Climate Adaptation with Reinforcement Learning: Experiments with Flooding and Transportation in Copenhagen

Miguel Costa, Morten W. Petersen, Arthur Vandervoort,
Martin Drews, Karyn Morrissey, Francisco C. Pereira
Department of Technology, Management and Economics
Technical University of Denmark
2800, Kgs. Lyngby, Denmark
{migcos, s184985, apiva, mard, kamorr, camara}@dtu.dk

Abstract

Due to climate change the frequency and intensity of extreme rainfall events, which contribute to urban flooding, are expected to increase in many places. These floods can damage transport infrastructure and disrupt mobility, highlighting the need for cities to adapt to escalating risks. Reinforcement learning (RL) serves as a powerful tool for uncovering optimal adaptation strategies, determining how and where to deploy adaptation measures effectively, even under significant uncertainty. In this study, we leverage RL to identify the most effective timing and locations for implementing measures, aiming to reduce both direct and indirect impacts of flooding. Our framework integrates climate change projections of future rainfall events and floods, models city-wide motorized trips, and quantifies direct and indirect impacts on infrastructure and mobility. Preliminary results suggest that our RL-based approach can significantly enhance decision-making by prioritizing interventions in specific urban areas and identifying the optimal periods for their implementation. Our framework is publicly available: `https://github.com/MLSM-at-DTU/floods_transport_rl`.

1 Introduction

As climate change continues to impact our world, the frequency and intensity of high impact weather events are expected to rise [1]. In Denmark, extreme rainfall is projected to become more severe and occur more frequently [2]. As rainfall increases, so does the risk of urban pluvial flooding. Floods can significantly disrupt social and economic activities, including transportation, causing delays and loss of vehicle control [3]. To effectively address these challenges, cities must enhance their resilience.

In this work, we address the challenge of identifying the most effective adaptation measures to minimize the impacts of pluvial floods on transportation. Using Copenhagen, Denmark, as our case study, we frame our problem using a reinforcement learning (RL) approach. We build an environment that incorporates current climate projections of rainfall and consequent floods. Concurrently, we model trips which are disrupted by varying water levels, affecting mobility and transport infrastructure. To our knowledge, this is the first comprehensive framework designed to identify the best adaptation measures for enhancing transportation resilience to urban floods using such approach.

1.1 Related Work

Urban pluvial flooding occurs when large volumes of water accumulate in streets or roads due to insufficient drainage and infiltration capacity as a result of heavy precipitation events [4]. Transporta-

Tackling Climate Change with Machine Learning workshop at NeurIPS 2024

tion is significantly impacted by such floods, both directly and indirectly [5]. These impacts include road deterioration [5], travel delays and congestion [6, 7], and loss of accessibility [8, 9].

RL has previously been applied to a few aspects of flood management and transportation. For instance, it has been used to design emergency routing systems [10], control urban drainage and stormwater systems [11, 12], and study travel behaviors to inform response strategies [13]. However, these applications generally focus on reactive strategies – responding to events as they occur – rather than proactively determining the best adaptation measures to minimize future flood impacts.

On the other hand, evaluating adaptation measures typically relies on expert knowledge or limited simulations, assessing how these measures can prevent or mitigate the impacts of floods [14, 15, 16, 17]. To the best of our knowledge, there has been no comprehensive study using reinforcement learning of how different adaptation measures should be implemented over time to proactively minimize both the direct and indirect impacts of floods on transportation.

2 Modelling Framework

We frame our approach as an Integrated Assessment Model (IAM) that connects: 1) a rainfall projection model, 2) a flood model, 3) a transportation model, and 4) a transport infrastructure and mobility impact model. Figure 1 provides an overview of our IAM framework, which we now detail.

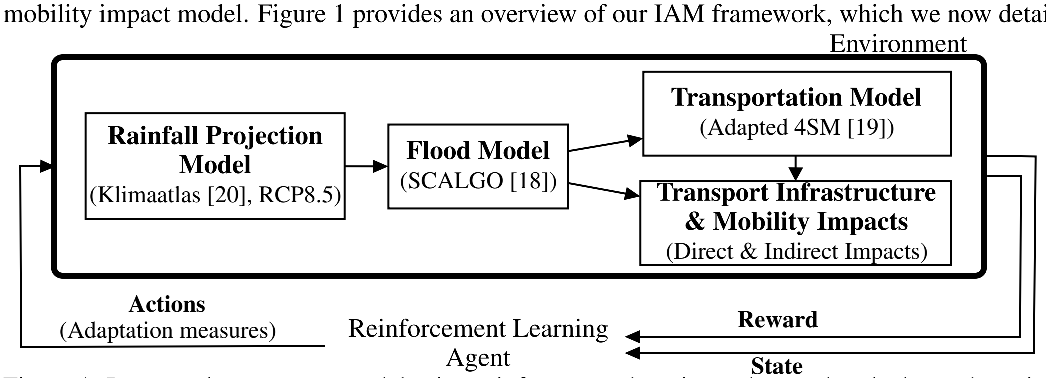

Figure 1: Integrated assessment model using reinforcement learning to learn what the best adaptation measures are that minimize transportation infrastructure and mobility impacts.

2.1 Rainfall Projection Model

Future daily rainfall statistics under the high RCP8.5 scenario [21] were retrieved from the Danish Meteorological Institute's Climate Atlas [20] for the periods 2011-2040, 2041-2070, and 2071-2100. For each time slice, we assumed stationarity and formed the associated cumulative density function (CDF). Based on the CDF we sampled one heavy rainfall event per year. Urban pluvial flooding is often caused by intense precipitation of short duration (cloudbursts, from minutes to a few hours), typically associated with warm and moist conditions. For simplicity and as proof-of-concept, we assumed the projected heavy rainfall intensity (amount of precipitation) to be equal to the accumulated daily rainfall. Jointly with the choice of climate scenario, the resulting CDF is likely to overestimate the rainfall intensities and therefore represent a worst-case scenario for our periodical estimates.

2.2 Flood Model

After sampling a particular rainfall event (i.e., amount of rainfall), we model the associated urban flood using SCALGO Live [18]. SCALGO Live is a simplified interactive event-based tool for watershed delineation, and for fast modelling of flood depths and flow direction based on high-resolution digital terrain data. For Denmark, SCALGO employs the Danish Elevation Model, which is one of the world's best national elevation models. It comprises 415 billion point data, which are used to map height differences for terrains and areas in a 0.4m grid for the entire country [22]. The model does not include a representation of the urban drainage system. We assumed a uniform rainfall event all over Copenhagen of unspecified duration, i.e., the water accumulates at all locations at once. Water is further distributed according to the terrain properties and filled any depressions or holes. If the volume of water exceeded the depression volume, it overflowed and continued downstream. In sum, the accumulation of heavy precipitation was mapped to water depths all around Copenhagen for identification of flooded areas.

2.3 Transportation Model

For the transportation component, we used a simplified version of the popular Four Step Model (4SM) [19], focusing exclusively on road network and motorized trips. We began by dividing Copenhagen in Traffic Analysis Zones (TAZs) following the Danish National Transport Model [23]. Then, we generated and distributed trips following the distribution of trips in Copenhagen [24], which were aggregated within each TAZ to map the supply and demand for each zone. In essence, this distribution reflects the underlying travel demand. Supply and demand distribution marginals were assigned used an iterative proportional fitting procedure [25] with distance as a travel impedance. Lastly, we mapped routes between TAZs. TAZs' centroids were defined as nodes and edges were created between neighboring TAZs. Paths between origin and destination TAZs were then defined as the shortest travel time paths for each trip, which can be found using Dijkstra's algorithm [26]. We used this network to route all trips and estimate volumes and travel times.

2.4 Transportation and Mobility Impacts

Finally, we computed transport and mobility impacts as three types of impacts: direct road infrastructure damage impacts and indirect impacts due to increase travel delays.

Road infrastructure damage: We began by downloading road network data from OpenStreetMap [27] using `osmnx` [28] from Copenhagen and computing their total construction costs per road type, number of lanes, presence of light posts and traffic lights van Ginkel et al. [29]. Damage was then computed using depth-damage functions [29], effectively mapping the percentage of damage on a road according to the water depth at its location. This damage accounts for reconstruction, repair, cleaning, and resurfacing works needed to restore roads to their original state. We aggregate direct impacts as the monetary losses at the i-th TAZ as R_i.

Travel delays: As water levels increase, travel speed is reduced, resulting in travel delays. To account for these effects we used a depth-disruption function [30], mapping decreased vehicle speeds to water depth. The speed reduction and consequent increased travel times were then modelled as economic losses using the danish travel delays value of time [31], which we aggregated as D_i for each TAZ.

2.5 Reinforcement Learning

Under the current climate uncertainty, we posit to uncover the best adaptation measures that minimize the impact of flooding events on transportation using RL. RL is a sub-field of machine learning that uses an agent-based approach to interact with an environment and achieve a certain goal [32]. Through training and by maximizing a pre-defined function (reward), the agent learns what is the best action (adaptation measure) to take. By default, the environment (as defined by the above IAM) is defined as a Markov Decision Process [33], where each state is independent.

Albeit many adaptation measures can be devised, in this first work, we defined one possible measure: elevate roads by 1 meter (i.e., increase the minimum water depth needed to affect roads). At each time step, our RL agent takes an action on a TAZ and and collects information about the state of Copenhagen (e.g., precipitation event, period of time, direct and indirect impacts per TAZ, water depths on roads), effectively learning the best set of actions to take over time and space. We defined the reward function as an overall metric of economic loss, defined as:

$$R = \sum_{i \in \texttt{TAZ}} \beta_R R_i + \beta_D D_i + \beta_A A_i \tag{1}$$

where R_i, D_i are as previously defined, A_i is the cost of applying an action (i.e., cost of elevating roads by 1 meter), and weights β ($\beta_R = \beta_D = \beta_A = 1$) adjust for different component importances.

3 Preliminary Experiments & Discussion

We setup our IAM using Python, the Gymnasium interface [34], Stable-Baseline3 [35], and PPO [36, 37]. To showcase the application of our approach, we experiment with a preliminary case study. We begin by running our experiments in the city center of Copenhagen, consisting of 29 TAZ, and set the time horizon to 2023–2100. We now present preliminary results for 10 runs with distinct seeds to allow for different weather projections and increased robustness.

Table 1: Differences between average reward, action costs, and impacts between random and learned policies. Two time horizons are presented: 2023–2035 coinciding with Copehagen's current climate adaptation plan [38] and until 2100. Values represent the mean ± standard deviation across 10 runs.

	2023-2035		2023-2100	
	Random Policy	Optimal Policy	Random Policy	Optimal Policy
Reward (M DKK) ↑	**-49.19** ± 4.91	**-36.48** ± 3.02	**-92.09** ± 20.72	**-69.56** ± 2.64
Cumulative Cost of Measures (M DKK) ↓	**20.08** ± 2.35	**22.93** ± 1.35	**52.94** ± 0.00	**51.79** ± 1.47
Cumulative Cost of Impacts (M DKK) ↓	**29.10** ± 4.07	**13.56** ± 2.32	**39.15** ± 20.72	**17.77** ± 2.12
Cumulative travel delays (k h) ↓	**5.17** ± 0.79	**2.15** ± 0.46	**6.91** ± 3.98	**2.67** ± 0.39

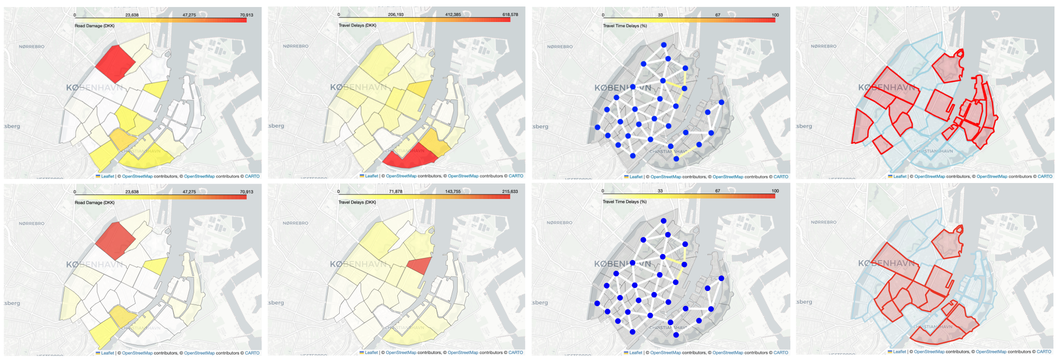

Figure 2: Costs of floods impacts on transportation and mobility in Copenhagen in 2035. Top row shows results with random adaptation measures deployed over time and space, while bottom row shows impacts using optimal adaptations over time. From left to right: direct road infrastructure impacts, indirect impacts as travel delays, percentage of travel time difference for travel between TAZ, and where adaptation measures were deployed (red).

Table 1 compares the performance of the trained RL agent against a random agent. The results demonstrate that our agent consistently outperforms the random policy, achieving significantly better outcomes overall. By 2035, our agent incurs additional costs but achieves a substantial reduction in direct and indirect impacts, lowering them by 47%. By 2100, although both policies result in similar adaptation costs, our agent's strategic deployment of optimal measures over time leads to significantly reduced travel delays (by 39%) and impacts (by 45%). Figure 2 illustrates the impacts in 2035 for a single run. As shown, at this point, our agent has implemented specific adaptation measures in certain TAZs, resulting in lower road damages and reduced travel delays. This highlights the agent's ability to prioritize interventions that would otherwise lead to greater losses.

These results underscore the effectiveness of using RL to identify optimal adaptation measures for Copenhagen over time, enhancing the city's ability to address climate change more efficiently. The proposed IAM introduces a novel framework for accurately simulating future rainfall, subsequent pluvial urban floods, and their impacts on transportation and mobility. Looking ahead, we suggest that this approach could be valuable for researchers and authorities in making more efficient and informed decisions and improving urban resilience.

In future work, we plan to further develop our IAM by extending the rainfall projection model, and by expanding the case study to encompass the entire city of Copenhagen and its metropolitan area. This expansion would include additional adaptation measures (e.g., constructing permeable roads or enhancing road durability), other modes of transport (e.g., cycling and walking), broader impact categories (e.g., electric vehicle charging infrastructure, public transportation accessibility, and subjective wellbeing), and comparative analyses (e.g., expert knowledge or participatory design). Additionally, we aim to enhance our transportation simulation by more accurately modeling supply and demand dynamics and their fluctuations during flood events, which can lead to trip cancellations and increased congestion levels [3]. Including such enhancements can further refine our IAM, making it a more comprehensive and practical tool for urban resilience planning in the face of climate change.

Acknowledgments and Disclosure of Funding

This work was supported by a research grant (VIL57387) from VILLUM FONDEN.

References

[1] IPCC, "Section 3: Long-Term Climate and Development Futures," in *Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, H. Lee and J. Romero (eds.)]*. IPCC, Geneva, Switzerland, doi: 10.59327/IPCC/AR6-9789291691647, 2023, pp. 35–115.

[2] Danmarks Meteorologiske Institut, "Adaptation to the future climate in Denmark," https://en.klimatilpasning.dk/media/7863/klimatilpasningshÃęfte%20uk%20web.pdf, 2011, accessed: 2024-06-14.

[3] T. Wang, Z. Qu, Z. Yang, T. Nichol, G. Clarke, and Y.-E. Ge, "Climate change research on transportation systems: Climate risks, adaptation and planning," *Transportation Research Part D: Transport and Environment*, vol. 88, p. 102553, 2020.

[4] M. Borowska-Stefańska, A. Bartnik, M. A. Dulebenets, M. Kowalski, A. Sahebgharani, P. Tomalski, and S. Wiśniewski, "Changes in intra-city transport accessibility accompanying the occurrence of an urban flood," *Transportation Research Part D: Transport and Environment*, vol. 126, p. 104040, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1361920923004376

[5] X. Lu, F. K. Shun Chan, W.-Q. Chen, H. K. Chan, and X. Gu, "An overview of flood-induced transport disruptions on urban streets and roads in chinese megacities: Lessons and future agendas," *Journal of Environmental Management*, vol. 321, p. 115991, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S030147972201564X

[6] F. J. Shahdani, M. Santamaria-Ariza, H. S. Sousa, M. Coelho, and J. C. Matos, "Assessing Flood Indirect Impacts on Road Transport Networks Applying Mesoscopic Traffic Modelling: The Case Study of Santarém, Portugal," *Applied Sciences*, vol. 12, no. 6, 2022. [Online]. Available: https://www.mdpi.com/2076-3417/12/6/3076

[7] W. Ding and J. Wu, "Interregional economic impacts of an extreme storm flood scenario considering transportation interruption: A case study of Shanghai, China," *Sustainable Cities and Society*, vol. 88, p. 104296, Jan. 2023. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S221067072200600X

[8] M. Li, M.-P. Kwan, J. Yin, D. Yu, and J. Wang, "The potential effect of a 100-year pluvial flood event on metro accessibility and ridership: A case study of central Shanghai, China," *Applied Geography*, vol. 100, pp. 21–29, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0143622818302716

[9] T. Papilloud and M. Keiler, "Vulnerability patterns of road network to extreme floods based on accessibility measures," *Transportation Research Part D: Transport and Environment*, vol. 100, p. 103045, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1361920921003424

[10] D. Li, Z. Zhang, B. Alizadeh, Z. Zhang, N. Duffield, M. A. Meyer, C. M. Thompson, H. Gao, and A. H. Behzadan, "A reinforcement learning-based routing algorithm for large street networks," *International Journal of Geographical Information Science*, vol. 38, no. 2, pp. 183–215, 2024.

[11] W. Tian, K. Xin, Z. Zhang, M. Zhao, Z. Liao, and T. Tao, "Flooding mitigation through safe & trustworthy reinforcement learning," *Journal of Hydrology*, vol. 620, p. 129435, 2023.

[12] B. D. Bowes, A. Tavakoli, C. Wang, A. Heydarian, M. Behl, P. A. Beling, and J. L. Goodall, "Flood mitigation in coastal urban catchments using real-time stormwater infrastructure control and reinforcement learning," *Journal of Hydroinformatics*, vol. 23, no. 3, pp. 529–547, 2021.

[13] C. Fan, X. Jiang, and A. Mostafavi, "Evaluating crisis perturbations on urban mobility using adaptive reinforcement learning," *Sustainable Cities and Society*, vol. 75, p. 103367, 2021.

[14] H. Vajjarapu, A. Verma, and H. Allirani, "Evaluating climate change adaptation policies for urban transportation in india," *International journal of disaster risk reduction*, vol. 47, p. 101528, 2020.

[15] W. Wang, S. Yang, J. Gao, F. Hu, W. Zhao, and H. E. Stanley, "An Integrated Approach for Assessing the Impact of Large-Scale Future Floods on a Highway Transport System," *Risk analysis*, vol. 40, no. 9, pp. 1780–1794, 2020.

[16] M. Pregnolato, A. Ford, and R. Dawson, "Disruption and adaptation of urban transport networks from flooding," in *E3s Web of conferences*, vol. 7. EDP Sciences, 2016, p. 07006.

[17] V. H. S. de Abreu, A. S. Santos, and T. G. M. Monteiro, "Climate change impacts on the road transport infrastructure: A systematic review on adaptation measures," *Sustainability*, vol. 14, no. 14, p. 8864, 2022.

[18] SCALGO, "SCALGO Live," 2024, Accessed: 2024-06-07. [Online]. Available: https://scalgo.com/live/denmark

[19] M. G. McNally, "The four-step model," in *Handbook of transport modelling*. Emerald Group Publishing Limited, 2007, vol. 1, pp. 35–53.

[20] Danmarks Meteorologiske Institut, "Klimaatlas," 2023, Accessed: 2024-08-26. [Online]. Available: https://www.dmi.dk/klima-atlas/data-i-klimaatlas

[21] D. P. Van Vuuren, J. Edmonds, M. Kainuma, K. Riahi, A. Thomson, K. Hibbard, G. C. Hurtt, T. Kram, V. Krey, J.-F. Lamarque *et al.*, "The representative concentration pathways: an overview," *Climatic change*, vol. 109, pp. 5–31, 2011.

[22] Styrelsen for Dataforsyning og Infrastruktur, "Danmarks Højdemodel," 2024, accessed: 2024-06-07. [Online]. Available: https://sdfi.dk/data-om-danmark/vores-data/danmarks-hoejdemodel

[23] Vejdirektoratet, "Grøn Mobilitetsmodel (GMM)," 2022, Accessed: 2024-06-06. [Online]. Available: https://www.vejdirektoratet.dk/segment/groen-mobilitetsmodel

[24] H. Christiansen and O. Baescu, "The danish national travel survey: 0623v1," https://doi.org/10.11581/dtu:00000034, 2024, Accessed: 2024-06-01.

[25] W. E. Deming and F. F. Stephan, "On a least squares adjustment of a sampled frequency table when the expected marginal totals are known," *The Annals of Mathematical Statistics*, vol. 11, no. 4, pp. 427–444, 1940.

[26] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische mathematik*, vol. 1, no. 1, pp. 269–271, 1959.

[27] OpenStreetMap contributors, "Openstreetmap foundation. available as open data under the open data commons open database license (odbl) at openstreetmap.org," 2024.

[28] G. Boeing, "Modeling and Analyzing Urban Networks and Amenities with OSMnx," 2024.

[29] K. C. H. van Ginkel, F. Dottori, L. Alfieri, L. Feyen, and E. E. Koks, "Flood risk assessment of the european road network," *Natural Hazards and Earth System Sciences*, vol. 21, no. 3, pp. 1011–1027, 2021. [Online]. Available: https://nhess.copernicus.org/articles/21/1011/2021/

[30] M. Pregnolato, A. Ford, S. M. Wilkinson, and R. J. Dawson, "The impact of flooding on road transport: A depth-disruption function," *Transportation Research Part D: Transport and Environment*, vol. 55, pp. 67–81, 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1361920916308367

[31] Transportministeriet, "Transportøkonomiske enhedspriser, version 2.0," https://www.man.dtu.dk/myndighedsbetjening/teresa-og-transportoekonomiske-enhedspriser, 2022.

[32] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.

[33] R. Bellman, "A markovian decision process," *Journal of Mathematics and Mechanics*, pp. 679–684, 1957.

[34] M. Towers, J. K. Terry, A. Kwiatkowski, J. U. Balis, G. d. Cola, T. Deleu, M. Goulão, A. Kallinteris, A. KG, M. Krimmel, R. Perez-Vicente, A. Pierré, S. Schulhoff, J. J. Tai, A. T. J. Shen, and O. G. Younis, "Gymnasium," Mar. 2023. [Online]. Available: https://zenodo.org/record/8127025

[35] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dormann, "Stable-Baselines3: Reliable Reinforcement Learning Implementations," *Journal of Machine Learning Research*, vol. 22, no. 268, pp. 1–8, 2021. [Online]. Available: http://jmlr.org/papers/v22/20-1364.html

[36] S. Huang and S. Ontañón, "A closer look at invalid action masking in policy gradient algorithms," *arXiv preprint arXiv:2006.14171*, 2020.

[37] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.

[38] City of Copenhagen, "Climate plan 2035," 2023, Accessed: 2024-08-27. [Online]. Available: https://international.kk.dk/about-copenhagen/liveable-green-city/2035-climate-plan