Multi-task deep learning of daily streamflow and water temperature

J. M. Sadler[1], A. P. Appling[2], J. S. Read[3], S. K. Oliver[4], X. Jia[5], J. A. Zwart[6], V. Kumar[7]

[1]U.S. Geological Survey, jsadler@usgs.gov
[2]U.S. Geological Survey, aappling@usgs.gov
[3]U.S. Geological Survey, jread@usgs.gov
[4]U.S. Geological Survey, soliver@usgs.gov
[5]U.S. Geological Survey, XIAOWEI@pitt.edu
[6]University of Pittsburgh, jzwart@usgs.gov
[7]University of Minnesota, kumar001@umn.edu

Corresponding author: Jeffrey Sadler (jsadler@usgs.gov)

Key Points:
- A single deep learning model is used to predict both water temperature and streamflow
- The multi-task approach improves overall performance for most sites tested
- A naïve use of multi-task learning is detrimental to water temperature predictions

Abstract
Deep learning models can accurately predict many hydrologic variables including streamflow and water temperature; however, these models have typically predicted hydrologic variables independently. This study explored the benefits of modeling two interdependent variables, daily average streamflow and daily average stream water temperature, together using multi-task deep learning. A multi-task scaling factor controlled the relative contribution of the auxiliary variable's error to the overall loss during training. Our experiments examined the improvement in prediction accuracy of the multi-task approach using paired streamflow and temperature data from sites across the conterminous United States. Our results showed that the best performing multi-task models performed better overall than the single-task models in terms of Nash-Sutcliffe efficiency. The improvement of the multi-task models relative to the single-task models had a seasonal trend with the multi-task models making larger improvements in the high-flow seasons. The multi-task scaling factor was consequential in determining to what extent the multi-task approach was beneficial, and a naïve selection of this factor led to worse-performing multi-task models for stream temperature. Our findings indicate that, when configured properly, a multi-task approach could make more accurate predictions of interdependent hydrologic variables.

1   Introduction

Accurately predicting the quality and quantity of water in streams is extremely consequential (Hamlet et al., 2002). Among the many hydrologic variables, water temperature and streamflow are among the most important. Water temperature is an ecological master factor that affects fish abundance, the biogeochemical fate of in-stream nutrients and contaminants, and the emergence of harmful algal blooms (HABs) (Barnett, 1972; Zeng et al., 2009). Streamflow is another critical factor in stream ecology affecting water quality and stream ecology (Bain et al., 1988; Carlisle et al., 2011). Streamflow forecasts are also important for understanding and forecasting the availability of water for human use, and flood and drought forecasting (Hamlet et al., 2002; Pappenberger et al., 2015; Robertson et al., 2013). Accurate streamflow and water temperature predictions can inform management decisions that have large economic and environmental effects such as: how much water should be released from a reservoir? how much water is expected to be available for agriculture? will a stream become too warm for fish to thrive or even survive? (Alemu et al., 2011).

Process-based models and simple statistical models have been used to make predictions of hydrologic variables including streamflow and water temperature for many years (Campolo et al., 1999); however, deep learning (DL) models have recently been shown to be a powerful and accurate alternative (Bai et al., 2016; Besaw et al., 2010; P. A. Chen et al., 2013; Kratzert et al., 2018; Kratzert, Klotz, Herrnegger, et al., 2019; Kratzert, Klotz, Shalev, et al., 2019; Read et al., 2019; Shen, 2018; Shen et al., 2018; Stammler et al., 2013) advancing beyond simpler machine learning (ML) methods (Hsu et al., 1995). DL models can model complex, non-linear spatial and temporal dependencies in hydrologic processes (Shen, 2018) and have even shown substantial predictive skill with out-of-sample predictions (Kratzert, Klotz, Herrnegger, et al., 2019; Read et al., 2019). Furthermore, like other ML models, once DL models are trained they run more quickly than their detailed, process-based counterparts, a characteristic that can be important for management timelines (Soleimani et al., 2019; Zahura et al., 2020).

One of the distinct properties of DL models (as with all statistical and ML models) is that an out-of-the-box DL model has no awareness of the physical laws that govern the hydrologic processes

2

being modeled. Human knowledge of these laws gained over years of scientific study, although limited, can be used to guide DL models resulting in improved predictive performance. DL models are especially suited for the injection of process understanding because of their flexibility. The use of human knowledge and understanding of the underlying physical, chemical, or other processes to guide DL models is a promising and growing field of research (Willard et al., 2020). The implementation of knowledge-guided DL (also known as process-guided DL) can take many forms. For example, Read et al. (2019) penalized a DL lake-temperature model for making predictions that violated energy conservation and used a process-based model to pretrain the DL model. Jia et al. (2020) incorporated spatial connections in modeling a network of stream segments so that, as happens in reality, the states and fluxes at given stream segment inform the predictions of the states and fluxes of its neighboring segments.

In this paper, we suggest another way to leverage human understanding of hydrologic processes to guide our DL models: multi-task learning of related hydrologic variables. In multi-task learning, a single DL model is trained to predict two related variables. Multi-task models have been shown to better generalize and reduce overfitting (Ruder, 2017) and have proven effective in several applications including natural language processing (Chen et al., 2014; Seltzer & Droppo, 2013) and in computer vision (Girshick, 2015). In the hydrologic sciences, there are many variables that are related to each other, and a multi-task model could learn and use those relationships to make better predictions. Most of the models in the literature, to this point, have been trained to perform only one task or, in other words, predict only one variable (e.g., streamflow (Hu et al., 2018; Kratzert, Klotz, Shalev, et al., 2019), water level (Bowes et al., 2019), or water temperature (Read et al., 2019)). By selecting appropriate variables for multi-task learning based on our understanding of the underlying processes, we could guide a DL model to better represent the hydrologic process and therefore better predict the hydrologic variable that we are interested in. In this paper, we will focus on the variables of streamflow and water temperature.

In addition to a means of incorporating process understanding, multi-task modeling may also be a way to address, at least in part, another limitation of DL models: the need for large observational datasets with which a model can be trained. In the earth sciences, these data can be sparse given the expense with which many observations are obtained. By implementing multi-task learning we may be able to leverage the information from one variable to help improve the prediction of another related variable when it is less observed. An imbalance in amounts of observations can happen in cases where one variable is more expensive to observe compared to another. For example, long-term stream discharge compared to water temperature.

Some studies have leveraged the relationship between hydrologic variables using data-driven models. Rahmani et al. (2020) and VanVliet et al. (2011) both found that including streamflow as an input for a ML model and statistical model (respectively) for temperature prediction improved model performance. Our contribution relative to these is that in multi-task learning streamflow is not used as a model input but as an output along with water temperature. Furthermore, we examine the benefit of multi-task learning for improving streamflow where the Rahmani et al. (2020) and Van Vliet et al. (2011) focused on improvements in water temperature predictions. Kraft et al. (2020) used a multi-task model to predict different parts of the water

balance equation including evapotranspiration, snow water equivalent, and groundwater recharge but did not formally assess the benefits of the multi-task approach.

Although our focus is a multi-task approach (one model predicts both variables), there are other ways of leveraging the information of one variable to benefit the other. For example, Rahmani et. al (2020) used the auxiliary variable (streamflow in their case) as an input variable to predict their main variable (water temperature). This approach is less heavy-handed than multi-task learning; when used as an input variable, the model parameters would only use the auxiliary variable it if it is helpful. However, this applies in the reverse as well. When used as an input variable, there is no guarantee that the auxiliary variable will be influential in making the final predictions. Multi-task learning ensures that the auxiliary variable will make an impact on the main variable predictions, though, as our results showed, that impact can be detrimental depending on the statistical relationship between the two variables and the multi-task scaling factor. Furthermore, one limitation of using the auxiliary information as an input variable is that it requires an observation of that variable for every time step for both the training and testing time periods. The multi-task method does not have this assumption.

To test the value of multi-task learning for streamflow and water temperature, we evaluated whether training a single model to predict the two variables together improved upon our predictions from a single model trained to predict each variable alone. We selected streamflow and stream temperature, in part, because changes in streamflow and temperature are part of the energy transfer process. Most generally, heat is transferred to or from a stream segment via advective and non-advective processes. Advective heat transfer manifests in changes to both streamflow and stream temperature. Changes in streamflow can change water temperature and, similarly, changes in water temperature can indicate changes in streamflow. A model trained using streamflow and stream temperature observations has the potential to better learn about the advective energy transfer process and, therefore, improve predictions of both variables.

In this paper we evaluate model performance gain for simultaneously predicting streamflow and water temperature compared to streamflow or water temperature alone using sites distributed across the conterminous United States. Additionally, we test whether multi-task performance shows improved predictions under data sparse conditions for streamflow, which is common given the relatively high cost of streamflow observations. The results of our experiments shed light on the potential of multi-task learning as a means of using process knowledge and making use of available data to improve streamflow and stream temperature predictions.

2   Materials and Methods
2.1   Data
The data that we used for our experiments came from the CAMELS dataset (Addor et al., 2017; Newman et al., 2015). For each site in the CAMELS dataset, there are nearly complete daily observations of streamflow and daily basin-averaged weather forcing data from 01-01-1980 through 12-31-2014. Of the 671 total sites in the CAMELS dataset, 109 had at least one water temperature observation collected by the U.S. Geological Survey (USGS) at the site in the time period of 1980-2014. These basins ranged from 5.4 to 14,269 square kilometers in area. We used all sites that had any water temperature data during the training period of our experiments. The

amount of water temperature daily observations ranged from 2 to 12,515 with a median of 779 observations. The sites were distributed fairly evenly across the conterminous United States except for a cluster of sites in the Mid-Atlantic. The input data we used were the seven basin-averaged Daymet forcing variables (also known as features in DL terminology) from CAMELS (Table 1). All data were at a daily timestep.

Table 1. Input features for single-task and multi-task DL models

Input variable	Description
dayl (s)	Day light in seconds
prcp (mm/day)	Precipitation in millimeters per day
srad (W/m2)	Short-wave radiation
swe (mm)	Snow water equivalent
tmax (C)	Maximum air temperature
tmin (C)	Minimum air temperature
vp (Pa)	Vapor pressure

2.2   The DL model

The DL model we used for our experiments consisted of three main parts: a Long Short-term memory (LSTM) layer (Hochreiter & Schmidhuber, 1997) and two parallel, densely-connected output layers (one for streamflow output and one for water temperature output) (see Figure 1). Our model was implemented using the TensorFlow package (Martin et al., 2015). The LSTM is a type of recurrent neural network (RNN). RNNs have explicit representations of sequences where the information of one member of the sequence (i.e., one time step) is part of the input for the next sequence member (i.e., the next time step). This is especially important for streamflow and water temperature where the temperature and streamflow of one time step have a strong correlation with the temperature and streamflow in the next time step. Compared to a standard RNN, LSTMs are even better suited to model time series of hydrologic variables because LSTM's memory cells can store long-term temporal dependencies such as snow-pack (Kratzert et al., 2018). The "memory" of an LSTM model is stored in two state vectors: the hidden state, *h,* and the carry-over state, *c.*
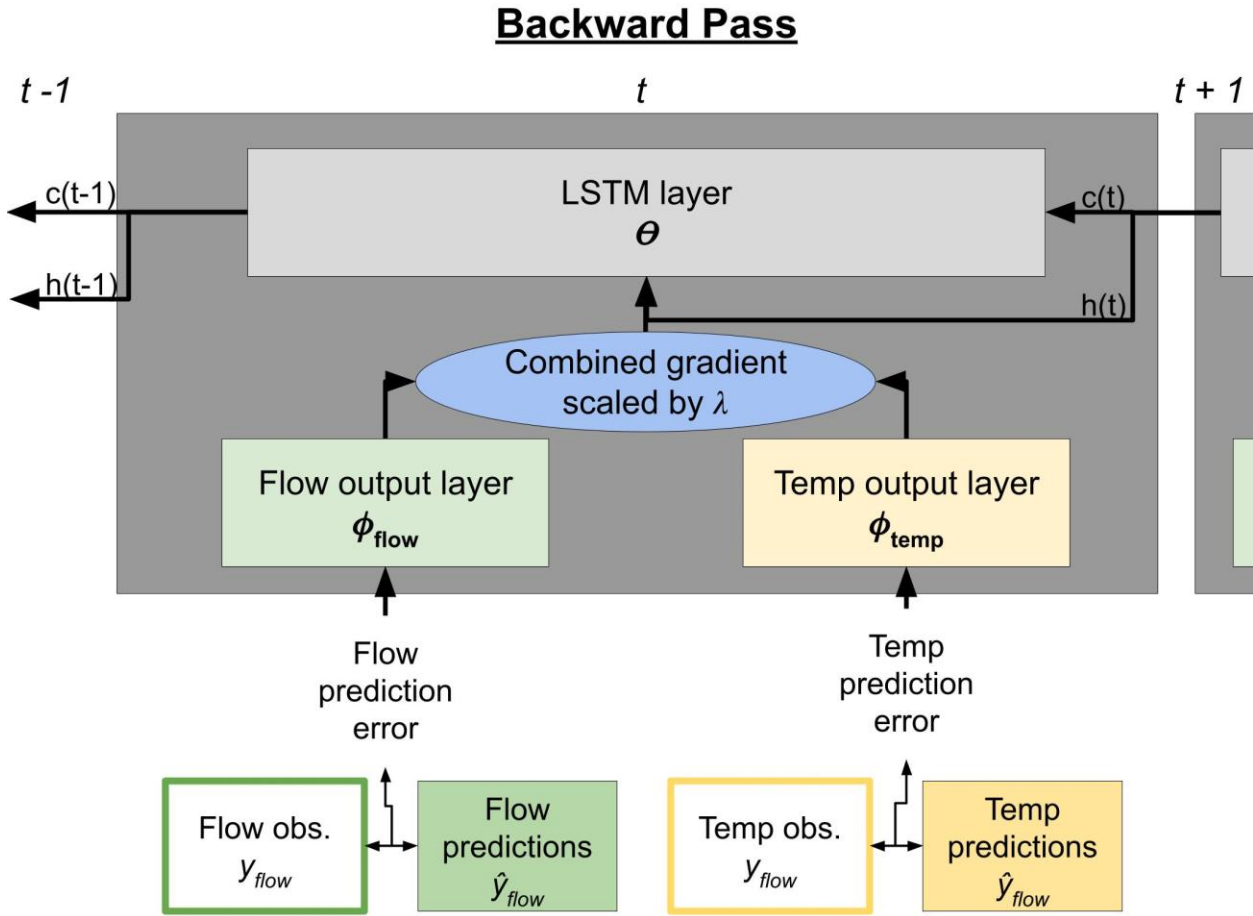
Backward Pass

Figure 1. Conceptual figure of the backward pass for a multi-task model. h *and* c *are the hidden and carry-over state of the LSTM, respectively (Hochreiter & Schmidhuber, 1997), and* t *is the timestep in the sequence being modeled.*

The LSTM layer in our model takes sequences of weather forcing inputs (see Table 1) and converts these into intermediate output sequences. The output from the LSTM layer at each timestep is passed as input to the two output layers. The two output layers produce the predictions: one output layer predicts streamflow and the other output layer predicts water temperature. Both output layers are given the same intermediate values from the LSTM layer. Although one model could be trained with multiple sites' information, we trained separate models for each site. The model's parameters are weights and biases that are adjusted in model training. The model hyperparameters are architectural and training process decisions that are typically tuned manually.

The difference between the single-task and multi-task version of the model is seen when the model parameters are being updated (or trained) based on the prediction errors. Passing the errors through the model to update the parameters is also known as a "backward pass" as opposed to the "forward pass" in which the model is given inputs and predictions are made. The equations for updating the DL model parameters for the output layers are

$$\phi_{main}(i+1) \coloneqq \phi_{main}(i) - \alpha \nabla_{\phi_{main}} \mathcal{L}_{main}\big(\theta(i), \phi_{main}(i)\big) \qquad (1)$$

$$\phi_{aux}(i+1) \coloneqq \phi_{aux}(i) - \alpha \nabla_{\phi_{aux}} \mathcal{L}_{aux}\big(\theta(i), \phi_{aux}(i)\big) \qquad (2)$$

where ϕ_{main} is the parameters for the main variable output layer, ϕ_{aux} is the parameters for the auxiliary variable output layer, $\mathcal{L}$ is the loss in either the main or the auxiliary variable, ∇ is the corresponding gradient for a given set of parameters, α is the learning rate hyperparameter, and i is the training step. These equations are written generically. If streamflow were the "main" variable, the ϕ_{main} would be replaced with ϕ_{flow}. The same would be done with the gradient and loss terms. Furthermore, the "aux" terms would be for the stream temperature variable. If stream temperature were the "main" variable, streamflow would be the "aux" variable. Our experiments tested multi-task modeling when streamflow was the main variable and stream temperature the auxiliary variable and vice versa.

In the backward pass, the gradient with respect to the loss in only the relevant variable is used to update the parameters of the respective output layer (Equations 1 and 2). For example, only errors in the streamflow predictions will be used to update the weights in the streamflow output layer.

The loss function was based on the Nash Sutcliffe Efficiency (NSE), a common metric for measuring streamflow prediction performance (Nash & Sutcliffe, 1970) :

$$NSE = 1 - \frac{\sum_{t=1}^{T}\big(\hat{Y}^t(X^t, \theta, \phi) - Y^t\big)^2}{\sum_{t=1}^{T}(Y^t - \bar{Y})^2} \qquad (3)$$

where $\hat{Y}$ is the predicted value, which is a function of the inputs (X) and the model parameters θ and either ϕ_{main} for the main variable or ϕ_{aux} for the auxiliary variable. Y is the observed value, t is a time step of the time series, and T is the full time series length. The unmodified NSE values range from negative infinity to one, with one being a perfect model. We used the normalized Nash Sutcliffe Efficiency (NNSE) which ranges from 0 to 1:

$$NNSE = \frac{1}{2 - NSE} \qquad (4)$$

Because the model training algorithm was minimizing the loss, our loss function was 1 minus the NNSE so that a smaller loss would correspond to a better model and vice versa:

$$\mathcal{L} = 1 - NNSE \qquad (5)$$

The equation to update the parameters of the shared LSTM layer is

$$\theta(i + 1) := \theta(i) - \alpha\left[\nabla_\theta \mathcal{L}_{main}(\theta(i), \phi_{main}(i))) + \lambda\nabla_\theta \mathcal{L}_{aux}\left(\theta(i), \phi_{aux}(i)\right)\right] \qquad (6)$$

where θ is the parameters for the LSTM layer, α is the learning rate (a hyperparameter), and λ is the hyperparameter that scales the auxiliary gradient. The gradients with respect to the loss from both variables are combined and used to update the parameters of the shared LSTM layer. When the main and auxiliary gradients are combined, λ weights the influence of the auxiliary gradient relative to the main gradient. Therefore, if $\lambda = 0$, the auxiliary gradient has no effect on the update of the LSTM parameters. In this case, the LSTM parameters are adjusted in proportion to the errors in the main variable alone, so the model is effectively a single-task model because the predictions of the auxiliary variable are ignored. If, on the other hand, $\lambda > 0$, the LSTM parameters are adjusted based on errors in both the main and auxiliary variables and the model is a multi-task model. For the remainder of the paper we refer to λ as the "multi-task scaling factor."

2.3   Experiments
In our experiments we compared a single-task model with multi-task models using a variety of multi-task scaling factors (see Equation 6 for the meaning of the multi-task scaling factor, λ). To test the ability of a multi-task model to improve streamflow and water temperature predictions, we ran three experiments.

The available data were split into three sets: a training set, a validation set, and a test set. The training data were used in updating the model parameters. The validation data were used for selecting hyperparameters. The only hyperparameter adjusted was the multi-task learning factor. The test set was used for final model evaluation and was not used to adjust any part of the model.

The amount of data used for training varied in each experiment (see Figure 2). For each of the experiments we trained individual models for each site in our dataset. All of our experiments were scripted using the Snakemake Python library (Köster & Rahmann, 2012).

2.3.1   Experiment A: Multi-task modeling for streamflow with 25 years of training data
In Experiment A we evaluated the multi-task model performances for predicting streamflow using 25 years of training data (1990-2014). In this experiment, stream temperature was the auxiliary variable. We used 1980-1985 as a validation set and 1985-1990 as a test set. We used data from the latter part of the dataset (1990-2014) for training because there is substantially more water temperature data available to train the multi-task model in later years compared to earlier years. This experiment was intended to represent predicting streamflow at a well-monitored location. There were 103 sites that had data in all three partitions for this experiment.

Given that there is a process-based relationship between streamflow and water temperature, it follows that the relationship between the two variables at a given site is expected to affect the performance of the multi-task approach at that site. Because there are catchment characteristics that influence the relationship between streamflow and water temperature, we hypothesized that we could use these attributes to predict and/or understand how beneficial a multi-task model

8

would be. To explore the predictability of the improvement of a multi-task model compared to a single-task model, we used a Random Forest model (Breiman, 2001) to predict the improvement of multi-task models compared to single-task models across the 103 sites. We used the catchment characteristics from the CAMELS dataset as input features. The details of our implementation of the Random Forest model are in the supplementary materials.

2.3.2   Experiment B: Multi-task modeling for streamflow with 2 years of training data
In Experiment B, we trained the model with only 2 years of streamflow data. In this experiment, stream temperature was the auxiliary variable. We used different 2-year periods to account for differences in events and flow regimes that may or may not be represented in one 2-year period compared to another. The data that we used to train the model in Experiment B were each non-overlapping 2-year period in the years 1990-2014 (i.e., Jan. 01 1990 to Jan. 01 1992, Jan. 01 1992 to Jan. 01 1994, Jan. 01 1994 to Jan. 01 1996, etc.). The validation and test sets were the same as in Experiment B. This experiment was designed to simulate the situation where only a small period of record is available for a site. Such situations are common, given the cost of maintaining a streamflow gage. There were 98 sites that had data in all three partitions for this experiment.

2.3.3   Experiment C: Multi-task modeling for temperature
In Experiment C, we evaluated the performance of the multi-task models for improving stream temperature predictions. In this experiment, streamflow was the auxiliary variable. Because there were fewer temperature observations available at these sites compared to streamflow observations, especially in the early years of the data record, we used the years 1980-2012 as the training period, 2013 as the validation period, and 2014 as the testing period. Temperature observations do not vary substantially year to year; therefore, we assumed that the performance in just one year of testing data is sufficiently representative of the model performance. There were 53 sites that had data in all three partitions for this experiment.
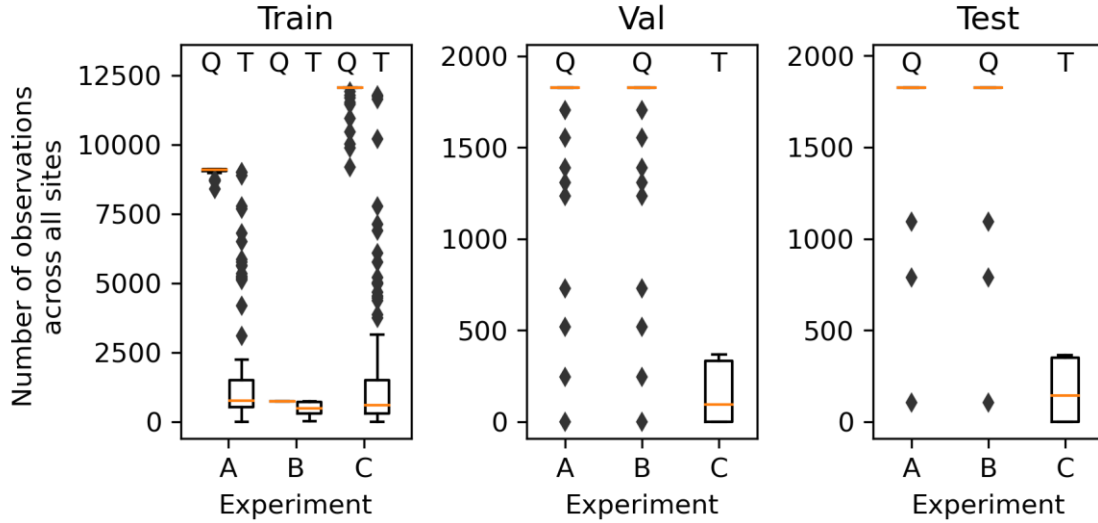
Figure 2. Boxplots showing number of observations for each experiment, variable, and data partition. The bottom of each box is the first quartile of the data (Q1), the orange bar is the median of the data, and the top of each box is the third quartile (Q3). The lower whisker is the lowest point within Q1 – 1.5*IQR, where IQR is the innerquartile range (Q3-Q1). The upper whisker is the highest point within Q3 + 1.5*IQR. The diamonds are points that lie outside of the whiskers.

2.4 Model configuration, training, and evaluation metrics

The model had 20 hidden units in the LSTM layer. The sequence length for training was 365 days. Having a 365-day sequence length allows the LSTM to build up states in its memory cells which is important because stream temperature and streamflow have long temporal dependencies that can extend back into previous seasons. Each 365-day sequence constituted a training example. The models were trained given the training examples in parallel as one batch. In each training epoch, the average of the gradients from each training example (a year-long period) was used to update the model weights. One training epoch is completed when the model parameters are adjusted based on the data from each training example one time.

Each model was trained for 400 epochs for Experiments A and C and 200 epochs for Experiment B. We trained the models for fewer epochs in Experiment B to reduce overfitting that was more likely with a much smaller dataset (using 2 years of training data instead of 25). For all the models, we used a dropout rate of 0.15 to reduce overfitting. The optimizer used for the model training was the Adam optimizer (Kingma & Ba, 2015). In each experiment, 30 replicates of each model were trained, each with different random initializations of the model parameters.

Both the weather input features (see Table 1) and target data (streamflow and stream temperature) were all scaled and centered by subtracting the mean from each variable and dividing by the standard deviation; in this way each scaled and centered variable had a mean of zero and a standard deviation of one. It was necessary to scale the target data as well as the input data so that both output streamflow and temperature predictions would be on the same scale in

10

model training. All the data were scaled and centered using the standard deviation and means from the training data only.

We compared the multi-task models and the single-task models using NSE. The comparisons were done on an overall level and on a monthly level using the testing data. Statistical comparisons were done on a site-wise basis using a one-sided Welch's t-test (Satterthwaite, 1946). In the t-test, our hypothesis was that, for a given site, a multi-task model is more accurate in terms of NSE than the single-task model. To calculate t in the Welch's t-test for a given site, n, the equation is

$$t_n = \frac{\overline{NSE}_{n,single_{task}} - \overline{NSE}_{n,multi_{task}}}{\sqrt{\dfrac{s^2_{n,multi_{task}} + s^2_{n,single_{task}}}{N}}} \tag{7}$$

where $\overline{NSE}$ is the NSE averaged across the 30 replicates, s the standard deviation, and N is the number of replicates. The t-score and p-value were calculated using Python's SciPy package (Virtanen et al., 2020).

3   Results

3.1   Experiment A: Multi-task modeling for streamflow with 25 years of training data

3.1.1   Influence of multi-task scaling factor, λ

For most of the model configurations, the multi-task model performed better than the single-task model at predicting streamflow (see Figure 3a and 3b). The multi-task scaling factor (λ from Equation 4) had a large influence on the accuracy of the multi-task models. As the multi-task scaling factor increased from 0 to 80, the performance of the model generally increased. In other words, as the emphasis that the model gave to predicting stream temperature grew, the accuracy of the streamflow predictions increased. However, when multi-task scaling factors exceeded 80, overall model performances in terms of median NSE decreased with further increases in the scaling factor.

The multi-task approach made significant improvements to the worst performing single-task models. This is seen when comparing the improvements to the mean NSEs (Figure 3b) to the median NSEs (Figure 3a) in the validation set. The minimum NSE for the single-task models was -11.6 whereas the minimum NSE for the multi-task model with a multi-task scaling factor of 80 was -5.3. The improvements to the worst performing models made the improvement of the mean NSEs much more pronounced than the improvements to the median NSEs. Furthermore, although the median NSE is lower than the single-task median when the multi-task scaling factor is 320, the mean NSE is much higher than the single-task's mean (although still lower than when the multi-task scaling factor peaks for mean performance at 160).

Although there was a clear overall pattern in the effect of increasing the multi-task scaling factor on model performance, there was also significant variability site to site (see Figure 3b). The multi-task scaling factor that performed best in terms of median NSE for the largest number of sites in the validation period was 80. However, for most of the sites, performance was best with a multi-task scaling factor other than 80. For each multi-task scaling factor we evaluated there

were at least six sites at which that multi-task scaling factor performed the best. For example, for six sites, the best-performing multi-task scaling factor was 10. For eight sites, the best multi-task scaling factor was 0, meaning that the multi-task approach did not improve overall streamflow predictions at those sites compared to the single-task model. Examining the geography and catchment characteristics of these sites showed no obvious reason for which the single-task models performed best at these sites.
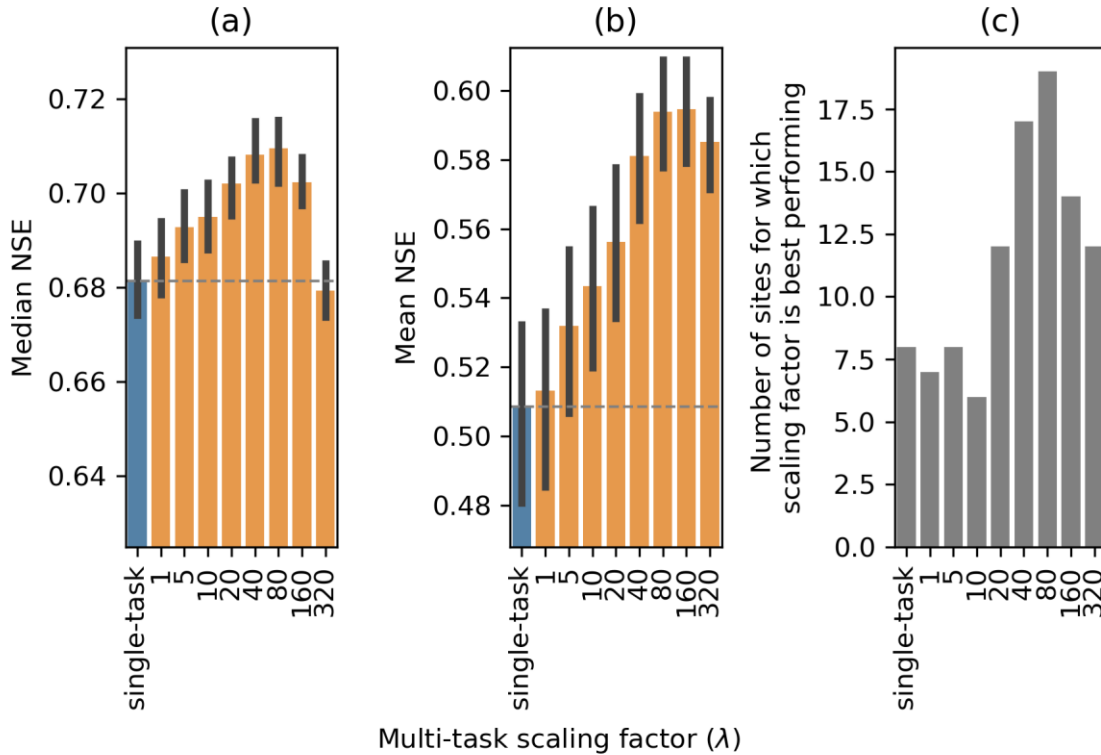
Figure 3. Influence of mutli-task scaling factor on streamflow prediction accuracy. (a) Performance of single-task and multi-tasks models with various values of multi-task scaling factors in terms of median NSE. (b) Performance of single-task and multi-tasks models with various values of multi-task scaling factors in terms of mean NSE. (c) The number of sites for which each multi-task scaling factor performed the best in terms of median NSE. Results shown for all panels are from the validation set. The dotted horizontal line in (a) and (b) is the median performance of the single task model. The error bars in (a) and (b) represent 95% confidence intervals.

Given the variability in the performance of the several multi-task scaling factors in the validation period across the 103 sites, we selected the best-performing, non-zero multi-task scaling factor for each site to evaluate in the test set. For the eight sites for which the best performing model was the single-task model, we used the second-best multi-task scaling factor. The best multi-task models performed better overall than the single-task models in the test set (see Figure 4). The median NSE of the best multi-task models in the test set was 0.66 compared to 0.61 for the single-task models.

With a median NSE of 0.61, the performance of our single-site models is comparable to that of Kratzert et al. (2019) whose single-site LSTM models that predicted runoff at CAMELS sites had a median of 0.59. The spatial distribution of the model performance was also similar to Kratzert et al. (2019) with the models in North Dakota performing the worst (see Figure 5). There was little spatial trend in the improvements made by multi-task models (with a multi-task factor of 80). The most and least improved sites in terms of t-score were two sites that were very near to each other in Texas.

For each multi-task scaling factor except for 320, the performance was better than the single-task models for more than half of the sites (see Figure 6), though the sites that improved for one multi-task scaling factor were not necessarily the same ones that improved for another. When the multi-task scaling factor was 20, 40, and 80 and when the best site-specific multi-task scaling factors were used, a majority of sites had statistically significant improvements (see Table 2). A larger multi-task scaling factor resulted in larger improvements at some sites and a larger number of significantly improved sites up until a factor of 160. A larger multi-task scaling factor also resulted in a larger magnitude decrease in performance for some sites. The best multi-task models performed the same in terms of median RMSE compared to when all the sites used the multi-task scaling factor of 80; however, with a multi-task scaling factor of 80, the majority of sites (57 sites) had significant improvements and 17 sites had significantly worse performance.
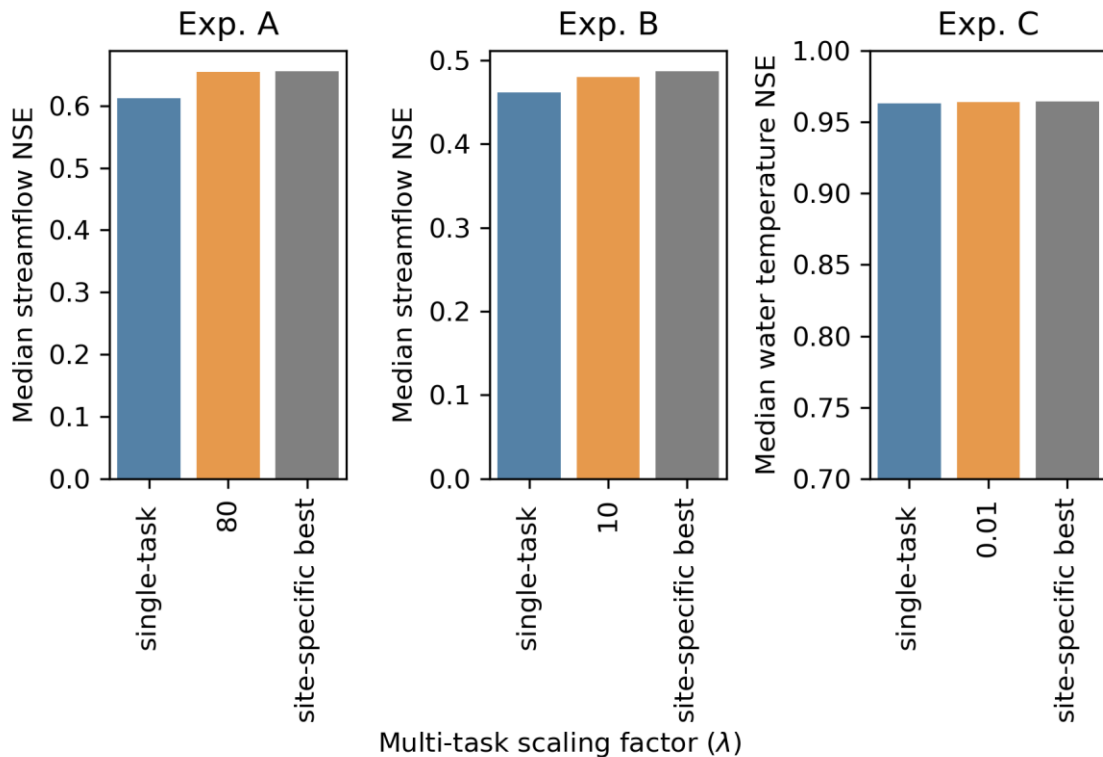

Figure 4. Cumulative distribution function of NSEs and median NSEs of the single-task models, the multi-task models with the scaling factor that performed best for the most sites (80), and the multi-task models with the best site-specific multi-task scaling factor. These results are for the test set.
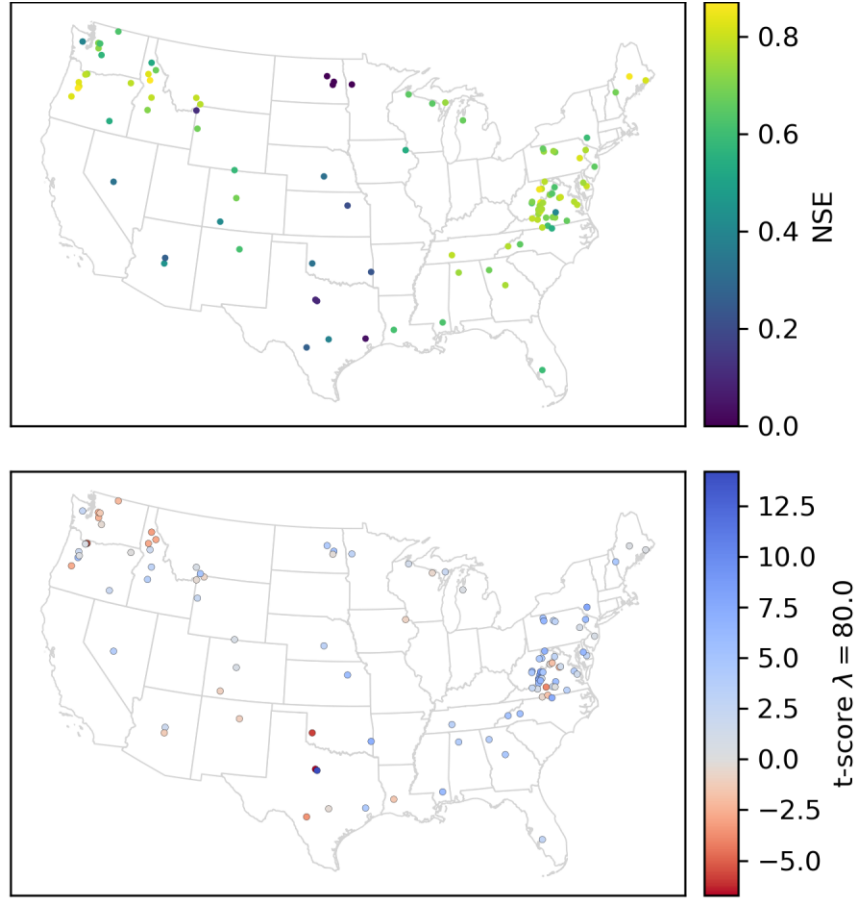
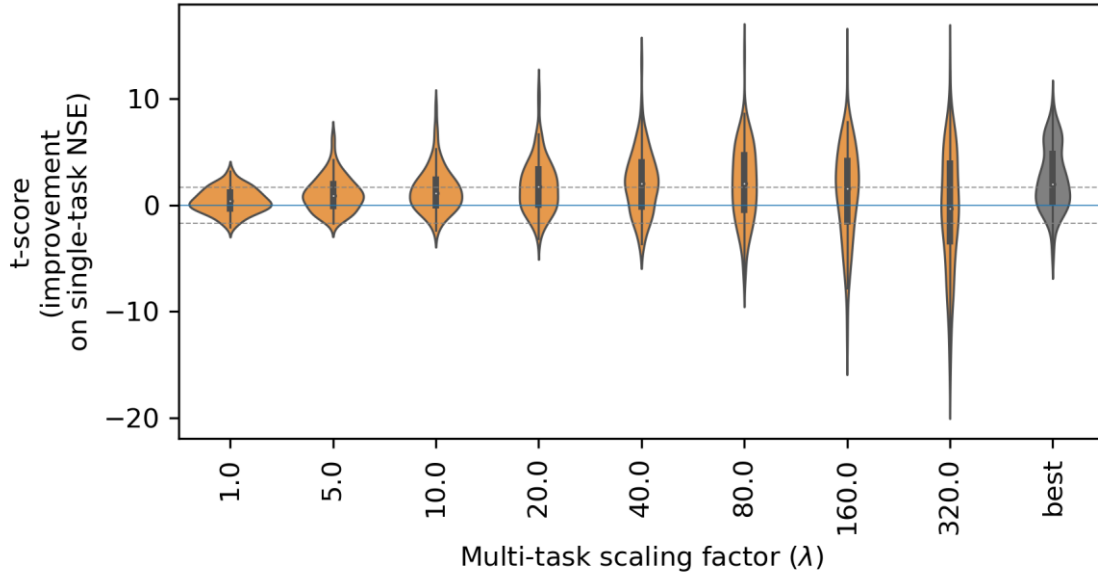*Figure 5. Spatial distribution of model NSEs and t-scores*

Figure 6. t-scores of multi-task model in test set. The "best" is the lambda that was best for each site based on the validation set. The distributions are across the 103 sites used in Experiment A. The bottom of each box is the first quartile of the data (Q1), the white circle is the median of the data, and the top of each box is the third quartile (Q3). The lower whisker is the lowest point within Q1 – 1.5*IQR, where IQR is the innerquartile range (Q3-Q1). The upper whisker is the highest point within Q3 + 1.5*IQR. The dotted lines mark where the t-scores are statistically significantly worse (below zero) or better (above zero) with p=0.05.

Table 2. Number of sites that were significantly better and significantly worse than the single-task models for several multi-task scaling factors. Significance was tested using the students t-test.

Multi-task scaling factor	Significantly better (n sites) out of 104 total (p < 0.05)	Significantly worse (n sites) out of 104 total (p<0.05)
1	16	4
5	32	1
10	40	3
20	54	4
40	59	7
80	57	17
160	51	24
320	39	39
best	56	4

3.1.2 Catchment characteristics

The Random Forest model was trained to predict the t-score of the multi-task models compared to the single-task models when the multi-task scaling factor was 80. The model had a mean R-squared value 0.2 (SD 0.29). This performance was measured on sites not present in the training data across 20 random data partitions. The model performance indicates that the model was able to learn general trends from the data that could help it learn to predict the multi-task improvements from the catchment characteristics. One likely limiting factor in the Random Forest model performance was the small dataset. There were only 103 sites to draw from. We used a 75/25 test/train split, so and only 77 of those could be used for training.

The importance of the Random Forest predictor variables was calculated based on the effect that perturbing a given input variable had on the prediction error. The most important variable was baseflow index (Figure 7). A description of each predictor variable is given in the supplemental material. Baseflow is an influential factor in both streamflow and water temperature. Sites with a larger baseflow index will have a different flow patterns and will have a thermal buffer against extremes in air temperature affecting water temperature compared to sites with a lower baseflow index. This buffer can alter the relative influence of different parts of the energy balance equation that governs the energy transfer process and that links the streamflow and water temperature variables. It is reasonable, therefore, that baseflow index is an important variable in predicting how much a multi-task deep learning approach would benefit a given site.
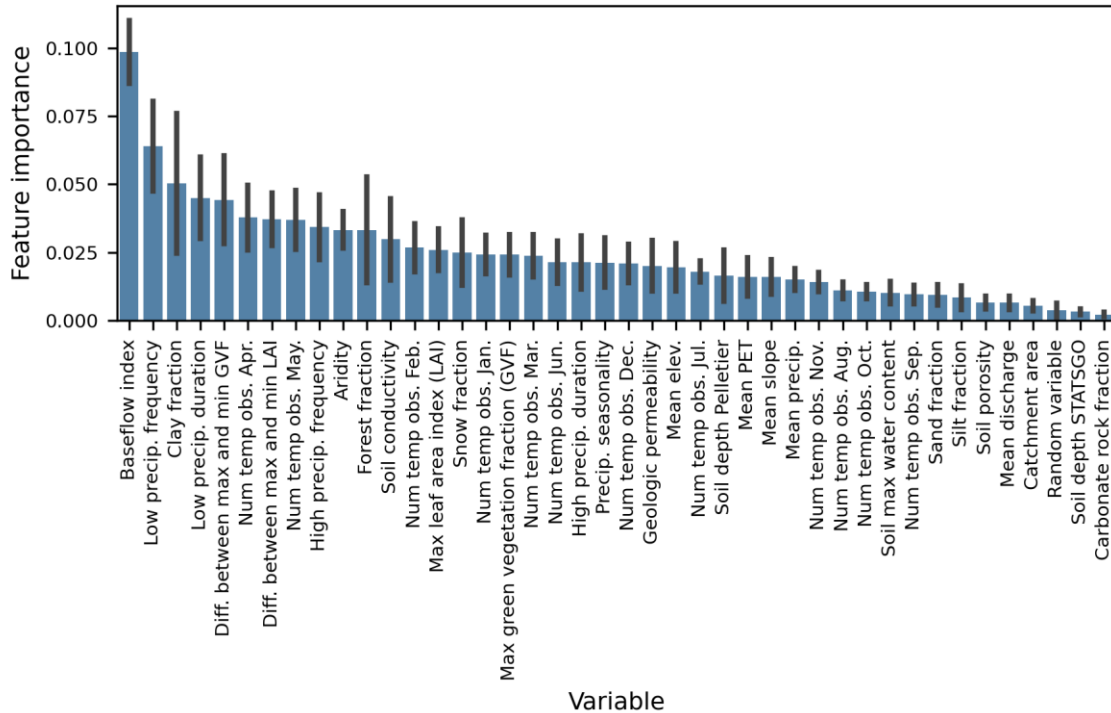

Figure 7. Feature importance for predicting t-scores for Experiment A when the multi-task scaling factor is 80. Error bars are the standard deviation across 20 random train/test splits.

16

3.1.3 Seasonal trends

The improvements on the single-task models made by the multi-task models had a noticeable seasonal trend with larger improvements in higher flow months and smaller improvements in lower flow months (Figure 8). A greater portion of the loss occurred in the high-flow periods. This is likely because the loss function is based on the squared error: a 10% error when the observed value is 40 cubic meters per secont (CMS) will have 10x the impact on the overall loss compared to a 10% error when the observed value is 4 CMS. Therefore, the multi-task (as well as the single-task) model parameters adjusted to reduce error in springtime. These results indicate that the multi-task models were able to exploit the relationship between streamflow and water temperature to further improve model performance in the spring months. The parameter values that adjusted to fit that relationship, however, did not seem to benefit the model in the summer months where, overall, the multi-task models performed worse than the single-task models. This difference indicates that the relationship learned in the high-flow months of the spring did not apply in the summertime.
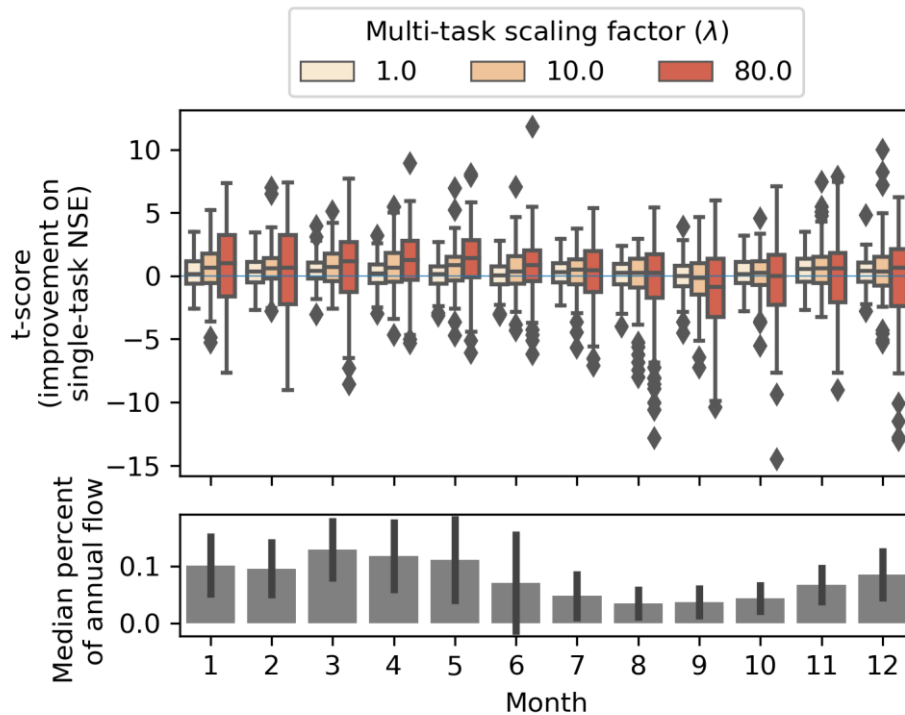

Figure 8. Monthly t-score distributions for predicting streamflow using 24 years of training data. The distributions are across the 103 sites used in the experiments. Also shown are the monthly median percent of annual flow across the 103 sites. In the boxplots, the bottom of each box is the first quartile of the data (Q1), the middle bar is the median of the data, and the top of each box is the third quartile (Q3). The lower whisker is the lowest point within Q1 − 1.5*IQR, where IQR is the innerquartile range (Q3-Q1). The upper whisker is the highest point within Q3 + 1.5*IQR. The diamonds are points that lie outside of the whiskers. Error bars in monthly median flows are the standard deviation across sites.

17

3.2   Experiment B: Multi-task modeling for streamflow with 2 years of training data

In general, the multi-task models with the best multi-task scaling factors improved upon the single task models in Experiment B (see **Error! Reference source not found.**4). However, the difference between the single-task and multi-task models was smaller in Experiment B compared to Experiment A. The model performance across the single-task models and all the multi-task scaling factors was worse in Experiment B (single-site median 0.45) compared to Experiment A (overall median of 0.61). This was expected given that there was a much smaller amount of training data from which the model could learn the flow dynamics (and temperature dynamics in the case of the multi-task models). There was more variability in the results and the effect of the multi-task model when the data were sparser. This variation was likely due to the variation in training data that the model used, because it had only 2 years of training data and there can be a lot of variability between a given 2-year period of streamflow record and any other 2-year period. The multi-task scaling factor that was the best for the most sites in Experiment B was much smaller compared to Experiment A: in Experiment A it was 80, in Experiment B it was 10.

Table 3. Number of sites that were significantly better and significantly worse than the single-task models for several multi-task scaling factors in Experiment B. Significance was tested using the students t-test.

Multi-task scaling factor	Significantly better (n sites) ($p < 0.05$) out of 98 total sites	Significantly worse (n sites) ($p<0.05$) out of 98 total sites
1	24	4
5	39	10
10	42	11
20	42	13
40	39	19
best	46	5

3.3   Experiment C: Multi-task modeling for temperature

The multi-task approach did not improve upon the single-task model as substantially for predicting water temperature as it did for predicting streamflow (see Figure 4 and **Error! Reference source not found.**). The multi-task scaling factor that was the best for the most sites in Experiment C was much smaller compared to Experiment A: in Experiment A it was 80, in Experiment B it was 0.01. Although the improvements were more modest with a multi-task scaling factor value of 0.01, there was improvement upon the single-task model for most of the

sites. When using the best site-specific multi-task scaling factor, the improvements were greater than when using the multi-task scaling factor for all sites.
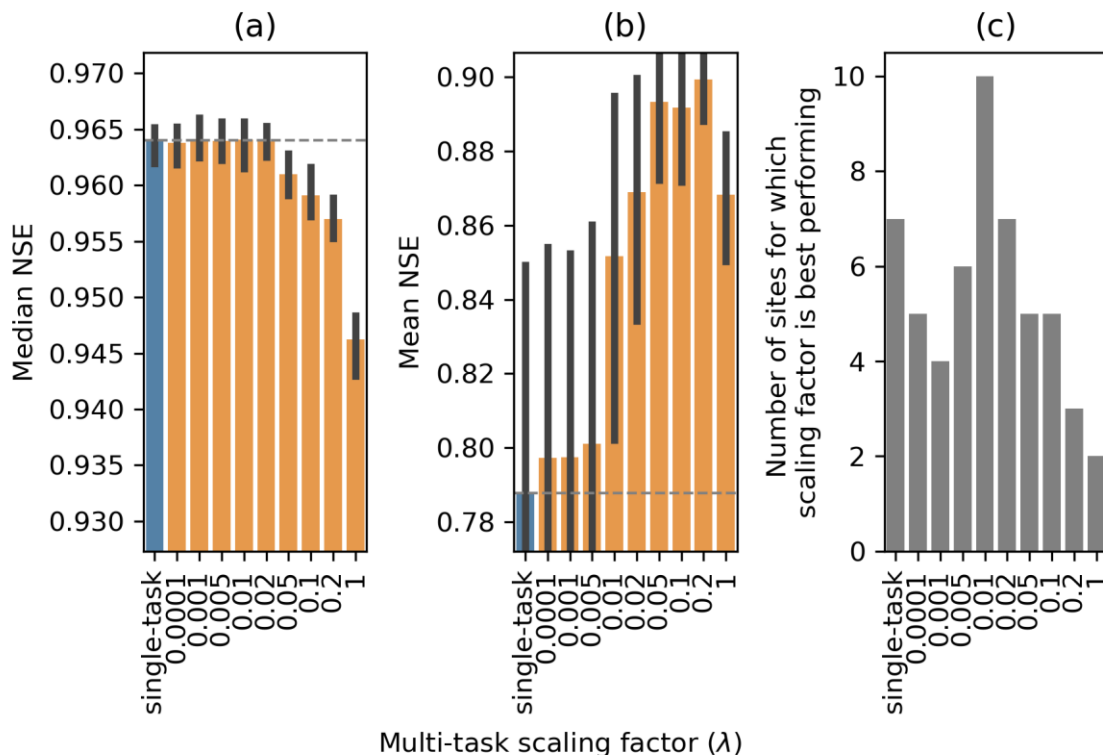

Figure 9. Influence of mutli-task scaling factor on streamflow prediction accuracy. (a) Performance of single-task and multi-tasks models with various values of multi-task scaling factors in terms of median NSE. (b) Performance of single-task and multi-tasks models with various values of multi-task scaling factors in terms of mean NSE. (c) The number of sites for which each multi-task scaling factor performed the best in terms of median NSE. Results shown for all panels are from the validation set. The dotted horizontal line in (a) and (b) is the median performance of the single task model. The error bars in (a) and (b) represent 95% confidence intervals.

4   Discussion

The goal of our experiments was to evaluate the performance of a multi-task modeling approach for predicting streamflow and water temperature. We compared the accuracy of the multi-task approach (predicting streamflow and water temperature simultaneously with one model) to a single-task modeling approach (predicting streamflow or water temperature with their own independent models). Because streamflow and water temperature are related, we hypothesized that a multi-task model would be more accurate than a single-task model. The results of our

experiments were nuanced. In some cases, the multi-task approach performed significantly better, in some there was no significant difference, and in some the multi-task model performed significantly worse. The improvements made by the multi-task approach seemed to depend on (1) seasonal factors, (2) which variable was being used as the primary variable, and (3) the site being modeled. Together these affected the appropriateness of the multi-task scaling factor. The nuance in our results indicates further work is needed to fully understand and be able to leverage multi-task learning for improved streamflow and water temperature prediction.

4.1   Seasonality of Experiment A results

There was an apparent seasonal pattern to the improvements made by the multi-task approach where, in general, the improvements were greater in the high flow months and lesser in the low flow months (see Figure 8). The influence of snowmelt and baseflow are two possible physical mechanisms that could cause differences in streamflow-water temperature relationship in spring compared to summer and therefore cause this seasonal pattern. Snowmelt and baseflow have seasonal patterns and can be strong influences on both streamflow and water temperature (Anderson, 2005; Constantz, 2008; Lisi et al., 2015). For a site that is significantly affected by snowmelt, a multi-task model might have parameters that are trained to predict a decrease in stream temperature for conditions in which an increase in streamflow is also predicted, as would happen during a snowmelt event. Those parameters could make performance worse, however, when increases in streamflow are not due to snowmelt but to rainfall as would happen in the summer months. Similarly, a multi-task model whose parameters capture the flow and temperature dynamics of high-flow months may be less well-suited to predict those dynamics in the summertime when the relative influence of baseflow on streamflow and water temperature changes.

The tendency to favor the high-flow periods would not be desirable for certain use cases such as drought. This behavior could be addressed in future work in several ways. As a first approach, the loss function could be changed to an average of the NSEs of each season. That way the errors in the model during the low-flow summer seasons would be considered separate from, and would not be overshadowed by, the errors during the higher-flow spring season. This could, however, make a model that is mediocre in performance for each of the seasons instead of better in springtime and worse in summertime. As a more extensive solution, one could train season-specific models where each model was only trained using data from a certain season. This approach would produce a model for each season. When making predictions, the season-specific model outputs would be concatenated together. In this way, the model would be able to focus on season-specific dynamics between streamflow and water temperature. One draw-back of this approach, however, could be the significantly smaller dataset that would result from dividing the training record by season. Another opportunity is to have lambda dynamically adjusted for different time steps. For example, the multi-task scaling factor could be a function of the day of year.

4.2   Smaller improvements in predicting water temperature

The multi-task models had much smaller improvements on the single-task models when the main variable was stream temperature compared to when the main variable was streamflow. One reason for this difference may be that, in terms of NSE, water temperature is easier to predict than streamflow so less improvement can be made in predicting water temperature. Both the

seasonal trends in water temperature and the constrained range (between 0 and about 30 degrees C), make water temperature predictions easier terms of annual NSE. In general, stream temperatures have a predictable sinusoidal pattern with peak temperatures in the summer and lowest temperatures in the winters. An approach that could prove more fruitful for improving the model for stream temperature predictions is to calculate season-specific NSE values and average those for the training loss. In that way, the model would need to do more than just predicting the seasonal trends to results in good performance.

The fact that the multi-task scaling factor was much smaller for predicting stream temperature compared to streamflow can be explained, at least in part, by the difference in the variability of the streamflow data compared to the variability of the water temperature data: water temperature rarely exceeds 30 degrees C and streamflow has no practical upper limit. Because streamflow is much more variable, without a small multi-task scaling factor, the model parameters will adjust to improve streamflow predictions because a larger portion of the error gradient will come from the streamflow error. A multi-task scaling factor much smaller than 1 is required to compensate for this difference in variability between the two variables. As an alternative to decreasing the multi-task scaling factor to such a large extent, one could train on the log of streamflow. This would greatly reduce the variability of the streamflow record and therefore reduce the contribution to the combined gradient relative to the gradient from the stream temperature errors. Training on the log of streamflow would still require the tuning of the multi-task scaling factor, though an appropriate value would likely be much larger.

4.3   Selecting an appropriate multi-task scaling factor

One of the main outcomes of our results is that they highlight the importance of choosing an appropriate multi-task scaling factor. The decrease in accuracy in terms of median NSE with multi-task scaling factors greater than 80 in Experiment A likely occurred because the combined gradient used to update the parameters became more influenced by the temperature loss contribution. With the gradient more influenced by the temperature loss, the parameter adjustments in gradient descent was directed more to reducing error in the temperature predictions. When the multi-task scaling factor was high enough, those adjustments reduced error in the temperature predictions at the expense of streamflow predictive performance. This is likely the reason that there was an overall decrease in streamflow prediction accuracy in terms of median NSE with the multi-task scaling factors of 160 and 320.

In our results, the appropriateness of a given multi-task scaling factor depended on which site we were predicting at, the availability of the training data, and whether streamflow or stream temperature was the main variable. In a general sense, we reason that the factors that contribute to an appropriate multi-task scaling factor boil down to the statistical relationship between the two variables. If one simply combines the losses of two related variables without considering the statistical relationship between the variables and/or without exploring how that will affect the performance of predicting the main variable, the model could end up doing worse than a single-task model. If the main variable has a larger variation than the auxiliary variable, the risk is much lower because the user is likely to select a multi-task scaling factor that is smaller than the optimal; this will result in a model that is no worse than the single-task model. If, however, the main variable has a smaller variation than the auxiliary variable, assuming a naïve multi-task scaling factor of 1 could result in model whose parameters will heavily favor reducing error in the auxiliary variable as the larger source of error. This can result in a much worse predictive

model for the main variable. This behavior was seen in our experiments when water temperature was the main variable.

One factor that affected which multi-task training factor was most effective for a given site was the ratio of training examples (365-day sequences) with data available for the main variable compared to training examples with data available the auxiliary variable. In model training, the examples without any temperature data produced gradients that were only informed by the errors in streamflow and were, therefore, in effect, single-task gradients. This was because in each training epoch, the average of the gradients from each example (a year-long period) was used to update the model weights. This is a typical approach in training a deep learning model (Le et al., 2011; Li et al., 2014). There were many years at various sites that had no temperature training data at all. We saw, in general, the inverse relationship between the number of training examples with water temperature data and the best multi-task scaling factor (see Figure 10 for the relationship in Experiment A). We hypothesize that when there were only a few training examples with the auxiliary data, a large multi-task scaling factor was needed to make an impact on the overall gradient used to update the model weights. The larger multi-task scaling factor amplified the effect of the gradient from the few examples with auxiliary data enough to counter-balance the influence from the larger number of gradients produced from the examples that had only data from the main variable. This explanation also holds for Experiment B where the best multi-task scaling factors were typically smaller (most commonly 10) and the ratio of training examples with the auxiliary variable to training examples with the main variable was larger (either 0.5 or 1).
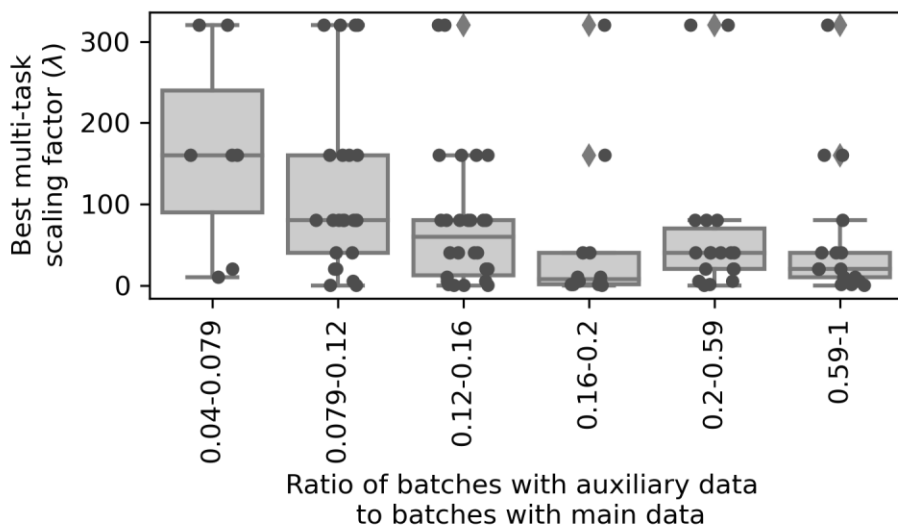

Figure 10. The best average multi-task scaling factor in Experiment A for sites with varying ratios of training examples with auxiliary (water temperature) data to training examples with main (streamflow) data. Circles are individual training examples. The bottom of each box is the first quartile of the data (Q1), the middle bar is the median of the data, and the top of each box is the third quartile (Q3). The lower whisker is the lowest point within Q1 – 1.5*IQR, where IQR is the innerquartile range (Q3-Q1). The upper whisker is the highest point within Q3 + 1.5*IQR. The diamonds are points that lie outside of the whiskers.

There may be more efficient methods of deciding on an appropriate multi-task scaling factor compared to just manual adjustment. For example, Kraft et al. (2020) used the uncertainty in the each variable to scale the influence of the variable in the loss function. One reason we did not explore this approach is that it does not allow for the explicit designation of a main variable of interest, whose performance should not be sacrificed at the expense of the other variables being predicted. We were instead interested in an approach that would only benefit the prediction of the main variable. There have been some more automated approaches for achieving this objective that could be explored in future research (Du et al., 2020; Liu et al., 2019).

It can reasonably be assumed that the appropriateness of a multi-task scaling factor is largely dependent on the statistical relationship between the two variables of interest. Acknowledging that the statistical relationship is influenced (if not determined) by physical characteristics and processes, it stands to reason that it should be possible to leverage our understanding of the physical processes to estimate an appropriate multi-task scaling factor for a given site and combination of variables. Although we did not discover a predictable pattern for estimating an appropriate multi-task scaling factor value based on catchment characteristics and data availability, that pattern may exist for streamflow and stream temperature with another set of catchment characteristics or different method of relating them to the multi-task scaling factor. If multi-task learning were being explored for other variables, there may be more easily discoverable patterns that would narrow the search for a beneficial multi-task scaling factor.

4.4 Limitations and potential future work

As with all data-driven approaches, the models we used were limited by the training data available. Likewise, the effectiveness of the approach rests on the assumption that the past will be a good predictor of the future. Evidence indicates that this assumption of stationarity is becoming less valid given the observed and expected climate change. Nevertheless, in our results the models performed well even when trained on data from certain decades and evaluated on data from other decades. Despite any changes to the climate that may have already occurred and that may occur in the future, the underlying physical, chemical, and biological mechanisms that drive the processes that hydrologists model will not change. This highlights the importance of the infusion of process understanding into our models rather than relying wholly on data that cannot adequately capture full process of the of the dynamics being modeled.

This paper examines multi-task learning for just two hydrologic variables; however, the approach could be applied to predicting other variables. We hypothesize that for multi-task learning to benefit, the variables would have to have some statistical relationship (stemming from a physical/chemical/biological relationship). There would also need to be sufficient data that for those variables from which the multi-task model could learn. For example, other parts of the water balance equation such as soil moisture that are observed could be used to train a multi-task model with a main objective of predicting streamflow. Another example is the modeling of dissolved oxygen and water temperature together. Multi-task learning could be used to predict any number of observed variables. For example, a model could be trained to predict streamflow, water temperature, and soil moisture. As we discussed above, the selection of the multi-task scaling factor would be an important part of a successful multi-task model. This selection would be even more interesting (and possibly more difficult) if multiple variables were used as auxiliary variables because there would need to be a multi-task scaling factor for each auxiliary variable.

Although the variables for multi-task learning must have enough data for training, that data could be simulated data. For example, it may be fruitful to train a multi-task model on all water balance outputs produced from a process-based hydrologic model such as the Soil and Water Assessment Tool (SWAT) model (e.g., evapotranspiration, runoff, infiltration). Then the model could fine-tune to just the data for which observations are available. The model would presumably learn the dynamics of the hydrologic cycle more completely from the process model output and could then better predict the variable of interest when given unseen inputs. This would also open the possibility of assimilating observations of multiple observed variables in a forecast application.

5   Conclusions

Our study examined the value of one approach for incorporating process-understanding into deep learning models for hydrologic prediction: multi-task deep learning (DL). In the multi-task approach, one DL model was trained to predict both streamflow and water temperature. The effectiveness of the approach was tested using three experiments: predicting streamflow with 25 years of training data, predicting streamflow with 2 years of training data, and predicting water temperature. For the experiments we used data from sites throughout the conterminous United States.

Compared to the single-task models, the multi-task approach significantly improved streamflow predictions for most of the monitoring sites with the full 25 years of training data. With only 2 years of training data, the multi-task approach significantly improved streamflow predictions at nearly half of the sites. When predicting water temperature, the multi-task approach improved predictions at only 7 sites. We believe the main reason that the multi-task approach had a smaller benefit for stream temperature predictions was that the stream temperature predictions were already very good (NSE > 0.95).

One of the most important factors in the effectiveness (or lack of effectiveness) of the approach was the multi-task scaling factor. This model parameter scaled the influence of the errors in the secondary (or auxiliary) variable in the overall error used to update the model parameters. When streamflow was the primary (or main) variable, the scaling factor had to be quite high (most commonly the best factor was 80). When the main variable was stream temperature, however, the scaling factor had to be very low (most commonly the best factor was 0.01). A naïve factor of 1 led to models that performed much worse for predicting stream temperature than the single-task models. The main reason for the difference in the scaling factors could be attributed to the difference in variability between the two variables. Seasonal trends in the effectiveness of the multi-task model as well as examining the characteristics of the catchments that benefitted most and least from the approach suggest that baseflow is an important factor in whether this approach is successful or not for improving streamflow predictions.

Streamflow and water temperature are but two among many variables that are part of the physical/chemical/biological processes that are of interest to hydrologic scientists and engineers. Our findings indicate that multi-task learning, when configured properly, can be an effective way to leverage our knowledge of the relationship between hydrologic variables in a deep learning model for improved predictions.

Acknowledgments and Data Availability

References

Addor, N., Newman, A. J., Mizukami, N., & Clark, M. P. (2017). The CAMELS data set: Catchment attributes and meteorology for large-sample studies. *Hydrology and Earth System Sciences*. https://doi.org/10.5194/hess-21-5293-2017

Alemu, E. T., Palmer, R. N., Polebitski, A., & Meaker, B. (2011). Decision support system for optimizing reservoir operations using ensemble streamflow predictions. *Journal of Water Resources Planning and Management*, *137*(1), 72–82. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000088

Anderson, M. P. (2005, November). Heat as a ground water tracer. *Ground Water*. https://doi.org/10.1111/j.1745-6584.2005.00052.x

Bai, Y., Chen, Z., Xie, J., & Li, C. (2016). Daily reservoir inflow forecasting using multiscale deep feature learning with hybrid models. *Journal of Hydrology*, *532*, 193–206. https://doi.org/10.1016/j.jhydrol.2015.11.011

Bain, M. B., Finn, J. T., & Booke, H. E. (1988). Streamflow regulation and fish community structure. *Ecology*, *69*(2), 382–392. https://doi.org/10.2307/1940436

Barnett, P. R. (1972). Effects of warm water effluents from power stations on marine life. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, *180*(61), 497–509. https://doi.org/10.1098/rspb.1972.0034

Besaw, L. E., Rizzo, D. M., Bierman, P. R., & Hackett, W. R. (2010). Advances in ungauged streamflow prediction using artificial neural networks. *Journal of Hydrology*, *386*(1–4), 27–37. https://doi.org/10.1016/J.JHYDROL.2010.02.037

Bowes, B. D., Sadler, J. M., Morsy, M. M., Behl, M., & Goodall, J. L. (2019). Forecasting groundwater table in a flood prone coastal city with long short-term memory and recurrent neural networks. *Water (Switzerland)*, *11*(5). https://doi.org/10.3390/w11051098

Breiman, L. (2001). Random Forests. *Machine Learning*, *45*(1), 5–32. https://doi.org/10.1023/A:1010933404324

Campolo, M., Andreussi, P., & Soldati, A. (1999). River flood forecasting with a neural network model. *Water Resources Research*, *35*(4), 1191–1197. https://doi.org/10.1029/1998WR900086

Carlisle, D. M., Wolock, D. M., & Meador, M. R. (2011). Alteration of streamflow magnitudes and potential ecological consequences: A multiregional assessment. *Frontiers in Ecology and the Environment*, *9*(5), 264–270. https://doi.org/10.1890/100053

Chen, D., Mak, B., Leung, C., & Sivadas, S. (2014). Joint acoustic modeling of triphones and trigraphemes by multi-task learning deep neural networks for low-resource speech recognition. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 5592–5596). https://doi.org/10.1109/ICASSP.2014.6854673

Chen, P. A., Chang, L. C., & Chang, F. J. (2013). Reinforced recurrent neural networks for multi-step-ahead flood forecasts. *Journal of Hydrology*, *497*, 71–79. https://doi.org/10.1016/j.jhrol.2013.05.038

Constantz, J. (2008). Heat as a tracer to determine streambed water exchanges. *Water Resources Research*, *46*(4). https://doi.org/10.1029/2008WR006996

Dask Development Team. (2016). Dask: Library for dynamic task scheduling. Retrieved from https://dask.org

Du, Y., Czarnecki, W. M., Jayakumar, S. M., Farajtabar, M., Pascanu, R., & Lakshminarayanan, B. (2020). Adapting auxiliary losses using gradient similarity. Retrieved from: https://arxiv.org/abs/1812.02224v2

Girshick, R. (2015). Fast R-CNN. In *Proceedings of the IEEE International Conference on Computer Vision*. Retrieved from https://github.com/rbgirshick/

Hamlet, A. F., Huppert, D., & Lettenmaier, D. P. (2002). Economic value of long-lead streamflow forecasts for Columbia River hydropower. *Journal of Water Resources Planning and Management*, *128*(2), 91–101. https://doi.org/10.1061/(ASCE)0733-9496(2002)128:2(91)

Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., et al. (2020). Array programming with NumPy. *Nature*, *585*, 357–362.

https://doi.org/10.1038/s41586-020-2649-2

Hochreiter, S., & Schmidhuber, J. (1997). Long Short-Term Memory. *Neural Computation*, *9*(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

Hoyer, S., & Hamman, J. (2017). xarray: N-D labeled arrays and datasets in Python. *Journal of Open Research Software*, *5*(1). https://doi.org/10.5334/jors.148

Hsu, K. -l, Gupta, H. V., & Sorooshian, S. (1995). Artificial neural network modeling of the rainfall-runoff process. *Water Resources Research*, *31*(10), 2517–2530. https://doi.org/10.1029/95WR01955

Hu, C., Wu, Q., Li, H., Jian, S., Li, N., & Lou, Z. (2018). Deep learning with a long short-term memory networks approach for rainfall-runoff simulation. *Water*, *10*(11), 1543. https://doi.org/10.3390/w10111543

Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, *9*(3), 90–95. https://doi.org/10.1109/MCSE.2007.55

Jia, X., Zwart, J., Sadler, J., Appling, A., Oliver, S., Markstrom, S., et al. (2020, September 26). Physics-guided recurrent graph networks for predicting flow and temperature in river networks. *ArXiv*. arXiv. Retrieved from http://arxiv.org/abs/2009.12575

Kingma, D. P., & Ba, J. L. (2015). Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*.

Köster, J., & Rahmann, S. (2012). Snakemake—a scalable bioinformatics workflow engine. *Bioinformatics*, *28*(19), 2520–2522. https://doi.org/10.1093/bioinformatics/bts480

Kraft, B., Jung, M., Körner, M., & Reichstein, M. (2020). Hybrid Modeling: fusion of a deep learning approach and a physics-based model for global hydrological modeling. *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, *XLIII-B2-2*, 1537–1544. https://doi.org/10.5194/isprs-archives-XLIII-B2-2020-1537-2020

Kratzert, F., Klotz, D., Brenner, C., Schulz, K., & Herrnegger, M. (2018). Rainfall–runoff modelling using Long Short-Term Memory (LSTM) networks. *Hydrology and Earth System Sciences*, *22*(11), 6005–6022. https://doi.org/10.5194/hess-22-6005-2018

Kratzert, F., Klotz, D., Herrnegger, M., Sampson, A. K., Hochreiter, S., & Nearing, G. S. (2019). Toward improved predictions in ungauged basins: Exploiting the power of machine learning. *Water Resources Research*, *55*(12), 11344–11354. https://doi.org/10.1029/2019WR026065

Kratzert, F., Klotz, D., Shalev, G., Klambauer, G., Hochreiter, S., & Nearing, G. (2019). Towards learning universal, regional, and local hydrological behaviors via machine learning applied to large-sample datasets. *Hydrology and Earth System Sciences*, *23*(12), 5089–5110.

https://doi.org/10.5194/hess-23-5089-2019

Le, Q. V, Ngiam, J., Coates, A., Lahiri, A., Prochnow, B., & Ng, A. Y. (2011). On optimization methods for deep learning. *Proceedings of the 28th International Conference on Machine Learning, ICML 2011 https://dl.acm.org/doi/10.5555/3104482.310451*

Li, M., Zhang, T., Chen, Y., & Smola, A. J. (2014). Efficient mini-batch training for stochastic optimization. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 661–670). New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/2623330.2623612

Lisi, P. J., Schindler, D. E., Cline, T. J., Scheuerell, M. D., & Walsh, P. B. (2015). Watershed geomorphology and snowmelt control stream thermal sensitivity to air temperature. *Geophysical Research Letters*, *42*(9), 3380–3388. https://doi.org/10.1002/2015GL064083

Liu, S., Davison, A. J., & Johns, E. (2019). Self-supervised generalisation with meta auxiliary Learning. Retrieved from https://arxiv.org/abs/1901.08933v3

Martin, A., Ashish, A., Paul, B., Eugene, B., Zhifeng, C., Craig, C., et al. (2015). TensorFlow: Large-scale machine learning on heterogeneous systems. Retrieved from https://www.tensorflow.org/

Nash, J. E., & Sutcliffe, J. V. (1970). River flow forecasting through conceptual models part I - A discussion of principles. *Journal of Hydrology*, *10*(3), 282–290. https://doi.org/10.1016/0022-1694(70)90255-6

Newman, A. J., Clark, M. P., Sampson, K., Wood, A., Hay, L. E., Bock, A., et al. (2015). Development of a large-sample watershed-scale hydrometeorological data set for the contiguous USA: Data set characteristics and assessment of regional variability in hydrologic model performance. *Hydrology and Earth System Sciences*. https://doi.org/10.5194/hess-19-209-2015

pandas development team, T. (2020, February). pandas-dev/pandas: Pandas. Zenodo. https://doi.org/10.5281/zenodo.3509134

Pappenberger, F., Cloke, H. L., Parker, D. J., Wetterhall, F., Richardson, D. S., & Thielen, J. (2015). The monetary benefit of early flood warnings in Europe. *Environmental Science and Policy*, *51*, 278–291. https://doi.org/10.1016/j.envsci.2015.04.016

Rahmani, F., Lawson, K., Ouyang, W., Appling, A., Oliver, S., & Shen, C. (2020). Exploring the exceptional performance of a deep learning stream temperature model and the value of streamflow data. *Environmental Research Letters*. https://doi.org/10.1088/1748-9326/abd501

Read, J. S., Jia, X., Willard, J., Appling, A. P., Zwart, J. A., Oliver, S. K., et al. (2019). Process-guided deep learning predictions of lake water temperature. *Water Resources Research*, *55*(11), 9173–9190. https://doi.org/10.1029/2019WR024922

Robertson, D. E., Shrestha, D. L., & Wang, Q. J. (2013). Post-processing rainfall forecasts from numerical weather prediction models for short-term streamflow forecasting. *Hydrol. Earth Syst. Sci*, *17*, 3587–3603. https://doi.org/10.5194/hess-17-3587-2013

Ruder, S. (2017). An overview of multi-task learning in deep neural networks. *ArXiv Preprint ArXiv:1706.05098.*

Sadler, J. M., Appling, A. P., Read, J. S., Oliver, S. K., Jia, X., Zwart, J. A., & Kumar, V. (2021). Data release: Multi-task deep learning for water temperature and streamflow prediction. U.S. Geological Survey. https://doi.org/10.5066/P9U0TG8L

Satterthwaite, F. E. (1946). An approximate distribution of estimates of variance components. *Biometrics Bulletin*, *2*(6), 110. https://doi.org/10.2307/3002019

Seltzer, M. L., & Droppo, J. (2013). Multi-task learning in deep neural networks for improved phoneme recognition. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing* (pp. 6965–6969). https://doi.org/10.1109/ICASSP.2013.6639012

Shen, C. (2018, November 1). A transdisciplinary review of deep learning research and its relevance for water resources scientists. *Water Resources Research*. https://doi.org/10.1029/2018WR022643

Shen, C., Laloy, E., Elshorbagy, A., Albert, A., Bales, J., Chang, F. J., et al. (2018). HESS Opinions: Incubating deep-learning-powered hydrologic science advances as a community. *Hydrology and Earth System Sciences*, *22*(11), 5639–5656. https://doi.org/10.5194/hess-22-5639-2018

Soleimani, S., Bozorg-Haddad, O., Saadatpour, M., & Loáiciga, H. A. (2019). Simulating thermal stratification and modeling outlet water temperature in reservoirs with a data-mining method. *Journal of Water Supply: Research and Technology - AQUA*, *68*(1), 7–19. https://doi.org/10.2166/aqua.2018.036

Stammler, K. L., Yates, A. G., & Bailey, R. C. (2013). Buried streams: Uncovering a potential threat to aquatic ecosystems. *Landscape and Urban Planning*, *114*, 37–41. https://doi.org/10.1016/J.LANDURBPLAN.2013.02.008

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., et al. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, *17*, 261–272. https://doi.org/10.1038/s41592-019-0686-2

Van Vliet, M. T. H., Ludwig, F., Zwolsman, J. J. G., Weedon, G. P., & Kabat, P. (2011). Global river temperatures and sensitivity to atmospheric warming and changes in river flow. *Water Resources Research*, *47*(2). https://doi.org/10.1029/2010WR009198

Waskom, M., & the seaborn development team. (2020, September). Seaborn. Zenodo. https://doi.org/10.5281/zenodo.592845

Willard, J. D., Jia, X., Xu, S., Steinbach, M., & Kumar, V. (2020, March 10). Integrating

physics-based modeling with machine learning: A survey. *ArXiv*. arXiv. Retrieved from https://doi.org/10.1145/1122445.1122456

Zahura, F. T., Goodall, J. L., Sadler, J. M., Shen, Y., Morsy, M. M., & Behl, M. (2020). Training machine learning surrogate models from a high-fidelity physics-based model: Application for real-rime street-scale flood prediction in an urban coastal community. *Water Resources Research*, *56*(10). https://doi.org/10.1029/2019WR027038

Zeng, L. Q., Cao, Z. D., Fu, S. J., Peng, J. L., & Wang, Y. X. (2009). Effect of temperature on swimming performance in juvenile southern catfish (Silurus meridionalis). *Comparative Biochemistry and Physiology - A Molecular and Integrative Physiology*, *153*(2), 125–130. https://doi.org/10.1016/j.cbpa.2009.01.013

S1.   Implementation of Random Forest model

The Random Forest model was implemented using the Scikit learn Python package. We used the RandomForestRegressor class. The training and testing data were split so that 75% of the sites were used in the training and 25% of the sites were used in the testing. We did the model training and evaluation for 20 random train/test splits. The meta parameters of the model are shown in Table S1. The list of the different features and their descriptions are listed in Table S2.

Table S1. Random Forest Hyperparameter values

Hyper-parameter	Value
Number of trees	1000
Max tree depth	2
Max features per tree	6

Table S2. Random Forest features and descriptions. These are taken from the CAMELS dataset. For feature units, data sources and references, see the CAMELS dataset (Addor et al., 2017; Newman et al., 2015).

Figure 7 display name	CAMELS attribute name	Description	Unit
Mean elev.	elev_mean	catchment mean elevation	meter above sea level
Mean slope	slope_mean	catchment mean slope	m/km
Catchment area	area_gages2	catchment area (GAGESII estimate)	km^2
Mean precip.	p_mean	mean daily precipitation	mm/day
Mean PET	pet_mean	mean daily PET [estimated by N15 using Priestley-Taylor formulation calibrated for each catchment]	mm/day
Aridity	aridity	aridity (PET/P, ratio of mean PET [estimated by N15 using Priestley-Taylor formulation calibrated for each catchment] to mean precipitation)	-
Precip. seasonality	p_seasonality	seasonality and timing of precipitation (estimated using sine curves to represent the annual temperature and precipitation cycles, positive [negative] values indicate that precipitation peaks in summer [winter], values close to 0 indicate uniform precipitation throughout the year)	-
Snow fraction	frac_snow	fraction of precipitation falling as snow (i.e., on days colder than 0°C)	-
High precip. frequency	high_prec_freq	frequency of high precipitation days (>= 5 times mean daily precipitation)	days/year
High precip. duration	high_prec_dur	average duration of high precipitation events (number of consecutive days >= 5 times mean daily precipitation)	days

Low precip. frequency	low_prec_freq	frequency of dry days (<1 mm/day)	days/year
Low precip. duration	low_prec_dur	average duration of dry periods (number of consecutive days <1 mm/day)	days
Baseflow index	baseflow_index	baseflow index (ratio of mean daily baseflow to mean daily discharge, hydrograph separation performed using Ladson et al. [2013] digital filter)	-
Forest fraction	frac_forest	forest fraction	-
Max leaf area index (LAI)	lai_max	maximum monthly mean of the leaf area index (based on 12 monthly means)	-
Diff. between max and min LAI	lai_diff	difference between the maximum and minimum monthly mean of the leaf area index (based on 12 monthly means)	-
Max green vegetation fraction (GVF)	gvf_max	maximum monthly mean of the green vegetation fraction (based on 12 monthly means)	-
Diff. between max and min GVF	gvf_diff	difference between the maximum and minimum monthly mean of the green vegetation fraction (based on 12 monthly means)	-
Soil depth Pelletier	soil_depth_pelletier	depth to bedrock (maximum 50 m)	m
Soil depth STATSGO	soil_depth_statsgo	soil depth (maximum 1.5 m, layers marked as water and bedrock were excluded)	m
Soil porosity	soil_porosity	volumetric porosity (saturated volumetric water content estimated using a multiple linear regression based on sand and clay fraction for the layers marked as USDA soil texture class and a default value [0.9] for layers marked as organic material, layers marked as water, bedrock, and "other" were excluded)	-
Soil conductivity	soil_conductivity	saturated hydraulic conductivity (estimated using a multiple linear regression based on sand and clay fraction for the layers marked as USDA soil texture class and a default value [36cm/hr] for layers marked as organic material, layers marked as water, bedrock, and "other" were excluded)	cm/hr
Soil max water content	max_water_content	maximum water content (combination of porosity and soil_depth_statgso, layers marked as water, bedrock, and "other" were excluded)	m
Sand fraction	sand_frac	sand fraction (of the soil material smaller than 2 mm, layers marked as organic	%

		material, water, bedrock, and "other" were excluded)	
Silt fraction	silt_frac	silt fraction (of the soil material smaller than 2 mm, layers marked as organic material, water, bedrock, and "other" were excluded)	%
Clay fraction	clay_frac	clay fraction (of the soil material smaller than 2 mm, layers marked as oragnic material, water, bedrock and "other" were excluded)	%
Carbonate rock fraction	carbonate_rocks_frac	fraction of the catchment area characterized as "Carbonate sedimentary rocks"	-
Geologic permeability	geol_permeability	subsurface permeability (log10)	m2
Mean discharge	discharge_cms	mean annual discharge	cubic meters per second
Random variable	random	a random number	NA
Num temp obs. Jan.	n_obs_1	Total number of temperature observations in the month of January for each site	NA
Num temp obs. Feb.	n_obs_2	Total number of temperature observations in the month of February for each site	NA
Num temp obs. Mar.	n_obs_3	Total number of temperature observations in the month of March for each site	NA
Num temp obs. Apr.	n_obs_4	Total number of temperature observations in the month of April for each site	NA
Num temp obs. May.	n_obs_5	Total number of temperature observations in the month of May for each site	NA
Num temp obs. Jun.	n_obs_6	Total number of temperature observations in the month of June for each site	NA
Num temp obs. Jul.	n_obs_7	Total number of temperature observations in the month of July for each site	NA
Num temp obs. Aug.	n_obs_8	Total number of temperature observations in the month of August for each site	NA
Num temp obs. Sep.	n_obs_9	Total number of temperature observations in the month of September for each site	NA
Num temp obs. Oct.	n_obs_10	Total number of temperature observations in the month of October for each site	NA
Num temp obs. Nov.	n_obs_11	Total number of temperature observations in the month of November for each site	NA
Num temp obs. Dec.	n_obs_12	Total number of temperature observations in the month of December for each site	NA

Any use of trade, firm, or product names is for descriptive purposes only and does not imply endorsement by the U.S. Government.