

Article

Mitigation of Flooding in Stormwater Systems Utilizing Imperfect Forecasting and Sensor Data with Deep Deterministic Policy Gradient Reinforcement Learning

Sami Saliba^{1,†} , Benjamin D. Bowes^{1,2,†} , Stephen Adams¹ , Peter A. Beling^{1,2} , Jonathan L. Goodall^{1,2,*} 

¹ Dept. of Engineering Systems and Environment, Univ. of Virginia, P.O. Box 400747, Charlottesville, VA 22904, USA; sms8fr@virginia.edu (S.M.S.); bdb3m@virginia.edu (B.D.B.); sca2c@virginia.edu (S.C.A.); beling@virginia.edu (P.A.B.); goodall@virginia.edu (J.L.G.)

² Link Lab, Univ. of Virginia, P.O. Box 400259, Charlottesville, VA 22904, USA

* Correspondence: goodall@virginia.edu

† These authors contributed equally to this work.

Abstract: Climate change and development have increased urban flooding, requiring modernization of stormwater infrastructure. Retrofitting standard passive systems with controllable valves/pumps is promising, but requires real-time control (RTC). One method of automating RTC is reinforcement learning (RL), a general technique for sequential optimization and control in uncertain environments. The notion is that an RL algorithm can use inputs of real-time flood data and rainfall forecasts to learn a policy for controlling the stormwater infrastructure to minimize measures of flooding. In real-world conditions, rainfall forecasts and other state information, are subject to noise and uncertainty. To account for these characteristics of the problem data, we implemented Deep Deterministic Policy Gradient (DDPG), an RL algorithm that is distinguished by its capability to handle noise in the input data. DDPG implementations were trained and tested against a passive flood control policy. Three primary cases were studied: (i) perfect data, (ii) imperfect rainfall forecasts, and (iii) imperfect water level and forecast data. Rainfall episodes (100) that caused flooding in the passive system were selected from 10 years of observations in Norfolk, Virginia, USA; 85 randomly selected episodes were used for training and the remaining 15 unseen episodes served as test cases. Compared to the passive system, all RL implementations reduced flooding volume by 70.5% on average, and performed within a range of 5%. This suggests that DDPG is robust to noisy input data, which is essential knowledge to advance the real-world applicability of RL for stormwater RTC.

Keywords: Real-time Control; Reinforcement Learning; Smart Stormwater Systems; Urban Flooding

1. Introduction

Flooding in urban areas is becoming more prevalent due to changing weather patterns, rising sea levels, and increases in impermeable surfaces [1–4]. The cost of stormwater flooding is estimated in the billions of dollars with the loss of property, potential for loss of life, and mobilization of emergency personnel to keep people safe [5]. Currently, the majority of stormwater systems are designed to passively manage rainfall and runoff. Although these systems are designed for planned growth and development, they often cannot accommodate the increased water volumes with new or changing construction of buildings and roads, which decrease permeable surfaces, or the increasing storm intensities associated with climate change [6,7].

Retrofitting stormwater systems for active monitoring and real time control (RTC), while becoming more common, is not yet the standard [8]. The utilization of these dynamic systems with heuristic control has shown significant performance improvements compared to their passive counterparts [9]. Heuristic control typically offers a generalized strategy to adjust of valves and pumps based on a prediction model. While shown to be successful on an individual storage pond, the repercussions of having multiple ponds in parallel are rarely considered. For example, when two of these heuristic controlled systems are used, each is designed to monitor its own storage pond, and thus may cause flooding downstream by releasing water to reduce their own flooding.

Reinforcement learning (RL) [10], a type of machine learning, is an emerging approach to stormwater system RTC that allows the creation of policies for flow control valves, pumps, and ponds within a stormwater system using simulations [11,12]. Machine learning has previously been utilized to predict flooding [13], but this strategy used perfect prediction rainfall data and was not used to reduce flooding. With all forms of RL, the ultimate goal of the agent is to learn a policy (π) for the given simulated environment. A policy acts as a map for actions to a given state. In the case of this project, sensor data (state information) is connected to flow valve opening (action) at any given time. RL functions similarly in principle to the way humans or animals learn or make behavioral modifications. An action occurs and is evaluated, then a reward is then applied. A greater reward incentivizes an action, and a lesser reward discourages an action. As an analogy, a treat serves as positive reinforcement for the subject, while punishment sends a negative message. With consistent administration, the human or animal “learns” to avoid negative consequences using this technique, while positive rewards reinforce good behaviors. The RL is designed with reward functions to provide a positive or negative reward.

Deep Deterministic Policy Gradients (DDPG) is a RL algorithm that uses a deep learning framework to create function approximators with neural networks [14]. DDPG is an actor-critic RL. The Actor and Critic are separate neural networks, but they collaborate through constant interaction. The action is performed by an actor, and the evaluation of the performance of the actor is done by a critic. The Actor takes actions in the environment based on the gradient of the Critic. The Critic receives information about the environment and “grades” the action made by the Actor. The relationship of the actor Actor and the Critic dictates how the overall agent learns based on a reward function that reinforces the learning [15]. DDPG is useful to learn policies in complex environments and problems in which the number of available control actions is large or a continuous space. This would make DDPG an good choice for stormwater systems in which a continuum of valve settings can be used to shift water to another location that may have upstream, downstream, or even parallel implications. In the models used here, control valves can be opened to any value from 0-100%. Finding the most effective reward function is a key component to fully automating the stormwater systems to mitigate flooding across broad areas in a complex environments.

Bowes et al. (2020) demonstrated an improvement over the passive system, reducing flooding using DDPG when the sum of flooding was used as a reward function to train the RL agent [16]. They demonstrated that RL could reduce flooding compared to an passive stormwater system and rule-based RTC. Although they were able to reduce flooding substantially, their RL agent was trained and tested with perfect state information: the dataset used exact water levels at a current time, and the

rainfall forecast was equal to the exact future rainfall amount. In reality however, rainfall forecasts are uncertain, and sensor data are noisy.

A simulation representing the real world must combine a multitude of sensors working in parallel, should accommodate imperfect data or imperfections in the sensors, and should account for uncertain weather conditions. A RL strategy that includes factors such as incorrect water measurements, and imperfect forecasts has not been examined. Considering that weather forecasts and rainfall prediction have poor reliability even a day in advance [17], robustness to these variables is essential. Our goal, therefore, is to advance previous work by studying the use of imperfect data that is more representative of true environmental conditions. In other applications, RL has been used to find superior control policies despite imperfect data [18]. But the use of RL with imperfect data has not been studied for stormwater RTC. Therefore, the purpose of this study is to examine the effect of imperfect state information on the performance of DDPG in reducing flooding. We hypothesize that RL can accommodate for the imperfect data and provide an optimal or near-optimal solution, similar to perfect rainfall forecasting and fully functional sensors.

This project evaluates the effectiveness of DDPG based stormwater control in reducing flooding based on real rainfall data from Norfolk Virginia. Expanding on previous work, we explore the robustness of a DDPG implementation by utilizing noisy forecasting and state data. Without exact measurements, the heuristic or algorithmic control could cause catastrophic effects – holding water as the system floods, or releasing the entirety of the storage. However, this project demonstrates that DDPG can overcome imperfect data to successfully reducing flooding in all cases.

The remainder of the paper is organized as follows in Sections 2-5. Section 2 presents the methods by detailing the modeling of the stormwater environment, examining the development and implementation of a DDPG agent and RL, and describing the data. Finally, we introduce the experiment design. Section 3 presents the results comparing the RL with perfect and imperfect data to the Passive system. Section 4 is a discussion of the study with further description of the utility of the design, and Section 5 presents a conclusion and proposes future work.

2. Methods

This study was designed to compare the effectiveness of RL policies in reducing flooding volume under the following conditions: (i) perfect state and forecast data, (ii) perfect state and noisy forecast data, (iii) noisy state and forecast data. The RL policies were compared to a passive system where all stormwater valves were 100% open.

2.1. Modeling The Stormwater Environment

The US Environmental Protection Agency has developed software to simulate drainage and water transportation within urban environments called the Storm Water Management Model (SWMM) [19]. SWMM allows for the modeling of stormwater systems, taking into account runoff quality, hydrology, and elevation changes within a specific environment. It has been used when designing stormwater management systems and when development within a city changes such as building new communities or roads, as the number of impermeable surfaces affects stormwater absorption [20]. The SWMM tool allows for the modeling of a system, such as the City of Norfolk, Virginia, which commonly floods even in low-rate rainfall events when concurrent with high tides. SWMM allows for the manipulation of specific variables such as the number of storm drains, junctions, valves, and storage areas (retention ponds) within a simulation.

Figure 1 demonstrates the SWMM model utilized in this research is based on the designs of previous machine learning implementations by Bowes et al. (2020) [16]. This model houses eight subcatchments (S_{1-8}), three major junctions (J_{1-3}), and three flow control valves – orifices – (R_{1-3}) in this system. The elevation changes are not indicated within the diagram, but are programmed into the model and the system drains downward toward the outfall (*Out*). The R valve position can be dynamically controlled in any range of openings between 0% and 100%.

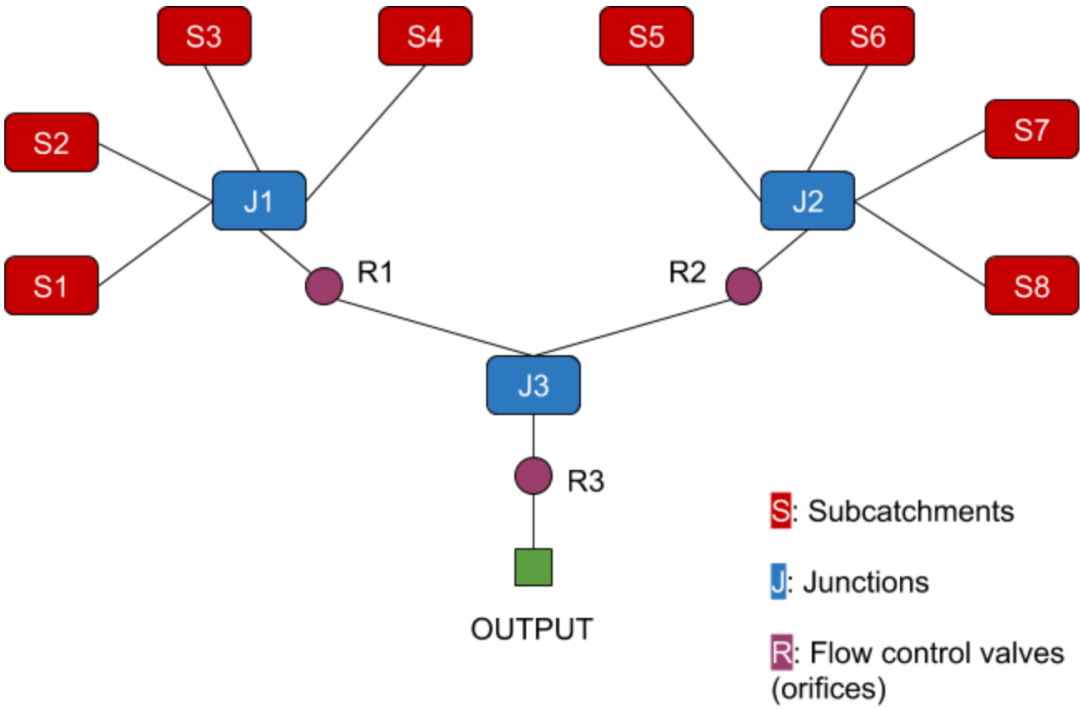


Figure 1. Graphic of the SWMM model used.

Table 1 shows SWMM setup parameters. Within the model, each subcatchment used the same rainfall data at each timestep, although each subcatchment has unique area, impermeability, and width. Although in the table and diagram there are junctions labeled, these junctions have a secondary function as storage ponds. Throughout, they are consistently denoted as junctions as although they do have a depth component, they act as the joint output of multiple connections.

Table 1. SWMM model setup parameters

Name	Gage	Outlet	Area(km2)	%Imperv	Width(m)	%Slope
S_1	Gage1	J_1	283.28	25	457.2	0.5
S_2		J_1	364.21	35	457.2	0.5
S_3		J_1	202.34	30	457.2	0.5
S_4		J_1	343.98	25	457.2	0.5
S_5		J_2	283.28	35	457.2	0.5
S_6		J_2	263.04	30	457.2	0.5
S_7		J_2	222.57	25	457.2	0.5
S_8		J_2	303.51	25	457.2	0.5
	Name	Elev.(m)	MaxDepth(m)	InitDepth(m)	Area(km2)	
	J_1	27.4	1.2	0.2286	0.0607	
	J_2	27.4	1.2	0.2286	0.0809	
	J_3	25.9	1.2	0.2286	0.0809	
Name	From Node	To Node	Length (m)	Roughness	Diameter (m)	
C	J_3	Out	121.92	0.01	0.2286	
Name	From Node	To Node	Type	Qcoeff	Gated	CloseTime
R_1	J_1	J_3	BOTTOM	0.65	NO	0.2
R_2	J_2	J_3	BOTTOM	0.65	NO	0.2
R_3	J_3	Out	BOTTOM	0.65	NO	0.2
			Name	Elevation (m)		
			Out	25.29		

SWMM allowed for measurement of depth at each junction, accounting for water height within the system and out of the system (flooding).

2.2. Reinforcement Learning

RL has been studied since the 1990s [10], but has only more recently become viable for complex problems with the advancement of computational technology. The sequential decision-making model for standard RL problems is a Markov decision process (MDP). An MDP is composed of a set of states (\mathcal{S}), a set of actions (\mathcal{A}), a reward function (\mathcal{R}), a transition function (\mathcal{P}), and a discount factor (γ). Let $s \in \mathcal{S}$ represent a specific state, and let $a \in \mathcal{A}$ represent a specific action. The reward function is generally a function of both state and action $\mathcal{R}(s, a)$. The transition function $\mathcal{P}(s'|s, a)$ defines the probability of moving to s' given s and a .

The core idea of RL is to learn from experience through an iterative process of taking actions in a given state and observing a reward and state change. The goal of the agent is to learn a policy that maximizes the sum of expected future reward

$$\begin{aligned}
 G_t &= R_t + \gamma R_{t+1} + \gamma^2 R_{t+2} + \dots \\
 &= \sum_{k=0}^{\infty} \gamma^k R_{t+k},
 \end{aligned} \tag{1}$$

where $R_t = \mathcal{R}(s_t, a_t)$ is the reward at time t .

Two key concepts in RL are the value function and the Q-function. The value function is the expected return of following a given policy starting in a state $V^\pi(s) = \mathbb{E}_\pi[G_t|s]$. The Q-function defines the value of taking an action in a state and then following a given policy $Q^\pi(s, a) = \mathbb{E}_\pi[G_t|s, a]$. An optimal policy maximizes either the value function or the Q-function. Significant advances in RL can be attributed to the integration of deep learning techniques to approximate the value function, the Q-function, or the policy.

Various types of RL algorithms have been developed to create policies for increasingly complex environments. In this case, the action is the opening or closing of the drainage valves (0%-100%) within the stormwater system based on measured water height at various aspects in and around the system when rainfall data was introduced. RL allows large numbers of simulations to be run to develop a generalized policy between many different conditions within the environment. In this case, the water height or total water volume above the system (flooding) was the outcome. Thus, it was not important to pre-determine the relationship or anticipate how the drains and valves might interact. RL tests thousands of policies and iteratively performs its adjustments based on the reward function during training and testing on an unseen dataset prevents overfitting the model or biasing results. Another aspect of RL is a concept of exploration. While the agent may learn some policy over time, without random exploration, it may learn a sub-optimal policy as the optimal may not even be discovered. To achieve this exploration, a random process generates the action of the agent for a certain amount of time during training. As training progresses, this random action becomes less common. One of the first types of deep RL algorithms was Q-learning, which allowed the creation of its own relationships among variables, rather than the programmer establishing a model [21] establishing a framework for future RL algorithms.

2.3. Deep Deterministic Policy Gradient

New RL techniques are capable of handling continuous states and actions, a requirement for this project. Deep Deterministic Policy Gradient (DDPG) is an Actor-Critic RL algorithm that combines some of the optimal characteristics of deep Q learning in a continuous action space. Although this algorithm was developed in 2015 and has been tested primarily on games – Atari games are the current industry standard for evaluating reinforcement learning [14] – DDPG has potential applicability for large, complex systems such as urban infrastructure because of its ability to integrate the continuous and streaming datasets with interactive variables. Interactions between SWMM and DDPG are processed through an interpreter [22], converting DDPG output into a context readable by SWMM and vice versa with SWMM output. This interpreter is loosely based around the OpenAI Gym environment [23], a standardized framework for RL and environment interaction and RL is implemented with the keras-rl python package [24]. The python PySWMM wrapper for SWMM allows for the SWMM model to be run incrementally, an essential requirement for RL [25].

DDPG can be divided into two major parts: the Actor network, $\mu(s|\theta^\mu)$ is a function which updates the current policy π mapping states to actions. The critic, $Q(s,a|\theta^Q)$ learns using the Bellman equation utilized by past RL such as Q learning [21] (see Figure 2).

Figure 2: DDPG model and its interactions with SWMM

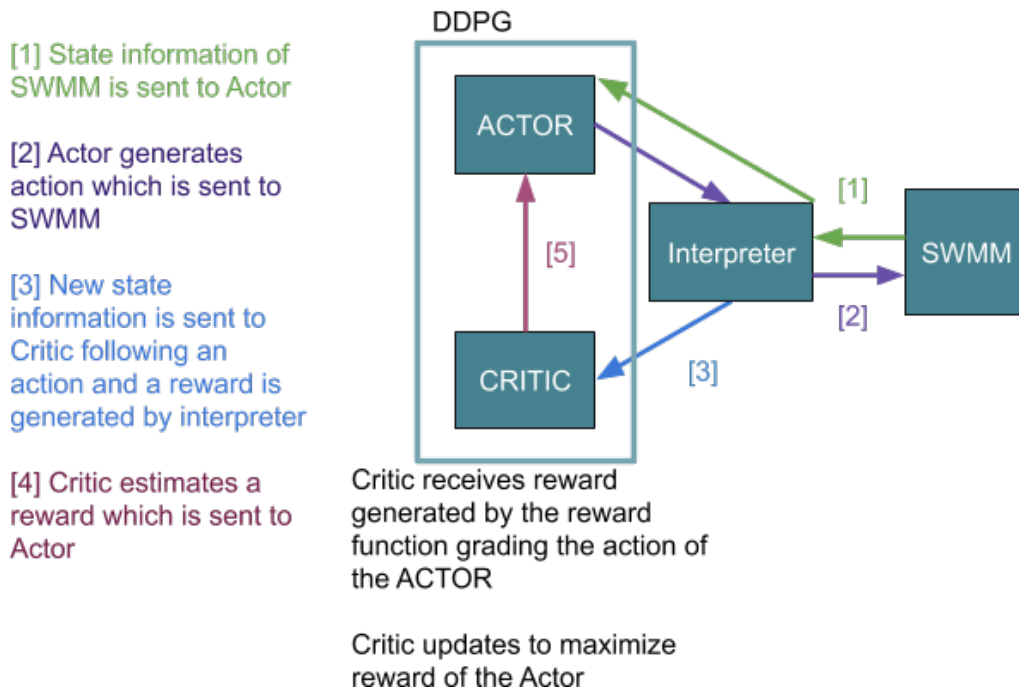


Figure 2. Communication between SWMM and DDPG

The actor and critic are initialized with weights θ^μ and θ^Q . With this, the target networks of Q' and μ' are also initialized with the weights of Q and μ creating $\theta^{\mu'}$ and $\theta^{Q'}$. Due to the constant updating of the actor and critic neural networks, and the actor being based on the critic's network, target networks are utilized to slowly update the networks. A replay buffer (\mathcal{B}) is also initialized. \mathcal{B} stores state, reward, and action information as the episode progresses and is sampled for training. \mathcal{B} and target networks improve the stability of DDPG as one imperfect event will not cause a drastic change to the model.

The simulations begin with an initial observation and a random process $\epsilon \sim \mathcal{N}(0, \sigma)$ for action exploration. As the simulations progress, an action is selected at each time step using the formula,

$$a_t = \mu(s|\theta^\mu) + \epsilon. \quad (2)$$

Then, for each step in the simulation: The action is performed resulting in a reward, and the next state, s_{t+1} is observed; all stored in \mathcal{B} .

As training begins, a random batch of s, a, r, s_{i+1} of size (N) is sampled from \mathcal{B} . Intermediary variable y_i , which is the expected future reward as predicted by the target networks, is created,

$$y_i = r_i + \gamma Q'(s_{i+1}, \mu'(s_{i+1}|\theta^{\mu'})|\theta^{Q'}). \quad (3)$$

The critic network is updated to minimize the loss; the sum of the difference between y_i and the critic with respect to the current state and actions where i represents the i th sample,

$$L(\theta^Q) = \frac{1}{N} \sum_i (y_i - Q(s_i, a_i|\theta^Q))^2. \quad (4)$$

The actor policy is updated using the sampled policy gradient of the critic with respect to the actions multiplied by the policy gradient of the actor with respect to the state. This is then averaged by the total size of the batch,

$$\nabla_{\theta^{\mu}} \approx \frac{1}{N} \sum_i \nabla_a Q(s, a | \theta^Q) |_{s=s_i, a=\mu(s_i)} \cdot \nabla_{\theta^{\mu}} \mu(s | \theta^{\mu}) |_{s_i}. \quad (5)$$

The weights of the target networks are then updated, discounted by learning rate multiplier τ ,

$$\theta^{Q'} \leftarrow \tau \theta^Q + (1 - \tau) \theta^{Q'}, \quad (6)$$

$$\theta^{\mu'} \leftarrow \tau \theta^{\mu} + (1 - \tau) \theta^{\mu'}. \quad (7)$$

This process iterates for the duration of the simulations.

2.4. MDP

The DDPG functioned on the following MDP state-action space:

1. The state is composed of several features defined below:

- the depth of the water at junction $i = \{1, 2, 3\}$ is defined as d_i ,
- the amount of flooding at junction $i = \{1, 2, 3\}$ is defined as f_i ,
- the current orifice setting at junction $i = \{1, 2, 3\}$ is defined as o_i ,
- the predicted rainfall within the current hour is defined as r_1 , and the predicted rainfall within the next hour is defined as r_2 .

$$s_t = [d_1, d_2, d_3, f_1, f_2, f_3, o_1, o_2, o_3, r_1, r_2]. \quad (8)$$

2. The action: at the next time step set valve position o_i to value $x \in [0, 1]$

3. The reward: $\mathcal{R}(F, D, E)$, where $F = f_1 + f_2 + f_3$, $D = d_1 + d_2 + d_3$, and E is the target depth of 0.6096

$$\mathcal{R}(F, D, E) = \begin{cases} -F^2 \cdot [\sigma(f_1), \sigma(f_2), \sigma(f_3)] & F > 0 \\ |D - E| & F = 0 \end{cases}$$

Where σ represents the sigmoid function:

$$\sigma = \frac{1}{1 + e^{-x}} \quad (9)$$

\mathcal{R} is a piece wise function, with flooding, the reward is the dot product of the average of the sigmoid of the flooding at different junctions multiplied by the negative sum of flooding. The sigmoid function utilized normalizes the total flooding amount at each node in a range of zero to one. This would cause the reward to be approaching the actual flooding value as flooding increases. An even more general multiplier for each node was created as the episode progresses due to the averaging. In other words, even if flooding was very large for a short period of the episode - potentially indicative of a flash flooding event - averaging it would not as negatively penalize the RL for unchangeable externalities such as an amount of rainfall that would greatly overwhelm any system passive or controlled. If there is no flooding, then the reward is dependent on the difference of water level from the expected height.

A secondary addition to the function is a minor reward related to the water depth and its difference from an expected depth (E) of 0.6096 meters, 50% of the junction max depth. In the context of real world applications, these junctions/storage ponds are not meant to run dry. As a result, this expected height would act as the intended safe operating depth for the pond. Application of this solution to a real world situation would require an examination of the results to promote greater distribution of water over the entire system where the total water volume could be managed. For the reward function, the sum of the total flooding F , and depths D are used.

After pilot testing, a reward function was selected for further training and testing. Pilot testing included five total rewards each evaluated against the passive and each other following 10,000 training steps with the same setup as this study. It was determined that a beneficial reward function should

include a dot product of an array method to prevent a policy that placed all of the flooding on one node. Although the total volume of water might be less, it would be a sub-optimal policy since one node would have a catastrophic amount of flooding.

2.5. Data

For experimentation, a rainfall data set was obtained from the mean of two gauges operated by the Hampton Roads Sanitation District in Norfolk, VA. This data set recorded rainfall in 15-minute increments from 2010 to 2018. From this data, 100 episodes in 24-hour periods were selected that were found to result in flooding using a passive system. This was done in order to reduce variability within the successes of the RL implementation - if the RL agent does not have to make any actions and no flooding occurs anyway, it takes longer for an optimal policy to be found.

2.6. Experimental Design

Training simulations were run 10,000 times, making measures and adjustments on a 15-minute basis, with each simulation corresponding to a 24 hour period within the data set. Training was repeated with each of the testing conditions. This means that in total, 240,000 hours were simulated for each condition, equating to a total of 16,000 15-minute periods.

For each condition, the episodes were processed into a simulated perfect forecast, giving the sum of the cm of rainfall received during each period in one hour increments. In other words, one value was created for each hour would represent the total rainfall for that particular hour. This was done for the current hour of the simulation, (f_1), and the next (f_2).

Forecasts were found to have an approximate 40% accuracy within a tertile of variation six days in advance, however the accuracy greatly increases as the timescale decreases (1 hour in advance increases accuracy) [26]. The perfect forecast generated has no imperfections, as such to evaluate more real world performance, imperfections are added. Thus to simulate imperfect forecasts, (f'_{1-2}), a uniform distribution (U) was sampled in order to generate a value between the upper (A) and lower bounds (B) of the accuracy,

$$X \sim U(A, B). \quad (10)$$

This uncertainty was simulated by the product of the perfect forecast and sampled value (X) with upper and lower bounds of 0.95 and 1.05 to create a possible spread of values in between +/- 5%.

$$f'_1 = f_1 * X \sim U(0.95, 1.05) \quad (11)$$

The bounds of the sample became greater the further the forecast was from the current timestep; an example being +/- 10% two hours in the future.

$$f'_2 = f_2 * X \sim U(0.90, 1.10) \quad (12)$$

During the same time period, the SWMM tool simulated water heights and flooding at each node.

For the perfect state data, the exact sensor and exact forecast data was used. For both the imperfect state and imperfect forecast data, the uniform sampling discussed above was utilized, however, additionally, a uniform distribution was again sampled +/- 10% for the state data to create the imperfect state condition.

A random sample of 15% of rain events was selected for test episodes (rainfall test episodes 1-15) to compare the performance of each trained RL to that of the passive condition. These test episodes were not used for training to prevent and test for overfitting - the policy is only optimal for the training data. The passive condition mimicked the traditional stormwater systems which are commonly used, and served as a control condition for the RL implementations.

As a starting parameter, $t = 0$, $o_i = 0.5$.

3. Results

3.1. Overall Comparison of RL Agents and Passive System

The performance of each DDPG implementation with perfect and noisy forecast and state information was evaluated against the passive system for the 15 randomly chosen test datasets (Table 2). The table lists the total flooding volume for each episode. Each implementation showed major reductions in overall flooding volume compared to the passive system (Figure 3). When compared to the passive system, total flood volumes had average reductions of 69.23%, 67.90%, and 73.02% for the RL implementations using perfect data, noisy forecast data, and noisy state data, respectively. The episodes 1, 5 and 10 were randomly selected to further explore the performance of the implementations on a step basis.

Table 2. Total flood volumes and percent change compared to the passive system for the 15 test events from each control method. Results in bold are further evaluated in Figures 5, 6, and 7. As shown in the lower table, each RL implementation has a mean positive reduction in total flood volume compared to the passive system.

Testing episode	Passive (m^3)	Perfect RL (m^3)	Noisy Forecast RL (m^3)	Noisy State RL
1	9,227,798	6,279,419	4,203,413	4,545,268
2	2,086,213	0	176,732	0
3	3,334,000	0	195,665	0
4	502,718	0	90,518	0
5	5,982,981	207,619	0	0
6	1,785,434	0	135,982	0
7	2,258,397	0	159,199	0
8	7,035,674	337,263	0	0
9	1,232,353	0	250,331	0
10	23,302,974	19,503,673	18,023,794	17,336,890
11	20,115,112	18,056,242	16,624,101	16,936,806
12	20,265,776	13,885,644	12,774,289	12,496,106
13	614,207	0	89,129	0
14	11,353,640	4,223,142	3,803,721	4,415,311
15	18,299,270	12,509,538	12,183,680	11,279,106
Passive Comparison		Perfect RL	Noisy Forecast RL	Noisy State RL
Mean % reduction		69.23%	67.90%	73.02%
Median % reduction		96.53%	81.99%	100.00%
Standard deviation % reduction		36.35%	29.34%	32.77%




Figure 3. Percent reduction of flooding volume of the RL implementations compared to the passive condition. The similar performance of the implementations result in the range of the mean % reductions to be within 4%.

3.2. Analysis of RL Agent Policies

Of the 15 test rainfall events used to evaluate the performance of the RL implementation, three events were randomly selected as a sample for further analysis. The rain events for testing episodes one, five, and ten are displayed in figure 4 .




Figure 4. Rainfall events further analyzed: Test Episode 1 - short duration high intensity storm peaking at 1.9177 cm of rainfall in a 15 minute period. Test Episode 5 - long duration low intensity storm peaking at 0.2667 cm of rainfall in a 15 minute period. Test Episode 10 - long duration low intensity storm peaking at 0.3175 cm of rainfall in a 15 minute period.

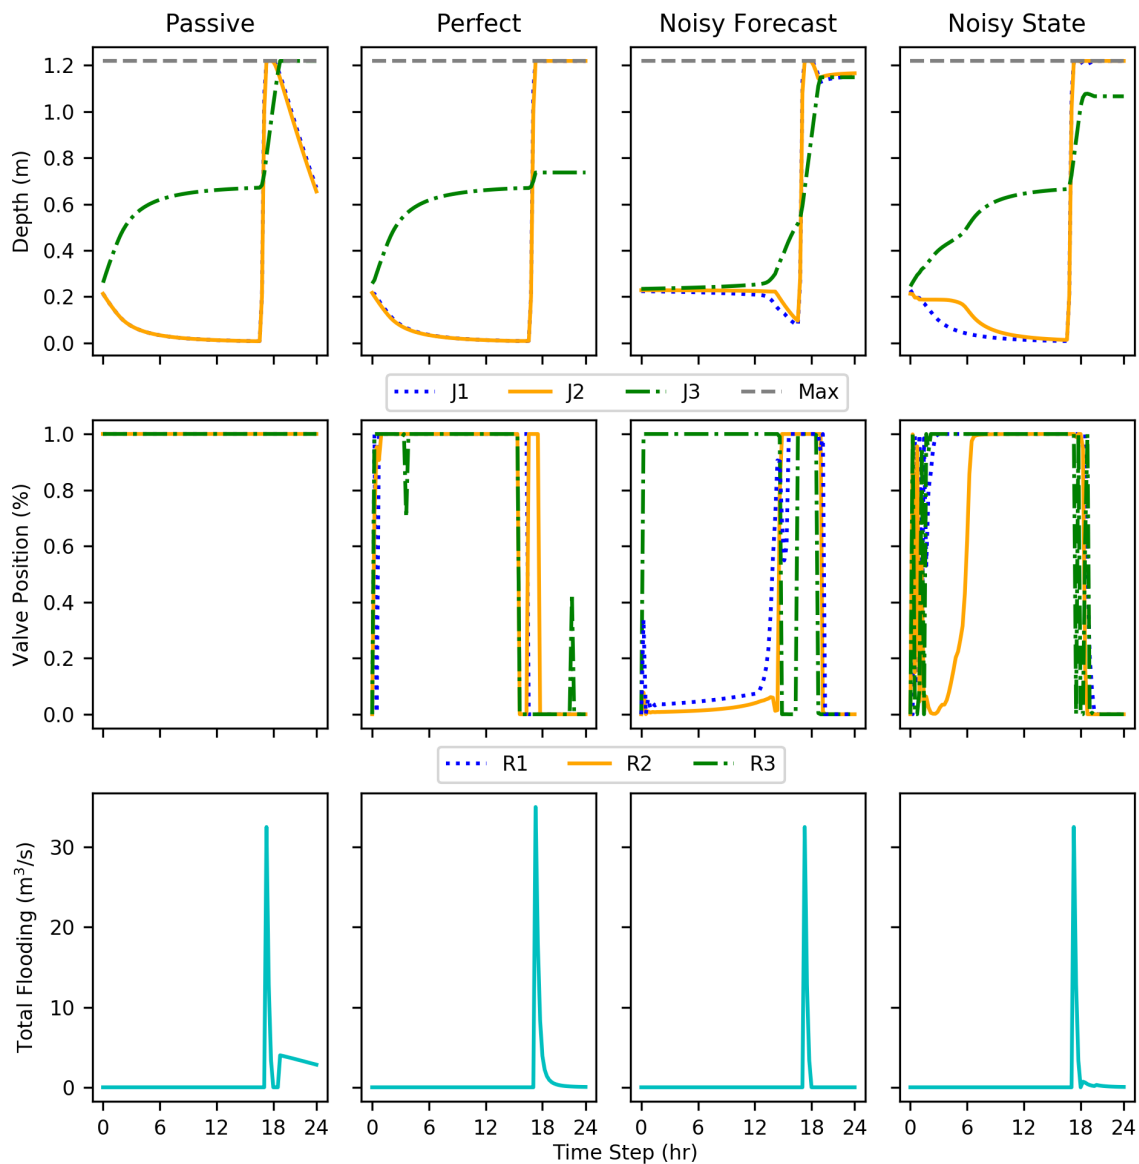


Figure 5. Comparison of the passive system to the RL implementations (episode 1). Each implementation preformed better than the passive condition. In this instance the Perfect RL preformed worse than the other two implementations only resulting in a 32% reduction opposed to the ~50% reduction of the other two.

For test episode 1 where there was approximately 2 cm of rainfall within a 15 minute period, all stormwater implementations were quickly overwhelmed. The passive system resulting in $9 \times 10^6 m^3$ of flooding, preformed the worst. As all of the valves are open, J_1 and J_2 do not see substantial volume, however J_3 is quickly overwhelmed resulting in the large flooding amounts. The RL implementation using perfect data was able to modulate the valves, resulting in a 32% reduction. Opposed to the passive system, the perfect RL caused flooding at J_1 and J_2 while keeping J_3 below flood stage. In a real world context, while still causing flooding, it would result in less damage distributing the flood volume across two different areas as opposed to $9 \times 10^6 m^3$ at one. Both of the noisy data variants preformed similarly resulting in a 50+% reduction. Similar to the perfect data implementation, both kept J_3 below flood stage, however they achieved this - noisy state especially - by rapidly cycling the orifices from 0 - 100%. While in a simulated environment this may be optimal, in a real world context, this may not be a desirable behavior. Further research, potentially with another reward function, would be needed to avoid this behavior.

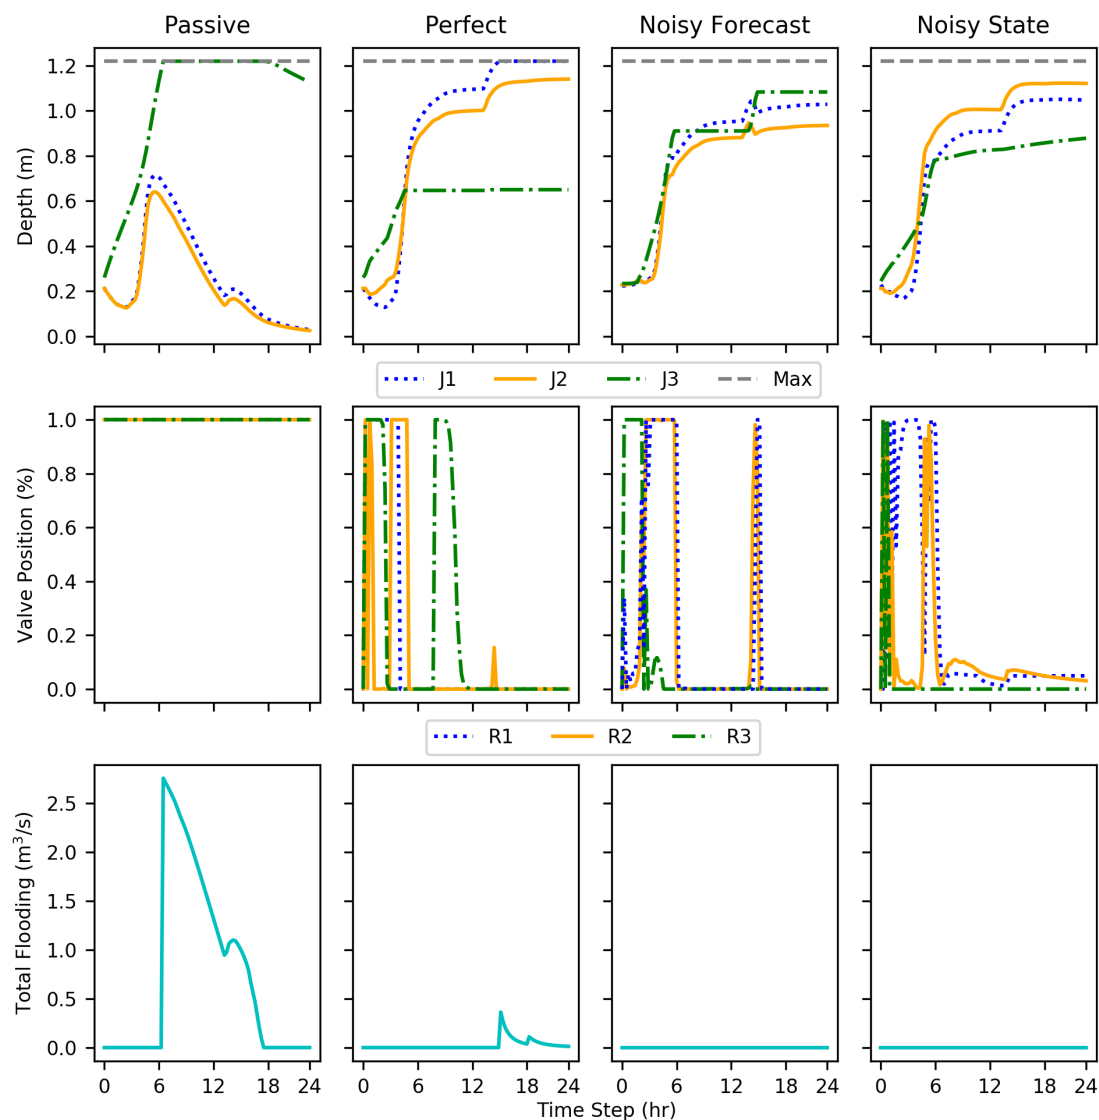


Figure 6. Comparison of the passive system to the RL implementations (episode 5). All implementations are able to reduce flooding, yet reducing a lesser amount than the first test episode.

For test episode 5, there was peak rainfall of approximately 0.4 cm of rainfall over a 15 minute period, however there was a continuous, yet minor rainfall for 24 hour period. While a seemingly small quantity of water, multiplied across the area of eight subcatchments, the perfect implementation was overwhelmed. Similar to the first episode, J_{1-2} were well below flood stage, however, J_3 flooded. In the case of this episode, flooding occurred for nearly 12 of the 24 hours resulting in $6 \times 10^6 m^3$, the passive performing the worst. The perfect data implementation was able to modulate the valves, resulting in a 96% reduction. Similarly to the first test episode, the perfect data implementation caused flooding at J_1 and J_2 while keeping J_3 below flood stage; again distributing the flooding volume to two areas. Both of the noisy data variants preformed similarly resulting in a 100% reduction. Each case was successfully able to keep each junction below flood stage.

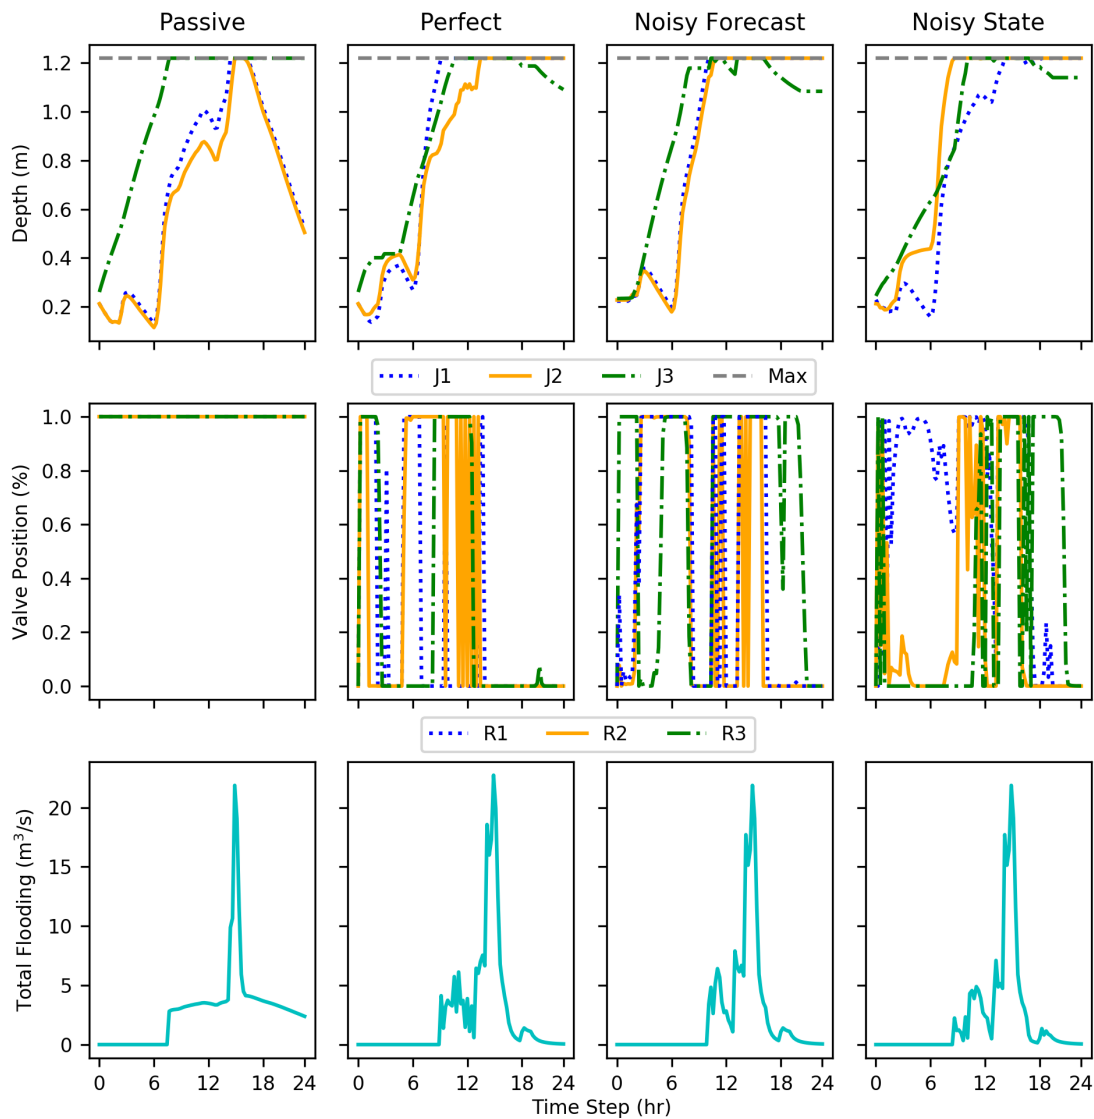


Figure 7. Comparison of the passive system to the RL implementation (episode 10). With this test episode, each implementation shows rapid cycling between 0 and 100% open. All show a reduction compared to the passive condition.

For test episode 10, there was a moderate intensity storm peaking at at 0.3 cm per 15 minutes. The storm fluctuated around that volume for the duration of the 24 hours. The passive system resulting in over $23 \times 10^6 m^3$ of flooding, performed the worst. All junctions had flooding, J_3 the most. Opposed to the other test episodes, all of the RL implementations resulted in flooding. The junctions are all overwhelmed in a similar way, yet minor differences result in differences in flooding volume. The perfect data implementation performs the worst with a 16% reduction, and the imperfect state resulting in a 25% reduction, a range of 9%.

4. Discussion

The DDPG implementations, even with added noise, showed reductions in flooding compared to the passive conditions, evaluated on 15 test conditions. The tests show that DDPG learned a policy that reduced total flood volume by an average of 69-73% compared to the passive system and several test events had 100% reductions. The implementations were evaluated using disadvantages more realistic in real world conditions such as imperfect forecast and state data. Nevertheless, each DDPG

implementation performed similarly to the others. This shows the ability to account for and even learn how to use data formerly unused due to unpredictability and uncertainty, such as forecasting.

Within the results, each implementation at some point showed rapid cycling between valve positions, this is particularly evident with test episode 10 (Figure 7). This action may be indicative of incomplete training, a fully optimal policy not being completely realized.

While only resulting in an average decrease in flooding volume by an additional one percent, the noisy state implementation performed better than the perfect condition. Without more extensive testing, any explanation may be speculation, but, the better performance can most likely be attributed to further exploration of the state space due to increased irregularity with the state information. Another explanation for the irregularity is simply variation within training. While ideally an optimal policy can be learned every time, there is an aspect of random exploration to RL.

This project acted as a first step showing that DDPG can perform optimally even with imperfect data. However, for an implementation such as this to be utilized in the real world, there are still a multitude of factors such as the consequences of control system or sensor failure need further exploration.

While imperfect data is more representative of reality and can often lead to unpredictability within a system, the loss or removal of data can have even graver consequences. Opposed to if a human were to look at operating levels of something such as a storage pond and could realize highly improbable values such as a depth of zero, a machine without vast training on a wide array of issues may never realize there is an error. RL or DDPG solutions are currently limited by what can be simulated. Without the knowledge of what can fail and how, one can only speculate as to the ramifications of a system continuously flooding a storage pond as it thinks it is below standard operating height.

Reinforcement learning is an evolving field, and computer systems are able to handle larger datasets. A multitude of factors go into the performance of a reinforcement learning implementation, DDPG being no exception [14]. This study utilized a reward function selected following pilot testing of 5 others; each using different methods of administering a negative reward. Other formulae may have produced less total flooding volume, but potentially may be catastrophic based on the concentration of where the flooding occurred. However, without further area specific evaluation; this remains speculation. The reward function chosen for further training and testing would mitigate flooding, but may have unforeseen consequences based on the current infrastructure model. Additional training cycles would likely improve the results.

The reported rainfall amounts and subsequent flooding, although simulated, were representative of historical events in Norfolk, Virginia. Norfolk has been especially affected by changing weather patterns and climate change due to a rising sea level and development in various surrounding areas. In cases where seawater rises impact stormwater outlets, even minor rain events can lead to so-called nuisance floods [27]. DDPG with the tested reward function, regardless of noise added showed the greatest reduction in flooding during these minor rainfall events, situations when nuisance floods would occur. Thus, communities such as Norfolk could benefit greatly from RL solutions. The SWMM tool helps with the design factors, but the ability to learn how to best adjust the valves during rain events could continually improve.

RL application to mitigate flooding in coastal cities is an emerging strategy for this important problem. Using authentic data sets, with the introduction of noise provides a realistic simulation of stormwater systems. DDPG was able to accommodate a complex model considering an integration of rain forecasting, each with imperfect or noisy data. DDPG and simulations can be adapted with different rainfall datasets and models indicative of other cities and as development of infrastructure continues. Thus, one could relatively quickly retrain the algorithm based on the new environment, and hope to see similar results. The application of DDPG, and, to a greater extent reinforcement learning, has application in a wide variety of civil infrastructure systems, not limited to just stormwater infrastructure. Using DDPG and other machine learning strategies to control civil structures has the ability to process multiple, simultaneous factors in near real time following sufficient training.

Furthermore, this paper describes the ability of the machine to accommodate imperfections that represent real-world factors and can reduce the loss of life and property by dynamically automating stormwater systems.

5. Conclusion

This study demonstrated the viability and success of a DDPG RL implementation in controlling stormwater systems with a reduction of flooding in a variety of storm events observed in Norfolk, VA. Three stormwater control conditions were compared with the passive system: (i) RL using perfect input data, (ii) RL using noisy forecasts, and (iii) RL using noisy state data. We hypothesized that RL would provide better flood mitigation than the passive system by accommodating for noisy forecast or state data that is more representative of the real world than having perfect knowledge of current and future states. Three DDPG agents were trained by interacting with a SWMM simulation for 10,000 time steps. Following training, the RL policies were tested on a set of previously unseen rainfall events. We found that the DDPG strategies were able to accommodate imperfect data resulting in similar flood mitigation ability between the three data conditions. Compared to the passive system, the three RL implementations reduced flooding by an average of 70.5% on the test events.

The results of this research demonstrate the promise of using RL to control stormwater systems in real-time, even under imperfect conditions. Additional work to further understand how RL performs for stormwater RTC includes using real rainfall forecasts that are longer than the forecasts used here. This would allow control decisions to be made even further in advance of storm events to prepare the system to prevent flooding. Longer forecasts, however, may include more uncertainty and increase the chance of making inappropriate control decisions. Another key topic to advance this research is to explore RL's ability to maintain overall system performance when a sensor or group of sensors is not functioning or reporting erroneous data. This will be especially critical as both the density of sensor networks and the adoption of stormwater RTC systems continues to increase [9].

Author Contributions: Conceptualization, Sami Saliba, Benjamin Bowes, Peter Beling and Jonathan Goodall; Data curation, Sami Saliba and Benjamin Bowes; Formal analysis, Sami Saliba; Funding acquisition, Jonathan Goodall; Investigation, Sami Saliba and Benjamin Bowes; Methodology, Sami Saliba, Benjamin Bowes and Stephen Adams; Project administration, Sami Saliba and Benjamin Bowes; Resources, Sami Saliba; Software, Sami Saliba; Supervision, Benjamin Bowes and Stephen Adams; Validation, Sami Saliba and Benjamin Bowes; Visualization, Sami Saliba and Benjamin Bowes; Writing – original draft, Sami Saliba; Writing – review editing, Sami Saliba, Benjamin Bowes, Stephen Adams, Peter Beling and Jonathan Goodall.

Funding: This work was funded as part of two U.S. National Science Foundation grants: Award #1735587 (CRISP-Critical, Resilient Interdependent Infrastructure Systems and Processes) and Award #1737432 (SCC-IRG Track 1: Overcoming Social and Technical Barriers for the Broad Adoption of Smart Stormwater Systems).

Acknowledgments: We gratefully acknowledge the Hampton Roads Sanitation District for access to their high quality rainfall data.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

RTC	Real Time Control
RL	Reinforcement Learning
DDPG	Deep Deterministic Policy Gradients
RL	Reinforcement Learning
SWMM	Storm Water Management Model
S	Subcatchment
J	Junction
R	Orifice Valve
MDP	Markov Decision Process
S	State
A	Action
P	Transition Function
γ	Discount factor
s	Specific State
a	Specific Action
o	Orifice Valve Setting
f	Forecast
D	Depths
E	Expected Depth
f'	Imperfect Forecast

References

1. Ezer, T.; Atkinson, L.P. Accelerated flooding along the US East Coast: On the impact of sea-level rise, tides, storms, the Gulf Stream, and the North Atlantic Oscillations. *Earth's Future* **2014**, *2*, 362–382.
2. O'Driscoll, M.; Clinton, S.; Jefferson, A.; Manda, A.; McMillan, S. Urbanization effects on watershed hydrology and in-stream processes in the southern United States. *Water* **2010**, *2*, 605–648.
3. Sweet, W.V.; Park, J. From the extreme to the mean: Acceleration and tipping points of coastal inundation from sea level rise. *Earth's Future* **2014**, *2*, 579–600.
4. Sörensen, J.; Persson, A.; Sternudd, C.; Aspegren, H.; Nilsson, J.; Nordström, J.; Jönsson, K.; Mottaghi, M.; Becker, P.; Pilesjö, P.; others. Re-thinking urban flood management—Time for a regime shift. *Water* **2016**, *8*, 332.
5. Smith, A.B.; Katz, R.W. US billion-dollar weather and climate disasters: data sources, trends, accuracy and biases. *Natural hazards* **2013**, *67*, 387–410.
6. Sadler, J.; Goodall, J.; Morsy, M.; Spencer, K. Modeling urban coastal flood severity from crowd-sourced flood reports using Poisson regression and Random Forest. *Journal of hydrology* **2018**, *559*, 43–55.
7. Sadler, J.M.; Haselden, N.; Mellon, K.; Hackel, A.; Son, V.; Mayfield, J.; Blase, A.; Goodall, J.L. Impact of sea-level rise on roadway flooding in the Hampton Roads region, Virginia. *Journal of Infrastructure Systems* **2017**, *23*, 05017006.
8. Kerkez, B.; Gruden, C.; Lewis, M.; Montestruque, L.; Quigley, M.; Wong, B.; Bedig, A.; Kertesz, R.; Braun, T.; Cadwalader, O.; others. Smarter stormwater systems, 2016.
9. Marchese, D.; Jin, A.; Fox-Lent, C.; Linkov, I. Resilience for Smart Water Systems. *Journal of Water Resources Planning and Management* **2020**, *146*, 02519002.
10. Sutton, R.S.; Barto, A.G. Reinforcement learning, a Bradford book. *MIT Press Cambridge, MA* **1998**, 2015, 2.
11. Mullapudi, A.; Lewis, M.J.; Gruden, C.L.; Kerkez, B. Deep reinforcement learning for the real time control of stormwater systems. *Advances in Water Resources* **2020**, p. 103600.
12. Wang, C.; Bowes, B.; Tavakoli, A.; Adams, S.; Goodall, J.; Beling, P. Smart Stormwater Control Systems: A Reinforcement Learning Approach. ISCRAM 2020 Conference Proceedings - 17th International Conference on Information Systems for Crisis Response and Management; Hughes, A.L.; McNeill, F.; Zobel, C., Eds.; , 2020; pp. 2–13.
13. Jafari, F.; Mousavi, S.J.; Yazdi, J.; Kim, J.H. Real-time operation of pumping systems for urban flood mitigation: single-period vs. multi-period optimization. *Water Resources Management* **2018**, *32*, 4643–4660.
14. Lillicrap, T.P.; Hunt, J.J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; Wierstra, D. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971* **2015**.

15. Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; others. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* **2018**, *362*, 1140–1144.
16. Bowes, B.D.; Tavakoli, A.; Wang, C.; Heydarian, A.; Behl, M.; Beling, P.A.; Goodall, J.L. Flood Mitigation in Coastal Urban Catchments using Real-time Stormwater Infrastructure Control and Reinforcement Learning. *Journal of Hydroinformatics* **in press**.
17. Rayner, S.; Lach, D.; Ingram, H. Weather forecasts are for wimps: why water resource managers do not use climate forecasts. *Climatic change* **2005**, *69*, 197–227.
18. Hartono, P.; Hashimoto, S. Learning from imperfect data. *Applied Soft Computing* **2007**, *7*, 353–363.
19. Gironás, J.; Roesner, L.A.; Rossman, L.A.; Davis, J. A new applications manual for the Storm Water Management Model(SWMM). *Environmental Modelling & Software* **2010**, *25*, 813–814.
20. Jang, S.; Cho, M.; Yoon, J.; Yoon, Y.; Kim, S.; Kim, G.; Kim, L.; Aksoy, H. Using SWMM as a tool for hydrologic impact assessment. *Desalination* **2007**, *212*, 344–356.
21. Gu, S.; Lillicrap, T.; Sutskever, I.; Levine, S. Continuous deep q-learning with model-based acceleration. International Conference on Machine Learning, 2016, pp. 2829–2838.
22. Choo, B.; Crannel, G.; Adams, S.; Dadgostari, F.; Beling, P.A.; Bolcavage, A.; McIntyre, R. Reinforcement learning from simulated environments: an encoder decoder framework. 2020 Spring Simulation Conference (SpringSim). IEEE, 2020, pp. 1–12.
23. Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; Zaremba, W. OpenAI Gym, 2016, [[1606.01540](https://arxiv.org/abs/1606.01540)].
24. Plappert, M. keras-rl, 2016.
25. McDonnell, B.E.; Ratliff, K.; Tryby, M.E.; Wu, J.J.X.; Mullanpudi, A. PySWMM: The Python Interface to Stormwater Management Model (SWMM). *Journal of Open Source Software* **2020**, *5*, 2292.
26. Hu, Q.S.; Skaggs, K. Accuracy of 6-10 day precipitation forecasts and its improvement in the past six years. 7th NOAA annual climate prediction application science workshop, 2009.
27. Sheridan, S.C.; Pirhalla, D.E.; Lee, C.C.; Ransibrahmanakul, V. Atmospheric drivers of sea-level fluctuations and nuisance floods along the mid-Atlantic coast of the USA. *Regional Environmental Change* **2017**, *17*, 1853–1861.