Prediction in Ungauged Basins with Long Short-Term Memory Networks

Frederik Kratzert[1], Daniel Klotz[1], Mathew Herrnegger[2], Alden K. Sampson[3], Sepp Hochreiter[1], Grey S. Nearing[4]

[1]LIT AI Lab & Institute for Machine Learning, Johannes Kepler University; Linz, Austria
[2]Institute for Hydrology and Water Management, University of Natural Resources and Life Sciences; Vienna, Austria
[3]Upstream Technologies, Public Benefit Corporation; Alameda, CA United States
[4]Department of Geological Sciences, University of Alabama; Tuscaloosa, AL United States

Key Points:

- Overall accuracy of LSTMs in ungauged basins is comparable to standard hydrology models in gauged basins
- There is sufficient information in catchment characteristics data to differentiate between catchment-specific rainfall-runoff behaviors

Corresponding author: Grey Nearing, `gsnearing@ua.edu`

Abstract

Long Short-Term Memory (LSTM) networks offer unprecedented accuracy for prediction in ungauged basins. We trained and tested an LSTM on the CAMELS basins (approximately 30 years of daily rainfall/runoff data from 531 catchments in the US of sizes ranging from 4 km^2 to 2,000 km^2) using k-fold validation, so that predictions were made in basins that supplied no training data. This effectively 'ungauged' model was benchmarked over a 15-year validation period against the Sacramento Soil Moisture Accounting (SAC-SMA) model and also against the NOAA National Water Model reanalysis. SAC-SMA was calibrated separately for each basin using 15 years of daily data (i.e., this is a 'gauged' model). The out-of-sample LSTM had higher median Nash-Sutcliffe Efficiencies across the 531 basins (0.69) than either the calibrated SAC-SMA (0.64) or the National Water Model (0.58). We outline several future research directions that would help develop this technology into a comprehensive regional hydrology model.

1 Introduction

We are firmly in the age of Machine Learning (ML). ML models currently out-perform state-of-the-art techniques at some of the most sophisticated domain problems across the Natural Sciences (e.g., AlQuraishi, 2019; He et al., 2019; Liu et al., 2016; Mayr, Klambauer, Unterthiner, & Hochreiter, 2016). In Hydrology, the first demonstration of ML out-performing a process-based model that we are aware of was by Hsu et al. (1995), who compared a calibrated Sacramento Soil Moisture Accounting Model (SAC-SMA) against a feed-forward artificial neural network across a range of flow regimes. More recently, Nearing et al. (2018) compared neural networks against the half-hourly surface energy balance of several hydrometeorological models used operationally by several international weather and climate forecasting agencies, and showed that the former generally out-performed the latter at out-of-sample FluxNet sites. Kratzert et al. (2019) showed that regionally-trained Long Short-Term Memory Networks (LSTMs) out-perform basin-specific calibrations of several traditional hydrology models.

There has been a long-standing discussion in the field about the relative merits of data-driven vs. process-driven models (e.g., Klemeš, 1986). In their summary of a recent workshop on 'Big Data and the Earth Sciences' Sellars (2018) noted that *"Many participants who have worked in modeling physical-based systems continue to raise caution about the lack of physical understanding of machine learning methods that rely on data-driven approaches."* It is often argued that data-driven models have the potential to under-perform relative to models that include explicit process representations in conditions that are dissimilar to training data (e.g., Kirchner, 2006; Milly et al., 2008; Vaze, Chiew, Hughes, & Andréassian, 2015). While this may be true, in any case where an ML model *does* out-perform against a process-based model we can conclude that the process-based model does not take advantage of the full information content of the input/output data (Nearing & Gupta, 2015). At the very least, such cases indicate that there is potential to improve the process-based model. In the Discussion section (*Section 5*) of this technical note we offer some thoughts about how the community might leverage the unprecedented ability of modern ML algorithms to find useful patterns and information in data with the decades of domain science that supports our current hydrological simulation models.

One of the situations where the accuracy of out-of-sample predictions matter is for Prediction in Ungauged Basins (PUB). PUB was the decadal problem of the International Association of Hydrological Sciences (IAHS) from 2003-2012 (Hrachowitz et al., 2013; Sivapalan et al., 2003). State-of-the-art regionalization, parameter transfer, catchment similarity, and surrogate catchment techniques (e.g., Parajka et al., 2013; Razavi & Coulibaly, 2012; Samaniego et al., 2017) result in streamflow predictions that are significantly less accurate than from models calibrated individually in gauged catchments.

Current community best-practices for PUB center fundamentally around obtaining detailed local knowledge of a particular basin (Blöschl, 2016), which is expensive for individual catchments and impossible for large-scale (e.g., continental) simulations like those from the US National Water Model (NWM) (Salas et al., 2018) or the streamflow component of the North American Land Data Assimilation System (NLDAS) (Xia et al., 2012). Moreover, reliable streamflow predictions from lumped catchment models typically require at least two to three years of gauge data for calibration (Vrugt et al., 2006). PUB remains an important challenge because the majority of streams in the world are either ungauged or poorly gauged (Goswami, Oconnor, & Bhattarai, 2007; Sivapalan, 2003), and the number of gauged catchments, even in the US, is shrinking (Fekete et al., 2015).

In this technical note, we demonstrate an ML strategy for PUB. Our results show that out-of-sample LSTMs out-perform, on average, a conceptual model calibrated independently for each catchment (SAC-SMA), and also a distributed process-based model (NWM). The purpose of this demonstration is twofold. First, to show that there is sufficient information in the available hydrological data record to provide meaningful predictions in ungauged basins - at least a significant portion of the time. Second, to show that ML offers a promising path forward for PUB. The current authors are unaware of any existing model that performs as well on average as the LSTMs that we demonstrate here in ungauged basins. At the end of this technical note we offer some thoughts - both philosophical and practical - about future work that could be done to advance the utility of ML in a complex systems science like Hydrology.

To state our primary findings succinctly, ML in *ungauged basins* out-performs, on average (i.e., in more catchments than not) a calibrated lumped model in *gauged basins*, and also a state-of-the-art distributed process-based model. This rapid correspondence is intended to highlight initial results that might motivate continued development of these and similar techniques - this is not intended to be a comprehensive analysis of the application of LSTMs to PUB, which will appear in an upcoming full-length manuscript.

2 Data

Experimental data for our analysis came from the publicly available Catchment Attributes and Meteorology for Large-Sample Studies (CAMELS) data set curated by National Center for Atmospheric Research (NCAR) (Addor, Newman, Mizukami, & Clark, 2017a; A. Newman et al., 2015, n.d.). CAMELS consists of 671 catchments in the continental US ranging in size from 4 km^2 to 2,000 km^2 (median basin size is 336 km^2). These catchments were chosen from the available gauged catchments in the US due to the fact that they are largely natural and have long gauge records (1980-2010) available from the United States Geological Survey National Water Information System. CAMELS includes daily forcing from Daymet, Maurer, and NLDAS, as well as several static catchment characteristics related to soils, climate, vegetation, topography, and geology (Addor et al., 2018). It is important to point out that these catchment characteristics were derived from maps, remote sensing products, and climate data that are generally available over the continental US and, either exactly or in close approximation, globally. For this project, we used the same 531 (of 671 total) catchments that were used for model benchmarking by A. J. Newman et al. (2017).

The CAMELS repository also includes daily streamflow values simulated by 10 SAC-SMA models calibrated separately in each catchment using Shuffled Complex Evolution (SCE) (Duan, Gupta, & Sorooshian, 1993) with 10 random seeds. Each SAC-SMA was calibrated on 15 years of data in each catchment (1980-1995). This calibration was performed as part of previous work at NCAR (A. Newman et al., 2015). We used this ensemble of SAC-SMA models as a benchmark for our LSTMs. In addition, we benchmarked against the NWM reanalysis, which spans the years 1993-2017 (`https://docs.opendata.aws/nwm-archive`). All performance statistics that we report are from the water years

117 1996-2010, so that the SAC-SMA models were tested out-of-sample in time but at the
118 same basins where they were calibrated.

3 Methods

120 LSTMs are a type of Recurrent Neural Network (RNN) first proposed by Hochre-
121 iter and Schmidhuber (1997). LSTMs have memory cells that are analgous to the states
122 of a traditional dynamical systems model, which make them potentially useful for sim-
123 ulating natural systems like watersheds. Kratzert, Klotz, Brenner, Schulz, and Herrneg-
124 ger (2018) applied LSTMs to the problem of streamflow forecasting, and later demon-
125 strated that the internal memory states of the network were highly correlated with ob-
126 served snow and soil moisture states without ever seeing snow or soil moisture data (Kratzert,
127 Herrnegger, Klotz, Hochreiter, & Klambauer, 2018).

128 The LSTMs used in this study take as inputs at each timestep the following NL-
129 DAS meteorological forcing data:

130 • 2 meter daily mean air temperature $[°C]$,
131 • precipitation $[mm/day]$,
132 • surface incident solar radiation $[W/m^2]$, and
133 • vapor pressure $[P_a]$

134 Additionally, at each timestep, the meteorological inputs were augmented with the fol-
135 lowing catchment attributes:

136 • soil depth (Pelletier),
137 • soil depth (STATSGO),
138 • soil porosity,
139 • soil conductivity,
140 • maximum water content,
141 • soil sand fraction,
142 • soil silt fraction,
143 • soil clay fraction,
144 • mean elevation,
145 • mean slope,
146 • catchment area,
147 • annual mean precipitation,
148 • annual mean potential evaporation,
149 • precipitation seasonality index,
150 • annual mean snow fraction,
151 • aridity index,
152 • frequency of high-intensity precipitation ($> 90^{th}$ percentile of annual flow),
153 • average duration of high-intensity precipitation events,
154 • frequency of low-intensity precipitation ($< 20^{th}$ percentile of annual flow),
155 • average duration of low-intensity precipitation events,
156 • forest cover fraction,
157 • annual maximum leaf area index,
158 • annual maximum greenness vegetation fraction,
159 • annual difference between maximum and minimum leaf area index,
160 • annual difference between maximum and minimum greenness vegetation fraction,
161 • carbonate rocks fraction,
162 • geological permeability.

These catchment attributes are described in detail by Addor et al. (2018); Addor, Newman, Mizukami, and Clark (2017b) and remain constant in time throughout the simulation (training and testing). In total we used 31 LSTM inputs at each daily timestep: 4 meteorological forcings and 27 catchment characteristics.

We trained and tested three types of LSTM models:

1. **Global LSTM without static features:** LSTMs with only meteorological forcing inputs, and without catchment attributes, trained on all catchments simultaneously (without k-fold validation).
2. **Global LSTM with static features:** LSTMs with both meteorological forcings and catchment characteristics as inputs, trained on all catchments simultaneously (without k-fold validation).
3. **PUB LSTM:** LSTMs with both meteorological forcings and catchment characteristics as inputs, trained and tested with k-fold validation ($k = 12$).

The third model is the one we want to test. The second model sets an upper benchmark for our PUB LSTMs - comparison against this model tells us how much information is lost due to prediction in out-of-sample basins. The first model lets us evaluate the value of adding catchment attributes to the model inputs, since these are what will, at least potentially, allow the model to be transferable between catchments. For each model type we trained and tested an ensemble of $N = 10$ LSTM models to match the 10 SCE restarts of the SAC-SMA model. All metrics reported in Section 4 are calculated from the mean of the 10-member ensembles, except for the NWM reanalysis.

All LSTM models were trained on the first 15 years of CAMELS data (1981-1995 water years) - this is the same data period used to calibrate SAC-SMA. And all models (LSTMs, SAC-SMA, and NWM) were evaluated on the last 15 years of CAMELS data (1996-2010 water years). LSTMs were trained and evaluated using a k-fold approach ($k = 12$). The training loss function was the average Nash-Sutcliffe Efficiency (NSE) over all training catchments; this is a squared-error loss function that, unlike a more traditional MSE loss funciton, does not overweight catchments with larger mean streamflow values (i.e., does not overweight large, humid catchments) (Kratzert et al., 2019).

4 Results

A comparison between interpolated frequency distributions over the NSE values from 531 CAMELS catchments from all three LSTM models and both benchmark models (SAC-SMA, NWM) is shown in **Figure 1**. Mean and median values of several performance statistics are given in **Table 1. Interpolation was done with kernel density estimation using Gaussian kernels and an optimized bandwidth.**

The primary result is that the out-of-sample PUB LSTM ensemble performed at least as well as both of the in-sample benchmarks in more than half of the catchments against all four performance metrics that we tested, except that the basin-calibrated SAC-SMA has a slightly lower average difference between 95^{th} percentile flows (both SAC-SMA and the PUB LSTM underestimated peak flows to some extent. The PUB LSTM had a higher NSE than SAC-SMA in 307 of 531 (58%) catchments, and higher than the NWM in 347 of 531 (66%) catchments. The PUB LSTM ensemble also had higher mean and maximum NSE scores than the the benchmark models, however SAC-SMA tended to out-perform the PUB LSTM in catchments with low NSE values (see the CDF plot in **Figure 1**).

There is some amount of stochasticity associated with training the LSTMs, especially through the random weight initialization of the LSTMs, but also by the weight optimization strategy (we used an ADAM optimizer (Kingma & Ba, 2014)). Because of this,

Figure 1. Frequencies of NSE values from 531 catchments given by 'gauged' and 'ungauged' LSTMs, calibrated (gauged) SAC-SMA, and the National Water Model reanalysis.

the LSTM-type models give better predictions when used as an ensemble. It is not necessarily the case that if one particular LSTM model performs poorly in one catchment that a different LSTM trained on *exactly the same data* will also perform poorly. In our case, we used an ensemble of $N = 10$ (the same size as the SAC-SMA ensemble developed by A. Newman et al. (2015) that was used here for benchmarking). **Figure 3** shows the NSE values for each ensemble member for the PUB LSTM models. In total, there were 103 basins with at least one PUB LSTM ensemble member with an NSE score of below zero. Only 9 of these 103 basins have all $N = 10$ ensemble members with NSE < 0, while 55 of the 103 have at least one ensemble member with NSE > 0.5. As an example, one of the basins (USGS basin ID: 01142500, which is basin number 232 in **Figure 3**) had 9 of 10 ensemble members with NSE < 0, but one ensemble member with NSE > 0.7. This indicates that a substantial portion of the uncertainty in these LSTM models is due to randomness, rather than to systematic model structural error.

The global LSTM model with static catchment attributes performs better than all other models against the metrics that we tested. **Figure 2** compares the performance of the Global LSTM with other benchmark models (all except the PUB LSTM) in all catchments. The Global LSTM performs better in most - but not all - catchments. This indicates two things. First, the comparison between the Global LSTM with and without static catchment attributes indicates that although there is useful information in the catchment attributes, in some catchments having this data actually hurts us. This indicates a need for future work to understand how uncertainty in catchment attributes data can be quantified and mitigated in this context. Second, the comparison between the Global LSTM and SAC-SMA indicates that there is substantial room to improve SAC-SMA overall. This is probably not a surprise to any working Hydrologist, but this analysis clearly shows that the LSTM is finding rainfall-runoff relationships in individual catchments that SAC-SMA cannot emulate. However, the fact that SAC-SMA performs better in some catchments indicates the potential value of having physical constraints in a hydrological model. The LSTM in these cases are either overfitted or are not able to behaviors of certain similar catchments in the training data set.

5 Discussion

The results illustrated in the previous section tell us three things:

1. The process-driven hydrology models that we used here as benchmarks could be improved. The LSTM often finds a better functional representation of rainfall-runoff behavior in most catchments than either SAC-SMA or the NWM.
2. The argument that process-driven models may be preferable in out-of-sample conditions may not hold water. Modern ML methods are quite powerful at extract-

Table 1. Summary of Benchmark Statistics for All Models across 531 Catchments

	Median	Mean	Minimum	Maximum
Nash Sutcliffe Efficiency:	$(-\infty, 1]$ – *values close to 1 are desirable.*			
SAC-SMA:	0.64	0.51	-12.28	0.88
NWM:	0.58	0.31	-20.28	0.89
Global LSTM (no statics):	0.63	0.45	-31.72	0.90
Global LSTM (with statics):	0.74	0.68	-1.78	0.93
PUB LSTM:	0.69	0.54	-13.02	0.90
Fractional Bias:	$(-\infty, 1]$ – *values close to 0 are desirable.*			
SAC-SMA:	0.04	0.02	-1.76	0.71
NWM:	0.05	-0.01	-4.80	1.00[c]
Global LSTM (no statics):	0.01	-0.03	-3.01	0.77
Global LSTM (with statics):	-0.01	-0.04	-2.19	0.49
PUB LSTM:	-0.02	-0.09	-4.86	0.72
Standard Deviation Ratio[a]:	$[0, \infty)$ – *values close to 1 are desirable.*			
SAC-SMA:	0.83	0.87	0.10	3.76
NWM:	0.86	0.93	0.00[c]	4.04
Global LSTM (no statics):	0.74	0.81	0.10	5.83
Global LSTM (with statics):	0.88	0.89	0.17	1.96
PUB LSTM:	0.86	0.91	0.10	3.23
95th Percentile Difference[b]:	$(-\infty, 1]$ – *values close to 0 are desirable.*			
SAC-SMA:	0.02	-0.05	-3.98	0.83
NWM:	0.07	-0.07	-8.59	1.00[c]
Global LSTM (no statics):	0.12	0.02	-4.97	0.81
Global LSTM (with statics):	0.03	-0.03	-3.30	0.63
PUB LSTM:	0.03	-0.08	-5.26	0.78

[a] *Ratio of the standard deviation of simulated vs. observed flows at each catchment.*
[b] *Difference between the values of the observed vs. simulated 95^{th} percentile flows divided by the observed 95^{th} percentile flows at each catchment.*
[c] *Values of zero and one in the NWM max/min statistics are due to rounding. In particular, for one basin (USGS basin ID: 2108000) the NWM simulates a 95^{th} flow percentile of $\sim 1 \times 10^{-3} [mm/day]$ whereas the 95^{th} percentile of observed flow is $\sim 4 [mm/day]$*

247   ing information from large, diverse data sets under a variety of hydrological con-
248   ditions.
249   3. The comparison between models with and without static catchment attributes as
250   inputs demonstrates that there is sufficient information contained in catchment
251   attribute data to distinguish between different rainfall-runoff relationships in at
252   least most of the US catchments that we tested.

253   Related to the third conclusion, the challenge going forward is about how to ex-
254 tract the useful information from catchment attributes data for regional modeling. One
255 of the historical reasons why this has been a hard problem is because we have often tried
256 to use observable catchment attributes or characteristics to identify parameters of con-
257 ceptual or process-based simulation models (e.g., Blöschl, 2016; Hrachowitz et al., 2013;
258 Prieto, Le Vine, Kavetski, Garca, & Medina, 2019). This is hard because of strong in-
259 teractions in high-dimensional parameter spaces (Beven & Freer, 2001; **?**). There are many
260 strategies for circumventing this issue - notably a family of methods that Razavi and Coulibaly
261 (2012) called 'model independent' methods, however we are unaware of any approach
262 that is as effective as LSTMs at extracting this information for streamflow simulation.
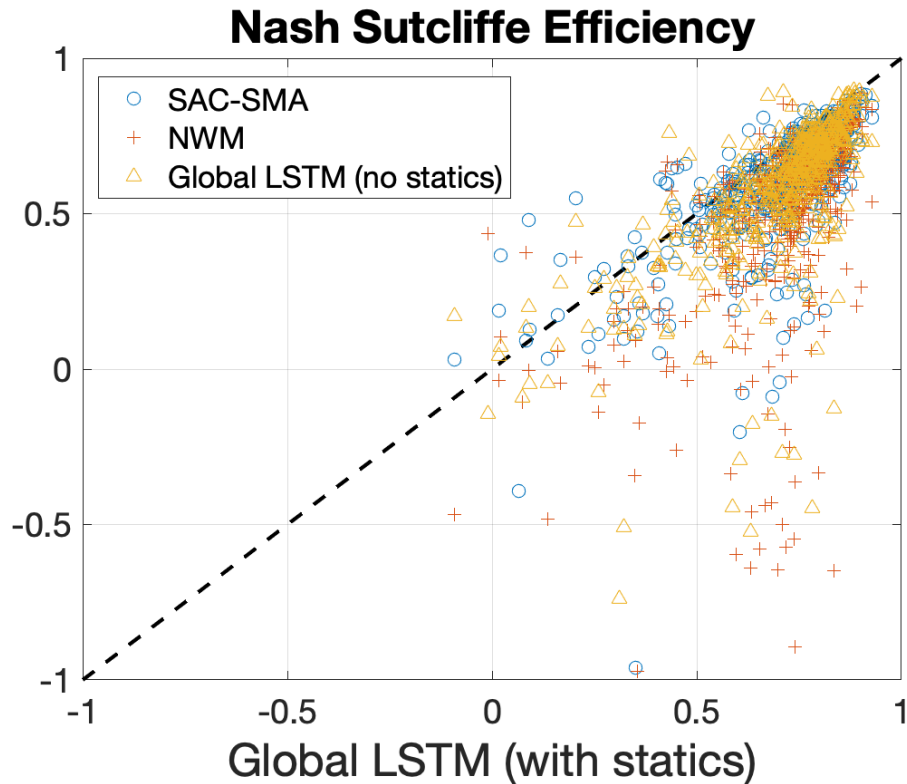263 Kratzert et al. (2019) compared similar LSTMs with a state-of-the-art regionalization

Figure 2. Comparison between the Global LSTM model and all other models used in this study (except the PUB LSTM).

method in gauged catchments. Again, the catchment attribute data we used here were derived from maps available over the entire continental US (Addor et al., 2017a).

Related to the first and second conclusion, the Hydrological Sciences community must take these results seriously. We have a long history in our community of preferring process-based insight and process-based models over black-box or data driven models (e.g., Kirchner, 2006; Klemeš, 1986; Milly et al., 2008; Vaze et al., 2015), but the reality of the situation is that these models often perform worse than data-driven models, even out-of-sample. We would argue that the hydrological modeling community does not currently have a robust understanding of where, when, and under what conditions process-based insight is useful or necessary for hydrological prediction. At present, the arguments are largely philosophical. For our discipline to continue to provide value to operational forecasting efforts, the onus is on us to clearly delineate (i) cases and situations where process-based insight is critical for accurate modeling (e.g., perhaps under climate change), and (ii) how to best use this type of process-based insight in modeling systems that *also* utilize the now-undeniable power of sophisticated methods for learning from large data sets.

Theory-guided data-science, sensu Karpatne et al. (2017), is the idea of using the scientific insights and knowledge to augment and guide data-driven methods like machine learning. Examples for this line of research are physically bounded black-boxes, hybrid-modelling approaches or, as shown here, model benchmarking and/or hypothesis testing. Optimally, a reciprocal relationship will form over time in which data-driven methods inform the development of "classical" methods and vice-versa. We encourage the hy-
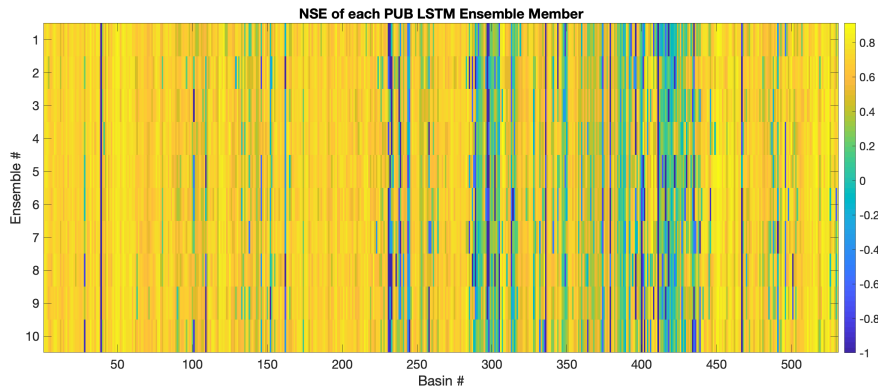
Figure 3. NSE scores for all PUB LSTM ensemble members. In some number of basins, certain ensemble members perform well and certain ensemble members perform poorly. This motivates the use of ensembles of LSTMs.

drological modelling community to adapt these techniques and hope to provide a motivating example for accelerating their use in Hydrology.

6 Code and Data Availability

CAMELS data, including SAC-SMA simulations, are available from NCAR at `https://ral.ucar.edu/solutions/products/camels`. National Water Model reanalysis data are available form the NOAA Big Data Repository at `https://registry.opendata.aws/nwm-archive/`. All code used for this project is available at `https://github.com/kratzert/lstm_for_pub`.

References

Addor, N., Nearing, G., Prieto, C., Newman, A., Le Vine, N., & Clark, M. (2018). A ranking of hydrological signatures based on their predictability in space. *Water Resources Research*, *54*(11), 8792–8812.

Addor, N., Newman, A., Mizukami, N., & Clark, M. P. (2017b). *Catchment attributes for large-sample studies.* `https://doi.org/10.5065/D6G73C3Q`. doi: 10.5065/D6G73C3Q

Addor, N., Newman, A. J., Mizukami, N., & Clark, M. P. (2017a). The camels data set: catchment attributes and meteorology for large-sample studies. *Hydrology and Earth System Sciences (HESS)*, *21*(10), 5293–5313.

AlQuraishi, M. (2019). End-to-end differentiable learning of protein structure. *Cell systems*, *8*(4), 292–301.

Beven, K., & Freer, J. (2001). Equifinality, data assimilation, and uncertainty estimation in mechanistic modelling of complex environmental systems using the

315 glue methodology. *Journal of Hydrology*, *249*(1), 11 - 29.

316 Blöschl, G. (2016). Predictions in ungauged basins–where do we stand? *Proceedings*
317 *of the International Association of Hydrological Sciences*, *373*, 57–60.

318 Duan, Q., Gupta, V. K., & Sorooshian, S. (1993). Shuffled complex evolution ap-
319 proach for effective and efficient global minimization. *Journal of optimization*
320 *theory and applications*, *76*(3), 501–521.

321 Fekete, B. M., Robarts, R. D., Kumagai, M., Nachtnebel, H.-P., Odada, E., & Zhuli-
322 dov, A. V. (2015). Time for in situ renaissance. *Science*, *349*(6249), 685–686.

323 Goswami, M., Oconnor, K., & Bhattarai, K. (2007). Development of regionalisation
324 procedures using a multi-model approach for flow simulation in an ungauged
325 catchment. *Journal of Hydrology*, *333*(2-4), 517–531.

326 He, S., Li, Y., Feng, Y., Ho, S., Ravanbakhsh, S., Chen, W., & Póczos, B. (2019).
327 Learning to predict the cosmological structure formation. *Proceedings of the*
328 *National Academy of Sciences*, 201821458.

329 Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural compu-*
330 *tation*, *9*(8), 1735–1780.

331 Hrachowitz, M., Savenije, H., Blöschl, G., McDonnell, J., Sivapalan, M., Pomeroy, J.,
332 . . . others (2013). A decade of predictions in ungauged basins (pub)a review.
333 *Hydrological sciences journal*, *58*(6), 1198–1255.

334 Hsu, K.-l., Gupta, H. V., & Sorooshian, S. (1995). Artificial neural network mod-
335 eling of the rainfall-runoff process. *Water resources research*, *31*(10), 2517–
336 2530.

337 Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing In Science*
338 *& Engineering*, *9*(3), 90–95.

339 Karpatne, A., Atluri, G., Faghmous, J. H., Steinbach, M., Banerjee, A., Ganguly,
340 A., . . . Kumar, V. (2017). Theory-guided data science: A new paradigm for
341 scientific discovery from data. *IEEE Transactions on Knowledge and Data*
342 *Engineering*, *29*(10), 2318–2331.

343 Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv*
344 *preprint arXiv:1412.6980*.

345 Kirchner, J. W. (2006). Getting the right answers for the right reasons: Linking
346 measurements, analyses, and models to advance the science of hydrology. *Wa-*
347 *ter Resources Research*, *42*(3).

348 Klemeš, V. (1986). Dilettantism in hydrology: Transition or destiny? *Water Re-*
349 *sources Research*, *22*(9S), 177S–188S.

350 Kratzert, F., Herrnegger, M., Klotz, D., Hochreiter, S., & Klambauer, G. (2018). Do
351 internals of neural networks make sense in the context of hydrology? In *Pro-*
352 *ceedings of the 2018 agu fall meeting*.

353 Kratzert, F., Klotz, D., Brenner, C., Schulz, K., & Herrnegger, M. (2018).
354 Rainfall–runoff modelling using long short-term memory (lstm) networks.
355 *Hydrology and Earth System Sciences*, *22*(11), 6005–6022. doi: 10.5194/
356 hess-22-6005-2018

357 Kratzert, F., Klotz, D., Shalev, G., Klambauer, G., Hochreiter, S., & Nearing, G.
358 (2019). Benchmarking a catchment-aware long short-term memory network
359 (lstm) for large-scale hydrological modeling. *arXiv preprint arXiv:1907.08456*.

360 Liu, Y., Racah, E., Correa, J., Khosrowshahi, A., Lavers, D., Kunkel, K., . . . oth-
361 ers (2016). Application of deep convolutional neural networks for detecting
362 extreme weather in climate datasets. *arXiv preprint arXiv:1605.01156*.

363 Mayr, A., Klambauer, G., Unterthiner, T., & Hochreiter, S. (2016). Deeptox: toxic-
364 ity prediction using deep learning. *Frontiers in Environmental Science*, *3*, 80.

365 McKinney, W. (2010). Data Structures for Statistical Computing in Python. *Pro-*
366 *ceedings of the 9th Python in Science Conference*, *1697900*(Scipy), 51–56.

367 Milly, P. C. D., Betancourt, J., Falkenmark, M., Hirsch, R. M., Kundzewicz, Z. W.,
368 Lettenmaier, D. P., & Stouffer, R. J. (2008). Stationarity is dead: Whither
369 water management? *Science*, *319*(5863), 573–574.

Nearing, G. S., & Gupta, H. V. (2015). The quantity and quality of information in hydrologic models. *Water Resources Research*, *51*(1), 524–538.

Nearing, G. S., Ruddell, B. L., Clark, M. P., Nijssen, B., & Peters-Lidard, C. (2018). Benchmarking and process diagnostics of land models. *Journal of Hydrometeorology*, *19*(11), 1835–1852.

Newman, A., Clark, M., Sampson, K., Wood, A., Hay, L., Bock, A., . . . others (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*, *19*(1), 209–223.

Newman, A., Sampson, K., Clark, M. P., Bock, A., Viger, R. J., & D., B. (n.d.). *A large-sample watershed-scale hydrometeorological dataset for the contiguous usa.* https://dx.doi.org/10.5065/D6MW2F4D. doi: 10.5065/D6MW2F4D

Newman, A. J., Mizukami, N., Clark, M. P., Wood, A. W., Nijssen, B., & Nearing, G. (2017). Benchmarking of a physically based hydrologic model. *Journal of Hydrometeorology*, *18*(8), 2215–2225.

Parajka, J., Viglione, A., Rogger, M., Salinas, J., Sivapalan, M., & Blöschl, G. (2013). Comparative assessment of predictions in ungauged basins–part 1: Runoff-hydrograph studies. *Hydrology and Earth System Sciences*, *17*(5), 1783–1795.

Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., . . . Lerer, A. (2017). Automatic differentiation in pytorch.

Prieto, C., Le Vine, N., Kavetski, D., Garca, E., & Medina, R. (2019). Flow prediction in ungauged catchments using probabilistic random forests regionalization and new statistical adequacy tests. *Water Resources Research*, *55*(5), 4364-4392.

Razavi, T., & Coulibaly, P. (2012). Streamflow prediction in ungauged basins: review of regionalization methods. *Journal of Hydrologic Engineering*, *18*(8), 958–975.

Salas, F. R., Somos-Valenzuela, M. A., Dugger, A., Maidment, D. R., Gochis, D. J., David, C. H., . . . Noman, N. (2018). Towards real-time continental scale streamflow simulation in continuous and discrete space. *JAWRA Journal of the American Water Resources Association*, *54*(1), 7–27.

Samaniego, L., Kumar, R., Thober, S., Rakovec, O., Zink, M., Wanders, N., . . . others (2017). Toward seamless hydrologic predictions across spatial scales. *Hydrology and Earth System Sciences*, *21*(9), 4323–4346.

Sellars, S. (2018). grand challenges in big data and the earth sciences. *Bulletin of the American Meteorological Society*, *99*(6), ES95–ES98.

Sivapalan, M. (2003). Prediction in ungauged basins: a grand challenge for theoretical hydrology. *Hydrological Processes*, *17*(15), 3163–3170.

Sivapalan, M., Takeuchi, K., Franks, S., Gupta, V., Karambiri, H., Lakshmi, V., . . . others (2003). Iahs decade on predictions in ungauged basins (pub), 2003–2012: Shaping an exciting future for the hydrological sciences. *Hydrological sciences journal*, *48*(6), 857–880.

Van Der Walt, S., Colbert, S. C., & Varoquaux, G. (2011). The NumPy array: A structure for efficient numerical computation. *Computing in Science and Engineering*, *13*(2), 22–30.

van Rossum, G. (1995). *Python tutorial, Technical Report CS-R9526* (Tech. Rep.). Amsterdam: Centrum voor Wiskunde en Informatica (CWI).

Vaze, J., Chiew, F., Hughes, D., & Andréassian, V. (2015). Preface: Hs02–hydrologic non-stationarity and extrapolating models to predict the future. *Proceedings of the International Association of Hydrological Sciences*, *371*, 1–2.

Vrugt, J. A., Gupta, H. V., Dekker, S. C., Sorooshian, S., Wagener, T., & Bouten, W. (2006). Application of stochastic parameter optimization to the sacramento

425   soil moisture accounting model. *Journal of Hydrology*, *325*(1-4), 288–307.

426   Xia, Y., Mitchell, K., Ek, M., Cosgrove, B., Sheffield, J., Luo, L., . . . others   (2012).

427   Continental-scale water and energy flux analysis and validation for north amer-

428   ican land data assimilation system project phase 2 (nldas-2): 2. validation of

429   model-simulated streamflow.   *Journal of Geophysical Research: Atmospheres*,

430   *117*(D3).