

Incorporating Quality of Life in Climate Adaptation Planning via Reinforcement Learning

Miguel Costa *Technical University of Denmark* migcos@dtu.dk
Arthur Vandervoort *Technical University of Denmark* apiva@dtu.dk
Martin Drews *Technical University of Denmark* mard@dtu.dk
Karyn Morrissey *Univ. of Galway* karyn.morrissey@universityofgalway.ie
Francisco C. Pereira *Technical University of Denmark* camara@dtu.dk

Abstract

Urban flooding is expected to increase in frequency and severity as a consequence of climate change, causing wide-ranging impacts that include a decrease in urban Quality of Life (QoL). Meanwhile, policymakers must devise adaptation strategies that can cope with the uncertain nature of climate change and the complex and dynamic nature of urban flooding. Reinforcement Learning (RL) holds significant promise in tackling such complex, dynamic, and uncertain problems. Because of this, we use RL to identify which climate adaptation pathways lead to a higher QoL in the long term. We do this using an Integrated Assessment Model (IAM) which combines a rainfall projection model, a flood model, a transport accessibility model, and a quality of life index. Our preliminary results suggest that this approach can be used to learn optimal adaptation measures and it outperforms other realistic and real-world planning strategies. Our framework is publicly available: https://github.com/MLSM-at-DTU/maat_qol_framework.

Keywords: Quality of Life, Climate Adaptation, Pluvial Flooding, Reinforcement Learning, Transportation

1 Introduction

The drive to measure Quality of Life (QoL) in cities has given urban planners and policymakers the tools to identify needs not typically captured by traditional economic metrics [Marans and Stimson, 2024, Atkinson, 2013, Lee and Sener, 2016]. We define QoL as the set of amenities accessible to a given person at a given time, measuring the “circumstances of a person’s life rather than his or her response to those circumstances” [Diener, 2006]. As urban flooding looks to become more prevalent in the coming century, adapting cities to flooding while maintaining QoL should be a critical priority for urban planners [IPCC, 2023].

Climate change will significantly impact mobility and access to locations and services. High-impact weather events, including extreme rainfall [IPCC, 2023], are projected to rise in frequency and intensity, increasing disruptions caused by urban flooding [Olesen et al., 2014]. Flooding disrupts economic and social activities [Hammond et al., 2015], restricting individuals’ access and,

consequently, affecting QoL. These disruptions have both immediate and long-term implications for urban resilience and accessibility.

Copenhagen is among Denmark’s most vulnerable cities to flooding [Prall et al., 2024]. To make the city resilient while maintaining or improving QoL, policymakers must choose effective adaptation strategies to implement over time. Given the complex and dynamic nature of climate change and urban flooding, reinforcement learning (RL) presents a promising approach to optimise policy decisions [Gilbert et al., 2022]. RL can be used to discover which adaptation strategies balance long- and short-term quality of life with economic costs.

Thus, we aim to answer the following question: “*What is the sequence of climate adaptation measures that lead to a higher quality of life in the long term?*” Seeking to minimise the impacts of climate-related pluvial flooding on quality of life, we explore how different climate adaptation strategies impact accessibility to key services and use RL as a tool to identify adaptation sequences that maintain current levels of QoL under uncertain climate projections. Though we are focusing on Copenhagen’s inner city as a case study, our framework can easily be expanded to other cities, climate-related impacts (e.g., coastal flooding or wildfires), and contexts (e.g., subjective wellbeing or mobility justice).

2 Methodology

We frame our approach as an Integrated Assessment Model (IAM) that includes a rainfall model, 2) a flood model, 3) a transport accessibility component, and 4) a quality of life index model. Figure 1 provides an overview of our framework.




Figure 1: IAM using RL to learn what the best sequence of adaptation policies are to maximise quality of life under climate changing scenarios.

2.1 Rainfall Projection and Flood Model

To simulate rainfall events, we sampled precipitation intensities based on probability distributions built from the Danish Meteorological Institute’s Climate Atlas [DMI, 2023] for 2023–2100. For simplicity, we assumed the projected rainfall intensity to be equal to the accumulated daily rainfall. After sampling a rainfall event (i.e., amount of rainfall), we modelled the associated urban flood. Urban pluvial flooding is often caused by intense short duration precipitation (cloudbursts, from minutes to a few hours). For this, we used SCALGO Live [SCALGO, 2024], a event-based tool for watershed delineation, flood depth, and

flow direction modelling based on digital terrain data. For any rainfall event, water was distributed according to its intensity and terrain properties. In the end, water depths of resulting floods were mapped in Copenhagen.

2.2 Transport Accessibility and Quality of Life Index

We define QoL as the weighted sum of the per-capita number of amenities accessible within a reasonable travel time from residential locations [Dobrowolska and Kopczewska, 2024]. Taking this approach enables us to model QoL at a granular level, which allows us to identify the effects of flooding across (and within) urban administrative boundaries where QoL indicators are typically gathered [Marans and Stimson, 2024].

Our QoL index is adapted from previous work on QoL by Dobrowolska and Kopczewska [2024]. First, we compute the per-capita number of accessible points of interest (POI) of a given category for a location. We then take the weighted sum of all POIs for a given location, where weights are determined using a logistic regression model that predicts satisfaction with life in cities from [EC/DG REGIO, 2023]. This yields an index that maps the relative density of amenities per capita for a given area, weighted by the relative importance of those amenities based on survey data.

2.3 Reinforcement Learning

We posit to learn the best sequence of adaptation policies that maximise QoL using RL. RL uses an agent-based approach to interact with the above environment by taking an action (adaptation measure) and maximising a (delayed) reward function [Sutton and Barto, 2018], learning to balance trade-offs between competing actions and input uncertainty in the environment.

In this work, we propose eight actions that can be implemented in each area of Copenhagen. These include increasing road drainage and permeable paving solutions. Actions change the environment, directly changing transport and indirectly affecting QoL. We define the reward function to optimize for as:

$$R = \sum_i \beta_Q Q_i + \beta_A A_i + \beta_M M_i \quad (1)$$

where Q_i corresponds to Quality of Life index at the i -th zone, A_i is the cost of applying an action and M_i its maintenance cost over time. This reward function is highly customisable. Users can choose different β weights depending on competing priorities and trade-offs between QoL and economic costs.

To learn which action to perform at a given time, our RL agent takes an action and collects information on the state of our digital city. Over time, the agent learns the best sequence of policies that maximise the cumulative reward.

3 Preliminary Results and Discussion

We setup our IAM using Python, Gymnasium interface [Towers et al., 2023], Stable-Baseline3 [Raffin et al., 2021], and PPO [Huang and Onta  n, 2020,

Schulman et al., 2017]. As a preliminary case study, we conduct an experiment in Copenhagen’s inner city (consisting of 29 zones) by setting the time horizon between 2023–2100, and setting β weights that prioritise QoL over economic costs. We now present preliminary results for five distinct seed runs to allow for different weather projections and robustness.

Figure 2 showcases the results comparing our learnt (Optimal Policy) with five feasible baselines (see Appendix A). Our results show that the RL framework achieves better rewards compared to all baselines. The similarity between the random and the learnt policy can be explained by our experimental setup, which places heavier emphasis on QoL. This reduces the relative importance of action costs, leading to a similar cumulative reward to the random policy. Crucially, the sequence of actions taken by the agent is different to the random policy, leading to small increases in cumulative QoL and small decreases in cumulative action costs. Future experiments with different weightings are likely to yield further improvements.




Figure 2: Top row: Side-by-side comparison of No Control and Learnt Policy result by city zone. Bottom row: Comparison of total reward between Learnt Policy and baselines (left), and between reward components (right).

Our results demonstrate the feasibility of using RL to construct IAMs and identify policy pathways focused on minimising QoL loss, in contrast to the cost-minimisation or welfare-optimisation approaches typical used in IAMs [Weyant, 2017]. Though these results are promising, IAMs (with or without RL) feature multiple quantifiable and unquantifiable uncertainties that are important to consider [Beck and Krueger]. Moving forward, we plan to continue to investigate the effects of normative choices in models such as these, including testing the

sensitivity of the model to different β weights for QoL (representing the 'value' of QoL relative to economic costs) and using sufficientarian metrics (e.g. Foster-Greer-Thorbecke indices [Karner et al., 2024]). Besides this, we also intend to extend our framework to the full city of Copenhagen.

References

- S. Atkinson. Beyond components of wellbeing: The effects of relational and situated assemblage. *Topoi*, 32(2):137–144, 2013.
- M. Beck and T. Krueger. The epistemic, ethical, and political dimensions of uncertainty in integrated assessment modeling. 7(5):627–645. ISSN 1757-7799. doi: 10.1002/wcc.415. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/wcc.415>.
- Danmarks Meteorologiske Institut (DMI). Klimaatlas, 2023. URL <https://www.dmi.dk/klima-atlas/data-i-klimaatlas>. Accessed: 2024-08-26.
- E. Diener. Guidelines for national indicators of subjective well-being and ill-being. *Applied Research in Quality of Life*, 1(2):151–157, 2006.
- E. Dobrowolska and K. Kopczewska. Mapping urban well-being with Quality Of Life Index (QOLI) at the fine-scale of grid data. *Scientific Reports*, 14(1): 9680, Apr. 2024. ISSN 2045-2322. doi: 10.1038/s41598-024-60241-0.
- EC/DG REGIO. Quality of Life in European cities survey, 2023.
- T. K. Gilbert, S. Dean, T. Zick, and N. Lambert. Choices, risks, and reward reports: Charting public policy for reinforcement learning systems, 2022. URL <https://arxiv.org/abs/2202.05716>.
- M. J. Hammond, A. S. Chen, S. Djordjević, D. Butler, and O. Mark. Urban flood impact assessment: A state-of-the-art review. *Urban Water Journal*, 12(1):14–29, 2015.
- S. Huang and S. Ontañón. A closer look at invalid action masking in policy gradient algorithms. *arXiv preprint arXiv:2006.14171*, 2020.
- 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.)]*, pages 35–115. IPCC, Geneva, Switzerland, doi: 10.59327/IPCC/AR6-9789291691647, 2023.
- A. Karner, R. H. Pereira, and S. Farber. Advances and pitfalls in measuring transportation equity. *Transportation*, pages 1–29, 2024. ISSN 0049-4488. Publisher: Springer.
- R. J. Lee and I. N. Sener. Transportation planning and quality of life: Where do they intersect? *Transport Policy*, 48:146–155, 2016. ISSN 0967-070X. doi: <https://doi.org/10.1016/j.tranpol.2016.03.004>. URL <https://www.sciencedirect.com/science/article/pii/S0967070X16300968>.

- R. W. Marans and R. J. Stimson. Quality of life in large-scale, big-city urban environments: A world perspective. In R. W. Marans, R. J. Stimson, and N. J. Webster, editors, *Handbook of Quality of Life Research*, pages 147–164. Edward Elgar Publishing, Mar. 2024. ISBN 978-1-78990-879-4 978-1-78990-878-7. doi: 10.4337/9781789908794.00018.
- M. Olesen, K. S. Madsen, C. A. Ludwigsen, F. Boberg, T. Christensen, J. Cappelen, O. B. Christensen, K. K. Andersen, and J. H. Christensen. Fremtidige klimaforandringer i danmark (danmarks klimacenter rapport nr. 6 2014). Technical report, Danmarks Meteorologiske Institut, 2014.
- M. C. Prall, U. S. Brandt, N. S. Halvorsen, M. U. Hansen, N. Dahlberg, and K. J. Andersen. A comprehensive approach for assessing social flood vulnerability and social flood risk: The case of Denmark. *International Journal of Disaster Risk Reduction*, 111:104686, Sept. 2024. ISSN 2212-4209. doi: 10.1016/j.ijdr.2024.104686.
- A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dornmann. Stable-Baselines3: Reliable Reinforcement Learning Implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021. URL <http://jmlr.org/papers/v22/20-1364.html>.
- SCALGO. SCALGO Live, 2024. URL <https://scalgo.com/live/denmark>. Accessed: 2024-06-07.
- J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- 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. URL <https://zenodo.org/record/8127025>.
- J. Weyant. Some Contributions of Integrated Assessment Models of Global Climate Change. *Review of Environmental Economics and Policy*, 11(1):115–137, Jan. 2017. ISSN 1750-6816. doi: 10.1093/reep/rew018. URL <https://www.journals.uchicago.edu/doi/full/10.1093/reep/rew018>. Publisher: The University of Chicago Press.

A Appendix A

We compare the learnt (Optimal Policy) reward versus five other baselines:

- No Control: no actions are implemented

- Cumulative Response: adaptation starts when 3 10-year return period rainfall events occur within a 10-year period
- Intense Event Response: adaptation starts after a 20-year return period event occurs
- Election Cycle: adaptation occurs every 4 years
- Random Control: actions are taken randomly