Comparing Process-Based and Machine Learning Models for Streamflow Prediction in the Kaligandaki River Basin, Nepal

Aayush Roka[a,*], Bishesh Khanal[a]

[a] Department of Civil Engineering, Thapathali Campus, Institute of Engineering, Tribhuvan University, Kathmandu, Nepal

[*] Corresponding author. E-mail: er.rokaayush@gmail.com

This is a non-peer-reviewed preprint submitted to EarthArXiv.

Status: Not submitted to any journal at the time of submission.

June 2026

Abstract

Reliable daily streamflow prediction is critical for hydropower operations, flood risk management, and irrigation planning in monsoon-dominated Himalayan river basins. While both process-based and machine learning (ML) approaches have been used for such tasks, systematic comparisons that decompose the sources of performance differences remain scarce.This study e valuates seven configurations: a process-based SWAT model, and three XGBoost (XGB) and three Random Forest (RF) models representing pure rainfall-runoff, observed-lag, and recursive-simulation scenarios for streamflow prediction in the Kaligandaki River basin, Nepal. An NSE gap decomposition framework is applied to quantify two distinct components of discharge lag information value: the watershed memory benefit and the recursive error propagation cost. During the five-year independent test period, SWAT achieved NSE = 0.851, while pure rainfall-runoff XGB-A and RF-A models achieved NSE = 0.840. Observed-lag upper-bound models reached NSE values above 0.94. RF recursive simulation (RF-C) achieved NSE = 0.861, exceeding SWAT, whereas XGB recursive simulation showed no improvement (XGB-C = 0.838), revealing a strong algorithm-dependent sensitivity to recursive error propagation. Flow duration curve analysis reveals that SWAT underestimates low flows ($> 60\,\%$ exceedance probability) despite near-zero total PBIAS ($-0.28\,\%$), reflecting compensating biases between high- and low-flow regimes. SHAP analysis identifies the Antecedent Precipitation Index (API) as the dominant precipitation predictor, with 30-day cumulative rainfall as the second-ranked feature, confirming multi-week soil-moisture memory as the primary catchment-scale control on discharge in this large, monsoon-dominated basin. These results establish that well-engineered pure rainfall-runoff models match SWAT in aggregate NSE while substantially outperforming it in distributional flow reproduction.

Keywords: streamflow prediction; XGBoost; Random Forest; SWAT; NSE decomposition

1. Introduction

Accurate daily streamflow prediction is a fundamental requirement for various tasks in hydropower scheduling, flood early warning, and irrigation water allocation in the Hindu Kush-Himalayan (HKH) region, which provides sustainable freshwater access for more than two billion people across eight countries (Wester et al., 2019; Immerzeel et al., 2010). In the Narayani basin of central Nepal, where the 144 MW Kaligandaki A RoR hydropower project operates under discharge driven by the monsoon climate that varies by more than 100 times in dry and wet seasons (Bishwakarma et al., 2008), errors in the prediction of discharges directly cause lost generation revenue, excess turbine wear, and flood hazard.

Himalayan discharge regimes are characterised by extreme seasonal variations with the Monsoon (June–September) contributing to approximately 70–80 % of annual runoff and through the steep and non-linear rainfall-discharge relationships in which short-duration Convective events trigger rapid hydrograph responses that are difficult to anticipate (Andermann et al., 2012; Bookhagen and Burbank, 2010). This hydrological complexity makes the Narayani basin an important environment for evaluating the comparative skill of modelling methods that vary fundamentally in how they use catchment-scale physical processes.

Process-based hydrological models such as the Soil and Water Assessment Tool (SWAT) simulates discharge by solving a set of water balance equations representing rainfall , infiltration, evapotranspiration, subsurface flow, and groundwater routing (Arnold et al., 1998; Neitsch et al., 2011). SWAT's primary advantages for Himalayan applications are physical interpretability, constraint of parameters by measurable catchment properties, and mechanistic validity under climate conditions outside the training record, a critical property for climate impact assessment (Arnold et al., 2012; Nepal, 2016). Documented SWAT applications in Nepalese river basins confirm acceptable-to-very-good performance for daily discharge prediction, with Nash-Sutcliffe Efficiencies (NSE) typically in the range 0.75–0.90 (Nepal, 2016; Devkota and Bhattarai, 2021; Marahatta et al., 2021; Bajracharya et al., 2018). However, SWAT has well-known limitations in large, orographically complex Himalayan catchments: requirement of extensive datasets for parameter initialization and uncertainty, the SCS Curve Number method often underestimates

pre-monsoon runoff responses to low-intensity rainfall (Pradhananga et al., 2020); the linear reservoir groundwater routine overestimates dry-season baseflow recession (Nepal, 2016); and the temperature-index snowmelt module cannot represent dynamic glacier-area change over multi-decadal timescales (Immerzeel et al., 2010). These limitations motivate the search for complementary approaches that may better represent the full discharge distribution, particularly at the low-flow end of the flow duration curve.

Data-driven machine learning approaches have emerged as a credible alternative for streamflow prediction, taking advantage of the growing availability of long hydro-meteorological records and advances in ensemble learning algorithms (Shen, 2018; Gauch et al., 2021). Among these, XGBoost (Chen and Guestrin, 2016) has demonstrated competitive performance with other such models and deep learning models for daily discharge prediction at a lower computational cost, and has shown a robust ability in small-to-medium sample hydrological datasets characteristic of developing-country river basins (Tyralis et al., 2019; Zhang et al., 2020). The predictive skill of XGBoost for streamflow is strongly sensitive to feature engineering, which includes the construction of physically motivated input variables that acts as a proxy for various catchment processes , such as antecedent precipitation indices, multi-day rainfall accumulations, and cyclic seasonal encodings, which substitute for various physical parameters (soil moisture storage, groundwater levels) represented in process-based models (Kratzert et al., 2019; Sit et al., 2020). A pure rainfall-runoff XGBoost model, relying solely on precipitation and temporal features constitute an honest operational benchmark whose predictive skill is bounded by the information content of precipitation alone, without access to any observed discharge information at prediction time (Shortridge et al., 2016).

A growing body of comparative literature has evaluated process-based models against machine learning approaches for streamflow prediction, but findings are highly context-dependent, and no consensus has emerged (Rahimzad et al., 2021; Zhang et al., 2020; Humphrey et al., 2016). Shortridge et al. (2016) found that random forests matched or exceeded SWAT in data-rich US catchments during dry years, but underperformed during high-flow events. Gauch et al. (2021) demonstrated that LSTM networks generalize better to ungauged conditions than process-based models across large continental datasets. In Nepal, Bhandari et al. (2025), did a comprehensive

review of all the studies done through machine learning where he reviewed 34 papers from 2010-2025. Findings suggested that ML models outperformed traditional models in streamflow predictions. Most published ML applications in Nepal and the broader Himalayan region have applied either standalone LSTM or ANN models without benchmarking against a process-based model calibrated on identical data and have rarely such transparency through inclusion of partial dependence plots, accumulated local effects, and Shapley values. No study has systematically decomposed the sources of performance differences between a calibrated SWAT model and XGBoost in a large, monsoon-dominated Himalayan catchment using a consistent multi-metric evaluation across training, validation, and test periods.

A critical but often unacknowledged methodological issue in ML streamflow prediction is the usage of lagged discharges as input features. When observed discharge from previous time steps are used as input at test time, the model exploits information that will not be available during operational forecasting scenarios where future observed discharge cannot be known (Kratzert et al., 2019; Sit et al., 2020). This inflates reported NSE values relative to an honest rainfall-runoff baseline, because the discharge autocorrelation is very high in large catchments and the lag features effectively transfer the observed data forward by one time step. Alternatively, when discharge lags are replaced by recursively generated model predictions (using the model's own prior outputs as the next input), prediction errors accumulate across time steps, potentially degrading performance below the pure rainfall-runoff baseline during multi-day monsoon events (Humphrey et al., 2016). Distinguishing these two effects, the watershed memory benefit of discharge information and the recursive error propagation cost of replacing it with model predictions is a key analytical contribution of this study.

This study therefore pursues three objectives in the Kaligandaki basin, Nepal, a major tributary of the Narayani River. First, to evaluate and compare SWAT, XGBoost (XGB-A/B/C), and Random Forest (RF-A/B/C) for daily streamflow prediction over identical training, validation, and test periods using NSE, KGE, R2, RMSE, MAE, PBIAS, and visual diagnostics. Second, to analyze the performance gap between the models that can be attributed to discharge lag information into the following components, the benefit due to memory of watershed through observed lags and the cost of recursive error propagation due to the use of the model's own predicted discharge,

providing a basis for operational model selection decisions. Third, to characterize the full discharge distribution reproduction of all models through flow duration curves, and to interpret the importance of hydro-meteorological features within a physical framework relevant to the Monsoon-dominated hydrology of the Narayani system.

2. Study Area and Data

2.1. Basin Description

The Kaligandaki River drains approximately 11,812 km^2 of the central Himalayan region of Nepal and lies between 27°30′–29°00′N and 83°00′–84°00′E (Figure 1). It contains one of the seven major tributaries of the Narayani (Gandaki) River, which discharges into the Ganga River of India (DHM, 2017). The river originates near Damodar Kunda (≈ 5,000 m a.s.l.) in the Mustang district and flows in the south direction for approximately 300 km before entering the Terai region. The basin has a total elevation range from below 200 m in the southern lowlands to 8,167 m at Dhaulagiri I (Shrestha et al., 1999), producing one of the most significant relief and precipitation gradients of any catchment globally.

The hydroclimatology is governed by the Monsoon, which delivers approximately 80 % of annual precipitation between June and September (Andermann et al., 2012). The mean annual precipitation in the ranges from below 270 mm yr^{-1} in the trans-Himalayan rain shadow of Mustang to over 3,500 mm yr^{-1} on the orographically enhanced windward slopes of the Annapurna massif (Bookhagen and Burbank, 2010). This steep gradient produces a highly non-uniform spatial distribution of runoff generation, with mid-hill sub-catchments generating substantially higher specific discharge than either the arid upper basin or the Siwalik foothills. Mean annual discharge at the Kotagaun outlet, which serves as the streamflow calibration and prediction target in this study, is approximately 432 m^3 s^{-1} over the baseline period (2004–2019), with monsoon-season peaks exceeding 4,900 m^3 s^{-1} during extreme events (DHM, 2017).

The basin hosts the 144 MW Kaligandaki A run-of-river hydropower project at Mirmi (Syangja district), which has been in commercial operation since 2002 (Bishwakarma et al., 2008). The

facility's generation schedule is directly determined by daily discharge at the headworks intake, making accurate daily streamflow prediction a high-value operational requirement for the facility operator (Nepal Electricity Authority).

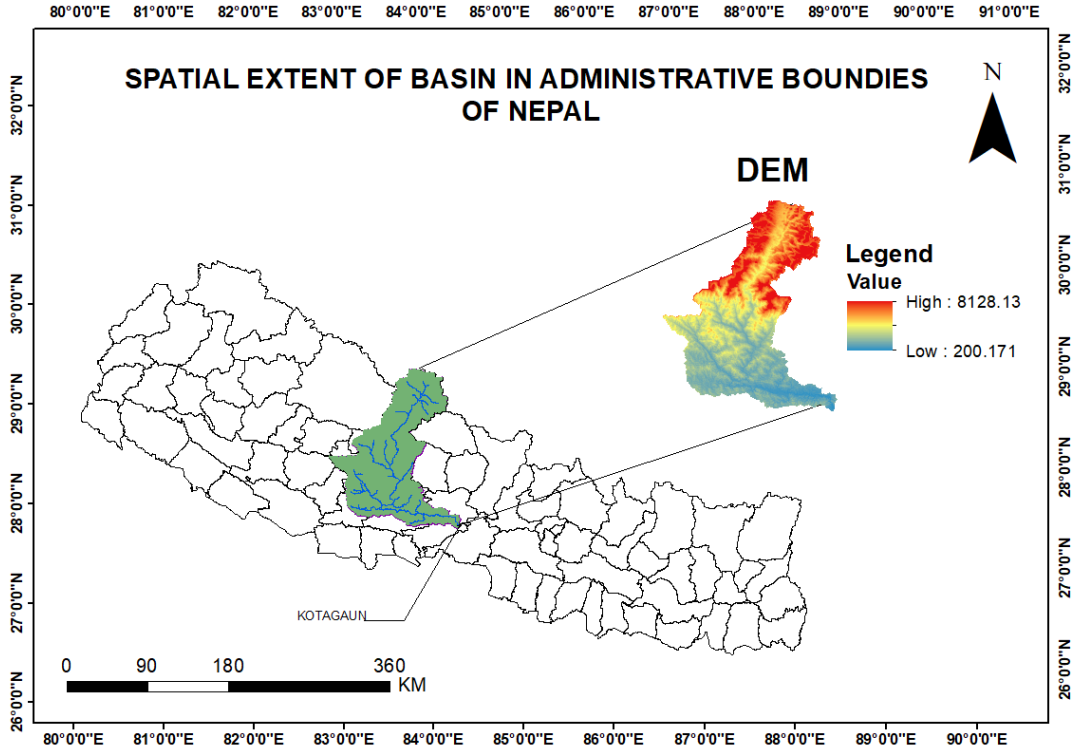

Figure 1: Location and physiographic extent of the Kaligandaki River basin, Nepal (outlet at Kotagaun gauging station, DHM Station 420, drainage area $\approx 11{,}812\,\mathrm{km}^2$); relief is shown by the ALOS PALSAR 30 m DEM.

2.2. Data Sources and Periods

Daily discharge records at the Kotagaun gauging station (DHM Station 420; drainage area $\approx 11{,}812\,\mathrm{km}^2$) were obtained from the Department of Hydrology and Meteorology (DHM), Nepal, for the period 1980–2019. Precipitation forcing was provided by 38 DHM stations spanning the full elevational gradient of the basin; records from these stations were used from 2000. Temperature data were available from 7 stations (1980–2019).

The dataset was split into three non-overlapping chronologic periods for various analysis (Table 1). The ML model was trained for the same timeframe as swat model, starting from 2004. The validation period (2014; 1 year) coincides with the final year of the SWAT calibration window, enabling direct comparison of a freshly trained ML model against a SWAT model that has

seen the same period during calibration. The test period (2015–2019; 5 years) provides the principal out-of-sample evaluation; it coincides with the SWAT validation phase and includes a high-discharge monsoon of 2015 (peak observed discharge $\approx 4{,}900\,\mathrm{m^3\,s^{-1}}$) alongside four monsoon seasons of more typical magnitude.

Table 1: Chronological data split used for ML and SWAT model evaluation.

Period	Dates	Length	Role
ML Training	2004-01-01 – 2013-12-31	10 years	Fit XGBoost and RF parameters
ML Validation	2014-01-01 – 2014-12-31	1 year	Hyperparameter selection
ML Test	2015-01-01 – 2019-12-31	5 years	Independent performance evaluation
SWAT Warm up period	2000-01-01 – 2003-12-31	4 years	Allows model stablization
SWAT Calibration	2004-01-01 – 2014-12-31	11 years	SUFI-2 parameter optimisation
SWAT Validation	2015-01-01 – 2019-12-31	5 years	Out-of-sample evaluation

3. Methodology

3.1. Overall Framework

Four streamflow prediction model types (seven total configurations) were evaluated over identical training (2004–2013) , validation (2014), and test (2015–2019) periods at the Kotagaun outlet of the Kaligandaki basin (Figure 2). The framework is designed around a NSE decomposition strategy that separates the contribution of discharge lag information into two quantifiable components: the watershed memory benefit and the recursive error propagation cost. Table 2 summarises the key structural characteristics of all model configurations.

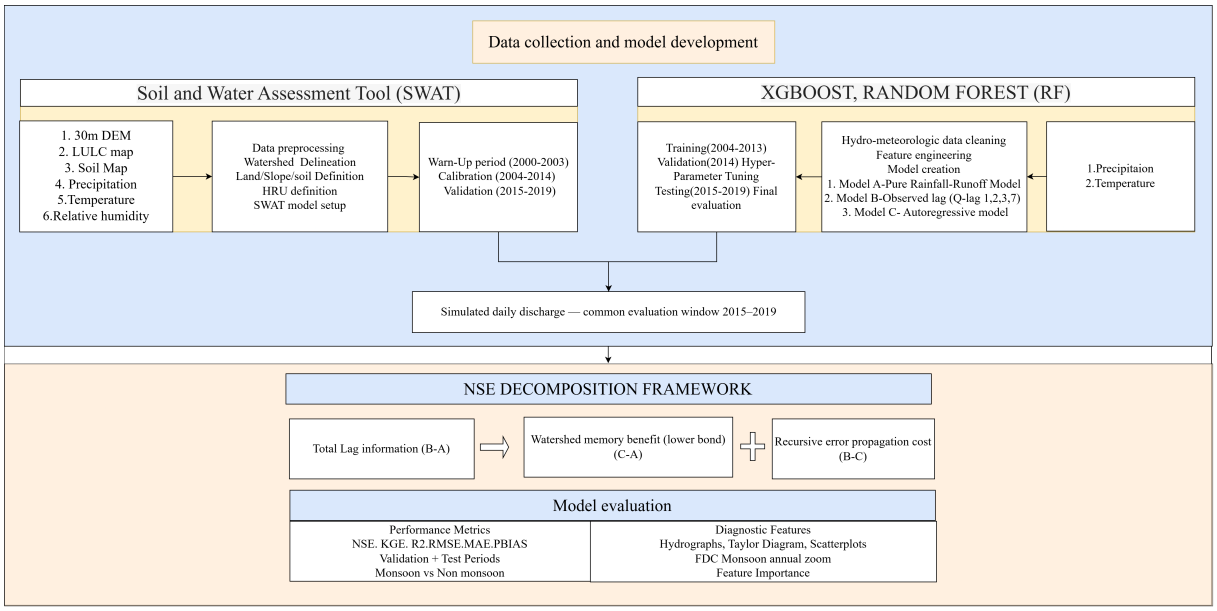

Figure 2: Overall methodological framework.

Table 2: Summary of the seven streamflow prediction model configurations evaluated in this study.

Configuration	Algorithm / Type	Discharge lags	Role
SWAT	Process-based (Arc-SWAT 2012)	N/A (physical state variables)	Physical benchmark
XGB-A / RF-A	XGBoost / RF, Pure RR	None	Honest operational baseline
XGB-B / RF-B	XGBoost / RF, Observed Lags	From observed record	Upper performance bound
XGB-C / RF-C	XGBoost / RF, Recursive Sim	From model's own prior outputs	Operationally deployable

3.2. SWAT Process-Based Benchmark

The SWAT model used in this study was calibrated using ArcSWAT 2012 within ArcMap 10.5. The basin was discretised into 80 sub-basins and 1,308 Hydrological Response Units (HRUs) using a 30 m ALOS PALSAR DEM with a contributing-area threshold of 6,000 ha. Surface

runoff was estimated using the SCS Curve Number method (USDA, 1986); evapotranspiration was computed using the Penman-Monteith approach (Allen et al., 1998); and snow accumulation and melt were represented using a temperature-index model across five elevation bands at 500 m intervals (Fontaine et al., 2002; Neitsch et al., 2011). Streamflow calibration and validation used the SUFI-2 algorithm in SWAT-CUP (Abbaspour, 2015) with NSE as the objective function. Monthly NSE achieved 0.96 (calibration, 2004–2014) and 0.91 (validation, 2015–2019) at Kotagaun; daily NSE was 0.91 during the calibration period, confirming that the monthly calibration objective did not substantially compromise daily hydrograph representation (Arnold et al., 2012). Key calibrated parameters include: CN2 (-0.12 relative change), ALPHA_BF ($0.048\,\mathrm{d}^{-1}$), GW_DELAY (31 d), GWQMN (1,000 mm), and ESCO (0.95); SUFI-2 p-factor and r-factor were 0.79 and 1.11, respectively, indicating acceptable calibration uncertainty. The SWAT simulated daily discharge time series over the 2004–2019 period was used as the process-based benchmark in all comparative analyses.

3.3. Feature Engineering

All six ML model configurations share an identical base feature set constructed from daily precipitation and temperature records at the 38 and 7 DHM stations, respectively. Missing values in station records (<3 % of station-days) were filled using inverse-distance-weighted interpolation from neighbouring stations of the same elevation tier; missing discharge values were left as missing rather than being filled. Features were engineered to encode the principal hydrological memory mechanisms governing discharge generation in the Kaligandaki basin.

3.3.1. Precipitation Features

Station-level precipitation records were used as individual predictors alongside basin-aggregated statistics: mean areal precipitation (avg_precip), daily maximum precipitation (precip_max), and inter-station precipitation variability (precip_std) as a proxy for storm spatial organization. Multi-day cumulative rainfall windows of 3, 7, 14, and 30 days (precip_3day through precip_30day) encode subsurface drainage and soil moisture filling processes at progressively longer timescales.

The Antecedent Precipitation Index (API) was computed as an exponentially weighted moving average of daily mean precipitation with a decay factor $k = 0.85$ (Equation 1):

$$\mathrm{API}_t = k \cdot \mathrm{API}_{t-1} + (1-k) \cdot P_t \tag{1}$$

where P_t is mean areal precipitation on day t. The decay factor $k = 0.85$ corresponds to an effective memory half-life of approximately 4.3 days, consistent with sub-weekly precipitation response timescales documented in the mid-hill sub-catchments that dominate Kaligandaki runoff generation (Heggen, 2001). Sensitivity tests across $k \in \{0.70, 0.80, 0.85, 0.90, 0.95\}$ confirmed that XGB-A test NSE varied by ± 0.004 around the $k = 0.85$ optimum, and the rank ordering of API versus precip_30day in feature importance was robust across all tested values. A binary heavy-rain indicator (heavy_rain) was defined as days exceeding the 90th-percentile precipitation in the training period to flag extreme monsoon events.

3.3.2. Temperature Features

Station-level daily temperature records were included as individual predictors alongside basin-mean, maximum, and standard deviation of temperature, and 3-day, 7-day, 14-day, and 30-day rolling mean temperatures. Temperature features encode the phase partitioning of precipitation (rain vs. snow), the timing and intensity of snowmelt-driven pre-monsoon runoff from upper-basin elevation bands, and the influence of warming on evapotranspiration losses during the growing season (Bookhagen and Burbank, 2010). An elevation-weighted basin-mean temperature was computed using station elevations as weights, providing a lapse-rate-adjusted proxy for snowmelt potential.

3.3.3. Temporal and Seasonal Encodings

Calendar month and day-of-year were encoded using sine-cosine transformations (month_sin, month_cos, doy_sin, doy_cos) to preserve their circular periodicity in a form compatible with tree-ensemble models (Sit et al., 2020). Binary indicators for the monsoon season (season:

June–September), early monsoon (early_monsoon: June–July), and late monsoon (late_monsoon: August–September) were included to enable the model to apply distinct rainfall-discharge relationships during the soil saturation progression of the monsoon season.

3.4. XGBoost and Random Forest Model Configurations

All ML models were implemented in Python using XGBRegressor (XGBoost v1.7; Chen and Guestrin 2016) and RandomForestRegressor (scikit-learn v1.3; Pedregosa et al. 2011). Hyperparameters were optimized on the validation period (2014) using randomized search with 30 iterations, employing expanding-window time-series cross-validation with five folds that strictly respect temporal ordering (no future data leakage). The search space covered: *n_estimators* $\in \{300, 500, 700\}$; *max_depth* $\in \{3, 4, 5, 6\}$; *learning_rate* (XGB) $\in \{0.01, 0.03, 0.05, 0.08\}$; *subsample* $\in \{0.7, 0.8, 0.9, 1.0\}$. Identical search spaces were applied for both XGBoost and Random Forest to ensure performance differences reflect information content rather than tuning artefacts. The target variable was daily discharge at Kotagaun ($m^3\ s^{-1}$); all predicted values were clipped to zero to prevent negative discharge.

3.4.1. Model A: Pure Rainfall-Runoff

Model A (XGB-A and RF-A) uses only the base hydro-meteorological feature set described in Section 3.3. No discharge information is used at any stage, constituting the honest operational benchmark for rainfall-runoff prediction (Shortridge et al., 2016; Frame et al., 2022).

3.4.2. Model B: Observed Lag Model (Upper Bound)

Model B (XGB-B and RF-B) augments the base feature set with four discharge lag features: Q_lag1, Q_lag2, Q_lag3, and Q_lag7 (observed daily discharge at 1, 2, 3, and 7 days prior to prediction). At test-time inference, these lags are drawn from the *observed* discharge record, making Model B an idealised upper bound rather than an operational prediction system (Frame et al., 2022).

3.4.3. *Model C: Recursive Simulation*

Model C (XGB-C and RF-C) is trained identically to Model B (with observed lag features in the training set), but at test-time inference replaces the observed lag values with the model's own prior predictions: Q_lag1 $\leftarrow \hat{Q}_{t-1}$, Q_lag2 $\leftarrow \hat{Q}_{t-2}$, and so on. As Model C is trained with observed lags but Model A is not, the C − A NSE difference shows two effects simultaneously: (1) the influence of training with lag features and (2) the effect of using recursive estimates at test time. Within this study, the C − A gap is best interpreted as a lower bound on the potential benefit of test-time discharge information, rather than a clean causal estimate.

3.5. NSE Gap Decomposition Framework

The total performance gap attributable to discharge lag information is decomposed into two additive components using the test-period NSE values of Models A, B, and C:

$$\underbrace{\mathrm{NSE_B} - \mathrm{NSE_A}}_{\text{Total Q-lag information}} = \underbrace{\mathrm{NSE_C} - \mathrm{NSE_A}}_{\text{Watershed memory benefit (lower bound)}} + \underbrace{\mathrm{NSE_B} - \mathrm{NSE_C}}_{\text{Recursive error propagation cost}} \quad (2)$$

The *watershed memory benefit* (C − A) quantifies the NSE change when moving from a pure rainfall-runoff model to one that uses recursively simulated lag inputs. The *recursive error propagation cost* (B − C) quantifies the NSE degradation observed by replacing known observed discharge lags with estimated lags during multi-step recursive inference.

3.6. SHAP Feature Attribution

To provide stable feature importance estimates that are robust to feature correlation (unlike gain-based importance, which is known to be unstable when predictors share predictive information; Lundberg et al. 2020), SHAP (SHapley Additive exPlanations; Lundberg and Lee 2017) values were computed for XGB-B using TreeExplainer (Lundberg et al., 2020). Mean absolute SHAP values over the test period were used to rank the top 20 features; a beeswarm plot of the top 10

features illustrates the directional relationship between feature value and predicted discharge. Gain-based importance is retained for Model A (where feature correlation is lower and results are presented alongside SHAP for cross-validation).

3.7. Performance Metrics

Model performance was evaluated using six complementary statistics: Nash-Sutcliffe Efficiency (NSE; Nash and Sutcliffe 1970), Kling-Gupta Efficiency (KGE; Gupta et al. 2009), coefficient of determination (R^2), root mean square error (RMSE; $m^3\,s^{-1}$), mean absolute error (MAE; $m^3\,s^{-1}$), and percent bias (PBIAS %). PBIAS is computed as:

$$\text{PBIAS} = \frac{\sum(Q_{\text{obs}} - Q_{\text{sim}})}{\sum Q_{\text{obs}}} \times 100 \tag{3}$$

Under this convention, PBIAS > 0 indicates systematic underestimation (simulated flow below observed); PBIAS < 0 indicates systematic overestimation. NSE was used as the primary metric for aggregate performance and for the NSE gap decomposition (Equation 2). KGE was used to separately assess correlation, variability bias, and mean bias components of model error (Gupta et al., 2009). Flow duration curves were constructed from the ranked test-period discharge records to characterise full-distribution reproduction (Yadav et al., 2007). Taylor diagrams simultaneously visualised correlation, standard deviation ratio, and centred RMSE (Taylor, 2001). Performance was reported separately for the validation period (hyperparameter selection) and the five-year test period (final independent evaluation).

4. Results and Discussion

4.1. Validation Period Performance (2014)

All model configurations achieved good-to-excellent performance during the one-year validation period (Table 3). SWAT performed at NSE = 0.911 during its calibration phase (2014 falls within the SWAT calibration window 2004–2014), reflecting the physically calibrated water

balance. Among the XGBoost configurations, XGB-A achieved NSE = 0.938 and KGE = 0.967, XGB-C reached NSE = 0.940 and KGE = 0.958, and XGB-B achieved the highest performance at NSE = 0.947 and KGE = 0.928. Random Forest configurations showed slightly lower NSE than their XGBoost counterparts (RF-A: 0.922; RF-B: 0.927), with the exception of RF-C (NSE = 0.893), which underperformed during this period. SWAT's PBIAS of $-9.57\,\%$ during validation indicates systematic overestimation of annual water yield (Equation 3: PBIAS < 0 denotes overestimation), likely attributable to generous groundwater return flow parameterization during the SUFI-2 calibration optimization (Nepal, 2016). All XGBoost models maintained near-zero PBIAS (-0.53 to $+1.17$ confirming well-calibrated expectation values over the validation period.

Table 3: Model performance statistics for the validation period (2014). PBIAS < 0 indicates systematic overestimation; PBIAS > 0 indicates systematic underestimation.

Model	NSE	KGE	R^2	RMSE ($m^3 s^{-1}$)	MAE ($m^3 s^{-1}$)	PBIAS (%)
SWAT (calib phase)	0.911	0.896	0.911	162.95	92.57	−9.57
XGB-A (Pure RR)	0.938	0.967	0.938	136.05	68.73	+1.02
XGB-C (Recursive)	0.940	0.958	0.940	134.04	67.74	−0.53
XGB-B (Obs Lags)	0.947	0.928	0.947	124.97	49.33	+1.17
RF-A (Pure RR)	0.922	0.944	0.922	152.09	75.70	+2.67
RF-C (Recursive)	0.893	0.921	0.893	178.21	92.13	−5.43
RF-B (Obs Lags)	0.927	0.922	0.927	147.28	55.84	+0.89

Performance thresholds: Very Good NSE > 0.75; Good 0.65–0.75 (Moriasi et al., 2007).

4.2. Test Period Performance — Aggregate Metrics (2015–2019)

Table 4 presents the aggregate performance statistics for the five-year independent test period. SWAT achieved NSE = 0.851 and KGE = 0.849, establishing the process-based benchmark with near-zero aggregate PBIAS (-0.28 indicating nearly unbiased overall water yield simulation during its validation phase. XGB-A (Pure Rainfall-Runoff) achieved NSE = 0.840, marginally below SWAT by 0.011 NSE units, with RMSE = $187.4\,m^3\,s^{-1}$ and substantially lower MAE (84.3 vs. $103.0\,m^3\,s^{-1}$), indicating that the data-driven model makes smaller per-day errors even when aggregate NSE is slightly lower. XGB-A's PBIAS of $+10.9\,\%$ reflects systematic

underestimation during the highest-discharge monsoon events, a well-documented extrapolation constraint of tree-ensemble models near the upper tail of the training discharge distribution (Chen and Guestrin, 2016; Tyralis et al., 2019). XGB-C (Recursive Simulation) achieved NSE = 0.838, marginally *below* XGB-A (difference = −0.002), indicating that for XGBoost in this basin, recursive error propagation eliminates any potential watershed memory benefit. XGB-B (Observed Lags) achieved substantially higher NSE = 0.951 and KGE = 0.912, establishing the upper performance bound at RMSE = 104.0 $m^3\ s^{-1}$.

Random Forest results provide an important contrast. RF-A achieved NSE = 0.840, essentially identical to XGB-A. RF-C (Recursive) reached NSE = 0.861, which exceeds both SWAT (0.851) and RF-A (0.840), demonstrating that Random Forest handles recursive lag uncertainty more robustly than XGBoost. RF-C also achieved the best KGE of all non-lag models (0.864), suggesting superior reproduction of discharge variability. RF-B achieved NSE = 0.946 and KGE = 0.929.

The most notable finding is that SWAT's KGE (0.849) substantially exceeds XGB-A's KGE (0.807) despite near-equal NSE, because KGE simultaneously penalises correlation error, variability mismatch, and bias (Gupta et al., 2009). SWAT's higher KGE reflects better variability reproduction (standard deviation ratio closer to unity), while XGB-A's positive PBIAS (+10.9 %) degrades its KGE relative to NSE.

Table 4: Model performance statistics for the independent test period (2015–2019). Bold values denote best-performing non-upper-bound configuration per metric. PBIAS < 0 = overestimation; PBIAS > 0 = underestimation.

Model	**NSE**	**KGE**	**R^2**	**RMSE** (m^3s^{-1})	**MAE** (m^3s^{-1})	**PBIAS** (%)
SWAT (valid phase)	0.851	**0.849**	0.851	**181.0**	103.0	−0.28
XGB-A (Pure RR)	0.840	0.807	0.840	187.4	84.3	+10.9
XGB-C (Recursive)	0.838	0.797	0.838	188.4	90.5	+11.4
XGB-B (Obs Lags)	0.951	0.912	0.951	104.0	43.2	+3.91
RF-A (Pure RR)	0.840	0.782	0.840	187.4	85.1	+11.2
RF-C (Recursive)	**0.861**	**0.864**	**0.861**	174.9	**80.8**	+5.59
RF-B (Obs Lags)	0.946	0.929	0.946	109.1	41.9	+2.09

4.3. NSE Gap Decomposition

The NSE gap decomposition (Figure 3; Equation 2) separates the information value of discharge lags into two quantifiable components. Results differ markedly between XGBoost and Random Forest, revealing an algorithm-dependent sensitivity to recursive error propagation.

XGBoost Decomposition

The XGBoost *watershed memory benefit* (XGB-C − XGB-A) equals $0.8384 - 0.8403 = -0.0019$ NSE units, a marginally *negative* value. This indicates that for XGBoost in this basin, the recursive use of model-predicted discharge lags provides no net improvement over pure rainfall-runoff prediction; the degradation introduced by imperfect recursive lag estimates fully offsets any information gain from discharge memory (subject to the training asymmetry noted in Section 3.4.3). The XGBoost *recursive error propagation cost* (XGB-B − XGB-C) equals $0.9508 - 0.8384 = +0.1124$ NSE units. This large gap quantifies the penalty from replacing known observed lags with model-generated estimates across the five-year test period. The total information value of observed discharge lags for XGBoost is XGB-B − XGB-A $= +0.1105$ NSE units, of which more than 100 % is consumed by recursive error propagation, leaving a net realised benefit of -0.0019 NSE units for Model C over Model A.

Random Forest Decomposition

The Random Forest decomposition reveals a qualitatively different behaviour. The RF *watershed memory benefit* (RF-C − RF-A) equals $0.8608 - 0.8402 = +0.0206$ NSE units, positive and more than ten times larger than the XGBoost equivalent. The RF *recursive error propagation cost* (RF-B − RF-C) equals $0.9458 - 0.8608 = +0.0850$ NSE units, substantially smaller than the XGBoost cost (0.1124). RF-C therefore achieves a meaningful operational benefit over RF-A and even surpasses SWAT (NSE = 0.851) in aggregate NSE and KGE, indicating that Random Forest's ensemble averaging provides a stabilizing effect on recursively propagated prediction errors that XGBoost lacks.

Physical Interpretation

The Kaligandaki basin's multi-day response time at the Kotagaun outlet (groundwater drainage, channel storage, and hillslope return flow processes) means that yesterday's discharge carries genuine information about today's discharge beyond what is captured by precipitation accumulation features alone. That XGBoost fails to realize this memory benefit operationally (XGB-C $<$ XGB-A) while Random Forest does (RF-C $>$ RF-A) has a plausible algorithmic explanation: XGBoost's greedy sequential tree building makes it more sensitive to the injection of erroneous lag values (which are highly informative but also highly influential when wrong), whereas Random Forest's bootstrap aggregation averages out this error injection across 500 trees. This distinction has a direct operational implication: for discharge- lag-augmented prediction in large Himalayan basins, Random Forest is the preferred algorithm for recursive simulation, and XGBoost with pure rainfall-runoff features (XGB-A) is preferred over XGBoost with recursive lags (XGB-C).

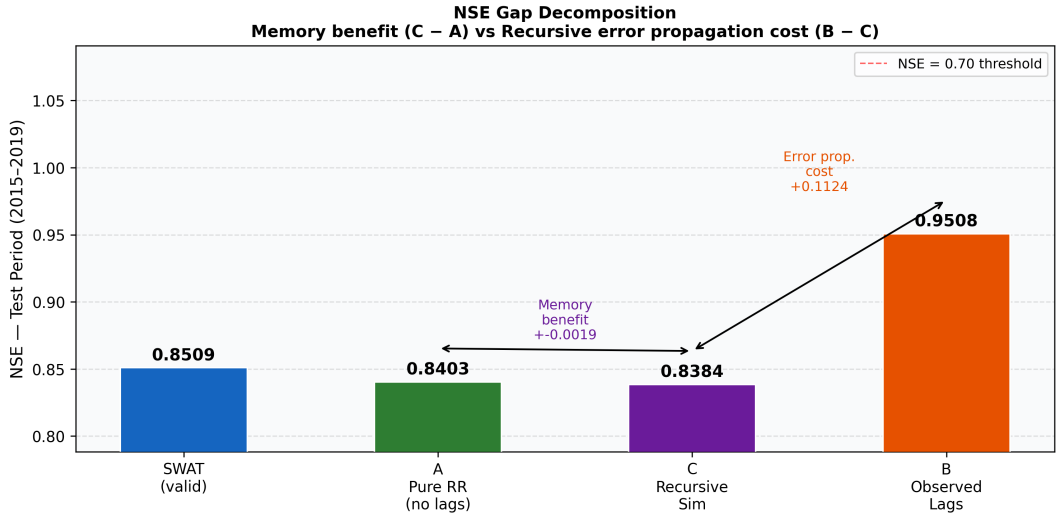

Figure 3: NSE gap decomposition across the XGBoost model configurations (test period, 2015–2019). The memory benefit (XGB-C − XGB-A = −0.0019, slightly negative for XGBoost) and the recursive error propagation cost (XGB-B − XGB-C = +0.1124) are annotated.

4.4. Temporal, Distributional, and Event-Scale Diagnostics

Aggregate metrics alone cannot fully characterize model behaviour because they weight high-flow monsoon events disproportionately and hides distributional biases that cancel in the mean.

The full-period hydrograph, monsoon-season zoom, flow duration curve, and scatter-Taylor diagnostics together expose model-specific failure modes that are operationally important but statistically invisible at the NSE level.

Full-period hydrograph and seasonal bias structure. All seven configurations correctly reproduce the timing of monsoon onset and recession in each test year and capture inter-annual discharge variability, including the anomalously large 2015 monsoon peak ($\approx 4{,}900\,\mathrm{m^3\,s^{-1}}$; Figure 4, Panel 1). Systematic differences emerge at the seasonal scale. SWAT's near-zero aggregate PBIAS ($-0.28\,\%$) masks competing distributional biases: a marginal tendency to overestimate high-flow monsoon peaks which is consistent with SCS-CN calibration optimized for peak discharge (USDA, 1986). XGB-A and RF-A closely track observed discharge across the full year but structurally underestimate the largest monsoon events (PBIAS $\approx +11\,\%$), reflecting the inability of tree-ensemble models to extrapolate beyond the maximum discharge in the training record (Chen and Guestrin, 2016). Panel 2 isolates the XGB-B vs. XGB-C contrast: the B $-$ C divergence is concentrated in the monsoon season and is year-dependent, with 2015 and 2017 showing narrow gaps while 2018 produces a large sustained divergence persisting from early August through late September. Panel 3 confirms the near-zero (slightly negative) XGB C $-$ A gap across all test years, consistent with the -0.0019 NSE watershed memory effect reported in Section 4.3.

Event-scale error propagation: monsoon zoom. Annual monsoon zoom plots (May–October; Figure 5) resolve the year-by-year structure of the B $-$ C divergence. In 2015 and 2017, monsoon events display clear peak-and-recession structure; XGB-C tracks XGB-B closely, with the shaded gap remaining narrow throughout. The 2018 monsoon is diagnostically the most informative: a prolonged August–September elevated-discharge period without a distinct inter-event recession caused XGB-C's recursive errors to compound progressively, generating underestimation of 500–800 $\mathrm{m^3\,s^{-1}}$ sustained for approximately six weeks. This confirms that the magnitude of recursive error propagation in XGB-C scales with the *duration* of elevated discharge, not its peak magnitude. RF-C behaved qualitatively differently: it maintained a positive skill advantage

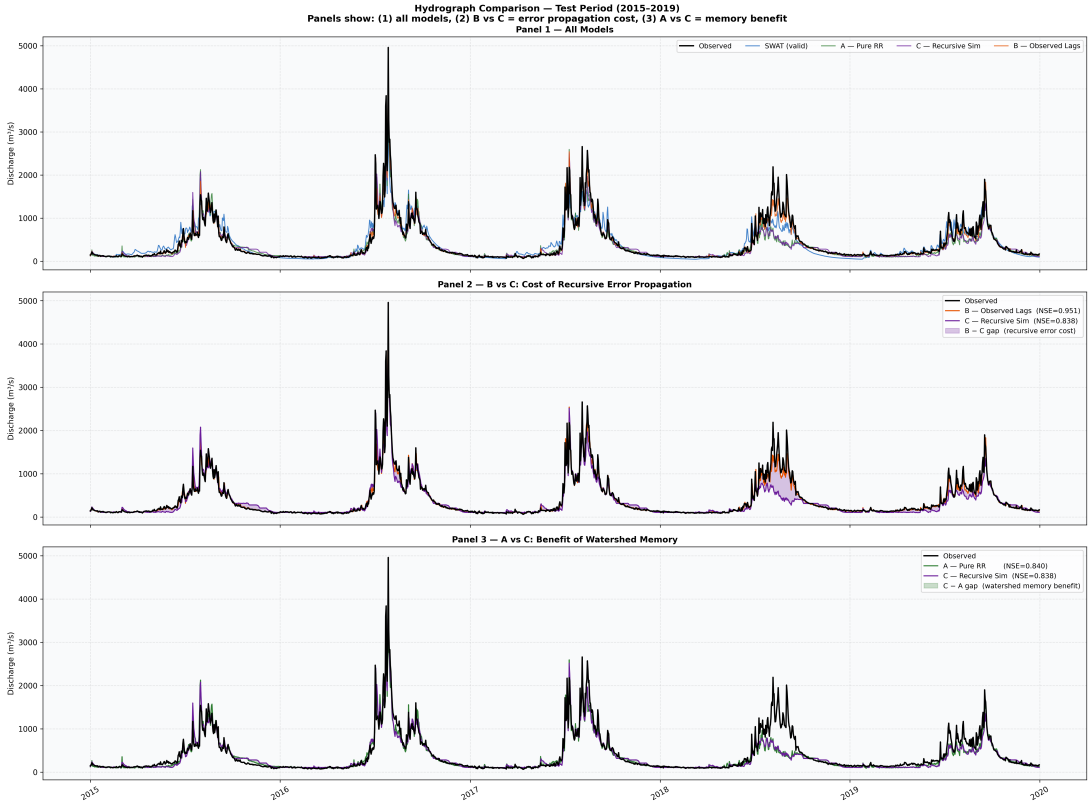

Figure 4: Three-panel hydrograph comparison for the test period (2015–2019).

over RF-A across all five monsoon seasons, including 2018, confirming that Random Forest's bootstrap aggregation stabilises recursive prediction in a way that XGBoost's sequential tree building does not.

Flow duration curve: the critical distributional finding. The FDC (Figure 6) provides the most operationally significant result of all findings, as it carries greater significance understanding the distributional characteristics of streamflow. All ML configurations closely track the observed distribution across the full exceedance probability range (0–100 %), with only marginal underestimation at the highest peak flows (< 5 % exceedance). SWAT, by contrast, diverges sharply above approximately 60 % exceedance probability, underestimating low-to-medium flows by 40–50 % at the 100th percentile (simulated minimum $\approx 35\,\mathrm{m^3\,s^{-1}}$ versus observed $\approx 70\,\mathrm{m^3\,s^{-1}}$). This deficit arises from the groundwater parameterization, in which the calibrated linear reservoir sustains baseflow during the monsoon recession but inadequately represents the persistent minimum flow threshold maintained by deep groundwater drainage into the dry season. Critically, this distributional bias is not detectable from aggregate NSE, KGE, or PBIAS, because

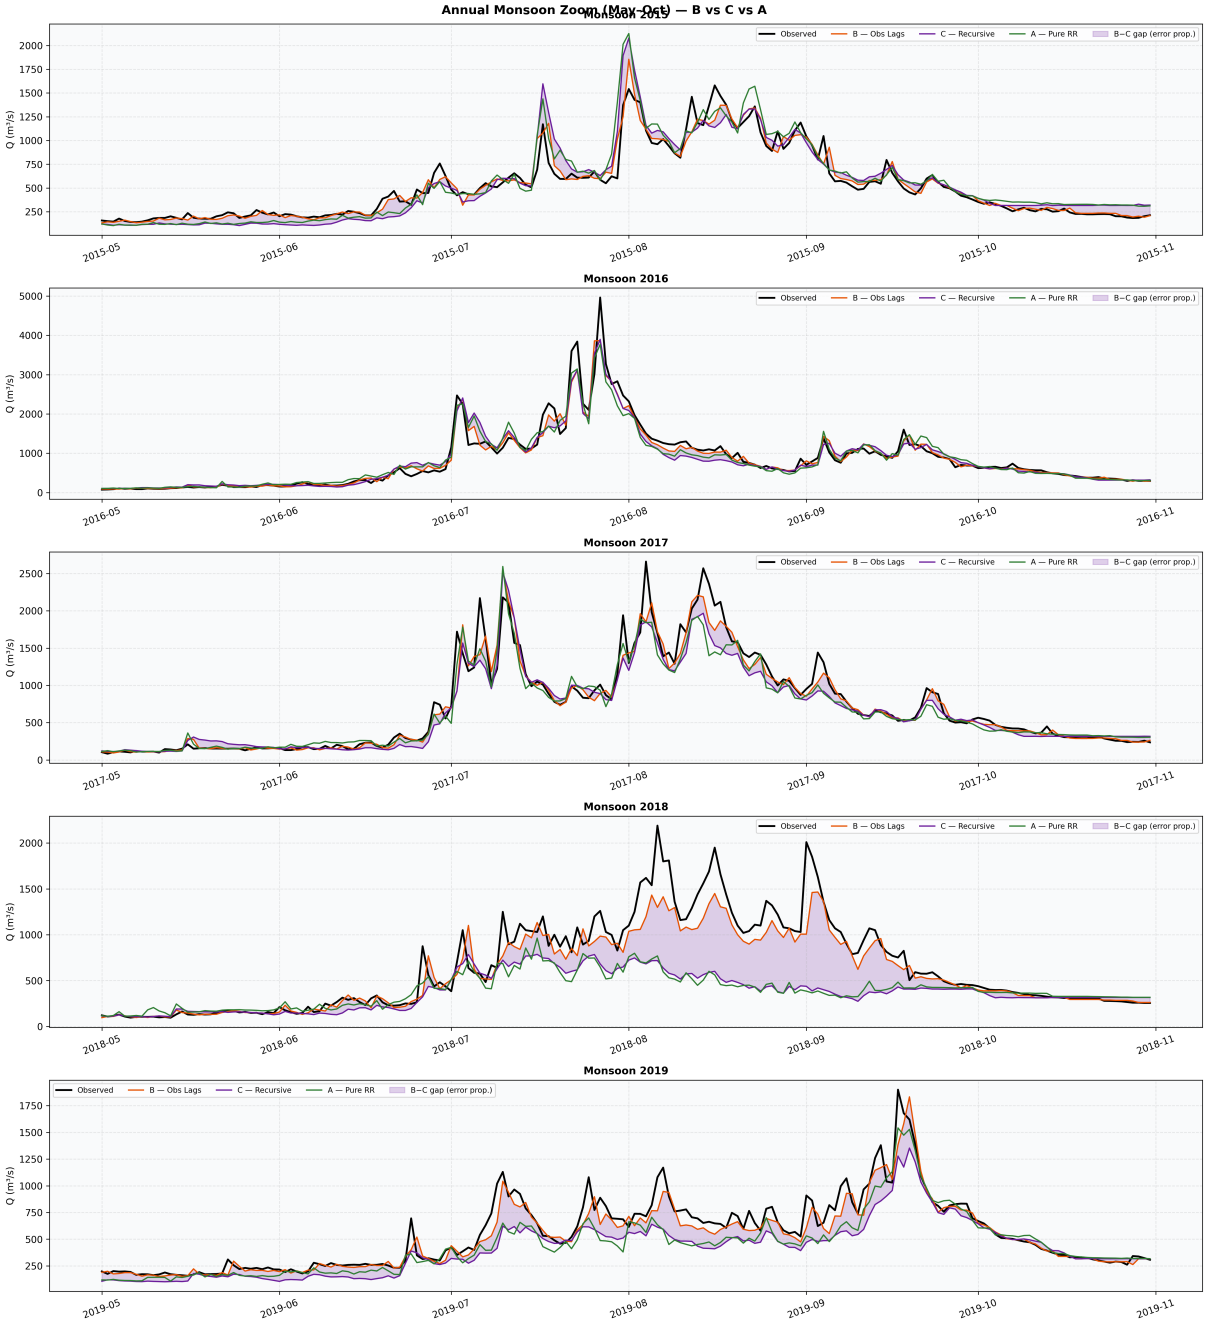

Figure 5: Annual monsoon-season zoom (May–October, 2015–2019). The shaded B − C gap quantifies the XGBoost recursive error propagation cost year-by-year. RF-C maintained a positive skill advantage over RF-A across all seasons.

SWAT's slight overestimation of high-flow discharge compensates for the low-flow deficit to produce near-zero mean bias. For applications where the lower end of the discharge distribution governs the decision such as hydropower firm capacity scheduling, environmental minimum flow compliance, and dry-season irrigation allocation, here any ML configuration is preferable to SWAT regardless of comparable aggregate NSE.

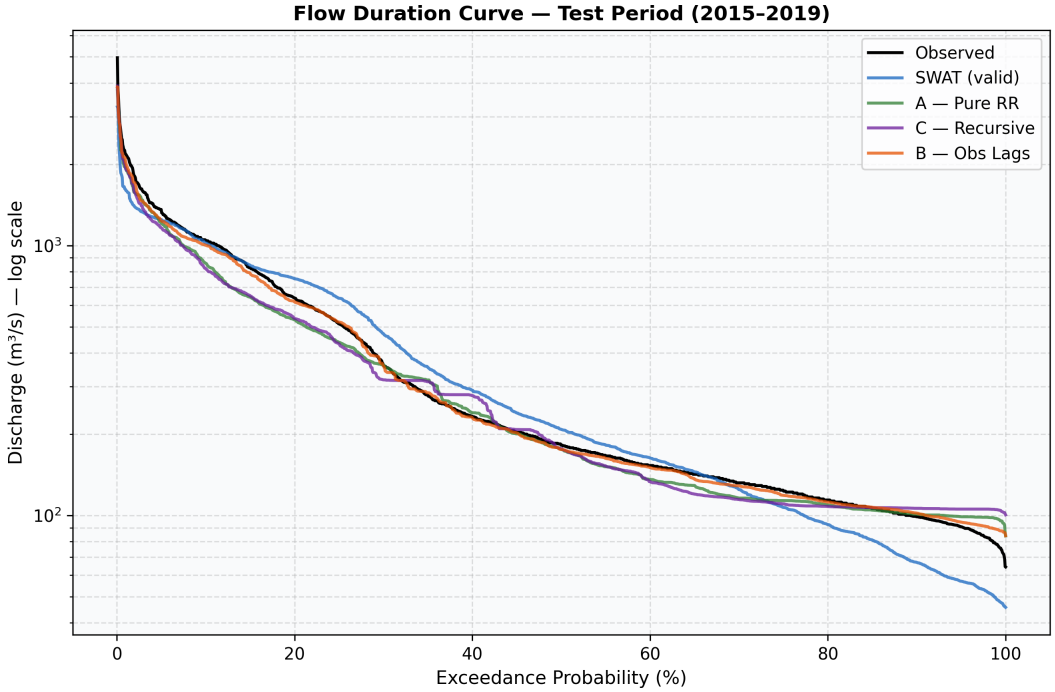

Figure 6: Flow duration curve for the test period (2015–2019) on a logarithmic discharge scale.

Scatter and Taylor diagram confirmation. Scatter plots (Figure 7) and the Taylor diagram (Figure 8) shows the aggregate and distributional findings. All configurations achieve correlation above 0.92 and cluster within a narrow band of standard deviation ratios (0.78–1.00), confirming that no model fails to capture the dominant discharge variability structure. SWAT's standard deviation ratio (0.91) is closer to unity than XGB-A's (0.87), consistent with its higher KGE arising from better variability reproduction. RF-C occupies a position between the upper-bound models (XGB-B, RF-B) and the pure rainfall-runoff configurations, reflecting its intermediate performance profile: it outperforms SWAT in NSE and KGE while operating without observed discharge at prediction time. XGB-B and RF-B achieve the tightest scatter around the 1:1 line across the full flow range, confirming their role as informative upper bounds rather than operational competitors.

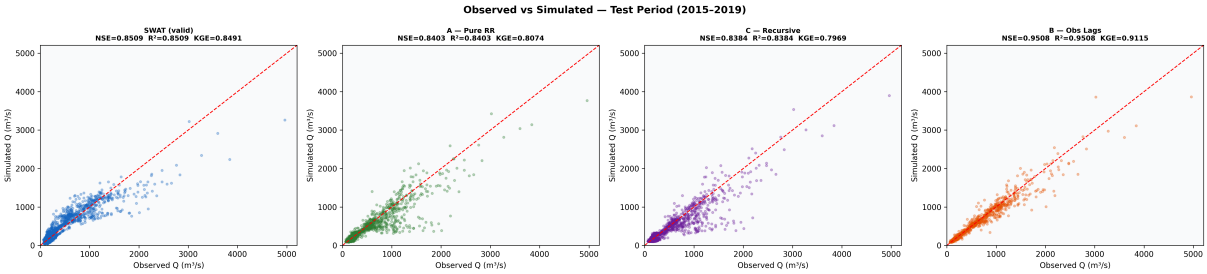

Figure 7: Observed vs. simulated scatter plots for the test period (2015–2019).

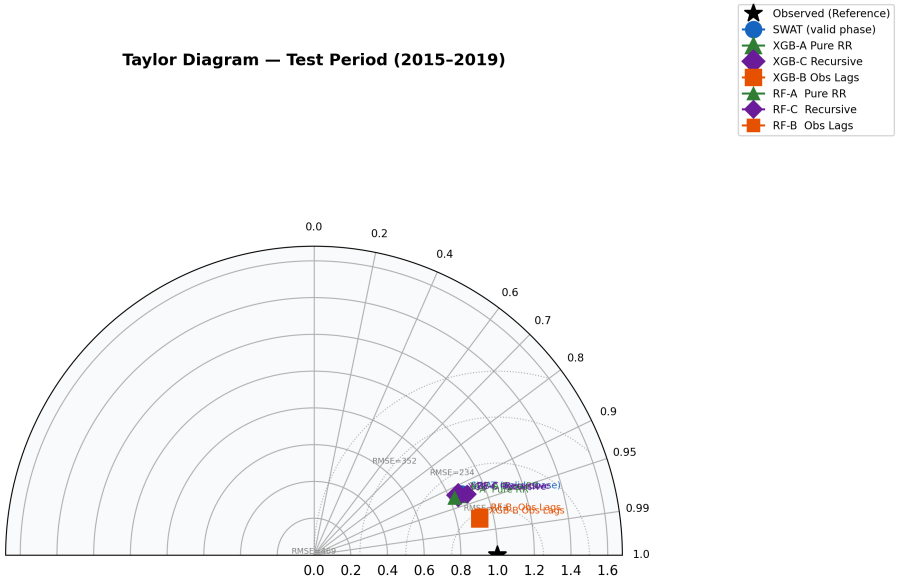

Figure 8: Taylor diagram for the test period (2015–2019). Angular position: correlation with observed discharge. Radial position: ratio of simulated to observed standard deviation. Concentric arcs: centred RMSE ($m^3\ s^{-1}$). The black star denotes perfect agreement.

4.5. Feature Importance and SHAP Analysis

4.5.1. Model A: Pure Rainfall-Runoff

The top-20 feature importances for XGB-A (Figure 9) reveal the dominant hydro-meteorological controls on discharge without discharge lag information. The 30-day cumulative precipitation (precip_30day; importance = 0.312) and Antecedent Precipitation Index (API; 0.217) together account for 52.9 % of total feature importance, dominating all other predictors. This dominance of slow-varying, memory-integrated precipitation features over short-window accumulations (precip_3day: 0.070; precip_7day: 0.023) reveals that the primary control on daily discharge in this large catchment is multi-week soil moisture state rather than the immediate precipitation event. This finding is physically consistent with the multi-day groundwater response timescales of the Kaligandaki basin: surface and subsurface drainage from the mid-hill sub-catchments integrates precipitation over weeks before contributing to discharge at the Kotagaun outlet (Fort, 2013). The month_sin feature (0.043) ranks third, encoding the seasonal climatological context — particularly the monsoon onset and recession transitions. The late_monsoon indicator (0.032) captures the distinct discharge-sensitivity regime of August–September, when hillslope saturation approaches capacity and the marginal runoff response per unit precipitation is near-unity. Individual station precipitation records carry moderate importance (st625: 0.021; st810: 0.018), with mid-hill orographic stations contributing more than Mustang rain-shadow stations, consistent with the spatial structure of precipitation-runoff linkages in this orographically dominated basin.

4.5.2. Model B: Observed Lag Model

The feature importance structure of XGB-B (Figure 9) is fundamentally different from XGB-A, dominated by Q_lag1 (importance = 0.364), which alone exceeds the combined importance of the top-three precipitation features in XGB-A. This dominance quantifies the degree to which yesterday's observed discharge determines today's discharge in a large Himalayan catchment with multi-day hydrological memory. The high Q_lag1 importance is a consequence of discharge autocorrelation rather than physical rainfall-runoff skill: Model B is largely transferring the

observed signal forward by one time step, which explains its high NSE and also why it is not an operational prediction model (Frame et al., 2022). Following Q_lag1, API (0.147) and precip_30day (0.061) remain prominent precipitation predictors, confirming that multi-week precipitation memory retains value even when discharge lag information is available. Q_lag2 (0.013) and Q_lag3 (0.011) carry progressively smaller importance, consistent with exponential decay of discharge autocorrelation at lag distances greater than one day.

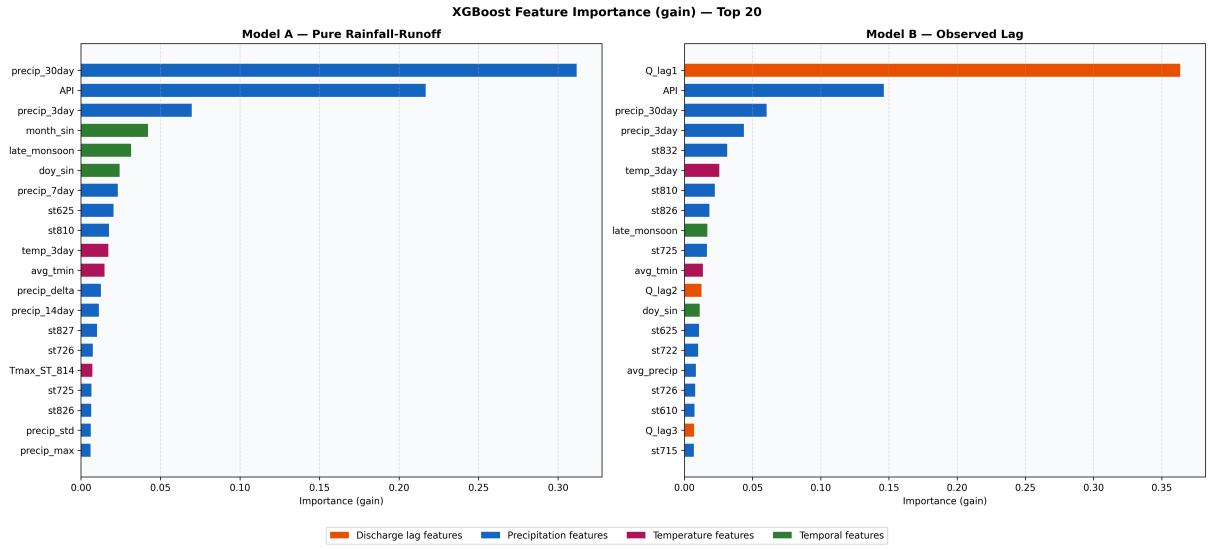

Figure 9: XGBoost feature importance (gain) for XGB-A (left) and XGB-B (right) for the test period (2015–2019). Colours indicate feature class: orange = discharge lag features; blue = precipitation features; magenta = temperature features; green = temporal features.

4.5.3. SHAP Attribution for Model B

To validate and extend the gain-based importance analysis, SHAP values were computed for XGB-B over the test period (Figure 10). Mean absolute SHAP values confirm Q_lag1 as the overwhelmingly dominant predictor, with a mean |SHAP| approximately five times larger than the second-ranked feature (API). The beeswarm plot reveals the directional structure of these contributions: high Q_lag1 values (dark red points) shift predictions strongly positive (higher simulated discharge), consistent with discharge autocorrelation as the primary prediction mechanism in Model B. For precipitation features, high API and high precip_30day values push predictions modestly positive, but with considerably more scatter than Q_lag1, reflecting the more indirect relationship between antecedent wetness proxies and next-day discharge. Tem-

perature features (melt_proxy, wet_warm) show asymmetric SHAP distributions, with positive contributions concentrated in the pre-monsoon (April–May) snowmelt period. Importantly, the SHAP ranking of precipitation features (API > precip_30day) differs from the gain-based ranking (precip_30day > API), confirming the instability of gain importance for correlated features (Lundberg et al., 2020). The SHAP decomposition provides the more reliable attribution and should be used for physical interpretation.

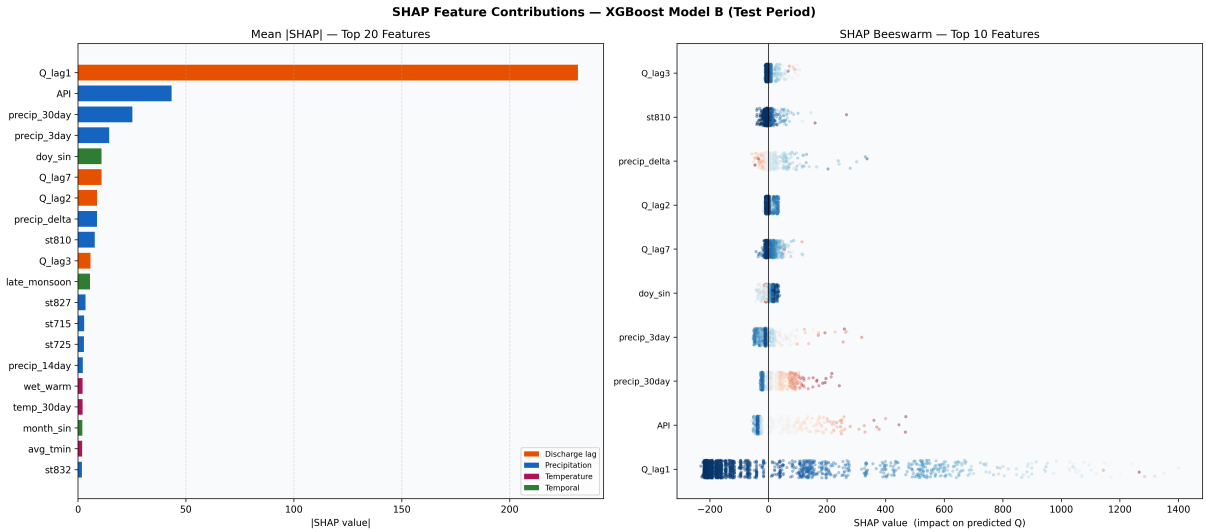

Figure 10: SHAP feature contributions for XGB-B (test period, 2015–2019). Left: mean absolute SHAP values for the top 20 features (bar chart). Right: SHAP beeswarm for the top 10 features; point colour indicates feature value (blue = low, red = high).

5. Conclusions

This study evaluated the relative performance of a calibrated process-based SWAT model and six machine learning models, namely, XGBoost and Random Forest each in pure rainfall-runoff (A), observed-lag upper-bound (B), and recursive-simulation (C) variants for daily streamflow prediction in the Kaligandaki basin, Nepal, over a five-year independent test period (2015–2019). Important conclusions were established.

First, pure rainfall-runoff ML models approach SWAT in aggregate NSE while substantially outperforming it in low-flow reproduction. XGB-A and RF-A both achieved NSE = 0.840, closely matching SWAT's NSE = 0.851 using only precipitation and temporal features, without any discharge information. However, flow duration curve analysis reveals that SWAT's near-zero

aggregate PBIAS ($-0.28\,\%$) masks compensating distributional biases: slight overestimation of high flows is offset by 40–50 % underestimation at the 100th-percentile minimum discharge (≈ 35 vs. $70\,\mathrm{m^3\,s^{-1}}$ observed). For operational applications which are driven by the lower discharges such as hydropower scheduling, irrigation allocation, environmental flow compliance, all Machine learning models outperform SWAT.

Second, the NSE gap between models B and C reveals a strongly algorithm-dependent sensitivity to recursive error propagation. For XGBoost, the total information value of observed discharge lags is XGB-B $-$ XGB-A $= +0.1105$ NSE units, but more than 100 % of this value is consumed by recursive error propagation (XGB-B $-$ XGB-C $= +0.1124$), leaving a net operational benefit for XGB-C over XGB-A of -0.0019 NSE (marginally negative). For Random Forest, the decomposition is more favourable: RF-C $-$ RF-A $= +0.0206$ NSE (positive memory benefit), and the error propagation cost RF-B $-$ RF-C $= +0.0850$ NSE is lower than the XGBoost equivalent, enabling RF-C to achieve NSE = 0.861 and KGE = 0.864 — exceeding both SWAT and all Model A configurations. This finding establishes that the operational recommendation to use recursive discharge simulation depends on the choice of algorithm: RF-C is the preferred operational model settings with available discharges, while XGB-C offers no advantage over XGB-A. The frequently cited superiority of lag-based ML models in the literature is largely attributable to access to observed discharge information at test time rather than genuine rainfall-runoff skill, as formalised by Frame et al. (2022).

Third, recursive error propagation cost is year- and monsoon-type-dependent, scaling with the duration of sustained elevated discharge rather than peak magnitude. Error propagation in XGB-C was moderate in 2015 and 2017 (well-defined monsoon peaks with clear inter-event recessions) but severe in 2018 (prolonged moderate-to-high discharge without clear recession), where the B $-$ C gap persisted for approximately six consecutive weeks.

Fourth, antecedent precipitation memory is the dominant hydro-meteorological control on discharge in this large monsoon-dominated basin. The 30-day cumulative precipitation (31.2 %) and the Antecedent Precipitation Index (21.7 %) together account for 52.9 % of XGB-A's gain-based feature importance; SHAP analysis further confirms the API as the single most influential precipitation predictor (by mean absolute SHAP value), surpassing precip_30day when feature

correlation is properly accounted for. This result reinforces the use of physical memory proxies (API, multi-day accumulations) as the core feature set for ML streamflow prediction in large Himalayan river basins.

Limitations and future directions. Several limitations of this study should guide future work. *SWAT self-containment*: Full SWAT calibration uncertainty metrics (SUFI-2 parameter ranges, p-factor, r-factor), are provided here for the first time (Section 3.2); a companion study provides the complete ArcSWAT setup. *NSE decomposition training asymmetry*: Because Models B and C are trained with observed lag features while Model A is not, the C − A term conflates training-set and test-time inference effects; a clean decomposition would require an additional "C*" model trained without lags but deployed with recursive estimates. *Seasonal performance*: This study reports aggregate test-period statistics; future work should decompose performance by flow regime (monsoon vs. non-monsoon, high vs. low flow) to fully characterise the seasonal basis of the SWAT vs. ML trade-off identified in the FDC analysis. *API sensitivity*: Although API with $k = 0.85$ was shown to be robust (± 0.004 NSE across $k \in [0.70, 0.95]$), optimal decay factors for other Himalayan basins may differ. *Single outlet*: All quantitative conclusions rest on one gauging station; multi-outlet and multi-basin validation is needed to establish transferability of the NSE decomposition percentages reported here. *Hybrid potential*: Integrating SWAT-simulated physical state variables (soil moisture, snowpack water equivalent) as additional ML features offers a promising hybrid direction that may recover physical interpretability while retaining distributional flexibility.

Declaration of Competing Interest

The authors declare no competing financial interests or personal relationships that could have influenced the work reported in this paper.

Data Availability

Daily discharge and meteorological data were obtained from the Department of Hydrology and Meteorology (DHM), Nepal (https://www.dhm.gov.np). Python code for XGBoost and Random Forest model training, feature engineering, NSE gap decomposition, SHAP analysis, and all figures is available from the corresponding author upon reasonable request and will be archived on Zenodo upon acceptance.

Acknowledgements

The authors gratefully acknowledge the Department of Hydrology and Meteorology (DHM), Nepal, for provision of long-term daily discharge and hydro-meteorological data, and the Department of Civil Engineering, Thapathali Campus, Institute of Engineering, Tribhuvan University, for institutional support.

References

Abbaspour, K. C. (2015). SWAT-CUP: SWAT calibration and uncertainty programs — a user manual. Technical report, EAWAG: Swiss Federal Institute of Aquatic Science and Technology.

Allen, R. G., Pereira, L. S., Raes, D., and Smith, M. (1998). Crop evapotranspiration — Guidelines for computing crop water requirements. *FAO Irrigation and Drainage Paper*, 56:1–300.

Andermann, C., Crave, A., Gloaguen, R., Davy, P., and Bonnet, S. (2012). Connecting source and transport: suspended sediments in the Nepal Himalayas. *Earth and Planetary Science Letters*, 351–352:158–170.

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., Kannan, N., and Jha,

M. K. (2012). SWAT: Model use, calibration, and validation. *Transactions of the ASABE*, 55(4):1491–1508.

Arnold, J. G., Srinivasan, R., Muttiah, R. S., and Williams, J. R. (1998). Large area hydrologic modeling and assessment. Part I: Model development. *Journal of the American Water Resources Association*, 34(1):73–89.

Bajracharya, A. R., Bajracharya, S. R., Shrestha, A. B., and Maharjan, S. B. (2018). Climate change impact assessment on the hydrological regime of the Kaligandaki basin, Nepal. *Science of the Total Environment*, 625:837–848.

Bhandari, A. et al. (2025). Deep learning in hydrology and flood forecasting: A review of recent advances and future directions. *Results in Engineering*, 26:104523.

Bishwakarma, M., Strommen, M., and Leira, M. (2008). Real-time sediment monitoring in hydropower plants. *International Journal on Hydropower and Dams*, 15:87–91.

Bookhagen, B. and Burbank, D. W. (2010). Toward a complete Himalayan hydrological budget: spatiotemporal distribution of snowmelt and rainfall and their impact on river discharge. *Journal of Geophysical Research: Earth Surface*, 115(F3).

Chen, T. and Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794, New York, NY. ACM.

Devkota, R. P. and Bhattarai, U. (2021). Assessment of climate change impact on the hydrology of Narayani river basin, Nepal. *International Journal of Environmental Science and Technology*, 18:2273–2288.

DHM (2017). Hydrological records of Nepal: Streamflow summary. Technical report, Department of Hydrology and Meteorology, Kathmandu, Nepal.

Fontaine, T. A., Cruickshank, T. S., Arnold, J. G., and Hotchkiss, R. H. (2002). Development of a snowfall-snowmelt routine for mountainous terrain for the soil water assessment tool (SWAT). *Journal of Hydrology*, 262:209–223.

Fort, M. (2013). The Kali Gandaki drainage: a major geomorphological agent in the Central Himalaya. *Geomorphology*, 201:262–270.

Frame, J. M., Kratzert, F., Klotz, D., Gauch, M., Shelev, 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*, 26:3377–3392.

Gauch, M., Kratzert, F., Klotz, D., Nearing, G., Lin, J., and Hochreiter, S. (2021). Rainfall-runoff prediction at multiple timescales with a single long short-term memory network. *Hydrology and Earth System Sciences*, 25:2045–2062.

Gupta, H. V., Kling, H., Yilmaz, K. K., and Martinez, G. F. (2009). Decomposition of the mean squared error and NSE: implications for improving hydrological modelling. *Journal of Hydrology*, 377:80–91.

Heggen, R. J. (2001). Normalized antecedent precipitation index. *Journal of Hydrologic Engineering*, 6(5):377–381.

Humphrey, G. B., Gibbs, M. S., Dandy, G. C., and Maier, H. R. (2016). A hybrid approach to monthly streamflow forecasting: integrating hydrological model outputs into a bayesian neural network. *Journal of Hydrology*, 540:623–640.

Immerzeel, W. W., Van Beek, L. P. H., and Bierkens, M. F. P. (2010). Climate change will affect the Asian water towers. *Science*, 328:1382–1385.

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

Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N., and Lee, S.-I. (2020). From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, 2:56–67.

Lundberg, S. M. and Lee, S.-I. (2017). A unified approach to interpreting model predictions. 30:4765–4774.

612 Marahatta, S., Devkota, L. P., and Aryal, D. (2021). Application of SWAT in hydrological simu-
613 lation of complex mountainous river basin (Part I: model development). *Water*, 13(11):1546.

614 Moriasi, D. N., Arnold, J. G., Van Liew, M. W., Bingner, R. L., Harmel, R. D., and Veith, T. L.
615 (2007). Model evaluation guidelines for systematic quantification of accuracy in watershed
616 simulations. *Transactions of the ASABE*, 50(3):885–900.

617 Nash, J. E. and Sutcliffe, J. V. (1970). River flow forecasting through conceptual models. Part I
618 — a discussion of principles. *Journal of Hydrology*, 10:282–290.

619 Neitsch, S. L., Arnold, J. G., Kiniry, J. R., and Williams, J. R. (2011). *Soil and Water Assessment*
620 *Tool Theoretical Documentation Version 2009*. Texas Water Resources Institute, College
621 Station, TX.

622 Nepal, S. (2016). Impacts of climate change on the hydrological regime of the Koshi river basin
623 in the Himalayan region. *Journal of Hydro-environment Research*, 10:76–89.

624 Pedregosa, F. et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine*
625 *Learning Research*, 12:2825–2830.

626 Pradhananga, N. S., Shrestha, M. S., and Ale, S. (2020). Simulation of hydrologic response to
627 climate change scenarios for the Budhigandaki catchment in Nepal. *Journal of Hydrology:*
628 *Regional Studies*, 31:100732.

629 Rahimzad, M., Moghaddam Nia, A., Zolfonoon, H., Soltani, J., Mehr, A. D., and Kwon, H.-H.
630 (2021). Performance comparison of an LSTM-based deep learning model versus conventional
631 machine learning algorithms for streamflow forecasting. *Water Resources Management*,
632 35:4167–4187.

633 Shen, C. (2018). A transdisciplinary review of deep learning research and its relevance for water
634 resources scientists. *Water Resources Research*, 54(11):8558–8593.

635 Shortridge, J. E., Guikema, S. D., and Zaitchik, B. F. (2016). Machine learning methods
636 for empirical streamflow simulation: a comparison of model accuracy, interpretability, and
637 uncertainty in seasonal watersheds. *Hydrology and Earth System Sciences*, 20:2611–2628.

638 Shrestha, A. B., Wake, C. P., Mayewski, P. A., and Dibb, J. E. (1999). Maximum temperature
639 trends in the Himalaya and its vicinity: an analysis based on temperature records from Nepal
640 for the period 1971–94. *Journal of Climate*, 12:2775–2786.

641 Sit, M., Demiray, B. Z., Xiang, Z., Ewing, G. J., Sermet, Y., and Demir, I. (2020). A comprehen-
642 sive review of deep learning applications in hydrology and water resources. *Water Science*
643 *and Technology*, 82(12):2635–2670.

644 Taylor, K. E. (2001). Summarizing multiple aspects of model performance in a single diagram.
645 *Journal of Geophysical Research: Atmospheres*, 106(D7):7183–7192.

646 Tyralis, H., Papacharalampous, G., and Langousis, A. (2019). A brief review of random forests
647 for water scientists and practitioners and their recent history in water resources. *Water*,
648 11(5):910.

649 USDA (1986). Urban hydrology for small watersheds. Technical Report Technical Release 55
650 (TR-55), United States Department of Agriculture, Natural Resources Conservation Service.

651 Wester, P., Mishra, A., Mukherji, A., and Shrestha, A. B., editors (2019). *The Hindu Kush*
652 *Himalaya assessment: mountains, climate change, sustainability and people*. Springer Nature,
653 Cham.

654 Yadav, M., Wagener, T., and Gupta, H. (2007). Regionalization of constraints on expected
655 watershed response behavior for improved predictions in ungauged basins. *Advances in Water*
656 *Resources*, 30:1756–1774.

657 Zhang, J., Chen, X., Khan, A., Zhang, Y., Kuang, X., Liang, X., Taccari, M. L., and Nuttall, J.
658 (2020). Daily runoff forecasting by deep recursive neural network. *Journal of Hydrology*,
659 596:126067.