
Latent ODEs for Irregularly-Sampled Time Series

Yulia Rubanova, Ricky T. Q. Chen, David Duvenaud
University of Toronto and the Vector Institute
{rubanova, rtqichen, duvenaud}@cs.toronto.edu

Abstract

Time series with non-uniform intervals occur in many applications, and are difficult to model using standard recurrent neural networks (RNNs). We generalize RNNs to have continuous-time hidden dynamics defined by ordinary differential equations (ODEs), a model we call ODE-RNNs. Furthermore, we use ODE-RNNs to replace the recognition network of the recently-proposed Latent ODE model. Both ODE-RNNs and Latent ODEs can naturally handle arbitrary time gaps between observations, and can explicitly model the probability of observation times using Poisson processes. We show experimentally that these ODE-based models outperform their RNN-based counterparts on irregularly-sampled data.

1 Introduction

Recurrent neural networks (RNNs) are the dominant model class for high-dimensional, regularly-sampled time series data, such as text or speech. However, they are an awkward fit for irregularly-sampled time series data, common in medical or business settings. A standard trick for applying RNNs to irregular time series is to divide the timeline into equally-sized intervals, and impute or aggregate observations using averages. Such preprocessing destroys information, particularly about the timing of measurements, which can be informative about latent variables [Lipton et al., 2016, Che et al., 2018].

An approach which better matches reality is to construct a continuous-time model with a latent state defined at all times. Recently, steps have been taken in this direction, defining RNNs with continuous dynamics given by a simple exponential decay between observations [Che et al., 2018, Cao et al., 2018, Rajkomar et al., 2018, Mei and Eisner, 2017].

We generalize state transitions in RNNs to continuous-time dynamics specified by a neural network, as in Neural ODEs [Chen et al., 2018]. We call this model the ODE-RNN, and use it to construct two distinct continuous-time models. First, we use it as a standalone autoregressive model. Second, we refine the Latent ODE model of Chen et al. [2018] by using the ODE-RNN as a recognition network. Latent ODEs define a generative process over time series based on the deterministic evolution of an initial latent state, and can be trained as a variational autoencoder [Kingma and Welling, 2013]. Both models

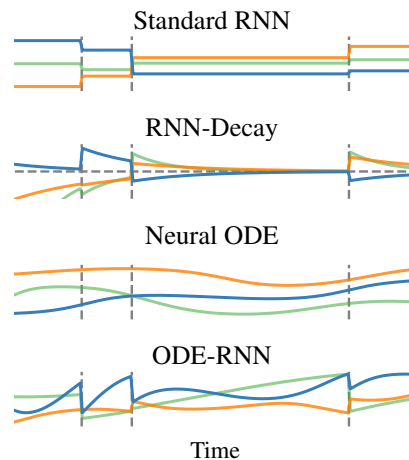


Figure 1: Hidden state trajectories. Vertical lines show observation times. Lines show different dimensions of the hidden state. Standard RNNs have constant or undefined hidden states between observations. The RNN-Decay model has states which exponentially decay towards zero, and are updated at observations. States of Neural ODE follow a complex trajectory but are determined by the initial state. The ODE-RNN model has states which obey an ODE between observations, and are also updated at observations.

naturally handle time gaps between observations, and remove the need to group observations into equally-timed bins. We compare ODE models to several RNN variants and find that ODE-RNNs can perform better when the data is sparse. Since the absence of observations itself can be informative, we further augment Latent ODEs to jointly model times of observations using a Poisson process.

2 Background

Recurrent neural networks A simple way to handle irregularly-timed samples is to include the time gap between observations $\Delta_t = t_i - t_{i-1}$ into the update function of the RNN:

$$h_i = \text{RNNCell}(h_{i-1}, \Delta_t, x_i) \quad (1)$$

However, this approach raises the question of how to define the hidden state h between observations. A simple alternative introduces an exponential decay of the hidden state towards zero when no observations are made [Che et al., 2018, Cao et al., 2018, Rajkomar et al., 2018, Mozer et al., 2017]:

$$h_i = \text{RNNCell}(h_{i-1} \cdot \exp\{-\tau\Delta_t\}, x_i) \quad (2)$$

where τ is a decay rate parameter. However, Mozer et al. [2017] found that empirically, exponential-decay dynamics did not improve predictive performance over standard RNN approaches.

Neural Ordinary Differential Equations Neural ODEs [Chen et al., 2018] are a family of continuous-time models which define a hidden state $h(t)$ as a solution to ODE initial-value problem:

$$\frac{dh(t)}{dt} = f_\theta(h(t), t) \quad \text{where} \quad h(t_0) = h_0 \quad (3)$$

in which the function f_θ specifies the dynamics of the hidden state, using a neural network with parameters θ . The hidden state $h(t)$ is defined at all times, and can be evaluated at any desired times using a numerical ODE solver:

$$h_0, \dots, h_N = \text{ODESolve}(f_\theta, h_0, (t_0, \dots, t_N)) \quad (4)$$

Chen et al. [2018] used the adjoint sensitivity method [Pontryagin et al., 1962] to compute memory-efficient gradients w.r.t. θ for training ODE-based deep learning models using black-box ODE solvers. They also conducted toy experiments in a time-series model in which the latent state follows a Neural ODE. Chen et al. [2018] used time-invariant dynamics in their time-series model: $\frac{dh(t)}{dt} = f_\theta(h(t))$, and we follow the same approach, but adding time-dependence would be straightforward if necessary.

3 Method

In this section, we use neural ODEs to define two distinct families of continuous-time models: the autoregressive ODE-RNN, and the variational-autoencoder-based Latent ODE.

3.1 Constructing an ODE-RNN Hybrid

Following Mozer et al. [2017], we note that an RNN with exponentially-decayed hidden state implicitly obeys the following ODE $\frac{dh(t)}{dt} = -\tau h$ with $h(t_0) = h_0$, where τ is a parameter of the model. The solution to this ODE is the pre-update term $h_0 \cdot \exp\{-\tau\Delta_t\}$ in (2). This differential equation is time-invariant, and assumes that the stationary point (i.e. zero-valued state) is special. We can generalize this approach and model the hidden state using a Neural ODE. The resulting algorithm is given in Algorithm 1. We define the state between observations to be the solution to an ODE: $h'_i = \text{ODESolve}(f_\theta, h_{i-1}, (t_{i-1}, t_i))$ and then at each observation, update the hidden state using a standard RNN update $h_i = \text{RNNCell}(h'_i, x_i)$. Our model does not explicitly depend on t or Δ_t when updating the hidden state, but does depend on time implicitly through the resulting dynamical system. Compared to RNNs with exponential decay, our approach allows a more flexible parameterization of the dynamics. A comparison between the state dynamics of these models is given in table 2.

Autoregressive Modeling with the ODE-RNN The ODE-RNN can straightforwardly be used to probabilistically model sequences. Consider a series of observations $\{x_i\}_{i=0}^N$ at times $\{t_i\}_{i=0}^N$. Autoregressive models make a one-step-ahead prediction conditioned on the history of observations, i.e. they factor the joint density $p(x) = \prod_i p_\theta(x_i | x_{i-1}, \dots, x_0)$. As in standard RNNs, we can use an ODE-RNN to specify the conditional distributions $p_\theta(x_i | x_{i-1} \dots x_0)$ (Algorithm 1).

Algorithm 1 The ODE-RNN. The only difference, highlighted in blue, from standard RNNs is that the pre-activations h' evolve according to an ODE between observations, instead of being fixed.

Input: Data points and their timestamps $\{(x_i, t_i)\}_{i=1..N}$
 $h_0 = \mathbf{0}$
for i in $1, 2, \dots, N$ **do**
 $h'_i = \text{ODESolve}(f_\theta, h_{i-1}, (t_{i-1}, t_i))$ ▷ Solve ODE to get state at t_i
 $h_i = \text{RNNCell}(h'_i, x_i)$ ▷ Update hidden state given current observation x_i
end for
 $o_i = \text{OutputNN}(h_i)$ for all $i = 1..N$
Return: $\{o_i\}_{i=1..N}; h_N$

3.2 Latent ODEs: a Latent-variable Construction

Autoregressive models such as RNNs and the ODE-RNN presented above are easy to train and allow fast online predictions. However, autoregressive models can be hard to interpret, since their update function combines both their model of system dynamics, and of conditioning on new observations. Furthermore, their hidden state does not explicitly encode uncertainty about the state of the true system. In terms of predictive accuracy, autoregressive models are often sufficient for densely sampled data, but perform worse when observations are sparse.

An alternative to autoregressive models are latent-variable models. For example, [Chen et al. \[2018\]](#) proposed a latent-variable time series model, where the generative model is defined by ODE whose initial latent state z_0 determines the entire trajectory:

$$z_0 \sim p(z_0) \tag{5}$$

$$z_0, z_1, \dots, z_N = \text{ODESolve}(f_\theta, z_0, (t_0, t_1, \dots, t_N)) \tag{6}$$

$$\text{each } x_i \stackrel{\text{indep.}}{\sim} p(x_i | z_i) \quad i = 0, 1, \dots, N \tag{7}$$

We follow [Chen et al. \[2018\]](#) in using a variational autoencoder framework for both training and prediction. This requires estimating the approximate posterior $q(z_0 | \{x_i, t_i\}_{i=0}^N)$. Inference and prediction in this model is effectively an encoder-decoder or sequence-to-sequence architecture, in which a variable-length sequence is encoded into a fixed-dimensional embedding, which is then decoded into another variable-length sequence, as in [Sutskever et al. \[2014\]](#).

Encoder-decoder models	Encoder	Decoder
Latent ODE (ODE enc.)	ODE-RNN	ODE
Latent ODE (RNN enc.)	RNN	ODE
RNN-VAE	RNN	RNN

Table 1: Different encoder-decoder architectures.

[Chen et al. \[2018\]](#) used an RNN as a recognition network to compute this approximate posterior. We conjecture that using an ODE-RNN as defined above for the recognition network would be a more effective parameterization when the datapoints are irregularly sampled. Thus, we propose using an ODE-RNN as the encoder for a latent ODE model, resulting in a fully ODE-based sequence-to-sequence model. In our approach, the mean and standard deviation of the approximate posterior $q(z_0 | \{x_i, t_i\}_{i=0}^N)$ are a function of the final hidden state of an ODE-RNN:

$$q(z_0 | \{x_i, t_i\}_{i=0}^N) = \mathcal{N}(\mu_{z_0}, \sigma_{z_0}) \quad \text{where} \quad \mu_{z_0}, \sigma_{z_0} = g(\text{ODE-RNN}_\phi(\{x_i, t_i\}_{i=0}^N)) \tag{8}$$

Where g is a neural network translating the final hidden state of the ODE-RNN encoder into the mean and variance of z_0 . To get the approximate posterior at time point t_0 , we run the ODE-RNN encoder backwards-in-time from t_N to t_0 . We jointly train both the encoder and decoder by maximizing the

Table 2: Definition of hidden state $h(t)$ between observation times t_{i-1} and t_i in autoregressive models. In standard RNNs, the hidden state does not change between updates. In ODE-RNNs, the hidden state is defined by an ODE, and is additionally updated by another network at each observation.

Model	State $h(t_i)$ between observations
Standard RNN	$h_{t_{i-1}}$
RNN-Decay	$h_{t_{i-1}} e^{-\tau \Delta t}$
GRU-D	$h_{t_{i-1}} e^{-\tau \Delta t}$
ODE-RNN	$\text{ODESolve}(f_\theta, h_{i-1}, (t_{i-1}, t))$




Figure 2: The Latent ODE model with an ODE-RNN encoder. To make predictions in this model, the ODE-RNN encoder is run backwards in time to produce an approximate posterior over the initial state: $q(z_0|\{x_i, t_i\}_{i=0}^N)$. Given a sample of z_0 , we can find the latent state at any point of interest by solving an ODE initial-value problem. Figure adapted from [Chen et al. \[2018\]](#).

evidence lower bound (ELBO):

$$\text{ELBO}(\theta, \phi) = \mathbb{E}_{z_0 \sim q_\phi(z_0|\{x_i, t_i\}_{i=0}^N)} [\log p_\theta(x_0, \dots, x_N)] - \text{KL} [q_\phi(z_0|\{x_i, t_i\}_{i=0}^N) || p(z_0)] \quad (9)$$

This latent variable framework comes with several benefits: First, it explicitly decouples the dynamics of the system (ODE), the likelihood of observations, and the recognition model, allowing each to be examined or specified on its own. Second, the posterior distribution over latent states provides an explicit measure of uncertainty, which is not available in standard RNNs and ODE-RNNs. Finally, it becomes easier to answer non-standard queries, such as making predictions backwards in time, or conditioning on a subset of observations.

3.3 Poisson process likelihoods

The fact that a measurement was made at a particular time is often informative about the state of the system [[Che et al., 2018](#)]. In the ODE framework, we can use the continuous latent state to parameterize the intensity of events using an *inhomogeneous* Poisson point process [[Palm, 1943](#)] where the event rate $\lambda(t)$ changes over time. Poisson point processes have the following log-likelihood:

$$\log p(t_1, \dots, t_N | t_{\text{start}}, t_{\text{end}}, \lambda(\cdot)) = \sum_{i=1}^N \log \lambda(t_i) - \int_{t_{\text{start}}}^{t_{\text{end}}} \lambda(t) dt$$

Where t_{start} and t_{end} are the times at which observations started and stopped being recorded.

We augment the Latent ODE framework with a Poisson process over the observation times, where we parameterize $\lambda(t)$ as a function of $z(t)$. This means that instead of specifying and maximizing the conditional marginal likelihood $p(x_1, \dots, x_N | t_1, \dots, t_N, \theta)$, we can instead specify and maximizing the joint marginal likelihood $p(x_1, \dots, x_N, t_1, \dots, t_N, | \theta)$. To compute the joint likelihood, we can evaluate the Poisson intensity $\lambda(t)$, precisely estimate its integral, and the compute latent states at all required time points, using a single call to an ODE solver.

[Mei and Eisner \[2017\]](#) used a similar approach, but relied on a fixed time discretization to estimate the Poisson intensity. [Chen et al. \[2018\]](#) showed a toy example of using Latent ODEs with a Poisson process likelihood to fit latent dynamics from observation times alone. In section 4.4, we incorporate a Poisson process likelihood into a latent ODE to model observation rates in medical data.

3.4 Batching and computational complexity

One computational difficulty that arises from irregularly-sampled data is that observation times can be different for each time series in a minibatch. In order to solve all ODEs in a minibatch in sync, we must we must output the solution of the combined ODE at the union of all time points in the batch.

Taking the union of time points does not substantially hurt the runtime of the ODE solver, as the adaptive time stepping in ODE solvers is not sensitive to the number of time points ($t_1 \dots t_N$) at which




Figure 3: Visualization of the inferred Poisson rate $\lambda(t)$ (green line) for two selected features of different patients from the Physionet dataset. Vertical lines mark observation times.




Figure 4: (a) A Latent ODE model conditioned on a small subset of points. This model, trained on exactly 30 observations per time series, still correctly extrapolates when more observations are provided. (b) Trajectories sampled from the prior $p(z_0) \sim \text{Normal}(z_0; 0, I)$ of the trained model, then decoded into observation space.

the solver outputs the state. Instead, it depends on the length on the time interval $[t_1, t_N]$ and the complexity of the dynamics. (see suppl. figure 3). Thus, ODE-RNNs and Latent ODEs have a similar asymptotic time complexity to standard RNN models. However, as the ODE must be continuously solved even when no observations occur, the compute cost does not scale with the sparsity of the data, as it does in decay-RNNs. In our experiments, we found that the ODE-RNN takes 60% more time than the standard GRU to evaluate, and the Latent ODE required roughly twice the amount of time to evaluate than the ODE-RNN.

3.5 When should you use an ODE-based model over a standard RNN?

Standard RNNs ignore the time gaps between points. As such, standard RNNs work well on regularly spaced data, with few missing values, or when the time intervals between points are short.

Models with continuous-time latent state, such as the ODE-RNN or RNN-Decay, can be evaluated at any desired time point, and therefore are suitable for interpolation tasks. In these models, the future hidden states depend on the time since the last observation, also making them better suited for sparse and/or irregular data than standard RNNs. RNN-Decay enforces that the hidden state converges monotonically to a fixed point over time. In ODE-RNNs the form of the dynamics between the observations is learned rather than pre-defined. Thus, ODE-RNNs can be used on sparse and/or irregular data without making strong assumptions about the dynamics of the time series.

Latent variable models versus autoregressive models We refer to models which iteratively compute the joint distribution $p(x) = \prod_i p_\theta(x_i | x_{i-1}, \dots, x_0)$ as autoregressive models (e.g. RNNs and ODE-RNNs). We call models of the form $p(x) = \int \prod_i p(x_i | z_0) p(z_0) dz_0$ latent-variable models (e.g. Latent ODEs and RNN-VAEs).

In autoregressive models, both the dynamics and the conditioning on data are encoded implicitly through the hidden state updates, which makes them hard to interpret. In contrast, encoder-decoder models (Latent ODE and RNN-VAE) represent state explicitly through a vector z_t , and represent dynamics explicitly through a generative model. Latent states in these models can be used to compare different time series, for e.g. clustering or classification tasks, and their dynamics functions can be examined to identify the types of dynamics present in the dataset.

4 Experiments

4.1 Toy dataset

We tested our model on a toy dataset of 1,000 periodic trajectories with variable frequency and the same amplitude. We sampled the initial point from a standard Gaussian, and added Gaussian noise to the observations. Each trajectory has 100 irregularly-sampled time points. During training, we subsample a fixed number of points at random, and attempt to reconstruct the full set of 100 points.

Conditioning on sparse data Latent ODEs can often reconstruct trajectories reasonably well given a small subset of points, and provide an estimate of uncertainty over both the latent trajectories and

predicted observations. To demonstrate this, we trained a Latent ODE model to reconstruct the full trajectory (100 points) from a subset of 30 points. At test time, we conditioned this model on a subset of 10, 30 or 50 points. Conditioning on more points results in a better fit as well as smaller variance across the generated trajectories (fig. 4). Figure 4(b) demonstrates that the trajectories sampled from the prior of the trained model are also periodic.

Extrapolation Next, we show that a time-invariant ODE can recover stationary periodic dynamics from data automatically. Figure 5 shows a Latent ODE trained to condition on 20 points in the $[0; 2.5]$ interval (red area) and predict points on $[2.5; 5]$ interval (blue area). A Latent ODE with an ODE-RNN encoder was able to extrapolate the time series far beyond the training interval and maintain periodic dynamics. In contrast, a Latent ODE trained with RNN encoder as in Chen et al. [2018] did not extrapolate the periodic dynamics well.




Figure 5: **(a)** Approximate posterior samples from a Latent ODE trained with an RNN recognition network, as in Chen et al. [2018]. **(b)** Approximate posterior samples from a Latent ODE trained with an ODE-RNN recognition network (ours). At training time, the Latent ODE conditions on points in red area, and reconstruct points in blue area. At test time, we condition the model on 20 points in red area, and solve the generative ODE on a larger time interval.

4.2 Quantitative Evaluation

We evaluate the models quantitatively on two tasks: interpolation and extrapolation. On each dataset, we used 80% for training and 20% for test. See the supplement a detailed description.

Baselines In the class of autoregressive models, we compare ODE-RNNs to standard RNNs. We compared the following autoregressive models: (1) ODE-RNN (proposed) (2) A classic RNN where Δ_t is concatenated to the input (RNN- Δ_t) (3) An RNN with exponential decay on the hidden states $h \cdot e^{-\tau\Delta_t}$ (RNN-Decay) (4) An RNN with missing values imputed by a weighted average of previous value and empirical mean (RNN-Impute), and (5) GRU-D [Che et al., 2018] which combines exponential decay and the above imputation strategy. Among encoder-decoder models, we compare the Latent ODE to a variational autoencoder in which both the encoder and decoder are recurrent neural nets (RNN-VAE). The ODE-RNN can use any hidden state update formula for the RNNCell function in Algorithm 1. Throughout our experiments, we use the Gated Recurrent Unit (GRU) [Cho et al., 2014]. See the supplement for the architecture details.

Interpolation The standard RNN and the ODE-RNN are straightforward to apply to the interpolation task. To perform interpolation with a Latent ODE, we encode the time series backwards in time, compute the approximate posterior $q(z_0|\{x_i, t_i\}_{i=0}^N)$ at the first time point t_0 , sample the initial state of ODE z_0 , and generate mean observations at each observation time.

Extrapolation In the extrapolation setting, we use the standard RNN or ODE-RNN trained on the interpolation task, and then extrapolate the sequence by re-feeding previous predictions. To encourage extrapolation, we used scheduled sampling [Bengio et al., 2015], feeding previous predictions instead of observed data with probability 0.5 during training. One might expect that directly optimizing for extrapolation would perform best at extrapolation. Such a model would resemble an encoder-decoder model, which we consider separately below (the RNN-VAE). For extrapolation in encoder-decoder models, including the Latent ODE, we split the timeline in half. We encode the observations in the first half forward in time and reconstruct the second half.

4.3 MuJoCo Physics Simulation

Next, we demonstrated that ODE-based models can learn an approximation to simple Newtonian physics. To show this, we created a physical simulation using the ‘‘Hopper’’ model from the Deepmind Control Suite [Tassa et al., 2018]. We randomly sampled the initial position of the hopper and initial velocities such that hopper rotates in the air and falls on the ground (figure 6). These trajectories are deterministic functions of their initial states, which matches the assumptions made by the Latent ODE. The dataset is 14-dimensional, and we model it with a 15-dimensional latent state. We generated 10,000 sequences of 100 regularly-sampled time points each.

We perform both interpolation and extrapolation tasks on the MuJoCo dataset. During training, we subsampled a small percentage of time points to simulate sparse observation times. For evaluation, we measured the mean squared error (MSE) on the full time series.

Table 3: Test Mean Squared Error (MSE) ($\times 10^{-2}$) on the MuJoCo dataset.

Model	Interpolation (% Observed Pts.)				Extrapolation (% Observed Pts.)				
	10%	20%	30%	50%	10%	20%	30%	50%	
Autoreg	RNN Δ_t	2.454	1.714	1.250	0.785	7.259	6.792	6.594	30.571
	RNN GRU-D	1.968	1.421	1.134	0.748	38.130	20.041	13.049	5.833
	ODE-RNN (Ours)	1.647	1.209	0.986	0.665	13.508	31.950	15.465	26.463
Enc-Dec	RNN-VAE	6.514	6.408	6.305	6.100	2.378	2.135	2.021	1.782
	Latent ODE (RNN enc.)	2.477	0.578	2.768	0.447	1.663	1.653	1.485	1.377
	Latent ODE (ODE enc, ours)	0.360	0.295	0.300	0.285	1.441	1.400	1.175	1.258

Table 3 shows mean squared error for models trained on different percentages of observed points. Latent ODEs outperformed standard RNN-VAEs on both interpolation and extrapolation. Our ODE-RNN model also outperforms standard RNNs on the interpolation task. The gap in performance between RNN and ODE-RNN increases with sparser data. Notably, the Latent ODE (an encoder-decoder model) shows better performance than the ODE-RNN (an autoregressive model).

All autoregressive models performed poorly at extrapolation. This is expected, as they were only trained for one-step-ahead prediction, although standard RNNs performed better than ODE-RNNs. Latent ODEs outperformed RNN-VAEs on the extrapolation task.

Interpretability of the latent state Figure 6 shows how the norm of the latent state time-derivative $f_\theta(z)$ changes with time for two reconstructed MuJoCo trajectories. When the hopper hits the ground, there is a spike in the norm of the ODE function. In contrast, when the hopper is lying on the ground, the norm of the dynamics is small.

Figure 7 shows the entropy of the approximate posterior $q(z_0|\{x_i, t_i\}_{i=0}^N)$ of a trained model conditioned on different numbers of observations. The average entropy (uncertainty) monotonically decreases as more points are observed. Figure 8 shows the latent state z_0 projected to 2D using UMAP [McInnes et al., 2018]. The latent state corresponds closely to the physical parameters of the true simulation that most strongly determine the future trajectory of the hopper: distance from the ground, initial velocity on z-axis, and relative position of the leg of the hopper.

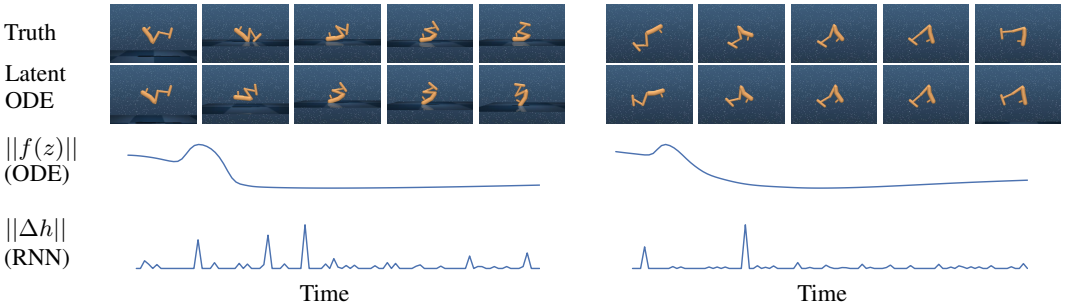


Figure 6: *Top row*: True trajectories from MuJoCo dataset. *Second row*: Trajectories reconstructed by a latent ODE model. *Third row*: Norm of the dynamics function f_θ in the latent space of the latent ODE model. *Fourth row*: Norm of the hidden state of a RNN trained on the same dataset.




Figure 7: Entropy of the approximate posterior over z_0 versus number of observed time points. The line shows the mean; shaded area shows 10% and 90% percentiles estimated over 1000 trajectories

Table 4: Test MSE (mean \pm std) on PhysioNet. **Autoregressive** models.

Model	Interp ($\times 10^{-3}$)
RNN Δ_t	3.520 \pm 0.276
RNN-Impute	3.243 \pm 0.275
RNN-Decay	3.215 \pm 0.276
RNN GRU-D	3.384 \pm 0.274
ODE-RNN (Ours)	2.361 \pm 0.086

4.4 Physionet

We evaluated our model on the PhysioNet Challenge 2012 dataset [Silva et al., 2012], which contains 8000 time series, each containing measurements from the first 48 hours of a different patient’s admission to ICU. Measurements were made at irregular times, and of varying sparse subsets of the 37 possible features.

Most existing approaches to modeling this data use a coarse discretization of the aggregated measurements per hour [Che et al., 2018], which forces the model to train on only one-twentieth of measurements. In contrast, our approach, in principle, does not require any discretization or aggregation of measurements. To speed up training, we rounded the observation times to the nearest minute, reducing the number of measurements only 2-fold. Hence, there are still 2880 (60*48) possible measurement times per time series under our model’s preprocessing, while the previous standard was to use only 48 possible measurement times. We used 20 latent dimensions in the latent ODE generative model. See supplement for more details on hyperparameters. Tables 4 and 5 report mean squared error averaged over runs with different random seeds, and their standard deviations. We run one-sided t-test to establish a statistical significance. Best models are marked in bold. ODE-based models have smaller mean squared error than RNN baselines on this dataset.

Finally, we constructed binary classifiers based on each model type to predict in-hospital mortality. We passed the hidden state at the last measured time point into a two-layer binary classifier. Due to class imbalance (13.75% samples with positive label), we report test area under curve (AUC) instead of accuracy. Table 6 shows that the ODE-RNN, Latent ODE and GRU-D achieved the similar classification AUC. A possible explanation is that modelling dynamics between time points does not make a difference for binary classification of the full time series.

We also included a Poisson Process likelihood on observation times, jointly trained with the Latent ODE model. Figure 3 shows the inferred measurement rate on a patient from the dataset. Although the Poisson process was able to model observation times reasonably well, including this likelihood term did not improve classification accuracy.

4.5 Human Activity dataset

We trained the same classifier models as above on the Human Activity dataset, which contains time series from five individuals performing various activities: walking, sitting, lying, etc. The




Figure 8: Nonlinear projection of latent space of z_0 from a Latent ODE model trained on the MuJoCo dataset). Each point is the encoding of one time series. The points are colored by the (a) initial height (distance from the ground) (b) initial velocity in z-axis (c) relative initial position of the hip of the hopper. The latent state corresponds closely to the physical parameters of the true simulation.

Table 5: Test MSE (mean \pm std) on PhysioNet. **Encoder-decoder** models.

Model	Interp ($\times 10^{-3}$)	Extrap ($\times 10^{-3}$)
RNN-VAE	5.930 \pm 0.249	3.055 \pm 0.145
Latent ODE (RNN enc.)	3.907 \pm 0.252	3.162 \pm 0.052
Latent ODE (ODE enc.)	2.118 \pm 0.271	2.231 \pm 0.029
Latent ODE + Poisson	2.789 \pm 0.771	2.208 \pm 0.050

Table 6: **Per-sequence classification.**
AUC on Physionet.

Method	AUC
RNN Δ_t	0.787 ± 0.014
RNN-Impute	0.764 ± 0.016
RNN-Decay	0.807 ± 0.003
RNN GRU-D	0.818 ± 0.008
RNN-VAE	0.515 ± 0.040
Latent ODE (RNN enc.)	0.781 ± 0.018
ODE-RNN	0.833 ± 0.009
Latent ODE (ODE enc.)	0.829 ± 0.004
Latent ODE + Poisson	0.826 ± 0.007

Table 7: **Per-time-point classification.**
Accuracy on Human Activity.

Method	Accuracy
RNN Δ_t	0.797 ± 0.003
RNN-Impute	0.795 ± 0.008
RNN-Decay	0.800 ± 0.010
RNN GRU-D	0.806 ± 0.007
RNN-VAE	0.343 ± 0.040
Latent ODE (RNN enc.)	0.835 ± 0.010
ODE-RNN	0.829 ± 0.016
Latent ODE (ODE enc.)	0.846 ± 0.013

data consists of 3d positions of tags attached to their belt, chest and ankles (12 features in total). After preprocessing, the dataset has 6554 sequences of 211 time points (details in supplement). The task is to classify each time point into one of seven types of activities (walking, sitting, etc.). We used a 15-dimensional latent state (more details in the supplement). Table 7 shows that the Latent ODE-based classifier had higher accuracy than the ODE-RNN classifier on this task.

5 Related work

Standard RNNs treat observations as a sequence of tokens, not accounting for variable gaps between observations. One way to accommodate this is to discretize the timeline into equal intervals, impute missing data, and then run an RNN on the imputed inputs. To perform imputation, Che et al. [2018] used a weighted average between the empirical mean and the previous observation. Others have used a separate interpolation network [Shukla and Marlin, 2019], Gaussian processes [Futoma et al., 2017], or generative adversarial networks [Luo et al., 2018] to perform interpolation and imputation prior to running an RNN on time-discretized inputs. In contrast, Lipton et al. [2016] used a binary mask to indicate the missing measurements and reported that RNNs performs better with zero-filling than with imputed values. They note that such methods can be sensitive to the discretization granularity.

Another approach is to directly incorporate the time gaps between observations into RNN. The simplest approach is to append the time gap Δ_t to the RNN input. However, Mozer et al. [2017] suggested that appending Δ_t makes the model prone to overfitting, and found empirically that it did not improve predictive performance. Another solution is to introduce the hidden states that decay exponentially over time [Che et al., 2018, Cao et al., 2018, Rajkomar et al., 2018].

Mei and Eisner [2017] used hidden states with exponential decay to parametrize neural Hawkes processes, and explicitly modeled observation intensities. Hawkes processes are self-exciting processes whose latent state changes at each observation event. This architecture is similar to our ODE-RNN. In contrast, the Latent ODE model assumes that observations do not affect the latent state, but only affect the model’s posterior over latent states, and is more appropriate when observations (such as taking a patient’s temperature) do not substantially alter their state. Ayed et al. [2019] used a Neural-ODE-based framework to learn the initial state and ODE parameters from a physical simulation. Concurrent work by De Brouwer et al. [2019] proposed an autoregressive model with ODE-based transitions between observation times and Bayesian updates of the hidden states.

6 Discussion and conclusion

We introduced a family of time series models, ODE-RNNs, whose hidden state dynamics are specified by neural ordinary differential equations (Neural ODEs). We first investigated this model as a standalone refinement of RNNs. We also used this model to improve the recognition networks of a variational autoencoder model known as Latent ODEs. Latent ODEs provide relatively interpretable latent states, as well explicit uncertainty estimates about latent states. Neither model requires discretizing observation times, or imputing data as a preprocessing step, making them suitable for the irregularly-sampled time series data common in many applications. Finally, we demonstrate that continuous-time latent states can be combined with Poisson process likelihoods to model the rates at which observations are made.

Acknowledgments

We thank Chun-Hao Chang, Chris Cremer, Quaid Morris, and Ladislav Rampasek for helpful discussions and feedback. We thank the Vector Institute for providing computational resources.

References

- Ibrahim Ayed, Emmanuel de Bézenac, Arthur Pajot, Julien Brajard, and Patrick Gallinari. Learning Dynamical Systems from Partial Observations. *arXiv e-prints*, art. arXiv:1902.11136, Feb 2019.
- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. Scheduled sampling for sequence prediction with recurrent neural networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, NIPS’15, pages 1171–1179, Cambridge, MA, USA, 2015. MIT Press.
- Cao, Wei, Wang, Dong, Li, Jian, Zhou, Hao, Li, Lei, and et al. Brits: Bidirectional recurrent imputation for time series, May 2018. URL <https://arxiv.org/abs/1805.10572>.
- Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. Recurrent Neural Networks for Multivariate Time Series with Missing Values. *Scientific Reports*, 8(1):6085, 2018. ISSN 2045-2322. doi: 10.1038/s41598-018-24271-9.
- Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems 31*, pages 6571–6583. Curran Associates, Inc., 2018.
- Cho, Kyunghyun, van Merriënboer, Bart, and Yoshua. On the properties of neural machine translation: Encoder-decoder approaches, Oct 2014. URL <https://arxiv.org/abs/1409.1259>.
- Edward De Brouwer, Jaak Simm, Adam Arany, and Yves Moreau. GRU-ODE-Bayes: Continuous modeling of sporadically-observed time series. *arXiv e-prints*, art. arXiv:1905.12374, May 2019.
- Joseph Futoma, Sanjay Hariharan, and Katherine Heller. Learning to detect sepsis with a multitask Gaussian process RNN classifier. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1174–1182, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Zachary C Lipton, David Kale, and Randall Wetzel. Directly modeling missing data in sequences with rnns: Improved classification of clinical time series. In Finale Doshi-Velez, Jim Fackler, David Kale, Byron Wallace, and Jenna Wiens, editors, *Proceedings of the 1st Machine Learning for Healthcare Conference*, volume 56 of *Proceedings of Machine Learning Research*, pages 253–270, Children’s Hospital LA, Los Angeles, CA, USA, 18–19 Aug 2016. PMLR.
- Yonghong Luo, Xiangrui Cai, Ying ZHANG, Jun Xu, and Yuan xiaojie. Multivariate time series imputation with generative adversarial networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 1596–1607. Curran Associates, Inc., 2018.
- Leland McInnes, John Healy, and James Melville. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv e-prints*, art. arXiv:1802.03426, Feb 2018.
- Hongyuan Mei and Jason M Eisner. The neural hawkes process: A neurally self-modulating multivariate point process. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 6754–6764. Curran Associates, Inc., 2017.
- Mozer, Michael C., Kazakov, Denis, Lindsey, and Robert V. Discrete event, continuous time rnns, Oct 2017. URL <https://arxiv.org/abs/1710.04110>.
- Conny Palm. Intensitätsschwankungen im fernsprecher. *Ericsson Technics*, 1943.

- Lev Semenovich Pontryagin, EF Mishchenko, VG Boltyanskii, and RV Gamkrelidze. The mathematical theory of optimal processes. 1962.
- Alvin Rajkomar, Eyal Oren, Kai Chen, Andrew M. Dai, Nissan Hajaj, Peter J. Liu, Xiaobing Liu, Mimi Sun, Patrik Sundberg, Hector Yee, Kun Zhang, Gavin Duggan, Gerardo Flores, Michaela Hardt, Jamie Irvine, Quoc Le, Kurt Litsch, Jake Marcus, Alexander Mossin, and Jeff Dean. Scalable and accurate deep learning for electronic health records. *npj Digital Medicine*, 1, 01 2018. doi: 10.1038/s41746-018-0029-1.
- Satya Narayan Shukla and Benjamin Marlin. Interpolation-prediction networks for irregularly sampled time series. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=r1efr3C9Ym>.
- Ikaro Silva, George Moody, Daniel J Scott, Leo A Celi, and Roger G Mark. Predicting In-Hospital Mortality of ICU Patients: The PhysioNet/Computing in Cardiology Challenge 2012. *Computing in cardiology*, 39:245–248, 2012. ISSN 2325-8861. URL <https://www.ncbi.nlm.nih.gov/pubmed/24678516><https://www.ncbi.nlm.nih.gov/pmc/PMC3965265/>.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014.
- Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, Timothy Lillicrap, and Martin Riedmiller. DeepMind Control Suite. *arXiv e-prints*, art. arXiv:1801.00690, Jan 2018.