MSc thesis in Geomatics

Leveraging Street Level Predictive Modelling with Green Infrastructure for Urban Resilience

Shawn Tew

2025

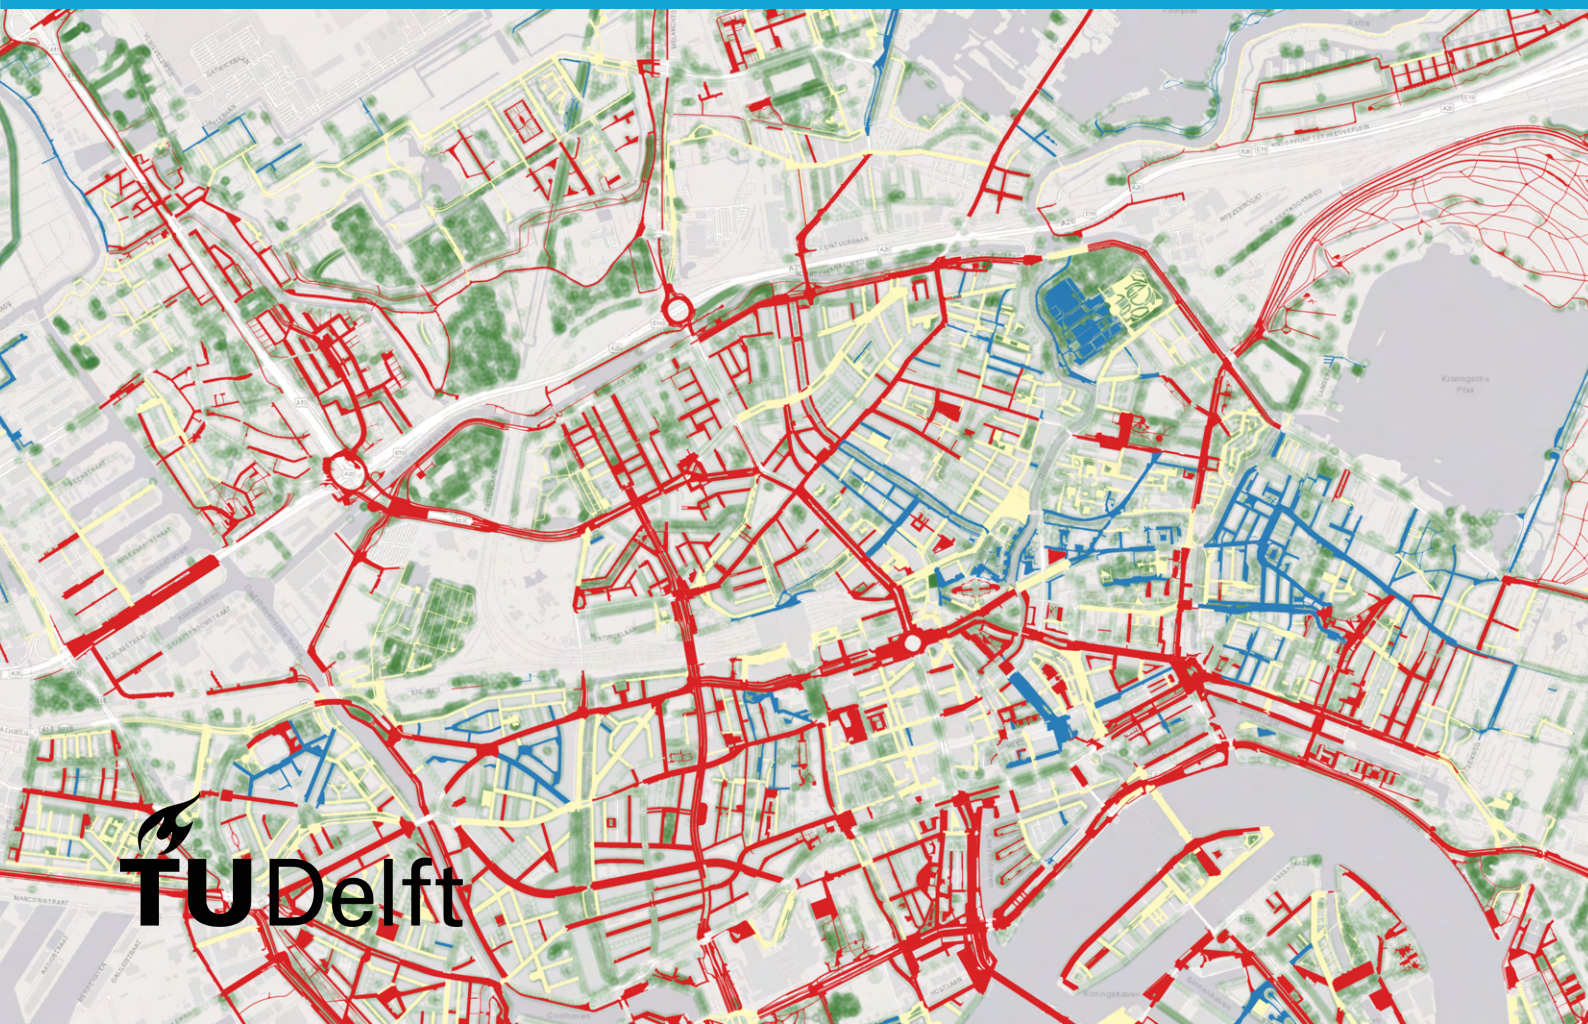

TU Delft

MSc thesis in Geomatics

Leveraging Street Level Predictive Modelling with Green Infrastructure for Urban Resilience

Shawn Tew

June 2025

A thesis submitted to the Delft University of Technology in partial fulfillment of the requirements for the degree of Master of Science in Geomatics

The work in this thesis was carried out in the:

3D geoinformation group
Delft University of Technology

Supervisors:	Dr.Daniele Cannatella
	Dr. Claudiu Forgaci
Co-reader:	Dr. Giorgio Agugiaro

Abstract

Urban resilience is an evergrowing issue as cities face environmental, social and infrastructural challenges. Green Infrastructure (GI) is currently being used as one of the main approaches in achieving urban resilience by providing multifunctional benefits, such as improving microclimate, enhancing biodiversity, and increasing accessibility to nature. However, through this process the evaluation and optimisation of green infrastructure at the street level remains a methodological challenge, particularly in the case of integrating multi dimensional indicators into predictive models using nature based solutions.

This study develops a machine learning based framework in order to assess and optimise green infrastructure in urban areas at a street level scale. Through the use of publicly available spatial datasets on environmental, biodiversity and morphological factors, the project constructs a comprehensive dataset incorporating factors such as urban heat islands, green infrastructure distribution, and green space accessibility.

Modelling techniques such as gradient boosting regression and random forest regression are employed as a regression technique to predict urban resilience related targets. These will provide insights into how resilient different areas of Rotterdam can be. By focusing on the study at a street level approach, this study aims to offer a comprehensive understanding of green infrastructure effectiveness at a high resolution, providing urban planners with data driven recommendations for optimising and implementing green infrastructure solutions to areas in need.

The findings of this research aim to contribute by bridging gaps in green infrastructure assessment by integrating geospatial data, network based spatial analysis, with machine learning, to assist in decision making to make urban areas more resilient.

Acknowledgements

I would like to express my sincerest gratitude to my supervisors , Dr. Daniele Cannatella and Dr. Claudiu Forgaci, for their invaluable guidance, constructive feedback, and continuous support throughout the course of this thesis. The insights given from their experience in this field helped shaped the direction and aim of this research. I would also like to extend my thanks to all the lecturers in the Geomatics programme at TU Delft. The teachings in all courses laid the foundation for this work by providing the knowledge and skills used to complete this thesis.

Contents

List of Figures

List of Tables

Acronyms

1. Introduction

As many cities today face escalating urban challenges such as heat stress, biodiversity loss, as well as ageing infrastructure, cities worldwide have increased pressure to become increasingly resilient. Urban resilience refers to a cities ability to absorb, adapt, and recover from environmental, social and economic deterrents while retaining functionality [Gkontzis et al., 2024]. Although efforts have been made to combat these urban challenges, significant gaps remain in tackling the multidimensional drivers that contribute to urban resilience, specifically using GI.

Green infrastructure, as described by Ghofrani et al. [2017], is defined as "an interconnected network of waterways, wetlands, wildlife habitats, and other natural areas; greenways, parks and other conservation lands; working farms, ranches and forests; and other open spaces that support species, maintain natural ecological processes, sustain air and water resources, and contribute to the health and quality of life for communities and people." GI has become increasingly common strategies for enhancing urban resilience by delivering ecosystem services, such as heat mitigation, storm water management, and biodiversity conservation. However, challenges persist in quantifying the multi-functional benefits of GI as highlighted by Korkou et al. [2023]. This includes the identification of priority intervention areas, and integrating diverse indicators of resilience into urban planning frameworks.

Rotterdam, as a densely populated and climate sensitive city provides valuable context in for examining how GI can contribute to enhancing urban resilience, particularly by mitigating heat stress, managing flood risks as well as equitable green space access. Although the city has made significant progress in executing advanced planning strategies, Rotterdam continues to face persistent challenges, including rising temperatures, increased flood risk due to its geographical location, and disparities in green space accessibility in differing neighbourhoods. This makes Rotterdam a reasonable case study for studying the advantages or disadvantages of green infrastructure in urban resilience.

The aim of this thesis is to develop a modular, spatial workflow, that utilises spatial data, network analysis and machine learning to predict diverse street level resilience outcomes related to GI and the overall morphology of the built environment. Compared to summarising resilience into a single measure, the framework attempts to predict real world outcomes such as night time heat stress, flood vulnerability and air quality based on commonly found indicators. This allows for a reproducible and flexible outcome oriented approach that can be tailored to different urban climate resilience challenges and priorities.

Environmental factors can be crucial in determining the effect of GI in mitigation strategies by improving thermal regulation and flood risks. Biodiversity indicators, such as tree species diversity and distribution, are essential in understanding green space connectivity as well as ecological importance of an urban area. Morphological factors, including angular betweenness and attraction reach may influence how effectively GI can be integrated into the ongoing urban fabric. Through the combination of these indicators within a predictive modelling framework, this study aims to provide a data driven approach to optimising the implementation or integration of GI in terms of urban climate resilience.

Through this research, it will be attempted to show how using data driven analysis to execute GI related intervention strategies can be utilised to create sustainable, equitable and resilient urban environments. The findings aim to output a framework for predicting, analysing and comparing different resilience outcomes. This allows for evidence based planning in Rotterdam and hopefully other cities.

1.1. Research objectives

The main objective of this research is to explore different machine learning techniques in the context of GI and their effect on urban resilience. The main research question of the thesis is:

"How can predictive modelling techniques be used to evaluate the effectiveness of GI in improving urban resilience at street level across different dimensions of the built environment"

To fully examine this research question, the project will be broken down into smaller sub questions, where each will be attempted to be explored in detail.

- What aspects of GI can be assessed in a multifunctional way to evaluate their contributions to urban resilience across multiple dimensions?
- How can environmental, biodiversity, and morphological features be combined to comprehensively assess and enhance urban resilience?
- How can predictive modelling techniques, such as machine learning, be used to evaluate and optimise GI for improving urban resilience?
- How can the presence of spatial autocorrelation in the residuals affect the reliability of the model's predictions?
- How can the proposed framework be applied to assess and prioritise GI interventions in Rotterdam?

1.2. Scope

The scope of this project will focus on exploring machine learning techniques for assessing the role of GI in urban resilience. The key objectives include:

- Core tasks
  - Investigate different machine learning techniques and different types of GI
  - Integrate multi-dimensional urban resilience indicators (e.g., environmental, morphological, biodiversity) into a unified predictive framework.
  - Preprocess and engineer features from geospatial datasets (e.g. imperviousness, tree data, accessibility measures).
- Key Analyses:

  - Predict resilience performance across Rotterdam neighbourhoods using trained models.
  - Analyse the relative importance of GI indicators in resilience outcomes.
  - Assess spatial disparities and correlations between resilience and socio-environmental features.
- Potential Additional Considerations
  - Assess social economic benefits of GI
  - Provide policy recommendations based on spatial analysis findings
- Exclusions
  - No development of real time predictive systems
  - No focus on private GI data

1.3. Outline

This thesis is structured around 5 main chapters. Chapter 1 introduces the research and focuses on explaining the role of GI in enhancing urban resilience. Here, the main research question and supporting sub questions to highlight the motivation for using machine learning techniques to explore spatial urban data. Chapter 2 then reviews existing literature on GI benefits, environmental, biodiversity and morphological urban resilience indicators and previous uses of machine learning in these types of studies. Chapter 3 details the methodological framework, including data collection, feature extraction, and selection of relevant indicators. The three training models used are also described along with the evaluation strategies used to assess each. Chapter 4 presents the implementation and results of each model across all three cases, including performance metrics, feature importance, and spatial visualisations of predictions and residuals. Finally chapter 5 interprets key findings in relation to the research questions, discusses the contributions and limitations of the study before finally proposing recommendations for future urban research in this area.

2. Related work

This section presents an overview of the relevant literature studied, giving an overall outlook of key findings and methodologies from previous work that inform and support the foundation of this thesis. The reviewed work addresses the roles and impacts of GI in urban resilience, modelling approaches for evaluating GI, and the how Rotterdam can act as a case study.

2.1. Green Infrastructure for Urban resilience

2.1.1. Green infrastructure for the environment

Green infrastructure has a large potential in contribution to the improvements of urban resilience, addressing many environmental challenges. Works done by Staddon et al. [2018] place emphasis on the environmental importance of GI in the case of mitigating Urban Heat Island (UHI) intensity, storm water management, and biodiversity protection. Similarly, as demonstrated by Pamukcu-Albers et al. [2021] GI may function to aid in climate adaptation processes and enhance capacity in withstanding and recovery of environmental related stresses.

Stormwater management is another key environmental utilisation of GI. Traditional grey infrastructure interventions can be seen as suboptimal in cities with the geographical landscape similar to Rotterdam. Research done by Salih and Báthoryné Nagy [2024] demonstrates interventions using GI provide a cleaner nature based solution for urban flood mitigation by increasing the imperviousness of an area and reducing the overall runoff. Furthermore, the concept of ecological corridors in the study highlights that green space connectivity may additionally support thermal comfort by promoting more natural airflow passages to aid in humidity regulation.

Additionally, biodiversity protection is an essential function of GI. In highly urbanised environments green space fragmentation poses significant challenges, such as integrating diverse vegetation types. This creates habitats for various species and supports pollination populations thereby aiding in ecosystem health. For instance, Filazzola et al. [2019] conducted a study to show that constructed GI significantly enhances urban biodiversity by providing habitats that can support a larger variety of species, including pollinators. This benefits urban areas in comparison to the conventionally used impervious grey infrastructure. However, despite these advantages, effective implementation of GI remains as a challenge due to the competing urban land use demands and spatial constraints.

2.1.2. Social benefits of Green Infrastructure

While much of the focus on GI has been the environmental benefits, recent studies have brought other GI contributions to light, including social well being as well as economic sustainability. Well integrated GI can be seen to improve public health by increasing access to nature, reducing stress and promoting active mobility through walking and cycling [Kabisch and Van Den Bosch, 2017]. Staddon et al. [2018] also highlights that green spaces may provide psychological relief, improve cultural inclusivity by encouraging social cohesion with offering spaces for recreation and community activities.

Furthermore, urban green spaces play a vital role in education and recreational opportunities for children and young adults. Green spaces such as parks provide spaces for learning experiences to foster environmental awareness and sustainability knowledge. For example, Dadvand et al. [2015] shows that green space exposure may positively influence cognitive development in primary school children in aspects such as attention, memory and overall academic performance. Similarly, Chawla [2015] showcases how GI in urban areas can promote curiosity and problem solving skills in addition to the overall emotional support and well being among children. Overall, these green spaces provide recreational opportunities, allowing urban residents to engage in outdoor activities that serve as an essential element in creating multifunctional spaces that support the growth and development of both children and young adults.

Following education, green spaces have been associated with reductions in crime rates and overall urban safety. The inclusion of trees and vegetation in urban areas has been linked to lower levels of property crime and violet acts as there is a larger number of natural surveillance, coming from community engagement [Kuo and Sullivan, 2001]. This study found that residents living in green environments showed better social cohesion, leading to crime prevention through community driven efforts.

Finally, ensuring equitable GI access is an essential issue in urban planning as disparities in green space allocation often reflect social and economic inequalities within the area [Anguelovski et al., 2022]. Green gentrification is also seen within urban areas in the same study done by Anguelovski et al. [2022] where it is seen that new GI developments lead to increased properly values, creating displacements for lower income families and reducing overall access to the environmental benefits of green spaces. Addressing these inequities requires the development of policies that prioritise GI to ensure communities regardless of socioeconomic status can access and benefit green spaces.

2.1.3. The role of Urban Morphology for Green Infrastructure

Apart from environmental and social factors, the spatial configuration of the built environment may additionally influence the effectiveness of GI and the assumed functions. This may include the urban morphology, arrangement of buildings, streets and open spaces and the effect on airflow, heat retention, water management [Stewart and Oke, 2012]. Similarly, Ziter et al. [2019] examines how street canyons, building density and road typologies impact the cooling and stormwater management capacity of GI. Compact urban layouts that include tall buildings and narrow streets induce canyon effects and may also lead to localised heat retention. In contrast, streets that are designed well with GI integration can facilitate better air circulation and reduce overall urban heat stress [Stewart and Oke, 2012].

Dense urban areas benefit from street oriented GI such as greenery lined streets and green corridors in order to alleviate general heat stress. Studies by Bowler et al. [2010] reveal that surface temperatures, air temperatures and overall thermal comfort in a region can be improved with the aid of higher tree canopy coverage. Moreover, integrating permeable surfaces are also known to improve groundwater recharge and reduce flooding in urban areas that are highly covered with impermeable surfaces [Andrés-Doménech et al., 2018]. Green infrastructure efficiency is also impacted by building density. For instance, areas that are high in building density but are low in open space usually have high dependence on alternative GI implementations such as green walls and roofs to improve the ecological biodiversity of the urban area [Pamukcu-Albers et al., 2021].

In terms of city wide interventions, [van der Berg, 2023] and [Brears and Brears, 2018] have innovative strategies to integrate GI into the urban landscape of the Rotterdam. For example, the blue-green infrastructure approach, which incorporates water and vegetative based solutions. This has recently been seen more globally in turn, as a form of engineering for biodiversity [Perrelet et al., 2024]. However, gaps remain in spatial equity and connectivity of green spaces across neighbourhoods, reinforcing the need for continued research in this field.

Overall, the general morphology of the built environment plays a defining role in how GI can be effectively implemented and maintained. Research has demonstrated that GI can be affected by multidimensional factors as well as how GI can meaningfully contribute towards urban resilience [Ghofrani et al., 2017].

2.2. Modelling Approaches

Machine learning is the study of algorithms in which a computer system would use to perform a task without being programmed. This is the main way in which computers are taught to handle data more efficiently [Mahesh et al., 2020]. Modern modelling approaches have been applied in many recent research to evaluate and optimise GI for urban resilience. Machine learning techniques such as Random Forest Regression (RFR), Gradient Boosting Regression (GBR) as well as Support Vector Regression (SVR) are some that have been widely employed due to their ability to handle large datasets and identify complex relationships between environmental, and morphological indicators.

This is demonstrated by [Kutty et al., 2022] where machine learning was used to predict urban resilience outcomes based on multi dimensional datasets. Their study applied multiple models ranging from supervised to unsupervised. The features used roughly 30 liveability related indicators and 38 resilience related indicators and the data from 35 leading European smart cities. Using this, the study was able to highlight spatial patterns of vulnerability and resilience, providing decision makers with targeted insights on specific areas requiring intervention. The study also emphasised the interpretability of model outputs. This enables practical use in planning scenarios and improving the use of model driven policy support.

Similarly, Cong and Inazumi [2024] applied ensemble learning techniques and geo data with an aim to improve urban resilience in Setagaya, Tokyo. With data from over 400 geotechnical surveys, the study predicted the depth different soil layers to assess earthquake risks. The predictive modelling models used were mainly bagging based ensemble methods similar to random forest that will be used in this project. These outputs from the models have direct relevance for disaster preparation and infrastructure planning, aiding in reinforcing

the value of machine learning in spatially targeted resilience strategies. The study provides an example of how machine learning can support resilience planning through actionable geospatial insights.

Furthermore, a study by Mason et al. [2023] used machine learning to analyse data from internet connected sensors to measure water levels in GI installations in Detriot, USA. The work done by Mason et al. [2023] applied automated storm segmentation methods to identify key factors influencing GI performance. These factors include depth to groundwater and impervious surface area. This approach allowed for optimisation of GI design and placement to enhance resilience to flooding related events. The integration of real-time sensor data with predictive analytics demonstrates the growing potential of machine learning in operationalising GI performance evaluation at scale.

In addition, using machine learning jointly with GIS-based spatial analysis techniques have also been used to map and assess GI distribution and accessibility. [Wanghe et al., 2020] showcase how GIS tools can be used to prioritise ecological corridors by analysing spatial connectivity and least cost paths. Similarly, [Gkontzis et al., 2024] integrate GIS with digital twins to simulate highly accurate spatial impacts of GI implementations.

2.2.1. Random Forest

Random forest is a machine learning algorithm that uses a collection of decision trees to classify or predict the value of a variable [Rodriguez-Galiano et al., 2015]. Random Forest builds decision trees by selecting random subsets of data and features to ensure robustness and reduce risk of overfitting. For regression tasks, random forests aggregate the predictions of individual trees by averaging their outputs, while classification uses majority voting to determine the final prediction[Rodriguez-Galiano et al., 2015].

Figure 2.1 illustrates the overall structure of a random forest algorithm. The input dataset is used to generate multiple decision trees in this case, Decision tree 1, Decision tree 2, up until decision tree N. Each tree makes a decision independent to each other and the final output is determined through averaging or majority voting. The advantage of this approach is that it reduces the bias and variance that come from individual decision trees, resulting in overall more stable predictions [Breiman, 2001].

Mathematically, the RF regression prediction for a new street segment x is given by:

$$\hat{Y} = \frac{1}{M} \sum_{m=1}^{M} f_m(x) \tag{2.1}$$

where:

- M is the total number of decision trees
- $f_m(x)$ represents the predition from the m-th tree
- $\hat{Y}$ is the final predicted target variable.

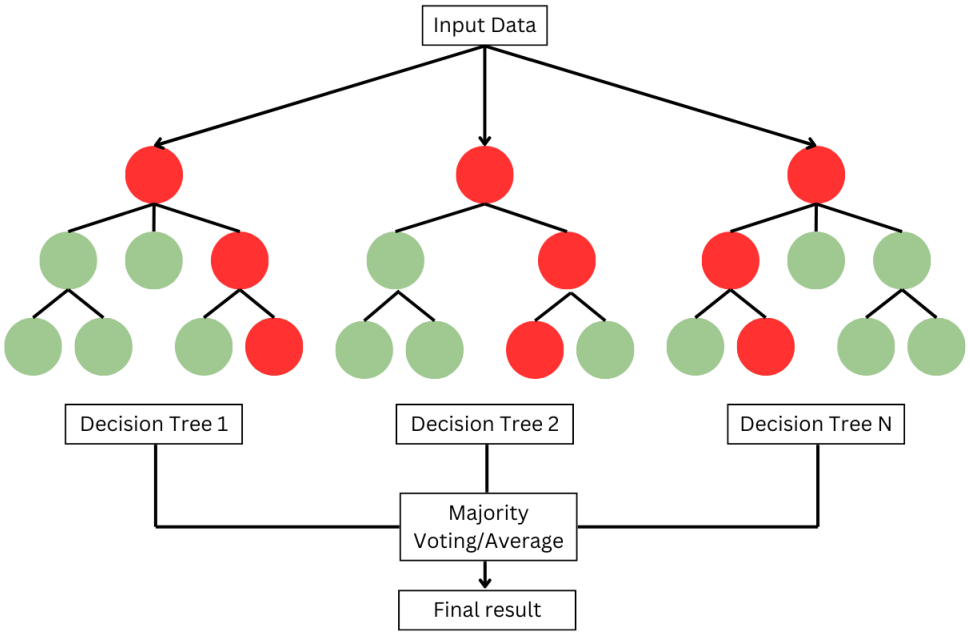

Figure 2.1.: Random Forest Diagram

In the case of this project, Random Forest can be applied as a framework to predict urban resilience performance by analysing environmental, biodiversity and morphological indicators, such as UHI effects, tree species distribution, and green space connectivity. Breiman [2001] highlights that it has the ability to handle complex, non linear relationships between variables. This may seem suited for the project, allowing for GI with multi dimensional impacts on urban resilience to be assessed appropriately. Moreover, random forests models can be used to rank importance of the different features [Breiman, 2001]. This helps identify which indicator may have the most significant impact on overall resilience outcomes. The following outcomes allow for interpretability for urban planners and policy makers as it creates actionable insights into most effective GI implementations to enhance urban resilience.

2.2.2. Gradient Boosting

Gradient boosting is a similar machine learning technique that uses decision trees. However, unlike random forest, each new tree will correct errors of previous ones working sequentially. Gradient Boosting will construct trees iteratively, optimising predictions by minimising the loss function at each step. This process allows the model to learn from the mistakes made in the previous step and continuously refine its predictions, overall making it effective for complex non linear relationships in our data [Friedman, 2001].

Figure 2.2 illustrates a basic workflow of the gradient boosting algorithm. The algorithm starts with an initial weak model that provides the first prediction. The residual errors are computed, and the new tree is then trained on the residuals from the first model, where a prediction on the residuals is then made. This process is repeated for multiple iterations with each new tree reducing the errors made by the combined previous models. The final

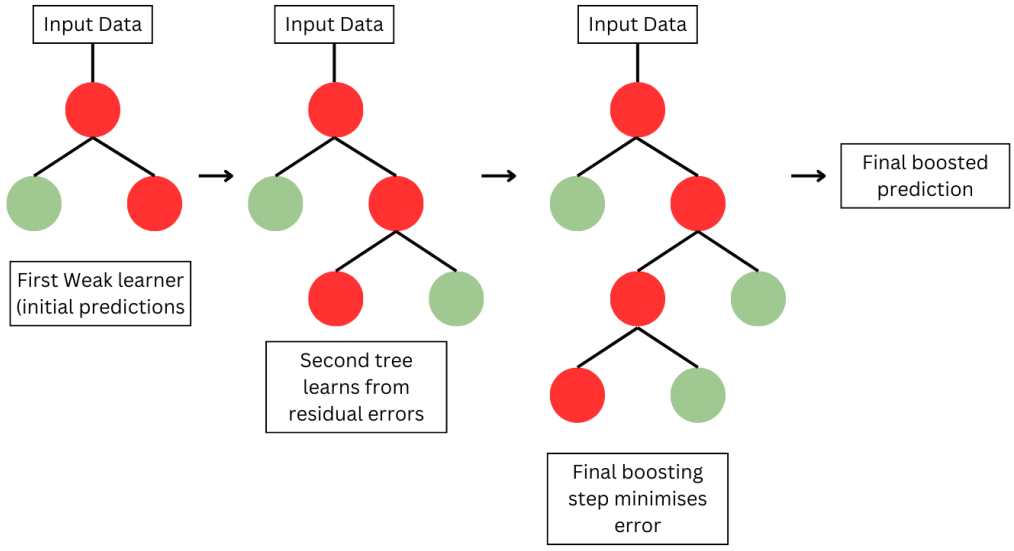

Figure 2.2.: Gradient Boosting Machine Diagram

output is obtained by summing the weighted contributions of all trees, leading to a more accurate prediction [Friedman, 2001].

Given the set of features in this study, the GBR model is trained to predict the target variable by fitting an additive model:

$$\hat{Y} = \sum_{m=1}^{M} \gamma_m f_m(x), \tag{2.2}$$

where:

- M is the total number of boosting iterations (trees),
- $f_m(x)$ represents the prediction from the m-th weak learner (decision tree),
- γ_m is the weight assigned to the m-th learner to minimize the residual error,
- $\hat{Y}$ is the final predicted Urban Resilience Score.

One of the key strengths of gradient boosting, is the ability to assign feature influence or importance scores [Natekin and Knoll, 2013]. This helps urban planners and policy makers to identify and prioritise distinct factors that contribute to urban resilience. For example, if tree count is identified as a dominant predictor of resilience, urban planners may allocate resources to expanding and maintaining existing linear tree corridors or urban forests.

In this study, gradient boosting may seem well suited for predicting urban resilience related outcomes by analysing GI indicators. Urban resilience is again influenced by multi dimensional factors. Gradient boosting allows these interactions between the factors to be studied, making it a solid choice in understanding which indicators or features are most effected by GI interventions. This data driven approach can aid in optimising current GI strategies.

2.2.3. Support Vector Regression

Support Vectors is another type of learning model that can be used for both classification and regression tasks. This is done by finding the optimal hyperplane that best separates data points in high dimensional space [Fletcher, 2009]. In the case of this study, SVR is an extension of support vectors, where the objective will be to fit a function that stays within a specified margin of tolerance ϵ, rather than minimising the errors directly as in an ordinary least squares regression. [Fletcher, 2009].

Figure 2.3 illustrates the core components of SVR, where the hyperplane represents the regression function that aims to predict the target variable based on the input features. The maximum margin sets a range around the prediction line where errors are acceptable, helping to ignore outliers. Support vectors are the key data points that shape this line. SVR relies these points,this makes it good at handling outliers and complex data.

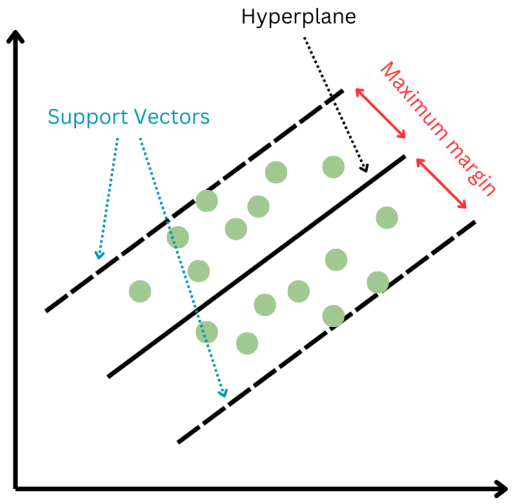

Figure 2.3.: Support Vector Machine Diagram

Mathematically, SVR aims to solve the following optimisation function:

$$\min_{w,b} \frac{1}{2}\|w\|^2 \tag{2.3}$$

where:

- w is the weight vector defining the hyperplane
- b is the bias term
- $\frac{1}{2}\|w\|^2$ represents the margin that we aim to minimize.

In this study, SVR will be applied to model and predict urban resilience related outcomes based on pre described indicators. Since urban resilience is again influenced by multi dimensional factors, SVR, is appropriate as it can handle high dimensional and non linearly

separable data. This is also appropriate to compare against RFR and GBR as it is not a tree based model.

2.3. Rotterdam as a case study

Rotterdam is located in the Rhine-Meuse delta and hosts the largest port in Europe [Frantzeskaki and Tilie, 2014]. Rotterdam serves as an appropriate case study for investigating the role of GI in Urban resilience, given its unique environmental vulnerabilities, and geographical landscape. The city is built on low lying land having roughly 90% of its surface area below sea level [Derkzen et al., 2017], giving exposure to flood risks with a high density built environment that requires deeper care to fully enhance climate resilience.

Implementation of GI in Rotterdam can be complicated due to the changing urban morphology. The city centre of Rotterdam was heavily reconstructed following the second world war, resulting in higher importance of grey infrastructure and separation of spatial functions [Tillie et al., 2018]. This in turn has left a disjointed urban landscape, leaving heavy industrial zones, green and urban areas separate. Incorporating data driven GI implementations into the fragmented environment requires careful planning that should sensitive to both historical and current social needs.

In response to this, Rotterdam has attempted many nature based solutions that combine water and green space management with public space design, such as the Rotterdam Adaptation Strategy [Dircke and Molenaar, 2015]. These interventions include multifunctional spaces such as water squares, green roofs, and urban parks that help address climate related resilince issues. However, scaling these solutions across the city remains a challenge, particularly in high density or socioeconomically disadvantaged areas where implementation may be more complex [Frantzeskaki and Tilie, 2014].

Beyond spatial and morphological challenges, implementation of GI in Rotterdam is further complicated by governance. Responsibilities for climate adaptation, spatial planning and the environment are distributed across multiple municipal departments and external stakeholders, making immediate interventions or implementations difficult. The institutional fragmentation can slow decision making and hinder GI implementation in appropriate time frames. Frantzeskaki et al. [2014] notes effective implementation of GI requires careful cross department collaboration with shared political commitment.

2.4. Research Gaps and Opportunities

The reviewed literature highlights the multifunctionality of GI and the relevance to urban climate resilience. Still, critical research gaps remain. Firstly, most existing studies show focus on either environmental, social and economic outcomes in isolation. There were few that assess the combined multidimensional effects of GI within a single framework. This limits the overall understanding of how GI contributes meaningfully across all dimensions to climate resilience at local scales.

Furthermore, Rotterdam is seen as a highly developed city with high quality open data. However, there is limited research using street level implementations of GI using predictive modelling. Existing models are aimed at neighbourhood or city level scales throughout

the Netherlands, this often overlooks the existing microclimates within each city which are critical in local level planning.

Following this, while spatial disparities and accessibility to GI have been recognised, it is not commonly seen in predictive modelling tasks. The use of integrating these morphological features, biodiversity indicators and environmental metrics at street level can help in better informing GI interventions in dense, complex urban environments such as Rotterdam.

Finally, a gap remains in translating predictive outputs into actionable insights for urban planners. This includes identifying specific street segments that may benefit most from targeneted GI implementations. This study addresses these gaps by applying three regression models to analyse how different features of the built environment influence urban resileince at the street level.

3. Methodology

This chapter outlines the methodology applied in the development of a machine learning-based framework for assessing and optimising GI at the street level in Rotterdam. The workflow as illustrated in figure 3.1 began with the collection of environmental, biodiversity, and morphological data, followed by data cleaning and integration, feature engineering to derive meaningful indicators which then served as inputs for training machine learning algorithms. The features created are evaluated using statistical and conceptual methods. Prediction targets are also selected based a conceptual basis with relevance to urban resilience. Using the selected features and targets, three different models are trained and optimised through hyperparameter tuning. The model results are evaluated using the metrics: R^2, Root Mean Squared Error (RMSE), and Mean Absolute Error (MAE). Finally, a case by case evaluation is done. Predictions are compared to observed values and spatial outputs are mapped to identify areas for GI interventions.

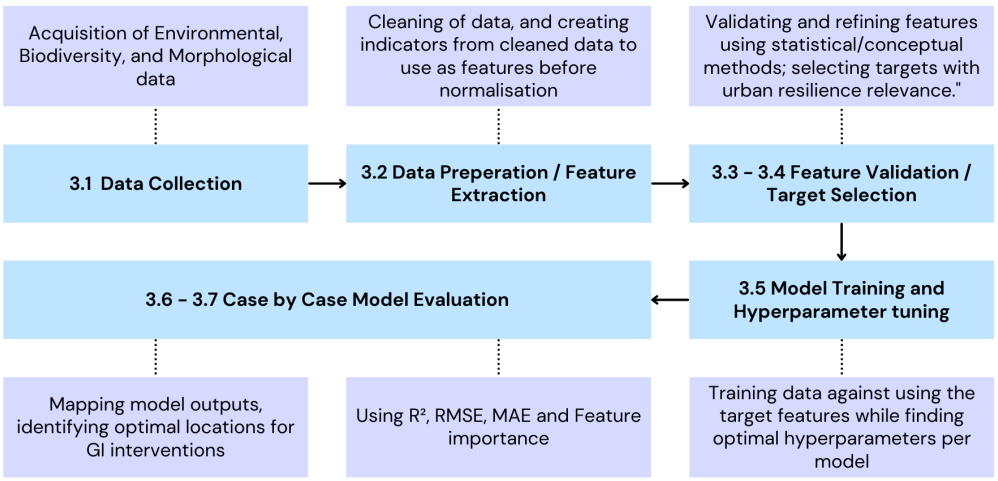

Figure 3.1.: Methodology Workflow

3.1. Data Collection

The first step in the methodology was the collection of raw data. The data collected in this study was obtained from multiple sources. Table 3.1 below shows the raw data type and source it was obtained from.

Table 3.1.: Raw Data Sources

Raw Data	Data Source	Relevance to the study
Urban Heat Island Raster (10m)	The Klimaat Effect Atlas	Used to analyse urban heat island effects by mapping temperature variations across different city areas. It helps in understanding thermal hotspots and urban cooling strategies.
Perceived Temperature Raster (5m)	The Klimaat Effect Atlas	Perceived temperature shows what the current air tempereature feels like to residents. It again helps in understanding thermal hotspots and urban cooling strategies.
Tree Locations (.shp)	Rotterdam Municipality	Used to analyse tree species biodivesity as well as overall tree distribution. This helps in understanding biodiversity of an area as well as the overall flora population.
Satellite Imagery Raster (Band 04 and 08) (20m)	Copernicus	Used to calculate Normalized Difference Vegetation Index (NDVI) to analyse the density and health of flora within the area. This helps in understanding green hotspots as well as areas that may need immediate attention.
Imperviousness Raster (10m)	Copernicus	Used to show the percentage of built up areas compared to permeable surfaces. This aids in highlighting areas where GI implementations can benefit.
Road Centrelines (.shp)	NWB / OSM	This is used in network analysis, where features such as angular betweeness and attraction reach are calculated.
Road Polygons (.shp)	Rotterdam Municipality / BGT	This is the level of resolution that this study will be based on. Street functions will be used as features and street IDs are used for spatial joins.
Digital Surface Model (5m)	AHN4 / BGT	This is used to calculate skyview factor which informs on the relationships between morphological factors and environemntal factors.

Data Completeness and Coverage

In general, almost all data sets used in this study offer full spatial coverage of the study area. However, there are some limitations related to completness and temporal consistency. UHI and Perceived Temperature (PET) rasters, both sourced from the Klimaateffectatlas provide complete coverage but are from 2017 and 2022 respectively. Similarly, the satellite imagery from Copernicus to calculate the NDVI and the imperviousness rasters were also captured at differnt times. Imperviousness was captured in 2018, while the satellite imagery was more up to date where a date within 2025 was selected to ensure the most up to date representation of vegetation cover. This temporal inconsistency may introduce minor discrepancies when comparing the two features, particularly in areas undergoing recent development or land-use change.

The tree location dataset, obtained from the Rotterdam Municipality represents public tree inventories and is not complete ion regards to coverage of urban greenery in public spaces. It covers the general pattern and contains most trees. However, private trees within residential areas and parks are not included. This may result in an underestimation of overall tree density and biodiversity in residential or commercial areas.

The Digital Surface Model (DSM), used for the computation of the Sky View Factor, was derived from the AHN4 dataset. While the AHN4 provides high spatial resolution, it is updated in a four year cycle, meaning that parts of Rotterdam may be based on data collected in different months. This temporal variation could affect consistency, particularly when comparing it with more current layers.

Finally, the road polygon and centreline datasets from the Rotterdam Municipality and NWB/OSM are complete for the urban area of interest and should provide sufficient resolution for street level analysis. These datasets serve as the base for spatial joins and network analysis to align well with the overall scope of the study.

3.2. Data Preperation

After collecting the raw data, some data processing steps were needed to create indicators that would later serve as features for the machine learning algorithm. Each feature will be discussed in the following subsections.

The main dataset used for this analysis is the road polygon dataset obtained from the Rotterdam Municipality. Compared to the road dataset from the BGT, this dataset only includes urban areas. As this study focuses only on urban areas in Rotterdam, excluding industrial and port zones, the dataset lines up well with the aim of this thesis. Following feature extraction in the next section, the features will be aggregated so that a dataset at street level resolution will be created. This allows for the data per street segment to be used as samples in the machine learning tasks described in section 4.1.

3.2.1. Feature extraction

The features extracted for machine learning will be showcased in this section by category. The indicators have been split into 3 main categories, environmental, biodiversity, and morphological.

Environmental Features

Environmental features are important in the context of urban resilience as these highlight the overall state of the built environment which may have climate related concerns and ecosystem functions. In this study, four features were selected, UHI, PET, NDVI and imperviousness. Collectively, these metrics aim to capture different environmental aspects of the built environment. UHI can be directly related to heat retention while PET being a calculated metric includes humidity and wind. NDVI and imperviousness show the quality and extent of vegetation cover and built up surfaces. High impervious surface cover is associated with

elevated flood risk and thermal discomfort, while high NDVI values indicate greener areas and higher cooling potential [Kabisch and Van Den Bosch, 2017].

Table 3.2.: Environmental Features

Feature	**Figure**	**Relevance to Urban resilience**
Urban Heat Island	3.2a	Measures temperature differences between urban and rural areas due to heat retention by built surfaces. Important for identifying heat stress zones and planning GI interventions.
PET	3.2b	A metric incorporating temperature, humidity, wind, and radiation to reflect how hot it ”feels” to humans. Useful for assessing thermal comfort and urban livability.
NDVI	3.2c	Measures vegetation health and density using satellite imagery. High NDVI indicates lush greenery, contributing to cooling and biodiversity.
Imperviousness	3.2d	Quantifies the extent of paved and built-up surfaces, which influence runoff, flooding, and local climate. Essential for assessing stormwater management needs.

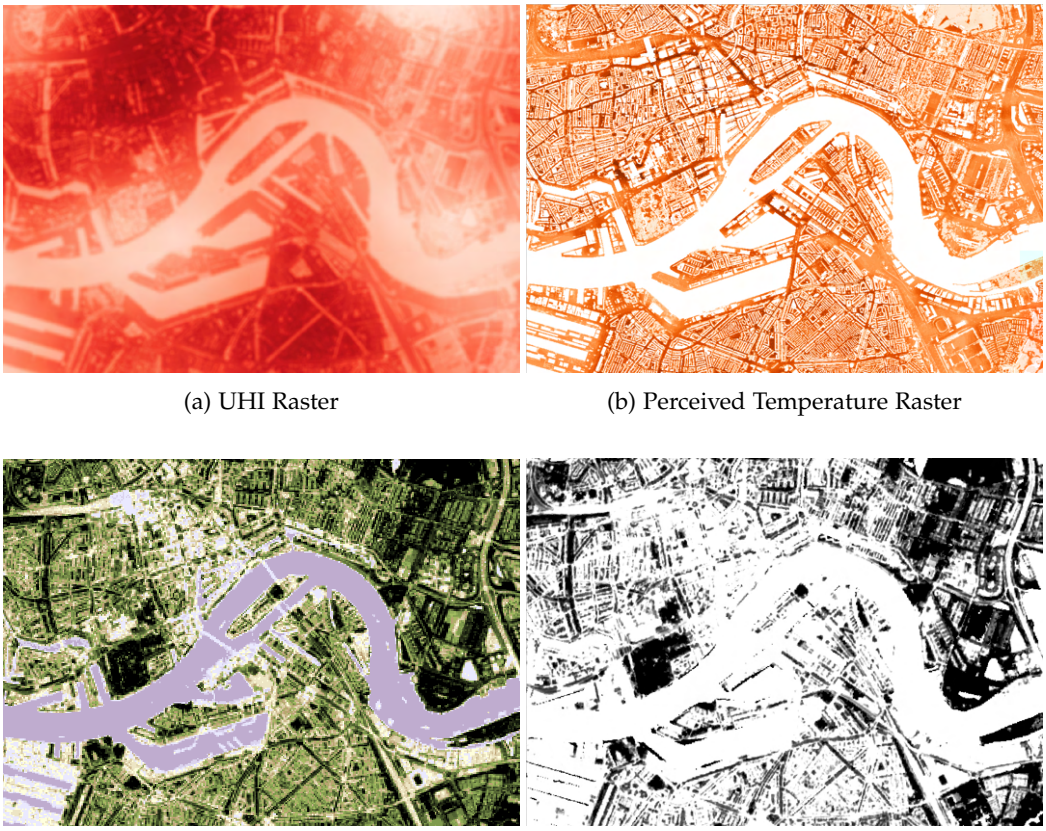

(a) UHI Raster

(b) Perceived Temperature Raster

(c) NDVI Raster

(d) Imperviousness Raster

Figure 3.2.: Overview of environmental indicator rasters

Urban Heat Island, Perceived Temperature, Imperviousness

To clean and extract data from the UHI, PET, and Imperviousness rasters, a series of steps were followed. Firstly, using QGIS the Coordinate Reference Systems (CRS) was converted form WGS:84 into EPSG:28992 in order to align the Dutch coordinate system and ensuring consistency across all datasets. Following this, the **Zonal Statistics** tool was used again in QGIS in order to compute the average value of the raster per street segment, using the street polygons as the base layer. This approach enabled the extraction of spatially aggregated values for each street. For reference, figure 3.3 shows an overlay of the road polygons on the perceived temperature raster to show how the the areas in which the values per street would be calculated.

NDVI

NDVI, on the other hand, had to be generated manually from satellite imagery. Firstly, sentinel-2 imagery was downloaded from the Copernicus open data platform, specifically Band 8 (Near Infrared) and Band 4 (Red). Once acquired, and converted to CRS EPSG:28992 using QGIS, the NDVI was calculated using equation 3.1 in the field calculator of QGIS.

Figure 3.3.: Road Polygons overlaid on Perceived Temperature Raster

Subsequently, similar to the other rasters, the **Zonal Statistics** tool in QGIS was used to calculate the average NDVI per street segment. Figures 3.4a and 3.4b show the rasters of band 4 and 8 respectively.

$$NDVI = \frac{NIR - RED}{NIR + RED} \tag{3.1}$$

(a) Band 4 (RED) (b) Band 8 (NIR)

Figure 3.4.: Overview of raster bands to calculate NDVI

Biodiversity Features

Urban biodiversity is a critical component of resilient and sustainable cities, offering many benefits that range from ecological stability to improved social well being [Filazzola et al., 2019]. This study includes two biodiversity features, specifically, tree species count which captures the diverse flora composition and total tree count, representing the density and extent of vegetative cover. These features reflect heterogeneity and potential species interactions within the urban space. The features are also linked to ecosystem services such as urban cooling and air purification [Kabisch and Van Den Bosch, 2017].

To extract the biodiversity related GI features at street level the total tree count per street and total tree species per street were calculated. Starting with the raw data, the public tree dataset, shown in figure 3.5b was used along with the buffered road polygon dataset in figure 3.5a. The **Join Attributes by Location (one-to-many)** tool was used with the **intersects** option to assign trees to each buffer they fell within. Since buffer overlaps could result in the same tree being linked to multiple street segments, duplicates were removed using the **Remove duplicate features by attribute** tool, with the unique tree ID field (ID) selected as the attribute. The **Statistics by Categories** tool was then used to calculate the total number of trees and the number of distinct species per street, based on the 'BOOMSORTIM' field. This allowed us to obtain the trees within the buffered areas while also selecting **count** and **unique** as the output to receive the number of trees that belong to each street as well as the number of distinct trees in each street.

Table 3.3.: Biodiversity Features

Feature	Figure	Relevance to Urban resilience
Overall Tree Count	3.5b	This measures the number of trees per street segment. This is useful in considering the GI density per area
Tree Species Count	3.5b	This measures the number of different trees per street segment. This is useful in considering the overall biodiversity of the area.

(a) Road Polygon Buffered by 10m (b) Tree Dataset

Figure 3.5.: Data needed to calculate biodiversity features

Figure 3.6.: Tree Polygons overlaid on Buffered Roads

Morphological Features

Urban morphology is critical in assessing Urban Resilience. The morphology for a city can be influenced by the arrangement of buildings, streets and open spaces. These in turn have

effects on urban climate resilience from factors such as airflow, heat retention and stormwater management [Stewart and Oke, 2012]. The morphological features extracted were Angular Betweenness and Attraction Reach. Angular Betweenness shows the significance of each street segment within Rotterdam by calculating how often the street is used between locations. Attraction Reach measures the number of set locations (green spaces in this case) within a specified distance. These reflect the overall accessibility and connectivity of Rotterdam between green spaces. Green spaces in this study are categorised as linear street trees, parks and urban forests.

Looking further than street networks, Sky View Factor (SVF) incorporates more dimensions in urban morphology. SVF quantifies the openness of the sky from any given point. This includes buildings and vegetations into the calculation by using a Digital Surface Model (DSM). The DSM used in this study obtained from the AHN, includes both green and grey infrastructure such as buildings and vegetation. SVF reflects how enclosed or exposed a location is using urban form, making it a valuable metric in assessing urban climate resilience.

The data extraction process was done using the Place Syntax Tool (PST) which is a plug-in in QGIS.

Figure 3.7.: Road Centrelines Displayed in Red

Firstly, Angular Betweenness was calculated using the PST and the `Angular Betweenness` function. This measures the importance of a street segment in a network, considering the angular deviation as the primary cost function instead of euclidean distance [Ståhle, 2012]. The equation the Angular Betweenness function uses is as follows:

Table 3.4.: Morphological Features

Feature	Relevance to Urban resilience
Angular Betweenness	This measures the importance of a street segment in the urban transport network by evaluating how often it falls on the shortest paths between locations. Higher connectivity values indicate streets that play a crucial role in linking green spaces, facilitating movement, and enhancing accessibility
Attraction reach	This measures the accessibility of green infrastructure (GI) by evaluating the number or presence of green spaces within a specified distance (e.g., 1 km) along the street network. It reflects how well-connected a segment is to nearby green spaces, supporting walkability, recreational access, and urban cooling.
SVF	This measures the ratio of the visible sky from a street segment, indicating how open or closed an area is. Lower values suggest a more dense urban form or tree canopy. This can reduce overall cooling efficiency and increase heat retention in the area.

$$B_c(x) = \frac{\log(B(x)+1)}{\log(2+\sum_{i\neq x} D(x,i))} \tag{3.2}$$

Where:

- A = the set of reachable attractions within given radius.
- D(o,a) = shortest distance from origin o to attraction a.

The **PST tool** outputs several variations of Angular Betweeness:

- AB: Raw Betweeness Score
- ABN: Normalised Betweeness
- ABT: Total betweeness across all paths
- ABM: Mean betweeness per shortest path

For this study, the Raw Betweeness Score was selected as a morphological indicator for green space connectivity. This helps to quantify how well different street segments allow for movements to and from green spaces. Higher values indicate streets that are well connected between green spaces indicating higher levels of accessibility across this street.

Following this, to evaluate proximity to green spaces, the `Attraction Reach` function function from the PST in QGIS was used. Compared to direct euclidean distance, using Attraction reach takes the street network into account. The road centrelines from OpenStreetMap were

used as the network with green space entrances selected as the destination points. A 1km walking distance threshold was applied to reflect a general walking range in urban areas. The attraction reach function calculates, for each object in a set of origin objects, the sum of all attractions from a set of attraction objects that can be reached by the network within a user defined radius. The function applies the following formula:

$$AR(o) = \sum_{a \in A} (f(a) \cdot w(D(o,a))) \tag{3.3}$$

Where:

- A = the set of reachable attractions within given radius
- f(a) = attraction value associated with attraction a, or 1 if no attraction value is used
- D(o,a) = shortest distance from origin o to attraction a
- w(x) = attenuation function

Before using the tool, it is important to note, the PST inputs require the use of road centrelines as the street networks. Consequently, these road centrelines need to be broken into single line segments at each intersection before use. To handle this condition, the **explode** tool was used to seperate the road network into single line segments compared to multi line segments or poly lines. Once angular betweenness and attraction reach were calculated, the tool **Join by Location (Summary)** was used to join the newly processed data back to the base road network. This approach provided a spatially aware method of measuring accessibility. Streets with low attraction reach values indicated areas where GI was lacking, making them potential priorities for intervention or investment in future planning strategies.

SVF as shown in figure 3.8 was calculated using the UMEP plugin in QGIS. The required input uses a DSM of the study area. The DSM was obtained from the AHN4 dataset, where the raster was created at a 5m resolution. As this is a computationally heavy task, the large raster was turned into a grid by firstly using the **create grid** tool before being tiled using the **Zonal Statistics** tool to apply the full SVF raster to the grid. The Sky View Factor function was then ran in batches. Once the sky view factor raster was available, it was then merged to the main dataset through the use of **Zonal Statistics** again in QGIS, taking the mean of each street segment and appending it per street.

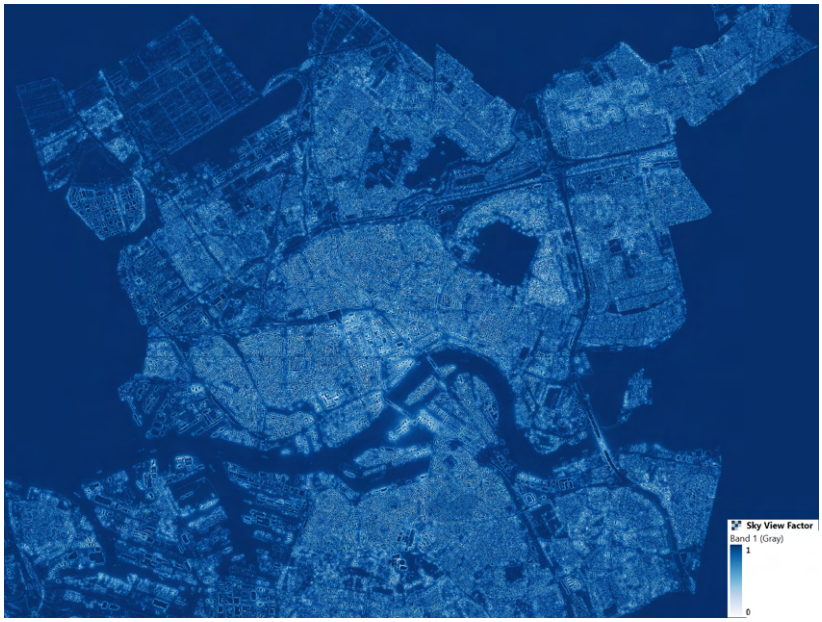

Figure 3.8.: Sky View Factor Raster Merged

3.2.2. Normalisation of features

In order to ensure that the features were comparable, normalisation is first applied to transform the variables to ensure the scales are the same. This is done because, the raw extracted features are all on different scales. Feature normalisation is necessary to ensure there is not a single overpowering feature.

The feature scaling method used in this paper is min max scaling. This method transformed all features to a range between 0 and 1, while also preserving the original distribution but ensuring no variable is outside of the range [de Amorim et al., 2023]. The formula used for the Min-Max scaling is given as:

$$X' = \frac{X - X_{\text{min}}}{X_{\text{max}} - X_{\text{min}}} \tag{3.4}$$

Where:

- **X'** is the normalised value,
- **X** is the original feature value,
- X_{min} and X_{max} are the minimum and maximum values of the feature.

Although Random Forest and Gradient Boosting are not particularly affected by normalisation as they are ensemble techniques using decision trees, SVR benefits significantly from this normalisation process especially when using kernel functions that may be sensitive to scales [de Amorim et al., 2023]. Using it in all models ensures consistency across the different models.

3.3. Feature Validation

Feature validation in this study followed a two step process. First, using statistical methods, before using a conceptual validation to ensure the features are theoretically relevant to urban resilience.

Feature selection was important as it reduced the overall number of input variable that the model has to account for. This in turn lowered the computational cost of training and made feature importance results more interpretable. Furthermore, model performance was potentially be improved as the algorithm did not focus on fitting irrelevant features.

The primary statistical tool used for determining feature selection was the correlation matrix. A correlation matrix visualises pairwise relationships between all normalised features. Features with higher correlation coefficients may contain redundant information and risk inflating the influence of specific variables within the model. By filtering out features with strong manual correlation, the remaining feature set achieved greater independence and interpretability [Katrutsa and Strijov, 2017].

To enforce this process, a Variance Inflation Factor (VIF) was also applied after correlation matrices filtering. VIF quantifies how much the variance of a regression coefficient is inflated due to multicollinearity. Features with VIF values higher than a common threshold were considered for removal to ensure greater independence within the feature set [O'brien, 2007].

The second stage of feature validation involved more of a conceptual evaluation. This was where selected features are assessed based on their relevance to established resilience theory and frameworks. Conceptual validation on top of statistical validation aids in capturing diverse and meaningful aspects of resilience beyond statistical behaviour.

3.4. Target Selection

In this study, resilience related metrics were used as dependent variables. The targets selected shown in figure 3.9 were aimed at representing real world cases that are linked to the effects or performance of GI within the built environment. The selected targets were chosen based on the data availability as well as the conceptual relevance to urban resilience. Each dataset aims to reflect a different dimension of resilience, including thermal comfort, ecological health, and climate vulnerability. All target layers were spatially joined to street polygons using either zonal statistics in QGIS, if the dataset was a raster, if not, it was joined by the function **Join by Location Summary** as a mean per street polygon. The three outcomes selected in this study are summarised as follows:

Table 3.5.: Dependent Variable Selection

Target	Source	Relevance to Urban Resilience
Cooling Effect of Green/Blue Space (Verkoelend effect van groen en water)	Klimaate Effect Atlas	Represents the relative cooling (°C) provided by green and blue infrastructure. Useful for identifying the effectiveness of existing vegetation and water in reducing local temperatures.
Pluvial Flood Depth (Water op straat bij extreme bui)	Klimaate Effect Atlas	Simulates depth of surface water accumulation during an extreme rainfall event. A proxy for stormwater resilience and GI based infiltration potential.
Nitrous Dioxide Concentration (NO_2)	Overheid.nl	Indicates the concentration of traffic-related air pollution, which impacts respiratory health and environmental quality. Useful for identifying areas where green buffers or pollution-mitigating strategies may support urban resilience.

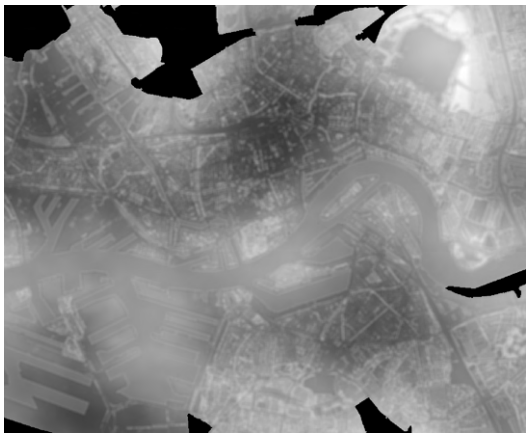

(a) Cooling Effects of Blue Green Infrastructure Raster

(b) Max Flood Depth Raster

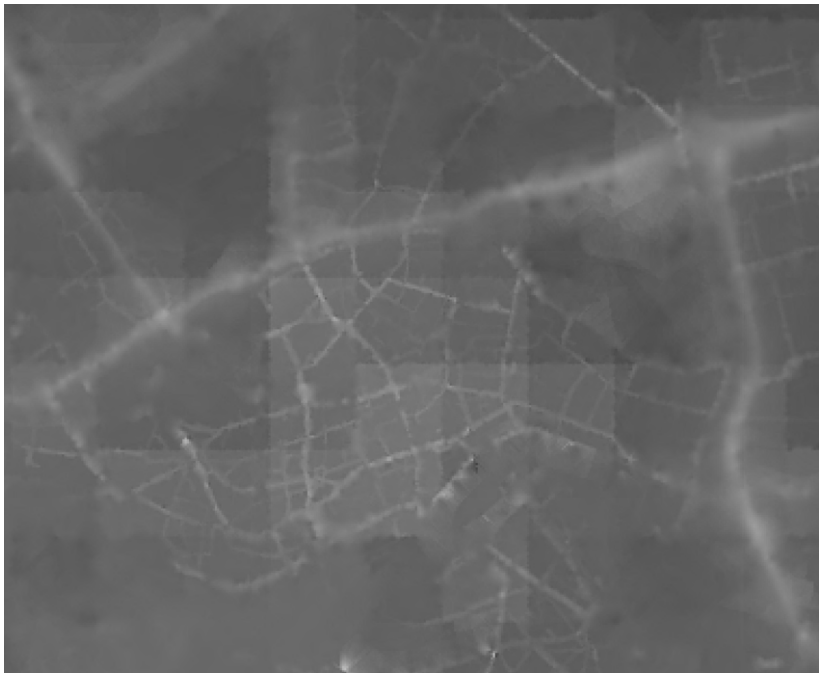

(c) NO_2 Distribution Raster

Figure 3.9.: Target Rasters

3.5. Model Training and Hyperparameter Tuning

After preprocessing and feature extraction, the features and data were trained on 3 models. RFR, GBR and SVR. All models were implemented using the Scikit-learn library in Python.

In order to enhance each models performance and ensure model robustness, hyperparam-

eter tuning was be performed for all three regression models: Random Forest, Gradient Boosting, and Support Vector Regression. Using default model parameters would have been inefficient for the complexity of the data. The random search method was utilised to systematically investigate combinations of key hyperparameters.

The random search was used to try a specified number (100) of parameter combinations using cross validation. A 10 fold cross validation approach has been used in this case to prevent overfitting and offering generalisability. The model was trained against 90% of the dataset and validated on 10% for every fold. The details of this can be found in section 3.7. The parameters with the best average R^2 value across folds were selected as the best setting.

Random Forest and Gradient Boosting are tree based models. These share several core hyperparameters related to decision tree construction. However, they are different in the way which the trees are built and combined [Friedman, 2001]. The shared hyperparameters are as follows:

- `n_estimators`: Number of trees in the ensemble
- `max_depth`: Maximum depth of individual trees
- `min_samples_split`: Minimum samples required to split an internal node
- `min_samples_leaf`: Minimum samples required at a leaf node
- `max_features`: Maximum number of features considered at each split

These parameters aid in regulating the balance between underfitting and overfitting. Trees that are shallow with more constraints can be better for generalisation but may not account for complexity in the data. Trees that are deeper can capture more variance, but have risk of overfitting [Bartz-Beielstein, 2023].

In RFR, bootstrapping is the option to determine whether each decision tree is trained on a random sample drawn with replacement from the original dataset. This introduces diversity among trees by training them on slightly different subsets of data [Breiman, 2001]. The focus of tuning in RF was to select a sufficient number of diverse trees while limiting depth control to prevent overfitting.

Conversely, gradient boosting builds trees sequentially, using each tree to try and correct the residual errors of previous trees. This in turn is more sensitive to overfitting and parameter choices [Friedman, 2001]. In addition to the base tree parameters that are shared with RF, `learning_rate` was also tuned. The learning rate is an important hyperparameter as it controls the contribution of each new tree to the overall model. It is used as a scaling factor on the predictions provided before they are added to the final output. Lower learning rates require more trees but may produce better generalised results [Friedman, 2001].

SVR tuning utilises different hyperparameters compared to RFR and GBR. Instead, these hyperparamters focus on margin control and kernel based transformations [Awad et al., 2015]. The tuning was centered around these three essential hyperparamters:

- `C`: The regularisation parameter controlling the trade-off between smoothness and accuracy. Higher values penalise training errors more strongly.
- `epsilon`: Defines the width of the margin of tolerance around the true values within which no penalty is applied. Tuning `epsilon` affects model sensitivity to small deviations.

- `kernel`: Specifies the function used to map data into higher-dimensional space. The grid included `rbf`, `poly`, and `sigmoid` kernels to explore linear and non-linear relationships.

SVR is sensitive to feature scales, which is why all features were normalised before model fitting as as discussed in Section 3.2.2.

3.6. Model Evaluation

In order to assess the performance of these models, the study applied three standard evaluation metrics. The coefficient of determination R^2, RMSE and MAE. These metrics provide insight into overall model accuracy and reliability. Additionally, k fold cross validation was utilised to ensure robustness of the model across different subsets of data.

3.6.1. R^2 Calculation

R^2 is used to measure the proportion of variance in the dependent variable that can be predicted from independent variables. Higher R^2 values indicate higher accuracy and reliability. In the context of this study, the higher R^2 score suggests the model captures a greater proportion of the spatial or statistical variation within the range of our indicators [Plonsky and Ghanbar, 2018]. The R^2 formula is represented by:

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2} \tag{3.5}$$

Where:

- y_i: the observed value
- $\hat{y}_i$: the predicted value
- $\bar{y}_i$: the mean of observed values
- n: the number of observations

3.6.2. RMSE Calculation

RMSE represents the average magnitude of prediction error by measuring the square root of the average squared differences between observed and predicted values [Chai and Draxler, 2014]. Lower RMSE values show better predictive accuracy. This is useful in detection of poor model fitting in areas with inconsistent data. RMSE penalises larger errors more heavily than other metrics such as Mean Absolute Error[Chai and Draxler, 2014]. The RMSE formula is represented by:

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{3.6}$$

Where:

- y_i: the observed value
- $\hat{y}_i$: the predicted value
- n: the number of observations

3.6.3. MAE Calculation

MAE, measures the average magnitude of prediction errors in a set of predictions. This represents the average absolute difference between observed and predicted values [Willmott and Matsuura, 2005].

The MAE formula is represented by:

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \tag{3.7}$$

Where:

- y_i: the observed value
- $\hat{y}_i$: the predicted value
- n: the number of observations

3.7. Case based evaluation

In this section, the outline of the process followed to evaluate the model performance individually for each case is shown. This involved a multi step analysis to assess the predictive accuracy, interpretability and spatial applicability of each model variant.

Firstly, performance evaluation per model is studied. For each case, the performance of 3 machine learning models are compared, using the coefficient of determination score (R^2), MAE, and RMSE. To ensure generalisability and detection of potential of under and overfitting, k-fold cross validation is applied during training so that a cross validation score was available for comparison.

K fold cross validation is a model validation technique that enhance reliability of performance estimates by partitioning the dataset into k sized subsets [Fushiki, 2011]. In each iteration, one fold is used for testing while the remaining folds are used for training. In the case of this study, 10 fold cross validation will be used, meaning 9 folds will be used in training and 1 fold will be used for testing. The average performance across folds delivers a

more robust estimate of model accuracy because it reduces risk of overfitting and the result is not swayed by a single biased train test split [Arlot and Celisse, 2010]. Figure 3.10 gives an overview of the validation partitioning process.

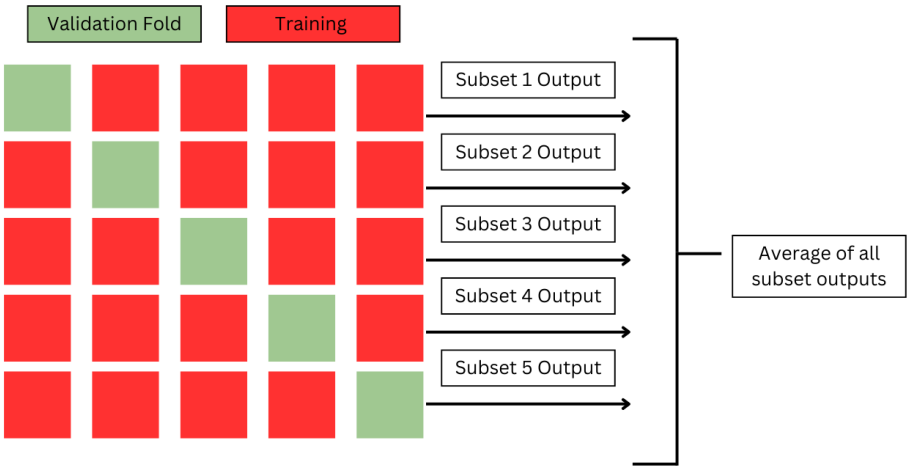

Figure 3.10.: Cross Validation Diagram adapted from [Scikit-learn, 2021a]

Ten folds were selected as a practical compromise between computational efficiency and evaluation reliability. While increasing the number of folds can lead to more precise estimates, in turn it increases computation cost. Given the complexity of the dataset, it was determined that ten folds would be a suitable balance for this study.

Secondly, as each model went through hyperparamter tuning using random searches, the best performing hyperparameters per case are recorded for reproducibility.

Following Model training, feature importance is analysed to identify which features most influence the model predicitions per case. This is done to evaluate model behavior and understand how different climate resilience related features contribute per target.

For the tree based models, Random Forest and Gradient boosting, feature importance was extracted using the `feature_importances` attribute from the `sklearn` package. This method is based on how much each feature reduces prediction error across the trees created in the decision making process. Features that contribute more to splitting the data and reducing overall variance in the results may receive higher feature importance scores [Scikit-learn, 2021b].

SVR on the other hand, has no native feature importance function. In this case, permutation importance was used instead. This works by randomly shuffling the values of each feature and measuring the decrease in the models R^2 score. The higher the decrease, the more important that feature is to the prediction. This learns a global relationship and does not reply on the splits of decision trees. By combining both built in and permutation based methods, each model could be interpreted in its own terms while still offering a consistent way to compare feature influence across models [Altmann et al., 2010].

Finally, spatial maps are generated from the results. The best performing model for each case is selected to generate spatial predictions. The chosen visualisations are, the observed values which are used as the ground truth to compare to, the predicted values, which are the model output, and finally residual maps, where the difference between the observed and predicted values are plotted. This allows the assessment of both predictive strength and spatial distribution of model accuracy. These steps provide a structured and spatially aware comparison of model outcomes across cases.

4. Model Implementation and Results

The use of the predictive modelling framework within the context of the case study is detailed in this chapter. The complete data set which was developed following the steps in section 3.2 is cleaned and preprocessed to prepare training and testing data segments. Prior to training, the data undergoes necessary cleaning and preprocessing, including normalisation and feature transformation, which is important to some algorithms like SVR as discussed in 3.2.2.

Once the data is prepared, three machine learning models are trained on the selected resilience-related outcome variables on a case by case basis. To improve the model performance and ensure the model is robust, hyperparameter tuning is performed using grid searches and cross validation. This allows for the evaluation of different values of model parameters to be tested in order to find the best fitting parameters to accomplish the highest prediction accuracy.

4.1. Data Cleaning and Model Training

4.1.1. Data Cleaning

To begin the model application process, the data was cleaned and processed in order to have consistent inputs for each model. This began with duplicates and null value detection followed by the extraction of relevant features such as UHI, tree species count and imperviousness. Outlier detection was not applied in this step as this study aims to represent the full extent of urban conditions across all urban areas within Rotterdam. Excluding street segments with higher or lower values could conceal relevant spatial patterns that may hinder the identification of areas in current need of GI interventions. Retaining all data points ensures a comprehensive and inclusive assessment of urban resilience across the entire study area.

4.1.2. Statistical Feature Validation

Feature validation was performed to remove redundancy from the dataset. A correlation matrix of normalised features was generated to identify multicollinearity and overlapping information between indicators. This is shown in figure 4.1. High pair correlations suggest redundancy and may skew model results due to multicollinearity or unfairly amplify a particular feature.

The features were normalised using min max scaling as detailed in section 3.2.2. This was done in order to ensure comparability between features of different units and ranges. However, features that are negatively associated with urban resilience, such as UHI, PET, and imperviousness were inverted after normalisation by subtracting their values from one. This

ensured that higher values consistently represent more resilient urban conditions. Table 4.1 shows the examples of the data being normalised.

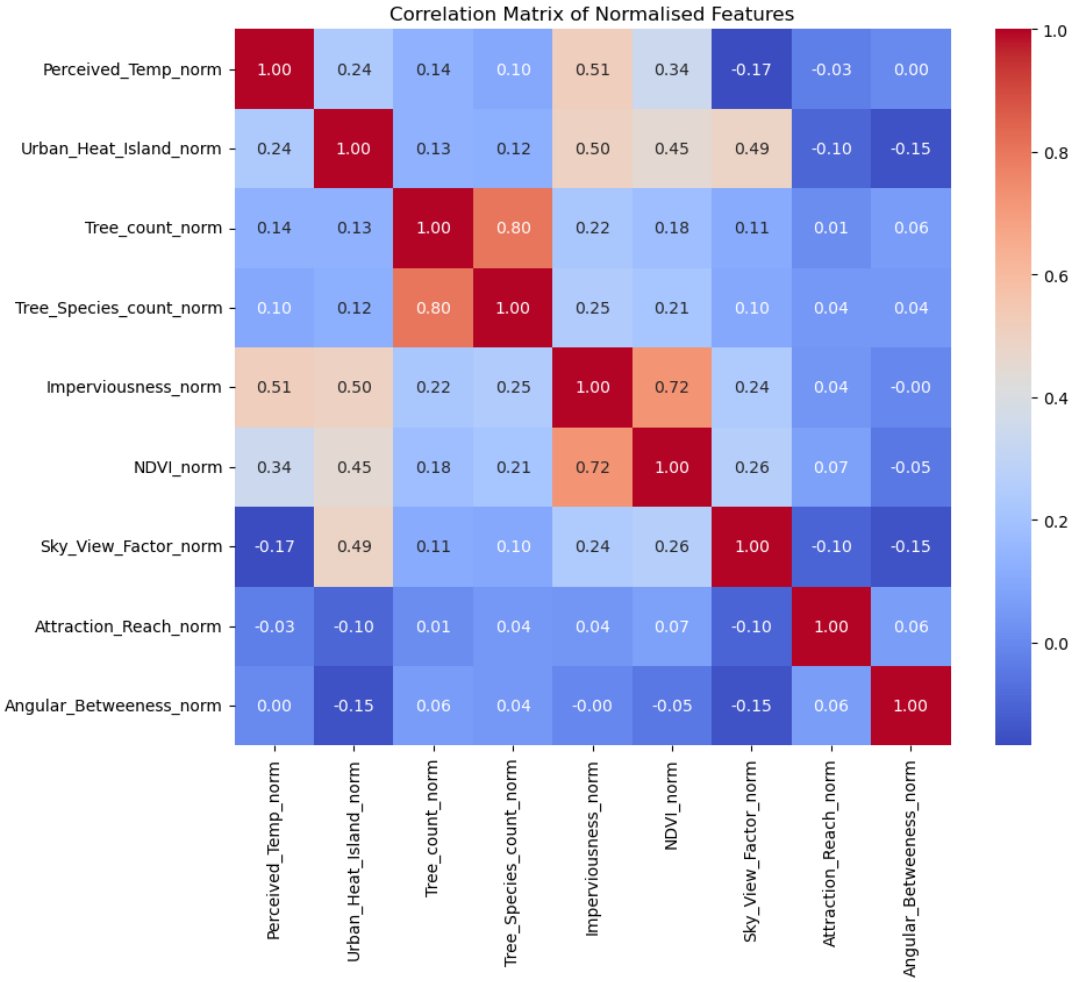

Figure 4.1.: Correlation Matrix of Normalised Features

Table 4.1.: Original vs. Normalised Feature Values Examples

Feature	Original Value	Normalised Value
Perceived Temperature	41.61904	0.29206
Urban Heat Island	0.80650	0.69267
Tree Count	5.00000	0.00222
Tree Species Count	4.00000	0.01948
Imperviousness	56.44431	0.43555
NDVI	0.37036	0.63644
Attraction Reach	10.60000	0.31718
Angular Betweeness	35891.10000	0.02232

Table 4.2.: Variance Inflation Factor (VIF) values for normalised features before feature reduction

Feature	VIF
NDVI	29.09
Sky View Factor	20.93
Urban Heat Island	12.94
Perceived Temp	10.65
Imperviousness	4.74
Tree Species count	4.31
Tree count	3.48
Attraction Reach	3.20
Angular Betweenness	1.34

Table 4.3.: Highly correlated feature pairs identified in the correlation matrix.

Feature Pair	Correlation	Notes
Tree Count & Tree Species Count	0.8	Extremely high – likely redundant
Imperviousness & NDVI	0.76	High correlation

Following the correlation matrices and VIF analysis, NDVI was removed from the final feature set due to its high multicollinearity with imperviousness. Although NDVI is commonly used as a vegetation indicator, its conceptual overlap with impervious surface coverage and a VIF value exceeding 29 suggested redundancy. Since both features capture surface characteristics, imperviousness was kept as it provided more distinct statistical behaviour

across the dataset. Removing NDVI led to a noticeable reduction in overall multicollinearity, improving the balance and stability of the remaining feature set.

Tree Count was also excluded from the final feature set due to its strong correlation with Tree Species Count. After normalisation, both variables followed similar spatial patterns, making the inclusion of both statistically redundant. Tree Species Count was prioritised as it captures both the presence of vegetation and ecological richness, aligning more closely with biodiversity-related resilience outcomes. Removing Tree Count improved the independence of features without compromising the conceptual coverage of GI in the analysis.

By reducing feature redundancy, the remaining variables capture distinct aspects of urban resilience (e.g. thermal comfort, vegetation health, imperviousness) without inflating certain indicators due to overlap. This step improves both model stability and interpretability.

4.1.3. Conceptual Feature Validation

After statistical validation to select features taking into account multicollinearity and redundancy, conceptual feature validation is performed. This is essential in ensuring that not only are the features statistically valid, but the indicators also align with theoretical understandings of urban resilience. This study uses a categorised framework adapted from Meerow et al. [2016] on the creation urban systems for conceptualising urban resilience.

From list 4.3, two features were initially dropped. Following this, considering conceptual validation, The remaining features were retained. While the correlation coefficient may suggest mid level potential redundancy, species richness aids in capturing essential dimensions in biodiversity to support ecology and function of the environment. This is supported by Livesley et al. [2016] where it is highlighted that both species richness as well as population is needed to support the ecology of an area. This is further supported by Ahern [2011] where it is stated that biodiversity aids in the overall resilience capacity of a city.

The final feature set was therefore selected not only for statistical independence but also for conceptual validity across the three categories:

- **Environmental**
  - **Mean Urban Heat Island value** – Represents localised heat stress and exposure [Amado, 2022].
  - **Mean Imperviousness value** – Captures surface sealing and runoff potential [Xue et al., 2024].
  - **Mean Perceived Temperature value** – Reflects thermal comfort, relevant to urban liveability and health [Amado, 2022].
  - **Sky View Factor** – Describes openness of urban form, linked to ventilation, thermal comfort, and radiation exposure [Dirksen et al., 2019].
- **Biodiversity**
  - **Tree Species Count** – Captures species richness and ecological diversity [Livesley et al., 2016].
- **Morphological**

- **Angular Betweenness** – Reflects spatial integration and connectivity potential within the street network [Pamukcu-Albers et al., 2021].
- **Attraction Reach** – Represents access to urban services and opportunities, reinforcing socio-spatial resilience [Pamukcu-Albers et al., 2021].
- **SVF** - Represents ratio of the visible sky, showing how open or closed an area can be, taking into account buildings and vegetation [Dirksen et al., 2019].

To ensure there is minimal redundancy, figure 4.2 and table 4.4 shows the final feature set. These features were retained based on the conceptual relevance as well as the statistical contribution to the correlation matrix. Although the VIF of Sky View Factor and Urban heat island are above 10, they were kept as multicollinearity does not have a significant effect on 2 out of the 3 models as they are tree based models. This aims to leave a balanced set of predictors to capture multidimensional aspects of urban systems, including environmental (e.g UHI), biodiversity (e.g tree species count), and morphological systems (e.g green attraction reach and angular betweeness).

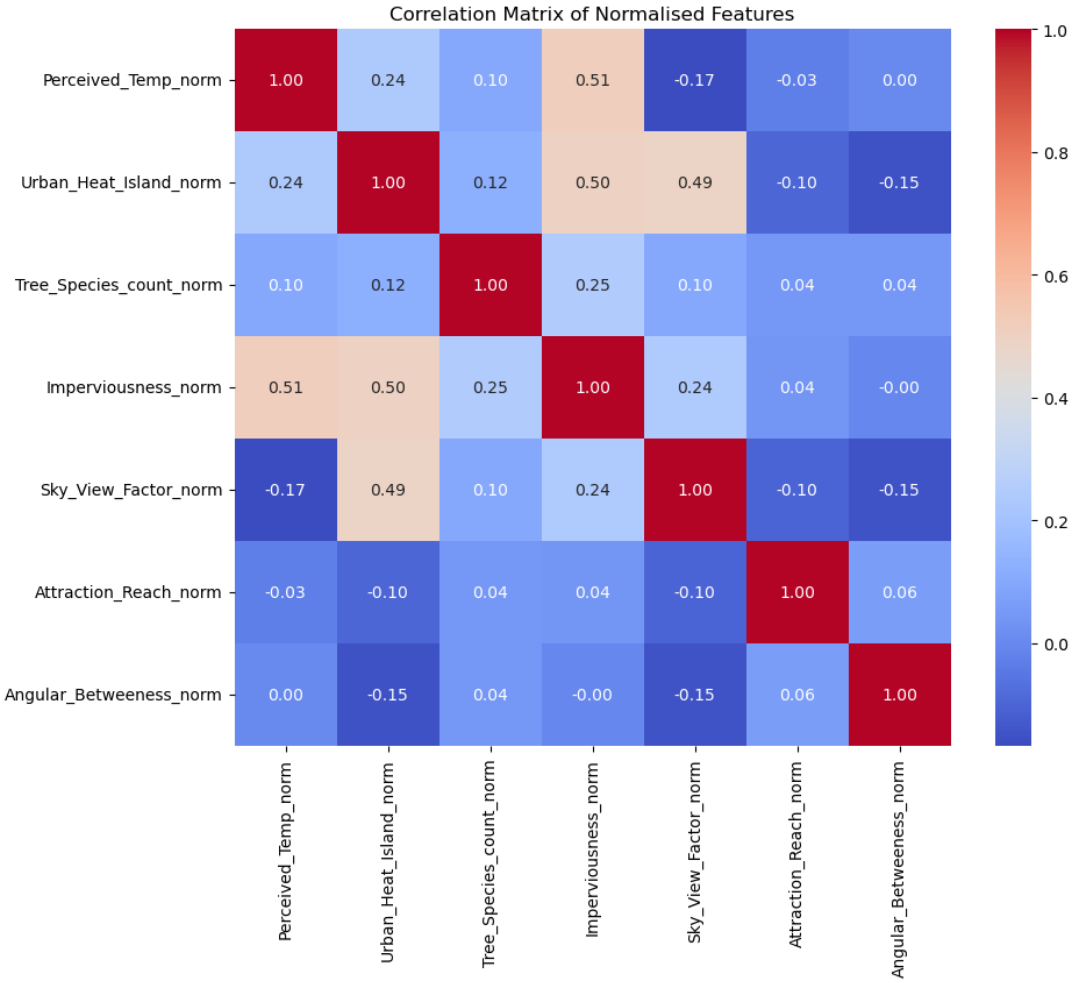

Figure 4.2.: Correlation Matrix of Selected Normalised Features

Table 4.4.: Variance Inflation Factor (VIF) values for normalised features after feature reduction

Feature	VIF
Sky View Factor	14.27
Urban Heat Island	12.82
Perceived Temp	9.12
Imperviousness	3.79
Tree Species count	3.01
Attraction Reach	1.60
Angular Betweenness	1.33

4.2. Hyperparameter Tuning Strategy

To keep the modelling consistent and allow fair comparison across different target variables, the same tuned hyperparameter grids were used for each model type in all three cases. This was done to support the aim of testing how well each urban resilience outcome could be predicted using the same framework, without the risk of introducing extra variation from model specific tuning grids. The hyperparamter grids were created with a focus on generalisability and robustness.

Random forest and Gradient boosting are both tree based models. However, they learn differently, which is why tuning these 2 models are slightly different even though these have the same hyperparameters.

Random Forest combines several relatively strong decision trees, it benefits from tighter regularisation. For example, the maximum depth is capped at relatively lower levels (4-12) to avoid the model learning small, unreliable patterns. Similarly, increasing the number of sample sizes required to split a node (5-25) or form a leaf (2-10) encourages more generalised trees. The range used to consider features at each split is from 2-4. This helps promote diversity across trees and reduce correlations within the ensemble. Overall, the hyperparameter grid aims to balance performance with generalisation, assuming that the spatial data may contain inconsistencies or be noisy. The parameters used to tune the Random Forest model can be fully seen in table 4.5.

On the other hand, gradient boosting builds trees one after another, with each one correcting errors from previous trees. Because of this, the algorithm is built upon a larger set of weaker learners. The number of estimators should be much higher compared to Random Forest (200-700 compared to 25 to 200). Similarly, the max depth per tree is also reduced to enforce the weak learning concept. Additionally, the learning rate is added to control how much each tree can contribute. A lower learning rate would slow the model down, encouraging it to learn more carefully and generalise better. Parameters such as minimum samples to split a node, minimum samples in a leaf, and max features are kept the same across both models to maintain some level of consistency and allow for fair comparison. Even though the algorithms behave differently, these parameters still play a role in reducing overfitting by limiting how deep or specific each tree can grow. Using the same values

ensures that any differences in model performance are due to the learning strategy itself, rather than additional variation introduced through tuning. Although both are tree based models, gradient boosting has a slightly higher risk of overfitting on complex data, this is the reason more regularisation is built into the parameter grid with larger number of estimators and a lower learning rate. The parameters used to tune the Gradient Boosting model can be fully seen in table 4.6.

Support Vector Regression is different from tree based models as it fits a global function to the data allowing for a margin of tolerance (epsilon) around the predictions. Using smaller epsilon values lead to more precise fits while larger ones are more forgiving. Because the precision and reliability of the urban resilience indicators can vary, the range (0.01 to 0.2) was used to allow the model flexibility in handling potential noise or uncertainty. The C parameter controls the trade off between regularisation and training error. Smaller values encourage the model to generalise better by allowing more training errors, and higher values prioritise accuracy on the training set. As spatial relationships in real time may be noisy or only partially explained by the subset of features, a wide range of C (0.1 - 500) was used on a log scale to test show strictly the model should respond to the data. The RBF kernel was chosen to allow the model to capture non-linear relationships in the input space. The gamma parameter plays a key role in defining how far each training points influence reaches. Low values result in smoother or broader decision boundaries, but higher values allow the model to react to local variations. The options scale, auto, 0.001, 0.01, 0.1 was used to provide a flexibility due to the risk of overfitting to noise in the dataset. The parameters used to tune the Support Vector Regression can be fully seen in table 4.7.

Table 4.5.: Hyperparameter Grid Used for Random Forest Regression Tuning

Hyperparameter	**Values Considered**
`n_estimators`	25, 50, 75, 100, 125, 150, 175, 200
`max_depth`	4, 6, 8, 10, 12
`min_samples_split`	5, 10, 15, 20, 25
`min_samples_leaf`	2, 4, 6, 8, 10
`max_features`	2, 3, 4
`boostrap`	True, False

Table 4.6.: Hyperparameter Grid Used for Gradient Boosting Regression Tuning

Hyperparameter	Values Considered
n_estimators	200, 300 ,400 ,500, 600, 700
max_depth	1, 2, 3, 4, 5
min_samples_split	5, 10, 15, 20, 25
min_samples_leaf	2, 4, 6, 8, 10
max_features	2, 3, 4
subsample	0.7, 0.8, 0.9, 1.0
learning_rate	0.01, 0.02, 0.03, 0.04, 0.05, 0.1

Table 4.7.: Hyperparameter Grid Used for Support Vector Regression Tuning

Hyperparameter	Values Considered
C	0.1, 1, 10, 50, 100, 250, 500
epsilon	0.01, 0.03, 0.05, 0.1, 0.15, 0.2
kernel	rbf
gamma	'scale', 'auto', 0.001, 0.01, 0.1

4.3. Case by Case Results

4.3.1. Predicting Cooling Effects of Green Blue Spaces (Case 1)

The cooling effect of green and blue spaces represents the extent to which vegetation and water bodies reduce local temperatures in urban areas. The map is derived from subtracting the actual Urban Heat Island from the Maximum Urban Heat Island. This shows the degree of cooling in °C that green and blue infrastructure provides. The map takes account for cooling effects from both local street trees to broader scaled parks. Only areas with a minimum surface sealing of 20% are include to ensure a focus on urban environments where cooling is most needed. This data is helpful in identifying areas where vegetation is currently contributing to thermal comfort and where future greening strategies could be prioritised [Rijksinstituut voor Volksgezondheid en Milieu (RIVM), 2020].

Model Performance

All 3 models were able to predict the cooling effects of green and blue infrastructure with moderate success. The results of the untuned metrics can be seen in table 4.8, and the results from the tuned parameters can be seen in table 4.9. Random Forest performed the best with the base parameters given by the `sklearn` package with an R^2 score of 0.45276, followed by Gradient Boosting (0.4317) and Support Vector (0.3975). The difference in R^2 may not seem

large. However, tree based models generally produced the lower MAE and RMSE, this may suggest that in case 1, these were better suited to capturing the underlying spatial variation in this target variable.

Moreover, after hyperparameter tuning, GBR produced the highest overall performance, improving to an R^2 of 0.47945, also with the lowest MAE and lowest RMSE in all models. This indicates the model was able to generalise more effectively after adjusting the learning process through the learning rate and max tree depth. After computing a training score and cross validation score to also determine if the model is under fitting, the training R^2 score had a value of 0.83847 and the validation R^2 score showed at 0.44911. The gap between the training and testing performance may imply that the model is learning noise or artefacts in the data that does not generalise well which is could be common in complex urban systems.

Support Vector showed the weakest performance both before and after tuning, with no real improvement. However, it produced the smallest gap between training R^2 (0.38725) and testing R^2 (0.39753). In case 1, SVR may have been limited by its assumption of a global function which might not aligned well, with non linear patterns in urban cooling.

Overall, Gradient boosting showed the best balance between accuracy and generalisation. random forest performed second best but leaned more towards overfitting while Support Vector was more stable but less accurate. These results highlight the importance of both hyperparameter tuning and model selection, particularly when working with spatially variable targets like cooling effects form urban green and blue infrastructure.

The learning curves for the best models per case can be seen in D.

Table 4.8.: Results from Untuned Parameters (Case1)

Validation Metric	**Random Forest**	**Gradient Boosting**	**Support Vector**
R^2	0.45276	0.43166	0.39753
MAE	0.04386	0.04768	0.06707
RMSE	0.13472	0.13729	0.14136

Table 4.9.: Best Hyperparameters and Performance Across Models (Case 1)

Random Forest		Gradient Boosting		Support Vector	
Param	**Value**	**Param**	**Value**	**Param**	**Value**
n_estimators	200	n_estimators	400	C	1
max_depth	10	max_depth	4	epsilon	0.1
max_features	3	max_features	4	kernel	rbf
min_samples_leaf	4	min_samples_leaf	4	gamma	scale
min_samples_split	5	min_samples_split	10	-	
boostrap	False	learning_rate	0.05	-	
-		subsample	1	-	
R^2	0.44911	R^2	0.47945	R^2	0.39753
MAE	0.04332	MAE	0.04399	MAE	0.06701
RMSE	0.13517	RMSE	0.13139	RMSE	0.14135

Feature Importance

Table 4.10.: Feature Importance by Model (Case1)

Feature	Random Forest	Gradient Boosting	Support Vector
Urban Heat Island	0.47571	0.48181	0.59311
Imperviousness	0.13748	0.11997	0.20291
Sky View Factor	0.12368	0.12108	0.10446
Angular Betweeness	0.05912	0.06416	0.04773
Attraction Reach	0.10636	0.11158	0.03743
Perceived Temperature	0.06268	0.06032	0.01215
Tree species Count	0.03493	0.04104	0.00221

Case 1 focused on predicting cooling effects of green blue spaces. Using feature importance, there was seen to be a clear pattern across the three models. UHI was consistently the most important predictor. This may be expected as the target variable is related to heating and cooling of green spaces.

In Random Forest, UHI, had a feature importance of 0.47571, followed by an imperviousness of 0.1375 and a SVF of 0.12369. These top 3 align with known urban cooling dynamics where impervious surfaces retain heat, SVF reflects exposure to open sky which can influence cooling. Additionally, Angular betweenness and atttraction reach show moderate contributions at 0.05913 and 0.10637 respectively. These help in reflecting how spatial movement patterns and accessibility may affect microclimate exposure.

Gradient Boosting showed a very similar ranking with a dominant UHI, followed by SVF and imperviousness. The distribution of feature importance was slightly more even in this

model. This is possibly due to the fact that gradient boosting is an iterative learning process where the less prominent features such as perceived temperature and Attraction reach are allowed to contribute meaningfully. This may suggest that gradient boosting was more sensitive to the diversity of features.

In contrast, SVR showed similar results but differ slightly as UHI was even more dominant with a feature importance score of 0.59311. Although this feature was much more dominant, the pattern remains the same in which SVF and imperviousness are close second. This reaffirms the models relation to environmental and morphological indicators in this case.

Overall, while UHI consistently led across all models, the tree based models were better at picking up a wider spread of spatial variables. SVR, though may have been more stable from the hyperparamter tuning results, seemed to capture a narrow band of predictors in this case.

Additional feature importance plots can be seen in B.3

Spatial Distribution

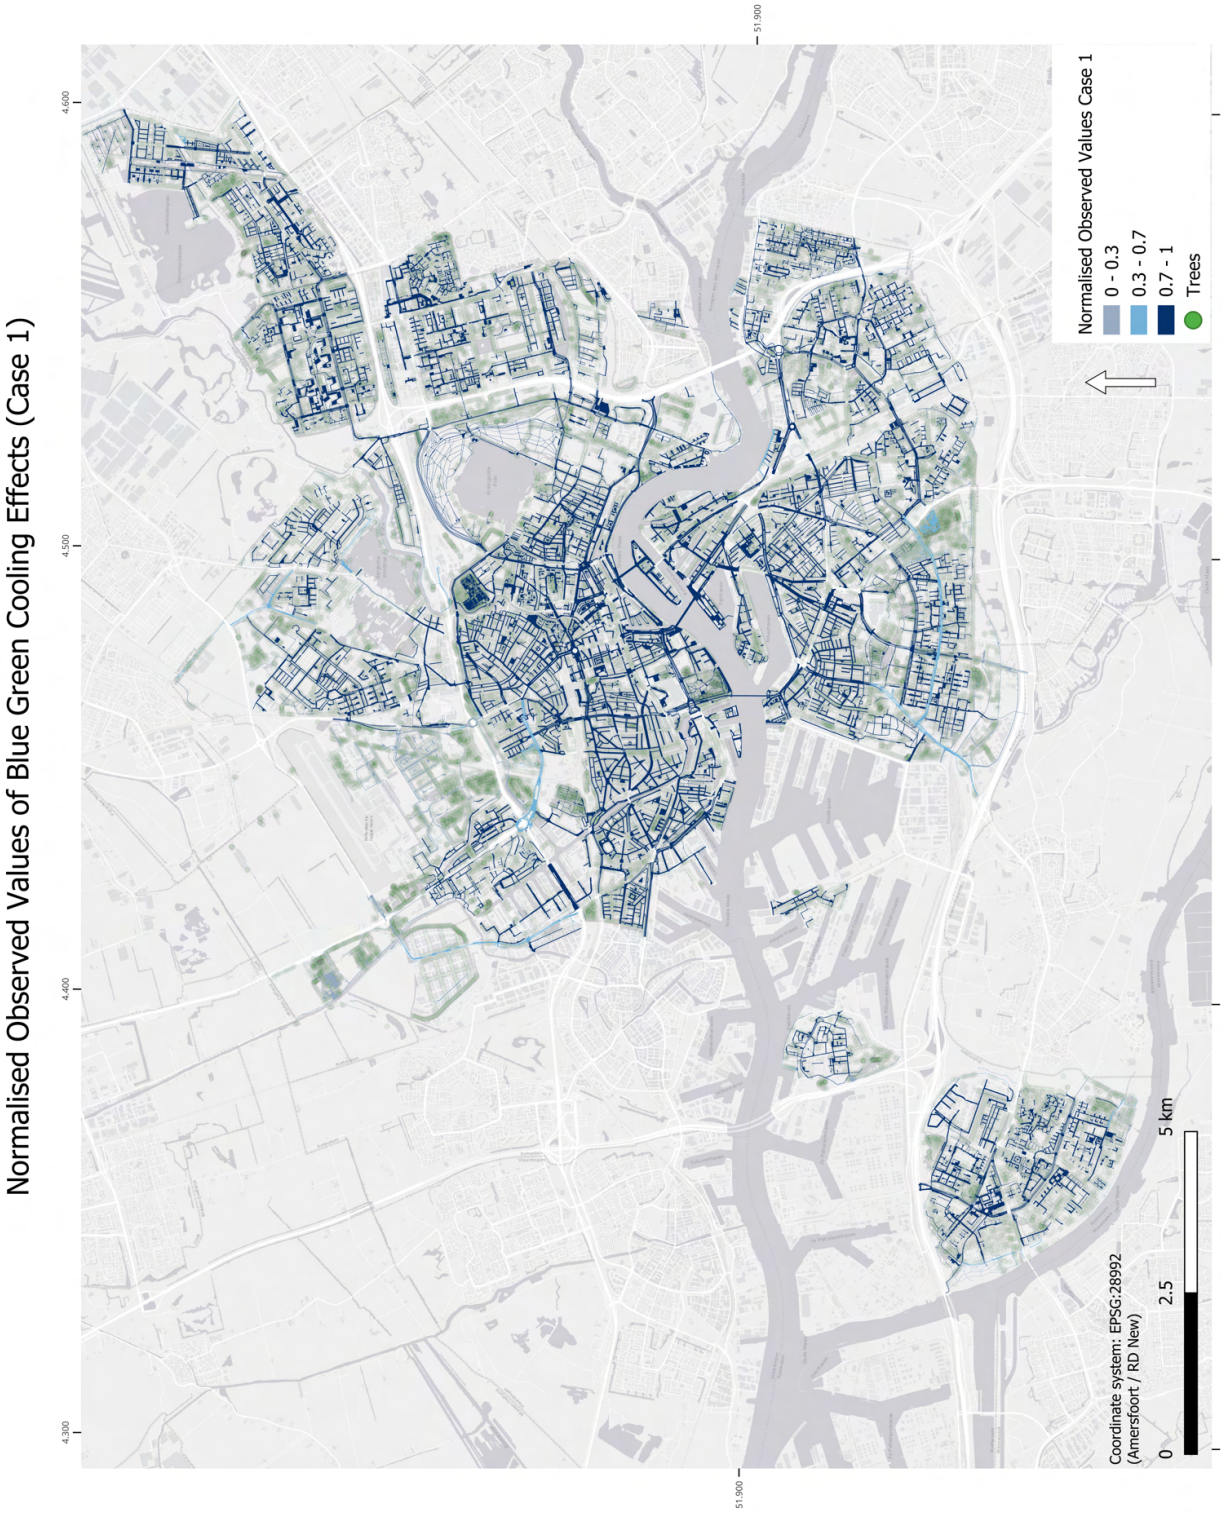

Figure 4.3.: Observed Normalised Cooling Effect of Green and Blue Infrastructure (Case 1)

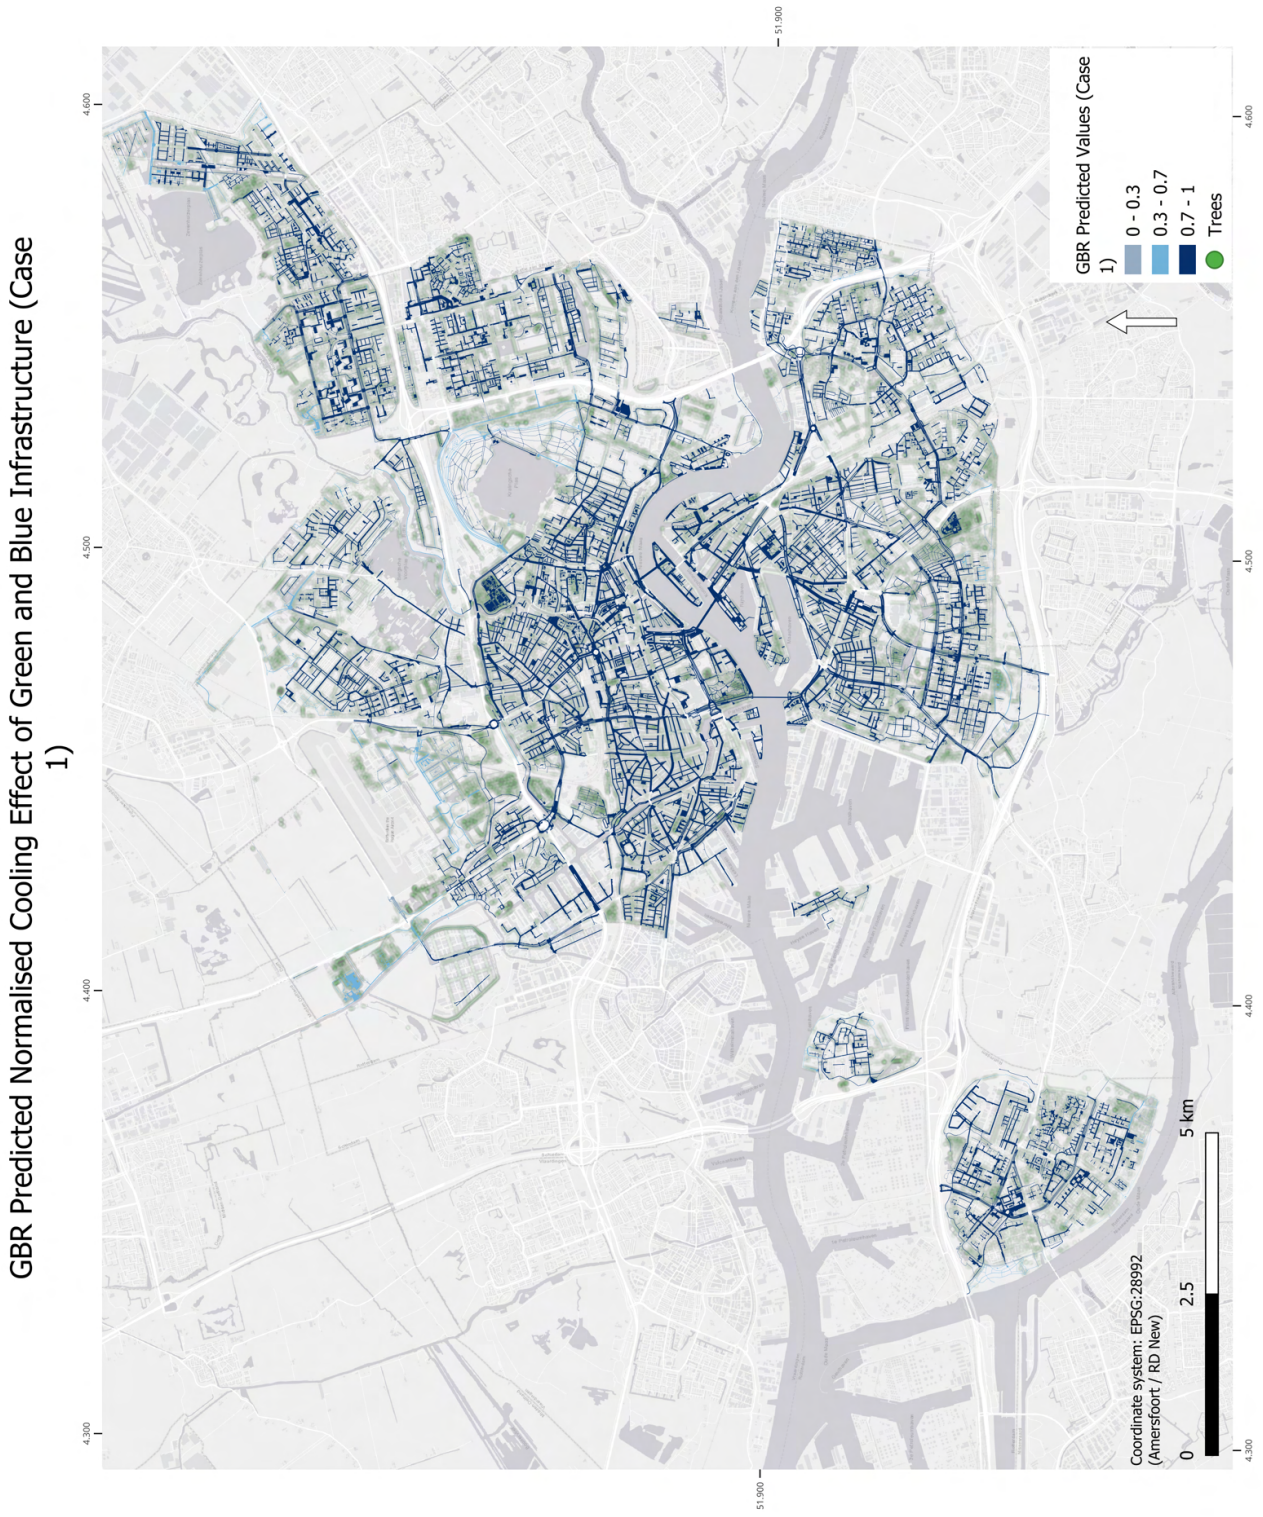

Figure 4.4.: GBR Predicted Normalised Cooling Effect of Green and Blue Infrastructure (Case 1)

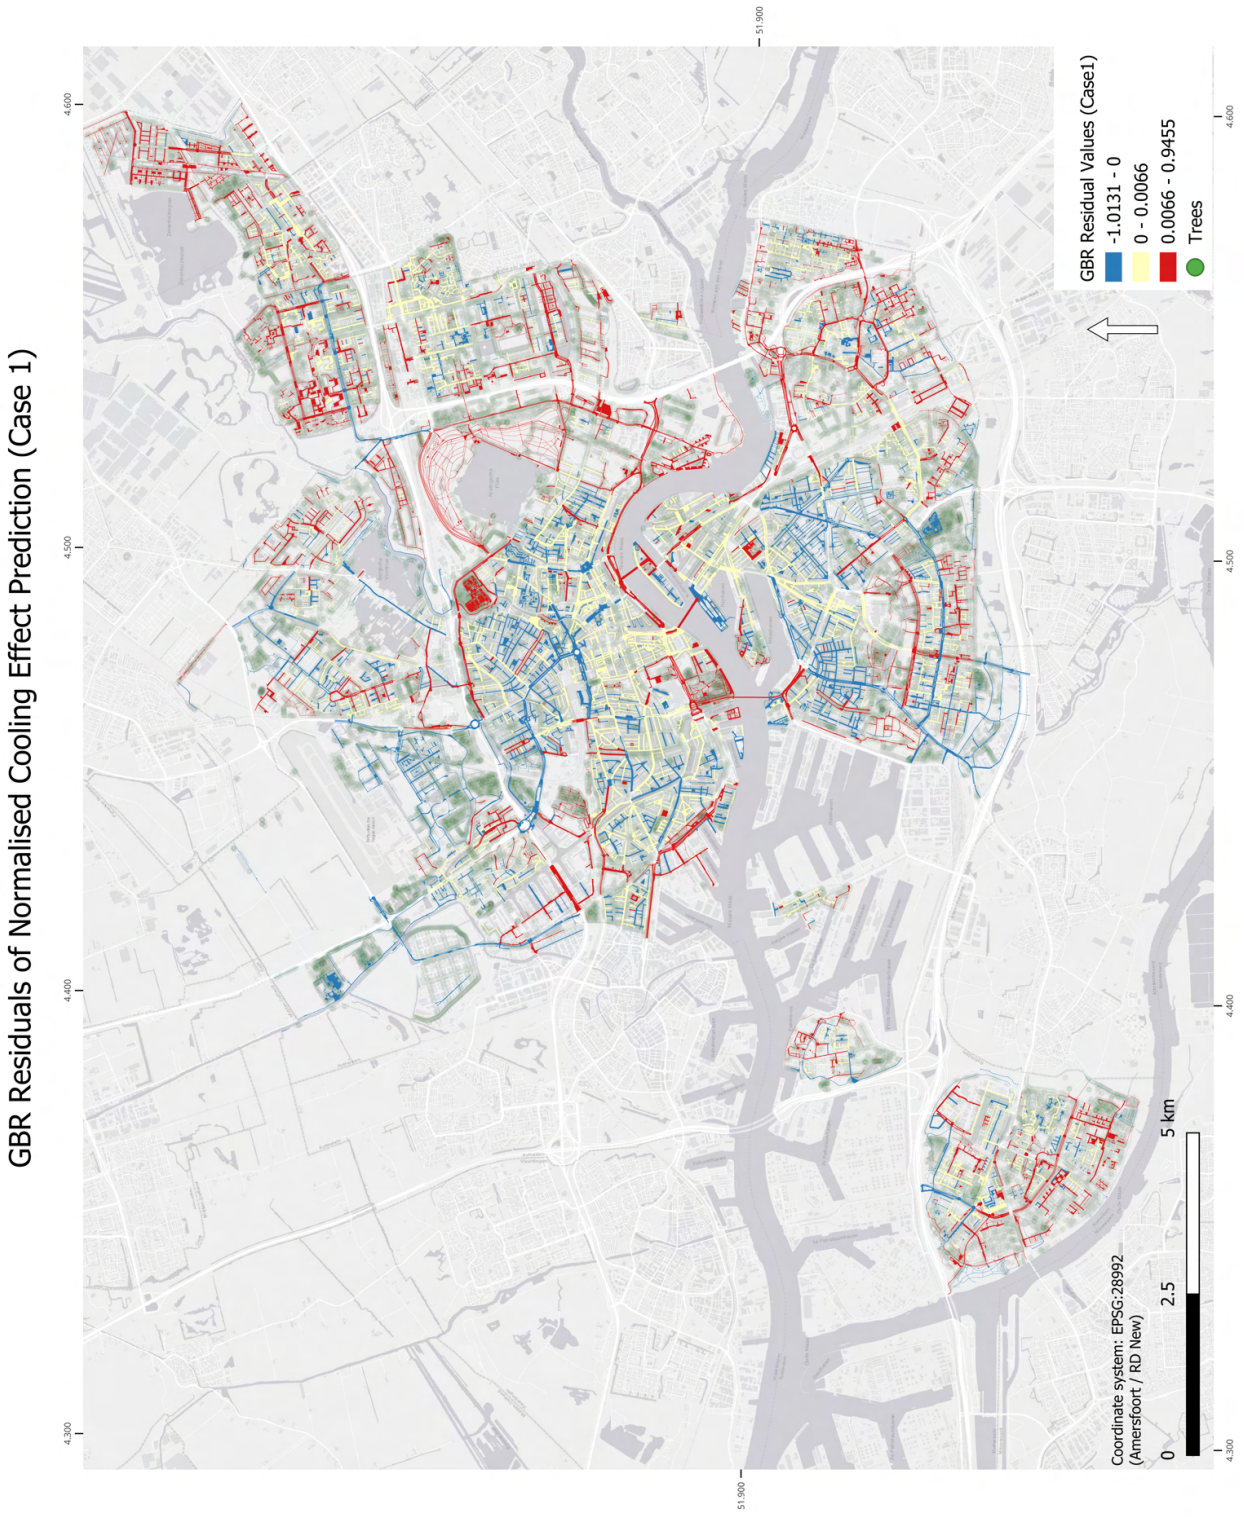

Figure 4.5.: GBR Residuals of Normalised Cooling Effect Prediction (Case 1)

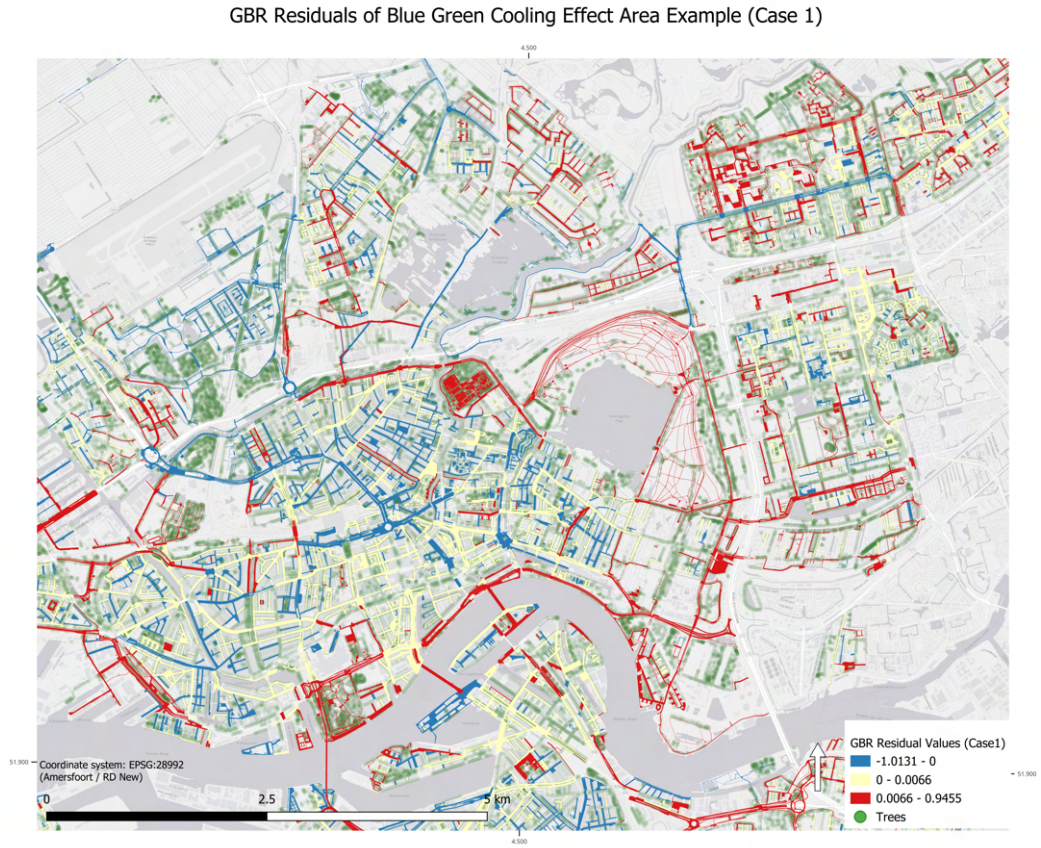

Figure 4.6.: GBR Residuals of Normalised Cooling Effect Prediction Area Example

From the results in section 4.3.1, GBR produced the best R^2 score. The results of these will be discussed below. Figure 4.3 shows the observed normalised cooling effect across our study area. Strong cooling zones can be seen in darker blue. These are concentrated around green corridors and dense vegetation patches, seen more around the central and western parts of Rotterdam. Lighter segments segments indicate where cooling effects may be weaker with more roads and built up environments.

The out of fold predictions by the Gradient Boosting Model in figure 4.4 appears to cover most of the spatial patterns seen in the observed data. It captures prominent cooling zones in central Rotterdam and other areas with vegetation. However, the model slightly underestimates cooling in some areas with dense vegetation or near water. This may be due to limited feature resolution or missing private tree data.

The residual map in figure 4.5 gives us an an understanding of which areas are being under predicted (red) and overpredicted (blue). The under predicted areas are mainly in dense vegetated zones, such as densely populated green areas and areas close to water bodies. This can be seen in figure 4.6. We can see that the areas where the trees are highly populated such as the middle to top right are red in colour, similar to the streets near the large bodies of

water right beside it. This suggest that the model under estimates the full extent of cooling in these areas by the water. This could be due to the under representation of trees as it may not cover trees not managed by the municipality, urban landscape inconsistencies and the overall microclimate of the area. The over predicted areas are more fragmented and seem concentrated in street segments with minimal vegetation. The model over predicts cooling in places that may have visible but low performing tree coverage.

In case 1, there is seen to be clear spatial correlation in the residuals. This may mean that model errors are not randomly distributed but come in the form of clusters. This implies that the model may have trouble with areas with distinct typologies.

The results for other models and city wide maps can be found in appendix C.

4.3.2. Predicting Flood Depth (Case 2)

Flood depth in this case is represented by the estimated maximum depth of water accumulation during extreme flood events in urban areas. This dataset is created via simulations of individual flooding events based from the National Water Model. The dataset highlights where floodwaters are likely to accumulate and display the max depth. Depth values are expressed in metres and this dataset is important in identifying urban zones that are highly exposed to flooding and may aid in supporting urban resilience planning against flood related risks [Rijkswaterstaat, 2023].

Model Performance

All 3 models were able to predict the flood depth values with limited success. The untuned results are presented in table 4.11. and the results with by the tuned hyperparameters are displayed in table 4.12. With the base parameters, Random Forest again showed the highest R^2 score at 0.332685, followed by Gradient Boosting with 0.28799 and Support Vector with 0.25919. despite the relatively small differences in R^2, the tree based models again had lower MAE and RMSE values.

However, in this case, Random Forest retained the best performance with an R^2 of 0.32554 as well as the lowest MAE and RMSE values with 0.12573 and 0.15879 respectively. Gradient Boosting followed very closely with an R^2 of 0.32106 and slightly higher error values. Although the difference is small, Random Forest appears to be the better suited to capture the spatial flood variation in case 2. The training R^2 for gradient boosting was 0.63483 while the test R^2 was 0.32106. This result was also similar to case 1, where although there may be a gap between the train and test values, it may be due to the spatial complexity of the flood depth as a target variable.

Support Vector again showed the weakest performance, both before and after tuning. The highest that it obtained was an R^2 score of 0.25918 after tuning, with no changes to the MAE and RMSE, meaning that the base parameters for this case were the most optimal.

In summary for case 2, all models struggled more with this target compared to case 1, which may reflect the more complex and less predictable nature of urban flood behaviours. Random forest showed the strongest predictive power, as well as the lowest errors.

The learning curves for the best models per case can be seen in D.

Table 4.11.: Results from Untuned Parameters (Case2)

Validation Metric	Random Forest	Gradient Boosting	Support Vector
R^2	0.32685	0.28799	0.25918
MAE	0.12480	0.13162	0.13562
RMSE	0.15864	0.16315	0.16642

Table 4.12.: Best Hyperparameters and Performance Across Models (Case 2)

Random Forest		Gradient Boosting		Support Vector	
Param	**Value**	**Param**	**Value**	**Param**	**Value**
n_estimators	150	n_estimators	600	C	1
max_depth	12	max_depth	5	epsilon	0.1
max_features	3	max_features	4	kernel	rbf
min_samples_leaf	4	min_samples_leaf	4	gamma	scale
min_samples_split	5	min_samples_split	20	-	
boostrap	False	learning_rate	0.03	-	
-		subsample	1	-	
R^2	0.32554	R^2	0.32106	R^2	0.25918
MAE	0.12573	MAE	0.12698	MAE	0.13562
RMSE	0.15879	RMSE	0.15932	RMSE	0.16642

Feature Importance

Table 4.13.: Feature Importance by Model (Case2)

Feature	Random Forest	Gradient Boosting	Support Vector
Urban Heat Island	0.21551	0.21561	0.17995
Imperviousness	0.09871	0.09140	0.15159
Sky View Factor	0.18333	0.18830	0.21712
Angular Betweeness	0.09180	0.08881	0.06664
Attraction Reach	0.25684	0.25750	0.20884
Perceived Temperature	0.10470	0.11858	0.14590
Tree species Count	0.04907	0.03977	0.02995

In case 2, the goal was to predict estimated flood depth from extreme events. Unlike case 1, the patterns shown here were that feature importance was less sharply skewed towards a single variable, with feature importance score appearing more even distributed. This may

reflect the complexity of flood simulations which are influenced by multiple environmental and morphological factors.

Random Forest showed that attraction reach (0.25684), UHI, (0.21551), and SVF (0.18333) were the most influential variables. The higher score of attraction reach might be affected by the vulnerability of built up or accessible zones, where water may accumulate such as flat urban areas. SVF's stronger contribution may indicate an indirect relationship with surface permeability or exposure to rainfall. Imperviousness with 0.09871 showed less contribution, surprisingly, given that there is a correlation between imperviousness and general surface runoff.

Gradient boosting presented a similar pattern, with Attraction Reach (0.25750), UHI (0.21561), and SVF (0.18830) still in the lead. This reinforces the idea that morphology, heat retention and open space may collectively contribute to water retention capacity. Notably, gradient boosting places more emphasis on PET (0.11858), potentially capturing overlap between hotter urban zones and poor draining performance.

In contrast, Support Vector distributed importance more evenly. SVF had the highest value (0.1583), followed by imperviousness (0.1074) and PET (0.1034). UHI (0.1275) still contributed was was less dominant this time.

Overall, this case showed higher reliance of a larger variety of features rather than single dominating features. This is consistent with the idea that urban flood depth is shaped by multidimensional urban features. Moreover, while the same features were used as in Case 1, their relative contributions shifted significantly.

Additional feature importance plots can be seen in B.3

Spatial Distribution

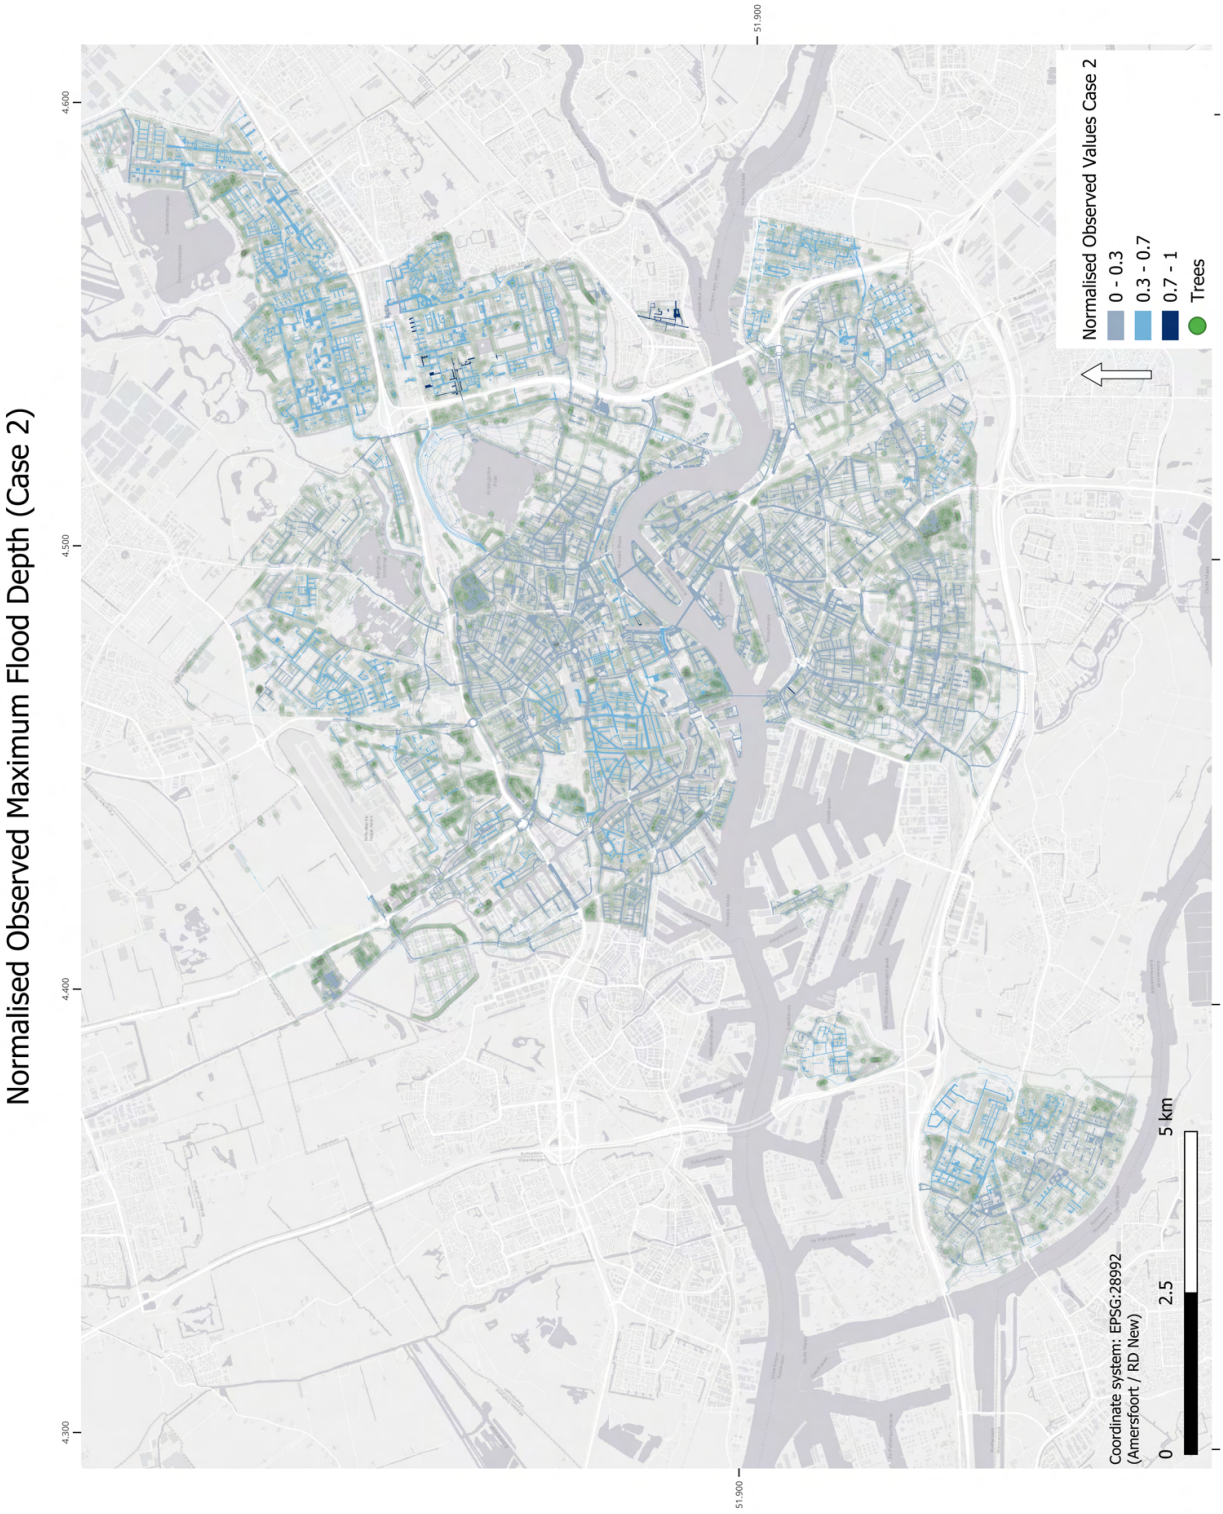

Figure 4.7.: Observed Normalised Max Flood Depth (Case 2)

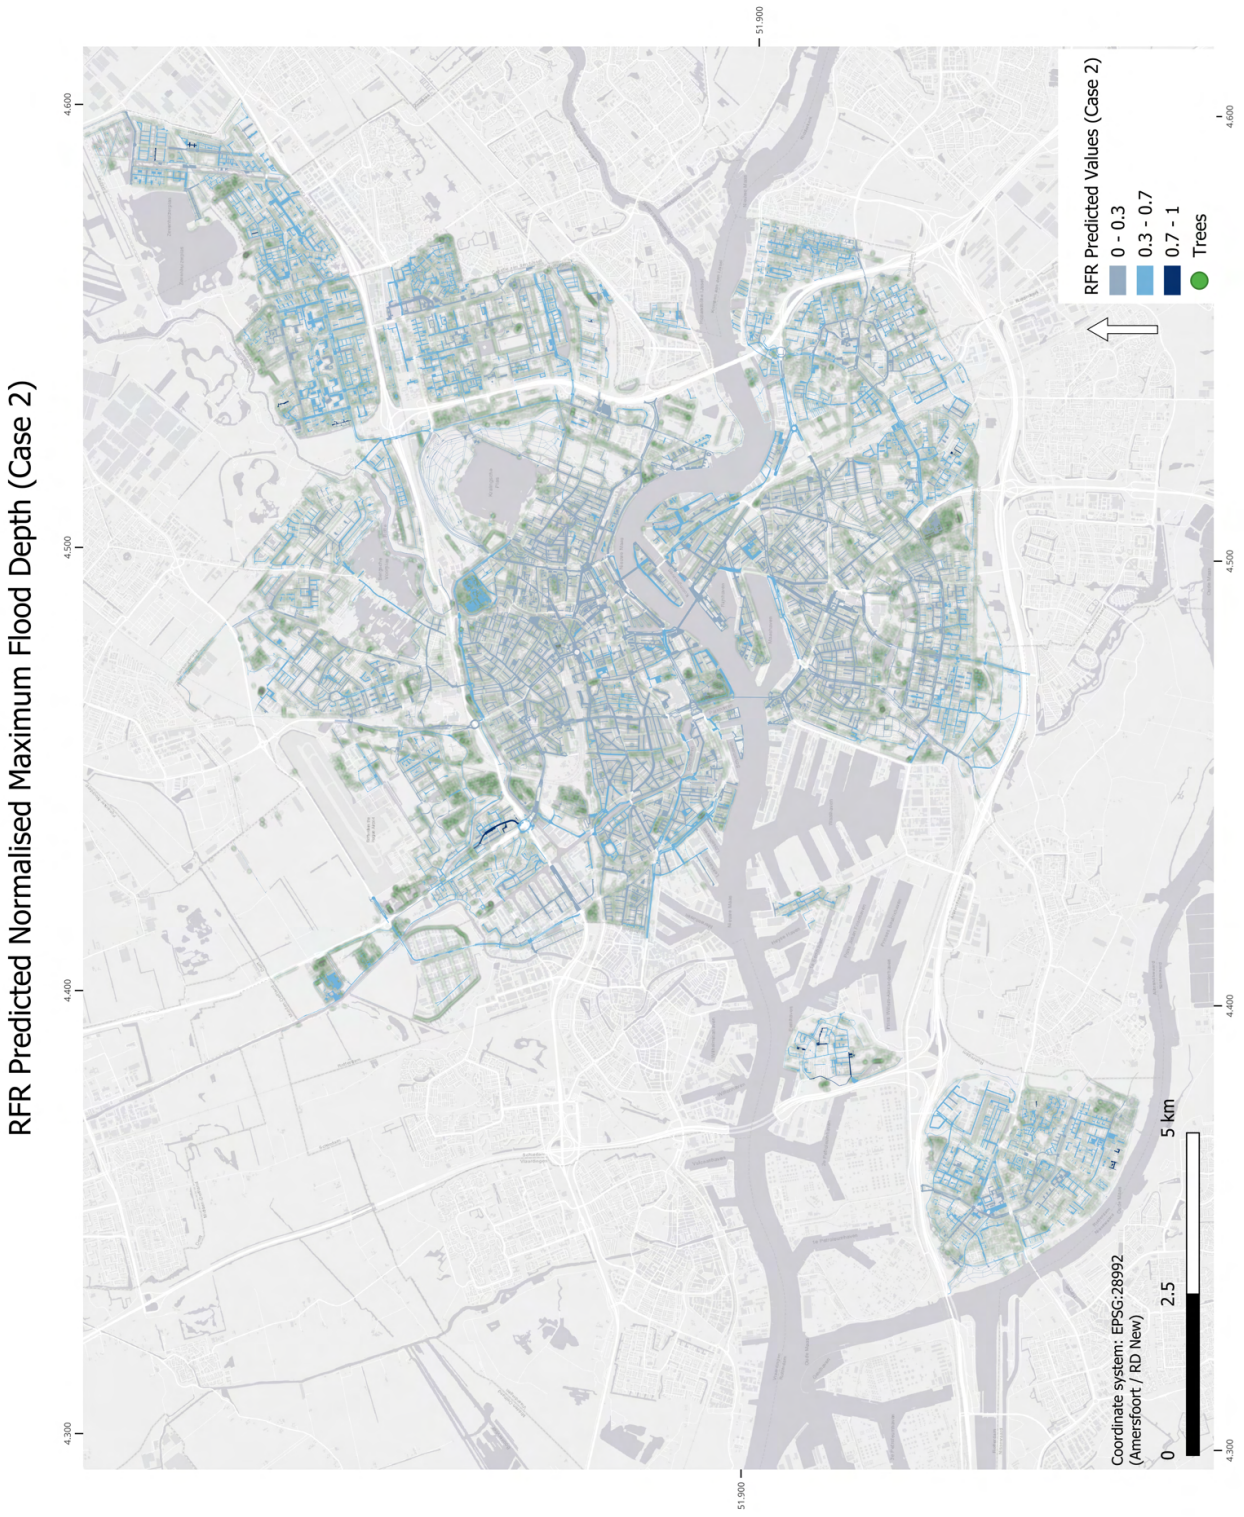

Figure 4.8.: RFR Predicted Normalised Max Flood Depth (Case 2)

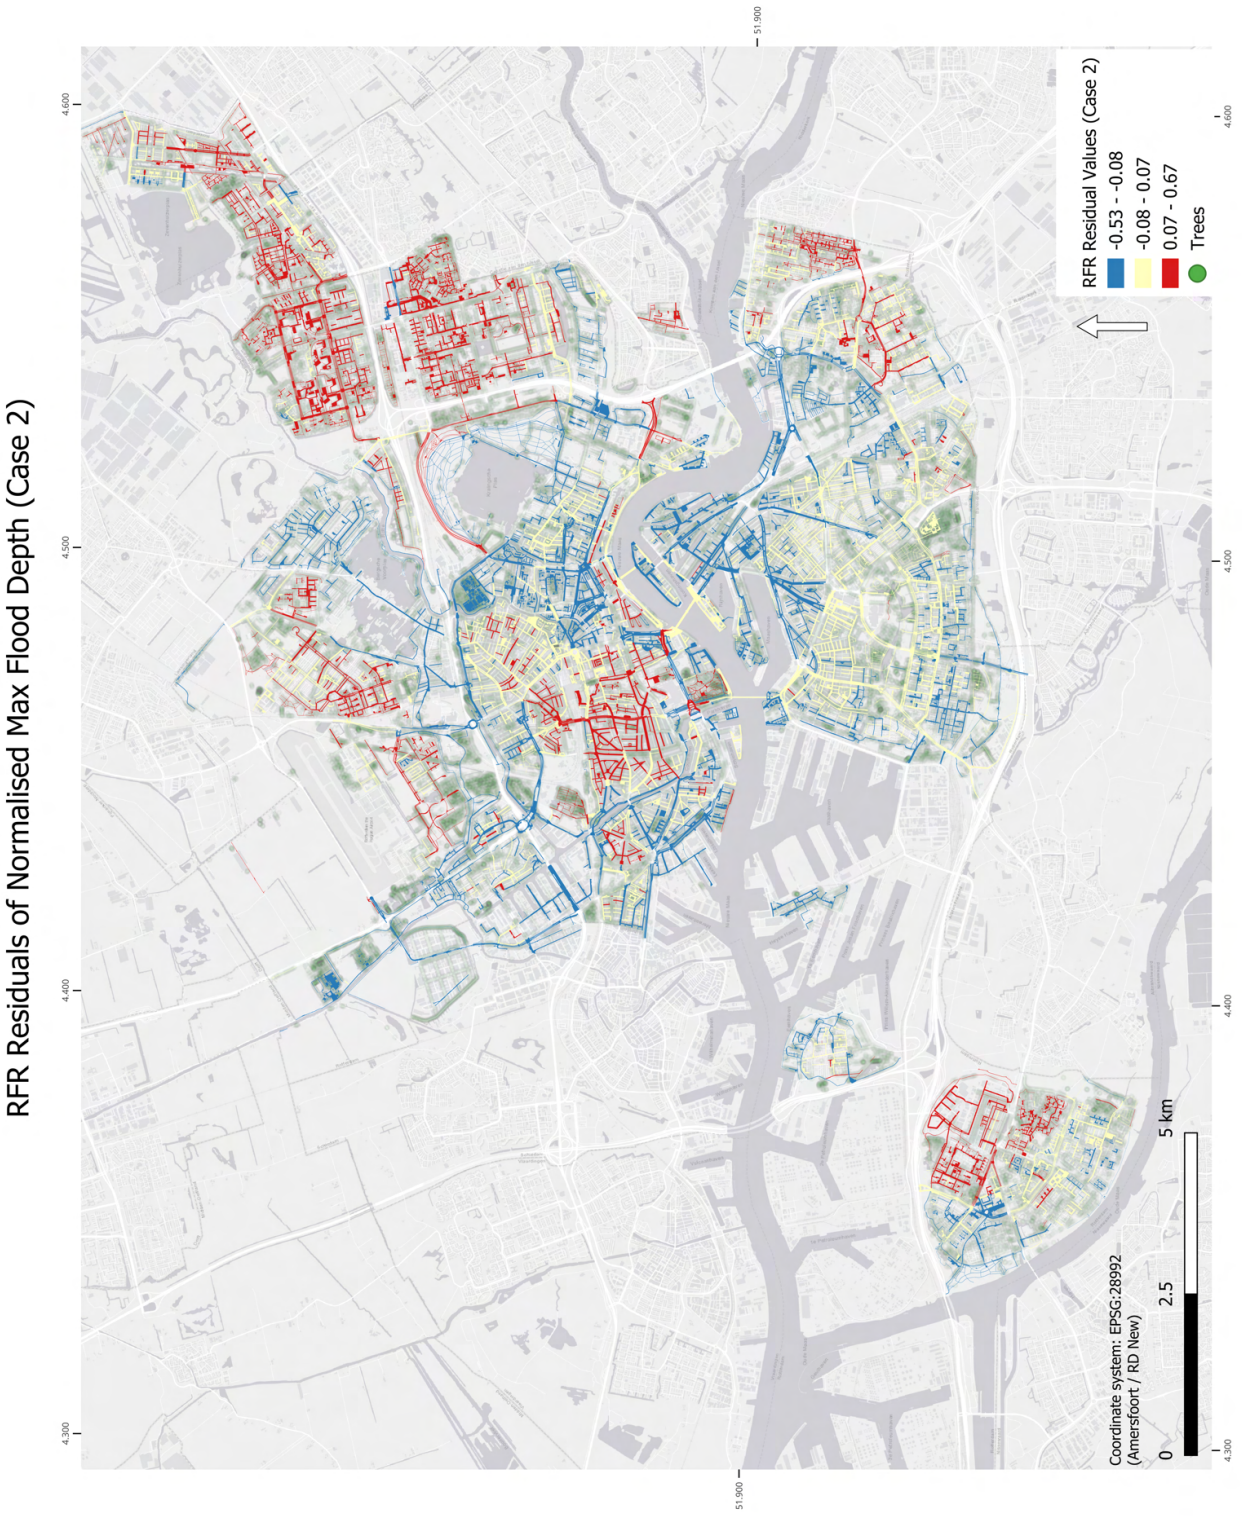

Figure 4.9.: RFR Residuals of Normalised Max Flood Depth (Case 2)

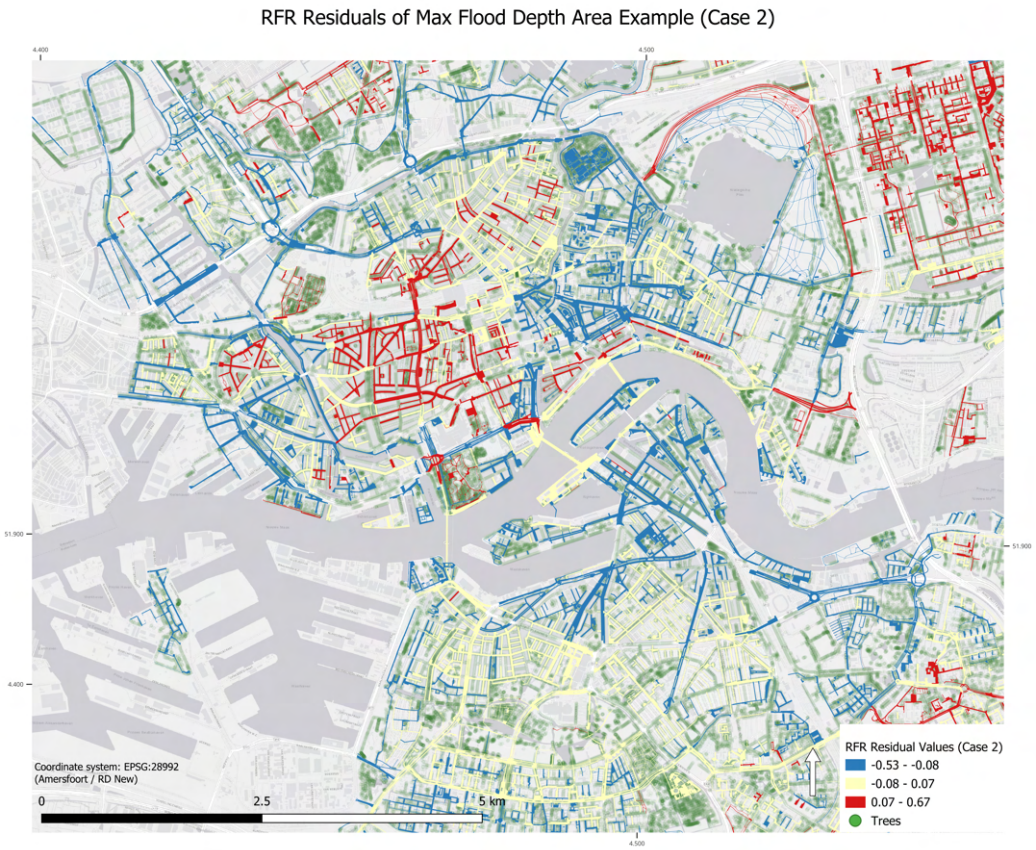

Figure 4.10.: RFR Residuals of Normalised Max Flood Depth Area Example

From the results in section 4.3.2, the Random Forest model produced the best R^2 score. The spatial outcomes of this model are visualised and analysed below. Figure 4.7 shows the normalised maximum flood depth across the study area. The dark blue segments reflect road sections highly exposed to flooding. These patterns seem to be concentrated around low lying areas near water bodies such as southern Rotterdam and the areas by the river.

The Random Forest predictions in figure 4.8 show a reasonably close match with the observed spatial patterns, suggesting that the model may have identified general flood prone areas. Higher risk zones such as the southern river adjacent areas as well as the western industrial areas and areas surrounding canals are captured. However, the prediction map is slightly less varied, suggesting there is some smoothing in the prediction process. This could be due to the ensemble nature of Random Forest, where the aggregated output could mute more extreme value, reducing overall contrast in spatial predictions.

The residual map in figure 4.9 shows underprediction in red (positive residuals), meaning the model estimated lower flood depths than were observed. As seen in figure 4.10 these areas are mostly concentrated inland, where actual flood depths may be driven by local topography or drainage limitations that weren't well captured in the model. Overprediction,

shown in blue (negative residuals), is more scattered and sometimes occurs closer to water bodies, where the model expected high flood exposure, but observed values were lower.

Similar to case 1, spatial correlation can be seen in the residuals to an extent. Clusters of model errors are not randomly distributed. These results may be due to feature presentation as well as quality of training data.

The results for other models and city wide maps can be found in appendix C.

4.3.3. Predicted NO_2 Concentration (Case 3)

Nitrogen Dixoide (NO_2) concentration is used as an indicator of air quality and environmental health risk in urban areas. The dataset represents the average NO_2 levels for 2022 and was produced by the National Institute for Public Health and the Environment (RIVM) [Rijksinstituut voor Volksgezondheid en Milieu (RIVM), 2025]. NO_2 is a pollutant often associated with traffic emissions and industrial activity. Higher concentrations are observed in heavy impervious road areas and low green cover. Using NO_2 patterns can be relevant for urban planning as long term exposure to poor air quality can impact overall public health.

Model Performance

Similar to Case 1, the models in case 3 were able to predict NO_2 concentration with moderate success. As shown in table 4.14, Random forest achieved the highest R^2 score of 0.46392 using untuned parameters, followed by gradient boosting, with 0.43457, and Support Vector with 0.41801. Similar again to the previous cases, tree based models produced much lower MAE and RMSE values.

After hyperparamter tuning, it is again similar to Case 1, where the Gradient boosting turned out to have the highest R^2 score overall at 0.46845, outperforming Random Forest which actually had its R^2 score drop by a small amount. By a small margin, the results show that the hyperparameters tuned aided in adapting to the underlying structure of the data. This case showed also showed that the MAE and RMSE values were tightly grouped across all 3 models, suggesting that in this case the model differences seem to come from the ability to generalise rather than overall accuracy prediction.

Gradient boosting recorded the highest training R^2 score at 0.66784, compared to random Forest at 0.66784, and Support Vector at 0.4504. The moderate gap between training and test in R^2 in the tree based models suggest there may be small amounts of overfitting in case 3 similar to case 1. However, these gaps are likely due to the complexity of urban NO_2 concentration patterns. Support Vector Regression again showed most consistent generalisation with its training and test R^2 scores being nearly identical (0.4504 and 0.4276), which reflects a stable model but is less accurate in this case.

Overall in case 3, Gradient Boosting achieved the best balance between accuracy and generalisation. Random Forest was a competitive second but leaned more towards overfitting from the training and validation test scores. Support Vector, similar to case 2 offered the most stable results but at a cost of predictive power.

The learning curves for the best models per case can be seen in D.

Table 4.14.: Results from Untuned Parameters (Case3)

Validation Metric	Random Forest	Gradient Boosting	Support Vector
R^2	0.46392	0.43457	0.41801
MAE	0.07503	0.07866	0.07939
RMSE	0.10082	0.10317	0.10468

Table 4.15.: Best Hyperparameters and Performance Across Models (Case 3)

Random Forest		Gradient Boosting		Support Vector	
Param	**Value**	**Param**	**Value**	**Param**	**Value**
n_estimators	200	n_estimators	400	C	1
max_depth	12	max_depth	5	epsilon	0.01
max_features	4	max_features	4	kernel	rbf
min_samples_leaf	6	min_samples_leaf	4	gamma	scale
min_samples_split	15	min_samples_split	10	-	
boostrap	True	learning_rate	0.03	-	
-		subsample	0.7	-	
R^2	0.46011	R^2	0.46845	R^2	0.42757
MAE	0.07618	MAE	0.07571	MAE	0.07742
RMSE	0.10082	RMSE	0.10004	RMSE	0.10382

Feature Importance

Table 4.16.: Feature Importance by Model (Case3)

Feature	Random Forest	Gradient Boosting	Support Vector
Urban Heat Island	0.41596	0.37409	0.55818
Imperviousness	0.07900	0.08712	0.11349
Sky View Factor	0.15151	0.15271	0.11337
Angular Betweeness	0.15750	0.15748	0.09257
Attraction Reach	0.09981	0.11084	0.05733
Perceived Temperature	0.06138	0.07782	0.04895
Tree species Count	0.03479	0.03990	0.01609

In case 3, the aim was to predict spatial variation in nitrogen dioxide concentration (NO_2). Similar to case 1, feature importance results showed that it had a stronger dominance of a

single variable. UHI was the most important feature across all three model types, suggesting stronger relationships between heated surfaces and pollution dispersion in urban areas.

RFR assigned the highest importance to UHI (0.41596), followed by angular betweeness (0.15750), and SVF (0.15151). This distribution indicates that both heat intensity and the morphology of the street network may play a role in trapping pollutants. The influence of angular betweenness suggests that street connectivity may affect how air pollutants circulate or become trapped in denser parts of the city. The lower importance of imperviousness (0.07900) and attraction reach (0.09981) suggests that urban form features related to surface runoff may be less directly linked to NO_2 levels.

GBR showed a similar pattern, again ranking UHI highest (0.37409), followed by angular betweenness (0.15748) and SVF (0.15271). The consistency across Random Forest and Gradient Boosting adds confidence to the relevance of these features. Interestingly, perceived temperature also gained more influence in gradient boosting (0.07782), potentially pointing to an overlap between thermal discomfort and areas with higher NO_2 exposure.

SVR showed an even stronger emphasis on UHI (0.55325), by far the most dominant feature in this model. This may reflect the global nature of SVR, where broad-scale temperature patterns significantly shape pollutant distribution. Other features like Imperviousness (0.11249), SVF (0.11137), and PET (0.04852) contributed to a lesser extent, but again highlight the role of imperviousness and exposure in shaping air quality.

Overall, this case displayed a clearer dominance of thermally related variables, particularly UHI. Unlike flood depth or cooling effect, NO_2 levels appear to be more strongly associated with heat retaining and morphologically enclosed areas. The results suggest that vegetation, while still included in the model, played a much more minor role in this case, hinting that maybe GI in this case alone may not be sufficient in improving urban resilience of areas with higher NO_2 values.

Additional feature importance plots can be seen in B.3

Spatial Distribution

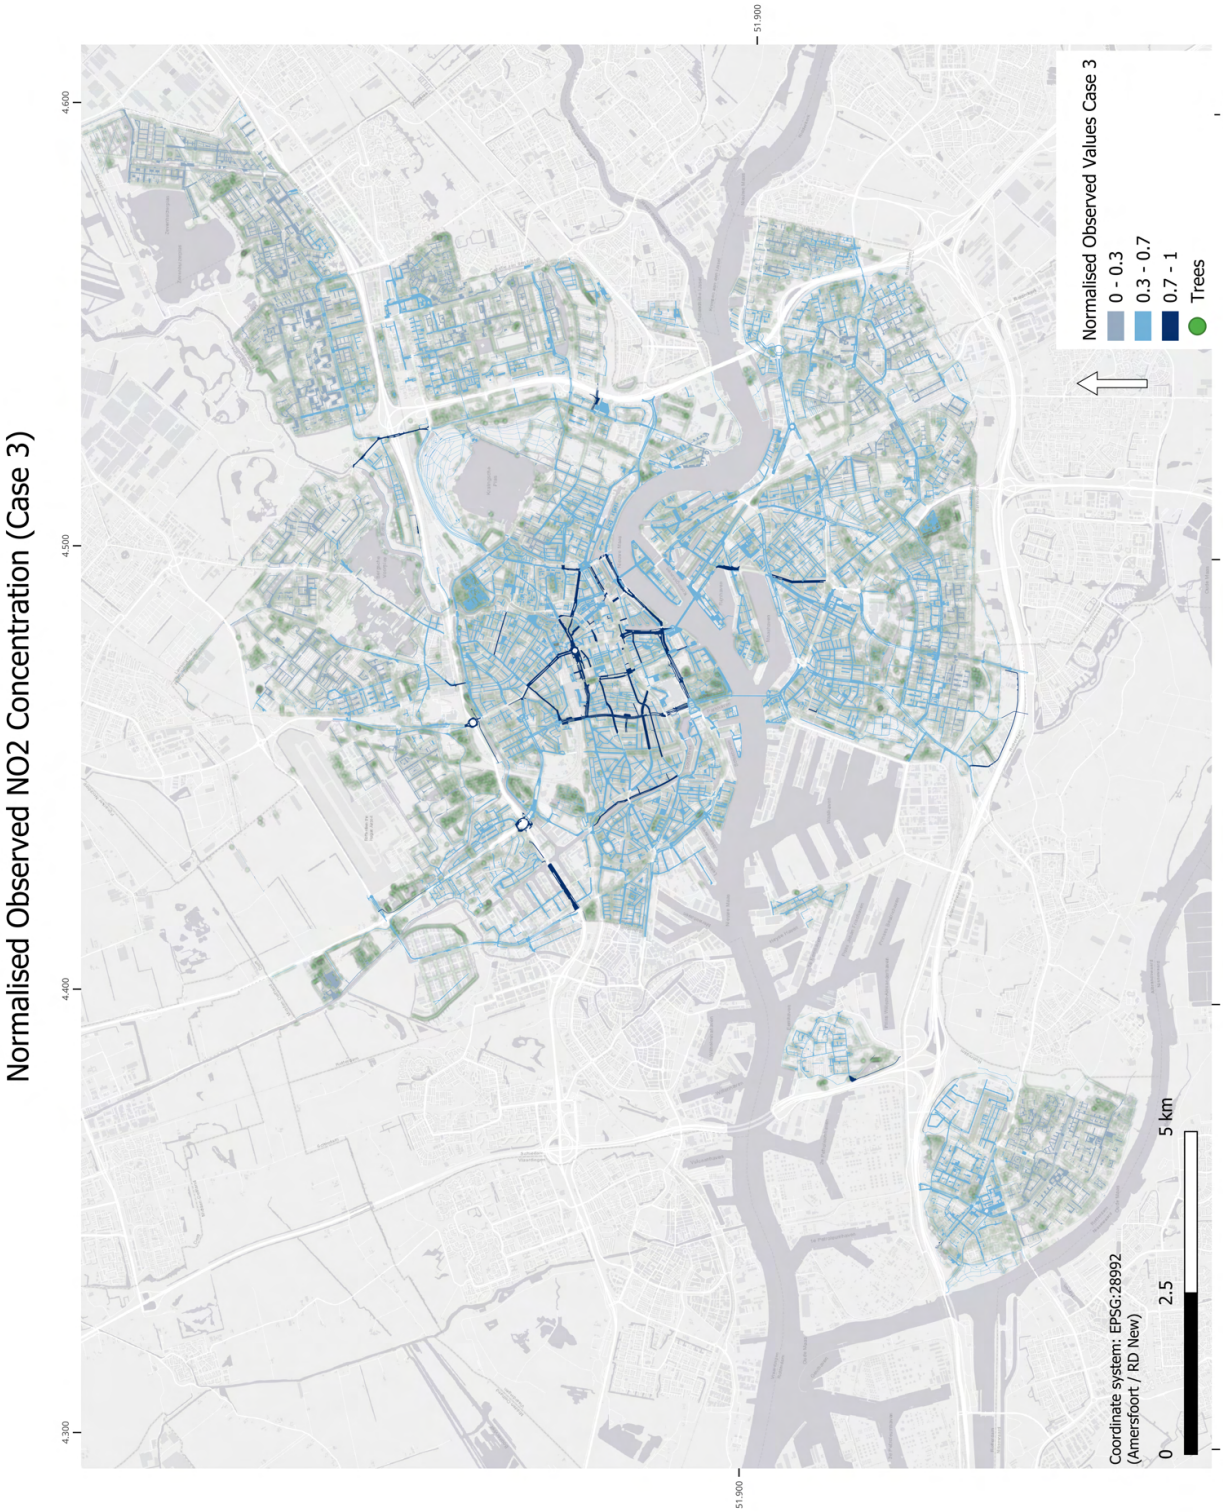

Figure 4.11.: Observed Normalised NO2 Concentration (Case 3)

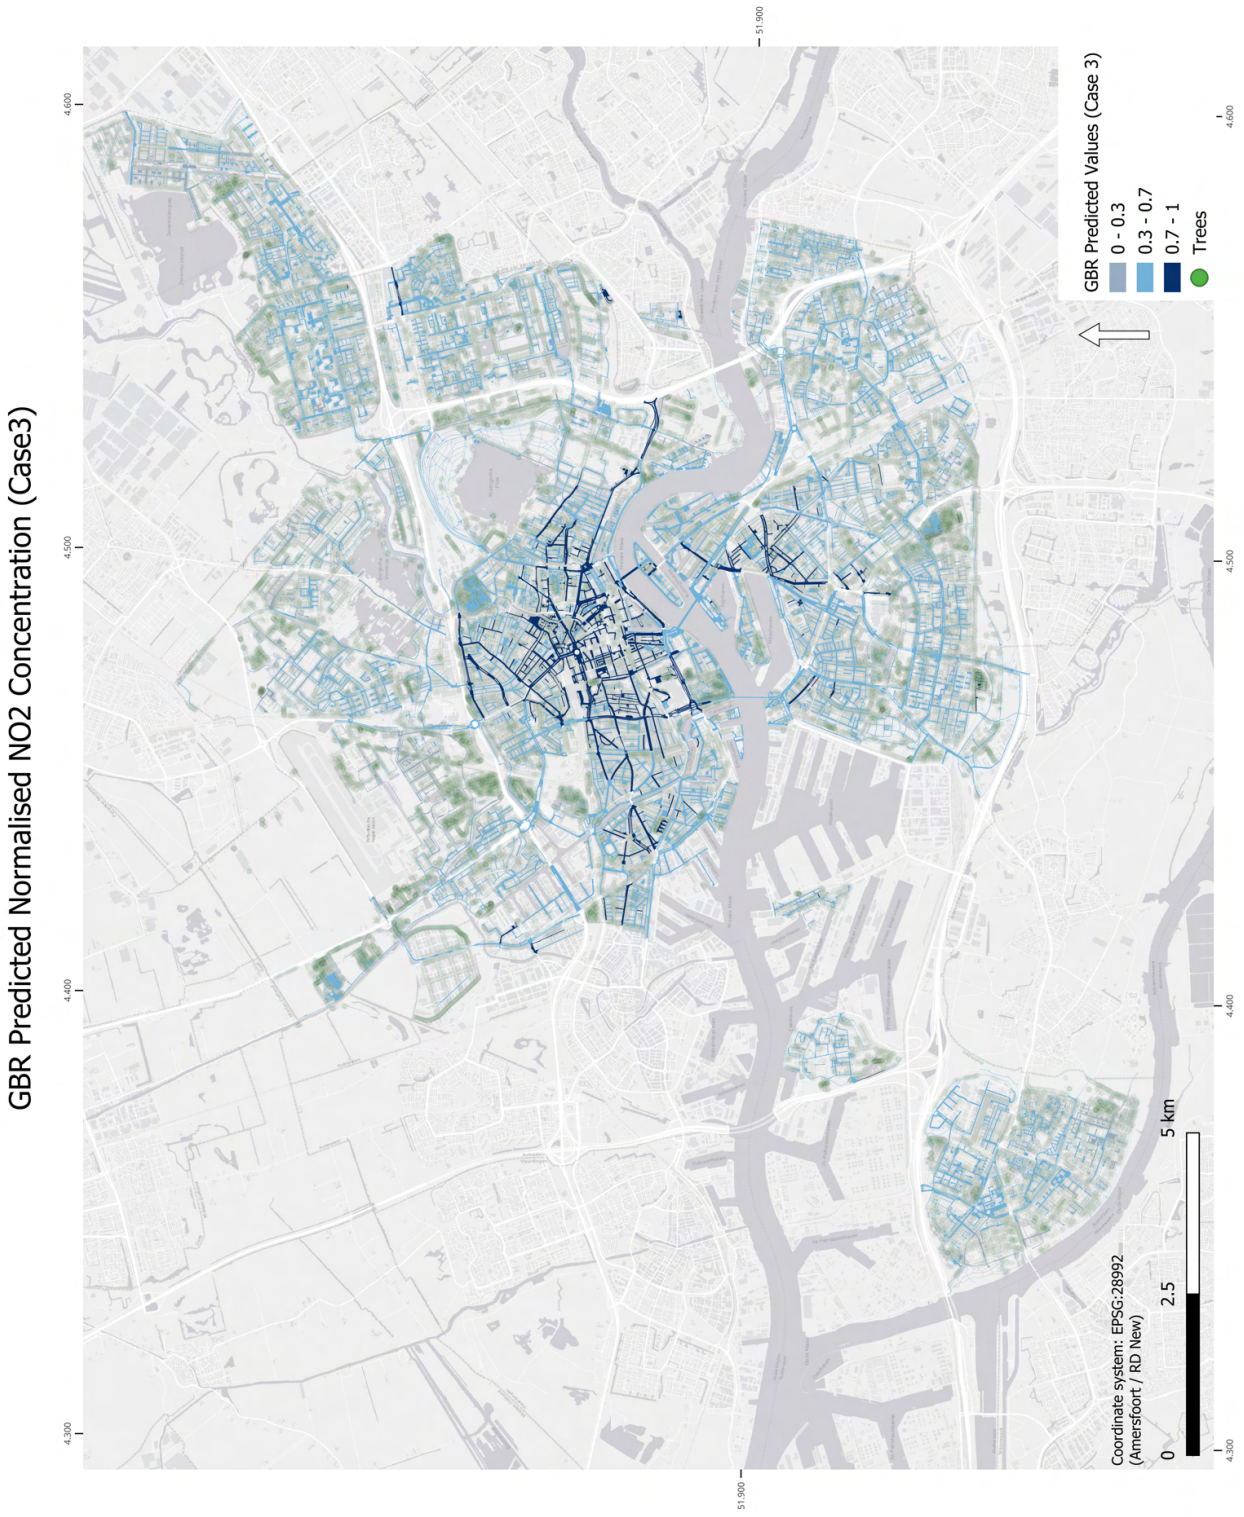

Figure 4.12.: GBR Predicted NO_2 Normalised Concentration (Case 3)

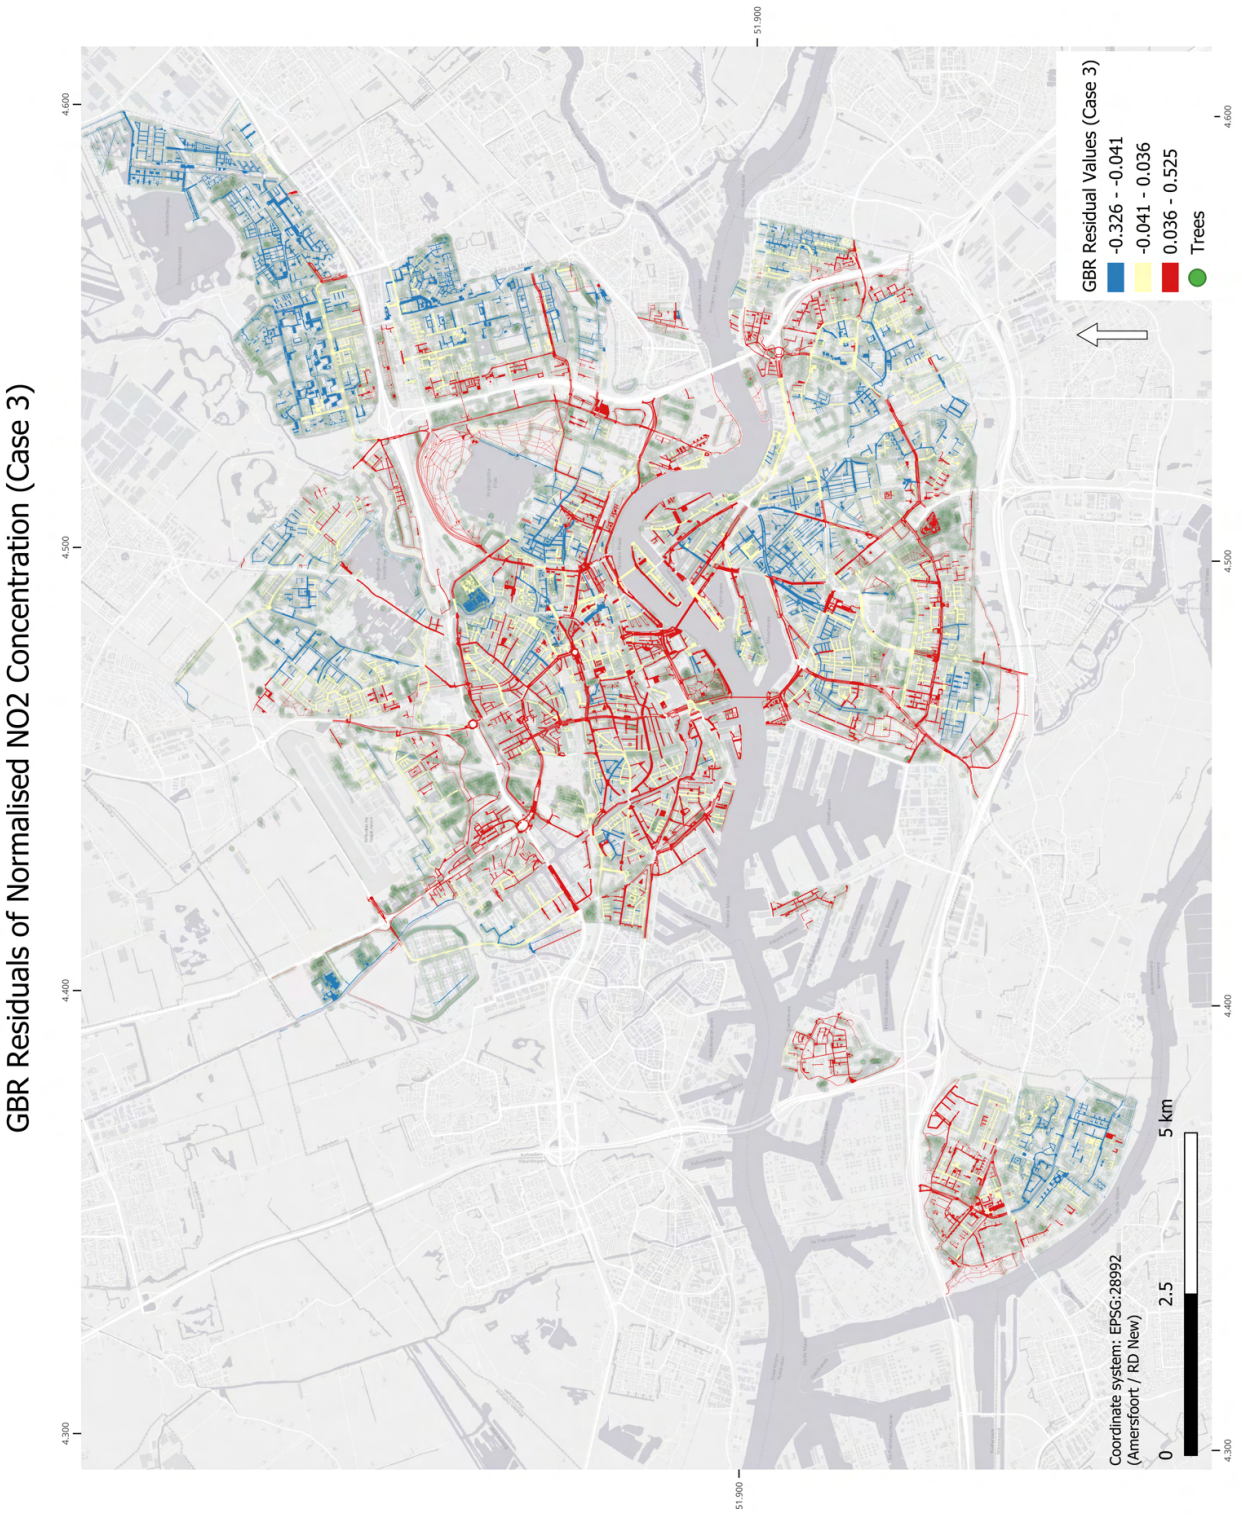

Figure 4.13.: GBR Residuals of Normalised NO_2 Concentration (Case 3)

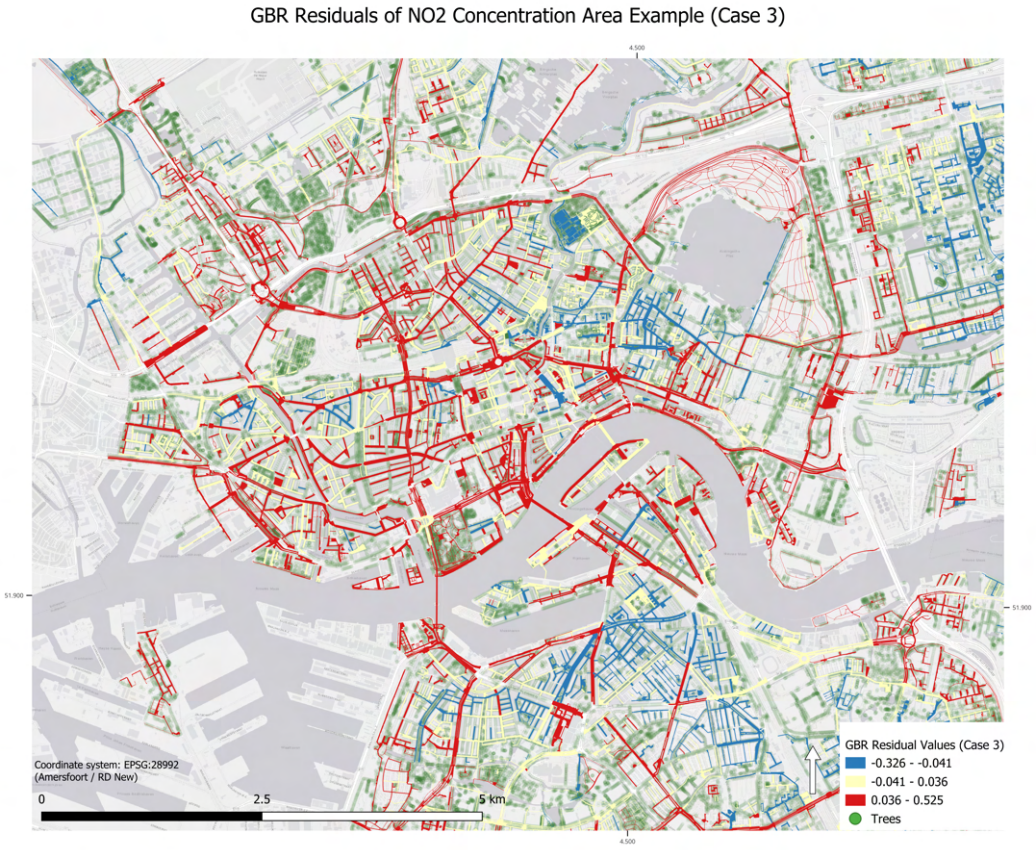

Figure 4.14.: GBR Residuals of NO_2 Area Example

From the results in section 4.3.3 Gradient boosting was found to perform best with the highest R^2 score among the models. The following figures below display the observed, predicted and residual spatial patterns of NO_2 concentrations across the study area.

Figure 4.11 shows the observed normalised NO_2 concentrations. Higher concentrations are shown in darker blue and are clearly concentrated in the city centre, with noticeable hotspots following larger roadways. These patterns align with traffic related air pollution sources. Suburban zones and residential green areas associated with lower NO_2 levels are shown in lighter tones.

Figure 4.12 presents the predictions made by the Gradient Boosting Model. the model manages to capture broad spatial trends of higher NO_2 in central and eastern Rotterdam and lower values outside of this. The general pattern of high pollution zones is captured, mostly around the urban centre as well as key transport areas.

The residual map in figure 4.13 provides an overview of the prediction errors. In figure 4.14 under are shown in red and are in this case more dispersed and appear along the southern and eastern areas where concentrations were generally lower in the observed data. Over-predictions, shown in blue, are found mostly in the densest traffic corridors and compact

urban zones in the north and centre regions. This might indicate that sharp NO_2 peaks were missed or underestimated.

Unlike case 1, residuals in this case appear less clustered and more scattered. This could indicate fewer issues with spatial autocorrelation of errors. Still, some biases may remain where localised extremes do not seem to do well in the model. This may be due to absence of high resolution traffic data or pollution related feature inputs.

In summary, the model captures the broad spatial trends in NO_2 distribution but its performance is lower at finer scales where local sources and environmental features may have stronger impacts.

The results for other models and city wide maps can be found in appendix C.

4.3.4. Comparison with Baseline (No Feature Extraction)

To assess the impact of feature extraction on model performance, baseline models were trained using the raw features prior to any refinement. These results were then compared with the performance of models trained on the extracted and refined feature set described in Section 3.2.1. The comparison aimed to evaluate the extent to which spatially informed indicators improved predictive accuracy and model interpretability. As part of the feature selection process, two features were removed: NDVI and total tree count per street. Although the objective was to reduce redundancy and enhance spatial relevance, these excluded features may have contained complementary or interacting information that contributed to predictive performance in the baseline models. NDVI was highly correlated with imperviousness and exhibited seasonal instability, often returning near-zero values in winter when trees lose their leaves. Similarly, total tree count showed high correlation with tree species richness and, given that the features were normalised, was considered redundant in favour of the more conceptually meaningful species count.

The performance results showed that feature extraction was important in improving performance across most models. In Case 1, for Gradient Boosting, the R^2 score improved from 0.473 (without extraction) to 0.479 (with extraction), and RMSE decreased slightly from 0.1321 to 0.1314. For Random Forest, however, performance slightly declined, with the R^2 decreasing from 0.473 to 0.449 and RMSE increasing from 0.1321 to 0.1352. Support Vector Regression showed a marginal improvement in R^2, increasing from 0.391 to 0.398, and a small decrease in RMSE from 0.1421 to 0.1414.

In Case 2, all models either showed stable or slightly reduced performance when trained on the extracted features. Gradient Boosting's R^2 decreased slightly from 0.329 to 0.321, and Random Forest also dropped from 0.336 to 0.326. SVR remained relatively consistent, with a very small increase from 0.256 to 0.260.

In Case 3, Gradient Boosting again benefitted from feature extraction, improving from an R^2 of 0.468 to 0.479. Random Forest saw a slight increase from 0.467 to 0.468, while SVR improved modestly from 0.426 to 0.427.

These results suggest that while spatial feature extraction enhances domain relevance and interpretability, its impact on predictive performance varies depending on the model and the specific case. Gradient Boosting consistently showed stable or improved results, indicating its robustness to different feature representations. In contrast, Random Forest and SVR

exhibited more variability, reinforcing the idea that the effectiveness of geospatial feature engineering is both model- and context-dependent.

Despite the removal of these features, it is possible that these features could have added value in ways not fully captured by the refined set. NDVI, for example, may have provided additional information on vegetation cover, particularly in areas where private or unregistered trees were not included in the municipal tree dataset. This broader dataset could have improved the model's ability to detect green infrastructure presence across different land types. Similarly, total tree count may have contributed to capturing structural characteristics of green space density, which are not entirely represented by species richness alone. These additional dimensions of greenery could have supported more accurate predictions in certain cases.

5. Discussion and Conclusion

In this section, a reflection on the research findings and their implications will be provided. The main research question along with it's associated sub questions will be revisited to see how well the results have satisfied the objectives created. The discussion is structured into sections where a research overview will answer each key research question, followed by an interpretation of model outcomes across cases, highlighting spatial patterns, model performance and important feature contributions. This is followed by the contributions that this study has provided to the field of urban resilience and green infrastructure assessment. Additionally, limitations of the methodology and recommendations are highlighted for future work before the final conclusion, summarising the study's key takeaways.

5.1. Research Overview

The central research question posed by this thesis was:

> "How can predictive modelling techniques be used to evaluate the effectiveness of green infrastructure in improving urban resilience at street level across different dimensions of the built environment?"

To address this, the study was guided by five sub-questions, each of which explored a different aspect of the research problem. The following section discusses each sub-question individually, drawing on the results from the three case studies, cooling effect, flood depth, and NO_2 concentration to present key findings and insights.

- What aspects of green infrastructure can be assessed in a multifunctional way to evaluate their contributions to urban resilience across multiple dimensions?

This study attempted to have a multifunctional approach to evaluate green infrastructure by using features that could influence multiple aspects of urban resilience at the same time. Compared to focusing on individual green infrastructure types such as singular parks or trees, the analysis captured a range of green infrastructure related indicators that help reflect urban resilience on a street level scale.

Tree species richness and tree population were first used to capture both vegetation density and ecological diversity. However, as these two features were highly correlated, only species count was kept. This was because from the correlations in this data, higher amounts species also proved to have higher amounts of trees. As the features were normalised, tree population was removed in the final feature set to remove redundancy. The SVF, may not be strictly a green infrastructure feature, but indirectly represents street level enclosure and potential shading effects from vegetation. Similarly, perceived temperature and UHI effect allows for analysis of thermal regulation or capacity of GI. Features such as attraction reach and angular betweenness were also included to capture how accessible and integrated green infrastructure such as parks and urban forests is within the city.

By incorporating these features, GI is not only assessed with vegetation cover, but also in it's ability to reduce heat, manage water and improve accessibility which may be seen as key components urban resilience.

- How can environmental, biodiversity, and morphological features be combined to comprehensively assess and enhance urban resilience?

Environmental features such as UHI, PET, and imperviousness were included to assess how the built environment responds to heat contributes to the local climate. These factors are seen to be important in exposure to thermal discomfort and heat related risks. The SVF, while more morphological in context, aids in bridging how morphological features can also influence the built environment.

Biodiversity was captured through the inclusion of tree species richness. Tree species richness is important as is captures the overall ecology of the study area. Not only does it include species types, but also the tree population in our case seemed almost directly correlated to tree species. This suggests that areas with a higher number of trees also tend to support greater species diversity, reinforcing the value of this metric as a proxy for urban biodiversity. By incorporating tree species richness, the models account for ecological diversity and the potential benefits of a diverse vegetated areas, such as enhanced habitat quality, resilience to pests and diseases, and improved ecosystem stability.

Morphological features such as angular betweenness and attraction reach were used to capture the spatial and functional accessibility and proximity of the street network and access to green infrastructure. These elements help assess how well green infrastructure is embedded within the study area, and if contributes meaningfully to connectivity and accessibility at the local level.

By combining these three categories, the study was able to study urban resilience in a way that goes beyond individual indicators. It allows for the identification of areas that are not only exposed to risks but may also lack ecological diversity or connectivity, providing a more targeted basis for green infrastructure planning and intervention.

- How can predictive modelling techniques, such as machine learning, be used to evaluate and optimise green infrastructure for improving urban resilience?

The use of predictive modelling in this study was to aid in understanding how features related to green infrastructure contribute more to urban resilience. Rather than focusing on description correlations alone, machine learning was used to estimate the predictive power of resilience based indicators across different cases.

Three regression models were used to evaluate performance across three resilience related target variables. The targets used were cooling effect, flood depth, and NO_2 concentration. These represent environmental comfort, flood mitigation, as well as air quality respectively. These were chosen as they are generally shaped by green infrastructure related features and are important features regarding urban resilience.

Using these techniques allowed for the analysis to show where green infrastructure is effective under different dimensions of urban resilience. Furthermore, feature importance rankings provided insight into which aspects of the urban environment mostly influenced each target variable. This adds value by allowing planners to focus not only on where greening is lacking, but what kind of interventions may have the most impact.

- How can the presence of spatial autocorrelation in the residuals affect the reliability of the model's predictions?

In Case 1 (Cooling Effects), areas with dense vegetation, green corridors and proximity to water bodies were predicted to show strong cooling benefits, these patterns were generally captured well by the Gradient Boosting Model, although residual maps showed small tendencies to underpredict in zones with higher green quality and overpredict in areas with sparse vegetation. This suggests that while the model captures broader spatial trends, it struggles with capturing locally complex areas. The reasons for this could be because of data completion. The tree dataset obtained from the municipality of Rotterdam did not include private trees. This meant that the model could have missed areas that may have contained more densely vegetated areas compared to what was fed into the model.

Case 2 (Max Flood Depth), predictions were less spatially concentrated around single variables, as seen from the feature importance rankings. Multiple factors such as angular betweenness, SVF, and urban heat patterns contributed more to the complex spatial patterns produced. Prediction accuracy was moderate which reflects the difficulty of urban flood modelling from surface features alone. To improve the accuracy of this case, additional features such as including a digital elevation model for direct elevation data and using Attraction Reach again, including water bodies such as canals and rivers.

Case 3 (NO_2 concentration) presented another unique challenge. While accessibility related features and green coverage played a role, observed air pollution may have been highly influenced by features not included in our feature set such as wind direction or traffic flow. Spatial predictions here were more disjointed, with clusters of under predictions more in dense streets and over predicted in greener zones.

Across all cases, spatial patterns highlighted the potential of GI in enhancing urban resilience, but also showed the limitations of using surface level features only. The models were good at picking up general spatial tendencies but struggled where there were specific local conditions not picked up by the feature set provided. The distribution of spatial autocorrelation in the residuals suggest errors were not random but in areas where the feature set may have been lacking or incomplete. This implies that model errors may be location dependent and affects the reliability of the model as the errors were more influenced by local context rather than noise in the data. As a result, spatial auto correlation in the residuals suggests the model may under or over predict in certain areas. This reduces the reliability of predictions where the local microclimate or landscape is absract.

- How can the proposed framework be applied to assess and prioritise green infrastructure interventions in Rotterdam

The framework developed in this study aims to offer a practical and data driven approach to evaluate GI from multiple resilience perspectives. By combining the multidimensional indicators to predict resilience outcomes, the framework enables planners to visualise areas that may benefit from targeted green interventions.

In Case 1 (cooling effect), the results show that areas with moderate to low cooling potential could be prioritised for new green infrastructure interventions. The predictive map in combination with feature importance provides a valuable tool to identify gaps such as these spatially.

In Case 2 (flood depth), the framework can help to highlight streets or urban areas where flood risk is prominent. While the model was limited by lack of hydrological data, the

ability to produce reasonable spatial predictions from surface features alone suggests this approach could support broader screening exercises to prioritise nature based stormwater interventions like permeable surfaces or rain gardens.

For Case 3 (NO_2 concentration), the framework points to urban areas where green coverage and accessibility are not sufficiently mitigating air quality issues. While predictions were more fragmented, the insights still help in flagging areas where vegetation buffers or green facades might be strengthened to support public health and air quality.

Overall, the spatial and model based insights can be translated into green infrastructure planning decisions. The street scale nature of the analysis means that the model outputs may potentially aid in municipal level urban greening strategies.

5.2. Interpretation of results

Across all 3 cases, the results show that GI related features can contribute to predicting different dimensions of urban resilience. However, the performance of each model varied depending on the type of resilience indicator being predicted due to our feature set. This shows how well individual features capture relevant processes, as well as how different model architectures handle each targets spatial differences.

Case 1 in section 4.3.1, which aimed to predict the cooling effects of green and blue infrastructure, showed the highest overall accuracy with an R^2 score of 0.47945. Spatial patterns seemed to be captured well, particularly in zones with dense vegetation and green corridors. Here, tree species richness and urban heat indicators were highly influential, and the most effective model (Gradient Boosting) was able to generalise well across the study area, despite some underpredictions in the higher performing cooling zones.

Case 2 in section 4.3.2, the target variable (Maximum Flood Depth) was built by a more complex and multidimensional set of features. Compared to unbalanced feature importance in Case 1, there was not a single dominant feature, several environmental and morphological variables contributed to the model performance. Case 2 ended up with a much lower accuracy score. This was likely due to the difficulty in capturing hydrological processes using surface level features at street level.

Case 3 in section 4.3.3 focused on NO_2 concentrations. This was the most difficult to predict, with moderate R^2 scores in both training and cross fold testing. While features such as imperviousness and accessibility contributed in predicting the observed data, it suggested that additional features would have helped with predictions. As a result, the spatial models were less stable and showed scattered error clusters.

Overall, the tree based ensemble models used (RFR and GBR) tended to perform better in cases of accuracy but may show slight signs of overfitting in some cases. SVR, while less accurate may have produced more stable generalisation, particularly when dealing with noisy or sparse data. Feature importance rankings also shifted on a case by case basis, showing no single factor consistently dominated across all resilience types. This shows the influence that case specific features may need.

5.3. Contributions of This Study

This study contributes to the growing body of research on green infrastructure and urban resilience by presenting a Geomatics driven framework that links spatial data that can be available at street level with machine learning based prediction of resilience indicators. Instead focusing on a single dimension such as vegetation cover or urban heat, this approach combines environmental, morphological and biodiversity related attributes with an aim to produce a coherent view of urban resilience at a street level scale.

One of the main contributions is the use of predictive modelling within a geo analytical context by incorporating techniques to explore how green infrastructure characteristics contribute to various resilience outcomes. By applying three regression models, the study assessed performance of different models and also highlighted the influence of different features through model explanation tools like feature importance, permutation importance and residual mapping.

Another key contribution is the street scale analytical resolution. Compared to using aggregated neighbourhood or city level indicators, this study utilises a street level resolution. This fine scale approach aims to make the findings more relevant for direct planning and intervention. This is useful in cases where decisions are needed to be made around specific roads, corridors or blocks.

Additionally, the study demonstrated the value of combining statistical validity, such as correlation matrices in addition with conceptual validation. The final feature set was selected based not only on independence but also on relevance to cooling, access, stormwater management, and ecological function.

Finally, the workflow developed in this study aims to make use of key Geomatics principles, including data integration, spatial analysis, and geo computational modelling. With the use case of Rotterdam, the approach aims to be scalable and transferable to other urban areas using similar datasets. The structured workflow, encompassing data preprocessing, spatial feature extraction, model training, and cartographic visualisation demonstrates how Geomatics methods can transform raw spatial data into meaningful insights for urban resilience assessment, aiding in the support of informed decision making in climate adaptation and green infrastructure planning

5.4. Limitations

Several limitations are acknowledged in this study. The first of which was the availability and quality of street level spatial data. Many urban resilience indicators were either unavailable at street resolution or rely on interpolated values from larger datasets. This made it difficult to capture the environmental, biodiversity or morphological processes that influence urban resilience outcomes. For instance, factors such as wind direction or traffic congestion could not be fully included in the analysis, which may have overall weakened the predictive capacity in certain cases. Similarly, public datasets were not always complete. Firstly, tree datasets available in Rotterdam do not include private trees which may have increased the total tree count or increase species richness.

Furthermore, the use of static datasets limits the temporal aspect of the analysis. Environmental resilience factors such as flood depth or urban heat may vary with the seasons as

well as weather conditions. However, this study used average values based on single time frames, which may not fully represent what is going on in real time as the seasons change. Similarly, another temporal issue with using static datasets was that some datasets were not from the same time frame. Different datasets had varying availabilities. For example, the DSM used in the SVF calculation was obtained from the AHN4 dataset. The AHN collects data over a four year period repeating. This meant that the data would only be renewed every four years, and not all areas within within the study region may not be captured exactly on the same day or month. As a result, the DSM used for SVF calculations may not fully reflect the overall current morphology.

The second example here is the imperviousness dataset. The imperviousness dataset was obtained from the Copernicus open data platform. The specific imperviousness dataset used was from the year 2018. This was selected because it was the latest dataset available for this feature. The next dataset for imperviousness is the 2021 dataset, which had not been published. Using this as a feature in combination with other features from different time frames may decrease overall accuracy in predictions or be taken into account when interpreting results.

Differences in the time frames of input datasets may have negatively impacted the overall model performance. Machine learning models rely on consistent relationships between features and target variables. When datasets are drawn from different years or seasons, changes in land use, vegetation cover, or urban morphology may introduce noise or outdated information. For example, imperviousness from 2018 may not align with tree data or heat data from more recent years, reducing the coherence of the input space. This temporal discrepancy can weaken the model's ability to learn accurate patterns and reduce predictive reliability.

Lastly, the target variables used are representative of key resilience functions. However, only cover a subset of possible urban resilience indicators. Other relevant indicators such as social vulnerability or economic pressures were not included due to data limitations. This narrows down the generalisation of the models, although it may still provide a meaningful baseline for urban system analysis at street levels.

Overall, these limitations highlight the challenges of working with spatially complex systems using machine learning. This reinforces the need for improved data infrastructure when applying this framework to a real world context.

5.5. Recommendations for Future Work

The first recommendation for future work would be to improve data resolution and availability. Several key features such as traffic intensity, wind direction or vegetation quality and age are missing and may have improved overall performance, especially for more complex targets such as NO_2 where current predictors may not have fully explained the observed variance.

Secondly, including temporal aspects could be valuable as this study currently uses year based datasets, but incorporating seasonal, daily or real time data may improve model performance and provide more meaningful insights for urban planners.

Additionally, in order to address the spatial autocorrelation in residuals, future work could be to apply spatial cross validation or location based weighted models. This could help

better generalise to different spatial contexts and reduce the potential clusters of model errors.

Moreover, using only supervised learning could have reduced the potential maximum accuracy achieved. Strong unsupervised models such as convolutional neural networks may have been able to capture localised spatial patterns, particularly in specific areas or microclimates that the generalised models were not able to.

Another recommendation that could be done is include scenario based modelling. This framework could be extended into simulating greening interventions, such as additional tree planting or redesigning urban corridors. This would allow prediction of resilience gains and support more actionable use in urban planning and design contexts.

Finally, testing the framework in other cities and applying the method beyond just the urban areas of Rotterdam would test its transferability and reveal which features are consistently relevant across diverse urban environments. this comparative perspective would improve its overall robustness and wider applicability.

5.6. Conclusion

This study explored how machine learning can be applied to assess green infrastructure's impact on urban resilience at a street level scale in Rotterdam. Using three target indicators, Cooling effects of blue green infrastructure, maximum flood depth and NO_2, an attempt was made to cover different aspects of urban resilience. A set of interpretable, spatially derived features were used. These included perceived temperature, UHI, imperviousness, tree species richness, SVF, attraction reach and angular betweenness.

Among the 3 models, gradient boosting consistently showed the strongest predictive performance, particularly for cooling effects. Model generalisability varied across the different cases. Random Forest showed to be a close second where predictions and generalisably were similar to gradient boosting, only having slightly more R^2 values.

The predictions made by the model aligned with existing patterns to an extent. For example, there are strong cooling effects near dense vegetation and water bodies. Similarly, flooding patterns were not easy to predict and produced a lower R^2 score as flooding simulations are influenced by many more morphological and environmental features than what was included in the feature set. NO_2 concentrations were less predictable as well. This was likely due to missing variables such as traffic congestion and wind direction.

While reviewing the prediction results, residual map plots revealed spatial autocorrelation, suggested that errors were not all random but related to localised landscapes. This meant that the model was having trouble with predicting cases where the urban morphology or environment was unique or uncommon.

Additionally, using this framework showed that street scale features may offer value in modelling urban resilience outcomes. This is because using a finer resolution than neighbourhood or city level assessments provide smaller samples sizes. Originally, using neighbourhoods as training samples was considered. However, the sample size of neighbourhoods in Rotterdam could not compare to using the number of streets.

The framework from this study may be useful for local authorities, policy makers, and urban planners who aim to make data driven decisions on where and how to implement green

infrastructure. The features used in the model are spatially interpretable and derived from accessible datasets. This makes the feature importance clearer and allows for policy makers to indicate what specifically can help the urban resilience related target per case in terms of practical planning strategies. For example, areas with lower residual errors may benefit from targeted interventions by looking at feature importance of predictions. By identifying the street segments with lower resilience performance, resources can be better prioritised for smaller scaled and more location specific solutions to support climate adaptation goals.

Finally, the study contributes a practical workflow that supports greener, more climate resilient cities. This helps highlight the potential of predictive modelling for spatial planning and green infrastructure prioritisation. The framework can also provide valuable early-stage insight for spatial planning, helping to identify priority areas and inform preliminary decisions before more detailed assessments or design interventions are undertaken.

A. Reproducibility self-assessment

A.1. Marks for each of the criteria

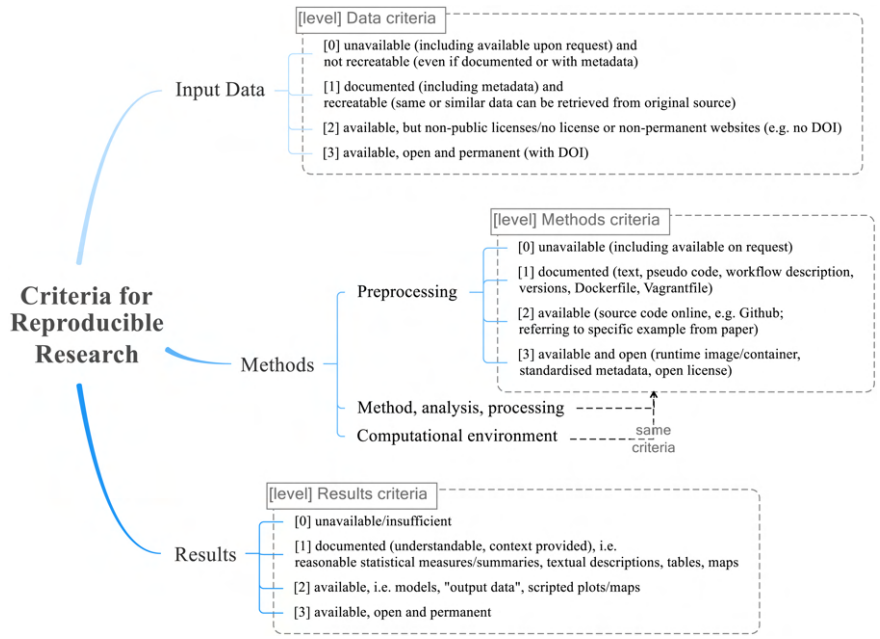

Figure A.1.: Reproducibility criteria to be assessed.

Grade/evaluate yourself for the 5 criteria (giving 0/1/2/3 for each):

1. input data - 2
2. preprocessing - 1
3. methods - 1
4. computational environment - 1
5. results - 2

A.2. Self-reflection

In terms of input data, this thesis mostly relied on openly accessible datasets such as those from the Atlas Natuurlijk Kapitaal and Klimaateffectatlas. These sources are openly published and have standardised formats as well as detailed metadata, making them easy to

retrieve. The only exception was the tree inventory dataset. This was requested from the municipality of Rotterdam's open data portal.

To preprocess the data, most of the spatial operations were performed in QGIS. These steps included spatial joins, zoning buffers and calculating zonal statistics. In python, preprocessing was kept minimal. This mostly involved dropping nulls, aligning indexes and loading preprocessed data into our models. Although the full process was not containerised or automated, it was made using a logical structure with an aim for reproducibility for someone who may be familiar with the tools used.

The methods used in this study involved widely available machine learning algorithms from scikitlearn. The scripts were also structured so that they can be rerun with minimal changes. Hyperparameter tuning grids were documented and evaluation metrics used such as MAE and RMSE were also used across all the models to ensure consistency. A fixed random state (42) was used throughout model training and evaluation to ensure reproducibility of results.

The computational environment used follows a standard Python setup using popular packages such as geopandas, scikitlearn, and matplotlib. These libraries are basic to install using pip or conda. An environment using basic packages should be easily reproduced with small effort.

In regards to results, and the outputs, spatial maps of predicted values, residuals as well as performance tables are clearly documented and labelled. The figures use standard symbology and colour scales. They follow a logical naming structure that connects them to the corresponding cases and models.

B. Feature Importance plots

This chapter shows the feature importance plots by case as described in section 4.3.1.

B.1. Case 1 Feature Importance

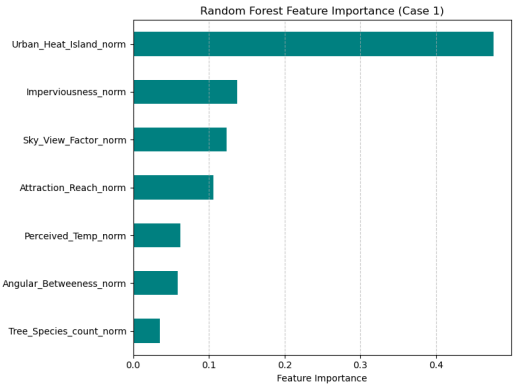

(a) Random Forest Feature Importance (Case 1)

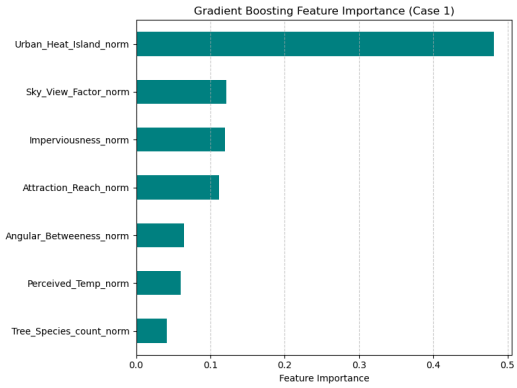

(b) Gradient Boosting Feature Importance (Case 1)

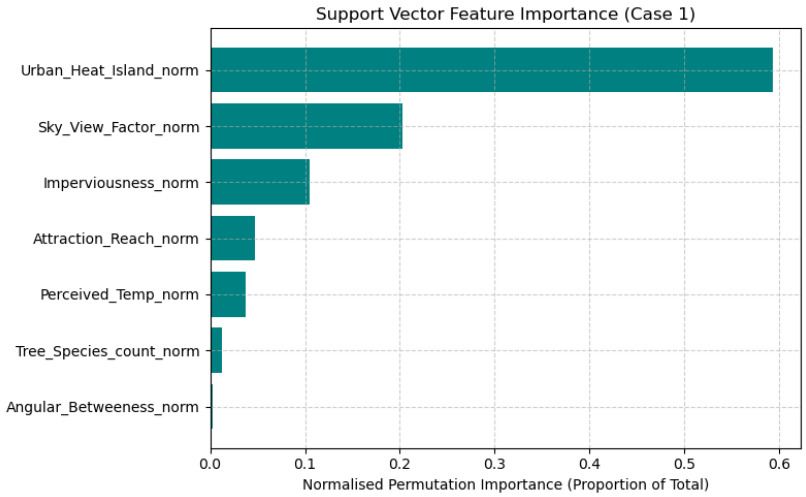

(c) Support Vector Feature Importance (Case 1)

Figure B.1.: Feature Importance for each model (Case 1)

B.2. Case 2 Feature Importance

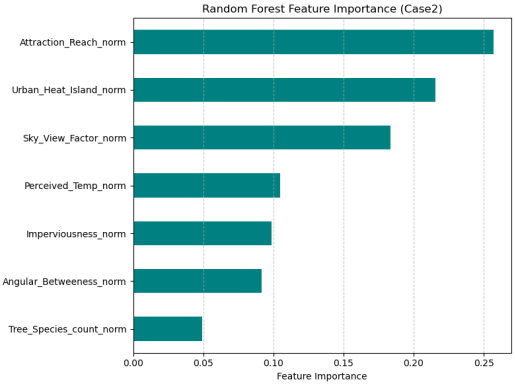

(a) Random Forest Feature Importance (Case 2)

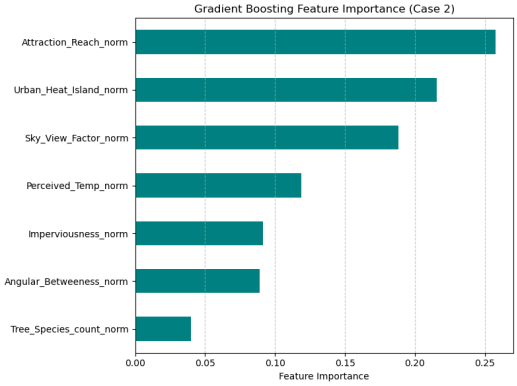

(b) Gradient Boosting Feature Importance (Case 2)

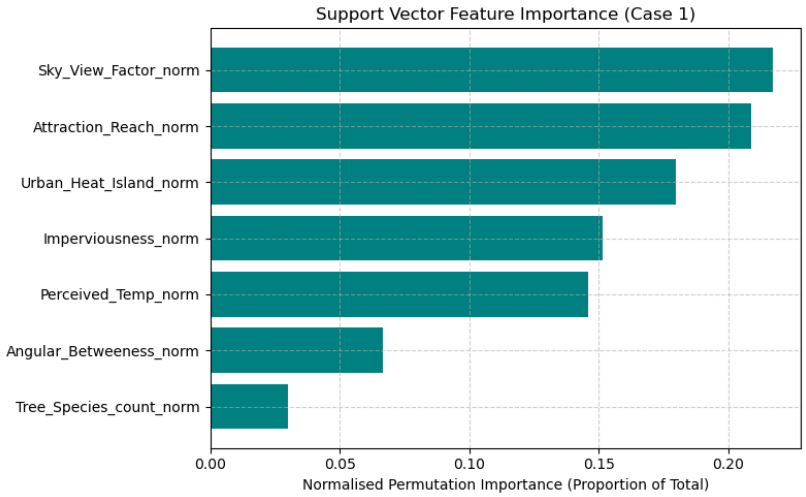

(c) Support Vector Feature Importance (Case 2)

Figure B.2.: Feature Importance for each model (Case 2)

B.3. Case 3 Feature Importance

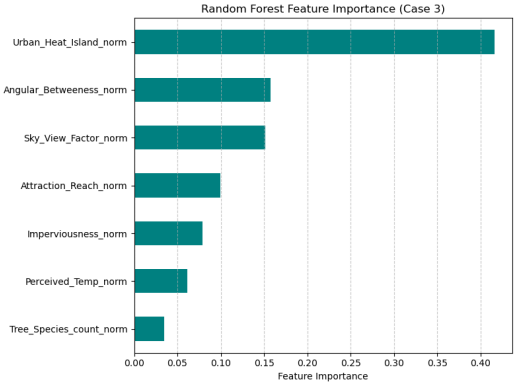

(a) Random Forest Feature Importance (Case 3)

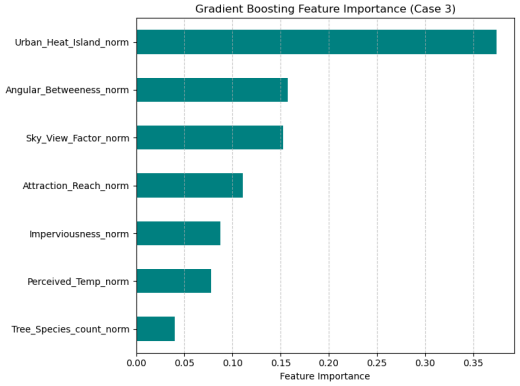

(b) Gradient Boosting Feature Importance (Case 3)

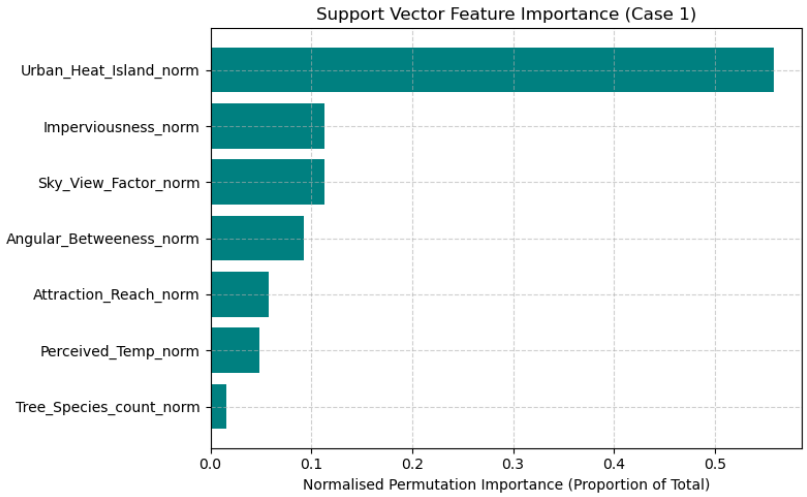

(c) Support Vector Feature Importance (Case 3)

Figure B.3.: Feature Importance for each model (Case 3)

C. Spatial Distribution Plots

This section highlights the remaining predicted and residual spatial distribution plots from the losing models in sections 4.3.1, 4.3.2, and 4.3.3.

C.1. Case 1 Predicted and Residual Plots

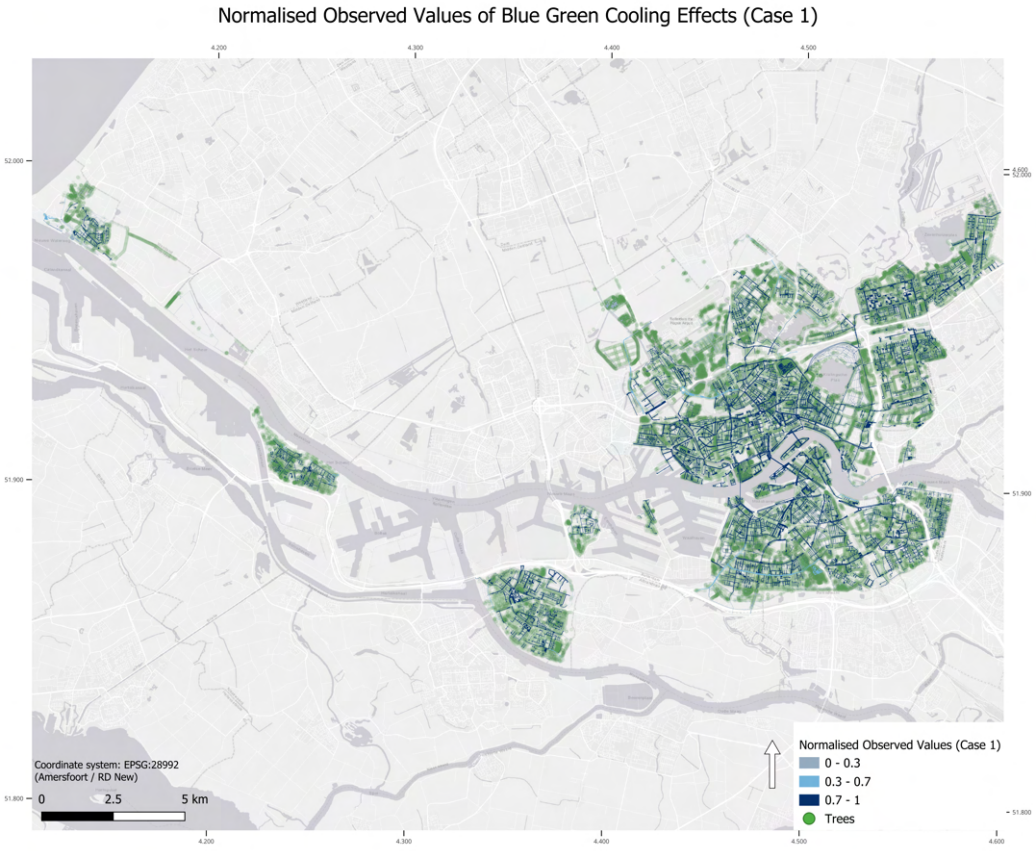

Figure C.1.: City Wide Normalised Observed Values of Blue Green Cooling Effects (Case 1)

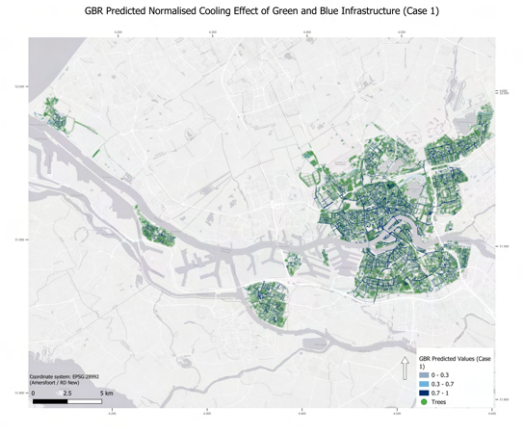

(a) City Wide GBR Predicted Normalised Cooling Effect of Green and Blue Infrastructure (Case 1)

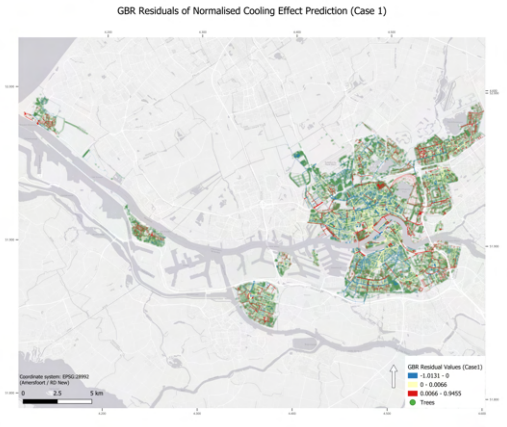

(b) RFR Residuals of Cooling Effect (Case 1)

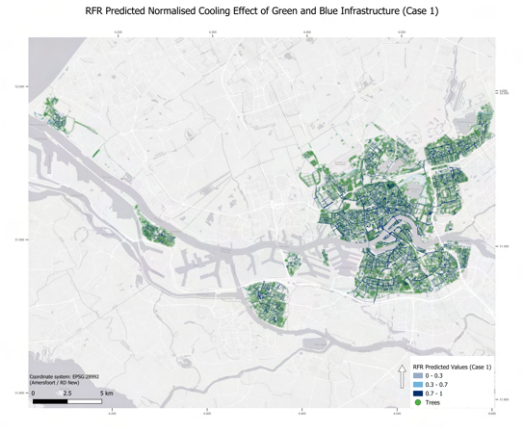

(c) RFR Predicted Cooling Effect (Case 1)

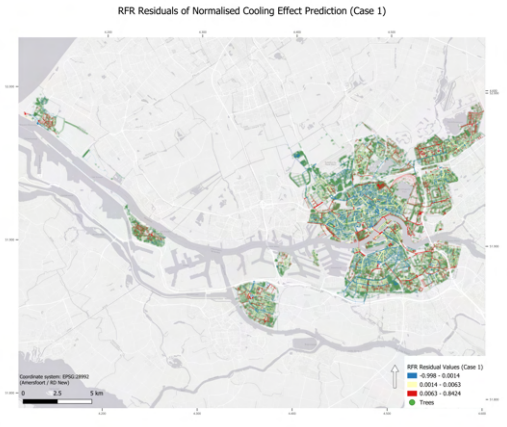

(d) RFR Residuals of Cooling Effect (Case 1)

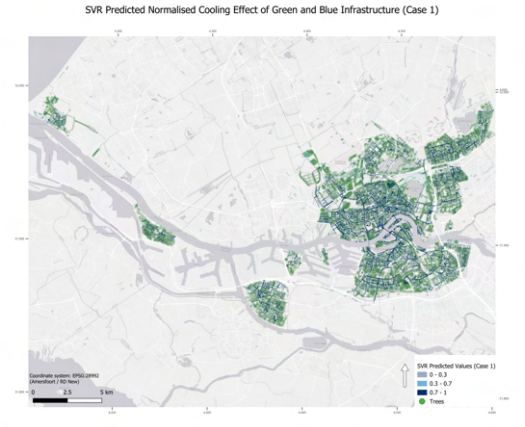

(e) SVR Predicted Cooling Effect (Case 1)

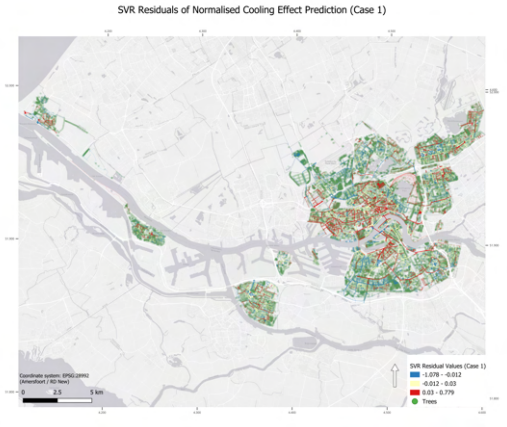

(f) SVR Residuals of Cooling Effect (Case 1)

Figure C.2.: Predicted and residual spatial patterns of cooling effect for RFR and SVR models (Case 1).

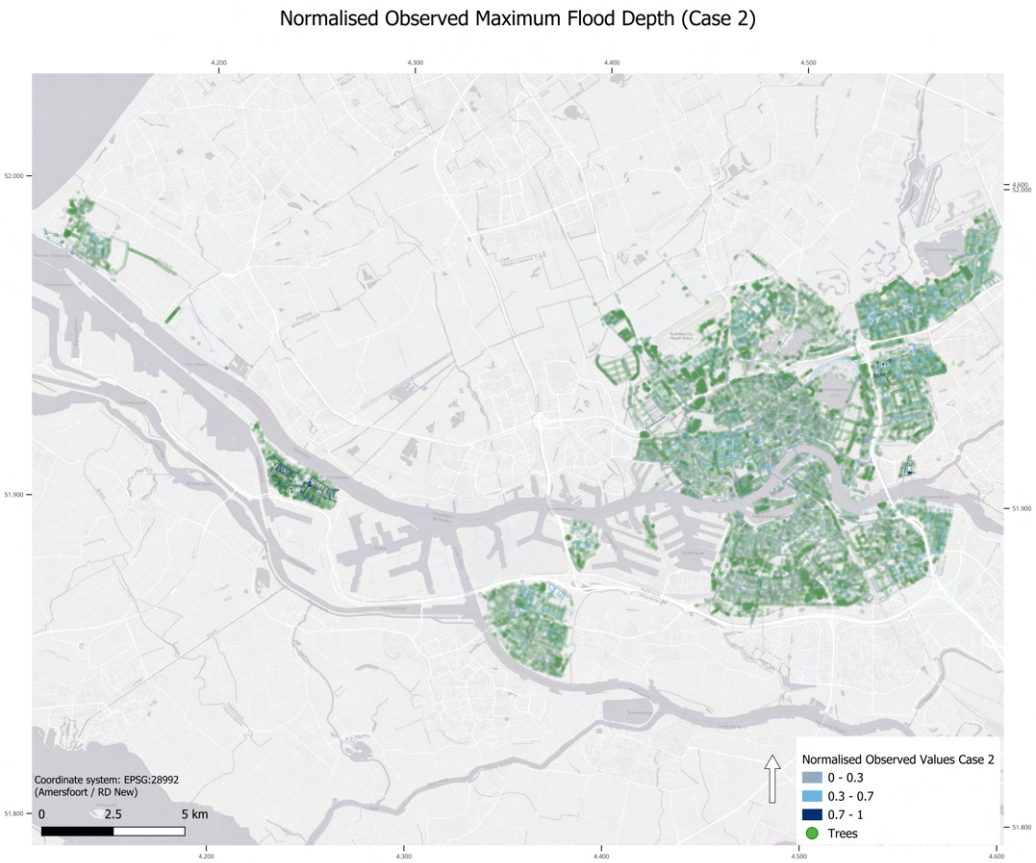

Figure C.3.: City Wide Normalised Observed Maximum Flood Depth (Case 2)

C.2. Case 2 Predicted and Residual Plots

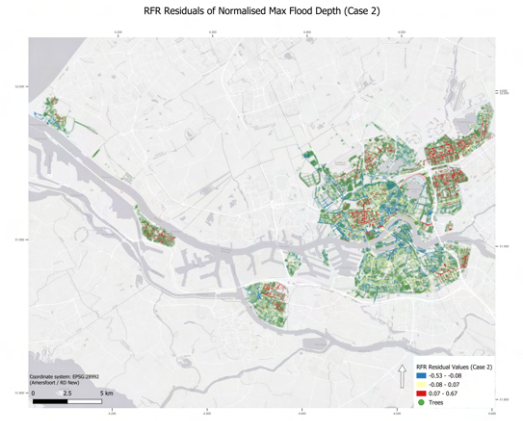

(a) RFR Predicted Max Flood Depth (Case 2)

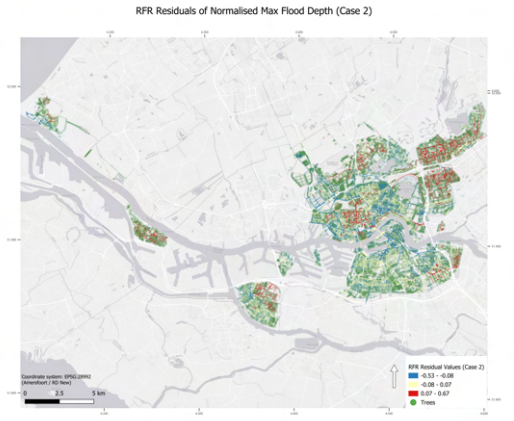

(b) RFR Residuals of Max Flood Depth (Case 2)

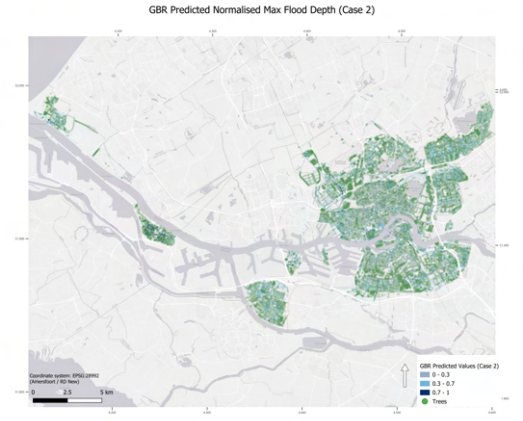

(c) GBR Predicted Max Flood Depth (Case 2)

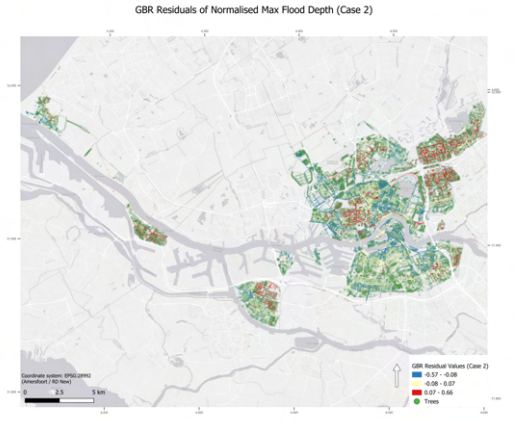

(d) GBR Residuals of Max Flood Depth (Case 2)

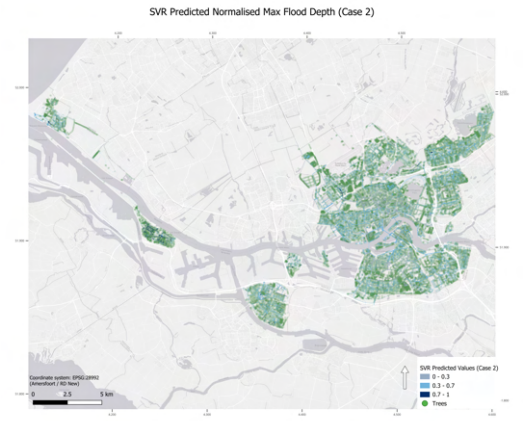

(e) SVR Predicted Max Flood Depth (Case 2)

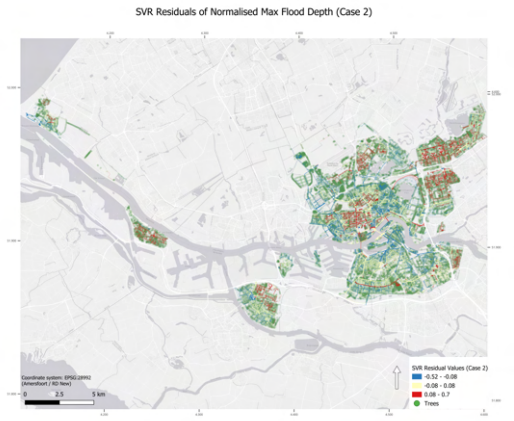

(f) SVR Residuals of Max Flood Depth (Case 2)

Figure C.4.: Predicted and residual spatial patterns of flood depth from GBR and SVR models (Case 2).

C.3. Case 3 Predicted and Residual Plots

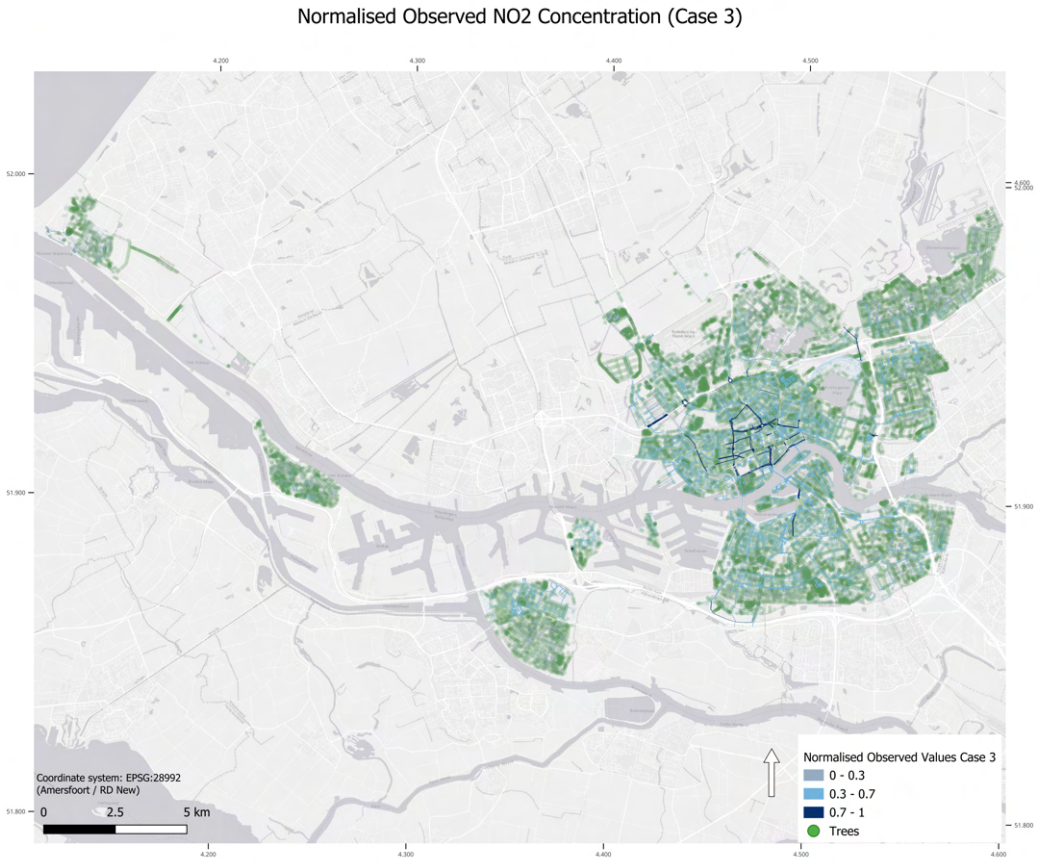

Figure C.5.: City Wide Normalised Observed NO2 Concentration (Case 3)

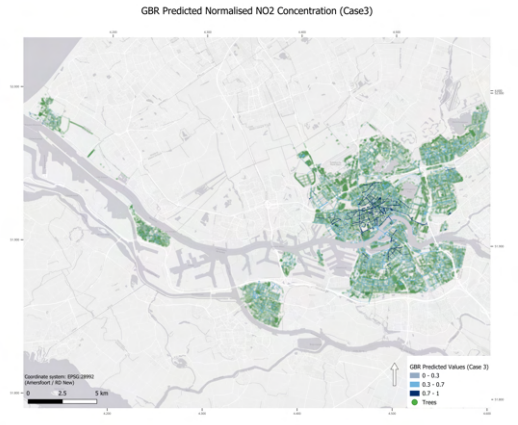

(a) GBR Predicted NO_2 Concentration (Case 3)

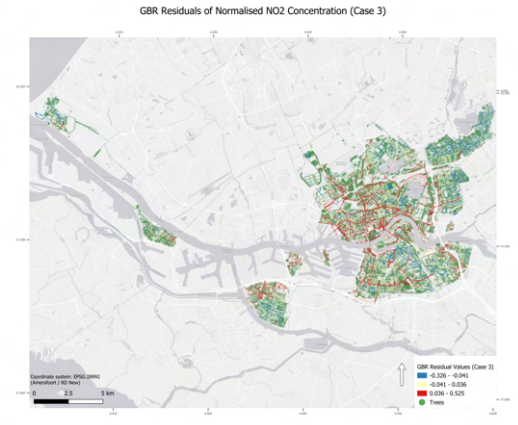

(b) GBR Residuals of NO_2 Concentration (Case 3)

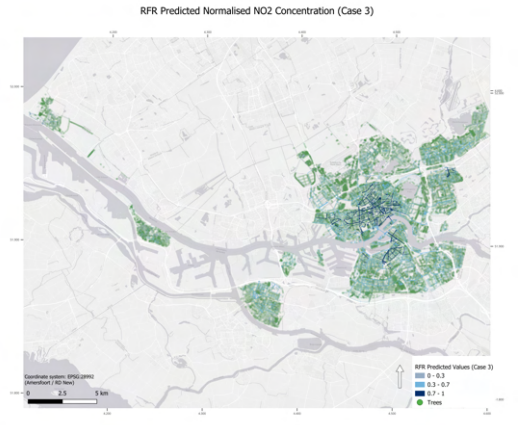

(c) RFR Predicted NO_2 Concentration (Case 3)

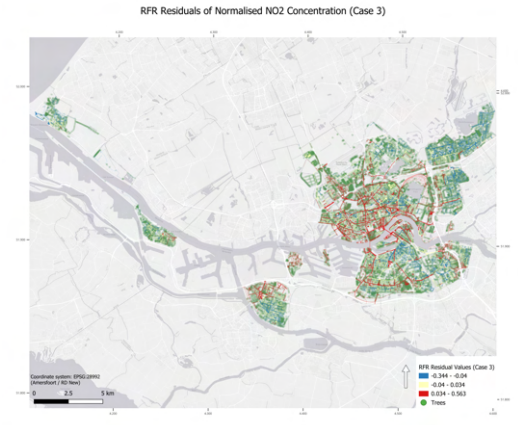

(d) RFR Residuals of NO_2 Concentration (Case 3)

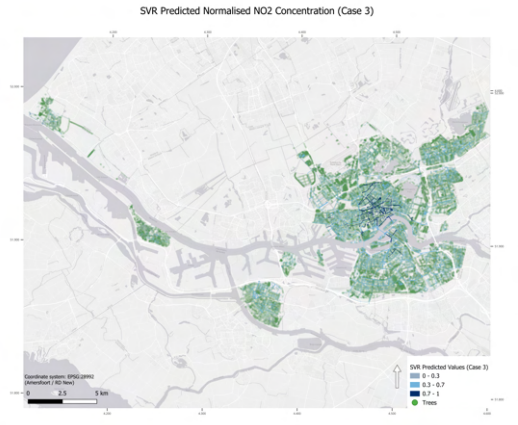

(e) SVR Predicted NO_2 Concentration (Case 3)

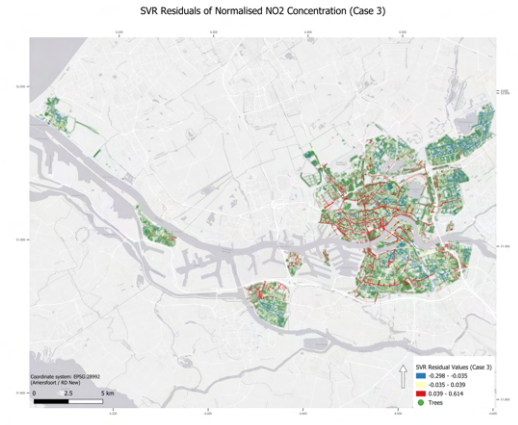

(f) SVR Residuals of NO_2 Concentration (Case 3)

Figure C.6.: Predicted and residual spatial patterns of NO_2 concentrations from RFR and SVR models (Case 3).

D. Learning Curves From best Models

D.1. Case 1 Best Models

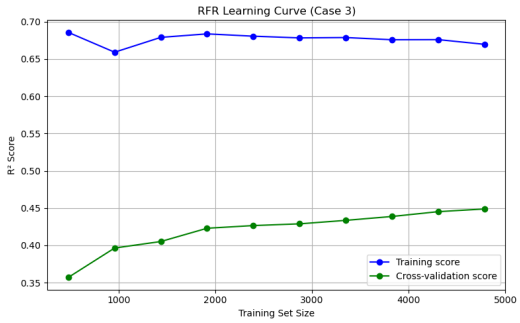

(a) RFR Learning Curve (Case 1)

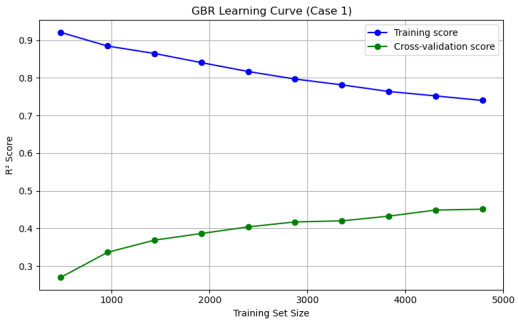

(b) GBR Learning Curve (Case 1)

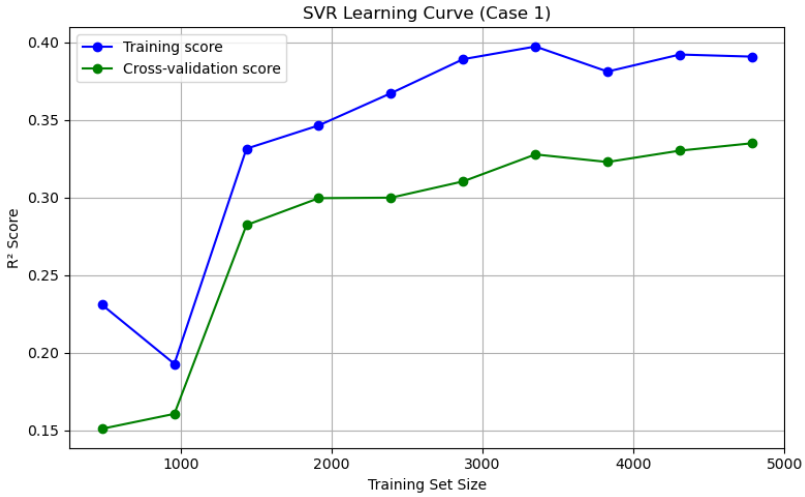

(c) SVR Learning Curve (Case 1)

Figure D.1.: Learning curves for each model in Case 1.

D.2. Case 2 Best Models

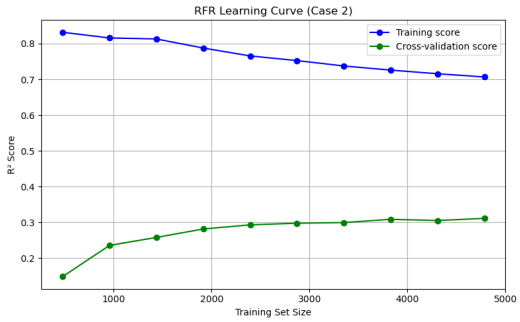

(a) RFR Learning Curve (Case 2)

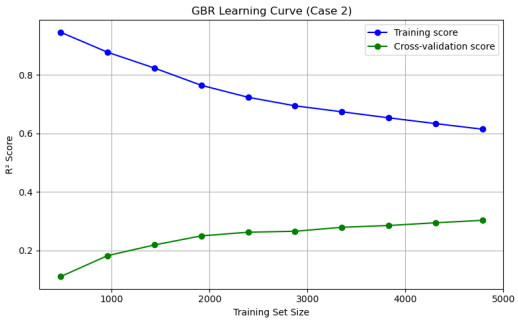

(b) GBR Learning Curve (Case 2)

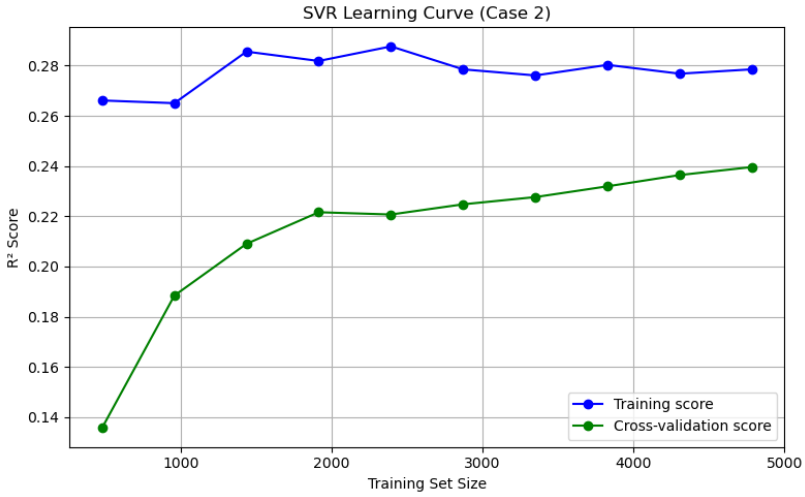

(c) SVR Learning Curve (Case 2)

Figure D.2.: Learning curves for each model in Case 2

D.3. Case 3 Best Models

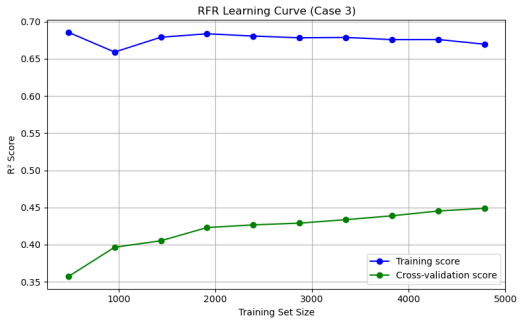

(a) RFR Learning Curve (Case 3)

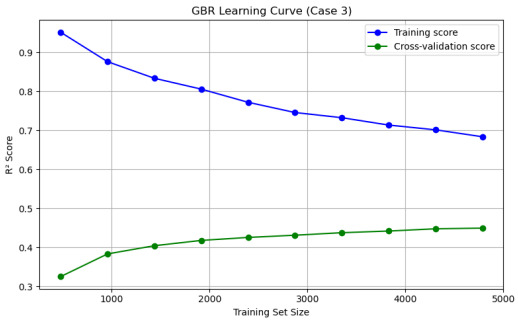

(b) GBR Learning Curve (Case 3)

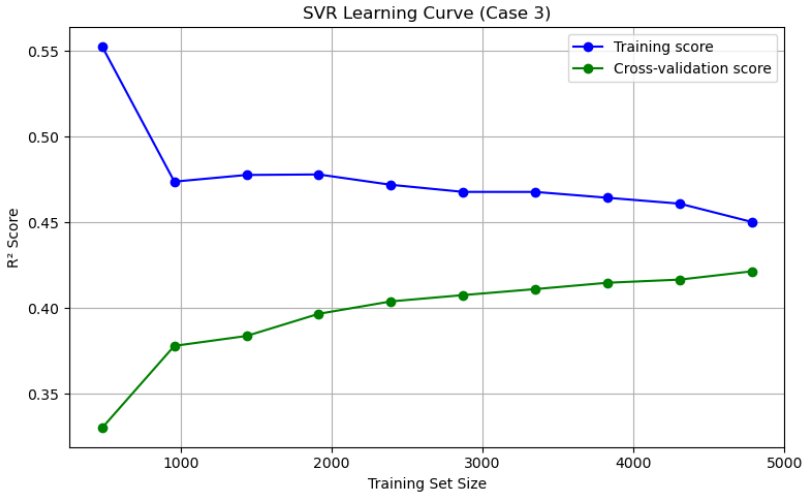

(c) SVR Learning Curve (Case 3)

Figure D.3.: Learning curves for each model in Case 3

Bibliography

Ahern, J. (2011). From fail-safe to safe-to-fail: Sustainability and resilience in the new urban world. *Landscape and urban Planning*, 100(4):341–343.

Altmann, A., Toloşi, L., Sander, O., and Lengauer, T. (2010). Permutation importance: a corrected feature importance measure. *Bioinformatics*, 26(10):1340–1347.

Amado, C. (2022). Unravelling the urban heat island phenomenon in the netherlands. a multicity spatial analysis on the distributive component of environmental justice, analysing the urban green infrastructure, and the urban heat island effect. *A Multicity Spatial Analysis on the Distributive component of Environmental Justice, analysing the Urban Green Infrastructure, and the Urban Heat Island Effect.(January 12, 2022)*.

Andrés-Doménech, I., Hernández-Crespo, C., Martín, M., and Andrés-Valeri, V. C. (2018). Characterization of wash-off from urban impervious surfaces and suds design criteria for source control under semi-arid conditions. *Science of The Total Environment*, 612:1320–1328.

Anguelovski, I., Connolly, J. J., Cole, H., Garcia-Lamarca, M., Triguero-Mas, M., Baró, F., Martin, N., Conesa, D., Shokry, G., Del Pulgar, C. P., et al. (2022). Green gentrification in european and north american cities. *Nature communications*, 13(1):3816.

Arlot, S. and Celisse, A. (2010). A survey of cross-validation procedures for model selection.

Awad, M., Khanna, R., Awad, M., and Khanna, R. (2015). Support vector regression. *Efficient learning machines: Theories, concepts, and applications for engineers and system designers*, pages 67–80.

Bartz-Beielstein, T. (2023). Hyperparameter tuning cookbook: A guide for scikit-learn, pytorch, river, and spotpython. *arXiv preprint arXiv:2307.10262*.

Bowler, D. E., Buyung-Ali, L., Knight, T. M., and Pullin, A. S. (2010). Urban greening to cool towns and cities: A systematic review of the empirical evidence. *Landscape and urban planning*, 97(3):147–155.

Brears, R. C. and Brears, R. C. (2018). Rotterdam becoming a blue-green city. *Blue and Green Cities: The Role of Blue-Green Infrastructure in Managing Urban Water Resources*, pages 183–204.

Breiman, L. (2001). Random forests. *Machine learning*, 45:5–32.

Chai, T. and Draxler, R. R. (2014). Root mean square error (rmse) or mean absolute error (mae)?–arguments against avoiding rmse in the literature. *Geoscientific model development*, 7(3):1247–1250.

Chawla, L. (2015). Benefits of nature contact for children. *Journal of planning literature*, 30(4):433–452.

Cong, Y. and Inazumi, S. (2024). Enhancing urban resilience with smart technologies and predictive analytics in setagaya, tokyo.

Dadvand, P., Nieuwenhuijsen, M. J., Esnaola, M., Forns, J., Basagaña, X., Alvarez-Pedrerol, M., Rivas, I., López-Vicente, M., De Castro Pascual, M., Su, J., et al. (2015). Green spaces and cognitive development in primary schoolchildren. *Proceedings of the national academy of sciences*, 112(26):7937–7942.

de Amorim, L. B., Cavalcanti, G. D., and Cruz, R. M. (2023). The choice of scaling technique matters for classification performance. *Applied Soft Computing*, 133:109924.

Derkzen, M. L., Van Teeffelen, A. J., and Verburg, P. H. (2017). Green infrastructure for urban climate adaptation: How do residents' views on climate impacts and green infrastructure shape adaptation preferences? *Landscape and urban planning*, 157:106–130.

Dircke, P. and Molenaar, A. (2015). Climate change adaptation; innovative tools and strategies in delta city rotterdam. *Water Practice and Technology*, 10(4):674–680.

Dirksen, M., Ronda, R., Theeuwes, N., and Pagani, G. (2019). Sky view factor calculations and its application in urban heat island studies. *Urban Climate*, 30:100498.

Filazzola, A., Shrestha, N., and MacIvor, J. S. (2019). The contribution of constructed green infrastructure to urban biodiversity: A synthesis and meta-analysis. *Journal of Applied Ecology*, 56(9):2131–2143.

Fletcher, T. (2009). Support vector machines explained. *Tutorial paper*, 1118:1–19.

Frantzeskaki, N. and Tilie, N. (2014). The dynamics of urban ecosystem governance in rotterdam, the netherlands. *Ambio*, 43:542–555.

Frantzeskaki, N., Wittmayer, J., and Loorbach, D. (2014). The role of partnerships in 'realising'urban sustainability in rotterdam's city ports area, the netherlands. *Journal of Cleaner Production*, 65:406–417.

Friedman, J. H. (2001). Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232.

Fushiki, T. (2011). Estimation of prediction error by using k-fold cross-validation. *Statistics and Computing*, 21:137–146.

Ghofrani, Z., Sposito, V., and Faggian, R. (2017). A comprehensive review of blue-green infrastructure concepts. *International Journal of Environment and Sustainability*, 6(1):15–36.

Gkontzis, A. F., Kotsiantis, S., Feretzakis, G., and Verykios, V. S. (2024). Enhancing urban resilience: smart city data analyses, forecasts, and digital twin techniques at the neighborhood level. *Future Internet*, 16(2):47.

Kabisch, N. and Van Den Bosch, M. A. (2017). Urban green spaces and the potential for health improvement and environmental justice in a changing climate. *Nature-Based Solutions to Climate Change Adaptation in Urban Areas: Linkages between Science, Policy and Practice*, pages 207–220.

Katrutsa, A. and Strijov, V. (2017). Comprehensive study of feature selection methods to solve multicollinearity problem according to evaluation criteria. *Expert Systems with Applications*, 76:1–11.

Korkou, M., Tarigan, A. K., and Hanslin, H. M. (2023). The multifunctionality concept in urban green infrastructure planning: A systematic literature review. *Urban Forestry & Urban Greening*, 85:127975.

Kuo, F. E. and Sullivan, W. C. (2001). Environment and crime in the inner city: Does vegetation reduce crime? *Environment and behavior*, 33(3):343–367.

Kutty, A. A., Wakjira, T. G., Kucukvar, M., Abdella, G. M., and Onat, N. C. (2022). Urban resilience and livability performance of european smart cities: A novel machine learning approach. *Journal of Cleaner Production*, 378:134203.

Livesley, S. J., Escobedo, F. J., and Morgenroth, J. (2016). The biodiversity of urban and peri-urban forests and the diverse ecosystem services they provide as socio-ecological systems.

Mahesh, B. et al. (2020). Machine learning algorithms-a review. *International Journal of Science and Research (IJSR).[Internet]*, 9(1):381–386.

Mason, B. E., Schmidt, J., and Kerkez, B. (2023). Measuring city-scale green infrastructure drawdown dynamics using internet-connected sensors in detroit. *Environmental Science: Water Research & Technology*, 9(12):3213–3226.

Meerow, S., Newell, J. P., and Stults, M. (2016). Defining urban resilience: A review. *Landscape and urban planning*, 147:38–49.

Natekin, A. and Knoll, A. (2013). Gradient boosting machines, a tutorial. *Frontiers in neurorobotics*, 7:21.

O'brien, R. M. (2007). A caution regarding rules of thumb for variance inflation factors. *Quality & quantity*, 41:673–690.

Pamukcu-Albers, P., Ugolini, F., La Rosa, D., Grădinaru, S. R., Azevedo, J. C., and Wu, J. (2021). Building green infrastructure to enhance urban resilience to climate change and pandemics. *Landscape ecology*, 36(3):665–673.

Perrelet, K., Moretti, M., Dietzel, A., Altermatt, F., and Cook, L. M. (2024). Engineering blue-green infrastructure for and with biodiversity in cities. *npj Urban Sustainability*, 4(1):27.

Plonsky, L. and Ghanbar, H. (2018). Multiple regression in l2 research: A methodological synthesis and guide to interpreting r2 values. *The Modern language journal*, 102(4):713–731.

Rijksinstituut voor Volksgezondheid en Milieu (RIVM) (2020). Verkoelend effect van groen en blauw. https://www.atlasnatuurlijkkapitaal.nl/. Raster dataset showing cooling effects (in °C) of vegetation and water on urban temperatures. Source date: 2020-01-01. License: CC0 1.0. Based on data from CBS, RIVM, and others.

Rijksinstituut voor Volksgezondheid en Milieu (RIVM) (2025). Nitrogen dioxide 2022 (no). https://data.rivm.nl/. Accessed: 2025-04-14.

Rijkswaterstaat (2023). Maximum water depth map – extremely small chance scenario. https://www.liwo.nl. Accessed: 2025-04-14.

Rodriguez-Galiano, V., Sanchez-Castillo, M., Chica-Olmo, M., and Chica-Rivas, M. (2015). Machine learning predictive models for mineral prospectivity: An evaluation of neural networks, random forest, regression trees and support vector machines. *Ore geology reviews*, 71:804–818.

Salih, K. and Báthoryné Nagy, I. R. (2024). Review of the role of urban green infrastructure on climate resiliency: A focus on heat mitigation modelling scenario on the microclimate and building scale. *Urban Science*, 8(4):220.

Scikit-learn (2007–2021a). Cross-validation: Evaluating estimator performance. https://scikit-learn.org/stable/modules/cross_validation.html. Accessed: 2025-03-25.

Scikit-learn (2007–2021b). Feature importances with a forest of trees. https://scikit-learn.org/stable/auto_examples/ensemble/plot_forest_importances.html. Accessed: 2025-04-14.

Staddon, C., Ward, S., De Vito, L., Zuniga-Teran, A., Gerlak, A. K., Schoeman, Y., Hart, A., and Booth, G. (2018). Contributions of green infrastructure to enhancing urban resilience. *Environment Systems and Decisions*, 38:330–338.

Ståhle, A. (2012). Place syntax tool (pst). *Accessibility instruments for planning practice*, pages 173–178.

Stewart, I. D. and Oke, T. R. (2012). Local climate zones for urban temperature studies. *Bulletin of the American Meteorological Society*, 93(12):1879–1900.

Tillie, N., Borsboom-van Beurden, J., Doepel, D., and Aarts, M. (2018). Exploring a stakeholder based urban densification and greening agenda for rotterdam inner city—accelerating the transition to a liveable low carbon city. *Sustainability*, 10(6):1927.

van der Berg, A. (2023). Climate adaptation planning for resilient and sustainable cities: Perspectives from the city of rotterdam (netherlands) and the city of antwerp (belgium). *European Journal of Risk Regulation*, 14(3):564–582.

Wanghe, K., Guo, X., Wang, M., Zhuang, H., Ahmad, S., Khan, T. U., Xiao, Y., Luan, X., and Li, K. (2020). Gravity model toolbox: An automated and open-source arcgis tool to build and prioritize ecological corridors in urban landscapes. *Global Ecology and Conservation*, 22:e01012.

Willmott, C. J. and Matsuura, K. (2005). Advantages of the mean absolute error (mae) over the root mean square error (rmse) in assessing average model performance. *Climate research*, 30(1):79–82.

Xue, F., Luan, B., Fan, Y., Xie, S., Yang, X., Luo, J., and Zheng, R. (2024). Assessing the lifecycle environmental resilience of urban green infrastructures coping with acute disturbances and chronic stresses. *Water*, 16(8):1162.

Ziter, C. D., Pedersen, E. J., Kucharik, C. J., and Turner, M. G. (2019). Scale-dependent interactions between tree canopy cover and impervious surfaces reduce daytime urban heat during summer. *Proceedings of the National Academy of Sciences*, 116(15):7575–7580.

Colophon

This document was typeset using LaTeX, using the KOMA-Script class `scrbook`. The main font is Palatino.