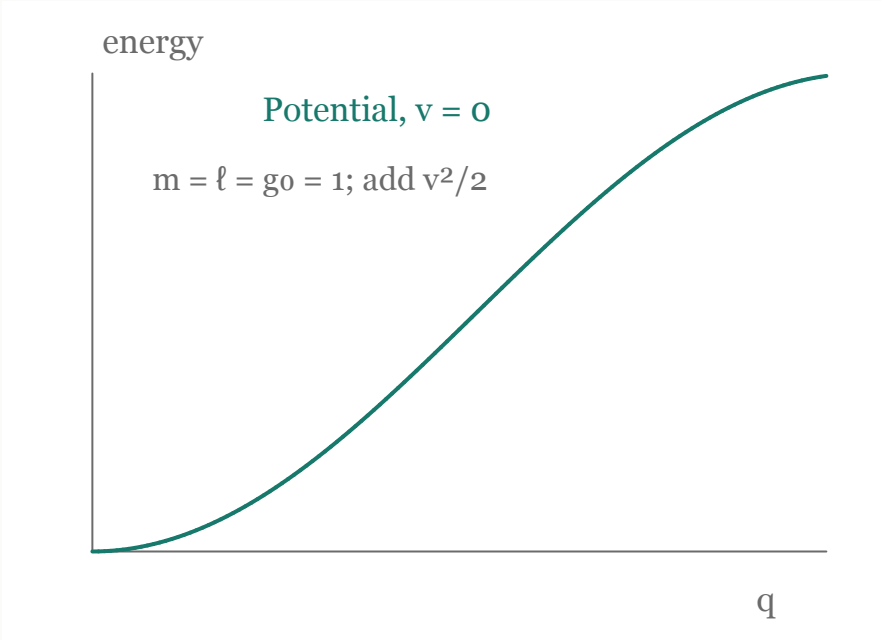
Pendulum Simulation

Mechanical energy is a diagnostic

$$E(q, v) = \tfrac{1}{2} m \ell^2 v^2 + m g_0 \ell (1 - \cos q)$$

`q, v`	angle and angular velocity
`m, ℓ, g₀`	mass, length and gravity
`two terms`	kinetic and potential energy

The zero of potential energy is a convention; energy differences have meaning.



Mechanical energy — [1] Tedrake, ch.2

DERIVATION

Pendulum Simulation

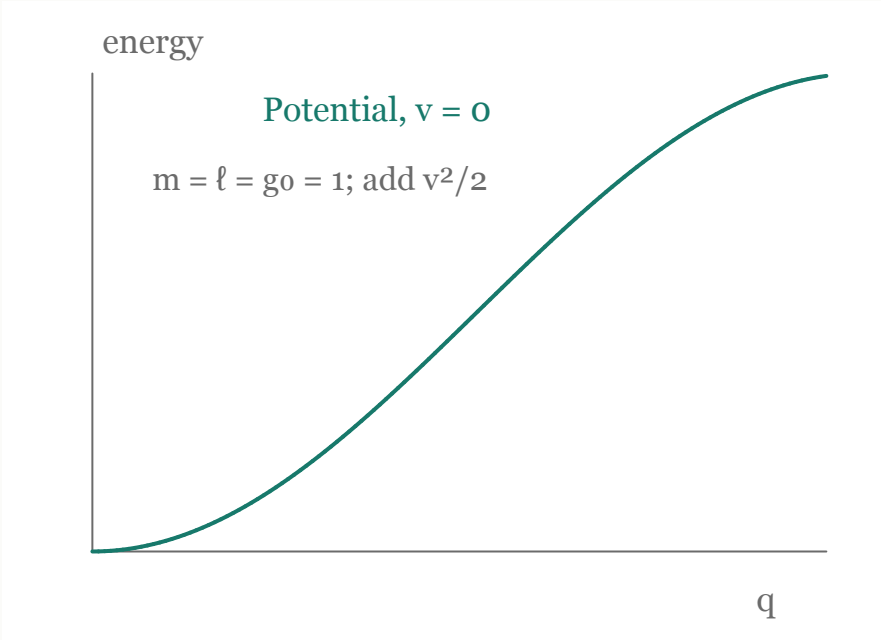
The exact dynamics conserve energy

$$E = \tfrac{1}{2} m \ell^2 v^2 + m g_0 \ell (1 - \cos q)$$

$$\dot{E} = m \ell^2 v \dot{v} + m g_0 \ell \sin q \, \dot{q} = 0$$

- $\dot{q}=v$ kinematic identity
- $\dot{v}$ $-(g_0/\ell)\sin q$
- $\dot{E}$ rate of mechanical energy

The two terms cancel only for the stated unforced, undamped model.



Conservation of mechanical energy — [1] Tedrake, ch.2

YOUR TURN

Integrated Pendulum Problem

Choose the state and the check first

$$\ddot{q} = -\sin q, \quad q_0 = 1, \quad v_0 = 0$$

`q₀,v₀` angle and angular velocity

`model` unforced, undamped; m=ℓ=g0=1

Which state and field give a first-order ODE for this model?

A. x=q; f=−sin q

B. $x=(q,v)^T$; $f=(v,-\sin q)^T$

C. $x=(q,v)^T$; $f=(-\sin q,v)^T$

Pendulum IVP; numerical integration — [1] Tedrake, ch.2 · [4] FNC §6.4 · Course example

YOUR TURN

Integrated Pendulum Problem

Compare states at t = 0.1

$$x_0 = (1, 0)^T,\ h = .1,\ f(q, v) = (v, -\sin q)^T$$

`x=(q,v)ᵀ` angle and angular velocity

`h` full time step

Why can midpoint give $q_1 \neq q_0$ even though $v_0 = 0$?

A. Nonzero velocity at the trial midpoint.

B. Twice the physical time.

C. A different initial velocity.

Euler and explicit midpoint methods — [4] FNC §6.2 · [4] FNC §6.4 · Course example

YOUR TURN

Integrated Pendulum Problem

Repeat to the same final time

$$\dot{x} = (v, -\sin q)^T, \quad x(0) = (1, 0)^T, \quad T = 10$$

N	step count
E	$v^2/2+1-\cos q$
ΔE	E at T minus E at 0

Halve h: how do we still end at time T=10?

A. Keep N.

B. Stop at equal E.

C. Double N.

Global error; step refinement — [4] FNC §6.2 · [4] FNC §6.4 · Course example

YOUR TURN

Integrated Pendulum Problem

Compare two energy-matched states

$$E(q, v) = v^2/2 + 1 - \cos q$$

`q,v` angle and angular velocity

`E` energy

What does equal energy of two pendulum states establish?

A. Their angles and velocities are equal.

B. They can still be different states.

C. Both numerical trajectories are exact.

Energy level; state error — [1] Tedrake, ch.2 · Course example

YOUR TURN

Independent Transfer

Transfer: a moving pendulum near the top

$$q_0 = \pi, \quad v_0 = 1, \quad h = .1, \quad m = \ell = g_0 = 1$$

`q₀,v₀`	initial angle and angular velocity
`f(q,v)`	$(v, -\sin q)^T$
`E(q,v)`	$v^2/2 + 1 - \cos q$

One call: which method and energy claim is correct?

A. Euler; this step decreases energy.

B. Midpoint; one call is enough.

C. Euler; every step increases energy.

Pendulum IVP; explicit Euler method — [1] Tedrake, ch.2 · [4] FNC §6.2 · Course example

YOUR TURN

Simulation Checklist

Simulation workflow

$$\dot{x} = f(x, t), \qquad x(t_0) = x_0$$

`x`	complete state
`model`	unforced, undamped pendulum; $\ddot{q} = -\sin q$
`h`	chosen integration step

For this pendulum simulation, which workflow distinguishes errors?

A. Use q alone; shrink h until the curve looks smooth.

B. Use f as the next state; inspect only energy.

C. Use (q,v), integrate rates, compare states and energy.

Verification; validation; step refinement — [5] NASA, V&V glossary · [4] FNC §6.2 · Course checklist

CONNECTION

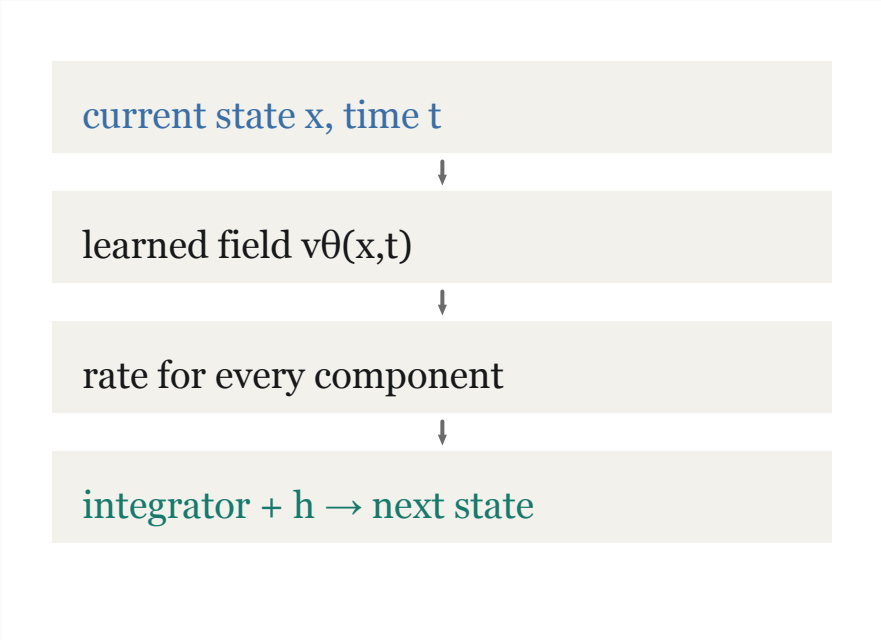
Simulation Checklist

A learned field still needs an integrator

$$\dot{x} = v_\theta(x, t)$$

- `vθ` rate model with learned parameters θ
- `x` generated state
- `t` continuous evolution coordinate

A learned field supplies rates; Euler or Runge–Kutta advances the state.



Neural ODE; numerical integration — [6] Chen et al. (2018), §1

References

Sources for models and exact dynamics

[1] Tedrake, R. Underactuated Robotics, ch.2. MIT notes, 2024.

[2] Lynch, K. M. & Park, F. C. Modern Robotics, §8.1. Cambridge, 2017; official supplement.

[3] Strang, G. Differential Equations and exp(At). MIT 18.06, lecture 23; OCW 2010.

[7] NIST Digital Library of Mathematical Functions, §1.2(vi): defective and nondefective matrices.

References

Sources for integration and stability

[4] Driscoll, T. A. & Braun, R. J. Fundamentals of Numerical Computation. Online text, §§6.1–6.4, 11.3; accessed 2026-10-20.

[5] NASA Glenn. Glossary of Verification and Validation Terms. Online tutorial; accessed 2026-10-20.

[6] Chen, R. T. Q., Rubanova, Y., Bettencourt, J. & Duvenaud, D. Neural Ordinary Differential Equations. NeurIPS, 2018, §1.