CS**5643**

03 Solving ordinary differential equations

Steve Marschner
Cornell University
Spring 2025

Ordinary differential equation

An equation involving an unknown function and its derivatives

- but with only one independent variable (typically time)
- general form $f(t, y(t), y'(t), y''(t), \ldots, y^{(k)}) = 0$ for all t

In an *initial value problem* we know what is happening now and want to know the future

- boundary conditions are all at $t = 0$: $y(0), y'(0), \ldots, y^{(k-1)}(0)$
- goal: find $y(t)$ for all $t > 0$
- (notice that we need starting values for the derivatives less than the highest one involved)

In this course usually $k = 2$ (but sometimes 1)

Systems of ODEs

Typically there are multiple unknown functions

- e.g. the x and y coordinates of a particle, or of many particles, …

Can think of this as a system of interdependent ODEs…

…or simply as an ODE with a vector-valued unknown

- $\mathbf{f}(t, \mathbf{y}(t), \dot{\mathbf{y}}(t), \ddot{\mathbf{y}}(t)) = \mathbf{0}$ where $\mathbf{y} : \mathbb{R} \to \mathbb{R}^N$ and $\mathbf{f} : \ldots \to \mathbb{R}^N$

In this setting the solution is a path though $\mathbb{R}^N$

- an N-dimensional parameterized curve
- solving $\mathbf{f}$ tells you how to continue this curve by looking at the position, tangent, curvature, etc. at the end

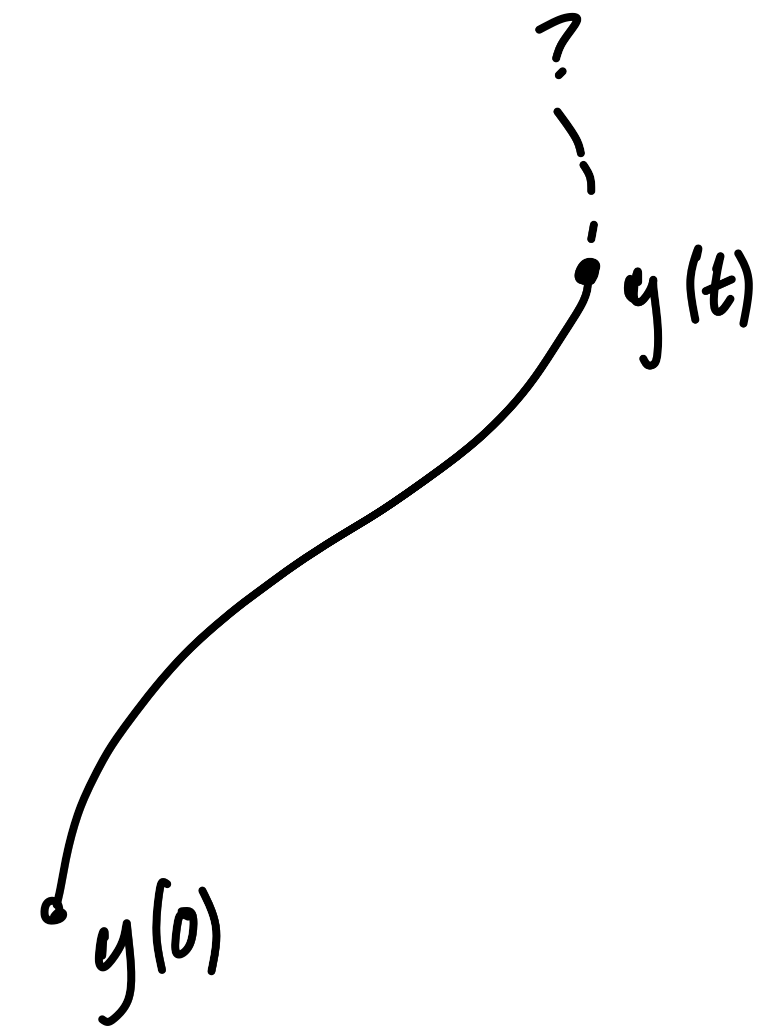

Some simplifications

Most often we work with ODEs that are solved for the highest derivative:

- this is called an *explicit* ODE
- $\mathbf{y}^{(k)}(t) = \mathbf{f}(t, \mathbf{y}(t), \ldots, \mathbf{y}^{(k-1)})$
- or in the $k = 2$ case: $\ddot{\mathbf{y}}(t) = \mathbf{f}(t, \mathbf{y}(t), \dot{\mathbf{y}}(t))$

Also we can choose to work only with:

- first-order systems ($k = 1$)
- autonomous systems ($\mathbf{f}$ independent of t)
- (next slides)

Reduction to first order

Someone gave me an ODE $\mathbf{y}^{(k)}(t) = \mathbf{f}(t, \mathbf{y}(t), \ldots, \mathbf{y}^{(k-1)})$ in N variables

I'll give back a first-order ODE

- unknown functions $\mathbf{y}(t)$, $\mathbf{y}_1(t)$, ..., $\mathbf{y}_{k-1}(t) : \mathbb{R} \to \mathbb{R}^N$

equation: $$\begin{bmatrix} \dot{\mathbf{y}} \\ \dot{\mathbf{y}}_1 \\ \vdots \\ \dot{\mathbf{y}}_{k-2} \\ \dot{\mathbf{y}}_{k-1} \end{bmatrix}(t) = \begin{bmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \\ \vdots \\ \mathbf{y}_{k-1} \\ \mathbf{f}(t, \mathbf{y}, \mathbf{y}_1, \ldots, \mathbf{y}_{k-1}) \end{bmatrix}(t)$$

- this is a single first-order ODE in kN variables with the **same solution**

2nd order:

$$M\ddot{x}(t) = f(t, x(t), \dot{x}(t))$$

$$\Downarrow$$

$$\begin{bmatrix} \dot{x}(t) \\ \dot{v}(t) \end{bmatrix} = \begin{bmatrix} v(t) \\ M^{-1} f(t, x(t), v(t)) \end{bmatrix}$$

So at the highest level of abstraction the order doesn't matter

- but sometimes can get better results by remembering it started as a higher-order system

Autonomous vs. non-autonomous

Sometimes you see time as an explicit parameter, sometimes not

- $\dot{\mathbf{y}}(t) = \mathbf{f}(\mathbf{y}(t))$ is "autonomous"
- $\dot{\mathbf{y}}(t) = \mathbf{f}(t, \mathbf{y}(t))$ is "non-autonomous"

If we want to do math or write code without the t, we can make a simple conversion:

$$u(t) = \begin{bmatrix} y(t) \\ \tau \end{bmatrix}$$

$$\text{ODE: } \begin{bmatrix} \dot{y}(t) \\ \dot{\tau} \end{bmatrix} = \begin{bmatrix} f(\tau, y(t)) \\ 1 \end{bmatrix} \; ; \quad \tau(0) = 0$$

- and just relabel the axis to τ

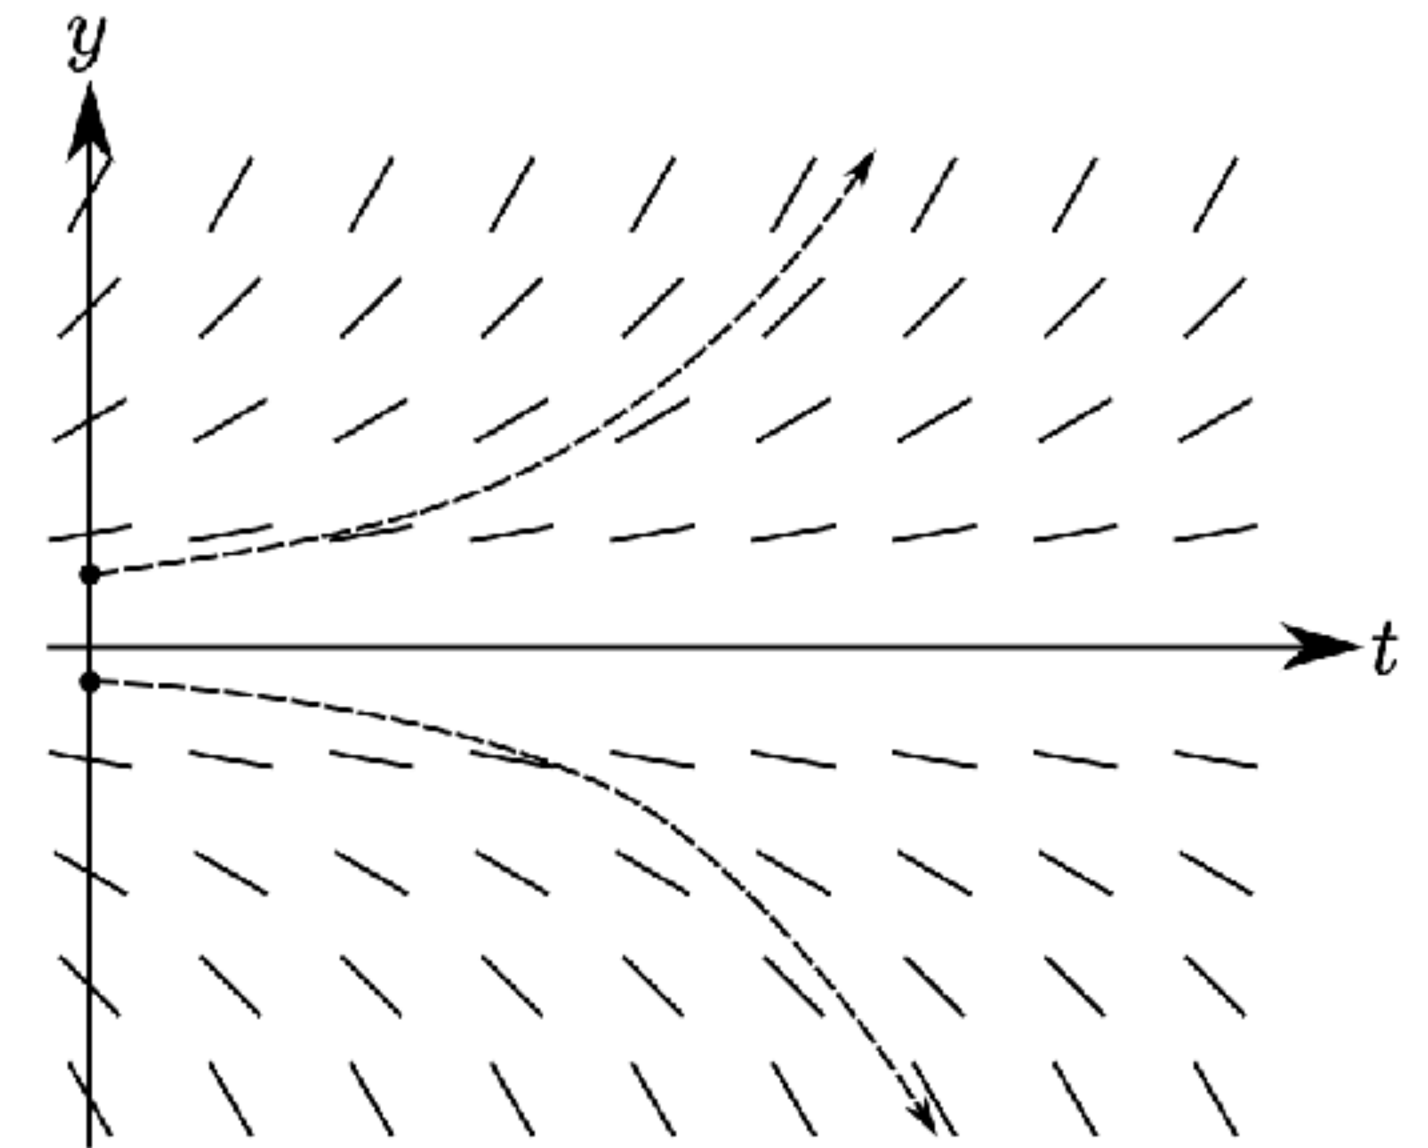

Time-independent

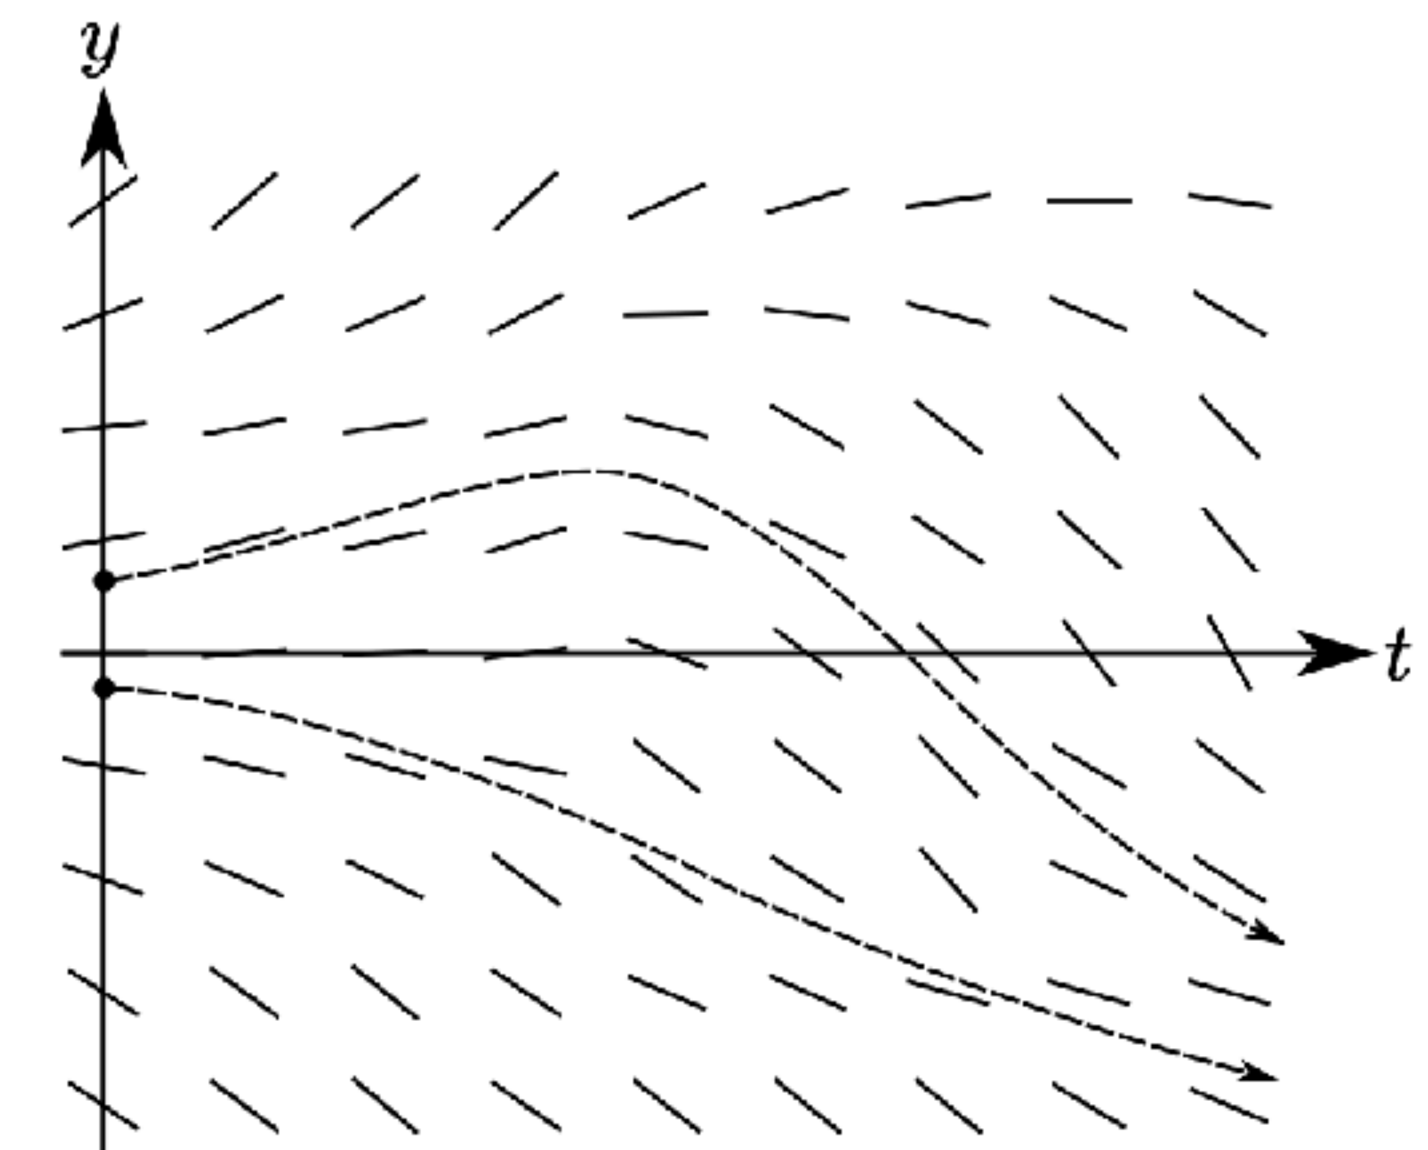

Time-dependent

from Justin Solomon, *Numerical Algorithms*

Vector field picture

Now that we only have systems of the form $\dot{\mathbf{y}}(t) = \mathbf{f}(\mathbf{y}(t))$ there is a simple mental model:

- $\mathbf{y}(t)$ is the path of a point through the *state space* of the system
- remember $\mathbf{y}$ here is after a reduction to first order, so for instance in a Newtonian particle system
  - $\mathbf{y}$ includes both the position $\mathbf{x}$ and the velocity $\mathbf{v}$
- $\mathbf{f}$ is a vector field in that state space that tells the particle which way to go
- so the process reduces to advection through a flow field (though in many dimensions)

canonical example: harmonic oscillator in 1D, $m\ddot{x} = -kx$

- $\begin{bmatrix} \dot{x} \\ \dot{v} \end{bmatrix} = \begin{bmatrix} v \\ -(k/m)x \end{bmatrix}$ or with appropriate choice of units $\begin{bmatrix} \dot{x} \\ \dot{v} \end{bmatrix} = \begin{bmatrix} v \\ -x \end{bmatrix}$
- aka $\dot{\mathbf{y}} = R\mathbf{y}$ where R is a rotation by –90 degrees

Canonical example: harmonic oscillator

e.g. mass on a spring, plucked rubber band, tuning fork

1D ODE $m\ddot{x} = -kx$

$$\begin{bmatrix}\dot{x}\\ \dot{v}\end{bmatrix} = \begin{bmatrix} v \\ -(k/m)x \end{bmatrix}$$

- or with appropriate units:

$$\begin{bmatrix}\dot{x}\\ \dot{v}\end{bmatrix} = \begin{bmatrix} v \\ -x \end{bmatrix}$$

- aka $\dot{\mathbf{y}} = R\mathbf{y}$ where R is a rotation by –90 degrees
- solutions are like $x(t) = \sin t, v(t) = \cos t$

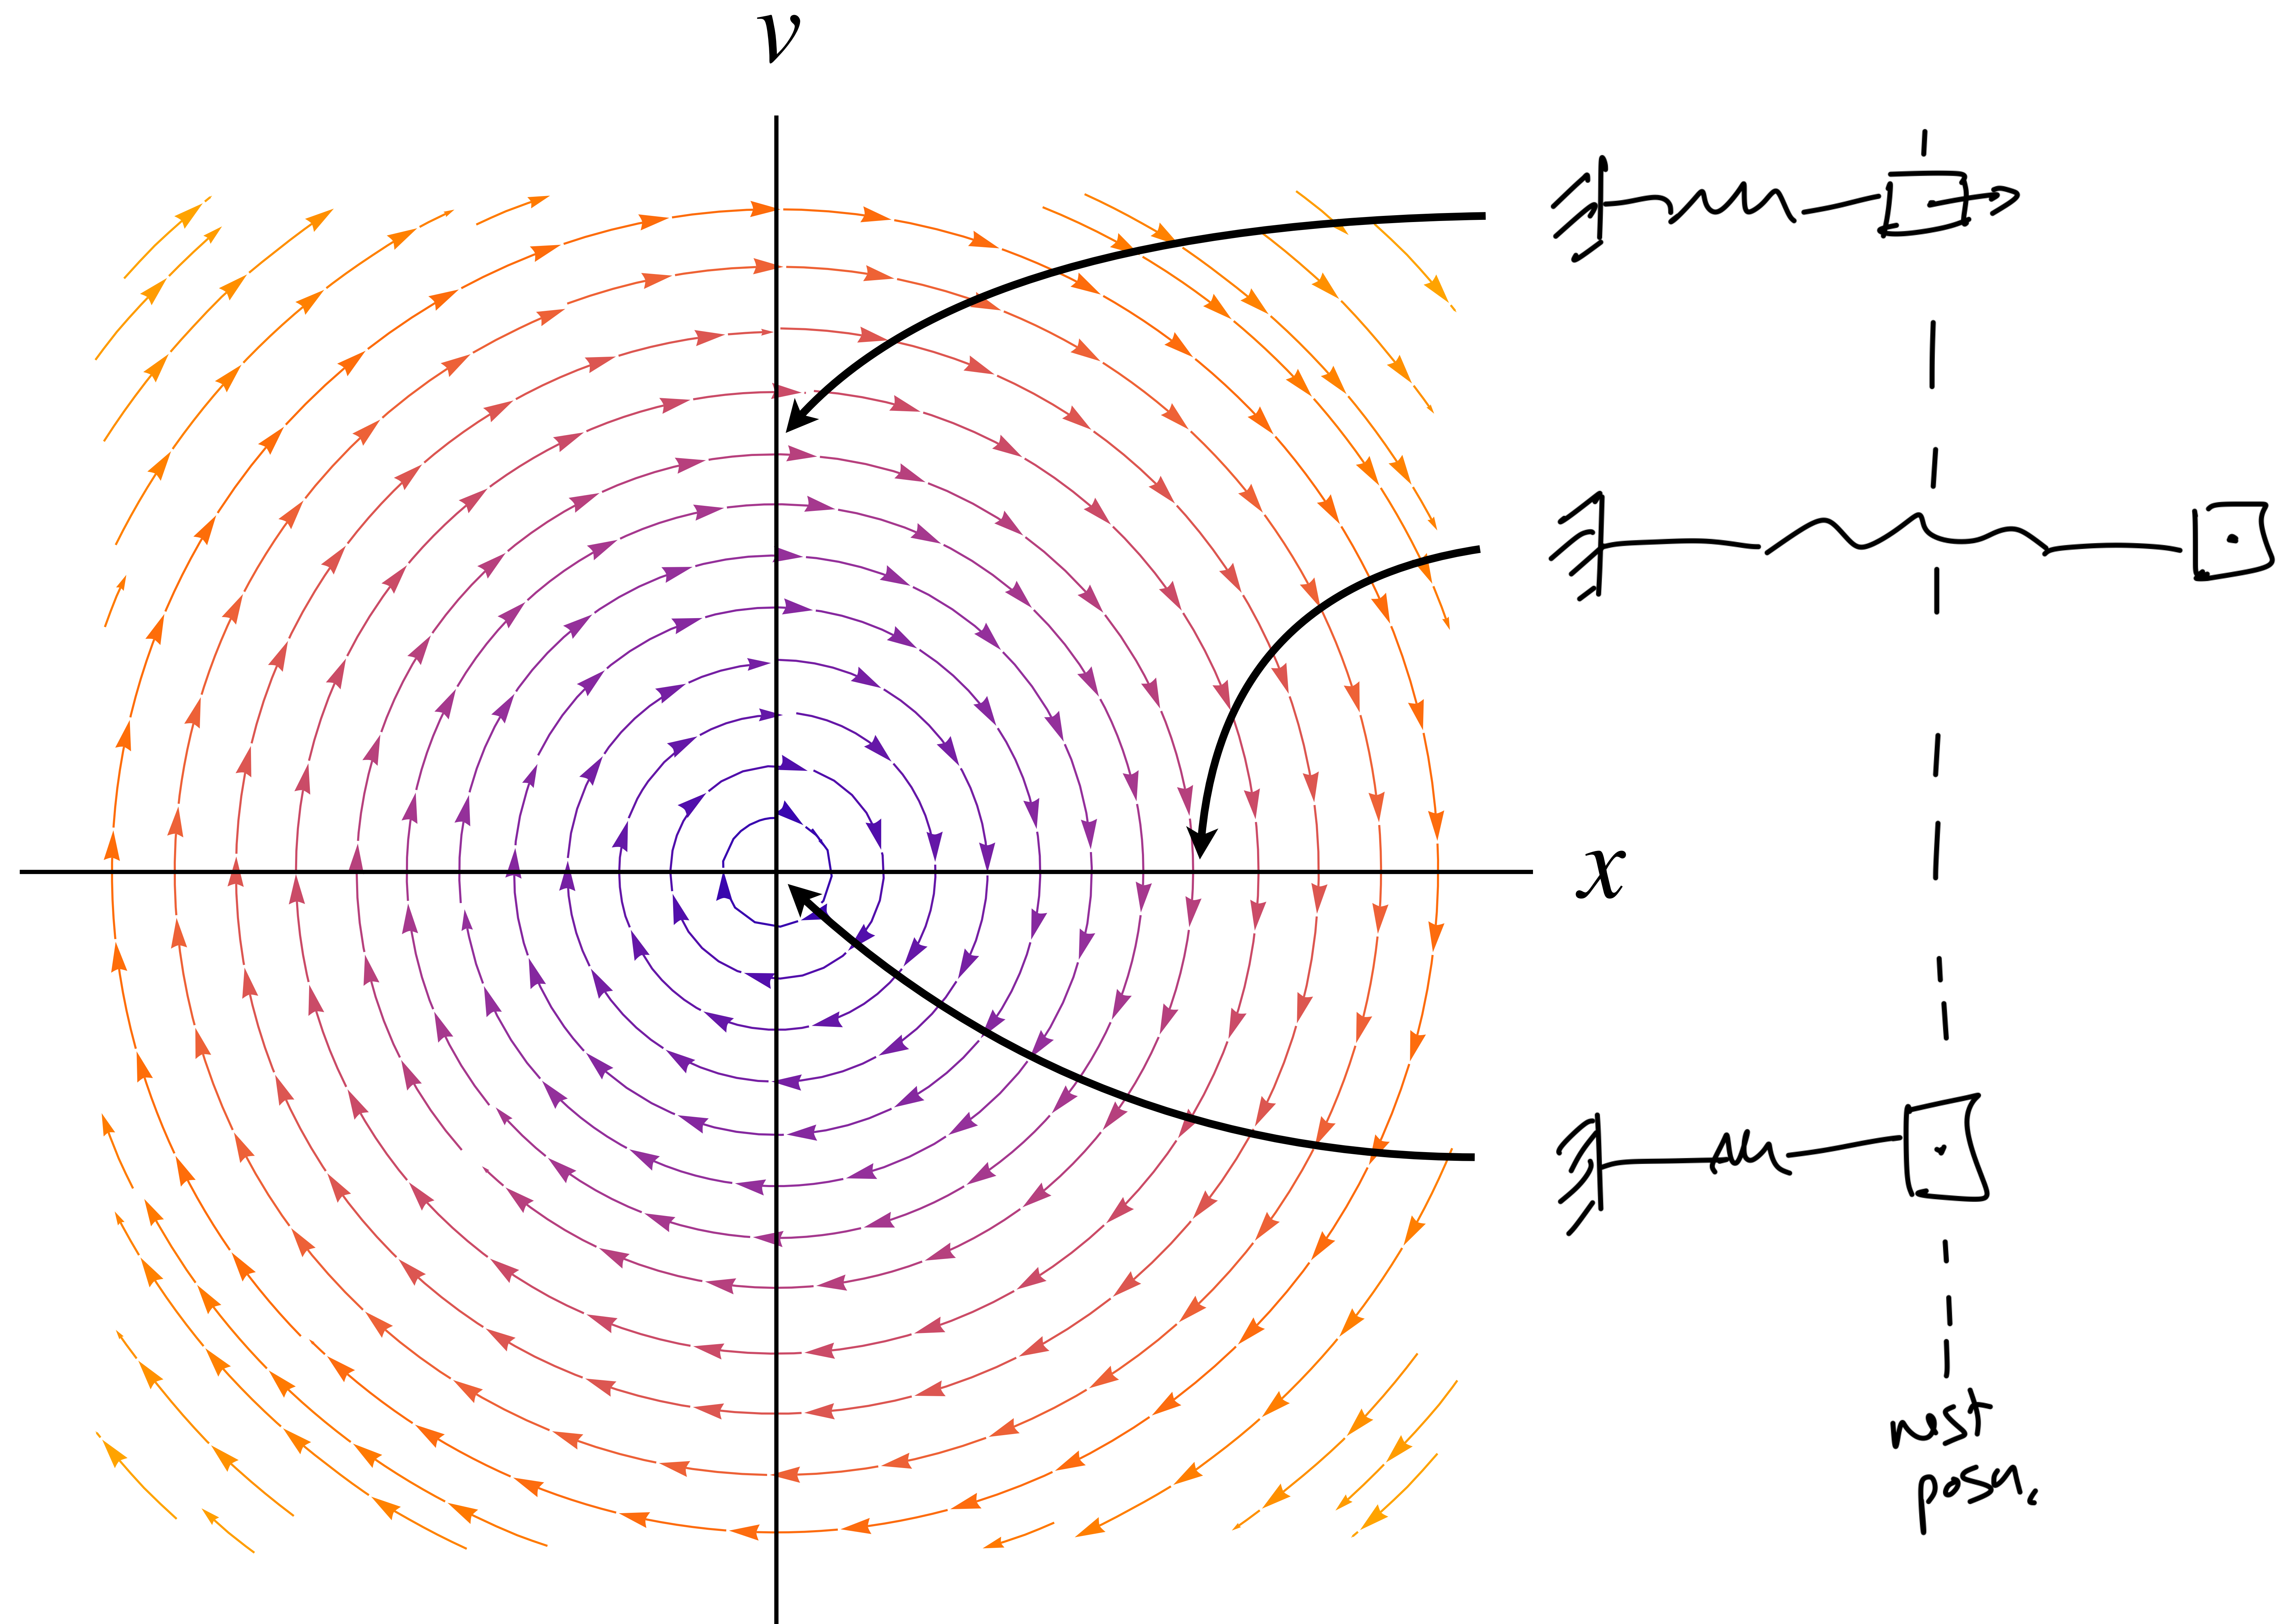

Canonical example: exponential decay

E.g. cup of tea cooling off
or particle slowing in fluid

1D ODE: $\dot{y} = -ky$

- solutions are like $y(t) = \exp(-t)$

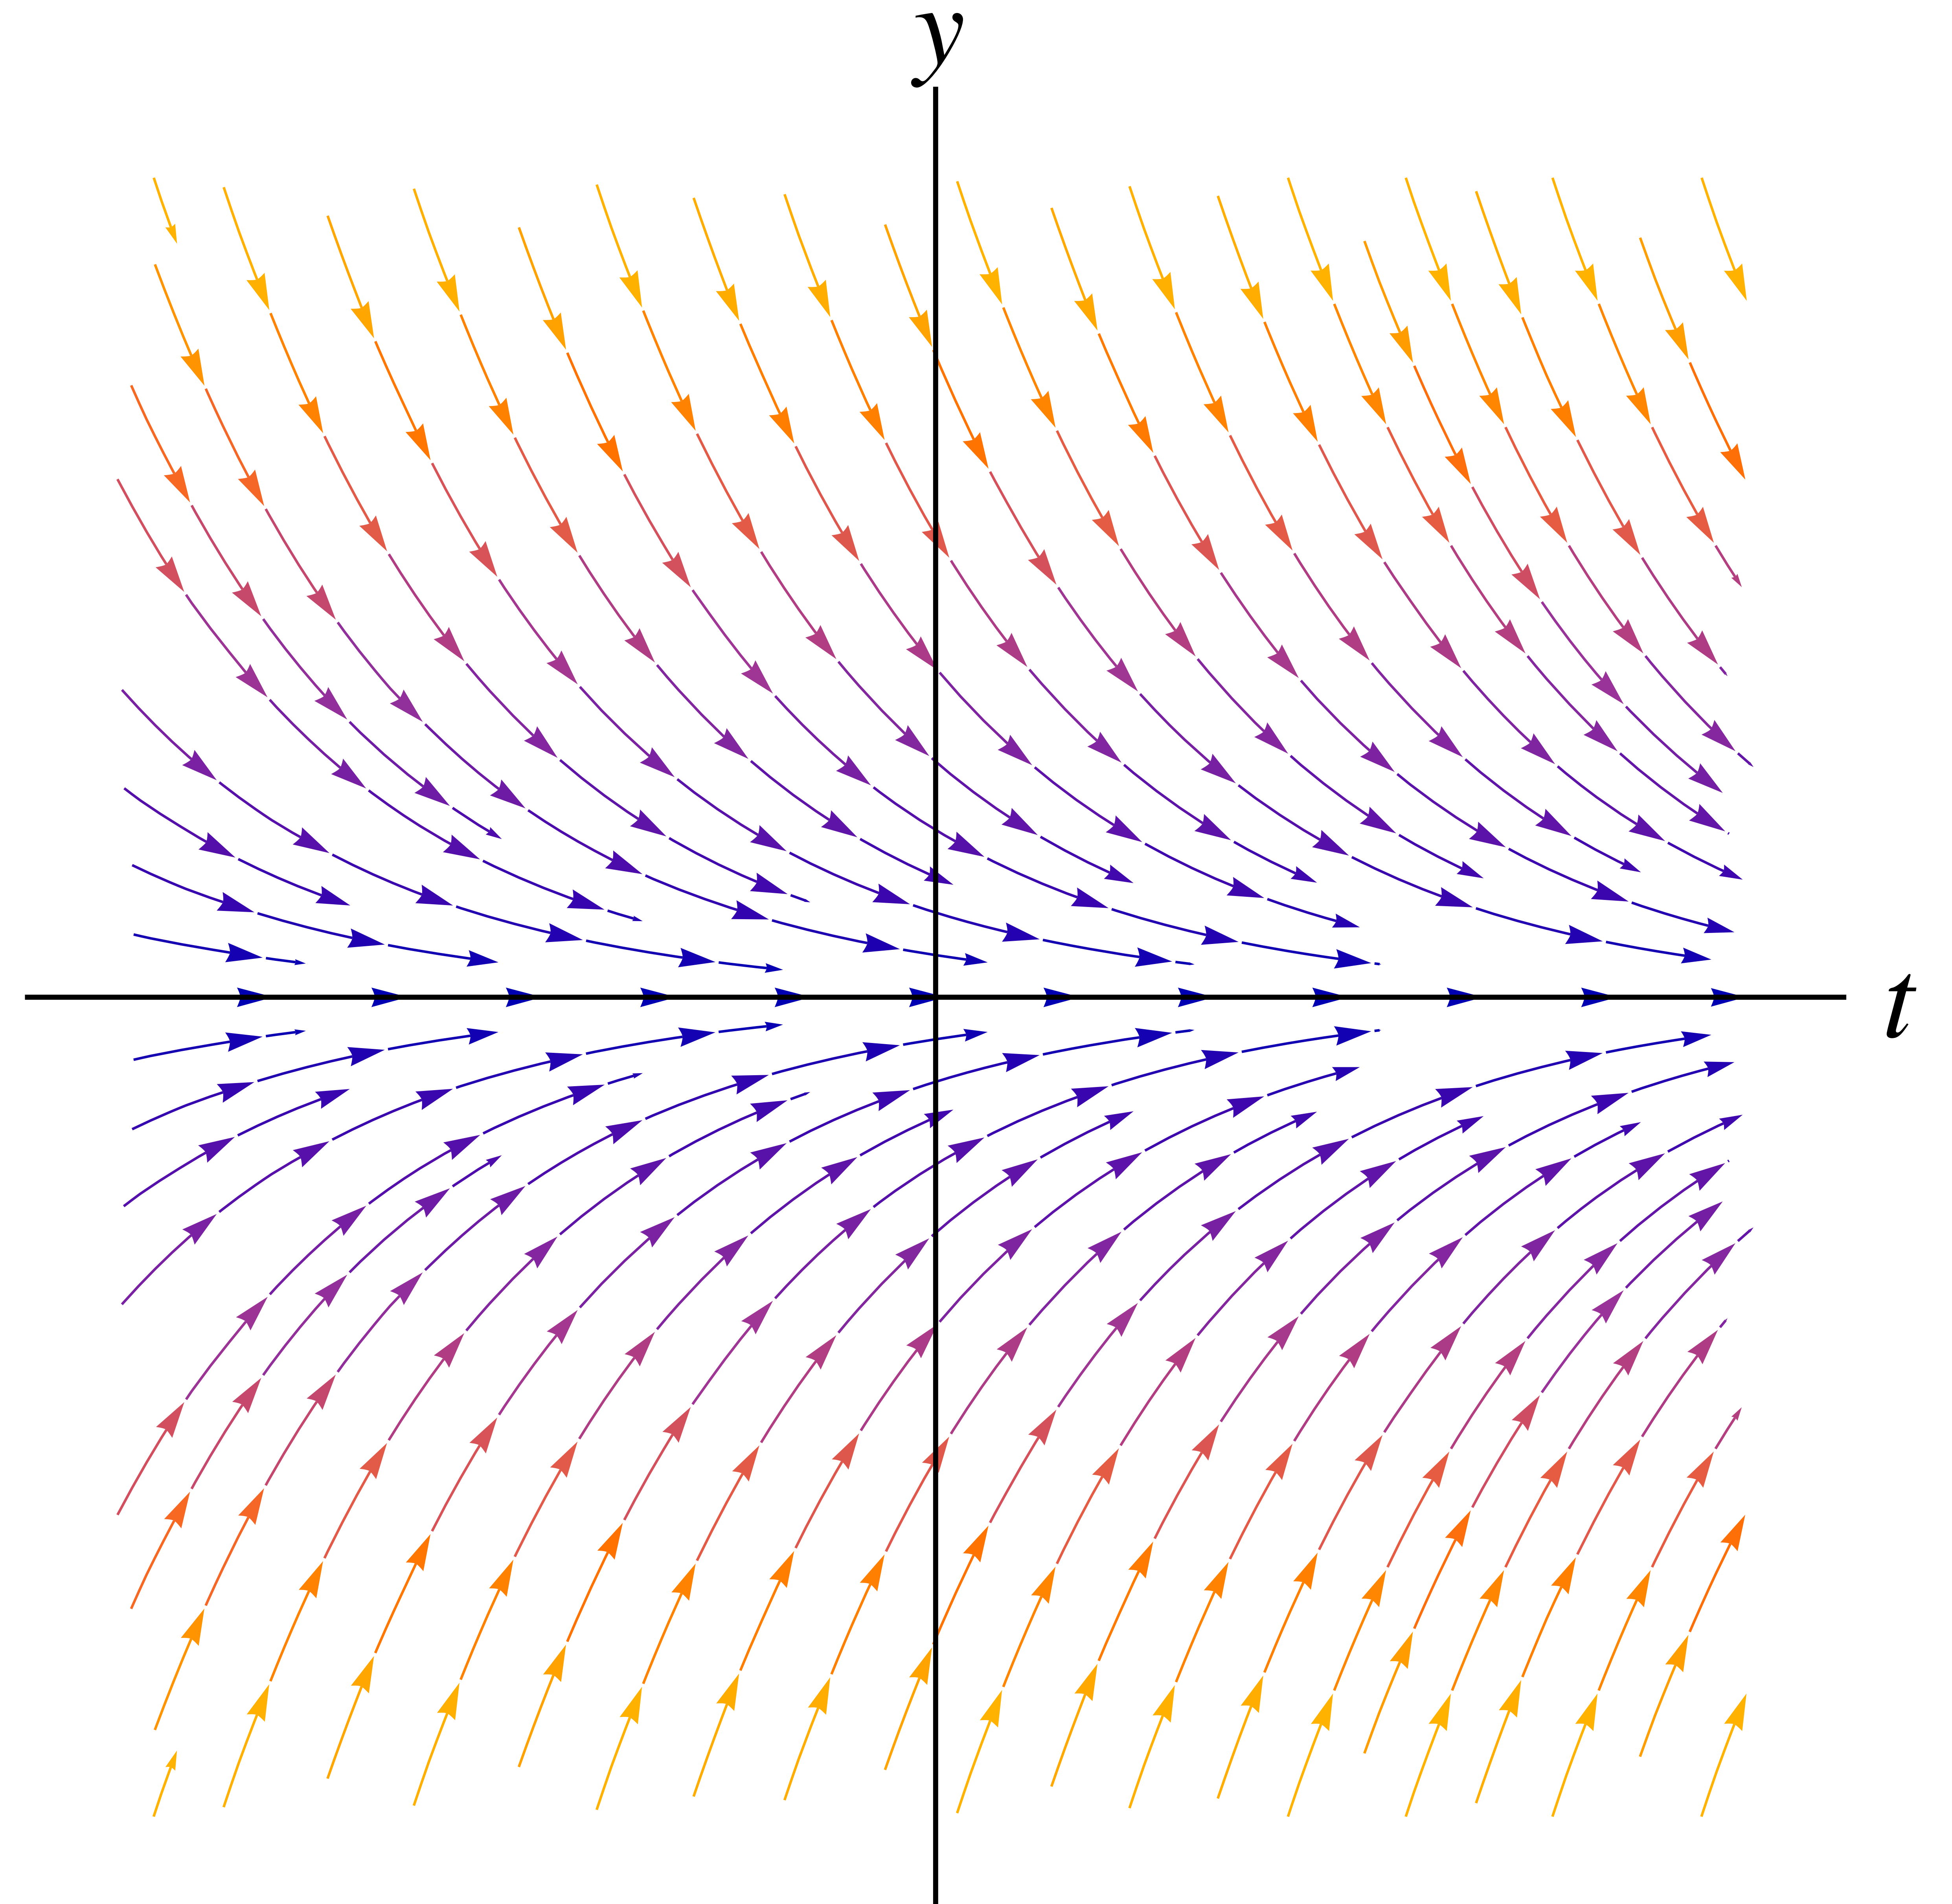

Numerical solution methods

Most ODEs don't have closed form solutions
so we resort to numerical approximation

- the only thing we know how to compute is the $\mathbf{f}$ in $\dot{\mathbf{y}}(t) = \mathbf{f}(\mathbf{y}(t))$

Want to compute approximate values of the unknown $\mathbf{y}(t)$
for desired values of t

- to do this we compute $\mathbf{y}(t_k)$ for a series of *time steps*
  - from where we know $\mathbf{y}$ (canonically at $t = 0$)
  - to where we want $\mathbf{y}$ (e.g. at the time of each animation frame)
- compute each step from the results of previous steps
  using a *local approximation* to $\mathbf{y}$
- different local approximations lead to different time stepping algorithms, known as numerical integration methods or ODE solvers or just "integrators" or "solvers."

Setup for simple integration methods

Start with a constant step size h

- time steps are equally spaced, $t_{k+1} = t_k + h$
- if we start at $t_0 = 0$ then $t_k = kh$ and the number of steps to reach time T is T/h

We want some equation we can solve to approximate $\mathbf{y}(t_{k+1})$ assuming we know $\mathbf{y}(t_k)$

- in practice we don't know $\mathbf{y}(t_k)$ exactly; we just have the approximation from the previous step
- I will use $\mathbf{y}_k$ for the approximation we computed at step k and $\mathbf{y}(t_k)$ for the actual value
- the goal of our method is to ensure $\mathbf{y}_k \approx \mathbf{y}(t_k)$ so that the points $(t_k, \mathbf{y}_k)$ are a good approximation to the solution function $\mathbf{y}(t)$
- an important question: how to quantify how accurately $\mathbf{y}_k$ approximates $\mathbf{y}(t_k)$

Euler's integrators

Most integrators can be derived from a Taylor expansion

- after all it's the first tool we reach for when we want a local approximation

E.g. let's expand $\mathbf{y}$ around $t = t_k$:

$$\mathbf{y}(t) = \mathbf{y}(t_k) + \dot{\mathbf{y}}(t_k)(t - t_k) + O((t - t_k)^2)$$

- evaluate at $t_{k+1} = t_k + h$ and substitute the ODE $\dot{\mathbf{y}}(t) = \mathbf{f}(\mathbf{y}(t))$

$$\mathbf{y}(t_{k+1}) = \mathbf{y}(t_k) + h\mathbf{f}(\mathbf{y}(t_k)) + O(h^2)$$

- leading to the *timestep equation* $\mathbf{y}_{k+1} = \mathbf{y}_k + h\mathbf{f}(\mathbf{y}_k)$ known as "Euler's method" or "forward Euler"

This is a *first order accurate*, *explicit* integration method

- "explicit" because the timestep equation is already solved for $\mathbf{y}_{k+1}$; it is an explicit formula
- "first order accurate" because the error is proportional to h^2

Euler's integrators

Alternatively we could expand $\mathbf{y}$ around $t = t_{k+1}$:

$$\mathbf{y}(t) = \mathbf{y}(t_{k+1}) + \dot{\mathbf{y}}(t_{k+1})(t - t_{k+1}) + O((t - t_{k+1})^2)$$

- evaluate at $t_k = t_{k+1} - h$ and substitute the ODE $\dot{\mathbf{y}}(t) = \mathbf{f}(\mathbf{y}(t))$

$$\mathbf{y}(t_k) = \mathbf{y}(t_{k+1}) - h\mathbf{f}(\mathbf{y}(t_{k+1})) + O(h^2)$$

- leading to the *timestep equation* $\mathbf{y}_{k+1} = \mathbf{y}_k + h\mathbf{f}(\mathbf{y}_{k+1})$ known as "backward Euler's method"

This is a first order accurate, *implicit* integration method

- "implicit" because the timestep equation needs to be solved to find $\mathbf{y}_{k+1}$
- "first order accurate" because the error is still proportional to h^2

How does your error shrink?

If things are working at all, we can get any accuracy we need by decreasing h

- that is, $\lim_{h \to 0} [\mathbf{y}_k - \mathbf{y}(t_k)] = 0$

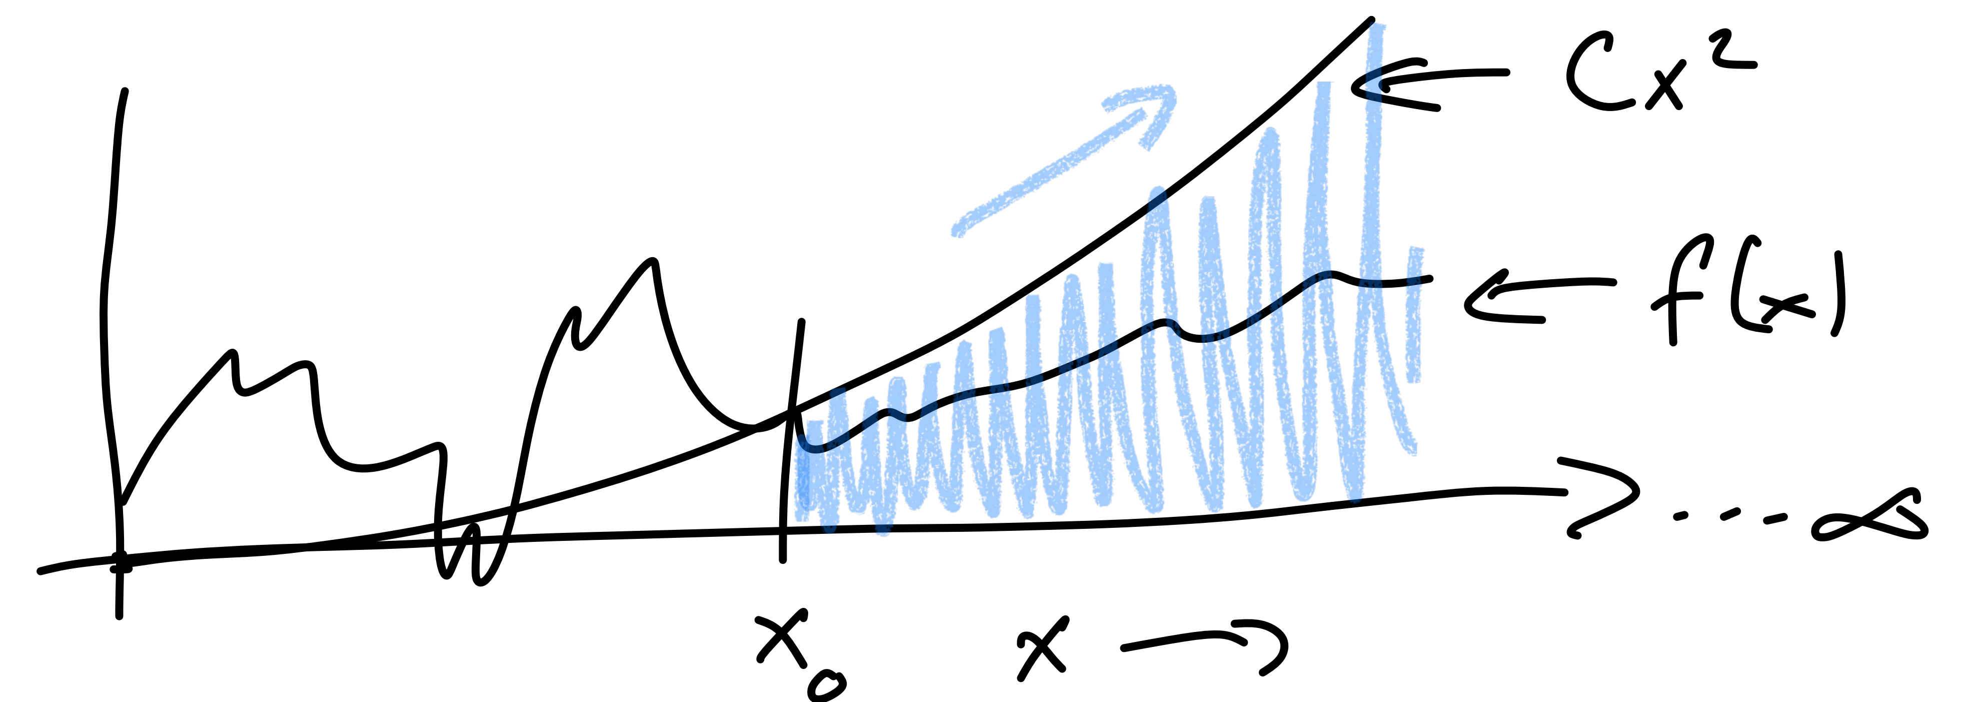

we compare integrators' accuracy in terms of *asymptotic rate* of convergence

- recall big-O notation $f(x) \in O(x^2)$ as $x \to \infty$ means there are constants C and x_0 such that $x > x_0 \implies f(x) \leq Cx^2$
- we can use the same idea for asymptotics as $x \to 0$: $f(x) \in O(x^2)$ as $x \to 0$ means there exist constants C and δ such that $x < \delta \implies f(x) \leq Cx^2$

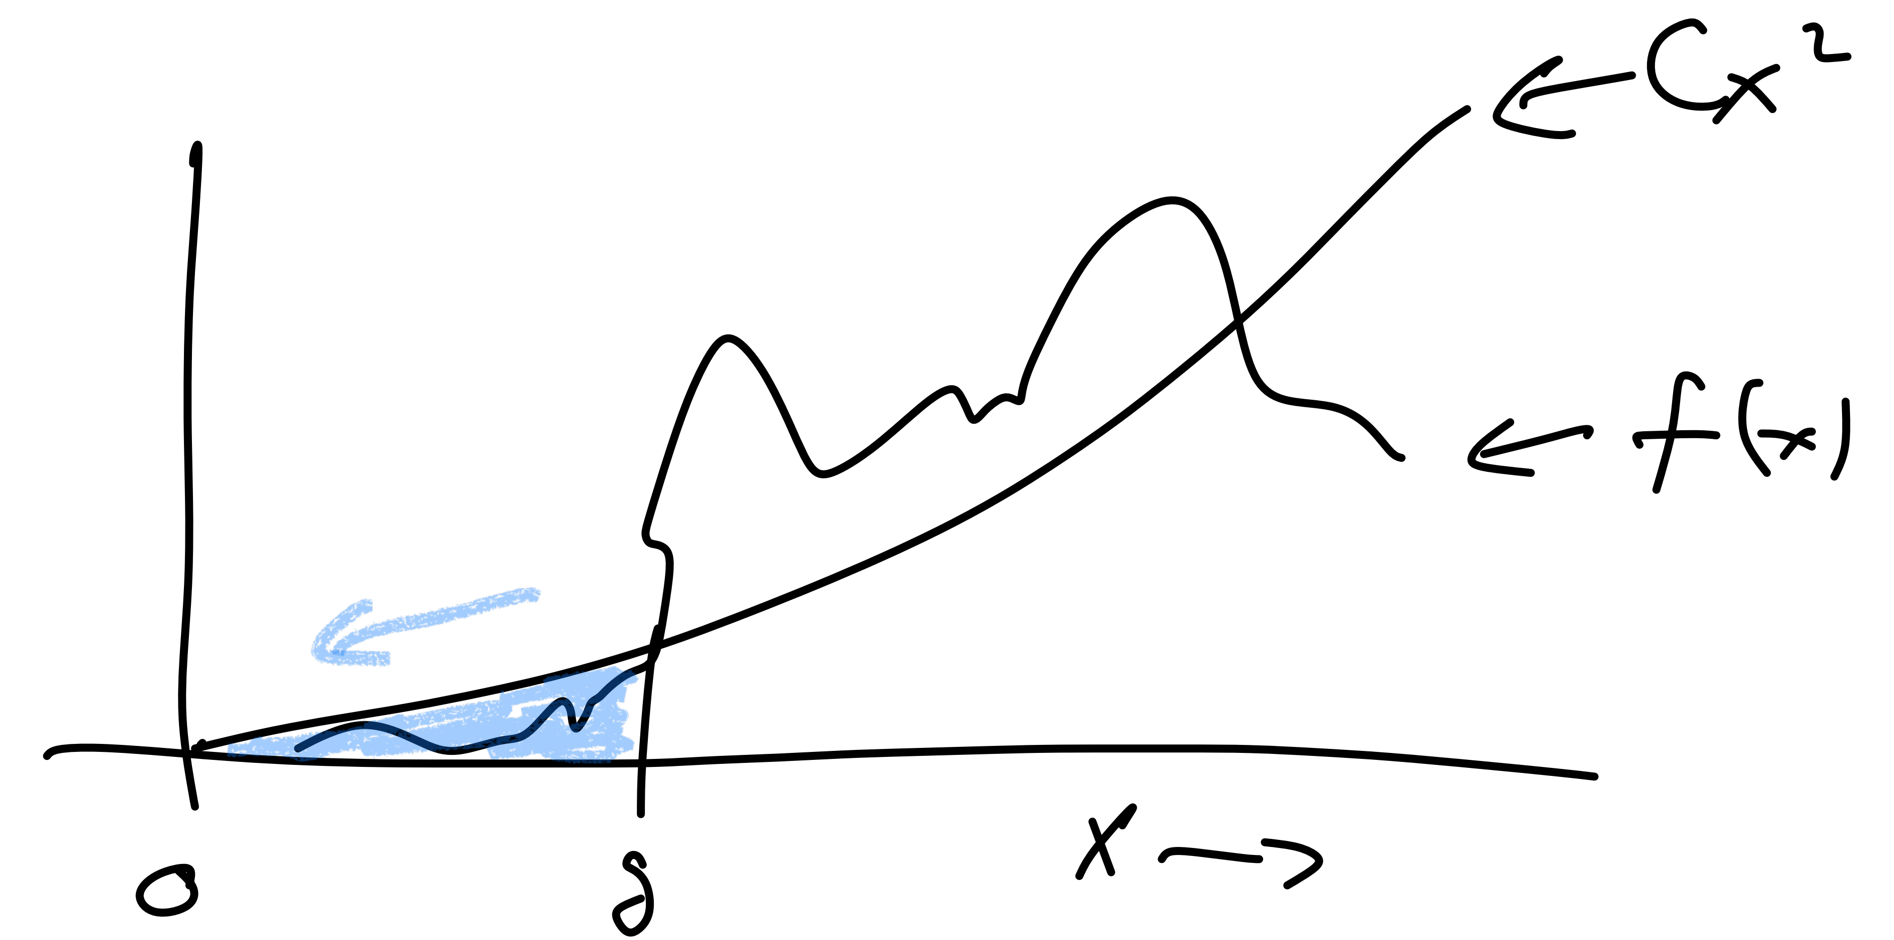

How does your error grow?

The error in a time-stepped approximation accumulates

- in worst case (sadly not so uncommon) all the errors point the same way so the error after N steps is N times the error in one step
- to get to time T requires $N \approx T/h$ steps
- so if error in one step is $O(h^p)$ then error after N steps is $O(h^{p-1})$

Nomenclature for integrators works two ways

- *p*th order integrator is "accurate to *p*th order" in one step, meaning the error is $O(h^{p+1})$
- *p*th order integrator has order-*p* error after a fixed time, meaning the error is $O(h^p)$

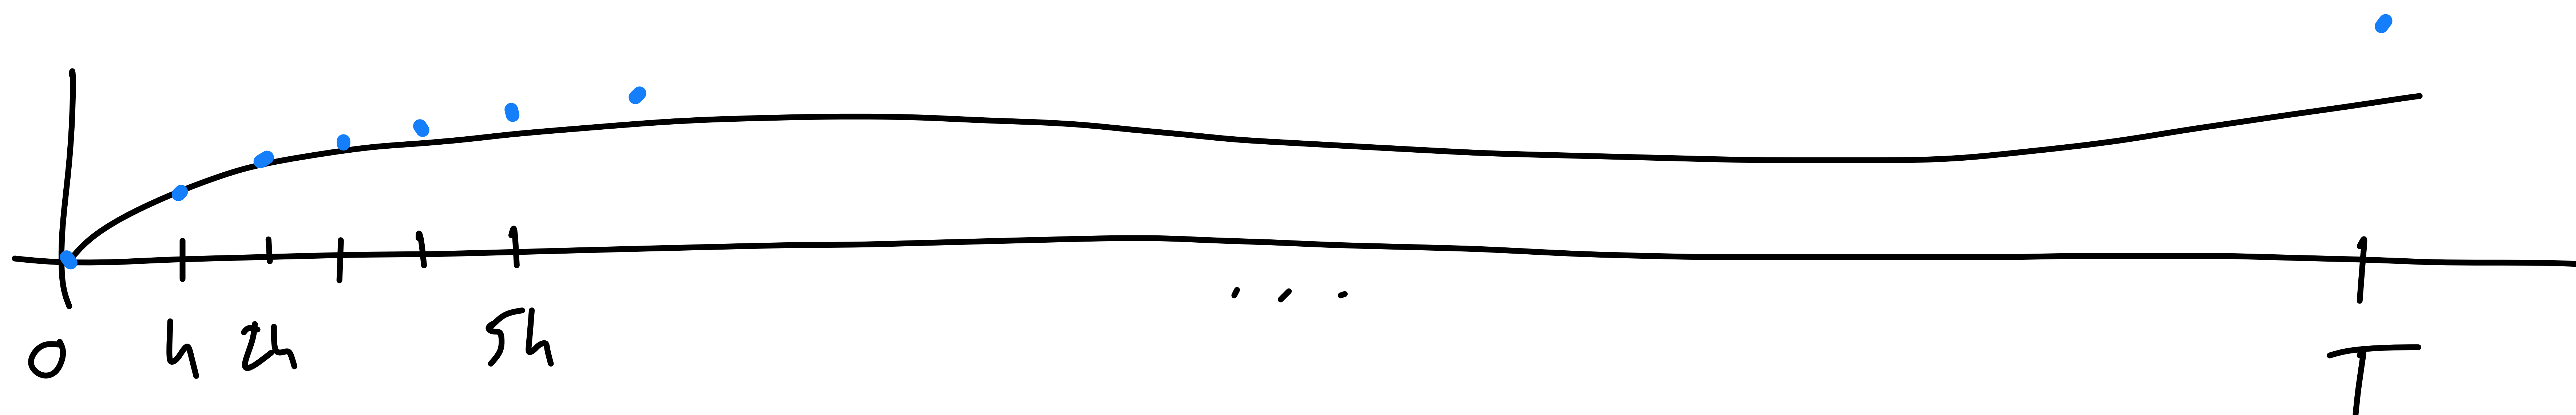

Behavior of Euler integrators

v

x

forward

v

x

backward

Behavior of Euler integrators

y

t

forward

y

t

backward

Towards higher order

Let's try expanding around $t = (t_k + t_{k+1})/2$; call this time t_m for "midpoint"

- $\mathbf{y}(t) = \mathbf{y}(t_m) + \dot{\mathbf{y}}(t_m)(t - t_m) + \frac{1}{2}\ddot{\mathbf{y}}(t_m)(t - t_m)^2 + O((t - t_m)^3)$

evaluate at t_k and t_{k+1} to compute the step increment

- $\mathbf{y}(t_{k+1}) - \mathbf{y}(t_k) = h\dot{\mathbf{y}}(t_m) + O(h^3)$ (try it yourself to see the canceling h^2 term)

…if only we knew $\mathbf{y}(t_m)$! But we can use Forward Euler to estimate it

- $\mathbf{y}(t_m) = \mathbf{y}(t_k) + \frac{h}{2}\dot{\mathbf{y}}(t_k) + O(h^2)$, so let $\mathbf{y}_m = \mathbf{y}_k + \frac{h}{2}\mathbf{f}(\mathbf{y}_k)$
- $\mathbf{f}(\mathbf{y} + O(h^2)) = \mathbf{f}(\mathbf{y}) + \mathbf{f}'(\mathbf{y})O(h^2) + O(h^2) = \mathbf{f}(\mathbf{y}) + O(h^2)$, so $\mathbf{f}(\mathbf{y}_m) = \dot{\mathbf{y}}(t_m) + O(h^2)$
- then $\mathbf{y}(t_{k+1}) = \mathbf{y}(t_k) + h(\mathbf{f}(\mathbf{y}_\mathbf{m}) + O(h^2)) + O(h^3) = \mathbf{y}(t_k) + h\mathbf{f}(\mathbf{y}_\mathbf{m}) + O(h^3)$
- so let $\mathbf{y}_{k+1} = \mathbf{y}_k + h\mathbf{f}(\mathbf{y}_m)$ and $\mathbf{y}_{k+1}$ is a second-order estimate of $\mathbf{y}(t_{k+1})$

Midpoint method

Timestep equations

$$\mathbf{y}_m = \mathbf{y}_k + \frac{h}{2}\mathbf{f}(\mathbf{y}_k)$$

$$\mathbf{y}_{k+1} = \mathbf{y}_k + h\mathbf{f}(\mathbf{y}_m)$$

This is

- an explicit integrator
- a *two-step* integrator (requires two evaluations of $\mathbf{f}$)
- accurate to second order

It’s also the first in a family of higher order integrators

- Runge-Kutta methods achieve order p accuracy with at least p function evaluations
- RK4 is a popular fourth-order scheme, good for smooth problems requiring high accuracy
- animation = not-so-smooth problems requiring low accuracy, hence we rarely go past second order

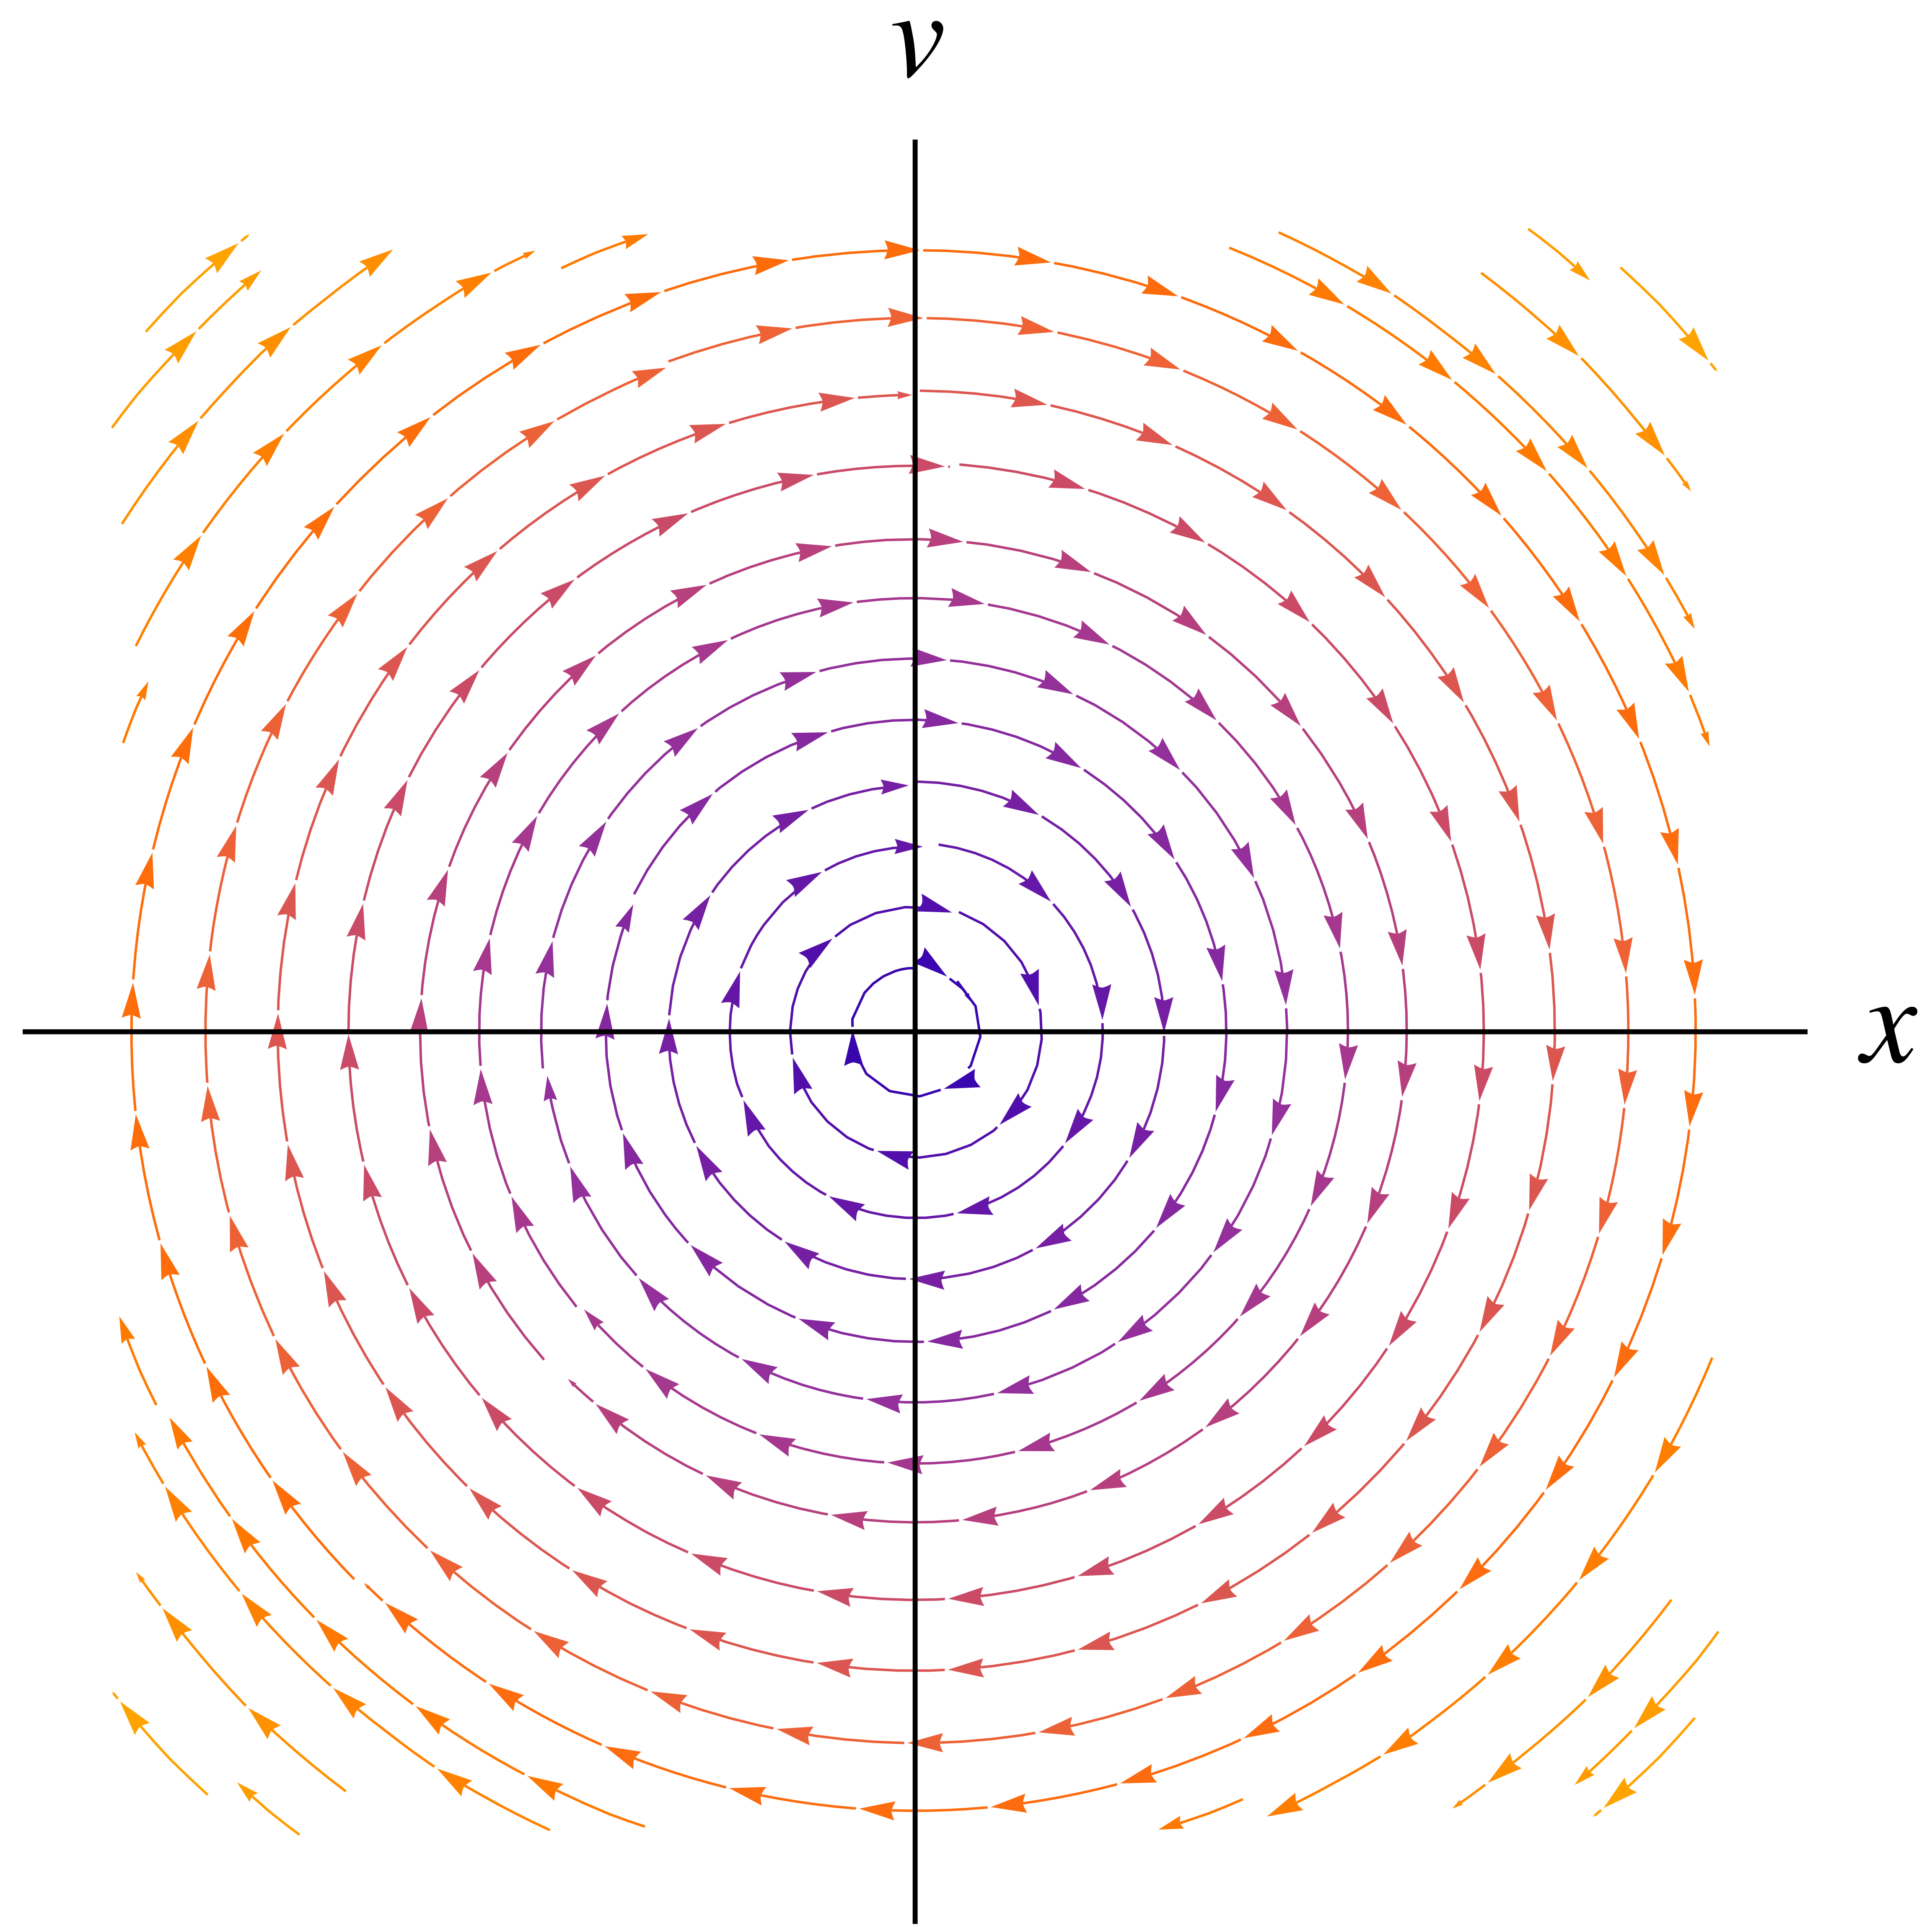

v
x

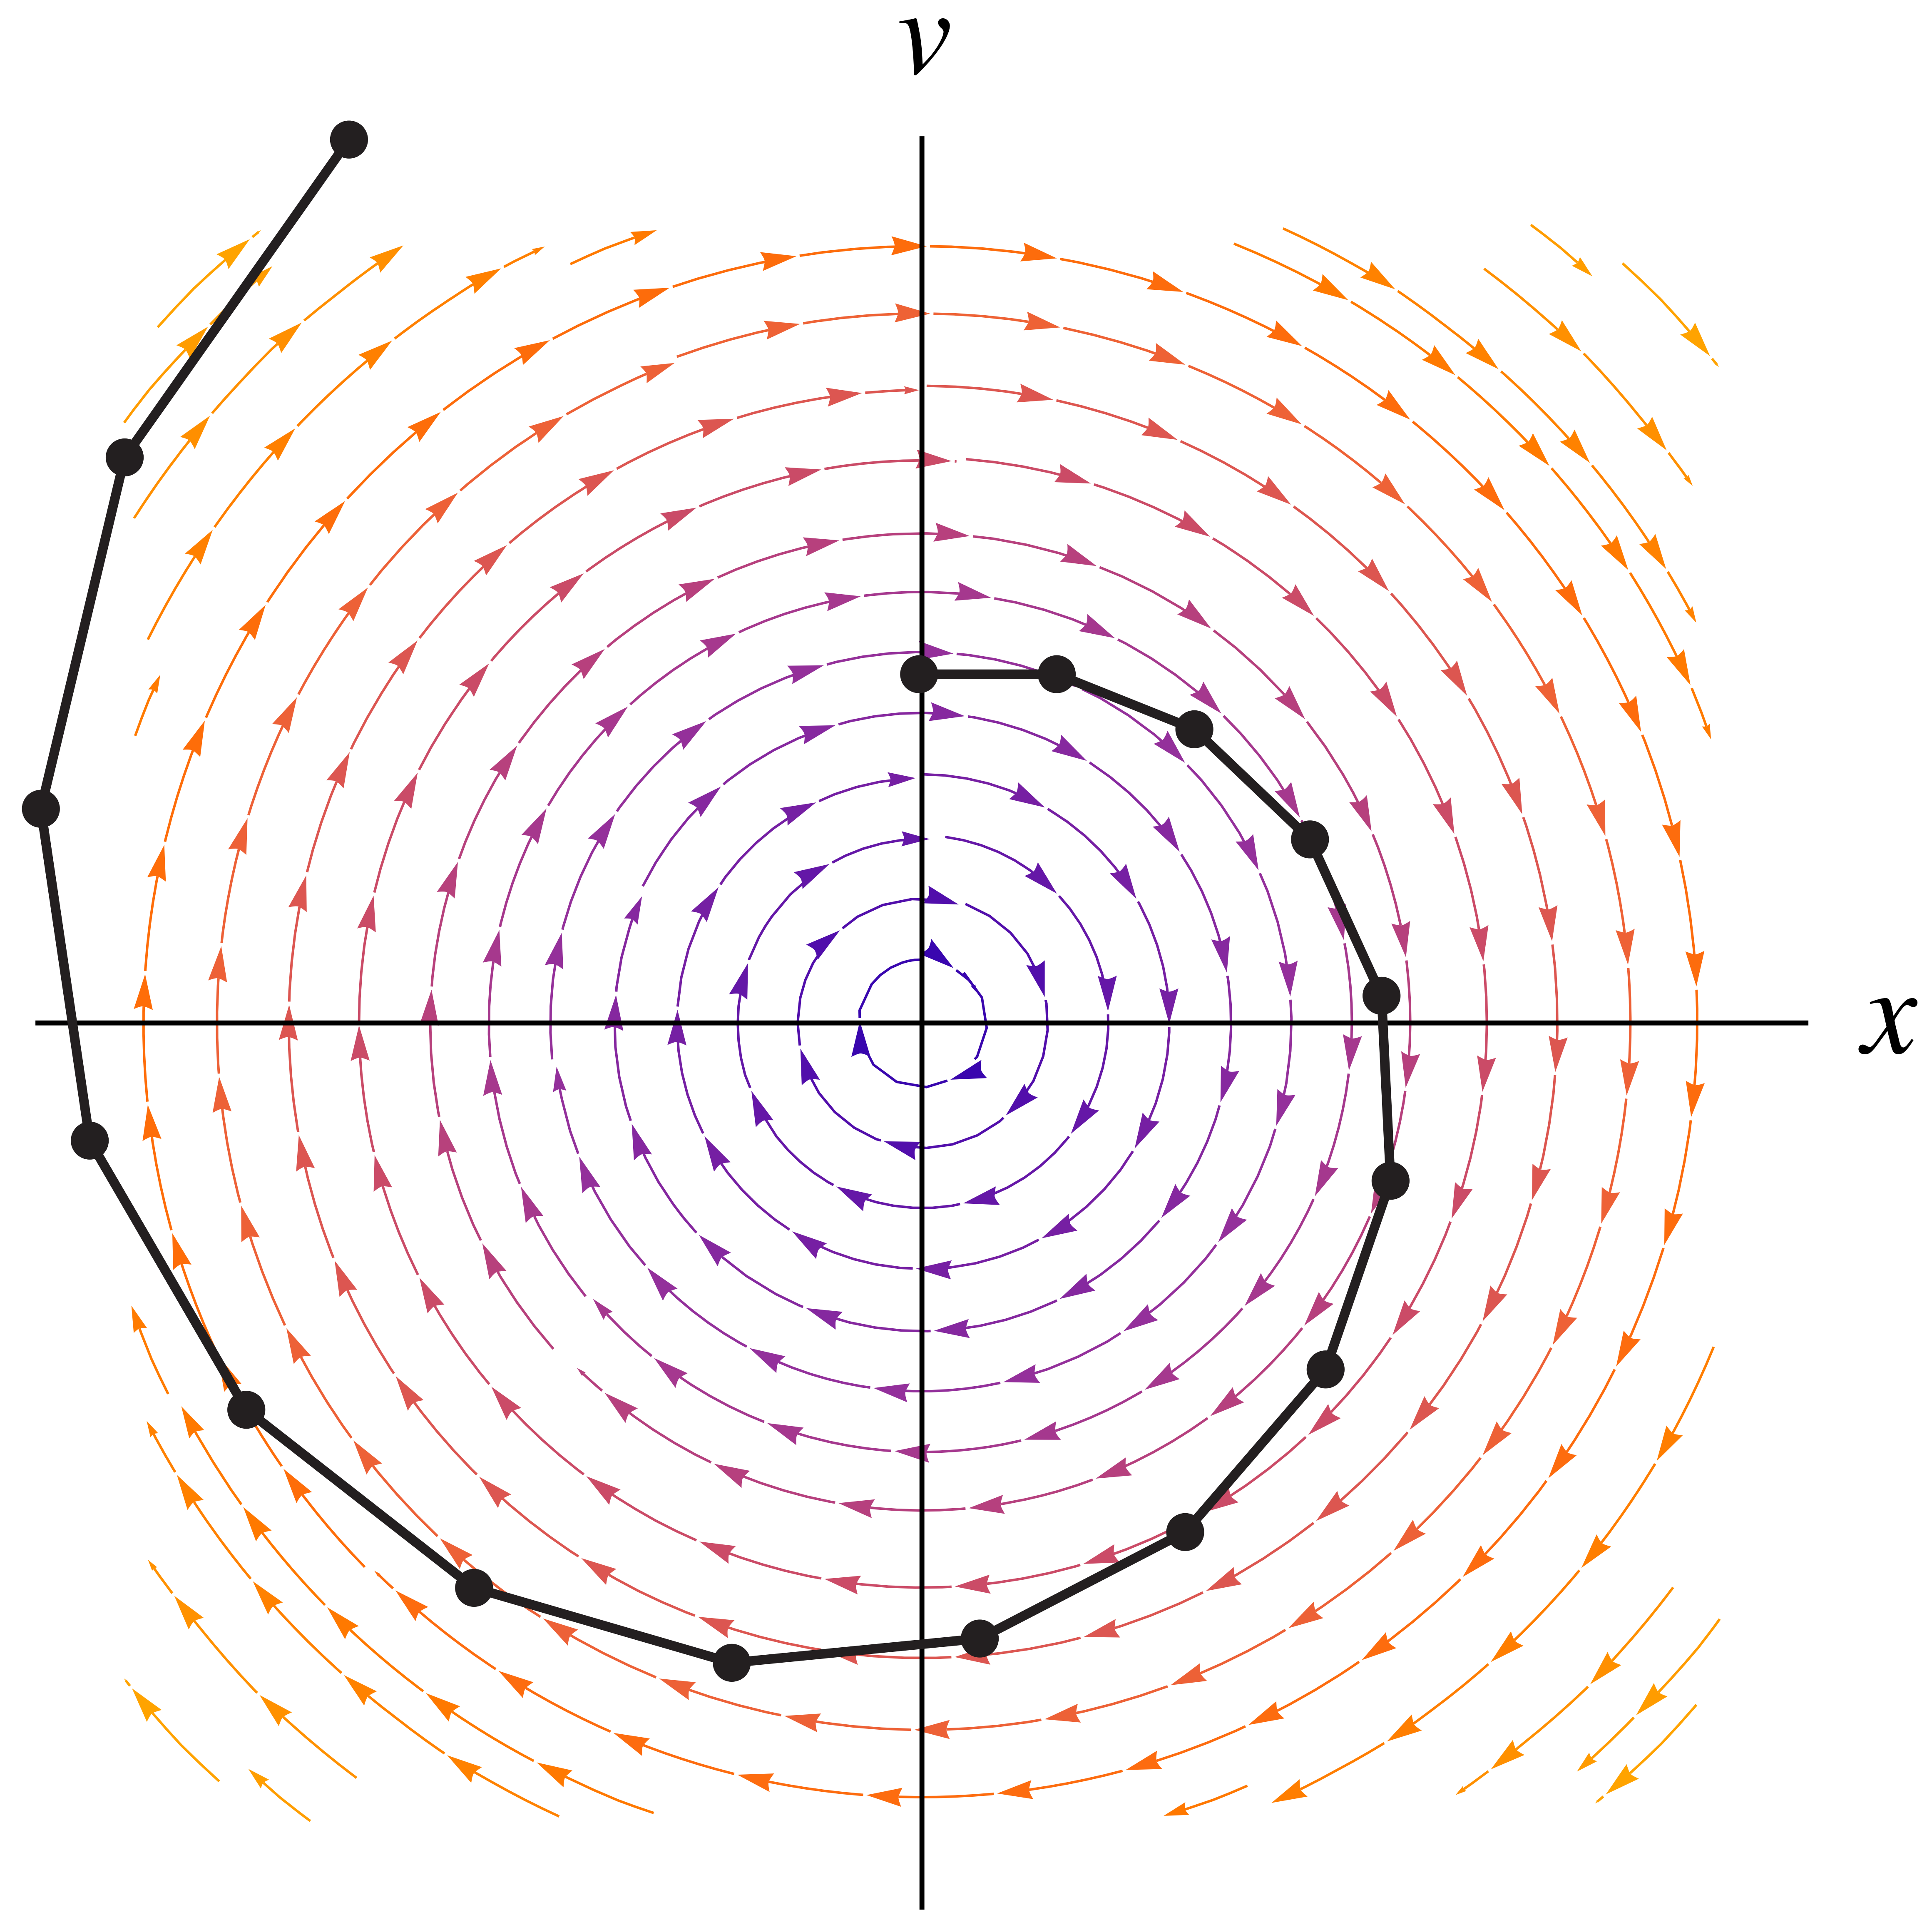

forward Euler

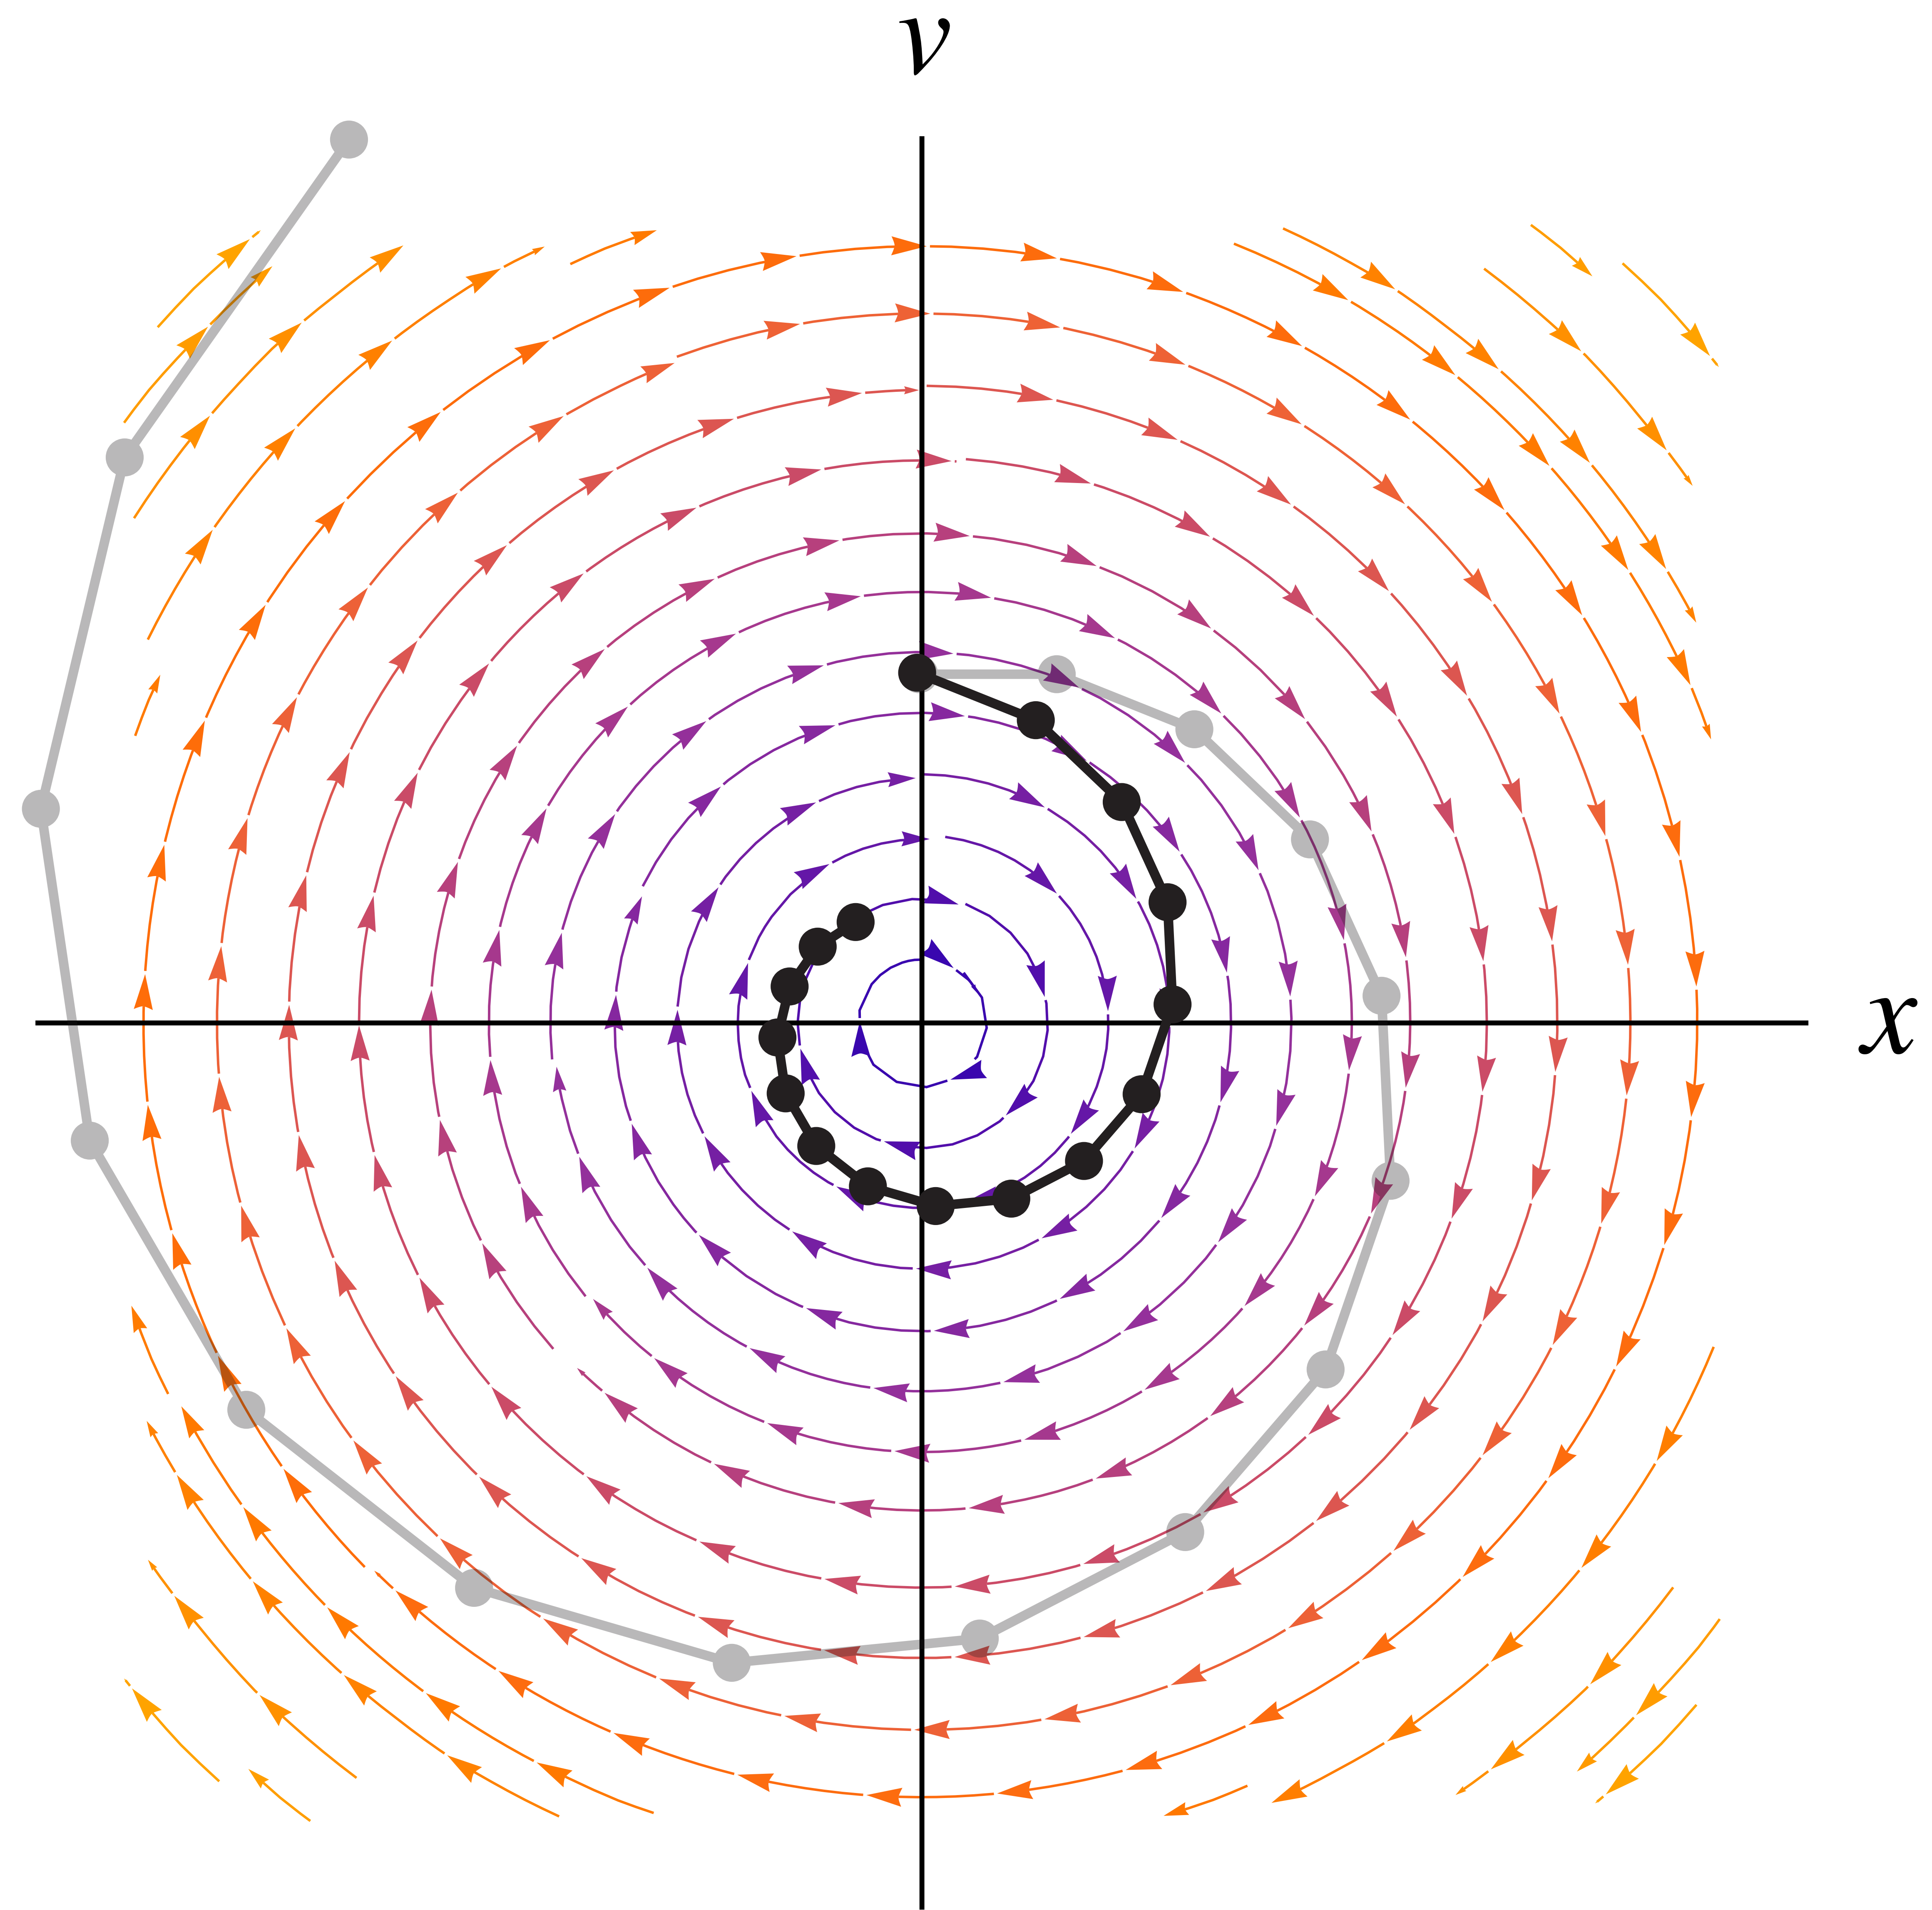

backward Euler

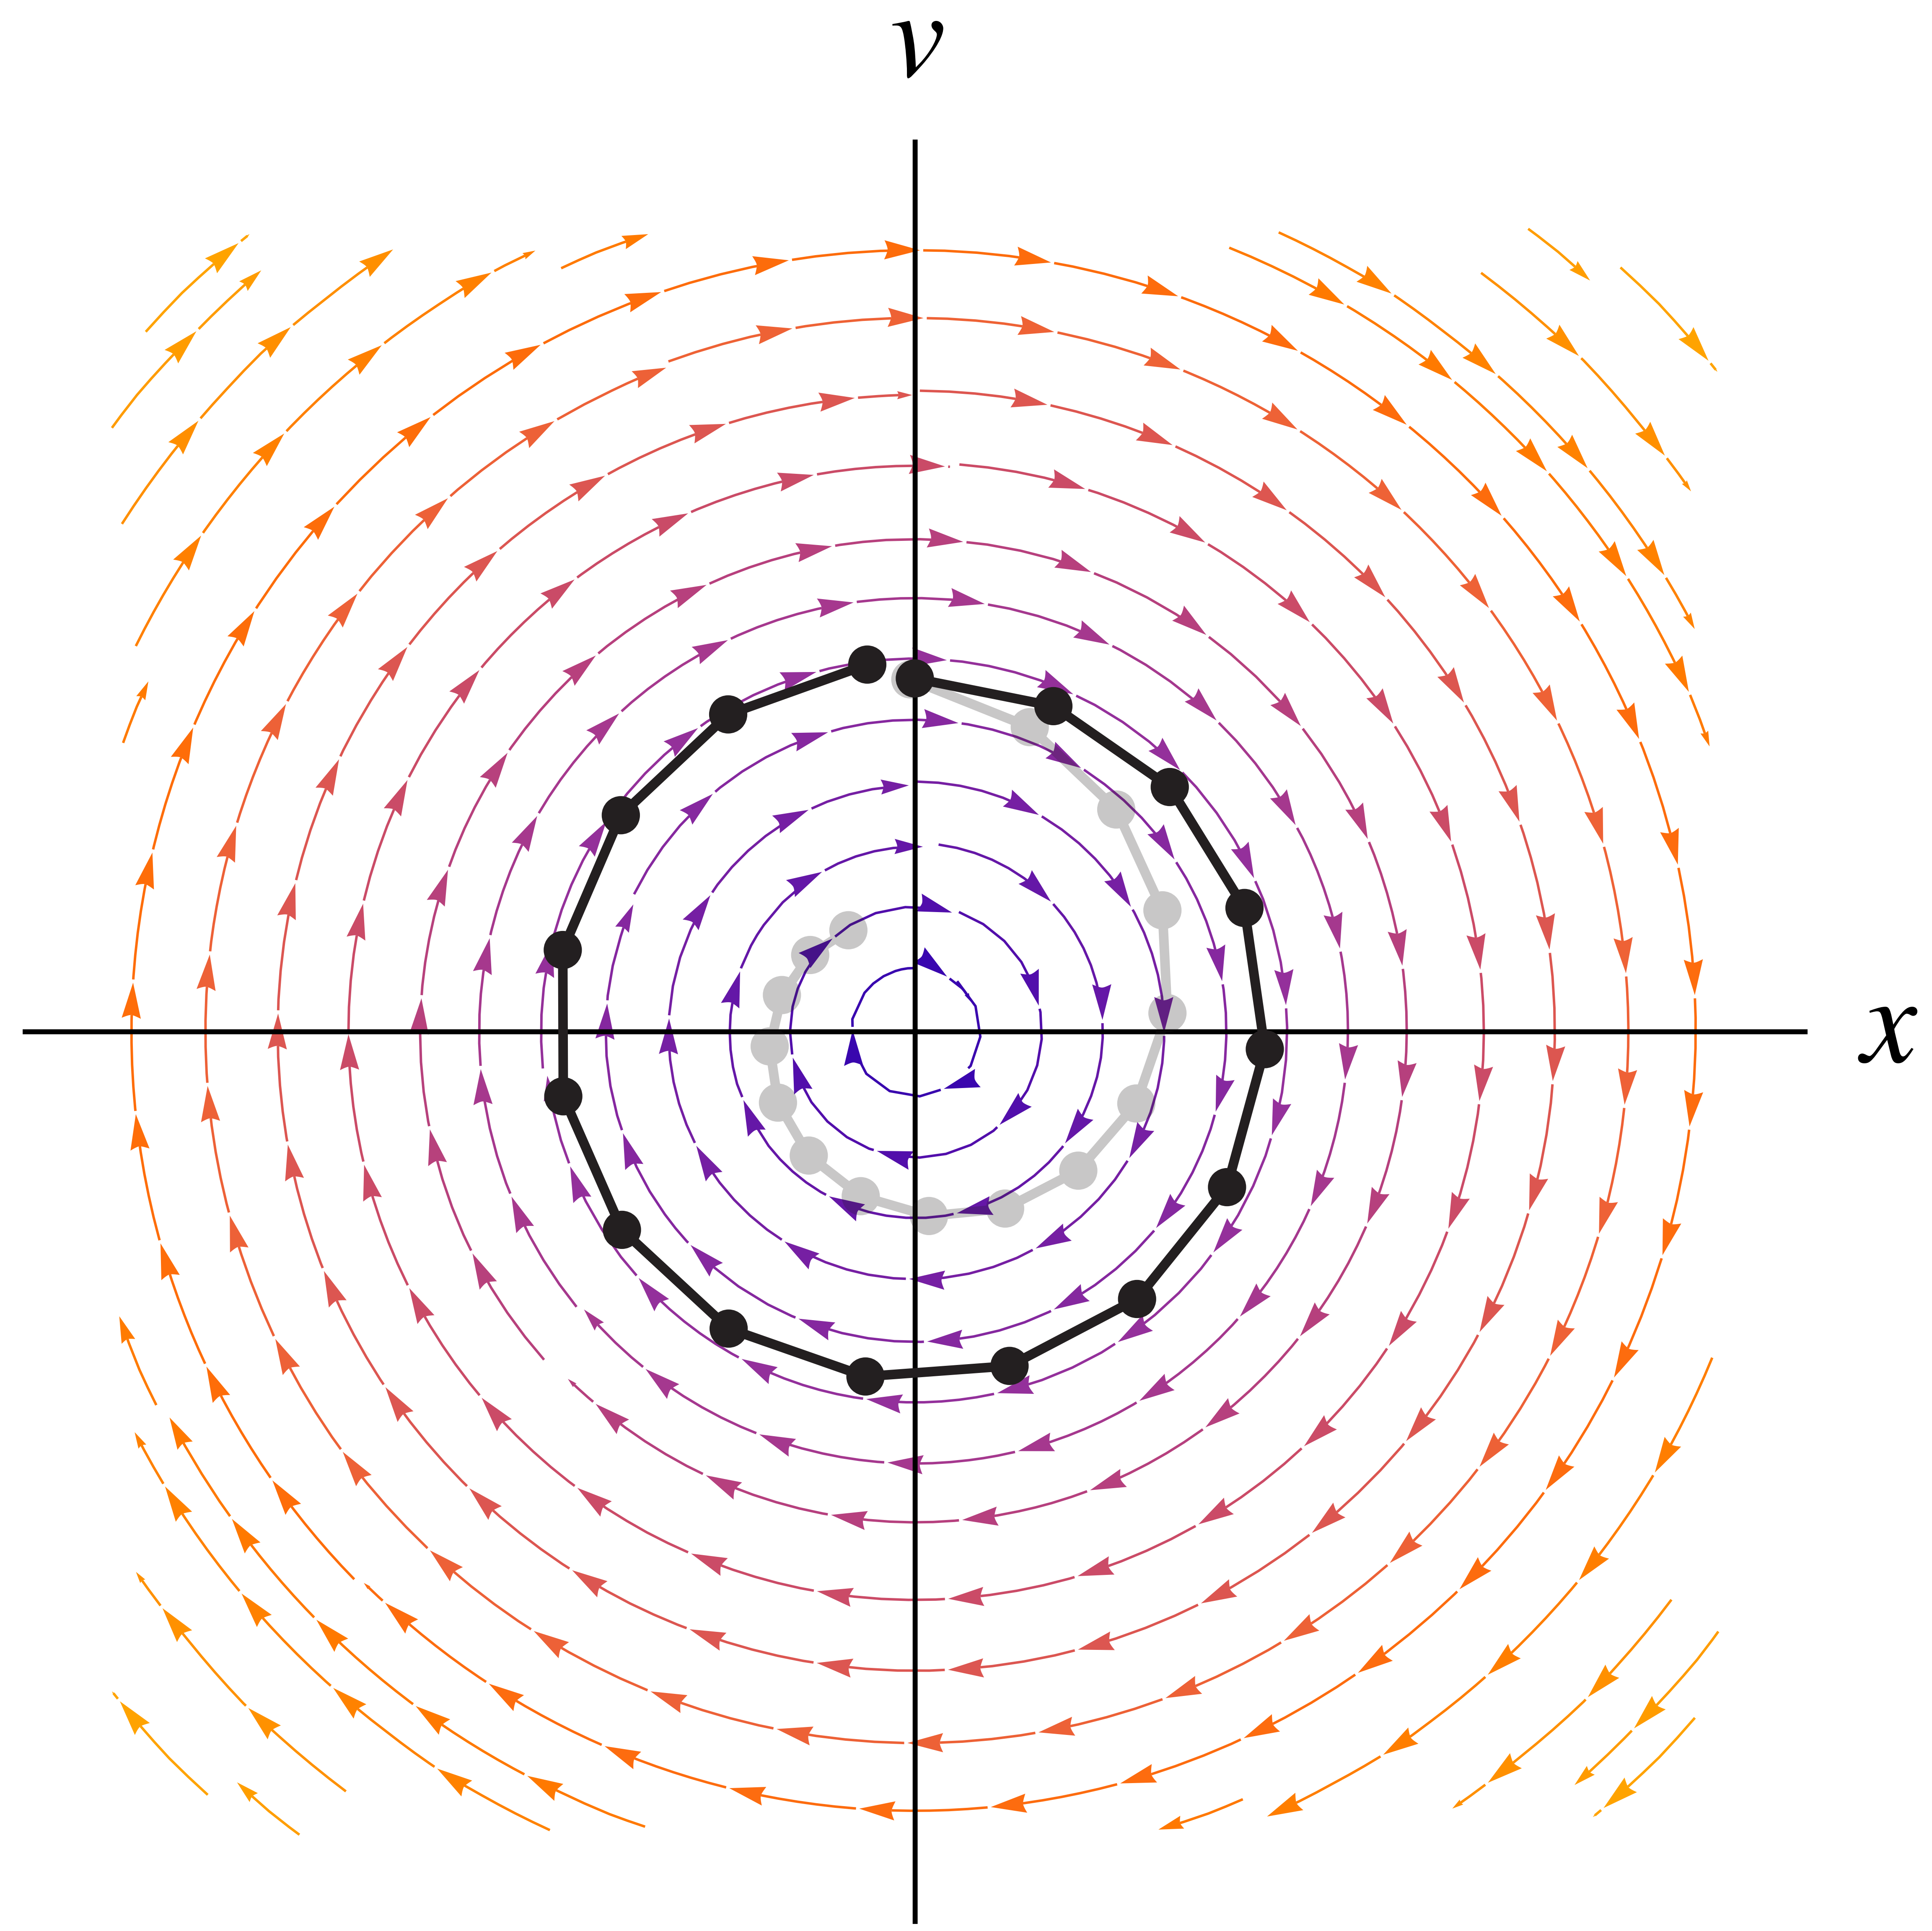

midpoint method

Demo!

accuracy of integration along circular paths

- Euler vs. midpoint

Integrators for second-order systems

Many useful systems have the form $\ddot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t))$

- note this equation skips over $\dot{\mathbf{x}}$; acceleration does not depend on velocity, only position.

Look at what the second step of the midpoint method does

- $\mathbf{y}_{k+1} = \mathbf{y}_k + h\mathbf{f}(\mathbf{y}_m)$ translates to (naming $\mathbf{y}_m$ as $\mathbf{y}_{k+0.5}$)

$$\mathbf{x}_{k+1} = \mathbf{x}_k + h\mathbf{v}_{k+0.5}$$
$$\mathbf{v}_{k+1} = \mathbf{v}_k + \mathbf{f}(\mathbf{x}_{k+0.5})$$

← updating $\mathbf{x}$ only requires $\mathbf{v}_{k+0.5}$, and updating $\mathbf{v}$ only requires $\mathbf{x}_{k+0.5}$

- if we stagger the grids then we can have these values already!

$$\mathbf{x}_{k+1} = \mathbf{x}_k + h\mathbf{v}_{k+0.5}$$
$$\mathbf{v}_{k+1.5} = \mathbf{v}_{k+0.5} + h\mathbf{f}(\mathbf{x}_{k+1})$$

- this is an explicit method, and it's second order accurate for both position and velocity
- known as the Leapfrog integrator — elegant but prohibits velocity dependent forces

Symplectic Euler's method (aka. semi-implicit)

Leapfrog is nice but doesn't work for $\ddot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{v})$

- practical problem: can't evaluate $\mathbf{f}$ without knowing $\mathbf{x}$ and $\mathbf{v}$ at the same time
- a practical solution: give up the interleaved steps but keep the timestep equations

$$\begin{aligned} \mathbf{x}_{k+1} &= \mathbf{x}_k + h\mathbf{v}_k \\ \mathbf{v}_{k+1} &= \mathbf{v}_k + h\mathbf{f}(\mathbf{x}_{k+1}) \end{aligned}$$

this looks just like Forward Euler except for the last +1

- or: use the position update from Forward Euler and the velocity update from Backward Euler
- this integrator shares a very nice property with Leapfrog: each timestep preserves area in the $(\mathbf{x}, \mathbf{v})$ picture (really in position–momentum space)

$$\begin{bmatrix} \mathbf{x}_{k+1} \\ \mathbf{v}_{k+1} \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & h \\ -h & 1 - h^2 \end{bmatrix}}_{\det = 1} \begin{bmatrix} \mathbf{x}_k \\ \mathbf{v}_k \end{bmatrix}$$

- this property holds for any Hamiltonian (roughly, energy conserving) system

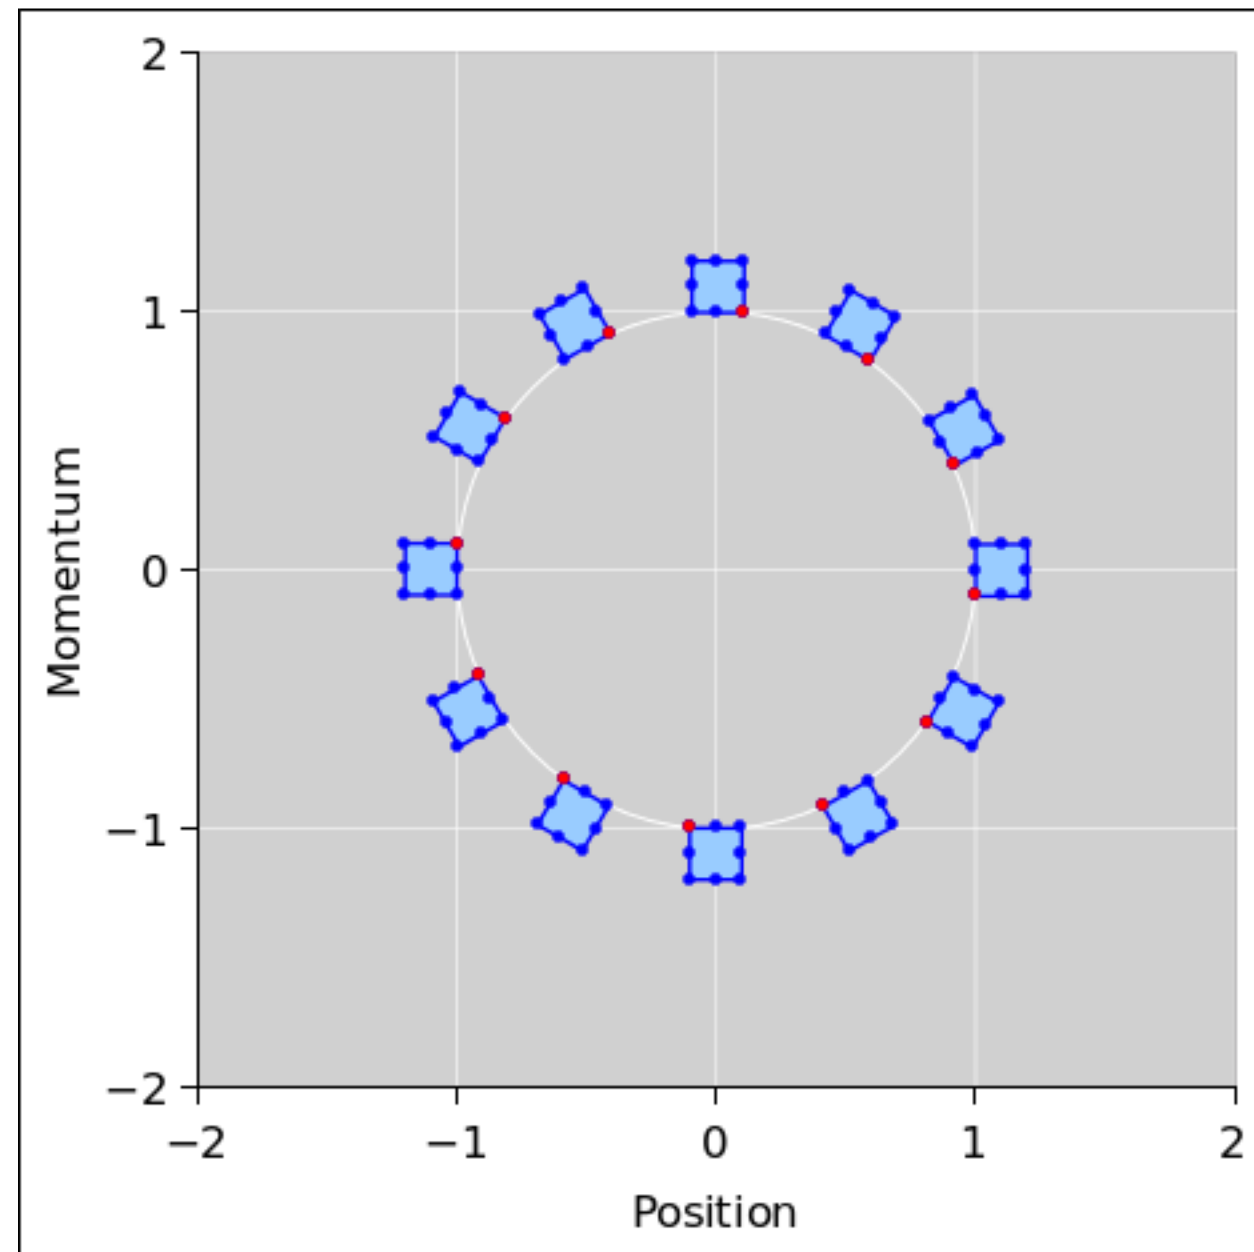

exact

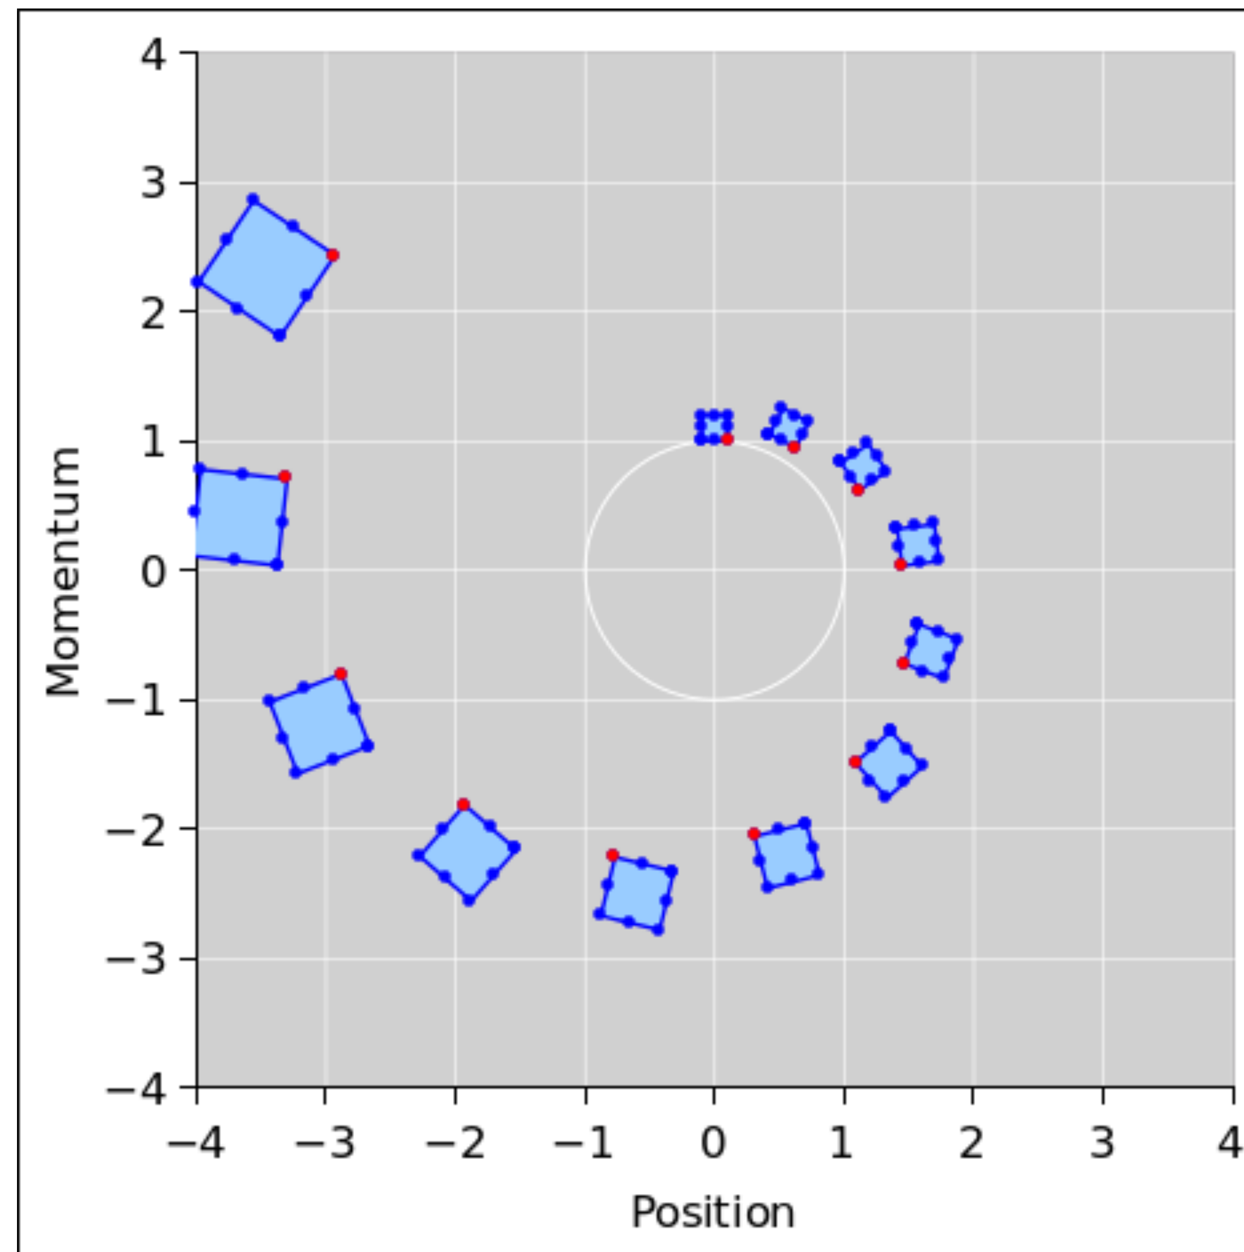

forward Euler

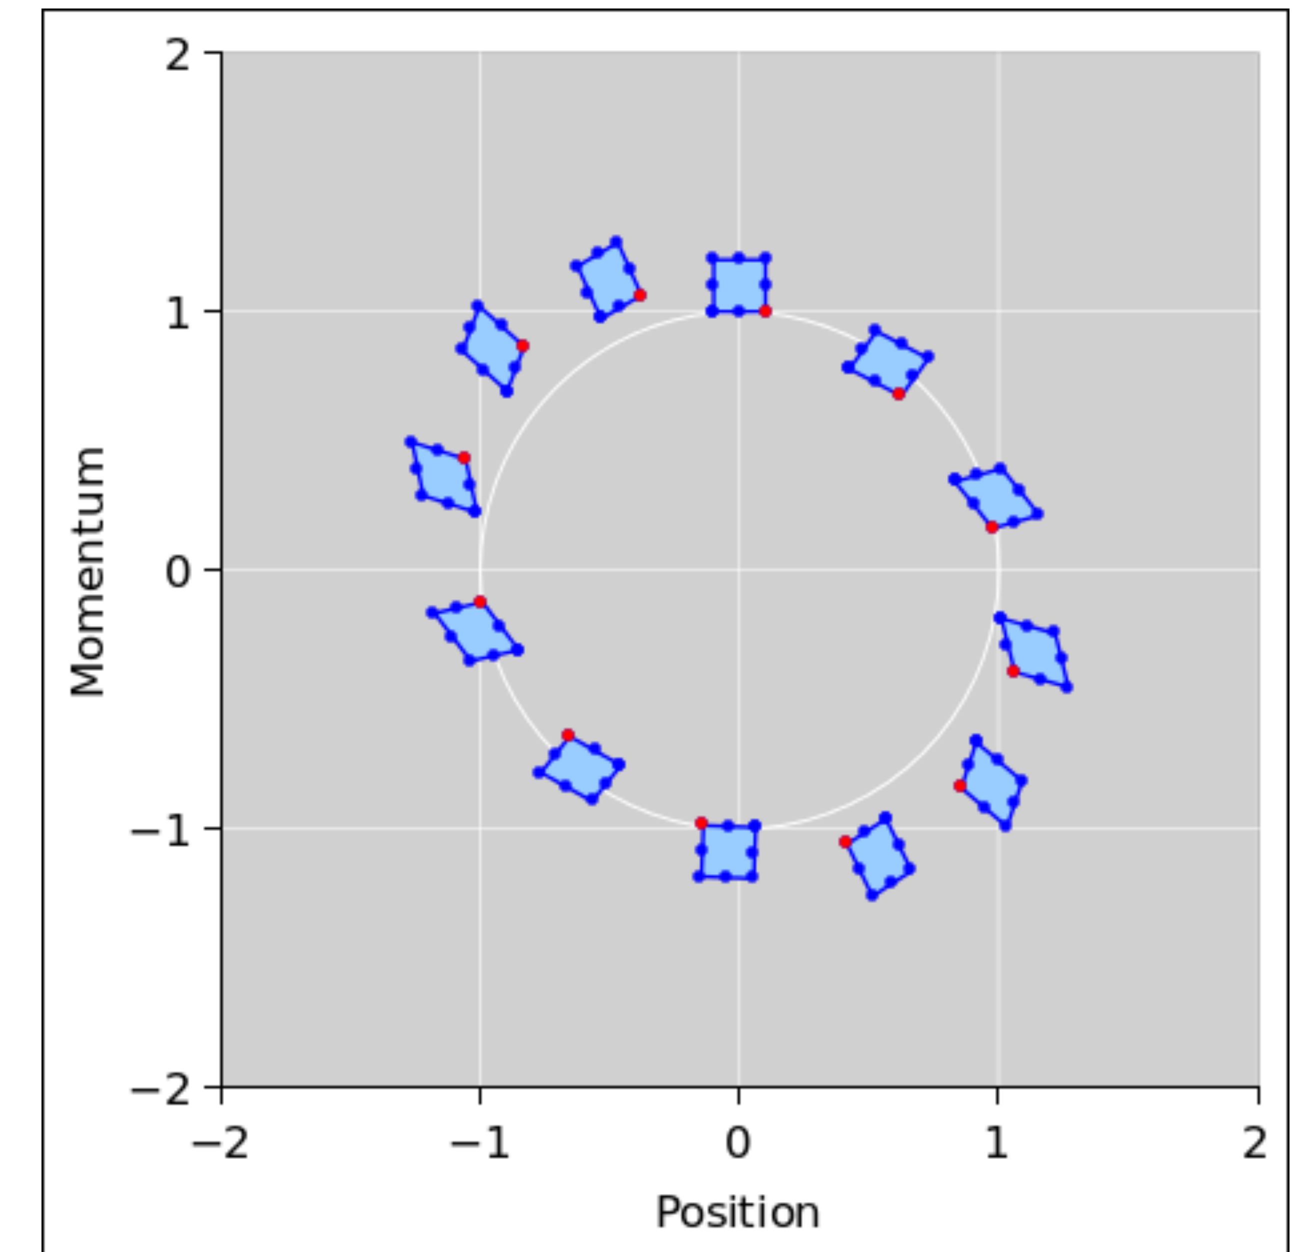

symplectic Euler

https://www.av8n.com/physics/symplectic-integrator.htm

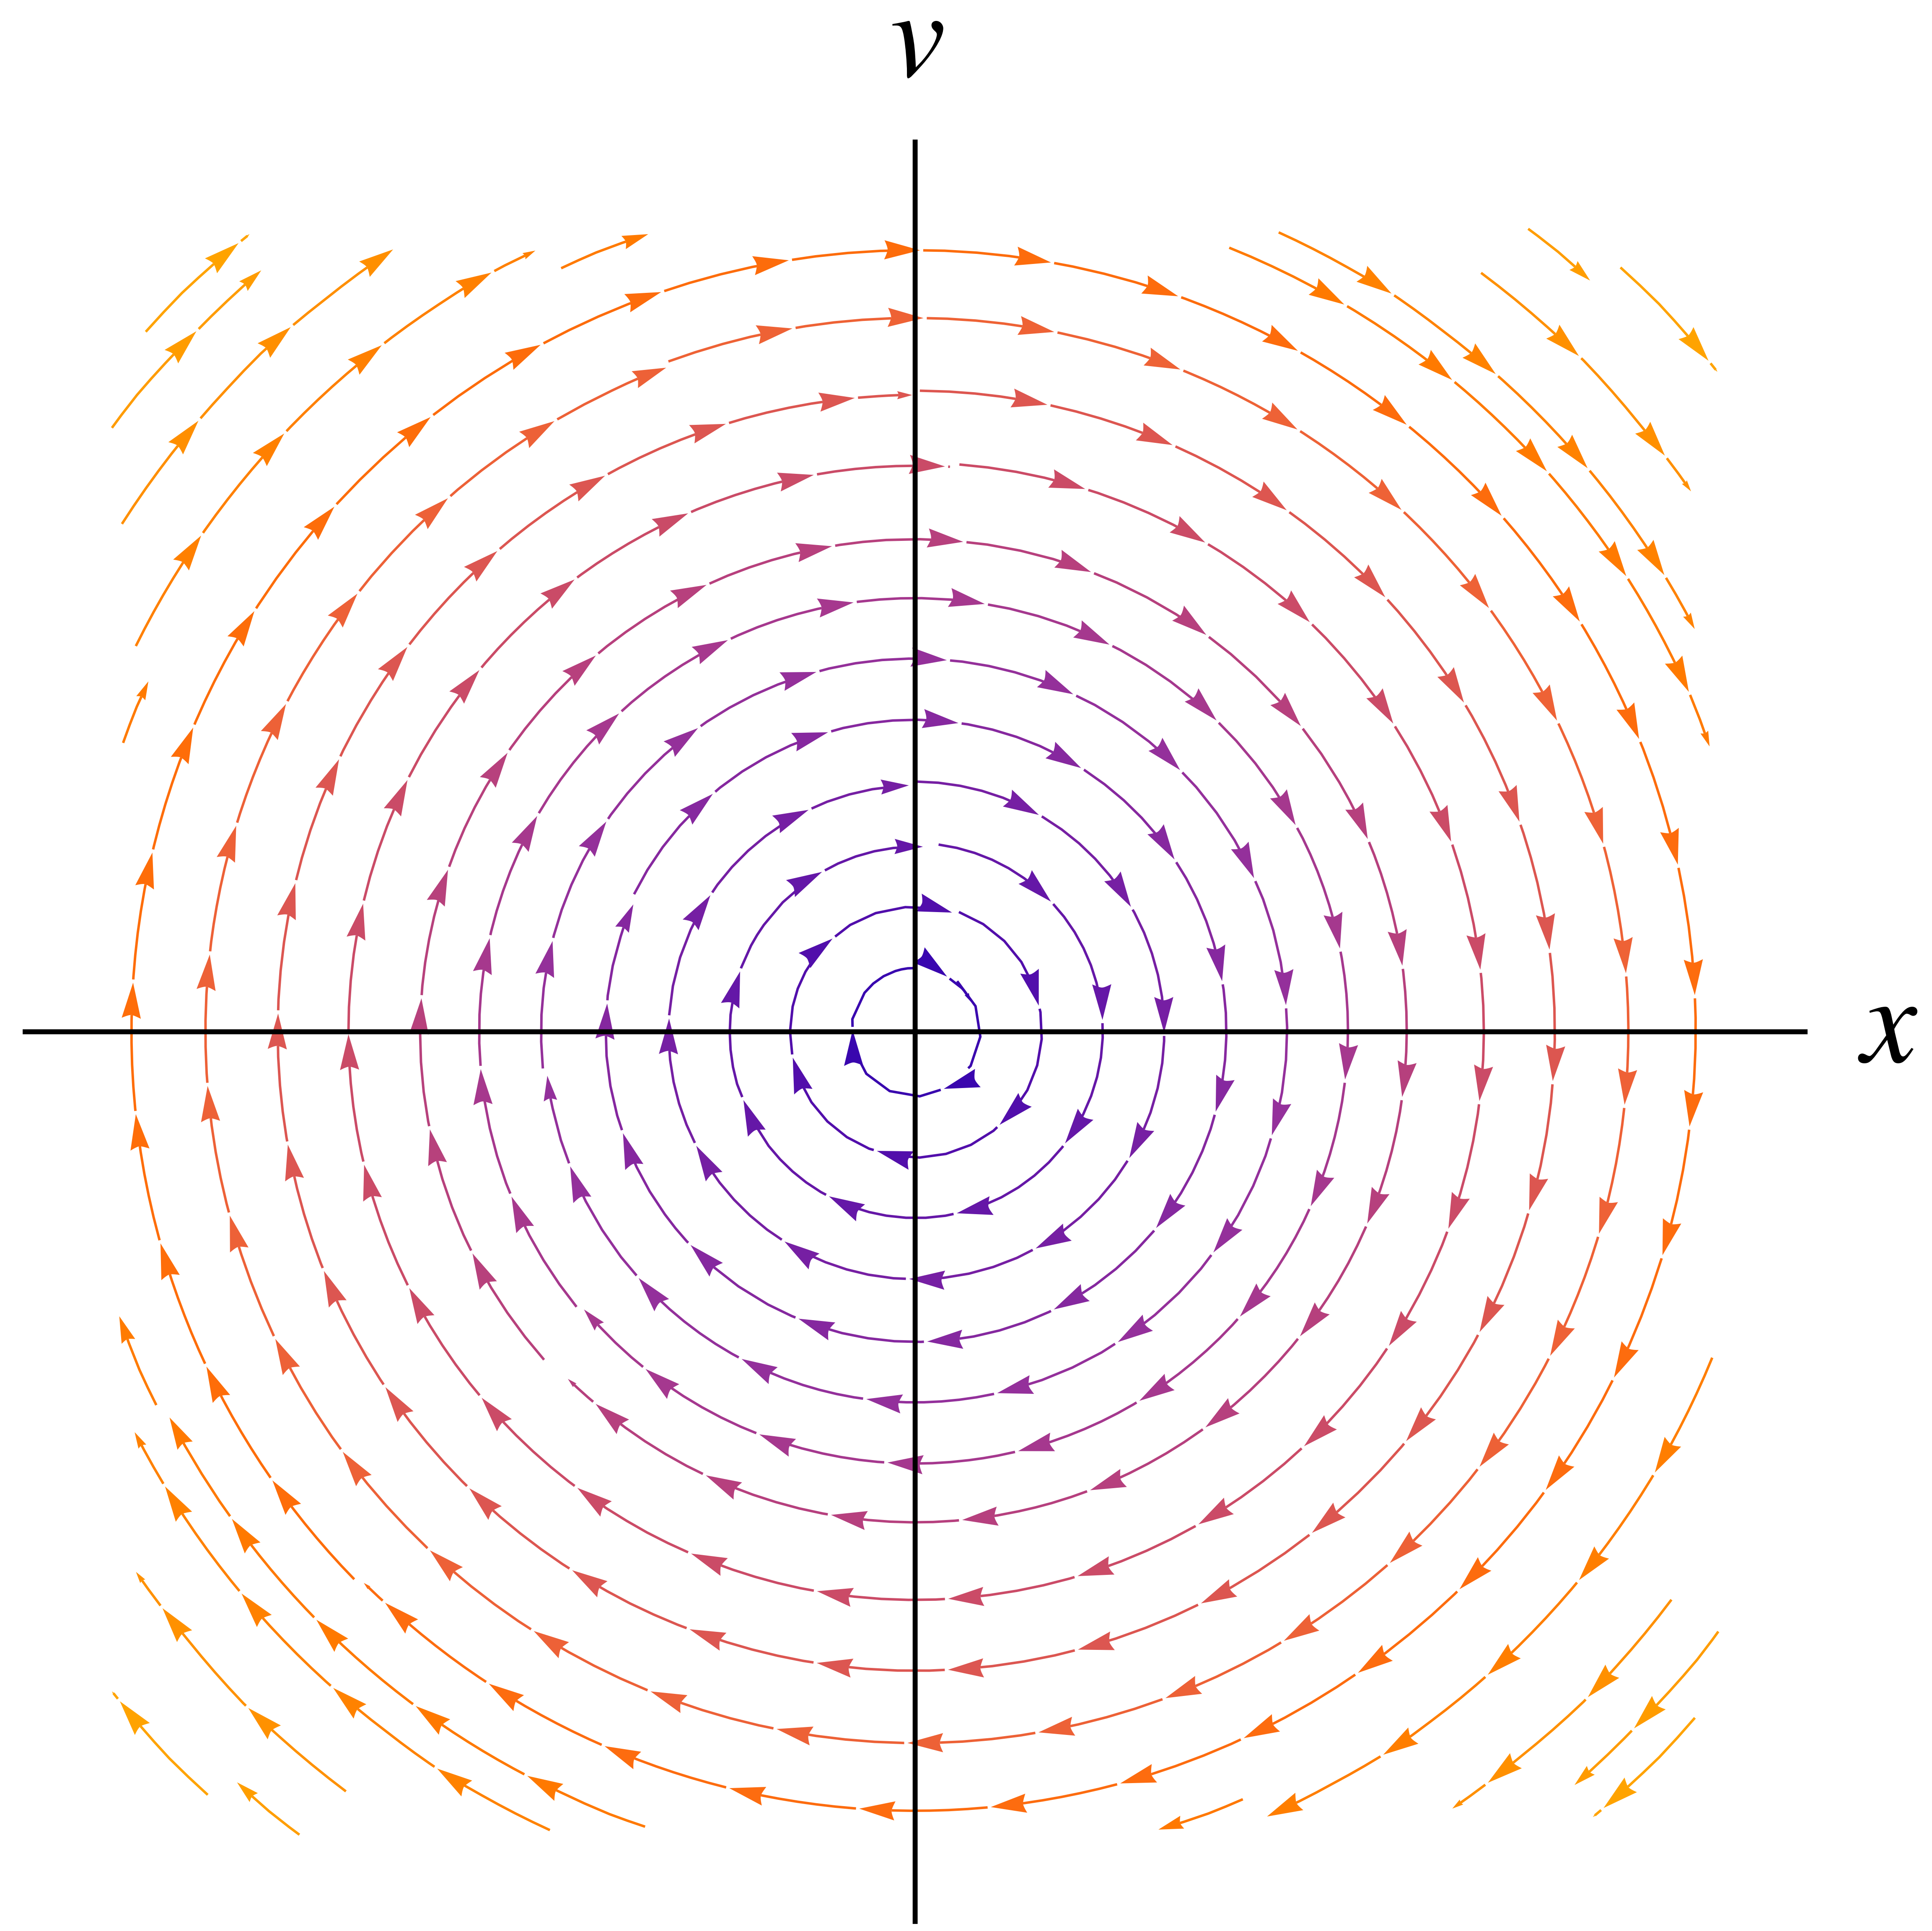

v
x

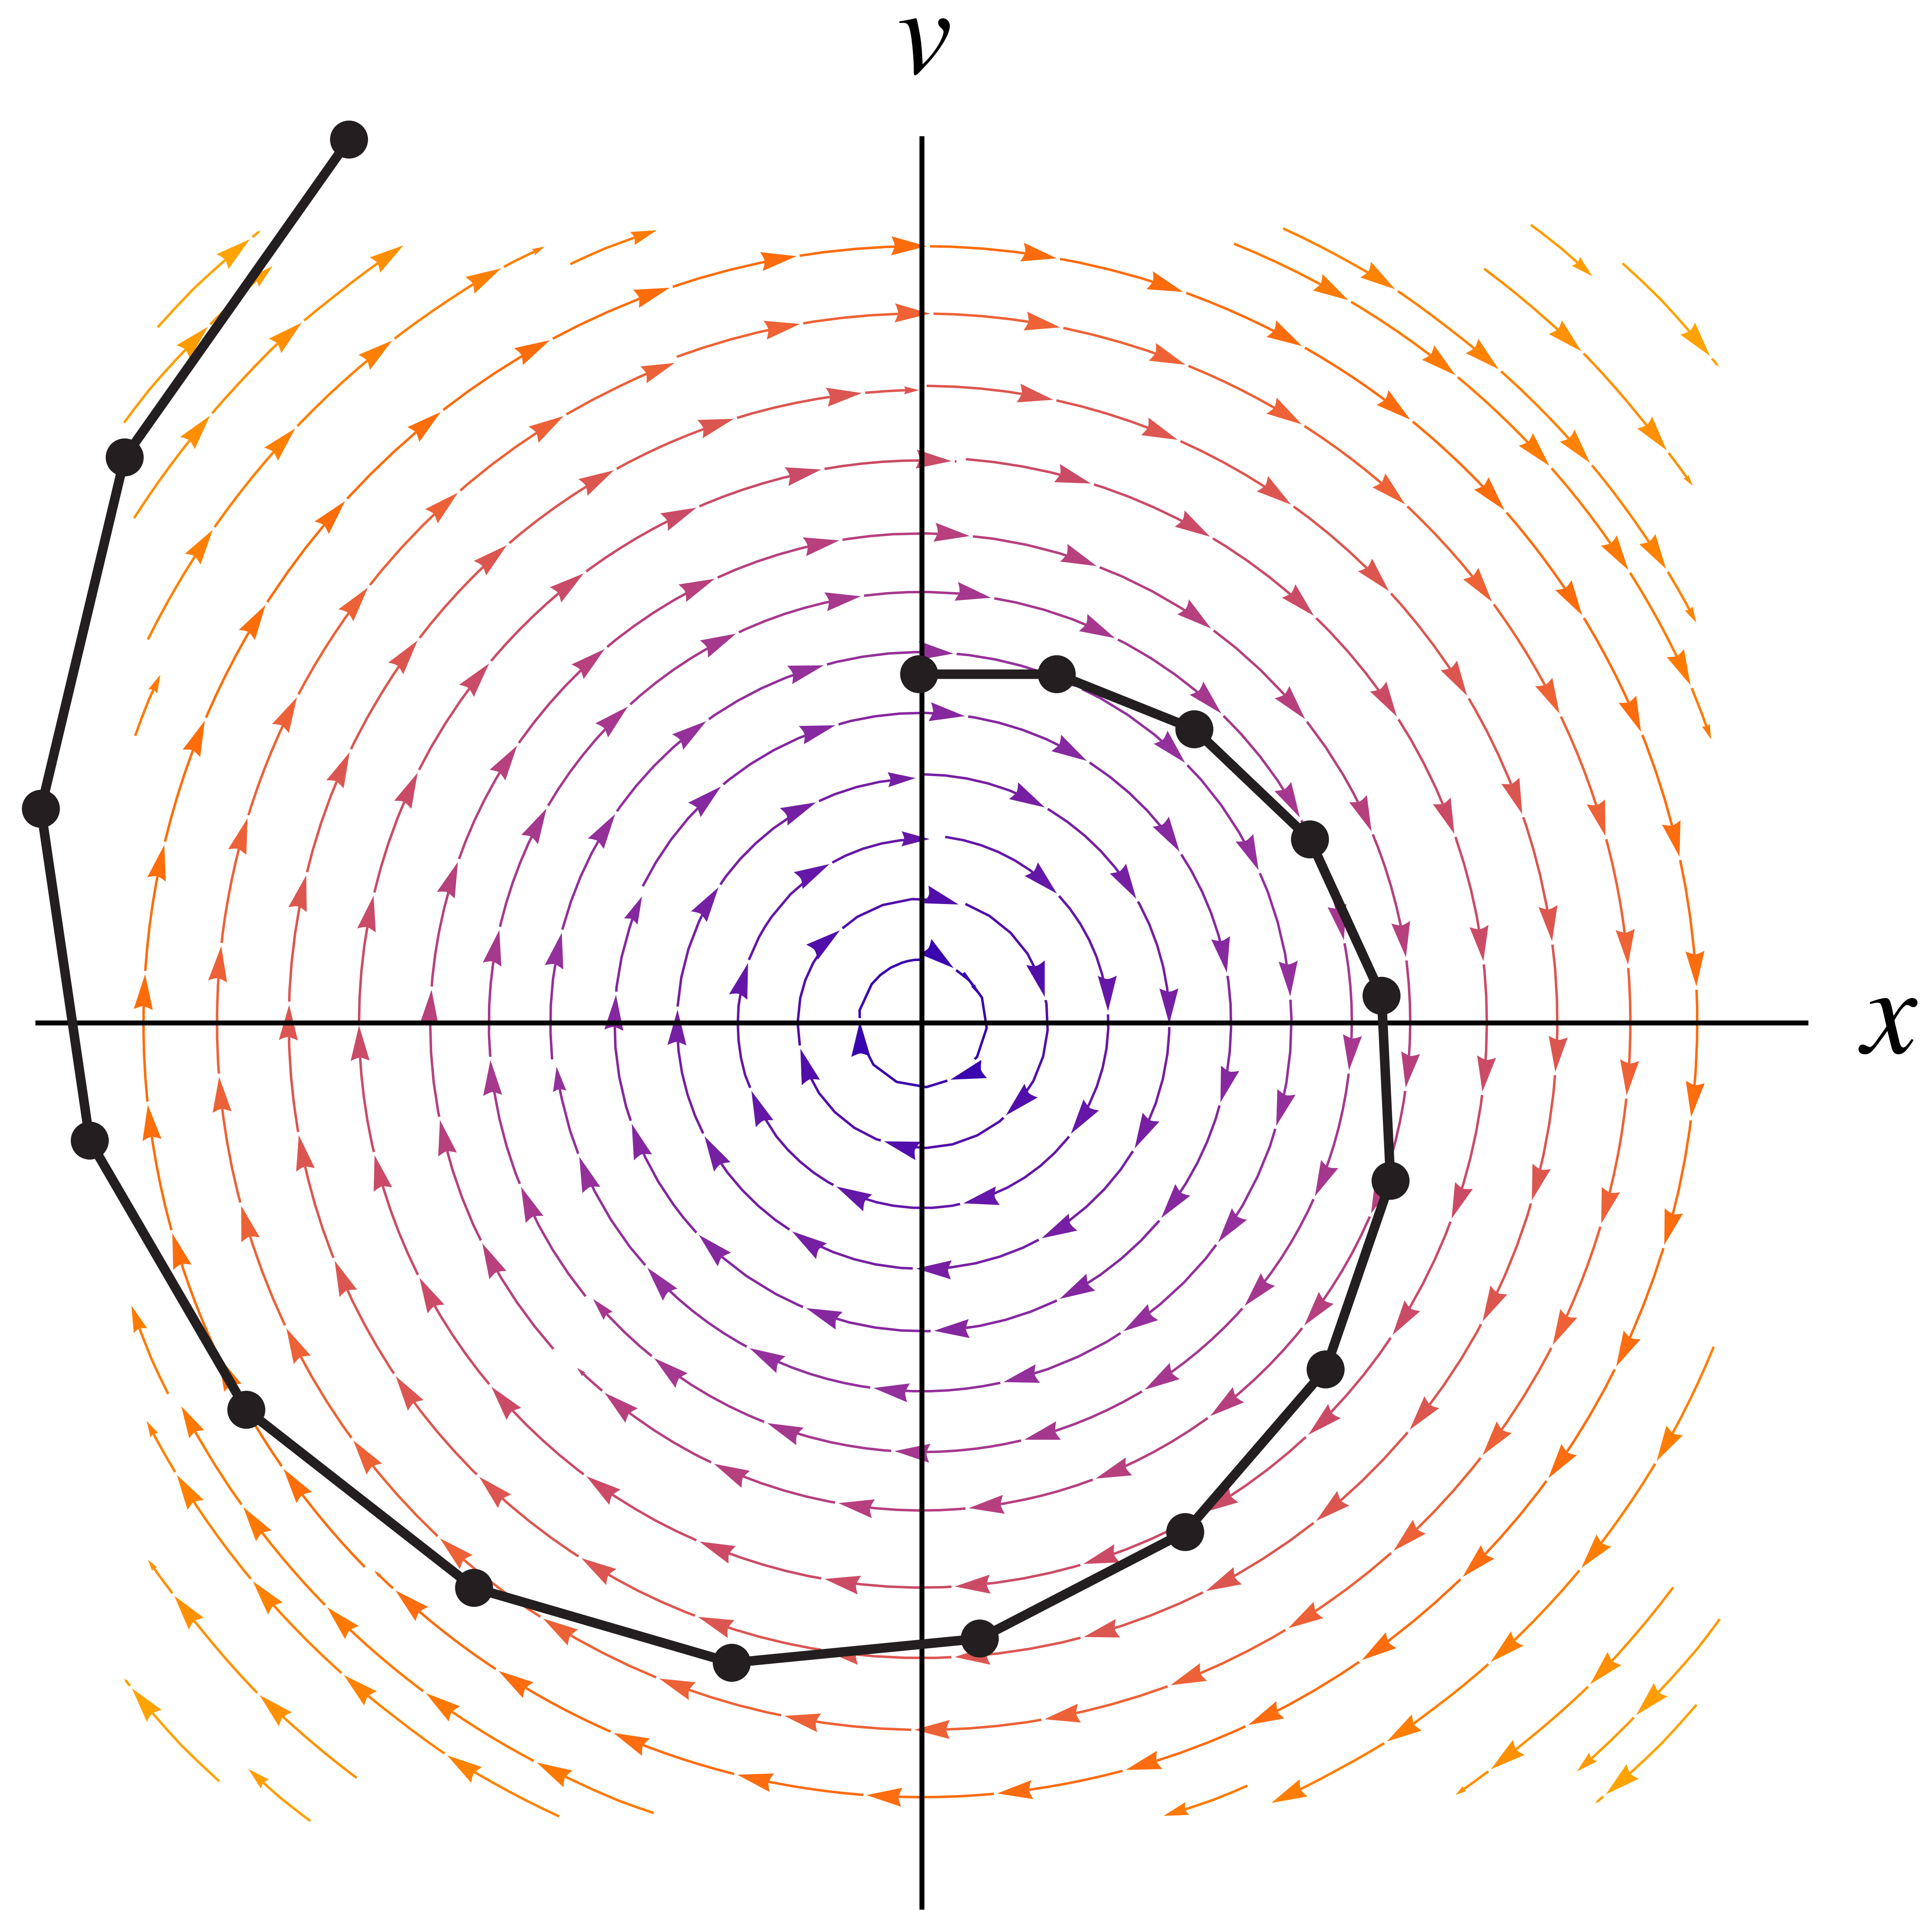

forward Euler

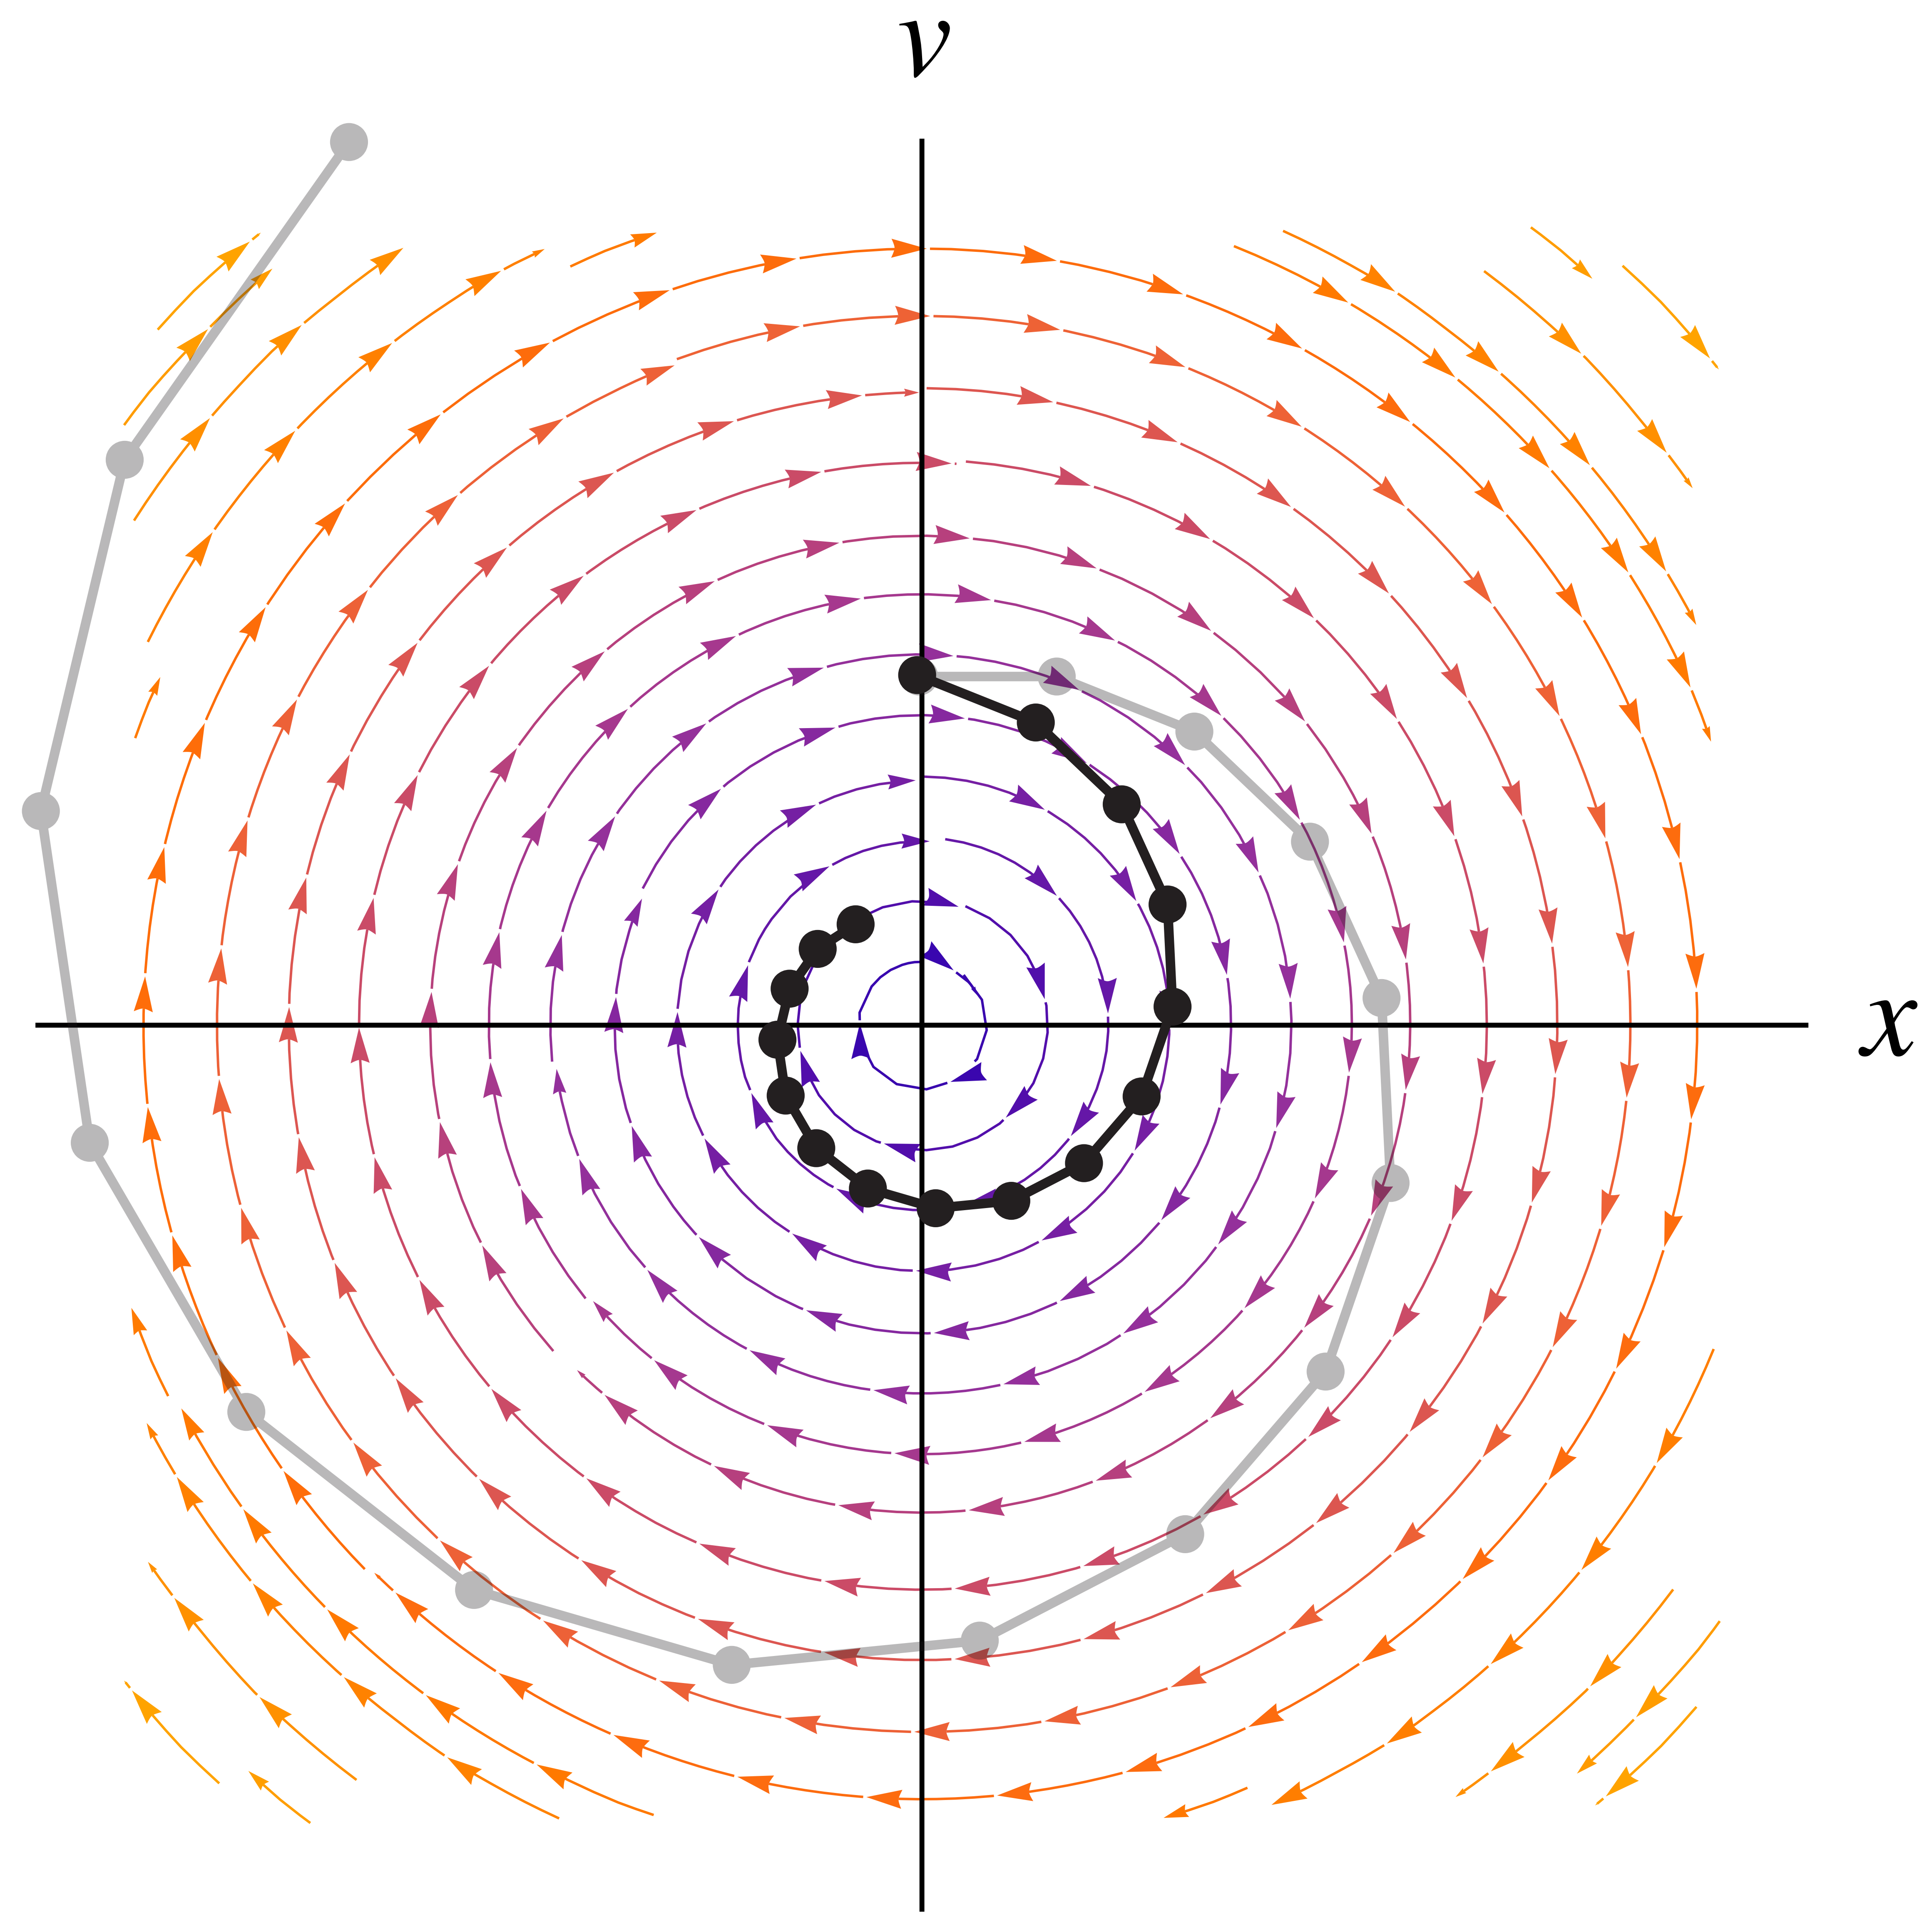

backward Euler

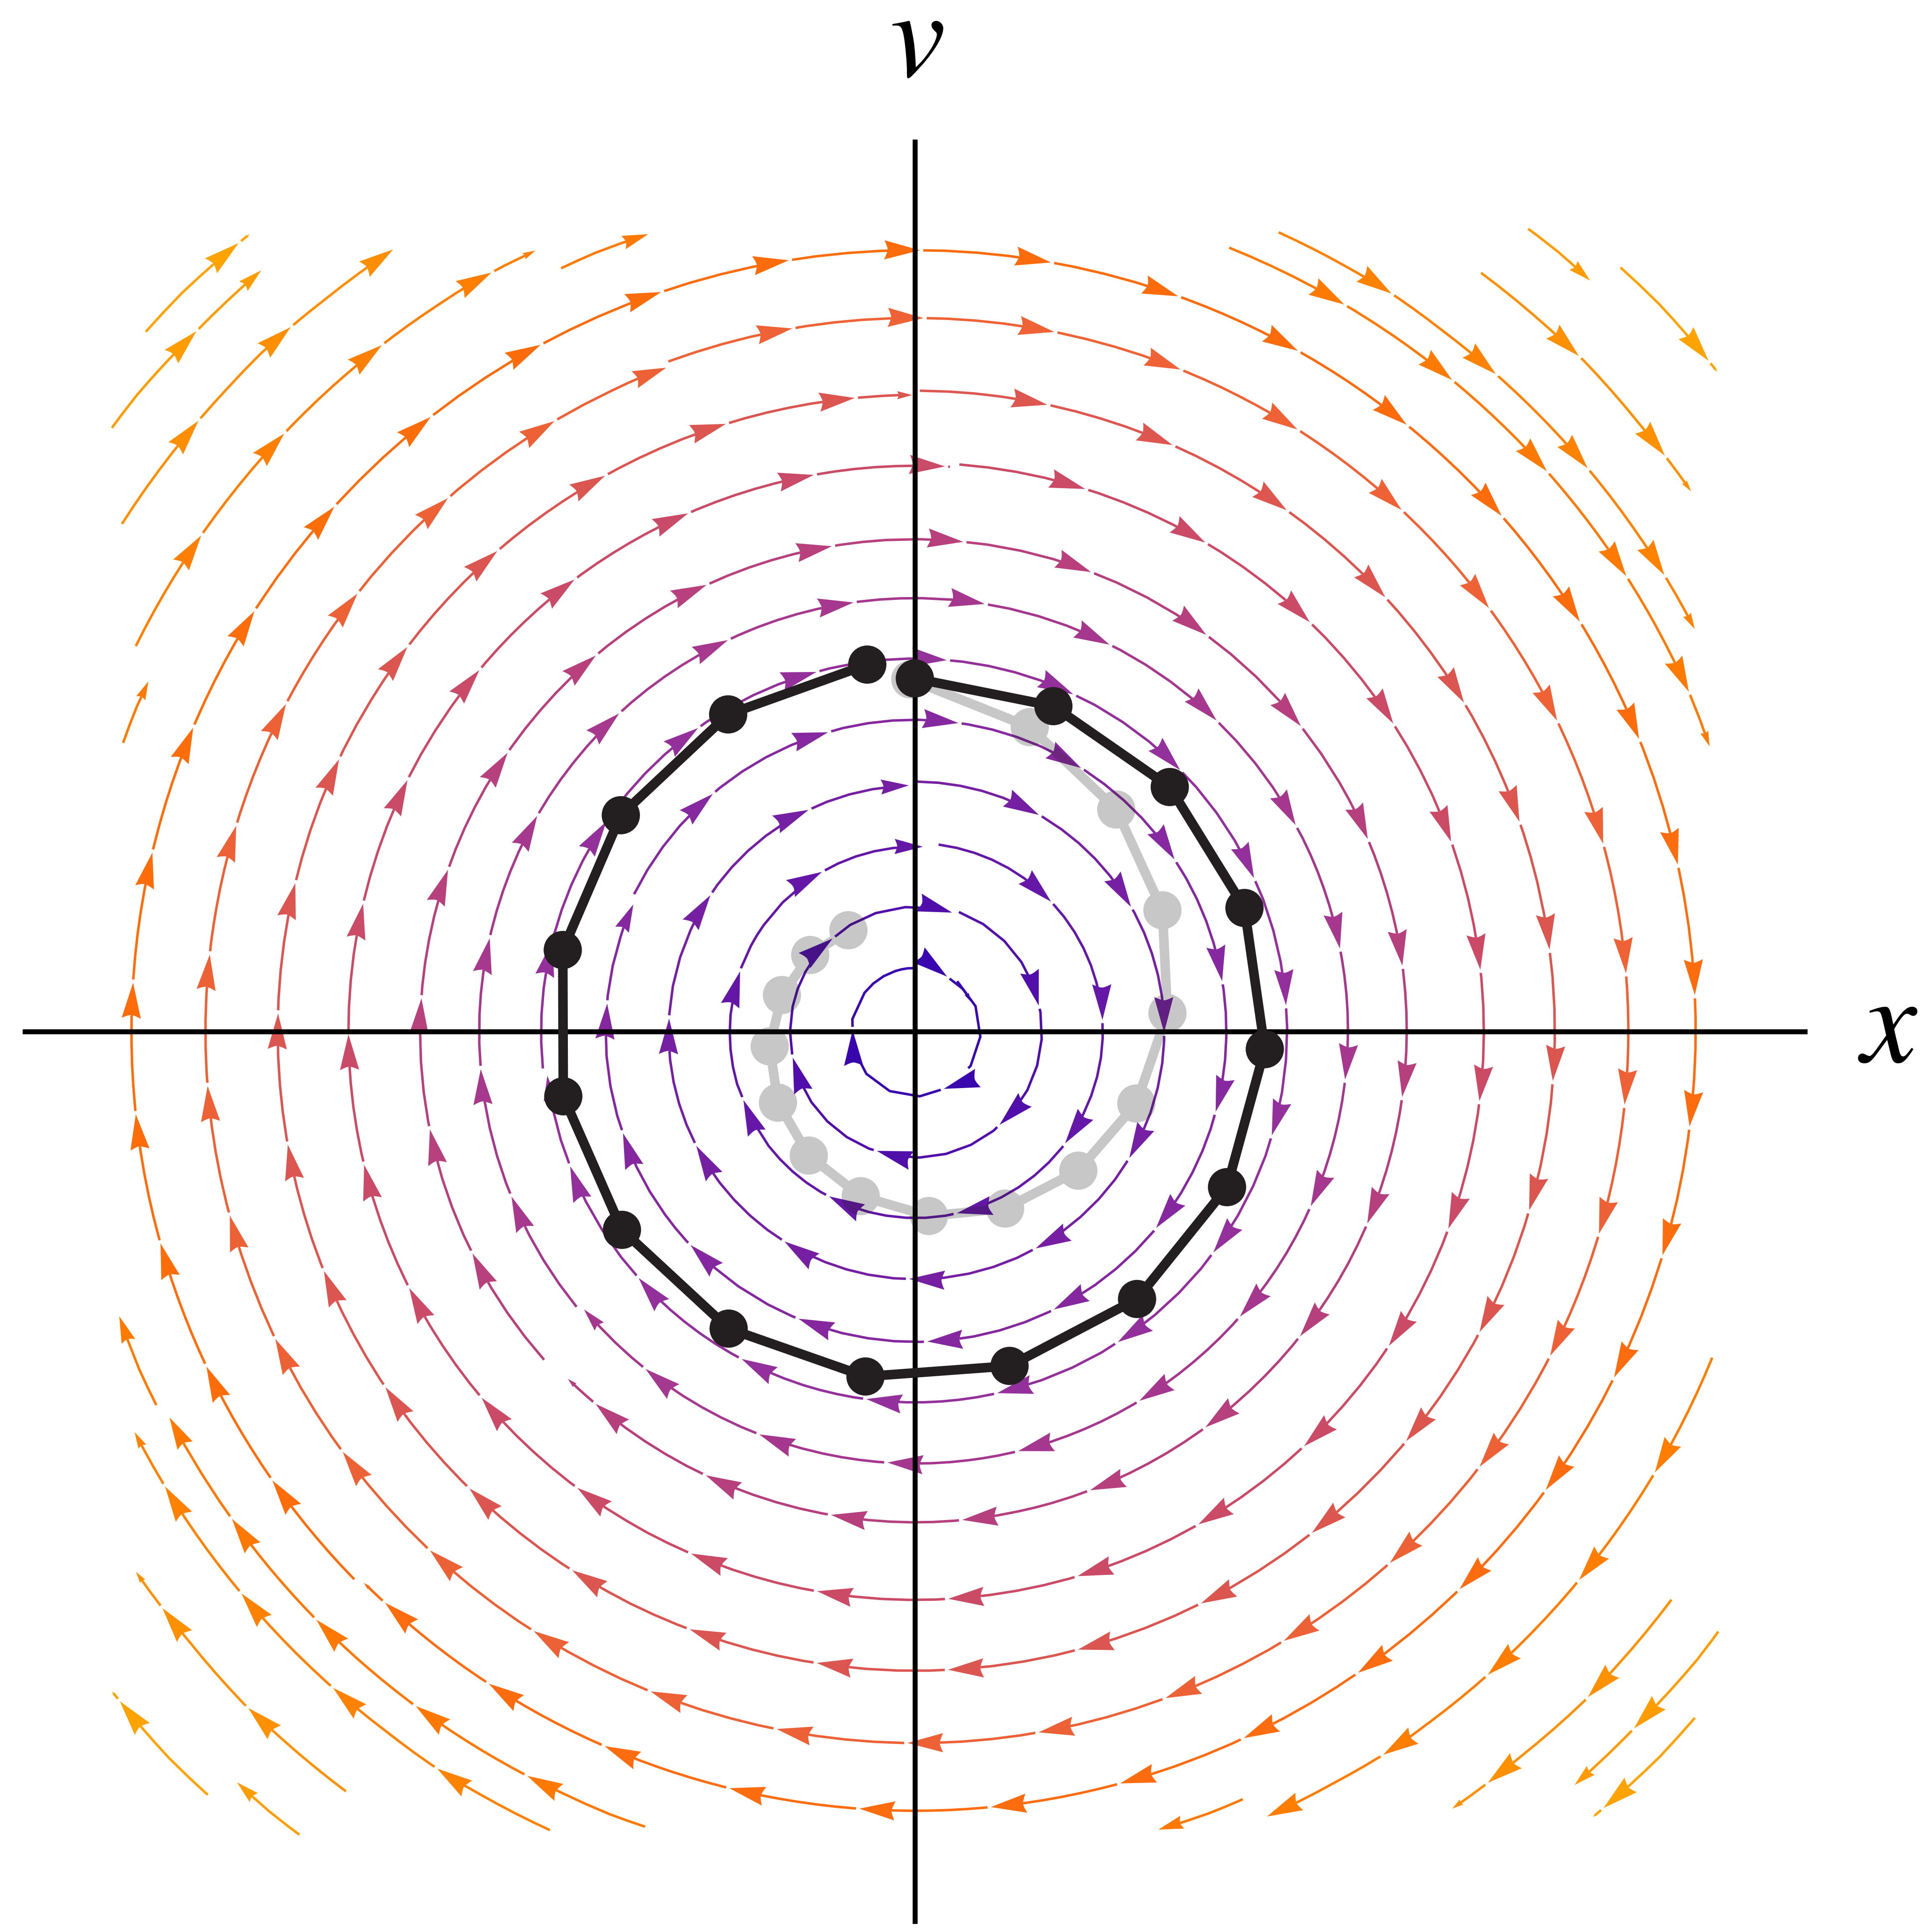

midpoint method

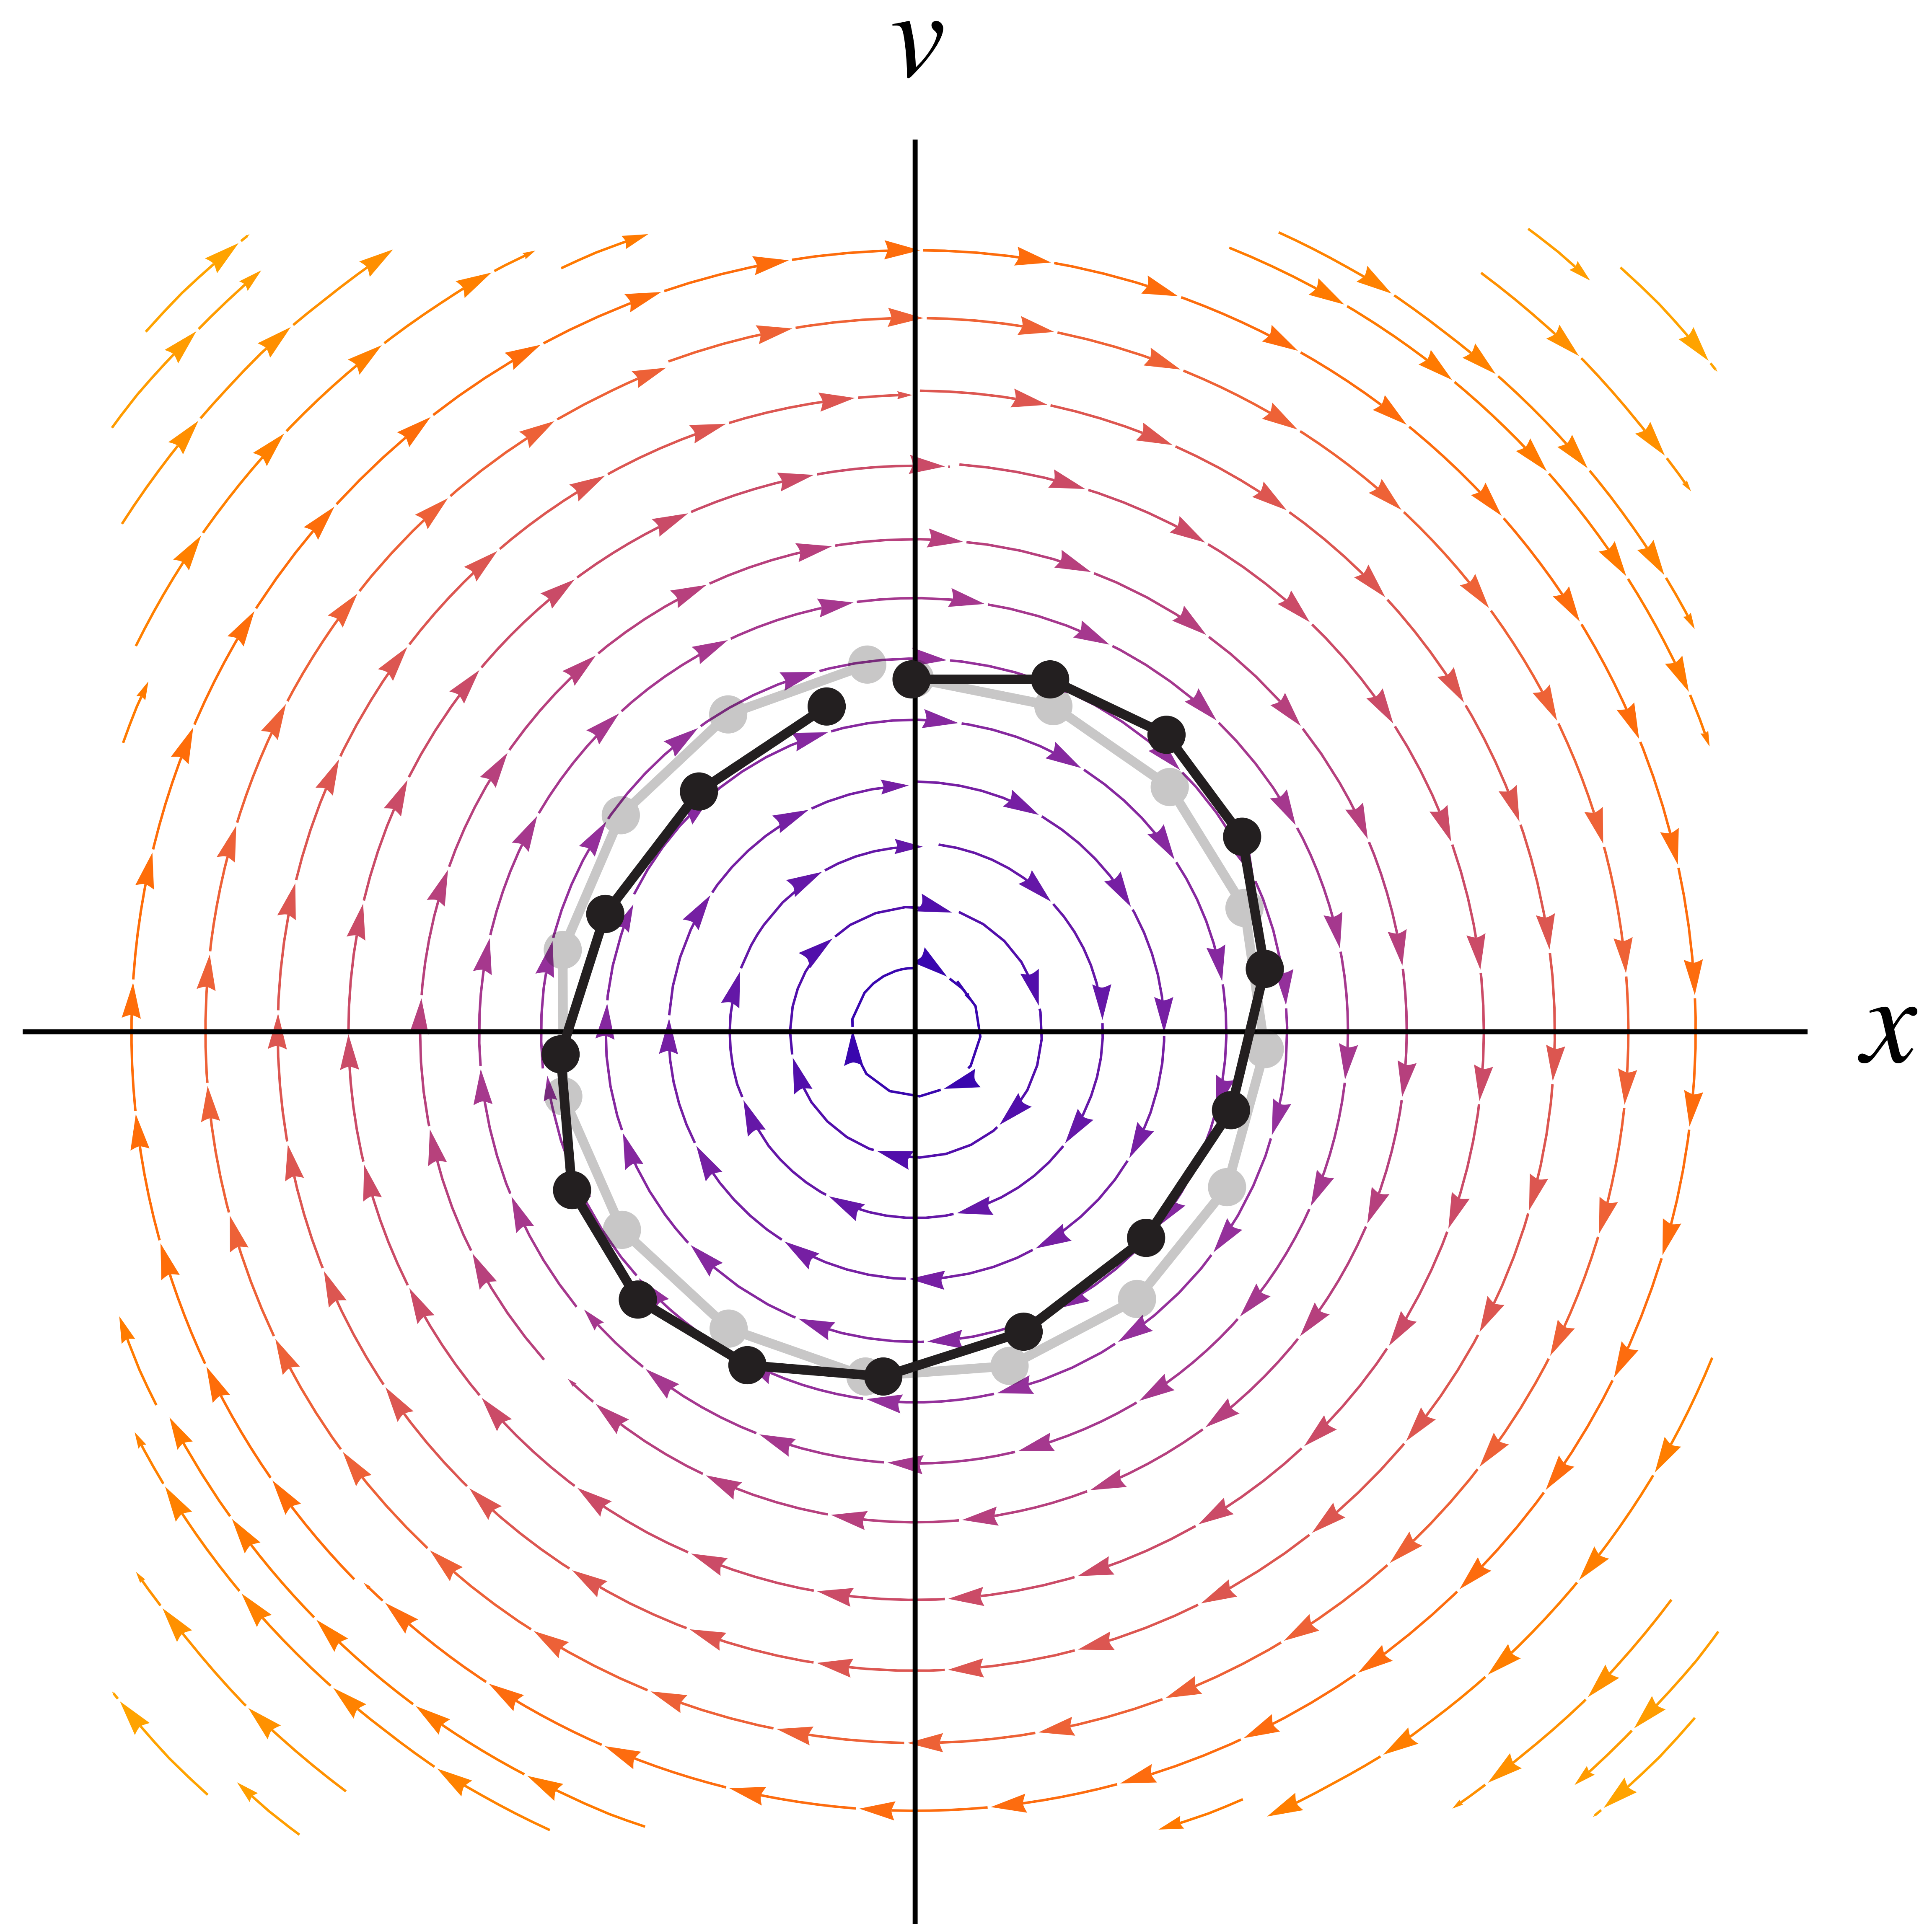

symplectic Euler

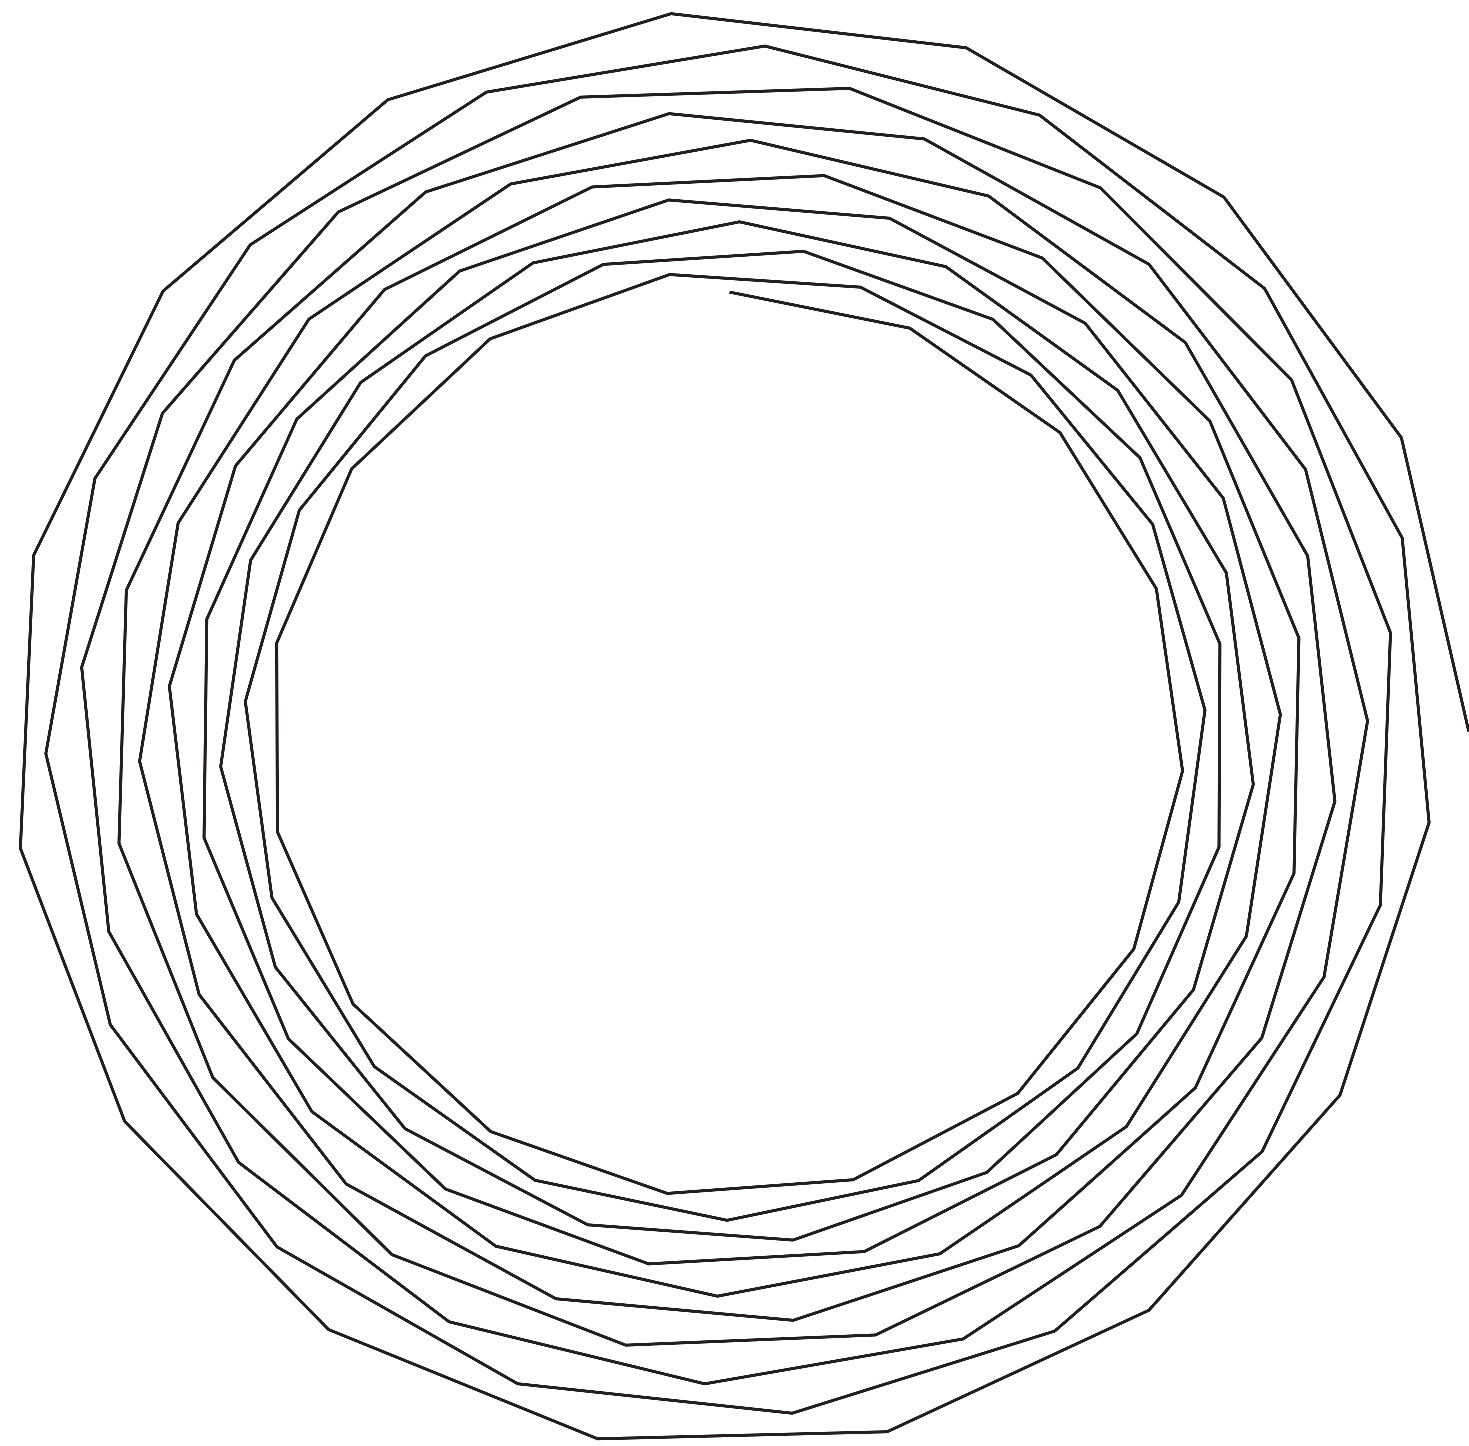

midpoint for 10 laps

v

x

symplectic Euler

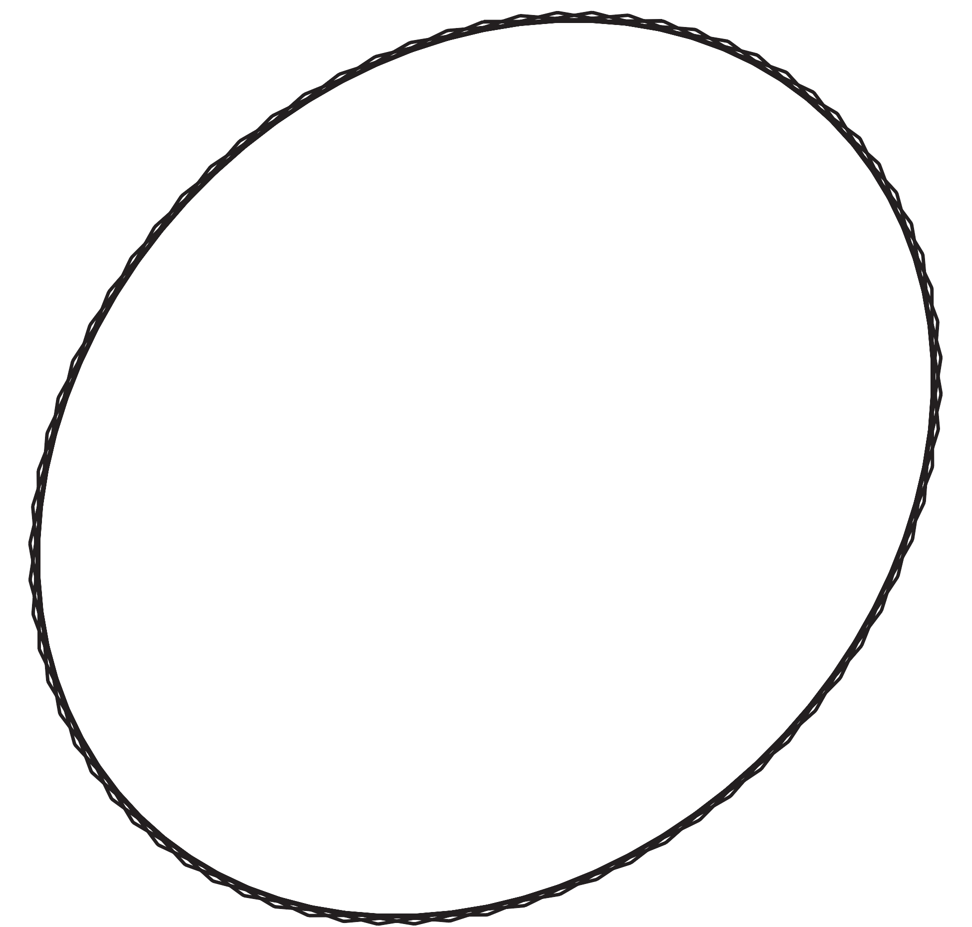

symplectic Euler for 10 laps