

Machine Learning for Model Error Inference and Correction

Massimo Bonavita¹ and Patrick Laloyaux¹

¹European Centre for Medium-Range Weather Forecasts, Reading, United Kingdom

Key Points:

- Model error is currently the biggest barrier to improve forecast accuracy in Weather and Climate Prediction
- An Artificial Neural Network (ANN) is trained to simulate the error in the ECMWF operational model
- Using the model error from the ANN inside weak-constraint 4D-Var extends the model error correction to the whole atmospheric column

Corresponding author: Massimo Bonavita, massimo.bonavita@ecmwf.int

Abstract

Model error is one of the main obstacles to improved accuracy and reliability in state-of-the-art analysis and forecasting applications, both in Numerical Weather Prediction (NWP) and in Climate Prediction, conducted with comprehensive high resolution General Circulation Models. In a data assimilation framework, recent advances in the context of weak constraint 4D-Var have shown that it is possible to estimate and correct for a large fraction of systematic model error which develops in the stratosphere over short-range forecast ranges. The recent explosion of interest in Machine Learning/Deep Learning technologies has been driven by their remarkable success in disparate application areas. This raises the question of whether model error estimation and correction in operational NWP and Climate Prediction can also benefit from these techniques. In this work, we aim to start to give an answer to this question. Specifically, we show that Artificial Neural Networks (ANN) can reproduce the main results obtained with weak constraint 4D-Var in the operational configuration of the IFS model of ECMWF. We show that the use of ANN models inside the weak-constraint 4D-Var framework has the potential to extend the applicability of the weak constraint methodology for model error correction to the whole atmospheric column. Finally, we discuss the potential and limitations of the Machine Learning/Deep Learning technologies in the core NWP tasks. In particular, we reconsider the fundamental constraints of a purely data driven approach to forecasting and provide a view on how to best integrate Machine Learning technologies within current data assimilation and forecasting methods.

Plain Language Summary

Model error is one of the main obstacles to improved accuracy and reliability in current Numerical Weather Prediction and in Climate Prediction. Recent advances in Data Assimilation at ECMWF indicate that it is possible to estimate and correct for a large fraction of systematic model error in the stratosphere. The question we address is whether Machine Learning techniques can be used alone and in conjunction with standard Data Assimilation methods to improve on those results. We show that it is indeed possible to extend current Data Assimilation capabilities in operational state-of-the-art forecast systems using Machine Learning tools and we discuss the potential and limitations of future applications of these ideas to other core NWP tasks.

1 Introduction

Numerical Weather Prediction (NWP) can be seen as an initial value problem where a numerical model is integrated in time to forecast the future state of the atmosphere and, increasingly, of the other components of the Earth System that interact with it. Like any other forecasting enterprise, NWP forecasts are affected by errors. In the data assimilation community, forecast errors are traditionally partitioned between errors from evolved erroneous initial conditions and model errors. This distinction is, for example, formalised in the evolution equation of the state error covariance in the Kalman Filter (Kalman, 1960)

$$\mathbf{P}_t^b = \mathbf{M}\mathbf{P}_{t-1}^a\mathbf{M}^T + \mathbf{Q}_t, \quad (1)$$

where the state error covariance matrix of the background forecast state \mathbf{P}_t^b is written as the sum of the evolved analysis errors from the previous analysis update ($\mathbf{M}\mathbf{P}_{t-1}^a\mathbf{M}^T$ where \mathbf{M} is the linear/linearised model) and a zero-mean stochastic model error of covariance \mathbf{Q}_t . This distinction has proved useful, as most data assimilation algorithms in current use can be seen as variations/extensions of the Kalman Filter, but it is also limited by significant assumptions: a) model error is assumed additive; b) model error is assumed to be white in time and c) model error is assumed to be zero-mean.

52 Assumptions a) and b) are somewhat relaxed in operational settings. For example,
 53 at ECMWF the model error parameterisations used in the Ensemble of Data Assimilations
 54 (EDA) to simulate model error evolution are based on a multiplicative ansatz
 55 (Buizza et al., 1999) and spatial model error correlations are cycled from one assimilation
 56 update to the next (Leutbecher et al., 2017). The third assumption (zero-mean errors)
 57 is probably the most important as it effectively makes any Kalman Filter based
 58 data assimilation system blind to the presence of systematic model errors (Dee, 2005).
 59 Note that we here use the term bias in a wider sense than it is typically used in the
 60 meteorological literature: Biases are systematic errors that can vary in space, time and
 61 prevalent meteorological conditions. Thus, we can encounter different model biases in
 62 different locations, at different times of day or year, in different meteorological conditions
 63 and they can be also influenced by systematic errors arising from the interaction
 64 with other components of the Earth System

65 In many data assimilation systems used in operational NWP, model bias is not accounted
 66 for explicitly. Rather, common strategies aim at reducing the impact of model biases
 67 on the performance of the assimilation system. Recognising that the impact of model
 68 biases on the assimilation algorithm mainly comes through the observation-minus-background
 69 (O-B) residuals, these strategies typically involve a combination of: a) debiasing the O-B
 70 residuals, for example through variational bias correction techniques (Auligné et al., 2007),
 71 and b) inflating the estimates of the background forecast errors sampled from an ensemble
 72 data assimilation system run in parallel to the main, higher resolution, analysis system
 73 (Bonavita et al., 2012; Whitaker & Hamill, 2012). Both techniques have proved effective
 74 in improving the performance of the data assimilation and forecast systems, but it is
 75 obvious that they are partial, sub-optimal solutions to the model bias problem. In fact,
 76 bias correction of the O-B residuals implicitly assumes that all the systematic components
 77 of these residuals are due to observation (and observation operator) biases. While this
 78 can be a reasonable working assumption for a large number of satellite radiances, the
 79 fact that we still see systematic O-B errors in largely unbiased observing systems (e.g.,
 80 radiosondes, radio occultation observations from the Global Positioning System, a.k.a.
 81 GPS-RO) in operational data assimilation statistics shows that this is not the case in
 82 general. This effect is also visible in modern reanalyses (Hersbach et al., 2020) where
 83 long-term temperature trends in the stratospheric analysis show discontinuities connected
 84 to the introduction or withdrawal of specific observing systems. Inflating the background
 85 errors is a standard tool in ensemble data assimilation to deal with all components of
 86 the forecast error that are not properly sampled by the assimilation system (Houtekamer
 87 & Zhang, 2016). This technique has also proved effective in reducing the total analysis
 88 mean square error (Raanes et al., 2019), but it is clearly a blunt tool for dealing with
 89 model error. More importantly, any change to the Kalman Gain matrix in a bias-blind
 90 assimilation system will still result in a biased, sub-optimal analysis (Dee, 2005).
 91

Weak Constraint 4D-Var (WC-4DVar) is an extension of 4D-Var which explicitly
 attempts to take model error into account in the solution of the 4D-Var assimilation
 problem (Tremolet, 2006). In the forcing formulation of WC-4DVar implemented at
 ECMWF, this is done by extending the 4D-Var control variable with a model error
 tendency term which is evaluated during the 4D-Var minimisation and used in the
 subsequent first-guess integration to de-bias the model trajectory:

$$\mathbf{x}_k = \mathbf{M}_{k,k-1}(\mathbf{x}_{k-1}) + \boldsymbol{\eta}, \quad k = 1, \dots, N \quad (2)$$

$$J_{WC}(\mathbf{x}_0, \boldsymbol{\eta}) = \frac{1}{2} (\mathbf{x}_0 - \mathbf{x}_0^b)^T \mathbf{B}^{-1} (\mathbf{x}_0 - \mathbf{x}_0^b) + \frac{1}{2} \sum_{k=0}^N \left((H(\mathbf{x}_k) - \mathbf{y}_k)^T \mathbf{R}_k^{-1} (H(\mathbf{x}_k) - \mathbf{y}_k) \right) + \frac{1}{2} (\boldsymbol{\eta} - \boldsymbol{\eta}^b)^T \mathbf{Q}^{-1} (\boldsymbol{\eta} - \boldsymbol{\eta}^b), \quad (3)$$

92 where $\boldsymbol{\eta}$ is the model error forcing (this is kept constant over the assimilation window,
 93 which is the main approximation of the IFS implementation of WC-4DVar), $\boldsymbol{\eta}^b$ is the
 94 prior estimate of the model error forcing and \mathbf{Q} is the model error covariance matrix.

95 While this WC-4DVar formulation has been used at ECMWF since 2009, it is
 96 only very recently (IFS Cycle 47R1, scheduled to become operational from July 2020)
 97 that WC-4DVar has been shown to be effective at correcting stratospheric model biases
 98 (Laloyaux, Bonavita, Dahoui, et al., 2020). The key insight of this revised WC-4DVar
 99 implementation has been to impose scale separation between the error covariance mat-
 100 rices describing the spatial structures of background error \mathbf{B} and of model biases \mathbf{Q}
 101 (see Laloyaux, Bonavita, Dahoui, et al., 2020; Laloyaux, Bonavita, Chrust, & Gürol,
 102 2020, for a detailed explanation). The scale separation allows to successfully de-alias
 103 initial state and model error corrections during the 4D-Var minimisation, and is con-
 104 sistent with a view that model biases represent a type of errors that take place on
 105 larger spatial and longer temporal scales than background errors. It is also apparent
 106 from Equations (2) and (3) that WC-4DVar estimates a model error tendency term
 107 which is then applied as an additional forcing term in the prognostic equations of the
 108 model. Thus, it can be viewed as a data-driven algorithm to estimate (some of) the
 109 missing physical forcing in the model prognostic equations. In other words, WC-4DVar
 110 as described in Equations (2) and (3) is a type of on-line machine learning algorithm.

Machine learning (ML) methods, and more specifically the Deep Learning (DL)
 implementations of ML, have seen a remarkable resurgence over the past decade (Chol-
 let, 2018). This was driven by the unrivalled results obtained through ML/DL tech-
 nologies in a vast range of problems in computer vision, speech recognition, natural
 language processing and translation, among others (Goodfellow et al., 2016). At a fun-
 damental level, most of the successful ML applications in use today implement a type of
 supervised statistical learning where we aim to learn from a dataset of examples (\mathbf{X}, \mathbf{Y})
 a (possibly) nonlinear mapping between “features” $\mathbf{X} = \mathbf{x}^1, \dots, \mathbf{x}^m$ and a correspond-
 ing set of observed targets (“labels”) \mathbf{Y} . This is usually done by assuming a parametric
 model for the conditional distribution of the observations, $p_{\text{model}}(\mathbf{Y}|\mathbf{X}, \boldsymbol{\theta})$, and max-
 imising the likelihood of the model over the empirical data distribution $p_{\text{obs}}(\mathbf{Y}|\mathbf{X})$

$$\boldsymbol{\theta}_{ML} = \arg \max_{\boldsymbol{\theta}} p_{\text{model}}(\mathbf{Y}|\mathbf{X}, \boldsymbol{\theta}). \quad (4)$$

Under standard i.i.d. (independent and identically distributed) conditions for the fea-
 tures and observations distributions, Equation (4) can be transformed in the equivalent
 optimisation problem of maximising the log-likelihood of the predictive model under
 the observed distribution

$$\boldsymbol{\theta}_{ML} = \arg \max_{\boldsymbol{\theta}} \sum_{i=1}^m \log (p_{\text{model}}(\mathbf{y}^i|\mathbf{x}^i, \boldsymbol{\theta})), \quad (5)$$

where i is the index running over the m members of the examples’ dataset. This is
 equivalent to minimising the cross-entropy between the two distribution (Goodfellow et
 al., 2016). For our purposes, we are interested in discovering a statistical regression law
 between model error (or, to be precise, available estimates of model error) and a set of
 predictors (features) to be defined based on physical intuition and experimental results.
 The simplest approach is assuming a linear relationship between (\mathbf{Y}, \mathbf{X}) represented
 by the affine transformation

$$\mathbf{Y} = \mathbf{W}\mathbf{X} + \mathbf{b}. \quad (6)$$

This is equivalent to assuming a Gaussian predictive model of the form

$$p_{\text{model}}(\mathbf{Y}|\mathbf{X}, \boldsymbol{\theta}) = \mathcal{N}(\mathbf{Y}|\mathbf{W}\mathbf{X} + \mathbf{b}, \mathbf{I}), \quad (7)$$

111 where the general set of learnable parameters $\boldsymbol{\theta}$ has been particularised to the sets of
 112 weights \mathbf{W} and bias coefficients \mathbf{b} of a generic neural network. Maximising the log-
 113 likelihood (or, more commonly, minimising the negative log-likelihood) of this model

114 leads to the standard “Normal” equations. Adding constraints on the size of the
 115 regression coefficients matrix \mathbf{W} (known in different communities as Tikhonov regu-
 116 larisation, ridge regression, weight decay) or the sparsity of said matrix (Tibshirani,
 117 1996) can be seen as ways of improving the generalisation properties of the estimator
 118 by trading increased bias for reduced variance.

119 The main limitation of the regression model in Equation (6) lies in its limited
 120 capacity. If the underlying relation between (\mathbf{Y}, \mathbf{X}) is nonlinear, then the maximum
 121 likelihood estimator in Equation (5) will be sub-optimal. In our problem of model
 122 error estimation, it is a priori unclear how much of an issue this is. The WC-4DVar of
 123 Equation (3) is implemented at ECMWF in an incremental formulation, so it can deal
 124 with moderate nonlinearities through repeated re-linearisation steps (Bonavita et al.,
 125 2018). In Deep Learning, the nonlinearity problem is solved by introducing multiple
 126 additional layers in the regression that implement nonlinear transformations between
 127 their inputs and outputs (hidden layers). Even in their simplest algorithmic form,
 128 these nonlinear regressors variously known as Feedforward Neural Networks, Artificial
 129 Neural Networks (ANNs), or MultyLayer Perceptrons (MLPs) have the remarkable
 130 property of being universal function approximators (Cybenko, 1989). Thus, an ANN
 131 of sufficient capacity can theoretically learn any nonlinear mapping to any desired level
 132 of accuracy, given a sufficiently large and representative training dataset.

133 Attempts to use DL techniques to estimate and correct for model errors have
 134 already been documented in the geophysical literature. For example, Watson (2019)
 135 uses ANN to estimate model error tendencies in the Lorenz '96 system and uses them
 136 to correct short and long range forecasts with significant improvements both in fore-
 137 cast skill and model climate statistics. In that work an approximate (coarser) version
 138 of the Lorenz '96 model was still used for prediction, and the ANN was used to “fill
 139 in” the gaps with respect to the high resolution, “true” version of the model. This
 140 idea of hybridising machine learning methods with knowledge-based models is also ex-
 141 ploited in the influential paper of Pathak et al. (2018), where a different ML technique
 142 is employed (Jaeger, 2001), but also very good results are obtained in two low-order
 143 models. In a similar vein, Bolton & Zanna (2019) present an oceanographic applica-
 144 tion of hybrid forecasting using Convolutional Neural Networks (CNN) in a simplified
 145 Ocean model. Again, the goal was to reproduce the effects of unresolved physical
 146 processes in a coarser version of their reference model. More recently, Brajard et al.
 147 (2020) demonstrate a way to combine ML with data assimilation of noisy and partial
 148 observations. In their scheme, DA and ML alternate in producing progressively more
 149 accurate estimates of the state and of the surrogate predictive model. This idea has
 150 been framed into a unifying Bayesian formalism by Bocquet et al. (2020), which allows
 151 to develop approximations and alternative algorithms.

152 In the works described above and, to the Authors’ knowledge, in other recent
 153 relevant literature in the geophysical domain, the application of ML techniques for
 154 model error inference and correction has been studied in the context of low-order,
 155 simplified models. Thus, while the reported results appear encouraging, questions
 156 remain about the extent to which those results are applicable and relevant for high
 157 resolution, operational level data assimilation and forecasting applications. These
 158 applications pose a new set of additional challenges. Firstly, in real world applications
 159 the true state is typically unknown. What is known are incomplete and noisy estimates
 160 of the true state, either directly through observations (which are affected by random
 161 and systematic errors of their own) or indirectly through analyses produced by a data
 162 assimilation system (which are themselves affected by model and observation errors).
 163 Secondly, and possibly more importantly, the dimensions of the analysis and forecast
 164 system in operational NWP are very large. In the current Integrated Forecasting
 165 System (IFS) used at ECMWF, the size of the model state vector is $\mathcal{O}(10^{10})$ and the
 166 size of the analysis control vector is $\mathcal{O}(10^8)$. These numbers are orders of magnitude

167 larger than those for typical low or intermediate complexity models discussed in the
 168 literature and they pose a new set of practical and conceptual questions. The aim
 169 of this work is to give some initial and, at this stage, necessarily tentative answers
 170 to these questions. The main conclusion that we derive from the results presented
 171 in the following is that, while a considerable amount of work still needs to be done,
 172 there is a concrete prospect to successfully integrate ML solutions inside the 4D-Var
 173 machinery of state-of-the-art operational NWP systems like the IFS and, by doing so,
 174 of significantly improving their analysis accuracy and their forecast skill.

175 This paper is organised as follows. In Section 2 we describe the ML methodology
 176 used in this work and the results achieved in terms of the predictive properties of
 177 the ML model. In Section 3 we examine the structure of the model error tendency
 178 predictions of the ML model and compare them to the predictions of the forthcoming
 179 operational version of WC-4DVar. In Section 4 we examine the results of using the
 180 ML-derived model error tendency predictions in cycled 4D-Var experiments, both as
 181 a stand-alone replacement of the WC-4DVar estimates and in conjunction with WC-
 182 4DVar. In Section 5 we discuss these results further in terms of their implications for
 183 our future research and, more generally, in the context of the current research effort
 184 to integrate ML tools in the NWP chain. Conclusions are offered in Section 6.

185 2 Machine Learning Methodology and Results

186 2.1 Set-up of the regression problem

187 The first task in a regression setting is to identify the set of predictors and
 188 predictands that are most relevant (in ML terminology, the examples (\mathbf{X}, \mathbf{Y}) of the
 189 supervised learning problem). As remarked in the introduction, in a real-world set-
 190 ting we do not have access to the true model error predictands (\mathbf{Y}), thus we need to
 191 find suitable substitutes. Generally speaking, the fundamental sources of information
 192 about model error are observations. In a data assimilation context, we can access
 193 this information directly in observation space (through background, O-B, and analy-
 194 sis, O-A, departures) or mediated by an analysis (through analysis increments fields,
 195 A-B). In this work we have chosen the second option, mainly because it is technically
 196 easier to implement, the increments have global, homogeneous coverage and are al-
 197 ready available in the space of the IFS model variables: temperature (t), logarithm
 198 of surface pressure ($\ln sp$), vorticity (vo), divergence (d), specific humidity (q). We
 199 still think, however, that a direct use of observation departures would be a direction
 200 worth pursuing in the future. We remark here that this idea of using timeseries of
 201 analysis increments' fields to estimate the predictable component of model error is not
 202 new in the meteorological literature. For example, one of the algorithms proposed in
 203 Dee (2005) for the correction of model bias in a cycled data assimilation framework
 204 explicitly involves using an online model error estimate based on a running mean over
 205 past analysis increments (e.g., Eqs (43, 44) in Dee, 2005).

206 We can broadly consider two classes of predictors (\mathbf{X}). The first, which we
 207 call “climatological”, comprises predictors that do not depend on the state of the
 208 flow. In this work, our climatological predictors are the set: (latitude, longitude,
 209 time_of_the_day, month). This set of predictors aims at capturing that part of model
 210 error which is related to geographical location, to the diurnal cycle and to the seasonal
 211 cycle. The other class of predictors used in this work are called “state” predictors.
 212 These are predictors that are meant to represent the part of model error linked to
 213 the large scale state of the flow, e.g. oceanic stratocumulus areas, Intertropical Con-
 214 vergence Zone, extra-tropical cyclonic areas, etc. In this first implementation, and
 215 with an operational application in mind, we have chosen the vertical columns of the
 216 background forecast fields of the subset of state variables of the model whose analysis
 217 increments are also used as predictands (i.e., t , $\ln sp$, vo , d , q). This choice is practical,

218 but it can be potentially extended to other state variables and also to the use of state
 219 variables valid at different times. An example of possible avenues for expanding the
 220 set of state predictors is discussed in Section 2.2.

221 Connected with the choice of analysis and forecast fields as (a component of the)
 222 predictors and predictands, is the choice of horizontal spatial resolution for the fields
 223 whose vertical columns are used in the regression (see Figure 1 for a schematic of
 224 the ANN structure). In this work we have selected a resolution in spectral space of
 225 Triangular spectral truncation 21 (T21), which corresponds to an approximate grid
 226 spacing of 900 km on a reduced quadratic Gaussian grid. This choice is motivated by
 227 both practical and fundamental reasons. On the practical side, the coarse resolution
 228 chosen here facilitates the training phase of the ANN as it keeps its memory and com-
 229 putational requirements at a manageable level (in this work we did not have access to
 230 supercomputing resources for the training of the ANN). On the science side, this choice
 231 is motivated by the findings in Laloyaux, Bonavita, Dahoui, et al. (2020); Laloyaux,
 232 Bonavita, Chrust, & Gürol (2020) that only large scale model errors are predictable in
 233 a weak constraint 4D-Var framework. Additionally, there are fundamental arguments
 234 from ergodic dynamical systems theory that suggest that only large scale features of
 235 model error can be learned statistically. We will come back to these arguments in the
 236 discussion in Section 5.

237 2.2 Training the ANN

238 The training of the artificial neural networks (ANN) was conducted on a ded-
 239 icated dual GPU workstation (NVIDIA Quadro GV100) using the open source deep
 240 learning backend Tensorflow (version 1.14.0, Abadi et al., 2016) and its high-level
 241 Python interface Keras. Initial experiments were conducted on an Intel i5 CPU-based
 242 workstation. The relative speed-up in the training phase achieved on the GPU sys-
 243 tem was a factor of approx. 3. The training dataset consisted of operational analysis
 244 increments and background forecasts collected over the whole year of 2018 every 36h
 245 (i.e., using one every three of the available analyses of the ECMWF operational 12-
 246 hourly assimilation cycle) at T21 resolution. The climatological predictors defined in
 247 Section 2.1 were extracted from the grib headers of the state predictors fields. The
 248 validation dataset was composed in the same way using a short two-months period
 249 from 1 January 2019. This dataset was used to get an indication of appropriate hyper-
 250 parameters' values. The test dataset used for verifying the performance of the ANN
 251 was composed by the analysis increments and background forecasts of a three and half
 252 month period starting on the 1 April 2019.

The statistical regressions have been computed separately for the three set of
 predictands and the related state predictors: mass (t , l_{insp}), wind (v_0 , d) and humid-
 ity (q), leading to three separate ANN models. The reason was again to reduce the
 computational and memory cost of the learning phase. This will be reviewed in the
 future, but we do not expect to see large benefits from performing a combined regres-
 sion on the whole set of predictands as the statistical signatures of mass-wind error
 cross-correlations are typically small (Hamrud et al., 2015). We have tested two types
 of regression models. One is a standard linear regression with full connections between
 predictors and predictands. This implements the regression model in Equation 6. The
 number of trainable parameters is the number of input predictors times the number
 of output predictands (dimension of weight matrix \mathbf{W}) plus the number of predic-
 tands (bias vector \mathbf{b}). Considering for example the case of the full neural network
 for the mass variables (t , l_{insp}) with the current IFS number of vertical levels (137)
 this implies a number of trainable parameters equal to $142 \times 138 + 138 = 19734$. The
 number of vertical profiles in the training dataset is $\mathcal{O}(10^6)$, which, as we will see, is
 enough to prevent over-fitting. The other regression model is a nonlinear model where
 a nonlinear transformation is applied element-wise to the output of Equation (6) on

ANNs of increasing depth. The nonlinear transformation is modelled by the REctified Linear Unit (RELU) function, expressed by the function:

$$\text{Relu}(\mathbf{x}) = \max(0, \mathbf{x}) \quad (8)$$

253 The nonlinear transformation in Equation (8) is applied to all layers of the ANN except
 254 the output layer. In the terminology we adopt in the following, `relu_one_layer` is the
 255 fully-connected ANN composed of two layers: an input layer where we use the nonlinear
 256 transform Equation 8 and a linear output layer. Similarly, `relu_two(three)_layers` refer
 257 to the fully connected ANN derived from `relu_one_layer` ANN through the addition of
 258 one (two) hidden nonlinear layer between input and output. In our terminology, the
 259 cardinality refers thus to the number of nonlinear layers in the ANN and not to the
 260 number of hidden layers of the model, as it is more common in the ML literature.

261 The minimiser used in the training is Adam (Kingma & Ba, 2014), which is
 262 an adaptive version of stochastic gradient descent (SGD). We found it to be gener-
 263 ally able to show more monotonous convergence properties and require less tuning of
 264 its hyper-parameters (learning rate and decay rates) than standard SGD and other
 265 adaptive methods available in the Tensorflow toolbox. Regularisation is also an im-
 266 portant aspect of deep learning methodology. In our case we found regularisation to be
 267 only moderately helpful, mainly improving monotonicity of convergence and slightly
 268 improving generalisation power of the model. After some trials, we have settled on
 269 weight decay for the linear regression model and dropout (Srivastava et al., 2014) with
 270 a dropout rate of 20% for the nonlinear models. This relative lack of sensitivity to
 271 regularisation methods is likely due to the relative shallowness of the ANN we have
 272 used and the fact that the size of our training dataset is one to two orders of magnitude
 273 larger than the number of model parameters.

As it is standard in regression settings, the Mean Square Difference between predicted and actual model errors is minimised during training. In order to give a more expressive view of the predictive capability of the ANN we present training results in terms of the coefficient of determination, which represents the proportion of total variance in the training sample that is explained by the model

$$R2 = 1 - \frac{SS_{red}}{SS_{tot}} \quad (9)$$

274 where $SS_{tot} = \sum_{i=1}^m (\mathbf{y}^i - \bar{\mathbf{y}})^2$ is the total Sum of Squares (proportional to the sample
 275 variance) and $SS_{res} = \sum_{i=1}^m (\mathbf{y}^i - \mathbf{f}^i)^2$ is the residual Sum of Squares. In a perfect model
 276 scenario where the model is able to accurately predict every instance of the sampling
 277 dataset, $R2 = 1$. As our model error generating processes are inherently stochastic,
 278 even a perfect model, i.e. a model that makes predictions sampling from the true
 279 error generating distribution, will produce some error, so that $R2$ will in general be
 280 smaller than 1 (this irreducible error is sometimes called Bayes error (Goodfellow et
 281 al., 2016)). Note also that a baseline model that always predicts the average value of
 282 the sampled predictand $\bar{\mathbf{y}}$ has a $R2 = 0$ and models that do worse than this baseline
 283 will have negative $R2$. The $R2$ coefficient has also been used in this work as stopping
 284 criterion in the training to avoid overfitting (i.e., training is stopped when $R2$ has not
 285 increased over the previous 20 epochs). In Figures 2, we present the results of the
 286 ANN training for the three sets of model error tendency predictands: (t, lns_p), (vo,
 287 d) and q. State and climatological predictors are used which means that the column
 288 background forecast fields as well as the metadata (latitude, longitude, time of the
 289 day, month of year) are selected as the input of the neural network. From this set of
 290 training and test results we draw the following conclusions:

- 291 • The mass errors (t, lns_p) are the most predictable: approximately 14% of the
 292 variance of the analysis increments of the test dataset is predicted by the best
 293 ANN model;

- 294 • The wind errors (vo, d) have lower predictability, with the best ANN accounting
295 for $\sim 5\%$ of the variance of the analysis increments of the test dataset;
- 296 • The humidity errors (q) have the lowest predictability. Even the best ANN
297 has a R^2 not significantly larger than zero. This implies that it has no better
298 predictive skill than a baseline model using the mean analysis increment of the
299 training dataset;
- 300 • The predictive power of the ANNs increases going from linear to nonlinear mod-
301 els of increasing depth. The improvements are very large ($\sim 100\%$) going from
302 the linear to the nonlinear regression with one nonlinear layer and saturate with
303 the `relu_three_layers` model. Adding more nonlinear layers does not produce
304 further improvements in test dataset R^2 (not shown).

305 These results confirm that estimating model error in the IFS at the rather coarse
306 scales we are considering here is a mildly nonlinear problem, which can partly explain
307 the success of WC-4DVar in its current configuration (Laloyaux, Bonavita, Dahoui,
308 et al., 2020; Laloyaux, Bonavita, Chrust, & Gürol, 2020). In the current WC-4DVar
309 configuration only mass and (to a lesser extent) wind model errors are estimated and
310 corrected, which also seems a good choice based on the results in Figure 2.

311 An interesting aspect of any regression model is to understand which of the pre-
312 dictors have the most predictive power. We have not looked into this aspect in great
313 detail, but we have trained two separate regression models, one using only climatolog-
314 ical predictors, the other only using state predictors. In Figure 3, we present results
315 for the (t, l_{insp}) predictands. From this plot we can conclude that the state predictors
316 are more informative than climatological predictors ($R^2_{\text{state}} \sim 10\%$, $R^2_{\text{climat}} \sim 8\%$)
317 but both set of predictors contribute independent information to the final regression
318 model ($R^2_{\text{full}} \sim 14\%$).

319 2.3 Training the ANN with an Augmented State Predictor Set

320 There are several possible avenues for extending the set of predictors in our
321 regression problem. One way would be to use the whole set of state variables con-
322 sidered in this work (t, l_{insp}, vo, d, q) as predictors in each regression problem. This
323 would amount to try to leverage the cross-variable correlations in the model error esti-
324 mates. In practice, mass-wind error cross-correlations are found to be small on average
325 ($\sim 10\%$, Hamrud et al., 2015), so a considerably larger training dataset would likely
326 be required to estimate these small covariances. In a similar vein, the set of state pre-
327 dictors could be augmented to include any other state variables that could potentially
328 co-vary with the predictands, for example vertical velocity, precipitation rate, liquid
329 water content, etc. Another option which we have started to investigate, is to extend
330 the set of state predictors in time. The intuition here is to try to extract information
331 on current errors not only from the forecast state valid at the same time but on its
332 recent evolution. A simple and relatively inexpensive way of achieving this result is
333 to augment the set of state predictors, which are 12-hour background forecast fields,
334 with the analysis fields from which they were forecasted. This implies an approximate
335 doubling of the size of the predictors (e.g., for (t, l_{insp}) from 142 to 280). An example
336 of results from the training of ANNs using this flavour of augmented predictor set
337 is shown in Figure 4. This plot suggests that the ANN trained on the augmented
338 predictor set has more predictive power than the ANN trained on the standard set
339 ($R^2 \sim 15\%$ vs 14%). A similar improvement is seen for the wind error predictands
340 ($R^2 \sim 6\%$ vs 5%), while no improvement is visible in the humidity ANN results (not
341 shown).

3 Predicting Model Error with Artificial Neural Networks

In this section we present a series of diagnostic results in order to give a first impression of what the model error tendencies predicted by the trained ANNs look like and how they compare with those estimated by WC-4DVar and also visible in observation departures. The plots presented in the following refer to one week of data but are indicative of the ANN results over the test period. Results shown here refer to `relu_three_layers` ANNs trained with the standard set of climatological and state predictors, not with the extended set described in Section 2.2. For comparison, we show weak-constraint 4D-Var model error estimates for an experiment run over the same period and initialised from the operational IFS.

In Figures 5a and 5b, we present a weekly average of the temperature model error tendencies estimated by the ANN (left) and by WC-4DVar (right). To be consistent with current IFS WC-4DVar practice, the ANN model error tendencies are derived from the ANN predicted analysis increments divided by the length (in hours) of the assimilation window. The current version of WC-4DVar is not active below model level 60 (approx. 100 hPa); the ANN is active everywhere and in the troposphere (below 100 hPa) it shows patterns of warm and cold error layers with larger intensities in the boundary layer. In the layer between model level 60 to 30 (approx. 100 to 10 hPa) both WC-4DVar and the ANN show a general tendency to warm the atmosphere, more noticeably in the tropics. This is consistent with the cold model bias seen in radiosonde temperature measurements in the lower troposphere (see below). Above model level 30 both ANN and WC-4DVar show a generally negative (cooling) tendency, which is also consistent with the warm model bias with respect to radiosonde measurements in this layer of the model atmosphere. In the vorticity and divergence model error plots (Figures 5c-f) the corrections estimated by the ANN and WC-4DVar are smaller and more homogeneous. It is difficult to see clear physically interpretable patterns apart from a general tendency to decrease both parameters and the hint of a coherent negative-positive-negative divergence pattern in the tropical troposphere (Figure 5e). This last pattern appears to be a robust feature of the ANN regression (it is present in all the weekly averages computed over the test period, not shown), and thus it likely points to local issues in the current parameterised convection scheme, the data assimilation system, or both.

To obtain further insight in the spatial variability of the model error tendencies predicted by the ANN, we present in Figure 6 the weekly averaged plots of temperature model error tendencies from the ANN and WC-4DVar at model level 24 (approx. 5 hPa) and 50 (approx. 50 hPa). While the globally averaged values agree, the spatial structures are different: in particular, the ANN tendencies are larger scale and less intense than those of WC-4DVar. This is to be expected, as the WC-4DVar estimates are online estimates, and thus more sensitive to existing flow conditions than the ANN estimates.

It is also interesting to see the geographical distribution of the ANN-derived model error tendencies for model levels where the current WC-4DVar does not produce an estimate (i.e., below 100 hPa). An example is given in Figure 7 where the ANN estimates are shown for model levels close to 100 (a), 500 (b) and 850 hPa (c). From these plots one can see clear signatures of errors connected to downstream flow from the main mountain ranges (Rockies (b), Himalayas (a)), to convectively active areas (Amazons (b), Maritime Continent (b)), to storm tracks regions (south hemisphere storm tracks (b)) and to marine stratocumulus areas off the western seabords of continental land masses (c). These features all point to predictable, flow-dependent errors in the model which the ANN regression tries to correct and they can be viewed as an additional model diagnostic tool.

393 Surface pressure is another component of the state vector whose model errors
 394 are not estimated by the current version of WC-4DVar but is an output of the ANN
 395 regression. Two examples of the ANN estimates of surface pressure errors are presented
 396 in Figure 8, one from a weekly average in July 2019 (a), the other from a week in
 397 November 2019 (b). It is interesting that while some features of the estimated surface
 398 pressure error appear stationary (e.g., oceanic western boundaries, convective areas
 399 like the maritime continent and the Amazons, etc.), seasonal variability is visible in
 400 other parts of the Globe where the underlying meteorology is significantly different
 401 and more sensitive to the seasonal cycle (e.g., Antarctic region, Siberian landmass).
 402 Again, each of these signals can provide potentially valuable diagnostic indications
 403 as they point to systematic problems in the short-range forecast and/or the use of
 404 observations in those areas. How these diagnostics produced by the applications of
 405 ANN and WC-4DVar compare with those derived from more traditional approaches
 406 based on the accumulation of data assimilation statistics (Rodwell & Jung, 2008) is
 407 an interesting line of research that we defer to future work.

408 4 Testing ANN in the IFS 4D-Var

409 While the analysis of model error predictions produced by an ANN can pro-
 410 vide useful diagnostics of model error patterns and hint at their underlying drivers, a
 411 more stringent test of the ANN potential is whether its model error predictions can
 412 be gainfully used in a data assimilation context. We recall from Equations (2) and (3)
 413 that the current ECMWF WC-4DVar works by estimating a constant in time model
 414 error correction during the 4D-Var minimisation and then using that correction as a
 415 model forcing during the successive first guess trajectory integration. As explained in
 416 Section 2, the ANN used in this study has been trained to predict the analysis incre-
 417 ments (A-B) of the operational IFS, whose systematic component can be viewed as an
 418 estimate of the cumulated model error over the 12 hour integration from one analysis
 419 update to the next. By dividing this quantity by the number of timesteps used in
 420 the 12 hour integration we obtain an estimate of the model error tendencies, under
 421 the same assumptions of time constancy over the assimilation window length of the
 422 ECMWF operational WC-4DVar. There are at least two main ways in which one can
 423 use the offline model error tendencies produced by the ANN. The first option, called
 424 NN_SC in the following, is based on a strong-constraint 4D-Var where the model first
 425 guess trajectories are corrected by the ANN model error tendencies (see Algorithm 1).
 426 The second option, called NN_WC, is based on the weak-constraint 4D-Var presented
 427 in the Equations (2) and (3) where the model forcing $\boldsymbol{\eta}$ is initialised by the ANN
 428 tendency (see Algorithm 2). The model forcing that comes out of the minimisation
 429 of Equation (3) which contains the ANN tendency updated by the weak-constraint
 430 4D-Var minimisation is not carried forward in time. In the successive assimilation
 431 window, the model forcing $\boldsymbol{\eta}$ is initialised with the corresponding ANN tendency valid
 432 12 hours later. It is also important to note that since current WC-4DVar is active only
 433 above 100 hPa, the model error tendencies below 100 hPa derive entirely from the
 434 ANN estimates in both NN_SC and NN_WC experiments. The results shown in the
 435 following come from assimilation and forecast experiments conducted with the latest
 436 available IFS cycle at the time of writing (Cycle 47R1, May 2020) and at the opera-
 437 tional configuration for both the 4D-Var analysis step and the forecast step (TC01279
 438 spectral truncation, approx. 9 km grid spacing), over the period 16-07-2019 to 24-08-
 439 2019. The two experiments making use of the ANN model error estimates (NN_SC
 440 and NN_WC) are compared with an experiment using the standard operational weak
 441 constraint configuration (denoted “WC”) and a strong-constraint 4D-Var used as a
 442 baseline.

Algorithm 1 NN_SC

Loop over the assimilation cycles from 16-07-2019 to 24-08-2019 :

Loop over the physical variables (T,lnsp), (vo,d) and q :

Compute the model error forcing with the trained ANN at valid times
 Concatenate the outputs in a vector η^b
 Minimise the strong constrain 4D-Var

$$J_{SC}(\mathbf{x}_0) = \frac{1}{2} (\mathbf{x}_0 - \mathbf{x}_0^b)^T \mathbf{B}^{-1} (\mathbf{x}_0 - \mathbf{x}_0^b) + \frac{1}{2} \sum_{k=0}^N (\mathcal{H}_k(\mathbf{x}_k) - \mathbf{y}_k)^T \mathbf{R}_k^{-1} (\mathcal{H}_k(\mathbf{x}_k) - \mathbf{y}_k)$$

where the model trajectories are computed as

$$\mathbf{x}_k = \mathcal{M}_{k,k-1}(\mathbf{x}_{k-1}) + \boldsymbol{\eta}^b \quad \text{for} \quad k = 1, \dots, N$$

Algorithm 2 NN_WC

Loop over the assimilation cycles from 16-07-2019 to 24-08-2019 :

Loop over the physical variables (T,lnsp), (vo,d) and q :

Compute the model error forcing with the trained ANN at valid times
 Concatenate the outputs in a vector η^b
 Minimise the weak-constrained 4D-Var

$$J_{WC}(\mathbf{x}_0, \boldsymbol{\eta}) = \frac{1}{2} (\mathbf{x}_0 - \mathbf{x}_0^b)^T \mathbf{B}^{-1} (\mathbf{x}_0 - \mathbf{x}_0^b) + \frac{1}{2} \sum_{k=0}^N (\mathcal{H}_k(\mathbf{x}_k) - \mathbf{y}_k)^T \mathbf{R}_k^{-1} (\mathcal{H}_k(\mathbf{x}_k) - \mathbf{y}_k) + \frac{1}{2} (\boldsymbol{\eta} - \boldsymbol{\eta}^b)^T \mathbf{Q}^{-1} (\boldsymbol{\eta} - \boldsymbol{\eta}^b)$$

where the model trajectories are computed as

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

443

4.1 Time evolution of the model error estimates

444

445

446

447

448

449

450

451

452

453

454

455

Figure 9a shows the timeseries of the global mean model error correction estimated with weak-constraint 4D-Var between 15 and 24 August 2019. The model error is initialised at the beginning of the experiment from operations and is then cycled over the 12-hour assimilation windows. Weak constraint 4D-Var corrects the warm bias in the upper stratosphere and the cold bias in the mid/lower stratosphere. It correctly captures the transition layer (20 to 10 hPa) where the model bias changes from cold to warm. Figures 9b and 9c present the same diagnostic for NN_SC and NN_WC respectively. The transition level between the cold and the warm bias layers is estimated at the same pressure level as in weak-constraint 4D-Var. The main difference is in the upper stratosphere where the neural network produces a positive correction around 2hPa as this signal is present in the analysis increments used to train the neural network.

456

4.2 Evaluation of mean errors

457

458

459

460

461

462

463

464

465

466

467

468

469

470

471

472

473

474

475

476

477

478

479

480

One of the main successes of the new WC-4DVar introduced in IFS Cycle 47R1 has been the drastic reduction (up to 50%) of temperature biases in the ECMWF stratospheric analyses (Laloyaux, Bonavita, Dahoui, et al., 2020). The first set of diagnostics presented in Figure 10 aims at understanding if and in what measure these results can be achieved using an ANN (alone or as a first guess). The general impression from these plots is that the two ANN-driven WC-4DVar experiments produce similar results to one another and manage to broadly replicate the effects of current IFS WC-4DVar (though a closer look points to a relative better behaviour of WC-4DVar at certain heights/pressure levels, e.g. radiosonde temperature at 5 hPa). Apart from temperatures, current WC-4DVar also corrects for wind model errors in the stratosphere. The mean wind observation departures presented in Figure 10c confirm that WC-4DVar is able to correct for systematic wind model errors above 50 hPa, where it is fully active, while results are mixed in the transition layer between 100 and 50 hPa. The two ANN driven experiments show similar results in the layer above 50 hPa, but are also able to substantially reduce model biases in the atmospheric column down to approx. 700 hPa. In Figure 10e we present results for surface observations. The only significant differences are seen here in the departures for surface pressure observations. For these pressure observations, background departures for strong and weak constraint 4D-Var are almost identical, which is not surprising since the mass adjustments that current WC-4DVar performs in the stratosphere have small impact on the total column weight. On the other hand, the ANN driven WC-4DVar experiments show an approximate halving of the surface pressure observed biases, which confirms that the ANN derived corrections are also effective in the troposphere and for the surface pressure field.

481

4.3 Evaluation of random errors

482

483

484

485

486

487

488

489

490

491

492

493

The ability of WC-4DVar and its ANN-driven variants to effectively debias the first guess trajectories should in principle improve the successive analyses and forecasts by allowing the assimilation to make better use of the available observations. To investigate this aspect of the assimilation system performance we start by showing in Figure 11 the normalised standard deviation of analysis (O-A) and background (O-B) departures for AMSU-A radiances (a) and radiosonde temperature observations (b). The AMSU-A radiometer on board multiple operational and research meteorological satellites is a microwave radiometer whose channels are sensitive to deep layers of the atmosphere. The channels used in the experiments have weighting functions that peak in the troposphere (ch. 5 to 8) and the stratosphere (ch. 9 to 14). From the AMSU-A plot it is apparent that current WC-4DVar is more effective in the stratosphere and upper troposphere than either of the ANN-driven 4DVars, and equivalent

494 to the NN_WC set-up in the middle to lower troposphere (the NN_SC setup performs
 495 consistently worse). The picture is more nuanced for radiosonde observations, where
 496 the ANN 4DVars appear to perform significantly better than current weak and strong
 497 constraint 4DVar in the lower troposphere and comparably above. The results from
 498 other independent observing systems sensitive to atmospheric temperature not shown
 499 here (e.g., hyperspectral sounders) appear to confirm the advantage of the current WC-
 500 4DVar in the middle and lower stratosphere and suggest an improved performance of
 501 NN_WC in the middle and lower troposphere. We note, additionally, that the results
 502 from the stratospheric-peaking satellite radiances need to be taken with some caution,
 503 as they are influenced by the evolution of the corresponding bias correction coeffi-
 504 cients. For the experiments reported in this work the bias correction coefficients have
 505 been initialised by a long-running pre-operational WC-4DVar, and thus are likely to
 506 be sub-optimal for the other configurations over most or all of the test period. In fact,
 507 no degradations are apparent for the ANN driven experiments when the verification
 508 is conducted against non-bias corrected observing systems (radiosondes, GPS-RO).
 509 For conventional wind observations (Figure 11c) there is hardly any significant dif-
 510 ference among the four experiments. For satellite atmospheric motion vector winds
 511 (Figure 11d) the differences are clearer: the two ANN driven experiments improve
 512 results over standard weak and strong constraint 4D-Var in the boundary layer and
 513 in the Upper Troposphere Lower Stratosphere (UTLS) layer. This is consistent with
 514 results seen in the mean wind errors plots (Figure 10c 10d). The statistics for random
 515 errors affecting humidity sensitive observations are not presented as they generally
 516 do not show appreciable differences among the experiments. This is to be expected
 517 because neither the current WC-4DVar nor the two ANN-driven versions apply any
 518 additional model forcing for humidity, thus any change in behaviour would only be an
 519 indirect effect of changes to temperature/wind evolution.

520 On the other hand, there are significant changes in the diagnosed background
 521 error standard deviations for surface pressure observations (Figure 11e). Consistently
 522 with the results for the mean errors, the NN_WC version of ANN-driven 4D-Var sig-
 523 nificantly reduces random errors for surface pressure observations with respect to the
 524 reference strong and weak constraint 4D-Var. It is also to be noticed that NN_WC
 525 uses 6% more Dribu (Drifting Buoys) observations than the reference SC, due to the
 526 fact that more observations pass first guess quality control checks (This is what the
 527 red caption in the plot refers to). This is an indication that the surface pressure model
 528 error correction is particularly useful in the Ocean, where the observing system is
 529 significantly sparser than over land and thus model errors play a bigger role in the
 530 forecast error budget.

531 4.4 Evaluation of forecast skill

532 Here we concentrate on two aspects of the forecast performance of the ANN
 533 driven WC-4DVar experiments. The first aspect is whether they are able to replicate
 534 the improvements in stratospheric temperature reductions of forecast bias produced
 535 by the recent version of WC-4DVar (Laloyaux, Bonavita, Dahoui, et al., 2020). 10-
 536 day forecasts are initialised using the analysis from strong-constraint 4D-Var, weak-
 537 constraint 4D-Var, NN_SC and NN_WC between 10 and 24 August 2019. The model
 538 used to compute these forecasts is not corrected by any forcing estimated in weak-
 539 constraint 4D-Var or neural networks. Given the possible problems of correlated anal-
 540 ysis and forecast errors that a standard own-analysis verification is likely to cause,
 541 we present forecast verification results against independent GPS-RO derived temper-
 542 ature profiles. Figure 12 shows the difference in temperature forecast RMSE after 72
 543 hours between forecasts initialised by NN_SC and strong-constraint 4D-Var (a) and
 544 by NN_WC and strong constraint 4D-Var (b) and by weak-constraint 4D-Var and
 545 strong-constraint 4D-Var (c). The improvements obtained by weak-constraint 4D-Var
 546 are replicated by the two neural networks to a large extent. Degradations observed

at different pressure levels and latitudes are mainly not statistically significant. Comparing the two neural network approaches, one can see that weak-constraint 4D-Var used in NN_WC mitigates the degradation observed in NN_SC. Longer experiments are currently running to improve the statistical robustness of these results.

The other main question that we would like to answer is whether the introduction of model error forcing in the troposphere in the ANN driven 4DVar experiments is beneficial or not in terms of synoptic performance of the forecast. This is of particular interest in light of the fact that previous attempts to extend the current WC-4DVar formulation to the full atmospheric column resulted in significant degradations of various aspects of tropospheric forecast skill for the reasons explained in Laloyaux, Bonavita, Dahoui, et al. (2020). In Figure 13, we present two standard measures of synoptic performance for the mass field. Both 500 hPa geopotential (a-b) and Mean Sea Level Pressure (c-e) forecast errors for either of the ANN configurations appear slightly better than the reference strong and weak constraint 4D-Var, though statistical significance is only reached sporadically in the relatively short test period used here. Forecast performance for the wind field (not shown) is similar to that of the mass field presented earlier: No significant degradation is apparent, and some localised improvements consistent with the positive indications coming from the observation space assimilation diagnostics presented in Section 4.1 and 4.2 are also visible.

5 Discussion and research perspectives

The work presented in this paper and recent developments in 4DVar methodology (Laloyaux, Bonavita, Dahoui, et al., 2020; Laloyaux, Bonavita, Chrust, & Gürol, 2020) are based on the idea that an effective strategy to deal with model error in NWP is to partition it in two components: a) a stochastic, small scale (temporally and spatially) component and b) a predictable component active on larger and longer spatial/temporal scales. The random component of model error is typically represented with physically-based model error simulation models (Leutbecher et al., 2017) which are derived from an understanding of the approximations done in the development of the forecast model and an attempt to sample from those sources of uncertainties. These stochastic models of model error are then applied both in an ensemble data assimilation framework (Bonavita et al., 2012; Bowler, 2017) and ensemble forecast mode (Leutbecher et al., 2017). In ensemble data assimilation, their net effect is to produce a flow-dependent increase of ensemble spread and an associated improvement in the ensemble forecast reliability budget, which is usually under-dispersive (Houtekamer & Zhang, 2016; Bonavita et al., 2012; Bowler, 2017). These model error parameterisations are targeted at improving the ensemble estimate of the second order moment of the forecast error pdf. They might also be able to indirectly affect the ensemble forecast mean due to nonlinear effects arising during the model integrations, but these effects are small in a data assimilation cycling context (and often explicitly discarded through re-centring techniques). The second component of model error, which we have considered in this work, is the large scale error that evolves slowly over the time scale of the assimilation window length. We posit that this error is predictable, i.e. we can estimate the first moment of its distribution through statistical estimation techniques. Weak constraint 4DVar is an on-line example of a statistical estimation technique for dealing with the systematic errors of the model. The machine learning models described in this paper are examples of off-line statistical models aimed at achieving similar goals. A variety of hybrid configurations with a combination of on-line and off-line estimators are also possible: the WC-4DVar configuration where the ANN model error estimate is used as first guess and background for the WC-4DVar minimisation (NN_WC) is just an initial, proof-of-concept attempt. Similarly to the stochastic model error parameterisations, the use of WC-4DVar or its ML hybrids can improve reliability in an ensemble assimilation and forecasting system, but through

599 a different mechanism, namely reducing the total error budget by reducing/removing
600 the systematic error components.

601 5.1 Research perspectives

602 The preliminary results presented in the previous section show that combining
603 ANN models and WC-4DVar holds promise of improving on each technique used in
604 isolation. In particular, it appears that a hybrid ANN-WC-4Dvar setup can be con-
605 figured to effectively reduce model error throughout the atmospheric column and not
606 only in the stratosphere. The specific configuration of this hybrid ANN-WC-4DVar
607 is being currently investigated and the findings of this research will be reported in a
608 follow-up paper. Other aspects of the methodology presented in this work can be fur-
609 ther improved. Of fundamental importance is the choice of predictands and predictors
610 for the model error regression problem. In terms of model error predictands, we have
611 chosen to use analysis increments in state space. This idea is not new in NWP (Dee,
612 2005) and stems from the somewhat obvious consideration that only observations can
613 (directly or indirectly) tell us something useful about model error with respect to the
614 real atmosphere. This idea has been more recently revived in an ensemble data assim-
615 ilation and forecasting context by Bowler (2017), following earlier work by (Piccolo
616 & Cullen, 2015). With respect to these later works, our application differs in two
617 important aspects: a) it can also be applied in a deterministic, perfect model assim-
618 ilation and forecasting system, and b) its estimates are derived from flow-dependent
619 regressions. This second property is important not only because the flow-dependent
620 component adds more predictive power than that coming from climatological predic-
621 tors (Section 2), but because it opens the perspective of using ANN models as an
622 online, flow-dependent model error correction forcing term. This will be potentially
623 interesting for improving predictions at longer forecast ranges than those considered
624 in this work, as the accuracy and reliability of ensemble forecast predictions at long
625 forecast scales are notoriously affected by the model systematic errors. At the current
626 stage of research it is unclear whether this application of ANN, WC-4DVar or their
627 hybrids will be practically successful. This is also in view of the complex and typically
628 non-linear model error interactions that arise between the various components of a
629 coupled Earth System model during extended integrations. We note however, that
630 recent results in both medium range NWP (Laloyaux, Bonavita, Dahoui, et al., 2020)
631 and seasonal prediction (Ham et al., 2019) have already shown that the introduction of
632 pure or hybrid ML/DL models can lead to significant improvements in specific aspects
633 of forecast performance.

634 The choice of analysis increments as model error predictands was mainly dic-
635 tated by reasons of convenience and practicality. Another option is to directly use
636 observation departures. This would have the advantage of avoiding another source of
637 errors from the data assimilation system. On the other hand, one is limited to the
638 relatively small subset of observations which are thought not to be affected by signifi-
639 cant systematic errors themselves (e.g., radiosondes, GPS-RO), issues connected with
640 their spatial and temporal homogeneity and more complex relations of the observed
641 to the state variables. Still, we believe these issues can be addressed to some extent
642 and a separate research effort is ongoing in this direction at ECMWF. Another area
643 of development regards the type and choice of predictors used for our regression prob-
644 lem. As shown in Section 2.2 a judicious choice of additional predictors can further
645 improve the predictive power of the ANN model and, likely, its impact on the IFS anal-
646 yses and forecasts. Additional predictors can be envisaged which exploit additional
647 sources of predictability of the atmospheric flow, especially those coming from fixed
648 or slow-evolving boundary conditions (e.g., orography, land use, sea surface tempera-
649 tures, etc.). In the context of choosing appropriate sets of predictors and predictands,
650 their geometry and spatial resolution, the issues connected to overfitting and the so-

651 called “curse of dimensionality” become prominent. We discuss them in the following
 652 sub-section.

653 5.2 Statistical regression and the curse of dimensionality

As a type of statistical learning, machine learning is exposed to the problem of the “curse of dimensionality”. Loosely speaking, this means that for systems with a large number of degrees of freedom, the number of available training examples will always be much smaller than the number of possible configurations in state space. Standard results from ergodic theory of dynamical systems (Cecconi et al., 2012) show that, for a dissipative nonlinear dynamical system like the atmosphere (or a state of the art NWP model), the minimum length M of the time series of past observed states (i.e., the size of the training dataset) necessary to find an analogue of the current state within an error distance measure ϵ has a scaling law of the form:

$$M \sim \left(\frac{L}{\epsilon}\right)^{D_a} \quad (10)$$

654 where L represents a measure of the variability of the system and D_a is the
 655 effective dimension of the system attractor, which can be a non integer number (i.e.,
 656 a strange attractor). The exponential dependence of the training dataset size on the
 657 effective attractor dimension makes a fully statistical approach to forecasting unfeasible
 658 (Van den Dool, 1994). For machine learning applications to NWP and climate they
 659 indicate that an acritical application of ML tools is not likely to give good results
 660 unless effective mitigating strategies are put in place. There are at least two possible
 661 avenues to combat the curse of dimensionality. One is obviously to try to reduce
 662 the size of the regression or classification problem. This has motivated our choice to
 663 deal with the model error estimation problem, which can be framed as the problem
 664 of trying to identify a residual model which fills the gap between the actual forecast
 665 model and reality (cp. Eq. 2). It is reasonable to assume that the attractor dimension
 666 of the residual model is much smaller than that of the full model, as suggested, for
 667 example, in studies using reduced order models (Watson, 2019). The applicability
 668 of this assumption is further strengthened by the choice of using a coarse spatial
 669 resolution for the training dataset. This limits the modes of variability allowed in
 670 the regression and allows to train the ML model on a relatively small dataset and
 671 achieve good generalisation performance. The other standard tool to beat the curse
 672 of dimensionality is to use prior knowledge about the data generating distribution to
 673 suitably restrict the choice of the model space in which the machine learning algorithm
 674 is allowed to search for solutions (this is called the “hypothesis space” in machine
 675 learning literature). This is where expert knowledge of the problem at hand becomes
 676 valuable as there is no machine learning algorithm that is universally better on all
 677 possible tasks (The so-called “No free lunch” theorem Wolpert, 1996). With this in
 678 mind, we have chosen the avenue of training a fully connected ANN over atmospheric
 679 columns of predictor and predictand examples. The insight here is that it is important
 680 for the regression model to learn vertically balanced increments to avoid introducing
 681 spurious unphysical instabilities in the model evolution. This is also consistent with
 682 the way standard NWP and Climate Prediction models are currently formulated: the
 683 equations governing the model physical tendencies are typically formulated over model
 684 columns. Other approaches are possible, e.g. use Convolutional Neural Networks of
 685 the type that are currently popular in image recognition applications to try to learn
 686 spatial patterns on model levels. We leave this for future investigation. An additional
 687 advantage of our choice of predictors and predictands geometry is that it helps to
 688 drastically reduce the number of learnable parameters of the ANN model and thus the
 689 risk of overfitting the training dataset.

6 Conclusions

Machine Learning and Deep Learning technologies have been applied successfully in many disparate fields. These remarkable success stories have in the past few years generated interest in the NWP and climate communities to understand whether there is scope to apply ML/DL techniques in their respective fields. However, while a number of visionary, speculative papers have been published explaining the case for the application of ML/DL to NWP and climate, and an even greater number have investigated the use of ML/DL techniques in a variety of low-order models, very little work seems yet to have been undertaken to apply ML/DL methods to state-of-the-art, high resolution global circulation models such as those used in operational global NWP and climate. The work presented here aims at starting to fill this gap. The results presented in this paper show a first application of ML/DL tools to the problem of model error estimation and correction in a data assimilation context. Building on recent results obtained in a weak constraint 4D-Var framework (Laloyaux, Bonavita, Dahoui, et al., 2020; Laloyaux, Bonavita, Chrust, & Gürol, 2020) we show that the use of ANN-derived model error forecasts potentially allows to extend the benefits of the weak constraint formulation of 4D-Var to the troposphere, which had been an intractable problem since the introduction of WC-4DVar at ECMWF more than ten years ago. While these results need to be validated over longer testing periods and the technical infrastructure is not yet fully in place for reliable operational use, we believe these results to be promising enough to warrant actively pursuing this line of research further.

From the vantage point of data assimilation in the geosciences, ML/DL do not introduce completely new or revolutionary ideas. In fact, ML/DL techniques and their theoretical underpinnings have much in common with the standard toolbox of variational data assimilation, though these similarities are partly obfuscated by the different nomenclature (Geer, 2020). What the most recent ML/DL wave of interest has brought about is the availability of a set of powerful, easy to use, open source libraries which greatly facilitate the application of these techniques to disparate fields; and a renewed awareness of the effectiveness of these techniques in many different contexts. These newly available tools are undoubtedly helping the adoption of ML/DL techniques in the NWP community beyond already well-established areas such as e.g. NWP output post-processing and nowcasting (McGovern et al., 2017; Gagne et al., 2017; Rasp & Lerch, 2018). As discussed in Section 5, enthusiasm for adopting these new tools in core NWP tasks needs to be tempered by a careful appreciation of their fundamental limitations. Even with the huge increase in modern computational resources and the size of available training datasets, it is unlikely, for example, that a fully ML/DL-based forecast model will supersede the current knowledge-based forecast models. On the other hand, it is possible and even probable that knowledge-based models of the not-too-distant future will integrate ML/DL components for reasons of computational efficiency and possibly improved performance. At the same time, it is likely that ML/DL tools for model error estimation and correction like those presented in this work will play a major role in a variety of data assimilation and forecasting applications.

Acknowledgments

We gratefully acknowledge constructive discussions with a number of colleagues at ECMWF and beyond. In particular, Marc Bocquet, Alban Farchi, Peter Dueben, Alan Geer, Peter Bauer, Nils Wedi and Elias Hölm have reviewed an earlier version of the manuscript and provided many useful suggestions for its improvement. We also wish to express our gratitude to our ECMWF colleague Xavier Abellan for his patient and precious technical support during the course of this project. Data availability statement: The input and output data of the experiments described in the paper is

742 freely available for research purposes from ECMWF and can be requested following
 743 the procedures described in <https://www.ecmwf.int/en/forecasts/datasets>

744 References

- 745 Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., ... Zheng, X.
 746 (2016). Tensorflow: A system for large-scale machine learning. In *12th usenix*
 747 *symposium on operating systems design and implementation (osdi 16)* (pp. 265–
 748 283). Retrieved from [https://www.usenix.org/system/files/conference/
 749 osdi16/osdi16-abadi.pdf](https://www.usenix.org/system/files/conference/osdi16/osdi16-abadi.pdf)
- 750 Auligné, T., McNally, A. P., & Dee, D. P. (2007). Adaptive bias correction for satel-
 751 lite data in a numerical weather prediction system. *Quarterly Journal of the Royal*
 752 *Meteorological Society*, *133*(624), 631-642.
- 753 Bocquet, M., Brajard, J., Carrassi, A., & Bertino, L. (2020). Bayesian inference of
 754 chaotic dynamics by merging data assimilation, machine learning and expectation-
 755 maximization. *Foundations of Data Science*, *2*, 55–80.
- 756 Bolton, T., & Zanna, L. (2019). Applications of deep learning to ocean data infer-
 757 ence and subgrid parameterization. *Journal of Advances in Modeling Earth Sys-*
 758 *tems*, *11*(1), 376-399.
- 759 Bonavita, M., Isaksen, L., & Holm, E. (2012). On the use of EDA background error
 760 variances in the ecmwf 4D-Var. *Quarterly Journal of the Royal Meteorological So-*
 761 *cety*, *138*, 1540–1559.
- 762 Bonavita, M., Lean, P., & Holm, E. (2018). Nonlinear effects in 4d-var. *Nonlinear*
 763 *Processes in Geophysics*, *25*(3), 713–729.
- 764 Bowler, N. E. (2017). On the diagnosis of model error statistics using weak-
 765 constraint data assimilation. *Quarterly Journal of the Royal Meteorological*
 766 *Society*, *143*(705), 1916-1928.
- 767 Brajard, J., Carrassi, A., Bocquet, M., & Bertino, L. (2020). Combining data assim-
 768 ilation and machine learning to emulate a dynamical model from sparse and noisy
 769 observations: a case study with the lorenz 96 model. *Journal of Computational*
 770 *Science*, 101–171.
- 771 Buizza, R., Milleer, M., & Palmer, T. N. (1999). Stochastic representation of model
 772 uncertainties in the ecmwf ensemble prediction system. *Quarterly Journal of the*
 773 *Royal Meteorological Society*, *125*(560), 2887-2908.
- 774 Cecconi, F. M., Cencini, M., Falcioni, & Vulpiani, A. (2012). The prediction of fu-
 775 ture from the past: an old problem from a modern perspective. *American Journal*
 776 *of Physics*, *80*(11), 1001-1008.
- 777 Chollet, F. (2018). *Deep learning with python*. Manning Publicationa, New-York.
- 778 Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Con-*
 779 *trol Signal Systems*, *2*, 303–314.
- 780 Dee, D. P. (2005). Bias and data assimilation. *Quarterly Journal of the Royal Mete-*
 781 *orological Society*, *131*(613), 3323–3343.
- 782 Gagne, I., David John, McGovern, A., Haupt, S. E., Sobash, R. A., Williams, J. K.,
 783 & Xue, M. (2017, 09). Storm-Based Probabilistic Hail Forecasting with Machine
 784 Learning Applied to Convection-Allowing Ensembles. *Weather and Forecasting*,
 785 *32*(5), 1819-1840.
- 786 Geer, A. (2020). Learning earth system models from observations: machine learning
 787 or data assimilation? *Philosophical Transactions of the Royal Society*, Submitted.
 788 doi: 10.21957/7fyj2811r
- 789 Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press, Cam-
 790 bridge.
- 791 Ham, Y.-G., Kim, J.-H., & Luo, J.-J. (2019, 09). Deep learning for multi-year enso
 792 forecasts. *Nature*, *573*. doi: 10.1038/s41586-019-1559-7
- 793 Hamrud, M., Bonavita, M., & Isaksen, L. (2015, 11). EnKF and Hybrid Gain En-

- semble Data Assimilation. Part I: EnKF Implementation. *Monthly Weather Review*, 143(12), 4847-4864.
- Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., . . . Thépaut, J.-N. (2020). The era5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*. doi: 10.1002/qj.3803
- Houtekamer, P. L., & Zhang, F. (2016). Review of the Ensemble Kalman Filter for Atmospheric Data Assimilation. *Monthly Weather Review*, 144(12), 4489-4532.
- Jaeger, H. (2001). *The "echo state" approach to analysing and training recurrent neural networks-with an erratum note* (GMD Technical Report No. 148). German National Research Center for Information Technology.
- Kalman, R. E. (1960). A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, 82(1), 35-45.
- Kingma, D., & Ba, J. (2014, 12). Adam: A method for stochastic optimization. *International Conference on Learning Representations*.
- Laloyaux, P., Bonavita, M., Chrust, M., & Gürol, S. (2020). Exploring the potentials and limitations of weak-constraint 4d-var. *Quarterly Journal of the Royal Meteorological Society*.
- Laloyaux, P., Bonavita, M., Dahoui, M., Farnan, J., Healy, S., Hólm, E., & Lang, S. T. K. (2020). Towards an unbiased stratospheric analysis. *Quarterly Journal of the Royal Meteorological Society*, <https://doi.org/10.1002/qj.3798>.
- Leutbecher, M., Lock, S.-J., Ollinaho, P., Lang, S. T. K., Balsamo, G., Bechtold, P., . . . Weisheimer, A. (2017). Stochastic representations of model uncertainties at ecmwf: state of the art and future vision. *Quarterly Journal of the Royal Meteorological Society*, 143(707), 2315-2339.
- McGovern, A., Elmore, K. L., Gagne, I., David John, Haupt, S. E., Karstens, C. D., Lagerquist, R., . . . Williams, J. K. (2017, 10). Using Artificial Intelligence to Improve Real-Time Decision-Making for High-Impact Weather. *Bulletin of the American Meteorological Society*, 98(10), 2073-2090.
- Pathak, J., Wikner, A., Fussell, R., Chandra, S., Hunt, B. R., Girvan, M., & Ott, E. (2018). Hybrid forecasting of chaotic processes: Using machine learning in conjunction with a knowledge-based model. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 28(4), 041101. doi: 10.1063/1.5028373
- Piccolo, C., & Cullen, M. (2015, 12). Ensemble Data Assimilation Using a Unified Representation of Model Error. *Monthly Weather Review*, 144(1), 213-224.
- Raanes, P. N., Bocquet, M., & Carrassi, A. (2019). Adaptive covariance inflation in the ensemble kalman filter by gaussian scale mixtures. *Quarterly Journal of the Royal Meteorological Society*, 145(718), 53-75.
- Rasp, S., & Lerch, S. (2018, 10). Neural Networks for Postprocessing Ensemble Weather Forecasts. *Monthly Weather Review*, 146(11), 3885-3900.
- Rodwell, M., & Jung, T. (2008, 07). Understanding the local and global impacts of model physics changes: An aerosol example. *Quarterly Journal of the Royal Meteorological Society*, 134, 1479 - 1497. doi: 10.1002/qj.298
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014, January). Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1), 1929-1958.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1), 267-288.
- Tremolet, Y. (2006). Accounting for an imperfect model in 4d-var. *Quarterly Journal of the Royal Meteorological Society*, 132(621), 2483-2504.
- Van den Dool, H. M. (1994). Searching for analogues, how long must we wait? *Tellus A*(46), 314-324.
- Watson, P. A. G. (2019). Applying machine learning to improve simulations of a chaotic dynamical system using empirical error correction. *Journal of Advances in Modeling Earth Systems*, 11(5), 1402-1417.

- 848 Whitaker, J. S., & Hamill, T. M. (2012). Evaluating Methods to Account for System
849 Errors in Ensemble Data Assimilation. *Monthly Weather Review*, *140*(9), 3078-
850 3089.
- 851 Wolpert, D. H. (1996). The lack of a priori distinctions between learning algorithms.
852 *Neural Computation*, *8*(7), 1341-1390.

Figure.

List of Figures

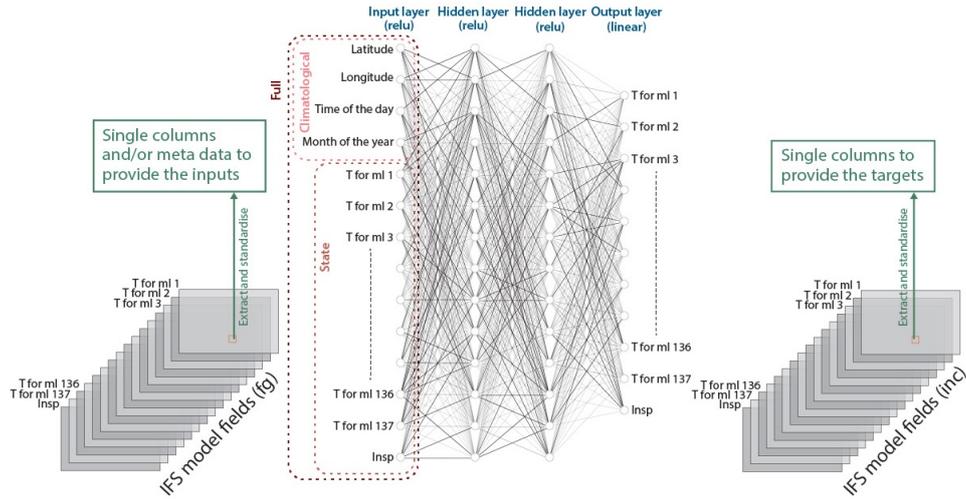


Figure 1. Diagram representing how the `relu_three_layers` ANN is built for the regression over temperature and logarithm of surface pressure. Single columns plus metadata (latitude, longitude, time of the day and month of the year) are extracted from the first guess and analysis increment gridded fields to produce the input and the target of the neural network. Climatological neural network uses only the metadata as input while state neural network uses only the temperature and logarithm of surface pressure values as input. The full neural network combines both information. All neural networks used in this work contain a certain number of fully connected nonlinear layers and an output linear layer.

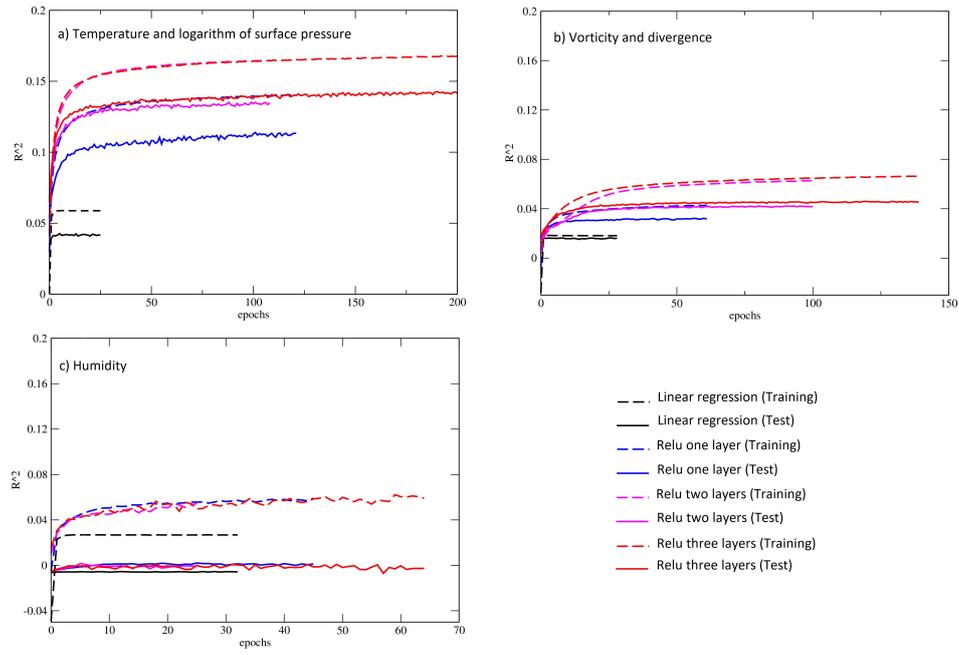


Figure 2. Evolution of the R² coefficient against the number of epochs during the training of different ANN for the (t, lns_p) predictands (a), for the wind (v_o, d) error set of predictands (b) and for the humidity (q) error set of predictands (c). Dashed lines refer to R² values over the training dataset, continuous lines to R² values over the test dataset.

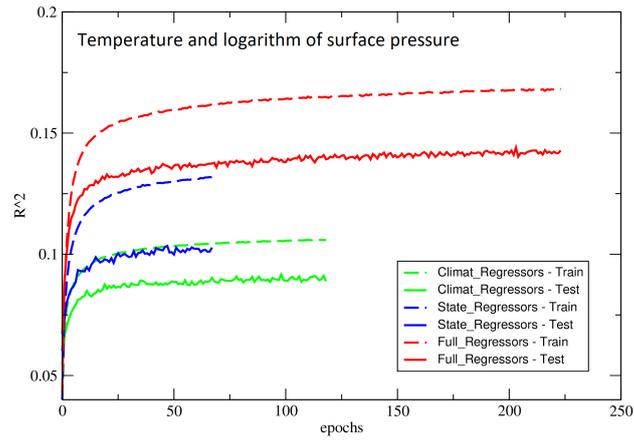


Figure 3. R² coefficients computed during the training phase of the `relu_three_layers` ANNs using only climatological predictors (green curves); only state predictors (blue curves); both sets of predictors (red curves). Dashed curves show R² coefficients for the training dataset, continuous curves for the test dataset.

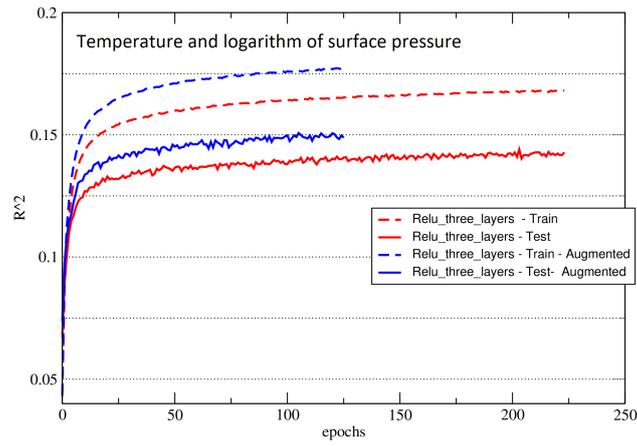


Figure 4. R² coefficients computed during the training phase of a `relu_three_layers` ANN for the prediction of $(t, \ln sp)$ model errors using the standard set of predictors (red lines) and the augmented set (blue lines). Dashed lines for the training dataset, continuous lines for the test dataset.

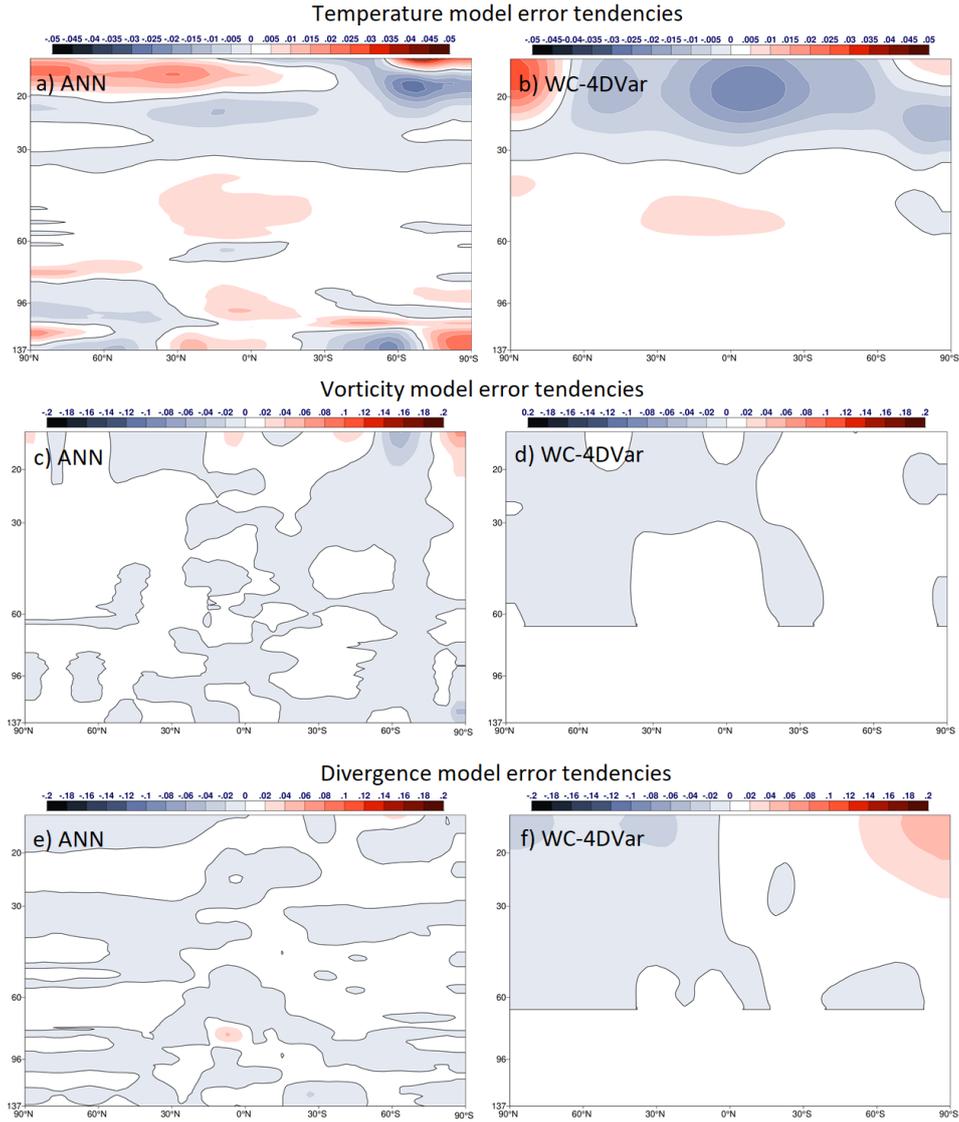


Figure 5. Vertical profiles from model level 15 (approx. 1 hPa) to model level 137 (bottom model level) of longitudinal average of the temperature (a-b), vorticity (c-d) and divergence (e-f) model error tendencies estimated by the ANN described in the text (left panels) and by WC-4DVar from a pre-operational version of IFS cycle 47R1 (right panels). IFS model levels (20, 30, 60, 96, 137) correspond approx. to pressure levels (2 hPa, 10 hPa, 100 hPa, 500 hPa, surface). Values are averaged over a one week period starting on 2019-07-20. The thick black line separates negative tendencies (shades of blue) to positive tendencies (shades of red).

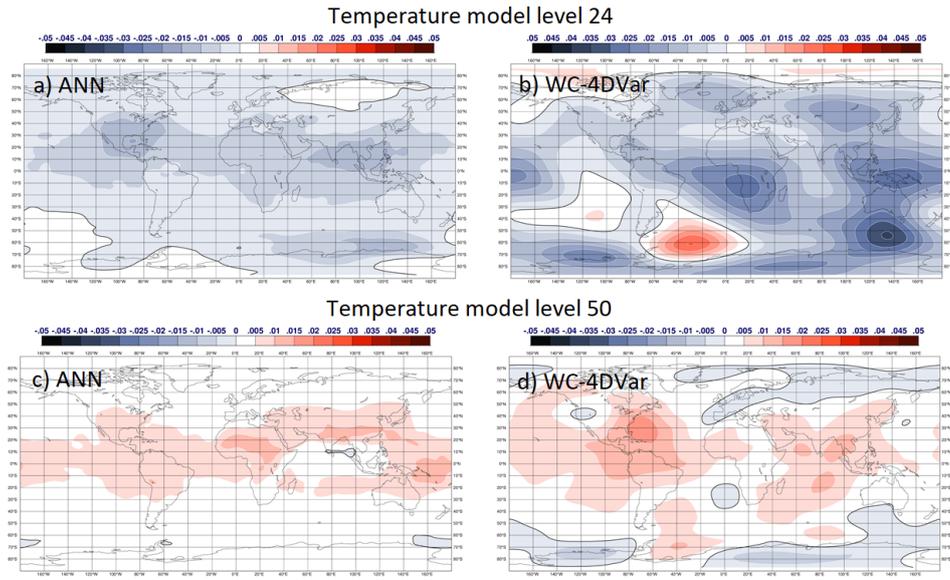


Figure 6. Geographical maps at model level 24, approx. 5 hPa (a-b) and model level 50, approx. 50 hPa (c-d) of the temperature model error tendencies estimated by the ANN (left panels) and by WC-4DVar from a pre-operational version of IFS cycle 47R1 (right panels). Values are averaged over a one week period starting on 2019-07-20. Units in kelvin/h. The thick black line separates negative tendencies (shades of blue) to positive tendencies (shades of red)

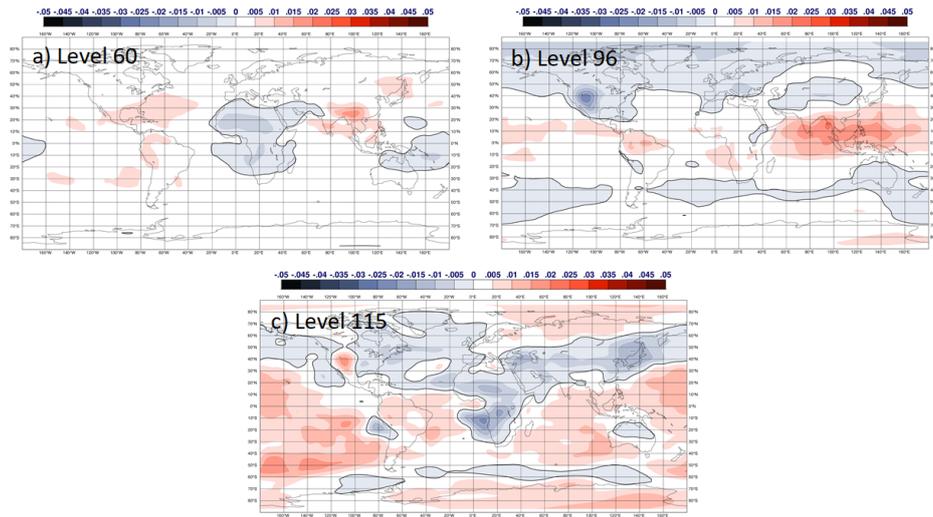


Figure 7. As in Figure 6 for model level 60, approx. 100 hPa (a), model level 96, approx. 500 hPa (b) and model level 115, approx. 850 hPa (c).

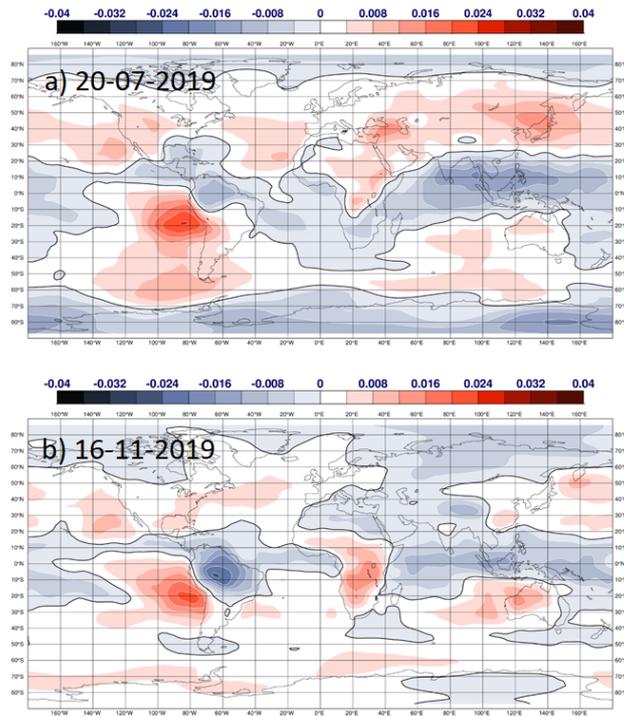


Figure 8. Geographical maps of the surface pressure model error tendencies estimated by the ANN. Values are averaged over two one-week periods starting on 2019-07-20 (a) and on 2019-11-16 (b). Units in hPa/hour. The thick black line separates negative tendencies (shades of blue) to positive tendencies (shades of red)

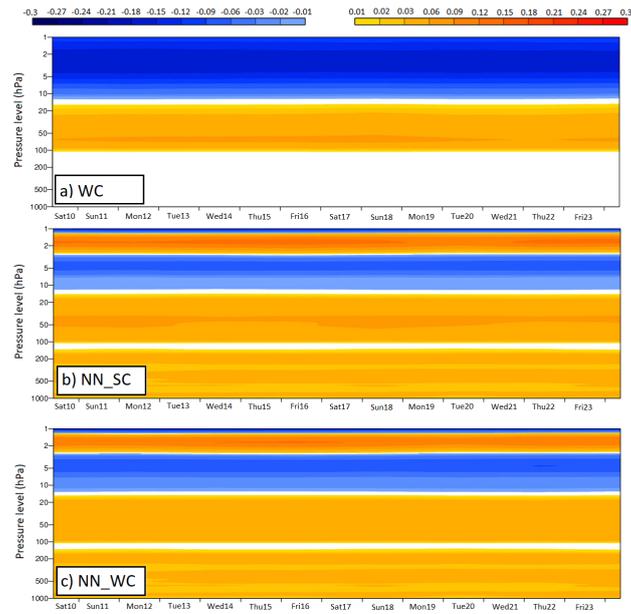


Figure 9. Time series of the global mean model error correction estimated with weak-constraint 4D-Var (top), NN_SC (middle) and NN-WC (bottom) from 15 to 24 August 2019.

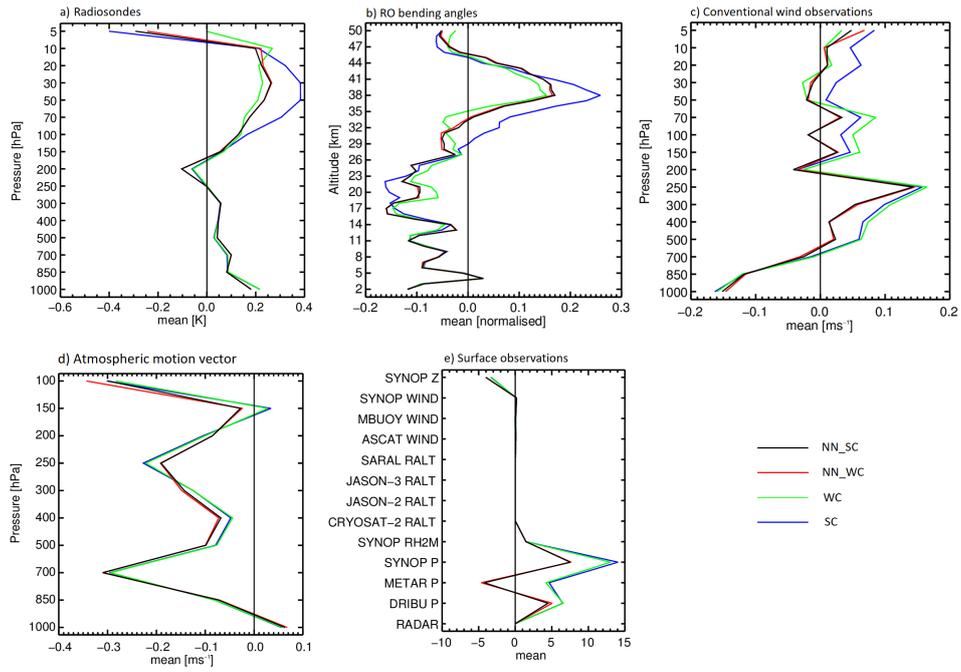


Figure 10. Mean background departures for radiosonde temperature (a), radio occultation bending angles (b), conventional wind observations (c), atmospheric motion vectors (d) and surface observations (e) for the four experiments described in the main text (SC: blue; WC: green; NN_SC: black; NN_WC: red). Values averaged over the 16-07-2019 to 24-08-2019 period. The legend entries SYNOP P, METAR P and DRIBU P stand for surface pressure observations from Synop stations, Metar stations and Drifting Buoys.

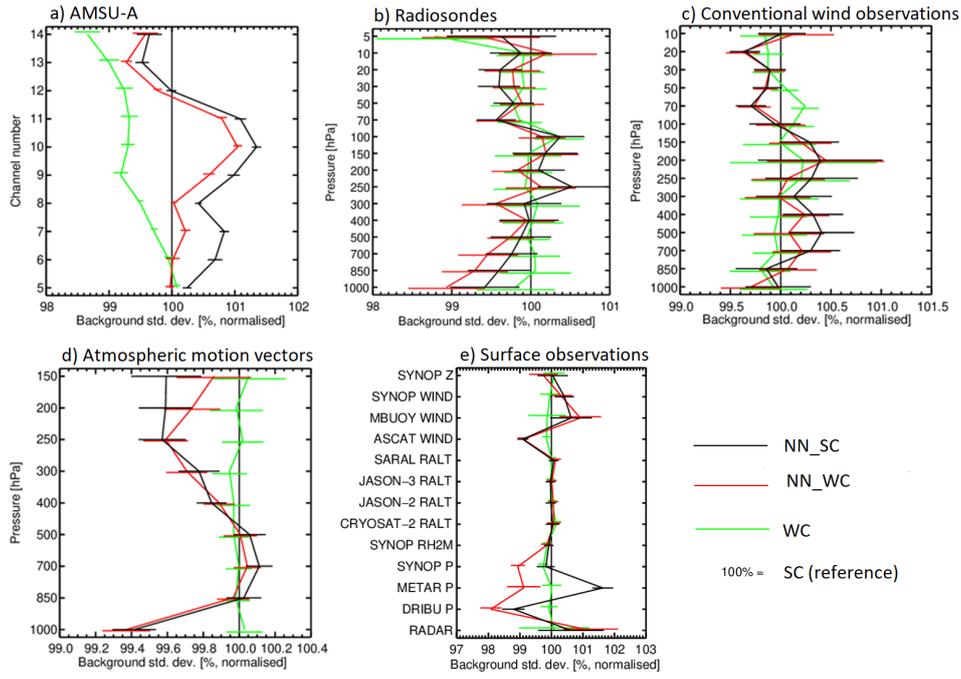


Figure 11. Normalised standard deviation of background departures and background departures for AMSU-A radiances (a), radiosonde temperatures (b), conventional wind observations (c), atmospheric motion vectors (d) and surface observations (e) for the four experiments described in the main text (SC: reference 100%; WC: green; NN_SC: black; NN_WC: red). Values averaged over the 16-07-2019 to 24-08-2019 period. The legend entries SYNOP P, METAR P and DRIBU P stand for surface pressure observations from Synop stations, Metar stations and Drifting Buoys).

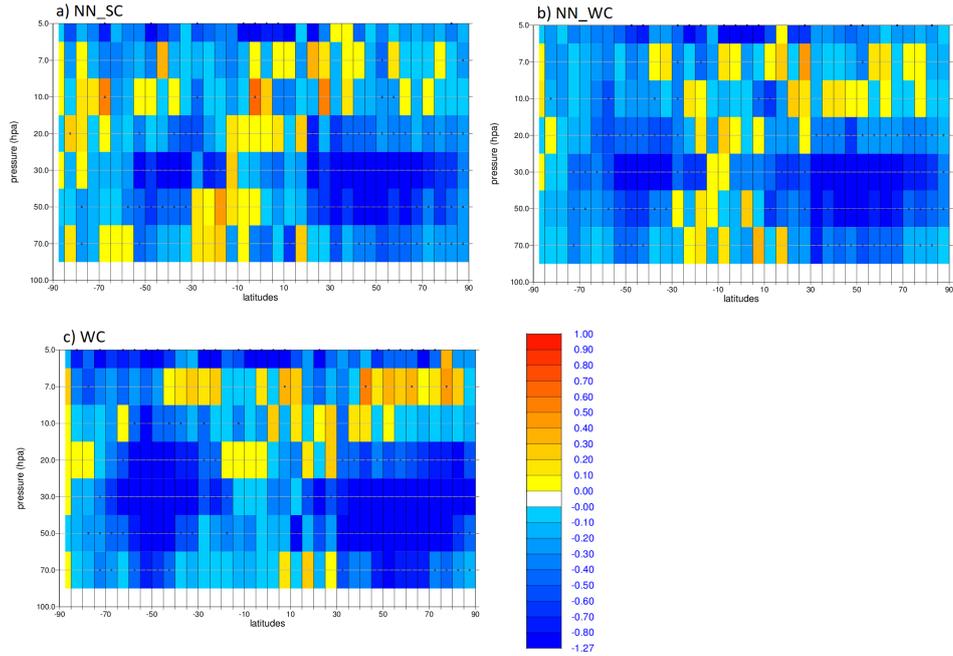


Figure 12. Difference in temperature forecast RMSE after 72 hours between forecasts initialised by NN_SC and strong-constraint 4D-Var (a) and by NN_WC and strong constraint 4D-Var (b) and by weak-constraint 4D-Var and strong-constraint 4D-Var (c). RMSE is computed using radio occultation temperature retrievals and averaged between 10 and 24 August 2019. A negative (positive) difference means that the new system reduces (increases) the forecast error with respect to strong-constraint 4D-Var. Black dots indicate areas where the signal is significant at the 95% confidence level.

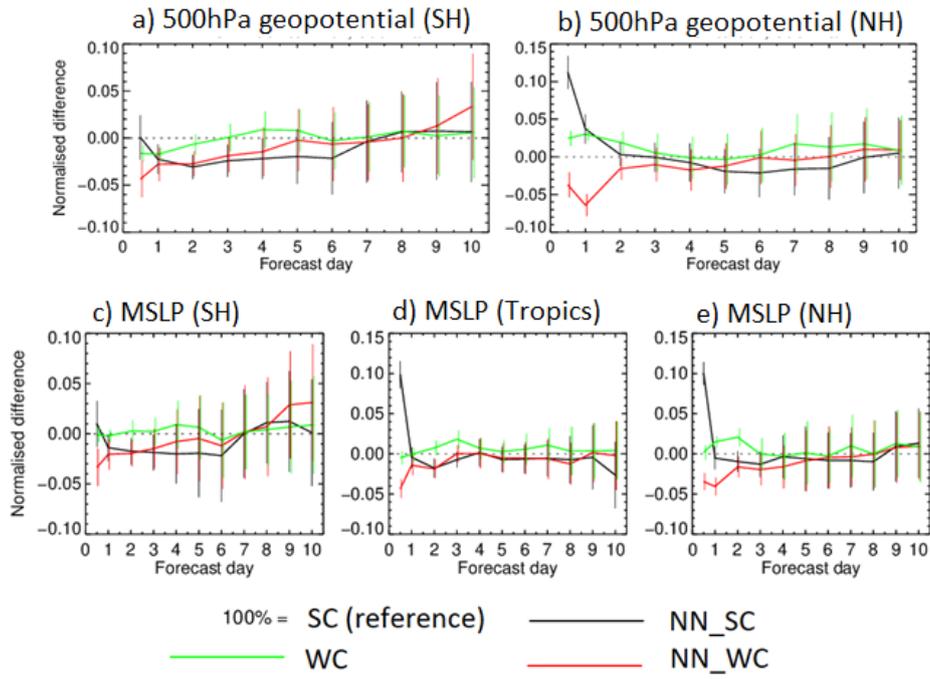


Figure 13. Left panel: Normalised forecast root mean square error of the 500 hPa geopotential in the southern (a) and northern (b) hemisphere, as well as normalised forecast root mean square error of the mean sea level pressure field in the southern hemisphere (c), tropics (d) and northern hemisphere (e). Reference zero line is the strong constraint 4DVar experiment error level. Values averaged over the 16-07-2019 to 24-08-2019 period. Verification is against own analysis. Error bars denote 95% confidence levels.