



Data-Driven Dynamics: Learning Vector Fields for Epidemic Forecasting with Neural ODEs

Vasilis Tsilidis, Department of Mathematics, University of Patras

PCMA-2025

20 December, 2025

Mechanistic

$$\text{First principles} \Rightarrow \begin{cases} \frac{dS}{dt} = -\frac{\beta SI}{N} \\ \frac{dI}{dt} = \frac{\beta SI}{N} \end{cases}$$

$\beta, \frac{\beta SI}{N}$
correspond
to real
processes

Unknown
mechanisms or
modeling
simplifications

\Rightarrow

\Rightarrow

Allows for
explanation
and verification
of assumptions

High bias

Data Driven

$$\text{Data} \Rightarrow f(\mathbf{x}) = \sigma(\mathbf{Wx} + \mathbf{b})$$

Availability of
high-quality
data

Overreliance on
data

\Rightarrow

\Rightarrow

Capable of
capturing
complex,
nonlinear
relations

No intrinsic
way of explaining
mechanisms

Hybrid

$$\frac{dz}{dt} = f_\theta(z(t), t), \quad \text{where} \quad f_\theta(\mathbf{x}) = \sigma(\mathbf{Wx} + \mathbf{b})$$

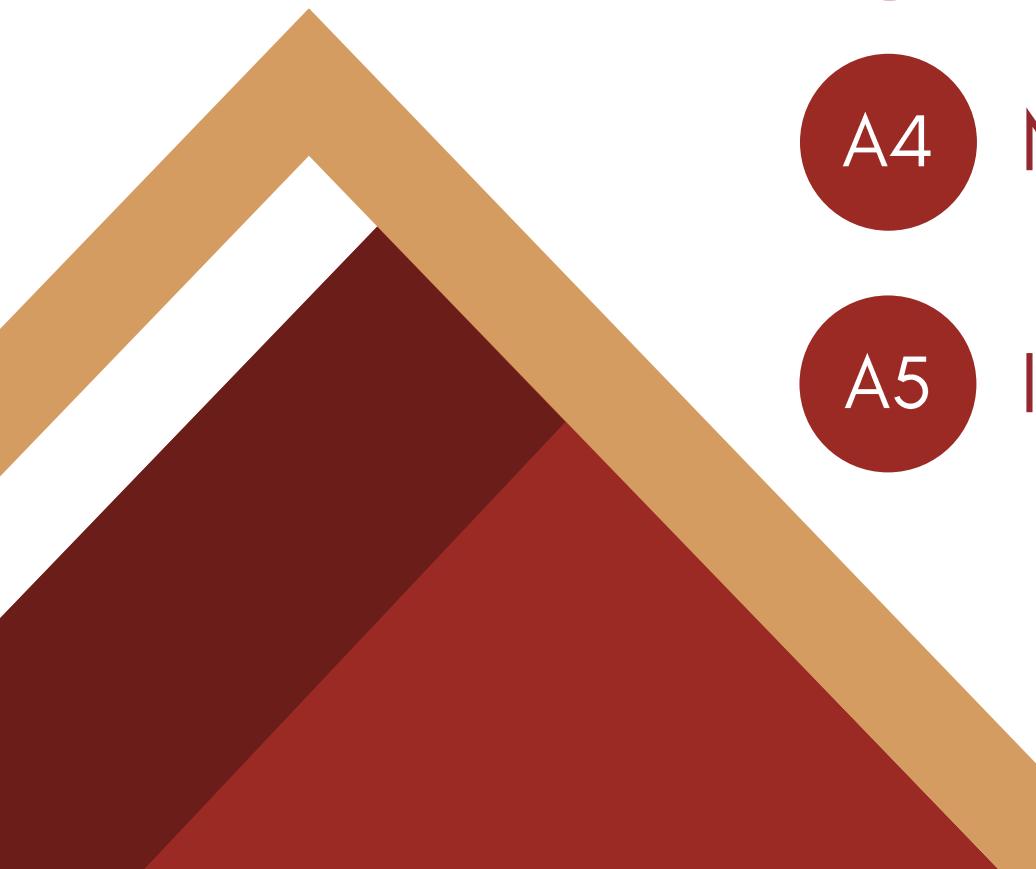
Agenda Overview

A Theory

- A1 MLP
- A2 ResNet
- A3 The Bridge
- A4 Neural ODEs
- A5 Intuition

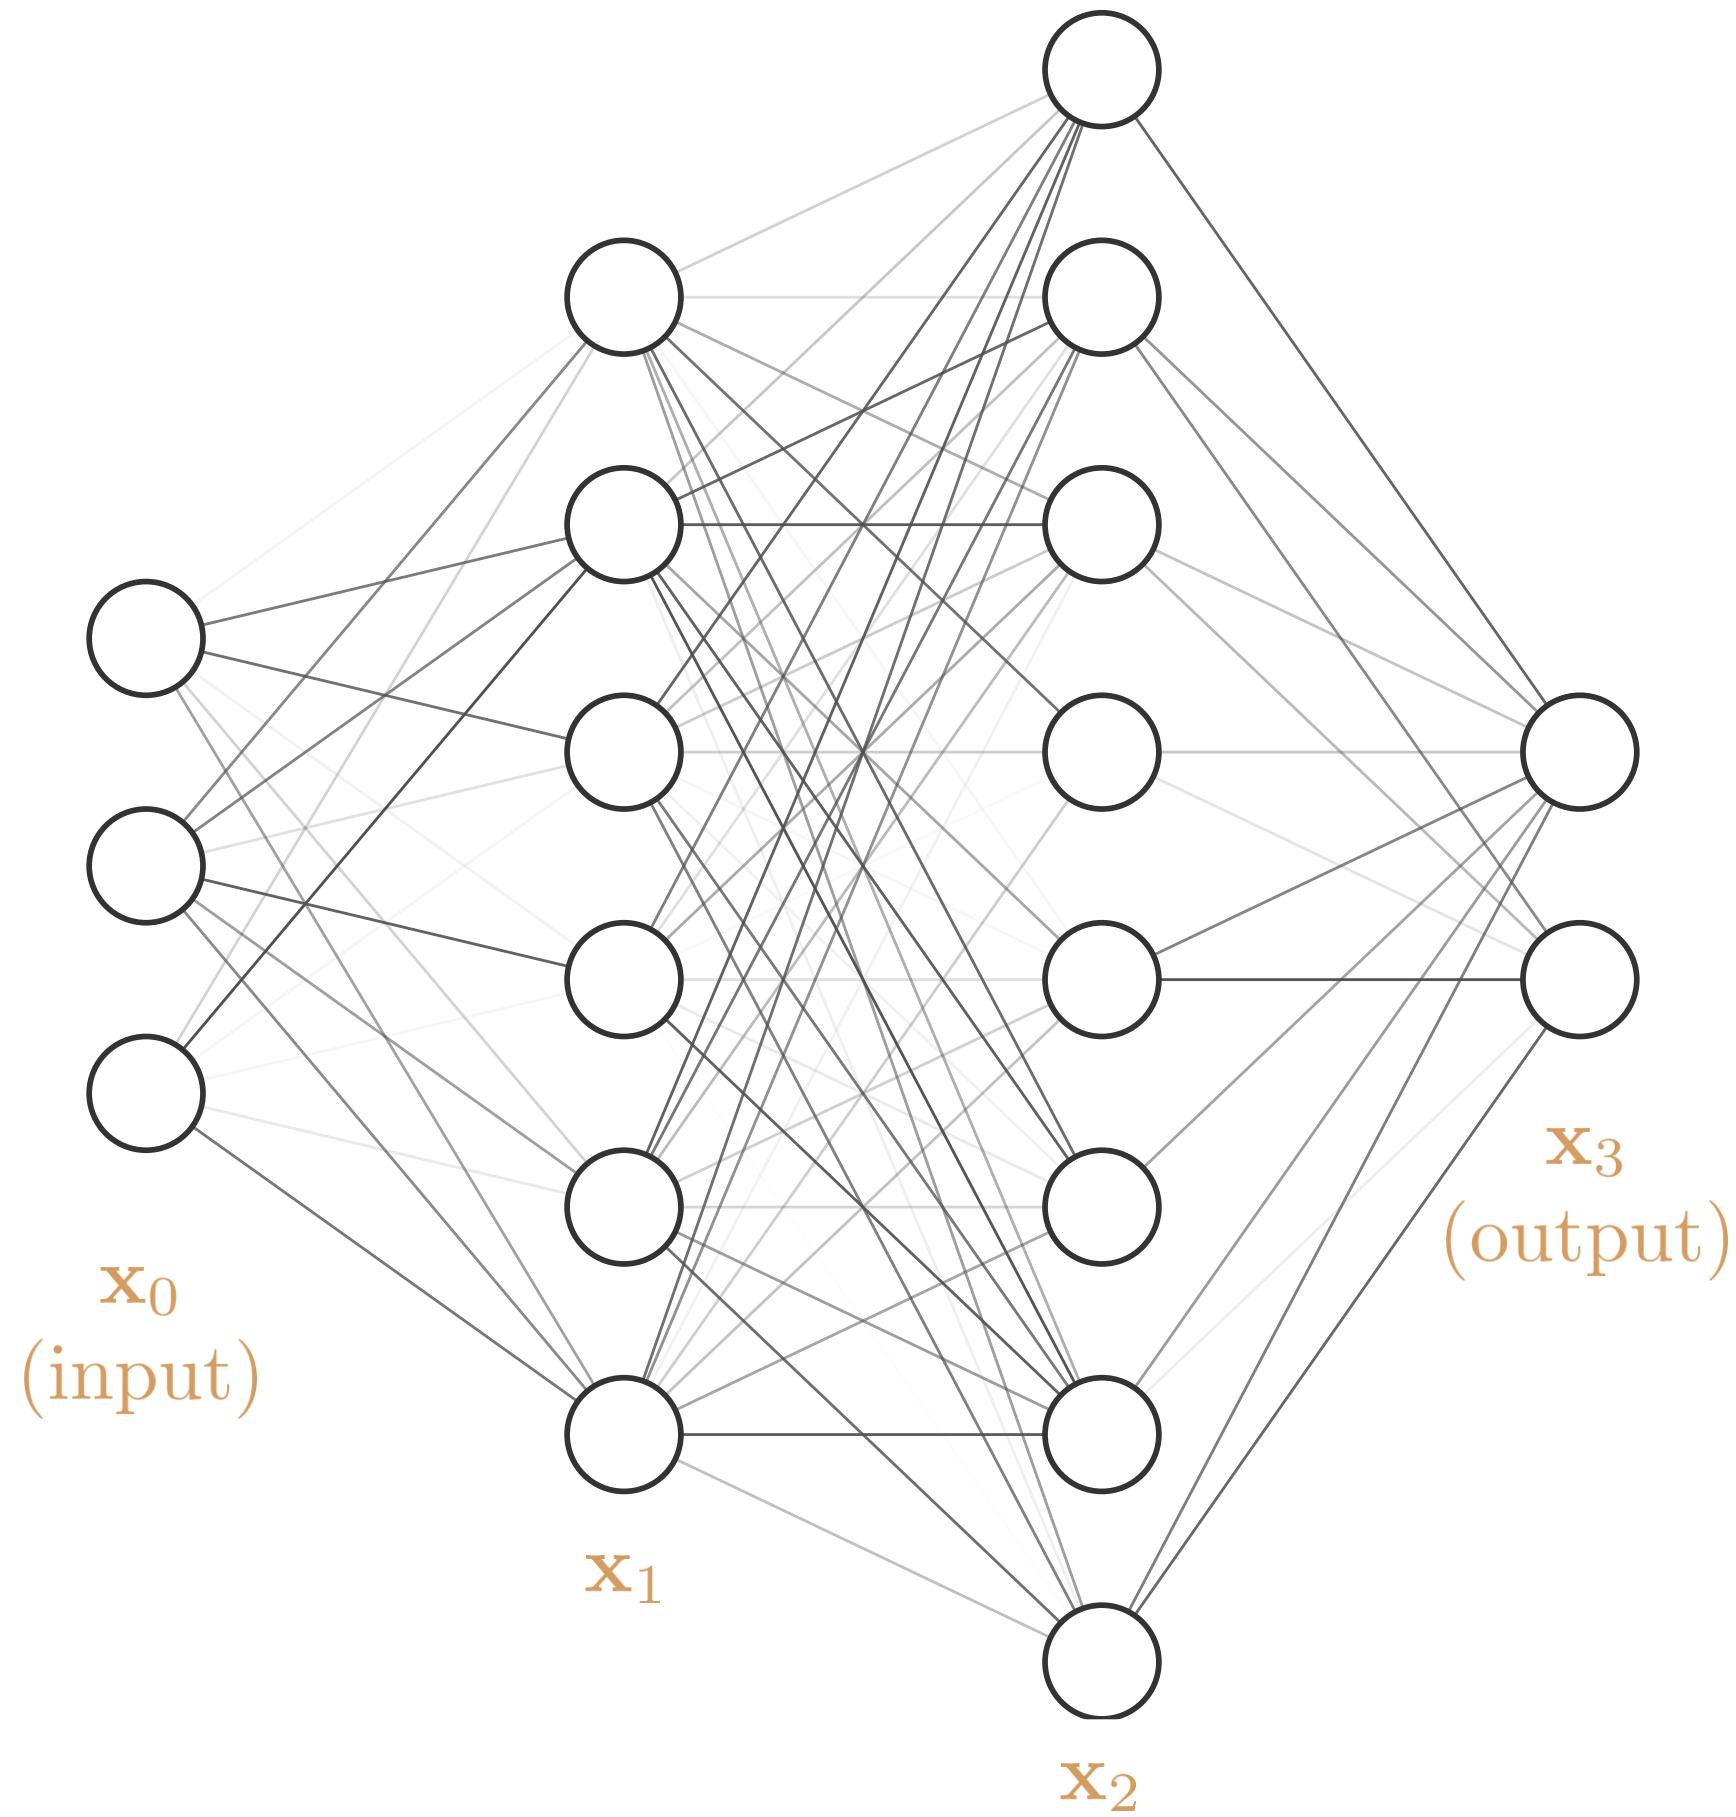
B Epidemic Forecasting

- B1 Data
- B2 Forecasting
- B3 Interpolation
- B4 Results
- B5 Takeaways



A | Theory

The Multilayer Perceptron (MLP)



$$\mathbf{x}_{t+1} = \sigma(\mathbf{W}_t \mathbf{x}_t + \mathbf{b}_t), \quad 0 \leq t \leq L - 1$$

$$\parallel \\ g_t(\mathbf{x}_t)$$

Where:

- $L + 1$ is the number of layers

- $\mathbf{x}_t = \begin{cases} \text{input vector,} & t = 0 \\ \text{hidden vector,} & 0 < t < L, \\ \text{output vector,} & x = L \end{cases}$

- $\mathbf{W}_t, \mathbf{b}_t$ are learnable parameters
- σ is a nonlinear function

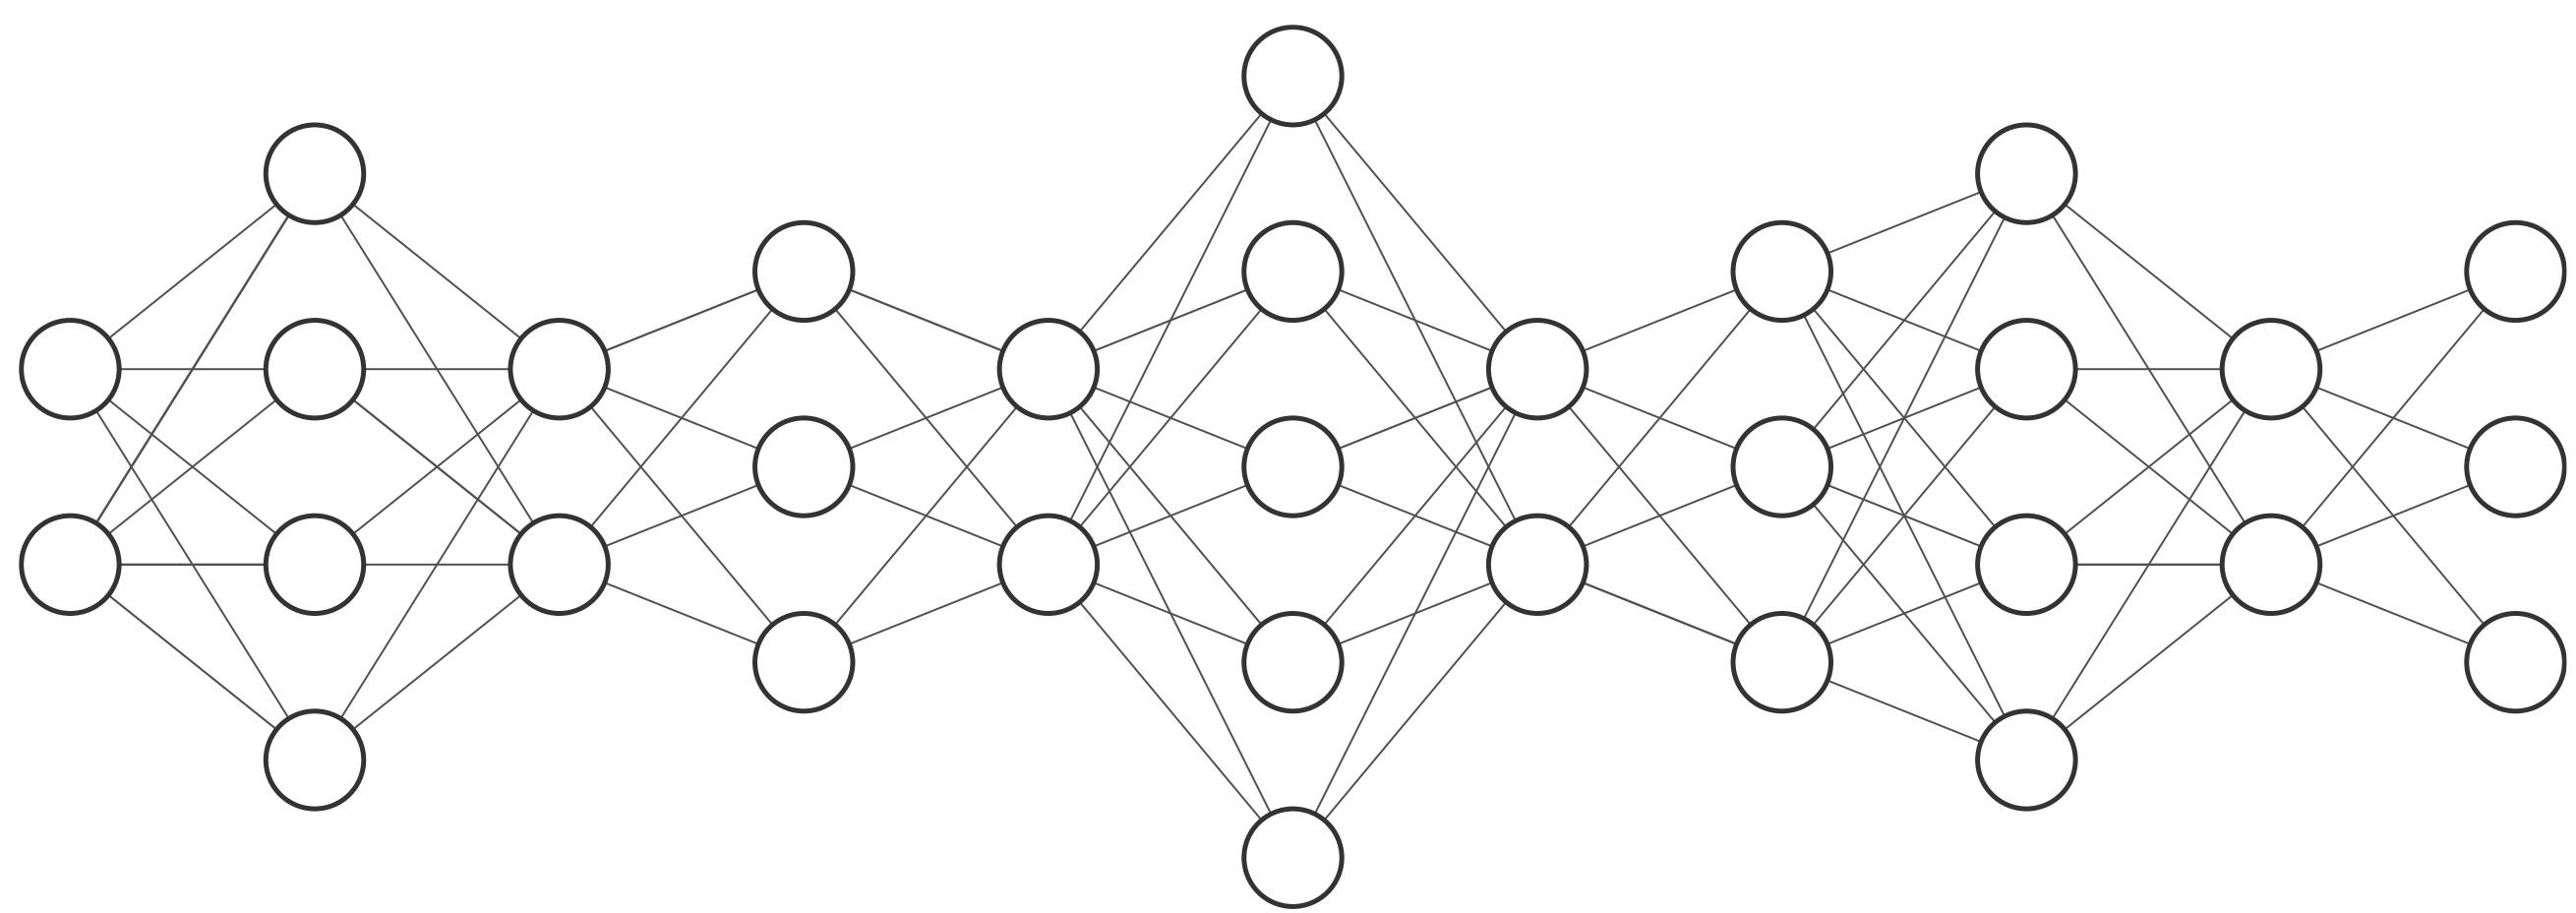
For the output to be produced, an input vector is passed through a sequence of discrete layers.

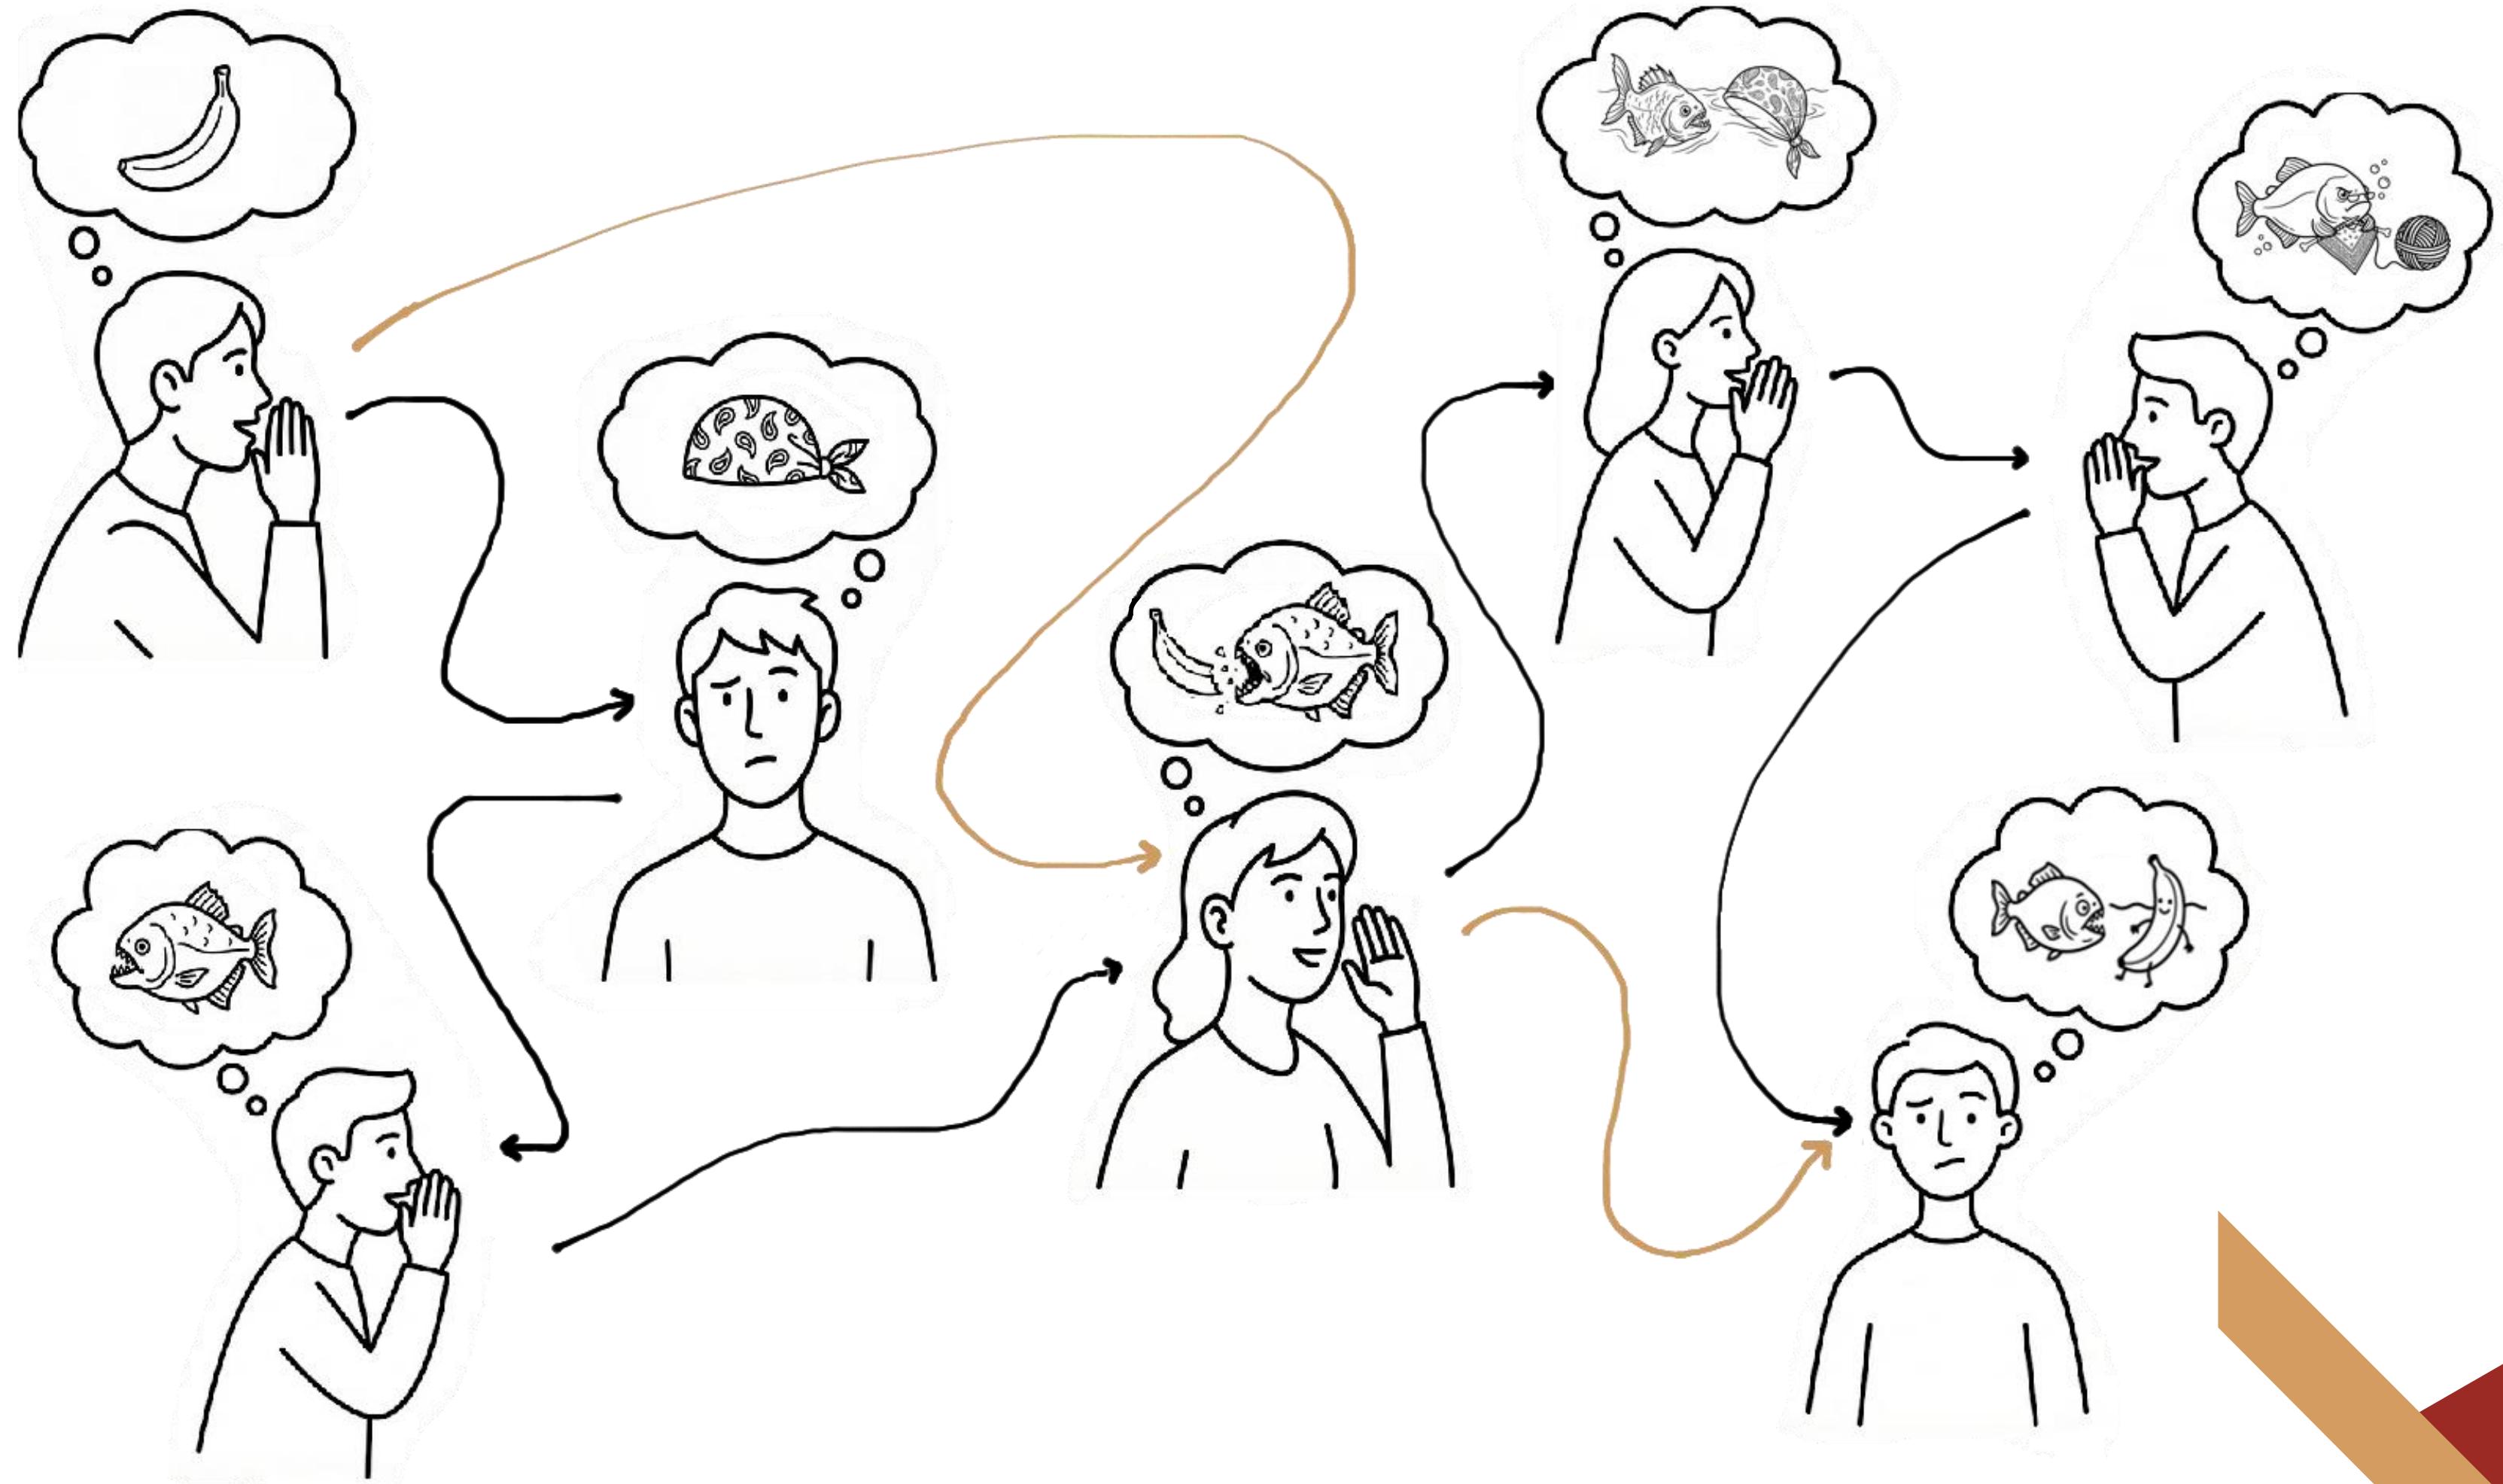
The output of each layer is recursively produced, using the output of the previous layer as input.

Each layer applies an affine transformation followed by a nonlinear transformation.

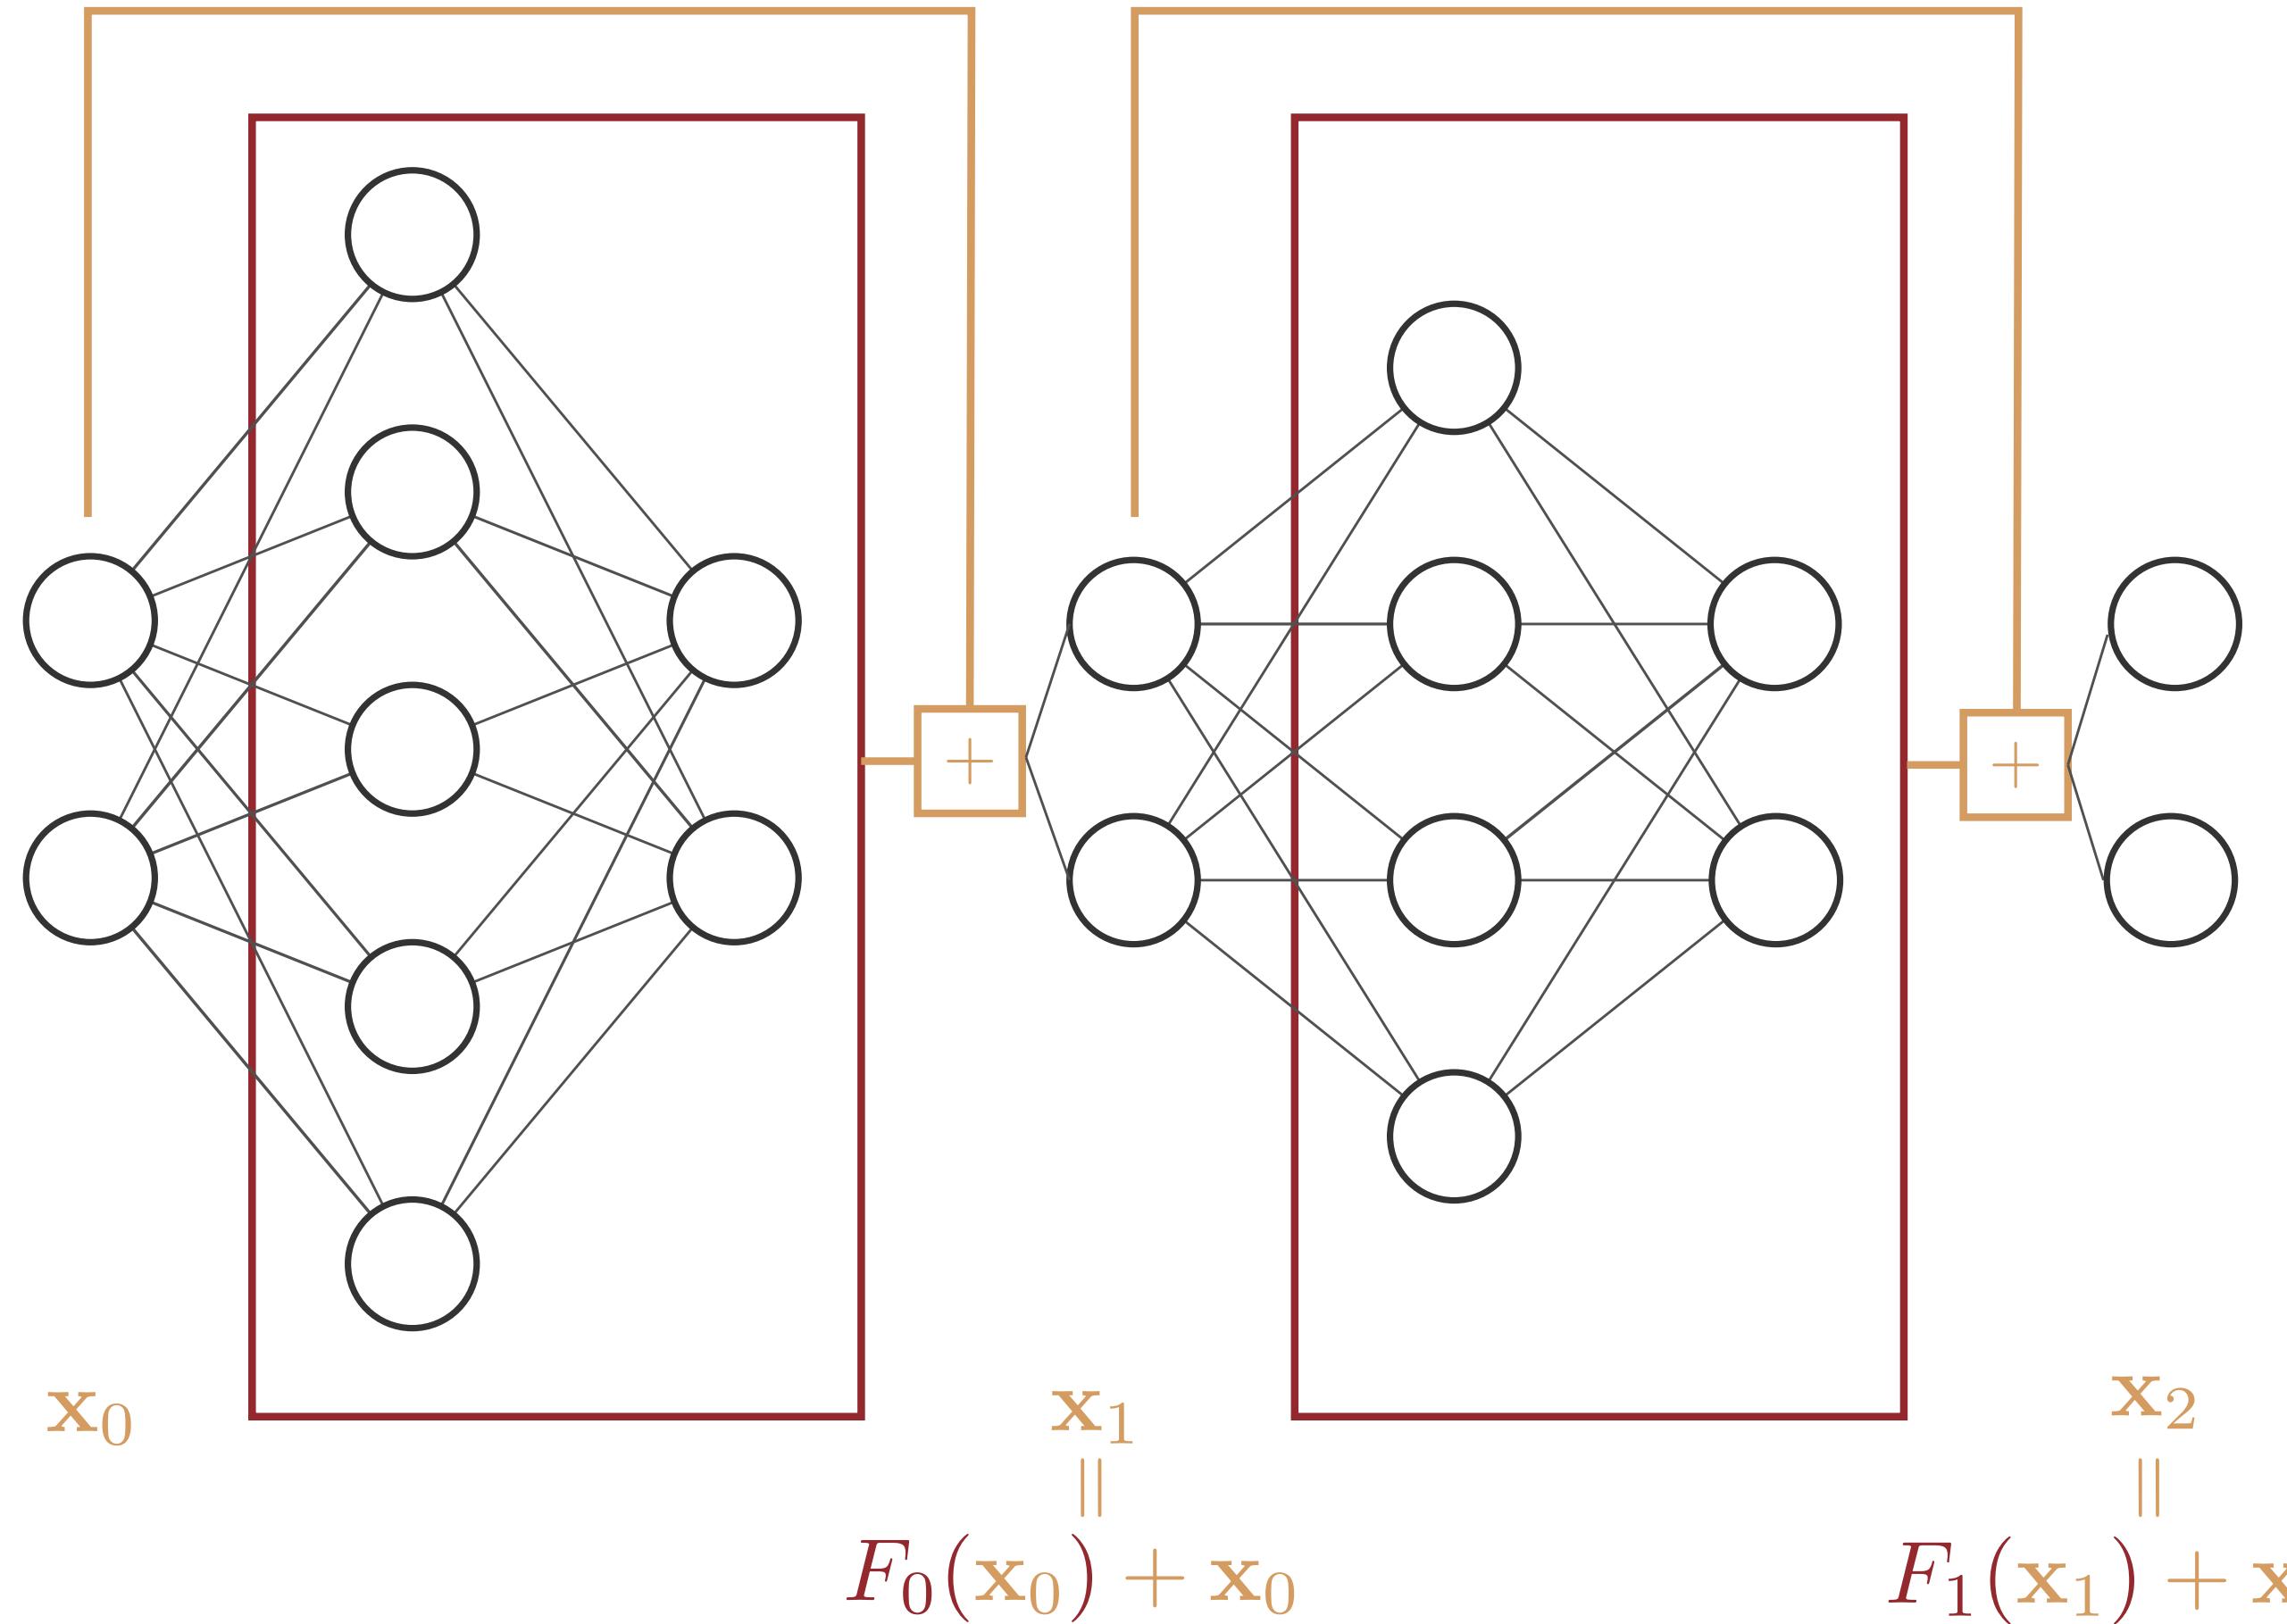
Hence, an MLP is a composition of functions:

$$f(\mathbf{x}) = g_{L-1} \circ g_{L-2} \circ \cdots \circ g_0(\mathbf{x})$$





The Residual Network (ResNet)



$$\mathbf{x}_{t+1} = F_t(\mathbf{x}_t) + \mathbf{x}_t, \quad 0 \leq t \leq L - 1$$

Where:

- $F_t : \mathbb{R}^d \mapsto \mathbb{R}^d$ are MLPs
- $\mathbf{x}_t \in \mathbb{R}^d$ are state vectors

The skip connection is modeled through a small algebraic change in the recurrence relation. Instead of learning the new state from scratch, ResNet learns the additive change to the current state. This allows the information to propagate easier.

The Bridge

Euler Method

To solve the IVP $\frac{dx}{dt} = f(x, t)$, $x(0) = x_0$
we discretize using Euler's method:

$$x_{t+1} = x_t + h \cdot f(x_t, t),$$

- y_t : state at step t .
- h : step size.
- f : vector field.

$$h = 1$$


ResNet

To transform input x_0 , we apply layers:

$$x_{t+1} = x_t + F(x_t),$$

- x_t : state at layer t .
- F : MLP.



$$x(t) = x_0 + \int_0^T f(x(s), s) \, ds$$

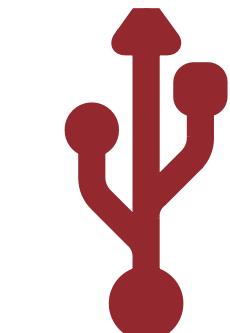
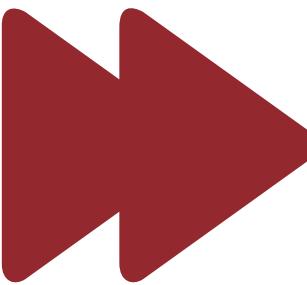
more
layers

Neural ODEs

$$\begin{cases} \frac{d\mathbf{x}}{dt} = f(\mathbf{x}(t), t) \\ \mathbf{x}(0) = \mathbf{x}_0 \end{cases}$$

Where:

- \mathbf{x}_0 is the input data
- f is an MLP
- t is a continuous depth variable



The Forward Pass

The output of the model is the trajectory of the solution of the IVP. The computation is performed using any standard numerical ODE solver

Adaptive Computation

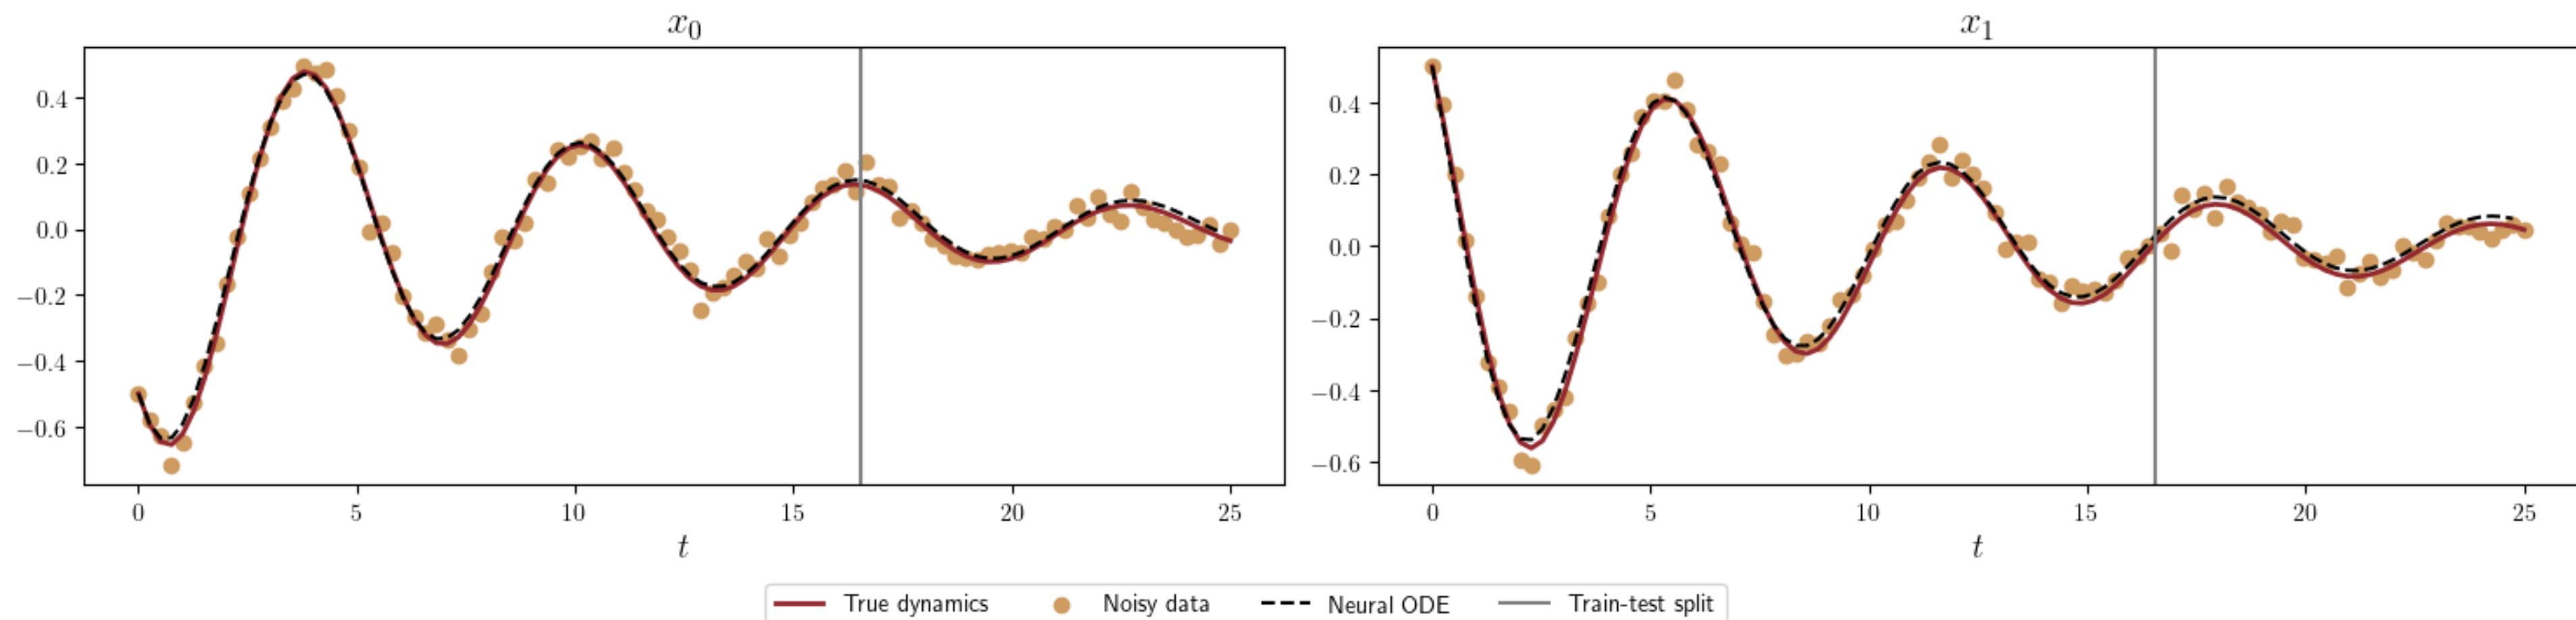
Unlike ResNets the solver can adapt the step size based on the complexity of the trajectory (error tolerance)

Topology Preservation

Assuming the MLP is Lipschitz continuous, the trajectories cannot cross (uniqueness), preserving the topology of the input space

Intuition - The Stable Spiral

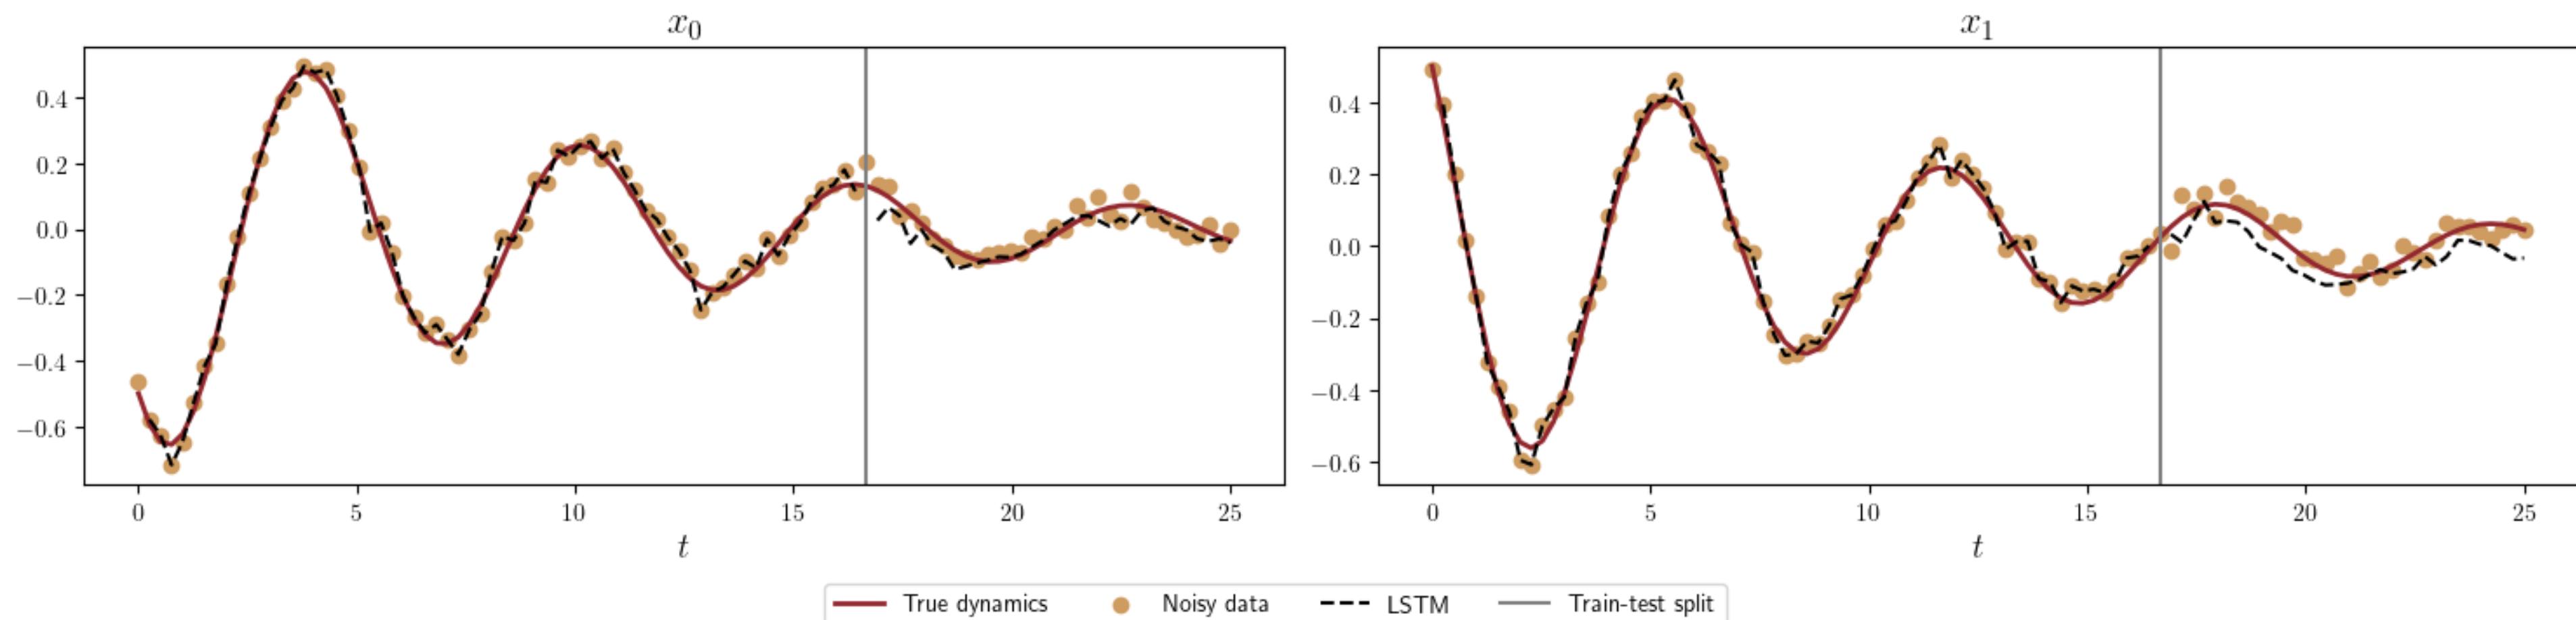
$$\begin{cases} \frac{d\mathbf{x}}{dt} = \begin{pmatrix} -0.1 & 1 \\ 1 & -1 \end{pmatrix} \mathbf{x} \\ \mathbf{x}(0) = \begin{pmatrix} -0.5 & .5 \end{pmatrix}^\top \end{cases}$$



The neural ODE learns the vector field of a stable spiral.

The trajectory is smooth, continuous, and physically plausible.

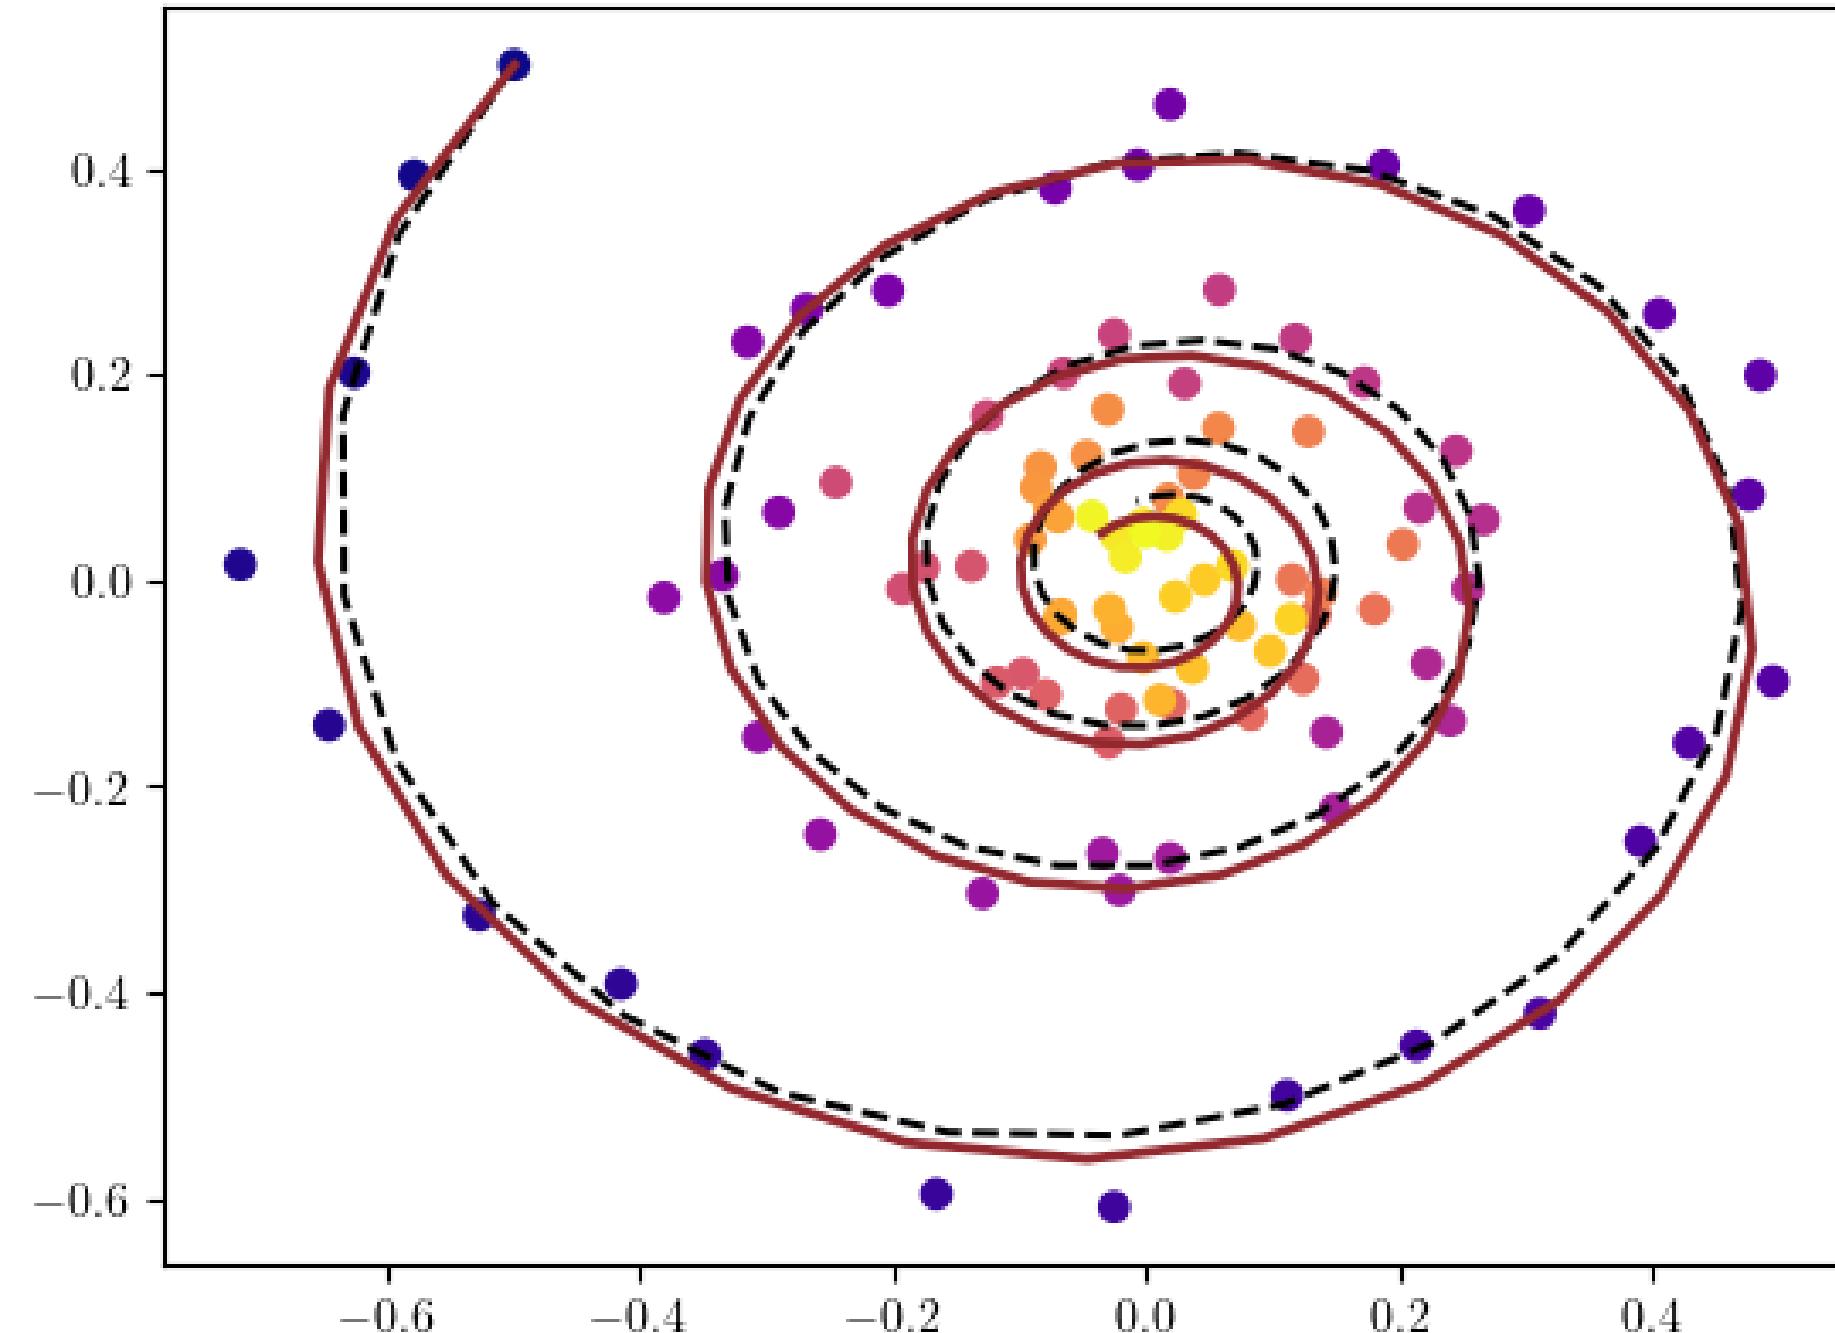
$$\begin{cases} \frac{d\mathbf{x}}{dt} = \begin{pmatrix} -0.1 & 1 \\ 1 & -1 \end{pmatrix} \mathbf{x} \\ \mathbf{x}(0) = \begin{pmatrix} -0.5 & .5 \end{pmatrix}^\top \end{cases}$$



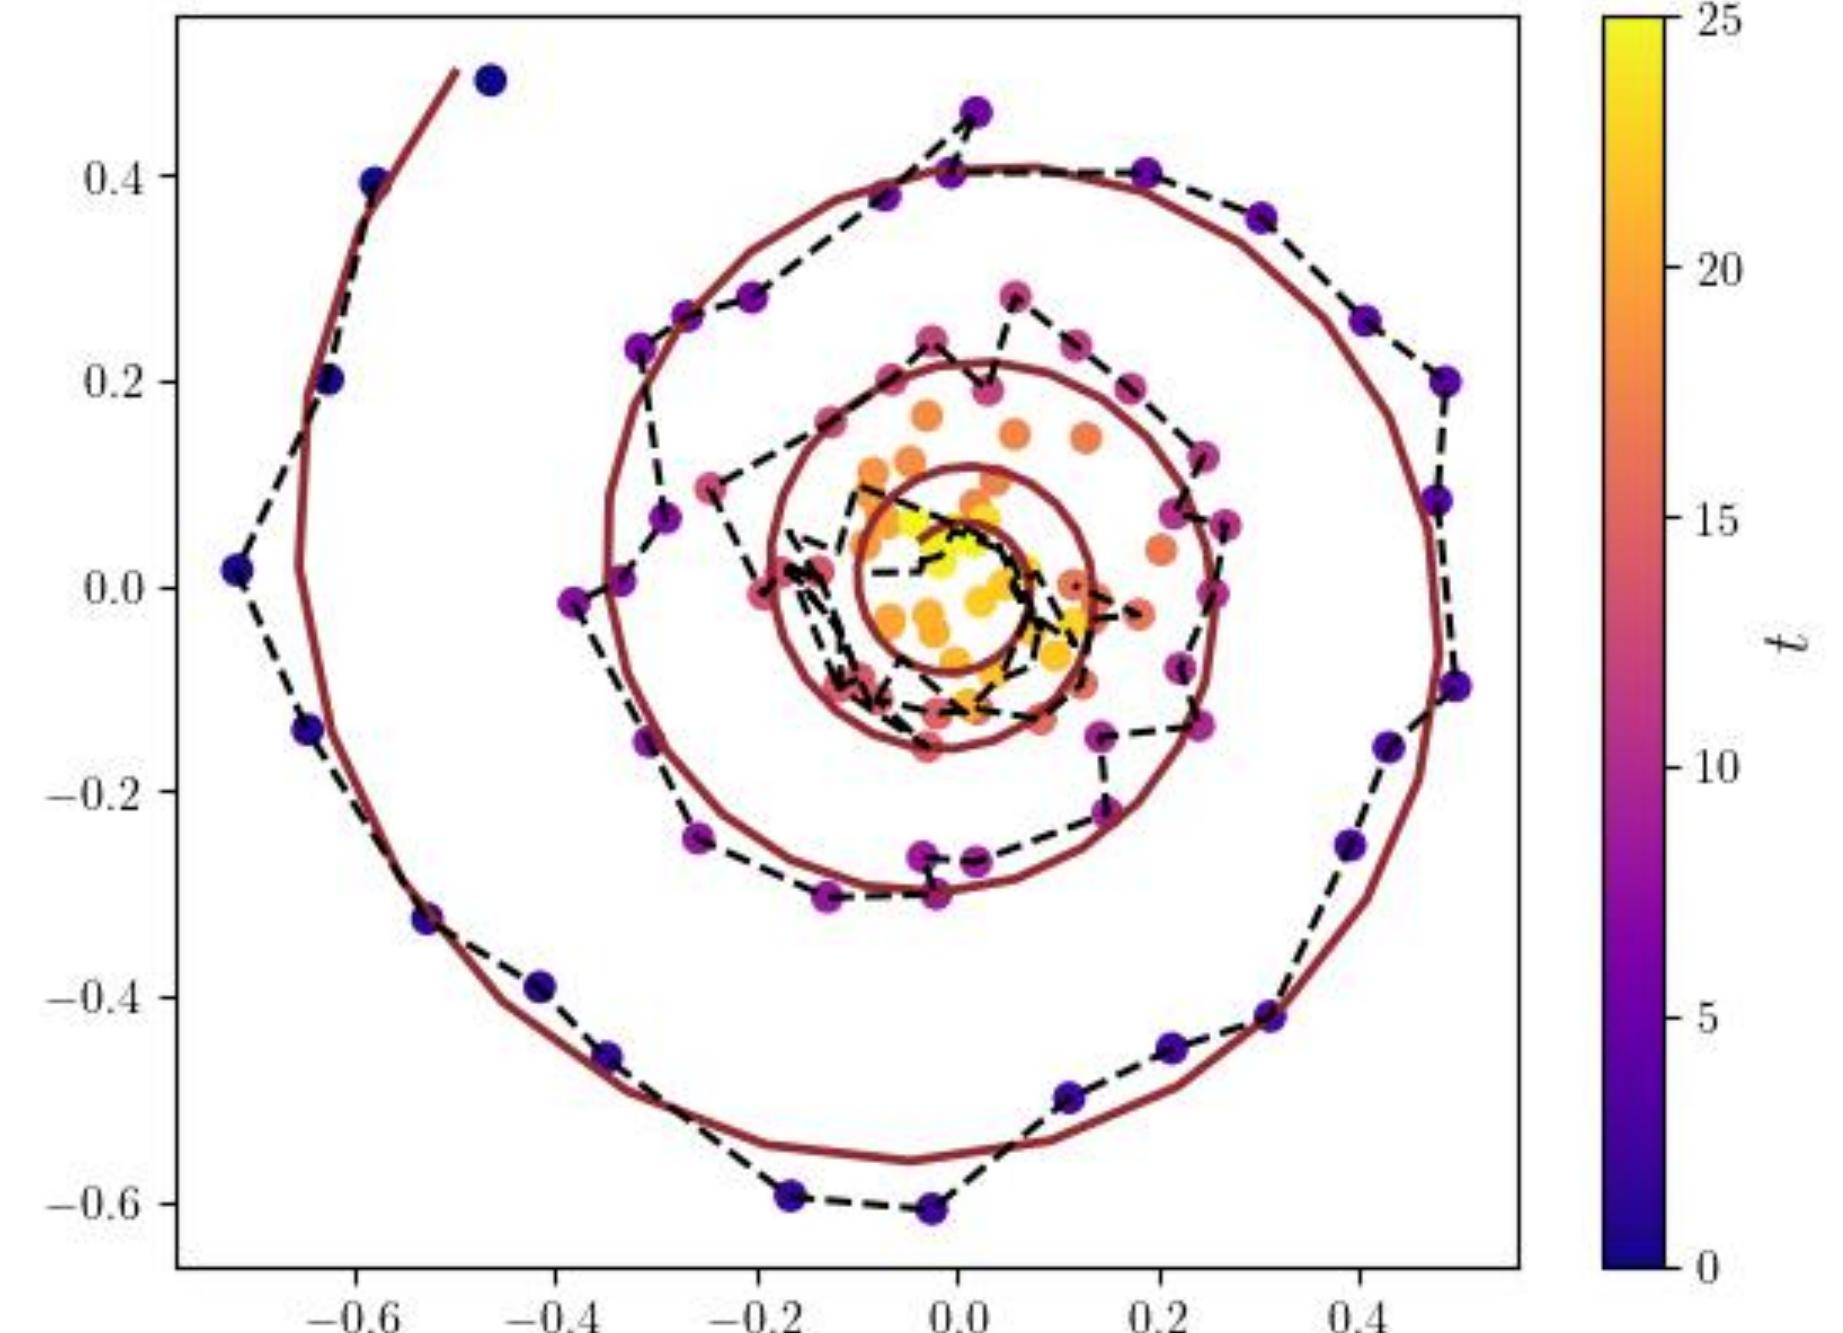
The LSTM tries to predict the next point of a stable spiral.

The trajectory is “jagged” and fails to capture the smooth dynamics.

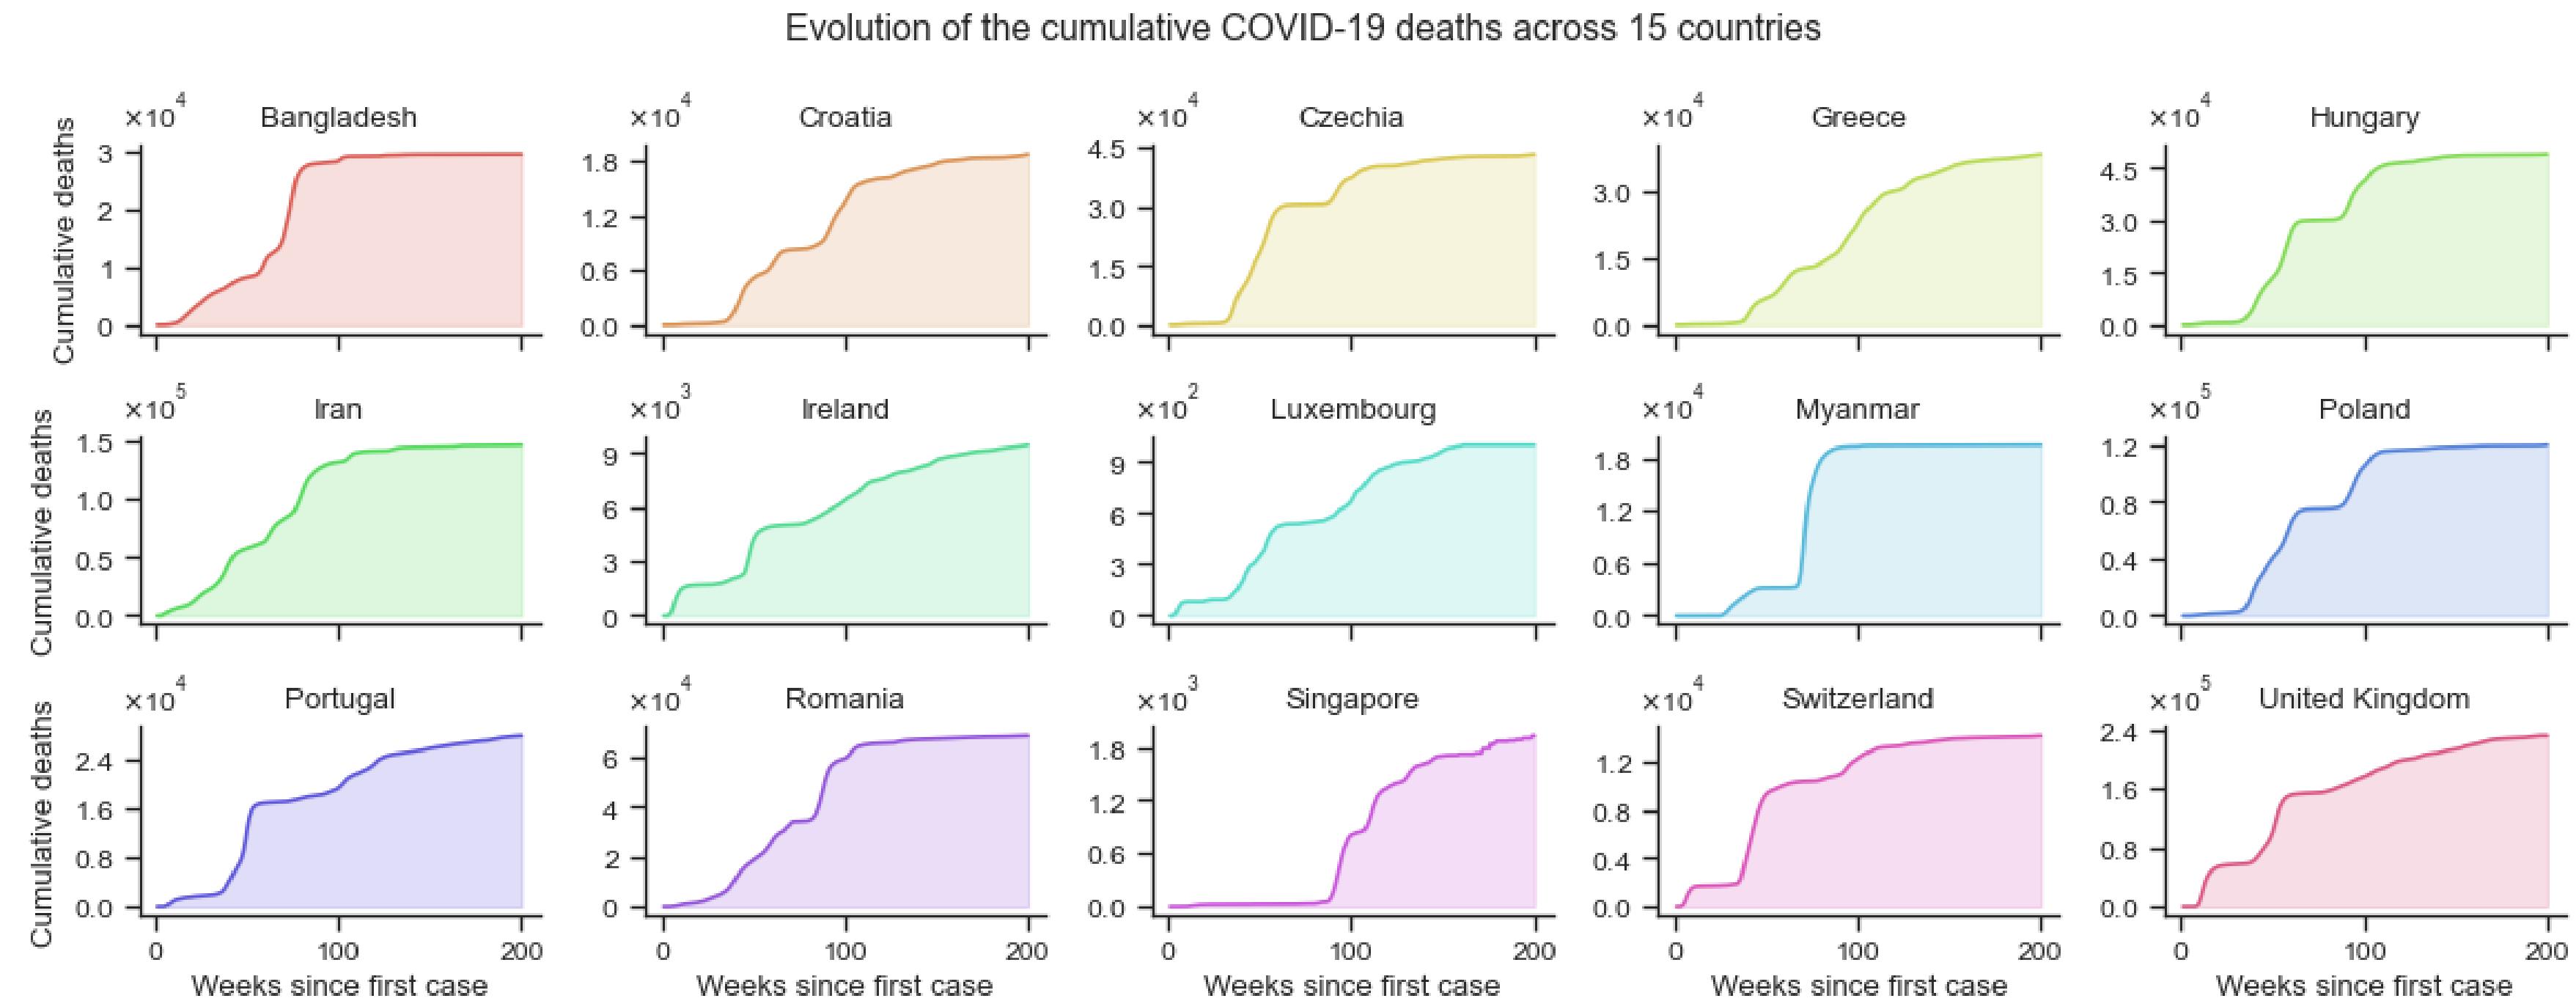
Neural ODE



LSTM



B Epidemic Forecasting



The time series for each country has been divided into four datasets, each containing 25%, 50%, 75%, or 100% of the full time series.

Models

ANODE

$$\frac{d\mathbf{z}}{dt} = f_{\theta}(\mathbf{z}(t))$$

Where:

$$\mathbf{z}(t) = [D(t), \mathbf{h}(t)]$$

Seasonal ANODE

$$\frac{d\mathbf{z}}{dt} = f_{\theta}(\mathbf{z}(t), \mathbf{s}(t))$$

Where:

$$\mathbf{s}(t) = \bigoplus_{k=1}^K \begin{bmatrix} \sin(\omega_k t + \phi_k) \\ \cos(\omega_k t + \phi_k) \end{bmatrix}$$

SIRD

$$\begin{cases} \frac{dS}{dt} &= -\beta \frac{IS}{N} \\ \frac{dI}{dt} &= \beta \frac{IS}{N} - \gamma I - \mu I \\ \frac{dR}{dt} &= \gamma I \\ \frac{dD}{dt} &= \mu I \\ N &= S + I + R + D \end{cases}$$

SIRD w/ reinfection

$$\begin{cases} \frac{dS}{dt} &= -\beta \frac{IS}{N} + \delta R \\ \frac{dI}{dt} &= \beta \frac{IS}{N} - \gamma I - \mu I \\ \frac{dR}{dt} &= \gamma I - \delta R \\ \frac{dD}{dt} &= \mu I \\ N &= S + I + R + D \end{cases}$$

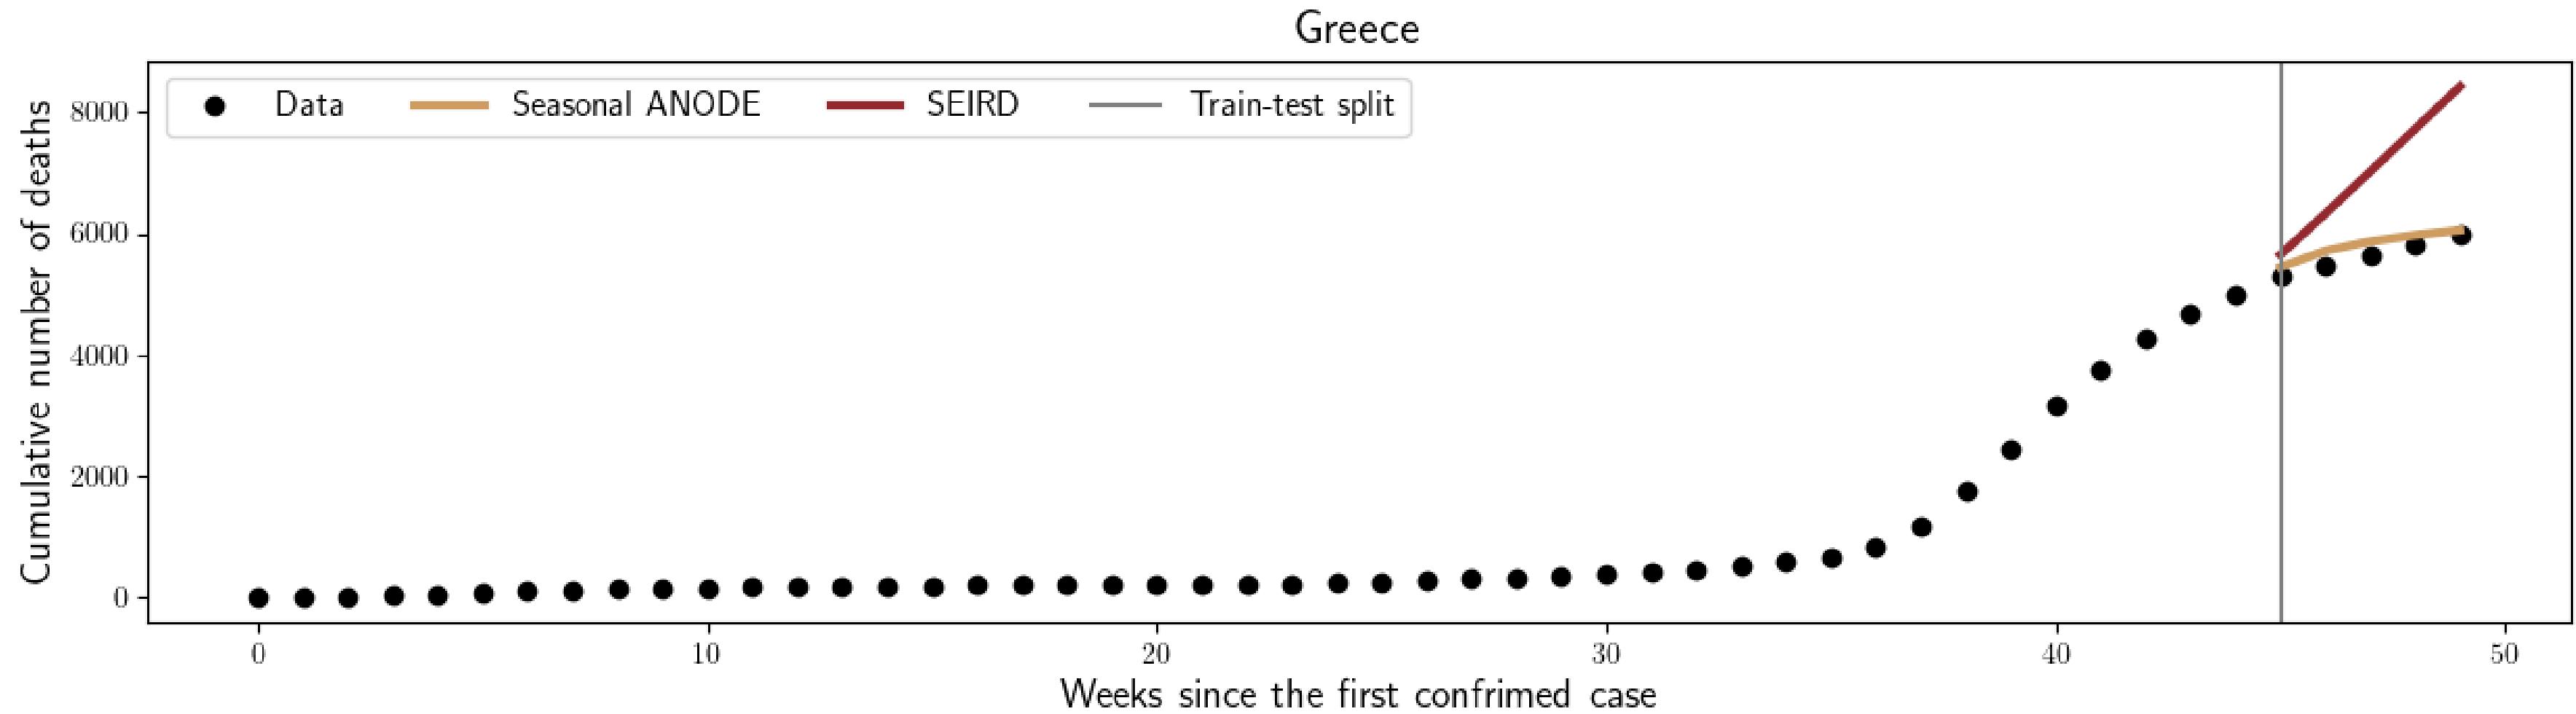
SEIRD

$$\begin{cases} \frac{dS}{dt} &= -\beta \frac{IS}{N} \\ \frac{dE}{dt} &= \beta \frac{IS}{N} - \sigma E \\ \frac{dI}{dt} &= \sigma E - \gamma I - \mu I \\ \frac{dR}{dt} &= \gamma I \\ \frac{dD}{dt} &= \mu I \\ N &= S + E + I + R + D \end{cases}$$

SEIRD w/ vital dynamics

$$\begin{cases} \frac{dS}{dt} &= \Lambda - \frac{\beta SI}{N} - \mu S \\ \frac{dE}{dt} &= \frac{\beta SI}{N} - \sigma E - \mu E \\ \frac{dI}{dt} &= \sigma E - (\gamma + \delta + \mu)I \\ \frac{dR}{dt} &= \gamma I - \mu R \\ \frac{dD}{dt} &= \delta I \end{cases}$$

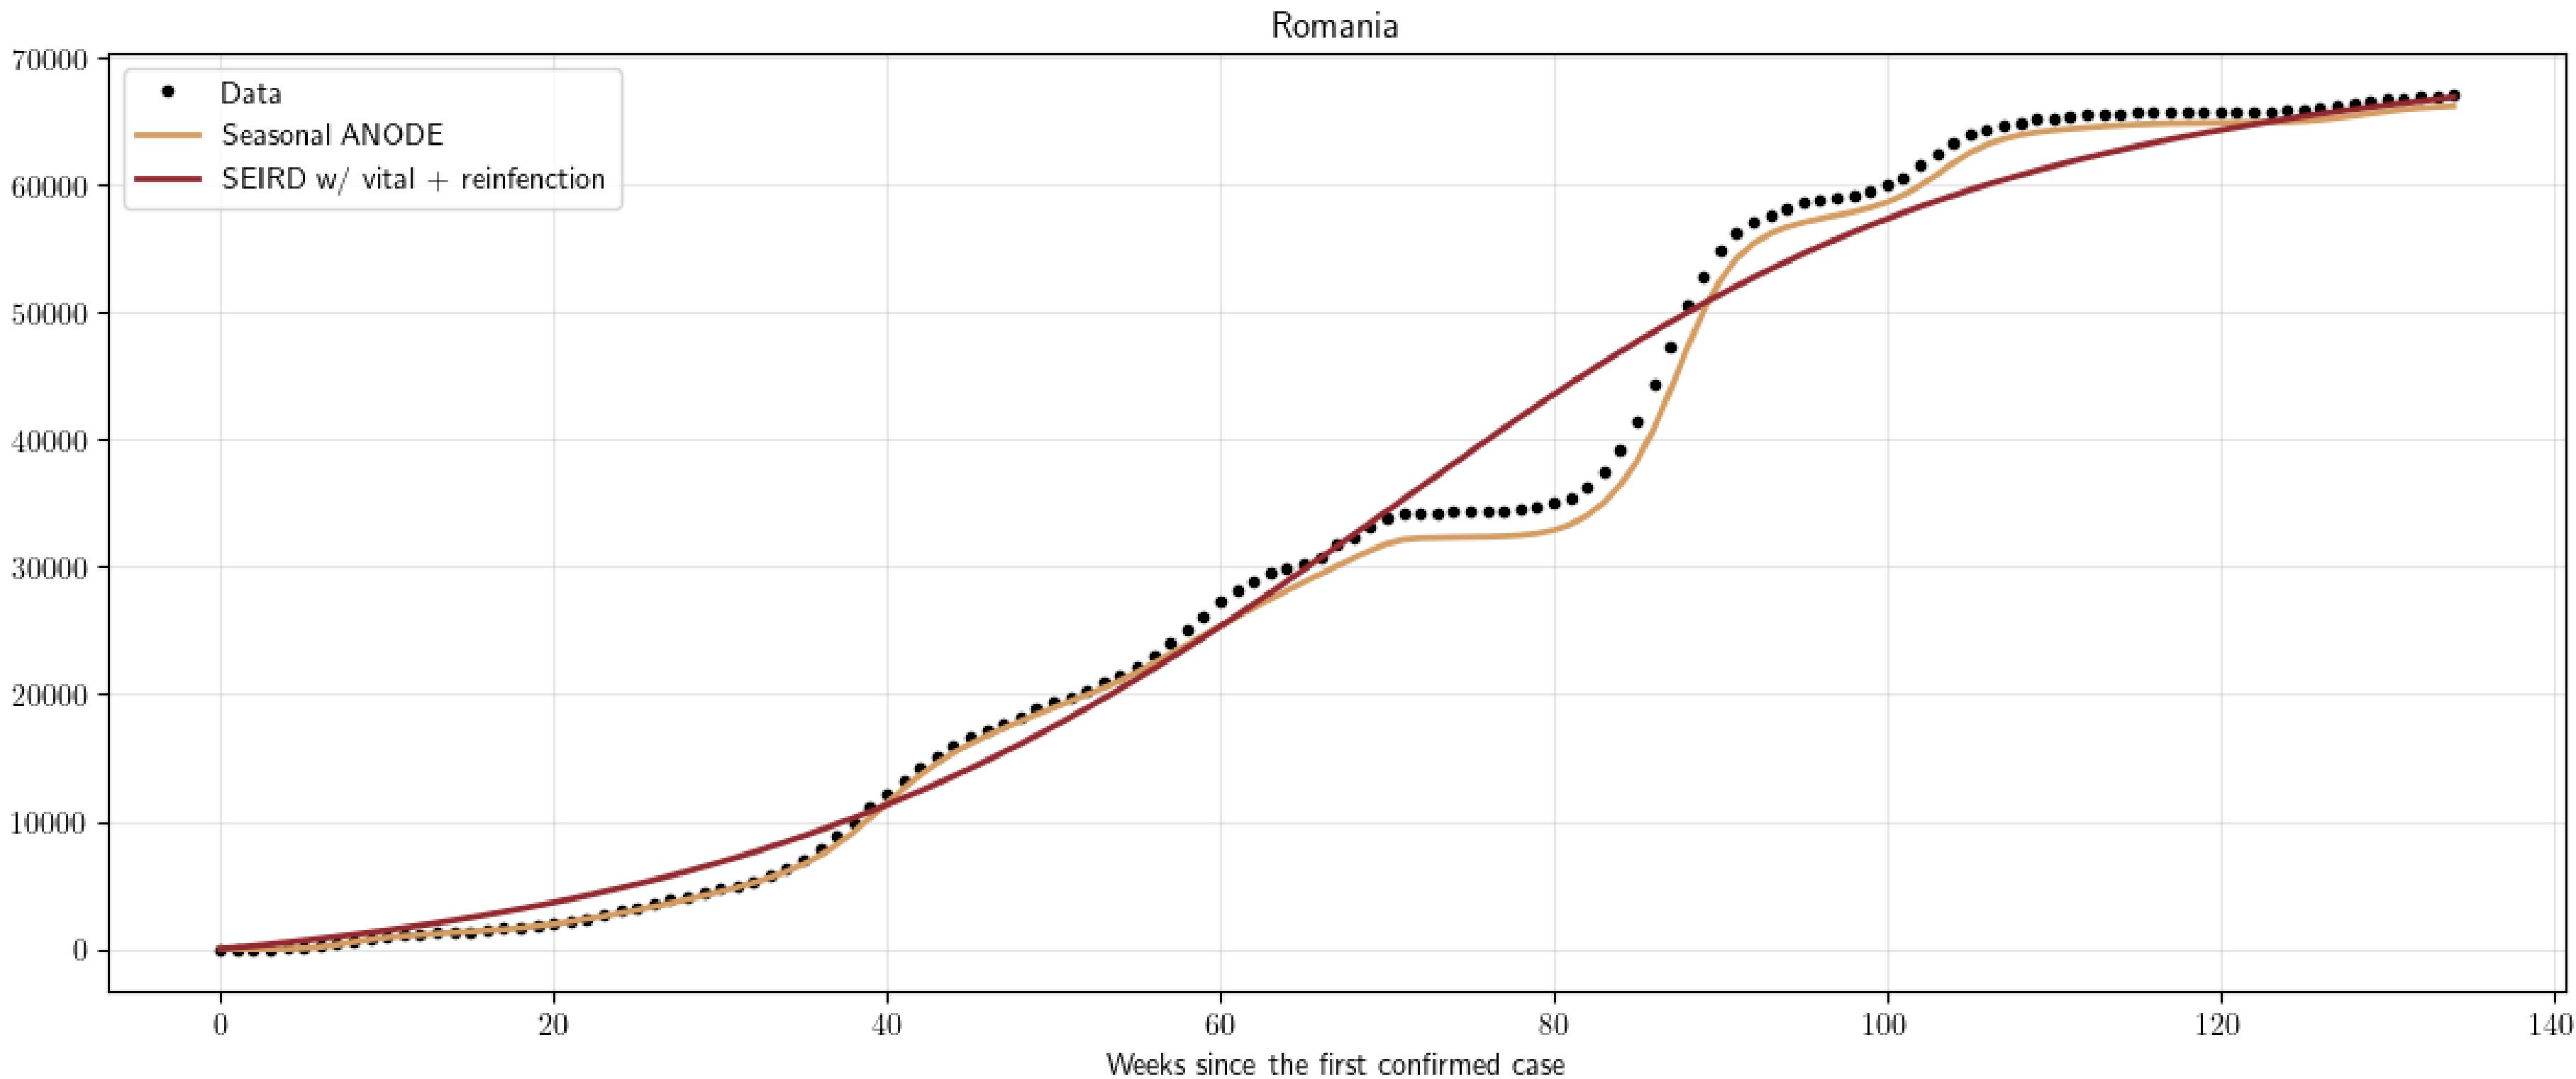
Forecasting



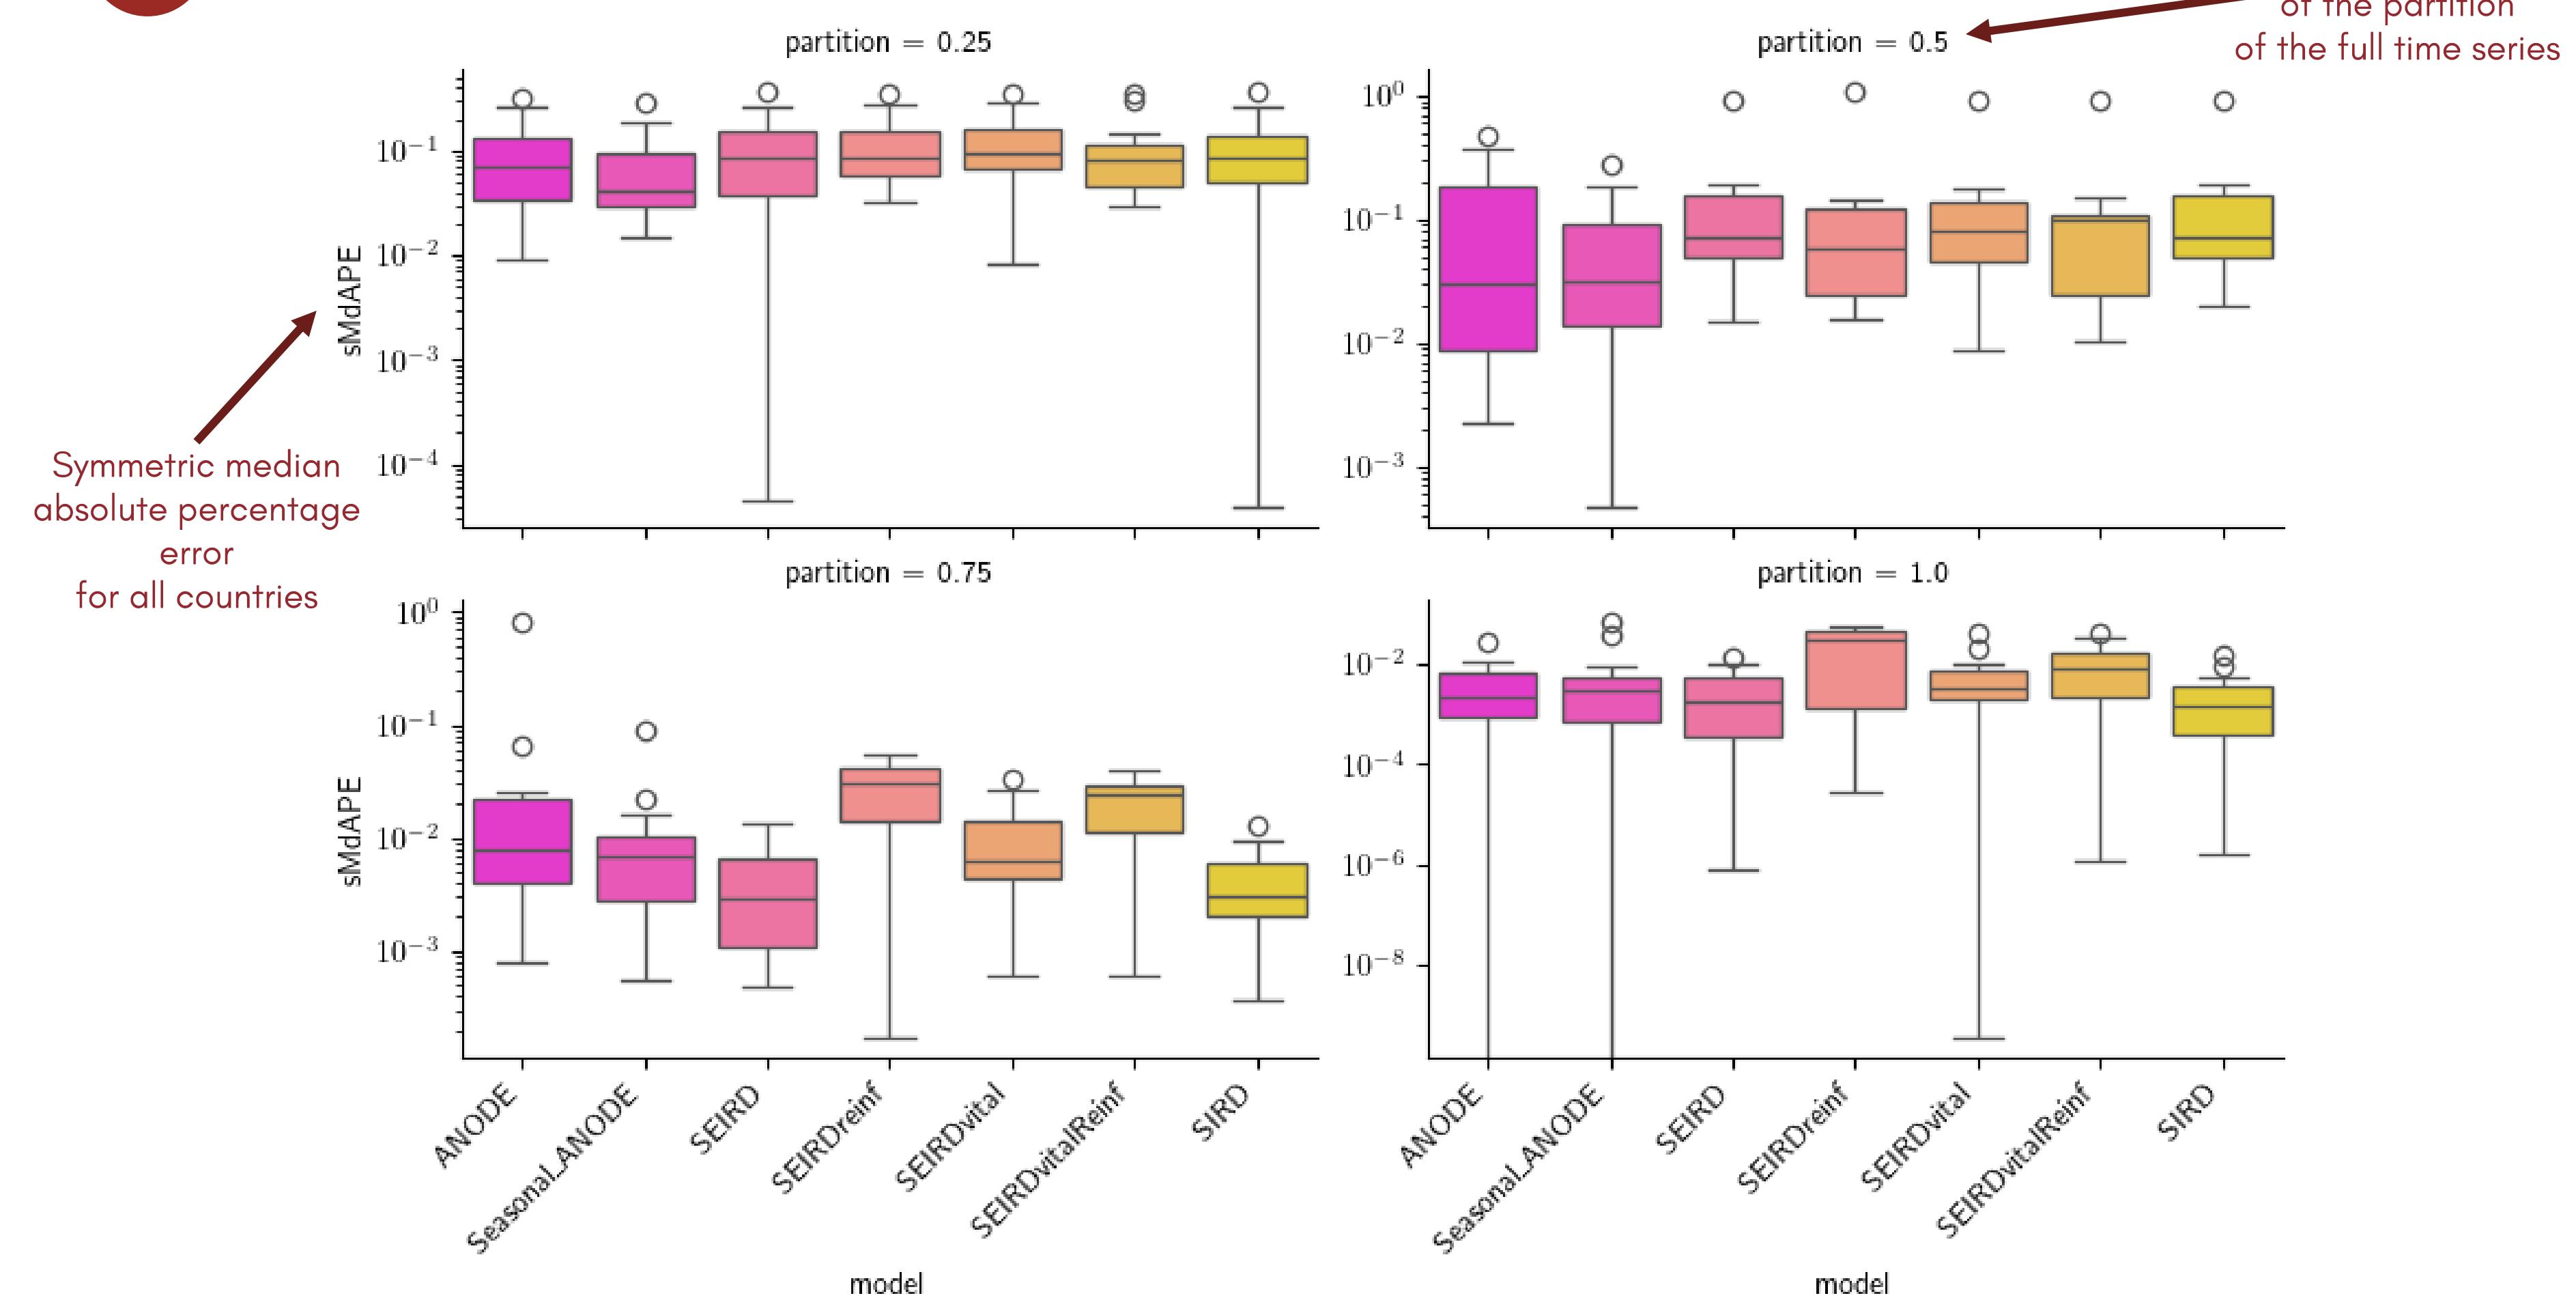
The SEIRD model thinks the number of deaths is still increasing.

The Neural ODEs model successfully captures the slow-down in deaths.

Interpolation



Results



Takeaways

Forecasting

Computational time

Classical ODEs

Fail to capture complex waves or policy shifts without manual tweaking

Simple functional forms means fast integration

Neural ODEs

Can easily learn time-varying dynamics and complex interactions

Requires many function evaluations (NFE) during the solver steps and possibly hyperparameter tuning



UNIVERSITY OF
PATRAS
ΠΑΝΕΠΙΣΤΗΜΙΟ ΠΑΤΡΩΝ

Thank You!