

Deep Reinforcement Learning for Adaptive Control of Urban Drainage and Green Infrastructure Under Climate Change

RAJKISHOR, ABHIJEET RAJ, AJEET KUMAR,
AMAN KUMAR, AJIT DAS

Abstract: *Climate-driven intensification of short-duration rainfall is eroding the design margins of urban drainage networks, and static operating rules for storage basins, gates, pumps and green infrastructure (GI) can no longer cope with non-stationary storm regimes. This study develops a deep reinforcement learning (DRL) framework for the coordinated, real-time control of conventional drainage assets and distributed GI under projected climate change. The control problem is formulated as a Markov decision process in which an agent observes water levels, rainfall nowcasts and GI storage states and issues continuous gate, pump and orifice commands. A proximal policy optimization (PPO) agent with a multi-objective reward that penalizes flooding, outfall exceedance and control effort was trained on randomized synthetic storms in a hydraulic environment representing a 4.2 km² urban catchment, and was evaluated on held-out events against passive, rule-based, model predictive control (MPC) and deep Q-network (DQN) alternatives under three scenarios of rainfall intensification (0%, +15% and +30%). The proposed controller reduced total flood volume by 42.6–45.7% relative to passive operation and by 11.4–14.3% relative to MPC, while keeping control actions smooth and decision latency below 5 ms. Performance degraded gracefully as rainfall intensified, and the learned policy exploited GI retention ahead of peak inflow. The results indicate that DRL is a credible option for climate-adaptive operation of hybrid grey–green systems, provided that safety constraints, uncertainty handling and field validation are addressed.*

Keywords: *Deep Reinforcement Learning; Real-Time Control; Urban Drainage; Green Infrastructure; Climate Adaptation; Flood Mitigation; Proximal Policy Optimization*

Rajkishor, Department of Civil Engineering, Bhagalpur College of Engineering, Bhagalpur, India
Abhijeet Raj, Department of Geology, Government Engineering College Vaishali, Bihar, India
Ajeet Kumar, Department of Civil Engineering, Government Engineering College Vaishali, Bihar, India
Aman Kumar, Department of Civil Engineering, Government Engineering College Vaishali, Bihar, India
Ajit Das, Department of Civil Engineering, Government Engineering College Vaishali, Bihar, India
*Corresponding author: Rajkishor (rk.bce22@gmail.com). Co-authors: abhijeetraj571@gmail.com;
ajeet.gecv@gmail.com; amangecv@gmail.com; ajit.gecv@gmail.com

Introduction

Urban flooding is among the most frequent and costly consequences of rapid urbanization combined with a warming climate. Higher atmospheric moisture capacity increases the intensity of short-duration rainfall extremes, and observational and modelling evidence shows that sub-daily extremes are intensifying faster than daily totals in many regions (Westra et al., 2014; Fowler et al., 2021). The Sixth Assessment Report of the Intergovernmental Panel on Climate Change concludes that heavy precipitation events will continue to become more intense and frequent with additional warming (IPCC, 2021). In South Asia, where monsoon convection already produces severe short bursts, increasing trends in extreme rain events have been documented for several decades (Goswami et al., 2006). Drainage systems designed on the assumption of stationary rainfall statistics therefore face a growing mismatch between hydraulic capacity and hydrological demand.

Two broad responses are available. The first is structural: enlarging pipes and tanks, which is costly, disruptive in dense cities and prone to becoming undersized again as the climate shifts. The second is to extract more performance from existing assets through real-time control (RTC) and to complement them with distributed green infrastructure such as bioretention cells, permeable pavements and rain gardens (Fletcher et al., 2015). RTC has been studied for decades in wastewater and combined sewer systems (Schütze et al., 2004; García et al., 2015), and recent work on smart stormwater systems has highlighted the potential of low-cost sensing and actuation to turn passive networks into adaptive ones (Kerkez et al., 2016).

Most operational RTC schemes rely on rule-based logic or model predictive control. Rule-based schemes are transparent but tuned to a narrow set of conditions. MPC can optimize over a forecast horizon, but it requires a fast and sufficiently accurate internal model, solves an optimization problem at every control step and scales poorly when many nonlinear assets, including GI units with stage-dependent outflow, are coordinated (Lund et al., 2018). Deep reinforcement learning (DRL) offers a complementary route: the agent learns a control policy by interacting with a simulator, and once trained the policy is evaluated with a single forward pass of a neural network (Sutton & Barto, 2018; Mnih et al., 2015). Early applications show that DRL can reduce flooding in stormwater networks (Mullapudi et al., 2020) and in coastal catchments affected by sea-level rise (Bowes et al., 2021; Sadler et al., 2020).

Despite this progress, three gaps remain. First, most studies control conventional storage and conveyance assets only and treat GI as a fixed, passive land-use element, although adjustable outlets and pre-event drawdown can make GI an active component of the control system. Second, climate non-stationarity is rarely built into the training process, so the policies may be brittle when rainfall exceeds the range seen in training. Third, the comparison with strong optimization-based benchmarks is often limited. This paper addresses these gaps by (i) formulating a unified DRL problem for hybrid grey–green drainage with continuous actuation of gates, pumps and GI outlets, (ii) training the agent on storms with randomized intensification so that it encounters a climate-informed distribution, and (iii) benchmarking against passive, rule-based, MPC and DQN controllers under three climate scenarios.

Background and Related Work

2.1 Real-time control of urban drainage

Real-time control of drainage systems aims to use storage and conveyance capacity dynamically in space and time. Reviews of the field indicate that benefits are largest in networks with unevenly loaded sub-catchments and underused storage, where coordination can redistribute volume before overflow occurs (Schütze et al., 2004; García et al., 2015). Model predictive control has become the reference optimization-based approach, but its practical deployment is limited by model calibration effort, computational cost and sensitivity to forecast errors (Lund et al., 2018). Simplified digital twins, such as the Pipedream model, have been developed to make fast simulation and control experiments feasible (Bartos & Kerkez, 2021).

2.2 Reinforcement learning for water systems

Reinforcement learning (RL) frames control as sequential decision-making under uncertainty. Value-based methods such as the deep Q-network handle discrete actions (Mnih et al., 2015), whereas policy-gradient and actor–critic methods, including deep deterministic policy gradient, soft actor–critic and proximal policy optimization, support continuous actions and more stable training (Lillicrap et al., 2016; Haarnoja et al., 2018; Schulman et al., 2017). In stormwater applications, Mullanpudi et al. (2020) demonstrated DRL control of valves in a simulated network, and Bowes et al. (2021) showed flood mitigation in a coastal catchment under tidal and rainfall forcing. Reproducibility is a recognized difficulty in DRL because results are sensitive to seeds and hyperparameters (Henderson et al., 2018), which motivates multi-seed evaluation as adopted here.

2.3 Green infrastructure as a controllable asset

Green infrastructure reduces runoff volume and peak flow through infiltration, evapotranspiration and detention. Its effect is normally represented through static parameters in hydrologic models. In this study, GI clusters are equipped with adjustable outlet orifices so that stored water can be released ahead of an approaching storm, thereby recovering retention capacity at the moment it is most valuable. This operating concept is analogous to pre-storm drawdown in detention basins and creates a coupling between rainfall nowcasts and GI storage that a learned policy can exploit.

Methodology

3.1 Problem formulation

The drainage control task is modelled as a Markov decision process defined by the tuple (S, A, T, R, γ) , where S is the state space, A the action space, T the transition dynamics given by the hydraulic environment, R the reward function and γ the discount factor. At each control step t ($\Delta t = 5$ min), the agent observes the state s_t and selects an action a_t to maximize the expected discounted return:

$$J(\pi) = E[\sum_t \gamma^t r(s_t, a_t)] \quad (1)$$

The state vector concatenates (i) water depths at monitored nodes and storage units, normalized by node capacity; (ii) current rainfall intensity and a 30-min rainfall nowcast aggregated into three 10-min bins; (iii) fill fractions of GI clusters; (iv) the previous action, which allows the policy to learn smoothness; and (v) a scalar climate-context indicator equal to the intensification factor of the forecast storm relative to the historical baseline. The inclusion of the last element is optional in deployment because the nowcast already carries the information, but it accelerated learning in preliminary tests.

The action vector $a_t \in [0, 1]^{11}$ contains the opening fractions of three controllable gates, the normalized speeds of two pump stations and the outlet opening of six GI clusters. Opening fractions are mapped to flow through orifice and pump-curve relations inside the environment.

3.2 Hydraulic and hydrologic environment

The environment couples a runoff generation module with a nonlinear reservoir routing scheme for the conduit and storage network. Subcatchment runoff is computed from rainfall, imperviousness and infiltration losses; node continuity is

updated at each step; and link flows follow a simplified dynamic relation constrained by pipe capacity and downstream depth. The formulation follows the structure used in widely adopted stormwater management models (Rossman, 2015) but is implemented as a lightweight vectorized environment so that thousands of episodes can be simulated within practical time. GI clusters are represented as storage units with infiltration to the underlying soil, evapotranspiration losses and a controllable outlet.

3.3 Reward design

The reward balances flood prevention, receiving-water protection, actuation smoothness and effective use of GI:

$$r_t = -w_1 F_t - w_2 O_t - w_3 \|a_t - a_{t-1}\|^2 + w_4 G_t \quad (2)$$

where F_t is the flood volume at all nodes during step t , O_t is outfall discharge exceeding a permitted threshold, the third term penalizes abrupt actuation that may damage equipment, and G_t is the volume retained in GI. The weights (Table 1) were chosen so that avoided flooding dominates, which is consistent with reward-shaping principles whereby secondary terms guide learning without altering the primary objective (Ng et al., 1999). The GI term is capped when GI is full to avoid rewarding unnecessary filling.

3.4 Learning algorithm

Proximal policy optimization was selected because it handles continuous action spaces, uses a clipped surrogate objective that limits destructive policy updates and has shown robust performance across domains (Schulman et al., 2017). The actor and critic are separate multilayer perceptrons with two hidden layers of 128 units. The clipped objective is

$$L(\theta) = E[\min(\rho_t(\theta) \hat{A}_t, \text{clip}(\rho_t(\theta), 1-\epsilon, 1+\epsilon) \hat{A}_t)] \quad (3)$$

where $\rho_t(\theta)$ is the probability ratio between the new and old policies, $\hat{A}_t$ is the generalized advantage estimate and ϵ is the clipping parameter. Optimization used the Adam algorithm (Kingma & Ba, 2015). Training episodes were generated by sampling rainfall events from a synthetic generator that produces Chicago-type hyetographs with return periods of 2–100 years, durations of 1–3 h and random temporal skewness. Each event was multiplied by an intensification factor drawn uniformly from 1.00 to 1.30, which exposes the agent to a range of climate futures rather than a single projection. Five independent training seeds were run, and results are reported as means across seeds.

3.5 Benchmark controllers

Four benchmarks were implemented. The passive configuration keeps gates fully open, pumps off and GI outlets at their default setting. The rule-based controller applies depth-threshold logic, opening gates and starting pumps when upstream depth exceeds 70% of capacity. The MPC benchmark solves a receding-horizon optimization with a 30-min horizon using the same nowcast supplied to the agent, a linearized internal model and a gradient-based solver. The DQN benchmark uses the same state vector but a discretized action set of five levels per actuator, reduced through grouped actuation to keep the action space tractable.

Case Study and Experimental Design

The study uses a synthetic urban catchment of 4.2 km² with 62% imperviousness, 24 network nodes, three controllable gates, two pump stations and six GI clusters (bioretention, permeable pavement and rain gardens) that together treat roughly 9% of the contributing area. The configuration was designed to resemble a medium-density Indian city district with limited conveyance headroom and low-lying outfall zones, but it does not represent a specific monitored site. Using a synthetic benchmark allows controlled comparison and full reproducibility, but it also means that absolute values should not be transferred directly to real systems.

Climate scenarios were defined by scaling design-storm intensity. A scaling of +15% corresponds approximately to two degrees of warming and +30% to about four degrees if short-duration extremes scale near the Clausius–Clapeyron rate of roughly 7% per degree (Fowler et al., 2021; IPCC, 2021). Each controller was evaluated on 200 held-out storm events per scenario that were not used in training. Performance was assessed using total flood volume, peak outfall discharge, mean GI utilization and per-step decision time. Hyperparameters and environment settings are given in Table 1.

Table 1. Experimental configuration and training hyperparameters.

Component	Setting
Catchment	4.2 km ² ; 62% impervious; 24 nodes
Controllable assets	3 gates; 2 pump stations; 6 GI outlets (11 actions)
Control time step	5 min; 6 h episode (72 steps)

Training storms	Chicago-type; 2–100 yr; 1–3 h; intensification 1.00–1.30
Algorithm	PPO; actor and critic 2×128 (tanh)
Learning rate / optimizer	3×10^{-4} / Adam
Discount γ / GAE λ / clip ε	0.99 / 0.95 / 0.2
Rollout batch size	2,048 steps; 3,000 training episodes; 5 seeds
Reward weights (w_1 , w_2 , w_3 , w_4)	10, 3, 0.1, 0.5
Test events	200 held-out events per scenario

Results and Discussion

5.1 Training behaviour

Figure 1a shows the mean episodic reward during training. The PPO agent improved rapidly during the first 600 episodes and plateaued after approximately 2,000 episodes at a reward that exceeded both the MPC and rule-based benchmarks. The DQN agent converged more slowly and to a lower plateau, a pattern attributable to its coarse discretization of a high-dimensional action space. Variability across seeds was small for PPO after convergence, indicating that the reward and state design were not unduly sensitive to initialization, although more seeds would be needed for formal statistical claims (Henderson et al., 2018).




Figure 1. (a) Mean episodic reward during training for PPO and DQN agents with MPC and rule-based benchmark levels; (b) total flood volume for each controller under baseline, +15% and +30% rainfall scenarios.

5.2 Flood mitigation under climate scenarios

Table 2 and Figure 1b summarize flood performance. In the baseline scenario, passive operation produced $48.6 \times 10^3 \text{ m}^3$ of flooding, and PPO reduced it to $26.4 \times 10^3 \text{ m}^3$, a reduction of 45.7%. With +15% rainfall, the reduction was 44.9%, and with +30% it was 42.6%. All controllers lost effectiveness as intensification increased because the system approached its physical limits: when inflow volume exceeds the combined storage and conveyance capacity, no control strategy can prevent flooding. The slow decline of the PPO benefit, however, indicates that the policy retained most of its advantage under conditions at the upper end of its training distribution.

Table 2. Mean total flood volume (10^3 m^3) on held-out events (200 per scenario, mean of five seeds) for each controller and rainfall scenario.

Controller	Baseline (0%)	Moderate (+15%)	High (+30%)	Reduction passive (%)* vs.
Passive	48.6	63.9	82.4	—
Rule-based	39.2	53.1	70.7	17.3–19.3
MPC	29.8	40.6	55.2	32.9–38.7
DQN	31.5	43.9	59.8	27.4–35.2
PPO (proposed)	26.4	35.2	47.3	42.6–45.7

*Range across the three scenarios, computed as $(\text{passive} - \text{controller})/\text{passive}$.

The learned policy was 11.4%, 13.3% and 14.3% better than MPC in the three scenarios, and the margin widened as rainfall intensified. A plausible explanation is that the MPC internal model, being linearized, becomes less accurate under surcharged and highly nonlinear conditions, whereas the neural policy was trained on the full nonlinear environment. The DQN controller outperformed rule-based control in all scenarios but trailed MPC, which suggests that the continuous action formulation matters for coordinating many actuators.

5.3 Role of green infrastructure and control behaviour

Inspection of simulated events showed that the PPO agent opened GI outlets 20–40 min before the peak of the nowcast hyetograph in events with high forecast intensity, creating extra storage in advance and then closing the outlets as the network filled. Mean GI utilization at peak inflow was highest for PPO (about 78%), compared with

52% for the rule-based controller, which does not use GI actively. This anticipatory behaviour is not prescribed in the reward and emerges from the combination of nowcast inputs and the GI retention term, which illustrates the main advantage of learning-based control: strategies that are difficult to encode manually are discovered through interaction.

Peak outfall discharge was also reduced relative to passive operation (by 31–38% across scenarios), which is important for downstream channel erosion and water quality. The smoothness penalty limited gate oscillation, with the mean absolute change in gate opening per step below 6%, a level compatible with typical actuator limits. Once trained, the policy required less than 5 ms per decision on a standard laptop CPU, far lower than the several hundred milliseconds to seconds typically needed to solve the MPC problem at each step, which favours deployment on low-power edge devices.

5.4 Practical implications

For municipal agencies, the results suggest that retrofitting actuators and sensors to existing structures, together with a trained policy, can delay or reduce the need for capacity expansion under moderate climate change. The approach is most attractive where storage is available but underused and where GI is already planned, since adding controllable outlets to GI is inexpensive compared with new tunnels or basins. The inclusion of climate-randomized storms in training provides a pragmatic way of preparing the controller for non-stationarity, although it does not replace periodic retraining as observations and projections evolve.

Limitations and Future Work

Several limitations should be acknowledged. The environment is a reduced-order model of a synthetic catchment; hydraulic phenomena such as backwater, surcharge dynamics and sediment deposition are simplified, and an agent trained in a simplified simulator may exploit its inaccuracies. Rainfall nowcasts were generated by perturbing true rainfall with noise, which may underestimate real nowcast error during convective events. Safety was handled through the reward and action bounds rather than formal guarantees, which is insufficient for critical infrastructure. Future work should (i) validate the framework on a calibrated model of a real catchment and with hardware-in-the-loop tests, (ii) incorporate safe RL techniques such as constrained policy optimization and shielding, (iii) represent forecast and sensor

uncertainty explicitly, (iv) explore multi-agent formulations for large networks and (v) integrate water-quality objectives and life-cycle costs.

Conclusions

This study presented a deep reinforcement learning framework for the adaptive, coordinated control of urban drainage assets and green infrastructure under climate change. By training a PPO agent on climate-randomized storms in a hybrid grey–green hydraulic environment, the controller reduced flood volume by 42.6–45.7% relative to passive operation and by 11.4–14.3% relative to model predictive control across three rainfall-intensification scenarios, with smooth actuation and negligible decision latency. The learned policy exhibited anticipatory use of green infrastructure storage, demonstrating that GI can be an active control asset rather than a passive land-use feature. These findings support further development of DRL-based control for climate-resilient stormwater management, with priority given to field validation, safety assurance and uncertainty-aware training before operational deployment.

References

- Bartos, M., & Kerkez, B. (2021). Pipedream: An interactive digital twin model for natural and urban drainage systems. *Environmental Modelling & Software*, 144, 105120. <https://doi.org/10.1016/j.envsoft.2021.105120>
- Bowes, B. D., Tavakoli, A., Wang, C., Heydarian, A., Behl, M., Beling, P. A., & Goodall, J. L. (2021). Flood mitigation in coastal urban catchments using real-time stormwater infrastructure control and reinforcement learning. *Journal of Hydroinformatics*, 23(3), 529–547. <https://doi.org/10.2166/hydro.2020.080>
- Fletcher, T. D., Shuster, W., Hunt, W. F., Ashley, R., Butler, D., Arthur, S., Trowsdale, S., Barraud, S., Semadeni-Davies, A., Bertrand-Krajewski, J.-L., Mikkelsen, P. S., Rivard, G., Uhl, M., Dagenais, D., & Viklander, M. (2015). SUDS, LID, BMPs, WSUD and more – The evolution and application of terminology surrounding urban drainage. *Urban Water Journal*, 12(7), 525–542. <https://doi.org/10.1080/1573062X.2014.916314>
- Fowler, H. J., Lenderink, G., Prein, A. F., Westra, S., Allan, R. P., Ban, N., Barbero, R., Berg, P., Blenkinsop, S., Do, H. X., Guerreiro, S., Haerter, J. O., Kendon, E. J., Lewis, E., Schaer, C., Sharma, A., Villarini, G., Wasko, C., & Zhang, X. (2021). Anthropogenic intensification of short-duration rainfall extremes. *Nature Reviews Earth & Environment*, 2(2), 107–122. <https://doi.org/10.1038/s43017-020-00128-6>

García, L., Barreiro-Gomez, J., Escobar, E., Téllez, D., Quijano, N., & Ocampo-Martinez, C. (2015). Modeling and real-time control of urban drainage systems: A review. *Advances in Water Resources*, 85, 120–132. <https://doi.org/10.1016/j.advwatres.2015.08.007>

Goswami, B. N., Venugopal, V., Sengupta, D., Madhusoodanan, M. S., & Xavier, P. K. (2006). Increasing trend of extreme rain events over India in a warming environment. *Science*, 314(5804), 1442–1445. <https://doi.org/10.1126/science.1132027>

Haarnoja, T., Zhou, A., Abbeel, P., & Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *Proceedings of the 35th International Conference on Machine Learning*, 1861–1870.

Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., & Meger, D. (2018). Deep reinforcement learning that matters. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), 3207–3214.

Intergovernmental Panel on Climate Change. (2021). *Climate change 2021: The physical science basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* (V. Masson-Delmotte et al., Eds.). Cambridge University Press. <https://doi.org/10.1017/9781009157896>

Kerkez, B., Gruden, C., Lewis, M., Montestruque, L., Quigley, M., Wong, B., Bedig, A., Kertesz, R., Braun, T., Cadwalader, O., Poresky, A., & Pak, C. (2016). Smarter stormwater systems. *Environmental Science & Technology*, 50(14), 7267–7273. <https://doi.org/10.1021/acs.est.5b05870>

Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. *Proceedings of the 3rd International Conference on Learning Representations (ICLR 2015)*. <https://arxiv.org/abs/1412.6980>

Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., & Wierstra, D. (2016). Continuous control with deep reinforcement learning. *Proceedings of the 4th International Conference on Learning Representations (ICLR 2016)*. <https://arxiv.org/abs/1509.02971>

Lund, N. S. V., Falk, A. K. V., Borup, M., Madsen, H., & Mikkelsen, P. S. (2018). Model predictive control of urban drainage systems: A review and perspective towards smart real-time water management. *Critical Reviews in Environmental Science and Technology*, 48(3), 279–339. <https://doi.org/10.1080/10643389.2018.1455484>

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533. <https://doi.org/10.1038/nature14236>

Mullapudi, A., Lewis, M. J., Gruden, C. L., & Kerkez, B. (2020). Deep reinforcement learning for the real-time control of stormwater systems. *Advances in Water Resources*, 140, 103600. <https://doi.org/10.1016/j.advwatres.2020.103600>

Ng, A. Y., Harada, D., & Russell, S. (1999). Policy invariance under reward transformations: Theory and application to reward shaping. *Proceedings of the 16th International Conference on Machine Learning*, 278–287.

Rossman, L. A. (2015). Storm Water Management Model user's manual, version 5.1 (EPA/600/R-14/413b). U.S. Environmental Protection Agency.

Sadler, J. M., Goodall, J. L., Behl, M., Bowes, B. D., & Morsy, M. M. (2020). Exploring real-time control of stormwater systems for mitigating flood risk due to sea level rise. *Journal of Hydrology*, 583, 124571. <https://doi.org/10.1016/j.jhydrol.2020.124571>

Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv*. <https://arxiv.org/abs/1707.06347>

Schütze, M., Campisano, A., Colas, H., Schilling, W., & Vanrolleghem, P. A. (2004). Real time control of urban wastewater systems—where do we stand today? *Journal of Hydrology*, 299(3–4), 335–348. <https://doi.org/10.1016/j.jhydrol.2004.08.010>

Sutton, R. S., & Barto, A. G. (2018). Reinforcement learning: An introduction (2nd ed.). MIT Press.

Westra, S., Fowler, H. J., Evans, J. P., Alexander, L. V., Berg, P., Johnson, F., Kendon, E. J., Lenderink, G., & Roberts, N. M. (2014). Future changes to the intensity and frequency of short-duration extreme rainfall. *Reviews of Geophysics*, 52(3), 522–555. <https://doi.org/10.1002/2014RG000464>