



HESS Opinions: Never train an LSTM on a single basin

Frederik Kratzert¹, Martin Gauch², Daniel Klotz³, and Grey Nearing⁴

¹Google Research, Vienna, Austria

²Google Research, Zurich, Switzerland

³Helmholtz Centre for Environmental Research - UFZ, Leipzig, Germany

⁴Google Research, Mountain View, California, USA

Correspondence: Frederik Kratzert (kratzert@google.com)

Abstract. Machine learning (ML) has played an increasing role in the hydrological sciences. In particular, certain types of time series modeling strategies are popular for rainfall–runoff modeling. A large majority of studies that use this type of model do not follow best practices, and there is one mistake in particular that is common: training deep learning models on small, homogeneous data sets (i.e., data from one or a small number of watersheds). In this position paper, we show that Long Short
5 Term Memory (LSTM) streamflow models are best when trained with a large amount of hydrologically diverse data.

1 Machine learning requires different intuitions about hydrological modeling

Regionalizing rainfall–runoff models across multiple watersheds is a longstanding problem in the hydrological sciences (Guo et al., 2021). The most accurate streamflow predictions from conceptual and process-based hydrological models generally require calibration to long data records in individual watersheds. Hydrology models based on machine learning (ML) are
10 different – ML models work best when trained on data from many watersheds (Nearing et al., 2021). In fact, this is one of the main benefits of ML-based streamflow modeling. Because ML models are trained with data from multiple watersheds, they are able to learn hydrologically diverse rainfall–runoff responses (Kratzert et al., 2019b) in a way that is useful for example for prediction in ungauged basins (Kratzert et al., 2019a).

Prediction in ungauged basins is not the only reason to train ML models on data from many watersheds. Models trained this
15 way have better skill even in individual, gauged watersheds with long training data records (Nearing et al., 2021), and they are also better at predicting extreme events (Frame et al., 2022).

The purpose of this position paper is to suggest a change in intuition. ML requires a top-down modeling approach, in contrast to traditional hydrological modeling that is usually most effective with a bottom-up approach. We do not mean top-down vs. bottom-up in the sense discussed by Hrachowitz and Clark (2017). Instead, we mean that traditional hydrology models (both
20 lumped conceptual models and process-based models) are typically developed, calibrated, and verified at a local scale, ideally using long and comprehensive data records from experimental watersheds. Then, in the bottom-up approach, after a model is developed, we might work on regionalization strategies to extrapolate parameters and parameterizations to larger areas (e.g., Samaniego et al., 2010; Beck et al., 2016). With ML modeling, the best approach is to start by training on all available data from



as many watersheds as possible, and then work to fine tune models for individual catchments. The effort goes into localizing
25 large scale models, instead of regionalizing small scale models.

We focus on rainfall–runoff modeling with Long Short Term Memory (LSTM) networks because this is currently the most
common type of ML model in surface hydrology. The use of LSTMs for rainfall–runoff modeling has increased exponentially
in the last several years (see Fig. 1). We see no reason to suspect that the lessons learned about big data with this type of model
are not general, and several reasons to suspect that they are (Sutton, 2019).

30 To understand the current state of practice with LSTM-based rainfall runoff modeling, we collected the top 100 papers
returned by a keyword search on Google Scholar for "*rainfall-runoff modeling LSTM streamflow*" with publication year 2021
sorted by relevance, and skipping papers that did not involve training models or developing systems for training models. Of
those 100 papers surveyed, 79 trained models on individual catchments, and 6 of the remaining 21 papers (that trained models
on multiple catchments) were co-authored by one or more of the authors of the paper that you are currently reading. We
35 collected these 100 papers for review in September, 2022, nearly three years after the original regional LSTM rainfall–runoff
modeling papers (Kratzert et al., 2019a, b) were published. The list of 100 papers is included in the published data repository.




Figure 1. Number of hydrological publications related to rainfall–runoff modeling with LSTMs over time, based on data retrieved from
Google Scholar in September 2022.

It is important to recognize that there is usually no reason in practice to train LSTM streamflow models using data from only
a small number of watersheds. There is enough publicly available streamflow data to train robust ML models. For example,
the various CAMELS datasets such as CAMELS-US (Newman et al., 2015), the Global Runoff Data Center (BAFG), or the
40 Caravan dataset and its extensions (Kratzert et al., 2023). The NeuralHydrology software package (Kratzert et al., 2022) is set
up to train ML models on multiple basins by default. It is possible to fine tune large-sample models to individual locations
and/or for specific purposes (e.g., Ma et al., 2021), however fine tuning is outside the scope of this paper. It is sufficient for our

purpose to show that training large-scale ML models is better than training small-scale ML models for streamflow prediction and fine tuning would only widen this difference.

45 2 Skill gaps between local and regional models

Figure 2 shows differences in performance between LSTM models trained on single basins vs. multiple basins (regional). The single-basin models were optimized and trained separately for 531 basins of the CAMELS data set (Addor et al., 2017) using a standard train/validation/test data split with approximately 10 years of data in each split. Details about how LSTM models were hypertuned, trained, and tested can be found in Appendix A.

50 The takeaway from Fig. 2 is that large-sample models are more accurate overall than models trained for specific locations. Other hydrological metrics are reported in Appendix C, however the main point of our argument holds regardless of which hydrological metric(s) is used for evaluation.




Figure 2. Cumulative Density Functions (CDF) of Nash–Sutcliffe Efficiencies (NSE) over watersheds in the CAMELS data set (Addor et al., 2017) for models trained on individual basins (basin) vs. on multiple basins (regional). The left-hand subplot shows a conceptual model (mHM) and a process-based model (VIC), both calibrated using data from 400+ watersheds by (different) research groups that are familiar with each model – VIC single basins: Newman et al. (2017), VIC regional: Mizukami et al. (2017), mHM single basins: Mizukami et al. (2019), and mHM regional: Rakovec et al. (2019). The right-hand subplot shows the NSE CDFs over 531 CAMELS basins for ten regional LSTM models and ten single-basin LSTM models. Randomness in the LSTM repetitions is due to randomness in the initial weights prior to training, and we recommend averaging hydrographs from this kind of repetition, as was done by Kratzert et al. (2019a) and Kratzert et al. (2019b).

3 Why this matters for extreme events

Training on large-sample data sets with hydrologic diversity means that the training envelope is larger, making it less likely
55 that any new prediction will be an extrapolation. Intuitively, the training envelope refers to the region of data where model
performance is well-supported by the training process. If the training set includes a very humid basin, then the model is more
likely to have seen large precipitation events, so that a new extreme precipitation event seen during inference is less likely to
be outside of the training envelope. As an example of this, Nearing et al. (2019) discussed how watersheds can move within
the training envelope as (e.g., climate) conditions within a catchment change, and how this causes changes in the modeled
60 rainfall–runoff response in individual watersheds.

We can look at the target data to see an example of how this diversity in training data helps. The LSTM model used by
Kratzert et al. (2019a) and Kratzert et al. (2019b) have a linear “head” layer that condenses the LSTM hidden state into a scalar
estimate of streamflow at each timestep. The LSTM hidden state is a vector that is bounded bounded to 1 (-1) and has a size
equal to the number of cell states. This means that the maximum (limiting) value of the scalar streamflow estimate from the
65 model is defined by the sum of the absolute values of weights in the linear head layer (see also Appendix B. More diversity in
training data (here, training targets) causes the model to expand the range of weights in the head layer to accommodate higher
flow values.

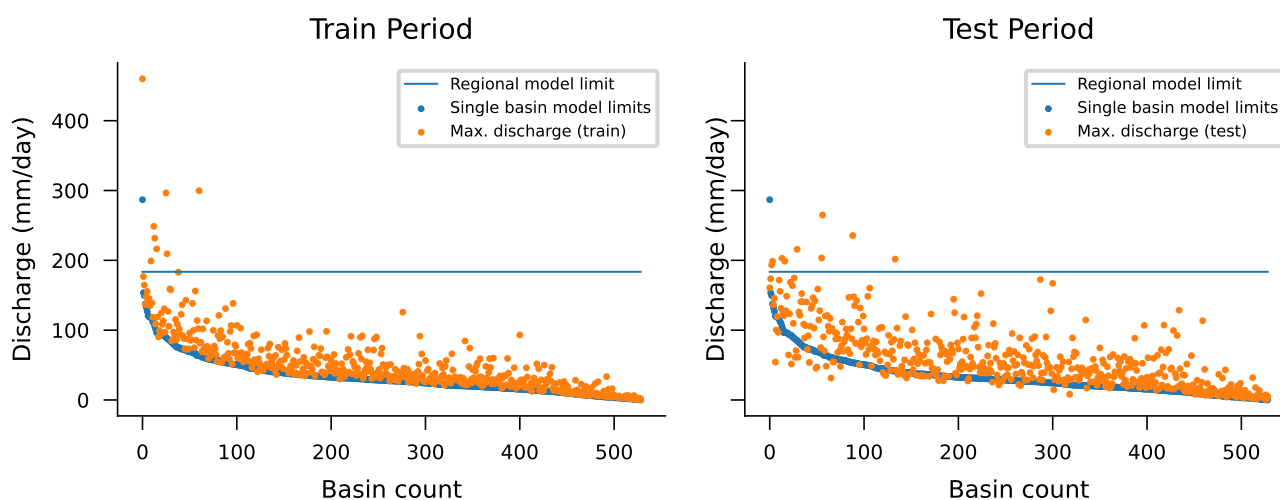


Figure 3. The theoretical maximum streamflow prediction of an LSTM model with a linear head layer when trained per basin (blue dots) and for all 531 CAMELS basins together (blue line). Orange dots represent the maximum streamflow values per basin in the train period (left) and test period (right). In the test period, there are 10 flow values above the maximum prediction of the regional model, while there are more than 6,346 flow values above the theoretical maximums of their respective single basin models across all basins and all timesteps.



This effect can be seen in Fig. 3, which shows the theoretical maximum prediction from each of 531 single basin LSTM models and from a single LSTM model trained on all 531 CAMELS basins. During inference (test period) there are in total 10 timesteps of streamflow observations across all 531 catchments that are above the regional model's theoretical maximum when trained on data from all 531 watersheds. However when separate models are trained per catchment, there are more than 6,000 timesteps of streamflow observations that are above the theoretical maximums for each model in its respective catchment. Notice that no model captures all of the extreme events, even in the training data set (which is common for physically-based models as well; Frame et al. (2022)). Figure 4 shows how this effect manifests in an example hydrograph from one particular basin (not chosen at random).

All of the LSTM models described in this paper were trained individually, and we also chose the model hyperparameters individually per model (e.g., per basin for basin-specific models), as described in Appendix A2.

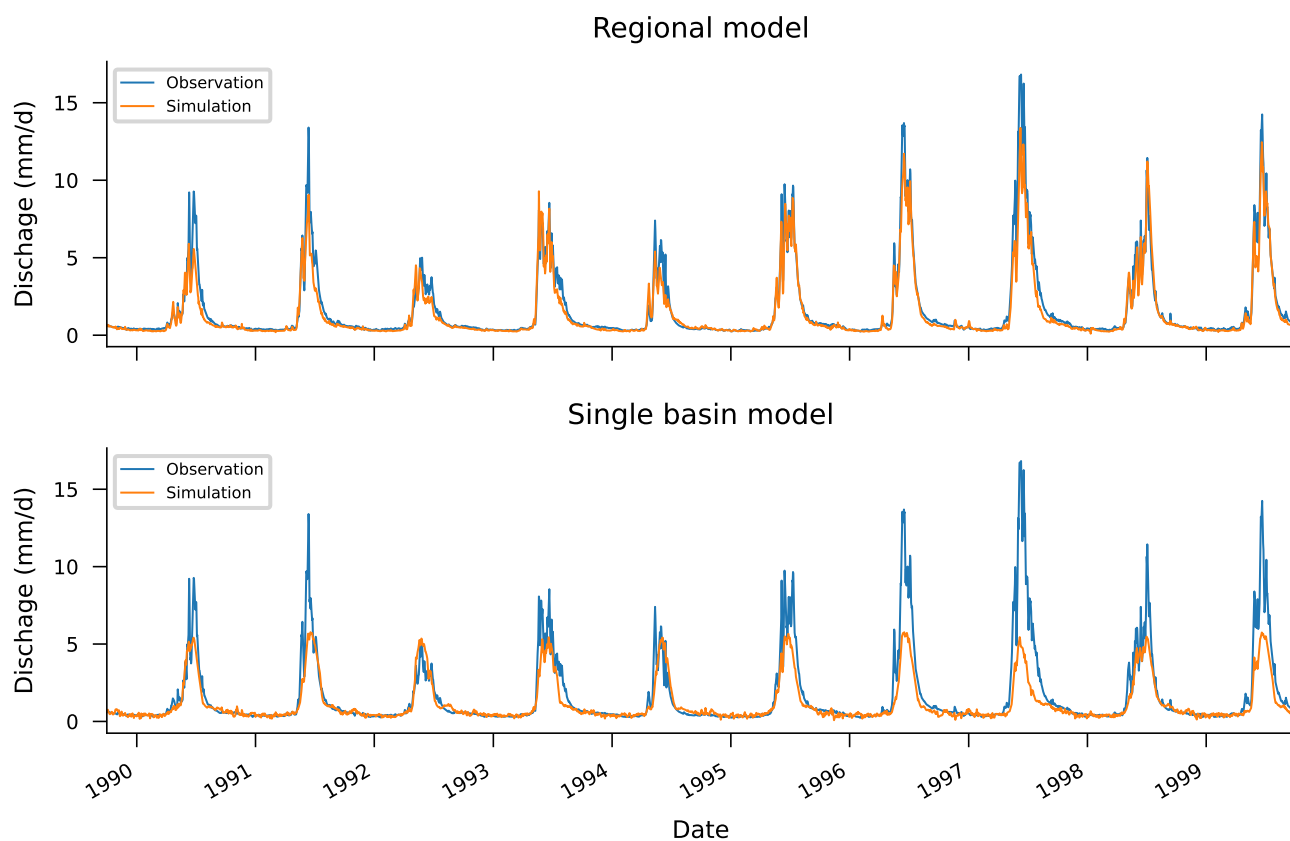


Figure 4. Observed and simulated hydrographs in a particular basin (13011900). Notice how the high-flow effect outlined in Fig. 3 manifests in the differences between hydrographs predicted by a regional LSTM vs. a single basin LSTM. This example was chosen to highlight this effect (not randomly), however the effect is similar in most basins.



4 How many basins are necessary?

Figure 5 shows how test period performance increases as more basins are added to the training set. The blue line in this figure
80 was created by splitting 531 CAMELS basins into different sized training and test sets randomly without replacement. Each
point on the blue line shows the median test-period Nash–Sutcliffe Efficiency (NSE) over all 531 basins from training and
testing models on groups of basins of a certain size (details can be found in Appendix A3). Performance continues to increase
up to the maximum size of the CAMELS data set (531 basins). In other words, even these 531 basins are most likely not enough
to train optimal LSTM models for streamflow.

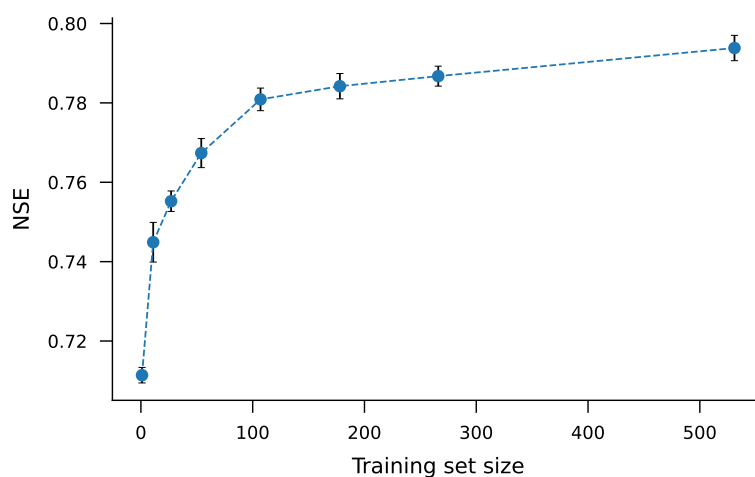


Figure 5. Median NSE scores over 531 CAMELS basins for LSTM models trained and tested on splitting the 531 CAMELS basins into different sized groupings for training and testing. All models were tested on the same basins where they were trained (during different time periods). Grouping sizes were chosen randomly without replacement so that all 531 basins are modeled exactly once for each size split. The dots represent the median NSE score across all basins, average over the ten separately-trained LSTM models. The error bars represent standard deviation of the medians over the ten repetitions. The ten separately-trained LSTMs all use the same basin splits, so that variability is due to randomness in the initial weights and biases before training.

85 5 Is hydrological diversity always an asset?

There are at least two factors at play when choosing training data: volume and variety. Volume refers to the total amount of data used for training (more is always better, as far as we have seen), and variety refers to the (hydrologic) diversity of data. Diversity might be in the form of different geophysical catchment attributes, different types and magnitudes of events, or different hydrological behaviors.



- 90 We have some evidence that there *might* be ways to construct training sets that could result in better models than simply training on all available streamflow data. We do not have results that support this directly – separating the full CAMELS data set into hydrologically similar groups results in smaller training sets that result in models that do not perform as well as a single model trained on all CAMELS data. *However*, separating the training set into hydrologically similar groups of basins results in models that perform better than models trained on random basin groups of similar size.
- 95 **It is an open question as to whether a larger data set than CAMELS (e.g., Kratzert et al., 2023) might be divisible into hydrologically similar groups that individually perform better than a model trained on all available data.** This could happen if, for example, the curve in Fig. A2 becomes asymptotic at some point beyond the size of the CAMELS data set, and if the performance of models trained on hydrologically-informed basin groups continues to increase with sample size. Note that this analysis does not account for the value of hydrologic-diversity for prediction in ungauged basins.
- 100 Figure 6 shows this effect for two basin grouping strategies. The orange line in Figure 6 shows the median NSE (over 531 CAMELS basins) from training and testing models on individual USGS hydrological unit codes (HUCs). There are 18 HUCs represented in the CAMELS data set, with between 2 and 79 basins per HUC. The green line in Figure 6 shows the median NSE (over 531 CAMELS basins) from training and testing models on basin groups derived from k-means clustering on static catchment attributes. The CAMELS data set includes catchment attributes related to climate, vegetation, pedology, geology, and topography, and we clustered using 25 catchment attributes described in Appendix A5. We selected a k-means cluster
- 105 model based on a maximin criterion on silhouette scores, which resulted in a model with 6 clusters ranging from 59 to 195 basins per cluster. Details of HUC-based training and attribute clustering are in Appendices A4 and A5.




Figure 6. Median NSE scores over 531 CAMELS basins for LSTM models trained and tested on different groupings of basins. The blue line shows results for different training set sizes selected randomly without replacement (so that all 531 basins are modeled exactly once for each size split), and is identical to the blue line in Figure A2. The orange and green lines show results from splitting the 531 basins into sets based on USGS hydrological unit codes (orange), and based on k-means clustering of basin attributes (green). Dots on the orange and green lines indicate the sizes of the (18 and 6) basin groups used in those splits, since there are a different number of CAMELS basins in each HUC and a different number of basins in each attribute cluster.



What we see from Figure 6 is that separating basins into groups based on hydrologically-relevant information causes models trained on small amounts of data to perform better than models trained on random basin groups of a similar size. However, this effect is not enough to make up for the fact that these hydrologically-informed training sets are necessarily smaller, since they require dividing up the full data set. As we said previously, it is an open question whether increasing the overall size of the data set would result in models trained on hydrologically-informed groups to become better than models trained on the largest data set possible.

That being said, it is not the case that the best model in each basin is always the model trained on all 531 basins. Figure 7 shows the number of basins for which models trained on each grouping (size, HUC, attributes cluster) perform statistically better than (green), not statistically different than (orange), or statistically worse than (blue) the regional model trained on all 531 basins. All models perform worse than the full regional model in more basins than they perform better.

We have not found a way to (reliably) predict which model will perform best in a particular basin. It is not possible to use train period or validation period metrics to (reliably) choose the best model in the test period, and we have not been able to construct a predictor model (e.g., a random forest) that uses catchment attributes and/or hydrological signatures to predict which training set will result in the best model in specific basins. The details of these predictability experiments are out-of-scope for this paper.




Figure 7. Counts of basins for which models trained on each grouping (sizes, HUC, attributes cluster) perform statistically better than (green), not statistically different than (orange), or statistically worse than (blue) the regional model trained on all 531 basins. Significance was assessed using a two-sided Wilcoxon signed-rank test over ten repetitions of each model with a significance level of $\alpha = 0.05$.

The take away is that even if enough basins exist to divide your training data into hydrologically-informed training sets, one is likely better off simply training a single model with all available data. At least, one should perform an analysis like what is



125 shown in Figure 6 to understand whether splitting the training set helps or hurts. We are interested to see (through future work) what these tradeoffs look like with larger training sets.

6 Conclusion

The main point that we would like for readers to take from this opinion paper is that training good LSTMs for rainfall-runoff modeling requires using data from many basins. We've seen a number of papers that train large ML models (LSTMs or similar) on very small data sets, and many of these papers then go on to test some type of adaptations that seems to offer improvement. Of course, it is trivial (but most likely uninteresting) to beat improperly trained models. It *would* be interesting to show that adding physics to a well-trained ML model adds information – so far, to our knowledge, this has not been done in streamflow hydrology.

135 Whatever goal a researcher might have for training an ML-based rainfall-runoff model, there is no reason not to train the model with a large-sample data set. There is enough publicly available streamflow data that there should never be an excuse not to use at least hundreds of basins for training. This is true even if the focus of a particular study is on one or a small number of watersheds.

Code and data availability. The run directories of all experiments, including model weights, simulations and pre-computed metrics are available at <https://zenodo.org/records/10139249>. The code that was used for analysing all experiments and to create all figures, based on the run directories, can be found at <https://github.com/kratzert/never-paper>. We used the open source Python package NeuralHydrology (Kratzert et al., 2022) to run all experiments. The forcing and streamflow data as well as the catchment attributes used in this manuscript are from the publicly available CAMELS dataset by Newman et al. (2015) and Addor et al. (2017). The simulations from the two hydrological models that were used to create Fig. 2 are available at <https://doi.org/10.4211/hs.474ecc37e7db45baa425cdb4fc1b61e1>.

Appendix A: Hyperparameter tuning, training, and testing

145 All models in this paper were trained using data from the CAMELS data set (Newman et al., 2015; Addor et al., 2017). Building on the community benchmarking experiment proposed by Newman et al. (2017), and used by many LSTM modeling studies (e.g., Kratzert et al., 2019a, b, 2021; Frame et al., 2021, 2022; Klotz et al., 2022; Nearing et al., 2022), we trained and tested models on 531 CAMELS basins using time periods for training (1 October 1999 through 30 September 2000), validation (1 October 1980 through 30 September 1989), and testing (1 October 1989 through 30 September 1999). All models were trained and evaluated using NeuralHydrology v.1.3.0 (Kratzert et al., 2022) with an NSE loss function. All LSTMs consist of a single layer LSTM with a linear head layer.

A1 Regional LSTM

The regional LSTM uses hyperparameters from Kratzert et al. (2021). The most important hyperparameters are:



Hidden size 256.

155 **Dropout** 40% dropout in the linear output layer.

Optimizer Adam.

Number of epochs 30.

Learning rate Initial learning rate $1e-3$, reduced to $5e-4$ at epoch 20, further reduced to $1e-4$ at epoch 25.

Sequence length 365.

160 **Loss function** Adapted NSE loss, see Kratzert et al. (2019b).

After training, we picked the weights from the epoch with the highest validation metric (median NSE across all basins) and evaluated the model with these weights on test period data from all 531 CAMELS catchments. Validation curves for all ten regional models over 30 training epochs are shown in Figure A1




Figure A1. Validation scores of all ten repetitions of the regional model over training epochs.

A2 Single basin LSTMs

165 Single basin LSTMs were trained for each basin individually, and use the same basic architecture as described in Section A1: a single layer LSTM followed by a linear head layer. Hyperparameters were tuned specifically for each basin for the experiments in this paper using the following two-step procedure. Both steps were done with a grid search:

First step: We used 3 repetitions of each hyperparameter setting for each basin ($n = 531$) with different random seeds for initializing the weights. All models were run for 100 epochs using the Adam optimizer with a learning rate of $5e-3$ and a



170 batch size of 256. During training, the model was validated after every 4 epochs on validation period data. Hyperparameters were chosen using the model settings with the highest median NSE scores over the 3 repetitions in any validation epoch.

Hidden size (8, 16, 32)

Dropout rate on the head layer (0.0, 0.2, 0.4, 0.5)

For all models, we used the same sequence length ($n = 365$) as for the regional model.

175 **Second step:** Using the hyperparameters chosen from the first step, we tuned the learning rate and batch size in a similar way, maximizing over the median NSE over 3 model repetitions:

Learning rate ($5e - 3, 1e - 3, 5e - 4, 1e - 4$)

Batch size (128, 256, 512)

180 For each basin separately, we picked model weights from the best validation epoch of the model with the highest NSE score over all validation epochs from all models in each basin.

Final training and evaluation: Given the set of per-basin optimized parameters, we trained ten models per basin, each with a distinct random seed. All statistics reported in this paper for all models are from test period data, except where otherwise noted.

A3 Random size basin splits

185 Models reported in Fig. 5 were tuned (hyperparameters chosen), trained, and tested in a way that is similar to the single basin LSTMs described in Appendix A2. For random basins splits we divided the 531 CAMELS basins into random sets without replacement using 6 different sizes of splits that were chosen by (approximately) dividing the full 531 basin group into [50, 20, 10, 5, 3, and 2] basin groups. An example of one of these random splits with 5 groups (approximately 107 basins per split) is shown in Fig. A2.

190 Choosing hyperparameters was done as described in Appendix A2, except that for these splits we did not use 3 random repetitions and only trained up to 30 epochs to reduce computational expense. We also expanded the hyperparameter search slightly due to our experience training larger models – the hyperparameter ranges for the two grid search stages were:

First step:

Hidden size (8, 16, 64, 128, 256)

195 **Dropout rate on the head layer** (0.0, 0.2, 0.4, 0.5)

Second step: In the second stage, we used a batch size of 256 always (as with the regional model), and tuned over the following multi-stage learning rates, where the index is the epoch on which the learning rate switched to the listed value:

– 0: $5e-3$, 10: $1e-3$, 25: $5e-4$

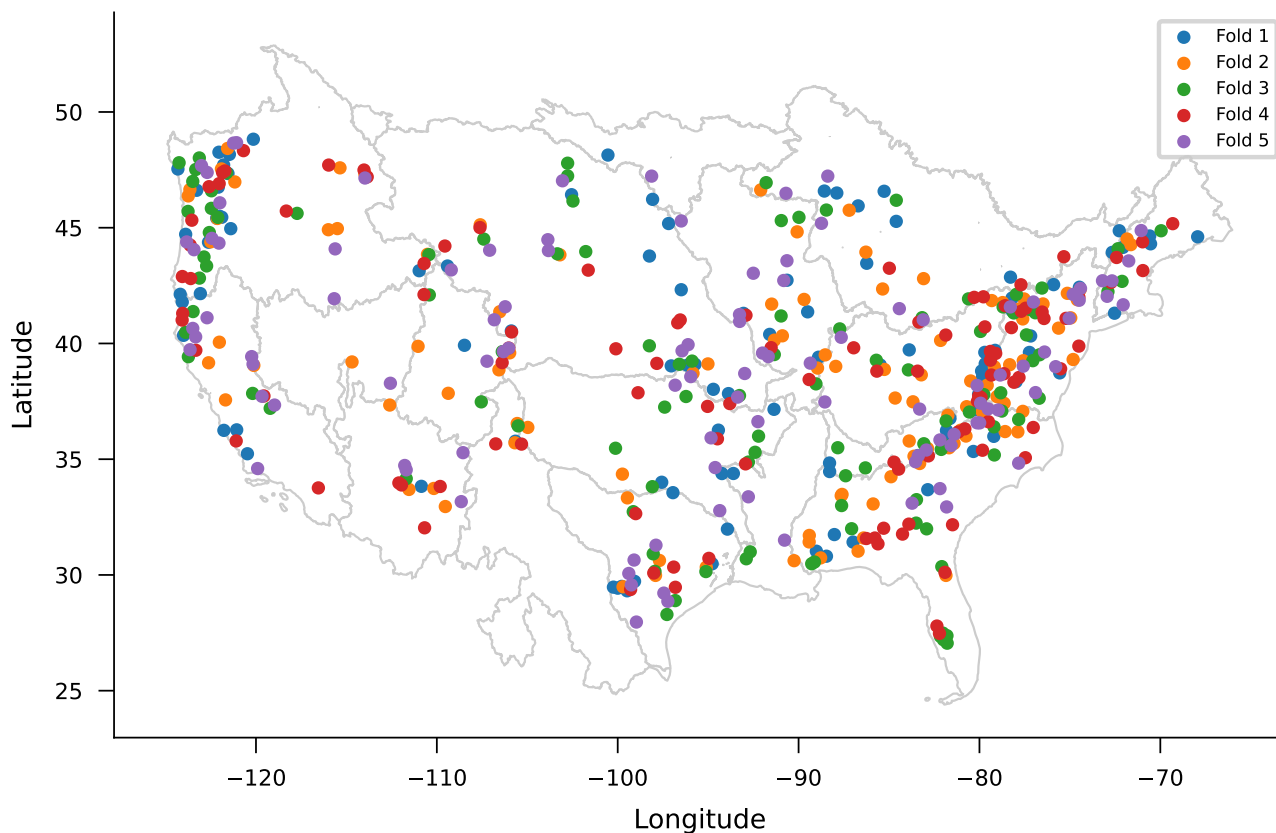


Figure A2. Basin location of random split with 5 groups of approximately 107 basins per group.

– 0: $1e-3$, 10: $5e-4$, 25: $1e-4$

200 – 0: $5e-4$, 10: $1e-4$, 25: $5e-5$

– 0: $1e-4$, 10: $5e-5$, 25: $1e-5$

A4 Hydrological unit code splits

The orange curve in Fig. 6 shows median NSE scores over CAMELS basins that result from LSTM models trained on basin groups defined by USGS HUCs. Figure A3 shows the locations of the 531 CAMELS basins by HUC region. The set of 205 531 basins was divided according to these geographical regions and a separate model was trained on all basins from each. Hyperparameter tuning was done as described in Appendix A3. Testing was done as described in Appendix A2.

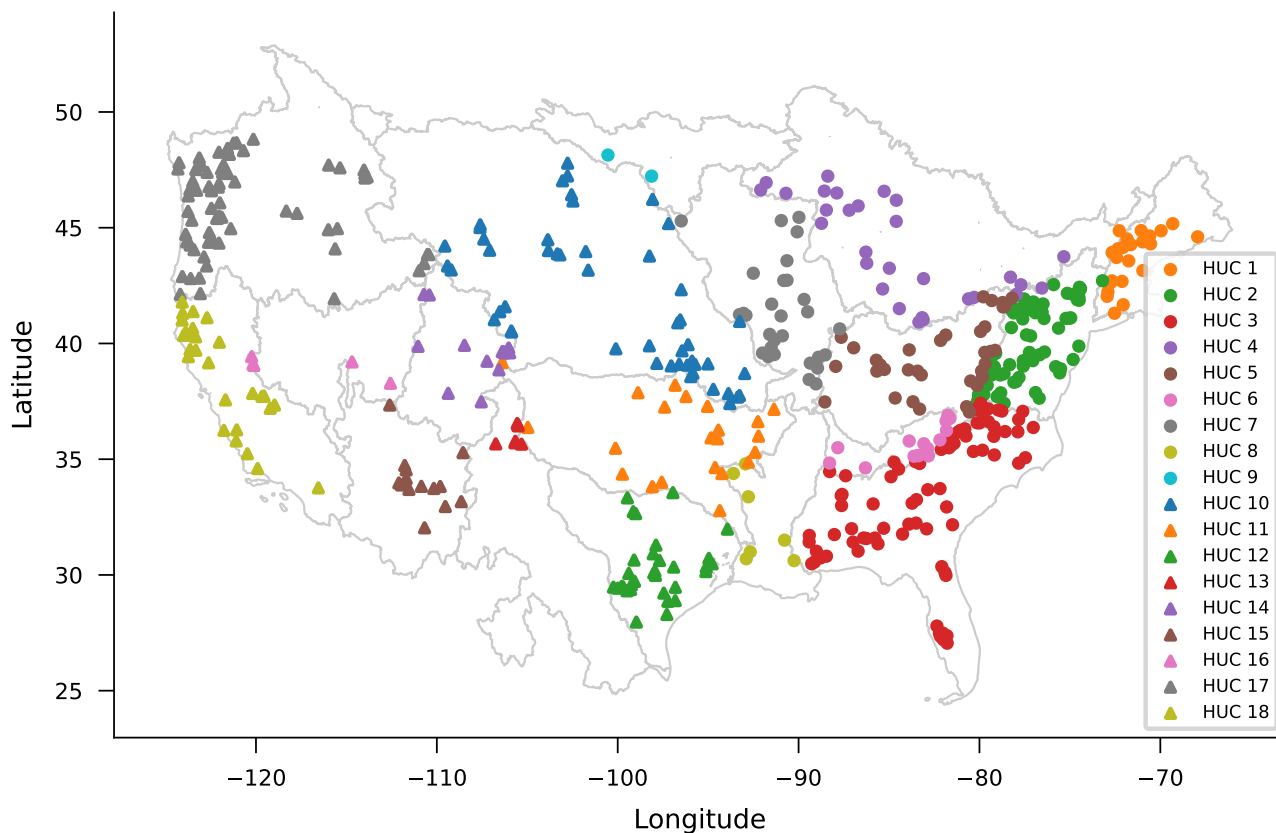


Figure A3. Spatial location of basins split by USGS hydrological unit code 02.

A5 Attribute clusters splits

The green curve in Fig. 6 shows median NSE scores over CAMELS basins that result from LSTM models trained on basin groups defined by k-means clustering based on static catchment attributes. The catchment attributes used for clustering are described in Tab. A1. These are almost the same attributes that were used by Kratzert et al. (2019b) but without carbonate rocks fraction and the seasonality of precipitation (the former is often zero and the latter is categorical, both of which make clustering slightly more difficult).

We performed k-means clustering on these 25 basin attributes (all attributes were normalized), using 300 iterations and 10 random initializations. Using a maximin criterion on silhouette scores for between 3 and 100 clusters, we chose to divide basins into 5 groups with sizes of 83, 195, 67, 61, 66, and 59 basins. These clusters are mapped in Fig. A4.

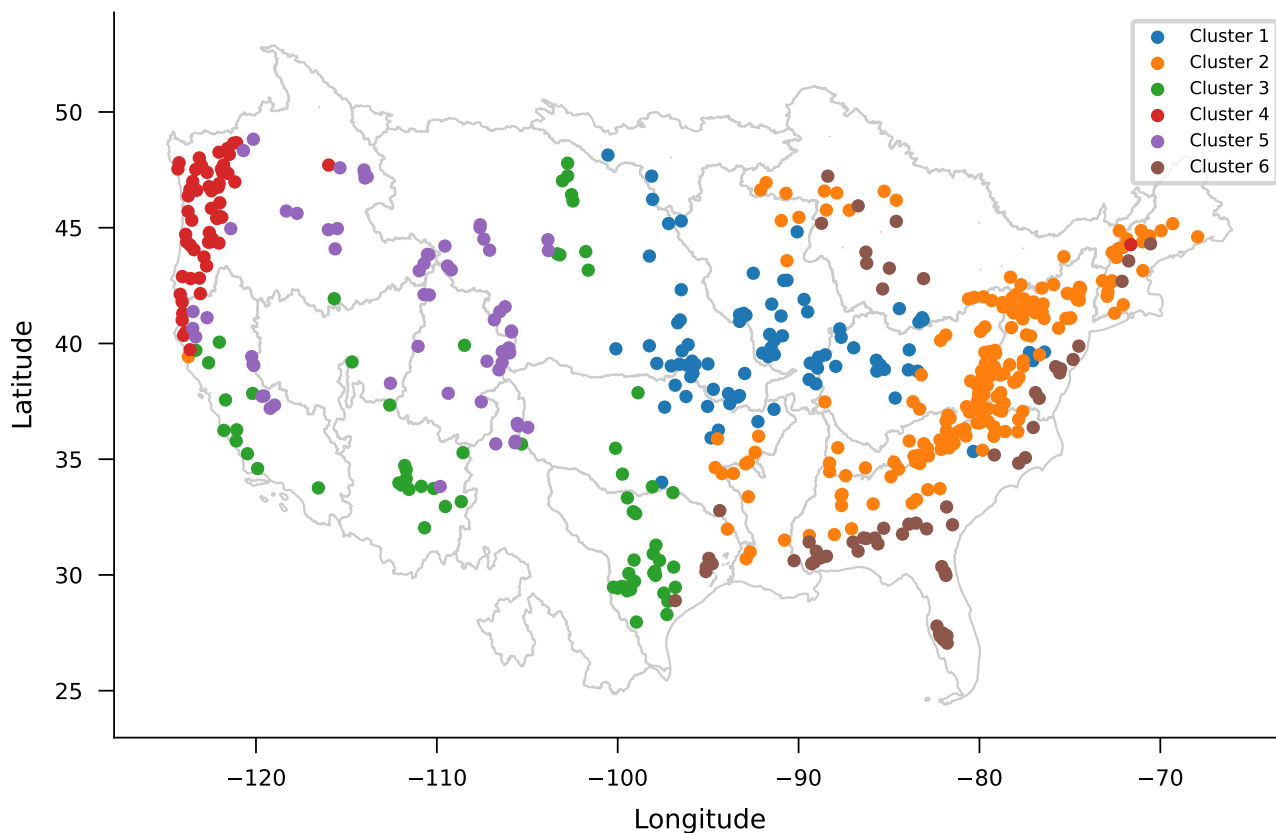


Figure A4. Spatial location of basins split by k-means clustering on the basin attributes.

Appendix B: Theoretical prediction limit

Figure 3 shows the theoretical maximum prediction limits for regional and single-basin LSTMs. To understand how those limits were derived, it is important to understand how the output of the LSTM layer is computed and how this output translates into the model prediction.

220 The output of the LSTM layer, \mathbf{h}_t , is computed according to the following equation:

$$\mathbf{h}_t = \mathbf{o}_t \odot \tanh(\mathbf{c}_t), \quad (\text{B1})$$

where \mathbf{o}_t is the output gate at time step t , $\tanh()$ is the hyperbolic tangent function and \mathbf{c}_t the LSTM cell state of time step t . The output gate is computed according to the following equation:

$$\mathbf{o}_t = \text{sigmoid}(\mathbf{W}\mathbf{x}_t + \mathbf{V}\mathbf{h}_{t-1} + \mathbf{b}), \quad (\text{B2})$$



Table A1. Table of catchment attributes used in this experiments. Description taken from the data set Addor et al. (2017).

| | |
|----------------------|--|
| p_mean | Mean daily precipitation. |
| pet_mean | Mean daily potential evapotranspiration. |
| aridity | Ratio of mean PET to mean precipitation. |
| frac_snow_daily | Fraction of precipitation falling on days with temperatures below 0°C. |
| high_prec_freq | Frequency of high precipitation days (≥ 5 times mean daily precipitation). |
| high_prec_dur | Average duration of high precipitation events (number of consecutive days with ≥ 5 times mean daily precipitation). |
| low_prec_freq | Frequency of dry days (< 1 mm/day). |
| low_prec_dur | Average duration of dry periods (number of consecutive days with precipitation < 1 mm/day). |
| elev_mean | Catchment mean elevation. |
| slope_mean | Catchment mean slope. |
| area_gages2 | Catchment area. |
| forest_frac | Forest fraction. |
| lai_max | Maximum monthly mean of leaf area index. |
| lai_diff | Difference between the max. and min. mean of the leaf area index. |
| gvf_max | Maximum monthly mean of green vegetation fraction. |
| gvf_diff | Difference between the maximum and minimum monthly mean of the green vegetation fraction. |
| soil_depth_pelletier | Depth to bedrock (maximum 50m). |
| soil_depth_statsgo | Soil depth (maximum 1.5m). |
| soil_porosity | Volumetric porosity. |
| soil_conductivity | Saturated hydraulic conductivity. |
| max_water_content | Maximum water content of the soil. |
| sand_frac | Fraction of sand in the soil. |
| silt_frac | Fraction of silt in the soil. |
| clay_frac | Fraction of clay in the soil. |
| geol_permeability | Surface permeability (log10). |

225 where $\text{sigmoid}()$ is the logistic function, \mathbf{x}_t are the input features of time step t , \mathbf{h}_{t-1} the hidden state (or LSTM output) from the previous time step $t - 1$, and \mathbf{W} , \mathbf{V} , and \mathbf{b} are learnable model parameters.

Given that $\text{sigmoid}()$ is bounded by $[0, 1]$ and $\text{tanh}()$ is bounded by $[-1, 1]$, the output of the LSTM, \mathbf{h}_t , is also bounded by $[-1, 1]$.



Finally, in the case of our model architecture, the output of the LSTM is passed through a linear layer that maps from
230 the hidden size of the LSTM to 1, the model's prediction. More formally, the model prediction \hat{y}_t at time step t is computed
according to the following equation:

$$\hat{y}_t = \mathbf{W}\mathbf{h}_t + b, \quad (\text{B3})$$

where \mathbf{W} and b are another set of learnable model parameters, specific to this linear layer. Since our model maps to a single
output value, \mathbf{W} is of shape [hidden size, 1]. With \mathbf{h}_t of shape [hidden size], we can write Eq. B3 as:

$$235 \quad \hat{y}_t = b + \sum_{i=0}^n \mathbf{W}_i * \mathbf{h}_{i,t} \quad (\text{B4})$$

Knowing that each element of \mathbf{h}_t is bounded by $[-1, 1]$ and that \mathbf{W} and b are fixed after training, the maximum possible
value a trained LSTM (of our architecture) can predict can be computed by

$$upperlimit = b + \sum_{i=0}^n \text{abs}(\mathbf{W}_i), \quad (\text{B5})$$

where $\text{abs}()$ is the function that returns the absolute value. Note that this value is in the space of training labels and if the
240 labels were normalized for training, the *upperlimit* needs to be re-transformed into discharge space to get the *upperlimit* in
e.g. mm per day.

Appendix C: Other hydrological metrics

There are a large number of metrics that hydrologists use to assess hydrograph simulations Gupta et al. (2012); Gauch et al.
(2023). Several of these metrics are described in Table C1, including bias, correlation, Nash–Sutcliffe Efficiency (NSE) (Nash
245 and Sutcliffe, 1970), Kling–Gupta Efficiency (KGE) (Gupta et al., 2009), and metrics related to hydrograph peaks. Figure C1
shows differences between regional and single-basin models for these metrics, similar to the NSE comparison shown in Fig. 2.

The takeaway from this figure is that the main message of this paper (do not train a rainfall-runoff LSTM on data from a single
basin) holds regardless of the metric(s) that we focus on. The skill differences are in correlation-based metrics (NSE, KGE,
and Pearson-R) as well as variance-based metrics (Alpha-NSE). The latter is an artifact of what we saw in Fig. ??, that training
250 on more, diverse data improves the ability of the model to predict high flows. Figure C1 shows only small improvements in the
timing and capture of hydrograph peaks (the Missed-Peaks metric measures whether a peak in the hydrograph was captured at
all, not whether the magnitude of the peak was predicted accurately). And we see little or no difference in the two bias metrics
(Beta-NSE, Beta-KGE), meaning that improvements to catchment-specific mean discharge is not strongly affected by using
training data from multiple catchments (i.e., we don't strongly bias one type of catchment by using other types of catchments
255 in training).

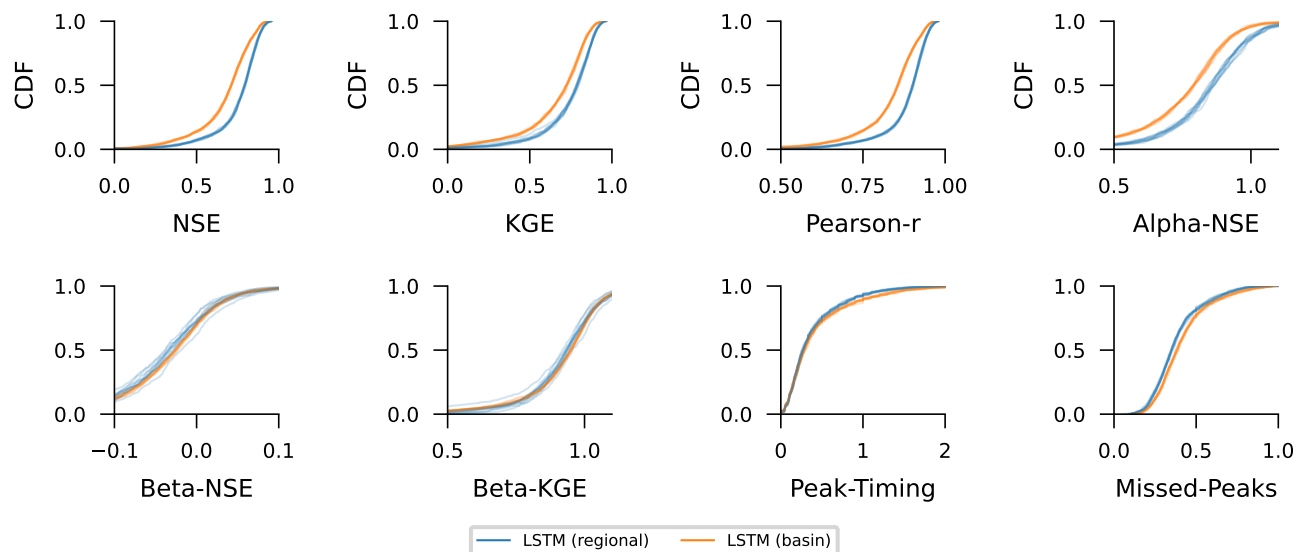


Figure C1. Comparisons between CDFs over 531 CAMELS basins of regional vs. single-basin LSTMs. This is similar to Figure 2, but for the hydrograph metrics listed in Table C1.

Table C1. A selection of standard hydrograph evaluation metrics.

| Name | Description | Reference |
|--------------|---|------------------------------------|
| NSE | Nash–Sutcliffe efficiency | Eq. 3 in Nash and Sutcliffe (1970) |
| KGE | Kling–Gupta efficiency | Eq. 9 in Gupta et al. (2009) |
| Pearson-r | Pearson correlation | Pearson (1895) |
| Alpha-NSE | Ratio of standard deviations of observed and simulated flow | From Eq. 4 in Gupta et al. (2009) |
| Beta-NSE | Bias scaled by standard deviation of observations | From Eq. 4 in Gupta et al. (2009) |
| Beta-KGE | Bias ratio: ratio of mean simulated and mean observed flow | From Eq. 10 in Gupta et al. (2009) |
| Peak-Timing | Mean time lag between observed and simulated peaks | Appendix A in Gauch et al. (2021) |
| Missed-Peaks | Fraction of hydrograph peaks that were missed | Nearing et al. (2022) |

Author contributions. FK had the initial idea for this paper. FK and GN set up all experiments. All authors contributed to the analysis of the experiment results. All authors contributed to the writing process with GN doing the majority of the writing.

Competing interests. The authors declare no competing interests.



References

- 260 Addor, N., Newman, A. J., Mizukami, N., and Clark, M. P.: The CAMELS data set: catchment attributes and meteorology for large-sample studies, *Hydrology and Earth System Sciences*, 21, 5293–5313, 2017.
BAFG: <https://www.bafg.de/GRDC>, "The Global Runoff Data Centre, 56068 Koblenz, Germany".
- Beck, H. E., van Dijk, A. I., De Roo, A., Miralles, D. G., McVicar, T. R., Schellekens, J., and Bruijnzeel, L. A.: Global-scale regionalization of hydrologic model parameters, *Water Resources Research*, 52, 3599–3622, 2016.
- 265 Frame, J. M., Kratzert, F., Raney, A., Rahman, M., Salas, F. R., and Nearing, G. S.: Post-processing the national water model with long short-term memory networks for streamflow predictions and model diagnostics, *JAWRA Journal of the American Water Resources Association*, 57, 885–905, 2021.
- Frame, J. M., Kratzert, F., Klotz, D., Gauch, M., Shalev, G., Gilon, O., Qualls, L. M., Gupta, H. V., and Nearing, G. S.: Deep learning rainfall–runoff predictions of extreme events, *Hydrology and Earth System Sciences*, 26, 3377–3392, 2022.
- 270 Gauch, M., Kratzert, F., Klotz, D., Nearing, G., Lin, J., and Hochreiter, S.: Rainfall–runoff prediction at multiple timescales with a single Long Short-Term Memory network, *Hydrology and Earth System Sciences*, 25, 2045–2062, 2021.
- Gauch, M., Kratzert, F., Gilon, O., Gupta, H., Mai, J., Nearing, G., Tolson, B., Hochreiter, S., and Klotz, D.: In Defense of Metrics: Metrics Sufficiently Encode Typical Human Preferences Regarding Hydrological Model Performance, *Water Resources Research*, 59, e2022WR033918, 2023.
- 275 Guo, Y., Zhang, Y., Zhang, L., and Wang, Z.: Regionalization of hydrological modeling for predicting streamflow in ungauged catchments: A comprehensive review, *Wiley Interdisciplinary Reviews: Water*, 8, e1487, 2021.
- Gupta, H. V., Kling, H., Yilmaz, K. K., and Martinez, G. F.: Decomposition of the mean squared error and NSE performance criteria: Implications for improving hydrological modelling, *Journal of hydrology*, 377, 80–91, 2009.
- Gupta, H. V., Clark, M. P., Vrugt, J. A., Abramowitz, G., and Ye, M.: Towards a comprehensive assessment of model structural adequacy, *Water Resources Research*, 48, 2012.
- 280 Hrachowitz, M. and Clark, M. P.: HESS Opinions: The complementary merits of competing modelling philosophies in hydrology, *Hydrology and Earth System Sciences*, 21, 3953–3973, 2017.
- Klotz, D., Kratzert, F., Gauch, M., Keefe Sampson, A., Brandstetter, J., Klambauer, G., Hochreiter, S., and Nearing, G.: Uncertainty estimation with deep learning for rainfall–runoff modeling, *Hydrology and Earth System Sciences*, 26, 1673–1693, 2022.
- 285 Kratzert, F., Klotz, D., Herrnegger, M., Sampson, A. K., Hochreiter, S., and Nearing, G. S.: Toward improved predictions in ungauged basins: Exploiting the power of machine learning, *Water Resources Research*, 55, 11344–11354, 2019a.
- 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, *Hydrology and Earth System Sciences*, 23, 5089–5110, 2019b.
- Kratzert, F., Klotz, D., Hochreiter, S., and Nearing, G. S.: A note on leveraging synergy in multiple meteorological data sets with deep learning for rainfall–runoff modeling, *Hydrology and Earth System Sciences*, 25, 2685–2703, 2021.
- 290 Kratzert, F., Gauch, M., Nearing, G., and Klotz, D.: NeuralHydrology—A Python library for Deep Learning research in hydrology, *Journal of Open Source Software*, 7, 4050, 2022.
- Kratzert, F., Nearing, G., Addor, N., Erickson, T., Gauch, M., Gilon, O., Gudmundsson, L., Hassidim, A., Klotz, D., Nevo, S., et al.: Caravan—A global community dataset for large-sample hydrology, *Scientific Data*, 10, 61, 2023.



- 295 Ma, K., Feng, D., Lawson, K., Tsai, W.-P., Liang, C., Huang, X., Sharma, A., and Shen, C.: Transferring hydrologic data across continents—
leveraging data-rich regions to improve hydrologic prediction in data-sparse regions, *Water Resources Research*, 57, e2020WR028 600,
2021.
- Mizukami, N., Clark, M. P., Newman, A. J., Wood, A. W., Gutmann, E. D., Nijssen, B., Rakovec, O., and Samaniego, L.: Towards seamless
large-domain parameter estimation for hydrologic models, *Water Resources Research*, 53, 8020–8040, 2017.
- 300 Mizukami, N., Rakovec, O., Newman, A. J., Clark, M. P., Wood, A. W., Gupta, H. V., and Kumar, R.: On the choice of calibration metrics
for “high-flow” estimation using hydrologic models, *Hydrology and Earth System Sciences*, 23, 2601–2614, 2019.
- Nash, J. E. and Sutcliffe, J. V.: River flow forecasting through conceptual models part I—A discussion of principles, *Journal of hydrology*,
10, 282–290, 1970.
- Nearing, G. S., Pelissier, C. S., Kratzert, F., Klotz, D., Gupta, H. V., Frame, J. M., and Sampson, A. K.: Physically informed machine learning
305 for hydrological modeling under climate nonstationarity, *UMBC Faculty Collection*, 2019.
- Nearing, G. S., Kratzert, F., Sampson, A. K., Pelissier, C. S., Klotz, D., Frame, J. M., Prieto, C., and Gupta, H. V.: What role does hydrological
science play in the age of machine learning?, *Water Resources Research*, 57, e2020WR028 091, 2021.
- Nearing, G. S., Klotz, D., Frame, J. M., Gauch, M., Gilon, O., Kratzert, F., Sampson, A. K., Shalev, G., and Nevo, S.: Data assimilation
and autoregression for using near-real-time streamflow observations in long short-term memory networks, *Hydrology and Earth System*
310 *Sciences*, 26, 5493–5513, 2022.
- Newman, A. J., Clark, M. P., Sampson, K., Wood, A., Hay, L. E., Bock, A., Viger, R. J., Blodgett, D., Brekke, L., Arnold, J., et al.: Develop-
ment 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*, 19, 209–223, 2015.
- Newman, A. J., Mizukami, N., Clark, M. P., Wood, A. W., Nijssen, B., and Nearing, G.: Benchmarking of a physically based hydrologic
315 model, *Journal of Hydrometeorology*, 18, 2215–2225, 2017.
- Pearson, K.: Note on Regression and Inheritance in the Case of Two Parents, *Proceedings of the Royal Society of London Series I*, 58,
240–242, 1895.
- Rakovec, O., Mizukami, N., Kumar, R., Newman, A. J., Thober, S., Wood, A. W., Clark, M. P., and Samaniego, L.: Diagnostic evaluation
of large-domain hydrologic models calibrated across the contiguous United States, *Journal of Geophysical Research: Atmospheres*, 124,
320 13 991–14 007, 2019.
- Samaniego, L., Kumar, R., and Attinger, S.: Multiscale parameter regionalization of a grid-based hydrologic model at the mesoscale, *Water*
Resources Research, 46, 2010.
- Sutton, R.: The bitter lesson, *Incomplete Ideas (blog)*, 13, 2019.