Geoscientific
Model Development

Model experiment description paper

Representing chemical history in ozone time-series predictions – a model experiment study building on the MLAir (v1.5) deep learning framework

Felix Kleinert[1,2], **Lukas H. Leufen**[1,2], **Aurelia Lupascu**[3], **Tim Butler**[3], **and Martin G. Schultz**[1]

[1]Forschungszentrum Jülich GmbH, Jülich Supercomputing Centre (JSC), Jülich, Germany
[2]Institute of Geosciences, Rheinische Friedrich-Wilhelms-Universität Bonn, Bonn, Germany
[3]Institute for Advanced Sustainability Studies, Potsdam, Germany

Correspondence: Felix Kleinert (f.kleinert@fz-juelich.de)

Abstract. Tropospheric ozone is a secondary air pollutant that is harmful to living beings and crops. Predicting ozone concentrations at specific locations is thus important to initiate protection measures, i.e. emission reductions or warnings to the population. Ozone levels at specific locations result from emission and sink processes, mixing and chemical transformation along an air parcel's trajectory. Current ozone forecasting systems generally rely on computationally expensive chemistry transport models (CTMs). However, recently several studies have demonstrated the potential of deep learning for this task. While a few of these studies were trained on gridded model data, most efforts focus on forecasting time series from individual measurement locations. In this study, we present a hybrid approach which is based on time-series forecasting (up to 4 d) but uses spatially aggregated meteorological and chemical data from upstream wind sectors to represent some aspects of the chemical history of air parcels arriving at the measurement location. To demonstrate the value of this additional information, we extracted pseudo-observation data for Germany from a CTM to avoid extra complications with irregularly spaced and missing data. However, our method can be extended so that it can be applied to observational time series. Using one upstream sector alone improves the forecasts by 10 % during all 4 d, while the use of three sectors improves the mean squared error (MSE) skill score by 14 % during the first 2 d of the prediction but depends on the upstream wind direction. Our method shows its best performance in the northern half of Germany for the first 2 prediction days. Based on the data's

seasonality and simulation period, we shed some light on our models' open challenges with (i) spatial structures in terms of decreasing skill scores from the northern German plain to the mountainous south and (ii) concept drifts related to an unusually cold winter season. Here we expect that the inclusion of explainable artificial intelligence methods could reveal additional insights in future versions of our model.

1   Introduction

Near-surface ozone (O_3) is a secondary air pollutant which is harmful to living beings (WHO, 2013; Fleming et al., 2018) and crops (Avnery et al., 2011; Mills et al., 2018). The first Tropospheric Ozone Assessment Report (TOAR; https://igacproject.org/activities/TOAR/TOAR-I, last access: 22 February 2022) provided the first globally consistent analysis of the global distribution and trends of tropospheric ozone. Key aspects of the assessment were, among others, changes in the tropospheric ozone burden and its budget (Archibald et al., 2020), observed long-term trends and their uncertainties (Tarasick et al., 2019), present-day tropospheric ozone distribution and trends of metrics that are relevant to health (Fleming et al., 2018), vegetation (Mills et al., 2018) and climate (Gaudel et al., 2018). Furthermore, the capabilities of current atmospheric chemistry models were reviewed (Young et al., 2018).

   Field and time-series forecasting are two examples where earth system scientists are starting to pick up deep learning

(DL) models to enhance the quality or performance of air pollution forecasts or explore novel analyses of air quality, weather and climate data. The success of these DL models is largely due to two factors: (1) improved model architectures that can capture spatio-temporal relations in the data and (2) the increasing amount of data that has become available in recent years. While new DL methods in application areas like image or speech recognition or video frame prediction are typically developed with the help of specific benchmark data sets, thus greatly accelerating the adoption of new concepts, such benchmark data sets are only now beginning to be developed for atmospheric applications. For example, Rasp et al. (2020) developed the first meteorological benchmark data set called *WeatherBench* for medium-range weather forecasting based on ERA5 data. In terms of air quality, Betancourt et al. (2021) developed a benchmark data set called *AQ-Bench* focusing on long-term ozone metrics from time-independent local features of ozone measurement sites using the TOAR database (Schultz et al., 2017).

Concerning the problem of air quality and specifically ozone forecasts, researchers have tried out different DL models for a range of lead times in hourly (for example Eslami et al., 2020; Sayeed et al., 2020) or daily (Kleinert et al., 2021) resolution. Hybrid forecasting models combining chemical transport models with DL (Sayeed et al., 2021) and a data-driven Bayesian neural network ensemble method for (numerical) geophysical models (Sengupta et al., 2020) were developed. He et al. (2022) used a deep neural network to evaluate NO_x emissions by forecasting ozone concentrations over several years and showed that their model reproduced ozone concentrations in low emission regions best when using satellite-derived NO_x trends. Sayeed et al. (2022) recently developed a DL model for post-processing by mapping ozone precursors and meteorological information from models to observed ozone concentrations at monitoring stations. The available ozone forecasting studies employed different selections of input variables and different methods to preprocess the input data in order to help the DL methods extract the most relevant information. Often, environmental scientists use their knowledge of atmospheric processes to select variables or design the preprocessing strategies. As one example out of many, a decomposition of input time series can help neural networks to learn features like seasonality much easier – especially when the amount of training data is limited (Leufen et al., 2022).

The present study focuses on the extraction of spatio-temporal features in the context of time-series predictions. We reuse the DL set-up of Kleinert et al. (2021), who developed a time-series prediction model based on an inception architecture (Szegedy et al., 2015; Ismail Fawaz et al., 2020). That model was trained on daily aggregated ozone data from more than 300 German air quality measurement stations for a lead time of up to 4 d. In contrast to this earlier study, the present work uses a U-shaped architecture and is based on simulation output from the chemical transport model (CTM) WRF-Chem (Grell et al., 2005) instead of the TOAR observations to demonstrate the added value of upstream wind sector information without having to cope with irregularly spaced and sometimes missing observations. Using WRF-Chem data also as target data y avoids representation problems when comparing gridded model data and point measurements.

Through the adoption of an aggregated upstream wind sector approach following Yi et al. (2018), we are aiming to capture the *chemical history* of air masses, i.e. the fact that air pollutant concentrations at a given observation site are a result of emission and sink processes, mixing and chemical transformations along the transport pathways of air. Yi et al. (2018) defined multiple wind sectors around a measurement station and used the spatially aggregated information as additional input for their time-series model. In our study we condense this method using one or three upstream sectors, and we identify the influence of those sectors on the air quality forecast accuracy.

In reality, the chemical history of air parcels must be expressed as a multi-dimensional integral over a wide spectrum of chemical ages (i.e. the product of loss rates and time) and including different mixing rates. Lagrangian particle models such as FLEXPART (Pisso et al., 2019) have been used to disentangle these processes and allow for attribution of air pollutant concentrations to specific source regions (see Stohl et al., 1998, 2013; Wenig et al., 2003; Yu et al., 2020; Aliaga et al., 2021). However, such simulations are not straightforward to run and would have added a lot of complexity to this study. The much simpler aggregation of input variables in one or three upstream wind sectors should at least capture the first-order effects of the air parcel history and thus add valuable information to the input data for our DL network.

This article is structured as follows: in Sect. 2 we briefly introduce the WRF-Chem model data. In Sect. 3 we introduce our preprocessing methods (Sect. 3.1) and the *Machine Learning on Air data* (MLAir) framework (Sect. 3.2) which we use for our study. In Sect. 3.3 we present the DL model architecture and training procedure, followed by a description of our reference models (Sect. 3.4). We present our results in Sect. 4 and discuss them in Sect. 5 with a special focus on the DL method's loss (Sect. 5.1), the benefits of using additional upstream information (Sect. 5.2), the sensitivity of our DL model with respect to input variables (Sect. 5.3) and the implications of an unusually cold winter season for the evaluation of our DL model (Sect. 5.4). Finally, Sect. 6 provides conclusions.

2 Data: WRF-Chem

We use numerical model data from WRF-Chem (Grell et al., 2005), version 3.9, as the foundation of our study. We carried out a simulation for 2009 and spring 2010 over Europe because a model set-up existed for this period (Galmarini et al.,

2021). The simulation domain covers 400 grid points in a west–east direction and 360 grid points in a south–north direction across 35 vertical levels from the surface to 50 hPa. This corresponds to a horizontal grid spacing of ~ 12 km. Initial and boundary conditions for the meteorological fields are taken from the ERA-Interim reanalyses (Dee et al., 2011), and the chemical initial and boundary conditions were derived from the CAMS reanalysis (Inness et al., 2019). To force the model toward the spatial and temporal analyses, the model was run with grid nudging. Single simulations were performed for 3.5-month periods, leaving out the first 15 d. Thus we ensure that the model does not deviate from the observed synoptic events, and the impact of meteorological errors on atmospheric chemistry simulations is reduced. The anthropogenic emissions data are taken from the TNO-MACC III emission inventory (Kuenen et al., 2014), and the MEGAN model (Guenther et al., 2006) was employed to estimate biogenic species emissions. As in Kuik et al. (2016), land cover classes have been updated with the CORINE data set (CLC, 2012, Copernicus Land Monitoring Service). The emissions from open burning are based on the Fire INventory from NCAR (FINNv1.5) (Wiedinmyer et al., 2011). The main physics options are described in Lupaşcu and Butler (2019). The gas-phase mechanism is MOZART (Emmons et al., 2010; Knote et al., 2014) coupled with the MOSAIC aerosol module (Zaveri et al., 2008).

To train and evaluate the deep learning model, we extracted WRF-Chem data at the locations of the air quality measurement stations of the German Umweltbundesamt as in Kleinert et al. (2021) (see also Sect. 3.1).

Winter 2009/10 in Europe was special in the sense that it was an unusually cold winter. The average winter temperature in Germany was $-1.3\,°C$, which is $1.5\,°C$ below the average winter temperatures from 1961 to 1990 (DWD, 2022). Figure 1 shows that the anomalies in northern Germany were stronger than those in southern Germany. We will discuss the implications of this anomaly for the evaluation of our DL model in Sect. 5.4.

3   Method

Data-driven machine learning applications require substantial effort to select and preprocess the data to be used. In the case of environmental data, researchers have to find a compromise between available data and the independence of data used for training, validation and testing, mostly due to limitations in data availability (Schultz et al., 2021). To incorporate the information from one or three upstream wind sectors, the preprocessing of data had to be expanded compared to Kleinert et al. (2021) and is presented in Sect. 3.1. Moreover, we briefly introduce the MLAir framework (Leufen et al., 2021), which we have used to carry out the experiments (Sect. 3.2), and the model architecture and training procedure (Sect. 3.3).
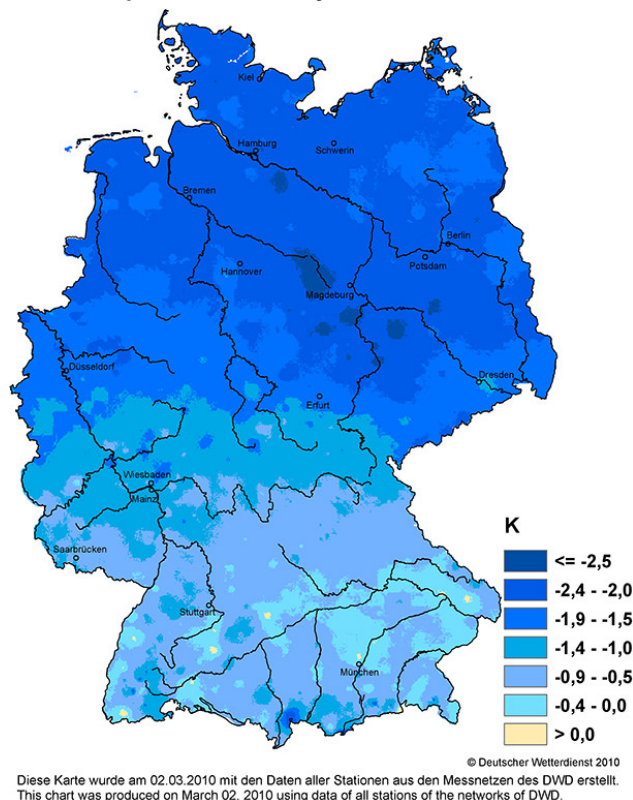
Figure 1. Temperature anomaly in winter 2009/10 with respect to the multi-year average from 1961 to 1990. Figure created by Deutscher Wetterdienst (German weather service, DWD) and downloaded from https://www.dwd.de/EN/ourservices/klimakartendeutschland/klimakartendeutschland.html?nn=495490 (last access: 1 April 2022).

3.1   Data preprocessing

The main focus of this study is to assess the added value of incorporating upstream information in the sense of a chemical air parcel history on the characteristics of air quality, i.e. ozone, forecasts with a deep neural network. For this, we define three sets of experiments: firstly, a baseline (NNb) experiment with no upstream information, secondly, a single sector (NN1s) and thirdly a multi-sector (NN3s) approach. These sets of experiments necessitate different preprocessing steps, which we introduce in the following (see also Table 3 for a summary of acronyms). As the ultimate goal of our experiments and methods is to apply them to air quality observations, we selected the locations of German air quality stations as our input and target data and applied nearest-neighbour sampling to extract the specific WRF-Chem model grid box "representing" the location of a given measurement site. From here on, we will refer to those grid boxes as pseudo-observations and pseudo-measurement stations.

All in all, we use data from 332 pseudo-measurement stations for training, validation and testing of the neural network by following Kleinert et al. (2021). Thereby, 247 grid boxes contain exactly one, 35 grid boxes contain two and five grid boxes contain three pseudo-stations, respectively. As Schultz et al. (2021) and Kleinert et al. (2021) propose, we use a consecutive data split. We train the model with data ranging from 1 January to 15 October 2009 and use a short validation period ranging from 16 October to 14 December 2009. The final test period, which we used for all results presented below, covers 16 December 2009 to 31 March 2010. This split is motivated by the fact that ozone concentrations during winter are primarily determined by transport processes as opposed to summertime, when photochemistry plays a more active role. Hence, the effects of incorporating upstream information should become more apparent if we evaluate (i.e. "test") our model for winter months (DJF(M)). However, the model is trained with data from all seasons and thus needs to generalise sufficiently well to capture the strong seasonal cycle of ozone concentrations. An overview of our data split is shown in Fig. 2 and Table 1.

We use the following meteorological variables as model input ($\mathbf{X}$): 2 m temperature ($T2$), 2 m water vapour mixing ratio ($Q2$), surface pressure (PSFC), planetary boundary layer height (PBLH), 10 m horizontal wind speed (wspd10ll) and 10 m wind direction (wdir10ll). We convert the horizontal wind components from the Lambert conformal conic projection of WRF-Chem to the geographic coordinate system. We further convert the winds' u and v component to wind speed and wind direction to determine the upstream wind direction. Moreover, we use the following chemical variables from the lowest model level: ozone (O_3), nitrogen oxide (NO), nitrogen dioxide (NO_2) and carbon monoxide (CO). We aggregate the hourly model values to daily means (meteorological variables) and maximum daily 8 h means (dma8eu) (chemical variables[1]; see also Table 2) according to the EU definitions (European Parliament, 2008) using the *toarstats* Python package (Selke et al., 2021). Afterwards, we z-standardise all data from the training set except the wind direction to mean zero with unit variance. The wind direction is min–max-scaled. Subsequently, we apply the scaling parameters obtained from the training set to the validation and testing data sets. This approach allows for the most rigorous evaluation of the model's generalisation capability because no implicit information of the validation and testing sets contaminates the training set. Table 2 summarises all variables, their aggregated statistics and the applied scaling method.

For the baseline experiment (NNb), we take the daily meteorological and chemical variables (see also Sect. 2) at

---

[1] We decided to use dma8eu for all chemical variables to sample all chemical quantities during the same time periods. dma8eu is calculated based on data starting at 17:00 local time (LT) on the previous day, while the mean is calculated based on data starting at 00:00 LT at the current day.

a pseudo-measurement station of the previous $N = 6$ time steps (t_{-6} to t_0) to create the input tensor ($\mathbf{X}$) for our neural network. Accordingly, we use the next $M = 4$ time steps (t_1 to t_4) at the same pseudo-measurement station to create the labels ($\mathbf{y}$). We repeat this procedure for all $K = 332$ pseudo-measurement stations.

For NN1s and NN3s we use additional chemical and meteorological information of the surrounding area as inputs to the neural network. We divide the wind directions into eight sectors corresponding to (i) north, (ii) north-east, (iii) east, (iv) south-east, (v) south, (vi) south-west, (vii) west and (viii) north-west. Consequently, each section covers 45°. For each of our pseudo-observations, we identify the upstream wind sector at time t_0. We then select all WRF grid boxes (centre points) which fall into this sector and have a maximal distance of 200 km to the central point of interest. Next, we compute the spatial mean for all chemical and meteorological variables in this set of grid boxes and use this aggregated information as additional input on top of the local pseudo-observational input that is used in NNb. For NN3s experiments, the same processing is applied to data in the sectors to the left and right of the upstream wind sector. Algorithm 1 depicts the preprocessing strategy.

Figure 3 shows the distribution of upstream sectors, i.e. number of samples, within the test set. The south-western sector dominates ($\sim$ 8200 samples), followed by the western sector ($\sim$ 4800 samples). The northern and north-western sectors contain fewer than 2000 samples each.

Algorithm 1 Data preprocessing: NNb, NN1s, NN3s.

1: Apply toar_stats to all variables
2: Convert horizontal wind components from Lambert conformal to geographic coordinate system
3: Rescale time series with parameters from training set
4: **for all** Pseudo-stations: Create samples **do**
5:     Identify grid box representing pseudo-station
6:     **if** NN1s **or** NN3s **then**
7:         detect upstream wind sector at t_0
8:         average across all grid boxes within the sector
9:         **if** NN3s **then**
10:            detect "left" and "right" upstream wind sector at t_0
11:            average across all grid boxes within each sector
12:        **end if**
13:    **end if**
14:    Create inputs $\mathbf{X}$ with variables of 7 d ($-6\,d$) to ($0\,d$)).
       # Shape of $\mathbf{X}$: $7 \times 1 \times (10 + 10s)$: number of days, 1, number of pseudo-measurement variables + aggregated sector variables
       # s(NNb) = 0, s(NN1s) = 1, s(NN3s) = 3
15:    Create labels $\mathbf{y}$ with ozone (dma8eu) concentrations for the next 4 d (1 to 4 d).
       # Shape of $\mathbf{y}$: 1×4 (1, lead time)
16: **end for**
17: Permute samples in the training set
18: Create batches of size 256

Table 1. Number of pseudo-stations and number of samples for the training (train), validation (val) and testing (test) data sets.

	No. ps. stations	×	No. days	=	No. samples
train	332		278		92 296
val	332		50		16 600
test	332		97		32 204

Table 2. Input and target variables with applied daily aggregation and scaling method.

	Variable	Daily metric	Scaling method
Input	NO	dma8eu	z-standardised
	NO_2	dma8eu	z-standardised
	O_3	dma8eu	z-standardised
	CO	dma8eu	z-standardised
	2 m temperature ($T2$)	mean	z-standardised
	2 m water vapour ($Q2$)	mean	z-standardised
	10 m wind direction (wdir10ll)	mean	min–max-scaled
	10 m wind speed (wspd10ll)	mean	z-standardised
	surface pressure (PSFC)	mean	z-standardised
	planetary boundary layer height (PBLH)	mean	z-standardised
Target	O_3	dma8eu	z-standardised

3.2 MLAir framework

We use the *Machine Learning on Air data* (MLAir, version 1.5) framework (Leufen et al., 2021) as the backbone for our experiments. MLAir provides a workflow framework for machine-learning-based atmospheric forecasts with easily extensible modules for data preprocessing, training, hyperparameter optimisation and evaluation. MLAir uses the TensorFlow (Abadi et al., 2015), dask (Rocklin, 2015) and xarray (Hoyer and Hamman, 2017) libraries. For each of the mentioned preprocessing methods (see Sect. 3.1), we implemented individual *DataHandler*s (see Leufen et al., 2021, Sect. 2.5) that allow us to modify the data preprocessing steps while maintaining the same training and evaluation procedures for all experiments in spite of the different data structures.

3.3 Model architecture and training procedure

In contrast to our previous study (Kleinert et al., 2021), we use a U-shaped convolutional neural network (CNN) with two parallel long short-term memory (LSTM) cells in the lowest level. Ronneberger et al. (2015) first introduced the U-Net architecture for biomedical image segmentation, but many other disciplines picked up this architecture. U-shaped or U-Net architectures have already been used to successfully tackle spatio-temporal problems (He et al., 2022). As outlined in Sect. 3.1, we train three different models according to the three preprocessing variants: (i) baseline (NNb), (ii) one upstream sector (NN1s) and (iii) three upstream sectors (NN3s). The model architecture for the NN3s model is

shown in Fig. 4. The first layer is the input layer. During training the individual samples in **X**, representing short 7 d time series (t_{-6} to t_0) for each variable, are passed to the network. We use the variables as channel dimension. We use 2D convolutional layers with a degenerated width dimension. Thus, the overall dimension of our training set is no. samples × prev. N days × 1 × no. (sector) variables. We use a kernel size of 3×1 for convolutions and a pooling kernel size of 2×1. We apply exponential linear units (ELUs; Clevert et al., 2016) as an activation function for inner layers and a linear activation function for the output layer. Due to the different input data structures between the three experiment sets, the model architectures differ from each other in terms of the number of channels. Other design parameters – like the number of layers and kernel sizes – are identical across all runs. As the time series are relatively short, we use symmetric padding to ensure that the temporal dimension does not shrink. In the U structure's lowest level, we use two LSTM branches after the third convolution block. While we used a max-pooling layer before the first LSTM branch to capture highly dominant features, the second LSTM branch uses the whole temporal dimension of size seven. We use an upsampling layer to expand the temporal dimension of the first LSTM branch and concatenate the resulting tensor with the symmetric padded skip connection of the third convolution block. We apply an additional convolutional layer for information compression before we append (concatenate) the second LSTM branch. Afterwards, we use altering concatenation and convolution layers to reconstruct the U's right slope. Finally, we use a simple dense layer with four nodes as the
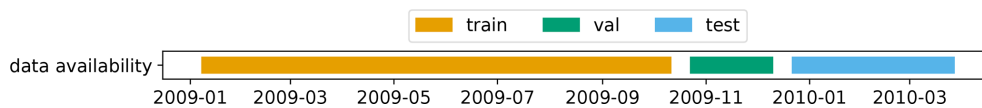
Figure 2. Data availability and split into training (orange), validation (green) and testing (blue) data sets. Each day contains data from all 332 pseudo-stations.
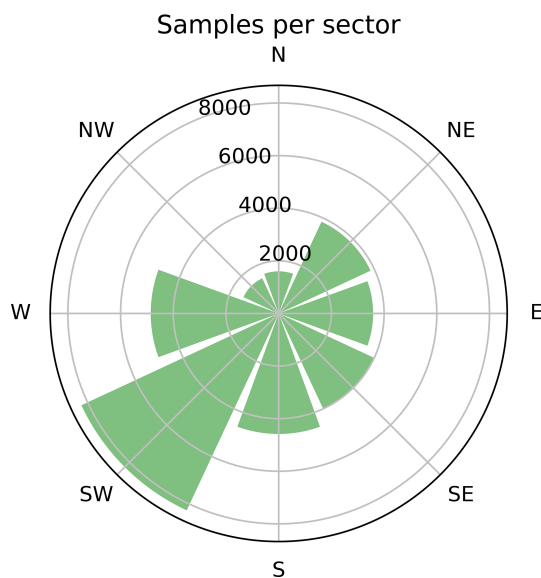
Figure 3. Number of samples per (main) upstream wind sector at time step t_0 within the testing set.

output layer. Here each node corresponds to the prediction for time step t_1 to t_4.

The original U-Net proposed by Ronneberger et al. (2015) does not include padding as their input images are large enough for multiple convolutions. For each convolutional operation with a kernel size of k, the processed input data's shape shrinks by $k-1$ (assuming no padding and strides $s = 1$). To prevent the inputs from becoming too small, we implement the paddings mentioned and use the standard *Up-Conv* layer after the unpadded LSTM branch only. We train all models for 300 epochs using Adam (Kingma and Ba, 2014) as the optimiser and the mean squared error as the loss function. Detailed information on the choice of hyperparameters can be found in Table A1.

3.4 Reference models

As in Kleinert et al. (2021), we use persistence and an ordinary least-squares (OLS) model as references to provide a meaningful baseline evaluation of our machine learning models. The persistence forecast is built simply using the latest observation at t_0 as forecast for all lead times t_1 to t_4. The persistence model serves as a relatively strong competitor on short lead times and should be outperformed by all other models which add any value to the forecasting solution.

To train the OLS model, we use the same data set as for the NN3s network. The OLS model serves as a linear competitor and therefore as an indicator on how well the neural networks differ from a "simple" linear forecast. A detailed description of these two reference forecasts is provided in Leufen et al. (2021) (MLAir) and Kleinert et al. (2021) (IntelliO3-ts). We also reran the experiments from this study with the IntelliO3-ts (Kleinert et al., 2021) model architecture, which is based on inception layers (Szegedy et al., 2015), for comparison. As described in Sect. 3.1, three different variants of the IntelliO3 architecture had to be built because of the different input data dimensions. All models were trained from scratch using the same data as in our main experiments.

Table 3 summarises the experiments described in this paper and introduces the labels that are used to denote these experiments in the Results section.

4 Results

We first present the results of the two sectorial approaches (NN1s and NN3s) in comparison to the baseline method (NNb), which does not use any upstream information. These three models are compared to OLS, persistence and the IntelliO3_[b, 1s, 3s] variants. Secondly, we compare the individual losses of NN3s based on the training, validation and testing loss. Afterwards, we present more detailed results for the multi-sector approach (NN3s), including an exemplary sensitivity analysis for input variables.

Figure 5 shows the overall mean squared error (MSE) for all models and reference models across all lead times. We can clearly distinguish between different experiments ($p < 0.001$; see below). The U-Net architecture networks (NN3s, NN1s, and NNb) show the lowest root mean square error (RMSE). Conversely, the multi-sector, baseline IntelliO3 and persistence experiments (IntelliO3_3s, IntelliO3_b and persi) show the largest RMSE. We estimate the uncertainty by performing a block-wise bootstrapping with a block length of 7 d and 1000 draws with replacement. Additionally, we perform a non-parametric two-sided Mann–Whitney u test (5 % significance level) for NN3s and all competitors.

The NN3s approach shows the lowest RMSE, and the null hypothesis of the performed u test can be rejected ($p < 0.001$) for all competitors (detailed test statistics and corresponding p values are shown in Appendix Table B1). NN3s performs better than the baseline (NNb) and the single sec-
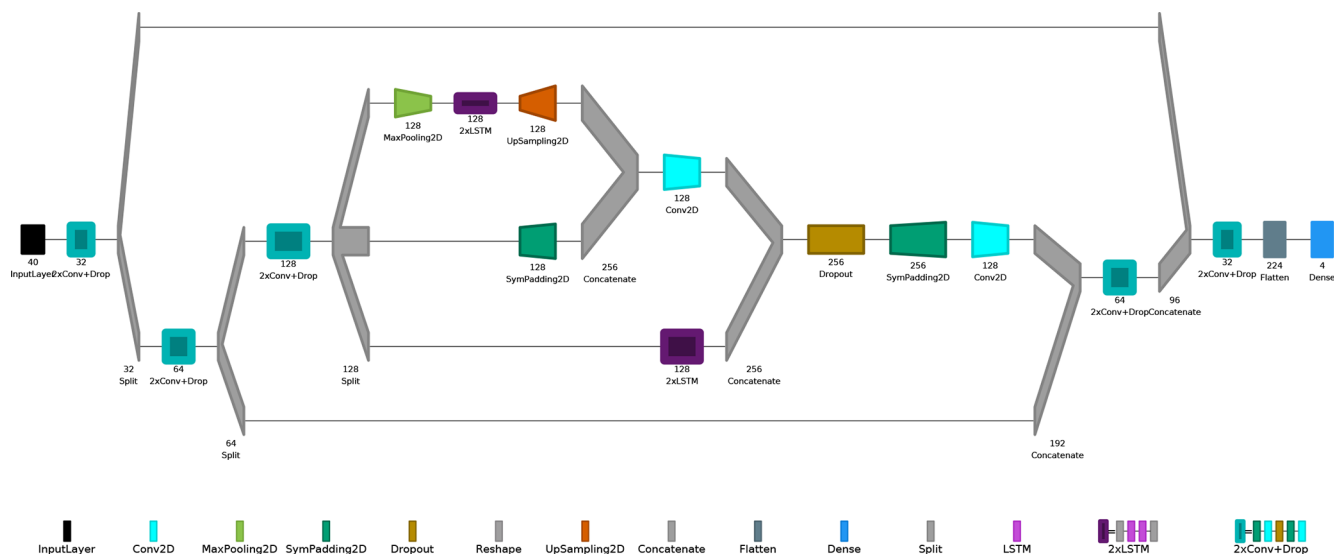
Figure 4. Summary of NN3s model architecture with input (black, left), hidden, and output (blue, right) layer. Several layers are grouped for better visualisations. Colour-coding of individual layers and groups is shown on the right. This figure was created with Net2Vis (Bauerle et al., 2021).

Table 3. Naming of deep learning experiments and reference models described in this paper.

Label	Description
NNb	neural network baseline experiment; U-Net architecture, no upstream information
NN1s	as NNb but with aggregated information from one upstream wind sector
NN3s	as NN1b but with aggregated information from three upstream wind sectors (Fig. 4)
persi	persistence as reference model
ols	ordinary least-squares regression model as reference
IntelliO3-b	as NNb but with IntelliO3 network architecture
IntelliO3-1s	as NN1s but with IntelliO3 network architecture
IntelliO3-3s	as NN3s but with IntelliO3 network architecture

tor (NN1s) approach, which in turn exhibits a lower RMSE compared to OLS, the InitelliO3 variants and persistence.

Figure 6 shows a monthly comparison of the forecast distributions per lead time for the pseudo-observation. The results are shown for the test data set period ranging from December 2009 to March 2010 (see also Sect. 3.1). The forecasts show a narrower distribution when compared to the pseudo-observations' distribution. While we can observe that the NN3s model captures the changing monthly structure in terms of varying mean and median concentrations, the network is not able to adequately reproduce the variability and forecasts converging towards the monthly means. Such a behaviour was already found in Kleinert et al. (2021).

More insights can be gained from evaluating the skill scores of the various experiments for the individual lead times (i.e. 1 d to 2 d). As we base the comparison of our models on the mean squared error (MSE), the skill scores (S) take the form of

$$S(\boldsymbol{m}, \boldsymbol{r}, \boldsymbol{o}) = 1 - \frac{\mathrm{MSE}(\boldsymbol{m}, \boldsymbol{o})}{\mathrm{MSE}(\boldsymbol{r}, \boldsymbol{o})}, \qquad (1)$$

where $\boldsymbol{m}$ is a vector containing the model's forecast, $\boldsymbol{o}$ is a vector containing the corresponding observations and $\boldsymbol{r}$ is a vector containing the reference forecast (Murphy, 1988). Here a positive value of $S > 0$ corresponds to an improvement of the model over the reference. Consequently, a negative value ($S < 0$) corresponds to a deterioration of skill with respect to the reference forecast.

Figure 7 shows the skill scores separated for all lead times. Here the uncertainties (boxes and whiskers) are calculated based on the individual pseudo-stations. As expected, the NN3s approach shows an increasing skill score with increasing lead time when the persistence forecast is used as reference (mean from ~ 0.13 on 1 d to ~ 0.38 on 4 d). When we compare NN3s with respect to the OLS model, the mean skill score is positive throughout (overall mean ~ 0.10) and has its maximum for a lead time of 2 d (mean ~ 0.13), with the smallest interquartile range (IQR). For 3 and 4 d lead times, the skill score decreases. For the comparison of NN3s with respect to the IntelliO3 variants (_b, _1s _3s), we can observe a general pattern with the highest skill score on 1 d, which
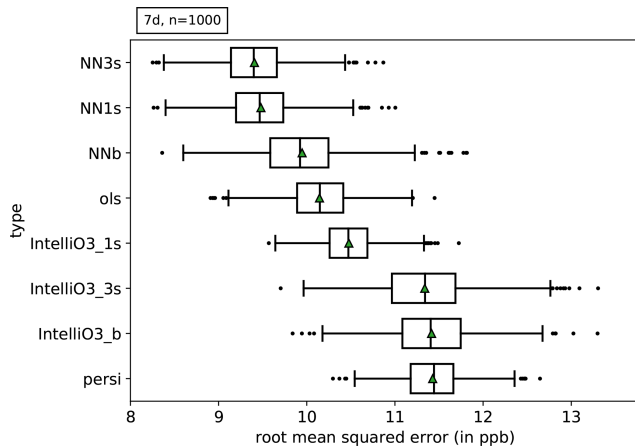
Figure 5. Estimated uncertainty of the mean squared error using block-wise bootstrapping (1000 realisations with replacement with block length of 7 d). IntelliO3_* corresponds to the model architecture as presented in Kleinert et al. (2021) (IntelliO3-ts v1.0).
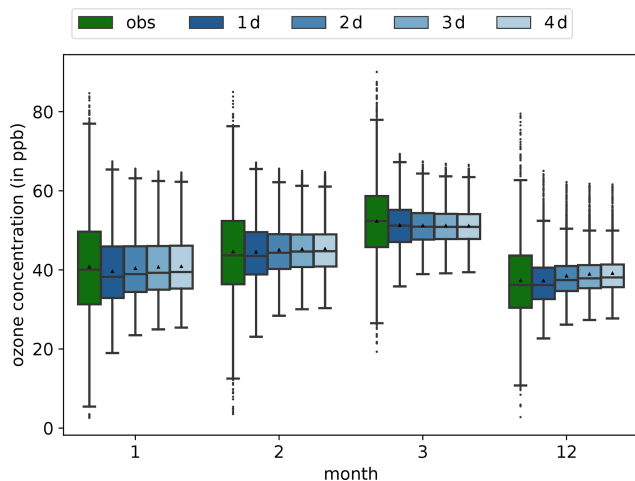
Figure 6. Monthly dma8eu ozone concentration for all test pseudo-stations as box plots. Pseudo-observations are denoted by obs (green), and forecasts are denoted by 1 d (dark blue) to 4 d (light blue). Triangles display the arithmetic means.

decreases with increasing lead time. The mean skill score of NN3s vs. IntelliO3_3s decreases from ~ 0.34 to ~ 0.27, NN3s vs. IntelliO3_1s decreases from ~ 0.24 to ~ 0.12 and NN3s vs. IntelliO3_b from ~ 0.35 to ~ 0.26, respectively. The additional two upstream sectors of IntelliO3_3s do not add any additional information compared to its single sector approach (IntelliO3_1s).

Even though NN3s provides a significant performance improvement over NN1s as shown in Fig. 5, the skill score of the NN3s vs. NN1s comparison is close to zero during all 4 forecast days (mean ~ 0.03 on 1 d, ~ 0.02 on 2 d, ~ 0.01 on 3 d and ~ 0.0 on 4 d). The added value of neighbouring upstream sectors is apparently lost after 1 d.
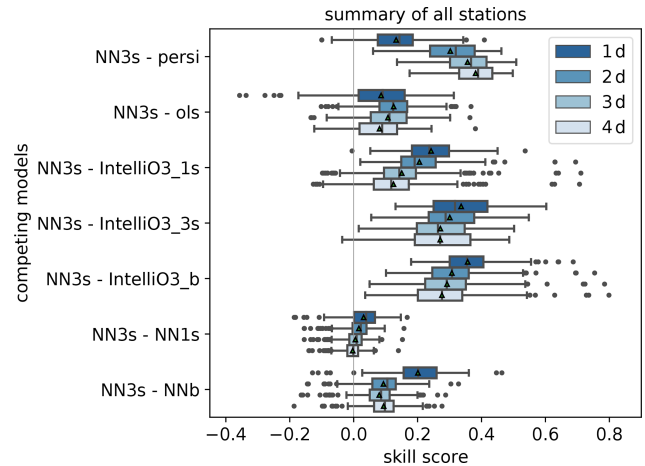
Figure 7. Skill scores of the NN3s model versus the reference models persistence (persi), ordinary least-squares (ols), single upstream sector model (NN1s) and pseudo-station model (NNb) based on the mean squared error, separated for all lead times (1 d (dark blue) to 4 d (light blue)). Positive values denote that NN3s performs better than the given references. Triangles display the arithmetic means.

For the comparison of NN3s with respect to NNb, we see the largest skill score for a lead time of 1 d (mean ~ 0.2) which decreases up to 3 d (~ 0.09, ~ 0.08) and slightly increases again at 4 d (~ 0.09).

As we originally expected a greater benefit of using upstream wind sector information for ozone forecasts, we have further analysed these skill scores and looked at their spatial distribution. From Fig. 7 we can identify some pseudo-stations for which the ols model performs better then the NN3s model. To capture potential differences in spatial properties, Fig. 8 shows the skill NN3s vs. ols skill scores for each pseudo-station. The skill score (NN3s vs. ols) is mostly positive in the northern part of Germany on 1 d and is mostly negative in mountainous regions in the south of Germany. With increasing lead time, the skill score also becomes positive in the mountainous regions but tends to turn negative in the north-east.

5   Discussion

Based on the results presented in Sect. 4, we can observe that NN3s outperforms the (simplistic) persistence and OLS forecasts as well as NNb, which only uses the pseudo-observations of a specific grid box. The increase of the NN3s-persi skill score with increasing lead time (Fig. 7) is in line with expected behaviour as the persistence forecast has its most valuable predictions for short lead times. Consequently, the increased skill score for NN3s-persi is mostly caused by the worsening of the persistence forecast with increasing lead time. However, the positive skill score for a lead time of 1 d shows that the NN3s model has a genuine added value over
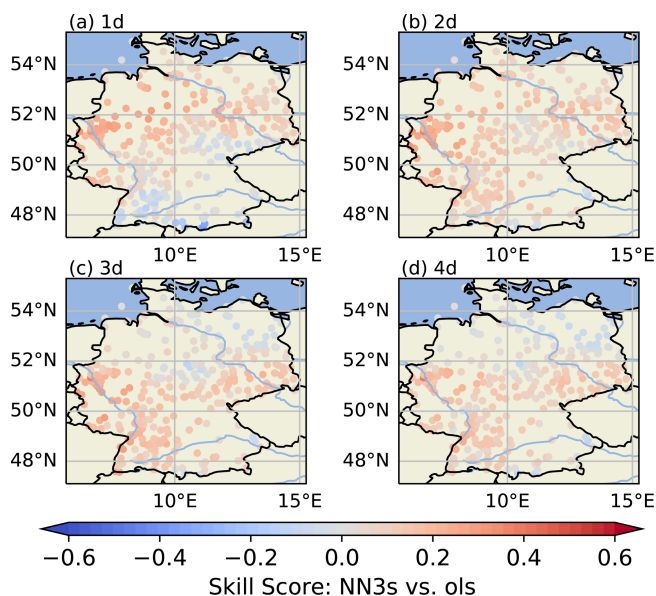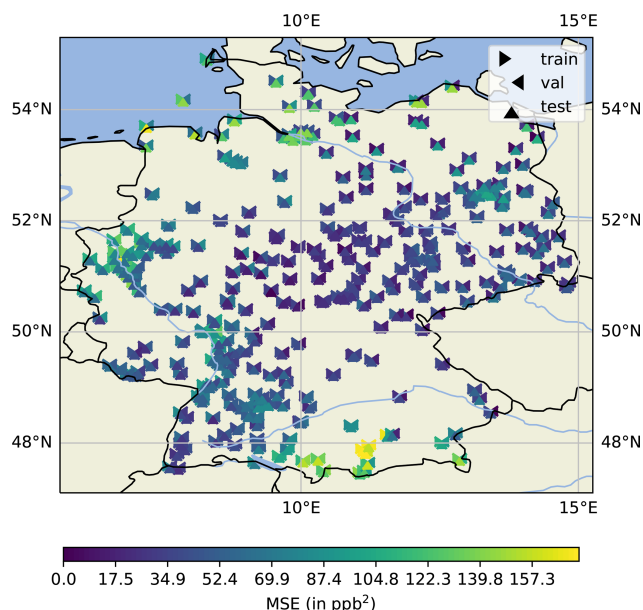
Figure 8. Skill score (NN3s vs. ols) per station for lead times of 1 d **(a)** to 4 d **(d)**. This figure was created with Cartopy (Met Office, 2010–2015). Map data © OpenStreetMap contributors 2022. Distributed under the Open Data Commons Open Database License (ODbL) v1.0.

Figure 9. MSE of the training (left triangle), validation (right triangle) and testing (bottom triangle) data sets for each station. This figure was created with Cartopy (Met Office, 2010–2015). Map data © OpenStreetMap contributors 2022. Distributed under the Open Data Commons Open Database License (ODbL) v1.0.

the persistence forecast on short lead times. When comparing NN3s and NN1s, we see that the skill score's lower bound of the IQR is close to zero for the first 2 d of prediction and that the skill score's mean and median converge towards zero for the remaining lead times, meaning that NN3s behaves like NN1s for 3 and 4 d. Consequently, NN3s can not extract any additional helpful information from the input fields for 3 and 4 d. Most likely, this effect is caused by the definition of neighbouring sectors as described in Sect. 3.1, where we use the upstream wind sector and the two adjacent sectors with a radius of 200 km at time step t_0 to calculate the spatial means for all input time steps t_{-6} to t_0. Thus, depending on the average wind speed, the static upstream wind sectors cannot provide all relevant information. Moreover, the sectorial approach is a cruder approximation of a streamline than a backwards trajectory, and differences between both of them tend to increase with increasing lead time. Based on the NN3s and NN1s sample uncertainty estimate in Fig. 5 and the competitive skill score per lead time in Fig. 7 we can conclude that the statistical significance of the u test (Appendix Table B1) is related to the better performance on the first two lead times (1 and 2 d).

In the following, we discuss differences in the training, validation and testing loss of NN3s, the influence of the upstream wind direction and implications based on our data split in more detail.

5.1 Evaluation of training, validation and testing differences

We use different parts of a year to train, validate and test our models (see Sect. 3.1). Consequently, the network encounters different meteorological and chemical conditions. Therefore, we compare the differences in the MSE (which we used as the loss function during training) for the training, validation and testing sets for each station individually. Averaged over all stations, the mean rescaled losses for NN3s are 64.08 ppb^2 (train, scaled: 1.04), 63.34 ppb^2 (val, scaled: 0.98) and 64.62 ppb^2 (test, scaled: 1.08). Figure 9 shows how the losses are distributed geographically. At first glance, we can identify regions with high MSE ($>$ 120 ppb^2, yellowish colours) in the mountainous south and in the northern coastal area. The triangles denote the losses of each station separated for the training (left part of the symbol), validation (right part of the symbol) and test (lower part of the symbol) set. When focussing on the set's loss differences at individual pseudo-stations, we can identify the following pattern: the test set's loss is mostly lower than the validation and train loss in Germany's western and south-western parts. The high individual validation loss in the northern part is directly related to the large temperature anomaly, as shown in Fig. 1. Thus, the feature combination is not explicitly present in the training set, resulting in larger discrepancies in forecasts and observations.

5.2 Sectorial results – NN3s

In the following we further analyse the influence of the upstream wind direction on the skill score. As mentioned in Sect. 2, the SW direction is the dominant upstream sector in the testing set, while the fewest cases are found in the NW sector.

Figure 10 shows the skill score separated for each wind sector (NN3s vs. NNb). We can observe that the skill score is mostly negative for the northern sector (~ -0.2), having a low number of samples (Fig. 3). On the contrary, the southwestern sector's skill score is – on average – positive but close to zero, indicating that NNb performs equally well in the dominant upstream direction. On average the skill scores for E, SE, S and W are positive, indicating that the NN3s model can extract some useful information in contrast to using the pseudo-observation at the pseudo-location only. However, it also becomes obvious that this is not the case for all stations (negative skill scores indicated by dots). It seems that the additional information provided by the upstream sectors might sometimes confuse the network and therefore does not have the desired effect at all instances.

5.3 Non-linearities

As we saw in Fig. 8, there are pseudo-stations in Germany, where a simpler OLS model produces a better forecast than NN3s. Naively, one might expect a neural network to generate better forecasts because of its ability to capture non-linear dependencies between variables. To better understand these non-linearities, we iteratively modify the input variables of each sample (in all sectors and the pseudo-observation). Afterwards, we feed the modified input samples into the trained NN3s model and detect how the ozone forecasts for all lead times change. Figure 11 shows the sensitivity of ozone inputs for one example station. For the O_3 sensitivity test, the resulting sample distribution is very narrow. We can identify two regions from 0 to ~ 50 ppb and from 80 ppb to the maximum value of 170 ppb, with a transition region in between. The shallower increase in the second region aligns with the findings from Fig. 6, where we see that NN3s cannot adequately reproduce high ozone levels. Figure 12 shows the sensitivity towards 10 m wind speed. The wind speed analysis results in a broader distribution with a flat shape from 0 to $10\,\mathrm{m\,s}^{-1}$ and a linear increase for larger wind speeds.

An increase of NO levels leads to an increase of the resulting ozone level (Fig. 13). The increase is strongest for low NO levels and flattens for larger NO levels. Conversely, an increase on NO_2 leads to a decrease of the resulting ozone levels (Fig. 14) forecasted by NN3s.

Disentangling the influences of NO_x and volatile organic compound (VOC) species on ozone concentrations is a complex endeavour. Investigations into these relations with CTMs show that the results depend on the details of the chemical mechanisms as well as other model parame-

terisations (e.g. for deposition or biogenic VOC emissions). Fast et al. (2014) showed that the reduction of all anthropogenic emissions by 50 % only has a relatively minor effect on the O_3 bias in California, with differences of up to ~ 5 ppb at four selected sites compared to a simulation with default emissions. Abdi-Oskouei et al. (2020) performed several sensitivity studies to assess the impact of meteorological boundary conditions and the land surface model on modelled O_3 concentrations, and they showed that these changes led to a minor sensitivity of average ozone concentration (< 2 ppb). Georgiou et al. (2018) analysed the impact of the different chemical schemes on predicted trace gases and aerosol concentrations, and they found that CBMZ-MOSAIC and MOZART-MOSAIC have similar O_3 biases (10.9 and 11.6 ppb), while the use of the RADM2-MADE/SORGAM mechanism led to a smaller bias of 4.25 ppb. These differences were attributed to the difference in VOC treatment. Gupta and Mohan (2015) have also analysed the sensitivity of ozone concentrations to the choice of the chemical mechanism and noted that the CBMZ performed better than the RACM mechanism due to the revised rates of inorganic reactions in the CBMZ mechanism. Mar et al. (2016) found that the absolute concentration of ozone predicted by the MOZART-4 chemical mechanism is up to $20\,\mu\mathrm{g\,m}^{-3}$ greater than RADM2 in summer. This is explained by the different representations of VOC chemistry and different inorganic rate coefficients. Liu et al. (2022) developed a deep learning model for correcting O_3 biases in a global chemistry–climate model and demonstrated that temperature and geographic variables showed the strongest interaction with O_3 biases in their model.

5.4 Concept drift

We trained all our models and reference models with data ranging from 1 January to 15 October 2009. We selected a short validation period from mid-autumn to early winter and finally based our analysis on data from the winter and early spring seasons. Thus, as depicted in Fig. 1 the testing set contains data from an unusually cold winter, which was not reflected adequately in the training set. Therefore, the network encounters some input patterns and combinations of features during testing, which are not similar to any data used for training, and thus the DL network has to generate predictions outside of its "comfort zone" (see, for example, Leonard et al., 1992, and Pastore and Carnini, 2021, or Ziyin et al., 2020, for periodic data). As presented in Sect. 2, the temperature anomaly during winter 2009/10 was largest in northern Germany and less, but still negative, in the southern parts of Germany. We can observe a similar north–south pattern for the skill scores (NN3s vs ols) per station (Fig. 8) and a skill scores' change of sign with increasing lead time. Nonetheless, our analysis cannot attribute these phenomena and the geographic height to each other, and further analysis using explainable machine learning techniques like, for ex-
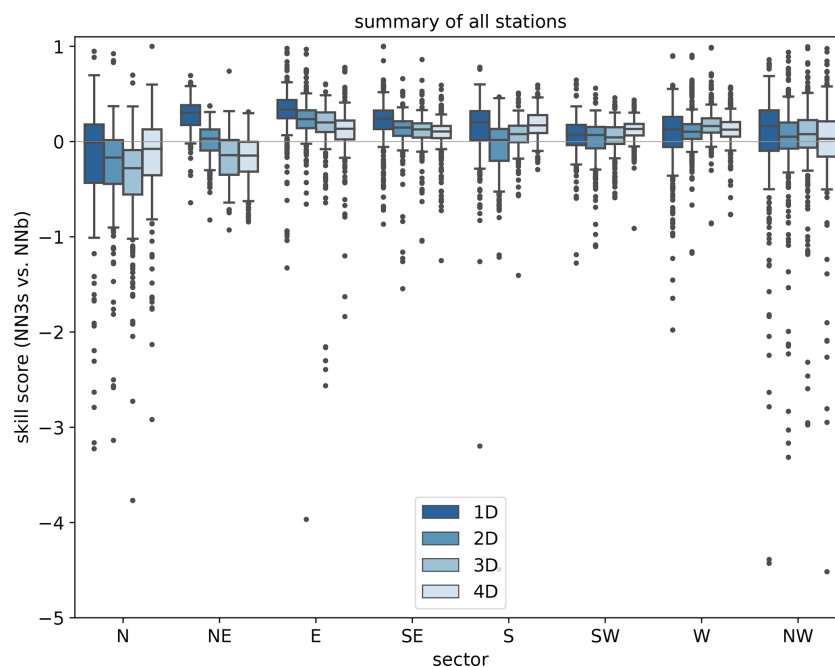
**Figure 10.** Skill scores (NN3s vs. NNb) based on the upstream wind direction at time step t_0.
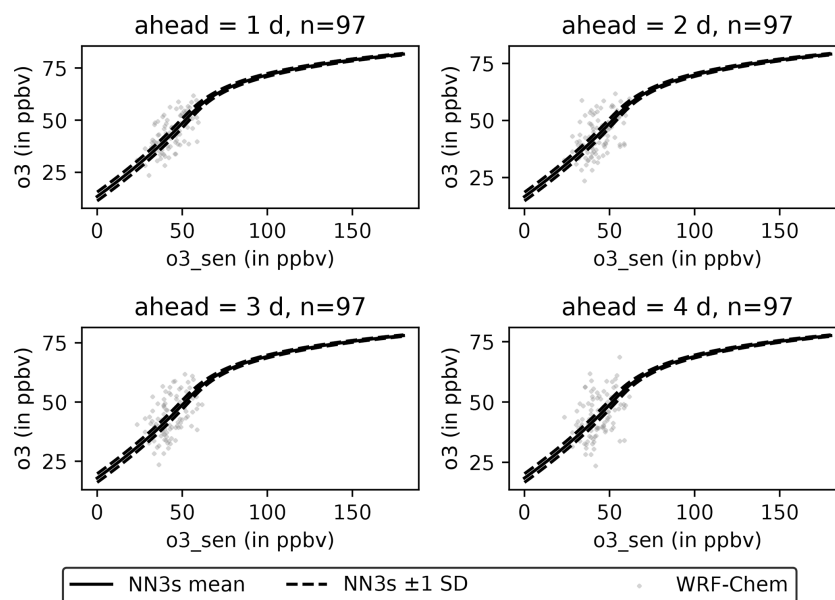
Figure 11. Sensitivity of NN3s forecasts for an exemplary station located at 50.770382° N and 9.459403° E with respect to ozone. Light grey dots (97) represent each sample's daily aggregated WRF-Chem values at time step t_0. The solid black line represents the mean sensitivity across all samples, and the dashed black lines represent ± 1 standard deviation, respectively. Sensitivities are shown individually for all lead times (1 to 4 d).

ample, in Stadler et al. (2022) would be required. In order to reduce the concept drift's effect, extending each data set to multiple years would be beneficial in upcoming studies. This would allow the network to operate on more robust data distributions and thus minimise the risk of out-of-sample predictions.

5.5 Alternative neural network approaches

This study indicates that the forecast quality of air pollutant concentrations can be increased by taking into account the spatial patterns of the meteorological and chemical fields. While our work has focused on time-series predictions to ex-
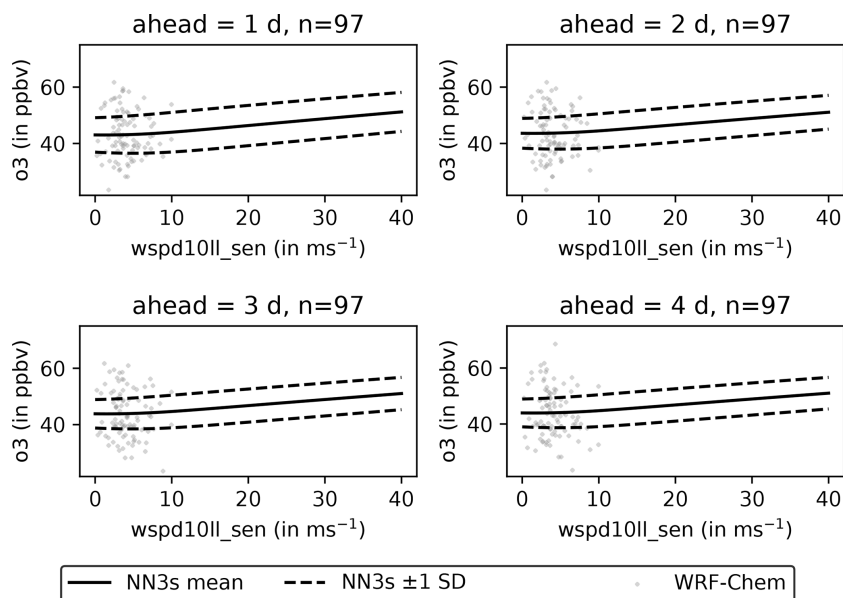
Figure 12. Same as Fig. 11 but for 10 m wind speed instead of ozone input.
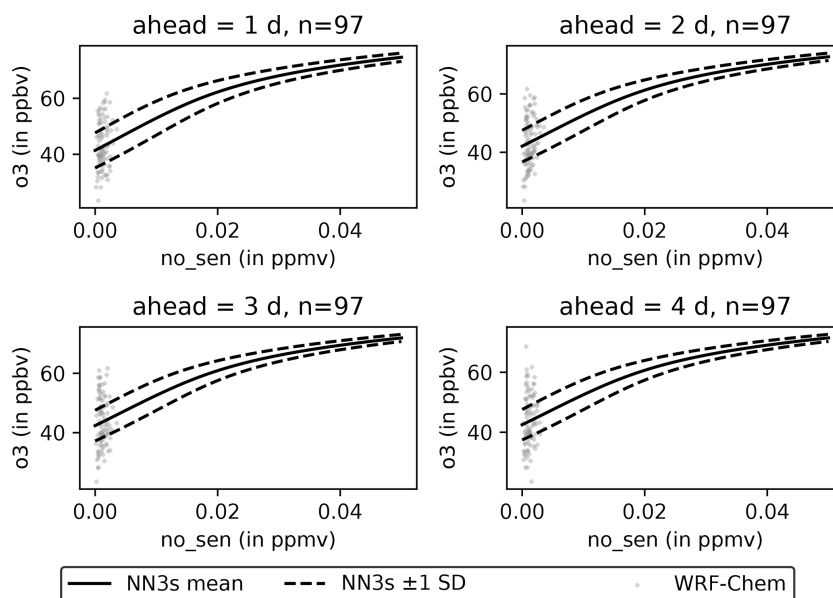
Figure 13. Same as Fig. 11 but for NO instead of ozone input.

plore machine learning methods that can be trained on observational data, there are other studies which employ machine learning to forecast spatio-temporal fields. For example, Steffenel et al. (2021) used the PredRNN++ (Wang et al., 2018) model to forecast the total column ozone for the southern part of South America and parts of Antarctica. Gong et al. (2022) recently used convolutional long short-term memory models (ConvLSTM; Shi et al., 2015) and a generative Stochastic Adversarial Video Prediction (SAVP; (Lee et al., 2018)) model to forecast the 2 m temperature for a lead time of up to 12 h over Europe with some success. As a general out-

come of these studies, it can be said that generative models are better suited to preserve the multi-scale structures of atmospheric fields. While the development of such models for weather applications proceeds rapidly and adopts the latest deep learning methods (see for example Pathak et al., 2022), the machine learning models for atmospheric chemistry have so far been comparatively simpler.

One option to further enhance the approach presented in this study would be the use of Lagrangian particle modelling to derive the area of influence and chemical history for observation sites (see, for example, Yu et al., 2020). Furthermore,
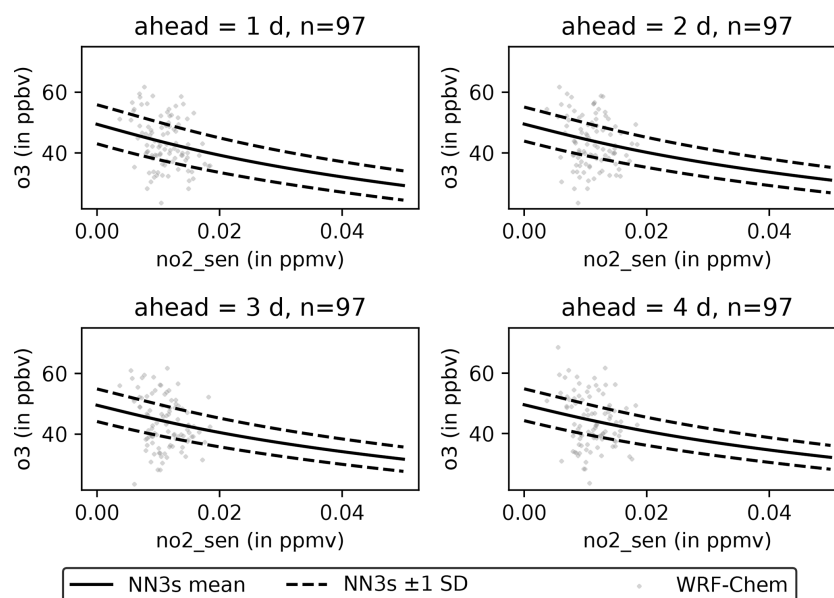
Figure 14. Same as Fig. 11 but for NO_2 instead of ozone input.

the application of Bayesian network architectures can help to characterise data and model uncertainties in future studies (for gridded model examples, see, for example, Sengupta et al., 2020; Sun et al., 2022; Ren et al., 2022).

6 Conclusions

In this study, we explored the potential benefit of using spatially aggregated upstream information to improve pointwise predictions of near-surface ozone concentrations. Even though this analysis was based on pseudo-observations sampled from chemistry transport model, the results should apply to real observation data as well if they are not confounded by the inhomogeneous spatial distribution and missing data occurrences in observational data.

The first result from this study is that a U-Net architecture with a combination of convolutional and LSTM cells is superior to the inception block architecture presented in Kleinert et al. (2021) in this explicit forecasting setting. The second and main result is that the additional information provided by the central upstream wind sector (NN1s) improves the forecast for all lead times (1 to 4 d) with respect to the baseline model (NNb) by 10 %, which has not seen any upstream information during training. Moreover, we show that further information provided by the left and right upstream sectors (NN3s) only improves the forecasts for the first 2 d (∼ 14 %) and that there is no further improvement on the remaining days (3 and 4 d) with respect to the central upstream model (NN1s). NN3s outperforms the ols model at most of the pseudo-stations that we trained on exactly the same data. Thus, we can conclude that the non-linearity provided by

the neural network is essential to extracting meaningful upstream information.

Nonetheless, we showed that the NN3s model does not outperform the ols model at all pseudo-stations and lead times. This can in part be attributed to a sampling bias because the winter in 2010 (test data) was much colder than the winter of 2009 (part of the training data). Besides these limitations, we conservatively evaluated the generalisation capability in the sense that the testing set has the largest possible "distance" to the training data. The improvement in forecast quality arising from the upstream sector information can therefore be seen as a lower limit.

The previous study of Yi et al. (2018) showed that their deep neural network model (*DeepAir*) that makes use of spatially aggregated information outperforms several other network architectures. In complement to their study, we investigated different preprocessing variants. Our variants use various amounts of data to encapsulate the influence on the forecasting performance of one U-shaped network type.

Appendix A: NN3s – hyperparameters

Table A1 shows the hyperparameters used to train the NN3s model. The table is automatically pooled and generated by MLAir (spelling according to TensorFlow).

Appendix B: Additional statistics

Table B1 summarises the results of two-sided Mann–Whitney u tests (5 % significance level) for NN3s and all our competitive models.

Table A1. Hyperparameters used for NN3s.

Parameter	Model setting
_input_shape	(7, 1, 40)
_output_shape	4
activation	elu
amsgrad	False
beta_1	0.9
beta_2	0.999
bias_regularizer	tf.python.keras.regularizers.L1L2
count params	885 252
decay	0.0
epsilon	1e-07
first_filter_size	32
initial_lr	0.001
kernel_initializer	he_normal
kernel_regularizer	tf.python.keras.regularizers.L1L2
kernel_size	(3, 1)
learning_rate	0.001
loss	mean_squared_error
lstm_units	128
metrics	mse, mae
model_name	NN3s
name	Adam
optimizer	tf.python.keras.optimizer_v2.adam.Adam
pool_size	(2, 1)

Table B1. Results of two-sided Mann–Whitney u test (5 % significance level) of NN3s and listed competitor models.

Type	Statistics	p value
NN1s	550 094.0	0.0001
NNb	7.965250×10^5	$\ll 0.0001$
ols	9.024260×10^5	$\ll 0.0001$
IntelliO3_1s	9.790170×10^5	$\ll 0.0001$
IntelliO3_3s	998 110.0	$\ll 0.0001$
IntelliO3_b	998 892.0	$\ll 0.0001$
persi	999 849.0	$\ll 0.0001$

References

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke,

V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X.: TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, https://www.tensorflow.org/ (last access: 1 December 2022), 2015.

Abdi-Oskouei, M., Carmichael, G., Christiansen, M., Ferrada, G., Roozitalab, B., Sobhani, N., Wade, K., Czarnetzki, A., Pierce, R., Wagner, T., and Stanier, C.: Sensitivity of Meteorological Skill to Selection of WRF-Chem Physical Parameterizations and Impact on Ozone Prediction During the Lake Michigan Ozone Study (LMOS), J. Geophys. Res.-Atmos., 125, e2019JD031971, https://doi.org/10.1029/2019JD031971, 2020.

Aliaga, D., Sinclair, V. A., Andrade, M., Artaxo, P., Carbone, S., Kadantsev, E., Laj, P., Wiedensohler, A., Krejci, R., and Bianchi, F.: Identifying source regions of air masses sampled at the tropical high-altitude site of Chacaltaya using WRF-FLEXPART and cluster analysis, Atmos. Chem. Phys., 21, 16453–16477, https://doi.org/10.5194/acp-21-16453-2021, 2021.

Archibald, A. T., Neu, J. L., Elshorbany, Y. F., Cooper, O. R., Young, P. J., Akiyoshi, H., Cox, R. A., Coyle, M., Derwent, R. G., Deushi, M., Finco, A., Frost, G. J., Galbally, I. E., Gerosa, G., Granier, C., Griffiths, P. T., Hossaini, R., Hu, L., Jöckel, P., Josse, B., Lin, M. Y., Mertens, M., Morgenstern, O., Naja, M., Naik, V., Oltmans, S., Plummer, D. A., Revell, L. E., Saiz-Lopez, A., Saxena, P., Shin, Y. M., Shahid, I., Shallcross, D., Tilmes, S., Trickl, T., Wallington, T. J., Wang, T., Worden, H. M., and Zeng, G.: Tropospheric Ozone Assessment Report: A critical review of changes in the tropospheric ozone burden and budget from 1850 to 2100, Elementa, 8, 034, https://doi.org/10.1525/elementa.2020.034, 2020.

Avnery, S., Mauzerall, D. L., Liu, J., and Horowitz, L. W.: Global crop yield reductions due to surface ozone exposure: 1. Year 2000 crop production losses and economic damage, Atmos. Environ., 45, 2284–2296, https://doi.org/10.1016/j.atmosenv.2010.11.045, 2011.

Bauerle, A., van Onzenoodt, C., and Ropinski, T.: Net2Vis – A Visual Grammar for Automatically Generating Publication-Tailored CNN Architecture Visualizations, IEEE T. Vis. Compu. Gr., 27, 2980–2991, https://doi.org/10.1109/TVCG.2021.3057483, 2021.

Betancourt, C., Stomberg, T., Roscher, R., Schultz, M. G., and Stadtler, S.: AQ-Bench: a benchmark dataset for machine learning on global air quality metrics, Earth Syst. Sci. Data, 13, 3013–3033, https://doi.org/10.5194/essd-13-3013-2021, 2021.

CLC: Copernicus Land Monitoring Service: Corine Land Cover, http://land.copernicus.eu/pan-european/corine-land-cover/clc-2012/ (last access: 1 December 2022), 2012.

Clevert, D.-A., Unterthiner, T., and Hochreiter, S.: Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs), arXiv [preprint], https://doi.org/10.48550/arXiv.1511.07289, 2016.

Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P., Kobayashi, S., Andrae, U., Balmaseda, M. A., Balsamo, G., Bauer, P., Bechtold, P., Beljaars, A. C. M., van de Berg, L., Bidlot, J., Bormann, N., Delsol, C., Dragani, R., Fuentes, M., Geer, A. J., Haimberger, L., Healy, S. B., Hersbach, H., Hólm, E. V., Isaksen, L., Kållberg, P., Köhler, M., Matricardi, M., McNally, A. P., Monge-Sanz, B. M., Morcrette, J.-J., Park, B.-K., Peubey, C., de Rosnay, P., Tavolato, C., Thépaut, J.-N., and Vitart, F.: The ERA-Interim reanalysis: configuration and performance of the data assimilation system, Q. J. Roy. Meteor. Soc., 137, 553–597, https://doi.org/10.1002/qj.828, 2011.

DWD: Monthly description, https://www.dwd.de/EN/ourservices/klimakartendeutschland/klimakartendeutschland_monatsbericht.html?nn=495490#buehneTop (last access: 1 December 2022), 2022.

Emmons, L. K., Walters, S., Hess, P. G., Lamarque, J.-F., Pfister, G. G., Fillmore, D., Granier, C., Guenther, A., Kinnison, D., Laepple, T., Orlando, J., Tie, X., Tyndall, G., Wiedinmyer, C., Baughcum, S. L., and Kloster, S.: Description and evaluation of the Model for Ozone and Related chemical Tracers, version 4 (MOZART-4), Geosci. Model Dev., 3, 43–67, https://doi.org/10.5194/gmd-3-43-2010, 2010.

Eslami, E., Choi, Y., Lops, Y., and Sayeed, A.: A real-time hourly ozone prediction system using deep convolutional neural network, Neural Comput. Appl., 32, 8783–8797, https://doi.org/10.1007/s00521-019-04282-x, 2020.

European Parliament, C. o. t. E. U.: Directive 2008/50/EC of the European Parliament and of the Council of 21 May 2008 on ambient air quality and cleaner air for Europe, http://data.europa.eu/eli/dir/2008/50/oj (last access: 1 December 2022), 2008.

Fast, J. D., Allan, J., Bahreini, R., Craven, J., Emmons, L., Ferrare, R., Hayes, P. L., Hodzic, A., Holloway, J., Hostetler, C., Jimenez, J. L., Jonsson, H., Liu, S., Liu, Y., Metcalf, A., Middlebrook, A., Nowak, J., Pekour, M., Perring, A., Russell, L., Sedlacek, A., Seinfeld, J., Setyan, A., Shilling, J., Shrivastava, M., Springston, S., Song, C., Subramanian, R., Taylor, J. W., Vinoj, V., Yang, Q., Zaveri, R. A., and Zhang, Q.: Modeling regional aerosol and aerosol precursor variability over California and its sensitivity to emissions and long-range transport during the 2010 CalNex and CARES campaigns, Atmos. Chem. Phys., 14, 10013–10060, https://doi.org/10.5194/acp-14-10013-2014, 2014.

Fleming, Z. L., Doherty, R. M., von Schneidemesser, E., Malley, C. S., Cooper, O. R., Pinto, J. P., Colette, A., Xu, X., Simpson, D., Schultz, M. G., Lefohn, A. S., Hamad, S., Moolla, R., Solberg, S., and Feng, Z.: Tropospheric Ozone Assessment Report: Present-day ozone distribution and trends relevant to human health, Elementa, 6, 12, https://doi.org/10.1525/elementa.273, 2018.

Galmarini, S., Makar, P., Clifton, O. E., Hogrefe, C., Bash, J. O., Bellasio, R., Bianconi, R., Bieser, J., Butler, T., Ducker, J., Flemming, J., Hodzic, A., Holmes, C. D., Kioutsioukis, I., Kranenburg, R., Lupascu, A., Perez-Camanyo, J. L., Pleim, J., Ryu, Y.-H., San Jose, R., Schwede, D., Silva, S., and Wolke, R.: Technical note: AQMEII4 Activity 1: evaluation of wet and dry deposition schemes as an integral part of regional-scale air quality models, Atmos. Chem. Phys., 21, 15663–15697, https://doi.org/10.5194/acp-21-15663-2021, 2021.

Gaudel, A., Cooper, O. R., Ancellet, G., Barret, B., Boynard, A., Burrows, J. P., Clerbaux, C., Coheur, P.-F., Cuesta, J., Cuevas, E., Doniki, S., Dufour, G., Ebojie, F., Foret, G., Garcia, O., Granados-Muñoz, M. J., Hannigan, J. W., Hase, F., Hassler, B., Huang, G., Hurtmans, D., Jaffe, D., Jones, N., Kalabokas, P., Kerridge, B., Kulawik, S., Latter, B., Leblanc, T., Le Flochmoën, E., Lin, W., Liu, J., Liu, X., Mahieu, E., McClure-Begley, A., Neu, J. L., Osman, M., Palm, M., Petetin, H., Petropavlovskikh, I., Querel, R., Rahpoe, N., Rozanov, A., Schultz, M. G., Schwab, J., Siddans, R., Smale, D., Steinbacher, M., Tanimoto, H., Tarasick, D. W., Thouret, V., Thompson, A. M., Trickl, T., Weath-

erhead, E., Wespes, C., Worden, H. M., Vigouroux, C., Xu, X., Zeng, G., and Ziemke, J.: Tropospheric Ozone Assessment Report: Present-day distribution and trends of tropospheric ozone relevant to climate and global atmospheric chemistry model evaluation, Elementa, 6, 39, https://doi.org/10.1525/elementa.291, 2018.

Georgiou, G. K., Christoudias, T., Proestos, Y., Kushta, J., Hadjinicolaou, P., and Lelieveld, J.: Air quality modelling in the summer over the eastern Mediterranean using WRF-Chem: chemistry and aerosol mechanism intercomparison, Atmos. Chem. Phys., 18, 1555–1571, https://doi.org/10.5194/acp-18-1555-2018, 2018.

Gong, B., Langguth, M., Ji, Y., Mozaffari, A., Stadtler, S., Mache, K., and Schultz, M. G.: Temperature forecasting by deep learning methods, Geosci. Model Dev. Discuss. [preprint], https://doi.org/10.5194/gmd-2021-430, in review, 2022.

Grell, G. A., Peckham, S. E., Schmitz, R., McKeen, S. A., Frost, G., Skamarock, W. C., and Eder, B.: Fully coupled "online" chemistry within the WRF model, Atmos. Environ., 39, 6957–6975, https://doi.org/10.1016/j.atmosenv.2005.04.027, 2005.

Guenther, A., Karl, T., Harley, P., Wiedinmyer, C., Palmer, P. I., and Geron, C.: Estimates of global terrestrial isoprene emissions using MEGAN (Model of Emissions of Gases and Aerosols from Nature), Atmos. Chem. Phys., 6, 3181–3210, https://doi.org/10.5194/acp-6-3181-2006, 2006.

Gupta, M. and Mohan, M.: Validation of WRF/Chem model and sensitivity of chemical mechanisms to ozone simulation over megacity Delhi, Atmos. Environ., 122, 220–229, https://doi.org/10.1016/j.atmosenv.2015.09.039, 2015.

He, T., Jones, D. B. A., Miyazaki, K., Huang, B., Liu, Y., Jiang, Z., White, E. C., Worden, H. M., and Worden, J. R.: Deep Learning to Evaluate US NO_x Emissions Using Surface Ozone Predictions, J. Geophys. Res.-Atmos., 127, e2021JD035597, https://doi.org/10.1029/2021JD035597, 2022.

Hoyer, S. and Hamman, J.: xarray: N-D labeled arrays and datasets in Python, J. Open Res. Softw., 5, p. 10, https://doi.org/10.5334/jors.148, 2017.

Inness, A., Ades, M., Agustí-Panareda, A., Barré, J., Benedictow, A., Blechschmidt, A.-M., Dominguez, J. J., Engelen, R., Eskes, H., Flemming, J., Huijnen, V., Jones, L., Kipling, Z., Massart, S., Parrington, M., Peuch, V.-H., Razinger, M., Remy, S., Schulz, M., and Suttie, M.: The CAMS reanalysis of atmospheric composition, Atmos. Chem. Phys., 19, 3515–3556, https://doi.org/10.5194/acp-19-3515-2019, 2019.

Ismail Fawaz, H., Lucas, B., Forestier, G., Pelletier, C., Schmidt, D. F., Weber, J., Webb, G. I., Idoumghar, L., Muller, P.-A., and Petitjean, F.: InceptionTime: Finding AlexNet for time series classification, Data Min. Knowl. Disc., 34, 1936–1962, https://doi.org/10.1007/s10618-020-00710-y, 2020.

Kingma, D. P. and Ba, J.: Adam: A Method for Stochastic Optimization, arXiv [preprint], https://doi.org/10.48550/arXiv.1412.6980, 2014.

Kleinert, F., Leufen, L. H., and Schultz, M. G.: IntelliO3-ts v1.0: a neural network approach to predict near-surface ozone concentrations in Germany, Geosci. Model Dev., 14, 1–25, https://doi.org/10.5194/gmd-14-1-2021, 2021.

Kleinert, F., Leufen, L. H., Lupaşcu, A., Butler, T., and Schultz, M. G.: Representing chemical history in ozone time-series predictions – a model experiment study building on the MLAir (v1.5) deep learning frame-

work: Experiments and source code, b2share [code], https://doi.org/10.34730/19c94b0b77374395b11cb54991cc497d, 2022a.

Kleinert, F., Leufen, L. H., Lupaşcu, A., Butler, T., and Schultz, M. G.: Representing chemical history in ozone time-series predictions – a model experiment study building on the MLAir (v1.5) deep learning framework: Data 1/4, b2share [data set], https://doi.org/10.34730/c799f04beb644e38a575fa20c2dd8d40, 2022b.

Kleinert, F., Leufen, L. H., Lupaşcu, A., Butler, T., and Schultz, M. G.: Representing chemical history in ozone time-series predictions – a model experiment study building on the MLAir (v1.5) deep learning framework: Data 2/4, b2share [data set], https://doi.org/10.34730/d5f34ae6a8e34d4c8ac33f75b993e8a9, 2022c.

Kleinert, F., Leufen, L. H., Lupaşcu, A., Butler, T., and Schultz, M. G.: Representing chemical history in ozone time-series predictions – a model experiment study building on the MLAir (v1.5) deep learning framework: Data 3/4, b2share [data set], https://doi.org/10.34730/a423ec9003194209989726a95a1a490c, 2022d.

Kleinert, F., Leufen, L. H., Lupaşcu, A., Butler, T., and Schultz, M. G.: Representing chemical history in ozone time-series predictions – a model experiment study building on the MLAir (v1.5) deep learning framework: Data 4/4, b2share [data set], https://doi.org/10.34730/718262bd2c894fd6aadce19a08040f69, 2022e.

Knote, C., Hodzic, A., Jimenez, J. L., Volkamer, R., Orlando, J. J., Baidar, S., Brioude, J., Fast, J., Gentner, D. R., Goldstein, A. H., Hayes, P. L., Knighton, W. B., Oetjen, H., Setyan, A., Stark, H., Thalman, R., Tyndall, G., Washenfelder, R., Waxman, E., and Zhang, Q.: Simulation of semi-explicit mechanisms of SOA formation from glyoxal in aerosol in a 3-D model, Atmos. Chem. Phys., 14, 6213–6239, https://doi.org/10.5194/acp-14-6213-2014, 2014.

Kuenen, J. J. P., Visschedijk, A. J. H., Jozwicka, M., and Denier van der Gon, H. A. C.: TNO-MACC_II emission inventory; a multi-year (2003–2009) consistent high-resolution European emission inventory for air quality modelling, Atmos. Chem. Phys., 14, 10963–10976, https://doi.org/10.5194/acp-14-10963-2014, 2014.

Kuik, F., Lauer, A., Churkina, G., Denier van der Gon, H. A. C., Fenner, D., Mar, K. A., and Butler, T. M.: Air quality modelling in the Berlin–Brandenburg region using WRF-Chem v3.7.1: sensitivity to resolution of model grid and input data, Geosci. Model Dev., 9, 4339–4363, https://doi.org/10.5194/gmd-9-4339-2016, 2016.

Lee, A. X., Zhang, R., Ebert, F., Abbeel, P., Finn, C., and Levine, S.: Stochastic Adversarial Video Prediction, arXiv [preprint], https://doi.org/10.48550/arXiv.1804.01523 2018.

Leonard, J., Kramer, M., and Ungar, L.: A neural network architecture that computes its own reliability, Comput. Chem. Eng., 16, 819–835, https://doi.org/10.1016/0098-1354(92)80035-8, 1992.

Leufen, L. H., Kleinert, F., and Schultz, M. G.: MLAir (v1.0) – a tool to enable fast and flexible machine learning on air data time series, Geosci. Model Dev., 14, 1553–1574, https://doi.org/10.5194/gmd-14-1553-2021, 2021.

Leufen, L. H., Kleinert, F., and Schultz, M. G.: Exploring decomposition of temporal patterns to facilitate learn-

ing of neural networks for ground-level daily maximum 8-hour average ozone prediction, Environ. Data Sci., 1, e10, https://doi.org/10.1017/eds.2022.9, 2022.

Liu, Z., Doherty, R. M., Wild, O., O'Connor, F. M., and Turnock, S. T.: Correcting ozone biases in a global chemistry–climate model: implications for future ozone, Atmos. Chem. Phys., 22, 12543–12557, https://doi.org/10.5194/acp-22-12543-2022, 2022.

Lupaşcu, A. and Butler, T.: Source attribution of European surface O_3 using a tagged O3 mechanism, Atmos. Chem. Phys., 19, 14535–14558, https://doi.org/10.5194/acp-19-14535-2019, 2019.

Mar, K. A., Ojha, N., Pozzer, A., and Butler, T. M.: Ozone air quality simulations with WRF-Chem (v3.5.1) over Europe: model evaluation and chemical mechanism comparison, Geosci. Model Dev., 9, 3699–3728, https://doi.org/10.5194/gmd-9-3699-2016, 2016.

Met Office: Cartopy: a cartographic python library with a Matplotlib interface, Exeter, Devon, https://scitools.org.uk/cartopy (last access: 1 December 2022), 2010–2015.

Mills, G., Pleijel, H., Malley, C. S., Sinha, B., Cooper, O. R., Schultz, M. G., Neufeld, H. S., Simpson, D., Sharps, K., Feng, Z., Gerosa, G., Harmens, H., Kobayashi, K., Saxena, P., Paoletti, E., Sinha, V., and Xu, X.: Tropospheric Ozone Assessment Report: Present-day tropospheric ozone distribution and trends relevant to vegetation, Elementa, 6, 47, https://doi.org/10.1525/elementa.302, 2018.

Murphy, A. H.: Skill Scores Based on the Mean Square Error and Their Relationships to the Correlation Coefficient, Mont. Weather Rev., 116, 2417–2424, https://doi.org/10.1175/1520-0493(1988)116<2417:SSBOTM>2.0.CO;2, 1988.

Pastore, A. and Carnini, M.: Extrapolating from neural network models: a cautionary tale, J. Phys. G, 48, 084001, https://doi.org/10.1088/1361-6471/abf08a, 2021.

Pathak, J., Subramanian, S., Harrington, P., Raja, S., Chattopadhyay, A., Mardani, M., Kurth, T., Hall, D., Li, Z., Azizzadenesheli, K., Hassanzadeh, P., Kashinath, K., and Anandkumar, A.: FourCastNet: A Global Data-driven High-resolution Weather Model using Adaptive Fourier Neural Operators, arXiv [preprint], https://doi.org/10.48550/arXiv.2202.11214, 2022.

Pisso, I., Sollum, E., Grythe, H., Kristiansen, N. I., Cassiani, M., Eckhardt, S., Arnold, D., Morton, D., Thompson, R. L., Groot Zwaaftink, C. D., Evangeliou, N., Sodemann, H., Haimberger, L., Henne, S., Brunner, D., Burkhart, J. F., Fouilloux, A., Brioude, J., Philipp, A., Seibert, P., and Stohl, A.: The Lagrangian particle dispersion model FLEXPART version 10.4, Geosci. Model Dev., 12, 4955–4997, https://doi.org/10.5194/gmd-12-4955-2019, 2019.

Rasp, S., Dueben, P. D., Scher, S., Weyn, J. A., Mouatadid, S., and Thuerey, N.: WeatherBench: A Benchmark Data Set for Data-Driven Weather Forecasting, J. Adv. Model. Earth Sy., 12, e2020MS002203, https://doi.org/10.1029/2020MS002203, 2020.

Ren, X., Mi, Z., Cai, T., Nolte, C. G., and Georgopoulos, P. G.: Flexible Bayesian Ensemble Machine Learning Framework for Predicting Local Ozone Concentrations, Environ. Sci. Technol., 56, 3871–3883, https://doi.org/10.1021/acs.est.1c04076, 2022.

Rocklin, M.: Dask: Parallel computation with blocked algorithms and task scheduling, in: Proceedings of the 14th python in science conference, Citeseer [paper descr. code], 130–136, https://conference.scipy.org/proceedings/scipy2015/pdfs/matthew_rocklin.pdf (last access: 1 December 2022), 2015.

Ronneberger, O., Fischer, P., and Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation, in: Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, edited by: Navab, N., Hornegger, J., Wells, W. M., and Frangi, A. F., Springer International Publishing, Cham, 234–241, https://doi.org/10.48550/ARXIV.1505.04597, 2015.

Sayeed, A., Choi, Y., Eslami, E., Lops, Y., Roy, A., and Jung, J.: Using a deep convolutional neural network to predict 2017 ozone concentrations, 24 hours in advance, Neural Networks, 121, 396–408, https://doi.org/10.1016/j.neunet.2019.09.033, 2020.

Sayeed, A., Choi, Y., Eslami, E., Jung, J., Lops, Y., Salman, A. K., Lee, J.-B., Park, H.-J., and Choi, M.-H.: A novel CMAQ-CNN hybrid model to forecast hourly surface-ozone concentrations 14 days in advance, Sci. Rep., 11, 10891, https://doi.org/10.1038/s41598-021-90446-6, 2021.

Sayeed, A., Eslami, E., Lops, Y., and Choi, Y.: CMAQ-CNN: A new-generation of post-processing techniques for chemical transport models using deep neural networks, Atmos. Environ., 273, 118961, https://doi.org/10.1016/j.atmosenv.2022.118961, 2022.

Schultz, M. G., Schröder, S., Lyapina, O., Cooper, O., Galbally, I., Petropavlovskikh, I., Von Schneidemesser, E., Tanimoto, H., Elshorbany, Y., Naja, M., Seguel, R., Dauert, U., Eckhardt, P., Feigenspahn, S., Fiebig, M., Hjellbrekke, A.-G., Hong, Y.-D., Christian Kjeld, P., Koide, H., Lear, G., Tarasick, D., Ueno, M., Wallasch, M., Baumgardner, D., Chuang, M.-T., Gillett, R., Lee, M., Molloy, S., Moolla, R., Wang, T., Sharps, K., Adame, J. A., Ancellet, G., Apadula, F., Artaxo, P., Barlasina, M., Bogucka, M., Bonasoni, P., Chang, L., Colomb, A., Cuevas, E., Cupeiro, M., Degorska, A., Ding, A., Fröhlich, M., Frolova, M., Gadhavi, H., Gheusi, F., Gilge, S., Gonzalez, M. Y., Gros, V., Hamad, S. H., Helmig, D., Henriques, D., Hermansen, O., Holla, R., Huber, J., Im, U., Jaffe, D. A., Komala, N., Kubistin, D., Lam, K.-S., Laurila, T., Lee, H., Levy, I., Mazzoleni, C., Mazzoleni, L., McClure-Begley, A., Mohamad, M., Murovic, M., Navarro-Comas, M., Nicodim, F., Parrish, D., Read, K. A., Reid, N., Ries, L., Saxena, P., Schwab, J. J., Scorgie, Y., Senik, I., Simmonds, P., Sinha, V., Skorokhod, A., Spain, G., Spangl, W., Spoor, R., Springston, S. R., Steer, K., Steinbacher, M., Suharguniyawan, E., Torre, P., Trickl, T., Weili, L., Weller, R., Xu, X., Xue, L., and Zhiqiang, M.: Tropospheric Ozone Assessment Report: Database and Metrics Data of Global Surface Ozone Observations, Elementa, 5, 58, https://doi.org/10.1525/elementa.244, 2017.

Schultz, M. G., Betancourt, C., Gong, B., Kleinert, F., Langguth, M., Leufen, L. H., Mozaffari, A., and Stadtler, S.: Can deep learning beat numerical weather prediction?, Philos. T. Roy. Soc. A, 379, 20200097, https://doi.org/10.1098/rsta.2020.0097, 2021.

Selke, N., Schröder, S., and Schultz, M. G.: toarstats, https://gitlab.jsc.fz-juelich.de/esde/toar-public/toarstats (last access: 1 December 2022), 2021.

Sengupta, U., Amos, M., Hosking, J. S., Rasmussen, C. E., Juniper, M., and Young, P. J.: Ensembling geophysical models with Bayesian Neural Networks, arXiv [preprint], https://doi.org/10.48550/arXiv.2010.03561,, 2020.

Shi, X., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-k., and Woo, W.-c.: Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting, in: Proceedings of the 28th

International Conference on Neural Information Processing Systems – Volume 1, NIPS'15, MIT Press, Cambridge, MA, USA, event-place: Montreal, Canada, 802–810, 2015.

Stadtler, S., Betancourt, C., and Roscher, R.: Explainable Machine Learning Reveals Capabilities, Redundancy, and Limitations of a Geospatial Air Quality Benchmark Dataset, Machine Learning and Knowledge Extraction, 4, 150–171, https://doi.org/10.3390/make4010008, 2022.

Steffenel, L. A., Anabor, V., Kirsch Pinheiro, D., Guzman, L., Dornelles Bittencourt, G., and Bencherif, H.: Forecasting upper atmospheric scalars advection using deep learning: an O_3 experiment, Special Issue on Machine Learning for Earth Observation Data, Mach. Learn., 1–24, https://doi.org/10.1007/s10994-020-05944-x, 2021.

Stohl, A., Hittenberger, M., and Wotawa, G.: Validation of the lagrangian particle dispersion model FLEXPART against large-scale tracer experiment data, Atmos. Environ., 32, 4245–4264, https://doi.org/10.1016/S1352-2310(98)00184-8, 1998.

Stohl, A., Klimont, Z., Eckhardt, S., Kupiainen, K., Shevchenko, V. P., Kopeikin, V. M., and Novigatsky, A. N.: Black carbon in the Arctic: the underestimated role of gas flaring and residential combustion emissions, Atmos. Chem. Phys., 13, 8833–8855, https://doi.org/10.5194/acp-13-8833-2013, 2013.

Sun, H., Shin, Y. M., Xia, M., Ke, S., Wan, M., Yuan, L., Guo, Y., and Archibald, A. T.: Spatial Resolved Surface Ozone with Urban and Rural Differentiation during 1990–2019: A Space–Time Bayesian Neural Network Downscaler, Environ. Sci. Technol., 56, 7337–7349, https://doi.org/10.1021/acs.est.1c04797, 2022.

Szegedy, C., Wei Liu, Yangqing Jia, Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A.: Going deeper with convolutions, in: 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, Boston, MA, USA, 1–9, https://doi.org/10.1109/CVPR.2015.7298594, 2015.

Tarasick, D., Galbally, I. E., Cooper, O. R., Schultz, M. G., Ancellet, G., Leblanc, T., Wallington, T. J., Ziemke, J., Liu, X., Steinbacher, M., Staehelin, J., Vigouroux, C., Hannigan, J. W., García, O., Foret, G., Zanis, P., Weatherhead, E., Petropavlovskikh, I., Worden, H., Osman, M., Liu, J., Chang, K.-L., Gaudel, A., Lin, M., Granados-Muñoz, M., Thompson, A. M., Oltmans, S. J., Cuesta, J., Dufour, G., Thouret, V., Hassler, B., Trickl, T., and Neu, J. L.: Tropospheric Ozone Assessment Report: Tropospheric ozone from 1877 to 2016, observed levels, trends and uncertainties, Elementa, 7, 39, https://doi.org/10.1525/elementa.376, 2019.

Wang, Y., Gao, Z., Long, M., Wang, J., and Yu, P. S.: PredRNN++: Towards A Resolution of the Deep-in-Time Dilemma in Spatiotemporal Predictive Learning, in: Proceedings of the 35th International Conference on Machine Learning, edited by: Dy, J. and Krause, A., vol. 80 of Proceedings of Machine Learning Research, 5123–5132, https://proceedings.mlr.press/v80/wang18b.html (last access: 1 December 2022), 2018.

Wenig, M., Spichtinger, N., Stohl, A., Held, G., Beirle, S., Wagner, T., Jähne, B., and Platt, U.: Intercontinental transport of nitrogen oxide pollution plumes, Atmos. Chem. Phys., 3, 387–393, https://doi.org/10.5194/acp-3-387-2003, 2003.

WHO: Health risks of air pollution in Europe – HRAPIE project. Recommendations for concentration–response functions for cost–benefit analysis of particulate matter, ozone and nitrogen dioxide, Tech. rep., WHO Regional Office for Europe, UN City, Marmorvej 51 DK-2100 Copenhagen Ø, Denmark, https://www.euro.who.int/__data/assets/pdf_file/0006/238956/Health_risks_air_pollution_HRAPIE_project.pdf (last access: 1 December 2022) 2013.

Wiedinmyer, C., Akagi, S. K., Yokelson, R. J., Emmons, L. K., Al-Saadi, J. A., Orlando, J. J., and Soja, A. J.: The Fire INventory from NCAR (FINN): a high resolution global model to estimate the emissions from open burning, Geosci. Model Dev., 4, 625–641, https://doi.org/10.5194/gmd-4-625-2011, 2011.

Yi, X., Zhang, J., Wang, Z., Li, T., and Zheng, Y.: Deep Distributed Fusion Network for Air Quality Prediction, in: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, ACM, 965–973, https://doi.org/10.1145/3219819.3219822, 2018.

Young, P. J., Naik, V., Fiore, A. M., Gaudel, A., Guo, J., Lin, M. Y., Neu, J. L., Parrish, D. D., Rieder, H. E., Schnell, J. L., Tilmes, S., Wild, O., Zhang, L., Ziemke, J., Brandt, J., Delcloo, A., Doherty, R. M., Geels, C., Hegglin, M. I., Hu, L., Im, U., Kumar, R., Luhar, A., Murray, L., Plummer, D., Rodriguez, J., Saiz-Lopez, A., Schultz, M. G., Woodhouse, M. T., and Zeng, G.: Tropospheric Ozone Assessment Report: Assessment of global-scale model performance for global and regional ozone distributions, variability, and trends, Elementa, 6, 10, https://doi.org/10.1525/elementa.265, 2018.

Yu, C., Zhao, T., Bai, Y., Zhang, L., Kong, S., Yu, X., He, J., Cui, C., Yang, J., You, Y., Ma, G., Wu, M., and Chang, J.: Heavy air pollution with a unique "non-stagnant" atmospheric boundary layer in the Yangtze River middle basin aggravated by regional transport of $PM_{2.5}$ over China, Atmos. Chem. Phys., 20, 7217–7230, https://doi.org/10.5194/acp-20-7217-2020, 2020.

Zaveri, R. A., Easter, R. C., Fast, J. D., and Peters, L. K.: Model for Simulating Aerosol Interactions and Chemistry (MOSAIC), J. Geophys. Res., 113, D13204, https://doi.org/10.1029/2007JD008782, 2008.

Ziyin, L., Hartwig, T., and Ueda, M.: Neural Networks Fail to Learn Periodic Functions and How to Fix It, in: Advances in Neural Information Processing Systems, edited by: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H., 33, 1583–1594, https://proceedings.neurips.cc/paper/2020/file/1160453108d3e537255e9f7b931f4e90-Paper.pdf (last access: 1 December 2022), 2020.