

Comments/Text of Anonymous Referee 1 (AR1) posted in blue, our text in black with old passages in red and the new passage in green.

Summary: This paper utilizes a data-driven approach, based on recurrent neural networks, to model rainfall-runoff relationships. A novel method is applied to model runoff in catchments in the continental U.S where gage data and meteorological forcings are available, and results are compared with existing process-based model results which are used as a benchmark. The LSTM method presented is tested through various experiments where the network is either trained for individual catchments, large aggregated regional catchments, or a combination approach where models are initialized based on large catchments and then “fine-tuned” to smaller catchments. This study is presented to introduce LSTM as an efficient hydrological modelling approach that is shown to provide similar quality predictions as an existing process-based model.

Novelty: The novelty of this paper is in the LSTM network approach, which is an improvement over other types of data driven approaches in its capacity to retain longer time dependencies. The results indicate that this type of model, when adequately trained, provides similar results as a benchmark model and may be useful to estimate runoff in ungauged catchments. The experiments are generally well-described and organized. Overall this is an interesting study that is appropriate for the journal, but I have several comments and suggestions detailed below. They involve the description and advantages of the methodology, linking with existing knowledge of the basins in the study, and suggestions for re-organization.

We thank Anonymous Referee 1 (AR1) for the general evaluation and feedback. The constructive comments and remarks have made us rethink and reflect our results and findings in more detail and we sincerely believe that this has significantly improved the revised manuscript. We will respond to specific questions and comments in some detail and will indicate how we are going to make changes to the manuscript in the following.

Comments:

1. In Section 2.1, it is mentioned that the LSTM overcomes the weakness of traditional RNNs to learn long-term dependencies. This seems to be addressed in the additional cell state that stores or “forgets” long-term dependencies. However, it is not clear what the difference would be, for example in a hydrological application, between the two methods. It would be helpful to include a “traditional” or more simple RNN model to the LSTM model on the study dataset to show how this capacity for long-term storage comes into play.

In general, I recommend to expand the description of the methods, particularly the significance of the forget, input, output gates, and hidden states. As it is, readers will have to dig back through 2 cited papers or further on the LSTM method, and I think that a few sentences within this section could go a long way to help interpret what is going on.

Regarding the remark, concerning a comparison of LSTM and RNN: We did not include any comparison in the first submission, because it is a proven and known fact that the traditional RNN can not learn dependencies of more than approx. 10 time steps (the phenomenon is referred to as “vanishing-” or “exploding gradients”, see Bengio et al. 1994 and Hochreiter and Schmidhuber 1997). However, we agree that it is interesting to see, what this means for hydrological applications, since they were already applied in some studies in the field of hydrology (Carriere et al., 1996; Hsu et al., 1997; Kumar et al., 2004): We know from hydrological science that there are many catchment processes, which can have dependencies of far more than 10 days (which corresponds to 10 time steps here), e.g. snow accumulation and snow melt. Modelling these processes correctly is inevitable for the correct prediction of the river discharge, at least for traditional hydrological modelling. However, in principle it is not said that this must be similar for a data driven approach.

We therefore added a comparison of RNN vs LSTM at the beginning of the results and discussion section, showing the effect of (not) learning long-term dependencies with an explicit example. We believe that adding the following new section and additionally a pseudo-code to the manuscript (see answer to comment 5, AR2) also highlights the significance of the forget, input, output gates, and hidden states as mentioned by the reviewer.

New Section:

3.1 The effect of (not) learning long-term dependencies

As stated in Sect. 2.1, the traditional RNN can only learn dependencies of 10 or less time steps. The reason for this is the so-called “vanishing or exploding gradients” phenomenon (see Bengio et al. (1994) and Hochreiter and Schmidhuber (1997)), which manifests itself in an error signal during the backward pass of the network training that either diminishes towards zero or grows against infinity, preventing the effective learning of long-term dependencies. However, from the perspective of hydrological modelling a catchment contains various processes with dependencies well above 10 days (which corresponds to 10 time steps in the case of daily streamflow modelling), e.g. snow accumulation during winter and snow melt during spring and summer. Traditional hydrological models need to reproduce these processes correctly in order to be able to make accurate streamflow predictions. This is in principle not the case for data-driven approaches.

To empirically test the effect of (not) being able to learn long-term dependencies, we compared the modelling of a snow influenced catchment (basin 13340600 of the Pacific Northwest region) with a LSTM and a traditional RNN. For this purpose we adapted the number of hidden units of the RNN to be 41 for both layers (so that the number of learnable parameters of the LSTM and RNN is approximately the same). All other

modelling boundary conditions , e.g. input data, the number of layers, dropout rate, number of training epochs, are kept identical.

Figure 6a shows two years of the validation period of observed discharge as well as the simulation by LSTM and RNN. We would like to highlight three points: (i) The hydrograph simulated by the RNN has a lot more variance compared to the smooth line of the LSTM. (ii) The RNN underestimates the discharge during the melting season and early summer, which is strongly driven snow melt and by the precipitation that has fallen through the winter months. (iii) In the winter period, the RNN systematically overestimates observed discharge, since snow accumulation is not accounted for. These simulation deficits can be explained by the lack of the RNN to learn and store long-term dependencies, while especially the last two points are interesting and connected. Recall that the RNN is trained to minimize the average RMSE between observation and simulation. The RNN is not able to store the amount of water which has fallen as snow during the winter and is, in consequence, also not able to generate sufficient discharge during the time of snow melt. The RNN, trained to minimize the average RMSE, therefore overestimates the discharge most time of the year by a constant bias and underestimates the peak flows, thus being closer to predicting the mean flow. Only for a short period at the end of the summer, it is close at predicting the low flow correctly.

In contrast, the LSTM seems to have (i) no or less problems with predicting the correct amount of discharge during the snowmelt season and (ii) the predicted hydrograph is much smoother and fits the general trends of the hydrograph much better. Note that both networks are trained with the exact same data and have the same data available for predicting a single day of discharge.

Here we have only shown a single example for a snow influenced basin. We also compared the modelling behavior in one of the arid catchments of the Arkansas-White-Red region (HUC 11), and found that the trends and conclusion where similar.

To conclude, although only based on an illustrative example, it shows very well the problem RNNs have with learning long-term dependencies and why they shouldn't be used if (e.g. daily) discharge is predicted only from meteorological observations.

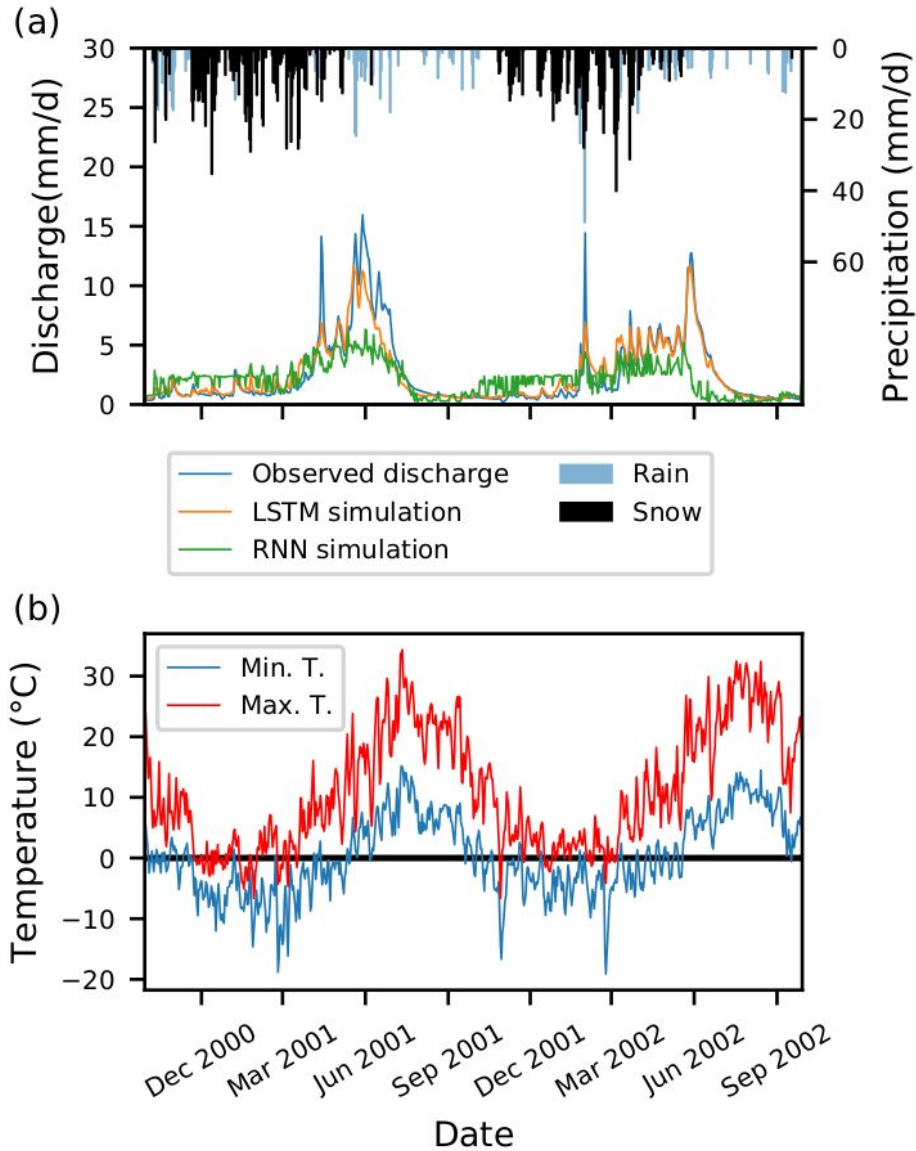


Figure caption:

- a) Two years of observed as well as the simulated discharge of the LSTM and RNN from the validation period of basin 13340600. The precipitation is plotted from top to bottom and days with minimum temperature below zero are marked as snow (black bars).
- b) The corresponding daily maximum and minimum temperature.

2. Page 6, Line 25: This is not specific and should be more detailed, "... were varied and found to work well in a number of preceding tests" – what values or ranges worked well, and how is "worked well" defined? I think this "initial screening" is also referred to in the conclusion and should be more clearly addressed as to how it was done.

The entire study arose from a free-time project some of us spent working on over the last 1.5 years. We did some experiments with data from (seasonally influenced) Austrian catchments. These were investigated in previous studies at our institute and calibrated hydrological models exist as reference (e.g. Herrnegger et al., 2018). We did some manual hyperparameter tuning, in which we mainly tried out one- and two-layer LSTMs with various size of hidden units. The architecture we used in the experiment of this manuscript is at the upper end of what we tested, in terms of number of learnable parameters. The capability of being able to model the rainfall-runoff process is given for all hyperparameter combinations we tried. The one used in this paper was one of the best we found at that time for the Austrian catchments. "Worked well" at this time was defined as "the LSTM achieved similar model performance as the hydrological models". We agree that some more of this information can be added into the manuscript and therefore adapt it in the following way.

Old passage (P6 L24ff):

The specific design of the network architecture, i.e. the number of layers, cell/hidden state lengths, dropout rate and input sequence length were varied and found to work well in a number of preceding tests.

New passage (Placed now under "Experimental design"):

The specific design of the network architecture, i.e. the number of layers, cell/hidden state length, dropout rate and input sequence length were found through a number of experiments in several Austrian, seasonal influenced, catchments. In these experiments, different architectures (e.g. one or two LSTM layer or 5, 10, 15, 20 cell/hidden units) were varied manually. The architecture used in this study, proved to work well for these catchments (in comparison to a calibrated hydrological model we had available from previous studies; Herrnegger et al., 2018) and was therefore chosen to be applied here without further tuning. A systematic sensitivity analysis of the effects of different hyper-parameters was however not done and is something to do in the future.

3. Section 2.1.1: The hydrological interpretation was not very useful until I got to the very end of the paper (Figure 14) where the evolution of a cell state is compared to temperature variables. Since Figure 14 and its associated discussion seem to be an afterthought in the conclusion section, I would recommend folding this example into section 2.1.1 instead, as they both relate to a "hydrological interpretation" of the data-driven network. Also in Figure 14, some vertical lines through the figure would be

useful to better link to the narrative about the thresholds between temperature and cell state.

We agree that vertical lines will enhance this figure and its interpretability and will adapt the figure in the revised manuscript.

Regarding the hydrological interpretation and Fig. 14 (also mentioned by AR #2 minor comment #6): We agree with AR2 that section 2.1.1 should be moved out of the method section, since it is more of a discussion or hypothesis. We also agree with AR1 that it is beneficial to link Fig. 14 directly into this section. In the revised paper we will update this paragraph accordingly.

4. Section 2.2: The definition of epoch is not quite clear to me – for example, is it the same as the “next iteration” loop in Figure 3, or something different? If the same, the idea of the epoch could be illustrated in Figure 3. It makes sense that a higher number of “epochs” in this sense would lead to improvement of the simulation as shown in Figure 4.

We see that this explanation/definition can be misleading, since an epoch is not the same as the “next iteration” in Figure 3. We will therefore adapt the passage in the revised manuscript so that it contains an example as illustration for the difference between epoch and iteration step in the context of neural networks:

Old passage (P7 L29ff):

The corresponding term for neural networks is called epoch. One epoch is defined as the period, in which each training sample is used once for updating/training the model parameters. This means, each time step of the discharge time series in the training data is simulated exactly once (which is similar to one iteration in classical hydrological model calibration).

New passage:

The corresponding term for neural networks is called epoch. One epoch is defined as the period, in which each training sample is used once for updating the model parameters. For example, if the data set consists of 1000 training samples and the batch size is 10, one epoch would consist of 100 iteration steps (number of training samples divided by the number of samples per batch). In each iteration step, 10 of the 1000 samples are taken without replacement, until all 1000 samples are used once. In our case this means, each time step of the discharge time series in the training data is simulated exactly once. This is somewhat similar to one iteration in the calibration of a classical hydrological model, however with the significant difference that every sample is generated independently from each other.

5. Section 2.3: In Figure 5 and discussion throughout the experiments and results sections, it would be useful to refer to the HUC basins (01,03,11,17) by the names of the watersheds or the regions (e.g. Pacific Northwest, Northeast, etc). This may make the results more interpretable for many readers, especially those familiar with climatology in the U.S.

We agree to the comment of AR1 and will adapt the namings throughout the manuscript accordingly. We also adapted Fig. 5, 6, 9 and 12 (all plots showing the contiguous united states in the background) to not show the US state borders, but rather the borders of the hydrological units (as suggested by the author of the short comment).

6. Section 2.4.2: In Line 19, the statement “in our case, the network has to learn the entire hydrological model purely from available data” – should specify that this is true of any data-driven approach, not specific to this case. Also in this section, comment on why fewer epochs were needed for Experiment 2 compared to 1?

AR1 is right, this is true for any data-driven approach and we will adapt the passage in the revised manuscript accordingly:

Old passage (P10 L19f):

In our case, the network has to learn the entire “hydrological model” purely from the available data (see Fig. 4).

New passage:

As for all data-driven approaches, the network has to learn the entire “hydrological model” purely from the available data (see Fig. 4).

Regarding the number of epochs: In our view this question/statement might be related to comment 4. Please read our answer here, together with the one provided therein. It is true that the number of epochs is lower compared to the models in experiment 1. The reason is that the total number of parameter updates is much higher since the training set of experiment 2 has a much higher number of samples (remember that all basins within a HUC contribute to the calibration). And, because the batch size is the same for the models in experiment 1 and experiment 2, the number of iteration steps per epoch has increased (see also answer to comment 4). For example: The HUC 01 has 27 basins, which means the number of available training samples per epoch are 27 times higher than for the models in experiment 1. Because the batch size is the same, this implies that one epoch of the HUC 01 model of experiment 2 has 27 times more iteration steps (= parameter updates) per epoch, compared to a single basin model of experiment 1.

From this observation it follows that the models in experiment 2 have seen a specific data point of a specific basin within the region less often during training, compared to a single basin model in experiment 1. This is because each data point is used once and only once during one epoch and maybe highlights also the cross-basin learning of the models in experiment 2.

To reduce the confusion, we will adapt the section in the revised manuscript as follows:

Old passage (P11 L 4ff):

Across all catchments, the highest mean NSE was achieved after 20 epochs in this case. Thus, for the final training, we train one LSTM for each of the four used HUCs for 20 epochs with the full 15-year long calibration period of all catchments within the specific HUC.

New passage:

Across all catchments, the highest mean NSE was achieved after 20 epochs in this case. Although the number of epochs is smaller compared to experiment 1, the number of weight updates is much larger. This is because the number of available training samples has increased and the same batch size as in experiment 1 is used (see Sect. 2.2 for an explanation of the connection of number of iterations, number of training samples and number of epochs). Thus, for the final training, we train one LSTM for each of the four used HUCs for 20 epochs with the entire 15-year long calibration period.

7. Section 2.6: This section breaks the flow of the paper between the description of the experiments and their results. I suggest placing this information earlier in the paper before the experiment descriptions or as an appendix.

We see the point of AR1 comments and agree that Section 2.6 (Open source software) may be a break in the flow of the story. However, we see the software we used as an essential tool/method for our work and we would therefore prefer to keep this section in the methods section. We propose to place this section before the description of the experiments in the revised version of the manuscript.

8. Page 12, Line 24: From Figure 6b, this claim is not very apparent to me, that LSTM outperforms the benchmark for more dry catchments (in HUC 11, it seems like it outperforms in the western part but not the eastern part, but the NSE is higher in the eastern part).

We agree that this statement seems unclear and not very apparent in the first version of the manuscript. We missed to state, that the arid basins are located in the western part of HUC 11 (see image below) which matches the location of basins, for which the LSTM

performs better (see Fig 6b of the original submission). We therefore adapted the passage in the revised manuscript as follows.

Old passage (P12 L24):

The performance deteriorates in the more arid catchments in the center of the CONUS, where no discharge is observed for longer periods of the year.

New passage:

The performance deteriorates in the more arid catchments, which are located in the western part of the Arkansas-White-Red region, where no discharge is observed for longer periods of the year (see Fig. 5b).

Furthermore, we added a second map to Fig. 5 (in the original manuscript it only showed the mean annual precipitation of each catchment). This additional map (see (b) in the figure below) shows the aridity index of all basins, and will hopefully be an aid for readers to understand the given statement with more ease.




Figure caption:

Overview of the location of the four hydrological units from the CAMELS data set used in this study including all their basins. (a) Shows the mean annual precipitation of each basin, whereas the type of marker symbolizes the snow influence of the basin. (b) shows the aridity index of each basin, calculate as PET/P (see Addor et al. 2017a)

- Page 12, Line 27: Why is this result surprising, since the LSTM is posed as a method to retain longer-term dependencies? This is a place where it would be advantageous to show how a traditional RNN would not capture these dependencies to prove its capabilities in this area.

To us it was surprising because it practically demonstrated the theoretical capacity for learning long-term relationships of the LSTM. At least for us it seemed not clear that it would work as good as it did for complex processes such as snow accumulation/melt. Regarding the comparison of LSTMs and RNNs see our answer to comment #1 of this review.

10. Figure 11 and associated discussion in Section 3.2: This may be expected since gages in the Northeast are more closely spaced and homogeneous compared to the Central Plains region, where there is a large wet-to-dry gradient between Missouri and Colorado. Some discussion on the characteristics of the regions of interest would be beneficial here (linking back to annual precipitation, other climate characteristics). Also, I don't think the Basin numbers in Figure 11 are ever defined so there is no way to interpret Figure 11 spatially (e.g. there is no way to look at a certain correlation for a pair of basins and understand why they are very different from each other). Possibly a better way to create this figure would be to order basins by longitude?

We agree that a discussion on the characteristics of each regions would be beneficial for the reader. Therefore we will add a table with some key attributes (see image below), as well as a textual description in the revised manuscript.

Old passage:

In our study, we used 241 catchments from the HUCs 01, 03, 11, 17 (see Fig. 5) in order to cover a wide range of different hydrological conditions on one hand and to limit the computational costs on the other hand. The selected catchments contain snow-driven catchments as well as catchments without any influence of snow. In addition, the four units cover a wide range of climates, containing rather dry catchments with less than 400 mm/year of mean precipitation, as well as catchments with mean precipitation up to 3260 mm/year.

New passage:

In our study, we used 4 out of the 18 hydrological units with their 241 catchments (see Fig. 5 and Table 1) in order to cover a wide range of different hydrological conditions on one hand and to limit the computational costs on the other hand. The New England region in the North-East contains 27 more or less homogeneous basins (e.g. in terms of snow-influence, aridity). The Arkansas-White-Red region in the center of CONUS has a comparable number of basins, namely 32, but is completely different elsewhere. Within this region, attributes e.g. aridity and mean annual precipitation have a high variance and strong gradient from East to West (see Fig. 5). Also comparable in size but with very different hydro-climatic conditions are the South Atlantic-Gulf region (92 basins) and the Pacific Northwest region (91 basins). The latter spans from the Pacific coast till the Rocky Mountains and also exhibits a high variance of attributes across the basins,

comparable to the Arkansas-White-Red region. For example, there are very humid catchments with more than 3000 mm/yr precipitation close to the Pacific coast and very arid (aridity index 2.17, mean annual precipitation 500 mm/yr) basins in the South-East of this region. The relatively flat South Atlantic-Gulf region contains more homogeneous basins (similar to the New England region), but is in contrast not influenced by snow.

(Screenshot of the new table)

Table 1. Overview of the HUCs considered in this study and some region statistics averaged over all basins in that region. For each variable mean and standard deviation is reported.

| HUC | Region Name | # Basins | Mean precipitation [mm/d] | Mean aridity ¹ [-] | Mean altitude [m] | Mean snow frac. ² [-] | Mean seasonality ³ [-] |
|-----|---------------------|----------|------------------------------|----------------------------------|----------------------|-------------------------------------|--------------------------------------|
| 01 | New England | 27 | 3.61 ± 0.26 | 0.60 ± 0.03 | 316 ± 182 | 0.24 ± 0.06 | 0.10 ± 0.08 |
| 03 | South Atlantic-Gulf | 92 | 3.79 ± 0.49 | 0.87 ± 0.14 | 189 ± 179 | 0.02 ± 0.02 | 0.12 ± 0.26 |
| 11 | Arkansas-White-Red | 31 | 2.86 ± 0.89 | 1.18 ± 0.50 | 613 ± 713 | 0.08 ± 0.13 | 0.25 ± 0.29 |
| 17 | Pacific Northwest | 91 | 5.22 ± 2.03 | 0.59 ± 0.40 | 1077 ± 589 | 0.33 ± 0.23 | -0.72 ± 0.17 |

¹: PET/P, see Addor et al. (2017a)

²: Fraction of precipitation falling on days with temperatures below 0°C

³: Positive values indicate that precipitation peaks in summer, negative values that precipitation peaks in the winter month and values close to 0 that the precipitation is uniform throughout the year (see Addor et al. (2017a))

The intention of Figure 11 was not to link specific basins within the confusion matrix to basins in the map. Our goal was to show the overall picture of the correlation between basins within one HUC, which is reflected by the overall color appearance. However, we agree that ordering the basins by longitude does enhance this figures, because the overall image is still the same, while at the same time the correlation plot is spatially interpretable. Thus, we change this plot as suggested by AR1 in the revised manuscript.

11. Section 4: In the conclusion, it would help to come back to the broad topic the introduction of hydrological modeling in general, and a discussion of process based models and other types of data-driven models in the context of the results, instead of re-iterating the results. As mentioned previously, Page 20 Lines 18 onward seem to be tacked-on to the end, and would be better placed earlier in the paper and referred back to here.

Regarding P20 Line 18ff: As stated in our answer to comment #3, we will move this section together with the hydrological interpretation and Fig. 14 to a new section under results and discussions.

We rewrote and restructured the entire conclusion and reduced the amount of summarization of our experiments. Furthermore, we added some additional discussion about limitations and advantages of our approach and possible future studies.

The new version of our conclusion is added to our answer to comment 3 of AR2.

12. Finally, a general comment regarding the results: It was found that the regional model performed better for regions with correlated discharge (e.g. the Northeast). However, the basis for the regional model was that more scenarios are present in the dataset (i.e. stated that long dry periods or extreme events may be observed in one catchment in the training, which may help to simulate similar types of events in another catchment). This makes it seem like the regional model should actually benefit for places where discharge is not correlated between stations (i.e. in the Central Plains rather than the Northeast) and spans a wider range of behaviors, whereas the opposite results are found in the study. I think this is linked to the catchment processes, in that in the Central Plains, rainfall-runoff processes occur differently between basins, so that a set of inputs and outputs for one basin cannot translate to model outputs in another. Meanwhile in the Northeast, climate is very similar between catchments, so while the regional model may not include so many disparate events (input samples are relatively similar), it still serves to improve the overall model of a given catchment. This may be somewhat addressed in the results and discussion, but could be expanded upon and help to discuss the model in a “hydrological process” context.

We agree with AR1 on a conceptual level. Although, for the case at hand the low MSE-values of the dry regions are the dominating factor for the distortion. What is happening here is that the LSTM learns to predict all basins well in average and thus arid basins (with very few error signals due to low MSE in dry periods) do not force the neural network to specifically adapt to these cases. Instead, the LSTM will adapt the parameters to fit the basins with large errors signals. However, if there would be a sufficiently larger number of arid basins (compared to semi-arid/humid basins), the LSTM would most likely learn to adapt to arid basins as well as to non-arid basins. For further studies or applications one could try to introduce weights to the objective function to compensate for the low MSE-values of arid basins.

Minor line by line comments and typos:

1. Page 7, Line 6: “as well as”: Thank you, this will be changed in the revised manuscript.
2. Page 7, Line 19: “iteration”: Thank you, this will be changed in the revised manuscript.

3. Page 10, Line 12: typo in “each the model”: Thank you, this will be changed in the revised manuscript.
4. Page 10, Line 17: would expand acronym to “deep learning”: Thank you, this will be changed in the revised manuscript.
5. Page 10, Line 20: “would help to obtain”? Thank you, this will be changed in the revised manuscript.
6. Page 10, Line 21: remove “e.g.”: Thank you, this will be changed in the revised manuscript.
7. Page 10, Line 30: remove comma after “analyze”: Thank you, this will be changed in the revised manuscript.
8. Page 12, Line 21: This makes sense that many zero-values would lead to worse predictions, since there are effectively “fewer” data points (in that many samples correspond to zero-flow values) in those training data sets. Could comment here on whether more epochs (greater than 50) would have benefited the model or not for this region?

This is indeed an interesting observation. Without further tests, we think that more epochs would not be beneficial but harmful. A central dichotomy of data-driven methods is the balance between generalisation and overfitting. The reason why more epochs might be harmful is that they would increase the probability of overfitting (an already serious problem for the models in experiment 1). Having “fewer” data points available, as AR1 correctly mentions, leads to an increase of the effect of overfitting. Thus training for more epochs would result in a model that is even more overfitted on the training data and generalizes even worse on the validation data. The opposite might be true that in arid catchments fewer epochs could be beneficial.

9. The acronyms FHV, FMS, FLV should be re-defined in this figure caption.: Thank you, this will be changed in the revised manuscript.
10. Figure 9 (and Figure 12): tiny text in the insets, should be able to read axis values: Thank you, this will be changed in the revised manuscript.
11. Page 13, Line 6: “more strongly”: Thank you, this will be changed in the revised manuscript
Page 13, Line 8: can barely see this from Figure 7a:
There is a mistake in the very next sentence, which might have made it more confusing. We reported the lowest NSE not for the calibration period but for the validation period. This will be changed in the revised manuscript.

We agree that even then it is difficult to see this statement in the empirical CDF. This is why we added the next sentence with the number of the lowest values to the original manuscript in the first place.

Old passage (P13 L7ff):

Regarding the performance in terms of the NSE, the LSTM shows fewer negative outliers and thus seems to be more robust. The poorest model performance in the calibration period is an NSE of -0.42 compared to -20.68 of the SAC-SMA + Snow-17.

New passage:

Regarding the performance in terms of the NSE, the LSTM shows fewer negative outliers and thus seems to be more robust. The poorest model performance in the validation period is an NSE of -0.42 compared to -20.68 of the SAC-SMA + Snow-17.

References:

1. Bengio, Y., Simard, P., and Frasconi, P.: Learning Long-Term Dependencies with Gradient Descent is Difficult, 1994.
2. Carriere, P., Mohaghegh, S., and Gaskar, R.: Performance of a Virtual Runoff Hydrographic System, *Water Resources Planning and Management*, 122, 120–125, 1996.
3. Herrnegger, M; Senoner, T; Nachtnebel, HP. Adjustment of spatio-temporal precipitation patterns in a high Alpine environment. *J HYDROL*. 2018; 556: 913-921.
4. Hochreiter, S. and Schmidhuber, J.: Long Short-Term Memory, *Neural Computation*, 9, 1735–1780, 1997.
5. Hsu, K.-I., Gupta, H. V., and Sorooshian, S.: Application of a recurrent neural network to rainfall-runoff modeling, in: *Proceedings of the 1997 24th Annual Water Resources Planning and Management Conference.*, ASCE, 1997.
6. Kumar, D. N., Raju, K. S., and Sathish, T.: River Flow Forecasting using Recurrent Neural Networks, *Water Resources Management*, 18,143–161, 2004.