



HAL
open science

Evaluation of Machine Learning Techniques for Forecast Uncertainty Quantification

Maximiliano Sacco, Juan Ruiz, Manuel Pulido, Pierre Tandeo

► **To cite this version:**

Maximiliano Sacco, Juan Ruiz, Manuel Pulido, Pierre Tandeo. Evaluation of Machine Learning Techniques for Forecast Uncertainty Quantification. Quarterly Journal of the Royal Meteorological Society, In press, 10.1002/qj.4362 . hal-03685523

HAL Id: hal-03685523

<https://imt-atlantique.hal.science/hal-03685523>

Submitted on 2 Jun 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

RESEARCH ARTICLE

Evaluation of Machine Learning Techniques for Forecast Uncertainty Quantification

Maximiliano A. Sacco^{1,2} | Juan J. Ruiz^{2,4,5} | Manuel Pulido^{5,6} | Pierre Tandeo³

¹Servicio Meteorológico Nacional, Buenos Aires, Argentina

²Departamento de Ciencias de la Atmósfera y los Océanos, Facultad de Ciencias Exactas y Naturales, Universidad de Buenos Aires, Buenos Aires, Argentina

³Lab-STICC, UMR CNRS 6285, IMT Atlantique, Plouzané, France

⁴Centro de Investigaciones del Mar y la Atmósfera, Facultad de Ciencias Exactas y Naturales, Universidad de Buenos Aires, CONICET-UBA, Buenos Aires, Argentina

⁵UMI-IFAECI (CNRS-CONICET-UBA), Buenos Aires, Argentina

⁶Departamento de Física - Facultad Ciencias Exactas y Naturales y Agrimensura, Universidad Nacional del Nordeste, Corrientes, Argentina

Correspondence

msacco@smn.gob.ar
Email: msacco@smn.gob.ar

Funding information

Producing an accurate weather forecast and a reliable quantification of its uncertainty is an open scientific challenge. Ensemble forecasting is, so far, the most successful approach to produce relevant forecasts along with an estimation of their uncertainty. The main limitations of ensemble forecasting are the high computational cost and the difficulty to capture and quantify different sources of uncertainty, particularly those associated with model errors. In this work proof-of-concept model experiments are conducted to examine the performance of ANNs trained to predict a corrected state of the system and the state uncertainty using only a single deterministic forecast as input. We compare different training strategies: one based on a direct training using the mean and spread of an ensemble forecast as target, the other ones rely on an indirect training strategy using a deterministic forecast as target in which the uncertainty is implicitly learned from the data. For the last approach two alternative loss functions are proposed and evaluated, one based on the data observation likelihood and the other one based on a local estimation of the error. The performance of the networks is examined at different lead times and in scenarios with and without model errors. Experiments using the Lorenz'96 model show that the ANNs are able to emulate some of the properties of en-

semble forecasts like the filtering of the most unpredictable modes and a state-dependent quantification of the forecast uncertainty. Moreover, ANNs provide a reliable estimation of the forecast uncertainty in the presence of model error.

KEYWORDS

neural networks, chaotic dynamic models, uncertainty quantification, observation likelihood loss function, forecast

1 | INTRODUCTION

Uncertainty in weather forecasting is attributed to errors in the initial conditions and model errors. The first ones are associated with imperfect knowledge of the initial state of the system. The second ones are produced by truncation errors and an inaccurate representation of small-scale processes in numerical forecast models. While the model error can be reduced by a higher horizontal resolution and a better physical parametrization (i.e., improving the model), initial condition errors can be reduced by improving observation networks and data assimilation systems (Carrassi et al., 2018).

Currently, the uncertainty associated with weather forecasts is usually quantified by ensemble forecasts. To produce an ensemble forecast, several numerical simulations are performed using different initial conditions appropriately selected (e.g. Kalnay, 2003) and using different approaches such as parametric error representation and stochastic processes to represent the imperfection in the model (e.g. Hagedorn et al., 2012). One of the main limitations of ensemble forecasting systems lies in its high computational cost and in the complexity of designing them to adequately capture and quantify the different sources of uncertainty, particularly model errors.

Some works propose producing forecasts and quantifying their uncertainty by replacing numerical models using machine learning techniques. These works focus on obtaining a surrogate model of a dynamical system directly from the data using, for example, neural networks (Scher and Messori, 2021) or analog regression (Grooms, 2021; Tandeo et al., 2015; Lguensat et al., 2017). Then the data-driven surrogate model can be used in combination of ensemble forecasting techniques for the quantification of the forecast uncertainty, or even for data assimilation as in (Tandeo et al., 2015; Lguensat et al., 2017).

Another approach is to rely on numerical weather prediction to produce a good estimate of the state of the system in combination with machine learning techniques that can assist in reducing the forecast error and providing a more reliable estimation of its uncertainty. There are several works that use machine learning techniques as post-processing tools. For instance, analog regression (Hamill and Whitaker, 2006), random forests (Herman and Schumacher, 2018), neural networks (Rasp and Lerch, 2018) and fuzzy neural networks (Lu et al., 2021) have been successfully applied in this context.

Artificial neuronal networks (ANNs) are used to reduce systematic errors, particularly in situations in which the systematic error relationship with the state of the system is highly nonlinear (e.g. Haupt et al. (2021)). Farchi et al. (2021) use machine learning to represent model errors and use data assimilation to update the machine learning model parameters as new observations became available. Recently, neural network representations have been proposed to jointly estimate the systematic component of the forecast error and the forecast uncertainty (e.g. Rasp and Lerch, 2018; Scher and Messori, 2018; Camporeale, 2018; Wang et al., 2018; Grönquist et al., 2019).

Scher and Messori (2018) designed a methodology based on a convolutional neural network that learned the

standard deviation of the error from an existing ensemble system. The convolutional neural networks predict the spatial distribution of the forecast error variance using a deterministic forecast as input. The learning target is the error variance estimated from an ensemble forecast. The skill of the network on forecasting the uncertainty is lower than the one obtained with an ensemble of forecasts, but the network provides a reasonable quantification of the uncertainty at a much lower computational cost than those of ensemble forecasts. Furthermore, the network performs better than other low-computational cost approaches, such as a simple estimation of the uncertainty based on a multiple linear regression. Along the same lines, [Grönquist et al. \(2021\)](#) use a neural network to improve the uncertainty quantification obtained with a limited size ensemble. In this case, a deep neural network is used to model the ensemble spread generated with a large ensemble using a small-sized ensemble as input. They show a significant improvement in the skill of probabilistic forecasts and particularly for extreme weather events. One disadvantage of these approaches is that the network models the uncertainty estimated by an ensemble of forecasts, which can be biased, particularly in the presence of model errors.

A different approach estimates the forecast uncertainty directly from the time series data without requiring an ensemble of forecasts. [D'Isanto and Polsterer \(2018\)](#); [Camporeale et al. \(2019\)](#); [Rasp and Lerch \(2018\)](#) and [Wang et al. \(2018\)](#) introduced neural network architectures and training methodologies that directly quantify the uncertainty taking a deterministic forecast as input and performing the training over a large data-base of deterministic forecasts and their corresponding observations.

The key aspect is the definition of a loss function which should have a well-defined sensitivity to the state-dependent forecast uncertainty. [Wang et al. \(2018\)](#) achieved this by defining a loss function inspired in the observation likelihood, while [Camporeale \(2018\)](#), [D'Isanto and Polsterer \(2018\)](#) and [Rasp and Lerch \(2018\)](#) define a loss function based on the continuous ranked probability score (CRPS) and the Brier score. These works conclude that this is a feasible approach for a cost-effective estimation of the forecast uncertainty for idealized and real-data cases. [Clare et al. \(2021\)](#) propose a slightly different way to estimate forecast uncertainty purely from data. They discretize forecasted variables into bins and predict the probability of the forecasted value to fall within each bin providing a discrete approximation of the probability density function.

This work aims to contribute to the development and evaluation of machine learning-based uncertainty quantification schemes in the context of the prediction of complex chaotic systems. In particular, we want to estimate the deterministic (i.e, bias) and stochastic (i.e., variance) terms of forecast uncertainty, using ANNs trained with a dataset of deterministic forecast and analysis data as inputs and targets respectively. The main objective of this work is to compare different methods to predict error and uncertainty, and to evaluate them in different model error scenarios using a simple dynamical chaotic model to identify how neural networks cope with different sources of error. We focus our attention on the network architecture and loss function that enable a computationally efficient and robust methods of uncertainty quantification.

This work is organized as follows: Section 2 presents an introduction to the problem of uncertainty quantification in the forecasting of chaotic systems, and the notation used in this work. The architecture of the neural networks and the proposed training methodology is described in section 3. A detailed description of the design of all the experiments involved in this work is presented in section 4. Section 5 analyzes the results obtained, and the conclusions of this work are finally presented in section 6.

2 | FORECAST UNCERTAINTY QUANTIFICATION

The evolution of a dynamical system, like the atmosphere or the ocean, is represented in this work via a Markov process

$$\mathbf{x}_k = \mathcal{M}_{k:k-1}(\mathbf{x}_{k-1}) + \boldsymbol{\eta}_k, \quad (1)$$

where \mathbf{x}_k is a S -dimensional state vector of the system at time k , $\mathcal{M}_{k:k-1}$ is a known surrogate model of the system dynamics (e.g., a numerical weather prediction model or a climate simulation model) that propagates the state of the system from time $k - 1$ to time k . This model is assumed to be nonlinear with chaotic properties. The random model error, $\boldsymbol{\eta}_k$, are samples from an unknown probability distribution. Since we consider systematic model errors, this distribution is expected to have a non-zero mean. The evolution of the system results from the recursive application of the model over several time steps. In this case, model errors interact with the system dynamics:

$$\mathbf{x}_{k+l} = \underbrace{\mathcal{M}(\mathcal{M}(\dots(\mathcal{M}(\mathbf{x}_k) + \boldsymbol{\eta}_{k+1}) + \boldsymbol{\eta}_{k+2}) + \dots)}_{l\text{-times}} + \boldsymbol{\eta}_{k+l} \quad (2)$$

In this work, this is simplified to a representation with an additive state-dependent error term

$$\mathbf{x}_{k+l} = \mathcal{M}_{k+l:k}(\mathbf{x}_k) + \boldsymbol{\delta}_{k+l;k}(\mathbf{x}_k, \boldsymbol{\eta}_{k+1}, \dots, \boldsymbol{\eta}_{k+l}), \quad (3)$$

where $\boldsymbol{\delta}_{k+l;k}$ represents the accumulated effect of model errors and its evolution under the nonlinear system dynamics between times k and $k + l$ and $\mathcal{M}_{k+l;k}$ is the nonlinear operator resulting from the l -times recursive application of the surrogate model.

For numerical weather prediction a deterministic forecast (\mathbf{x}^d) is usually obtained by integrating the dynamical model starting from the most probable initial condition usually referred as the analysis (\mathbf{x}^a), such that

$$\mathbf{x}_{l,k+l}^d = \mathcal{M}_{k+l;k}(\mathbf{x}_k^a), \quad (4)$$

where $\mathbf{x}_{l,k+l}^d$ is the forecasted state at lead time l (first subindex) valid at time $k + l$ (second subindex) and \mathbf{x}_k^a is the analysis at time k (i.e., a pointwise estimation of the system state at time k given all the available observations up to time k). This estimation is often obtained using data assimilation (e.g. Kalnay, 2003; Carrassi et al., 2018).

The resulting forecast error is defined as:

$$\boldsymbol{\epsilon}_{l,k+l}^d = \mathbf{x}_{k+l}^t - \mathbf{x}_{l,k+l}^d, \quad (5)$$

where \mathbf{x}_{k+l}^t is the actual state of the system usually referred to as the true state. The forecast error arises due to the presence of errors in the initial conditions and in the model formulation. Initial condition errors are driven by the dynamical system. Furthermore, model error is also partially driven by the dynamical system as explicitly noted in Eq. 2. This effect is state dependent, meaning that error growth rate depends on the state of the system which is a particularly challenging aspect of short and medium range prediction of atmospheric phenomena. The main goal of this

work is to model the state dependent mean ($\bar{e}_{l,k+l}$) and standard deviation ($\sigma_{l,k+l}$) of the forecast error distribution using machine learning techniques. Hereinafter, we drop the time index (second subscript) whenever possible and only retain the forecast lead time as subindex to simplify the notation (i.e., $\mathbf{x}_l^d = \mathbf{x}_{l,k+l}^d$).

Currently, the most common way to estimate the state dependent mean and covariance of the forecast error distribution is through ensemble forecasting. In this approach, the initial state is assumed uncertain and, therefore, given by a probability distribution. This initial distribution is represented through a sample, i.e., an ensemble of initial states. Then, an ensemble of model integrations is performed initialized from the different initial conditions. In the case of nonlinear dynamical models, these integrations spread from each other. The magnitude of the spread is used to estimate the forecast standard deviation at a given lead time. The n -th ensemble member ($\mathbf{x}^{e,(n)}$) is generated by integrating the model from the n -th initial condition, $\mathbf{x}^{a,(n)}$,

$$\mathbf{x}_l^{e,(n)} = \mathcal{M}_{l;0}^{(n)}(\mathbf{x}^{a,(n)}) + \hat{\delta}_{l;0}^{(n)}, \quad (6)$$

where $\mathbf{x}_l^{e,(n)}$ is the forecasted state vector given by the n -th ensemble member at lead time l , $\hat{\delta}_{l;0}^{(n)}$ is a representation of the effect of model errors, and $\mathcal{M}_{l;0}^{(n)}$ is the numerical model. Note that a different numerical model can be used for the integration of each ensemble member as in stochastically perturbed parameters or parametrization outputs.

Given an ensemble of forecast states (i.e., a sample of possible future states of the system), a pointwise estimate of the state of the system at lead time l can be obtained from the ensemble mean state vector:

$$\bar{\mathbf{x}}_l^e = \frac{1}{N} \sum_{n=1}^N \mathbf{x}_l^{e,(n)}, \quad (7)$$

where N is the number of ensemble members. This estimate is statistically more accurate in terms of the root mean square error, than the one provided by a deterministic forecast (Kalnay, 2003). The forecast uncertainty is usually quantified in high-dimensional systems through the second moment by computing the sample error covariance matrix,

$$\Sigma_l^e = \frac{1}{(N-1)} \sum_{n=1}^N (\mathbf{x}_l^{e,(n)} - \bar{\mathbf{x}}_l^e) (\mathbf{x}_l^{e,(n)} - \bar{\mathbf{x}}_l^e)^\top. \quad (8)$$

While the full covariance matrix is commonly used in data assimilation applications (see Carrassi et al. (2018) for more details), most forecast applications are only concerned with the error associated with the forecasted variables so that there is no need to quantify the correlation between the state variables and only the diagonal elements of the covariance matrix, $\sigma_l^{e^2} = \text{diag}(\Sigma_l^e)$ are of interest. The accuracy of the sample covariance estimate depends on the ensemble size and on the adequacy of model error representation.

In the imperfect model scenario forecast errors are larger due to the contribution of model errors. Moreover, the state-dependent forecast error mean is usually non zero, i.e., there is a systematic error component.

3 | METHODOLOGY

Our goal is to design and train an ANN which takes as input a deterministic forecast and provides a better estimate of the state of the system ($\bar{\mathbf{x}}$, i.e. by removing the systematic component of model errors or by smoothing out the unpredictable components of the state) and its uncertainty ($\bar{\sigma}$, i.e. the standard deviation of the error in $\bar{\mathbf{x}}$). In a broad sense we would like to use an ANN to emulate the advantages of a well tuned ensemble forecast without the

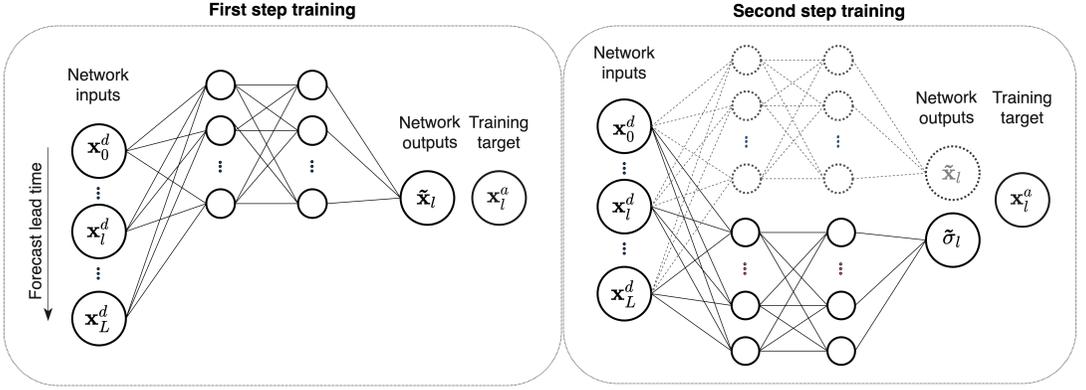


FIGURE 1 Schematic of the network architecture and the two-stage training approach. The neural network architecture is divided in two parts, one for the estimation of the corrected forecast and one for the estimation of its uncertainty. The left panel represents the first training step in which the network that predicts the corrected forecast is trained. On the right panel the network that predicts the uncertainty is trained. The network providing the corrected forecast is used but not trained during this step. Note that networks that train uncertainty directly from the ensemble must incorporate this data (ensemble spread) into the target. This figure does not show this case.

computational burden of performing several model integrations.

To jointly estimate the corrected forecast \bar{x} and its standard deviation $\bar{\sigma}$ at a certain forecast lead time l , we propose to use multi-layer perceptron. The input of the network is a deterministic forecast of the state of the system at different lead times $\{\mathbf{x}_0^d, \dots, \mathbf{x}_l^d, \dots, \mathbf{x}_L^d\}, 0 \leq l \leq L$ providing information about the dynamic evolution of the system over a time window surrounding lead time l . The outputs of the network are a corrected forecast (\bar{x}_l) and its corresponding standard deviation ($\bar{\sigma}_l$) at lead time l .

The network is composed of two sub-networks, one for the corrected forecast (\bar{x}_l) and the other for the estimation of its associated standard deviation ($\bar{\sigma}_l$), working together in a tandem scheme as shown in Figure 1. Taking advantage of the proposed architecture, we perform the training of the network in two steps. Firstly, we train the sub-network that produces the corrected forecast using the analysis states valid at the same time of the forecast (\mathbf{x}_l^a) as targets and a mean squared error (MSE) based loss function,

$$\mathcal{L}_{MSE}^x(\bar{\mathbf{x}}_l; \mathbf{x}_l^a) = (\bar{\mathbf{x}}_l - \mathbf{x}_l^a)^\top (\bar{\mathbf{x}}_l - \mathbf{x}_l^a). \quad (9)$$

The analysis states are used as targets, rather than the available observations, because they provide a complete estimation of the state vector on the model grid and at regular times and with an error which is statistically smaller than the observational error (Carrassi et al., 2018). This scheme directly provides an estimation of the state in a supervised manner.

Secondly, the sub-network that estimates the state dependent standard deviation of the corrected forecast error is trained. This network is the main focus of this study. For this second step, we compare three strategies, which are described below, that differ from each other in the data used as target and in the formulation of the loss function.

3.1 | Training with the ensemble spread (NN-mse)

In this strategy, we follow a similar approach as in [Scher and Messori \(2018\)](#) where $\bar{\sigma}_I$ is estimated with a network trained using a dataset of σ_I^e (i.e. the sample forecast standard deviation provided by an ensemble forecast at lead time I). In this scheme the networks learns the dependence between the deterministic forecast and forecast error standard deviation in a supervised manner.

Network training is performed using a MSE-based loss function which is defined as follows:

$$\mathcal{L}_{MSE}^{\sigma}(\bar{\sigma}_I; \sigma_I^e) = (\bar{\sigma}_I - \sigma_I^e)^{\top} (\bar{\sigma}_I - \sigma_I^e). \quad (10)$$

The main drawback of using direct training is that it requires an *explicit uncertainty* data set (i.e., σ_I^e) which in turn needs a full ensemble forecast integration to provide an estimation of the forecast error variance for each element in the training sample. Furthermore, it will inherit the limitations of the ensemble forecast.

3.2 | Training with a locally estimated error (NN-ext)

In this approach the only information used for training are the analysed states (\mathbf{x}_I^a). We refer to this kind of approach as *indirect training*, because no a-priori estimation of the standard deviation of the forecast error is required since this is estimated directly from the data (e.g. [Camporeale, 2018](#); [Wang et al., 2018](#)). Indirect training approaches, require a loss function definition that is sensitive to the standard deviation of the forecast error (i.e. an uncertainty aware loss function). To construct such a loss function, first a local proxy of the corrected forecast error is obtained,

$$\bar{\epsilon}_I = \bar{\mathbf{x}}_I - \mathbf{x}_I^a. \quad (11)$$

Then, the element-wise absolute value of the error (i.e. $|\bar{\epsilon}_I|$) is used as a local estimate of the standard deviation of the corrected forecast. In this case the loss function is defined as,

$$\mathcal{L}_{eMSE}^{\sigma}(\bar{\sigma}_I; \bar{\epsilon}_I) = (\bar{\sigma}_I - |\bar{\epsilon}_I|)^{\top} (\bar{\sigma}_I - |\bar{\epsilon}_I|). \quad (12)$$

We named this loss function *extended MSE (eMSE)*, since we use the MSE metric in combination with a target based on a local estimation of the absolute value of the error ($|\bar{\epsilon}_I|$). The eMSE loss function is expected to be used in a batch of observations. Note that in the computation of $\mathcal{L}_{eMSE}^{\sigma}$ we use $\bar{\mathbf{x}}_I$ and \mathbf{x}_I^a to compute the target ($|\bar{\epsilon}_I|$), this requires the evaluation of the sub-network that provides the corrected forecast.

3.3 | Training based on the sample likelihood (NN-lik)

Another form of indirect training can be obtained assuming that the error in the corrected forecast ($\bar{\mathbf{x}}$) follows a Gaussian distribution with zero mean and covariance $\bar{\Sigma}$. We approximate the covariance matrix as a diagonal matrix,

$$\bar{\Sigma}_I \approx (\bar{\sigma}_I^{\top} I_S \bar{\sigma}_I), \quad (13)$$

where I_S is the $S \times S$ identity matrix. The diagonal assumption significantly simplifies the definition of $\tilde{\Sigma}$ and the computation. This is taken because only the quantification of the error variance is usually required for forecasting applications.

The loss function based on the negative of the log-likelihood is given by

$$\mathcal{L}_{lik}^{\sigma}(\tilde{\sigma}_l; \tilde{\epsilon}_l) = \frac{\log(\det(\tilde{\Sigma}_l))}{2} + 2\tilde{\epsilon}_l^T (\tilde{\Sigma}_l)^{-1} \tilde{\epsilon}_l \quad (14)$$

Then $\tilde{\sigma}$ can be estimated as the one that minimizes the loss function, and so it maximizes the log-likelihood of the analyzed states over the entire dataset (Wang et al., 2018; Pulido et al., 2018; Tandeo et al., 2020).

By introducing this loss function, we are assuming independence of the errors at different model variables and times for this forecasting application. Moreover, the analysis uncertainty is not explicitly taken into account. As with the NN-ext (sec. 3.2), this loss function has the advantage of using an *implicit uncertainty* dataset and then there is no need for an a priori estimation of the forecast uncertainty.

In this section, we defined different loss functions considering a single sample of the dataset. However, it should be kept in mind that the optimization of the ANNs aims to minimize the loss function of the entire dataset, which is given by the sum over all the samples in the training dataset.

4 | EXPERIMENTAL DESIGN

Experiments using a simple chaotic dynamical system are conducted to evaluate the performance of the ANNs. Two different scenarios are considered, a scenario in which the model used to estimate the state of the system and to produce forecasts is a perfect representation of the true system dynamics (i.e., $\eta_k = 0$ in Eq.1) and another scenario in which the model is a surrogate model in which small-scale dynamics are not represented and, therefore, forecasts are affected by model errors.

4.1 | Perfect model scenario

In the perfect-model scenario (PMS) we use the single-scale version of Lorenz'96 equations (Lorenz, 1995) given by

$$\frac{dx_{(i)}}{dt} = -x_{(i-2)}x_{(i-1)} + x_{(i-1)}x_{(i+1)} - x_{(i)} + F, \quad (15)$$

where $x_{(i)}$ is the i -th component of \mathbf{x} and F is a constant that represents an external forcing. Cyclic boundary conditions are imposed so that this set of equations resembles an advective-dissipative forced system over a circle of latitude.

Firstly, a long integration is performed with the model, which is assumed to be the unknown true evolution of the system. The model is integrated during 500 time units using the 4th order Runge-Kutta method and a forcing of $F = 8$, a time-step of 0.0125 time units and $S = 8$ state variables. This simulation will be referred to as the *nature* integration. Secondly, a set of noisy observations are generated from the *nature* integration by adding an uncorrelated Gaussian noise with zero mean and covariance $\mathbf{R} = \sigma_R^2 \mathbf{I}$ with $\sigma_R = 1$. All the state variables are observed every 4 time-steps (0.05 time units).

Initial conditions for the numerical forecasts are generated using an ensemble-based data assimilation system (the Local Ensemble Transform Kalman Filter, LETKF; Hunt et al., 2007), which provides an analysis every time observations

are available (i.e., every 0.05 time units). An ensemble of 50 members is used to generate the analysis. The data assimilation is conducted over the length of the *nature* integration giving a total of 13000 estimations of the state of the system and its associated uncertainty (7000 used as training set, 3000 for validation and 3000 for testing). Finally, two sets of forecasts are generated, an ensemble forecast in which each member is initialized from the members of the analysis ensemble (to be used in the direct training approach) and a deterministic forecast initialized from the mean of the analysis ensemble (to be used as input to the neural network). The target used for the training of the ANNs is the mean of the ensemble analysis states. The forecasts are integrated up to a maximum lead time of 3.5 time units (280 time steps). This maximum lead time is taken because the forecast error is close to its nonlinear saturation at this lead time.

4.2 | Imperfect model scenario

For the generation of the *nature* integration in the imperfect model scenario (IMS), we use the two-scale Lorenz' model (Lorenz, 1995) which is defined by the following system of coupled differential equations:

$$\frac{dx_{(i)}}{dt} = -x_{(i-1)}(x_{(i-2)} - x_{(i-1)}) - x_{(i)} + F - \frac{hc}{b} \sum_{j=J(i-1)+1}^{iJ} y_{(j)} \quad (16)$$

$$\frac{dy_{(j)}}{dt} = -cb y_{(j+1)}(y_{(j+2)} - y_{(j-1)}) - c y_{(j)} + \frac{hc}{b} x_{(\text{int}[(j-1)/J]+1)}, \quad (17)$$

where $x_{(i)}$ are the state variables associated to the slow dynamics, while $y_{(j)}$ are the variables associated with the faster dynamics. J is the number of y variables for each x variable, and h , c , and b are time-independent parameters controlling the coupling strength between the two systems. Cyclic boundary conditions are applied to both sets of equations, namely $x_{(1)} = x_{(S+1)}$, and $y_{(1)} = y_{(J \cdot S+1)}$. This two-scale system is a simple representation of systems with multiple spatio-temporal scales like the atmosphere or the ocean. In our experiments, $J=32$ and $S=8$ (i.e., the y vector has a total of 256 variables) and the forcing term F is set to 20 to obtain a chaotic behavior.

The model used for data assimilation and forecasting is given by

$$\frac{dx_{(i)}}{dt} = -x_{(i-1)}(x_{(i-2)} - x_{(i-1)}) - x_{(i)} + F + G_{(i)}, \quad (18)$$

where $G_{(i)}$ is a state dependent parametrization term that approximates the effect of the missing dynamics (i.e., the effect of fast variables y). As in Pulido et al. (2016), $G_{(i)}$ is assumed to be a linear function of the state variable $x_{(i)}$

$$G_{(i)} = \alpha x_{(i)} + \beta, \quad (19)$$

with $\alpha = 19.16$ and $\beta = -0.81$ constant parameters whose optimal values are taken from Scheffler et al. (2019). The general experimental setting including the observation frequency, observation noise, and length of the assimilation experiment are the same as in the PMS. The ensemble forecasts generated in the imperfect model scenario do not include an explicit representation of model errors (i.e., $\delta_{i,0}^{(n)} = 0$ in Eq.3), all ensemble members are integrated with the same model configuration version.

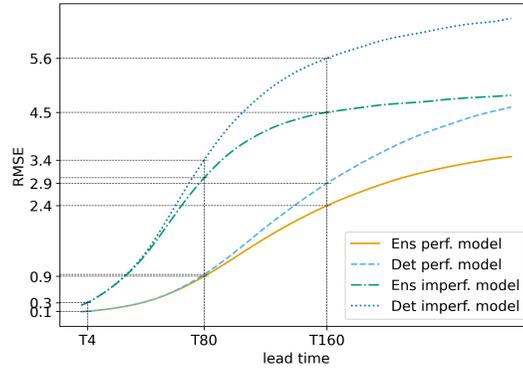


FIGURE 2 RMSE as a function of lead time in the perfect and imperfect model scenarios using the evolution of a single model integration (Det) and the mean of an ensemble of 50 integrations of the model (Ens).

4.3 | Network architecture and training

As a result of a design experiment process, the networks used in our experiments are fully-connected multilayered networks mainly because these networks are able to account for high nonlinearities.

As has been already mentioned in section 3, the architecture consists of two sub-networks that can be trained separately in two steps (see Fig. 1). Each network has two hidden layers with 50 neurons each and softplus activation functions (Dugas et al., 2001). The first half of the output corresponds to the estimation of the system state and must generate outputs within the range of possible state values. Therefore, we use a linear activation function in the first half of the output layer. The second half corresponds to the estimation of the forecast standard deviation and, consequently, must be positive. Hence, we use a softplus activation function in the second half of the output layer.

The network architecture was defined based on preliminary experiments varying the size of the hidden layers (from 16 to 200 neurons) and the activation functions of the hidden layers (sigmoid, relu and softplus). All the networks are trained with the Adam optimizer (Kingma and Ba, 2017) and using minibatches with a batchsize of 50 elements. The loss function is evaluated over the validation set every 20 training epochs. The training stops when the loss function evaluated over the validation set stops decreasing or starts to increase (early stop). The hyper-parameters values were selected by tuning them by hand in preliminary experiments (not shown here).

Note that the input for both sub-networks is the same, and that the output of the first network is used in the computation of the loss function during the training of the second network when either eMSE or the likelihood loss functions are used. Two-step training is not strictly necessary, as both networks can be trained together in a single step, but we found that two-steps training facilitates and improves the convergence of the weights. We believe this is because a two-step training helps to improve the conditioning of the minimization problem thus leading to a faster convergence to the global minimum. Also, preliminary experiments indicate that the learning capability of this architecture is very similar to that of a network in which the two sub-networks are connected and trained in a single step.

For the evaluation of the performance of the ANNs in the experiments, three different lead times were taken as representative, 4, 80 and 160 time steps, on the two scenarios, PMS and IMS. Figure 2 shows the evolution of deterministic and ensemble mean forecast errors for the PMS and IMS as a function of lead time. The first lead time (T4) is of interest because it roughly corresponds to a very short range forecast (4 time steps), like those routinely used in the data assimilation cycle. At this lead time, the forecast error is comparable to the error present in the

analyses used as training target by the ANN. The second lead time (T80) is within the linear error growth regime, particularly in the perfect model experiment (see Fig. 2). At this lead time, the difference between the error in the deterministic forecast and the ensemble forecast is small in the PMS but is important in the IMS. Unlike T4, the initial error propagated through the dynamics of the system has grown enough to become independent of the error present in the analyses. The third lead time (T160) is within the nonlinear error growth regime and it is also close to the error saturation in the IMS. In this lead time, the initial error has evolved through the dynamics of the system and grown large enough to dominate the state of the system. Therefore, the system is at the limit of predictability regardless of the model error.

For each scenario (PMS and IMS) and for each selected lead time (T4, T80, T160), we conducted a deterministic forecast, an ensemble forecast (baseline for comparison) and three ANNs predictions which were trained using the three approaches described in Section 3 (i.e. NN-mse, NN-ext and NN-lik).

4.4 | Verification metrics

To compare different aspects of the performance of the different forecasting system (deterministic, ensemble, and ANNs with different training strategies) we use three metrics, root-mean-square error, coverage probability and standard deviation-error correlation. They are used as diagnostics in which the corresponding *nature* state, \mathbf{x}^t , is used as ground truth.

4.4.1 | Root mean square error

The purpose of this metric is to evaluate the accuracy in estimating the state of the system. The closer the value is to zero, the more accurate the forecast estimate. The RMSE is computed as

$$\text{RMSE} = \sqrt{\frac{1}{S \cdot M} \sum_{k=1}^M (\bar{\mathbf{x}}_k - \mathbf{x}_k^t)^\top (\bar{\mathbf{x}}_k - \mathbf{x}_k^t)}, \quad (20)$$

where \mathbf{x}_k^t is the true state of the system corresponding to the forecast time, S is the state vector dimension and M is the total number of samples in the verification set. The RMSE is a deterministic-oriented metric [Murphy \(1993\)](#) and does not take into account the forecast uncertainty, so that in this work the main purpose of this metric is to measure the accuracy of the corrected forecast ($\bar{\mathbf{x}}$).

4.4.2 | Coverage Probability

To evaluate the quality of the estimation of the standard deviation made by the network we use the coverage probability that consists of, given a confidence interval for the estimator, calculating the number of times that the confidence interval contains the true value

$$\text{CP} = \frac{\sum_{i=0}^M \mathbb{1}(\ell_i < x_i^a < u_i)}{M}, \quad (21)$$

where ℓ_i/u_i is the lower/upper bound of the confidence interval for an estimator with normal distribution $\mathcal{N}(\bar{x}_i, \bar{\sigma}_i)$, $\mathbb{1}$ is the indicator function (takes the value of 1 if the condition is true and 0 otherwise) and M the size of the dataset. The perfect value for CP is equal to the confidence level used to construct the confidence interval, in our case 90% or

(0.9). A CP value higher/lower than the confidence level would indicate an over/underestimation of the uncertainty. In this work we use the coverage probability as a global measure of the reliability of the quantified uncertainty.

4.4.3 | Standard deviation-absolute error correlation

Since we aim to estimate the state dependent standard deviation of the forecast error, the Pearson's correlation coefficient is computed between the estimated forecast standard deviation and the absolute value of the actual forecast error ($|\epsilon_l| = |\bar{x}_l - x_l^f|$),

$$\text{CORR}(\bar{\sigma}_l, |\epsilon_l|) = \frac{\text{cov}(\bar{\sigma}_l, |\epsilon_l|)}{\text{std}(\bar{\sigma}_l)\text{std}(|\epsilon_l|)} \quad (22)$$

where cov and std are the spatio-temporal covariance and standard deviation respectively. The purpose of this metric is to evaluate the correspondence in the temporal evolution of the uncertainty and the error produced in the forecast. A good estimate of uncertainty should be correlated with the forecast error. Negative values or values very close to zero could indicate poor estimation.

5 | RESULTS

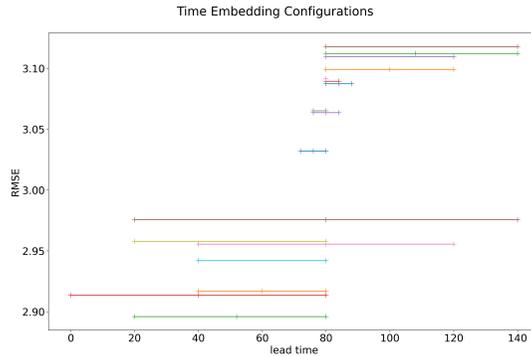


FIGURE 3 RMSE of the corrected forecast for experiments with different inputs. The RMSE is indicated in the y axis, while the x axis indicates the lead time. Each experiment is represented by a horizontal line with a different color, indicating the lead time range spanned by the input. The dots indicate the specific lead times aggregated to construct the input to the ANN. For example, the dark green line at the bottom of the figure has dots at lead times 20, 50 and 80 meaning that the input to the network resulted from the aggregation of the deterministic forecast at these three lead times. In this case, the RMSE achieved over the testing sample corresponding to this network is 2.9.

5.1 | Sensitivity to network input

We conducted experiments to investigate the sensitivity of the network performance to the forecast lead times included in the input. The input to the network consists of an aggregation of L model states at different forecast lead times, with $L = 1, 2, 3$; i.e. $\{x_{l_1}\}$, $\{x_{l_1}, x_{l_2}\}$ or $\{x_{l_1}, x_{l_2}, x_{l_3}\}$, in which l_j is representing the forecast lead time.

We performed a set of 17 experiments combining different number of model states (from 1 to 3 lead times), lead time window lengths (from 4 to 60 time steps) and lead time shifts with respect to the output lead time, to evaluate how these different inputs influence the accuracy of the corrected forecast (\bar{x}). In these experiments, the output variable is the model state at T80 (i.e., $l = 80$) lead time. The experiments are conducted in the imperfect model scenario using the same hyper-parameter and network architecture. For each experiment, we repeat the ANN training ten times using different initial weights to reduce the impact of weight initialization upon the resulting performance of the networks.

Figure 3 shows the RMSE of the predicted state \bar{x}_l obtained with different inputs. The RMSE is indicated in the y axis, while the x axis indicates the lead time. Each experiment is represented by an horizontal line indicating the lead time range spanned by the input. The dots indicate the specific lead times aggregated to construct the input to the ANN.

In general, all the inputs performed reasonably well. However, the ones including information from lead times prior to the output lead time (l) perform better than the ones relying on lead times after l . This may be related with the fact that forecast errors, and particularly model errors, evolve according to the system dynamics (see Eq. 2), thus including information about previous states can contribute to better describe the state dependent forecast mean error at lead time l . The range of lead times spanned by the input also produces a significant impact upon the ANN performance. Networks spanning a wider range of lead times (and using information from lead times prior to l) perform better. Furthermore, the performance of the ANN with $L = 3$ improves over the ones using $L = 2$ and $L = 1$. Best performance in the experiments is obtained for the ANN with three input variables at lead times T20, T50, and T80, i.e. $\{x_{20}^d, x_{50}^d, x_{80}^d\}$.

Based on these results, in the rest of this work, all the ANNs use the following inputs: for T4 the input consists of $\{x_0^d, x_4^d\}$ (only two input lead times because not intermediate data would be available in realistic systems), for T80 we use $\{x_0^d, x_{40}^d, x_{80}^d\}$; and for T160, we use $\{x_0^d, x_{80}^d, x_{160}^d\}$.

5.2 | PMS experiments

In the perfect model experiments, forecast errors are the result of errors in the specification of the initial conditions. This particular scenario allows us to evaluate how well the ANNs capture error growth due to the chaotic behavior. Furthermore, these experiments are designed to evaluate whether ANNs are able to filter the most uncertain modes present in the forecasted state as in the case of the forecast ensemble mean (e.g. Kalnay, 2003). In ensemble forecasting, the ensemble mean \bar{x}_l^e exerts a filtering of the forecast error associated to predictability loss as l increases. Moreover, the ensemble spread (σ_l^e) in the PMS is a reliable estimate of the state-dependent standard deviation of the forecast error. This is particularly true when the initial ensemble perturbations are generated using an ensemble Kalman filter (Candille et al., 2007).

Figure 4a shows the time evolution of the ensemble mean forecast at T80 for the first state variable (blue solid line) and the ensemble standard deviation (blue shadow), the black solid line is the truth and the red dashed line is the deterministic forecast. Figure 4b shows the neural network prediction obtained in the NN-lik approach. The green line shows the time series of the first state variable and the green shade its forecast standard deviation as predicted by the network. At first glance, both the ensemble and the networks provide good estimations of the state of the system. Since we are in the linear error regime, in general, no big differences are found between the performance of the deterministic forecast and those from the ensemble mean and the network. In terms of forecast uncertainty, overall, both systems provide a good estimation of the state dependent forecast uncertainty (i.e., the black curve is within the shaded area most of the time). The uncertainty estimated from the ANN shows a smaller time variability

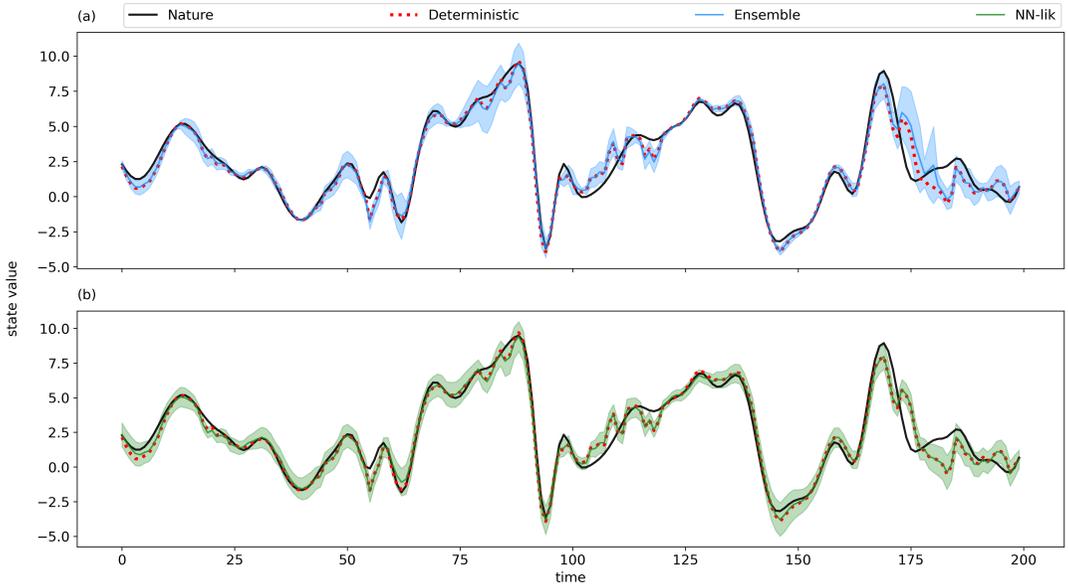


FIGURE 4 Time evolution of the first state variable in the T80 forecast experiment for 200 consecutive time steps taken from the testing set. (a) Mean of the ensemble forecasts (solid blue line) and the predicted spread (shaded). (b) Prediction with the NN-lik approach (solid green line) and the predicted spread (shaded). In both panels, the deterministic model forecast (red dotted lines) and the ground truth (solid black lines) are shown.

than the one from the ensemble. The relationship between the forecast and its uncertainty is not trivial. Similar values of $x_{(1)}^d$ are associated with different uncertainties. However, the periods of rapid increase or decrease in $x_{(1)}$ with time appears to be associated with lower forecast uncertainty in both, the ensemble and the ANN.

Figure 5 shows the spatio-temporal evolution of the absolute error ($|\epsilon|$, left panels) and the estimated standard deviation of the forecast (right panels) for the ensemble, the NN-lik approach and the deterministic forecast. In the case of the deterministic forecast, its uncertainty is estimated as the standard deviation of the difference between the forecast and the analysis ($x_{l,k+l}^d - x_{0,k+l}^a$) over the training sample, and thus is independent of time (Figure 5f). Absolute errors for the three experiments show a relatively large variability and are rather similar in terms of magnitude, spatio-temporal evolution and location of the extreme values. This confirms the strong dependence of forecast errors with the state of the system.

In terms of standard deviation prediction (Figs 5b, 5d and 5f), areas with maximum spread in the ensemble shows are in good agreement with the location of larger forecast errors. It is important to note that we do not expect a perfect match between these two quantities since the forecast error is one realization from a probability density function whose standard deviation is being approximated from the forecast ensemble. The ANN also provides an estimation of the forecast error standard deviation in good agreement with the evolution of the forecast error. As has been previously noted, the forecast standard deviation estimated from the network is smoother and shows less variability than the one estimated from the ensemble.

These results are quite encouraging. They show that the NN-lik approach is able to estimate the forecast uncertainty as a function of the state of the system. Moreover, this estimation is achieved using a deterministic forecast as input whose computational cost is much cheaper than integrating an ensemble of forecasts.

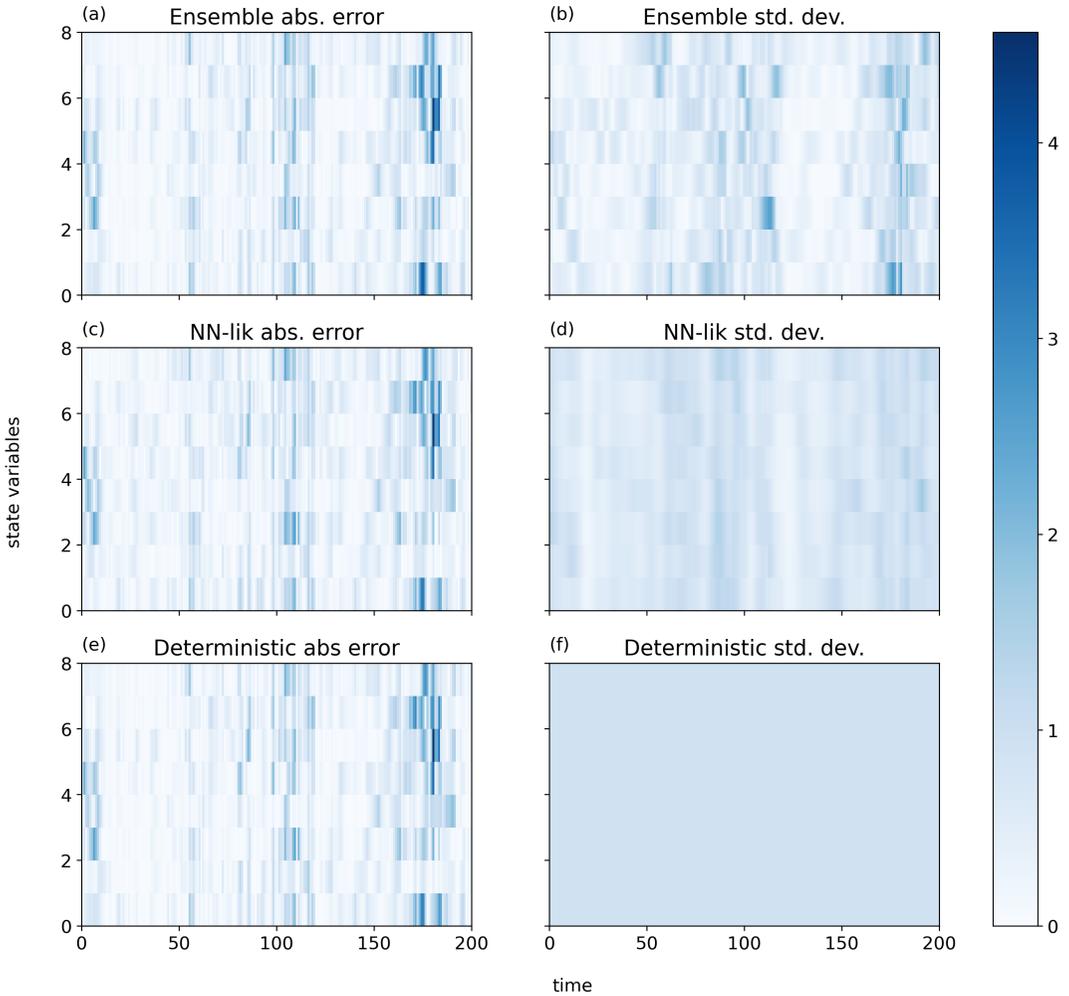


FIGURE 5 Left panels (a,c,e): Time evolution of the absolute error for each state variable (y-axis) of the dynamic model at T80 lead time for the ensemble mean (a), the NN-lik approach (c), and the deterministic forecast (e). Right panels (b,d,f): the forecast error standard deviation as estimated from the ensemble (b), the NN-lik approach (d) and the deterministic forecast (f).

So far we have qualitatively analyzed the skill of the ANN in quantifying the forecast uncertainty from a deterministic forecast at lead time T80. Figure 6 summarizes the results obtained for the different output lead times considered (T4, T80 and T160) and for the different scores described in Section 3. Confidence intervals for the different scores are computed using a bootstrap approach in which 500 samples are obtained from the original sample using random selection with replacement. Forecast at valid times which are more than 20 time steps apart from each other are selected for the computation of the scores to increase the independence between different sample elements.

Regarding the RMSE (first row of Fig. 6), all forecast systems show a similar performance for T4 and T80 with the ensemble mean performing slightly better at T80. All the ANNs have the same RMSE since the training of the

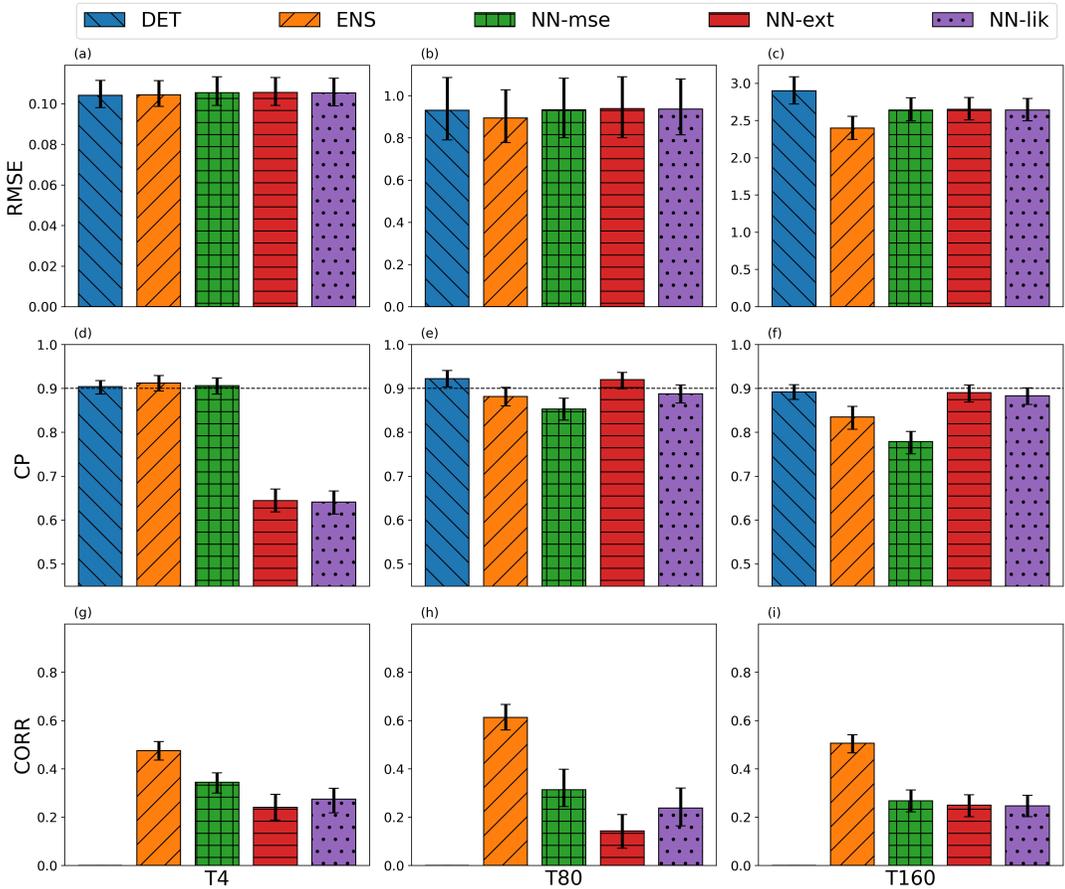


FIGURE 6 Scores corresponding to the deterministic forecast (blue bar), ensemble forecast (yellow bar), and the NN-mse (green bar), NN-ext (red bar) and NN-lik (purple bars) approaches in the perfect model scenario. The first row corresponds to the RMSE, the second to the 90% coverage probability and the third to the standard deviation-absolute error correlation coefficient. The first, second and third columns corresponds to the T4, T80 and T160 lead times respectively. The error line on top of the bars indicates the 95% confidence interval calculated using bootstrap technique. The horizontal line in the second row panels indicate the perfect value associated with the 90% coverage probability.

sub-network producing the corrected forecast is the same for all of them.

For time T160 (Fig. 6c) error growth is affected by nonlinear effects, the ensemble is able to adequately filter some of the unpredictable modes resulting in a lower RMSE with respect to the deterministic forecasts. One important result is that the ANNs outperforms the deterministic forecast for this lead time. This suggests that ANNs are able to filter at least part of the unpredictable modes, although it is not as efficient as the ensemble mean in doing that.

The second row in Fig. 6 shows the coverage probability (CP). A strong underestimation of the uncertainty is present in the estimation of the NN-ext and NN-lik approaches for the T4 forecast times. At lead time T4, the magnitude of the forecast error in PMS is highly correlated with the analysis error, namely the forecast error proxy used for training is underestimating the magnitude of the forecast error. This affects the NN-ext and NN-lik approaches which

rely on this proxy for the estimation of the forecast uncertainty but not the MSE network which is trained against the uncertainty estimated by the ensemble. To verify this conjecture, we performed additional experiments training a network with the true error (i.e. the difference between the deterministic forecast and the true state). As a result, an improvement in terms of RMSE and CP is found. In particular the sub-estimation of the forecast uncertainty at T4 is not present in these experiments. CP values went from 0.63 when training with analysis to 0.87 when training with ground truth and RMSE values varies from 0.12 to 0.10 respectively.

At longer lead times (T80 and T160) the ensemble and the NN-mse network under-estimate the forecast standard deviation. The reason for this is not clear but it may be the result of a sub-optimal estimation of the analysis uncertainty in the ensemble-based data assimilation system. In any case, this behaviour shows a possible limitation of the NN-mse approach which can inherit biases in the quantification of the forecast uncertainty derived from a sub-optimal formulation of the ensemble used for the training of the network.

The third row in Fig. 6 shows the spatio-temporal correlation between the estimated forecast standard deviation and the absolute forecast error. A low correlation between the error and the estimated forecast standard deviation indicates that the system is unable to capture the state-dependent nature of the forecast uncertainty. Figures 6g, h and i show positive correlations for the ensemble forecast and the ANNs. As expected the time independent standard deviation of the error associated with the deterministic forecast is uncorrelated with the actual forecast error. The ensemble forecast outperforms the ANNs for all the considered lead times. This is an expected result since the ensemble forecasts is based on direct sampling from the probability density function of the state of the system, particularly in these perfect model experiments. The NN-mse network which is trained directly from an ensemble of forecasts performs slightly better than the indirectly trained networks at T4 and T80, but at T160 all ANNs shows similar performances.

5.3 | IMS experiments

In the imperfect model scenario forecast errors are larger due to the contribution of model errors. Moreover, the forecast error mean is usually non zero, i.e., there is a systematic error component. In this scenario the ANNs aim to "remove" the systematic error component from state forecast and to capture state-dependent contributions of model errors to the forecast standard deviation.

As shown, the error grows more rapidly in the IMS than in the PMS (see Figure 2) meaning that T80 and T160 are representative of different model error regimes. In the IMS experiment, the T80 lead time is at the beginning of the nonlinear regime and T160 is near the error saturation.

Figure 7 shows the spatio-temporal evolution of the absolute error and its standard deviation estimated by the ensemble and the NN-lik approach. Although the error patterns are not as similar as those observed in PMS (Fig. 5), there is a strong similarity in shape and intensity (Fig. 7a and 7c). The standard deviation panels, Fig. 7b and 7d, show that both, the network and the ensemble, are able to represent the state-dependent forecast uncertainty as in the PMS experiment. The estimated forecast standard deviation is larger for the NN-lik approach than for the ensemble. This indicates a sub-estimation of the error standard deviation in the ensemble forecasting approach that may be partially attributed to the lack of a model error representation in the ensemble forecast experiment.

A summary of the performance of the different forecast systems in the IMS scenario is shown in Fig. 8 for lead times T4, T80 and T160. In terms of the RMSE, a noticeable advantage in the accuracy of the state estimate is achieved by the ANNs at all lead times. The ANN corrected forecasts are even better than the ensemble mean forecast. This suggests that the contribution of model errors is significant and comparable or even larger to the error associated with the imperfect knowledge of the initial conditions. The ANNs are able to partially capture the state-dependent

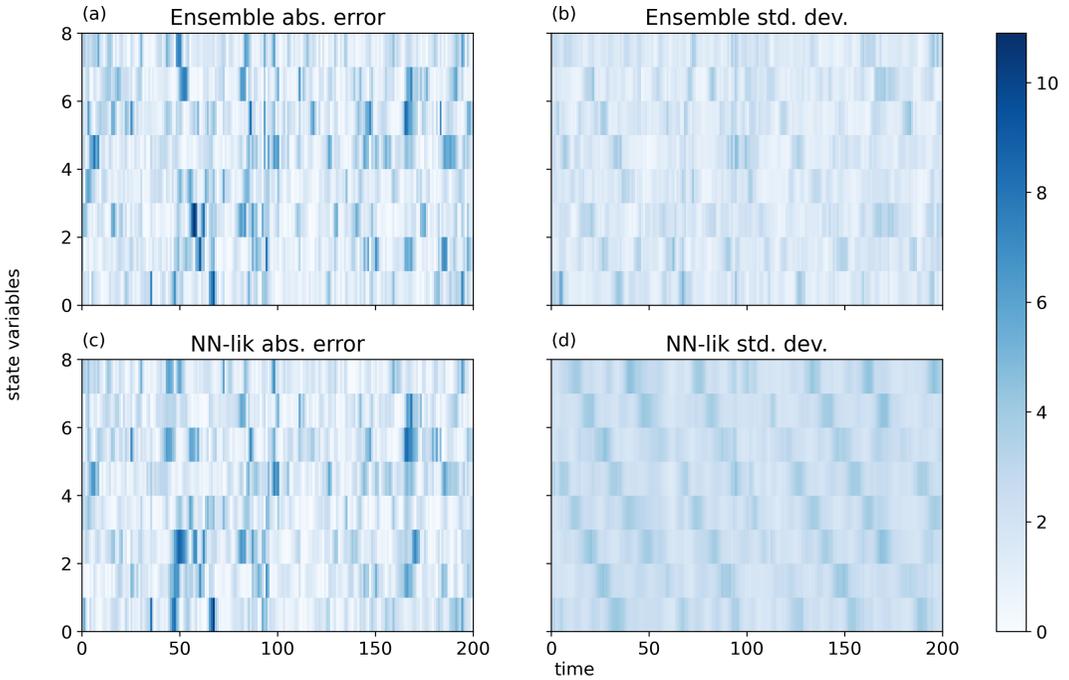


FIGURE 7 Left panels(a,c):Time evolution of the absolute error for each state variable (y-axis) of the dynamic model at T80 lead time for the ensemble mean (a) and the NN-lik approach (c). Right panels (b,d): the forecast error standard deviation as estimated from the ensemble (b) and the NN-lik approach (d)

model errors while the deterministic forecast and the ensemble forecast are not.

The CP (second row in Fig. 8) shows a clear advantage of ANNs over the deterministic and ensemble forecasts, particularly at T80 and T160. At T4 a similar behaviour is found as in the PMS with NN-ext and NN-lik performing worse than NN-mse. This is again a direct consequence of using the analysis as a learning target for the estimation of the forecast uncertainty. Once the forecast error has become larger (e.g. T80), this behavior is reversed: the networks that learn the uncertainty from the data perform better than the ensemble in terms of CP shown in Figs. 7, 9 and 10. It is interesting to note that the NN-mse learning technique, which learns directly from the ensemble spread performs better than the ensemble itself in terms of CP. This is partially because, the network improves the forecasted state by removing the state-dependent systematic forecast error component. Thus, the estimated forecast standard deviation is more consistent with the actual forecast errors than in the case of the ensemble forecast. This suggests that even an ensemble with a sub-standard representation of model errors can be a reasonable target for the training of an ANN, although it is not clear whether this particular result generalizes to more complex model errors as the ones present in state of the art numerical weather prediction or climate models. The ensemble shows the worst performance in terms of CP for T80 and T160, even worse than the deterministic forecast. This is because the ensemble mean is affected by model errors which are not accounted for in the formulation of the ensemble leading to an under-estimation of the forecast standard deviation.

The correlation between the absolute error and the forecast standard deviation (third row in Fig. 8) shows that the ANNs in the IMS perform as well as or slightly worse than the ensemble forecast.

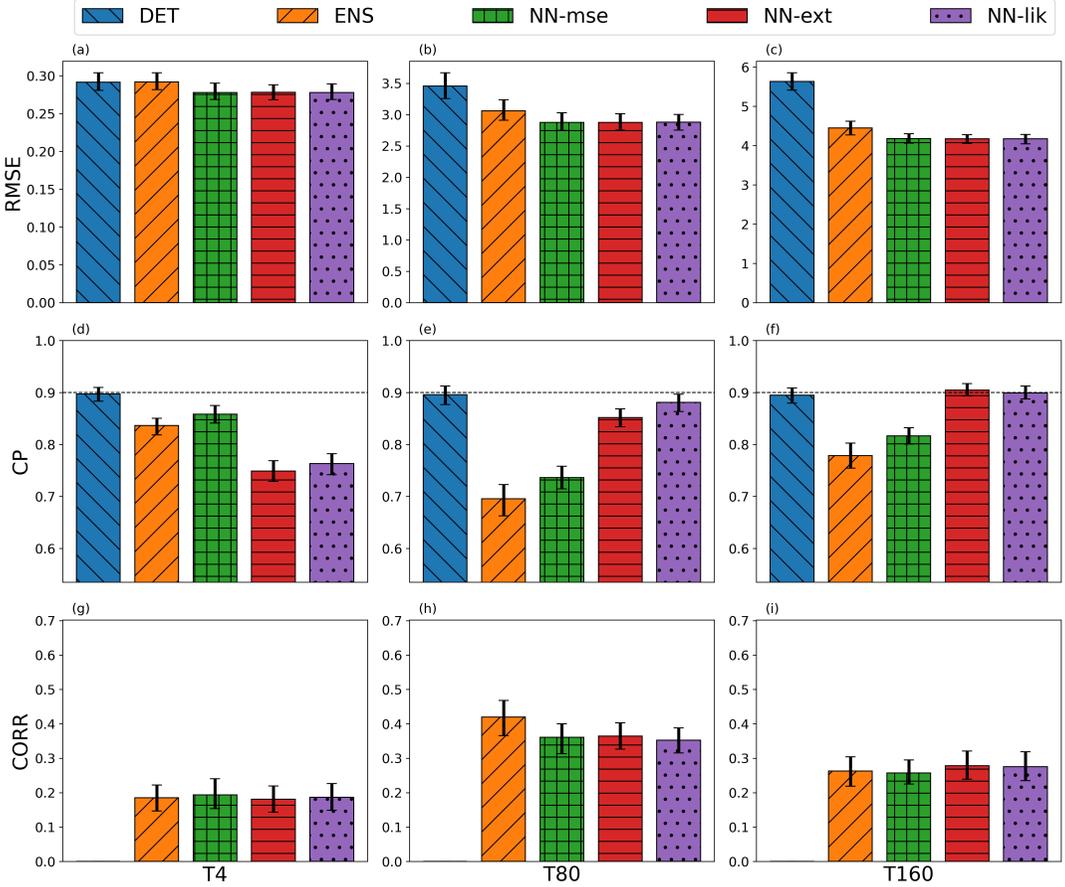


FIGURE 8 As in Fig. 6 but for the imperfect model scenario.

5.4 | ANN performance analysis

To gain more insight into how forecast errors depend on the state variables and to better understand to what extent these dependencies are represented by the ANNs and the ensemble, we examine some aspects of the relationship between state variables and the forecast error. We focus on the IMS experiments discussed in the previous section. Left panels in Fig. 9 show 2D histograms of the forecast error distribution at each particular variable (i.e. $\epsilon_{(i)}$) as a function of the value of the forecasted state variable for the deterministic forecast ($x_{(i)}^d$) and the ensemble mean forecast ($\bar{x}_{(i)}^e$). Since for this simple model all state variables share the same error statistics, results from different state variables are pooled together to construct a single sample. We investigate the relationship between the error and the forecast state at the corresponding variable under two different regimes: One in which the trend at the corresponding state variable is positive (i.e. $\frac{dx_{(i)}}{dt} > 0$) and one in which the trend at the corresponding state variable is negative (i.e. $\frac{dx_{(i)}}{dt} < 0$). This distinction is motivated by the work of Pulido et al. (2016), that shows, in a similar experimental setting, that the magnitude of model errors has a non-Markovian behavior which depends on the time derivative of the state variables. The solid lines over the 2D histogram in Figs. 9a and 9c, represent the mean error as a function

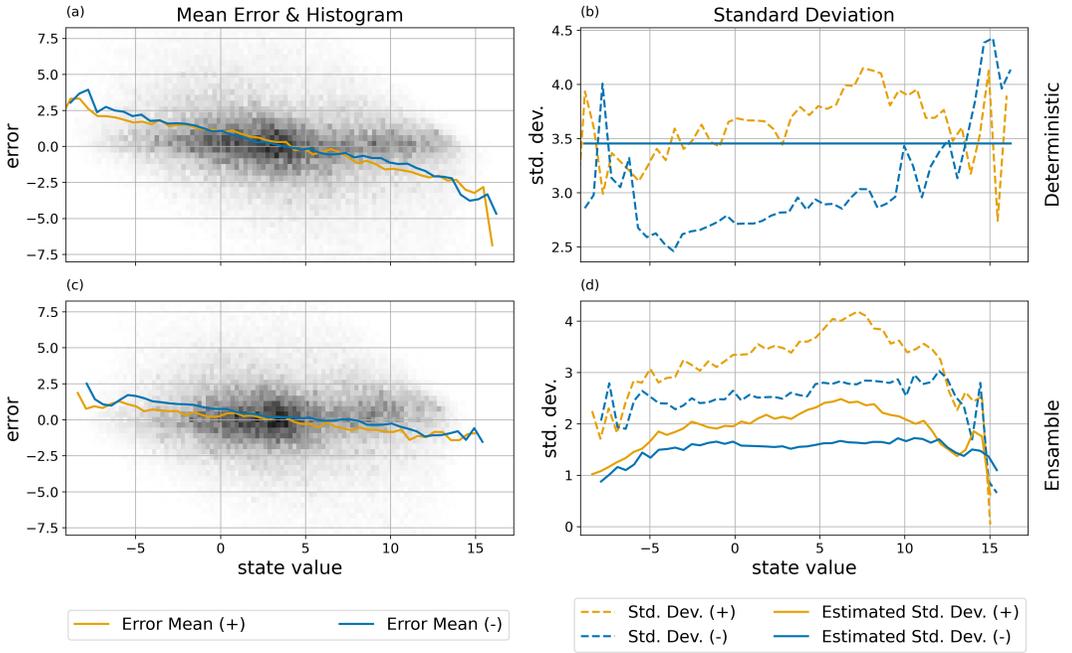


FIGURE 9 Panels (a) and (c) show the 2D density of the forecast error and forecasted state values (grey shades). The yellow and blue solid lines indicates the mean forecast error as a function of the forecasted state value for the positive trend and negative trend sub-samples respectively (see the text for details). Panels (c) and (d) show the standard deviation of the error as function of the forecasted state variable for the positive trend and negative trend sub-samples (dashed yellow line and dashed blue line respectively). The solid yellow and blue lines represents the estimated standard deviation for the positive trend and negative trends sub-samples respectively. Panels (a) and (b) corresponds to the deterministic forecast and (c) and (d) to the ensemble forecast. All panels corresponds to the imperfect model scenario.

of the state value for the positive trend sub-sample (yellow line) and the negative trend sub-sample (blue line). To construct these curves, the range of forecasted state values is divided into 50 bins, and then the mean of the forecast errors corresponding to each bin is computed. This procedure is conducted independently for the positive trend and negative trend sub-samples.

For the deterministic forecast experiment shown in Fig. 9a, the mean error strongly depends on the state variable, although the globally averaged error is close to 0. There is no clear difference in terms of the error magnitude between the positive trend and negative trend sub-samples. The slope of the mean error associated with the ensemble mean forecast (Fig. 9c) is less steep than for the deterministic forecast. This slope is likely associated with the following mechanism: when a high value of the state variable is forecasted, it is more likely to have a negative error than a positive one (related to a subestimation of the value). The opposite is true for low values of x . This produces an asymmetry in the shape of the error probability distribution that becomes more evident for extreme state values. The ensemble mean partially account for this effect, thus reducing the slope of the mean error as a function of x .

Figs. 9b and 9d show the error standard deviation as a function of the forecasted state value (dashed lines). These curves are obtained independently for the positive trend sub-sample (yellow dashed line) and for the negative trend sub-sample (blue dashed line) following a similar procedure as for the mean error curves. The different estimations of

the error standard deviation are shown as solid lines in Figs. 9b and 9d. In this case we compute the average of the estimated error standard deviation at each of the bins covering the forecasted state variable range. This procedure is also applied independently to the positive trend (yellow solid line) and negative trend (blue solid line) independently.

A clear difference between the positive and negative trend sub-samples is found in Fig. 9b. The forecast error standard deviation in the positive trend sub-sample is larger than in the negative trend sub-sample. The difference between the two sub-samples becomes smaller at the extremes of the forecast range. This result is consistent with the findings of Pulido et al. (2016). The forecast error standard deviation also depends on the forecasted value. For example, the standard deviation of the ensemble mean errors for the positive trend sub-sample reaches a maximum approximately around a forecasted state value of 7 (Fig. 9b). The time independent forecast error standard deviation associated with the deterministic forecast (horizontal solid line in Fig. 9b) is, as expected unable to represent the complexity of the dependence between the forecast error standard deviation and the forecasted state. The ensemble spread, on the other hand (Fig. 9c) underestimates the standard deviation of the forecast error but adequately represents its dependence with the forecasted value. The ensemble spread captures both, the difference between the positive and negative trends sub-samples and the dependence on the forecasted state value.

The results for the NN-ext and NN-lik approaches are shown in Figure 10. The mean error of the corrected forecast is almost independent of the forecasted state value for both ANNs. This behaviour comes at the expense of a narrower forecast range. Note that the histograms in Fig. 10a and 10c, span a smaller range in state values than in Fig. 9a and 9c. This means that extreme values are less frequent in the corrected forecasts. A similar behavior is observed for T4 and T160 (not shown) for the IMS and the PMS. In these IMS experiments, Figs. 10b and 10d show that the NN-ext and NN-lik provide a more reliable estimation of the forecast standard deviation. The indirect training strategy allows these networks to accurately estimate the forecast uncertainty including the effect of model errors. They can also capture the relationship between the forecast standard deviation and the forecasted state value. NN-ext performs slightly better at modelling this relationship, particularly for the positive trend sub-sample, in which the local maximum in the error standard deviation is better represented. However, the performance of both networks is more similar for other lead times (not shown).

6 | CONCLUSION AND PERSPECTIVES

In this work we evaluate different strategies for the joint estimation of state-dependent systematic forecast error and its associated uncertainty using machine learning. We compare the performance of artificial neural networks (ANNs) that estimate the forecast uncertainty and that are directly trained from an ensemble forecast or indirectly trained using loss function formulations that implicitly take into account the uncertainty. We implement two different approaches for the implicit training, one based on mean square error loss function (referred to as NN-ext) and the other based on a Gaussian likelihood loss function (referred to as NN-lik). The input to the ANNs is a set of deterministic forecasts at different lead times. We evaluate the performance of the ANN from experiments using a simple dynamical model, and compare them with a deterministic forecast and an ensemble forecast in perfect model (PMS) and imperfect model scenarios (IMS) at different forecast lead times.

Overall the results obtained in this work in a simple proof-of-concept experimental setting are quite encouraging about the possibility of increasing the accuracy of a deterministic forecast and providing a reliable estimation of its uncertainty at a much lower computational cost than that required by an ensemble forecast. In particular, we show that the ANNs are able to filter part of the uncertainty associated to unpredictable modes, particularly in the nonlinear error growth regime. In this aspect, the ANNs are able to emulate the filtering effect of the ensemble mean, although

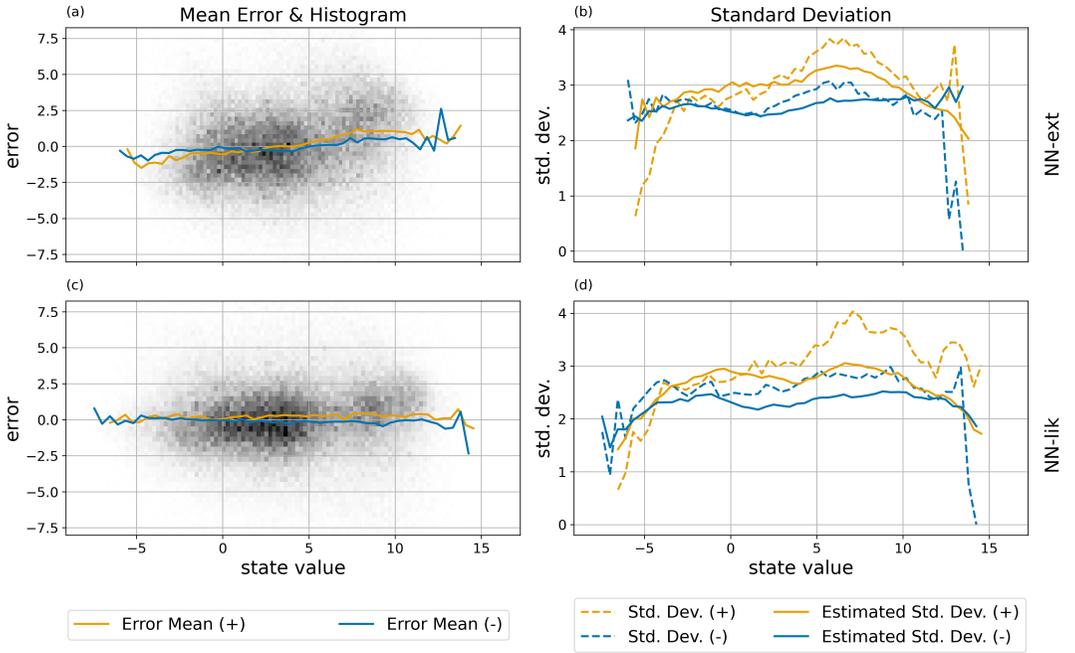


FIGURE 10 As in Fig. 9 for the NN-ext (upper panels, a-b) and the NN-lik approaches (lower panels, c-d).

with less accuracy. The ANNs are also able to accurately represent the state-dependent forecast error standard deviation resulting from initial conditions errors and their amplification due to the chaotic dynamics of the system. A systematic sub-estimation of the forecast standard deviation is found at short lead times in the ANNs due to our training strategy in which the analysis is used as a proxy of the true state of the system.

When the model is imperfect, the ANNs are capable of estimating the state-dependent systematic error component associated with the effect of model errors. The estimated forecast uncertainty by the ANNs also captured this effect. This result is particularly encouraging since a proper representation of model errors and its impact upon the forecast remains a significant scientific challenge.

The network directly trained with an ensemble of forecasts performs better in the perfect model scenario. However, the indirect training approaches (NN-ext and NN-lik) produce better results in the imperfect model scenario. This is mainly because in the direct training, the network inherits the uncertainty estimation achieved by the ensemble. Since model errors are usually sub-represented, this may lead to worse results in comparison to a network which is indirectly trained and able to learn the effect of model errors from the data.

The two loss functions used for indirect training produce similar results in terms of forecast uncertainty quantification. One possible advantage of the mean squared error based training is that no assumptions of the underlying error distribution are required, while the likelihood based loss function assumes that errors are normally distributed. Further investigation is required to conclude if this can result in an advantage of the mean squared error based loss function (e.g. under circumstances in which forecast errors departs significantly from the Gaussian).

One direct application of our results is the development of hybrid data assimilation systems in which ANNs can produce an estimation of the uncertainty of the forecast at a lower computational cost than that required by ensemble forecasts. This is in fact a more challenging scenario since an estimation of the full error covariance matrix is required

in this context which implies a significant increase in the dimensionality of the estimation problem. The development of a hybrid machine-learning based data assimilation approach based on the results presented in this paper is part of a work in progress.

In this work we present a proof of concept based on a simple dynamical model and a low dimension space. Numerical models in geoscience are usually characterized by a very large number of state variables as well as more complex dynamics and error sources. More research is required to investigate how more advanced network architectures (e.g. deep convolutional or recurrent networks) are potentially able to achieve better results in more realistic scenarios, and which strategies are required to reduce the dimensionality of the network input.

references

- Camporeale, E. (2018) Accuracy-reliability cost function for empirical variance estimation.
- Camporeale, E., Chu, X., Agapitov, O. V. and Bortnik, J. (2019) On the generation of probabilistic forecasts from deterministic models. *Space Weather*.
- Candille, G., Côté, C., Houtekamer, P. and Pellerin, G. (2007) Verification of an ensemble prediction system against observation. *Monthly Weather Review*, **135**.
- Carrassi, A., Bocquet, M., Bertino, L. and Evensen, G. (2018) Data assimilation in the geosciences: An overview of methods, issues, and perspectives. *WIREs Climate Change*, **9**, e535.
- Clare, M., Jamil, O. and Morcrette, C. (2021) A computationally efficient neural network for predicting weather forecast probabilities.
- Dugas, C., Bengio, Y., Elisle, F. and Nadeau, C. (2001) Incorporating second-order functional knowledge for better option pricing. *Cirano Working Papers*.
- D'Isanto, A. and Polsterer, K. L. (2018) Photometric redshift estimation via deep learning. *Astronomy & Astrophysics*, **609**, A111.
- Farchi, A., Bocquet, M., Laloyaux, P., Bonavita, M. and Malartic, Q. (2021) A comparison of combined data assimilation and machine learning methods for offline and online model error correction. *Journal of Computational Science*, **55**, 101468. URL: <https://www.sciencedirect.com/science/article/pii/S187750321001435>.
- Grooms, I. (2021) Analog ensemble data assimilation and a method for constructing analogs with variational autoencoders. *Quarterly Journal of the Royal Meteorological Society*, **147**, 139–149.
- Grönquist, P., Ben-Nun, T., Dryden, N., Dueben, P., Lavarini, L., Li, S. and Hoefler, T. (2019) Predicting weather uncertainty with deep convnets.
- Grönquist, P., Yao, C., Ben-Nun, T., Dryden, N., Dueben, P., Li, S. and Hoefler, T. (2021) Deep learning for post-processing ensemble weather forecasts. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, **379**, 20200092.
- Hagedorn, R., Buizza, R., Hamill, T. M., Leutbecher, M. and Palmer, T. N. (2012) Comparing tigre multimodel forecasts with reforecast-calibrated ecmwf ensemble forecasts. *Quarterly Journal of the Royal Meteorological Society*, **138**, 1814–1827.
- Hamill, T. and Whitaker, J. (2006) Probabilistic quantitative precipitation forecasts based on reforecast analogs: Theory and application. *Monthly Weather Review*, **134**.
- Haupt, S. E., Chapman, W., Adams, S. V., Kirkwood, C., Hosking, J. S., Robinson, N. H., Lerch, S. and Subramanian, A. C. (2021) Towards implementing artificial intelligence post-processing in weather and climate: proposed actions from the oxford 2019 workshop. *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, **379**.

- Herman, G. R. and Schumacher, R. S. (2018) Money Doesn't Grow on Trees, but Forecasts Do: Forecasting Extreme Precipitation with Random Forests. *Monthly Weather Review*, **146**, 1571–1600.
- Hunt, B. R., Kostelich, E. J. and Szunyogh, I. (2007) Efficient data assimilation for spatiotemporal chaos: A local ensemble transform kalman filter. *Physica D: Nonlinear Phenomena*, **230**, 112 – 126. Data Assimilation.
- Kalnay, E. (2003) *Atmospheric Modeling, Data Assimilation and Predictability*. Cambridge University Press.
- Kingma, D. P. and Ba, J. (2017) Adam: A method for stochastic optimization.
- Lguensat, R., Tandeo, P., Ailliot, P., Pulido, M. and Fablet, R. (2017) The analog data assimilation. *Monthly Weather Review*, **145**, 4093 – 4107. URL: <https://journals.ametsoc.org/view/journals/mwre/145/10/mwr-d-16-0441.1.xml>.
- Lorenz, E. (1995) Predictability: a problem partly solved.
- Lu, H., Ou, Y., Qin, C. and Jin, L. (2021) A fuzzy neural network bagging ensemble forecasting model for 72-h forecast of low-temperature chilling injury. *Natural Hazards: Journal of the International Society for the Prevention and Mitigation of Natural Hazards*, **105**, 2147–2160.
- Murphy, A. H. (1993) What is a good forecast? an essay on the nature of goodness in weather forecasting. *Weather and Forecasting*, **8**, 281–293.
- Pulido, M., Scheffler, G., Ruiz, J. J., Lucini, M. M. and Tandeo, P. (2016) Estimation of the functional form of subgrid-scale parametrizations using ensemble-based data assimilation: a simple model experiment. *Quarterly Journal of the Royal Meteorological Society*, **142**, 2974–2984.
- Pulido, M., Tandeo, P., Bocquet, M., Carrassi, A. and Lucini, M. (2018) Stochastic parameterization identification using ensemble kalman filtering combined with maximum likelihood methods. *Tellus A: Dynamic Meteorology and Oceanography*, **70**, 1–17.
- Rasp, S. and Lerch, S. (2018) Neural networks for postprocessing ensemble weather forecasts. *Monthly Weather Review*, **146**, 3885–3900.
- Scheffler, G., Ruiz, J. and Pulido, M. (2019) Inference of stochastic parametrizations for model error treatment using nested ensemble kalman filters. *Quarterly Journal of the Royal Meteorological Society*, **145**, 2028–2045.
- Scher, S. and Messori, G. (2018) Predicting weather forecast uncertainty with machine learning. *Quarterly Journal of the Royal Meteorological Society*.
- (2021) Ensemble methods for neural network-based weather forecasts. *Journal of Advances in Modeling Earth Systems*, **13**.
- Tandeo, P., Ailliot, P., Bocquet, M., Carrassi, A., Miyoshi, T., Pulido, M. and Zhen, Y. (2020) A review of innovation-based methods to jointly estimate model and observation error covariance matrices in ensemble data assimilation. *Monthly Weather Review*, **148**, 3973–3994.
- Tandeo, P., Pulido, M. and Lott, F. (2015) Offline parameter estimation using enkf and maximum likelihood error covariance estimates: Application to a subgrid-scale orography parametrization. *Quarterly Journal of the Royal Meteorological Society*, **141**, 383–395. URL: <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.2357>.
- Wang, B., Lu, J., Yan, Z., Luo, H., Li, T., Zheng, Y. and Zhang, G. (2018) Deep uncertainty quantification: A machine learning approach for weather forecasting.