








GeoAI Unveiled: Case Studies in Explainable GeoAI for Environmental Modeling

Courage Kamusoko



GeoAI Unveiled: Case Studies in Explainable GeoAI for Environmental Modeling

Courage Kamusoko
Machida, Tokyo, Japan
Ai.Geolabs Publishing

Copyright© Courage Kamusoko, 2024

This work is subject to copyright. All rights are solely and exclusively licensed by the author, whether the whole or part of the material is concerned, specifically, the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc., in this publication does not imply, even without a specific statement, that such names are exempt from the relevant protective laws and regulations and, therefore, free for general use.

The author is safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. The author gives a warranty, expressed or implied, concerning the material contained herein or for any errors or omissions that may have been made. The author and publisher remain neutral regarding jurisdictional claims in published maps and institutional affiliations.

Preface

Geospatial Artificial Intelligence (GeoAI) is rapidly reshaping how we interact with the environment, opening new possibilities for understanding and managing the Earth's resources. In this new landscape, where AI and geospatial technologies converge, the ability to explain and trust AI-driven decisions is becoming a priority. *GeoAI Unveiled: Case Studies in Explainable GeoAI for Environmental Modeling* bridges the critical gap between cutting-edge GeoAI models and the transparency required to deploy them responsibly in environmental contexts.

In the past decade, AI techniques like machine learning (ML) and deep learning (DL) have driven remarkable advancements in geospatial applications. These technologies are now at the forefront of solutions for some of the world's most pressing environmental challenges—deforestation, climate change, urbanization, and natural disasters. However, while AI brings unprecedented predictive power, it also introduces complexity. The often opaque nature of AI models—commonly called "black boxes"—creates barriers to trust, particularly in high-stakes scenarios like forest conservation, flood disaster management, urban development, and climate monitoring.

Explainable AI (XAI) offers a solution by transforming complex AI systems into more transparent and understandable models. XAI also allows users to see what a model predicts and why it makes those predictions. This capability is especially vital in geospatial contexts, where AI outputs must inform environmental stewardship, public safety, and resource management decisions.

This introductory book explores the powerful intersection of XAI and GeoAI through real-world case studies. Chapter by chapter, it dives into critical applications of explainable GeoAI in environmental modeling, starting with forest structure and extending to flood inundation mapping. Each case study demonstrates the value of XAI techniques like SHAP (Shapley Additive Explanations) and Grad-CAM (Gradient-weighted Class Activation Mapping).

GeoAI Unveiled is not just for geospatial researchers or environmental scientists but for all stakeholders engaged in deploying GeoAI for environmental solutions. Whether you are a graduate student, researcher, policymaker, urban planner, or disaster response manager, this book empowers you to harness the full potential of GeoAI. It also ensures that GeoAI models remain transparent, explainable, and accountable. The book equips you with the knowledge and tools to confidently navigate the complex landscape of GeoAI in environmental monitoring and management.

In the pages ahead, you will discover how explainable GeoAI can help overcome the trust and transparency challenges that have long limited the use of AI in environmental modeling. More than a technical guide, this book offers a forward-looking vision of how responsible, explainable AI can reshape our ability to solve complex environmental problems. This makes a tangible difference in managing and protecting the Earth's natural resources.

As we embark on this journey through GeoAI, I invite you to think beyond the "black box" and embrace the promise of an explainable future. This journey in which explainable GeoAI serves as a powerful tool and a trusted partner in addressing the environmental challenges of our time.

How is this Book Organized?

GeoAI Unveiled: Case Studies in Explainable GeoAI for Environmental Modeling is structured to guide readers through the core concepts, techniques, and applications of XAI within GeoAI. The book is divided into four key chapters. The aim is to provide a comprehensive understanding of how GeoAI

can be applied to solve environmental challenges while ensuring the transparency and explainability of AI models. The book is organized as follows.

Chapter 1: Introduction to Explainable GeoAI

This chapter lays the foundation for the book by introducing the concept of GeoAI and its increasing relevance in geospatial technologies. It explores the rapid growth of AI-driven geospatial applications and emphasizes the importance of explainability in these models. The chapter introduces the core principles behind XAI, covering techniques like SHAP (Shapley Additive Explanations), LIME (Local Interpretable Model-agnostic Explanations), and Grad-CAM (Gradient-weighted Class Activation Mapping). Readers will gain a solid understanding of why explainability is crucial, particularly in environmental modeling.

Chapter 2: Case Study 1: Explainable Machine Learning for Modeling Forest Structure

This chapter presents the first case study, which focuses on modeling forest structure. The case study demonstrates how to apply random forest regression to predict forest canopy height using Earth observation data such as GEDI canopy height data and Sentinel-2 imagery. A thorough explanation of the importance of the predictors is provided using SHAP values. Readers will learn how different predictors—such as Sentinel-2 bands and spectral indices—impact the model's outputs. This case study highlights the importance of explainable ML for assessing forest structure.

Chapter 3: Case Study 2: Explainable Deep Learning for Flood Inundation Mapping

In the second case study, deep learning (DL) techniques are applied to map flood inundations using high-resolution imagery. The chapter introduces the U-Net model and demonstrates how it can be used to identify flooded areas. Grad-CAM is employed to visualize which parts of the satellite images influenced the model's predictions. This chapter highlights the importance of XAI in disaster management.

Chapter 4: The Power of Explainable AI: Why It's the Future of Geospatial Modeling

The final chapter summarizes the book and reflects on the strengths and limitations of current explainable GeoAI techniques. It emphasizes the need for transparency in GeoAI-driven decisions. This chapter also looks ahead, exploring emerging research directions and the importance of developing inherently interpretable models that address technical and ethical challenges. The role of explainable GeoAI in ensuring responsible and equitable use of AI is a central theme in this book.

Appendix

The book also includes an appendix that provides additional resources, such as further reading, for those interested in applying the techniques discussed. These resources are designed to help readers deepen their understanding and practical application of explainable GeoAI in their projects.

Table of Contents

Chapter 1. Introduction to Explainable GeoAI.....	1
1.1 The Rise of GeoAI.....	1
1.2 Why Explainability Matters in GeoAI.....	3
1.3 XAI Techniques.....	3
1.3.1 Overview of machine and deep learning techniques.....	3
1.3.2 Overview of XAI techniques.....	4
1.3.3 XAI techniques for GeoAI models.....	5
1.4 Applications of Explainable GeoAI for Environmental Modeling.....	6
1.5 Summary.....	7
Chapter 2. Case Study 1: Explainable Machine Learning for Modeling Forest Structure.....	8
2.1 Overview.....	8
2.2 Exploratory Data Analysis.....	9
2.2.1 GEDI canopy height data.....	9
2.2.2 Spectral reflectance analysis using boxplots.....	9
2.2.3 Spectral reflectance analysis using density plots.....	10
2.3 Evaluating Model Performance.....	11
2.4 Interpreting SHAP Values.....	14
2.4.1 Interpreting global SHAP values.....	14
2.4.2 Interpreting Local SHAP Values.....	18
2.4 Spatial Analysis of Predictor Contributions using SHAP.....	20
2.5 Summary.....	21
Chapter 3. Case Study 2: Explainable Deep Learning for Flood Inundation Mapping.....	23
3.1 Introduction.....	23
3.2 U-Net Model.....	24
3.3 Flood Inundation Maps.....	25
3.4 Interpreting Grad-CAM.....	26
3.5 Summary.....	27

Chapter 4. The Power of Explainable AI: Why It's the Future of Geospatial Modeling.....29

4.1 Summary.....29

4.2 Strengths and Limitations of Explainable GeoAI for Modeling Canopy Height.....29

4.3 Strengths and Limitations of Explainable GeoAI for Mapping Flood Inundation.....30

4.4 Future Work.....30

4.5 Next Steps: How to Unlock the Full Potential of Explainable GeoAI.....31

References.....32

Appendix.....34

Chapter 1. Introduction to Explainable GeoAI

1.1 The Rise of GeoAI

Undoubtedly, geospatial technology has become an integral part of our daily lives. In most parts of the world, geospatial data and technology contribute to our lives. We interact with it constantly—from Google Maps guiding our commutes to tracking a pizza delivery or simply checking the weather on our mobile phones. A few decades ago, however, processing geospatial datasets was a significant challenge due to limited storage capacity and computing power. Only experts had the capability to work with large geospatial datasets at national, regional or global scales. Now, thanks to advances in cloud computing and freely accessible geospatial data, generating even countrywide land cover maps has become far more feasible and efficient. Building on these advancements, geospatial technology is entering a new era—one driven by artificial intelligence (AI).

The new frontier for geospatial data and technology is Geospatial Artificial Intelligence (GeoAI), which is transforming the capabilities of geospatial applications. GeoAI integrates AI with geospatial data, science, and technology to solve complex geospatial problems and enhance understanding (VoPham et al., 2018; Li et al., 2024). Using AI techniques, GeoAI generates valuable geospatial insights from geospatial data (Pierdicca & Paolanti, 2022). Recent advancements in AI, particularly machine learning (ML) and deep learning (DL), combined with increased computing power, have propelled GeoAI forward. These developments have enabled diverse applications such as land cover mapping, deforestation monitoring, urban planning, and climate modeling (Liu & Biljecki, 2022).

The growth of GeoAI is also driven by the abundance of geospatial data and the swift progress in data handling and storage technology. Massive amounts of geospatial data—from high-resolution satellite imagery to real-time sensor observations—demand robust, scalable, and efficient storage and analytical capabilities. These capabilities are crucial for compiling vital datasets about the environment and human activities, helping us better understand dynamic environmental, ecological, and social systems (Gao et al., 2023). Simultaneously, advancements in AI technologies and the open science movement facilitate the development of new learning strategies.

Incorporating spatial characteristics into AI models is a key aspect of GeoAI (Goodchild, 2004). Geospatial data and processes have unique properties, such as spatial autocorrelation and spatial heterogeneity. Spatial autocorrelation, as described by Tobler's First Law of Geography, states that things geographically closer are more similar than distant things. That is, spatial autocorrelation measures the similarity of values for a given variable based on geographic proximity. This principle guides the design of spatial algorithms like spatial clustering and interpolation. Spatial heterogeneity, often called spatial non-stationarity, describes the variation in relationships, processes, or characteristics across different geographic locations (Anselin, 1989). It means that environmental, socio-economic, or ecological phenomena do not remain constant across space. Instead, these phenomena show differences that vary from location to location due to the influence of various local factors. Newer GeoAI models increasingly embed this property to address spatial heterogeneity (Georganos et al., 2021; Xie et al., 2021). For example, Xie et al. (2021) developed a spatial-heterogeneity-aware deep learning framework that uses a multivariate scan statistic to identify spatial

heterogeneity in input data. This approach guides hierarchical spatial partitioning of the study area to train multiple deep learning models, resulting in more robust predictions than other AI models.

A recent development in GeoAI is Prithvi, a geospatial foundation model developed by NASA and IBM (Jakubik et al., 2023; Li et al., 2024). Prithvi addresses large-scale image analysis challenges by leveraging massive geospatial datasets like Harmonized Landsat and Sentinel-2 (HLS) imagery. Unlike traditional models, Prithvi uses self-supervised learning, reducing the need for labeled data and enabling more scalable analyses. Its applications include crop mapping, flood mapping, and wildfire prediction. Prithvi's significance lies in its ability to create an efficient and adaptable framework for processing large-scale Earth observation data. This transformer-based geospatial foundation model reduces the need for extensively labeled data, thereby accelerating solutions in areas with scarce labeled geospatial data (Jakubik et al., 2023). However, Prithvi's generalized nature can struggle with local-specific phenomena without further fine-tuning. For example, accurately classifying crops in small agricultural fields in sub-Saharan Africa requires specialized adjustments due to unique regional variations and field characteristics. The model's lack of transparency, often called the "black-box" nature, complicates validation and stakeholder trust. Furthermore, the high computational resources required for training and processing large-scale data limit accessibility, particularly for institutions with restricted budgets in developing countries.

Another notable GeoAI tool that is gaining attention is GeoSAM. GeoSAM is an advanced adaptation of the Segment Anything Model (SAM), designed explicitly for multi-class segmentation of aerial or satellite imagery (Sultan et al., 2024). GeoSAM automates the segmentation process by leveraging sparse and dense prompts, significantly improving performance over traditional CNN and ViT-based models. Sultan et al., 2024 tested the model using high-resolution orthorectified imagery from Washington DC and Cambridge (MA). Their results showed a 26% improvement in road infrastructure segmentation and a 7% improvement in pedestrian infrastructure segmentation. The authors recommend future research to enhance further GeoSAM's scalability and applicability across various geospatial tasks and regions. While GeoSAM surpasses traditional models, it struggles with accurately segmenting narrow features like roads and sidewalks, especially when their textures resemble other elements in the imagery. Additionally, interference from surrounding objects, such as trees, buildings, and vehicles, can hinder the model's ability to generate accurate segmentation maps. The model's performance declines when applied to images from different regions compared to those used during training. The model's "black-box" nature also complicates validation and stakeholder trust. Moreover, its reliance on the large-scale computational resources needed for training and inference may limit its accessibility for specific users or applications.

While GeoAI offers significant strengths—such as handling large, multimodal datasets and producing detailed geospatial insights—it also faces challenges. These challenges include interpretability and explainability issues inherent in the machine and deep learning models, difficulties with reproducibility across different geographical areas, and potential ethical concerns related to privacy and fairness (Li et al., 2024). As GeoAI becomes instrumental in processing and analyzing geospatial data, researchers can tackle complex environmental issues at scales and speeds beyond human capability. However, this increased capability brings a heightened need for transparency. Therefore,

stakeholders must understand the reasoning behind the decisions made by these powerful GeoAI models.

1.2 Why Explainability Matters in GeoAI

The high performance of ML and DL models often leads them to be seen as opaque or "black box" systems due to their complexity. For example, random forests and neural networks are typical opaque ML and DL models. Explainable AI (XAI) techniques attempt to transform these "black box" models into more transparent "glass box" models by providing insights into why the predictors influence the model outcome. The importance of explainability in GeoAI is amplified in environmental modeling, where outcomes can influence policy and resource management decisions at large scales. Trust in these AI models is crucial to enable ethical deployment in public sectors—such as environmental management and public health—and ensure that decisions are based on sound and explainable logic. Explainable GeoAI enhances transparency and accountability by allowing stakeholders to understand the rationale behind GeoAI outputs. This is crucial for advancing GeoAI applications in land cover mapping, forest carbon stock monitoring, disaster preparedness, climate change, and other scenarios where actionable insights must be accurate and defensible.

Explainability is especially crucial in geospatial applications due to the distinctive nature of spatial data, which inherently involves location and context. Unlike traditional datasets, geospatial data often depends on spatial relationships and geographical complexities such as variations in terrain, climate, etc. Therefore, ML and DL models applied to analyze geospatial data must consider these spatial dynamics. Consequently, explainability in GeoAI models ensures that decisions reflect real-world spatial intricacies. For instance, understanding how predictors affect the model output can help make the model's predictions more informative for stakeholders (scientists, decision-makers, and policymakers).

1.3 XAI Techniques

1.3.1 Overview of machine and deep learning techniques

GeoAI has leveraged various AI techniques such as random forests, support vector machines (SVMs), and convolutional neural networks (CNNs). Researchers have successfully used these methods to process geospatial data and extract meaningful information and insights. Random forests (Rodriguez-Galiano et al., 2012) and SVMs (Huang et al., 2002; Pal & Mather., 2005) are extensively used for land cover classification due to their strong ability to model nonlinear relationships. These models classify satellite imagery into land cover types based on spectral and textural features. However, one limitation of these ML techniques is that they require extensive feature engineering to optimize model performance. Feature engineering often requires domain expertise to manually craft features that represent the underlying patterns in the data, which can be time-consuming and error-prone.

On the other hand, convolutional neural networks (CNNs) have gained widespread use for extracting spatial features from high-resolution satellite and aerial imagery (Yang et al., 2018). Unlike traditional ML techniques, CNNs eliminate the need for feature engineering entirely, making them more efficient for complex image segmentation and classification. CNNs automatically learn spatial hierarchies of features from raw input data, making it possible to capture patterns and relationships that might otherwise be challenging to design manually. This capability has made CNNs particularly valuable in applications like mapping building footprints, identifying tree species, and extracting other landscape features. Their proficiency in automatically detecting features has facilitated urban planning, infrastructure monitoring, and environmental management, where detailed and fine-resolution maps are critical.

1.3.2 Overview of XAI techniques

Researchers have developed several XAI techniques to address the 'black box' and trust challenges of AI models. These XAI techniques can generally be categorized into global and local explanation methods. Global explanation methods provide insights into the overall behavior of the model by analyzing its entire dataset. These methods include feature importance scores, partial dependence plots, and global surrogate models. Global explanation methods offer valuable insights into overall model behavior, such as identifying which predictors are most influential in making decisions. For instance, in a canopy height regression model, global methods can determine predictor variables (Sentinel-2 bands, spectral indices) influencing the model outcome (Figure 1.1).




Figure 1.1. Example of global SHAP feature importance.

In contrast, local explanation methods aim to provide an understanding of individual predictions. These methods focus on explaining the outcome for a specific instance, making them ideal for interpreting particular cases. For example, local explanation techniques such as SHAP (SHapley Additive Explanations) or LIME (Local Interpretable Model-agnostic Explanations) can help determine why a specific area was classified as high risk for flooding based on factors like elevation, slope, land cover, or soil texture (Figure 1.2). These local insights are particularly useful for stakeholders making decisions about specific locations. For example, emergency planners must understand the basis for classifying an area susceptible to natural disasters. Adopting XAI techniques in GeoAI applications is critical for fostering trust and ensuring accountability in decision-making processes. By providing insights into the model predictors, these techniques allow stakeholders to better understand model predictions.

Moreover, XAI techniques also contribute to the iterative refinement of models. Developers can identify potential biases, refine feature selection, and enhance model accuracy by understanding which predictor variables drive predictions. This capability is critical in geospatial contexts where diverse factors—such as climate, topography, and human activities—interact in complex ways. Thus, explainability supports informed decision-making and plays a vital role in improving the robustness and reliability of GeoAI models.




Figure 1.2. LIME plot showing predictors that influence flood susceptibility mapping for a specific instance.

1.3.3 XAI techniques for GeoAI models

Researchers have developed various XAI techniques, some of which have been adopted for GeoAI applications. These XAI techniques offer different explanations depending on the model and the specific geospatial application involved. In this section, we will briefly review popular XAI techniques such as SHapley Additive exPlanations (SHAP), Local Interpretable Model-agnostic Explanations (LIME), decision tree surrogate models, and Gradient-weighted Class Activation Mapping (Grad-CAM).

SHapley Additive exPlanations (SHAP) is one of the most widely used XAI techniques and originates from cooperative game theory (Shapley, 1953). SHAP values provide local and global explanations by attributing each feature's contribution to the model's prediction. In GeoAI, SHAP can help explain which predictor variables, such as land cover type or elevation, influence the model's predictions for flood risk mapping or deforestation modeling. One of SHAP's key strengths is its consistency and completeness. This strength guarantees that the sum of the feature contributions matches the model's output, ensuring an accurate breakdown of how each feature affects the prediction. SHAP is also model-agnostic, meaning it can be applied across various machine and deep learning models, making it versatile for diverse GeoAI applications.

Local Interpretable Model-agnostic Explanations (LIME) is another popular XAI technique for explaining individual predictions. Generally, LIME approximates complex models with simpler, more interpretable models. LIME works by perturbing the input data and observing the resulting changes in the model's output, creating a local surrogate model that is easier to understand compared to the original complex model. In GeoAI models, LIME can be used to explain how different predictor variables, such as elevation or soil moisture, influence the model's decision in a specific instance. This makes LIME a powerful tool for understanding individual, localized predictions. However, a limitation of LIME is that its explanations are guaranteed to be valid only locally. As a result, they may not generalize the entire model well. This interpretation could lead to misunderstandings if interpreted beyond their intended scope.

Decision tree surrogate models (Figure 1.3) are another XAI technique that can approximate complex models with simpler, more interpretable decision tree models. These surrogate models are built to

mimic the behavior of a more complex model, providing an overall understanding of how the model makes its decisions. In GeoAI applications, decision tree surrogates can offer insights into how predictor variables like Sentinel-2 bands, spectral indices, and elevation influence predictions at a broader level. However, like LIME, surrogate models may oversimplify complex relationships and thus might not capture all the details of the original model's behavior.

Gradient-weighted Class Activation Mapping (Grad-CAM) is a visual explanation technique specifically designed for deep learning models (Selvaraju et al., 2017), particularly convolutional neural networks (CNNs). Grad-CAM produces heatmaps that visually depict which regions of an input image were most influential in the model's decision-making process. For example, Grad-CAM can highlight the parts of a satellite image that contributed the most to classifying a region as urban, water, or forest. These visual explanations are highly valuable for domain experts and stakeholders, providing intuitive insights into why the model made a particular classification. Therefore, visualization instills greater confidence in the decision-making process. However, Grad-CAM is limited to CNNs and does not apply to other DL models.

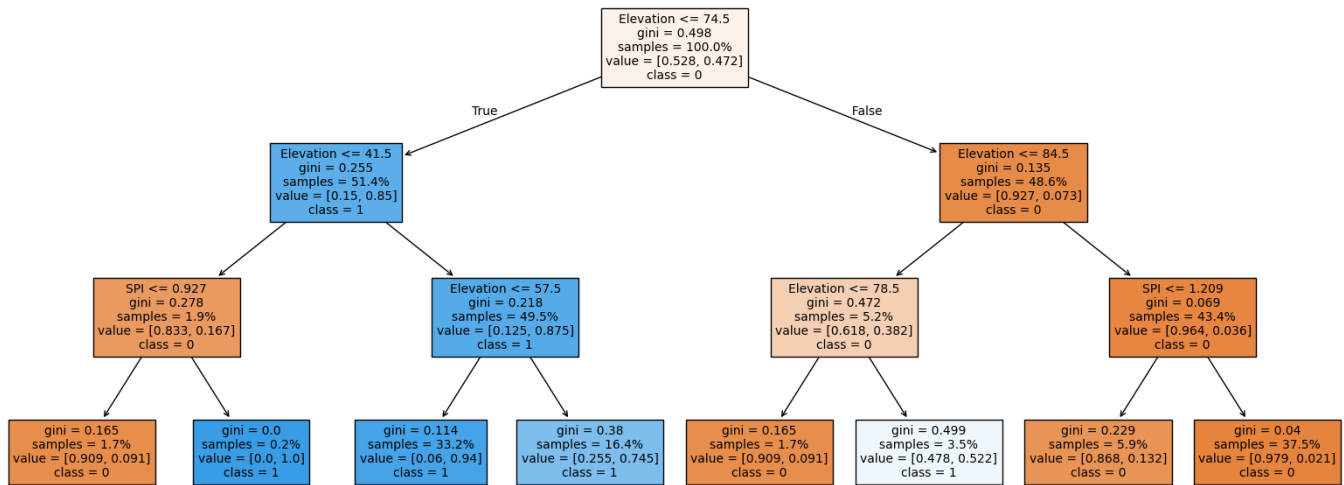


Figure 1.3. Example of a decision tree surrogate model.

1.4 Applications of Explainable GeoAI for Environmental Modeling

Explainable GeoAI holds immense potential to enhance the reliability and transparency of GeoAI models in environmental modeling. By increasing the explainability of these models, explainable GeoAI ensures that the insights produced are valid, understandable, and accessible to stakeholders. The use of GeoAI fosters trust and facilitates informed decision-making. This book will explore two key case studies highlighting the benefits of explainable GeoAI by modeling forest structure and mapping flood inundation.

The first case study focuses on modeling forest structure using random forests in conjunction with Earth observation (EO) data (spaceborne LiDAR, Sentinel-2 imagery). The objective is to model canopy height—a critical metric for monitoring biodiversity, carbon sequestration, and overall forest health. By applying SHAP, we aim to identify which predictor variables significantly influence canopy height predictions. This level of insight is particularly valuable for scientists aiming to enhance canopy height prediction models. It also benefits other stakeholders who seek to understand how different predictors contribute to the model's outcome.

The second case study will utilize a U-Net model to map flood inundation using high-resolution imagery. U-Net is well-suited for this type of task because of its ability to segment complex spatial patterns. We will employ Grad-CAM to generate heatmaps to make the U-Net model more explainable. These

heatmaps will visually illustrate which regions of the input imagery were most influential in the model's decision-making process. For example, we will analyze areas that were classified as inundated. Grad-CAM will help stakeholders, such as disaster response teams and urban planners, better understand the basis of the model's predictions.

These two case studies demonstrate the power of explainable GeoAI in providing meaningful transparency to environmental modeling. By applying explainable methods such as SHAP and Grad-CAM, we can uncover the reasoning behind model predictions, which is especially crucial for tasks that directly impact environmental policy and management. Explainable GeoAI helps bridge the gap between GeoAI methods and practical insights for stakeholders.

1.5 Summary

GeoAI integrates AI with spatial data, science, and geospatial technologies to solve spatial problems and enhance understanding. This field has grown rapidly due to advances in AI technologies, particularly machine learning (ML) and deep learning (DL), and the increasing availability of geospatial data. GeoAI is applied in diverse areas, including land cover mapping, deforestation monitoring, urban planning, and climate modeling. However, due to the inherent complexity of ML and DL models, they are often seen as "black boxes," which can hinder their acceptance, especially in sensitive fields like environmental management and public health. This creates a critical need for explainable AI (XAI) techniques to improve model transparency and stakeholder trust.

Explainable GeoAI uses XAI techniques such as SHAP, LIME, decision tree surrogate models, and Grad-CAM to provide explanations of global and local model behavior. These methods help stakeholders understand what a model predicts and why it makes those predictions. SHAP, for example, attributes feature contributions to a model's prediction, while LIME approximates complex models with simpler ones to explain individual predictions. Grad-CAM is particularly useful in visualizing which regions of input imagery influenced a model's classification decisions, making it ideal for DL models like convolutional neural networks (CNNs).

Explainable GeoAI has significant applications in environmental modeling, where transparency is crucial for policy-making and resource management. For instance, modeling forest structures using random forests combined with SHAP can help identify the key factors influencing canopy height, thus aiding conservation planning. Similarly, mapping flood inundation using a U-Net model along with Grad-CAM helps visualize which parts of imagery influenced the classification of flooded areas, supporting informed disaster management. Explainable GeoAI fosters greater trust in advanced machine and deep learning by offering clear insights into AI-driven geospatial models. The inclusion of explainability into GeoAI models makes it a powerful tool for sustainable environmental management.

Chapter 2. Case Study 1: Explainable Machine Learning for Modeling Forest Structure

2.1 Overview

This chapter employs a random forest approach to model forest canopy height based on NASA's Global Ecosystem Dynamics Investigation (GEDI) observations and Sentinel-2 satellite data. The GEDI's Level 2A Geolocated Elevation and Height Metrics Product (GEDI02_A) and Sentinel-2 data are available in Google Earth Engine (GEE). GEDI Level 2A dataset offers various metrics related to vegetation vertical structure, such as canopy height (Dubayah et al., 2020). This dataset includes relative height (RH) metrics, which represent the height at a given energy quantile (ranging from the 1st to the 100th) relative to the elevation of the lowest waveform mode, indicating ground level (Hofton & Blair, 2019; Potapov et al., 2021). In this book, we utilize RH98 to represent top canopy height, as it has been considered a more robust and less noise-sensitive metric than RH100 (Milenković et al., 2022; Li et al., 2023). Figure 2.1 shows GEDI Level 2A footprint data overlaid on Sentinel-2 imagery.




Figure 2.1. GEDI Level 2 footprint data (red lines) overlaid on Sentinel-2 imagery in the Mavuradonha Wilderness area, Zimbabwe.

This chapter uses eight Sentinel-2 spectral bands, the normalized difference vegetation index (NDVI), and the canopy chlorophyll content index (CCCI). These spectral indices—widely utilized in remote sensing—are instrumental in monitoring vegetation health, vigor, and productivity. NDVI, in particular, is one of the most commonly used indicators for assessing vegetation health. It measures the difference between the near-infrared (NIR) and red bands. Higher NDVI values typically indicate healthy, green vegetation since healthy plants strongly reflect NIR while absorbing red light. Conversely, lower NDVI values suggest stressed vegetation or land with no plant cover.

CCCI, on the other hand, is a more specialized index for evaluating chlorophyll content within plant canopies. The chlorophyll content is a critical indicator for assessing vegetation health and productivity. One of the key strengths of CCCI is its incorporation of the red edge band, which is sensitive to variations in chlorophyll content. Therefore, CCCI can reveal finer variations in chlorophyll concentration, providing a more detailed understanding of vegetation conditions. The CCCI is calculated as follows:

$$\text{CCCI} = \frac{[(\text{NIR} - \text{Red Edge}) / (\text{NIR} + \text{Red Edge})]}{[(\text{NIR} - \text{Red}) / (\text{NIR} + \text{Red})]}.$$

2.2 Exploratory Data Analysis

2.2.1 GEDI canopy height data

Figure 2.2 presents box and density plots, showing the GEDI Level 2A canopy height distribution within the Mavuradonha Wilderness area. Miombo woodlands and grasslands primarily cover the protected area. However, recent increases in human activities—such as tree cutting, farming, and mining—pose significant threats to the integrity of this ecosystem. The box plot shows that the median canopy height falls between approximately 7 m and 15 m, consistent with the expected canopy height of Miombo woodlands, typically ranging from 7 m to 20 m. The interquartile range (IQR), which represents the middle 50% of canopy height data, also falls within this expected range, indicating that a significant portion of the canopy height remains typical of the undisturbed Miombo woodlands. However, the box plot also reveals outliers extending beyond 25 m. These outliers may not accurately represent actual canopy height and could be attributed to measurement errors. Due to the high elevation and steep slopes of the Mavuradonha Wilderness area, the GEDI sensor might have overestimated the canopy height, leading to the presence of these outliers.

The density plot further reveals that the canopy height distribution is right-skewed, with many canopy heights between 5 m and 15 m. The density peak at lower canopy heights suggests that the region is experiencing some level of forest degradation, particularly in areas where human activities have encroached on natural forested zones. The density plot also shows a long tail extending beyond 30 m, though these taller canopy heights are less frequent. This suggests an artifact of measurement inaccuracies on high terrain.




Figure 2.2. Distribution of GEDI canopy height in the study area.

In summary, the plots suggest that while much of the canopy height distribution remains within the expected range for Miombo woodlands, there is clear evidence of disturbance, likely due to human encroachment and farming activities. Additionally, outliers and the long tail observed in the plots may be attributed to overestimation by the GEDI sensor, particularly in areas of high elevation and steep slopes. These observations highlight the need to interpret LiDAR data in such terrains carefully.

2.2.2 Spectral reflectance analysis using boxplots

Figure 2.3 shows the distribution of spectral reflectance values for eight Sentinel-2 spectral bands—specifically, B3 (Green), B4 (Red), B5 (Vegetation Red Edge 1), B6 (Vegetation Red Edge 2), B7 (Vegetation Red Edge 3), B8 (Near Infrared), B11 (Shortwave Infrared 1), and B12 (Shortwave Infrared 2)—along with CCCI and NDVI. Each boxplot presents key statistical measures such as median, interquartile range, and outliers, providing insights into the variability of spectral reflectance

across different wavelengths. The median values and the presence of outliers are significant when interpreting these distributions.

For bands 3 to 5, the median spectral reflectance is generally below 0.4, which is consistent with the typical reflectance of vegetation in these parts of the electromagnetic spectrum. Specifically, B3 (Green) and B4 (Red) are lower, likely due to their sensitivity to chlorophyll absorption. This results in lower reflectance for healthy vegetation. The boxplots also show that these bands have numerous outliers, indicating a certain degree of heterogeneity in spectral reflectance. Band 5 (VRE1) also has a median value below 0.4, which is slightly higher than bands 3 and 4. The higher median value in the red edge band is significant, as red edge bands are often used to detect changes in vegetation health and stress levels. The red edge bands located between the red and near-infrared bands show a rapid increase in reflectance caused by chlorophyll absorption and leaf structure reflection. This characteristic explains why B5 has a higher median than the B3 and B4 bands.




Figure 2.3. Box plot showing the distribution of Sentinel-2 spectral reflectance values.

Moving on to bands 6 (VRE2), 7 (VRE3), and 8 (NIR), there is a gradual increase in median spectral reflectance. Band 8 (NIR) shows slightly higher spectral reflectance values. This is typical for healthy vegetation, which strongly reflects near-infrared light, contributing to higher reflectance in this band. The median for band 8 is higher than 0.5, indicating strong NIR reflectance. Bands 6 and 7 also show moderate increases in median reflectance compared to bands 3 to 5, reflecting their role in capturing subtle variations in chlorophyll content. Bands 11 (SWIR1) and 12 (SWIR1) display lower median spectral reflectance than the NIR band. These bands are sensitive to moisture content in vegetation and soil, and their lower reflectance suggests varying moisture levels in the landscape.

Finally, the NDVI and CCCI display higher median values. NDVI shows a high median value, indicating overall healthy vegetation within the study area. The CCCI, which incorporates the red edge, also presents a high median value and appears to capture more subtle variations in chlorophyll content. Generally, NDVI and CCCI correlate to vegetation canopy cover and biomass more than the Red or NIR bands alone.

2.2.3 Spectral reflectance analysis using density plots

Figure 2.4 shows the distribution of spectral reflectance values for eight Sentinel-2 spectral bands. We observe that the spectral reflectance is mostly concentrated between 0.15 and 0.3 for B3 and between 0.1 and 0.35 for B4, with peaks at approximately 0.2. This is typical for vegetation, which reflects moderate amounts of green light while absorbing a significant amount of red light. The narrow

peaks and the relatively lower maximum reflectance values indicate that vegetation cover dominates these observations.

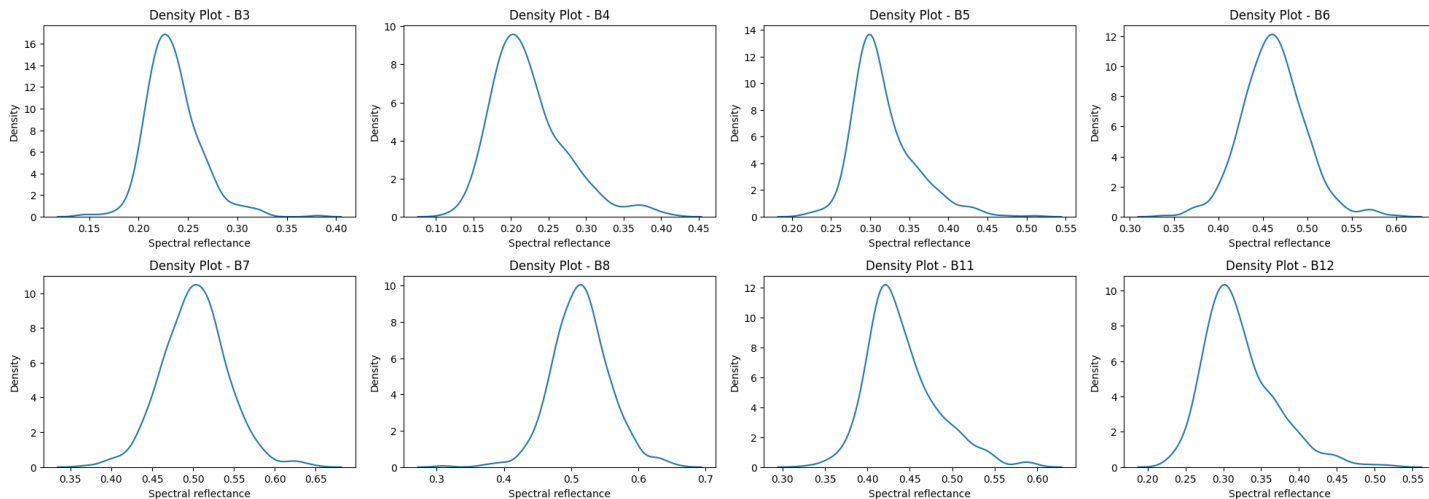


Figure 2.4. Density plot showing the distribution of spectral reflectance values for Sentinel-2 bands.

We observe a shift towards higher reflectance values for bands 5, 6, and 7. Specifically, B5 (VRE1) has a peak around 0.3, while bands 6 (VRE2) and 7 (VRE3) have higher peaks near 0.45 and 0.5, respectively. The red edge bands are sensitive to changes in chlorophyll content and provide crucial information on plant health and stress. The higher and broader reflectance distributions indicate the characteristic response of healthy vegetation in the red edge region, reflecting more as chlorophyll content increases.

Band 8 (NIR) shows a clear peak at around 0.5, indicating strong reflectance in this range. Healthy vegetation strongly reflects NIR light due to internal leaf structures, so the peak is significantly higher than the visible bands. The higher density at these reflectance levels signifies the presence of healthy, green vegetation in the observed data.

For bands 11 (SWIR1) and 12 (SWIR2), the peaks are observed at approximately 0.45 and around 0.3, respectively. SWIR bands are sensitive to water content in vegetation and soil. Therefore, the lower spectral reflectance values are consistent with SWIR's role in detecting moisture levels. Band 11 shows a relatively high peak density, suggesting moderate reflectance indicative of moisture variation across vegetation and soil surfaces. In contrast, band 12 shows a broader distribution that captures a range of reflectance values possibly influenced by variations in water stress and soil conditions. Overall, these density plots provide insights into how different wavelengths interact with vegetation and soil. Each band offers unique information for monitoring plant health, stress, and moisture content.

2.3 Evaluating Model Performance

We employed a cross-validation approach with 10 folds ($cv=10$) to assess the random forest regression model's prediction accuracy. Figure 2.5 illustrates that the RMSE scores range between approximately 5 m and 6 m across different iterations, highlighting model performance variability. Specifically, some iterations exhibit relatively low RMSE scores, around 5 m, which indicates better prediction accuracy and lower model errors. Conversely, other iterations demonstrate higher RMSE values, approximately 6 m, suggesting that the model experienced more significant prediction errors in those cases.

The lowest RMSE values for models 6 and 7 indicate that these particular iterations provided more reliable and accurate predictions. On the other hand, the highest RMSE values are recorded for models 3 and 9, highlighting instances where the model's performance was suboptimal with higher error rates. This variability indicates that the model's predictive performance changes across iterations. Data characteristics, model settings, and variability in training and testing splits during cross-validation contribute to this inconsistency.




Figure 2.5. RMSE scores for random forest regression model across 10-fold cross-validation.

Further model refinement is recommended to achieve a more consistent and lower RMSE across all iterations. This inconsistency points to a need for enhancements such as better hyperparameter tuning to optimize model parameters or more effective data preprocessing to ensure robust performance. Addressing these aspects could help stabilize the model's prediction accuracy, thereby achieving more reliable and generalized results across all folds of the cross-validation process. Ultimately, the goal would be to reduce variability and consistently achieve lower RMSE values, indicating improved overall model performance.

Figure 2.6 shows a scatter plot depicting the relationship between the GEDI reference canopy height on the x-axis and the estimated canopy height on the y-axis. The regression line, characterized by $y=0.04x+11.29$, indicates a weak positive linear relationship between the reference and estimated canopy heights. The plot reveals several model prediction problems. Note that the data points are widely scattered around the regression line, pointing to substantial inconsistency in the model's predictions.

Another critical observation is the bias present in the estimation. The intercept value of 11.29 implies that the model predicts heights of around 11.29 m, even when the reference height approaches zero. This consistent overestimation is particularly evident for lower reference canopy heights, suggesting that the model systematically overpredicts the height when reference values are low. This bias could lead to inaccurate conclusions. Li et al. (2023) observed that GEDI performed well for trees between 3 m and 15 meters in height in African Savannah woodlands. However, GEDI was less reliable for canopy heights below 2.34 m due to limitations in its pulse width, making it unsuitable for estimating shrub heights (Li et al., 2023).

The influence of outliers (see Figure 2.2) is evident in the scatter plot, as they contribute to a weak correlation and distort the accuracy of the random forest model's predictions. The increased spread of data points, especially at the extremes of the reference height range, indicates that these outliers have negatively affected the model's ability to generalize. As a result, the model fails to provide accurate estimates across different height values. Therefore, data cleaning through outlier detection and removal should be done to improve the model's performance. In particular, it is critical to analyze and clean GEDI canopy heights, especially in areas with high elevations or steep slopes, which are known to introduce inaccuracies in LiDAR measurements. In addition, it is also important to adjust the model parameters to minimize the impact of outliers.




Figure 2.6. Relationship between the GEDI reference canopy height and the estimated canopy height.

The color bar in the figure represents the absolute difference between the reference and estimated canopy heights, providing further insight into the accuracy of the model's predictions. The darker (purple) points indicate lower absolute differences, signifying more accurate predictions, whereas the lighter (yellow) points represent larger discrepancies between reference and estimated canopy heights. The lighter points are more common at higher reference canopy heights, suggesting the model struggles to accurately predict taller trees. This issue is particularly concerning because it implies that the model's performance declines as canopy height increases, potentially limiting its effectiveness in estimating the heights of mature forests.

Finally, the distribution of the data points reveals a pattern in the spread of the predictions. Most estimated heights are clustered between 10 m and 15 m, irrespective of the reference canopy height. This clustering highlights the model's tendency to predict within a narrow range, reflecting an inability to adequately capture the variability in high reference canopy heights. This limitation further contributes to the poor performance of the model, especially for taller trees or more complex canopy structures.

The model generally exhibits a weak linear relationship between reference and estimated canopy heights, consistent overestimation at lower reference values, and significant inaccuracy at higher reference heights. The color-coded absolute differences emphasize the model's struggle to predict high canopy height accurately. Furthermore, the clustering of predictions within a narrow range demonstrates the model's limited responsiveness to variations in the reference data. Addressing these issues requires data cleaning, extensive model tuning, and additional predictors to better capture the underlying complexity in the canopy height data. Figure 2.7 shows the predicted canopy height.




Figure 2.7. Predicted canopy height.

2.4 Interpreting SHAP Values

2.4.1 Interpreting global SHAP values

2.4.4.1 Feature importance analysis using SHAP

Figure 2.8 shows the feature importance plot, which uses global SHAP values to explain the contribution of each Sentinel-2 band and spectral indices to the model's predictions. The plot shows the mean absolute SHAP values for each predictor, representing each predictor's average impact on the model's output. This type of analysis is valuable in understanding which predictor variables are most influential in determining the model's predictions, thus providing insight into the predictor contributions within the regression model.

Figure 2.8 shows that the CCCI has the highest mean SHAP value, which indicates that it has the most significant influence on the model's predictions. This suggests that CCCI plays a vital role in explaining the variation in the target variable. Its effectiveness is likely due to its sensitivity to vegetation health and chlorophyll levels. Following CCCI, NDVI is the second most impactful predictor, though its influence is notably less compared to CCCI.

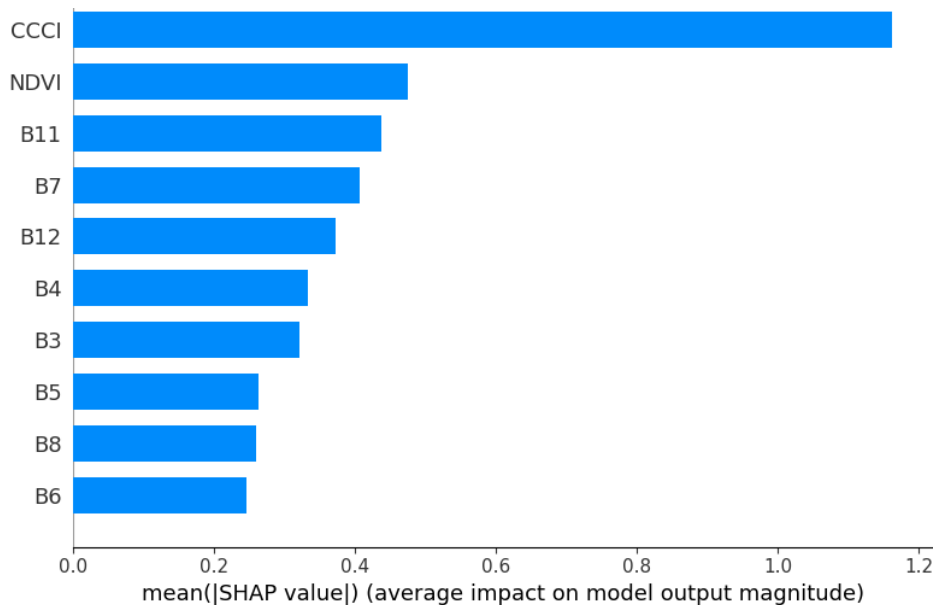


Figure 2.8. Mean SHAP values representing feature importance for Sentinel-2 bands and spectral indices.

Among the Sentinel-2 spectral bands, bands 11 (SWIR 1), 7 (VRE 3), and 12 (SWIR 2) are the most impactful, with relatively high mean SHAP values. This highlights the importance of these bands in

capturing information related to soil and vegetation moisture content and the subtle variations in chlorophyll along the red-edge region. In contrast, bands 6 (VRE 2) and 8 (NIR) have the least influence on the model's predictions, as indicated by their lower mean SHAP values. This suggests that these bands contribute less to the model's overall explanatory power.

The feature importance plot, derived from SHAP values, helps identify the most critical predictors that drive the model's accuracy, providing valuable information for understanding model behavior. However, it is important to note that this analysis does not provide information about the direction of the influence—whether each predictor contributes positively or negatively to the predictions. Instead, it solely provides the magnitude of each predictor's impact. Therefore, further analysis is required to understand how changes in these features affect the model predictions specifically.

2.4.4.2 Analysis of SHAP value distribution using beeswarm plot

Figure 2.9 shows the beeswarm plot derived using SHAP values. The beeswarm plot shows the underlying distribution of data while avoiding overlap between individual points. The beeswarm plot provides a detailed view of how individual predictor contributions vary across all predictions made by the model. Each dot on the beeswarm plot represents a single prediction, with the dot color indicating the feature value—red for high feature values and blue for low feature values. This visualization allows us to observe the influence of each predictor on the model's output, revealing both the magnitude and the direction of their impact.

The canopy chlorophyll content index (CCCI) has the most significant impact on the model's output. High values of CCCI, represented by red dots, generally push the model's predictions positively, suggesting that increased chlorophyll content leads to higher predictions for the target variable. Conversely, low CCCI values, shown in blue, tend to push the predictions downwards, indicating a negative impact on the model outcome. This highlights CCCI as a strong driver of model predictions, with distinct and consistent patterns depending on the value of this feature.




Figure 2.9. SHAP beeswarm plot illustrating feature value impact on model predictions.

A similar pattern is observed for NDVI and the other Sentinel-2 bands, although their overall impact is less pronounced compared to CCCI. High values of these predictors also contribute positively, while low values contribute negatively to the predictions. However, the spread and density of SHAP values

are generally smaller. This suggests that while NDVI and individual spectral bands influence model predictions, their impact is not as strong or consistent as that of CCCI.

The spread of SHAP values for each predictor indicates the variability they introduce into the model's predictions. A wide spread of SHAP values means that the predictor contributes with a high degree of variability, influencing the prediction significantly in both positive and negative directions, depending on the feature value. For CCCI, the wide spread of SHAP values emphasizes its role in contributing substantial variability to the model, making it a critical feature for determining the predicted outcomes. In contrast, predictors with a narrower spread, such as some of the spectral bands, indicate a more moderate and consistent effect on model predictions.

Overall, the beeswarm plot provides an insightful visual representation of the distribution and influence of SHAP values across predictors. This helps to understand the relative importance and dynamic nature of predictor contributions in the model. It reveals that while multiple predictors contribute to the predictions, CCCI has the strongest influence with a distinct and consistent effect, whereas other features like NDVI and spectral bands have a less pronounced but still relevant impact on the model's predictions.

2.4.4.3 Analysis of SHAP value distributions using violin plots

Figure 2.10 represents the distribution of each predictor's impact on the model's output. Each "violin" shows the range and concentration of SHAP values for a particular predictor. The violin plot provides insights into the magnitude and variability of a predictor's influence on model predictions. The width of each violin plot at different points along the SHAP value axis indicates the density of predictions with similar SHAP values. Therefore, the violin plot shows the predictor's contribution to the model's performance.

Figure 2.10 shows that the CCCI has the widest distribution of SHAP values. This suggests that CCCI introduces the greatest variability in influencing the model's output, with SHAP values ranging widely from negative to positive impacts. This variability reinforces CCCI as a significant driver of the model's predictions, where high feature values can strongly increase predictions while low values can equally decrease them. The broad and varied distribution emphasizes that CCCI is crucial in shaping model behavior, reflecting positive and negative contributions across the dataset.

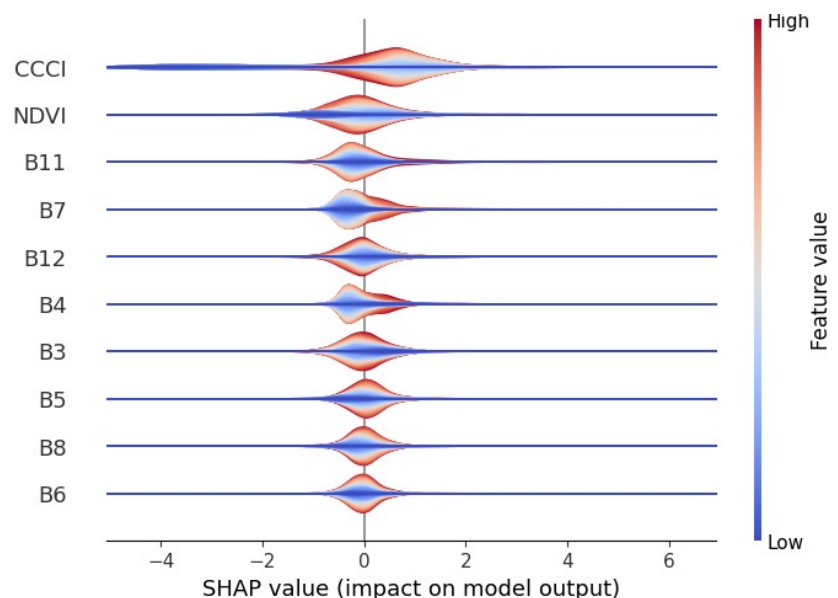


Figure 2.10. Violin plot of SHAP value distributions for Sentinel-2 bands and spectral indices.

For bands 11 (SWIR 1), 7 (VRE 3), and 12 (SWIR 2), the violin plots show moderate spread in their SHAP value distributions. This indicates that while these bands notably impact the model predictions, their influence is less dominant or variable than CCCI's. The narrower distribution also suggests moderate contributions. Therefore, the bands do not swing the model output as significantly as CCCI.

In contrast, features like bands 6 (VRE 2) and 8 (NIR) show a much narrower distribution of SHAP values, indicating a limited role in the model's prediction process. The narrow shape of these violins suggests that these features contribute minimally to the overall predictions, with most of their SHAP values clustered around zero. This limited spread implies that changes in these features do not substantially influence the model output.

Overall, the violin plot helps us understand which features are important and how their contributions vary across all predictions. The broader distributions, such as seen with CCCI, indicate predictors with more dynamic and influential impacts. In contrast, the narrower distributions reflect predictors with more limited, consistent roles. This interpretation aids in refining feature selection by identifying the predictor that contributes most variably and significantly to the model's output. This information can be leveraged for model improvements or further analysis.

2.4.4.3 Scatter plot analysis of SHAP value distributions

Figure 2.11 provides a detailed examination of how the predictor values relate to their respective SHAP values. This analysis offers insights into the contribution of each predictor to the model's predictions. The scatter plots help visualize whether a predictor positively or negatively impacts the model's output and how this impact changes across different predictor values.

For bands 3, 4, 5, and 8, the SHAP values are mainly concentrated between -0.5 and 0.5, with a slight spread to both positive and negative sides. This suggests that these predictors have a relatively weak influence on the model's output. The absence of a clear upward or downward trend in the SHAP values across the range of predictor values indicates that the influence of these bands is relatively consistent, with minimal impact on driving the model predictions in a specific direction. Higher or lower reflectance values in these bands do not markedly shift the model predictions, highlighting their limited role.

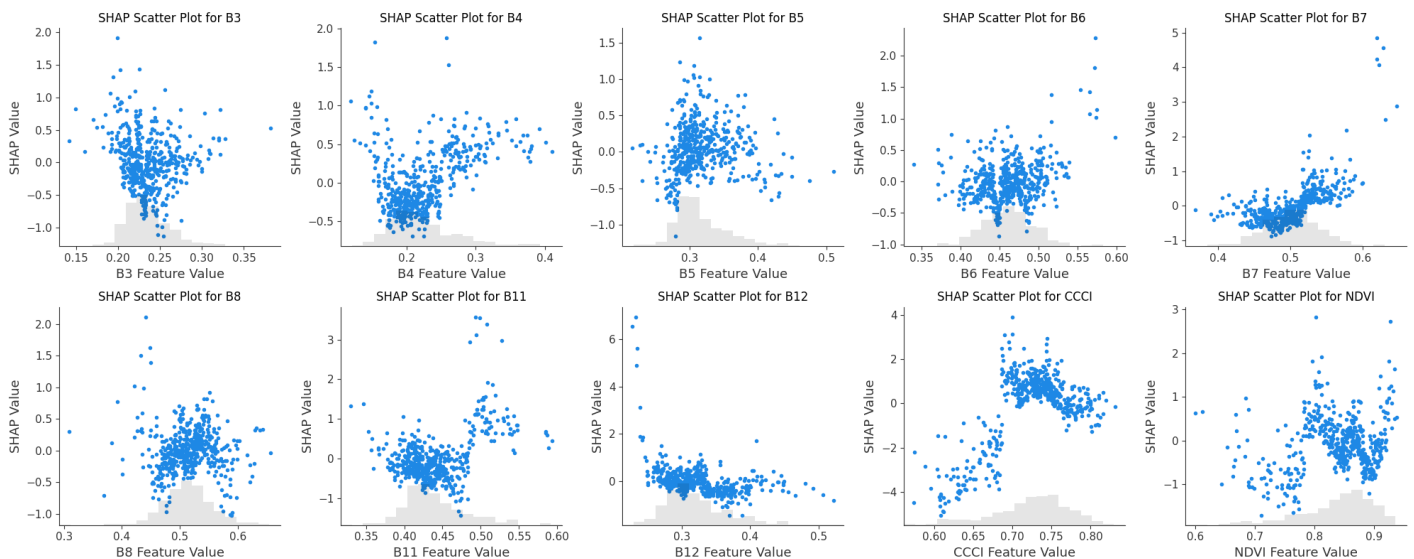


Figure 2.11. Scatter plots of SHAP values vs. predictor values for Sentinel-2 bands and spectral indices.

Bands B7 and B11 show a distinct positive correlation between predictor and SHAP values. This means that higher reflectance in these bands tends to have a stronger positive impact on the model's predictions. Band 12 (SWIR2) also shows a positive relationship between spectral reflectance and SHAP values. However, this trend is less linear than B7 and B11, implying a more complex influence.

The canopy chlorophyll content index (CCCI) has a distinctive pattern, with a broad range of SHAP values. High CCCI values consistently contribute positively to the model's predictions, pushing the SHAP values upward. Conversely, lower CCCI values have a negative impact. This strong separation between high and low feature values emphasizes that CCCI plays a critical role in the model, significantly influencing the prediction direction and magnitude. This observation aligns with earlier analyses, where CCCI was identified as a major driver of model predictions.

The normalized difference vegetation index (NDVI) also displays a relatively strong positive relationship with SHAP values. Higher NDVI values correspond to higher positive SHAP values, indicating a substantial positive influence on model outputs. This reflects NDVI's importance in the model. Although the spread of SHAP values for NDVI is more compact than for CCCI, it still significantly impacts prediction accuracy, underscoring its importance in capturing vegetation characteristics.

Overall, CCCI has the strongest and most consistent impact on the model's predictions, with higher values resulting in higher positive SHAP values. This suggests that CCCI is a crucial predictor in the model, significantly influencing the output. NDVI, along with bands B7 and B11, also shows a clear positive contribution, meaning higher values of these features push the model predictions positively. In contrast, bands B3 to B6 exhibit weaker effects, with their SHAP values mostly centered around zero, suggesting a moderate and less directional influence on the model's behavior.

2.4.2 Interpreting Local SHAP Values

2.4.2.1 Local SHAP value analysis

Figure 2.12 presents a bar plot of local SHAP values. The bar plot provides a detailed interpretation of how each predictor contributes to the model's prediction for a specific observation. This local interpretation helps us understand how each predictor's value increases or decreases the model output for an individual prediction. The analysis gives insights into a single data point's result.




Figure 2.12. Local SHAP values show feature contributions to model prediction for a single observation.

Figure 2.12 shows that B7 (VRE 3) makes the largest positive contribution to the model's prediction, adding approximately +1.02 to the output. This indicates that, for this specific observation, B7 significantly drives the prediction upward. Closely following B7 is the CCCI, which also adds a considerable amount to the model's output, with a contribution of +0.86. The substantial positive contributions from B7 and CCCI indicate that high values of these features are likely pushing the prediction upward, making them dominant factors in determining the output for this instance.

In contrast, other predictors, such as B11 and B8, contribute negatively to the prediction. Specifically, B11 reduces the model output by -0.46, while B8 has a smaller negative impact of -0.17. This suggests that the values of B11 and B8 for this particular observation are counteracting the positive contributions from other predictors, thereby pulling down the prediction. These negative contributions are important as they show that not all predictors support the prediction in the same direction. Rather, they provide opposing influences based on their individual values.

Additionally, other predictors, including B6, NDVI, B12, and B3, contribute positively to a smaller extent. Their SHAP values are much lower compared to B7 and CCCI, suggesting a more moderate influence on the final prediction. Band 6 adds +0.22, NDVI adds +0.12, and other features have contributions under +0.1, implying that these features support the prediction. Therefore, their influence is relatively minor compared to the dominant effects of B7 and CCCI.

Overall, the interpretation of this local SHAP plot reveals that B7 and CCCI are the primary drivers of the prediction for this specific observation, substantially boosting the model's output. Meanwhile, B11 and B8 serve as counterbalances, reducing the output value. This localized view of feature contributions provides useful insights into why the model made a particular prediction. It also highlights the interplay between predictors that contribute positively and those that oppose the prediction.

2.4.2.2 Local SHAP waterfall plot

Figure 2.13 shows a detailed breakdown of how the model arrived at a final prediction of 13.596 for a specific observation. The analysis shows the cumulative contributions of each predictor relative to the base value, which is the average model output of 11.762. Each step in the waterfall plot shows how much a particular predictor is either added to or subtracted from the base value. This provides insights into the model's decision-making process for this particular instance.

The waterfall plot shows that B7 (VRE3) makes the largest positive contribution to the prediction, adding +1.02 to the final value. This indicates that B7 is the most influential feature driving the prediction upward, highlighting its importance for this observation. CCCI is the next most impactful predictor, contributing an additional +0.86 to the prediction, reinforcing its critical role in increasing the model's output for this observation. These contributions from B7 and CCCI show that higher values for these predictors strongly impact the prediction.

In contrast, B11 (SWIR1) negatively influences the prediction, decreasing the value by -0.46. This indicates that B11 is counteracting some of the positive contributions from other predictors for this specific instance, thus reducing the model output. This negative contribution suggests that the conditions associated with B11 for this observation are less favorable, effectively pulling the prediction down.

Other predictors, such as B6 (VRE2), NDVI, and B8 (NIR), contribute more modestly to the prediction, positively or negatively. For instance, B6 adds +0.22 to the model output, while NDVI adds +0.12. Conversely, B8 contributes negatively, subtracting -0.17 from the prediction. These smaller

contributions reflect the fine-tuning of the prediction. This indicates that while these predictors play a role, their influence is less significant than B7 and CCCI.




Figure 2.13. SHAP waterfall plot demonstrating predictor contributions to a specific prediction.

Ultimately, the final prediction of 13.596 is reached by summing these individual contributions from the base value of 11.762. The SHAP waterfall plot demonstrates how each predictor, positively or negatively, impacts the model output for this particular observation. This analysis offers an intuitive way to understand the model's decision-making process for an individual prediction.

2.4 Spatial Analysis of Predictor Contributions using SHAP

Figure 2.14 shows how CCCI, NDVI, B7 (VRE3), and B8 (NIR) contribute to the model's prediction. Each subset of the SHAP map shows the localized influence of these predictors, with a color gradient ranging from red to blue. The color range reflects the magnitude and direction of the predictor's impact on the model's predictions. Red areas indicate positive contributions, suggesting that these predictor values contribute to higher canopy height estimates at those locations. Conversely, blue areas indicate negative contributions, meaning the predictors are associated with lower canopy height predictions.

Figure 2.14 shows that CCCI and NDVI are the most influential predictors for estimating forest canopy height. The SHAP maps for these predictors show larger red areas, suggesting that high chlorophyll content and healthy vegetation are positively associated with increased canopy height. This pattern confirms that regions with higher CCCI and NDVI values will likely have healthy vegetation, which strongly drives the model's canopy height prediction. The positive contributions of these spectral indices are crucial for accurately representing vegetative growth dynamics.

On the other hand, bands 8 (NIR) and 7 (VRE 3) show fewer areas with high influence, as evidenced by their limited patches of red in the SHAP maps. The impact of these bands on the prediction of canopy height is less pronounced compared to CCCI and NDVI. While they still play a role, the spatial extent of their contribution is less widespread. Therefore, these bands do not consistently drive the

prediction of forest canopy height across the study area. This may suggest that their relevance is more context-specific.

The spatial visualization offered by these SHAP maps is invaluable for understanding the localized contribution of each predictor in the random forest model. The SHAP maps provide essential insights to stakeholders such as forest managers by highlighting the most important predictor variables for predicting forest canopy height. In other words, the SHAP maps show where the influence of predictors is most significant. Moreover, these SHAP maps are instrumental for researchers and model developers, providing critical feedback on model performance given its spatial nature. Understanding how features contribute across different regions can help refine model training and improve spatial accuracy, potentially leading to better predictive power and generalization. While the benefits of these visualizations are significant, it is worth noting that generating SHAP maps is computationally intensive. This particular analysis took approximately three hours to compute the SHAP maps for a subset of the study area. Thus, there is a pressing need to optimize the computation of spatial SHAP values to make this valuable tool more accessible for broader applications without the substantial computational burden.

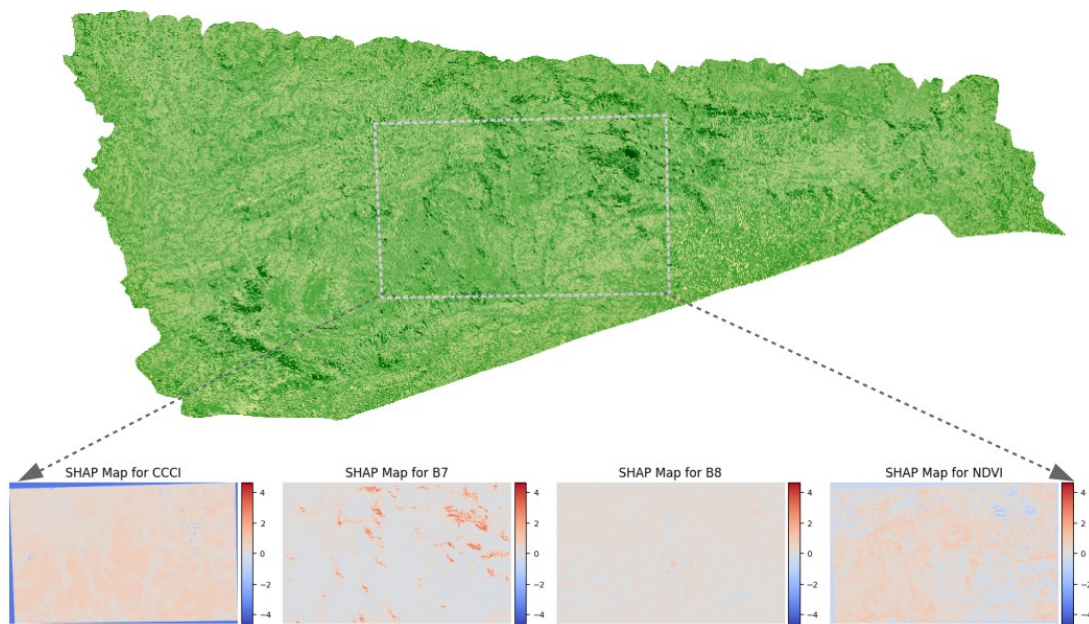


Figure 2.14. A subset of SHAP Maps showing contributions of predictors to forest canopy height predictions.

2.5 Summary

Chapter 2 explored a case study on explainable machine learning techniques for modeling forest structures. The chapter utilized a random forest model. GEDI canopy height data, eight Sentinel-2 bands, and two spectral indices (NDVI and CCCI). These spectral indices are essential in assessing vegetation health and productivity. We also performed exploratory data analysis using box and density plots. The exploratory data analysis using boxplots and density plots highlighted significant insights, which are critical before training a machine learning model.

We also evaluated the random forest regression model's performance using a 10-fold cross-validation approach. The variability in RMSE scores across different iterations highlighted the need for further model improvements. In particular, data cleaning is required, given the outliers present in the training data. The model's predictions of canopy height were compared with GEDI reference canopy heights, revealing a weak linear relationship and consistent overestimation. This indicates the model's

challenges in accurately capturing the full variability of forest canopy heights, especially for taller trees. The chapter emphasized the importance of improving model accuracy to enhance the reliability of forest height predictions for environmental applications.

Finally, we used the SHAP values to understand the contributions of different predictors to the model's output. Both global and local SHAP value analyses were performed, providing insights into which predictors most influenced the model's predictions. CCCI and NDVI emerged as key drivers, with a significant and consistent positive influence on canopy height predictions. SHAP value visualizations, including bar plots, beeswarm plots, scatter plots, violin plots, and waterfall plots, highlighted the critical features and their influence across individual and global contexts. Moreover, spatial SHAP maps provided a valuable visualization of how these features contribute across different regions, offering stakeholders a clearer understanding of which predictors are most important and where their influence is strongest. The chapter concluded by discussing the computational intensity of generating SHAP maps, suggesting a need for optimization to make this tool more practical for large-scale spatial analyses.

Chapter 3. Case Study 2: Explainable Deep Learning for Flood Inundation Mapping

3.1 Introduction

Flooding remains among the most destructive and widespread natural hazards worldwide, leading to significant human, economic, and environmental losses (Bentivoglio et al., 2022). Consequently, mapping flood inundation is critical for effective disaster management, urban planning, and mitigation strategies. Accurate flood maps enable policymakers and emergency responders to develop informed action plans, prioritize high-risk areas, and reduce the long-term impacts of flood events. As climate change continues to intensify both the frequency and severity of floods, the need for accurate flood mapping has become even more pressing to protect lives, infrastructure, and ecosystems (Lee & Li, 2024).

Over the past several decades, researchers have increasingly turned to Earth observation data to enhance flood inundation mapping capabilities. Earth observation data, particularly satellite imagery, offer a unique ability to capture large-scale spatial dynamics of flood events, making it a valuable resource for flood risk assessment. However, traditional approaches to flood inundation mapping have faced several limitations. One of the primary challenges was the reliance on low to medium-spatial resolution imagery, which often lacked the spatial detail needed to accurately represent flood impacts on critical infrastructure such as buildings and roads (Bentivoglio et al., 2022). Additionally, traditional processing techniques for high-resolution imagery were computationally intensive and relied heavily on manual interpretation. This made the process laborious and time-consuming, leading to significant delays in generating flood maps during emergencies when timely information is crucial (Sanderson et al., 2023). These methods' computational demands and time constraints also limited their scalability, making them less effective for responding to large-scale or fast-evolving flood events (Lin et al., 2023).

Recent advancements in deep learning (DL), particularly developing and applying convolutional neural networks (CNNs), have significantly transformed mapping using aerial or high-resolution satellite imagery. CNNs are particularly well-suited for processing high-resolution imagery, as they excel at recognizing spatial patterns, making them highly effective for detecting flood extents (Bentivoglio et al., 2022). These DL models reduce the manual effort required for feature engineering. They can efficiently process large volumes of satellite imagery, generating accurate flood maps in near real-time. Furthermore, the fusion of multiple data sources, such as synthetic aperture radar (SAR) and optical imagery, has enabled models to overcome data quality issues like cloud cover (Sanderson et al., 2023).

In addition to improving accuracy and efficiency, DL models have introduced enhanced scalability to flood mapping. Advanced CNN architectures, such as U-Net, have demonstrated their capacity to handle large datasets and produce high-resolution flood inundation maps across diverse geographical regions (Lin et al., 2023). These capabilities are critical for ensuring rapid response to flood events on both local and regional scales. Therefore, DL models are an indispensable tool for flood risk management.

Despite these advancements, challenges remain. One persistent issue is model generalization—ensuring that a model trained in one region or on one type of flood event can effectively generalize to different regions or flood scenarios. Flood events can vary greatly in nature, influenced by diverse

climatic, topographic, and hydrological conditions. This variability makes developing a single model that performs consistently across all scenarios difficult. Moreover, incorporating uncertainty into predictions remains a critical area of research, particularly as decision-makers need a clear understanding of the reliability of model outputs to make informed decisions during emergencies (Bentivoglio et al., 2022). Additionally, there is a growing need for explainability in DL models. Explainability is crucial for improving model trust and transparency. For example, scientists and policymakers can understand the predictors influencing the model outcome.

In this chapter, we address these gaps using free data from Kaggle (Figure 3.1) to map flood inundations using DL and XAI techniques. Specifically, we utilize a U-Net model for automated feature engineering and mapping. Furthermore, we employ Gradient-weighted Class Activation Mapping (Grad-CAM) to visualize the areas of input imagery most influential in the model's decision-making process. Grad-CAM provides an essential step towards explainable DL, allowing stakeholders to better interpret and understand model predictions. This will lead to more informed and transparent flood disaster management and mitigation strategies.




Figure 3.1. Sample flood inundation data set. Note that flooded areas are in white, while non-flooded areas are in black.

3.2 U-Net Model

The U-Net model, an innovative U-shaped encoder-decoder convolutional neural network (CNN) architecture, was originally designed for biomedical image segmentation (Ronneberger et al., 2015). This study trained the U-Net model using the Kaggle dataset, carefully tuning key parameters and hyperparameters to optimize performance. These adjustments included fine-tuning the initial learning rate, batch size, and number of training epochs, alongside selecting appropriate elements such as the loss function, optimization algorithm, and evaluation metrics. Data augmentation techniques were also applied to enhance the diversity of the training set and improve the model's robustness, ultimately leading to better generalization in flood inundation mapping tasks.

Figure 3.2 provides insight into the training and validation performance of the U-Net model over 70 epochs, using both accuracy and loss as metrics to evaluate its effectiveness. The model accuracy

plot illustrates the progression of training and validation accuracy over time. Initially, training accuracy (blue line) starts at approximately 0.6 and steadily increases, reaching around 0.87 by the final epoch. This consistent rise in accuracy suggests that the model effectively learns to classify the input data. Meanwhile, the validation accuracy (orange line) exhibits some fluctuations during the early epochs, which is typical as the model adapts to different data sets. However, it gradually stabilizes, converging to around 0.86. The validation accuracy closely follows the training accuracy, and both curves converge towards the end of training. This convergence indicates that the model generalizes well without significant overfitting and performs similarly on both seen (training) and unseen (validation) data.




Figure 3.2. U-Net model results.

The model loss plot, on the other hand, tracks the evolution of loss throughout the training process. The training loss (blue line) decreases sharply during the initial epochs, reflecting the model's ability to learn from the training data quickly. As training continues, the loss declines more gradually but consistently, indicating steady improvements in the model's predictive capability. The validation loss (orange line) follows a similar trajectory, decreasing substantially at the beginning and remaining close to the training loss across the epochs. The convergence of training and validation loss near similar values further suggests that the model is balanced and generalizes well, successfully capturing the underlying patterns in the data without becoming overly tailored to the training set.

The U-Net model demonstrates strong performance across both metrics. The convergence of the training and validation curves for accuracy and loss without significant divergence indicates effective learning and good generalization. The absence of any significant gap between the training and validation curves confirms that the model is neither underfitting nor overfitting, achieving steady improvements in its capability to predict flood inundation effectively.

3.3 Flood Inundation Maps

Figure 3.3 shows the input image, the true mask, and the predicted flood mask. The input image displays a flooded area with water covering large landscape parts. The true mask, which serves as the ground truth, correctly identifies the flooded regions in white and the non-flooded regions in black. This mask provides a reference for evaluating the accuracy of the U-Net model's segmentation.

The predicted mask, generated by the model, captures the general structure of the flooded areas and is reasonably good at matching the true mask. However, there are some discrepancies, particularly in the upper portion of the image, where the model overestimates the extent of flooding, resulting in

extra white areas that were not flooded. Additionally, some smaller flooded regions, especially in the center and lower parts of the image, are missed or inaccurately segmented by the model.

In general, the U-Net model has performed well in identifying the larger flood areas. However, there are areas where the segmentation is less accurate, particularly in over-predicting flooded regions and missing finer details. Furthermore, model tuning may be necessary to improve the precision of the predictions, reduce false positives, and improve the delineation of smaller flood zones.



Figure 3.3. U-Net model input, mask, and predicted flood inundation.

3.4 Interpreting Grad-CAM

Figure 3.4 presents a series of images to visualize and interpret the Grad-CAM results, including the original flood images (left column), Grad-CAM heatmaps (middle column), and Grad-CAM overlays (right column). The Grad-CAM visualizations provide essential insights into which areas of the input images the DL model focused on when predicting flood zones. These heatmaps highlight critical regions that the model considered for its decisions, enhancing the transparency of its inner workings and understanding the rationale behind its predictions.

In the Grad-CAM heatmaps (middle column), the bright red and yellow regions signify areas where the model placed significant attention to make its predictions. In contrast, blue regions indicate areas of minimal or no importance in the decision-making process. For instance, in the first row (Image 5), the model strongly emphasizes the flooded streets in the middle and bottom right, as highlighted by the intense red regions. This suggests that these areas played a crucial role in the model's ability to detect flooding in the image. Similarly, in the second and third rows (Images 1022 and 3045), the heatmaps focus on the areas most affected by flooding, such as inundated structures, waterlogged vegetation, and large expanses of water. The intensity of these areas demonstrates the model's focus on the most impacted sections.

The Grad-CAM overlay images (right column) provide a merged view of the original images and the heatmaps, offering a more comprehensive understanding of the model's focus within the scene. In the first row (Image 5), the overlay indicates that the model concentrated primarily on the flooded streets and areas near buildings, especially around waterlogged zones. Similarly, in the second row (Image 1022), the model highlights areas of inundation near structures and vegetation, suggesting a strong association between these features and flood prediction. The third row (Image 3045) reveals that the model focused extensively on large flooded areas, as seen by the pronounced red regions overlaying the wide water bodies. These overlays are instrumental in visually assessing the relevance

and accuracy of the model's focus, helping to verify whether the model is paying attention to the correct regions, i.e., the areas affected by flooding.

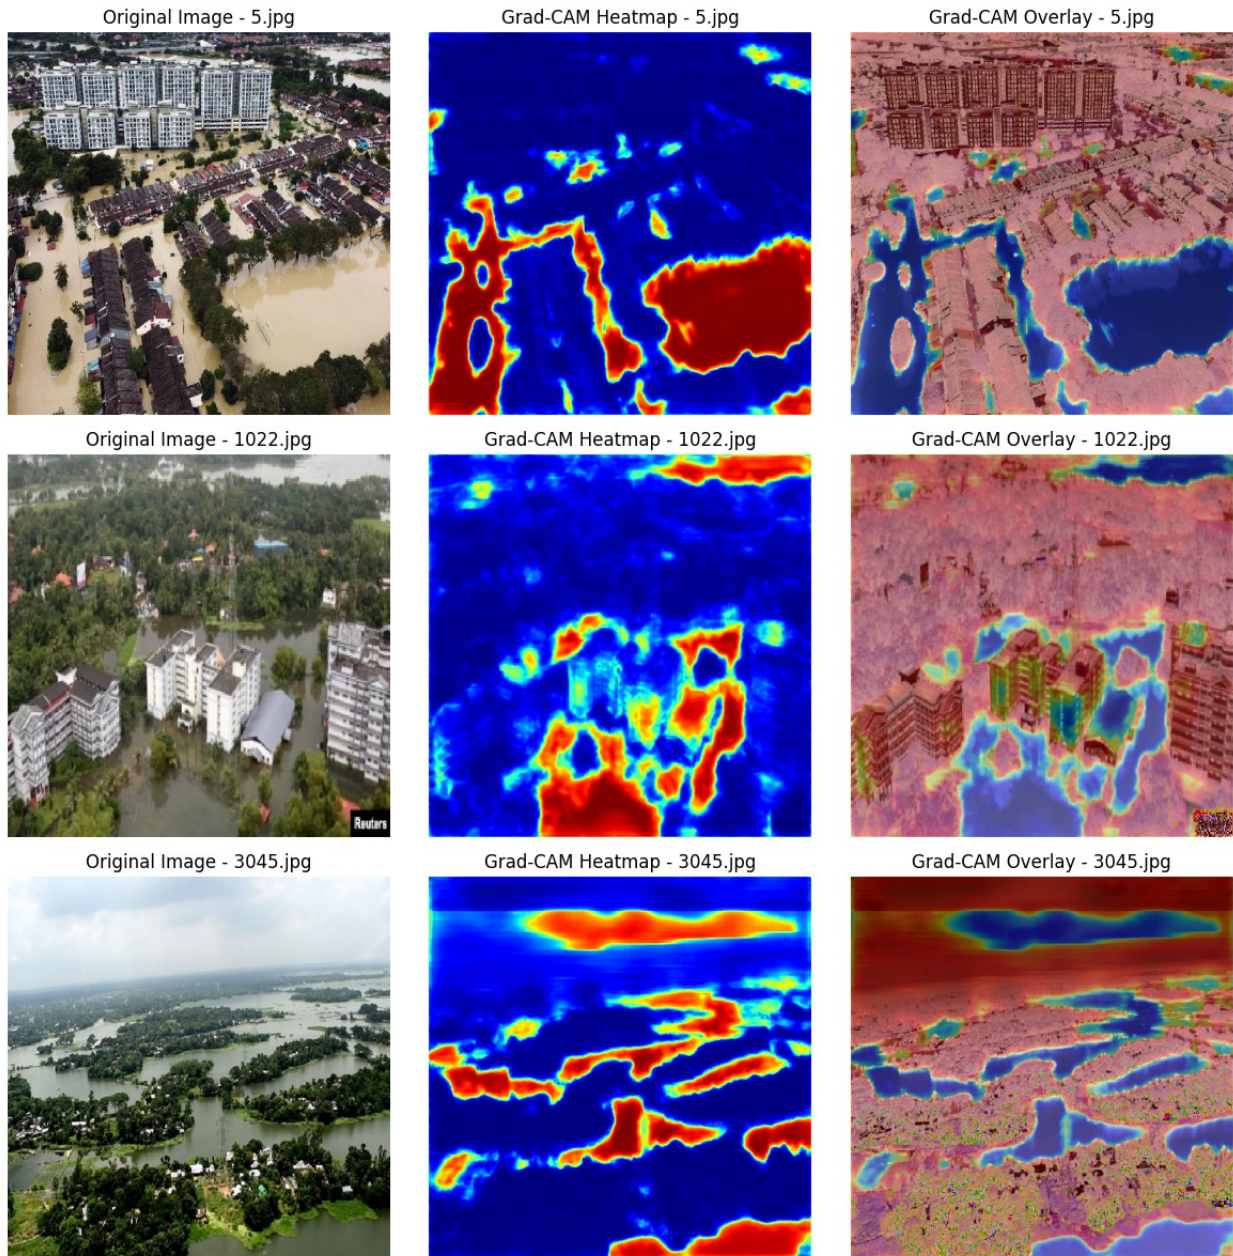


Figure 3.4. Original image, Grad-CAM image, and Grad-CAM overlay.

Generally, the Grad-CAM results demonstrate that the model effectively identified the flooded regions, focusing on key areas such as water bodies and inundated structures. The heatmaps provide insights into the model's decisions, revealing which parts of the image most influenced the model's predictions. These visualizations play an important role in building trust in the model's outputs by making the DL process more transparent and understandable. This is particularly critical in high-stakes applications like flood mapping, where accurate and reliable results are essential for decision-making.

3.5 Summary

Chapter 3 explores explainable DL for flood inundation mapping, emphasizing the utility of the U-Net model and Grad-CAM. Flood mapping is essential for disaster preparedness, response, and mitigation, particularly as the impacts of climate change intensify the frequency and severity of such events globally. While helpful, traditional methods for mapping floods were often constrained by low resolution, slow processing times, and labor-intensive manual interpretation. Recent advances in DL, specifically CNNs, have enabled more efficient and accurate flood mapping at finer resolutions, overcoming many of these traditional challenges.

In this case study, the U-Net model was used to detect and map flooded regions. The U-Net model achieved promising results by training the model on a flood dataset obtained from Kaggle and optimizing hyperparameters such as learning rate, batch size, and epochs. The model's accuracy steadily improved over 70 training epochs, ultimately reaching a near-optimal value of approximately 0.87 for both training and validation accuracy, with minimal evidence of overfitting. The loss curves for training and validation also converged, indicating effective generalization to unseen data. Despite its good overall performance, the U-Net model exhibited some limitations, particularly in identifying smaller or more complex flood areas. The segmentation results showed that the model overestimated the extent of flooding in certain sections while missing some smaller inundated regions, highlighting areas where further refinement is needed.

Grad-CAM was employed to help interpret the U-Net model's predictions by visually understanding which regions of the input images influenced the model's decision-making process. This technique generated heat maps showing areas of high model focus, such as water bodies, submerged infrastructure, and other inundated areas. The bright red and yellow areas on the heatmaps indicated the regions where the model concentrated the most when determining flood extents, giving insights into the decision-making logic behind each prediction. By overlaying these heatmaps onto the original flood images, the visualizations allowed stakeholders to assess the relevance of the model's focus. This helped confirm whether the model paid attention to the correct parts of the scene, ensuring transparency in its decision-making process.

This chapter underscores the transformative role of explainable DL models in flood inundation mapping. Combining U-Net for accurate segmentation and Grad-CAM for interpretability provides a powerful framework for creating reliable and explainable flood maps. This approach can significantly aid policymakers, emergency responders, and scientists by ensuring the model's predictions are accurate and interpretable. While the model showed strong results, further improvements, such as enhanced tuning and additional data augmentation, could help address challenges related to overestimation and missed details, particularly in more complex or localized flooding scenarios. The use of Grad-CAM also demonstrates the importance of transparency in DL, especially in critical applications like flood management.

Chapter 4. The Power of Explainable AI: Why It's the Future of Geospatial Modeling

4.1 Summary

The introductory book explores how AI and geospatial technology can be leveraged for environmental applications. Case studies illustrate the power of GeoAI for forest structure modeling and flood inundation mapping. The book highlights the strengths and limitations of advanced ML and DL techniques. It also showcases how explainable methods like SHAP, LIME, and Grad-CAM can bridge the gap between complex model predictions and explainability. These XAI techniques allow stakeholders to see the models' predictions and the predictors influencing the model outcome.

In critical fields such as environmental management, urban planning, and disaster response, where AI-driven decisions can have far-reaching implications, transparency is key to fostering trust and ensuring ethical, informed decision-making. This transparency is particularly important as these applications directly impact natural resources, communities, and ecosystems. Explainable GeoAI can mitigate the risks associated with "black box" models. Furthermore, it makes AI solutions more understandable to non-experts, including policymakers, community leaders, and other stakeholders.

This forward-looking book introduces the use of AI for geospatial modeling. It also provides a path toward more responsible and transparent applications of GeoAI in managing the Earth's natural resources and responding to environmental challenges. The book also encourages the integration of XAI as a standard practice in GeoAI workflows. This empowers users to leverage advanced GeoAI models in a way that upholds principles of transparency and accountability.

4.2 Strengths and Limitations of Explainable GeoAI for Modeling Canopy Height

Explainable AI (XAI) techniques are crucial in clarifying which spectral bands or indices (NDVI and CCCI) drive canopy height predictions. By providing a transparent breakdown of how these features influence model outputs, XAI ensures that models do not function as mere "black boxes." Instead, they produce meaningful and informative model insights. This transparency makes the predictions trustworthy and actionable for stakeholders. For example, forest managers need reliable insights for conservation planning and management strategies that support carbon sequestration.

Explainability also provides a deeper understanding of why certain areas are predicted to have high or low canopy height. For example, scientists or researchers can use these insights to improve the training data or model structure. XAI can enhance the applicability of random forest models at larger scales by illuminating the factors driving canopy height predictions. This scalability allows for more extensive use of data sources, such as Sentinel-2 satellite imagery, enabling regional or even global canopy height modeling.

However, while SHAP values and other feature attribution methods excel in identifying the importance of individual predictors, they often lack consideration of the spatial context—a critical component in geospatial applications. The relationships between predictor variables are often spatially dependent, meaning that predictors may not have the same impact across different locations. To address this shortcoming, this book introduces the concept of SHAP maps, which provide a spatial representation of predictor contributions. By integrating spatial context, SHAP maps enhance the interpretability of predictor contributions. This ensures that stakeholders understand the significance of predictors and where and why those predictors matter within the geographical landscape.

4.3 Strengths and Limitations of Explainable GeoAI for Mapping Flood Inundation

Grad-CAM has significantly enhanced DL methods such as U-Net by providing visual explanations in heatmaps. These heatmaps highlight regions of the input image that the model relied upon during decision-making, which can be extremely helpful in determining whether the model correctly identified flooded areas and focused on the appropriate parts of the satellite image. Such insights are essential for verifying that the model makes accurate and reliable predictions, particularly in critical applications like flood inundation mapping. When properly tuned, U-Net models often deliver high accuracy for flood segmentation tasks, distinguishing between flooded and non-flooded areas with minimal errors, even under real-world conditions.

Explainability further ensures that the model generalizes well across diverse areas by helping identify consistent flooding patterns across different regions. This makes DL models, like U-Net, particularly useful for large-scale disaster management, where models must adapt to various landscapes while maintaining robustness and reliability. Using XAI techniques provides a layer of trustworthiness for emergency responders and disaster management agencies. In addition, it serves as a validation mechanism, demonstrating that the model's focus is aligned with real-world flooding phenomena.

However, despite the capabilities of U-Net models, explainability techniques can sometimes reveal issues, such as the model overpredicting flooded areas in certain regions, leading to false alarms. Such overestimation is often a consequence of overfitting, particularly when models are trained on limited data or fail to capture the finer details of environmental variability, like subtle changes in elevation or localized vegetation. These false alarms can lead to inefficient use of resources during flood response and erode trust in AI-driven decision-making, especially if the model frequently misidentifies areas as flooded.

While Grad-CAM heatmaps provide valuable visual insights for deep learning practitioners, they can be challenging for non-experts to interpret. Stakeholders who lack extensive knowledge of geospatial analysis or AI may struggle to understand what these heat maps indicate about model performance or reliability. This complexity presents a significant limitation in practical applications, where XAI should ideally communicate insights that are accessible and actionable for a broad range of users. To maximize the utility of explainable GeoAI, it is crucial to develop simpler, more intuitive ways to convey these visual explanations to stakeholders. This simplification ensures that decision-makers without specialized technical expertise can understand and use critical insights.

4.4 Future Work

In the future, GeoAI will become more seamlessly integrated with explainability techniques. This will enable the development of refined and transparent geospatial models that are both powerful and explainable. Several aspects stand out as promising areas for future research and practical application.

First, integrating XAI techniques with core geographic principles such as spatial autocorrelation and spatial heterogeneity can provide insights into GeoAI model performance. Embedding such geographic principles into explainability methods will also enhance the explainability of GeoAI models. In addition, it will also contextualize XAI in a spatial framework, offering richer, more meaningful insights into how spatial relationships influence model outputs.

Another major avenue for future research lies in optimizing the efficiency of current explainability techniques for geospatial data. For example, computing spatial SHAP maps, as used in this book, is

computationally intensive and often time-prohibitive for large-scale applications. Therefore, developing more efficient algorithms for calculating spatial SHAP values is essential. Faster, more computationally feasible methods will be crucial for scaling up explainability efforts, particularly for regional or global high-resolution data applications.

Future research directions should also focus on developing inherently interpretable models. Instead of merely applying explainability techniques as a post hoc solution, creating interpretable models by design will foster greater transparency. Moreover, improving spatial generalization remains a critical research challenge, particularly ensuring that GeoAI models are reproducible across diverse geographical contexts. Enhanced spatial generalization would allow models trained in one region to be effectively applied in others, ensuring broader applicability and reducing biases that might arise from regional specificities.

Another emerging area of interest is enhancing the social responsibility of GeoAI models. As these models are increasingly used in urban planning, environmental monitoring, and disaster management, they can significantly influence public policies and resource allocation. Therefore, ensuring that GeoAI models are developed and deployed responsibly is paramount. This involves improving their technical robustness and embedding fairness, equity, and inclusivity principles into the modeling process to minimize unintended consequences and biases (Li et al., 2024).

Collaboration across disciplines—such as geography, computer science, environmental science, and social sciences—is essential to advancing these objectives. Embedding spatial thinking directly into AI models will ensure that GeoAI continues to evolve in a way that has lasting scientific and societal impacts. By prioritizing interdisciplinary research, we can develop GeoAI models that are not only technically advanced but also aligned with broader human and environmental values, making them more socially responsible and impactful.

4.5 Next Steps: How to Unlock the Full Potential of Explainable GeoAI

GeoAI is revolutionizing environmental mapping and modeling, but significant hurdles remain. As explored in this book, challenges like data quality, model accuracy, and overfitting are still prevalent. In the forest structure modeling example, it became clear that achieving robust accuracy requires meticulous data cleaning, model tuning, and a deeper understanding of the factors that drive predictions. Deep learning models, too, are prone to pitfalls like overfitting—problems that demand an intimate grasp of how models function and how predictor variables influence their outcomes. This is precisely where XAI steps in, bridging the gap between complex model behavior and meaningful, actionable insights. However, the journey does not end here—there is much more to uncover, understand, and refine.

Are you ready to take the next step and truly master Explainable GeoAI? I am thrilled to announce my upcoming ebook, 'Unlocking the Black Box: A Practical Guide to Explainable GeoAI for Environmental Modeling'. This comprehensive guide will take you beyond the basics and into the depths of applying explainable GeoAI techniques confidently. Whether you are just starting or aiming to elevate your expertise, this book is crafted to help you bridge the gap between understanding theory and achieving hands-on mastery. You will find in-depth tutorials, practical exercises, and ready-to-use code tailored to help you tackle real-world challenges in forest structure modeling, flood risk assessment, and other geospatial machine learning applications.

The future of GeoAI belongs to those who are not only willing to use it but can also explain it, trust it, and improve it. Do not get left behind, as this field evolves at breakneck speed. Imagine

understanding your model's predictions, justifying your results to stakeholders, and gaining an edge in leveraging GeoAI responsibly and effectively for environmental challenges.

This is your chance to unlock Explainable GeoAI's true potential—to understand your models and master them, improving accuracy, reliability, and trust in every output. Pre-order now and secure your copy with an exclusive early-bird discount. Sign up or email me today!

Elevate your skills. Expand your impact. Embrace the power of Explainable GeoAI.

References

- Bentivoglio, R., Isufi, E., Jonkman, S. N., & Taormina, R. (2022). Deep learning methods for flood mapping: A review of existing applications and future research directions. *Hydrology and Earth System Sciences*, 26(8), 4345–4378. <https://doi.org/10.5194/hess-26-4345-2022>
- Dubayah, R., Blair, J.B., Goetz, S., Fatoyinbo, L., Hansen, M., Healey, S., Hofton, M., Hurtt, G., Kellner, J., Luthcke, S., et al. (2020). The global ecosystem dynamics investigation: High-resolution laser ranging of the Earth's forests and topography. *Science Remote Sensing*, 1, 100002.
- Gao, S., Hu, Y., & Li, W. (Eds.). (2023). *Handbook of Geospatial Artificial Intelligence*. CRC Press. <https://doi.org/10.1201/9781003308423>
- Georganos, S., Grippa, T., Niang Gadiaga, A., Linard, C., Lennert, M., Vanhuyse, S., Mboga, N., Wolff, E., & Kalogirou, S. (2021). Geographical random forests: A spatial extension of the random forest algorithm to address spatial heterogeneity in remote sensing and population modelling. *Geocarto International*, 36(2), 121–136. <https://doi.org/10.1080/10106049.2019.1688980>
- Goodchild, M. F. (2004). GIScience, geography, form, and process. *Annals of the Association of American Geographers*, 94(4), 709–714. <https://doi.org/10.1111/j.1467-8306.2004.00426.x>
- Hofton, M., & Blair, J. B. (2019). Algorithm theoretical basis document (ATBD) for GEDI transmit and receive waveform processing for L1 and L2 products
- Hofton, M., & Blair, J. B. (2019). Algorithm Theoretical Basis Document (ATBD) for GEDI Transmit and receive waveform processing for L1 and L2 products.
- Huang, C., Davis, L. S., & Townshend, J. R. G. (2002). An assessment of support vector machines for land cover classification. *International Journal of Remote Sensing*, 23(4), 725–749. <https://doi.org/10.1080/0143116011004032>
- Jakubik, J., Roy, S., Phillips, C. E., Fraccaro, P., Godwin, D., Zadrozny, B., Szwarcman, D., Gomes, C., Nyirjesy, G., Edwards, B., Kimura, D., Simumba, N., Chu, L., Mukkavilli, S. K., Lambhate, D., Das, K., Bangalore, R., Oliveira, D., Muszynski, M., Ankur, K., Ramasubramanian, M., Gurung, I., Khallaghi, S., Li, H. (S.), Cecil, M., Ahmadi, M., Kordi, F., Alemohammad, H., Maskey, M., Ganti, R., Weldemariam, K., & Ramachandran, R. (2023). Foundation models for generalist geospatial artificial intelligence. Preprint available on arXiv. <https://doi.org/10.48550/arXiv.2310.18660>
- Lee, H., & Li, W. (2024). Improving interpretability of deep active learning for flood inundation mapping through class ambiguity indices using multi-spectral satellite imagery. *Journal Name*, Volume(Issue), Page Numbers
- Li, W., Arundel, S. T., Gao, S., Goodchild, M. F., Hu, Y., Wang, S., & Zipf, A. (2024). GeoAI for science and the science of GeoAI. *Journal of Spatial Information Science*, 29, 1–17. <https://doi.org/10.5311/JOSIS.2024.29.349>
- Li, X., Wessels, K., Armston, J., Hancock, S., Mathieu, R., Main, R., et al. (2023). First validation of GEDI canopy heights in African savannas. *Remote Sensing of Environment*, 285, 113402. <https://doi.org/10.1016/j.rse.2022.11>
- Lin, J., Sreng, C., Oare, E., & Batarseh, F. A. (2023). NeuralFlood: An AI-driven flood susceptibility index. *Frontiers in Water*, 5, 1291305. <https://doi.org/10.3389/frwa.2023.1291305>

- Liu, P., & Biljecki, F. (2022). A review of spatially-explicit GeoAI applications in urban geography. *International Journal of Applied Earth Observation and Geoinformation*, 112, 102936. <https://doi.org/10.1016/j.jag.2022.102936>
- Milenković, M., Reiche, J., Armston, J., Neuenschwander, A., De Keersmaecker, W., Herold, M., & Verbesselt, J. (2022). Assessing Amazon rainforest regrowth with GEDI and ICESat-2 data. *Science of Remote Sensing*, 5, 100051. <https://doi.org/10.1016/j.srs.2022.100051>
- Pal, M., & Mather, P. M. (2005). Support vector machines for classification in remote sensing. *International Journal of Remote Sensing*, 26(5), 1007–1011.
- Pierdicca, R., & Paolanti, M. (2022). GeoAI: A review of artificial intelligence approaches for the interpretation of complex geomatics data. *Geosci. Instrum. Method. Data Syst.*, 11(195–218). <https://doi.org/10.5194/gi-11-195-2022>
- Potapov, P., Li, X., Hernandez-Serna, A., Tyukavina, A., Hansen, M. C., Kommareddy, A., Pickens, A., Turubanova, S., Tang, H., Silva, C. E., Armston, J., Dubayah, R., Blair, J. B., & Hofton, M. (2021). Mapping global forest canopy height through integration of GEDI and Landsat data. *Remote Sensing of Environment*, 253, 112165. <https://doi.org/10.1016/j.rse.2020.112165>
- Rodriguez-Galiano, V. F., Chica-Olmo, M., Abarca-Hernandez, F., Atkinson, P. M., & Jeganathan, C. (2012). Random forest classification of Mediterranean land cover using multi-seasonal imagery and multi-seasonal texture. *Remote Sensing of Environment*, 121, 93–107. <https://doi.org/10.1016/j.rse.2011.12.021>
- Sanderson, J., Mao, H., Abdullah, M. A. M., Al-Nima, R. R. O., & Woo, W. L. (2023). Optimal fusion of multispectral optical and SAR images for flood inundation mapping through explainable deep learning. *Information*, 14(660). <https://doi.org/10.3390/info14120660>
- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision* (pp. 618–626)
- Shapley, L. S. (1953). A value for n-person games. In H. W. Kuhn & A. W. Tucker (Eds.), *Contributions to the Theory of Games II* (pp. 307–317). Princeton: Princeton University Press.
- Sultan, R. I., Li, C., Zhu, H., Khanduri, P., Brocanelli, M., & Zhu, D. (2024). GeoSAM: Fine-tuning SAM with sparse and dense visual prompting for automated segmentation of mobility infrastructure. *arXiv preprint arXiv:2311.11319v2*
- VoPham, T., Hart, J. E., Laden, F., & Chiang, Y. (2018). Emerging trends in geospatial artificial intelligence (geoAI): Potential applications for environmental epidemiology. *Environmental Health*, 17(40). <https://doi.org/10.1186/s12940-018-0386-x>
- Xie, Y., He, E., Jia, X., Bao, H., Zhou, X., Ghosh, R., & Ravirathinam, P. (2021). A statistically-guided deep network transformation and moderation framework for data with spatial heterogeneity. In *2021 IEEE International Conference on Data Mining (ICDM)* (pp. 767–776). IEEE. <https://doi.org/10.1109/ICDM50108.2021.00088>
- Yang, H. L., Yuan, J., Lunga, D., Laverdiere, M., Rose, A., & Bhaduri, B. (2018). Building extraction at scale using convolutional neural network: Mapping of the United States. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 11(8), 2600–2614. <https://doi.org/10.1109/JSTARS.2018.2835370>

Appendix

Here are some valuable resources and websites that can help deepen your understanding of Explainable GeoAI and XAI in general.

A. General XAI Resources

1. SHAP (SHapley Additive exPlanations)

SHAP is one of the most widely used tools for explaining machine learning models. It provides both global and local interpretability for model predictions.

Website: <https://shap.readthedocs.io/en/latest/>

2. LIME (Local Interpretable Model-agnostic Explanations)

LIME is a popular tool for understanding individual predictions made by any machine learning model.

Website: <https://github.com/marcotcr/lime>

3. Google's Explainable AI (XAI) Tools

Google provides Vertex explainable AI tools that can be used to interpret and explain ML models.

Website: <https://cloud.google.com/vertex-ai/docs/explainable-ai/overview>

B. GeoAI and Geospatial Modeling Resources

1. Prithvi - NASA & IBM Geospatial Foundation Model

Prithvi is a geospatial foundation model developed by NASA and IBM. It leverages Harmonized Landsat and Sentinel-2 (HLS) imagery and uses self-supervised learning for large-scale image analysis. Prithvi is particularly useful for applications like flood mapping, crop mapping, and wildfire prediction, and is designed to handle Earth observation datasets.

Website: <https://huggingface.co/ibm-nasa-geospatial/Prithvi-100M>

2. GeoSAM (Segment Anything Model for Geospatial Data)

GeoSAM is an advanced adaptation of the Segment Anything Model (SAM) developed for segmentation of satellite and aerial imagery. It uses sparse and dense prompts to automate segmentation tasks, significantly improving upon traditional CNN-based models for geospatial image analysis.

Website: <https://samgeo.gishub.org/>

3. Google Earth Engine

Google Earth Engine is one of the most widely used platforms for analyzing geospatial datasets. It supports large-scale geospatial data processing and includes tutorials for integrating machine learning models.

Website: <https://earthengine.google.com/>

C. GeoAI and XAI Research and Educational Resources

1. Explainable AI for Earth Observation (XAI4EO)

A platform that hosts research and tools related to explainable AI in Earth observation.

Website: <https://www.lancaster.ac.uk/scc/research/ai4eo/>

2. The Alan Turing Institute – Explainable AI Research

The Turing Institute offers insights and research publications on XAI, focusing on applications that include geospatial and environmental data.

Website: <https://www.turing.ac.uk/research/research-projects/project-explain>

D. Open Source Libraries for GeoAI and XAI

1. GeoPandas

A Python library for working with geospatial data, commonly used in conjunction with machine learning models for GeoAI applications.

Website: <https://geopandas.org/en/stable/>

2. PyTorch Geometric (PyG)

A library designed for graph-based deep learning, often used to analyze spatial data.

Website: <https://pytorch-geometric.readthedocs.io/en/latest/>

3. TensorFlow with Keras

TensorFlow's Keras API is frequently used for building and training deep learning models on geospatial data. You can apply XAI techniques like Grad-CAM directly in TensorFlow.

Website: <https://www.tensorflow.org/>

4. GeoShapley

A game theory approach to measuring spatial effects from machine learning models. GeoShapley is built on Shapley value and Kernel SHAP estimator.

Website: <https://github.com/Ziqi-Li/geoshapley>

E. Tutorials and Learning Platforms

1. GeoAI tutorials at Ai.Geolabs

Ai.Geolabs offers many free courses on GeoAI applications.

Website: <https://aigeolabs.com/courses/>

2. Google AI Education – Explainability

Google's dedicated portal for learning AI, which includes explainability tutorials and resources.

Website: <https://explainability.withgoogle.com/>

About the Author

Courage Kamusoko is an independent geospatial consultant based in Japan. His expertise includes land use/cover change modeling and the design and implementation of geospatial database management systems. Courage Kamusoko's primary research involves analyses of remotely sensed images, land use/cover modeling, modeling aboveground biomass, machine learning, and deep learning. In addition to his focus on geospatial research and consultancy, he has dedicated time to teaching practical machine learning for geospatial data analysis and modeling.

Books published:

1. Kamusoko, C. (In press). *Explainable Machine Learning for Geospatial Data Analysis: A Data-Centric Approach*. Taylor and Francis.
2. Kamusoko, C. (2021). *Optical and SAR Remote Sensing of Urban Areas: A Practical Guide*. Springer.
3. Kamusoko, C. (2019). *Remote Sensing Image Classification in R*. Springer.
4. Kamusoko, C. (2021). *Data-centric Explainable Machine Learning for Land Cover Classification: A Practical Guide in R*. Self-published (<https://aigeolabs.com/books/>).
5. Kamusoko, C. (2021). *Explainable Machine Learning for Land Cover Classification: An Introductory Guide*. Self-published (<https://aigeolabs.com/books/>).

Books Edited:

1. Murayama, Y., Kamusoko, C., Yamashita, A., Estoque, R.C. eds. (2017). *Urban Development in Asia and Africa: Geospatial Analysis of Metropolises*. Dordrecht: Springer, 424p.
2. Kamusoko, C., Mundia, C. N. and Murayama, Y. eds. (2011). *Recent Advances in Remote Sensing and GIS in Sub-Sahara Africa*. New York: Nova Publishers, 211p.

Contribute to the Book!!

Thank you for reading my book on 'GeoAI Unveiled: Case Studies in Explainable GeoAI for Environmental Modeling.' The book is under continuous development. I will improve the text and add more chapters. If you find a mistake or something is missing, please email me at cou.kamusoko@aigeolabs.com.