

# HopCast: Calibration of Autoregressive Dynamics Models

Anonymous authors

Paper under double-blind review

## Abstract

Deep learning models are often trained to approximate dynamical systems that can be modeled using differential equations. Many of these models are optimized to predict one step ahead; such approaches produce calibrated one-step predictions if the predictive model can quantify uncertainty, such as Deep Ensembles. At inference time, multi-step predictions are generated via autoregression, which needs a sound uncertainty propagation method to produce calibrated multi-step predictions. This work introduces an alternative Predictor-Corrector approach named HOPCAST that uses Modern Hopfield Networks (MHN) to learn the errors of a deterministic Predictor that approximates the dynamical system. The Corrector predicts a set of errors for the Predictor’s output based on a context state at any timestep during autoregression. The set of errors creates sharper and well-calibrated prediction intervals with higher predictive accuracy compared to baselines without uncertainty propagation. The calibration and prediction performances are evaluated across a set of dynamical systems. This work is also the first to benchmark existing uncertainty propagation methods based on calibration errors.

## 1 Introduction

Approximating dynamical systems that can be modeled using differential equations has many applications in systems biology (Wang et al., 2021; Brauer & Kribs, 2016), control (Moerland et al., 2023; Lu et al., 2021), economics (Tu, 2012), and cyber-physical systems (Bilal Shahid & Fleming, 2025; Shahid et al., 2024; Robison et al., 2024). Many deep learning approaches to approximate dynamical systems optimize the model to predict one step ahead during training (Khansari-Zadeh & Billard, 2011; Coates et al., 2008; Ko et al., 2007). The multi-step predictions are generated via autoregression at inference time (Lu et al., 2021). In doing so, the model error accumulates over time, resulting in multi-step predictions diverging from the ground truth (Venkatraman et al., 2015; Janner et al., 2021). Hence, it is imperative to develop methods to generate accurate and calibrated multi-step predictions.

State-of-the-art methods for generating accurate and calibrated predictions include Probabilistic Ensembles (Deep Ensembles) (Lakshminarayanan et al., 2017). To approximate the dynamics typically associated with differential equations, these models are optimized to predict one step ahead during training and then autoregressively generate multi-step predictions at inference time. To generate calibrated multi-step predictions, Probabilistic Ensembles need sound uncertainty propagation methods such as Trajectory Sampling (Chua et al., 2018).

In contrast, we introduce a Predictor-Corrector mechanism to produce accurate and calibrated multi-step predictions for dynamical systems without the need for uncertainty propagation. The Predictor is optimized to predict a one-step ahead point estimate of the next state of the system during training and produce multi-step predictions via autoregression at inference time. Our Corrector is based on the idea that similar states of a system lead to similar errors. This idea was leveraged by Auer et al. (2023) to build calibrated prediction intervals for one-step prediction utilizing Modern Hopfield Networks (MHN). However, the same state may not lead to the same error for multi-step predictions for dynamical systems. In such situations, the error also depends on how the system arrives at that state, i.e., where the system started and when it reached the state. To that end, we introduce the concept of *context state*, assuming that similar context states lead to similar errors. The Corrector, utilizing MHN, learns similarity between context states in terms of errors. At

any timestep during autoregression, the Predictor predicts, and the Corrector generates prediction intervals for the prediction. To control the width of prediction intervals for calibration, we introduce the concept of Attention Span. We name our approach HOPCAST.

Our contributions are:

- the introduction of the Predictor-Corrector mechanism, where the Predictor autoregressively generates the next state for a dynamical system, and the Corrector predicts a set of errors for the full forecasting horizon. The expected error corrects the prediction, while the set of errors generates the calibrated prediction intervals.
- a method – HOPCAST – that has no assumptions about the Predictor except that it predicts a point estimate of the next state of the system. It is therefore a general framework that can provide multi-step calibrated predictions irrespective of the form or capability of the underlying Predictor.
- lower calibration and prediction error across several benchmarks, without the use of complex uncertainty propagation techniques. Unlike uncertainty propagation methods, HOPCAST creates sharper prediction intervals based on context state instead of accumulating uncertainty over time during autoregression.

## 2 Related Work

Since we propose a mechanism to quantify the predictive uncertainty and generate calibrated prediction intervals, the related works include those from the uncertainty quantification (UQ) and model calibration. In UQ, there are two lines of work: Bayesian and Ensemble methods.

**Bayesian Methods.** In Bayesian models, a distribution over the parameters of the underlying model is learned, called posterior distribution. At inference time, multiple samples of parameters can be taken from the posterior distribution using an accurate sampling method, such as Hamiltonian Monte Carlo (HMC) (Betancourt, 2017). These samples serve as multiple predictive models and give us the desired diversity in the predictive variable of interest. A notable method from this domain is Bayes by Backprop (Blundell et al., 2015), which learns an approximate posterior distribution over the weights of a feedforward model via Variational Inference (VI) by optimizing the evidence lower bound (ELBO) (Kingma et al., 2015).

**Ensemble Methods.** In classical *ensemble* methods, an ensemble of models is trained with variations in the data for each model in the ensemble to train robust predictors (Davison & Hinkley, 1997). Deep Ensembles, however, are constructed using modern deep learning models with random initializations in the parameter space. These over-parametrized models, when initialized randomly, converge to different local minima, giving us the desired diversity in the predictive variable (Fort et al., 2019). These models were first presented by Lakshminarayanan et al. (2017) in the context of predictive uncertainty estimation. A notable extension of that is Neural Ensemble Search (Zaidi et al., 2021), which uses different architectures for models in an ensemble and improves upon deep ensembles. The Deep Ensemble is a scalable way of quantifying uncertainty compared to Bayesian methods (Wilson & Izmailov, 2020). Hence, these are used as baseline models with different uncertainty propagation methods in the current work.

**Model Calibration.** Calibration has been studied widely for classification models. In binary classification, Platt Scaling (Platt et al., 1999) and isotonic regression (Niculescu-Mizil & Caruana, 2005) have been used successfully for recalibration. There are extensions of such works to multi-class classification problems (Zadrozny & Elkan, 2002). In the context of regression, Gneiting & Raftery (2007) proposed several proper scoring rules to evaluate the calibration of a probabilistic model for continuous variables. Those scoring rules have been used in the literature as loss functions, for example, continuous ranked probability score (CRPS) (Gasthaus et al., 2019). Calibration has also been discussed in the literature on probabilistic forecasting, mainly in the context of meteorology (Gneiting & Raftery, 2005), resulting in specialized calibration systems (Raftery et al., 2005). An approach called calibrated regression was proposed by Kuleshov et al. (2018), which used isotonic regression to recalibrate Bayesian models. In contrast, our work turns a deterministic Predictor into a calibrated Predictor via a Predictor-Corrector mechanism. Also, we focus on Predictors

that *autoregressively* generate the next state of a dynamical system, a setting rarely addressed in the model calibration literature to the best of our knowledge.

### 3 Problem Description

Consider a dynamical system that can be defined using differential equations, such as Lotka-Volterra (LV). The LV system is used throughout the following discussion for illustrative purposes without loss of generality; subsequent sections show results for a variety of systems. LV dynamics can be written as:

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

$$\frac{dy}{dt} = \delta xy - \gamma y \quad (2)$$

The state of this system can be defined as:  $\tilde{\mathbf{s}} = (x, y)$ . We assume a set of ground-truth  $N$  trajectories of  $T$  timesteps  $I_N = \{\tau_n\}_{n=1}^N$  where  $\tau_n = \{(\mathbf{s}_t^n, \mathbf{s}_{t+1}^n)\}_{t=1}^T$  and  $\mathbf{s}_t^n = (x_t^n, y_t^n)$ . This dataset can be generated by solving LV with a numerical ODE solver such as `solve_ivp` from `scipy` with random initial conditions. To simulate measurement noise, we perturb each state post-integration with additive Gaussian noise, i.e.,  $\mathbf{s} = \tilde{\mathbf{s}} + \epsilon$ , where  $\epsilon \sim \mathcal{N}(0, \sigma * \Sigma)$  and  $\Sigma = \text{diag}(\sigma_x, \sigma_y)$ . The  $\sigma$  is the scaling factor. We name  $I_N$  an *Integrated* dataset. We assume that there is a Predictor  $B$  (trained on  $I_N$ ) that predicts  $\tilde{\mathbf{s}}_{t+1}$ <sup>1</sup> given  $\tilde{\mathbf{s}}_t$ , autoregressively generating a predicted set of  $N$  trajectories starting from the same initial conditions. Precisely,  $A_N = \{\bar{\tau}_n\}_{n=1}^N$  where  $\bar{\tau}_n = \{(\bar{\mathbf{s}}_t^n, \bar{\mathbf{s}}_{t+1}^n)\}_{t=1}^T$  and  $\bar{\mathbf{s}}_t^n = (\bar{x}_t^n, \bar{y}_t^n)$ . We call  $A_N$  an *Autoregressive* dataset. The error ( $\mathbf{e}_t^n$ ) is the difference between  $t^{\text{th}}$  timestep of  $n^{\text{th}}$  trajectory from *Integrated* and *Autoregressive*, that is,  $\mathbf{e}_t^n = (e_{x_t^n}, e_{y_t^n})$  where  $\mathbf{e}_t^n = \mathbf{s}_t^n - \bar{\mathbf{s}}_t^n$ . Therefore, the set of errors of  $N$  trajectories is  $F_N = \{\{\mathbf{e}_t^n\}_{t=1}^T\}_{n=1}^N$ .

For the LV example, we assume two correction models,  $M_x$  and  $M_y$ , corresponding to each output of the Predictor  $B$ .  $M_x$  and  $M_y$  will predict a set of errors  $E_{\bar{x}} = \{\varepsilon_{\bar{x}_{t+1}}^i\}_{i=1}^s$  and  $E_{\bar{y}} = \{\varepsilon_{\bar{y}_{t+1}}^i\}_{i=1}^s$  associated with  $t^{\text{th}}$  timestep of a trajectory at inference time, respectively. The  $z\%$  prediction interval, i.e.,  $\text{PI}_{\bar{x}_{t+1}}^z$  and  $\text{PI}_{\bar{y}_{t+1}}^z$  for  $\bar{x}_{t+1}$  and  $\bar{y}_{t+1}$  at  $t^{\text{th}}$  timestep can be generated as follows:

$$\text{PI}_{\bar{x}_{t+1}}^z = \bar{x}_{t+1} + [Q_\alpha(E_{\bar{x}}), Q_{1-\alpha}(E_{\bar{x}})] \quad (3)$$

$$\text{PI}_{\bar{y}_{t+1}}^z = \bar{y}_{t+1} + [Q_\alpha(E_{\bar{y}}), Q_{1-\alpha}(E_{\bar{y}})] \quad (4)$$

where  $\alpha = \frac{1-z}{2}$ . For example, a 90% prediction interval corresponds to  $\alpha = 0.05$ . The  $Q_\alpha$  denotes the  $\alpha$ -quantile. We generate prediction intervals for the outputs of Predictor  $B$  at  $t^{\text{th}}$  timestep during autoregression solely based on the set of errors predicted by correction models ( $M_x$  and  $M_y$ ), thereby avoiding uncertainty propagation.

### 4 HopCast

The proposed methodology, HOPCAST, consists of a Predictor-Corrector mechanism shown in Fig. 1. At any timestep ( $t$ ) during autoregression, the Predictor  $B$  predicts the next state of the system ( $\bar{x}_{t+1}, \bar{y}_{t+1}$ ) given previous state ( $\bar{x}_t, \bar{y}_t$ ). The Corrector consists of correction models ( $M_x$  and  $M_y$ ), one for each output of the Predictor  $B$ . Particularly, each correction model ( $M$ ) consists of an Encoder ( $m$ ) and an MHN (Ramsauer et al., 2020; Auer et al., 2023). The Encoder ( $m$ ) is a fully connected feedforward model. Architectural details are included in Appendix D.1. We frame correction as a pattern retrieval task (Ramsauer et al., 2020). Precisely,  $M_x$  and  $M_y$  take in the context state ( $\bar{x}_1, \bar{y}_1, \bar{x}_t, \bar{y}_t, t$ ) and outputs  $E_{\bar{x}}$  and  $E_{\bar{y}}$ , respectively.  $E_{\bar{x}}$  and  $E_{\bar{y}}$  contains errors of similar context states from past. The context state is the input of Predictor  $B$  ( $\bar{x}_t, \bar{y}_t$ ) augmented with the initial condition ( $\bar{x}_1, \bar{y}_1$ ) and time step ( $t$ ). The  $E_{\bar{x}}$  &  $E_{\bar{y}}$  are used to derive prediction intervals.

The rest of this section has three sub-sections. Section 4.1 discusses the shuffling of  $I_N$ ,  $A_N$ , &  $F_N$  datasets to train correction models for pattern retrieval tasks. Section 4.2 then discusses the training of the correction model ( $M_x$ ), and section 4.3 describes the inference with  $M_x$  where we generate the set of errors ( $E_{\bar{x}}$ ) given the context state.

<sup>1</sup>The overhead bar over a variable shows that it belongs to *Autoregressive* dataset.

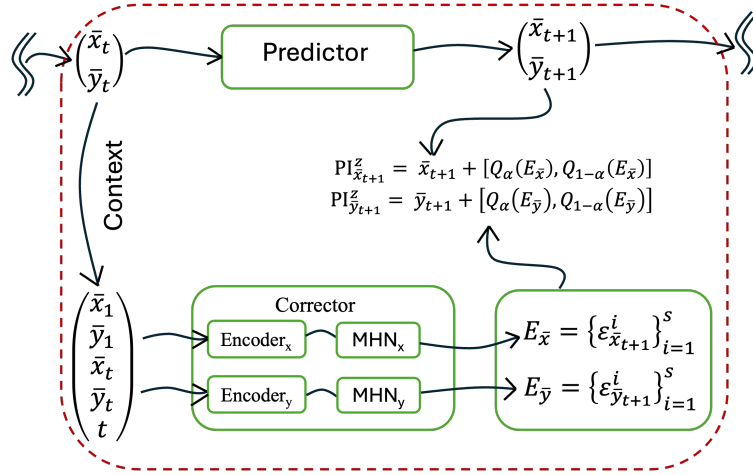


Figure 1: Predictor-Corrector mechanism

#### 4.1 Shuffled Data

**Shuffled Queries ( $\mathcal{S}_Q$ ).** Before training the correction model ( $M_x$ ), the states of the system from *Autoregressive* ( $A_N$ ) dataset are augmented with context information, that is, initial condition and time step. Appendix D.2 emphasizes the importance of adding initial condition as a context. Specifically, the  $t^{\text{th}}$  timestep of  $n^{\text{th}}$  trajectory  $\bar{\tau}_n$  from  $A_N$ , i.e.,  $(\bar{x}_t^n, \bar{y}_t^n)$ , will be  $\bar{c}_t^n = (\bar{x}_1^n, \bar{y}_1^n, \bar{x}_t^n, \bar{y}_t^n, t)$  after adding context information. The tuple  $(\bar{x}_1^n, \bar{y}_1^n)$  is the initial condition of trajectory  $\bar{\tau}_n$ . The  $\bar{c}_t^n$  denotes the context state associated with the  $t^{\text{th}}$  timestep of  $n^{\text{th}}$  trajectory. The set of all context states constructed from dataset  $A_N$  is  $\{\{\bar{c}_t^n\}_{t=1}^T\}_{n=1}^N$ . This nested set of context states is flattened and shuffled randomly to form a new set  $\mathcal{S}_Q = \text{Shuffle} \left( \bigcup_{n=1}^N \{\bar{c}_t^n\}_{t=1}^T \right)$ . Appendix D.5 highlights the importance of shuffling the data.

**Shuffled Keys ( $\mathcal{S}_K$ ).** Likewise, we construct the context states  $\{\{c_t^n\}_{t=1}^T\}_{n=1}^N$  for *Integrated* ( $I_N$ ) dataset as well, where  $c_t^n = (x_1^n, y_1^n, x_t^n, y_t^n, t)$ . Similar to  $\mathcal{S}_Q$ , we flatten the nested set  $\{\{c_t^n\}_{t=1}^T\}_{n=1}^N$  and randomly shuffle it to form  $\mathcal{S}_K = \text{Shuffle} \left( \bigcup_{n=1}^N \{c_t^n\}_{t=1}^T \right)$ .

**Shuffled Values ( $\mathcal{S}_{V_x}$  &  $\mathcal{S}_{V_y}$ ).** Given errors  $F_N = \{\{\mathbf{e}_t^n\}_{t=1}^T\}_{n=1}^N$  where  $\mathbf{e}_t^n = (e_{x_t^n}, e_{y_t^n})$ ,  $F_N$  can be split into two sets of errors  $F_N^x = \{\{e_{x_t^n}\}_{t=1}^T\}_{n=1}^N$  and  $F_N^y = \{\{e_{y_t^n}\}_{t=1}^T\}_{n=1}^N$ . Recall that we train separate correction models ( $M_x$  and  $M_y$ ) for each output of Predictor  $B$ .  $F_N^x$  and  $F_N^y$  are flattened and shuffled to form  $\mathcal{S}_{V_x} = \text{Shuffle} \left( \bigcup_{n=1}^N \{e_{x_t^n}\}_{t=1}^T \right)$  and  $\mathcal{S}_{V_y} = \text{Shuffle} \left( \bigcup_{n=1}^N \{e_{y_t^n}\}_{t=1}^T \right)$ .

Once we have  $\mathcal{S}_Q$ ,  $\mathcal{S}_K$ ,  $\mathcal{S}_{V_x}$ ,  $\mathcal{S}_{V_y}$ , they are split into  $\mathcal{S}_Q^{\text{Train}}$ ,  $\mathcal{S}_K^{\text{Train}}$ ,  $\mathcal{S}_{V_x}^{\text{Train}}$ ,  $\mathcal{S}_{V_y}^{\text{Train}}$  and  $\mathcal{S}_Q^{\text{Inference}}$ ,  $\mathcal{S}_K^{\text{Inference}}$ ,  $\mathcal{S}_{V_x}^{\text{Inference}}$ ,  $\mathcal{S}_{V_y}^{\text{Inference}}$  with an 80/20 split. To train  $M_x$ , we need dataset  $\mathcal{D}_x = \{(\mathbf{Q}, \mathbf{K}, \mathbf{V}_x)_i\}$  comprising triplets  $(\mathbf{Q}, \mathbf{K}, \mathbf{V}_x)$ . Each triplet of  $\mathbf{Q}$ ,  $\mathbf{K}$ , and  $\mathbf{V}_x$  are constructed by sampling  $\mathbf{S}_L$  (Sequence Length) number of elements from sets  $\mathcal{S}_Q^{\text{Train}}$ ,  $\mathcal{S}_K^{\text{Train}}$ , and  $\mathcal{S}_{V_x}^{\text{Train}}$ , respectively. For instance, one triplet with  $\mathbf{S}_L = 5$  is:  $\mathbf{Q} = (\bar{c}_4^3, \bar{c}_5^1, \bar{c}_7^5, \bar{c}_2^3, \bar{c}_1^7)^\top$ ,  $\mathbf{K} = (c_4^3, c_5^1, c_7^5, c_2^3, c_1^7)^\top$ , and  $\mathbf{V}_x = (e_{x_3^3}, e_{x_5^1}, e_{x_7^5}, e_{x_2^3}, e_{x_1^7})^\top$ . The error  $e_{x_4^3} = x_4^3 - \bar{x}_4^3$  denotes the error of Predictor  $B$  at 4<sup>th</sup> timestep of 3<sup>rd</sup> trajectory. Likewise, we construct  $\mathcal{D}_y = \{(\mathbf{Q}, \mathbf{K}, \mathbf{V}_y)_i\}$  using  $\mathcal{S}_Q^{\text{Train}}$ ,  $\mathcal{S}_K^{\text{Train}}$  and  $\mathcal{S}_{V_y}^{\text{Train}}$ . The 20% split with inference data is used for final evaluation.

#### 4.2 Training

The correction model ( $M_x$ ) consists of an Encoder ( $m_x$ ) and an MHN ( $\text{MHN}_x$ ). The Encoder  $m_x$  takes in as input the context state, such as  $c_4^3$ , and outputs a  $d$ -dimensional embedding. We choose  $d = 4$  in our



experiments, though other values of  $d$  are expected to perform comparably [Appendix D.4]. We use the same  $m_x$  to construct embedded queries  $\mathbf{Q}_\psi$  and keys  $\mathbf{K}_\psi$  from  $\mathbf{Q}$  and  $\mathbf{K}$ , respectively. The  $\text{MHN}_x$  uses an attention mechanism (Vaswani, 2017) to construct an association matrix ( $\mathbf{A}_{\mathbf{S}_L \times \mathbf{S}_L}$ ) based on the similarity of elements in  $\mathbf{Q}_\psi$  with elements in  $\mathbf{K}_\psi$ . Mathematically,

$$\mathbf{A}_{\mathbf{S}_L \times \mathbf{S}_L} = \text{softmax} \left( \frac{\mathbf{Q}_\psi \mathbf{K}_\psi^\top}{\sqrt{d}} \right) \quad (5)$$

$$\hat{\mathbf{V}}_x = \mathbf{A}_{\mathbf{S}_L \times \mathbf{S}_L} \cdot \mathbf{V}_x \quad (6)$$

$$\mathcal{L} = \frac{1}{\mathbf{S}_L} \|\mathbf{V}_x - \hat{\mathbf{V}}_x\|_2^2 \quad (7)$$

The association matrix ( $\mathbf{A}_{\mathbf{S}_L \times \mathbf{S}_L}$ ) is shown in the *Training* block of Fig. 2 with  $\mathbf{S}_L = 5$ . We mask the association from *Autoregressive* context state to its *Integrated* context state. The  $(i, j)^{th}$  entry of  $\mathbf{A}_{\mathbf{S}_L \times \mathbf{S}_L}$  shows the similarity of  $i^{th}$  element of  $\mathbf{Q}_\psi$  with  $j^{th}$  element of  $\mathbf{K}_\psi$ . We use **softmax** over  $\mathbf{A}_{\mathbf{S}_L \times \mathbf{S}_L}[i, :]$  to get association weights for  $i^{th}$  element of  $\mathbf{Q}_\psi$ . Therefore, each row of  $\mathbf{A}_{\mathbf{S}_L \times \mathbf{S}_L}$  sums up to one. The loss function  $\mathcal{L}$  is used to learn the parameters of the  $m_x$  via backpropagation, where the  $\hat{\mathbf{V}}_x$  denotes the predicted errors.

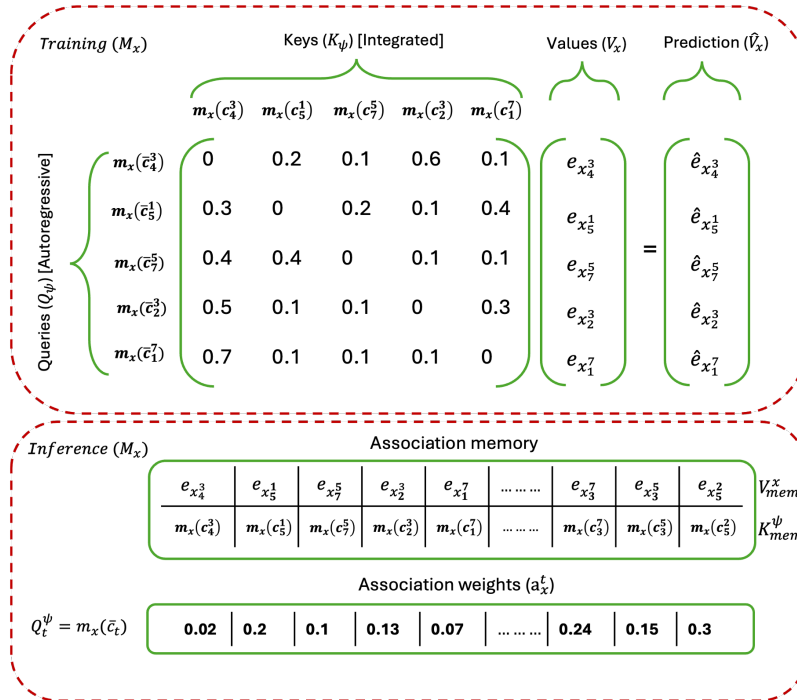


Figure 2: Training and Inference with Correction Model ( $\text{Encoder}_x + \text{MHN}_x$ ).

### 4.3 Inference

At inference time, a set of  $K$  keys ( $\mathbf{K}_{mem}$ ) with the corresponding values ( $\mathbf{V}_{mem}^x$ ) is sampled from  $\mathcal{S}_K^{\text{Train}}$  and  $\mathcal{S}_{V_x}^{\text{Train}}$ , respectively. For instance, a randomly sampled set of keys and values would be  $\mathbf{K}_{mem} = (c_4^3, c_5^1, c_7^5, \dots, c_5^2)$  and  $\mathbf{V}_{mem}^x = (e_{x_4^3}, e_{x_5^1}, e_{x_7^5}, \dots, e_{x_5^2})$ , respectively. The trained  $m_x$  is used to encode the  $\mathbf{K}_{mem}$  as  $\mathbf{K}_{mem}^\psi$ . The  $\mathbf{K}_{mem}^\psi$  and  $\mathbf{V}_{mem}^x$  constitute the Association memory of MHN. It is shown in the *Inference* block in Fig. 2. We loaded 2000 keys in memory for evaluation [Appendix D.3]. Once the memory is set up, the correction model ( $M_x$ ) takes in the context state ( $\bar{x}_1, \bar{y}_1, \bar{x}_t, \bar{y}_t, t$ ) as a query  $\mathbf{Q}_t$  at  $t^{th}$  time step during autoregression and encodes it as  $\mathbf{Q}_t^\psi$ . The association weights ( $\mathbf{a}_x^t$ ) at  $t^{th}$  timestep are,

$$\mathbf{a}_x^t = \text{softmax} \left( \frac{\mathbf{Q}_t^\psi \mathbf{K}_{mem}^\psi}{\sqrt{d}} \right) \quad (8)$$

where  $\mathbf{a}_x^t = \{p_i\}_{i=1}^K$ ,  $\sum_{i=1}^K p_i = 1$ , and  $K$  denotes the number of keys in memory. A set of  $s$  categories is sampled as:  $\{b_i\}_{i=1}^s \sim \text{Categorical}(\mathbf{a}_x^t)$ . A set of errors ( $E_{\bar{x}}$ ) is selected from  $\mathbf{V}_{mem}^x$  based on sampled categories  $\{b_i\}_{i=1}^s$ , i.e.,  $E_{\bar{x}} = \{V_{mem}^x[b_i]\}_{i=1}^s$ . A sampled set of errors  $E_{\bar{x}}$  is shown in Fig. 3 for the x-output of the Predictor  $B$  for the Lorenz system.




Figure 3: An example set of errors ( $E_x$ ) sampled from Association Memory at Timestep 106 during autoregression for the x-output of the Predictor for the Lorenz system, where the red dotted line denotes the ground truth error.

We have discussed the training of  $M_x$  and how to utilize it to obtain a set of errors ( $E_{\bar{x}}$ ) for the x-output of the Predictor  $B$  at the inference time. To obtain  $E_{\bar{y}}$  by training  $M_y$ , we repeat the same procedure stated in section 4.2 and section 4.3 with dataset  $\mathcal{D}_y$ , and this can be done for an arbitrary number of dimensions depending on the system dynamics.

## 5 Experiments

### 5.1 Evaluation Metrics

We propose three metrics to evaluate the quality of HOPCAST against the benchmarks.

**Calibration Error (Kuleshov et al., 2018) (CE).** A calibrated model has observed fractions ( $\hat{p}$ ) of its predicted random variable match with the expected fractions ( $p$ ). The difference between these two is calibration error. To evaluate CE,  $m(=9)$  equally spaced prediction intervals are chosen from 10% to 90%. For each prediction interval, we count the number of times the observed variable falls between intervals. Mathematically,  $\text{CE} = \sum_{i=1}^m (\hat{p}_i - p_i)^2$ .

**Prediction Interval Width (Auer et al., 2023) (PI-Width).** A calibrated model should also be sharp. Sometimes, it is trivial to reduce the CE by always predicting the expected value of a random variable (Kuleshov et al., 2018). A good predictor should be calibrated and sharp. Hence, we propose to use PI-Width as a measure of sharpness. Mathematically,

$$\text{PI-Width} = \frac{1}{m} \frac{1}{v} \sum_{i=1}^m \sum_{j=1}^v |U_j^i - L_j^i|, \quad (9)$$

where  $m = 9$  is for equally spaced PI, and  $v$  denotes the number of validation data points; i.e.,  $v = |\mathcal{S}_Q^{\text{Inference}}|$ .

**Mean Squared Error (MSE).** MSE is used to evaluate the predictive accuracy of the proposed approach against the baselines.

### 5.2 Datasets

We have discussed one dynamical system, i.e., LV, in section 3. Other dynamical systems include Lorenz, FitzHugh-Nagumo (FHN), Lorenz95, and the Glycolytic Oscillator. To generate datasets, we randomly

sample  $N$  initial conditions from within a specified range for the state variables of each system. The `solve_ivp` method from `scipy` is used to integrate the dynamics with the adaptive-step RK45 solver. To produce uniformly sampled trajectories, system states are extracted at fixed time intervals  $\Delta t$  as mentioned in Table 1. For Lorenz and LV, the dynamics of both system states and their derivatives are modeled, whereas the dynamics of states are modeled for the rest of the systems. The mathematical forms of each dynamical system, ranges of initial conditions, and parameter values are given in Appendix B.

Table 1: Details about data generation

Model	$\Delta t$	Timesteps	Trajectories
Lotka Volterra	0.1	300	500
Lorenz	0.01	300	1000
FHN	0.5	400	350
Lorenz95	0.01	300	666
Glycolytic	0.01	400	750

### 5.3 Baselines

The baseline methods include three uncertainty propagation approaches, i.e., Expectation, Moment Matching, and Trajectory Sampling. These approaches are used to propagate uncertainty with Probabilistic Ensembles (Chua et al., 2018). The details of Probabilistic Ensembles’ training and uncertainty propagation methods are presented in Appendix A.

## 6 Results

**Overall Comparison** As shown in Table 2, HOPCAST outperforms the baseline approaches in terms of CE and MSE in the vast majority of cases. It achieves the lowest average CE of 0.046 among all, with the second-best approach being the Trajectory Sampling with CE of 0.075. In terms of MSE, HOPCAST outperforms in 11 out of 15 cases, showing the effectiveness of our Predictor-Corrector mechanism. The Expectation and Moment Matching do not perform well in terms of CE and MSE. These two approaches generate conservative intervals with average PI-Widths of 7.56 and 3.63, respectively, resulting in overconfident models. The overconfidence is evident in the calibration curves of Expectation and Moment Matching in Fig. 6. The average CE of these two approaches is 0.21 and 0.68, respectively, showing high miscalibration compared to Trajectory Sampling and HOPCAST. PI-Width is only useful when discussed in conjunction with CE as it is trivial to reduce the PI-Width while being overconfident and miscalibrated.

**Attention Span** HOPCAST lets us control the confidence of the model with precise control over the width of PI based on a concept we introduce called *Attention Span*. The *Attention Span* can be increased/decreased by increasing/decreasing the  $S_L$  of  $\mathbf{Q}$  and  $\mathbf{K}$  in  $\mathbf{A}_{S_L \times S_L}$ . To demonstrate that, we collect a dataset  $H = \{(x_i, y_i)\}_{i=1}^{6000}$  using the sine function with heteroscedastic noise (Chua et al., 2018) shown in equation 10 and Fig. 4(a). The  $x_i$ ’s are uniformly sampled from the following domain  $[-\frac{5\pi}{2}, -\pi] \cup [\pi, \frac{5\pi}{2}]$ .

$$(x, y) \rightarrow \left( x, y + \mathcal{N} \left( 0, 0.2 \left| \sin \left( \frac{3}{2}x + \frac{\pi}{8} \right) \right| \right) \right) \quad (10)$$

We use the same correction model ( $M_x$ ) consisting of Encoder ( $m_x$ ) and  $MHN_x$  with the same training and inference procedures described in section 4.2 and 4.3, respectively. The only exception is that the set of triplets  $\{(\mathbf{Q}, \mathbf{K}, \mathbf{V})_i\}$  are constructed from dataset  $H$ . For instance, a triplet with  $S_L = 4$  is:  $\mathbf{Q} = (x_{10}, x_{13}, x_{14}, x_1)$ ,  $\mathbf{K} = (x_{10}, x_{13}, x_{14}, x_1)$ , and  $\mathbf{V} = (y_{10}, y_{13}, y_{14}, y_1)$ . We trained two  $M_x$  models with  $S_L$  of 3 & 8. At the inference time, the  $M_x$  is given a data point  $x = 3.28$  as a query to be associated with the keys in memory based on similarity in terms of  $y$ . The keys from memory are sampled based on their

Table 2: Prediction Interval Widths (**PI-Widths**) and Calibration Error (**CE**) metrics of HOPCAST and baseline approaches with different scaling factor ( $\sigma$ ) across various dynamical systems. The results are averaged across 3 runs of the same experiment. All of the CEs within 5% of the best one in a row are highlighted. The PI-Width is only used as a tie-breaker if there are different models with CEs within 5% of the best. In that case, the best-performing model based on CE is denoted with an asterisk and PI-Width is highlighted. For MSE, we simply highlight anything within 5% of the best one in a row.

Model		HopCast			Probabilistic Ensemble								
					Expectation			Moment Matching			Trajectory Sampling		
Metrics		MSE	PI-Width	CE	MSE	PI-Width	CE	MSE	PI-Width	CE	MSE	PI-Width	CE
Lotka Volterra	$\sigma = 0.05$	<b>6.87 <math>\pm</math> 0.32</b>	<b>2.93 <math>\pm</math> 0.17</b>	0.022 $\pm$ 0.004	20.58 $\pm$ 1.25	4.01 $\pm$ 0.40	<b>0.017 <math>\pm</math> 0.009</b>	28.9 $\pm$ 1.04	0.65 $\pm$ 0.06	0.56 $\pm$ 0.12	19.45 $\pm$ 0.36	5.43 $\pm$ 0.10	0.22 $\pm$ 0.11
	$\sigma = 0.1$	<b>6.10 <math>\pm</math> 0.20</b>	1.88 $\pm$ 0.05	<b>0.0070 <math>\pm</math> 0.002</b>	24.1 $\pm$ 0.15	0.83 $\pm$ 0.02	0.47 $\pm$ 0.17	23.9 $\pm$ 0.04	0.96 $\pm$ 0.04	0.41 $\pm$ 0.14	21.01 $\pm$ 0.27	5.87 $\pm$ 0.59	0.18 $\pm$ 0.10
	$\sigma = 0.3$	<b>9.69 <math>\pm</math> 0.23</b>	3.19 $\pm$ 0.12	<b>0.0051 <math>\pm</math> 0.004</b>	25.6 $\pm$ 0.09	2.13 $\pm$ 0.02	0.21 $\pm$ 0.05	25.4 $\pm$ 0.08	2.93 $\pm$ 0.02	0.059 $\pm$ 0.02	23.3 $\pm$ 0.05	6.19 $\pm$ 0.22	0.13 $\pm$ 0.07
Lorenz	$\sigma = 0.05$	<b>1126 <math>\pm</math> 7.97</b>	28.76 $\pm$ 0.74	0.011 $\pm$ 0.006	<b>1101.9 <math>\pm</math> 28.3</b>	30.37 $\pm$ 0.80	<b>0.002 <math>\pm</math> 0.001</b>	1991.3 $\pm$ 4.03	4.45 $\pm$ 0.009	1.23 $\pm$ 0.07	<b>1080.10 <math>\pm</math> 5.59</b>	38.79 $\pm$ 0.49	0.16 $\pm$ 0.04
	$\sigma = 0.1$	<b>1248.17 <math>\pm</math> 15</b>	38.66 $\pm$ 0.38	0.019 $\pm$ 0.03	1413.1 $\pm$ 38.83	35.49 $\pm$ 1.11	<b>0.011 <math>\pm</math> 0.015</b>	2579.7 $\pm$ 48.7	8.78 $\pm$ 0.01	1.21 $\pm$ 0.08	1314.60 $\pm$ 3.80	45.28 $\pm$ 0.69	0.078 $\pm$ 0.03
	$\sigma = 0.3$	<b>1681 <math>\pm</math> 46.36</b>	<b>40.44 <math>\pm</math> 0.91</b>	<b>0.019 <math>\pm</math> 0.004</b>	1848.5 $\pm$ 15.8	25.9 $\pm$ 1.43	0.38 $\pm$ 0.10	2098.5 $\pm$ 18.08	23.13 $\pm$ 0.06	0.58 $\pm$ 0.30	1788.6 $\pm$ 10.9	48.52 $\pm$ 0.35	<b>0.018 <math>\pm</math> 0.016</b>
FHN	$\sigma = 0.05$	<b>0.076 <math>\pm</math> 0.002</b>	0.20 $\pm$ 0.005	0.091 $\pm$ 0.004	0.68 $\pm$ 0.13	1.13 $\pm$ 0.061	<b>0.044 <math>\pm</math> 0.034</b>	0.95 $\pm$ 0.45	0.28 $\pm$ 0.007	0.91 $\pm$ 0.42	0.66 $\pm$ 0.04	1.25 $\pm$ 0.007	0.054 $\pm$ 0.026
	$\sigma = 0.1$	<b>0.17 <math>\pm</math> 0.03</b>	1.03 $\pm$ 0.05	0.32 $\pm$ 0.05	1.39 $\pm$ 0.17	0.76 $\pm$ 0.064	0.38 $\pm$ 0.19	2.21 $\pm$ 0.13	0.38 $\pm$ 0.003	1.42 $\pm$ 0.086	0.83 $\pm$ 0.005	1.35 $\pm$ 0.003	<b>0.051 <math>\pm</math> 0.02</b>
	$\sigma = 0.3$	<b>0.57 <math>\pm</math> 0.04</b>	1.27 $\pm$ 0.03	0.11 $\pm$ 0.02	1.53 $\pm$ 0.086	1.009 $\pm$ 0.046	0.39 $\pm$ 0.13	1.25 $\pm$ 0.003	1.52 $\pm$ 0.010	0.108 $\pm$ 0.019	1.12 $\pm$ 0.008	1.55 $\pm$ 0.004	<b>0.061 <math>\pm</math> 0.011</b>
Lorenz95	$\sigma = 0.05$	10.44 $\pm$ 0.08	4.41 $\pm$ 0.13	<b>0.007 <math>\pm</math> 0.003</b>	10.11 $\pm$ 0.21	3.25 $\pm$ 0.015	0.13 $\pm$ 0.017	14.11 $\pm$ 0.48	1.84 $\pm$ 0.009	0.62 $\pm$ 0.04	<b>9.30 <math>\pm</math> 0.031</b>	4.90 $\pm$ 0.023	0.022 $\pm$ 0.007
	$\sigma = 0.1$	13.98 $\pm$ 0.11	5.78 $\pm$ 0.018	<b>0.0028 <math>\pm</math> 0.007</b>	13.31 $\pm$ 0.20	3.29 $\pm$ 0.05	0.27 $\pm$ 0.025	17.38 $\pm$ 1.14	3.24 $\pm$ 0.03	0.35 $\pm$ 0.06	<b>11.63 <math>\pm</math> 0.11</b>	5.44 $\pm$ 0.04	0.004 $\pm$ 0.003
	$\sigma = 0.3$	16.44 $\pm$ 0.66	6.08 $\pm$ 0.06	<b>0.0015 <math>\pm</math> 0.001</b>	17.46 $\pm$ 0.11	4.61 $\pm$ 0.04	0.13 $\pm$ 0.029	16.37 $\pm$ 0.23	6.01 $\pm$ 0.03	0.0036 $\pm$ 0.003	<b>15.13 <math>\pm</math> 0.075</b>	6.49 $\pm$ 0.047	0.009 $\pm$ 0.006
Glycolytic Oscillator	$\sigma = 0.05$	<b>0.03 <math>\pm</math> 0.001</b>	0.20 $\pm$ 0.007	<b>0.043 <math>\pm</math> 0.018</b>	0.10 $\pm$ 0.006	0.24 $\pm$ 0.02	0.07 $\pm$ 0.04	0.13 $\pm$ 0.019	0.056 $\pm$ 0.002	1.35 $\pm$ 0.32	0.09 $\pm$ 0.004	0.36 $\pm$ 0.006	0.06 $\pm$ 0.026
	$\sigma = 0.1$	<b>0.072 <math>\pm</math> 0.003</b>	0.25 $\pm$ 0.007	<b>0.018 <math>\pm</math> 0.006</b>	0.12 $\pm$ 0.005	0.21 $\pm$ 0.01	0.10 $\pm$ 0.03	0.15 $\pm$ 0.005	0.08 $\pm$ 0.004	0.99 $\pm$ 0.18	0.10 $\pm$ 0.001	0.41 $\pm$ 0.01	0.07 $\pm$ 0.03
	$\sigma = 0.3$	<b>0.11 <math>\pm</math> 0.001</b>	<b>0.37 <math>\pm</math> 0.008</b>	<b>0.016 <math>\pm</math> 0.003</b>	0.18 $\pm$ 0.001	0.23 $\pm$ 0.01	0.62 $\pm$ 0.28	0.25 $\pm$ 0.002	0.21 $\pm$ 0.003	0.54 $\pm$ 0.07	0.17 $\pm$ 0.005	0.46 $\pm$ 0.003	<b>0.015 <math>\pm</math> 0.01</b>
Average	-	-	9.03	0.046	-	7.56	0.21	-	3.63	0.68	-	11.48	0.075

association weights. The frequency plots of sampled keys for  $S_L$  of 3 & 8 are in Fig. 4(b) & (c), respectively. For  $S_L = 3$ , the sampled keys are clustered around 3.28 in Fig. 4(b), which are the most similar keys to the query  $x = 3.28$ . As we increase  $S_L$  to 8, the  $M_x$  expands its *Attention Span* by picking other keys from memory around  $x = -3.2$  & 6.5 that have nearly similar  $y$  values as shown in Fig. 4(c).




Figure 4: (a) Upper center: Sine function with heteroscedastic noise (b) Lower left: Attention Span for  $x = 3.28$  as a query with  $S_L = 3$  (c) Lower right: Attention Span for  $x = 3.28$  as a query with  $S_L = 8$

**Calibration Performance** The reason HOPCAST performs well in terms of CE compared to baselines is owing to the  $S_L$  that can be increased/decreased to increase/decrease the diversity of sampled errors  $E_{\bar{x}}$  from MHN memory. As we increase the  $S_L$ , the  $M_x$  starts to expand its *Attention Span* by associating other keys from memory ( $K_{men}^\psi$ ) to the query ( $Q_t^\psi$ ) based on their similarity in terms of errors, effectively controlling the width of prediction intervals. We train separate correction models (e.g.,  $M_x$  and  $M_y$ ) for each output of the Predictor and tune their  $S_L$  separately. The  $S_L$  hyperparameter and number of models in a Probabilistic Ensemble for each setting are shown in Appendix C. For the baseline approaches, we tune the number of models in an ensemble. In general, we observed that the Probabilistic Ensembles tend to increase the uncertainty (or widen their prediction intervals) as we increase the number of models, and the uncertainty plateaus at some point. Therefore, if the uncertainty plateaus and the model is overconfident, adding more models to the ensemble will not make it a calibrated model. This was observed in the calibration curves of Expectation and Moment Matching in Fig. 6. Trajectory Sampling, in contrast, provided calibrated uncertainty with few models in the ensemble (i.e., 3 or 4) in most cases. It is owing to its nature of taking predicted uncertainty into account during propagation.




Figure 5: (a). Upper left: HOPCAST (b). Upper right: Expectation (c). Lower left: Moment Matching (d). Lower right: Trajectory Sampling. A comparison of prediction intervals generated by HOPCAST and three uncertainty propagation approaches. The x-axis denotes the time steps, and the y-axis shows the  $x$  output of the Lorenz system. The Upper PI and Lower PI show the  $m = 9$  equally spaced prediction intervals from 10% to 90%.

**Predictive Performance** HOPCAST outperforms in most cases in terms of MSE because we proposed a *Corrector* that predicts a set of errors, and prediction intervals are a side effect of the approach. Also, our model generates a smaller PI-width in most cases because we generate prediction intervals based on the context state  $(x_0, y_0, x_t, y_t, t)$  at any timestep during autoregression. On the other hand, uncertainty propagation approaches generate intervals based on uncertainty accumulated due to propagation up to a certain time step ( $t = t'$ ) starting from  $t = 1$ . A manifestation of that is shown in Fig. 5 (b) & (d). For Fig. 5(a), the output of Predictor  $B$  is shown in the dark, and the Mean indicates the prediction after adding the expected error ( $\sum_{i=1}^s \mathbf{a}_x^t[b_i]E_x[i]$ ) from the Corrector. For Fig. 5(b), (c), & (d), the Mean indicates the expected prediction from all models in a Probabilistic ensemble. Fig. 5(a) shows that HOPCAST reduces its PI-Width around time step 270 due to smaller error, resulting in sharper intervals. On the other hand, Expectation and Trajectory Sampling have wider PI due to accumulated uncertainty up to that time step, even though the Mean is close to the ground truth.

**Trajectory Sampling vs. Counterparts** On average, Trajectory Sampling performs better than its counterpart baselines—Expectation and Moment Matching—in terms of CE, owing to how it propagates uncertainty. It takes into account the predicted uncertainty of the model while propagating uncertainty from time step  $t$  to  $t+1$ . It samples the state for the next step from the predicted Multivariate Normal distribution rather than propagating just the predicted mean like the Expectation. Moment Matching also samples the next state from the predicted Multivariate Normal distribution, just like Trajectory Sampling. However, it recasts prediction from all models in an ensemble at every timestep ( $t$ ) to a Multivariate Normal distribution and then resamples next states before propagation. This results in overly conservative PI, as shown in Fig. 5(c), and evidenced by the PI-Width of 3.63 in Table 2.

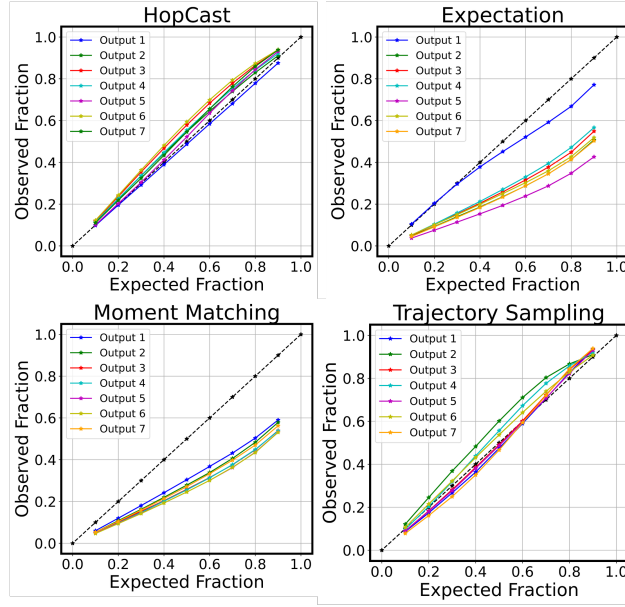


Figure 6: The calibration curves of HOPCAST and three uncertainty propagation methods for Glycolytic Oscillator at  $\sigma = 0.3$ . The dark dotted line corresponds to the perfect calibration. The calibration curves are shown separately for seven states of the system.

## 7 Hyperparameter Tuning

$S_L$  (Sequence Length) is a hyperparameter that must be tuned for calibrated uncertainty. We propose Algorithm 1 to hand-tune it based on the intuition that the uncertainty increases as we increase  $S_L$  as discussed in section 6 and shown in Fig. 7. We pick a relatively smaller (e.g., 10) and larger (e.g., 1000)  $S_L$  and expect the model to show overconfidence and underconfidence on the calibration curve, respectively. Next, we expect three cases. First, we pick  $S_L = 500$ , which is the optimal  $S_L$  because it is very close to perfect calibration (dark dotted line in Fig. 7). Second, our choice leads to overconfidence, e.g.,  $S_L = 400$ . Third, there is an underconfident model, e.g.,  $S_L = 600$ . In the first case, we have the calibrated uncertainty. In the second case, an  $S_L$  slightly bigger than 400 would be a good choice. In the third case, an  $S_L$  relatively smaller than 600 is a good choice. If it is the second or third case, the  $S_L$  needs to be re-tuned and we expect three cases again. We iteratively do this with the help of the calibration curve to find optimal  $S_L$  within or less than five to six trials. Appendix D.7 has additional results highlighting the impact of  $S_L$  on uncertainty.




Figure 7: Calibration curves of y-output of Lorenz at different  $S_L$ . The dark dotted line corresponds to perfect calibration.

The number of models in an ensemble is a hyperparameter. As discussed in the section 6, the ensembles tend to build up the uncertainty until it plateaus [Appendix D.6]. We keep adding models to the ensemble until it provides calibrated uncertainty. In some cases, the uncertainty saturates while the model is overconfident. In those cases, the number of optimal models is very high, e.g. 8 or 7. Appendix E contains additional details about hyperparameters.

---

**Algorithm 1**  $S_L$  Tuning for Calibration

---

```

1: Initialize OptimalSL  $\leftarrow$  True
2:  $\triangleright$  Pick a large  $S_L$  & a small  $S_L$ 
3:  $\triangleright$  Pick an  $S_L$  between large and small  $S_L$ 
4: while OptimalSL do
5:    $\triangleright$  Three possibilities on Calibration Curve
6:   if Calibrated Uncertainty then
7:     OptimalSL  $\leftarrow$  False
8:   else if Underconfidence then
9:     Decrease  $S_L$ 
10:  else if Overconfidence then
11:    Increase  $S_L$ 
12:  end if
13: end while

```

---

## 8 Conclusions

We proposed a Predictor-Corrector mechanism for autoregressive dynamics models to correct the Predictors' predictions and generate calibrated prediction intervals for it at any timestep during autoregression. HOPCAST performs competitively well against the alternative approach (i.e., uncertainty propagation), giving accurate and calibrated autoregressive dynamics models. Out of three uncertainty propagation approaches, i.e., Trajectory Sampling, Moment Matching, and Expectation, Trajectory Sampling performs competitively well against ours. This is due to its nature of taking predicted uncertainty into account during propagation. HOPCAST offers a lower calibration and prediction error with sharper prediction intervals. The sharper prediction intervals result directly from modeling context-specific errors, and the lower prediction error is the consequence of modeling errors in the form of a Corrector. In addition, we introduce a concept called *Attention Span* that gives precise control over the width of prediction intervals with a hyperparameter we introduce called Sequence Length ( $S_L$ ).

**Limitations and Future Directions** One of the limitations of our work is due to the use of explicit timestep IDs (i.e.,  $t=1,2,3$ , etc.) in the context state ( $c_t^n$ ). At inference time, the Corrector will only generate calibrated prediction intervals for a trajectory length equal to or less than what it was trained on. Moreover, the Corrector was trained on trajectories sampled at a fixed time interval  $\Delta t$  and does not generalize to other sampling intervals. A modified version of the Corrector could be made discretization-invariant. We hand-tuned the  $S_L$  hyperparameter, training multiple models to get a lower calibration error. An alternative way to get calibrated uncertainty with HOPCAST without tuning  $S_L$  hyperparameter is an interesting future work.

## References

- Andreas Auer, Martin Gauch, Daniel Klotz, and Sepp Hochreiter. Conformal prediction for time series with modern hopfield networks. *Advances in Neural Information Processing Systems*, 36:56027–56074, 2023.
- Michael Betancourt. A conceptual introduction to hamiltonian monte carlo. *arXiv preprint arXiv:1701.02434*, 2017.
- Muhammad Bilal Shahid and Cody Fleming. Toward robust car-following dynamics modeling via blackbox models: Methodology, analysis, and recommendations. *Transportation Research Record*, 2679(4):114–128, 2025.

- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *International conference on machine learning*, pp. 1613–1622. PMLR, 2015.
- Fred Brauer and Christopher Kribs. *Dynamical systems for biological modeling*. CRC press, 2016.
- Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the national academy of sciences*, 113(15):3932–3937, 2016.
- Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. *Advances in neural information processing systems*, 31, 2018.
- Adam Coates, Pieter Abbeel, and Andrew Y Ng. Learning for control from multiple demonstrations. In *Proceedings of the 25th international conference on Machine learning*, pp. 144–151, 2008.
- Bryan C Daniels and Ilya Nemenman. Efficient inference of parsimonious phenomenological models of cellular dynamics using s-systems and alternating regression. *PloS one*, 10(3):e0119821, 2015.
- Anthony Christopher Davison and David Victor Hinkley. *Bootstrap methods and their application*. Number 1. Cambridge university press, 1997.
- Stanislav Fort, Huiyi Hu, and Balaji Lakshminarayanan. Deep ensembles: A loss landscape perspective. *arXiv preprint arXiv:1912.02757*, 2019.
- Jan Gasthaus, Konstantinos Benidis, Yuyang Wang, Syama Sundar Rangapuram, David Salinas, Valentin Flunkert, and Tim Januschowski. Probabilistic forecasting with spline quantile function rnns. In *The 22nd international conference on artificial intelligence and statistics*, pp. 1901–1910. PMLR, 2019.
- Tilmann Gneiting and Adrian E Raftery. Weather forecasting with ensemble methods. *Science*, 310(5746):248–249, 2005.
- Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007.
- E. M. Izhikevich and R. FitzHugh. FitzHugh-Nagumo model. *Scholarpedia*, 1(9):1349, 2006. doi: 10.4249/scholarpedia.1349. revision #123664.
- Michael Janner, Qiyang Li, and Sergey Levine. Offline reinforcement learning as one big sequence modeling problem. *Advances in neural information processing systems*, 34:1273–1286, 2021.
- S Mohammad Khansari-Zadeh and Aude Billard. Learning stable nonlinear dynamical systems with gaussian mixture models. *IEEE Transactions on Robotics*, 27(5):943–957, 2011.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. URL <https://arxiv.org/abs/1412.6980>.
- Durk P Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameterization trick. *Advances in neural information processing systems*, 28, 2015.
- Jonathan Ko, Daniel J Klein, Dieter Fox, and Dirk Haehnel. Gp-ukf: Unscented kalman filters with gaussian process prediction and observation models. In *2007 IEEE/RSJ international conference on intelligent robots and systems*, pp. 1901–1907. IEEE, 2007.
- Volodymyr Kuleshov, Nathan Fenner, and Stefano Ermon. Accurate uncertainties for deep learning using calibrated regression. In *International conference on machine learning*, pp. 2796–2804. PMLR, 2018.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.



- Edward N Lorenz. Predictability: A problem partly solved. In *Proc. Seminar on predictability*, volume 1. Reading, 1996.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization, 2019. URL <https://arxiv.org/abs/1711.05101>.
- Cong Lu, Philip J Ball, Jack Parker-Holder, Michael A Osborne, and Stephen J Roberts. Revisiting design choices in offline model-based reinforcement learning. *arXiv preprint arXiv:2110.04135*, 2021.
- Thomas M Moerland, Joost Broekens, Aske Plaat, Catholijn M Jonker, et al. Model-based reinforcement learning: A survey. *Foundations and Trends® in Machine Learning*, 16(1):1–118, 2023.
- Alexandru Niculescu-Mizil and Rich Caruana. Predicting good probabilities with supervised learning. In *Proceedings of the 22nd international conference on Machine learning*, pp. 625–632, 2005.
- John Platt et al. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers*, 10(3):61–74, 1999.
- Adrian E. Raftery, Tilmann Gneiting, Fadoua Balabdaoui, and Michael Polakowski. Using bayesian model averaging to calibrate forecast ensembles. *Monthly Weather Review*, 133(5):1155–1174, 2005. doi: 10.1175/MWR2906.1. URL <https://journals.ametsoc.org/view/journals/mwre/133/5/mwr2906.1.xml>.
- Hubert Ramsauer, Bernhard Schöfl, Johannes Lehner, Philipp Seidl, Michael Widrich, Thomas Adler, Lukas Gruber, Markus Holzleitner, Milena Pavlović, Geir Kjetil Sandve, et al. Hopfield networks is all you need. *arXiv preprint arXiv:2008.02217*, 2020.
- Robert Robison, Tom Shafer, Victor Diloroto, Natalia M Alexandrov, Muhammad Bilal Shahid, and Cody H Fleming. Differential equation approximation using gradient-boosted quantile regression. In *AIAA SCITECH 2024 Forum*, pp. 0107, 2024.
- Muhammad Bilal Shahid, Robert Robison, Tom Shafer, Victor Diloroto, Natalia M Alexandrov, and Cody H Fleming. Uncertainty quantification using deep ensembles for decision making in cyber-physical systems. In *AIAA SCITECH 2024 Forum*, pp. 0108, 2024.
- Pierre NV Tu. *Dynamical systems: an introduction with applications in economics and biology*. Springer Science & Business Media, 2012.
- A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.
- Arun Venkatraman, Martial Hebert, and J Bagnell. Improving multi-step prediction of learned time series models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.
- Ying Wang, Bo Hu, Yuxue Zhao, Guofang Kuang, Yaling Zhao, Qingwei Liu, and Xiuli Zhu. Applications of system dynamics models in chronic disease prevention: a systematic review. *Preventing Chronic Disease*, 18:E103, 2021.
- Peter J Wangersky. Lotka-volterra population models. *Annual Review of Ecology and Systematics*, 9:189–218, 1978.
- Andrew G Wilson and Pavel Izmailov. Bayesian deep learning and a probabilistic perspective of generalization. *Advances in neural information processing systems*, 33:4697–4708, 2020.
- Yuan Yao, Lorenzo Rosasco, and Andrea Caponnetto. On early stopping in gradient descent learning. *Constructive approximation*, 26(2):289–315, 2007.
- Bianca Zadrozny and Charles Elkan. Transforming classifier scores into accurate multiclass probability estimates. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 694–699, 2002.
- Sheheryar Zaidi, Arber Zela, Thomas Elsken, Chris C Holmes, Frank Hutter, and Yee Teh. Neural ensemble search for uncertainty estimation and dataset shift. *Advances in Neural Information Processing Systems*, 34:7898–7911, 2021.

## A Baseline Methods

This section will discuss how we propagate uncertainty with Probabilistic Ensembles using three uncertainty propagation methods, i.e., Expectation, Moment Matching and Trajectory Sampling. These will be discussed with reference to the Lotka-Volterra (LV) system, whose state at time step  $(t + 1)$  is  $\mathbf{s}_{t+1} = (x_{t+1}, y_{t+1})$ . We will see how to generate  $z\%$  prediction intervals associated with  $\mathbf{s}_{t+1}$  ( $\text{PI}_{\mathbf{s}_{t+1}}^z$ ) at any timestep during autoregression with these approaches.

**Probabilistic Ensembles** To construct Probabilistic Ensembles, we train a population of fully connected feedforward models, i.e.,  $\Omega = \{f_i\}_i$ , where  $f_i$  represents the  $i^{\text{th}}$  model within the population. Each model is trained with Gaussian negative log-likelihood (NLL) loss as demonstrated in Lakshminarayanan et al. (2017). Each model predicts a multivariate normal (MVN) over the next state  $\mathbf{s}_{t+1}$  given the previous state of the system  $\mathbf{s}_t$ . Formally,  $\mathbf{s}_{t+1} \sim \text{MVN}(\boldsymbol{\mu}_{t+1}, \boldsymbol{\Sigma}_{t+1})$ . For the LV system,  $\boldsymbol{\mu}_{t+1} = (\mu_{t+1}^x, \mu_{t+1}^y)$  and  $\boldsymbol{\Sigma}_{t+1} = \text{diag}(\sigma_{t+1}^x, \sigma_{t+1}^y)$ .

**Uncertainty Propagation** At inference time, a sample of  $M$  models  $\{f_m\}_{m=1}^M$  can be taken from the population ( $\Omega$ ) to form a Probabilistic Ensemble. All three propagation approaches use a particle-based approach, where the core idea is to generate *multiple trajectories* via autoregression starting from the *same initial condition* (Chua et al., 2018). Let  $\mathbf{s}_t^{p,m}$  denote the state of the system associated with  $p^{\text{th}}$  particle from  $m^{\text{th}}$  model at  $t^{\text{th}}$  time step. Each model  $f_m$  receives  $P$  copies of the same initial condition  $\{\mathbf{s}_{t=0}^{p,m}\}_{p=1}^P$ , where  $P$  denotes the number of particles assigned to model  $f_m$ . The propagation of uncertainty starts with  $P * M$  copies of the same initial conditions across  $\{f_m\}_{m=1}^M$  models and results in  $P * M$  number of trajectories. The diversity in these trajectories will come from two sources, i.e.,  $M$  different models and their probabilistic predictions. These diverse trajectories will be used to derive prediction intervals. In contrast, HOPCAST used a set of errors (e.g.,  $E_{\bar{x}}$  &  $E_{\bar{y}}$ ) to derive prediction intervals.

The rest of the section discusses the algorithmic details of three propagation approaches and the construction of prediction intervals using the outputs of three algorithms.

### A.1 Expectation

The Expectation method always uses  $P = 1$  particle for each model  $f_m$ . This is because the predicted uncertainty ( $\boldsymbol{\Sigma}_{t+1}^{p,m}$ ) is ignored, and the predicted mean ( $\boldsymbol{\mu}_{t+1}^{p,m}$ ) is propagated to the next time step. Therefore, assigning more than one particle to a model  $f_m$  will not generate diverse trajectories from that model. Even though this approach ignores predicted uncertainty ( $\boldsymbol{\Sigma}_{t+1}^{p,m}$ ) during propagation, it will be used to generate prediction intervals as explained in section A.3. The overall procedure of this approach is in Algorithm 2.

---

#### Algorithm 2 Expectation

---

- 1: **Input:** Initial states  $\{\mathbf{s}_{t=0}^{1,m}\}_{m=1}^M$ , models  $f_1, \dots, f_M$ , timesteps  $T$
  - 2: **for**  $t = 0$  **to**  $T - 1$  **do**
  - 3:   **for**  $m = 1$  **to**  $M$  **do**
  - 4:     Predict  $\boldsymbol{\mu}_{t+1}^{1,m}, \boldsymbol{\Sigma}_{t+1}^{1,m} \leftarrow f_m(\mathbf{s}_t^{1,m})$
  - 5:      $\mathbf{s}_{t+1}^{1,m} = \boldsymbol{\mu}_{t+1}^{1,m} \triangleright$  *Propagate mean*
  - 6:   **end for**
  - 7:    $\{\mathbf{s}_{t+1}^{1,m}\}_{m=1}^M \triangleright$  *inputs to the next time step*
  - 8: **end for**
  - 9: **return**  $\{\mathbf{s}_{1:T}^{1,m}\}_{m=1}^M \triangleright$  *the  $1 * M$  number of trajectories of length  $T$*
- 

### A.2 Trajectory Sampling

This method uses  $P$  particles for each model  $f_m$  since it considers predicted uncertainty ( $\boldsymbol{\Sigma}_{t+1}^{p,m}$ ) while propagating uncertainty to the next step. The overall procedure of this algorithm is summarized in Algorithm 3.

**Algorithm 3** Trajectory Sampling

---

```

1: Input: Initial states  $\{\{\mathbf{s}_{t=0}^{p,m}\}_{p=1}^P\}_{m=1}^M$ , models  $f_1, \dots, f_M$ , timesteps  $T$ , particles  $P$ 
2: for  $t = 0$  to  $T - 1$  do
3:   for  $m = 1$  to  $M$  do
4:     for  $p = 1$  to  $P$  do
5:       Predict  $\boldsymbol{\mu}_{t+1}^{p,m}, \boldsymbol{\Sigma}_{t+1}^{p,m} \leftarrow f_m(\mathbf{s}_t^{p,m})$ 
6:       Sample  $\mathbf{s}_{t+1}^{p,m} \sim \text{MVN}(\boldsymbol{\mu}_{t+1}^{p,m}, \boldsymbol{\Sigma}_{t+1}^{p,m})$   $\triangleright$  considers predicted uncertainty
7:     end for
8:   end for
9:    $\{\{\mathbf{s}_{t+1}^{p,m}\}_{p=1}^P\}_{m=1}^M$   $\triangleright$  inputs to the next time step
10: end for
11: return  $\{\{\mathbf{s}_{1:T}^{p,m}\}_{p=1}^P\}_{m=1}^M$   $\triangleright$  the  $P * M$  number of trajectories of length  $T$ 

```

---

**A.3 Moment Matching**

This method uses  $P$  particles for each model  $m$  since it considers predicted uncertainty ( $\boldsymbol{\Sigma}_t^{p,m}$ ) while propagating uncertainty to the next step. This is different from Trajectory Sampling in that it fits a Gaussian distribution to  $P * M$  predictions  $\{\{\mathbf{s}_{t+1}^{p,m}\}_{p=1}^P\}_{m=1}^M$  at  $t^{\text{th}}$  timestep, that is,  $\text{MVN}(\boldsymbol{\mu}_{t+1}, \boldsymbol{\Sigma}_{t+1})$ . We assume independence between all  $P * M$  predictions. A set of  $P * M$  samples  $\{\{\mathbf{s}_{t+1}^{p,m}\}_{p=1}^P\}_{m=1}^M$  is taken from that distribution which becomes input at the next time step. The overall procedure is summarized in Algorithm 4.

**Algorithm 4** Moment Matching

---

```

1: Input: Initial states  $\{\{\mathbf{s}_0^{p,m}\}_{p=1}^P\}_{m=1}^M$ , models  $f_1, \dots, f_M$ , timesteps  $T$ , particles  $P$ 
2: for  $t = 0$  to  $T - 1$  do
3:   for  $m = 1$  to  $M$  do
4:     for  $p = 1$  to  $P$  do
5:       Predict  $\boldsymbol{\mu}_{t+1}^{p,m}, \boldsymbol{\Sigma}_{t+1}^{p,m} \leftarrow f_m(\mathbf{s}_t^{p,m})$ 
6:       Sample  $\mathbf{s}_{t+1}^{p,m} \sim \text{MVN}(\boldsymbol{\mu}_{t+1}^{p,m}, \boldsymbol{\Sigma}_{t+1}^{p,m})$ 
7:     end for
8:   end for
9:    $\{\{\mathbf{s}_{t+1}^{p,m}\}_{p=1}^P\}_{m=1}^M$   $\triangleright$  fit a Gaussian distribution
10:   $\boldsymbol{\mu}_{t+1} \leftarrow \frac{1}{P \cdot M} \sum_{m=1}^M \sum_{p=1}^P \mathbf{s}_{t+1}^{p,m}$   $\triangleright$  Mean
11:   $\boldsymbol{\Sigma}_{t+1} \leftarrow \frac{1}{P \cdot M} \sum_{m=1}^M \sum_{p=1}^P (\mathbf{s}_{t+1}^{p,m} - \boldsymbol{\mu}_{t+1})(\mathbf{s}_{t+1}^{p,m} - \boldsymbol{\mu}_{t+1})^T$   $\triangleright$  Variance
12:   $\{\{\mathbf{s}_{t+1}^{p,m}\}_{p=1}^P\}_{m=1}^M \sim \text{MVN}(\boldsymbol{\mu}_{t+1}, \boldsymbol{\Sigma}_{t+1})$   $\triangleright$  sample inputs for the next time step
13: end for
14: return  $\{\{\mathbf{s}_{1:T}^{p,m}\}_{p=1}^P\}_{m=1}^M$   $\triangleright$  the  $P * M$  number of trajectories of length  $T$ 

```

---

**Prediction intervals using uncertainty propagation** All three approaches output  $P * M$  number of trajectories of length  $T$   $\{\{\mathbf{s}_{1:T}^{p,m}\}_{p=1}^P\}_{m=1}^M$ . At any time step  $t$  during propagation, the  $(\text{PI}_{\mathbf{s}_{t+1}}^z)$  are generated using  $\{\{\mathbf{s}_t^{p,m}\}_{p=1}^P\}_{m=1}^M$ . Specifically, we construct a  $\text{MVN}_{\text{net}}$  whose parameters are:  $\boldsymbol{\mu}_{\text{net}} = \frac{1}{P * M} \sum_{m=1}^M \sum_{p=1}^P \mathbf{s}_t^{p,m}$  and  $\boldsymbol{\Sigma}_{\text{net}} = \frac{1}{P * M} \sum_{m=1}^M \sum_{p=1}^P [\boldsymbol{\Sigma}_t^{p,m} + (\mathbf{s}_t^{p,m} - \boldsymbol{\mu}_{\text{net}})(\mathbf{s}_t^{p,m} - \boldsymbol{\mu}_{\text{net}})^T]$ . For LV system, the  $\boldsymbol{\mu}_{\text{net}} = (\mu_{\text{net}}^x, \mu_{\text{net}}^y)$  and  $\boldsymbol{\Sigma}_{\text{net}} = \text{diag}(\sigma_{\text{net}}^x, \sigma_{\text{net}}^y)$ . The prediction intervals for  $x$  and  $y$  are derived as follows:

$$\text{PI}_{\mathbf{x}_{t+1}}^z = \mu_{\text{net}}^x \pm \Phi^{-1}(1 - \alpha) \sigma_{\text{net}}^x \quad (11)$$

$$\text{PI}_{\mathbf{y}_{t+1}}^z = \mu_{\text{net}}^y \pm \Phi^{-1}(1 - \alpha) \sigma_{\text{net}}^y \quad (12)$$

where  $\alpha = \frac{1-z}{2}$  for  $z\%$  prediction intervals. The  $\Phi^{-1}$  denotes the inverse cumulative distribution function of the standard normal distribution. In contrast, equations 3 & 4 were used to generate prediction intervals with HOPCAST.

## B Dynamical Systems

There are five dynamical systems studied in this paper. One of those is Lotka-Volterra (LV) equations, which are already discussed in section 3 along with its closed form, parameters, and ranges of initial conditions. Here, we repeat the description of the LV system and then describe the rest of the systems. We add zero-mean additive Gaussian noise to each differential equation post-integration as discussed in section 3, with the variance equal to the standard deviation of the variable modelled by the differential equation scaled by a factor  $\sigma$ .

### B.1 Lotka-Volterra (Wangersky, 1978)

$$\frac{dx}{dt} = \alpha x - \beta xy \quad (13)$$

$$\frac{dy}{dt} = \delta xy - \gamma y \quad (14)$$

**Parameters:**  $\alpha = 1.1$ ;  $\beta = 0.4$ ;  $\gamma = 0.4$ ;  $\delta = 0.1$

**Initial Condition Ranges:**  $x \in [5, 20]$ ;  $y \in [5, 10]$

### B.2 Lorenz (Brunton et al., 2016)

$$\frac{dx}{dt} = \sigma(y - x) \quad (15)$$

$$\frac{dy}{dt} = x(\rho - z) - y \quad (16)$$

$$\frac{dz}{dt} = xy - \beta z \quad (17)$$

**Parameters:**  $\sigma = 10$ ;  $\rho = 28$ ;  $\beta = \frac{8}{3}$

**Initial Condition Ranges:**  $x \in [-20, 20]$ ;  $y \in [-20, 20]$ ;  $z \in [0, 50]$

### B.3 FitzHugh-Nagumo (FHN) (Izhikevich & FitzHugh, 2006)

$$\frac{dv}{dt} = v - \frac{v^3}{3} - w + I \quad (18)$$

$$\frac{dw}{dt} = \epsilon(v + a - bw) \quad (19)$$

**Parameters:**  $a = 0.7$ ;  $b = 0.8$ ;  $\epsilon = 0.08$ ;  $I = 0.5$

**Initial Condition Ranges:**  $v \in [-1.5, 1.5]$ ;  $w \in [-1.5, 1.5]$

### B.4 Lorenz95 (Lorenz, 1996)

$$\frac{dX_1}{dt} = (X_2 - X_5)X_4 - X_1 + F \quad (20)$$

$$\frac{dX_2}{dt} = (X_3 - X_1)X_5 - X_2 + F \quad (21)$$

$$\frac{dX_3}{dt} = (X_4 - X_2)X_1 - X_3 + F \quad (22)$$

$$\frac{dX_4}{dt} = (X_5 - X_3)X_2 - X_4 + F \quad (23)$$

$$\frac{dX_5}{dt} = (X_1 - X_4)X_3 - X_5 + F \quad (24)$$

**Parameters:**  $F = 8$

**Initial Condition Ranges:**  $X_i \in [-10.5, 10.5]$  where  $i \in \{1, 2, 3, 4, 5\}$

## B.5 Glycolytic Oscillator (Daniels & Nemenman, 2015)

$$\frac{dS_1}{dt} = J_0 - \frac{k_1 S_1 S_6}{1 + (S_6/K_1)^q} \quad (25)$$

$$\frac{dS_2}{dt} = 2 \frac{k_1 S_1 S_6}{1 + (S_6/K_1)^q} - k_2 S_2 (N - S_5) - k_6 S_2 S_5 \quad (26)$$

$$\frac{dS_3}{dt} = k_2 S_2 (N - S_5) - k_3 S_3 (A - S_6) \quad (27)$$

$$\frac{dS_4}{dt} = k_3 S_3 (A - S_6) - k_4 S_4 S_5 - \kappa (S_4 - S_7) \quad (28)$$

$$\frac{dS_5}{dt} = k_2 S_2 (N - S_5) - k_4 S_4 S_5 - k_6 S_2 S_5 \quad (29)$$

$$\frac{dS_6}{dt} = -2 \frac{k_1 S_1 S_6}{1 + (S_6/K_1)^q} + 2k_3 S_3 (A - S_6) - k_5 S_6 \quad (30)$$

$$\frac{dS_7}{dt} = \psi \kappa (S_4 - S_7) - k S_7 \quad (31)$$

**Parameters:**  $J_0 = 2.5; k_1 = 100; k_2 = 6; k_3 = 16; k_4 = 100; k_5 = 1.28; k_6 = 12; k = 1.8; \kappa = 13; q = 4; K_1 = 0.52; \psi = 0.1; N = 1; A = 4$

**Initial Condition Ranges:**  $S_1 \in [0.15, 1.60]; S_2 \in [0.19, 2.16]; S_3 \in [0.04, 0.20]; S_4 \in [0.10, 0.35]; S_5 \in [0.08, 0.30]; S_6 \in [0.14, 2.67]; S_7 \in [0.05, 0.10]$

## C Hyperparameters

In Table 3, we provide the Sequence Length ( $S_L$ ) for each output of the Predictor and a number of models in Probabilistic Ensembles for each experimental setting in Table 2.

## D Ablation Studies

### D.1 Encoder Architecture

As discussed in section 4, each correction model utilizes an Encoder with MHN. The architecture of the Encoder (e.g.,  $m_x$ ) is a fully connected feedforward network with one layer and 100 neurons for all experiments. Here, we show the impact of adding more layers of fully connected neurons to the Encoder on the proposed metrics. Table 4 shows that adding more layers doesn't significantly impact the MSE and CE. However, the MSE and CE increase considerably when the Encoder is removed from the correction models.

### D.2 Initial condition and context state ( $c_t^n$ )

As discussed in section 1, our context state ( $c_t^n$ ) comprises the initial condition, the current state of the system, and the timestep ID. Here, we show the results on two systems without the initial condition a part of the context state in Table 5. The MSE increases significantly for both cases, while the CE exhibits minor changes. This shows the significance of making the initial condition a part of the context state.

### D.3 Size of Association Memory

In section 4.3, we discussed the inference with the correction model ( $M_x$ ). A set of  $K$  keys ( $\mathbf{K}_{mem}$ ) with the corresponding values ( $\mathbf{V}_{mem}^x$ ) is randomly sampled from  $\mathcal{S}_K^{\text{Train}}$  and  $\mathcal{S}_{V_x}^{\text{Train}}$ , respectively and loaded into the Association memory. Here, the impact of the size of Association memory on the proposed metrics will

Table 3: Sequence lengths ( $S_L$ ) for all outputs of Predictor and number of models in Probabilistic ensemble across various dynamical systems and noise levels

Model		HopCast	Probabilistic Ensemble		
			Expectation	Moment Matching	Trajectory Sampling
Hyperparameter		Sequence lengths	Models	Models	Models
Lotka Volterra	$\sigma = 0.05$	30,30,50,60	4	5	3
	$\sigma = 0.1$	70,30,70,40	5	5	3
	$\sigma = 0.3$	70,20,60,70	6	5	4
Lorenz	$\sigma = 0.05$	30,35,35,35,50,50,100	6	8	3
	$\sigma = 0.1$	500,500,500,300,300,300	7	7	4
	$\sigma = 0.3$	25,25,25,500,500,500	8	7	3
FHN	$\sigma = 0.05$	35,55	5	5	3
	$\sigma = 0.1$	15,15	4	5	3
	$\sigma = 0.3$	15,15	5	4	3
Lorenz95	$\sigma = 0.05$	2200 for all	6	7	3
	$\sigma = 0.1$	2000 for all	6	5	4
	$\sigma = 0.3$	2000 for all	7	6	4
Glycolytic Oscillator	$\sigma = 0.05$	35,50,35,100,35,50,30	8	6	3
	$\sigma = 0.1$	50,50,35,35,800,35,100	6	7	3
	$\sigma = 0.3$	50 for all	7	8	3

Table 4: Effect of different Encoding network architectures on the proposed metrics. The results without the Encoder are also included. The  $FC(100)_1$  means the fully connected feedforward model with one layer of 100 neurons. The results are reported over three runs of the same experiment at random seeds.

Dynamical System	Lorenz			Lotka Volterra		
Noise Level	$\sigma = 0.3$			$\sigma = 0.3$		
Metrics	MSE	PI-Width	CE	MSE	PI-Width	CE
Without Encoder	$1907.07 \pm 9.33$	$38.69 \pm 0.42$	$0.06 \pm 0.005$	$16.12 \pm 0.36$	$3.37 \pm 0.23$	$0.025 \pm 0.012$
$FC(100)_1$	$1681 \pm 46.36$	$40.44 \pm 0.91$	$0.019 \pm 0.004$	$9.69 \pm 0.23$	$3.19 \pm 0.12$	$0.0051 \pm 0.004$
$FC(100)_2$	$1717.362 \pm 30.3$	$41.84 \pm 0.68$	$0.014 \pm 0.002$	$9.43 \pm 0.23$	$3.08 \pm 0.09$	$0.006 \pm 0.002$
$FC(100)_3$	$1710.18 \pm 46.19$	$42.45 \pm 0.55$	$0.018 \pm 0.002$	$9.37 \pm 0.007$	$2.98 \pm 0.07$	$0.0035 \pm 0.001$

be discussed. Table 6 shows the variation in proposed metrics as the number of keys in memory varies from 10 to 2000. For all systems, the MSE and CE improve significantly as the keys increase from 10 to 50. The MSE and CE that showed a considerable change as the number of keys increases are shown in bold. As the keys increase beyond 50, not all systems' MSE and CE showed a considerable increase. The MSE and CE show insignificant changes as the number of keys increases beyond 100 until 2000. Thus, we fix the number of keys to 2000 for our final evaluation of results.

Table 5: Comparison of proposed metrics with and without initial condition included in the context state. The results are averaged over three runs.

Dynamical System	Glycolytic Oscillator			Lotka Volterra		
Noise Level	$\sigma = 0.1$			$\sigma = 0.1$		
Metrics	MSE	PI-Width	CE	MSE	PI-Width	CE
<b>With IC</b>	$0.072 \pm 0.003$	$0.25 \pm 0.007$	$0.018 \pm 0.006$	$6.10 \pm 0.20$	$1.88 \pm 0.05$	$0.0070 \pm 0.002$
<b>Without IC</b>	$0.176 \pm 0.003$	$0.416 \pm 0.003$	$0.027 \pm 0.008$	$43.90 \pm 1.09$	$5.17 \pm 0.14$	$0.011 \pm 0.002$

Table 6: The results on proposed metrics with varying numbers of keys in memory. The MSE and CE that showed a considerable improvement with the increase in the number of keys in memory are shown in bold. The results are averaged over three runs.

Keys in Memory	Metric	Lorenz	Glycolytic Oscillator	Lotka Volterra	FHN	Lorenz95
		$\sigma = 0.1$	$\sigma = 0.3$	$\sigma = 0.05$	$\sigma = 0.05$	$\sigma = 0.3$
<b>10</b>	MSE	<b><math>1765.48 \pm 170.73</math></b>	<b><math>0.15 \pm 0.009</math></b>	<b><math>26.86 \pm 4.80</math></b>	<b><math>0.15 \pm 0.03</math></b>	<b><math>24.81 \pm 3.11</math></b>
	PI-Width	$34.59 \pm 5.69$	$0.37 \pm 0.006$	$3.61 \pm 1.06$	$0.24 \pm 0.03$	$6.59 \pm 0.69$
	CE	<b><math>0.11 \pm 0.03</math></b>	<b><math>0.080 \pm 0.009</math></b>	<b><math>0.10 \pm 0.04</math></b>	<b><math>0.24 \pm 0.19</math></b>	<b><math>0.039 \pm 0.05</math></b>
<b>50</b>	MSE	<b><math>1419.29 \pm 103.78</math></b>	$0.13 \pm 0.001$	<b><math>9.66 \pm 1.51</math></b>	$0.072 \pm 0.001$	$18.10 \pm 1.75$
	PI-Width	$37.64 \pm 0.55$	$0.36 \pm 0.008$	$2.90 \pm 0.03$	$0.21 \pm 0.012$	$5.98 \pm 0.11$
	CE	<b><math>0.025 \pm 0.009</math></b>	$0.012 \pm 0.002$	<b><math>0.038 \pm 0.003</math></b>	$0.12 \pm 0.02$	<b><math>0.016 \pm 0.003</math></b>
<b>100</b>	MSE	$1285.90 \pm 30.9$	$0.12 \pm 0.002$	$7.76 \pm 0.67$	$0.076 \pm 0.001$	$18.57 \pm 0.51$
	PI-Width	$36.83 \pm 0.39$	$0.36 \pm 0.015$	$2.99 \pm 0.06$	$0.21 \pm 0.01$	$6.09 \pm 0.15$
	CE	$0.015 \pm 0.007$	$0.018 \pm 0.007$	$0.026 \pm 0.006$	<b><math>0.111 \pm 0.003</math></b>	<b><math>0.006 \pm 0.005</math></b>
<b>500</b>	MSE	$1249 \pm 11.03$	$0.12 \pm 0.001$	$6.89 \pm 0.30$	$0.076 \pm 0.002$	$17.06 \pm 0.86$
	PI-Width	$38.19 \pm 1.39$	$0.36 \pm 0.001$	$2.85 \pm 0.08$	$0.21 \pm 0.006$	$6.02 \pm 0.10$
	CE	$0.020 \pm 0.004$	$0.013 \pm 0.003$	$0.022 \pm 0.003$	$0.093 \pm 0.03$	$0.0021 \pm 0.003$
<b>1000</b>	MSE	$1241.72 \pm 7.55$	$0.12 \pm 0.003$	$7.08 \pm 0.01$	$0.077 \pm 0.001$	$16.91 \pm 0.74$
	PI-Width	$38.37 \pm 1.91$	$0.37 \pm 0.005$	$2.91 \pm 0.093$	$0.21 \pm 0.001$	$6.13 \pm 0.09$
	CE	$0.018 \pm 0.001$	$0.017 \pm 0.004$	$0.021 \pm 0.004$	$0.092 \pm 0.01$	$0.0020 \pm 0.002$
<b>2000</b>	MSE	$1248.17 \pm 15$	$0.11 \pm 0.001$	$6.87 \pm 0.32$	$0.076 \pm 0.002$	$16.44 \pm 0.66$
	PI-Width	$38.66 \pm 0.38$	$0.37 \pm 0.008$	$2.93 \pm 0.17$	$0.20 \pm 0.005$	$6.08 \pm 0.06$
	CE	$0.019 \pm 0.03$	$0.016 \pm 0.003$	$0.022 \pm 0.004$	$0.091 \pm 0.004$	$0.0015 \pm 0.001$

#### D.4 Embedding dimension

As discussed in section 4.2, the Encoder outputs a  $d$ -dimensional embedding given context state ( $c_t^n$ ), where  $d = 4$ . Here, we want to show that the performance of correction models is insensitive to the embedding dimension provided that the  $S_L$  is retuned accordingly. In Table 7, the results on metrics and  $S_L$  for LV at  $\sigma = 0.1$  in the first row are copied from Tables 2 & 3, respectively. Table 7 shows that the CE shows a considerable increase from 0.0070 to 0.065 as we change the  $d$  to 5 from 4, while the MSE shows minor changes. The last row shows the results at  $d = 5$  when the  $S_L$  is retuned, where the CE (i.e., 0.0098) is close to what we had earlier at  $d = 4$  (i.e., 0.0070).

Table 7: The impact of change in embedding dimension on the proposed metrics. The last row shows the results when  $S_L$  is retuned for the new embedding dimension ( $d = 5$ ). The results are reported over three runs.

Embedding Dimension ( $d$ )	$S_L$ of all outputs	Lotka Volterra		
		$\sigma = 0.1$		
		MSE	PI-Width	CE
$d = 4$	[70,30,70,40]	$6.10 \pm 0.20$	$1.88 \pm 0.05$	$0.0070 \pm 0.002$
$d = 5$	[70,30,70,40]	$6.75 \pm 0.065$	$1.94 \pm 0.02$	$0.065 \pm 0.005$
$d = 5$	[5,5,70,5]	$6.98 \pm 0.115$	$1.69 \pm 0.01$	$0.0098 \pm 0.003$

### D.5 Impact of shuffled data

As discussed in section 4.1, the dataset is shuffled before forming queries, keys, and values for training. Here, we show that this shuffling brings two major benefits. One, the MSE and CE show significant improvement. Second, CE and MSE show changes with respect to the number of keys in memory that are consistent across different systems.

Table 8 shows the results on the proposed metrics for FHN and LV with and without shuffling the data. The MSE and CE increase considerably for both systems when the data is not shuffled, irrespective of the number of keys in memory. For LV, the MSE starts to decrease and CE starts to increase as we reduce the number of keys in memory from 2000 to 10. For FHN, in contrast, MSE and CE begin to grow as the number of keys goes down to 10 from 2000. On the other hand, the MSE and CE remain largely unchanged for both systems as we reduce the number of keys from 2000 to 500 when the data is shuffled. For keys less than 500, the MSE and CE show considerable change for both systems.

When the data is shuffled, the changes in MSE and CE as we reduce the number of keys show a generalizable pattern for both systems. Hence, the guidelines regarding the number of keys in the memory at the inference time are generalizable.

### D.6 Number of models in Probabilistic Ensembles

As discussed in section 7, the Probabilistic Ensembles tend to saturate uncertainty after a certain number of models in the ensemble. Here, we demonstrate that for all propagation approaches via calibration curves. As shown in Fig. 8(a) for Expectation, the uncertainty goes up from overconfidence to underconfidence as we increase the number of models in the Probabilistic Ensemble. The last three increments in the number of models (13 to 15) don't impact uncertainty much. For Moment Matching (Fig. 8(b)), the uncertainty increases with the number of models and saturates later while the model is still overconfident. For Trajectory Sampling (Fig. 8(c)), increasing the number of models leads to a marginal change in uncertainty and saturates after four models.

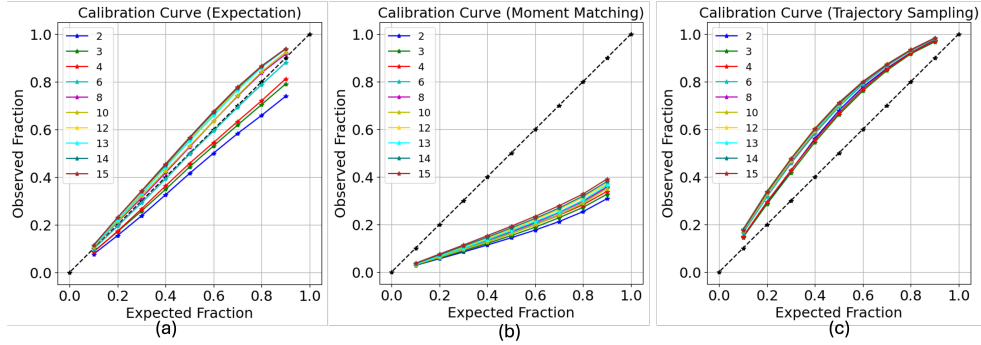
### D.7 Sequence length ( $S_L$ ) & Calibration

The section 7 discusses the impact of varying  $S_L$  on the model's confidence for the y-output of Lorenz. Here, we show the same for all outputs of the two systems, i.e., Lorenz and Glycolytic Oscillator. The calibrated confidence of both systems with their optimal  $S_L$  is shown in the middle plot of Fig. 9a & 9b. To show underconfidence and overconfidence, we pick a small  $S_L = 10$  and a large  $S_L = 1000$ . At small  $S_L = 10$ , we expect all outputs of both systems to be overconfident as shown in the leftmost plots of Fig. 9a & 9b. At large  $S_L = 1000$ , we expect all outputs of both systems to be underconfident as shown in the rightmost plots




Table 8: Performance on proposed metrics at different numbers of keys in memory, with and without data shuffling. The results are averaged over three runs.


Keys in Memory	Metric	Lotka-Volterra		FHN	
		$\sigma = 0.05$		$\sigma = 0.05$	
		Shuffled Data	Non-Shuffled Data	Shuffled Data	Non-Shuffled Data
10	MSE	$26.86 \pm 4.80$	$33.09 \pm 0.009$	$0.15 \pm 0.03$	$0.16 \pm 0.01$
	PI-Width	$3.61 \pm 1.06$	$0.54 \pm 0.001$	$0.24 \pm 0.03$	$0.17 \pm 0.005$
	CE	$0.10 \pm 0.04$	$0.98 \pm 0.007$	$0.24 \pm 0.19$	$1.15 \pm 0.05$
50	MSE	$9.66 \pm 1.51$	$32.91 \pm 0.03$	$0.072 \pm 0.001$	$0.12 \pm 0.02$
	PI-Width	$2.90 \pm 0.03$	$0.49 \pm 5.42$	$0.21 \pm 0.012$	$0.18 \pm 0.01$
	CE	$0.038 \pm 0.03$	$0.74 \pm 0.005$	$0.12 \pm 0.02$	$0.38 \pm 0.06$
100	MSE	$7.76 \pm 0.67$	$32.63 \pm 0.09$	$0.076 \pm 0.001$	$0.10 \pm 0.02$
	PI-Width	$2.99 \pm 0.06$	$0.51 \pm 0.001$	$0.21 \pm 0.01$	$0.16 \pm 0.01$
	CE	$0.026 \pm 0.006$	$0.72 \pm 0.001$	$0.111 \pm 0.003$	$0.46 \pm 0.08$
500	MSE	$6.89 \pm 0.30$	$42.14 \pm 22.03$	$0.076 \pm 0.002$	$0.12 \pm 0.03$
	PI-Width	$2.85 \pm 0.08$	$3.32 \pm 0.09$	$0.21 \pm 0.006$	$0.18 \pm 0.009$
	CE	$0.022 \pm 0.003$	$0.12 \pm 0.06$	$0.093 \pm 0.03$	$0.41 \pm 0.09$
1000	MSE	$7.08 \pm 0.01$	$34.80 \pm 16.27$	$0.077 \pm 0.001$	$0.092 \pm 0.01$
	PI-Width	$2.91 \pm 0.093$	$3.36 \pm 0.07$	$0.21 \pm 0.001$	$0.27 \pm 0.006$
	CE	$0.021 \pm 0.004$	$0.083 \pm 0.053$	$0.092 \pm 0.01$	$0.12 \pm 0.12$
2000	MSE	$6.87 \pm 0.32$	$60.84 \pm 11.17$	$0.076 \pm 0.002$	$0.091 \pm 0.013$
	PI-Width	$2.93 \pm 0.17$	$4.77 \pm 0.10$	$0.20 \pm 0.005$	$0.37 \pm 0.005$
	CE	$0.022 \pm 0.004$	$0.063 \pm 0.015$	$0.091 \pm 0.004$	$0.16 \pm 0.042$

Figure 8: The calibration curves of three propagation approaches for x-output of Lorenz ( $\sigma = 0.05$ ) to delineate the saturation of uncertainty as the number of models in an ensemble goes up from 2 to 15.

of Fig. 9a & 9b. It shows that the concept of the increase in uncertainty with the increase in  $S_L$  generalizes well across different outputs of different systems.



(a) Lorenz ( $\sigma = 0.1$ ). The leftmost plot shows the overconfidence of all six outputs of Lorenz at small  $S_L$ . The middle plot shows the calibrated confidence for all outputs at optimal  $S_L$ . The rightmost plot shows the underconfidence of all outputs at large  $S_L$ .



(b) Glycolytic Oscillator ( $\sigma = 0.3$ ). The leftmost plot shows the overconfidence of all seven outputs at small  $S_L$ . The middle plot shows the calibrated confidence for all outputs at optimal  $S_L$ . The rightmost plot shows the underconfidence of all outputs at large  $S_L$ .

Figure 9: Calibration curves of HOPCAST for two dynamical systems to demonstrate the impact of  $S_L$  on calibration. At the bottom of each plot,  $S_L$  of all outputs are shown in ascending order such that the first value corresponds to the first output, second to the second output, and so on.

## E Implementation Details

This section contains details regarding the implementation of baselines and HOPCAST. All experiments were run on NVIDIA A100-SXM4-80GB. PyTorch is used to implement baselines and HOPCAST.

### E.1 Baselines Implementation

We train a population of 15 models for the Probabilistic Ensembles. Regarding architecture, two layers of a fully connected feedforward model with 400 neurons each were used for LV and FHN, and three layers with 400 neurons for Lorenz, Lorenz95, and Glycolytic Oscillator. The train/test datasets were prepared following an 80/20 split. The batch size, learning rate, optimizer, and epochs were kept the same for all experiments, and are 128, 0.001, Adam (Kingma & Ba, 2017), and 1000, respectively. The early stopping was used as a criterion to train the Probabilistic Ensembles (Yao et al., 2007). The number of particles  $P$  for each model in Probabilistic Ensembles was 1 for Expectation, and 20 for Moment Matching & Trajectory Sampling (Chua et al., 2018).

### E.2 HopCast Implementation

The Encoder was a fully connected feedforward model with one layer and 100 neurons. The section 7 contains details about tuning an important hyperparameter, i.e.,  $S_L$ , of HOPCAST. Table 3 has  $S_L$  for each

output of all systems at various noise scaling factors  $\sigma$ . The learning rate of 0.001 and the optimizer **AdamW** (Loshchilov & Hutter, 2019) were kept the same for all experiments. The batch size and epochs were different for each experiment, and are provided in the form of `yml` files as a supplementary material along with other hyperparameters for each system and noise scaling factor ( $\sigma$ ).

The deterministic Predictor was a fully connected feedforward model of two layers with 400 neurons each for LV and FHN, and three layers with 400 neurons for Lorenz, Lorenz95, and Glycolytic Oscillator.