

1 **Deep learning to estimate model biases in an**
2 **operational NWP assimilation system**

3 **Patrick Laloyaux¹, Thorsten Kurth², Peter Dominik Dueben¹ and David Hall²**

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

5 ²Nvidia Corporation, Santa Clara, California, United States

6 **Key Points:**

- 7 • Temperature retrievals from radio occultation measurements can be used as
8 ground truth to measure stratospheric model biases
9 • 3D convolutional neural networks are suitable for model bias estimation but do
10 not outperform weak-constraint 4D-Var
11 • Transfer learning can help to mitigate data limitations when the atmospheric
12 model is upgraded

Corresponding author: Patrick Laloyaux, European Centre for Medium-Range
Weather Forecasts, Shinfield Park, Reading RG2 9AX, United Kingdom, Email:
patrick.laloyaux@ecmwf.int

Abstract

Model error is one of the main obstacles to improved accuracy and reliability in numerical weather prediction (NWP) conducted with state-of-the-art atmospheric models. To deal with model biases, a modification of the standard 4D-Var algorithm, called weak-constraint 4D-Var, has been developed where a forcing term is introduced into the model to correct for the bias that accumulates along the model trajectory. This approach reduced the temperature bias in the stratosphere by up to 50% and is implemented in the ECMWF operational forecasting system.

Despite different origins and applications, Data Assimilation and Deep Learning are both able to learn about the Earth system from observations. In this paper, a deep learning approach for model bias correction is developed using temperature retrievals from Radio Occultation (RO) measurements. Neural Networks require a large number of samples to properly capture the relationship between the temperature first-guess trajectory and the model bias. As running the IFS data assimilation system for extended periods of time with a fixed model version and at realistic resolutions is computationally very expensive, we have chosen to train, the initial Neural Networks are trained using the ERA5 reanalysis before using transfer learning on one year of the current IFS model. Preliminary results show that convolutional neural networks are adequate to estimate model bias from RO temperature retrievals. The different strengths and weaknesses of both deep learning and weak constraint 4D-Var are discussed, highlighting the potential for each method to learn model biases effectively and adaptively.

Plain Language Summary

The state of the Earth system is estimated via a combination of information from both previous weather predictions and Earth system observations. This complex, mathematical procedure is called data assimilation. Weather predictions could be improved if the error of the numerical models that are used could be reduced. Recent advances in data assimilation at the European Centre for Medium-Range Weather Forecasts (ECMWF) indicate that it is possible to estimate and correct for a large fraction of systematic model errors of those models. During data assimilation, the forecast model and Earth system observations are representing the same situation of the global atmosphere. A direct comparison between models and observations during the short time interval of the data assimilation process can be used to diagnose model errors.

Deep learning is a comparably new method from machine learning that can be used to learn complex mapping procedures. The question we address in this paper is whether deep learning techniques can be used to predict model errors when they are trained to predict the mapping between the global temperature and the model error that was diagnosed during data assimilation.

1 Introduction

Machine learning (ML) has made rapid progress in many domains including natural language processing, computer vision, autonomous driving, healthcare and finance (Goodfellow et al., 2016). Machine learning applications can be very complex, and neural networks (NN) can consist of millions to billions of trainable parameters, large numbers of layers, and specialised architectures. In recent years, the weather and climate modelling community has started to explore machine learning techniques with many applications in Numerical Weather Predictions (NWP) (Dueben et al., 2021). In general, these applications can be divided into three groups: methods that improve computational efficiency, methods that improve the quality of the prediction system,

62 and methods that help improve our understanding of the Earth system, for example
63 via unsupervised learning and causal discovery. This paper belongs to the group that
64 aims to improve prediction quality. In particular, we will use deep learning to learn
65 the systematic error of weather forecast models. Attempts to use DL techniques to
66 estimate and correct for model errors have already been documented in the geophysical
67 literature. For example, Watson (2019) uses an Artificial Neural Network (ANN) to
68 estimate model error tendencies in the Lorenz-96 system. Predicting the error via
69 deep learning is appealing, as errors can often be measured but are typically not
70 easily described by a formula or theory, which makes them difficult to approach using
71 conventional methods.

72 While there are several papers that learn the error during post-processing of
73 the model output (Rasp & Lerch (2018); Groenquist et al. (2021)), this paper will
74 investigate learning model error within the data assimilation (DA) framework of the
75 European Centre for Medium-Range Weather Forecasts (ECMWF). DA is the process
76 that involves merging information from observations with previous model predictions
77 to generate initial conditions for weather forecasts that are both close to the obser-
78 vations and consistent with the state of the forecast models, in order to avoid shocks
79 during model initialisation.

80 Weather observations make a crucial contribution to the quality of today’s nu-
81 merical weather forecasts. Satellites carry passive instruments (e.g. infrared or mi-
82 crowave) to measure natural radiation, while active instruments (e.g. scatterometer
83 or lidar) probe the surface, clouds, and winds by sending out signals and measuring
84 the backscatter (Saunders, 2021). Radio occultation observations evaluate signals sent
85 from one satellite to another (Kursinski et al. (1997)). This array of satellite obser-
86 vations is complemented by a network of in-situ measurements coming from various
87 platforms (e.g. surface stations, aircraft or radiosondes) with a rather inhomogeneous
88 distribution compared to satellite data (Haiden et al., 2018). However, observations
89 are inadequate to provide a complete and accurate picture of the state of the Earth
90 system across the globe at a given point in time. The current model used in opera-
91 tions at the ECMWF contains almost one billion grid points that are updated several
92 times per hour, while only 40 million observations are processed every 12 hours. For
93 this reason, the DA community came up with methods to estimate the most likely
94 state of a system by combining different imperfect sources of information. On the one
95 hand, most observations are unevenly distributed in space and time. They come with
96 errors, and they do not measure the prognostic model variables directly. Instead, they
97 measure quantities linked to these variables, such as radiances or radar echoes. On the
98 other hand, NWP models include the dynamics of the atmosphere and the physical
99 processes that occur. DA combines observations and models in a way that accounts
100 for the uncertainties in each. A popular DA algorithm is the four-dimensional varia-
101 tional (4D-Var) method that iteratively adjusts the initial conditions of a short-range
102 forecast to bring it into closer agreement with meteorological observations in space
103 and time (Rabier et al., 2000).

104 4D-Var is particularly well-suited to satellite data assimilation as it include a
105 radiative transfer model that simulates the top of atmosphere radiances, which are
106 compared to the observed radiances from a specific instrument. This enables the
107 direct application of satellite measurements and extracts the maximum amount of in-
108 formation in clear-sky or all-sky conditions (A. J. Geer et al., 2018). Dealing with
109 random and systematic errors in observations and models is critical for computing an
110 accurate and unbiased estimate. For this reason, an observation error covariance ma-
111 trix is introduced in the 4D-Var formulation to take into account stochastic observation
112 errors arising from the instruments and from the observation operator (Janjic et al.,
113 2018). The error covariance matrix can also represent spatial and inter-channel cross-
114 correlations between observation errors (Waller et al., 2014). Similarly, a background

115 error covariance matrix is implemented to represent flow-dependent, spatially-random
 116 errors in the short-range forecast used in 4D-Var (Bonavita et al., 2016). This matrix
 117 weights the importance of the a-priori state and distributes information horizontally
 118 and vertically in space as well as between model variables (Bannister, 2008a,b). To
 119 deal with systematic observation errors, ECMWF played a pioneering role in the de-
 120 velopment of the Variational Bias Correction (VarBC) schemem, which is embedded
 121 in 4D-Var and automatically removes biases coming from observations and radiative
 122 transfer models. Similarly, the short-range forecast used in 4D-Var also contains sys-
 123 tematic errors which grow over time. A weak-constraint 4D-Var formulation has been
 124 proposed to estimate these model biases within the assimilation process and to correct
 125 the dynamical model accordingly (Laloyaux, Bonavita, Dahoui, et al., 2020).

126 There are strong mathematical similarities between the 4D-Var formulation in
 127 data assimilation and the training of NNs. Both use gradient descent techniques, and
 128 the adjoint method for calculating gradients in 4D-Var is mathematically identical to
 129 the standard backpropagation method used in NN training. From a broad enough
 130 viewpoint, DA and ML may be viewed as two flavours of inverse methods that can be
 131 united under Bayesian statistics (A. Geer, 2020). Brajard et al. (2020) demonstrated
 132 a way to combine ML with DA when observations are noisy and partial. In their
 133 scheme, DA and ML alternate and compute progressively more accurate estimates of
 134 the state and of the surrogate predictive model. Following this idea, (Farchi, Laloyaux,
 135 et al., 2021) used a dataset of analysis increments to train a ML statistical/empirical
 136 model that complements the original dynamical model. The resulting hybrid surrogate
 137 model significantly improves the accuracy of the analysis and produces better short-
 138 and mid-range forecasts in a two-layer, two-dimensional, quasi-geostrophic channel
 139 model. These encouraging results with a simplified system have been confirmed to
 140 a certain extent in the operational atmospheric Integrated Forecasting System (IFS)
 141 model developed at ECMWF (Bonavita & Laloyaux, 2020). The idea of using time
 142 series of analysis increments fields to estimate the predictable component of model error
 143 is not new in the meteorological literature. For example, one of the algorithms proposed
 144 in (Dee, 2005) for the correction of model bias in a cycled data assimilation framework
 145 explicitly involves using an online model error estimate based on a running mean
 146 over past analysis increments. The increments have global, homogeneous coverage
 147 and are already available in the space of the dynamical model variables which makes
 148 the method easy to implement. However, this approach also has some limitations,
 149 as increments can contain signals that are not induced by model biases but by other
 150 error sources that have not been properly accounted for in the DA system. A well-
 151 known illustration is the positive temperature increment in the ERA-Interim reanalysis
 152 coming from aircraft temperature biases that have not been corrected properly by
 153 VarBC (Dee & Uppala, 2009).

154 This paper will focus on the estimation and correction of temperature system-
 155 atic errors (bias) in the stratosphere using satellite temperature retrievals as ground
 156 truth. But how important are these biases for NWP? In a global NWP model, the
 157 troposphere may be viewed as a turbulent boundary layer for the atmosphere, and
 158 the stratosphere as being comparatively isolated from the surface of the Earth. To
 159 a first approximation, the global-mean stratosphere is in radiative equilibrium, with
 160 long-wave cooling balancing solar heating through ozone absorption (Fomichev et al.,
 161 2002). The latitudinal temperature structure is affected by the meridional circulation
 162 which is driven by breaking and dissipating planetary and gravity waves in the strato-
 163 sphere. To quantify how stratospheric biases influence the troposphere, we ran a denial
 164 experiment and blacklisted observations that are important for stratospheric variables.
 165 This includes stratospheric observations from radiosondes, aircraft, RO bending an-
 166 gles above 100hPa, as well as the microwave and infrared stratospheric channels (see
 167 details in Table1). It is not possible to remove all observations that are sensitive to the
 168 stratospheric conditions, as microwave instruments measure radiances that reflect con-

| Type | Pressure/Altitude/Channels |
|------------|--|
| Radiosonde | above 100hPa |
| Aircraft | above 100hPa |
| RO | above 17km |
| AMSU-A | 9,10,11,12,13,14 |
| ATMS | 10,11,12,13,14,15 |
| AIRS | 7, 15, 20, 21, 22, 27, 28, 40, 52, 69, 72, 92, 93, 98, 99, 104, 105, 110, 111, 116, 117, 123, 128, 129 |
| IASI | 16, 38, 49, 51, 55, 57, 59, 61, 63, 66, 70, 72, 74, 79, 81, 83, 85, 87, 89, 101, 104, 106, 109, 111, 113, 116, 119, 122, 125, 128, 131, 133,138, 135, 141, 144, 146, 148, 151, 154, 157, 159, 161, 163, 165, 167, 170, 176, 178, 183, 189, 191, 195, 197, 201, 203, 301, 303 |
| CrIS | 20, 23, 26, 33, 36, 39, 42, 45, 48, 51, 54, 57, 60, 61, 62, 63, 64, 65, 66, 68, 69, 70, 71, 73, 74, 113, 114 |

Table 1. List of all the observations considered as sensitive to stratospheric conditions and withheld in the denial experiments

169 conditions in a deep layer of the atmosphere. This means that some tropospheric-peaking
170 channels could still have a slight impact on the stratosphere.

171 The data-denial experiment runs over two months, between the 25th of January
172 2020 and the 25th of March 2020. The top panel of Figure1 shows the impact on the
173 analysis mean error when stratospheric observations are withheld. The large biases
174 developed in the stratosphere over these two months are transferred to the troposphere,
175 especially over the Southern pole. The bottom panel of Figure1 shows how these
176 biases present in the analysis evolution during forecasts. The impact of the missing
177 stratospheric observations can still be observed after 48 hours. The impact shrinks
178 with the forecast lead time as the model drifts towards its climatology and forgets
179 about the information present in the initial conditions. This experiment shows the
180 importance of tackling residual stratospheric temperature biases as they can descend
181 into the troposphere.

182 It is important to note, that the model bias changes when the IFS model is
183 upgraded, on a regular basis. The most recent improvements to the stratospheric
184 physics are the implementation of a new radiation scheme and ozone climatology in
185 cycle 46r1 (Hogan et al., 2017; Shepherd et al., 2018). Furthermore, a quintic vertical
186 interpolation has been implemented in the semi-Lagrangian advection in cycle 47r1
187 (Polichtchouk et al., 2019) to resolve a larger fraction of gravity waves in the vertical
188 direction. These changes reduced the temperature bias in the stratosphere, but the
189 residual bias is still significant. It consists of a global-mean cold bias in the lower/mid
190 stratosphere of -0.5C and a global-mean warm bias in the upper stratosphere of 0.5C
191 that accumulate over a 12-hour data assimilation window.

192 This article develops a deep learning solution for estimating the three-dimensional
193 stratospheric temperature bias in the IFS. State-of-the-art NNs are trained to learn
194 the mapping from three-dimensional fields of stratospheric temperature to the three-
195 dimensional bias diagnosed via Radio Occultation (RO) temperature retrievals. As a
196 first step, we use information from ERA5 reanalysis to show that deep learning can
197 indeed learn to predict the three-dimensional temperature bias of short-term forecasts
198 when using a large training data set spanning several years. In a second step, we
199 study the use of transfer learning to adjust the trained model when only one year of
200 training data is available for a new model cycle. Finally, we perform tests that apply

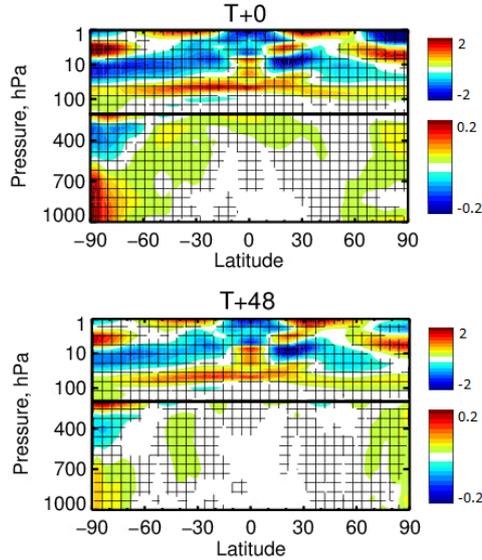


Figure 1. Difference in the forecast mean error across zonal bands at lead time +0h (top) and +48h (bottom) when all the stratospheric observations valid above 100hPa are withheld. Scores have been computed against the operational analysis between January 25th, 2020 and March 25th, 2020. Different colorbars are used for the stratosphere and for the troposphere.

201 the NN bias correction within 4D-Var DA experiments and compare results against
 202 weak-constraint 4D-Var which serves as a benchmark.

203 A description and assessment of the RO temperature retrieval dataset are pre-
 204 sented in Section 2. The design and the training of various NN solutions are sum-
 205 marised in Section 3. Section 4 describes results obtained when the NN temperature
 206 correction is applied to the model in an assimilation experiment. This NN approach is
 207 then compared with the weak-constraint formulation used in operations at ECMWF in
 208 Section 5. We finally discuss various aspects of weak-constraint 4D-Var that are also
 209 essential for ML such as learning rate and NN retraining, in Section 6. We summarize
 210 the paper in Section 7 and provide a perspective for future developments.

211 2 Temperature retrieval datasets

212 It is very challenging to produce a ground-truth database for Numerical Weather
 213 Prediction (NWP) as all weather measurements and weather simulations contain errors
 214 that cannot be ignored. However, some types of observations are more accurate than
 215 others and can therefore serve as a reasonable proxy for the true atmospheric state.
 216 This is the case for the GNSS Radio Occultation (RO) measurements in the strato-
 217 sphere, which offer a spatially homogeneous observing system. These measurements
 218 consist of high-quality bending-angle profiles that are sensitive to the stratospheric
 219 temperature. It has been shown that RO profiles reduce NWP analysis and forecast
 220 temperature biases in the lower and middle stratosphere for most NWP centres (Healy
 221 & Thépaut, 2006; Poli et al., 2009; Rennie, 2010; Cucurull et al., 2013).

222 The RO measurement technique is described in Kursinski et al. (1997). The
 223 GPS signals propagating between the GPS transmitter and a receiver on a low earth
 224 orbiting (LEO) satellite are bent by gradients of the refractive index in the ionosphere
 225 and neutral atmosphere, as they pass through the limb of the Earth. The ionospheric

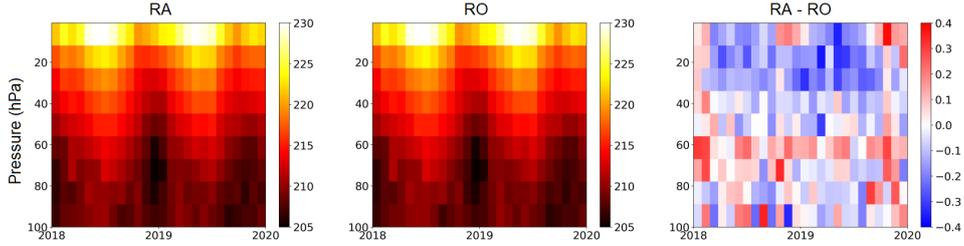


Figure 2. Timeseries of collocated radiosondes (left), RO temperature retrievals (middle) and the difference between the two (right). Observations are collocated on a 5-degree grid every hour within a 1hPa pressure difference between 2018 and 2020

226 bending can be removed with a simple correction (Vorobev & Krasilnikova, 1994). The
 227 ray bending as function of “impact parameter” can be determined, as a result of the
 228 motion of the LEO satellite. The impact parameter defines the height of the tangent
 229 point of the ray path above the surface. The ray-bending angle values as a function of
 230 impact parameter can be inverted to provide information about the atmospheric state,
 231 such as temperature. RO measurements are distributed globally, have good vertical
 232 resolution, and RO bending angles can be assimilated without bias correction into the
 233 NWP model (Healy & Thépaut, 2006). However, in the context of this work, profiles of
 234 mean bending angle departures can be difficult to interpret since a given bending angle
 235 can have both positive and negative sensitivity to temperature biases in the vertical
 236 profile (see section 5.3 Eyre, 1994). We have therefore mapped the bending angle
 237 profiles to temperature using a simple implementation of the widely used temperature
 238 retrieval algorithm described by Kursinski et al. (1997). Refractive index profiles
 239 are derived from bending angles with an Abel transform. There is no measurement
 240 information to enable the separation of the effects of temperature and water vapor,
 241 and therefore these quantities can be retrieved only using prior information (ERA5
 242 reanalysis in our case). Although this retrieval method provides temperature values
 243 up to the top of the atmosphere, the retrieval noise increases with height and most of
 244 the information comes from the prior above 3hPa. Therefore, we only use the retrieved
 245 temperature values between model level 20 (3hPa) and model level 65 (125hPa) out
 246 of a total of 137 vertical levels in the IFS.

247 It is important to evaluate the quality and accuracy of the temperature retrievals,
 248 as they are used as ground truth in our study. RO temperature retrievals can be collocated
 249 with conventional temperature observations from radiosondes (RA) to quantify
 250 the error characteristics of the observing system (Sun et al., 2010). RO and RA profiles
 251 are not available at exactly the same vertical and horizontal location. For comparison,
 252 profiles have been collocated on a 5-degree grid, every hour, and within a 1hPa pressure
 253 difference. Figure 2 shows a timeseries of the collocated observations from RA (left)
 254 from RO (middle) and the difference RA-RO (right). This has been averaged over
 255 pressure levels and for every month between 2018 and 2020 to reduce the collocation
 256 errors introduced through spatial and temporal mismatch between RA and RO that
 257 could influence the accuracy of the obtained statistics. The RA and RO observations
 258 present a very similar seasonal signal when the stratosphere is warming up during the
 259 Northern hemisphere summer, or cooling down during the Northern hemisphere winter.
 260 This pattern arises from the inhomogeneous distribution of radiosondes, mainly
 261 sampling the Northern hemisphere. The difference between RA and RO (right panel of
 262 Figure 2) shows that the average discrepancies between the two types of observations
 263 in the mid/lower stratosphere are smaller than 0.2C and confirms what has been found
 264 in other collocation studies (Sun et al., 2010, 2019). In the upper stratosphere (above

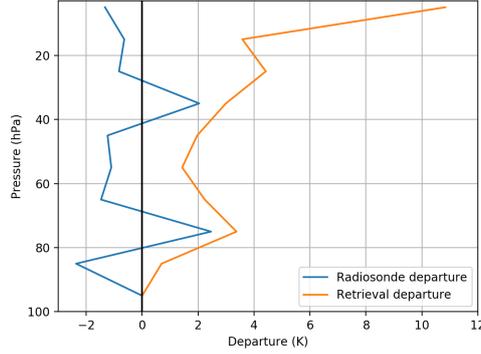


Figure 3. Vertical profile of the ERA5 first-guess departure from a collocated radiosonde (blue) and RO temperature retrieval (orange). Both profiles are measured over the USA (36N, 93W) on 16/10/2020.

265 30hPa), there is a systematic difference where RO observations are warmer than RA
 266 by approximately 0.3C, especially in summer. This shows the intrinsic challenge of
 267 finding the ground truth in NWP as each observing system will be sensitive to different
 268 sources of error (e.g. solar elevation angle, dry temperature adjustment, ...).

269 During the collocation study, a small fraction (less than 1%) of profiles showed
 270 very large discrepancies. One example is illustrated in Figure 3 for a collocated profile
 271 over the USA (36N, 93W) on 16/10/2020. The RA profile agrees roughly with the
 272 ERA5 first-guess trajectory, presenting a small first-guess departure. However, the
 273 RO profile shows very large differences with respect to the trajectory of ERA5 (over
 274 5 degrees in the upper stratosphere). Future work could include an improved quality
 275 control procedure to detect and automatically remove outlier RO profiles with lower
 276 quality. The current QC is based on the parameters used in the bending angle assim-
 277 ilation, but the bending angle assimilation is more robust to measurement noise than
 278 the RO temperature retrievals used here.

279 The purpose of the NN is to learn a function representing the model bias that
 280 develops in the data assimilation system over the 12-hour assimilation window. A
 281 natural choice for the input of the NN is the temperature first-guess trajectory, as it
 282 contains the state of the model. The output of the NN is the model bias estimated as
 283 the difference between the temperature first-guess trajectory and the RO retrievals.
 284 The spatial and temporal structure of the stratospheric temperature bias has been
 285 studied in Laloyaux, Bonavita, Dahoui, et al. (2020) and presents large scale patterns
 286 that evolve slowly over time. For this reason, the first-guess trajectory and first-guess
 287 departure are averaged over a 10-degree regular grid for all the model levels between
 288 130hPa (level 65) and 3hPa (level 20). This means that we have 31,635 inputs and the
 289 same number of outputs (19 latitude grid points x 37 longitude grid points x 45 vertical
 290 levels). Unfortunately, the observations are not available at every point in space and
 291 time. To reduce the number of missing data points, we average the input/output
 292 samples over 10 days. The averaging also helps capture slowly varying signals. Linear
 293 interpolation is used to fill the remaining observational gaps (representing 5% missing
 294 values when using the 10-day average).

295 Machine Learning requires a large number of samples to properly capture the
 296 relationship between input and output variables. To run a dedicated assimilation
 297 system with the current IFS model for a long time period is computationally expensive
 298 and serial in time, and therefore very slow. It is thus prohibitive to train the networks
 299 within the assimilation framework. Therefore, to obtain training data for a long time

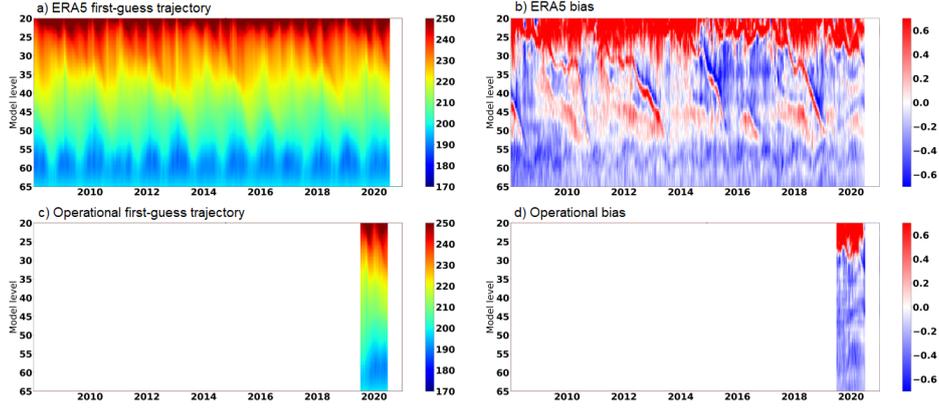


Figure 4. Timeseries of ERA5 temperature first-guess (top left) and departure with RO temperature retrievals (top right) for the different stratospheric model levels (level 20 is 3hPa and level 65 is 130hPa) averaged between 5N and 5S available between 2008 and 2020. The bottom panels show a similar timeseries from the operational dataset that is available only between June 2019 and June 2020.

300 period, we use data from the ERA5 reanalysis as the first-guess trajectories, and
 301 departures have been archived over the entire period for which good RO coverage is
 302 available (from 1st of January 2008 until 1st of June 2020). ERA5 is based on an
 303 IFS model version (cycle 41r2) implemented in 2015. As we also want to study how
 304 a trained bias correction tool can be adjusted to a new model cycle, we also estimate
 305 the bias of the model used in operations between June 2019 and June 2020. We will
 306 refer to this dataset as the "operational dataset". It consists of one year of first-guess
 307 trajectories from cycle 46r1, which improves several aspects of the dynamics and the
 308 physics of the model, compared with the ERA5 dataset. The spatial resolution of the
 309 two datasets is identical and equal to 18km (the control member of the Ensemble Data
 310 Assimilation system is used for the operational dataset, instead of the high-resolution
 311 system). Figure 4 shows a timeseries of inputs and outputs produced from the ERA5
 312 (top) and the operational (bottom) dataset, averaged over the Tropics between 5S
 313 and 5N. The ERA5 model exhibits a cold bias in the mid/lower stratosphere and a
 314 warm bias in the upper stratosphere that propagates down during QBO events. The
 315 operational dataset has a similar vertical structure, although the amplitude is much
 316 larger.

317 ML studies generally divide the available data into three different datasets to
 318 train, develop, and evaluate an ML model. The training set is the largest and is
 319 used to learn the relationship between input and output variables. The second set,
 320 referred to as the validation set, is used exclusively for tuning model hyper-parameters
 321 set manually by the model developer (e.g. activation function, learning rate). A
 322 key goal of the hyper-parameter tuning process is the optimization of the network's
 323 generalization capabilities, in order to avoid overfitting and ensure that the network
 324 will function well on previously unseen data. The third dataset is the test set, a
 325 collection of previously unseen data, which is used to evaluate the network. The three
 326 datasets should be independent of each other, but at the same time they should reflect
 327 the same statistical distribution. Several strategies are discussed by Schultz et al.
 328 (2021) to achieve this with meteorological time series that are usually auto-correlated.
 329 A block sampling strategy is used for our application in order to mitigate this issue.
 330 For the ERA5 datasets, we assign the first 10 days of every month and the days after
 331 the 20th of each month in 2019 to the validation set and the remaining samples to the

| Dataset | training | validation | test |
|-------------|-----------------|------------|----------------|
| ERA5 | 2008-2019 (412) | 2019 (26) | 2020-2021 (42) |
| operational | 2019 (15) | 2019 (5) | 2020 (18) |

Table 2. Partition details for ERA5 and operational data sets. Shown are only the years and total number of samples, in parenthesis. For overlapping years, the data is split into date ranges for each month, as described in the text, in order to create disjoint sets.

332 training set. The test set is comprised of the entire year 2020 and the first half of 2021.
 333 For the operational dataset, we use a similar strategy for splitting it into training and
 334 validation sets. Since only data for half of 2019 is available, we assign June 21 to July
 335 1st, July 21st to 31 and August 10 to the validation set, and the rest of the 2019 data
 336 to the training set. We assign the full set of 2020 operational data to the test set.
 337 Table 2 summarizes the various splits for the two datasets.

338 3 Design and training of neural networks

339 3.1 Data representation

340 The machine learning problem at hand is a multi-dimensional regression problem:
 341 the state of the IFS model is used as the input for our network, and the bias computed
 342 from the departures between the temperature retrievals and the IFS model is our
 343 prediction target. This means, that we are aiming to learn the departure values and
 344 not the RO ground-truth data itself, which typically produces a more stable learning
 345 process.

346 The raw data is available as tuples in the form (longitude, latitude, level, T),
 347 where T represents the temperature at these coordinates. In this paper, we make
 348 use of data regression on structured grids using convolutional neural networks, after
 349 converting the data into multi-channel images using suitable projections and inter-
 350 polations, with the vertical model level mapped to the feature/channel dimension.
 351 We examine two possible interpretations of this data: they can be treated as three-
 352 dimensional fields, consisting of (projected longitude, latitude and level) with a single
 353 feature (temperature)¹, or as two-dimensional fields of (projected longitude, latitude)
 354 with a vector of features (temperatures at different altitudes). We will discuss the
 355 implication of these two different views below.

356 In order to stabilize training, we rescale the data using mean-variance normaliza-
 357 tion. For the 2D case we perform a separate normalization per altitude/level, whereas
 358 for the 3D case we perform a single normalization across all levels. This is important
 359 in order to preserve vertical gradients in the data. In each case, we normalize the input
 360 and target datasets separately.

361 3.2 Network architecture

362 Image regression is similar to image segmentation, and NNs which are designed
 363 for segmentation can be adjusted to work for regression simply by dropping the final
 364 per-pixel softmax function. Therefore, we can choose a suitable network architec-
 365 ture from a plethora of available image segmentation networks, such as the UNet

¹ Note that model levels can be transformed into physical altitudes by applying an exponential map-
 ping. However, in the 3D approach we treat them as equidistant and rely on the network to learn a rea-
 sonable set of filters.

366 (Ronneberger et al. (2015)) or DeepLab (Chen et al. (2018)) architectures. Both of
 367 these architectures employ an encoder, which is responsible for extracting features at
 368 multiple length scales. The encoder typically consists of convolutional blocks, with
 369 skip connections added to improve training stability. The output of the encoder is
 370 passed to a decoder which combines the extracted features to generate a prediction.
 371 DeepLab architectures employ an additional step between the encoder and decoder,
 372 the so-called atrous spatial pyramid pooling (ASPP) (Chen et al. (2018)) designed to
 373 improve feature combination at different scales. The DeepCAM (Kurth et al. (2018))
 374 NN architecture has been successfully applied to the identification of extreme weather
 375 phenomena in climate simulations. Therefore, we use a modified variant of the origi-
 376 nal architecture in which the ResNet-50 (He et al. (2016)) backbone is replaced by an
 377 XCception (Chollet (2017)) backbone. Also, instead of relying on interpolated upsam-
 378 pling, we employ a fully convolutional decoder. These two improvements lead to the
 379 network architecture which forms the basis of the MLPerf HPC DeepCAM benchmark
 380 (see *MLPerf HPC DeepCAM website* (2021)). In order to reduce checkerboard arti-
 381 facts produced by convolutional upsampling, we furthermore insert average pooling
 382 layers after the convolutions with a pooling kernel size equal to the convolutional up-
 383 sampling stride. It has been shown by Kinoshita & Kiya (2020) that this is an effective
 384 technique for reducing such artifacts in the generated images.

385 For the 2D representations of the data, we simply adjust the number of input
 386 channels in the previously described architecture. For the 3D representation, we con-
 387 vert all 2D operations (convolutions, batch-normalizations, pooling) into their respec-
 388 tive 3D counterparts. A significant difference between these two approaches is that in
 389 the 2D case, all altitude levels are combined in an all-to-all fashion through the matrix
 390 multiplication along the feature dimension in the 2D convolutional kernel. In contrast,
 391 the 3D convolutions only correlate neighbouring levels. Therefore, they have are bet-
 392 ter suited to capturing temperature gradients between levels, whereas 2D convolutions
 393 might be better at capturing long distance correlations spanning multiple levels. Both
 394 architectures are reasonable choices for solving the bias prediction problem at hand,
 395 and thus we pursued both approaches.

396 3.3 Training process, R2-score and hyper parameter optimization

397 We employ the AdamW optimizer (Loshchilov & Hutter (2019)) and apply weight
 398 decay regularization in order to reduce overfitting, which is particularly important
 399 when training on the smaller, operational dataset. For the loss function, we use either
 400 the L2 distance or a smooth version of the L1 distance between network output and
 401 prediction target.

The R2 score is used as a validation metric for hyper parameter tuning. It is defined as

$$R2 = 1 - \frac{\sum_{i=1}^m (y^{(i)} - f^{(i)})^2}{\sum_{i=1}^m (y^{(i)} - \bar{y})^2}, \quad \text{where } \bar{y} = \frac{1}{m} \sum_{i=1}^m y^{(i)} \quad (1)$$

402 Here, $y^{(i)}$ is the NN prediction for sample i and $f^{(i)}$ is the corresponding ground truth,
 403 i.e. in our case the model bias. The R2 score compares the prediction accuracy with
 404 the intrinsic variance of the data: if prediction accuracy is high, then the numerator
 405 in equation (1) is small, which leads to $R2 \approx 1$. If the prediction accuracy does not
 406 outperform the intrinsic noise, then the numerator and denominator in equation (1)
 407 will be of similar magnitude and we find that $R2 \approx 0$. For predictions of even lower
 408 accuracy we have $R2 < 0$ which signals a failure of the model. In order to obtain
 409 a scalar score, we perform a summation over all the pixels and levels in the output
 410 image. However, a more detailed qualitative analysis is possible by computing the R2
 411 score per level and/or per longitude/latitude coordinate.

We tune hyper parameters (HPO) using the ray.tune package (*Ray Tune website* (2021)) with HyperOpt (*Hyperopt website* (2021)), running 128 instances for both the 2d and 3d models. Tuneable hyper parameters in our model include learning rate, weight decay, learning rate schedules (selection of multi-step with different milestones, cosine annealing with different choices for decay frequency), loss definition (smooth L1 vs. L2) and batch size. Our hyperparameter optimization target is the maximisation of the R2 value, as described above. Each model is trained for about 150 epochs.

3.4 Computational Performance

We use a single NVIDIA DGX-2 system for training and run a single instance on each GPU concurrently. This means, we can train 16 instances in parallel. Training a single instance on an NVIDIA V100 GPU takes about 30 minutes for the 3D model. Therefore, training 128 instances does not take longer than 4 hours in total. It is unlikely that training more instances would lead to the discovery of better hyperparameters, because many good hyperparameter choices² perform equally well and it is hard to define a quantitative criterion which configuration to prefer over the others.

3.5 Training, using a small operational dataset

In order to produce the most accurate weather forecast possible, we would like to construct a NN bias-correction model based on data from the latest IFS model cycle. While there is plenty of ERA5 data available (based on the 2015 cycle), the dataset for the current cycle is much smaller. In our case, we had only 15 training samples and 5 validation samples available. We examined several approaches in an attempt to build the most useful tool for this scenario.

1. Finetuning: in this approach, the model does not begin using random initial weights. Rather, it begins with weights that have been pre-trained on a related dataset. Specifically, we pre-trained the model using the ERA5 dataset and then fine-tuned the entire model using only the operational dataset, but with a much smaller learning rate. Using this approach, we found that the NN quickly overfit to the operational data. Therefore, we did not pursue this approach further.
2. Training from scratch: In this case we trained the model using only data from the latest IFS cycle. While it appeared more promising than finetuning for the first few epochs, this approach broke down rapidly as well, heavily overfitting the training dataset for all hyper parameter configurations tried. Hence we abandoned this approach as well.
3. No retraining: in this, simplest approach. we used only the existing model, trained exclusively on the ERA5 dataset, with no fine-tuning. This model was then applied directly to the shorter, operational dataset. This approach is promising if the underlying intrinsic features of both datasets are similar. It turns out that this approach yields reasonable results, producing R2 values *only* about $\sim 20\%$ lower than the original ERA5 test dataset.
4. Training on both datasets simultaneously: for this transfer learning strategy, we implemented a data loader which can feed the NNs samples from either dataset. The two datasets have a relative sample imbalance of about 27:1 (ERA5:operational). In order to help the NN learn the features of the operational dataset, the dataloader selects samples from the both datasets, but with inverted frequencies. This means that the NN is presented samples from both datasets with almost equal probability. In practice, we chose a final ratio of

² *good* means that they deliver a high R2 score on the validation set

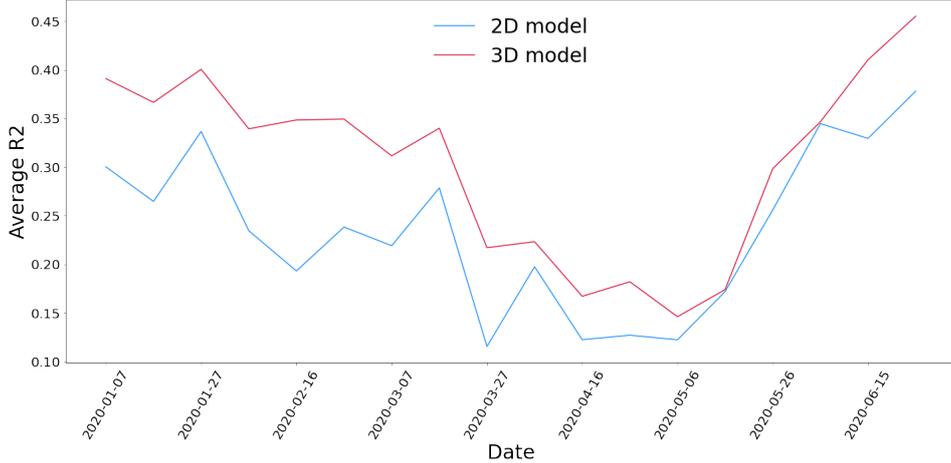


Figure 5. Timeseries of R2 values averaged over all levels for 2D + features (blue) and 3D NN (red) architectures on the 41r2 test dataset.

459 27:29 (ERA5 : operational data) in order to provide a small emphasis on the im-
 460 portance of the operational data. (This is a tunable hyperparameter). For the
 461 validation dataset, we use only samples from the operational set, as we are inter-
 462 ested only in the operational model performance. We also use only operational
 463 data to compute the R2 score when performing hyper-parameter optimisation.
 464 It turns out that training the network with this approach is very stable.

465 4 Results for bias correction with deep learning

466 In this section, we present results for temperature bias correction based on deep
 467 learning. The first and second subsections present results for offline bias correction,
 468 for the ERA5 and operational datasets respectively. The third subsection discusses
 469 the use of bias correction within data assimilation experiments.

470 4.1 Performance comparison of 2D and 3D models

471 We trained the 2D and 3D models with their respective best known hyperpa-
 472 rameters on the ERA5 training dataset and compared their performance on the ERA5
 473 test dataset. Figure 5 displays the R2 value, averaged over all levels, for the test set
 474 . The plot demonstrates that the 3D model outperforms the 2D model consistently.
 475 This is likely due to the importance of gradients and other local co-variances in the
 476 vertical direction, and the inherent advantage convolutions provide for learning such
 477 relationships in a data-efficient fashion. Therefore, we decided to conduct subsequent
 478 studies exclusively using 3D network architectures. The variability of the first-guess
 479 trajectory and of the model bias is larger between March and May, as the Northern
 480 hemisphere warms up and the Southern hemisphere cools down. There is a drop in
 481 the R2 value for both models as they struggle to accurately capture the model bias
 482 over that period.

483 Figure 6 shows the target bias (left) and the prediction of the 3D model (right)
 484 for two different vertical levels on February 6 2020, from the ERA5 test set. The NN
 485 clearly learns to reproduce important features, such as the negative bias correction
 486 around the equator for level 40. It also learns to reproduce the region of stronger

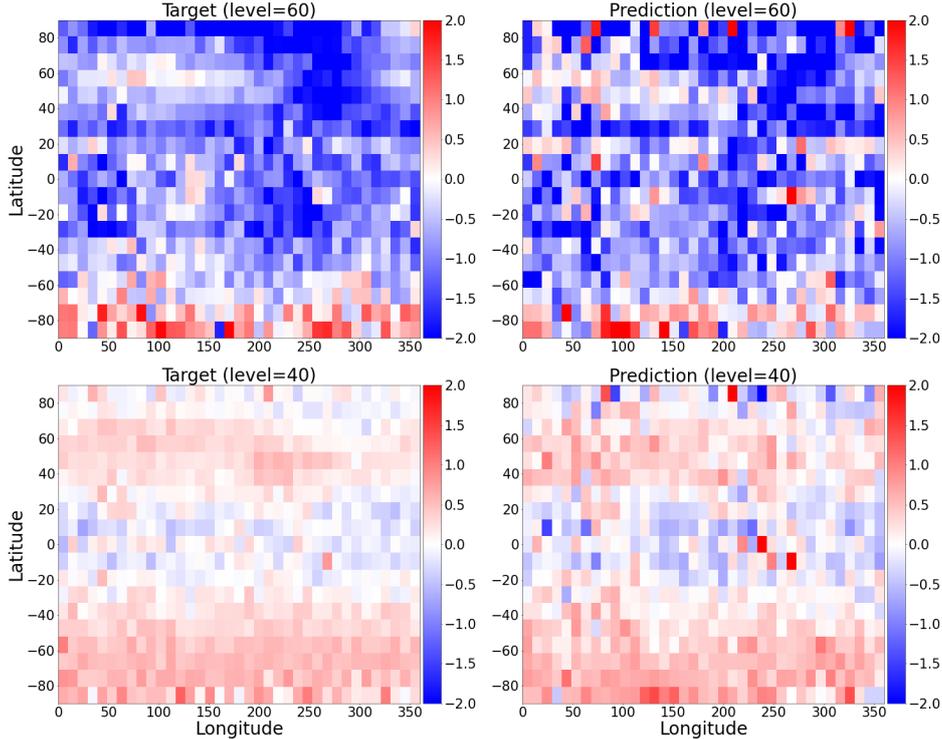


Figure 6. Target bias (left) and bias prediction from the NN model (right) for two levels on February 6 2020.

487 negative bias near (60° lat, 275° lon) as well as in the latitude band between 20° and
 488 40° .

489 **4.2 Performance of 3D models on ERA5 and operational datasets**

490 In this section, we compare three test cases, each using the 3d convolutional
 491 architecture: (i) the original ERA5 model evaluated on the ERA5 test set, (ii) the
 492 original ERA5 model evaluated on the operational test set, and (iii) the model retrained
 493 on both ERA5 and operational training data, and evaluated on the operational test
 494 set. Note that test cases (ii) and (iii) correspond to training scenarios 3 and 4 from
 495 section 3.5.

496 Figure 7 shows a vertical profile of the globally-averaged R2 scores for each of
 497 the three test cases. We observe a steep drop in prediction quality when testing on
 498 the operational dataset. The retrained model produces a better prediction than the
 499 original ERA5-only model on the operational dataset, except for the top-most levels.
 500 Comparing the model biases from the ERA5 and operational datasets (panel b and d
 501 in Figure4), we see that the retrained model struggles to capture the larger warm bias
 502 in the top levels of the operational dataset.

503 Figure 8 shows the target and predicted biases for the original ERA5 model on the
 504 ERA5 targets and the retrained model on the operation targets. Before April 2020,
 505 both the original and retrained predictions correctly capture the patterns observed in
 506 the ERA5 targets, although the predictions have a somewhat larger amplitude than
 507 the target values. After April 2020, we see that the operational target values differ
 508 significantly, where the Northern hemisphere is nearly bias free and the Southern hemi-

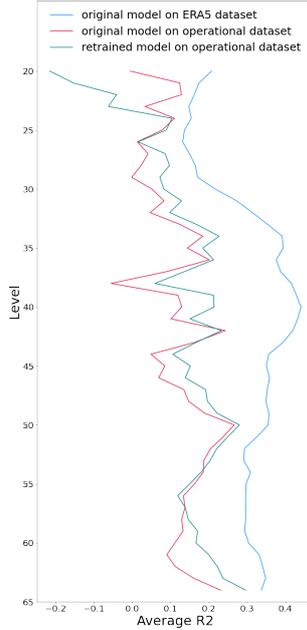


Figure 7. Vertical profile of globally averaged R2 values for the original model on the ERA5 target (blue), for the same model on the operational target (red) and for a retrained model using the sample balancing technique described above, on the operational target (green).

509 sphere exhibits a cold bias that was not present for the ERA5 target set. This feature
 510 is not capture by the NN prediction and may explain the much of the performance
 511 drop for these models in this time window.

512 4.3 NN bias correction in 4D-Var data assimilation

513 The bias predicted by our NNs can, in principle, be used to correct the model
 514 tendencies of the IFS within data assimilation experiments, in order to produce a better
 515 analysis. Unfortunately, it is technically challengng to introduce the bias correction
 516 tools into the workflow of the 4D-Var data assimilation experiments. Not only is it
 517 difficult to couple the machine learning tools with the IFS workflow, using our NN bias
 518 correction models to correct the IFS tendencies also requires one to re-gridding the
 519 model fields from the reduced-gaussian model grid of the IFS to the regular gaussian
 520 grid at the coarse resolution used to predict the bias. It is therefore beyond the scope
 521 of this paper to perform "online" simulations that calculate and correct the bias within
 522 4D-Var experiments.

523 However, we are able to predict the model bias "offline" using the retrained
 524 NN on the first-guess trajectories contained in the operational test dataset. We can
 525 run a 4D-Var experiment where the model is corrected with the respective offline
 526 correction valid for the same date. The correction is applied as an integrated term
 527 between each model timestep. Using this framework, the machine learning approach is
 528 evaluated in 4D-Var over the test period between 1st January 2020 and 1st March 2020.
 529 Figure 9 shows the first-guess mean error with respect to RO temperature retrievals
 530 for different 4D-Var experiments. The red line is the control experiment, where the
 531 dynamical model is not corrected. The dotted blue line shows the first-guess mean
 532 error, where the dynamical model is corrected using the actual target from the RO
 533 datasets. This provides a estimate of how much the bias could actually be reduced if

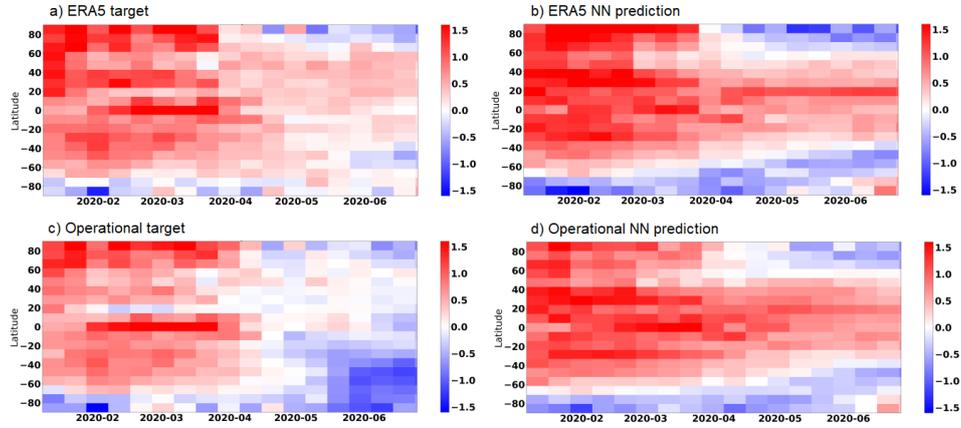


Figure 8. Zonal timeseries of the ERA5 targets (a) and the bias predictions from the original NN for model level 25 (6hPa). The bottom panels show similar timeseries for the operational targets (c) and the bias predictions from the retrained NN (d)

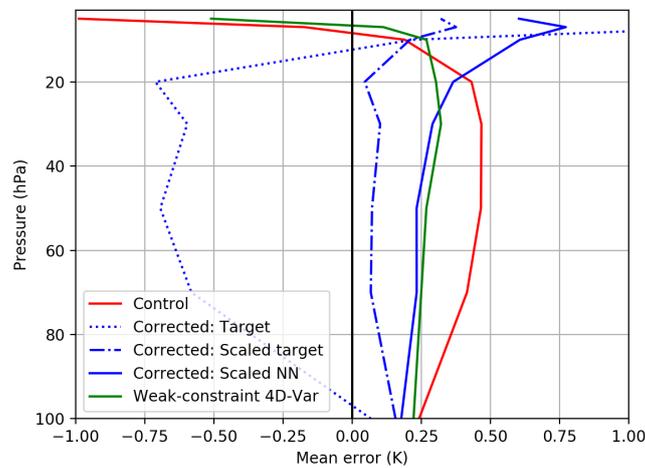


Figure 9. First-guess mean error with respect to RO temperature retrievals for the control (red), for weak-constraint 4D-Var (green), for the model corrected with the target (dotted blue), with the scaled target (dash-dot blue) and with the scaled prediction of the NN (solid blue). Statistics are averaged over the globe between 01/01/2020 and 01/03/2020.

534 the NNs provided a perfect fit to the training data. One can see that the model bias is
 535 over-corrected, for example, the original 0.4C cold bias at 60hPa became a -0.6C warm
 536 bias. This over-correction is due to the intrinsic cycling principle of data assimilation
 537 where the analysis valid at the beginning of the assimilation window is integrated
 538 forward in time to produce the background at the beginning of the next assimilation
 539 window. The first-guess departures used to diagnose the model bias contain not only
 540 the bias that develops over a single assimilation window but also includes the bias
 541 accumulated over the previous assimilation cycles contained in the background. A
 542 study of the background and analysis departures shows that only one quarter of the
 543 total bias comes from the current assimilation cycle while the other three quarters
 544 are carried forward in time from the previous cycle (e.g. at 50hPa, the global-mean
 545 analysis departure is equal to 0.38 and the global-mean background departure is equal
 546 to 0.5). Following these findings, another 4D-Var experiment was run where the model
 547 is corrected by a scaled target where all the values are reduced by a factor of 4. This
 548 approach, plotted in dash-dot blue, is able to efficiently correct the model bias for the
 549 entire stratosphere. The last experiment plotted in solid blue shows the results when
 550 the model is corrected by the offline NN predictions with the same 1/4 scaling. The
 551 NN is able to capture and correct a large fraction of the actual model bias. The first-
 552 guess mean error is reduced by almost 0.2C in the mid/lower stratosphere. The poor
 553 performance around 5hPa where the model is over-corrected is likely due to the small
 554 size of the operational dataset. The ERA5 warm bias at 5hPa is well captured by the
 555 initial NN (comparing left and right top plots in Figure 8). However, the operational
 556 model presents a smaller bias that is not well represented in the NN, which retains too
 557 much of the structure learned from ERA5. This means that the NN will cool the top
 558 of the atmosphere too aggressively, over-correcting the model warm bias.

559 5 Comparison with weak-constraint 4D-Var

Weak-constraint 4D-Var has been introduced by several authors to denote a fam-
 ily of algorithms which relax the perfect model assumption (Wergen, 1992; Zupanski,
 1993; Bennett et al., 1996; Vidard et al., 2004; Dee, 2005; Trémolet, 2006). In the
 forcing formulation of weak-constraint 4D-Var (Trémolet, 2006) a forcing is estimated
 and then applied in the model’s equations to represent the error which gradually en-
 ters into the model trajectory. The model is then treated in the same manner as
 other sources of information, taking into account that there is a degree of uncertainty
 about the information it can provide on the evolution of the atmospheric state over the
 analysis cycle. Mathematically, the weak-constraint 4D-Var formulation that has been
 implemented at ECMWF introduces a forcing η to represent the error which gradually
 enters into the model trajectory

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

The model error forcing is assumed to be additive and constant within the 12-hour
 assimilation window (Laloyaux, Bonavita, Dahoui, et al., 2020; Laloyaux, Bonavita,
 Chrust, & Gürol, 2020). It contains temperature, vorticity and divergence. We also
 assume that the model error η follows a Gaussian distribution with no cross-correlation
 with the background error. This is justified if we assume that the model error that
 we want to estimate and the background errors act on different spatial and temporal
 scales. This set of assumptions allows one to write the weak-constraint 4D-Var cost
 function as

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

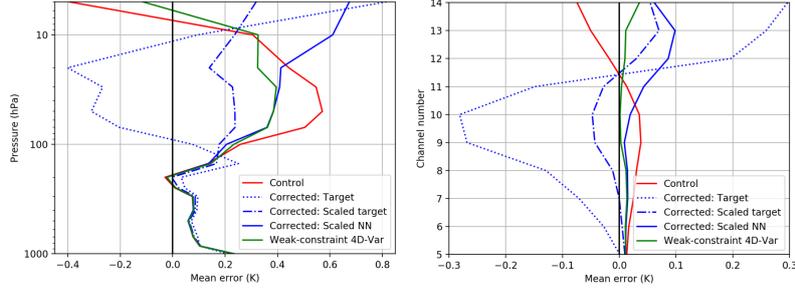


Figure 10. Same as Figure 9 but respect to radiosondes (left) and AMSU-A (right).

560 where η^b is the prior estimate of the model forcing estimated from the previous as-
 561 simulation cycle. This forcing formulation of weak-constraint 4D-Var simultaneously
 562 estimates the initial state \mathbf{x}_0 and model forcing η that best fit the observations and
 563 the background information with respect to their error covariance matrices.

564 A weak-constraint 4D-Var experiment was run from the 1st of January 2020 and
 565 is presented in green in Figure 9. In this experiment, the model error estimate is set to
 566 zero initially as no a priori knowledge of the model error is assumed. After processing
 567 only one month of data, weak-constraint 4D-Var is able to correct for a significant
 568 fraction of the model bias without requiring the computation of a large dataset for
 569 offline training.

570 Weak-constraint 4D-Var can be seen as a specific machine learning algorithms
 571 that learns the model error by estimating the parameters in the forcing vector (Farchi,
 572 Bocquet, et al., 2021). However, there are several conceptual differences with the
 573 machine learning approach described in Section 4. Weak-constraint 4D-Var is an online
 574 learning algorithm which simultaneously estimates the model state and the model
 575 error while the NN approach is estimating the model error offline before estimating
 576 the model state. An online NN is feasible but it would require a stronger interaction
 577 between NN tools and the IFS model to exchange data at each assimilation cycle.
 578 This work would require a substantial effort, given the current software infrastructure.
 579 Another difference is the amount of information used to estimate the model bias.
 580 Weak-constraint 4D-Var uses the information from all observations (conventional and
 581 satellites) as all of these are actively assimilated thanks to the radiative transfer scheme
 582 included in the 4D-Var cost function. The NN has learned the model bias using only
 583 the RO temperature retrievals which represents a small subset of the whole observing
 584 system. It is therefore interesting to study how the two approaches will fit other
 585 conventional and satellite instruments. Figure 10 shows the first-guess mean error
 586 with respect to radiosondes (left) and AMSU-A (right). In the weak-constraint 4D-Var
 587 experiment, these observations have been actively used in the observation term of the
 588 cost function (see Equation 3). This means that the data assimilation algorithm finds
 589 the optimal state that fits all of the observations with respect to their uncertainties.
 590 In the NN approach, only RO retrievals have been used to estimate the model bias as
 591 radiosondes and AMSU-A observations have not been introduced during the training.
 592 The scaled NN shows a similar improvement than weak-constraint 4D-Var in the lower
 593 and mid stratosphere (i.e. radiosondes below 30hPa and AMSU-A channel numbers
 594 below 11). This is an excellent news that can possibly be explained by the fact that
 595 RO, radiosonde and AMSU-A observing systems are consistent between each others,
 596 highlighting a similar model bias. We chose to illustrate this point using AMSU-A
 597 observations, but a similar conclusion can be drawn for other microwave instruments
 598 (e.g. ATMS) or infrared instruments (e.g. AIRS or Cris). The performance of the NN
 599 approach is not as good in the upper stratosphere (i.e. radiosondes above 30hPa and

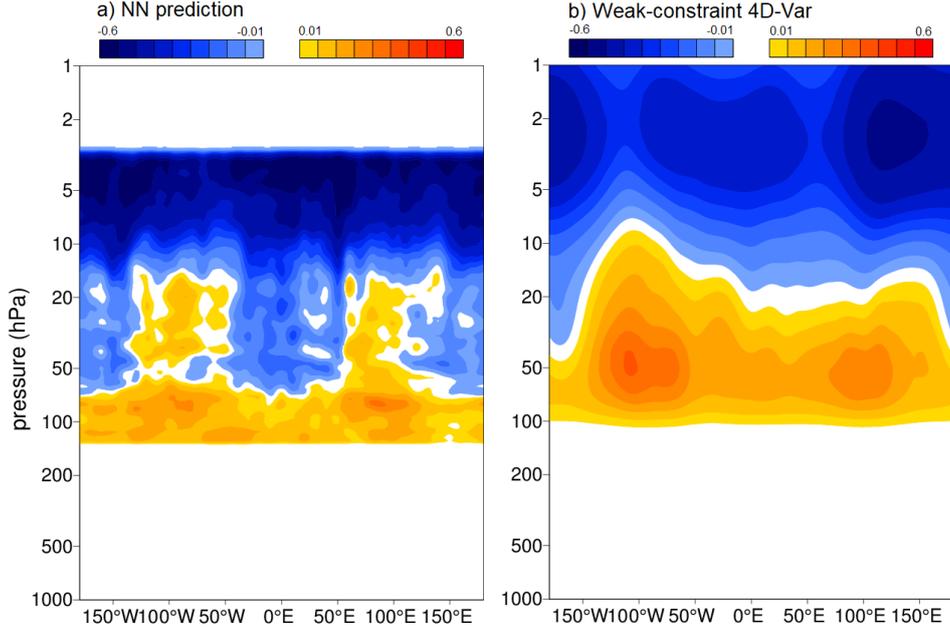


Figure 11. Meridional cross-section temperature error correction from the NN prediction (a) and from weak-constraint 4D-Var (b) averaged over the tropics (10N-10S) between 1st January 2020 and 1st March 2020.

600 AMSU-A channel numbers above 11) which confirms what has been noticed in Figure 8
 601 against RO retrievals. We have run 10-day forecasts initialised with weak-constraint
 602 4D-Var and NN analyses to study the impact on medium-range weather forecasting.
 603 We found that signals in the analysis are retained throughout the forecast and are still
 604 present after five days which confirms the results presented in Laloyaux, Bonavita,
 605 Dahoui, et al. (2020). In the lower and mid stratosphere, weak-constraint 4D-Var and
 606 NN forecasts show similar improvements at day five. The only difference happens in
 607 the upper stratosphere where the NN forecasts are degraded due to the poorer quality
 608 of the NN analysis above 20hPa (see Figure 9 and Figure 10).

609 Developing methods that estimate model biases should eventually help modellers
 610 improve their models by providing more complete knowledge of the bias structure. This
 611 will fulfil the synergies between better observations, sophisticated DA algorithms and
 612 improved physical models. It is therefore informative to study the model biases high-
 613 lighted by both approaches. Figure 11 shows a meridional cross-section temperature
 614 error correction from the NN prediction (left) and from weak-constraint 4D-Var (right)
 615 averaged over the tropics (10N-10S) between 1st January 2020 and 1st March 2020.
 616 Both approaches warm up the atmosphere over areas of strong convection (e.g. In-
 617 donesia and Southern America). The weak-constraint 4D-Var model error estimate is
 618 smoother, due to the specification of the model error covariance matrix Q which retains
 619 only large-scale patterns. This could be linked to an insufficient representation of the
 620 effects of sub-gridscale gravity wave activity, which leads to missing momentum from
 621 the troposphere to the stratosphere (Plichtchouk et al., 2019). The NN prediction is
 622 also larger for the top of the stratosphere compared the the weak-constraint 4D-Var
 623 correction. This larger NN correction is the reason for the degradation observed in
 624 Figure 9 and 10 for the top of the stratosphere.

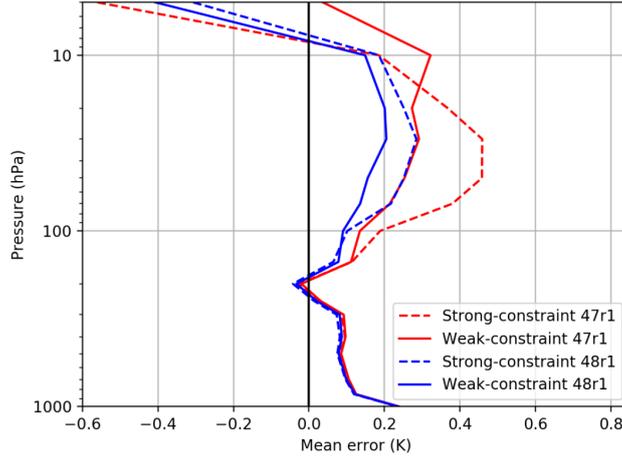


Figure 12. Vertical profile of first-guess departure with respect to radiosondes for 47r1 strong-constraint (dashed red), for 48r1 strong-constraint (dashed blue), for 47r1 weak-constraint (solid red) and for 48r1 strong-constraint (solid blue). Statistics are averaged over the globe between 20/01/2020 and 20/02/2020.

6 Weak-constraint 4D-Var learning rate

We already discussed the question of retraining the NN in Section 4 as new IFS models are made available on a regular basis with improved dynamical and physical processes of the atmosphere. We have shown that this is a challenge for the NN as the training dataset with the new model is usually relatively small (less than a year) as it is expensive to run the assimilation system for a longer period. We illustrate here how weak-constraint 4D-Var handles model upgrades using, as an example, the package of changes that is currently being tested as a possibility for the implementation of the next cycle (tentative 48r1). It contains the hybrid linear ozone, the semi-lagrangian vertical filter and a new solar spectrum. The impact of these model changes is assessed in the strong-constraint 4D-Var formulation where no model bias correction is computed. This allows one to accurately quantify how much the model upgrade reduces the model bias. Figure 12 shows the vertical profile of first-guess departure with respect to radiosondes for strong-constraint experiments with 47r1 (in dashed red) and tentative 48r1 model (in dashed blue). The improvements proposed for 48r1 significantly reduce the stratospheric model biases. At 50hPa, the original bias of 0.45 is brought down to 0.2. Weak-constraint 4D-Var aims to correct the residual model bias. The dashed red and blue lines in Figure 12 show the results of weak-constraint 4D-Var with the 47r1 and 48r1 model respectively. Although the structure of the bias is different for the two models, weak-constraint 4D-Var reduces the first-guess mean error in both situations. The weak-constraint 4D-Var cost function depends on a number of parameters that are estimated offline (e.g. standard deviation and correlation in Q). It is important to note that these parameters have not been retuned in the experiments. This shows the robustness of weak-constraint 4D-Var and its fast learning rate.

The initialisation of the model error correction at the beginning of an experiment can be compared to the challenge of initialising the weights of a NN. The middle panel in Figure 13 shows a timeseries of the model bias correction with the tentative 48r1 model when weak-constraint 4D-Var has been cold started (i.e. setting the model error correction to zero at the beginning of the experiment). It takes a couple of weeks for the model errors estimate to be properly spun-up. This is mainly because weak-constraint 4D-Var aims to correct model biases that are evolving slowly over time. To

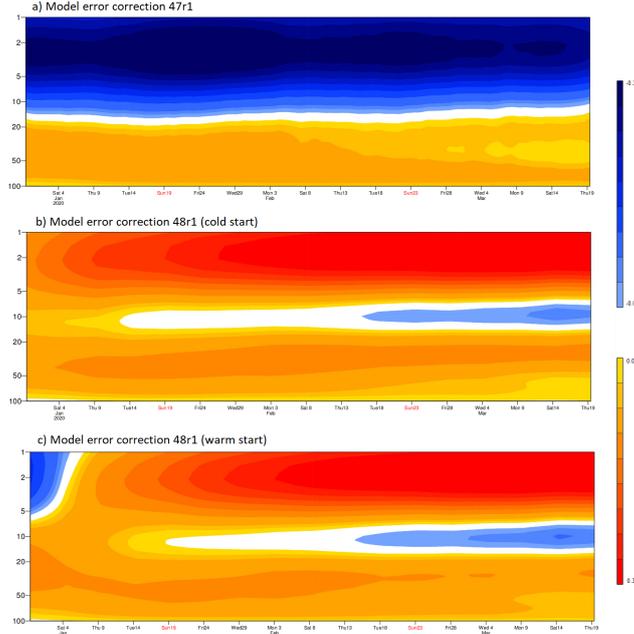


Figure 13. Timeseries of model error corrections estimated by weak-constraint 4D-Var for the 47r1 model (top), for the 48r1 model initialised from zero (middle) and for the 48r1 model initialised from the 47r1 bias estimate (bottom). Statistics are averaged between 70S and 30S.

656 study the sensitivity of the initialisation, a weak-constraint 4D-Var experiment was
 657 run where the model error correction is initialised from the previous 47r1 model
 658 estimate. This timeseries is presented at the bottom panel in Figure 13 and shows that
 659 weak-constraint 4D-Var converges towards the same solution although the behaviour
 660 is different during the spin-up period. This is a reassuring result demonstrating that
 661 weak-constraint 4D-Var is not very sensitive to the way it has been initialised. This
 662 can be explained by having a number of observations assimilated in weak-constraint
 663 4D-Var and the model error covariance Q , which are sufficient to constrain the model
 664 error correction.

665 Finally, it is important to understand how efficiently the model bias can be
 666 estimated during extreme events. The stratospheric sudden warming (SSW) is the
 667 most dramatic meteorological phenomenon to take place in the stratosphere, usually
 668 occurring over the north pole. As the temperature drops during winter, low-pressure
 669 (cyclonic) circulation begins to develop across the polar stratosphere. A strong polar
 670 vortex usually means strong polar circulation even at the lower levels. It can lock the
 671 cold air into the Polar regions, resulting in milder winters for most of the United States
 672 and Europe. If this vortex is disturbed, the winds can reverse and the temperature
 673 can rapidly increase by up to 50 degrees Celsius over a few days, in the vertical region
 674 between 1hPa and 10hPa. This can create a chain reaction, which can disrupt the
 675 jet stream, creating a high-pressure area over the Arctic circle. This, in turn, can
 676 release the cold arctic air into Europe and the United States (Polichtchouk et al., 2018;
 677 Mariotti et al., 2020). SSWs happen every-other year or so, with the most recently
 678 event recorded in January 2021. The top panel of Figure 14 shows a timeseries of
 679 first-guess departure with respect to RO temperature retrievals, averaged over the
 680 Northern pole (70N 90N) between September 24, 2020 and February 24, 2021. At the
 681 beginning of the SSW event (1st of January 2021), the structure of the model bias

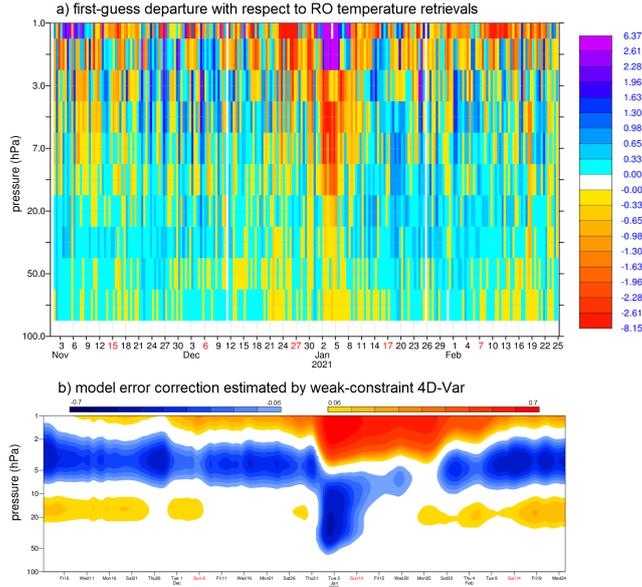


Figure 14. Timeseries of first-guess departure with respect to RO temperature retrievals (top) and timeseries of model error correction estimated by weak-constraint 4D-Var (bottom). Statistics are averaged over the Northern pole (70N 90N) between 24th September 2020 and 24th February 2021.

682 changes significantly as stratospheric dynamics are disrupted. There is a model cold
 683 bias above 3 hPa and a model warm bias between 50 hPa and 3 hPa. The bottom panel
 684 of Figure 14 shows the model error correction estimated by weak-constraint 4D-Var.
 685 The model bias change is captured quickly as weak-constraint 4D-Var warms up the
 686 stratosphere above 3hPa and cools down between 50hPa and 3hPa. This illustrates
 687 the efficient learning rate of weak-constraint 4D-Var when an extreme event occurs.
 688 A similar study could not be done for the NN approach as the test dataset (June
 689 2019 to June 2020) does not contain such an event. This is however a critical aspect
 690 that will be studied in the future, as extreme events occur infrequently in the training
 691 dataset and it might be challenging for the NN to correctly represent the model error
 692 structure.

693 7 Summary and perspectives

694 Artificial intelligence and machine learning are entering the domain of Earth
 695 system predictions in parallel with the development of more heterogeneous High-
 696 Performance Computing (HPC) architectures. This changing context presents new
 697 development opportunities that ECMWF is considering, with the ambition of retain-
 698 ing leadership in global medium- and extended range weather forecasting. 4D-Var
 699 data assimilation and machine learning share a common theoretical foundation and
 700 use similar computational tools. This has driven the work presented in this paper,
 701 which compares how each method is able to estimate and correct systematic errors in
 702 the IFS atmospheric model developed at ECMWF model.

703 The results of this paper show that convolutional NNs are adequate to learn to
 704 estimate three-dimensional model bias from RO temperature retrievals. While large
 705 datasets containing several years of data are required for the training to achieve optimal
 706 results, transfer learning can help to mitigate data limitations if only a small quantity

707 of training data is available. Still, when used to perform bias correction in data
 708 assimilation experiments for a recent IFS model cycle and with a single year of training
 709 data for re-training, the deep learning tools of this paper were not able to outperform
 710 the current weak-constraint 4D-Var formulation that is in operational use at ECMWF.

711 However, direct comparison between the two methods has one main limitation.
 712 Weak-constraint 4D-Var can be seen as an "online" machine learning method, where
 713 observations over the last 12 hours are used to update the previous weather forecasts.
 714 The machine learning tool of this paper is based on an "offline" training. Further-
 715 more, the deep learning bias correction was computed "offline" before the assimilation
 716 experiment was started. It is difficult to estimate how much results would change if
 717 an update of the bias correction was calculated during the assimilation experiment
 718 which is – for technical reasons – beyond the scope of this paper. Another difference
 719 between the two approaches lies in the physical variables that are corrected. Weak-
 720 constraint 4D-Var estimates a forcing field for temperature, vorticity and divergence.
 721 Although very few stratospheric wind observations are available, these variables are
 722 linked through the model's equation in the 4D-Var cost function. This means that
 723 wind corrections are made in conjunction with temperature adjustments. The NN
 724 approach corrected only temperature biases. The weak-constraint 4D-Var also in-
 725 cludes a model error covariance matrix Q that represents separately the statistics of
 726 the model error for temperature, vorticity and divergence. Cross-correlation between
 727 variables are not taken into account at the moment. Diagnostics in the IFS model
 728 show that the stratospheric temperature model biases evolve on larger spatial scales
 729 and longer timescales than background errors (Laloyaux, Bonavita, Chrust, & Gürol,
 730 2020). This information is contained in the Q matrix and helps weak-constraint 4D-
 731 Var to correctly attribute the different sources of errors. A similar approach could be
 732 investigated in the NN approach introducing a similar regularization term in the loss
 733 function. Finally, the jump from the model cycle used in ERA5 and in operations as
 734 performed in this paper represents a significant change in the temperature bias as it
 735 represents a transition over several years of model development.

736 The deep learning approach has room for improvement, for example by extending
 737 the dataset to encompass more observation types. However, this is challenging as most
 738 observations do not measure model prognostic variables on a given grid point but a
 739 radiance that is sensitive to a broad vertical level. The development of machine learned
 740 observation operators to project observations onto model fields would be mandatory.
 741 The use of deep learning methods could also be extended further to include estimates of
 742 background and observation error covariance matrices, and to represent uncertainties
 743 explicitly, for example via Generative Adversarial Networks (Leinonen et al. (2021)).
 744 The treatment for sparsity observations could also be improved further, for example
 745 via the use of graph-NNs, which could evaluate observations at the points in space and
 746 time when they are available, and even respect spherical symmetry of the globe (cf. e.g.
 747 Defferrard et al. (2020)). Graph-NNs would also allow for the use of unstructured grids
 748 potentially including the native grid of the IFS and could better exploit the sparsity of
 749 the data by replacing the interpolation step with a NN based extrapolation. An online
 750 NN could be implemented in the future to study the full potential of a ML solution
 751 in the 4D-Var framework. However, this is work in progress and will require further
 752 developments regarding software infrastructure and more research to find the best way
 753 to update NN weights in a 4D-Var cycling environment. One of the key aspects of
 754 ECMWF business is the Research-to-Operations (R2O) process, which is followed to
 755 upgrade the software used in forecast production (Buizza et al., 2017). R2O includes
 756 a series of actions that could be summarized in 6 activities: planning, development,
 757 testing, evaluation, communication and implementation. The IFS model is upgraded
 758 at every cycle to better represent physical processes or introduce new ones that were
 759 missing. This paper illustrated the strength of weak-constraint 4D-Var that is able to
 760 estimate the bias of a new model with no need to construct a new training dataset

761 or to retune parameters. Specific solutions are required to achieve a similar flexibility
762 with a NN.

763 8 Acknowledgements

764 PD gratefully acknowledge funding from the Royal Society for his University Re-
765 search Fellowship as well as the ESiWACE project funded under Horizon 2020 No.
766 823988 and the MAELSTROM EuroHPC-JU project (JU) funded under No 955513.
767 The JU receives support from the European Union’s Horizon research and innova-
768 tion programme and United Kingdom, Germany, Italy, Luxembourg, Switzerland, and
769 Norway.

770 9 Open research

771 Data availability statement: The input and output data of the experiments de-
772 scribed in the paper is freely available for research purposes from ECMWF and can be
773 requested following the procedures described in <https://www.ecmwf.int/en/forecasts/datasets>

774 References

- 775 Bannister, R. N. (2008a). A review of forecast error covariance statistics in at-
776 mospheric variational data assimilation. I: Characteristics and measurements of
777 forecast error covariances. *Quarterly Journal of the Royal Meteorological Society*,
778 *134*(637), 1951-1970.
- 779 Bannister, R. N. (2008b). A review of forecast error covariance statistics in atmo-
780 spheric variational data assimilation. II: Modelling the forecast error covariance
781 statistics. *Quarterly Journal of the Royal Meteorological Society*, *134*(637), 1971-
782 1996.
- 783 Bennett, F., Chua, A., & Leslie, B. (1996, 03). Generalized inversion of a global nu-
784 merical weather prediction model. *Meteorology and Atmospheric Physics*, *60*, 165-
785 178. doi: 10.1007/BF01029793
- 786 Bonavita, M., Hólm, E., Isaksen, L., & Fisher, M. (2016). The evolution of the
787 ECMWF hybrid data assimilation system. *Quarterly Journal of the Royal Meteoro-
788 logical Society*, *142*(694), 287-303.
- 789 Bonavita, M., & Laloyaux, P. (2020). Machine learning for model error inference and
790 correction. *Journal of Advances in Modeling Earth Systems*, *12*(12).
- 791 Brajard, J., Carrassi, A., Bocquet, M., & Bertino, L. (2020). Combining data assim-
792 ilation and machine learning to emulate a dynamical model from sparse and noisy
793 observations: A case study with the lorenz 96 model. *Journal of Computational
794 Science*, *44*.
- 795 Buizza, R., Andersson, E., Forbes, R., & Sleigh, M. (2017). *The ECMWF research
796 to operations (R2O) process* (Technical Memorandum No. 806). Reading, UK:
797 ECMWF.
- 798 Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., & Yuille, A. L. (2018).
799 DeepLab: Semantic image segmentation with deep convolutional nets, atrous con-
800 volution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and
801 Machine Intelligence*, *40*(4), 834-848.
- 802 Chollet, F. (2017). *Xception: Deep learning with depthwise separable convolutions*.
- 803 Cucurull, L., Derber, J. C., & Purser, R. J. (2013). A bending angle forward opera-
804 tor for global positioning system radio occultation measurements. *Journal of Geo-
805 physical Research*, *118*.
- 806 Dee, D. P. (2005). Bias and data assimilation. *Quarterly Journal of the Royal Mete-
807 orological Society*, *131*(613), 3323-3343.
- 808 Dee, D. P., & Uppala, S. (2009). Variational bias correction of satellite radiance data

- 809 in the ERA-Interim reanalysis. *Quarterly Journal of the Royal Meteorological So-*
810 *ciety*, 135(644), 1830–1841.
- 811 Defferrard, M., Milani, M., Gusset, F., & Perraudin, N. (2020). *Deepsphere: a graph-*
812 *based spherical cnn*.
- 813 Dueben, P., Modigliani, U., Geer, A., Siemen, S., Pappenberger, F., Bauer, P., ...
814 Baousis, V. (2021). *Machine learning at ECMWF: A roadmap for the next 10*
815 *years* (Technical Memorandum No. 878). Reading, UK: ECMWF.
- 816 Eyre, J. R. (1994). *Assimilation of radio occultation measurements into a numer-*
817 *ical weather prediction system* (Technical Memorandum No. 199). Reading, UK:
818 ECMWF.
- 819 Farchi, A., Bocquet, M., Laloyaux, P., Bonavita, M., & Malartic, Q. (2021). *A com-*
820 *parison of combined data assimilation and machine learning methods for offline*
821 *and online model error correction*.
- 822 Farchi, A., Laloyaux, P., Bonavita, M., & Bocquet, M. (2021). Using machine learn-
823 ing to correct model error in data assimilation and forecast applications. *Quarterly*
824 *Journal of the Royal Meteorological Society*, 147(739), 3067-3084.
- 825 Fomichev, V. I., Ward, W. E., Beagley, S. R., McLandress, C., McConnell, J. C.,
826 McFarlane, N. A., & Shepherd, T. G. (2002). Extended canadian middle atmo-
827 sphere model: Zonal-mean climatology and physical parameterizations. *Journal of*
828 *Geophysical Research: Atmospheres*, 107(D10).
- 829 Geer, A. (2020). *Learning earth system models from observations: machine learning*
830 *or data assimilation?* (Technical Memorandum No. 863). Reading, UK: ECMWF.
- 831 Geer, A. J., Lonitz, K., Weston, P., Kazumori, M., Okamoto, K., Zhu, Y., ...
832 Schraff, C. (2018). All-sky satellite data assimilation at operational weather fore-
833 casting centres. *Quarterly Journal of the Royal Meteorological Society*, 144(713),
834 1191-1217.
- 835 Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. Cambridge: MIT
836 Press.
- 837 Groenquist, P., Yao, C., Ben-Nun, T., Dryden, N., Dueben, P., Li, S., & Hoefler, T.
838 (2021). Deep learning for post-processing ensemble weather forecasts. *Philosoph-*
839 *ical Transactions of the Royal Society A: Mathematical, Physical and Engineering*
840 *Sciences*, 379(2194).
- 841 Haiden, T., Dahoui, M., Ingleby, B., de Rosnay, P., Prates, C., Kuscu, E., ... Jones,
842 L. (2018). *Use of in situ surface observations at ecmwf* (Technical Memorandum
843 No. 834). Reading, UK: ECMWF.
- 844 He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image
845 recognition. In *2016 IEEE conference on computer vision and pattern recognition*
846 *(CVPR)* (p. 770-778). doi: 10.1109/CVPR.2016.90
- 847 Healy, S. B., & Thépaut, J.-N. (2006). Assimilation experiments with CHAMP
848 GPS radio occultation measurements. *Quarterly Journal of the Royal Meteorologi-*
849 *cal Society*, 132, 605–623.
- 850 Hogan, R., Ahlgrimm, M., Balsamo, G., Beljaars, A., Berrisford, P., Bozzo, A., ...
851 Wedi, N. (2017). *Radiation in numerical weather prediction* (Technical Memoran-
852 dum No. 816). Reading, UK: ECMWF.
- 853 *Hyperopt website*. (2021, 11). Retrieved from <http://hyperopt.github.io/hyperopt/>
- 854 *hyperopt/*
- 855 Janjic, T., Bormann, N., Bocquet, M., Carton, J. A., Cohn, S. E., Dance, S. L., ...
856 Weston, P. (2018). On the representation error in data assimilation. *Quarterly*
857 *Journal of the Royal Meteorological Society*, 144(713), 1257-1278.
- 858 Kinoshita, Y., & Kiya, H. (2020). Fixed smooth convolutional layer for avoiding
859 checkerboard artifacts in cnns. In *Icassp 2020 - 2020 IEEE international conference*
860 *on acoustics, speech and signal processing (icassp)* (p. 3712-3716).
- 861 Kursinski, E., Hajj, G., Schofield, J., Linfield, R., & Hardy, K. (1997). Observ-
862 ing earth's atmosphere with radio occultation measurements using the Global

- 863 Positioning System. *Journal of Geophysical Research*, *102*, 23.429–23.465.
- 864 Kurth, T., Treichler, S., Romero, J., Mudigonda, M., Luehr, N., Phillips, E. H.,
865 ... Houston, M. (2018). Exascale deep learning for climate analytics. *CoRR*,
866 *abs/1810.01993*.
- 867 Laloyaux, P., Bonavita, M., Chrust, M., & Gürol, S. (2020). Exploring the poten-
868 tial and limitations of weak-constraint 4d-var. *Quarterly Journal of the Royal Me-*
869 *teorological Society*, *146*(733), 4067-4082.
- 870 Laloyaux, P., Bonavita, M., Dahoui, M., Farnan, J., Healy, S., Holm, E., & Lang,
871 S. T. K. (2020). Towards an unbiased stratospheric analysis. *Quarterly Journal of*
872 *the Royal Meteorological Society*, *146*(730), 2392-2409.
- 873 Leinonen, J., Nerini, D., & Berne, A. (2021). Stochastic super-resolution for down-
874 scaling time-evolving atmospheric fields with a generative adversarial network.
875 *IEEE Transactions on Geoscience and Remote Sensing*, *59*(9), 7211-7223.
- 876 Loshchilov, I., & Hutter, F. (2019). *Decoupled weight decay regularization*.
- 877 Mariotti, A., Baggett, C., Barnes, E. A., Becker, E., Butler, A., Collins, D. C., ...
878 Albers, J. (2020). Windows of opportunity for skillful forecasts subseasonal to
879 seasonal and beyond. *Bulletin of the American Meteorological Society*, *101*(5),
880 E608 - E625.
- 881 *Mlperf hpc deepcam website*. (2021, 11). Retrieved from [https://github.com/](https://github.com/mlcommons/hpc/tree/main/deepcam)
882 [mlcommons/hpc/tree/main/deepcam](https://github.com/mlcommons/hpc/tree/main/deepcam)
- 883 Poli, P., Moll, P., Puech, D., Rabier, F., & Healy, S. B. (2009). Quality control,
884 error analysis, and impact assessment of FORMOSAT-3/COSMIC in numerical
885 weather prediction. *Terrestrial, Atmospheric and Oceanic Sciences*, *20*, 101–113.
- 886 Polichtchouk, I., Shepherd, T. G., Hogan, R. J., & Bechtold, P. (2018). Sensitivity
887 of the brewer–dobson circulation and polar vortex variability to parameterized
888 nonorographic gravity wave drag in a high-resolution atmospheric model. *Journal*
889 *of the Atmospheric Sciences*, *75*(5), 1525-1543.
- 890 Polichtchouk, I., Stockdale, T., Bechtold, P., Diamantakis, M., Malardel, S., Sandu,
891 I., ... Wedi, N. (2019). *Control on stratospheric temperature in IFS: resolution*
892 *and vertical advection* (Technical Memorenda No. 847). Shinfield Park, Reading
893 RG2 9AX, United Kingdom: ECMWF.
- 894 Rabier, F., Järvinen, H., Klinker, E., Mahfouf, J.-F., & Simmons, A. (2000). The
895 ECMWF operational implementation of four-dimensional variational assimilation.
896 i: Experimental results with simplified physics. *Quarterly Journal of the Royal*
897 *Meteorological Society*, *126*(564), 1143-1170.
- 898 Rasp, S., & Lerch, S. (2018). Neural networks for postprocessing ensemble weather
899 forecasts. *Monthly Weather Review*, *146*(11).
- 900 *Ray tune website*. (2021, 11). Retrieved from <https://www.ray.io/ray-tune>
- 901 Rennie, M. P. (2010). The impact of GPS radio occultation assimilation at the Met
902 Office. *Quarterly Journal of the Royal Meteorological Society*, *136*, 116–131.doi:
903 10.1002/qj.521.
- 904 Ronneberger, O., Fischer, P., & Brox, T. (2015). *U-net: Convolutional networks for*
905 *biomedical image segmentation*.
- 906 Saunders, R. (2021). The use of satellite data in numerical weather prediction.
907 *Weather*, *76*(3), 95-97.
- 908 Schultz, M. G., Betancourt, C., Gong, B., Kleinert, F., Langguth, M., Leufen, L. H.,
909 ... Stadler, S. (2021). Can deep learning beat numerical weather prediction?
910 *Philosophical Transactions of the Royal Society A: Mathematical, Physical and*
911 *Engineering Sciences*, *379*(2194).
- 912 Shepherd, T., Polichtchouk, I., Hogan, R., & Simmons, A. (2018). *Report on strato-*
913 *sphere task force* (Technical Memorenda No. 824). Shinfield Park, Reading RG2
914 9AX, United Kingdom: ECMWF.
- 915 Sun, B., Reale, A., Seidel, D. J., & Hunt, D. C. (2010). Comparing radiosonde
916 and cosmic atmospheric profile data to quantify differences among radiosonde

- 917 types and the effects of imperfect collocation on comparison statistics. *Journal of*
918 *Geophysical Research: Atmospheres*, 115(D23).
- 919 Sun, B., Reale, T., Schroeder, S., Pettey, M., & Smith, R. (2019). On the accuracy
920 of vaisala rs41 versus rs92 upper-air temperature observations. *Journal of Atmo-*
921 *spheric and Oceanic Technology*, 36(4), 635 - 653.
- 922 Trémolet, Y. (2006). Accounting for an imperfect model in 4D-Var. *Quarterly Jour-*
923 *nal of the Royal Meteorological Society*, 132(621), 2483-2504.
- 924 Vidard, P., Piacentini, A., & Dimet, F.-X. L. (2004). Variational data analysis with
925 control of the forecast bias. *Tellus A*, 56(3), 177-188.
- 926 Vorobev, V. V., & Krasilnikova, T. G. (1994). Estimation of the accuracy of the
927 atmospheric refractive index recovery from doppler shift measurements at frequen-
928 cies used in the NAVSTAR system. *USSR Phys. Atmos. Ocean, Engl. Transl.*, 29,
929 602-609.
- 930 Waller, J. A., Dance, S. L., Lawless, A. S., & Nichols, N. K. (2014). Estimating cor-
931 related observation error statistics using an ensemble transform kalman filter. *Tel-*
932 *lus A: Dynamic Meteorology and Oceanography*, 66(1), 23294.
- 933 Watson, P. A. G. (2019). Applying machine learning to improve simulations of a
934 chaotic dynamical system using empirical error correction. *Journal of Advances in*
935 *Modeling Earth Systems*, 11(5), 1402-1417.
- 936 Wergen, W. (1992). The effect of model errors in variational assimilation. *Tellus A:*
937 *Dynamic Meteorology and Oceanography*, 44(4), 297-313.
- 938 Zupanski, M. (1993, 08). Regional four-dimensional variational data assimilation in
939 a quasi-operational forecasting environment. *Monthly Weather Review*, 121, 2396-
940 2408.