

# Continual Learning of Multi-modal Dynamics with External Memory

**Abdullah Akgül**

*University of Southern Denmark, Odense, Denmark*

AKGUL@IMADA.SDU.DK

**Gozde Unal**

*Istanbul Technical University, Istanbul, Turkey*

GOZDE.UNAL@ITU.EDU.TR

**Melih Kandemir**

*University of Southern Denmark, Odense, Denmark*

KANDEMIR@IMADA.SDU.DK

**Editors:** A. Abate, M. Cannon, K. Margellos, A. Papachristodoulou

## Abstract

We study the problem of fitting a model to a dynamical environment when new modes of behavior emerge sequentially. The learning model is aware when a new mode appears, but it cannot access the true modes of individual training sequences. The state-of-the-art continual learning approaches cannot handle this setup, because parameter transfer suffers from catastrophic interference and episodic memory design requires the knowledge of the ground-truth modes of sequences. We devise a novel continual learning method that overcomes both limitations by maintaining a *descriptor* of the mode of an encountered sequence in a neural episodic memory. We employ a Dirichlet Process prior on the attention weights of the memory to foster efficient storage of the mode descriptors. Our method performs continual learning by transferring knowledge across tasks by retrieving the descriptors of similar modes of past tasks to the mode of a current sequence and feeding this descriptor into its transition kernel as control input. We observe the continual learning performance of our method to compare favorably to the mainstream parameter transfer approach.

**Keywords:** continual learning, multi-modality, external memory, dynamics modeling.

## 1. Introduction

Continual Learning (CL) aims to develop a versatile model that is capable of solving multiple prediction tasks which are presented to the model one task at a time. The model is then expected to learn the latest task as accurately as possible while preserving its excellence at the previous ones. Performance drop caused by the newly learned task is called *catastrophic forgetting*. CL is essential for developing intelligent agents that can adapt to new environmental conditions not encountered during training. For instance, an autonomous driving controller may improve its policy based on new experiences collected during its customer-side lifetime.

There exists a solid body of work that adopts parameter transfer across tasks as the key element of task memorization (Kirkpatrick et al., 2017; Nguyen et al., 2018; Singh et al., 2019; Zenke et al., 2017). There also appear preliminary studies on building attentive memories to capture tasks (Garnelo et al., 2018; Fraccaro et al., 2018). Principles that yield memory mechanisms optimal for CL are yet to be discovered. There has been prior work that focuses on CL for Recurrent Neural Networks (RNN) (Cossu et al., 2021) but either classification or instance forecasting for time series. Complementarily, we study CL in the context of dynamical system identification. Probabilistic State-Space Models (SSM) are the gold standard methodology for the inference of complex latent dynamics. SSMs are widely applicable to forecasting impactful quantities such as weather, currency exchange, equity prices, and sales trends. SSM research gains significance also in robot learning parallel to the growing interest in model-based reinforcement learning (Hafner et al., 2019, 2020).

Our main contributions are two-fold. First is a novel problem setup where dynamical system modeling tasks emerge sequentially and a probabilistic SSM is expected to learn them cumulatively. We assume each task to follow multi-modal dynamics, where each individual sequence of a task follows one of the possible modes that describe the task. Successful CL in such a setup presupposes maximally efficient encoding of tasks into long-term memory and their accurate retrieval. Our second contribution is a novel CL model tailored for addressing this challenging problem. The state of the art in CL uses either parameter transfer or end-to-end differentiable attentive episodic memory for knowledge transfer across tasks. The multimodality of the sequences

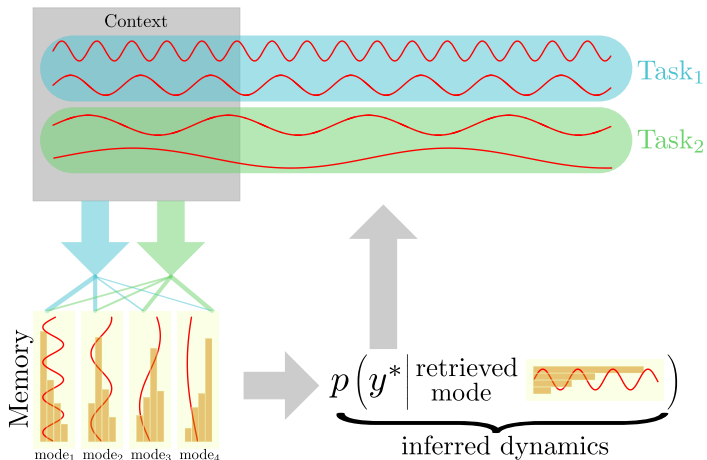


Figure 1: Our research shows that inferring and using mode descriptors from observed sequences significantly enhances CL performance. Storing these descriptors in external neural episodic memory and integrating them into subsequent tasks drives improved learning outcomes.

would undermine the parameter transfer approach due to catastrophic interference (the effect of global parameter updates on unintended regions of the input space) and it would require the ground-truth modality labels for the memory-based approaches to be applicable. Our approach sidesteps these limitations by capturing the characteristics of unknown modes of sequences of the present task into fixed-sized vectors, called *mode descriptors*, and storing these descriptors in an external neural episodic memory addressable via a learnable attention mechanism. Our approach represents multiple task modes by feeding these mode descriptors into the state transition kernel as an additional input. We place a Dirichlet Process (DP) prior on the attention weights of the memory to encourage the explanation of the data with a minimum number of modes. Figure 1 illustrates our model with external memory and the problem with two tasks and four modes. Our resulting Bayesian model can be efficiently trained using a straightforward adaptation of existing variational SSM inference techniques.

We evaluate our model in two time series prediction data sets and three synthetic data sets generated from challenging nonlinear multi-modal dynamical systems. The performance of our method improves consistently over the established parameter transfer approach, verifying the importance of parsimonious use of neural episodic memory for efficient acquisition of knowledge within tasks and effective transfer of knowledge across different tasks.

## 2. Continual Multi-Modal Dynamics Learning

We assume a learning agent that observes a dynamical environment via sequences  $y_{1:T} = \{y_1, \dots, y_T\}$  consisting of time-indexed measurements  $y_t \in \mathcal{Y}$  living in a measurable space  $\mathcal{Y}$ . We denote a snippet of a sequence  $y_{t:t'} = \{y_t, \dots, y_{t'}\}$  for an arbitrary time interval  $[t, t']$ . Modes refer to general

distinguishable properties of a dynamic system such as states with different characteristics of a sine wave (amplitude, frequency) or semantically different dynamics such as different character trajectories. We define the *mode* of a sequence  $y_{1:T}$  as a fixed-sized vector and  $m$  as an element of a  $K$ -dimensional embedding space  $\mathcal{X}$ . A dynamical system can potentially operate within a large number of modes. Making an analogy to the real world, an autonomous vehicle accounts for different environmental characteristics when planning and controlling for different weather conditions, countries, and times of a day. Only a few of the modes are active for a particular time point and each mode instantly activates and deactivates for limited time periods.

We search for a learning algorithm that enables the agent to fit a dynamical environment that has perpetually changing global characteristics. *We cast the corresponding learning problem as CL of multi-modal dynamical systems, where each task is defined as a group of modes that modulate a specific dynamical system.* In formal terms, a task  $\mathcal{T}_i$  is defined as a group of mode descriptors sampled from a mode generating oracle  $P(\mathcal{X})$ , that is

$$\mathcal{T}_i = \{m_r^i | m_r \sim P(\mathcal{X}), r = 1, \dots, R\}, \quad i = 1, \dots, U, \quad (1)$$

where  $R$  is the number of modes per task and  $U$  is the number of tasks. We assume that each sequence of a task follows dynamics modulated by a mode sampled from  $\mathcal{T}_i$

$$\mathcal{D}_{\mathcal{T}_i} = \left\{ y_{1:T}^n \mid m_n \sim P(\mathcal{T}_i), y_{1:T}^n \sim p(y_{1:T} | m_n), n = 1, \dots, N \right\}, \quad (2)$$

where  $N$  is number of instances per task,  $P(\mathcal{T}_i)$  is a probability mass function defined on the modes of task  $\mathcal{T}_i$ , and  $p(y_{1:T} | m_r)$  is the probability measure that describes the true behavior of the environment dynamics under mode  $m_r$  within a time interval  $[1, T]$ . The marginal distribution of a task with respect to its modes is given as  $p(y_{1:T} | \mathcal{T}_i) = \sum_{r=1}^R p(y_{1:T} | m_r) P(m_r | \mathcal{T}_i)$ . The agent observes a task via a data set  $\mathcal{D}_{\mathcal{T}_i}$  that contains only the sequences  $y_{1:T}^n$  *but not their corresponding modes*.

We are interested in learning a model that minimizes the true CL risk function below

$$\mathbb{R}_{\mathbb{L}}^{CL}(h_\theta) = \lim_{U \rightarrow +\infty} \sum_{i=1}^U \mathbb{E}_{\mathcal{T}_i \sim P(\mathcal{X})} \left[ \mathbb{E}_{y_{1:T} \sim p(y_{1:T} | \mathcal{T}_i)} [\mathbb{L}(h_\theta(\cdot | y_{1:C}), y_{C+1:T})] \right], \quad (3)$$

which amounts to the limit of the cumulative risk of individual tasks as they appear one at a time. Above,  $h_\theta : \mathcal{Y}^C \rightarrow \mathcal{Y}^{T-C}$  is a stochastic process that can map an observed sequence of an arbitrary length  $C$  to the subsequent  $T - C$  time steps. We call the first  $C$  observations in the sequence  $y_{1:C}$  as the *context* and define  $\mathbb{L}$  as a sequence-specific loss function defined on the future values of the sequence  $y_{C+1:T}$ . We evaluate the performance of predictions  $\hat{y}$  via two scores: *Normalized Mean Squared Error (NMSE)* as a measure of the prediction accuracy of  $h_\theta$  when used as a Gibbs predictor, and *Negative Log-likelihood (NLL)* as a measure of Bayesian model fit that quantifies the model's own assessment on the uncertainty of its assumptions. We approximate the true CL risk by its empirical counterpart:  $\hat{\mathbb{R}}_{\mathbb{L}}^{CL} = \frac{1}{N} \sum_{i=1}^U \sum_{n=1}^N \mathbb{L}(h_\theta(y_{C+1:T}^n | y_{1:C}^n), y_{C+1:T}^n | \mathcal{T}_i)$  for a finite number  $U$  of tasks presented to  $h_\theta$  one at a time.

### 3. Novel Baseline: Variational Continual Learning for Bayesian State-Space Models

We build our target model on a Bayesian treatment of state-space modeling, which is proven to be effective in learning under high uncertainty and knowledge transfer across tasks, as practiced in the

seminal prior art of CL (Kirkpatrick et al., 2017; Nguyen et al., 2018). We perform approximate inference using variational Bayes due to its multiple successful applications to state-space models (Frigola et al., 2014; Doerr et al., 2018; Jalongo et al., 2019) and its favorable computational properties. *As there does not exist any prior work tailored specifically towards CL for dynamical systems, we curate our own baseline.* We adopt the established practice of setting the posterior of the learned parameters of the previous task as the prior of the next one and determine Variational Continual Learning (VCL) (Nguyen et al., 2018) as the state-of-the-art representative of this approach. Elastic Weight Consolidation (EWC) (Kirkpatrick et al., 2017) also follows the same approach but uses a simpler posterior inference scheme.

**Bayesian State Space Models (BSSM):** are characterized by the data generating process below

$$\theta \sim p(\theta), \quad x_0 \sim p(x_0), \quad x_t|x_{t-1}, \theta \sim p(x_t|x_{t-1}, \theta), \quad y_t|x_t \sim p(y_t|x_t), \quad (4)$$

where  $x_t$  and  $y_t$  correspond to the latent and observed state variables for time step  $t$ , respectively. The system dynamics are modeled by the first-order Markovian transition kernel  $p(x_t|x_{t-1}, \theta)$  parameterized by  $\theta$  that in turn follows a prior distribution  $p(\theta)$ . The latent states are mapped to the observation space via a probabilistic observation model  $p(y_t|x_t)$ . We formulate the initial latent state  $x_0$  as another random variable that follows the prior distribution  $p(x_0)$ .

**Variational Inference** is required to approximate the posterior  $p(x_{0:T}, \theta|y_{1:T})$ , which will be intractable for many choices of distribution families for the data generating process in Eq. 4. Following Yildiz et al. (2019), we choose the variational distribution to be mean-field across the parameters of the dynamics and the latent states

$$q_{\xi, \psi}(x_{0:T}, \theta|y_{1:T}) = q_{\xi}(x_0|y_{1:C}) \prod_{t=1}^T p(x_t|x_{t-1}, \theta) q_{\psi}(\theta), \quad (5)$$

where  $q_{\xi, \psi}$  is an approximation to the true posterior  $p(x_{0:T}, \theta|y_{1:T})$  with variational free parameters  $(\xi, \psi)$ . This formulation has multiple advantages. Firstly, modeling the marginal posterior on the initial latent state  $x_0$  by amortizing on the context observations  $y_{1:C}$  makes the Evidence Lower Bound (ELBO) calculation invariant to the context length. Secondly, adopting the prior transition kernel avoids duplicate learning of environment dynamics with twice as many free parameters and prevents training from instabilities caused by the inconsistencies between prior and posterior dynamics. Applying Jensen’s inequality in a conventional way, the corresponding ELBO will be

$$\begin{aligned} \log p(y_{1:T}) &\geq \mathbb{E}_{p(x_{1:T}|\theta, x_0)q_{\xi}(x_0|y_{1:C})q_{\psi}(\theta)} [\log p(y_{1:T}|x_{1:T})] \\ &\quad - KL(q_{\xi}(x_0|y_{1:C})||p(x_0)) - KL(q_{\psi}(\theta)||p(\theta)), \end{aligned} \quad (6)$$

where  $KL(\cdot||\cdot)$  stands for the Kullback-Leibler (KL) divergence between the two distributions on its arguments, and  $X = \{x_0, \dots, x_T\}$ .

**VCL for BSSMs** can be implemented as follows. Having fitted the ELBO (Eq. 6) on the data set for the first task which has observed sequences  $\mathcal{D}_{\mathcal{T}_1}$ , we attain  $(\xi_1^*, \psi_1^*) = \operatorname{argmax}_{\xi, \psi} \mathcal{L}(\xi, \psi, \mathcal{D}_{\mathcal{T}_1})$ . When the next task arrives with data  $\mathcal{D}_{\mathcal{T}_2}$ , we assign  $p(\theta) \leftarrow q_{\psi_1^*}(\theta)$  and maximize the ELBO again  $(\xi_2^*, \psi_2^*) = \operatorname{argmax}_{\xi, \psi} \mathcal{L}(\xi, \psi, \mathcal{D}_{\mathcal{T}_2})$ . We repeat this process continually for every new task. We refer to this newly curated baseline in the rest of the paper as *VCL-BSSM*. We neglect the coreset extension of VCL since its application to BSSMs is tedious and its advantage is not demonstrated with sufficient significance in static prediction tasks studied in original work.

#### 4. Target Model: The Continual Dynamic Dirichlet Process

The commonplace Bayesian approach to CL transfers knowledge across tasks by assigning the learned posterior on the parameters of the previous task as the prior on the parameters of the current task. This is an effective approach when the subject of transfer is a feed-forward model, such as a classifier in a supervised learning setup (Nguyen et al., 2018) or a policy network in reinforcement learning (Kirkpatrick et al., 2017). *We conjecture that the existing parameter transfer Ansatz would not be sufficient for the transfer of more complex task properties such as modes of dynamical systems.* We address this problem by tailoring a novel CL approach from an original combination of an aged statistical machine learning tool, the DP, with modern neural episodic memory and attention mechanisms.

**Dirichlet Processes** are stochastic processes defined on a countably infinite number of categorical outcomes, every finite subset of which follows a multinomial distribution drawn from a Dirichlet prior (Teh et al., 2006). A DP follows a Griffiths-Engen-McCloskey (GEM) distribution (Pitman et al., 2002)

$$\pi'_r | \alpha_0 \sim \text{Beta}(1, \alpha_0), \quad \pi_r = \pi'_r \prod_{j=1}^{r-1} (1 - \pi'_j), \quad \pi = [\pi_1, \dots, \pi_R], \quad (7)$$

which we denote in short hand as  $\pi \sim \text{GEM}(\alpha_0, R)$ . The GEM distribution can also be viewed as a stick-breaking process (Sethuraman, 1994; Fox et al., 2011) where  $\alpha_0 > 0$  is a scalar hyper-parameter. The data generation process below is called a DP for a base measure  $G_0(\mathcal{X})$  defined on a  $\sigma$ -algebra  $\mathbb{B}$  of  $\mathcal{X}$

$$m_r \sim G_0(\mathcal{X}), \quad \pi \sim \text{GEM}(\alpha_0, R), \quad G|\pi = \sum_{r=1}^R \pi_r \delta_{m_r}, \quad (8)$$

where  $\delta_x(A)$  is Dirac delta measure that takes value 1 if  $x \in A$  and 0 otherwise for any measurable set  $A \in \mathbb{B}$  and  $G$  is a categorical distribution with parameters  $\pi$ .

**Neural Episodic Memory.** We assume that the environment dynamics can be expressed within a measurable latent embedding space  $x_t \in \mathcal{X}$ , and a sequence encoder  $e_\lambda(x_{t:t'})$  parameterized by  $\lambda$  for  $t \leq t'$  that maps a sequence of latent embeddings into a fixed dimensional vector, as well as a neural memory  $M = \{m_1, \dots, m_R\}$  consisting of  $R$  elements  $m_r \in \mathcal{X}$  that live in the same space as latent embeddings  $x_t$ . We can construct a probability measure for  $\mathbb{B}$  from the memory  $M$  by updating

$$m_r \leftarrow (1 - w_r(y_{1:C}, m_r))m_r + w_r(y_{1:C}, m_r)e_\lambda(y_{1:C}) \quad (9)$$

for each sequence where  $w_r(y_{1:C}, m_r) = \frac{e^{\langle m_r, e_\lambda(y_{1:C}) \rangle}}{\sum_{j=1}^R e^{\langle m_j, e_\lambda(y_{1:C}) \rangle}}$  for some similarity function  $\langle \cdot, \cdot \rangle : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^+$ .

This construction imposes a memory attention mechanism, where the encoded mode descriptors attend to the memory elements. Here we make the fair assumption that the modality of a sequence can be identified also from the observation space, while we need to infer the latent representations accurately to model the mode dynamics in detail. We choose an uninformative base measure that assigns equal prior probabilities to memory elements  $G_0(M) = \sum_{k=1}^R \frac{1}{R} \delta_{m_r}$ .

**The full model.** We complement the BSSM in Eq. 4 with an external neural episodic memory  $M$  that is updated for each observed sequence with the rule in Eq. 9. We place a DP prior on the retrieval of mode descriptors  $m_r \in M$  to encourage the model to generate a minimum number of modes. We also feed the retrieved mode descriptor into the transition kernel  $p(x_t|x_{t-1}, m, \theta)$  as control input. The resultant model, which we call as the *Continual Dynamic Dirichlet Process (CDDP)*, follows the generative process

$$\pi \sim GEM(\alpha_0, R), \quad m|\pi \sim \sum_{r=1}^R \pi_r \delta_{m_r}, \quad \theta \sim p(\theta), \quad (10)$$

$$x_0 \sim p(x_0), \quad x_t|x_{t-1}, m, \theta \sim p(x_t|x_{t-1}, m, \theta), \quad y_t|x_t \sim p(y_t|x_t), \quad (11)$$

where memory capacity  $R$  is set to a bigger number than the expected upper limit of the mode count.

**Inference.** Since  $p(x_{0:T}, \theta|y_{1:T})$  is intractable, we approximate it by variational inference. We inherit the advantages of the BSSM inference scheme by choosing the variational distribution as

$$q_{\xi, \psi}(x_{0:T}, \theta, \pi|y_{1:T}) = q_{\xi}(x_0|y_{1:C}, \pi) \prod_{t=1}^T \sum_{r=1}^R q_{\xi}(\pi = r|y_{1:C}) p(x_t|x_{t-1}, m_r, \theta) q_{\psi}(\theta), \quad (12)$$

where

$$q_{\psi}(\theta) = \mathcal{N}(\theta|\mu, \Sigma), \quad q_{\xi}(\pi|y_{1:C}) = \text{Cat}(w_1(m_1, y_{1:C}), \dots, w_R(m_R, y_{1:C})), \quad (13)$$

$$q_{\xi}(x_0|y_{1:C}, \pi) = \sum_{r=1}^R \pi_r \mathcal{N}(x_0|v_1(m_r, e_{\lambda}(y_{1:C})), v_2(m_r, e_{\lambda}(y_{1:C}))) \quad (14)$$

with  $x_0 \sim \mathcal{N}(0, I)$ . In the expressions above,  $v_1$  and  $v_2$  refer to dense layers. The corresponding ELBO is then calculated with

$$\begin{aligned} \log p(y_{1:T}) &\geq \sum_{r=1}^R q_{\xi}(\pi = r|y_{1:C}) \mathbb{E}_{p(x_{1:T}|\theta, x_0, m_r) q_{\xi}(x_0|y_{1:C}) q_{\psi}(\theta)} [\log p(y_{1:T}|x_{1:T})] \\ &\quad - KL(q_{\xi}(x_0|y_{1:C})||p(x_0)) - KL(q_{\psi}(\theta)||p(\theta)) - \sum_{r=1}^R w_r(m_r, y_{1:C}) \log(w_r(m_r, y_{1:C})/\pi_r). \end{aligned} \quad (15)$$

**Prediction.** Having trained the model on the latest task  $\mathcal{T}_i$ , its posterior predictive distribution for new sequence  $Y^*$  and corresponding latent embeddings  $X = \{x_0, \dots, x_T\}$  reads

$$\begin{aligned} p(Y_{C+1:T}^*|\mathcal{D}_{\mathcal{T}_i}, Y_{1:C}^*) &= \\ &\sum_{r=1}^R q(\pi = r|Y_{1:C}^*) \int_{X, \theta} q_{\xi}(x_0|Y_{1:C}^*, \pi = r) q_{\psi}(\theta) \left[ \prod_{t=C+1}^T p(y_t^*|x_t) p(x_t|x_{t-1}, \theta, m_r) \right]. \end{aligned} \quad (16)$$

## 5. Related Work

**State Space Models.** There exists a considerable body of work on Bayesian versions of SSMs that employ Gaussian Process as transition kernels (Ialongo et al., 2019) and perform variational

inference. Another vein of work named Recurrent SSMs (Hafner et al., 2019, 2020) models the transition dynamics as RNN that map a state to the next time step deterministically while admitting a random state variable from the previous time step as input and feeding its output to the distribution of this variable at the subsequent time step.

**Attention and Memory in Neural Nets.** Neural Turing Machines (Graves et al., 2014) are the first examples of attention-based neural episodic memory use with external updates. Attentive Neural Process (Kim et al., 2019) employs an attention network to build a neural stochastic process that is consistent over observed predictions. Evidential Turing Process (Kandemir et al., 2022) maintains an external memory that learns to feed a Dirichlet prior on class distribution with informative concentration parameters inferred during minibatch training updates.

**Continual Learning** is an instance of meta-learning (Finn et al., 2017) where new tasks are introduced one at a time and a base model is expected to learn the newest task without forgetting the previous ones. Early approaches to CL such as EWC (Kirkpatrick et al., 2017) transfer knowledge by the transfer of either deterministic parameters or their inferred distribution using Fisher information. VCL (Nguyen et al., 2018) improves on EWC with a more comprehensive inference scheme. Generalized VCL (Loo et al., 2021) maximizes the same ELBO as VCL but uses  $\beta$ -VAE to prevent the training instability caused by the dominance of the KL-divergence term. We do not use  $\beta$ -VAE since our setup does not have this problem.

Knoblauch et al. (2020) explains the superior performance of memorization-based CL algorithms based on experience replay, core set, and episodic memory (Shin et al., 2017; Lopez-Paz and Ranzato, 2017; Lüders et al., 2016) over regularization-based algorithms (Kirkpatrick et al., 2017). For the first time, our CDDP studies probabilistic multi-modal dynamics in a CL setting by knowledge transfer via learned mode descriptors maintained in external memory.

**Continual Learning with Episodic Memory.** Rios and Itti (2019) builds a memory by maximum sample diversity in order to reduce catastrophic forgetting by remembering samples of previous tasks. In Guo et al. (2020), a memory keeps sample random examples from previous tasks and these samples are used in the training of the new tasks. There are other works Lopez-Paz and Ranzato (2017); Chaudhry et al. (2019) that use memory for CL however, all of them build the memory with the help of data points provided with the ground truth but in our setup mode labels are not provided. Hence our work is not comparable to theirs. As neither of them studies dynamics modeling, their adaption to our setup is not straightforward.

## 6. Experiments

We evaluate the performance of CDDP rigorously on five challenging applications of CL to time series forecasting. We provide a supplementary that includes a comprehensive PyTorch implementation of the entire experimental process, encompassing the studied models, the data generation process of the used synthetic environments, as well as the performance evaluation procedures.

**Baseline Selection.** As we are the first to investigate CL for BSSM inference, there is no prior work we can take as a baseline without significant adaptation. We determine knowledge transfer across tasks by setting the parameter posterior of the previous task as the prior of the next as the state-of-the-art approach in CL. We adapt VCL, the most established variant of this approach, to BSSMs in §3 as a baseline that can maximally challenge our CDDP. Baseline methods must be



generative due to the nature of the dynamic modeling problem. Furthermore, as stated in the problem setup (§2), the correct mode labels are not provided therefore baseline methods need to be unsupervised generative. No method meets these requirements but we adapt the VCL method.

## Data Sets

**Synthetic Data Sets.** We evaluate our CDDP on three non-linear dynamical systems, where forecasting is challenging, while ground-truth task similarity is controllable. Too much task similarity would make CL unnecessary, while too much task difference would make it infeasible. When tasks share a reasonable degree of similarity, a successful CL algorithm is expected to capture, encode, and memorize these similarities and discard their differences.

We generate modes that share similar dynamical properties, as their time evolution follows the same set of differential equations. However, modes differ from each other in the choice of the free parameters that govern the dynamical system. **i) Sine Waves** data set consists of signals grouped into modes described by different magnitudes and frequencies. Sine waves are created from the function  $A\sin(2\pi ft)$  where  $A$  is the magnitude,  $f$  is the frequency, and  $t$  is time. We generate different modes by changing the choices for magnitude and frequency levels. For **ii) Lotka-Volterra** and **iii) Lorenz Attractor**, we follow the prior works [Haußmann et al. \(2021\)](#); [Satorras et al. \(2019\)](#) and generate different modes by changing the free parameters of the systems. We perturb all three dynamical systems with Gaussian white noise for training splits.

**Real-World Data Sets.** We evaluate our CDDP in two real-world time series classification data sets from [Dua and Graff \(2017\)](#). **i) Libras Movement Data Set** is the official Brazilian sign language and the data set consists of sequences of  $(x, y)$  coordinates of hand movements from 15 different classes. **ii) Character Trajectories** data set consists of sequences of velocities of  $(x, y)$  coordinates and pen force collected from hand-writings of 20 English alphabet characters, which can be written by a single pen-down segment. Treating each class as a mode, we create multimodal time series forecasting tasks from these two data sets that satisfy the learning setup in Eq. 3.

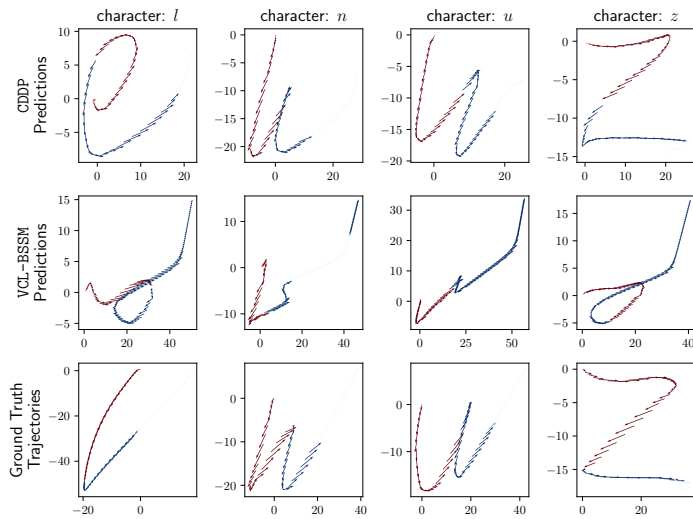


Figure 2: After learning is finished, CDDP and VCL-BSSM predictions on Character Trajectories are displayed alongside ground-truth trajectories. Context and predictions are plotted in red and blue respectively. CDDP exhibits superior memory of the initial task and precise adaptation to the last task due to its external episodic memory.



Table 1: Model performance on five datasets: Mean  $\pm$  standard error over 10 repetitions. Area Under Curve (AUC) represents the averaged score across tasks.

Data Set	#modes	Model	Score	AUC	1 Task	2 Tasks	3 Tasks	4 Tasks	5 Tasks
Sine Waves	15	VCL-BSSM	NMSE	1.00 $\pm$ 0.04	<b>0.13 <math>\pm</math> 0.02</b>	0.97 $\pm$ 0.14	1.22 $\pm$ 0.15	1.27 $\pm$ 0.09	1.39 $\pm$ 0.08
			NLL	3.57 $\pm$ 0.09	<b>1.76 <math>\pm</math> 0.22</b>	3.75 $\pm$ 0.24	3.84 $\pm$ 0.13	4.20 $\pm$ 0.15	4.29 $\pm$ 0.10
		CDDP	NMSE	<b>0.91 <math>\pm</math> 0.03</b>	0.16 $\pm$ 0.02	<b>0.88 <math>\pm</math> 0.11</b>	<b>1.07 <math>\pm</math> 0.15</b>	<b>1.12 <math>\pm</math> 0.14</b>	<b>1.31 <math>\pm</math> 0.11</b>
			NLL	<b>3.50 <math>\pm</math> 0.09</b>	2.09 $\pm$ 0.27	<b>3.43 <math>\pm</math> 0.22</b>	<b>3.77 <math>\pm</math> 0.17</b>	<b>4.05 <math>\pm</math> 0.14</b>	<b>4.15 <math>\pm</math> 0.10</b>
Lotka-Volterra	8	VCL-BSSM	NMSE	<b>0.58 <math>\pm</math> 0.04</b>	0.17 $\pm$ 0.03	<b>0.71 <math>\pm</math> 0.08</b>	<b>0.92 <math>\pm</math> 0.15</b>	<b>0.53 <math>\pm</math> 0.03</b>	
			NLL	1.50 $\pm$ 0.05	0.56 $\pm$ 0.20	1.97 $\pm$ 0.28	<b>2.02 <math>\pm</math> 0.19</b>	<b>1.46 <math>\pm</math> 0.12</b>	
		CDDP	NMSE	0.60 $\pm$ 0.06	<b>0.13 <math>\pm</math> 0.03</b>	0.72 $\pm$ 0.20	0.96 $\pm$ 0.18	0.60 $\pm$ 0.05	
			NLL	<b>1.32 <math>\pm</math> 0.08</b>	<b>0.32 <math>\pm</math> 0.18</b>	<b>1.35 <math>\pm</math> 0.21</b>	2.05 $\pm$ 0.15	1.54 $\pm$ 0.21	
Lorenz Attractor	12	VCL-BSSM	NMSE	0.26 $\pm$ 0.00	0.22 $\pm$ 0.02	0.25 $\pm$ 0.03	0.27 $\pm$ 0.02	0.28 $\pm$ 0.01	
			NLL	4.42 $\pm$ 0.04	4.22 $\pm$ 0.16	4.41 $\pm$ 0.08	4.55 $\pm$ 0.09	4.52 $\pm$ 0.04	
		CDDP	NMSE	<b>0.24 <math>\pm</math> 0.01</b>	<b>0.21 <math>\pm</math> 0.02</b>	<b>0.24 <math>\pm</math> 0.02</b>	<b>0.25 <math>\pm</math> 0.02</b>	<b>0.26 <math>\pm</math> 0.02</b>	
			NLL	<b>4.35 <math>\pm</math> 0.06</b>	<b>4.06 <math>\pm</math> 0.19</b>	<b>4.33 <math>\pm</math> 0.02</b>	<b>4.53 <math>\pm</math> 0.06</b>	<b>4.46 <math>\pm</math> 0.03</b>	
Libras	15	VCL-BSSM	NMSE	<b>0.14 <math>\pm</math> 0.00</b>	<b>0.13 <math>\pm</math> 0.01</b>	0.14 $\pm$ 0.00	<b>0.13 <math>\pm</math> 0.01</b>	0.15 $\pm$ 0.00	<b>0.14 <math>\pm</math> 0.00</b>
			NLL	-0.37 $\pm$ 0.02	-0.41 $\pm$ 0.04	-0.35 $\pm$ 0.03	<b>-0.42 <math>\pm</math> 0.04</b>	-0.32 $\pm$ 0.03	-0.36 $\pm$ 0.03
		CDDP	NMSE	<b>0.14 <math>\pm</math> 0.00</b>	0.14 $\pm$ 0.00	<b>0.13 <math>\pm</math> 0.01</b>	0.14 $\pm$ 0.01	<b>0.14 <math>\pm</math> 0.01</b>	<b>0.14 <math>\pm</math> 0.01</b>
			NLL	<b>-0.39 <math>\pm</math> 0.04</b>	<b>-0.42 <math>\pm</math> 0.03</b>	<b>-0.40 <math>\pm</math> 0.04</b>	-0.38 $\pm$ 0.05	<b>-0.40 <math>\pm</math> 0.03</b>	<b>-0.37 <math>\pm</math> 0.04</b>
Character Trajectories	20	VCL-BSSM	NMSE	0.87 $\pm$ 0.04	<b>0.42 <math>\pm</math> 0.04</b>	0.80 $\pm$ 0.04	0.90 $\pm$ 0.02	1.11 $\pm$ 0.08	1.11 $\pm$ 0.10
			NLL	0.14 $\pm$ 0.02	<b>-0.23 <math>\pm</math> 0.06</b>	0.06 $\pm$ 0.05	0.20 $\pm$ 0.05	0.31 $\pm$ 0.08	0.35 $\pm$ 0.09
		CDDP	NMSE	<b>0.64 <math>\pm</math> 0.01</b>	0.52 $\pm$ 0.05	<b>0.56 <math>\pm</math> 0.02</b>	<b>0.69 <math>\pm</math> 0.04</b>	<b>0.71 <math>\pm</math> 0.03</b>	<b>0.71 <math>\pm</math> 0.02</b>
			NLL	<b>-0.19 <math>\pm</math> 0.03</b>	-0.05 $\pm$ 0.05	<b>-0.26 <math>\pm</math> 0.04</b>	<b>-0.17 <math>\pm</math> 0.03</b>	<b>-0.23 <math>\pm</math> 0.04</b>	<b>-0.26 <math>\pm</math> 0.04</b>

**Context Length.** We select context length empirically as one-third of the sequence length amounting to five for Sine Waves, eight for Lotka-Volterra, 16 for the Lorenz Attractor, 15 for Libras, and 35 for Character Trajectories. The reason for that is, generally, one-third of the sequence captures the general characteristic of the dynamics and gives indications about modes.

**Neural Network Architecture Details.** Our CDDP has four main architectural elements: **Encoder:** The sequence encoder  $e_\lambda(x_{t:t'})$  governs the mean of our normal distributed recognition model  $q_\psi(x_0|y_{1:C}, \pi)$ . We feed  $C$  observations as a stacked set of values into the encoder. The sequence encoder is a single dense layer for Sine Waves, Lotka-Volterra, Libras; and a multi-layer perceptron for the Lorenz Attractor, and Character Trajectories with two hidden layers. The perceptron uses the  $\tanh(\cdot)$  activation function followed by layer normalization. **Decoder:** The likelihood function  $p(y_t|x_t)$  of the base model serves as a probabilistic decoder that maps the latent state  $x_t$  to the observed state  $y_t$ . We choose the emission distribution to be normal with mean governed by a single dense layer for Sine Waves, Lotka-Volterra, Libras; and a multilayer perceptron with two hidden layers for the rest. The perceptron uses the  $\tanh(\cdot)$  activation function followed by layer normalization. **Transition Kernel:** We choose the transition kernel  $p(x_t|x_{t-1}, m, \theta)$  of CDDP to be a normal distribution with mean governed by a plain RNN that receives a concatenation of the previous hidden state and the mode descriptor as input. The RNN on the mean is a multilayer perceptron with one hidden layer. The perceptron uses the  $\tanh(\cdot)$  activation function followed by layer normalization. The transition kernel of the base model of VCL  $p(x_t|x_{t-1}, \theta)$  follows the same RNN architecture except that its input does not contain a mode descriptor. **External Memory:** We set the memory size to 20 for the Sine Wave environment, 10 for Lotka-Volterra, 15 for the Lorenz Attractor, 20 for Libras, and 30 for Character Trajectories.

**Main Results.** Table 1 presents model performance throughout the whole CL process as the area under the learning curve. Our CDDP outperforms the parameter transfer-based VCL-BSSM baseline consistently in nearly all cases. Storing mode descriptors of the learned dynamics in an external memory, retrieving them in the subsequent tasks, and feeding them into the state transition kernel prevents catastrophic forgetting more effectively than plain parameter transfer. As seen in Figure 2, in the challenging character trajectories dataset, the prediction accuracy of our CDDP significantly outperforms VCL-BSSM in both the early and late stages of the CL period.

**Ablation Study.** We investigate the contribution of individual design choices to the total performance of our target model. We study the effect of three design choices: i) knowledge transfer via parameters  $\theta$  of the learned transition dynamics, ii) quantifying the uncertainty of the parameters of transition dynamics by a distribution  $q_\psi(\theta)$ , and iii) maintaining an external

Table 2: Ablation study results given as the AUC on the Sine Waves data set. The reported numbers are mean  $\pm$  standard error over 5 repetitions.

	Parameter Transfer	Probabilistic Parameters	Memory Content	NMSE	NLL
RNN	✓	×	N/A	$1.03 \pm 0.05$	$3.50 \pm 0.10$
VCL-BSSM	✓	✓	N/A	$0.93 \pm 0.03$	$3.59 \pm 0.16$
CDDP Variants	×	✓	Zeros	$0.94 \pm 0.04$	$3.56 \pm 0.12$
	×	✓	Ones	$1.12 \pm 0.13$	$3.75 \pm 0.14$
	×	✓	Twos	$1.35 \pm 0.23$	$3.80 \pm 0.18$
	✓	✓	Learned	$0.91 \pm 0.04$	$3.56 \pm 0.12$
CDDP Target	×	✓	Learned	<b><math>0.87 \pm 0.03</math></b>	<b><math>3.35 \pm 0.15</math></b>

memory with learned or unlearned content. Table 2 shows a map of model variants corresponding to the activation status of these three design choices, as well as the corresponding numerical results on the Sine Waves data set over five repetitions. We observe a performance increase when knowledge transfer is done via the external memory *instead of—but not together with*—parameter transfer, supporting the central assumptions of our target model. Setting the memory content to unlearned values causes rapid performance deterioration as values diverge from the learned values. This outcome demonstrates the essential role of the external memory in the CL performance of CDDP.

## 7. Conclusion

**Summary.** We report the first study on CL of multi-modal dynamical systems. We curate a competitive baseline for this new problem setup from an adaptation of VCL to BSSMs. We introduce a novel alternative to the parameter transfer approach of VCL for within-task knowledge acquisition and cross-task knowledge transfer using an original combination of neural episodic memory, DPs, and BSSMs. We observe in CL of challenging multi-modal dynamics modeling environments that our alternative approach compares favorably to the established parameter transfer approach.

**Broad Impact.** Our work can be used for varying applications such as: i) weather forecasting, where features can be transferred from one climate to another, ii) autonomous driving, where driving patterns can be adapted across different countries, and iii) model-based reinforcement learning algorithms, when the environment changes due either to the actions of the ego agent or to external factors. The memory architecture of CDDP may be improved by alternative embedding, update, and attention mechanisms. Our formulation of the transition rule is agnostic to the architecture that governs the transition kernel.

## References

- A. Chaudhry, M. Ranzato, M. Rohrbach, and M. Elhoseiny. Efficient lifelong learning with A-GEM. In *ICLR*, 2019.
- A. Cossu, A. Carta, V. Lomonaco, and D. Bacciu. Continual learning for recurrent neural networks: An empirical evaluation. *Neural Networks*, 143:607–627, 2021. ISSN 0893-6080.
- A. Doerr, C. Daniel, M. Schiegg, N. Duy, S. Schaal, M. Toussaint, and T. Sebastian. Probabilistic recurrent state-space models. In *ICML*, 2018.
- D. Dua and C. Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.
- E.B. Fox, E.B. Sudderth, M.I. Jordan, and A.S. Willsky. A sticky HDP-HMM with application to speaker diarization. *The Annals of Applied Statistics*, pages 1020–1056, 2011.
- M. Fraccaro, D.J. Rezende, Z. Zwols, A. Pritzel, S.M.A. Eslami, and F. Viola. Generative temporal models with spatial memory for partially observed environments. In *ICML*, 2018.
- R. Frigola, Y. Chen, and C. E. Rasmussen. Variational Gaussian process state-space models. In *NeurIPS*, 2014.
- M. Garnelo, J. Schwarz, D. Rosenbaum, F. Viola, D.J. Rezende, S.M. Eslami, and Y.W. Teh. Neural processes. *arXiv preprint arXiv:1807.01622*, 2018.
- A. Graves, G. Wayne, and I. Danihelka. Neural Turing Machines. *arXiv preprint arXiv:1410.5401*, 2014.
- Y. Guo, M. Liu, T. Yang, and T. Rosing. Improved schemes for episodic memory-based lifelong learning. *NeurIPS*, 2020.
- D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson. Learning latent dynamics for planning from pixels. In *ICML*, 2019.
- D. Hafner, T. Lillicrap, J. Ba, and M. Norouzi. Dream to control: Learning behaviors by latent imagination. In *ICLR*, 2020.
- M. Haußmann, S. Gerwinn, A. Look, B. Rakitsch, and M. Kandemir. Learning partially known stochastic dynamics with empirical PAC Bayes. In *AISTATS*, 2021.
- A.D. Ialongo, M. van der Wilk, J. Hensman, and C.E. Rasmussen. Overcoming mean-field approximations in recurrent Gaussian process models. In *ICML*, 2019.
- M. Kandemir, A. Akgül, M. Haussmann, and G. Unal. Evidential Turing processes. In *ICLR*, 2022.
- H. Kim, A. Mnih, J. Schwarz, M. Garnelo, A. Eslami, D. Rosenbaum, O. Vinyals, and Y.W. Teh. Attentive neural processes. In *ICLR*, 2019.

- J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A.A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, D. Hassabis, C. Clopath, D. Kumaran, and R. Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13), 2017.
- J. Knoblauch, H. Husain, and T. Diethe. Optimal continual learning has perfect memory and is np-hard. In *ICML*, 2020.
- N. Loo, S. Swaroop, and R.E. Turner. Generalized variational continual learning. In *ICLR*, 2021.
- D. Lopez-Paz and M. Ranzato. Gradient episodic memory for continual learning. In *NeurIPS*, 2017.
- B. Lüders, M. Schläger, and S. Risi. Continual learning through evolvable neural turing machines. 2016.
- C.V. Nguyen, Y. Li, T.D. Bui, and R.E. Turner. Variational continual learning. In *ICLR*, 2018.
- J. Pitman et al. Combinatorial stochastic processes. Technical report, 2002.
- A. Rios and L. Itti. Closed-loop memory GAN for continual learning. In *IJCAI*, 2019.
- V.G. Satorras, Z. Akata, and M. Welling. Combining generative and discriminative models for hybrid inference. In *NeurIPS*, 2019.
- J. Sethuraman. A constructive definition of Dirichlet priors. *Statistica Sinica*, 4(2):639–650, 1994.
- H. Shin, J.K. Lee, J. Kim, and J. Kim. Continual learning with deep generative replay. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *NeurIPS*, 2017.
- G. Singh, J. Yoon, Y. Son, and S. Ahn. Sequential neural processes. In *NeurIPS*, 2019.
- Y.W. Teh, M.I. Jordan, M.J. Beal, and D.M. Blei. Hierarchical Dirichlet processes. *Journal of the American Statistical Association*, 101(476):1566–1581, 2006.
- C. Yildiz, M. Heinonen, and H. Lähdesmäki. ODE<sup>2</sup>VAE: Deep generative second order ODEs with Bayesian neural networks. In *ICML*, 2019.
- F. Zenke, B. Poole, and S. Ganguli. Continual learning through synaptic intelligence. In *ICML*, 2017.