

Toward Improved Real-Time Rainfall Intensity Estimation Using Video Surveillance Cameras

Zheng, Feifei; Yin, Hang; Ma, Yiyi; Duan, Huan Feng; Gupta, Hoshin; Savic, Dragan; Kapelan, Zoran

DOI

[10.1029/2023WR034831](https://doi.org/10.1029/2023WR034831)

Publication date

2023

Document Version

Final published version

Published in

Water Resources Research

Citation (APA)

Zheng, F., Yin, H., Ma, Y., Duan, H. F., Gupta, H., Savic, D., & Kapelan, Z. (2023). Toward Improved Real-Time Rainfall Intensity Estimation Using Video Surveillance Cameras. *Water Resources Research*, 59(8), Article e2023WR034831. <https://doi.org/10.1029/2023WR034831>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Water Resources Research®

RESEARCH ARTICLE

10.1029/2023WR034831

Key Points:

- A two-stage algorithm is proposed to provide good quality rainfall intensity estimates using surveillance camera imagery
- The generalization capability of the proposed algorithm is demonstrated under different imagery conditions
- This work is an important step toward achieving real-time, high spatiotemporal resolution of urban rainfall data at low cost

Correspondence to:

Y. Ma,
yiyima@zju.edu.cn

Citation:

Zheng, F., Yin, H., Ma, Y., Duan, H.-F., Gupta, H., Savic, D., & Kapelan, Z. (2023). Toward improved real-time rainfall intensity estimation using video surveillance cameras. *Water Resources Research*, 59, e2023WR034831. <https://doi.org/10.1029/2023WR034831>

Received 6 MAR 2023
Accepted 11 AUG 2023

Author Contributions:

Conceptualization: Feifei Zheng, Hang Yin, Yiyi Ma
Funding acquisition: Feifei Zheng, Huan-Feng Duan
Investigation: Feifei Zheng, Hang Yin
Methodology: Feifei Zheng, Hang Yin, Dragan Savic, Zoran Kapelan
Supervision: Feifei Zheng, Yiyi Ma, Huan-Feng Duan, Hoshin Gupta, Dragan Savic, Zoran Kapelan
Validation: Hang Yin
Writing – original draft: Feifei Zheng, Hang Yin, Yiyi Ma
Writing – review & editing: Feifei Zheng, Yiyi Ma, Huan-Feng Duan, Hoshin Gupta, Dragan Savic, Zoran Kapelan

Toward Improved Real-Time Rainfall Intensity Estimation Using Video Surveillance Cameras

Feifei Zheng¹ , Hang Yin¹, Yiyi Ma¹ , Huan-Feng Duan² , Hoshin Gupta³, Dragan Savic^{4,5,6} , and Zoran Kapelan^{5,7} 

¹College of Civil Engineering and Architecture, Zhejiang University, Hangzhou, China, ²Department of Civil and Environmental Engineering, The Hong Kong Polytechnic University, Kowloon, Hong Kong, ³Department of Hydrology and Atmospheric Sciences, The University of Arizona, Tucson, AZ, USA, ⁴KWR Water Research Institute, Nieuwegein, The Netherlands, ⁵Centre for Water Systems, University of Exeter, Exeter, UK, ⁶Faculty of Engineering and Built Environment, Universiti Kebangsaan Malaysia, Bangi, Malaysia, ⁷Department of Water Management, Delft University of Technology, Nieuwegein, The Netherlands

Abstract Under global climate change, urban flooding occurs frequently, leading to huge economic losses and human casualties. Extreme rainfall is one of the direct and key causes of urban flooding, and accurate rainfall estimates at high spatiotemporal resolution are of great significance for real-time urban flood forecasting. Using existing rainfall intensity measurement technologies, including ground rainfall gauges, ground-based radar, and satellite remote sensing, it is challenging to obtain estimates of the desired quality and resolution. However, an approach based on processing distributed surveillance camera network imagery through machine learning algorithms to estimate rainfall intensities shows considerable promise. Here, we present a novel approach that first extracts raindrop information from the surveillance camera images (rather than using the raw imagery directly), followed by the use of convolutional neural networks to estimate rainfall intensity from the resulting raindrop information. Evaluation of the approach on 12 rainfall events under both daytime and nighttime conditions shows that generalization ability, and especially nighttime predictive performance, is significantly improved. This represents an important step toward achieving real-time, high spatiotemporal resolution, measurement of urban rainfall at relatively low cost.

1. Introduction

The human costs of flooding are huge, resulting in about 220,000 deaths worldwide and cumulative direct economic losses estimated at about US\$1 trillion, for the period 1980 to 2014 (Winsemius et al., 2016). Focusing just on urban flooding, during 2004–2014 the United States alone experienced direct economic losses of about US\$9 billion per year (National Academies of Sciences and Medicine, 2019). Similarly, the Ministry of Water Resources (China.com, 2019) reported an average annual economic loss of RMB 229.87 billion (during 2013–2018) due to floods across China with, on average, 153 cities suffering from urban flooding each year.

Urban flooding is a consequence of both natural and social factors (Zheng et al., 2016), including extreme rainfall, urban topography and terrain, insufficient flow capacity of existing urban drainage systems (Vojinovic et al., 2014), increase in extent of impermeable surfaces (Lin et al., 2019), and river floods (O'Donnell & Thorne, 2020). Of these, the dominant factor is rainfall, and especially high-intensity rainstorms of short-duration, which are becoming increasingly more frequent due to global climate change (Field et al., 2012). Uncertainty about the spatial and temporal distribution of rainfall can severely challenge our ability to predict the severity and extent of urban flooding (Berne & Krajewski, 2013). Therefore, access to accurate real-time estimates of rainfall, with high spatial and temporal resolution, is particularly important for short lead-time forecasting of urban hydrological processes, as well as for real-time monitoring and warning to mitigate the effects of urban flooding.

At present, the main strategies commonly used to monitor urban rainfall are rainfall gauges, ground-based radar and satellite-based remote sensing. While interpolation between rainfall gauges can provide estimates of the spatial distribution of urban rainfall, the availability of rain gauges in urban areas is usually insufficient. Ground-based radar can provide local rainfall estimates with fairly high spatiotemporal resolution (about 1 × 1 km) (Zhou et al., 2019), but accuracy can be compromised by numerous factors including uneven vertical distributions of rainfall intensity, abnormal propagation of electromagnetic waves, ground clutter, blockage due to high-rise buildings, and variations in the surface wind field (Qin et al., 2014). Satellite-based remote sensing

can provide much larger spatial coverage (Long et al., 2015; Yong et al., 2012) but often at lower spatiotemporal resolution (about 20×20 km) (Zorretto & Marani, 2019). Therefore, using such methods, it is difficult to obtain real-time rainfall estimates with a sufficiently fine spatial resolution, ideally on the order about 10×10 m (Gires et al., 2012).

In recent years, some novel approaches to estimating rainfall have been proposed. One promising approach is to exploit the potential power of crowdsourcing (Zheng et al., 2018) and the Internet of Things (IoT), which enables decentralization of the data collection mission by use of distributed sensor technologies. For example, the sensors used for rainfall monitoring can include smart wipers (Rabiei et al., 2013), intelligent umbrellas (Hut et al., 2014), optical rainfall intensity sensors, and microwave receivers (Uijlenhoet et al., 2018). However, this approach relies upon the availability of large numbers of sensors, and can therefore be costly and complicated to implement.

Alternatively, with the rapid pace of urbanization, the density of networks of surveillance cameras also increases, providing a potentially more viable economic solution to the distributed sensor problem since such networks are already in place (Jiang et al., 2019; Yin et al., 2023). The idea is, that rainfall images captured by surveillance cameras can be processed through machine learning algorithms in order to estimate rainfall intensities. For example, Yin et al. (2023) trained a Convolutional Neural Network (CNN), named the *irCNN*, to estimate rainfall from daytime rainfall imagery. However, reported limitations of that approach include: (a) the need for large amounts of training data, and (b) poor generalization ability to different conditions (and especially nighttime).

A possible solution to the above-mentioned problems is to first extract *raindrop* information from the rainfall images, and then use this more detailed information as input into a CNN that estimates the rainfall intensity. This task is (of course) not simple and algorithms that have been designed to extract raindrops information from images suffer from various kinds of limitations. Much of the relevant research is based on the premise that raindrops in images are noise that affects the image quality and must therefore be removed (Wang et al., 2019). Algorithms have been developed to remove rain from videos (Garg & Nayar, 2007) and still images (Xu et al., 2012) based on recognizing their characteristic properties (e.g., color, texture, etc.) (Zhang et al., 2006), and on detecting the changes induced between adjacent video frames (Garg & Nayar, 2007). But when doing so, other random movements such as oscillations of branches and leaves, and changes in brightness, also tend to get removed. Garg and Nayar (2007) proposed a method for extracting raindrops from videos by assuming that the background remains still between consecutive video frames and identifying the changes in gray scales of three consecutive images. However, due to the oversimplified nature of their algorithm, other kinds of noise were extracted as well. More recently, Jiang et al. (2019) attempted to extract rainfall drop information from surveillance camera images by means of a convex optimization algorithm that exploits the statistical differences between real and synthetic rainfall images. A challenge associated with their approach is that real raindrops are considerably more complex and variable than the synthetic ones used; the method fails under the high-wind speed conditions that can accompany high-intensity rainfall events.

In this work, we propose and test a new two-stage algorithm that estimates rainfall intensity from rainfall images. In the first stage, instead of directly processing the rainfall images through a CNN, the image is first divided into three layers, namely background, rainfall and noise, from which the raindrop image information is extracted based on the different characteristics of each layer. In the second stage, an *irCNN* model is trained to estimate rainfall intensity using the extracted raindrop image information as input data. We show that the proposed two-stage algorithm enhances the generalization ability of the model, and especially its predictive performance under nighttime conditions.

2. Methodology

2.1. Model Architecture

Our proposed two-stage algorithm for predicting rainfall intensity from video imagery is shown in Figure 1. As illustrated, each video frame is divided into a background layer, rainfall layer, and a noise layer. The background layer represents the main body of the image, which is assumed to remain unchanged in consecutive frames of the video. The rainfall layer represents the raindrops, which are characterized by frame-to-frame changes and spatial sparsity. The noise layer is characterized by randomness, and occupies a relatively small proportion of the signal compared to the other two layers. Based on the characteristic features of each layer, a function for extracting raindrop information is designed. The *irCNN* model architecture is then used to perform the rainfall intensity




Figure 1. Two-stage algorithm for rainfall intensity estimation.

calculation, using the extracted raindrop information as inputs. Two evaluation methods are utilized to assess the effectiveness of the proposed approach under different conditions. In the following sections, the raindrop information extraction algorithm and the method for estimating rainfall intensity are described.

2.2. Raindrop Extraction Algorithm

Conventional algorithms for removing raindrops from imagery often do so along with removal of other kinds of “random” noise (e.g., changes in background, etc.). In this study we develop a new raindrop extraction algorithm for obtaining accurate images of raindrops, with the key novelty being the partitioning of the image into three layers, consisting of background, rainfall, and noise, which facilitates raindrop extraction based on different characteristic features in each of the layers.

Consider the source grayscale image to be a two-dimensional matrix $X \in R^{h \times w}$ (h, w = height and width of image), which can be expressed in terms of the decomposition:

$$X = B + R + N \quad (1)$$

where $B \in R^{h \times w}$, $R \in R^{h \times w}$, and $N \in R^{h \times w}$ represent the background layer, rainfall layer and noise layer, respectively, as shown in Figure 1. Assuming that, except for changes in brightness and due to moving objects, the background changes are negligible between consecutive frames, a low-rank matrix decomposition (Li et al., 2018) can be used to approximate the rainfall image as follows. Based on the singular value decomposition theorem (SVD theorem), any arbitrary real matrix $A \in R^{m \times n}$ can be decomposed as:

$$A = C \Gamma D^T \quad (2)$$

where $C \in R^{m \times m}$, $\Gamma \in R^{m \times n}$, $D \in R^{n \times n}$, such that Γ is a diagonal matrix whose diagonal elements Γ_{ii} ($0 \leq i \leq \min(m, n)$) are the singular values of the matrix A , the column vectors of C are the left singular vectors, and the column

vectors of D are the right singular vectors. By means of a low-rank matrix decomposition, the matrix A can be approximated by choosing $k < \text{rank}(A)$, such that,

$$A^* \approx A = E\Lambda F^T \quad (3)$$

where $E \in R^{m \times k}$, $\Lambda \in R^{k \times k}$, and $F \in R^{n \times k}$, such that E and F consist of the first k columns of singular vectors in C and D , respectively, and Λ is the diagonal matrix obtained after preserving the first k singular values.

Now, each image from the video X ($X \in R^{h \times w}$) can be transformed into a one-dimensional vector x ($x \in R^{hw}$) by means of an *Unfold* operation, by taking the first column of X , followed by the second column, and then, the third column, etc. Next, the transformed one-dimensional vectors for each successive image can be arranged in time order to form a new two-dimensional matrix $X' \in R^{hw \times t}$. Using the low-rank matrix decomposition described above, and choosing $r < \text{rank}(X')$, we can obtain:

$$X'^* \approx X' = U\Sigma V \quad (4)$$

where $U \in R^{hw \times r}$, $\Sigma \in R^{r \times r}$, and $V \in R^{n \times r}$. The background is the main part of the rainfall image, which can be estimated using the singular value decomposition theorem (SVD) as did in Li et al. (2018). More specifically, a *Fold* operation (the inverse operation of the *Unfold* operation) is used to convert the 2D matrix X'^* to a 3D matrix B , which is expressed as:

$$B \approx \text{Fold}(X'^*) \approx \text{Fold}(U\Sigma V) \quad (5)$$

where $B \in R^{h \times w \times t}$. To ensure that the estimate B is a good approximation to (close to) $\text{Fold}(U\Sigma V)$, we can minimize the Frobenius norm of the difference expressed as:

$$\min \|B - \text{Fold}(U\Sigma V)\|_F \quad (6)$$

The next step is to characterize the rainfall layer R . To estimate R , we take advantage of the fact that the sparsity of raindrops is the major feature that distinguishes them from other objects in an image. We assume that we can represent the rainfall layer as $R \approx HR \circ X$, where $\circ$ represents the Hadamard product (also known as the element-wise product, namely, the product of two matrixes element by element). Now, we can use Markov random fields (Wei et al., 2017) to represent the positions of raindrops, by creating a matrix $HR_{ijk} \in R^{h \times w \times t}$ whose elements consist of:

$$HR_{ijk} = \begin{cases} 1, & \text{if the position is a raindrop} \\ 0, & \text{if the position is not a raindrop} \end{cases} \quad (7)$$

where to ensure sparsity of the matrix, we can minimize the L1 norm (e.g., $\|A\|_1 = \sum_i \sum_j |a_{ij}|$) expressed as (Li & Zhu, 2008):

$$\min \|HR\|_1 \quad (8)$$

Finally, we need to characterize the noise layer N , which represents random perturbations in the image due to wind and changes in brightness etc. From Equation 1 we can write:

$$N = X - B - R = X - B - HR \circ X = (1 - HR) \circ X - B \quad (9)$$

To control the degree of the noise to be small enough for successful image decomposition, we can minimize the Frobenius norm of N , expressed as:

$$\min \|(1 - HR) \circ X - B\|_F \quad (10)$$

Finally, we combine Equations 6, 8 and 10 to express the desired properties of all three layers B , R , and N , by defining the objective function:

$$\min (\lambda_1 \|B - \text{Fold}(U\Sigma V)\|_F + \lambda_2 \|HR\|_1 + \lambda_3 \|(1 - HR) \circ X - B\|_F) \quad (11)$$

where, λ_1 , λ_2 , and λ_3 represent weighting parameters that must be specified. Of course, the specific values chosen for λ_1 , λ_2 , and λ_3 will affect the raindrop extraction results.

We note that small variations of the weighting parameter values may produce systematic errors in estimating raindrops. However, these errors can be accounted for in the second stage of the entire method (see Section 2.3), wherein a deep learning algorithm is implemented. This is because the strong fitting ability of the deep learning algorithm enables it to learn and compensate for such systematic patterns in the errors, making the entire approach robust to small variations in the parameters. In this study, we fixed the values of λ_1 , λ_2 , and λ_3 at 1, 0.1, and 1, respectively.

Our task now is to find the values of r and $\mathbf{HR}$ that give rise to the matrices $\mathbf{B} = \text{Fold}(\mathbf{X}'^*)$, $\mathbf{R} \approx \mathbf{HR} \cdot \mathbf{X}$ and $\mathbf{N} = (\mathbf{I} - \mathbf{HR}) \cdot \mathbf{X} - \mathbf{B}$ that minimize the objective function defined by Equation 11. Since the matrix to be optimized consists of hundreds of thousands of elements, it is necessary to use a sparse optimization algorithm, to transform the nonconvex function into a convex function and then solve it by stepwise iteration. Here, we use the Alternating Direction Method of Multipliers (Boyd et al., 2011). Noting that the SVD method is able to identify raindrops from the background noise, which can include the moving cars/people or trees; this has been further verified in this study, with details given in Section 4.1. Further, the folding and unfolding operations are widely used in computer science for image processing (Li et al., 2018) and hence these are adopted in this study.

2.3. Rain Intensity Estimation

After extracting raindrop information from the rainfall images, we use it as input data when training the irCNN machine learning model (Yin et al., 2023) to estimate rainfall intensity. This works, because rainfall intensity is strongly related to the density and size of raindrops in the image, expressed mathematically as:

$$I = f(Z(d, s)) \quad (12)$$

where, I = rainfall intensity (mm/h), Z = rainfall image, d = the number of raindrops, s = raindrop size, and f = underlying nonlinear relationship between rainfall image and rainfall intensity. Noting that Equation 12 used for the deep learning algorithm may be different to the integration equation as the deep learning method exploits information about the number and the size of the raindrops for model training (Yin et al., 2023).

The irCNN model is based on the ResNet34 model that has been widely used in computer science and engineering (He et al., 2016). The irCNN model structure is different from that of the original CNN, in that a regression layer is included for rainfall intensity estimation:

$$\hat{I} = \mathbf{W}^T \mathbf{X} + b \quad (13)$$

where $\hat{I}$ = estimated rainfall intensity, $\mathbf{W}$ = parameter matrix (512×1) of the linear regression layer, $\mathbf{X}$ = output of the previous layer, T = transpose operation, and b = bias term.

Due to the large number of irCNN model parameters that must be trained, a correspondingly large number of rainfall images will be needed if training the model completely from scratch. To circumvent this problem, we start with an existing irCNN model (Deng et al., 2009) that was (pre)trained using an open-source ImageNet data set containing 1.28 million images of 1,000 object categories that are not necessarily related to rainfall. The pre-trained model was then fine-tuned using the rainfall information extracted in stage one (describe above) to determine the final irCNN model parameter values. This pre-training approach is commonly used for CNN models since images about different objects share many common features that are relevant to their detection and classification.

3. Data Description and Performance Metrics

For this study, we analyzed 12 rainfall events captured between June and July 2020. Six of these events (Event 1–6) occurred during the daytime and six at night (Event 7–12) as listed in Table 1. Imagery for these events was

Table 1
Details of the 12 Rainfall Events, Where Events 1–6 Occurred in Daytime and Events 7–12 Occurred in the Nighttime

Rainfall events	Date	Duration (min)	Average rainfall intensity (mm/h)	Maximum rainfall intensity (mm/h)
1	21/06/2020	17	16.9	42.0
2	21/06/2020	69	19.0	66.0
3	26/06/2020	21	22.3	60.0
4	10/07/2020	12	11.0	36.0
5	16/07/2020	18	13.7	42.0
6	26/07/2020	33	23.6	60.0
7	29/06/2020	21	27.4	96.0
8	29/06/2020	18	9.7	24.0
9	30/06/2020	39	22.2	60.0
10	02/07/2020	55	39.3	156.0
11	03/07/2020	29	11.2	30.0
12	05/07/2020	25	10.1	24.0

captured using a fixed camera with a shutter speed of 1 frame per second (fps). Rainfall intensity was recorded with a time resolution of 1 min and an intensity resolution of 0.1 mm/min (6 mm/hr), using a tipping-bucket rain gauge installed near the camera (within 1–2 m). Although the sampling error of tipping bucket gauges across different temporal scales (especially at the one-minute scale) can be large (Ciach, 2003), this does not affect the applicability of the proposed method, since accurate rainfall records can be easily incorporated for actual model implementation.

3.1. Determination of Instantaneous Rainfall Intensity

The exposure time of a single video frame taken once per second by the camera, was about 1/200 s. However, the data recorded by the rain gauge was cumulative rainfall volume over a period of 1 min. To address this mismatch, we assumed a linear variation of rainfall intensity within each one-minute interval. Since the change in rainfall intensity during any one-minute period is typically relatively small, this assumption seems reasonable, and enables us to use linear interpolation to estimate the instantaneous rainfall intensity from the rain gauge records. By further assuming that the value at the midpoint of each recording time interval of the rain gauge is exactly the rainfall intensity for that time period, we can compute the instantaneous rainfall intensity I_t at any time t during that time interval as:

$$I_t = I_L + \frac{t - t_L}{T}(I_R - I_L) \quad (14)$$

where T is the recording interval of the rain gauge ($T = 1$ min in this study), and I_L and I_R are the rainfall intensities at the moments t_L and t_R , respectively.

3.2. Model Performance Metrics

We computed several metrics that are commonly used to measure model predictive performance (here the estimated rainfall intensity). These include the standard statistical indices of mean absolute error (MAE), mean absolute percentage error (MAPE), and coefficient of determination (R^2), and two metrics that are widely applied in hydrology—the Nash-Sutcliffe model efficiency (NSE) and the Kling-Gupta efficiency (KGE). These metrics are selected as they represent different aspects of the model performance, where MAE and MAPE respectively shows the model's errors (absolute and percentage) relative to observations, R^2 focuses on the overall correlation between simulations and observations, NSE indicates the model's ability in fitting the observations and KGE simultaneously considers the bias, variance and correlation. The expressions for MAE and MAPE are:

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |Y_i - \hat{Y}_i| \quad (15)$$

$$\text{MAPE} = \frac{1}{n} \sum_{i=1}^n \left| \frac{Y_i - \hat{Y}_i}{Y_i} \right| \quad (16)$$

where n is the total number of data points, Y_i is the i th observation, and $\hat{Y}_i$ is the i th prediction. Lower values of MAE and MAPE indicate better model performance. The NSE, and KGE measures of goodness of fit are:

$$\text{NSE} = 1 - \frac{\sum_{i=1}^n (Y_i - \hat{Y}_i)^2}{\sum_{i=1}^n (Y_i - \bar{Y})^2} \quad (17)$$

$$\text{KGE} = 1 - \sqrt{(r - 1)^2 + \left(\frac{\sigma_{\text{pred}}}{\sigma_{\text{obs}}} - 1 \right)^2 + \left(\frac{\mu_{\text{pred}}}{\mu_{\text{obs}}} - 1 \right)^2} \quad (18)$$

where $\bar{Y}$ is the average value of the observations, r represents the linear correlation between the observations and predictions, σ_{pred} and μ_{pred} are the standard error and average value of the predictions, σ_{obs} and μ_{obs} are the standard




Figure 2. Sample images, illustrating the results of raindrop extraction under different scenes: Image (a) is a daytime scene taken from Jiang et al. (2019), while image (c) shows the corresponding extracted raindrop information. Image (c) shows a typical night scene, while image (d) shows the corresponding extracted raindrop information.

error and average value of the observations. Larger values of R^2 , NSE, and KGE indicate a better fit of the model. The model fits all data perfectly when R^2 , NSE, or KGE equals 1.

4. Results and Discussion

4.1. Raindrop Extraction

Sample images of raindrop information extracted by the proposed algorithm under two different conditions (daytime and nighttime) are shown in Figure 2. These represent an illustration of typical results obtained, and further quantitative analysis of the method's performance is presented below. An original daytime rainfall image, taken from Jiang et al. (2019), is shown in Figure 2a and the corresponding extracted raindrop image in Figure 2b. A nighttime rainfall image, and the corresponding extracted raindrop image are shown in Figures 2c and 2d, respectively. The figure demonstrates that the raindrop extraction algorithm works well under both daytime and nighttime conditions.

Computational efficiency of the raindrop extraction algorithm was improved by using a GPU. The specific hardware configuration used was a PC with Intel Core i9-9820X @ 3.3 GHz processor, 32 GB RAM, and the GPU was NVIDIA RTX 2080Ti 11 GB GPU. Computational time to process 100 rainfall video frames at $1,920 \times 1,080$ pixels resolution was about 24 s. In comparison, the computational time of the algorithm proposed by Jiang et al. (2019) to process 100 frames of rainfall images at much lower 320×480 pixels resolution was about 26.4 s. Overall, the speed of the proposed raindrop extraction algorithm was ~ 18 times faster than that reported by Jiang et al. (2019) to process the same rainfall video. The reason for this improved efficiency is that the new raindrop extraction algorithm has a simpler structure and avoids the need to solve complex partial differential equations.

4.2. Evaluation of Results by Random Sampling

To assess the performance of the proposed two-stage algorithm, and to compare with the previous algorithm of Yin et al. (2023) that directly applied irCNN model to the unprocessed rainfall images, six different scenarios were created as shown in Table 2, using the images from the 12 rainfall events listed in Table 1. The groups labeled "Original" use images directly obtained from the camera videos, while those labeled "Raindrop" use the raindrop images extracted by the proposed algorithm.

Table 2
Scenarios of Randomly Selected Images

Case no.	Rainfall scenario	Image type	Training set (80% data)	Validation set (20% data)
1	Daytime	Original	3,784	946
2	Daytime	Raindrop	3,784	946
3	Nighttime	Original	4,820	1,205
4	Nighttime	Raindrop	4,820	1,205
5	Daytime & nighttime	Original	8,604	2,151
6	Daytime & nighttime	Raindrop	8,604	2,151

For example, in Case 1, original daytime rainfall images were used from which 3784 (80%) were randomly selected for training and 946 (20%) utilized for evaluation (see results in Table 3). To minimize estimation uncertainties, each evaluation case was repeated five times, and Table 3 reports the average performance obtained over the five runs. Note that the irCNN model performance reported in Table 3 differs from that presented in Yin et al. (2023) which used only daytime rainfall images, whereas the results reported also include nighttime images.

In Cases 1 and 2, only daytime imagery was used to train the irCNN model. From Table 3, we see that, for most of the performance metrics, Case 1 (using original imagery) obtains slightly better values than Case 2 (using extracted raindrop information). More specifically, the MAE and MAPE for Case 1 are 3.29 mm/hr and 17.06% respectively, which is lower than the corresponding

Table 3
Average Performance Metric Values When Evaluated Using Random Sampling

Case no.	MAE (mm/h)	MAPE (%)	R^2	NSE	KGE
1	3.29	17.06	0.91	0.89	0.86
2	4.07	19.58	0.82	0.80	0.89
3	6.60	28.35	0.87	0.87	0.93
4	4.76	19.73	0.91	0.91	0.96
5	7.54	29.50	0.80	0.73	0.77
6	5.77	22.69	0.87	0.85	0.87

values obtained in Case 2; the R^2 and NSE for Case 1 are 0.91 and 0.89 respectively which is overall larger than the corresponding values obtained in Case 2. From Figure 3, it can be seen that, for the daytime scenes, the irCNN model shows better performance in estimating rainfall intensity directly from the original rainfall images than using the raindrop information. This is because the extra step of first extracting raindrop information from the images can increase the error rate. Nonetheless, the MAPE performance of Case 2 still remains within 20%.

In Cases 3 and 4, only nighttime imagery was used to train the irCNN model. From Table 3, we see the reverse of the daytime imagery results, with Case 4 (using extracted raindrop information) obtaining better values for all the performance metrics compared with Case 3 (using original imagery).

For example, the MAE in Case 3 was 6.60 mm/hr, which is significantly higher than in Case 4 with an MAE of 4.76 mm/hr. This is because nighttime imaging relies mainly on the raindrops reflecting other (chaotic) light sources, which makes it more difficult to detect and capture the raindrops clearly. However, at night, the extraction algorithm does not have to contend with ambient environmental brightness. The net effect is that, although overall performance declines at night, the two-stage method achieves better overall performance than the original (Yin et al., 2023) methodology.

Finally, Cases 5 and 6 used images randomly selected from both daytime and nighttime. From Table 3, we see that Case 6 (using extracted raindrop information) achieves significant performance improvement over all performance metrics compared with Case 5 (using original imagery). In terms of NSE and KGE, Case 6 values are 0.85 and 0.87 respectively, which is all significantly greater than the corresponding values in Case 5. When exposed to both daytime and nighttime imagery the estimation problem is much more complicated, due to changes in the background and brightness of the images. Overall, our results indicate that the two-stage method is more robust (than when video imagery is used directly) when dealing with complex rainfall scenarios.

Because urban flooding is often caused by short-duration high-intensity rainstorms, we next assess whether the model can predict intense rainfall events accurately. For each of the Cases 1–6 in Table 2, we further classified the events as being light rainfall (<12 mm/hr), moderate rainfall (12–17.9 mm/hr), heavy rainfall (18–23.9 mm/hr), rainstorm (24–35.9 mm/hr), intense rainstorm (36–60 mm/hr), and extraordinary rainstorm (>60 mm/hr). The results are presented in Table 4 with for different categories of rainfall intensities.

As can be seen from Table 4, for daytime imagery (Cases 1 and 2), both methods perform quite well across the entire range with the original approach (Case 1) being better in the intermediate intensity categories but the two-stage approach (Case 2) performing better for the “extraordinary rainstorm” category. For nighttime imagery




Figure 3. Validation results of irCNN models by random sampling.

Table 4
Average Values of Mean Absolute Percentage Error (%) When Validated Using Random Sampling for Different Categories of Rainfall (Case Number is Given in Table 2)

Case No.	Light rainfall (%)	Moderate rainfall (%)	Heavy rainfall (%)	Rainstorm (%)	Intense rainstorm (%)	Extraordinary rainstorm (%)
1	25.31	15.22	16.59	15.03	18.42	16.20
2	24.73	15.29	22.34	19.52	18.70	11.44
3	36.49	32.04	20.65	27.84	22.68	18.09
4	27.57	16.53	13.79	19.11	17.11	15.27
5	38.00	16.60	25.01	31.10	35.41	34.64
6	27.01	19.13	22.96	20.03	24.00	22.16

(Cases 3 and 4), the two-stage approach (Case 4) is clearly better across the entire range of intensities. When using mixed day and night time imagery (Cases 5 and 6), the two-stage approach is significantly better for both very light and higher intensity rain events and both methods are comparable in the intermediate range. In particular, the original method performs noticeably poorly at both extremes (light and intense/extraordinary). Further, for the two-stage approach, there is a generally flat-to-improving trend from light to higher intensity rainfall cases, and much less variability between daytime, nighttime and mixed cases.

4.3. Evaluation on Independent Rainfall Events

To investigate the predictive performance of the two-stage algorithm for new rainfall events, several independent rainfall events under both daytime and nighttime conditions were selected from the 12 rainfall events in Table 1 for evaluation. When dividing rainfall events into the training set and evaluation set, the following conditions were used: (a) the rainfall intensity of the independent rainfall events for evaluation is medium compared to other events, and (b) the duration of the rainfall events used for training is sufficiently long to provide adequate data (Yin et al., 2023). Thus, rainfall event 1 in Table 1 was selected as an independent evaluation event for the daytime scenario, and event 11 was selected for the nighttime scenario. A total of 8 cases for independent events evaluation were created (see details in Table 5). For example, in Case 7, model training was conducted using the original daytime images from events 2–6 (see Table 1) those from event 1 were used for evaluation.

The evaluation results are shown in Table 6. As before, each evaluation test was repeated five times, and average performance values are reported. As can be seen, the evaluation results are similar to those reported in Table 3. More specifically, when comparing Cases 7 and 8, the evaluation results using original daytime rainfall images (Case 7) are slightly better than the corresponding results obtained using raindrop extraction (Case 8), while both cases are slightly worse than those reported in Table 3. Similar results can be observed for the nighttime scenes. Again, as shown in Figure 4 for the daytime imagery, raindrop extraction (Case 8) results in worse performance than the original method (Case 7).

Table 5
Details of the Selected Independent Evaluation Events

Case No.	Rainfall scenario	Image type	Training set (no. of images)	Validation set (no. of images)
7	Daytime	Original	Event 2–6 (4,058)	Event 1 (672)
8	Daytime	Raindrop	Event 2–6 (4,058)	Event 1 (672)
9	Nighttime	Original	Event 7–10 and 12 (5,505)	Event 11 (520)
10	Nighttime	Raindrop	Event 7–10 and 12 (5,505)	Event 11 (520)
11	Daytime scene for training, nighttime scene for evaluation	Original	Event 1–6 (4,730)	Event 11 (520)
12	Daytime scene for training, nighttime scene for evaluation	Raindrop	Event 1–6 (4,730)	Event 11 (520)
13	Daytime scene for evaluation, nighttime scene for training	Original	Event 7–12 (6,025)	Event 1 (672)
14	Daytime scene for evaluation, nighttime scene for training	Raindrop	Event 7–12 (6,025)	Event 1 (672)

Table 6
Average Performance Metric Values When Validated by Independent Rainfall Events

Case No.	MAE (mm/h)	MAPE (%)	R^2	NSE	KGE
7	2.56	18.64	0.93	0.92	0.90
8	3.70	20.95	0.81	0.80	0.85
9	3.52	26.24	0.64	0.63	0.70
10	2.98	20.88	0.81	0.77	0.82
11	Fail	—	—	—	—
12	4.19	28.74	0.46	0.44	0.62
13	Fail	—	—	—	—
14	5.39	32.89	0.71	0.64	0.74

Looking at the nighttime imagery results (Cases 9 and 10), we see that Case 9 (using original rainfall imagery) performs consistently worse than Case 10 (using extracted raindrop information). Figure 5, shows that the scatterplot for Case 9 is more dispersed than for Case 10.

To further investigate generalization ability of the two-stage algorithm, we also examined Cases 11 and 12, where daytime imagery was used for training but nighttime imagery was used for evaluation, and Cases 13 and 14, where nighttime imagery was used for training but daytime imagery was used for evaluation. From Figure 6, we see that the original approach (cases 11 and 13) fails when using models trained on daytime (night-time) images but tested under nighttime (daytime) conditions (Figures 6a and 6c). In contrast, the two-stage approach continues to function, demonstrating some degree of ability to generalize across conditions (Figures 6b and 6d).

Finally, Cases 7–10 were further classified into rainfall intensities of various degrees, and the MAPE evaluation results are listed in Table 7. As can be seen from this table, model predictions for relatively small rainfall intensities (e.g., light rainfall and moderate rainfall) exhibit lower accuracy than those for larger intensities (e.g., heavy rain, rainstorm, intense rainstorm, and extraordinary rainstorm). Further, the two-stage algorithm (Cases 8 and 10) achieves generally better results under larger rainfall intensity cases.

4.4. Practical Application

Although to some extent the mismatch of rainfall intensity information provided by the rain gauge records and the instantaneous rainfall images can be addressed by the linear interpolation method, the real situation can be more complex. To examine this situation in more depth, the rain gauge records at 1 min intervals were interpolated once, then compared with the average rainfall intensities computed from the corresponding 30-s-video (30 frames). The results, for daytime and nighttime conditions are shown in Figure 7, where the histograms represent the average rainfall intensity computed using the two-stage algorithm, the red dots are the values obtained using gauges every 1 minute, and the blue dots indicate the interpolated values. As can be seen, the two-stage algorithm exhibits a high degree of accuracy (small errors) at most of the time steps.

Finally, the average rainfall intensity was computed using the two-stage algorithm for all of the video frames within a 1 min interval, for Cases 7–10, and compared with the rainfall intensity measurements provided by the rain gauge (at 1-min temporal resolution). The results are shown in Table 8. We can compare these




Figure 4. Evaluation results for daytime scenes in (a) Case 7 using original rainfall imagery, and (b) Case 8 using extracted raindrop information.




Figure 5. Evaluation results for nighttime scenes in (a) Case 9 using original rainfall imagery, and (b) Case 10 using extracted raindrop information.

interval averaged results with those provided in Table 6 for instantaneous rainfall intensity and in Table 8 for average rainfall intensity within 1 min), for Cases 7–10. As can be seen the estimation is further reduced when we consider the average rain intensity in a 1 min period. For the two-stage approach, the 1-min average MAPE of rainfall intensity under daytime and nighttime conditions is 10.86% and 10.62%, respectively (significantly reduced compared to the corresponding results in Table 6). So, the use of 1-min average rainfall intensity rather than instantaneous rainfall intensity, can provide a viable approach for many practical applications.

5. Conclusions

Real-time rainfall intensity can be difficult to ascertain from video imagery due to the confounding influences of the background, sources of noise, and variations in brightness. The proposed two-stage algorithm for rainfall intensity estimation from surveillance images works by first extracting raindrop information from the surveillance images, and then using a deep learning irCNN method to estimate the rainfall intensity using the extracted raindrop information as input data.

We tested the proposed two-stage algorithm using video imagery collected using real-life surveillance cameras in Hangzhou, China, and compared it with the original single-stage approach reported by Yin et al. (2023). Based on the results obtained, our main findings can be summarized as:

1. The two-stage rainfall estimation algorithm can provide good quality rainfall intensity estimates using surveillance camera imagery.
2. The two-stage model has improved prediction and generalization capability when compared to the original single-stage approach, particularly under nighttime conditions, or when a mix of daytime and nighttime images is used for training.
3. In general, nighttime performance tends to be worse than daytime performance due to the fact that raindrops are usually easier to detect in daytime images (at night the light is dimmer so that other sources of noise can affect predictive performance). However, under these conditions the two-stage approach becomes especially beneficial.
4. By aggregating the rainfall intensity estimates up to a 1-min resolution (as opposed to 1-s resolution), the MAPE prediction error of the two-stage algorithm can be cut in half (from ~20% to ~11%). These 1-min aggregate estimates should be suitable for many practical applications, such as urban flood modeling and prediction.

From the results obtained, it is clear that the two-stage approach proposed in this paper can effectively improve the generalization ability of the irCNN model, especially in complex rainfall scenes with insufficient




Figure 6. Evaluation results for (a) Case 11 with original rainfall images, (b) Case 12 with extracted raindrop information, (c) Case 13 with original rainfall images, and (d) Case 14 with extracted raindrop information.

Table 7
Average Values of Mean Absolute Percentage Error When Validated Using Independent Rainfall Events of Various Rainfall Categories

Case No.	Light rainfall (%)	Moderate rainfall (%)	Heavy rainfall (%)	Rainstorm (%)	Intense rainstorm (%)
7	28.28	22.95	13.96	12.06	11.25
8	21.97	23.48	18.44	19.69	13.74
9	37.83	17.41	20.96	19.98	—
10	29.21	16.96	19.10	13.88	—

illumination (such as nighttime scenes). This is important for the prevention and control of urban flooding, especially when heavy rainfall events and associated flooding events take place during nighttime periods. Another merit of the proposed method is that it can use existing cameras for rainfall intensity estimation without the need for additional sensors, and hence can provide urban rainfall data at low cost. It is noted that some potential limitations of the proposed two-stage model exist for practical applications, such as the mismatch of rainfall data recorded by rainfall gauges (cumulative rainfall data) and rainfall images (instantaneous rainfall intensity). This is due to difference in measurement frequency, and the errors caused by the raindrop extraction algorithm, both of which require further investigations. While practical implementation of the irCNN model may require the availability of a large amount of training data, this can be obtained from traffic cameras that are widely distributed in urbanized regions. However,




Figure 7. Independent rainfall events validation results of two-stage algorithm applied to 30-s resolution for (a) daytime scenes, and (b) nighttime scenes.

Table 8
Average Performance Metric Values of the Two-Stage Algorithm Using Independent Rainfall Events of 1-Min Resolution

Case No.	MAE (mm/h)	MAPE (%)	R^2	NSE	KGE
7	2.14	10.40	0.98	0.96	0.83
8	2.85	10.86	0.95	0.90	0.84
9	3.82	18.03	0.67	0.62	0.62
10	2.54	10.62	0.88	0.73	0.79

the data privacy issue needs to be carefully addressed when accessing data from public sources. Further, while model performance can be affected by lighting, background motion and imaging parameters, proper training of the raindrop extraction algorithm can effectively reduce such impacts.

Finally, the use of the video data (sequences of images) rather than individual still images can be a way to further improve the model's performance, and this should be explored in future studies. In this regards, while this study aimed to demonstrate the utility of the proposed method using data from a single camera, future work should explore the performance of the proposed model using information provided by networks of cameras.

Data Availability Statement

The rainfall event videos and the measured rainfall data can be downloaded from <https://doi.org/10.6084/m9.figshare.22122500.v1>.

Acknowledgments

This study was supported by the NSFC-RGC Joint Research Scheme (JRS) under project (52261160379, N_PolyU599/22).

References

- Berne, A., & Krajewski, W. F. (2013). Radar for hydrology: Unfulfilled promise or unrecognized potential? *Advances in Water Resources*, 51, 357–366. <https://doi.org/10.1016/j.advwatres.2012.05.005>
- Boyd, S., Parikh, N., Chu, E., Peleato, B., & Eckstein, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1), 1–122. <https://doi.org/10.1561/22000000016>
- China.com. (2019). 2018 China flood and drought disaster bulletin [in Chinese].
- Ciach, G. J. (2003). Local random errors in tipping-bucket rain gauge measurements. *Journal of Atmospheric and Oceanic Technology*, 20(5), 752–759. [https://doi.org/10.1175/1520-0426\(2003\)20<752:lreitb>2.0.co;2](https://doi.org/10.1175/1520-0426(2003)20<752:lreitb>2.0.co;2)
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *Paper presented at 2009 IEEE conference on computer vision and pattern recognition*. IEEE.
- Field, C. B., Barros, V., Stocker, T. F., & Dahe, Q. (2012). *Managing the risks of extreme events and disasters to advance climate change adaptation: Special report of the intergovernmental panel on climate change*. Cambridge University Press.
- Garg, K., & Nayar, S. K. (2007). Vision and rain. *International Journal of Computer Vision*, 75(1), 3–27. <https://doi.org/10.1007/s11263-006-0028-6>
- Gires, A., Onof, C., Maksimovic, C., Schertzer, D., Tchiguirinskaia, I., & Simoes, N. (2012). Quantifying the impact of small scale unmeasured rainfall variability on urban runoff through multifractal downscaling: A case study. *Journal of Hydrology*, 442, 117–128. <https://doi.org/10.1016/j.jhydrol.2012.04.005>
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Paper presented at Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Hut, R., de Jong, S., & van de Giesen, N. (2014). Using umbrellas as mobile rain gauges: Prototype demonstration. In *Paper presented at EGU general assembly conference abstracts*.
- Jiang, S., Babovic, V., Zheng, Y., & Xiong, J. (2019). Advancing opportunistic sensing in hydrology: A novel approach to measuring rainfall with ordinary surveillance cameras. *Water Resources Research*, 55(4), 3004–3027. <https://doi.org/10.1029/2018wr024480>
- Li, M., Xie, Q., Zhao, Q., Wei, W., Gu, S., Tao, J., & Meng, D. (2018). Video rain streak removal by multiscale convolutional sparse coding. In *Paper presented at Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Li, Y., & Zhu, J. (2008). L 1-norm quantile regression. *Journal of Computational & Graphical Statistics*, 17(1), 163–185. <https://doi.org/10.1198/106186008x289155>
- Lin, L., Wu, Z., & Liang, Q. (2019). Urban flood susceptibility analysis using a GIS-based multi-criteria analysis framework. *Natural Hazards*, 97(2), 455–475. <https://doi.org/10.1007/s11069-019-03615-2>
- Long, D., Longuevergne, L., & Scanlon, B. R. (2015). Global analysis of approaches for deriving total water storage changes from GRACE satellites. *Water Resources Research*, 51(4), 2574–2594. <https://doi.org/10.1002/2014wr016853>
- National Academies of Sciences, E., and Medicine. (2019). *Framing the challenge of urban flooding in the United States*. National Academies Press.
- O'Donnell, E. C., & Thorne, C. R. (2020). Drivers of future urban flood risk. *Philosophical Transactions of the Royal Society A*, 378(2168), 20190216. <https://doi.org/10.1098/rsta.2019.0216>
- Qin, Y., Chen, Z., Shen, Y., Zhang, S., & Shi, R. (2014). Evaluation of satellite rainfall estimates over the Chinese Mainland. *Remote Sensing*, 6(12), 11649–11672. <https://doi.org/10.3390/rs6111649>
- Rabiei, E., Haberlandt, U., Sester, M., & Fitzner, D. (2013). Rainfall estimation using moving cars as rain gauges—laboratory experiments. *Hydrology and Earth System Sciences*, 17(11), 4701–4712. <https://doi.org/10.5194/hess-17-4701-2013>
- Uijlenhoet, R., Overeem, A., & Leijnse, H. (2018). Opportunistic remote sensing of rainfall using microwave links from cellular communication networks. *Wiley Interdisciplinary Reviews: Water*, 5(4), e1289. <https://doi.org/10.1002/wat2.1289>
- Vojinovic, Z., Sahlu, S., Torres, A., Seyoum, S., Anvarifar, F., Matungulu, H., et al. (2014). Multi-objective rehabilitation of urban drainage systems under uncertainties. *Journal of Hydroinformatics*, 16(5), 1044–1061. <https://doi.org/10.2166/hydro.2014.223>
- Wang, H., Wu, Y., Li, M., Zhao, Q., & Meng, D. (2019). A survey on rain removal from video and single image. arXiv preprint arXiv:1909.08326.
- Wei, W., Yi, L., Xie, Q., Zhao, Q., Meng, D., & Xu, Z. (2017). Should we encode rain streaks in video as deterministic or stochastic? In *Paper presented at proceedings of the IEEE international conference on computer vision*.
- Winsemius, H. C., Aerts, J. C., Van Beek, L. P., Bierkens, M. F., Bouwman, A., Jongman, B., et al. (2016). Global drivers of future river flood risk. *Nature Climate Change*, 6(4), 381–385. <https://doi.org/10.1038/nclimate2893>
- Xu, J., Zhao, W., Liu, P., & Tang, X. (2012). Removing rain and snow in a single image using guided filter. In *Paper presented at 2012 IEEE international conference on computer science and automation engineering (CSAE)*. IEEE.
- Yin, H., Zheng, F., Duan, H.-F., Savic, D., & Kapelan, Z. (2023). Estimating rainfall intensity using an image-based deep learning model. *Engineering*, 21, 162–174. <https://doi.org/10.1016/j.eng.2021.11.021>
- Yong, B., Hong, Y., Ren, L. L., Gourley, J. J., Huffman, G. J., Chen, X., et al. (2012). Assessment of evolving TRMM-based multisatellite real-time precipitation estimation methods and their impacts on hydrologic prediction in a high latitude basin. *Journal of Geophysical Research*, 117(D9), D09108. <https://doi.org/10.1029/2011jd017069>
- Zhang, X., Li, H., Qi, Y., Leow, W. K., & Ng, T. K. (2006). Rain removal in video by combining temporal and chromatic properties. In *Paper presented at 2006 IEEE international conference on multimedia and expo*. IEEE.
- Zheng, F., Tao, R., Maier, H. R., See, L., Savic, D., Zhang, T., et al. (2018). Crowdsourcing methods for data collection in geophysics: State of the art, issues, and future directions. *Reviews of Geophysics*, 56(4), 698–740. <https://doi.org/10.1029/2018rg000616>
- Zheng, Z., Gao, J., Ma, Z., Wang, Z., Yang, X., Luo, X., et al. (2016). Urban flooding in China: Main causes and policy recommendations. *Hydrological Processes*, 30(7), 1149–1152. <https://doi.org/10.1002/hyp.10717>
- Zhou, Q., Zhang, Y., Li, B., Li, L., Feng, J., Jia, S., et al. (2019). Cloud-base and cloud-top heights determined from a ground-based cloud radar in Beijing, China. *Atmospheric Environment*, 201, 381–390. <https://doi.org/10.1016/j.atmosenv.2019.01.012>
- Zorzetto, E., & Marani, M. (2019). Downscaling of rainfall extremes from satellite observations. *Water Resources Research*, 55(1), 156–174. <https://doi.org/10.1029/2018wr022950>