
HYBRID DATA ASSIMILATION TECHNIQUES USING THE ADJOINT METHOD IN A COUPLED LORENZ SYSTEM

A PREPRINT

Philip David Kennedy^{1,*} 

Abhirup Banerjee¹ 

Armin Köhl¹ 

Detlef Stammer¹ 

¹Institute of Oceanography, Center for Earth System Research and Sustainability (CEN), Universität Hamburg, Hamburg, Germany.

*philip.david.kennedy@uni-hamburg.de

5th March 2024

ABSTRACT

A hybrid 4D-variational data assimilation method for numerical climate models is introduced using the Lorenz '63 model. This new approach has the potential to optimise a high complexity Earth system model (ESM) by utilising the adjoint equations of an intermediate complexity ESM. The method is conceptually demonstrated by consecutively synchronising two Lorenz '63 systems to observations before optimisation. The first represents a 'high complexity' model and the second an 'intermediate complexity' model which has adjoint equations. This method will save computational power for a full ESM and has negligible error and uncertainty change compared to the optimisation of a single model with adjoint equations. A similar setup can be applied to sparse observations. An alternative assimilation setup, with two identical models, is used to filter noisy data. This reduces optimised parametric model uncertainty by approximately one third. Such a precision gain could prove valuable for seasonal, annual, and decadal predictions.

Keywords Lorenz '63, data assimilation, adjoint, precision, uncertainty, climate models, error estimates

1 Introduction

The Earth system can be realistically described using numerical models that capture the system's processes and time evolution. ESMs can be used to forecast future states of the system provided that the past and present are known. Data assimilation is a powerful tool that combines observations and a numerical model representing the dynamic principles governing the system to generate an estimation of its state [Wunsch and Heimbach, 2006, Nichols, 2010] in an attempt to further improve the ESMs predictive skills.

Generally, there are two common approaches to incorporate observations into a model, a *sequential data assimilation scheme* [Bertino et al., 2003] and the *variational approach* [Le Dimet and Talagrand, 1986]. In this work, we use the *variational data assimilation* approach, which involves minimising a scalar, also called *cost function*, defined as the quadratic misfit between the observational and model data within an assimilation time window. This method is commonly known as *four-dimensional variational assimilation* (4D-Var.) [Rabier and Liu, 2003], which utilises an adjoint of the model to iteratively minimise the model-data misfit, by adjusting control variables [Tett et al., 2017, Lyu et al., 2018, Köhl and Willebrand, 2002, Allaire, 2015, Navon, 2009]. Aside from data assimilation purposes, adjoint models have been widely used for sensitivity analysis in meteorology and oceanography [Hall et al., 1982, Hall and Cacuci, 1983, Hall, 1986, Marotzke et al., 1999, Stammer et al., 2016]; for instance, calculating sensitivity with respect to lateral boundary condition [Gustafsson et al., 1998], estimating the sensitivity of the 2m surface temperature with respect to the sea surface temperature, sea ice, sea surface salinity [Stammer et al., 2018], and many more.

Due to the non-linearities within ESMs, the relevance of adjoint method, and thus the gradient of the cost function in reducing model-data misfits, can be limited by the predictability time scale. This can lead to spikes in the adjoint

sensitivities which results in multiple local minima in the cost-function. Under such circumstances spikes occur in the estimated gradients and the cost function becomes very rough by showing an increasing number of local minima [Köhl and Willebrand, 2002, Lea et al., 2000]. Fortunately the problem can be mitigated by synchronisation which removes the non-linear or chaotic dynamics from the adjoint model leading to a smooth cost function [Abarbanel et al., 2010, Sugiura et al., 2014]. This method allows to extent the assimilation window beyond the predictability time-scale, provided that sufficient observations are available. Another challenge of applying an adjoint model to any state-of-the-art system is that these ESMs have a very large number of state variables $\mathcal{O}(10^7 - 10^8)$, and the computation of its adjoint would be very time consuming and require large volumes of memory [Stammer et al., 2018].

To mitigate both problems, we propose a novel framework where we use two climate models both coupled through synchronisation, one with a high resolution and complex parametrisation scheme and the other with coarse resolution and intermediate complexity but for which an adjoint exist. In this context we only use the adjoint of the low-resolution model to estimate the parameters that optimise the high-resolution model; i.e., that brings the model closer to the assimilated observations. A pre-requisite of this method is that the parameters being optimised must be the same in both models. The objective of this paper is to quantify the precision and the benefit of such a synchronised data assimilation approach. We perform this test conceptually using a Lorenz '63 model system.

The Lorenz '63 system [Lorenz, 1963] is a well established proxy model to study chaotic fluid systems, such as the atmosphere [Gauthier, 1992, Miller et al., 1994, Pires et al., 1996, Stensrud and Bao, 1992, Kravtsov and Tsonis, 2021, Huai et al., 2017, Yang et al., 2006, Daron and Stainforth, 2015, Errico, 1997]. The advantage is that it can be used to quantitatively evaluate the parameter dependence of the system prior to application in a full model. New modelling techniques can thus be trialled in relatively fast experiments [Pasini and Pelino, 2005, Tandeo et al., 2015, Goodliff et al., 2020, Marzban, 2013, Yin et al., 2014]. It can also be used in a wide range of other applications [Du and Shiue, 2021, Cameron and Yang, 2019, Pelino and Maimone, 2007]. The system generates a three dimensional, time varying trajectory which with variation of both model parameters and/or initial conditions will produce very different trajectories. Thus, it is an ideal test bed for testing non-linear modelling in a number of fields [Hirsch et al., 2013]. In a previous study [Lyu et al., 2018] using the Lorenz '63 model, the assimilation primarily focused on the fit of a single parameter ρ and the initial conditions (x, y, z) in a single model with an adjoint. This study will expand on this to fit all three model parameters simultaneously and use a model with adjoint to optimise the parameters of a model without one.

The structure of the remaining paper is as follows: In Section 2 we introduce the model, outline the methodology of how these models are synchronised, show how the adjoint method is applied, describe our proposed multi-model setup, and detail the fitting procedure, including details on the minimisation algorithm. Section 3 shows and discusses the results of our multi-model setups, using a single model setup as a baseline for comparison. The results of introducing a mismodelling term to the adjoint model are also included along with a study of the setup's response to sparse data conditions. Our results are summarised and conclusions discussed in Section 4.

2 Methodology

2.1 Lorenz '63 model

In this study, we use the coupled Lorenz '63 system as described by Yang et al. [2006]. The model is defined by the equations:

$$\frac{dx}{dt} = \sigma(y - x), \quad (1a)$$

$$\frac{dy}{dt} = \rho x - y - xz, \quad (1b)$$

$$\frac{dz}{dt} = xy - \beta z \quad (1c)$$

where (x, y, z) are the state variables at each given time step and σ , ρ , and β are the model parameters. Throughout this article, we integrate all our models using the fourth-order Runge-Kutta method with a step size of $\Delta t = 0.01$ and total time period of 100 time units [TUs]. This system of equations will be referred to as the *true* model for which we will apply the standard values for $\sigma_t = 10$, $\rho_t = 28$, and $\beta_t = 8/3$. This *true* model is used to generate pseudo-observation which will be used to synchronise our physical models in Sub-section 2.3. Noise is included in these pseudo-observations by adding random values from a Gaussian distribution centred at zero to the *true* values. The random noise value magnitudes are bounded by a given percentage relative to the systems' standard distribution. These pseudo-observations will be labelled as (x_o, y_o, z_o) .

2.2 Adjoint method

For the Lorenz '63 system outlined in Eq. (1) the adjoint matrix can be derived by transposing its tangent liner matrix (TLM). The TLM follows from Eq. (1) as:

$$M^T \equiv \frac{\partial \dot{\vec{x}}_i}{\partial \vec{x}_j} = \begin{pmatrix} -\sigma & \sigma & 0 \\ \rho - z_a & -1 & -x_a \\ y_a & x_a & -\beta \end{pmatrix}. \quad (2)$$

The adjoint M^* is then given by $M^* \equiv M^T$. As part of the adjoint model assimilation process, the adjoint equations are integrated backwards in time to calculate the gradient of the cost function with respect to control parameters which are used subsequently in an iterative process to adjust the control parameters such that the system is brought into consistency with observations.

2.3 Synchronisation

A fundamental limitation of the adjoint method arises when integrating over periods that are longer than the predictability time scale of a system, leading to spikes in the estimated gradients and a cost function with an increasing number of local minima [Köhl and Willebrand, 2002, Lea et al., 2000]. The problem can be mitigated by synchronisation which removes the non-linear or chaotic dynamics from the adjoint model [Abarbanel et al., 2010, Sugiura et al., 2014]. To incorporate a synchronisation technique we expand the Lorenz '63 model by so called nudging terms which then reads as:

$$\frac{dx_a}{dt} = \sigma(y_a - x_a) + \alpha(x_o - x_a), \quad (3a)$$

$$\frac{dy_a}{dt} = \rho x_a - y_a - x_a z_a + \alpha(y_o - y_a), \quad (3b)$$

$$\frac{dz_a}{dt} = x_a y_a - \beta z_a. \quad (3c)$$

Here α is the synchronisation constant, sub-script a denotes the model with an adjoint, (x_o, y_o, z_o) are the pseudo-observations generated from the *true* model. For this synchronised model we will apply the standard values used in the *true* system plus a 10%-error, which gives $\sigma_a = 11$, $\rho_a = 30.8$, and $\beta_a = 44/15$. These will act as our initial conditions for the parametric fit.

The adjoint matrix for the synchronised model in Eq. (3) is given by:

$$M_a^* = \begin{pmatrix} -(\sigma + \alpha) & \rho - z_a & y_a \\ \sigma & -(1 + \alpha) & x_a \\ 0 & -x_a & -\beta \end{pmatrix}. \quad (4)$$

The synchronisation constant is the only difference from the adjoint of Eq. (2). The significance of this change will be shown in Section 3, as α has a critical role in the precision with which parameters can be estimated, due to its influence on both the cost function and its gradient.

There are seven possible combinations of the three state variables which can be synchronised. The effect of each of the possible choices on the root mean squared error (RMSE) between the *true* and adjoint systems is shown in Fig. 1 by varying the synchronisation constant α from 0 to 30. The figure demonstrates that synchronising the z -component is ineffective at reducing the RMSE [Yang et al., 2006]. In contrast, synchronising both x and y prove effective, with y leading the lowest RMSE values of the single variable for all values of α . Synchronising xyz and xy achieve the most effective reduction in RMSE for the lowest value of α . It can be seen that synchronising z can lead to model instability. Thus, we choose to only synchronise x and y in the following, as seen in Eq. 3, to achieve more stable results with negligible precision loss.

The Lorenz '63 attractors for the trajectories of the *true* model and that with an adjoint are shown in Fig. 2a without synchronisation. A large divergence is visible between the attractors. However, if synchronisation is introduced, as shown in Fig. 2b, the attractors become very similar. There is now significant overlap between their kernel density estimations (KDEs). We choose to use KDEs throughout this paper as they represent a smoothed estimate of the PDF for the attractor. This allows for convenient visual comparison of attractors. A more numerically rigorous method to check for effective synchronisation will be discussed in section 3.

2.4 Multi-model data assimilation

This technique consecutively synchronises two forward models before running the adjoint of the second model backward in time. For this purpose, Eq. 3 must be modified to incorporate a consecutive synchronisation. A schematic of this




Figure 1: Different coupling schemes are trialled. RMSE is calculated between the *truth* and *model* trajectories after synchronisation to pseudo-observations. Noise was added (with zero mean and $\sqrt{2}$ standard deviation) to the *truth* when constructing the pseudo-observations. The synchronisation constant α is varied from 0 to 30 in steps of 0.5.

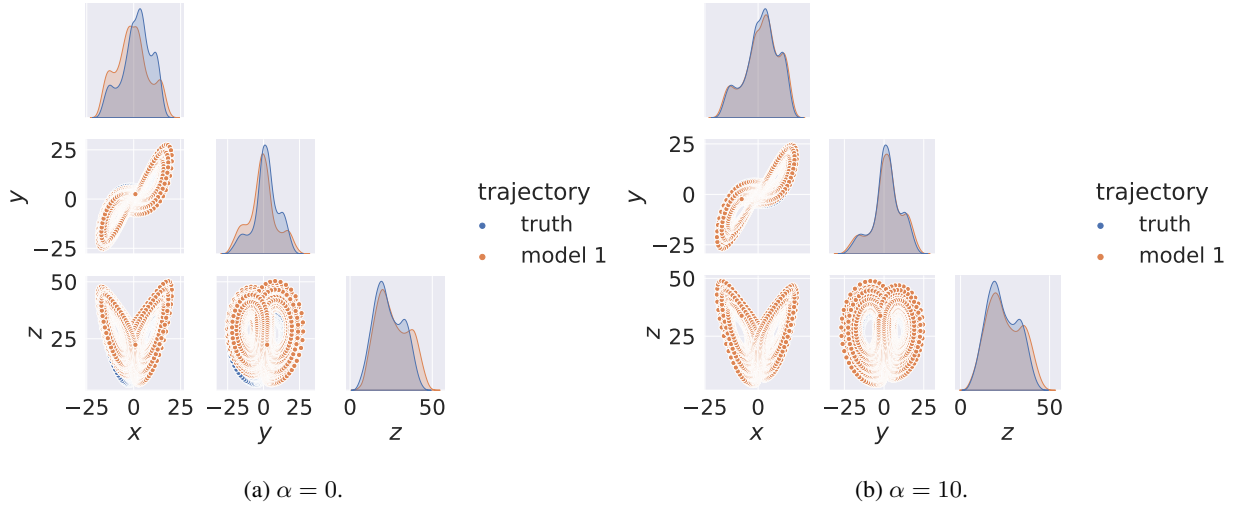


Figure 2: The bottom left quadrants show the Lorenz '63 *truth* and *model* attractors from the main three variable orientations. The diagonal plots show kernel density estimations (KDEs). Fig. 2a shows the attractors without synchronisation. Fig. 2b shows the attractors with synchronisation.

setup is provided in Fig.3 and the implications of the two possible ways to calculate the cost function are discussed in

the subsequent subsections. The equations of model 1, which is run only in forward mode, are

$$\frac{dx_f}{dt} = \sigma(y_f - x_f) + \alpha(x_o - x_f), \quad (5a)$$

$$\frac{dy_f}{dt} = \rho x_f - y_f - x_f z_f + \alpha(y_o - y_f), \quad (5b)$$

$$\frac{dz_f}{dt} = x_f y_f - \beta z_f \quad (5c)$$

where the sub-script f denotes the forward run of model 1 and sub-script o denotes observations generated from *truth* model. The system of equations for the model 2 which has an adjoint will now be modified to synchronise with the forward-only model and not the observations:

$$\frac{dx_a}{dt} = \sigma(y_a - x_a) + \alpha(x_f - x_a), \quad (6a)$$

$$\frac{dy_a}{dt} = \rho x_a - y_a - x_a z_a + \alpha(y_f - y_a), \quad (6b)$$

$$\frac{dz_a}{dt} = x_a y_a - \beta z_a \quad (6c)$$

where the sub-script a denotes model 2 which has an adjoint. This model synchronises with the model 1 but never directly with the observations.




Figure 3: Illustration of the multi-model setup where each pseudo-observation generated from the truth includes random additive Gaussian noise. The cost function can measure the difference between the observations and either model 1 or 2 depending on the assumptions made. Both options are discussed in the text.

To assimilate the data, we fit one of the synchronised models to the observations by optimising the model parameters. A cost function to calculate the misfit between observations and the model can be attached to model 1 or 2. The gradient of the cost function, with respect to the model parameters, is always calculated using the adjoint method associated with

model 2. The adjoint model is numerically evaluated by automatic differentiation of the forward model 2. The process of synchronising all models, calculating the cost function and its gradient, and then adjusting model parameters is carried out iteratively by our chosen minimisation algorithm. Throughout all steps the parameter value of forward-only and adjoint models are identical and optimised simultaneously. The procedure stops when the algorithm-specific stopping criteria is met. The minimum of the cost function is the point at which the model provides the maximum likelihood estimate.

For the present study we have chosen to use the minimisation package `iminuit` [Dembinski and et al., 2020]. The version currently used is the `Minuit2` algorithm `MIGRAD` which is a variable-metric method including an inexact line search, a stable metric updating scheme, and it checks for positive-definiteness. This is chosen as it is a robust and stable minimiser that both utilises the gradient we derived using adjoint model in Eq. (11), and calculates accurate uncertainty estimates for all fitted parameters. Further details on the specific minimisation algorithm are given in James and Roos [1975].

2.4.1 Setup 1 - state filtered data assimilation (SFDA)

Assuming that both Lorenz models can be thought of as representing two identical climate models, the cost function can be placed on the model 2 which has adjoint equations. This allows model 1 to filter out some of the background noise on the observations before they are given to the cost function attached to model 2. We will subsequently refer to this setup as state filtered data assimilation (SFDA.) This filtering should act to increase the sensitivity of the cost function to the parameters by reducing the impact of observational noise which favourably increases the signal to background ratio synchronised into model 2. Such a filtering setup would theoretically reduce parametric uncertainty below that of traditional single model data assimilation because model 1 should act to reduce the amount of noise synchronised into model 2. This produces a steeper, more-pronounced cost function increasing the parametric sensitivity which reduces the uncertainty on each fitted parameter.

In SFDA the cost function acts to constrain model 2 which has an adjoint. The cost function is

$$J_{\text{SFDA}} = \frac{1}{2N} \sum_{t=0}^N (\vec{x}_o(t) - \vec{x}_a(t))^T \frac{1}{\sigma_{\vec{x}_o}^2} (\vec{x}_o(t) - \vec{x}_a(t)) \quad (7)$$

N is the total number of time steps of the assimilation window, and $\sigma_{\vec{x}_o}$ is the uncertainty associated with the observation noise. The gradient will be calculated subsequently using the adjoint of model 2.

2.4.2 Setup 2 - hybrid data assimilation (HDA)

Here we want to explore if using an existing adjoint from one model could be utilised to optimise a different target model without adjoint. This will be referred to as hybrid data assimilation (HDA.) In HDA we assume that both models differ in their equations and spatial resolution. Instead of interpolating or transforming the original model onto the adjoint model grid, formulation of the adjoint model through synchronisation would provide a simpler means to do this, as only essential parameters need to be optimised. Auxiliary variables and parameters, including those that may not exist in the target model, will be adjusted via the model dynamics.

The cost function of HDA is

$$J_{\text{HDA}} = \frac{1}{2N} \sum_{t=0}^N (\vec{x}_o(t) - \vec{x}_f(t))^T \frac{1}{\sigma_{\vec{x}_o}^2} (\vec{x}_o(t) - \vec{x}_f(t)). \quad (8)$$

N is the total number of time steps of the assimilation window and $\sigma_{\vec{x}_o}$ is the uncertainty associated with the observation noise. This measures the quadratic misfit between the forward-only model 1 and the observations. Model 1 \vec{x}_f will be constrained by this cost function and its gradient will be calculated using the adjoint of model 2 \vec{x}_a .

2.4.3 Cost function gradient

The gradient of the cost function, with respect to the parameters, will be identical for both methods above. This is because adjoint equations will always be associated with model 2. The gradient of the cost function, with respect to the state variables at $t = 0$, is given by

$$\vec{\nabla}_{\vec{x}_a} J = \frac{1}{N} \vec{\lambda}_{\vec{x}_a}(0), \quad (9)$$

where $\vec{\lambda}$ is the adjoint vector with respect to the state variables calculated in the reverse time direction:

$$\vec{\lambda}_{\vec{x}_a}(N+1) = \vec{0}, \quad (10a)$$

$$\begin{aligned} \vec{\lambda}_{\vec{x}_a}(t) &= \vec{x}_o^*(t) \frac{1}{\sigma_{\vec{x}_o}} (\vec{x}_o(t) - \vec{x}_a(t)) \\ &\quad - M^*(t) \vec{\lambda}_{\vec{x}_a}(t+1). \text{ for } t = N, \dots, 0 \end{aligned} \quad (10b)$$

M^* is the adjoint matrix, calculated by automatic differentiation, and $\vec{x}_o^*(t)$ is the adjoint of the observation. These equations were derived using the method detailed in Talagrand [2010]. This gradient can then be calculated with respect to the parameters (σ, ρ, β) notated by the subscript \vec{p} . This yields

$$\vec{\nabla}_{\vec{p}} J = \frac{1}{N} \sum_{t=N}^0 \vec{\lambda}_{\vec{x}_a}(t) \cdot \begin{pmatrix} y_a(t) - x_a(t) \\ x_a(t) \\ -z_a(t) \end{pmatrix} \quad (11)$$

where the vector derived using the chain rule is multiplied by each term from Eq. (10) before being summed.

3 Results

Throughout the following section we will use the single model described in Lyu et al. [2018] as our benchmark to compare the new setups against. To understand the behaviour of the setups at different operating extremes, assimilations are carried out for variations of observational noise and α . This will help establish the optimal synchronisation strength dependent on the noise amplitude. We will also be able to compare the errors and uncertainties of the single model with our multi-model setups.

To get a more accurate picture of our setup's behaviour, it is necessary to repeat our study over a number data sets to calculate medians and percentile intervals. This allows us to examine general traits of our model without an individual noise event obscuring trends and features of significance. Here this is done by generating 100 pseudo-data sets and assimilating each set independently. The plotting package is then applied directly to these 100 outputs to plot the median and 68% percentile intervals. The percentile intervals are included to illustrate the statistical spread of the results and reproducibility, not to explicitly indicate uncertainty. The mean percentage error and uncertainty are plotted separately.^{1,2} The error is simply calculated by percentage difference between the fitted and *true* values. The uncertainty is calculated by the minimisation algorithm using a Hessian estimate.

3.1 SFDA (setup 1)

The results from a scan of α are shown in Fig. 4. The single model scan has two main regions. The first, $\alpha \leq 7.5$, is where the system is poorly synchronised giving an inaccurate fit of the parameters. The second, above $\alpha > 7.5$, is where the system is fully synchronised and recovers the *true* model parameters very effectively. SFDA has a higher onset of effective synchronisation than the single model setup, beginning at $\alpha = 11$. Above $\alpha = 12.5$, SFDA has consistently more accurate parameter recovery than the single model setup.

Fig. 5 shows the results of two fits carried out with noise of 25%. The mean percentage uncertainty over the three parameters is plotted for both setups. Noticing in particular the spread of the percentile intervals, the single model setup is found to be synchronised and have a high precision from $\alpha = 7.5$ and SFDA from $\alpha = 11.5$. Once the SFDA setup is synchronised, it is found to have a reduced uncertainty compared with the single model. SFDA is found to be approximately one third more precise better than the single model setup for all value of α investigated.

Fig. 6 shows the results of varying the noise levels on the fitted parameter values. For all noise levels studied the quality of the fit can be considered good as the median of the mean percentage uncertainty on the parameters remains below 0.5% even with noise levels of up to 50%. SFDA is found to have similar mean error performance to the single model system across the range of noise levels tested. The spread of the error is slightly improved in the double model setup at low noise likely due to the forward-only model 'smoothing' outlying observation better than a single model setup. The parametric uncertainty is found to be consistently reduced in the double model system for all noise levels. This demonstrates the precision improvement achieved by running the forward model twice to 'smooth' the observations before carrying out data assimilation. The consequences of this are that for smaller or localised climate models, where computational resources are available and improved precision is desirable, SFDA could reduce error and particularly uncertainty.

¹mean %-error = $100\% \cdot \frac{1}{3} \cdot \left(\left| \frac{\sigma - \sigma_t}{\sigma_t} \right| + \left| \frac{\rho - \rho_t}{\rho_t} \right| + \left| \frac{\beta - \beta_t}{\beta_t} \right| \right)$
²mean %-uncertainty = $100\% \cdot \frac{1}{3} \cdot \left(\left| \frac{\Delta\sigma}{\sigma_t} \right| + \left| \frac{\Delta\rho}{\rho_t} \right| + \left| \frac{\Delta\beta}{\beta_t} \right| \right)$




Figure 4: The percentage error between the *true* values of (σ, ρ, β) and the fitted value from SFDA. A single model assimilation is included for comparison. An ensemble of 100 assimilations is carried out over 100 different data sets. The median (lines) and 68% percentile intervals (shaded areas) are plotted. The noise level is 25%.





Figure 5: The average percentage uncertainty on the three parameters (σ, ρ, β) after SFDA from the minimisation algorithm for different values of α . A single model assimilation is included for comparison. An ensemble of 100 assimilations is carried out over 100 different data sets. The median (line) and 68% percentile intervals (shaded areas) are plotted.



(a) Percentage error.

(b) Percentage uncertainty.

Figure 6: The percentage error between the *true* values of (σ, ρ, β) and those from SFDA, as well as average percentage uncertainty on the SFDA parameters. A single model assimilation is included for comparison. An ensemble of 100 assimilations is carried out over 100 different data sets. The median (line) and 68% percentile intervals (shaded areas) are plotted. The noise level varies between 5% and 50% in steps of 5%.

3.2 HDA (setup 2)

The results from a scan of α are shown in Fig. 7. The single model scan has two main regions. The first, $\alpha \leq 7.5$, is where the system is poorly synchronised giving an inaccurate fit of the parameters. The second, $\alpha \geq 7.5$, is where the system is synchronising very effectively and recovers the true model parameters very accurately. The HDA scan follows the behaviour of the primary model very closely.

Fig. 8 shows the results of two fits carried out with noise of 25%. The mean percentage uncertainty over the three parameters is plotted for both setups. HDA is found to have almost identical uncertainty to the single model. The plot consists of two regions. The first, for $\alpha < 7.5$ is the region where the model is not yet consistently synchronised producing high variability depending on the specific noise. The second, $\alpha \geq 7.5$ is where the system is consistently synchronised. The minimum median of the mean parametric uncertainty, after consistent synchronisation begins, is $\approx 0.35\%$ and achieved at $\alpha = 7.5$. The subsequent increase in uncertainty is due to the increased α both reducing and flattening the cost function which reduces the parametric sensitivity of the fit.

As the uncertainty of the parameters is independent of knowledge of the true parameters, this post-fit observable can be plotted for all climate models. Thus, we propose a quantitative procedure for verifying the optimal coupling strength, beyond simply observing synchronisation in the time series, in variational data assimilation. This procedure involves finding point of minimum parametric uncertainty in the linear region of the graph. It occurs at the first point the system is successfully synchronised, in this case $\alpha = 7.5$. We acknowledge that such a procedure is only applicable on models with lower numbers of parameters where the global minimum of the cost function has been successfully found and running an assimilation multiple times is not adversely computationally costly. Running pre-fit scans of the cost function against a varying parameter would provide a more computationally efficient alternative to check if the cost function is smooth with no local minima for a fixed value of α .

Fig. 9 shows the results of varying the noise levels on the fitted parameter values. For all noise levels studied the quality of the fit can be considered good as the mean percentage error on the parameters remains below 1% even with noise levels of 50%. The increased spread of the error results at low noise is thought to be the adjoint experiencing reduced effects of the noise compared to the single model. The length of the time window used is 100TUs, but if the window were increased the accuracy of the results would improve. The HDA setup is found to have extremely consistent uncertainty compared to the single model system.




Figure 7: The percentage error between the *true* values of (σ, ρ, β) and those from HDA. A single model assimilation is included for comparison. An ensemble of 100 assimilations is carried out over 100 different pseudo-data sets. The median (lines) and 68% percentile intervals (shaded areas) are plotted. The noise level is 25%.

4 Conclusion

In this paper we have demonstrated the ability to constrain a Lorenz '63 model using a second model with an adjoint by 4D-Var data assimilation. Such an approach removes the need to generate an adjoint for a forward model, if such an adjoint already exists for a separate, yet similar system. An important application of this technique in Earth system modelling would be a situation where a low-resolution ESM with an adjoint shares a parametrisation with a high-resolution ESM without an adjoint. The low-resolution ESM could provide all necessary variables, through synchronisation, to run its tangent linear adjoint model. This can then be utilised to estimate parameters in the high-resolution ESM without an adjoint. Moreover, using a simpler lower resolution version could make data assimilation quicker and would use fewer computational resources than an adjoint of the high-resolution ESM. We have also shown that running a forward model twice before beginning data assimilation can act to smooth the data and reduce the parametric uncertainty by roughly one third compared to a single forward run. Our overarching attention in this approach is optimising the parameters of a full Earth System model which will be tested as a next step. Nevertheless, it would also be possible to optimise the state variables which are more applicable to weather forecasting techniques. Future work will evaluate the resilience of hybrid data assimilation to mismodelling in the lower resolution model with an adjoint and examine how all setups in this research behave when assimilating sparse observations.

Acknowledgements

The figures in this research were plotted by the seaborn plotting package [Waskom, 2021] utilising our output which was managed using pandas [pandas development team, 2020]. Automatic differentiation of our model was carried out using JAX [Bradbury et al., 2018]. The minimisation of our cost function and uncertainty evaluation was done by iminuit [Dembinski and et al., 2020]. This research was supported, in part, through the Koselleck grant *Earth^{RA}* from the Deutsche Forschungsgemeinschaft. Contribution to the Centrum für Erdsystemforschung und Nachhaltigkeit (CEN) of Universität Hamburg. This work used resources of the Deutsches Klimarechenzentrum (DKRZ) granted by its Scientific Steering Committee (WLA).

References

Carl Wunsch and Patrick Heimbach. Estimated decadal changes in the north atlantic meridional overturning circulation and heat flux 1993–2004. *Journal of Physical Oceanography*, 36(11):2012 – 2024, 2006.




Figure 8: The percentage uncertainty, from the minimisation algorithm, averaged over all three parameters (σ, ρ, β) after HDA. A single model assimilation is included for comparison. An ensemble of 100 assimilations is carried out over 100 different data sets. The median (lines) and 68% percentile intervals (shaded areas) are plotted.

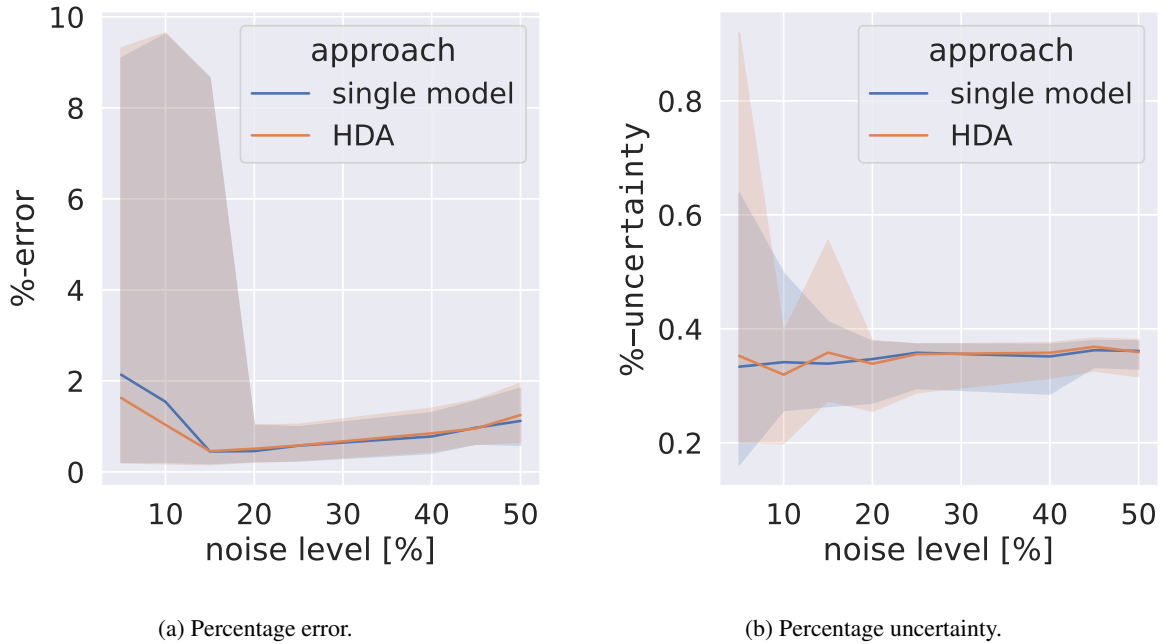


Figure 9: The percentage error between the *true* values of (σ, ρ, β) and those from HDA. A single model assimilation is included for comparison. An ensemble of 100 assimilations is carried out over 100 different data sets. The median (lines) and 68% percentile intervals (shaded areas) are plotted. The noise level varies between 5% and 50% in steps of 5%.

- N. K. Nichols. *Mathematical Concepts of Data Assimilation*, pages 13–39. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010. ISBN 978-3-540-74703-1. doi:10.1007/978-3-540-74703-1_2. URL https://doi.org/10.1007/978-3-540-74703-1_2.
- Laurent Bertino, Geir Evensen, and Hans Wackernagel. Sequential data assimilation techniques in oceanography. *International Statistical Review*, 71(2):223–241, 2003.
- François-Xavier Le Dimet and Olivier Talagrand. Variational algorithms for analysis and assimilation of meteorological observations: theoretical aspects. *Tellus A: Dynamic Meteorology and Oceanography*, 38(2):97–110, 1986.
- Florence Rabier and Zhiqian Liu. Variational data assimilation: theory and overview. In *Proc. ECMWF Seminar on Recent Developments in Data Assimilation for Atmosphere and Ocean, Reading, UK, September 8–12*, pages 29–43, 2003.
- Simon FB Tett, Kuniko Yamazaki, Michael J Mineter, Coralia Cartis, and Nathan Eizenberg. Calibrating climate models using inverse methods: case studies with hadam3, hadam3p and hadcm3. *Geoscientific Model Development*, 10(9):3567–3589, 2017.
- Guokun Lyu, Armin Köhl, Ion Matei, and Detlef Stammer. Adjoint-based climate model tuning: Application to the planet simulator. *Journal of Advances in Modeling Earth Systems*, 10(1):207–222, 2018.
- Armin Köhl and Jürgen Willebrand. An adjoint method for the assimilation of statistical characteristics into eddy-resolving ocean models. *Tellus A: Dynamic meteorology and oceanography*, 54(4):406–425, 2002.
- Grégoire Allaire. A review of adjoint methods for sensitivity analysis, uncertainty quantification and optimization in numerical codes. *Ingénieurs de l’Automobile*, 836:33–36, 2015.
- Ionel M Navon. Data assimilation for numerical weather prediction: a review. *Data assimilation for atmospheric, oceanic and hydrologic applications*, pages 21–65, 2009.
- Matthew CG Hall, Dan G Cacuci, and Michael E Schlesinger. Sensitivity analysis of a radiative-convective model by the adjoint method. *Journal of Atmospheric Sciences*, 39(9):2038–2050, 1982.
- Matthew CG Hall and Dan G Cacuci. Physical interpretation of the adjoint functions for sensitivity analysis of atmospheric models. *Journal of the atmospheric sciences*, 40(10):2537–2546, 1983.
- Matthew CG Hall. Application of adjoint sensitivity theory to an atmospheric general circulation model. *Journal of Atmospheric Sciences*, 43(22):2644–2652, 1986.
- Jochem Marotzke, Ralf Giering, Kate Q Zhang, Detlef Stammer, Chris Hill, and Tong Lee. Construction of the adjoint MIT ocean general circulation model and application to atlantic heat transport sensitivity. *Journal of Geophysical Research: Oceans*, 104(C12):29529–29547, 1999.
- Detlef Stammer, M Balmaseda, P Heimbach, A Köhl, and A Weaver. Ocean data assimilation in support of climate applications: status and perspectives. *Annual review of marine science*, 8:491–518, 2016.
- Nils Gustafsson, Erland Källén, and Sigurdur Thorsteinsson. Sensitivity of forecast errors to initial and lateral boundary conditions. *Tellus A: Dynamic Meteorology and Oceanography*, 50(2):167–185, 1998.
- Detlef Stammer, Armin Köhl, A Vlasenko, I Matei, F Lunkeit, and S Schubert. A pilot climate sensitivity study using the cen coupled adjoint model (cesam). *Journal of Climate*, 31(5):2031–2056, 2018.
- Daniel J Lea, Myles R Allen, and Thomas WN Haine. Sensitivity analysis of the climate of a chaotic system. *Tellus A: Dynamic Meteorology and Oceanography*, 52(5):523–532, 2000.
- Henry DI Abarbanel, Mark Kostuk, and William Whartenby. Data assimilation with regularized nonlinear instabilities. *Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography*, 136(648):769–783, 2010.
- Nozomi Sugiura, Shuhei Masuda, Yosuke Fujii, Masafumi Kamachi, Yoichi Ishikawa, and Toshiyuki Awaji. A framework for interpreting regularized state estimation. *Monthly Weather Review*, 142(1):386–400, 2014.
- Edward N Lorenz. Deterministic nonperiodic flow. *Journal of atmospheric sciences*, 20(2):130–141, 1963.
- Pierre Gauthier. Chaos and quadri-dimensional data assimilation: A study based on the lorenz model. *Tellus A: Dynamic Meteorology and Oceanography*, 44(1):2–17, 1992.
- Robert N Miller, Michael Ghil, and Francois Gauthiez. Advanced data assimilation in strongly nonlinear dynamical systems. *Journal of Atmospheric Sciences*, 51(8):1037–1056, 1994.
- Carlos Pires, Robert Vautard, and Olivier Talagrand. On extending the limits of variational assimilation in nonlinear chaotic systems. *Tellus A*, 48(1):96–121, 1996.

- David J Stensrud and Jian-Wen Bao. Behaviors of variational and nudging assimilation techniques with a chaotic low-order model. *Monthly weather review*, 120(12):3016–3028, 1992.
- Sergey Kravtsov and Anastasios A Tsonis. Lorenz-63 model as a metaphor for transient complexity in climate. *Entropy*, 23(8):951, 2021.
- Xiao-Wei Huai, Jian-Ping Li, Rui-Qiang Ding, Jie Feng, and De-Qiang Liu. Quantifying local predictability of the lorenz system using the nonlinear local lyapunov exponent. *Atmospheric and Oceanic Science Letters*, 10(5):372–378, 2017.
- Shu-Chih Yang, Debra Baker, Hong Li, Katy Cordes, Morgan Huff, Geetika Nagpal, Ena Okereke, Josue Villafane, Eugenia Kalnay, and Gregory S Duane. Data assimilation as synchronization of truth and model: Experiments with the three-variable lorenz system. *Journal of the atmospheric sciences*, 63(9):2340–2354, 2006.
- JD Daron and David A Stainforth. On quantifying the climate of the nonautonomous lorenz-63 model. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 25(4):043103, 2015.
- Ronald M Errico. What is an adjoint model? *Bulletin of the American Meteorological Society*, 78(11):2577–2592, 1997.
- Antonello Pasini and Vinicio Pelino. Can we estimate atmospheric predictability by performance of neural network forecasting? the toy case studies of unforced and forced lorenz models. In *CIMSA. 2005 IEEE International Conference on Computational Intelligence for Measurement Systems and Applications, 2005.*, pages 69–74. IEEE, 2005.
- Pierre Tandeo, Pierre Ailliot, Juan Ruiz, Alexis Hannart, Bertrand Chapron, Anne Cuzol, Valérie Monbet, Robert Easton, and Ronan Fablet. Combining analog method and ensemble data assimilation: application to the lorenz-63 chaotic system. In *Machine Learning and Data Mining Approaches to Climate Science: proceedings of the 4th International Workshop on Climate Informatics*, pages 3–12. Springer, 2015.
- Michael Goodliff, Steven Fletcher, Anton Kliewer, John Forsythe, and Andrew Jones. Detection of non-gaussian behavior using machine learning techniques: A case study on the lorenz 63 model. *Journal of Geophysical Research: Atmospheres*, 125(2):e2019JD031551, 2020.
- Caren Marzban. Variance-based sensitivity analysis: An illustration on the lorenz’63 model. *Monthly Weather Review*, 141(11):4069–4079, 2013.
- Xudong Yin, Bin Wang, Juanjuan Liu, and Xiaowei Tan. Evaluation of conditional non-linear optimal perturbation obtained by an ensemble-based approach using the lorenz-63 model. *Tellus A: Dynamic Meteorology and Oceanography*, 66(1):22773, 2014.
- Yi Juan Du and Ming-Cheng Shiue. Analysis and computation of continuous data assimilation algorithms for lorenz 63 system based on nonlinear nudging techniques. *Journal of Computational and Applied Mathematics*, 386:113246, 2021.
- Maria Cameron and Shuo Yang. Computing the quasipotential for highly dissipative and chaotic sdes an application to stochastic lorenz’63. *Communications in Applied Mathematics and Computational Science*, 14(2):207–246, 2019.
- Vinicio Pelino and Filippo Maimone. Energetics, skeletal dynamics, and long-term predictions on kolmogorov-lorenz systems. *Physical Review E*, 76(4):046214, 2007.
- Morris W. Hirsch, Stephen Smale, and Robert L. Devaney. 14 - the lorenz system. pages 305–328. Academic Press, Boston, third edition edition, 2013. ISBN 978-0-12-382010-5. doi:<https://doi.org/10.1016/B978-0-12-382010-5.00014-2>. URL <https://www.sciencedirect.com/science/article/pii/B9780123820105000142>.
- Hans Dembinski and Piti Ongmongkolkul et al. scikit-hep/iminuit. Dec 2020. doi:10.5281/zenodo.3949207. URL <https://doi.org/10.5281/zenodo.3949207>.
- F. James and M. Roos. Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations. *Comput. Phys. Commun.*, 10:343–367, 1975. doi:10.1016/0010-4655(75)90039-9.
- Olivier Talagrand. Variational assimilation. *Data assimilation: making sense of observations*, pages 41–67, 2010.
- Michael L. Waskom. seaborn: statistical data visualization. *Journal of Open Source Software*, 6(60):3021, 2021. doi:10.21105/joss.03021. URL <https://doi.org/10.21105/joss.03021>.
- The pandas development team. pandas-dev/pandas: Pandas, February 2020. URL <https://doi.org/10.5281/zenodo.3509134>.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.