A variational approach at uncertainty estimation in data-driven rainfall-runoff modeling

Manuel Álvarez Chaves[1,*], Eduardo Acuña Espinoza[2], Daniel Klotz[3], Hoshin Gupta[4], Uwe Ehret[2], Anneli Guthke[1]

[1]Stuttgart Center for Simulation Science, Cluster of Excellence EXC 2075, University of Stuttgart, Stuttgart, Germany
[2]Institute of Water and Environment, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany
[3]Interdisciplinary Transformation University Austria (IT:U), Linz, Austria
[4]Hydrology and Atmospheric Sciences, The University of Arizona, Tucson, AZ, USA

[*]Corresponding author: manuel.alvarez-chaves@simtech.uni-stuttgart.de

This is a non-peer-reviewed preprint submitted to EarthArXiv and to *Machine Learning: Earth* (IOP Science) for peer review.

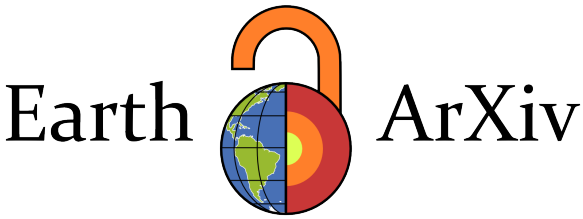

Version 0.1.4 — January 20, 2026

Journal Name

Crossmark

RECEIVED
dd Month yyyy

REVISED
dd Month yyyy

PAPER

A variational approach at uncertainty estimation in data-driven rainfall-runoff modeling

Manuel Álvarez Chaves[1,*], Eduardo Acuña Espinoza[2], Daniel Klotz[3], Hoshin V. Gupta[4], Uwe Ehret[2] and Anneli Guthke[1]

[1]Stuttgart Center for Simulation Science, Cluster of Excellence EXC 2075, University of Stuttgart, Stuttgart, Germany [2]Institute of Water and Environment, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany [3]Interdisciplinary Transformation University Austria (IT:U), Linz, Austria [4]Hydrology and Atmospheric Sciences, The University of Arizona, Tucson, AZ, USA *Author to whom any correspondence should be addressed.

E-mail: manuel.alvarez-chaves@simtech.uni-stuttgart.de

Keywords: rainfall-runoff modeling, lstm, mixture-density network, variational inference

Abstract

Reliable uncertainty estimation is essential for decision making, evaluating model performance, and defining the limits of what can be inferred from data. While uncertainty estimation typically requires specifying prior assumptions about distributional form, we introduce an approach to learn the structure of uncertainty directly from data. Specifically, we introduce a variational long short-term memory network (vLSTM) that uses variational inference to enable flexible, non-parametric probabilistic predictions. The vLSTM is assessed against deep learning baseline models for probabilistic rainfall–runoff prediction. We discuss training dynamics of probabilistic models, including concerns of overfitting, and compare predictive strategies that emphasize coverage versus point-wise accuracy. Results demonstrate that the vLSTM achieves state-of-the-art performance when evaluated using log-likelihood, while offering a distinct approach to uncertainty estimation that lets uncertainty patterns emerge instead of prescribing them. In our case study, the learned predictive distributions closely resemble that of the current baseline approach, which prescribes a mixture of asymmetric Laplacian distributions. This finding validates our approach, but also points to its fundamental strength: our variational approach to learning uncertainty structure has the potential to provide a more fundamental understanding of predictive uncertainty in arbitrary types of dynamic models and applications across scientific disciplines, enabling progress especially in fields where a priori assumptions seem hard to justify. In general, the vLSTM serves as a valuable approach for exploring uncertainty structures before transitioning to more computationally efficient models once the emerging patterns of uncertainty are better understood.

1 Introduction

Model evaluation and uncertainty estimation are fundamental to scientific advancement, as they allow researchers to quantify reliability, identify limitations, and establish the boundaries of what can be confidently concluded from their work (Chalmers, 2013). When researchers conduct hypothesis testing, they are essentially evaluating models that represent alternative theories about system behavior. Within a research program (Lakatos, 1970), each model contributes valuable understanding regardless of its apparent "success" or "failure" in fitting data. When a model fails to adequately represent observations, this failure itself becomes informative, revealing limitations in our current theoretical framework as represented by the model and pointing toward necessary refinements. Rather than discarding models that don't perfectly fit data, we learn from both their successes and shortcomings to iteratively improve our theoretical understanding. This approach acknowledges the Duhem-Quine thesis that perfect falsification is impossible, as individual hypotheses are embedded within broader theoretical frameworks and cannot be tested in isolation (Curd and Cover, 1998). The distinction between model performance and hypothesis validity is crucial; while we quantify outcomes using performance metrics and uncertainty bounds, we must interpret these results in terms of how informative they are to our evolving understanding (Nearing et al., 2020). The ultimate goal is not acceptance or rejection of a hypotheses, but rather the

continuous refinement of our models and theories through systematic learning from empirical evidence.

This philosophical framework becomes particularly important when models are used to inform real-world decisions, where the consequences of inadequate uncertainty quantification extend beyond scientific or academic considerations (Krzysztofowicz, 2001). As an example, in operational flood forecasting, accurate uncertainty bounds are essential for determining safe evacuation zones and issuing timely flood warnings (Nevo et al., 2022). Emergency managers, water utilities, and insurance companies depend not only on point predictions but also on uncertainty bounds to make informed risk-based decisions. Without proper uncertainty quantification, decision-makers are forced to rely on arbitrary thresholds or subjective judgment which potentially compromises the safety of the public or critical infrastructure (Todini, 2017).

In the context of hydrology, uncertainty quantification has mainly relied on propagating uncertainty in a Bayesian framework since very early work (Kuczera, 1983). Although the focus was initially on parameter uncertainty of hydrological models, the field has evolved by developing methods that also account for input and output uncertainty (Kavetski et al., 2006), enabled by advances that speed up sampling in a Markov Chain Monte Carlo (MCMC) framework (Vrugt, 2016). Bayesian model selection and averaging has been used to account for (parts of) conceptual uncertainty (Schöniger et al., 2014).

The use of model ensembles has also been a common approach to uncertainty quantification in hydrology, particularly in the subfield of flood forecasting (Demeritt et al., 2007). Model ensembles can be viewed through a Bayesian lens, where uncertainty is addressed by marginalizing not only over parameters but also over model structures. This approach is particularly useful when no single model structure is well specified, as it allows for the incorporation of structural uncertainty alongside parametric uncertainty.

Nevertheless, there is no consensus about the best approach to uncertainty quantification in hydrology (Beven, 2025). This is not surprising, as there are many sources of uncertainty that interact and affect hydrological predictions in often intractable ways. Further, it remains both a theoretical challenge to a priori specify all sources of uncertainty in a statistically rigorous manner and a practical challenge to propagate all known sources of uncertainty comprehensively with a limited computational budget. However, it is clear that uncertainty quantification is an essential part of hydrological modeling and should be considered in any model evaluation framework, as rigorously and transparently as possible.

Data-driven approaches have gained prominence in rainfall-runoff modeling, particularly those based on Long Short-Term Memory (LSTM) networks (Kratzert et al., 2018). Recently, there has been a growing interest in models that predict probability distributions instead of only point predictions based on these methods. Specifically, Klotz et al. (2022) proposed the use of Mixture Density Networks (MDNs) for this purpose, a standard method from deep learning (Bishop and Bishop, 2024). Alternatives such as Bayesian and Bayesian-like approaches (Blundell et al., 2015; Gal, 2016) have also been successfully applied in hydrology (Vasheghani Farahani et al., 2025), with the key distinction that these methods can capture both epistemic (from limited knowledge or data) and aleatoric (from inherent system variability) uncertainty, whereas MDNs primarily capture aleatoric uncertainty. This reflects an important difference between approaches: the MDN framework takes an "end-of-pipeline" perspective by learning the distribution that optimally covers the training data; it does not follow a source-model-output pathway of propagating uncertainty through a conceptual or physics-based model. In that sense, uncertainty quantification also becomes less mechanistic and more data-driven instead.

However, a challenge in implementing the MDN framework lies in the specification of probability distributions of the components. This choice requires researchers to make a priori assumptions about the functional constituents of uncertainty (whether they follow a Gaussian, log-Normal, Exponential, or other theoretical distribution functions). While a Gaussian Mixture Model (GMM) with sufficient components can (theoretically) approximate any distribution, the practical challenge remains: how can we choose an appropriate uncertainty structure without first completely understanding the system we are modeling? In other words, how can we anticipate which patterns of uncertainty will persist in the final model?

Ideally, uncertainty should emerge from the task itself rather than being imposed externally. Rather than prescribing a particular distributional form, the modeling framework should be flexible enough to learn and represent the uncertainty patterns that naturally arise from the data that is being modeled (Nearing et al., 2016). In the context of Bayesian inference, there are a number of likelihood-free methods that allow for the estimation of uncertainty without imposing a specific distributional form (Beven and Binley, 1992; Vrugt and Sadegh, 2013; Fenicia et al., 2018), yet

they still require the specification of prior distributions for parameters. Recent improvements to these methods include the use of a surrogate model when the simulator is too expensive (Reiser et al., 2025) and neural networks for parameter estimation (Hull et al., 2024). A different approach is to take a more data-centric perspective and use the data itself to derive the best-matching predictive distribution (Koutsoyiannis and Montanari, 2022; Auer et al., 2024).

In this study, we propose a novel approach to uncertainty estimation in rainfall-runoff modeling using a variational LSTM network (vLSTM). The vLSTM extends the traditional LSTM architecture by incorporating a variational layer (Blei et al., 2017) that is used to sample from and generate probability distributions as predictions. This approach allows the model to learn the uncertainty structure directly from the data without imposing a specific distributional form. The vLSTM is trained using a variational inference framework that enables it to generate probabilistic predictions while providing insights into the combined effects of both model error and observation error.

To evaluate the vLSTM's capabilities and validate our approach, we benchmark our model using the same experimental setup as the MDN models from Klotz et al. (2022). This cross-study comparison (Donoho, 2017) serves two objectives: First, it directly compares the performance of our sampling-based uncertainty estimation approach against the baseline mixture distribution methods under the same conditions. Second, to investigate the degree to which the distributions emerging from our data-driven vLSTM approach resemble the a priori defined distributions (Gaussian mixtures and asymmetric Laplacian mixtures) that have proven successful in previous work.

We put an emphasis on evaluating a model's performance by comparing the likelihood of the observed data under the model's predicted distributions, rather than using other scoring rules like the Winkler score (Winkler, 1972) or the Continuous Ranked Probability Score (CRPS) (Gneiting and Katzfuss, 2014). This evaluation approach maintains consistency with our training objective (maximum likelihood), ensuring that the metric used for assessment in testing directly reflects what the model was optimized to achieve. This approach also aligns with our philosophical framework of letting the model inform us about uncertainty structure as the likelihood directly measures how probable the observed data are under the model's learned distribution, without imposing external criteria for what constitutes "good" probabilistic predictions. Moreover, recent work has shown that many proper scoring rules, including CRPS, exhibit asymmetric penalization properties that can introduce unintended biases in model evaluation (Buchweitz et al., 2025). Note that the log-likelihood is equivalent to evaluating a model with the logarithmic score (Gneiting and Raftery, 2007). To be clear, we are not advocating for the use of likelihood as the sole metric for model evaluation. In fact, models should be evaluated using a diverse set of diagnostic tools (Gupta et al., 2008; Gauch et al., 2023). However, we argue that likelihood is a fundamental metric that should always be reported for probabilistic predictions; yet this is seldom done in practice. Even if the likelihood has no closed form, as is the case for the vLSTM, we demonstrate how it can still be approximated and reported (c.f. Appendix 5.2).

The paper is structured as follows: Section 2 describes the data and methods used in this study, including the CAMELS-US dataset, the vLSTM architecture, and the likelihood functions used for training. Section 3 presents the results of the vLSTM model, including its performance compared to an established benchmark as well as specific findings about the learned distributions of the model. Finally, Section 5 concludes with a discussion of the implications of this work for future research on uncertainty estimation in rainfall-runoff modeling and beyond.

2 Methods and Data

2.1 Variational LSTM (vLSTM)

The vLSTM extends the LSTM architecture by incorporating a variational layer that transforms the point predictions of a traditional LSTM into probability distributions in order to perform variational inference. In Section 2.1.1, we first highlight key concepts of traditional LSTMs. Afterwards, Section 2.1.2 describes the uncertainty propagation decoder of the vLSTM, which is also the first contribution of this study. Finally, we describe how uncertainty is propagated through the vLSTM in Section 2.1.3.

2.1.1 LSTMs LSTM networks (Hochreiter and Schmidhuber, 1997) are a type of recurrent neural network that are particularly well-suited for tasks related to sequences or time-series. The LSTM architecture consists of a series of cells that are connected sequentially. Each cell is composed by a cell state and a hidden state and contains three gates: the input gate, the forget gate, and the output gate. The input and forget gates control the flow of information in and out of the cell state, and the output gate controls the flow of information from the cell state to a hidden state. The

hidden state is then used to make predictions by passing it through a linear layer. These last two operations, the final update of the hidden state and prediction, are shown in Eq. 1 and Eq. 2:

$$h_t = o_t \odot \tanh(c_t), \tag{1}$$

$$\hat{y}_t = W_d h_t + b_d, \tag{2}$$

where h_t is the vector of hidden states at time t, o_t is the output gate at time t, c_t is the cell state at time t, $\odot$ is the element-wise multiplication operation, tanh is the hyperbolic tangent activation function, $\hat{y}_t$ is the prediction at time t, W_d is the weight matrix of the linear layer used for prediction, and b_d is the bias of the same linear layer.

In a more general sense, one can view the typical approach for creating predictions with an LSTM as an encoder-decoder approach. The encoder is the LSTM itself, which takes the input sequence and generates a hidden state that is used to make predictions. The decoder is then the layer that takes the hidden state and generates the prediction. Therefore, Eq. 2 can also be expressed as:

$$\hat{y}_t = f(h_t), \tag{3}$$

where f represents the operation of passing the hidden state through a function to generate a prediction and, in the special case above, f is a linear layer.

In this section we have focused on the internal operations of the LSTM that are of interest for the vLSTM. For a more in-depth description of LSTMs and, in particular, how they are applied in rainfall-runoff modeling, we refer to the work of Kratzert et al. (2018).

2.1.2 The uncertainty propagation decoder The main idea behind the vLSTM is to extend the traditional deterministic LSTM model to be able to make probabilistic predictions. This is achieved by incorporating a variational layer that is used to generate deviations from the point predictions of the LSTM. Using the hidden state shown in Eq. 1 as input, the variational layer generates one additional vector σ_t that is a vector of standard deviations which matches the number of dimensions of the hidden state. This vector is used to generate a random sample of noise for each hidden state. This approach of adding noise to the hidden state has been suggested before as an effective regularization technique (Dieng et al., 2018). Here, the noise is then treated differently depending on the uncertainty propagation decoder specified for the vLSTM.

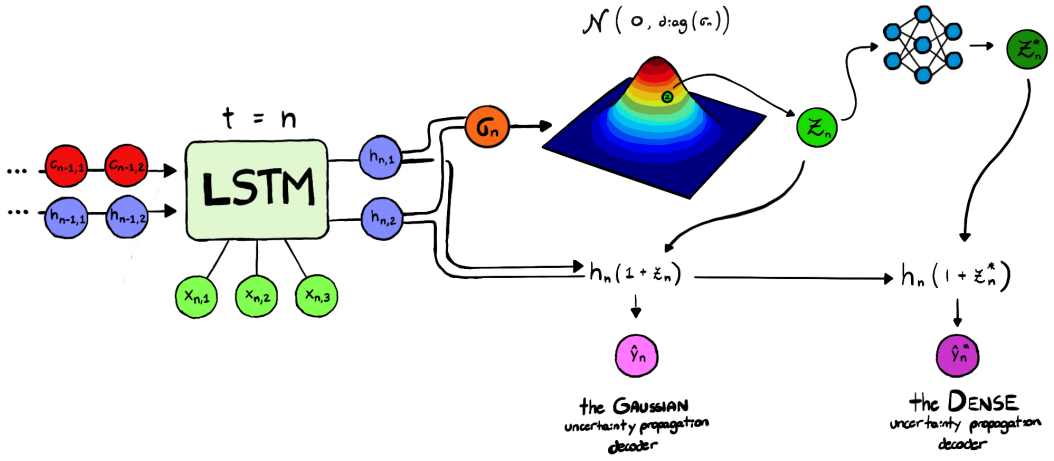

Figure 1. Illustration of the variational LSTM. The core idea is to use the hidden states of the LSTM layer to fit a multivariate Gaussian distribution with mean of 0 and a covariance matrix of σ_n in its diagonal. The noise z_n is then sampled from this distribution and used to generate the final prediction $\hat{y}_n$. The uncertainty decoder defines how the noise is applied to the hidden state to generate the final prediction.

For this study we introduce two uncertainty propagation decoders: the Gaussian decoder and the dense decoder. These are two extreme cases of how the noise can be treated to generate probability distributions; the proposed architecture in general is flexible to allow for any kind of transformation to be applied to the noise.

Gaussian uncertainty propagation decoder The Gaussian decoder scales the hidden state by a factor that is proportional to the noise generated by the variational layer before decoding to get $\hat{y}_t$:

$$\hat{y}_t = f(h_t(1 + z_t)) \tag{4}$$

In this case, as the noise z_t added to the hidden state comes from a multivariate Gaussian distribution and f is linear, the shape of the outcome $\hat{y}_t$ is also Gaussian.

Dense uncertainty propagation decoder The dense decoder applies a transformation g to the noise generated by the variational layer and then scales the hidden state by a factor that is proportional to this transformed noise:

$$\hat{y}_t = f(h_t(1 + g(z_t))) \tag{5}$$

The transformation g is a dense layer that takes the noise as input and outputs a vector of the same size.

As mentioned before, these are two extremes of the uncertainty propagation decoder. The Gaussian decoder propagates noise in a way that the predicted distribution matches the Gaussian prior of the latent noise. In fact, one wouldn't need a variational approach to accomplish the Gaussian predictive distribution but could just specify it a priori and learn the respective parameters of the theoretical distribution function; yet, here, we wish to demonstrate the spectrum of possible transformations within the vLSTM approach. The dense decoder allows for a completely flexible transformation to be applied to the noise. Importantly, the dense decoder allows the predictive distribution to be *learned* based on the characteristics and "predictability" of the training data. This property provides useful insights into the uncertainty that is representative of the underlying physical system to be modeled. Ultimately, arbitrary formulations relating $\hat{y}_t$ to h_t and $g(z_t)$ could be used and we leave it up to future work to explore different formulations further — especially for their effectiveness and identifiability during training.

2.1.3 Uncertainty propagation Uncertainty propagation is achieved by sampling. At the time of prediction, the model calculates the hidden state h_t and then takes a number of samples z_t from the variational layer as described in Section 2.1.2. The hidden state is then copied to match the number of samples from the variational layer and both are decoded through the uncertainty propagation decoder. The result is a number of (in our case: streamflow) predictions that matches the number of samples taken from the variational layer.

Figure 2a) exemplarily visualizes the vLSTM samples in prediction space, and their representation as a PDF via kernel density estimation (KDE, see Section 2.3). The predictive distributions of the models in Figures 2b) and 2c) are introduced in Section 2.2.

Sampling as a means for uncertainty propagation is what makes the vLSTM distinct from other approaches and enables its capability to learn the structure of uncertainty in the target. This capability comes at the cost of computational time. In comparison to a traditional LSTM trained to make point predictions, training a vLSTM with a Gaussian uncertainty propagation decoder is 1.9× more expensive in our case study (cf. Sections 2.5 and 3), while the dense uncertainty propagation decoder is 2.3× more expensive. A more detailed comparison is presented in Appendix 5.

2.2 Mixture density networks

Mixture Density Networks (MDNs) are a type of neural network whose outputs are treated as parameters for a mixture of probability distributions (Bishop and Bishop, 2024). More precisely, they model a conditional probability distribution $p(y|\mathbf{x})$, where y is the output and $\mathbf{x}$ is the input. Since the conditional probability distribution is modeled as a mixture of distributions, the MDN's output can be used to model complex distributions that are not easily captured by a single theoretical distribution function.

The MDN architecture uses a neural network to process inputs $\mathbf{x}$ and output parameters that define a mixture distribution. These outputs include mixture weights and distribution-specific parameters for each mixture component, which together define the PDF of the complete mixture.

We adapt the MDN framework of Klotz et al. (2022) which uses an LSTM encoder to process the input time series data. From the encoded hidden states, a head layer predicts the parameters and weights for different mixture models. We focus on two head configurations: a Gaussian Mixture Model (GMM) and the Countable Mixture of Asymmetric Laplacians (CMAL). Figure 3

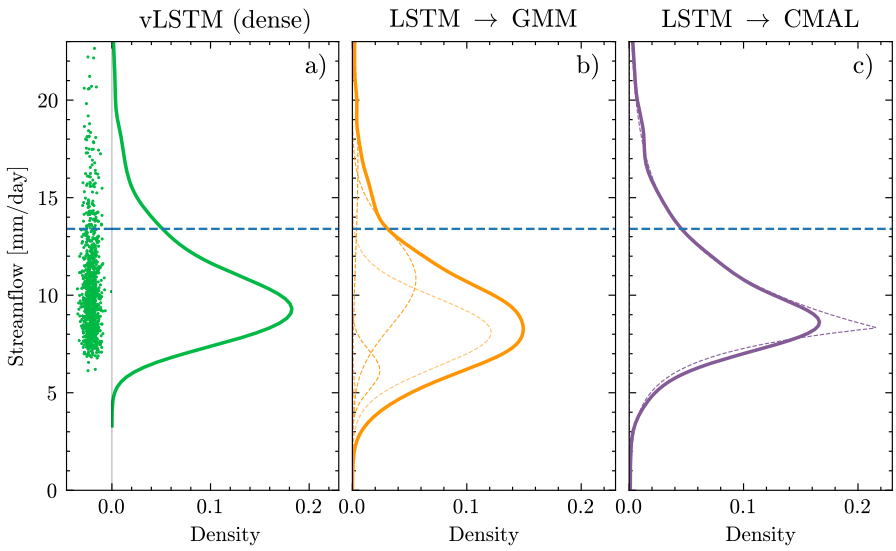

Figure 2. Illustration of predictive distributions obtained from a) dense vLSTM, b) an LSTM with a GMM head and c) LSTM with a CMAL head. Streamflow predictions are made for a specific basin and day of the case study, cf. Figure 4d). Points represent samples, thin dashed lines represent the components of the final mixture, thick lines represent a KDE-based approximation of the output distribution and the horizontal thick dashed line represents the observed value.

shows the MDN architecture with a GMM head as an example, and Figure 2b) and c) illustrate the resulting shapes of the mixture distributions.

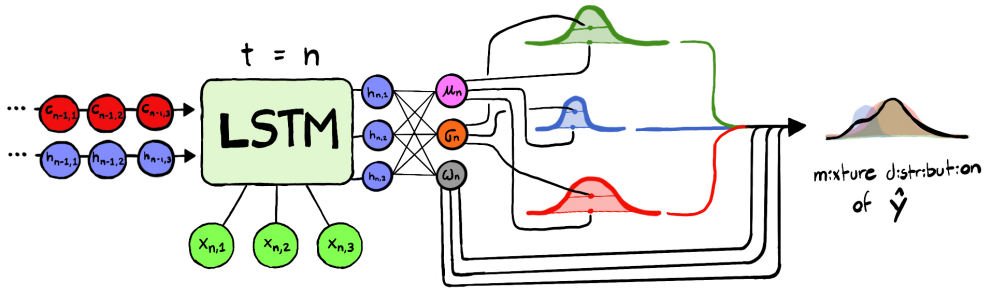

Figure 3. Illustration of a Gaussian mixture density network with three components. In this case, an LSTM is used as the specific type of neural network used to predict the weights and parameters of the components of the mixture. As the mixture consists of Gaussians, the predicted parameters are vectors of means μ_n and the standard deviations σ_n of each component. The resulting mixture distribution can be multi-modal and skewed.

2.3 Loss functions

In probabilistic modeling, the likelihood function is used to evaluate the adequacy of the predicted distribution as the probability of the observed data given the model structure and its inputs and parameters. This aspect should not be conceptualized as an independent methodological choice, but rather as the mathematical consequence of evaluating the model's predictive distribution against observed data. What is commonly seen as selecting a likelihood is actually the specification of a prior distribution embedded within the model's framework for predictions (Nearing et al., 2016).

This distinction reframes the relationship between model definition and computing a likelihood (Reuschen et al., 2021). When evaluating model performance during training, we simply read off the probability that the model assigns to the observed data according to its inherent predictive distribution. For the vLSTM with a Gaussian decoder, this evaluation is mathematically equivalent to applying a Gaussian likelihood function. In contrast, for the vLSTM with the dense decoder, which has a non-parametric predictive distribution, evaluating a predefined likelihood function is not possible. Instead, an empirical likelihood must be approximated via KDE. This perspective

emphasizes that likelihood functions are emergent properties of probabilistic model specification rather than independent modeling choices.

In this section we adapt the notation of Klotz et al. (2022). Runoff $y \in \mathbb{R}$ at a given time step is modeled as a probability distribution $p(\cdot)$ of the input $\mathbf{x} \in \mathbb{R}^{M \times T}$ where M is the number of input features and T is the number of time steps. The likelihood function is then defined as $p(y|\mathbf{x})$.

Gaussian likelihood The Gaussian likelihood assumes that the observed data is the outcome a Gaussian data-generating process:

$$p(y|\mathbf{x}) = \mathcal{N}(y|\mu(\mathbf{x}), \sigma(\mathbf{x})) \tag{6}$$

Where $\mathcal{N}$ is the probability density function of a Gaussian distribution shown in Eq. 7, $\mu(\mathbf{x})$ is the mean, and $\sigma(\mathbf{x})$ is the standard deviation of the Gaussian distribution. Note that both μ and σ are functions of the input $\mathbf{x}$ and, in the case of the vLSTM with the Gaussian decoder, they are both calculated from samples taken from the model.

$$\mathcal{N}(x|\mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \cdot e^{\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)} \tag{7}$$

For the GMM head, likelihood is calculated using Eq. 8 where α_k is the weight of the k-th Gaussian distribution and K is the number of Gaussian distributions in the mixture. This closely resembles Eq. 6 but instead of a single Gaussian distribution, the likelihood represents a mixture of K Gaussian distributions as shown in Eq. 8. The CMAL head is similar but instead uses a mixture of asymmetric Laplacian distributions. More details about the loss function for the CMAL head can be found in Klotz et al. (2022).

$$p(y|\mathbf{x}) = \sum_{k=1}^{K} \alpha_k \cdot \mathcal{N}(y|\mu_i(\mathbf{x}), \sigma_k(\mathbf{x})) \tag{8}$$

Empirical likelihood The KDE-based empirical likelihood is a non-parametric likelihood function that makes no assumptions about the distribution of the observed data. Instead, the likelihood of the observed data is approximated based on a kernel fitted to the predictions made by the model; with the choice of the kernel being the only a priori specification to make and its impact vanishing if chosen well (Jones et al., 1996). The KDE-based likelihood function is given by:

$$p(y|\mathbf{x}) = \frac{1}{n} \sum_{i=1}^{n} K(y_i - \hat{y}_i(\mathbf{x})) \tag{9}$$

Where n is the number of samples taken from the model, K is the kernel function, y is the observed runoff, and $\hat{y}_i(\mathbf{x})$ are the predictions of runoff made by the model.

There is a large range of kernel functions that can be used for the KDE-based likelihood. For simplicity, in this study we have chosen to use a Gaussian kernel function. With this assumption, Eq. 9 can be rewritten as:

$$p(y|\mathbf{x}) = \frac{1}{n\,h\,\sigma} \frac{1}{\sqrt{2\pi}} \sum_{i=1}^{n} \exp\left(-\frac{(y - \hat{y}_i(\mathbf{x}))^2}{2h^2\sigma^2}\right) \tag{10}$$

Where h is the bandwidth of the kernel function and σ is the standard deviation of the predictions $\hat{q}$. For the bandwidth, we have chosen to use the Silverman's rule of thumb (Silverman, 1998) which is shown in Eq. 11. Here $\hat{\sigma}$ is the standard deviation, IQR is the interquartile range and n is the number of sampled predictions $\hat{y}_i(\mathbf{x})$.

$$h = 0.9 \min\left(\hat{\sigma}, \frac{\text{IQR}}{1.34}\right) n^{-1/5} \tag{11}$$

Using a KDE-based likelihood function is not new. For instance, the Kernel Mixture Network (KMN) (Ambrogioni et al., 2017) uses a KDE-based likelihood function to calculate the likelihood of the observed data. This approach differs from ours as they fit a kernel over the observed data with no sampling involved.

Finally, in the context of machine learning, both Eq. 8 and Eq. 10 can be used as loss functions L during training by minimizing the negative log-likelihood (NLL) because this is equivalent to maximizing the likelihood of the observed data under the model:

$$L = -\log(p(y|\mathbf{x})) \tag{12}$$

NLL has a distinct connection to information theory as it relates to the cross-entropy $H(p, q)$, which measures the discrepancy between the true distribution of the data p and the predicted distribution by a model q. Furthermore, cross-entropy can be written in terms of the Kullback-Leibler divergence D_{KL} of p in relation to q. This is shown in Eq. 13:

$$H(p, q) = H(p) + D_{KL}(p||q) = \int_{\mathcal{X}} P(x) \log Q(x)\, dx, \tag{13}$$

where $H(p)$ is the entropy of the distribution of the observed data, and P and Q are probability density functions of p and q, respectively. This is a key insight because, from the perspective of information theory, any type of model evaluation which is not distribution-based will result in a loss of information (Vrugt, 2024). This principle applies broadly to any model of real-world systems where observations are imperfect and arise from complex chains of physical processes and measurement procedures, making the probabilistic characterization of uncertainty fundamental rather than optional.

For the vLSTM specifically, an issue arises from σ_t predicted by the variational layer. As σ_t is learned from the training data, it often causes problems for the model to generalize outside of the training set, ostensibly because too much of the specific characteristics of the training data are captured in σ_t rather than the underlying patterns of uncertainty. To address this, during training a second term L_{KLD} is included in the loss function that penalizes the deviation of σ_t from a constant value of 1.0. This term is the Kullback-Leibler divergence between a multivariate Gaussian distribution with mean of 0 and a covariance matrix of σ_t in its diagonal and a multivariate Gaussian distribution with mean of 0 and a covariance matrix of I in its diagonal:

$$L_{KLD} = \frac{1}{2} \sum_{i=1}^{n} (\sigma_{t,i}^2 - 1 - \log(\sigma_{t,i}^2)) \tag{14}$$

2.4 Model evaluation

The overall performance of each model is evaluated on the testing set as the log-likelihood of the observed data under the predicted distribution made by the model. Evaluating the adequacy of the predictive distribution using log-likelihood ultimately matches our training objective and focuses on what the model was encouraged to learn. For the MDN-based models, log-likelihood is evaluated according to their specific heads, either GMM or CMAL. For the vLSTM models, log-likelihood is evaluated as a Gaussian or empirical likelihood depending on the specific uncertainty propagation decoder.

While the adequacy of the predicted uncertainty intervals is our primary goal, we also measure the accuracy of traditional point-wise predictions made by probabilistic models with the Nash-Sutcliffe Efficiency (NSE; Nash and Sutcliffe, 1970):

$$NSE = 1 - \frac{\sum_i^n (y_i - \hat{y}_i)^2}{\sum_i^n (y_i - \overline{y})^2}. \tag{15}$$

Because NSE uses the mean of observed data as a baseline ($\overline{y}$), it can be misleading when calculated for batches, as we do in both our training and validation sets (Klotz et al., 2024). Nevertheless, NSE is only used to monitor the training process, and our actual training target, NLL, is not affected by this behavior.

Since for the vLSTM models we can only obtain samples of our target y_i (and since there is no closed-form probability density function for the distribution learned by the model), we calculate for each time step the mean of the samples drawn from the model and use this mean as $\hat{y}$ to calculate NSE. Similarly, for the MDN-based models we also use the mean prediction to calculate NSE. However, in their case, there is a closed-form solution for both the mean of the GMM head and the mean of the CMAL head.

2.5 Case study setup and implementation

The case study for this paper is based on 531 basins in the daily CAMELS-US dataset (Addor et al., 2017). As with other studies which utilize this dataset (Kratzert et al., 2018; Klotz et al., 2022), we focus on the prediction of streamflow at the outlet of each basin using dynamic meteorological and static basin attributes as input features. More information about the inputs and target can be found in Appendix 5. The training period was from 1 October 1999 to 30 September 2008, the validation period from 1 October 1980 to 30 September 1989, and the testing

period from 1 October 1989 to 30 September 1999. Further details on the data set can be found in the above mentioned previous works.

The training procedure follows closely the benchmark of Klotz et al. (2022). Specifically, the MDN-based models directly adopt the optimal hyperparameters from Klotz et al. (2022) (listed in Table 2), i.e., 10 components in the case of GMM and 3 components in the case of CMAL, with an LSTM hidden size of 250 in both cases. For the vLSTM model, we conducted hyperparameter search as described in Appendix 5 with the optimal hyperparameters shown in Table 3, i.e., the same hidden size of 250, a sample size of 1 000, and a two-layer DNN for the dense decoder. Additionally, each model was trained ten times to account for the effect of initialization of the model's weights and biases based on a random seed. Results are reported for all of the seeds except when otherwise mentioned.

Each model was trained for 30 epochs; then, we identified the epoch in which the lowest loss function (NLL) was achieved in the validation set (Goodfellow et al., 2016) and used the respective weights and biases for producing predictions in the testing set which correspond to the results shown in Section 3. This model selection approach differs from the original benchmark, which used different criteria for selecting their final models. Nevertheless, to validate our methodology, we provide a comparison with the benchmark results in Appendix 5.3.

3 Results and Discussion

In this section, we evaluate the different approaches to uncertainty estimation. First, we investigate and compare the resulting predictive distributions of the models in Section 3.1. Then, we summarize the overall performance of the models on the large sample data set in Section 3.2, before evaluating specific aspects such as coverage across quantiles vs. accuracy of point predictions in Section 3.3. Finally, we comment on the characteristic training behavior of MDN vs. vLSTM in Section 3.4.

3.1 Learned distributions

Before summarizing each model's performance in terms of metrics, the most novel and attractive aspect of the proposed dense vLSTM is revealing the structure of the distributions that it can learn. As such, while metrics can capture human preference (Gauch et al., 2023), visualizing the learned predictions and accompanying uncertainty ranges in a hydrograph remains essential. Therefore we begin with visual analysis before evaluating overall model performance.

3.1.1 Dense vLSTM Figure 4 shows the hydrograph representing a typical peak flow for the basin of the Buffalo Fork river measured at gauge 13011900, located under East Highway 287 in Wyoming, from late February to early November 1991. The June peak coincides with the floods in Central Wyoming that year (Jordan and Combs, 1996).

In the figure, we see the predicted hydrograph along with the uncertainty ranges provided by one specific seed of the dense vLSTM model. The model captures the observed flows well, but there are notable differences in the predicted distributions over time, particularly at five different points of the hydrograph that are highlighted within the figure.

Subplots b) and f) show the predicted distributions during periods of low flow, where the model is highly (and rightly) confident in its predictions, producing a very narrow distribution with very high density at the point of the observed data. Subplots c) to e) show the predicted distributions during three specific peak flows in the analyzed period. At these points, the shape of the learned distribution can be clearly seen. This learned distribution is not Gaussian but resembles a peaked distribution with a tail toward higher values (right skew). This shape reflects the fact that the model tends to underestimate rather than overestimate peak flows, with greater uncertainty about how high the streamflow will rise than about potential underestimation, even when inputs suggest high-flow conditions.

At the point of maximum flow (subplot d), the model has to widen the distribution in order to place a reasonable probability density on the observed data point. Probabilistic models still struggle to adequately capture high flow events, likely due to the consecutive nature of these three peaks (Brunner, 2023), but they provide valuable information about the uncertainty associated with them.

3.1.2 Comparison with Gaussian vLSTM and MDN The previous analysis made for the predictions of the dense vLSTM can be extended to all other models as shown in Figure 5. Column-wise results are for predictions on the same day made by different models. We observe the effect of the Gaussian vLSTM being constrained to a symmetric predictive distribution. Given the

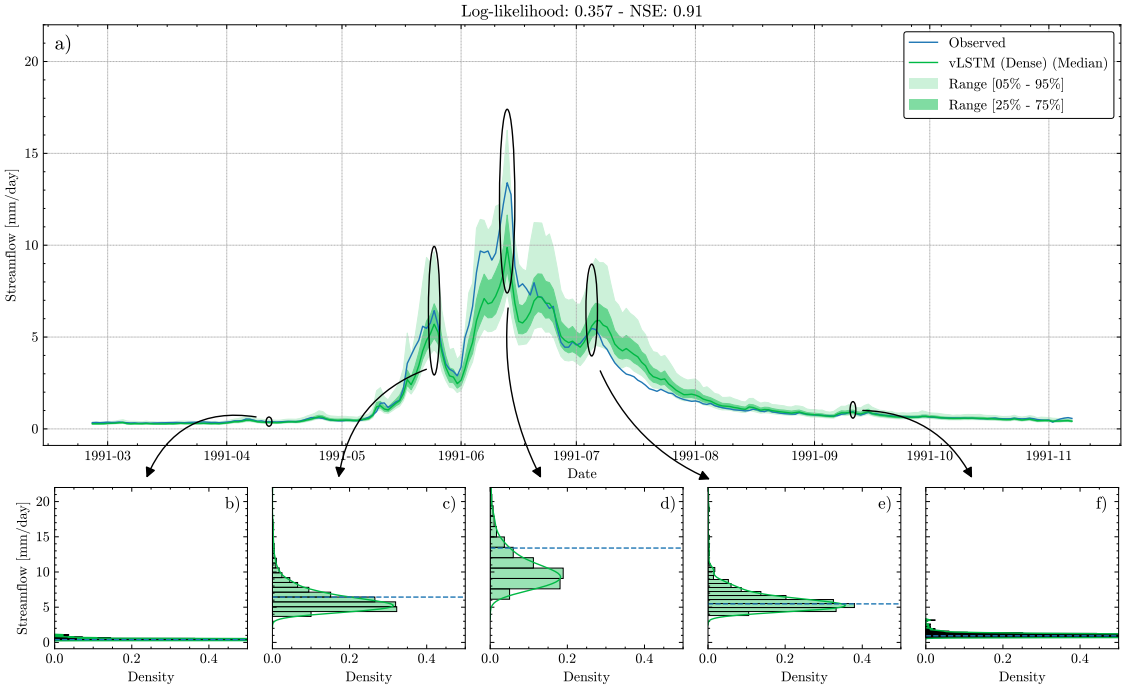

Figure 4. Subplot a) Observed hydrograph of basin 13011900 and the predictions made by the dense vLSTM model showing the median prediction and the 50% and 90% range of predictions. Subplots b-f) show the distribution of the predictions made by the dense vLSTM model at different days as indicated in Subplot a).

difficulty of getting peak flows right, the Gaussian vLSTM is forced to produce very wide distributions, which results in a lower density at observations, as can be seen in the top row of Figure 5. In contrast, the skewed learned distribution by the dense vLSTM closely resembles that of both MDN-based models with the closest match being the MDN with the CMAL head. This is not a coincidence as it is noted in Klotz et al. (2022) as their best model and it has succesfully been adopted in operational settings (Nevo et al., 2022) and to generate new hydrological data (Klotz et al., 2025).

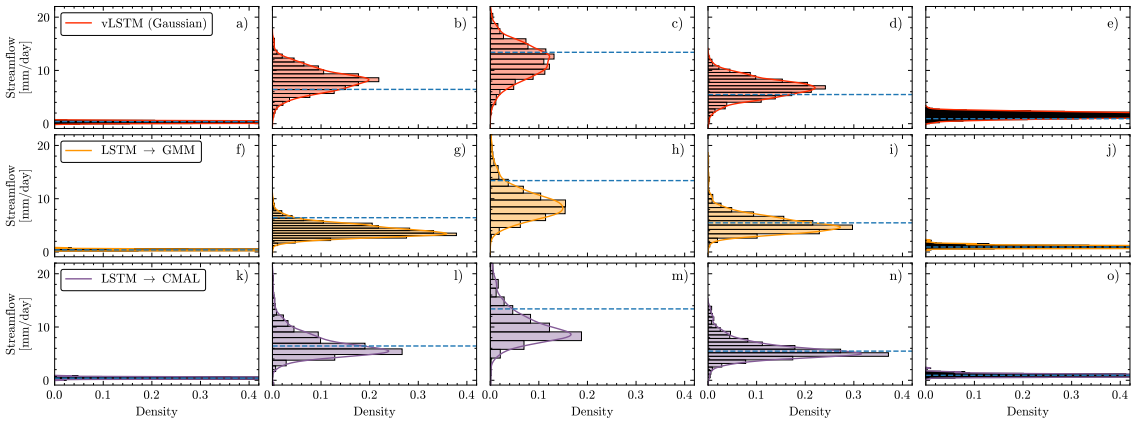

Figure 5. Row-wise comparison of the predicted distributions made by the models other than the dense vLSTM for the same days in Figure 4.

The shape of the learned distributions showing a strong match means that we find the same answer through two very different approaches. The vLSTM learns the distribution by sampling from a latent space, while the MDN-based models learn the distribution by fitting a mixture of a priori specified distributions to the data. This supports our hypothesis that the appropriate shape of the distribution is an emergent property of the combination of data and model (here: the LSTM) and can be learned. It also confirms that CMAL as an LSTM head was well chosen.

3.2 Overall performance

As is typical in large sample case studies, we report the overall model performance as a cumulative density function (CDF) of performance scores for each basin in the dataset during the testing period. Here, we present the CDFs of log-likelihood in Figure 6 and NSE in Figure 7. The median of the CDF is shown as a solid line, while the shaded area represents the 95% interval of the median, calculated using the 2.5% and 97.5% quantiles for each metric across the ten distinct random initial seeds.

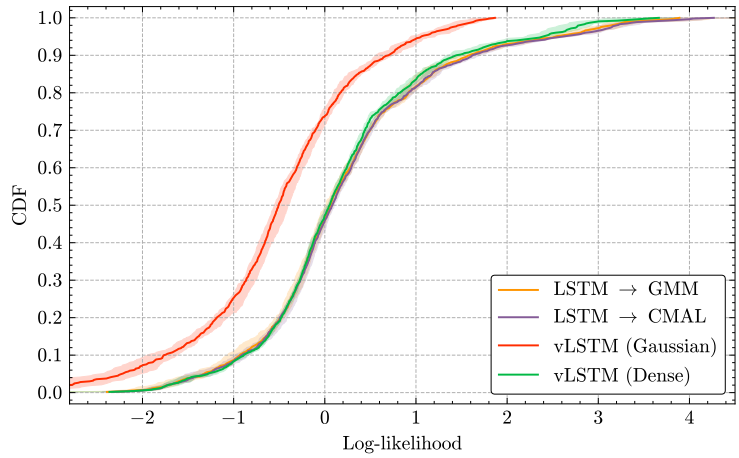

Figure 6. Cumulative log-likelihood distribution for the testing set. Log-likelihood is calculated for each basin in the testing set and the result for each basin is accumulated along the x-axis. The solid line represents the median for each model and a 95% interval is shown representing the variability in the result due to random initialization of the model.

In terms of log-likelihood, i.e., the probability of the observed data under the model, we see that the Gaussian vLSTM model (which propagates noise proportionally from its latent space as a Gaussian distribution) performs significantly worse than all other models in this study. This is due to the model being constrained to a symmetric predictive distribution and hence "diluting" its probability mass unrealistically (cf. Figure 5). The same is true for NSE (cf. Figure 7), as a symmetric distribution will have a mean forced away from the observations to "counteract" the natural skewness of the data to be modeled by the Gaussian vLSTM. This shows that predictive uncertainty does not correspond to a simple Gaussian distribution, and the model struggles with both coverage and point-wise accuracy.

In contrast, the LSTMs with both the GMM and CMAL heads and the dense vLSTM perform significantly better because they share the fundamental property of being more flexible in the shape of distribution they learn. In the case of MDN, this flexibility comes from being able to use more than one component to create a mixed distribution and allows these models to capture a wider range of potential outcomes and uncertainties, as evident from Figure 6.

These results similarly translate to the NSE scores, as shown in Figure 7. The arrangement of the models in Figs. 6 and 7 is the same, with one exception: the dispersion due to initialization is larger for the vLSTM in terms of point-wise predictions. This is likely due to how the vLSTM propagates noise (cf. Section 3.4) and that NSE was not the training or model selection target.

The additional model in Figure 7 is a pure LSTM trained specifically to make the best point-predictions possible. This represents a fundamentally different objective than all other models because it focuses on point-wise accuracy rather than density and coverage, which would immediately disqualify this model from the present study. Yet, it is included here to demonstrate how different loss functions ultimately guide what the final model can or cannot do.

The MDN models have been shown to match the point-wise accuracy of the LSTM in Klotz et al. (2022) when used as part of an ensemble: while in the original benchmark the authors calculated performance metrics using the median predictions across multiple random seeds (more natural for their operational focus), here the analysis was done for the individual performance of each seed separately. This approach was used to examine the full distribution of model performance for a single model, rather than collapsing results to a single value. For a visual and quantitative comparison, we refer to Appendix 5.3. Nevertheless, the fundamental limitation of the pure LSTM remains: it is incapable of making probabilistic predictions by itself.

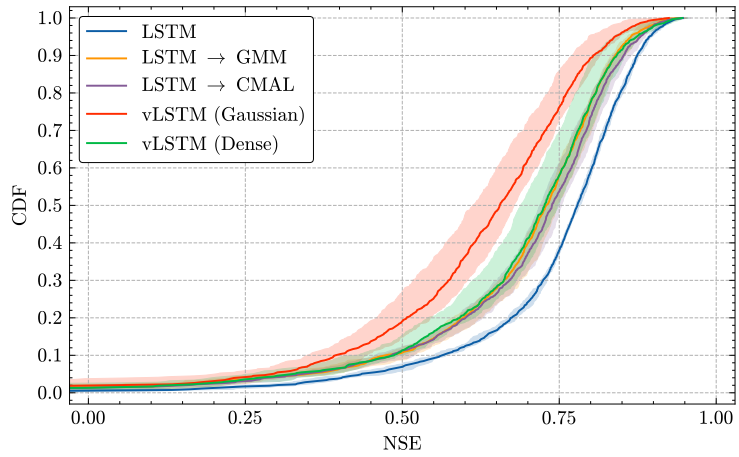

Figure 7. Cumulative NSE distribution for the testing set. NSE is calculated for each basin in the testing set and the result for each basin is accumulated along the x-axis. The solid line represents the median for each model and a 95% interval is shown representing the variability in the result due to random initialization of the model.

3.3 Model calibration

To assess how well the model's predictions cover the observed data, we use the Probability Integral Transform (PIT) method, as proposed by Laio and Tamea (2007). This method evaluates probabilistic predictions by transforming each observed value using the model's predicted CDF, producing values between 0 and 1. For a well-calibrated model, these PIT values should follow a uniform distribution, which can be checked by plotting their empirical CDF against a uniform CDF. A perfect match lies along the 1:1 line, indicating that predicted probabilities align with observed frequencies. Deviations from this line represent over- vs. underconfidence, and over- vs. underprediction. While the original method uses a continuous threshold window, we adopt a discrete version, evaluating at ten thresholds between 0 and 1 (Klotz et al., 2022).

Figure 8 shows the mean probability plots for all models on the testing set. The Gaussian vLSTM exhibits a typical case of over-prediction, where observations predominantly fall toward the lower tail of the model's predictive distribution. All other models, both MDN-based models and the dense vLSTM, appear well-calibrated with a slight tendency to produce what Laio and Tamea (2007) describe as "narrow" predictions (overconfidence). This reflects a training behavior issue, as models can overfit the training set and become overly confident in their predictions, which subsequently harms coverage on the testing set. This is discussed further in Section 3.4.

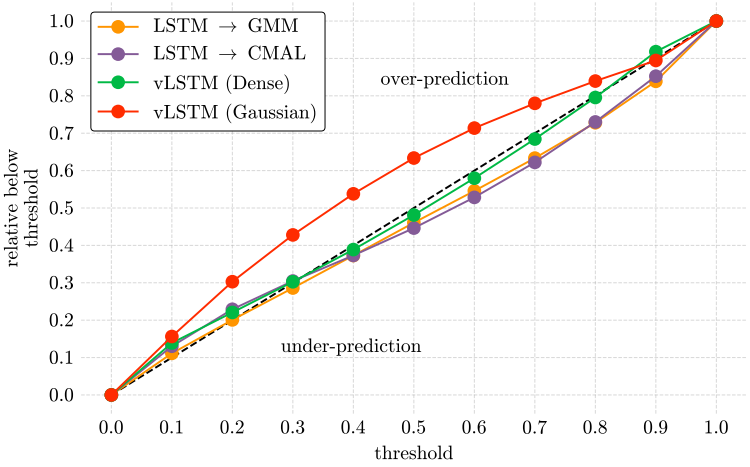

Figure 8. Probability plot to assess the calibration of each model in the testing set.

Albeit correct from a probabilistic point of view, it could be argued that the analysis from Figure 8 is "too marginal" from a hydrological standpoint, since it presents the average expected behavior across all basins. Figure 9 shows the actual spread of the reliability plots for every basin in the dataset, displayed as deviations from the 1:1 line to simplify interpretation. Here we observe that, while there is substantial variability in the reliability plots across basins, the MDN-based

models and the dense vLSTM exhibit very similar behaviors.

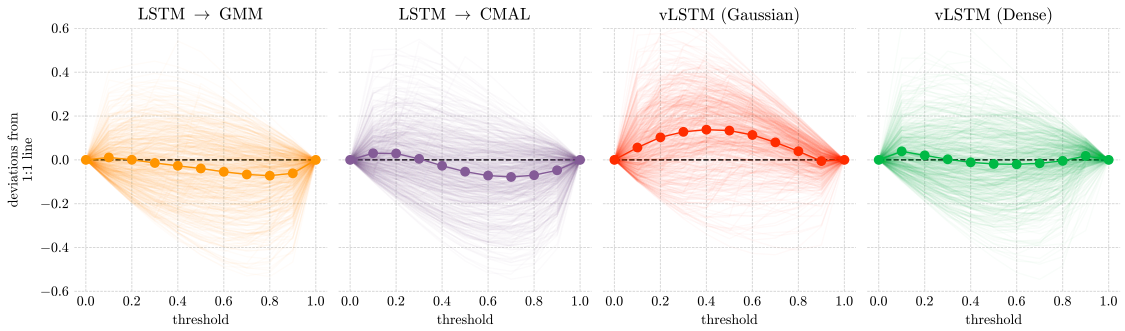

Figure 9. Probability plot shown as deviations from the 1:1 line. Background lines represent the reliability plot for each individual basin, while the foreground line shows the mean across all basins, as in Figure 8.

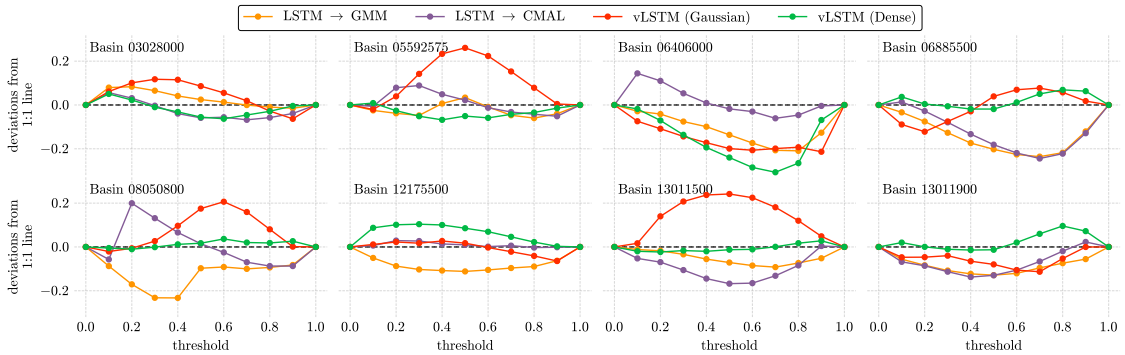

Figure 10. Probability plots (as deviations from the 1:1 line) for individual basins illustrating model behavior across different hydrological regimes.

Figure 10 shows calibration plots for selected individual basins. Basin 13011900, which was used previously in the example for Section 3.1.1, presents a case where most models under-predict, meaning that the observed data predominantly falls in the upper tail of the predictive distribution. This can be observed in some of the subplots of Figure 13, where the predictive distribution fails to cover some of the higher values of the observed hydrograph. Nevertheless, a single hydrograph is not representative of the behavior across the entire 'testing' set, and Figure 10 provides a more comprehensive view.

In general, Figure 10 shows large variability in model behavior across basins. For example, basin 12175500 shows that the MDN-based model with the CMAL head is perfectly calibrated, while in basin 13011500 it largely under-predicts. Also there are cases like basin 08050800 where all models but the dense vLSTM seem uncalibrated. Nevertheless, the average behavior across all basins aligns with what was described for Figure 8.

3.4 Training dynamics and overfitting

Figure 11 compares the training dynamics of the MDN with a GMM head and the vLSTM with the dense decoder. The MDN with a CMAL head and the vLSTM with the Gaussian decoder exhibit similar training dynamics to those shown and are therefore omitted for clarity. Looking at the training dynamics of the MDN with a GMM head in subplot a), we can immediately see one of the disadvantages of this particular architecture. The loss function steadily decreases in the training set, but the validation set shows a clear increase in loss after a few epochs, indicating overfitting as the loss steadily increases from this epoch on. This overfitting occurs due to the uncertainty in the model's predictions collapsing around the observed data, a behavior that is apparent precisely because we are minimizing NLL for specific data points (that is, the model can always put more probability mass on the training data by reducing the spread around them). Looking at NSE in subplot b), we see that overfitting does not negatively affect accuracy when using a point-wise distance between the expected value of the predicted distribution and the observed data. In fact, NSE continues to increase slightly during training for both the training and validation sets. This indicates that the model is learning to make better point predictions (i.e., the

uncertainty collapses around the predicted points), but it also shows that the model is not learning (or rather: forgetting how) to capture uncertainty adequately.

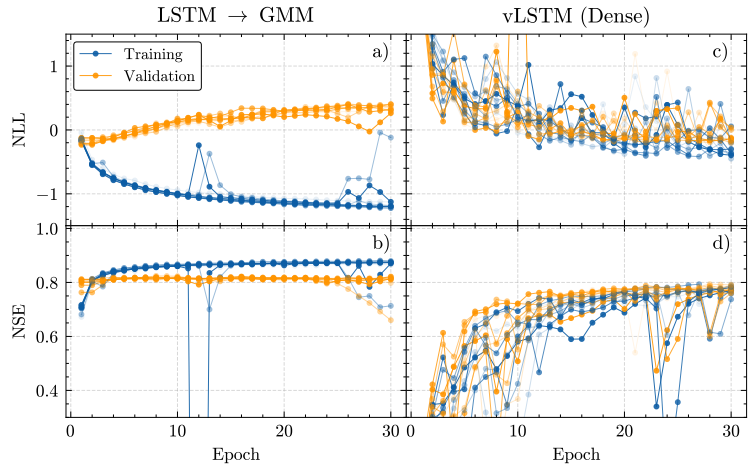

Figure 11. Training dynamics for the mixture density network with a Gaussian head and the vLSTM with the dense decoder. Subplots a) and c) show the evolution of the loss function (NLL) during training. Subplots b) and d) show the evolution of NSE as a monitored metric during training. Validation results are also shown as a function of the model weights and biases after finishing training for an epoch.

The training dynamics of the vLSTM shown in subplots c) and d) are, on the other hand, much more chaotic, with the loss function exhibiting significant fluctuations throughout training. This instability is due to the propagation of noise from the latent distribution used to make predictions and hence somewhat expected. Nevertheless, we observe a slower yet steady decrease in the loss function for both the training and validation sets. This slower training process is ultimately the main tradeoff made by choosing this particular architecture. Nevertheless, both the MDN and vLSTM models ultimately achieve similar performance in both NLL and NSE in validation, with the vLSTM avoiding the over-fitting during training and hence not reaching the minimum NLL and maximum NSE of MDN in training.

The trade-off between optimizing for coverage vs. point predictions is illustrated further with the training experiment in Figure 12. Here we see different versions of the MDN models. Two MDN models that were selected for coverage (log-likelihood) from an early epoch in Fig.11 and as in Figure 12, while the other two were selected for point predictions (NSE) from a later epoch in Fig. 11, and use weights and biases from a later training epoch when the model is overfitting in NLL but not NSE, as described above.

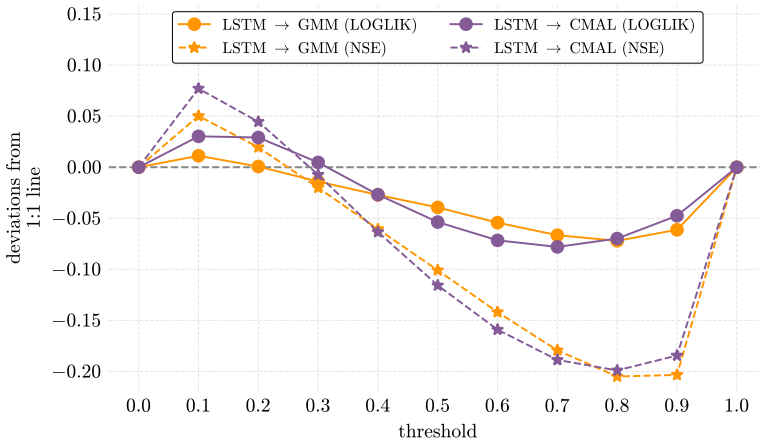

Figure 12. Comparison between probability plots of models that focus on coverage (LOGLIK) and those who focus on point accuracy (NSE) as deviations from the 1:1 line.

When the uncertainty bounds collapse around the observed data, the model overfits the 'training' set and becomes overconfident on the 'testing' set. The probability plot shows that both models exhibit greater overconfidence. This observation also contributes to explaining the results

shown in Figure 7, where the deterministic LSTM outperformed the probabilistic models in terms of NSE. Since the probabilistic models were trained and selected based on metrics that emphasize uncertainty quantification rather than point prediction accuracy, their parameters have been optimized to capture the characteristics of an output distribution at the expense of predictive precision of a mean estimate.

4 Summary

In summary, the vLSTM's ability to learn and represent uncertainty patterns directly from data with minimal assumptions about their structure has been demonstrated through a series of experiments and comparisons with existing methods, showing the effectiveness of this approach in capturing the complexities of uncertainty in modeling hydrological data. The main tradeoff users face is computational speed.

The results indicate that the vLSTM with a dense uncertainty propagation decoder is able to match the performance of state-of-the-art models for this dataset when evaluated using the same objective as the training loss function, which emphasizes consistent coverage of the quantiles of the observed data. Based on point-prediction accuracy, the vLSTM shows competitive performance, but we argue that this evaluation is misguided. The issue lies in the fact that point-prediction accuracy does not adequately capture aspects of predictions made using probability distributions, which are crucial for understanding uncertainty. If the predictive distribution was Gaussian, this would not be such an issue because maximizing log-likelihood for it is equivalent to minimizing the sum of squared errors as in NSE. However, as we saw in Section 3.1, the learned distribution is definitely not Gaussian. This issue can be further exemplified by analyzing the training dynamics and probability plots of the MDN-based models, which become better at making point predictions if they overfit the training data, yet become much more miscalibrated in a reliability assessment.

More importantly, the methods explored in this paper allow us to study the structure of uncertainty in a data-driven workflow. For hydrological modeling and specifically for streamflow data, a pattern of a peaked distribution with a tail toward higher values was identified, and this pattern was found consistently using two very distinct methods: the sampling-based, assumption-free approach of the dense vLSTM and the a priori specification approach of the MDNs. Notably, while this research focuses on combined model and observed data errors and uncertainty as a whole, we observe similar patterns to studies that have focused on observed data alone (De Oliveira and Vrugt, 2022). The shape of this distribution can also be related to the known issue of extrapolation in purely data-driven models, where performance deteriorates at higher flow magnitudes (Frame et al., 2022; Acuña Espinoza et al., 2025; Baste et al., 2025).

Overall, the proposed vLSTM can be seen as a valuable initial approach for exploring the structure of uncertainty. Given its computational demands, findings from this initial exploration can be used to transition to more efficient MDN-based models once these patterns are better understood.

5 Conclusions

What defines a model? More than just its model class or specific architecture, a model is also defined by the training data from which it shall learn and the objective function it must accomplish. Many mistakenly reduce models to their architecture alone, overlooking these other essential components. Given these foundations, we must carefully consider what we expect from our statistical golems (McElreath, 2020).

Immanuel Kant argued that our understanding of reality is fundamentally shaped by what he called the "categories of understanding" (Kant, 2007). Just as humans can only perceive and understand the world through these mental frameworks, models can only "know" what their architecture allows them to process, what their training data has exposed them to, and what their optimization criteria have encouraged them to learn. This parallel suggests that models, like humans, are constrained by the boundaries of their fundamental design.

In this paper, we have introduced the vLSTM, a model designed with minimal assumptions about the structure of uncertainty, allowing it to learn and represent uncertainty patterns directly from data. These findings suggest that, while multiple theories might explain the same data, a data-driven approach significantly constrains the space of plausible theoretical frameworks. By first extracting maximum information from the available data, we can eliminate theories that are inconsistent with empirical patterns and focus on those that align with the structure of our observed uncertainty.

This methodology aligns with Lakatos' philosophy of scientific research programs, where auxiliary hypotheses; in our case, assumptions about a specific choice of distribution should be

modified based on empirical evidence while preserving the core theoretical framework. Our success validates previous work using Mixture Density Networks. Specifically, our results corroborate the use of a mixture of asymmetric Laplacian distributions. What was a modelers prior choice in Klotz et al. (2022) now reveals itself as an emergent property that data-driven methods can discover systematically. By allowing models to inform us about uncertainty with minimal assumptions, we open pathways to more robust and generalizable uncertainty quantification in hydrology and beyond, especially in application cases where a priori assumptions about model-data uncertainty seem hard to justify.

Appendix A

Additional details of the experimental setup

As mentioned in Section 2, the models in this study were trained on the daily CAMELS-US dataset (Addor et al., 2017). This dataset contains meteorological time-series forcing data, originally from Newman et al. (2015), and static basin attributes for 671 catchments in the contiguous United States (CONUS) that have minimal human impact. A subset of 531 basins was chosen based on considerations regarding basin area and the methods used to calculate these areas (Kratzert et al., 2019).

Precipitation (*prcp*) forcing data comes from three distinct gridded products: Daymet (1 × 1 km), National Land Data Assimilation System (NLDAS), and Maurer (12 × 12 km, both). Four additional variables were derived from the Daymet forcing time series: shortwave radiation (*srad*), minimum and maximum temperature (*tmin* and *tmax*), and vapor pressure (*vp*). The static basin attributes are shown in Table 1. As target, observed streamflow at a gauge (*QObs*) was used.

Table 1. Basin attributes from the CAMELS-US dataset used to train all models. Adapted from Addor et al. (2017).

Type	Attribute	Description	Unit
Topographic	elev_mean	Basin mean elevation	m.a.s.l
Topographic	slope_mean	Basin mean slope	m km^{-1}
Topographic	area_gages2	Basin estimated area	km^2
Climate	p_mean	Mean daily precipitation	mm day^{-1}
Climate	pet_mean	Mean daily PET, estimated using the Priestley-Taylor formulation	mm day^{-1}
Climate	aridity	Ratio of estimated PET and precipitation PET/P	-
Climate	frac_snow	Fraction of precipitation falling as snow	-
Climate	high_prec_freq	Frequency of high precipitation days ($\geq$ 5 times mean daily precipitation)	days yr^{-1}
Climate	high_prec_dur	Number of consecutive high precipitation days	days
Climate	low_prec_freq	Frequency of dry days ($<$ 1 mm day^{-1})	days yr^{-1}
Climate	low_prec_dur	Number of consecutive dry days	days
Land Cover	frac_forest	Forest fraction	-
Land Cover	lai_max	Maximum monthly mean of the leaf area index (LAI)	-
Land Cover	lai_diff	Difference between the maximum and minimum monthly mean of the LAI	-
Land Cover	gvf_max	Maximum monthly mean of the green vegetation fraction (GVF)	-
Land Cover	gvf_diff	Difference between the maximum and minimum monthly mean of the GVF	-
Soil	soil_depth_pelletier	Depth to bedrock (maximum 50 m)	m
Soil	soil_depth_statsgo	Soil depth (maximum 1.5 m; layers marked as water and bedrock were excluded)	m
Soil	soil_porosity	Volumetric porosity (estimate based on sand and clay fraction)	-
Soil	soil_conductivity	Saturated hydraulic conductivity (estimate based on sand and clay fraction)	cm h^{-1}
Soil	max_water_content	maximum water content (combination of porosity and soil_depth_statsgo)	m
Soil	sand_frac	sand fraction (of the soil material smaller than 2 mm)	%
Soil	silt_frac	silt fraction (of the soil material smaller than 2 mm)	%
Soil	clay_frac	clay fraction (of the soil material smaller than 2 mm)	%
Geology	carbonate_rocks_frac	fraction of the catchment area characterized as “carbonate sedimentary rocks”	-
Geology	geol_permeability	subsurface permeability (log10)	m^2

Hyperparameters

Mixture density networks For the current study, we adopted the hyperparameters used in Klotz et al. (2022). The hyperparameters of the MDN models in this study are shown in Table 2.

Table 2. Hyperparameters of the MDN models used in this study.

Hyperparameter	MDN → GMM	MDN → CMAL
LSTM hidden size	250	250
Number of components	10	3
Dropout	0.4	0.5
Learning rate	$1e^{-3}$	$5e^{-4}$

vLSTMs Table 3 summarizes the hyperparameter search space for the dense and Gaussian vLSTM model, with the final choices indicated in bold. This space contains 72 unique combinations, and we conducted an exhaustive search where each combination was evaluated using the same initialization seed (42). The results were tabulated, and the parameters of the model with the lowest validation NLL were selected for the final runs presented in Section 3.

Table 3. Search space for hyperparameters of the dense vLSTM model. Final choice of hyperparameters in bold.

Hyperparameter	Choices
LSTM hidden size	**250**, 500
Num. of layers for the DNN	1, **2**
Num. of samples	100, **1000**, 5000
Dropout	0.4, **0.5**
Learning rate	$1e^{-3}$, $5e^{-4}$, $\mathbf{1e^{-4}}$

Sampling from the vLSTM
As mentioned in Section 2.1.2, the variational layer in the vLSTM is sampled to generate deviations from the point predictions of the traditional LSTM. The variational layer is a multivariate Gaussian distribution centered at 0 with a covariance fitted by the predicted σ_t from the LSTM. Sampling from this learned distribution using σ_t presents a challenge for optimization using gradient descent because gradients cannot be tracked during a sampling procedure. To overcome this, we use a variation of the reparametrization trick (Kingma and Welling, 2022) and generate an additional vector ϵ_t that is sampled from a standard multivariate Gaussian distribution $\mathcal{N}(0, I)$ and multiply it by σ_t to generate the noise z_t.

Computational time
As discussed in the main text, learning the structure of uncertainty requires additional time. In Table 4, we report the time it takes for each model to finish a single training epoch on an NVIDIA RTX A6000. Since we trained the models in a many-to-one setting, one training epoch refers to one pass through the model to make predictions for each day in the training period across all 531 basins in the CAMELS-US data. This is equivalent to making 1 745 928 single-day predictions.

Table 4. Training time of one epoch for each of the models in this study is reported in mm:ss format (minutes and seconds). Times are reported as the mean and standard deviation for independent training runs where no other jobs were allocated to the GPU.

Model	Time
LSTM	05:49 ± 00:07
LSTM → GMM	07:03 ± 00:05
LSTM → CMAL	07:34 ± 00:06
Gaussian vLSTM	11:00 ± 00:06
Dense vLSTM	13:13 ± 00:12

Comparing the average times and using the traditional LSTM trained to make point predictions as a baseline, we see that the MDN models are 1.2× and 1.3× more expensive for the GMM and CMAL heads, respectively, while the vLSTM models are 1.2× more expensive for the Gaussian uncertainty model and the dense uncertainty model, respectively.

Appendix B

5.1 Additional hydrographs

In Section 3.1.2 we compared the learned distributions of models other than the dense vLSTM for specific days. Figure 13 shows the hydrographs of the models other than the dense vLSTM for the period between February and November 1991, as discussed in Section 4.

5.2 Truth and approximation

In Section 3.2, the MDN models were evaluated using their theoretical log-likelihoods and compared to the vLSTM models. This comparison is somewhat unfair, as the log-likelihoods of vLSTM models have no closed-form solution and must be approximated using a KDE. If the MDN models lacked a closed form for their log-likelihoods, one would similarly need to sample from their predictive distributions and use the same approximation. This comparison is presented in Figure 14.

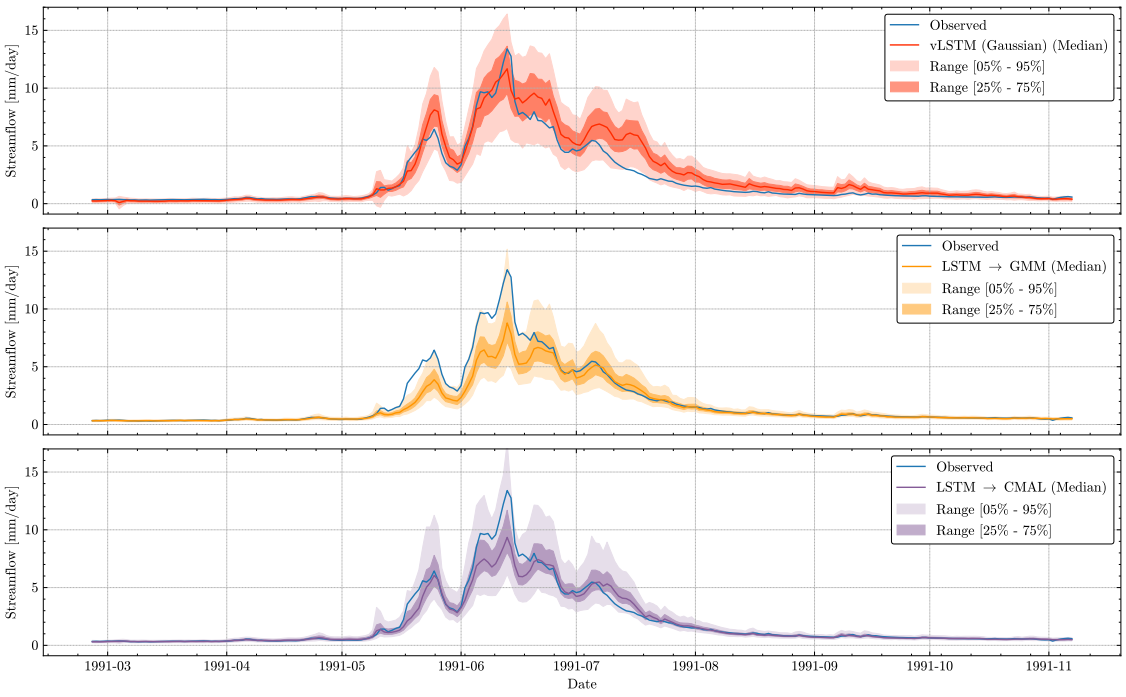

Figure 13. Hydrographs for Gaussian vLSTM and both MDN models for the period matching Figure 4.

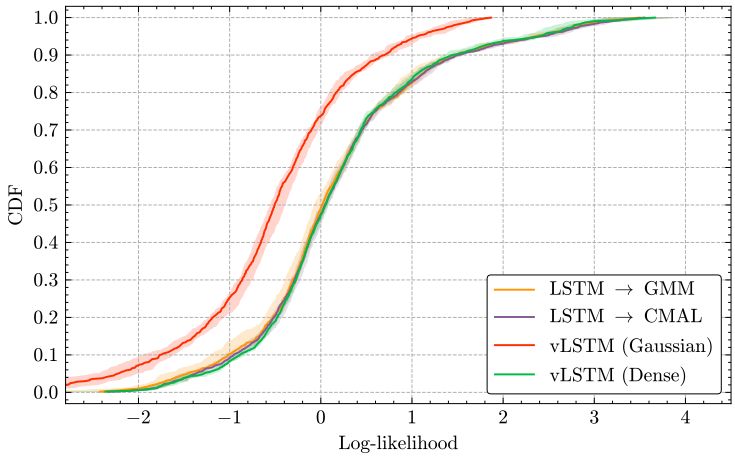

Figure 14. Approximated log-likehood CDF of the models in this article.

From Figure 14, we can see that the results for overall performance are qualitatively the same. The overlap between the dense vLSTM and the MDN-based models is larger in this case, but the general conclusion: that all models perform similarly, remains unchanged. Even if a model's log-likelihood has no closed form, one can sample from its predictive distribution and compute an approximation to a reasonable degree of accuracy.

5.3 Matching a benchmark

In the main text, we referenced how the MDN models presented in this study differ from those of Klotz et al. (2022) in both training methodology and approach to evaluation. Figure 15 compares the models in this paper with the original benchmark using an NSE CDF, where the benchmark results are shown as dashed lines and labeled "ensemble". Two differences distinguish the approach used in this article: First, different stopping criteria were used during training (as discussed in the main text). Second, models were evaluated differently: while in the original benchmark the authors calculated performance metrics using the median predictions across multiple random seeds (more natural for their operational focus), here the analysis was done for the individual performance of each seed separately. This approach was used to examine the full distribution of model performance for a single model, rather than collapsing results to a single value. In Table 5, apart

from showing only the median NSE, the area under the curve (AUC) of the median is also shown. AUC serves as an overall metric to assess the CDF and has been effectively used for model comparison (Gauch et al., 2021).

We hope that this comparison is useful to clarify the differences between the models presented in this study and those of our main reference.

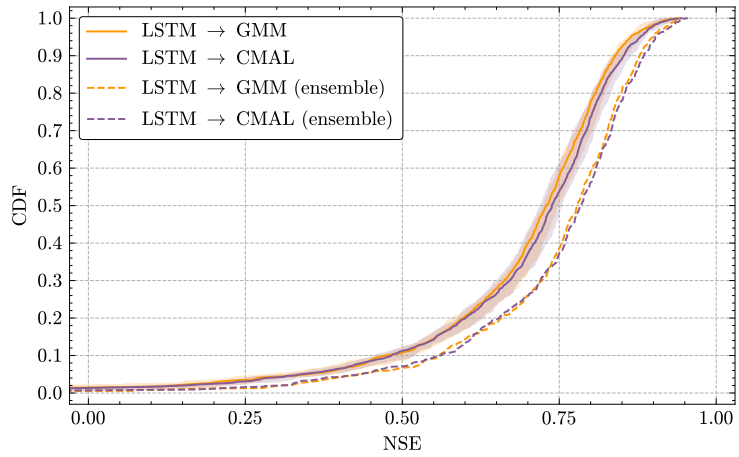

Figure 15. Comparison of the MDN models shown in the main manuscript with the original benchmark.

Table 5. Comparison of the median and AUC of NSE for the models shown in the main manuscript with the original benchmark. The values are read from Figure 15.

Model	Median	AUC
LSTM → GMM	0.729	0.327
LSTM → CMAL	0.739	0.320
LSTM → GMM (ensemble)	0.779	0.266
LSTM → CMAL (ensemble)	0.785	0.265

Funding

We acknowledge funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy—EXC 2075–390740016 and project 507884992. We also acknowledge support by the state of Baden-Württemberg through the use of bwHPC.

Author contributions

The idea of the vLSTM was developed by MAC, HVG and AG. DK provided advice on the MDN models and training practices for probabilistic models. The codes were written by EAE and MAC. The simulations were conducted by MAC. Results were discussed by all authors. The draft was prepared by MAC and AG. Reviewing and editing was provided by all authors. Funding was acquired by AG and UE. All authors have read and agreed to the current version of the paper.

Code and data availability

The code for the models and specific loss functions used in this study can be found as part of the Hy2DL library (Acuna et al., 2025).

The CAMELS US dataset can be found at https://doi.org/10.5065/D6G73C3Q. All the code for this project, model state dictionaries, model configurations, training logs and netCDF files of the results of this study have been archived at the data repository of the University of Stuttgart (DaRUS) and can be found at this link: https://doi.org/10.18419/DARUS-5118.

References

E. Acuna, M. Álvarez Chaves, A. Dolich, and A. M. J. Hy2dl: Hybrid hydrological modeling using deep learning methods, 2025. URL https://doi.org/10.5281/zenodo.17251944.

E. Acuña Espinoza, R. Loritz, F. Kratzert, D. Klotz, M. Gauch, M. Álvarez Chaves, and U. Ehret. Analyzing the generalization capabilities of a hybrid hydrological model for extrapolation to extreme events. *Hydrology and Earth System Sciences*, 29(5):1277–1294, 2025. doi: 10.5194/hess-29-1277-2025.

N. Addor, A. J. Newman, N. Mizukami, and M. P. Clark. The CAMELS data set: Catchment attributes and meteorology for large-sample studies. *Hydrology and Earth System Sciences*, 21 (10):5293–5313, 2017. doi: 10.5194/hess-21-5293-2017.

L. Ambrogioni, U. Güçlü, M. A. J. van Gerven, and E. Maris. The Kernel Mixture Network: A Nonparametric Method for Conditional Density Estimation of Continuous Random Variables, 2017.

A. Auer, M. Gauch, F. Kratzert, G. Nearing, S. Hochreiter, and D. Klotz. A data-centric perspective on the information needed for hydrological uncertainty predictions. *Hydrology and Earth System Sciences Discussions*, pages 1–37, 2024. doi: 10.5194/hess-2024-64.

S. Baste, D. Klotz, E. A. Espinoza, A. Bardossy, and R. Loritz. Unveiling the Limits of Deep Learning Models in Hydrological Extrapolation Tasks. *EGUsphere*, pages 1–24, 2025. doi: 10.5194/egusphere-2025-425.

K. Beven and A. Binley. The future of distributed models: Model calibration and uncertainty prediction. *Hydrological Processes*, 6(3):279–298, 1992. doi: 10.1002/hyp.3360060305.

K. J. Beven. A short history of philosophies of hydrological model evaluation and hypothesis testing. *WIREs Water*, 12(1):e1761, 2025. doi: 10.1002/wat2.1761.

C. M. Bishop and H. Bishop. *Deep Learning: Foundations and Concepts*. Springer International Publishing, Cham, 2024. ISBN 978-3-031-45467-7 978-3-031-45468-4. doi: 10.1007/978-3-031-45468-4.

D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational Inference: A Review for Statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017. doi: 10.1080/01621459.2017.1285773.

C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra. Weight Uncertainty in Neural Networks, 2015.

M. I. Brunner. Floods and droughts: A multivariate perspective. *Hydrology and Earth System Sciences*, 27(13):2479–2497, 2023. doi: 10.5194/hess-27-2479-2023.

E. Buchweitz, J. V. Romano, and R. J. Tibshirani. Asymmetric Penalties Underlie Proper Loss Functions in Probabilistic Forecasting, 2025.

A. F. Chalmers. *What Is This Thing Called Science?* Hackett Publishing Company, Inc, Indianapolis, 4. ed edition, 2013. ISBN 978-1-62466-038-2 978-1-62466-039-9.

M. Curd and J. A. Cover, editors. *Philosophy of Science: The Central Issues*. Norton, 1998.

D. Y. De Oliveira and J. A. Vrugt. The Treatment of Uncertainty in Hydrometric Observations: A Probabilistic Description of Streamflow Records. *Water Resources Research*, 58(11): e2022WR032263, 2022. doi: 10.1029/2022WR032263.

D. Demeritt, H. Cloke, F. Pappenberger, J. Thielen, J. Bartholmes, and M.-H. Ramos. Ensemble predictions and perceptions of risk, uncertainty, and error in flood forecasting. *Environmental Hazards*, 7(2):115–127, 2007. doi: 10.1016/j.envhaz.2007.05.001.

A. B. Dieng, R. Ranganath, J. Altosaar, and D. M. Blei. Noisin: Unbiased Regularization for Recurrent Neural Networks, 2018.

D. Donoho. 50 years of data science. *Journal of Computational and Graphical Statistics*, 26(4): 745–766, 2017. doi: 10.1080/10618600.2017.1384734. URL https://doi.org/10.1080/10618600.2017.1384734.

F. Fenicia, D. Kavetski, P. Reichert, and C. Albert. Signature-Domain Calibration of Hydrological Models Using Approximate Bayesian Computation: Empirical Analysis of Fundamental Properties. *Water Resources Research*, 54(6):3958–3987, 2018. doi: 10.1002/2017WR021616.

J. M. Frame, F. Kratzert, D. Klotz, M. Gauch, G. Shelev, O. Gilon, L. M. Qualls, H. V. Gupta, and G. S. Nearing. Deep learning rainfall–runoff predictions of extreme events. *Hydrology and Earth System Sciences*, 26(13):3377–3392, 2022. doi: 10.5194/hess-26-3377-2022.

Y. Gal. *Uncertainty in Deep Learning*. PhD thesis, University of Cambridge, Cambridge, 2016.

M. Gauch, J. Mai, and J. Lin. The proper care and feeding of CAMELS: How limited training data affects streamflow prediction. *Environmental Modelling & Software*, 135:104926, 2021. doi: 10.1016/j.envsoft.2020.104926.

M. Gauch, F. Kratzert, O. Gilon, H. Gupta, J. Mai, G. Nearing, B. Tolson, S. Hochreiter, and D. Klotz. In Defense of Metrics: Metrics Sufficiently Encode Typical Human Preferences Regarding Hydrological Model Performance. *Water Resources Research*, 59(6):e2022WR033918, 2023. doi: 10.1029/2022WR033918.

T. Gneiting and M. Katzfuss. Probabilistic Forecasting. *Annual Review of Statistics and Its Application*, 1(Volume 1, 2014):125–151, 2014. doi: 10.1146/annurev-statistics-062713-085831.

T. Gneiting and A. E. Raftery. Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007. doi: 10.1198/016214506000001437.

I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, Massachusetts, 2016. ISBN 978-0-262-03561-3.

H. V. Gupta, T. Wagener, and Y. Liu. Reconciling theory with observations: Elements of a diagnostic approach to model evaluation. *Hydrological Processes*, 22(18):3802–3813, 2008. doi: 10.1002/hyp.6989.

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

R. Hull, E. Leonarduzzi, L. De La Fuente, H. Viet Tran, A. Bennett, P. Melchior, R. M. Maxwell, and L. E. Condon. Simulation-based inference for parameter estimation of complex watershed simulators. *Hydrology and Earth System Sciences*, 28(20):4685–4713, 2024. doi: 10.5194/hess-28-4685-2024.

M. C. Jones, J. S. Marron, and S. J. Sheather. A Brief Survey of Bandwidth Selection for Density Estimation. *Journal of the American Statistical Association*, 91(433):401–407, 1996. doi: 10.1080/01621459.1996.10476701.

P. R. Jordan and L. J. Combs. *Summary of Floods in the United States during 1990 and 1991*. U.S. G.P.O. ; For sale by the U.S. Geological Survey, Branch of Information Services, Washington, Denver, CO, 1996. ISBN 978-0-607-86217-1.

I. Kant. *Critique of Pure Reason*. Penguin Classics. Penguin Books, London, 2007. ISBN 978-0-14-044747-7.

D. Kavetski, G. Kuczera, and S. W. Franks. Bayesian analysis of input uncertainty in hydrological modeling: 1. Theory. *Water Resources Research*, 42(3), 2006. doi: 10.1029/2005WR004368.

D. P. Kingma and M. Welling. Auto-Encoding Variational Bayes, 2022.

D. Klotz, F. Kratzert, M. Gauch, A. Keefe Sampson, J. Brandstetter, G. Klambauer, S. Hochreiter, and G. Nearing. Uncertainty estimation with deep learning for rainfall–runoff modeling. *Hydrology and Earth System Sciences*, 26(6):1673–1693, 2022. doi: 10.5194/hess-26-1673-2022.

D. Klotz, M. Gauch, F. Kratzert, G. Nearing, and J. Zscheischler. Technical Note: The Divide and Measure Nonconformity. *Hydrology and Earth System Sciences Discussions*, pages 1–25, 2024. doi: 10.5194/hess-2024-59.

D. Klotz, P. Miersch, T. V. M. Do Nascimento, F. Fenicia, M. Gauch, and J. Zscheischler. EARLS: A runoff reconstruction dataset for Europe, 2025.

D. Koutsoyiannis and A. Montanari. Bluecat: A Local Uncertainty Estimator for Deterministic Simulations and Predictions. *Water Resources Research*, 58(1):e2021WR031215, 2022. doi: 10.1029/2021WR031215.

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

F. Kratzert, D. Klotz, G. Shalev, G. Klambauer, S. Hochreiter, and G. Nearing. Towards learning universal, regional, and local hydrological behaviors via machine learning applied to large-sample datasets. *Hydrology and Earth System Sciences*, 23(12):5089–5110, 2019. doi: 10.5194/hess-23-5089-2019.

R. Krzysztofowicz. The case for probabilistic forecasting in hydrology. *Journal of Hydrology*, 249 (1):2–9, 2001. doi: 10.1016/S0022-1694(01)00420-6.

G. Kuczera. Improved parameter inference in catchment models: 1. Evaluating parameter uncertainty. *Water Resources Research*, 19(5):1151–1162, 1983. doi: 10.1029/WR019i005p01151.

F. Laio and S. Tamea. Verification tools for probabilistic forecasts of continuous hydrological variables. *Hydrol. Earth Syst. Sci.*, 2007.

I. Lakatos. Falsification and the Methodology of Scientific Research Programmes. In A. Musgrave and I. Lakatos, editors, *Criticism and the Growth of Knowledge: Proceedings of the International Colloquium in the Philosophy of Science, London, 1965*, volume 4, pages 91–196. Cambridge University Press, Cambridge, 1970. ISBN 978-0-521-09623-2. doi: 10.1017/CBO9781139171434.009.

R. McElreath. *Statistical Rethinking: A Bayesian Course with Examples in R and Stan*. CRC Texts in Statistical Science. Taylor and Francis, CRC Press, Boca Raton, 2 edition, 2020. ISBN 978-0-367-13991-9.

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

G. S. Nearing, Y. Tian, H. V. Gupta, M. P. Clark, K. W. Harrison, and S. V. Weijs. A philosophical basis for hydrological uncertainty. *Hydrological Sciences Journal*, 61(9):1666–1678, 2016. doi: 10.1080/02626667.2016.1183009.

G. S. Nearing, B. L. Ruddell, A. R. Bennett, C. Prieto, and H. V. Gupta. Does Information Theory Provide a New Paradigm for Earth Science? Hypothesis Testing. *Water Resources Research*, 56(2):e2019WR024918, 2020. doi: 10.1029/2019WR024918.

S. Nevo, E. Morin, A. Gerzi Rosenthal, A. Metzger, C. Barshai, D. Weitzner, D. Voloshin, F. Kratzert, G. Elidan, G. Dror, G. Begelman, G. Nearing, G. Shalev, H. Noga, I. Shavitt, L. Yuklea, M. Royz, N. Giladi, N. Peled Levi, O. Reich, O. Gilon, R. Maor, S. Timnat, T. Shechter, V. Anisimov, Y. Gigi, Y. Levin, Z. Moshe, Z. Ben-Haim, A. Hassidim, and Y. Matias. Flood forecasting with machine learning models in an operational framework. *Hydrology and Earth System Sciences*, 26(15):4013–4032, 2022. doi: 10.5194/hess-26-4013-2022.

A. J. Newman, M. P. Clark, K. Sampson, A. Wood, L. E. Hay, A. Bock, R. J. Viger, D. Blodgett, L. Brekke, J. R. Arnold, T. Hopson, and Q. Duan. 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, 2015. doi: 10.5194/hess-19-209-2015.

P. Reiser, J. E. Aguilar, A. Guthke, and P.-C. Bürkner. Uncertainty quantification and propagation in surrogate-based Bayesian inference. *Statistics and Computing*, 35(3):66, 2025. doi: 10.1007/s11222-025-10597-8.

S. Reuschen, W. Nowak, and A. Guthke. The Four Ways to Consider Measurement Noise in Bayesian Model Selection—And Which One to Choose. *Water Resources Research*, 57(11): e2021WR030391, 2021. doi: 10.1029/2021WR030391.

A. Schöniger, T. Wöhling, L. Samaniego, and W. Nowak. Model selection on solid ground: Rigorous comparison of nine ways to evaluate Bayesian model evidence. *Water Resources Research*, 50(12):9484–9513, 2014. doi: 10.1002/2014WR016062.

B. W. Silverman. *Density Estimation for Statistics and Data Analysis*. Number 26 in Monographs on Statistics and Applied Probability. Chapman & Hall/CRC, Boca Raton, 1998. ISBN 978-0-412-24620-3.

E. Todini. Flood Forecasting and Decision Making in the new Millennium. Where are We? *Water Resources Management*, 31(10):3111–3129, 2017. doi: 10.1007/s11269-017-1693-7.

E. Vasheghani Farahani, A. R. Massah Bavani, and A. Roozbahani. Enhancing reservoir inflow forecasting precision through Bayesian Neural Network modeling and atmospheric teleconnection pattern analysis. *Stochastic Environmental Research and Risk Assessment*, 39(1):205–229, 2025. doi: 10.1007/s00477-024-02858-1.

J. A. Vrugt. Markov chain Monte Carlo simulation using the DREAM software package: Theory, concepts, and MATLAB implementation. *Environmental Modelling & Software*, 75:273–316, 2016. doi: 10.1016/j.envsoft.2015.08.013.

J. A. Vrugt. Distribution-Based Model Evaluation and Diagnostics: Elicitability, Propriety, and Scoring Rules for Hydrograph Functionals. *Water Resources Research*, 60(6):e2023WR036710, 2024. doi: 10.1029/2023WR036710.

J. A. Vrugt and M. Sadegh. Toward diagnostic model calibration and evaluation: Approximate Bayesian computation. *Water Resources Research*, 49(7):4335–4345, 2013. doi: 10.1002/wrcr.20354.

R. L. Winkler. A Decision-Theoretic Approach to Interval Estimation. *Journal of the American Statistical Association*, 67(337):187–191, 1972. doi: 10.1080/01621459.1972.10481224.