



Calculating the turbulent fluxes in the atmospheric surface layer with neural networks

Lukas Hubert Leufen^{1,a} and Gerd Schädler¹

¹Institute of Meteorology and Climate Research – Department Troposphere Research,
Karlsruhe Institute of Technology, Karlsruhe, Germany

^anow at: Forschungszentrum Jülich GmbH, Jülich, Germany

Correspondence: Gerd Schädler (gerd.schaedler@kit.edu)

Received: 24 October 2018 – Discussion started: 21 November 2018

Revised: 21 March 2019 – Accepted: 17 April 2019 – Published: 24 May 2019

Abstract. The turbulent fluxes of momentum, heat and water vapour link the Earth's surface with the atmosphere. Therefore, the correct modelling of the flux interactions between these two systems with very different timescales is vital for climate and weather forecast models. Conventionally, these fluxes are modelled using Monin–Obukhov similarity theory (MOST) with stability functions derived from a small number of field experiments. This results in a range of formulations of these functions and thus also in differences in the flux calculations; furthermore, the underlying equations are non-linear and have to be solved iteratively at each time step of the model. In this study, we tried a different and more flexible approach, namely using an artificial neural network (ANN) to calculate the scaling quantities u_* and θ_* (used to parameterise the fluxes), thereby avoiding function fitting and iteration. The network was trained and validated with multi-year data sets from seven grassland, forest and wetland sites worldwide using the Broyden–Fletcher–Goldfarb–Shanno (BFGS) quasi-Newton backpropagation algorithm and six-fold cross validation. Extensive sensitivity tests showed that an ANN with six input variables and one hidden layer gave results comparable to (and in some cases even slightly better than) the standard method; moreover, this ANN performed considerably better than a multivariate linear regression model. Similar satisfying results were obtained when the ANN routine was implemented in a one-dimensional stand-alone land surface model (LSM), paving the way for implementation in three-dimensional climate models. In the case of the one-dimensional LSM, no CPU time was saved when using the ANN version, as the small time step of the standard version required only one it-

eration in most cases. This may be different in models with longer time steps, e.g. global climate models.

1 Introduction

The turbulent fluxes of momentum, heat, water vapour and trace gases link the atmosphere with the Earth's surface. Therefore, the faithful representation of these fluxes is essential for climate and weather forecast models to function properly. In these models, the fluxes are parameterised as momentum flux $\tau = \rho u_*^2$ and heat flux $H = -\rho c_p u_* \theta_*$ (where ρ is air density and c_p is air heat capacity), using a velocity scale u_* and a (potential) temperature scale θ_* . u_* and θ_* depend on near-surface wind and temperature, their gradients, surface roughness and atmospheric stability. In the framework of the almost exclusively used Monin–Obukhov similarity theory (MOST; Monin and Obukhov, 1954), one has to determine stability functions for momentum and heat which depend on a single stability parameter (for details, see e.g. Arya, 2001). These stability functions must be determined empirically and have been obtained by different authors from regressions on observations from a small number of field experiments. As shown in Högström (1996), the results vary considerably, especially in the very stable and the very unstable regimes, due to a lack of and/or a large scatter of the observations and possibly violations of the assumptions of MOST. Furthermore, the underlying non-linear equations must be solved iteratively at each time step of a model run which can be time consuming.

In the present study, artificial neural networks (ANN) and their ability to simulate a wide range of relationships between input and output variables as a universal approximator (Hornik et al., 1989) are used to model the stability functions. Our goals in this study are (a) to see how well ANNs can approximate the stability relationships, (b) to possibly increase accuracy by using larger data sets, (c) to use the more flexible ANN approach instead of function fitting and (d) to possibly speed up the calculations. With positive outcomes, we ultimately want to replace the relevant subroutines in a climate model with ANNs in order to improve overall model performance.

A good overview of various applications of ANNs in different disciplines can be found in Zhang (2008). Several studies (e.g. Gardner and Dorling, 1999; Elkamel et al., 2001; Kolehmainen et al., 2001) describe applications of ANNs to meteorological and air quality problems. In these studies, long time series of observational data are available for ANN training and only one station is involved in the training and validation process. Comrie (1997) compares ozone forecasts using ANNs with forecasts using standard linear regression models and finds that ANNs are “somewhat, but not overwhelmingly” better than the regression models. The best performance is obtained with an ANN incorporating time lagged data. Gomez-Sanchis et al. (2006) use a multilayer perceptron (MLP) to predict ozone concentrations near Valencia based on meteorological and traffic information. Different model architectures are tested and good agreement with observations is found. However, for different years different model architectures are required for optimal results, which they attribute to the varying relative importance of the input variables. Elkamel et al. (2001) use a one hidden layer ANN and meteorological and precursor concentrations to predict ozone levels in Kuwait. They find that the ANN gives consistently better predictions than both linear and non-linear (log output) multivariate regression models. Kolehmainen et al. (2001) compare the ability of self-organising maps and the MLP to predict NO₂ concentrations when combined with different methods to preprocess the data. They find that direct application of the MLP gives the best results. In all of these studies, just one hidden layer is sufficient, and it is pointed out that careful selection of the input data is crucial.

Some papers deal with the idea of replacing whole models or model components with ANNs. For example, Knutti et al. (2003) teach a neural network to simulate certain output variables of a global climate model and use the result to establish probability density functions as well as to enlarge a global climate model ensemble considerably. Gentile et al. (2018) use an ANN to parameterise the effects of sub-grid-scale convection in a global climate model. The ANN learns the combined effects of turbulence, radiation and cloud microphysics from a convection resolving sub-model. They find that using the ANN, many of these processes can be predicted skilfully, but spatial variability is reduced compared with the original climate model; they attribute this to chaotic

dynamics accounted for in the original model, but not in the version using the ANN, which is deterministic by construction. Sarghini et al. (2003) and Volland et al. (2017) use an ANN trained with direct numerical simulation data as a sub-grid-scale model in a large-eddy simulation model. Sarghini et al. (2003) find that the ANN is able to reproduce the non-linear behaviour of the turbulent flows, whereas Volland et al. (2017) find that the ANN performs well for the flow cases the ANN was trained for, but that it can fail for other flow configurations.

This paper is structured as follows: in Sect. 2, we give a short overview over Monin–Obukhov similarity theory and artificial neural networks, introduce cross-validation, present the data used (including important quality checks) and describe our strategy to find the best network. Thereafter, trained ANNs (which are in fact MLPs, but we will stick to the generic name ANN here) are validated and results are discussed (Sect. 3). In Sect. 4, the ANN that shows the best performance is implemented in a one-dimensional land surface model (LSM), and the results are compared with those of the standard version. A summary is given in Sect. 5.

2 Methods and data

2.1 Monin–Obukhov similarity theory (MOST)

In weather forecast and climate models, the turbulent fluxes of momentum, heat, water vapour and trace gases between the Earth’s (land and water) surface and the atmosphere are usually calculated on the basis of Monin–Obukhov similarity theory (MOST, Monin and Obukhov, 1954). Here, we give a very brief overview of MOST, focussing on momentum and heat fluxes; details can be found in Arya (2001). The main assumptions of MOST are as follows: horizontally homogeneous terrain (in particular, flow characteristics are independent of wind direction), stationarity, fair (i.e. dry) weather conditions and low terrain roughness, i.e. no or low vegetation; for tall vegetation, the last assumption is usually circumvented by introducing the concept of displacement height $d \approx 0.67z$. MOST postulates that turbulence in the surface layer (also called the Prandtl or constant flux layer) only depends on four quantities: the height above the ground z (for tall canopies $z-d$), a velocity scale u_* , a temperature scale θ_* and a buoyancy term g/θ , where g is gravitational acceleration and θ denotes potential temperature. The velocity and temperature scales depend on the respective velocity and temperature gradients as well as on atmospheric stability, and this dependence will be used later to build the neural networks. According to the Buckingham Pi theorem, these four quantities based on length, time and temperature can be combined to a single non-dimensional quantity $\zeta = z/L$, where $L = u_*^2 \theta / (\kappa g \theta_*)$ is the Obukhov length and $\kappa \approx 0.40$ is the von Kármán constant; other dimensionless quantities such as dimensionless wind and temperature gradients can

be expressed as functions of ζ . The Obukhov length L measures the stratification of the surface layer: large (positive or negative) values (i.e. $\zeta \approx \pm 0$) indicate neutral stratification, positive values indicate stable stratification and negative values indicate unstable stratification. As momentum flux is expressed as $\tau = \rho u_*^2$, and heat flux as $H = -\rho c_p u_* \theta_*$ (ρ is air density and c_p is air heat capacity), our goal is to determine u_* and θ_* from known quantities, which are modelled or observed wind and temperature gradients in the surface layer in our case. Non-dimensional wind shear ϕ_m and the non-dimensional gradient of the potential temperature ϕ_h (also called stability functions) can be written as

$$\phi_m(\zeta) = \frac{\kappa z}{u_*} \frac{\partial u}{\partial z}, \quad \phi_h(\zeta) = \frac{\kappa z}{\theta_*} \frac{\partial \theta}{\partial z} \quad (1)$$

respectively, where u is the mean wind speed at height z .

The “universal” functions ϕ_m and ϕ_h can be obtained from simultaneously measured values of the wind speed and temperature gradients and the momentum and heat fluxes (providing u_* and θ_*). Conversely, u_* and θ_* can be calculated from these universal functions, given the wind speed and temperature gradients; this is how these functions are used in weather and climate models. Data from field experiments, notably the Kansas experiment in 1968, have been used to derive these universal functions by Businger et al. (1971). Generally, the stability functions obtained in this manner have the following form:

$$\phi_{m,h}(\zeta) = (\alpha_{m,h} + \beta_{m,h}\zeta)^{\gamma_{m,h}}, \quad (2)$$

with the coefficients depending on $\zeta > 0$ or $\zeta \leq 0$. An overview of these functions can be found in Högström (1988); Högström (1988) shows that there is considerable scatter in the data (especially under very stable and very unstable conditions) and, as a result, also in the derived universal functions.

In applications, differences are known rather than gradients. Integrating the functions (Eq. 1) between a reference height z_r and z yields

$$\begin{aligned} \kappa(u(z) - u(z_r)) / u_* &= \ln(z/z_r) - \Psi_m(z/L), \\ \kappa(\theta(z) - \theta(z_r)) / \theta_* &= \ln(z/z_r) - \Psi_h(z/L), \end{aligned} \quad (3)$$

where

$$\Psi_{m,h}(z/L) = \int_{z_r/L}^{z/L} (1 - \phi_{m,h}(x)) dx/x. \quad (4)$$

For the purpose of climate modelling, i.e. obtaining fluxes from simulated wind and temperature profiles, u_* and θ_* need to be derived from the respective wind and temperature data at two heights using Eq. (1) or Eq. (3). As ζ itself depends on u_* and θ_* , this amounts to solving a system of two non-linear equations; we will call this traditional method the MOST method.

2.2 Neural networks

In this section, we describe only those aspects of neural networks which are relevant to our study; for more information on neural networks, the reader is referred the literature, e.g. Rojas (2013); Kruse et al. (2016). Neural networks, or more precisely artificial neural networks (ANNs), are a widely used technique to solve classification and regression problems as well as to analyse time series (Zhang, 2008). The building blocks of an ANN are the so-called neurons, arranged in different layers. An ANN has at least an input and an output layer; between these layers, there can be so-called hidden layers. The neurons in successive layers (but not within the same layer) are connected via weights (see Fig. 7). A neuron processes input data as follows:

$$o_j = f\left(\sum_i^N o_i \cdot w_{ij}\right), \quad (5)$$

where o_j is the output of the neuron j , N is the number of neurons in the preceding layer (including the bias neuron, see below), o_i is the output of the i th neuron in the preceding layer and w_{ij} is corresponding weight. Non-linear behaviour of the network is induced by using non-linear activation functions f . Each neuron belongs to a unique layer in a directed graph. Here, we use so-called multilayer perceptrons (MLP), also known as feed-forward networks due to the unidirectional information flow. Each MLP consists of an input, an output and at least one hidden layer with an arbitrary number of neurons. The input layer takes (normalised) input data and the output data returns the (also normalised) results of the MLP. Normalisation is essential for equal weighting of the input and for consistency with the domain and range of the activation functions. Input information is propagated from layer to layer while each neuron responds to the signal. Bias neurons are used to adjust the activation level.

All free parameters (i.e. weights) of a MLP need to be determined by a training process. In the case of supervised learning, the MLP knows its deviation from target values at every time and an error can be calculated using this deviation (Zhang, 2008). The aim of the training is to minimise an error metric by adjusting the network’s weights. Here we use the mean squared error (MSE):

$$MSE = \frac{1}{|P|} \sum_{p \in P} \frac{1}{N^\Omega} \sum_j^{N^\Omega} (t_{jp} - o_{jp})^2, \quad (6)$$

where P is the total number of data points, N^Ω is the number of neurons in the output layer, t_{jp} is the target value of data point p and o_{jp} is the output of the MLP for data point p . In the study described here, we use a MLP with hyperbolic tangent activation functions in the hidden layer(s) (here one or two) and linear functions in the output layer trained by the Broyden–Fletcher–Goldfarb–Shanno (BFGS) quasi-Newton backpropagation algorithm (Broyden, 1970; Fletcher, 1970; Goldfarb, 1970; Shanno, 1970).

2.3 Regression model

To compare the performance of the neural networks described in Sect. 2.2 with a standard regression model, we used the multivariate linear regression (MLR) model as implemented in the “mvregress” MATLAB routine:

$$y_j = \sum_i \beta_{ij} \cdot x_i + \epsilon_j, \quad (7)$$

where the β_{ij} are the regression coefficients and the ϵ_j are the residual errors with a multivariate normal distribution. The model uses a multivariate normal maximum likelihood estimation. The resulting values for β_{ij} maximise the log-likelihood function $\log L(\beta, \epsilon | y, x)$. We used the same six-element input vector and two-element target vector as for the ANN (both described in Sect. 2.6), as well as the same training and independent test data sets (from DE-Keh station; see Sect. 2.4 and 2.5).

2.4 Data

To train and validate the neural network, data from 20 meteorological towers in Europe, Brazil and Russia spread over different land use types including forest, grassland and crop fields were collected. All data were measured after the year 2000 and observation periods range from a few months to several years. Figure 1 shows a map of the sites that provided data. Stations varied widely with respect to their environmental surrounding, instrumental set-up and measurement heights. The tower configuration of the sites is shown schematically in Fig. 2. For our purposes, we required temperature and wind speed at two measurement heights as well as the momentum and sensible heat fluxes to calculate the scaling quantities u_* and θ_* (see Sect. 2.6). The fluxes at the sites used were all measured using the eddy covariance method. If this information was not available, density was calculated from the ideal gas equation using virtual temperature when humidity data were available, otherwise the temperature of dry air was utilised. For forests, all observations had to be above the canopy, and all vertical distances were reduced by the displacement height, which was assumed to be two-thirds of the canopy height. The original temporal resolution of the data was either 10 or 30 min; these data were aggregated to 1 h averages.

An important step before using data as input for the ANN was to check if the data were compatible with Monin–Obukhov theory, i.e. if an (at least approximate) functional relationship between ζ and the right-hand sides of Eq. (1) was present and if so, how well they were represented by the universal stability functions in Eq. (1). It was found that no relationship existed for some sites. This may have been due to a violation of the assumptions of the Monin–Obukhov theory, such as inhomogeneous terrain around the site or the dependence of the roughness length on the wind direction. Data from these sites were not used further, except for data

from the DE-Tha site (see Sect. 4). The remaining stations (see Table 1), which comprised about 113 500 hourly averaged data points in total (see Table 2), were used to train and validate the networks. For these stations, agreement was generally better for temperature than for wind; furthermore, agreement was better for unstable than for stable stratification, an observation which is often mentioned in the literature.

Data were preprocessed before they were presented to the ANN. Input and output data were normalised according to their extrema to the interval $[0, 1]$ (see Table 3). Furthermore, weak wind situations with wind speeds below 0.3 m s^{-1} were filtered out. Because of the large scatter of wind and temperature gradients under atmospheric conditions with absolute heat fluxes below 10 W m^{-2} or small scaling wind speeds ($u_* < 0.1 \text{ m s}^{-1}$), such data were excluded. Finally, the signs of the temperature scale θ_* and the potential temperature gradient had to be the same; this meant excluding counter-gradient fluxes which can be observed over forests (Denmead and Bradley, 1985) and ice (Sodemann and Foken, 2005) but violate the assumptions of MOST (Foken, 2017a, b).

2.5 Cross-validation and generalisation

Trained networks were validated using k -fold cross-validation (Kohavi, 1995; Andersen and Martinez, 1999) to prevent overfitting (Domingos, 2012). Overfitting originates from the trade-off between minimising the error for given data and maximising performance for new unknown data (Chicco, 2017). In the first experiment, the full data set is divided into $k = 6$ subsets using a random data split with approximately equal size first. Cyclically, one subset is kept for independent testing, the remaining $k - 1$ subsets are used for training and validation. Using this experiment, we can show that ANNs are able to learn from the data and to represent their characteristics. In the second experiment, we go one step further and check if the ANNs found can handle not only unknown data but also completely new stations that were not previously used, i.e. if they are able to generalise. For this experiment, we decided to validate trained models using the NL-Cab station and to then test the best ANNs on the DE-Keh station, which had been left out in the training and validation phases of this experiment (see stations details in Sect. 2.4). For these two stations, the MOST method performed best; thus, they present a strong challenge for the ANNs with respect to achieving similar quality.

2.6 ANN set-up and the selection of the best ANN

Neural networks are very flexible in terms of the number of layers, the number of nodes, the error metrics, the training method, the activation function and so on; thus, a series of sensitivity runs were performed, which always consisted of a training and a validation phase. To find an optimal network architecture, we varied the following parameters:

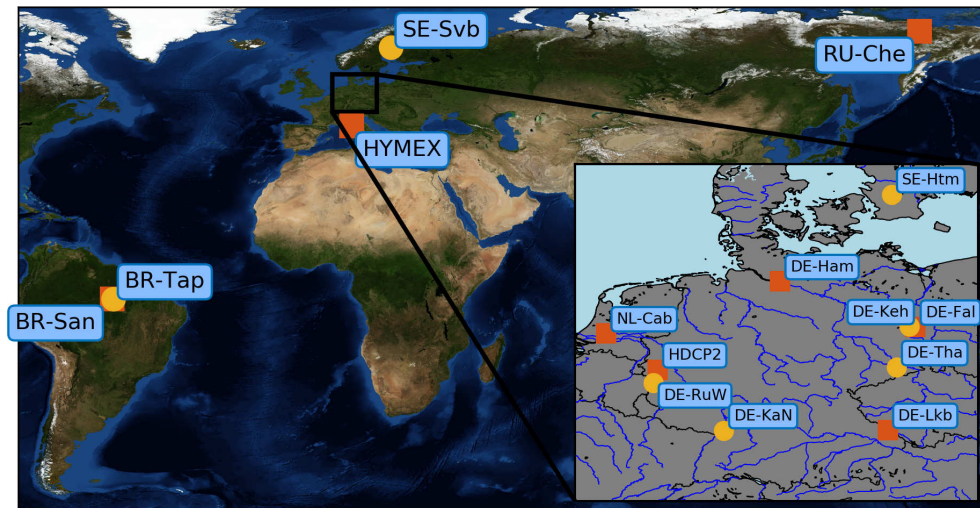


Figure 1. Location of the stations that provided data for this study. Station symbols represent low (red square; grasslands, croplands and wetlands) and tall (yellow circle; forest) vegetation. HDCP2 includes the DE-Nie07, DE-Nie13 and DE-Was06 stations, and HYMEX includes the FR-CorX and FR-GiuX stations. Further information can be found in Table A1.




Figure 2. The schematic set-up of the meteorological towers used for this study. Available measurements for wind velocity (black, left arm) and temperature (black, right arm) are shown as well as the final measurement height that was used for wind (blue), temperature (yellow) and turbulent fluxes (red). Vegetation height is illustrated in green, and towers with a total height above 80 m are clipped. (“Left arm” and “right arm” in this caption refer to horizontal arms that the instruments are mounted to on a real mast.)

- the number and type of input variables,
- the number of hidden layers (one or two) and
- the total number of nodes in the hidden layer(s) (between 1 and 14).

To avoid an excessive number of sensitivity runs, the parameters listed in Table 3 were kept fixed based on recommendations in the literature (Zhang, 2008; Kruse et al.,

2016). Training was carried out in batch mode; therefore, the network’s weights were adjusted after each epoch. Training ended at most after 1000 epochs or if the error on the validation data increased for 50 successive epochs (early stopping). In the latter case, the state of the trained network with the lowest error for the validation data (and not the early stopping state) was set as final state. We tested network architectures with six- and seven-element input vectors. The six-element input vector consisted of the wind speed and potential temperature averages over the two heights, the vertical gradients of wind and potential temperature, and their ratio and a classifier to distinguish between low ($c_{\text{veg}} = 0$) and tall ($c_{\text{veg}} = 1$) vegetation. For the seven-element input vector, we replaced the temperature gradient by its absolute value and added an additional input node describing the sign of the potential temperature gradient. In both cases, the target vector remained a two-element vector consisting of the wind scale u_* and the temperature scale θ_* . As mentioned above, we experimented with ANNs with one and two hidden layers. For the ANNs with one hidden layer, we varied the number of neurons in the hidden layer from one to twice the size of the input layer. For ANNs with two hidden layers, the number of neurons in each layer was increased up to the number of input neurons.

All networks were trained to minimise the overall (sum of u_* and θ_*) MSE on normalised data from Eq. (6). To compare the different ANNs, we used the root-mean-squared error (RMSE) $\text{RMSE} = \sqrt{\text{MSE}}$, the mean absolute error (MAE)

$$\text{MAE} = \frac{1}{|P|} \sum_{p \in P} \frac{1}{N} \sum_{j=1}^{N^{\Omega}} |t_{jp} - y_{jp}| \quad (8)$$

Table 1. Station information for the meteorological towers selected for training and validation (see Sect. 2.4); a list of all stations is given in Table A1. Land usage classification follows the International Geosphere–Biosphere Programme (IGBP) standards: evergreen needle-leaf forests (ENF), grasslands (GRA), permanent wetlands (WET) and croplands (CRO).

Station	Complete station name	Lat (°)	Long (°)	Height (m a.s.l.)	IGBP	Tower height (m)
BR-San	Santarem Pasture Tower Site (Para, Brazil)	−3.02	−54.89	100	GRA/CRO	20
DE-Fal	Grenzschiichtmessfeld Falkenberg	52.17	14.12	73	GRA	10
DE-KaN	KIT CN Messmast	49.09	8.43	110	ENF	200
DE-Keh	Messstation Forst Kehrigk	52.18	13.95	49	ENF	30
NL-Cab	CESAR Observatory	51.97	4.93	−0.7	GRA	213
RU-Che	Cherkssii Tower	68.61	161.34	6	WET	5
SE-Svb	Svartberget ICOS Sweden	64.25	19.77	270	ENF	150

Table 2. Time series information for the meteorological towers selected for training and validation. Count and availability are measured at hourly intervals and not at the original resolution of each time series.

Station	From	To	Availability	Count
BR-San	1 Jan 2001	22 Sep 2005	62 %	25 503
DE-Fal	1 Jan 2008	21 Dec 2009	70 %	12 118
DE-KaN	1 Mar 2015	30 Dec 2016	78 %	12 541
DE-Keh	1 Jan 2008	29 Dec 2009	70 %	12 207
NL-Cab	1 Jan 2014	30 Nov 2017	94 %	32 337
RU-Che	26 May 2014	14 Oct 2016	40 %	8283
SE-Svb	18 Jan 2015	1 Nov 2016	68 %	10 707
Total				113 696

and the Pearson correlation coefficient (r)

$$r = \frac{1}{N^{\Omega}} \sum_{j=1}^{N^{\Omega}} \frac{\sum_p (y_{jp} - \bar{y}_j) (t_{jp} - \bar{t}_j)}{\sqrt{\sum_p (y_{jp} - \bar{y}_j)^2} \cdot \sqrt{\sum_p (t_{jp} - \bar{t}_j)^2}} \quad (9)$$

$$\in [-1, 1],$$

where \bar{y}_j and \bar{t}_j are the averages of the j th net output and the target value with $\bar{y}_j = \frac{1}{|P|} \sum_p y_{jp}$ and $\bar{t}_j = \frac{1}{|P|} \sum_p t_{jp}$.

When ANNs are to be used in climate models, one has to find a trade-off between two aspects: on the one hand, the model should perform well according to the quality metrics described above, on the other hand, a superior model in terms of small errors but with higher computational demands may not be the best choice for use in climate models where reducing computing time is a very high-priority criterion. For ANNs, computing time normally increases with the complexity of a network, i.e. with its size. Therefore, we also tested ANNs with smaller-than-optimal numbers of neurons in view of this trade-off. To find smaller networks that possibly required less computing time, we looked at networks that met the requirement that the size of each hidden layer n_{h_i} was less or equal to the size of the input layer n_I minus 1:

$$n_I - 1 \geq n_{h_1} (\geq n_{h_2}). \quad (10)$$

This condition was found after some experimentation and is somewhat arbitrary, but there is no hard rule defining the simplicity of a model. We will refer to ANNs that satisfy this condition as “simple networks”.

3 Results

As described in Sect. 2.5, ANNs are always trained on the training data set only and validated on a disjoint validation data set. If the MSE on the validation set rises continuously, training is stopped to prevent overfitting (early stopping). Following this training and validation stage, the ability of the selected ANNs to generalise is tested on data that are completely new to the ANNs. All in all, more than 100 000 networks were trained and tested this way.


3.1 Effect of data splitting

The validation results from ANNs with six inputs and one single hidden layer trained under six-fold cross-validation with random data splitting are shown in the box-and-whiskers plot in Fig. 3 as a function of the number of hidden neurons. One can see that the validation MSE decreases with the increasing number of hidden neurons and has already reached an asymptotic value of about 0.008 with six to seven neurons. Furthermore, the scatter of the MSE is quite small, meaning that the quality of the results from ANNs trained on different sets varies only slightly.

If the training data are not split randomly but undergo a station-wise split, a larger MSE and a considerably larger scatter of the MSE results are found. Comparing Fig. 4 with Fig. 3 shows that the MSE roughly doubles, whereas scatter increases by about a factor of 10, almost independent of the network architecture. Conversely, increasing the network size does not necessarily imply a lower MSE. Using two hidden layers slightly reduces the median and error minimum, but also increases the MSE spread. The comparison of Fig. 3 with Fig. 4 also shows that the station-wise error minima are comparable to those obtained from a random data split. In

Table 3. Fixed network parameters for training (following Zhang, 2008; Kruse et al., 2016).

Normalisation	$\tilde{x}_i = (x_i - \min_{\text{Data}}(x_i)) / (\max_{\text{Data}}(x_i) - \min_{\text{Data}}(x_i))$
Activation function	Hyperbolic tangent
Activation function output	Linear
Training algorithm	BFGS quasi-Newton backpropagation
Error metric	MSE
Early stopping after ... epochs	50
Maximum number of epochs	1000
Training mode	Batch

**Figure 3.** Network with six inputs and one hidden layer under six-fold cross-validation: the MSE of the network trained on the validation data set using random data splitting as a function of the hidden layer size. Whiskers indicate the interquartile range, and each box summarises the results from 750 single networks.

both types of validation, ANNs with one and two hidden layers are not significantly different.

All in all, comparing Fig. 3 with Fig. 4 shows that the station-wise data split substantially reduces the ANN performance. This implies that using not enough stations as well as station-wise training impairs the generalisation of learned relationships between inputs and target values. Possible reasons for this may be the tendency of the ANNs to overfit training data by memorising relationships and local effects contaminating the validity of MOST, such as unideal positioning of sites or unideal atmospheric conditions. These findings confirm the need for independent testing with data that is unknown to the ANN in order to estimate the ANNs real ability to generalise. This will be discussed in the next section.

3.2 Generalisation to unknown data

After showing that ANNs are able to extract u_* and θ_* from training data successfully, our next step is to assess how the ANNs found in the previous section can handle input from stations which were not used for training or validation, i.e.

data completely unknown to them; this simulates the situations in which ANNs would be used in climate models (where grid points play the role of stations). To test this, we choose the NL-Cab station for validation and DE-Keh as the unknown station. We selected these two stations because the MOST method performed best for these stations; therefore it is a strong challenge for the ANNs to produce equivalent results. The results of the networks that perform best on the validation set are summarised in Table 4, where we compare the ANNs according to the increasing complexity of their network architecture. For comparison, and in view of reducing CPU time, we also show the results of the best simple networks (as defined in Sect. 2.6) in this table. Table 4 shows that all ANNs perform better than the MOST method on the validation data set (NL-Cab), in terms of the MSE and correlation coefficient (r). Applying these ANNs to the test data set (DE-Keh) results in an increased MSE and a lower correlation coefficient, whereas the MOST method performs better on the test data set. Among the ANNs, the 6–5–3–2 ANN displayed the best test performance with a MSE of 0.68×10^{-2} , but the simpler 6–3–2 ANN was second best (also in terms of the MSE); thus, simple networks can be almost as good as larger networks. Networks with seven inputs have no substantial advantage over networks with six inputs in our research. ANNs with two hidden layers perform slightly better on the test data than ANNs with a single hidden layer. The overall correlation between network outputs and target values is quite high ($r \geq 0.85$) in all cases.

We also carried out a comparison for the turbulent momentum and heat fluxes $\tau = \rho u_*^2$ and $H = -\rho c_p u_* \theta_*$, which are the quantities ultimately needed in climate simulations. Results for the momentum and heat fluxes of three networks that performed well as well as for the MOST method are shown in Figs. 5 and 6 and in Tables 5 and 6, respectively. In the tables we also show the results of the multivariate linear regression (MLR) described in Sect. 2.3. Both ANNs, MLR and the standard method tend to underestimate larger momentum fluxes, but differences among ANNs are quite small. The best agreement is achieved with the 6–5–3–2 ANN, which is almost as good as the standard method.

Regarding the heat flux, the differences between the ANNs are again relatively small, but the ANNs as well as the standard method tend to overestimate the heat fluxes, whereas

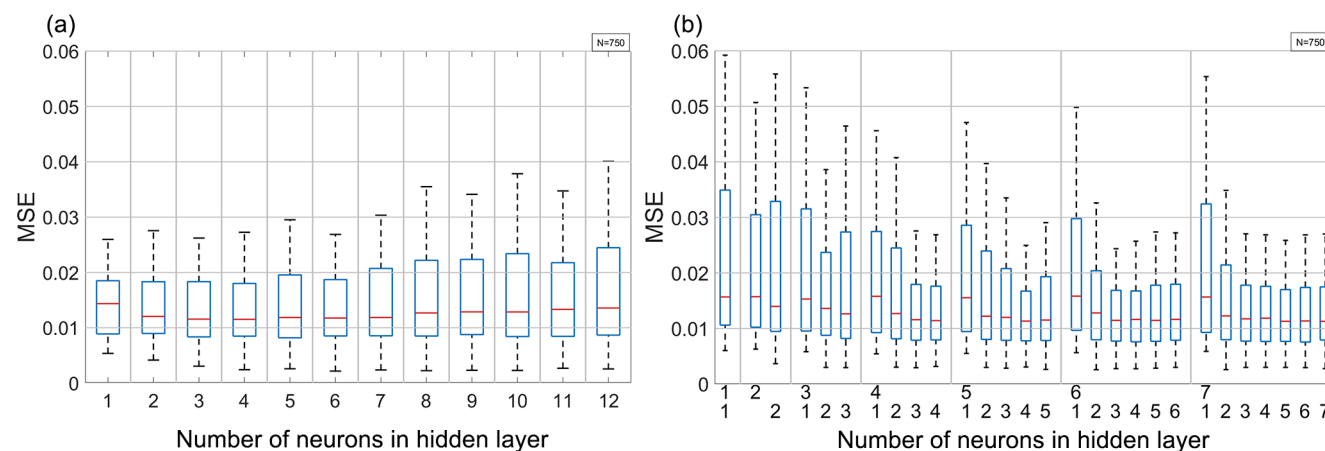


Figure 4. The validation MSE of trained networks using a station-wise data split as a function of the hidden layer size for (a) the network with six inputs and one hidden layer (left), and (b) the network with seven inputs and two hidden layers (right). The numbers on the bottom axis of (b) indicate the number of neurons in the first (top row) and second (bottom row) hidden layer. Values for the other networks considered are similar. Whiskers indicate the length of interquartile range, and each box summarises the results from 750 single networks.

Table 4. Performance results of the overall best network and the best simple network. MSE and r are measured on normalised data and are non-dimensional. MSE_v and r_v are calculated on validation data, and MSE_t and r_t are calculated on test data. The performance of the MOST method (the “Benchmark”) is also shown.

Condition	Net. structure	No. weights	MSE_v (10^{-2})	r_v	MSE_t (10^{-2})	r_t
Overall best net.	6–5–2	47	0.17	0.94	0.90	0.89
	7–11–2	112	0.18	0.92	0.96	0.86
	6–5–3–2	61	0.20	0.93	0.68	0.88
	7–5–2–2	58	0.19	0.92	0.79	0.88
Best simple net.	6–3–2	29	0.38	0.92	0.74	0.87
	7–4–2	42	0.21	0.92	1.36	0.87
	6–3–3–2	41	0.27	0.91	0.84	0.85
	7–4–2–2	48	0.22	0.90	1.01	0.86
Benchmark	–	–	0.92	0.85	0.58	0.92

MLR underestimates them (not shown). The best results are obtained with the 6–3–2 ANN. For heat flux, the 7–5–2–2 ANN behaves quite differently to the other ANNs. It produces two distinct states, one around -30 W m^{-2} and the other from 50 to 200 W m^{-2} ; as a result, r is reduced but the MAE is lowest for this 7–5–2–2 ANN. Thus, the 7–5–2–2 ANN acts more like a dichotomous classifier of stability rather than the continuous regression we are looking for. As for the momentum fluxes, the ANNs shows considerably better performance than the regression model. These results reiterate that smaller networks can be as good as or even better than larger networks.

All ANNs show considerably better performance than the multivariate linear regression model. This is not really surprising, as the scaling quantities to be approximated are non-linear functions of stability (Arya, 2001), meaning that an ANN with a non-linear activation function would be ex-

Table 5. Performance of networks vs. multivariate linear regression (MLR) and the MOST method (“Benchmark”) for momentum flux at the DE-Keh station.

Net. structure	MSE ($10^{-2} \text{ N}^2 \text{ m}^{-4}$)	RMSE (N m^{-2})	MAE (N m^{-2})	r
6–5–3–2	2.11	0.15	0.09	0.90
7–5–2–2	2.44	0.16	0.10	0.89
6–3–2	2.56	0.16	0.09	0.87
MLR	5.81	0.24	0.17	0.89
Benchmark	1.72	0.13	0.08	0.90

pected to perform better than any linear model; as Tables 5 and 6 show, this is the case, even for the small 6–3–2 ANN with one hidden layer.

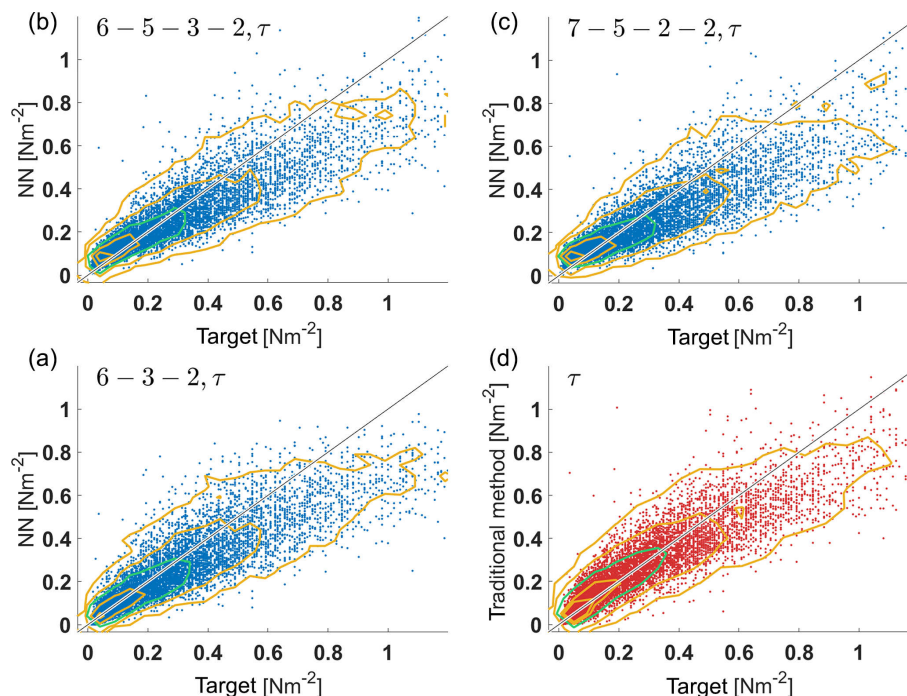


Figure 5. Plots of network output vs. target values for momentum flux on unknown test data (DE-Keh). The 6–3–2 ANN (a), 6–5–3–2 ANN (b), 7–5–2–2 ANN (c) and the MOST method (d) are shown. Contours represent the kernel density estimates of the two-dimensional probability density distribution with the 95th, 75th, 25th and 5th percentiles (yellow contours – starting from the outermost contour) and the 50th percentile (green contour).

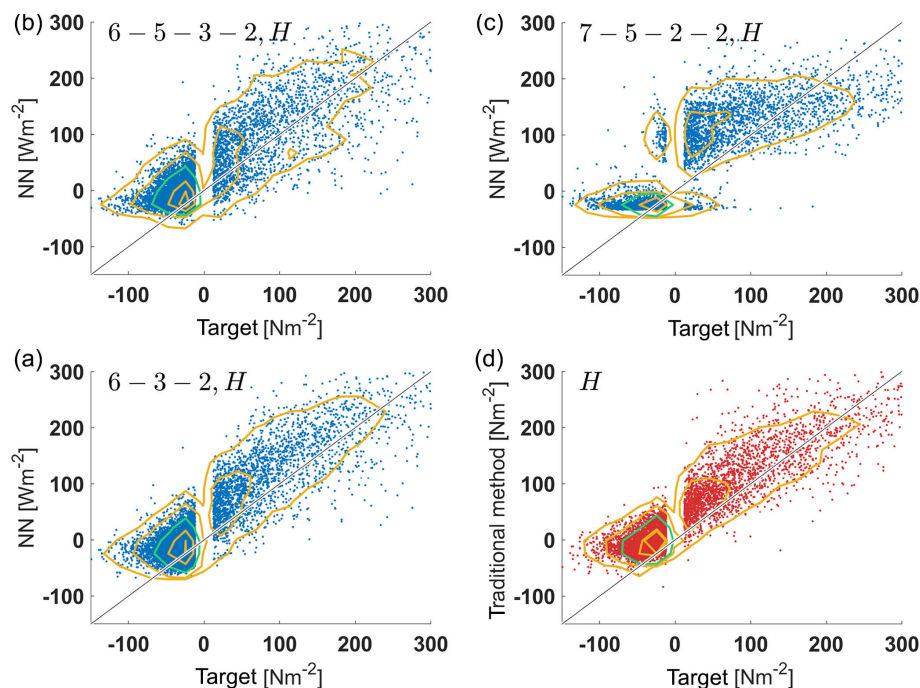


Figure 6. Plots of network output vs. target values for heat flux on unknown test data (DE-Keh). The 6–3–2 ANN (a), 6–5–3–2 ANN (b), 7–5–2–2 ANN (c) and the MOST method (d) are shown. Contours represent the kernel density estimates of the two-dimensional probability density distribution with the 95th, 75th, 25th and 5th percentiles (yellow contours – starting from the outermost contour) and the 50th percentile (green contour). The vertical gap is due to the exclusion of heat fluxes between $\pm 10 \text{ W m}^{-2}$.

Table 6. Performance of networks vs. multivariate linear regression (MLR) and the MOST method (“Benchmark”) for heat flux at the DE-Keh station.

Net. structure	MSE ($\text{W}^2 \text{m}^{-4}$)	RMSE (W m^{-2})	MAE (W m^{-2})	r
6–5–3–2	2461	49.6	37.6	0.85
7–5–2–2	2329	48.3	31.4	0.82
6–3–2	2092	45.8	35.1	0.88
MLR	4447	66.7	53.3	0.65
Benchmark	1915	43.8	34.4	0.90

Table 7. Relative computational demand of the ANNs discussed in the text.

Net. structure	No. of weights	CPU time (relative to 6–3–2 ANN)
6–3–2	29	1
6–5–2	47	1.6
7–11–2	112	3.7
6–5–3–2	61	2.5
7–5–2–2	58	2.4

A comparison of the CPU time required by the different ANNs relative to 6–3–2 ANN is shown in Table 7. The table shows that the increase in computational demand is approximately proportional to the number of weights (as could be expected), and therefore increases considerably when two-layer networks are used. As the discussion above shows, these costs are not reflected in a substantially higher quality of results.

We can conclude that generalisation entails a reduced performance of the ANNs with quite small differences between the various ANNs. The performance of the ANNs is comparable to the MOST method, and the simplest 6–3–2 network has the best score in terms of accuracy and computational efficiency.

4 Implementation of an ANN in a land surface model

As already mentioned, our goal is to replace the MOST method for calculating fluxes with an ANN in the land surface component of climate models; in doing so, we expect more flexibility, accuracy and to possibly save CPU time. The results presented in the previous section indicate that from an accuracy and computational efficiency point of view, the 6–3–2 ANN seems to be most suitable for implementation into a land surface model (LSM). This ANN is shown in Fig. 7.

We implemented the 6–3–2 ANN with weights as obtained in the previous sections in a stand-alone version of the one-dimensional LSM Veg3d (Braun and Schädler, 2005); this




Figure 7. The architecture of the 6–3–2 ANN implemented in the land surface model. Input is described in Sect. 2.6. Purple circles are bias neurons.

replaced the routine using the MOST method to calculate the scaling quantities u_* and θ_* . Here, we refer to the LSM version with the original MOST version as the reference version. Input data for the ANN and data normalisation were the same as described in Sect. 2.4 and output was analogously de-normalised. As the LSM requires the moisture flux in addition to the momentum and heat fluxes, we calculated the scaling specific humidity q_* as proportional to θ_* following the standard procedure used in boundary layer meteorology (Arya, 2001). Meteorological input for the LSM consisted of 30 min values of short- and longwave radiation, wind speed, temperature, specific humidity and air pressure at two heights; soil type and land use were also additionally prescribed. In the present study, the meteorological data were only available for the DE-Fal station for the year 2011 and for the DE-Tha site for the year 1998. For comparison with observations, time series of heat and moisture fluxes as well as soil temperature and soil moisture in the upper soil layers were available, so that the effect of the ANN on the soil component could also be assessed. We performed the comparison with data from the DE-Fal (grassland, year 2011) and DE-Tha (evergreen needle-leaf forest, year 1998) stations for years which had not yet been used for training or validation; thus, the data were new to the ANN in the sense that time periods were used which had not been previously used for training and validation. The DE-Tha station had not been used at all up until this point, as the other sites selected in Sect. 2.4 were more consistent with MOST than DE-Tha and because the DE-Tha time series covered only 1 year. We compared the RMSE and the correlation coefficient of the calculated values with those observed for the reference version and the ANN version. Additionally, we compared the required CPU times. The results of the comparison are shown in Tables 8 and 9.

Especially for grassland, the results of the reference version are very good in terms of RMSE and correlation coefficient.

Table 8. Comparison of the reference version with the ANN version of Veg3d for the DE-Fal grassland station. H denotes the heat flux, M is moisture flux, T_s is soil temperature and w_s is soil moisture.

	Reference	ANN
CPU time	10.83	10.65
RMSE H (W m^{-2})	16.8	27.3
rH	0.87	0.81
RMSE M (W m^{-2})	15.1	20.5
rM	0.91	0.86
RMSE T_s ($^{\circ}\text{C}$)	0.8	1.3
rT_s	0.99	0.99
RMSE w_s (%)	4.8	5.5
rw_s	0.87	0.89

Table 9. Same as Table 8, but for the forest station DE-Tha.

	Reference	ANN
CPU time	95.47	97.74
RMSE H (W m^{-2})	39.0	40.9
rH	0.52	0.57
RMSE M (W m^{-2})	27.9	33.1
rM	0.78	0.71
RMSE T_s ($^{\circ}\text{C}$)	2.4	2.2
rT_s	0.98	0.98
RMSE w_s (%)	5.3	3.7
rw_s	0.53	0.75

cients, and it is difficult for the ANN version to outperform this. However, the results show that the ANN version is able to produce results of a similar quality to the reference version for the fluxes as well as for soil temperature and soil moisture. For tall vegetation, RMSEs are larger and the correlation coefficients are lower; but the differences between the ANN version and the reference version are even smaller than for grassland, and the ANN version even outperforms the reference version for soil moisture. In terms of fluxes, the reference version is generally slightly better. Regarding CPU time, there are only minor differences, although we expected the ANN version to be faster. However, due to the small prognostic time step used, once initialised, the reference version does not need to do more than one iteration to find a solution to the non-linear equation and to update the scaling quantities in most cases; hence, the expensive iteration is reduced considerably. In summary, as a result of this first comparison, it can be concluded that the ANN version works as well as the reference version.

5 Summary

We used an ANN (more precisely, a MLP) to obtain the scaling quantities u_* and θ_* as defined in MOST; these parameters are used in weather and climate models to calculate the

turbulent fluxes of heat and momentum in the atmospheric surface layer. To train, validate and test the neural network, a large set of worldwide observations was used, which represented tall (forests) and low vegetation (grassland and agricultural terrain). A quality assessment of the data sets showed that not all of them were compatible with MOST, so only 7 of the 20 initial data sets could be used.

Sensitivity studies were performed with different sets of input parameters, data sampling methods and network architectures; validation was undertaken using 6-fold cross validation. An important part of the overall network validation was to assess the ability of the network to generalise, i.e. to produce acceptable output if input were data from stations that were completely unknown to the network. These studies showed that even a relatively small 6–3–2 network with six input parameters and one hidden layer yields satisfying results in terms of the RMSE and correlation coefficient. With respect to the trade-off between the quality of results and the computational efficiency, this network performed best.

We could show that the results of the ANN were equivalent to the standard method in all of the tests we performed. A final validation with the heat and momentum fluxes instead of the scaling quantities showed that the MOST method and the ANN approach were also almost equal in terms of quality in this case, with the 6–3–2 ANN performing best. Furthermore, we could show that the ANNs outperform a multivariate linear regression model with the same input and output variables and training and test data. This could be expected, as the stability functions are non-linear functions; therefore, even a small ANN with one hidden layer and a non-linear activation function could be expected to perform better than any linear model. An implementation of the 6–3–2 ANN into an existing LSM showed that the ANN version gives results equivalent to the standard implementation, sometimes even with even higher correlations. However, no decrease in the required CPU time was found.

In summary, it could be shown that even at this stage, an ANN gives results comparable in quality to the MOST method. Some obvious improvements will include more and better differentiated land use classes (e.g. water and urban areas) and more situations of strong stratification. Next steps will include more experiments with the input parameters (e.g. including a time lag) and some fine tuning to improve the computational efficiency (e.g. using different activation functions). We intend to implement and test the neural network routine in a three-dimensional regional climate model (RCM). We expect to save about 5 % of the CPU time, taking parallelisation into account. This may not seem like much, but RCMs in particular are very expensive to run (climatologically relevant multidecadal simulations at high resolution can take several tens of weeks on a high performance system), so every saving counts. The implementation will require the ANN to learn some additional land use types, such as urban areas or water surfaces. If these tests are positive, this would pave the way for replacing other “uncertain”

components of climate models (e.g. cloud microphysics, sea ice) with neural network subroutines, similar to the work described in Sarghini et al. (2003) and Volland et al. (2017), which would increase flexibility and save CPU time. The main hindrance to this undertaking is the current lack of suitable training and validation data. An alternative to “real” data may be the use of data from more detailed models such as LES or urban climate models.

Code availability. A MATLAB script (run.m) that runs the 6–3–2 network with a sample data set (DE-KaN.dat) can be found at <http://doi.org/10.23728/b2share.36ef510c515c4a00bb963113647e44a9>.

Data availability. The data for this study were obtained from the sources mentioned in the acknowledgements.

Appendix A

Table A1. Station information for all of the meteorological towers utilised in the study. Land use classification follows the International Geosphere–Biosphere Programme (IGBP) standards: evergreen needle-leaf forests (ENF), evergreen broadleaf forests (EBF), grasslands (GRA), permanent wetlands (WET) and croplands (CRO).

Station	Complete station name	Lat (°)	Long (°)	Height (m a.s.l.)	IGBP	Tower height (m)
BR-San	Santarem Pasture Tower Site (Para, Brazil)	−3.02	−54.89	100	GRA/CRO	20
BR-Tap	Tapajos National Forest (Santarem, Para, Brazil)	−3.01	−54.58	100	EBF	67
DE-Fal	Grenzsichtmessfeld Falkenberg	52.17	14.12	73	GRA	10
DE-Ham	Wettermast Hamburg	53.52	10.10	0.3	GRA	300
DE-KaN	KIT CN Messmast	49.09	8.43	110	ENF	200
DE-Keh	Messstation Forst Kehrigk	52.18	13.95	49	ENF	30
DE-Lkb	Lackenberg Messstation	49.10	13.30	1300	GRA	9
DE-Nie07	HDGP2 Flux Station 07 Hambach Niederzier	50.90	6.46	110	GRA	5
DE-Nie13	HDGP2 Tower 13 Hambach Niederzier	50.90	6.46	110	GRA	30
DE-RuW	Wüstebach	50.50	6.33	621	ENF	38
DE-Tha	Anchor Station Tharandt	50.96	13.57	380	ENF	42
DE-Was06	HDGP2 Flux Station 06 Wasserwerk	50.89	6.43	96	CRO	5
FR-Cor02	HYMEX Flux Station 02 Corte	43.30	9.17	369	GRA	5
FR-Cor13	HYMEX Tower 13 Corte	43.30	9.17	369	GRA	20
FR-Giu04	HYMEX Flux Station 04 San-Giuliano	42.27	9.52	39	GRA	5
FR-Giu07	HYMEX Flux Station 07 San-Giuliano	42.27	9.52	39	GRA	5
NL-Cab	CESAR Observatory	51.97	4.93	−0.7	GRA	213
RU-Che	Cherkii Tower	68.61	161.34	6	WET	5
SE-Htm	Hyltemossa ICOS Sweden	56.10	13.42	115	ENF	150
SE-Svb	Svartberget ICOS Sweden	64.25	19.77	270	ENF	150

Author contributions. LL was responsible for data collection, quality checks and data preprocessing. LL also trained, validated and generalised the ANNs, performed the MLR analysis and made the comparisons with the MOST method. GS implemented the 6–3–2 network in a land surface model and carried out performance measurements in terms of result quality and CPU time. GS prepared the paper with contributions from LL.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. The authors would like to thank following people and institutions for providing station data: Martin Kohler and Rainer Steinbrecher (Karlsruhe Institute for Technology), Mathias Göckede, Olaf Kolbe and Fanny Kittler (Max Planck Institute for Biogeochemistry Jena), Frank Beyrich (German Weather Service), Ingo Lange (University of Hamburg), Clemens Drüe (University of Trier), Marius Schmidt (Forschungszentrum Jülich) and Thomas Grünwald (TU Dresden). In addition, data from the following sources were collected and used: the Integrated Carbon Observation System Sweden (ICOS), the Cabauw Experimental Site for Atmospheric Research (CESAR) database, the Oak Ridge National Laboratory Distributed Active Archive Center (ORNL DAAC) and the University Corporation for Atmospheric Research (UCAR). We thank the two anonymous reviewers and the editor for their helpful comments and suggestions. Finally, we acknowledge support from the KIT-Publication Fund of the Karlsruhe Institute of Technology.

Financial support. The article processing charges for this open-access publication were covered by a Research Centre of the Helmholtz Association.

Review statement. This paper was edited by Chiel van Heerwaarden and reviewed by two anonymous referees.

References

- Andersen, T. and Martinez, T.: Cross validation and MLP architecture selection, in: IJCNN'99. International Joint Conference on Neural Networks. Proceedings, Washington, DC, USA, 10–16 July 1999, IEEE, 3, 1614–1619, 1999.
- Arya, P. S.: Introduction to micrometeorology, in: International Geophysics Series, San Diego, Calif., Academic Press, vol. 79, 2001.
- Braun, F. and Schädler, G.: Comparison of Soil Hydraulic Parameterizations for Mesoscale Meteorological Models., *J. Appl. Meteorol.*, 44, 1116–1132, 2005.
- Broyden, C. G.: The Convergence of a Class of Double-rank Minimization Algorithms 1. General Considerations, *IMA J. Appl. Math.*, 6, 76–90, <https://doi.org/10.1093/imamat/6.1.76>, 1970.
- Businger, J. A., Wyngaard, J. C., Izumi, Y., and Bradley, E. F.: Flux-Profile Relationships in the Atmospheric Surface Layer, *J. Atmos. Sci.*, 28, 181–189, [https://doi.org/10.1175/1520-0469\(1971\)028<0181:FPRITA>2.0.CO;2](https://doi.org/10.1175/1520-0469(1971)028<0181:FPRITA>2.0.CO;2), 1971.
- Chicco, D.: Ten quick tips for machine learning in computational biology, *BioData Min.*, 10, 35, <https://doi.org/10.1186/s13040-017-0155-3>, 2017.
- Comrie, A. C.: Comparing neural networks and regression models for ozone forecasting, *JAPCA J. Air Waste Ma.*, 47, 653–663, 1997.
- Denmead, O. T. and Bradley, E. F.: Flux-Gradient Relationships in a Forest Canopy, in: *The Forest-Atmosphere Interaction: Proceedings of the Forest Environmental Measurements Conference held at Oak Ridge, Tennessee, 23–28 October 1983*, edited by: Hutchison, B. A. and Hicks, B. B., Springer Netherlands, Dordrecht, 421–442, 1985.
- Domingos, P.: A few useful things to know about machine learning, *Commun. ACM*, 55, 78–87, 2012.
- Elkamel, A., Abdul-Wahab, S., Bouhamra, W., and Alper, E.: Measurement and prediction of ozone levels around a heavily industrialized area: a neural network approach, *Adv. Environ. Res.*, 5, 47–59, 2001.
- Fletcher, R.: A new approach to variable metric algorithms, *Comput. J.*, 13, 317–322, 1970.
- Foken, T.: *Micrometeorology*, SpringerLink: Bücher, Springer, Berlin, Heidelberg, 2nd edn., 2017a.
- Foken, T.: *Energy and Matter Fluxes of a Spruce Forest Ecosystem*, vol. 229, Springer, International Publishing, 2017b.
- Gardner, M. and Dorling, S.: Neural network modelling and prediction of hourly NO_x and NO₂ concentrations in urban air in London, *Atmos. Environ.*, 33, 709–719, 1999.
- Gentine, P., Pritchard, M., Rasp, S., Reinaudi, G., and Yacalis, G.: Could machine learning break the convection parameterization deadlock?, *Geophys. Res. Lett.*, 45, 5742–5751, <https://doi.org/10.1029/2018GL078202>, 2018.
- Goldfarb, D.: A family of variable-metric methods derived by variational means, *Math. Comput.*, 24, 23–26, 1970.
- Gomez-Sanchis, J., Martín-Guerrero, J. D., Soria-Olivas, E., Vila-Francés, J., Carrasco, J. L., and del Valle-Tascón, S.: Neural networks for analysing the relevance of input variables in the prediction of tropospheric ozone concentration, *Atmos. Environ.*, 40, 6173–6180, 2006.
- Högström, U.: Non-dimensional wind and temperature profiles in the atmospheric surface layer: A re-evaluation, *Bound.-Lay. Meteorol.*, 42, 55–78, <https://doi.org/10.1007/BF00119875>, 1988.
- Högström, U.: Review of some basic characteristics of the atmospheric surface layer, *Bound.-Lay. Meteorol.*, 78, 215–246, <https://doi.org/10.1007/BF00120937>, 1996.
- Hornik, K., Stinchcombe, M., and White, H.: Multilayer feedforward networks are universal approximators, *Neural Networks*, 2, 359–366, 1989.
- Knutti, R., Stocker, T., Joos, F., and Plattner, G.-K.: Probabilistic climate change projections using neural networks, *Clim. Dynam.*, 21, 257–272, 2003.
- Kohavi, R.: A study of cross-validation and bootstrap for accuracy estimation and model selection, in: *IJCAI'95 Proceedings of the 14th international joint conference on Artificial intelligence*, Montreal, Quebec, Canada, 20–25 August 1995, 2, 1137–1145, 1995.

- Kolehmainen, M., Martikainen, H., and Ruuskanen, J.: Neural networks and periodic components used in air quality forecasting, *Atmos. Environ.*, 35, 815–825, 2001.
- Kruse, R., Borgelt, C., Braune, C., Mostaghim, S., and Steinbrecher, M.: *Computational intelligence: a methodological introduction*, Springer, London, 2016.
- Monin, A. S. and Obukhov, A. M.: Basic laws of turbulent mixing in the surface layer of the atmosphere, *Tr. Akad. Nauk SSSR Geophiz. Inst.*, 24, 163–187, 1954.
- Rojas, R.: *Neural networks: a systematic introduction*, Springer Science & Business Media, Berlin, 2013.
- Sarghini, F., de Felice, G., and Santini, S.: Neural networks based subgrid scale modeling in large eddy simulations, *Comput. Fluids*, 32, 97–108, 2003.
- Shanno, D. F.: Conditioning of quasi-Newton methods for function minimization, *Math. Comput.*, 24, 647–656, 1970.
- Sodemann, H. and Foken, T.: Special characteristics of the temperature structure near the surface, *Theor. Appl. Climatol.*, 80, 81–89, 2005.
- Vollant, A., Balarac, G., and Corre, C.: Subgrid-scale scalar flux modelling based on optimal estimation theory and machine-learning procedures, *J. Turbul.*, 18, 854–878, <https://doi.org/10.1080/14685248.2017.1334907>, 2017.
- Zhang, G. P.: Neural Networks For Data Mining, in: *Soft Computing for Knowledge Discovery and Data Mining*, edited by: Maimon, O. and Rokach, L., Springer US, Boston, MA, chap. 21, 17–44, https://doi.org/10.1007/978-0-387-69935-6_2, 2008.