

NUMERICS OF MACHINE LEARNING

LECTURE 06


SOLVING ORDINARY DIFFERENTIAL EQUATIONS

Nathanael Bosch & Jonathan Schmidt
24 November 2022

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN



FACULTY OF SCIENCE
DEPARTMENT OF COMPUTER SCIENCE
CHAIR FOR THE METHODS OF MACHINE LEARNING



Where are we in the course?

- ▶ Last week: **State-space models and extended Kalman filters/smoother**
 (“How to estimate the *state* of a dynamical system from *observations*”)
- ▶ This week: **Ordinary differential equations and how to solve them**
 (“How to *simulate*, approximately, the evolution of a deterministic dynamical system”)

Today:

- ▶ What is an ordinary differential equation (ODE) and why should we care?
- ▶ **How to numerically solve an ODE:** From Euler (forward and backward) to Runge–Kutta
- ▶ **Parameter inference in ODEs** (and *neural* ODEs)



Ordinary differential equation:


$$\dot{x}(t) = f(x(t), t), \quad t \in \mathbb{T} \subset \mathbb{R},$$

where

- ▶ $x : \mathbb{T} \rightarrow \mathbb{R}^d$ is the unknown function
- ▶ $f : \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}^d$ is the *vector field*
- ▶ \mathbb{T} is the time domain; typically $\mathbb{T} = [0, T]$

Differential Equations can be found *everywhere*

- **Diffusion Models**
ODEs and SDEs for generative modeling




←
<https://developer.nvidia.com/blog/improving-diffusion-models-as-an-alternative-to-gans-part-1/>



Differential Equations can be found *everywhere*

- ▶ **Diffusion Models**
ODEs and SDEs for generative modeling
- ▶ **Normalizing Flows**
ODEs as bijectors to model distributions





[https://docs.pymc.io/en/v3/pymc-examples/
examples/variational_inference/normalizing_
Flows_overview.html](https://docs.pymc.io/en/v3/pymc-examples/examples/variational_inference/normalizing_flows_overview.html)



Differential Equations can be found *everywhere*


- ▶ **Diffusion Models**
ODEs and SDEs for generative modeling
- ▶ **Normalizing Flows**
ODEs as bijectors to model distributions
- ▶ **Neural ODEs**
ResNets as discretized ODEs



Chen et al, "Neural Ordinary Differential Equations",
NeurIPS 2018

Differential Equations can be found *everywhere*

- ▶ **Diffusion Models**
ODEs and SDEs for generative modeling
- ▶ **Normalizing Flows**
ODEs as bijectors to model distributions
- ▶ **Neural ODEs**
ResNets as discretized ODEs
- ▶ **Optimization Theory**
Gradient descent follows ODE dynamics




<https://francisbach.com/gradient-flows/>



Differential Equations can be found *everywhere*


- ▶ **Diffusion Models**
ODEs and SDEs for generative modeling
- ▶ **Normalizing Flows**
ODEs as bijectors to model distributions
- ▶ **Neural ODEs**
ResNets as discretized ODEs
- ▶ **Optimization Theory**
Gradient descent follows ODE dynamics
- ▶ **Parameter Inference** (later this lecture!)
ODEs as inductive bias



Tronarp, Bosch, Hennig, "Fenrir: Physics-Enhanced Regression for Initial Value Problems", ICML 2022

Differential Equations can be found *everywhere*

- ▶ **Diffusion Models**
ODEs and SDEs for generative modeling
- ▶ **Normalizing Flows**
ODEs as bijectors to model distributions
- ▶ **Neural ODEs**
ResNets as discretized ODEs
- ▶ **Optimization Theory**
Gradient descent follows ODE dynamics
- ▶ **Parameter Inference** (later this lecture!)
ODEs as inductive bias
- ▶ **Probabilistic Numerics** (next lecture!)
ODE solving as *learning*



<https://raw.githubusercontent.com/nathanaelbosch/ProbNumDiffEq.jl/main/examples/banner.svg>

Ordinary differential equation :

$$\dot{x}(t) = f(x(t), t), \quad t \in \mathbb{T} \subset \mathbb{R},$$

where

- ▶ $x : \mathbb{T} \rightarrow \mathbb{R}^d$ is the unknown function
- ▶ $f : \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}^d$ is the *vector field*
- ▶ \mathbb{T} is the time domain; typically $\mathbb{T} = [0, T]$

Ordinary differential equation :

$$\dot{x}(t) = f(x(t), t), \quad t \in \mathbb{T} \subset \mathbb{R},$$

where

- ▶ $x : \mathbb{T} \rightarrow \mathbb{R}^d$ is the unknown function
- ▶ $f : \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}^d$ is the *vector field*
- ▶ \mathbb{T} is the time domain; typically $\mathbb{T} = [0, T]$

Solution (fundamental theorem of calculus):

$$x(t) = x(0) + \int_0^t f(x(\tau), \tau) \, d\tau$$

Ordinary differential equation :

$$\dot{x}(t) = f(x(t), t), \quad t \in \mathbb{T} \subset \mathbb{R},$$

where

- ▶ $x : \mathbb{T} \rightarrow \mathbb{R}^d$ is the unknown function
- ▶ $f : \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}^d$ is the *vector field*
- ▶ \mathbb{T} is the time domain; typically $\mathbb{T} = [0, T]$

Solution (fundamental theorem of calculus):

$$x(t) = x(0) + \int_0^t f(x(\tau), \tau) \, d\tau$$

⇒ Solutions depend on the initial value $x(0)$

Ordinary differential equation **initial value problem**:

$$\dot{x}(t) = f(x(t), t), \quad t \in \mathbb{T} \subset \mathbb{R}, \quad x(0) = x_0,$$

where

- ▶ $x : \mathbb{T} \rightarrow \mathbb{R}^d$ is the unknown function
- ▶ $f : \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}^d$ is the *vector field*
- ▶ \mathbb{T} is the time domain; typically $\mathbb{T} = [0, T]$
- ▶ x_0 is the initial value

Solution (fundamental theorem of calculus):

$$x(t) = x(0) + \int_0^t f(x(\tau), \tau) \, d\tau$$

⇒ Solutions depend on the initial value $x(0)$

A simple example: Modeling population growth

The logistic ODE

► Logistic ODE:

$$\dot{P}(t) = rP(t) \left(1 - \frac{P(t)}{K}\right),$$

where P is the population size, r is the growth rate, and K is the carrying capacity (bottleneck).



Pierre-François Verhulst

A simple example: Modeling population growth

The logistic ODE

► **Logistic ODE:**

$$\dot{P}(t) = rP(t) \left(1 - \frac{P(t)}{K} \right),$$

where P is the population size, r is the growth rate, and K is the carrying capacity (bottleneck).

► **Solution:**

$$P(t) = \frac{K}{1 + \frac{K-P(0)}{P(0)} e^{-rt}}.$$

(You can verify for yourself by taking its derivative!)



Pierre-François Verhulst

A simple example: Modeling population growth



The logistic ODE

https://upload.wikimedia.org/wikipedia/commons/0/04/Pierre_Francois_Verhulst.jpg

► Logistic ODE:

$$\dot{P}(t) = rP(t) \left(1 - \frac{P(t)}{K} \right),$$


where P is the population size, r is the growth rate, and K is the carrying capacity (bottleneck).

► Solution:

$$P(t) = \frac{K}{1 + \frac{K-P(0)}{P(0)} e^{-rt}}.$$

(You can verify for yourself by taking its derivative!)

► Example with $K = 1$ and $r = 1$:



A simple example: Modeling population growth



The logistic ODE

https://upload.wikimedia.org/wikipedia/commons/0/04/Pierre_Francois_Verhulst.jpg

► **Logistic ODE:**

$$\dot{P}(t) = rP(t) \left(1 - \frac{P(t)}{K} \right),$$

where P is the population size, r is the growth rate, and K is the carrying capacity (bottleneck).


► **Solution:**

$$P(t) = \frac{K}{1 + \frac{K-P(0)}{P(0)} e^{-rt}}.$$

(You can verify for yourself by taking its derivative!)

► **Initial condition:** $P(0) = 0.01$.

► **Example with $K = 1$ and $r = 1$:**



A simple example: Modeling population growth



The logistic ODE

https://upload.wikimedia.org/wikipedia/commons/0/04/Pierre_Francois_Verhulst.jpg

► Logistic ODE:

$$\dot{P}(t) = rP(t) \left(1 - \frac{P(t)}{K} \right),$$

where P is the population size, r is the growth rate, and K is the carrying capacity (bottleneck).


► Solution:

$$P(t) = \frac{K}{1 + \frac{K-P(0)}{P(0)} e^{-rt}}.$$

(You can verify for yourself by taking its derivative!)

► Initial condition: $P(0) = 0.01$.

► Example with $K = 1$ and $r = 1$:



Next: **How can we solve ODEs in general?**



How to *numerically* solve ODEs – the general case

Is there a way we can solve ODEs in general?

Recall: The initial value problem

$$\dot{x}(t) = f(x(t), t), \quad t \in [0, T], \quad x(0) = x_0,$$

has the solution

$$x(t) = x(0) + \int_0^t f(x(\tau), \tau) \, d\tau.$$

How to *numerically* solve ODEs – the general case

Is there a way we can solve ODEs in general?

Recall: The initial value problem

$$\dot{x}(t) = f(x(t), t), \quad t \in [0, T], \quad x(0) = x_0,$$

has the solution

$$x(t) = x(0) + \int_0^t f(x(\tau), \tau) \, d\tau.$$

Numerical solvers extrapolate step by step: If we know $x(t)$, then $x(t + h)$ is given by

$$x(t + h) = x(t) + \int_t^{t+h} f(x(\tau), \tau) \, d\tau.$$

⇒ At each step, we need to approximate $\int_t^{t+h} f(x(\tau), \tau) \, d\tau$.

How to *numerically* solve ODEs – the general case

Is there a way we can solve ODEs in general?

Recall: The initial value problem

$$\dot{x}(t) = f(x(t), t), \quad t \in [0, T], \quad x(0) = x_0,$$

has the solution

$$x(t) = x(0) + \int_0^t f(x(\tau), \tau) \, d\tau.$$

Numerical solvers extrapolate step by step: If we know $x(t)$, then $x(t + h)$ is given by

$$x(t + h) = x(t) + \int_t^{t+h} f(x(\tau), \tau) \, d\tau.$$

⇒ At each step, we need to approximate $\int_t^{t+h} f(x(\tau), \tau) \, d\tau$.

How?

How to *numerically* solve ODEs

Taylor series expansions to the rescue

$$\text{Recall: } x(t+h) = x(t) + \int_t^{t+h} f(x(\tau), \tau) d\tau.$$

$$\text{Recall: } x(t+h) = x(t) + \int_t^{t+h} f(x(\tau), \tau) d\tau.$$

Definition (Taylor Series Expansion)

Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be a function. Then the *Taylor series expansion* of g at t_0 is given by

$$g(\tau) = g(t_0) + g^{(1)}(t_0)(\tau - t_0) + \frac{1}{2}g^{(2)}(t_0)(\tau - t_0)^2 + \dots = \sum_{n=0}^{\infty} \frac{g^{(n)}(t_0)}{n!} (\tau - t_0)^n.$$

Recall:
$$x(t+h) = x(t) + \int_t^{t+h} f(x(\tau), \tau) d\tau.$$

Definition (Taylor Series Expansion)

Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be a function. Then the *Taylor series expansion* of g at t_0 is given by

$$g(\tau) = g(t_0) + g^{(1)}(t_0)(\tau - t_0) + \frac{1}{2}g^{(2)}(t_0)(\tau - t_0)^2 + \dots = \sum_{n=0}^{\infty} \frac{g^{(n)}(t_0)}{n!}(\tau - t_0)^n.$$

- Zeroth-order Taylor series approximation around $\tau = t$:

$$\text{Recall: } x(t+h) = x(t) + \int_t^{t+h} f(x(\tau), \tau) d\tau.$$

Definition (Taylor Series Expansion)

Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be a function. Then the *Taylor series expansion* of g at t_0 is given by

$$g(\tau) = g(t_0) + g^{(1)}(t_0)(\tau - t_0) + \frac{1}{2}g^{(2)}(t_0)(\tau - t_0)^2 + \dots = \sum_{n=0}^{\infty} \frac{g^{(n)}(t_0)}{n!} (\tau - t_0)^n.$$

- Zeroth-order Taylor series approximation around $\tau = t$:

$$f(x(\tau), \tau) \approx f(x(t), t) + \mathcal{O}(h)$$

How to *numerically* solve ODEs

Taylor series expansions to the rescue

$$\text{Recall: } x(t+h) = x(t) + \int_t^{t+h} f(x(\tau), \tau) d\tau.$$

Definition (Taylor Series Expansion)

Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be a function. Then the *Taylor series expansion* of g at t_0 is given by

$$g(\tau) = g(t_0) + g^{(1)}(t_0)(\tau - t_0) + \frac{1}{2}g^{(2)}(t_0)(\tau - t_0)^2 + \dots = \sum_{n=0}^{\infty} \frac{g^{(n)}(t_0)}{n!} (\tau - t_0)^n.$$

- Zeroth-order Taylor series approximation around $\tau = t$:

$$\begin{aligned}
 f(x(\tau), \tau) &\approx f(x(t), t) + \mathcal{O}(h) \\
 \Rightarrow \int_t^{t+h} f(x(\tau), \tau) d\tau &\approx \int_t^{t+h} f(x(t), t) d\tau = h \cdot f(x(t), t).
 \end{aligned}$$

How to *numerically* solve ODEs

Taylor series expansions to the rescue

$$\text{Recall: } x(t+h) = x(t) + \int_t^{t+h} f(x(\tau), \tau) d\tau.$$

Definition (Taylor Series Expansion)

Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be a function. Then the *Taylor series expansion* of g at t_0 is given by

$$g(\tau) = g(t_0) + g^{(1)}(t_0)(\tau - t_0) + \frac{1}{2}g^{(2)}(t_0)(\tau - t_0)^2 + \dots = \sum_{n=0}^{\infty} \frac{g^{(n)}(t_0)}{n!} (\tau - t_0)^n.$$

- Zeroth-order Taylor series approximation around $\tau = t$:

$$\begin{aligned} f(x(\tau), \tau) &\approx f(x(t), t) + \mathcal{O}(h) \\ \Rightarrow \int_t^{t+h} f(x(\tau), \tau) d\tau &\approx \int_t^{t+h} f(x(t), t) d\tau = h \cdot f(x(t), t). \end{aligned}$$

(Explicit) Forward Euler: $\hat{x}(t+h) = \hat{x}(t) + h \cdot f(\hat{x}(t), t).$

$$\text{Recall: } x(t+h) = x(t) + \int_t^{t+h} f(x(\tau), \tau) d\tau.$$

Definition (Taylor Series Expansion)

Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be a function. Then the *Taylor series expansion* of g at t_0 is given by

$$g(\tau) = g(t_0) + g^{(1)}(t_0)(\tau - t_0) + \frac{1}{2}g^{(2)}(t_0)(\tau - t_0)^2 + \dots = \sum_{n=0}^{\infty} \frac{g^{(n)}(t_0)}{n!} (\tau - t_0)^n.$$

- Zeroth-order Taylor series approximation around $\tau = t+h$:

$$\begin{aligned} f(x(\tau), \tau) &\approx f(x(t+h), t+h) + \mathcal{O}(h) \\ \Rightarrow \int_t^{t+h} f(x(\tau), \tau) d\tau &\approx \int_t^{t+h} f(x(t+h), t+h) d\tau = h \cdot f(x(t+h), t+h). \end{aligned}$$

(Implicit) Backward Euler: $\hat{x}(t+h) = \hat{x}(t) + h \cdot f(\hat{x}(t+h), t+h).$

$$\text{Recall: } x(t+h) = x(t) + \int_t^{t+h} f(x(\tau), \tau) d\tau.$$

Definition (Taylor Series Expansion)

Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be a function. Then the *Taylor series expansion* of g at t_0 is given by

$$g(\tau) = g(t_0) + g^{(1)}(t_0)(\tau - t_0) + \frac{1}{2}g^{(2)}(t_0)(\tau - t_0)^2 + \dots = \sum_{n=0}^{\infty} \frac{g^{(n)}(t_0)}{n!} (\tau - t_0)^n.$$

- Zeroth-order Taylor series approximation around $\tau = t+h$:

$$\begin{aligned} f(x(\tau), \tau) &\approx f(x(t+h), t+h) + \mathcal{O}(h) \\ \Rightarrow \int_t^{t+h} f(x(\tau), \tau) d\tau &\approx \int_t^{t+h} f(x(t+h), t+h) d\tau = h \cdot f(x(t+h), t+h). \end{aligned}$$

(Implicit) Backward Euler: $\hat{x}(t+h) = \hat{x}(t) + h \cdot f(\hat{x}(t+h), t+h)$.

```
1 using Plots
```

```
2  
3 f(x, t) = x * (1 - x)
```

```
4 x0, tspan = 0.1, (0, 5)
```

```
5  
6 h = 1 // 10
```

```
7 x, out = x0, [x0]
```


```
8 for t in tspan[1]:h:(tspan[2]-h)
```

```
9     x = x + h * f(x, t)
```

```
10     push!(out, x)
```

```
11 end
```

```
12  
13 plot(tspan[1]:h:tspan[2], out, marker=:o)
```



```
1 using Plots, Roots
```

```
2  
3 f(x, t) = x * (1 - x)
```

```
4 x0, tspan = 0.1, (0, 5)
```

```
5  
6 h = 1 // 10
```

```
7 x, out = x0, [x0]
```


```
8 for t in (tspan[1]+h):h:tspan[2]
```

```
9     x = find_zero(y -> y - (x + h*f(y, t)), x)
```

```
10     push!(out, x)
```

```
11 end
```

```
12  
13 plot(tspan[1]:h:tspan[2], out, marker=:o)
```




Stability: The difference between forward and backward Euler

Sometimes explicit methods are just not that great

Consider the following scalar ODE (test equation)

$$\dot{x}(t) = \lambda x(t).$$

How small do we have to make the steps, depending on λ ?
(here $\lambda = -21$).



Stability: The difference between forward and backward Euler

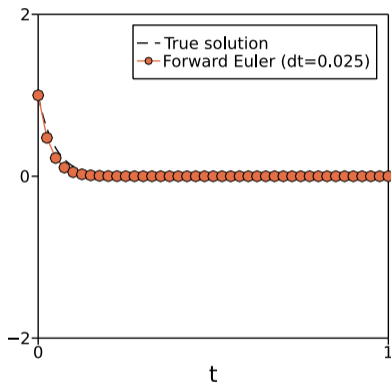
Sometimes explicit methods are just not that great

Consider the following scalar ODE (test equation)

$$\dot{x}(t) = \lambda x(t).$$

How small do we have to make the steps, depending on λ ?
(here $\lambda = -21$).

- **Forward Euler:** $\hat{x}(t+h) = \hat{x}(t) + h \cdot \lambda \hat{x}(t)$
 $\Rightarrow \hat{x}(t+h) = (1 + h\lambda) \cdot \hat{x}(t)$
 \Rightarrow For $\hat{x}(t)$ to remain bounded, we need $|1 + h\lambda| \leq 1$.



Stability: The difference between forward and backward Euler


Sometimes explicit methods are just not that great

Consider the following scalar ODE (test equation)

$$\dot{x}(t) = \lambda x(t).$$

How small do we have to make the steps, depending on λ ?
(here $\lambda = -21$).

- **Forward Euler:** $\hat{x}(t+h) = \hat{x}(t) + h \cdot \lambda \hat{x}(t)$
 $\Rightarrow \hat{x}(t+h) = (1 + h\lambda) \cdot \hat{x}(t)$
 \Rightarrow For $\hat{x}(t)$ to remain bounded, we need $|1 + h\lambda| \leq 1$.



Stability: The difference between forward and backward Euler


Sometimes explicit methods are just not that great

Consider the following scalar ODE (test equation)

$$\dot{x}(t) = \lambda x(t).$$

How small do we have to make the steps, depending on λ ?
(here $\lambda = -21$).

- ▶ **Forward Euler:** $\hat{x}(t+h) = \hat{x}(t) + h \cdot \lambda \hat{x}(t)$
 $\Rightarrow \hat{x}(t+h) = (1+h\lambda) \cdot \hat{x}(t)$
 \Rightarrow For $\hat{x}(t)$ to remain bounded, we need $|1+h\lambda| \leq 1$.
- ▶ **Backward Euler:** $\hat{x}(t+h) = \hat{x}(t) + h \cdot \lambda \hat{x}(t+h)$
 $\Rightarrow \hat{x}(t+h) = \frac{1}{(1-h\lambda)} \cdot \hat{x}(t)$
 $\Rightarrow \frac{1}{|1-h\lambda|} \leq 1$



Stability: The difference between forward and backward Euler

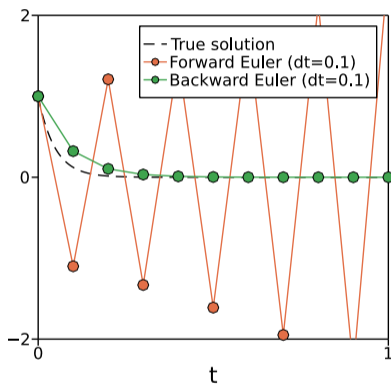
Sometimes explicit methods are just not that great

Consider the following scalar ODE (test equation)

$$\dot{x}(t) = \lambda x(t).$$

How small do we have to make the steps, depending on λ ?
(here $\lambda = -21$).

- ▶ **Forward Euler:** $\hat{x}(t+h) = \hat{x}(t) + h \cdot \lambda \hat{x}(t)$
 $\Rightarrow \hat{x}(t+h) = (1+h\lambda) \cdot \hat{x}(t)$
 \Rightarrow For $\hat{x}(t)$ to remain bounded, we need $|1+h\lambda| \leq 1$.
- ▶ **Backward Euler:** $\hat{x}(t+h) = \hat{x}(t) + h \cdot \lambda \hat{x}(t+h)$
 $\Rightarrow \hat{x}(t+h) = \frac{1}{(1-h\lambda)} \cdot \hat{x}(t)$
 $\Rightarrow \frac{1}{|1-h\lambda|} \leq 1$





Stability: The difference between forward and backward Euler


Sometimes explicit methods are just not that great

Consider the following scalar ODE (test equation)

$$\dot{x}(t) = \lambda x(t).$$

How small do we have to make the steps, depending on λ ?
 (here $\lambda = -21$).

- ▶ **Forward Euler:** $\hat{x}(t+h) = \hat{x}(t) + h \cdot \lambda \hat{x}(t)$
 $\Rightarrow \hat{x}(t+h) = (1 + h\lambda) \cdot \hat{x}(t)$
 \Rightarrow For $\hat{x}(t)$ to remain bounded, we need $|1 + h\lambda| \leq 1$.
- ▶ **Backward Euler:** $\hat{x}(t+h) = \hat{x}(t) + h \cdot \lambda \hat{x}(t+h)$
 $\Rightarrow \hat{x}(t+h) = \frac{1}{(1-h\lambda)} \cdot \hat{x}(t)$
 $\Rightarrow \frac{1}{|1-h\lambda|} \leq 1$



\Rightarrow Different algorithms have different stability properties!

Next: Runge–Kutta solvers



How to *numerically* solve ODEs – continued

Building better solvers via numerical quadrature

$$\text{Recall: } x(t+h) = x(t) + \int_t^{t+h} f(x(\tau), \tau) d\tau.$$

How to *numerically* solve ODEs – continued

Building better solvers via numerical quadrature

Recall:
$$x(t+h) = x(t) + \int_t^{t+h} f(x(\tau), \tau) d\tau.$$

Numerical Quadrature

Let $g : \mathbb{R} \rightarrow \mathbb{R}^d$ be a function. Then *numerical quadrature* (or *numerical integration*) approximates

$$\int_l^r g(\tau) d\tau \approx \sum_{i=1}^n w_i g(t_i),$$

where t_i are the quadrature nodes and w_i are the quadrature weights.

How to *numerically* solve ODEs – continued

Building better solvers via numerical quadrature

$$\text{Recall: } x(t+h) = x(t) + \int_t^{t+h} f(x(\tau), \tau) d\tau.$$

Numerical Quadrature

Let $g : \mathbb{R} \rightarrow \mathbb{R}^d$ be a function. Then *numerical quadrature* (or *numerical integration*) approximates

$$\int_l^r g(\tau) d\tau \approx \sum_{i=1}^n w_i g(t_i),$$

where t_i are the quadrature nodes and w_i are the quadrature weights.

- This motivates (explicit) **Runge–Kutta**:

$$\int_t^{t+h} f(x(\tau), \tau) d\tau \approx h \cdot \sum_{i=1}^s w_i f(\hat{x}(\tau_i), \tau_i).$$

Recall:
$$x(t+h) = x(t) + \int_t^{t+h} f(x(\tau), \tau) d\tau.$$

Numerical Quadrature

Let $g : \mathbb{R} \rightarrow \mathbb{R}^d$ be a function. Then *numerical quadrature* (or *numerical integration*) approximates

$$\int_l^r g(\tau) d\tau \approx \sum_{i=1}^n w_i g(t_i),$$

where t_i are the quadrature nodes and w_i are the quadrature weights.

- This motivates (explicit) **Runge–Kutta**:

$$\int_t^{t+h} f(x(\tau), \tau) d\tau \approx h \cdot \sum_{i=1}^s w_i f(\hat{x}(\tau_i), \tau_i).$$

Next: How to choose weights w_i and nodes τ_i ? And how to construct $\hat{x}(\tau_i)$?

Definition ((Explicit) Runge–Kutta method)

An explicit Runge–Kutta method is given by

$$\hat{x}(t+h) = \hat{x}(t) + h \sum_{i=1}^s b_i k_i,$$

where

$$k_1 = f(\hat{x}(t), t),$$

$$k_2 = f(\hat{x}(t) + h(a_{21}k_1), t + hc_2),$$

$$k_3 = f(\hat{x}(t) + h(a_{31}k_1 + a_{32}k_2), t + hc_3),$$

⋮

$$k_s = f\left(\hat{x}(t) + h \sum_{j=1}^{s-1} a_{sj} k_j, t + hc_s\right).$$

Definition ((Explicit) Runge–Kutta method)

An explicit Runge–Kutta method is given by

$$\hat{x}(t+h) = \hat{x}(t) + h \sum_{i=1}^s b_i k_i,$$

where

$$k_1 = f(\hat{x}(t), t),$$

$$k_2 = f(\hat{x}(t) + h(a_{21}k_1), t + hc_2),$$

$$k_3 = f(\hat{x}(t) + h(a_{31}k_1 + a_{32}k_2), t + hc_3),$$

⋮

$$k_s = f\left(\hat{x}(t) + h \sum_{j=1}^{s-1} a_{sj} k_j, t + hc_s\right).$$

“Butcher tableau”: A compact representation of a *specific* Runge–Kutta method

0				
c ₂	a ₂₁			
c ₃	a ₃₁	a ₃₂		
⋮	⋮	⋮	⋱	
c _s	a _{s1}	a _{s2}	⋯	a _{s,s-1}
	b ₁	b ₂	⋯	b _s

Runge–Kutta Methods – Examples

Forward Euler is a Runge–Kutta method!

Forward Euler is a Runge–Kutta method!

Runge–Kutta in general:

$$\hat{x}(t+h) = \hat{x}(t) + h \sum_{i=1}^s b_i k_i, \quad \text{with} \quad k_i = f \left(\hat{x}(t) + h \sum_{j=1}^{i-1} a_{ij} k_j, t + hc_i \right).$$

Runge–Kutta Methods – Examples

Forward Euler is a Runge–Kutta method!

Runge–Kutta in general:

$$\hat{x}(t+h) = \hat{x}(t) + h \sum_{i=1}^s b_i k_i, \quad \text{with } k_i = f \left(\hat{x}(t) + h \sum_{j=1}^{i-1} a_{ij} k_j, t + hc_i \right).$$

Turns out **forward Euler is actually a Runge–Kutta method**:

$$\begin{aligned} k_1 &= f(\hat{x}(t), t), \\ \hat{x}(t+h) &= \hat{x}(t) + hk_1. \end{aligned}$$

Runge–Kutta Methods – Examples

Forward Euler is a Runge–Kutta method!

Runge–Kutta in general:

$$\hat{x}(t+h) = \hat{x}(t) + h \sum_{i=1}^s b_i k_i, \quad \text{with } k_i = f \left(\hat{x}(t) + h \sum_{j=1}^{i-1} a_{ij} k_j, t + hc_i \right).$$

Turns out **forward Euler is actually a Runge–Kutta method**:

$$k_1 = f(\hat{x}(t + 0h), t + 0h),$$

$$\hat{x}(t+h) = \hat{x}(t) + h \cdot 1k_1.$$

Butcher tableau:

0	0
	1

Runge–Kutta Methods – Examples

Backward Euler is a Runge–Kutta method as well!

Runge–Kutta in general:

$$\hat{x}(t+h) = \hat{x}(t) + h \sum_{i=1}^s b_i k_i, \quad \text{with } k_i = f \left(\hat{x}(t) + h \sum_{j=1}^{i-1} a_{ij} k_j, t + hc_i \right).$$

Turns out **backward Euler is actually a (implicit) Runge–Kutta method:**

$$k_1 = f(\hat{x}(t + 1h), t + 1h),$$

$$\hat{x}(t+h) = \hat{x}(t) + h \cdot 1k_1.$$

Butcher tableau:

$$\begin{array}{c|c} 1 & 1 \\ \hline & 1 \end{array} \quad (\text{the } 1 \text{ makes it implicit!})$$

Runge–Kutta Methods – Examples

Improving on forward Euler with the *explicit midpoint rule*

The **explicit midpoint rule** aims to improve the accuracy of the forward Euler method by selecting

$$\hat{x}(t+h) = \hat{x}(t) + hf\left(\hat{x}\left(t + \frac{h}{2}\right), t + \frac{h}{2}\right).$$

But how to choose $\hat{x}\left(t + \frac{h}{2}\right)$?

Runge–Kutta Methods – Examples

Improving on forward Euler with the *explicit midpoint rule*

The **explicit midpoint rule** aims to improve the accuracy of the forward Euler method by selecting

$$\hat{x}(t+h) = \hat{x}(t) + hf\left(\hat{x}\left(t + \frac{h}{2}\right), t + \frac{h}{2}\right).$$

But how to choose $\hat{x}\left(t + \frac{h}{2}\right)$?

With another Euler step!

Runge–Kutta Methods – Examples

Improving on forward Euler with the *explicit midpoint rule*

The **explicit midpoint rule** aims to improve the accuracy of the forward Euler method by selecting

$$\hat{x}(t+h) = \hat{x}(t) + hf\left(\hat{x}\left(t + \frac{h}{2}\right), t + \frac{h}{2}\right).$$

But how to choose $\hat{x}\left(t + \frac{h}{2}\right)$?

With another Euler step!

This leads to the scheme:

$$k_1 = f(\hat{x}(t), t + 0h),$$

$$k_2 = f\left(\hat{x}(t) + h \cdot \frac{1}{2}k_1, t + \frac{1}{2}h\right),$$

$$\hat{x}(t+h) = \hat{x}(t) + h(0k_1 + 1k_2).$$

0	
$\frac{1}{2}$	$\frac{1}{2}$
	0 1

Runge–Kutta Methods – Examples

The original Runge–Kutta methods of order 4 ($s = 4$)

The **classic fourth-order Runge–Kutta method** selects

$$k_1 = f(\hat{x}(t), t + 0h),$$

$$k_2 = f\left(\hat{x}(t) + h \cdot \frac{1}{2}k_1, t + \frac{1}{2}h\right),$$

$$k_3 = f\left(\hat{x}(t) + h \cdot \frac{1}{2}k_2, t + \frac{1}{2}h\right),$$

$$k_4 = f(\hat{x}(t) + h \cdot k_3, t + h),$$

and then

$$\hat{x}(t + h) = \hat{x}(t) + h \left(\frac{1}{6}k_1 + \frac{1}{3}k_2 + \frac{1}{3}k_3 + \frac{1}{6}k_4 \right).$$

0			
$\frac{1}{2}$	$\frac{1}{2}$		
$\frac{1}{2}$	0	$\frac{1}{2}$	
1	0	0	1
	$\frac{1}{6}$	$\frac{1}{3}$	$\frac{1}{3}$
	$\frac{1}{6}$	$\frac{1}{3}$	$\frac{1}{6}$

(Further reading: “Solving Ordinary Differential Equations I” by Hairer, Norsett and Wanner, Chapter II.1; includes derivations for the coefficients!)

The Dormand–Prince method

This is one of the most used ODE solvers of all time! A.k.a. `DOPRI5` in SciPy, `DP5` in Julia, or `ode45` in MATLAB

The Dormand–Prince method

This is one of the most used ODE solvers of all time! A.k.a. DOPRI5 in SciPy, DP5 in Julia, or ode45 in MATLAB

DOPRI5 has a much more complicated Butcher tableau:

0							
$\frac{1}{5}$	$\frac{1}{5}$						
$\frac{3}{10}$	$\frac{3}{40}$	$\frac{9}{40}$					
$\frac{4}{5}$	$\frac{44}{45}$	$-\frac{56}{15}$	$\frac{32}{9}$				
$\frac{8}{9}$	$\frac{19372}{6561}$	$-\frac{25360}{2187}$	$\frac{64448}{6561}$	$-\frac{212}{729}$			
1	$\frac{9017}{3168}$	$-\frac{355}{33}$	$\frac{46732}{5247}$	$\frac{49}{176}$	$-\frac{5103}{18656}$		
$\frac{1}{2}$	$\frac{35}{384}$	0	$\frac{500}{1113}$	$\frac{125}{192}$	$-\frac{2187}{6784}$	$\frac{11}{84}$	
	$\frac{35}{384}$	0	$\frac{500}{1113}$	$\frac{125}{192}$	$-\frac{2187}{6784}$	0	$\frac{11}{84}$
	$\frac{5179}{57600}$	0	$\frac{7571}{16695}$	$\frac{393}{640}$	$-\frac{92097}{339200}$	$\frac{187}{2100}$	$\frac{1}{40}$

The Dormand–Prince method

This is one of the most used ODE solvers of all time! A.k.a. DOPRI5 in SciPy, DP5 in Julia, or ode45 in MATLAB

DOPRI5 has a much more complicated Butcher tableau:

0							
$\frac{1}{5}$	$\frac{1}{5}$						
$\frac{3}{10}$	$\frac{3}{40}$	$\frac{9}{40}$					
$\frac{4}{5}$	$\frac{44}{45}$	$-\frac{56}{15}$	$\frac{32}{9}$				
$\frac{8}{9}$	$\frac{19372}{6561}$	$-\frac{25360}{2187}$	$\frac{64448}{6561}$	$-\frac{212}{729}$			
1	$\frac{9017}{3168}$	$-\frac{355}{33}$	$\frac{46732}{5247}$	$\frac{49}{176}$	$-\frac{5103}{18656}$		
$\frac{1}{2}$	$\frac{35}{384}$	0	$\frac{500}{1113}$	$\frac{125}{192}$	$-\frac{2187}{6784}$	$\frac{11}{84}$	
	$\frac{35}{384}$	0	$\frac{500}{1113}$	$\frac{125}{192}$	$-\frac{2187}{6784}$	0	$\frac{11}{84}$
	$\frac{5179}{57600}$	0	$\frac{7571}{16695}$	$\frac{393}{640}$	$-\frac{92097}{339200}$	$\frac{187}{2100}$	$\frac{1}{40}$

This is the reason for the SciPy code:

```

055 ..... C2=0.200
056 ..... C3=0.300
057 ..... C4=0.800
058 ..... C5=0.00/9.00
059 ..... A21=0.200
060 ..... A31=3.00/40.00
061 ..... A32=9.00/40.00
062 ..... A41=44.00/45.00
063 ..... A42=-56.00/15.00
064 ..... A43=32.00/9.00
065 ..... A51=19372.00/6561.00
066 ..... A52=-25360.00/2187.00
067 ..... A53=64448.00/6561.00
068 ..... A54=-212.00/729.00
069 ..... A61=9017.00/3168.00
070 ..... A62=-355.00/33.00
071 ..... A63=46732.00/5247.00
072 ..... A64=49.00/176.00
073 ..... A65=-5103.00/18656.00
074 ..... A71=35.00/384.00
075 ..... A73=500.00/1113.00
076 ..... A74=125.00/192.00
077 ..... A75=-2187.00/6784.00
078 ..... A76=11.00/84.00
079 ..... E1=71.00/57600.00
080 ..... E3=-71.00/16695.00
081 ..... E4=71.00/1920.00
082 ..... E5=-17253.00/339200.00
083 ..... E6=22.00/525.00
084 ..... E7=-1.00/40.00

```

<https://github.com/scipy/scipy/blob/main/scipy/integrate/dop/dopri5.f>

The Dormand–Prince method

This is one of the most used ODE solvers of all time! A.k.a. DOPRI5 in SciPy, DP5 in Julia, or ode45 in MATLAB

DOPRI5 has a much more complicated Butcher tableau:

0								
$\frac{1}{5}$	$\frac{1}{5}$							
$\frac{3}{10}$	$\frac{3}{40}$	$\frac{9}{40}$						
$\frac{4}{5}$	$\frac{44}{45}$	$-\frac{56}{15}$	$\frac{32}{9}$					
$\frac{8}{9}$	$\frac{19372}{6561}$	$-\frac{25360}{2187}$	$\frac{64448}{6561}$	$-\frac{212}{729}$				
1	$\frac{9017}{3168}$	$-\frac{355}{33}$	$\frac{46732}{5247}$	$\frac{49}{176}$	$-\frac{5103}{18656}$			
$\frac{1}{2}$	$\frac{35}{384}$	0	$\frac{500}{1113}$	$\frac{125}{192}$	$-\frac{2187}{6784}$	$\frac{11}{84}$		
	$\frac{35}{384}$	0	$\frac{500}{1113}$	$\frac{125}{192}$	$-\frac{2187}{6784}$	0	$\frac{11}{84}$	
	$\frac{5179}{57600}$	0	$\frac{7571}{16695}$	$\frac{393}{640}$	$-\frac{92097}{339200}$	$\frac{187}{2100}$	$\frac{1}{40}$	

(the two bottom lines are there because of “error estimation” which is not covered in this lecture; if interested, check out Chapter II.4 in “Solving Ordinary Differential Equations I” by Hairer et al.)

This is the reason for the SciPy code:

```

655 ..... C2=0.200
656 ..... C3=0.300
657 ..... C4=0.800
658 ..... C5=0.00/9.00
659 ..... A21=0.200
660 ..... A31=3.00/40.00
661 ..... A32=9.00/40.00
662 ..... A41=44.00/45.00
663 ..... A42=-56.00/15.00
664 ..... A43=32.00/9.00
665 ..... A51=19372.00/6561.00
666 ..... A52=-25360.00/2187.00
667 ..... A53=64448.00/6561.00
668 ..... A54=-212.00/729.00
669 ..... A61=9017.00/3168.00
670 ..... A62=-355.00/33.00
671 ..... A63=46732.00/5247.00
672 ..... A64=49.00/176.00
673 ..... A65=-5103.00/18656.00
674 ..... A71=35.00/384.00
675 ..... A73=500.00/1113.00
676 ..... A74=125.00/192.00
677 ..... A75=-2187.00/6784.00
678 ..... A76=11.00/84.00
679 ..... E1=71.00/57600.00
680 ..... E3=-71.00/16695.00
681 ..... E4=71.00/1920.00
682 ..... E5=-17253.00/339200.00
683 ..... E6=22.00/525.00
684 ..... E7=-1.00/40.00
    
```

<https://github.com/scipy/scipy/blob/main/scipy/integrate/dop/dopri5.f>

Intermediate summary on classical numerical ODE solvers:

- ▶ **There are A LOT of numerical ODE solvers!**



Intermediate summary on classical numerical ODE solvers:

- ▶ **There are A LOT of numerical ODE solvers!** Differences between solvers include:
 - ▶ *Stability:* Explicit vs implicit methods.



Intermediate summary on classical numerical ODE solvers:

- ▶ **There are A LOT of numerical ODE solvers!** Differences between solvers include:
 - ▶ *Stability*: Explicit vs implicit methods.
 - ▶ *Order* and convergence rates:
 A Runge–Kutta method has *order* p if the local truncation error is of order $O(h^{p+1})$. Examples:



Intermediate summary on classical numerical ODE solvers:

▶ **There are A LOT of numerical ODE solvers!** Differences between solvers include:

- ▶ *Stability*: Explicit vs implicit methods.
- ▶ *Order* and convergence rates:

A Runge–Kutta method has *order* p if the local truncation error is of order $O(h^{p+1})$. Examples:

- ▶ Forward Euler: $p = 1$ (exercise sheet)



Intermediate summary on classical numerical ODE solvers:

▶ **There are A LOT of numerical ODE solvers!** Differences between solvers include:

- ▶ *Stability*: Explicit vs implicit methods.
- ▶ *Order* and convergence rates:

A Runge–Kutta method has *order* p if the local truncation error is of order $O(h^{p+1})$. Examples:

- ▶ Forward Euler: $p = 1$ (exercise sheet)
- ▶ Explicit midpoint method: $p = 2$



Intermediate summary on classical numerical ODE solvers:

▶ **There are A LOT of numerical ODE solvers!** Differences between solvers include:

- ▶ *Stability*: Explicit vs implicit methods.
- ▶ *Order* and convergence rates:

A Runge–Kutta method has *order* p if the local truncation error is of order $O(h^{p+1})$. Examples:

- ▶ Forward Euler: $p = 1$ (exercise sheet)
- ▶ Explicit midpoint method: $p = 2$
- ▶ The classical fourth-order Runge–Kutta method: $p = 4$.



Intermediate summary on classical numerical ODE solvers:

▶ **There are A LOT of numerical ODE solvers!** Differences between solvers include:

- ▶ *Stability*: Explicit vs implicit methods.
- ▶ *Order* and convergence rates:

A Runge–Kutta method has *order* p if the local truncation error is of order $O(h^{p+1})$. Examples:

- ▶ Forward Euler: $p = 1$ (exercise sheet)
- ▶ Explicit midpoint method: $p = 2$
- ▶ The classical fourth-order Runge–Kutta method: $p = 4$.
- ▶ The Dormand–Prince method: $p = 5$.



Intermediate summary on classical numerical ODE solvers:

- ▶ **There are A LOT of numerical ODE solvers!** Differences between solvers include:

- ▶ *Stability*: Explicit vs implicit methods.
- ▶ *Order* and convergence rates:

A Runge–Kutta method has *order* p if the local truncation error is of order $O(h^{p+1})$. Examples:

- ▶ Forward Euler: $p = 1$ (exercise sheet)
- ▶ Explicit midpoint method: $p = 2$
- ▶ The classical fourth-order Runge–Kutta method: $p = 4$.
- ▶ The Dormand–Prince method: $p = 5$.

- ▶ There is a lot of stuff happening under the hood when calling `scipy.integrate.ode` or similar:



Intermediate summary on classical numerical ODE solvers:

▶ **There are A LOT of numerical ODE solvers!** Differences between solvers include:

- ▶ *Stability*: Explicit vs implicit methods.
- ▶ *Order* and convergence rates:

A Runge–Kutta method has *order* p if the local truncation error is of order $O(h^{p+1})$. Examples:

- ▶ Forward Euler: $p = 1$ (exercise sheet)
- ▶ Explicit midpoint method: $p = 2$
- ▶ The classical fourth-order Runge–Kutta method: $p = 4$.
- ▶ The Dormand–Prince method: $p = 5$.

▶ There is a lot of stuff happening under the hood when calling `scipy.integrate.ode` or similar:

- ▶ *Step-size selection*: Discretize on the fly instead of using a fixed step size. (exercise sheet)



Intermediate summary on classical numerical ODE solvers:

▶ **There are A LOT of numerical ODE solvers!** Differences between solvers include:

- ▶ *Stability*: Explicit vs implicit methods.
- ▶ *Order* and convergence rates:

A Runge–Kutta method has *order* p if the local truncation error is of order $O(h^{p+1})$. Examples:

- ▶ Forward Euler: $p = 1$ (exercise sheet)
- ▶ Explicit midpoint method: $p = 2$
- ▶ The classical fourth-order Runge–Kutta method: $p = 4$.
- ▶ The Dormand–Prince method: $p = 5$.

▶ There is a lot of stuff happening under the hood when calling `scipy.integrate.ode` or similar:

- ▶ *Step-size selection*: Discretize on the fly instead of using a fixed step size. (exercise sheet)
- ▶ (sometimes) *Automatic solver selection*: Use heuristics to decide which solver to use.



Intermediate summary on classical numerical ODE solvers:


- ▶ **There are A LOT of numerical ODE solvers!** Differences between solvers include:
 - ▶ *Stability*: Explicit vs implicit methods.
 - ▶ *Order* and convergence rates:

A Runge–Kutta method has *order* p if the local truncation error is of order $O(h^{p+1})$. Examples:

 - ▶ Forward Euler: $p = 1$ (exercise sheet)
 - ▶ Explicit midpoint method: $p = 2$
 - ▶ The classical fourth-order Runge–Kutta method: $p = 4$.
 - ▶ The Dormand–Prince method: $p = 5$.
- ▶ There is a lot of stuff happening under the hood when calling `scipy.integrate.ode` or similar:
 - ▶ *Step-size selection*: Discretize on the fly instead of using a fixed step size. (exercise sheet)
 - ▶ (sometimes) *Automatic solver selection*: Use heuristics to decide which solver to use.

Next block: **What if we don't know f but instead have to estimate it from data?**

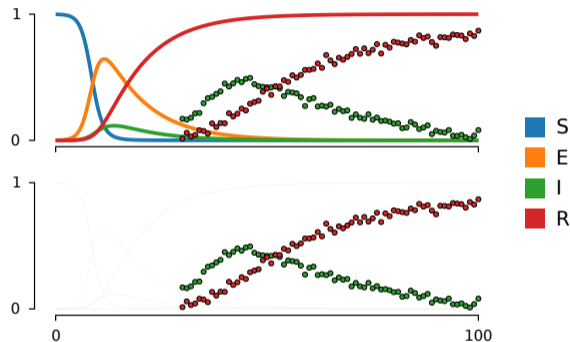




Tronarp, Bosch, Hennig, "Fenrir: Physics-Enhanced Regression for Initial Value Problems", ICML 2022

Parameter Inference

Learning unknown dynamics from data.



- ▶ **Typical goal:** “Fit the data”.
- ▶ **Parameter inference:** Learn the parameters of a *Mechanistic model*, e.g. here the SEIR model

$$\dot{S} = -\beta_E SE/N$$


$$\dot{E} = \beta_E SE/N - \gamma E$$

$$\dot{I} = \gamma E - \lambda I$$

$$\dot{R} = \lambda I$$

such that the solution fits the data.

Tronarp, Bosch, Hennig, “Fenrir: Physics-Enhanced Regression for Initial Value Problems”, ICML 2022



Tronarp, Bosch, Hennig, "Fenrir: Physics-Enhanced Regression for Initial Value Problems", ICML 2022

- ▶ **Typical goal:** "Fit the data".
- ▶ **Parameter inference:** Learn the parameters of a *Mechanistic model*, e.g. here the SEIR model

$$\dot{S} = -\beta_E SE/N$$

$$\dot{E} = \beta_E SE/N - \gamma E$$


$$\dot{I} = \gamma E - \lambda I$$

$$\dot{R} = \lambda I$$

such that the solution fits the data.
⇒ *Provides interpretable results*

Parameter Inference

Learning unknown dynamics from data.



- ▶ **Typical goal:** “Fit the data”.
- ▶ **Parameter inference:** Learn the parameters of a *Mechanistic model*, e.g. here the SEIR model

$$\dot{S} = -\beta_E SE/N$$

$$\dot{E} = \beta_E SE/N - \gamma E$$

$$\dot{I} = \gamma E - \lambda I$$

$$\dot{R} = \lambda I$$

Tronarp, Bosch, Hennig, “Fenrir: Physics-Enhanced Regression for Initial Value Problems”, ICML 2022

such that the solution fits the data.
⇒ *Provides interpretable results*

Up to this point f was always given – *now it needs to be estimated!*

Setup: Consider an initial value problem

$$\dot{x}(t) = f(x(t), t, \theta), \quad x(0) = x_0(\theta), \quad t \in [0, T],$$

where the parameters $\theta \in \mathbb{R}^d$ are unknown.

Setup: Consider an initial value problem

$$\dot{x}(t) = f(x(t), t, \theta), \quad x(0) = x_0(\theta), \quad t \in [0, T],$$

where the parameters $\theta \in \mathbb{R}^d$ are unknown.

Assume noisy observations $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where

$$y_i = H \cdot x(t_i) + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, \Sigma).$$

Setup: Consider an initial value problem

$$\dot{x}(t) = f(x(t), t, \theta), \quad x(0) = x_0(\theta), \quad t \in [0, T],$$

where the parameters $\theta \in \mathbb{R}^d$ are unknown.

Assume noisy observations $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where

$$y_i = H \cdot x(t_i) + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, \Sigma).$$

Goal: Estimate θ from \mathcal{D} :

$$p(\theta | \mathcal{D}) = \frac{p(\mathcal{D} | \theta)p(\theta)}{p(\mathcal{D})}.$$

Setup: Consider an initial value problem

$$\dot{x}(t) = f(x(t), t, \theta), \quad x(0) = x_0(\theta), \quad t \in [0, T],$$

where the parameters $\theta \in \mathbb{R}^d$ are unknown.

Assume noisy observations $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where

$$y_i = H \cdot x(t_i) + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, \Sigma).$$

Goal: Estimate θ from \mathcal{D} :

$$p(\theta | \mathcal{D}) = \frac{p(\mathcal{D} | \theta)p(\theta)}{p(\mathcal{D})}.$$

Cheaper goal: Compute the *maximum-likelihood estimate*

$$\hat{\theta}_{\text{ML}} = \arg \max_{\theta} p(\mathcal{D} | \theta).$$

Assuming i.i.d. data, the likelihood is given by

$$p(\mathcal{D} \mid \theta) = \prod_{i=1}^n \mathcal{N}(y_i; Hx_\theta(t_i), \Sigma).$$

Issue: The likelihood is intractable, since it depends on the *true* solution $x_\theta(t)$ of the ODE.

Assuming i.i.d. data, the likelihood is given by

$$p(\mathcal{D} \mid \theta) = \prod_{i=1}^n \mathcal{N}(y_i; Hx_\theta(t_i), \Sigma).$$

Issue: The likelihood is intractable, since it depends on the *true* solution $x_\theta(t)$ of the ODE.

Solution: Use a numerical ODE solver to approximate the solution:

$$x_\theta(t) \approx \hat{x}_\theta(t),$$

where $\hat{x}_\theta(t)$ is the numerical solution of the ODE with parameters θ .

Assuming i.i.d. data, the likelihood is given by

$$p(\mathcal{D} \mid \theta) = \prod_{i=1}^n \mathcal{N}(y_i; Hx_\theta(t_i), \Sigma).$$

Issue: The likelihood is intractable, since it depends on the *true* solution $x_\theta(t)$ of the ODE.

Solution: Use a numerical ODE solver to approximate the solution:

$$x_\theta(t) \approx \hat{x}_\theta(t),$$

where $\hat{x}_\theta(t)$ is the numerical solution of the ODE with parameters θ .

Then

$$p(\mathcal{D} \mid \theta) \approx \prod_{i=1}^n \mathcal{N}(y_i; H\hat{x}_\theta(t_i), \Sigma),$$

which is tractable.

Assuming i.i.d. data, the likelihood is given by

$$p(\mathcal{D} \mid \theta) = \prod_{i=1}^n \mathcal{N}(y_i; Hx_\theta(t_i), \Sigma).$$

Issue: The likelihood is intractable, since it depends on the *true* solution $x_\theta(t)$ of the ODE.

Solution: Use a numerical ODE solver to approximate the solution:

$$x_\theta(t) \approx \hat{x}_\theta(t),$$

where $\hat{x}_\theta(t)$ is the numerical solution of the ODE with parameters θ .

Then

$$p(\mathcal{D} \mid \theta) \approx \prod_{i=1}^n \mathcal{N}(y_i; H\hat{x}_\theta(t_i), \Sigma),$$

which is tractable.

Maximizing the likelihood is equivalent to minimizing the *negative log-likelihood*:


$$L(\theta) = \frac{1}{2} \sum_{i=1}^n (H\hat{x}_\theta(t_i) - y_i)^T \Sigma^{-1} (H\hat{x}_\theta(t_i) - y_i).$$

Example: Simulating epidemics with the SIRD model

► **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.



Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.





Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.



Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.





Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.
- **Optimize:** with the optimizer of your choice (e.g. L-BFGS)





Parameter Inference with numerical ODE solvers

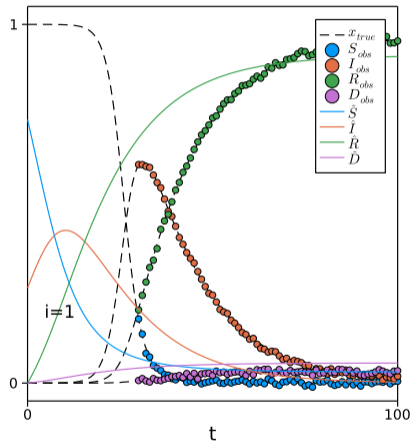
Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.
- **Optimize:** with the optimizer of your choice (e.g. L-BFGS)



Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.
- **Optimize:** with the optimizer of your choice (e.g. L-BFGS)



Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.
- **Optimize:** with the optimizer of your choice (e.g. L-BFGS)





Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.
- **Optimize:** with the optimizer of your choice (e.g. L-BFGS)





Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.
- **Optimize:** with the optimizer of your choice (e.g. L-BFGS)





Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.
- **Optimize:** with the optimizer of your choice (e.g. L-BFGS)





Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.
- **Optimize:** with the optimizer of your choice (e.g. L-BFGS)



Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.
- **Optimize:** with the optimizer of your choice (e.g. L-BFGS)





Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.
- **Optimize:** with the optimizer of your choice (e.g. L-BFGS)





Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.
- **Optimize:** with the optimizer of your choice (e.g. L-BFGS)





Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.
- **Optimize:** with the optimizer of your choice (e.g. L-BFGS)





Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.
- **Optimize:** with the optimizer of your choice (e.g. L-BFGS)





Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.
- **Optimize:** with the optimizer of your choice (e.g. L-BFGS)





Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.
- **Optimize:** with the optimizer of your choice (e.g. L-BFGS)





Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.
- **Optimize:** with the optimizer of your choice (e.g. L-BFGS)



Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.
- **Optimize:** with the optimizer of your choice (e.g. L-BFGS)
- **Result:** $\hat{\theta}_{ML} = (9.46 \cdot 10^{-6}, 0.5014, 0.0603, 0.0020)$. (took ~ 3.5 seconds)





Parameter Inference with numerical ODE solvers


Example: Simulating epidemics with the SIRD model

- **Dynamics:** (simplified) SIRD model

$$\dot{S} = -\beta SI \quad \dot{I} = \beta SI - \gamma I - \eta I \quad \dot{R} = \gamma I \quad \dot{D} = \eta I$$

for $t \in [0, 100]$, with $SIRD(0) = [1 - I_0, I_0, 0, 0]$ and four unknown parameters $I_0, \beta, \gamma, \eta \in \mathbb{R}$.

- **Data:** $\mathcal{D} = \{(y_i, t_i)\}_{i=1}^n$, where $y_i \sim \mathcal{N}(x(t_i), 0.1 \cdot I)$; generated with $\theta^* = (10^{-5}, 0.5, 0.06, 0.002)$.
- **Loss** (as in last slide): $L(\theta) = \frac{1}{2} 0.1 \sum_{i=1}^n \|\hat{x}_\theta(t_i) - y_i\|_2^2$.
- **Initial guess:** $\theta_0 = (0.1, 0.1, 0.1, 0.1)$.
- **Optimize:** with the optimizer of your choice (e.g. L-BFGS)
- **Result:** $\hat{\theta}_{ML} = (9.46 \cdot 10^{-6}, 0.5014, 0.0603, 0.0020)$. (took ~ 3.5 seconds)



We can learn system parameters from data via (local) optimization!

You saw this example in the first lecture - so let's revisit it!

Figure from: Schmidt, Krämer, Hennig, NeurIPS2021

ODE dynamics as before, but this time with **time-varying contact rate** $\beta(t)$:

$$\dot{S} = -\beta(t)SI/N, \quad \dot{I} = \beta(t)SI/N - \gamma I - \eta I, \quad \dot{R} = \gamma I, \quad \dot{D} = \eta I.$$

Data are the real COVID counts from Germany.


You saw this example in the first lecture - so let's revisit it!

Figure from: Schmidt, Krämer, Hennig, NeurIPS2021

ODE dynamics as before, but this time with **time-varying contact rate $\beta(t)$** :

$$\dot{S} = -\beta(t)SI/N, \quad \dot{I} = \beta(t)SI/N - \gamma I - \eta I, \quad \dot{R} = \gamma I, \quad \dot{D} = \eta I.$$

Data are the real COVID counts from Germany. **Related result** as shown in lecture 1:





Parameter Inference on real COVID data

You saw this example in the first lecture - so let's revisit it!

Figure from: Schmidt, Krämer, Hennig, NeurIPS2021


ODE dynamics as before, but this time with time-varying contact rate $\beta(t)$:

$$\dot{S} = -\beta(t)SI/N, \quad \dot{I} = \beta(t)SI/N - \gamma I - \eta I, \quad \dot{R} = \gamma I, \quad \dot{D} = \eta I.$$

Data are the real COVID counts from Germany.

Idea today: Just model $\beta(t)$ with a neural network β_{θ}^{NN} , and do parameter inference on θ as before!

Result:





Parameter Inference on real COVID data

You saw this example in the first lecture - so let's revisit it!

Figure from: Schmidt, Krämer, Hennig, NeurIPS2021

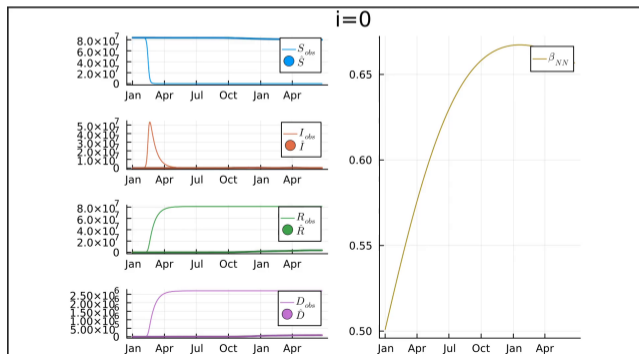
ODE dynamics as before, but this time with time-varying contact rate $\beta(t)$:

$$\dot{S} = -\beta(t)SI/N, \quad \dot{I} = \beta(t)SI/N - \gamma I - \eta I, \quad \dot{R} = \gamma I, \quad \dot{D} = \eta I.$$

Data are the real COVID counts from Germany.

Idea today: Just model $\beta(t)$ with a neural network β_{θ}^{NN} , and do parameter inference on θ as before!

Result:





Parameter Inference on real COVID data

You saw this example in the first lecture - so let's revisit it!

Figure from: Schmidt, Krämer, Hennig, NeurIPS2021


ODE dynamics as before, but this time with time-varying contact rate $\beta(t)$:

$$\dot{S} = -\beta(t)SI/N, \quad \dot{I} = \beta(t)SI/N - \gamma I - \eta I, \quad \dot{R} = \gamma I, \quad \dot{D} = \eta I.$$

Data are the real COVID counts from Germany.

Idea today: Just model $\beta(t)$ with a neural network β_{θ}^{NN} , and do parameter inference on θ as before!

Result:



Disclaimer: I only had limited time and it might very well be possible to do this much better!

Parameter Inference *on real COVID data*

You saw this example in the first lecture - so let's revisit it!


Figure from: Schmidt, Krämer, Hennig, NeurIPS2021

ODE dynamics as before, but this time with **time-varying contact rate $\beta(t)$** :

$$\dot{S} = -\beta(t)SI/N, \quad \dot{I} = \beta(t)SI/N - \gamma I - \eta I, \quad \dot{R} = \gamma I, \quad \dot{D} = \eta I.$$

Data are the real COVID counts from Germany.

Next week: $\beta(t) \sim \mathcal{GP}$!



Summary

- ▶ ODEs play an important role in machine learning.
- ▶ In general, solving an ODE requires a *numerical* solver, e.g. Euler or Runge–Kutta
- ▶ ... of which there are many! With different properties (stability, order, ...).
- ▶ We can *learn* ODE parameters via (local) optimization, even neural networks!

Please cite this course, as

```
@techreport{NoML22,
  title = {Numerics of Machine Learning},
  author = {N. Bosch and J. Grosse
and P. Hennig and A. Kristiadi
and M. Pförtner and J. Schmidt
and F. Schneider and L. Tatzel
and J. Wenger},
  series = {Lecture Notes in Machine Learning},
  year = {2022},
  institution = {Tübingen AI Center},
}
```



Summary

- ▶ ODEs play an important role in machine learning.
- ▶ In general, solving an ODE requires a *numerical* solver, e.g. Euler or Runge–Kutta
- ▶ ... of which there are many! With different properties (stability, order, ...).
- ▶ We can *learn* ODE parameters via (local) optimization, even neural networks!

Please cite this course, as

```
@techreport{NoML22,
  title = {Numerics of Machine Learning},
  author = {N. Bosch and J. Grosse
and P. Hennig and A. Kristiadi
and M. Pförtner and J. Schmidt
and F. Schneider and L. Tatzel
and J. Wenger},
  series = {Lecture Notes in Machine Learning},
  year = {2022},
  institution = {Tübingen AI Center},
}
```

Next week: ***Probabilistic*** numerical ODE solvers!
Combining ODEs and Bayesian state estimation.

