Imperial College
London

Fast ocean data assimilation and forecasting using a neural-network reduced-space regional ocean model of the North Brazil Current

César Andrés Quilodrán Casas

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Physics and
the Diploma of Imperial College, May 2018

Declaration of originality

The work presented in this thesis is all my own, except where referenced. It is the outcome of work that I have done over the course of my studentship, and any collaboration with others to achieve this outcome is specified in the text and Acknowledgements. Some of the work presented in this thesis has been submitted for publication:

C. A. Quilodrán Casas, N. Sparks, R. Toumi. Fast ocean data assimilation using a neural-network reduced-space regional ocean model of the North Brazil Current. *Ocean Modelling.*

Copyright declaration

Abstract

Data assimilation is computationally demanding, typically many times slower than model forecasts. Fast and reliable ocean assimilation methods are attractive for multiple applications such as emergency situations, search and rescue, and oil spills. A novel framework which performs fast data assimilation with sufficient accuracy is proposed for the first time for the open ocean. Speed improvement is achieved by performing the data assimilation on a reduced-space rather than on a full-space. A surface 10km resolution hindcast of the North Brazil current from the Regional Ocean Modelling System (ROMS) serves as the full-space state. The target variables are sea surface height, sea surface temperature, and surface currents. A dimension reduction of the full-state is made by an Empirical Orthogonal Function analysis while retaining most of the explained variance. The dynamics are replicated by a state-of-the-art neural network trained on the truncated principal components of the full-state. An Ensemble Kalman filter assimilates the data in the reduced-space, where the trained neural network produces short-range forecasts from perturbed ensembles. The Ensemble Kalman filter of the reduced-space is successful in reducing the root mean squared error by $\sim 45\,\%$ and increases the correlations between state variables and data. The performance is similar to other full-space data assimilation studies. However, the computations are three to four orders of magnitude faster than for other full-space data assimilation schemes.

The forecast of ocean variables is a computationally demanding task in terms of speed and accuracy. This framework manages to create fast forecasts in ~ 30 seconds, once data have been assimilated. The forecasts are obtained using the trained neural network. We performed additional experiments using data and forecasts from July 2015 and January 2016. The analysis and forecasts in our framework yield a higher skill score and high spatial correlation when compared to the operational dataset Global Ocean Physics Analysis and Forecast by the UK MetOffice. Forcing the neural network with 10 m surface winds in order to improve the total surface currents forecast was considered. There is no additional skill in the forecasts using wind forcing because of the low Ekman component compared to the dominant geostrophic currents. The reduced model approach could be a useful tool when full physics regional models are not available to make a forecast.

Acknowledgements

This research was supported by BECAS CHILE, a governmental Chilean scholarship from the Comisión Nacional de Investigación Científica y Tencnológica (CONICYT).

I would like to thank my supervisor Ralf Toumi for his guidance and patience, and for giving me enough independence and support throughout this journey. I would also like to thank Nathan Sparks for being a friend and a bottomless bucket of statistical knowledge.

To my colleagues in SPAT and my 708 office-mates for being the most amicable and supportive group of people that I have come across.

To my flatmates, for proofreading, putting up with me, for their infinite patience and their ever enjoyable company.

To my parents and family, who have always been and always will be my support even if we are 7682 miles away.

Finally, I would like to thank my beautiful wife, Marissa. You taught me how to smile and every day you keep on showing me a thousand more reasons why I should not stop doing it. You kept me going when I wanted to give up.

A mi madre, Graciela y a mi padre, Reinaldo.

Contents

List of Tables

List of Figures

Acronyms

CFSR Climate Forecast System Reanalysis

CMEMS Copernicus - Marine Environment Monitoring Services

DA+NN The framework used in this study

DA Data assimilation

DJF December, January, March

EnKF Ensemble Kalman Filter

EOF Empirical Orthogonal Functions

ITCZ Inter Tropical Convergence Zone

JJA July, August, September

MAM March, April, May

MDT Mean Dynamical Topography

MLP Multi layer perceptron

NBC North Brazil Current

NECC North Equatorial Countercurrent

NN Neural network

PC Principal Components

RMSE Root Mean Squared Error

RMS Root Mean Square

ROMS Regional Ocean Modelling System

SEC South Equatorial Current

SON September, October, November

SSH Sea surface height

SST Sea surface temperature

SVD Singular Value Decomposition

UKMO UK MetOffice reanalysis and forecast profuct

U Eastward component of surface velocities

V Northward component of surface velocities

Chapter 1

Introduction

1.1 Background and motivation

Heinlein (1987) expressed: "Climate is what we expect and weather is what we get". Climate is considered a collection of the long term statistics of weather (Lorenz 1970), where its variability is the result of nonlinear interactions with several degrees of freedom (Hannachi et al. 2007). In the past few decades, there have been major advances in the development of climate models. The numerical model of the global climate system is arguably one of the most important developments in climate sciences (Palmer 1999). These models have transformed the daily weather prediction, where 3 to 10-day forecast skill has been increasing by about 1 day per decade; the impact of El Niño is successfully predicted 3-4 months in advance; and projections of climate change due to anthropogenic causes provide policy makers the tools for major international protocols (Palmer 2001, Bauer et al. 2015). These models help us understand the current state and forecast future states of the climate. Classical predictability, as Pohlmann & Keenlyside (2004) explains, performs ensemble experiments with a single coupled model perturbing only the initial conditions. Before performing the forecast model, it is essential to obtain information from a range of current observations and input them into the numerical model, in order

to produce the best possible estimate of the initial state of the system (Swinbank et al. 2012).

Making predictions of ocean variables is particularly challenging. It is difficult to obtain and maintain observations of the state of the ocean because satellite information is limited to the surface layer (Anderson et al. 1996). Moreover, ocean models present high-dimensionality and high-non-linearity. Despite the high-dimensionality and non-linearity, one proven way to improve ocean models representation is to blend information from observations with numerical models using Data Assimilation (DA). DA is arguably the best practicable technique available to make use of the large variety and volume of data accessible to numerical models. However, the application of DA can be very computationally demanding and costly. Speeding up the DA process can be very beneficial in emergency situations in the ocean where a fast forecast of the future states are needed. This research aims to employ a fast DA framework to a regional ocean setup and study ocean weather patterns of the North Brazil Current (NBC) over the North Brazil - Trinidad and Tobago area.

The improvement, compared to previous studies in the area, is to employ the fast DA and fast forecasts of surface ocean variables. Fast DA is achieved by using a model surrogate instead of a full model. This model surrogate is a projection of the full space into a reduced dimensionality space with further training by using Artificial Neural Networks (NNs). In consequence, the DA and forecasts are less computationally expensive.

1.1.1 The Lorenz attractor

The behaviour of the Earth's atmosphere, governed by physical laws, can be expressed as a system of differential equations (Kocarev et al. 2009). The problem of weather forecasting, as expressed by Lorenz (1969), is identified with the discovery, through different means, of a particular solution of these differential equations whose initial conditions define the

atmosphere's current state. The Lorenz attractor was introduced in 1963, in a remarkable paper, by Lorenz (1963) and is defined by the following three-dimensional non-linear system of differential equations:

$$\begin{aligned} \frac{dx}{dt} &= \sigma(y - x) \\ \frac{dy}{dt} &= x(\rho - z) - y \\ \frac{dz}{dt} &= xy - \beta z \end{aligned} \tag{1.1}$$

with $\sigma \equiv \frac{\nu}{\kappa}$, $\rho \equiv \frac{Ra}{Ra_c}$ and $\beta \equiv \frac{4}{1+a^2}$, all real positive parameters, where σ is the Prandtl number[1], Ra is the Rayleigh number[2], Ra_c is the critical Rayleigh number, and β is a geometric factor (Tabor 1989, Weisstein 2015). When studying two-dimensional flow of fluid, representing a basic model of atmospheric dynamics, Lorenz accidentally discovered the chaotic behaviour of this system. These equations led Lorenz to the discovery of sensitive dependence of initial conditions, which is an essential factor of unpredictability in many systems. This experiment represented one of the earliest discoveries of the phenomenon known as the Butterfly effect (Kocarev et al. 2009).

The original derivation of these differential equations can be briefly described as the study of a two-dimensional cell of fluid being cooled from above and warmed from below, with a convective motion described by a partial differential equation. The variable x measures the convective overturning rate, whereas y and z represent the horizontal and vertical temperature variation, respectively (Sparrow 2012).

Notice that eq. (1.1) is a deterministic set, therefore contains no noisy, random or stochastic terms, and it will determine an unique flow. The classical set of parameter values chosen by Lorenz are $\sigma = 10$, $\beta = 8/3$ and $\rho = 28$ and suggests that almost all points in phase space tend to the Lorenz attractor (Sparrow 2012). A minimal change in these parameters may lead to chaotic dynamics. As a result, chaotic systems can be defined as

[1] Prandlt number is the ratio between momentum diffusivity and thermal diffusivity
[2] Rayleigh number is a dimensionless number associated with buoyancy-driven flow

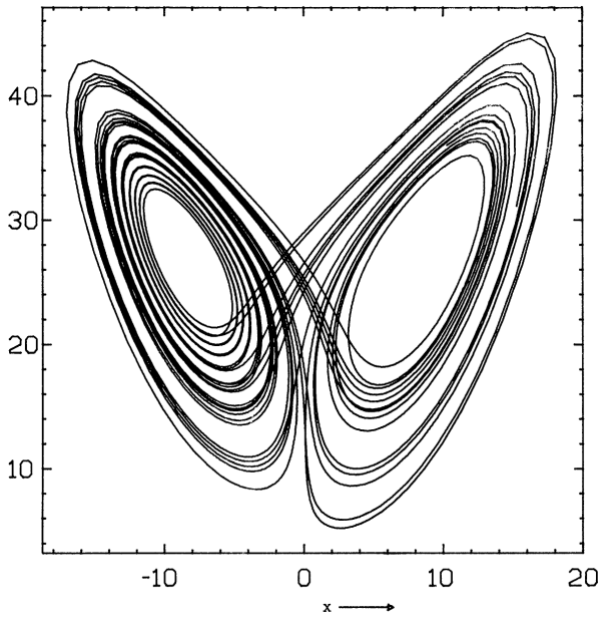

Figure 1.1: Computationally obtained solutions to the Lorenz set of equations, projected into the xz plane, using the classical parameters $\sigma = 10$, $\beta = 8/3$ and $\rho = 28$. From Sparrow (2012).

dynamical systems that defy synchronisation (Pecora & Carroll 1990).

1.2 Reduced models

Despite climate being the long-term statistics of weather (Hannachi et al. 2007), as aforementioned climate is still characterised by non-linearity and high dimensionality. Consequently, and as expressed by Hannachi (2004) a challenging task is to find ways to reduce the dimensionality of the system, allowing us to identify the most important patterns explaining its variations. In the past decades, numerous studies have been performed with the aim of recognising patterns or to reduce the dimensionality of complex models like Thompson & Wallace (1998, 2000), Ångström (1935), Bjerknes (1969), Wallace & Gutzler (1981), Frolov et al. (2009), and Weare & Nasstrom (1982) among others.

Because of its simplicity and analytic derivation, Empirical Orthogonal Functions (EOFs) (Obukhov 1947, Fukuoka 1951, Lorenz 1956) have become a widely popular statistical tool

in atmosphere, ocean and climate science (Kutzbach 1967, Preisendorfer & Mobley 1988). The EOF analysis is also known as Principal Component Analysis (PCA) or Karnuhen-Loève basis functions in mathematical literature. Analysis of EOFs can be used to identify structures in geophysical data (observations or general circulation models) (Monahan et al. 2009) which hold a large part of the variance in the study field. Nevertheless, its caveat remains on the physical interpretability of its results due to the strong constraints satisfied by the EOFs: the orthogonality in space and time (Hannachi 2004). Physical modes are not in general orthogonal, since they are not uncorrelated (Simmons et al. 1983, Hannachi et al. 2006). Essentially, this method is exploratory (i.e., non-model orientated). Empirical Orthogonal Functions are often used in dimensionality reduction and patterns extraction (Hannachi et al. 2007). As stated by North et al. (1982): "We cannot know the exact EOFs for climate but must be satisfied with estimates of them based upon a finite number of independent realisations of the instantaneous state of the field".

1.2.1 Empirical orthogonal functions to gridded data

In order to obtain the EOFs, we need to manipulate the data. The following procedure works as an example. Gridded data in a space-time field is simpler and less memory consuming when projected and represented by a two-dimensional array $X(t, s)$:

$$X = \begin{pmatrix} x_{1,1} & x_{1,2} & \cdots & x_{1,p} \\ x_{2,1} & x_{2,2} & \cdots & x_{2,p} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n,1} & x_{n,2} & \cdots & x_{n,p} \end{pmatrix} \tag{1.2}$$

where the value of X at discrete time step t_i and grid point s_j is denoted by x_{ij}, with $i = 1, \ldots, n$ and $j = 1, \ldots, p$. The value of the field at time t is represented by $\bar{x} =$

$(\bar{x}_{t1}, \cdots, \bar{x}_{tp})$. The time average of X at the i^{th} grid point (i.e., climatology of the field) is:

$$\bar{x}_i = \frac{1}{n} \sum_{k=1}^{n} x_{ki} \tag{1.3}$$

The anomaly field (Hannachi et al. 2007), or departure from the climatology, can be expressed in both linear and matrix form as:

$$x'_{ts} = x_{ts} - \bar{x}_s; X' = X - 1_n \bar{x} = \left(I - \frac{1}{n} 1_n 1_n^T \right) X \tag{1.4}$$

where $\mathbf{1_n} = (1, \ldots, 1)^T$ is a column vector with n ones and I is the $n \times n$ identity matrix. Once the anomaly field matrix eq. (1.4) is obtained, the covariance matrix is then given by:

$$S = \frac{1}{n} X'^T X' \tag{1.5}$$

which contains the covariances between the time series of X' at any pair of grid points. The objective of an analysis through EOFs is to find the linear combination of all the grid points that successively have maximum variance (Hannachi 2004). Therefore, it is necessary to obtain a direction vector $u = (u_1, \ldots, u_p)$ that maximises the variability of $X'\mathbf{u}$.

The EOFs are the eigenvectors of S sorted in decreasing order of eigenvalues, and are derived as the solution of:

$$S\mathbf{u} = \lambda^2 \mathbf{u} \tag{1.6}$$

where λ is the eigenvalue. Singular Value Decomposition (SVD), a linear algebra tool,

allows us to decompose any $n \times p$ matrix in:

$$X' = A\Lambda U^T \tag{1.7}$$

where $A^T A = I$ and $U^T U = I$ (i.e., A and U are unitary matrices) with dimensions $n \times r$ and $r \times p$, respectively, with $r \leq min(n, p)$; and $\Lambda = diag(\lambda_1, \ldots, \lambda_r)$ with diagonal elements (i.e., singular values of X) $\lambda_1 \geq \lambda_2 \geq, \ldots, \lambda_r \geq 0$.

A more rigorous explanation on the formulation and computation of EOFs can be found in Hannachi et al. (2007).

1.2.2 Artificial Neural Networks

Our framework also uses NNs, which are mathematical tools inspired by the biological nerves system which aim to model the functionality of the human brain (Knutti et al. 2003). The basic unit is the artificial neuron and the network consists of a set of interconnected neurons by links and weights.

The NNs field of research is relatively young, with the work of McCulloch & Pitts (1943) being one of the first to develop this technique. Nevertheless, it was not until the early 1980's that NNs have been extensively applied. New discoveries, along with increased computational capacities, have renewed the interest in the NNs field (Knutti et al. 2003, Goodfellow et al. 2016).

They are now utilised to solve, successfully, problems like pattern recognition and classification, tracking, signal processing, fault detection, data compression, etc. where they may have some advantages over the traditional techniques (Hippert et al. 2001, Goodfellow et al. 2016, LeCun et al. 2015). Applications of NNs in atmospheric science include prediction, function approximation and pattern classification. Artificial NNs have been used in prediction of air-quality (Yi & Prybutok 1996, Prybutok et al. 2000); atmospheric

sulphur dioxide concentrations (Boznar et al. 1993); ozone forecasts (Comrie 1997); tornadoes (Marzban & Stumpf 1996, Santoso et al. 2015); Indian monsoons (Navone & Ceccatto 1994, Dwivedi & Pandey 2011); Brazilian rainfall anomalies (Hastenrath & Greischar 1993, Greischar & Hastenrath 2000), among others. In the field of function approximations NNs have been used to model the relationship between hourly ozone concentrations and local meteorological variables (Gardner & Dorling 1996), as well as nonlinear transfer functions (e.g., in the retrieval of surface wind speed from microwave sensors) (Krasnopolsky et al. 1995). In pattern classification, (Kim et al. 2011) used NNs for classification of typhoon tracks, meanwhile (Peak & Tag 1992) utilised NNs method to classify clouds.

Artificial NNs are skilful tools in extracting patterns from past events and extrapolating these into future time steps (Knutti et al. 2003). There are two reasons for this. Firstly, NNs are capable of numerically approximate any given continuous function with required accuracy. They are also expected to perform better than traditional linear models when modelling complex nonlinear relationships. Secondly, NNs are driven by data, therefore it is not mandatory to provide either model or parameters, just using input and output vectors NNs can find, learn, and store the relationship between them. This learning is particularly helpful when substantial quantities of data are available, but there is little previous knowledge about the laws that govern the system (Hippert et al. 2001). Artificial NNs are able to learn from and generalise from experience (Zhang et al. 1998).

There are many different types of NNs, each with a specific application. These NNs are characterised by their differing architectures (Dibike & Coulibaly 2006) (i.e., the way the neurons are organised). Here the multi-layer perceptron (MLP) type will be described, in which the neurons are organised in layers. The layers are connected, generally, in a feed-forward fashion with no connections between neurons in the same layer. If the feed-forward architecture (feed-forward network) is utilised the outputs of one layer are inputs into the next one (Hippert et al. 2001). An N-layer feed-forward network consists

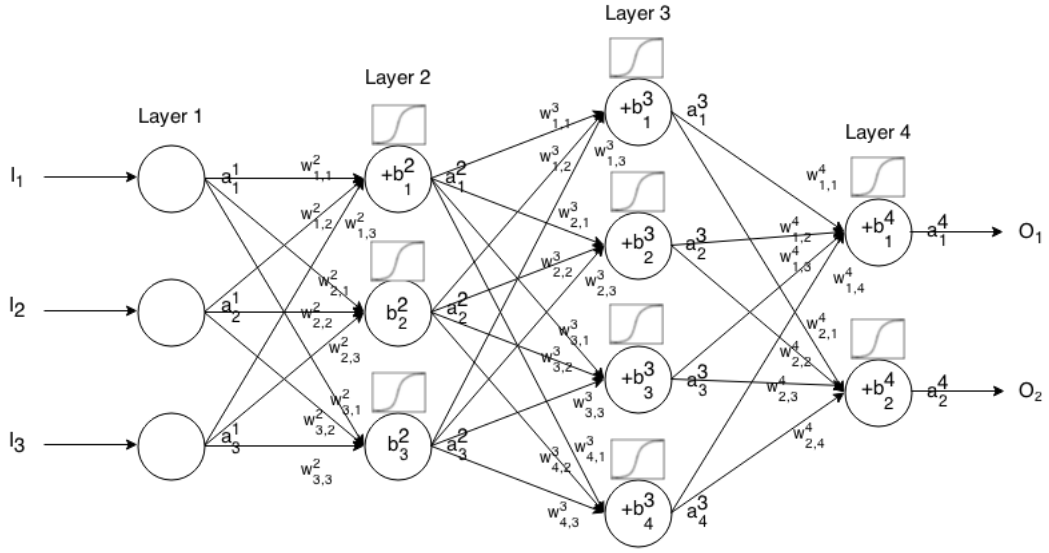

Figure 1.2: Schematic example of a feed-forward NN with two hidden layers. The network has 3 input nodes, 3 and 4 hidden nodes in the hidden layers and 2 output nodes. From Nielsen (2015).

of an input layer, $N-2$ hidden layers, and an output layer, and it is shown in Fig. 1.2. The neuron receives information through the input nodes and produces a response. The processing is done by linearly combining the inputs and subsequently using the result as a nonlinear activation function. The activation functions are nondecreasing and differentiable. The linear function $y = x$ and the sigmoid $y = 1/(1+e^{-x})$ are among the most popular.

The NNs go through a training process, where the weights w_i are estimated by minimising a cost function, so that the network is able to perform the required task. It is important to select an appropriate set of training data. Too few or poorly selected training datasets will prevent the ANN to learn the input-output relationship (Knutti et al. 2003). The contrary happens when a large dataset is considered, since the size of the dataset may increase the computational costs. The back-propagation algorithm was the first training algorithm to be discovered and it computes the gradient of the cost function with respect to the NN parameters, therefore as stated above, the activation functions must be differentiable. The conjugate gradient is a widely used training algorithm for its efficiency and fast convergence toward the minimisation of the cost function.

From Fig. 1.2, the value of the j^{th} node at the i^{th} layer a_j^i is defined by:

$$a_j^i = \sigma(\sum_k (w_{jk}^i \cdot a_k^{i-1}) + b_j^i) \tag{1.8}$$

where σ is the activation function; w_{jk}^i is the weight of the j^{th} node of the $(i-1)^{th}$ layer to the k^{th} neuron of the i^{th} layer; b_k^i and a_k^i represent the bias and activation value, respectively, of the k^{th} neuron in the i^{th} layer.

The cost function (or loss function, error function) is a measure of how well the network performs at predicting the targeted values. The cost function C depends on:

$$C(W, B, S^r, E^r) \tag{1.9}$$

where W are the NN weights, B are the NN biases, S^r are the inputs of a single training sample, and E^r is the target value or desired output of that training sample. The cost function is used in the backpropagation to compute the error of the output layer, δ^L like:

$$\delta^L = \frac{\partial C}{\partial a_j^L} \sigma'(z^i) \tag{1.10}$$

The cost function used throughout this study is the quadratic cost function or Mean Squared Error (MSE), defined as:

$$C_{MSE}(W, B, S^r, E^r) = \frac{1}{2} \sum_j (a_j^L - E_j^r)^2 \tag{1.11}$$

The NN training is tested over a dataset that has not been shown to the network before. The dataset is used only for testing the final solution in order to confirm the actual predictive power of the network.

The sensitivity of an ANN is defined by its size (i.e., the number of neurons in the input and hidden layers) (Dibike & Coulibaly 2006, Knutti et al. 2003). A small network

will have too few degrees of freedom and it will perform poorly in the approximation of the input-output relationship. If the network is too large, the computational costs will increase and furthermore can lead to *overfitting*. A way to prevent *overfitting* is a method called *early stopping*, which stops the training process before *overfitting* starts (Hippert et al. 2001, Knutti et al. 2003).

There is a large number of different feed-forward designs. One of particular interest in this research is the autoregressive NN. The general approach used in the NNs field is to use historical data from time series to train the NN and test it with new values (Rojas 2013). As explained by Dibike & Coulibaly (2006) on temporal problems, measurements from physical systems are no longer an independent set of input samples, but functions of time from where the ANN must have access to. These functions of time add a memory structure to the network, and its size depends on the number of past samples. The delay, or lag, with memory depth k can be represented by $x(n) = [x(n), x(n-1), \ldots, x(n-k+1)]$. The autoregressive approach has been used in climate science by Frolov (2007), Kim & Valdés (2003), Trigo & Palutikof (1999), among others.

1.3 Data assimilation

Numerical weather prediction is an initial/boundary value problem. Given an estimate of the initial conditions, present state of the atmosphere, and suitable boundary conditions, the model will forecast the time evolution of the atmosphere. Data assimilation is the procedure that produces initial conditions based on a "statistical combinations of observations and short-range forecast" (Kalnay 2003). It uses all the available information to determine, as accurately as possible, the atmospheric or oceanic state (Talagrand 1997). Several reviews (Ghil & Malanotte-Rizzoli 1991, Talagrand 1997, Ghil 1997, Zupanski & Kalnay 1999) and textbooks (Daley 1991, Kalnay 2003, Swinbank et al. 2012) have been written on DA.

In early stages of numerical weather prediction, Richardson (1922) and Charney et al. (1950) implemented interpolations and digitisation by hand of observations into a regular grid. Since these methods were extremely time consuming, new interpolation methods were shortly developed (Charney & Eliassen 1949, Panofsky 1949).

1.3.1 Kalman filter

The Kalman Filter (Kalman 1960) is an estimator for the problem of approximating the instantaneous state of a linear dynamic system with linearly related measurements, both perturbed by white noise (Grewal & Andrews 2014), based on the relative magnitudes of the covariances of both, the model system state estimate and the observations (Swinbank et al. 2012). The advantage of the Kalman filter is that the entire system is updated with covariances and thus representing the reliability of the observations and model prediction (Houser 2003). The equations for evolving the system states, in a series of forecasting and update steps, must be in a linear space. The following description is based on Mandel (2009). The probability density function (pdf) of a n-dimensional state vector $\mathbf{x}$ with a Gaussian distribution with mean μ and covariance Q is

$$p(\mathbf{x}) \; \alpha \; \mathbf{exp}(-\frac{\mathbf{1}}{\mathbf{2}}(\mathbf{x}-\mu)^{\mathbf{T}}\mathbf{Q}^{-\mathbf{1}}(\mathbf{x}-\mu)) \tag{1.12}$$

where α means proportional. This probability distribution is called the prior and it is evolved in time by running the model in order to account for new data $\mathbf{d}$. It is necessary that the data comes with an associated error distribution. Data $\mathbf{d}$ is assumed to have also a Gaussian distribution, mean $\mathbf{Hx}$ ($\mathbf{H}$ is called the observation matrix), and covariance $\mathbf{R}$. The covariance $\mathbf{R}$ is an estimate of the error of the data. If the errors of $\mathbf{d}$ are independent then $\mathbf{R}$ is diagonal. The diagonal elements of $\mathbf{R}$ are the squares of the standard deviation of the errors associated to $\mathbf{d}$. The pdf $p(\mathbf{d}|\mathbf{x})$ of $\mathbf{d}$ conditional to $\mathbf{x}$ is:

$$p(\mathbf{d}|\mathbf{x})\ \alpha\ \mathbf{exp}(-\frac{\mathbf{1}}{\mathbf{2}}(\mathbf{d}-\mathbf{Hx})^{\mathbf{T}}\mathbf{R}^{-\mathbf{1}}(\mathbf{d}-\mathbf{Hx})) \tag{1.13}$$

Using the Bayes theorem and eq. (1.12) and eq. (1.13) we can obtain the pdf of the system state $\mathbf{x}$ conditional to $\mathbf{d}$, called the posterior:

$$p(\mathbf{x}|\mathbf{d})\ \alpha\ \mathbf{p}(\mathbf{d}|\mathbf{x})\mathbf{p}(\mathbf{x}) \tag{1.14}$$

If we denote the posterior state and posterior pdf by $\hat{\mathbf{x}}$ and $p(\mathbf{x}\hat{|}\mathbf{d})$, the equation can be re-written as:

$$p(\hat{\mathbf{x}})\ \alpha\ \mathbf{exp}(-\frac{\mathbf{1}}{\mathbf{2}}(\hat{\mathbf{x}}-\hat{\mu})^{\mathbf{T}}\hat{\mathbf{Q}}^{-\mathbf{1}}(\hat{\mathbf{x}}-\hat{\mu})) \tag{1.15}$$

where:

$$\begin{aligned}&\hat{\mu}=\mu+\mathbf{K}(\mathbf{d}-\mathbf{H}\mu),\ \hat{\mathbf{Q}}=(\mathbf{I}-\mathbf{KH})\mathbf{Q}\\&\mathbf{K}=\mathbf{QH}^{T}(\mathbf{HQH}^{T}+\mathbf{R})^{-1}\end{aligned} \tag{1.16}$$

come from the Kalman update formulas and K is the Kalman gain.

When the system, and therefore the equations, are non-linear, then the filter is called extended Kalman filter (EKf) (Kalnay 2003). The EKf is an approximation, by linearisation, of the nonlinear system.

1.3.2 Ensemble Kalman filter

The Ensemble Kalman filter (EnKF) was introduced by Evensen (1994) and is an alternative to the extended Kalman filter. The EnKF is a Monte Carlo approximation of the Kalman filter and avoids evolving the covariance matrix of the pdf of $\mathbf{x}$. The EnKF is based on an ensemble of states, which determine the initial state if it is not known and assumes a Gaussian distribution. The covariance is obtained when the ensemble of states is updated inside the filter. Gauthier (2003) explains that the EnKF is a practical way to implement Kalman filter in complex models without having to develop the adjoint

of a numerical model. Since the EnKF uses a Monte Carlo approach it is a subject of debate about how many ensembles are required (Swinbank et al. 2012). The distribution is represented by ensembles:

$$\mathbf{X} = [\mathbf{x}_1, \cdots, \mathbf{x}_N] = [\mathbf{x}_i],\ \mathbf{D} = [\mathbf{d}_1, \cdots, \mathbf{d}_N] = [\mathbf{d}_i] \tag{1.17}$$

where $\mathbf{X}$ is the prior ensemble and $\mathbf{D}$ is the data matrix where $\mathbf{d}_i$ are the data vector $\mathbf{d}$ plus a perturbance from a Gaussian distribution with zero mean and covariance $\mathbf{R}$. Then, the posterior state $\hat{\mathbf{X}}$ can be expressed as:

$$\hat{\mathbf{X}}=\mathbf{X} + \mathbf{K}(\mathbf{D} - \mathbf{HX}),\ \mathbf{K} = \mathbf{QH}^T(\mathbf{HQH}^T + \mathbf{R})^{-1} \tag{1.18}$$

Finally, the EnKF can be obtained by replacing $\mathbf{Q}$ by $\mathbf{B}$ which is the covariance computed from the ensemble members.

Some advantages of the EnKF over the extended Kalman filter include a) less computational cost; b) it does not require a linearisation of the evolution of the forecast error covariances; and c) it may provide excellent initial perturbations (Kalnay 2003).

1.4 Study Area

The North Brazil Current (NBC) (Fig. 1.3) was chosen as the testbed for our fast DA scheme, because it is a complex eddy-dominated area. The eddies dominating the area are the retroflection of the NBC into the North Equatorial Counter Current (NECC) and the anticyclonic eddies that are detached from the retroflection of the NBC into the NECC. These anticyclonic rings are called NBC rings. Several studies have been performed on the NBC (Johns et al. 1990, Fratantoni et al. 1995, Flagg et al. 1986, Didden & Schott 1993, Fratantoni & Richardson 2006, Johns et al. 2003, Cochrane et al. 1979, Newinger & Toumi 2015, Jochum & Malanotte-Rizzoli 2003). The NBC is a strong low-latitude

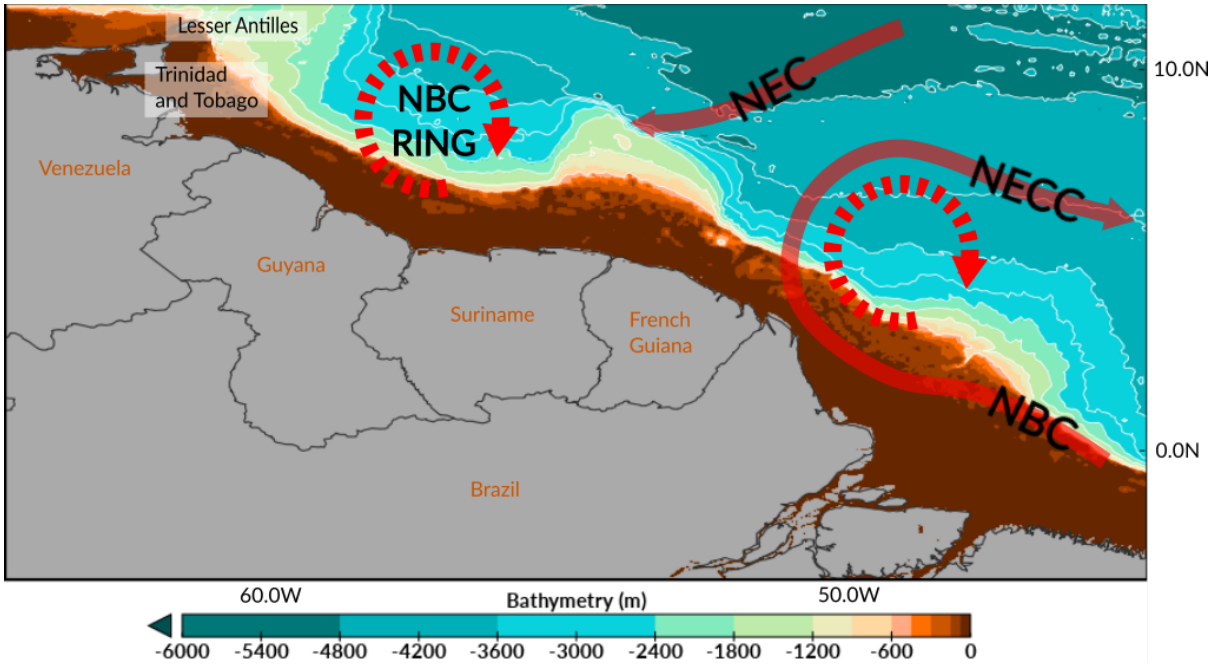

Figure 1.3: Schematic representation of the North Equatorial Current (NEC), North Equatorial Countercurrent (NECC), North Brazil Current (NBC) and eddy generation (NBC rings) at the western boundary of the tropical Atlantic with warm-water pathways in solid red.

western boundary current along the East and North coast of South America carrying water northward supplied by the South Equatorial Current (SEC) from the tip of Brazil near 5°S across the Equator and into the northern hemisphere (Didden & Schott 1993). Low-latitude western boundary currents are vital in the transport across the Equator and recirculation of water within zonal equatorial systems (Johns et al. 1990, Fratantoni et al. 1995). The NBC has a width between 100 and 200 km, with a 500 m depth extension of its flow, and total volume transport of around 50 Sv ($1Sv = 10^6 m^3 s^{-1}$) across 2°N (Flagg et al. 1986). During June, July and August (JJA) and September, October and November (SON), the NBC surface flow drifts away from the coast at about 6°N and feeds the eastward NECC. There is observational evidence for seasonal flow variability above the thermocline that connects the NECC, NBC and SEC (Csanady 1985). The separation process of the NBC is more complex and greatly layered. Flow underneath the thermocline splits apart near 3°N and feeds an eastward thermocline current, while the flow in the thermocline separates just north of the Equator and supplies the Equatorial

undercurrent (Johns et al. 1990).

Observational data suggest that the NBC rings are large long-lived (100 days), anticyclonic rings that contribute to the mean circulation (Johns et al. 1990). These features appear between 4°N and 8°N (Bruce et al. 1985). The detachment of these rings is observed during November through January where the NBC retroflection into the NECC weakens. Figure 1.3 shows a schematic diagram of these eddies along with the NBC, NECC, and North Equatorial Current. These rings can exceed 450 km in overall diameter with strong azimuthal velocities (Fratantoni & Richardson 2006). Annually, 5 to 6 NBC rings are formed around 50°W and travel to the southeastern Caribbean, near the Lesser Antilles, where they dissolve. In the context of the larger North Atlantic circulation, the transport of the NBC rings has been estimated as 13 - 15 Sv of warm water transport to close the Atlantic Meridional Overturning Circulation (Johns et al. 2003).

1.5 Thesis plan

Using the NBC as a testbed, this study develops a fast data assimilation and forecasting framework for ocean surface variables. The framework and numerical model used is thoroughly explained in Chapter 2 and is based on Frolov et al. (2009). The fast framework and model is used in the three results Chapters (3, 4, and 5) with different applications.

In Chapter 3, we assess the fast data assimilation framework by assimilating data of SSH and SST to a model that includes SSH, SST, U and V. The computing times against similar full-space data assimilation systems are reported, and we compared the fast DA performance against the ROMS model solution performance, a persistence experiment and a least squares overfit. These experiments set up a benchmark for estimating a set of optimal initial conditions for forecasting.

Chapter 4 goes beyond the results of the previous chapter and adds an extra step to

the framework: forecasting. Here, we compare our data assimilation performance and forecasts against an operational product. The comparison is valid for July 2015 and January 2016. Different combinations of assimilating and including SSH, SST, U and V in our model are tried.

The framework is modified a step further in Chapter 5 to allow the input of external forcing from near-surface winds. Using the data assimilation and forecast of U and V in the previous chapter as benchmark, winds from a reanalysis product are used in order to add an additional physical constraint for the prediction of surface currents.

Finally, the key findings are summarised in Chapter 6 along with future work.

Chapter 2

Methods and data

2.1 ROMS Model

To undertake this research project, the ocean model ROMS (Regional Ocean Modelling System) will be utilised. ROMS is a three-dimensional, free-surface, and primitive equations ocean model and it is comprehensively described by Shchepetkin & McWilliams (2005). The Regional Oceanic Modelling System is a state-of-the-art "time-stepping and mode-splitting suitable for a high-resolution, free-surface, terrain-following coordinate oceanic model" (Shchepetkin & McWilliams 2005), capable of high resolution descriptions of basin-wide and coastal oceanic flows (Di Lorenzo et al. 2007). ROMS had been used broadly to study the dynamics of ocean circulation over a varied range of locations such as the North Atlantic Basin (Haidvogel et al. 2000, Malanotte-Rizzoli et al. 2000), California Current System (Di Lorenzo & Ohman 2013, Marchesiello et al. 2003), Arctic (Budgell 2005) and Antartic Sea (Robertson et al. 2003), among others. Previous river and inner-shelf circulation studies have used ROMS (Warner et al. 2005, Choi & Wilkin 2007, White & Toumi 2014). ROMS uses the Arakawa C-grid, where the velocities are calculated at the grid boundaries and the variables of state are in the middle of the grid box (Newinger 2015). An extensive explanation of the following ROMS model setup can

be found in Newinger (2015) and Newinger & Toumi (2015).

The domain of the ROMS model (Fig. 1.3) is enclosed by 2.5°S - 11.5°N and 63.5°W - 46°W, and covers about 2000 km × 1500 km including the mouths of the Amazon and Orinoco Rivers. The domain area covers: North-eastern South America including Brazil, French Guyana, Suriname, Guyana and Venezuela, and the Lesser Antilles. The model has a horizontal resolution of 10 km and 50 vertical levels.

2.1.1 Ocean and atmospheric boundary conditions

The monthly mean 0.25-degree Mercator Ocean (GLORYS2V3) global ocean reanalysis (Lellouche et al. 2013) provides the open boundary and initial conditions of temperature, three-dimensional velocity fields, salinity and non-tidal sea level (Newinger 2015). GLORYS2V3 reanalysis uses the NEMO v3.1 ocean model. The reanalysis is data assimilated with SST and sea level from satellites; temperature and salinity from floats and drifters; and surface velocities from the Global drifter program (Newinger 2015). There is also a tidal forcing from the 0.083-degree resolution Atlantic Ocean Atlas Solution (Egbert et al. 1994). However, the ROMS model runs were performed without tides.

The atmospheric boundary conditions are provided by the National Center for Environment Prediction (NCEP), Climate Forecast System Reanalysis (CFSR) (Saha et al. 2010) from 2001 to 2010. The variables used for the atmospheric boundary conditions are the 10 m winds, air temperature, specific humidity at 2 m, precipitation rate, downward longwave and shortwave radiation at the ground, and total cloud coverage (Newinger 2015). The variables are 6-hourly and available at a horizontal 0.312-degree resolution. The atmospheric pressure at mean sea level has a 0.5-degree horizontal resolution.

In ROMS, the river input occurs as freshwater sources at the river mouths (Newinger 2015). The discharge of the Amazon and Orinoco Rivers are obtained from the monthly discharge measured by the Runoff Data Center (Global Runoff Data Center 2014)

2.2 The Framework

The framework presented here aims to produce fast DA with sufficiently accurate results. The handling of the model and DA occurs in the reduced-space making it faster and less computationally expensive. This substitute is a projection of the model solution from full-space into a reduced-space. The latter is used to train a NN. To reduce the model's dimensionality, we use an EOF analysis. The aim of the EOF analysis is to reduce the number of variables of the high-dimensional system while capturing most of the variance (Hannachi 2004). The NN attempts to replicate the dynamics of the full-space model solution, and it is trained using state-of-the-art software packages. NNs, used in time series, are skillful tools in extracting patterns from past events and extrapolating these into future time-steps (Knutti et al. 2003). They approximate any given continuous function with required accuracy and are data-driven. These features are particularly helpful when substantial quantities of data are available (Hippert et al. 2001).

The main framework is represented in Fig. 2.1 and can be outlined in three steps: (a) dimension reduction and NN training, (b) DA, and (c) reconstruction to the full-space.

2.2.1 Dimension reduction algorithm

The surface layer of daily hindcast ROMS simulation of the North Brazil region is utilised as the original full-space state (FS). The length of the training and test datasets from ROMS varies throughout Chapters 3, 4 and 5. The variables extracted from the ROMS model solution were Sea Surface Height (SSH), Sea Surface Temperature (SST), Eastward and Northward horizontal velocities (U and V, respectively) of the surface currents. The ROMS model solution can provide additional variables, but this study is limited and targeted to the previously mentioned variables.

We used the `eofs` Python package to perform the EOF analysis which is computationally

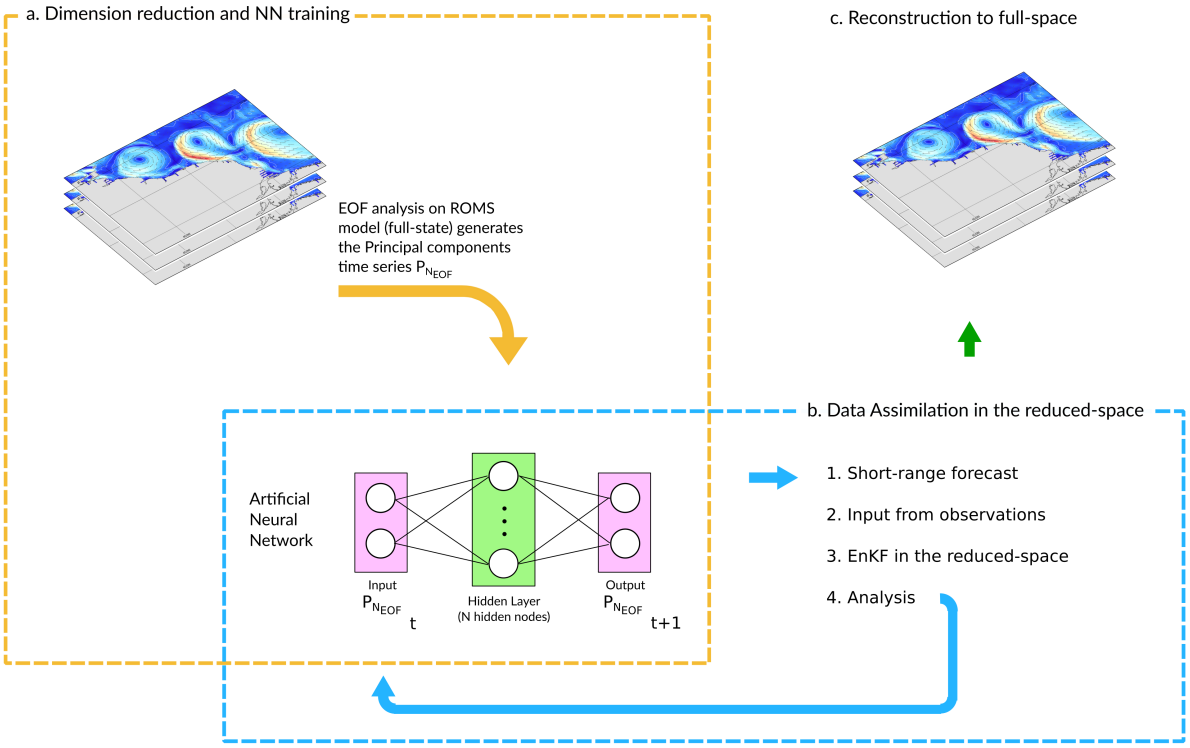

Figure 2.1: Framework diagram to perform DA in a reduced-space. **a**. Dimension reduction and NN training, **b**. DA, and **c**. Analysis reconstruction to full-space. The first step includes the EOF analysis to reduce the dimensionality of the surface layer of a ROMS model solution and obtain the reconstructed full-space state (RFS) at every time step. Via an EOF analysis we obtain the truncated time series of the Principal Components ($\mathbf{P_{N_{eof}}}$) which are used to train a NN with a hidden layer and N hidden nodes. The NN is configured so the $\mathbf{t+1}$ timestep can be predicted from the $\mathbf{t}$ timestep of $\mathbf{P_{N_{eof}}}$. The second step is a sequential EnKF, and includes a short-range forecast from the NN followed by the input of the observational products into the system. This step creates a new updated state (analysis) which is used as a new input into the NN from which we obtain a new short-range forecast. Lastly, after the DA, the updated reduced-space analysis is reconstructed into the full-space.

efficient for large output datasets. The `eof` package uses a technique based on Singular Value Decomposition (SVD). A further description of this package and the solving methods can be found in Dawson (2016). The package is available at http://ajdawson.github.io/eofs/.

Before performing the EOF analysis on FS it is necessary to prepare the surface ocean

model state. Each variable is reshaped into a two-dimensional matrix $\mathbf{A}_i$:

$$\mathbf{A}_i = \begin{bmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,m} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n,1} & a_{n,2} & \cdots & a_{n,m} \end{bmatrix}, \; i = 1, \cdots, q \tag{2.1}$$

where a are values of the variable i at the ocean model state grid points, m is the number of data points of the i^{th} variable and n is the number of time steps. The landmask gridpoints were ignored and will be placed back in the outputs datasets of the EOF analysis. The horizontal concatenations of these matrices result in $\mathbf{D^{fs}}$, as follows:

$$\mathbf{D^{fs}} = [\mathbf{A}_1, \cdots, \mathbf{A}_i, \cdots, \mathbf{A}_q], \; M = m \times q \tag{2.2}$$

where M is the length of the state vector of $\mathbf{D^{fs}} \in \mathbb{R}^{n \times M}$. $\mathbf{D^{fs}}$ is the original full-space matrix. Because the sub-matrices do not have the same units, it is necessary to standardise them. The dimensionless anomaly matrix $\bar{\mathbf{D}}^{\mathbf{fs}}$ is then defined by:

$$\bar{\mathbf{D}}^{\mathbf{fs}} = [\bar{\mathbf{A}}_1, \cdots, \bar{\mathbf{A}}_i, \cdots, \bar{\mathbf{A}}_q] \tag{2.3}$$

where $\bar{\mathbf{A}}_i$ is the standardised data matrix of the i^{th} variable. The standardisation is done by removing the column-wise mean and dividing by the i^{th} variable standard deviation $\sigma_{\mathbf{A}_i}$. Other possible normalisations include log-transform and scaling from 0 to 1. The log-transform normalisation is usually used when variables are always positive or when their values are in different orders of magnitude. Scaling to values from 0 to 1 could skew the data distribution towards outliers. The standardisation used in this study is therefore more appropriate since our variables have different units, but they are not represented in drastically different orders of magnitude (Shlens 2014). A disadvantage would be that if we were analysing variables that have the same units and values within the same order of magnitude, then it might not be necessary to do any standardisation.

After the EOF analysis is performed, $\bar{\mathbf{D}}^{\mathbf{fs}}$ can be expressed as:

$$\bar{\mathbf{D}}^{\mathbf{fs}} = \mathbf{P}\mathbf{\Pi} \tag{2.4}$$

where $\mathbf{P} \in \mathbb{R}^{n \times n}$ contains the time series of the principal components (PCs) and $\mathbf{\Pi} \in \mathbb{R}^{n \times M}$ are the EOFs of $\bar{\mathbf{D}}^{\mathbf{fs}}$. The PCs nor the EOFs are scaled by the eigenvalues as it is the default output used in `eofs` by Dawson (2016). The total number of PCs is $n - 1$ because the dimensionless anomaly matrix has rank smaller or equal to $n-1$ (i.e, the sum of all rows is equal to zero vector, and consequently the rows are linearly dependent).

The dimension reduction of the system is then obtained by truncating $\mathbf{P}$ and $\mathbf{\Pi}$. Let N_{eof} be the number of selected EOFs, then the reconstructed full-space anomaly matrix $\bar{\mathbf{D}}^{\mathbf{rs}}$ is defined by:

$$\begin{aligned} \bar{\mathbf{D}}^{\mathbf{rs}} &= \mathbf{P}_{\mathbf{N_{eof}}}\mathbf{\Pi}_{\mathbf{N_{eof}}} \\ &= [\bar{\mathbf{A}}^{\mathbf{rs}}_1, \cdots, \bar{\mathbf{A}}^{\mathbf{rs}}_i, \cdots, \bar{\mathbf{A}}^{\mathbf{rs}}_q] \end{aligned} \tag{2.5}$$

with $\mathbf{P}_{\mathbf{N_{eof}}} \in \mathbb{R}^{n \times N_{eof}}$ and $\mathbf{\Pi}_{\mathbf{N_{eof}}} \in \mathbb{R}^{N_{eof} \times M}$.

To reconstruct $\bar{\mathbf{D}}^{\mathbf{rs}}$ back to ocean model state it is necessary to unstandardise the matrix by multiplying each submatrix $\mathbf{A}^{\mathbf{rs}}_i$ by $\sigma_{\mathbf{A}_i}$ and adding the previously calculated column-wise means. Finally, the non-standardised reduced-dimension matrix $\mathbf{D}^{\mathbf{rs}}$ is reshaped and landmask nodes are placed back at every time step. The reconstruction will be referred to as the reconstructed full-space state (RFS).

2.2.2 Neural Network for short-range forecast within the EnKF

At this stage, it is necessary to replicate or emulate the temporal evolution of the full-space ROMS solution. Our approach consists of training a NN that will achieve the replication with an acceptable degree of accuracy. The chosen architecture is an autoregressive feed-forward neural network based on the MLP, which consists of three parts: an input layer, a set of hidden layers, and an output layer (Fig. 2.1b). Here, we trained the NN to

receive $\mathbf{P}_{N_{eof}}$ at time-step t in the input layer. The target of the output layer is to predict $\mathbf{P}_{N_{eof}}$ at time-step $t+1$ given the information entered in the input layer. The hidden layers act as the connection between input and output layers through a series of weights and hidden nodes. A single hidden layer can also be used.

The NN needs to be properly trained before being used in the DA step. Training the NNs finds the optimal hyperparameters which minimise the difference between the target and predicted value in the output layer. These hyperparameters include, but are not limited to, activation functions (linked to the hidden nodes), learning rates, training methods, the number of hidden layers and the number of hidden nodes. Different combinations of these hyperparameters determine the prediction ability of the NN. The dataset of $\mathbf{P}_{\mathrm{N_{eof}}}$ is divided into subsets of training (70 %), validation (15 %) and testing (15 %). The NN will find the optimal weights that minimise the cost function (eq. (1.11)) within the training subset and validate them in the validation subset. Our cost function is the Mean Squared Error (MSE) between the target and the predicted value. If the cost function can be minimised further, the weights are modified during the validation stage. Since the validation of a NN evaluates whether the chosen weights for training were optimal, the validation can be considered part of the training method. Therefore, 7.65 years were considered for training/validation with a testing period of 1.35 years. During the training and validation stages the target output is known. During the testing stage we obtain predicted values at the output layer with the weights fixed, and the predicted values are not adjusted to the targets accordingly. If the MSE between the predicted values and the target values of the test dataset differ greatly, it means the trained NN is overfitted to the training values and it is not capable to generalise. Thus, it is necessary to modify the hyperparameters and train the NN again.

Different combinations of the hyperparameters can be tested using the function `GridSearch` in Sci-kit learn. The hyperparameters tested include:

- Learning rate: 10 combinations in the range of 10^{-7} to 10^{2}.

- Activation functions: sigmoid, logarithmic sigmoid, linear, ReLU (Nair & Hinton 2010).
- Training method: ADAM (Kingma & Ba 2014), scalar conjugate gradient (SCG) (Møller 1993), gradient descent (Bengio et al. 1994).
- Number of Hidden layers: 1 to 5.
- Number of Hidden Nodes: 70 combinations in the range of 1 to 1000.

We limited the epochs to 2000 per combination. Once the search for the hyperparameters is finished, these are used in MATLAB. Along with the MSE cost functions during the training and testing stages, we aimed for a correlation coefficient $r \geq 0.9$ between the target and predicted values.

The first step including the dimension reduction and the optimal NN training needs to be performed once and then is ready and available to be used in the reduced-space DA.

2.2.3 Data assimilation using the EnKF

The second step consists of a sequential DA on the reduced-space. Any appropriate method of DA can be applied at this stage. We chose the EnKF (Evensen 1994, 2003). The EnKF sequentially assimilate time-distributed observations. The EnKF involves a forecast step, where a previous state estimate evolves forward to the time of the observations, and an update step in which the evolved state estimate is updated with information from the observations. First, we create a 50-member ensemble ($\mathbf{X}(t)$). As an example, if we start the EnKF DA at January, 1st 2010, we compute the mean of the time steps of $\mathbf{P}_{N_{eof}}$ related to this day (i.e., January 1st, 2001, ..., January 1st, 2009). This mean is perturbed by adding Gaussian noise with zero mean and creates a perturbed ensemble. Each member of the perturbed ensemble is then forecasted one time-step forward using

the trained NN creating $\mathbf{X_{forecast}}$. The update of $\mathbf{X}$ is given by:

$$\mathbf{X_{analysis}} = \mathbf{X_{forecast}} + \mathbf{K}(\mathbf{y} - \tilde{\mathbf{H}}\mathbf{X_{forecast}}) \tag{2.6}$$

where $\mathbf{X_{analysis}}$ is the updated ensemble state ($\mathbf{X}(t+1)$) of the ensemble $\mathbf{X_{forecast}}$. $\mathbf{K}$ is the Kalman gain, $\mathbf{y}$ is the ensemble of standardised observations and $\tilde{\mathbf{H}}$ is the reduced-space observation matrix. The analysis is the mean of this updated ensemble.

The ensemble $\mathbf{y}$ is obtained by adding Gaussian perturbations with zero mean and data error covariance $\mathbf{R}$. The reduced-space observation matrix $\tilde{\mathbf{H}}$ is defined as:

$$\tilde{\mathbf{H}} = \mathbf{H}\mathbf{\Pi}^{\mathbf{T}}, \tilde{\mathbf{H}} \in \mathbb{R}^{O \times N_{eof}} \tag{2.7}$$

where $\mathbf{H}$ is the full-space observation matrix and O is the number of observations.

The Kalman gain, $\mathbf{K}$, is defined by:

$$\begin{aligned} \mathbf{K} &= \tilde{\mathbf{B}}\tilde{\mathbf{H}}\underbrace{(\tilde{\mathbf{H}}\tilde{\mathbf{B}}\tilde{\mathbf{H}}^{\mathbf{T}} + \mathbf{R})^{-1}}_{\mathbf{W}^{-1}} \\ \tilde{\mathbf{B}} &= (1+\alpha)\mathbf{B} \end{aligned} \tag{2.8}$$

where the background error covariance matrix $\mathbf{B}$ is calculated directly as the covariance of $\mathbf{X_{forecast}}$, and $\tilde{\mathbf{B}}$ is the inflated $\mathbf{B}$ using a multiplicative inflation factor of $(1+\alpha)$. The inflation is done to increase the spread of the background error. The gridded data is highly available and the same error variance to observations in coastal areas and around the mouth of the Amazon River is prescribed. Thus, the EnKF does not know that these are troublesome regions and the EOFs might not reproduce data properly in these regions due to smaller scales of variability. Therefore, localisation would not help in troublesome areas.

The inversion of $\mathbf{W}$ becomes very expensive for high-dimensional systems or for a large dataset of observations, like gridded data. Since the computational cost contradicts

the goal of performing fast DA, it is necessary to optimise the inversion. Considering uncorrelated data errors, $\mathbf{R}$ can be written as a diagonal matrix. The Sherman-Morrison-Woodbury formula is used to decompose $\mathbf{W^{-1}}$ as:

$$\begin{aligned} \mathbf{W^{-1}} &= (\mathbf{R} + \beta\gamma^{\mathbf{T}})^{-\mathbf{1}} \\ &= \mathbf{R^{-1}} - \mathbf{R^{-1}}\beta(\mathbf{I} + \gamma^{\mathbf{T}}\mathbf{R^{-1}}\beta)^{-\mathbf{1}}\gamma^{\mathbf{T}}\mathbf{R^{-1}} \\ \beta &= \tilde{\mathbf{H}}\mathbf{B} \\ \gamma &= \tilde{\mathbf{H}}^{\mathbf{T}} \end{aligned} \tag{2.9}$$

This way, the computation of $\mathbf{W^{-1}}$ is expected to be inexpensive as it depends on the inversion of the diagonal matrix $\mathbf{R}$ and of a square matrix with dimension N_{eof}. In order to build $\mathbf{R}$, a gridded map of data errors was created using the same landmask; and surface ocean grid points for every variable were replaced by the measurement error associated to their respective observational products. The data errors were standardised as the gridded values of the FS. The standardised error across the grid helps to create $\mathbf{R}$ as follows:

$$\mathbf{R} = \begin{bmatrix} diag(\frac{\omega}{\sigma})_1^2 & & & & \\ & \ddots & & & \\ & & diag(\frac{\omega}{\sigma})_i^2 & & \\ & & & \ddots & \\ & & & & diag(\frac{\omega}{\sigma})_q^2 \end{bmatrix} \tag{2.10}$$

where ω is the measurement data error of the observation dataset linked to the variable i and σ is the aforementioned standard deviation of every i^{th} variable. The term "*diag*" indicates diagonal matrices containing the standardised measurement errors for each variable, where the diagonal has length O. R is a dimensionless non-singular square diagonal matrix that can be used with eq. (2.9).

2.2.4 Reconstruction of the reduced-space analysis

The last step involves the reconstruction of the sequential reduced-space analyses back into the full-space. The outputs of the EnKF can only be analysed after the reconstruction. It is important to note that the reconstruction can be performed at any time (e.g., to assess the dimension reduction or the performance of the NNs) and not strictly after the DA is implemented. Thus:

$$\mathbf{X^{RA}}(t) = \mathbf{X}\bar{(t)}_{\mathbf{analysis}}\mathbf{\Pi}_{N_{eof}} \tag{2.11}$$

where, $\mathbf{X^{RA}}(t)$ is the reconstructed reduced-space analysis at time-step t. $\mathbf{X^{RA}}(t)$ is subsequently unstandardised and the landmask nodes of the ROMS model solution are put back into place.

2.3 Assimilated data and observational datasets

Throughout this study Sea Surface Temperature (SST), Sea Surface Height (SSH), Eastward and Northward component of surface velocities (U and V) are assimilated and validated using the following datasets:

Gridded SST data from GHRSST (Group for High-Resolution Sea Surface Temperature) (GHRSST 2010) by the National Oceanic and Atmospheric Administration (NOAA). GHRSST-L4 is a gap-free gridded satellite product, with complimentary in-situ observations, with a 10 km resolution. The GHRSST dataset uses optimal interpolation from the Advanced Very High Resolution Radiometer (AVHRR) Pathfinder version 5 data (Newinger 2015, Reynolds et al. 2007). The dataset is bias adjusted by the in-situ data, but Reynolds et al. (2002, 2007) suggest that small biases are low in the Amazon River plume region.

The mean sea level anomaly (MSLA) in delayed time was obtained from the AVISO Ssalto/Duacs database. The latter is a 0.25-degree resolution gridded satellite repository. The Ssalto/Duacs altimeter products were produced and distributed by the Copernicus Marine and Environment Monitoring Service (CMEMS). The AVISO dataset provides a temporal coverage since January 1993. The SSHs and derived variables are computed with respect to a 20-year mean. The delayed-time SLA is delivered within 2 months depending on the mission. To obtain SSH in this study we added the mean dynamical topography (MDT) from the ROMS climatology to the MSLA.

The Globcurrent (v.1, v.2, and v.3) is a 3-hourly 0.25-degree resolution dataset that includes the Ekman and geostrophic currents. While Globcurrent is available at two different depths (surface and 15 m depth), we used the total combined currents ($u_{combined}$) at the surface (Rio et al. 2014*b*):

$$u_{combined}(z = hs) = u_{geost} + u_{ek}(z = hs) \tag{2.12}$$

where u_{geost} and u_{ek} are the geostrophic and Ekman currents, respectively. The calculation of the Ekman and geostrophic currents are described in Rio et al. (2014*a*). The v.1 product are global maps of the tidal currents, Stokes drift, the geostrophic currents, and the Ekman currents and are available from 2010 to the end of 2012. The v.2 does not include tidal currents nor Stokes drift currents and is available for the 2012-2014 period. The v.3 is available from 1993 to 2017 for global currents at 0.25-degree and 0.50-degree resolution, and for the Mediterranean Sea at 0.12-degree resolution.

The data that will be assimilated is interpolated to the ROMS grid resolution before the DA, but preserves the original landmasks of the observational data. The measurement data errors considered in this study are 0.04 m for SSH, 0.5°C for SST, and 0.1 m s^{-1} for surface velocities U and V.

Chapter 3

Fast Data assimilation of ocean surface variables

3.1 Introduction

Data assimilation is a method to conciliate the large variety and volume of data with numerical models. It can also be used to find optimal initial conditions (Leeds et al. 2014) for better forecasts by numerical models which represent the physics. The high-dimensionality of these numerical models contributes to their high computational-cost. The application of DA schemes in such models increases this cost, because they typically require multiple runs of the numerical model making it hard to satisfy the requirements of speed and accuracy at the same time (Margvelashvili et al. 2016).

An approach that has gained some attention, is to use fast and cheap model substitutes, surrogates or emulators of more complex full-space models to perform DA (Frolov et al. 2009, Mattern et al. 2012, Margvelashvili & Campbell 2012). The idea is to reduce the high-dimensionality and/or complexity of the original full-space model and replace it with the model substitute which can run often orders of magnitude quicker and comes with an

acceptable loss in accuracy. A caveat is that a emulator created for a specific model region or application may not perform well in other regions or applications. Several studies have used emulators to replicate or approximate a full-space model. These applications range from idealised models of ocean eddies (Kondrashov & Berloff 2015), to real ocean applications using polynomial chaos expansion (Mattern et al. 2012), rank-reduced statistical emulation (Leeds et al. 2014, Hooten et al. 2011), and sequential emulator-based assimilation on coastal sediment (Margvelashvili & Campbell 2012, Margvelashvili et al. 2013)

Other approaches include the replacement of a high-dimensional physical model for one with reduced dimensionality by using EOF analysis and further training of the reduced-space models using neural networks (van der Merwe et al. 2007). For the latter, it is essential to have a long run of the hindcast model to provide a sufficiently large dataset, in this case the physics and dynamics, of the original high-dimensional model. Frolov et al. (2009) successfully expanded this idea into a DA framework using data from observations and combining them with an estuary model of the Columbia River. We apply a framework based on Frolov et al. (2009), but for the first time in a regional model of the open ocean.

3.2 Methods and observational data

In this study, we included SSH, SST, U and V in the full-space state (FS). Although, we only assimilate SSH and SST data using the AVISO and GHRSST datasets, respectively, with further validation using data of SSH, SST, U and V. Synthetic velocities from Globcurrent v.1 during 2010 are used to validate U and V.

The hindcast ROMS model solution is run for 10 years. Here, we used the first 9-years (2001-2009) to train and test the NN. We then used 2010 to validate our results. The sequential DA starts on January 1st, 2010 and it is performed until December 31st, 2010. The reconstructed full-space state (RFS) (as in section 2.2.1) corresponds to the

unstandardised reconstruction back to ocean model state, after applying the EOF analysis to the standardised FS. The RFS uses 100 EOFs to make the reconstruction.

3.2.1 Validation experiments

In order to address the role of the DA and of the one-day prediction given by the NN within the DA, we perform two additional experiments: 1) a projection of the gridded SSH and SST data onto the set of EOFs using a Least Squares (LS) solution, and 2) a persistence experiment.

The LS experiment involves estimating the coefficients of the time series of the PCs based on restricted vectors of the EOF matrix. These restricted column vectors correspond to vectors composed of EOF elements where data is available. These coefficients can be found by solving an LS problem on the reduced-dimension space and minimising at each time step i:

$$J = \|d_i - c_i e_i\| \tag{3.1}$$

where d are the data vectors in 2010, e are the restricted vectors to available data and c are the coefficients that minimise J. The LS solution should be an overfit to the observations and pose less physical constraint than the DA solution using the EnKF.

The persistence experiment, which entails no dynamics, consists of replacing the forward stepping from the NN with the ensemble persistence to subsequently assimilate the data. At the following time-step, new information is obtained from the observations and the process is repeated. The persistence solution should present a smaller spread of the ensemble during the forecast step than the NN forecast and therefore decouple from the truth state underperforming the NN.

3.3 Results

3.3.1 EOF analysis

The concatenated matrices of SSH, SST, U, and V resulted in a state vector of $\sim 10^5$ dimensions at a single time-step. All the landmask nodes were ignored in the process, reducing the state vector to almost half of its original dimensionality ($\sim 6\times10^4$). The chosen number of EOFs determines the dimensionality of the new system and the explained variance of RFS. Several test cases with different EOFs were performed.

The first 3 modes of SSH and SST are shown in Fig. 3.1. These account for 50.97% of the variability. The first mode is dominated by seasonal variability. The second and third modes are dominated by SSH in combination, and permit the movement of mesoscale waves across the domain.

Figure 3.2 shows how the explained variance changes with respect to the number of EOFs and the different combinations of SSH, SST, U, and V that can be included in the model. Clearly, SSH and SST are the more tractable variables, where the percentage of explained variance does not change greatly if we were to use 50, 100 or 200 EOFs. On the other hand, the surface currents U and V, are more difficult to model. A large range of explained variance exists if 50 or 200 EOFs are kept. The explained variance of the combination of all of the variables is expected to be in between the variance of SSH and SST, and the variance of the currents. In this study, the first 100 EOFs (N_{eof} in section 2.2) were retained explaining $\sim 93\,\%$ of the variance, and all of the presented results are based on this number of modes.

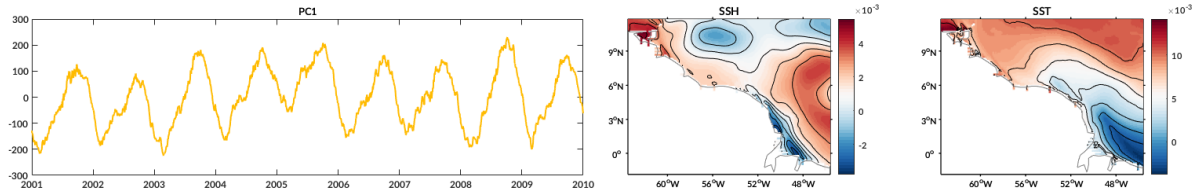

(a) PC1: 29.77%

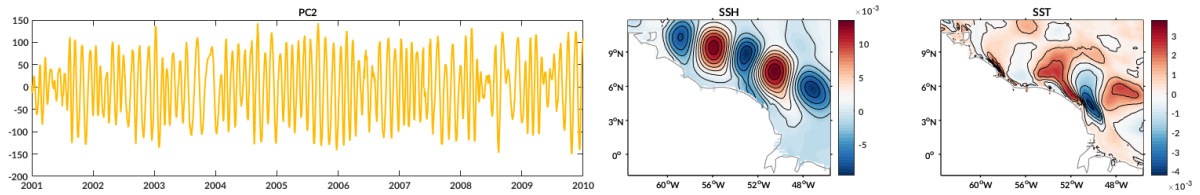

(b) PC2: 10.78%

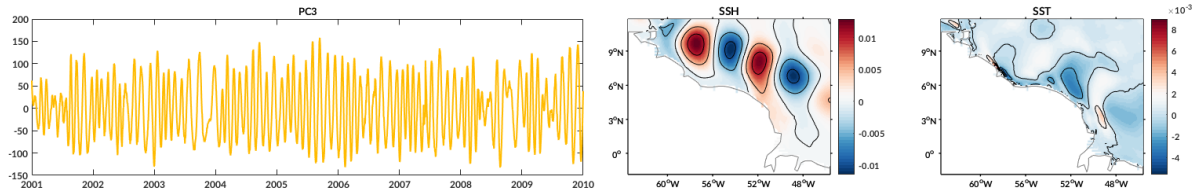

(c) PC3: 10.42%

Figure 3.1: Time series of the Principal Components (PC) and their respective EOFs of SSH and SST of the ROMS model solution from 2001 to 2009. From top to bottom: first, second and third principal components and modes, including the explained variance.

3.3.2 Seasonal correlations

Between RFS and FS, the seasonal correlations of SSH, SST, U and V are high with $r > 0.99$ and with negligible mean RMSE. Boreal Winter, Spring, Summer and Autumn are represented by DJF, MAM, JJA and SON, respectively.

For SSH the lowest correlations are observed during JJA in the West of the Caribbean islands. The r values along the coast, the mouth of the Amazon River and boundaries have the lowest values compared to the rest of the grid (not shown).

Seasonal correlation coefficients of the horizontal velocities, U and V, show similar behaviour between them. Similar to the SSH, both U and V have their lowest correlations in the NW part of the grid and along the coast. The seasonal values of r are larger in the

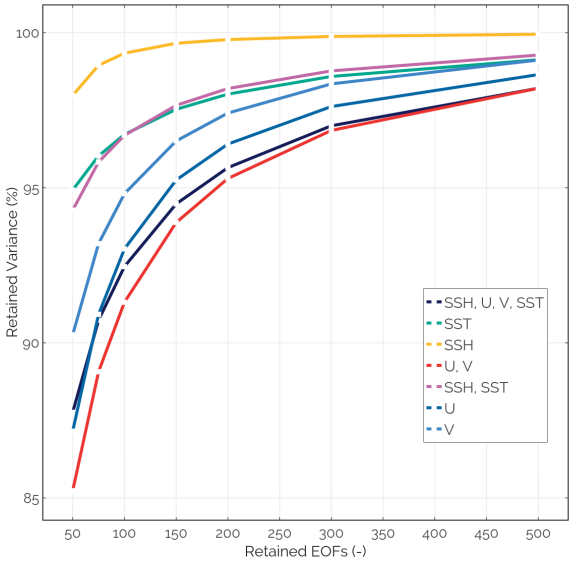

Figure 3.2: Explained variance versus number of EOFs retained. Different dashed lines represent the explained variability with respect to 50, 75, 100, 150, 200, 300 and 500 EOFs for different combinations of variables used in the model: SSH, SST, U, and V.

open ocean than in the rest of the domain (not shown). The discharge of the Amazon River and the NW movement of the NBC show an area of high correlation between the RFS and the FS, but lower than in the open ocean. This correlation is larger during DJF and MAM.

The low correlations along the coast and the NW areas appear in SST as well. The input of the Amazon River in JJA and SON induces large variability over the SST. The area with higher correlations is the open ocean and the North boundary of the domain (not shown). JJA exhibits the lowest seasonal correlations for SSH, SST, U and V.

3.3.3 Dimension reduction

Figure 3.3 illustrates a comparison between the RFS and the FS as a snapshot of December 1st, 2009 which clearly shows the retroflection of the NBC and formation of NBC rings.

Sea surface height (Fig. 3.3a and 3.3b) displays a clear representation of the NBC rings

aforementioned. Between 48°W - 52°W the FS is not distinguishable from the RFS. The NBC ring positioned around 56°W shows some difference between the RFS and the FS figures. In the FS and RFS it can be seen how the NBC ring is about to move in the NW direction. The main features of the NBC rings can be observed in both model solutions, thus the mesoscale features are well represented in both. The RFS and the FS exhibit the same SSH anomalies around the mouth of the Amazon River.

The FS snapshots of SST (Fig. 3.3c) show distinctive swirls of temperature anomalies around the NBC rings, while in the RFS (Fig. 3.3d) these features are hardly noticeable. Nevertheless, both the FS and the RFS show lower temperatures near the retroflection of the NBC, along the coast of Suriname and the west of Trinidad.

The currents (Fig. 3.3e and 3.3f) show a NBC ring moving SE-NW along the coastline. The formation of the NBC ring, between 48°W - 52°W, can be seen in both model solutions. Whereas the RFS illustrates these rings clearly, the FS presents a higher resolution of the mesoscale features. For example, in the retroflection of the NBC, near 48°W, the NBC ring advecting NW is already detached from the NBC in the RFS, while it is still attached in the FS. The FS generates the high azimuthal velocity of the full eddy near 56°W, while the RFS shows this feature only over the southern and western region of the ring. Furthermore, the RFS presents lower azimuthal velocity in the eddies.

3.3.4 Neural network training

In our NN training, 30000 combinations of hyperparameters were tested, which took approximately 50 hours using 12 cores of Intel Xeon (2.5 GHz) processors. MATLAB and Sci-kit learn, in Python, provided the necessary machine learning tools for this task. The targeted correlation coefficient between the targets and output layer predictions was $r \geq 0.9$. The best result was found for the NN with a single hidden layer, 151 hidden nodes with linear activation functions, a learning rate of 0.001, and using the SCG as a

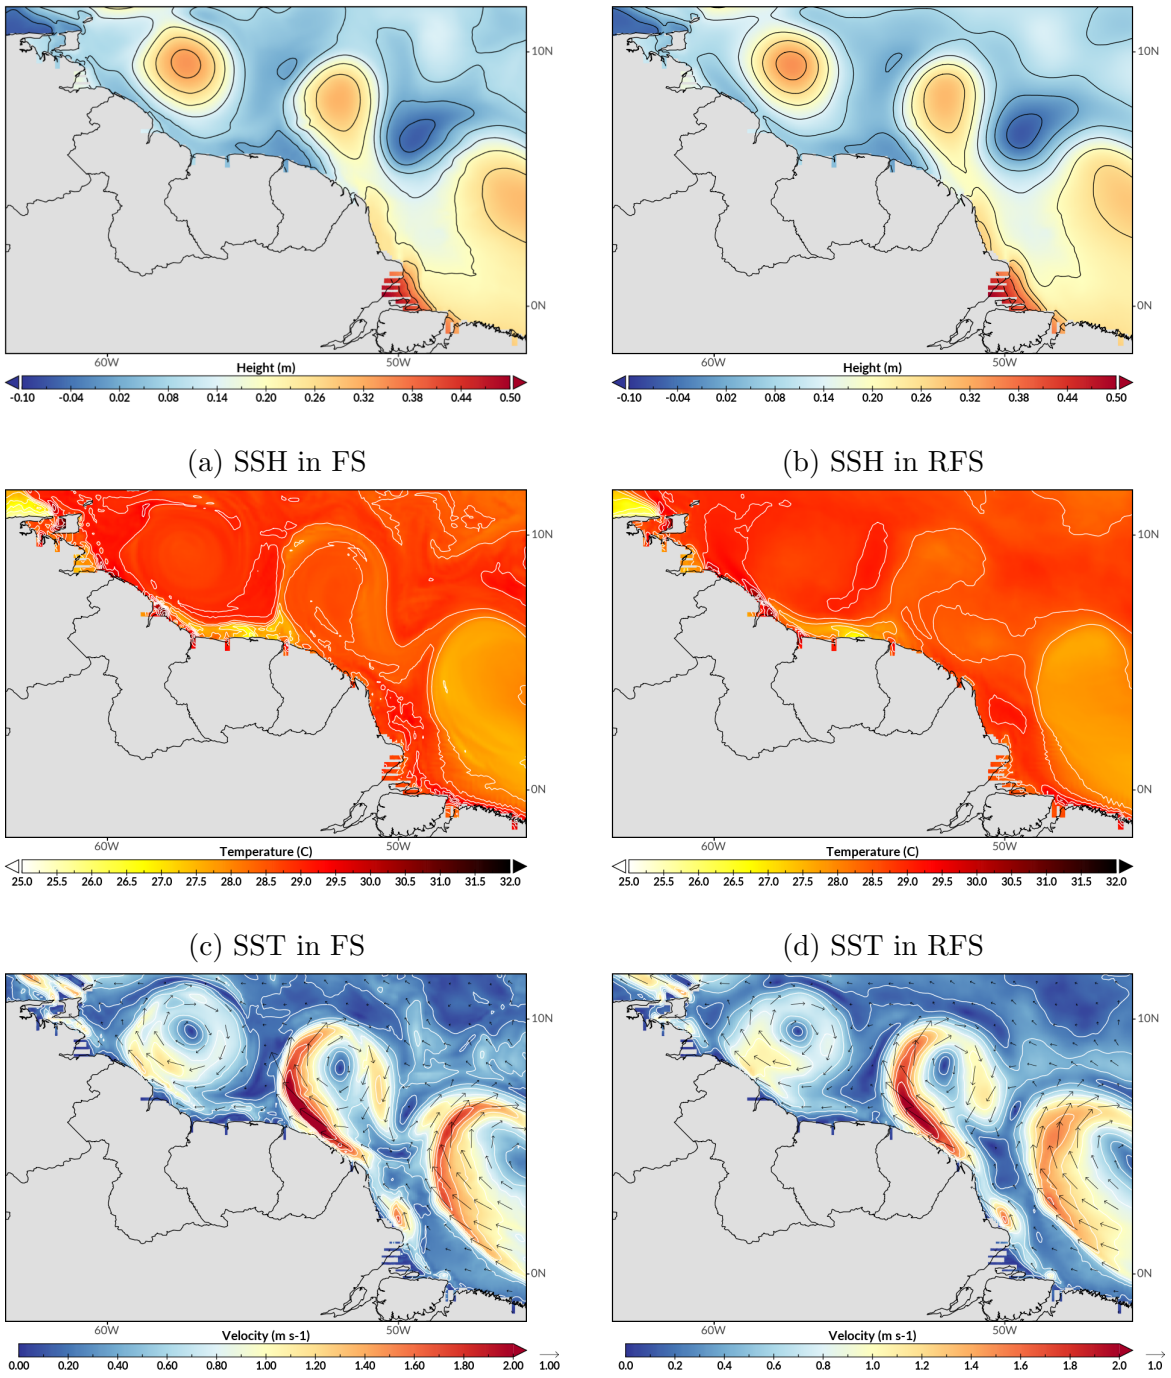

(a) SSH in FS

(b) SSH in RFS

(c) SST in FS

(d) SST in RFS

(e) Surface currents in FS

(f) Surface currents in RFS

Figure 3.3: Full-space state (FS) and reconstructed full-space state (RFS). Snapshot of December 1st, 2009 showing the genesis and advection of the North Brazil Current rings travelling northwestward and the retroflection of the NBC to supply the North Equatorial Counter Current

training method. (Fig. 2.1).

Figure 3.4: RMSE time series (from January 1st, 2010) of the validation with respect to GHRSST, AVISO altimetry, and Globcurrent datasets to validate (from top to bottom): SST, SSH, U, and V. Blue corresponds to the reconstructed full-space state (RFS), green is the Least Squares (LS) solution, red represents the ensemble-mean (analysis) after the EnKF was applied using the NN, and black is the analysis using the EnKF and persistence (P). Yellow and pink shades are centered on the snapshot days in Fig. 3.9 and 3.8.

3.3.5 DA analysis

Daily data for 2010 were assimilated. We used an inflation factor of $\alpha = 3$ to increase the spread of the ensemble of the background covariance error in both the NN and persistence

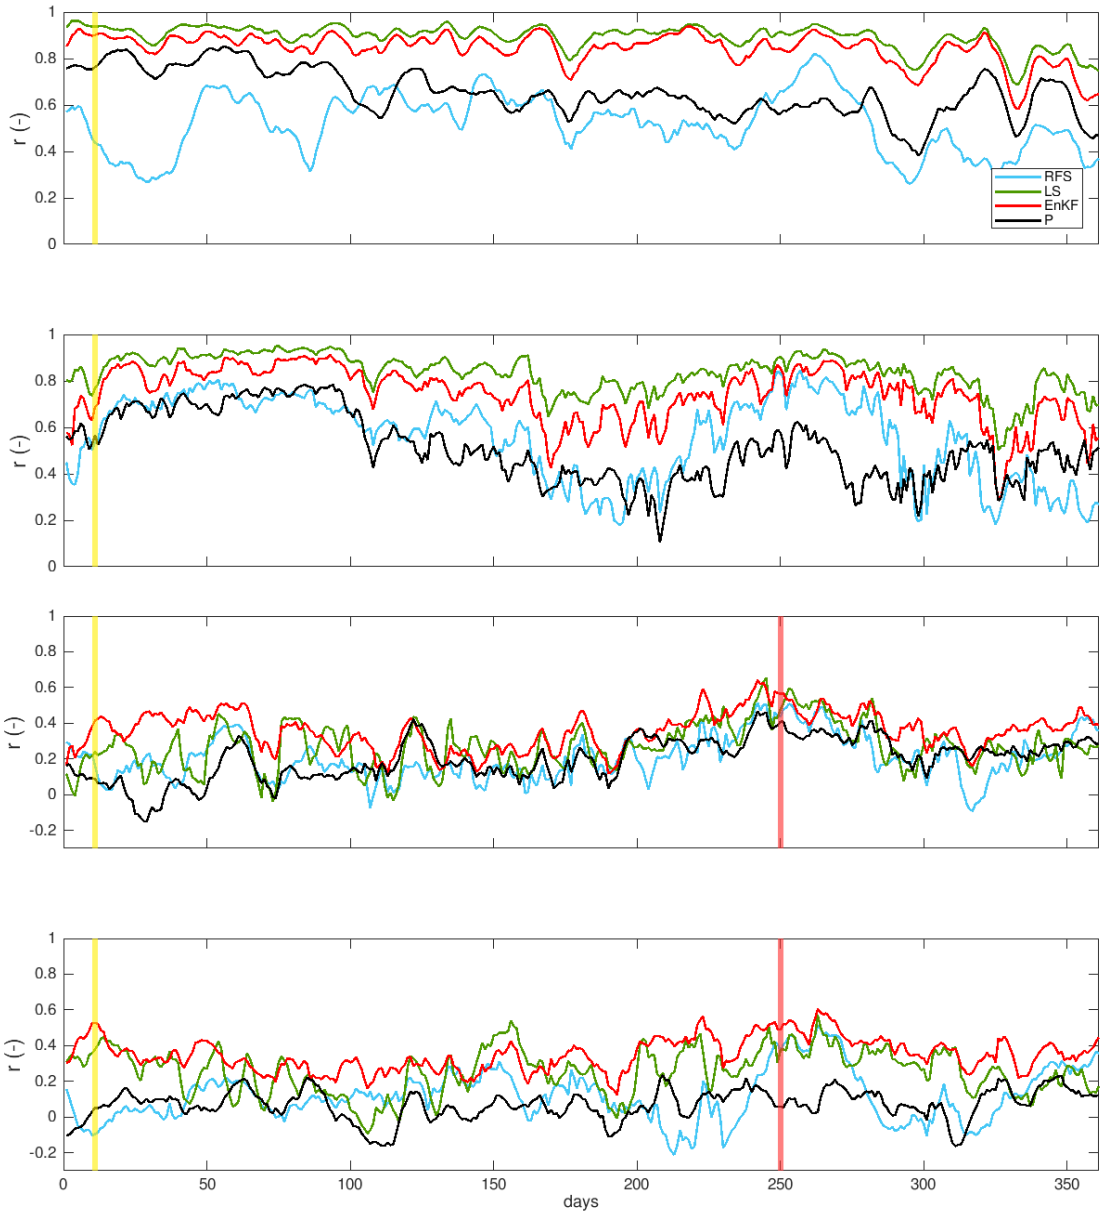

Figure 3.5: Correlation coefficient (r) time series (from January 1st, 2010) of the validation with respect to GHRSST, AVISO altimetry, and Globcurrent datasets to validate (from top to bottom): SST, SSH, U, and V. Blue corresponds to the reconstructed full-space state (RFS), green is the Least Squares (LS) solution, red represents the ensemble-mean (analysis) after the EnKF was applied using the NN, and black is the analysis using the EnKF and persistence (P). Yellow and pink shades are centered on the snapshot days in Fig. 3.9 and 3.8.

experiment. The mean of the diagonal of $\mathbf{H\tilde{B}H^T}$, using the NN, degrades from 0.49 and converges within 0.16 and 0.19. The mean of the diagonal of $\mathbf{R}$ is 0.24. Figures 3.4 and 3.5 show the results of the DA, where the coloured lines represent the RMSE and

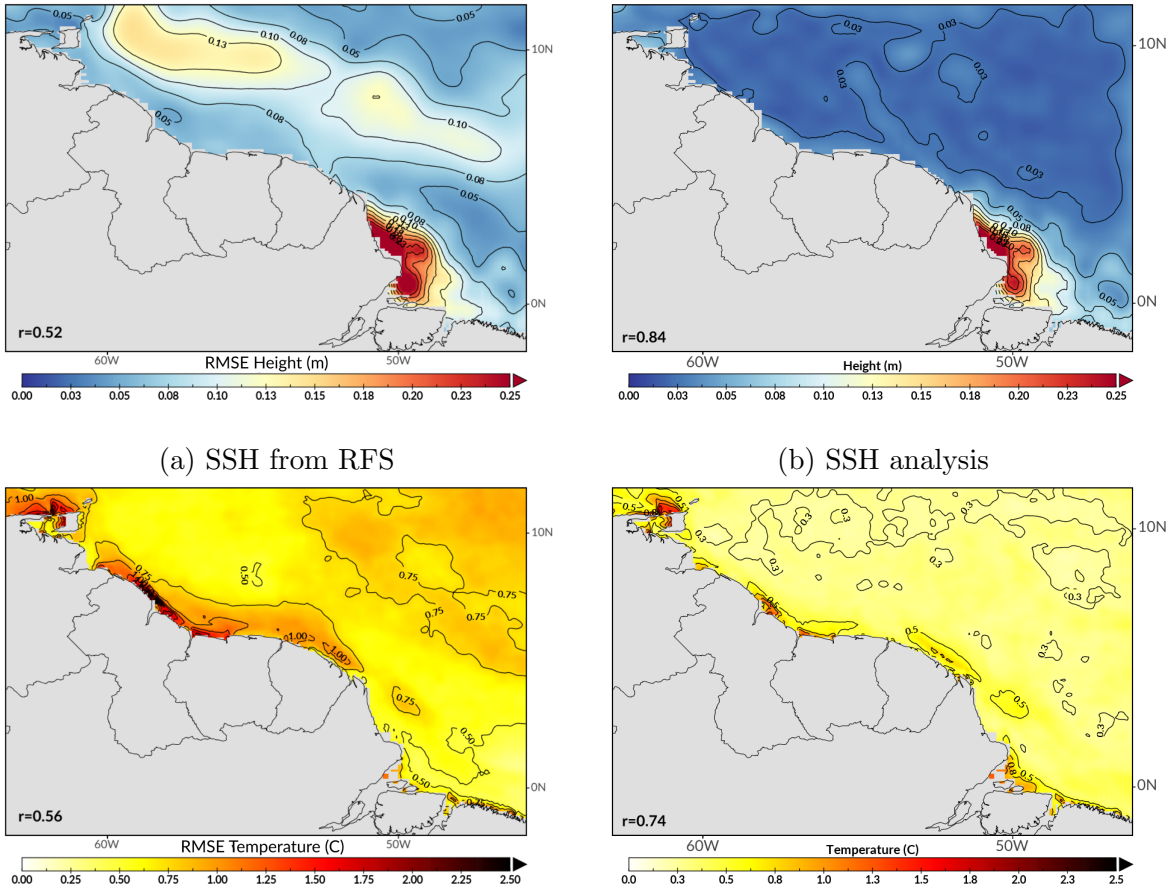

(a) SSH from RFS

(b) SSH analysis

(c) SST from RFS

(d) SST analysis

Figure 3.6: Annual average spatial RMSE for SSH and SST. Right figures represent the average spatial RMSE of the reconstructed full-space state (RFS) during 2010, while left figures represent the average spatial RMSE after the DA with an EnKF. The correlation coefficients (r) with respect to the validation datasets AVISO and GHRSST for SSH and SST, respectively, are depicted in the bottom left corner of each subfigure.

correlation coefficients of the RFS, LS and analysis with respect to the validation datasets of SSH from AVISO, SST from GHRSST, and U and V from Globcurrent. The RMSE and correlation time series of the FS (not shown) and the RFS are very similar. The annual mean RMSE of RFS and FS with the gridded datasets used for their validation are: 0.09 m for SSH (r = 0.52), 0.73 °C for SST (r = 0.56), 0.54 m s^{-1} for U (r = 0.21), and 0.62 m s^{-1} (r = 0.12) for V. It is clear that the EnKF DA using NN does decrease considerably the mean RMSE across the time series. The 365-day mean RMSEs after the EnKF using the NN are 0.05 m (r = 0.84), 0.35 °C (r = 0.74), 0.46 m s^{-1} (r = 0.36), and 0.53 m s^{-1} (r = 0.35) for SSH, SST, U, and V, respectively.

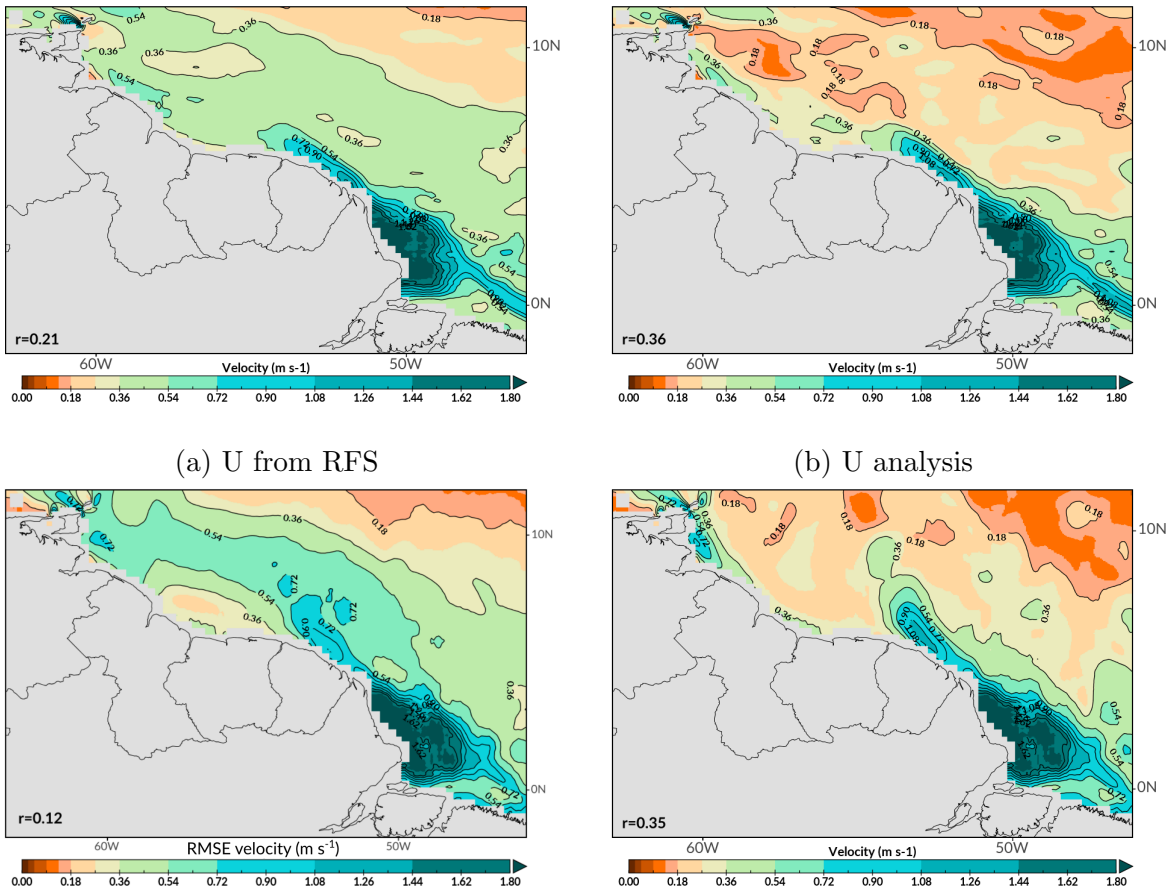

(a) U from RFS

(b) U analysis

(c) V from RFS

(d) V analysis

Figure 3.7: Annual average spatial RMSE for U and V. Right figures represent the average spatial RMSE of the reduced-space model solution (RFS) during 2010, while left figures represent the average spatial RMSE after the DA with an EnKF. The correlation coefficients (r) with respect to the validation dataset Globcurrent for U and V are depicted in the bottom left corner of each subfigure.

An additional metric to assess the improvement of the fit of SSH and SST to the observations is the unitless measure J_{fit} (da Rocha Fragoso et al. 2016):

$$J_{fit} = \frac{1}{m} \sum_{i=1}^{m} \left| \frac{x_i^{model} - y_i^{obs}}{e^{obs}} \right| \tag{3.2}$$

where m is the number of data points, x_i^{model} and y_i^{obs} are the model solution at the location of the i^{th} observation and the i^{th} observation, respectively, and e^{obs} is the observation error of the assimilated variable. J_{fit} determines if the fit of the observations to the model is within one standard deviation of the observation error. If J_{fit} is less than

1, then the analysis residual is within the observation error (Phillipson & Toumi 2017). The annual mean J_{fit} of SSH and SST in the RFS is 1.38 and 1.09, respectively, while is only 0.70 and 0.53 when SSH and SST data are assimilated, respectively, using the NN. These values indicate that the analysis error is of the same magnitude of their respective observation error.

The LS solution of the SSH, SST, U and V correspond to the reconstruction into the full-space of the EOF matrix using the PC coefficients (c) estimated by the Least Squares solution. The LS experiment yields an annual RMSE of 0.04 m (r = 0.89), 0.28 °C (r = 0.84), 0.64 m s^{-1} (r = 0.27), and 0.68 m s^{-1} (r = 0.26) for SSH, SST, U, and V, respectively. The persistence underperforms the NN with an annual RMSE of 0.08 m (r = 0.65), 0.71 °C (r = 0.51), 0.76 m s^{-1} (r = 0.20), and 0.76 m s^{-1} (r = 0.07) for SSH, SST, U and V, respectively. The statistics of both DA experiments, including the NN and persistence, are summarised in Table 3.1.

We introduce G_R as a unitless ratio, between the Root Mean Square (RMS) of the geostrophic and the total currents, that provides some measure of how much of the flow is geostrophic:

$$RMS = \sqrt{\frac{1}{m}\sum_{i=1}^{m}(x_U)_i^2 + \frac{1}{m}\sum_{i=1}^{m}(x_V)_i^2} \tag{3.3}$$

$$G_R = \frac{RMS_{geo}}{RMS_{total}} \tag{3.4}$$

where we obtain the joint RMS of the U and V components from the geostrophic (RMS_{geo}) and total currents (RMS_{total}), and we calculate the annual G_R for LS, DA using the NN, and the Globcurrent observations. The velocities of the geostrophic currents were obtained from the obtained SSH gradients:

$$u_g = -\frac{g}{f}\frac{\partial P}{\partial y},\ v_g = -\frac{g}{f}\frac{\partial P}{\partial x} \tag{3.5}$$

$$f = 2\Omega sin\phi,\ \Omega = \frac{2\Pi}{day} \tag{3.6}$$

where P is the horizontal pressure field from the SSH, g is the gravity constant, f is the Coriolis parameter, ϕ is the latitude, and Ω is the angular velocity of the Earth. The annual RMS_{geo} ($m^2\ s^{-2}$) means are: 0.42 ($\sigma = 0.03$), 0.42 ($\sigma = 0.02$), and 0.44 ($\sigma = 0.03$) for LS, DA, and Observations, respectively. The annual G_R (-) means are: 0.73 ($\sigma = 0.09$), 0.92 ($\sigma = 0.08$), and 0.95 ($\sigma = 0.05$) for LS, DA, and Observations, respectively. The RMS did not include the area comprising 3°S to 3°N in the vicinity of the Equator.

The following results correspond only to the DA using the NN. Figure 3.6 shows the annual spatial mean RMSE, with respect to the validation datasets of SST and SSH for the RFS and analysis. After the assimilation, the differences in the advection route of the NBC and NBC rings have been clearly minimised but less so in the coastal regions or the area near the Lesser Antilles (Fig. 3.6a and 3.6b). In these coastal regions, the assimilation of SST is better in those areas, but they remain the most challenging regions (Fig. 3.6c and 3.6d). The annual mean correlation coefficient improves from 0.52 to 0.80 in SSH when compared to AVISO, and from 0.56 to 0.72 in SST when compared to GHRSST.

The improvements of the time averaged RMSE for the fast DA framework are 0.04 m (41 %) and 0.36 °C (49 %) for SSH and SST, respectively. It is encouraging to note that the time averaged RMSE of the assimilated fields are similar to those of global reanalysis, including DA, products like HYbrid Coordinate Ocean Model (HYCOM) (Fox et al. 2002, Cummings 2005).

The annual spatial mean RMSE of the horizontal U and V components of the currents is illustrated in Fig. 3.7. Although, U and V were not assimilated, the overall spatial RMSE is reduced and the time averaged RMSE improves by 0.07 m s^{-1} (13 %) in U and by 0.09 m s^{-1} (15 %) in V. The reduction of RMSE is more noticeable in the open ocean.

Less significant improvement can be seen around the mouth of the Amazon River for the four variables.

Table 3.1: DA annual means of RMSE, correlation coefficient and J_{fit} for the LS, NN and Persistence (P) experiments. The correlation coefficient r and J_{fit} are unitless and only shown for the assimilated variables and the DA experiments.

Variable	**RMSE**			r **(-)**			J_{fit} **(-)**		
	LS	NN	P	LS	NN	P	LS	NN	P
SSH (m)	0.04	0.05	0.08	0.89	0.84	0.65	-	0.70	1.38
SST (°C)	0.28	0.35	0.71	0.84	0.74	0.51	-	0.53	1.09
U (m s^{-1})	0.64	0.46	0.76	0.27	0.36	0.20	-	-	-
V (m s^{-1})	0.68	0.53	0.77	0.26	0.35	0.07	-	-	-

Figure 3.8 shows a snapshot of the surface currents for September 7th, 2010 for three cases: observations (Globcurrent); the ensemble-mean analysis using the NN; the LS solution. The annual RMSE for U are 0.41 m s^{-1} (r = 0.56) and 0.77 m s^{-1} (r = 0.48) for DA using NN and LS, respectively. The annual RMSE for V are 0.51 m s^{-1} (r = 0.49) and 0.64 m s^{-1} (r = 0.38), for DA using NN and LS, respectively.

Figure 3.9 is a snapshot of January 11th, 2010 comparing observations, RFS, and the analysis. This day was chosen as it portrays the main features of the NBC and how the DA changes the fields of the variables. The velocity fields of total and geostrophic currents are represented by the vector magnitude of their respective eastward and northward components. The spatial representation of the DA of SSH shows the bigger similarities with the altimetry dataset, where the main features like the NBC ring can be easily recognised along with the retroflection of the NBC in the eastern part of the grid. The SST from RFS seems to capture enough information from the observations in order to change its spatial patterns. The response and modification of U and V is the result of the assimilation of SSH and SST within the reduced-space. The retroflection of the NBC and its rings appear in the same location for the geostrophic currents and total surface currents. This positioning means that the SSH is also correlated to the currents since the geostrophic currents are extracted from the SSH gradient. The NBC is captured and well represented by the currents of the assimilated system when compared to Globcurrents.

The main difference between the variables in the assimilated system and their validation datasets are in the Amazon River region of influence. In general, the variables present great variability in the mouth of the Amazon River. After the analysis, the representation of the dynamics of the system is physically plausible showing features like the retroflection of the NBC and the signal of the NBC ring in the NW part of the grid across the four variables in the analysis.

3.3.6 Speed of the framework

A critical aspect of this framework is the gain in computer time. The EOF analysis on FS is performed once in 20 cores of an Intel Xeon E5-4650v2 (2.4 GHz) processor and takes $\sim$ 90 seconds. Typically, the assimilation of SSH and SST takes $\sim$ 20 seconds per model day, including the forward stepping with the NN and daily EnKF DA. A year-long simulation takes between 2-3 hours.

Other sensitivity runs showed that the assimilation of one variable takes $\sim$ 10 seconds, while the joint assimilation of four variables takes up to 60 seconds per model day. While the ROMS simulation runs on 256 cores of Intel Xeon (2.5 GHz) processors, the fast DA runs on a single core of an Intel Xeon E5-4650v2 (2.4 GHz) processor.

In relative times, running one model day of the 10km resolution NBC ROMS model on one processor takes 30.52 hours. Our framework has a speed up of $\sim$ 1800 times when assimilating four variable and up to $\sim$ 11000 when assimilating one variable. On a daily time-step basis, the fast DA framework presented here is computationally more efficient outperforming the full-space ROMS run of the NBC region by 3 to 4 orders of magnitude.

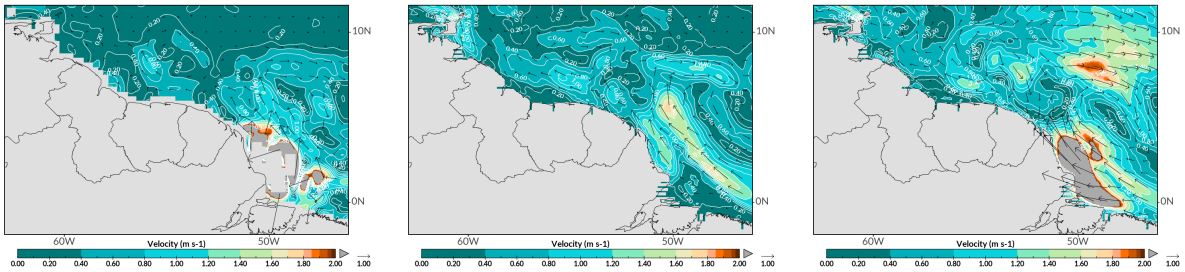

(a) Globcurrent (b) EnKF(NN) (c) LS

Figure 3.8: The surface currents represent the magnitude and direction of U and V. Globcurrent total currents at the surface (left column), ensemble-mean of the assimilated reduced-space using EnKF and NN (middle column), and reconstructed Least Squares (LS) solution (right column). Snapshot of September 7^{th}, 2010 (pink shade in Fig. 3.4 and 3.5). The validation dataset is Globcurrent (0.25-degree resolution).

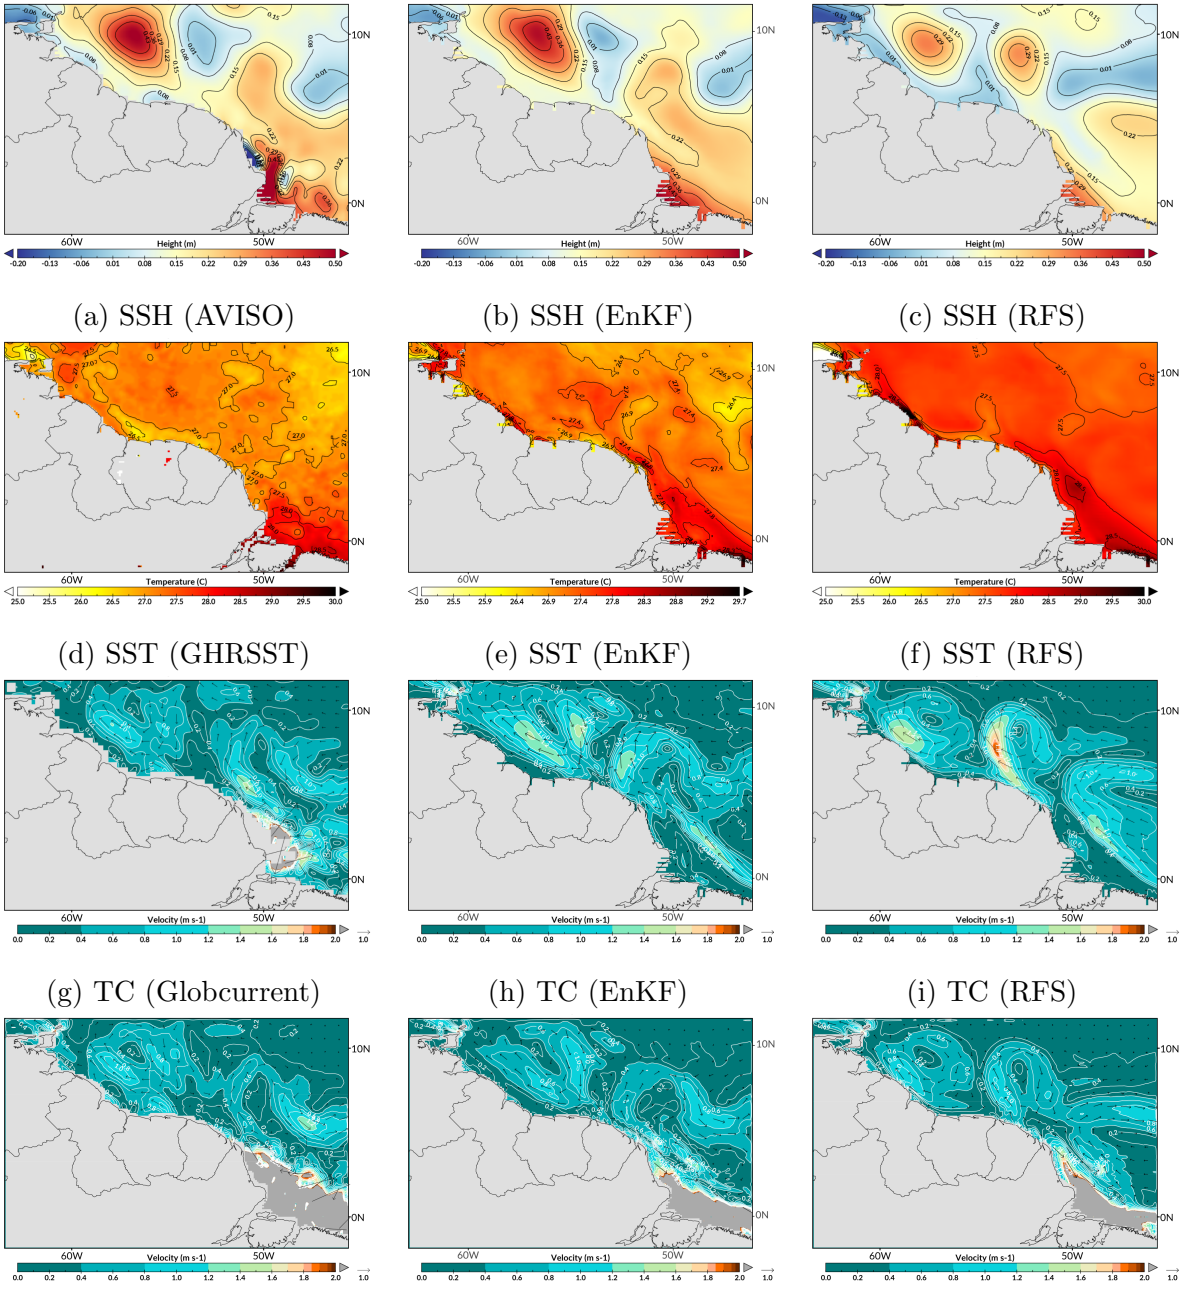

(a) SSH (AVISO) (b) SSH (EnKF) (c) SSH (RFS)

(d) SST (GHRSST) (e) SST (EnKF) (f) SST (RFS)

(g) TC (Globcurrent) (h) TC (EnKF) (i) TC (RFS)

(j) GC (Globcurrent) (k) GC (EnKF) (l) GC (RFS)

Figure 3.9: Validation datasets (left column), ensemble-mean of the assimilated reduced-space using EnKF (middle column), and reconstructed full-space model (RFS). Snapshot of January 11th, 2010 (yellow shade in Fig. 3.4 and 3.5). This day shows the comparison between the validation datasets and the ensemble-mean of the EnKF on the reduced-space. The validation datasets are SSH from AVISO (0.25-degree resolution), SST from GHRSST (10 km resolution), total currents (TC) from Globcurrent (0.25-degree resolution) and geostrophic currents (GC) from Globcurrent (10km resolution). The surface currents represent the vector magnitude and direction of U and V.

3.4 Discussion

The results can be discussed in terms of the effectiveness of the dimension reduction algorithm, the performance of the DA using the EnKF and how the characteristics of the NBC region are represented in this scheme, and finally the execution speed and robustness of the fast DA framework.

The DA diminishes the RMSE. This diminishing is most clear when we observe SSH and SST since these are the state variables directly constrained by data. A spatial representation of the analysis shows that it keeps enough information from the reduced-space to be physically plausible. The reconstructed analysis also presents a correlation across variables. The correlation suggests that the reduced-space DA solution has value. There are three distinguishable problematic sections in the study area when we perform DA: The Lesser Antilles islands, the coastline of Northern South America and the outlet of the Amazon River into the Atlantic Ocean.

The islands located in the Lesser Antilles are a problematic area since these serve as a natural barrier to the NBC and NBC rings, where the latter collide and dissipate. The DA in this area does not perform as well as in the open ocean, due to the coastline of the islands where the satellite data are limited and prone to larger measurement errors. The same phenomenon can be observed along the coastline of the Northern region of South America, where is particularly true for the availability of altimetry observations since the satellite missions do not perform well close to the coast (Fig. 3.6). Additionally, the altimetry dataset used in this study has a lower resolution (0.25-degree) than the product used to assimilate SST (10 km).

The freshwater discharge affects the area near the mouth of the Amazon River. The discharge also affects the performance of obtaining observations around this area. The mouth of the Amazon River is a troublesome area where the DA performs poorly, not being able to reduce the annual spatial mean RMSE considerably. The poor performance

might also be due to the EOF not capturing the variability around these areas. The LS solution also does not capture the high variability of the Amazon River mouth. The DA of SST performs better in this area because of the higher original resolution of the observations. The DA of SST performs specifically better in the coastlines, but overall the annual spatial mean of assimilation of SSH is remarkably better. The altimetry sea level is a tractable variable, with less noise, displaying strong features such as the formation of the NBC rings, and retains more explained variability when the dimension reduction is applied.

Despite being part of the reduced-space state vector and the NN dynamics, U and V data were not assimilated in this study. With the selected number of EOFs, these variables are only affected by the changes in the reduced-space during the DA, and when the analysis is reconstructed back to the full-space. The currents simulation only improves marginally. The average RMSE and spatial r improve compared to the Globcurrent dataset, but the same problems in the coastlines, mouth of the Amazon River and the Lesser Antilles appear. Globcurrent products are also derived from altimetry (Rio et al. 2014*b*). The impact of assimilating SSH and SST on U and V is in concordance with previous studies (Phillipson & Toumi 2017), where the currents are corrected once SSH is assimilated, but the influence of assimilating of SST is less. The geostrophic currents obtained from the SSH analysis present similar spatial patterns to the analysis surface currents.

The result of the application of the EnKF is dependent on the construction of the error covariance matrix $\mathbf{R}$. The importance of $\mathbf{R}$ has been discussed previously in Stewart et al. (2007). Here, we assume that the error associated with the observations is uncorrelated and that is the same for the entire grid and works as a simplification, meaning that $\mathbf{R}$ might be under or overestimating the error covariance with more or less weight, respectively, compared to the error covariance of the model. Moreover, recent studies have shown that introducing observation error correlations improve DA results (Waller et al. 2016, Miyoshi et al. 2013). If information of detailed errors of the gridded obser-

vational datasets were available these could be easily applied to this framework and into the assembly of $\mathbf{R}$. The goal of this framework of reduced dimensionality models is to perform fast DA. Therefore, we need to avoid a bottleneck when calculating the inversion in the Kalman gain using of the Sherman-Morrison-Woodbury formula (eq. (2.9)).

In this framework the background error covariance matrix $\mathbf{B}$ in the reduced-space can be explicitly calculated and this allows us to easily apply a multiplicative inflation to increase the spread of the ensemble of the NN forecast within the EnKF. Explicit calculation of $\mathbf{B}$ in high-dimensional systems can restrict the resolution of the DA. The ensemble method also helps to provide a better simulation of the forecasts errors (Bannister 2008), and in this case are computationally cheap to generate. This is only valid in the reduced-space as a correct representation of $\mathbf{B}$ in the full-space will be computationally expensive and would defeat the purpose of having a fast DA framework.

The reduced-space allows us to perform a fast DA. The comparison between the observations and the model state occurs in the observational space, thus more observations will slow down the procedure. Working with a reduced-space solution is a cheaper task than working with high-dimensionality systems and therefore powerful processing power (supercomputer) is not necessary. Certain features may not appear in the reduced-space but are present in the full-space. More EOFs could be retained but with the caveat that the framework will slow down because of the internal multiplications of a now bigger matrix. A quick calculation of the DA times assimilating SSH and SST with 50, 100, 150 and 200 EOFs shows that the scaling of the calculations time is $\sim N_{eof}^2$. If we were to include another field like surface salinity or 3D fields, then the scaling of the calculations will not change as adding more information from the ROMS model solution would only increase the number of EOFs that retain enough variance from the original model. For example, if we add surface salinity and then we apply a EOF analysis. With this addition, we would need more than 100 EOFs to retain the same amount of retained variance that SSH, SST, U and V retain with 100 EOFs. Therefore, the scaling depends on the

number of EOFs.

The DA also depends on the short-range forecast to create the forecast estimates. A comprehensive search for the optimal hyperparameters of the NN was performed. Any other configurations will change this forecast. If another number of EOFs were to be considered, it would produce a new reduced-space which would also change the configuration of the NN, modifying it from the number of hidden layers to the time necessary to train a new system. The length of the ROMS data also alters the training and configuration of the NN. If there is not enough data, the NN would not have enough information from the full-space and therefore the ability of replicating the physics and dynamics of the former will be compromised. There is therefore a great deal of effort required to establishing the optimal reduced-space DA model. However, once this effort has been undertaken the assimilation is fast.

For a daily time-step basis, the fast DA framework presented here outperforms the full-space ROMS run of this region by 3 to 4 orders of magnitude in computational efficiency. In terms of computing time, Frolov et al. (2009) reports a speed-up of two orders of magnitude when comparing the DA on the full-space estuary model of the Columbia River to the one performed on the model surrogate.

The robustness of the application of this framework can be assessed by comparing the results to Frolov et al. (2009) and by comparing the results of the DA with other models performing DA in the same region. The RMSE results of the assimilation of altimetry and temperature are similar to the ones presented by Frolov et al. (2009) for an estuary: between 0.04 to 0.07 m in SSH and 0.4 to 1.4 °C for SST. HYCOM has a 10 km resolution, the same as the ROMS model solution employed in this study. A study of the NBC area includes the DA of sea level anomaly, temperature and salinity into HYCOM around the western tropical and South Atlantic (Costa & Tanajura 2015). The reported RMSE of SSH and SST validation datasets range from 0.05 to 0.15 m, and 0.5 to 3.5 °C. Our fast DA framework has a similar magnitude of errors compared to these full-space systems.

Finally, having a NN for the forecasting task is useful because these systems can detect subtle features of the input data when the information travels through the different layers. It is also simple to train and there is an increasing availability of state-of-the-art and optimised libraries ready to be used.

3.4.1 Least squares and persistence experiments

The LS solution sets up a reference for the comparison of the DA solution using NN. The LS solution presents the largest misfits in the area around the Amazon river mouth and coastlines (not shown). It outperforms the DA solutions in RMSE and correlation with SSH and SST data because of the large availability of gridded data. Nonetheless, if we reconstruct using all the column vectors of D^{rs}, which include U and V, and the c coefficients found with the least squares problem, then is clear that this solution is an overfit to the restricted vectors associated to SSH and SST data. The comparison of spatial U and V between the DA using the NN and LS solutions shows that the surface velocities from the LS solution present a larger error (Fig. 3.8). The larger error of the LS solution is supported by the RMSE and r time series of U and V in Fig. 3.4 and 3.5. Here, the DA using the NN outperforms the LS solution.

The G_R is a measure of how geostrophic a flow is. The annual means of G_R show that the DA solution is closer to the Observations, where this is 3 % larger than the DA solution. When compared to the LS, the observations present a 23 % larger G_R annual mean. The values of G_R suggest that the DA solution preserves this ratio much better than the LS solution and demonstrates that there is a meaningful physical constraint in the model surrogate in the DA method rather than in the LS solution. The values of G_R greater than 1 are related to the high geostrophic velocities obtained from the horizontal pressure fields near the Equator. In this region, the Coriolis parameter, f, is very small which might result in unrealistically high geostrophic currents. We did not include information from 3°S to 3 °N, in order to avoid values greater than 1 that can be obtained if the area

close of the Equator is considered.

The persistence experiment allows us to assess the role of the NN in the forecast step within the DA. Since the persistence does not increase the spread of the ensemble from the previous analysis, apart from the inflation, we would expect the NN to grow the spread of the ensemble in the one-day forecast. As aforementioned, the mean of the diagonal of $\mathbf{H}\tilde{\mathbf{B}}\mathbf{H}^{\mathbf{T}}$, using the NN, converges and oscillates to a value in the same order of magnitude as the diagonal mean of $\mathbf{R}$. A bigger spread of the ensemble means that the EnKF synchronises better with the true system trajectory (Hunt et al. 2007, Whitaker & Hamill 2002). From Table 3.1 it is clear that the DA using NN outperforms the persistence experiment resulting in lower RMSE and higher correlations with the validation data, and a J_{fit} less than 1. Thus, the NN also outperforms the persistence experiment, and yields better results after the DA, since it can better simulate the dynamics of the full-space state.

3.4.2 Optimal hyperparameters

We used grid search to find the optimal hyperparameters to tune the NN. Grid search is among the most widely used methods for hyperparameter tuning and optimisation (Bergstra & Bengio 2012). The choice of using grid search over random search is the computational power available. By using grid search in a high-performance computer we can find the optimal hyperparameterisation by simply running a large number of combinations of hyperparameters. Additionally, grid search is simple to implement and is capable to find a better minimum than manual searching with access to a computer cluster (Bergstra & Bengio 2012). Throughout this chapter we used 30000 combinations. Larochelle et al. (2007) explains that a random search of hyperparameters is able to locate the minimum of the cost function using ~ 100 trials. Therefore, it is likely that the same or similar optimisation will be found using 20000 or 40000 combinations with grid search, and therefore the performance of the short-range forecasts would be really similar.

In certain cases, random search has been proven to outperform grid search using less combinations of hyperparameters (Bergstra et al. 2011). The use of random search or other optimisation methods can be explored in further detail in future work, as these would reduce the computational cost of training the NN. Additionally, future work can include a sensitivity analysis to the different methods and number of combinations used to find the optimal hyperparameters.

3.5 Conclusion

We present the first application in the open ocean of a novel fast DA framework using a reduced-space. The framework proved to be much faster than full-space DA schemes by three to four orders of magnitude, with similar accuracy to similar global reanalyses. The dimension reduction managed to capture most of the explained variance. Since the NN is data-driven, a longer training dataset will be beneficial as it might include extreme events that are not present in the current training dataset.

Only SSH and SST data were assimilated. There was a clear improvement when these variables were validated against satellite gridded data. Although U and V were not assimilated, the reconstructed reduced-space analyses present a higher correlation coefficient and lower time-averaged RMSE when compared to the Globcurrent synthetic velocity product.

Finally, this framework is not limited to outputs from ROMS nor to the four variables studied here, and therefore any global analysis product like HYCOM or CMEMS could be used as the underlying full-state. The framework can be applied to other regional ocean areas around the globe, making it the first step to fast forecasting applicable in real emergency situations where new data is quickly available and an immediate prediction of, for example, currents, is necessary. It could be a powerful complimentary tool to full models when they are not readily available.

Chapter 4

Fast forecast of sea surface variables using neural networks

4.1 Introduction

The prediction of ocean variables like SSH, SST, and U and V, can be studied using physics-based numerical models, or more recently with data-driven methods. While numerical methods can be more accurate, they are also computationally expensive and slow for some applications. Data-driven methods are of particular interest due to the latest advances in machine learning and artificial intelligence. Studies that have explored a data-driven approach for ocean variables prediction include pattern searching (Bakun & Broad 2003), regression (McDonnell & Holbrook 2004), and support vector machines (Lins et al. 2013) among others. NNs are of particular interest since they provide easiness to setup, and can learn subtle features from the data fed into them (Goodfellow et al. 2016).

The speed of simulation is also important. Ocean numerical models are limited by the computation resources available, and reduced models are thus a complementary tool. In

the last few decades, reduced ocean models working as surrogates or emulators of the full model have been used as a computationally fast and cheap option with an acceptable loss of accuracy. The approach provides a surrogate without the high-dimensionality or complexity of the full model. Several studies have used emulators to replicate or approximate a full-space model. Fast forecasts are needed in emergency situations like search and rescue missions and oil spills. To achieve this, the forecast is initialised from optimal initial conditions.

This study presents a data-driven approach using a reduced regional ocean model; it builds on Frolov et al. (2009) as it is applied for the first time in an open ocean setting, and goes beyond Chapter 3 by adding a fast forecasting step using a trained NN. The forecasting step follows the DA step and creates a 6-day prediction of surface ocean variables.

4.2 Methods

4.2.1 The modified framework including forecast

The methods are based on and they are a continuation of the framework described in Chapter 3 where a fast DA scheme on a reduced regional ocean model is described.

Firstly, a dimension-reduction and NN training is needed. A full-scale dynamical ocean model is chosen. The selected ocean model is based on ROMS. The model solution is a 10-year (2001-2010), daily output run with a horizontal resolution of 1/12-degree. The year 2010 used as a validation year in the previous chapter is added to the analysed time period in order to create the full-scale model. The full-scale model outputs several physical variables as salinity, ocean velocity, sea surface height and sea surface temperature, among others. To use this framework, it is necessary to reduce the high dimensionality of the full-space state (FS). The dimension reduction is achieved by performing an EOF analysis

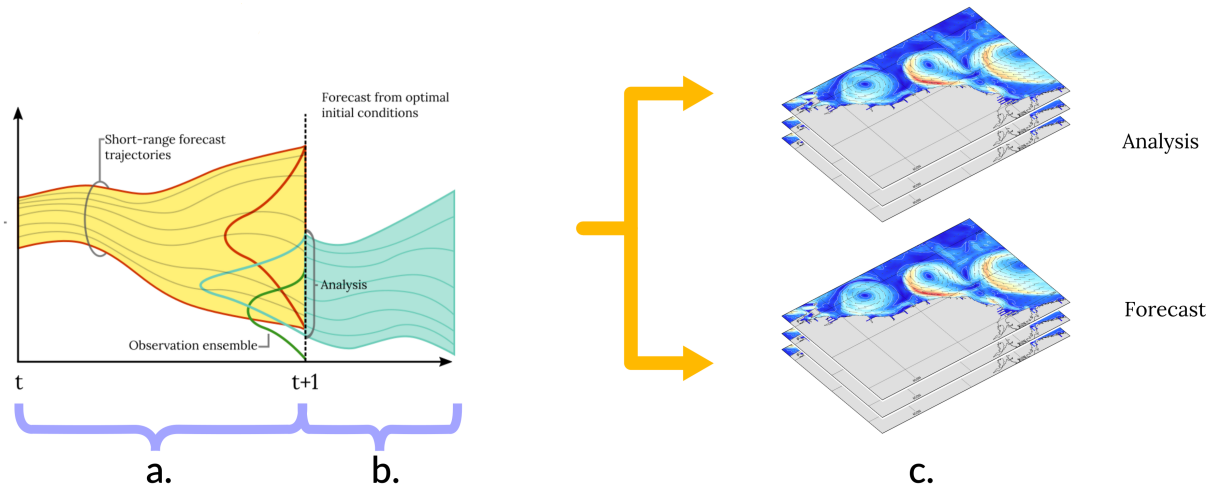

Figure 4.1: Schematic representation of the DA+NN framework. **a.** Data assimilation. The EnKF is used to perform the DA. An ensemble of forecasts (yellow shaded area) is propagated to the next time step where the forecast (prior, red) probability is compared to the observations ensemble and produces the analysis (or posterior, blue). **b.** Forecasts using the trained NN. The forecasts are created quickly by the NN from the mean of the analysis as the initial state, and (schematic modified from Bauer et al. (2015). **c.** Reconstruction of the newest update reduced-space state into a full-space state. Finally, the updated reduced-space state can be reconstructed from either the analysis provided by the DA or after creating the forecasts with the NN.

on the FS. Frolov et al. (2009) present an algorithm to divide the large EOF problem into smaller EOF analysis into sub-grids of the original system. In this study, the EOF analysis was performed over the entire data matrix using the Python software package `eofs` (Dawson 2016). The standardised zero mean FS $\bar{\mathbf{D}}^{\mathbf{fs}}$ is:

$$\begin{aligned} \bar{\mathbf{D}}^{\mathbf{fs}} &= \mathbf{P}\mathbf{\Pi} \\ \bar{\mathbf{D}}^{\mathbf{rs}} &= \mathbf{P}_{N_{eof}}\mathbf{\Pi}_{N_{eof}} \end{aligned} \tag{4.1}$$

where: $\mathbf{P}$ are the Principal Components (PCs); $\mathbf{\Pi}$ are the EOFs; N_{eof} is the chosen number of EOFs; $\bar{\mathbf{D}}^{\mathbf{rs}}$ is the zero mean standardised reconstruction of the full-space state truncated at N_{eof}. The resulting reduced-space state consists of the first N_{eof} = 100 EOFs. The $\mathbf{P}_{N_{eof}}$ is then fed into a feed-forward NN. The NN is necessary to extract enough information from the full-space state and emulate its dynamics. A detailed explanation of the configuration of the NN is discussed in section 4.2.2.

Secondly, a DA step is required. We use the EnKF (Evensen 2003) to assimilate observa-

tions into the system. The DA is fast because it is performed in the reduced-dimension space. While the DA occurs in the reduced-dimension space, the comparison of observational values to the model values occurs in the observational space. The dimensions of the observational space depend on the number of observations that are to be assimilated. Therefore, more observations will slow down the calculations. The EnKF needs a short-range forecast to generate the prior and to be compared with the observations ensemble. The forecast is obtained from the previously trained NN. The EnKF is sequential and it does not require a supercomputer to execute.

Here, we introduce a step between the DA and the reconstruction: forecast. The forecast is generated using the previously trained NN. The NN is initialised (input layer) with the analysis obtained from the DA. The forecast is free-running meaning that the output is then used to initialise the following time step. We produce a 6-day forecast from each DA analysis. The first iteration, Day+0 (assimilated day) is used as the input layer and produces a 1-day forecast as the output. Subsequently, the 1-day forecast (Day+1) is used as the input layer to generate the 2-day (Day+2) forecast as the output, and progressively until we obtain a 6-day (Day+6) forecast. Therefore, the NN does not require any external inputs to produce the forecast, making it a rapid forecast generating procedure.

Finally, we need to reconstruct to the full-space. The results from the DA are represented in the reduced-dimension space. Thus, it is necessary to reconstruct back into the full-space in order to make performance assessment and comparison with real data using the function $\mathbf{\Pi}_{N_{eof}}$. A schematic diagram of the framework is shown in Fig. 4.1. From here, our framework will be referred to as DA+NN.

4.2.2 Neural networks for forecasts

It is necessary to replicate the dynamics of the full-space in the reduced-space. We achieve this by training an NN. The chosen architecture is the Multi-Layer Perceptron which is divided into three layers: an input layer, hidden layers, and output layer (Fig. 2.1). The NN is trained to predict the time-step $t+1$ of $\mathbf{P}_{N_{eof}}$ feeding the time-step t into the input layer. To find the best configuration of the hyperparameters of the NN, an exhaustive search was performed using the Scikit-learn library in Python. The search includes activation functions, learning rates, number of hidden layers and nodes and training algorithms. The training divides the time series in 7, 1.5, and 1.5 years for training, validation and testing, respectively. After the optimal parameters have been found, the NN is trained in MATLAB. The training finds the minimum of the cost function or stops it to prevent overfitting. The cost function is the MSE between the targets and predicted values. Initially, the NN is trained using an open-loop where it constantly receives feedback (input and target) from the training dataset. It is followed by a closed-loop training where the NN does not have external input rather than the NN itself. Having a trained NN is useful when rapid forecasts are needed, as the NN and the input depends only on the previous output.

The trained NN is used inside the EnKF for the short-range forecast, and for the free-running forecast after the analysis.

4.2.3 Reanalysis data and forecast: CMEMS and UK MetOffice

Data assimilated and forecasting data were obtained from the archived repository of forecasts of the Global Analysis Forecast Phys 001 015 product distributed by Copernicus Marine and Environment Monitoring Service (CMEMS) (EU Copernicus Marine Service 2015). The product is a daily-mean 0.25-degree resolution dataset that includes 43 depth levels including, among others, sea surface height, temperature, and sea water velocity.

The 3D and 0.25-degree resolution global ocean forecasts are provided by the UK MetOffice weakly coupled DA and forecast system, and it is updated daily. The reanalysis uses the MetOffice Unified Model v10.6 atmosphere configuration at 40 km resolution coupled to a NEMO v3.4 ocean model and to CICE v4.1 sea ice model. The atmospheric DA uses 4D-VAR, while the ocean DA uses a specifically developed 3D-VAR scheme for NEMO. The processing level of this dataset is L4. Data of July 2015 and January 2016 were obtained from the archive of forecasts. Since the global analysis forecast is produced by the UK MetOffice, this dataset will be referred to as UKMO.

Here, we use a skill score (Carrier et al. 2014) relative to the UKMO forecast to validate the forecast generated with the DA+NN framework. The skill score (SS) is defined by:

$$SS = 1 - \frac{RMSE_{DA+NN}}{RMSE_{UKMO}} \tag{4.2}$$

where $\mathrm{RMSE_{DA+NN}}$ and $\mathrm{RMSE_{UKMO}}$ are the spatial RMSE of the forecasts using the DA+NN framework and UKMO, respectively, in regard to the observational datasets. When the DA+NN has a lower RMSE than UKMO, then SS is greater than zero; when the RMSE is the same or higher, SS is equal or less than zero. The SS is compared from Day+0 to Day+6.

4.2.4 Experiments

The list of experiments is presented in Table 4.1. All the listed experiments use the surface sub-set 10-year (01/01/2001 - 31/12/2010) daily hindcast ROMS simulation of the North Brazil region as numerical data. Therefore, the EOF analysis is applied to the 10-year simulation and $\mathbf{P}_{N_{eof}}$ obtained from this reduction is used to train the NN. Each of the experiments were performed for July 2015 and January 2016 separately. For assimilation and validation, we use the data and measurement errors described in

Section 2.3. The mean dynamical topography (MDT) was calculated from the ROMS model climatology and added to the MSLA previous to the assimilation. The synthetic velocities from Globcurrent are 3-hourly and were averaged to obtain a daily mean. The dataset has a 0.25-degree resolution and includes the geostrophic and Ekman currents at the surface level.

Table 4.1: Experiments names and description of variables included and assimilated. All the observational data is gridded. MDT is the mean dynamical topography from the ROMS model solution. All experiments were conducted independently for July 2015 and January 2016.

Experiment name	**Variables included**	**Data assimilated**
HT	SSH	AVISO+MDT
	SST	GHRSST
HTc	SSH	AVISO+MDT
	SST	GHRSST
	U	-
	V	-
UV	U	Globcurrent v.2 and v.3
	V	Globcurrent v.2 and v.3
All	SSH	AVISO+MDT
	SST	GHRSST
	U	Globcurrent v.2 and v.3
	V	Globcurrent v.2 and v.3

4.3 Results

For the HT experiment, the exhaustive search (31500 combinations) of the hyperparameters of the NN suggests using ADAM (Kingma & Ba 2014) as the training method, a single hidden layer with 450 hidden nodes, a learning rate of 0.0137, and a rectified linear (ReLu) (Bengio et al. 2015) activation function in the hidden layer.

For the HT and ALL experiment, the exhaustive search (31500 combinations) of the hyperparameters of the NN suggests using scalar conjugate gradient (Shewchuk et al. 1994) as the training method, a single hidden layer with 151 hidden nodes, a learning

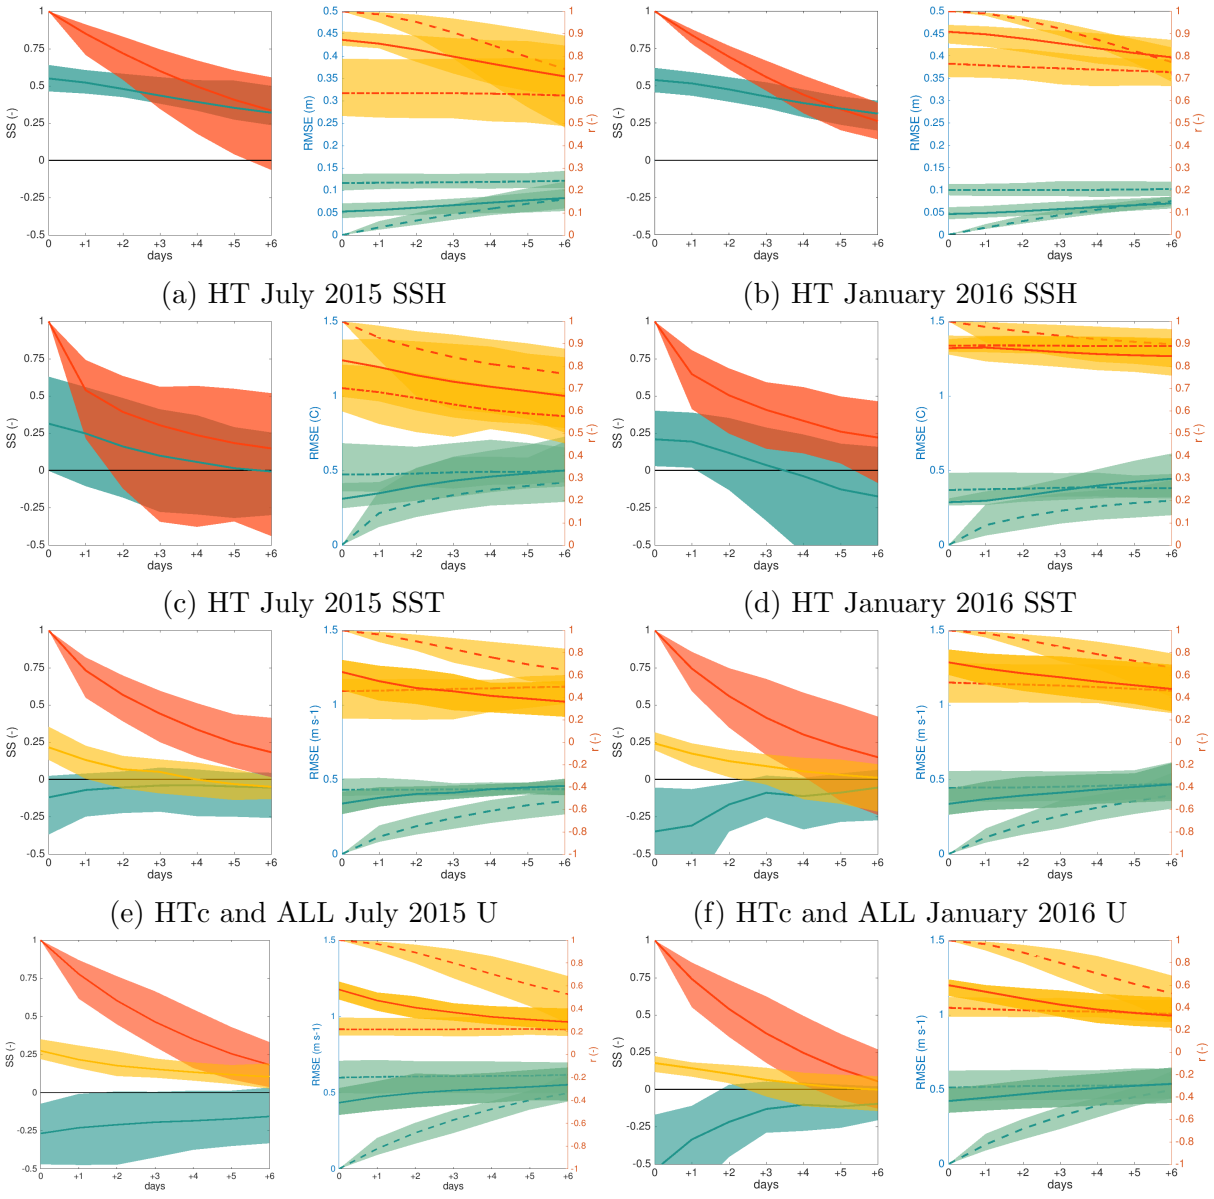

(a) HT July 2015 SSH

(b) HT January 2016 SSH

(c) HT July 2015 SST

(d) HT January 2016 SST

(e) HTc and ALL July 2015 U

(f) HTc and ALL January 2016 U

(g) HTc and ALL July 2015 V

(h) HTc and ALL January 2016 V

Figure 4.2: Time series of the ensemble spread of forecasts for SSH, SST, U and V. First column: July 2015; second column: January 2016. Shaded areas represent the 95% interquantile range. On each column, **Left-hand side:** Ensemble spread of the skill score of DA+NN forecasts (blue); and persistence of the observations from Day+0 (red). Solid lines are the ensemble means. **Right hand side:** ensemble spread of RMSE (green, left axis) and correlation coefficients (yellow, right axis). Lines represent the ensemble means. DA+NN (solid line); UKMO (dash point line); observations persistence (dash line). Day+0 is the day of the analysis, while Day+1 to +6 are the forecasts. Addition to the U and V figures the ensemble spread of the skill score of the ALL experiment (yellow); the ensemble spread of the RMSE and correlation coefficients (solid line) of the ALL experiment. The original spread includes the 31 days of sequential assimilation.

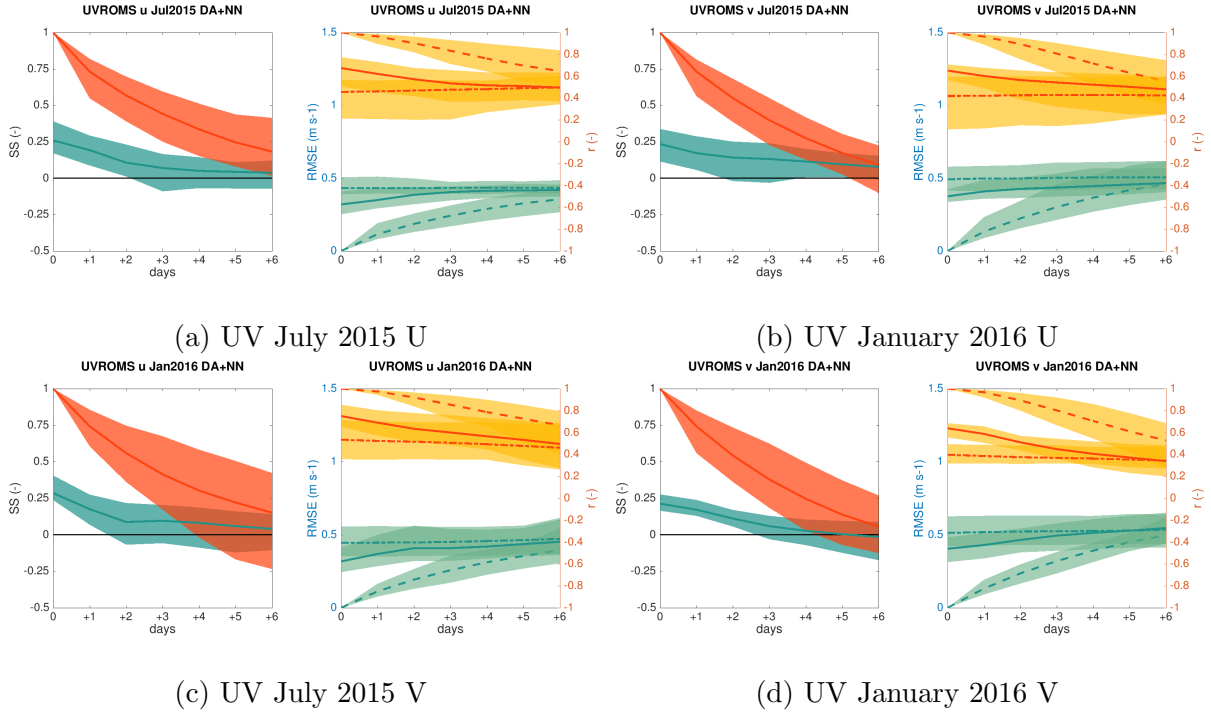

(a) UV July 2015 U

(b) UV January 2016 U

(c) UV July 2015 V

(d) UV January 2016 V

Figure 4.3: Time series of the ensemble spread of forecasts for U and V in the UV experiment. First column: July 2015; second column: January 2016. Shaded areas represent the 95% interquantile range. On each column, **Left-hand side:** Ensemble spread of the skill score of DA+NN forecasts (blue); and persistence of the observations from Day+0 (red). Solid lines are the ensemble means. **Right hand side:** ensemble spread of RMSE (green, left axis) and correlation coefficients (yellow, right axis). Lines represent the ensemble means. DA+NN (solid line); UKMO (dash point line); observations persistence (dash line). Day+0 is the day of the analysis, while Day+1 to +6 are the forecasts. The original spread includes the 31 days of sequential assimilation.

rate of 0.001 and a linear activation function in the hidden layer.

The search of hyperparameters for the UV experiment suggests using scalar conjugate gradient (Shewchuk et al. 1994) as the training method, a single hidden layer with 230 hidden nodes, a learning rate of 0.0001 and a linear activation function in the hidden layer.

The time to create the reduced model, the DA per day and a 6-day forecast are $\sim$ 90 seconds, $\sim$ 60 seconds and $\sim$ 20 seconds, respectively, using a single core of an Intel Xeon E5-4650v2 (2.4 GHz) processor. The optimisation of the hyperparameters of the NN is performed in 12 cores of Intel Xeon (2.5 GHz) processors taking up to 50 hours. The framework is ready to receive new information after the NN is trained, meaning that a

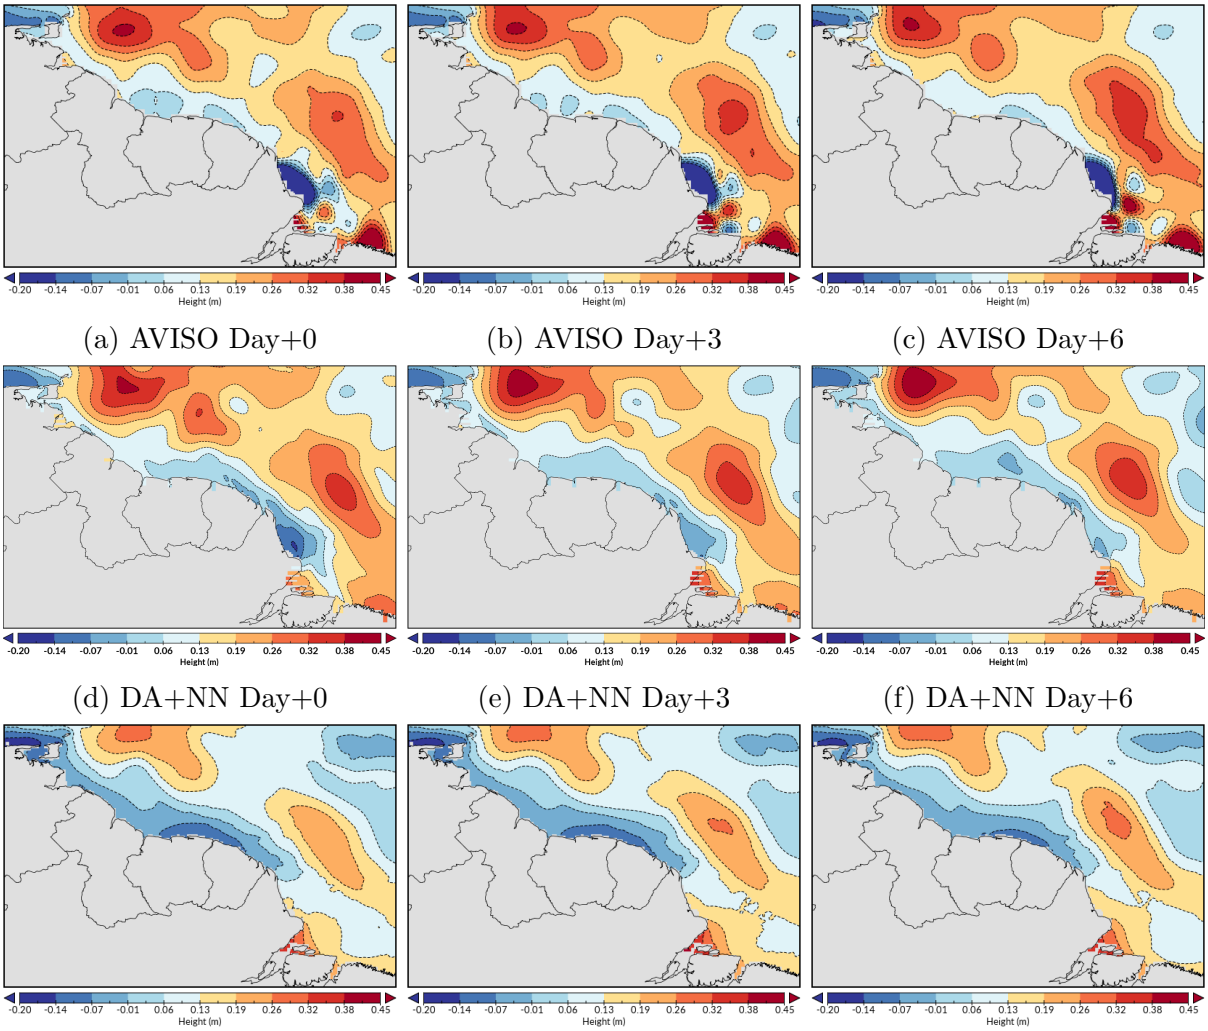

(a) AVISO Day+0 (b) AVISO Day+3 (c) AVISO Day+6

(d) DA+NN Day+0 (e) DA+NN Day+3 (f) DA+NN Day+6

(g) UKMO Day+0 (h) UKMO Day+3 (i) UKMO Day+6

Figure 4.4: Sea surface height forecast after the DA (HTc experiment) of July 12th, 2015, Day+0 corresponds to the day of assimilation, while Day+3 and Day+6 correspond to the 3-day and 6-forecast from July 12th, 2015. Rows, from top to bottom, present the AVISO (0.25-degree resolution) observational data, the DA+NN framework forecasts, and the forecasts of UKMO. The latter represents July 12th, 15th, and 18th, 2015.

forecast can be produced in a matter of seconds.

Figure 4.2 shows the time series of the ensemble spreads of forecasts of SSH, SST, U and V including the day of DA (Day+0) and the 1 to 6-day forecasts (Day+1 to Day+6). The shaded area shows the 95% interquantile range that includes the 31 days of assimilation during July 2015 and during January 2016, from now on referred to as Summer and Winter. The mean of the skill scores of the persistence of the observations (red shaded area) against the UKMO is always above the zero reference line up to Day+6. The

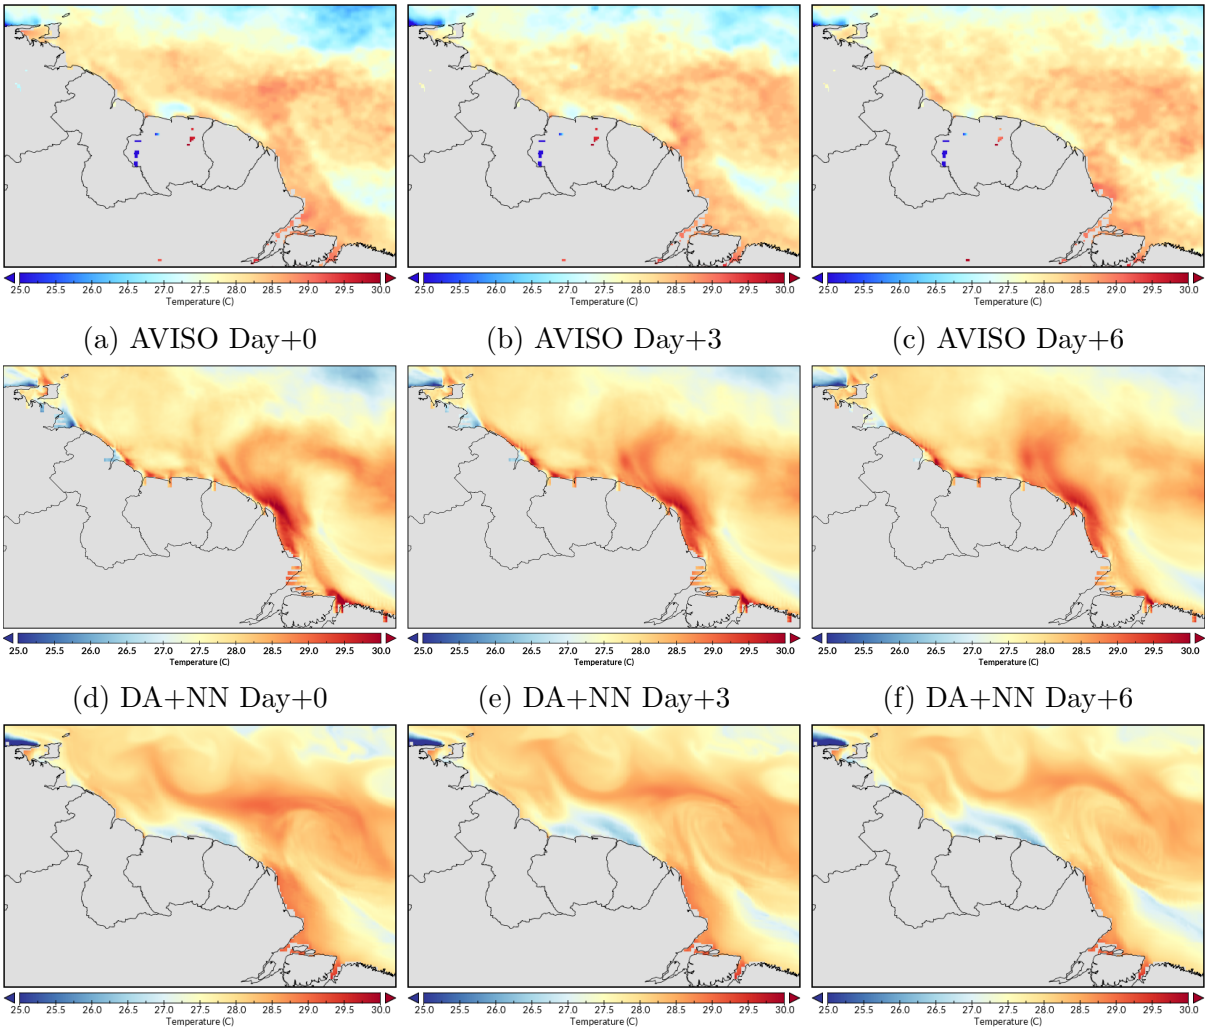

(a) AVISO Day+0 (b) AVISO Day+3 (c) AVISO Day+6

(d) DA+NN Day+0 (e) DA+NN Day+3 (f) DA+NN Day+6

(g) UKMO Day+0 (h) UKMO Day+3 (i) UKMO Day+6

Figure 4.5: Sea surface temperature forecast after the DA (HTc experiment) of July 12th, 2015, Day+0 corresponds to the day of assimilation, while Day+3 and Day+6 correspond to the 3-day and 6-day forecast from July 12th, 2015. Rows, from top to bottom, present the GHRSST (10 km resolution) observational data, the DA+NN framework forecasts, and the forecasts of UKMO. The latter represents July 12th, 15th, and 18th, 2015

positive skill score shows that the mean of the RMSE spread of the persistence is lower than the mean of the RMSE spread of the UKMO forecast. If the SS is under the zero reference line, the RMSE is higher than the RMSE of the UKMO dataset; if SS is positive, then the RMSE of DA+NN and the persistence of the observations are lower.

The HT experiment only includes and assimilates SSH and SST data. In SSH (Fig. 4.2a and 4.2b), DA+NN show a positive and decreasing SS from Day+0 to Day+6, in both Summer and Winter. At Day+0, Winter shows a mean SS 11% larger than Summer.

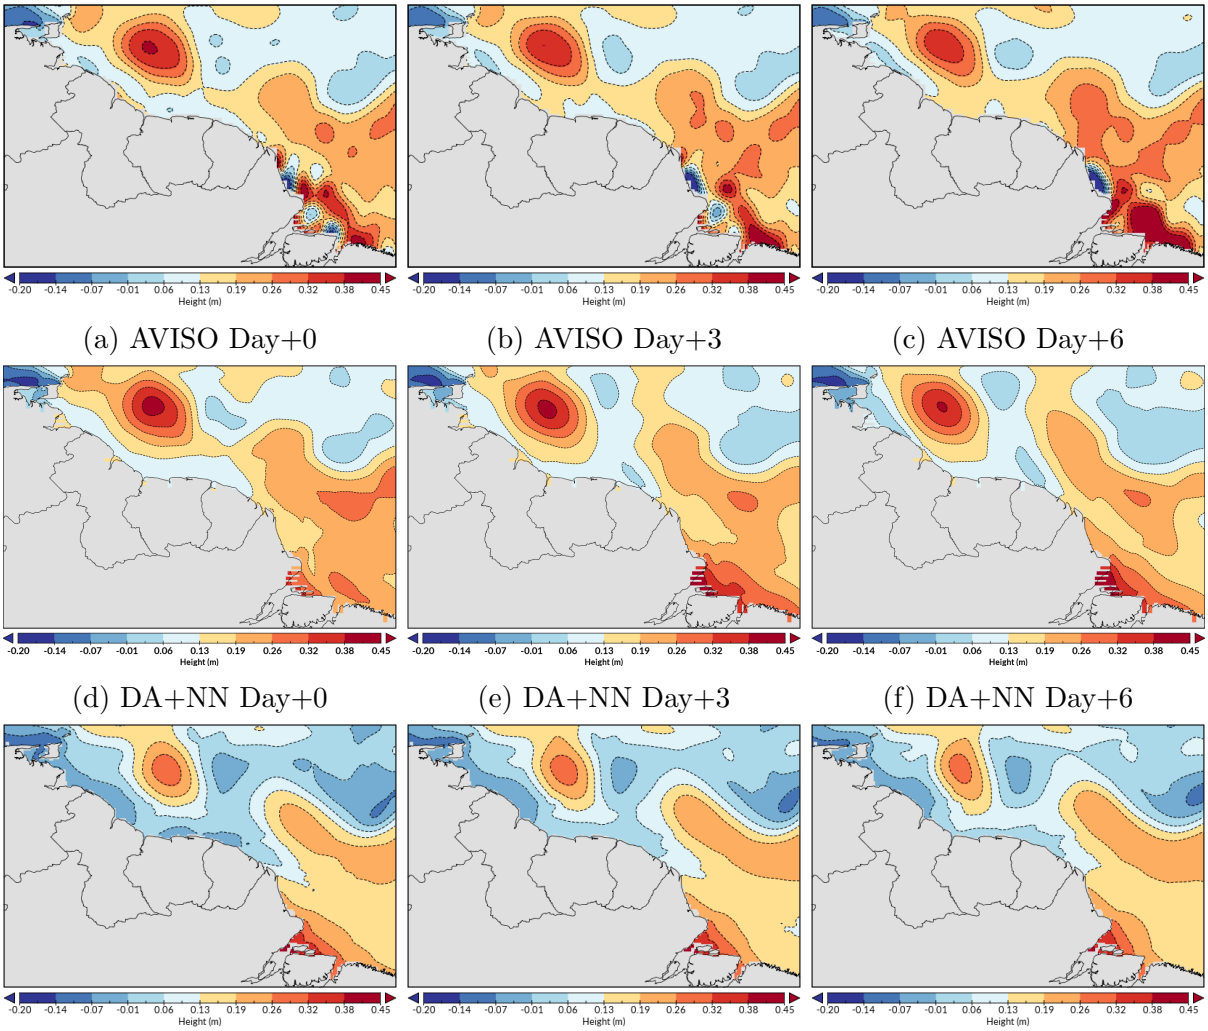

(a) AVISO Day+0 (b) AVISO Day+3 (c) AVISO Day+6

(d) DA+NN Day+0 (e) DA+NN Day+3 (f) DA+NN Day+6

(g) UKMO Day+0 (h) UKMO Day+3 (i) UKMO Day+6

Figure 4.6: Sea surface height forecast after the DA (HTc experiment) of January 29$^{\text{th}}$, 2016, Day+0 corresponds to the day of assimilation, while Day+3 and Day+6 correspond to the 3-day and 6-day forecast from January 29$^{\text{th}}$, 2016. Rows, from top to bottom, present the AVISO (0.25-degree resolution) observational data, the DA+NN framework forecasts, and the forecasts of UKMO. The latter represents January 29$^{\text{th}}$, and February 1$^{\text{st}}$, and 4$^{\text{th}}$, 2016.

At Day+6, this is 17% larger. The DA+NN RMSE spreads seem to be close to the observations and at Day+6 shows only $\sim$ 1% difference and is slightly better during Winter. At Day+0 during Summer, DA+NN shows a mean spatial correlation 37% larger than the UKMO forecasts, while at Day+6 this is only 12% larger. In Winter, these differences between DA+NN and UKMO are 27% and 13% at Day+0 and Day+6, respectively. In SST (Fig. 4.2c and 4.2d), DA+NN shows a decreasing SS from Day+0

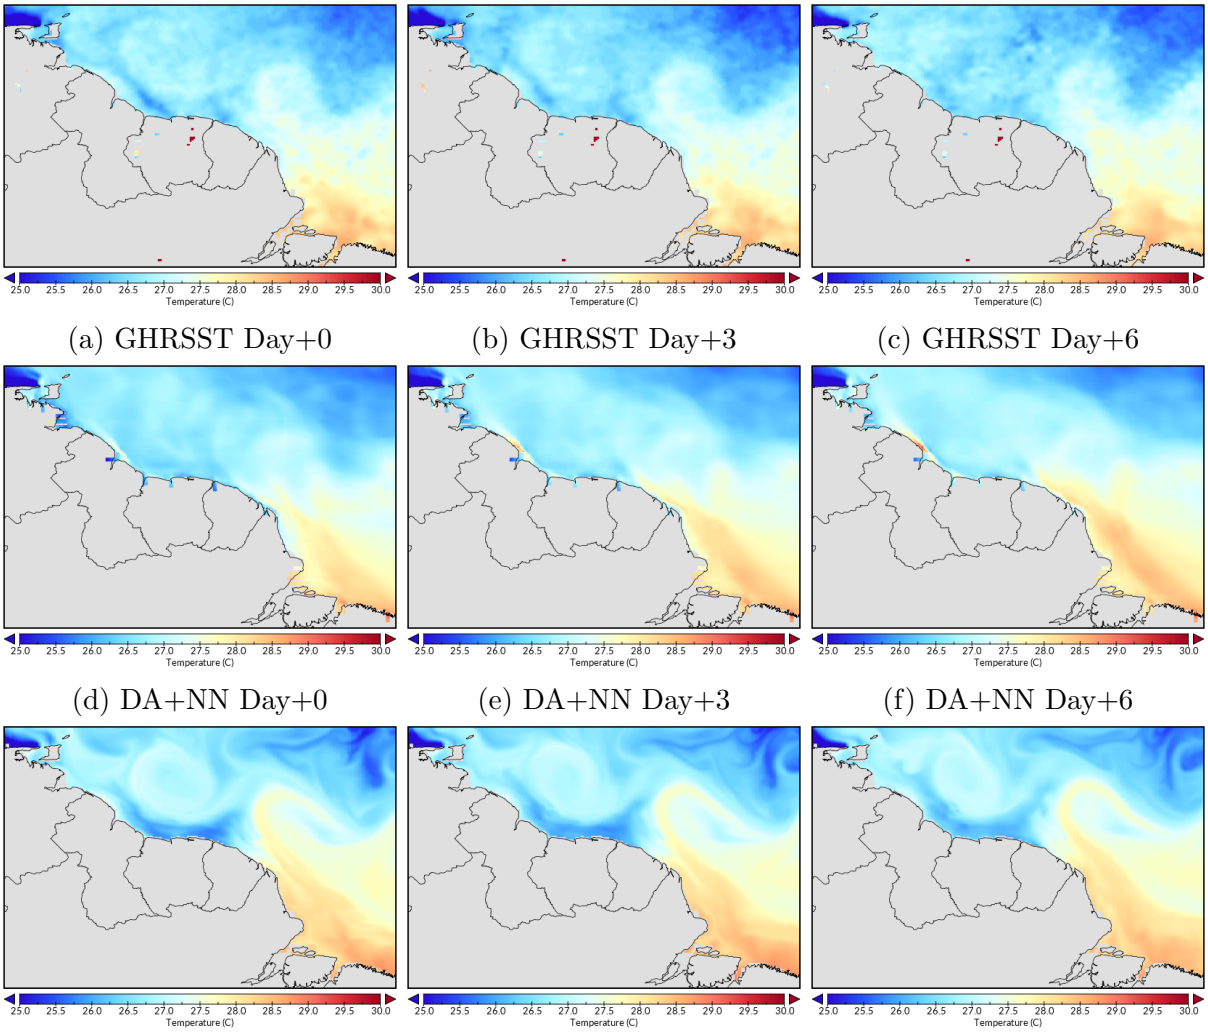

(a) GHRSST Day+0 (b) GHRSST Day+3 (c) GHRSST Day+6

(d) DA+NN Day+0 (e) DA+NN Day+3 (f) DA+NN Day+6

(g) UKMO Day+0 (h) UKMO Day+3 (i) UKMO Day+6

Figure 4.7: Sea surface temperature forecast after the DA (HTc experiment) of January 29$^{\text{th}}$, 2016, Day+0 corresponds to the day of assimilation, while Day+3 and Day+6 correspond to the 3-day and 6-day forecast from January 29$^{\text{th}}$, 2016. Rows, from top to bottom, present the AVISO (0.25-degree resolution) observational data, DA+NN framework forecasts, the forecasts of UKMO. The latter represents January 29$^{\text{th}}$, and February 1$^{\text{st}}$, and 4$^{\text{th}}$, 2016.

to Day+6 in both Summer and Winter, and the spread mean crosses the reference line at Day+6 and Day+4, respectively. At Day+0, Winter shows a mean SS 10% lower than Summer. At Day+6, this is 16% lower. The DA+NN RMSE spreads seem to be close to UKMO and at Day+6 present a negligible difference with it while being $\sim$ 2% higher than UKMO in Winter. At Day+0 during Summer, the mean spatial correlations are 0.81 and 0.70 for DA+NN and UKMO, respectively. The difference is more or less preserved until

Day+6. In Winter, these differences between DA+NN and UKMO are much smaller, with both spreads above r=0.75, from the analysis until the 6-day forecast. In general, the SST analysis, forecasts and persistence of the observations present a bigger spread for SSH. Several of the time series of the spread cross the reference line at Day+2, except for the persistence of the observations during Winter. We did not include the SS, RMSE and r results of the HTc experiment here as the results yielded by HT are very similar. The HTc results present slightly worse SS, RMSE and r across the analysis and the forecasts.

Figures 4.2e, 4.2f, 4.2g and 4.2h show the results of the HTc and ALL experiment for U and V. Firstly, for U there is a clear improvement at Day+0, from the HT to the ALL experiment, of almost 35% and 65% in Summer and Winter, respectively. But at Day+6 there is a small difference between the means of the spread of the HT and ALL experiment. Secondly, for V there is a 58% increase during Summer and 75% during Winter.

In all cases of SS of U and V in the ALL experiment, there is a decreasing trend, where the ensemble performs worse with time. The decreasing trend includes the mean of the spread of U crossing the reference line at Day+4 and overlapping with the ensemble of the HTc experiment. It is also noticeable that in all cases of U and V during the HTc experiment, the SS increases with time. A few ensemble members have a positive SS. In all cases, the RMSE increases with time and seems to converge towards the persistence of the observations and UKMO. The correlation spreads decrease with time in the ALL experiment. The HTc experiment yield higher RMSE at the analysis and therefore the forecasts increase this value with time. The spatial correlations are also worse and seem to stay constant with time. The UV experiment present some improvement over the ALL experiment yielding higher correlations and lower RMSEs.

Two representative days of Summer and Winter were chosen to illustrate the results. These days clearly show the main features of the NBC like the retroflection into the NECC and the NBC rings. Fig. 4.4 and Fig. 4.5 show snapshots of the Summer SSH

Table 4.2: Summary of root mean squared error and correlation coefficient between assimilation and forecasts of SSH and SST (UKMO and DA+NN) and observations, shown in Fig. 4.4, 4.5, 4.6, and 4.7

	UKMO		DA+NN	
Date	**SSH**	**SST**	**SSH**	**SST**
Jul 12th 2015 (Day+0)	0.12 (r=0.58)	0.43 (r=0.71)	0.06 (r=0.85)	0.43 (r=0.71)
Jul 15th 2015 (Day+3)	0.14 (r=0.52)	0.46 (r=0.60)	0.09 (r=0.77)	0.41 (r=0.69)
Jul 18th 2015 (Day+6)	0.14 (r=0.53)	0.44 (r=0.56)	0.10 (r=0.70)	0.41 (r=0.69)
Jan 29th 2016 (Day+0)	0.09 (r=0.77)	0.30 (r=0.94)	0.04 (r=0.90)	0.28 (r=0.91)
Feb 1st 2016 (Day+3)	0.09 (r=0.73)	0.37 (r=0.93)	0.05 (r=0.82)	0.38 (r=0.89)
Feb 4th 2016 (Day+6)	0.11 (r=0.71)	0.37 (r=0.93)	0.07 (r=0.78)	0.40 (r=0.89)

and SST spatial maps of the DA+NN and UKMO forecasts, while the AVISO observations are represented daily and are not forecasts. The depicted days are: July 12th, 2015 (day of assimilation), July 15th, 2015 (3-day forecast), and July 18th, 2015 (6-day forecast).

Fig. 4.6 and Fig. 4.7 show: January 29th, 2016 (day of assimilation), February 1st, 2016 (3-day forecast), and February 4th, 2016 (6-day forecast).The RMSE and correlation coefficients between UKMO and DA+NN for these figures are summaried in Table 4.2. The three datasets present the same characteristic features like the weakening of the NBC retroflection into the NECC, a warm pool of water near the Amazon River outlet and the detachment of an NBC ring advecting to the NW.

The skill score of DA+NN against the persistence of observations (SS_p) is shown in Fig. 4.8 and it is defined as:

$$SS_p = 1 - \frac{RMSE_{Persistence}}{RMSE_{DA+NN}} \tag{4.3}$$

The SS is calculated using the ALL experiment results, where SSH, SST, U and V were assimilated. A positive SS means that the persistence of the observations have a lower RMSE than DA+NN, while a negative SS means that DA+NN have a lower RMSE. Here, we included Day+0 until Day+9, which is the 9-day forecast. We extended this by three days in order to see if there is a cross over time between SS of DA+NN and the persistence. The SS ensemble spread present a decreasing trend for SSH, SST, U and V in all the cases shown. There is a cross over time, by the spread mean, in 5 cases: SSH

July 2015 at Day+6, SST July 2015 at Day+5, V July 2015 at Day+6, SSH January 2016 at Day+6, and V January 2016 at Day+7.

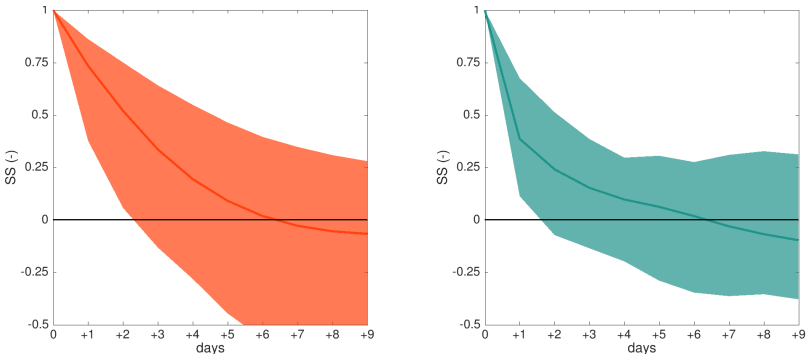

(a) ALL July 2015 SSH

(b) ALL July 2015 SST

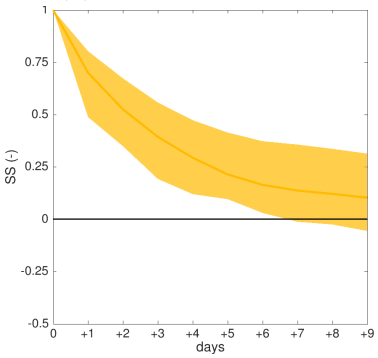

(c) ALL July 2015 U

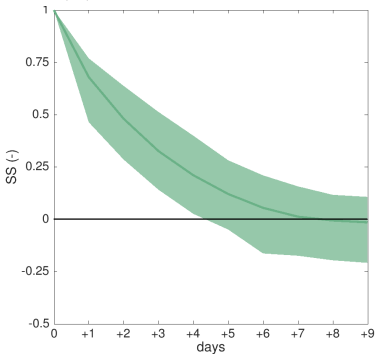

(d) ALL July 2015 V

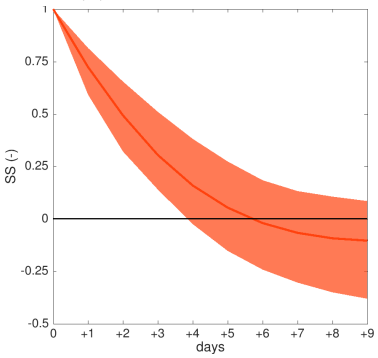

(e) ALL January 2016 SSH

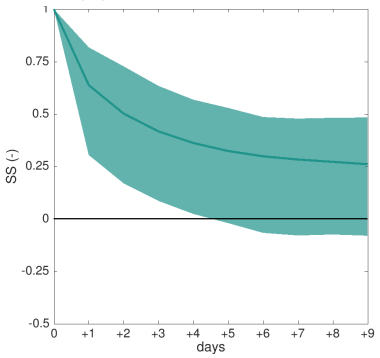

(f) ALL January 2016 SST

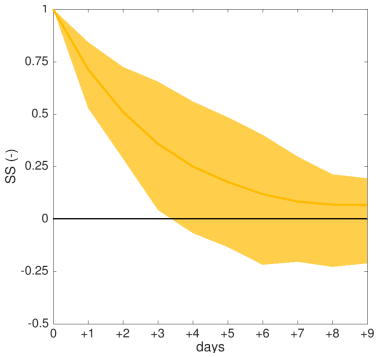

(g) ALL January 2016 U

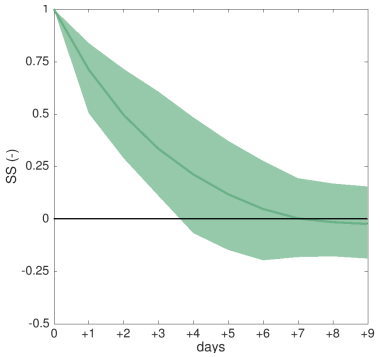

(h) ALL January 2016 V

Figure 4.8: Ensemble spread of the skill score of DA+NN against the persistence of the observations. Left to right: SSH, SST, U and V. Results of the ALL experiment. Top: July 2015, experiment; January 2016. Shaded areas represent the 95% interquantile range. The observations datasets are AVISO mean level anomaly, GHRSST, and Globcurrent. The original spread includes the 31 days of sequential assimilation. The solid lines are the ensemble means and the black line is the zero-reference line. Day+0 is the day of the analysis, while Day+1 to Day+9 are the forecasts.

4.4 Discussion

The DA+NN framework has been applied successfully producing fast DA and forecasts. The addition of a forecast step to the methodology in Chapter 3 is the logical next step as the fast framework allow us to initialise the free-running forecasts from optimal initial conditions. DA+NN also permits the input of new information via DA and does not require that the initial ensemble of the EnKF is produced from a specific day of the reduced-space model solution. An advantage is that an ensemble can be created from any day using the climatology of the reduced-space model solution. The reduced-space model solution provides the climatology of the region and also serves as the input to train the NN for predicting the next time-step.

Including fewer variables in the model provide a better DA result. We would expect that this targeted framework would perform better with fewer variables being included and assimilated. In the HT experiment only SSH and SST are included and assimilated and thus a higher skill score and correlation could be achieved, compared to HTc and ALL experiments. During the HTc experiment, SSH and SST are assimilated but U and V are not, yielding slightly worse forecasts and analysis of SSH and SST than during the HT experiment. When U and V are assimilated too during the ALL experiment, the performance of SSH and SST forecasts decreases. But this also increases the SS, RMSE and correlations of U and V.

The skill score is favourable to DA+NN compared to full physics models such as those produced by UKMO. In the three experiments, the mean of the skill score ensemble spread is positive meaning that the spatial RMSE of the analysis and the subsequent forecasts are lower than the ones produced by UKMO reanalysis and forecasts. The majority of these ensemble spreads (Fig. 4.2 and 4.3) present a decreasing trend in time, with the 6-day forecast presenting the lower skill score. While, the U and V forecasts of the HTc experiment, where the skill score increases with time. The increase with

time can be due to the NN trying to go back to the dynamics it has learned from the principal components of the ROMS model solution. Nonetheless, it is expected that the NN increases the RMSE over time as errors propagate to produced forecasts with no other external constraint input (Bishop 2006). The error created with the 1-day forecast is then propagated into the prediction of the 2-day forecast, and subsequently until producing the 6-day forecast. While the spread crosses the zero-reference line, even starting with the analysis at Day+0, the mean of these spreads only presents a negative skill score after the 4-day forecast of SST during the Winter. This outcome is a favourable result for DA+NN as it is only trained and tested with information from 2001 to 2010. The DA updates the system but the forecasts are based on the NN trained in the aforementioned period. The NN can produce skilful forecasts from new initial conditions that can comfortably be outside the 2001-2010 training period, making it useful. The spatial correlation coefficient spread has higher values with a decreasing trend in time. These findings go substantially beyond those of Patil & Deo (2017), Lins et al. (2013), Garcia-Gorriz & Garcia-Sanchez (2007) since we present spatial information rather than local point-wise time series using NNs. The decreasing trend can also be explained by the recursion of the NN at the time of producing the forecasts. Because there are no external inputs it might mean that the NN has no physical constraints and its behaviour comes from what it has learned from the ROMS model solution.

DA+NN captures the main features of the NBC region. Summer and Winter in the NBC are very distinct. During Summer, warm waters of the strong retroflective current from the NBC lead into the NECC. During Winter, this current weakens and produces the NBC rings and this process is captured by DA+NN during the analysis. An optimal set of initial conditions is necessary for the NN to produce accurate forecasts. In Summer (Fig. 4.4) it reproduces the sea surface height of the region, considering the AVISO dataset as the truth, with a raised sea surface around the retroflection of the NBC. The same is observed in Summer SST (Fig. 4.5), but in some coastal areas it increases the temperature with time. The response might be due to the NN not being able to learn

how to resolve these areas in the reduced-space. During Winter (Fig. 4.6 and 4.7), the NBC ring is perfectly captured by DA+NN, which advects north-westward and starts dissipating in the Lesser Antilles. This subtle feature is learned by the NN and it is able to reproduce these results even with a new initialisation outside the training data. Thus, it is important for the NN not to be over-fitted. The Winter SST has a larger patch of warm water closer to the one depicted by the UKMO forecast, but the eddies of the NBC are present in all three datasets. Problematic areas include the mouth of the Amazon River and coastal areas. The analysis cannot capture the high variability of the observations in this area. This issue was discussed in Chapter 3, where an over-fit to the observations was assessed. The over-fit was obtained by solving a least squares problem. The over-fitted solution did not manage to capture the high variability of the Amazon River mouth. The regression was used as a benchmark for the DA solution and it struggles resolving this area too. The poor resolution might be due to the high uncertainty of the satellite gridded observations of altimetry around coastal areas and specifically in the outlet of the Amazon River. The same problem can be observed for SST, but this data set includes an L4 processing, which includes in-situ measurements and has a higher resolution.

The persistence experiment yields high spatial correlations. However, persistence has obvious limitations for faster dynamical systems. For example, if the analysis starts with a noticeable NBC ring commencing to dissipate into the Lesser Antilles, the persistence of the observations will not predict where these eddies will be. Although, the time scale of these eddies is much longer than 6 to 9 days and this might be why the persistence of the observations has a better skill score against DA+NN. We would expect DA+NN to outperform the persistence since it tries to replicate the dynamics of the region, but as shown before in Phillipson & Toumi (2018), it is hard for even the current generation of ocean models to beat persistence in this timescale. There is a skill crossover time for certain variables, but this happens after 5 days. In general, the forecasts created by DA+NN do not diverge greatly from their respective validation datasets and are

mostly superior to the forecast of the current UKMO system and thus very likely to other operational systems too.

More work needs to be done to improve the forecasting of the surface currents and needs to be tackled in future research. A better forecasting may be achieved by using new architectures of deep NNs like recurrent NN, specifically a Long Short-Term Memory (LSTM) network, which is able to find deeper dependences of time-steps with previous ones within a time-series. This Chapter included the assessment of 6-day forecast against UKMO and 9-day forecasts against persistence. Nonetheless, the timescale of the eddies is approximately 100 days. This issue can be resolved by training the NN with weekly or monthly information in order to make a forecast with longer timescales, as it can be very difficult for daily forecasts to have any real skill after a 100-day forecast.

4.5 Conclusion

DA+NN is a fast reduced-space framework that yields a similar performance to an operational ocean forecast product, UKMO. The DA+NN framework performs fast DA and forecasts when given new information through the EnKF. During the forecast, the NN is able to learn the main features of the region like the NBC rings and their direction of advection but also learns subtle features like the dissipation of these rings in the Lesser Antilles. Generally, DA+NN has a positive skill score, compared to UKMO. However, this skill score decreases with time and it is expected as the errors propagate. The degradation happens with the persistence of the observations.

The framework can be applied to any region or full model with enough information to train the NN. The experiments in this Chapter are a successful attempt to apply this methodology to open ocean where prediction of ocean variables are necessary.

Chapter 5

Forecasting surface ocean currents with external wind data

5.1 Introduction

An accurate short-term forecast of surface currents can improve search and rescue missions, responses to oil spills and marine operations in general (Frolov et al. 2012). Studies towards a better estimation of surface currents include the use of primitive equations models (Shulman & Paduan 2009), tidal harmonics (Egbert et al. 2010, Treguier et al. 2017) and High Frequency (HF) radars (Gurgel et al. 1999, Paduan & Shulman 2004). HF radars measure the surface currents directly up to 200 km off-shore from coastal areas (Barrick 1977). While these are hindcasts, it is important to predict future surface currents in order to be applied to emergency situations. In order to improve the forecasts, previous studies have used the DA of HF radar (Marmain et al. 2014) or satellite surface velocities (Phillipson & Toumi 2017) and empirical models for short-range forecast of surface currents (Röhrs et al. 2012, Frolov et al. 2012). As mentioned in previous chapters, our framework is based on Frolov et al. (2009), and Frolov et al. (2012) is an extension of their previous work. There are several important differences between our framework and

Frolov et al. (2012). While Frolov et al. (2012) uses historical information of HF radar in the Monterey Bay, we used a 10-year ROMS run of a regional ocean model of the NBC and we cover a larger area of study. Initially, Frolov et al. (2009) uses an NN to make a short-range forecast and Frolov et al. (2012) replaced this with a linear autoregressive model that includes predicted wind stress from a regional atmospheric model. Our framework still includes a NN to make the short-term forecast. Finally, our framework includes DA while Frolov et al. (2012) does not.

In Chapter 4 we set up the UV experiment as a benchmark of an optimal data assimilated system of surface currents, where no additional ocean variables were included. In this chapter we added wind data as forcing and as a physical constraint for the surface currents. The methodology used in this Chapter is based on the one described in Chapter 3 and 4. Making use of the skill score (eq. (4.3)) we compare the results of our framework forced with winds against UKMO and also against the experiments without winds.

5.1.1 Wind regime in the Tropical Atlantic

The surface circulation is strongly forced by the transfer of momentum from atmospheric winds across the air-sea interface (Williams & Follows 2011, Josey et al. 2002). The surface winds have a strong seasonal cycle where the contrast in temperature between the pole and the equator is greater in Winter and this leads to a stronger eastward flow in the atmosphere (Josey et al. 2002). Thus, this seasonal strengthening of the winds results in stronger surface stress to the ocean during Winter rather than in Summer (Williams & Follows 2011). In the NBC region there is a characteristic pattern of easterly Trade winds, present at the tropics. The Inter Tropical Convergence Zone (ITCZ) largely determines the climate and meteorology of the NBC region. The seasonal migration of the ITCZ fluctuates between 15°S to 5°N in boreal Winter and boreal Summer, respectively. Thus due to the extent of our region, we expect the ITCZ to have a bigger influence during July rather than during January. The NBC has large annual variations due to the seasonality

of the Trade winds (Newinger 2015).

5.2 Methods and data

5.2.1 Neural network with external wind forcing

Unlike previous chapters, here we introduce a new type of NN: A feed-forward NN, similar to the ones previously used, that accepts external input data along with the information from the PC time series of the full-space system. The external input of data in this case is surface wind from the CFSR reanalysis. The schematic diagram of these NNs is shown in Fig. 5.1. We preprocessed the external wind data before inputting them into the networks. The preprocessing consists of an EOF analysis on the wind data. The wind data correspond to the u- and v-directions of the winds 10m above the sea level surface. The external wind data is used in two ways:

- To train the NN. We use the PC of the winds dataset from 2001 to 2010.
- During the DA using the EnKF. Here, we used the re-analysis winds data as an external forcing during the short-range forecast necessary within the EnKF. Because the DA occurs outside the 2001-2010 period, the external inputs of wind come from data of 2015 and 2016. These inputs were not included during the NNs training.
- During the forecast after the DA. Here, we used data from 2015 and 2016 as well, but it is fed into the system while it creates the free-running forecasts.

5.2.2 Wind data

The wind data was obtained from CFSR reanalysis (Saha et al. 2010). We decided to use CFSR winds as our ROMS model is originally forced with CFSR data. The CFSR

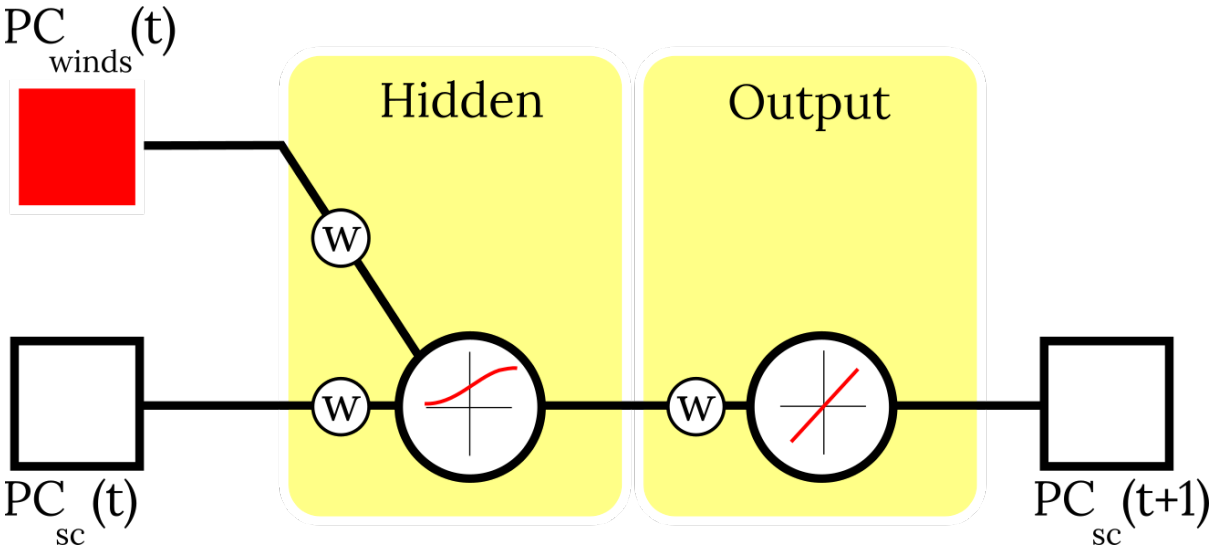

Figure 5.1: Schematic representation of the NN with external input. The external input (red square) are the principal components of the 10 m winds (u and v-direction components), while the input and output (white) are the principal components of the total surface currents at time-step t and $t+1$, respectively. The hidden nodes have a logarithmic sigmoid as activation function and the output layer has a linear regression.

is a third generation, global, and high-resolution reanalysis product. It is a coupled atmosphere-ocean-land surface-sea ice system designed to provide the best estimate of the state of coupled domains. The CFSR global atmosphere resolution is $\sim$ 38 km (T382) with 64 levels. The global ocean is 0.25 at the equator, extending to a global 0.5°beyond the tropics, with 40 levels. The global land surface model has 4 soil levels and the global sea ice model has 3 levels. With these variable parameters, the analysed state will include estimates of changes in the Earth system climate due to changes in aerosols and other trace gases. We used the CFSR magnitude of the u-direction and v-direction surface winds (10 m above sea surface) as an external input for the DA and forecast of surface ocean currents.

In order to prepare the wind data, we interpolated it to the same ROMS grid and added the same landmask that the ROMS model solution uses. The interpolation and land masking of data allow us to eliminate strong dependencies from the wind data over the land and therefore the PCs are only related to wind data over the ocean. By doing this, most of the variance explained by the PCs will not come from the wind over the land,

but rather just from the wind over the ocean and therefore affect the surface currents.

5.2.3 Ekman currents

We used Ekman currents to assess the performance of DA+NN. Since Ekman currents depend strongly on wind, these currents along with the geostrophic currents, give us an indication of the contribution of wind on the total currents of the NBC region. The observed Ekman currents are obtained from the Globcurrent dataset and a further description on how they are obtained can be found in Rio et al. (2014*a*). The Ekman currents uses reanalysis winds from ERA-Interim and a parameterised empirical model at the surface and 15 m depth that depends on latitude, longitude and month (Rio et al. 2014*a*). The dataset corresponds to version 3 of Globcurrent, is 3-hourly and it is valid at the surface with a 0.25-degree resolution.

5.2.4 Data assimilation setup

Again, we chose the EnKF as the DA system. We only assimilate U and V data, while the wind data is used as external forcing. The setup of the DA includes:

- a multiplicative inflation (α) of 3.
- Number of ensembles: 200.
- Measurement error of the observations for U and V: 0.1 m s-1.

The wind data is not perturbed and it remains the same for each ensemble. We expect that this external forcing increases the growth of the spread during the short-range forecast step within the EnKF, on top of the artificial spread growth given by the dynamics learned by the NNs.

5.2.5 Experiments

We performed two DA and forecast experiments: for July 2015 and January 2016. The NN uses winds data as external input for both months as aforementioned. The experiment will be referred to as UV-winds. The UV experiment from Chapter 4 will be referred to as UV-no-winds. We use data from July 2015 and January 2016 because we have available operational forecast datasets from UKMO during those months providing a benchmark for the performance of our framework. The skill score metric used in Chapter 4 is used again as a metric against the performance of the UKMO product.

5.2.6 Validation

The validation of the analysis days and their respective forecast will be against the u and v-directions of the total currents from the Globcurrent dataset. The experiment UV-no-winds from Chapter 4 is used as a benchmark for the performance of the framework using winds for training, DA and forecast.

5.3 Results

5.3.1 Dimension reduction using EOFs

Figure 5.2 shows the first three EOFs and the first three PCs (PC1, PC2, and PC3) time series of U and V from the ROMS model solution. These three modes explain 41.03% of the variance. There is a decoupling between PC1 and PC2 suggesting that there is a high-frequency variability allowing the passage of waves in the mesoscale. EOF1 and EOF2 show that this movement occurs in the NW-SE axis and also suggest that the main features within the region are the NBC rings.

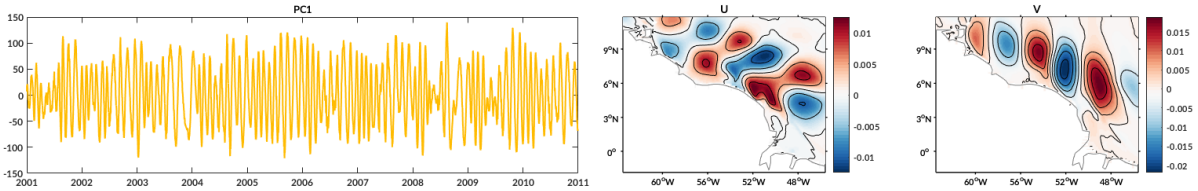

(a) PC1: 16.92%

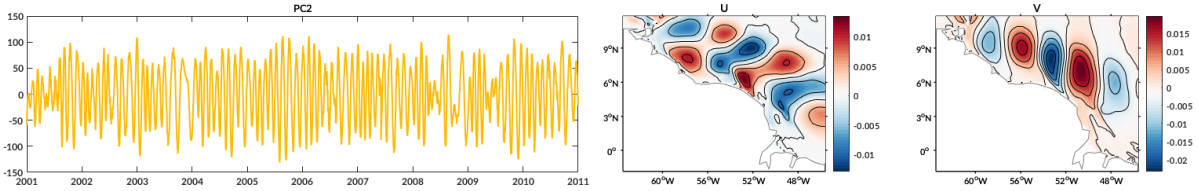

(b) PC2: 16.39%

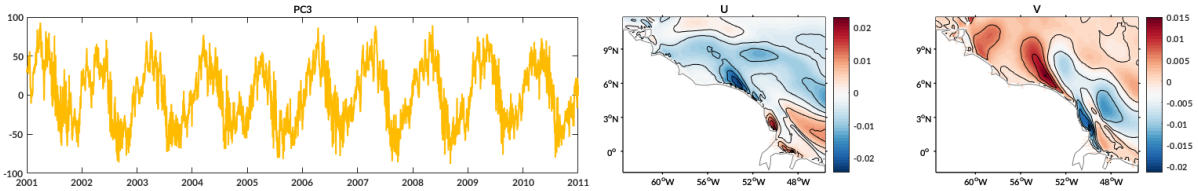

(c) PC3: 7.72%

Figure 5.2: Time series of the Principal Components (PC) and their respective EOFs of U and V of the ROMS model solution from 2001 to 2010. From top to bottom: first, second and third principal components and modes, including the explained variance.

The EOF dimension reduction of the surface winds (Fig. 5.3) of CFSR explains 95.17% in the first 13 modes. Therefore, we used only the first 13 PCs to train and obtain the forecasts from the NN. Wind is dominated by seasonality dominated by the prevailing winds. Changes in the prevailing winds with latitude can be due to the shift of the ITCZ.

5.3.2 Training of the NNs

We would expect the feed-forward NN to have similar hyperparameters as in Chapter 2 during the UV-no-wind experiment, but since the winds data are also used as an input it is necessary to modify these. For the UV-wind experiment, the exhaustive search (10000 combinations) of the hyperparameters of the NN suggests using ADAM (Kingma & Ba 2014) as the training method, a single hidden layer with 356 hidden nodes, a learning

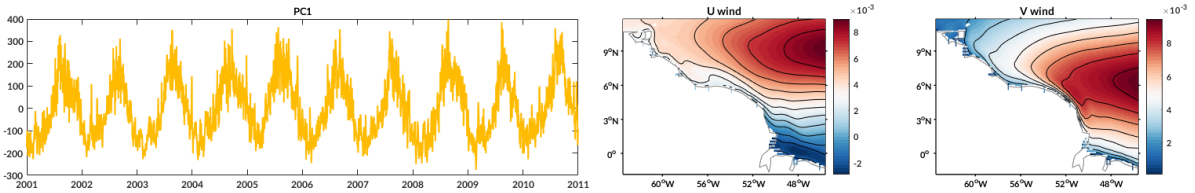

(a) PC1: 58.67%

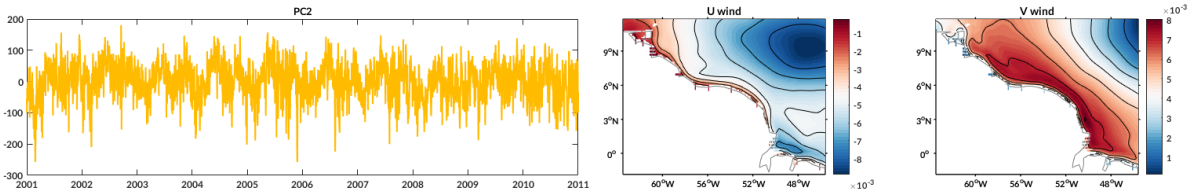

(b) PC2: 12.62%

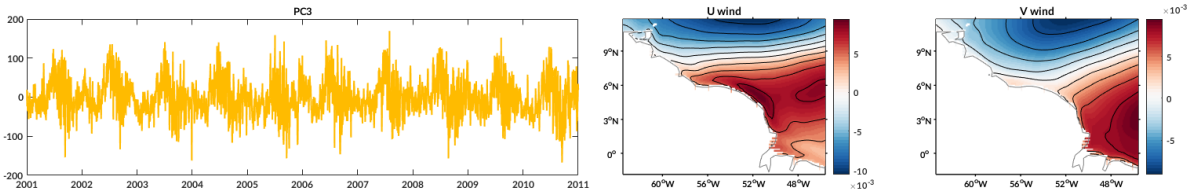

(c) PC3: 7.24%

Figure 5.3: Time series of the Principal Components (PC) and their respective EOFs of u and v-direction of the surface winds from 2001 to 2010. From top to bottom: first, second and third principal components and modes, including the explained variance.

rate of 0.000145, and a logarithmic sigmoid function (logsig) activation function in the hidden nodes.

The speed of training these two NNs was 22 hours for the feed-forward NN. The training uses one node of 12 cores of Intel Xeon (at 2.4 GHz) processors.

5.3.3 Skill scores

Figure 5.4 shows the time series of the ensemble spread of the analysis (Day+0) and 1 to 6-day forecasts (Day+1 to Day+6) of U and V. The shaded area shows the 95% interquantile range that includes the 31 days of assimilation for each month. In general, the spread of the persistence of the Globcurrent observations outperforms the forecasts

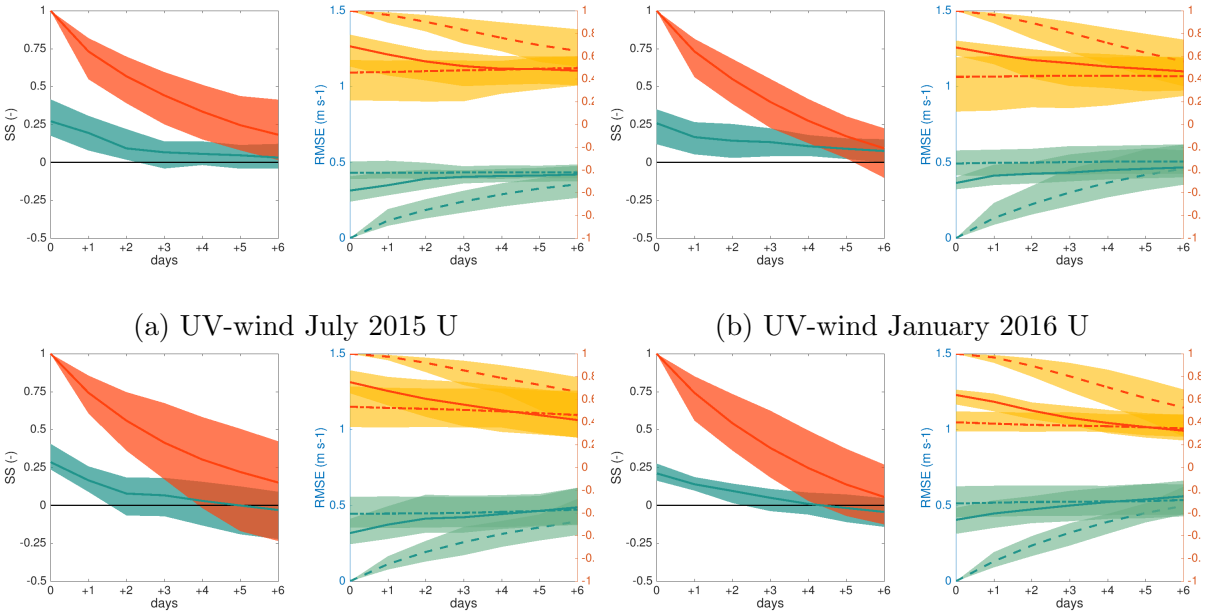

(a) UV-wind July 2015 U

(b) UV-wind January 2016 U

(c) UV-wind July 2015 V

(d) UV-wind January 2016 V

Figure 5.4: Time series of the ensemble spread of forecasts for U and V with external wind forcing. First column: July 2015; second column: January 2016. Shaded areas represent the 95% interquantile range. On each column, **Left-hand side:** Ensemble spread of the skill score of DA+NN forecasts (blue); and persistence of the observations from Day+0 (red). Solid lines are the ensemble means. **Right hand side:** ensemble spread of RMSE (green, left axis) and correlation coefficients (yellow, right axis). Lines represent the ensemble means. DA+NN (solid line); UKMO (dash line); observations persistence (dash point line). Day+0 is the day of the analysis, while Day+1 to +6 are the forecasts. The original spread includes the 31 days of sequential assimilation.

of our framework. Our framework presents a mostly positive skill score against UKMO, where the mean of the spread does not cross the zero-reference line. The earliest that the mean of the spread is outperformed by UKMO is at Day+4 of the forecast of V during January 2016, and Day+5 of the forecast of V during July 2015. The benchmark for these skill scores is the UV-wind experiment in Chapter 4 (Fig. 4.3). Both experiments present very similar performance on a spatial average for the SS, RMSE and correlation coefficient.

Figure 5.5 and 5.6 show the spatial maps of the total currents including the DA day (July 17$^{\text{th}}$, 2015 and January 17$^{\text{th}}$, 2016) (Day+0), 3-day forecast (Day+3), and 6-day forecast (Day+6). The observations are the total currents from Globcurrent. It also presents a

snapshot of the 7-day span average of the near surface winds from CFSR. Globcurrent and CFSR are not forecasts but rather the observations at specific days. The third row includes the results of the experiment without winds from the previous chapter, including analysis and forecasts. The snapshot of July 17$^{\text{th}}$, 2015 and its respective forecasts was chosen because it clearly shows the retroflection of the NBC into the NECC across the four datasets. Day+0 between UV-winds and UV-no-winds present differences in the mesoscale that propagate too into their respective forecasts. We chose January 17$^{\text{th}}$, 2016 as it shows a clear NBC and the weakening of the retroflection of the NBC into the NECC. These features can be observed in all four datasets. The analysis day between UV-winds and UV-no-winds is very similar and there are only a few spatial changes in the mesoscale during their forecasts.

The Ekman and geostrophicity ratio of the surface currents in the NBC region are summarised in Table 5.1. The geostrophicity ratio is calculated using eq. (3.4) and the Ekman ratio is obtained as:

$$Ek_R = \frac{RMS_{Ekman}}{RMS_{total}} \tag{5.1}$$

where RMS_{Ekman} is the root mean square of the Ekman and total currents using eq. (3.3). Here, we present the mean of the time series of Ek_R during July 2015 and January 2016 using Globcurrent. To show how these ratios are preserved by UV-wind and Uv-no-wind, we also include the ratios of the Ekman currents against the total currents obtained by the analysis. The spatial representation of the ratio between Ekman and total surface currents from the Globcurrent dataset is represented is shown in Fig. 5.7. As shown, this is not a domain dominated by Ekman transport, as the total currents depend greatly on the geostrophic currents.

The spatial representation of the average skill score of UV-no-winds and UV-winds is

Table 5.1: Mean of the Ekman ratio using the total surface currents of Globcurrent, and the experiments with and without winds. For the latter, this is the mean of the day of the analysis.

	July 2015	January 2016
Globcurrent	0.29 ± 0.05	0.22 ± 0.07
UV-winds	0.37 ± 0.07	0.30 ± 0.09
UV-no-winds	0.39 ± 0.08	0.31 ± 0.09

depicted in Fig. 5.8 and is calculated as:

$$SS_{Day+i} = \frac{1}{m}\sum^{m}(1 - \frac{RMSE_{winds}}{RMSE_{nowinds}})_{Day+i},\ i = 0, \cdots, 6 \tag{5.2}$$

where SS_{Day+i} is the average skill score per grid point at Day+i, and m are the 31 days of July 2015 and January 2016. The average is the mean of each of the sets of 31 days of Day+0, Day+1, Day+3, and Day+6. Positive values means that the experiment with winds has a higher skill score, while negative values mean that the experiment without winds has a higher skill score. The figure shows the average skill score of the magnitude of the currents. At Day+0, there is a similarly distributed average spatial skill score between UV-winds and UV-no-winds. In July 2015, there is a higher skill score by UV-no-winds in the Northeast that turns into a higher skill score by UV-winds promptly after only one-day forecast and it is maintained until Day+6. In January 2016, there is no distinguishable skill score between the two experiments and this is only visible from Day+1 onwards. The area near the Lesser Antilles is mainly dominated by a higher skill score by UV-no-winds. The latter is also more skillful around the area comprising the NBC ring but this changes in Day+3 and Day+6 where UV-winds turns out to have more skill. In all cases, none of the experiments are noticeable more skillful around the Amazon River mouth.

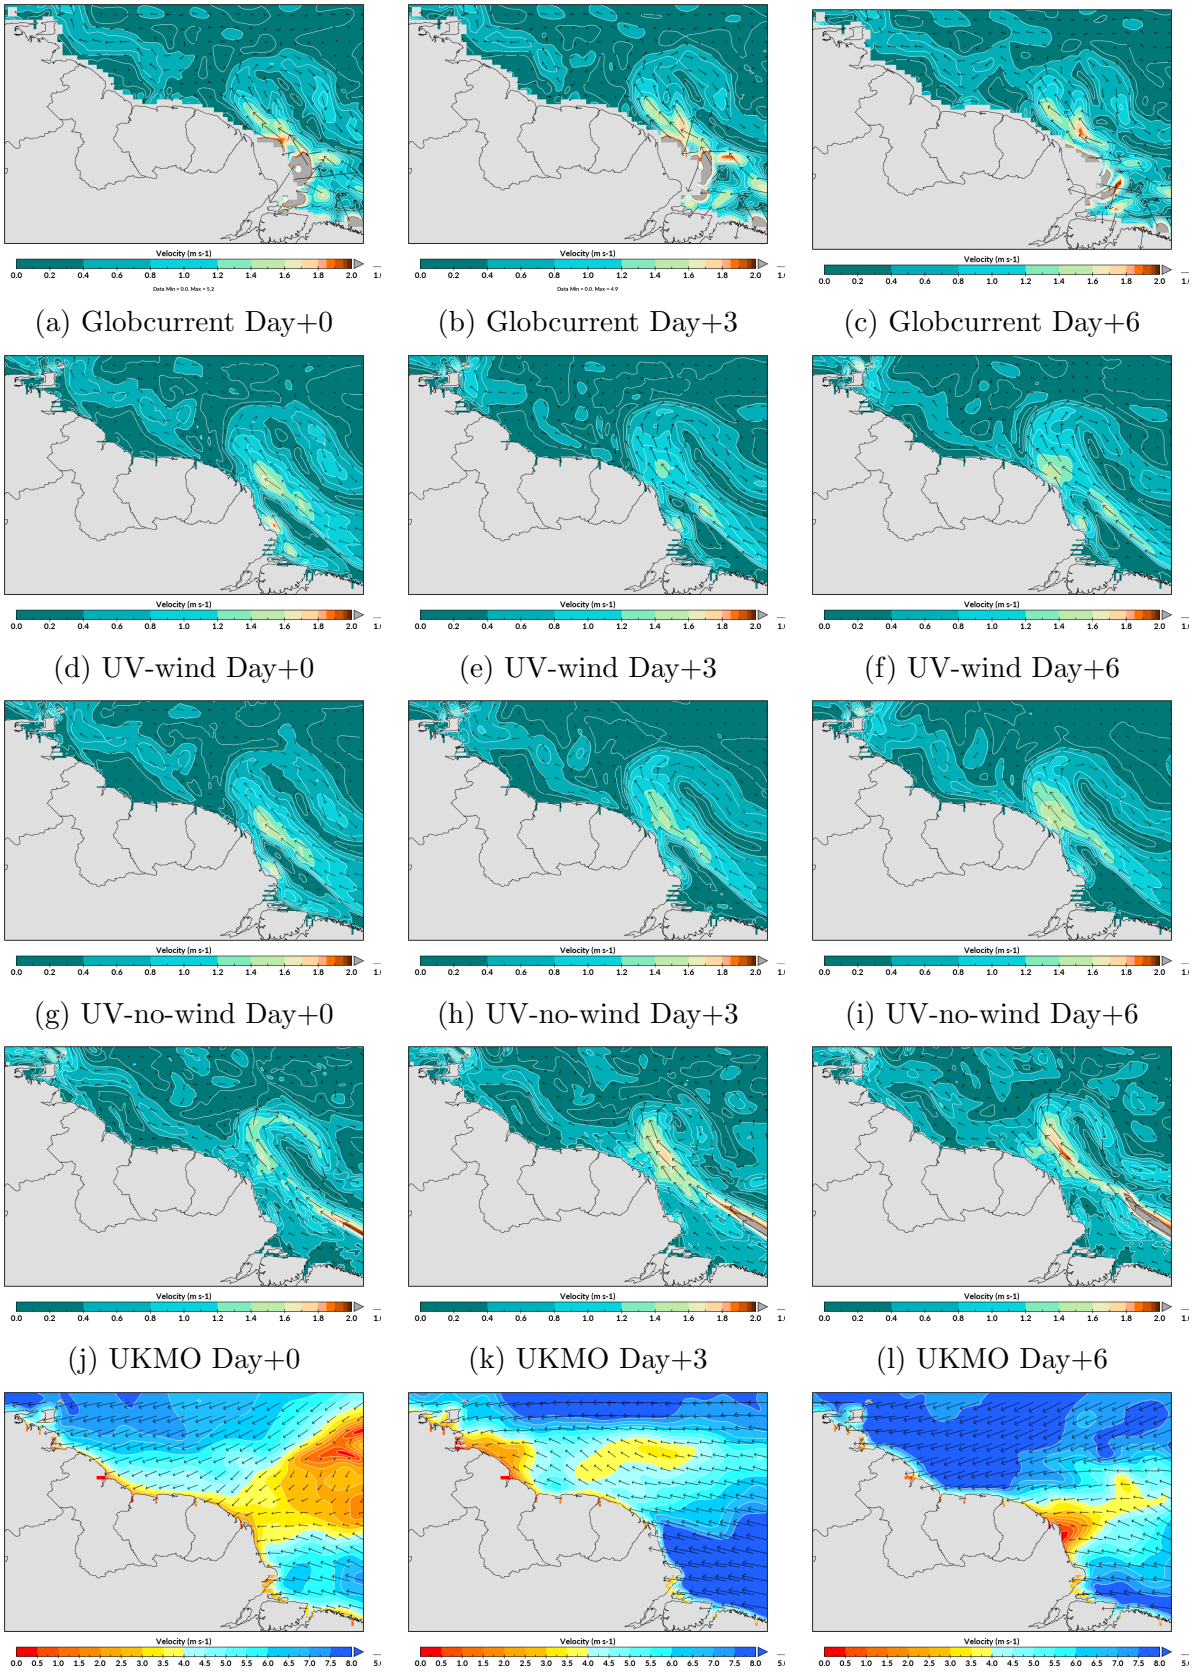

(a) Globcurrent Day+0 (b) Globcurrent Day+3 (c) Globcurrent Day+6

(d) UV-wind Day+0 (e) UV-wind Day+3 (f) UV-wind Day+6

(g) UV-no-wind Day+0 (h) UV-no-wind Day+3 (i) UV-no-wind Day+6

(j) UKMO Day+0 (k) UKMO Day+3 (l) UKMO Day+6

(m) Wind at 10 m Day+0 (n) Wind at 10 m Day+3 (o) Wind at 10 m Day+6

Figure 5.5: Total surface currents forecast after the DA of July 17$^{\text{th}}$, 2015. Day+0 is the day of assimilation, while Day+3 and Day+6 correspond to the 3-day and 6-day forecast from July 17$^{\text{th}}$, 2015. Rows, from top to bottom, present the Globcurrent (0.25-degree resolution) observational data, the DA+NN framework forecasts of the UV-wind experiment, the DA+NN framework forecasts of the UV-no-wind experiment, forecasts of UKMO and winds. The latter represents July 17$^{\text{th}}$, 20$^{\text{th}}$ and 23$^{\text{th}}$, 2015

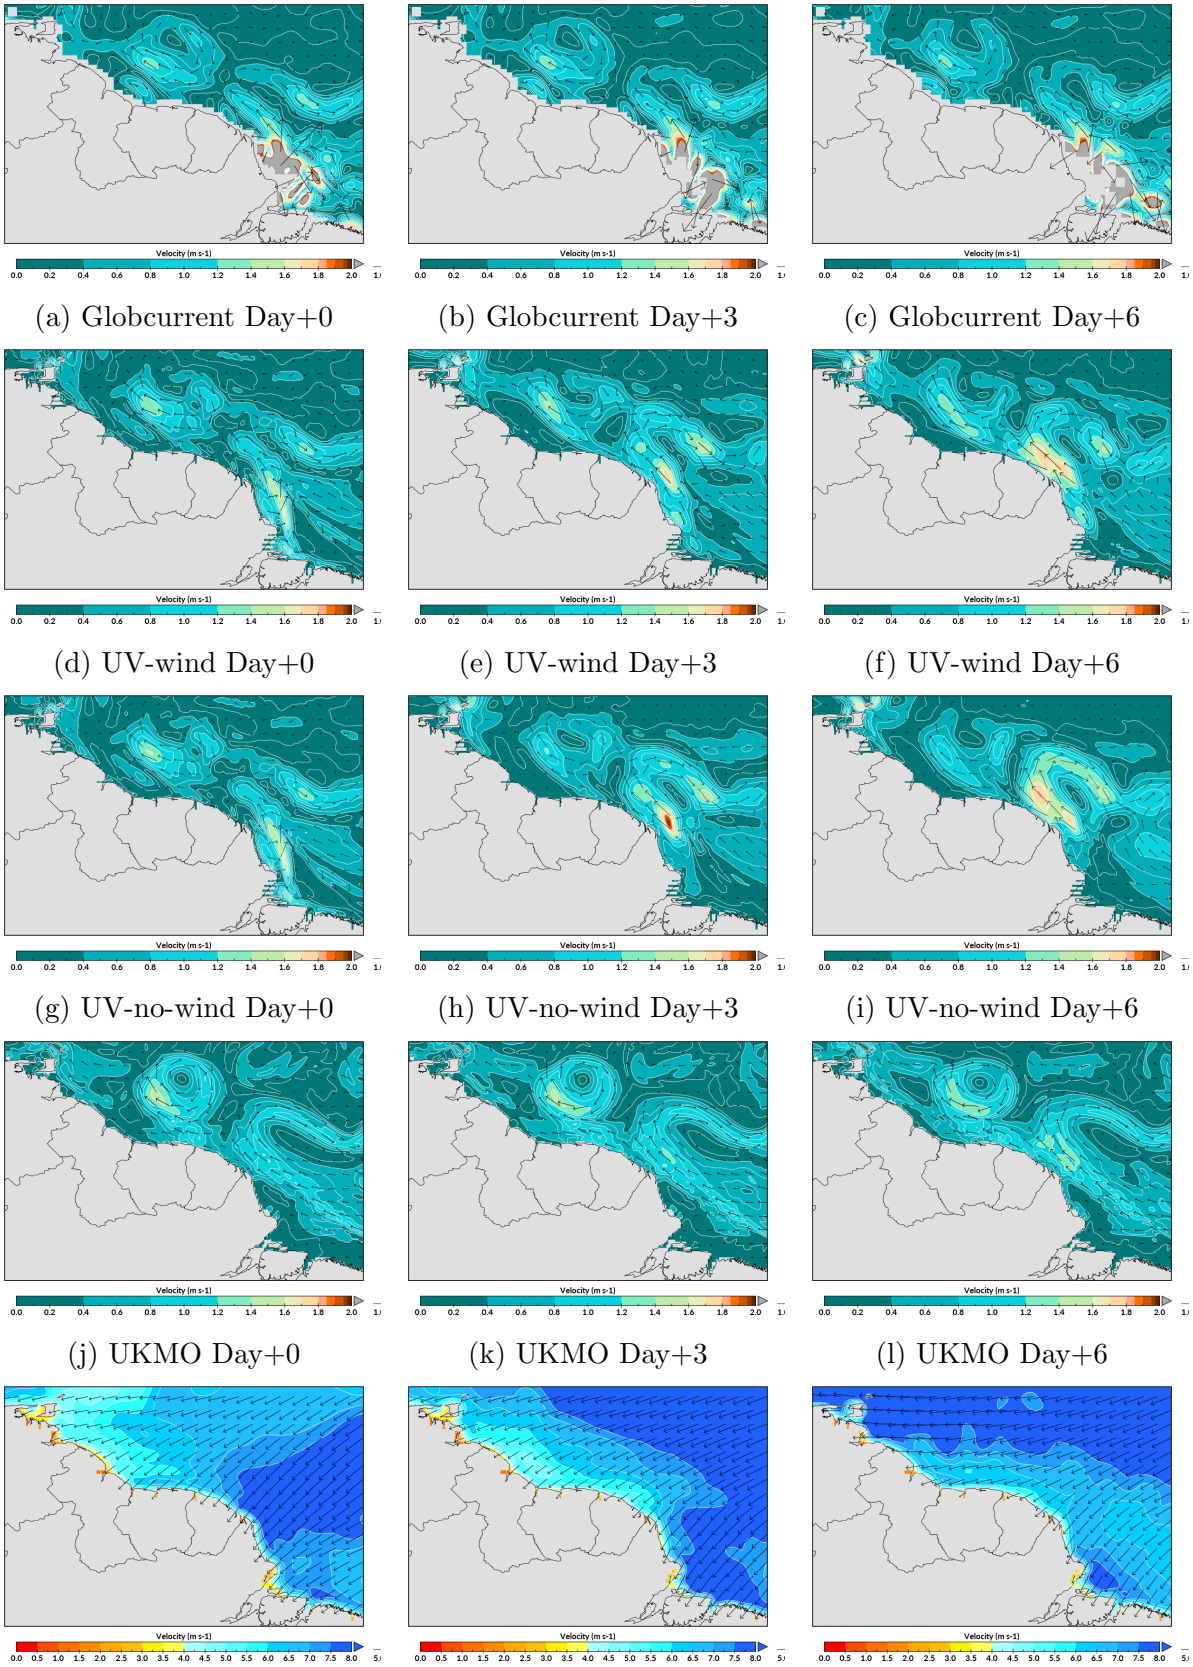

(a) Globcurrent Day+0 (b) Globcurrent Day+3 (c) Globcurrent Day+6

(d) UV-wind Day+0 (e) UV-wind Day+3 (f) UV-wind Day+6

(g) UV-no-wind Day+0 (h) UV-no-wind Day+3 (i) UV-no-wind Day+6

(j) UKMO Day+0 (k) UKMO Day+3 (l) UKMO Day+6

(m) Wind at 10 m Day+0 (n) Wind at 10 m Day+3 (o) Wind at 10 m Day+6

Figure 5.6: Total surface currents forecast after the DA of January 17th, 2016, Day+0 is the day of assimilation, while Day+3 and Day+6 correspond to the 3-day and 6-day forecast from January 17th, 2016. Rows, from top to bottom, present the Globcurrent (0.25-degree resolution) observational data, the DA+NN framework forecasts of the UV-wind experiment, the DA+NN framework forecasts of the UV-no-wind experiment, forecasts of UKMO and winds. The latter represents January 17th, 20th and 23th, 2016

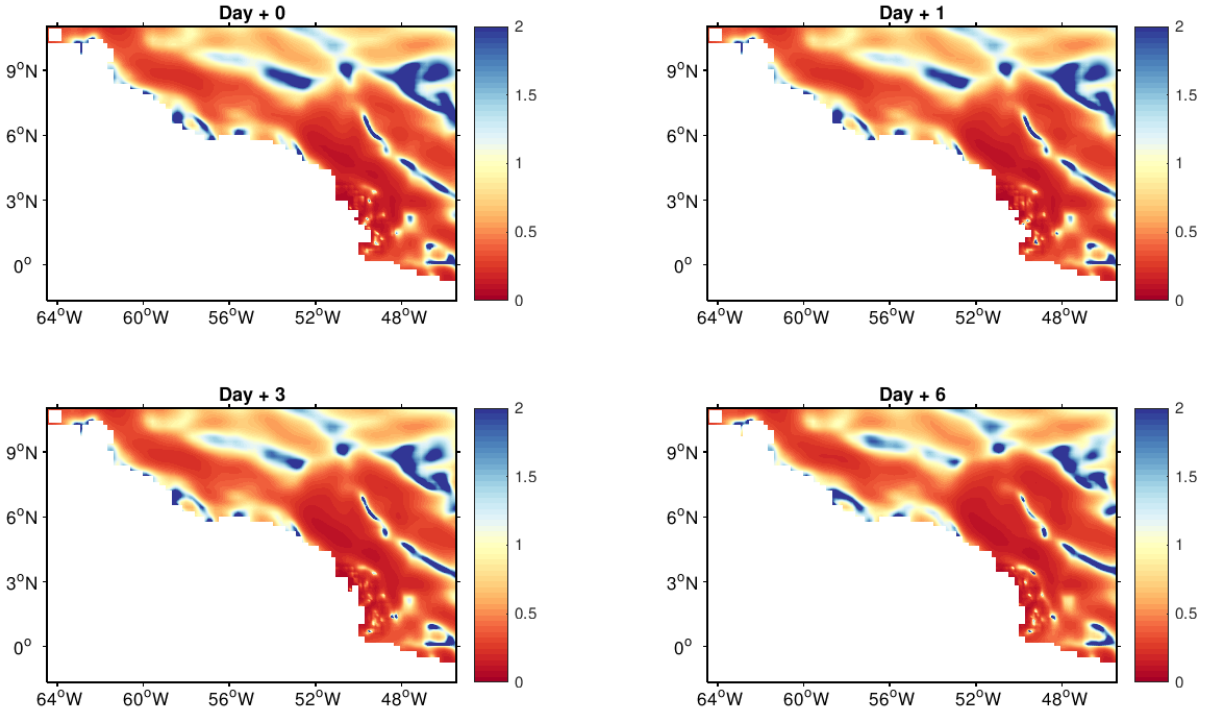

(a) Ekman to total currents ratio (Ek_R) July 2015

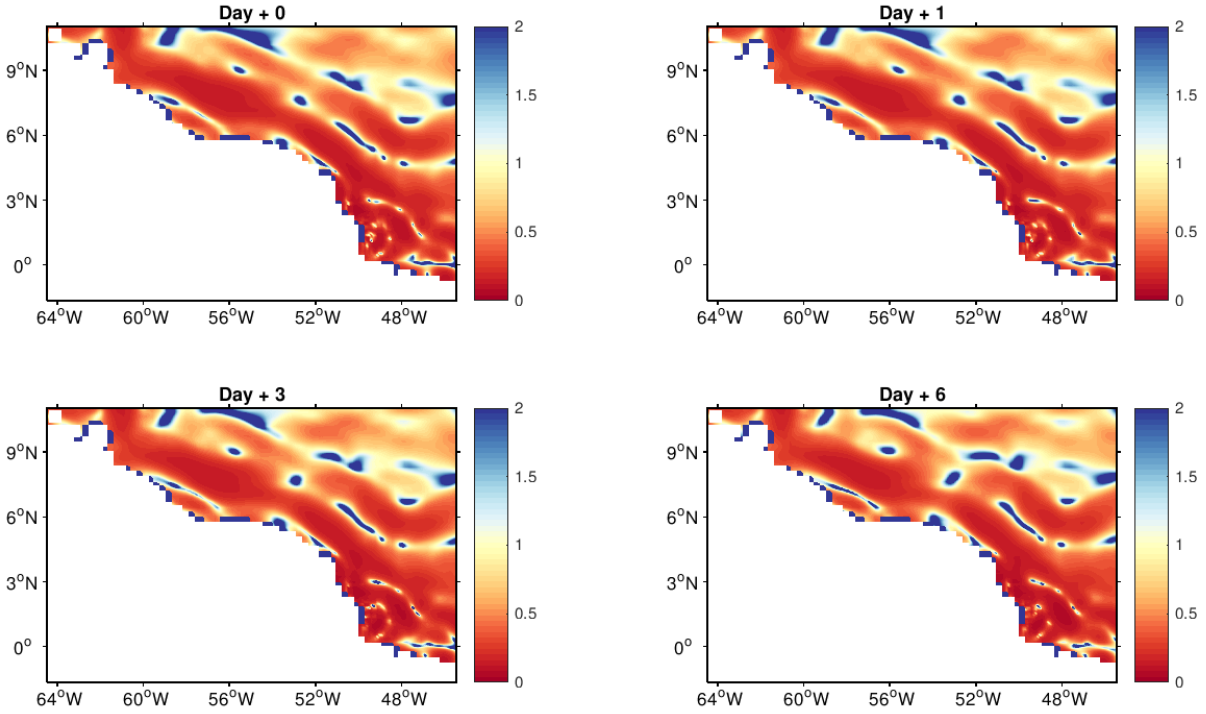

(b) Ekman to total currents ratio (Ek_R) January 2016

Figure 5.7: Spatial Ekman to total currents ratio (Ek_R) over the NBC. Including Day+0, +1, +3, and +6. The ratio between the Globcurrent datasets. Ratio > 1 means that the Ekman currents have a bigger contribution, while ≥ 0 and < 1 means that the major contribution come from the total currents dominated by geostrophic currents.

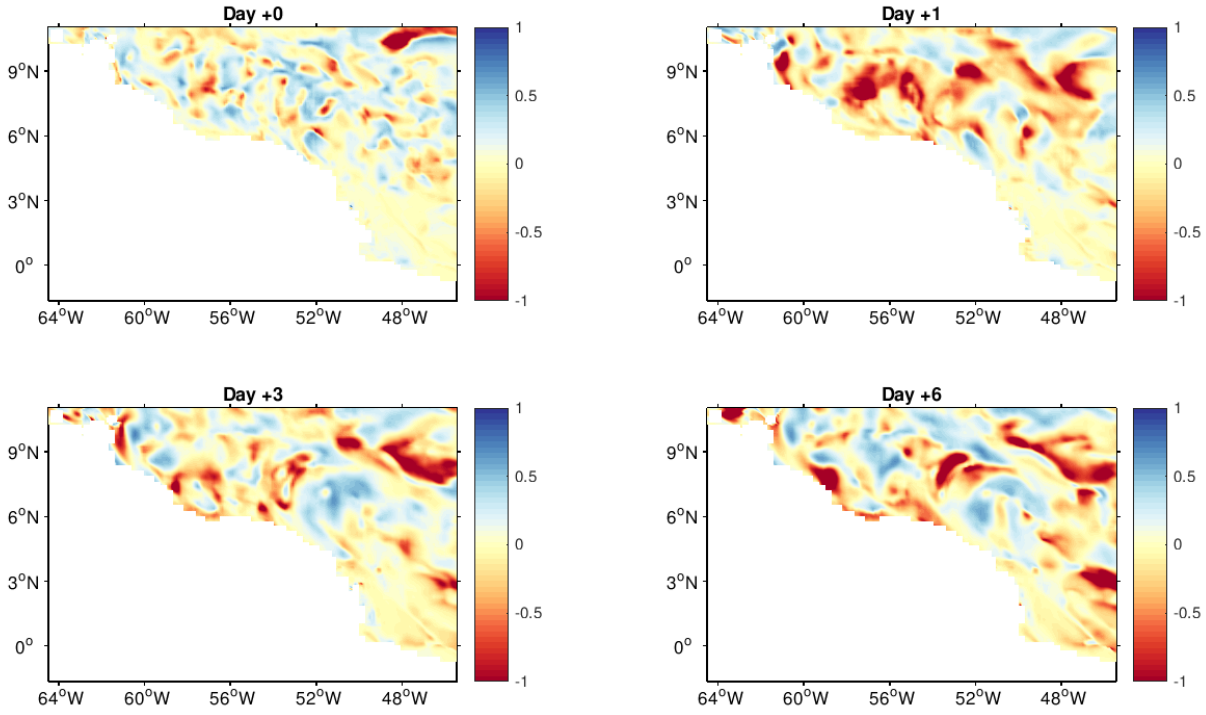

(a) Average spatial skill score during July 2015

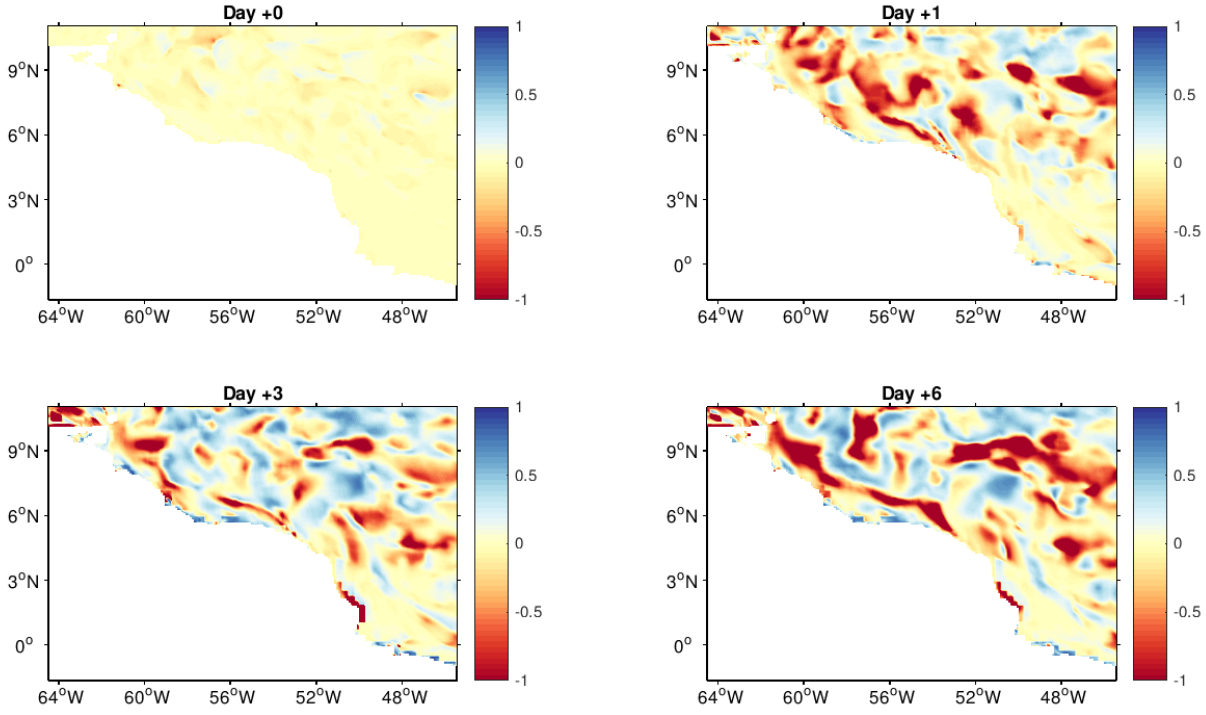

(b) Average spatial skill score during January 2016

Figure 5.8: 31-day average spatial skill score. Average of all the analyses, and 1, 3, and 6-day forecasts. Positive values indicate that the UV-wind experiment has a higher skill score than the UV-no-wind experiment and vice versa.

5.4 Discussion

The forecasts that we obtained from UV-winds were compared to the UKMO forecasts and while UV-winds does not perform better than UV-no-winds, on average, these do perform better than those of UKMO. This performance is mostly valid until a 4-day forecast, where the mean of the spread has less skill than UKMO.

The prevailing currents are the geostrophic currents. The total surface currents from Globcurrent are the sum of Ekman and geostrophic currents, but it is clear that the dominant currents in the NBC region are the geostrophic currents as shown in the snapshot in Fig. 5.7. The NBC is located in the equator and while the ITCZ converges within the study area there is also a strong geostrophic component in the currents. The first two EOF fields of U and V show that there is a steady advection of the surface currents throughout the 10-year ROMS model solution, and this can be mostly due to the strong geostrophic currents in the region and the NBC rings. The third EOF field can be related to the seasonal patterns of the surface currents.

There is a successful application of the NN with external input. The NN is harder to train since the external input makes it more difficult to find the minimum of the cost function. Therefore, it was necessary to preprocess the surface winds data since this dataset contained also information over the landmass. The continental winds were clearly slower than those over the open ocean. By masking this data, we introduce less variation and we also get rid of data that either would not improve the forecast and does not act directly over the surface currents. Because of the external input, it was necessary to apply a logarithmic sigmoid activation function. The logarithmic sigmoid training function allowed us to accelerate the training of the NN and also permits the training to not be trapped in high local minima due to the high values of the Principal Components of the surface winds. Nonetheless, it was not necessary to retain more than 13 principal components of the surface winds as this explains more than 95% of the variance. We

would expect the wind forcing within the short-range forecast in the EnKF to increase the grow of the spread, so the background covariance does not degrade and decouples from the truth. The forecast of the experiment with winds are obtained using the forecast of the previous day and we would also expect these to increases the error from prediction to prediction. However, by adding the wind forcing this gives a physical constraint that does not greatly degrade the performance when compared to the experiment without winds and it does improve it in some areas. While the NN with external forcing can be harder to train, it is not hard to implement and it is encouraging to see that generally does not degrade the performance of our framework.

The spatial skill score gives us an indication as to where the wind forcing is adding skill to the forecast. During July 2015, the main feature of the region is the retroflection of the NBC in the western part of the region. It seems that on average the experiment with wind forcing help the surface currents in the loop of the retroflection to be better predicted and closer to the Globcurrent observations. Even though the retroflection is geostrophic, this might be a non-linear response of the NN and the wind forcing. Along the coastline, the experiment without wind have either the same or a higher skill score than the experiment with winds, and could be due to the lower wind regime on the coast during July. During January, there is a spatial skill score close to 0 during the analysis day. Since the short-range forecast within the EnKF uses the wind forcing, we would expect that the influence of the winds would be reflected in the day of the DA. However, in this case it seems that the wind forcing only helps the spread of the background ensemble to grow. Nonetheless, the growth of the spread using the wind forcing does not improve the DA greatly since the experiment without winds seems to yield very similar results. Like aforementioned, it does make localised improvements during July 2015. In average, using wind forcing is as good as not utilising them to force the system and does not decrease the performance of the DA or forecasts greatly.

The findings of this study can be compared to the ones found by Frolov et al. (2012).

Nonetheless, their framework is based on HF radars and a relatively small study area in the Monterey Bay (up to 200 km offshore). Thus, our study goes beyond as it is applied in the open ocean of a larger region as the NBC. They have also found that when the wind forecast is not available, their model degrades slightly when compared to drifter trajectories. As seen in this study, the wind forcing only makes localised improvements. Unfortunately, we did not find any drifters passing through this region during July 2015 or January 2016 and thus a comparison of their trajectories to our predicted trajectories would not be possible. Lack of drifter data can be a limitation and it can be solved by assessing a different period of time where several drifters pass through this region. The selection of July 2015 and January 2016 is to show how our DA and forecasts compared to UKMO. We consider our wind forcing to be perfect winds from the reanalysis rather than the actual forecasted surface winds from CFSR.

While there is a localised improvement over the Ekman currents it would clearer if the framework is applied on a high-wind mid-latitude regional ocean model or against a drifter-rich observation period.

5.4.1 Wind stress curl

While Richardson et al. (1994) suggests that weak eddies might be generated in the ocean interior by instability of the near equatorial currents driven by wind, as future work we should consider the application of the wind stress curl rather than the surface winds. Some previous studies (Philander & Pacanowski 1986) have mentioned that the NBC is influenced by the wind stress curl across the interior of the basin. And while it has been reported that at midlatitudes, the response of western boundary currents to annual changes in the wind stress curl is small (Lee et al. 1996), this is not true for low latitudes. Johns et al. (1998) and (Philander & Pacanowski 1980) explain that there is a lagged response of the currents to the interior wind stress curl and this influences the seasonality of the NBC. Remote wind forcing is overpowered by the thermohaline flow by the NBC

(Johns et al. 1998). Considering additional EOFs does not guarantee that the wind stress curl is included, and is not certain how many EOFs would have to be considered before the wind stress curl is well represented.

The fact that the wind stress curl is more important than the surface winds or the wind stress means that ideally this framework should include the wind stress curl as an input to train the NN, perform the DA, and produce the forecasts. Despite the influence of the wind stress curl in the seasonality of the NBC, Johns et al. (1998) also reports a rapid impact and no appreciable lag between the interior wind stress curl and the NBC currents during early summer. Thus, we could expect a better response of the wind stress curl forcing during the July 2015 experiments. During the January 2016 experiment, the forecasts might not benefit from the wind stress curl forcing if the lagged response is in a longer timescale than a reasonable 10-day forecast.

5.5 Conclusion

We have studied how an external forcing of winds in our framework can improve the forecast of surface currents. The results of this new experiment were compared to a previous experiment without winds, using the latter as a benchmark. In general, we only found small localised improvements in the skill score against the benchmark experiment, but these are more clear during the forecasts rather than during the DA. The small improvement of using wind forcing can be due to the application of this framework over an equatorial region like the NBC that is mostly dominated by geostrophic currents and not Ekman currents, which are driven by wind. Because of this targeted framework, it is not necessary to include any other variables rather than the surface currents. Further work can include the use of this framework over high-wind mid-latitude ocean regions where Ekman currents are more dominant and therefore we should not discard it completely.

Chapter 6

Conclusion

6.1 Summary of key findings

A framework of fast data assimilation (based on Frolov et al. (2009)) and fast forecasting of ocean variables have been applied on a reduced-space regional ocean model of the North Brazil current. This model is based on a 10-year (2001-2010) Regional Ocean Modelling system simulation and includes different combinations of SSH, SST, and U and V. The reduction of dimensions of the ROMS model solution was obtained via an EOF analysis. The analysis divides the model into time series of PCs and EOF fields. The time series of the PCs are used to train an artificial neural network, NN, to emulate the physics of the original full-space ROMS model solution. The NN is used inside the EnKF that we use as a data assimilation scheme. Inside the EnKF there is a short-term forecast that in this study is obtained using a NN rather than using a supercomputer to perform an additional run in ROMS. In general, the framework runs on a single core processor and is three to four orders of magnitude faster than other full-space data assimilation schemes. The framework was successfully applied for the first time in a large area of the open ocean.

In Chapter 3, we studied the data assimilation performance of this framework. The first application of the fast data assimilation used the first 9 years of the ROMS model solution, and uses the final year, 2010, as a validation year. While our model includes SSH, SST, U and V only SSH and SST were assimilated. As a benchmark we performed two extra experiments: an overfit to observations using a least squares solution, and another where the short-term forecast by the NN was replaced by persistence. The LS solution represents an overfit to the assimilated data but does not translate well into the variables not assimilated like U and V. The persistence experiment does not increase the spread of the ensembles and yields a poor performance data assimilation using the EnKF. Our framework outperforms both of these experiments giving value to the use of the NN as a forecasting system, and maintains the geostrophicity ratio between geostrophic and total currents while the least squares solution does not. Another benchmark is the original ROMS model solution for the year 2010. Here, our framework beats the ROMS average spatial RMSE of SSH, SST, U and V against validation datasets by 41, 49, 13, and 15%, respectively. These assimilation errors are similar to state-of-the-art physical model analysis at orders of magnitude lower in the computational cost.

Forecasting of the ocean variables was added in Chapter 4. Additionally to performing data assimilation we forecast the ocean variables included in our system. Here, we performed four different experiments using different combinations of ocean variables that are included in the system and variables that are assimilated, but always being limited to SSH, SST, U and V. For this, we trained specific neural networks for each case. After the data assimilation, we performed a 1 to 6-day forecasts and compared those to an operational product, UKMO. UKMO provided us with archived forecasts from July 2015 and January 2016 and therefore our analysis is limited to these months. We used 10 years of the ROMS model solution to train the specific NNs. Our framework generally outperforms the skill score and correlation coefficient of the UKMO forecasts. When compared to the persistence of the observations, only our forecasted SSH beats persistence after a 6-day forecast in both July and January. Using different combinations of ocean variables

shows that using targeted variables improve the data assimilation and forecasts, rather than using any additional variables which are surplus.

A further study of the regional ocean surface currents was performed in Chapter 5. Using the experiment of assimilating surface currents from the previous chapter as a benchmark, we added an additional physical forcing from 10 m winds in order to improve the surface currents forecasts. For this we introduce a modified NN that includes external input of surface winds data. Comparing average spatial RMSE, correlation coefficient and skill score there is no clear improvement if we used an external wind forcing compared to the experiment without winds. The NBC region is highly dominated by geostrophic currents, and the influence of winds over the Ekman currents do not show a noticeable improvement during the forecasts of the total surface currents. The geostrophic currents can be obtained directly from SSH and since these are a large component of the total currents in this region, a way to improve the forecast of surface currents here is just assimilate and forecast SSH. SSH proved to yield very high skill score compared to UKMO and it is the only variable that outperforms the persistence of the observations within 9 days of forecasts as seen in Chapter 4.

This fast framework takes between $\sim$ 20 to 60 seconds per assimilation day. The speed depends on the number of assimilated variables and the number of assimilated observations. Once the data assimilation is performed, the forecasts can be obtained in a matter of seconds and this will depend on how long we want our forecasts to be. It takes $\sim$ 300 seconds to perform the dimension reduction of the ROMS model solution and up to $\sim$ 50 hours to train and find the optimal hyperparameters of the different neural networks used throughout this study. Since the dimension reduction and the NN training needs to be performed only once, this framework greatly depends on the data assimilation step. Because of this, the framework can be used in emergency situations like oil spills and search and rescue missions where a fast prediction of ocean variables is needed. If there is further model or training data a brand new reduced model can also be rapidly created.

6.2 Future work

There are three clear improvements than can be applied to this framework. While we chose the NBC region because it is an eddy-rich area, this is a relatively low-wind equatorial region and works as a proof of concept. Nonetheless, it would be necessary to test the framework in another regional ocean, maybe a high-wind mid latitude system where the surface currents are not strongly dominated by geostrophic currents and in that way the wind forcing can help to improve the surface currents forecast. Also, this framework is not limited to the ocean variables used throughout this study or to a specific ocean model. Here we used a run of ROMS but other operational models like HYCOM or UKMO can be used to obtain the reduced model and train the NN.

The NN can be improved with more recent architectures. Recently, there has been a significant increase of the use of neural networks and deep learning in classification and, more importantly to us, in time series forecasting. The use of Long Short-Term Memory networks can improve the extraction of more subtle features from the Principal Components and because of their memory it can learn from several previous time steps and also forget information that it is not necessary.

This framework is targeted at the surface layer, and further work would require to extend its capabilities to deeper ocean levels. Some work on the reconstruction of the state of deeper ocean levels based only on the state of the surface ocean variables has been done by Wang et al. (2013). The approach could be useful as it would not be necessary to extract information of other ocean levels and include them in the data assimilation making the process slower. Rather, the DA can be still be performed at the surface and quickly extrapolate the deeper ocean states from the new surface state.

Finally, development towards a real-time application of this novel framework in an ocean emergency situation would be highly valuable.

Bibliography

Anderson, D., Sheinbaum, J. & Haines, K. (1996), 'Data assimilation in ocean models', *Reports on Progress in Physics* **59**(10), 1209.

Ångström, A. (1935), 'Teleconnections of climatic changes in present time', *Geografiska Annaler* pp. 242–258.

Bakun, A. & Broad, K. (2003), 'Environmental loopholes and fish population dynamics: comparative pattern recognition with focus on El Nino effects in the Pacific', *Fisheries Oceanography* **12**(4-5), 458–473.

Bannister, R. N. (2008), 'A review of forecast error covariance statistics in atmospheric variational data assimilation. II: Modelling the forecast error covariance statistics', *Quarterly Journal of the Royal Meteorological Society* **134**(637), 1971–1996.

Barrick, D. E. (1977), 'Extraction of wave parameters from measured HF radar sea-echo Doppler spectra', *Radio Science* **12**(3), 415–424.

Bauer, P., Thorpe, A. & Brunet, G. (2015), 'The quiet revolution of numerical weather prediction', *Nature* **525**(7567), 47.

Bengio, Y., Goodfellow, I. J. & Courville, A. (2015), 'Deep learning', *Nature* **521**, 436–444.

Bengio, Y., Simard, P. & Frasconi, P. (1994), 'Learning long-term dependencies with gradient descent is difficult', *IEEE transactions on neural networks* **5**(2), 157–166.

Bergstra, J. & Bengio, Y. (2012), 'Random search for hyper-parameter optimization', *Journal of Machine Learning Research* **13**(Feb), 281–305.

Bergstra, J. S., Bardenet, R., Bengio, Y. & Kégl, B. (2011), Algorithms for hyper-parameter optimization, *in* 'Advances in neural information processing systems', pp. 2546–2554.

Bishop, C. M. (2006), *Pattern recognition and machine learning*, springer.

Bjerknes, J. (1969), 'Atmospheric teleconnections from the equatorial Pacific 1', *Monthly Weather Review* **97**(3), 163–172.

Boznar, M., Lesjak, M. & Mlakar, P. (1993), 'A neural network-based method for short-term predictions of ambient SO 2 concentrations in highly polluted industrial areas of complex terrain', *Atmospheric Environment. Part B. Urban Atmosphere* **27**(2), 221–230.

Bruce, J., Kerling, J. & Beatty, W. (1985), 'On the North Brazilian eddy field', *Progress in oceanography* **14**, 57–63.

Budgell, W. (2005), 'Numerical simulation of ice-ocean variability in the Barents Sea region', *Ocean Dynamics* **55**(3-4), 370–387.

Carrier, M. J., Ngodock, H., Smith, S., Jacobs, G., Muscarella, P., Ozgokmen, T., Haus, B. & Lipphardt, B. (2014), 'Impact of assimilating ocean velocity observations inferred from Lagrangian drifter data using the NCOM-4DVAR', *Monthly Weather Review* **142**(4), 1509–1524.

Charney, J. G. & Eliassen, A. (1949), 'A numerical method for predicting the perturbations of the middle latitude westerlies', *Tellus A* **1**(2).

Charney, J. G., Fjörtoft, R. & Von Neumann, J. (1950), 'Numerical integration of the barotropic vorticity equation', *Tellus A* **2**(4).

Choi, B.-J. & Wilkin, J. L. (2007), 'The effect of wind on the dispersal of the Hudson River plume', *Journal of Physical Oceanography* **37**(7), 1878–1897.

Cochrane, J. D., Kelly Jr, F. J. & Olling, C. R. (1979), 'Subthermocline countercurrents in the western equatorial Atlantic Ocean', *Journal of Physical Oceanography* **9**(4), 724–738.

Comrie, A. C. (1997), 'Comparing neural networks and regression models for ozone forecasting', *Journal of the Air & Waste Management Association* **47**(6), 653–663.

Costa, F. & Tanajura, C. (2015), 'Assimilation of sea-level anomalies and Argo data into HYCOM and its impact on the 24 hour forecasts in the western tropical and South Atlantic', *Journal of Operational Oceanography* **8**(1), 52–62.

Csanady, G. (1985), 'A zero potential vorticity model of the North Brazilian Coastal Current', *Journal of Marine Research* **43**(3), 553–579.

Cummings, J. A. (2005), 'Operational multivariate ocean data assimilation', *Quarterly Journal of the Royal Meteorological Society* **131**(613), 3583–3604.

da Rocha Fragoso, M., de Carvalho, G. V., Soares, F. L. M., Faller, D. G., de Freitas Assad, L. P., Toste, R., Sancho, L. M. B., Passos, E. N., Böck, C. S., Reis, B. et al. (2016), 'A 4D-variational ocean data assimilation application for Santos Basin, Brazil', *Ocean Dynamics* **66**(3), 419–434.

Daley, R. (1991), 'Atmospheric data analysis, Cambridge atmospheric and space science series', *Cambridge University Press* **6966**, 25.

Dawson, A. (2016), 'eofs: A library for eof analysis of meteorological, oceanographic, and climate data', *Journal of Open Research Software* **4**(1).

Di Lorenzo, E., Moore, A. M., Arango, H. G., Cornuelle, B. D., Miller, A. J., Powell, B., Chua, B. S. & Bennett, A. F. (2007), 'Weak and strong constraint data assimilation in

the inverse Regional Ocean Modeling System (ROMS): Development and application for a baroclinic coastal upwelling system', *Ocean Modelling* **16**(3), 160–187.

Di Lorenzo, E. & Ohman, M. D. (2013), 'A double-integration hypothesis to explain ocean ecosystem response to climate forcing', *Proceedings of the National Academy of Sciences* **110**(7), 2496–2499.

Dibike, Y. B. & Coulibaly, P. (2006), 'Temporal neural networks for downscaling climate variability and extremes', *Neural Networks* **19**(2), 135–144.

Didden, N. & Schott, F. (1993), 'Eddies in the North Brazil Current retroflection region observed by Geosat altimetry', *Journal of Geophysical Research: Oceans (1978–2012)* **98**(C11), 20121–20131.

Dwivedi, S. & Pandey, A. C. (2011), 'Forecasting the Indian summer monsoon intraseasonal oscillations using genetic algorithm and neural network', *Geophysical Research Letters* **38**(15).

Egbert, G. D., Bennett, A. F. & Foreman, M. G. (1994), 'TOPEX/POSEIDON tides estimated using a global inverse model', *Journal of Geophysical Research: Oceans* **99**(C12), 24821–24852.

Egbert, G. D., Erofeeva, S. Y. & Ray, R. D. (2010), 'Assimilation of altimetry data for nonlinear shallow-water tides: Quarter-diurnal tides of the Northwest European Shelf', *Continental Shelf Research* **30**(6), 668–679.

EU Copernicus Marine Service (2015), 'Global ocean- in-situ near-real-time observations'. **URL:** *http://marine.copernicus.eu/*

Evensen, G. (1994), 'Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics'.

Evensen, G. (2003), 'The ensemble Kalman filter: Theoretical formulation and practical implementation', *Ocean dynamics* **53**(4), 343–367.

Flagg, C. N., Gordon, R. L. & McDowell, S. (1986), 'Hydrographic and current observations on the continental slope and shelf of the western equatorial Atlantic', *Journal of Physical Oceanography* **16**(8), 1412–1429.

Fox, D., Teague, W., Barron, C., Carnes, M. & Lee, C. (2002), 'The modular ocean data assimilation system (MODAS)', *Journal of Atmospheric and Oceanic Technology* **19**(2), 240–252.

Fratantoni, D. M., Johns, W. E. & Townsend, T. L. (1995), 'Rings of the North Brazil Current: Their structure and behavior inferred from observations and a numerical simulation', *Journal of Geophysical Research: Oceans (1978–2012)* **100**(C6), 10633–10654.

Fratantoni, D. M. & Richardson, P. L. (2006), 'The Evolution and Demise of North Brazil Current Rings*', *Journal of Physical Oceanography* **36**(7), 1241–1264.

Frolov, S. (2007), Enabling technologies for fast, nonlinear data assimilation in a coastal margin observatory, PhD thesis, Oregon Health and Science University.

Frolov, S., Baptista, A. M., Leen, T. K., Lu, Z. & van der Merwe, R. (2009), 'Fast data assimilation using a nonlinear Kalman filter and a model surrogate: An application to the Columbia River estuary', *Dynamics of Atmospheres and Oceans* **48**(1), 16–45.

Frolov, S., Paduan, J., Cook, M. & Bellingham, J. (2012), 'Improved statistical prediction of surface currents based on historic HF-radar observations', *Ocean Dynamics* **62**(7), 1111–1122.

Fukuoka, A. (1951), 'The Central Meteorological Observatory, A study on 10-day forecast (A synthetic report)', *Geophysical Magazine* **22**(3), 177–208.

Garcia-Gorriz, E. & Garcia-Sanchez, J. (2007), 'Prediction of sea surface temperatures in the western Mediterranean Sea by neural networks using satellite observations', *Geophysical research letters* **34**(11).

Gardner, M. & Dorling, S. (1996), Neural network modelling of the influence of local meteorology on surface ozone concentrations, *in* 'Proceedings 1st International Conference on GeoComputation', pp. 359–370.

Gauthier, P. (2003), Operational implementation of variational data assimilation, *in* 'Data Assimilation for the Earth System', Springer, pp. 167–176.

Ghil, M. (1997), 'Advances in sequential estimation for atmospheric and oceanic flows', *Meteorological society of Japan Series 2* **75**, 179–194.

Ghil, M. & Malanotte-Rizzoli, P. (1991), 'Data assimilation in meteorology and oceanography', *Adv. Geophys* **33**, 141–266.

GHRSST, J. M. M. P. (2010), 'GHRSST Level 4 MUR Global Foundation Sea Surface Temperature Analysis. Ver. 2. PO.DAAC, CA, USA'. Available from: http://dx.doi.org/10.5067/GHGMR-4FJ01 [Accessed 29th November 2016].

Global Runoff Data Center (2014), 'Federal Institute of Hydrology/Germany, and International Hydrological Programme/United Nations Educational/Scientific and Cultural Organization (2001), Monthly Discharge Data for World Rivers (except former Soviet Union), Research Data Archive at the National Center for Atmospheric Research, Computational and Information Systems Laboratory'.

Goodfellow, I., Bengio, Y. & Courville, A. (2016), *Deep learning*, MIT press.

Greischar, L. & Hastenrath, S. (2000), 'The rainy seasons of the 1990s in northeast Brazil: Real-time forecasts and verification', *Journal of Climate* **13**(21), 3821–3826.

Grewal, M. S. & Andrews, A. P. (2014), *Kalman filtering: Theory and Practice with MATLAB*, John Wiley & Sons.

Gurgel, K.-W., Antonischki, G., Essen, H.-H. & Schlick, T. (1999), 'Wellen Radar (WERA): a new ground-wave HF radar for ocean remote sensing', *Coastal engineering* **37**(3-4), 219–234.

Haidvogel, D. B., Arango, H. G., Hedstrom, K., Beckmann, A., Malanotte-Rizzoli, P. & Shchepetkin, A. F. (2000), 'Model evaluation experiments in the North Atlantic Basin: simulations in nonlinear terrain-following coordinates', *Dynamics of Atmospheres and Oceans* **32**(3), 239–281.

Hannachi, A. (2004), 'A primer for EOF analysis of climate data', *Department of Meteorology, University of Reading* pp. 1–33.

Hannachi, A., Jolliffe, I. & Stephenson, D. (2007), 'Empirical orthogonal functions and related techniques in atmospheric science: A review', *International Journal of Climatology* **27**(9), 1119–1152.

Hannachi, A., Jolliffe, I., Stephenson, D. & Trendafilov, N. (2006), 'In search of simple structures in climate: simplifying EOFs', *International journal of climatology* **26**(1), 7–28.

Hastenrath, S. & Greischar, L. (1993), 'Further work on the prediction of northeast Brazil rainfall anomalies', *Journal of Climate* **6**(4), 743–758.

Heinlein, R. A. (1987), *Time enough for love*, Penguin.

Hippert, H. S., Pedreira, C. E. & Souza, R. C. (2001), 'Neural networks for short-term load forecasting: A review and evaluation', *Power Systems, IEEE Transactions on* **16**(1), 44–55.

Hooten, M. B., Leeds, W. B., Fiechter, J. & Wikle, C. K. (2011), 'Assessing first-order emulator inference for physical parameters in nonlinear mechanistic models', *Journal of Agricultural, Biological, and Environmental Statistics* **16**(4), 475–494.

Houser, P. R. (2003), Assimilation of land surface data, *in* 'Data Assimilation for the Earth System', Springer, pp. 331–343.

Hunt, B. R., Kostelich, E. J. & Szunyogh, I. (2007), 'Efficient data assimilation for

spatiotemporal chaos: A local ensemble transform Kalman filter', *Physica D: Nonlinear Phenomena* **230**(1-2), 112–126.

Jochum, M. & Malanotte-Rizzoli, P. (2003), 'On the generation of North Brazil Current rings', *Journal of Marine research* **61**(2), 147–173.

Johns, W. E., Lee, T. N., Schott, F. A., Zantopp, R. J. & Evans, R. H. (1990), 'The North Brazil Current retroflection: Seasonal structure and eddy variability', *Journal of Geophysical Research: Oceans (1978–2012)* **95**(C12), 22103–22120.

Johns, W. E., Zantopp, R. J. & Goni, G. J. (2003), 'Cross-gyre transport by North Brazil Current rings', *Elsevier Oceanography Series* **68**, 411–441.

Johns, W., Lee, T., Beardsley, R., Candela, J., Limeburner, R. & Castro, B. (1998), 'Annual cycle and variability of the north brazil current', *Journal of Physical Oceanography* **28**(1), 103–128.

Josey, S. A., Kent, E. C. & Taylor, P. K. (2002), 'Wind stress forcing of the ocean in the SOC climatology: Comparisons with the NCEP–NCAR, ECMWF, UWM/COADS, and Hellerman and Rosenstein datasets', *Journal of Physical Oceanography* **32**(7), 1993–2019.

Kalman, R. E. (1960), 'A new approach to linear filtering and prediction problems', *Journal of basic Engineering* **82**(1), 35–45.

Kalnay, E. (2003), *Atmospheric modeling, data assimilation, and predictability*, Cambridge University Press.

Kim, H.-S., Kim, J.-H., Ho, C.-H. & Chu, P.-S. (2011), 'Pattern classification of typhoon tracks using the fuzzy c-means clustering method', *Journal of Climate* **24**(2), 488–508.

Kim, T.-W. & Valdés, J. B. (2003), 'Nonlinear model for drought forecasting based on a conjunction of wavelet transforms and neural networks', *Journal of Hydrologic Engineering* **8**(6), 319–328.

Kingma, D. P. & Ba, J. (2014), 'Adam: A method for stochastic optimization', *arXiv preprint arXiv:1412.6980* .

Knutti, R., Stocker, T., Joos, F. & Plattner, G.-K. (2003), 'Probabilistic climate change projections using neural networks', *Climate Dynamics* **21**(3-4), 257–272.

Kocarev, L., Galias, Z. & Lian, S. (2009), *Intelligent computing based on chaos*, Vol. 184, Springer.

Kondrashov, D. & Berloff, P. (2015), 'Stochastic modeling of decadal variability in ocean gyres', *Geophysical Research Letters* **42**(5), 1543–1553.

Krasnopolsky, V., Breaker, L. & Gemmill, W. (1995), 'A neural network as a nonlinear transfer function model for retrieving surface wind speeds from the special sensor microwave imager', *Journal of Geophysical Research: Oceans (1978–2012)* **100**(C6), 11033–11045.

Kutzbach, J. E. (1967), 'Empirical eigenvectors of sea-level pressure, surface temperature and precipitation complexes over North America', *Journal of Applied Meteorology* **6**(5), 791–802.

Larochelle, H., Erhan, D., Courville, A., Bergstra, J. & Bengio, Y. (2007), An empirical evaluation of deep architectures on problems with many factors of variation, *in* 'Proceedings of the 24th international conference on Machine learning', ACM, pp. 473–480.

LeCun, Y., Bengio, Y. & Hinton, G. (2015), 'Deep learning', *nature* **521**(7553), 436.

Lee, T. N., Johns, W. E., Zantopp, R. J. & Fillenbaum, E. R. (1996), 'Moored observations of western boundary current variability and thermohaline circulation at 26.5 in the subtropical north atlantic', *Journal of Physical Oceanography* **26**(6), 962–983.

Leeds, W. B., Wikle, C. K. & Fiechter, J. (2014), 'Emulator-assisted reduced-rank ecological data assimilation for nonlinear multivariate dynamical spatio-temporal processes', *Statistical Methodology* **17**, 126–138.

Lellouche, J.-M., Le Galloudec, O., Drévillon, M., Régnier, C., Greiner, E., Garric, G., Ferry, N., Desportes, C., Testut, C.-E., Bricaud, C. et al. (2013), 'Evaluation of global monitoring and forecasting systems at Mercator Ocean', *Ocean Science* **9**(1), 57.

Lins, I. D., Araujo, M., das Chagas Moura, M., Silva, M. A. & Droguett, E. L. (2013), 'Prediction of sea surface temperature in the tropical Atlantic by support vector machines', *Computational Statistics & Data Analysis* **61**, 187–198.

Lorenz, E. N. (1956), 'Empirical orthogonal functions and statistical weather prediction'.

Lorenz, E. N. (1963), 'Deterministic nonperiodic flow', *Journal of the atmospheric sciences* **20**(2), 130–141.

Lorenz, E. N. (1969), 'Atmospheric predictability as revealed by naturally occurring analogues', *Journal of the Atmospheric Sciences* **26**(4), 636–646.

Lorenz, E. N. (1970), 'Climatic change as a mathematical problem', *Journal of Applied Meteorology* **9**(3), 325–329.

Malanotte-Rizzoli, P., Hedstrom, K., Arango, H. & Haidvogel, D. B. (2000), 'Water mass pathways between the subtropical and tropical ocean in a climatological simulation of the North Atlantic ocean circulation', *Dynamics of Atmospheres and Oceans* **32**(3), 331–371.

Mandel, J. (2009), 'A brief tutorial on the ensemble Kalman filter', *arXiv preprint arXiv:0901.3725* .

Marchesiello, P., McWilliams, J. C. & Shchepetkin, A. (2003), 'Equilibrium structure and dynamics of the California Current System', *Journal of Physical Oceanography* **33**(4), 753–783.

Margvelashvili, N., Andrewartha, J., Herzfeld, M., Robson, B. J. & Brando, V. E. (2013), 'Satellite data assimilation and estimation of a 3D coastal sediment transport model using error-subspace emulators', *Environmental modelling & software* **40**, 191–201.

Margvelashvili, N. & Campbell, E. (2012), 'Sequential data assimilation in fine-resolution models using error-subspace emulators: Theory and preliminary evaluation', *Journal of Marine Systems* **90**(1), 13–22.

Margvelashvili, N. Y., Herzfeld, M., Rizwi, F., Mongin, M., Baird, M. E., Jones, E., Schaffelke, B., King, E. & Schroeder, T. (2016), 'Emulator-assisted data assimilation in complex models', *Ocean Dynamics* **66**(9), 1109–1124.
URL: *http://dx.doi.org/10.1007/s10236-016-0973-8*

Marmain, J., Molcard, A., Forget, P., Barth, A. & Ourmières, Y. (2014), 'Assimilation of HF radar surface currents to optimize forcing in the northwestern Mediterranean Sea', *Nonlinear Processes in Geophysics* **21**, 659–675.

Marzban, C. & Stumpf, G. J. (1996), 'A neural network for tornado prediction based on Doppler radar-derived attributes', *Journal of Applied Meteorology* **35**(5), 617–626.

Mattern, J. P., Fennel, K. & Dowd, M. (2012), 'Estimating time-dependent parameters for a biological ocean model using an emulator approach', *Journal of Marine Systems* **96**, 32–47.

McCulloch, W. S. & Pitts, W. (1943), 'A logical calculus of the ideas immanent in nervous activity', *The bulletin of mathematical biophysics* **5**(4), 115–133.

McDonnell, K. A. & Holbrook, N. J. (2004), 'A Poisson regression model of tropical cyclogenesis for the Australian-southwest Pacific Ocean region', *Weather and Forecasting* **19**(2), 440–455.

Miyoshi, T., Kalnay, E. & Li, H. (2013), 'Estimating and including observation-error correlations in data assimilation', *Inverse Problems in Science and Engineering* **21**(3), 387–398.

Møller, M. F. (1993), 'A scaled conjugate gradient algorithm for fast supervised learning', *Neural networks* **6**(4), 525–533.

Monahan, A. H., Fyfe, J. C., Ambaum, M. H., Stephenson, D. B. & North, G. R. (2009), 'Empirical orthogonal functions: The medium is the message', *Journal of Climate* **22**(24), 6501–6514.

Nair, V. & Hinton, G. E. (2010), Rectified linear units improve restricted boltzmann machines, *in* 'Proceedings of the 27th International Conference on Machine Learning (ICML-10)', pp. 807–814.

Navone, H. & Ceccatto, H. (1994), 'Predicting Indian monsoon rainfall: a neural network approach', *Climate Dynamics* **10**(6-7), 305–312.

Newinger, C. (2015), The barrier layer and ocean colour in the Amazon and Orinoco plume: Competing for the oceanic control on tropical cyclone intensity, PhD thesis, Imperial College London.

Newinger, C. & Toumi, R. (2015), 'Potential impact of the colored Amazon and Orinoco plume on tropical cyclone intensity', *Journal of Geophysical Research: Oceans* **120**(2), 1296–1317.

Nielsen, M. A. (2015), *Neural networks and deep learning*, Determination Press.

North, G. R., Bell, T. L., Cahalan, R. F. & Moeng, F. J. (1982), 'Sampling errors in the estimation of empirical orthogonal functions', *Monthly Weather Review* **110**(7), 699–706.

Obukhov, A. (1947), 'Statistically homogeneous fields on a sphere', *Usp. Mat. Nauk* **2**(2), 196–198.

Paduan, J. D. & Shulman, I. (2004), 'HF radar data assimilation in the Monterey Bay area', *Journal of Geophysical Research: Oceans* **109**(C7).

Palmer, T. N. (1999), 'A nonlinear dynamical perspective on climate prediction', *Journal of Climate* **12**(2), 575–591.

Palmer, T. N. (2001), 'A nonlinear dynamical perspective on model error: A proposal for non-local stochastic-dynamic parametrization in weather and climate prediction models', *Quarterly Journal of the Royal Meteorological Society* **127**(572), 279–304.

Panofsky, R. (1949), 'Objective weather-map analysis', *Journal of Meteorology* **6**(6), 386–392.

Patil, K. & Deo, M. C. (2017), 'Prediction of daily sea surface temperature using efficient neural networks', *Ocean Dynamics* **67**(3-4), 357–368.

Peak, J. E. & Tag, P. M. (1992), 'Toward automated interpretation of satellite imagery for navy shipboard applications', *Bulletin of the American Meteorological Society* **73**(7), 995–1008.

Pecora, L. M. & Carroll, T. L. (1990), 'Synchronization in chaotic systems', *Physical review letters* **64**(8), 821.

Philander, S. & Pacanowski, R. (1980), 'The generation of equatorial currents', *Journal of Geophysical Research: Oceans* **85**(C2), 1123–1136.

Philander, S. & Pacanowski, R. (1986), 'A model of the seasonal cycle in the tropical atlantic ocean', *Journal of Geophysical Research: Oceans* **91**(C12), 14192–14206.

Phillipson, L. & Toumi, R. (2017), 'Impact of assimilating OSCAR, altimetry, and limited drifter data on ocean current forecasting in the Angola Basin', *Ocean Modelling* .

Phillipson, L. & Toumi, R. (2018), 'The Crossover Time as an Evaluation of Ocean Models Against Persistence', *Geophysical Research Letters* **45**(1), 250–257.

Pohlmann, H. & Keenlyside, N. (2004), Review: Decadal-Multidecadal Climate Predictability, *in* 'Proceedings of the Workshop on Decadal Climate Variability, Hawaii', pp. 23–26.

Preisendorfer, R. W. & Mobley, C. D. (1988), *Principal component analysis in meteorology and oceanography*, Vol. 425, Elsevier Amsterdam.

Prybutok, V. R., Yi, J. & Mitchell, D. (2000), 'Comparison of neural network models with ARIMA and regression models for prediction of Houston's daily maximum ozone concentrations', *European Journal of Operational Research* **122**(1), 31–40.

Reynolds, R. W., Rayner, N. A., Smith, T. M., Stokes, D. C. & Wang, W. (2002), 'An improved in situ and satellite SST analysis for climate', *Journal of climate* **15**(13), 1609–1625.

Reynolds, R. W., Smith, T. M., Liu, C., Chelton, D. B., Casey, K. S. & Schlax, M. G. (2007), 'Daily high-resolution-blended analyses for sea surface temperature', *Journal of Climate* **20**(22), 5473–5496.

Richardson, L. (1922), 'Weather prediction by numerical process Cambridge University Press', *Cambridge* .

Richardson, P., Hufford, G., Limeburner, R. & Brown, W. (1994), 'North brazil current retroflection eddies', *Journal of Geophysical Research: Oceans* **99**(C3), 5081–5093.

Rio, M.-H., Mulet, S. & Picot, N. (2014*a*), 'Beyond GOCE for the ocean circulation estimate: Synergetic use of altimetry, gravimetry, and in situ data provides new insight into geostrophic and Ekman currents', *Geophysical Research Letters* **41**(24), 8918–8925.

Rio, M., Mulet, S. & Picot, N. (2014*b*), 'Beyond GOCE for the ocean circulation estimate: Synergetic use of altimetry, gravimetry, and in situ data provides new insight into geostrophic and Ekman currents', *Geophys. Res. Lett.* **41(24)**, 8918–8925.

Robertson, R., Beckmann, A. & Hellmer, H. (2003), 'M 2 tidal dynamics in the Ross Sea', *Antarctic Science* **15**(01), 41–46.

Röhrs, J., Christensen, K. H., Hole, L. R., Broström, G., Drivdal, M. & Sundby, S. (2012), 'Observation-based evaluation of surface wave effects on currents and trajectory forecasts', *Ocean Dynamics* **62**(10-12), 1519–1533.

Rojas, R. (2013), *Neural networks: a systematic introduction*, Springer Science & Business Media.

Saha, S., Moorthi, S., Pan, H., Wu, X., Wang, J., Nadiga, S., Tripp, P., Kistler, R., Woollen, J., Behringer, D. et al. (2010), 'NCEP climate forecast system reanalysis (CFSR) 6-hourly products, January 1979 to December 2010', *Research Data Archive at the National Center for Atmospheric Research, Computational and Information Systems Laboratory, Boulder, CO* .

Santoso, A. J., Dewi, F. K. S. & Sidhi, T. A. P. (2015), Natural disaster detection using wavelet and artificial neural network, *in* 'Science and Information Conference (SAI), 2015', IEEE, pp. 761–764.

Shchepetkin, A. F. & McWilliams, J. C. (2005), 'The regional oceanic modeling system (ROMS): a split-explicit, free-surface, topography-following-coordinate oceanic model', *Ocean Modelling* **9**(4), 347–404.

Shewchuk, J. R. et al. (1994), 'An introduction to the conjugate gradient method without the agonizing pain'.

Shlens, J. (2014), 'A tutorial on principal component analysis', *arXiv preprint arXiv:1404.1100* .

Shulman, I. & Paduan, J. D. (2009), 'Assimilation of HF radar-derived radials and total currents in the Monterey Bay area', *Deep Sea Research Part II: Topical Studies in Oceanography* **56**(3-5), 149–160.

Simmons, A., Wallace, J. & Branstator, G. (1983), 'Barotropic wave propagation and instability, and atmospheric teleconnection patterns', *Journal of the Atmospheric Sciences* **40**(6), 1363–1392.

Sparrow, C. (2012), *The Lorenz equations: bifurcations, chaos, and strange attractors*, Vol. 41, Springer Science & Business Media.

Stewart, L. M., Dance, S. L. & Nichols, N. K. (2007), 'Correlated observation errors in data assimilation', *Int. J. Numer. Meth. Fluids* (56), 1521–1527.

Swinbank, R., Shutyaev, V. & Lahoz, W. A. (2012), *Data assimilation for the earth system*, Vol. 26, Springer Science & Business Media.

Tabor, M. (1989), *Chaos and integrability in nonlinear dynamics: an introduction*, Wiley-Interscience.

Talagrand, O. (1997), 'Assimilation of observations, an introduction', *Meteorological Society of Japan - Series 2* **75**, 81–99.

Thompson, D. W. & Wallace, J. M. (1998), 'The Arctic Oscillation signature in the wintertime geopotential height and temperature fields', *Geophysical Research Letters* **25**(9), 1297–1300.

Thompson, D. W. & Wallace, J. M. (2000), 'Annular modes in the extratropical circulation. Part I: month-to-month variability*', *Journal of Climate* **13**(5), 1000–1016.

Treguier, A. M., Chassignet, E. P., Boyer, A. L. & Pinardi, N. (2017), 'Modeling and forecasting the weather of the ocean at the mesoscale', *Journal of Marine Research* **75**(3), 301–329.

Trigo, R. M. & Palutikof, J. P. (1999), 'Simulation of daily temperatures for climate change scenarios over Portugal: a neural network model approach', *Climate Research* **13**(1), 45–59.

van der Merwe, R., Leen, T. K., Lu, Z., Frolov, S. & Baptista, A. M. (2007), 'Fast neural network surrogates for very high dimensional physics-based models in computational oceanography', *Neural Networks* **20**(4), 462–478.

Wallace, J. M. & Gutzler, D. S. (1981), 'Teleconnections in the geopotential height field during the Northern Hemisphere winter', *Monthly Weather Review* **109**(4), 784–812.

Waller, J. A., Dance, S. L. & Nichols, N. K. (2016), 'Theoretical insight into diagnosing observation error correlations using observation-minus-background and observation-minus-analysis statistics', *Quarterly Journal of the Royal Meteorological Society* **142**(694), 418–431.

Wang, J., Flierl, G. R., LaCasce, J. H., McClean, J. L. & Mahadevan, A. (2013), 'Reconstructing the ocean's interior from surface data', *Journal of Physical Oceanography* **43**(8), 1611–1626.

Warner, J. C., Geyer, W. R. & Lerczak, J. A. (2005), 'Numerical modeling of an estuary: A comprehensive skill assessment', *Journal of Geophysical Research: Oceans* **110**(C5).

Weare, B. C. & Nasstrom, J. S. (1982), 'Examples of extended empirical orthogonal function analyses', *Monthly Weather Review* **110**(6), 481–485.

Weisstein, E. (2015), 'Lorenz Attractor'. Available from: `http://mathworld.wolfram.com/LorenzAttractor.html` [Accessed 29th November 2015].

Whitaker, J. S. & Hamill, T. M. (2002), 'Ensemble data assimilation without perturbed observations', *Monthly Weather Review* **130**(7), 1913–1924.

White, R. & Toumi, R. (2014), 'River flow and ocean temperatures: The Congo River', *Journal of Geophysical Research: Oceans* **119**(4), 2501–2517.

Williams, R. G. & Follows, M. J. (2011), *Ocean dynamics and the carbon cycle: Principles and mechanisms*, Cambridge University Press.

Yi, J. & Prybutok, V. R. (1996), 'A neural network model forecasting for prediction of daily maximum ozone concentration in an industrialized urban area', *Environmental Pollution* **92**(3), 349–357.

Zhang, G., Patuwo, B. E. & Hu, M. Y. (1998), 'Forecasting with artificial neural networks: The state of the art', *International journal of forecasting* **14**(1), 35–62.

Zupanski, M. & Kalnay, E. (1999), 'Principles of data assimilation', *Global energy and water cycles. Cambridge Univ Press, Cambridge* pp. 48–54.