
Segmenting Hybrid Trajectories using Latent ODEs

Ruian Shi^{1 2} Quaid Morris^{1 2 3}

Abstract

Smooth dynamics interrupted by discontinuities are known as hybrid systems and arise commonly in nature. Latent ODEs allow for powerful representation of irregularly sampled time series but are not designed to capture trajectories arising from hybrid systems. Here, we propose the Latent Segmented ODE (LatSegODE), which uses Latent ODEs to perform reconstruction and change-point detection within hybrid trajectories featuring jump discontinuities and switching dynamical modes. Where it is possible to train a Latent ODE on the smooth dynamical flows between discontinuities, we apply the pruned exact linear time (PELT) algorithm to detect changepoints where latent dynamics restart, thereby maximizing the joint probability of a piece-wise continuous latent dynamical representation. We propose usage of the marginal likelihood as a score function for PELT, circumventing the need for model-complexity-based penalization. The LatSegODE outperforms baselines in reconstructive and segmentation tasks including synthetic data sets of sine waves, Lotka Volterra dynamics, and UCI Character Trajectories.

1. Introduction

The complexity of modelling time-series data increases when accounting for discontinuous changes in dynamical behavior. As a motivational example, consider the Lotka-Volterra equations, a simplified model of predator-prey interactions. The system is described by the pair of ordinary differential equations (ODEs):

$$\frac{dx}{dt} = \alpha x - \beta xy \quad \frac{dy}{dt} = \delta xy - \gamma y \quad (1)$$

where x and y are the population size of predators and prey respectively. Coefficients $\alpha, \beta, \delta, \gamma$ describe interac-

tion characteristics, such as the rate of encounter, and rate of successful predation per encounter. When these parameters are fixed, modelling this system from observed population trajectories is straightforward. However, external factors can perturb the system. Additional predators can suddenly be introduced via migration midway in an observed population trajectory, causing a jump discontinuity in the trajectory. The coefficients describing predator-prey interaction may also abruptly change, instantaneously changing the dynamical mode of the system. Systems featuring smooth dynamical flows (SDFs) interrupted by discontinuities are known as hybrid systems (Van Der Schaft & Schumacher, 2000). These discontinuities can arise as discrete jumps or instantaneous switches in dynamical mode (Ackerson & Fu, 1970), shown in Figure 1 at times (a) and (b) respectively. We propose a method to model the hybrid trajectories which arise from hybrid systems.




Figure 1. A Lotka-Volterra hybrid trajectory composed of three smooth dynamical flows. The plot shows populations of predators and prey over time. At time (a), a jump discontinuity occurs. At time (b), a distributional shift in dynamical coefficients occurs.

Recently, the Latent ODE architecture (Rubanova et al., 2019) has been introduced to represent time series using latent dynamical trajectories. However, Latent ODEs are not designed to model discontinuous latent dynamics and, thus, represent hybrid trajectories poorly. Here, we propose the Latent Segmented ODE (LatSegODE), an extension of a Latent ODE explicitly designed for hybrid trajectories. Given a base model Latent ODE trained on the segments of SDFs between discontinuities, we apply the Pruned Exact Linear Time (PELT) search algorithm (Killick et al., 2012) to model hybrid trajectories as a sequence of samples from the base model, each with a different initial state. The LatSegODE detects the positions where the latent ODE dynamics are restarted with a new initial state, thus modelling hybrid trajectories using a piece-wise continuous latent trajectory. We provide a novel way to use deep architectures in conjunction

¹University of Toronto ²Vector Institute, Toronto ³Memorial Sloan Kettering Cancer Center. Correspondence to: Ruian Shi <ian.shi@mail.utoronto.ca>.

with offline changepoint detection (CPD) methods. Using the marginal likelihood under the Latent ODE as a score function, we find the Bayesian Occam’s Razor (MacKay, 1992) effect automatically prevents over-segmentation in CPD methods.

We evaluate LatSegODE on data sets of 1D sine wave hybrid trajectories, Lotka-Volterra hybrid trajectories, and a synthetically composed UCI Character Trajectories data set. We demonstrate that the LatSegODE interpolates, extrapolates, and finds the changepoints in hybrid trajectories with high accuracy compared to current baseline methods.

2. Background

2.1. Latent ODEs

The Latent ODE architecture (Rubanova et al., 2019) is an extension of the Neural ODE method (Chen et al., 2018), which provides memory-efficient gradient computation without back-propagation through ODE solve operations. Neural ODEs represent trajectories as the solution to the initial value problem:

$$\frac{dh(t)}{dt} = f_{\theta}(h(t), t) \quad (2)$$

$$h_{0:N} = \text{ODESolve}(f_{\theta}, h_0, t_{0:N}) \quad (3)$$

where f_{θ} is parameterized by a neural network, and $h(t)$ represents hidden dynamics. The continuous dynamical representation allows Neural ODEs to natively incorporate irregularly sampled time series.

Latent ODEs arrange Neural ODEs in an encoder-decoder architecture. Observed trajectories are encoded using a GRU-ODE architecture (Brouwer et al., 2019; Rubanova et al., 2019). The GRU-ODE combines a Neural ODE with a gated recurrent unit (GRU) (Cho et al., 2014). Observed trajectories are encoded by the GRU into a hidden state, which is continuously evolved between observations by a Neural ODE parameterized by neural network f_{θ} . The GRU-ODE encodes the observed data sequence into parameters for a variational posterior. Using the reparameterization trick (Kingma & Welling, 2014), a differentiable sample of the latent initial state z_0 is obtained. A Neural ODE parameterized by neural network f_{Ψ} deterministically solves a latent trajectory from the latent initial state. Finally, a neural network f_{Φ} decodes the latent trajectory into data space. The Latent ODE architecture can thus be represented as:

$$\mu_{z_0}, \sigma_{z_0}^2 = \text{GRUODE}_{f_{\theta}}(x_{1:N}, t_{1:N}) \quad (4)$$

$$z_0 \sim q(z_0|x_{1:N}) = \mathcal{N}(\mu_{z_0}, \sigma_{z_0}^2) \quad (5)$$

$$z_{1:N} = \text{ODESolve}(f_{\Psi}, z_0, t_{1:N}) \quad (6)$$

$$x_i \sim \mathcal{N}(f_{\Phi}(z_i), \sigma^2) \quad \text{for } i = 1, \dots, N \quad (7)$$

where σ^2 is a fixed variance term. The Latent ODE is trained by maximizing the evidence lower-bound (ELBO). Letting

$X = x_{1:N}$, the ELBO is:

$$\mathbb{E}_{z_0 \sim q(z_0|X)}[\log p(X)] - \text{KL}[q(z_0|X) || p(z_0)] \quad (8)$$

2.2. Representational Limitations of the Neural ODE

Latent ODEs use Neural ODEs to represent latent dynamics, and thus inherit their representational limitations. The accuracy of an ODE solver used by a Neural ODE depends on the smoothness of the solution; the local error of the solution can exceed ODE solver tolerances when a jump discontinuity occurs (Calvo et al., 2008). At a jump, adaptive ODE solvers will continuously reduce step size in response to increased error, possibly until numerical underflow occurs. Even if integration is possible across the jump, it is slow, and the global error of the solution can be adversely affected (Calvo et al., 2003). Typically, these issues can be easily avoided by restarting ODE solutions at the discontinuity but this requires these positions to be known. Classical methods use the increase in local error or adaptive rejections associated with jump discontinuity as criteria to restart solutions (Calvo et al., 2008). Recently, Neural Event ODEs (Chen et al., 2020) uses a similar paradigm of discontinuity detection, using an event function parameterized by a neural network to detect locations to restart the ODE solution. With all event detection approaches, failure to accurately detect jump discontinuity will cause the local error bound decrease to a lower order (Stewart, 2011). Hybrid trajectories with discontinuous change in the dynamical coefficients present different but still hard modeling challenges due to the representational limitations of Neural ODEs.

Latent ODEs do not circumvent these limitations, and cannot generalize in hybrid trajectories. When a hybrid trajectory is encountered, the Latent ODE can only encode the exact sequence of SDFs into a single latent representation. Should a permutation of these SDFs arise at test time, the Latent ODE will not be able to reconstruct the test trajectory.

3. Method

The LatSegODE detects positions of jump discontinuity or switching dynamical mode by representing a hybrid trajectory as a piece-wise combination of samples from a learned base model Latent ODE. At each changepoint, the latent dynamics of the base model are restarted from a new initial state. We apply the PELT algorithm to efficiently search through all possible positions to restart ODE dynamics, and return changepoints that correspond to the positions of restart which maximize the joint probability of a hybrid trajectory. This avoids the need to train an event detector, and guarantees optimal segmentation, but the LatSegODE requires the availability of a training data set of SDFs on which the base model can be trained.




Figure 2. Schematic of the LatSegODE reconstructing a hybrid trajectory. Arrows indicate computation flow. Data in each segment is encoded into parameters for the variational posterior, from which a latent initial state is sampled. Each latent segment is solved using shared latent dynamic f_Φ , which continues until the next point of change. The latent trajectory is decoded into data space. At evaluation time, an arbitrary number of changepoints can be detected by the PELT algorithm. Plot adapted from (Rubanova et al., 2019).

3.1. Extension to Hybrid Trajectories

We first define the class of hybrid trajectories which can be represented by the LatSegODE. Consider a sequential series of data $X = x_1, x_2, \dots, x_N$ and associated times of observation $T = t_1, t_2, \dots, t_N$. We represent a hybrid trajectory as a piece-wise sequence of C continuous dynamical segments. Each observed data point can only belong to a single segment. Each segment is bounded by starting index s_i and ending index e_i , where $0 \leq i \leq C$, $s_0 = 1$, and $e_C = N$. Segments are sequential and do not intersect, i.e., $s_{i+1} = e_i + 1$. The boundaries of segments represent locations of jump discontinuity or switch in dynamical mode. The trajectory within each segment is represented by a sample from the base model Latent ODE.

The LatSegODE can be applied to hybrid trajectories containing an unknown number and order of SDFs. The LatSegODE aims to approximate each SDF using a segment. Using offline CPD, the LatSegODE detects positions of jump discontinuity or switching dynamical mode, and introduces a latent discontinuity at the timepoint indexed by s_i . At these timepoints, indexed by s_i , the latent dynamics are restarted from a new latent initial condition $z0_i$, which is obtained from the Latent ODE encoder network acting on segment data points $x_{s_i:e_i}$. The latent dynamics are solved using the same latent Neural ODE parameterized by f_Φ . We provide a schematic visualizing LatSegODE hybrid trajectory reconstruction in Figure 2. The example hybrid trajectory is represented by a sequence of base model Latent ODE reconstructions, each starting from a new initial latent state which can discontinuously jump from the previous dynamic. An arbitrary number of restarts can be detected at test time.

To finish the problem formulation, we define \mathcal{I} as the unknown ground truth set of segment boundaries and latent initial states, such that each hybrid trajectory is associated with set:

$$\mathcal{I} = (s_i, e_i, z0_i) \quad 0 \leq i \leq C \quad (9)$$

Where $Z = z_{0:C}$, the joint log probability of an observed

hybrid trajectory can be represented as:

$$\log p(X, Z | s_{1:C}, e_{1:C}) = \sum_{i=0}^C \log p(x_{s_i:e_i}, z0_i) \quad (10)$$

This formulation assumes independence between observations in separate segments, such that $x_{s_i:e_i} \perp (X \setminus x_{s_i:e_i})$. While this assumption can be limiting in trajectories with long term dependencies, it also allows for increased reconstruction performance in the absence of inter-segment dependency. In these situations, given a trajectory with two dynamical modes, allowing latent dynamics to completely restart at the time of modal change allows for a better representation. In comparison, methods which cannot account for shifts in latent dynamics will be forced to adopt an averaged representation between the two dynamical modes. This intuition is later demonstrated in the experimental section.

We note that the LatSegODE does not represent the location of changepoints using a random process. Since event detection is non-probabilistic, the method is not suitable for hybrid trajectories which self-excite or otherwise change dynamical mode past the observed trajectory.

3.2. Optimal Segmentation

Given this formulation of hybrid trajectories, the key challenge is finding the unknown set \mathcal{I} which maximizes the joint probability of an observed hybrid trajectory. We propose application of optimized search algorithms from the field of offline changepoint detection (CPD) to recover locations of jump discontinuity and switches in dynamical mode, and consequently \mathcal{I} . Through complexity penalization, these search algorithms can automatically determine the optimal number and location of segments without prior specification.

Offline CPD methods attempt to discover changepoints which define segment boundaries. A combination of segments which reconstruct a trajectory is referred to as a segmentation. We allow each observed timepoint to be a potential changepoint. Thus, the space of all possible segmentations is formed by all combinations of an arbitrary

number of changepoints. At either extremes, placing no changepoints or a changepoint at each time of observation are both valid segmentations. The space of all possible segmentations grows exponentially (2^N) with the number of observations (N).

The optimal partitioning method (Jackson et al., 2005) uses dynamical programming to search through this large space of solutions. Where \mathcal{C} is a cost function, m is the number of changepoints, and τ is a set of changepoints such that $\tau_0 = 0, \tau_{m+1} = n$, it minimizes

$$\sum_{i=1}^{m+1} \mathcal{C}(x_{\tau_{i-1}+1:\tau_i}) + \beta \quad (11)$$

with respect to τ using dynamic programming. Of all possible segmentations up to data index t , we let $F(t)$ represent the one which results in the minimal cost. This result is memoized. For a new data index $s > t$, we can extend the optimal solution via recursion

$$F(s) = \min_t F(t) + \mathcal{C}(x_{(t+1):s}) + \beta \quad (12)$$

Thus, we begin by solving for $F(1)$, and incrementally extend the solution until $F(N)$, at which point the optimal segmentation is returned. The memoization of previous optimal sub-solutions allows a quadratic runtime with respect to number of observations. The full algorithm is provided in Appendix A. The β term penalizes over-segmentation, and typically scales with the number of parameters introduced by each additional changepoint. When a maximum likelihood cost function is used without a β penalty, optimal partitioning degenerates by placing a changepoint at each possible index. The presence of β enforces a trade-off between accuracy and model complexity. With an appropriate β , this formulation also conveniently recovers the segmentation with the minimized Bayesian Information Criterion (BIC) (Schwarz et al., 1978) through minimization of equation (11).

Choice of β is a key challenge in using CPD methods with deep architectures. It is not always clear how many effective parameters are introduced by each additional segment, though this number is upper bounded by the dimensionality of the latent initial state. Additionally, the theoretical assumptions required by the BIC are violated by neural network architectures (Watanabe, 2013). The LatSegODE circumvents these challenges by using the marginal likelihood under the Latent ODE as the score function for each segment.

We compute a Monte Carlo estimate of the marginal likelihood by importance sampling using a variational approxi-

mation to the posterior over the initial state:

$$\log p(x_{s:e}) = \log \int p(x_{s:e}|z_0) p(z_0) dz_0 \quad (13)$$

$$= \mathbb{E}_{z_0 \sim q(z_0|x_{s:e})} \left[p(x_{s:e}|z_0) \frac{p(z_0)}{q(z_0|x_{s:e})} \right] \quad (14)$$

$$= \frac{1}{M} \sum_{j=1}^M \mathcal{N}(\overline{x_{s:e}}|x_{s:e}, \sigma^2) \frac{\mathcal{N}(z_{0j}|0, 1)}{\mathcal{N}(z_{0j}|\mu_{z_0}, \sigma_{z_0}^2)} \quad (15)$$

where $\overline{x_{s:e}}$ is the output of the Latent ODE base model, $\mu_{z_0}, \sigma_{z_0}^2$ is obtained by the GRU-ODE encoder, and z_{0j} is sampled as $\mathcal{N}(\mu_{z_0}, \sigma_{z_0}^2)$. The variance σ^2 is fixed, and set to the same value used to compute the ELBO during training. We take M samples for the Monte Carlo estimate.

Because we use the marginal likelihood, the complexity of the recovered segmentation is implicitly regularized by the Bayesian Occam’s Razor (MacKay, 1992). Reflecting this, in our experiments, we show that the penalization term β can be set to 0 without over-segmentation. Thus, we can simply set \mathcal{C} in equation (11) to be the marginal likelihood computed by equation (15), and solve for the set of changepoints τ which maximize the joint probability of the entire trajectory using optimal partitioning (the original objective is a minimization, but this can trivially be switched to maximization).

The quadratic runtime of optimal partitioning can be reduced to between $\mathcal{O}(N)$ and $\mathcal{O}(N^2)$ through the pruned exact linear time (PELT) (Killick et al., 2012) algorithm. Using an identical search algorithm, PELT introduces a pruning condition which allows removal of sub-solutions from consideration. Given the existence of K such that for all changepoint indexes s, t, T such that $t < s < T$:

$$\mathcal{C}(x_{(t+1):s}) + \mathcal{C}(x_{(s+1):T}) + K \leq \mathcal{C}(x_{(t+1):T}) \quad (16)$$

Then if

$$F(t) + \mathcal{C}(x_{(t+1):s}) + K \geq F(s) \quad (17)$$

we are able to discard the changepoint t from future consideration, asymptotically reducing the number of operations required. Due to noise in the estimates of the score function, finding an analytic method to determine K is an area for further research. If K is set too low, sub-optimal solutions are recovered. In practice, this issue is not limiting, as setting K to a sufficiently high value allows for near-optimal solutions at the cost of higher runtime. This trade-off is documented in Appendix B.

The computation of $F(t)$, the optimal segmentation up to length t , and Monte Carlo estimate of the marginal likelihood can all be batch parallelized using GPU computation. An implementation is available at: <https://github.com/IanShi1996/LatentSegmentedODE>.

3.3. When can I use this method?

The LatSegODE requires a Latent ODE base model trained on a family of SDFs. We propose two scenarios where SDFs may be available. First, the LatSegODE is applicable when a training set of hybrid trajectories with labelled change-points exists. In this case, given a training set of N hybrid trajectories $X = (x^{(i)}, t^{(i)})_{i=1}^N$ each with C labelled SDF boundaries $(s_k, e_k)_{k=0}^C$, we treat each $x_{s_j:e_j}^{(i)}$ as an independent training trajectory, and train on the union of all SDFs. The LatSegODE can also be applied when physical simulation is available. In these scenarios, the base model can be trained on trajectories which are simulated in the range of dynamical modes which we expect in hybrid trajectories at test time. These two use cases are illustrated in the first two experiments.

4. Related work

Switching Dynamical Systems: Hybrid trajectories have previously been modelled as Switching Linear Dynamical Systems (SLDS). We provide a non-exhaustive summary of these methods. Typically, trajectories are represented by a Bayesian network containing a sequence of latent variables, from which observations are emitted. Latent variables are updated linearly, while a higher order of latent variable represents the current dynamical mode. Structured VAEs (Johnson et al., 2016) introduce a discrete latent variable to control dynamical mode, and use a VAE observation model. GPHSMMs (Nakamura et al., 2017) uses a Gaussian Process observation model within a hidden semi-Markov model. Kalman VAEs integrate a Kalman Filter with a VAE observation model (Fraccaro et al., 2017). Models in this class are generally trained via an inference procedure (Dong et al., 2020), while several are fully differentiable (Kipf et al., 2019). These methods are unsupervised, requiring no training data with labelled changepoint locations.

In contrast, the LatSegODE requires a base model to be trained on SDFs. It does not model dependency between segments unlike methods such as rSLDS (Linderman et al., 2017). At evaluation time, the LatSegODE operates without specification of the number of segments or dynamical modes. This is an advantage compared to previously discussed works, where performance is sensitive to these hyperparameters (Dong et al., 2020).

The Neural Event ODE (Chen et al., 2020) is closely related to the LatSegODE. It represents observed dynamics using a Neural ODE and trains a neural network to detect the positions and update values of a switching dynamical system. The Neural Event ODE can be trained in an unsupervised fashion, without prior knowledge of change point locations in training data. When extrapolating past observed data, it is able to introduce additional change points, which

the LatSegODE cannot model. However, the Neural Event ODE inherits the same limitations as the Neural ODE: it cannot model a data set which cannot be described by a single ODE function in data space. So, for example, two different dynamics cannot start from the same observed point. This issue is elaborated in Appendix C. The LatSegODE circumvents these limitation by modelling the data using an ODE in latent space.

Offline Changepoint Detection: The LatSegODE closely relates to offline CPD, and we refer to Truong et al. (2020) for an in-depth review. The LatSegODE leverages search algorithms from offline CPD, but represents the behavior within segments using a complex generative model, as opposed to a simple statistical cost function. The use of the Latent ODE allows for higher representational power and extrapolation/interpolation within segments. However, training data is required to fit the base model and, as such, its total runtime is significantly higher. Other methods have incorporated deep architectures with CPD search methods (Lee et al., 2018), but use a sliding window search with predefined window size, and use a feature distance metric to determine boundaries as opposed to the marginal likelihood used by LatSegODE.

Miscellaneous: A distantly related class of methods classify individual observations into class labels, which can be seen as segmentation (Supratak et al., 2017). These approaches are distinct as they do not explicitly model dynamics, and require a fixed segment size and trajectory length, a limitation which the LatSegODE does not have. The LatSegODE does not treat positions of jump discontinuity or switching dynamical mode as a random variable, unlike methods that model these jumps as a random process (Mei & Eisner, 2017; Jia & Benson, 2019).

5. Experiments

Here we investigate the LatSegODE’s ability to simultaneously perform accurate reconstruction and segmentation on synthetic and semi-synthetic data sets.

When training the base model, we mask observations from the last 20% of the timepoints and 25% of internal timepoints, this 25% is shared across all training and test examples. When evaluating the model on the test set, we use the 55% of unmasked timepoints to infer the initial states and perform segmentation, and then attempt to reconstruct the observations at the masked timepoints. We report the mean squared error (MSE) between ground truth and predicted observations on test trajectories. We benchmark against auto-regressive and vanilla Latent ODE baselines for reconstructive tasks. We augment the input data for the vanilla Latent ODE with a binary-valued time series denoting changepoint positions. This ensures it has access

to the same information as the LatSegODE. We report performance on an extrapolation region which assumes the last observed dynamical mode continues. We attempted to benchmark against Neural ODEs and Neural Event ODEs, but found that their training did not converge on any of our benchmarks (see Appendix C).

We benchmark the segmentation performance of the LatSegODE against classic CPD algorithms using Gaussian kernelized mean change (Arlot et al., 2019), auto regressive (Bai et al., 2000), and Gaussian Process (Lavielle & Teyssiere, 2006) cost functions. These are denoted RPT-RBF, RPT-AR, and RPT-NORM respectively.

Segmentation performance is measured using the Rand Index (Rand, 1971), the Hausdorff metric (Rockafellar & Wets, 2009), and the F1 score. The Rand Index measures the overlap between the predicted segmentation and the ground truth segmentation. Given data points $x_{1:N}$, a membership matrix A is defined such that $A_{ij} = 1$ if x_i and x_j are in the same segment. Otherwise, $A_{ij} = 0$. Membership matrices are generated for the ground truth segmentation (A) and the predicted segmentation (\tilde{A}). Using these two matrices, the Rand Index is calculated as:

$$\frac{\sum_{i < j} \mathbb{1}[A == \tilde{A}]}{N(N-1)/2} \quad (18)$$

The Hausdorff metric is a measure of the maximal error between the predicted segmentation and the ground truth segmentation. Given a set of ground truth changepoints \mathcal{T} and predicted changepoints \mathcal{P} , the Hausdorff metric is computed as:

$$\max \left(\max_{\tau \in \mathcal{T}} \min_{\rho \in \mathcal{P}} |\tau - \rho|, \max_{\rho \in \mathcal{P}} \min_{\tau \in \mathcal{T}} |\tau - \rho| \right) \quad (19)$$

We use the `ruptures` library (Truong et al., 2020) implementation of these baseline methods and metrics.

We found that the segmentation baselines performed extremely poorly when using penalized detection of changepoints. In response, we simplified the problem for them by providing the correct number of changepoints, so that they only needed to choose the correct locations. In contrast, we did not provide LatSegODE with the number of changepoints, thus the evaluation was biased in favor of the baselines. Also, we excluded trajectories with zero changepoints from this benchmark because they are trivially correct. Irregular locations of data observation is handled by applying linear interpolation prior to segmentation. An extended description of baselines, metrics, and experimental set up is provided in Appendix D.

5.1. Sine Wave Hybrid Trajectories

We evaluate the LatSegODE on a benchmark data set of 1D sine wave hybrid trajectories. Here, we assume access to

trajectories with labelled changepoint positions, one of the situations where the LatSegODE can be realistically applied. We generate 7500 hybrid trajectories each containing up to two changepoints. Between each changepoint, segment trajectories are sine waves generated under random parameters. We hold out 300 validation trajectories, 150 test trajectories, and train the LatSegODE base model on the SDFs contained in the remaining trajectories. Data parameters, model architecture and hyper-parameters are reported in Appendix E. In Figure 3, we provide a visual comparison of the LatSegODE against baselines on an example test set trajectory. The LatSegODE outperforms baselines in both reconstruct-

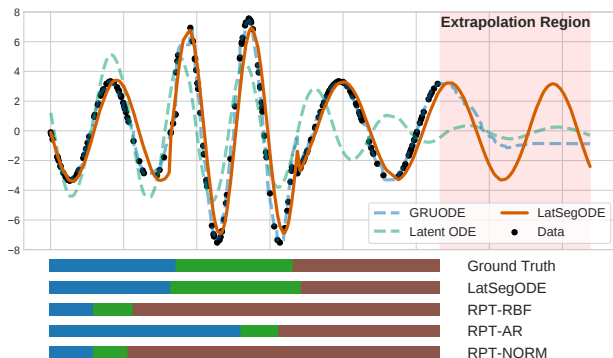


Figure 3. Comparison against baselines in a sample 1D Sine Wave hybrid trajectory. Top: Reconstructed trajectories are shown. Data in the extrapolation region is held out from all models during training. Bottom: Segmentation results are shown. Each distinctly colored region represents a segment.

tion and segmentation tasks. The presence of discontinuities prevent vanilla Latent ODEs from learning accurate representations. Although Latent ODEs can represent the initial SDFs, they lack the ability to represent switches to the dynamical mode. As time progresses, the Latent ODE reconstruction collapses near zero, a local minima which minimizes error given its reconstructive limitations. In contrast, because the LatSegODE can restart latent dynamics, it can represent trajectories with jump discontinuities. The LatSegODE provides an accurate reconstruction, and we see the periodic solution is cleanly captured in the extrapolation region. The GRU-ODE method can fit observed data well, but yields poor interpolations and extrapolations. The LatSegODE recovers the segmentation closest to the ground truth segmentation. The trends observed in this example trajectory are reflected in the overall test results, where the LatSegODE outperforms all baselines. These results are reported in Appendix F. We found that inclusion of the binary-valued changepoint location time series did not result in significant improvement, and we omit this feature from further experiments. We report the effects of the training set size and the number of samples per training trajectory on LatSegODE performance in Appendix K.

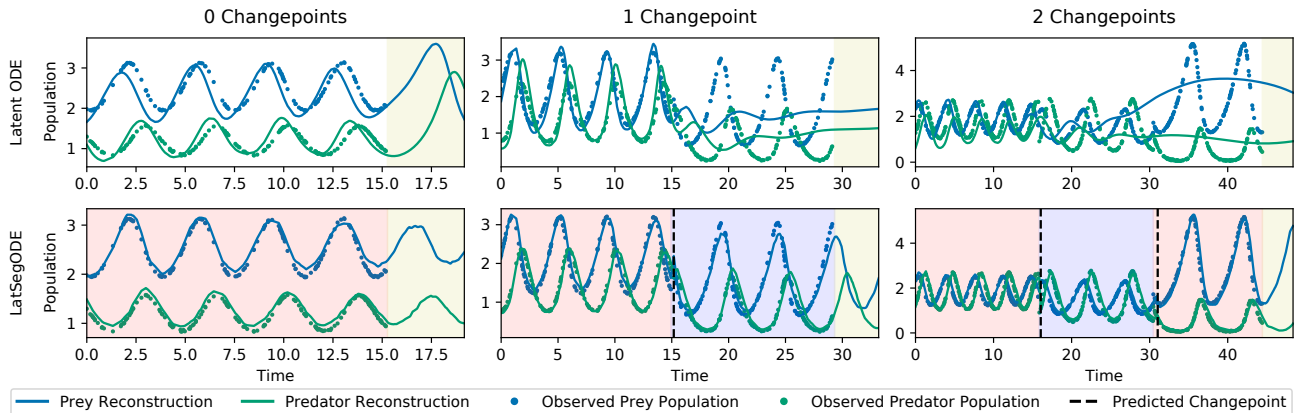


Figure 4. Comparison of reconstructions of Lotka-Volterra hybrid trajectories. Top row contains baseline reconstruction by Latent ODE. Bottom row shows reconstruction by LatSegODE. Sample hybrid trajectories contain the same number of ground truth changepoints in each column. Ground truth segments are shown as a contiguous background color block. Yellow background indicates extrapolation region which assumes that the last observed dynamical mode continues. Visualization inspired by ruptures package (Truong et al., 2020).

5.2. Lotka-Volterra Hybrid Trajectories

Next, we evaluate the LatSegODE on hybrid trajectories whose SDFs are the Lotka-Volterra dynamics described in equation (1). We simulate 34000/600/150 hybrid trajectories for the training/validation/test set. Lotka-Volterra dynamics are generated by randomly sampling coefficients $(\alpha, \beta, \delta, \gamma)$ from ranges $[(0.5, 1.5), (0.5, 1.5), (1.5, 2.5), (0.5, 1.5)]$ respectively. Each trajectory contains up to two changepoints, and at each changepoint we restart dynamics from new initial values sampled from $[(0.5, 1.5), (1.5, 2.5)]$. We resample the coefficient vector at changepoints, so the trajectories feature both jump discontinuity and switch of dynamical mode. We train the LatSegODE base model on the SDFs in the generated training trajectories. The vanilla Latent ODE baseline is trained on full hybrid trajectories, while other baselines were separately trained on both full trajectories and SDFs, with the best performing result reported. The data generation procedure, and model architectures/training is documented in Appendix G.

Results are reported in Table 1, where metrics are averages over 150 test trajectories. The LatSegODE outperforms baselines in both segmentation and reconstruction. An expanded evaluation with additional metrics and experiments is provided in Appendix H.

In Figure 4, we show sample trajectory reconstructions from the LatSegODE versus the vanilla Latent ODE baseline. All vanilla Latent ODE reconstructions over-fit to the change-point locations observed in training data. It is difficult for vanilla Latent ODEs to generalize on permutations of the piece-wise hybrid training trajectories, because they need to encode all sequence information into a single latent initial state. When a permutation in the sequence of SDFs is encountered, the non-robust latent representation predicts ar-

Table 1. Results on Lotka-Volterra hybrid trajectories. Metrics generated using 150 test trajectories. Best result is bolded.

METHOD	TEST MSE	RAND INDEX	HAUSDORFF METRIC
LATSEGOE	0.068	0.9464	47.67
GRU Δt	0.1718	-	-
GRU-ODE	0.2747	-	-
LATENT ODE	0.6155	-	-
RPT-RBF	-	0.7956	84.7
RPT-AR	-	0.6994	164.65
RPT-NORM	-	0.7693	105.92

bitrary dynamical shifts. The vanilla Latent ODE performs badly even in the zero change point reconstruction in Figure 4 where one might expect it to do well, likely because it is anticipating potential change points. In contrast, the structured nature of the LatSegODE bypasses this need to learn a complex latent representation. Segmenting trajectories into SDFs allows for a complex hybrid trajectories to be represented by a sequence of simpler dynamics, yielding the accurate reconstructions shown.

While the base model Latent ODE can powerfully represent SDFs, the LatSegODE method also inherits limitations of the architecture. We visualize two common failure modes in Figure 6. We observed that learning limit cycles was challenging for the base model Latent ODE. In the top trajectory, imprecise base model representations cause deviation from the true periodic solution as time progresses. Eventually, enough error accumulates such that the accuracy gain from introducing a new segment overcomes the complexity cost of this action, resulting in over-segmentation. In the bottom trajectory, a failure mode is caused by the inability for

Segmenting Hybrid Trajectories using Latent ODEs




Figure 5. Example reconstruction on a long hybrid trajectory synthetically generated from UCI Character Trajectory data set. While most characters are accurately detected, an erroneous change point is introduced at $t \approx 6$. As segments are independent under PELT, future segments are not affected by this error, and reconstruction quality recovers after the introduction of the change point at $t \approx 9.5$.




Figure 6. Example failure modes encountered in Lotka Volterra modelling. See Figure 4 for legend.

the base model to generalize. Over-segmentation occurs if test trajectories contain SDFs which start outside of initial values found in training data, such as at the second true changepoint. The base model cannot generalize well to unseen dynamical modes or initial values, so changepoints are erroneously introduced to improve fit. In Appendix I, we report data augmentation tricks which slightly improve generalization, remedying these issues.

This experiment also shows how the LatSegODE can be used in conjunction with physical simulators in a paradigm similar to simulator based inference (Cranmer et al., 2020). We train a MLP to map latent initial states from a trained base model to the labelled Lotka-Volterra coefficients of training SDFs. On test trajectories where the correct number of changepoints were predicted, we could recover the dynamical coefficients with a MSE of 0.08 ± 0.01 . In contexts such as Wright-Fisher population dynamics (Fisher, 1923; Wright, 1931), where forward simulation is available but cannot be expressed in closed form, the LatSegODE could be applied to solve inverse parameter estimation problems.

5.3. UCI Character Trajectories

Finally, we apply the LatSegODE to the UCI Character Trajectory data set (Dua & Graff, 2017). This data set contains

2858 pen tip trajectories collected while writing letters of the alphabet. The trajectories are three dimensional, corresponding to x/y coordinates and pen pressure while writing one character. The data set is pre-processed by normalization and smoothing. Trajectories are regularly sampled with a maximum of 205 observations. We sanitized the data set by removing sections at the beginning and end of trajectories where no movement occurs. We use 5% of the data for validation, and hold out 5% for testing. The LatSegODE base model is trained on the remaining data, using each character trajectory as a SDF. Model architecture and hyper-parameters are reported in Appendix J.

We synthetically construct hybrid test trajectories by composing character trajectories. We randomly sampled a base character trajectory from the test set, then append up to two further randomly sampled character trajectories. To increase task difficulty, we add independent Gaussian noise with standard deviation of 0.2. We also sub-sample the test trajectories to reduce number of observations the Latent ODE base model is able to condition upon. Using this method, we generate 75 synthetic test hybrid trajectories, each containing zero to two changepoints. We report LatSegODE’s segmentation performance on this synthetic test set in Table 2.

Table 2. Segmentation results on UCI Character Trajectories.

METHOD	RAND INDEX	HAUSDORFF METRIC	F1 SCORE
LATSEGODE	0.9732	4.493	0.977
RPT-RBF	0.7956	84.7	0.656
RPT-AR	0.6994	164.65	0.738
RPT-NORM	0.7693	105.92	0.611

In Figure 5, we provide an example reconstruction of a hybrid trajectory constructed by composing six character trajectories sampled from the test set. In both this figure and Table 2, the LatSegODE performs well in reconstructing long sequences of realistic data with noise, and accurately

detects position of change in dynamical mode.

6. Scope and Limitations

Data Labelling: The LatSegODE requires SDF training data, typically obtained by splitting hybrid trajectories using labelled changepoints. This can be hard to obtain, so ideally, LatSegODE could be extended so it could be trained directly on hybrid trajectories. One approach would be marginalizing over changepoints during training using an inference procedure or a iterated-conditional-modes-like procedure that iterates between estimating an optimal segmentation given the current base model, and updating the base model given the segmentation.

Dependency on Dynamical Models: The LatSegODE relies on a Latent ODE base model to capture SDF behavior. Thus, it inherits many limitations of Latent ODEs, but any future advancements in the architecture and training of Latent ODEs can be directly integrated. While we chose to use Latent ODEs due its powerful representational ability, it could be replaced with any model for which marginal likelihood can be computed. Bayesian approaches to Neural ODEs such as the ODE2VAE (Yildiz et al., 2019) and Neural ODE Process (Norcliffe et al., 2021), as well as the Latent SDE (Li et al., 2020) method, could replace the Latent ODE base model with modifications. Thus, our framework can be used a paradigm for an expanded family of methods which combine PELT and dynamical models.

Runtime: The runtime of the LatSegODE can be improved. The current implementation naively computes the ODE solution for the union of batch timepoints. Chen et al. (2020) provide a change of variables method to solve ODEs with irregular timepoints in parallel. This can reduce the memory bottleneck of the current approach, allowing additional parallelism to decrease evaluation runtime. The LatSegODE can be integrated with recent methods to regularize ODE dynamics (Kelly et al., 2020), (Finlay et al., 2020), which decrease evaluation runtime.

7. Conclusion

Here, we present the LatSegODE which leverages Latent ODEs to represent hybrid trajectories. Using a Latent ODE base model trained on SDFs and the PELT changepoint detection algorithm, we identify positions of jump discontinuity and switching dynamical mode, and restart latent dynamics from new initial states at these points. We provide a novel integration of Latent ODEs and CPD methods that uses the marginal likelihood of segments as a scoring function. We find that this Bayesian Occam’s Razor effect prevents over-segmentation. We compared LatSegODE to baselines on synthetic and semi-synthetic benchmarks. Through qualitative analysis of example reconstructions,

we highlight LatSegODE’s ability to represent hybrid trajectories, and demonstrate common failure modes. The LatSegODE outperforms all baselines in both reconstruction and segmentation, supporting it as a novel approach to modelling hybrid trajectories governed by hybrid systems.

Acknowledgements

We thank Tianxing Li and David Duvenaud for their helpful feedback and preliminary reviewing. We also thank Haoran Zhang, Yulia Rubanova, and other members of the Morris Lab for many helpful suggestions. Finally, we thank the ICML reviewers of this paper for their insightful feedback. Resources used in preparing this research were provided, in part, by the Memorial Sloan Kettering Cancer Center, Province of Ontario, the Government of Canada through CIFAR, and companies sponsoring the Vector Institute www.vectorinstitute.ai/partners.

References

- Ackerson, G. and Fu, K. On state estimation in switching environments. *IEEE transactions on automatic control*, 15(1):10–17, 1970.
- Arlot, S., Celisse, A., and Harchaoui, Z. A kernel multiple change-point algorithm via model selection. *Journal of Machine Learning Research*, 20(162):1–56, 2019.
- Bai, J. et al. Vector autoregressive models with structural changes in regression coefficients and in variance-covariance matrices. Technical report, China Economics and Management Academy, Central University of Finance and Economics, 2000.
- Bengio, S., Vinyals, O., Jaitly, N., and Shazeer, N. Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pp. 1171–1179, 2015.
- Brouwer, E. D., Simm, J., Arany, A., and Moreau, Y. GRU-ODE-Bayes: Continuous modeling of sporadically-observed time series. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 7377–7388, 2019.
- Calvo, M., Montijano, J., and Rández, L. On the solution of discontinuous IVPs by adaptive Runge–Kutta codes. *Numerical Algorithms*, 33, 2003. doi: 10.1023/A:1025507920426.
- Calvo, M., Montijano, J., and Rández, L. The numerical

- solution of discontinuous IVPs by Runge-Kutta codes: A review. *SeMA Journal*, 44, 2008.
- Chen, R. T., Amos, B., and Nickel, M. Learning neural event functions for ordinary differential equations. *arXiv preprint arXiv:2011.03902*, 2020.
- Chen, T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 6572–6583, 2018.
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1724–1734, Doha, Qatar, 2014. Association for Computational Linguistics. doi: 10.3115/v1/D14-1179.
- Cranmer, K., Brehmer, J., and Louppe, G. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, 2020.
- Dong, Z., Seybold, B. A., Murphy, K., and Bui, H. H. Collapsed amortized variational inference for switching nonlinear dynamical systems. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 2638–2647. PMLR, 2020.
- Dua, D. and Graff, C. UCI machine learning repository, 2017.
- Dupont, E., Doucet, A., and Teh, Y. W. Augmented neural ODEs. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 3134–3144, 2019.
- Finlay, C., Jacobsen, J., Nurbekyan, L., and Oberman, A. M. How to train your Neural ODE: the world of jacobian and kinetic regularization. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 3154–3164. PMLR, 2020.
- Fisher, R. A. XXI.—on the dominance ratio. *Proceedings of the royal society of Edinburgh*, 42:321–341, 1923.
- Fraccaro, M., Kamronn, S., Paquet, U., and Winther, O. A disentangled recognition and nonlinear dynamics model for unsupervised learning. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 3601–3610, 2017.
- Fu, H., Li, C., Liu, X., Gao, J., Celikyilmaz, A., and Carin, L. Cyclical annealing schedule: A simple approach to mitigating KL vanishing. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 240–250, Minneapolis, Minnesota, 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1021.
- Jackson, B., Scargle, J. D., Barnes, D., Arabhi, S., Alt, A., Gioumoussis, P., Gwin, E., Sangtrakulcharoen, P., Tan, L., and Tun Tao Tsai. An algorithm for optimal partitioning of data on an interval. *IEEE Signal Processing Letters*, 12(2):105–108, 2005. doi: 10.1109/LSP.2001.838216.
- Jia, J. and Benson, A. R. Neural jump stochastic differential equations. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 9843–9854, 2019.
- Johnson, M. J., Duvenaud, D., Wiltchko, A. B., Adams, R. P., and Datta, S. R. Composing graphical models with neural networks for structured representations and fast inference. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 2946–2954, 2016.
- Kelly, J., Bettencourt, J., Johnson, M. J., and Duvenaud, D. Learning differential equations that are easy to solve. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- Kidger, P., Chen, R. T., and Lyons, T. “Hey, that’s not an ODE”: Faster ODE adjoints with 12 lines of code. *arXiv preprint arXiv:2009.09457*, 2020.
- Killick, R., Fearnhead, P., and Eckley, I. A. Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association*, 107(500): 1590–1598, 2012.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- Kingma, D. P. and Welling, M. Auto-encoding variational Bayes. In *2nd International Conference on Learning*

- Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.
- Kipf, T., Li, Y., Dai, H., Zambaldi, V. F., Sanchez-Gonzalez, A., Grefenstette, E., Kohli, P., and Battaglia, P. W. CompILE: Compositional imitation learning and execution. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 3418–3428. PMLR, 2019.
- Lavielle, M. and Teysriere, G. Detection of multiple change-points in multivariate time series. *Lithuanian Mathematical Journal*, 46(3):287–306, 2006.
- Lee, W.-H., Ortiz, J., Ko, B., and Lee, R. Time series segmentation through automatic feature learning. *arXiv preprint arXiv:1801.05394*, 2018.
- Li, X., Wong, T. L., Chen, R. T. Q., and Duvenaud, D. Scalable gradients for stochastic differential equations. In *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020, 26-28 August 2020, Online [Palermo, Sicily, Italy]*, volume 108 of *Proceedings of Machine Learning Research*, pp. 3870–3882. PMLR, 2020.
- Linderman, S., Johnson, M., Miller, A., Adams, R., Blei, D., and Paninski, L. Bayesian learning and inference in recurrent switching linear dynamical systems. In *Artificial Intelligence and Statistics*, pp. 914–922. PMLR, 2017.
- MacKay, D. J. Bayesian interpolation. *Neural computation*, 4(3):415–447, 1992.
- Mei, H. and Eisner, J. The neural Hawkes process: A neurally self-modulating multivariate point process. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 6754–6764, 2017.
- Nakamura, T., Nagai, T., Mochihashi, D., Kobayashi, I., Asoh, H., and Kaneko, M. Segmenting continuous motions with hidden semi-Markov models and Gaussian processes. *Frontiers in neurorobotics*, 11:67, 2017.
- Norcliffe, A., Bodnar, C., Day, B., Moss, J., and Lio, P. Neural ODE processes. *arXiv preprint arXiv:2103.12413*, 2021.
- Rackauckas, C., Ma, Y., Martensen, J., Warner, C., Zubov, K., Supekar, R., Skinner, D., and Ramadhan, A. Universal differential equations for scientific machine learning. *arXiv preprint arXiv:2001.04385*, 2020.
- Rand, W. M. Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical association*, 66(336):846–850, 1971.
- Rockafellar, R. T. and Wets, R. J.-B. *Variational analysis*, volume 317. Springer Science & Business Media, 2009.
- Rubanova, Y., Chen, R. T. Q., and Duvenaud, D. K. Latent ordinary differential equations for irregularly-sampled time series. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Schwarz, G. et al. Estimating the dimension of a model. *The annals of statistics*, 6(2):461–464, 1978.
- Stewart, D. *Dynamics with Inequalities*, chapter 8, pp. 283–306. Society for Industrial and Applied Mathematics, 2011. doi: 10.1137/1.9781611970715.ch8.
- Supratak, A., Dong, H., Wu, C., and Guo, Y. DeepSleepNet: A model for automatic sleep stage scoring based on raw single-channel EEG. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 25(11):1998–2008, 2017.
- Truong, C., Oudre, L., and Vayatis, N. Selective review of offline change point detection methods. *Signal Processing*, 167:107299, 2020.
- Van Der Schaft, A. J. and Schumacher, J. M. *An introduction to hybrid dynamical systems*, volume 251. Springer London, 2000.
- Watanabe, S. A widely applicable Bayesian information criterion. *Journal of Machine Learning Research*, 14 (Mar):867–897, 2013.
- Wright, S. Evolution in Mendelian populations. *Genetics*, 16(2):97, 1931.
- Yildiz, C., Heinonen, M., and Lähdesmäki, H. ODE2VAE: deep generative second order ODEs with Bayesian neural networks. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 13412–13421, 2019.