



Mathematical  
Institute

# Neural Controlled Differential Equations

PATRICK KIDGER

*Mathematical Institute  
University of Oxford*

Unlocking data streams, March 2021

Oxford  
Mathematics



Differential equations have arguably been the dominant modelling paradigm for centuries.

Differential equations have arguably been the dominant modelling paradigm for centuries.

Neural networks have more recently joined them as another dominant modelling paradigm.

Differential equations have arguably been the dominant modelling paradigm for centuries.

Neural networks have more recently joined them as another dominant modelling paradigm.

Neural differential equations:

Differential equations have arguably been the dominant modelling paradigm for centuries.

Neural networks have more recently joined them as another dominant modelling paradigm.

Neural differential equations: best of both worlds. Memory efficiency,

Differential equations have arguably been the dominant modelling paradigm for centuries.

Neural networks have more recently joined them as another dominant modelling paradigm.

Neural differential equations: best of both worlds. Memory efficiency, strong priors on model space,

Differential equations have arguably been the dominant modelling paradigm for centuries.

Neural networks have more recently joined them as another dominant modelling paradigm.

Neural differential equations: best of both worlds. Memory efficiency, strong priors on model space, high capacity function approximation,

Differential equations have arguably been the dominant modelling paradigm for centuries.

Neural networks have more recently joined them as another dominant modelling paradigm.

Neural differential equations: best of both worlds. Memory efficiency, strong priors on model space, high capacity function approximation, draw on a deep well of theory on both sides.



Differential equations have arguably been the dominant modelling paradigm for centuries.

Neural networks have more recently joined them as another dominant modelling paradigm.

Neural differential equations: best of both worlds. Memory efficiency, strong priors on model space, high capacity function approximation, draw on a deep well of theory on both sides.

Applications to traditional mathematical modelling (finance, physics, engineering, ...),

Differential equations have arguably been the dominant modelling paradigm for centuries.

Neural networks have more recently joined them as another dominant modelling paradigm.

Neural differential equations: best of both worlds. Memory efficiency, strong priors on model space, high capacity function approximation, draw on a deep well of theory on both sides.

Applications to traditional mathematical modelling (finance, physics, engineering, ...), and to modern machine learning (time series, generative modelling).

## Neural differential equations

## Neural differential equations Intuition

## Neural differential equations

Intuition

Applications

Neural differential equations

Intuition

Applications

Neural controlled differential equations

Neural differential equations

Intuition

Applications

Neural controlled differential equations

Continuous time recurrent neural networks

Neural differential equations

Intuition

Applications

Neural controlled differential equations

Continuous time recurrent neural networks

Irregular time series



Neural differential equations

Intuition

Applications

Neural controlled differential equations

Continuous time recurrent neural networks

Irregular time series

Neural stochastic differential equations

Neural differential equations

Intuition

Applications

Neural controlled differential equations

Continuous time recurrent neural networks

Irregular time series

Neural stochastic differential equations

Software

# What is a neural differential equation anyway?

(And why you might already be using them.)



These are differential equations for which the vector field is parameterised as a neural network.

# What is a neural differential equation anyway?

(And why you might already be using them.)



These are differential equations for which the vector field is parameterised as a neural network.

Canonical example - neural ordinary differential equation (Chen et al. 2018):

$$y(0) = y_0 \quad \frac{dy}{dt}(t) = f(t; y(t))$$

# What is a neural differential equation anyway?

(And why you might already be using them.)



These are differential equations for which the vector field is parameterised as a neural network.

Canonical example - neural ordinary differential equation (Chen et al. 2018):

$$y(0) = y_0 \quad \frac{dy}{dt}(t) = f(t; y(t))$$

$f$  is any standard network - for our purposes today, often a feedforward network:

# What is a neural differential equation anyway?

(And why you might already be using them.)




These are differential equations for which the vector field is parameterised as a neural network.

Canonical example - neural ordinary differential equation (Chen et al. 2018):

$$y(0) = y_0 \quad \frac{dy}{dt}(t) = f(t; y(t))$$

$f$  is any standard network - for our purposes today, often a feedforward network:



# What is a neural differential equation anyway?

(And why you might already be using them.)



Classical example of a 'neural' differential equation: the SIR model.

$$\frac{d}{dt} \begin{pmatrix} s(t) \\ i(t) \\ r(t) \end{pmatrix} = \begin{pmatrix} b s(t) i(t) \\ -b s(t) i(t) - k i(t) \\ k i(t) \end{pmatrix}$$

# What is a neural differential equation anyway?

(And why you might already be using them.)

---

Classical example of a 'neural' differential equation: the SIR model.

$$\frac{d}{dt} \begin{pmatrix} s(t) \\ i(t) \\ r(t) \end{pmatrix} = \begin{pmatrix} b s(t) i(t) \\ -b s(t) i(t) - k i(t) \\ k i(t) \end{pmatrix}$$

$b$  and  $k$  are parameters, controlling infectivity and removal rates respectively.



# What is a neural differential equation anyway?

(And why you might already be using them.)

---

Classical example of a 'neural' differential equation: the SIR model.

$$\frac{d}{dt} \begin{pmatrix} s(t) \\ i(t) \\ r(t) \end{pmatrix} = \begin{pmatrix} b s(t) i(t) \\ -b s(t) i(t) - k i(t) \\ k i(t) \end{pmatrix}$$

$b$  and  $k$  are parameters, controlling infectivity and removal rates respectively.

Crucially for our purposes: they are parameters learnt from data.

# What is a neural differential equation anyway?

(And why you might already be using them.)

---

Classical example of a 'neural' differential equation: the SIR model.

$$\frac{d}{dt} \begin{pmatrix} s(t) \\ i(t) \\ r(t) \end{pmatrix} = \begin{pmatrix} b s(t) i(t) \\ -b s(t) i(t) - k i(t) \\ k i(t) \end{pmatrix}$$

$b$  and  $k$  are parameters, controlling infectivity and removal rates respectively.

Crucially for our purposes: they are parameters learnt from data.

# What is a neural differential equation anyway?

(And why you might already be using them.)

---

Another example: the Heston model.

$$\begin{aligned}dS_t &= S_t \left( \mu - \frac{\rho \sigma \sqrt{v_t}}{2} \right) dt + \sigma \sqrt{v_t} S_t dW_t^1 \\dv_t &= \kappa(\theta - v_t) dt + \sigma \sqrt{v_t} dW_t^2\end{aligned}$$

# What is a neural differential equation anyway?

(And why you might already be using them.)

---

Another example: the Heston model.

$$\begin{aligned}dS_t &= S_t \left( \mu - \frac{\rho \sigma \sqrt{v_t}}{2} \right) dt + \sigma \sqrt{v_t} S_t dW_t^1 \\dv_t &= \kappa(\theta - v_t) dt + \xi \sqrt{v_t} dW_t^2\end{aligned}$$

# What is a neural differential equation anyway?

(And why you might already be using them.)

---

Consider solving

$$\frac{dy}{dt}(t) = f(t; y(t))$$

with the explicit Euler method.

# What is a neural differential equation anyway?

(And why you might already be using them.)

---

Consider solving

$$\frac{dy}{dt}(t) = f(t; y(t))$$

with the explicit Euler method. Approximate

$$\frac{y_{t_{n+1}} - y_{t_n}}{t} \approx \frac{dy}{dt}$$

# What is a neural differential equation anyway?

(And why you might already be using them.)

---

Consider solving

$$\frac{dy}{dt}(t) = f(t; y(t))$$

with the explicit Euler method. Approximate

$$\frac{y_{t_{n+1}} - y_{t_n}}{t} \approx \frac{dy}{dt}$$

which rearranges to

$$y_{t_{n+1}} = y_{t_n} + \Delta t f(t; y_{t_n}):$$

This is now the same as a residual network, one of the main neural network architectures.

# Benefits

---



From traditional mathematical modelling:

From traditional mathematical modelling neural networks are very high capacity models

From traditional mathematical modelling neural networks are very high capacity models easy to train

From traditional mathematical modelling, neural networks are very high capacity models, easy to train, the NDE formalism means they easily augment traditional techniques.

From traditional mathematical modelling neural networks are very high capacity models, easy to train, the NDE formalism means they easily augment traditional techniques.

From machine learning:

From traditional mathematical modelling neural networks are very high capacity models, easy to train, the NDE formalism means they easily augment traditional techniques.

From machine learning differential equations are good priors on model space

From traditional mathematical modelling neural networks are very high capacity models easy to train, the NDE formalism means they easily augment traditional techniques.

From machine learning differential equations are good priors on model space able to handle irregular data

From traditional mathematical modelling neural networks are very high capacity models easy to train, the NDE formalism means they easily augment traditional techniques.

From machine learning differential equations are good priors on model space able to handle irregular data memory efficient



From traditional mathematical modelling neural networks are very high capacity models, easy to train, the NDE formalism means they easily augment traditional techniques.

From machine learning differential equations are good priors on model space, able to handle irregular data, memory efficient, give theoretical grounding to many neural networks.

Neural differential equations have (to my knowledge) three main uses:

Neural differential equations have (to my knowledge) three main uses:

Physical ( financial, ...) modelling, for which a differential equation is explicitly desired.

Neural differential equations have (to my knowledge) three main uses:

Physical (financial, ...) modelling, for which a differential equation is explicitly desired.

Generative modelling: continuous normalising flows; neural SDEs.

Neural differential equations have (to my knowledge) three main uses:

Physical (financial, ...) modelling, for which a differential equation is explicitly desired.

Generative modelling: continuous normalising flows; neural SDEs.

Time series applications: data may arrive irregularly sampled, partially observed and so on.

# Neural controlled differential equations

---

## Neural controlled differential equations

---

Consider the (neural) ODE

$$y(t) = y(0) + \int_0^t f(y(s)) ds \quad \text{for } t \in [0; T]:$$

## Neural controlled differential equations

---

Consider the (neural) ODE

$$y(t) = y(0) + \int_0^t f(y(s)) ds \quad \text{for } t \in [0; T]:$$

The integral over  $[0; T]$  is introduced and then integrated over, and is just an internal detail of the model.



## Neural controlled differential equations

---

Consider the (neural) ODE

$$y(t) = y(0) + \int_0^t f(y(s)) ds \quad \text{for } t \in [0; T]:$$

The integral over  $t \in [0; T]$  is introduced and then integrated over, and is just an internal detail of the model.

But given some ordered data  $(x_0; \dots; x_n)$ , we would like to align  $s \in [0; T]$  with the ordering of the data.

## Neural controlled differential equations

---

Consider the (neural) ODE

$$y(t) = y(0) + \int_0^t f(y(s)) ds \quad \text{for } t \in [0; T]:$$

The integral over  $t \in [0; T]$  is introduced and then integrated over, and is just an internal detail of the model.

But given some ordered data  $(x_0; \dots; x_n)$ , we would like to align  $s \in [0; T]$  with the ordering of the data.

Problem: it's an ODE. The solution is determined by its initial condition.

## Neural controlled differential equations

---

Consider the (neural) ODE

$$y(t) = y(0) + \int_0^t f(y(s)) ds \quad \text{for } t \in [0; T]:$$

The integral over  $[0; T]$  is introduced and then integrated over, and is just an internal detail of the model.

But given some ordered data  $(x_0; \dots; x_n)$ , we would like to align  $s \in [0; T]$  with the ordering of the data.

Problem: it's an ODE. The solution is determined by its initial condition.

Want a way to incorporate incoming information.

# Neural controlled differential equations

---

# Neural controlled differential equations

(Lyons 1998)

---

# Neural controlled differential equations

(Lyons 1998)

---

Have the local dynamics of the system depend upon some time-varying  $X : [0; T] \rightarrow \mathbb{R}^V$ !

# Neural controlled differential equations

(Lyons 1998)

---

Have the local dynamics of the system depend upon some time-varying  $X : [0; T] \rightarrow \mathbb{R}^v$ :

$$y(t) = y(0) + \int_0^t f(y(s)) dX(s) \quad \text{for } t \in [0; T]:$$

# Neural controlled differential equations

(Lyons 1998)

---

Have the local dynamics of the system depend upon some time-varying  $X : [0; T] \rightarrow \mathbb{R}^v$ :

$$y(t) = y(0) + \int_0^t f(y(s)) dX(s) \quad \text{for } t \in [0; T]:$$

which is a Riemann-Stieltjes integral.



# Neural controlled differential equations

(Lyons 1998)

---

Have the local dynamics of the system depend upon some time-varying  $X : [0; T] \rightarrow \mathbb{R}^V$ :

$$y(t) = y(0) + \int_0^t f(y(s)) dX(s) \quad \text{for } t \in [0; T]:$$

which is a Riemann-Stieltjes integral. Intuitively similar to

$$y(t) = y(0) + \int_0^t f(y(s); X(s)) ds \quad \text{for } t \in [0; T]:$$

# Neural controlled differential equations

(Lyons 1998)

---

Have the local dynamics of the system depend upon some time-varying  $X : [0; T] \rightarrow \mathbb{R}^v$ :

$$y(t) = y(0) + \int_0^t f(y(s)) dX(s) \quad \text{for } t \in [0; T]:$$

which is a Riemann-Stieltjes integral. Intuitively similar to

$$y(t) = y(0) + \int_0^t f(y(s); X(s)) ds \quad \text{for } t \in [0; T]:$$

(The first equation ends up being slightly more flexible, and a bit mathematically neater.)

## Neural controlled differential equations

(Lyons 1998)

---

Have the local dynamics of the system depend upon some time-varying  $X : [0; T] \rightarrow \mathbb{R}^V$ :

$$y(t) = y(0) + \int_0^t f(y(s)) dX(s) \quad \text{for } t \in [0; T]:$$

which is a Riemann-Stieltjes integral. Intuitively similar to

$$y(t) = y(0) + \int_0^t f(y(s); X(s)) ds \quad \text{for } t \in [0; T]:$$

(The first equation ends up being slightly more flexible, and a bit mathematically neater.)

Changes in  $X$  provoke changes in  $y$ .

---

# Neural controlled differential equations

Irregular time series

---

# Neural controlled differential equations

Irregular time series

---

Observe  $x = ((t_0; x_0); \dots; (t_n; x_n))$  with  $t_i \in \mathbb{R}$  and  $x_i \in \mathbb{R}^y$ .

# Neural controlled differential equations

Irregular time series

---

Observe  $x = ((t_0; x_0); \dots; (t_n; x_n))$  with  $t_i \in \mathbb{R}$  and  $x_i \in \mathbb{R}^v$ .  
Let  $X : [t_0; t_n] \rightarrow \mathbb{R}^{v+1}$  interpolate this data, so  $X(t_i) = (t_i; x_i)$ .

# Neural controlled differential equations

Irregular time series

---

Observe  $x = ((t_0; x_0); \dots; (t_n; x_n))$  with  $t_i \in \mathbb{R}$  and  $x_i \in \mathbb{R}^v$ .

Let  $X : [t_0; t_n] \rightarrow \mathbb{R}^{v+1}$  interpolate this data, so  $X(t_i) = (t_i; x_i)$ .

Learn functions  $\sigma, f$  and a linear map  $\phi$  such that

# Neural controlled differential equations

Irregular time series

---

Observe  $x = ((t_0; x_0); \dots; (t_n; x_n))$  with  $t_i \in \mathbb{R}$  and  $x_i \in \mathbb{R}^v$ .

Let  $X : [t_0; t_n] \times \mathbb{R}^{v+1}$  interpolate this data, so  $X(t_i) = (t_i; x_i)$ .

Learn functions  $\sigma, f$  and a linear map  $\phi$  such that

$$y(t_0) = \phi(t_0; x_0);$$



# Neural controlled differential equations

Irregular time series

---

Observe  $x = ((t_0; x_0); \dots; (t_n; x_n))$  with  $t_i \in \mathbb{R}$  and  $x_i \in \mathbb{R}^v$ .

Let  $X : [t_0; t_n] \rightarrow \mathbb{R}^{v+1}$  interpolate this data, so  $X(t_i) = (t_i; x_i)$ .

Learn functions  $\sigma, f$  and a linear map  $\delta$  such that

$$y(t_0) = (t_0; x_0); \quad y(t) = y(t_0) + \int_{t_0}^t \delta(y(s)) f(y(s)) dX(s);$$

# Neural controlled differential equations

Irregular time series

---

Observe  $x = ((t_0; x_0); \dots; (t_n; x_n))$  with  $t_i \in \mathbb{R}$  and  $x_i \in \mathbb{R}^v$ .

Let  $X : [t_0; t_n] \rightarrow \mathbb{R}^{v+1}$  interpolate this data, so  $X(t_i) = (t_i; x_i)$ .

Learn functions  $\sigma, f$  and a linear map  $\delta$  such that

$$y(t_0) = \sigma(t_0; x_0); \quad y(t) = y(t_0) + \int_{t_0}^t \delta(y(s)) f(y(s)) dX(s);$$

and the output is either  $\sigma(y(t_n))$  or  $\delta(y(t))$ .

# Neural controlled differential equations

Irregular time series

---

Observe  $x = ((t_0; x_0); \dots; (t_n; x_n))$  with  $t_i \in \mathbb{R}$  and  $x_i \in \mathbb{R}^v$ .  
Let  $X : [t_0; t_n] \rightarrow \mathbb{R}^{v+1}$  interpolate this data, so  $X(t_i) = (t_i; x_i)$ .

Learn functions  $\sigma, f$  and a linear map  $\delta$  such that

$$y(t_0) = \sigma(t_0; x_0); \quad y(t) = y(t_0) + \int_{t_0}^t f(y(s)) dX(s);$$

and the output is either  $\sigma(y(t_n))$  or  $\sigma(y(t))$ .

$\sigma$  and  $f$  are arbitrary neural networks (feed-forward, ...), and  $y$  is hidden state.

# Neural controlled differential equations

Irregular time series

---

Observe  $x = ((t_0; x_0); \dots; (t_n; x_n))$  with  $t_i \in \mathbb{R}$  and  $x_i \in \mathbb{R}^v$ .


Let  $X : [t_0; t_n] \rightarrow \mathbb{R}^{v+1}$  interpolate this data, so  $X(t_i) = (t_i; x_i)$ .

Learn functions  $\sigma$ ,  $f$  and a linear map  $A$  such that

$$y(t_0) = \sigma(t_0; x_0); \quad y(t) = y(t_0) + \int_{t_0}^t A(y(s)) f(y(s)) dX(s);$$

and the output is either  $\sigma(y(t_n))$  or  $\sigma(y(t))$ .

and  $f$  are arbitrary neural networks (feed-forward, ...), and  $y$  is hidden state.



# Neural controlled differential equations

Irregular time series


Observe  $x = ((t_0; x_0); \dots; (t_n; x_n))$  with  $t_i \in \mathbb{R}$  and  $x_i \in \mathbb{R}^v$ .  
Let  $X : [t_0; t_n] \rightarrow \mathbb{R}^{v+1}$  interpolate this data, so  $X(t_i) = (t_i; x_i)$ .

Learn functions  $\sigma$ ,  $f$  and a linear map  $\mathcal{L}$  such that

$$y(t_0) = \sigma(t_0; x_0); \quad y(t) = y(t_0) + \int_{t_0}^t f(y(s)) dX(s);$$

and the output is either  $\sigma(y(t_n))$  or  $\sigma(y(t))$ .

$\sigma$  and  $f$  are arbitrary neural networks (feed-forward, ...), and  $y$  is hidden state.



# Neural controlled differential equations

Irregular time series


Observe  $x = ((t_0; x_0); \dots; (t_n; x_n))$  with  $t_i \in \mathbb{R}$  and  $x_i \in \mathbb{R}^V$ .  
Let  $X : [t_0; t_n] \rightarrow \mathbb{R}^{V+1}$  interpolate this data, so  $X(t_i) = (t_i; x_i)$ .

Learn functions  $\sigma$ ,  $f$  and a linear map  $Z_t$  such that

$$y(t_0) = \sigma(t_0; x_0); \quad y(t) = y(t_0) + \int_{t_0}^t f(y(s)) dX(s);$$

and the output is either  $\sigma(y(t_n))$  or  $\sigma(y(t))$ .

and  $f$  are arbitrary neural networks (feed forward, ...), and  $y$  is hidden state.



# Neural controlled differential equations

Irregular time series


Observe  $x = ((t_0; x_0); \dots; (t_n; x_n))$  with  $t_i \in \mathbb{R}$  and  $x_i \in \mathbb{R}^V$ .  
Let  $X : [t_0; t_n] \rightarrow \mathbb{R}^{V+1}$  interpolate this data, so  $X(t_i) = (t_i; x_i)$ .

Learn functions  $\sigma, f$  and a linear map  $Z_t$  such that

$$y(t_0) = \sigma(t_0; x_0); \quad y(t) = y(t_0) + \int_{t_0}^t f(y(s)) dX(s);$$

and the output is either  $\sigma(y(t_n))$  or  $\sigma(y(t))$ .

and  $f$  are arbitrary neural networks (feed-forward, ...), and  $y$  is hidden state.



# Neural controlled differential equations

Irregular time series


Observe  $x = ((t_0; x_0); \dots; (t_n; x_n))$  with  $t_i \in \mathbb{R}$  and  $x_i \in \mathbb{R}^v$ .  
Let  $X : [t_0; t_n] \rightarrow \mathbb{R}^{v+1}$  interpolate this data, so  $X(t_i) = (t_i; x_i)$ .

Learn functions  $\sigma$ ,  $f$  and a linear map  $Z_t$  such that

$$y(t_0) = \sigma(t_0; x_0); \quad y(t) = y(t_0) + \int_{t_0}^t f(y(s)) dX(s);$$

and the output is either  $\sigma(y(t_n))$  or  $\sigma(y(t))$ .

and  $f$  are arbitrary neural networks (feed-forward, ...), and  $y$  is hidden state.





# Neural controlled differential equations

Irregular time series


Observe  $x = ((t_0; x_0); \dots; (t_n; x_n))$  with  $t_i \in \mathbb{R}$  and  $x_i \in \mathbb{R}^V$ .  
 Let  $X : [t_0; t_n] \rightarrow \mathbb{R}^{V+1}$  interpolate this data, so  $X(t_i) = (t_i; x_i)$ .

Learn functions  $\sigma, f$  and a linear map  $\mathcal{Z}_t$  such that

$$y(t_0) = \sigma(t_0; x_0); \quad y(t) = y(t_0) + \int_{t_0}^t f(y(s)) dX(s);$$

and the output is either  $\sigma(y(t_n))$  or  $\sigma(y(t))$ .

and  $f$  are arbitrary neural networks (feed-forward, ...), and  $y$  is hidden state.



# Neural controlled differential equations

Irregular time series


Observe  $x = ((t_0; x_0); \dots; (t_n; x_n))$  with  $t_i \in \mathbb{R}$  and  $x_i \in \mathbb{R}^V$ .  
Let  $X : [t_0; t_n] \rightarrow \mathbb{R}^{V+1}$  interpolate this data, so  $X(t_i) = (t_i; x_i)$ .

Learn functions  $\sigma, f$  and a linear map  $Z_t$  such that

$$y(t_0) = \sigma(t_0; x_0); \quad y(t) = y(t_0) + \int_{t_0}^t f(y(s)) dX(s);$$

and the output is either  $\sigma(y(t_n))$  or  $\sigma(y(t))$ .

and  $f$  are arbitrary neural networks (feed-forward, ...), and  $y$  is hidden state.



# Neural controlled differential equations

Irregular time series


Observe  $x = ((t_0; x_0); \dots; (t_n; x_n))$  with  $t_i \in \mathbb{R}$  and  $x_i \in \mathbb{R}^v$ .  
 Let  $X : [t_0; t_n] \rightarrow \mathbb{R}^{v+1}$  interpolate this data, so  $X(t_i) = (t_i; x_i)$ .

Learn functions  $\sigma, f$  and a linear map  $Z_t$  such that

$$y(t) = \sigma(t_0; x_0); \quad y(t) = y(t_0) + \int_{t_0}^t f(y(s)) dX(s);$$

and the output is either  $\sigma(y(t_n))$  or  $\sigma(y(t))$ .

and  $f$  are arbitrary neural networks (feed-forward, ...), and  $y$  is hidden state. Directly analogous to an RNN  $y_{n+1} = f(y_n; x_n)$ .



# Neural controlled differential equations

Irregular time series

---

We call this a neural controlled differential equation

# Neural controlled differential equations

Irregular time series

---

We call this a neural controlled differential equation

Examples: the time series  $x = ((t_0; x_0); \dots; (t_n; x_n))$  can correspond to:

# Neural controlled differential equations

Irregular time series

---

We call this a neural controlled differential equation

Examples: the time series  $x = ((t_0; x_0); \dots; (t_n; x_n))$  can correspond to:

The x; y-position as a pen is moved over paper, to draw a character.

# Neural controlled differential equations

Irregular time series

---

We call this a neural controlled differential equation

Examples: the time series  $x = ((t_0; x_0); \dots; (t_n; x_n))$  can correspond to:

The x;y-position as a pen is moved over paper, to draw a character.

Or patient hospital records: heart rate, lab results, and so on.  
A classic example of messy irregular data.

# Neural controlled differential equations

Irregular time series

---

We call this a neural controlled differential equation

Examples: the time series  $x = ((t_0; x_0); \dots; (t_n; x_n))$  can correspond to:

The x;y-position as a pen is moved over paper, to draw a character.

Or patient hospital records: heart rate, lab results, and so on.  
A classic example of messy irregular data.

Spoken audio.



# Neural controlled differential equations

Irregular time series

---

We call this a neural controlled differential equation

Examples: the time series  $x = ((t_0; x_0); \dots; (t_n; x_n))$  can correspond to:

The x; y-position as a pen is moved over paper, to draw a character.

Or patient hospital records: heart rate, lab results, and so on.  
A classic example of messy irregular data.

Spoken audio.

Weather data, e.g. temperature and pressure as they change over time.

# Neural controlled differential equations

Irregular time series

---

We call this a neural controlled differential equation

Examples: the time series  $x = ((t_0; x_0); \dots; (t_n; x_n))$  can correspond to:

The x;y-position as a pen is moved over paper, to draw a character.

Or patient hospital records: heart rate, lab results, and so on.  
A classic example of messy irregular data.

Spoken audio.

Weather data, e.g. temperature and pressure as they change over time.

Physics: the movement of a double pendulum.

# Neural controlled differential equations

Irregular time series

---

## Advantages:

Using a continuous-time theory pushes the problem of messy data into the interpolation, which handles it easily.

# Neural controlled differential equations

Irregular time series

---

## Advantages:

Using a continuous-time theory pushes the problem of messy data into the interpolation, which handles it easily.

The equation  $y(t) = y(0) + \int_0^t f(y(s)) dX(s)$  may (with a bit of mathematical rearranging) be solved as a (neural) ODE:

# Neural controlled differential equations

Irregular time series

---

## Advantages:

Using a continuous-time theory pushes the problem of messy data into the interpolation, which handles it easily.

The equation  $y(t) = y(0) + \int_0^t f(y(s)) dX(s)$  may (with a bit of mathematical rearranging) be solved as a (neural) ODE:

So we can solve it using existing software.

# Neural controlled differential equations

Irregular time series

---

## Advantages:

Using a continuous-time theory pushes the problem of messy data into the interpolation, which handles it easily.

The equation  $y(t) = y(0) + \int_0^t f(y(s)) dX(s)$  may (with a bit of mathematical rearranging) be solved as a (neural) ODE:

So we can solve it using existing software.

So we can utilise the memory-efficient adjoint method.

## Advantages:

Using a continuous-time theory pushes the problem of messy data into the interpolation, which handles it easily.

The equation  $y(t) = y(0) + \int_0^t f(y(s)) dX(s)$  may (with a bit of mathematical rearranging) be solved as a (neural) ODE:

So we can solve it using existing software.

So we can utilise the memory-efficient adjoint method.

Drawing on the existing theory of CDEs gives strong theoretical guarantees. For example, neural CDEs are universal approximators.

## Advantages:

Using a continuous-time theory pushes the problem of messy data into the interpolation, which handles it easily.

The equation  $y(t) = y(0) + \int_0^t f(y(s)) dX(s)$  may (with a bit of mathematical rearranging) be solved as a (neural) ODE:

So we can solve it using existing software.

So we can utilise the memory-efficient adjoint method.

Drawing on the existing theory of CDEs gives strong theoretical guarantees. For example, neural CDEs are universal approximators.

Neural CDEs demonstrate state-of-the-art performance.



# Neural controlled differential equations

Irregular time series

---

Model	Test Accuracy						Memory usage
	(mean		std, computed across			ve runs)	
	30% dropped	50% dropped	70% dropped				
GRU-ODE	92.6%	1.6%	86.7%	3.9%	89.9%	3.7%	1.5
GRU- t	93.6%	2.0%	91.3%	2.1%	90.4%	0.8%	15.8
GRU-D	94.2%	2.1%	90.2%	4.8%	91.9%	1.7%	17.0
ODE-RNN	95.4%	0.6%	96.0%	0.3%	95.3%	0.6%	14.8
Neural CDE (ours)	98.7%	0.8%	98.8%	0.2%	98.6%	0.4%	1.3

---

# Neural controlled differential equations

Irregular time series

---

Comparison to control theory problems (speaking broadly):

# Neural controlled differential equations

Irregular time series

---

Comparison to control theory problems (speaking very broadly):

Control theory: System is fixed; try to find optimal  $X$  producing a desired response

# Neural controlled differential equations

Irregular time series

---

Comparison to control theory problems (speaking very broadly):


Control theory: System  $f$  is fixed; try to find optimal  $X$  producing a desired response

(Neural) CDEs: Input  $X$  is fixed; try to find optimal  $f$  producing a desired response


# Neural stochastic differential equations

---

# Neural stochastic differential equations




# Neural stochastic differential equations



Neural SDE / CDE form a generator/discriminator pair.

# Neural stochastic differential equations




Neural SDE / CDE form a generator/discriminator pair.

Arbitrary drift and diffusions are admissible.



# Neural stochastic differential equations




Neural SDE / CDE form a generator/discriminator pair.

Arbitrary drift and diffusions are admissible.

In the infinite data limit any SDE may be learnt.

# Neural stochastic differential equations



Neural SDE / CDE form a generator/discriminator pair.

Arbitrary drift and diffusions are admissible.

In the infinite data limit any SDE may be learnt.

We've come a long way from models like the Heston model...

$$dS_t = \mu(t; S_t) dt + \sigma(t; S_t) dW_t^1$$

$$d_t = \mu(t; S_t) dt + \sigma(t; S_t) dW_t^2$$

# Neural stochastic differential equations

---

Google/Alphabet Stocks:

Metric	Neural SDE		CTFP		Latent ODE	
Classification	0.357	0.045	0.165	0.087	0.000239	0.000086
Prediction	0.144	0.045	0.725	0.233	46.2	12.3
MMD	1.92	0.09	2.70	0.47	60.4	35.8

Beijing Air Quality:

Metric	Neural SDE		CTFP		Latent ODE	
Classification	0.589	0.051	0.764	0.064	0.392	0.011
Prediction	0.395	0.056	0.810	0.083	0.456	0.095
MMD	0.000160	0.000029	0.00198	0.00001	0.000242	0.000002

SGD dynamics:

Metric	Neural SDE		CTFP		Latent ODE	
Classification	0.507	0.019	0.676	0.014	0.0112	0.0025
Prediction	0.00843	0.00759	0.0808	0.0514	0.127	0.152
MMD	5.28	1.27	12.0	0.5	23.2	11.8

## Software

---

We've got standardised software for solving neural differential equations.

## Software

---

We've got standardised software for solving neural differential equations.

PyTorch:

We've got standardised software for solving neural differential equations.

PyTorch:

`torchdiffeq` (ODEs)

We've got standardised software for solving neural differential equations.

PyTorch:

`torchdiffeq` (ODEs)

`torchcde` (CDEs)

We've got standardised software for solving neural differential equations.

PyTorch:

`torchdiffeq` (ODEs)

`torchcde` (CDEs)

`torchsde` (SDEs)



We've got standardised software for solving neural differential equations.

PyTorch:

`torchdiffeq` (ODEs)

`torchcde` (CDEs)

`torchsde` (SDEs)

Julia:

We've got standardised software for solving neural differential equations.

PyTorch:

- `torchdiffeq` (ODEs)
- `torchcde` (CDEs)
- `torchsde` (SDEs)

Julia:

- `DifferentialEquations.jl`

We've got standardised software for solving neural differential equations.

PyTorch:

`torchdiffeq` (ODEs)

`torchcde` (CDEs)

`torchsde` (SDEs)

Julia:

`DifferentialEquations.jl`

We've got standardised software for solving neural differential equations.

PyTorch:

- torchdiffeq (ODEs)
- torchcde (CDEs)
- torchsde (SDEs)

Julia:

- DifferentialEquations.jl

Autodifferentiable, GPU-capable, composable.

# Summary

---

## Summary

---

Parameterised differential equations are already ubiquitous throughout mathematical modelling.

## Summary

---

Parameterised differential equations are already ubiquitous throughout mathematical modelling.

Neural differential equations are a powerful extension to this paradigm.

## Summary

---

Parameterised differential equations are already ubiquitous throughout mathematical modelling.

Neural differential equations are a powerful extension to this paradigm.

Widespread applications to improve traditional mathematical modelling (science, physics, engineering, etc.), and to modern machine learning problems (time series, generative modelling).



## Summary

---

Parameterised differential equations are already ubiquitous throughout mathematical modelling.

Neural differential equations are a powerful extension to this paradigm.

Widespread applications to improve traditional mathematical modelling (finance, physics, engineering, etc.), and to modern machine learning problems (time series, generative modelling).

Convince the rest of the world that they should be using neural differential equations too.

## References

---

R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, "Neural Ordinary Differential Equations", in *Advances in Neural Information Processing Systems* 31, pp. 6571-6583, Curran Associates, Inc., 2018

T. J. Lyons, "Differential equations driven by rough signals", *Revista Matematica Iberoamericana*, vol. 14, no. 2, pp. 215-310, 1998

P. Kidger, J. Morrill, J. Foster, and T. Lyons, "Neural Controlled Differential Equations for Irregular Time Series", *Advances in Neural Information Processing Systems* 33, 2020a

P. Kidger, J. Foster, X. Li, H. Oberhauser, and T. Lyons, "Neural SDEs as Infinite-Dimensional GANs", *Machine Learning and the Physical Sciences*, *NeurIPS* 2020b

# References

(Software)

---

torchdiffeq : <https://github.com/rtqichen/torchdiffeq>

torchsde : <https://github.com/google-research/torchsde>

torchcde : <https://github.com/patrick-kidger/torchcde>

DifferentialEquations.jl:

<https://github.com/SciML/DifferentialEquations.jl>

(A version of) these slides are available on my website.  
(<https://kldger.site>)

(A version of) these slides are available on my website.  
(<https://kidger.site>)

Feel free to send me an email / contact me on Twitter  
(@PatrickKidger) if you have any questions later.

(A version of) these slides are available on my website.  
(<https://kidger.site>)

Feel free to send me an email / contact me on Twitter  
(@PatrickKidger) if you have any questions later.

I'm a developer for most of the major PyTorch libraries so you  
can get in touch about those too.

(A version of) these slides are available on my website.  
(<https://kidger.site>)

Feel free to send me an email / contact me on Twitter  
(@PatrickKidger) if you have any questions later.

I'm a developer for most of the major PyTorch libraries so you  
can get in touch about those too.

Any questions now?