



Machine learning methods for stream water temperature prediction

Moritz Feigl^{1,*}, Katharina Lebedzinski^{1,*}, Mathew Herrnegger¹, and Karsten Schulz¹

¹Institute for Hydrology and Water Management, University of Natural Resources and Life Sciences, Vienna, Austria

*These authors contributed equally to this work.

Correspondence: Moritz Feigl (moritz.feigl@boku.ac.at)

Abstract. Water temperature in rivers is a crucial environmental factor with the ability to alter hydro-ecological as well as socio-economic conditions within a catchment. The development of modelling concepts for predicting river water temperature is and will be essential for an effective integrated water management and the development of adaptation strategies to future global changes (e.g. climate change). This study tests the performance of 6 different machine learning models: step-wise linear regression, Random forest, eXtreme Gradient Boosting (XGBoost), Feedforward neural networks (FNN), and two types of Recurrent neural networks (RNN). All models are applied using different data inputs for daily water temperature prediction in 10 Austrian catchments ranging from 200 km² to 96000 km² and exhibiting a wide range of physiographic characteristics. The evaluated input data sets include combinations of daily means of air temperature, runoff, precipitation and global radiation. Bayesian optimization is applied to optimize the hyperparameters of all applied machine learning models. To make the results comparable to previous studies, two widely used benchmark models are applied additionally: linear regression and air2stream.

With a mean root mean squared error (RMSE) of 0.55 °C the tested models could significantly improve water temperature prediction compared to linear regression (1.55 °C) and air2stream (0.98 °C). In general, the results show a very similar performance of the tested machine learning models, with a median RMSE difference of 0.08 °C between the models. From the 6 tested machine learning models both FNNs and XGBoost performed best in 4 of the 10 catchments. RNNs are the best performing models in the largest catchment, indicating that RNNs are mainly performing well when processes with long-term dependencies are important. Furthermore, a wide range of performance was observed for different hyperparameter sets for the tested models, showing the importance of hyperparameter optimization. Especially the FNN model results showed an extremely large RMSE standard deviation of 1.60 °C due to the chosen hyperparameters.

This study evaluates different sets of input variables, machine learning models and training characteristics for daily stream water temperature prediction, acting as a basis for future development of regional multi-catchment water temperature prediction models. All preprocessing steps and models are implemented into the open source R package `wateRtemp`, to provide easy access to these modelling approaches and facilitate further research.

1 Introduction

Water temperature in rivers should not be considered only as a physical property, since it is a crucial environmental factor and a substantial key element for water quality and aquatic habitats. In particular, it influences riverine species by governing e.g. metabolism (Álvarez and Nicieza, 2005), distribution (Boisneau et al., 2008), abundance (Wenger et al., 2011), community



composition (Dallas, 2008) and growth (Imholt et al., 2010), thus aquatic organisms have a specific range of river temperature they are able to tolerate (Caissie, 2006). Due to the impact of water temperature on chemical processes (Hannah et al., 2008) and other physical properties such as density, vapour pressure and viscosity (Stevens et al., 1975), stream temperature indirectly influences key ecosystem processes such as primary production, decomposition and nutrient cycling within rivers (Friberg et al., 2009). These parameters and processes affect the level of dissolved oxygen (Sand-Jensen and Pedersen, 2005) and, of course, have a major influence on water quality (Beaufort et al., 2016).

Besides its ecological importance, river temperature is also of socio-economic interest for electric power and industry (cooling), drinking water production (hygiene, bacterial pollution) and fisheries (fish growth, survival and demographic characteristics) (Hannah and Garner, 2015). Hence, a changing river temperature can strongly alter the hydro-ecological and socio-economic conditions within the river and its neighboring region. Assessing alterations of this sensitive variable and its drivers is essential for the management of impacts and enabling prevention measurements.

Direct temperature measurements are often scarce and rarely available. For a successful integrated water management, it will be essential to derive how river temperature will be developing in the future, in particular when considering relevant global change processes (e.g. climate change), but also on shorter time scales. The forecast, for example, of river temperature with a lead time of a few days can substantially improve or even allow the operation of thermal power plants. Two aspects are important: The efficiency of cooling depends on the actual water temperature. On the other hand, legal constraints regarding maximum allowed river temperatures due do ecological reasons can be exceeded, when warmed up water is directed into the river after the power plant. This is especially relevant during low flow conditions in hot summers. The knowledge of the to be expected water temperature in the next few days is therefore of advantage. An important step in this context is the development of appropriate modelling concepts to predict river water temperature, to describe thermal regimes and to investigate the thermal development of a river. In Austria there are 210 river water temperature measurement stations available, sometimes with 30+ years of data. This large number of available date in Austria is highly advantageous for developing new modelling concepts. Additionally, a wide range of catchments with different physiographic properties are available, ranging from high-alpine, glacier dominated catchments to low land rivers, with meandering characteristics.

In the past, various models were developed to investigate thermal heterogeneity at different temporal and spatial scales, the nature of past availability and likely future trends (Laizé et al., 2014; Webb et al., 2008). In general, water temperature in rivers is modelled by process-based models, statistical/machine learning models or a combination of both approaches. Process-based models are representing physical processes controlling river temperature. According to Dugdale et al. (2017), these models are based on two key steps: First, calculating energy fluxes to or from the river and then determining the temperature change in a second step. Calculating the energy fluxes means to solve the energy balance equation for a river reach by considering the heat fluxes at the air-water and riverbed-water interface (Beaufort et al., 2016). These demanding energy budget components are derived either by field measurements or by approximations (Caissie and Luce, 2017; Dugdale et al., 2017; Webb and Zhang, 1997), highlighting the complexity and parametrisation of this kind of models. Although it is not feasible to monitor these components over long periods or at all points along a river network and contributing catchments (Johnson et al., 2014), they provide clear benefits: (i) give insights into the drivers of river water temperature, (ii) inform about metrics, which can be used



in larger statistical models and (iii) different impact scenarios (Dugdale et al., 2017). These arguments are also the reasons why data-intensive process-based models are widely used despite of their high complexity.

Statistical and machine learning models are grouped into parametric approaches, including regression (e.g. Mohseni and Stefan, 1999) and stochastic (e.g. Ahmadi-Nedushan et al., 2007) models, and non-parametric approaches based on computational algorithms like neural networks or k-nearest-neighbours (Benyahya et al., 2007). In contrast to process-based models, statistical models cannot inform about energy transfer mechanisms within a river (Dugdale et al., 2017). Another main concern is that parametric statistical models showed higher prediction performances on weekly, monthly or seasonal time scales in the past (Caissie, 2006) leading to a loss of temporal variation (DeWeber and Wagner, 2014). However, unlike process-based models, they do not require a large number of input variables, which are unavailable in many cases. Non-parametric statistical models have gained attention in the past few years. Especially machine learning techniques are proved to be a useful tool in river temperature modelling already (Zhu and Piotrowski, 2020).

For this study we chose a set of state of the art machine learning models that showed promising results for water temperature prediction or in similar time series prediction tasks. The 6 chosen models are: step-wise linear regression, Random Forest, extreme Gradient Boosting (XGBoost), Feedforward neural networks (FNN) and two types of recurrent neural networks (RNN). Step-wise linear regression models combine an iterative variable selection procedure with linear regression models. To our knowledge only one previous study by Neumann et al. (2003) already applied this method for predicting daily maximum river water temperature. The Random Forest model (RF) (Breiman, 2001) is an ensemble learning model that averages the results of multiple regression trees. To the authors knowledge, RF has not been applied for river water temperature prediction yet. Zhu et al. (2019d) used bootstrap aggregated decision trees, which are similar, but do not include the random variable sampling for splitting the tree nodes, which is an important characteristics of the RF model. XGBoost (Chen and Guestrin, 2016) is also a regression tree based ensemble learning model. However, instead of averaging multiple trees, XGBoost builds complementary trees for prediction. To the authors knowledge, XGBoost has not been applied for river water temperature predictions yet. However, results from short term water quality parameters predictions, which are also including water temperature, show promising performances (Lu and Ma, 2020; Joslyn, 2018).

Feedforward neural networks (FNN) (White and Rosenblatt, 1963) are the first and simplest type of neural networks. FNNs have already been applied in numerous stream water temperature prediction studies (e.g. Risley et al., 2003; Bélanger et al., 2005; Chenard and Caissie, 2008; Sahoo et al., 2009; Wehrly et al., 2009; McKenna et al., 2010; Westenbroek et al., 2010; Hadzima-Nyarko et al., 2014; DeWeber and Wagner, 2014; Piotrowski et al., 2015; Rabi et al., 2015; Abba et al., 2017; Zhu et al., 2018; Temizyurek and Dadaser-Celik, 2018; Zhu et al., 2019b, d; Graf et al., 2019; Zhu et al., 2019a). In contrast to FNN, Recurrent Neural Networks (RNN) are networks developed specifically to process sequences of inputs. This is achieved by having internal (hidden) states. While there are many different types of RNNs, we focused on the two most widely known, the Long short-term memory (LSTM) (Hochreiter and Schmidhuber, 1997) and the Gated recurrent unit (GRU) (Cho et al., 2014). To the authors knowledge, RNNs have been used in one study by Stajkowski et al. (2020) in which a LSTM in combination with a genetic algorithm hyperparameter optimization was used to forecast hourly urban river water temperature. However, LSTMs



have recently been applied in a wide range of hydrological studies and showed promising results for time-series prediction tasks (e.g. Kratzert et al., 2018, 2019; Xiang et al., 2020; Li et al., 2020).

To make findings comparable with other studies investigating this approach, we apply two benchmark models as baseline: linear regression and `air2stream` (Toffolon and Piccolroaz, 2015). Linear regression models are widely used for river water temperature studies. While earlier studies used mainly air temperature as a regressor to predict river water temperature (e.g., 100 Smith, 1981; Crisp and Howson, 1982; Mackey and Berrie, 1991; Stefan and Preud'homme, 1993), more recent publications are using a wider range of input variables or some modification to the standard linear regression model (e.g., Caldwell et al., 2013; Li et al., 2014; Segura et al., 2015; Arismendi et al., 2014; Naresh and Rehana, 2017; Jackson et al., 2018; Trinh et al., 2019; Piotrowski and Napiorkowski, 2019). `Air2stream` is a hybrid model for predicting river water temperature, which combines 105 a physical based structure with a stochastic parameter calibration. It was already applied in multiple studies over a range of catchments and generally had an improved performance compared to linear regression and other machine learning models (e.g., Piccolroaz et al., 2016; Yang and Peterson, 2017; Piotrowski and Napiorkowski, 2018; Zhu et al., 2019d; Piotrowski and Napiorkowski, 2019; Tavares et al., 2020).

Most studies mainly use air temperature and discharge as inputs for water temperature prediction (e.g. Piccolroaz et al., 2016; 110 Naresh and Rehana, 2017; Sohrabi et al., 2017), while others use additional information from precipitation (e.g. Caldwell et al., 2013) and/or solar radiation (e.g. Sahoo et al., 2009). Additionally, air temperature can either be included as mean, maximum or minimum daily temperature (e.g. Piotrowski et al., 2015). To further investigate which meteorological and hydrological inputs are important and necessary for water temperature prediction, we here use multiple sets of input data and compare their outcome. Especially knowing how simple models with few data inputs perform in comparison with more complex input 115 combinations can give insight on how to plan applications of water temperature modelling for a range of purposes.

Machine learning models are generally parameterized by a set of hyperparameters that have to be chosen by the user to maximize performance of the model. Depending on the model, hyperparameters can have a large impact on model performance (Claesen and De Moor, 2015), but are still most often chosen by rules-of-thumb (Hinton et al., 2012; Hsu et al., 2003) or by testing sets of hyperparameters on a predefined grid (Pedregosa et al., 2011). In this study we apply a hyperparameter 120 optimization using the Bayesian optimization method (Kushner, 1964; Zhilinskis, 1975; Močkus, 1975; Močkus et al., 1978; Močkus, 1989) to minimize the possibility of using unsuitable hyperparameters for the applied models and to investigate the spread in performance depending on the chosen hyperparameters.

This publication presents a thorough investigation of models, input data and model training characteristics for daily stream water temperature prediction. It consists of the application of 6 types of machine learning models on a range of different 125 catchments using multiple sets of data inputs. This includes novel approaches on input data preparation and hyperparameter optimization. The resulting performance of all models is compared to two widely applied benchmark models to make the presented results comparable. Finally, all methods and models are incorporated into an open source R library to make these approaches available for researchers and industries.



2 Methods

130 2.1 Study sites and data

For this study, 10 catchments with a wide range of physiographic characteristics, human impacts (e.g. hydropower, river regulation) and available observation period length were selected. Including study sites with diverse properties allows for validating the applicability and performance of the introduced modelling approach. The catchments are situated in Austria, Switzerland and Germany, with outlets located in the Austrian Alps or adjacent flatlands. All catchments and gauging stations
135 are shown in Figure 1 and their main characteristics are summarized in Table 1.

The gauging stations are operated by the Austrian Hydrographical Service (HZB) and measure discharge (Q) in 15 minute intervals and water temperature (T_w) in a range of different time intervals (daily mean - 1 minute). The temperature sensors are situated in a way that complete mixing can be assumed, e.g. after a bottom ramp. Consequently, the measured water temperature should reflect the water temperature of the given cross section.

140 The meteorological data used in this study are daily mean air temperature (T_a), daily max air temperature (T_{max}), daily min air temperature (T_{min}), precipitation sum (P) and global radiation (GL). T_a , T_{max} , T_{min} and P were available from the SPARTACUS project (Hiebl and Frei, 2016, 2018) on a 1x1 km grid from 1961 onward. The SPARTACUS data was generated by using observations and external drift kriging to create continuous maps. GL data were available from the INCA analysis (Integrated Nowcasting through Comprehensive Analysis) (Haiden et al., 2011, 2014) from 2007 onward. The INCA analysis
145 used numerical weather simulations in combination with observations and topographic information to provide meteorological analysis and nowcasting fields of several meteorological parameters on a 1x1 km grid in 15-60 minute time steps. For the presented study, the 15 minute INCA GL analysis fields were aggregated to daily means. The catchment means of all variables are shown in Table 1. By using high resolution spatially distributed meteorological data as basis for our inputs, we aim to better represent the main drivers of water temperature changes in the catchments. Similar data sets are available for other parts of
150 the world, e.g. globally (Hersbach et al., 2020), for North America (Thornton et al., 2014; Werner et al., 2019), for Europe (Brinckmann et al., 2016; Razafimaharo et al., 2020) and for China (He et al., 2020).

2.2 Data preprocessing

The applied data preprocessing consists of aggregation of gridded data, feature engineering (i.e. deriving new features from existing inputs) and splitting the data into multiple sets of input variables. Since river water temperature is largely controlled
155 by processes within the catchment, variables with an integral effect on water temperature over the catchment (i.e. T_a , T_{max} , T_{min} , P and GL) are aggregated to catchment means.

Computing additional features from a given data set (i.e. feature engineering) and therefore having additional data representation, can significantly improve the performance of machine learning models (Bengio et al., 2013). Previous studies have shown that especially time information is important for water temperature prediction. This includes time expressed as day of
160 the year (e.g., Hadzima-Nyarko et al., 2014; Li et al., 2014; Jackson et al., 2018; Zhu et al., 2018, 2019c, d), the content of the Gregorian calendar (i.e. year, month, day) (Zhu et al., 2019b), or expressed as the declination of the sun (Piotrowski et al.,




Figure 1. Study sites in Austria, Germany and Switzerland. All gauging station IDs refer to the IDs in Table 1.

Table 1. Overview of study catchment characteristics, including means of meteorological values of catchment means, catchment areas, available data time periods and number of years with data. IDs are referring to the IDs used in Figure 1.

| ID | Catchment | Gauging Station | Time period | Years | Area (km^2) | T_w ($^{\circ}C$) | Q (m^3/s) | T_a ($^{\circ}C$) | P (mm) | GL (W/m^2) |
|----|-------------|-----------------|-------------|-------|--------------------|--------------------------|--------------------|--------------------------|-----------------|---------------------|
| 1 | Kleine Mühl | Obermühl | 2002-2015 | 14.0 | 200.2 | 8.87 | 3.12 | 9 | 2.73 | 135 |
| 2 | Aschach | Kropfmühle | 2004-2015 | 11.9 | 312.2 | 10.78 | 3.80 | 10 | 2.50 | 136 |
| 3 | Erlauf | Niederndorf | 1980-2015 | 35.3 | 604.9 | 9.42 | 15.27 | 8 | 3.59 | 127 |
| 4 | Traisen | Windpassing | 1998-2015 | 17.7 | 733.3 | 9.83 | 14.88 | 8 | 3.33 | 131 |
| 5 | Ybbs | Greimpersdorf | 1981-2015 | 34.7 | 1 116.6 | 9.87 | 31.50 | 8 | 3.77 | 127 |
| 6 | Saalach | Siezenheim | 2000-2015 | 16.0 | 1 139.1 | 8.50 | 39.04 | 7 | 4.60 | 135 |
| 7 | Enns | Liezen | 2006-2015 | 10.0 | 2 116.2 | 1.19 | 67.56 | 6 | 3.60 | 137 |
| 8 | Inn | Kajetansbrücke | 1997-2015 | 18.8 | 2 162.0 | 6.00 | 59.26 | 0 | 2.56 | 153 |
| 9 | Salzach | Salzburg | 1977-2015 | 39.0 | 4 425.7 | 7.63 | 178.11 | 5 | 4.16 | 136 |
| 10 | Donau | Kienstock | 2005-2015 | 11.0 | 95 970.0 | 10.77 | 1 798.31 | 10 | 2.13 | 131 |

2015), which is a function of the day of the year. Nevertheless, using cyclical features like day of the year as an integer variable will most likely reduce model performance, since the days 1 and 365 are as close together as 1 and 2. To translate time information into a more suitable format we chose to transform months and days of months into trapezoidal fuzzy sets, called fuzzy months. Similar to dummy encoding, the values of a fuzzy month are between 0 and 1. They are equal 1 on the 15th of the corresponding month and linearly decreasing each day until they are zero on the 15th day of the previous and following month.



Therefore, the values of two adjacent months will be around 0.5 at the turn of the month. By encoding the categorical variable "month" into these 12 new fuzzy variables, it should be possible to represent time of the year influence more smoothly, as no jumps in monthly influence are possible. Initial test showed that the advantage of this representation exceeds the disadvantage of using 12 variables instead of one or two. A similar approach for encoding time variables was already applied by Shank et al. (2008).

Beside time variables, a previous study by Webb et al. (2003) showed that lag information are significantly associated with water temperature and can improve model performance. Therefore, to allow for information of previous days to be used by the models, the lags of all variables for the 4 previous days are computed and used as additional features.

Using these input variables six experiments with different sets of inputs considering different levels of data availability are defined. The variable compositions of all experiments are shown in Table 2. All features include 4 lags and each experiment also includes fuzzy months as inputs. Experiment 0 (Tmean) acts as another simple benchmark in which only mean water temperature and fuzzy months are used for predictions. Experiment 1 (T) will be able to show the benefit of including T_{max} and T_{min} . Experiments 2-4 consists of combinations of experiment 1 with precipitation and discharge data. Experiments 5-6 include combinations with GL and therefore include only data of the time period 2007-2015 in which GL data was available.

Table 2. Overview of available meteorological and hydrological variables and the composition of the different input data set experiments. If an input variable is used in a data set the lags for the 4 previous days are included as well. Additional to the shown variables, all experiments use fuzzy months as input.

| Experiment | T_a | T_{max} | T_{min} | P | Q | GL |
|------------|-------|-----------|-----------|-----|-----|------|
| 0 (Ta) | X | | | | | |
| 1 (T) | X | X | X | | | |
| 2 (TP) | X | X | X | X | | |
| 3 (TQ) | X | X | X | | X | |
| 4 (TQP) | X | X | X | X | X | |
| 5 (TPGL) | X | X | X | X | | X |
| 6 (TQPGL) | X | X | X | X | X | X |

180

2.3 Benchmark models

Two widely applied models for stream water temperature prediction are used as benchmark for all models tested in this study: multiple linear regression models (LM) and air2stream (Toffolon and Piccolroaz, 2015). By including these two models, it will be possible to compare this study's results to a wider range of previous studies, which were investigating models for stream water temperature prediction.

185



2.3.1 Linear regression

Linear regression models are widely used for river water temperature studies. Earlier studies used mainly air temperature as a regressor to predict river water temperature (e.g., Smith, 1981; Crisp and Howson, 1982; Mackey and Berrie, 1991; Stefan and Preud'homme, 1993). More recent publications are using a wider range of input variables or some modification to the standard
190 linear regression model (e.g., Caldwell et al., 2013; Li et al., 2014; Segura et al., 2015; Arismendi et al., 2014; Naresh and Rehana, 2017; Jackson et al., 2018; Trinh et al., 2019; Piotrowski and Napiorkowski, 2019).

The ordinary least square linear regression model is defined as:

$$\mathbf{Y} = \beta\mathbf{X} + \epsilon \quad (1)$$

where \mathbf{Y} denotes the vector of the dependent variable (river water temperature), \mathbf{X} denotes the matrix of independent variables
195 (e.g. daily mean air temperature, global radiation), β denotes the vector of model coefficients and ϵ denotes the error term. ϵ is assumed to be normal distributed with a diagonal covariance matrix. The estimates for the model coefficients and the independent variable, which minimize the sum of squared errors are given by

$$\hat{\mathbf{Y}} = \hat{\beta}\mathbf{X} \quad (2)$$

$$\hat{\beta} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y} \quad (3)$$

200 where $\hat{\mathbf{Y}}$, and $\hat{\beta}$ represent estimated values. The linear regression model applied in this study includes an intercept and the variables T_a and Q as independent variables to predict T_w .

2.3.2 air2stream

Air2stream (Toffolon and Piccolroaz, 2015) is a hybrid model for predicting river water temperature, which combines a physical based structure with a stochastic parameter calibration. It was already applied in multiple studies over a range of catchments
205 and generally had an improved performance compared to linear regression models (e.g., Piccolroaz et al., 2016; Yang and Peterson, 2017; Piotrowski and Napiorkowski, 2018; Zhu et al., 2019d; Piotrowski and Napiorkowski, 2019; Tavares et al., 2020). Air2stream uses the inputs T_a and Q and was derived from simplified physical relationships expressed as ordinary differential equations for heat budgeted processes. Due to this simplification it may be applied like a data driven model, which depends on parameter calibration. The 8-parameter version of air2stream is defined as

$$210 \frac{dT_w}{dt} = \frac{1}{\theta^{a_4}} [a_1 + a_2 T_a - a_3 T_w + \theta (a_5 + a_6 \cos(2\pi(\frac{t}{t_y} - a_7))) - a_8 T_w] \quad (4)$$

where t is the time in days, t_y is the number of days per year, \bar{Q} the mean discharge, $\theta = Q/\bar{Q}$ is the dimensionless discharge and a_1, \dots, a_8 the model parameters. This differential equation is numerically integrated at each time step using the Crank-Nicolson numerical scheme (Crank and Nicolson, 1947) and the model parameters are calibrated using the Particle Swarm Optimization (Kennedy and Eberhart, 1995).



215 2.4 Machine Learning Models

In this study we compare 6 different machine learning models: step-wise linear regression (step-LM), Random Forest (RF), eXtreme Gradient Boosting (XGBoost), Feedforward neural networks (FNN) and two Recurrent neural networks (RNN) - the a Long short-term network (RNN-LSTM) and the Gated recurrent unit (RNN-GRU). An overview and simple depiction of the models is shown in Figure 2.

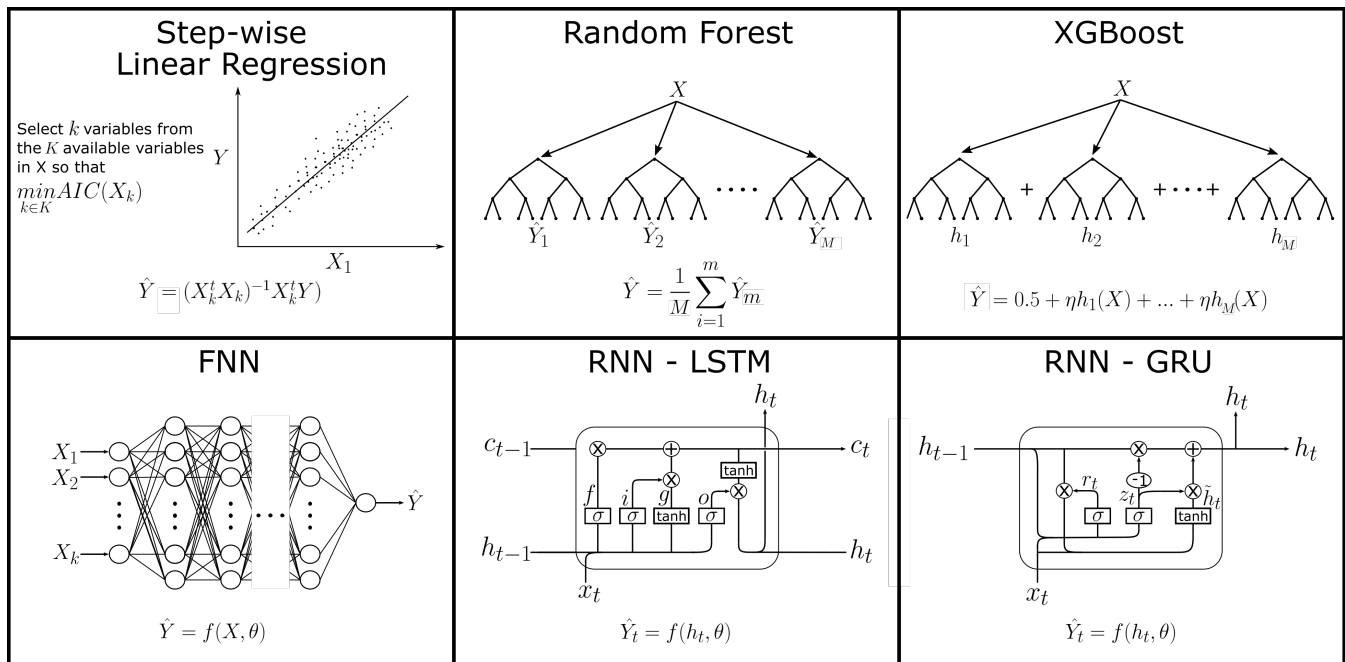


Figure 2. Overview of the applied models with \hat{Y} denoting estimated water temperatures and X the matrix of observed variables. $\hat{Y}_1, \dots, \hat{Y}_M$ are the predictions from individual RF trees. h_1, \dots, h_M are the predicted residuals from individual XGBoost trees. $f(X, \theta)$ denotes a mapping from a FNN with the parameters θ . For a given time step, h_t denotes the hidden internal state of a RNN cell and c_t the internal cell state of a LSTM cell and $f(h_t, \theta)$ the mapping from a RNN with the parameters θ . RNNs consist of a cascade of cells, each feeding their internal states into the next cell, finally resulting in a single feedforward layer estimating \hat{Y} from h_t .

220 2.4.1 Step-wise linear regression

Step-wise linear regression models combine an iterative variable selection procedure with linear regression models. The step-wise variable selection starts at an initial model (e.g. all variables) and removes or adds at each iteration based on a prespecified criterion. We applied the step-wise variable selection starting with an initial model including all variables and using the Akaike Information Criterion (AIC) (Akaike H., 1973). The AIC for a linear regression model is given by:

$$225 \quad AIC = n \times \ln\left(\frac{\sum_{i=1}^{i=n} (Y_i - \hat{Y}_i)^2}{n}\right) + 2k \quad (5)$$



where n is the number of samples, $\ln(\cdot)$ the natural logarithm, Y and \hat{Y} the observed and predicted water temperatures and k the number of selected input variables. The step-wise variable selection is iteratively applied until AIC is at a minimum. Additional to the variables given in section 2.2, interactions terms between T_a , Q , GL and P are included.

2.4.2 Random Forest

230 The Random Forest model (RF) (Breiman, 2001) is an ensemble learning model based on the idea of bagging (bootstrap aggregating) (Breiman, 1996). Bagging predictors average multiple model predictions, where each model is trained on a bootstrapped sample instead of the full observed sample. This randomness introduced by bootstrapping increases the models ability to generalize and to produce stable prediction results.

RF models are bagging predictors which use classification and regression trees (CARTs) as a base learner. RF CARTs
235 recursively apply binary splits to the data to minimize entropy in the tree nodes. This is done until each node reaches a minimum node size or a previously defined maximum tree depth is reached. Breiman (2001) showed that adding further randomness to the bagging method improves prediction accuracy. In random forests this is achieved by only selecting a random subset of available variables for the split at each node. The estimate for the independent variable is given by

$$\hat{Y} = \frac{1}{M} \sum_{m=1}^M f_m(\mathbf{X}) \quad (6)$$

240 where f_m denotes a single fitted CART, M the number of used CARTs, \mathbf{X} the matrix of regressors and \hat{Y} the vector of estimated water temperature. A simplified depiction of the RF algorithm is shown in Figure 2. RF has 2 important hyperparameters: the number of predictors sampled at each node (mtry) and the minimum size of nodes (min node size). The number of trees was chosen to be constant with 500 trees.

2.4.3 XGBoost

245 XGBoost (Chen and Guestrin, 2016) is a tree boosting algorithm that was developed based on the already existing concept of boosting, which was further enhanced to increase efficiency, scalability and reduced overfitting. Similar to bagging, boosting methods combine the prediction of an ensemble of weak learner to improve prediction accuracy. However, while bagging ensemble members are trained in parallel, boosting iteratively trains new ensemble members and adds them to the existing ensemble. Boosting was first introduced by Schapire (1990) and then widely applied after the introduction of the Adaboost
250 algorithm (Freund and Schapire, 1995). Friedman (2001) further enhanced boosting by adding gradient decent optimization for the boosting iterations. This resulted in the development of gradient tree boosting (Friedman, 2002), which uses CART as weak learners.

XGBoost is an implementation of gradient tree boosting with further enhancements in the form of added stochasticity and regularization. The XGBoost estimated for the independent variable is given by

$$255 \hat{Y} = 0.5 + \sum_{m=1}^M \eta f_m(\mathbf{X}) \quad (7)$$



where f_1, \dots, f_M is a sequence of CARTs, $\eta \in [0, 1]$ is the learning rate, M is the number of used CARTs, \mathbf{X} is the matrix of input features and $\hat{\mathbf{Y}}$ is the vector of estimated water temperatures. The m -th tree is trained to predict the residuals of a model of the form given in equation 7, which uses the previous $m-1$ CARTs. The loss function used to train each tree includes a regularization term to prevent overfitting. Additionally, overfitting is reduced by only allowing a random subset of samples and variables, be used for constructing trees and tree nodes at each iteration. A simplified depiction of the XGBoost algorithm is shown in Figure 2.

XGBoost has multiple important hyperparameters that have to be chosen before fitting the model: the maximum number of iterations (nrounds), the learning rate (η), the maximum depth of a tree (max depth), the minimum sum of instance weight needed in a child (min node size), the ratio of random subsamples used for growing a tree (subsample) and the random fraction of variables used for growing a tree (colsample bytree).

2.4.4 Feedforward Neural Network

Feedforward neural networks (FNN) (White and Rosenblatt, 1963) are the first and simplest type of neural networks. FNNs consists of multiple layers of nodes, where each node is connected to all nodes of the previous and following layer. A node applies a linear and a non-linear (activation-) function to its input to produce an output. The general structure of a FNN is shown in Figure 2.

Piotrowski et al. (2020) showed that adding dropout (Hinton et al., 2012; Srivastava et al., 2014; Baldi and Sadowski, 2014) to FNNs for stream water temperature prediction improved performance of single layer FNNs. Dropout refers to randomly dropping nodes from a layer during training, which can prevent overfitting and potentially improve generalization. We added a dropout to every FNN layer and defined the dropout rate as a hyperparameter.

While the parameters (θ) of the linear function get optimized using backpropagation (Rumelhart et al., 1986), FNN have multiple hyperparameters that need to be predefined before training. These hyperparameters include the activation functions, the number of layers, the number of nodes per layer and the dropout ratio. After initial tests, in which a large set of different activation function were applied, we chose the Scaled Exponential Linear Unit (SELU) activation function (Klambauer et al., 2017) for all nodes in the network. SELU includes a normalization, which enhances convergence and avoids both vanishing and exploding gradients during backpropagation. The other hyperparameters are optimized as described in section 2.5.

The here presented hyperparameter optimization approach differs from previous studies, which generally assume a set of fixed number of layers and/or nodes per layer that were derived by a trial and error approach (e.g., Bélanger et al., 2005; Hadzima-Nyarko et al., 2014; Piotrowski et al., 2015; Zhu et al., 2018, 2019d).

2.4.5 Recurrent Neural Networks

In contrast to FNN, Recurrent Neural Networks (RNN) are able to process sequences of inputs. This is achieved by having internal (hidden) states. While there are many different types of RNNs, we focused on the two most widely known, the Long short-term memory (LSTM) (Hochreiter and Schmidhuber, 1997) and the Gated recurrent unit (GRU) (Cho et al., 2014). Each layer of an RNN consists of a sequence of cells that share a common set of weights. The cells of both LSTM and GRU are



shown in Figure 2 and are described in the next sections. A single RNN cell consists of multiple gates, which refers to the
 290 nodes of a cell where non-linear transformations are applied to the inputs and states. Each RNN contains a FNN layer with a
 single node at its end, which is used to compute the predicted values from the hidden states of the last time step (\mathbf{h}_T). Both
 types of RNNs have the same set of hyperparameters that need to be specified before training the model: the number of used
 RNN layers, the number of units per layer, the numbers of timesteps, the dropout ratio, and the batch size.

Due to their internal states and the usage of multiple time steps for prediction, it can be assumed that RNNs do not need time
 295 information (here in form of fuzzy months) for predicting water temperature data. To test this assumption both RNN variants
 are also trained without fuzzy months to check the influence of these additional variables on model performance. Being able to
 achieve equally good results without fuzzy months, would reduce training time considerably due to decreasing the input data
 by 12 dimensions (columns).

Long short-term memory cells

300 Given a sequence of inputs for T time steps x_1, \dots, x_T , where each $x_t \in \mathbb{R}^d$ is a vector of d input features, the forward pass
 of a single LSTM cell with h hidden units is given by following equations:

$$\mathbf{f}_t = \sigma(\mathbf{W}_f \mathbf{x}_t + \mathbf{U}_f \mathbf{h}_{t-1} + \mathbf{b}_f), \quad (8)$$

$$\mathbf{i}_t = \sigma(\mathbf{W}_i \mathbf{x}_t + \mathbf{U}_i \mathbf{h}_{t-1} + \mathbf{b}_i), \quad (9)$$

$$\mathbf{o}_t = \sigma(\mathbf{W}_o \mathbf{x}_t + \mathbf{U}_o \mathbf{h}_{t-1} + \mathbf{b}_o), \quad (10)$$

305 $\tilde{\mathbf{c}}_t = \tanh(\mathbf{W}_g \mathbf{x}_t + \mathbf{U}_g \mathbf{h}_{t-1} + \mathbf{b}_g), \quad (11)$

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tilde{\mathbf{c}}_t, \quad (12)$$

$$\mathbf{h}_t = \mathbf{o}_t \odot \tanh(\mathbf{c}_t) \quad (13)$$

where \mathbf{f}_t , \mathbf{i}_t , and $\mathbf{o}_t \in \mathbb{R}^h$ are the forget gate, input gate and output gate, $\tilde{\mathbf{c}}_t \in \mathbb{R}^h$ is the cell input activation, $\mathbf{h}_t \in \mathbb{R}^h$ is
 the hidden state, $\mathbf{c}_t \in \mathbb{R}^h$ is the cell state and all $\mathbf{W} \in \mathbb{R}^{h \times d}$, $\mathbf{U} \in \mathbb{R}^{h \times h}$ and $\mathbf{b} \in \mathbb{R}^h$ are trainable weights. σ is the sigmoid
 310 function and \tanh the hyperbolic tangent function and \odot is element-wise multiplication.

The hidden state (h_t) is computed from the current input (x_t) and the previous hidden state (h_{t-1}). The amount of informa-
 tion that is passed through the current cell is regulated by the input gate (i_t) and the forget gate (f_t). The cell state (c_t) regulates
 how much memory will be stored in the hidden state (h_t). The output gate (o_t) controls how much information is passed to the
 next cell.

315 Gated Recurrent Units cells

The GRU cell is similar to a LSTM cell, but much simpler. It combines the forget and input gate into a single update gate
 and also merges the cell state and the hidden state. Given a sequence of inputs for T time steps x_1, \dots, x_T , where each $x_t \in \mathbb{R}^d$



is a vector of d input features, the forward pass of a single GRU cell with h hidden units is given by following equations:

$$\mathbf{z}_t = \sigma(\mathbf{W}_z \mathbf{x}_t + \mathbf{U}_z \mathbf{h}_{t-1} + \mathbf{b}_z) \quad (14)$$

320 $\mathbf{r}_t = \sigma(\mathbf{W}_r \mathbf{x}_t + \mathbf{U}_r \mathbf{h}_{t-1} + \mathbf{b}_r) \quad (15)$

$$\hat{\mathbf{h}}_t = \tanh(\mathbf{W}_h \mathbf{x}_t) + \mathbf{U}_h (\mathbf{r}_t \odot \mathbf{h}_{t-1} + \mathbf{b}_h) \quad (16)$$

$$\mathbf{h}_t = (1 - \mathbf{z}_t) \odot \mathbf{h}_{t-1} + \mathbf{z}_t \odot \hat{\mathbf{h}}_t \quad (17)$$

where $\mathbf{z}_t \in \mathbb{R}^h$ is the update gate, $\mathbf{r}_t \in \mathbb{R}^h$ is the reset gate, $\hat{\mathbf{h}}_t \in \mathbb{R}^h$ is the candidate activation, $\mathbf{h}_t \in \mathbb{R}^d$ is the output and all $\mathbf{W} \in \mathbb{R}^{h \times d}$, $\mathbf{U} \in \mathbb{R}^{h \times h}$ and $\mathbf{b} \in \mathbb{R}^h$ are trainable weights.

325 The reset gate (r_t) determines how much information from the previous state will be forgotten when computing the candidate activation ($\hat{\mathbf{h}}_t$). The update gate is the amount of information used from the candidate activation ($\hat{\mathbf{h}}_t$) for computing the current output h_t .

2.5 Bayesian hyperparameter optimization

Choosing adequate hyperparameters for a machine learning model can have a large impact on its performance. Therefore, it is necessary to apply some sort of optimization procedure. While it might be possible to apply a grid search over the range of all possible parameter value combinations for a small set of hyperparameters, it is usually not feasible due to available computational resources. For that reason, we chose to optimize the hyperparameters of nearly all machine learning models in this study with the bayesian optimization method. Only Random forest with 3 hyperparameters is optimized using a grid search. Step-wise linear regression does not have hyperparameters that need optimization.

335 Bayesian optimization is a global optimization method for black-box functions (i.e. lacks known structure and is derivative-free) that is often applied in cases where the objective function is computational expensive to evaluate. It originates from work by Kushner (1964), Zhilinskis (1975), Močkus (1975), Močkus et al. (1978), Močkus (1989) and was later popularized by Jones et al. (1998). It got especially well known for being suitable for optimizing machine learning hyperparameters after a study by Snoek et al. (2012).

340 Bayesian optimization consists of two parts: a method for statistical inference and an acquisition function for deciding the next sample point. The method for statistical inference is usually a Gaussian process (GP) which provides an estimated posterior distribution at each iteration that is an estimate for the function that should be optimized. The acquisition function is used to find the next point to evaluate during each optimization step and was chosen to be the Upper Confidence Bound (UCB) (Srinivas et al., 2009) in this study. In summary, bayesian optimization constructs a surrogate model at each iteration during optimization to choose a suitable next point. The hyperparameters of all optimized models and their chosen bounds are given in the Appendix A.



2.6 Evaluation metrics

The objective function for all models and the hyperparameter optimization is the mean squared error (MSE):

$$MSE = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2 \quad (18)$$

350 where n is the number of samples (days) and y_i the observed and \hat{y}_i the predicted water temperatures. For comparing the performance of different models the root mean squared error (RMSE) and the mean absolute error MAE are used:

$$RMSE = \sqrt{MSE} \quad (19)$$

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (20)$$

2.7 Experimental setup

355 To be able to objectively compare all applied models, the available data sets are split into two parts: the first 80% of the time series were used for training/validation and the last 20% were used for testing. We deliberately did not choose a random split, because predicting water temperatures for a future time period is a more adequate test for models. This is especially relevant for water temperature, which is characterized by non-stationarity due to climate change (Van Vliet et al., 2013). The training/validation and test time series are compared to assess the difference of water temperature distribution of all catchments.

360 The step-wise linear regression model, RF and XGBoost are optimized using cross-validation (CV). Two kinds of CV are applied: a 5 times repeated 10-fold CV and a time series CV. While the 10-fold CV splits the data randomly, the time series CV gradually adds data to an initial part of the time series while evaluating the performance of each step. The time-series CV started with an initial window of 730 days for training the following 90 days for validation. The training set is increased by 90 days at each different cross validation set until the full time series except for the last 90 days was used. Therefore, instead of
365 10 folds, the number of folds for the time-series CV depends on the time series length.

Due to computational and time constrains hyperparameter optimization for all neural networks was done by using a 60%/20% of the training/validation data. This allows model validation performance estimation by training a model once, while a 5 times repeated 10-fold CV would require training a model 50 times. This is the standard way of training neural networks for real-world applications.

370 Bayesian hyperparameter optimization consists of 20 random parameter samples and 40 iterations of optimization. The data inputs for all neural networks were standardized by subtracting the mean and dividing by the standard deviation of the training data. The optimized neural networks hyperparameter sets are used to create 5 independently trained models, from which an ensemble for prediction is created by taking the average of all 5 prediction results. Using ensembles of networks is a way to significantly increase a neural networks ability to generalize and is an often applied approach which was first introduced by the
375 work of Hansen and Salamon (1990). In addition, early stopping with patience = 5 was applied to all neural networks to avoid overfitting.



The best performing model for each model type and experiment is chosen using the validation RMSE. Test RMSE and MAE results are only compared after choosing the models with minimum validation RMSE. Consequently, it might be possible that some models have a superior test performance, but are not chosen as the best performing model for a specific model type and/or experiment. This should reflect a real-world application, where test data acts as a previously unknown future time series.

2.8 Statistical tests

The Kruskal-Wallis test (Kruskal and Wallis, 1952) was used to test for differences in overall model performances, different training/model characteristics and different data inputs. Dunn's test for multiple comparison (Dunn, 1964) was used for pairwise comparisons between model performances. To investigate the association of model types, experiments and catchments with test RMSE an ordinary least square linear regression model was used. Level of significance was set to $p = 0.05$ for all statistical tests.

2.9 Open-source R package

All preprocessing steps and models were implemented in the open source R package *wateRtemp*, which is available under github.com/MoritzFeigl/wateRtemp or from Feigl (2021a). This provides easily applicable modelling tools for the water temperature community and allows all results of this study to be replicated. All programming was done in R (R Core Team, 2020), where the model development relied heavily on Caret (Kuhn, 2020), xgboost (Chen et al., 2020) and TensorFlow (Allaire and Tang, 2020), and the visualizations on ggplot2 (Wickham, 2016).

3 Results

3.1 Time period characteristics

Due to climate change induced steadily increasing air temperatures, also water temperature is not a stationary variable and exhibits warming trends. This is clearly visible when comparing the change in number of extreme warm days and the increase of mean water temperature in all catchments with time. For this we compared the training/validation and test time data in each catchment. Since test data consists of the last 20% of the overall data, the exact length of these time series is depending on the catchment where the test time consists of years in 2008-2015. We can observe an increase of 138% of the median number of days with water temperature above the 90% quantile between training/validation and test time period in all catchments. This increase ranges from 69%, or from 32 days to 54 days, in the Donau catchment and up to 285%, or from 26 days to 100 days, in the Salzach catchment. This change is even more pronounced when comparing the last year of test data (2015) to all other available years, where the median number of days with water temperatures above the 90% quantile (computed for the overall timeseries) of all catchments increases by 273%. Figure 3 shows the corresponding boxplots of days with stream temperature above the 90% quantile for each catchment in training/validation and in test time period. A similar pattern can be observed in the changes of mean yearly stream temperatures. The median increase of mean yearly water temperature of all catchments is



0.48 °C when comparing training/validation with test time period and 0.77 °C when comparing the last year of the test period (2015) with all other years. Since the test period is, as shown here regarding extremes, different from the training/validation period, the models are also, at least to some extent, tested on how they perform under instationary conditions. This is a test, where environmental models often fail (e.g. Kling et al., 2015).




Figure 3. Boxplots showing the distribution of numbers of days with stream temperatures above the 90% quantile per year for all study catchments for the training/validation and the test time period, where the test time period consists of the last 20% of data in each catchment. The 90% quantile values were estimated using the full time series for each catchment.

410

3.2 Overall performance comparison

Table 3 gives an overview of the best applied machine learning models and the two benchmark models LM and air2stream. The mean test RMSE of LM is 1.55 °C with an overall range of [1.25, 2.15] °C, while air2stream has a mean test RMSE of 0.98 °C with an overall range of [0.74, 1.17] °C. The performance results for each catchment show that air2stream always outperformed LM and consequently results in a significant lower test RMSE ($p < 0.001$). The mean test RMSE of the best machine learning models per catchment is 0.55 °C with an overall range of [0.42, 0.82] °C and always outperformed the air2stream benchmark. Based on the RMSE means, the highest performing ML model is 64% and 43% better, compared to LM and air2stream. This results in a significantly lower test RMSE of the tested machine learning models compared to the air2stream benchmark ($p < 0.001$).

420 Both XGBoost and FNN were found to be the best performing model in 4 of 10 analysed catchments each. RF was the best performing model in the Saalach catchment and RNN-LSTM in the Donau catchment. Step-LM and RNN-GRU did not



outperform the other models in any of the study catchments. Experiment 3, which only includes air temperature and discharge input features, resulted in the best performing model in four catchments. Experiment 6, which included all available input features, also produced the best performing model in four catchments. Experiment 4, which includes air temperature, discharge and precipitation input features, performed best in two catchments.

Table 3. Overview of model performance of the best machine learning model for each catchment and the two reference models. The best performing model results in each catchments are shown in bold font. The best machine learning model for each catchment was chosen by comparing validation RMSE values, while test RMSE and test MAE values were never part of any selection or training procedure. The shown values all refer to the test time period.

| Catchment | Model | Best ML model results | | | LM | | air2stream | |
|-------------|----------|-----------------------|----------------------|---------------------|----------------------|---------------------|----------------------|---------------------|
| | | Experiment | RMSE ($^{\circ}C$) | MAE ($^{\circ}C$) | RMSE ($^{\circ}C$) | MAE ($^{\circ}C$) | RMSE ($^{\circ}C$) | MAE ($^{\circ}C$) |
| Kleine Mühl | XGBoost | 4 (TQP) | 0.740 | 0.578 | 1.744 | 1.377 | 0.908 | 0.714 |
| Aschach | XGBoost | 6 (TQPGL) | 0.815 | 0.675 | 1.777 | 1.408 | 1.147 | 0.882 |
| Erlauf | XGBoost | 6 (TQPGL) | 0.530 | 0.419 | 1.354 | 1.057 | 0.911 | 0.726 |
| Traisen | FNN | 3 (TQ) | 0.526 | 0.392 | 1.254 | 0.970 | 0.948 | 0.747 |
| Ybbs | RF | 3 (TQ) | 0.576 | 0.454 | 1.787 | 1.415 | 0.948 | 0.756 |
| Saalach | XGBoost | 6 (TQPGL) | 0.527 | 0.420 | 1.297 | 1.062 | 0.802 | 0.646 |
| Enns | FNN | 6 (TQPGL) | 0.454 | 0.347 | 1.425 | 1.166 | 1.168 | 0.671 |
| Inn | FNN | 3 (TQ) | 0.422 | 0.329 | 1.376 | 0.098 | 1.097 | 0.949 |
| Salzach | FNN | 4 (TQP) | 0.430 | 0.338 | 1.327 | 1.077 | 0.743 | 0.595 |
| Donau | RNN-LSTM | 3 (TQ) | 0.521 | 0.415 | 2.145 | 1.721 | 1.099 | 0.910 |
| | | Mean: | 0.554 | 0.437 | 1.549 | 1.235 | 0.977 | 0.760 |

Figure 4 shows the results of all models, catchments and experiment combinations. The boxplots in Figure 4a show the range of model performances depending on the model type. Kruskal-Wallis-Test results show no significant difference ($p = 0.11$) of the test RMSE of different model types. Figure 4b shows boxplots of model performance for all experiments. Kruskal-Wallis-Test results show a highly significant difference of test RMSE of the different experiments ($p < 10^{-14}$). The results in Figure 4b show an increase in median performance with increasing number of input features until experiment 4 (TQP). When adding global radiation as additional input parameter, the median performance is not further increasing. This could be explained by a reduced time series length of experiment 5 (TPGL) and 6 (TQPGL), since global radiation was only available from 2007 on. A comparison between experiments with equal time series lengths (experiments 0-4 and experiments 5-6) also indicate that runoff information improves the modelling performance.

Figure 4c illustrates the RMSE performance results for each catchment shown as boxplots. A corresponding figure of the MAE results is shown in A1. The boxplots are overlaid with scatterplot points adding an overview of the individual performance of each model and experiment combination. To account for a better visibility, the scatterplot points are shifted in



horizontal direction randomly. The difference in performance between catchments is clearly visible and ranges from a median RMSE of around $0.93\text{ }^{\circ}\text{C}$ in catchments Kleine Mühl and Aschach down to a median RMSE of $0.58\text{ }^{\circ}\text{C}$ in the Inn catchment.

440 Figure 4c also includes the air2stream benchmark performance shown as a grey line for each catchment. Nearly all tested experiments and model combinations showed improved performance compared to the air2stream benchmark. Only in 5 catchments we could observe models in combination with experiments 0, 1, 5 and one time with experiment 6 that predicted worse than air2stream. This are surprisingly few models considering the fact that both experiment 0, 1, 5 and 6 are heavily constrained due to the amount of information that is available for prediction. Experiment 0 and 1, which only use air temperature, are still

445 able to improve predictions compared to air2stream for all model types in 7 catchments. Similarly, experiments 5 and 6 with only 6 years of training data are able to improve predictions compared to air2stream for all model types in 5 catchments.

From the results in Figure 4a,b,c it seems likely that performance is in general influenced by the combination of model, data inputs (experiment) and catchment, while the influence of different experiments and catchments is larger than the influence of model types on test RMSE. The linear regression model for test RMSE with catchment, experiment and model type as

450 regressors is able to explain most of the test RMSE variance with a coefficient of determination of $R^2 = 0.988$. Furthermore, it resulted in significant association of all catchments ($p < 10^{-15}$), all experiments ($p < 0.005$) and with the FNN model type ($p < 0.001$). The estimated coefficient of the FNN is -0.05 , giving evidence to an prediction improvement when applying the FNN model. All other model types do not show a significant association. However, this might be due to a lack of statistical power, as the estimated coefficients of the model types (mean: -0.01 , range: $[-0.05, 0.02]$) are generally small compared to

455 catchment coefficients (mean: 0.86 , range: $[0.69, 1.06]$) and experiment coefficients (mean: -0.12 , range: $[-0.2, -0.04]$).

Multiple experiments often result in very similar RMSE values for a single model type. Furthermore, the best performing experiment of different model types are always very close in performance. This results in a median test RMSE difference of the best experiments of different model types of $0.08\text{ }^{\circ}\text{C}$ and a median test RMSE difference of the best performing model and the second best model of another model type of $0.035\text{ }^{\circ}\text{C}$. On the other hand, the median difference between the tested

460 machine learning model RMSE and the air2stream RMSE is $-0.39\text{ }^{\circ}\text{C}$.

3.3 Detailed analysis for a single catchment

To further investigate the difference in performance, the prediction results for the last year of the test data (2015) of the Inn catchment are examined. The year 2015 was chosen for comparison, since it has an extraordinary large number of days with high water temperatures and therefore can be used to give a robust estimate of model performance. It is a strong test under

465 instationary conditions. The time period 1997-2014 has a median of 30 days per year with water temperatures over $11\text{ }^{\circ}\text{C}$, whereas 102 days with such high water temperature could be observed in the year 2015. Figure 5 shows the prediction results of each model (red lines) compared to the observation (blue line) and all other model predictions (grey lines) for the year 2015 and the corresponding RMSE and MAE result for that year.

The two benchmark models (LM and air2stream) show large differences between prediction and observations and show in

470 general a very different behaviour than all tested machine learning models. While the largest prediction errors of the tested

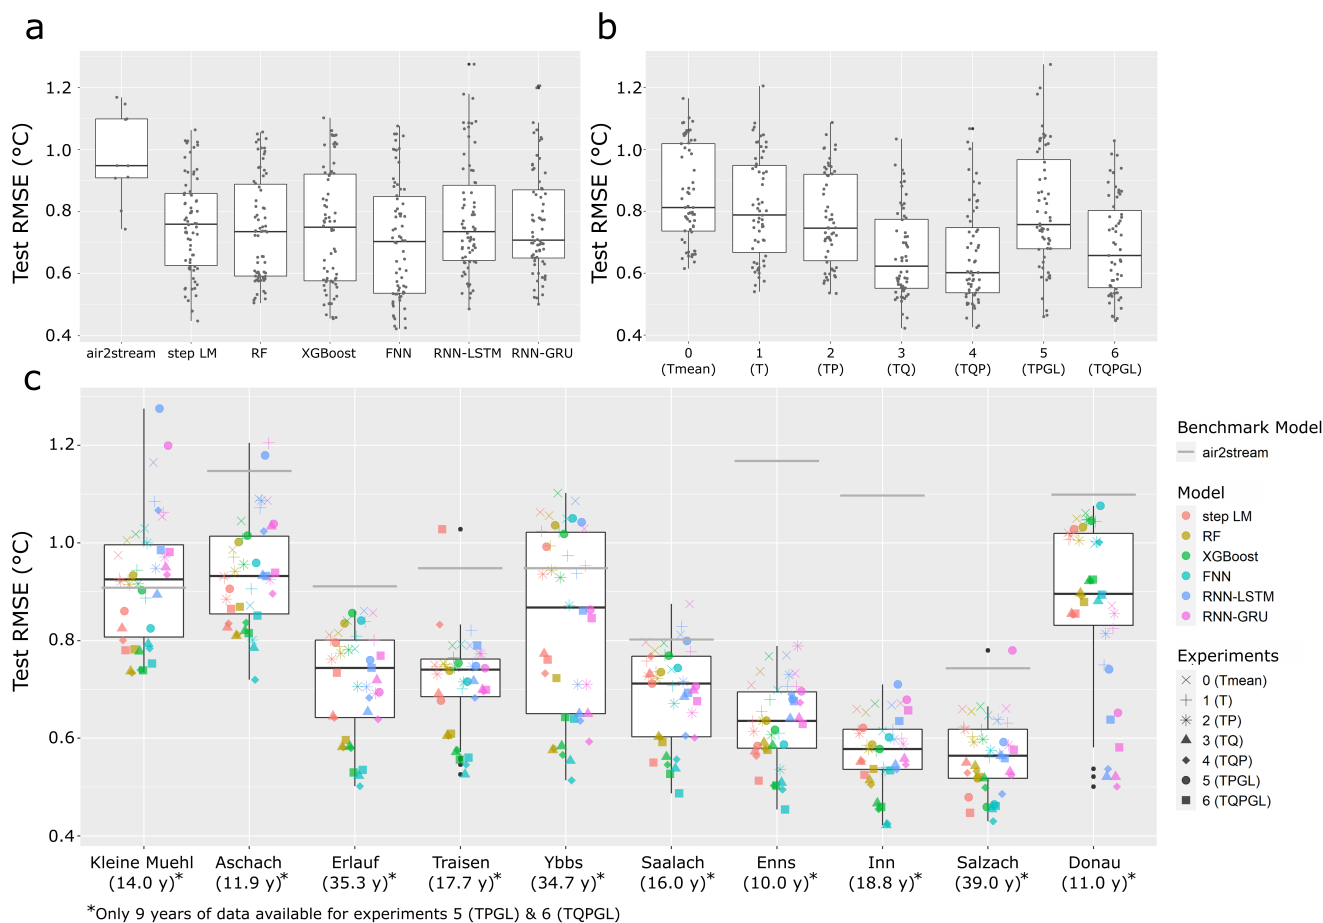


Figure 4. Boxplots of model performance comparing **a** the different machine learning models, **b** the different experiments and **c** model performance in each catchment with additional scatter plot overlay to show performance of individual combinations of catchments, models and experiments. The catchments are ordered by catchment size from smallest (left) to largest (right) with additional information of the available time series length in parenthesis below. The air2stream benchmark performance is shown as grey line for each catchment. Due to the much larger test RMSE values, LM performance is not shown to account for a better visibility.

machine learning models occur during similar time periods, large deviations can be observed over the whole year in both benchmark models.

The largest prediction errors of all machine learning models occur during warmer periods and peaks in the summer months and during periods of low water temperature in November - December. This is clearly visible in all tested models. Therefore, differences in RMSE and MAE mainly result from their performance during these periods and consequently can be quite large even though the actual numerical difference is rather small. This can be observed when comparing the results of best performing model FNN and RNN-GRU in Figure 5. Both models produce similar prediction results for the largest part of

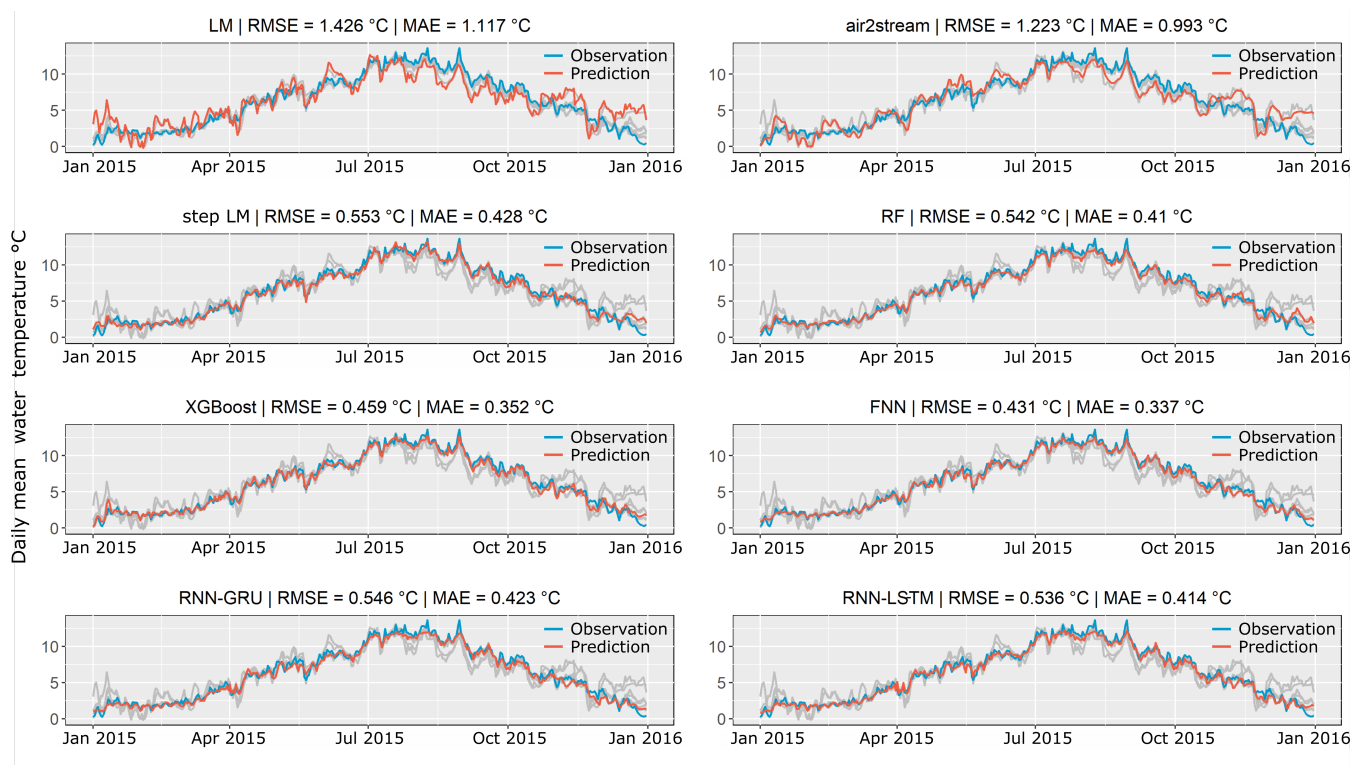


Figure 5. Comparison of the prediction of all tested model types for the Inn catchment for the year 2015. Data from 2015 was not used for training and validation. Prediction results for each model are shown with red lines, while the observations are shown in blue lines. The predictions of all other models are illustrated with grey lines.

the year, but the FNN model is better able to predict the peaks with high water temperatures in the summer months, which results in a RMSE and MAE difference of 0.115 and 0.086, respectively. Very small differences in RMSE and MAE as seen
480 between two best performing models, FNN and XGBoost, results in only very subtle differences in the predicted time series. Very similar observations can be made when analyzing the prediction results in the other catchments. The only exception can be observed in the largest catchment Donau (Figure A2), where the time series is much smoother with relatively few peaks in water temperature. This results in the RNN models being the best performing models with a large performance difference compared to all other models.

485 3.4 Influence of time variables for RNNs, cross validation methods

Removing time information in the form of fuzzy months from the training data of RNNs does not significantly change the catchment test RMSE ($p = 0.17$). However, the optimal number of time steps estimated by the hyperparameter optimization is significantly increased ($p = 0.02$). By removing time information from the inputs, the optimal timesteps are 37.78 days longer

than when using time information as additional input. This significantly increases model training time ($p = 0.034$) with a mean
490 difference of 132.45 minutes.

The different CV schemes applied on step LM, RF and XGBoost showed no significant difference in performance ($p = 0.91$).

3.5 Influence of hyperparameters on model results

The influence of different sets of hyperparameters on model performance is shown in Figure 6. This figure shows the validation
495 RMSE for all parameter sets which were used during hyperparameter optimization. A large difference in the range of perfor-
mance can be observed for different models. Validation RMSE means, standard deviations, minimum and maximum of all
models are shown in Table 4. The largest variability is apparent in the FNN results with a validation RMSE standard deviation
of $\sigma_{FNN} = 1.60$ °C and a overall RMSE range of [0.41, 16.6] °C. This is followed by XGBoost which has multiple outliers
in each catchment that increase the performance spread, resulting in $\sigma_{XGBoost} = 1.07$ °C and the RMSE range [0.40, 9.15]
°C. Both RNNs show a very similar performance distributions with a RMSE range of around [0.45, 6.3] °C. Compared to all
500 other tested models, the RF model has a much smaller spread in performance resulting from different hyperparameter sets with
 $\sigma_{RF} = 0.16$ °C and a resulting RMSE range of [0.45, 1.14] °C.




Figure 6. Boxplots showing the validation RMSE distribution for different hyperparameter sets for all model types, catchments and experi-
ments. The catchments are ordered by catchment size from smallest (left) to largest (right) with additional information of the available time
series length in parenthesis below.



Table 4. Validation RMSE means μ , standard deviations σ , maxima and minima for all model types resulting from hyperparameter optimization.

| Model | Validation RMSE ($^{\circ}C$) | | | |
|----------|---------------------------------|----------|------|------|
| | μ | σ | min | max |
| RF | 0.70 | 0.16 | 0.45 | 1.14 |
| XGBoost | 0.95 | 1.07 | 0.40 | 9.15 |
| FNN | 1.70 | 1.60 | 0.41 | 16.6 |
| RNN-LSTM | 0.97 | 0.53 | 0.46 | 6.4 |
| RNN-GRU | 0.91 | 0.44 | 0.45 | 6.3 |

4 Discussion

In this study, we show the stream water temperature prediction performance of 6 machine learning models with a range of input data sets in 10 catchments and compared them to two widely used benchmark models. The results show generally a very similar performance of the tested machine learning models with a median test RMSE difference of $0.08^{\circ}C$ between models. In contrast, the models had a significantly improved performance when compared to the air2stream benchmark model, with a mean test RMSE decrease of $0.42^{\circ}C$ (42%). Results showed that nearly all of test RMSE variance ($R^2 = 0.99$) can be explained by the catchment, the input data set and the model type. This also showed that the performance is significantly influenced by the type of input data, where more inputs generally performed better and that of all models, only the FNN model had a significant association with lower test RMSE values. Furthermore, a wide range of performance was observed for different hyperparameter sets for the tested models, with extremely large RMSE standard deviation ($1.60^{\circ}C$) observed in the FNN results.

Except for very few model types and experiment combinations, all tested machine learning models showed an improved performance when compared to the two benchmark models. The difference between the benchmark and the tested models was not only visible in the resulting test RMSE and MAE values, but also clearly visible in the range and time of occurrence of large prediction errors in the predicted time series (see Figure 5). Given the range of estimated coefficients of the catchments ($[0.69, 1.06]$), data inputs ($[-0.2, -0.04]$) and model types ($[-0.05, 0.02]$) in the regression model for test RMSE, we can state that: given an adequate model setup and selected hyperparameters, the influence of different data inputs and different catchments is much larger than the influence of the model types. However, there seems to be an advantage of using the FNN model, as it was the only model that had a significant association with lower RMSE values and also the largest estimated coefficient of all model types (-0.05).

The here presented result show that FNN and XGBoost perform best in 8 of 10 catchments and are therefore a first choice for water temperature prediction tasks. For modelling large catchments with comparable size to the Donau catchment (96 000



km^2), where long-term dependencies seem to be more relevant, RNNs are the best choice. Both RNN architectures, GRU
525 and LSTM, produce very similar results in the Donau catchment, with a best test RMSE of approximately $0.52\text{ }^\circ C$. This is
considerably lower than the median test RMSE of the other models ($0.90\text{ }^\circ C$) and the air2stream benchmark ($1.10\text{ }^\circ C$). The
RF model has the lowest standard deviation in resulting RMSE depending on the chosen hyperparameter ($0.16\text{ }^\circ C$) and thus
might be the most reasonable choice in situations with limited computational resources. More input data is generally better, but
the combination of air temperature and discharge input data already produces prediction results with a median RMSE of 0.62
530 $^\circ C$. This can be further enhanced by adding precipitation data, which decreases the median RMSE further to $0.60\text{ }^\circ C$. Adding
GL data can potentially increase performance as well, as experiment 6 shows a similar performance range to experiment 3
while using only 6 years of training data. Results of experiment 2 (TP), which is most relevant for practical application as it
uses inputs that are general available for most regions and from climate models, show a median test RMSE of $0.75\text{ }^\circ C$. This is
only a 19% reduction in RMSE performance compared to the experiment with the lowest median RMSE and an improvement
535 of 21% compared to air2stream. Thus, application with this set of widely available data inputs is able to produce prediction
performance improving the current state of the art and could be used for short term forecasts and climate change studies.

The presented machine learning approaches could considerably improve prediction results compared the current state-of-
the-art air2stream model. This stands in contrast to the findings of Zhu et al. (2019d), which assessed the performance of a
suit of machine learning models for daily stream water temperature. Zhu et al. (2019d) results showed that air2stream had an
540 improved performance when compared to FNNs, Gaussian process regression and decision tree models in 8 catchments using
water temperature, discharge and day-of-year as model inputs. The here presented air2stream results have a test RMSE range
of $[0.74, 1.17]\text{ }^\circ C$, which is comparable to results of Zhu et al. (2019d) with $[0.64, 1.16]\text{ }^\circ C$ and also to other studies applying
air2stream, e.g. Piotrowski and Napiorkowski (2018) with a range of $[0.625, 1.31]\text{ }^\circ C$. This leads us to the conclusion that our
benchmark performance is in line with other air2stream applications and therefore provides a consistent reference, even though
545 air2stream was originally set up for the use of point source data and not the catchment means that we used to make results
comparable to the tested machine learning models. Consequently, our presented approaches show a significant improvement
compared to existing machine learning daily stream water temperature prediction models.

Due to the lack of physical restraints, statistical modelling approaches are often suspected to fail when extrapolating outside
their training data range (Benyahya et al., 2007). However, machine learning methods are more powerful and flexible than
550 previous modelling approaches and give spatial and temporal information at different scales (Reichstein et al., 2019). This
is especially important for climate change studies, where increasing air temperature might change the statistical relationships
between meteorological drivers and stream water temperature. To investigate the extrapolation performance of considered ML
methods, we selected the much warmer recent years of the time series as test period and analyzed the year with the most
frequent days of extreme temperatures in detail. All tested models were able to produce predictions with a performance close
555 to the training performance in the test time period and in the year with the most temperature anomalies. These results show
that it is still possible to produce robust prediction results at least for short time predictions (1-8 yrs) under a changing climate.
Successful extrapolation for short term periods suggest that also mid-to long term predictions might produce reasonable results.
However, this can only be evaluated based on future observations. It is clear that the ML approaches will fail in extrapolation,



when catchment properties change with time. In the context of high alpine, glacier dominated catchments, for example, it can
560 be assumed that the water temperature characteristics will change, when glaciers vanish. As a consequence, the underlying
processes leading the water temperature in the stream change. These changes are not reflected in the ML approaches. It would
need more physically or process-based approaches. For example, air2stream would not have an advantage in this respect.

Depending on the machine learning model, our results varied considerably with the chosen hyperparameters. Especially the
two best performing models, XGBoost and FNNs, show an extreme variance in performance due to the chosen hyperparame-
565 ters. This leads to the conclusion that flexibility might be necessary for a well performing model, but that it is also a possible
source of error or reduced model performance. These findings highlight the importance of hyperparameter optimization of
machine learning models and might be a possible explanation of the fact that especially FNNs did not perform equally well
in other studies. Most publications reporting findings regarding FNN performance for stream water temperature tested only
a small set of FNN hyperparameter combinations, mostly chosen by trial and error (e.g. Piotrowski et al., 2015; Rabi et al.,
570 2015; Abba et al., 2017; Zhu et al., 2018; Temizyurek and Dadaser-Celik, 2018; Zhu et al., 2019d). Our results show the
extremely large influence of hyperparameters and therefore rendering any trial and error approach as insufficient and certainly
non-optimal.

RNNs are successfully applied in current rainfall-runoff modelling studies (e.g. Kratzert et al., 2018, 2019; Xiang et al., 2020;
Li et al., 2020), and are thus a promising candidate for stream water prediction. However, our results show a below average
575 performance in most catchments when compared to the other tested machine learning models. This is especially relevant, since
compared to the other methods, RNNs use a range of previous timesteps (optimized hyperparameter) for prediction, which
contains much more information than the 4 previous timesteps available for the other models. RNNs are the best performing
models in the largest catchment indicating that RNNs are especially strong when processes with long-term dependencies
have to be described. These long-term dependencies result most likely from increased concentration times, which is generally
580 dependent on catchment size (McGlynn et al., 2004). For all other catchments in this study, the 4 day lagged variables seem to
be sufficient and RNNs are not able to predict the corresponding fast changes in water temperature. Our results also show the
importance of using time information as input for RNNs. RNNs are generally able to learn the corresponding information from
data, since there is no significant difference in performance for the RNNs with and without time information. However, RNNs
optimized with time information inputs needed a significantly lower number of timesteps for the same prediction performance,
585 thus decreasing computation time and increasing the number of data points available for training.

This study has some limitations. Firstly, the selected catchments are all central European catchments with humid condi-
tions. Testing these approaches on Mediterranean or more dynamic hydro-climatological conditions could potentially result
in different importance of input variables (e.g. discharge in arid climates) and performance ranking of models. By selecting
catchments with a wide range of physiographic characteristics this potential bias should be kept at a minimum. Furthermore,
590 the performance of the air2stream benchmark is similar to the performance range of other studies, allowing for comparison.
Secondly, we trained all models only for individual catchments and did not try to produce a global model that could predict
water temperatures in multiple catchments, or even in a prediction of ungauged basin setting. While this is a relevant problem,



we found it necessary to have a comprehensive evaluation of different data inputs, model types and training characteristics before combining all of this in a multi-catchment water temperature prediction model.

595 5 Conclusions

Current standard methods in daily stream water prediction are able to model 10 Austrian study catchments with a mean test RMSE of $1.55^{\circ}C$ (linear regression) and $0.98^{\circ}C$ (air2stream). We tested 6 machine learning models with different data inputs and could produce predictions with a mean RMSE of $0.55^{\circ}C$, an improvement of 64% and 43%. Of these tested models, the FNN model using air temperature, discharge and precipitation and, if available, radiation as inputs, produce the best performing
600 models. With only 6 years of training data state of the art prediction models results can be achieved.

One major influence on performance are model hyperparameter. The variability in performance for different hyperparameters is much larger than for different model types or data inputs. Thus hyperparameter optimization is extremely important for a well performing model. In situations where computing resources are limited and hyperparameter optimization is not possible, the RF model seems to be a reasonable choice for application, because it has the lowest variance in prediction RMSE resulting
605 from the chosen hyperparameters.

RNNs with their internal states and ability to process long time series, are the best performing model type for very large catchments. This is most likely a result from increased concentration times in the catchment. Consequently, estimating concentration times of a catchment for adequately choosing a model type or relevant lags of variables should be included in future research. Applying variable importance estimation methods are also another way to further enhance the understanding of the in-
610 teractions between variables and model performance and could help deciding on the relevant number of variable lags. Applying these methods however, especially for neural networks, is out of scope for this study and will be part of future research.

The study catchment where chosen to have a wide range of physiographic characteristics, but are all located in Central Europe. Thus the range of characteristics is still limited and testing these model approaches in a wider range of catchments is still necessary and should also be included in future research. This will be especially important for developing multi-catchment
615 water temperature prediction models for regional prediction, which is an important next step and topic of current research. The presented machine learning methods, driven with observed meteorological inputs, seem to represent the system in an appropriate manner for applying them to predict river water temperature in changing conditions and may be promising for short time or real time forecasting approaches. The resulting prediction uncertainties in such systems will be mainly related to uncertainties in the meteorological forecasts. By implementing all methods into the open source R package wateRtemp,
620 we hope to further contribute to reproducible research and make the presented methods available and easily applicable for management issues, scientists and industries and to facilitate research on these next steps.

Code and data availability. The R code used to generate all results for this publication can be found in Feigl (2021b). This includes the version of the wateRtemp R package providing all machine learning methods and code that was used for producing the results of this



manuscript. A maintained and continuously updated version of the waterRtemp package can be found at github.com/MoritzFeigl/waterRtemp or in Feigl (2021a).

We do not have the permission for further distribution of the data used in this study. All input data can, however, be acquired from the rights holders of these data sets. The water temperature and discharge data used in this study can be requested from the Central Hydrographical Bureau (HZB) at <https://www.ehyd.gv.at>. The rights for the meteorological data from the INCA and the SPARTACUS datasets belong to the Zentralanstalt für Meteorologie und Geodynamik (ZAMG) and can be acquired from <https://www.zamg.ac.at>.

630 Appendix A: Model hyperparameter bounds

RF: min.node.size: 2-10, mtry: 3-(number of inputs -1)

XGBoost: nrounds: 300-3000, eta: 0.001-0.3, max_depth: 3-12, min_child_weight: 1-10, subsample: 0.7-1, colsample_bytree: 0.7-1, gamma: 0-5

FNN: layers: 1-5, units: 5-200, dropout: 0-0.2, batch_size: 5-150, epochs: 100, early_stopping_patience: 5

635 **RNNs:** layers: 1-5, units: 5-300, dropout: 0-0.4, batch_size: 5-150, timesteps: 5-200, epochs: 100, early_stopping_patience: 5

A1




Figure A1. Boxplots of model performance comparing model MAE values in each catchment with additional scatter plot overlay to show performance of individual combinations of catchments, models and experiments. The catchments are ordered by catchment size from smallest (left) to largest (right) with additional information of the available time series length in parenthesis below. The air2stream benchmark performance is illustrated as grey line for each catchment.

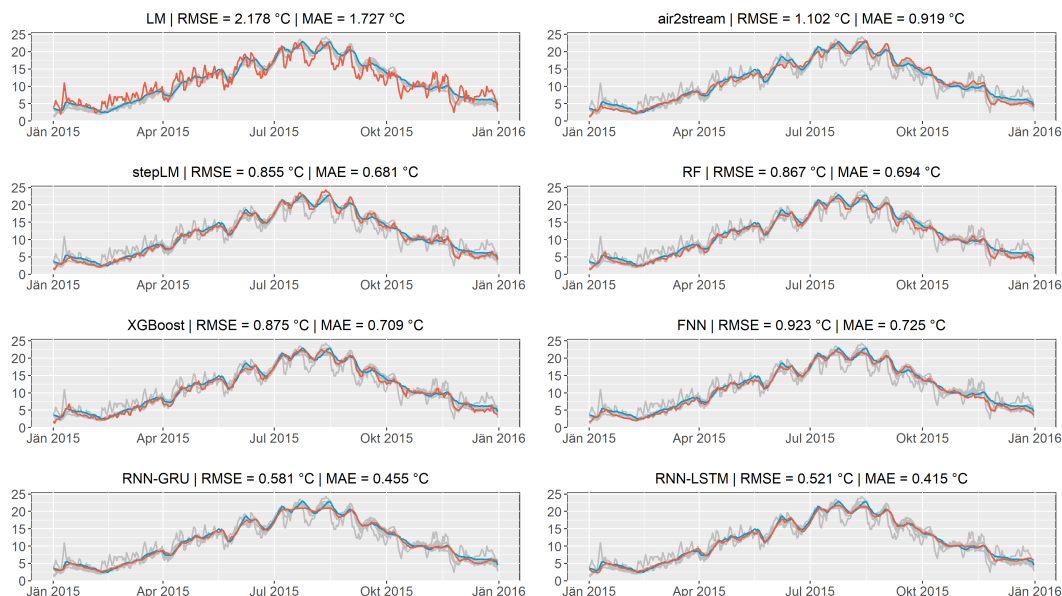


Figure A2. Comparison of the prediction of all tested model types for the Donau catchment for the year 2015. Prediction results for each model are shown with red lines, while the observations are shown in blue lines. The predictions of all other models are shown with grey lines.

Author contributions. KL, MF and MH designed the study and acquired and processed the input data. MF and KL performed all analyses and prepared the figures. MF developed the software published with this work. MH and KS contributed to the methodological framework. MF prepared the paper with contributions from KL, MH and KS.

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

Acknowledgements. This work was partly funded by the Austrian Science Fund FWF, project number P 31213, and the Austrian Academy of Sciences (ÖAW) projects Rechout and Poco-Flood. The computational results presented have been achieved using the Vienna Scientific Cluster (VSC). We also thank Ignacio Martin Santos for providing data from the upper Danube catchment and many valuable discussions about seasonal forecast and team spirit during the Covid-19 pandemic.



645 References

- Abba, S. I., Hadi, S. J., and Abdullahi, J.: River water modelling prediction using multi-linear regression, artificial neural network, and adaptive neuro-fuzzy inference system techniques, in: *Procedia Computer Science*, vol. 120, pp. 75–82, Elsevier B.V., <https://doi.org/10.1016/j.procs.2017.11.212>, 2017.
- Ahmadi-Nedushan, B., St-Hilaire, A., Ouarda, T. B. M. J., Bilodeau, L., Robichaud, É., Thiémondge, N., and Bobée, B.: Predicting river water temperatures using stochastic models: case study of the Moisie River (Québec, Canada), *Hydrological Processes*, 21, 21–34, <https://doi.org/10.1002/hyp.6353>, <http://doi.wiley.com/10.1002/hyp.6353>, 2007.
- Akaike H.: Information theory as an extension of the maximum likelihood principle–In: *Second International Symposium on Information Theory* (Eds) BN Petrov, F. B.N. Petrov, and F. Csáki (Eds.). 2nd International Symposium on Information Theory, 1973.
- Allaire, J. J. and Tang, Y.: tensorflow: R Interface to 'TensorFlow', <https://github.com/rstudio/tensorflow>, 2020.
- 655 Álvarez, D. and Nicieza, A. G.: Compensatory response 'defends' energy levels but not growth trajectories in brown trout, *Salmo trutta* L., *Proceedings of the Royal Society B: Biological Sciences*, 272, 601–607, <https://doi.org/10.1098/rspb.2004.2991>, /pmc/articles/PMC1564075/?report=abstract<https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1564075/>, 2005.
- Arismendi, I., Safeeq, M., Dunham, J. B., and Johnson, S. L.: Can air temperature be used to project influences of climate change on stream temperature?, *Environmental Research Letters*, 9, <https://doi.org/10.1088/1748-9326/9/8/084015>, <https://iopscience.iop.org/article/10.1088/1748-9326/9/8/084015/meta>, 2014.
- 660 Baldi, P. and Sadowski, P.: The dropout learning algorithm, *Artificial Intelligence*, 210, 78–122, <https://doi.org/10.1016/j.artint.2014.02.004>, 2014.
- Beaufort, A., Moatar, F., Curie, F., Ducharme, A., Bustillo, V., and Thiéry, D.: River Temperature Modelling by Strahler Order at the Regional Scale in the Loire River Basin, France, *River Research and Applications*, 32, 597–609, <https://doi.org/10.1002/rra.2888>, <http://doi.wiley.com/10.1002/rra.2888>, 2016.
- 665 Bélanger, M., El-Jabi, N., Caissie, D., Ashkar, F., and Ribí, J. M.: Water temperature prediction using neural networks and multiple linear regression, *Revue des Sciences de l'Eau*, 18, 403–421, <https://doi.org/10.7202/705565ar>, 2005.
- Bengio, Y., Courville, A., and Vincent, P.: Representation learning: A review and new perspectives, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35, 1798–1828, <https://doi.org/10.1109/TPAMI.2013.50>, <http://ieeexplore.ieee.org/document/6472238/>, 2013.
- 670 Benyahya, L., Caissie, D., St-Hilaire, A., Ouarda, T. B., and Bobée, B.: A Review of Statistical Water Temperature Models, *Tech. Rep. 3*, <https://doi.org/10.4296/cwrj3203179>, 2007.
- Boisneau, C., Moatar, F., Bodin, M., and Boisneau, P.: Does global warming impact on migration patterns and recruitment of Allis shad (*Alosa alosa* L.) young of the year in the Loire River, France?, in: *Fish and Diadromy in Europe (ecology, management, conservation)*, pp. 179–186, Springer Netherlands, https://doi.org/10.1007/978-1-4020-8548-2_14, https://link.springer.com/chapter/10.1007/978-1-4020-8548-2_{_}14, 2008.
- 675 Breiman, L.: Bagging predictors, *Machine Learning*, 24, 123–140, <https://doi.org/10.1007/bf00058655>, 1996.
- Breiman, L.: Random forests, *Machine Learning*, 45, 5–32, <https://doi.org/10.1023/A:1010933404324>, 2001.
- Brinckmann, S., Krähenmann, S., and Bissolli, P.: High-resolution daily gridded data sets of air temperature and wind speed for Europe, *Earth System Science Data*, 8, 491–516, <https://doi.org/10.5194/essd-8-491-2016>, <https://essd.copernicus.org/articles/8/491/2016/>, 2016.



- 680 Caissie, D.: The thermal regime of rivers: A review, <https://doi.org/10.1111/j.1365-2427.2006.01597.x>, <https://onlinelibrary.wiley.com/doi/full/10.1111/j.1365-2427.2006.01597.x><https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1365-2427.2006.01597.x><https://onlinelibrary.wiley.com/doi/10.1111/j.1365-2427.2006.01597.x>, 2006.
- Caissie, D. and Luce, C. H.: Quantifying streambed advection and conduction heat fluxes, *Water Resources Research*, 53, 1595–1624, <https://doi.org/10.1002/2016WR019813>, <https://onlinelibrary.wiley.com/doi/10.1002/2016WR019813>, 2017.
- 685 Caldwell, R. J., Gangopadhyay, S., Bountry, J., Lai, Y., and Elsner, M. M.: Statistical modeling of daily and subdaily stream temperatures: Application to the Methow River Basin, Washington, *Water Resources Research*, 49, 4346–4361, <https://doi.org/10.1002/wrcr.20353>, <http://doi.wiley.com/10.1002/wrcr.20353>, 2013.
- Chen, T. and Guestrin, C.: XGBoost: A scalable tree boosting system, *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 13-17-Aug, 785–794, <https://doi.org/10.1145/2939672.2939785>, 2016.
- 690 Chen, T., He, T., Benesty, M., Khotilovich, V., Tang, Y., Cho, H., Chen, K., Mitchell, R., Cano, I., Zhou, T., Li, M., Xie, J., Lin, M., Geng, Y., and Li, Y.: xgboost: Extreme Gradient Boosting, <https://cran.r-project.org/package=xgboost>, 2020.
- Chenard, J.-F. and Caissie, D.: Stream temperature modelling using artificial neural networks: application on Catamaran Brook, New Brunswick, Canada, *Hydrological Processes*, 22, 3361–3372, <https://doi.org/10.1002/hyp.6928>, <http://doi.wiley.com/10.1002/hyp.6928>, 2008.
- 695 Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y.: Learning phrase representations using RNN encoder-decoder for statistical machine translation, in: *EMNLP 2014 - 2014 Conference on Empirical Methods in Natural Language Processing*, *Proceedings of the Conference*, pp. 1724–1734, Association for Computational Linguistics (ACL), <https://doi.org/10.3115/v1/d14-1179>, 2014.
- Claesen, M. and De Moor, B.: Hyperparameter Search in Machine Learning, <http://arxiv.org/abs/1502.02127>, 2015.
- 700 Crank, J. and Nicolson, P.: A practical method for numerical evaluation of solutions of partial differential equations of the heat-conduction type, *Mathematical Proceedings of the Cambridge Philosophical Society*, 43, 50–67, <https://doi.org/10.1017/S0305004100023197>, 1947.
- Crisp, D. and Howson, G.: Effect of air temperature upon mean water temperature in streams in the north Pennines and English Lake District, *Freshwater Biology*, 12, 359–367, <https://doi.org/10.1111/j.1365-2427.1982.tb00629.x>, <http://doi.wiley.com/10.1111/j.1365-2427.1982.tb00629.x>, 1982.
- 705 Dallas, H.: Water temperature and riverine ecosystems: An overview of knowledge and approaches for assessing biotic responses, with special reference to South Africa, <https://doi.org/10.4314/wsa.v34i3.180634>, 2008.
- DeWeber, J. T. and Wagner, T.: A regional neural network ensemble for predicting mean daily river water temperature, *Journal of Hydrology*, 517, 187–200, <https://doi.org/10.1016/j.jhydrol.2014.05.035>, 2014.
- Dugdale, S. J., Hannah, D. M., and Malcolm, I. A.: River temperature modelling: A review of process-based approaches and future directions, *Earth-Science Reviews*, 175, 97–113, <https://doi.org/10.1016/j.earscirev.2017.10.009>, 2017.
- 710 Dunn, O. J.: Multiple Comparisons Using Rank Sums, *Technometrics*, 6, 241–252, <https://doi.org/10.1080/00401706.1964.10490181>, 1964.
- Feigl, M.: MoritzFeigl/wateRtemp: paper code, <https://doi.org/10.5281/ZENODO.4438575>, <http://doi.org/10.5281/zenodo.4438575>, 2021a.
- Feigl, M.: MoritzFeigl/ML_methods_for_stream_water_temperature_prediction: paper code, <https://doi.org/10.5281/ZENODO.4438582>, <http://doi.org/10.5281/zenodo.4438582>, 2021b.
- 715 Freund, Y. and Schapire, R. E.: A decision-theoretic generalization of on-line learning and an application to boosting, in: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 904, pp. 23–37, Springer Verlag, https://doi.org/10.1007/3-540-59119-2_166, https://link.springer.com/chapter/10.1007/3-540-59119-2_{_}166, 1995.



- Friberg, N., Dybkjær, J. B., Olafsson, J. S., Gislason, G. M., Larsen, S. E., and Lauridsen, T. L.: Relationships between structure and function in streams contrasting in temperature, *Freshwater Biology*, 54, 2051–2068, <https://doi.org/10.1111/j.1365-2427.2009.02234.x>,
720 <https://onlinelibrary.wiley.com/doi/full/10.1111/j.1365-2427.2009.02234.x><https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1365-2427.2009.02234.x><https://onlinelibrary.wiley.com/doi/10.1111/j.1365-2427.2009.02234.x>, 2009.
- Friedman, J. H.: Greedy function approximation: A gradient boosting machine, *Annals of Statistics*, 29, 1189–1232, <https://doi.org/10.1214/aos/1013203451>, 2001.
- Friedman, J. H.: Stochastic gradient boosting, *Computational Statistics and Data Analysis*, 38, 367–378, [https://doi.org/10.1016/S0167-9473\(01\)00065-2](https://doi.org/10.1016/S0167-9473(01)00065-2), 2002.
725
- Graf, R., Zhu, S., and Sivakumar, B.: Forecasting river water temperature time series using a wavelet–neural network hybrid modelling approach, *Journal of Hydrology*, 578, 124–115, <https://doi.org/10.1016/j.jhydrol.2019.124115>, <https://doi.org/10.1016/j.jhydrol.2019.124115>, 2019.
- Hadzima-Nyarko, M., Rabi, A., and Šperac, M.: Implementation of Artificial Neural Networks in Modeling the Water–Air Temperature Relationship of the River Drava, *Water Resources Management*, 28, 1379–1394, <https://doi.org/10.1007/s11269-014-0557-7>, 2014.
730
- Haiden, T., Kann, A., Wittmann, C., Pistotnik, G., Bica, B., and Gruber, C.: The integrated nowcasting through comprehensive analysis (INCA) system and its validation over the Eastern Alpine region, *Weather and Forecasting*, 26, 166–183, <https://doi.org/10.1175/2010WAF2222451.1>, <http://journals.ametsoc.org/doi/abs/10.1175/2010WAF2222451.1>, 2011.
- Haiden, T., Kann, A., and Pistotnik, G.: Nowcasting with INCA During SNOW-V10, *Pure and Applied Geophysics*, 171, 231–242, <https://doi.org/10.1007/s00024-012-0547-8>, <http://link.springer.com/10.1007/s00024-012-0547-8>, 2014.
735
- Hannah, D. M. and Garner, G.: River water temperature in the United Kingdom, *Progress in Physical Geography: Earth and Environment*, 39, 68–92, <https://doi.org/10.1177/0309133314550669>, <http://journals.sagepub.com/doi/10.1177/0309133314550669>, 2015.
- Hannah, D. M., Webb, B. W., and Nobilis, F.: River and stream temperature: dynamics, processes, models and implications, *Hydrological Processes*, 22, 899–901, <https://doi.org/10.1002/hyp.6997>, <http://doi.wiley.com/10.1002/hyp.6997>, 2008.
- 740 Hansen, L. K. and Salamon, P.: Neural Network Ensembles, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12, 993–1001, <https://doi.org/10.1109/34.58871>, 1990.
- He, J., Yang, K., Tang, W., Lu, H., Qin, J., Chen, Y., and Li, X.: The first high-resolution meteorological forcing dataset for land process studies over China, *Scientific Data*, 7, 25, <https://doi.org/10.1038/s41597-020-0369-y>, <https://doi.org/10.1038/s41597-020-0369-y>, 2020.
- Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D.,
745 Simmons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo, G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., De Chiara, G., Dahlgren, P., Dee, D., Diamantakis, M., Dragani, R., Flemming, J., Forbes, R., Fuentes, M., Geer, A., Haimberger, L., Healy, S., Hogan, R. J., Hólm, E., Janisková, M., Keeley, S., Laloyaux, P., Lopez, P., Lupu, C., Radnoti, G., de Rosnay, P., Rozum, I., Vamborg, F., Villaume, S., and Thépaut, J. N.: The ERA5 global reanalysis, *Quarterly Journal of the Royal Meteorological Society*, 146, 1999–2049, <https://doi.org/10.1002/qj.3803>, <https://rmets.onlinelibrary.wiley.com/doi/full/10.1002/qj.3803><https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.3803><https://rmets.onlinelibrary.wiley.com/doi/10.1002/qj.3803>, 2020.
750
- Hiebl, J. and Frei, C.: Daily temperature grids for Austria since 1961—concept, creation and applicability, *Theoretical and Applied Climatology*, 124, 161–178, <https://doi.org/10.1007/s00704-015-1411-4>, 2016.
- Hiebl, J. and Frei, C.: Daily precipitation grids for Austria since 1961—development and evaluation of a spatial dataset for hydroclimatic monitoring and modelling, *Theoretical and Applied Climatology*, 132, 327–345, <https://doi.org/10.1007/s00704-017-2093-x>, 2018.



- 755 Hinton, G. E., Srivastava, N., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. R.: Improving neural networks by preventing co-adaptation of feature detectors, <http://arxiv.org/abs/1207.0580>, 2012.
- Hochreiter, S. and Schmidhuber, J.: Long Short-Term Memory, *Neural Computation*, 9, 1735–1780, <https://doi.org/10.1162/neco.1997.9.8.1735>, <http://www.mitpressjournals.org/doi/10.1162/neco.1997.9.8.1735>, 1997.
- Hsu, C.-W., Chang, C.-C., and Lin, C.-J.: A Practical Guide to Support Vector Classification, Tech. rep., <http://www.csie.ntu.edu.tw/~cjlin>,
760 2003.
- Imholt, C., Gibbins, C. N., Malcolm, I. A., Langan, S., and Soulsby, C.: Influence of riparian cover on stream temperatures and the growth of the mayfly *Baetis rhodani* in an upland stream, *Aquatic Ecology*, 44, 669–678, <https://doi.org/10.1007/s10452-009-9305-0>, <https://link.springer.com/article/10.1007/s10452-009-9305-0>, 2010.
- Jackson, F. L., Fryer, R. J., Hannah, D. M., Millar, C. P., and Malcolm, I. A.: A spatio-temporal statistical model of maximum daily river
765 temperatures to inform the management of Scotland’s Atlantic salmon rivers under climate change, *Science of the Total Environment*, 612, 1543–1558, <https://doi.org/10.1016/j.scitotenv.2017.09.010>, 2018.
- Johnson, M. F., Wilby, R. L., and Toone, J. A.: Inferring air-water temperature relationships from river and catchment properties, *Hydrological Processes*, 28, 2912–2928, <https://doi.org/10.1002/hyp.9842>, 2014.
- Jones, D. R., Schonlau, M., and Welch, W. J.: Efficient Global Optimization of Expensive Black-Box Functions, *Journal of Global Optimiza-*
770 *tion*, 13, 455–492, <https://doi.org/10.1023/A:1008306431147>, <https://link.springer.com/article/10.1023/A:1008306431147>, 1998.
- Joslyn, K.: Water Quality Factor Prediction Using Supervised Machine Learning Water Quality Factor Prediction Using Supervised Machine Learning, REU Final Reports Research Experiences for Undergraduates on Computational Modeling Serving the City, <https://pdxscholar.library.pdx.edu/reu{ }reports>, 2018.
- Kennedy, J. and Eberhart, R.: Particle swarm optimization, in: Proceedings of ICNN’95 - International Conference on Neural Networks, vol. 4, pp. 1942–1948, IEEE, <https://doi.org/10.1109/ICNN.1995.488968>, <http://ieeexplore.ieee.org/document/488968/>, 1995.
775
- Klambauer, G., Unterthiner, T., Mayr, A., and Hochreiter, S.: Self-normalizing neural networks, in: *Advances in Neural Information Processing Systems*, vol. 2017-Decem, pp. 972–981, <http://arxiv.org/abs/1706.02515>, 2017.
- Kling, H., Stanzel, P., Fuchs, M., and Nachtnebel, H.-P.: Performance of the COSERO precipitation–runoff model under non-stationary conditions in basins with different climates, *Hydrological Sciences Journal*, 60, 1374–1393, <https://doi.org/10.1080/02626667.2014.959956>,
780 <http://www.tandfonline.com/doi/full/10.1080/02626667.2014.959956>, 2015.
- Kratzert, F., Klotz, D., Brenner, C., Schulz, K., and Herrnegger, M.: Rainfall – runoff modelling using Long Short-Term Memory (LSTM) networks, pp. 6005–6022, 2018.
- Kratzert, F., Klotz, D., Shalev, G., Klambauer, G., Hochreiter, S., and Nearing, G.: Towards Learning Universal, Regional, and Local Hydrological Behaviors via Machine-Learning Applied to Large-Sample Datasets, <http://arxiv.org/abs/1907.08456>, 2019.
- 785 Kruskal, W. H. and Wallis, W. A.: Use of Ranks in One-Criterion Variance Analysis, *Journal of the American Statistical Association*, 47, 583–621, <https://doi.org/10.1080/01621459.1952.10483441>, 1952.
- Kuhn, M.: caret: Classification and Regression Training, <https://cran.r-project.org/package=caret>, 2020.
- Kushner, H. J.: A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise, *Journal of Fluids Engineering*, *Transactions of the ASME*, 86, 97–106, <https://doi.org/10.1115/1.3653121>, 1964.
- 790 Laizé, C. L., Acreman, M. C., Schneider, C., Dunbar, M. J., Houghton-Carr, H. A., Flörke, M., and Hannah, D. M.: Projected flow alteration and ecological risk for pan-European rivers, *River Research and Applications*, 30, 299–314, <https://doi.org/10.1002/rra.2645>, <http://doi.wiley.com/10.1002/rra.2645>, 2014.



- Li, H., Deng, X., Kim, D.-Y., and Smith, E. P.: Modeling maximum daily temperature using a varying coefficient regression model, *Water Resources Research*, 50, 3073–3087, <https://doi.org/10.1002/2013WR014243>, <http://doi.wiley.com/10.1002/2013WR014243>, 2014.
- 795 Li, W., Kiaghadi, A., and Dawson, C.: High temporal resolution rainfall–runoff modeling using long-short-term-memory (LSTM) networks, *Neural Computing and Applications*, pp. 1–18, <https://doi.org/10.1007/s00521-020-05010-6>, <https://link.springer.com/article/10.1007/s00521-020-05010-6>, 2020.
- Lu, H. and Ma, X.: Hybrid decision tree-based machine learning models for short-term water quality prediction, *Chemosphere*, 249, 126–169, <https://doi.org/10.1016/j.chemosphere.2020.126169>, 2020.
- 800 Mackey, A. P. and Berrie, A. D.: The prediction of water temperatures in chalk streams from air temperatures, *Hydrobiologia*, 210, 183–189, <https://doi.org/10.1007/BF00034676>, 1991.
- McGlynn, B. L., McDonnell, J. J., Seibert, J., and Kendall, C.: Scale effects on headwater catchment runoff timing, flow sources, and groundwater-streamflow relations, *Water Resources Research*, 40, <https://doi.org/10.1029/2003WR002494>, <https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2003WR002494><https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2003WR002494><https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2003WR002494>, 2004.
- 805 McKenna, J. E., Butryn, R. S., and McDonald, R. P.: Summer Stream Water Temperature Models for Great Lakes Streams: New York, *Transactions of the American Fisheries Society*, 139, 1399–1414, <https://doi.org/10.1577/t09-153.1>, <https://www.tandfonline.com/doi/abs/10.1577/T09-153.1>, 2010.
- Močkus, J.: On Bayesian Methods for Seeking the Extremum, in: *Optimization Techniques IFIP Technical Conference*, pp. 400–404, https://doi.org/10.1007/978-3-662-38527-2_55, 1975.
- 810 Močkus, J.: *Bayesian Approach to Global Optimization*, vol. 37 of *Mathematics and Its Applications*, Springer Netherlands, Dordrecht, <https://doi.org/10.1007/978-94-009-0909-0>, 1989.
- Močkus, J., Tiesis, V., and Zilinskas, A.: The application of Bayesian methods for seeking the extremum, *Towards Global Optimisation*, https://doi.org/10.1007/978-94-009-0909-0_8, 1978.
- 815 Mohseni, O. and Stefan, H. G.: Stream temperature/air temperature relationship: A physical interpretation, *Journal of Hydrology*, 218, 128–141, [https://doi.org/10.1016/S0022-1694\(99\)00034-7](https://doi.org/10.1016/S0022-1694(99)00034-7), 1999.
- Naresh, A. and Rehana, S.: Modeling Stream Water Temperature using Regression Analysis with Air Temperature and Streamflow over Krishna River, *Rehana International Journal of Engineering Technology Science and Research*, 4, 2017.
- Neumann, D. W., Rajagopalan, B., and Zagona, E. A.: Regression model for daily maximum stream temperature, *Journal of Environmental Engineering*, 129, 667–674, [https://doi.org/10.1061/\(ASCE\)0733-9372\(2003\)129:7\(667\)](https://doi.org/10.1061/(ASCE)0733-9372(2003)129:7(667)), 2003.
- 820 Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Cournapeau, D., Passos, A., Brucher, M., Perrot, M., Perrot, M.: *Scikit-learn: Machine Learning in Python*, Tech. rep., <https://hal.inria.fr/hal-00650905v2>, 2011.
- Piccolroaz, S., Calamita, E., Majone, B., Gallice, A., Siviglia, A., and Toffolon, M.: Prediction of river water temperature: a comparison between a new family of hybrid models and statistical approaches, *Hydrological Processes*, 30, 3901–3917, <https://doi.org/10.1002/hyp.10913>, <http://doi.wiley.com/10.1002/hyp.10913>, 2016.
- 825 Piotrowski, A. P. and Napiorkowski, J. J.: Performance of the air2stream model that relates air and stream water temperatures depends on the calibration method, *Journal of Hydrology*, 561, 395–412, <https://doi.org/10.1016/j.jhydrol.2018.04.016>, 2018.
- Piotrowski, A. P. and Napiorkowski, J. J.: Simple modifications of the nonlinear regression stream temperature model for daily data, *Journal of Hydrology*, 572, 308–328, <https://doi.org/10.1016/j.jhydrol.2019.02.035>, 2019.
- 830



- Piotrowski, A. P., Napiorkowski, M. J., Napiorkowski, J. J., and Osuch, M.: Comparing various artificial neural network types for water temperature prediction in rivers, *Journal of Hydrology*, 529, 302–315, <https://doi.org/10.1016/j.jhydrol.2015.07.044>, 2015.
- Piotrowski, A. P., Napiorkowski, J. J., and Piotrowska, A. E.: Impact of deep learning-based dropout on shallow neural networks applied to stream temperature modelling, <https://doi.org/10.1016/j.earscirev.2019.103076>, 2020.
- 835 R Core Team: R: A Language and Environment for Statistical Computing, R Foundation for Statistical Computing, Vienna, Austria, <https://www.r-project.org/>, 2020.
- Rabi, A., Hadzima-Nyarko, M., and Šperac, M.: Modelling river temperature from air temperature: case of the River Drava (Croatia), *Hydrological Sciences Journal*, 60, 1490–1507, <https://doi.org/10.1080/02626667.2014.914215>, 2015.
- Razafimaharo, C., Krähenmann, S., Höpp, S., Rauthe, M., and Deutschländer, T.: New high-resolution gridded dataset of daily mean, minimum, and maximum temperature and relative humidity for Central Europe (HYRAS), *Theoretical and Applied Climatology*, 142, 1531–1553, <https://doi.org/10.1007/s00704-020-03388-w>, <https://doi.org/10.1007/s00704-020-03388-w>, 2020.
- 840 Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., and Prabhat: Deep learning and process understanding for data-driven Earth system science, *Nature*, 566, 195–204, <https://doi.org/10.1038/s41586-019-0912-1>, 2019.
- Risley, J. C., Roehl Jr., E. A., and Conrads, P. A.: Estimating Water Temperatures in Small Streams in Estimating Water Temperatures in Small Streams in Western Oregon Using Neural Network Models, Tech. rep., USGS Water- Resources Investigation Report 02-4218., <https://doi.org/10.3133/wri024218>, <http://pubs.er.usgs.gov/publication/wri024218>, 2003.
- 845 Rumelhart, D. E., Hinton, G. E., and Williams, R. J.: Learning representations by back-propagating errors, *Nature*, 323, 533–536, <https://doi.org/10.1038/323533a0>, 1986.
- Sahoo, G. B., Schladow, S. G., and Reuter, J. E.: Forecasting stream water temperature using regression analysis, artificial neural network, and chaotic non-linear dynamic models, *Journal of Hydrology*, 378, 325–342, <https://doi.org/10.1016/j.jhydrol.2009.09.037>, 2009.
- 850 Sand-Jensen, K. and Pedersen, N. L.: Differences in temperature, organic carbon and oxygen consumption among lowland streams, *Freshwater Biology*, 50, 1927–1937, <https://doi.org/10.1111/j.1365-2427.2005.01436.x>, <http://doi.wiley.com/10.1111/j.1365-2427.2005.01436.x>, 2005.
- Schapire, R. E.: The Strength of Weak Learnability, *Machine Learning*, 5, 197–227, <https://doi.org/10.1023/A:1022648800760>, 1990.
- 855 Segura, C., Caldwell, P., Sun, G., McNulty, S., and Zhang, Y.: A model to predict stream water temperature across the conterminous USA, *Hydrological Processes*, 29, 2178–2195, <https://doi.org/10.1002/hyp.10357>, 2015.
- Shank, D. B., Hoogenboom, G., and McClendon, R. W.: Dewpoint temperature prediction using artificial neural networks, *Journal of Applied Meteorology and Climatology*, 47, 1757–1769, <https://doi.org/10.1175/2007JAMC1693.1>, 2008.
- Smith, K.: The prediction of river water temperatures, *Hydrological Sciences Bulletin*, 26, 19–32, <https://doi.org/10.1080/02626668109490859>, 1981.
- 860 Snoek, J., Larochelle, H., and Adams, R. P.: Practical Bayesian optimization of machine learning algorithms, in: *Advances in Neural Information Processing Systems*, vol. 4, pp. 2951–2959, 2012.
- Sohrabi, M. M., Benjankar, R., Tonina, D., Wenger, S. J., and Isaak, D. J.: Estimation of daily stream water temperatures with a Bayesian regression approach, *Hydrological Processes*, 31, 1719–1733, <https://doi.org/10.1002/hyp.11139>, <http://doi.wiley.com/10.1002/hyp.11139>, 2017.
- 865 Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M.: Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design, *IEEE Transactions on Information Theory*, 58, 3250–3265, <https://doi.org/10.1109/TIT.2011.2182033>, <http://arxiv.org/abs/0912.3995><http://dx.doi.org/10.1109/TIT.2011.2182033>, 2009.



- Srivastava, N., Hinton, G., Krizhevsky, A., and Salakhutdinov, R.: Dropout: A Simple Way to Prevent Neural Networks from Overfitting, Tech. Rep. 56, <http://jmlr.org/papers/v15/srivastava14a.html>, 2014.
- 870 Stajkowski, S., Kumar, D., Samui, P., Bonakdari, H., and Gharabaghi, B.: Genetic-algorithm-optimized sequential model for water temperature prediction, *Sustainability (Switzerland)*, 12, <https://doi.org/10.3390/su12135374>, 2020.
- Stefan, H. G. and Preud'homme, E. B.: Stream temperature estimation from air temperature, *JAWRA Journal of the American Water Resources Association*, 29, 27–45, <https://doi.org/10.1111/j.1752-1688.1993.tb01502.x>, 1993.
- 875 Stevens, H., Ficke, J., and Smoot, G.: Techniques of water-resources investigations of the u.s. geological survey, U.S. Government printing office, p. 65, <http://water.usgs.gov/pubs/twri/twri4a3/>, 1975.
- Tavares, M. H., Cunha, A. H. F., Motta-Marques, D., Ruhoff, A. L., Fragoso, C. R., Munar, A. M., and Bonnet, M. P.: Derivation of consistent, continuous daily river temperature data series by combining remote sensing and water temperature models, *Remote Sensing of Environment*, 241, 111 721, <https://doi.org/10.1016/j.rse.2020.111721>, 2020.
- 880 Temizyurek, M. and Dadaser-Celik, F.: Modelling the effects of meteorological parameters on water temperature using artificial neural networks, *Water Science and Technology*, 77, 1724–1733, <https://doi.org/10.2166/wst.2018.058>, 2018.
- Thornton, P., Thornton, M., Mayer, B., Wilhelm, N., Wei, Y., Devarakonda, R., and Cook, R.: Daymet: Daily Surface Weather Data on a 1-km Grid for North America, Version 2. Data set., [https://daac.ornl.gov/cgi-bin/dsvviewer.pl?ds\[_\]id=1328http://daymet.ornl.gov](https://daac.ornl.gov/cgi-bin/dsvviewer.pl?ds[_]id=1328http://daymet.ornl.gov), 2014.
- Toffolon, M. and Piccolroaz, S.: A hybrid model for river water temperature as a function of air temperature and discharge, *Environmental Research Letters*, 10, 1–10, <https://doi.org/10.1088/1748-9326/10/11/114011>, 2015.
- 885 Trinh, N. X., Trinh, T. Q., Phan, T. P., Thanh, T. N., and Thanh, B. N.: Water Temperature Prediction Models in Northern Coastal Area, Vietnam, *Asian Review of Environmental and Earth Sciences*, 6, 1–8, <https://doi.org/10.20448/journal.506.2019.61.1.8>, 2019.
- Van Vliet, M. T., Franssen, W. H., Yearsley, J. R., Ludwig, F., Haddeland, I., Lettenmaier, D. P., and Kabat, P.: Global river discharge and water temperature under climate change, *Global Environmental Change*, 23, 450–464, <https://doi.org/10.1016/j.gloenvcha.2012.11.002>,
- 890 2013.
- Webb, B. W. and Zhang, Y.: Spatial and seasonal variability in the components of the river heat budget, *Hydrological Processes*, 11, 79–101, [https://doi.org/10.1002/\(sici\)1099-1085\(199701\)11:1<79::aid-hyp404>3.0.co;2-n](https://doi.org/10.1002/(sici)1099-1085(199701)11:1<79::aid-hyp404>3.0.co;2-n), 1997.
- Webb, B. W., Clack, P. D., and Walling, D. E.: Water-air temperature relationships in a Devon river system and the role of flow, *Hydrological Processes*, 17, 3069–3084, <https://doi.org/10.1002/hyp.1280>, 2003.
- 895 Webb, B. W., Hannah, D. M., Moore, R. D., Brown, L. E., and Nobilis, F.: Recent advances in stream and river temperature research, <https://doi.org/10.1002/hyp.6994>, 2008.
- Wehrly, K. E., Brenden, T. O., and Wang, L.: A Comparison of Statistical Approaches for Predicting Stream Temperatures Across Heterogeneous Landscapes, *JAWRA Journal of the American Water Resources Association*, 45, 986–997, <https://doi.org/10.1111/j.1752-1688.2009.00341.x>, <http://doi.wiley.com/10.1111/j.1752-1688.2009.00341.x>, 2009.
- 900 Wenger, S. J., Isaak, D. J., Dunham, J. B., Fausch, K. D., Luce, C. H., Neville, H. M., Rieman, B. E., Young, M. K., Nagel, D. E., Horan, D. L., and Chandler, G. L.: Role of climate and invasive species in structuring trout distributions in the interior Columbia River Basin, USA, *Canadian Journal of Fisheries and Aquatic Sciences*, 68, 988–1008, <https://doi.org/10.1139/f2011-034>, <https://cdnsiencepub.com/doi/abs/10.1139/f2011-034>, 2011.
- Werner, A. T., Schnorbus, M. A., Shrestha, R. R., Cannon, A. J., Zwiers, F. W., Dayon, G., and Anslow, F.: A long-term, temporally consistent, gridded daily meteorological dataset for northwestern North America, *Scientific Data*, 6, 1–16, <https://doi.org/10.1038/sdata.2018.299>, www.nature.com/sdata/, 2019.
- 905



- Westenbroek, S., Stewart, J. S., Buchwald, C. A., Mitro, M., Lyons, J. D., and Greb, S.: A Model for Evaluating Stream Temperature Response to Climate Change Scenarios in Wisconsin, pp. 1–12, American Society of Civil Engineers (ASCE), [https://doi.org/10.1061/41143\(394\)1](https://doi.org/10.1061/41143(394)1), 2010.
- 910 White, B. W. and Rosenblatt, F.: Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms, *The American Journal of Psychology*, 76, 705, <https://doi.org/10.2307/1419730>, 1963.
- Wickham, H.: *ggplot2: Elegant Graphics for Data Analysis*, Springer-Verlag New York, <https://ggplot2.tidyverse.org>, 2016.
- Xiang, Z., Yan, J., and Demir, I.: A Rainfall-Runoff Model With LSTM-Based Sequence-to-Sequence Learning, *Water Resources Research*, 56, <https://doi.org/10.1029/2019WR025326>, <https://onlinelibrary.wiley.com/doi/abs/10.1029/2019WR025326>, 2020.
- 915 Yang, D. and Peterson, A.: River water temperature in relation to local air temperature in the Mackenzie and Yukon basins, *Arctic*, 70, 47–58, <https://doi.org/10.14430/arctic4627>, 2017.
- Zhilinskas, A. G.: Single-step Bayesian search method for an extremum of functions of a single variable, *Cybernetics*, 11, 160–166, <https://doi.org/10.1007/BF01069961>, 1975.
- Zhu, S. and Piotrowski, A. P.: River/stream water temperature forecasting using artificial intelligence models: a systematic review, <https://doi.org/10.1007/s11600-020-00480-7>, <https://doi.org/10.1007/s11600-020-00480-7>, 2020.
- 920 Zhu, S., Nyarko, E. K., and Hadzima-Nyarko, M.: Modelling daily water temperature from air temperature for the Missouri River, *PeerJ*, 2018, e4894, <https://doi.org/10.7717/peerj.4894>, 2018.
- Zhu, S., Hadzima-Nyarko, M., Gao, A., Wang, F., Wu, J., and Wu, S.: Two hybrid data-driven models for modeling water-air temperature relationship in rivers, *Environmental Science and Pollution Research*, 26, 12 622–12 630, <https://doi.org/10.1007/s11356-019-04716-y>,
925 2019a.
- Zhu, S., Heddam, S., Nyarko, E. K., Hadzima-Nyarko, M., Piccolroaz, S., and Wu, S.: Modeling daily water temperature for rivers: comparison between adaptive neuro-fuzzy inference systems and artificial neural networks models, *Environmental Science and Pollution Research*, 26, 402–420, <https://doi.org/10.1007/s11356-018-3650-2>, 2019b.
- Zhu, S., Heddam, S., Wu, S., Dai, J., and Jia, B.: Extreme learning machine-based prediction of daily water temperature for rivers, *Environmental Earth Sciences*, 78, 1–17, <https://doi.org/10.1007/s12665-019-8202-7>, 2019c.
- 930 Zhu, S., Nyarko, E. K., Hadzima-Nyarko, M., Heddam, S., and Wu, S.: Assessing the performance of a suite of machine learning models for daily river water temperature prediction, *PeerJ*, 7, e7065, <https://doi.org/10.7717/peerj.7065>, 2019d.