



Università degli Studi di Napoli Federico II
Ph.D. Program in
Information Technology and Electrical Engineering
XXXVI Cycle

THESIS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

Empowering Sustainability through Artificial Intelligence and Earth Observation: Learning from the Past, Monitoring the Present, Building the Future

by
GIOVANNI GIACCO

Advisor: Prof. Carlo Sansone



SCUOLA POLITECNICA E DELLE SCIENZE DI BASE
DIPARTIMENTO DI INGEGNERIA ELETTRICA E DELLE TECNOLOGIE DELL'INFORMAZIONE

To my family

EMPOWERING SUSTAINABILITY THROUGH ARTIFICIAL INTELLIGENCE AND EARTH OBSERVATION: LEARNING FROM THE PAST, MONITORING THE PRESENT, BUILDING THE FUTURE

Ph.D. Thesis presented
for the fulfillment of the Degree of Doctor of Philosophy
in Information Technology and Electrical Engineering
by

GIOVANNI GIACCO

March 2024



Approved as to style and content by

A handwritten signature in black ink, appearing to read 'Carlo Sansone', written over a horizontal line.

Prof. Carlo Sansone, Advisor

Università degli Studi di Napoli Federico II

Ph.D. Program in Information Technology and Electrical Engineering

XXXVI cycle - Chairman: Prof. Stefano Russo



<http://itee.dieti.unina.it>

Candidate's declaration

I hereby declare that this thesis submitted to obtain the academic degree of Philosophiæ Doctor (Ph.D.) in Information Technology and Electrical Engineering is my own unaided work, that I have not used other than the sources indicated, and that all direct and indirect sources are acknowledged as references.

Parts of this dissertation have been published in international journals and/or conference articles (see list of the author's publications at the end of the thesis).

Napoli, March 11, 2024



Giovanni Giacco

Abstract

Sustainability is a concept that has gained significant attention in recent years due to the growing concerns about the impact of human activities on the environment. At its core, sustainability is about meeting the needs of the present without compromising the ability of future generations to meet their own needs. While the data-driven model has traditionally prioritized the accumulation and availability of vast datasets in the past, now we need a shift from a data-driven Earth Observation (EO) to a User-Centric Earth Observation to focus on user's needs. This thesis grounds its approach on the pillars of Learning from the Past, Monitoring the Present, and Building the Future. *Learning from the Past* explores retrospective analysis through applications like Land Consumption Assessment and Aboveground Biomass Estimation, integrating Remote Sensing data with novel deep learning models to generate historical and contemporary maps. These applications demonstrate the ability to monitor temporal changes in urban expansion and ecological biomass, providing a foundational understanding for future sustainability measures. *Monitoring the Present* addresses the need to continuously assess ongoing sustainability initiatives through accessible, shareable, and usable data. The proposed GeoAI Processing Block embodies this concept by enabling the serverless execution of AI algorithms on EO Data within a Data Space ecosystem. The efficacy of this block is showcased through a case study on monitoring Thermal Comfort for Soft mobility. In *Building the Future*, the thesis illustrates two methodologies for user-centric action planning. The first leverages historical insights to guide future urban interventions, while the second employs scenario thinking to foresee the effects of urban interventions before their actual implementation by utilizing Generative Adversarial Networks to generate synthetic multispectral satellite imagery. These methods equip users with the tools not only to conceptualize but also to shape the future strategically, ensuring that actions taken today are informed by their long-term implications. This work contributes to the field by demonstrating a multifaceted approach that combines the use of Earth observation data and Artificial Intelligence, facilitating a comprehensive process that spans from understanding historical patterns to real-time monitoring and strategic planning for sustainable development.

Keywords: Sustainability, AI, Earth Observation, Remote Sensing

Sintesi in lingua italiana

Il concetto di sostenibilità ha acquisito notevole attenzione negli ultimi anni a causa delle crescenti preoccupazioni riguardo l'impatto delle attività umane sull'ambiente. La sostenibilità si concentra sul soddisfare i bisogni del presente senza compromettere i bisogni delle generazioni future. Sebbene in passato il modello basato sui dati abbia tradizionalmente privilegiato l'accumulo e la disponibilità di vasti dataset, oggi si rende necessario un cambiamento da un approccio data-driven a un approccio user-driven che metta al centro l'utente e le sue necessità. Questa tesi propone un approccio basato sul *Imparare dal Passato*, *Monitorare il Presente* e *Costruire il Futuro*. *Imparare dal Passato* esplora l'analisi retrospettiva attraverso applicazioni come la *Valutazione del Consumo del Suolo* e la *Stima della Biomassa*, utilizzando l'integrazione dei dati di Telerilevamento con algoritmi di Deep Learning per generare mappe sia storiche che aggiornate. Queste applicazioni dimostrano la capacità di monitorare i cambiamenti nell'espansione urbana e nella biomassa terrestre, fornendo una comprensione fondamentale per le future misure di sostenibilità. Il *Monitoraggio del Presente* affronta la necessità di un monitoraggio continuo attraverso dati accessibili e condivisibili. E' introdotto in questa tesi un modulo software che consente l'esecuzione di algoritmi di Intelligenza Artificiale su dati geospaziali dell'ecosistema dei Data Space. L'efficacia di questo modulo è dimostrata attraverso un caso di studio sul monitoraggio del Comfort Termico per la Mobilità dolce. In *Costruire il Futuro* la tesi presenta due metodologie per supportare gli utenti nella pianificazione di azioni per massimizzare l'impatto sociale futuro e nella valutazione preventiva di azioni potenziali. Il primo approccio sfrutta le analisi storiche per guidare gli interventi urbani futuri, mentre il secondo impiega l'analisi di scenari ipotetici per analizzare e comprendere le implicazioni di varie azioni prospettiche attraverso l'impiego di immagini satellitari sintetiche generate con Generative Adversarial Networks. Questo lavoro contribuisce con un approccio che combina l'uso dei dati di Osservazione della Terra e dell'Intelligenza Artificiale al monitoraggio in tempo reale e alla pianificazione predittiva per lo sviluppo sostenibile, facilitando un processo completo che spazia dalla comprensione dei modelli storici.

Parole chiave: Sostenibilità, AI, Osservazione della Terra

Acknowledgements

I extend my deepest gratitude to my supervisor, Carlo Sansone, whose expertise, guidance, and unwavering support have been pivotal throughout this research journey. His insightful feedback, encouragement, and patience have profoundly shaped this work, and for that, I am incredibly thankful. Additionally, I wish to express my heartfelt thanks to Stefano Marrone, for his invaluable contributions to my research. His perspectives and critiques have been instrumental in refining my work, and his support throughout my doctoral studies has been immensely beneficial. I would also like to express my sincere appreciation to the coordinator of the doctoral program, Stefano Russo, for his support during these three years.

I express my profound gratitude to my employers at Latitudo 40 for their understanding and support, which allowed me to complete this academic journey while fulfilling my professional duties. Their flexibility and encouragement have been crucial in managing the demands of both my work and research. Their belief in the value of continuous learning and professional development has been a source of motivation and has played a significant role in my ability to pursue and conclude this doctoral program. My heartfelt thanks to the entire team for their support. I am deeply thankful to my colleagues Donato Amitrano, Mattia Rigioli and Paolo de Piano at Latitudo 40 for their intellectual companionship. Working alongside them on various research themes that have found their way into this thesis has been both a privilege and a pleasure. Their insights, combined efforts, and shared moments of breakthroughs have significantly enriched my research experience.

Furthermore, my deepest appreciation goes to my family. I owe a profound debt of gratitude to my parents and my brother, as well as my sister-in-law, for their boundless support and understanding. They have more than compensated for my occasional lapses in attention and time

with their extraordinary grace and light-heartedness.

Special thanks are also due to my friends Andrea and Giovanna. Their consistent support and unwavering belief in my abilities have been a constant source of strength. I am profoundly grateful for their friendship.

Finally, to my wife Raffaella, for being my pillar of strength, providing comfort and understanding through the highs and lows. She has gracefully filled the gaps left by my countless hours dedicated to research, understanding that the weekends and weekdays spent working on this doctorate were necessary sacrifices. Her love has been the gentle force that propelled me forward and the quiet comfort that replenished my spirit after long days. Her presence and care ensured that our life continued with warmth and joy, even when my attention was divided. This accomplishment is as much a testament to her love as it is to my perseverance, and for this, I am eternally grateful.

Contents

Abstract	i
Sintesi in lingua italiana	iii
Acknowledgements	vi
List of Acronyms	xi
List of Figures	xix
List of Tables	xxii
List of Symbols	xxiii
1 Introduction	1
1.1 Sustainability and Earth Observation	1
1.2 Earth Observation and Artificial Intelligence	7
1.2.1 AI for Sustainability	7
1.2.2 Advanced processing and computing	9
1.2.3 AI Data Spaces for the Green Deal	11
1.3 Proposed approach	18
1.4 Thesis structure	21
2 Learning from the Past	23
2.1 Land Consumption Assessment	23
2.1.1 Introduction	23
2.1.2 Related Works	27

2.1.3	Material and Methods	30
2.1.4	Experimental Setup	43
2.1.5	Results	45
2.1.6	Conclusion	48
2.2	Aboveground Biomass and Carbon Storage Estimation . . .	52
2.2.1	Introduction	52
2.2.2	Related Works	55
2.2.3	ReUse: REgressive U-Net	58
2.2.4	Machine Learning Ensembles: Active Learning Strategies for Model Transfer and Field Sampling Reduction	69
2.2.5	Conclusion	80
3	Monitoring the Present	85
3.1	Introduction	85
3.2	GeoAI Processing Block: A serverless AI Module for data space ecosystems	89
3.3	Monitoring Soft-Mobility Thermal Comfort	101
3.3.1	Downscaling Land Surface Temperature	101
3.3.2	The SPOTTED Use Case	104
3.4	Conclusion	105
4	Building the Future	111
4.1	Introduction	111
4.2	Planning Urban Forest as Nature-based Solutions	113
4.2.1	Introduction	113
4.2.2	The case study of Turin	118
4.2.3	The case study of Asti	132
4.3	Scenario Thinking through GAN-Based Synthetic Satellite Imagery	140
4.3.1	Introduction	140
4.3.2	Materials and Methods	143

4.3.3 Results	145
4.4 Conclusion	148
5 Conclusions	153
Bibliography	157
Author's publications	181

List of Acronyms

The following acronyms are used throughout the thesis.

AGB	Aboveground Biomass
AI	Artificial Intelligence
ALS	Airborne Laser Scanning
CNN	Convolutional Neural Network
DL	Deep Learning
ESA	European Space Agency
EO	Earth Observation
GAN	Generative Adversarial Network
GIS	Geographic Information System
IS	Impervious Surface
LST	Land Surface Temperature
LULC	Land Use Land Cover
ML	Machine Learning

NIR	Near-Infrared
NN	Neural Network
RQ	Research Question
RS	Remote Sensing
SDG	Sustainable Development Goal
SIT	Territorial Information System
UF	Urban Forestry

List of Figures

- 1.1 Geospatial Information and Earth Observations: Supporting Official Statistics in Monitoring the SDGs (March, 2016). Credit: GEO 6
- 1.2 Development and Operational Phase of an AI algorithm . . 12
- 1.3 European Common Data Space implemented through sectorial data space initiatives 13
- 1.4 Interoperability level from the New European Interoperability Framework [1] 16
- 1.5 Data spaces technical building blocks [2]. 16
- 1.6 The three pillars of the proposed approach: Learning from Past, Monitoring the Present, Building the Future. 18
- 1.7 Thesis structure 22

- 2.1 High-resolution satellite image of a dense urban area in Rome, Italy (2.1a), and the corresponding imperviousness maps from ISPRA (2.1b), ESRI (2.1c) and DynamicWorld (2.1d). For imperviousness maps, red indicates impervious areas while green corresponds to pervious ones. 25
- 2.2 Example of the area of interest selection in the proposed tool. 30

2.3	Inference Pipeline scheme. On the left is the process of the Sentinel-2 satellite imagery retrieval. On the right is the execution of the prediction. As detailed in Section 2.1.3, the inference is managed through a Docker image. Dask makes the solution scalable by parallelizing the execution. .	31
2.4	Schematic representation of the impervious surface dataset generation.	32
2.5	Localization of the Sentinel-2 tiles used in dataset generation.	35
2.6	The proposed <i>ReFuse</i> encoder–decoder architecture: the encoder consists of two parallel ResNet-50 networks where every residual block output from the second branch is fused into the main one by a feature map summation. Feature maps from the encoder are then connected with the decoding part by feature map concatenation, such as in the U-Net cross-connections.	37
2.7	Schematic representation of the inference pipeline workflow.	39
2.8	Workflow Management System layers. In bold, the choices made in this thesis.	40
2.9	Flow execution process for the imperviousness map generation. A Docker image stores library dependencies. The model, weights, and flow code are stored in object-based storage. During flow execution, the orchestrator pulls the image from the Docker registry and, using it, starts a Dask cluster. Finally, the prediction flow is executed on the Dask cluster.	42

2.10	Learning Rate (LR) tuning analysis for the ReFuse model. (a) The output of the LR range test. The figure highlights the maximum learning rate before loss increases. (b) Learning rate values during training with the one-cycle policy [3]. The one-cycle strategy increases the learning rate from an initial value to a maximum learning rate. It then decreases the learning rate from the maximum value to some minimum, substantially lower than the starting learning rate.	45
2.11	Example of an imperviousness mask obtained using ReFuse. The top left is the considered path (only RGB bands have been reported for visualization reasons). The top right is the corresponding ISPRA Ground Truth. The bottom left is the imperviousness mask obtained by the proposed ReFuse architecture. The bottom right is the ReFuse segmentation errors. (a) RGB patch. (b) Ground truth. (c) ReFuse. (d) Error mask: FP (blue), FN (red).	49
2.12	Semantic mapping with ReFuse model for two areas in the city of Turin. (a,e) Optical images extracted from Google Earth. (b,f) Sentinel-2 RGB images. (c,g) Ground truth extracted from ISPRA Soil consumption map. In red, imperviousness surfaces. (d,h) Predicted maps with the ReFuse model. Despite errors in the ground-truth labels (see image (c)), the network correctly detects the green infrastructure within the area, i.e., trees, grass, and parks (see image (d)).	50
2.13	Example of the imperviousness map visualization into a GIS web application.	51
2.14	ReUse architecture for pixel-wise regression. The input is the Sentinel-2 image with dimensions (patch size, patch size, number of channels); the output is the AGB image with dimensions (patch size, patch size, 1).	59

2.15	Base-map of the test area in Central Europe.	65
2.16	ReUse’s AGB predictions with raw bands in Central Europe.	65
2.17	ReUse’s AGB predictions with raw bands and feature ex- traction in Central Europe.	65
2.18	Predictions of AGB using the machine learning approach [4] in Central Europe.	65
2.19	Predictions of AGB using the machine learning approach [5] in Central Europe.	65
2.20	On the left is the predicted aboveground biomass raster of the Astroni nature reserve before the July 2017 fire; on the right is the predicted aboveground biomass raster after a major fire event for the same area.	68
2.21	Composition of the exploited dataset for some sample clips. From left to right, S1 SAR image, natural color S2 image and AGB map. The size of each clip is 256×256 pixels. The pixel spacing of the maps is 10 m.	70
2.22	Proposed workflow for AGB model transfer. Colored blocks identify input/output data.	71
2.23	Schematic representation of histogram calculations for ac- tive sampling.	73
2.24	Some scatterplots relevant to the performed experiments. (a) Clip 0, bootstrap ($R^2 = 0.798$, $RMSE = 9.98t/ha$), (b) Clip 4, ensemble regression $k = 5$ ($R^2 = 0.589$, $RMSE =$ $27.6t/ha$), (c) Clip 4 neural network ($R^2 = 0.731$, $RMSE$ $= 27.6t/ha$). (d) Clip 7, ensemble regression, all samples ($R^2 = 0.515$, $RMSE = 29.1t/ha$), (e) Clip 7 ensemble re- gression, $k = 5$ ($R^2 = 0.423$, $RMSE = 31.8t/ha$), (f) en- semble regression results in the plane (R^2 , $RMSE$) for the different sampling settings. Overall means reported with solid markers.	78

3.1	Data spaces technical building blocks. Into a yellow box building blocks, the GeoAI Processing Block deals with. . .	90
3.2	SPOTTED Architecture.	91
3.3	SPOTTED Data Flow.	92
3.4	GeoAI Processing Block Architecture.	94
3.5	Integration detail of the GeoAI Processing Block with the FiWARE Context-Broker.	97
3.6	Comparison between several machine learning regressors for the task of LST downscaling.	103
3.7	Surface Heat Island Analysis for the Soft-Mobility in the Municipality of Milan, Italy. The hottest bike lanes are indicated in red, while the coolest ones are marked in blue.	106
4.1	Location of the city of Turin and subdivision into neighborhoods.	117
4.2	The research workflow methodology.	120
4.3	Differentiation of buildings according to compliance with the <i>3 trees from every home</i> rule (left); differentiation of neighborhoods according to compliance with the <i>3 trees from every home</i> rule (right).	126
4.4	Analysis of the % Tree Canopy Cover of the city of Turin (left); the subdivision of neighborhoods according to the % Tree Canopy Cover (right)	127
4.5	Proximity map concerning urban green areas at urban level (left); Subdivision of neighborhoods according to their potential to comply with the '300 m from the nearest park or green space' rule (right).	128
4.6	The results of the Urban Cooling Index at the urban level (left). The subdivision of neighborhoods which have different heat mitigation values (right).	129

4.7	Subdivision of neighborhoods according to Population Density values.	131
4.8	Identification of neighborhoods needing NbS implementation precedence (left). Detail on the neighborhoods to identify those in which the municipality should start investing immediately by implementing NbS (right)	132
4.9	Location of the city of Asti and its division into voting sections or wards. (Sources: Google Maps and OpenStreetMap for the maps; SIT Municipality of Asti for the voting sections. Available online at [6].	133
4.10	The presence of a minimum of three trees within a 30 – m radius of buildings. The left image identifies the buildings, while the right image classifies the electoral sections based on the number of trees surrounding each building. The score inside the voting district is computed by averaging the scores of each building.	136
4.11	A percentage analysis of tree canopy cover (TCC) at the city and electoral section level.	137
4.12	Distance of buildings from green spaces and parks of 0.5 hectares or more.	137
4.13	Analysis of the greenery vigour quality at the city and electoral section.	138
4.14	Classification of areas according to the potential need for implementing nature-based solutions (NbS). Classes with lower values indicate a greater need for intervention.	139
4.15	Proposed approach for the generation of Synthetic Sentinel-2 Multispectral image from Simulation Scenarios.	144

4.16 Synthetic Multispectral Sentinel-2 data. (left) Input scenario, (center) Step 1 output: synthetic RGB image with conditional GAN, (right) NDVI computed on the Synthetic Multispectral Sentinel-2 image generated from the step 2. .	147
4.17 What-if analysis for four scenarios to study the impact of different future Nature-based solution installations in terms of Land Surface Temperature	148

List of Tables

2.1	List of Sentinel-2 products used to generate the dataset. . .	34
2.2	Results obtained by the proposed approach (last row) and the considered competitors for the imperviousness map extraction. For the proposed approach (ReFuse), using brackets under the “bands” column highlights the ability of the proposed approach to use different band types without the need for resizing.	47
2.3	List of the file names downloaded from the ESA Biomass Climate Change Initiative of the Global aboveground forest biomass for 2018, v3 in Geotiff format containing AGB rasters. In each GeoTiff file, a study area was cut out and reported in Well-known text (WKT) format.	63
2.4	The results of the experiments performed on Vietnam and Myanmar study areas of [4] and [5] and Central Europe study area. At each iteration of the eight-fold cross-validation, six folds are used for training, one for validation, and one for testing. The averages and standard deviations of the Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and R^2 metrics are calculated on the test set. In bold, the best value for each metric per study area.	66

2.5	AGB statistics for the considered area.	69
2.6	Aggregated AGB estimation results for the different tested sampling settings and benchmarking techniques. The column "Inc" refers to the incremental sampling setting.	76
2.7	VIP variables extracted from the regression of the reference area.	77
3.1	Comparison of the Machine Learning requirements between the pillars <i>Learning from Past</i> and <i>Monitoring the Present</i> .	108
3.2	Results for the task of LST downscaling	109
4.1	Detail of neighborhoods, number of inhabitants, neighborhood surface (data elaboration from WorldPop project [7], 2021).	119
4.2	Detail of neighborhoods and related Population Density. . .	130
4.3	Results for step two of the proposed approach: generate multispectral content from RGB bands.	146

List of Symbols

The following symbols are used within the thesis

c Speed of light in a vacuum

h Planck constant

Chapter 1

Introduction

Life is divided into three terms - that which was, which is, and which will be. Let us learn from the past to profit by the present, and from the present, to live better in the future.

William Wordsworth

1.1 Sustainability and Earth Observation

Sustainability is a concept that has gained significant attention in recent years due to the growing concerns about the impact of human activities on the environment. At its core, sustainability is about meeting the needs of the present without compromising the ability of future generations to meet their own needs [8].

The environmental and sustainability problem became evident from 1960 onwards, as the world began to realize the consequences of rapid industrialization and the depletion of natural resources. One of the critical contributions to understanding these challenges was the study "The Limits to Growth" [9] [10], which highlighted the finite nature of resources and the potential consequences of unchecked economic and population growth on the environment. the concept of

The findings underscored the necessity of integrating environmental, social, and economic considerations into development strategies. It served

as a wake-up call for governments, organizations, and individuals to reevaluate their approach to development and prioritize sustainability in their decision-making processes, emphasizing the importance of balancing economic growth with environmental protection and social equity, laying the foundation for the concept of sustainable development [11] [12]. The concept of sustainable development is crucial for addressing the global challenges we face today, including climate change, resource depletion, and social injustice. It provides a framework for decision-making and action that considers the long-term implications of our choices to create a sustainable and prosperous future for all. Furthermore, sustainable development calls for implementing policies and practices that promote a resilient and sustainable economy while also addressing social challenges such as poverty, inequality, and access to education and healthcare. Sustainable development is about creating a balanced and inclusive world where economic prosperity, social well-being, and environmental health go hand in hand.

In 2015, the United Nations adopted the 17 Sustainable Development Goals (SDGs), aiming to end poverty, protect the planet, and ensure peace and prosperity [13].

The goals address a broad range of interconnected global challenges, including those related to poverty, inequality, climate, environmental degradation, prosperity, and peace and justice. They recognize that ending poverty and other deprivations must go hand-in-hand with strategies that improve health and education, reduce inequality, and spur economic growth – all while tackling climate change and working to preserve our oceans and forests. The SDGs were defined to provide a universal set of targets for all countries to work towards, acknowledging that the urgent environmental, political, and economic challenges faced by the world today require coordinated global action. The 17 interconnected goals and their associated targets are designed to address the root causes of poverty and the universal need for development that works for all people.

The Sustainable Development Goals (SDG) framework encompasses 169 objectives to be achieved by 2030, monitored through 232 quantitative indicators. They address the three interconnected elements of sustainable development: economic growth, social inclusion, and environmental sustainability. These indicators are categorized by the Inter-agency and Expert Group on SDG Indicators into three tiers, distinguished by their

methodological maturity and the availability of global data. Tier 1 indicators possess a globally recognized methodology for calculation and are supported by data regularly provided by a minimum of 50 of countries. Tier 2 indicators, while methodologically sound, suffer from a data availability shortfall. Notably, no indicators are classified under Tier 3, which is reserved for those lacking an internationally approved methodology. Initially, the methodologies for calculating these indicators predominantly relied on data from statistical agencies.

Monitoring progress toward the SDGs and their associated targets necessitates a comprehensive global Indicator framework. This framework should encapsulate the complex and ambitious objectives for ongoing national and societal development. Effective tracking of these indicators demands the utilization of diverse data types. This includes traditional data sources such as national accounts, household surveys, routine administrative data, and emerging data sources external to national statistical systems. Notably, EO and geospatial information (GI) play a significant role, especially when coupled with advanced data processing techniques suitable for handling large volumes of EO data. The fusion of these varied data sources has the potential to enhance the way development progress is monitored and societal well-being is advanced. Since Earth observation and geospatial information are often continuous in their spatial and temporal resolutions, their use in SDG monitoring can prove essential in capturing the sustainability of developments underpinning the SDG framework [14].

EO encompasses a range of methodologies, such as utilizing drones, aircraft, and satellites. Notably, over 2000 active EO satellites, managed by space agencies, government institutions, and commercial entities [15], have significantly increased the availability of detailed information about Earth's conditions and characteristics [16].

EO satellite imagery is broadly categorized into two types based on the sensors used for image capture: passive sensors, which detect natural radiation emitted or reflected by the Earth's surface, and active sensors, which emit radiation and then capture the reflected or refracted echoes from the surface. These sensors provide data across four key resolutions: spectral, spatial, radiometric, and temporal. Spectral resolution refers to the sensor's ability to differentiate between ranges of radiation wave-

lengths, with different spectral bands offering unique spectral signatures for various land cover types like soil, water, or buildings. Spatial resolution describes the surface area represented by each pixel in an image, while radiometric resolution is the sensor's capacity to distinguish between different light intensities. Temporal resolution, conversely, pertains to the revisit time or the frequency at which a sensor passes over a specific area on Earth. Additionally, the nature of the data satellites provide varies according to their orbital paths. Satellites in geostationary orbits remain fixed over a specific area of Earth, whereas those in Low Earth orbits move relative to the planet's surface, tracking different areas as they orbit.

Despite the diverse nature of EO-based measurements across different mission types, their applications lead to various characteristics that support the demand for global governance and the derivation of indicator information:

- *Scale*: satellites have the capability to gather information across various scales, ranging from local to global. They are often the primary source of global data for many parameters. Depending on the specific use and level of detail required, satellites can produce large-scale datasets within short timeframes - spanning from daily to annual updates as necessary and feasible with current technology. This enables rapid updating of indicator information regardless of time or weather conditions.
 - *Long time series and continuity*: satellites have been continuously gathering data over extended periods, with missions spanning from the 1970s and projected beyond 2030. This extensive dataset enables governments to monitor progress, establish baselines for future trend analysis, ensure compliance with agreements, improve predictive capabilities, and effectively manage and mitigate various issues.
 - *Consistency and comparability*: satellites enable effective comparison of results between different countries, addressing potential issues caused by lack of standardization in measurements or methods. This can improve the accuracy and relevance of regional and global statistics.
 - *Diversity of measurements*: advancements in instrumentation have
-

led to a growing variety of Earth observation satellite missions. These missions measure numerous geophysical parameters daily and operate from various types of satellite orbits.

- *Complementarity with traditional statistical methods*: some specific SDG indicators can be directly monitored using EO datasets. Additionally, these datasets can serve as a valuable and supplementary source of information to validate in-situ data measurements, communicate the geographic dimensions and context of the indicators, and enable disaggregation of the indicators when necessary.
- *Free and open data is increasing*: not all countries can build and send their own Earth observation satellites. The United States has made its mission data freely accessible for a long time, and with Europe's Copernicus program offering multiple satellite data streams under a free and open data policy, there is now an improved outlook for developing countries to access the necessary EO data. This makes data from these missions essential for all nations, as it significantly impacts a global scale.

Group on Earth Observations (GEO) and the Committee on Earth Observation Satellites (CEOS) have conducted an analysis that has identified particular objectives and measures that can be aided by observations of the Earth, as outlined in Figure 1.1.

From Figure 1.1, it is clear that Earth Observation has the capacity to offer a wide array of parameters for the SDG framework. This includes data on atmospheric conditions, ocean health, agricultural output, forest cover, climate patterns, occurrences of natural disasters, and availability of natural resources. Additionally, EO provides insights into urbanization trends, biodiversity status, and human welfare indicators. Among these indicators, population distribution and mapping of cities/infrastructure stand out as particularly crucial since they contribute to all the SDGs. Furthermore, zero hunger (SDG 2), clean water and sanitation (SDG 6), climate action (SDG 13), life below water (SDG 14), and partnership for goals (SDG 17) are among the SDGs that benefit from various EO indicators.

SUSTAINABLE DEVELOPMENT GOALS

	Population distribution	Cities and infrastructure mapping	Elevation and topography	Land cover and use mapping	Oceanographic observations	Hydrological and water quality observations	Atmospheric and air quality monitoring	Biodiversity and ecosystem observations	Agricultural monitoring	Hazards, disasters and environmental impact monitoring
1 No poverty										
2 Zero hunger										
3 Good health and well-being										
4 Quality education										
5 Gender equality										
6 Clean water and sanitation										
7 Affordable and clean energy										
8 Decent work and economic growth										
9 Industry, innovation and infrastructure										
10 Reduced inequalities										
11 Sustainable cities and communities										
12 Responsible consumption and production										
13 Climate action										
14 Life below water										
15 Life on land										
16 Peace, justice and strong institutions										
17 Partnerships for the goals										

Figure 1.1: Geospatial Information and Earth Observations: Supporting Official Statistics in Monitoring the SDGs (March, 2016). Credit: GEO

1.2 Earth Observation and Artificial Intelligence

1.2.1 AI for Sustainability

A substantial portion of Earth Observation (EO) data is composed of densely distributed images across the Earth's surface, encompassing various forms like optical (panchromatic, multi-spectral, hyper-spectral), thermal imagery, and techniques like synthetic aperture radar (SAR), as well as derived raster layers such as elevation models.

In this context, Computer Vision (CV) - focused on extracting information from images through computational methods - is highly applicable to EO. Both fields share key tasks: object recognition and detection, image segmentation and labeling, change detection, time series analysis, and spatial regression for surface variable retrieval. Some of the earliest instances of computer vision algorithms were created for EO [17].

CV has embraced AI techniques, particularly Machine Learning, early, which has benefited Earth observation. The field of Earth Observation has significantly evolved with the incorporation of Artificial Intelligence technologies. AI has revolutionized the processing and analysis of EO data, enabling more efficient and accurate extraction of valuable information for various applications. By leveraging AI algorithms, the interpretation of vast amounts of EO satellite imagery has become more sophisticated, leading to enhanced capabilities in monitoring and understanding our planet [18] [19].

In the field of Earth Observation, applying various learning paradigms has become increasingly pivotal. These paradigms each bring unique strengths and face distinct challenges in their application to EO:

- *Supervised Learning*: the most prevalent method in EO relies heavily on labeled data to train algorithms for tasks such as land-cover mapping. However, the complexity of EO data, characterized by intricate physical phenomena and spectral variability, often makes the collection of high-quality, labeled samples both costly and time-consuming, especially for large-scale applications [20].
 - *Unsupervised Learning*: Unsupervised learning methods utilize the input data in the training set solely, as their objective is to identify
-

concealed patterns within the data and can be managed using clustering or dimensionality reduction techniques [21] [22].

- *Semi-supervised Learning*: Semi-supervised learning integrates elements of both supervised and unsupervised learning. Semi-supervised learning alternatively uses data with labels to learn a representation function to the target objective (e.g. classification) and data only to get the underlying data structure [23] [24] [25].

EO gets advantages from different ML techniques such as classification, semantic segmentation, regression, object detection, clustering, and dimensionality reduction. EO has utilized classification to categorize land cover types, monitor changes in vegetation, identify urban areas, and classify various environmental features. Semantic segmentation techniques can aid in the detailed identification and delineation of specific objects within satellite imagery, such as land cover classification at a finer scale, urban infrastructure mapping, and environmental monitoring. Regression models can be employed to predict surface variables, such as temperature, precipitation, and soil moisture, based on EO data, enabling the generation of valuable environmental and climate forecasts. Object detection algorithms can contribute to identifying and monitoring specific objects of interest, such as infrastructure, natural resources, or changes in land use over time. Clustering algorithms offer the capability to group similar Earth observation data points, enabling the identification of patterns and trends within large datasets. Dimensionality reduction techniques can assist in visualizing and understanding complex EO datasets by reducing the number of input variables while preserving essential information.

Machine learning has a significant impact on achieving the Sustainable Development Goals through its applications in Earth Observation. [12] illustrates several studies that utilize Machine Learning techniques in the field of Earth Observation and their connection to the Sustainable Development Goals. It provides a strong indication of the potential of machine learning in using Earth observation to track Sustainable Development Goals, demonstrating the diverse applications of ML in EO and how they contribute to addressing specific SDGs.

So, the seamless integration of machine learning techniques into the Earth Observation field has brought great potential for advancement. How-

ever, it also brings about unique challenges that must be addressed in the EO sector.

A key difficulty lies in the process of data fusion, where information from multiple, often disparate sources such as satellite imagery, radar data, and ground-based observations must be integrated. This integration necessitates creating a cohesive representation from varied inputs, a task that remains largely unresolved. Additionally, the EO field grapples with the dynamic and diverse range of data sources, including non-traditional ones like social media posts. This diversity calls for flexible, user-centric solutions capable of processing various data types to deliver accurate geospatial information.

Another significant challenge is marrying the data-driven approaches of modern AI with the rich, physically grounded domain knowledge of Earth systems. While AI models are adept at predicting patterns in complex data, physical models excel in describing well-understood phenomena. A hybrid approach that combines these strengths could yield more consistent, explainable, and efficient results. However, developing such hybrid models is not straightforward.

Compounding these challenges is the issue of reference data critical for training ML models. Obtaining large, well-labeled datasets for EO can be daunting, with fieldwork often required for data collection. The reference data in EO often suffers from problems like misalignment, sparse and context-less observations, and highly unbalanced distributions, which hinder learning accurate spatial patterns. Moreover, the outputs in remote sensing often have additional structures, such as taxonomic hierarchies, adding another layer of complexity to model training and prediction.

In essence, applying ML to EO involves navigating a multifaceted landscape, where integrating diverse data types, combining statistical and physical modeling, and managing complex reference data all play critical roles in the successful analysis and interpretation of Earth observation data.

1.2.2 Advanced processing and computing

Recent advances in the Earth Observation field, thanks to Machine Learning, are creating unprecedented opportunities to evaluate and monitor policy decisions as well as humanitarian initiatives. Despite the advantages of using ML techniques, it may require greater computational

resources. From this viewpoint, the administration of large-scale EO data combined with Machine Learning (ML) forms the foundational backbone that fuels the ongoing surge of advancements in Artificial Intelligence (AI) applied to EO. This synergy is instrumental in driving forward the field of EO, leveraging the power of AI and ML to unlock new capabilities and insights.

The necessity of developing a new system architecture to support Earth Observation is underscored by the increasing complexity and heterogeneity of EO data. Remote sensing is experiencing rapid development, leading to significant growth in the diversity and volume of available data. This presents new challenges for effectively processing large and varied datasets to support emerging Earth observation applications.

EO data's increasing complexity and heterogeneity have highlighted the need for a new system architecture to support these advancements effectively. Leveraging this data to its full potential requires significant expertise and sophisticated data management systems. Recent EO exploitation platforms, like the Google Earth Engine [26], have simplified large-scale analyses through some level of processing and data abstraction. Despite their benefits, these platforms are built on diverse technologies with different interfaces and data formats, complicating their federated use. EO ecosystems face the challenge of managing assets that range from datasets and processing tools to ML algorithms and high-performance computing clusters. The high *heterogeneity* of these assets makes their management an open problem, particularly in storing, querying, and composition. Moreover, most EO analytics are exploratory in nature, necessitating quick query results, which becomes challenging when data assets are scattered across multiple sites and platforms.

The *decentralized* nature of EO data also poses a challenge to standard ML algorithms, which typically require centralizing data before training. This centralization is often impractical due to the large volume of EO data, privacy regulations, and other legal constraints. Federated learning emerges as a solution to these governance and privacy concerns by learning a local model on each distributed site and aggregating these into a shared global model. This approach minimizes data transfer costs and can lead to more accurate and less biased models. However, applying federated learning to geo-distributed EO data is challenging. EO data rarely meet

the assumption of being independent and identically distributed, and the inherent heterogeneity of the data complicates this further.

Another aspect of building an effective EO system architecture is *interoperability* – the ability to seamlessly integrate and communicate data and models across different platforms for common goals like building a global model. The lack of interoperability presents a significant barrier to efficient and effective collaboration in EO ecosystems.

It's crucial to revisit the entire processing stack to support emerging applications in EO. This includes redesigning data ingestion into processing systems, developing new indexing techniques for fast querying of EO data, and adapting data analyses to run on diverse processing hardware. Particularly pressing is the shift of satellite data processing from Earth to space, necessitating onboard payload data processing. This shift requires flexible, robust solutions with a low computation burden, where AI models can play a vital role.

In summary, finding a new system architecture to support EO is critical in managing EO data's growing complexity and heterogeneity. This architecture must address challenges in data management, analytics, federated learning, interoperability, and onboard processing to unlock the full potential of EO data for global applications.

1.2.3 AI Data Spaces for the Green Deal

The sections before highlighted how the application of Artificial Intelligence and Machine Learning in Earth Observation has brought about significant advancements, particularly in enhancing the sustainability and effectiveness of EO practices. These AI technologies are playing a crucial role in revolutionizing the analysis and interpretation of EO data.

In delineating the life cycle of most AI algorithms, we can discern two primary stages in the application process: firstly, the *Development Phase*, which entails utilizing data or knowledge for training and testing/validating the algorithm, and secondly, the *Operational Phase*, where the algorithm is employed to make decisions or predictions based on the acquired knowledge. Figure 1.2 offers a simplified depiction of how an algorithm progresses from its untrained state to deployment. After deployment, the algorithm must return to the Development Phase (as indicated by the grey dashed lines) for retraining and validation using new data in order to

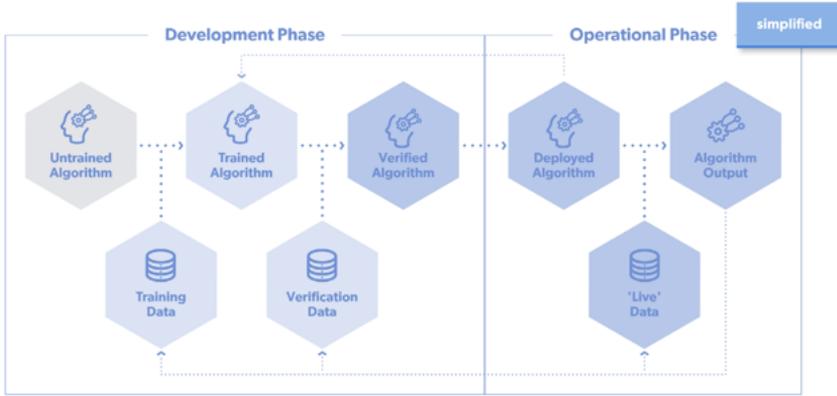


Figure 1.2: Development and Operational Phase of an AI algorithm

enhance the accuracy of the algorithmic system and mitigate the risks associated with concept drift. Furthermore, *Live* data, in conjunction with the algorithm's output, can also serve as valuable resources for future training and validation datasets. It is important to note that there are instances of incremental learning wherein Live data is continuously employed to enhance the algorithm's proficiency further.

It is clear that data stands as a pivotal and strategic resource that determines the success and impact of AI algorithms. To achieve the full potential of AI in Earth observation, it is imperative to address several challenges in the data management process [27]. In contrast to many physical assets, data exhibits the unique trait of being usable repeatedly and concurrently by various stakeholders. The same data can offer endless benefits when shared and not confined to isolated silos, inaccessible to necessary applications.

Creating AI data spaces that allow for seamless and secure sharing of Earth observation data among different stakeholders is crucial. The concept of *Data Spaces* has emerged as a promising approach to address the challenges in managing Earth Observation data. Data Spaces offers a new system architecture that is specifically tailored to meet the increasing complexity and heterogeneity of EO data. Data Spaces provides a solution for integrating diverse data types, combining statistical and physical modeling, and managing complex reference data in EO. This architecture

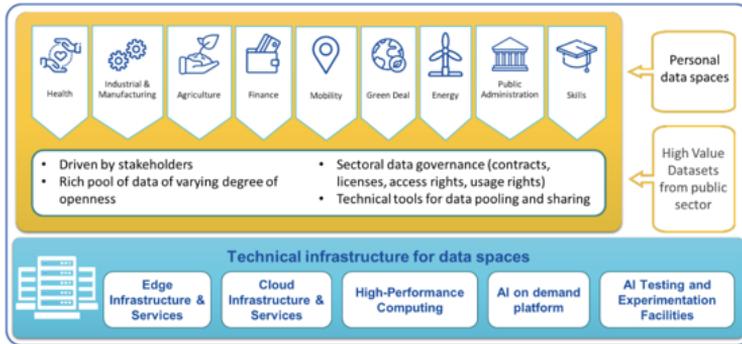


Figure 1.3: European Common Data Space implemented through sectorial data space initiatives

enables the seamless integration and communication of data and models across different platforms, facilitating efficient and effective collaboration within EO ecosystems.

The EU Commission has underscored the significance of data spaces as a pivotal enabler in achieving the Sustainable Development Goals (SDGs). Recognizing the transformative potential of data-driven technologies and solutions, the EU Commission has emphasized that data spaces can play a critical role in addressing various global challenges outlined in the SDGs. Data and data sharing are fundamental elements that the European Commission is actively focusing on, particularly in the context of AI [28, 29].

The European Commission adopted a European strategy for data [30] in 2020, creating a unified data market to secure Europe’s global competitiveness and data sovereignty. The Data Strategy includes the initial support for nine data spaces: *Health*, *Industrial Manufacturing*, *Agriculture*, *Finance*, *Mobility*, *Green Deal*, *Energy*, *Public Administration*, and *Skills* (Figure 1.3).

The different data spaces will be linked together to create a cohesive unique European space. The Green Deal data space is a relevant enabler for the European Green Deal (EGD) [31]. The Green Deal represents the EU’s commitment to tackling climate change and environmental degradation, marking a significant shift towards a more sustainable and inclusive European economy. The primary objective is to achieve climate neutrality

by 2050. It encompasses a broad range of policy areas, including clean energy, sustainable industry, building and renovating, sustainable mobility, biodiversity, and pollution.

The European Green Deal data space aims to provide access to a wide range of environmental data, making it more accessible and beneficial. It will enable data providers to maintain control over their information while facilitating access to numerous applications focused on advancing sustainability in society. It will make many environmental data accessible, usable, and useful following the *FAIR* principles:

- **F**indable: Data should be easy to find for both humans and computers through well-described metadata and unique identifiers.
- **A**ccessible: Once found, data should be accessible with clear and standardized access mechanisms, while considering necessary security or privacy restrictions.
- **I**nteroperable: Data should be compatible with other datasets, tools, and applications for analysis, storage, and processing.
- **R**eusable: Data should be well-documented and preserved to enable reuse in various contexts, ensuring long-term research and application value.

The big change is the shift from a centralized data space operated by a single organization, e.g., a platform, to a *federation model*.

Data space is defined as "a federated, open infrastructure for sovereign data sharing, based on common policies, rules, and standards" [32].

In a federation of data spaces, each specific data space has considerable independence to create and implement its own internal arrangements and ICT structure. However, these individual instances collectively work towards the shared objective of facilitating trusted data sharing.

Accessing the required data is just one aspect of the challenge. The distribution of both data and AI algorithms among multiple entities adds complexity to executing these algorithms, necessitating the implementation of distributed execution mechanisms. This introduces additional intricacies beyond simply obtaining access to the data, such as ensuring compliance with usage policies from all parties involved.

Additionally, regulatory limitations exist regarding the sharing of data and AI, including the General Data Protection Regulation (GDPR) and the forthcoming AI Act. Hence, it is imperative for both organizations and individuals to have authority over who can access the data, its intended purposes, and under what specific conditions (e.g. ethical considerations).

To address the issue of providing sufficient and high-quality data to power AI algorithms, it is essential for organizations to collaborate effectively in order to enable (i) data sharing for AI, (ii) execution of AI algorithms, and (iii) locally run data apps.

Numerous data environments have already been created or are currently being developed. Organizations have the choice to participate in data spaces in various capacities: as *data providers*, potentially for a fee, as *data consumers* aiming to create and implement improved AI algorithms, or as *service providers* facilitating the inclusion of both data suppliers and consumers or supplying components for constructing the data space.

This dynamic environment can operate effectively only when there is a specific level of standardization and effectiveness in the operation of these data areas. So, interoperability is key for the federation of data spaces to interconnect seamlessly. However, achieving data space interoperability encompasses more than just ensuring the technical components are compatible. The new European Interoperability Framework, as developed by the European Commission, distinguishes four interoperability levels: technical, semantic, organizational, and legal interoperability. Figure 1.4 gives an overview of the four layers.

In the end, Figure 1.5 summarizes the technology Building Blocks of a Data Space that ensure:

- **Data interoperability.** These building blocks should be deployed by all data providers and data consumers participating in a data space. This way, each data provider can be sure that any data published can be technically consumed by any data consumer entitled to do so. In contrast, each data consumer can be sure they can technically access and use any data available by any data provider selected [2]. The building block includes:
 - *Data Models and Formats:* This building block establishes a common format for data model specifications and data repre-
-

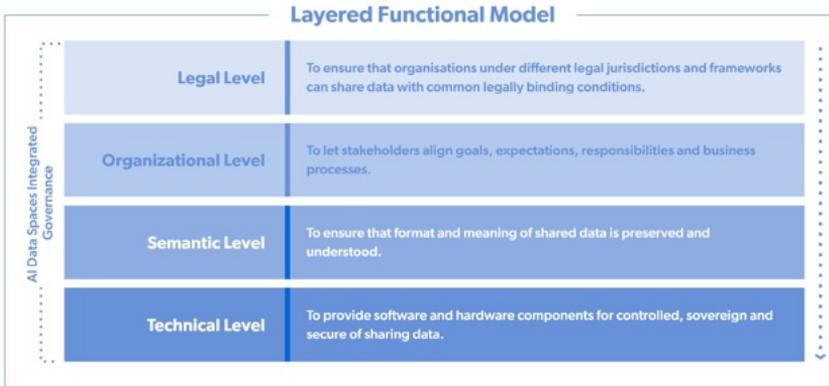


Figure 1.4: Interoperability level from the New European Interoperability Framework [1]

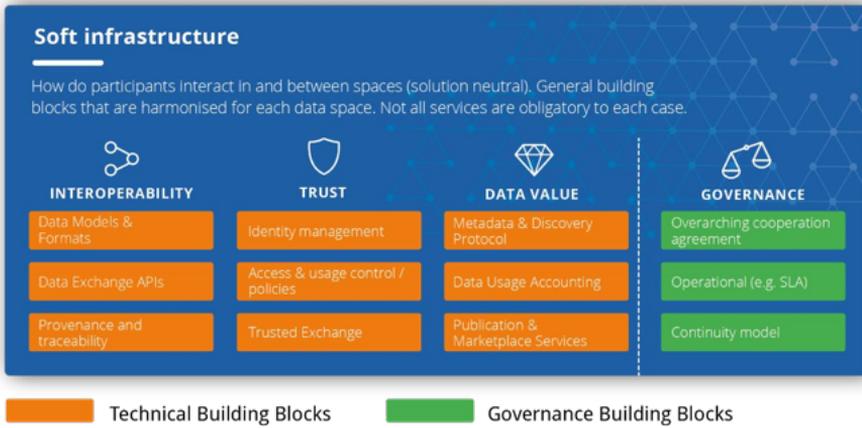


Figure 1.5: Data spaces technical building blocks [2].

sentation in data exchange payloads [2];

- *Data Exchange APIs*: This building block facilitates the sharing and exchange of data (i.e., data provision and data consumption/use) between data space participants. Combined with the Data Models building block, this ensures full interoperability among participants [2].
 - *Data Provenance and Traceability*: This building block provides the means for tracing and tracking the process of data provision and data consumption/use [2].
- **Data sovereignty and trust.** Data Spaces aim to provide the technical infrastructure to ensure that participants within the Data Space can have confidence in each other and maintain control over their shared data. Achieving this goal involves implementing shared standards for participant identity management, verifying the accuracy of their claims, and enforcing agreed-upon policies for controlling data access and usage.
 - **Data value creation.** Data Spaces are essential for facilitating the development of multi-sided markets where participants can derive value from sharing data. To achieve this, there is a need to implement standardized mechanisms that allow for the specification of terms and conditions related to data offerings, including pricing.

Finally, in addition to a shared technology foundation, Data Spaces necessitate **governance** through various business, operational, and organizational agreements among participants. Business agreements encompass terms and conditions for data sharing and the legal framework governing Data Space contracts. Operational agreements address policies such as GDPR compliance or sector-specific regulations like PSD2 in finance, including implementing auditing tools and cybersecurity practices. Organizational agreements establish governance bodies, outlining product specifications and the adoption of technology Building Blocks, as well as business and operational agreements.

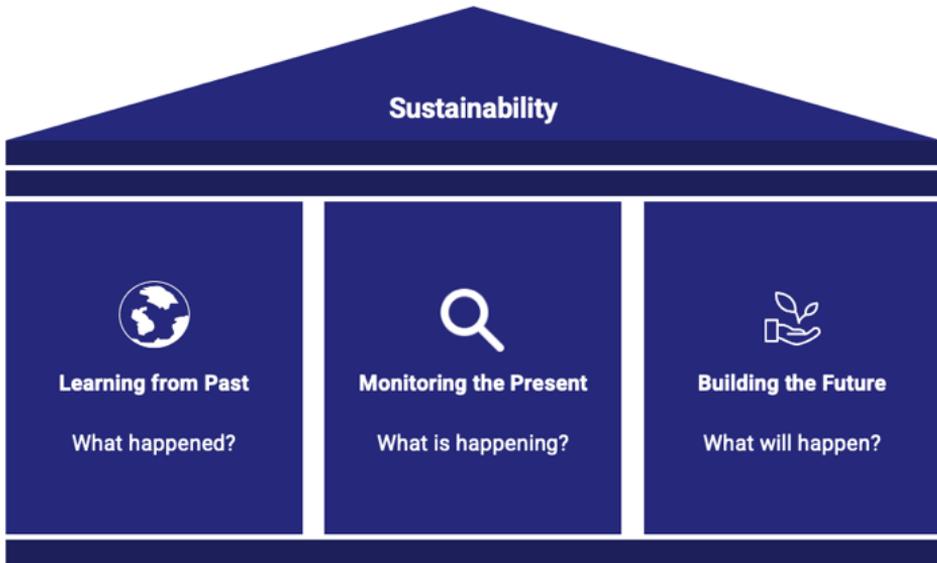


Figure 1.6: The three pillars of the proposed approach: Learning from Past, Monitoring the Present, Building the Future.

1.3 Proposed approach

Users of Earth observations have a wide range of data requirements and priorities that depend on their specific applications [33]. From the raw observational datasets sought after by researchers to the more refined forecasts and products essential for decision-making, the range of data needs is extensive. Within this spectrum, some delve deeply into the raw, unprocessed data, seeking the unvarnished truths of direct observation. Others, perhaps less technically inclined or constrained by time, might opt for processed data products that still offer a level of detail sufficient for their analyses. At the other end of the spectrum, there are end users whose requirements are met by highly processed products, sophisticated tools, or comprehensive forecasts designed for immediate application.

Earth observation data users extend across multiple domains—public and private sectors. These users require data reflective of various observational timescales to suit many applications, such as retrospectives on historical trends, immediate operational decisions, and long-term strategic

forecasting.

Making EO data accessible in catalogs is not enough. We need a shift from a data-driven Earth Observation to a *User-Centric* Earth Observation.

While the data-driven model has traditionally prioritized the accumulation and availability of vast datasets, this approach can often fall short in meeting diverse end-users specific and nuanced needs. Data providers cannot make informed decisions regarding Earth observation systems without knowing users' observation requirements and associated priorities.

A user-driven model also facilitates the democratization of Earth observation data, ensuring that it is accessible and usable to non-specialist users, thereby expanding its utility beyond traditional scientific domains. By actively engaging with the user community to understand their needs, Earth observation providers can develop more intuitive interfaces, create more relevant data products, and ensure that the data serves the broader societal goals it intends to support.

Irrespective of the diversity of users and their applications, three categories of inquiries universally emerge as central to the effective use of Earth observation data:

1. What has happened in the past?

2. What is happening now?

3. What will happen in the future if I take this action? Or, How can I perform an action to maximize its future impact?

Starting from these user's requirements, this thesis grounds its approach on three pivotal pillars (Figure 1.6) :

- **Learning from the Past** provides a historical context, allowing users to analyze and understand previous events and patterns. This retrospective analysis is critical for sustainability as it informs us about the effectiveness of past environmental policies and practices, helping to identify successful strategies and cautionary tales.
 - **Monitoring the Present** addresses the second category of questions, offering real-time data and observations. This constant vigi-
-

lance allows for the immediate assessment of ongoing sustainability initiatives and the health of ecosystems, ensuring that current practices align with sustainable objectives.

- **Building the Future** reflects foresight and predictive analytics. By modeling potential outcomes of proposed actions, this pillar empowers users to strategize and optimize future interventions. Building the Future is the visionary aspect of sustainability. It involves using predictive models to envision the consequences of today's actions on future generations. This forward-thinking approach is essential for planning sustainable development, as it allows us to simulate the long-term effects of various policies and interventions, striving for a balance between growth and environmental stewardship.

For the pillar of "Learning from the Past", this thesis presents two examples of applications: Land Consumption Assessment and Aboveground Biomass Estimation. These cases exemplify how integrating Remote Sensing data with Deep Learning algorithms facilitates the generation of both retrospective and up-to-date maps. The thesis proposes ReFuse and ReUse networks for the Land Consumption Assessment and the Aboveground Biomass and Carbon Storage Estimation, respectively. Through the Land Consumption Assessment, the thesis demonstrates the ability to trace the historical expansion of urban areas, quantifying the encroachment on natural landscapes over time. The Aboveground Biomass Estimation further showcases the utility of historical Remote Sensing data, combined with advanced deep learning techniques, in assessing changes in biomass across various ecosystems.

For the pillar of "Monitoring the Present," the thesis emphasizes the critical need to make data readily accessible to users through the concept of Data Space. This approach advocates for an environment where data can be easily shared, accessed, and utilized, facilitating real-time decision-making and monitoring. To operationalize this, the dissertation introduces the *GeoAI Processing Block*, a software solution designed to execute AI algorithms on Remote Sensing data and Open Data and interface seamlessly within the Data Spaces Ecosystem.

The thesis illustrates this module's capabilities by implementing a continuous monitoring service for Thermal Comfort specifically tailored for

soft mobility in Milan. This service leverages the GeoAI Processing Block to analyze various data streams, including satellite imagery and open data, to assess and map thermal comfort levels across the city. By doing so, it provides valuable insights that support the promotion of soft mobility solutions, such as cycling and walking, contributing to the city's sustainability goals and enhancing the well-being of its residents.

Finally, for the "Building the Future" pillar, the thesis proposes two approaches to meet the user's need for action planning with an eye on maximizing future impact and understanding the consequences of an action before its execution. The first approach revolves around the idea of learning from the past to understand where to take action to maximize the social impact. The second approach focuses on scenario analysis. Here, various 'what-if' scenarios are created and analyzed to understand the possible outcomes of different actions. The proposed approach uses Generative Adversarial Networks (GANs) to generate synthetic multispectral satellite imagery. This method is particularly useful for users who need to assess the impact of a proposed action before its implementation. By simulating different scenarios, users can gain insights into how various decisions might play out, helping them choose the course of action that best aligns with their desired future objectives.

Both approaches provide users with tools not only to envision but also strategically shape the future, ensuring that actions taken today are informed by an understanding of their potential effects.

1.4 Thesis structure

The structure of the thesis is depicted in the Figure 1.7

Learning from the Past (Chapter 2) delves into historical analysis through two main sections: Land Consumption (2.1) and Aboveground Biomass Estimation (2.2.1). Within the Aboveground Biomass Estimation, further specialized discussions are presented on ReUse (deep-learning) (2.2.3) and Model Transfer with Active Learning (2.2.4), showcasing different deep-learning methodologies applied to past data for ecological assessment.

Monitoring the Present (Chapter 3) focuses on current data analysis and introduces the GeoAI Processing Block (3.2), a tool designed for real-

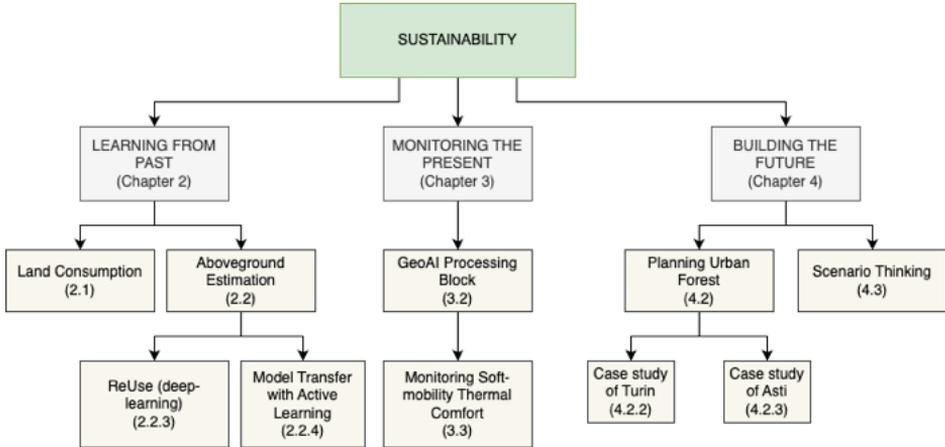


Figure 1.7: Thesis structure

time data assessment. A specific application of this tool is detailed in Monitoring Soft-mobility Thermal Comfort (3.3), highlighting the use of GeoAI Processing Block for urban environmental monitoring.

Building the Future (Chapter 4) projects into forward-looking planning and simulation. It encompasses Planning Urban Forest (4.2) as a strategy for urban greening and two case studies, Turin (4.2.2) and Asti (4.2.3), which provide practical examples of future-oriented urban planning. Additionally, Scenario Thinking (4.3) is outlined as a method for envisaging and evaluating potential future developments.

Each pillar and its corresponding sections are structured to reflect the multifaceted approach to achieving sustainability through the use of Earth observation data and Artificial Intelligence, encompassing historical evaluation, current condition monitoring, and future scenario planning.

Chapter 2

Learning from the Past

Whatever worked in the past, build on it; whatever didn't work in the past, break the chain that binds you to it.

Marianne Williamson

2.1 Land Consumption Assessment

2.1.1 Introduction

Human beings have been altering the face of the Earth for the last few centuries. This process has seen a fast boost after the introduction of machines, resulting in drastic changes to land cover. Identifying the physical aspect of the Earth's surface (land cover) and how we exploit the land (land use) is an essential task. Indeed, land-cover changes may significantly influence several processes that can eventually lead to the degradation of local ecosystems. By definition, Impervious Surfaces (IS) are artificial surfaces (such as roads, driveways, sidewalks, parking lots, and rooftops) through which water cannot infiltrate into the soil [34]. Rapid urbanization has greatly expanded urban impervious surfaces, decreasing previous surfaces, such as forests, green spaces, bare soils, and wetlands. Consequently, in recent years, *impervious surface analysis and monitoring* have emerged as an indicator of urbanization degree and a significant indi-

cator of environmental quality since IS cover quickly measures the impact of human activities on alterations of the environment. Therefore, accurate methods for determining impervious surface distribution are fundamentals for monitoring changes to urban areas and achieving sustainable urban development [35]. IS cover monitoring can be done through on-place surveys (by experienced and specialized personnel) or by analyzing satellite images (also known as remote sensing). Although carrying out on-place surveys produces more comprehensive and authoritative outcomes, performing it is an expensive and time-consuming process involving the movement of people and tools. Therefore, automating this process is extremely useful for reducing the amount of work and to limit the associated costs.

One of the aspects making automatic analysis non-trivial is that impervious areas are usually made of different construction materials, resulting in significantly variegated spectral signatures and spatial patterns [36]. As for many other computer vision tasks, in recent years, Deep-Learning (DL) algorithms, and in particular Convolutional Neural Networks (CNNs), have been showing promising results in land-cover classification [37, 38]. This has been firmly pushed by satellite images usually existing also in visible-light (RGB) channels, thus allowing for the leverage of DL models designed for tasks requiring the analysis of natural images. However, the situation is dramatically different regarding imperviousness analysis due to (i) the lack of high-quality labeled datasets and (ii) the fact that RGB spectral data are usually not the most suited for the task. In addition, current solutions providing land-cover maps through automatic algorithms show some issues in addressing the high granularity of the data in an urban context, losing the essential details needed for accurate analysis. To better understand the extent of this problem in urban areas, in Figure 2.1, different imperviousness maps are compared for the identical land segment encompassing a small area in the city center of Rome., Italy. In particular, Figure 2.1a shows the area as it appears from a very high-resolution (50 cm) Pleiades satellite image; Figure 2.1b shows the imperviousness map at 10 m resolution realized by on-plane analysis performed by the Italian National Institute for Environmental Protection and Research (ISPRA); Figure 2.1c reports a 10 m imperviousness map realized by the Environmental Systems Research Institute (ESRI) using artificial intelligence algorithms; finally, Figure 2.1d shows a 10 m imperviousness map currently available on the

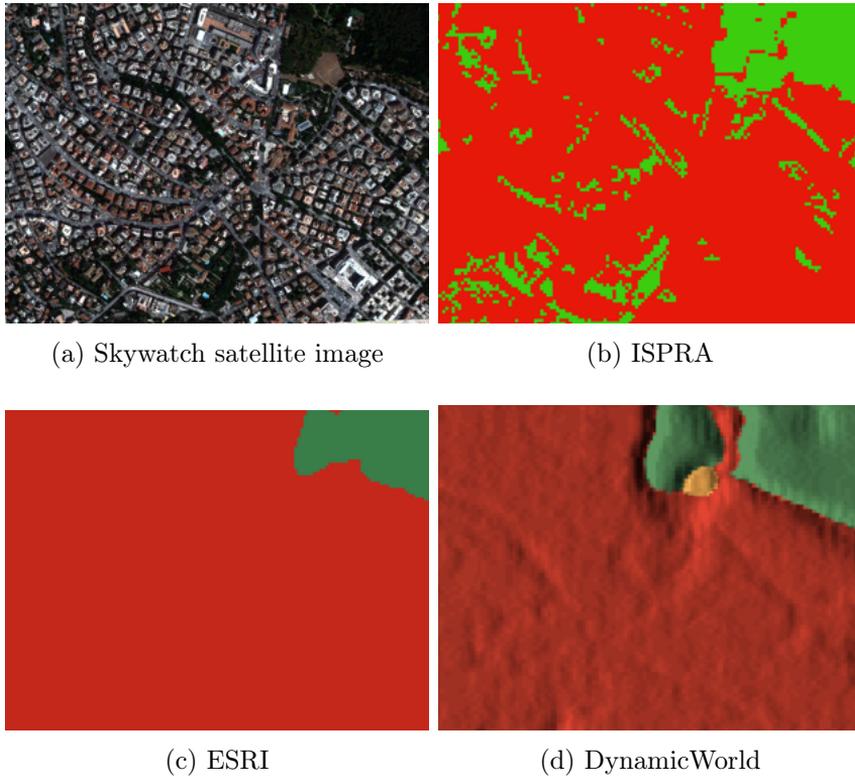


Figure 2.1: High-resolution satellite image of a dense urban area in Rome, Italy (2.1a), and the corresponding imperviousness maps from ISPR (2.1b), ESRI (2.1c) and DynamicWorld (2.1d). For imperviousness maps, red indicates impervious areas while green corresponds to pervious ones.

Google Earth Engine under the name of "DynamicWorld", automatically generated using deep learning on Sentinel-2 imagery [39].

Besides this, recent years have seen an enormous increase in web-based applications leveraging techniques derived from geographic information systems (GIS). Despite it should be good news as having access to spatial data as well as advanced mapping and spatial analysis over the internet is a critical point to pursue to reduce the distance between data, information, and decision-makers, it has often been observed that many map layers publicly available are accessible only by people skilled with GIS. This thesis addresses these problems by compiling a novel dataset, conducting experiments with a DL architecture tailored to leverage Sentinel-2 multi-spectral data, and incorporating the entire process into a proof-of-concept web application. In particular, the key contributions can be summarized as follows:

- To address the difficulties of the current solutions in the urban context, a new dataset has been gathered using an authoritative imperviousness map (ISPRA) as ground truth. ISPRA is a public institute, part of the Italian Ministry for the Environment, Territory, and Sea, promoting and supporting scientific, technical, and research functions as well as assessment, monitoring, and control activities. Among other functions, it provides several land cover maps, including the imperviousness map. These maps are produced semi-automatically from data provided by several European projects and authoritative data available for the Italian territory and are generated and released yearly;
 - Introducing *ReFuse*, a new DL architecture for impervious surface extraction based on a U-Net backbone [40], residual blocks (**Re**) [41] and the FuseNet principle (**Fuse**) [42] to take advantage of Sentinel-2 multi-spectral bands despite their different spatial resolution. Furthermore, a performance comparison was conducted between the proposed approach and several state-of-the-art CNNs.;
 - Imperviousness inference and visualization were incorporated into a GIS web application featuring a user-friendly interface tailored for individuals lacking specialized GIS expertise. The implementation involved an inference pipeline that utilized contemporary distributed
-

parallel computing and adhered to MLOPs best practices. This facilitates rapid deployment of the solution on HPC or cloud computing systems, ensuring robust scalability.

Section 2.1.2 reports some related works; the process of the dataset generation, the proposed DL approach for impervious surface extraction and the web tool are described in Section 2.1.3; Section 2.1.4 describes the experimental setup; Section 2.1.5 reports and analyses obtained results while Section 2.1.6 provides some conclusions and future perspectives.

2.1.2 Related Works

Computer vision and machine learning strongly contribute to satellite image classification. Focusing on remote-sensing methods for impervious surface extraction, machine-learning approaches can be divided into three groups: (i) pixel-based, (ii) texture-based, and (iii) semantic segmentation algorithms [43].

Pixel-wise classifiers typically exploit the spectral signature by relying on ad hoc features. Usually, they leverage a similarity measure to measure the spectral differences between impervious surfaces and other ground objects. Most commonly used indexes are Normalised Difference Built-up Index (NDBI) [44, 45], Normalised Difference Vegetation Index (NDVI) [44, 46], Index-based Built-up Index (IBI) [47], Normalised Impervious Surface Index (NISI) [48], Combination establishment index (CBI) [49] and Corrected Normalised Difference Impervious Surface Index (MNDISI) [50]. These methods are usually computationally not demanding. The flip side is that pixel-based approaches ignore spatial context information. This implies that they can easily be misled by the noise and the within-class variability, causing a salt-and-pepper effect within the classification result. Instead, the *texture-based approaches* do not rely on the spectral information of imagery but rather exploit spatial information among neighboring pixels to overcome the noise better and to capture different types of spatial structures. Given a set of features to consider, different classification methods have been used to divide pixels into impervious and permeable surfaces. Commonly used classifiers include Support Vector Machines [51, 52], artificial neural networks [53], decision trees [54] and random forest [55].

The wide variety of solutions so far described highlights the difficulty in finding the best combination of features that suit the classification task due to the high variability of impervious surfaces' appearance on remote-sensing imagery [34]. Therefore, this task can benefit from data-driven feature-learning approaches and end-to-end model training provided by *semantic segmentation* algorithms [56]. In the past decade, deep learning has proved to be effective on this task, with Convolutional Neural Networks (CNN) outperforming many traditional machine-learning solutions. Nearly all state-of-the-art architectures for semantic segmentation follow principles stated in [57], where semantic segmentation using Fully Convolutional Networks (FCN) is demonstrated to achieve impressive results. The main idea consists of modifying traditional CNN so that the output is no longer a probability vector but rather a probability map. That was possible by replacing standard fully connected layers of CNNs with fully convolutional layers to "densify" the single-vector output of a traditional CNN. A second feature was the use of transposed convolutions, also called deconvolutions. A deconvolution layer is used for up-sampling a feature map and obtaining a prediction of the same size as the input image [58]. The third feature was the skip connections to combine dense prediction at shallow layers and coarse predictions at deep layers, improving segmentation details.

Several architectures have been proposed along this line. The U-Net architecture [40], designed for biomedical image segmentation, introduced the encoder–decoder paradigm for up-sampling gradually from lower-size features to the original image size. Since then, almost all CNN models for semantic segmentation have some form of encoder–decoder structure. The encoder reduces the spatial resolution of the input and creates lower-resolution feature mappings that are highly effective at classifying objects. The decoder increases the resolution of the feature representations to create a full-resolution segmentation map. U-Net added several skip connections, which concatenate the feature maps of the encoder part with the mirrored feature maps in the contracting path.

In [59], Sun et al. experimented with a CNN to extract impervious surfaces through Worldview-2 and airborne LiDAR. The findings showed that 3D-CNN had more ability to extract features than SVM since it used pixel-level spatial information and texture. [60] uses a deep-learning ap-

proach to automatically extract impervious surfaces from WorldView 2 and Pléiades-1A datasets. In [61], the authors conducted a comparative study for the impervious surface estimation mixing optical and SAR data; experimental results indicated the effectiveness of the proposed deep convolutional network, which exhibited a better accuracy outperforming other benchmark methods. In [62], Fu et al. proposed a solution based on a deep CNN to map impervious surfaces in town–rural areas using China’s GF-2 Imagery. They showed the effectiveness of deep models and how transfer learning could significantly boost overall accuracy.

Fewer studies are available on Sentinel-2 imagery for built-up ISA extraction, and even fewer with a deep-learning approach. In [63], the authors assessed the feasibility of using Sentinel-2 images for this task through an artificial neural network. In [64], using Sentinel-2 satellite imagery, a CNN was employed as a deep feature extractor, and the classification was made using a random forest classifier. Similarly, in [65], the authors compared different machine-learning and deep-learning algorithms for the land-cover classification, with a CNN showing the best performance in classifying impervious areas.

Deep learning has also been investigated for multi-sensor and multi-modal image segmentation. Multi-modal fusion strategies are of great interest to RS classification since satellite images usually contain multi-spectral content. Similarly, extensive research has also been conducted to combine heterogeneous data (multi-sensor fusion), such as optical images with Synthetic Aperture Radar (SAR) and LiDAR data. In both cases, fusion models help reduce confusion from spectral heterogeneity in landscapes and enhance classification accuracy. On this line, in [66], the authors use existing CNNs (FCN or SegNet) as a base network to experiment with different data fusion strategies in early and late fusions.

In the same way, [67] explores how deep, fully convolutional networks can be modified to handle multi-modal and multi-scale remote-sensing data for semantic labeling. To this aim, the authors extended the FuseNet architecture [42] by considering two branches, one trained with IR-R-G bands and one with Normalised Digital Surface Model (NDSM), Digital Surface Model (DSM) and NDVI data. The proposed approaches outperformed a SegNet trained only on IR-R-G bands, thus proving the effectiveness of using multi-spectral data for remote-sensing classification.

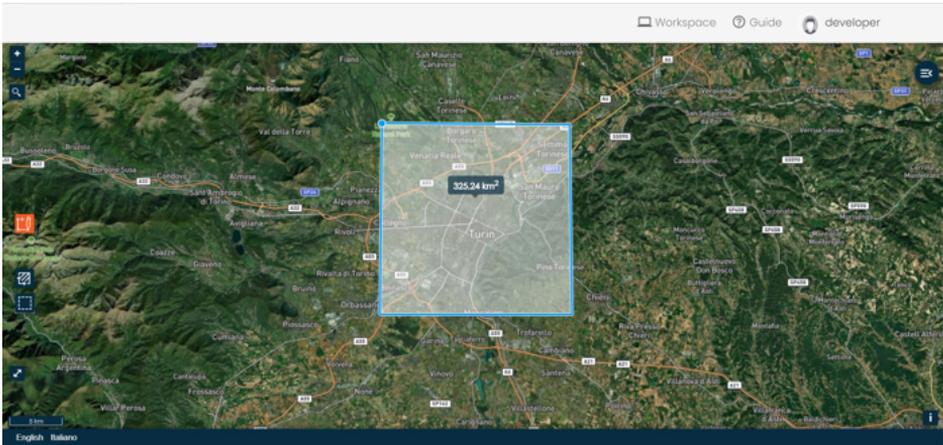


Figure 2.2: Example of the area of interest selection in the proposed tool.

2.1.3 Material and Methods

As described in Section 2.1, a web-based system that utilizes a new deep-learning model to generate imperviousness maps from Sentinel-2 satellite images is introduced. The resulting system has been integrated within a GIS web-based application to support non-expert operators in quickly generating imperviousness maps. Figure 2.2 shows the interface presented to the user for requesting a map of imperviousness in a specific part of the globe. Users only need to draw their area of interest and then select the date range in which they are interested. Once the area and date range have been chosen, a map calculation phase for the specific area starts. An asynchronous process handles the whole computation. All these processes happen in the background and are transparent to the user. Figure 2.3 reports a logical diagram of the inference pipeline: after the user requests, satellite images are collected from the stores and sent to the inference pipeline, which produces the imperviousness map as the final output. It is worth noting that such a web application requires an inference process that can handle and scale toward enormous quantities of satellite data.

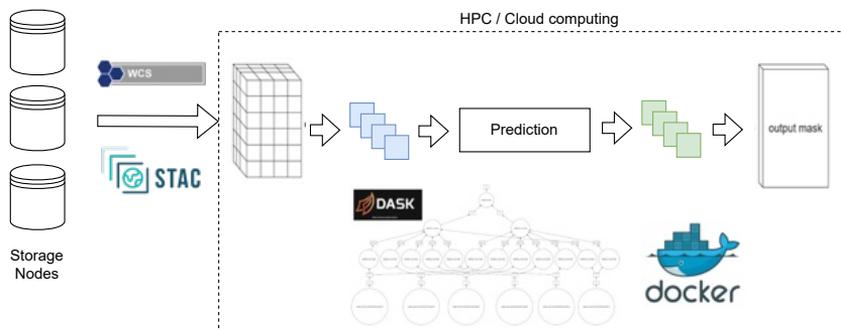


Figure 2.3: Inference Pipeline scheme. On the left is the process of the Sentinel-2 satellite imagery retrieval. On the right is the execution of the prediction. As detailed in Section 2.1.3, the inference is managed through a Docker image. Dask makes the solution scalable by parallelizing the execution.

Imperviousness Dataset Generation

This section describes the methodology to create the dataset to be utilized to train the deep-learning model for generating imperviousness maps from Sentinel-2 satellite images (dataset openly available in Zenodo at <https://doi.org/10.5281/zenodo.7058860>). The *Sentinel-2* platform consists of two satellites equipped with sensors able to acquire images with 13 spectral bands ranging from the visible range to the short-wave infrared. Each sensor has a different spatial resolution, with RGB bands and near-infrared (NIR) showing the highest (i.e., 10 m). With a 12-bit radiometric resolution, the image can be collected from 0 to 4095 possible light intensity values, enabling the satisfactory identification of minor variations in reflected or emitted energy.

The lack of labeled data seriously hinders the development of deep-learning algorithms that detect impervious surfaces. Most available imperviousness maps typically have a coarse spatial resolution that does not adequately emphasize green spaces in urban settings. The soil consumption map covering Italy provided by ISPRA for 2017 with a 10 m spatial resolution was used. These data were built by merging regional Land Use Land Cover (LULC) maps, in situ data provided to ISPRA by

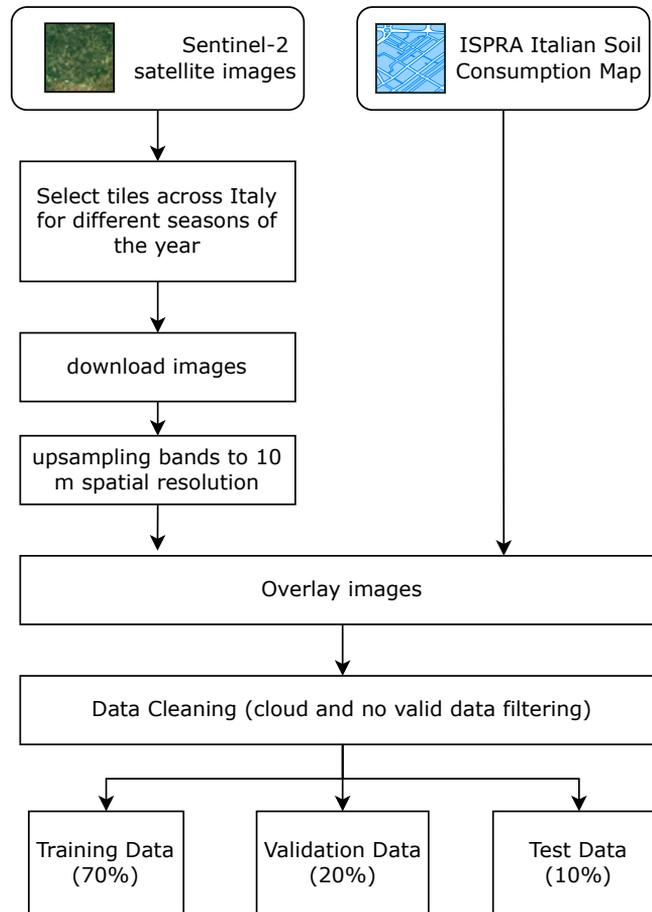


Figure 2.4: Schematic representation of the impervious surface dataset generation.

Regional Environmental Agencies, Copernicus HRL Imperviousness products, OpenStreetMap, and local supplementary data. The map consists of a hierarchical classification ID with each pixel having up to three digits class: the first digit, starting from the left, describes whether a point is consumed soil (i.e., a value of 1) or not consumed soil (i.e., a value of 2), while the other two optional digits can specify the class with more detail (e.g., 112 stands for “soil consumed by asphalt roads”). Since the goal is the segmentation between impervious and non-impervious surfaces, only the first digit was considered to label the data.

The flowchart in Figure 2.4 illustrates the process for generating labelled training, evaluation, and testing data. Sentinel-2 granules are organized in a tiling, partially overlapping grid. As for other image-processing domains, having a dataset representative of the population is crucial. Sentinel-2 tiles covering particular zones of the Italian peninsula containing a variegated distribution of hills, waters, plains, mountains, etc., both in urban and suburban realms, were selected. To increase the variance, images recorded throughout the year were extracted, acquiring at least one image per season. Nonetheless, it is worth noting that only images for 2017 were considered to minimize labelling mismatch with the available ground truth. Table 2.1 lists Sentinel-2 products used for the dataset generation, while Figure 2.5 shows selected patch distribution along the Italian peninsula.

Sentinel-2 bands with lower spatial resolution were up-sampled to 10 m per pixel using cubic convolution. When necessary, ground-truth raster and acquired Sentinel-2 images have also been reprojected to a common coordinate system to have two perfectly stackable files. It is worth noting that a typical Sentinel-2 tile has a size of $10,980 \times 10,980$ pixels for the 10 m spatial resolution. Since these dimensions are computationally infeasible, non-overlapping patches of 244×244 pixels from each ground-truth image were extracted. Finally, a data-cleaning process excludes chips without corresponding ground truth, i.e., chips containing no data values or chips with clouds. To this aim, a cloud mask is computed using the Scene Classification Layer within Sentinel-2 Level-2A products.

Table 2.1: List of Sentinel-2 products used to generate the dataset.

Month	Product Name Identifier
March	S2A_MSIL2A_20170329T095021_N0204 _R079_T33TVF_20170329T095024
June	S2A_MSIL2A_20170613T101031_N0205 _R022_T32TQR_20170613T101608
July	S2A_MSIL2A_20170720T100031_N0205 _R122_T33TUH_20170720T100027
July	S2A_MSIL2A_20170706T102021_N0205 _R065_T32TMQ_20170706T102301
August	S2A_MSIL2A_20170806T095031_N0205 _R079_T33TWF_20170806T095744
August	S2A_MSIL2A_20170802T101031_N0205 _R022_T32TQP_20170802T101051
October	S2A_MSIL2A_20171014T102021_N0205 _R065_T32TNR_20171014T102235
October	S2A_MSIL2A_20171015T095031_N0205 _R079_T33SVB_20171015T095944
October	S2A_MSIL2A_20171028T100121_N0206 _R122_T33TUG_20171028T134729
December	S2A_MSIL2A_20171220T101431_N0206 _R022_T32TPR_20171220T122057
December	S2A_MSIL2A_20171224T095421_N0206 _R079_T33TVG_20171224T122256



Figure 2.5: Localization of the Sentinel-2 tiles used in dataset generation.

Multi-Spectral Bands Fusion Network

Semantic labeling of satellite images requires a dense pixel-wise image classification. In this thesis, some popular neural networks' capabilities for pixel-wise extraction of impervious surfaces have been explored. However, it is crucial to consider that Remote-Sensing (RS) image data are more than a picture since they include electromagnetic wavelengths beyond the three RGB bands of natural images. A CNN can take an arbitrary number of spectral bands as input, modifying the depth of the first input convolutional layer. However, exploiting the multi-spectral content of RS images is not as simple as presenting more bands as input to the network. Indeed, increasing the number of spectral input bands eventually results in the need for more extensive training datasets, consisting of satellite images and corresponding ground-truth data, to avoid overfitting. Moreover, this approach does not allow leveraging *transfer learning*, which proved effective for remote-sensing image analysis [68]. In a multi-sensor setup, with more than three bands available, a possible approach is to employ two or more neural network branches to analyze some of the bands separately, fusing the features at a later stage in the network. However, the main drawback of this late fusion procedure is that the number of weights doubles, thus requiring more computation time for both the training and inference phase. Considering the extension of satellite images, this limitation might be too demanding to make the approach feasible in a real production environment.

The thesis thus focus on designing a CNN that can efficiently combine features from multi-spectral bands, providing a good balance between the number of parameters and segmentation performance. The proposed architecture, shown in Figure 2.6, comprises three state-of-the-art ideas:

1. A **U-Net** as backbone architecture. Its encoder-decoder paradigm with cross-connections for pixel-wise labeling and skip connections between same-sized parts in down-sampling\up-sampling paths help to address the loss of fine detail during up-sampling [40];
 2. To achieve better results as the depth of the network increases, *the building blocks of the standard U-Net encoder part were replaced with residual blocks*. More in detail, a ResNet-50 model has been used to replace the U-Net encoder down-sampling section. The idea is to
-

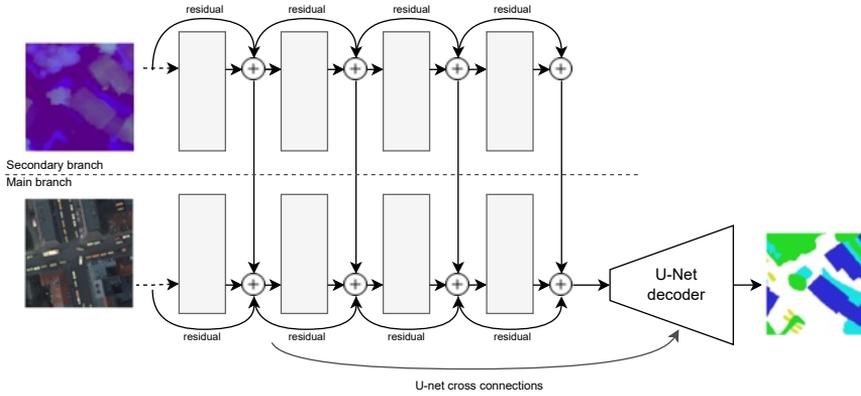


Figure 2.6: The proposed *ReFuse* encoder-decoder architecture: the encoder consists of two parallel ResNet-50 networks where every residual block output from the second branch is fused into the main one by a feature map summation. Feature maps from the encoder are then connected with the decoding part by feature map concatenation, such as in the U-Net cross-connections.

leverage residual blocks' ability to reduce the problems associated with the vanishing gradients strongly;

3. To exploit multi-spectral content beyond classic RGB wavelengths, a **FuseNet** [42] approach has been used. The FuseNet model jointly encodes the RGB and depth information using two encoders (in this case, two ResNet-50, as described in the previous point) whose contributions are summed after each convolutional block. The approach adapted the standard fusion method to residual networks by summing contributions from different branches after each residual block.

The result is a U-Net-like network with the encoder replaced by two parallel ResNet-50 networks where the main branch inputs the RGB bands while the second uses bands B07, B08, and B11. Every residual block output from the second branch is fused into the main branch using feature map summation. After that, the fused map is connected to a convolutional layer for the decoding part through concatenation, implementing the classical U-Net cross-connections. *All these solutions allow the proposed architecture*

to benefit from the combination of short (i.e., residual skip connections) and long skip connections (i.e., U-Net cross-connections) during the training strategy. This thesis refers to the approach with the name **ReFuse** after its two core components: **RE**sidual blocks and **Fuse**Net.

Some minor changes were applied to the ResNet-50 encoder, following the minor adjustments presented in [69]. At first, the down-sampling block of a ResNet-50 was modified, changing the stride of the first convolutional layer from 2 to 1. The idea is that a stride of 2 with a kernel size 1×1 ignores three-quarters of the input feature map. To leave the spatial size unchanged, the size of the strides in the first two convolutions was switched so that the second convolutional layer would have a stride of 2. Second, for the same reasons, the 2-stride convolution was replaced with an average pooling layer followed by a one-stride convolution, keeping intact the output dimensions: adding a 2×2 average pooling layer with a stride of 2 before the convolution 1×1 with a stride of 1 the model will not overlook any information. Finally, the first convolutional layer of the ResNet (i.e., a 7×7 convolution) was changed with three 3×3 convolutions since the replacement will make the model easier to train [69].

Finally, to leverage transfer learning, a three-step training approach was adopted: (i) the encoder weights are initialized from a ResNet50 pre-trained on ImageNet; (ii) the training is started with all the encoder layers frozen (with the aim of training only the decoder ones); (iii) the network is trained again by considering all the layers trainable. This procedure helps the network converge while giving the decoder enough information to produce helpful segmentation masks properly.

Distributed Inference Pipeline

Serving a model in production is crucial, especially when serving a web-based application. Building highly available, scalable, distributed systems for machine-learning data pipelines is complex. In this thesis, the whole process was managed as a workflow, a sequence of tasks representing units of business logic. Figure 2.7 shows the workflow implemented for the prediction pipeline. The workflow is a Directed Cyclic Graph (DAG) where each node is a task and arrows are dependencies between tasks. In the extraction step, from the AoI and date range provided by the user, a first task searches and downloads Sentinel-2 data. Any images with clouds are

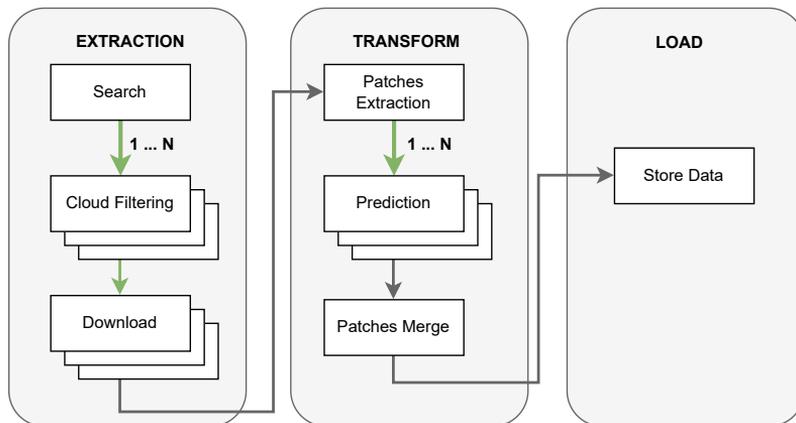


Figure 2.7: Schematic representation of the inference pipeline workflow.

appropriately removed from the process. The prediction is performed on the downloaded images after a subdivision into patches of the size required by the deep-learning model. Finally, the map is assembled and stored for later visualization by merging predicted patches.

The whole approach was implemented through a Workflow Management System (WfMS) (Source is available at <https://github.com/priamus-lab/ReFuse>). A WfMS provides an infrastructure for a defined workflow's setup, performance, and monitoring. Therefore, involving a WfMS gives flexibility and extendibility to the approach because it implements off-the-shelf features such as data-sharing between tasks, recovery from failure, retrying failed tasks, task scheduling for batch runs, flow versioning, and more. In this thesis, Prefect (<https://www.prefect.io/> accessed on 31 July 2022), an open-source orchestration and workflow tool, was used. However, despite the choice, this thesis aims to show a general approach that is easily replicable with similar tools. Five essential layers can be highlighted in every WfMS (Figure 2.8): (i) the *orchestration layer*, which is responsible for the workflow's life cycle; (ii) the *agents*, daemon-like processes that look for tasks and run them if any are available; (iii) the *flow run layer*, which is responsible for defining where the flow run; (iv) the *execution layer*, which manages where and how single tasks within the flow run; (v) the *storage layer*, which defines where the flow's code is stored,

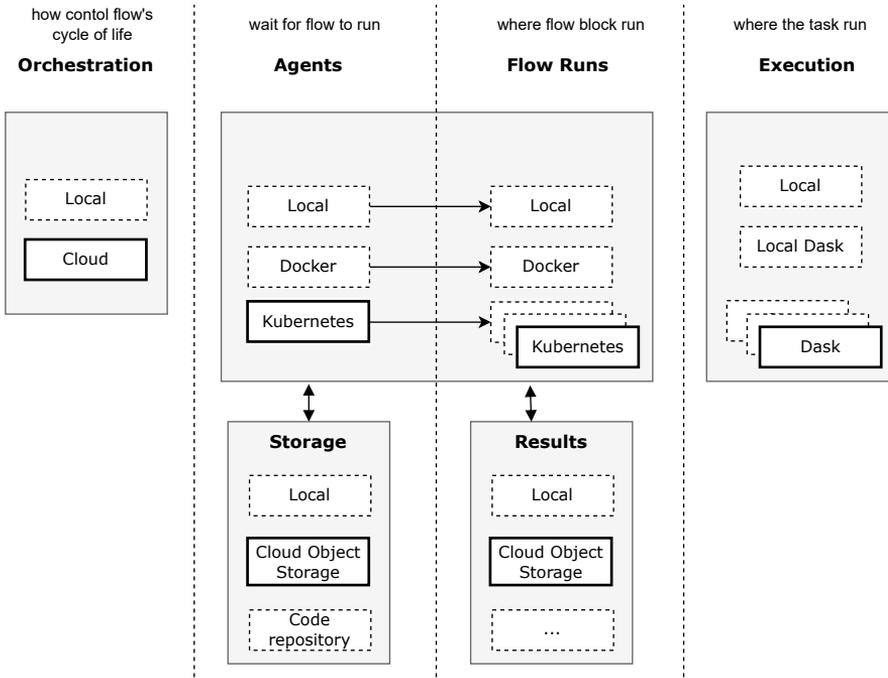


Figure 2.8: Workflow Management System layers. In bold, the choices made in this thesis.

to be collected when needed. In addition, when tasks need to pass data between them, a *result layer* defines and manages where to store task results.

It is worth pointing out that in a fully distributed system such as the one defined here, the layers and components can be chosen from different types and deployed elsewhere. This aspect makes the solution extremely flexible: for Example, in a development and test phase, one could choose to deploy an agent on a personal computer but let the flow run in the cloud. Figure 2.8 highlights in bold choices made in this thesis. Both agents and flows ran on a Kubernetes cluster; Prefect Cloud, a cloud-managed service, performed the orchestration. Flow’s code and task results were stored in an object-based storage, particularly an AWS S3 bucket. Kubernetes is a system that manages containers where a container could be explained as

a lightweight virtual machine. Containers encapsulate an application with all its dependencies, including system libraries, binaries, and configuration files, making it portable across different hosts. Kubernetes can automatically create and scale these containerized applications and manage storage among all the containers. Instead, object-based storage is a strategy that manages and manipulates data storage as distinct units called objects. Data are kept in separate storehouses versus files in folders and are bundled with associated metadata and a unique identifier to form a storage pool. Object-based storage effectively manages unstructured distributed content, such as this use case. This solution is adaptable to different scenarios because the business logic is separate from the execution methods. Changing the configuration of one of the components, e.g., where to store the code or execute the flow, does not require any change to the business logic of the tasks. A Dask cluster was used for the execution layer. This is the most critical choice in the pipeline, as the choice of this executor allowed us to parallelize task execution and potentially scale the approach indefinitely in a distributed environment. Dask is an open-source Python library for parallel computing. In particular, a temporary Dask cluster was created, i.e., a cluster that scales up and down when needed and executes tasks on them. Despite an initial latency time to start the cluster, an ephemeral cluster allows for the leveraging of several machines but releases them when the workload is complete. The WfMS orchestrated all the tasks together, respecting dependencies and data flow between them. Figure ?? highlights parallelizable computations within the inference pipeline. It shows with green arrows the outputs of tasks that produce a list of elements on which subsequent tasks can proceed in parallel. The designed workflow involves parallelization in processing the identified Sentinel-2 images and inferring over the patches into which the image is divided. For example, if the initial search step identifies 100 Sentinel-2 images, subsequent processing can proceed in parallel with these. A workflow like this can be optimized with a *map-reduce* approach. Map-reduce is a powerful two-stage programming paradigm famous in the big-data ecosystem that can be used to distribute and parallelize work (the “map” phase) before collecting and processing all the results (the “reduce” phase). Tasks can be executed dynamically across an iterable input with a map-reduce model. This, in turn, allows us to execute mapped tasks in a distributed and par-

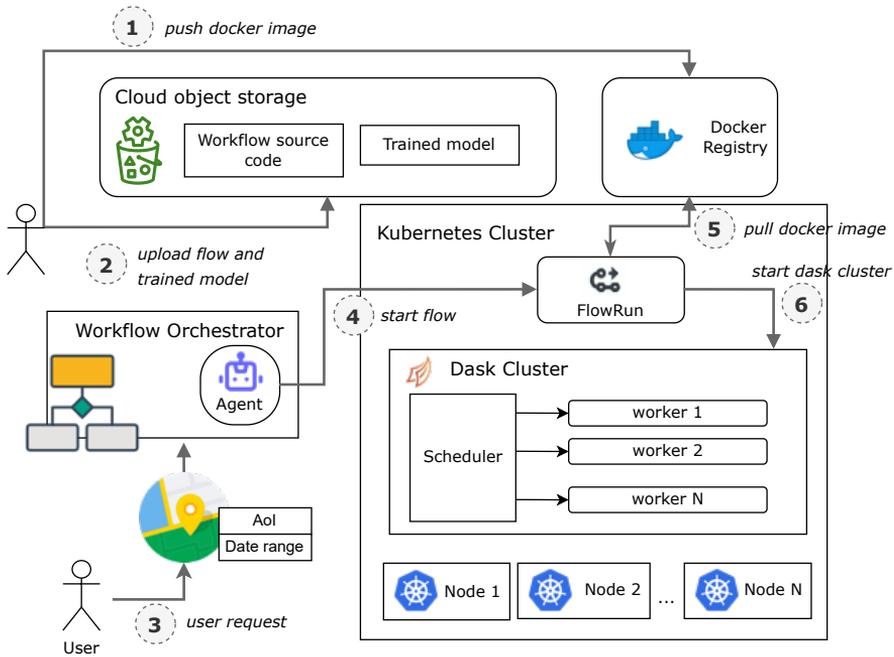


Figure 2.9: Flow execution process for the imperviousness map generation. A Docker image stores library dependencies. The model, weights, and flow code are stored in object-based storage. During flow execution, the orchestrator pulls the image from the Docker registry and, using it, starts a Dask cluster. Finally, the prediction flow is executed on the Dask cluster.

allel manner on a Dask cluster, drastically reducing the total execution time.

Figure 2.9 depicts the whole process. The first step is releasing the Docker image for the flow execution into a Docker registry. A Docker image is an immutable template file containing the source code, libraries, dependencies, tools, and other files needed to create a container where the application will run. In this way, nodes that execute the inference pipeline will have all software dependencies correctly in place. The model-serving strategy adopted is straightforward: together with weights, the model was stored in object-based storage. This method permits the download of the model for inference using a URL accessible via the Internet, a mandatory

requirement for a distributed data pipeline. In addition, it enables a fast and easy replacement with newer versions because the model and business logic are decoupled. Finally, these steps can easily be automatized within a Continuous Deployment flow, e.g., by initiating automatic uploads following a code committed into a code repository. At the user’s request to calculate the imperviousness map on a new area, a request is sent for a new execution of the inference flow. The agent, therefore, upon receiving the request, starts the flow. The flow is executed within a Kubernetes cluster as a Kubernetes job. A Kubernetes job is a workload controller that performs one or more finite tasks in the cluster. At startup, the flow pulls the Docker image from the Docker registry for machine instantiation and then deploys an ephemeral Dask cluster. After the Dask cluster is up, flow tasks can execute appropriately on the cluster. A Dask cluster is composed of one scheduler node and N worker nodes. By increasing the number of workers, the number of maximum tasks executable in parallel can be scaled up, giving the solution great flexibility and scalability. Although not mandatory, such a solution fits well with the serverless infrastructure made available by most cloud providers today. Serverless computing is an execution model in which the cloud provider allocates machine resources on demand, allowing customers to pay only when computational power is needed. Creating a Dask cluster when required and deploying it on a serverless infrastructure dramatically reduces operational costs while maintaining a virtually infinite ability to scale. It is worth remembering that the type of machine used to instantiate Dask workers can be defined during configuration. For Example, machines with GPU enabled can be selected to reduce inference time further.

2.1.4 Experimental Setup

The network was implemented in PyTorch, and the training execution was on ad hoc AWS EC2 instances. Models were trained using the Adam optimizer [70] with $\beta_1 = 0.9$, $\beta_2 = 0.99$, $\epsilon = 1^{-5}$, a weight decay of 1^{-2} and a batch size of 8. The considered loss function is a combination of dice and pixel-wise cross-entropy loss:

$$L_{total} = L_{dice} + L_{ce}$$

with

$$L_{ce} = -(y \log(p) + (1 - y) \log(1 - p))$$

$$L_{dice} = 1 - \frac{2|A \cap B|}{|A| + |B|}$$

where A is the predicted segmentation, mask B is the ground truth, y is the ground truth, and p is the probability for that class. The choice of also including the dice loss is to help to regularise results in the case of unbalanced data chips (e.g., 95% bare soil and 5% impervious surface in a single chip).

One of the most critical hyperparameters is the learning rate (LR): a big LR causes the model to diverge, while a small LR causes the model to converge slowly or stack in an unfavorable solution. Two methods were used, the *learning rate range test* [71] and the *one-cycle policy* [3], to determine the ideal LR value and train the model for accurate fine-tuning. The LR range test is a method for determining what are considered to be the acceptable minimum and maximum boundary values for LR. It entails running the model over several epochs while letting the LR rise linearly between low and high LR values after each mini-batch until the loss value increases enormously. When plotting accuracy trends that vary LR values, the LR one order lower than the point where loss is minimum can be chosen. Figure 2.10a shows the ReFuse model learning range test output: a maximum LR of 1^{-3} was adopted for the experiments. On the other hand, the one-cycle policy is a technique similar to the simulated annealing algorithm [72] designed for varying the LR during the training. The method needs an initial interval of values: the maximum value using the LR range test and the lower one as 1/10th of the maximum LR was chosen. The algorithm moves from the lower to the higher value during the first half of the cycle and from the higher back to the lower during the latter half of the cycle. Finally, in the last few iterations, this method anneals the learning rate way below the lower learning rate value for final fine-tuning. Conventionally, the learning rate decreases as the learning starts to converge, but the idea behind this approach is that a higher learning rate may help overcome saddle points. In [3], the author shows that when the learning rate is higher, during the middle of learning, the learning rate works as a regularisation method and keeps the network from overfitting because it avoids steep areas of loss and finds a better flatter minimum. Figure 2.10b shows LR values used during training across the considered

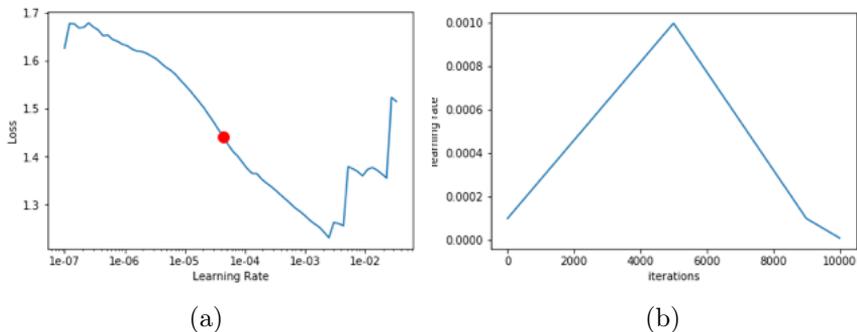


Figure 2.10: Learning Rate (LR) tuning analysis for the ReFuse model. (a) The output of the LR range test. The figure highlights the maximum learning rate before loss increases. (b) Learning rate values during training with the one-cycle policy [3]. The one-cycle strategy increases the learning rate from an initial value to a maximum learning rate. It then decreases the learning rate from the maximum value to some minimum, substantially lower than the starting learning rate.

iterations. When the model performance on the validation dataset ceased improving, an early-stopping criterion was used to stop the training.

The GIS web application and the inference pipeline were executed on cloud infrastructure using a Kubernetes cluster. The inference pipeline was released on a serverless infrastructure to reduce the solution’s running costs.

2.1.5 Results

The solution was compared against some state-of-the-art CNNs and variants to measure the effectiveness of the proposed approach. In particular, the comparison was made against some variants of three main architectures:

- A FCN-8s, a variant of an FCN model introduced in Section ??.
- The classical FCN architecture consists of a series of convolutional and pooling layers, with FCN-8s also implementing a fusing strategy

between predictions of the shallower layer Pool3 with the twice-up-sampled sum of the two predictions derived from Pool4 and the last layer. The stride 8 predictions are subsequently up-sampled back to the image;

- A standard U-Net using different pre-trained CNNs as encoders. In particular, the thesis explored the use of VGG16 [73], ResNet [41], and EfficientNet [74] architectures pre-trained on ImageNet. The reasons behind this choice are the high generalization ability demonstrated over the years by VGG, the ability to deal with gradient vanishing of ResNet, and the high efficiency/performance trade-off of EfficientNet;
- The thesis used also DeepLabv3+ [75], an architecture introducing changes to the encoder–decoder structure, such as using dilated convolutions [76], to preserve most of the spatial input information.

These architectures were tested with different bands as input (i.e., RGB, RGB + NIR, all 13 bands resized to have the same spatial resolution). All the experiments used, as evaluation metrics, the pixel-wise segmentation accuracy and the Intersection over Union (IoU), defined as:

$$IoU = |Predicted \cap GT| / |Predicted \cup GT|$$

where GT is the ground-truth imperviousness mask.

Table 2.2 reports the results for the considered analysis, reporting for each configuration the base architecture, the used encoder (if any), the used input bands, and the obtained performance. The table shows that the proposed approach outperforms all competitors by a large margin. Moreover, analyzing the table, there are a few points worth highlighting:

- Models trained on RGB bands and using transfer learning (with weights pre-trained on ImageNet) tend to perform better than those using a different combination for the bands;
- Results obtained using all bands at 10 m spatial resolution, i.e., RGB and NIR, are slightly comparable with results obtained with RGB bands only;

Table 2.2: Results obtained by the proposed approach (last row) and the considered competitors for the imperviousness map extraction. For the proposed approach (ReFuse), using brackets under the “bands” column highlights the ability of the proposed approach to use different band types without the need for resizing.

Network	Encoder	Bands	Accuracy	IoU
FCN-8s	-	R, G, B	89.54%	69.76%
FCN-8s	-	R, G, B, NIR	88.25%	69.55%
FCN-8s	-	All 13 bands	84.80%	60.35%
U-Net	VGG16	R, G, B	87.45%	70.03%
U-Net	ResNet-34	R, G, B	90.13%	70.54%
U-Net	ResNet-50	R, G, B	92.39%	73.50%
U-Net	ResNet-50	R, G, B, NIR	92.07%	71.37%
U-Net	ResNet-50	All 13 bands	89.37%	70.32%
U-Net	ResNet-101	R, G, B	90.39%	70.57%
U-Net	EfficientNetB7	R, G, B	94.48%	74.61%
DeepLabv3+	ResNet-50	R, G, B	92.19%	71.35%
DeepLabv3+	ResNet-50	R, G, B, NIR	91.32%	71.29%
DeepLabv3+	ResNet-50	All 13 bands	88.25%	68.50%
ReFuse	ResNet-50	(R, G, B) + (B7, B8, B11)	95.72%	75.85%

- ResNet-50 tends to be the most effective encoder;
- Using all 13 bands causes the worst results both in terms of accuracy and IoU. This confirms the claim that using more bands does not necessarily result in better performance.

To better understand the effectiveness of the proposed approach, Figure 2.11 reports an inference example on a single patch, highlighting false positives (FP) and false negatives (FN). Interestingly, ReFuse produces a segmentation mask that is very close to the ground truth, with errors almost all located along the borders of urban areas or in shadow zones (misclassified as impervious areas). Similarly, Figure 2.12 reports the extraction results from the test set of two representative regions in the city of

Turin in Italy, including small, dense residential buildings and large high buildings in urban commercial areas. For both areas, the model produces excellent segmentation masks. Both images in Figure 2.12d,h reveal some difficulty in segmenting small buildings and noise along the edges of streets. This result can be explained in the first instance by the spatial resolution of the Sentinel-2 data. A spatial resolution of 10 m is too low to capture such levels of detail. Indeed, even observing Sentinel-2 images (Figure 2.12b–f) by the naked eye, there is the same difficulty in discriminating between impervious and non-impervious pixels.

It is worth noting, finally, the presence of mislabelled data shown in Figure 2.12c. Despite the presence of trees and gardens, everything is labeled as impervious. Even more interesting here is the result obtained. Despite errors in the ground truth, the network correctly detects the green infrastructure within the area, i.e., trees, grass, and parks, showing an ability to generalize the result that is sometimes superior to that of the training data. This result is of fundamental interest in this thesis because, as stated in Section 2.1, one of the objectives was to address the difficulties of current solutions in extracting imperviousness maps with a high level of detail within the urban context.

2.1.6 Conclusion

In this section, a deep-learning-based method for extracting imperviousness maps from multi-spectral Sentinel-2 pictures was introduced, leveraging bands with different spatial resolutions without the need for rescaling or other adaptations. The proposed approach has also been made available through a portable and scalable inference pipeline, easily pluggable within a web-based GIS application. The aim is to support the generation of imperviousness maps as soon as new satellite images are available for a fast, effective, and reliable analysis of human environmental impact. To achieve this, one of the biggest challenges was the lack of a labeled dataset with temporal and spatial granularity and precision suited for the task. To address this problem, the proposed approach gathered a new dataset using the ISPRA imperviousness map as the ground-truth raster. The proposed approach is a deep-learning architecture designed for impervious surface extraction based on a U-Net backbone and leveraging residual blocks and the FuseNet principles (here, the name) to effectively take advantage of

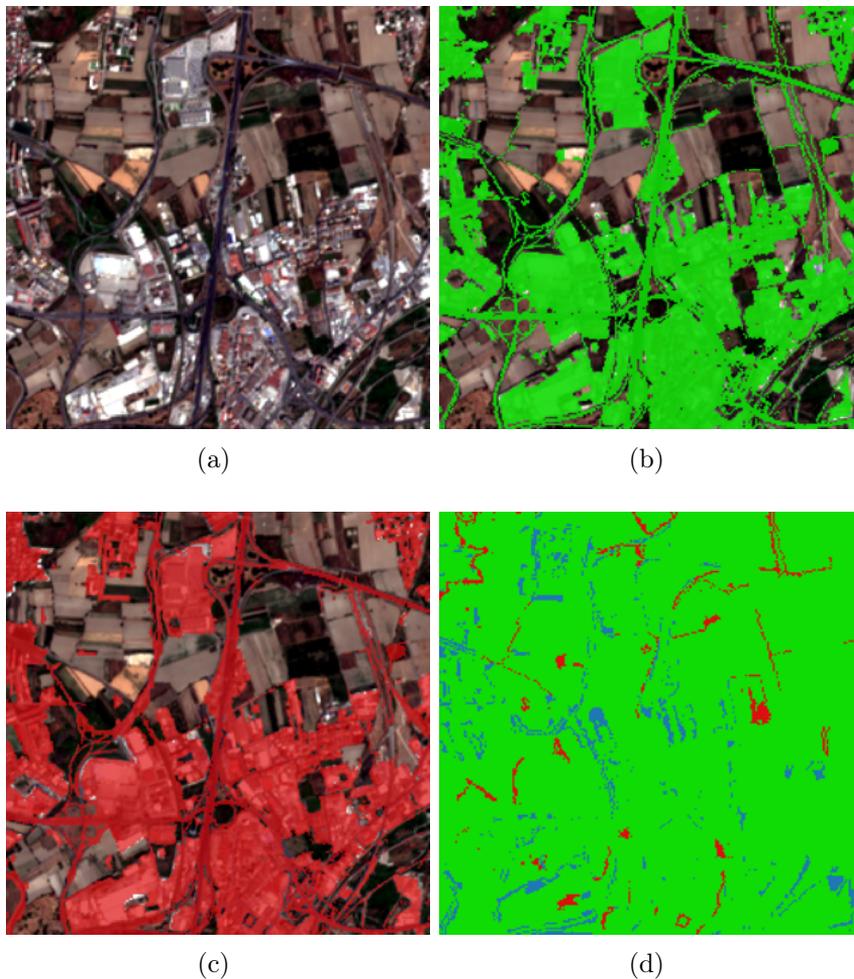


Figure 2.11: Example of an imperviousness mask obtained using ReFuse. The top left is the considered path (only RGB bands have been reported for visualization reasons). The top right is the corresponding ISPRA Ground Truth. The bottom left is the imperviousness mask obtained by the proposed ReFuse architecture. The bottom right is the ReFuse segmentation errors. (a) RGB patch. (b) Ground truth. (c) ReFuse. (d) Error mask: FP (blue), FN (red).

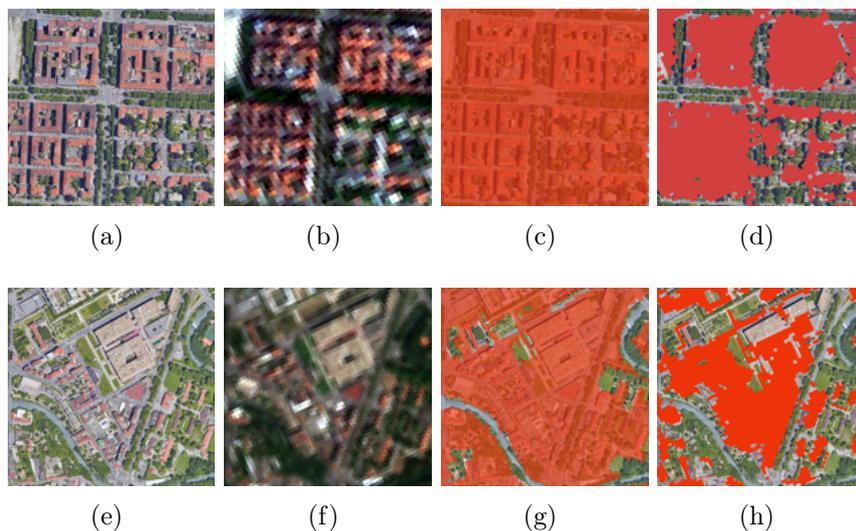


Figure 2.12: Semantic mapping with ReFuse model for two areas in the city of Turin. (a,e) Optical images extracted from Google Earth. (b,f) Sentinel-2 RGB images. (c,g) Ground truth extracted from ISPRA Soil consumption map. In red, imperviousness surfaces. (d,h) Predicted maps with the ReFuse model. Despite errors in the ground-truth labels (see image (c)), the network correctly detects the green infrastructure within the area, i.e., trees, grass, and parks (see image (d)).

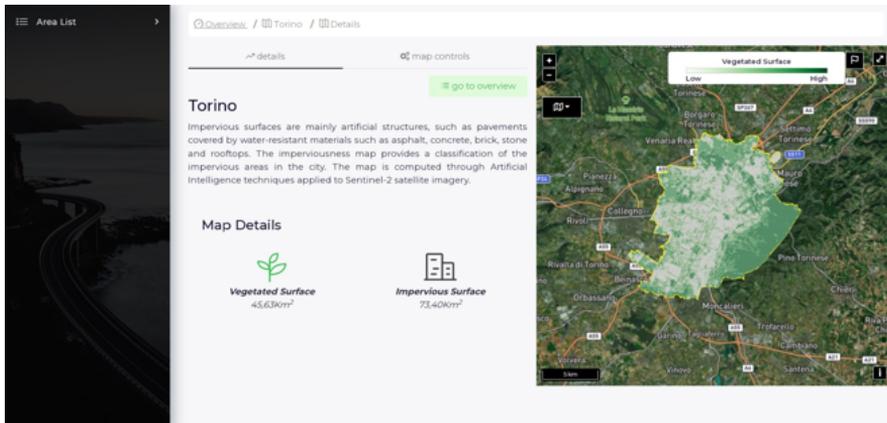


Figure 2.13: Example of the imperviousness map visualization into a GIS web application.

Sentinel-2 multi-spectral bands despite their different spatial resolutions. To evaluate the effectiveness of the proposed approach, the thesis compared the performance of the ReFuse architecture against some state-of-the-art CNNs. For the sake of completeness, the effectiveness of some variants obtained by changing the used encoder and/or considered bands were analyzed. Results show how using residual connections and the selected bands ensures the best performance. Nonetheless, all the considered competitors are outperformed by ReFuse.

Additionally, as the thesis aims to realize a simple and effective tool, the proposed approach was integrated into a GIS web application. Figure 2.2 shows how requesting data for a specific area of interest can be simple, even for non-experts. Figure 2.13 shows the implemented web page where the impervious data are presented to the end user using a map with a hexagonal grid. The grid was computed using a zonal statistics process applied to the imperviousness map computed through the proposed approach. The whiter areas represent the areas with the most impervious surfaces; vice versa, greener hexagons are the zones with the highest presence of trees, parks, and gardens. The image shows the map of impervious surfaces for the city of Turin in Italy. In particular, it demonstrates how the areas with the highest content of impervious surfaces in the city center can be easily identified, i.e., white hexagons.

In conclusion, the proposed approach shows how deep learning, MLOPS, and web-based applications can effectively be used simply and intuitively for a social-good application, such as imperviousness classification. The applicability of the proposed approach to different land-cover classification tasks will be tested to analyze its generalization capability across various applications. Finally, other experiments will be conducted, considering the time variable and extending the proposed dataset with new data sources, such as Sentinel-1 SAR satellite imagery.

2.2 Aboveground Biomass and Carbon Storage Estimation

2.2.1 Introduction

Forests cover about 40% of the Earth's ice-free land surface, although, recently, vast forested areas have been converted to agricultural or urban land uses. They constitute the vital core of the whole planet because they are the sentinels of the carbon cycle and the custodians of its biodiversity [77]. The expected increase in the global population is a further threat, representing a potential trigger for increasing deforestation and overexploitation of resources. The challenge to be faced by both politicians and scientists is the definition of protocols and strategies that allow for the sustainable management of such critical areas [78]. In this context, it is essential to support decision-makers with timely information about several biophysical parameters to implement adequate planning activities [79]. Among them, the aboveground biomass (AGB) is particularly interesting due to its impact on the carbon cycle and, consequently, on the global climate and its changes [80]. Accurate assessment of AGB, which in this thesis is defined as the mass expressed as oven-dry weight of the woody parts (stem, bark, branches, and twigs) of all living trees, excluding stump and roots, and related carbon stock, is essential for the sustainable management of forests. Recently, the United Nations Framework Convention for Climate Change (UNFCCC) established the Reducing Emissions from Deforestation and Forest Degradation (REDD+) that requires countries to report their carbon emissions and sink estimates through national greenhouse gas inventories (NGHGI) [81][82]. Furthermore, Sustainable Devel-

opment Goal 15 aims to "protect, restore and promote the sustainable use of terrestrial ecosystems, sustainably manage forests, combat desertification, halt and reverse land degradation and halt biodiversity loss"[83]. Subsequently, it is paramount to conduct an explicit investigation into the methods and procedures for quantifying carbon sinks. Generally, the aboveground dry biomass holds about 50%

Traditionally, this parameter is estimated through field sampling [84] or the exploitation of remote sensing data [85]. Undoubtedly, traditional methods are more accurate [86]. Still, they are way more time-consuming, laborious, unfeasible for forest-scale applications, challenging to implement in inaccessible areas, and destructive in nature[87]. Conversely, remote sensing technologies can deliver synoptic pictures of large areas with high revisit time. Therefore, as widely demonstrated by the literature, they represent a unique data source to feed methodologies, allowing for AGB estimation at a large scale [88][89][90]. Among all, the Sentinel-2 satellite system, launched by the European Space Agency (ESA) and equipped with a multispectral instrument (MSI), presents a high potential for applications in land management, agricultural industry (food security), forestry (AGB) disaster control, and humanitarian relief operations [91]. Sentinel-2 is a polar-orbiting satellite system comprised of two satellites, each carrying an MSI characterized by a 290-km swath width, offering a multi-purpose design of 13 spectral bands traversing from visible and near-infrared (NIR) wavelengths to shortwave infra-red wavelengths at refined (10, 20 m) and coarse (60 m) spatial resolution. Furthermore, the presence of four bands within the red-edge region, centered at 705 (band 5), 740 (band 6), 783 (band 7), and 865 nm (band 8a), gives the satellite system the potential for mapping various vegetation characteristics [92].

This section introduces a deep-learning approach for estimating Above-ground Biomass (AGB) using high-resolution Sentinel-2 satellite imagery. The architecture called the *ReUse* framework is grounded in a pixel-wise regressive U-Net model. This U-Net model is proficient in generating pixel masks for AGB predictions, offering substantial computational advantages, particularly when monitoring expansive geographical regions.

One notable advantage of the *ReUse* framework is its contrast with classical machine learning algorithms. Traditional methods necessitate a labor-intensive feature extraction process to derive indices that capture

spectral and spatial information content [4, 5]. Similarly, solutions relying on convolutional neural networks (CNNs) for AGB estimation, such as those utilizing the commercial Worldview-2 satellite and visible spectrum images captured by unmanned aerial vehicles [93, 94], typically produce a solitary AGB prediction value. In contrast, the U-Net model stands out due to its computational efficiency.

Compared to simpler CNNs, the U-Net model's distinctive strength lies in its operational approach. While CNNs leverage a pixel's neighborhood and associated spectral bands to generate a single AGB prediction per input pixel, the U-Net framework operates differently. It associates a patch of input pixels with an equally sized patch of output pixels, effectively producing a comprehensive map of AGB estimations as its output. This fundamental distinction enhances its utility for large-scale AGB monitoring and spatial analysis.

However, the significance of fine-tuning Aboveground Biomass (AGB) models for new and distinct geographical areas cannot be underestimated, primarily due to the critical challenge of generalization. While developing robust AGB estimation models is imperative, effectively applying them to diverse regions or across different periods presents substantial hurdles.

Ecosystems and forest types exhibit profound variations from one location to another. Differences in environmental conditions, species compositions, and vegetation structures are often substantial. Tailoring models for specific geographic areas enables them to better adapt to these regional intricacies, resulting in more precise AGB predictions.

Forests are dynamic environments characterized by continuous changes in vegetation density, growth, and decay over time. Moreover, models calibrated for a specific period may struggle to capture the evolving dynamics of AGB in subsequent years. Fine-tuning ensures the model's continued relevance and effectiveness over time.

Lastly, many regions face limitations in the availability of ground truth data, which can introduce sampling bias during model training. Fine-tuning facilitates the incorporation of region-specific reference data, reducing the risk of bias and enhancing the model's overall reliability. The problem of updating an existing model to make it work when applied to other areas or the same area at different times has been scarcely addressed.

In this regard, the following sections introduce a strategy to tackle the

typical generalization issues in the past AGB estimation literature. In particular, active learning principles [95] are exploited to build a novel workflow suitable for transferring a calibrated regression model to new areas by selecting relevant re-calibration samples. To this end, a novel technique for selecting the most informative samples based on their entropy rank is introduced. Unlike commonly adopted solutions, in which training data are distributed over large areas and used to train a single model, the principal novelty here is the exploitation of several regression models generated from a reference one, retrained with few samples collected on specific areas. The aim is to extend the forecasting capability of an existing training dataset by adding a limited number of samples made available by new acquisitions through active learning. Finally, particular attention is paid to validation. This is one of the weaknesses found in the literature, which tends to assess the performance of the estimation over vast areas using a few tens of plots. The results have been evaluated against extended ground truth, accounting for more than 850,000 validation points derived from the LiDAR processing of data acquired over Finnish forests.

This section is organized as follows. Related works are introduced in 2.2.2. Section 2.2.3 illustrates ReUse. Section 2.2.4 illustrates the active learning methodology for AGB estimation. Conclusion are drawn in 2.2.5.

2.2.2 Related Works

The literature on remote sensing for vegetation detection is vast. However, Aboveground Biomass detection is a different and more crucial task, especially for detailed estimation of other indicators, such as the carbon sequestration ability of a portion of land. Thus, this section focuses solely on works to estimate aboveground biomass and carbon storage.

To this end, several techniques have been proposed in the literature. Regression-based models are among the most exploited [96][97][98]. They have many advantages, like the possibility of combining heterogeneous data, the low computational complexity, and the suitability for dealing with sparse sampling, which is the typical case in forest studies [84]. Machine learning methodologies have been exploited to better deal with nonlinearities between measurements and remote sensing data. Previous studies demonstrated the effectiveness of solutions like random forests [99], gradient boost algorithms [100], and support vector machines [101].

In [102], the authors attempted to examine the prospects of Sentinel-2 spectral data for quantifying the carbon stock in a reforested urban landscape, using Random Forest and, as predictors, 10 Sentinel bands plus 15 spectral indices that summarise the spectral content without taking into account in any way the spatial correlations typical of an image. Similarly, in [4], the authors applied Sentinel-2 satellite images combined with field-measured biomass using Random Forest (RF) to estimate aboveground biomass in Yok Don National Park, Vietnam. 132 spectral and texture variables were extracted from the Sentinel-2 images; the grey level co-occurrence matrix (GLCM) method [103] was used to compute the texture variables.

In [104], Sentinel-2 performance was evaluated for a buffer zone community forest in Parsa National Park, Nepal, using field-based AGB as a dependent variable, as well as spectral band values and spectral-derived vegetation indices as independent variables in the Random Forest algorithm; in this approach, no features were extracted from the spatial dimensions, but indicators were only extracted from the spectral dimension of the input tensor. Spectral bands, vegetation indices (VIs), and texture variables derived from processed Sentinel-2 data and topographic parameters have been used in [5] to statistical link with field-based AGB by implementing random forest and stochastic gradient boosting (SGB) algorithms. The grey level co-occurrence matrix method [103] and wavelet decomposition were applied using the first principal component of the Sentinel-2 multi-spectral tensor.

In [105], to estimate the AGB from remotely sensed data, parametric and non-parametric methods, including Multiple Regression (MR), k-nearest Neighbour (kNN), Random Forest, the multi-layer perceptron, which performed best among the various methods, were applied to a single Sentinel-2 image using spectral bands and derived indices. Similarly, in [106] the authors explored the capability of spectral and texture features from the Sentinel-2 Multispectral Instrument (MSI) for modeling grassland AGB using random forest (RF) and extreme gradient boosting (XGBoost) algorithms in Shengjin Lake wetland (a Ramsar site), showing that the RF and XGBoost models had a robust and efficient performance and that the introduction of eight grey-level co-occurrence matrix (GLCM) textures moderately improved the accuracy of modeling AGB. The texture is also

the core of [107], in which texture metrics were derived based on different working window sizes (3×3 , 5×5 , 7×7 , and 9×9). The results were compared with those obtained using raw traditional bands (Band 2, 3, 4, 8, 11, and 12), raw traditional and red edge bands (Band 5, 6, 7, and 8A), and red edge bands only; the use of texture with a 7×7 window size and vegetation indices (VIs) yielded higher biomass estimates.

In [108], the authors proposed an innovative and dynamic architecture based on the generative neural network that extracts target-oriented generative features for forest prediction AGB using satellite data. The architecture exploits its generative capacity to produce variables in a latent space to predict AGB, exploiting only the spectral dimension and not the spatial correlations of the images. In the same year, in [109], the authors present an automated machine learning (AutoML) framework for modeling, evaluating, and stacking multiple base models for AGB prediction. This thesis incorporates a hyperparameter optimization procedure for automatically extracting targeted features from multitemporal Sentinel-2 data that minimizes human bias. Also, in this context, the automatic feature extraction only considered the spectral dimension. Finally, a recent work [110] mainly discusses three non-parametric models: the artificial neural network (ANN) and the random forests, and in particular, explores the quantile regression neural network (QRNN), using spectral index and texture features as variables.

These works use machine learning techniques combined with intensive feature extraction, with some focusing only on the spectral dimension and others involving both the spatial and spectral dimensions. Besides the used features, some work [108, 109] also exploited generative networks and AutoML pipelines to minimize the human bias in the feature extraction phase. Moreover, as for other image processing-related domains, researchers [93, 94] are also working on the use of Convolutional Neural Networks (CNNs) designed to produce numerical values for AGB prediction (one for each input image), using input from commercial satellites such as Worldview-2 or visible spectrum images captured by an unmanned aerial vehicle instead of Sentinel-2 open data.

However, some concerns have been raised against the parametrization of these machine-learning models, as tuning can significantly affect the performance of these models [111]. Working at the forest scale, the main

issue related to the exploitation of algorithms requiring the acquisition of calibration points is probably generalization, as sampling is usually implemented in specific areas [112].

Therefore, a concept of interest is the introduction of methodologies allowing for the transfer of a model from one area to another, sharing similar characteristics [113]. This problem is particularly of interest, for example, in precision agriculture, since the variations of the AGB are typically higher due to crop cycles and their sensitivity to environmental conditions. In this context, Reference [113] exploited active sampling to select significant samples to transfer a calibrated model for ryegrass AGB and nitrogen content estimation across different seasons. [114] exploited the Kennard-Stone algorithm to select re-calibration samples acquired in other years to extend the forecasting capabilities of a rice yield model. Reference [115] proposed an active sampling framework to implement multi-year incremental grassland biomass and nitrogen content monitoring.

Dealing with forests, the problem of updating an existing model to make it work when applied to other areas or to the same place at different times has been scarcely addressed, as the dominant approach is the validation of new techniques fed by variously collected data [80, 88, 96, 99, 116] - the most relevant literature on the topic concerns processing airborne laser scanning (ALS) data. Reference [116] exploited repeated ALS acquisitions and ground measures to update models predicting aboveground carbon, basal area, stand density index, and total stem volume within a study area in northern Idaho (USA). [117] demonstrated the feasibility of transferring between different areas of forest attributes derived using ALS data [117]. Reference [118] proposed a methodology to generate new maps of forest structural attributes, integrating ALS samples within pre-fitted models, pointing out that the direct transfer of an available model elsewhere may involve extrapolations or application to conditions not included in the training dataset.

2.2.3 ReUse: REgressive U-Net

Materials and Methods

The following approach introduces ReUse, a U-Net network trained to perform a pixel-wise regression task to map Sentinel-2 images into AGB

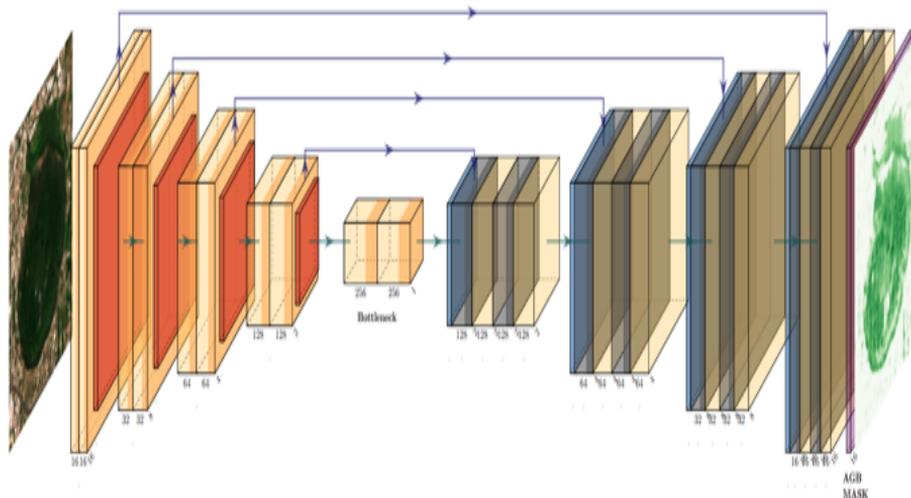


Figure 2.14: ReUse architecture for pixel-wise regression. The input is the Sentinel-2 image with dimensions (patch size, patch size, number of channels); the output is the AGB image with dimensions (patch size, patch size, 1).

rasters. The U-Net was developed by [119] for biomedical semantic segmentation. In the original proposal, the architecture contains two paths. The first is a contraction path (also known as the encoder) designed to capture the context in the image. Several structures are possible but usually involve convolutional and max pooling layer sequences. The second path is a symmetric expanding path (the decoder) designed to produce a pixel-wise prediction using transposed convolutions. These two paths are connected by some skipping connections, designed to improve the localization ability of the network by combining the high-resolution features from the contracting path with the corresponding one in the expanding path. A final convolution layer can then be used to learn and assemble a more precise output based on this information.

The main difference introduced to the original U-Net architecture is that the network has been trained not for performing semantic segmentation but to produce a pixel-wise regression map. This has been obtained by

omitting the softmax in the last layer, thus forcing the optimizer to minimize the loss function based on the actual values predicted by the network, directly comparing them with the AGB ground truth (Fig. 2.14). To the author's knowledge, this is the first time such an approach has been used to estimate AGB. The main advantage is that the proposed variant can extract spatial and spectral features from the satellite multispectral images using an end-to-end paradigm. In particular, a patch-wise approach has been used [120], dividing each Sentinel-2 input and AGB-raster into non-overlapping patches of 16*16 pixels.

One of the problems with the evaluation of remote-sensing applications is the lack of a standard experimental protocol. Indeed, given the high spatial and temporal variability associated with images captured by satellite sensors, selecting different Earth zones and acquisition times can result in different performances, making any comparison with other approaches in the literature less fair. To cope with this, different evaluations were performed on different Earth zones, comparing ReUse against two literature proposals [4, 5] both leveraging a Random Forest [121] classifier trained on spatial and spectral features in an end-to-end paradigm. It is worth noting that in [5], the authors use topographical parameters, such as altitude, as a variable. However, as that data is not always available, only Sentinel-2 data were used as input.

Comparisons between ReUse and the [4] and [5] approaches were made on study areas used by the latter, located in the Central Highlands of Vietnam, and the Yinmar Forest (YM) situated in the northern and central-eastern part of Myanmar, respectively. In [4], a total of 132 spectral and texture variables were extracted from the Sentinel-2 images; the grey level co-occurrence matrix (GLCM) method [103] was used to compute mean, variance, homogeneity, contrast, dissimilarity, entropy, second moment and correlation. In [5], principal component analysis (PCA) [122] was used to eliminate correlated information in satellite images and simultaneously reduce their dimensionality. The first principal component (PC1) was used for texture extraction. When extracting textural features, the grey level co-occurrence matrix method [103] was used to compute mean, variance, homogeneity, contrast, dissimilarity, entropy, second moment and correlation, and wavelet decomposition was also applied considering their usefulness for the representation of relevant features [123]. The wavelet analysis

produces four essential components: the approximation image, horizontal detail, vertical detail, and diagonal detail images. The latter three are regarded as helpful textural measures. In particular, the Coiflet discrete wavelet function was chosen. Thus, based on the first principal component, a three-level decomposition strategy was implemented to generate nine detailed images as independent textural variables for AGB modeling. Finally, two types of textures derived from GLCM-based and wavelet analysis were included in the AGB modeling in combination with 11 spectral indices.

The proposed ReUse architecture has instead been trained using two different setups: one relying only on the raw Sentinel-2 bands, suitably normalized; the other leveraging raw band together with textural and spectral indices as in [5] and [4] respectively. Focusing on the latter setup, concerning the texture variables, the GLCM method was adopted to compute mean, variance, homogeneity, contrast, dissimilarity, entropy, second moment, and correlation to which the nine detail images obtained by applying the Coiflet wavelets have been added by adopting a three-level strategy. The texture variables were computed using the first principal component of the raw Sentinel bands. A 5×5 kernel to construct the GLCM-based features was adopted. Indeed, considering that training was conducted at a spatial resolution of 100 meters (because this is the spatial resolution of ESA's AGB data), such a kernel involves an area of 500 meters by 500 meters.

For all the considered approaches (including ours and competitors), the images have been rescaled to a spatial resolution of 100 meters to match the ESA CCI Biomass Project AGB data. All the experiments have been run using an 8-fold cross-validation strategy, where each fold contains only patches associated with the same zone on Earth. In contrast, different folds refer to other portions of Earth (thus ensuring that there are no data leaks between different folds). For each iteration, 6 of the eight folds were used for training, 1 for validation, and 1 for testing. The validation set was used to optimize the number of epochs for ReUse with an early stopping procedure [124] and the number of trees for the two competitors using two possible values: 250 or 500. In this study, concerning the early stopping procedure, training is stopped when the monitored validation loss has stopped improving after 35 epochs; the maximum number of epochs is

set at 500. The optimizer used is Adam [125] with the default parameters. For the learning rate, following the approach introduced in [126], if no improvement in validation loss is seen for 25 epochs, the learning rate is reduced by a factor of 0.2. The Mean Absolute Error is used as a loss function for ReUse. The trees in the Random Forest are maximally grown, and the number of variables each tree can choose at each split is equal to the square root of the number of features, as suggested by [127]. The source code is available at GitHub repository¹ for reproducibility purposes.

Image Acquisition and Pre-Processing

Starting from the Global Dataset of aboveground biomass of 2018 version 3 of the ESA CCI BIOMASS Project, three study areas were downloaded and used separately to compare ReUse with the competitors [4, 5]. The file "N20E100" contains the AGB of the study area of [4] in Vietnam, while the file "N30E90" contains the AGB of the study area of Yinmar (YM) forest of [5] in Myanmar. It is again emphasized that [4] and [5] use field AGB data, while in this thesis, public AGB data provided by ESA are used. The third study area cut from the "N60E00" file in central Europe was used to compare the approaches in a western area and for the Astroni use case (section 4.1). Table 2.3 shows the names of the downloaded AGB files in GeoTIFF format and the clipped areas of interest in Well-known text (WKT), a text markup language for representing vector geometry objects. From the WKTs, it was possible to download the corresponding Sentinel-2 L2A satellite multispectral images acquired during available cloud-free days. The dates of the downloaded Sentinel-2 images of Vietnam, Myanmar, and Europe are three April, seven March, and twenty-seven July 2018, respectively.

The Sentinel-2 satellite system acquires images with 13 spectral channels at variable spatial resolutions of 10, 20, and 60 meters. This satellite system covers the red-edge region (i.e., b5, 6, 7, and 8A), strategically positioned in the electromagnetic spectrum with unique band settings critical for vegetation modeling [128]. For ReUse, bands 1, 9, and 10 were eliminated due to their coarse spatial resolution, resulting in only 10-band images. The values in the retrieved rasters are digital numbers (DN) that

¹<https://github.com/priamus-lab/ReUse>

Table 2.3: List of the file names downloaded from the ESA Biomass Climate Change Initiative of the Global aboveground forest biomass for 2018, v3 in Geotiff format containing AGB rasters. In each GeoTiff file, a study area was cut out and reported in Well-known text (WKT) format.

AGB file name	Well-Known-Text Area of interest
N60E00	Polygon ((6.41116296045328227 50.733179027500789, 7.47810311940281469 50.733179027500789, 7.47810311940281469 51.57252282701965385, 6.41116296045328227 51.57252282701965385, 6.41116296045328227 50.733179027500789))
N20E100	Polygon ((107.10766723771612874 12.51511413035953346, 107.83207927588165376 12.51511413035953346, 107.83207927588165376 13.26250192346686596, 107.10766723771612874 13.26250192346686596, 107.10766723771612874 12.51511413035953346))
N30E90	Polygon ((96.00143548690360262 22.97297258442770485, 96.49905818628494103 22.97297258442770485, 96.49905818628494103 23.42872411816529876, 96.00143548690360262 23.42872411816529876, 96.00143548690360262 22.97297258442770485))

must be transformed into reflectance by dividing them by the quantification value. The quantification value in the Sentinel-2 product metadata is equal to 10000 [129]. The infrastructure provided by the company Lattitudo 40 was used to download and prepare the Sentinel-2 data described above.

Concerning the AGB data, the dataset [130] comprises estimates of forest aboveground biomass for 2010 [131], 2017, and 2018. They are derived from a combination of Earth observation data, depending on the year, from the Copernicus Sentinel-1 mission, Envisat’s ASAR instrument, and JAXA’s Advanced Land Observing Satellite (ALOS-1 and ALOS-2), along with additional information from Earth observation sources. The Biomass CCI team has produced the data as part of the European Space Agency’s Climate Change Initiative program. The mapping is at 100 m grid spacing with a target relative error of less than 20%

Results

As described in the previous section, the proposed ReUse approach has been compared against two machine learning approaches [4, 5] using Random Forest and leveraging both spatial and spectral dimensions during the feature-engineering phase. The study areas used are Vietnam and Myanmar, which contain the areas of the articles mentioned above, respectively, to compare ReUse with competitors in the areas where these methodologies were initially designed. Moreover, an area in Central Europe was taken into account to test the models on the territorial characteristics of a Western country. Eight-fold cross-validation was used to estimate the error of the models. Table 2.4 presents the experiments for Vietnam, Myanmar, and Central Europe, highlighting that ReUse performs better than its competitors regarding MAE, RMSE, and R^2 .

Figures 2.15 to 2.19 shows a base map of the test area in Central Europe and the AGB predictions of the different approaches. It can be seen that the city areas where there is no greenery are appropriately set to 0, and the tree areas are all highlighted. In particular, solutions (Fig. 2.16 and 2.17) show more marked differences between green and non-green areas than the classical machine learning approaches (Fig. 2.18 and 2.19). Furthermore, the experiments show that ReUse with feature extraction, 42 spectral indexes as in [4] in conjunction with texture variables obtained

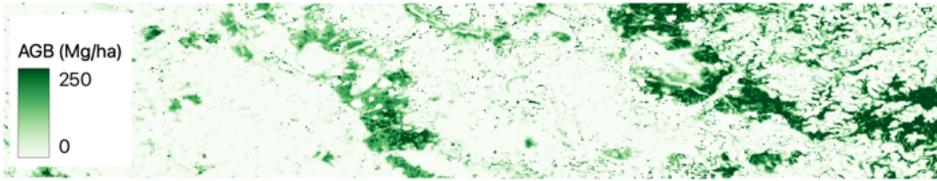


Figure 2.15: Base-map of the test area in Central Europe.

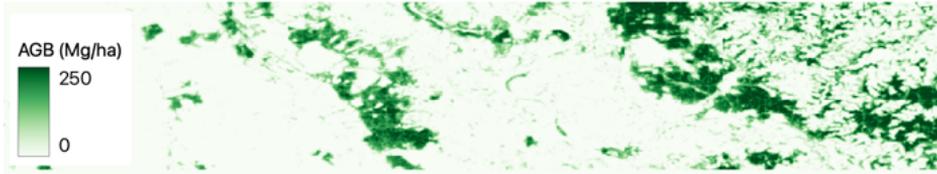


Figure 2.16: ReUse's AGB predictions with raw bands in Central Europe.

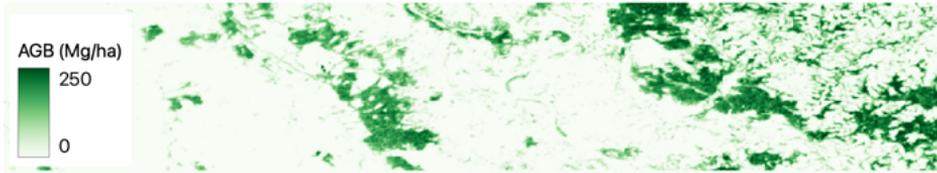


Figure 2.17: ReUse's AGB predictions with raw bands and feature extraction in Central Europe.

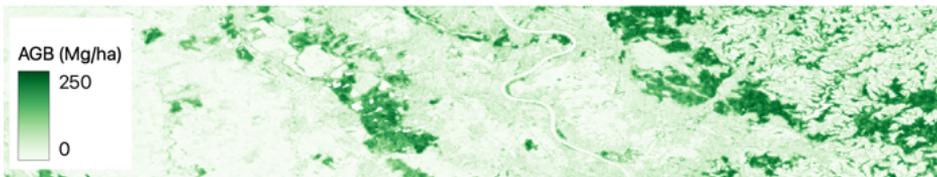


Figure 2.18: Predictions of AGB using the machine learning approach [4] in Central Europe.

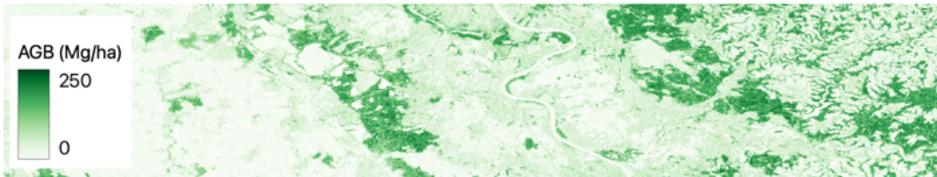


Figure 2.19: Predictions of AGB using the machine learning approach [5] in Central Europe.

Table 2.4: The results of the experiments performed on Vietnam and Myanmar study areas of [4] and [5] and Central Europe study area. At each iteration of the eight-fold cross-validation, six folds are used for training, one for validation, and one for testing. The averages and standard deviations of the Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and R^2 metrics are calculated on the test set. In bold, the best value for each metric per study area.

Area	Model	MAE	RMSE	R^2
Vietnam	ReUse raw bands	42.0 ±6.6	57.7 ±7.3	0.4 ±0.2
	ReUse feature extraction	44.4 ±6.0	59.5 ±4.7	0.4 ±0.2
	Competitor 1 [4]	60.1 ±8.3	73.0 ±9.4	0.2 ±0.2
	Competitor 2 [5]	58.9 ±8.6	72.0 ±9.7	0.2 ±0.2
Myanmar	ReUse raw bands	10.8 ±2.0	15.0 ±2.4	0.7 ±0.1
	ReUse feature extraction	10.7 ±2.2	14.9 ±2.6	0.7 ±0.1
	Competitor 1 [4]	15.7 ±1.9	20.2 ±2.3	0.4 ±0.1
	Competitor 2 [5]	15.5 ±1.5	20.1 ±1.8	0.4 ±0.1
Europe	ReUse raw bands	24.5 ±3.3	46.6 ±5.2	0.6 ±0.1
	ReUse feature extraction	24.1 ±3.4	46.9 ±4.2	0.6 ±0.1
	Competitor 1 [4]	32.5 ±3.1	48.0 ±4.4	0.5 ±0.5
	Competitor 2 [5]	34.8 ±3.1	51.1 ±3.9	0.5 ±0.5

with GLCM method and wavelet analysis as in [5], does not significantly improve compared to ReUse with raw bands. This suggests that such a deep approach can avoid the feature engineering phase in the contest of AGB prediction. Notably, errors reported for the competitors are higher than those reported in the corresponding papers. This is because the ground truth in this thesis comes from ESA, while the original works used field measurements, which were not released for reproducibility purposes.

A case study: The WWF Astroni Nature Reserve

A case study for the Astroni nature reserve in southern Italy showed how ReUse, Sentinel-2, and ESA’s AGB public data can help estimate

CO₂ in forest areas and monitor deforestation downstream of events such as fires. ReUse is adopted with raw Sentinel-2 bands alone without using other extracted features that would not bring decisive benefits, as demonstrated in the previous section.

The area of central Europe contained in the file 'N60E00' was chosen to monitor the Astroni reserve to train ReUse because it is certainly an area with characteristics closer to those of southern Italy than the other two datasets containing the areas of Vietnam and Myanmar. In [132], research was conducted within the Astroni Crater World Wildlife Fund (WWF) Reserve in the volcanic area of the Campi Flegrei in the urban area of Naples, Italy. The Reserve (247 hectares) lies within the caldera of an extinct volcano with a maximum altitude of 255 m above sea level and an elliptical shape (2×1.6 km). The inner part of the crater has a deep depression containing a lake, where a minimum altitude of 9 m above sea level is reached. In the central part of the crater, near the largest lake, two other small lakes and three hills rise to the bottom at 45, 74, and 82 m above sea level, respectively. Throughout the crater, [132] focused on the area of holm oak forest (127 ha) and the location of mixed forest (104 ha) to define the two main ecosystems and found that the total carbon stocks of the phytomass of these two ecosystems were 22173 ± 7054 tonnes using sampling from April to October 2016.

To make inferences, the Sentinel-2 images downloaded on 31 May 2017 were upscaled to a spatial resolution of 10 meters, assuming that the spatial correlations learned from the network at 100 meters are also reproducible at 10 meters; this improves the resolution of the predictions compared to the resolution of the AGB raster of the ESA CCI Biomass project which is at 100 meters. The predictions of Carbon stocks were made by creating the predicted raster of AGB ten by ten non-overlapping patches, corresponding to one hectare each. Then, for each patch, the average AGB value expressed in tonnes per hectare was taken, which, when multiplied by 1 hectare, which is the extent of the patch, yields a value in tonnes. These values were summed over all patches, and the final result was multiplied by 0.5 [133][134] to obtain the value of absorbed carbon in tonnes. From this procedure, the Carbon stock estimate for Astroni on 31 May 2017 was 18748 tons, in line with the forecast of [132], which is 22173 ± 7054 tons for phytomass in 2016. Please note that the latter includes the roots of

the plants, which is not contained in the estimation described here but can be considered around 30%. To confirm the fact that the trained network recognizes a decrease in AGB downstream of a fire, the AGB raster of Astroni on 24 August 2017 downstream of the fire on July 2017 is shown in Fig. 2.20 compared with the AGB raster of the same area before the fire. The estimated aboveground biomass carbon stock on 24 August 2017 for the nature reserve downstream of the fire on July 2017 is 10104 tonnes, confirming the above.

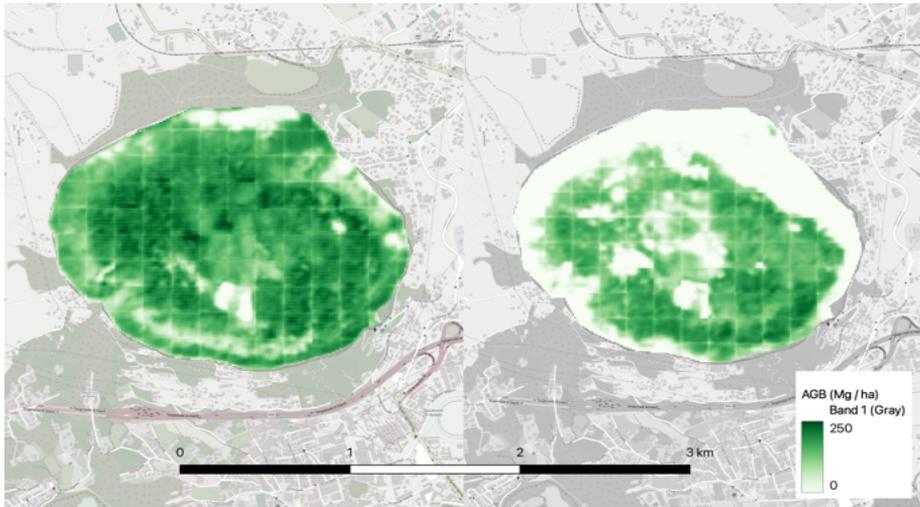


Figure 2.20: On the left is the predicted aboveground biomass raster of the Astroni nature reserve before the July 2017 fire; on the right is the predicted aboveground biomass raster after a major fire event for the same area.

It is stressed that the proposed approach is based entirely on open data. The fact that a prediction based on public data agrees with ranges given by a ground truth provides worth to a tool that can be obtained without field measurements, which could help monitor carbon stocks in forest areas. The prediction after the fire cannot be verified in any way; however, the fact that the prediction before the fire is in line with the ground truth and that the forecast after this event shows a decrease due

Table 2.5: AGB statistics for the considered area.

Min AGB Clip (t/ha)	Max AGB Clip (t/ha)	Mean AGB (t/ha)	Std AGB (t/ha)	Total Area (km ²)
22.1	81.3	69.5	43.1	85.1

to the fire is encouraging.

2.2.4 Machine Learning Ensembles: Active Learning Strategies for Model Transfer and Field Sampling Reduction

Materials and Methods

ALS data are systematically collected in Finland for digital terrain mapping and forest parameters estimation [135], including AGB. These data, opportunely integrated with satellite images and historical field sampling, have been exploited by the Finnish Forest Centre to produce the reference AGB data used in this thesis [136], which analyzes an area of about 85 km², partitioned in 13 clips of 6.55 km² extension each. The clips, whose global statistics are reported in Table 2.5, are quite inhomogeneous regarding the average AGB amount and its distribution. The average AGB ranges between 37.1t/ha and 81.3t/ha. Its standard deviation ranges between 22.2t/ha and 63.2t/ha. The overall AGB mean and standard deviation are in the order of 69.5t/ha and 43.1t/ha. The area is characterized by almost flat topography, with elevation values ranging approximately 100-200 m above sea level, with a lack of significant urban areas and agricultural lands. It belongs predominantly to the middle boreal vegetation zone, with a dominant presence of coniferous [137].

Satellite data, used to predict the AGB, belong to a publicly available dataset ² composed of several clips. It is presented as a benchmarking tool for AGB studies. Therefore, data are provided without any geographic and/or temporal attributes. For each clip, the dataset includes, beyond AGB data, one year of Sentinel-1 SAR and Sentinel-2 MS observations,

²<https://huggingface.co/datasets/nascetti-a/BioMassters> (accessed on 18 July 2023). Further information about data can be found at <https://neurips.cc/virtual/2023/poster/73499> (accessed on 18 July 2023).

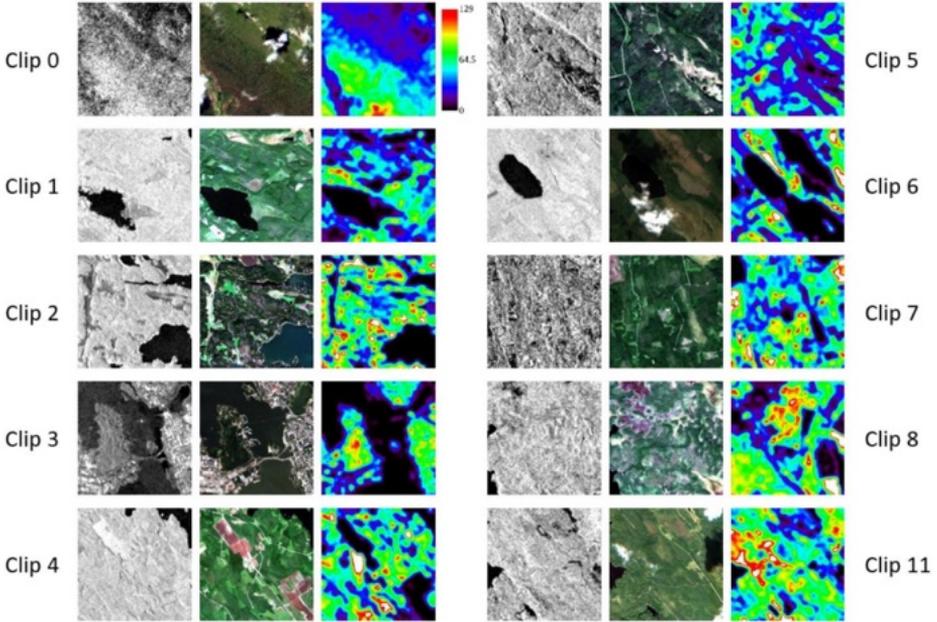


Figure 2.21: Composition of the exploited dataset for some sample clips. From left to right, S1 SAR image, natural color S2 image and AGB map. The size of each clip is 256×256 pixels. The pixel spacing of the maps is 10 m.

made approximately every month (see Figure 2.21). In particular, SAR data have been acquired regularly once per month due to their insensitivity to atmospheric conditions. MS acquisitions are mostly concentrated during spring and summer to cope with massive cloud coverage typically registered in the other periods of the year.

A methodology is introduced to transfer a calibrated partial least squares regression (PLSR) [138] model to areas with similar characteristics. In other words, it is assumed that a model calibrated using an extended ground truth is available. The question to be addressed is about how to select relevant samples for fitting this model with new scenes.

To this end, the workflow depicted in Figure 2.22 is adopted. The inputs are a calibrated regression model and two-time series of data, one MS and one SAR, which are exploited to calculate a set of predictors, a total of

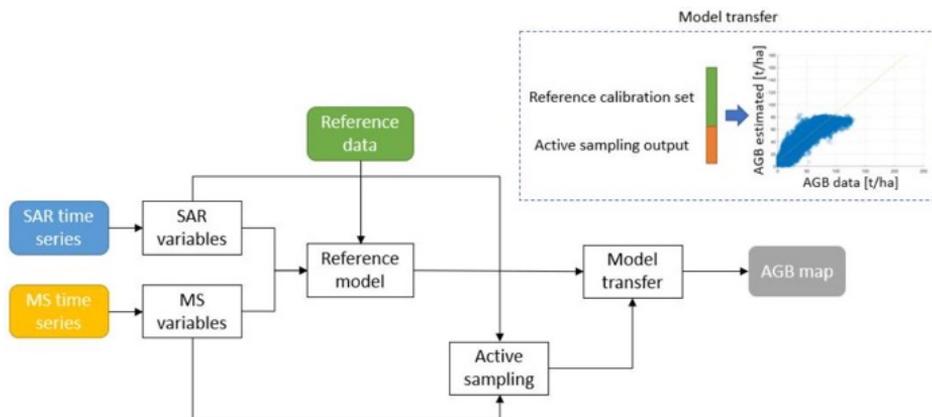


Figure 2.22: Proposed workflow for AGB model transfer. Colored blocks identify input/output data.

159, to be exploited for regression. MS predictors are mostly constituted by vegetation indices. SAR ones are constituted by both intensity-based and texture variables derived from the Haralick co-occurrence matrix [139].

The predictors are calculated for each available acquisition and then averaged on a yearly basis. Following this operation, the processing forks in two branches. The first one is dedicated to the retrieval of the model to be transferred, which is obtained via bootstrap, i.e., a regression calibrated with a random sample selection. Bootstrapping is implemented as follows. It is assumed that the AGB is known in a specific area (Clip 0 in Table S1). Then, among 1000 PLSR experiments calibrated with random sample selection, the one providing the lowest RMSE is selected as reference for transferring.

Model transferring requires that the same predictors are used for each regression. However, the literature highlighted that the large amount of data usually ingested into PLSR can contain irrelevant information that can cause the reduction in the model performance [140]. Therefore, a set of potentially informative predictors is selected through the calculation of the variable importance in prediction (VIP) score [141].

As the VIP score is an output of PLSR, a tentative regression is implemented with this purpose. In other words, this will not be used for AGB estimation. In this phase, the number of latent variables, i.e., the set of

orthogonal factors having the best predictive power, is set to 10 .

The VIP score is defined, for each variable j , as the sum, over latent variables f , of its PLS-weight value w_j weighted by the percentage of explained variance of the specific latent variable SSY_f . For the j -th variable, it holds [141]:

$$VIP_j = \sqrt{\frac{J \cdot \sum_{f=1}^F w_{jf}^2 \cdot SSY_f}{F \cdot \sum_{f=1}^F SSY_f}} \quad (2.1)$$

where F is the total number of latent variables and J the total number of variables. The relation for the calculation of SSY_f is given by [141]:

$$SSY_f = \mathbf{b}_f^2 \mathbf{t}_f^T \mathbf{t}_f, \quad (2.2)$$

where \mathbf{b}_f and \mathbf{t}_f are the PLS inner relation vector of coefficients and the score vector relevant with the f -th latent variable, respectively. It is common practice in the literature to assume that informative variables are associated with a VIP score larger than one. This means that those variables have an above-average influence on the building of the model explaining the observations [142].

According to Figure 2.22, following the definition of the model to be transferred, active sampling is implemented to select re-calibration data. To this end, the stack composed of all the predictors for the new area is processed with a principal component analysis (PCA) for dimensionality reduction. Then, an ISODATA clustering is implemented to obtain labelled regions. In this case, the maximum number of allowed classes is set to 10 .

The classified image provided by ISODATA clustering is treated with connected component labelling [143] to retrieve elementary regions, i.e., the image segments. They are used to calculate the spatial average of the regression variables, which constitute the input for active learning.

As described in [144], active learning techniques can be categorized based on uncertainty [145] or diversity criteria [146]. In the first case, samples are ranked according to their uncertainty. The higher the uncertainty, the better their rank. Among these techniques, those based on variance-based pool of regressors (PAL) are probably the most interesting [147]. Active learning techniques based on diversity criteria, instead, select

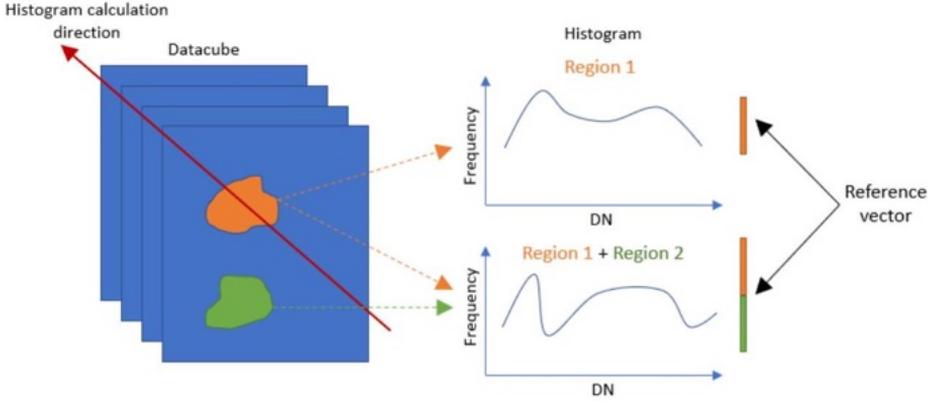


Figure 2.23: Schematic representation of histogram calculations for active sampling.

samples based on the dissimilarities they introduce in the training datasets [144]. Different metrics can be used to assess such dissimilarity like the Euclidean distance [148], the cosine angle distance [149] or the entropy [150], which is the one used in this thesis according to the definition provided by Shannon [151]:

$$H = - \sum_{i=1}^N P_n \log_2 P_n \quad (2.3)$$

where P_n is the normalized probability of the n -th histogram quantization level and N is the total number of bins. The entropy gives information about the shape of the histogram and measures the quantity of information carried by a signal. The higher the entropy, i.e., the flatter the histogram, the higher its informative content.

According to Reference [115], the histogram is calculated on the vector constituted by the region-wise average of the predictors. In other words, given the N predictors for the generic orange region provided by the ISO-DATA clustering depicted in Figure 2.23, the reference vector for histogram calculation is composed of the N average values of the predictors. The entropy H_1 is calculated according to Equation 2.3 and stored as metrics for active sampling selection.

The second step of the active sampling process, the vector representative of the green region of Figure 2.23, is appended to the one previously considered. The entropy H_2 of this new vector, constituted by $2N$ elements is calculated. If the condition $H_2 > H_1$ is verified, it is retained. Otherwise, the vector corresponding to the green region is discarded. In any case, the entropy value used for the acceptance test is the last one computed. The process continues until all the regions have been tested. Those determining an increase in the entropy against the last available one are appended to the selection, and the corresponding regions are marked as informative.

1. Consider each retrieved cluster and make the spatial average for each spectral band;
2. Starting from a randomly selected cluster, calculate its histogram considering all the elements of the average spectral response and its associated entropy H_1 according to Equation 2.3;
3. Add another cluster to the dataset by appending its (average) spectral response to the vector constituted by the one previously considered. Calculate the new histogram and its entropy H_2 ;
4. If $H_2 > H_1$ mark the cluster as informative and continue the process by adding new clusters to the dataset. Do not delete those marked as not informative. Clusters marked as informative are those to be sampled to retrieve model calibration data.

The average number of samples, i.e., regions, selected by the adopted active sampling methodology is around 301 within areas of about 6.5 km^2 . In an operational environment, this number can be considered high. Therefore, a simple strategy for reducing the number of samples to be considered for model transfer has been implemented.

The output of the above-discussed procedure is the list of the regions marked as informative. Each of them is associated with an entropy value, calculated by adding the vector representative of the i -th region to the one collecting those already accepted. Moreover, a class ID coming from the ISODATA clustering is available.

Informative regions are then separated, based on the ISODATA class ID, and sorted in descending entropy order. The first k samples of the chart are retained for model calibration. In such a way, the maximum number of calibration samples is fixed to $n_{\max} = kt$, where t is the maximum number of allowed classes in ISODATA clustering. In the following experiments, t has been set to 10. Two different settings for k have been tested, i.e., $k = 10$, yielding $n_{\max} = 100$, and $k = 5$, yielding $n_{\max} = 50$. In the following, these samples will be referred to as very informative samples (VIS). The adoption of this strategy allowed for the reduction in the areas to be sampled to 73 per clip, in the case of $k = 10$, and to 36 per clip, in the case of $k = 5$.

Regressions for test areas are implemented using different strategies. The first one is a PLSR. It is a statistical method aiming at finding a linear regression model between the observations and the independent variables by projecting the predicted variables and the observed ones to a new space [138]. This technique is particularly suited when the matrix of predictors has more variables than observations, and there is multicollinearity among predictors, as in the case study at hand.

The second regression technique used for AGB predictions is gradient boosting (GB). It consists of a finite set of weak learners (typically decision trees) that minimize an arbitrary differentiable loss function. A meta-learner assigns weights to each learner and combines their predictive results through voting methods to provide better predictive performance for regression problems. In particular, the least-squares boosting algorithm has been used [152]. This means that the loss function to be minimized is the mean squared error of the learners. The ensemble of GB and PLSR represents the third regression strategy. In this case, the final result is given by the average of the two predictions.

Results

The aggregated results of the proposed AGB estimation methodology for the 12 considered clips are reported in Table 2.6, together with benchmarking results.

Using all the samples selected by the adopted active sampling technique (see Table S2), the obtained RMSE, using PLSR, ranges between 20.8t/ha and 47.2t/ha with an average of 28.8t/ha. GB provided a min-

Table 2.6: Aggregated AGB estimation results for the different tested sampling settings and benchmarking techniques. The column "Inc" refers to the incremental sampling setting.

Samples	Proposed						Benchmark		NN
	PLSR		GB		Ensemble		Bootstrap		
	Area	Inc	Area	Inc	Area	Inc	RMSE *	RMSE	
All	28.8	31.3	33.7	36.3	28.8	30.6	26.2	46.8	
$k = 10$	30.7	32.1	34.2	36.7	30.0	31.5	27.7	47.7	30.4
$k = 5$	31.6	32.2	34.0	37.9	30.7	32.3	28.5	49.7	

imum RMSE of 24.6t/ha and a maximum one of 48.5t/ha with an average of 33.7t/ha. The ensemble resulted in RMSE values ranging between 21.4t/ha and 44.7t/ha with an average of 28.8t/ha.

Reducing the samples to VIS does not significantly affect the estimate (see also Table S2). In the case of $k = 10$, the RMSE ranges between 21.5t/ha and 51.6t/ha with an average of 30.7t/ha, using PLSR. GB estimates provided RMSE values ranging between 25.1t/ha and 52.1t/ha, with an average of 34.2t/ha. The estimates obtained using the ensemble range between 21.4t/ha and 46.2t/ha, with an average of 30.0t/ha.

When the setting is $k = 5$ (see Table S2), the obtained RMSE ranges between 21.7t/ha and 50.3t/ha, with an average of 31.6t/ha, using PLSR. GB exploitation provided RMSE values ranging between 24.8t/ha and 51.4t/ha with an average of 34.0t/ha. Using the ensemble, the RMSE ranges between 21.4t/ha and 49.3t/ha with an average 30.7t/ha.

The column named "Inc" reports the results obtained using an incremental model. In this configuration, all the points used to calibrate clips previously analyzed are included within the calibration set for the current regression. In other words, as an example, the calibration set for the third regression includes the data used to train the models used for the first and the second regression. Indeed, the exploitation of more calibration points, referring to different areas, seems to not be beneficial for the estimate, as the average RMSE is higher than the one retrieved using the standard setting in most of the experiments (see Table S2).

Some scatterplots relevant to selected experiments (see Table S2 for clip indexing) are depicted in Figure 2.24. In particular, the figure reports

Table 2.7: VIP variables extracted from the regression of the reference area.

Name	Formula	Ref
Band 2	R_{492}	
Band 5	R_{704}	
Band 8	R_{832}	
Aerosol free vegetation index	$AFRI_{1600} = R_{830} - 0.66 \frac{R_{1613}}{R_{832} + 0.66R_{1613}}$	[153]
Aerosol free vegetation index	$AFRI_{2100} = R_{832} - 0.5 \frac{R_{832} + R_{2202} + 0.6R_{1613}}{R_{832} + 0.5R_{270}}$	[153]
Ashburn vegetation index	$AVI = 2R_{832} - R_{560}$	[154]
Chlorophyll absorption ratio index	$CARI = \frac{R_{704}}{R_{664}} \sqrt{\frac{(670a + R_{664} + b)^2}{(a^2 + 1)^{0.5}}}$	[155]
Difference 800/550	$D_{800/550} = R_{832} - R_{560}$	[156]
Green difference vegetation index	$GDVI = R_{864} - R_{560}$	[157]
Differenced vegetation index MSS	$DVIMSS = 2.4R_{832} - R_{664}$	[154]
Global environment monitoring index	$GEMI = n(1 - 0.25n) - \frac{R_{664} - 0.125}{1 - R_{644}}$	[154]
Misra soil brightness index	$MSBI = 0.406R_{492} + 0.600R_{560} + 0.645R_{704} + 0.243R_{832}$	[154]
Misra yellow vegetation index	$MYVI = 0.723R_{492} + 0.597R_{560} + 0.206R_{704} + 0.278R_{832}$	[154]
Modified chlorophyll absorption in reflectance index	$MCARI = [(R_{704} - R_{832}) - 0.2(R_{704} - R_{560})] \frac{R_{704}}{R_{664}}$	[158]
Nonlinear vegetation index	$NLI = R_{782}^2 - R_{664}$	[159]
Reflectance at the inflexion point	$RRE = \frac{R_{664} + R_{782}}{2}$	[160]
Simple ratio 833/1649	$SR_{800/1649} = \frac{R_{832}}{R_{1613}}$	[161]

the results of AGB estimation for Clip 0 (Figure 2.24a) and Clip 4, relative to ensemble regression with the setting $k = 5$ (Figure 2.24b) and the neural network output (Figure 2.24c). The regressions shown in the second row of the graphics concern Clip 7. The one reported in Figure 2.24d has been obtained using ensemble regression calibrated with all the samples selected by active learning. The plot shown in Figure 2.24e concerns ensemble regression with $k = 5$. All the other plots are omitted for brevity, but aggregated results in the plane (R^2 , RMSE) are shown in Figure 2.24f, relative to ensemble regression.

The variables used for regression, determined through VIP score thresholding, are all derived from the MS dataset. In particular, the 17 indices reported in Table 2.7 were informative for model training. Remarkably, none of them are derived from the SAR dataset.

The adopted benchmarking techniques are randomly calibrated PLSR

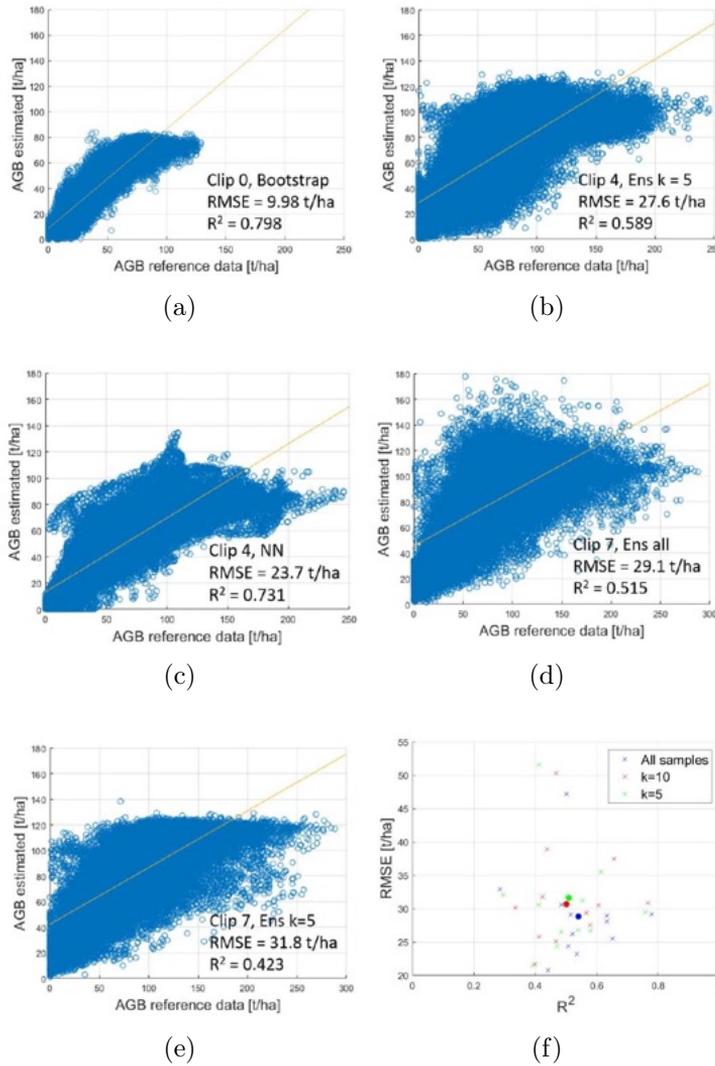


Figure 2.24: Some scatterplots relevant to the performed experiments. (a) Clip 0, bootstrap ($R^2 = 0.798$, $RMSE = 9.98t/ha$), (b) Clip 4, ensemble regression $k = 5$ ($R^2 = 0.589$, $RMSE = 27.6t/ha$), (c) Clip 4 neural network ($R^2 = 0.731$, $RMSE = 27.6t/ha$). (d) Clip 7, ensemble regression, all samples ($R^2 = 0.515$, $RMSE = 29.1t/ha$), (e) Clip 7 ensemble regression, $k = 5$ ($R^2 = 0.423$, $RMSE = 31.8t/ha$), (f) ensemble regression results in the plane (R^2 , $RMSE$) for the different sampling settings. Overall means reported with solid markers.

and NN-based estimation. In the first case, a number of samples equal to the one determined via active sampling has been adopted for a fair comparison. Bootstrapping has been implemented to retrieve the minimum RMSE and the average one, for a total of 1000 experiments.

Bootstrapping returned a minimum RMSE (indicated as RMSE* in Table 2) ranging between 21.0t/ha and 39.4t/ha (see Table S2). The average of the minimum RMSE is 26.2t/ha. The column reporting the mean RMSE values ($\overline{\text{RMSE}}$ in Table 2), obtained via bootstrap, is useful to aid in pointing out that this method can fail, as not all the experiments can be considered successful whatever the combination of samples is used for model calibration (see Table S2).

NN-based estimation has been implemented using ReUse [112]. This is a UNet network trained to perform a pixel-wise regression task, mapping Sentinel-1 and Sentinel-2 images, concatenated along the spectral axis into a single tensor, into AGB raster. The architecture consists of two pathways. The first is a contracting pathway (the encoder) that captures the image context. The second pathway is a symmetric expanding pathway (the decoder) designed to produce pixel-wise predictions using transposed convolutions. Several skip connections connect the two pathways.

It is worthwhile to remark that the deep learning benchmark was carried out without the usage of training samples belonging to the clips under investigation for AGB estimation. The training phase has been implemented using the same number of clips exploited for validation. As for the bootstrap benchmark, which is typically based on sparse sampling, the training on points belonging to the clips under testing allows the model to learn the characteristics of the specific clip effectively. This approach is unsuitable for deep learning, which is typically trained using patches, here retrieved from clips not considered for AGB estimates.

The network generates a pixel-wise regression map and can extract spatial and spectral features from satellite images using an end-to-end paradigm. In particular, a patch-wise approach is adopted [158]. Each input image and ground truth AGB map is divided into nonoverlapping 32×32 pixel patches. The training has been implemented using a ground truth with comparable dimensions of the area to be estimated. It stops when the monitored validation loss stops to improve after 35 epochs. The maximum number of epochs is set to 200. The optimizer used is Adam [159]

with default parameters. The learning rate is reduced by a factor of 0.2 if no improvement in validation loss is observed for 25 epochs. The adopted loss function is the mean absolute error. The RMSE values obtained using this technique range between 20.6t/ha and 43.9t/ha (see Table S2), with an average of 30.4 t/ha.

2.2.5 Conclusion

Sustainable Development Goal 15 aims to protect, restore and promote the sustainable use of terrestrial ecosystems, sustainably manage forests, combat desertification, halt and reverse land degradation and halt biodiversity loss' [83]. Furthermore, REDD+ projects aim to reduce Greenhouse Gases (GHG) concentrations in the atmosphere and contribute to climate change mitigation through various activities, including carbon stock enhancement [162]. Developing systems that can estimate the carbon absorbed by forests globally and monitor losses associated with deforestation phenomena such as fires on time is essential. Forests cover about 40% of the Earth's ice-free land surface and constitute the vital core of the whole planet since they represent the sentinels of the carbon cycle. Monitoring their biomass is key, as this parameter is the principal indicator of the carbon sequestration capacity of a vegetation ecosystem.

Section 2.2.3 introduced the ReUse architecture. ReUse is based on a pixel-wise regressive U-Net, able to generate a pixel mask of AGB predictions with computational advantages, particularly when monitoring large areas. This is a great advantage over classical machine learning algorithms that require feature extraction work to derive indices that capture both spectral and spatial information content. Result demonstrated that this approach can dispense with the feature engineering phase and work directly with the Sentinel-2 raw bands appropriately normalized. The combined use of Sentinel-2 data and ESA's AGB data with a U-Net approach could be suitable for estimating the carbon absorbed in urban and rural areas and help monitor deforestation events without field measurements.

Instead, Section 2.2.4 described a solution to the problem of above-ground estimation using active learning principles, allowing for the transfer of a calibrated regression model through a data-driven selection of the most informative re-calibration samples. In this regard, a technique exploiting Shannon's entropy as a diversity criterion has been discussed, with a par-

ticular focus on how it allows for the dramatic reduction in the number of samples to be collected, to a few samples per square kilometer, which is a quantity suitable for operational environments. The proposed framework requires one reference area, for which an extended ground truth should be available, to calibrate the model to be transferred to the rest of the forest, assuming that its structure does not exhibit significant variations. Reference AGB data can be obtained, as an example, with a LiDAR acquisition. The obtained results testified the reliability of the workflow, which restituted an average RMSE of 28.8t/ha, in case all the samples selected by the adopted active sampling methodology, on average 46/km², are exploited. The adoption of the ISODATA-based selection allowed for the reduction in the samples to be collected to about 10/km² (when the setting was $k = 10$) and 6/km² (when the setting was $k = 5$) with negligible impact on the AGB estimation, whose RMSE increased to 30.0t/ha when the setting was $k = 10$ and to 30.7t/ha when the setting was $k = 5$. The registered values of the coefficient of determination, R^2 , as expected, follow the trend of the RMSE, i.e., they tend to increase as the RMSE decreases, as shown in Figure 2.24f. As a general comment, it is remarkable that the incremental framework, in which all the samples collected for previous estimates are involved in the current regression, did not add value to the estimates. The aforementioned RMSE values refer to the most performing regression configuration which, interestingly, is provided by the ensemble of PLSR and GB. This configuration outperforms the single regression technique in 18 experiments out of 32. This trend increases as the number of calibration points is reduced. The analysis of the errors reveals that they are strongly correlated with the AGB standard deviation of the clip. The higher the standard deviation, the higher the RMSE. This result is in line with the literature, as reported, as an example, in [163]. As shown in the previous section, the variables selected for regression through VIP scoring belong to the MS dataset. Despite the well-known relation between SAR backscattering and the AGB [89], spectral information in the near-infrared and red edge frequencies, which appear in almost all the indices listed in Table 2.7, resulted dominant. This is due to the scarce sensitivity of C-band frequencies to AGB within the resolution cell [164]. However, the contribution of SAR features is expected to be more significant using L-band data, as this wavelength typically shows higher correlation with the AGB [89]. This

result can be found, as an example, in [163], where the authors highlighted how AGB estimates can benefit from the presence of L-band SAR features among regression variables. The RMSE values obtained adopting the proposed methodology are comparable with those declared in the literature for the specific application. As an example, Reference [88] claimed an RMSE of 18.9t /ha on an area having an average AGB of 49.4t/ha, with a standard deviation of 33.8t/ ha using extreme gradient boost algorithm. Reference [99] reported an RMSE of about 30.0t /ha in the study of an area with an average AGB of 55.5t/ ha and a standard deviation of 18.2t/ha, using extreme gradient boosting as well. Vafaei et al. [163] reported an RMSE of 38.7t/ha, using support vector regression in an area with an average AGB of 206.5t/ ha and a standard deviation of 78.0t/ha.

All these studies were characterized by the same approach, consisting of sparse sampling of a few plots (149 in Reference [163] and 367 in Reference [88]) distributed over large areas. These data are used in part to train the model and in part to validate it. In other words, the authors used a few samples to calibrate a single model to fit a large area. The validation was performed on a small subset of the ground truth, as the training phase tends to absorb the majority of the collected samples.

The approach proposed here is different. First, the AGB estimation problem is tackled via active learning, which allows for the selection of highly informative calibration samples, thus introducing a further guide for field operations which can be accordingly significantly reduced. It requires extended reference data about a small area, which can be retrieved with a single LiDAR flight, useful to calibrate a machine learning model via bootstrapping, which, as demonstrated in the previous section, is the most effective solution in the presence of extended ground data. This model can be effectively transferred elsewhere through the collection of a few new calibration samples, according to active learning principles. This means that the objective was to demonstrate that a reliable selection of re-calibration samples feeding simple regression models can provide results comparable with methodologies exploiting more sophisticated solutions tuned on the specific case study and/or requiring extended ground data.

In this regard, the proposed active learning technique dramatically reduces the necessity of field sampling, which is the principal limitation in large-scale AGB mapping [84], up to about 6 samples /km². The concept

of active sampling for model transfer also opens the possibility of the exploitation of historical field data in combination with newly acquired ones with the purpose of updating and/or expanding past AGB maps [137]. Differently from the literature, using, as aforementioned, a few plots to retrieve estimation quality metrics, the validation has been performed on an extended ground truth for a total of more than 850,000 validation points.

The reliability of the proposed methodology has been tested through comparison against some of the most popular techniques for AGB estimation, like bootstrap and NNbased estimation. The benchmark revealed that random sampling is not a viable option if an extended ground truth is not available.

NNs are popular tools for AGB estimation allowing for reliable estimates, provided that an extended ground truth is available for model training [165][166]. Dealing with AGB estimates, this represents a significant limitation, since reference data are usually retrieved via sparse plot sampling. The obtained RMSE values, using ReUse, testify that NNs are suitable for model transfer, i.e., it is possible to train the net using reference data acquired on a test area and then use the weights for making inference elsewhere. However, the amount of reference data should be comparable with the extension of the area to be predicted.

The performed experiments revealed that the average RMSE values obtained using ReUse are comparable with those provided by the proposed framework and by the best bootstrap performance. This means that the choice of the best solution for the specific AGB estimation problem to be solved can be a function of the available data and computational power. In other words, the usage of machine learning ensembles can be a viable solution to cope with the sparsity of reference data. On the other hand, when extended ground truths can be collected, NNs and bootstrapping, or even a combination of the two, could represent an option, as they can slightly reduce the estimation error.

Chapter 3

Monitoring the Present

The future depends on what you do today.

Mahatma Gandhi

3.1 Introduction

Continuous monitoring of the Sustainable Development Goals is crucial for tracking progress and ensuring accountability. Chapter 2 introduced the concept of "learning from the past," demonstrating its feasibility with various use cases. These examples highlight how combining Remote Sensing data with Machine Learning techniques facilitates extracting critical information, showcasing the robust potential of historical data analysis in environmental and geographical studies. This approach involves training algorithms on historical satellite data and existing reference data. The retrospective analysis facilitated by this method is pivotal in understanding long-term trends and variations in Earth's ecosystems. By leveraging past data, end users can decipher the dynamics of environmental changes, know the status quo, evaluate historical land use alterations, and gain insights into climatic shifts over extended periods. This historical perspective is invaluable for contextualizing current observations and informing future predictions, enriching our understanding of Earth's complex systems and their evolution. Understanding the status quo provides a snapshot of cur-

rent conditions but is insufficient. Continuous monitoring is crucial for several reasons:

- **Dynamic Environmental Changes:** The environment is constantly evolving. Continuous monitoring captures these changes, providing insights into patterns, trends, and anomalies over time.
- **Predictive Analysis:** Ongoing data collection allows for more accurate forecasting. Understanding past and present dynamics enables better prediction of future events or changes.
- **Immediate Response:** Continuous monitoring allows for the timely detection of critical changes, such as natural disasters, enabling quicker response and mitigation efforts.
- **Long-term Studies:** Continuous monitoring contributes to long-term environmental studies, essential for understanding broader ecological and climatic trends.
- **Policy and Decision-Making:** Continuous data flow aids in informed decision-making and policy development, ensuring strategies are based on current information. While understanding the current state is necessary, continuous monitoring provides a comprehensive, evolving picture for informed decision-making, timely response, and long-term environmental stewardship.

The methodologies elucidated in Chapter 2 exemplify how the effective integration of Earth observation (EO) with AI enables the generation of up-to-date geographical data spanning extensive regions, thereby supporting urban planning and evidence-based policymaking. Algorithms trained on past data can also be effectively utilized for continuous analysis without modification. This method is cost-effective and efficient, maintaining consistency in data interpretation. However, it's not without challenges. The primary concern is the algorithm's ability to adapt to new environmental variables or changes, as models trained on historical data may not accurately reflect evolving conditions. Additionally, advancements in data collection and sensor technology could potentially reduce the effectiveness of these algorithms. Therefore, while this approach benefits its efficiency

and consistency, assessing its ongoing relevance and adaptability to current and future environmental scenarios is crucial.

Although not exhaustive, the following is a list of requirements to ensure effective and reliable continuous monitoring of environmental and geographical changes with Remote Sensing data:

- **Near Real-Time Data Acquisition:** This involves the continuous capture of satellite imagery, allowing for the immediate observation of current conditions.
- **Near Real-Time Analysis and Response:** Unlike historical data, monitoring the present is geared towards immediate analysis and action. This is particularly important for disaster response, environmental monitoring, and urban planning. Cities are dynamic, with rapid changes in infrastructure, population, and land use. Real-time data allows planners to track these changes as they happen, ensuring that planning decisions are based on the most current information.
- **High Temporal Resolution:** Satellites used for monitoring the present often have a high revisit time, meaning they frequently capture images of the same area. This allows for consistent tracking of ongoing changes and phenomena.
- **Integration with Other Real-Time Data Sources:** Modern monitoring often involves integrating satellite data with other real-time data sources such as sensors, weather data, and social media feeds, providing a comprehensive view of current conditions.

Table 3.1 compares Machine Learning requirements between the pillars Learning from Past and Monitoring the Present, highlighting the differences in the type of data used, the objectives of the analysis, and the methodologies applied. In particular, requirements differences in data volume and Processing need some attention. "Learning from the Past" involves analyzing large volumes of historical data, which can be vast and varied, thus requiring substantial storage and processing capabilities. On the other hand, "Monitoring the Present" focuses on the real-time or near-real-time Processing of data, necessitating rapid data handling and analysis capabilities to address immediate needs. This distinction underscores

the importance of leveraging cloud computing for continuous monitoring. Cloud computing offers scalable and flexible resources capable of handling the high-volume, high-velocity data characteristic of real-time monitoring. This provides an efficient solution for Processing and storing large datasets, ensuring timely insights and responses in dynamic environmental contexts. However, the use of cloud computing brings with it several challenges to take into account. In the "Learning from the Past" phase, since this analysis is often one-shot, computational resources can be deactivated post-analysis, offering some efficiency in resource utilization. In contrast, the "Monitoring the Present" phase requires continuous operation of computational resources to handle real-time data streams. Despite this constant operation, significant computation bursts are infrequent, leading to potential underutilization of resources during idle periods. This scenario highlights the need for mechanisms to optimize resource use and manage costs effectively, especially during low-demand intervals. With their ability to scale resources dynamically, cloud computing solutions present a viable approach to addressing these challenges, allowing for cost-effective management of varying computational loads.

In this chapter, the focus is on the technological aspects of rendering analyses accessible for continuous monitoring. Starting with historical data, ongoing updates can highlight significant changes in the data. This approach underscores the importance of leveraging cutting-edge technology to ensure that urban planning and environmental management are adaptive and proactive. By integrating advanced computational models and AI-driven insights, the goal is to provide a framework that interprets the past landscape, monitors present changes and anticipates future shifts, ensuring sustainable urban development and effective response to environmental challenges. Section 3.2 illustrates the GeoAI Processing Block, a software solution designed to execute AI algorithms on Remote Sensing data and Open Data and interface seamlessly within the Data Spaces Ecosystem. Section 3.3 outlines the use of the GeoAI Processing Block in monitoring thermal comfort for soft mobility infrastructure in Milan, demonstrating the module's practical application. Conclusions are drawn in the section 3.4.

3.2 GeoAI Processing Block: A serverless AI Module for data space ecosystems

The Implementing Act on *High-Value Datasets*[167] from the European Commission mandates public sector bodies to publish open, machine-readable data belonging to specific thematic categories, including Earth Observation. The Act establishes the protocols for the publication and reuse of high-value datasets, as well as the dissemination of data via APIs. This is one of many initiatives that gravitate around the concept of data space introduced in 1.2.3 and enforced by the EU Commission.

In this context, this section introduces a module for geospatial data processing through AI to integrate into a federated data space ecosystem. Section 1.2.3 described the technical building blocks of data spaces, reported in Figure 1.5. Figure 3.1 reports the same diagram by emphasizing the critical building blocks the proposed module tries to fulfill, delineated within yellow boxes for clarity. This section will go through each building block, explaining the proposed solutions for each requirement. The module's functionalities must encompass a comprehensive suite designed to enhance the processing and analysis of Earth Observation (EO) data within federated data space ecosystems. This block is tailored to execute AI algorithms on remote sensing data and open data, interfacing seamlessly within the data spaces ecosystem. The core functionalities include:

- **Near Real-Time Data Acquisition and/or Analysis:** It facilitates the continuous capture and immediate observation of satellite imagery, allowing for the analysis of current environmental conditions and urban dynamics. This capability is crucial for monitoring dynamic environmental changes and providing timely insights for decision-making.
 - **High Temporal Resolution Monitoring:** the processing block needs to ensure frequent and consistent updates as soon as possible when new data comes. This is vital for tracking ongoing changes and phenomena accurately.
 - **Integration with Other Data Sources:** It must support integrating satellite data with other real-time sources, such as sensors, weather
-



Figure 3.1: Data spaces technical building blocks. Into a yellow box building blocks, the GeoAI Processing Block deals with.

data, and social media feeds. This integration offers a more holistic view of current conditions and enhances the comprehensiveness of environmental monitoring.

- **Event-Based Operation:** the block must support asynchronous data flows. The block receives event notifications to stay updated on changes. This event-driven model enables the system to automatically trigger processes like map updates, model training, and automated Processing based on the data or configuration changes.
- **Interoperability and Standardization:** the GeoAI Processing Block must ensure compatibility and interoperability within the federated data space ecosystem.
- **AI Models Management:** This component needs to manage the use, versioning, and updating of AI models, managing different versions of model weights and configurations.
- **Cost-efficient deployment:** The deployment strategy must emphasize cost-efficiency, scalability, and real-time data processing capabilities, taking into account the event-based nature of the data processing flow.

The effectiveness of the implemented module has been tested within the SPOTTED project [168]. The Spotted project aims to take initial steps towards enabling the implementation of High-Value Datasets within EU

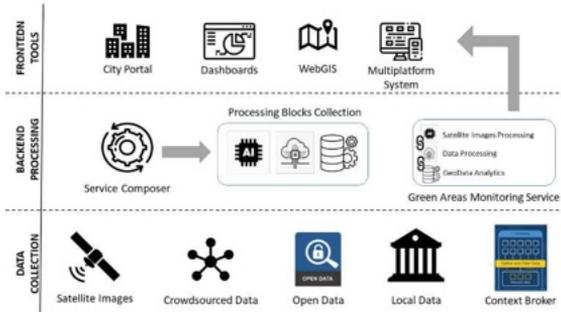


Figure 3.2: SPOTTED Architecture.

territories to benefit local communities. Specifically, the project seeks to streamline the management of green areas by supporting Public Administrations. It aims to assist policymakers in analyzing various data types, including geospatial, earth observation, environmental, meteorological, and statistical High-Value Datasets, thus enhancing their decision-making capabilities.

The SPOTTED architecture follows the three-layer approach (Figure 3.2). The lowermost layer is dedicated to the task of data collection, encompassing the aggregation of data from various sources such as city open data portals, geospatial data derived from Copernicus, and privately sourced data from municipalities. Moving up the hierarchy, the middle layer constitutes the backend processing phase, where the collected data undergoes comprehensive analysis and Processing. Within this stratum, the GeoAI processing block plays a pivotal role. Finally, the upper layer encompasses the frontend tool, exclusively designed to visualize processed datasets.

To adhere to a high level of abstraction, the SPOTTED project has adopted the strategy of furnishing cities with the pertinent datasets, with each city assuming responsibility for rendering these datasets accessible through their respective open data city portals. City portals can integrate data through the European Data Portal [169] into which results are published. This approach affords cities a distinctive access point, enabling stakeholders such as citizens, organizations, businesses, and policymakers to review and visualize the outcomes of the analytical processes.

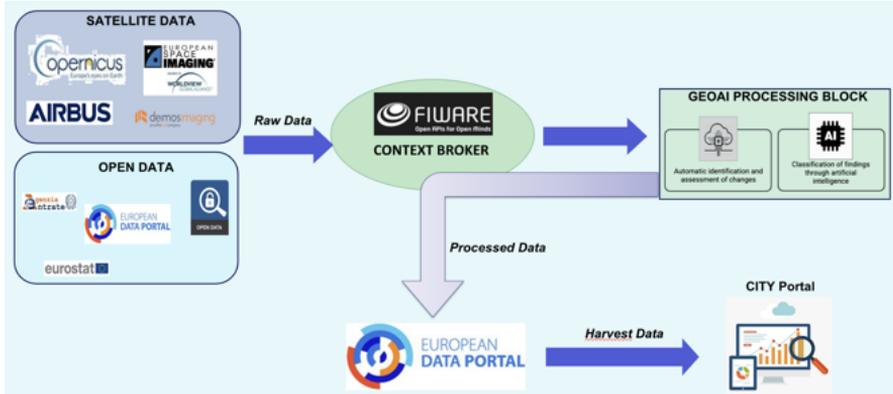


Figure 3.3: SPOTTED Data Flow.

As clearly stated in Section 1.2.3, Technology Building Blocks for Data Spaces (Figure 3.1) must be domain-agnostic. They should rely on open standards, allowing multiple infrastructure and global service providers to emerge and support Data Spaces without getting locked by any particular provider. A middleware layer that acts as a data integration hub is essential since it can collect data from different sources, transform it into a standard format, and make it available to downstream applications. In the context of the Spotted project, the middleware was implemented through FIWARE technologies [170]. FiWare is a proven technology in implementing Data Spaces [171, 172, 173, 174].

For simplicity, this section will not provide an exhaustive examination of all the components that constitute the FiWare solution but will focus on the essential concepts required to understand the implementation of the GeoAI Processing Block. It is recommended that the reader refer to the official documentation for a comprehensive understanding of the technology's specific details [170].

Figure 3.3 illustrates the data flow in the Spotted project. Central to the whole process is the FiWare Context Broker. The Context Broker is the sole mandatory component required for creating an application labeled as "powered by FIWARE." Satellite data and open data repositories serve as primary input sources, feeding raw data into the FIWARE Context Broker. This integration layer is pivotal for mediating and processing

Earth Observation data, which is then channeled to a GeoAI Processing Block. The processed data is disseminated through the European Data Portal to the City Portal.

Any software architecture "powered by FIWARE" revolves around managing a **Digital Twin** data representation of the physical world. This representation uses information from various sources, such as sensors, cameras, and information systems. It is continuously updated and made available in near real-time ("right-time" is also commonly used to indicate that the interval between data collection and accessibility allows for timely reactions). Applications consistently analyze this data (current values and historical records) to automate specific tasks or support informed decisions by end users. Aggregating all Digital Twins modeling the physical world represents the **Context**, while the associated data with Digital Twin attributes is called **context information**. In FIWARE, a Digital Twin is an entity that digitally represents a real-world physical asset (e.g., a road in a city) or a concept (e.g., a weather forecast or a product analysis). Each Digital Twin (i) is universally identified with a URI (Universal Resource Identifier), (ii) belongs to a well-known type (e.g., the road type) also universally identified by a URI, and (iii) is characterized by several attributes.

Attributes of a Digital Twin may vary, ranging from pretty static attributes (e.g., road name) to attributes that change dynamically (e.g., road surface temperature) to attributes that still change but not that often.

The Context Broker component is at the core of the architecture, keeping a Digital Twin representation of the real-world objects and concepts relevant to the problem tackled.

As reported in Figure 3.1, to enable efficient data integration, two pivotal aspects require standardization: the API for accessing Digital Twin data and the data models elucidating the attributes and meanings linked to various Digital Twin types under consideration. FiWare solves the two aspects with:

- The *NGSI API* offers a straightforward yet robust RESTful interface for retrieving context or Digital Twin data. The initial mature iteration of this API was *NGSIv2*. The *NSI-LD API* is the data integration API employed by the Context Broker component.
-

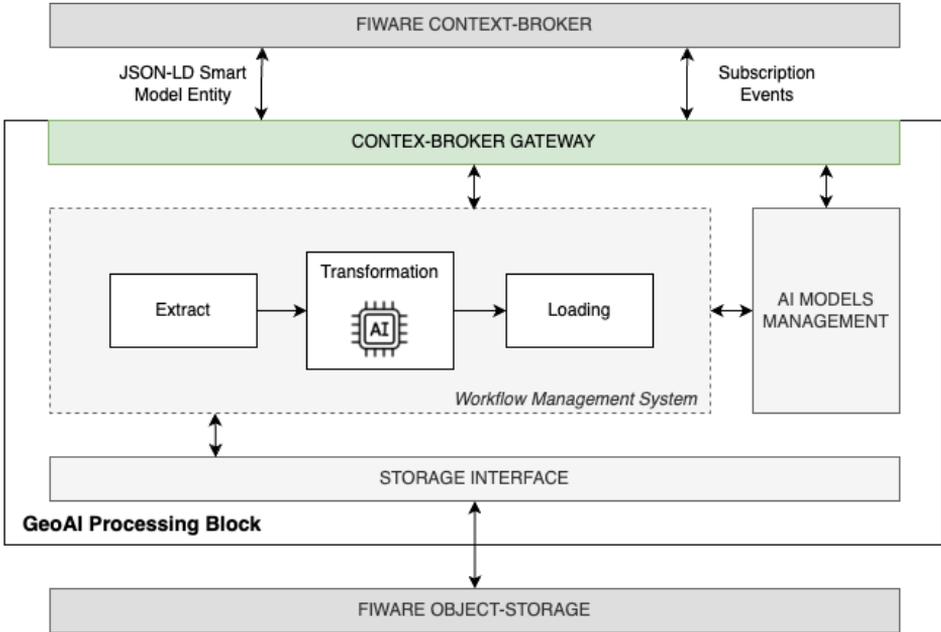


Figure 3.4: GeoAI Processing Block Architecture.

- The *Smart Data Models* initiative by the FIWARE Foundation offers JSON/JSON-LD Data Models compatible with NGSIv2/NGSI-LD APIs. This initiative addresses a significant challenge developers encounter: the various valid but diverse mappings of data model specifications into JSON/JSON-LD. The Smart Data Models initiative alleviates this problem by providing concrete mappings into JSON/JSON-LD that are fully compatible with the NGSIv2/NGSI-LD APIs. This ensures Interoperability and helps developers avoid compatibility issues arising from alternative mappings.

So, the GeoAI Processing Block should be designed to interface with the FIWARE Context Broker and be adept at handling and interpreting the Digital Twin data representation. It should be capable of interacting with the Context Broker via the NGSI API, which enables efficient retrieval and updating of context information. The AI module will utilize this API to access the current state and historical data of Digital Twins, thereby

supporting near real-time Processing.

It must also adhere to the Smart Data Models initiative. This involves using standardized JSON/JSON-LD data models aligned with the NGSIv2/NGSI-LD APIs. By conforming to these models, the module ensures that its processing logic and output are universally understandable and integrable across different systems.

Figure 3.7 shows the architecture of the GeoAI Processing Block within a FIWARE-enabled architecture. This block interacts with the FIWARE Context Broker through a *Context-Broker Gateway*. In essence, the gateway includes the ability to:

- Interface with the NGSI API for data exchange;
- Interpret and process data accessed via entities in the Context Broker effectively;
- Employ standardized data models for compatibility and Interoperability.

In other words, it fulfils the requirements of Interoperability defined in Figure 3.1 (most left column). It manages NGSI-LD Smart Model Entity data and interacts with NGSI API to collect information about geospatial data to collect and process.

The processing part is depicted as an Extraction-Transformation-Loading process. It processes incoming data in a uniform format suitable for machine learning, includes data cleaning, normalization, and augmentation capabilities, and ensures data compatibility with various machine learning algorithms. A workflow management system (WfMS), e.g., Airflow or Prefect, manages the whole process. A WfMS guarantees flexibility in developing and deploying the various data pipelines.

The Storage Interface manages the output generated from models and interfaces with federated data spaces to store these outputs securely. Includes mechanisms for tagging and organizing outputs for easy retrieval and analysis.

The AI Models Management component oversees AI models' use, versioning, and updating. It manages different versions of model weights and configurations, integrates with federated data spaces for storing and

retrieving model weights, ensures up-to-date models, and tracks model evolution over time. In particular, when a model is trained, its metadata and weight file are stored. The metadata is registered in the Context Broker, and the weight file is uploaded to the object storage. Versioning is handled during this phase, with each training session resulting in a new version. During model deployment, the models are deployed based on the metadata and version information stored in an entity in the Context Broker. The module uses this information to pull the correct model weights from object storage for deployment.

The GeoAI Processing Block is event-based. The system effectively utilizes FIWARE subscriptions to receive event notifications and stay updated about changes. The module can subscribe to changes in any entity in the context. Although not exhaustive, below is a list of some implemented uses of subscriptions:

- **Trigger Model Training:** Set up subscriptions to monitor for new or updated data in the federated data spaces. The system can automatically retrain AI models when new data arrives or specific criteria are met (like a particular volume of new data).
- **Version Control:** Automatically update model or workflow versions based on training data or performance metrics changes, ensuring that the latest, most effective model is in use.
- **Configuration Changes:** Trigger reconfiguration of the inference engine or related components when a new model version is available or when there are changes in the model's metadata.
- **Automated Processing:** When new data enters the system, subscriptions can trigger the start of the data pipeline to process new data.

Figure 3.5 illustrates the integration of the GeoAI Processing Block with the FiWARE Context-Broker. In the middle of the diagram is the FiWARE Context-Broker itself. It is designed around a RESTful architecture, meaning all interactions with it are conducted through REST APIs. It is the central node that handles "Context Information," acting as a mediator between context producers and consumers. Inside the Context-Broker, there's a section for *Subscriptions*. These subscriptions are created

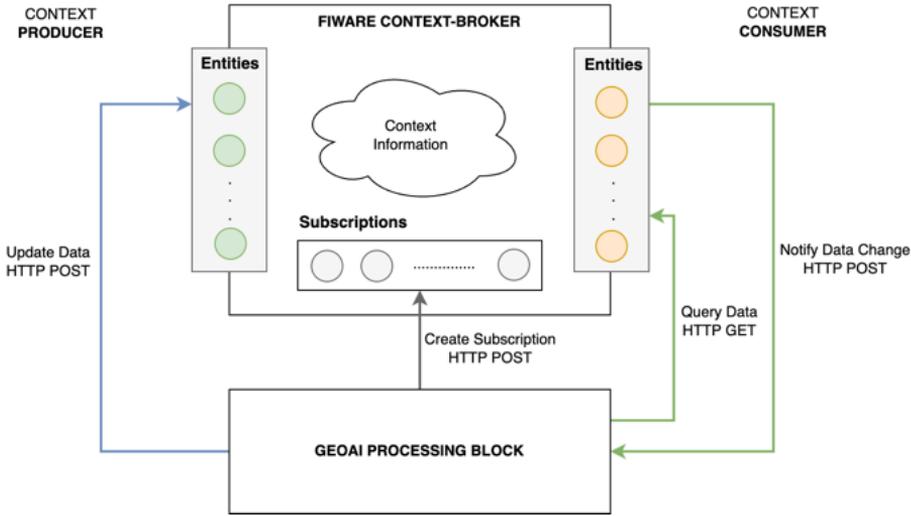


Figure 3.5: Integration detail of the GeoAI Processing Block with the FiWARE Context-Broker.

by making an HTTP POST request to the Context-Broker. They enable the system to listen for specific changes or events related to the entities. On the left side, the context producer flow is depicted. The module sends updated data entities to the FiWARE Context-Broker via an HTTP POST request. The right side shows the context consumer flow. The processing block can query the Context Broker to retrieve data using an HTTP GET request. It can also receive notifications of data changes from the Context Broker through an HTTP POST request. This occurs when a subscription is triggered due to a change in the entities' data, and the Context Broker pushes this information to the context consumer. Overall, the figure showcases a dynamic and interactive system that allows for the real-time processing and exchange of context information, with the GeoAI Processing Block providing advanced data processing capabilities within the ecosystem.

All the mechanisms described are tightly coupled with the entities managed through the context broker. So, to ensure Interoperability, it is necessary to define entity schemas a priori and, where possible, use some of the existing Smart Models. However, to date, no entities within the Smart

Model catalog can manage the lifecycle of Artificial Intelligence algorithms, so as part of this research, new Smart Models were proposed to achieve the goal. The following is an example of a smart model for managing the lifecycle of AI model weights:

```
{
  "id": "weightfile-12345",
  "modelName": "LandCoverClassifier",
  "version": "v2.3",
  "createdAt": "2023-05-10T12:00:00Z",
  "updatedAt": "2024-01-20T12:00:00Z",
  "fileSize": 455000,
  "fileFormat": ".h5",
  "sourceUri": "https://cloudstorage.example.com/
models/landcover/v2.3/weights.h5",
  "checksum": "4f76dbdf8",
  "trainingDataReference": "dataset-7890",
  "performanceMetrics": {
    "accuracy": 0.95,
    "precision": 0.94,
    "recall": 0.93
  },
  "description": "Improved accuracy in urban areas.",
  "status": "active",
  "associatedProject": "UrbanExpansionStudy2024",
  "tags": ["urban", "AI", "satellite-imaging"],
  "license": "GPL-3.0",
  "owner": "EO-Dept-23",
  "validUntil": "2025-05-10T12:00:00Z"
}
```

Where:

- id (string): A unique identifier for the weight file.
 - modelName (string): The name of the AI model this weight file belongs to.
-

- `version` (string): Version of the weight file, indicating updates or changes.
 - `createdAt` (dateTime): Timestamp of when the weight file was initially created.
 - `updatedAt` (dateTime): Timestamp of the most recent update to the weight file.
 - `fileSize` (number): Size of the weight file in bytes.
 - `fileFormat` (string): Format of the weight file (e.g., .h5, .pb, .pt).
 - `sourceUri` (string): URI where the weight file is stored. This could be a link to a cloud storage location, a database, or a file system.
 - `checksum` (string): A checksum value for verifying the integrity of the weight file.
 - `trainingDataReference` (string): Reference to the dataset used for training the model, possibly linked to a dataset identifier.
 - `performanceMetrics` (object): An object containing key performance metrics of the model with this weight file (e.g., accuracy, precision, recall).
 - `description` (string): A brief description of the weight file, including notes on specific changes or updates in this version.
 - `status` (string): Current status of the weight file (e.g., 'active', 'deprecated', 'in-testing').
 - `associatedProject` (string): Identifier for the project or initiative this weight file is associated with.
 - `tags` (array of strings): Tags for categorizing or adding metadata (e.g., 'environmental', 'urban-analysis').
 - `license` (string): If applicable, license information for the weight file.
 - `owner` (string): Identifier of the owner or responsible party for the weight file.
-

- `validUntil (dateTime)`: A timestamp indicates when the weight file is considered valid or up-to-date.

This Smart Data Model, integrated into the GeoAI processing module, enables effective management and tracking of AI model weight files throughout their lifecycle.

Finally, all metadata for the produced dataset has been managed by the FiWARE Context-Broker using the Data Catalogue Application Profile (DCAT-AP) data model. The DCAT-AP data model is a metadata model developed by the European Commission to support data exchange between different data portals and catalogs. It provides a standard way to describe datasets, making it easier to discover, access, and reuse data across various organizations and applications. Using the DCAT-AP data model was a natural choice since the European Union recommends it as part of its open data strategy. It helps promote data interoperability and sharing across different organizations and applications, making exchanging and reusing data for various purposes easier. Using the DCAT-AP data model, the proposed solution assured the data value building block highlighted in figure 3.1, highlighting the solution's effectiveness to be easily integrated into a dataspace ecosystem.

This concluding section explores the deployment strategy for the GeoAI Processing Block. Several decisive factors drove the adoption of a serverless architecture underpinned by a Kubernetes cluster. This architectural decision promises cost-efficiency, as it aligns financial outlay directly with resource consumption, a boon for workloads with unpredictable patterns. Scalability is another hallmark of this approach, with the serverless framework offering responsive scaling to match data processing demands. An event-driven computational model enhances this architecture's aptitude for real-time data processing, a vital requirement for managing the dynamic data influx within federated data spaces. Finally, this architecture not only ensures cost efficiency and automatic scalability but also supports, at least in theory, the strategic placement of computational resources "in proximity to data sources," i.e., code-to-data paradigm. This geographical alignment minimizes latency and accelerates the processing pipeline, which is essential for real-time or near-real-time data analysis applications. By allowing deployment adjacent to data repositories, the serverless approach further enhances the system's responsiveness and reduces the overheads

associated with data transfer and management.

3.3 Monitoring Soft-Mobility Thermal Comfort

3.3.1 Downscaling Land Surface Temperature

The Land Surface Temperature (LST) is the radiative skin temperature of the land surface, as measured in the direction of the remote sensor. LST is a mixture of vegetation and bare soil temperatures. Its estimation further depends on the albedo, the vegetation cover and the soil moisture. Because both respond rapidly to changes in incoming solar radiation due to cloud cover and aerosol load modifications and diurnal variation of illumination, the LST displays quick variations too. In turn, the LST influences the energy partition between ground and vegetation and determines the surface air temperature. Land surface temperature, described as one of the most important parameters for atmospheric and land surface interactions, material cycles and energy exchanges in the terrestrial ecosystem at regional and global scales, plays an essential role in the energy balance of land surfaces and the solution of biophysical parameters.

LST has been widely used in hydrological equilibrium assessments, global warming studies, urban heat island effect assessments and surface evapotranspiration calculations. A spatial resolution of 10 meters for LST data is crucial for urban analysis due to urban environments' complex and heterogeneous nature.

Algorithms for statistical downscaling of land surface temperature, utilizing various scaling factors and regression models, have been employed to generate LST at higher spatial resolutions using data derived from Landsat-8/9 LST. Various methodologies for downscaling land surface temperature (LST) have been discussed in scholarly works, highlighting physical and statistical techniques. Statistical downscaling, particularly, leverages the relationship between LST and supplementary data such as land cover indicators, emissivity, and additional variables. Commonly, many methods employ the normalized difference vegetation index (NDVI) at fine resolution (FR) for enhancing LST resolution, although NDVI alone may not fully capture the LST variability within urban settings.

This thesis illustrates a regression-based statistical LST downscaling

approach [175, 176] that utilizes a variety of spectral indices across diverse urban areas.

The downscaling approach is based on a regression model using spectral indices derived from Landsat 8/9 bands as predictors and LST as the predicted variable. After that, with the same model, the equivalent spectral indices derived from Sentinel-2 satellite bands are used to predict LST at 10 m resolution.

After the regression, a residual correction is applied to the estimated values. The regression output, LST_{10m} , incorporates the addition of regression residuals, ΔLST_{10m} , back into the downscaled LST map. This process enables the restoration of the original broad-scale temperature distribution by means of reaggregation while also accounting for the spatial variability in LST that arises from factors beyond the predictors used. The residual is computed as the difference between the model estimation and the correspondent observed LST, resampled to 10 m spatial resolution with a Gaussian kernel of 30 m size.

Given the scarcity of high-resolution Land Surface Temperature (LST) benchmark data, a widely recognized *upsampling-downsampling* [176, 177, 178] method for validation was adopted. In this method, the original LSTs with a spatial resolution of 30 m and scaling factors for 100 m LSTs were upscaled to 200m and 1000m resolutions. The LSTs upscaled to 200m served as benchmarks to evaluate the accuracy of LSTs downscaled from 1000m. It's important to note that the original 100m LSTs were not used directly for validation. This decision was due to the slight discrepancies between the upscaled scaling factors and the original 100m resolution LSTs, a phenomenon attributed to the point spread function effect [179]. The root-mean-square-error (RMSE) metric was used to evaluate the downscaled LSTs. RMSE is widely valued for its effectiveness and reliability in measuring the absolute differences between downscaled LSTs and their benchmarks.

The indices used for regression in this study were as follows:

- NDVI (Normalized Difference Vegetation Index);
 - NDMI (Normalized Difference Moisture Index);
 - NDWI (Normalized Difference Water Index);
-

Model	MAE	MSE	RMSE	R2	RMSLE	MAPE
0 CatBoost Regressor	1.785600	5.869000	2.420100	0.739200	0.007800	0.005800
1 Extra Trees Regressor	1.814600	6.010200	2.448200	0.732700	0.007900	0.005900
2 Gradient Boosting Regressor	1.819900	6.028800	2.453400	0.732200	0.007900	0.005900
3 Light Gradient Boosting Machine	1.809500	6.047900	2.456100	0.731200	0.007900	0.005900
4 Random Forest	1.820000	6.074400	2.461500	0.730000	0.007900	0.005900
5 Extreme Gradient Boosting	1.826700	6.099300	2.467300	0.729000	0.008000	0.005900
6 Linear Regression	2.033900	7.628700	2.759400	0.661000	0.008900	0.006600
7 Bayesian Ridge	2.039400	7.642400	2.762000	0.660400	0.008900	0.006600
8 Random Sample Consensus	1.996700	7.932200	2.812000	0.647400	0.009100	0.006500
9 Ridge Regression	2.096100	8.054300	2.835900	0.642300	0.009200	0.006800
10 K Neighbors Regressor	2.230800	8.983500	2.994700	0.601100	0.009700	0.007200
11 AdaBoost Regressor	2.505400	9.400200	3.063800	0.582700	0.009900	0.008200
12 Huber Regressor	2.179800	9.754300	3.118500	0.567000	0.010100	0.007100
13 Orthogonal Matching Pursuit	2.512100	11.183900	3.342000	0.504000	0.010800	0.008100
14 Decision Tree	2.488400	11.400400	3.373800	0.493600	0.010900	0.008100
15 Support Vector Machine	2.994100	16.107800	4.009500	0.286200	0.012900	0.009700
16 Elastic Net	3.364600	18.481600	4.297200	0.181100	0.013900	0.010900
17 Lasso Regression	3.370100	18.486600	4.297900	0.180900	0.013900	0.010900
18 TheilSen Regressor	2.385500	21.175100	4.578600	0.062600	0.014700	0.007800
19 Lasso Least Angle Regression	3.886700	22.610500	4.753800	-0.001500	0.015300	0.012600
20 Least Angle Regression	3.562500	28.669200	5.233800	-0.282600	0.017000	0.011600
21 Passive Aggressive Regressor	5.346400	59.288600	7.568100	-1.640200	0.024800	0.017400

Figure 3.6: Comparison between several machine learning regressors for the task of LST downscaling.

- MNDWI (Modified Normalized Difference Water Index);
- NDBI (Normalized Difference Built-up Index);
- SAVI (Soil-Adjusted Vegetation Index);
- NDDI (Normalized Difference Drought Index);

Using different indexes can increase the model's robustness to predict LST in different conditions and for different surfaces. Several machine learning regression algorithms have been compared. The dataset included the Municipality of Naples for the dates 12/05/2022 and 10/08/2023 and Milan for the dates 05/02/2023 and 15/07/2023. A hold-out technique to split the dataset into training and validation data with an 80%-20% ratio was adopted.

The result of the comparison between several machine learning regressors is depicted in Figure 3.6.

Finally, a Catboost regression model was chosen. After that, the model was tuned using a grid search. Table 3.2 shows the results compared with

state-of-the-art methods. Even though the improvement is marginal, the proposed approach outperforms the state of the art.

3.3.2 The SPOTTED Use Case

The SPOTTED project has initiated and conducted service-related experiments within the scope of varied use cases in Milan, Helsinki, and Naples. These services are powered by the SPOTTED Platform, which leverages open urban data and satellite imagery processed by the GeoAI Processing Block to extract actionable insights. Specifically, the initiative supports the city's green transition in Milan, focusing on sustainable urban development and environmental conservation. The GeoAI Processing Block's analysis for Milan aims to pinpoint optimal locales for green initiatives, thereby equipping policymakers with crucial data to drive informed decisions for the city's ecological advancement.

Utilizing satellite imagery alongside auxiliary data streams, the solution stands poised to furnish Milan's policymakers with precise, current environmental data. This vital information will underpin the formulation of bespoke sustainability initiatives, precisely aimed at measurable reductions in the city's environmental impact.

Within the discussed context, a solution has been proposed to monitor thermal comfort specifically for soft mobility infrastructures, focusing on bicycle paths. This solution harnesses environmental data, particularly the land surface temperature (LST), to assess the comfort levels experienced by cyclists, integrating open data pertaining to the bike paths. Monitoring the thermal comfort of bike paths is essential to bolster cycling as an eco-friendly transport option, fostering a healthier urban population and assisting in urban development strategies. It ensures cyclists' well-being by mitigating heat stress, which can lead to increased cycling frequency. The data gathered aids city planners in enhancing current infrastructure and formulating adaptive measures to climate change, ensuring bike paths remain inviting and practical in the face of rising temperatures.

This proposed methodology exemplifies a user-centric approach, as it imbues new value into the open data made available by users in a dataspace, in this instance, the municipality. Through this method, user-provided open data is enriched with additional layers of information derived from EO data, leading to enhanced insights that directly benefit the

user community.

The solution was implemented through a workflow executed through the GeoAI Processing Block described in the previous section. The first step was to acquire the required datasets, including Sentinel-2, Landsat-8 satellite imagery, and the dataset of biking paths provided by the Municipality of Milan via FiWare. Data were pre-processed to correct atmospheric distortions, radiometric calibration, and geometric correction. In addition, this step involved removing any cloud cover or shadows that may obscure the data. A pivotal step involved employing the Machine Learning algorithm described in 3.3 to refine land surface temperature measurements from Landsat-8 to a finer 10 m resolution using Sentinel-2 data, catering to the narrow expanse of bike paths. The temperature difference of each area from a rural reference area was calculated to extract areas affected by Urban Heat Islands. Finally, through zonal statistics, mixing the data with maps of bicycle and pedestrian paths provided by the city of Milan and population data collected from the Worldpop repository [7], areas with higher priority for intervention were identified.

Figure 2 exemplifies the results of the integrated approach, showcasing the significant value added by combining open city data, such as bike paths, with Earth Observation data and Machine Learning analysis. This figure illustrates how the synthesis of diverse data sources can yield comprehensive and insightful results, demonstrating the effectiveness of this multifaceted methodology in urban planning and environmental analysis.

3.4 Conclusion

This chapter has delved into the critical role of continuous monitoring in the context of Sustainable Development Goals (SDGs) and urban planning, illustrating the dynamic nature of environmental changes and the necessity for adaptive strategies to manage these changes effectively. Through a detailed examination of the SPOTTED project and the development and implementation of the GeoAI Processing Block within federated data space ecosystems, we have seen a practical application of combining Earth Observation (EO) data with advanced Artificial Intelligence (AI) techniques to monitor current conditions.

The chapter has highlighted the importance of leveraging serverless



Figure 3.7: Surface Heat Island Analysis for the Soft-Mobility in the Municipality of Milan, Italy. The hottest bike lanes are indicated in red, while the coolest ones are marked in blue.

architecture and cloud computing in processing large volumes of data in real-time, ensuring scalability, cost-efficiency, and the ability to respond swiftly to environmental changes. By integrating standardized data models and utilizing event-driven computational models, the GeoAI Processing Block stands as a testament to the power of technology in enhancing our ability to monitor and manage the environment effectively.

In the case of Milan, the application of these technologies to assess thermal comfort for soft mobility infrastructure like bicycle paths has shown how data-driven insights can inform urban planning decisions, promoting sustainable transportation options and enhancing urban liveability. This approach not only supports the immediate needs of urban planners but also contributes to the broader goals of sustainable urban development by providing a mechanism for continuous environmental monitoring and management.

In conclusion, the advancements discussed in this chapter represent a significant step forward in our ability to monitor the present. By harnessing the power of EO data, AI, and cloud computing, we can develop more resilient and sustainable urban environments. This journey towards integrating cutting-edge technology into environmental monitoring and urban planning is essential for achieving the SDGs and ensuring a sustainable future for our cities and their inhabitants.

Table 3.1: Comparison of the Machine Learning requirements between the pillars *Learning from Past* and *Monitoring the Present*

	Learning from Past	Monitoring the Present
Data Source and Nature	Historical satellite data. Data might have varying resolutions and qualities.	Near-real-time satellite imagery.
Objective and Application	Understand trends, patterns, and changes over time.	Focus on the occurred changes.
Model Training and Validation	Training on diverse datasets from different time periods. It may need calibration to account for variations in data quality and sensor technology over time.	Models trained on past data. They might be fine-tuned more frequently to adapt to the latest data characteristics.
Temporal Dynamics	Long-term temporal dynamics and changes, incorporating techniques to understand and predict trends over extended periods.	Emphasis on short-term dynamics, often requiring the models to adapt to rapidly changing conditions.
Predictive Modeling	Focus on forecasting long-term future scenarios based on historical trends.	Usually short-term, focusing on immediate future conditions or events.
Data Volume and Processing	Handling large volumes of historical data	While the volume of data can be substantial, the focus is on the efficient processing of current data streams for timely analysis
Generalization vs. Specificity	Models are often designed to generalize across different time periods and conditions.	Models may be more specialized, tailored to specific current conditions or events.

Table 3.2: Results for the task of LST downscaling

Area of Interest	Date	RMSE	RMSE[175]	RMSE [176]
Naples	12 May 2022	0.73	0.96	1.23
Naples	10 August 2023	1.18	1.25	1.22
Milan	5 February 2023	1.30	1.75	1.68
Milan	15 July 2023	1.75	2.50	2.62

Chapter 4

Building the Future

The best way to predict the future is to create it.

Peter Drucker

4.1 Introduction

Cities around the world are expanding at an unprecedented rate. This expansion is not just in terms of population growth but also urban areas' physical and infrastructural development. As cities grow, they face numerous challenges, including the need for efficient land use, infrastructure development, and transportation management. Alongside urbanization, environmental concerns such as climate change, biodiversity loss, and pollution are becoming increasingly critical. Urban areas are major contributors to environmental issues, yet they are also significantly affected by them. For instance, urban heat islands, air quality, and greenery management are significant challenges.

Urban planning decisions have long-lasting impacts. So, planners must consider the environmental impact of urban development and find ways to mitigate negative effects. With the increasing complexity of urban environments, there is a growing need for tools that provide precise and detailed insights into how different planning decisions might play out. Precision in planning helps make informed decisions that can minimize risks, optimize

resource use, and ensure the well-being of urban populations.

This section tries to address a solution for two pivotal questions that an urban planner consistently does:

Where is the most beneficial location for a planned urban intervention to maximize its "impact"?

Is it possible to foresee the effects of an urban intervention before its actual implementation?

For example, an urban planner tasked with the installation of a new park in the city might pose the following inquiries: (i) Where is the most advantageous location for the park to optimally enhance urban livability? and (ii) Is it possible to ascertain the sustainability impacts of introducing a new urban park in a specific area, before its practical realization?

Section 4.2 addresses the first question. It explores the methodology for identifying the most beneficial locations for implementing Nature-based Solutions, aiming to maximize social impact. This analysis begins by examining the area's existing environmental and social conditions. Section 4.3 then delves into the second question, exploring the possibility of predicting the sustainability impacts of urban interventions before their implementation.

At this juncture, a discerning reader may contend that the approach of Section 4.2 aligns more closely with the "Learning from Past" methodology rather than "Building the Future". However, the distinction lies in the intent behind the posed questions. A retrospective query like "What is the worst neighborhood in the city regarding a certain observational variable?" remains anchored in the past. In contrast, re-framing the question to "What is the worst neighborhood in the city based on certain variables to plan future actions?" shifts the focus to future-oriented planning. This perspective viewpoint justifies the inclusion of the use cases described in Section 4.2 in this chapter.

Section 4.3, instead, addresses the second question. It introduces a methodology to foresee the effects of urban interventions before their actual implementation by utilizing Generative Adversarial Networks (GANs) to generate synthetic multispectral satellite imagery. This innovative approach enables urban planners and decision-makers to visualize and an-

alyze the potential outcomes of various urban development scenarios, including nature-based solutions (NbS), without requiring physical changes to be made first.

4.2 Planning Urban Forest as Nature-based Solutions

4.2.1 Introduction

The requirement to reassess the present planning and management of cities and urban settlements from a resilience perspective is an essential topic of discussion in mitigation and adaptation policies at an international level, particularly concerning climate change debates. The United Nations 2030 Agenda for Sustainable Development highlights, in target 11.3, the need to "... enhance inclusive and sustainable urbanization and capacity for participatory, integrated and sustainable human settlement planning and management in all countries" [180]. Target 11.7 further emphasizes the necessity for "... access to safe, inclusive and accessible, green and public spaces, particularly for women and children, older persons and persons with disabilities" (ibid.). Specifically, the 2022 Sustainable Development Goals Progress Report highlights the necessity for policymakers and city officials to "... consider the distribution of open public spaces and green areas throughout the city" as they "... work to redesign and retrofit the spatial configuration of urban areas" [181]. Integrating urban green spaces as public infrastructure providing social and environmental services is crucial in addressing green inequality in urban planning [182]. Since Milton Keynes's Forest City project, several scholars have addressed developing and integrating green infrastructure in urban and landscape planning. These practices aim to embrace sustainability and resilience [183, 184], combining urban forestry planning with the management of urban green spaces [185].

Investigating the abundance, distribution, and types of green spaces is crucial in effectively guiding the design and planning of urban green spaces and forestry to ensure the provision of a range of human benefits [186, 187], related to public health [188] through ecosystem services [189]. These effects include mitigating ambient air pollutant levels by ur-

ban forestry [190] and providing food through urban food forestry, which combines urban agriculture, urban forestry, and agroforestry [191]. Others, such as Chen and Huang [192], indicated that the demand for green space could greatly vary across neighborhoods, and areas with a high population density would require more green space. The authors emphasized that future urban greening projects should aim to serve areas with low green space supply and high demand from people with limited access to green resources.

Defining the quantity and quality of green space accessibility is paramount, as it is often recognized as the difference between spatial heterogeneity and social differentiation [193]. Apart from accessibility aspects, the most advanced and up-to-date techniques used in planning urban green systems often involve satellite imagery and GIS tools focusing on greenness (e.g., NDVI data) [194].

In decision-making processes on the location of new urban green spaces [195], issues about equity, justice, and the reduction in socio-spatial inequalities at the city and neighborhood scale are surprisingly given only marginal consideration.

When focusing on urban forestry (UF), a comprehensive assessment of the existing tree canopy is often required for tree-planting programs to support the setting of specific community targets. Remote-sensing technologies such as GIS and satellite imagery, where optical imagery is combined with LiDAR data [196], are often helpful for capturing information on the tree canopy. Ideally, tree canopy maps should be accurate to the scale of individual trees, allowing for analysis of the tree canopy at different scales, both at the city level and at the individual plot level [197].

This section investigates the concept of applying Urban forestry as a Nature-based Solution (NbS). The Fifth Session of the United Nations Environment Assembly (UNEA-5) defines Nature-based Solutions (NbS) as "actions to protect, conserve, restore, sustainably use and manage natural or modified terrestrial, freshwater, coastal and marine ecosystems, which address social, economic and environmental challenges effectively and adaptively, while simultaneously providing human well-being, ecosystem services and resilience and biodiversity benefits" (UNEP, 2023). Specifically, it emerges that the objective of NbS is to promote and protect biodiversity, supporting the delivery of several Ecosystem Services (ES).

If considered superficially, the concept of NbS may present some overlaps with those of Green Infrastructure (GI), as they are both part of the group of metaphors that are mobilized to describe the role and the functions of natural components in urban environments [198]. However, there are key differences, as highlighted by [199]. Thus, while on the one hand, GI represents (mainly) a network of green/blue areas, the definition of NBS is broad, encompassing the use of nature to address environmental, social, and economic challenges. Specifically, the literature suggests that NbS should be conceived of as an umbrella concept [200, ?] that encompasses other concepts such as ES [198, 201] and GI. In addition, scholars recommend that NbS should not be merely a relabelling of previous practices but a unique concept targeting urban sustainability issues [202].

So, applying Urban forestry as Nature-based relies on the concept of Urban Forestry to address societal challenges and play an essential role in human well-being and biodiversity [203]. In this context, the need to consider and expand the range of socio-economic and environmental data and assess the social impacts of NBS scaling-up is emphasized [204, 205, 206]. The proposed research introduces an approach that overcomes this gap by explicitly considering social and environmental issues, proposing a method potentially adaptable to any urban context. Not considering exclusively green deficit distribution and including socio-economic data in the decision-making process about NBS's implementation, the proposal encompasses the main critiques that are moved to the NbS approach, which is considered as 'mechanistic and seemingly apolitical' [207]. In order to eschew the accusations of contributing to green gentrification and the reproduction of the urban neoliberal model addressed to NbS implementations, the integration of NbS in urban planning necessarily has to include a wide range of socio-environmental factors as criteria for the localization of urban green, echoing the call for 'just nature-based solutions' [208].

Reflecting on this theme, the research explores the main debate on the "greening city," trying to address the following question:

Where can NbS be implemented within the city in order to maximize its social impact?

Without adequate planning [209], the creation of urban green can increase urban social exclusion and stimulate gentrification, as it has been

widely discussed [210, 211, 212]. Surprisingly, issues related to equity, justice, and reduction of socio-spatial inequalities at the urban and neighborhood scale are only marginally considered in the decision-making processes concerning localizing new urban green areas [195]. This is also the case with NbS. While there is a long-lasting debate on GIS localization and planning in urban areas [213], the recent debate on NbS still lacks significant contributions on how to integrate NbS in urban planning, combining technical criteria with social objectives, with few exceptions [214, 215]. Planning for NbS at the urban and neighborhood scales is also essential. Nonetheless, this line of thinking can be found in European policies (e.g., EU Biodiversity Strategy for 2030, EU Green Infrastructure Strategy) aimed at ensuring the supply and quality of green spaces, as well as through the guidance provided by international and European organizations to ensure adequate access to green spaces (World Health Organization recommendation for access to 0.5 – 1.0 ha green space within 300 m).

Within the roster of recommendations from international organizations, there is a place for the International Union for Conservation of Nature (IUCN) 3-30-300 rule, proposed by Cecil Konijnendijk van den Bosch: at least three trees (of a 'decent size') visible from citizens' home; at least 30% of tree canopy cover in each neighborhood; 300 m from the nearest park or green space. The neighborhood characterizes the spatial unity of the rule. This is not the only rule that is applied to have a positive effect on the structure and diversity of urban forests. Among the most famous, used, and debated rules is the 10 – 20 – 30 rule [216]. However, that rule does not focus attention on the benefits provided by urban forests, as the one proposed by [217] does. The 3-30-300 rule [217] represents a message that policymakers can remember but takes up research and the recent debate about the planning and benefits of green spaces in urban settings. Among the various scientific literature, it appears necessary to note the importance of nearby, especially visible, green spaces for mental health and well-being [218]. This assertion was reinforced during the pandemic, emphasizing the importance of greenery near homes [219]. In addition to WHO research and recommendations, [220] emphasized the importance of proximity to high-quality green spaces that can be used for recreation, generally reachable within 5 min or 10 min, though often the traveled distance to the most used UGS is well beyond a 300 – 500 m buffer distance [221]. Although

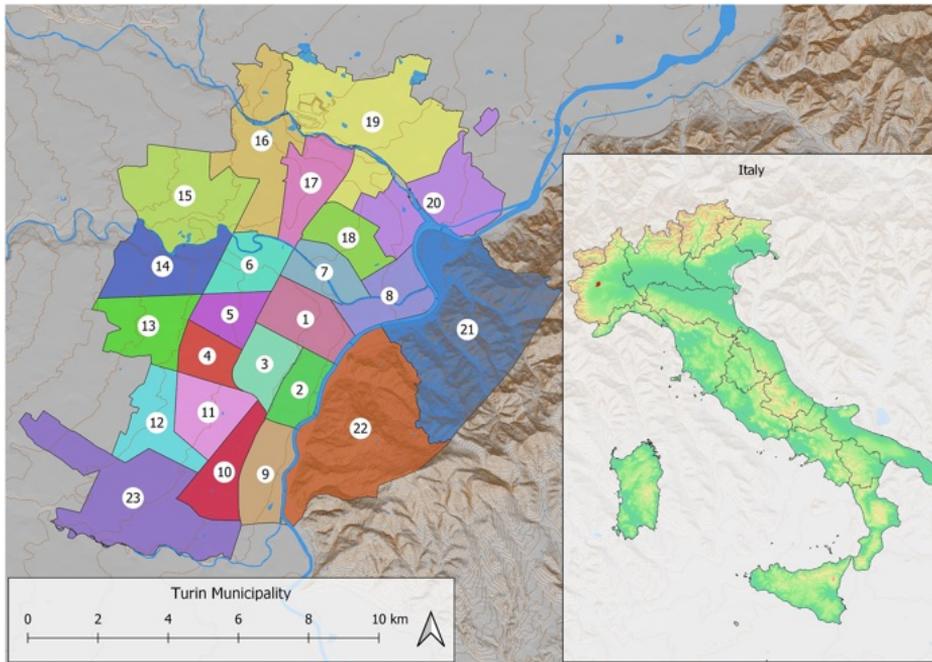


Figure 4.1: Location of the city of Turin and subdivision into neighborhoods.

this rule is proposed and promoted by the IUCN, no scientific articles on its application and modification have yet been published. This is probably due to the recent dissemination of this information.

This research moves the first step in this direction, linking identifying suitable areas for new NbS to implementing a "socially-oriented" UF localization framework, like the 3-30-300 rule, to more explicitly social and environmental criteria. The research meets some of the Sustainable Development Goals set by the United Nations, including Goal 3 (target 3.9), Goal 11 (targets 11.6, 11a, and 11b), Goal 12 (targets 12.2 and 12.8), and Goal 13 (targets 13.1 and 13.2) [180].

Expanding an existing method, Section 4.2.2 develops and tests - in the city of Turin, Italy - an index aiming to identify the parts of the city that could have precedence in the implementation of NbS. With an analogous approach, Section 4.2.3 presents preliminary results contributing to urban

forestry planning to support small and medium-sized administrations. The city of Asti (Piedmont, Italy), a compact urban center with a significant Roman and medieval past located in the northwest of Italy, serves as the case study. Conclusions are drawn in Section 4.4.

4.2.2 The case study of Turin

Materials and Methods

Turin is the capital of the Piedmont region in the northwest of Italy and is located on the western edge of the Po Valley. The city covers 130.2 km² in a flat area (239 m a.s.l.) and a hilly area reaching an altitude of 715 m a.s.l., where 878,074 inhabitants live (ISTAT, 2018). The 37% of the surface area of the City of Turin is composed of green areas (48 km² out of a total of 130 km²), with 55.43 m² of green space per inhabitant. Of this 37%, 5% is made up of public and private agricultural areas [222].

Turin is the capital of the Turin metropolitan area, which comprises 312 municipalities and is characterized by greenways, mainly concentrated along river strips, forming an essential ecological network.

Figure 4.1 shows the general framing of the study area concerning the Piedmont region and the entire Italian peninsula. For more information about the relationships between numerical identifiers and neighborhood names, please see Table 4.1.

In the implementation of the 3-30-300 rule, environmental and social characteristics have also been taken into account. This rule involves ensuring the following criteria:

1. Presence of a minimum of 3 trees of significant size visible from citizens' homes;
2. A tree canopy cover of at least 30% in each neighborhood;
3. Proximity to the nearest park or green space within 300 meters.

To better consider the climatic-environmental conditions of the city, the Urban Cooling model was considered [223]. Instead, for the social data, the population density of Turin was considered [224]. Figure 4.2 shows the workflow methodology applied.

Table 4.1: Detail of neighborhoods, number of inhabitants, neighborhood surface (data elaboration from WorldPop project [7], 2021).

Neighborhoods	N° inhabitants (in)	Neighborhood surface (Km ²)
1. Centro	39,651	3.77
2. San Salvario	35,901	2.33
3. Crocetta	34,176	2.77
4. San Paolo	34,955	2.21
5. Cenisia	39,114	2.33
6. San Donato	49,021	3.02
7. Aurora	39,300	2.67
8. Vanchiglia	30,558	3.38
9. Nizza Millefonti	28,207	3.50
10. Mercati Generali	48,803	3.46
11. Santa Rita	56,747	3.57
12. Mirafiori Nord	43,720	3.79
13. Pozzo Strada	57,154	4.22
14. Parella	47,221	4.91
15. Le Vallette	41,400	7.54
16. Madonna di Campagna	41,047	7.40
17. Borgata Vittoria	39,833	3.86
18. Barriera di Milano	48,081	2.83
19. Falchera	26,186	12.62
20. Regio Parco	28,171	6.92
21. Madonna del Pilone	14,167	15.50
22. Borgo Po e Cavoretto	19,541	13.60
23. Mirafiori Sud	35,334	11.43

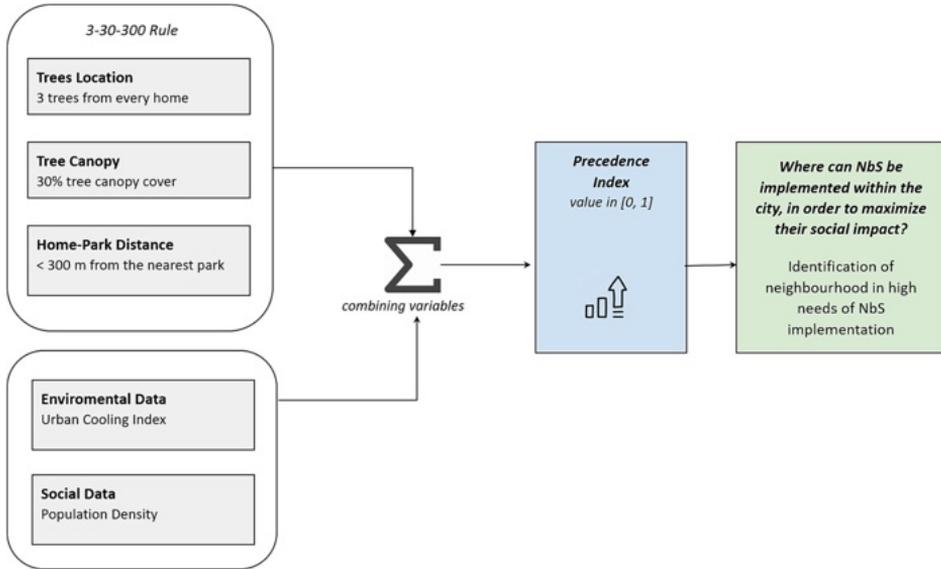


Figure 4.2: The research workflow methodology.

Specifically, the data acquired to conduct the research refer to geospatial datasets that are publicly accessible or provided by the municipal administration. Although the proposed methodology could be easily applied to commercial data, the thesis focused on using publicly accessible and globally available data to propose a practical approach easily scalable to other locations.

In addition, a precedence index has been calculated that combines the different parameters into a single score using additive weighting. Data analysis was first performed at the urban scale and then at the neighborhood level. Neighborhoods were ranked according to their precedence index scores, from lowest to highest. The neighborhoods with the lowest scores represent the areas with the most significant needs for NbS implementation.

All data used refer to the same year, namely 2018. The reasons are as follows: the trees visible or otherwise planted around the buildings must be of "decent size," which is the reason the experiment assumed that the trees should be at least five years old; the most complete and recent satellite information used for tree canopy cover assessment refers to 2018.

In this regard, to understand the methodology and the related results achieved, it should be noted that for the city of Turin, the neighborhoods of Madonna del Pilone (Figure 4.1, number 21) and Borgo Po and Cavoretto (Figure 4.1, number 22) will be excluded from the neighborhood focus. These are two hilly neighborhoods characterized by extensive forested areas and less by sealed and urbanized areas. In addition, the management of green spaces in these neighborhoods can be likened to forest/woodland management, thus strongly differing from the management of green areas in the remaining neighborhoods.

Three trees from every home

In order to verify the presence of three trees, visible or otherwise, around the buildings, the following procedure was carried out. Building footprints from OpenStreetMap were extracted, and the presence of three trees in a buffer of 30 m from each house in the City of Turin was subsequently verified. Trees present for a long time in the City of Turin were considered relevant, considering those that had spent > 5 years since planting (data based on [224]). It should be considered that in Turin, in the past years, urban forestation (with young plants 60 – 80 cm high) was rarely used, and the plants planted were at least two years old.

At the urban scale, a buffer area at the distance of 30 m from each building was identified to analyze the distribution of trees. The buffer function present in the geoprocessing tools of QGIS, which generates a polygon vector, was used. Then, using the Count points in the polygon function in Analysis Tools, a proper count was performed to identify the number of trees per building. Finally, using the Centroids function present in Vector Geometry, the centroids of the buildings were extracted, which were then thematized to make the data obtained by the count performed earlier more readily understandable. The Rule-based method was used for thematization, allowing each value class to be associated with a separate visualization.

The above centroids were interpolated for the neighborhood-scale analyses using the column with the tree count as the interpolation weight. The result of the interpolation is a grid-like raster surface. The interpolation algorithm used was Inverse Distance Weighting (IDW) .

The raster was then processed using zonal statistics to identify the

distribution of trees at the neighborhood scale using the neighborhoods downloaded from the OpenData portal of the Metropolitan City of Turin as a reference vector. The mean was employed as the analysis parameter of the zonal statistics. The data was thematized by dividing the output data into the previously exposed three intervals, by making the ratio between the values of polygons in which at least three trees fall and the number of buildings in the polygon.

30 % Tree Canopy Cover in every neighborhood

To calculate the Tree Canopy Cover in the city of Turin, the data present in the Pan-European High-Resolution Layers - Forests (2021) [225] were used, and in particular, the 2018 data concerning the Tree cover density in 10 m resolution (ranging from 0 to 100%). The primary sources are (since the 2018 reference year) Sentinel-2 and Sentinel-1. For the urban-scale analysis, values were themed into equal 10% intervals.

The zonal statistics algorithm was used for the neighborhood-scale analysis to calculate the average cover in each neighborhood. The thematization was performed by dividing the previously computed data into three quantile classes.

300 m from the nearest park or green space

The European Regional Office of the World Health Organization recommends a maximum distance of 300 m to the nearest green space (of at least 1ha). However, creating new public green spaces of 1 ha in size could be difficult. In these cases, a decent area of 0.5 ha should be a minimum [226]. In the case of Turin, it was considered appropriate to define the reference size of the green area to be reached in 300 m as 0.5 ha.

The data for parks and gardens was taken from OpenStreetMap. These data were reprocessed at both urban and neighborhood scales.

Specifically, at the urban scale, it was performed using the Proximity algorithm found in Raster Analysis.

The resulting raster was thematized into three classes to make it more intuitive. See the results section for a discussion of the classes.

Neighborhood-scale analysis was conducted by performing zonal statistics on the previously calculated neighborhood raster, using neighborhoods

as the reference vector.

Environmental data

An accurate assessment of the capacity of urban green spaces to reduce the heat island effect is crucial in planning decisions due to rising thermal pressures on both new and existing urban environments brought on by climate change. This frequently requires data that planners might not have. This dissertation used the InVEST Urban Cooling model, which solves this drawback by employing several readily available factors assigned to a land cover map to create a heat mitigation index (HMI) to estimate the vegetation's ability to cool an area.

The Urban Cooling model generates the heat mitigation index (HMI) based on evapotranspiration from vegetation, the cooling distance of significant urban parks, and albedo given to a land cover (LC) map to measure the average cooling capacity on air temperature.

The model initially calculates the city's evapotranspiration index (ETI) using a given LULC map:

$$ETI = \frac{K_c \times ET0}{ETmax} \quad (4.1)$$

It reflects the potential evapotranspiration from plants. Each pixel is determined by multiplying the reference evapotranspiration, $ET0$, the crop coefficient K_c related to the pixel's LULC type, and the highest value, $ETmax$, of the $ET0$ raster in the study area.

In particular, the evapotranspiration data provided for the computation was collected using MOD16 global evapotranspiration supplied by NASA. Using satellite remote sensing data, this data represents an evapotranspiration estimation from the earth's land surface.

Subsequently, the model computes each pixel's cooling capacity index (CC_i) based on local shade, albedo, and evapotranspiration.

$$CC_i = 0.6 \times shade + 0.2 \times albedo + 0.2 \times ETI \quad (4.2)$$

The percentage of tree canopy (> 2 m in height) connected to each land use/land cover (LULC) category is represented by the shade factor.

The model then determines the urban HM index: If there are no large

green spaces in the pixel's vicinity, HM is equal to CC; otherwise, it is set to a distance-weighted average of the CC values from the large green spaces and the pixel in question.

Socio-demographic data

Data collected by the WorldPop project (2021) [7] were analyzed to assess the socio-demographic characteristics of Turin neighborhoods, calculating the Population Density. This variable was calculated as the number of inhabitants on the neighborhood's surface [224]. To compare the neighborhoods, all the results were grouped into three categories: low, medium, and high values.

Prioritizing Neighborhoods for NbS Implementation

A *precedence index* (PI) is proposed to identify neighborhoods with a high need for NbS implementation.

$$\text{sum} = \sum_i w_i \times \text{variable}_i \quad (4.3)$$

$$\text{min}(\text{sum}) = \sum_i \text{min}(\text{variable}_i) \quad (4.4)$$

$$\text{max}(\text{sum}) = \sum_i \text{max}(\text{variable}_i) \quad (4.5)$$

$$PI = \frac{\text{sum} - \text{min}(\text{sum})}{\text{max}(\text{sum}) - \text{min}(\text{sum})} \quad (4.6)$$

The index was made by summing the variables reported in the mathematical formula, giving each variable a value of 0 (low), 1 (medium), or 2 (high). To apply it in different urban contexts, the index also provides a weight for each variable, i.e., w_i . In this experiment, all the identified factors have been equally weighted with a value of 1. Next, the index has been normalized, which consists of limiting the range of a set of values within a specific predefined range. After summing variable values (formula 4.3), a min-max normalization is performed to obtain a precedence index between 0 and 1. This procedure ensures that the values of the variables are comparable as they are reported on the same scale and, with the same

variables, it is possible to compare the values of precedence indices calculated on different neighborhoods.

By summing the variable values, the results were divided into three categories: $0 \leq x < 0.25$ (high intervention priority); $0.25 \leq x < 0.5$ (medium intervention priority); $0.5 \leq x < 1$ (low/no intervention priority).

This index provides an initial graphical representation of the neighborhoods that should be affected by interventions with NbS.

Neighborhoods falling in ranges characterized by numerical values close to zero have precedence for intervention with NbS. Subsequently, the other ranges have a decreasing need for intervention with NbS.

Should the number of variables increase, the methodology and approach would still apply in other contexts.

Results

The results of the variables analyzed at urban and neighborhood scales are presented below.

For the *3 trees from every home*, the results show the presence of the tree component within a buffer of 30 m from every building in Turin. The results were divided into three classes (low, medium, and high) to distinguish buildings better and, subsequently, neighborhoods that comply with the three trees from every home rule.

Figure 4.3 shows on the left the centroid for each building. The class 'low' is associated with buildings with several trees within 30 m radius < 3 ; the class 'medium' is related to buildings with number of trees equal to 3; the class 'high' is associated with buildings with number of trees > 3 .

Figure 4.3 shows on the right the results at the neighborhood level. The results were obtained by averaging the values of centroids falling within the individual neighborhood.

It is possible to highlight that about 40% of the neighborhoods can meet the rule of potentially having three trees visible from each building. In contrast, only one neighborhood has substantial deficiencies in this regard.

For the *30 percentage of Tree Canopy Cover in every neighborhood* rule, the results show the percentage of Tree Canopy Cover in Turin. To better distinguish neighborhoods that comply with the 30% of Tree Canopy Cover, the results were divided into three classes (low, medium, and high).

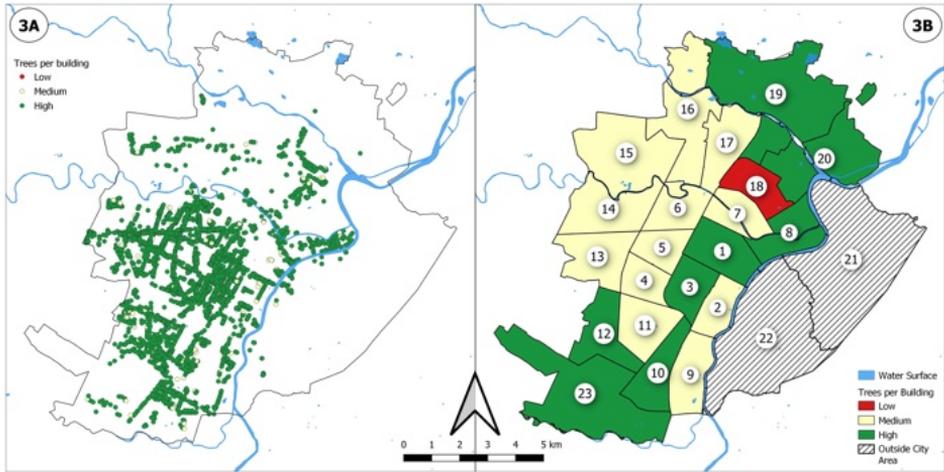


Figure 4.3: Differentiation of buildings according to compliance with the *3 trees from every home* rule (left); differentiation of neighborhoods according to compliance with the *3 trees from every home* rule (right).

Figure 4.4 shows on the left the values of tree canopy cover at the urban scale. The City of Turin has a Tree Canopy Cover percentage of 19.43%, also considering the hillside area. Focusing on the neighborhoods under investigation, the minimum value of percentage cover is 4.8%, while the maximum value is 14.7%. Given the results obtained, it was deemed inappropriate to graphically render a map showing a division of classes according to the benchmark of 30% tree canopy cover. However, it was necessary to investigate this further to understand whether any particular differences emerged between the various neighborhoods.

To differentiate the neighborhoods in terms of the % of Tree Canopy Cover, Figure 4.4 shows the classification of Tree Canopy Cover into the three classes (low, medium, and high) on the right. Precisely, the average value of tree canopy cover for each neighborhood was calculated, and the identification of the extremes for the three classes was done through a subdivision into three quantiles. The three classes emerge from a proposed subdivision into equal percentages. For the low class, the interval 0-5 % was used; for the medium class, the interval is 5 – 10%; for the high, the interval is 10 – 15%. Figure 4.4 presents on the right a division of the data

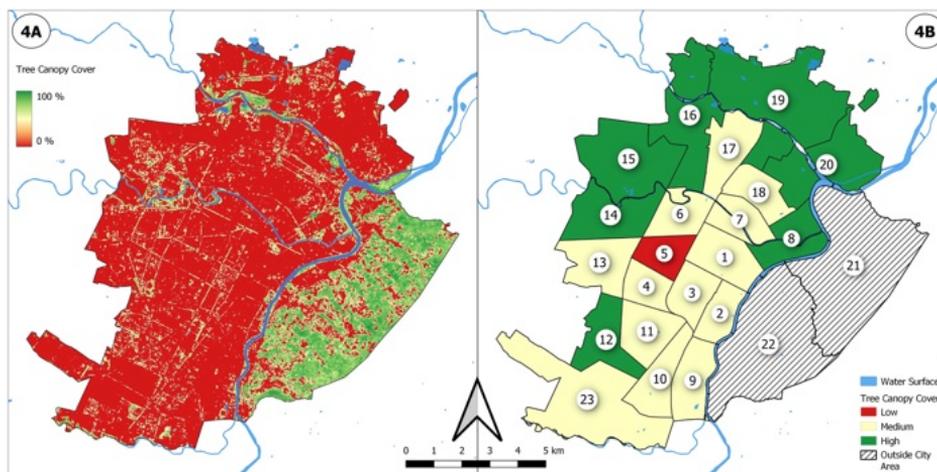


Figure 4.4: Analysis of the % Tree Canopy Cover of the city of Turin (left); the subdivision of neighborhoods according to the % Tree Canopy Cover (right)

into three classes, but the 'high' class refers to tree canopy cover values that are never higher than 14.7%.

The results underline that the City of Turin, to date, is far from being able to comply with the 30% Tree Canopy Cover rule in every neighborhood.

For the 300 m *from the nearest park or green space* rule, the results show the distance between buildings and green areas of at least 0.5 ha in Turin. To better distinguish neighborhoods that comply with the '300 m from the nearest park or green space' rule, the results were divided into three classes (low, medium, and high).

Figure 4.5 shows on the left the proximity map at the urban scale. A classification in low, medium, and high was applied, using the following intervals: low: (> 500); medium ($\geq 300; 500$); high ($0; 300$).

Following the same classification, Figure 4.5 shows on the right averaged data at the neighborhood level.

The results show that most neighborhoods in Turin have a green area of at least 0.5 ha within 300 m of buildings. Exceptions are the 'Crocetta' and 'Madonna di Campagna' districts, which have green areas within 500 m.

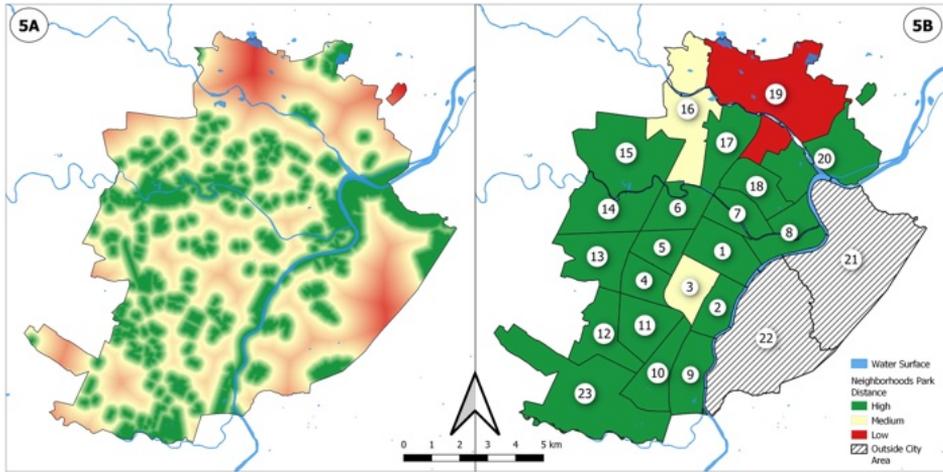


Figure 4.5: Proximity map concerning urban green areas at urban level (left); Subdivision of neighborhoods according to their potential to comply with the '300 m from the nearest park or green space' rule (right).

The district with the lowest values is 'Falchera.'

For the environmental variable, the results estimate the heat mitigation based on shade, evapotranspiration, albedo, and the distance from cooling islands (e.g., parks). The results were divided into three classes (low, medium, and high) to distinguish neighborhoods with different heat mitigation values better.

It can be highlighted from Figure 4.6 how about 80% of the area of Turin is in critical situations from an environmental point of view. This observation is represented by the red and yellow colors, which indicate low values of the Urban Cooling Index.

Only 5% of the area of the City of Turin has a high value of the Urban Cooling Index, found mainly in the hilly regions.

The Urban Cooling Index has its foundation in key scientific references. In particular, Zardo et al. (2017) identify scores/benchmarks useful for dividing areas with different index values. However, as can already be seen in Figure 4.6, the city of Turin has homogeneously low values. As in the case of the Tree Canopy Cover, the Urban Cooling Index subdivision has not been shown graphically because all the neighborhoods fall within

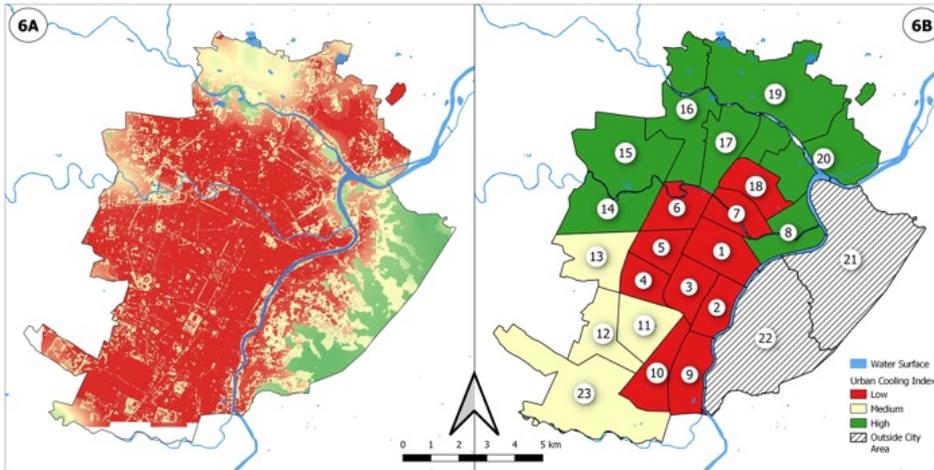


Figure 4.6: The results of the Urban Cooling Index at the urban level (left). The subdivision of neighborhoods which have different heat mitigation values (right).

the low class of values.

However, an attempt was made here to gain a deeper understanding of any differences between neighborhoods. Figure 4.6 highlights the Urban Cooling Index values referable to the City's neighborhoods. Specifically, the following division of values into quantile classes was used: low (0–0, 1); medium (0, 1 – 0, 15); high (0, 15 – 0, 3).

The results show the population density density of the Turin neighborhoods (Table 4.2). To better distinguish neighborhoods with different Population Densities, the results were divided into three classes (low, medium, and high), as reported in Figure 4.7.

Figure 4.7 shows the subdivision of neighborhoods according to the Population Density value. The values are in the range 913 – 20,000 [224] and have been divided into three classes: low (≥ 913 ; 5000); medium (> 5000 ; 10,000); and high ($> 10,000$ – 20,000).

It can be observed from the results that about 60% of the neighborhoods have a high Population Density value, mainly concentrated in the central area of the city. Interestingly, the three districts in red, which in the past were the industrial neighborhoods (numerous manufacturing es-

Table 4.2: Detail of neighborhoods and related Population Density.

Neighborhoods	Population Density (in/km ²)
1. Centro	10,517
2. San Salvario	15,371
3. Crocetta	12,335
4. San Paolo	15,759
5. Cenisia	16,764
6. San Donato	16,223
7. Aurora	14,733
8. Vanchiglia	9041
9. Nizza Millefonti	8043
10. Mercati Generali	14,101
11. Santa Rita	15,879
12. Mirafiori Nord	11,532
13. Pozzo Strada	13,526
14. Parella	9616
15. Le Vallette	5489
16. Madonna di Campagna	5543
17. Borgata Vittoria	10,309
18. Barriera di Milano	16,984
19. Falchera	2075
20. Regio Parco	4068
21. Madonna del Pilone	913
22. Borgo Po e Cavoretto	1435
23. Mirafiori Sud	3089

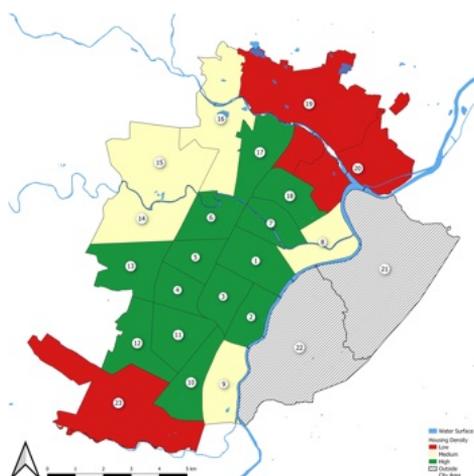


Figure 4.7: Subdivision of neighborhoods according to Population Density values.

tablishments, including the former FIAT plants) or neighborhoods where blue-collar workers lived, now have a lower population density. This could be due to less industrial/economic activity in these areas.

Finally, following the application of the proposed index, the neighborhoods in the City of Turin that need precedence of interventions with NbS are highlighted. Specifically, the results are returned in two different detail levels.

Figure 4.8 on the left shows the neighborhoods in Turin with different priorities for intervention with NbS. The results correspond to a division of values into three categories (low, medium, and high). However, the figure on the left shows the first graphical representation of the neighborhoods that should be affected by interventions with NbS.

The results, therefore, show the need for numerous NbS interventions in the city. To suggest to politicians the neighborhoods where they should focus their attention and funds promptly, Figure 4.8 on the right shows the three neighborhoods, Cenisia, Nizza Millefonti, Barriera di Milano, that have a precedence for NbS implementation.

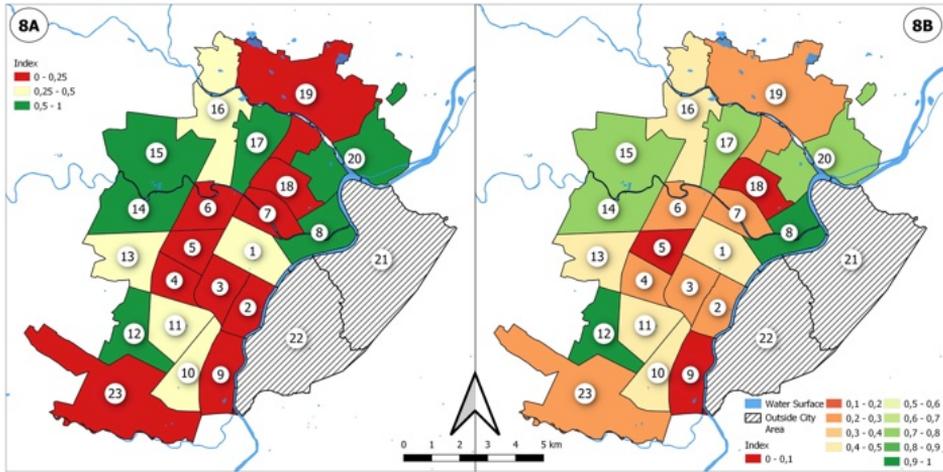


Figure 4.8: Identification of neighborhoods needing NbS implementation precedence (left). Detail on the neighborhoods to identify those in which the municipality should start investing immediately by implementing NbS (right)

4.2.3 The case study of Asti

Materials and Methods

The city of Asti (44.9°N, 8.206944°E) is located in the Piedmont region of northwestern Italy. It covers a total area of 151.31 km² with a population of 73,495 inhabitants, giving a population density of 485.72 people/km² as of 31 October 2022 [227]. The city has an average green area per inhabitant of 12.64 m², which is higher than the minimum value of 9.0 m² per person for public areas provided for greenery, play, and sport, as laid down in Italian Interministerial Decree no. 1444/1968, point 3.C. Specifically, it counts on an urban tree stock of about 12,500 specimens, of which 4500 are in tree-lined avenues and about 14,000 m of hedges, as of 2016 [228]. The most common species are linden, sycamore, hornbeam, ash, ginkgo biloba, and mulberry, including two metasequoias.

As of January 2023, the city of Asti does not seem to be divided into districts but rather into constituencies and 78 electoral sections (or wards). Therefore, to carry out appropriate analyses of the urban fabric of the city,

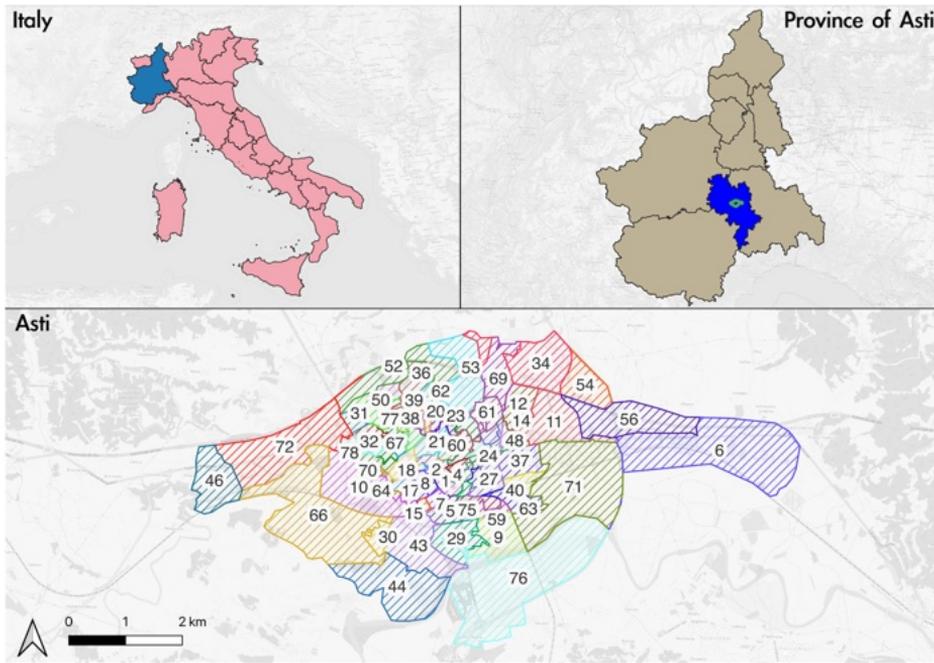


Figure 4.9: Location of the city of Asti and its division into voting sections or wards. (Sources: Google Maps and OpenStreetMap for the maps; SIT Municipality of Asti for the voting sections. Available online at [6].

it was decided to use the electoral sections (or wards) and their respective areas (Figure 4.9). The wards were discarded because in-depth studies cannot be carried out in the most urbanized part of the city.

This division was the most efficient way to analyze the city in relatively small areas and accurately identify risk factors and pressures. Electoral wards are ideal for small to medium-sized cities, including those with a historic urban layout, with no official administrative divisions into neighborhoods or districts.

As for demographics, as of 1 January 2022, the average age of citizens of Asti is 46.2 years, compared to 43.8 years in 2012 [229]. As Asti is indexed by ISTAT as the 75th city in the country in terms of population residing in the central city [230], this method has the potential to be replicated in another 7826 Italian municipalities, out of a total of 7901, as per data

updated on 1 January 2023 (ibid.).

The analysis follows an approach similar to the one described before. The study proposes a review of the application of the 3-30-300 rule to compact small and medium-sized towns and cities with a solid historical background that features Roman planning and has a medieval fabric.

In contrast to the proposed method for the Turin case study, differences in the case of the city of Asti pertain to the omission of environmental parameters, specifically the Urban Cooling Index analysis. Instead, to focus on Urban forestry, this parameter has been substituted with a variable that assesses the quality of green spaces within the city, computed through Earth Observation data.

The *Normalised Vegetation Vigour Index* (nVVI) [231] was computed and analyzed to assess the quality of the green spaces in Asti. The layer was generated utilizing a time series of Normalised Difference Vegetation Indexes (NDVIs), calculated from Sentinel-2 satellite images to assess the greenery vigor. This was achieved by analyzing the phenology of the area and its relation with trees' growth and health. For every date during the cloud-free periods in 2018, NDVI maps were created using the Sentinel-2 data. The NDVI profiles were regularised, applying splines with a time step of five days. Subsequently, they were filtered by FFT to minimize local fluctuations and emphasize the periodicity of the phenological trend. The Start of Season (SOS) and End of Season (EOS) were calculated using a 2nd derivative approach at the pixel level. The annual NDVI integral was then computed and divided by the length of the growing season to obtain an integrated NDVI. An integrated NDVI map was generated for 2022 with a Ground Sample Distance (GSD) of 10 m. The average integrated NDVI value for the given tree areas was computed using zonal statistics. The integrated NDVI was normalized to a range $[0, 1]$. It was assumed to be a predictor of tree decline, with a value of 0.5 used as the threshold for distinguishing between potentially declining trees and healthy ones. Finally, the $[0, 1]$ range was divided into five equal intervals.

Analyses were conducted at both the city level and the level of electoral sections (wards). At the qualitative level, findings were classified into five categories: very low, low, medium, high, and very high. This subdivision facilitated a simple and intuitive data comprehension, providing a clear framework for comparing and categorizing different variables.

One of the key objectives of this approach was to share the final analysis with the public administration and decision-makers, such as urban planners or policymakers, and assist them in making data-driven decisions. This was achieved by using descriptive terms such as 'very low,' 'low,' 'medium,' 'high,' and 'very high,' which make it easier to communicate the results to a broader audience. These featured categories provide a common language that is easily understandable and relatable to the audience, even if they are unfamiliar with the technical specifics of the analysis. This method highlights how the 3-30-300 rule serves as a helpful guideline for policy and planning purposes and stimulates debate about urban greenery and UF. Furthermore, this rule can be applied in various urban cities to analyze the status quo and develop future strategies for increasing or maintaining UF. It is important to note that the primary focus of this approach is not to compare results between different cities but to analyze differences within the same one to reduce possible social inequalities [217].

The area analyzed does not include the agricultural, wooded, and hilly regions surrounding the built-up zone, as their management is often entrusted to private owners, and the remaining public spaces are managed differently in urban areas. This decision was made after the author conducted three visits to the city in October 2022 and February and June 2023.

Results

The findings display the occurrence of tree species within a 30 m radius of each building in the urban area of Asti (Figure 4.10). The findings have been categorised into five classes (from 'Very Low' to 'Very High') at the electoral section level.

The left-hand map identifies the buildings with colours indicating the classification according to the number of trees within a 30-m buffer. Precisely, the 'High' category represents buildings surrounded by over three trees, the 'Medium' category corresponds exactly to three trees, and the 'Low' class stands for containing fewer than three trees.

On the other hand, the map on the right illustrates the analysis conducted within the voting sections (or wards), using the aforementioned intervals. This classification considers the number of buildings, electoral wards, and their respective average scores.

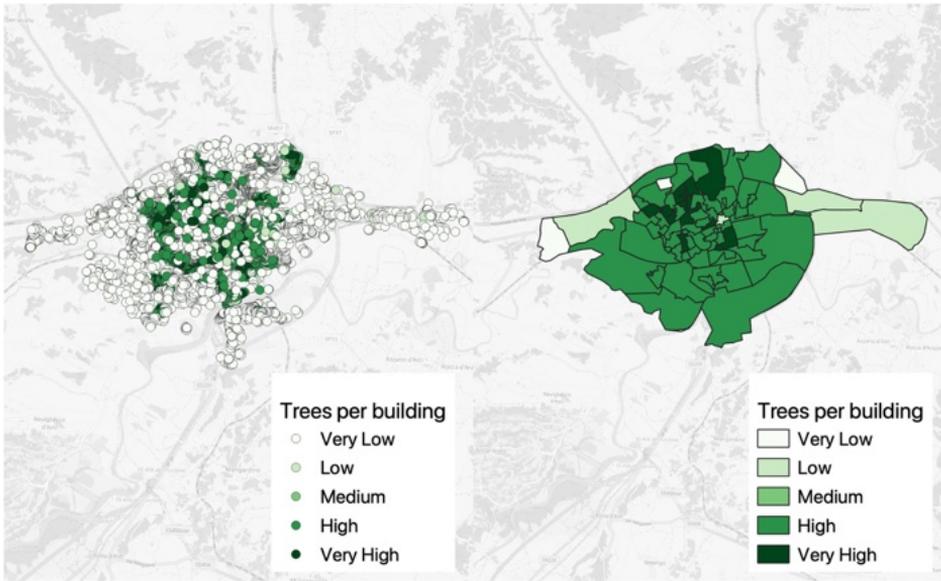


Figure 4.10: The presence of a minimum of three trees within a 30 – m radius of buildings. The left image identifies the buildings, while the right image classifies the electoral sections based on the number of trees surrounding each building. The score inside the voting district is computed by averaging the scores of each building.

Results indicate that 64 out of 70 electoral sections (or wards) had an average of at least three trees per building, demonstrating a high tree density throughout the surveyed areas.

Figure 4.11 shows the percentage of tree canopy cover (TCC) in the city of Asti. The results have been divided into five classes, ranging from 'Very low' ($\leq 10\%$) to 'Very high' ($> 40\%$), based on the electoral wards.

Figure 4.11's left-hand side displays the TCC values for Asti. In contrast, Figure 3's right-hand side analyses the spatial unity of electoral wards and displays the TCC value division into five classes, which comes from a subdivision into equal quantiles. Electoral wards' values range from roughly 0.6% to about 55%.

The results concerning the distance of buildings from green areas in Asti are presented in Figure 4.12, considering a radius of 300 m and a

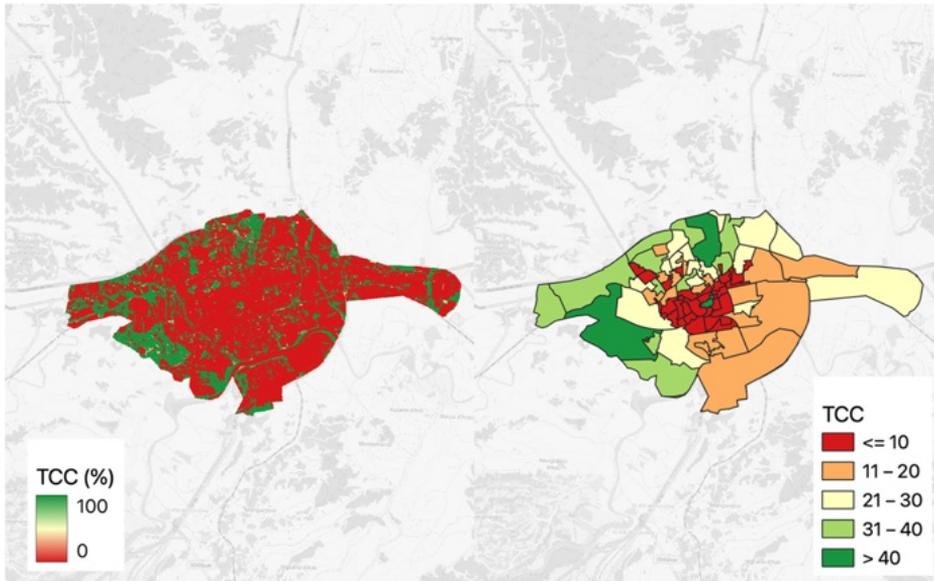


Figure 4.11: A percentage analysis of tree canopy cover (TCC) at the city and electoral section level.

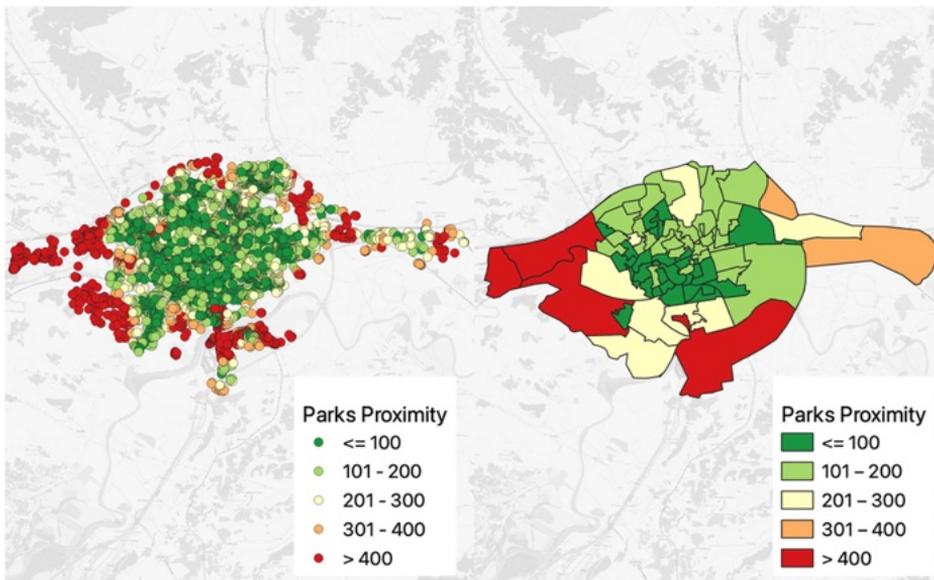


Figure 4.12: Distance of buildings from green spaces and parks of 0.5 hectares or more.

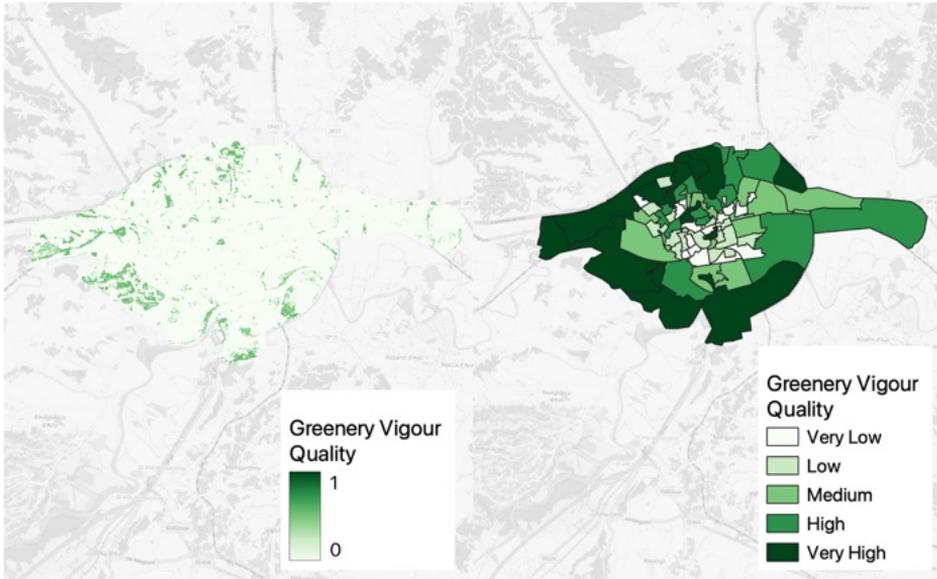


Figure 4.13: Analysis of the greenery vigour quality at the city and electoral section.

minimum size of green area of 0.5 ha. The analysis of this parameter was conducted at both the city level and the electoral section (or ward) level. Specifically, the analysis utilized the map of publicly owned green areas provided by the Territorial Information System (SIT) of the Municipality of Asti.

Figure 4.12 shows the proximity map results at the city level on its left-hand side. The division into five classes adheres to the following intervals: 'Very Low' (greater than or equal to 400 m); 'Low' (from 300 to 400 m); 'Medium' (between 200-300 m); 'High' (from 100 to 200 m); 'Very High' (less than or equal to 100 m). The classification for the electoral sections (or wards) was the same.

The results indicate that 64 out of 70 electoral wards satisfy the requirement of having at least 0.5 ha of green areas located 300 m away from buildings. However, 56 electoral wards fail to meet the 3-30-300 rule parameter. The primary reason for non-compliance is the requirement of a minimum 30% tree canopy cover. The size of the areas located far from

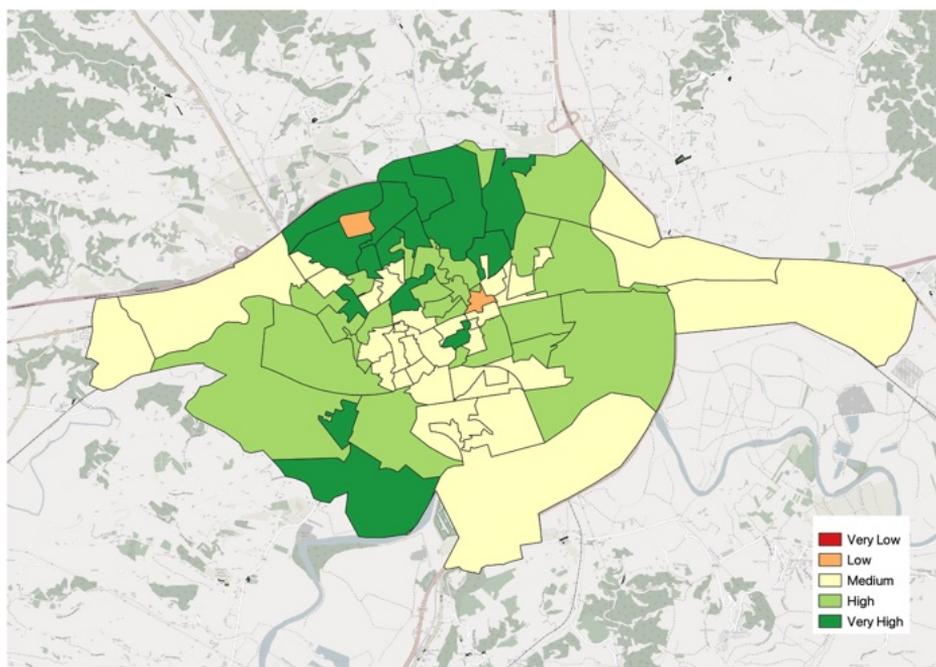


Figure 4.14: Classification of areas according to the potential need for implementing nature-based solutions (NbS). Classes with lower values indicate a greater need for intervention.

the city centre is a significant factor contributing to the non-compliance with the 30% tree canopy cover rule. The aerial sizes of these regions are larger than those of the central areas, which makes it more challenging to achieve the desired canopy cover. Furthermore, some wards may have an adequate number of trees, but often they do not meet the required canopy-cover value due to their youth and size. Juvenile trees have less canopy cover and require time to grow, thus diminishing their overall contribution to the total canopy cover. Moreover, smaller or elongated trees cover less surface area, further affecting the achievement of the target canopy-cover percentage.

Figure 4.13 displays the results regarding the quality of green spaces. The results at the urban level were categorised into five classes, ranging from 'Very low' to 'Very high'.

The results indicate a significant variation in green quality among different areas: more central areas have a lower to medium-low quality of green spaces.

The findings suggest that densely populated urban areas may significantly affect the availability and quality of green spaces. On the other hand, the surrounding areas, which may have experienced less intensive urbanization, show a gradual increase in green quality along the surrounding belt of the densely populated city center. The observation suggests that these outer regions provide more favorable conditions for the development and maintenance of green spaces, which leads to better green quality. Overall, the results emphasize the importance of urban planning strategies and spatial factors to enhance green quality and promote a healthier urban environment in the city center.

Potential areas for intervention can be identified in Figure 4.14. This Figure serves as a tool to initiate discussions. Based on the cross-analysis of the data and field visits, certain areas in the south/southeast of the city of Asti have the potential to host nature-based solutions (NbS). These areas have good connectivity to the existing ecological corridors associated with the Tanaro River, the WWF La Bula Oasis and the Belangero Ponds (44.878822°N, 8.209906°E), as well as to the western region where the Special Nature Reserve of Valle Andona, Valle Botto and Valle Grande (44.9409°N, 8.1151°E) is situated.

4.3 Scenario Thinking through GAN-Based Synthetic Satellite Imagery

4.3.1 Introduction

In the scientific discipline, scenario thinking is a comprehensive method that goes beyond mere speculation about the future. It systematically constructs a set of coherent and plausible future conditions under different assumptions and interpretations. The scenario thinking process typically involves several distinct phases:

- *Identifying the Scope*: This initial phase defines the specific focus and boundaries of the scenario exploration, such as particular aspects of urban development or environmental changes.
-

- *Gathering Information*: Involves collecting relevant data, including historical trends, current conditions, and expert opinions, to inform the scenario development.
- *Identifying Key Factors*: This step entails pinpointing the critical drivers and uncertainties that will influence future outcomes.
- *Developing Scenarios*: Multiple plausible scenarios are created based on the gathered information and critical factors. Each scenario represents a different potential future based on varying combinations of the identified drivers and uncertainties.
- *Analysis and Implication*: Each scenario is analyzed to understand its implications, challenges, and opportunities.
- *Strategy Development*: Using insights from the scenario analysis, strategies and plans are developed to navigate or leverage the potential futures.
- *Monitoring and Revision*: Continuously monitor real-world developments and revise the scenarios and strategies as necessary to stay relevant and practical.

This approach becomes particularly valuable in the context of Earth Observation for urban planning. Here, scenario thinking is utilized to develop a range of potential future urban developments. Planners can create detailed scenarios by integrating Earth Observation data, such as satellite imagery, climate models, and demographic trends. These scenarios might encompass varying degrees of urban expansion, green space allocation, infrastructure development, and their potential impacts on environmental sustainability and urban livability.

In applying scenario thinking to Earth Observation, urban planners can evaluate the implications of various developmental pathways. For instance, scenarios can be developed to show urban greenery's impact on microclimates or the effect of new infrastructure on traffic patterns and air quality. By comparing these scenarios, planners can identify strategies that align with sustainability goals, community needs, and resilience against environmental changes. This methodological approach is crucial

in rapid urbanization and climate change, offering a proactive tool for designing cities that can adapt and thrive in the face of future challenges and uncertainties.

This section introduces a novel urban planning and environmental simulation approach by leveraging Generative Adversarial Networks (GANs) to generate multispectral synthetic satellite imagery, particularly tailored for nature-based solutions.

Generative Adversarial Networks (GANs) have made significant strides in recent years, showing impressive capabilities in a range of applications [232]. This technique involves a pair of neural networks competing against each other: a generator that creates data aiming to mimic reality as closely as possible, and a discriminator that evaluates whether the data generated is fake or real. Conditional GANs (cGANs) incorporate additional conditions during the generation process. This enhanced model has been effectively applied in various tasks, including image colorization [233, 234] and style transfer [235, 236].

Several recent works have studied the generation of synthetic satellite imagery through GAN networks [237, 238, 239]. Satellite imagery possesses distinctive characteristics that set it apart from conventional images typically employed in computer vision. In particular, they have multispectral content that goes beyond the RGB bands. Consequently, standard GAN architectures cannot be readily employed for processing such imagery [237]. For this reason, significant research is dedicated to identifying methods to generate multispectral satellite images with consistent multispectral content [237, 238].

In the realm of this thesis, the primary goal was to develop a method capable of generating synthetic Sentinel-2 satellite imagery with a keen focus on maintaining the multispectral content's coherence in comparison to real data. Section 4.3.2 illustrates the proposed approach, which unfolds in two distinct phases, addressing the unique challenge posed by the multispectral nature of satellite imagery that distinguishes it from the standard RGB images prevalent in conventional computer vision tasks. Section 4.3.3 presents the results.

The proposed approach aims not only to contribute to the body of knowledge on synthetic imagery generation but to bolster the capabilities of urban simulation tools within the realms of sustainable urban develop-

ment and climate resilience. By providing a novel tool capable of simulating hypothetical urban scenarios, the goal is to offer urban planners, environmental scientists, and policy-makers an advanced instrument for envisioning and evaluating potential future developments and interventions. This forward-looking tool is designed to support decision-making processes by enabling the exploration of various urban growth patterns, land-use changes, and environmental mitigation strategies. By generating realistic synthetic satellite imagery, the assessment of urban development's impact on climate resilience and sustainability can be conducted with unprecedented accuracy and detail, thereby contributing to creating more sustainable and resilient urban landscapes.

4.3.2 Materials and Methods

The methodology extends upon existing frameworks by integrating advanced GAN architectures to generate high-fidelity, multispectral imagery that mimics the characteristics of the Sentinel-2 satellite constellation. Figure 4.15 illustrates the proposed approach. The solution generates Synthetic Sentinel-2 Multispectral images from Simulation Scenarios with a two-step process:

1. *Simulation Scenario to RGB:*

- Starts with a land cover mask representing a simulation scenario.
- This mask undergoes processing by a label-conditional Generative Adversarial Network (GAN).
- The GAN generates a realistic RGB image that reflects the textures corresponding to the different land cover types indicated by the mask.

2. *RGB to Multispectral:*

- The generated RGB image serves as input to the second phase.
 - A U-Net-like architecture processes it.
 - The U-Net model is designed to convert the RGB input into a Synthetic Sentinel-2 Multispectral image, predicting the values across the various spectral bands that Sentinel-2 would capture.
-

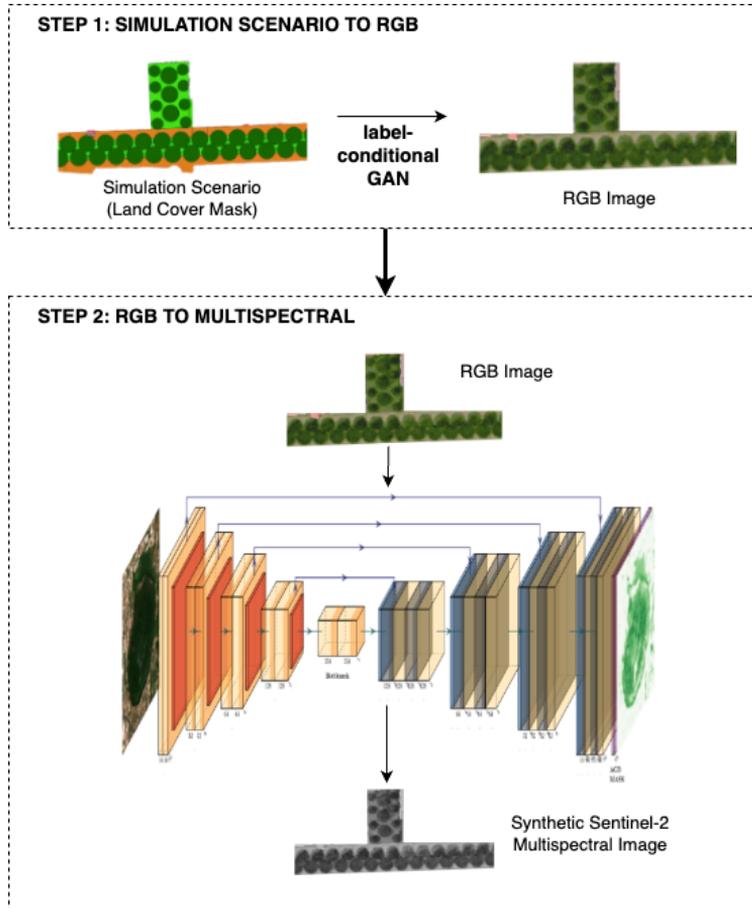


Figure 4.15: Proposed approach for the generation of Synthetic Sentinel-2 Multispectral image from Simulation Scenarios.

The adopted approach in the first step is based on the solution proposed in [240]. The model is a conditional GAN with spatially adaptive denormalization (SPADE) [239]. It consists of an encoder, a generator, and a discriminator. The encoder takes an input image and produces a mean and variance that is passed as an input to the generator. The generator uses a random vector z from the reparameterized output of the encoder and the class mask to generate synthetic images through a series of SPADE ResNet blocks.

Step 2 uses a Regressive U-Net, as described in Section 2.2.3. In particular, the network architecture was adapted to produce a feature mask of 7 channels as output. The seven channels correspond to the multispectral bands of Sentinel-2, excluding RGB bands and bands at 60 m spatial resolution that were ignored in this thesis.

This approach exploits the well-known inter-band dependencies that the network can learn [241]. It's efficient regarding computational resources and allows for shared feature learning, which could improve performance if the bands have significant correlations.

4.3.3 Results

For the experiment, step 1 was implemented using the already trained network made available from the authors in [240]. A dataset containing 2000 64x64 patches of Sentinel-2 satellite images was created for the second step. Patches were sampled across different areas of Europe, focusing on urban and peri-urban areas. The dataset was split into training, validation, and test sets with a 70-20-10 ratio. In order to create the dataset for step 2, step 1 was executed on all the patches, using a land cover mask derived from ESA WorldCover 2021 [242]. Thus, the output patches at step 1 for step 2 training were created.

Table 4.3 shows the results. The proposed solution was compared with machine-learning algorithms using RGB bands as features. The results presented in the table represent the average of the metrics obtained for the seven predicted bands. The comparison of the proposed solution with traditional machine-learning algorithms using RGB bands as features further supports the approach's effectiveness, with the Regressive U-Net demonstrating significantly lower Mean Absolute Error and Root Mean Square Error values. These preliminary results aim to show the feasibility of the

Table 4.3: Results for step two of the proposed approach: generate multispectral content from RGB bands.

Model	MAE	RMSE
Regressive U-Net	0.16	0.43
Linear Regression	0.37	0.91
Random Forest	0.34	0.87

proposed approach. Further experiments will need to be conducted, particularly employing Generative Adversarial Networks (GANs) for Step 2 as well.

Figure 4.16 shows the results of the proposed approach. The left column displays the input scenarios, which are masks of land cover layouts. These layouts indicate where specific land covers, such as vegetation or water bodies, are intended to be placed in an urban setting. The center column depicts the synthetic RGB images generated by a conditional Generative Adversarial Network (GAN) based on the input scenarios. The GAN has transformed the land cover masks into realistic-looking RGB images that mimic the appearance of natural features within an urban landscape. The right column shows the Normalized Difference Vegetation Index (NDVI) computed on the synthetic multispectral Sentinel-2 images produced from Step 2. The NDVI is a derived index that indicates the presence and condition of vegetation, with green tones representing higher vegetation density and orange to red tones indicating lower vegetation density or urbanized areas.

It is noteworthy to observe the consistency of the NDVI index produced in relation to the input scenario, which evidences that the generated synthetic image possesses coherent multispectral content. This coherence underscores the effectiveness of the synthetic generation process in accurately simulating the spectral characteristics associated with the various land covers depicted in the input scenario.

Figure 4.17 presents a what-if analysis for assessing the impact of various Nature-based Solution (NbS) installations on Land Surface Temperature (LST).

On the image's right side, each scenario's LST impact is visualized through thermal maps, with varying color intensities representing differ-

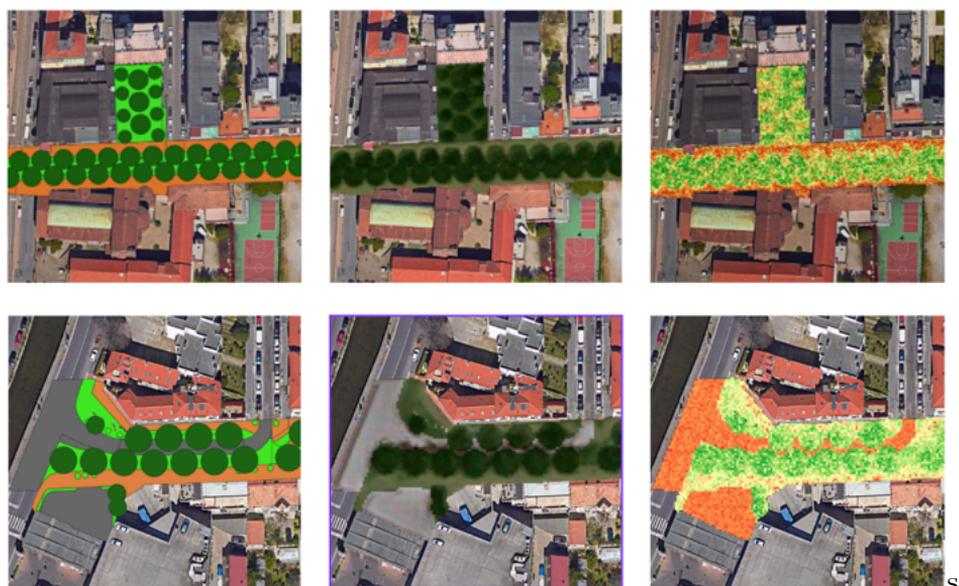


Figure 4.16: Synthetic Multispectral Sentinel-2 data. (left) Input scenario, (center) Step 1 output: synthetic RGB image with conditional GAN, (right) NDVI computed on the Synthetic Multispectral Sentinel-2 image generated from the step 2.

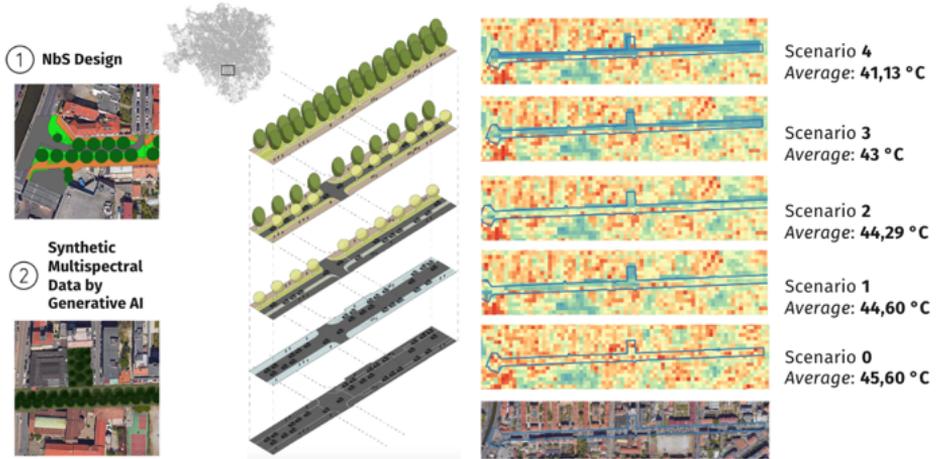


Figure 4.17: What-if analysis for four scenarios to study the impact of different future Nature-based solution installations in terms of Land Surface Temperature

ent temperatures. The scenarios are ranked from 0 to 4, with Scenario 0 depicting the current state and Scenarios 1 to 4 representing progressive NbS installations. Each scenario is accompanied by an average temperature value, allowing for a comparative analysis of the potential cooling effects offered by the NbS installations. Land Surface Temperature was computed using the approach described in 3.3.1. The Landsat-8 scene used for downscaling was the same for all the scenarios. To change was only the corresponding synthetic Multispectral Sentinel-2 of the scenario.

4.4 Conclusion

This chapter presents a comprehensive approach to integrating Nature-based Solutions (NbS) within urban planning frameworks to address the twin challenges of urban expansion and environmental sustainability.

Since the steps of planning NbS should follow five key guiding principles [243], this research focused on Place specificity, analyzing the green characteristics, together with the environmental and social characteristics of the city. The planning and management of green spaces in urban areas,

focusing on urban forests (UF), is gaining interest globally, especially in the USA and Canada [244]. Therefore, promoting NbS requires a preliminary investigation of the area available for the specific NbS to be implemented and scaled up [244].

This approach mainly aims to identify areas where UF can be implemented as Nature-based Solutions (NbS). It proposed a methodology for implementing NbS in urban settings, especially for climate-change adaptation, which also considers social aspects. To respond to these challenges, the thesis proposes implementing the 3-30-300 rule. The intent of the 'rule' seems more to serve as a guideline that prompts discussion and can help cities set targets. It, therefore, provides a discussion tool to understand where green spaces are adequately provided and where there are gaps. Thus, the rule is more of a tool to discuss the importance of trees with the ambition to be reported, for example, in a policy brief. It is not a benchmarking tool per se and certainly not between cities, but a touch more within cities to reduce social inequalities between neighborhoods.

The proposed revision of the rule is intended to trigger a debate on including the concept of UF as NbS [203]. Thus, the research proposes a PI to identify city neighborhoods that need NbS implementation, or more precisely, a method of suggesting to politicians where to direct funds for implementing new UF as NbS in the city. Therefore, the section discusses the data based on the research carried out and points out some synergies/differences with different methods [204] used in implementing and upscaling NbS in urban areas.

Focusing on the Turin and Asti case studies, the approach allowed the identification of those neighborhoods to upscale NbS, demonstrating its effectiveness in identifying neighborhoods where NbS could significantly enhance urban livability and environmental quality.

However, it should be noted that typology and data quality differences can lead to results variations [245]. Precisely for this reason, since a single approach for all cities under analysis does not work [246], the proposed PI allows the inclusion of several aspects (at least social and environmental), depending on the city under consideration. This aspect makes it possible to replicate the approach almost everywhere, finding an essential basis in the rule proposed by [217] and promoted by the IUCN. However, it is important to emphasize that the possibility of applying the PI in different

places does not mean that it is correct to compare the results of very different cities with different climates. Instead, it serves as a guideline to direct policy choices concerning the planning of UF as NbS within cities.

Focusing on the NbS of urban forestry, the research addressed the question: Where can NbS be implemented within the city to maximize their social impact? This regards explicitly the upscaling of those punctual NbS implemented and tested by European projects to amplify their impact. In particular, neighborhoods that should have precedence in NbS implementation were highlighted. These findings could help local governments implement the City's Green Plan with a helpful methodology to achieve environmental objectives.

On the other hand, the research proposed an approach to address the question: Can the impacts of urban development strategies be predicted before their execution? Starting from urban forests as pivotal components of NbS, the study leverages Generative Adversarial Networks (GANs) to produce synthetic multispectral imagery. This methodological advancement not only aids in visualizing the potential outcomes of implementing NbS within urban environments but also enhances the decision-making process by providing a scientifically grounded basis for predicting the environmental impacts of urban interventions before their actual realization.

The importance of strategically planning and managing urban green spaces cannot be overstated in the context of increasingly complex urbanization challenges and climate change imperatives. The novel approach discussed in this chapter underscores the synergy between technological innovation and urban ecological planning. By simulating various scenarios of NbS implementation and their consequent impacts, the study offers valuable insights into how urban areas can evolve toward sustainability and resilience.

The results of the proposed approach demonstrate the effectiveness of the synthetic generation process in accurately simulating the spectral characteristics associated with various land covers. Using the Generative Adversarial Network and the Regressive U-Net, realistic RGB images have been successfully generated and transformed into Synthetic Sentinel-2 Multispectral images with high spatial and spectral fidelity.

The what-if analysis presented for assessing the impact of Nature-based Solution installations on Land Surface Temperature showcases the practi-

cal utility of the generated synthetic multispectral data. Visualizing and comparing the LST impacts of different future scenarios provides valuable insights for urban planning and environmental management.

In conclusion, the approach contributes significantly to the discourse on sustainable urban development planning, providing a robust framework for employing NbS as a strategic tool for enhancing urban ecosystems. Integrating GAN-based synthetic satellite imagery into urban planning processes marks a step forward in utilizing advanced technologies for environmental sustainability. This approach supports the informed planning and implementation of urban green spaces and facilitates the proactive management of urban environments in the face of future challenges, aligning with global sustainability goals and fostering the well-being of urban populations.

Chapter 5

Conclusions

This thesis synthesizes a comprehensive approach towards Earth observation that is underpinned by the tenets of sustainability, designed to serve a wide array of user needs across temporal dimensions.

Embarking on a journey from past to future, the thesis posits that a user-centric model is imperative to harness the full potential of Earth observation data. Through its three pillars—Learning from the Past, Monitoring the Present, and Building the Future—the research encapsulates a complete cycle of understanding our planet’s dynamics.

The pillar "Learning from the Past" serves as a cornerstone for understanding how historical data, when analyzed through the lens of modern Artificial Intelligence (AI) and Earth Observation (EO) technologies, can inform current and future sustainability efforts. This phase of the research delves into the analysis of historical satellite imagery and environmental datasets, utilizing advanced AI algorithms to unearth patterns and trends in land use, vegetation cover, and biomass changes over time. By systematically studying the transformation of landscapes, urban sprawl, and the evolution of natural resources, significant insights are garnered into the impact of human activities on the environment.

One of the critical achievements under this pillar has been the development of methodologies capable of estimating aboveground biomass and understanding land consumption patterns. These methodologies leverage the temporal depth of EO data, providing a unique perspective on the trajectory of environmental degradation and recovery. Such insights are

invaluable for assessing the effectiveness of past conservation strategies, understanding the dynamics of ecosystem services, and informing policy and planning decisions.

Furthermore, "Learning from the Past" underscores the importance of historical context in sustainability science. By establishing a baseline of environmental conditions, researchers can better gauge the magnitude of changes and identify underlying causes. This historical perspective is crucial for setting realistic and achievable sustainability goals, as it considers the natural variability and resilience of ecosystems.

In essence, this pillar not only bridges the past with the present but also lays a solid foundation for future interventions. It exemplifies how the integration of AI and EO can transform vast and complex datasets into actionable knowledge, guiding efforts towards more sustainable and resilient environmental management practices.

The "Monitoring the Present" pillar emphasizes the critical role of real-time or near-real-time data analysis in understanding and managing current environmental and urban challenges. The application of Artificial Intelligence (AI) to Earth Observation (EO) data has revolutionized the capacity to observe, quantify, and respond to ongoing changes in the environment, urban settings, and climate patterns.

A key contribution under this pillar is the development of the GeoAI Processing Block, a sophisticated tool that enhances the processing capabilities for EO data. This innovation allows for more accurate and timely assessments of various phenomena, including urban heat islands, deforestation rates, water quality, and air pollution levels. Such capabilities are instrumental in deploying immediate responses to environmental crises, optimizing resource management, and improving urban planning strategies to enhance sustainability and resilience.

Furthermore, this pillar showcases the power of combining AI with EO for monitoring biodiversity, ecosystem health, and the impacts of climate change. By providing detailed, up-to-date information on environmental conditions, stakeholders are better equipped to make informed decisions, implement corrective measures, and track the effectiveness of sustainability interventions over time.

"Monitoring the Present" also highlights the importance of accessibility and integration of EO data into decision-making processes. The democ-

ratization of EO data, facilitated by AI-driven tools, enables a broader range of stakeholders, including governments, NGOs, researchers, and the public, to engage with and contribute to sustainability efforts.

In summary, this pillar underscores the importance of harnessing cutting-edge technologies to maintain a vigilant watch over the planet's health. It demonstrates how AI-enhanced EO can serve as a dynamic and powerful tool for real-time environmental stewardship, empowering societies to navigate the complexities of present-day sustainability challenges.

Lastly, "Building the Future" presents two strategic approaches that align with the predictive needs of users. By leveraging past data for informed action planning and employing scenario analysis, the thesis equips users with the tools not only to envision but also to strategically influence the future. It lays out a path for users to predict the outcomes of potential actions, enabling them to make decisions that optimally impact the future.

A notable aspect of this pillar is its emphasis on Nature-based Solutions (NbS) for urban planning. Through detailed case studies, including the cities of Turin and Asti, the research demonstrates how AI and EO can be instrumental in identifying potential areas for implementing NbS, such as urban forestry, green roofs, and water management systems. These solutions are recognized for their ability to enhance urban biodiversity, improve air and water quality, reduce urban heat islands, and contribute to the well-being of urban populations.

Moreover, the "Building the Future" pillar introduces innovative methodologies for scenario planning that incorporate the multispectral capabilities of EO data and the analytical prowess of AI. These methodologies allow urban planners, policymakers, and communities to simulate the impacts of different development pathways on urban greenery, biodiversity, and climate resilience. Such scenario planning is invaluable for assessing the trade-offs and synergies of various urban development strategies, ensuring that future growth aligns with sustainability goals and climate adaptation needs.

This pillar also underscores the importance of integrating AI and EO into the participatory planning processes, ensuring that the development of future strategies is inclusive, data-informed, and grounded in a deep understanding of ecological and social systems.

In conclusion, "Building the Future" represents a proactive approach to

sustainability, where AI and EO technologies serve as enablers for designing cities and environments that are resilient, equitable, and harmonious with nature. It highlights the potential of these technologies to inform and inspire the sustainable transformation of urban and natural landscapes for generations to come.

Throughout this thesis, the synergy between AI and EO emerges as a powerful enabler of sustainability. The novel approaches and methodologies developed have not only advanced the understanding of environmental and urban systems but also equipped stakeholders with the tools to monitor, assess, and influence the trajectory of sustainable development.

As the focus turns to the future, the integration of AI and EO will continue to play a pivotal role in addressing the complex challenges of sustainability. The insights gained and the technologies developed through this research provide a solid foundation for further innovation and implementation. The journey towards sustainability is ongoing, and the work presented in this thesis represents a significant step forward in leveraging AI and EO for the betterment of the planet and future generations.

In closing, this thesis reaffirms the transformative potential of Artificial Intelligence and Earth Observation in empowering sustainability efforts across the globe. The challenges ahead are significant, but with continued innovation and collaboration, it is possible to harness the power of these technologies to create a more sustainable, resilient, and equitable world.

Bibliography

- [1] Angelina Kouroubali and Dimitrios G Katehakis. The new european interoperability framework as a facilitator of digital transformation for citizen empowerment. *Journal of biomedical informatics*, 94:103166, 2019.
- [2] Opende design principles for data spaces. <https://design-principles-for-data-spaces.org/>. Accessed: 2024-01-12.
- [3] Leslie N Smith. A disciplined approach to neural network hyperparameters: Part 1—learning rate, batch size, momentum, and weight decay. *arXiv preprint arXiv:1803.09820*, 2018.
- [4] An Thi Ngoc Dang, Subrata Nandy, Ritika Srinet, Nguyen Viet Luong, Surajit Ghosh, and A Senthil Kumar. Forest aboveground biomass estimation using machine learning regression algorithm in yok don national park, vietnam. *Ecological Informatics*, 50:24–32, 2019.
- [5] Phyto Wai, Huiyi Su, and Mingshi Li. Estimating aboveground biomass of two different forest types in myanmar from sentinel-2 data with machine learning and geostatistical algorithms. *Remote Sensing*, 14(9):2146, 2022.
- [6] Sit comune di asti. <http://sit.comune.asti.it/site/?map=Mappa>. Accessed: 2024-01-12.
- [7] Worldpop. www.worldpop.org. Accessed: 2024-01-12.
- [8] Robert O. Vos. Defining sustainability: a conceptual orientation. 82(4):334–339, 04 2007.
- [9] Donella H. Meadows, Jørgen Randers, and Donella H. Meadows. The limits to growth (1972), 12 2017.

-
- [10] Graham Turner. A comparison of the limits to growth with 30 years of reality. *Global Environmental Change*, 18(3):397–411, 08 2008.
- [11] Ankur Dixit and Rajiv Chaudhary. Sustainability: an essential concept for current and upcoming technological practices. *IOP Conference Series: Materials Science and Engineering*, 804(1):012023–012023, 04 2020.
- [12] Bruno Ferreira, Muriel Iten, and Rui Silva. Monitoring sustainable development by means of earth observation data and machine learning: a review. *Environmental Sciences Europe*, 32(1), 09 2020.
- [13] Sustainable development goals - an overview | sciencedirect topics, 01 2023.
- [14] Katherine Anderson, Barbara J. Ryan, William E. Sonntag, Argyro Kavvada, and Lawrence Friedl. Earth observation in service of the 2030 agenda for sustainable development. *Geo-spatial Information Science*, 20(2):77–96, 04 2017.
- [15] Space debris by the numbers. https://www.esa.int/Space_Safety/Space_Debris/Space_debris_by_the_numbers. Accessed: 2024-01-12.
- [16] Ana Andries, Stephen Morse, Richard J. Murphy, J. M. Lynch, Emma Woolliams, and John Fonweban. Translation of earth observation data into sustainable development indicators: An analytical framework. *Sustainable Development*, 27(3):366–376, 11 2018.
- [17] Horst Bischof, Werner X. Schneider, and Axel Pinz. Multispectral classification of landsat-images using neural networks. *IEEE Transactions on Geoscience and Remote Sensing*, 30(3):482–490, 05 1992.
- [18] Bruno Ferreira, Muriel Iten, and Rui Silva. Monitoring sustainable development by means of earth observation data and machine learning: a review. *Environmental Sciences Europe*, 32(1), 09 2020.
- [19] Devis Tuia, Konrad Schindler, Begüm Demir, Gustau Camps-Valls, Xiao Xiang Zhu, Mrinalini Kochupillai, Sašo Džeroski, Jan N. van Rijn, Holger H. Hoos, Fabio Del Frate, Mihai Datcu, Jorge-Arnulfo Quiané-Ruiz, Volker Markl, Bertrand Le Saux, and Rochelle Schneider. Artificial intelligence to advance earth observation: a perspective. *arXiv (Cornell University)*, 05 2023.
- [20] Lorenzo Bruzzone and Begüm Demir. A review of modern approaches to classification of remote sensing data. *Remote sensing and digital image processing*, pages 127–143, 01 2014.
- [21] DeLiang L. Wang. Unsupervised learning: Foundations of neural computation. *Ai Magazine*, 22(2):101–102, 06 2001.
-

-
- [22] Benjamín Lucas, Charlotte Pelletier, Daniel F. Schmidt, Geoffrey I. Webb, and François Petitjean. Unsupervised domain adaptation techniques for classification of satellite image time series. 09 2020.
- [23] Javiera Castillo-Navarro, Bertrand Le Saux, Alexandre Boulch, Nicolas Audebert, and Sébastien Lefèvre. Semi-supervised semantic segmentation in earth observation: The minifrance suite, dataset analysis and multi-task network study. *arXiv (Cornell University)*, 10 2020.
- [24] Pablo Gómez and Gabriele Meoni. Msmatch: Semi-supervised multispectral scene classification with few labels. *arXiv (Cornell University)*, 03 2021.
- [25] Danfeng Hong, Naoto Yokoya, Gui-Song Xia, Jocelyn Chanussot, and Xiao Xiang Zhu. X-modalnet: A semi-supervised deep cross-modal network for classification of remote sensing data. 167:12–23, 09 2020.
- [26] Haifa Tamiminia, Bahram Salehi, Masoud Mahdianpari, Lindi J. Quackenbush, Sarina Adeli, and Brian Brisco. Google earth engine for geo-big data applications: A meta-analysis and systematic review. *Isprs Journal of Photogrammetry and Remote Sensing*, 164:152–170, 06 2020.
- [27] Claudio Persello, Jan Dirk Wegner, Ronny Hänsch, Devis Tuia, Pedram Ghamisi, Mila Koeva, and Gustau Camps-Valls. Deep learning and earth observation to support the sustainable development goals: Current approaches, open challenges, and future opportunities, 06 2022.
- [28] European Commission. On artificial intelligence—a european approach to excellence and trust, 2020.
- [29] European Commission. A european approach to artificial intelligence, 2022.
- [30] A european strategy for data, 02 2020.
- [31] The european green deal - european commission, 09 2022.
- [32] Abel Reiberg, Crispin Niebel, and Peter Kraemer. What is a data space. Technical report, Technical Report. Gaia-X Hub Germany, 2022.
- [33] Erica Zell, Amy K Huff, Adam T Carpenter, and Lawrence A Friedl. A user-driven approach to determining critical earth observation priorities for societal benefit. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 5(6):1594–1602, 2012.
- [34] Qihao Weng. Remote sensing of impervious surfaces in the urban areas: Requirements, methods, and trends. *Remote Sensing of Environment*, 117:34–49, 2012.
-

-
- [35] Hong Fang, Yuchun Wei, and Qiuping Dai. A novel remote sensing index for extracting impervious surface distribution from landsat 8 oli imagery. *Applied Sciences*, 9(13):2631, 2019.
- [36] Dengsheng Lu, Guiying Li, Wenhui Kuang, and Emilio Moran. Methods to extract impervious surface areas from satellite images. *International Journal of Digital Earth*, 7(2):93–112, 2014.
- [37] Nataliia Kussul, Mykola Lavreniuk, Sergii Skakun, and Andrii Shelestov. Deep learning classification of land cover and crop types using remote sensing data. *IEEE Geoscience and Remote Sensing Letters*, 14(5):778–782, 2017.
- [38] Grant J Scott, Matthew R England, William A Starms, Richard A Marcum, and Curt H Davis. Training deep convolutional neural networks for land–cover classification of high-resolution imagery. *IEEE Geoscience and Remote Sensing Letters*, 14(4):549–553, 2017.
- [39] Christopher F Brown, Steven P Brumby, Brookie Guzder-Williams, Tanya Birch, Samantha Brooks Hyde, Joseph Mazzariello, Wanda Czerwinski, Valerie J Pasquarella, Robert Haertel, Simon Ilyushchenko, et al. Dynamic world, near real-time global 10 m land use land cover mapping. *Scientific Data*, 9(1):1–17, 2022.
- [40] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [41] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [42] Caner Hazirbas, Lingni Ma, Csaba Domokos, and Daniel Cremers. Fusetnet: Incorporating depth into semantic segmentation via fusion-based cnn architecture. In *Asian conference on computer vision*, pages 213–228. Springer, 2016.
- [43] Yuliang Wang and Mingshi Li. Urban impervious surface detection from remote sensing images: A review of the methods and challenges. *IEEE Geoscience and Remote Sensing Magazine*, 7(3):64–93, 2019.
- [44] Youshui Zhang, Inakwu OA Odeh, and Chunfeng Han. Bi-temporal characterization of land surface temperature in relation to impervious surface area, ndvi and ndbi, using a sub-pixel image analysis. *International Journal of Applied Earth Observation and Geoinformation*, 11(4):256–264, 2009.
-

-
- [45] Abhisha Garg, Divyansu Pal, Hukum Singh, and Deepak Chander Pandey. A comparative study of ndbi, ndisi and ndii for extraction of urban impervious surface of dehradun [uttarakhand, india] using landsat 8 imagery. In *2016 International Conference on Emerging Trends in Communication Technologies (ETCT)*, pages 1–5. IEEE, 2016.
- [46] Joseph Knight and Margaret Voth. Mapping impervious cover using multi-temporal modis ndvi data. *Ieee journal of selected topics in applied earth observations and remote sensing*, 4(2):303–309, 2010.
- [47] Hanqiu Xu. A new index for delineating built-up land features in satellite imagery. *International Journal of Remote Sensing*, 29(14):4269–4276, 2008.
- [48] Genyun Sun, Xiaolin Chen, Xiuping Jia, Yanjuan Yao, and Zhenjie Wang. Combinational build-up index (cbi) for effective impervious surface mapping in urban areas. *IEEE Journal of selected topics in applied earth observations and remote sensing*, 9(5):2081–2092, 2015.
- [49] Zhongchang Sun, Cuizhen Wang, Huadong Guo, and Ranran Shang. A modified normalized difference impervious surface index (mndisi) for automatic urban mapping from landsat imagery. *Remote Sensing*, 9(9):942, 2017.
- [50] Wei Guo, Dengsheng Lu, and Wenhui Kuang. Improving fractional impervious surface mapping performance through combination of dmsp-ols and modis ndvi data. *Remote Sensing*, 9(4):375, 2017.
- [51] Xinwu Li, Huadong Guo, Zhongchang Sun, Guozhuang Shen, et al. Urban impervious surfaces estimation from radarsat-2 polarimetric data using svm method. *Progress In Electromagnetics Research*, 807, 2011.
- [52] Xi Cheng, Jiancheng Luo, Zhanfeng Shen, Changming Zhu, Xin Zhang, and Liegang Xia. Estimation of impervious surface based on integrated analysis of classification and regression by using svm. In *2011 IEEE International Geoscience and Remote Sensing Symposium*, pages 2809–2812. Ieee, 2011.
- [53] Nilanchal Patel and Rohit Mukherjee. Extraction of impervious features from spectral indices using artificial neural network. *Arabian Journal of Geosciences*, 8(6):3729–3741, 2015.
- [54] Lei Zhang and Qihao Weng. Annual dynamics of impervious surface in the pearl river delta, china, from 1988 to 2013, using time series landsat imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 113:86–96, 2016.
-

-
- [55] Hongsheng Zhang, Yuanzhi Zhang, and Hui Lin. Urban land cover mapping using random forest combined with optical and sar data. In *2012 IEEE International Geoscience and Remote Sensing Symposium*, pages 6809–6812. IEEE, 2012.
- [56] M Hamed Mozaffari and Won-Sook Lee. Semantic segmentation with peripheral vision. In *International Symposium on Visual Computing*, pages 421–429. Springer, 2020.
- [57] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [58] M Hamed Mozaffari and Won-Sook Lee. Encoder-decoder cnn models for automatic tracking of tongue contours in real-time ultrasound data. *Methods*, 179:26–36, 2020.
- [59] Zhongchang Sun, Xiangwei Zhao, Mengfan Wu, and Cuizhen Wang. Extracting urban impervious surface from worldview-2 and airborne lidar data using 3d convolutional neural networks. *Journal of the Indian Society of Remote Sensing*, 47(3):401–412, 2019.
- [60] Fenghua Huang, Ying Yu, and Tinghao Feng. Automatic extraction of impervious surfaces from high resolution remote sensing images based on deep learning. *Journal of Visual Communication and Image Representation*, 58:453–461, 2019.
- [61] Hongsheng Zhang, Luoma Wan, Ting Wang, Yinyi Lin, Hui Lin, and Zezhong Zheng. Impervious surface estimation from optical and polarimetric sar data using small-patched deep convolutional networks: A comparative study. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2374–2387, 2019.
- [62] Yongyong Fu, Kunkun Liu, Zhangquan Shen, Jinsong Deng, Muye Gan, Xinguo Liu, Dongming Lu, and Ke Wang. Mapping impervious surfaces in town–rural transition belts using china’s gf-2 imagery and object-based deep cnns. *Remote Sensing*, 11(3):280, 2019.
- [63] Deepak Kumar, Mandvi Misra, and Sulochana Shekhar. Assessing machine learning based supervised classifiers for built-up impervious surface area extraction from sentinel-2 images. *Urban Forestry & Urban Greening*, page 126714, 2020.
- [64] Khelifa Djerriri, Reda Adjouj, and Dalila Attaf. Convolutional neural networks for the extraction of built-up areas from sentinel-2 images. In *Proceedings of the 20th AGILE Conference on Geographic Information Science*, pages 9–12, 2017.
-

-
- [65] Deepank Verma and Arnab Jana. Lulc classification methodology based on simple convolutional neural network to map complex urban forms at finer scale: Evidence from mumbai. *arXiv preprint arXiv:1909.09774*, 2019.
- [66] Sankaranarayanan Piramanayagam, Eli Saber, Wade Schwartzkopf, and Frederick Koehler. Supervised classification of multisensor remotely sensed images using a deep learning framework. *Remote Sensing*, 10(9):1429, 2018.
- [67] Nicolas Audebert, Bertrand Le Saux, and Sébastien Lefèvre. Beyond rgb: Very high resolution urban remote sensing with multimodal deep networks. *ISPRS Journal of Photogrammetry and Remote Sensing*, 140:20–32, 2018.
- [68] Dimitrios Marmanis, Mihai Datcu, Thomas Esch, and Uwe Stilla. Deep learning earth observation classification using imagenet pretrained networks. *IEEE Geoscience and Remote Sensing Letters*, 13(1):105–109, 2015.
- [69] Tong He, Zhi Zhang, Hang Zhang, Zhongyue Zhang, Junyuan Xie, and Mu Li. Bag of tricks for image classification with convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 558–567, 2019.
- [70] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [71] Leslie N Smith. Cyclical learning rates for training neural networks. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 464–472. IEEE, 2017.
- [72] Scott Kirkpatrick, C Daniel Gelatt, and Mario P Vecchi. Optimization by simulated annealing. *science*, 220(4598):671–680, 1983.
- [73] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [74] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019.
- [75] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018.
- [76] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015.
- [77] Edward T. A. Mitchard. The tropical forest carbon cycle and climate change. *Nature*, 559(7715):527–534, July 2018.
-

-
- [78] Steven E. Franklin. Remote sensing for sustainable forest management, June 2001.
- [79] Louis R. Iverson, Robin Lambert Graham, and Elizabeth A. Cook. Applications of satellite remote sensing to forested ecosystems. *Landscape Ecology*, 3(2):131–143, December 1989.
- [80] Maurizio Santoro, Oliver Cartus, Johan E. S. Fransson, and Urs Wegmüller. Complementarity of x-, c-, and l-band sar backscatter observations to retrieve forest stem volume in boreal forest. *Remote Sensing*, 11(13):1563, July 2019.
- [81] Jorge Curiel-Esparza, Nuria Gonzalez-Utrillas, Julian Canto-Perello, and Manuel Martin-Utrillas. Integrating climate change criteria in reforestation projects using a hybrid decision-support system. *Environmental Research Letters*, 10(9):094022, 2015.
- [82] Ram K Deo, Matthew B Russell, Grant M Domke, Hans-Erik Andersen, Warren B Cohen, and Christopher W Woodall. Evaluating site-specific and generic spatial models of aboveground forest biomass based on landsat time-series and lidar strip samples in the eastern usa. *Remote Sensing*, 9(6):598, 2017.
- [83] United nations. sustainable development goals. Available online: <https://sdgs.un.org/goals>.
- [84] P. W. West. *Tree and Forest Measurement*. Springer International Publishing, 2015.
- [85] Dengsheng Lu. The potential and challenge of remote sensing-based biomass estimation. *International Journal of Remote Sensing*, 27(7):1297–1328, April 2006.
- [86] Dengsheng Lu. The potential and challenge of remote sensing-based biomass estimation. *International journal of remote sensing*, 27(7):1297–1328, 2006.
- [87] Dengsheng Lu, Qi Chen, Guangxing Wang, Lijuan Liu, Guiying Li, and Emilio Moran. A survey of remote sensing-based aboveground biomass estimation methods in forest ecosystems. *International Journal of Digital Earth*, 9(1):63–105, 2016.
- [88] Yingchang Li, Mingyang Li, Chao Li, and Zhenzhen Liu. Forest above-ground biomass estimation using landsat 8 and sentinel-1a data with machine learning algorithms. *Scientific Reports*, 10(1), June 2020.
-

-
- [89] Donato Amitrano, Gerardo Di Martino, Raffaella Guida, Pasquale Iervolino, Antonio Iodice, Maria Nicolina Papa, Daniele Riccio, and Giuseppe Ruello. Earth environmental monitoring using multi-temporal synthetic aperture radar: A critical review of selected applications. *Remote Sensing*, 13(4):604, February 2021.
- [90] E. T. A. Mitchard, S. S. Saatchi, I. H. Woodhouse, G. Nangendo, N. S. Ribeiro, M. Williams, C. M. Ryan, S. L. Lewis, T. R. Feldpausch, and P. Meir. Using satellite radar backscatter to predict above-ground woody biomass: A consistent relationship across four different african landscapes. *Geophysical Research Letters*, 36(23), December 2009.
- [91] Josef Aschbacher and Maria Pilar Milagro-Pérez. The european earth monitoring (gmes) programme: Status and perspectives. *Remote Sensing of Environment*, 120:3–8, 2012.
- [92] C Shoko and O Mutanga. Examining the strength of the newly-launched sentinel 2 msi sensor in detecting and discriminating subtle differences between c3 and c4 grass species. *ISPRS journal of photogrammetry and remote sensing*, 129:32–40, 2017.
- [93] Luofan Dong, Huaqiang Du, Ning Han, Xuejian Li, Di'en Zhu, Fangjie Mao, Meng Zhang, Junlong Zheng, Hua Liu, Zihao Huang, et al. Application of convolutional neural network on lei bamboo above-ground-biomass (agb) estimation using worldview-2. *Remote Sensing*, 12(6):958, 2020.
- [94] Lincoln Vinicius Schreiber, Joao Gustavo Atkinson Amorim, Leticia Guimarães, Debora Motta Matos, Celso Maciel da Costa, and Adriane Parraga. Above-ground biomass wheat estimation: Deep learning with uav-based rgb images. *Applied Artificial Intelligence*, 36(1):2055392, 2022.
- [95] Devis Tuia, Claudio Persello, and Lorenzo Bruzzone. Domain adaptation for the classification of remote sensing data: An overview of recent advances. *IEEE Geoscience and Remote Sensing Magazine*, 4(2):41–57, June 2016.
- [96] Dengsheng Lu, Qi Chen, Guangxing Wang, Lijuan Liu, Guiying Li, and Emilio Moran. A survey of remote sensing-based aboveground biomass estimation methods in forest ecosystems. *International Journal of Digital Earth*, 9(1):63–105, December 2014.
- [97] Nicholas S. Skowronski, Kenneth L. Clark, Michael Gallagher, Richard A. Birdsey, and John L. Hom. Airborne laser scanner-assisted estimation of aboveground biomass change in a temperate oak–pine forest. *Remote Sensing of Environment*, 151:166–174, August 2014.
-

-
- [98] Erik Næsset, Terje Gobakken, Ole Martin Bollandsås, Timothy G. Gregoire, Ross Nelson, and Göran Ståhl. Comparison of precision of biomass estimates in regional field sample surveys and airborne lidar-assisted surveys in hedmark county, norway. *Remote Sensing of Environment*, 130:108–120, March 2013.
- [99] Zizhao Li, Shoudong Bi, Shuang Hao, and Yuhuan Cui. Aboveground biomass estimation in forests with random forest and monte carlo-based uncertainty analysis. *Ecological Indicators*, 142:109246, September 2022.
- [100] Yingchang Li, Mingyang Li, and Yuehui Wang. Forest aboveground biomass estimation and response to climate change based on remote sensing data. *Sustainability*, 14(21):14222, October 2022.
- [101] Jean-Matthieu Monnet, Jocelyn Chanussot, and Frédéric Berger. Support vector regression for the estimation of forest stand parameters using airborne laser scanning. *IEEE Geoscience and Remote Sensing Letters*, 8(3):580–584, May 2011.
- [102] Mthembeni Mngadi, John Odindi, and Onesimo Mutanga. The utility of sentinel-2 spectral data in quantifying above-ground carbon stock in an urban reforested landscape. *Remote Sensing*, 13(21):4281, 2021.
- [103] Robert M Haralick, Karthikeyan Shanmugam, and Its' Hak Dinstein. Textural features for image classification. *IEEE Transactions on systems, man, and cybernetics*, (6):610–621, 1973.
- [104] Santa Pandit, Satoshi Tsuyuki, and Timothy Dube. Estimating above-ground biomass in sub-tropical buffer zone community forests, nepal, using sentinel 2 data. *Remote Sensing*, 10(4):601, 2018.
- [105] Fardin Moradi, Ali Asghar Darvishsefat, Manizheh Rajab Pourrahmati, Azade Deljouei, and Stelian Alexandru Borz. Estimating aboveground biomass in dense hyrcanian forests by the use of sentinel-2 data. *Forests*, 13(1):104, 2022.
- [106] Chunhua Li, Lizhi Zhou, and Wenbin Xu. Estimating aboveground biomass using sentinel-2 msi data and ensemble algorithms for grassland in the shengjin lake wetland, china. *Remote Sensing*, 13(8):1595, 2021.
- [107] Santa Pandit, Satoshi Tsuyuki, and Timothy Dube. Exploring the inclusion of sentinel-2 msi texture metrics in above-ground biomass estimation in the community forest of nepal. *Geocarto International*, 35(16):1832–1849, 2020.
- [108] Parth Naik, Michele Dalponte, and Lorenzo Bruzzone. Generative feature extraction from sentinel 1 and 2 data for prediction of forest aboveground
-

- biomass in the italian alps. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 15:4755–4771, 2022.
- [109] Parth Naik, Michele Dalponte, and Lorenzo Bruzzone. Automated machine learning driven stacked ensemble modelling for forest aboveground biomass prediction using multitemporal sentinel-2 data. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2022.
- [110] Lu Li, Boqi Zhou, Yanfeng Liu, Yong Wu, Jing Tang, Weiheng Xu, Leiguang Wang, and Guanglong Ou. Reduction in uncertainty in forest aboveground biomass estimation using sentinel-2 images: A case study of pinus densata forests in shangri-la city, china. *Remote Sensing*, 15(3):559, 2023.
- [111] Philipp Probst and Anne-Laure Boulesteix. To tune or not to tune the number of trees in random forest. *Journal of Machine Learning Research*, 18(181):1–18, 2018.
- [112] Tiago G. Morais, Ricardo F.M. Teixeira, Mario Figueiredo, and Tiago Domingos. The use of machine learning methods to estimate aboveground biomass of grasslands: A review. *Ecological Indicators*, 130:108081, November 2021.
- [113] Marston Franceschini, Rolf Becker, Florian Wichern, and Lammert Kooistra. Quantification of grassland biomass and nitrogen content through uav hyperspectral imagery—active sample selection for model transfer. *Drones*, 6(3):73, March 2022.
- [114] Liang Wan, Haiyan Cen, Jiangpeng Zhu, Jiafei Zhang, Yueming Zhu, Dawei Sun, Xiaoyue Du, Li Zhai, Haiyong Weng, Yijian Li, Xiaoran Li, Yidan Bao, Jianyao Shou, and Yong He. Grain yield prediction of rice using multi-temporal uav-based rgb and multispectral images and model transfer – a case study of small farmlands in the south of china. *Agricultural and Forest Meteorology*, 291:108096, September 2020.
- [115] Donato Amitrano, Luca Cicala, Marco De Mizio, and Francesco Tufano. Uav hyperspectral characterization of vegetation using entropy-based active sampling for partial least square regression models. *Applied Sciences*, 13(8):4812, April 2023.
- [116] Patrick A. Fekety, Michael J. Falkowski, and Andrew T. Hudak. Temporal transferability of lidar-based imputation of forest inventory attributes. *Canadian Journal of Forest Research*, 45(4):422–435, April 2015.
- [117] Piotr Tompalski, Joanne C. White, Nicholas C. Coops, and Michael A. Wulder. Demonstrating the transferability of forest inventory attribute
-

- models derived using airborne laser scanning data. *Remote Sensing of Environment*, 227:110–124, June 2019.
- [118] Francisco Mauro, Andrew T. Hudak, Patrick A. Fekety, Bryce Frank, Hailemariam Temesgen, David M. Bell, Matthew J. Gregory, and T. Ryan McCarley. Regional modeling of forest fuels and structural attributes using airborne laser scanning data in oregon. *Remote Sensing*, 13(2):261, January 2021.
- [119] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [120] Paolo Fazzini, Giuseppina De Felice Proia, Maria Adamo, Palma Blonda, Francesco Petracchini, Luigi Forte, and Cristina Tarantino. Sentinel-2 remote sensed image classification with patchwise trained convnets for grassland habitat discrimination. *Remote Sensing*, 13(12):2276, 2021.
- [121] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [122] Ian T Jolliffe and Jorge Cadima. Principal component analysis: a review and recent developments. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374(2065):20150202, 2016.
- [123] Asim Banskota, Randolph H Wynne, and Nilam Kayastha. Improving within-genus tree species discrimination using the discrete wavelet transform applied to airborne hyperspectral data. *International Journal of Remote Sensing*, 32(13):3551–3563, 2011.
- [124] Lutz Prechelt. Early stopping-but when? In *Neural Networks: Tricks of the trade*, pages 55–69. Springer, 1998.
- [125] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [126] Keras: Reducelronplateau. Available online: https://keras.io/api/callbacks/reduce_lr_on_plateau/.
- [127] Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. *An introduction to statistical learning*, volume 112. Springer, 2013.
- [128] Lauri Korhonen, Petteri Packalen, Miina Rautiainen, et al. Comparison of sentinel-2 and landsat 8 in the estimation of boreal forest canopy cover and leaf area index. *Remote sensing of environment*, 195:259–274, 2017.
-

-
- [129] Sentinel-2 products specification document. Available online: <https://sentinel.esa.int/documents/247904/685211/sentinel-2-products-specification-document>.
- [130] Esa's climate change initiative biomass project: Global datasets of forest above-ground biomass. Available online: <https://climate.esa.int/en/projects/biomass/data/>.
- [131] Maurizio Santoro, Oliver Cartus, Nuno Carvalhais, Danaë Rozendaal, Valerio Avitabile, Arnan Araza, Sytze De Bruin, Martin Herold, Shaun Quegan, Pedro Rodriguez-Veiga, et al. The global forest above-ground biomass pool for 2010 estimated from high-resolution satellite observations. *Earth System Science Data*, 13(8):3927–3950, 2021.
- [132] Michele Innangi, Rosario Balestrieri, Tiziana Danise, Francesco d'Alessandro, and Antonietta Fioretto. From soil to bird community: A partial least squares approach to investigate a natural wooded area surrounded by urban patchwork (astroni crater, southern italy). *Ecological Modelling*, 394:1–10, 2019.
- [133] E Ch Tooichi. Carbon sequestration: how much can forestry sequester co2. *Forestry Research and Engineering: International Journal*, 2(3):148–150, 2018.
- [134] Richard A Birdsey. *Carbon storage and accumulation in United States forest ecosystems*, volume 59. US Department of Agriculture, Forest Service, 1992.
- [135] Matti Maltamo, Hermanni Kinnunen, Annika Kangas, and Lauri Korhonen. Predicting stand age in managed forests using national forest inventory field data and airborne laser scanning. *Forest Ecosystems*, 7(1), July 2020.
- [136] Virpi Stenman. Finnish forest data-based metsään. fi-services. *Big Data in Bioeconomy: Results from the European DataBio Project*, pages 309–319, 2021.
- [137] Sakari Tuominen, Timo Pitkänen, Andras Balazs, and Annika Kangas. Improving finnish multi-source national forest inventory by 3d aerial imaging. *Silva Fennica*, 51(4), 2017.
- [138] Sijmen de Jong. Simpls: An alternative approach to partial least squares regression. *Chemometrics and Intelligent Laboratory Systems*, 18(3):251–263, March 1993.
- [139] Robert M. Haralick, K. Shanmugam, and Its'Hak Dinstein. Textural features for image classification. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-3(6):610–621, November 1973.
-

-
- [140] Jia Jin and Quan Wang. Evaluation of informative bands used in different pls regressions for estimating leaf biochemical contents from hyperspectral reflectance. *Remote Sensing*, 11(2):197, January 2019.
- [141] Mireia Farrés, Stefan Platikanov, Stefan Tsakovski, and Romà Tauler. Comparison of the variable importance in projection (vip) and of the selectivity ratio (sr) methods for variable selection and interpretation. *Journal of Chemometrics*, 29(10):528–536, July 2015.
- [142] Marina Cocchi, Alessandra Biancolillo, and Federico Marini. *Chemometric Methods for Classification and Feature Selection*, page 265–299. Elsevier, 2018.
- [143] Linda Shapiro and C George. Stockman g: computer vision. In *Prentice Hall*. 2002.
- [144] Katja Berger, Juan Pablo Rivera Caicedo, Luca Martino, Matthias Woche, Tobias Hank, and Jochem Verrelst. A survey of active learning for quantifying vegetation traits from terrestrial earth observation data. *Remote Sensing*, 13(2):287, January 2021.
- [145] Tianxu He, Shukui Zhang, Jie Xin, Pengpeng Zhao, Jian Wu, Xuefeng Xian, Chunhua Li, and Zhiming Cui. An active learning approach with uncertainty, representativeness, and diversity. *The Scientific World Journal*, 2014:1–6, 2014.
- [146] Xiaochen Lu, Junping Zhang, Tong Li, and Ye Zhang. Incorporating diversity into self-learning for synergetic classification of hyperspectral and panchromatic images. *Remote Sensing*, 8(10):804, September 2016.
- [147] Fouzi Douak, Farid Melgani, and Nabil Benoudjit. Kernel ridge regression with active learning for wind speed prediction. *Applied Energy*, 103:328–340, March 2013.
- [148] Biao Yuan, Zhongyuan Wu, Kai Zhang, Deqi Li, and Qiaoyu Ma. Application of active learning in carbonate lithologic identification. In *2021 4th International Conference on Artificial Intelligence and Big Data (ICAIBD)*. IEEE, May 2021.
- [149] Begüm Demir, Claudio Persello, and Lorenzo Bruzzone. Batch-mode active-learning methods for the interactive classification of remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing*, 49(3):1014–1031, March 2011.
- [150] Swarnajyoti Patra and Lorenzo Bruzzone. A fast cluster-assumption based active-learning technique for classification of remote sensing images. *IEEE*
-

- Transactions on Geoscience and Remote Sensing*, 49(5):1617–1626, May 2011.
- [151] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(4):623–656, October 1948.
- [152] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. Springer New York, 2009.
- [153] Arnon Karnieli, Yoram J. Kaufman, Lorraine Remer, and Andrew Wald. Afri — aerosol free vegetation index. *Remote Sensing of Environment*, 77(1):10–21, July 2001.
- [154] A. Bannari, D. Morin, F. Bonn, and A. R. Huete. A review of vegetation indices. *Remote Sensing Reviews*, 13(1–2):95–120, August 1995.
- [155] Russell Main, Moses Azong Cho, Renaud Mathieu, Martha M. O’Kennedy, Abel Ramoelo, and Susan Koch. An investigation into robust spectral indices for leaf chlorophyll estimation. *ISPRS Journal of Photogrammetry and Remote Sensing*, 66(6):751–761, November 2011.
- [156] G. le Maire, C. François, and E. Dufrêne. Towards universal broad leaf chlorophyll indices using prospect simulated database and hyperspectral reflectance measurements. *Remote Sensing of Environment*, 89(1):1–28, January 2004.
- [157] C.J. Tucker, J.H. Elgin, J.E. McMurtrey, and C.J. Fan. Monitoring corn and soybean crop development with hand-held radiometer spectral data. *Remote Sensing of Environment*, 8(3):237–248, August 1979.
- [158] D Haboudane. Hyperspectral vegetation indices and novel algorithms for predicting green lai of crop canopies: Modeling and validation in the context of precision agriculture. *Remote Sensing of Environment*, 90(3):337–352, April 2004.
- [159] Ruiliang Pu, Peng Gong, and Qian Yu. Comparative analysis of eo-1 ali and hyperion, and landsat etm+ data for mapping forest crown closure and leaf area index. *Sensors*, 8(6):3744–3766, June 2008.
- [160] J. G. P. W. Clevers, S. M. De Jong, G. F. Epema, F. D. Van Der Meer, W. H. Bakker, A. K. Skidmore, and K. H. Scholte. Derivation of the red edge index using the meris standard band setting. *International Journal of Remote Sensing*, 23(16):3169–3184, January 2002.
- [161] Michio Shibayama, Arto Salli, Tuomas Häme, Lasse Iso-Iivari, Saini Heino, Marjaana Alanen, Shinsuke Morinaga, Yoshio Inoue, and Tsuyoshi Akiyama. Detecting phenophases of subarctic shrub canopies by using
-

- automated reflectance measurements. *Remote Sensing of Environment*, 67(2):160–180, February 1999.
- [162] United nations. what is redd+? Available online: <https://unfccc.int/topics/land-use/workstreams/redd/what-is-redd>.
- [163] Sasan Vafaei, Javad Soosani, Kamran Adeli, Hadi Fadaei, Hamed Naghavi, Tien Pham, and Dieu Tien Bui. Improving accuracy estimation of forest aboveground biomass based on incorporation of alos-2 palsar-2 and sentinel-2a imagery and machine learning: A case study of the hyrcanian forest area (iran). *Remote Sensing*, 10(2):172, January 2018.
- [164] Emanuele Santi, Simonetta Paloscia, Simone Pettinato, Giacomo Fontanelli, Matteo Mura, Catherine Zolli, Fabio Maselli, Marta Chiesi, Lorenzo Bottai, and Gherardo Chirici. The potential of multifrequency sar images for estimating forest biomass in mediterranean areas. *Remote Sensing of Environment*, 200:63–73, October 2017.
- [165] Luofan Dong, Huaqiang Du, Ning Han, Xuejian Li, Di'en Zhu, Fangjie Mao, Meng Zhang, Junlong Zheng, Hua Liu, Zihao Huang, and Shaobai He. Application of convolutional neural network on lei bamboo above-ground-biomass (agb) estimation using worldview-2. *Remote Sensing*, 12(6):958, March 2020.
- [166] Antonio Elia Pascarella, Giovanni Giacco, Mattia Rigioli, Stefano Marrone, and Carlo Sansone. Reuse: Regressive unet for carbon storage and above ground biomass estimation. February 2023.
- [167] High value datasets - initiative details - european union. https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/12111-Open-data-availability-of-public-datasets_en. Accessed: 2024-01-12.
- [168] Spotted. <https://cef-spotted.eu/>. Accessed: 2024-01-12.
- [169] European data portal. <https://data.europa.eu/en>. Accessed: 2024-01-12.
- [170] Fiware. <https://www.fiware.org/>. Accessed: 2024-01-12.
- [171] Ulrich Ahle and Juan Jose Hierro. Fiware for data spaces. *Designing Data Spaces*, page 395, 2022.
- [172] Andres Munoz-Arcenales, Sonsoles López-Pernas, Alejandro Pozo, Álvaro Alonso, Joaquín Salvachúa, and Gabriel Huecas. Data usage and access control in industrial data spaces: Implementation using fiware. *Sustainability*, 12(9):3885, 2020.
-

-
- [173] Álvaro Alonso, Alejandro Pozo, José Manuel Cantera, Francisco De la Vega, and Juan José Hierro. Industrial data space architecture implementation using fiware. *Sensors*, 18(7):2226, 2018.
- [174] Gürkan Solmaz, Flavio Cirillo, Jonathan Fürst, Tobias Jacobs, Martin Bauer, Ernő Kovacs, Juan Ramón Santana, and Luis Sánchez. Enabling data spaces: Existing developments and challenges. In *Proceedings of the 1st International Workshop on Data Economy*, pages 42–48, 2022.
- [175] Katarína Onačillová, Michal Gallay, Daniel Paluba, Anna Péliová, Ondrej Tokarčík, and Daniela Laubertová. Combining landsat 8 and sentinel-2 data in google earth engine to derive higher resolution land surface temperature maps in urban environment. *Remote Sensing*, 14(16):4076, 2022.
- [176] Stefania Bonafoni. Downscaling of landsat and modis land surface temperature over the heterogeneous urban area of milan. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 9(5):2019–2027, 2016.
- [177] Kai Liu, Shudong Wang, Xueke Li, and Taixia Wu. Spatially disaggregating satellite land surface temperature with a nonlinear model across agricultural areas. *Journal of Geophysical Research: Biogeosciences*, 124(11):3232–3251, 2019.
- [178] F Gao, WP Kustas, and MC Anderson. A data mining approach for sharpening thermal satellite imagery over land. *remote sens.* 4, 3287–3319, 2012.
- [179] Wenfeng Zhan, Yunhao Chen, Ji Zhou, Jinfei Wang, Wenyu Liu, James Voogt, Xiaolin Zhu, Jinling Quan, and Jing Li. Disaggregation of remotely sensed land surface temperature: Literature survey, taxonomy, issues, and caveats. *Remote Sensing of Environment*, 131:119–139, 2013.
- [180] Assemblée General. *Transforming our world: The 2030 agenda for sustainable development*. UN, 2015.
- [181] United Nations. Department of Economic and Social Affairs. *The Sustainable Development Goals: Report 2022*. UN, 2022.
- [182] Chona Sister, Jennifer Wolch, and John Wilson. Got green? addressing environmental justice in park provision. *GeoJournal*, 75(3):229–248, September 2009.
- [183] N Vasiljević, B Radić, S Gavrilović, B Šljukić, M Medarević, and R Ristić. The concept of green infrastructure and urban landscape planning: a challenge for urban forestry planning in belgrade, serbia. *iForest - Biogeosciences and Forestry*, 11(4):491–498, August 2018.
-

-
- [184] *Ecosystem Services and Green Infrastructure: Perspectives from Spatial Planning in Italy*. Springer International Publishing, 2021.
- [185] Robert W Miller, Richard J Hauer, and Les P Werner. *Urban forestry: planning and managing urban greenspaces*. Waveland press, 2015.
- [186] Diana E. Bowler, Lisette Buyung-Ali, Teri M. Knight, and Andrew S. Pullin. Urban greening to cool towns and cities: A systematic review of the empirical evidence. *Landscape and Urban Planning*, 97(3):147–155, September 2010.
- [187] Silvija Krajter Ostoić and Cecil C. Konijnendijk van den Bosch. Exploring global scientific discourses on urban forestry. *Urban Forestry amp; Urban Greening*, 14(1):129–138, 2015.
- [188] Geoffrey H. Donovan. Including public-health benefits of trees in urban-forestry decision making. *Urban Forestry amp; Urban Greening*, 22:120–123, March 2017.
- [189] Fabio Salbitano, Simone Borelli, Michela Conigliaro, Yujuan Chen, et al. Guidelines on urban and peri-urban forestry. *FAO Forestry Paper*, (178), 2016.
- [190] P.J. Irga, M.D. Burchett, and F.R. Torpy. Does urban forestry have a quantitative effect on ambient air quality in an urban environment? *Atmospheric Environment*, 120:173–181, November 2015.
- [191] Kyle H. Clark and Kimberly A. Nicholas. Introducing urban food forestry: a multifunctional approach to increase food security and provide ecosystem services. *Landscape Ecology*, 28(9):1649–1669, July 2013.
- [192] Zhanghao Chen and Ganlin Huang. Greenspace to meet people’s demand: A case study of beijing in 2005 and 2015. *Remote Sensing*, 13(21):4310, October 2021.
- [193] Jiayu Wu, Yisong Peng, Ping Liu, Yubing Weng, and Jian Lin. Is the green inequality overestimated? quality reevaluation of green space accessibility. *Cities*, 130:103871, November 2022.
- [194] Wenyue Yang, Ruyue Yang, and Suhong Zhou. The spatial heterogeneity of urban green space inequity from a perspective of the vulnerable: A case study of guangzhou, china. *Cities*, 130:103855, November 2022.
- [195] Rebecca Leigh Rutt and Natalie Marie Gulsrud. Green justice in the city: A new agenda for urban green space research in europe. *Urban Forestry amp; Urban Greening*, 19:123–127, September 2016.
-

-
- [196] Xun Li, Wendy Y. Chen, Giovanni Sanesi, and Raffaele Laforteza. Remote sensing in urban forestry: Recent applications and future directions. *Remote Sensing*, 11(10):1144, May 2019.
- [197] Sulochana Shekhar and Jagannath Aryal. Role of geospatial technology in understanding urban green space of kalaburagi city for sustainable planning. *Urban Forestry amp; Urban Greening*, 46:126450, December 2019.
- [198] Francisco J. Escobedo, Vincenzo Giannico, C.Y. Jim, Giovanni Sanesi, and Raffaele Laforteza. Urban forests, ecosystem services, green infrastructure and nature-based solutions: Nexus or evolving metaphors? *Urban Forestry amp; Urban Greening*, 37:3–12, January 2019.
- [199] Hade Dorst, Alexander van der Jagt, Rob Raven, and Hens Runhaar. Urban greening through nature-based solutions – key characteristics of an emerging concept. *Sustainable Cities and Society*, 49:101620, August 2019.
- [200] Carsten Nesshöver, Timo Assmuth, Katherine N. Irvine, Graciela M. Rusch, Kerry A. Waylen, Ben Delbaere, Dagmar Haase, Lawrence Jones-Walters, Hans Keune, Eszter Kovacs, Kinga Krauze, Mart Külvik, Freddy Rey, Jiska van Dijk, Odd Inge Vistad, Mark E. Wilkinson, and Heidi Wittmer. The science, policy and practice of nature-based solutions: An interdisciplinary perspective. *Science of The Total Environment*, 579:1215–1227, February 2017.
- [201] João Paulo Fernandes and Nuno Guiomar. Nature-based solutions: The need to increase the knowledge on their potentialities and limits. *Land Degradation amp; Development*, 29(6):1925–1939, April 2018.
- [202] Nadja Kabisch, Niki Frantzeskaki, Stephan Pauleit, Sandra Naumann, McKenna Davis, Martina Artmann, Dagmar Haase, Sonja Knapp, Horst Korn, Jutta Stadler, Karin Zaunberger, and Aletta Bonn. Nature-based solutions to climate change mitigation and adaptation in urban areas: perspectives on indicators, knowledge gaps, barriers, and opportunities for action. *Ecology and Society*, 21(2), 2016.
- [203] Sebastian Scheuer, Jessica Jache, Martina Kičić, Thilo Wellmann, Manuel Wolff, and Dagmar Haase. A trait-based typification of urban forests as nature-based solutions. *Urban Forestry amp; Urban Greening*, 78:127780, December 2022.
- [204] Chiara Cortinovis, Peter Olsson, Niklas Boke-Olén, and Katarina Hedlund. Scaling up nature-based solutions for climate-change adaptation: Potential and benefits in three european cities. *Urban Forestry amp; Urban Greening*, 67:127450, January 2022.
-

-
- [205] Chiara Cortinovis and Davide Geneletti. A framework to explore the effects of urban planning decisions on regulating ecosystem services in cities. *Ecosystem Services*, 38:100946, August 2019.
- [206] Lorien Nesbitt, Michael J. Meitner, Cynthia Girling, Stephen R.J. Sheppard, and Yuhao Lu. Who has access to urban vegetation? a spatial analysis of distributional green equity in 10 us cities. *Landscape and Urban Planning*, 181:51–79, January 2019.
- [207] Panagiota Kotsila, Isabelle Anguelovski, Francesc Baró, Johannes Lange-meyer, Filka Sekulova, and James JT Connolly. Nature-based solutions as discursive tools and contested practices in urban nature’s neoliberalisation processes. *Environment and planning E: Nature and space*, 4(2):252–274, 2021.
- [208] Isabelle Anguelovski and Esteve Corbera. Integrating justice in nature-based solutions to avoid nature-enabled dispossession. *Ambio*, 52(1):45–53, August 2022.
- [209] Reza Ramyar, Aidan Ackerman, and Douglas M. Johnston. Adapting cities for climate change through urban green infrastructure planning. *Cities*, 117:103316, October 2021.
- [210] Isabelle Anguelovski, James JT Connolly, Melissa Garcia-Lamarca, Helen Cole, and Hamil Pearsall. New scholarly pathways on green gentrification: What does the urban ‘green turn’ mean and where is it going? *Progress in Human Geography*, 43(6):1064–1086, October 2018.
- [211] Anders Blok. Urban green gentrification in an unequal world of climate change. *Urban Studies*, 57(14):2803–2816, January 2020.
- [212] Dagmar Haase, Sigrun Kabisch, Annegret Haase, Erik Andersson, Ellen Banzhaf, Francesc Baró, Miriam Brenck, Leonie K. Fischer, Niki Frantzeskaki, Nadja Kabisch, Kerstin Krellenberg, Peleg Kremer, Jakub Kronenberg, Neele Larondelle, Juliane Mathey, Stephan Pauleit, Irene Ring, Dieter Rink, Nina Schwarz, and Manuel Wolff. Greening cities – to be socially inclusive? about the alleged paradox of society and ecology in cities. *Habitat International*, 64:41–48, June 2017.
- [213] João Monteiro, Nuno Sousa, Eduardo Natividade-Jesus, and João Coutinho-Rodrigues. Benchmarking city layouts—a methodological approach and an accessibility comparison between a real city and the garden city. *Sustainability*, 14(9):5029, April 2022.
- [214] Christian Albert, Barbara Schröter, Dagmar Haase, Mario Brillinger, Jennifer Henze, Sylvia Herrmann, Sarah Gottwald, Paulina Guerrero, Claire
-

- Nicolas, and Bettina Matzdorf. Addressing societal challenges through nature-based solutions: How can landscape planning and governance research contribute? *Landscape and Urban Planning*, 182:12–21, February 2019.
- [215] Judy Bush and Andréanne Doyon. Building urban resilience with nature-based solutions: How can urban planning contribute? *Cities*, 95:102483, December 2019.
- [216] Frank S Santamour Jr. Trees for urban planting: diversity uniformity, and common sense. *C. Elevitch, The Overstory Book: cultivating connections with trees*, pages 396–399, 2004.
- [217] Cecil C. Konijnendijk. Evidence-based guidelines for greener, healthier, more resilient neighbourhoods: Introducing the 3–30–300 rule. *Journal of Forestry Research*, 34(3):821–830, August 2022.
- [218] Ma.D. Velarde, G. Fry, and M. Tveit. Health effects of viewing landscapes – landscape types in environmental psychology. *Urban Forestry amp; Urban Greening*, 6(4):199–212, November 2007.
- [219] Federica Larcher, Enrico Pomatto, Luca Battisti, Paola Gullino, and Marco Devecchi. Perceptions of urban green areas during the social distancing period for covid-19 containment in italy. *Horticulturae*, 7(3):55, March 2021.
- [220] Matilda Annerstedt van den Bosch, Pierpaolo Mudu, Valdas Uscila, Maria Barndahl, Alexandra Kulinkina, Brigit Staatsen, Wim Swart, Hanneke Kruize, Ingrida Zurlyte, and Andrey I. Egorov. Development of an urban green space indicator and the public health rationale. *Scandinavian Journal of Public Health*, 44(2):159–167, November 2015.
- [221] Mirjam Schindler, Marion Le Texier, and Geoffrey Caruso. How far do people travel to use urban green space? a comparison of three european cities. *Applied Geography*, 141:102673, April 2022.
- [222] Torino vivibile, 2020. <https://www.torinovivibile.it/aree-tematiche/torino-green-print/>. Accessed: 2024-01-12.
- [223] L. Zardo, D. Geneletti, M. Pérez-Soba, and M. Van Eupen. Estimating the cooling capacity of green infrastructures to support urban planning. *Ecosystem Services*, 26:225–235, August 2017.
- [224] Luca Battisti, Lauranne Pille, Thomas Wachtel, Federica Larcher, and Ina Säumel. Residential greenery: State of the art and health-related ecosystem services and disservices in the city of berlin. *Sustainability*, 11(6):1815, March 2019.
-

-
- [225] High resolution layer forest type. <https://land.copernicus.eu/en/products/high-resolution-layer-forest-type>. Accessed: 2024-01-12.
- [226] WHO Europe. Urban green spaces: a brief for action. *World Health Organization*. Abgerufen von. Available from <https://apps.who.int/iris/handle/10665/344116>, 2017.
- [227] Italian national institute of statistics (istat). <https://demo.istat.it/app/?i=D7B&a=2022&l=en>. Accessed: 2024-01-12.
- [228] Città di asti. il verde cittadino. i dati del verde. https://www.comune.asti.it/pagina922_il-verde-cittadino.html. Accessed: 2024-01-12.
- [229] Italian national institute of statistics (istat) (2023, january 18). demographic indicators. <http://dati.istat.it/Index.aspx?lang=en&SubSessionId=d92a2215-ab06-4114-974a-cee80a24c909#>. Accessed: 2024-01-12.
- [230] Italian national institute of statistics (istat) (2023, april 22). resident population on 1st january. <http://dati.istat.it/Index.aspx?lang=en&SubSessionId=e60188e6-bdd8-408e-8c31-5057644d7a4e>. Accessed: 2024-01-12.
- [231] S De Petris, F Sarvia, and E Borgogno-Mondino. A new index for assessing tree vigour decline based on sentinel-2 multitemporal data. application to tree failure risk management. 2021.
- [232] Hamed Alqahtani, Manolya Kavakli-Thorne, and Gulshan Kumar. Applications of generative adversarial networks (gans): An updated review. *Archives of Computational Methods in Engineering*, 28:525–552, 2021.
- [233] Min Wu, Xin Jin, Qian Jiang, Shin-Jye Lee, Lin Guo, Yide Di, Shanshan Huang, and Jinfang Huang. Remote sensing image colorization based on multiscale senet gan. In *2019 12th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI)*, pages 1–6. IEEE, 2019.
- [234] Feimo Li, Lei Ma, and Jian Cai. Multi-discriminator generative adversarial network for high resolution gray-scale satellite image colorization. In *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*, pages 3489–3492. IEEE, 2018.
- [235] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.
-

-
- [236] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [237] Tharun Mohandoss, Aditya Kulkarni, Daniel Northrup, Ernest Mwebaze, and Hamed Alemohammad. Generating synthetic multispectral satellite imagery from sentinel-2. *arXiv preprint arXiv:2012.03108*, 2020.
- [238] Svetlana Illarionova, Dmitrii Shadrin, Alexey Trekin, Vladimir Ignatiev, and Ivan Oseledets. Generation of the nir spectral band for satellite images with convolutional neural networks. *Sensors*, 21(16):5646, 2021.
- [239] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2337–2346, 2019.
- [240] Van Anh Le, Varshini Reddy, Zixi Chen, Mengyuan Li, Xinran Tang, Anthony Ortiz, Simone Fobi Nsutezo, and Caleb Robinson. Mask conditional synthetic satellite imagery. *arXiv preprint arXiv:2302.04305*, 2023.
- [241] Jordi Inglada, Julien Michel, and Olivier Hagolle. Assessment of the usefulness of spectral bands for the next generation of sentinel-2 satellites by reconstruction of missing bands. *Remote Sensing*, 14(10):2503, 2022.
- [242] Esa worldcover 2021. <https://worldcover2021.esa.int/>. Accessed: 2024-01-12.
- [243] Christian Albert, Mario Brillinger, Paulina Guerrero, Sarah Gottwald, Jennifer Henze, Stefan Schmidt, Edward Ott, and Barbara Schröter. Planning nature-based solutions: Principles, steps, and insights. *Ambio*, 50(8):1446–1461, October 2020.
- [244] Tom Bradfer-Lawrence, Tom Finch, Richard B. Bradbury, Graeme M. Buchanan, Andrew Midgley, and Rob H. Field. The potential contribution of terrestrial nature-based solutions to a national ‘net zero’ climate target. *Journal of Applied Ecology*, 58(11):2349–2360, September 2021.
- [245] Marion Le Texier, Kerry Schiel, and Geoffrey Caruso. The provision of urban green space and its accessibility: Spatial data effects in brussels. *PLOS ONE*, 13(10):e0204684, October 2018.
- [246] Gabriele Manoli, Simone Fatichi, Markus Schlöpfer, Kailiang Yu, Thomas W. Crowther, Naika Meili, Paolo Burlando, Gabriel G. Katul, and Elie Bou-Zeid. Magnitude of urban heat islands largely explained by climate and population. *Nature*, 573(7772):55–60, September 2019.
-

Author's publications

1. Giacco, G., S. Marrone, G. Langella, and C. Sansone, ReFuse: Generating Imperviousness Maps from Multi-Spectral Sentinel-2 Satellite Imagery, *Future Internet* 14, no. 10 (2022): 278., DOI: <https://doi.org/10.3390/fi14100278>
2. Pascarella, A. E., G. Giacco, M. Rigioli, S. Marrone, and C. Sansone, ReUse: REgressive Unet for Carbon Storage and Above-Ground Biomass Estimation, *Journal of Imaging* 9, no. 3 (2023): 61, DOI: <https://doi.org/10.3390/jimaging9030061>
3. Battisti, L., F. Aimar, G. Giacco, and M. Devecchi, Urban Green Development and Resilient Cities: A First Insight into Urban Forest Planning in Italy, *Sustainability*, 15(15), 12085., DOI: <https://doi.org/10.3390/su151512085>
4. D. Amitrano, G. Giacco, S. Marrone, A. E. Pascarella, M. Rigioli, C. Sansone, Forest Aboveground Biomass Estimation Using Machine Learning Ensembles: Active Learning Strategies for Model Transfer and Field Sampling Reduction, *Remote Sensing* 15, no. 21 (2023): 5138., DOI: <https://doi.org/10.3390/rs15215138>
5. L. Battisti, G. Giacco, M. Moraca, G. Pettenati, E. Dansero, F. Larcher, Spatializing Urban Forests as Nature-based Solutions: a methodological proposal *Cities*, 144 (2024): 104629., DOI: <https://doi.org/10.1016/j.cities.2023.104629>
6. G. Giacco, G. Mariniello, S. Marrone, D. Asprone, C. Sansone Toward a system for post-earthquake safety evaluation of masonry buildings *International Conference on Image Analysis and Processing*, Lecce, Italy, May 17 (pp. 312-323), Springer International Publishing, DOI: https://doi.org/10.1007/978-3-031-06430-2_26
7. A. Filograna, G. Giacco; G. Di Caprio Leveraging cloud-based geospatial data to enhance public services. A case study of the SPOTTED project

- 2023 3rd International Conference on Electrical, Computer, Communications and Mechatronics Engineering (ICECCME) Tenerife, Spain, July 19, IEEE, DOI: <https://doi.org/10.1109/ICECCME57830.2023.10252358>
8. Pascarella, A. E, G. Giacco, M. Rigioli, B. Vento, S. Marrone, G. Langella, A. Coppola, R. Chirone, P. Salatino, and C. Sansone, AI and Sustainability: Territorial Monitoring and Waste Valorization, Proc. of the Ital-IA 2023 Thematic Workshops co-located with the 3rd CINI National Lab AIIS Conference on Artificial Intelligence (Ital IA 2023), Pisa, Italy, May 29-30, 2023, CEUR-WS.org, online <https://ceur-ws.org/Vol-3486/143.pdf>
 9. L. Battisti, G. Giacco, M. Moraca, F. Cuomo, G. Pettenati, E. Dansero Servizi Ecosistemici, Aree verdi urbane e dati spaziali: una formula vincente per città resilienti ed ecologicamente attente, Memorie Geografiche Vol. XXII “Geografia e Tecnologia”, 2023, Firenze, Società di studi geografici, 2023, pp. 779 – 784
-