NEURAL CDES FOR LONG TIME SERIES VIA THE LOG-ODE METHOD

Anonymous authors
Paper under double-blind review

ABSTRACT

Neural Controlled Differential Equations (Neural CDEs) are the continuous-time analogue of an RNN, just as Neural ODEs are analogous to ResNets. However just like RNNs, training Neural CDEs can be difficult for long time series. Here, we propose to apply a technique drawn from stochastic analysis, namely the log-ODE method. Instead of using the original input sequence, our procedure summarises the information over local time intervals via the log-signature map, and uses the resulting shorter stream of log-signatures as the new input. This represents a length/channel trade-off. In doing so we demonstrate efficacy on problems of length up to 17k observations and observe significant training speed-ups, improvements in model performance, and reduced memory requirements compared to the existing algorithm.

1 INTRODUCTION

Neural controlled differential equations (Neural CDEs) (Kidger et al., 2020) are the continuous-time analogue to a recurrent neural network (RNN), and provide a natural method for modelling temporal dynamics with neural networks.

Neural CDEs are similar to neural ordinary differential equations (Neural ODEs), as popularised by Chen et al. (2018). A Neural ODE is determined by its initial condition, without a direct way to modify the trajectory given subsequent observations. In contrast the vector field of a Neural CDE depends upon the time-varying data, so that the trajectory of the system is driven by a sequence of observations.

1.1 CONTROLLED DIFFERENTIAL EQUATIONS

We begin by stating the definition of a CDE. The only piece of theory that may be unfamiliar is that of a Riemann–Stieltjes integral.

Let $a, b \in \mathbb{R}$ with $a < b$, and let $v, w \in \mathbb{N}$. Let $\xi \in \mathbb{R}^w$. Let $X \colon [a, b] \to \mathbb{R}^v$ be a continuous function of bounded variation (which is for example implied by it being Lipschitz), and let $f : \mathbb{R}^w \to \mathbb{R}^{w \times v}$ be continuous.

Then we may define $Z : [a, b] \to \mathbb{R}^w$ as the unique solution of the *controlled differential equation*

$$Z_a = \xi, \qquad Z_t = Z_a + \int_a^t f(Z_s)\mathrm{d}X_s \quad \text{for } t \in (a, b], \tag{1}$$

where the integral is a Riemann–Stieltjes integral. The notation "$f(Z_s)\mathrm{d}X_s$" denotes a matrix-vector product, and if X is differentiable then $\int_a^t f(Z_s)\mathrm{d}X_s = \int_a^t f(Z_s)\frac{\mathrm{d}X}{\mathrm{d}s}(s)\mathrm{d}s$.

If in equation (1), $\mathrm{d}X_s$ was replaced with $\mathrm{d}s$, then the equation would just be an ODE. Using $\mathrm{d}X_s$ causes the solution to depend continuously on the evolution of X.

1.2 NEURAL CONTROLLED DIFFERENTIAL EQUATIONS

We recall the definition of a Neural CDE as introduced in Kidger et al. (2020).

Consider a time series $\mathbf{x}$ as a collection of points $x_i \in \mathbb{R}^{v-1}$ with corresponding time-stamps $t_i \in \mathbb{R}$ such that $\mathbf{x} = ((t_0, x_0), (t_1, x_1), ..., (t_n, x_n))$, and $t_0 < ... < t_n$.

Let $X \colon [t_0, t_n] \to \mathbb{R}^v$ be some interpolation of the data such that $X_{t_i} = (t_i, x_i)$. Kidger et al. (2020) use natural cubic splines. Here we will actually end up finding piecewise linear interpolation to be a more convenient choice. (We avoid issues with adaptive solvers as discussed in Kidger et al. (2020, Appendix A) simply by using fixed solvers.)

Let $\xi_\theta \colon \mathbb{R}^v \to \mathbb{R}^w$ and $f_\theta \colon \mathbb{R}^w \to \mathbb{R}^{w \times v}$ be neural networks. Let $\ell_\theta \colon \mathbb{R}^w \to \mathbb{R}^q$ be linear, for some output dimension $q \in \mathbb{N}$. Here θ is used to denote dependence on learnable parameters.

We define Z as the hidden state and Y as the output of a *neural controlled differential equation* driven by X if

$$Z_{t_0} = \xi_\theta(t_0, x_0), \quad \text{with} \quad Z_t = Z_{t_0} + \int_{t_0}^{t} f_\theta(Z_s) \mathrm{d}X_s \quad \text{and} \quad Y_t = \ell_\theta(Z_t) \quad \text{for } t \in (t_0, t_n]. \tag{2}$$

That is – just like an RNN – we have evolving hidden state Z, which we take a linear map from to produce an output. This formulation is a universal approximator (Kidger et al., 2020, Appendix B). The output may be either the time-evolving Y_t or just the final Y_{t_n}. This is then fed into a loss function (L^2, cross entropy, ...) and trained via stochastic gradient descent in the usual way.

The question remains how to compute the integral of equation (2). Kidger et al. (2020) let

$$g_{\theta,X}(Z, s) = f_\theta(Z) \frac{\mathrm{d}X}{\mathrm{d}s}(s), \tag{3}$$

where the right hand side denotes a matrix multiplication, and then note that the integral can be written as

$$Z_t = Z_{t_0} + \int_{t_0}^{t} g_{\theta,X}(Z_s, s) \mathrm{d}s. \tag{4}$$

This reduces the CDE to an ODE, so that existing tools for Neural ODEs may be used to evaluate this, and to backpropagate.

By moving from the discrete-time formulation of an RNN to the continuous-time formulation of a Neural CDE, then every kind of time series data is put on the same footing, whether it is regularly or irregularly sampled, whether or not it has missing values, and whether or not the input sequences are of consistent length.

Besides this, the continuous-time or differential equation formulation may be useful in applications where such models are explicitly desired, as when modelling physics.

1.3 Contributions

Neural CDEs, as with RNNs, begin to break down for long time series. Training loss/accuracy worsens, and training time becomes prohibitive due to the sheer number of forward operations within each training epoch.

Here, we apply the *log-ODE method*, which is a numerical method from stochastic analysis and rough analysis. It is a method for converting a CDE to an ODE, which may in turn be solved via standard ODE solvers. Thus this acts as a drop-in replacement for the original procedure that uses the derivative of the control path.

In particular, we find that this method is particularly beneficial for long time series (and incidentally does not require differentiability of the control path.) With this method both training time and model performance of Neural CDEs are improved, and memory requirements are reduced.

The resulting scheme has two very neat interpretations. In terms of numerical differential equation solvers, this corresponds to taking integration steps larger than the discretisation of the data, whilst incorporating substep information through additional terms[1]. In terms of machine learning, this corresponds to binning the data prior to running a Neural CDE, with bin statistics carefully chosen to extract precisely the information most relevant to solving a CDE.

[1]For the reader familiar with numerical methods for SDEs, this is akin to the additional correction term in Milstein's method as compared to Euler-Maruyama.

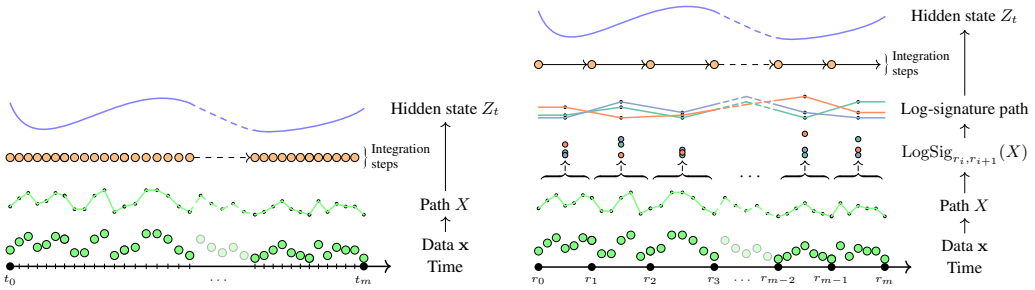

Figure 1: **Left:** The original Neural CDE formulation. The path X is quickly varying, meaning a lot of integration steps are needed to resolve it. **Right:** The log-ODE method. The log-signature path is more slowly varying (in a higher dimensional space), and needs fewer integration steps to resolve.

2 Theory

We begin with the motivating theory. Readers more interested in practical application should feel free to skip to section 3.

2.1 Signatures and Log-signatures

Signature transform Consider a path $X = (X^1, \ldots, X^v)\colon [a, b] \to \mathbb{R}^v$. We define its *signature transform* as the infinite collection of statistics

$$S_{a,b}(X) = \left(\left\{S_{a,b}^{i_1}(X)\right\}_{1\le i_1\le v}, \left\{S_{a,b}^{i_1,i_2}(X)\right\}_{1\le i_1,i_2\le v}, \left\{S_{a,b}^{i_1,i_2,i_3}(X)\right\}_{1\le i_1,i_2,i_3\le v}, \ldots\right), \tag{5}$$

where each term is a n-fold iterated integral of $\mathbf{x}$ with multi-index $i_1, ..., i_d$:

$$S_{a,b}^{i_1,\ldots,i_d}(X) := \int_a^b \int_a^{t_{d-1}} \cdots \int_a^{t_2} \mathrm{d}X_{t_1}^{i_1} \cdots \mathrm{d}X_{t_d}^{i_d} \in \mathbb{R}. \tag{6}$$

This produces a graded sequence of statistics associated with a path. Hambly & Lyons (2010) show that these statistics in fact characterise the path up to "tree-like equivalence".

As S is an infinite sequence of values, we also define the *depth-N signature transform* as

$$S_{a,b}^N(X) = \left(\left\{S_{a,b}^{i_1}(X)\right\}_{1\le i_1\le v}, \ldots, \left\{S_{a,b}^{i_1,\ldots,i_N}(X)\right\}_{1\le i_1,\ldots,i_N\le v}\right). \tag{7}$$

See Lyons et al. (2007) or Bonnier et al. (2019, Appendix A).

Log-signature transform However, the signature transform has some redundancy: a little algebra shows that for example $S_{a,b}^{1,2}(X) + S_{a,b}^{2,1}(X) = S_{a,b}^1(X)S_{a,b}^2(X)$, so that for instance we already know $S_{a,b}^{2,1}(X)$ provided we know the other three quantities.

The *log-signature transform* is then essentially obtained by computing the signature transform, and throwing out redundant terms, to obtain some minimal collection.

Such a minimal collection is nonunique, as we could have thrown out $S_{a,b}^{1,2}(X)$ instead of $S_{a,b}^{2,1}(X)$, for example. It turns out that the log-signature exhibits a free Lie algebra structure, and different possible minimal collections correspond to different possible bases of the free Lie algebra. A standard choice is the Lyndon basis, which corresponds to embedding the signature in a tensor algebra, performing a logarithm in the tensor algebra, keeping only those terms whose multi-index $i_1, \ldots, i_d$ is a Lyndon word, and then performing a particular sparse linear transformation (Reizenstein, 2017; Reizenstein & Graham, 2018). See Appendix A for a full exposition.

Starting from the depth-N signature transform and performing this procedure produces the *depth-N log-signature transform*.[2] We denote this $\mathrm{LogSig}_{a,b}^N$, which is a map from Hölder continuous paths

[2]Note that similar terminology such as "step-N log-signature" is also used in the literature.

$[a, b] \to \mathbb{R}^v$ into $\mathbb{R}^{\beta(v,N)}$, where

$$\beta(d, N) = \sum_{k=1}^{N} \frac{1}{k} \sum_{i|k} \mu\left(\frac{k}{i}\right) d^i$$

with μ the Möbius function.

Like the signature, the log-signature has a graded structure, where the depth-N+1 log-signature consists of the depth-N log-signature concatenated with additional terms.

The increment

$$\left\{S^{i_1}_{a,b}(X)\right\}_{1 \le i_1 \le v} = \left\{X^{i_1}_b - X^{i_1}_a\right\}_{1 \le i_1 \le v} = X_b - X_a$$

always appears as the lowest ($N = 1$) term in the log-signature.

2.2 The Log-ODE Method

First, recall the CDE of equation (1). Now let $N \in \mathbb{N}$ be a fixed depth. Decompose $f \colon \mathbb{R}^w \to \mathbb{R}^{w \times v}$ into vector fields $f_i \colon \mathbb{R}^w \to \mathbb{R}^w$ for $i \in \{1, \ldots, v\}$, and then let $\widehat{f} = (f_{i 1 \le i \le v}, [f_i, f_j]_{1 \le i < j \le v}, \ldots)$ be the collection of all (nested up depth N) Lie brackets whose foliage[3] is a Lyndon word, so that overall $\widehat{f} \colon \mathbb{R}^w \to \mathbb{R}^{w \times \beta(v,N)}$.

Recall for $X \colon [a, b] \to \mathbb{R}^v$ that $\mathrm{LogSig}^N_{a,b}(X) \in \mathbb{R}^{\beta(v,N)}$. Then the log-ODE method states

$$Z_b \approx \widehat{Z}_b \quad \text{where} \quad \widehat{Z}_u = \widehat{Z}_a + \int_a^u \widehat{f}(\widehat{Z}_s) \mathrm{LogSig}^N_{a,b}(X) \mathrm{d}s, \quad \text{and} \quad \widehat{Z}_a = Z_a. \tag{8}$$

The approximation becomes arbitrarily good[4] as $N \to \infty$.

That is, the solution of the CDE may be approximated by the solution to an ODE. In practice, we go further and pick some points r_i such that $a = r_0 < r_1 < \cdots < r_m = b$. We split up the CDE of equation (1) into an integral over $[r_0, r_1]$, an integral over $[r_1, r_2]$, and so on, and apply the log-ODE method to each interval separately.

See Appendix A for more details and Appendix B for a proof of convergence.

Also see Janssen (2011); Lyons (2014); Boutaib et al. (2014) for other discussions of the log-ODE method. See Gaines & Lyons (1997); Gyurkó & Lyons (2008); Flint & Lyons (2015); Foster et al. (2020) for applications of the log-ODE method to stochastic differential equations (SDEs).

The log-ODE method is a generalisation of the Magnus expansion for linear differential equations, see for example Blanes et al. (2009). There, the control path is obtained by integrating the time-varying vector field and the log-ODE solve reduces to exponentiating a matrix.

3 Method

We move on to discussing the application of the log-ODE method to Neural CDEs.

Recall that we observe some time series $\mathbf{x} = ((t_0, x_0), (t_1, x_1), ..., (t_n, x_n))$, and have constructed a piecewise linear interpolation $X \colon [t_0, t_n] \to \mathbb{R}^v$ such that $X_{t_i} = (t_i, x_i)$.

We now pick points r_i such that $t_0 = r_0 < r_1 < \cdots < r_m = t_n$. In principle these can be variably spaced but in practice we will typically space them equally far apart. The total number of points m should be much smaller than n.

In section 2 the log-signature transform was introduced. For $X \colon [t_0, t_n] \to \mathbb{R}^v$ and $t_0 \le r_i < r_{i+1} \le t_n$ then the depth-N log-signature of X over the interval $[r_i, r_{i+1}]$ is some collection of statistics $\mathrm{LogSig}^N_{r_i, r_{i+1}}(X) \in \mathbb{R}^{\beta(v,N)}$.

[3]The map produced by "ignoring the Lie brackets", for example $[f_1, [[f_2, f_1], f_3]] \mapsto (1, 2, 1, 3)$.

[4]With respect to any norm as $Z_b, \widehat{Z}_b$ exist in finite dimensions.

3.1 UPDATING THE NEURAL CDE HIDDEN STATE EQUATION VIA THE LOG-ODE METHOD

Recall how the Neural CDE formulation of equation (2) was solved via equations (3), (4). For the log-ODE approach we replace (3) with the piecewise

$$\widehat{g}_{\theta,X}(Z,s) = \widehat{f}_\theta(Z)\mathrm{LogSig}^N_{r_i,r_{i+1}}(X) \quad \text{for } s \in [r_i, r_{i+1}), \tag{9}$$

where $\widehat{f}_\theta \colon \mathbb{R}^w \to \mathbb{R}^{w \times \beta(v,N)}$ is an arbitrary neural network, and the right hand side denotes a matrix-vector product between $\widehat{f}_\theta$ and the log-signature. Equation (4) then becomes

$$Z_t = Z_{t_0} + \int_{t_0}^{t} \widehat{g}_{\theta,X}(Z_s, s)\mathrm{d}s. \tag{10}$$

This may now be solved as a (neural) ODE using standard ODE solvers.

3.2 DISCUSSION

Ease of Implementation This method is straightforward to implement using pre-existing tools. There are standard libraries available for computing the log-signature transform; we use Signatory (Kidger & Lyons, 2020b). Then, as equation (10) is an ODE, it may be solved directly using tools such as `torchdiffeq` (Chen, 2018).

As an alternative, we note that the formulation in equation (9) can be written in precisely the same form as equation (3), with the driving path taken to be piecewise linear in log-signature space. Computation of the log-signatures can therefore be considered as a preprocessing step, producing a sequence of logsignatures. From this we may construct a path in log-signature space, and apply existing tools for neural CDEs. This idea is summarised in figure 1. We make this approach available in the [redacted] open source project.

Lie Bracket Structure The description here aligns with the log-ODE scheme described in equation (8). There is one discrepancy: we do not attempt to model the Lie bracket structure of $\widehat{f}$. This is in principle possible (even straightforward with autodifferentiation), but is computationally expensive. Instead, we model $\widehat{f}$ as a neural network directly. This need *not* necessarily exhibit a Lie bracket structure, but as neural networks are universal approximators (Pinkus, 1999; Kidger & Lyons, 2020a) then this approach is at least as general from a modelling perspective.

Lossy Representation The log-signature transform can be thought of as a lossy representation for time series. This is made rigorous in Diehl et al. (2020), where it is shown that the log-signature can be obtained by iterating an "area" operation between paths. For CDEs, these geometric features precisely encode the interaction between the data and the system.

Length/Channel Trade-Off The sequence of log-signatures is now of length m, which was chosen to be much smaller than n. As such, it is much more slowly varying over the interval $[t_0, t_n]$ than the original data, which was of length n. The differential equation it drives is better behaved, and so larger integration steps may be used in the numerical solver. This is the source of the speed-ups of this method; we observe typical speed-ups by a factor of about 100.

Each element is a log-signature of size $\beta(v, N) \geq v$; the additional channels are higher-order corrections to compensate for the larger integration steps.

Generality of the Log-ODE Method If depth $N = 1$ and steps $r_i = t_i$ are used, then the above formulation exactly reduces onto the original Neural CDE formulation using linear interpolation. Thus the log-ODE method in fact generalises the original approach.

Applications In principle the log-ODE method may be applied to solve any Neural CDE. In practice, the reduction in length (from n to m), coupled with the loss of information (from using the log-signature as a summary statistic) makes this particularly useful for long time series.

Memory Efficiency Long sequences need large amounts of memory to perform backpropagation-through-time (BPTT). As with the original Neural CDEs, the log-ODE approach supports memory-efficient backpropagation via the adjoint equations, alleviating this issue. See Kidger et al. (2020).

Model	Step	Test Accuracy (%)	Time (Hrs)	Memory (Mb)
NCDE_1	1	62.4 ± 12.1	22.0	176.5
	8	64.1 ± 13.3	3.1	24.3
	32	64.1 ± 14.3	0.5	8.0
	128	48.7 ± 2.6	0.1	3.9
NCDE_2	8	$\mathbf{77.8 \pm 5.9}$	2.1	94.2
	32	67.5 ± 12.1	0.7	28.1
	128	$\mathbf{76.1 \pm 5.9}$	0.2	7.8
NCDE_3	8	70.1 ± 6.5	1.3	460.7
	32	$\mathbf{75.2 \pm 3.0}$	0.6	134.7
	128	68.4 ± 8.2	0.1	53.3

Table 1: Mean and standard deviation of test set accuracy (in %) over three repeats, as well as memory usage and training time, on the EigenWorms dataset for depths 1–3 and a small selection of step sizes. The bold values denote that the model was the top performer for that step size.

The Depth and Step Hyperparameters To solve a Neural CDE accurately via the log-ODE method, we should be prepared to take the depth N suitably large, or the intervals $r_{i+1} - r_i$ suitably small. Accomplishing this would realistically require that they are taken very large or very small, respectively. Instead, we treat these as hyperparameters. This makes use of the log-ODE method a modelling choice rather than an implementation detail.

Increasing step size will lead to faster (but less informative) training by reducing the number of operations in the forward pass. Increasing depth will lead to slower (but more informative) training, as more information about each local interval is used in each update.

4 Experiments

We investigate solving a Neural CDE with and without the log-ODE method on four real-world problems. Every problem was chosen for its long length. The lengths are in fact sufficiently long that adjoint-based backpropagation (Chen et al., 2018) was needed to avoid running out of memory at any reasonable batch size. Every problem is regularly sampled, so we take $t_i = i$.

We will denote a Neural CDE model with log-ODE method, using depth N and step s, as NCDE_N^s. Taking $N = 1$ (and any s) corresponds to not using the log-ODE method, with the data subsampled at rate $1/s$, as per section 3.2. Thus we use NCDE_1^1 as our benchmark: no subsampling, no log-ODE method.

Each model is run three times and we report the mean and standard deviation of the test metrics along with the mean training times and memory usages.

For each task, the hyperparameters were selected by performing a grid search on the NCDE_1^s model, where s was chosen so that the length of the sequence was 500 steps. This was found to create a reasonable balance between training time and sequence length. (Doing hyperoptimisation on the baseline NCDE_1^1 model would have been more difficult due to the larger training times.)

Precise details of the experiments can be found in Appendices C and D.

4.1 Classifying EigenWorms

Our first example uses the EigenWorms dataset from the UEA archive from Bagnall et al. (2017). This consists of time series of length 17 984 and 6 channels (including time), corresponding to the movement of a roundworm. The goal is to classify each worm as either wild-type or one of four mutant-type classes.

See Table 1. We see that the straightforward NCDE_1^1 model takes roughly a day to train. Using the log-ODE method (NCDE_2, NCDE_3) speeds this up to take roughly minutes. Doing so additionally improves model performance dramatically, and reduces memory usage. Naive subsampling

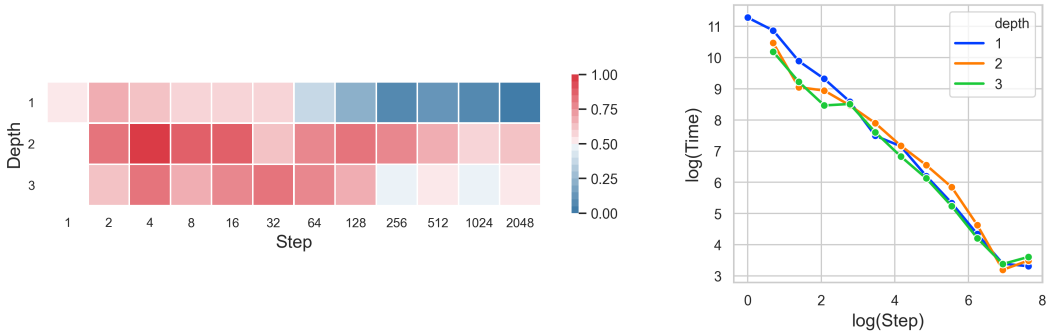

Figure 2: **Left:** Heatmap of normalised accuracies on the EigenWorms dataset for differing step sizes and depths. **Right:** Log-log plot of the elapsed time of the algorithm against the step size.

Depth	Step	L^2			Time (H)			Memory (Mb)
		RR	HR	SpO_2	RR	HR	SpO_2	
$NCDE_1$	1	2.79 ± 0.04	9.82 ± 0.34	2.83 ± 0.27	23.8	22.1	28.1	56.5
	8	2.80 ± 0.06	10.72 ± 0.24	3.43 ± 0.17	3.0	2.6	4.8	14.3
	32	2.53 ± 0.23	12.23 ± 0.43	2.68 ± 0.12	1.9	0.9	2.2	9.8
	128	2.64 ± 0.18	11.98 ± 0.37	2.86 ± 0.04	0.2	0.2	0.3	8.7
$NCDE_2$	8	2.63 ± 0.12	8.63 ± 0.24	2.88 ± 0.15	2.1	3.4	3.3	21.8
	32	1.90 ± 0.02	7.90 ± 1.00	1.69 ± 0.20	1.2	1.1	2.0	13.1
	128	1.86 ± 0.03	6.77 ± 0.42	1.95 ± 0.18	0.3	0.4	0.7	10.9
$NCDE_3$	8	$\mathbf{2.42 \pm 0.19}$	$\mathbf{7.67 \pm 0.40}$	$\mathbf{2.55 \pm 0.13}$	2.9	3.2	3.1	43.3
	32	$\mathbf{1.67 \pm 0.01}$	$\mathbf{4.50 \pm 0.70}$	$\mathbf{1.61 \pm 0.05}$	1.3	1.8	7.3	20.5
	128	$\mathbf{1.51 \pm 0.08}$	$\mathbf{2.97 \pm 0.45}$	$\mathbf{1.37 \pm 0.22}$	0.5	1.7	1.7	17.3

Table 2: Mean and standard deviation of the L^2 losses on the test set for each of the vitals signs prediction tasks (RR, HR, SpO_2) on the BIDMC dataset, across three repeats. Only mean times are shown for space. The memory usage is given as the mean over all three of the tasks as it was approximately the same for any task for a given depth and step. The bold values denote the algorithm with the lowest test set loss for a fixed step size for each task.

approaches ($NCDE_1^8$, $NCDE_1^{32}$, $NCDE_1^{128}$) only achieve speed-ups without performance improvements.

See also Figure 2, in which we summarise results for a larger range of step sizes.

4.2 Estimating Vitals Signs from PPG and ECG data

Next we consider the problem of estimating vital signs from PPG and ECG data. This comes from the TSR archive (Tan et al., 2020) using data from the Beth Israel Deaconess Medical Centre (BIDMC). We consider three separate tasks, in which we aim to predict a person's respiratory rate (RR), their heart rate (HR), and their oxygen saturation (SpO2). This data is sampled at 125Hz with each series having a length of 4 000. There are 7 949 training samples, and 3 channels (including time).

Figure 3: Heatmap of normalised losses on the three BIDMC datasets for differing step sizes and depths.

We train a model on each of the three vitals sign prediction tasks. The metric used to evaluate performance is the L^2 loss. The results over a range of step sizes are presented in table (2). We also provide heatmaps in Figure 3 for each dataset containing the loss values (normalised to $[0, 1]$) for each task. The full results over all step sizes may be found in Appendix D.

We find that the depth 3 model is the top performer for every task at any step size. What's more, it does so with a significantly reduced training time. We attribute the improved performance to the log-ODE model being better able to learn long-term dependencies due to the reduced sequence length. Note that the performance of the NCDE_2^s, NCDE_3^s models actually improves as the step size is increased. This is in contrast to NCDE_1^s, which sees a degradation in performance.

5 Limitations of the Log-ODE Method

Number of hyperparameters Two new hyperparameters – truncation depth and step size – with substantial effects on training time and memory usage must now also be tuned.

Number of input channels The log-ODE method is most feasible for low numbers of input channels, as the number of log-signature channels $\beta(v, N)$ grows exponentially in v.

6 Related Work

There has been some work on long time series for classic RNN (GRU/LSTM) models.

Wisdom et al. (2016); Jing et al. (2019) show that unitary or orthogonal RNNs can mitigate the vanishing/exploding gradients problem. However, they are expensive to train due to the need to compute a matrix inversion at each training step. Chang et al. (2017) introduce dilated RNNs with skip connections between RNN states, which help improve training speed and learning of long-term dependencies. Campos et al. (2017) introduce the 'Skip-RNN' model, which extend the RNN by adding an additional learnt component that skips state updates. Li et al. (2018) introduce the 'IndRNN' model, with particular structure tailored to learning long time series.

One important comparison is to hierarchical subsampling as in Graves (2012); De Mulder et al. (2015). There the data is split into windows, an RNN is run over each window, and then an additional RNN is run over the first RNN's outputs; we may describe this as an RNN/RNN pair. Liao et al. (2019) then perform the equivalent operation with a log-signature/RNN pair. In this context, our use of log-ODE method may then be described as a log-signature/NCDE pair. These connections are natural: at least within the signature/CDE/rough path community, it is a well-known but poorly-published fact that (log-)signatures, RNNs and (Neural) CDEs are all essentially equivalent; see for example Kidger et al. (2020) for a little exposition on this.

CNNs and Transformers have been shown to offer improvements over RNNs for modelling long-term dependencies (Bai et al., 2018; Li et al., 2019). However, both can be expensive in their own right; Transformers are famously $\mathcal{O}(L^2)$ in the length of the time series L. Whilst several approaches have been introduced to reduce this, for example Li et al. (2019) reduce this to $\mathcal{O}(L(\log L)^2)$, this can still be difficult with long series. Extensions specifically to long sequences do exist (Sourkov, 2018), but these typically focus on language modelling rather than multivariate time series data.

De Brouwer et al. (2019); Lechner & Hasani (2020) amongst others consider continuous time analogues of GRUs and LSTMs, going some way to improving the learning of long-term dependencies. Voelker et al. (2019); Gu et al. (2020) consider links with ODEs and approximation theory, to improve the long-term memory capacity of RNNs.

7 Conclusion

We demonstrate how to effectively apply Neural CDEs to long (17k) time series, via the log-ODE method. The model may still be solved via ODE methods and thus retains adjoint backpropagation and continuous dynamics. In doing so we see significant training speed-ups, improvements in model performance, and reduced memory requirements.

REFERENCES

A. Bagnall, J. Lines, A. Bostrom, J. Large, and E. Keogh. The great time series classification bake off: a review and experimental evaluation of recent algorithmic advances. *Data Mining and Knowledge Discovery*, 31:606–660, 2017.

Shaojie Bai, J Zico Kolter, and Vladlen Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*, 2018.

Sergio Blanes, Fernando Casas, Jose-Angel Oteo, and Jose Ros. The Magnus expansion and some of its applications. *Physics Reports*, 470(5-6):151–238, 2009.

Patric Bonnier, Patrick Kidger, Imanol Perez Arribas, Cristopher Salvi, and Terry Lyons. Deep Signature Transforms. *Advances in Neural Information Processing Systems*, 2019.

Youness Boutaib, Lajos Gergely Gyurkó, Terry Lyons, and Danyu Yang. Dimension-free Euler estimates of rough differential equations. *Revue Roumaine de Mathmatiques Pures et Appliques*, 59, 2014.

Víctor Campos, Brendan Jou, Xavier Giró-i Nieto, Jordi Torres, and Shih-Fu Chang. Skip RNN: Learning to Skip State Updates in Recurrent Neural Networks. *arXiv preprint arXiv:1708.06834*, 2017.

Shiyu Chang, Yang Zhang, Wei Han, Mo Yu, Xiaoxiao Guo, Wei Tan, Xiaodong Cui, Michael Witbrock, Mark A Hasegawa-Johnson, and Thomas S Huang. Dilated recurrent neural networks. In *Advances in Neural Information Processing Systems*, pp. 77–87, 2017.

Ricky T. Q. Chen. `torchdiffeq`, 2018. `https://github.com/rtqichen/torchdiffeq`.

Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural Ordinary Differential Equations. *Advances in Neural Information Processing Systems*, 2018.

Edward De Brouwer, Jaak Simm, Adam Arany, and Yves Moreau. Gru-ode-bayes: Continuous modeling of sporadically-observed time series. In *Advances in Neural Information Processing Systems*, pp. 7379–7390, 2019.

Wim De Mulder, Steven Bethard, and Marie-Francine Moens. A survey on the application of recurrent neural networks to statistical language modeling. *Computer Speech & Language*, 30(1): 61–98, 2015.

Joscha Diehl, Terry Lyons, Rosa Preiß, and Jeremy Reizenstein. Areas of areas generate the shuffle algebra. *arXiv preprint arXiv:2002.02338*, 2020.

Guy Flint and Terry Lyons. Pathwise approximation of SDEs by coupling piecewise abelian rough paths. *arXiv preprint arXiv:1505.01298*, 2015.

James Foster, Harald Oberhauser, and Terry Lyons. An optimal polynomial approximation of Brownian motion. *SIAM Journal on Numerical Analysis*, 58(3):1393–1421, 2020.

Peter K Friz and Nicolas B Victoir. *Multidimensional stochastic processes as rough paths: theory and applications*, volume 120. Cambridge University Press, 2010.

Jessica G Gaines and Terry J Lyons. Variable step size control in the numerical solution of stochastic differential equations. *SIAM Journal on Applied Mathematics*, 57(5):1455–1484, 1997.

Alex Graves. Supervised sequence labelling. In *Supervised sequence labelling with recurrent neural networks*, pp. 5–13. Springer, 2012.

Albert Gu, Tri Dao, Stefano Ermon, Atri Rudra, and Christopher Re. HiPPO: Recurrent Memory with Optimal Polynomial Projections. *arXiv:2008.07669*, 2020.

Lajos Gergely Gyurkó. *Numerical methods for approximating solutions to rough differential equations*. DPhil thesis, University of Oxford, 2008.

Lajos Gergely Gyurkó and Terry Lyons. Rough paths based numerical algorithms in computational finance. *Mathematics in Finance: UIMP-RSME Lluis A. Santaló Summer School*, 2008.

Ben Hambly and Terry Lyons. Uniqueness for the signature of a path of bounded variation and the reduced path group. *Annals of Mathematics*, 171, 2010.

Arend Janssen. *Order book models, signatures and numerical approximations of rough differential equations*. PhD thesis, University of Oxford, 2011.

Li Jing, Caglar Gulcehre, John Peurifoy, Yichen Shen, Max Tegmark, Marin Soljacic, and Yoshua Bengio. Gated orthogonal recurrent units: On learning to forget. *Neural computation*, 31(4): 765–783, 2019.

Patrick Kidger and Terry Lyons. Universal Approximation with Deep Narrow Networks. *COLT 2020*, 2020a.

Patrick Kidger and Terry Lyons. Signatory: differentiable computations of the signature and logsignature transforms, on both CPU and GPU. *arXiv:2001.00706*, 2020b. URL `https://github.com/patrick-kidger/signatory`.

Patrick Kidger, James Morrill, James Foster, and Terry Lyons. Neural controlled differential equations for irregular time series. *arXiv preprint arXiv:2005.08926*, 2020.

Mathias Lechner and Ramin Hasani. Learning long-term dependencies in irregularly-sampled time series. *arXiv preprint arXiv:2006.04418*, 2020.

Shiyang Li, Xiaoyong Jin, Yao Xuan, Xiyou Zhou, Wenhu Chen, Yu-Xiang Wang, and Xifeng Yan. Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting. In *Advances in Neural Information Processing Systems*, pp. 5243–5253, 2019.

Shuai Li, Wanqing Li, Chris Cook, Ce Zhu, and Yanbo Gao. Independently recurrent neural network (indrnn): Building a longer and deeper rnn. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5457–5466, 2018.

Shujian Liao, Terry Lyons, Weixin Yang, and Hao Ni. Learning stochastic differential equations using RNN with log signature features. *arXiv preprint arXiv:1908.08286*, 2019.

Terry Lyons. Rough paths, signatures and the modelling of functions on streams. *Proceedings of the International Congress of Mathematicians*, 4, 2014.

Terry Lyons, Caruana Michael, and Lévy Thierry. *Differential equations driven by rough paths*. In École d'été de probabilités de Saint-Flour XXXIV-2004, edited by J. Picard in Volume 1908 of Lecture Notes in Mathematics, Berlin, Springer, 2007.

A. Pinkus. Approximation theory of the MLP model in neural networks. *Acta Numer.*, 8:143–195, 1999.

Jeremy Reizenstein. Calculation of Iterated-Integral Signatures and Log Signatures. *arXiv preprint arXiv:1712.02757*, 2017.

Jeremy Reizenstein and Benjamin Graham. The iisignature library: efficient calculation of iterated-integral signatures and log signatures. *arXiv preprint arXiv:1802.08252*, 2018.

Raymond A. Ryan. *Introduction to Tensor Products of Banach Spaces*. Springer Monographs in Mathematics, Springer, 2002.

Vsevolod Sourkov. Igloo: Slicing the features space to represent sequences. *arXiv preprint arXiv:1807.03402*, 2018.

Chang Wei Tan, Anthony Bagnall, Christoph Bergmeir, Eamonn Keogh, Francois Petitjean, and Geoffrey I. Webb. Monash university, uea, ucr time series regression archive, 2020. `http://timeseriesregression.org/`.

Aaron Voelker, Ivana Kajić, and Chris Eliasmith. Legendre memory units: Continuous-time representation in recurrent neural networks. In *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 2019.

Scott Wisdom, Thomas Powers, John Hershey, Jonathan Le Roux, and Les Atlas. Full-capacity unitary recurrent neural networks. In *Advances in neural information processing systems*, pp. 4880–4888, 2016.

Eugene Wong and Moshe Zakai. On the Convergence of Ordinary Integrals to Stochastic Integrals. *Annals of Mathematical Statistics*, 36(5):1560–1564, 1965.