A Calibrated and Consistent Combination of Probabilistic Forecasts for the Exceedance of Several Precipitation Thresholds Using Neural Networks

P. SCHAUMANN,^a R. HESS,^b M. REMPEL,^b U. BLAHAK,^b AND V. SCHMIDT^a

^a *Institute of Stochastics, Ulm University, Ulm, Germany*

^b *Deutscher Wetterdienst, Offenbach, Germany*

(Manuscript received 15 October 2020, in final form 31 March 2021)

ABSTRACT: The seamless combination of nowcasting and numerical weather prediction (NWP) aims to provide a functional basis for very-short-term forecasts, which are essential (e.g., for weather warnings). In this paper we propose a statistical method for precipitation using neural networks (NN) that combines nowcasting data from DWD's radar-based RadVOR system with postprocessed forecasts of the high resolving NWP ensemble COSMO-DE-EPS. The postprocessing is performed by Ensemble-MOS of DWD. Whereas the quality of the nowcasting projections of RadVOR is excellent at the beginning, it declines rapidly after about 2 h. The postprocessed forecasts of COSMO-DE-EPS in contrast start with lower accuracy but provide meaningful information on longer forecast ranges. The combination of the two systems is performed for probabilities that the expected precipitation amounts exceed a series of predefined thresholds. The resulting probabilistic forecasts are calibrated and outperform both input systems in terms of accuracy for forecast ranges from 1 to 6 h as shown by verification. The proposed NN-model generalizes a previous statistical model based on extended logistic regression, which was restricted to only one threshold of 0.1 mm. The various layers of the NN-model are related to the conventional design elements (e.g., triangular functions and interaction terms) of the previous model for easier insight.

KEYWORDS: Europe; Precipitation; Neural networks; Forecast verification/skill; Probability forecasts/models/distribution; Short-range prediction

1. Introduction

Accurate and reliable forecasts of precipitation in the very-short-term range up to 6 h in terms of location and time are required in order to issue targeted warnings (Wang et al. 2017). Early warnings help to increase the lead time for decision makers in hydrological and emergency services and can help to diminish possible damages caused by flooding or debris flows. Commonly used as the basis of these warnings in current operational weather forecasting are nowcasting systems and numerical weather prediction (NWP). Both approaches can provide valuable warning guidance, however, for different forecast lead times (Heizenreder et al. 2015; Hess 2020).

Methods for nowcasting of precipitation usually rely on remote sensing observations from a radar network. In a processing step, the obtained radar reflectivities are transformed to estimates of the current rainfall rate. Based on the Lagrangian persistence approach the latest rainfall rates are extrapolated in space and time using a previously determined motion vector field (Germann and Zawadzki 2002). The dynamical evolutions of the convective cells are not considered. Hence, due to the spatial dependency of the lifetime of precipitation fields (Venugopal et al. 1999), such forecasts are skillful as long as the

Lagrangian persistence assumption is valid (Zawadzki et al. 1994). How far ahead of time weather events can be predicted depends on their size and reaches from hours at scales of several hundred kilometers down to minutes when considering developing thunderstorms (Foresti and Seed 2014).

In contrast, NWP models explicitly simulate the physical evolution of precipitation fields. Forecast errors result from initial and boundary conditions and from approximating physical equations and their inexact solutions due to finite resolutions in time and space. The simulation of cloud microphysics by subgrid parameterizations is especially important for precipitation forecasting (Nicolis et al. 2009). In Stephan et al. (2008), it is shown that deficiencies in the simulated rainfall intensities are attributed to shortcomings in the microphysics parameterization.

To estimate the intrinsic uncertainties accompanying NWP forecasts, ensembles were introduced. These ensembles consist of multiple realizations of a model run; diversity among members of the ensemble may be achieved by varying factors such as initial conditions, boundary conditions, model physics, and parameterizations (Palmer 2002; Gebhardt et al. 2011). Ensemble forecasting provides users with information on the possible range of weather scenarios to be expected.

Nevertheless, multiple runnings of an NWP model can only reduce the random errors and not the aforementioned structural model deficiencies. Therefore, ensemble members are not able to represent the whole spectrum of uncertainties (Scheuerer 2014). Hence, a statistical postprocessing is necessary to reduce systematic biases and an often experienced underdispersive behavior of ensemble forecasts (Gneiting et al. 2005).

Denotes content that is immediately available upon publication as open access.

Corresponding author: Peter Schaumann, peter.schaumann@uni-ulm.de

DOI: 10.1175/WAF-D-20-0188.1

© 2021 American Meteorological Society. For information regarding reuse of this content and general copyright information, consult the AMS Copyright Policy (www.ametsoc.org/PUBSReuseLicenses).

Even with a bias corrected and distributional improved NWP prediction, the NWP forecasts still exhibit various errors in shorter time and small spatial scales. Therefore, a statistical combination with extrapolated nowcasting can improve the skill. The so-called seamless combination aims to create a unique and consistent forecast regardless of location and lead time (Brunet et al. 2015).

Vannitsem et al. (2020) point out that the combination of nowcasting and NWP forecasts may take place in physical or probability space. NIMROD (Golding 1998) as one of the first combination schemes is based on a simple lead-time-dependent weighting function, where the weighting is based on a long-term comparative verification of both initial forecast systems. In INCA (Haiden et al. 2011), the weight for NWP forecasts linearly increases from the beginning until it reaches 1 at a lead time of +4 h.

The Short-Term Ensemble Prediction System (STEPS), see Bowler et al. (2006) and Seed et al. (2013), represents a more advanced combination method. Herein, tendencies in the latest observations and the NWP skill are quantified in real time and used to adjust weights combining the nowcast extrapolation and the NWP forecast, in dependence on lead time and spatial scale. Additionally, a forecast ensemble is generated due to the replacement of nonpredictable spatial scales with correlated stochastic noise. Due to the emergence of nowcasting ensembles, efforts were made to not only use the forecast skill as objective combination metric but also the ensemble spread. Recently, Nerini et al. (2019) utilize an ensemble Kalman filter to iteratively combine NWP forecasts with radar-based nowcasting extrapolations.

In reference to combination methods in probability space, the blending scheme of Kober et al. (2012) weights exceedance probabilities derived from output of the NWP model COSMO-DE-EPS with smoothed neighborhood probabilities computed from the deterministic nowcasting algorithm Rad-TRAM. The weights are based on a long-term verification. Combinations of multimodel ensembles are carried out in Johnson and Wang (2012) and Bouttier and Marchal (2020).

In a previous study, Schaumann et al. (2020) propose the LTI-model as a modified logistic regression model for precipitation rates higher than 0.1 mm h^{-1} . In addition to the logistic regression, triangular functions and interaction terms are introduced to take the possible differences in the individual initial probabilistic forecasts into account and, furthermore, to increase the flexibility to compensate a possible underestimation and overestimation. The logistic regression model is a common tool in the area of probabilistic weather forecasting and has been used for the calibration of forecasts in, e.g., Hamill et al. (2008).

The present study aims to generalize the LTI-model with the help of a neural network (NN). The network to be developed should not only satisfy the demands set for the LTI-model but should also provide consistent threshold exceedance probabilities, where consistency is understood such that the probabilities of exceeding a threshold monotonously decrease with increasing thresholds. This is not ensured when models are trained for each threshold independently. Further, the network should be able to represent forecast uncertainty with

increasing lead time. For other extensions to the logistic regression model in order to ensure consistency, see Wilks (2009); Ben Bouallègue (2013).

As in Schaumann et al. (2020), the training dataset is based on forecasts of RadVOR (Winterrath et al. 2012) and Ensemble-MOS (Hess 2020). RadVOR is a nowcasting system that provides deterministic extrapolation forecasts of radar-based rainfall estimates. Radar observations are obtained by the operational German radar network operated by Deutscher Wetterdienst (DWD). Exceedance probabilities from the deterministic extrapolation forecasts are derived by using the neighborhood approach described by Theis et al. (2005). Ensemble-MOS statistically postprocesses output from the ensemble of the convection-permitting COSMO-DE model, which was upgraded to COSMO-D2 on 15 May 2018. It provides probabilistic precipitation forecasts for various thresholds using logistic regression.

The present study is organized as follows. In section 2, a brief overview of the utilized training dataset is given. Section 3 provides a brief summary of the LTI-model. Afterward, the development of an NN architecture for the model generalization regarding the simultaneous consideration of multiple thresholds is described in detail. Since NNs react sensitive on the chosen hyper-parameters, a hyper-parameter optimization approach is provided in the appendix. Results are stated in section 4. Herein, sensitivities in the choice of hyper-parameters are investigated and a combination example is given. Finally, in section 5 conclusions are drawn and an outlook is given for possible future developments.

2. Data

As in Schaumann et al. (2020), we use forecasts of the DWD systems Ensemble-MOS and RadVOR as data sources. However, we now extend the considered training and forecast period by three months and have thus, altogether, 6 months of precipitation data from April to September 2016. Furthermore, in the previous study, we considered data for only one threshold (0.1 mm h^{-1}) whereas now we consider 9 precipitation thresholds $t_1 = 0.1, t_2 = 0.2, t_3 = 0.3, t_4 = 0.5, t_5 = 0.7, t_6 = 1, t_7 = 2, t_8 = 3$ and $t_9 = 5$, with mm h^{-1} as unit of measurement. It should be noted that the considered period was chosen because it contains severe weather events, see Piper et al. (2016). This makes the dataset especially well-suited for the consideration of higher precipitation thresholds.

The data used from both sources for the results derived in the present paper span across Germany and parts of neighboring countries. In this section, we briefly recall the main properties of the forecast systems Ensemble-MOS and RadVOR as well as the calibrated rainfall estimates used as ground truth. For further details regarding this dataset, we refer to our previous study.

a. Ensemble-MOS

The postprocessing system Ensemble-MOS of DWD is a model output statistics (MOS) system with the capability to statistically process ensemble forecasts resulting in hourly outputs of probabilistic forecasts. These forecasts are available

on a regular grid with a grid size of $20 \times 20 \text{ km}^2$, for lead times up to +21 h, and for various weather elements to support weather warnings (Hess 2020). From the latter output variables, we consider the exceedance of nine precipitation thresholds. The Ensemble-MOS forecasts used in the present paper are based on ensemble forecasts of DWD's 2016 operational convection-permitting NWP model COSMO-DE-EPS (Theis et al. 2012). In particular, the training of the Ensemble-MOS relies on COSMO-DE-EPS forecasts from 2011 to 2015.

It should be noted that while the considered grid size in this paper is $20 \times 20 \text{ km}^2$, the probabilities refer to the exceedance of a given threshold at these grid points, which is a good estimate for the probability of exceeding the threshold within an area of size $1 \times 1 \text{ km}^2$.

b. RadVOR

Additionally to the Ensemble-MOS forecasts, we use data from the nowcasting method RadVOR (Weigl and Winterrath 2010; Winterrath et al. 2012). RadVOR provides every 5 min deterministic precipitation forecasts for lead times up to +2 h on a regular grid of size $1 \times 1 \text{ km}^2$. These very-short-term forecasts consist of two components. In a first step, estimates of the current rainfall rate are derived from radar reflectivities. Thereafter, these rainfall rates are then advected in 5-min increments according to a previously estimated motion vector field. Note that the edges of the respective forecast domain are shifted accordingly.

For the combination with Ensemble-MOS using the model proposed in the present paper, the RadVOR forecasts are interpolated to the same grid and matching time intervals. For this, the 5-min precipitation amounts are added up for a 1-h sum. Next, we consider the grid points with precipitation rates larger than the given threshold on the grid with a grid size of $1 \times 1 \text{ km}^2$. Finally, to estimate the probability of threshold exceedance, we compute a local weighted average for the exceedance on the $1 \times 1 \text{ km}^2$ grid for each grid point on the $20 \times 20 \text{ km}^2$ grid.

c. Calibrated hourly rainfall estimates

As a ground truth for the training and validation of our models, we use calibrated rainfall estimates on a $1 \times 1 \text{ km}^2$ grid based on reflectivity measurements of DWD's operational radar network. These rainfall estimates are adjusted by about 1300 rain gauge measurements to reduce the error induced by the uncertainty of the relationship between radar reflectivities and precipitation amounts (Winterrath et al. 2012). For each grid point on the previously considered $20 \times 20 \text{ km}^2$ grid, we select the nearest neighbor on the $1 \times 1 \text{ km}^2$ grid as the corresponding ground truth. Note that, in comparison to the previous study, the filter algorithm proposed by Winterrath and Rosenow (2007) for removing pixel artifacts has not been applied.

3. Neural network architectures

In Schaumann et al. (2020) we used a modified logistic regression model for the combination of two different probabilistic forecasts. This model is referred to as LTI-model in the

following. Here, L stands for “logistic” while T and I refer to “triangular functions” and “interaction terms,” respectively, which are the modifications of the standard logistic regression, where the latter one is called the L-model.

To begin with, we briefly summarize the basic idea of the LTI-model and, later on in section 3b, we propose further modifications of it for the simultaneous consideration of multiple thresholds.

a. Modified logistic regression model

The introduction of the LTI-model aimed to develop a model for the seamless calibrated combination of the previously described forecast systems Ensemble-MOS and RadVOR, see section 2. Here, statistical calibration refers to an ideal reliability diagram, which is a desirable property of probabilistic forecasts (Murphy and Winkler 1977, 1987). The combined forecast should outperform the individual initial forecasts for all relevant lead times regarding the considered validation scores. For short lead times the extrapolated nowcasting of RadVOR outperforms Ensemble-MOS. However, with increasing lead time, forecast scores (Fig. 1) and reliability diagrams (Fig. 2) for RadVOR drop rapidly in comparison to those for Ensemble-MOS. Note that the green lines in Fig. 1 illustrate the notable improvements achieved by the NN-model introduced in section 3b below.

The LTI-model is based on the standard logistic regression model $f_L: [0,1]^n \rightarrow [0,1]$ for some $n > 1$, where n denotes the number of probabilistic input forecasts to be combined. From a mathematical point of view, the L-model estimates the conditional probability distribution of a dichotomous random variable $Y: \Omega \rightarrow \{0,1\}$, i.e., the occurrence of precipitation above a certain fixed threshold, conditioned on given realizations x_i of a family of random variables $X_i: \Omega \rightarrow [0,1]$ for $i \in \{1, \dots, n\}$, i.e., the probabilistic input forecast models. The random variables $Y, X_1, \dots, X_n$ are defined on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω contains the vectors $(y, x_1, \dots, x_n)$ consisting of all possible forecasts $\mathbf{x}_1, \dots, \mathbf{x}_n$ of the n input forecast models $X_1, \dots, X_n$, and the precipitation occurrence indicator $\mathbf{y}$. Thus,

$$f_L(\mathbf{x}_1, \dots, \mathbf{x}_n) \approx \mathbb{P}(Y=1 | X_1=\mathbf{x}_1, \dots, X_n=\mathbf{x}_n) \quad \text{for all } \mathbf{x}_1, \dots, \mathbf{x}_n \in [0,1].$$

Note that in Schaumann et al. (2020) the special case $n=2$ has been considered, where the probabilities x_1, x_2 originate from RadVOR and Ensemble-MOS, respectively. Then, the L-model combines these two input forecasts and provides calibrated forecast probabilities.

1) TRIANGULAR FUNCTIONS

For the case $n=2$ the L-model has three weights $w_0, w_1, w_2 \in \mathbb{R}$ and is given by

$$f_L(X_1, X_2) = \sigma(w_0 + w_1 X_1 + w_2 X_2), \quad (1)$$

where $\sigma(x) = e^x / (e^x + 1)$ is the logistic function. Note that there are individual weights for each forecast time step. Due to the small number of weights, the L-model is rather limited in how it combines the input forecast models X_1 and X_2 . The weights

FIG. 1. Validation scores for five of the nine considered thresholds for the NN-model (green), LTI-models (gray), Ensemble-MOS (yellow), and RadVOR (blue). (left) Bias, (center) Brier skill score, and (right) reliability. Note, that only one NN-model is trained on all thresholds, while separate LTI-models are trained for each threshold. In both cases modeling is performed for each lead time individually.

merely allow for enough flexibility to weight each forecast based on its overall forecast bias. However, the bias of X_i might vary across the range of possible predictions within the interval $[0, 1]$. This variation in forecast bias is expressed by the reliability diagram of X_i (see Fig. 2 for examples) indicating for which values $x_i \in [0, 1]$ the input forecast model X_i tends to over or underestimate the occurrence of the event that $Y = 1$. Therefore, we proposed to choose the weights for each model X_1, X_2 in dependence of the forecast values x_1 and x_2 . For this, a family of $m + 1$ triangular functions $\phi_j: [0, 1] \rightarrow [0, 1]$

has been defined for some integer $m > 0$ and all $j = 0, \dots, m$ with

$$\phi_j(x) = \max\left\{0, 1 - m\left|x - \frac{j}{m}\right|\right\} \quad \text{for all } x \in [0, 1], \quad (2)$$

and $\sum_{j=0}^m \phi_j(x) = 1$ for $x \in [0, 1]$. Thereby, for $i = 1, 2$, the input forecast model X_i can be encoded as a random vector $\boldsymbol{\phi}(X_i) = [\phi_0(X_i), \dots, \phi_m(X_i)] \in [0, 1]^{m+1}$. Thus, at most two consecutive triangular functions $\phi_j(X_i)$ and $\phi_{j+1}(X_i)$ are nonzero, which is

FIG. 2. Reliability diagrams for the NN-model (green), LTI-models (gray), Ensemble-MOS (yellow), and RadVOR (blue) for five of the nine considered thresholds (columns) and six lead times (rows). Bins with less than 100 exceedance probabilities are omitted. Note, that only one NN-model is trained on all thresholds, while separate LTI-models are trained for each threshold. In both cases modeling is performed for each lead time individually.

the case when the value of X_i falls between and $(j + 1)/m$. Now, instead of the forecast models X_1 and X_2 , we pass the random vectors $\phi(X_1)$ and $\phi(X_2)$ to the L-model and call that the LT-model. For some weights $w_{ij} \in \mathbb{R}$, the latter model is defined as

$$f_{\text{LT}}(X_1, X_2) = f_{\text{L}}[\phi(X_1), \phi(X_2)] = \sigma \left[\sum_{j=1}^m \sum_{i=1}^m w_{ij} \phi_j(X_i) \right]. \quad (3)$$

If we now consider a reliability diagram with $m + 1$ bins, then each weight for a triangular function corresponds to one bin and, therefore, allows the model to produce a calibrated forecast. Note that the LT-model does not use a weight w_0 like in Eq. (1), sometimes called “bias” or “intercept,” because w_0 is redundant as the triangular functions sum up to 1 for any $x \in [0, 1]$.

2) INTERACTION TERMS

The weights in the L- and LT-models are chosen for either one of the two input forecast models X_1 , X_2 , irrespective of the other forecast model. However, it might be sensible to choose different weights, depending on whether both forecasts agree or disagree on the probability of occurrence of the event $Y = 1$. For this, we consider four additional predictors $\gamma_1(X_1, X_2)$, $\dots$, $\gamma_4(X_1, X_2)$, called interaction terms, where $\gamma_1(X_1, X_2) = \sqrt{X_1 X_2}$, $\gamma_2(X_1, X_2) = \sqrt{(1 - X_1) X_2}$, $\gamma_3(X_1, X_2) = \sqrt{X_1 (1 - X_2)}$, and $\gamma_4(X_1, X_2) = \sqrt{(1 - X_1)(1 - X_2)}$.

Like in the previous section, where we have considered the vectors $\phi(X_i) = [\phi_0(X_i), \dots, \phi_m(X_i)]$ for $i = 1, 2$, we now apply triangular functions to the interaction terms and pass the random vectors $\phi[\gamma_i(X_1, X_2)] = \{\phi_0[\gamma_i(X_1, X_2)], \dots, \phi_m[\gamma_i(X_1, X_2)]\}$ for $i = 1, 2, 3, 4$ to the model:

$$\begin{aligned} f_{\text{LTI}}(X_1, X_2) &= f_{\text{L}}\{\phi(X_1), \phi(X_2), \phi[\gamma_1(X_1, X_2)], \dots, \\ &\quad \phi[\gamma_4(X_1, X_2)]\} \\ &= \sigma \left\{ \sum_{j=1}^m \sum_{i=1}^m w_{ij} \phi_j(X_i) + \sum_{k=1}^4 \sum_{j=1}^m w'_{ij} \phi_j[\gamma_k(X_1, X_2)] \right\}, \end{aligned}$$

where $w_{ij}, w'_{ij} \in \mathbb{R}$ are some weights.

3) MODEL TRAINING

For training the LTI-model, a rolling-origin with reoptimization scheme (Armstrong and Grohman 1972) is used in order to simulate the operational conditions. This updating scheme was chosen for several reasons: 1) The model is continuously updated on the newest available data, 2) The continuous updates require data from the past hour only, which makes the update process very fast and efficient in terms of storage space. Other schemes would require us to keep a backlog of old data for up to several months as training data. 3) The rolling-origin update works without a train/test split and therefore allows us to utilize the whole dataset for validation.

In a rolling-origin with reoptimization scheme, the available data are not split up in separate training and test datasets, but it is split by a point in time τ , which represents the “present,” into “past data” and “future data.” In each step of the rolling-origin

scheme the model is trained on “past data” and then validated based on predictions made for time steps ahead of τ , on which the model has not been trained yet. At the end of each step, τ is moved forward in time by one time step. This process is repeated until τ traversed the entire dataset.

b. Generalization of the LTI-model for multiple thresholds, using neural networks

1) OVERVIEW

In the present paper, we propose a number of modifications of the LTI-model, which allow for the combination of forecasts for several thresholds by one common model. Note that the LTI-model, being a modified logistic regression model, can be seen as a specific type of a NN, whereas the softmax layer is a generalization of the logistic regression model (Shah 2020). Therefore, more general variants of NNs are a natural choice to make further extensions of the LTI-model.

In NNs two types of parameters are distinguished: hyper-parameters and trainable parameters. Hyper-parameters determine the architecture of the NN-model (e.g., the numbers and types of layers and neurons). They have to be determined before fitting the trainable parameters to a dataset. The trainable parameters are the weights within each layer. The performance of a specific architecture as defined by the hyper-parameters is highly dependent on the problem to be solved. While there are some guidelines on how to design a NN, it is impossible to tell how the choice of a hyper-parameter affects the performance before fitting and validating the model (Bergstra et al. 2011). For the optimization of the hyper-parameters of our NN-model, we propose an algorithm in the appendix.

In the following sections we give a brief overview of the components of the proposed NN-model. Each component is either a generalization of a part of the LTI-model or a newly added modeling component. For each of them we introduce a number of hyper-parameters, which define the architecture of the NN-model and which have to be determined before training of the model can begin. In total there are 10 hyper-parameters. Each hyper-parameter has a set of possible configurations, which we denote by $H_i = \{c_i^1, \dots, c_i^{m_i}\}$ for $i = 1, \dots, 10$. Thus, the architecture of the NN-models considered in this paper is defined by a vector $\mathbf{h} \in H_1 \times \dots \times H_{10}$, which contains exactly one configuration selected among the valid configurations for each hyper-parameter. For each $\mathbf{h} \in H_1 \times \dots \times H_{10}$, the NN-model consists of the following layers: (i) zero to five convolutional layers, (ii) zero or one dense layer for interaction terms, (iii) zero or one layer consisting of triangular functions, and (iv) one softmax layer. For a schematic representation of the model architecture, see Fig. 3. For a list of the considered hyper-parameters and their configurations, see Table 1. With the exception of convolutional layers, each layer in this NN architecture corresponds to a component of the LTI-model, whereby the dense layer and softmax layer are more general than their LTI-counterparts. For an introduction to deep learning and explanations regarding topics like “activation function” or “convolutional layer,” see Chollet (2017).

FIG. 3. A schematic representation of the network architecture (green) and the input and output data (blue). The arrows indicate the flow of information.

2) CONVOLUTIONAL LAYERS FOR MODEL INPUT

With increasing lead time, it becomes more difficult to make accurate predictions due to increased forecast uncertainty. In particular, this might cause forecast models to be imprecise in the prediction of the location, intensity and duration of precipitation events. For probabilistic forecasts, this imprecision may manifest itself in two ways: (i) The precipitation events are predicted at wrong locations. Note that this is especially relevant for the RadVOR probabilities, which are based on a simple extrapolation and therefore do not take increased uncertainty for longer lead times into account. This results in equally sharp predictions for all lead times. (ii) The probability mass spreads out spatially. In other words, the forecast model predicts lower precipitation probabilities for a larger area to take the increased uncertainty into account. If the probabilistic forecast is based on an ensemble forecast, this effect can be caused by a higher spatial variation between ensemble members, which tends to increase with lead time and reflects the

ensemble forecasts uncertainty regarding the exact location of the precipitation event.

In both cases, it is advantageous for a combination model to be aware of the predictions made at adjacent locations in order to combine two forecasts at a given location. For this, we consider convolutional layers (Zhang et al. 1990), which are commonly used for analyzing image data or data arranged on a regular grid. A NN with convolutional layers takes the information from a neighborhood of adjacent grid points into account, whereas the LTI-model combines forecasts point-wise, i.e., the model output for each grid cell depends only on the individual input forecasts at this location. The size of this neighborhood is determined by the sizes of the convolutional kernels, which map the neighborhood data onto a vector of a specified length.

In the present paper, all kernels are square-shaped and therefore the considered sizes refer to both height and width, e.g., a kernel size of 3 refers to a neighborhood of 3×3 grid cells. Both the size of the kernel and the length of its output are

TABLE 1. Configurations of hyper-parameters considered in this paper (conv. = convolutional).

Layers	Parameter	Description	Configurations
Conv.	H_1	No. of convolutional layers	0, 1, 2, 3, 4, 5
Conv.	H_2	Kernel size	1, 2, 3, 5, 6, 7, 9, 11, 13
Conv.	H_3	No. of convolutional matrices	1, 2, 4, 6, 8, 10, 12, 14, 16
Conv.	H_4	Activation functions	$f_{\text{elu}}, f_{\text{exp}}, f_{\text{lin}}, f_{\text{sigmoid}}, f_{\text{relu}}, f_{\text{tanh}}$
Conv.	H_5	L_2 -regularization	$0, 10^{-7}, 5 \times 10^{-7}, 10^{-6}, 5 \times 10^{-6}, 10^{-5}$
Dense	H_6	No. of neurons	0, 2, 4, 6, 8, 10, 12
Dense	H_7	Activation functions	$f_{\text{elu}}, f_{\text{exp}}, f_{\text{lin}}, f_{\text{sigmoid}}, f_{\text{relu}}, f_{\text{tanh}}$
Dense	H_8	L_2 -regularization	$0, 10^{-7}, 5 \times 10^{-7}, 10^{-6}, 5 \times 10^{-6}, 10^{-5}$
Triangular	H_9	No. of triangular functions	0, 2, 4, 6, 8, 10, 12, 14
—	H_{10}	Optimizer	Adam, Adagrad, Adadelata, Adamax, Nadam

hyper-parameters of the layer. The input of a convolutional layer is a tensor $\mathbf{I} \in \mathbb{R}^{b_x \times b_y \times b_z}$, where b_x and b_y define the size of the forecast grid in the x and y direction, respectively. Furthermore, b_z depicts the number of probabilistic predictions of Ensemble-MOS and RadVOR for nine thresholds each. Thus, $b_z = 18$.

For technical reasons, the NN requires input data given on a rectangular grid. In cases where some of the grid cells are undefined (i.e., where no forecast is available), the region of input data used is restricted to the largest rectangular region free of missing data (i.e., containing no NaN values). Since the forecasts of RadVOR are based on radar measurements that are extrapolated according to a motion vector field, the edges of the respective forecast domain shift in time. Thus, the area containing data of both RadVOR and Ensemble-MOS depends on the magnitude of the motion vector field and may decrease with lead time. Another limitation is the total convolutional size, which has to fit inside the well-defined rectangle.

The total convolutional size is given by the formula: $c_1(c_2 - 1) + 1$ with $c_1 \in H_1$ (number of convolutional layers), $c_2 \in H_2$ (kernel size of each layer). Due to this, some combinations of H_1 and H_2 result in an oversized convolution and, therefore, only combinations with a total convolutional size less than or equal to 13 are used. The latter convolution corresponds to a neighborhood of 260 km, since the forecasts are given on a grid of size $20 \times 20 \text{ km}^2$.

The following hyper-parameters and configurations are considered: number of convolutional layers with $H_1 = \{0, 1, 2, 3, 4, 5\}$, kernel size with $H_2 = \{1, 2, 3, 5, 6, 7, 9, 11, 13\}$, length of the output vector, i.e., the number of convolution matrices used by the layer (Chollet 2017), with $H_3 = \{1, 2, 4, 6, 8, 10, 12, 14, 16\}$, activation functions with $H_4 = \{f_{\text{elu}}, f_{\text{exp}}, f_{\text{lin}}, f_{\text{sigmoid}}, f_{\text{relu}}, f_{\text{tanh}}\}$ and L_2 -regularization strength with $H_5 = \{0, 10^{-7}, 5 \times 10^{-7}, 10^{-6}, 5 \times 10^{-6}, 10^{-5}\}$.

3) DENSE LAYER FOR INTERACTION TERMS

For the LTI-model, the interaction terms were chosen by hand prior to model fitting. In the framework of NNs, the functionality of the interaction terms can be achieved with a densely connected layer of neurons, see e.g., Chollet (2017). In comparison to the LTI-model, a dense layer has the advantage that the shapes of the interaction terms are mostly determined

by trainable parameters. The hyper-parameters for this layer are the number of neurons with $H_6 = \{0, 2, 4, 6, 8, 10, 12\}$, activation functions with $H_7 \{f_{\text{elu}}, f_{\text{exp}}, f_{\text{lin}}, f_{\text{sigmoid}}, f_{\text{relu}}, f_{\text{tanh}}\}$, and L_2 -regularization strength with $H_8 = \{0, 10^{-7}, 5 \times 10^{-7}, 10^{-6}, 5 \times 10^{-6}, 10^{-5}\}$.

4) LAYER FOR TRIANGULAR FUNCTIONS

The role of this layer is the integration of triangular functions into the NN-model. For a definition of and the rationale behind triangular functions, see section 3a. The only hyper-parameter for this layer is the number of triangular functions with $H_9 = \{0, 2, 4, 6, 8, 10, 12, 14\}$. Thus, for each $c \in H_9 \setminus \{0\}$, this layer applies the triangular functions $\phi_0, \dots, \phi_c$ to the output of each neuron of the previous dense layer. In our case, for each $c' \in H_6$, the previous layer returns a vector of scalars $x_1, \dots, x_{c'}$. Hence the triangular functions layer has the output $\phi_0(x_1), \dots, \phi_c(x_1), \dots, \phi_0(x_{c'}), \dots, \phi_c(x_{c'})$ of size $(c + 1)c'$. To our knowledge, such functions are not used in conventional NNs [a similar concept are radial basis function networks (Park and Sandberg 1991)], but they can be manually constructed with the help of Keras backend functions (Keras 2020).

5) SOFTMAX LAYER TO ENSURE CONSISTENT PREDICTIONS FOR ALL THRESHOLDS

To obtain exceedance probabilities for all considered thresholds $t_1, \dots, t_m$ with $0 < t_1 < \dots < t_m < 1$ by means of LTI-models, a separate LTI-model would have to be trained for each threshold. However, this does not guarantee that the probabilities are decreasing monotonously for increasing thresholds, since the separate LTI-models have no knowledge about each other. Hence, an extended model is needed, of which the output is a vector of monotonously decreasing probabilities for the exceedance of the considered thresholds $t_1, \dots, t_m$. Additionally, the data available for one threshold might be useful for the combination of other thresholds, too, since each probability is a point on the discrete cumulative distribution function of the same event and, therefore, they are interlinked.

To ensure that the components of the vector of combined forecasts are decreasing monotonously, we train the neural network on a multilabel classification problem with a Softmax layer (Bridle 1990). This can be seen as a generalization of the logistic regression model, which has been utilized in the LTI-model. While the logistic regression model estimates the

conditional probability distribution of a dichotomous random variable $Y: \Omega \rightarrow \{0, 1\}$, the softmax layer allows for estimating the conditional probability distribution of a random variable $Y': \Omega' \rightarrow \{0, \dots, m\}$ where $m \geq 1$. In our case, Y' models the exceedance of the considered precipitation thresholds at a given location. For this, let $Y: \Omega' \rightarrow \{0, \infty\}$ be the precipitation amount and let $C_i = \{T \in [t_i, t_{i+1})\}$ denote the event that T takes values between the thresholds t_i and t_{i+1} , for $i = 0, \dots, m$. For this, we formally introduce two further thresholds t_0 and t_{m+1} , where $t_0 = 0$ and $t_{m+1} = \infty$. We then put $Y' = i$ if $T \in [t_i, t_{i+1})$, i.e., Y' indicates which of the events $C_0, \dots, C_m$ is occurring.

For the family of pairwise disjoint events $C = \{C_0, \dots, C_m\}$, the NN learns to predict a conditional discrete probability distribution $P_{C,I} = \{\mathbb{P}(C_0|I), \dots, \mathbb{P}(C_m|I)\}$, where $\mathbb{P}(C_i|I)$ is the conditional probability for the occurrence of event C_i given the model input I . Then, for each $j \in \{1, \dots, m\}$, the conditional probability of the event that $T \geq t_j$ can be computed by $\mathbb{P}(T \geq t_j|I) = 1 - \sum_{i=0}^{j-1} \mathbb{P}(C_i|I)$. Clearly, from this it follows by definition that $\mathbb{P}(T \geq t_j|I) \geq \mathbb{P}(T \geq t_k|I)$ if $t_j < t_k$, and therefore the predictions of the NN are consistent for all thresholds.

6) OPTIMIZER FOR TRAINABLE PARAMETERS

Another difference between the LTI-model introduced in Schaumann et al. (2020) and the NN-model considered in the present paper is the choice of the optimizer. Note that the optimizer controls how the trainable parameters change during the model training. This has a large influence on the model performance. For a more detailed introduction to the operating principle of optimizers, we refer to Chollet (2017).

In our previous paper, a stochastic gradient descent with a constant learning rate has been used as optimizer for the LTI-model. While this is sufficient for the (relatively small) number of parameters of the LTI-model, the NN-model considered in the present paper requires a more sophisticated optimizer, due to weaknesses of the classical stochastic gradient descent. Depending on the network architecture and the training dataset, gradients for some weights might “vanish” at some point in training, see Glorot et al. (2011). This means that parts of the NN might receive only small updates or a sparse number of updates leading to stagnation of the training process. More recent optimizers address this problem by various means, e.g., gradient descent with momentum (Sutskever et al. 2013) or adaptive learning rates. This ensures, first, to scale up small gradients back to a reasonable size or, second, to compensate for a sparse number of updates of a weight. In the present paper, five different optimizers are investigated. All of them are based on adaptive learning rates, which leads to a tenth hyper-parameter with $H_{10} = \{\text{Adam}, \text{Adagrad}, \text{Adadelata}, \text{Adamax}, \text{Nadam}\}$. For more details on how each optimizer works, see Kingma and Ba (2015), Duchi et al. (2011), Zeiler (2012), and Dozat (2016).

c. Training and validation

In the previous section, several modifications of the LTI-model have been discussed. Each of them introduces one or more hyper-parameters, needing to be determined before the NN can be trained. For this, a hyper-parameter optimization

algorithm is employed, which is explained in the appendix in more detail.

The training process of the NN consists of several steps: (i) Training of different model architectures for the hyper-parameter search, using data of the period from 1 April 2016 to 31 May 2016, (ii) performance evaluation of model architectures for the hyper-parameter search, using data from 1 June 2016 to 30 June 2016, (iii) picking the best architecture based on evaluation results, (iv) training of the best architecture (without validation), using data from 1 April 2016 to 31 May 2016, and (v) rolling origin update and validation with best architecture, using data from 1 July 2016 to 30 September 2016.

Note that the data are split such that the choice of the best performing architecture and its validation are based on two different time intervals. Otherwise, the validation results would be biased toward better scores since they would not reflect the uncertainty about the optimal choice of hyper-parameters.

4. Results

a. Influence of hyper-parameters

1) OPTIMIZER

For the hyper-parameter optimization, about 18000 model architectures have been evaluated for each considered lead time. The choice of the optimizer $c \in H_{10}$ has, by far, the largest impact on model performance (see section a in the appendix). The distribution of model performance for each optimizer is depicted in Fig. 4. Since Adam, Adamax, and Nadam outperform Adagrad and Adadelata for almost all model configurations, we will focus on results only with regard to Adam, Adamax, and Nadam in the following discussion.

2) CONVOLUTIONAL LAYERS

Furthermore, the model performance is highly affected by the number of convolutional layers (selected from the set H_1) and their kernel size (selected from H_2), see Fig. 5, where the results of the hyper-parameter optimization are visualized. While the difference between individual configurations is less pronounced for lead times +1 h, larger total convolutional sizes perform better for longer lead times. Thus, in situations of increased forecast uncertainty (e.g., for longer lead times), an improved forecast skill is achieved when considering more adjacent grid cells. This is in agreement with the ideas of Theis et al. (2005) and Schwartz and Sobash (2017).

The activation functions f_{elu} , f_{linear} , f_{relu} and f_{tanh} seem to perform similarly well when compared to each other. In contrast, the functions $f_{\text{exponential}}$ and f_{sigmoid} exhibit a much worse behavior, in particular for model architectures with many convolutional layers. As an exception, f_{sigmoid} does not show this behavior for the lead time +6 h and even performs best out of all considered activation functions.

Models with larger lengths of output vectors (selected from H_3) tend to perform better; however, the difference is clearly pronounced only up to a vector length of 4.

FIG. 4. Distribution of model performance (categorical cross-entropy) for the considered configurations of the first nine hyper-parameters and the five optimizers, for the following lead times: (left) +1 and (right) +6 h. The median performance is depicted by a blue line. Due to very long tails, the distributions are only shown between the respective 5th and 95th percentiles.

It should be noted that the number of convolutional layers (selected from H_1) and their kernel size (selected from H_2) affect the output size of the NN in two different ways: (i) As explained in section 3b, input data passed to the NN needs to be rectangular-shaped and defined on each grid cell in order to enable the NN to learn and to make a prediction. Note that the passed data domain underlies a large variability, due to irregular boundaries of and occasionally missing data within the input forecasts. The total convolutional size determines the possible minimum edge length of the data domain. (ii) Furthermore, the total convolutional size determines the size of the input area which is mapped to an output value. For example, for a total convolutional size of 5, a model input $\mathbf{I}$ of size $15 \times 30 \times 18$ is mapped to a model output of size $11 \times 26 \times 9$. Note that the third dimension contains the predictions for the nine thresholds considered in this paper. Since the model input $\mathbf{I}$ contains data from both initial forecasts, it is of length 18 along the third axis while the model output contains the combined forecast only and therefore it is of half the size.

Due to these effects, the amount of available data used for training and validation depends on the total convolutional size. This might be an explanation for the decreased performance of the four configurations (1, 13), (4, 4), (3, 5), (2, 7) of $H_1 \times H_2$ with the largest total convolutional size of 13, since they have much less output to be trained and validated on, see Fig. 5.

For a total convolutional size of 9 (used for lead times +1, +2, +4, and +5 h, see Table 2), the passed data domain $\mathbf{I}$ consists, on average, of about 292 grid cells in the considered time period (July, August, September 2016). This results in 636 126 data points in total. In comparison to this, for a total convolutional size of 11 (used for lead times +3 h and +6 h, see Table 2), the passed data domain

consists, on average, of about 160 grid cells with 344 431 data points in total.

3) TRIANGULAR FUNCTIONS AND INTERACTION TERMS

In Fig. 6 the effect of triangular functions and the dense layer (interaction terms) on the model performance is visualized. In general, one might expect that more neurons in the dense layer perform better than less. Thus, at first glance, it seems to be counterintuitive that two neurons perform worse than no neurons at all. A likely explanation of this is that few neurons act as a bottleneck that restricts the amount of information the NN can pass through the dense layer, whereas for zero neurons the layer and, therefore, the bottleneck is removed. Regarding triangular functions, typically 3–9 of them perform best for all lead times. However, some of these configurations can perform worse for specific lead times, e.g., 11 triangular functions for +4 h, and 5 triangular functions for +5 h, see Fig. 6. The results visualized in Fig. 6 show that model performance for lead time +1 h behaves differently in comparison to the model performance for longer lead times. See also Fig. 5, where this effect can be observed, too.

4) REMAINING HYPER-PARAMETERS

The remaining hyper-parameters are the activation functions (selected from H_7) for the dense layer and the L_2 -regularization strengths for the convolutional layers (selected from H_5) and the dense layer (selected from H_8). However, these parameters do not seem to affect the model performance in any significant way.

b. Model validation and comparison

The performances of the NN-model, LTI-model, EnsembleMOS and RadVOR have been evaluated on data for the months July, August, and September 2016. Within this period of time

FIG. 5. Median model performance (categorical cross-entropy, lower values are better) for different hyper-parameter settings regarding the convolutional layers, and for several lead times: (from top to bottom) +1, +3, +5, and +6 h. The color of each tile depicts the median model performance in dependence on number and kernel size of the convolutional layers (x axis) and activation function (y axis). Whereby the x-axis labels correspond to the kernel size of each layer, e.g., 5–5–5 stands for three layers with a kernel size of five each. The number in each tile indicates the number of evaluations made by the hyper-parameter optimization algorithm.

and the passed domain I (depending on the total convolutional size of the NN-model), the considered precipitation thresholds were exceeded as described below by the numbers of upcrossings (and corresponding relative frequencies in parentheses).

- (i) For a total convolutional size of 9 we have 39 303 (6.18%) for 0.1 mm, 31 284 (4.92%) for 0.2 mm, 26 402 (4.15%) for 0.3 mm, 20 300 (3.19%) for 0.5 mm, 16 459 (2.59%) for 0.7 mm, 12 677 (1.99%) for 1 mm, 6264 (0.98%) for 2 mm, 3494 (0.55%) for 3 mm, and 1375 (0.22%) for 5 mm.

TABLE 2. Selected configurations of hyper-parameters for different lead times (conv. = convolutional; reg. = regularization).

Lead time	No. of conv. layers	Kernel size	Conv. activation	Conv. reg.	Conv. output length
+1 h	4	3	elu	5×10^7	14
+2 h	4	3	elu	5×10^6	6
+3 h	5	3	relu	0	8
+4 h	1	9	elu	0	4
+5 h	1	9	elu	5×10^6	12
+6 h	2	6	sigmoid	0	4
Lead time	No. of neurons in dense layer	Dense activation	Dense reg.	No. of triangular functions	Optimizer
+1 h	12	sigmoid	0	9	Nadam
+2 h	12	tanh	0	5	Nadam
+3 h	12	exponential	1×10^6	3	Adamax
+4 h	6	tanh	0	5	Adamax
+5 h	10	sigmoid	0	3	Adamax
+6 h	10	relu	1×10^5	5	Adamax

(ii) For a total convolutional size of 11 we have 20 420 (5.93%) for 0.1 mm, 16 195 (4.70%) for 0.2 mm, 13 561 (3.94%) for 0.3 mm, 10 330 (3.00%) for 0.5 mm, 8358 (2.43%) for 0.7 mm, 6386 (1.85%) for 1 mm, 3046 (0.88%) for 2 mm, 1652 (0.48%) for 3 mm, 677 (0.20%) for 5 mm.

For each lead time, the parameter configurations are shown in Table 2, which were chosen by the hyper-parameter optimization algorithm explained in the appendix. For each chosen NN-architecture validation scores and reliability diagrams are shown in Figs. 1 and 2. Note that for some hyper-parameters several configurations perform equally well, which leads to random fluctuations between the choices for different lead times.

While one NN-model is trained on all thresholds for each lead time, the LTI-model combines probabilities for only one threshold and, therefore, its validation scores in Figs. 1 and 2 are based on separate model specifications for each combination of lead time and threshold ($6 \times 5 = 30$ in total). While the NN-model requires a NaN-free rectangular dataset, the LTI-model can be applied on any grid point without missing data. For the results in this paper, however, the LTI-model has been fitted on the same rectangular dataset as the NN-model to make the validation results of both models more comparable.

It can be seen that both combination models generate less biased and more calibrated predictions with a higher Brier skill score in comparison to both initial forecasts provided by Ensemble-MOS and RadVOR, for all lead times and thresholds. Although the RadVOR forecasts are only provided up to +2 h, the forecasts of both combination models have better scores up to +6 h. For more details on the used validation scores, see Wilks (2006).

Similar to the LTI-model, the NN-model has improved reliability diagrams in comparison to both initial forecasts, see Fig. 2. As expected, the reliability of forecasting models decreases with increasing lead times and increasing thresholds. To keep the reliability diagrams calibrated, the combination models learn to lower their predictions accordingly to not overestimate the occurrence of precipitation, which leads to shorter curves in the reliability diagram for longer lead

times and higher thresholds in comparison to both initial input models.

When comparing the NN-model with the LTI-model, it can be seen that the NN-model achieves better or equally good results for almost all considered lead times and validation scores. This improvement obtained by the NN-model is likely due to its more sophisticated architecture, which allows the NN-model to take data for all thresholds and also adjacent grid points into consideration. Moreover, this improvement is especially notable because in contrast to the LTI-model the NN-model produces consistent probabilities which is an additional constraint to be satisfied.

To test if it is actually necessary to run the hyper-parameter optimization algorithm for each lead time, we trained the chosen architecture for +1 h on the other five lead times, too. A visual comparison between the validations scores given in Fig. 1 and the results for the +1-h architecture showed only slight performance differences. However, the reliability diagrams seem to be less calibrated, see Fig. 2. We, therefore, decided to use a separate network architecture for each lead time.

c. Combination example

In Fig. 7, forecasting results obtained by the combination of input data from Ensemble-MOS and RadVOR are shown, as an example, for the hour 0600–0700 UTC 21 July 2016 and for three lead times (+1, +2, +3 h). To our knowledge, threshold probabilities have not been depicted in the literature with such kind of diagram before. Due to a larger total convolutional size of the NN-model for +3 h, the size of the output is smaller. For shorter lead times, the forecast of the NN-model closely resembles the forecast of RadVOR, while for increasing lead times, the predictions become more smooth and more dependent on the Ensemble-MOS prediction.

5. Conclusions

a. Summary of results

In this paper we presented NN architectures for the combination of two probabilistic forecasts, where we consider

FIG. 6. Median model performance (categorical cross-entropy, lower values are better) for different hyper-parameter settings regarding triangular functions (x axis) and interaction terms (y axis), and for the following lead times (top left) +1, (top right) +2, (middle left) +3, (middle right) +4, (bottom left) +5, and (bottom right) +6 h. The color depicts the median model performance. The number in each tile indicates the number of evaluations made by the hyper-parameter optimization algorithm.

several precipitation thresholds simultaneously. The architectures chosen by the hyper-parameter optimization algorithm show improvements for all considered validation scores across all thresholds, and calibrate the resulting probabilities. Like the previously developed LTI-model, the NN-model considered in the present paper improves forecast scores also for lead times longer than +2 h, although RadVOR forecasts were only provided up to +2 h.

The proposed hyper-parameter optimization algorithm worked as intended and yielded architectures with improved categorical cross-entropy compared to hand-picked architectures, which also led to improvements in all other validation scores considered in this paper, and to calibrated reliability diagrams in particular.

In a direct comparison between the LTI-model and the NN-model, the NN-model performs better than or equally well as

FIG. 7. Combination example in a NaN-free rectangle for (top) +1, (middle) +2, and (bottom) +3 h for the hour 0600–0700 UTC 21 Jul 2016. The columns correspond to the forecast models and the ground truth: (first column) Ensemble-MOS, (second column) RadVOR, (third column) neural net, and (fourth column) radar measurement. At each grid point (marked by a gray circle) several colored circles are stacked on top of each other. In the first three columns, the size of each colored circle corresponds to the exceedance probability for a

the LTI-model with respect to all considered validation scores. This is despite the fact that the NN-model must predict consistent exceedance probabilities for several thresholds, which is an additional constraint to be satisfied and should be kept in mind when comparing both models.

For practical purposes, it should be taken into account that while the NN-model outperforms the LTI-model, the LTI-model is not constrained to a NaN-free rectangular dataset. Additionally, the LTI-model has only a few hyper-parameters, due to its simpler design, which makes it much easier to train the LTI-model.

b. Outlook and possible next steps

According to the results of the hyper-parameter search performed in the present paper, some hyper-parameters seem to be much more important than others. Thus, it might be possible to further improve the architecture by adapting the search space. Since the optimizer and the number and size of convolutional layers have the largest influence on model performance, additional optimizers and convolutional layer combinations should be investigated. Thus, in a forthcoming paper, we will investigate the numerical stability of the hyper-parameter optimization algorithm and how the chosen architectures in Table 2 compare, e.g., to the second best architecture for each lead time.

Due to the restriction that the input of the NN must be rectangular and free of missing values, it should be considered to generate valid values by interpolation at grid points with missing values, or to pass an mask to the NN as additional input in order to specify, which values are valid. This would allow training without cropping of the data and also increase the area for which predictions can be made.

Furthermore, it would be interesting to investigate how additional information might affect the quality of combination, e.g., by increasing the resolution of the grid, by passing ensemble members directly to the NN without aggregation to probabilities, by adding an orography map to the input, or by including additional meteorological indicators. Given that a new dataset contains enough precipitation events for higher thresholds, the list of considered thresholds could be expanded.

APPENDIX

Hyper-Parameter Optimization

To choose hyper-parameters by means of a systematic approach, various optimization algorithms have been developed, attempting to find correlations between the hyper-parameters of a model and its performance by evaluating a number of

different network architectures. The following problems arise in such algorithms. (i) Curse of dimensionality: For each additional hyper-parameter, the size of the search space grows exponentially. (ii) Training time: Depending on the size of the model, the size of the training dataset, and the available hardware, the evaluation of a network architecture might take a considerable amount of computation time. (iii) Interactions between hyper-parameters: It is not enough to consider each hyper-parameter separately, because the best choice for some hyper-parameter might depend on the chosen configurations of other hyper-parameters. (iv) Nondeterministic model performance: The fitting of a NN is a nondeterministic process and the weights of a model might not converge to the same optimum in repeated runs. This means that a single evaluation of a network architecture might not reflect the actual performance of the architecture in general. For an introduction to hyper-parameter optimization in general and the concepts mentioned above in particular, see [Hutter et al. \(2019\)](#). To our knowledge, the following algorithm has not been proposed before.

To find a satisfactory model architecture despite the problems listed above, the proposed algorithm works according to the principle of Exploration and Exploitation, which is also explained in [Hutter et al. \(2019\)](#): At the beginning of the search, architectures across the whole search space are evaluated. With an increasing number of evaluations, promising candidates are prioritized.

In the following, we consider the search space $H = H_1 \times \dots \times H_n$ being the Cartesian product of a family of domains $H_1, \dots, H_n$ of n hyper-parameters for some integer $n \geq 1$. The set H_i consists of $m_i \geq 1$ available configurations of the i th hyper-parameter, i.e., $H_i = \{c_i^1, \dots, c_i^{m_i}\}$ for each $i = 1, \dots, n$.

a. Performance of an evaluation

In each iteration of the hyper-parameter optimization algorithm, the performance $f(\mathbf{h})$ of a model architecture specified by $\mathbf{h} = (c_1, \dots, c_n) \in H$ is evaluated, where $c_i \in H$ for each $i = 1, \dots, n$. The model architectures considered in this paper were trained for six epochs on a training dataset (April + May 2016) and validated after each epoch on a separate validation dataset (June 2016). Note that an epoch refers to one pass-through of the training dataset in the training process. Each batch consists of data for 1 h. We define the performance $f(\mathbf{h})$ of a configuration $\mathbf{h} \in H$ as the smallest model error achieved in any of the six epochs, whereas the model error is determined by the loss function ([Chollet 2017](#)) of the NN. Since we consider a classification problem in this paper, the loss function “categorical cross-entropy” is used ([Alla and Adari 2019](#)).

←

given threshold which is indicated by the color of the circle (see legend). In the fourth column, depicting the ground truth, the circle color indicates the highest threshold that has been exceeded. The maximum circle size corresponds to probability one while circles corresponding to probability zero vanish since they are of size zero. The corresponding probability is proportional to the circle area, not the radius. The gray solid lines indicate the borders of federal states of Germany with Hesse and Saxony-Anhalt in the center.

To find out which number of epochs is sufficient, the model errors for 20 epochs have been determined for a number of model architectures. However, for most model architectures the minimum model error converged within the first five epochs. Therefore, we considered six epochs in order to derive the results obtained in this paper.

b. Selection of new hyper-parameter configurations

A common strategy to pick model configurations for evaluation is the so-called random search method, where a certain probability distribution, e.g., the uniform distribution, is considered on the search space H from which new model architectures are sampled (Hutter et al. 2019).

The idea of the algorithm presented in this section is to start with a random search. However, after having made a number of evaluations, we can already estimate how the configuration $c_i \in H$ of a single hyper-parameter, for some $i \in \{1, \dots, n\}$, affects the performance of the model architecture $\mathbf{h} = (c_1, \dots, c_n)$. Based on this information, the probability distribution on H , from which further model architectures are sampled, can be adapted to favor model architectures which are more likely to perform well. With an increasing number of evaluations, the same concept can be applied to an increasing number of j hyper-parameters, where $j \in \{2, \dots, n\}$, in order to find out which configurations $\mathbf{h}' \in H'_j = H_{i_1} \times \dots \times H_{i_j}$ for some subset $J = \{i_1, \dots, i_j\} \subseteq \{1, \dots, n\}$ perform well in combination with each other. In the following, we sometimes write $H'_{\mathbf{h}'}$ instead of H'_j , in cases where we want to emphasize that a specific partial architecture $\mathbf{h}'$ belongs to the set H'_j . Furthermore, let $H^* = \bigcup_{J \in \mathcal{P}(\{1, \dots, n\})} H'_J$ denote the set of all (partial) model architectures.

1) DEFINITIONS

The $(k + 1)$ th choice of the model configuration $\mathbf{h}_{k+1} \in H$, which is to be evaluated next, is made based on the set of previously evaluated configurations $E = \{\mathbf{h}_1, \dots, \mathbf{h}_k\}$, for which the performances $f(\mathbf{h}_i), i = 1, \dots, k$ have been determined as described in section 4a. Let $\mathbf{h}' \in H'_j$ be a partial model architecture for a subset of hyper-parameters with indices $J = \{i_1, \dots, i_j\} \subseteq \{1, \dots, n\}$. Then define

$$E_{\mathbf{h}'} = \{\mathbf{h} \in E : \mathbf{h}' \subseteq \mathbf{h}\} \quad (\text{A1})$$

as the set of all evaluated hyper-parameter configurations $\mathbf{h}$ that share the same configurations with the partial architecture $\mathbf{h}'$. For a given integer $s \in \mathbb{N} = \{1, 2, \dots\}$ we define

$$E_s = \left\{ \mathbf{h}' \in H^* : |E_{\mathbf{h}'}| \geq s, \left| E_{\mathbf{h}' \cup \{c_i\}} \right| < s \text{ for all } c_i \in H_i, i \in \{1, \dots, n\} \right\}, \quad (\text{A2})$$

which is the set of partial architectures $\mathbf{h}'$ with at least s evaluations and for which all extensions $\mathbf{h}' \cup \{c_i\}$ have less evaluations than s . In other words, E_s contains the largest partial architectures for which a minimum number s of evaluations exists.

When considering a partial architecture $\mathbf{h}' \in H'_j$ for evaluation, we are not only interested whether E contains enough

data to estimate $f(\mathbf{h}')$, but also if all partial architectures in $H'_{\mathbf{h}'}$ have enough evaluations. Hence we define

$$E_s^p = \left\{ \mathbf{h}' \in E_s : |E_{\mathbf{h}'}| \leq \min_{\mathbf{g}' \in H'_{\mathbf{h}'}} |E_{\mathbf{g}'}|^p \right\} \quad (\text{A3})$$

for a given value $p \in [1, \infty)$, i.e., E_s^p is a subset of E_s containing all partial architectures $\mathbf{h}'$ with a number of evaluations $|E_{\mathbf{h}'}|$ below an upper bound, which depends on the partial architecture $\mathbf{g}' \in H'_{\mathbf{h}'}$ with the smallest number of evaluations $|E_{\mathbf{g}'}|$.

Furthermore, we define the median performance $M_E(\mathbf{h}')$ for a partial model architecture $\mathbf{h}'$ as

$$M_E(\mathbf{h}') = \text{median}(\{f(\mathbf{h}) : \mathbf{h} \in E_{\mathbf{h}'}\}), \quad (\text{A4})$$

and, finally, the set $\Delta_\delta(\mathbf{h}')$ of partial architectures $\mathbf{g}'$, which share $|\mathbf{h}'| - \delta$ configurations with $\mathbf{h}'$ as

$$\Delta_\delta(\mathbf{h}') = \{\mathbf{g}' \in H : |\mathbf{h}' \cap \mathbf{g}'| = |\mathbf{h}'| - \delta\} \quad \text{for } \delta \in \mathbb{N}. \quad (\text{A5})$$

2) THE ALGORITHM

In this section we explain how we determine an initial partial architecture $\mathbf{h}'_0$ and how we pick subsequent partial architectures $\mathbf{h}'_1, \dots, \mathbf{h}'_k$ until their union is a full architecture that can be evaluated next and added to the set E .

For given values of s and p , we sample from the set of partial architectures E_s^p . Initially, the set E_s^p is empty, because E is empty since no architectures have been evaluated, yet. Note that E_s^p can also be empty due to the upper bound determined by the parameter p . In these cases, a partial architecture $\mathbf{h}'_0 = \{c_i\}$ with a random configuration $c_i \in H_i$ for a random hyper-parameter H_i is chosen. Otherwise, we sample $\mathbf{h}'_0$ from the set E_s^p with probability $P_{E_s^p}$, where P_F is defined as

$$P_F(\mathbf{h}') = \frac{2^{-M_E(\mathbf{h}')/d}}{\sum_{\mathbf{g}' \in F} 2^{-M_E(\mathbf{g}')/d}} \quad (\text{A6})$$

for any partial architecture $\mathbf{h}' \in F$, a given set F of partial architectures and some $d > 0$, which is another parameter of the algorithm, along with s and p . Note that P_F is defined such that a partial architecture $\mathbf{h}'$ is twice as likely to be chosen than $\mathbf{g}'$ if $M_E(\mathbf{h}') = M_E(\mathbf{g}') + d$. Furthermore, once enough evaluations have been made such that E_s^p is nonempty, the algorithm will always pick a partial architecture from E_s^p . Without the upper bound, the first few partial architectures in E_s would be sampled ad infinitum, while other partial architectures, which have not been sampled often enough yet, would not be sampled at all. Therefore it is necessary to include the upper bound in E_s^p for the number of evaluations.

We now have the first part $\mathbf{h}'_0$ of the architecture $\mathbf{h}$, which we want to evaluate. Next, we successively determine more parts $\mathbf{h}'_1, \dots, \mathbf{h}'_k$ until their union is a complete architecture $\mathbf{h} = \mathbf{h}'_0 \cup \dots \cup \mathbf{h}'_k \in H$. For this, let $\mathbf{h}'_j = \mathbf{h}'_0 \cup \dots \cup \mathbf{h}'_j$ be the union of all partial architectures up to $\mathbf{h}'_j$.

For $j \in \{1, \dots, k\}$ we iteratively sample $\mathbf{h}'_j$ from $E_s^p \cap \Delta_\delta(\mathbf{h}'_{j-1})$ with probability $P_{E_s^p \cap \Delta_\delta(\mathbf{h}'_{j-1})}(\mathbf{h}'_j)$ with the smallest $\delta \in \mathbb{N}$ for which $E_s^p \cap \Delta_\delta(\mathbf{h}')$ is not empty. In other words,

we choose $\mathbf{h}'_j$ such that it has at least one new configuration and also the largest possible overlap with $\bar{\mathbf{h}}'_{j-1}$, and that the new configurations are likely to perform well in combination with the previously chosen configurations. If $E_s^p \cap \Delta_\delta(\bar{\mathbf{h}}'_{j-1})$ is empty for all $\delta \in \mathbb{N}$, the partial configuration $\mathbf{h}'_j = \{c_i\}$ consists of a single randomly chosen configuration $c_i \in H_i$, where H_i is a random hyper-parameter, for which it holds that $H_i \cap \bar{\mathbf{h}}'_{j-1} = \emptyset$.

Once enough architectures have been evaluated and we want to pick the best architecture based on E , we follow the same steps as described above, but instead of sampling partial architectures from E_s^p with probability $P_{E_s^p}$, we pick the partial architectures $\mathbf{h}' \in E_s$ with the lowest $M_E(\mathbf{h}')$ instead.

REFERENCES

- Alla, S., and S. K. Adari, 2019: *Beginning Anomaly Detection Using Python-Based Deep Learning*. Springer, 427 pp.
- Armstrong, J. S., and M. C. Grohman, 1972: A comparative study of methods for long-range market forecasting. *Manage. Sci.*, **19**, 211–221, <https://doi.org/10.1287/mnsc.19.2.211>.
- Ben Bouallègue, Z., 2013: Calibrated short-range ensemble precipitation forecasts using extended logistic regression with interaction terms. *Wea. Forecasting*, **28**, 515–524, <https://doi.org/10.1175/WAF-D-12-00062.1>.
- Bergstra, J. S., R. Bardenet, Y. Bengio, and B. Kégl, 2011: Algorithms for hyper-parameter optimization. *Advances in Neural Information Processing Systems*, J. Shawe-Taylor et al., Eds., Curran Associates, Inc., 2546–2554.
- Bouttier, F., and H. Marchal, 2020: Probabilistic thunderstorm forecasting by blending multiple ensembles. *Tellus*, **72A**, 1–19, <https://doi.org/10.1080/16000870.2019.1696142>.
- Bowler, N. E., C. E. Pierce, and A. W. Seed, 2006: STEPS: A probabilistic precipitation forecasting scheme which merges an extrapolation nowcast with downscaled NWP. *Quart. J. Roy. Meteor. Soc.*, **132**, 2127–2155, <https://doi.org/10.1256/qj.04.100>.
- Bridle, J. S., 1990: Probabilistic interpretation of feedforward classification network outputs, with relationships to statistical pattern recognition. *Neurocomputing*, F. F. Soulié and J. Héroult, Eds., Springer, 227–236.
- Brunet, G., S. Jones, and P. M. Ruti, 2015: *Seamless Prediction of the Earth System: From Minutes to Months*. WMO-1156, WMO, 471 pp.
- Chollet, F., 2017: *Deep Learning with Python*. Manning Publications, 384 pp.
- Dozat, T., 2016: Incorporating Nesterov momentum into Adam. *Int. Conf. on Learning Representations, ICLR 2016, Conf. Track Proc.*, San Juan, Puerto Rico, ICLR, 4 pp., <https://openreview.net/pdf/OM0jvwB8Jlp57ZJtNEZ.pdf>.
- Duchi, J., E. Hazan, and Y. Singer, 2011: Adaptive subgradient methods for online learning and stochastic optimization. *J. Mach. Learn. Res.*, **12**, 2121–2159.
- Foresti, L., and A. Seed, 2014: The effect of flow and orography on the spatial distribution of the very short-term predictability of rainfall from composite radar images. *Hydrol. Earth Syst. Sci.*, **18**, 4671–4686, <https://doi.org/10.5194/hess-18-4671-2014>.
- Gebhardt, C., S. Theis, M. Paulat, and Z. B. Bouallègue, 2011: Uncertainties in COSMO-DE precipitation forecasts introduced by model perturbations and variation of lateral boundaries. *Atmos. Res.*, **100**, 168–177, <https://doi.org/10.1016/j.atmosres.2010.12.008>.
- Germann, U., and I. Zawadzki, 2002: Scale dependence of the predictability of precipitation from continental radar images. Part I: Description of the methodology. *Mon. Wea. Rev.*, **130**, 2859–2873, [https://doi.org/10.1175/1520-0493\(2002\)130<2859:SDOTPO>2.0.CO;2](https://doi.org/10.1175/1520-0493(2002)130<2859:SDOTPO>2.0.CO;2).
- Glorot, X., A. Bordes, and Y. Bengio, 2011: Deep sparse rectifier neural networks. *Proc. 14th Int. Conf. on Artificial Intelligence and Statistics, JMLR Workshop and Conf. Proc.*, Fort Lauderdale, FL, PMLR, 315–323.
- Gneiting, T., A. E. Raftery, A. H. Westveld III, and T. Goldman, 2005: Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Mon. Wea. Rev.*, **133**, 1098–1118, <https://doi.org/10.1175/MWR2904.1>.
- Golding, B., 1998: Nimrod: A system for generating automated very short range forecasts. *Meteor. Appl.*, **5**, 1–16, <https://doi.org/10.1017/S1350482798000577>.
- Haiden, T., A. Kann, C. Wittmann, G. Pistotnik, B. Bica, and C. Gruber, 2011: The Integrated Nowcasting through Comprehensive Analysis (INCA) system and its validation over the eastern Alpine region. *Wea. Forecasting*, **26**, 166–183, <https://doi.org/10.1175/2010WAF2222451.1>.
- Hamill, T. M., R. Hagedorn, and J. S. Whitaker, 2008: Probabilistic forecast calibration using ECMWF and GFS ensemble reforecasts. Part II: Precipitation. *Mon. Wea. Rev.*, **136**, 2620–2632, <https://doi.org/10.1175/2007MWR2411.1>.
- Heizenreder, D., P. Joe, T. Hewson, L. Wilson, P. Davies, and E. de Coning, 2015: Development of applications towards a high-impact weather forecast system. *Seamless Prediction of the Earth System: From Minutes to Months*. G. Brunet, S. Jones, and P. M. Ruti, Eds., 419–443.
- Hess, R., 2020: Statistical postprocessing of ensemble forecasts for severe weather at Deutscher Wetterdienst. *Nonlinear Processes Geophys.*, **27**, 473–487, <https://doi.org/10.5194/npg-27-473-2020>.
- Hutter, F., L. Kotthoff, and J. Vanschoren, Eds., 2019: *Automated Machine Learning: Methods, Systems, Challenges*. Springer, 219 pp.
- Johnson, A., and X. Wang, 2012: Verification and calibration of neighborhood and object-based probabilistic precipitation forecasts from a multimodel convection-allowing ensemble. *Mon. Wea. Rev.*, **140**, 3054–3077, <https://doi.org/10.1175/MWR-D-11-00356.1>.
- Keras, 2020: Keras API reference: Backend utilities. Accessed 20 August 2020, https://keras.io/api/units/backend_utils/.
- Kingma, D. P., and J. Ba, 2015: Adam: A method for stochastic optimization. *ICLR 2015 Third Int. Conf. on Learning Representations*, San Diego, CA, ICLR, <http://arxiv.org/abs/1412.6980>.
- Kober, K., G. Craig, C. Keil, and A. Dörnbrack, 2012: Blending a probabilistic nowcasting method with a high-resolution numerical weather prediction ensemble for convective precipitation forecasts. *Quart. J. Roy. Meteor. Soc.*, **138**, 755–768, <https://doi.org/10.1002/qj.939>.
- Murphy, A. H., and R. L. Winkler, 1977: Reliability of subjective probability forecasts of precipitation and temperature. *J. Roy. Stat. Soc.*, **26C**, 41–47.
- , and —, 1987: A general framework for forecast verification. *Mon. Wea. Rev.*, **115**, 1330–1338, [https://doi.org/10.1175/1520-0493\(1987\)115<1330:AGFFV>2.0.CO;2](https://doi.org/10.1175/1520-0493(1987)115<1330:AGFFV>2.0.CO;2).
- Nerini, D., L. Foresti, D. Leuenberger, S. Robert, and U. Germann, 2019: A reduced-space ensemble Kalman filter approach for flow-dependent integration of radar extrapolation nowcasts and NWP precipitation ensembles. *Mon. Wea. Rev.*, **147**, 987–1006, <https://doi.org/10.1175/MWR-D-18-0258.1>.

- Nicolis, C., R. A. Perdigao, and S. Vannitsem, 2009: Dynamics of prediction errors under the combined effect of initial condition and model errors. *J. Atmos. Sci.*, **66**, 766–778, <https://doi.org/10.1175/2008JAS2781.1>.
- Palmer, T. N., 2002: The economic value of ensemble forecasts as a tool for risk assessment: From days to decades. *Quart. J. Roy. Meteor. Soc.*, **128**, 747–774, <https://doi.org/10.1256/0035900021643593>.
- Park, J., and I. W. Sandberg, 1991: Universal approximation using radial-basis-function networks. *Neural Comput.*, **3**, 246–257, <https://doi.org/10.1162/neco.1991.3.2.246>.
- Piper, D., M. Kunz, F. Ehmele, S. Mohr, B. Mühr, A. Kron, and J. Daniell, 2016: Exceptional sequence of severe thunderstorms and related flash floods in May and June 2016 in Germany—Part I: Meteorological background. *Nat. Hazards Earth Syst. Sci.*, **16**, 2835–2850, <https://doi.org/10.5194/nhess-16-2835-2016>.
- Schaumann, P., M. de Langlard, R. Hess, P. James, and V. Schmidt, 2020: A calibrated combination of probabilistic precipitation forecasts to achieve a seamless transition from nowcasting to very short-range forecasting. *Wea. Forecasting*, **35**, 773–791, <https://doi.org/10.1175/WAF-D-19-0181.1>.
- Scheuerer, M., 2014: Probabilistic quantitative precipitation forecasting using ensemble model output statistics. *Quart. J. Roy. Meteor. Soc.*, **140**, 1086–1096, <https://doi.org/10.1002/qj.2183>.
- Schwartz, C. S., and R. A. Sobash, 2017: Generating probabilistic forecasts from convection-allowing ensembles using neighborhood approaches: A review and recommendations. *Mon. Wea. Rev.*, **145**, 3397–3418, <https://doi.org/10.1175/MWR-D-16-0400.1>.
- Seed, A. W., C. E. Pierce, and K. Norman, 2013: Formulation and evaluation of a scale decomposition-based stochastic precipitation nowcast scheme. *Water Resour. Res.*, **49**, 6624–6641, <https://doi.org/10.1002/wrcr.20536>.
- Shah, C., 2020: *A Hands-On Introduction to Data Science*. Cambridge University Press, 459 pp.
- Stephan, K., S. Klink, and C. Schraff, 2008: Assimilation of radar-derived rain rates into the convective-scale model COSMO-DE at DWD. *Quart. J. Roy. Meteor. Soc.*, **134**, 1315–1326, <https://doi.org/10.1002/qj.269>.
- Sutskever, I., J. Martens, G. Dahl, and G. Hinton, 2013: On the importance of initialization and momentum in deep learning. *Proc. 30th Int. Conf. on Machine Learning*, Atlanta, GA, Proceedings of Machine Learning Research (PMLR), 1139–1147.
- Theis, S., A. Hense, and U. Damrath, 2005: Probabilistic precipitation forecasts from a deterministic model: A pragmatic approach. *Meteor. Appl.*, **12**, 257–268, <https://doi.org/10.1017/S1350482705001763>.
- , C. Gebhardt, and Z. B. Bouallegue, 2012: *Beschreibung des COSMO-DE-EPS und seiner Ausgabe in die Datenbanken des DWD*. Deutscher Wetterdienst, 71 pp., https://www.dwd.de/SharedDocs/downloads/DE/modelldokumentationen/nwv/cosmo_de_eps/cosmo_de_eps_dbbeschr_201208.pdf.
- Vannitsem, S., and Coauthors, 2020: Statistical postprocessing for weather forecasts—Review, challenges and avenues in a big data world. *Bull. Amer. Meteor. Soc.*, **102**, E681–E699, <https://doi.org/10.1175/BAMS-D-19-0308.1>.
- Venugopal, V., E. Foufoula-Georgiou, and V. Sapozhnikov, 1999: Evidence of dynamic scaling in space-time rainfall. *J. Geophys. Res.*, **104**, 31 599–31 610, <https://doi.org/10.1029/1999JD900437>.
- Wang, Y., and Coauthors, 2017: Guidelines for nowcasting techniques. WMO-1198, 82 pp., https://library.wmo.int/doc_num.php?explnum_id=3795.
- Weigl, E., and T. Winterrath, 2010: Radargestützte Niederschlagsanalyse und-vorhersage (RADOLAN, RADVOR-OP). *Promet (Zagreb)*, **35**, 78–86.
- Wilks, D. S., 2006: *Statistical Methods in the Atmospheric Sciences*. 2nd ed. International Geophysics Series, Vol. 100, Academic Press, 648 pp.
- , 2009: Extending logistic regression to provide full-probability-distribution MOS forecasts. *Meteor. Appl.*, **16**, 361–368, <https://doi.org/10.1002/met.134>.
- Winterrath, T., and W. Rosenow, 2007: A new module for the tracking of radar-derived precipitation with model-derived winds. *Adv. Geosci.*, **10**, 77–83, <https://doi.org/10.5194/adgeo-10-77-2007>.
- , —, and E. Weigl, 2012: On the DWD quantitative precipitation analysis and nowcasting system for real-time application in German flood risk management. *Weather Radar and Hydrology*, R. J. Moore, S. J. Cole, and A. J. Illingworth, Eds., IAHS Proceedings and Reports, 323–329.
- Zawadzki, I., J. Morneau, and R. Laprise, 1994: Predictability of precipitation patterns: An operational approach. *J. Appl. Meteor.*, **33**, 1562–1571, [https://doi.org/10.1175/1520-0450\(1994\)033<1562:POPPAO>2.0.CO;2](https://doi.org/10.1175/1520-0450(1994)033<1562:POPPAO>2.0.CO;2).
- Zeiler, M. D., 2012: Adadelta: An adaptive learning rate method. <https://arxiv.org/abs/1212.5701>.
- Zhang, W., K. Itoh, J. Tanida, and Y. Ichioka, 1990: Parallel distributed processing model with local space-invariant interconnections and its optical architecture. *Appl. Opt.*, **29**, 4790–4797, <https://doi.org/10.1364/AO.29.004790>.