Enhancing Hydrological Modeling with Transformers:
A Case Study for 24-Hour Streamflow Prediction

Bekir Z. Demiray[a]*, Muhammed Sit[a], Omer Mermer[a], Ibrahim Demir[a,b,c]

[a] IIHR—Hydroscience and Engineering, University of Iowa, Iowa City, Iowa, USA
[b] Civil and Environmental Engineering, University of Iowa, Iowa City, Iowa, USA
[c] Electrical and Computer Engineering, University of Iowa, Iowa City, Iowa, USA

* Corresponding Author, Email: bekirzahit-demiray@uiowa.edu

Abstract
In this paper, we address the critical task of 24-hour streamflow forecasting using advanced deep-learning models, with a primary focus on the Transformer architecture which has seen limited application in this specific task. We compare the performance of five different models, including Persistence, LSTM, Seq2Seq, GRU, and Transformer, across four distinct regions. The evaluation is based on three performance metrics: Nash-Sutcliffe Efficiency (NSE), Pearson's r, and Normalized Root Mean Square Error (NRMSE). Additionally, we investigate the impact of two data extension methods: zero-padding and persistence, on the model's predictive capabilities. Our findings highlight the Transformer's superiority in capturing complex temporal dependencies and patterns in the streamflow data, outperforming all other models in terms of both accuracy and reliability. The study's insights emphasize the significance of leveraging advanced deep learning techniques, such as the Transformer, in hydrological modeling and streamflow forecasting for effective water resource management and flood prediction.

Keywords:
Rainfall-runoff modeling; deep learning; flood forecasting; transformers; streamflow forecasting

1. Introduction

Globally, the incidence and catastrophic effects of natural disasters have increased dramatically. The World Meteorological Organization's analysis (2021) shows that, on average, each day for the past half-century, a weather, climate, or water-related disaster has led to a loss of $202 million and claimed 115 lives. Further, Munich Re's (2022) report indicated that natural catastrophes, encompassing hurricanes, floods, and other disaster types, have inflicted more than $280 billion in projected damage worldwide. Out of this total, disasters caused $145 billion in damages in the United States alone, along with thousands of fatalities and substantial damage to properties and infrastructure. Current research suggests that ongoing climate change is projected to cause an upsurge in extreme and intense natural disasters globally, leading to an increase in the number of victims and losses (WMO, 2021; Banholzer et al., 2014).

Floods are the most commonly occurring natural disaster, leading to billions in financial losses and innumerable fatalities over time (WHO, 2021). In the year 2020, over 60% of all reported natural disasters were flood-related, accounting for 41% of the overall death toll due to such events (NDRCC, 2021). Multiple studies suggest that climate change is causing an escalation in the frequency and severity of floods in specific areas (Davenport et al., 2021; NOAA, 2022; Tabari, 2020). This rise in flooding events can be attributed to factors like an increase in sea level (Strauss et al., 2016), the heightened occurrence of extreme rainfall (Diffenbaugh et al., 2017), or amplified rainfall during hurricanes (Trenberth et al., 2018). Hence, accurately forecasting streamflow and, as a result, potential flooding is essential for effectively mitigating the destructive consequences in terms of property damage and fatalities (Alabbad and Demir, 2022).

In addition, streamflow forecasting plays a vital role in numerous aspects of hydrology and water management, including watershed management (Demir and Beck, 2009), agricultural planning (Yildirim and Demir, 2022), flood mapping systems (Li and Demir, 2022), and other mitigation activities (Ahmed et al., 2021; Yaseen et al., 2018). Yet, achieving accurate and reliable predictions poses a challenge due to the inherent complexity of hydrological systems, which include nonlinearity, and unpredictability in the datasets (Honorato et al., 2018; Yaseen et al., 2017, Sit et al., 2023a).

Over time, a plethora of physical and data-driven methods have been introduced, each exhibits diverse characteristics such as employing different types of data, focusing on specific geographical areas, or offering varying levels of generalization (Salas et al., 2000; Yaseen et al., 2015). Physics-driven prediction models (Beven and Kirkby, 1979; Ren-Jun, 1992; Arnold, 1994; Lee and Georgakakos, 1996; Devia et al., 2015) have the capability to simulate the complex interactions among different physical processes, including atmospheric circulation and the long-term evolution of weather patterns in the world (Yaseen et al., 2019; Sharma and Machiwal, 2021). However, these models, while valuable, come with notable limitations. They demand extensive and precise hydrological and geomorphological data, increasing operational costs. The accuracy of these models tends to wane in long-term forecasting scenarios.

Furthermore, due to their computational intensity and high parameter counts, traditional physically-based hydrological models require substantial computing resources, leading to significant computational costs (Mosavi et al., 2018; Sharma and Machiwal, 2021; Liu et al., 2022; Castangia et al., 2023). As a result, recent research (Yaseen et al., 2015) has explored alternative approaches to streamflow forecasting, indicating that machine learning, especially deep learning models, can serve as viable alternatives and often outperform physically-based models in terms of accuracy. These deep learning models have shown promising results in enhancing the accuracy and reliability of streamflow predictions, presenting an opportunity to revolutionize hydrological modeling (Demiray et al., 2023; Sit et al., 2023b).

Many classical machine-learning approaches have been used in streamflow forecasting and environmental studies (Bayar et al., 2009; Li and Demir, 2023) including Support Vector Machines (SVMs) and Linear Regression (LR) (Granata et al., 2016; Yan et al., 2018; Sharma and Machiwal, 2021). However, advancements in artificial intelligence (AI) coupled with the increasing capabilities of graphics processing units (GPUs) have opened up new possibilities and accelerated the progress of deep learning techniques, which has led to the widespread usage of these techniques in streamflow forecasting as well (Sit et al., 2022a). Out of various neural network architectures explored for streamflow forecasting (Sit et al., 2021a; Xiang and Demir, 2022b; Chen et al., 2023), Recurrent Neural Networks (RNNs), especially the Long Short-Term Memory (LSTM) neural network and Gated Recurrent Units (GRUs), have emerged as the most extensively studied and researched models in this domain.

Kratzert et al. (2018) applied an LSTM model to predict daily runoff, incorporating meteorological observations, and demonstrated that the LSTM model outperformed a well-established physical model in their study area. In their study, Xiang et al. (2021) demonstrated that the LSTM-seq2seq model surpasses the performance of other linear models, such as linear regression, lasso regression, and ridge regression methods. The LSTM-seq2seq model outperformed these linear models in terms of predictive accuracy or other relevant evaluation metrics. Guo et al., (2021) compared LSTMs, GRUs, and SVMs over 25 different locations in China and found that while LSTMs and GRUs demonstrated comparable performance, GRUs exhibited faster training times. Since the research about the field is extensive, more detailed information about deep learning studies on streamflow prediction can be found in (Yaseen et al., 2015; Ibrahim et al., 2022).

In 2017, a group of researchers from Google introduced a new way to model longer sequences for language translation (Vaswani et al., 2017) and this new model, namely Transformers, was applied to various tasks since then including time series prediction (Zhou et al., 2021; Wu et al., 2021; Zhou et al., 2022; Lin et al., 2022). Despite attention from other fields, there is a limited number of studies that focus on the performance and usage of transformers in streamflow forecasting. Liu et al. (2022) introduced a Transformer neural network model for monthly streamflow prediction of the Yangtze River in China. Their approach utilized historical water levels and incorporated the El Niño-Southern Oscillation (ENSO) as additional input features. This allowed the model to capture the influence of ENSO on streamflow patterns and

improve the accuracy of monthly streamflow predictions for the Yangtze River. More recently, Castangia et al. (2023) used a Transformer based model to predict the water level of a river one day in advance, leveraging the historical water levels of its upstream branches as predictors. They conducted experiments using data from the severe flood that occurred in Southeast Europe in May 2014.

In this work, we investigate the performance of a Transformer model in streamflow forecasting for four different locations in Iowa, US. More specifically, we predict the upcoming 24-hour water levels using the previous 72-hour precipitation, evapotranspiration, and discharge values, then compare the results of Transformer based model with three deep learning models as well as the persistence method. According to experiment results, Transformer based model outperforms all tested methods.

The structure of the remaining sections of this paper is as follows: in the next section, the dataset that has been used in this research and study area will be introduced. Section 3 outlines the methods employed in this study. Following that, Section 4 presents the results of our experiments and provides a detailed discussion of the findings. Finally, in Section 5, we summarize the key findings of this study and discuss future prospects.

2. Dataset

WaterBench, developed by Demir et al. (2022), is a benchmark dataset explicitly created for flood forecasting research, adhering to FAIR (findability, accessibility, interoperability, and reuse) data principles. Its structure is designed for easy application in data-driven and machine-learning studies, and it also provides benchmark performance metrics for advanced deep-learning architectures, enabling comparative analysis. This dataset has been compiled by gathering streamflow, precipitation (Sit et al., 2021b), watershed area, slope, soil types, and evapotranspiration data from various federal and state entities, including NASA, NOAA, USGS, and the Iowa Flood Center. This consolidated resource is specifically geared towards studies of hourly streamflow forecasts.

The dataset's time-series spans from October 2011 to September 2018. In this work, four different U.S. Geological Survey (USGS) stations, each one of them located in a different watershed, are selected from WaterBench. More specifically, USGS 05387440 Upper Iowa River at Bluffton, USGS 05418400 North Fork Maquoketa River near Fulton, USGS 05454000 Rapid Creek near Iowa City, and USGS 06817000 Nodaway River at Clarinda are selected. Figure 1 illustrates the locations of the designated sites and their corresponding watersheds within the State of Iowa.

The data from October 2011 to September 2017 is selected for the training set. The rest of the data is used for evaluation and testing. As a preprocessing step, we followed the same methods in the original dataset paper (Demir et al., 2022) since we compared our results with the models provided in the WaterBench paper. The data and benchmark models can be accessible from https://github.com/uihilab/WaterBench. The statistical summary of streamflow values in used test data is provided in Table 1.
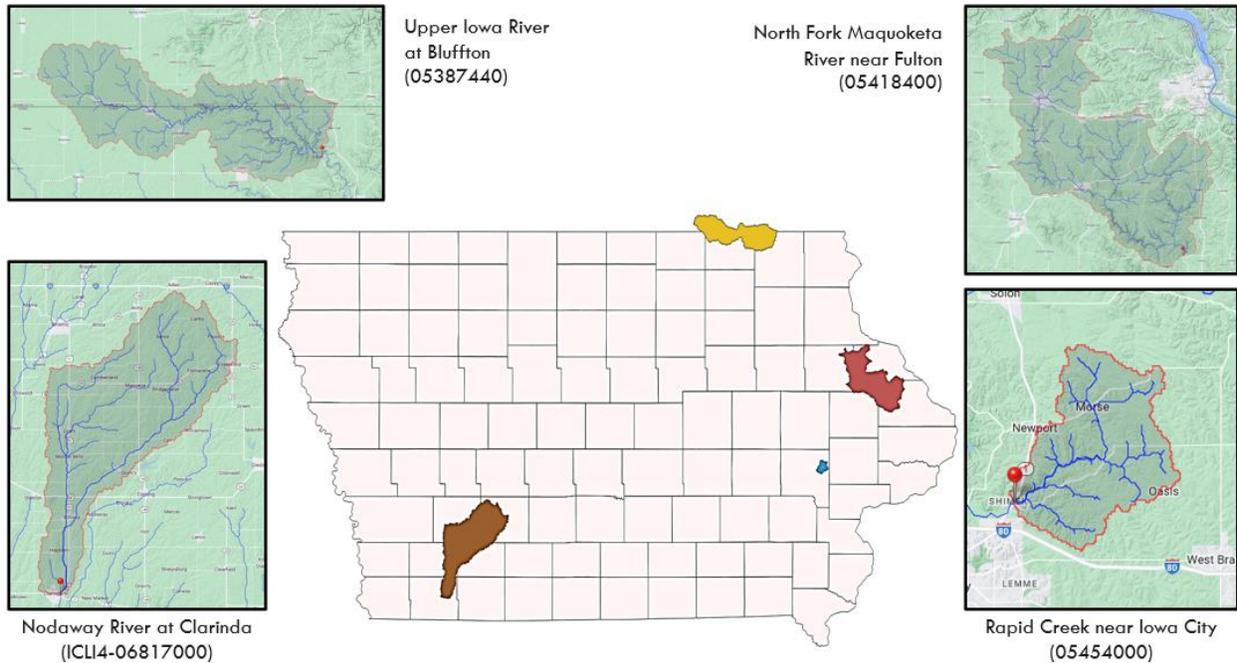
Figure 1: Selected Locations and Corresponding Watersheds in the State of Iowa

Table 1: Statistical Summary of Streamflow Values in Test Data (m³/s)

	Bluffton	Fulton	Iowa City	Clarinda
Max	13050.00	10075.00	2242.50	11575.00
Min	41.09	121.99	0.16	70.87
Mean	436.98	425.59	12.76	443.70
Median	246.00	308.00	4.28	256.00

3. Methods

In this study, we evaluated the Transformer-based model in streamflow prediction tasks and compared the results with the four models, (Persistence, GRU, LSTM, and Seq2Seq), that are mentioned and provided in the WaterBench dataset. In this section, we will provide the details of these methods as well as the Transformer-based approach.

3.1. Persistence Approach

Persistence (Eq. 1), also known as the nearest frame approach, is based on the principle that "tomorrow will be the same as today." In other words, persistence forecasts rely solely on the most recent available data and assume that future conditions will remain unchanged from the present. It is accepted as one of the baselines for hydrological studies including streamflow forecasting and several hydrological studies have indicated that the fundamental persistence model is challenging to surpass in terms of short-range predictions, especially when the forecasting lead time (n) is less than 12 hours (Krajewski et al., 2021; Demir et al., 2022).

$$\hat{Y}_{i+n} = Y_i \qquad\qquad\qquad \text{Eq. 1}$$

$$\hat{Y}_{i+n} = Predicted\ streamflow\ values\ between\ time\ t\ and\ t + n$$
$$Y_i = Observed\ streamflow\ values\ at\ time\ t$$

3.2. LSTM Model

In the context of time-series forecasting, Recurrent Neural Networks (RNNs) have proven effective in capturing temporal dependencies. However, they suffer from the vanishing gradient problem, where the gradient diminishes exponentially over time, hindering the model's ability to retain long-term dependencies. This limitation impacts the accuracy of time-series predictions, particularly for tasks that require memory of events far back in the past. To address these shortcomings, Long Short-Term Memory (LSTM) networks were introduced by Hochreiter and Schmidhuber (1997). LSTMs are designed to extend the lifespan of short-term memory and effectively capture long-term dependencies in the data. This makes them well-suited for time series problems consequently hydrological forecasting tasks as well that involve longer memory requirements, such as flood and rainfall forecasting (Kratzert et al., 2018; Feng et al., 2020; Frame et al., 2022; Sit et al., 2022b).

An LSTM node operates by receiving input tensor x_t and the hidden state tensor h_{t-1} from the previous LSTM node in the sequence. The LSTM cell consists of several gates, each performing specific operations to control the flow of information and memory retention. Input gate (i_t) decides how much of the new information should be stored in the cell state. It is calculated using a sigmoid function after linear transformations of the current input x_t and the previous hidden state tensor h_{t-1}, utilizing weight matrices $W(i)$ and $U(i)$. Forget gate (f_t) determines what information should be discarded from the cell state. Like the input gate, it employs a sigmoid function and weight matrices $W(f)$ and $U(f)$ to control the retention of information from the previous cell state c_{t-1}. Output gate (O_t) regulates how much of the cell state should be exposed as the hidden state of the LSTM node. It employs a sigmoid function and weight matrices $W(O)$ and $U(O)$ to control the influence of the current input x_t and the previous hidden state tensor h_{t-1} on the cell state. Candidate cell state ($\tilde{c}$) represents the new information that could be added to the cell state. It is calculated using the hyperbolic tangent function ($tanh$) after linear transformations of the current input x_t and the previous hidden state h_{t-1}, utilizing weight matrices $W(c)$ and $U(c)$.

The cell state c_t is updated based on the input, forget, and candidate cell state using element-wise operations as shown in Equation 2.

$$c_t = f_t \odot c_{t-1} + i_t \odot \tilde{c}_t \qquad\qquad \text{Eq. 2}$$

Finally, the hidden state h_t is computed by applying the output gate to the hyperbolic tangent of the updated cell state as shown in Equation 3.

$$h_t = O_t \odot \tanh(c_t) \qquad\qquad \text{Eq. 3}$$

The updated hidden state h_t and cell state c_t are then passed to the next LSTM node in the sequence and to subsequent layers in the neural network architecture. In hydrological forecasting tasks, LSTM networks have demonstrated superior performance compared to basic RNNs and other time-series forecasting models, making them a popular choice in the field. By effectively addressing the vanishing gradient problem and capturing long-term dependencies, LSTMs have proven to be valuable tools for accurate and reliable predictions in various hydrological applications.

3.3. GRU Model

While LSTM networks have been instrumental in addressing the vanishing gradient problem and achieving remarkable progress in Natural Language Processing and time-series prediction, their time complexity can be a concern, especially for large-scale applications. To mitigate this issue, Gated Recurrent Unit (GRU) networks were introduced by Cho et al. (2014) as an efficient alternative that retains the effectiveness of LSTM while reducing computational burden.

The detailed workings of a GRU are governed by a series of equations that outline the core operations of the model. The GRU cell starts with the computation of the updating gate, denoted as z_t, which determines the extent of information to be retained from the previous hidden state h_{t-1} and the current input x_t. This gate is calculated by applying a sigmoid function to the linear transformations of h_{t-1} and x_t, which are achieved using learnable weight parameters W_z and U_z (Equation 4).

$$z_t = \sigma(W_z \cdot x_t + U_z \cdot h_{t-1}) \qquad\qquad \text{Eq. 4}$$

Subsequently, the reset gate r_t is computed in the next stage, controlling the amount of information to be discarded from the prospective new cell state $\check{h}_t$. Similar to the updating gate, the reset gate r_t is derived through a sigmoid function applied to the linear transformations of h_{t-1} and x_t using weight parameters W_r and U_r (Equation 5).

$$r_t = \sigma(W_r \cdot x_t + U_r \cdot h_{t-1}) \qquad\qquad \text{Eq. 5}$$

In the third step, the candidate hidden state $\check{h}_t$ is calculated, representing the new information to be considered for the current timestep. It is obtained by applying the hyperbolic tangent (tanh) function to the linear transformations of x_t and the element-wise product of r_t and h_{t-1} (Equation 6). This element-wise product allows the GRU to control the extent to which the previous state influences the candidate state.

$$\tilde{h}_t = \tanh(W_h \cdot x_t + r_t \odot U_h \cdot h_{t-1}) \qquad\qquad \text{Eq. 6}$$

Finally, the actual cell state h_t at the current timestep is updated by blending the previous state h_{t-1} and the candidate state $\check{h}_t$, with the proportions determined by the updating gate z_t (Equation 7).

$$h_t = (1 - z_t) \odot h_{t-1} + z_t \odot \tilde{h}_t \qquad \text{Eq. 7}$$

The resulting cell state h_t is then passed on to the next timestep within the same layer and to the subsequent layers at the same timestep, facilitating information flow throughout the network. By employing these gating mechanisms and streamlined computations, the GRU model strikes a balance between computational efficiency and predictive performance, making it an appealing choice for various time-series forecasting tasks, including 24-hour streamflow prediction.

3.4. Seq2Seq Model

In addition to LSTM and GRU, a variant of the Seq2Seq model (Xiang and Demir, 2022a) is also employed as a baseline method in this study. The Seq2Seq model follows an encoder-decoder architecture and utilizes multiple TimeDistributed layers with a final dense layer. The encoder-decoder structure of the Seq2Seq model consists of two main components: an encoder and a decoder. The encoder processes the input time series data and encodes it into a fixed-size context vector, effectively capturing relevant temporal patterns and features. For this implementation, multiple Gated Recurrent Units (GRUs) are used as both the encoder and decoder, proven effective in modeling sequential data and handling long-range dependencies.

During the encoding process, the input time series data, including historical rainfall, streamflow, and evapotranspiration for the past 72 hours, along with 24-hour forecast data of rainfall and evapotranspiration, is passed through the multiple GRUs. The encoder generates a context vector that summarizes the important information from the input sequence. Next, the decoder takes the context vector produced by the encoder and predicts the future 24-hour streamflow. The decoder GRUs process the context vector along with the predicted streamflow values from the previous timestep, iteratively generating the streamflow predictions for the next 24 hours. To capture intricate patterns and temporal dynamics in the predictions, multiple TimeDistributed layers are employed, applying the same dense layer to each timestep of the output sequence. Finally, the Seq2Seq model concludes with a final dense layer that projects the output sequence to the desired format for 24-hour streamflow predictions. For comprehensive implementation details, we recommend referring to the works by Xiang and Demir (2022a) and Demir et al., (2022).

3.5. Transformer Model

The Transformer model represents a revolutionary neural network architecture that emerged as a seminal work by Vaswani et al. (2017) to tackle challenges in machine translation tasks. Its groundbreaking design subsequently found applications in various domains that deal with long input sequences, including time series forecasting (cites). The Transformer's key innovation lies

in the self-attention mechanism, which completely replaces traditional recurrent layers, enabling more efficient and effective analysis of extended input sequences.

The self-attention computation in the Transformer can be broken down into several stages to reveal its inner workings. Initially, each element in the input sequence is projected into three distinct representations - query (Q), key (K), and value (V) vectors of dimension d_{model}. The self-attention scores are then obtained by performing a dot-product operation between the query and key matrices, followed by scaling and applying a softmax function to capture the importance of each element in relation to others. Consequently, the weighted sum of the value vectors produces a new representation of the input sequence. By employing self-attention, the Transformer can dynamically adjust the representation of each element, factoring in the influence of all other elements in the sequence. This enables distant elements to contribute meaningfully to each other, fostering the capture of long-range dependencies that may be crucial for accurate time series forecasting. The self-attention mechanism can be mathematically represented as follows (Equation 8):

$$Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_{model}}}\right) V \qquad \text{Eq. 8}$$

To enhance the model's ability to capture diverse patterns, Transformer model employs multi-head attention. The query, key, and value vectors are divided into multiple chunks of dimension d_{model}/h, where h is the number of attention heads. Each head independently computes the self-attention process for each chunk, and the resulting representations are concatenated and subjected to a final linear transformation. The introduction of multi-head attention increases the potential combinations between elements, thereby enhancing the model's ability to capture intricate relationships within the input sequence.

As self-attention inherently lacks information about the order of elements in the input sequence, static positional encoding is introduced to provide positional awareness. The positional encoding is added to the initial input embedding, ensuring that the model distinguishes the positions of different elements. Each element's position is encoded using a specific formula, involving positional index and sine and cosine functions. The positional encoding can be mathematically represented as follows: (Equation 9).

$$PE(pos, 2i) = \sin\left(pos/10000^{2i/d_{model}}\right)$$
$$PE(pos, 2i + 1) = \cos\left(pos/10000^{2i/d_{model}}\right) \qquad \text{Eq. 9}$$

The Transformer model employed in this study deviates slightly from the original implementation by Vaswani et al. (2017). Notably, all decoder layers are removed, as the task specifically focuses on time series forecasting without the need for decoding. Additionally, a different input embedding technique is utilized, where the input sequence is transposed and passed through a convolutional layer before positional encoding. Dropout layers are incorporated

for regularization after adding positional encoding. A final linear layer is employed to reduce the feature size to 1 at the end of the last encoder, given the model has no decoders and a single output value required for forecasting. The model is depicted in Figure 2.
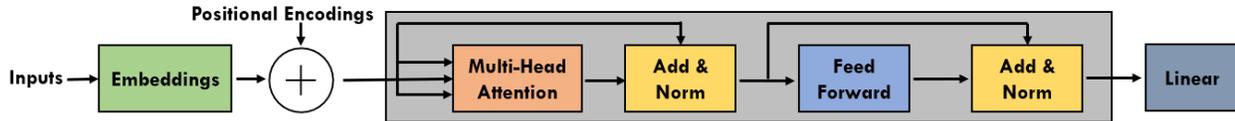

Figure 2: Transformer model architecture

The persistence, GRU, LSTM, and Transformer models are developed with Pytorch, whereas the Seq2Seq model is developed with Keras. Please refer to Demir et al. (2022) for further implementation details and model architectures of GRU, LSTM and Seq2Seq utilized in this study. During the training of each model, we used Mean Squared Error (MSE) as the loss function and Adam as the optimizer. In addition, we set the batch sizes to 32 for each model and the learning rate to 0.00001. The learning rate is divided by two if no improvement is noticed for 10 epochs and training is frozen if there was no improvement for 20 epochs.

4. Results and Discussion

In this section, we present and discuss the findings from our research into the 24-hour prediction of streamflow using different models, focusing primarily on the performance of the Transformer model we used. To assess its effectiveness, we compare it against four other models, three of which are deep learning models - LSTM, GRU, and Seq2Seq - and one is a classical approach known as Persistence.

Streamflow prediction holds immense significance in various domains such as water resource management, environmental monitoring, and decision-making processes. Deep learning models have demonstrated remarkable capabilities in time-series forecasting tasks, making them a natural choice for tackling streamflow prediction challenges. However, the application of Transformer models in this specific context is relatively new and deserving of detailed investigation. The Transformer's self-attention mechanism has shown great promise in sequence modeling tasks, making it an intriguing candidate for capturing temporal dependencies in streamflow data.

Our comparative analysis employs three metrics, namely: Nash-Sutcliffe Efficiency (NSE), Pearson's r, and Normalized Root Mean Square Error (NRMSE). Each of these metrics serves to facilitate a thorough and multidimensional understanding of each model's predictive capacities and the effectiveness of the Transformer model. The subsequent sections delve into a detailed exposition of the three-evaluation metrics and their relevance in streamflow prediction assessment. Following that, we meticulously present and analyze the results obtained from each model, highlighting their respective strengths and limitations. Through this thorough examination, we aim to uncover the effectiveness of the Transformer model in 24-hour

streamflow prediction and its potential implications for future research and real-world applications.

4.1. Performance Metrics

In the evaluation of streamflow prediction models, several performance metrics are commonly employed to assess the accuracy and reliability of the forecasts. In this study, we utilized three widely accepted metrics: Nash-Sutcliffe Efficiency (NSE), Pearson's r, and Normalized Root Mean Square Error (NRMSE). These metrics have been extensively applied in hydrological modeling and streamflow forecasting research due to their interpretability and ability to capture different aspects of model performance (Kratzert et al., 2018; Xiang and Demir, 2021; Liu et al., 2022).

Firstly, Nash-Sutcliffe Efficiency (Equation 8) is a widely used metric to quantify the predictive performance of hydrological models (Krause et al., 2005; Arnold et al., 2012). It provides a measure of how well the model predictions match the observed streamflow data, relative to the mean of the observed data. The NSE ranges from negative infinity to 1. A value of 1 indicates a perfect match, where the model predictions precisely align with the observations. Values greater than 0 denote better performance than using the mean of the observed data as a predictor. On the other hand, negative NSE values signify that the mean of the observed data outperforms the model, indicating poor predictive ability. Values greater than 0.5 are considered acceptable in hydrological modeling (Arnold et al., 2012).

$$NSE = 1 - \frac{\sum_{i=1}^{n}(Y_i - \hat{Y}_i)^2}{\sum_{i=1}^{n}(Y_i - \bar{Y}_i)^2} \qquad \text{Eq. 8}$$

$$Y_i = Observed\ streamflow\ value\ at\ time\ i$$
$$\hat{Y}_i = Predicted\ streamflow\ value\ at\ time\ i$$
$$\bar{Y}_i = Mean\ of\ all\ observations\ at\ time\ i$$

Secondly, Pearson's correlation coefficient (r), also known as Pearson's r or simply r, is a statistical measure used to assess the linear relationship between the model's predicted streamflow values and the observed streamflow data. Pearson's r directly quantifies the strength and direction of the linear association between the predicted and observed streamflow values. Ranging from -1 to 1, with 1 indicating a perfect positive linear relationship, a higher positive Pearson's r value signifies a more reliable and accurate model, capable of accurately capturing the patterns in the observed data and making precise predictions. Utilizing Pearson's r allows us to measure the accuracy and effectiveness of the forecasting models in capturing the variability of the observed streamflow for the 24-hour prediction horizon.

$$r = \frac{\sum_{i=1}^{n}(Y_i - \bar{Y}_i)(\hat{Y}_i - \bar{\hat{Y}}_i)}{\sqrt{\sum_{i=1}^{n}(Y_i - \bar{Y}_i)^2}\sqrt{\sum_{i=1}^{n}(\hat{Y}_i - \bar{\hat{Y}}_i)^2}} \qquad \text{Eq. 9}$$

$$Y_i = Observed\ streamflow\ value\ at\ time\ i$$
$$\hat{Y}_i = Predicted\ streamflow\ value\ at\ time\ i$$
$$\bar{Y}_i = Mean\ of\ all\ observations\ at\ time\ i$$
$$\bar{\hat{Y}}_i = Mean\ of\ all\ predicted\ values\ at\ time\ i$$

Lastly, Normalized Root Mean Square Error (NRMSE) measures the average error between the predicted and observed streamflow values, normalized by the mean of the observed data. It provides a relative measure of the model's predictive accuracy, enabling comparison across different datasets. Since this study uses different locations, it seems reasonable to use NRMSE. The NRMSE ranges from 0 to 1, with lower values indicating better predictive performance and higher values indicating larger errors relative to the mean of the observed streamflow.

$$NRMSE = \frac{\sqrt{\frac{\sum_{i=1}^{n}(Y_i - \hat{Y}_i)^2}{n}}}{\bar{Y}_i} \quad \text{Eq. 10}$$

$$Y_i = Observed\ streamflow\ value\ at\ time\ i$$
$$\hat{Y}_i = Predicted\ streamflow\ value\ at\ time\ i$$
$$\bar{Y}_i = Mean\ of\ all\ observations\ at\ time\ i$$

These metrics serve as essential tools in quantifying the predictive performance of our streamflow prediction models, enabling us to assess their effectiveness in capturing the underlying patterns and dynamics of streamflow behavior.

4.2. Experiment Results

In this section, we present the experiment results that address the core objective of our study: 24-hour streamflow prediction. For this investigation, we utilized a comprehensive dataset comprising historical data on precipitation, evapotranspiration, discharge values from the preceding 72 hours, as well as forecast data of 24-hour precipitation and evapotranspiration. Our investigation focused on evaluating the performance of the Transformer-based model, comparing it against three deep learning models (LSTM, GRU, and Seq2Seq), and a classical method (Persistence). To assess the predictive capabilities of these models, we employed three commonly used metrics in hydrological modeling and streamflow forecasting: Nash-Sutcliffe Efficiency (NSE), Pearson's r, and Normalized Root Mean Square Error (NRMSE). These metrics provide valuable insights into the accuracy and effectiveness of the models in capturing streamflow patterns.

In the experiments, a crucial aspect involved adjusting the dimensions of the input data and incorporating additional values to accommodate the implementation specifications of GRU, LSTM, and Transformer models. More specifically, input data for these networks is combination of previous values and forecast values. Previous values are 72-hours of precipitation, evapotranspiration, and discharge values, for the forecast values 24-hours of precipitation, and

evapotranspiration information are used. So, one has a shape of [batch size, 72, 3] and other has [batch size, 24, 2]. To merge and align these two input groups for the models, an extra dimension for forecast values needed to be introduced. According to experiment results, what is used as additional dimension affects the results dramatically. Two approaches are considered to handle this dimension discrepancy. One approach is zero-padding, wherein the forecast values are extended with zeros in the additional dimension. Alternatively, the persistence method can be adopted, wherein the historical values were extended into the forecast period by repeating the last available data. This method ensured consistency in the input data across time steps. Both techniques are employed to ensure compatibility between the input data and the specific model requirements. Once the additional dimension is added, past and forecast values merge and input with dimension of [batch size, 96, 3] is obtained for Transformer, GRU and LSTM models.

Table 2: Performance Comparison of Transformer Model for 24-hr streamflow forecasting using zero-padding and persistence approaches in four different regions (NSE scores)

	Bluffton		Fulton		Iowa City		Clarinda	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median
Transformer-zero-padding	0.71	0.67	0.53	0.51	0.48	0.48	0.55	0.52
Transformer-persistence	0.77	0.74	0.62	0.58	0.57	0.57	0.61	0.62

The results in Table 2 demonstrate the performance comparison of the Transformer model using zero-padding and persistence approaches for 24-hour streamflow forecasting in four different regions. The NSE scores reveal valuable insights into the model's predictive capabilities under each data extension method. Upon analysis, it becomes evident that the persistence method for data extension consistently outperforms zero-padding in capturing underlying streamflow patterns and dynamics for the Transformer model in all four regions. These findings emphasize the critical role of data extension techniques in improving the Transformer model's performance for streamflow forecasting tasks.

Table 3: Performance comparison of GRU and LSTM models for 24-hr streamflow forecasting using zero-padding and persistence approaches in four different regions (NSE scores)

	Bluffton		Fulton		Iowa City		Clarinda	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median
GRU-zero-padding	0.56	0.54	0.42	0.45	0.12	0.12	-0.40	-0.45
GRU-persistence	0.72	0.72	0.62	0.64	0.19	0.19	0.57	0.59
LSTM-zero-padding	0.77	0.76	0.45	0.44	-0.45	-0.52	0.51	0.53
LSTM-persistence	0.50	0.50	0.40	0.41	-1.50	-1.60	0.09	0.09

Similar to Table 2, Table 3 displays the NSE scores obtained from the predictions made by the GRU (Gated Recurrent Unit) and LSTM (Long Short-Term Memory) models under the zero-padding and persistence data extension methods for each region. Upon analysis, we observe variations in the models' performance across the four regions. Interestingly, for the LSTM model, the zero-padding approach yields higher mean and median NSE scores compared to the persistence method. Conversely, for the GRU model, the persistence method consistently outperforms the zero-padding approach, resulting in higher mean and median NSE scores. In summary, the different performance trends for the three models under the zero-padding and persistence approaches highlight the significance of selecting appropriate data extension techniques in streamflow forecasting tasks, as the effectiveness can vary depending on the model architecture.

Table 4: 24-hr streamflow prediction results (NSE scores)

Region	Metric	Persistence	LSTM	Seq2Seq	GRU	Transformer
Bluffton	NSE	0.58	0.77	0.66	0.72	0.77
	r	0.77	0.88	0.85	0.86	0.88
	NRMSE	1.26	0.92	1.13	1.01	0.93
Fulton	NSE	0.46	0.45	0.58	0.62	0.62
	r	0.73	0.67	0.76	0.78	0.78
	NRMSE	0.98	0.99	0.87	0.83	0.82
Iowa City	NSE	-0.30	-0.45	0.01	0.19	0.57
	r	0.34	0.29	0.16	0.44	0.76
	NRMSE	6.36	6.65	5.53	4.95	3.64
Clarinda	NSE	0.48	0.51	0.29	0.57	0.61
	r	0.74	0.84	0.56	0.90	0.86
	NRMSE	1.23	1.20	1.43	1.08	1.07

Table 4 presents the 24-hour streamflow prediction results for four different regions using five different models: Persistence, LSTM, Seq2Seq, GRU, and Transformer. The results are evaluated using three performance metrics: mean of 24 hours of NSE scores, Pearson's r, and NRMSE. In this study, the Persistence model serves as the baseline for comparison. While it exhibits moderate performance in some regions, it falls short in capturing the underlying dynamics of streamflow, leading to higher NRMSE values. As expected, it shows limited predictive capabilities compared to the advanced deep learning models. The LSTM and Seq2Seq models demonstrate mixed results across regions. While they achieve reasonably high NSE scores in certain regions, they struggle to consistently outperform the Persistence model, especially in regions Iowa City and Clarinda.

This indicates that their recurrent architecture might face challenges in capturing the complex temporal dependencies in streamflow data. The GRU model showcases competitive performance across all regions. With consistent NSE scores and relatively lower NRMSE values compared to LSTM and Seq2Seq models, it proves its capability to effectively model the temporal dynamics in streamflow data. However, it still falls behind the Transformer model's overall performance.

The Transformer model emerges as the top-performing model in 24-hour streamflow prediction across all regions. With the highest NSE scores and the lowest NRMSE values among all models, the Transformer demonstrates its efficacy in capturing and learning the long-range dependencies and patterns in the time series data. The self-attention mechanism, along with positional encoding, enables the Transformer to effectively process and utilize the sequential information, leading to its superior predictive capabilities.

In conclusion, the experimental results highlight the Transformer model's significant advantage over other models in 24-hour streamflow forecasting. Its powerful self-attention mechanism allows it to efficiently capture and utilize the temporal dependencies in the input time series, resulting in more accurate and reliable predictions compared to traditional LSTM and Seq2Seq models, as well as the GRU model. These findings underscore the importance of leveraging advanced deep learning architectures like the Transformer in hydrological modeling and streamflow forecasting tasks, offering valuable insights for the research community and practical applications in water resource management and flood forecasting.

5. Conclusion

In this study, we conducted an in-depth investigation of 24-hour streamflow forecasting using various deep learning models, with a particular focus on the Transformer architecture. Through extensive experimentation and analysis, we compared the performance of five different models across four distinct regions. The results demonstrate that the Transformer model consistently outperforms other models, including Persistence, LSTM, Seq2Seq, and GRU, in terms of accuracy and predictive capabilities. The Transformer's powerful self-attention mechanism, along with positional encoding, enables it to effectively capture long-range dependencies and underlying patterns in the input time series data. Consequently, the Transformer model excels in providing accurate and reliable streamflow predictions.

Furthermore, we explored the influence of two data extension methods: zero-padding and persistence, on the model's performance. The findings indicate that the persistence method, which incorporates historical streamflow data, consistently yields superior results compared to zero-padding. This underscores the importance of carefully considering data extension techniques to improve the model's forecasting accuracy.

Overall, our research contributes valuable insights into the field of hydrological modeling and streamflow forecasting. The superior performance of Transformer model highlights its potential as a promising tool for water resource management, flood prediction, and other hydrological applications. As future work, we suggest exploring the applicability of the Transformer in handling larger datasets and further investigating the impact of different hyperparameters on the model's performance. The knowledge gained from this study can significantly benefit water management practices, supporting sustainable decision-making and mitigation strategies in the face of increasingly unpredictable weather patterns and climate change.

References

Ahmed, A.M., Deo, R.C., Feng, Q., Ghahramani, A., Raj, N., Yin, Z. and Yang, L. (2021). Deep learning hybrid model with Boruta-Random forest optimiser algorithm for streamflow forecasting with climate mode indices, rainfall, and periodicity. *Journal of Hydrology*, *599*, p.126350.

Alabbad, Y., & Demir, I. (2022). Comprehensive flood vulnerability analysis in urban communities: Iowa case study. International journal of disaster risk reduction, 74, 102955.

Arnold, J., 1994. SWAT-soil and water assessment tool.

Arnold, J.G., Moriasi, D.N., Gassman, P.W., Abbaspour, K.C., White, M.J., Srinivasan, R., Santhi, C., Harmel, R.D., Van Griensven, A., Van Liew, M.W. and Kannan, N., 2012.

SWAT: Model use, calibration, and validation. *Transactions of the ASABE*, *55*(4), pp.1491-1508.

Banholzer, S., Kossin, J., & Donner, S. (2014). The impact of climate change on natural disasters. In *Reducing disaster: Early warning systems for climate change* (pp. 21-49). Springer, Dordrecht.

Bayar, S., Demir, I., & Engin, G. O. (2009). Modeling leaching behavior of solidified wastes using back-propagation neural networks. Ecotoxicology and environmental safety, 72(3), 843-850.

Beven, K.J. and Kirkby, M.J., 1979. A physically based, variable contributing area model of basin hydrology/Un modèle à base physique de zone d'appel variable de l'hydrologie du bassin versant. *Hydrological sciences journal*, *24*(1), pp.43-69.

Chen, Z., Lin, H. and Shen, G., 2023. TreeLSTM: A spatiotemporal machine learning model for rainfall-runoff estimation. *Journal of Hydrology: Regional Studies*, *48*, p.101474.

Cho, K., Van Merriënboer, B., Bahdanau, D. and Bengio, Y., 2014. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*.

Castangia, M., Grajales, L.M.M., Aliberti, A., Rossi, C., Macii, A., Macii, E. and Patti, E., 2023. Transformer neural networks for interpretable flood forecasting. *Environmental Modelling & Software*, *160*, p.105581.

Davenport, F. V., Burke, M., & Diffenbaugh, N. S. (2021). Contribution of historical precipitation change to US flood damages. *Proceedings of the National Academy of Sciences*, *118*(4).

Demir, I., & Beck, M. B., 2009, April. GWIS: a prototype information system for Georgia watersheds. In Georgia Water Resources Conference: Regional Water Management Opportunities, UGA, Athens, GA, US.

Demir, I., Xiang, Z., Demiray, B. and Sit, M., 2022. WaterBench-Iowa: a large-scale benchmark dataset for data-driven streamflow forecasting. *Earth system science data*, *14*(12), pp.5605-5616.

Demiray, B.Z., Sit, M. and Demir, I., 2023. EfficientTempNet: Temporal Super-Resolution of Radar Rainfall. arXiv preprint arXiv:2303.05552.

Devia, G.K., Ganasri, B.P. and Dwarakish, G.S., 2015. A review on hydrological models. *Aquatic procedia*, *4*, pp.1001-1007.

Diffenbaugh, N. S., Singh, D., Mankin, J. S., Horton, D. E., Swain, D. L., Touma, D., ... & Rajaratnam, B. (2017). Quantifying the influence of global warming on unprecedented extreme climate events. *Proceedings of the National Academy of Sciences*, *114*(19), 4881-4886.

Feng, D., Fang, K. and Shen, C., 2020. Enhancing streamflow forecast and extracting insights using long-short term memory networks with data integration at continental scales. *Water Resources Research*, *56*(9), p.e2019WR026793.

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

Granata, F., Gargano, R. and De Marinis, G., 2016. Support vector regression for rainfall-runoff modeling in urban drainage: A comparison with the EPA's storm water management model. *Water*, *8*(3), p.69.

Guo, Y., Yu, X., Xu, Y.P., Chen, H., Gu, H. and Xie, J., 2021. AI-based techniques for multi-step streamflow forecasts: application for multi-objective reservoir operation optimization and performance assessment. *Hydrology and Earth System Sciences*, *25*(11), pp.5951-5979.

Hochreiter, S. and Schmidhuber, J., 1997. Long short-term memory. *Neural computation*, *9*(8), pp.1735-1780.

Honorato, A.G.D.S.M., Silva, G.B.L.D. and Guimaraes Santos, C.A., 2018. Monthly streamflow forecasting using neuro-wavelet techniques and input analysis. *Hydrological Sciences Journal*, *63*(15-16), pp.2060-2075.

Ibrahim, K.S.M.H., Huang, Y.F., Ahmed, A.N., Koo, C.H. and El-Shafie, A., 2022. A review of the hybrid artificial intelligence and optimization modelling of hydrological streamflow forecasting. *Alexandria Engineering Journal*, *61*(1), pp.279-303.

Krajewski, W.F., Ghimire, G.R., Demir, I. and Mantilla, R., 2021. Real-time streamflow forecasting: AI vs. Hydrologic insights. *Journal of Hydrology X*, *13*, p.100110.

Krause, P., Boyle, D.P. and Bäse, F., 2005. Comparison of different efficiency criteria for hydrological model assessment. *Advances in geosciences*, *5*, pp.89-97.

Kratzert, F., Klotz, D., Brenner, C., Schulz, K. and Herrnegger, M., 2018. Rainfall–runoff modelling using long short-term memory (LSTM) networks. *Hydrology and Earth System Sciences*, *22*(11), pp.6005-6022.

Lee, T.H. and Georgakakos, K.P., 1996. Operational Rainfall Prediction on Meso-γ Scales for Hydrologic Applications. *Water Resources Research*, *32*(4), pp.987-1003.

Li, Z., & Demir, I., 2022. A comprehensive web-based system for flood inundation map generation and comparative analysis based on height above nearest drainage. Science of The Total Environment, 828, 154420.

Li, Z., & Demir, I. (2023). U-net-based semantic classification for flood extent extraction using SAR imagery and GEE platform: A case study for 2019 central US flooding. Science of The Total Environment, 869, 161757.

Lin, T., Wang, Y., Liu, X. and Qiu, X., 2022. A survey of transformers. *AI Open*.

Liu, C., Liu, D. and Mu, L., 2022. Improved transformer model for enhanced monthly streamflow predictions of the Yangtze River. *IEEE Access*, *10*, pp.58240-58253.

Mosavi, A., Ozturk, P. and Chau, K.W., 2018. Flood prediction using machine learning models: Literature review. *Water*, *10*(11), p.1536.

Munich Re. (2022). Hurricanes, cold waves, tornadoes: Weather disasters in USA dominate natural disaster losses in 2021. https://www.munichre.com/en/company/media-relations/media-information-and-corporate-news/media-information/2022/natural-disaster-losses-2021.html.

NDRCC. (2021). 2020 Global Natural Disaster Assessment Report. https://reliefweb.int/report/china/2020-global-natural-disaster-assessment-report

NOAA National Centers for Environmental Information (NCEI). (2022). US billion-dollar weather and climate disasters. https://www.ncei.noaa.gov/access/monitoring/billions/, DOI: 10.25921/stkw-7w73

Ren-Jun, Z., 1992. The Xinanjiang model applied in China. *Journal of hydrology*, *135*(1-4), pp.371-381.

Salas, J.D., Markus, M. and Tokar, A.S. (2000). Streamflow forecasting based on artificial neural networks. *Artificial neural networks in hydrology*, pp.23-51.

Sharma, P. and Machiwal, D. eds., 2021. *Advances in streamflow forecasting: from traditional to modern approaches*. Elsevier.

Sit, M., Demiray, B. and Demir, I., 2021a. Short-term hourly streamflow prediction with graph convolutional GRU networks. *arXiv preprint arXiv:2107.07039*.

Sit, M., Seo, B. C., & Demir, I., 2021b. Iowarain: A statewide rain event dataset based on weather radars and quantitative precipitation estimation. arXiv preprint arXiv:2107.03432.

Sit, M., Demiray, B.Z. and Demir, I., 2022a. A systematic review of deep learning applications in streamflow data augmentation and forecasting. EarthArxiv, 3617. https://doi.org/10.31223/X5HM08

Sit, M., Demiray, B.Z. and Demir, I., 2022b. A Systematic Review of Deep Learning Applications in Interpolation and Extrapolation of Precipitation Data. EarthArxiv, 4715. https://doi.org/10.31223/X57H2H

Sit, M., Seo, B.C., Demiray, B.Z. and Demir, I., 2023a. EfficientRainNet: Smaller Neural Networks Based on EfficientNetV2 for Rainfall Nowcasting. EarthArxiv, 5232. https://doi.org/10.31223/X5VQ1S

Sit, M., Demiray, B.Z. and Demir, I., 2023b. Spatial Downscaling of Streamflow Data with Attention Based Spatio-Temporal Graph Convolutional Networks. EarthArxiv, 5227. https://doi.org/10.31223/X5666M

Strauss, B. H., Kopp, R. E., Sweet, W. V., and Bittermann, K. (2016). Unnatural coastal floods: Sea level rise and the human fingerprint on US floods since 1950. Climate Central.

Tabari, H. (2020). Climate change impact on flood and extreme precipitation increases with water availability. *Scientific reports*, *10*(1), 1-10.

Trenberth, K. E., Cheng, L., Jacobs, P., Zhang, Y., & Fasullo, J. (2018). Hurricane Harvey links to ocean heat content and climate change adaptation. *Earth's Future*, *6*(5), 730-744.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł. and Polosukhin, I., 2017. Attention is all you need. *Advances in neural information processing systems*, *30*.

World Meteorological Organization (WMO). (2021). The Atlas of Mortality and Economic Losses from Weather, Climate and Water Extremes (1970–2019).

Wu, H., Xu, J., Wang, J. and Long, M., 2021. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *Advances in Neural Information Processing Systems*, *34*, pp.22419-22430.

Xiang, Z., Demir, I., Mantilla, R., & Krajewski, W. F., 2021. A regional semi-distributed streamflow model using deep learning. EarthArxiv, 2152. https://doi.org/10.31223/X5GW3V

Xiang, Z. and Demir, I., 2021. High-resolution rainfall-runoff modeling using graph neural network. *arXiv preprint arXiv:2110.10833*.

Xiang, Z., & Demir, I., 2022a. Real-Time streamflow forecasting framework, implementation and post-analysis using deep learning. EarthArxiv, 3162. https://doi.org/10.31223/X5BW6R

Xiang, Z. and Demir, I., 2022b. Fully distributed rainfall-runoff modeling using spatial-temporal graph neural network. EarthArxiv, 3018. https://doi.org/10.31223/X57P74

Yan, J., Jin, J., Chen, F., Yu, G., Yin, H. and Wang, W., 2018. Urban flash flood forecast using support vector machine and numerical simulation. *Journal of Hydroinformatics*, *20*(1), pp.221-231.

Yaseen, Z.M., El-Shafie, A., Jaafar, O., Afan, H.A. and Sayl, K.N. (2015). Artificial intelligence based models for stream-flow forecasting: 2000–2015. *Journal of Hydrology*, *530*, pp.829-844.

Yaseen, Z.M., Ebtehaj, I., Bonakdari, H., Deo, R.C., Mehr, A.D., Mohtar, W.H.M.W., Diop, L., El-Shafie, A. and Singh, V.P. (2017). Novel approach for streamflow forecasting using a hybrid ANFIS-FFA model. *Journal of Hydrology*, *554*, pp.263-276.

Yaseen, Z.M., Awadh, S.M., Sharafati, A. and Shahid, S. (2018). Complementary data-intelligence model for river flow simulation. *Journal of Hydrology*, *567*, pp.180-190.

Yaseen, Z.M., Sulaiman, S.O., Deo, R.C. and Chau, K.W., 2019. An enhanced extreme learning machine model for river flow forecasting: State-of-the-art, practical applications in water resource engineering area and future research direction. *Journal of Hydrology*, *569*, pp.387-408.

Yildirim, E. and Demir, I., 2022. Agricultural flood vulnerability assessment and risk quantification in Iowa. *Science of The Total Environment*, *826*, p.154165.

Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H. and Zhang, W., 2021, May. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 35, No. 12, pp. 11106-11115).

Zhou, T., Ma, Z., Wen, Q., Wang, X., Sun, L. and Jin, R., 2022, June. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *International Conference on Machine Learning* (pp. 27268-27286). PMLR.