



Data-Driven Estimation of the hydrologic response via Generalized Additive Models

Quentin Duchemin^{1,*}, Maria Grazia Zanoni^{2,*}, Marius G. Floriancic^{2,3}, Hansjörg Seybold², Guillaume Obozinski¹, James W. Kirchner^{3,4,5}, and Paolo Benettin⁶

Correspondence: Paolo Benettin (paolo.benettin@unil.ch)

Abstract.

Estimating the hydrologic response of watersheds to precipitation events is key to understanding streamflow generation processes. Impulse Response Functions, commonly referred to as the Instantaneous Unit Hydrograph (IUH) in hydrology, have been used for over 50 years to predict stormflow and compare catchment behaviors. These response functions are often strongly affected by modelers' choices of parameters and data preprocessing procedures, limiting their diagnostic utility and generalizability across different sites and time periods. With the increasing availability of compiled rainfall-runoff series, there is now a growing opportunity to develop new approaches that fully exploit such datasets. This paper introduces GAMCR, a novel datadriven approach for estimating impulse response functions using Generalized Additive Models. GAMCR is designed to capture the complex, nonlinear relationships between precipitation and runoff, offering a flexible and interpretable framework for the systematic analysis of hydrological responses. The model is successfully validated on synthetic data, where the true response functions are known. Additionally, we demonstrate the model's potential using real-world data from six Swiss basins with distinct hydrological behaviors. Results are fully consistent with those obtained from ERRA, another recent data-driven approach with a very different architecture, as well as with the climate and physical properties of the sites. Overall, GAMCR is a modern and effective tool for leveraging rainfall-runoff datasets to investigate the controls on hydrologic responses worldwide.

15 1 Introduction

Precipitation is generally the main water input to a landscape and the fundamental driver of streamflow generation. Quantifying how much streamflow is produced after a rain event is essential for water resources management and flood prevention, and is also useful to characterize watershed behavior. The hydrologic response (or runoff response) is usually defined as the change in streamflow induced by a given input of precipitation. Years of tracer studies have clarified that, apart from rare exceptions, such a response does not primarily consist of water that fell as precipitation during the same event, but rather by water already

^{*}These authors contributed equally to this work

¹Swiss Data Science Center, ETH Zürich & EPFL, Zürich & Lausanne, Switzerland

²Department of Environmental Systems Science, ETH Zürich, Zürich, Switzerland

³Department of Civil, Environmental and Geomatic Engineering, ETH Zürich, Zürich, Switzerland

⁴Swiss Federal Research Institute WSL, Birmensdorf, Switzerland

⁵Department of Earth and Planetary Science, University of California, Berkeley, California, USA

⁶Institute of Earth Surface Dynamics (IDYST), Faculty of Geosciences and Environment, Université de Lausanne



45



existing in the landscape (in the form of soil water and groundwater) that is quickly mobilized during the storm (Kirchner et al., 2000; McGuire and McDonnell, 2006; Botter et al., 2010; van der Velde et al., 2012; Kirchner, 2003; Knapp et al., 2024). The hydrologic response can be interpreted as reflecting the celerity with which increases in hydraulic potentials, induced by the new precipitation, propagate through the subsurface (McDonnell and Beven, 2014). Thus, stream water is generally much "older" than the most recent rainfall (McDonnell et al., 2010), although it may respond within minutes after the onset of precipitation.

The hydrologic response is a fundamental catchment signature, but its estimation is not straightforward, because catchment behavior is often nonlinear and nonstationary, meaning that the effects of precipitation inputs are not simply additive, and the same rain can generate different hydrological responses, depending on when it falls (Kirchner, 2024; Beven, 2001). The first approaches to characterize the hydrologic response came from the need to make streamflow predictions for engineering design. These approaches were based on instantaneous unit hydrographs (IUH, Sherman, 1932), analogous to the concepts of impulse response functions or transfer functions in signal processing, which are probability density functions describing how impulses of precipitation are transformed into runoff. The IUH has been typically modeled as a parametric curve like a Gamma or Weibull distribution. To cope with the complexities of runoff generation processes, the classic IUH approaches rely heavily on the concept of effective rainfall (or rainfall excess, J_e), which is the fraction of rainfall that effectively mobilizes runoff. The effective rainfall is typically modeled as a (nonlinear) function of antecedent wetness (e.g. through the popular SCS Curve Number approach, Soil Conservation Service, 1985) and acts as a filter that separates the rainfall volumes that effectively produce runoff from those that evaporate or that recharge subsurface storage. The IUH is then assumed to be linear and timeinvariant, enabling the use of standard convolution approaches to compute streamflow Q from an effective precipitation time series. The IUH theory, pioneered by the work of Sherman (1932) and further developed by Snyder (1955) and by Bruen and Dooge (1992), provided an effective and structured way to represent the relationship between (effective) rainfall and runoff. Several advances to IUH theory have been made over the years, including linking the IUH shape with basins' geomorphological properties (see Rigon et al., 2016). The IUH approach is also popular for teaching the rainfall-runoff transformation in many engineering programs (Mays, 2019).

Although IUH approaches are often successful at reproducing stormflow hydrographs, they typically require pre-processing steps to estimate effective precipitation, and to separate the hydrograph into stormflow vs. baseflow. These pre-processing steps limit the diagnostic capability of the IUH and its generality for comparing different sites and time periods. Rainfall-runoff data from hundreds of watersheds worldwide is increasingly available in harmonized databases that facilitate modeling and cross-site comparisons (e.g. Kratzert et al., 2023; do Nascimento et al., 2024). These emerging datasets create the possibility to characterize hydrological responses from many diverse watersheds, and thus to better understand their controlling factors. To characterize hydrological response without the constraints inherent in the IUH approach, Kirchner (2022) proposed a data-driven approach for estimating impulse response functions that account for nonlinear, nonstationary and heterogeneous system behavior. This approach was further developed for rainfall-runoff data and termed ensemble rainfall-runoff analysis, or ERRA (Kirchner, 2024). Although the ERRA approach shows considerable promise (e.g. Gao et al. (2025)), it is worth considering





whether other approaches can be developed to exploit the power of machine learning for innovative explorations of hydrological response.

Building on these advancements and on the widespread availability of rainfall-runoff data, here we introduce GAMCR, a data-driven approach that employs Generalized Additive Models (GAM) to estimate time-dependent Catchment Responses (CR). We present the general model architecture and provide a series of synthetic and real-world data examples to: 1) validate GAMCR and compare its performance with the ERRA approach, and 2) showcase the model's potential to estimate hydrological response at diverse watersheds, characterized by diverse properties and behaviors. The goal of GAMCR is to facilitate systematic comparisons of hydrological responses across sites where rainfall-runoff time series are available.

2 Model development

2.1 General convolution model

According to the classic convolution integral, streamflow Q is computed as the convolution of precipitation J with the stationary hydrologic response IUH, which in continuous time is expressed as:

$$Q(t) = \int_{0}^{\infty} J(t-\tau)IUH(\tau)d\tau \tag{1}$$

Here we use a discrete-time approximation to Equation (1), generalized to allow the IUH to vary with time:

$$y_{t} = \sum_{T=0}^{T_{max}} x_{t-T} h_{t-T}(T) \Delta T$$
 (2)

where y is the output flux (i.e., streamflow) at time t, x is the input flux (i.e., precipitation) T time steps earlier (i.e., at time t-T), and $h_{t-T}(T)$ is a time-variable and non-unitary response function that reflects the streamflow response to precipitation falling at time t-T, as a function of lag time T. The dependence of h on the precipitation time t-T incorporates any dependence on internal and external forcings, such as precipitation intensity and wetness conditions at the time that rain falls.

At this stage we make very few assumptions about the shape that h can take. It is not a probability density function, meaning that its area can be smaller or larger than one. While in principle h can take negative values (if this is what the system under consideration does, and is reflected in its data), we will assume that h is always non-negative (see Section 2.2). By design, h refers to the response to precipitation falling over a specific time step t-T. Any two time steps are generally expected to initiate different responses, but equation 2 is obviously ill-posed because the array $h_{t-T}(T)$ contains many more unknowns than can be constrained by the vectors y_t and x_{t-T} . Thus it is necessary to evaluate $h_{t-T}(T)$ as an average over one or more ensembles of time steps (for example during which the precipitation intensity or antecedent wetness is within a given range). In particular, the ensemble responses introduced by Kirchner (2024) can be readily obtained in a post-processing step. Given an ensemble of time points \mathcal{E} , the Runoff Response Distribution (RRD, units of 1/T) is the average response weighted by





precipitation intensity h over the selected time points \mathcal{E} :

$$RRD_{\mathcal{E}}(T) := \frac{\sum_{t' \in \mathcal{E}} x_{t'} h_{t'}(T)}{\sum_{t' \in \mathcal{E}} x_{t'}},\tag{3}$$

5 (where t' = t - T represents the time that precipitation falls), while the Nonlinear Response Function (NRF, units L/T²) is the average response multiplied by the corresponding precipitation intensity:

$$NRF_{\mathcal{E}}(T) := \frac{1}{|\mathcal{E}|} \sum_{t' \in \mathcal{E}} x_{t'} h_{t'}(T). \tag{4}$$

2.2 GAMCR model

100

110

GAMCR is a machine learning model that estimates transfer functions from flux data. GAMCR models the catchment's response to any single precipitation event as a weighted sum of spline basis functions. The time-varying coefficients of these basis functions are estimated using machine learning techniques, specifically Generalized Additive Models. As a result, we use more technical language in this section and the next, drawing terminology from the data science literature.

The problem of learning time-dependent transfer functions from rainfall-runoff data is ill-posed, meaning that considering a too large model class might result in zero training loss but with poor test error. In the machine learning community, the standard approach to cope with such badly conditioned inverse problems is to exploit prior knowledge on the structure of the studied system to either shrink the class of target functions or to regularize the optimization problem (Arridge et al., 2019). Following this approach, GAMCR is built on two core principles. First, GAMCR specifies a set of features that are assumed to be the main drivers of the catchment response to a given precipitation event. These features can be modified by the user if needed and should typically include information characterizing the catchment condition and the precipitation event considered. Second, we assume that the catchment response to a precipitation event will vary smoothly as a function of this feature vector, a structural assumption similar to the one implicitly used in the approach by Kirchner (2022).

Second, we expect the transfer functions, $T \mapsto h_{t'}(T)$, to exhibit sharp peaks for short time lags, that progressively smooth out as the lag time T increases.

With these guiding principles, we model the transfer functions as follows:

105
$$h_{t'}(T) = \sum_{\ell=1}^{L} g_{\ell}(\mathbf{z}_{t'}) b_{\ell}(T),$$
 (5)

where $(b_\ell)_{\ell\in[L]}$ are B-splines constructed by considering an irregular spacing of knots, \mathbf{z}_t' is a feature vector describing both the catchment conditions and the precipitation event at time t' and g_ℓ is a GAM. The basis functions $(b_\ell)_{\ell\in[L]}$ are illustrated in Figure 1, highlighting that the knot density is much higher for shorter lags, while the knots become more spaced out for longer lags. This design enables the model to capture the large variability of the transfer functions at short lags, while still accounting for potentially long recessions. The feature vectors $\mathbf{z}_{\mathbf{t}'}$ used in GAMCR are the intensity of the precipitation event at time t', the weighted averages of both the past precipitation and the past evapotranspiration over different time windows, and the sine and cosine of the fractional year.





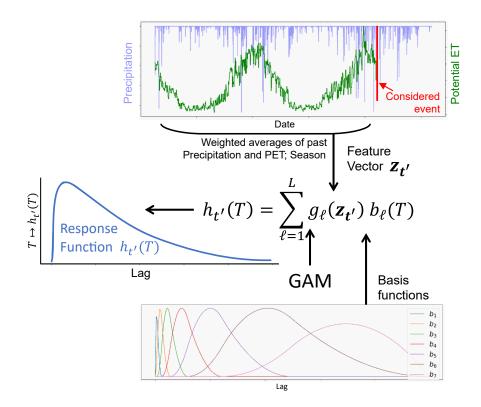


Figure 1. Overview of GAMCR. Given some precipitation event of interest occurring at time t', GAMCR computes a feature vector $\mathbf{z}_{t'}$ including information on the system up to time t'. The response function is expressed as a weighted sum of spline basis functions, $(b_{\ell})_{\ell \in [L]}$, where the weights are derived from $\mathbf{z}_{t'}$ through L distinct Generalized Additive Models $(g_{\ell}(\cdot))_{\ell \in [L]}$.

Since we model the functions $(g_{\ell})_{\ell \in [L]}$ using GAMs, one can write

$$g_{\ell}(\mathbf{z}_{t'}) = \boldsymbol{\xi}_{t'}^{\top} \boldsymbol{\gamma}_{\ell}, \tag{6}$$

where $(\boldsymbol{\xi}_{t'}^{\top})_{t'}$ is the design matrix of the GAM. Each entry of $\boldsymbol{\xi}_{t'}$ corresponds to one of the spline basis functions evaluated at a given feature (i.e. a specific entry of $\mathbf{z}_{t'}$). We have:

$$y_t = \sum_{\ell=1}^{L} \left(\sum_{T=0}^{T_{max}} x_{t'} b_{\ell}(T) \boldsymbol{\xi}_{t'} \right)^{\top} \boldsymbol{\gamma}_{\ell} = \sum_{\ell=1}^{L} \boldsymbol{W}_{t,\ell,:} \boldsymbol{\gamma}_{\ell} = \overrightarrow{\boldsymbol{w}}_{t}^{\top} \overrightarrow{\boldsymbol{\gamma}},$$
 (7)

where $\overrightarrow{\boldsymbol{\gamma}} := \operatorname{vec}(\boldsymbol{\gamma}_{\ell}, \ell \in [L])$ and $\overrightarrow{\boldsymbol{w}}_{t} := \operatorname{vec}(\boldsymbol{W}_{t,\ell,:}, \ell \in [L])$ given by $\boldsymbol{W} \in \mathbb{R}^{n \times L \times d_{s}}$ such that:

$$\mathbf{W}_{t,\ell,:} = \sum_{T=0}^{T_{max}} x_{t'} b_{\ell}(T) \boldsymbol{\xi}_{t'}. \tag{8}$$

120 Here, d_s is the number of features resulting from the GAM formulation.





2.3 Model training

With GAMs we can encode prior knowledge and control overfitting by using penalties and constraints during training. In our case, we consider two smoothness-inducing penalties. The first one promotes the smoothness of the functions g_{ℓ} by penaltzing the second order derivative, as commonly done in the GAM literature (see Hastie et al. (2017)). This penalty ensures that the coefficients of the transfer functions in the basis $(b_{\ell})_{\ell}$ smoothly evolve with respect to the catchment features $z_{t'}$. This penalty acts on the L time-dependent coefficients of the transfer functions in the basis $(b_{\ell})_{\ell}$ independently. The second regularization term promotes the smoothness of the transfer functions globally by adding a similar penalty on the model coefficients.

The final optimization problem considered is:

$$\min_{\left(\boldsymbol{\gamma}_{\ell}\right)_{\ell \in [L]} \geq \mathbf{0}} \frac{1}{n} \sum_{t=1}^{n} \left(y_{t} - \sum_{\ell=1}^{L} \boldsymbol{W}_{t,\ell,:} \boldsymbol{\gamma}_{\ell} \right)^{2} + \lambda_{1} \sum_{\ell=1}^{L} \boldsymbol{\gamma}_{\ell}^{\top} \boldsymbol{P}_{1} \boldsymbol{\gamma}_{\ell} + \lambda_{2} \sum_{1 \leq \ell, k \leq L} \frac{1}{n} \sum_{t'=1}^{n} (\boldsymbol{\xi}_{t'}^{\top} \boldsymbol{\gamma}_{\ell}) \left[\boldsymbol{P}_{2} \right]_{\ell,k} (\boldsymbol{\xi}_{t'}^{\top} \boldsymbol{\gamma}_{k}), \tag{9}$$

which can be equivalently written using a vectorized formulation as:

$$\min_{\overrightarrow{\gamma} \ge \mathbf{0}} \frac{1}{n} \sum_{t=1}^{n} (y_t - \overrightarrow{\boldsymbol{w}}_t^{\top} \overrightarrow{\gamma})^2 + \overrightarrow{\gamma}^{\top} [\lambda_1 \boldsymbol{P}_1' + \lambda_2 \boldsymbol{P}_2'] \overrightarrow{\gamma}, \tag{10}$$

where, denoting by \otimes the Kronecker product between two matrices we have defined

$$P'_1 := \operatorname{Id}_L \otimes P_1, \quad \text{and} \quad P'_2 := P_2 \otimes \left(\frac{1}{n} \sum_{t'=1}^n \boldsymbol{\xi}_{t'} \boldsymbol{\xi}_{t'}^{\top}\right)$$
 (11)

Provided that the hyperparameters λ_1 and λ_2 are not both zero, the optimization problem (10) has a strongly convex objective function with convex constraints. As a result, it admits a unique optimal solution, and the projected gradient descent algorithm is guaranteed to converge to this solution provided that the learning rate is set small enough (cf. Boyd (2004)). In practice, the parameters $\overrightarrow{\gamma}$ are initialized by solving the unconstrained version of the problem, which involves computing the minimum L^2 -norm solution via the pseudoinverse of a matrix. This initial solution is then projected onto the positive orthant, after which the projected gradient descent algorithm is applied. The learning rate starts at a large value and is gradually and adaptively reduced throughout the iterations to ensure a strict decrease in training loss at each step.

The matrix W is precomputed offline prior to running the projected gradient descent algorithm, and parallel computation can be employed to obtain W quickly. This precomputation significantly accelerates the training process by eliminating redundant calculations.

2.4 Software GAMCR v1.0 description

The model developed in Section 2.2 has been implemented in the Python language as the software GAMCR v1.0. To use GAMCR, the user must provide time series of precipitation, streamflow, potential evapotranspiration and the corresponding dates and times at equally spaced time intervals. The software operates in a series of steps to ensure accurate and efficient analysis. First, users can use a pre-defined notebook to ensure that their data has the proper format (e.g. column names that



150

155

160

175



conform to the software's requirements). Next, a predefined script is run to perform key precomputations, including the calculation of the matrix W, which significantly enhance the efficiency of the model training process. These precomputations are completed within a few seconds to a few minutes on a standard laptop for a decade of hourly data. Once these precomputations are completed, users can proceed to train the GAMCR model on their dataset. Let us stress once again that the number of basis functions L used by the model is automatically computed based on the maximum lag $T_{\rm max}$. With $T_{\rm max}=5$, 10 or 15 days the model uses 6, 7 or 8 basis functions, respectively. After the model has been trained, users can launch another predefined script to compute key statistics of interest, such as the NRFs over predefined ensembles (such as different precipitation quantiles) and the RRD. These results are automatically saved for further analysis. A detailed tutorial is provided in the online documentation of the GAMCR package, where users can reproduce the results of this paper for the Euthal catchment. The tutorial offers a step-by-step explanation of each stage, equipping users with the necessary tools to apply GAMCR effectively to their own datasets. Overall, GAMCR can be efficiently used on personal laptops, with model training on 20 years of hourly data typically taking around 30 minutes for $T_{\rm max}=10$ days.

3 Model testing

Developing strategies to rigorously quantify the performance of trained machine learning models is essential. In the case of the hydrologic response, the evaluation step is particularly important because the real-world impulse response functions cannot be measured directly and the model is trained on streamflow data only.

Below, we describe two datasets that serve two different purposes. A synthetic dataset (Section 3.1) is used to validate the model, because the estimated response can be compared against the benchmark "ground truth" response, which is exactly known (unlike in real-world systems). A real-world dataset (Section 3.2), which includes measurements from six diverse catchments across Switzerland, is used to showcase how the model can be used to estimate the hydrologic response at different locations.

170 3.1 Synthetic data

The synthetic dataset was generated using precipitation and air temperature measurements available from the Federal Office of Meteorology and Climatology (MeteoSwiss) for the station of Lugano, along with streamflow data from the nearby gauging station Chiasso, Ponte di Polenta, on the Breggia River. These real-world measurements were used to calibrate a lumped nonlinear and nonstationary conceptual model (Section S2), allowing us to create a synthetic streamflow time series (40 years at hourly resolution) that closely mirrors actual measurements (case A). To explore different hydrological responses, we adjusted the model parameters to represent both a more damped (case B) and a more flashy (case C) hydrologic system. By working with these synthetic yet realistic datasets, we can rigorously assess the model's performance, because the underlying mechanisms are exactly known and the data are free from disturbances such as dams, hydropeaking, or leakages. Details of the approach employed in the model and the parameters used are provided in section S2.





The generated synthetic time series are shown in Figure 2 over an example 4-month period. The figure shows clearly that, compared to the reference streamflow series (case A), the damped series (case B) has lower peaks and longer recessions, while the flashy series (case C) has higher peaks and similar recessions. The data also clearly show the nonlinearity and nonstationarity of hydrologic systems, as some precipitation events cause almost no streamflow response (e.g. in June 2010) while others may cause a sharp response (e.g. in late August 2010). To compute the response functions for the synthetic data

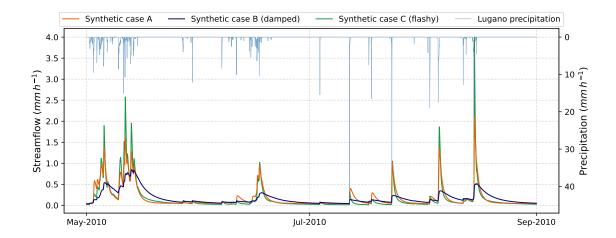


Figure 2. Example of the synthetic streamflow time series (for snow-free months in 2010): case A is the reference (orange curve); case B is a more damped response (purple curve), with lower peaks and longer recessions; case C is a flashier response with higher peaks (green curve). Additionally, the Lugano precipitation time series is shown (light blue curve) with an inverted y-axis for comparison.

185 (ground-truth response), we simply ran the lumped hydrological model as many time as there were time steps with nonzero precipitation. In every simulation, we masked a different time step by setting its precipitation to zero. The hydrologic response to precipitation occurring on a specific time step was then computed as the difference between the modeled series with and without precipitation over that time step. This approach provides responses for each event individually, which can be aggregated to compute ensemble responses over e.g. particular periods, precipitation events or antecedent conditions.

190 3.2 Real-world data

We compiled a 15-year record (2005–2019) of real-world, hourly precipitation-runoff data from six Swiss watersheds (Fig. 3). Streamflow time series were provided by the Federal Office for the Environment (FOEN). Precipitation data were sourced from the 'CombiPrecip' product, developed by MeteoSwiss (MeteoSwiss CombiPrecip). Potential evapotranspiration time series were computed based on air temperatures provided by MeteoSwiss, through the Hargreaves method from Hargreaves and Samani (1985) (implemented through the Python Pyeto package https://github.com/woodcrafty/PyETo) and then uniformly distributed across each day at hourly intervals. At each site we also extracted key catchment attributes (Table 1) and computed the mean monthly precipitation, streamflow, and potential evapotranspiration over the study period (Figure 4). We selected these sites because they are all medium-sized (between 34–185 km²) but with different hydrological regimes, elevation and





soil depths, which we expect will be reflected in substantially different hydrologic responses. The sites were also selected because they are not much affected by the presence of glaciers, they have a natural flow regime (no dams or major abstractions) and their data records are complete and reliable. Additional analyses of the catchment characteristics are provided in the Supplement, in Section S1.

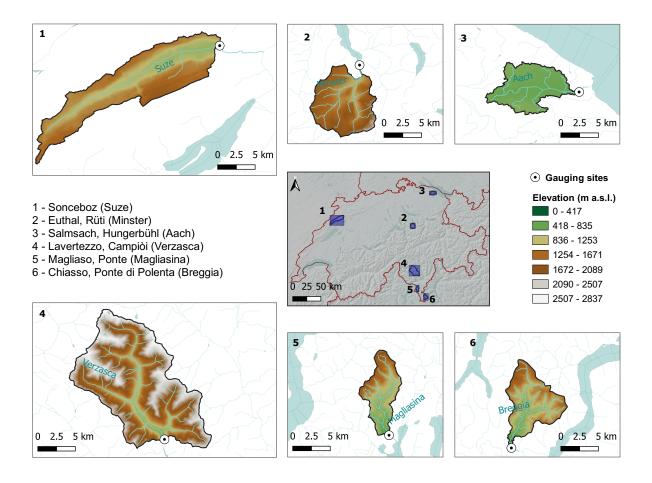


Figure 3. Map of Switzerland showing the six catchments analyzed, along with their corresponding gauging stations (listed in the left with the river names in brackets). Each catchment is displayed in a separate plot for a detailed view of its dimensions and elevation ranges. Numbers mark the catchments' locations within Switzerland and can be seen on the map in the center. The sixth gauging station (Chiasso, Ponte di Polenta) also provided the streamflow time series used to create the synthetic dataset with precipitation data from Lugano.

As snow introduces complexities in catchment response, such as delayed runoff generation and temperature-driven melt rates, we focused our analysis on snow-free periods only. We considered as snow-free periods the months from May to October, inclusive, except for two basins at the highest altitudes (Euthal and Lavertezzo, with maximum and mean elevations above 2200 and 1300 m a.s.l., respectively), for which we assumed that first snow-free month is June (Figure 4).





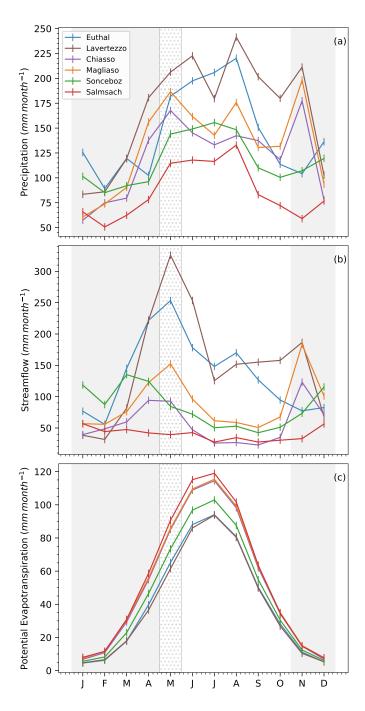


Figure 4. Hydrological regimes, in terms of monthly mean precipitation (a), streamflow (b), and potential evapotranspiration (c) for the six sites. The time series were averaged over the complete period of study (2005-2019). The light grey shadowed areas indicate what we considered as snowy periods with potential snow-melt effects on streamflow, including also May for Lavertezzo and Euthal (dotted grey shadowed band).



230



Table 1. Overview of the gauging stations and their catchment features: associated river; mean elevation; mean slope; area; mean soil depth; mean permeability; hydrological regime type; snow-free months considered in the study.

Station	River	Mean Alt. (m a.s.l.)	Mean Slope (°)	Area (km²)	Mean soil depth (cm)	Mean permeability (cm d ⁻¹)	Regime type	Snow-free months
Sonceboz	Suze	1036.95	14.77	127.25	49.12	72.09	Jura-nivopluvial	May-Oct
Euthal, Rüti	Minster	1346.91	22.20	59.13	33.87	54.27	Transition nival	June-Oct
Salmsach, Hungerbuühl	Aach	472.47	3.42	47.38	69.57	52.22	Pluvial	May-Oct
Lavertezzo, Campiòi	Verzasca	1655.71	38.49	185.12	17.16	81.56	Southern nivo-pluvial	June-Oct
Magliaso	Magliasina	928.25	28.44	34.38	29.44	96.77	Southern pluvio-nival	May-Oct
Chiasso, Ponte di Polenta	Breggia	934.40	33.21	47.10	20.87	75.98	Southern pluvio-nival	May-Oct

3.3 Implementation details

While our model is designed to estimate the hydrologic response to each precipitation event, we are primarily interested in the model's ability to reproduce the ensemble responses (RRD or NRF) over given conditions of precipitation intensity or antecedent wetness. Therefore, the model will be tested over ensemble responses. This also offers the opportunity to estimate the hydrologic response—and its main statistics—with ERRA and assess the consistency between GAMCR and ERRA.

We tested the need for optimization of the hyperparameters λ_1 and λ_2 through initial (and computationally expensive) cross-validation experiments. Since we obtained only minor improvements over the default values $\lambda_1=10^{-3}$, $\lambda_2=1$, we consistently used the defaults across all applications. Since we are only interested in the evaluation of the hydrologic response up to a few days after precipitation, we kept the hyperparameter $T_{max}=24\times10$ hours. The positions of the knots to get the B-splines basis functions b_ℓ follow an exponentially increasing sequence, starting at 0 with an initial step of 1. After each step, the step size doubles, leading to a pattern where knots are densely spaced at the beginning and become increasingly sparse as values grow. Following this procedure, the value of T_{max} automatically sets the number of basis functions to L=7 in our case.

In real-world data, the response to very small rainfall events may be easily hidden by measurement noise and other processes. While these events are not particularly relevant for the hydrologic response, they may corrupt the training phase. Hence it is convenient to set a precipitation intensity threshold J_{th} and train the model only for events that exceed J_{th} . We trained GAMCR using $J_{th} = 0.05$ mm/h.

The results from ERRA were obtained using the R scripts accessible at the following repository: https://doi.org/10.16904/envidat.529, as specified in Kirchner (2024). The RRD curves were computed considering a maximum lag of 40 hours. Initial estimates of precipitation bins were automatically generated by the algorithm, invoking six approximately even ranges, while ensuring a minimum threshold of 40 nonzero values in each precipitation bin. To improve comparability across models, the same precipitation ensembles were used to average the true transfer functions and the GAMCR estimates. Using a coarser input data resolution is beneficial to ERRA when the hydrologic response is long relative to the input temporal resolution (because in such cases, it can be difficult to separate the overprinted effects of input signals at closely spaced lag times). Using a coarser time step helps clarify these impacts. For this reason, after some initial testing, the flashy, base, and damped synthetic input time series are aggregated into 2-, 3-, and 6-hour time steps, respectively.





4 Results

235

240

255

260

4.1 Model validation

The model was trained (Section 2.3) on the synthetic data (Section 3.1), which consists of three cases: the reference response, a flashier response, and a more damped response. We validate the model by first computing the hydrologic response in the form of NRF over six quantiles of precipitation intensity, and comparing it against the ERRA estimates and the benchmark generated directly from the model Figure 5.

As a result of the different aggregation of the input time series for the three synthetic data sets, their precipitation intensities (and thus the bins used in Figure 5a-c) appear different, although the original hourly input data are the same.

Figure 5 shows that GAMCR accurately estimates the transfer functions on synthetic data, particularly in the flashy and damped scenarios, where their curves nearly overlap with the benchmark. In the base case, (panel b) the peak value and tail of the response are well captured, but the peak timing is systematically early compared to the benchmark. Overall, in these three cases characterized by very different responses (Figure 5d) ERRA and GAMCR provide generally consistent estimates.

We also computed the peak height, peak lag and runoff volume of the NRF, and explored their relationship with precipitation intensity (Figure 6). The results highlight GAMCR's ability to accurately estimate key quantities related to the magnitude of the catchment's response (peak height and runoff volume). These statistics are also very consistent with those estimated by ERRA. For both approaches, the flashy case remains the most sensitive for estimation, with GAMCR underestimating runoff volume for intermediate precipitation values (from 10 to 25 mm h ⁻¹) but accurately capturing peak height. In the base case, GAMCR slightly underestimates runoff volume while maintaining accurate peak height estimates. For the damped scenario, it closely matches ground truth values for peak height and produces nearly overlapping runoff volume estimates. Overall, both approaches show a strong consistency in their peak height and runoff volume estimates across different scenarios.

Despite the models' strong performance in estimating the magnitude of the catchment response, both face challenges in predicting peak lag, though in opposite ways. ERRA tends to produce more variability across the NRFs (as shown by dashed lines with triangles in all the (c) panels). By contrast, GAMCR tends to produce lag values that are much less variable than the benchmark across different precipitation ranges.

4.2 Estimation of real-world hydrological responses

When applying GAMCR to real-world data it is not possible to validate its accuracy in estimating hydrologic response, because the true response is not known. However, it is instructive to compare the modeled vs measured streamflow series, for both the training and test periods. Since the model was not developed for the purpose of reproducing streamflow, its performance should not be compared to hydrologic models that are designed to maximize fit, but the simulated hydrograph serves as a valuable diagnostic tool. For example, periods where the modeled hydrograph deviates significantly from the measurements could be flagged as unreliable and excluded from the analysis. To support this evaluation, Figure 7 shows the streamflow predictions generated by GAMCR.





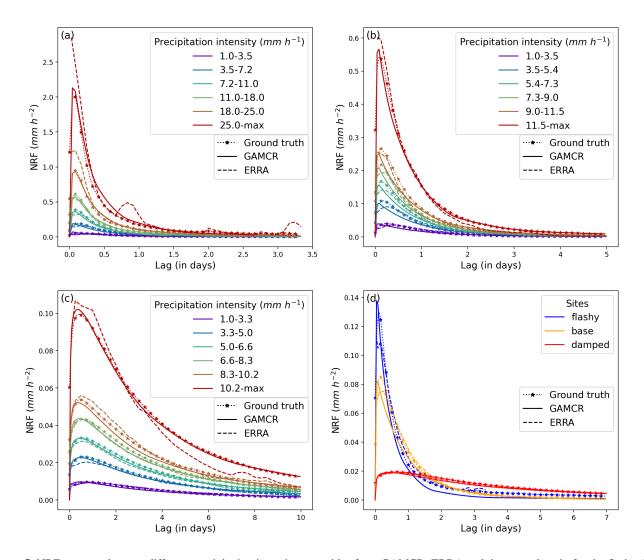


Figure 5. NRFs averaged across different precipitation intensity ensembles from GAMCR, ERRA and the ground truth, for the flashy (a), base (b), and damped (c) synthetic time series. Readers should note the different scales of NRFs between flashy, base and damped scenarios. Panel (d) combines the overall average NRFs for the three cases in a single plot.

The agreement between calibration and out-of-sample results varies across basins. Sonceboz and Magliaso show the strongest consistency, with only minor over- or underestimations. Lavertezzo follows a similar pattern but has some overestimated values in calibration. Salmsach and Chiasso, by contrast exhibit considerable dispersion and out-of-sample overestimation, suggesting lower predictive performance. Performance at Euthal is intermediate between these two groups, with overestimation of low out-of-sample streamflow values. These results suggest that the performance of GAMCR in reproducing streamflow is not directly correlated with the hydrological characteristics of the basins. This is even more visible when looking at the model performance aggregated over streamflow quantiles (Figure S2 in the Supplement), where the fit is consistently good across sites



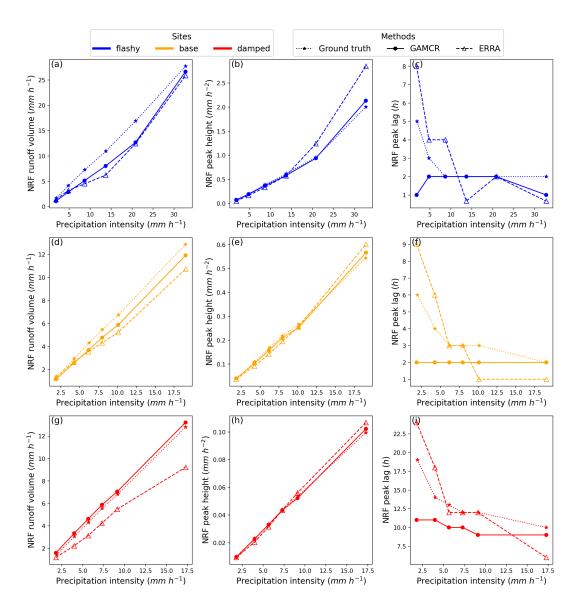


Figure 6. Different statistics computed on NRFs obtained from either GAMCR, ERRA, or the ground truth averaged across different precipitation intensity ensembles for the flashy, base, and damped datasets. Figures (a), (b) and (c) respectively depict the NRF runoff volume, NRF peak height and the peak lag.



280



and only a minor underestimation of the lowest flow conditions stands out. Timeseries plots for the validation period (Figure S3 in the Supplement) indicate that the temporal dynamics of the predicted hydrograph are appropriate and there are no periods that should be flagged and removed from the analysis.

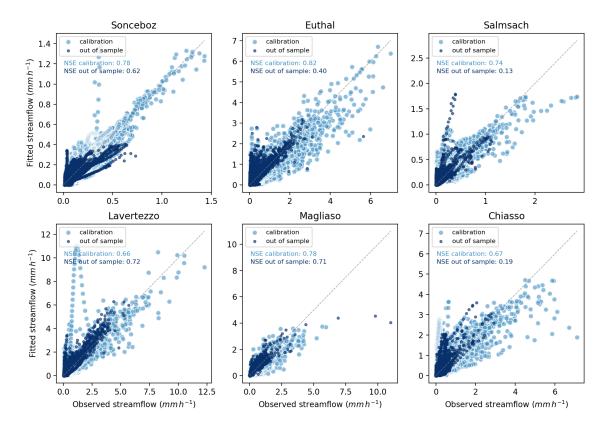


Figure 7. Fitted streamflow estimated through GAMCR for the six investigated sites. The larger light blue dots show measured and fitted discharges during 2005-2017 (calibration period); The smaller dark blue dots indicate measured and predicted discharges during 2018-2019 (out-of-sample predictions). 1:1 lines are shown in grey.

Figure 8 presents the weighted average RRDs and the peak heights of the NRFs estimated by ERRA and GAMCR for the six sites in the real-world dataset. Computations consider all events whose precipitation intensity exceeds 0.5 mm h⁻¹. The results align well with the hydrological regimes and characteristics of the basins (see Table 1 and Figures 3 and 4). The Sonceboz basin, in particular, shows a very flat runoff-response distribution, which is consistent with the relatively low mean slope, large area, and elongated shape of its basin. These features, along with its moderate permeability and location in the Jura's pluvio-nival region, contribute to the basin's very damped runoff response. A slightly flashier response is observed in the Salmsach basin, which has low mean slope, low permeability, and a pluvial hydrological regime. This results in a damped response, though less damped than Sonceboz's. The Chiasso and Magliaso basins exhibit similar peak values, but with different response shapes. Despite similarities in altitude and mean slope, Chiasso is larger than Magliaso and has lower permeability, consistent with the



285

290

295

300



larger area under its RRD (i.e., runoff coefficient). The flashier response in Magliaso is consistent with its high mean slope, in common with Lavertezzo and Euthal, where the flashiest responses (at least for GAMCR) are observed. Lavertezzo and Euthal are characterized by the highest altitudes, highest annual precipitation and lowest annual potential evapotranspiration. The higher RRD peak for Euthal compared to Lavertezzo is consistent with the lower permeability in the Euthal basin. Overall, the weighted average RRDs provided by both GAMCR and ERRA are broadly consistent with the distinctive characteristics of each basin. The results shown in Figure 8 demonstrate the general consistency between the GAMCR and ERRA approaches across the different sites. Only two basins exhibit some discrepancy: the Chiasso and Lavertezzo basins (purple and brown curves, respectively). In both these cases, GAMCR estimates a more pronounced RRD peak than ERRA within the first 7 hours, and a slightly lower tail after 10 hours. The estimated NRF peaks for different precipitation intensities for these sites (Figures 8d and 8f) are consistent between ERRA and GAMCR for most precipitation bins, but deviate slightly for the highest one. Overall, the responses estimated by GAMCR and ERRA are broadly similar, and since the models work very differently, consistency in their estimates increases our confidence in both approaches.

4.3 Effects of precipitation intensity and antecedent wetness

GAMCR can be used to investigate how variations in precipitation intensity and antecedent wetness affect the hydrologic response. Here we explore such effects at the six real-world sites. To characterize precipitation intensity, we use the same six precipitation intervals defined in Section 4.2 above. As a proxy for antecedent wetness, we use the values of streamflow during the timestep prior to the precipitation event under consideration, which we separate into five ranges. We then aggregate the individual response (RRD curves) over each class of precipitation intensity or antecedent wetness. Results are shown in Figure 9, where we plot the RRD peak height (not to be confused with the peak of the NRF shown in Figure 8) against precipitation intensity and antecedent wetness.

As Figure 9a shows, the RRD peak heights do not vary systematically with precipitation intensity. By contrast, Figure 9b demonstrates clear increasing trends in RRD peak heights with increasing antecedent wetness. Nearly all sites exhibit at least a threefold increase in peak heights across antecedent wetness levels, with the exception of Lavertezzo, which shows a rise in peak heights just for only the last two bins of antecedent wetness. Chiasso, in particular, displays the highest variability, with peak heights spanning almost an order of magnitude (from 0.006 to 0.05 h ⁻¹). Notably, for each site, the highest two antecedent wetness levels are widely separated, leading to a marked increase in RRD peak heights. These findings highlight a clear nonstationary response of the six catchments, strongly influenced by their antecedent wetness.

310 5 Discussion and Conclusions

We introduced a model based on GAMs to estimate the hydrologic response of watersheds based on precipitation-runoff data. The model was validated against three benchmark synthetic datasets and showed excellent agreement with the response curves of the underlying benchmark model, based only on its input and output time series (Figure 5). While the accurate reproduction of the individual responses goes beyond the scope of the model, the ensemble responses (RRD and NRF curves) proved





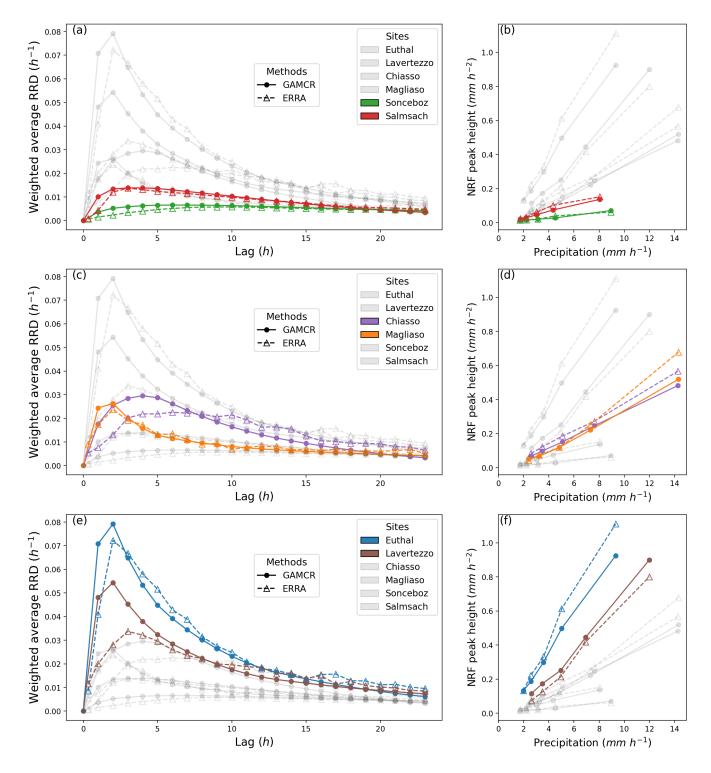


Figure 8. Hydrologic responses and their relationship with precipitation intensity for GAMCR and ERRA. Panels (a), (c) and (e): Weighted average RRD, where we keep time points with precipitation intensity above 0.5 mm h⁻¹. Panel (b), (d) and (f): NRF peak heights against precipitation intensity.

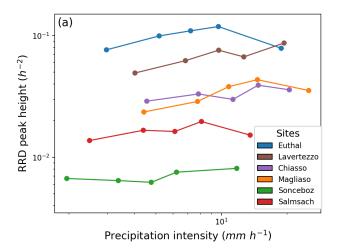


325

330

335





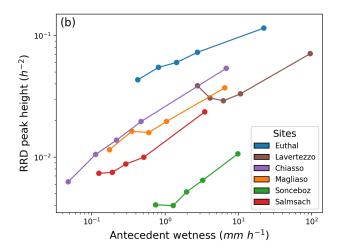


Figure 9. Peak height of RRDs when stratifying with respect to precipitation intensity (panel a) or antecedent wetness (panel b).

accurate. Closer inspection of the statistics of the responses (Figure 6) showed that GAMCR accurately estimated NRF peak height and volume across different precipitation bins. By contrast, the timing of the NRF peak was generally not very accurate, with GAMCR systematically underestimating the peak lag. While this behavior can likely be improved through a different organization of the basis functions that form the core of the response (Section 2.2), GAMCR should currently not be used to estimate the timing of the hydrologic response. Comparisons between GAMCR and ERRA highlight that these two models, despite their very different architectures, provide similar hydrologic responses that closely match the (synthetic) ground truth.

Additionally, we analyzed the runoff response during snow-free periods for six Swiss catchments with diverse climatic and physical characteristics (Section 4.2). Because the hydrologic response of a catchment is not directly measurable, verifying the accuracy of GAMCR is challenging in real-world settings. Among the diagnostic tools that help build confidence on the results (beyond the benchmark tests of Section 4.1), we verified that the modeled streamflow was generally realistic for both in-sample and out-of-sample data (Figure 7) and compared GAMCR's RRD and NRF statistics with those obtained from ERRA (Figure 8). GAMCR produced results that were closely aligned with ERRA and consistent with the properties of the catchments. For example, the Salmsach catchment, with flatter topography and deeper soils than Euthal, had a slower and less marked average response to rainfall. We conclude that GAMCR is a robust tool to study runoff response behavior in real-world catchments. As such, it enables advanced data-based analyses such as quantifying the effects of precipitation intensity and antecedent wetness on the average response peak (Figure 9).

Since we have often referred to the ERRA approach (Kirchner, 2024) in our analyses, it is worthwhile to clarify the differences and similarities between ERRA and GAMCR. Both methods aim to estimate runoff response to precipitation based on time series data, and they both can quantify nonlinear and nonstationary runoff responses to precipitation. Both methods also implicitly assume that precipitation intensity and catchment conditions are the main controls on the catchment response. However, the two approaches achieve their (common) objective in radically different ways. ERRA fundamentally works on



345

350

355

360

365



ensemble responses rather than single events. It extracts information from the entire precipitation-runoff time series or from portions of it that are selected to investigate different periods or conditions (provided that each portion has sufficient data). In contrast, GAMCR estimates the hydrologic response to each individual precipitation event using combinations of spline basis functions, with coefficients determined through machine learning techniques (Sections 2.2-2.3). These individual responses can then be aggregated to ensemble responses. These different starting points result in different ways to run the models. GAMCR is based on a single training phase to estimate all the responses. Then, users can simply aggregate such responses in various ways as a post-processing phase. Instead, ERRA runs instantly but any sub-setting of the time series (for periods or conditions of particular interest, for example) needs to be specified a priori and the code is re-run for each analysis. The way ERRA and GAMCR are parameterized limits the types of transfer functions they can estimate, embedding specific assumptions about their shape. ERRA produces piecewise linear transfer functions, which might take negative values, especially when the water balance in the data is not maintained. In contrast, GAMCR ensures strictly positive transfer functions and promotes smoothness by relying on smooth basis functions. Other minor operational differences between ERRA and GAMCR include the potential need to aggregate the temporal resolution of the data to improve the estimate (ERRA) and the need for potential evapotranspiration series, along with precipitation and runoff (GAMCR, although in the absence of potential evapotranspiration data the user may simply change the default set of features of GAMCR by removing the ones based on PET). Finally, ERRA not only estimates statistics and responses but also quantifies their uncertainty through standard errors. In contrast, GAMCR currently lacks this capability, highlighting the need for future integration of uncertainty quantification.

These first applications of GAMCR to synthetic and real-world data help us identify some current model limitations and encourage further model development. While the peak value and area of the aggregated response functions proved accurate, the timing of the response was not, systematically underestimating the peak lag (Figure 6). Next model developments could target different or denser basis functions capable of improving the estimation of the peak lag. We also stress that while the model simulates the response to every time step with precipitation, it was only evaluated on its capacity to reproduce ensemble behavior and so we do not recommend using it to evaluate individual responses or predict streamflow time series. Additional features that could be implemented in the future include an uncertainty estimation tool capable of providing accurate uncertainty bounds in the response, and the opportunity to integrate additional data (e.g. soil moisture series from sensors or remote sensing products) while training the model.

Catchments' capacity to mobilize water after storm events is a distinctive feature that is relevant for water resources management and useful to characterize catchment behavior. Quantifying the runoff response to precipitation using data-driven approaches is challenging due to the nonlinear and nonstationary nature of streamflow generation processes. GAMCR addresses these challenges by introducing a robust and flexible framework that leverages spline basis functions and Generalized Additive Models to learn the model's time-variable coefficients. Overall, GAMCR is a modern and effective tool for using the increasingly available rainfall-runoff series to investigate controls on hydrologic responses worldwide.





Code and data availability. The GAMCR code is available on the Zenodo archive: Duchemin (2025). Both the synthetic and real data are available on the Zenodo archive: Duchemin et al. (2025). All the material is published on the FAIR-compliant Zenodo repository: https://doi.org/10.5281/zenodo.15180911.

Author contributions. Q.D. developed the GAMCR model and performed the GAMCR computations. M.G.Z. processed all the observational data and performed the ERRA computations. M.G.F and H.S. gathered the real-world data. G.O. supervised the model development. P.B. and J.W.K. conceived the project and supervised this work. J.W.K. developed the ERRA method and the benchmark model. All authors discussed the results and contributed to the final manuscript.

375 Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. The authors acknowledge funding from the Swiss Data Science Center (Fifth Call for Data Science Projects, project C21-09).





References

- Arridge, S., Maass, P., Öktem, O., and Schönlieb, C.-B.: Solving inverse problems using data-driven models, Acta Numerica, 28, 1–174, 380 2019.
 - Beven, K.: How far can we go in distributed hydrological modelling?, Hydrology and Earth System Sciences, 5, 1–12, https://doi.org/10.5194/hess-5-1-2001, 2001.
 - Botter, G., Bertuzzo, E., and Rinaldo, A.: Transport in the hydrologic response: Travel time distributions, soil moisture dynamics, and the old water paradox, Water Resources Research, 46, https://doi.org/https://doi.org/10.1029/2009WR008371, 2010.
- 385 Boyd, S.: Convex optimization, Cambridge UP, 2004.
 - Bruen, M. and Dooge, J.: Unit hydrograph estimation with multiple events and prior information: I. Theory and a computer program, Hydrological Sciences Journal, 37, 429–443, 1992.
 - do Nascimento, T. V., Rudlang, J., Höge, M., van der Ent, R., Chappon, M., Seibert, J., Hrachowitz, M., and Fenicia, F.: EStreams: An integrated dataset and catalogue of streamflow, hydro-climatic and landscape variables for Europe, Scientific Data, 11, 879, 2024.
- 390 Duchemin, Q.: quentin-duchemin/GAMCR: GAMCR package, https://doi.org/10.5281/zenodo.15180816, 2025.
 - Duchemin, Q., Floriancic, M. G., Seybold, H., and Zanoni, M. G.: GAMCR paper: Data, https://doi.org/10.5281/zenodo.15180911, 2025.
 - Gao, H., Ju, Q., Zhang, D., Wang, Z., Hao, Z., and Kirchner, J. W.: Quantifying Dynamic Linkages Between Precipitation, Groundwater Recharge, and Streamflow Using Ensemble Rainfall-Runoff Analysis, Water Resources Research, 61, e2024WR037821, https://doi.org/10.1029/2024WR037821, e2024WR037821 2024WR037821, 2025.
- Hargreaves, G. H. and Samani, Z. A.: Reference crop evapotranspiration from temperature, Applied engineering in agriculture, 1, 96–99, 1985.
 - Hastie, T., Tibshirani, R., and Friedman, J.: The elements of statistical learning: data mining, inference, and prediction, 2017.
 - Kirchner, J. W.: A double paradox in catchment hydrology and geochemistry, Hydrological Processes, 17, 871–874, https://doi.org/https://doi.org/10.1002/hyp.5108, 2003.
- 400 Kirchner, J. W.: Impulse response functions for nonlinear, nonstationary, and heterogeneous systems, estimated by deconvolution and demixing of noisy time series, Sensors, 22, 3291, https://doi.org/10.3390/s22093291, 2022.
 - Kirchner, J. W.: Characterizing nonlinear, nonstationary, and heterogeneous hydrologic behavior using Ensemble Rainfall-Runoff Analysis (ERRA): proof of concept, Hydrology and Earth System Sciences Discussions, 2024, 1–42, https://doi.org/10.5194/hess-2024-103, 2024.
- Kirchner, J. W., Feng, X., and Neal, C.: Fractal stream chemistry and its implications for contaminant transport in catchments, Nature, 403, 524–527, https://doi.org/10.1038/35000537, 2000.
 - Knapp, J. L. A., Berghuijs, W. R., Floriancic, M. G., and Kirchner, J. W.: Catchment hydrological response and transport are affected differently by precipitation intensity and antecedent wetness, Hydrology and Earth System Sciences Discussions, 2024, 1–19, https://doi.org/10.5194/hess-2024-371, 2024.
- Kratzert, F., Nearing, G., Addor, N., Erickson, T., Gauch, M., Gilon, O., Gudmundsson, L., Hassidim, A., Klotz, D., Nevo, S., Shalev, G., and Matias, Y.: Caravan A global community dataset for large-sample hydrology, Scientific Data, 10, 61, https://doi.org/10.1038/s41597-023-01975-w, 2023.
 - Mays, L. W.: Water resources engineering, John Wiley & Sons, Inc., New York, 3rd edn., ISBN 978-1-119-59051-4, 2019.



415



- McDonnell, J. J. and Beven, K.: Debates—The future of hydrological sciences: A (common) path forward? A call to action aimed at understanding velocities, celerities and residence time distributions of the headwater hydrograph, Water Resources Research, 50, 5342–5350, https://doi.org/https://doi.org/10.1002/2013WR015141, 2014.
 - McDonnell, J. J., McGuire, K., Aggarwal, P., Beven, K. J., Biondi, D., Destouni, G., Dunn, S., James, A., Kirchner, J. W., Kraft, P., and others: How old is streamwater? Open questions in catchment transit time conceptualization, modelling and analysis, Hydrological Processes, 24, 1745–1754, publisher: Wiley Online Library, 2010.
- McGuire, K. J. and McDonnell, J. J.: A review and evaluation of catchment transit time modeling, Journal of Hydrology, 330, 543–563, https://doi.org/https://doi.org/10.1016/j.jhydrol.2006.04.020, 2006.
 - MeteoSwiss CombiPrecip: CombiPrecip precipitation data MeteoSwiss, https://www.meteoswiss.admin.ch/services-and-publications/service/weather-and-climate-products/combiprecip.html.
 - Rigon, R., Bancheri, M., Formetta, G., and de Lavenne, A.: The geomorphological unit hydrograph from a historical-critical perspective, Earth Surface Processes and Landforms, 41, 27–37, 2016.
- 425 Sherman, L. K.: Streamflow from rainfall by the unit-graph method, Eng. News Record, 108, 501–505, 1932.
 - Snyder, W. M.: Hydrograph analysis by the method of least square, in: Proceedings of the American Society of Civil Engineers, vol. 81, pp. 1–25, ASCE, 1955.
 - Soil Conservation Service: Section 4: Hydrology, in: SCS National Engineering Handbook, USDA, Washington D.C., 1985.
- van der Velde, Y., Torfs, P. J. J. F., van der Zee, S. E. A. T. M., and Uijlenhoet, R.: Quantifying catchment-scale mixing and its effect on time-varying travel time distributions, Water Resources Research, 48, https://doi.org/https://doi.org/10.1029/2011WR011310, 2012.