

Investigating the potential for mapping land
cover of the Caribbean Netherlands using
airborne LiDAR and passive optical data
Lisanne Sikam

Investigating the potential for mapping land cover of the Caribbean Netherlands using airborne LiDAR and passive optical data

Lisanne Sikam

Student Number: 4672453

Chair: Dr. R.C. Lindenbergh
Second Supervisor: Dr. ir. J. Timmermans
External Supervisor: Dr. F.S. Destra
Advisor: Dr. E. de Zeeuw-van Dalfsen

Project Duration: Jan, 2025 - Sept, 2025
Faculty: Faculty Of Civil Engineering and Geosciences, Delft
Discipline: Earth Observation

Abstract

Accurate and detailed land cover information is essential for effective biodiversity monitoring and conservation efforts. In the Caribbean Netherlands (consisting of Saba, St. Eustatius, and Bonaire), this need is particularly pressing, as the islands face growing ecological pressures from climate change, invasive species, and human activity. However, the islands' limited financial resources, steep terrain, and dense vegetation make systematic field surveys difficult, restricting the collection of consistent and detailed information about land cover. Remote sensing offers a powerful alternative, but previous satellite-based efforts for Saba were limited by their spatial and radiometric resolution, as well as the lack of vertical structural data, resulting in maps that lacked ecological specificity and taxonomic resolution. In early 2024, high-resolution airborne Light Detection and Ranging (LiDAR) and passive multispectral optical imagery were collected over the islands for the first time, providing a unique opportunity to assess their potential for land cover mapping. The objective of this study is therefore to *investigate the potential for mapping land cover of the Caribbean Netherlands using the recently acquired airborne LiDAR and passive optical data*. Specifically, this study investigates two main aspects, focused on Saba: (1) the quality of the LiDAR-derived Digital Terrain Model (DTM) released with the dataset, and (2) the combination of LiDAR structural parameters with optical remote sensing data for land cover mapping.

To assess the quality of the LiDAR-derived DTM, a novel data-driven reliability algorithm was developed. This algorithm combines several properties of the LiDAR data indicative of DTM reliability into a single pixel-based DTM reliability score R_{DTM} , which indicates how trustworthy the DTM is. Applying this method to Saba shows that 42% of the island has zero DTM reliability (no ground coverage at all), 9% low reliability, 16% moderate reliability, and 33% high reliability. Comparison of these reliability classes with vegetation indices confirmed that areas with dense vegetation tend to have lower DTM reliability.

In addition, several LiDAR-derived structural parameters were identified that describe both the topography and vertical structure on the surface, such as vegetation. These parameters were combined with products from passive multispectral optical satellite imagery to evaluate their potential for land cover mapping in two case studies. The first case study examined how LiDAR structural information can reveal variation within a single land cover class from a previous study, demonstrating that distinct vegetation structures can be distinguished within what was previously mapped as uniform forest. The second case study investigated whether structural parameters can explain the occurrence of the invasive vine *Coralita* (*Antigonon leptopus*), showing that its topographic occurrence was consistent with the literature.

The results demonstrate that the newly acquired airborne LiDAR dataset provides valuable structural information for land cover mapping. However, challenges remain in areas with dense vegetation and steep terrain, where limited ground penetration reduces DTM reliability. Furthermore, with the integration of field validation data (which was unavailable), the dataset's full potential for land cover mapping could be more fully realized.

Acknowledgments

I had the great fortune of not one, but two very dedicated supervisors, Roderik and Joris, whose enthusiasm during my thesis was truly contagious. Our weekly meetings often turned into brainstorming sessions about the endless possibilities of what could be done with the data, especially since it had never been used before. I often left meetings with new ideas that I wanted to try out. I'm very grateful for their feedback, guidance, and encouragement throughout.

I would also like to thank my external supervisor, Feven, and advisor Elske, who joined the important checkpoint meetings. Their more external perspective provided valuable reflections and helped steer the project in the right direction.

I want to thank the people I spoke with early on in my thesis who live and work on the islands and do such important work there. In particular, I want to thank Philip, who shared his insights on the practical challenges of limited resources on the islands and his perspective on the published dataset. Those conversations helped me understand the real-world context behind the data I was analyzing. Special thanks to Jochem and Lennard at Kadaster for helping me navigate the wonderfully complex world of coordinate reference systems. A standout moment was the meeting with Lennard, who was determined to make the manual implementation of the EPSG code work on my MacBook after I had spent a week (unsuccessfully) trying to resolve it myself. It was a fun debugging adventure, and while the original approach didn't quite pan out, we found an effective solution. Thank you to Adriaan for answering all my questions about the datasets, and to the team at Spheer AI for introducing me to their exciting platform. And a special mention to Akshit, a PHD students, who generously submitted an atmospheric correction for the Pleiades-Neo dataset so I could properly compare it to the Sentinel-2 data.

Finally, I want to thank my family and friends for their ongoing support and encouragement throughout this process. A special thanks goes to my boyfriend, who was also doing his thesis at the same time, so we could ride the thesis wave of ups and downs together. We swapped advice, did unprompted brain sessions, shared the chaos of doing a thesis, and proofread each other's theses when we read them one too many times ourselves.

All in all, I was pleasantly surprised by how much interest there was in this topic, and during the course of this work, I realized its broader importance. That realization motivated me to produce the best results I could. Thanks to the many people I interacted with along the way, the process felt far less lonely than I had expected.

*Lisanne Sikam
Delft, October 2025*

Contents

Abstract	iii
Preface	v
Nomenclature	xi
1 Introduction	1
1.1 Context and Motivation	1
1.1.1 Relevant Remote Sensing of Ecological states	1
1.2 Objective	2
2 Theoretical Background	3
2.1 Geography and Climate	3
2.2 Habitats	4
2.3 Biodiversity Threats	8
2.3.1 Climate Change	8
2.3.2 Roaming Livestock	9
2.3.3 Invasive Species	9
2.4 Advancing Conservation Monitoring	11
2.5 User Requirements and Land Cover Taxonomies	12
2.5.1 User Requirements	12
2.5.2 Land Cover Taxonomies for Saba	13
2.6 State of the Art	14
2.7 Research Gap	15
2.8 Research Questions	16
3 Data Description	17
3.1 Light Detection and Ranging (LiDAR)	17
3.1.1 Principles of LiDAR Remote Sensing	17
3.1.2 AHN LiDAR Data over Saba	19
3.2 Passive Optical Remote Sensing	21
3.2.1 Principles of Optical Remote Sensing	21
3.2.2 Beeldmateriaal NL Passive Multispectral ORS Data over Saba	22
3.2.3 Pleiades-Neo Passive Multispectral ORS Data over Saba	23
3.2.4 Sentinel-2 Passive Multispectral ORS Data over Saba	24
4 Methods	25
4.1 Qualitative AHN LiDAR Analysis	25
4.2 DTM Reliability Assessment	26
4.2.1 Identification of LiDAR Reliability Variables	26
4.2.2 A Priori Determined High DTM Reliability Group	31
4.2.3 Dimensionality Reduction	31
4.2.4 Variable Scoring	32
4.2.5 Vegetation Cover Influences on DTM Reliability	34
4.3 Assessing the Usability of LiDAR-Derived Structural Information for Land Cover Mapping	34
4.3.1 Vertical Structure Extraction	35
4.3.2 Case Study Analysis	36
5 Results	39
5.1 Qualitative AHN LiDAR Analysis	39
5.2 DTM Reliability Assessment	43
5.2.1 Identification of LiDAR Reliability Variables	44

5.2.2	A Priori Determined High DTM Reliability Group	45
5.2.3	Dimensionality Reduction	46
5.2.4	Variable Scoring	47
5.2.5	Vegetation Cover Influences on DTM Reliability	48
5.3	Assessing the Usability of LiDAR-Derived Structural Information for Land Cover Mapping	49
5.3.1	Vertical Structure Extraction	49
5.3.2	Case Study Analysis	52
6	Discussion	57
6.1	Potential of LiDAR information for Land Cover Mapping	57
6.1.1	The Reliability of the DTM	57
6.1.2	Case Study A: Level-2 Ecological Taxonomy	58
6.1.3	Case Study B: Level-3 Ecological Taxonomy	58
6.2	Workflow Design Impacts	59
6.2.1	Reflection on the DTM Quality Assessment	59
6.2.2	Reflection on the Vertical Structure Extraction	63
6.3	Societal Impacts of the Study	64
7	Conclusion and Recommendations	67
7.1	Conclusion	67
7.2	Recommendations	70
References		73
A	Flora of Saba	79
B	Invasive Species on Saba	83
C	Remote Sensing Techniques	85
C.1	LiDAR Remote Sensing	85
C.1.1	LiDAR System Components	85
C.1.2	LiDAR Resolution	85
C.1.3	LiDAR Data Quality	85
C.1.4	Extended LiDAR Applications	86
C.2	Passive Optical Remote Sensing	86
C.2.1	DN to Reflectance Conversion	86
C.2.2	Passive ORS Sensor Components	86
C.2.3	Passive ORS Resolution	87
C.2.4	Passive ORS Data Quality	87
C.2.5	Extended Passive ORS Applications	87
D	AHN and Beeldmateriaal Nederland Workflow and Accuracy Assessment for Saba	89
D.1	AHN LiDAR Data	89
D.1.1	RIEGL VQ-1560 II System Specifications	89
D.1.2	LiDAR Flight Execution	89
D.1.3	LiDAR Data Processing and Product Generation	91
D.1.4	LiDAR Data Quality	91
D.2	Beeldmateriaal Nederland Passive Multispectral ORS	92
D.2.1	UltraCam Eagle 4.1 Digital Aerial Mapping System Overview	92
D.2.2	Passive Multispectral ORS Data Processing and Product Generation	93
D.2.3	Passive Multispectral ORS Data Quality	93
E	Saba's Local Coordinate Reference System (CRS)	95
F	Cloud and Shadow Identification in the Pleiades-Neo Data	97
F.1	Calculation in Practice	97
G	Beeldmateriaal Nederland Passive Multispectral ORS Data Properties	99
H	Variable Distributions in the DTM Reliability Assessment	101
I	Scoring Functions R(x)	103

J GMM BIC scores	105
K GMM Fit of Normalized Height Distributions	107
L Local Incidence Angle	109
L.1 Definition	109
L.2 Calculation in Practice	109
M Dimensionality Reduction of the LiDAR Reliability Variables	111

Nomenclature

Abbreviations

Abbreviation	Definition
AHN	Actueel Hoogtebestand Nederland
ASPRS	American Society for Photogrammetry & Remote Sensing
BIC	Bayesian Information Criterion
CBD	Convention on Biological Diversity
CHM	Canopy Height Model
DN	Digital Numbers
DSM	Digital Surface Model
DTM	Digital Terrain Model
ECDF	Empirical Cumulative Distribution Function
GMM	Gaussian Mixture Model
GNSS	Global Navigation Satellite System
GPC	Ground Control Point
IUCN	International Union for Conservation of Nature
LAI	Leaf Area Index
LiDAR	Light Detection and Ranging
NDVI	Normalized Difference Vegetation Index
NIR	Near-Infrared
NSO	Netherlands Space Office
ORS	Optical Remote Sensing
PCA	Principal Component Analysis
PDF	Probability Distribution Function
PRF	Pulse Repetition Frequency
RAM	Random Access Memory
SIDS	Small Island States
TOF	Time of Flight

Introduction

1.1. Context and Motivation

The Caribbean Netherlands, consisting of Saba, St. Eustatius, and Bonaire (Figure 1.1), is highly vulnerable to the impacts of climate change. Recognized as Small Island States (SIDS) (United Nations Conference on Environment and Development, 1992), these islands have relatively large coastal areas and very small populations, which causes severe societal impacts from sea-level-rise, floods, heavy rainfall, and more frequent and severe tropical storms (A.O. Debrot, Henkens, and Verweij, 2018; KNMI, 2023). These climate-driven phenomena pose significant threats to the islands' often unique and rich biodiversity (Cherian, 2007). Indeed, most habitats and species in the Caribbean Netherlands are in a critical state with a high likelihood of further decline because of climate change and invasive species (A.O. Debrot, Henkens, and Verweij, 2018; KNMI, 2023). To combat these threats, adequate management and sufficient monitoring are required.

This stands in contrast to the obligations of the Netherlands to the special municipalities of Bonaire, Saba, and St. Eustatius (as of October 2010) for implementing, monitoring, and reporting of their ecological states. Per ratification of the Convention on Biological Diversity (CBD), the Netherlands commits to the CBD's 3 main objectives: conserving biodiversity, the sustainable use of biodiversity, and fairly sharing benefits from genetic resources (Convention on Biological Diversity, 2024). In response, the Nature and Environment Policy Plan Caribbean Netherlands 2020-2030 (NEPP) (Ministry of Agriculture, Nature and Food Quality, 2020) has been created to meet these obligations for a 10-year horizon. They include strategies to combat the ecological threats and assess the effectiveness of policy measures using a six-year nature report (A.O. Debrot, Henkens, and Verweij, 2018; A. Debrot et al., 2025).

Addressing these challenges requires a detailed and accurate understanding of the current land cover across the Caribbean Netherlands. Such information is critical for developing evidence-based management plans, implementing effective conservation strategies, and tracking ecological changes over time. However, ground-based measurements to establish such land cover maps at the required temporal frequency are complicated in the Caribbean Netherlands. Reasons for this are financial limitations, the dense vegetation on the island complicating estimations, and the spatial heterogeneity of nature on the island. Remote sensing techniques (including satellite and airborne observations) provide a means to circumvent these challenges (Timmermans and Daniel Kissling, 2023). Specifically, they offer frequent optical and radar observations with temporal resolutions ranging from days to weeks and spatial resolutions spanning a few meters to several kilometers (Saritha et al., 2025), which are relevant for monitoring ecological states.

1.1.1. Relevant Remote Sensing of Ecological states

Satellite data (from, for example, passive multispectral optical sensors) allowed land cover maps to be made available, such as the most recent attempt by Smith et al., 2013 for Saba. However, in their research, the developed landcover maps were limited in their ecological taxonomies. For example, they could only classify the landcover in broad classes, such as broadleaved evergreen forest, which covers mostly the entire island. Although their work provided a valuable first baseline, it is insufficient for adequate monitoring of the ecological state of the various habitats, as well as mapping the threats

(such as alien invasive species) to those habitats. Moreover, given the 2 meter spatial resolution and 8 spectral bands of these observations, the data lacked the spatial and radiometric detail needed to distinguish heterogeneous habitats or identify individual species. Additionally, the absence of vertical structural information at the time, combined with the aforementioned limitations, reduced the ecological specificity and taxonomic resolution of the resulting maps. To address these limitations, recent sensor technologies have emerged that provide higher resolution and richer information. One of the most important of these is LiDAR (Light Detection and Ranging), an active remote sensing technique that emits laser pulses and measures the time it takes for the signal to return. This enables the creation of detailed three-dimensional point clouds, from which structural information such as canopy height, vegetation density, and ground elevation can be derived. Unlike passive optical sensors, LiDAR is not limited to surface reflectance but provides vertical information about the landscape, making it especially valuable for applications where vertical information plays an important role.

In the beginning of 2024, Airborne LiDAR and high-resolution Passive Multispectral Optical Remote Sensing (ORS) data was acquired for the first time over the islands. This offers a high potential for mapping and monitoring land cover of the Caribbean Netherlands, which will support long-term efforts to mitigate biodiversity loss.



Figure 1.1: The Caribbean Netherlands. Source: Thayts, 2011

1.2. Objective

The objective of this study is to investigate the potential for mapping land cover of the Caribbean Netherlands using the recently acquired airborne LiDAR and passive optical data. The LiDAR data is provided by the Actueel Hoogtebestand Nederland (AHN), and the passive optical imagery is provided by Beeldmateriaal Nederland. In this context, “potential” refers to the extent to which these datasets can address limitations identified in prior research. The first sub-question, therefore, requires the definition of the user requirements: *What land cover taxonomies are relevant for biodiversity monitoring on Saba?* This sub-question is answered in Section 2.5. The rest of the sub-questions are defined in Section 2.8.

2

Theoretical Background

In this chapter, the theoretical background is provided, including a description of the Geography and Climate and Habitats of the study area Saba, as well as the Biodiversity Threats to this island. In Section 2.4, advancing conservation biodiversity monitoring methods are discussed, and in Section 2.5, the user requirements and land cover taxonomies are identified in order to monitor the island's biodiversity. In Section 2.6, the state of the art is presented, and in Section 2.7, the research gap is identified.

2.1. Geography and Climate

Saba (Figure 2.1) is the smallest island of the Caribbean Netherlands, covering about 13 km². Located in the northeastern Caribbean, it shares a lot of climatic and geographic similarities with St. Eustatius, though different because of the dominating Mount Scenery (*Volcanoes in the Dutch Caribbean* n.d.), an active but currently quiescent stratovolcano that reaches an elevation of 870 m above sea level. Like many volcanic islands in the Caribbean, its peak is often capped by clouds (Freitas et al., 2015), creating conditions for cloud and rain forests near the summit. These high-altitude forests play an important ecological role by capturing moisture, supplying freshwater to lower-lying areas, reducing landslide risk, and limiting soil erosion (A.O. Debrot, Henkens, and Verweij, 2018). Additionally, the presence of this volcano, characterized by its conical shape and steep elevation gradients, supports a wide variety of plant species through variations in temperature, sunlight, precipitation, and wind exposure. As a result, the island is highly heterogeneous in terms of vegetation.

The island is strongly influenced by persistent northeast trade winds, with a dominant eastern wind direction (Table 2.1). Saba's location within the Atlantic hurricane belt also exposes it to extreme weather



Figure 2.1: Saba. Source: Google Earth, 2025.

Table 2.1: Summary of meteorological data of Saba in 2024. The dry season ran from December 2023 to April 2024, and the wet season from May to November 2024. Data are averages from hourly KNMI observations (Koninklijk Nederlands Meteorologisch Instituut, n.d.).

Weather parameter (2024)	Value
Mean wind direction	100.1° (E)
Mean temperature – dry season	26.7°C
Mean temperature – wet season	29.1°C
Mean humidity – dry season	73.2%
Mean humidity – wet season	74.6%
Mean precipitation – dry season	68.5 mm/month
Mean precipitation – wet season	110.8 mm/month

events. These events pose a risk to the island's forest ecosystems through storm damage, landslides,

and forest destruction.

2.2. Habitats

The weather conditions in Table 2.1 vary significantly with elevation, slope orientation, and exposure to wind and cloud cover. Therefore, despite its small size, Saba is capable of supporting a large range of biodiversity. A 2019 survey documented approximately 772 vascular plant species (Press, 2021), of which a large number of endemic plant species (A. Debrot et al., 2025). The island also hosts diverse (endemic) animal species (A.O. Debrot, Henkens, and Verweij, 2018). The plant and animal species are divided among the following habitats identified on Saba: Elfin forest, Montane forest, dry tropical forests, dry shrubland and grassland, caves, beaches, bare rocks, seagrass and seaweed beds, coral reefs, and open ocean and deep sea (A. Debrot et al., 2025). Figure 2.2 shows the estimated distribution of these habitats as published in the 2018 state of nature report (A.O. Debrot, Henkens, and Verweij, 2018).

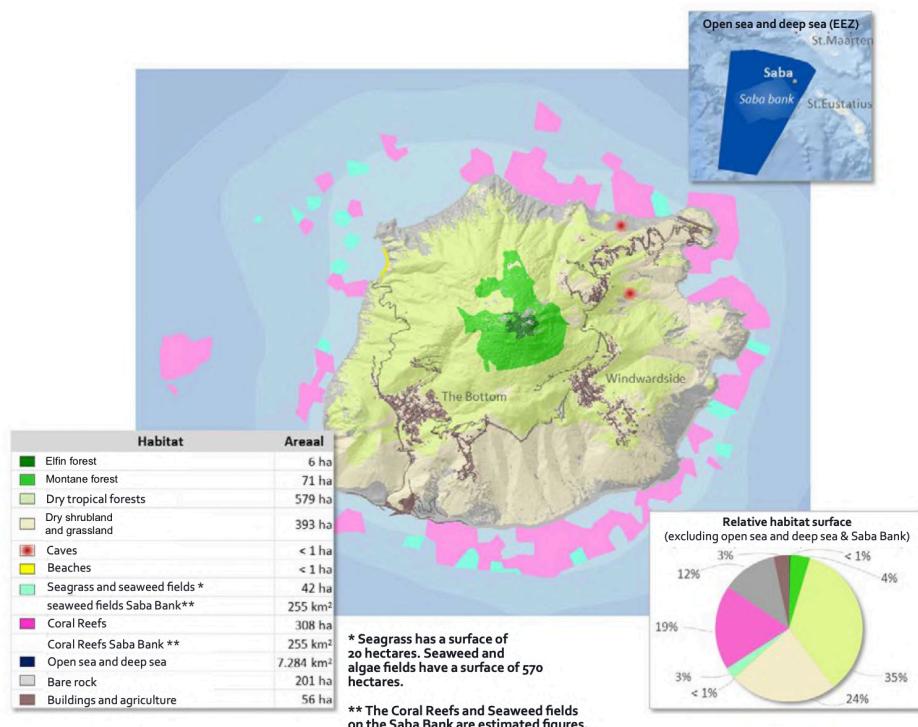


Figure 2.2: Estimation of spatial distribution of terrestrial and marine habitats on Saba as published in the 2018 state of nature report (A.O. Debrot, Henkens, and Verweij, 2018). The map is derived from Freitas et al., 2015 and Smith et al., 2013.

The main terrestrial habitats from high to low elevation are described below.

Elfin Forest

The Elfin forest, also known as a cloud forest, is located at the peak of Mt. Scenery. It starts around 750 m elevation extending to the top of the volcano, representing the smallest habitat of Saba's forest habitats. A representative view of this habitat is provided in Figure 2.3a. As mentioned earlier, the peak of the volcano is almost constantly capped in clouds (Figure 2.3b), which condense on foliage and branches, resulting in consistently high humidity levels. This persistent mist, combined with frequent rainfall, creates a cool, humid, and shaded environment that supports a lush understory with dense moss growth, abundant ferns and orchids, and a diverse community of epiphytes (plants growing on other plants) clinging to trunks and branches (A.O. Debrot, Henkens, and Verweij, 2018), as illustrated in Figure 2.3c. The dominant tree species that can be found in this habitat is the Lesser Antillean endemic *Freziera undulata*.



(a) Saba's Elfin forest. Source: Laughlin, 2009c

(b) Cloud formation at the peak of Mt. Scenery. Source: Gleason, 2011

(c) A branch in the Elfin forest heavily covered in moss and epiphytes. Source: Laughlin, 2009b

Figure 2.3: Examples of Flora Found in Saba's Elfin Forest.

This environment is unique compared to other Elfin forests in the Caribbean because this forest occurs at a relatively low elevation. This allows the *Freziera undulata* to grow up to 15 m, which is much taller compared to other Caribbean islands, where it typically reaches only about 6 m (A.O. Debrot, Henkens, and Verweij, 2018). In 1998, however, Hurricane Georges struck the island, causing significant damage to this habitat through canopy loss and a decrease in vegetation height. As a result, *Freziera undulata* endured a temporary decline in their dominance (Freitas et al., 2015). In the years since, they reach heights of around 9 m (A. Debrot et al., 2025).

Beyond its structural uniqueness, this cloud forest acts as a major water catchment and is also critical for biodiversity, as a lot of endemic plants can be found here, such as *Begonia retusa* and *Chromolaena macrantha* (Freitas et al., 2015), and it serves as nesting sites for seven restricted range bird species (United Nations Environment Program, 2018).

Montane Forest

Montane forests, also known as rainforests, occupy the middle and upper slopes of Mt. Scenery. It is characterized by high humidity, and can be found between 500 and 750 meters elevation (A. Debrot et al., 2025; United Nations Environment Program, 2018). This habitat surrounds the volcano, exposed to different wind conditions. For example, parts of the Montane forest that are sheltered from the wind (on the west side) are the most developed, as the forest is taller and can grow quite complex layers undisturbed by wind. Figure 2.4a shows the characteristic dense vegetation and lush understory of this habitat. In contrast, on the eastern side of the Montane forest, which is more exposed to wind, the vegetation is shorter and more compact (A. Debrot et al., 2025). Because the Montane forest borders the Elfin forest at higher elevations and the dry tropical forest at lower elevations, it shares many plant species and ecological traits with both. At higher elevations, for example, the Montane forest is characterized by epiphytic flora, although much less than in the Elfin forest.

Combined with the large elevation gradient, the Montane forest is the most diverse in plant species (United Nations Environment Program, 2018), and contains the tallest tree species on the island. Just

like the Elfin forest, the ample precipitation in this habitat is sufficient to sustain an evergreen tropical rainforest, meaning that the plants keep their leaves throughout the year (A.O. Debrot, Henkens, and Verweij, 2018). A typical tree species in the Montane forest is *Prestoea acuminata* var. *montana* (Figure 2.4b), while the endangered *Nectandra krugii* (Dutch Caribbean Species Register, 2018) also occurs in this habitat. Species part of the genus *philendron* such as the *Philendron giganteum* Figure 2.4c are common in the understory and often grow over other plants as hemiepiphytes. Additionally, several types of orchids, ferns and mosses thrive in this forest. The Montane forest also supports important fauna such as the endemic *Anolis sabanus* lizard and endangered *Alsophis rufiventris* racer snake (United Nations Environment Program, 2018).



(a) Image from the Sandy Cruz trail that runs through the Montane forest. Source: world-wide-mike, 2021

(b) *Prestoea acuminata* var. *montana* (Dutch Caribbean Species Register, n.d.[b]). Source: Observation.org, 2015

(c) *Philendron giganteum* (Observation.org, 2025c). Source: Dutch Caribbean Species Register, n.d.(a)

Figure 2.4: Examples of Flora Found in Saba's Montane Forest.

Dry Tropical Forest

Dry tropical forests are found at elevations roughly between 200 and 500 meters (A. Debrot et al., 2025). Similar to the Montane forest, the distribution and types of plant species in this habitat are influenced by environmental factors. This habitat consists of a mixture of evergreen and deciduous plant species. Deciduous species shed their leaves during the dry season and naturally occur at lower, drier elevations within the habitat. A view of this habitat at moderate elevation is shown in Figure 2.5a.

In areas where deciduous vegetation dominates, dense ground vegetation is present due to increased light availability (A. Debrot et al., 2025). Extending from the border with the Montane forest, many plant species found in this habitat also occur in the Montane forest, although their abundance is significantly lower (Freitas et al., 2015). In this region, there is a mix of tall and low trees, with ferns and shrubs in the understory. In the western part of this habitat, where slopes are gentle, the area is dominated by the endangered deciduous tree species *Swietenia mahagoni* (Freitas et al., 2015), which can be seen in Figure 2.5b. The northern region of this habitat is more exposed to wind, resulting in reduced vegetation cover and lower species diversity. This area consists mostly of the dominant evergreen tree species *Coccoloba swartzii* and *Guettarda scabra*, although some deciduous tree species can also be found here. The largest expanse of this habitat stretches from the northeast to the southwest of the island, an area that is most exposed to wind and includes degraded forest patches dominated by grasses and shrubs. A typical deciduous tree species is the *Bursera simaruba* (Figure 2.5c) and the endangered *Guaiacum officinale*. This habitat also plays an important role in the control of erosion at lower elevations.



(a) View of the dry tropical forest in the eastern part of this habitat at moderate elevation.
Source: Paul Illsley, 2023

(b) *Swietenia mahagoni*. Source: Observation.org, 2023a

(c) *Bursera simaruba*. Source: Observation.org, 2023b

Figure 2.5: Examples of Flora Found in Saba's Dry Tropical Forest.

Dry Shrubland and Grassland

Dry shrubland and grassland are primarily located on the lowest slopes of Saba, stretching from the northeast to the southwest and along the coastal areas. This habitat is dominated by open, low-lying vegetation, as shown in Figure 2.6. The shrubs on the lower slopes are ecologically important, as they



Figure 2.6: Saba's dry shrubland and grassland landscape. Source: Laughlin, 2009a

help stabilize the soil and reduce erosion (United Nations Environment Program, 2018). Besides having naturally occurring vegetation (referred to as primary vegetation), the habitat also includes degraded dry tropical forests, known as secondary vegetation. Primary vegetation is typically found on windward ridges, where shallow soils and strong winds limit tree growth. Under these conditions, species richness is lower compared to the higher-elevation habitats (A. Debrot et al., 2025). Typical primary vegetation species include the grass *Aristida adscensionis* (Figure 2.7a) and tree *Eugenia axillaris* (Figure 2.7b). Secondary vegetation often contains invasive species, such as *Bothriochloa pertusa* (Figure 2.7c) (United Nations Environment Program, 2018). In addition to grasses and shrubs, scattered rocks and patches of moss are also common features of this habitat.

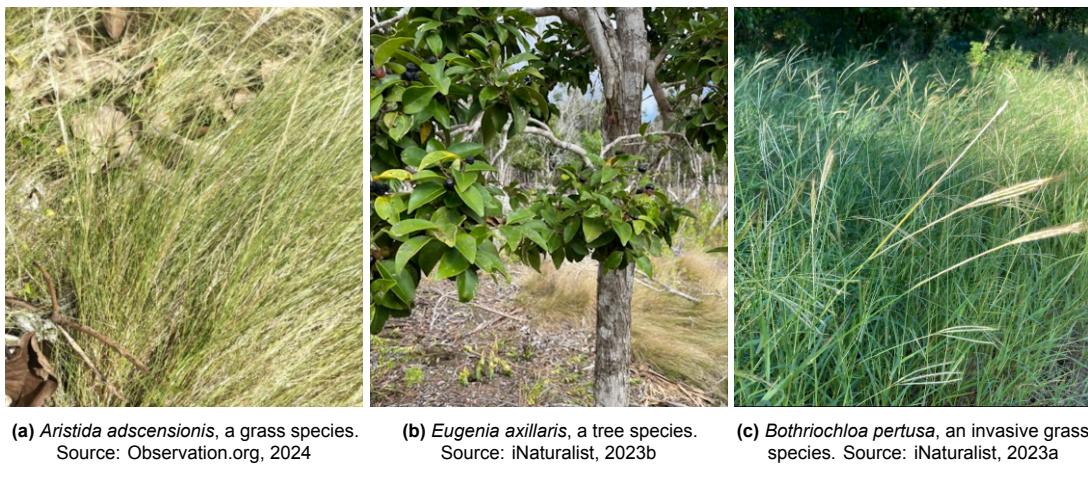


Figure 2.7: Examples of Flora Found in Saba's Dry Shrubland and Grassland.

A more extensive overview of flora found in these habitats can be found in Appendix A.

2.3. Biodiversity Threats

The biodiversity of Saba, as well as that of the wider Caribbean Netherlands, faces serious pressure (IPCC, 2019). Across Saba's identified habitats, numerous small and fragmented ecosystems exist, making them highly vulnerable to disturbances that can lead to habitat loss and the extinction of endemic plant and animal species. The three most significant threats to Saba's biodiversity are climate change (KNMI, 2023), roaming livestock (A. Debrot et al., 2025), and invasive species (Government of the Netherlands, 2025). These threats are often interconnected and intensify each other.

2.3.1. Climate Change

In 2024, the island experienced its highest recorded year-average temperature of 29 °C (KNMI, 2025). According to climate projections for the year 2050 (see Figure 2.8), even under the most optimistic scenario, temperatures are expected to rise by an additional 0.8 °C, precipitation during the dry season may decrease by 9.12 mm, and sea levels are projected to rise by 13 to 32 cm (KNMI, 2023).

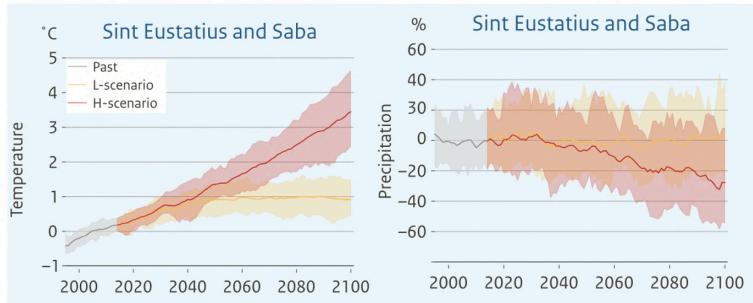


Figure 2.8: Projected changes in annual mean temperature and precipitation for St. Eustatius and Saba under two climate scenarios (L-scenario: low emissions, H-scenario: high emissions). Shaded areas indicate the uncertainty ranges around the scenario means. Source: KNMI, 2023.

Likewise, hurricanes and storms will become more frequent and intense (KNMI, 2023). Rainfall patterns are projected to shift toward prolonged dry periods, alternated by intense rainfall concentrated in shorter periods of time. This intense rainfall can overwhelm the island's capacity for water retention, leading to surface runoff and increased soil erosion. In addition, rising temperatures increase evaporation rates, which reduces soil moisture even in wetter habitats. Warmer conditions may also raise the cloud base, decreasing the frequency of mist immersion and gradually drying out the forests. These

climate change factors have an especially strong impact on the Elfin and Montane forests, which heavily depend on moist conditions. Without adequate moisture, the Elfin forest risks extinction (A. Debrot et al., 2025). The destructive potential of hurricanes on Saba has already been demonstrated. As previously mentioned, Hurricane Georges in 1998 severely damaged the Elfin forest, altering its structure and reducing species richness. Similar impacts have been observed in the Montane and dry tropical forests, where hurricane winds create light gaps that favor the spread of invasive plant species. In the dry grassland and shrubland, invasive grasses are typically the first species to re-establish themselves after a hurricane has cleared an area (A. Debrot et al., 2025). If hurricanes and tropical storms continue to intensify and occur more frequently as projected, these habitats may experience repeated damage before they can fully recover, leading to changes in species composition, reduced biodiversity, and diminished ecosystem services such as water catchment and soil stabilization.

Although SIDS like Saba contribute only a negligible share of global greenhouse gas emissions, they are among the most exposed to the impacts of global climate change (IPCC, 2019). Beyond ecological stress, climate change threatens infrastructure and economic instability. These effects can be worsened by poor land management. Unlike the European Netherlands, which possesses greater adaptive capacity due to its size and resources, Saba is far less prepared to withstand these pressures. Therefore, the island's most effective defenses against these threats are biodiversity conservation efforts and sound island management plans, requiring the need of clear legal and policy frameworks.

2.3.2. Roaming Livestock

Roaming livestock has been a long-time problem on Saba. In recent decades, insufficient control has allowed goats and other animals to enter the wild and reproduce freely. Goats pose the largest threat, as they are highly adaptable and reproduce quickly. Additionally, chickens, cats, non-native iguanas, and rats are also problematic. Goats threaten biodiversity by eating young seedlings and trees, preventing natural forest regeneration (A. Debrot et al., 2025). Their trampling compacts the soil, accelerates erosion, and leaves bare ground that becomes a breeding ground for invasive plant species to take hold. Compounding the problem, goats often avoid eating invasive plants, thereby giving the invasive plants a competitive advantage. Goats have been recorded at all elevations on the island, making every habitat susceptible to their impacts, though the effects of overgrazing are most pronounced on the lower slopes (A. Debrot et al., 2025). In these habitats, including the dry tropical forests and dry grassland and shrubland, roaming livestock is the main driving factor of biodiversity loss. In the Elfin and Montane forests, a significant decrease in understory species-richness has been observed since 1950, partly due to reduced light from canopy closure, but also as a result of sustained overgrazing.

To address the problem, Saba has implemented a livestock control project with measures such as public awareness campaigns, placement of fences, and hunting programs targeting free-roaming goats (Rijksdienst Caribisch Nederland, 2020). It is estimated that as a result of this campaign, approximately 90% of roaming livestock has been reduced since 2020 (Public Entity Saba, 2025). Simultaneously, a reforestation project was also started, focusing on erosion-prone areas. However, the time between reforestation and now is too short to see the impact of the removal of roaming livestock, as the replanted trees are still at a size vulnerable to goats. Despite the initial success in removing a large part of the roaming livestock, without consistent monitoring and action, goat populations can rebound quickly, reversing progress.

2.3.3. Invasive Species

The introduction and spreading of (alien) invasive species is slowly becoming the biggest threat to the functioning of Saba's ecosystems and habitats (A. Debrot et al., 2025). Invasive species are species non-native to the natural ecosystem and whose introduction (likely) causes economic or environmental harm or harm to human health (Government of the Netherlands, 2025). In contrast to other exotic species, these species have specific advantages over native plants and/or animals that enable them to rapidly expand and destabilize the habitats, thereby becoming an ecological threat. In particular, alien species may exhibit competitive traits that allow them to acquire nutrients or reproduce more efficiently, or they may be opportunistic, exploiting ecosystem disturbances such as forest fires or other environmental changes.

Once established, these species grow rapidly. At first, they may coexist with native vegetation without

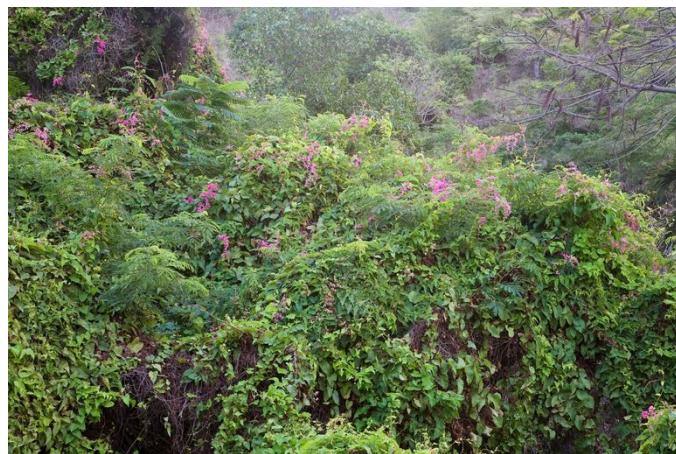
causing significant harm, but as their spread accelerates and becomes unmanageable, they invade natural habitats and displace native plants, causing biodiversity loss. Once the biodiversity declines, invasive species have more opportunities to expand, reinforcing a continuous feedback loop.

A 2012 study identified 41 exotic plant species on Saba that were in various stages of establishment, of which 14 were already invasive (van der Burg et al., 2012). A more recent update (provided in Appendix B) given by the Global Register of Introduced and Invasive Species (GRIIS), reports an increase to 46 invasive plant species (Debrot et al., 2025). These species are particularly threatening to the lower elevation areas where they can exploit disturbances caused by hurricanes (A. Debrot et al., 2025). For example, the invasive grass species *Bothriochloa pertusa* (Figure 2.7c), commonly known as "hurricane grass", tipped the ecosystem stability of some windward areas of the dry grassland and shrublands by becoming the dominant species by outcompeting native species and altering the open structure of the habitat (A. Debrot et al., 2025). Likewise, a naturalized plant species called *Cryptostegia madagascariensis* is also becoming dominant in parts of the dry grassland and shrubland by overgrowing native vegetation, contributing to one-third of the observed species loss (A. Debrot et al., 2025). These invasions are not limited however to only the lower and drier zones of Saba. In the Elfin and Montane forests, hurricanes and intense storms create canopy gaps that can be exploited by invasive species. For example, after Hurricane Georges in 1998, the increased light penetration in the Elfin forest facilitated the spread of *Rubus rosifolius* that now has a naturalized status (Dutch Caribbean Species Register, n.d.[c]). Even though the wetter conditions of these habitats slow the spread of invasive species, their expansion into these habitats is concerning because of the high proportion of endemic species present.

***Antigonon leptopus*, "Coralita"**

Of all these alien invasive species, *Bothriochloa pertusa* and *Antigonon leptopus* have been identified as the most problematic invasive species in Saba van der Burg et al., 2012. Not only is *Antigonon leptopus* considered a threat on more Dutch Caribbean islands, as mentioned on the current research and monitoring list from the Dutch Caribbean Nature Alliance (DCNA) (Dutch Caribbean Nature Alliance (DCNA), n.d.), it also has specific traits (such as being drought-resistant, allowing it to fully utilize the climate change on the islands), that give it a higher priority than other alien invasive species.

Antigonon leptopus, or commonly called Coralita, is a climbing vine and can be seen in Figure 2.9. This species was first introduced as an ornamental plant, but very quickly gained an invasive status. By climbing, it smothers other plants (even trees) and creates different lighting conditions, thereby outcompeting native vegetation (Figure 2.9a). Although climbing is Coralita's main way of spreading, it also spreads through its seeds and underground via its roots (Burke and DiTommaso, 2011).



(a) Coralita climbing and smothering vegetation. Source: Achsah Mitchell, 2023



(b) Coralita up close with its characteristic pink flowers. Source: Observation.org, 2025a

Figure 2.9: *Antigonon leptopus*, "Coralita": An invasive species on Saba.

The preferred conditions of Coralita are a dry to moist lowland environment (< 600 m), and it favors limestone as soil, although it can grow on any soil as long as it is well drained (Ernst and Ketner, 2007).

Coralita occurs across a broad range of soil pH values, but field data from Saba suggest a preference for moderately acidic soils, with the highest occurrence at a 5.0–5.5 pH, though the evidence remains inconclusive due to the small sample size (Ernst and Ketner, 2007). Importantly, Coralita also tolerates high pH soils, indicating that soil pH is not the sole driver of its distribution. Coralita is also drought-tolerant and grows in full sun to partial shade, but prefers full sun exposure. Hurricanes are notable events during and after which the plant rapidly spreads: during a storm, dispersal occurs through wind and runoff, while disturbed open ground is created for the plant to establish. Rainfall also influences its growth, as drought almost halts its expansion and flowering (but it can survive through defoliation), while wet periods cause extreme growth (Ernst and Ketner, 2007). Coralita favors disturbed areas such as roadsides, and abandoned land such as gardens or agricultural areas. From this established state, Coralita originally spread further into natural vegetation. On St. Eustatius, abandoned agricultural areas have been the main reason for rapid spread. On Saba, however, abandoned land is much less present, which limits its spread (A. Debrot et al., 2025). Here, Coralita primarily invades roadsides, areas disturbed by roaming livestock, erosion-affected sites (van der Burg et al., 2012).

When supported by vegetation or man-made structures, individual vines commonly extend 2.4 - 3 m in a single season and can achieve total lengths of 9 – 12 m (Ernst and Ketner, 2007; Gardenia.net, n.d.). Without support, Coralita spreads horizontally rather than climbing, and its vertical growth is minimal. The plant also has leaves about 2.5 – 7.5 cm long (Ernst and Ketner, 2007). Their underground stems, known as tubers, can reach lengths of up to 2 m (Ernst and Ketner, 2007). In heavily infested areas, densities of up to 280 tubers per m² have been recorded (A. Debrot et al., 2025), making removal extremely difficult. Depending on the water availability, the plant either has a distinct or continuous flowering period. The flowers are bright pink (occasionally white) and small, as seen in Figure 2.9b. On Saba, it is more likely that there is a continuous flowering period, although dry periods do occur in which fruiting is more prominent (Ernst and Ketner, 2007).

2.4. Advancing Conservation Monitoring

Given the above, frequent and accurate monitoring of Saba's biodiversity is crucial for tracking the effectiveness of conservation policies. On Saba, this is primarily accomplished through field surveys (for example, in Stoffers, 1956) with the aid of expert knowledge, as well as the compilation of inventories (for example, in Boeken, 2014; A.O. Debrot, Henkens, and Verweij, 2018). These methods have several disadvantages. Firstly, data collection has often been project-based, resulting in irregular time intervals between surveys, which is unfavorable for trend analysis. For example, the state of nature reports state that they are not sure about the circumference of habitats because they are using outdated habitat maps, and they draw comparisons between a vegetation map from 1950 to now when assessing trends in habitat distribution and area (A.O. Debrot, Henkens, and Verweij, 2018; A. Debrot et al., 2025). This makes it challenging to monitor changes that occur on a much smaller timescale, such as the spread of invasive species, which may persist at low densities for years before rapidly expanding once conditions allow (Crooks, 2005). Additionally, field surveys are labor-intensive, inconsistent across independent field surveys, and incomplete due to inaccessible areas, especially in higher-elevation habitats.

More recently, remote sensing techniques have been integrated on islands to solve limitations of relying solely on field surveys (Saritha et al., 2025). At the most basic level, satellite and aerial imagery can supplement field data to produce repeatable land cover maps. Satellite data allows monitoring over consistent and shorter time intervals, which is useful for detecting trends such as habitat circumference. Aerial imagery provides higher spatial resolution data, allowing the creation of more detailed maps. An example for Saba specifically, where aerial imagery serves as a basis for field sampling, is a research by Freitas et al., 2015 where they combined aerial imagery with field surveys to produce an ecological vegetation map. This research employs a method that combines remote sensing data and field data. A 2020 study, which placed less emphasis on field data, utilized satellite imagery of Aruba, combined with historical records and expert knowledge, to assess land cover changes from 1900 to 2020 and to produce a 1 m resolution land cover map for the year 2020 (Mucher et al., 2024). With these maps, they found a reduction in certain natural habitats driven by urban expansion, as well as the appearance of vegetation in previously bare areas, likely due to the removal of goats. To monitor fine-scale ecological processes via remote sensing, drones equipped with multispectral imaging capabilities can be deployed.

In St. Eustatius, for instance, such drones are being utilized to assess vegetation health and evaluate the impact of fencing interventions designed to exclude goats from specific areas.

2.5. User Requirements and Land Cover Taxonomies

To effectively monitor biodiversity, there must be a clear alignment between the conservation priorities defined by the NEPP and the data required to address them. As mentioned in Section 2.4, remote sensing techniques could offer a powerful tool to support biodiversity monitoring on Saba. To make this operational, the monitoring needs identified in the state of nature reports and by local organizations such as DCNA (Dutch Caribbean Nature Alliance (DCNA), n.d.) must be translated into remote sensing products. Therefore, user requirements and land cover taxonomies need to be defined to ensure that the products meet conservation objectives.

2.5.1. User Requirements

The monitoring priorities for Saba are summarized as: understanding the changing distribution, size, and condition of habitats and species (with a focus on endangered, key, and indicator species) as a result of the biodiversity threats mentioned in Section 2.3, as well as monitoring the effectiveness of mitigation efforts. The species consists of both flora and fauna, but the conditions of fauna depend directly on the conditions of flora. For example, as mentioned in Section 2.2, several bird species depend on the existence of the Elfin forest, as they are restricted-range species. Because of this dependency, remote sensing focused on flora can also serve as an indirect indicator of fauna health. Therefore, the focus here is on using remote sensing techniques to monitor flora. To meet the monitoring needs, it is essential to determine the spatial scale and frequency at which specific processes occur, allowing for the selection of an appropriate remote sensing method.

Spatial Scale

A spatial scale is needed that can monitor at both the habitat level and the species level. The spatial scale can be split into horizontal and vertical scales.

Horizontal Scale

The habitats on Saba range in size, with the smallest habitat being the Elfin forest (estimated 7.2 ha by A. Debrot et al., 2025) and the largest habitat being the dry shrubland and grassland (estimated 470 ha by A. Debrot et al., 2025). At the very least, the horizontal resolution should be able to resolve the smallest habitat. To monitor at the species level, a much finer horizontal scale is needed. The largest species can be considered trees and shrubs, which can grow a few meters wide, while the smallest species are considered to be vegetation in its early growing state, which can be about 5 - 10 cm at their beginning stage. Therefore, a horizontal spatial resolution of 10 cm or less is necessary to accommodate all of the above.

Vertical Scale

Vertical information is important to include due to the steep and elevated nature of this island, making several habitats and species strongly elevation-dependent. A vertical resolution of at least 100 m is needed to distinguish habitats in their elevation, but a much finer scale is required to distinguish vegetation patches and species. As the lowest vegetation types have an approximate height of about 0.1 m (Freitas et al., 2015), a vertical resolution of ≤ 0.1 m is needed to also distinguish between vegetation types.

Temporal Frequency

The frequency of acquisition must align with the processes being monitored. Long-term changes in habitat distribution can be monitored on an annual basis. However, to see the effects of increasing dry and more intense wet periods on vegetation health, more frequent acquisitions (i.e., biannually) are needed. Other processes that benefit from at least biannual, but preferably more, acquisitions are the monitoring of erosion, the spread of invasive species, or vegetation recovery after clearing roaming livestock. However, the latter also requires a very high spatial resolution to be able to detect vegetation in its early growing stages. At the species level, the monitoring frequency should align with the plant's life cycle and conservation status (IUCN, n.d.). Besides monitoring at consistent time intervals, event-based monitoring is also necessary in the case of hurricane damage and is expected to be needed

more frequently in the future, as the frequency and severity of hurricanes are expected to increase (KNMI, 2023). After such an event, close monitoring is also needed to see if any invasive species regenerate.

2.5.2. Land Cover Taxonomies for Saba

To translate the monitoring wishes and needs into measurable objects, land cover taxonomies are needed. In this section, a 3-level hierarchical classification system is presented, where level 1 is the broadest and level 3 is the most detailed, as shown in Table 2.2. Marine habitats and species are excluded from the scope of this study, which concentrates exclusively on terrestrial environments through land-focused remote sensing techniques. For the sake of consistency, similar taxonomies are taken as the most recent land cover classification study of Saba from Smith et al., 2013.

Table 2.2: Three-level hierarchical land cover taxonomy for Saba, linking broad land cover classes to habitats and species. Restricted species are based on habitat-specific flora listed in Appendix A, while additional invasive species are listed in Appendix B. IUCN-identified endangered species are indicated with (EN)

Level 1	Level 2	Level 3
Artificial surfaces	Man-made	-
	Invasive species	<i>Antigonon leptopus</i>
		<i>Bothriochloa pertusa</i>
Rangeland	Dry shrubland and grassland	Other invasive species (Table B.1)
		Restricted shrubland species
		Restricted grassland species
Forest land	Elfin forest	Restricted herb species
		Evergreen species restricted to the Elfin forest
	Montane forest	Other species restricted to the Elfin forest
		Evergreen species restricted to the Montane forest
		(Semi) deciduous species restricted to the Montane forest
		Other species restricted to the Montane forest
		<i>Nectandra krugii</i> (EN)
	Dry tropical forest	Evergreen species restricted to the dry tropical forest
		(Semi-)deciduous species restricted to the dry tropical forest
		Other species restricted to the dry tropical forest
		<i>Swietenia mahagoni</i> (EN)
		<i>Guaiacum officinale</i> (EN)
Barren land	Rocks	-
	Beaches	-

Level 1

At level 1, the taxonomy distinguishes among artificial surfaces, rangeland, forest land, and barren land, as also defined by Smith et al., 2013. This level can be used to track forest degradation by natural drivers, such as hurricanes, as well as the long-term influence of reforestation after the removal of roaming livestock.

Level 2

The second level refines the broad categories into habitats identified in the state of nature reports (A.O. Debrot, Henkens, and Verweij, 2018; A. Debrot et al., 2025) that require periodic reporting on the conservation state of these habitats. The only class from level 1 that is not split further is rangeland, as rangeland corresponds to dry shrubland and grassland. This level can be used to track the size,

condition, and distribution of the habitats.

Level 3

The third level gives the taxonomy at the species level. As can be seen, the species are restricted to their corresponding habitat in level 2. Since no official documentation of key and indicator plant species could be found, the focus on restricted species is an attempt to identify key and indicator species for each habitat; restricted and/or dominant species are likely to function as indicators of habitat condition and, in some cases, as key species whose presence or absence reflects the integrity of the ecosystem. For example, epiphytes are a good indicator for the health of the Elfin forest, as their presence and abundance indicate that the humidity is high enough to sustain the species, while pioneer species can be used as an indicator for barren land or degradation, since they are the first to appear in degraded areas. The restricted species that are referred to in the land cover taxonomy table can be found in Appendix A. It is important to note that some of the information used to create Table A.1 originates from Freitas et al., 2015, which relies on a 1999 survey that is now extremely outdated, although it remains the most recent island-wide inventory. Therefore, the occurrence and abundance of species may have changed since then. For the forest habitats, a distinction is made between evergreen and (semi) deciduous species, so that when taken together across habitats, these correspond to the broader forest classes used in the most recent classification system from Smith et al., 2013. For invasive species, the taxonomy shows the most problematic species for Saba (*Bothriochloa pertusa* and *Antigonon leptopus*, as mentioned in Subsection 2.3.3), while also including all other invasive plants listed in Appendix B. The International Union for Conservation of Nature (IUCN) identified endangered species of Saba are also included in the table, indicated with (EN). The man-made class does not require any more detail, as for biodiversity monitoring, the distinction between natural and artificial objects is enough. Rocks and beaches are also not split into further classes for the same reason.

Sub-Level: Remote Sensing Measurables

Each level includes a corresponding sub-level that translates categorical classes into measurable products derived from remote sensing data. The most commonly used techniques for biodiversity monitoring include passive hyperspectral imaging, Synthetic Aperture Radar (SAR), and Light Detection and Ranging (LiDAR) (Mulatu et al., 2017). Passive hyperspectral remote sensing enables the acquisition of spectral signatures at the species level, facilitating the differentiation of species and the identification of invasive species. It also provides biochemical indicators, such as chlorophyll content, and vegetation indices like the Normalized Difference Vegetation Index (NDVI), which are valuable for monitoring vegetation health. LiDAR provides elevation data that can be used to assess vertical habitat structure and species distribution. LiDAR can also capture the three-dimensional vegetation structure, making the detection and monitoring of understory vegetation possible. SAR provides backscatter intensity, which can be used to obtain surface characteristics such as surface roughness and soil moisture.

2.6. State of the Art

Remote sensing techniques are increasingly used in environmental monitoring, biodiversity assessment, and land cover classification (Saritha et al., 2025). Passive multispectral optical remote sensing, for example, is used in change detection to monitor agricultural development, natural habitat loss, and urban expansion, both at global (as seen in Song et al., 2018) and national scales. The results are used in support of climate change mitigation efforts. Spectral signatures derived from passive multispectral data allow for the identification of land surface properties, where classification is possible at the species level (Saritha et al., 2025), and habitat diversity and distribution can be assessed (Cavender-Bares et al., 2022). LiDAR remote sensing is widely used in forestry studies by capturing vertical forest structure, which cannot be achieved with traditional optical methods. Beyond canopy height models (CHM), LiDAR can derive forest variables such as biomass, Leaf Area Index (LAI), gap fraction, and structural parameters (Wang et al., 2024), which are crucial for habitat and biodiversity studies (Fassnacht et al., 2016a). At high point densities, individual trees can be segmented and classified, allowing for species-level analyses (Van Ewijk et al., 2014; Marselis et al., 2018). In terrestrial applications, even individual leaves can be resolved.

Despite the wide range of possibilities offered by remote sensing techniques, their application to Saba has been limited. The most recent study in which Saba is classified using remote sensing data is

from Smith et al., 2013. High-resolution satellite imagery from the WorldView-2 satellite, acquired in December 2010 and February 2011, was used to identify and map both natural and artificial land cover types. The data consists of a 0.5 m resolution panchromatic image, and 8-band multispectral images at 2 m resolution. Using the land cover map together with species-specific habitat requirements and niche-modeling techniques, they aimed to estimate the geographic distribution of key animal species and their habitats across the island. They used supervised and unsupervised classification with a hierarchical land cover typology, resulting in four levels, where level 1 is the least detailed, and level 4 is the most detailed. In Figure 2.10, the level 3 classification map can be seen.

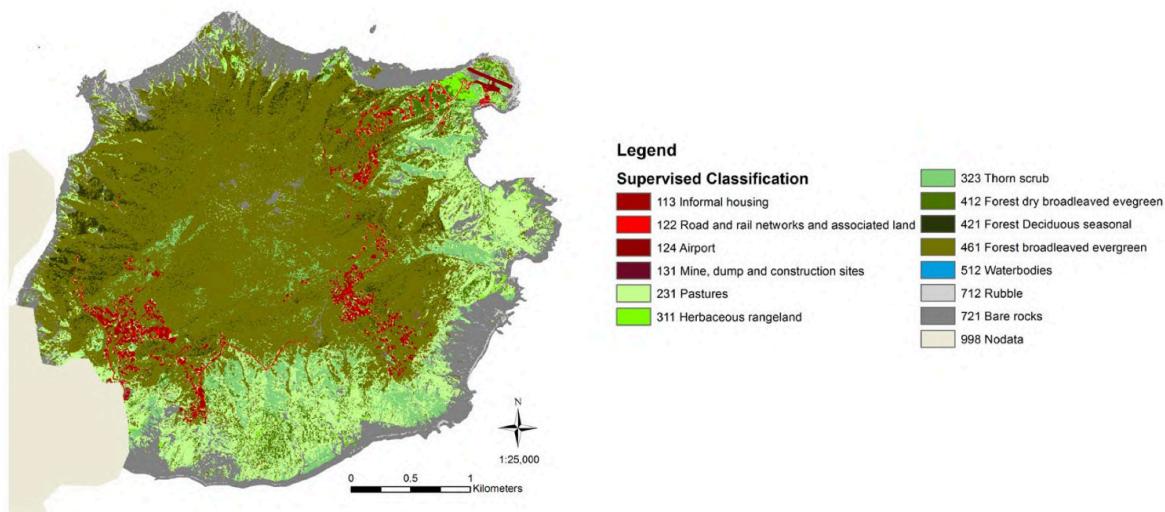


Figure 2.10: Level 3 land cover map of Saba from Smith et al., 2013

While they successfully identified some land cover types with their algorithm, the overall results were limited by the highly heterogeneous vegetation cover and the lack of ecological baseline knowledge. The high degree of heterogeneity led to significant spectral overlap, complicating the differentiation between vegetation types and forest classes. In particular, invasive species could not be reliably classified at the species level, as their spectral signatures were insufficient to separate them, and the resulting maps conflicted with expert judgment. Moreover, invasive species occurring under the canopy could not be detected, and cloud cover and shading from steep terrain further reduced classification accuracy. They suggested that the implementation of vertical information in the form of a high-resolution Digital Terrain Model (DTM) could improve results, but such data were unavailable at the time. Niche modeling was also not possible due to insufficient detail in the map and limited knowledge regarding the habitat requirements of particular species. In response, a follow-up study produced a DTM using stereo aerial imagery (Mücher et al., 2014). However, the resulting DTM had a resolution of only 5 m, a poor height accuracy of about 60 cm, particularly in densely vegetated areas, and a poor ecological relevance as it used observations from 1991. The study concluded that LiDAR altimetry is the preferred method for generating a DTM, due to its superior accuracy and enhanced vegetation penetration compared to photogrammetry. Specifically, it recommends adopting an approach similar to that used in the production of AHN in the European Netherlands. Since then, no further published studies have been found on land cover classification for Saba.

2.7. Research Gap

The limitations outlined above indicate a clear research gap. Currently, there is no recent, island-wide land cover map for Saba that meets the biodiversity monitoring and management needs described in Section 2.5, nor are the necessary data products available to produce such a map. Existing topographic information is outdated, as the only available DTM is derived from 1991 stereo aerial imagery with limited accuracy under dense vegetation. Recent state of nature reports further emphasize that little is known about the present (2024) extent and quality of habitats (A. Debrot et al., 2025), since mapping still relies on a combination of fieldwork-based vegetation surveys (Freitas et al., 2015) and older satellite

classifications (Smith et al., 2013). As a result, the available maps lack both temporal relevance and the spatial detail required for species-level monitoring.

A method that has not been explored yet is the use of airborne remote sensing, which offers higher spatial resolution compared to satellites, while providing full-island coverage in a single campaign. One major opportunity for Saba lies in the application of LiDAR data. To date, LiDAR remote sensing has not been utilized on the island, despite the need for elevation models and structural information, as it has never been acquired before. Similarly, multispectral aerial imagery is of great benefit as it can reach the desired spatial resolution to distinguish individual plants. These datasets can provide structural and spectral information at a spatial resolution much higher than what is currently available. Fortunately, in early 2024, airborne LiDAR and optical aerial imagery were acquired over Saba for the first time. This study investigates the potential of these recently acquired datasets to address the limitations of satellite-based monitoring and provide the detail and quality needed for land cover mapping on Saba.

2.8. Research Questions

The main research question is defined as:

- What is the potential for mapping land cover of the Caribbean Netherlands using airborne LiDAR and passive optical data?

The following sub-questions are identified to help answer the main research question:

1. What land cover taxonomies are relevant for biodiversity monitoring on Saba?
2. What are the properties of the AHN LiDAR dataset?
3. What is the quality of the AHN LiDAR-derived DTM?
4. What structural parameters can be derived from the AHN LiDAR data?
5. To what extent can the AHN LiDAR data be used to assess the land cover taxonomies?
6. What limitations remain, and how can they inform future research efforts?

3

Data Description

This chapter describes the principles of the remote sensing techniques used in this study, and provides an overview of the data. LiDAR remote sensing is discussed in Section 3.1, while passive multispectral optical remote sensing (ORS) is covered in Section 3.2. For further details on these techniques, additional information is available in Appendix C. A total of 4 datasets are used in this study: an airborne LiDAR dataset (described in Subsection 3.1.2), an airborne passive multispectral ORS dataset (described in Subsection 3.2.2), and two satellite passive multispectral ORS datasets (described in Subsection 3.2.3 and Subsection 3.2.4).

The airborne LiDAR data is obtained from Actueel Hoogtebestand Nederland (AHN), while the airborne passive multispectral ORS data is obtained from Beeldmateriaal Nederland. Both datasets are freely accessible through the AHN dataroom and the Beeldmateriaal NL dataroom. The data acquisitions took place at the end of 2023 and the beginning of 2024, marking the first time high-resolution aerial imagery and airborne LiDAR data were collected for the Caribbean Netherlands. The initiative was commissioned by the Ministry of the Interior and Kingdom Relations (BZK) and was carried out on behalf of the Samenwerkingsverband Beeldmateriaal Nederland. Het Waterschapshuis was responsible for processing and checking the data. Additional information on the processing workflow and accuracy assessment is available in Appendix D. All final products are provided in 1 km x 1 km tile format to facilitate more efficient processing. These tiles are referenced in Saba's local Coordinate Reference System (CRS), with each tile name corresponding to the upper-left x and y coordinate. Further details on Saba's local CRS are provided in Appendix E. In Figure 3.2, the names of the tiles can be seen. For example, the upper left tile has an x-coordinate of 2000 m and a y-coordinate of 4000 m in Saba's local CRS, therefore the name of the tile is 2000_4000.

The two additional passive multispectral ORS datasets from satellite sources are obtained from Pleiades-Neo and Sentinel-2. The Pleiades-Neo dataset is obtained from The Netherlands Space Office (NSO), which provides free access to satellite data through their satellite data portal. The Sentinel-2 dataset is obtained from the Copernicus Data Space Ecosystem, which also provides free access to Sentinel-2 data.

3.1. Light Detection and Ranging (LiDAR)

3.1.1. Principles of LiDAR Remote Sensing

LiDAR is an active remote sensing technique that measures the distance between a LiDAR instrument and a target surface by emitting laser pulses and detecting their reflections. These pulses, consisting of coherent light typically in the near-infrared wavelength range around 1064 nm (but 532 nm is typical for bathymetric surveying) (Wang et al., 2024), are emitted toward the surface, and when they encounter an object, part of the energy is reflected back to the instrument. By recording the time it takes for the signal to return, called Time of Flight (TOF), the system can calculate the distance between the sensor and the reflecting surface. By combining the TOF with the positioning and orientation of the platform on which the sensor is mounted, the 3D coordinates (x, y, z) of the reflecting surface can be determined. Besides geometry, LiDAR often records the intensity values of the returned signal, which represent the strength of the backscattered signal and can help distinguish between surfaces of different reflectivity.

There are two different ways a LiDAR system can record the backscattered signal: full-waveform LiDAR and discrete-return LiDAR. In full waveform LiDAR, the complete energy profile of the returned pulse is recorded as a continuous waveform. The waveform then shows the complete vertical structure of the scanned surface. In contrast, discrete return LiDAR records only a few locations in the energy profile by using detection thresholds, which significantly reduces the data volume compared to full-waveform LiDAR. The number of discrete returns is limited by the instrument's maximum discrete returns per pulse. For example, if a LiDAR instrument has a maximum of three discrete returns, a single laser pulse may first reflect off the top of a tree canopy (first return), then from branches or understory vegetation (intermediate returns), and finally from the ground surface (last return); each of these reflections is assigned a return number. In simpler cases, such as when the pulse encounters a flat surface like bare ground, only a single return is recorded.

The spatial resolution is measured by point density D , which is the number of points per m^2 . The flight parameters that contribute to the point density are the Pulse Repetition Frequency (PRF), number of beams N_{beam} , speed v , and swath width W of the platform, for which the latter is the horizontal distance covered by the laser scanner determined by the platform's altitude h and Field of View angle FOV through:

$$W = 2h \cdot \tan(FOV/2) \quad (3.1)$$

The area A covered in a certain time t can be defined as:

$$A = v \cdot W \cdot t$$

And the point density D is then related to its flight parameters through (Vosselman and Maas, 2010):

$$D = \frac{PRF \cdot t \cdot N_{beam}}{A} = \frac{PRF \cdot N_{beam}}{v \cdot W} = \frac{PRF \cdot N_{beam}}{v \cdot 2h \cdot \tan(FOV/2)} \quad (3.2)$$

LIDAR Products

The primary output of a LiDAR survey is a point cloud, which contains the 3D coordinates of the reflected points of the scanned surface. From these point clouds, several secondary products can be derived. One of the most common applications of LiDAR is topographic mapping, where essential products include DTMs and Digital Surface Models (DSMs). A DTM represents the bare Earth surface, excluding vegetation and man-made structures, whereas a DSM consists of all features present on the surface. Another widely used product is the CHM, which describes the height of vegetation measured from the ground to the top of the canopy. A CHM is obtained by subtracting the DTM from the DSM. The different models are illustrated in Figure 3.1.

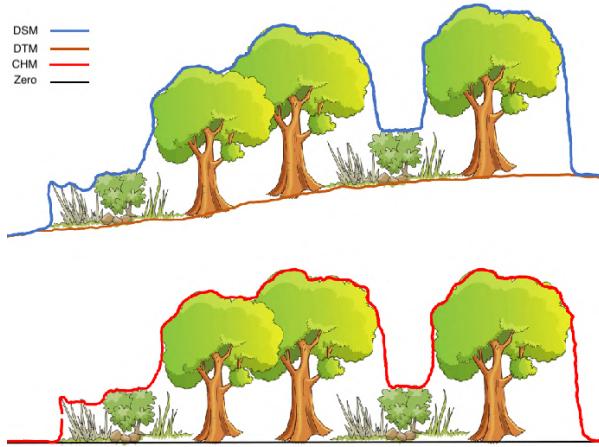


Figure 3.1: Different LiDAR-derived elevation models. The DSM represents the surface elevation, including features such as trees, while the DTM represents the bare ground elevation. The CHM is the resulting height of the vegetation, calculated by subtracting the DTM from the DSM. Source: Jafarbiglu and Pourreza, 2022

Advantages and Limitations

A major advantage of LiDAR compared to other remote sensing techniques is its ability to rapidly acquire direct three-dimensional information on surface objects. Because the light source is coherent

and concentrated in a narrow beam with little divergence, LiDAR has a high detection sensitivity and spatial resolution. Additionally, LiDAR is largely unaffected by electromagnetic interference, allowing it to be used in a wide range of conditions. One limitation is its reduced performance under adverse weather conditions, such as heavy rain or fog, which can significantly attenuate the signal. Another challenge is its penetration through very dense vegetation. Although multiple returns through canopy layers can be recorded, in extremely dense vegetation, the signal may attenuate before reaching the ground, leading to gaps or errors in the DTM. Lastly, steep or rugged terrain can also pose difficulties, as parts of the surface may fall outside the scanning geometry, leading to incomplete coverage. Another limitation is the lack of spectral detail. Conventional LiDAR primarily records the intensity of a single laser wavelength, although some specialized systems can capture limited spectral information, providing little information about surface reflectance properties. For this reason, LiDAR datasets are often complemented with optical remote sensing in applications such as land cover classification.

3.1.2. AHN LiDAR Data over Saba

Acquisition Specifications

The airborne LiDAR data was acquired using a RIEGL VQ-1560 II dual channel waveform processing laser scanning system (RIEGL Laser Measurement Systems GmbH, 2024). Both lasers have the same wavelength in the NIR at 1064 nm (Mandlburger, 2025). The RIEGL VQ-1560 II operates as a multi-return discrete waveform LiDAR system, providing up to 14 discrete returns per pulse. The two scanners generate parallel scan lines, and since the channels are tilted 28° with respect to one another, the scan lines overlap, creating a cross-fire pattern. More details on the scanner specifications can be found in Subsection D.1.1, and a visualization of the scanning configuration can be seen in Figure D.1.

The LiDAR survey over Saba was conducted on 20 December 2023 using a King Air B200 aircraft (PH-SLE) equipped with the RIEGL VQ-1560 II-S sensor. The flight lasted 45 minutes at an altitude of ≈ 2 km, and ground speed of ≈ 67 m/s. A full overview of the LiDAR acquisition parameters can be seen in Table 3.1. The effective swath width was ≈ 938 m, requiring seven parallel NW–SE flight

Table 3.1: Overview of the AHN LiDAR acquisition parameters for Saba on 20 December 2023 using the RIEGL VQ-1560 II-S sensor mounted on an aircraft. It should be noted that the altitude, flight strip overlap, and PRF mentioned in the flight plan do not correspond to those in the flight metadata; therefore, the flight parameters from the metadata are used.

Parameter	Value
speed v	130 kts, 241km/h
altitude h	6562 ft, 2000m
PRF	700 kHz
Laser power	100%
FOV	55°
Flight line overlap	55%
Number of flight lines	7 parallel + 1 cross-line

lines plus two cross-lines. Due to clouds, some lines were repeated, and line 102 was excluded. The full flight report can be found in Subsection D.1.2. For georeferencing, a permanent Global Navigation Satellite System (GNSS) base station (SABY) was utilized, supported by two height control fields and two connection fields, all of which were measured using GPS. The connection fields served as reference areas for block adjustment, while the control fields were used to assess vertical accuracy independently. A visualization of the flight lines and the locations of the connection and control fields can be found in Figure D.2.

Data Products

Point Clouds

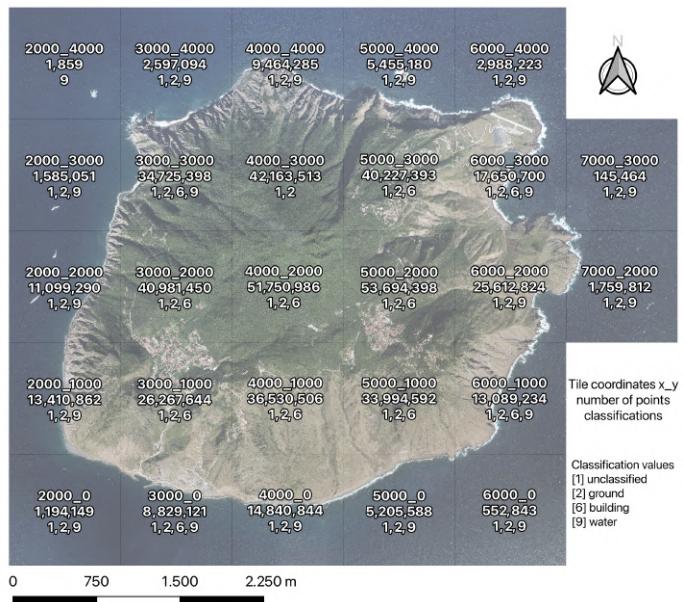
The dataset includes point clouds in .LAZ format, divided over 1 x 1 km tiles, as well as the original flight strip point clouds. The attributes of the points are listed in Table 3.2.

Figure 3.2 shows the number of points present in each tile, and the classification that occurs in each tile. The tile names are also indicated.

DTM and DSM

Table 3.2: Attributes and descriptions of the AHN LiDAR point cloud data of Saba

Attribute	Description
1 X, Y, Z	3D spatial coordinates of each point in Saba's local CRS (Appendix E).
2 Intensity	Strength of the returned laser pulse.
3 Return Number	Specifies which reflection in a multi-return sequence the point represents.
4 Number of Returns	Total number of returns detected for the pulse the point belongs to.
5 Scanner Channel	Identifies the scanner channel that emitted the laser pulse associated with the point. A value of "0" refers to the first scanner, while "1" indicates the second.
6 Scan Direction Flag	Direction in which the scanner mirror rotates. "1" indicates from left to right in the flight direction, "0" from right to left.
7 Edge of Flight Line	Indicates whether the point lies at the edge of a scan line, denoted by "1", else "0".
8 Classification	Coded value representing the type of surface defined by the ASPRS (The American Society for Photogrammetry & Remote Sensing, 2019). 1: unclassified, 2: ground, 6: building, 9: water.
9 Scan Angle	Angle between the laser beam and the nadir direction, where negative values indicate that they are on the left of the plane.
10 Point Source ID	Identifier for the flight line that generated the point.
11 GPS Time	Timestamp of when the pulse was emitted, in Adjusted Standard GPS Time.

**Figure 3.2:** A per-tile overview of the number of points and classifications in the LiDAR point cloud data. Each tile contains its tile name as published in the data, the number of points in that point cloud tile, and the classifications present in the tile. For example, point cloud tile 4000_3000 contains a total of 42,163,513 points, and these points are classified as either 1 (unclassified) or 2 (ground).

Additionally, a rasterized 50 cm resolution and resampled 5 m resolution DTM and DSM are included in the dataset, in the same 1 x 1 km tile format. The DTM was generated by dividing the point cloud into 50 cm x 50 cm grid cells. For each cell, the unweighted average of the z-values of the ground-classified points was calculated and assigned as the DTM elevation. The DTM was resampled to 5 m resolution by calculating the unweighted average of the DTM values within each 5 m x 5 m block, based on the corresponding 50 cm grid cells. If more than 60% of the 50 cm cells within a 5 m x 5 m block (i.e., 60

or more out of 100 cells) contained no data values due to the absence of ground-classified points, the corresponding 5 m grid cell was also assigned a no data value. This resampling rule was eventually dropped because of the extensive data gaps in the 50 cm resolution DTM. The DSM was generated using a similar approach, but then by taking the maximum z-value of all points in each cell. The 50 cm and 5 m resolution DTM can be seen in Figure 3.3 and Figure 3.4, where a self-written python script was used to stitch together the tiles. On the x and y-axis, the coordinates of Saba's local CRS are shown in meters.

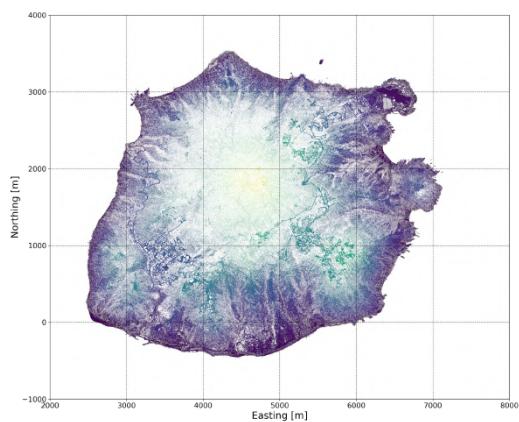


Figure 3.3: 50 cm resolution DTM as published, composed of the individual 1x1km DTM tiles in the AHN LiDAR dataset.

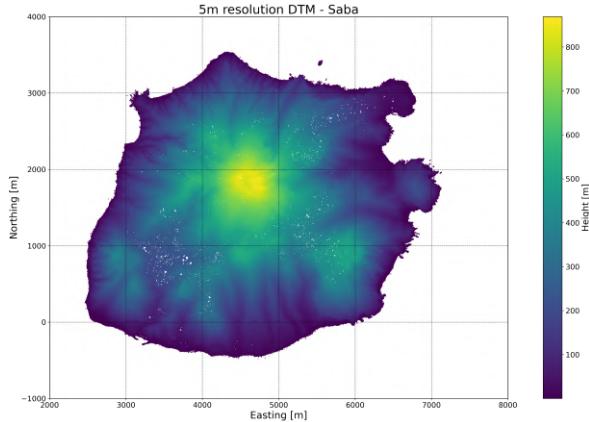


Figure 3.4: 5m resolution DTM - Saba, composed of the individual 1x1km DTM tiles in the AHN LiDAR dataset.

3.2. Passive Optical Remote Sensing

3.2.1. Principles of Optical Remote Sensing

Optical remote sensing (ORS) is a passive method that measures electromagnetic radiation reflected or emitted by the Earth's surface across discrete wavelength intervals, known as spectral bands. The spectrum captured usually ranges from visible (about 400-700 nm) to NIR regions (700-2500 nm), which lie in the optical range, hence the name. Typical bands are the red, green, and blue bands (RGB) from the visible spectrum, and several bands in the NIR. This method relies on the fact that each object on Earth interacts differently with electromagnetic radiation, depending on its physical and chemical properties, resulting in a unique spectral signature. This is illustrated in Figure 3.5, where the absorption coefficient for principal leaf constituents can be seen for the visible and NIR regions. For example, the concentration of chlorophyll can be recognized by its absorption of red light (high absorption coefficient around 600-700 nm in Figure 3.5), while it has a high reflectance in the NIR (zero absorption coefficient in the NIR range) (Saritha et al., 2025). Considering that each species contains specific levels of these vegetation functional traits, they exhibit a spectral footprint that can be utilized for land cover mapping.

For ORS, the resolution consists of three components: the spatial, spectral, and radiometric resolution. The spectral resolution describes the width of each spectral band that the sensor can distinguish. The radiometric resolution defines its ability to differentiate slight differences in energy, expressed as the number of bits used to encode the pixel value. Sensors can be grouped by their spectral coverage; for example, multispectral sensors capture a few broad spectral bands, whereas hyperspectral sensors measure many narrow, contiguous bands, offering much finer spectral resolution.

Optical Remote Sensing Products

The sensor of an ORS imaging system detects the incoming photons and translates their energy into unitless Digital Numbers (DN). The raw image data, therefore, consists of pixels with DN, which are sensor-specific. After calibration, the imagery is converted from raw DN to physically meaningful quantities such as top-of-atmosphere radiance or reflectance and, after atmospheric correction, surface reflectance. These calibrated data serve as the basis for deriving vegetation indices (e.g., NDVI) and for producing higher-level products such as land-cover or vegetation maps.

Advantages and Limitations

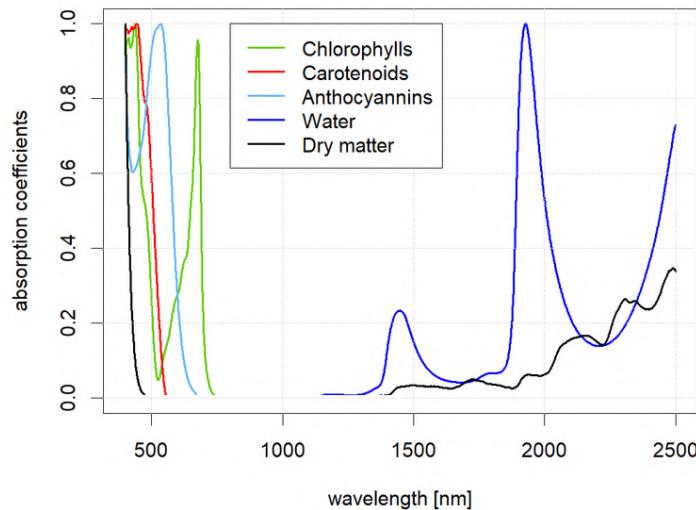


Figure 3.5: Absorption coefficients of leaf chlorophyll content [$\mu\text{g}/\text{cm}^2$], carotenoid content [$\mu\text{g}/\text{cm}^2$], anthocyanin content [$\mu\text{g}/\text{cm}^2$], leaf mass per area [g/cm^2], water content [g/cm^2]. For visualization purposes, all absorption coefficients were scaled between 0-1. Source: Kattenborn, 2018.

ORS enables the derivation of vegetation indices and other biophysical parameters that are difficult to obtain from ground measurements. As this is a passive remote sensing technique, it relies on external illumination, making it ineffective during nighttime or in low-light conditions, thereby limiting the observation time window. Additionally, this method is limited by several weather conditions. Clouds obscure the surface completely through absorption and scattering by cloud drops, especially in the visible, infrared, and submillimeter regions (Elachi and Zyl, 2021). Even thin clouds can scatter and absorb radiation. Other phenomena, such as fog, high humidity, and rain, similarly increase scattering and absorption, while the latter can also temporarily alter surface reflectance. In addition to clouds, shadows cast by terrain or clouds themselves pose a challenge, as they strongly reduce illumination and lead to artificially low reflectance values that do not represent the actual surface conditions. Finally, spectral limitations arise from the broad bands of multispectral ORS sensors, where the spectral resolution can be too coarse, resulting in different objects having similar reflectance characteristics across the available bands (as seen in Smith et al., 2013). In that case, one might need to switch to hyperspectral sensors, which can provide a higher spatial resolution (Saritha et al., 2025).

3.2.2. Beeldmateriaal NL Passive Multispectral ORS Data over Saba Acquisition Specifications

The passive multispectral ORS imagery was obtained using the UltraCam Eagle 4.1 digital aerial mapping system (Imaging, 2022). This camera captures data in the red (580–690 nm), green (490–580 nm), blue (420–500 nm), and near-infrared (690–880 nm) bands. Additional sensor specifications can be found in Subsection D.2.1. It should be noted that at the time of this study, the NIR band data was unavailable.

The imagery was collected over 3 days: 20 December 2023, 26 January 2024, and 28 January 2024 at 17:15-17:32, 14:38-15:50, and 15:47-16:14 local time, respectively. The lens of the camera has a focal length of $f = 120$ mm, which together with the average flight altitude of $h \approx 1586$ m results in a 5 cm spatial resolution. With the constrained altitude, overlap requirements and coverage goals, the total number of planned images in the flight plan was 701.

Data Products

The final dataset includes 651 raw aerial images, aerotriangulated GeoTIFFs before orthorectification, and fully orthorectified GeoTIFFs. The GeoTIFFs are provided in the in $1\text{ km} \times 1\text{ km}$ tile format. Aerotriangulation is the process of aligning overlapping the aerial images by determining their spatial orientation and position, while orthorectification removes geometric distortions to produce georeferenced images with consistent scale and true ground positioning. In these GeoTIFFs, each pixel contains a

DN for the red, green, blue, and NIR channels. The imagery is stored as 8-bit unsigned integer values (0–255 per band), meaning that each color channel is represented with 256 possible intensity levels. Figure 3.6 shows the orthorectified RGB image of Saba when the tiles are loaded in QGIS.

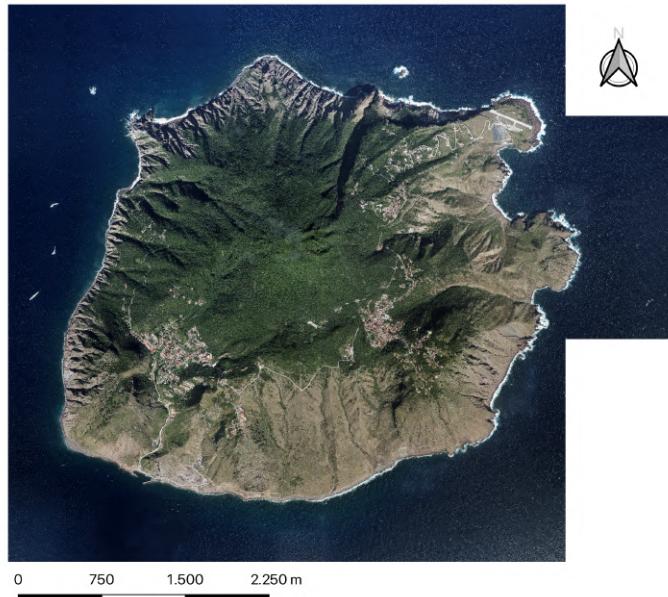


Figure 3.6: Orthorectified RGB composite image of Saba from Beeldmateriaal NL. The individual 1 km x 1 km GeoTIFF tiles are shown together, visualized in QGIS.

3.2.3. Pleiades-Neo Passive Multispectral ORS Data over Saba

As the NIR band was unavailable at the time of this study, NIR data from other sources are considered. The most recent acquisition of Saba with the highest resolution that is freely accessible was made by the Pleiades-Neo satellite. The instrument of Pleiades-Neo measures six spectral bands (Deep Blue, Blue, Green, Red, Red Edge, and NIR) in 1.2 m resolution and a panchromatic band in 30 cm resolution (Netherlands Space Office, n.d.). Data is provided in multiple processing stages, and the most processed data step is chosen (pansharpened RGBI imagery), which has undergone the same processing steps as the AHN imagery. The most recent acquisition, and coincidentally, the most cloud free, is 12-03-2024 14:50:20. Figure 3.7 shows the RGB composite image of this dataset. This dataset has not been atmospherically corrected yet. Therefore, atmospheric correction was done, and the final product contains the surface reflectance.



Figure 3.7: RGB composite of the Pleiades-Neo satellite acquisition of Saba on 12-03-2024 at 14:50:20, with a resolution of 30 cm.

3.2.4. Sentinel-2 Passive Multispectral ORS Data over Saba

Since the Pleiades-Neo data is not entirely free of clouds and shadows, supplementary imagery with additional spectral bands is used to aid in their identification. For this purpose, Sentinel-2 data is considered. Fortunately, the only acquisition with the least amount of clouds is close to the Pleiades-Neo data: 03-02-2024 at 14:57:31. To convert the DN of the optical bands into reflectance values, each DN must be divided by 10,000 (Sentinel Hub, 2024). The RGB composite image of the Sentinel-2 data can be seen in Figure 3.8. The process of cloud and shadow identification of the Pleiades-Neo data using the Sentinel-2 data is provided in Appendix F.



Figure 3.8: RGB composite of the Sentinel-2 satellite acquisition of Saba on 03-02-2024 at 14:57:31, with a resolution of 10 m.

4

Methods

While airborne LiDAR acquisitions have been carried out for the mainland of the Netherlands, and many analysis methods are available, the distinct differences in land cover between the Netherlands and Saba give rise to the aforementioned research (sub-) questions. To address these, a research methodology was created that spans three phases. First, an exploratory analysis of the LiDAR data is done (Section 4.1). A small exploratory analysis of the passive multispectral ORS data is provided in Appendix G. Next, the quality of the LiDAR-derived DTM by AHN is assessed by creating and applying a DTM reliability algorithm (Section 4.2), as the DTM forms the basis for many other relevant studies on Saba, and its quality determines its utility for these studies. Finally, the usability of the airborne LiDAR data for land cover applications is considered by extracting structural information from the point cloud, and applying it in two case studies where they are combined with products from passive multispectral optical satellite imagery (Section 4.3). A schematic overview of the workflow and the resulting products is presented in Figure 4.1 and explained in more detail in the following paragraphs.

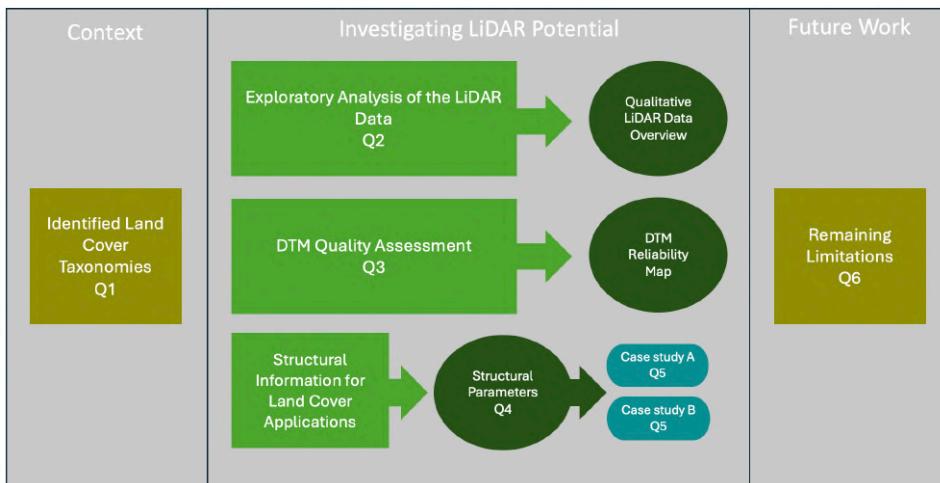


Figure 4.1: Schematic overview of the workflow and results in this study to help answer the main research question. The research sub-questions from Section 2.8 are also indicated in the corresponding box where the sub-question is addressed. Sub-question 1 is discussed in Section 2.5 and sub-question 6 is addressed in chapter 6.

4.1. Qualitative AHN LiDAR Analysis

To provide an initial understanding of the acquired AHN LiDAR dataset and get an initial sense of its strengths and weaknesses, a qualitative analysis is conducted. This serves two main purposes: (1) to visually examine the properties of the point cloud data, and (2) to qualitatively assess the point density and ground coverage.

First, to evaluate how the point cloud data captures Saba's landscape and features, three representative tiles are selected. For each tile, its point-cloud attributes (see Table 3.2) are analyzed to provide an initial

exploratory overview of the dataset. Next, it is qualitatively explored how the executed flight acquisition influenced the point density, which determined the spatial resolution of the dataset. The point density gives a first indication of how effectively this landscape is sampled by LiDAR under the given acquisition conditions. A point density map is generated at 1 m resolution by dividing the point cloud into $1 \text{ m} \times 1 \text{ m}$ grids. Within each m^2 grid, the number of last returns is counted to reflect the number of laser hits per square meter. Water-classified points are excluded to focus solely on terrestrial surfaces. To assess the impact of the flight geometry, the flight lines are included in the point density map to examine how the flight configuration affected the point density. These flight lines are derived from the LiDAR coverage polygons provided in the dataset, as exact flight path data is absent. Finally, the ground coverage of the LiDAR data is investigated. Ground coverage is only briefly addressed in the LiDAR data quality report (Het Waterschapshuis, 2024b), which notes that dense vegetation and steep terrain complicated the DTM derivation. Therefore, an additional analysis is conducted by comparing the AHN acquisition specifications for Saba with those of the European Netherlands, and by examining the published AHN DTM products of Saba, as seen in Figure 3.3 and Figure 3.4.

4.2. DTM Reliability Assessment

To apply LiDAR data for land cover mapping, it is crucial to understand the quality of the different datasets. For a DTM, this quality can be divided into two parts: accuracy and reliability. The accuracy refers to how well the point cloud coordinates correspond to their real-world coordinates, and in the context of a DTM, how well the ground point cloud coordinates correspond to their real-world coordinates. The accuracy has already been assessed by Het Waterschapshuis, and meets the data requirements (see Subsection D.1.4). However, the reliability of the DTM has not been assessed and therefore will be performed in this study.

The reliability refers to how well the DTM reflects the true ground. To assess this DTM reliability, an algorithm is created and applied to the LiDAR dataset. This algorithm is designed to separate areas where the DTM can be trusted from areas with reduced or misleading representation of the true ground surface. An understanding and quantification of this variability is necessary for studies that depend on reliable ground representation, such as land cover mapping. A schematic overview of the workflow of the DTM reliability algorithm is shown in Figure 4.2, and explained in the following sub-paragraphs.

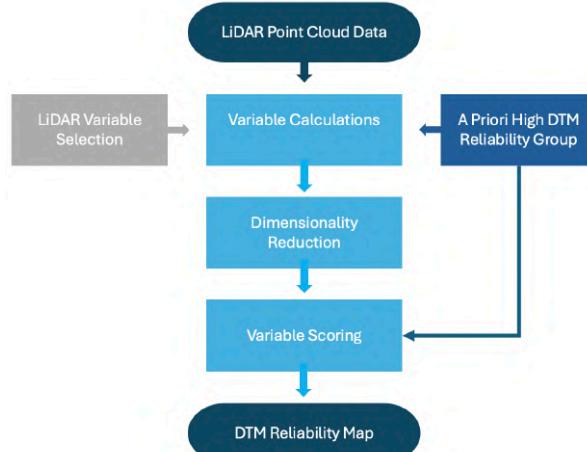


Figure 4.2: A schematic overview of the steps in the DTM reliability algorithm. The steps correspond to the section names of this method.

4.2.1. Identification of LiDAR Reliability Variables

In order to evaluate the DTM reliability, a set of LiDAR-derived variables is identified that are indicative of the DTM reliability. A summary of the set can be seen in Table 4.1. In this table, a distinction is made between variables that can be readily extracted from the LiDAR point cloud attributes (listed in Table 3.2), and the variables that require further computations, using one or more attributes, to be

generated (indicated as derived variables).

Table 4.1: LiDAR-derived variables used as indicators of DTM reliability. Unitless variables are indicated with ‘‘. Point density is normalized by pixel area (0.25 m² per 50 cm × 50 cm pixel).

LiDAR Point Cloud Attributes			
Symbol	Variable	Unit	Description
$I_{\mu,\text{all}}, I_{\sigma,\text{all}}$	Intensity	–	The mean μ and standard deviation σ of the strength of the backscattered signal, (“intensity”) of all LiDAR returns within a pixel (Wang et al., 2024).
$I_{\mu,\text{ground}}, I_{\sigma,\text{ground}}$	Ground Intensity	–	The mean μ and standard deviation σ of intensity of ground-classified LiDAR returns within a pixel.
Derived Variables			
Symbol	Variable	Unit	Description
D_{all}	Total Point Density	points/m ²	The number of LiDAR returns within a pixel, normalized by pixel area (Vosselman and Maas, 2010).
D_{ground}	Ground Point Density	points/m ²	The number of ground-classified LiDAR returns within a pixel, normalized by pixel area.
R_{first}	Ratio of First Returns Classified as Ground	–	Proportion of first returns in a pixel that are classified as ground (Hu et al., 2018).
R_{last}	Ratio of Last Returns Classified as Ground	–	Proportion of last returns in a pixel that are classified as ground (Hu et al., 2018).
R_{ground}	Ratio of Ground-Classified Points	–	Fraction of all returns in a pixel that are classified as ground (Wang et al., 2024).
H_{μ}, H_{σ}	Height Above Local Minimum (HALM)*	m	The mean μ and standard deviation σ of the vertical distance between ground-classified points and the lowest point in a pixel.
l_{μ}, l_{σ}	Path Length*	m	The mean μ and standard deviation σ of the 3D Euclidean distance between the first and corresponding ground return within a pixel.

* Defined in this study (Section 4.2.1).

Because these LiDAR-derived variables are meant to relate to the DTM reliability, they must be expressed in the same spatial framework as the DTM itself, namely a 2D grid of 50 cm × 50 cm pixels. Therefore, the information from the 3D point cloud must be converted into rasterized 2D LiDAR-derived variables on the same grid. Since each point in the point cloud contains its real-world coordinate (attribute 1 in Table 3.2), the pixel grid needs to be defined in real-world coordinates as well in order to divide the points among the pixels. To define the pixel grids for a specific tile, the DTM from that tile was used, and the transformation parameters were extracted from the tile (which contains the conversion from pixel indices to real-world coordinates). Using these parameters, the top-left corner coordinates of each pixel can be computed, allowing each pixel to be uniquely identified by a pixel ID based on its row and column index. Each LiDAR point was then assigned to a pixel ID by computing its relative offset to the origin in the X and Y directions, divided by the pixel resolution, and floored to obtain integer row and column indices. These indices were then flattened into a 1D pixel ID array, making it possible to calculate the pixel-based variables using NumPy’s `np.bincount` function. This function is a very efficient way to do pixel-based computations on large datasets while avoiding overloading system RAM, and significantly reducing computational time. For summative variables like total point density D_{all} , the `np.bincount` function aggregates values for each pixel by summing all values (in this case the count) that share the same pixel ID. For statistical variables, the mean μ and standard deviation σ are calculated as:

$$\mu = \frac{\sum_{i=1}^N x_i}{N}, \quad (4.1)$$

$$\sigma^2 = \frac{\sum_{i=1}^N x_i^2}{N} - \mu^2, \quad (4.2)$$

$$\sigma = \sqrt{\sigma^2}. \quad (4.3)$$

Here, $\sum_{i=1}^N x_i$ can be obtained by passing the values and their pixel ID to `np.bincount`, while $\sum_{i=1}^N x_i^2$

can be obtained by passing the squared values and their pixel ID to `np.bincount`. N is obtained by counting the number of values for each pixel ID.

Because the point clouds are provided in 1×1 km tiles, each tile is processed individually to further manage memory usage. For each tile, the variables are calculated and saved to a .TIF file, preserving the original spatial extent. Once all tiles are processed, each resulting .TIF file represents a complete 2D map of Saba with 50 cm x 50 cm pixels for the variables in Table 4.1, with pixel values corresponding to that variable at each location.

Intensity Variables

The intensity (of returned light) represents the strength of the backscattered signal received by the sensor (Wang et al., 2024). It is influenced by both the geometry and reflectivity of the surface, as well as the amount of energy from the laser that reaches the surface. Highly reflective surfaces such as roads and metals will typically yield a high return intensity, while surfaces like soil and vegetation reflect less energy and thus yield a lower return intensity (Wang et al., 2024). In multi-return signals, the first return will usually have the highest intensity as it originates from the first surface that the signal encounters. Part of the signal is then either absorbed and/or scattered in other directions, which reduces the energy of the signal in subsequent returns. This results in progressively weaker returns for underlying surfaces such as understory vegetation. The sensor contains a detection threshold, above which a reflected signal is recorded as a point. If the remaining energy is too weak to surpass this threshold, no point is recorded, and the signal is considered dissipated. In the context of DTM reliability, ground-classified points within a pixel that have consistently high intensities suggest a more reliable DTM as they correspond to solid and unobstructed surfaces. In contrast, weak and variable intensity ground returns may indicate partial reflections and signal attenuation, which leads to a lower DTM reliability (Goepfert, Soergel, and Brzank, 2008). To capture the difference in how land cover influences the backscattered signal, four intensity variables are derived on a per-pixel basis:

- **Mean intensity of all points $I_{\mu,all}$** : Represents the mean intensity of all points within a pixel. A higher mean indicates that the area generally reflects the laser signal well while a lower mean may indicate the presence of dense vegetation or surfaces with low reflectances.
- **Intensity standard deviation of all points $I_{\sigma,all}$** : Represents the variability in intensity among all points in a pixel.
- **Mean intensity of ground-classified points $I_{\mu,ground}$** : Indicates the mean intensity of points classified as ground in a pixel. A higher mean intensity suggests that the ground returns are stronger and more likely to represent solid and unobstructed (or easily penetrable) surfaces. A low mean intensity indicates weaker returns that arise from (partially) obstructed or poorly reflective ground surfaces.
- **Intensity standard deviation of ground-classified points $I_{\sigma,ground}$** : Indicates the variability in intensity among ground-classified points in a pixel. A low standard deviation suggests consistent intensities, while a higher standard deviation suggests mixed surface conditions.

Computation in Practice

The intensity can be directly calculated, as it is a point cloud attribute (attribute 2 in Table 3.2).

Point Density Variables

The point density is defined as the number of returns per unit area (Vosselman and Maas, 2010). It reflects how densely the terrain is sampled, and is influenced by flight parameters (such as altitude, scan angle, and overlap of flight lines) (RIEGL Laser Measurement Systems GmbH, 2024), as well as surface properties like flatness or the presence of vegetation. For example, the point density will increase in areas covered by multiple overlapping flight lines (which also increases the chance of true ground sampling (Guo et al., 2010)) or where the surface structure generates multiple returns. Total point density can be ambiguous: a vegetated patch may produce many returns from the canopy and understory, resulting in a high total point density, but that does not necessarily translate into good ground sampling. In contrast, a bare patch of ground might yield a similar total point density, but with all points representing reliable terrain. Therefore, the ground point density is also considered. Ground point density is a more direct indicator of how well the ground is sampled *according to a ground classification*

algorithm used by AHN. If the ground point density is low for a pixel, this reduces the trustworthiness of the DTM, while the confidence in the DTM is higher if a lot of points are classified as ground within a pixel. To characterize the difference in how landcover influences the LiDAR observations, two point density variables are derived on a per-pixel basis:

- **Total point density D_{all} :** Represents the total number of points within a pixel.
- **Ground point density D_{ground} :** Represents the number of points classified as ground within a pixel.

Computation in Practice

The total point density is calculated by counting the number of points that fall within a pixel. The ground point density is derived in the same way, but then considering only points that are classified as "2" (attribute 8 in Table 3.2), which corresponds to the ground.

Ratio Variables

Ratios provide insight into the proportional relationship between two variables. As mentioned above in the point density section, considering the combination of total point density and ground point density gives a more complete picture of how the signal interacts with the terrain and overlying features. Therefore, the following three pixel-based ratios are derived:

- **Ratio of first returns classified as ground out of all first returns in a pixel R_{first} :** This ratio reflects how many first returns in a pixel are classified as ground (Hu et al., 2018). In areas of bare ground (or minimal vegetation), a large portion of the first returns will be from the ground, resulting in a high ratio. In contrast, in densely vegetated areas, most first returns originate from the canopy, leading to a lower ratio.
- **Ratio of last returns classified as ground out of all last returns in a pixel R_{last} :** Similarly to R_{first} , this ratio represents how many last returns in a pixel are classified as ground (Hu et al., 2018). The interpretation of this ratio is ambiguous: for an area with penetrable vegetation, the ratio will be higher, while an area with bare ground can yield the same value as in this case, the first return is simultaneously the last.
- **Ratio of ground-classified points out of all returns R_{ground} :** This ratio gives a general measure of how many of the points in a pixel are classified as ground. In densely vegetated areas, this ratio gives the probability that the signal penetrates the canopy to reach the ground (Wang et al., 2024). In regions of bare ground, this ratio is expected to be (close to) 1, which increases the DTM reliability. In regions of dense vegetation, a high ratio can indicate the presence of canopy gaps, while a low ratio suggests poor ground visibility.

Computation in Practice

In the derivation of the ratios, several point cloud attributes are used. For the ratio of the first returns classified as ground out of all first returns in a pixel R_{first} , the numerator is the number of points that have a return number of 1 and a classification label of "2", while the denominator is the number of points that have a return number of 1 (using attributes 3 and 8 in Table 3.2). For the ratio of last returns classified as ground out of all last returns in a pixel R_{last} , the numerator is the number of points that have matching number of returns (how many returns there are in the laser pulse that generated the point) and return number, and are classified as ground, while the denominator is the number of the former (using attributes 3, 4, and 8 in Table 3.2). The ratio of ground-classified points out of all returns R_{ground} is the number of points that are classified as ground divided by the total number of points (using attribute 8 in Table 3.2).

Height Above Local Minimum (HALM) Variables

The Height Above Local Minimum (HALM) is defined as the vertical (z) distance between a ground-classified point and the lowest point in a pixel. This variable reflects the topographic variability and is inspired by the coefficient of variation of elevation in Guo et al., 2010, which considers the normalized variability in elevation of ground points. To characterize the topographic variability of the ground, two statistical HALM variables are derived on a per-pixel basis:

- **Mean HALM H_{μ} :** Represents the mean vertical distance between ground-classified points and the lowest point in a pixel, referred to as the local minimum. A low mean HALM suggests that

ground points are close to the pixel's local minimum, which can be expected in flat areas, while high mean HALM values are ambiguous: it can indicate a mix of ground and misclassified vegetation points in densely vegetated areas, while well-sampled steep areas can yield the same value.

- **HALM standard deviation H_σ :** Represents the variability in distance above the local minimum of ground-classified points. A low standard deviation indicates consistent elevation, while a high standard deviation, similar to the mean, can result from misclassified vegetation points or complex topography (Guo et al., 2010).

Computation in Practice

First, the number of ground points in a pixel is identified (using attribute 8 in Table 3.2). Then, the lowest point in the pixel is identified, and the z value of the ground points is used (using attribute 1 in Table 3.2) to calculate the distance from the ground point to the lowest point in the pixel. There exists special cases for which H_μ and H_σ cannot be computed:

- Case 1: The pixel only contains one ground-classified point that is simultaneously the lowest point. In this case, H_μ and H_σ cannot be computed as HALM is undefined.
- Case 2: The pixel contains one ground-classified point, and the lowest point is not a ground-classified point. In this case, H_μ can be computed (but it will be based on a single value) while H_σ cannot be computed.
- Case 3: The pixel contains two ground-classified points, and the lowest point is classified as ground. In this case, H_μ can be computed (but it will again be based on a single value) while H_σ cannot be computed.

For all three cases, a NaN value is assigned for the undetermined variables.

Path Length Variables

The path length is defined as the 3D Euclidean distance between the first return and the corresponding ground return. If a ground-classified point is also the first return, its path length is zero. The path length l is calculated as:

$$l = \sqrt{(x_1 - x_g)^2 + (y_1 - y_g)^2 + (z_1 - z_g)^2} \quad (4.4)$$

Where (x_1, y_1, z_1) are the coordinates of the first return, and (x_g, y_g, z_g) the coordinates of the ground-classified return. A large path length indicates that the signal traveled through tall vegetation or other structures before reaching the ground, suggesting a potential uncertainty in ground point classification (Goepfert, Soergel, and Brzank, 2008). Shorter path lengths, or a path length of 0, indicate regions with no or low vegetation, making the DTM more reliable in these areas. To characterize how land cover influences the distance of the traveled signal, two path length variables are derived on a per-pixel basis:

- **Mean path length l_μ :** Represents the mean path length of the ground-classified points within a pixel.
- **Standard deviation of path length l_σ :** Gives the variability among the path lengths of ground-classified points within a pixel.

Computation in Practice

To calculate the path length, the pulse ID of each point is needed. The pulse ID is the unique identifier assigned to each emitted pulse, so that it can be reconstructed which points were generated by the same pulse. However, the pulse ID is not a point cloud attribute for this dataset. To derive which points belong to the same pulse, the LiDAR points are grouped based on their scanner channel, point source ID, and GPS time (attributes 5, 10, and 11 in Table 3.2). These groups contain points that were generated at the same time, from the same scanner, and from the same flight line, therefore they originate from the same pulse. Grouping points is computationally heavy, and system Random Access Memory (RAM) can reach a maximum easily in this way, as each tile consists of tens of millions of points. Therefore, the following optimized workflow is applied to each tile to derive the path length:

1. If a tile consists of more than 20,000,000 points, the point cloud tile is split into chunks of 1,000,000 points.

2. Each chunk is split further into two groups: a first group of ground-classified points and a second group of points that are both first returns and come from a pulse for which the number of returns is larger than 2.
3. From the first group, the points that are first returns are assigned a path length of 0.
4. From the remainder of the ground-classified group, their first return partner is searched in group 2 based on equivalent GPS time, scanner channel, and point source ID.
5. The 3D Euclidean distance is calculated between the first return and the ground point using Equation 4.2.1.

As the path length is calculated per tile, a problem occurs for points that have a pulse crossing the border between two tiles. For example, the first return of a pulse lies in one tile, while the last return of a pulse lies in the adjacent tile. The extent of this problem is also investigated and discussed in Subsection 6.2.1.

4.2.2. A Priori Determined High DTM Reliability Group

A subgroup of pixels can be determined from the full dataset that is assumed to have the highest DTM reliability: areas of bare ground where the surface is directly visible to the LiDAR sensor. In these areas, the DTM quality depends primarily on the inherent accuracy of the LiDAR system, which has been assessed by Het Waterschapshuis (see Subsection D.1.4). Since these areas are free of vegetation and terrain complexity, the classification of ground returns is assumed to be representative of the true ground. Because of the lack of validation data, this group serves as a reference for high-confidence ground conditions that the algorithm should aim to recognize.

The pixel group is defined by constraining the following variables from Table 4.1:

- $R_{ground} = 1$: All points in the pixel are classified as ground, indicating no canopy or obstruction.
- $R_{first} = 1$: All first returns in the pixel are classified as ground.
- $l_\mu = 0$ (and thus $l_\sigma = 0$): All ground points are the first and only return in their pulse path, suggesting direct interaction with the surface and no prior scattering.

The characteristics of this group are later used in the scoring function (Subsection 4.2.4) to recognize similar conditions elsewhere in the dataset and score them accordingly.

4.2.3. Dimensionality Reduction

The interpretation of the individual variables listed in Table 4.1 is not only ambiguous, but it does not provide a single reliability metric that is easy to use. For example, a short nonzero path length can result from both favorable and unfavorable conditions: in densely vegetated areas, a short path length may indicate strong signal attenuation, while in sparsely vegetated regions, it simply reflects low vegetation cover. Therefore, combining the variables into a single DTM reliability score would solve the ambiguities because each variable contributes complementary information.

However, combining all variables in Table 4.1 into one single reliability metric also causes additional issues. For instance, two or more variables may carry redundant information. If this is the case, the final DTM reliability metric will be biased towards one or more factors that influence the DTM reliability. Therefore, it is important to address redundancy through dimensionality reduction by removing redundant variables. It is proposed to do this using a method of evaluating pairwise associations between variables using correlation metrics. The correlation metric that best suits the data is Spearman's rank correlation coefficient ρ_s , which measures the monotonic relationship between two variables (Zar, 2005). It is rank-based, making it less sensitive to outliers, works for discrete data, and does not assume a particular distribution or linearity. For two random variables x and y , Spearman's ρ_s is calculated as:

$$\rho_s = \frac{\text{cov}(R_x, R_y)}{\sigma_{R_x} \sigma_{R_y}} = \frac{\sum_{i=1}^n (R_{x,i} - \bar{R}_x)(R_{y,i} - \bar{R}_y)}{\sqrt{\sum_{i=1}^n (R_{x,i} - \bar{R}_x)^2} \sqrt{\sum_{i=1}^n (R_{y,i} - \bar{R}_y)^2}} \quad (4.5)$$

Here, R_x and R_y are the rank-transformed variables of x and y , \bar{R}_x and \bar{R}_y are their respective mean ranks, cov denotes covariance, and σ the standard deviation. ρ_s varies between -1 and 1, with 0 implying no correlation, and -1 and 1 implying a perfect negative and positive monotonic relationship, respectively. Since dimensionality reduction does not require the information of the direction of the relationship (positive or negative), the squared Spearman correlation coefficient is used. This way, $0 \leq \rho_s^2 \leq 1$, which implies that ρ_s^2 reflects the proportion of variability in one variable's ranks that can be explained by the ranks of another.

To identify and eliminate redundant variables, the ρ_s^2 values of each possible pair of variables are computed. The ρ_s^2 values will then be ranked, and pairs with high and moderate correlation are flagged. These pairs indicate that a large proportion of variability in one's variable's rank can be explained by the other, thereby being redundant. The threshold of the high and moderate correlation groups depends on the ρ_s^2 values found in the dataset. Then, using the ρ_s^2 matrix where the flagged pairs are indicated, variables from each pair will be eliminated considering their interpretability and ρ_s^2 correlation with other variables. Given RAM limitations that render full-dataset dimensionality reduction impractical, a random subsample of pixels from the dataset is extracted. The subsample size must be big enough so that the original distributions are still represented, but small enough not to overload system RAM. To accommodate for this, the subsample is drawn by selecting every n th pixel along the x and y direction, where n is chosen accordingly.

4.2.4. Variable Scoring

To translate the remaining LiDAR-derived variables after dimensionality reduction into a single DTM reliability score, it is necessary to normalize the individual variables first by bringing them to the same scale [0, 1], so that each value of a variable is associated with a "score" between 0 and 1. A score of 1 should correspond to conditions that are most indicative of high DTM reliability, and a score of 0 to the least favorable conditions. For example, the point density may range from 0 to 25 points/m², and a high point density is associated with a high DTM reliability, therefore the values should be scaled from 0 to 1. In contrast, the path length may range from 0 to 15 m, and a low path length is associated with a high DTM reliability, therefore, the values should be scaled the other way around from 1 to 0. The proposed method for normalizing the variables involves applying a scoring function defined as $R(x)$, which averages two functions: the Empirical Distribution Function (ECDF) $F(x)$, and the Logistic Function $S(x)$.

The Empirical Distribution Function (ECDF) $F(x)$

Because of the lack of validation data, a data-driven and generalizable normalization method is preferred. This avoids relying on fixed assumptions or knowledge only accessible through validation data; for example, manually defining thresholds for what counts as a "favorable" value for a given variable. The distributions of the variables are expected to differ from each other (as there are discrete variables, bimodal variables, and probably skewed variables). To accommodate this, the ECDF is used for normalizing the variables. The ECDF is a non-parametric estimate of the cumulative distribution function of a variable, which describes how much of the data lies at or below each value of the variable by using their ranks. It is calculated as:

$$F(x) = \frac{1}{n} \sum_{i=1}^n 1_{x_i \leq x} \quad (4.6)$$

Where n is the length of data, and $1_{x_i \leq x}$ the indicator function where it is 1 if $x_i \leq x$ and 0 otherwise. For some variables, the ECDF needs to be inverted because lower values are associated with higher reliability (path length for example), so the ECDF becomes:

$$F(x)_{\text{inverted}} = 1 - F(x) \quad (4.7)$$

In the case of equal values, these get treated as tied ranks. For example, many points may have the same point density, and in that case, their score must also be the same. All identical values are then assigned the lowest possible rank among their group.

The Logistic Function $S(x)$

A problem arises from solely using the ECDF, as the ECDF only describes how a value ranks relative to others in the dataset and provides no indication whether a value is inherently favorable or unfavorable

for DTM reliability. In addition, the ECDF is sensitive to the overall spread of the data, as extreme values can disproportionately stretch the lower tail of the distribution. As a result, the majority of values become compressed into a narrow interval of the ECDF. This means that values which lie very close together in the data can receive noticeably different scores, even though their underlying difference may be negligible. To overcome this, and to incorporate information about which values should be associated with high reliability, the ECDF is combined with the Logistic Function. The Logistic Function is defined as:

$$S(x) = \frac{L}{1 + e^{-k(x-x_0)}} \quad (4.8)$$

where x_0 is the midpoint of the curve (at $x = x_0$, $S(x) = 0.5$), k is the steepness parameter, and L the maximum value. The Logistic Function has the shape of a normal ECDF (see Figure 4.3b), and larger k makes the curve steeper so the transition from low to high S values is more rapid. L is set to 1, so that its maximum value is equal to that of the ECDF.

The parameters of the Logistic Function will be determined using the a priori determined high DTM reliability group (defined in Subsection 4.2.2), for which a visualization can be seen in Figure 4.3. For each variable remaining after dimensionality reduction, x_0 is set to either the 5th or 95th percentile of the distribution of the a priori high DTM reliability group for that variable. If low values of a variable correspond to high DTM reliability, the 95th percentile x_{95} is chosen, and if high values of a variable correspond to high DTM reliability, the 5th percentile x_5 will be chosen. By setting these percentiles as x_0 it means that 95% of the a priori determined high DTM reliability group will get a score of 0.5 or higher with the Logistic Function.

The steepness parameter k is chosen such that the Logistic Function rises from $S = 0.5$ to $S = 0.95$ over a distance d between x_5 and x_{95} of the distribution of the a priori high DTM reliability group. To clarify this further, the case is considered where high values of a variable should correspond to a higher value of the Logistic Function. In this case, x_0 is set at the 5th percentile (the red dashed line in Figure 4.3). k is subsequently chosen so that $S(x) = 0.95$ corresponds to the 95th percentile of the a priori high DTM reliability group distribution (the green dashed line in Figure 4.3). This means that 90% of the a priori high DTM reliability group has a value of S between 0.5 and 0.95 (the region between the red and green dashed line in Figure 4.3, corresponding to d), the extremes are mapped between 0.95 and 1 (the region to the right of the green line), and the lower extremes get a value between 0 and 0.5 (the region to the left of the red line). Since the Logistic Function is symmetric, it extends below x_0 to $S = 0.05$ over the same distance d , thereby avoiding a harsh down-weighting of values in the full dataset that fall outside the a priori high DTM reliability group.

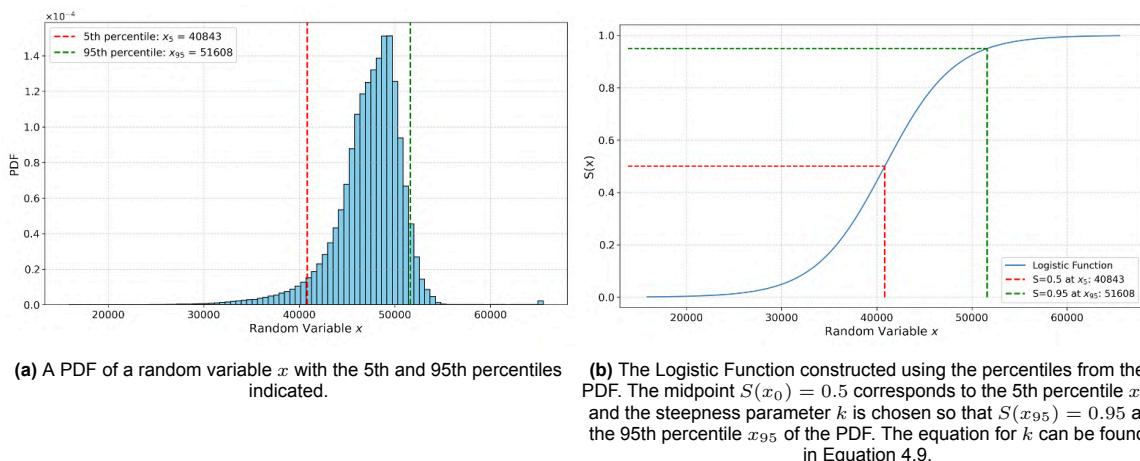


Figure 4.3: Comparison of (a) the PDF of a random variable and (b) the Logistic Function constructed from it.

To find the value of k , the following calculation is done: the midpoint x_0 is set at x_5 , and to ensure that at the 95th percentile the output of the Logistic Function is 0.95, the following needs to hold:

$$S(x_{95}) = 0.95 = \frac{19}{20}$$

Substituting this in Equation 4.8:

$$\frac{1}{1 + e^{-k(x_{95}-x_0)}} = \frac{19}{20} \Rightarrow e^{-k(x_{95}-x_0)} = \frac{20}{19} \Rightarrow -k(x_{95}-x_0) = \ln\left(\frac{1}{19}\right) \Rightarrow k = \frac{\ln(19)}{x_{95}-x_0}$$

Since $x_0 = x_5$, the final expression becomes:

$$k = \frac{\ln(19)}{x_{95}-x_5}, \quad x_0 = x_5 \quad (4.9)$$

In the case that lower values of the random variable x should be associated with higher outcomes of $S(x)$, the midpoint should be set to $x_0 = x_{95}$ and the steepness parameter becomes:

$$k = \frac{\ln(19)}{x_5-x_{95}}, \quad x_0 = x_{95} \quad (4.10)$$

The Scoring Function $R(x)$ for each Variable

The ECDF and Logistic Function are then combined by taking their mean, so the scoring function R for each variable x becomes $R(x) = \frac{F(x)+S(x)}{2}$. By averaging these functions, the final score balances the statistical description of the data with the assumption that the a priori determined high DTM reliability group should show high reliability.

The DTM Reliability Score R_{DTM}

After applying the scoring function $R(x)$ to all remaining variables after dimensionality reduction, the unweighted average of the individual scores is used to compute the final pixel-based DTM reliability R_{DTM} :

$$R_{DTM} = \frac{1}{n} \sum_{i=1}^n R(x_i) \quad (4.11)$$

Where n is the total number of variables left after dimensionality reduction. Based on the distribution of the R_{DTM} scores across the island, the pixels will be classified into reliability classes.

4.2.5. Vegetation Cover Influences on DTM Reliability

To assess how vegetation cover influences the DTM reliability algorithm proposed in Figure 4.2, the DTM reliability results are compared with the NDVI, which is a widely used proxy for vegetation density and health. The NDVI is calculated using the red and NIR bands as follows:

$$NDVI = \frac{NIR - RED}{NIR + RED} \quad (4.12)$$

This index ranges from -1 to 1. Values below zero typically indicate water or other non-vegetated surfaces, values near zero correspond to barren areas, low positive values suggest sparse vegetation, and high positive values indicate dense, healthy vegetation. By comparing the NDVI values across different DTM reliability classes, it can be assessed if the DTM reliability algorithm correctly identifies regions of dense canopy and scores the DTM reliability accordingly. To generate an NDVI map of Saba, data from Pleiades-Neo Passive Multispectral ORS Data over Saba was used. Prior to NDVI calculation, cloud and shadow pixels were removed from this dataset. The methodology for this removal process is described in Appendix F.

4.3. Assessing the Usability of LiDAR-Derived Structural Information for Land Cover Mapping

Land cover mapping is usually done with classification models that group areas of land into different categories (for example, forest, grassland, or urban). These models rely on input features (measurable parameters of the landscape such as vegetation height, canopy density, or reflectance values from

satellite imagery) that help distinguish one land cover type from another. Broadly, classification can be carried out through unsupervised or supervised methods. In unsupervised classification, the model automatically groups areas that look similar in terms of their input features, without needing any prior knowledge of what the groups represent. This does not require training data. In contrast, supervised classification uses examples where the correct land cover type is already known (training data with labels). The model learns from these examples and then predicts the land cover of new, unseen areas. However, in this study, training data is not available. This means that a fully supervised classification workflow cannot be carried out. Nonetheless, in order to assess the usability of structural information in land cover classification, two small case studies are carried out across two levels of the land cover taxonomy Table 2.2.

- (a) For land cover mapping at ecological taxonomy level 2 (habitat level), the land cover classification map from the 2013 Wageningen study (Smith et al., 2013) is used to assess structural variability within a single class in a case study A.
- (b) For land cover mapping at ecological taxonomy level 3 (species level), an AI classification model on the platform "Carto" is used to look at structural variability of the invasive species *Antigonon leptopus* (Coralita) in a case study B.

For these case studies, structural parameters are identified from the LiDAR data that describe both the topography (the shape of the terrain) and the vertical structure of whatever is present on the surface, such as vegetation. Even though developing a classification model is outside the scope of this thesis, these structural parameters can be used as input features in such a model.

4.3.1. Vertical Structure Extraction

Since land cover maps are commonly defined on a pixel basis, the LiDAR point cloud data needs to be translated into vertical structure descriptions at the same pixel resolution. As explained, the vertical structure of the land cover can be used as input features for land cover classification models. The full vertical structure of a pixel consists of the distribution of z values of the points in that pixel, normalized by subtracting the DTM. This normalized height distribution shows different layers of the surface structure, such as the canopy, understory, and ground returns. However, storing and analyzing the full normalized height distribution is computationally demanding. Therefore, the distribution can be reduced to a set of parameters that contains the majority of the information about the shape of the distribution. The pixel resolution is chosen as 5 m x 5 m. This provides a good resolution to investigate the vertical structures within the classes defined by Wageningen, even though their classification map has a resolution of 2 m. At a 2 m resolution, fluctuations in point density could result in undersampling of the vertical structure. When the point density is too low, the presence or absence of a single return can disproportionately affect the shape of the vertical profile, making it highly sensitive to random sampling variation. Therefore, a resolution of 5 m is chosen to ensure that the distribution reflects vegetation structure.

Normalized height Distribution Derivation

Firstly, the point cloud is divided into 5 m x 5 m grids by assigning a unique pixel ID to the points, using the same method described at the beginning of Subsection 4.2.1. Simultaneously, the DTM at 1 m resolution is used to normalize the points by subtracting the DTM value for a specific pixel from the z -value of each point within that pixel. This DTM resolution is selected because coarser resolutions, such as 5 m, could lead to a substantial amount of negative normalized height values in steep terrain. The 1 m DTM is generated by resampling the 50 cm DTM, using the same methodology as AHN for producing its 5 m resolution DTM. The resampled 1 m DTM contains gaps which were filled using the Fill Nodata tool in QGIS with a standard search distance of 10 pixels. The tool uses inverse distance weighting to estimate the missing values. After each point is normalized and assigned a pixel ID, a per-pixel distribution of normalized heights (the vertical profile) is obtained.

Parameterization of Normalized Height Distribution

The goal of parameterization is to obtain a set of parameters that is not computationally too complex, while still retaining sufficient information about the normalized height distributions. To parameterize this distribution, a curve is fitted. The chosen curve-fitting method is the Gaussian Mixture Model (GMM),

which represents the distribution as a weighted sum of n Gaussian distributions, called components. Each component is defined by its mean (μ_{GMM}), variance (σ_{GMM}^2), and weight (w), which together describe the location, spread, and relative contribution of each mode. The weights sum up to 1, and for n components, $3n - 1$ independent parameters are required: n means and n variances, and $n - 1$ independent weights. For example, a bimodal distribution modeled with two components would require five parameters to describe the distribution. The assumption for a GMM is that each component follows a Gaussian distribution. If some components of the distributions are skewed, this assumption is violated, and alternative curve types may provide a better fit. However, the procedure for fitting the distribution must be standardized in order to compare distributions across pixels. Allowing different distribution shapes per pixel would result in parameter sets that are no longer directly comparable. Therefore, the Gaussian assumption is a good starting point.

Selecting the Number of Components

First, a subset of representative pixels is selected using the Wageningen classification map (Figure 2.10). This subset includes five pixels from forest habitats (1 for forest deciduous seasonal, and 2 for forest broadleaved evergreen and forest dry broadleaved evergreen, as they are more abundant), three from the rangeland habitat (pastures, herbaceous rangeland, thorn scrub), and one pixel from the barren land (bare rocks). This subset is a comprehensive representation of the varying normalized height distributions one might encounter. GMMs with varying numbers of components are then fitted to the normalized height distributions, and the most suitable number of components is identified. Here, "suitable" means that the model should capture the main structure of the vertical profiles in the subset without being too simple or too complex. Additionally, the components need to be ecologically interpretable, for example, the first component corresponds to the ground, and the last component corresponds to the canopy layer. If the model uses too few components (underfitting), important features of the vertical structure may be missed, for example, the understory might be overlooked and merged into the canopy or ground layer. On the other hand, if the model uses too many components (overfitting), it may start describing small random fluctuations in the distribution that are not meaningful and have no ecological interpretation.

The suitable number of components is assessed through visual inspection of the Probability Distribution Function (PDF) and ECDFs. Since the curve of the GMM is fitted to the normalized height distribution, evaluating the fit in PDF space provides a direct comparison between the observed distribution and the GMM's ability to represent it. In ECDF space, the quality of the fit can be assessed, specifically in the tails of the normalized height distribution. In addition to visual inspection, the Bayesian Information Criterion (BIC) is calculated for each fitted GMM. The BIC score is a likelihood-based model that can be used in identifying which model (so the number of components) best explains the data while balancing model complexity and goodness of fit (GeeksforGeeks Contributors, 2025). The BIC is best suited for large datasets and tends to favor simpler models (fewer components), since models with a larger number of components are penalized. Usually, the model with the lowest BIC score is chosen as the best fit; however, when a distribution contains numerous small fluctuations, it can be beneficial to add more components to capture these, and the BIC score will be lower, potentially leading to overfitting. To prevent overfitting, the BIC is used to identify the number of components for which additional components provide only small improvements in model fit. Both the GMM fit and BIC score calculation are performed using the `GaussianMixture` function from the Python package `sklearn.mixture`.

Topographic Parameters

In addition to the parameters from the GMM, the DTM, slope, and aspect angle are also included. These topographic parameters give information about terrain characteristics, which are important drivers of vegetation distribution.

4.3.2. Case Study Analysis

In order to investigate the potential of the vertical structure information on land cover mapping, two case studies are performed focusing on a level-2 ecological taxonomy map and a level-3 ecological taxonomy map. In both cases, LiDAR data are combined with products derived from passive multispectral ORS imagery.

Table 4.2: Overview of LiDAR-derived structural parameters. Vegetation parameters are obtained by fitting an n -component Gaussian Mixture Model (GMM) to normalized height distributions. Unitless variables are indicated with ' \cdot '.

LiDAR Vegetation Structure Parameters (from GMM)			
Symbol	Variable	Unit	Description
μ_1, \dots, μ_n	Component Means	m	Means of the i^{th} GMM component ($i = 1, \dots, n$).
$\sigma_1, \dots, \sigma_n$	Component Standard Deviations	m	Standard deviations of the i^{th} GMM component.
w_1, \dots, w_n	Component Weights	—	Weight of the i^{th} GMM component, constrained by $\sum_{i=1}^n w_i = 1$.
LiDAR Topographic Parameters			
Symbol	Variable	Unit	Description
DTM	Digital Terrain Model	m	Average elevation of ground-classified points within a pixel.
α	Slope	°	Local slope angle derived from the DTM.
β	Aspect	°	Slope orientation relative to North.

Case Study A: Level-2 Ecological Taxonomy

For the level-2 taxonomy investigation, the extent of structural heterogeneity within a single vegetation class is investigated. A small subset of the study area is selected where the Wageningen land cover map (Figure 2.10) indicates the exclusive presence of broadleaved evergreen forest. The region is purposely chosen closer to the mountain peak where the Elfin and Montane forests are located, to ensure these forests did not degrade over time, as the Wageningen data dates from 2010 and the LiDAR data was collected at the end of 2023. For the selected region consisting of the broadleaved evergreen forest pixels, the GMM is fitted on the normalized height distributions of the pixels. The resulting parameters (from the LiDAR vegetation category in Table 4.2) are then compared across pixels to assess whether variation exists within the parameters of the same Wageningen land cover class. This analysis helps determine whether pixels belonging to the same land cover class exhibit consistent structural properties or if variation exists that could indicate further subdivisions of classes within the class.

Case Study B: Level-3 Ecological Taxonomy

For the level-3 taxonomy, the second case study investigates whether structural parameters can explain the occurrence of an invasive species on Saba. Here, the Carto AI classification tool, developed by Spheer AI (Spheer, 2025), is used to identify occurrences of *Antigonon leptopus* (Coralita), and is evaluated against the structural parameters.

Carto is a web-based geospatial AI platform developed by Spheer AI B.V. for nature, agriculture, water, and biodiversity monitoring (Spheer, 2025). On this platform, users can draw or import geospatial observations of the indicator they want to monitor. A small AI model is then trained, which is built on top of a foundation model. The foundation model is a geospatial AI model trained on Sentinel-2 satellite time series imagery using self-supervised learning. Once trained, Carto generates prediction maps of the indicators across the desired area at 10 m resolution. Carto serves both regression and classification models, where regression models predict a per-pixel percentage (0% is no coverage, 100% means full coverage) for a given indicator, while classification models assign a discrete class to each pixel.

The indicator chosen as a case study species is *Antigonon leptopus* (Coralita), which has been described in the section on Invasive Species. Carto readily provides a few Coralita observations. After the model is trained, it shows the predicted coverage for the years for which Sentinel-2 data is available, which is from 2017 to 2024. The year 2024 is chosen, as this is closest to the LiDAR observations. The goal of this case study is to investigate whether structural parameters can explain the occurrence of Coralita on Saba, and to evaluate which land cover classes identified by Wageningen are most susceptible to Coralita invasion. From the predicted coverage map, pixels with a Coralita coverage of at

least 50% are selected. This threshold is low enough to include more than just the pixels containing direct observations, and high enough to represent areas with substantial Coralita dominance. To assess which classes are most susceptible to Coralita, the spatial distribution of Coralita pixels is compared to the Wageningen land cover classification map (Figure 2.10). Since Carto provides a map with 10 m resolution, while the Wageningen classification map has a finer 2 m resolution, the following method is applied: for each 10-meter Coralita pixel, the proportion of its area occupied by each Wageningen class is calculated, as there are 25 2 m resolution pixels with a class in each Coralita pixel. Then, the coverage of Coralita is proportionally distributed over the classes within the pixel. For example, if a Coralita pixel has 60% coverage and consists of 10 herbaceous rangeland pixels and 15 thorn scrub pixels, that means that the Coralita pixel contributes $\frac{10}{25} \cdot 60\% = 24\%$ coverage to herbaceous rangeland and $\frac{15}{25} \cdot 60\% = 36\%$ coverage to thorn scrub. By summing these contributions across all pixels, the total Coralita coverage per class can be quantified. In addition, the structural parameters DTM, slope, and aspect from the LiDAR topography parameters in Table 4.2 are used to assess the structural conditions of the Coralita pixels.

5

Results

This chapter provides the results from the methods described in chapter 4. Section 5.1 shows the results from the qualitative AHN LiDAR analysis, Section 5.1 provides the results from the DTM reliability assessment, with its primary result shown in Figure 5.7. Lastly, Section 5.3 shows how the LiDAR data can be translated into structural information, which was then used in two case studies.

5.1. Qualitative AHN LiDAR Analysis

This section provides the results from the exploratory analysis (described in Section 4.1) of the dataset, consisting of the point clouds and the DTMs.

First, three tiles were selected for the point cloud analysis (i.e., 6000_3000 containing the airport strip and cliffs, 4000_2000 encompassing the peak of the volcano, and 3000_1000 containing a residential area located on flat terrain surrounded by steeper vegetated terrain). Figure 3.2 can be used to see where the tiles are located.

The tile containing the airport strip and cliffs illustrated by Figure 5.1 is located on the northeast coastline. It contains approximately 3 million points, which is significantly less compared to the other two tiles considered in this analysis. The lesser number of points is a result of the tile not being fully covered in the LiDAR survey, as it consists mainly of water, but also because the land that is included in the tile does not contain many features that would result in multi-returns. Figure 5.1 shows a somewhat

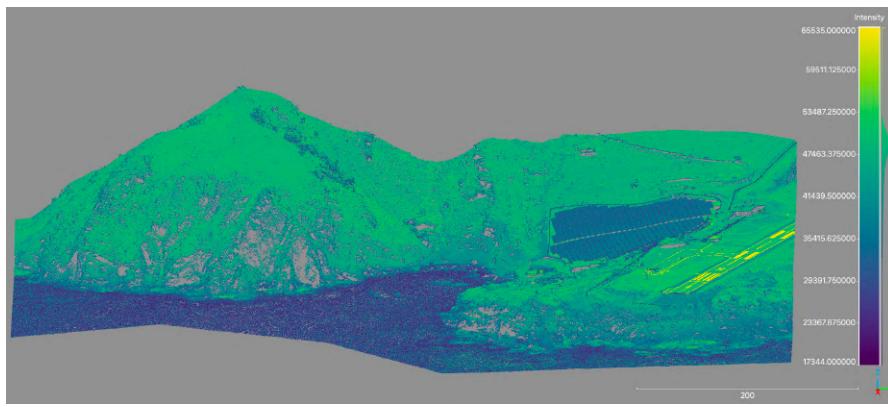
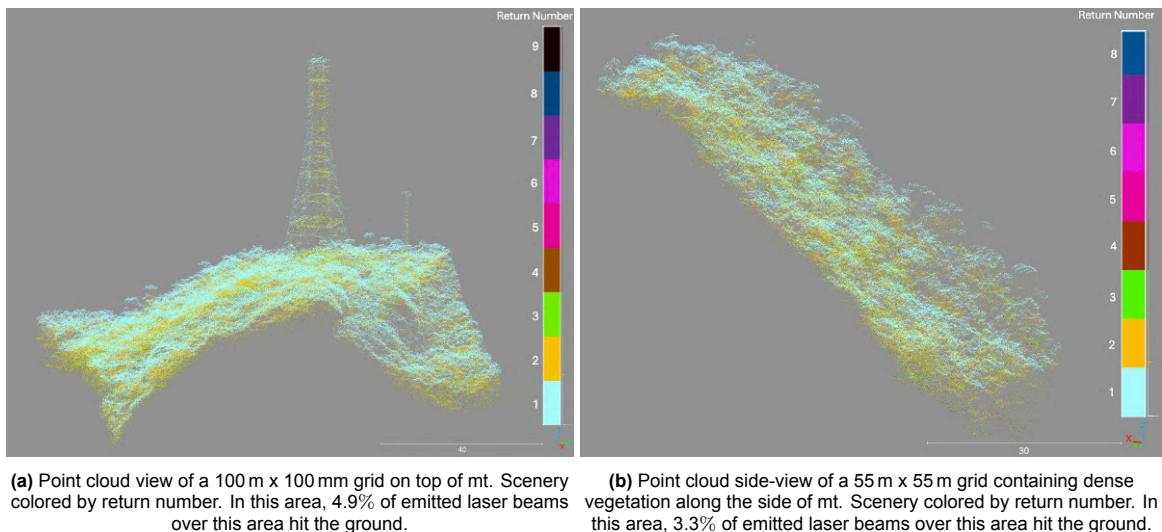


Figure 5.1: Point cloud view of the airport strip tile (6000_3000) colored by intensity (attribute 2 in Table 3.2). The point cloud covers approximately 700 m x 700 m.

smaller part of the point cloud (700 m x 700 m) colored by intensity (attribute 2 in Table 3.2), where yellow indicates a high intensity and dark blue a low intensity. The airport landing strip on the right side of the point cloud is clearly visible, especially the white crosswalk markings, which are shown in yellow in the figure. These stripes reflect the laser beam more effectively due to their white color, resulting in higher intensity values. Adjacent to the landing strip is a field of solar panels, which shows a lower intensity. This is expected, as solar panels are designed to absorb sunlight rather than reflect

it, leading to weaker LiDAR returns. The ocean also shows low intensity, as water bodies strongly absorb NIR wavelengths, which is the wavelength used in this survey. Sparse vegetation is visible in a strip along the edge of the cliffside, characterized by low intensity due to its complex structure. The cliffs themselves show noticeable gaps in the point cloud, highlighting one of the challenges of LiDAR, where steep terrain is difficult for the laser beams to reach. In contrast, the bare, more horizontal rock surfaces show higher intensity because they are flat and unobstructed, allowing for strong returns.

The mountain peak tile (4000_2000) is located at approximately the center of the island, where the forest habitats are located. It contains one of the highest numbers of points in the dataset, with approximately 52 million points. This is primarily due to the dense vegetation surrounding the summit of Mt. Scenery, which generates a high number of multi-returns. In Figure 5.2a, a point cloud covering a 100 m × 100 m area of the peak is shown, colored by return number (attribute 3 in Table 3.2). The transmission tower can also be seen in the point cloud and has a height of approximately 10 m (Jacobs, 2024). Its open structure, composed of numerous beams, promotes scattering and results in multi-return points. Low vegetation surrounds the transmission tower, with some individual species sticking out. These are most likely ferns, identified by their circular growth form. Points with a higher return number (mostly yellow and green) can be seen beneath the canopy layer, indicating deeper laser penetration through the vegetation. The highest recorded return number is 9, although this value occurs only once. Return numbers 4 through 8 are negligible, each accounting for less than 1% of the total. The portion of ground-classified corresponds to 4.9% of the number of last and only returns. This means that out of all emitted laser beams over this area, only 4.9% managed to reach the ground (according to the ground classification algorithm used to process the data).



(a) Point cloud view of a 100 m × 100 mm grid on top of mt. Scenery colored by return number. In this area, 4.9% of emitted laser beams over this area hit the ground.

(b) Point cloud side-view of a 55 m × 55 m grid containing dense vegetation along the side of mt. Scenery colored by return number. In this area, 3.3% of emitted laser beams over this area hit the ground.

Figure 5.2: Point cloud view of the areas on mt. Scenery (tile 4000_2000) showing the peak in (a) and a strip along the side of the volcano in (b), colored by return number (attribute 3 in Table 3.2).

In Figure 5.2b, a side view of a point cloud strip along the slopes of Mt. Scenery is shown, where taller vegetation should be present. Most noticeable is the high number of first returns at the canopy (54% of all points, indicated in light blue), with the number of points significantly decreasing for increasing return number. The canopy appears dense and continuous, making it difficult to distinguish individual trees. Towards the bottom of the point cloud layer, the points are sparse and scattered. In this strip, only 3.3% of the emitted laser pulses resulted in a ground-classified point, showing the challenge of ground sampling in areas with dense vegetation such as these. This is further demonstrated by the lack of a distinct ground layer in the lower part of the point cloud, where no clear line indicates successful ground detection.

Figure 5.3 shows the point cloud of the tile (3000_1000) containing the village in the western part of the island (the Bottom) that lies in a valley. The point cloud is colored by classification (attribute 8 in Table 3.2): orange represents buildings, green represents the ground, and blue is unclassified, primarily consisting of vegetation. The village is built on relatively flat terrain, where unobstructed ground is

easily identified by the dense concentration of green points, as these areas were readily accessible to the laser. In contrast, the steep hillsides are dominated by blue points, with occasional patches of green indicating less obstructed ground. This point cloud shows the clear separation of man-made infrastructure from the natural, vegetated landscape.

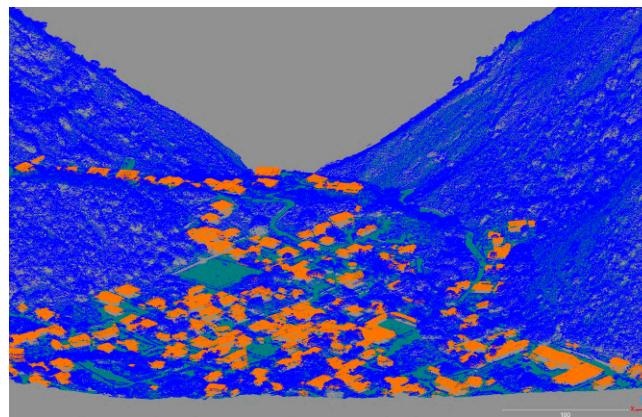


Figure 5.3: Point cloud view of the residential area tile (3000_1000) colored by classification (attribute 8 in Table 3.2). Buildings are colored orange, ground is colored green, and blue is unclassified.

Building on the initial understanding of the differences between the three representative tiles, the qualitative analysis focused on assessing point density, as illustrated in Figure 5.4. Point densities below the required point density of 10 points/m² are indicated in red. The figure also shows the estimated flight lines. From the map, the impact of the number of flight lines on the point density is clearly visi-

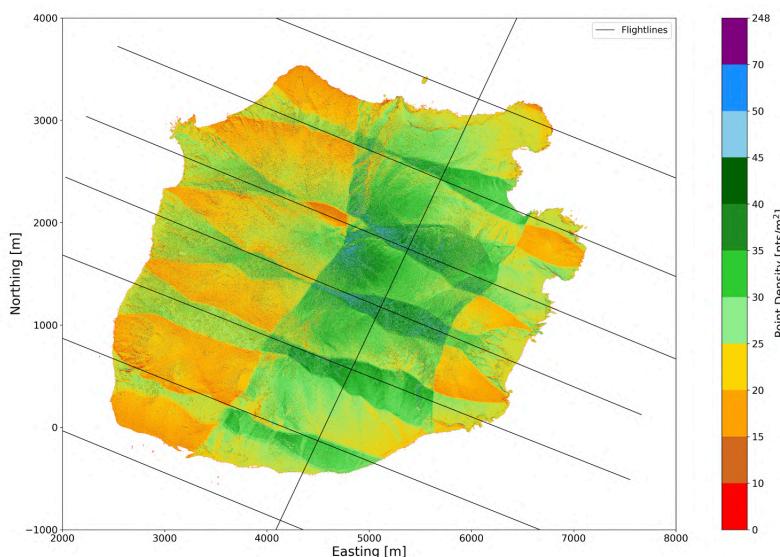


Figure 5.4: Map of the point density from the LiDAR Survey. The point density is defined as the number of only and/or last returns in a 1m x 1m pixel. Additionally, the flight lines are also shown to see the effect of flight overlap on point density.

ble, where higher point densities occur directly beneath the flight lines. Moreover, the most extreme point densities indicated in blue, are mostly concentrated beneath the flight path of the cross line 101. The mean point density is 26 points/m², while the lowest point density is 1 points/m², and the highest point density is 228 points/m². Notably, when looking up the location of the pixel with the highest point density, it is located on a ridge, with coordinates [6838,1494]. Other ridges with extremely high point densities can also be seen in the figure (for example, in the area around [2000, 7000]). Ridges often exhibit higher point densities due to favorable terrain geometry, as they are highly visible to the sensor due to their elevation and are unobstructed by vegetation or other objects. Another noteworthy point density of 213 points/m² belongs to pixel [3505, 907]. This pixel is located in the western village near a

house and contains a transmission tower. Due to its narrow size, it falls within one pixel, and its height caused many pulses to strike the tower. The lowest point density of 1 points/m² is located at pixel [4338,3538], which is along the northern coast. Moreover, a pattern is observed where most pixels with a point density below 10 points/m² are located along the coast. This is likely a result of excluding water classified points from the point density calculation, therefore only the part covered by land contributes to the point density. It can also be seen that pixels adjacent to ridges exhibit significantly lower point densities (for example, the ridge at approximately [6550, 1500] which shows both red and blue pixels), in some cases also falling below 10 points/m². This is due to the steep terrain, which, under unfavorable viewing geometries such as large scan angles, limits the sensor's ability to capture data beyond the ridge. As a result, the area behind the ridge becomes occluded and remains unsampled.

The final step in the qualitative analysis focused on the ground sampling criteria formulated in the AHN specifications. Specifically, while for vegetated areas in the European Netherlands a criterion was specified that for AHN products, at least 25% of laser pulses must return ground hits, this particular criterion was absent for Saba. While the document does not specify why this choice was made, from the analysis it can be observed that the requirement was likely not met. Specifically, from the aforementioned results from the return number analysis of the vegetated mt. Scenery tile (Figure 5.2a), it was found that for the two assessed point clouds located in forested areas (tile 4000_2000), only 4.9% and 3.3% of emitted laser pulses over these areas hit the ground, which is far below the minimum standards for the European Netherlands. This requirement helps ensure a complete DTM, and its absence from Saba's AHN specifications indicates either that the island's challenging sampling conditions made the 25% ground-return threshold unattainable or that, as a first acquisition, the constraints were intentionally kept flexible. In either case, this omission raises concerns about ground classification accuracy and the resulting DTM. To further visualize the implications of this requirement, Figure 5.5 shows a comparison between the DTM of an area in the southeast of Limburg in the Netherlands (Figure 5.5a), known to have forests (like the Vijlenerbos) and elevation differences, and the DTM of Saba (Figure 5.5b) at a similar scale. For the European Netherlands, these areas of elevation differences and vegetation

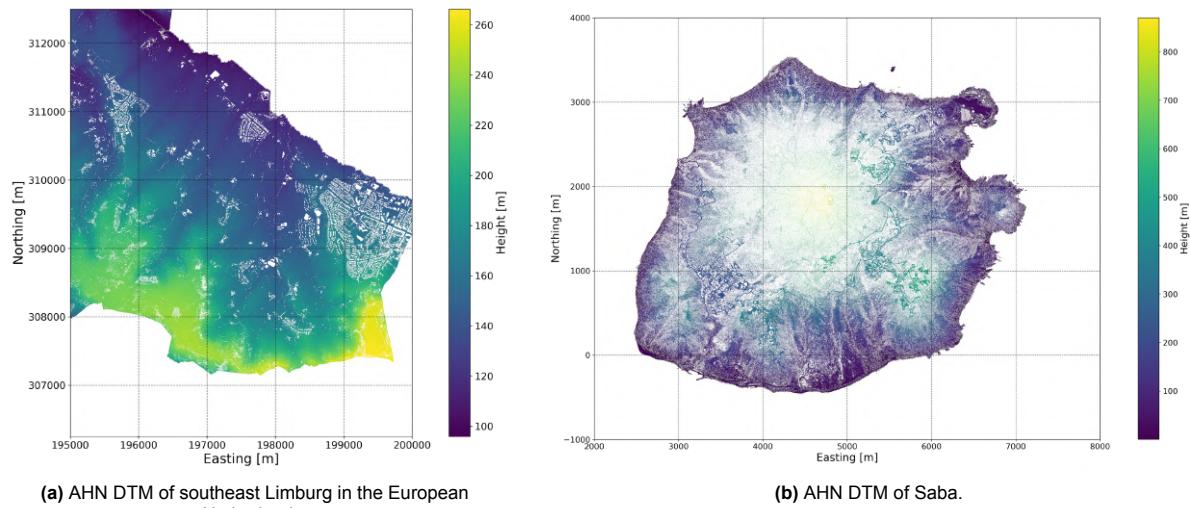


Figure 5.5: Comparison of the DTM of southeast of Limburg in the European Netherlands and Saba's DTM at similar scale. Both are derived with the same method and from the same data acquisitioner (AHN)

are going to be the most "challenging" to survey. Even then, compared to Saba, only data gaps can be seen in residential areas where buildings are masked out (i.e., in tile (310000_199000)). There are some small gaps in the vegetated areas (along the southern border at higher elevations indicated in yellow) but these are negligible, especially compared to the large data gap in the center of Saba's DTM.

Another important difference between the AHN data products for the European Netherlands and Saba lies in the generation of the 5 m resolution DTM. In contrast to the European Netherlands, Saba's 5 m DTM was generated without enforcing the original resampling rule that required at least 40% of the underlying 50 cm cells to contain valid DTM values. The 5 m DTM that was published in the dataset (Figure 3.4), therefore, deviates from the original product specifications. To evaluate the impact of this

decision, the 50 cm DTM was resampled to 5 m resolution using the original 40% threshold. The resulting DTM (Figure 5.6) reveals extensive data gaps across the central part of the island, and additional smaller gaps radiating outward. By dropping the resampling rule, AHN therefore effectively produced

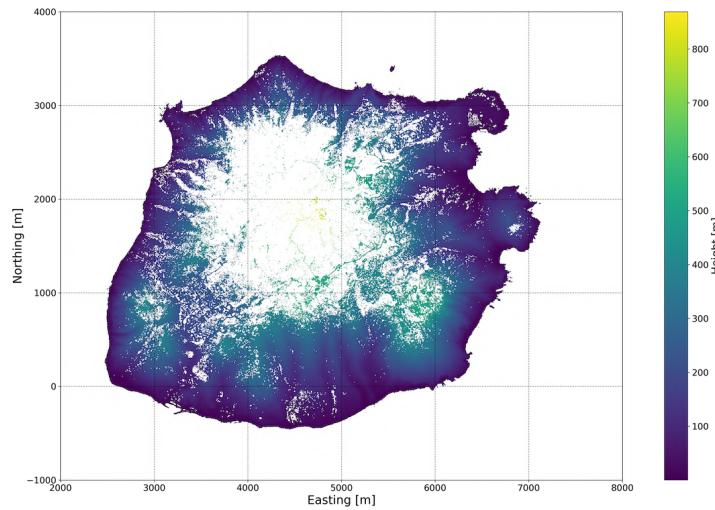


Figure 5.6: 5 m resolution resampled DTM using the resampling rule as originally intended in the specifications. This DTM was not published.

a 5 m resolution DTM that is more spatially complete, but at the cost of reduced adherence to quality thresholds.

5.2. DTM Reliability Assessment

Following the DTM reliability algorithm (explained in Section 4.2), a DTM reliability map was created, as shown in Figure 5.7. Of the island area, 42% falls within the no data category, while 9% is classified

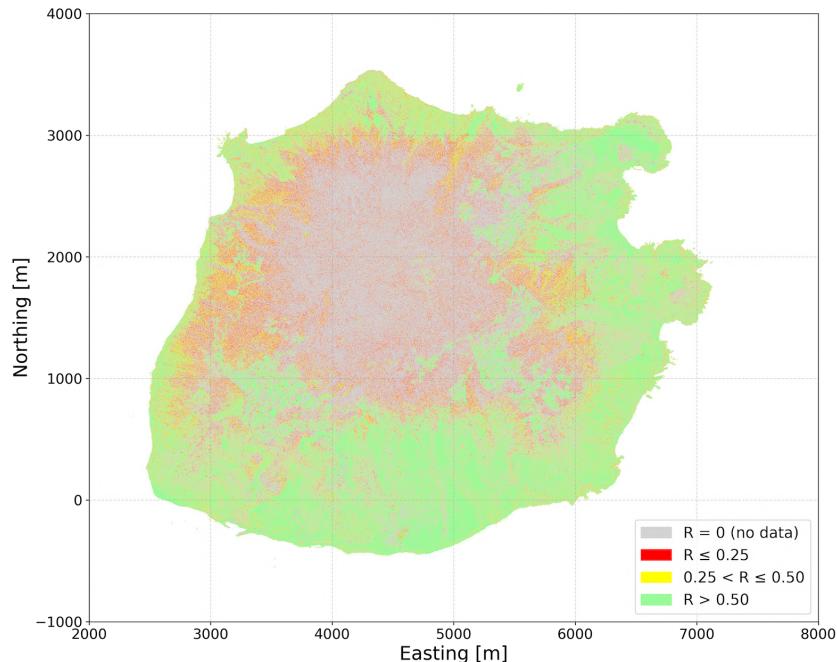


Figure 5.7: Map of the 50 cm resolution DTM reliability of Saba by reliability class.

as having low reliability, 16% as moderate reliability, and 33% as highly reliable.

5.2.1. Identification of LiDAR Reliability Variables

Figure 5.8, Figure 5.10, and Figure 5.12 show the spatial maps of three variables (Ground point density D_{ground} , mean ground intensity $I_{\mu,ground}$, and mean path length l_{μ}) out of the 13 variables listed in Table 4.1 that were selected in the process of identifying variables that are indicative of the DTM reliability. The ground point density map (Figure 5.8) shows how well the ground was sampled across Saba and

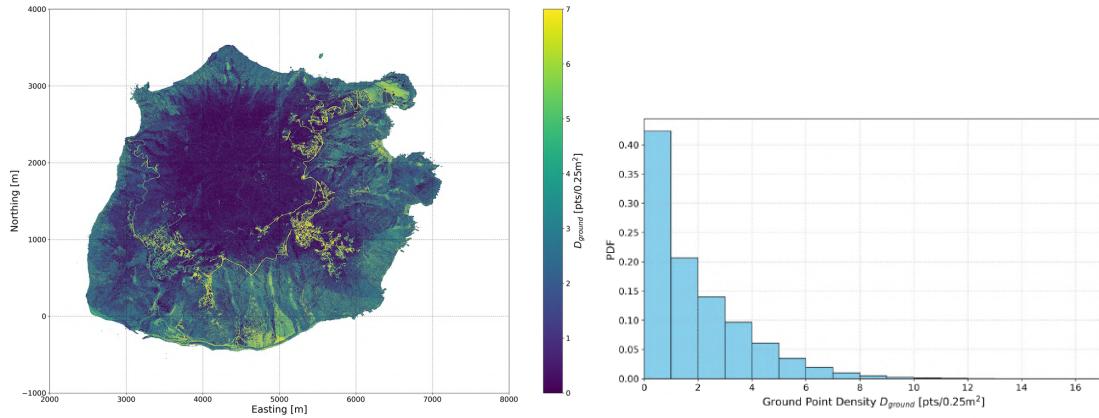


Figure 5.8: Spatial map of the Ground Point Density D_{ground} at 50 cm resolution. The map is clipped at the 99th percentile (7 points/0.25m²).

Figure 5.9: PDF of the Ground Point Density D_{ground} .

serves as the foundation for the 50 cm resolution DTM. There is a clear contrast between the center of the island where mt. Scenery is located (showing a low ground point density), and the surrounding area extending toward the island's edges (showing higher point densities). The highest ground point densities are found in the flattest and least obstructed regions, which are clearly identifiable as roads (shown in yellow). The lowest point densities are found in the central part of the island, where the forest habitats are located. A substantial area appears to have no ground sampling at all, with a ground point density of 0 pts/0.25m². This is further illustrated by the PDF in Figure 5.9, where the highest peak occurs at a ground point density of 0 pts/0.25m². A PDF shows how likely a certain value of a variable is to occur. It shows how this likelihood is spread across all possible values, so the curve of the PDF indicates which values occur more frequently and which ones are rarer. The ground point density PDF is highly right-skewed, with a median of 1 point/0.25m². Some areas with more intensive sampling show extreme values with a maximum of 17 points/0.25m².

The mean ground intensity map (Figure 5.10) shows the average signal strength of the ground-classified LiDAR returns, where higher intensities (indicated with yellow) correspond to stronger returns. Two dis-

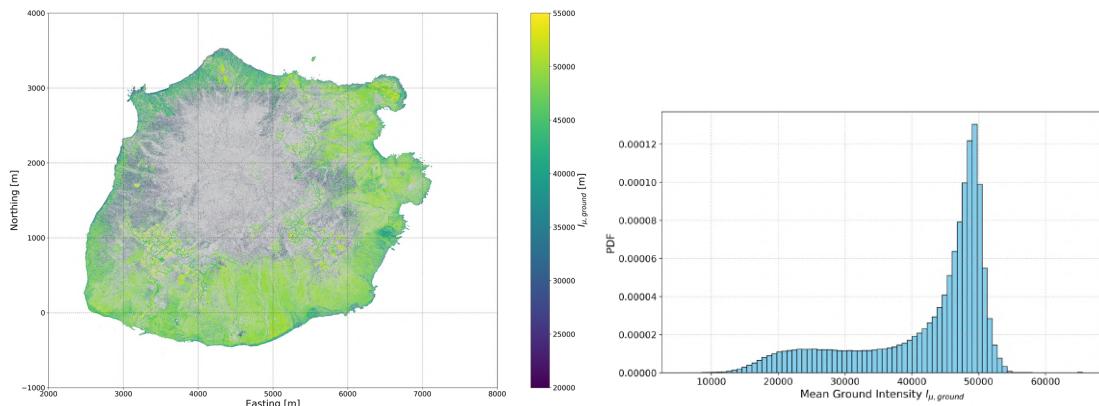


Figure 5.10: Spatial map of the Mean Ground Intensity $I_{\mu,ground}$. The map is clipped between approximately the 5th and 99th percentile. Pixels with no data are shown in gray.

Figure 5.11: PDF of the Mean Ground Intensity $I_{\mu,ground}$.

tinct groups are visible, one with high-intensity values and another with lower intensity values. The

high-intensity areas are primarily located on the lower slopes, mostly where exposed rock surfaces and residential areas with solid structures are present. In contrast, the center of the island shows significantly lower intensity values, likely corresponding to the heavily vegetated and steep areas surrounding mt. Scenery. The PDF of this variable (Figure 5.11) shows up as a binomial distribution, confirming the existence of these two distinct groups.

The mean path length l_μ gives structural information about the surface features within each pixel. In Figure 5.12, a clear border is visible where the path length is zero or close to zero everywhere, corresponding to areas with few surface features, such as bare ground. Although areas with larger path

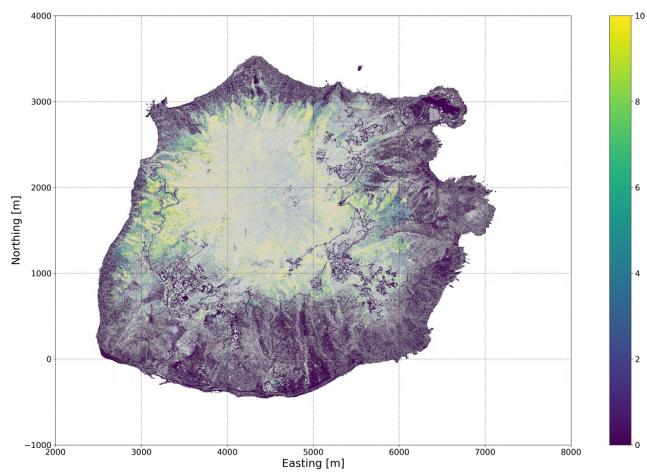


Figure 5.12: Spatial map of the Mean Path Length l_μ . The map is clipped at the 95th percentile. Pixels with no data are shown in gray.

lengths appear to correspond with regions of lower mean ground intensity in Figure 5.10, this variable seems to show more variation in these regions, especially in the center of the island. The central region shows moderately low path lengths, suggesting the presence of shorter vegetation or small patches of bare ground, or that the LiDAR signal did not fully penetrate the vegetation before reaching the ground. Surrounding this area is a ring with significantly higher mean path length, indicating the presence of taller vegetation. Beyond this ring, a mix of high and low path lengths emerges, likely corresponding to secondary rainforests, which are known to contain a mix of different tree heights with patches of bare ground. Additionally, it is noticeable that the area above the northeast to southwest diagonal shows more regions with a longer path length compared to the area below the diagonal. This pattern likely reflects the influence of prevailing easterly trade winds, with the wind-exposed southern side limiting vegetation height, while the more sheltered northern side allows for taller growth.

Appendix H includes the PDFs of all LiDAR derived variables listed in Table 4.1.

5.2.2. A Priori Determined High DTM Reliability Group

Figure 5.13 shows the location of the pixels that were identified as having the highest DTM reliability out of the whole dataset (indicated in black) by constraining the ratio of ground-classified points out of all returns in a pixel R_{ground} , the ratio of first returns classified as ground out of all first returns in a pixel R_{first} , and the mean path length l_μ (as described in Subsection 4.2.2). The group consists of a total of 1,963,926 pixels. This corresponds to 3.7% of the island, and 6.5% of all pixels containing at least one ground point. From Figure 5.13, it is clear that the high DTM reliability pixels are not evenly distributed across the island, as they are mostly concentrated in the lower slopes of the island, in residential areas, and coastal zones. The central part of the island shows very few high-reliability pixels, aligning with the expectation that, because they are located within the forest habitats, these pixels are going to be less reliable here.

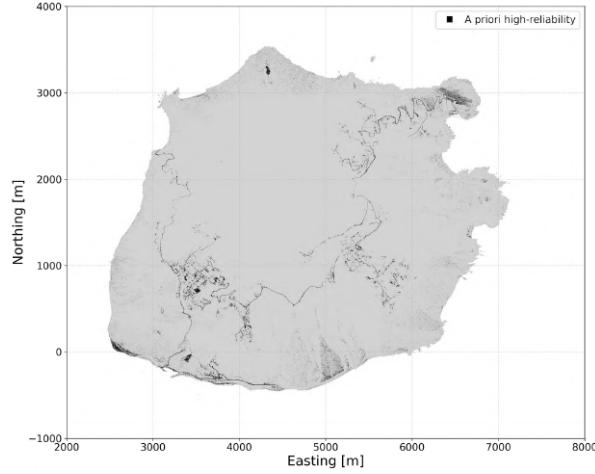


Figure 5.13: Overview of the a priori high DTM reliability group that are identified to have the highest DTM reliability, indicated in black. These pixels are identified through constraining some of the LiDAR variables in Table 4.1.

5.2.3. Dimensionality Reduction

As explained in Subsection 4.2.3, a subsample of the dataset must be drawn to perform dimensionality reduction. The best subsample size without overloading system RAM contained approximately 140,000 pixels. The distributions of the variables of the subsample were compared to the distributions of the total dataset shown in Figure H.1, and the original distributions are still broadly represented. Since there are a total of $N = 13$ variables (Table 4.1), there are $N(1 - N)/2 = 78$ possible pairs, and thus 78 squared Spearman's rank correlation ρ_s^2 values. The pairs were assessed from lowest to highest correlation. From this order, a secluded group of high ($\rho_s^2 > 0.75$) and a secluded group of moderate ($0.35 \leq \rho_s^2 \leq 0.75$) correlation pairs were identified. Figure 5.14 shows the Spearman's ρ_s^2 correlation matrix, where these groups are indicated with red and orange, respectively. For better visualization, the top half of the matrix is removed since it is symmetrical, and the diagonal is also removed because the correlation of a variable with itself is 1. From the high correlation group, l_σ , $I_{\mu,all}$, R_{first} , R_{last} , and R_{ground} were removed. From the moderate correlation group, $I_{\sigma,all}$ and H_σ were removed. The remaining variables are: the mean and std of the ground points $I_{\mu,ground}$, $I_{\sigma,ground}$, the total and ground point density D_{all} , D_{ground} , the mean HALM H_μ and the mean path length l_μ . A full assessment of why certain variables were kept or discarded can be found in Appendix M.

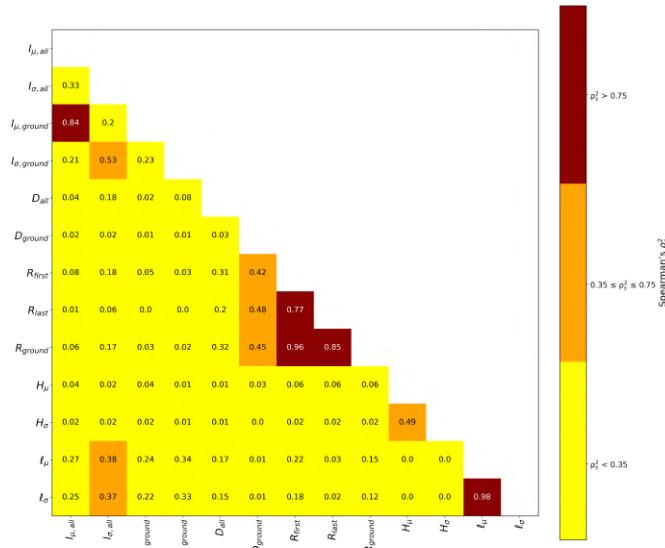


Figure 5.14: Spearman's ρ_s^2 correlation matrix of the LiDAR derived variables listed in Table 4.1 of the subsample data.

5.2.4. Variable Scoring

For the calculation of the ECDF and Logistic Function of the mean ground intensity, $I_{\mu,ground}$, *higher* values are associated with higher DTM reliability, so the ECDF and Logistic Function were calculated using Equation 4.6 and Equation 4.8 with Equation 4.9. The functions, including the scoring function R , can be seen in Figure 5.15. In contrast, the standard deviation of the ground intensity $I_{\sigma,ground}$ and the mean HALM H_{μ} , *lower* values are associated with higher DTM reliability, so the ECDF and Logistic Functions were determined using Equation 4.7 Equation 4.8, and Equation 4.10. For H_{μ} , the functions can be seen in Figure 5.16.

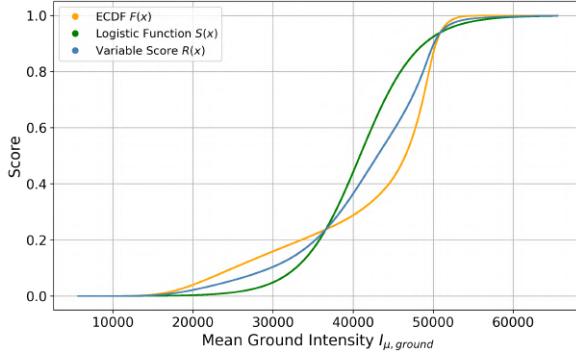


Figure 5.15: The ECDF F , Logistic Function S , and scoring function R for the mean ground intensity $I_{\mu,ground}$.

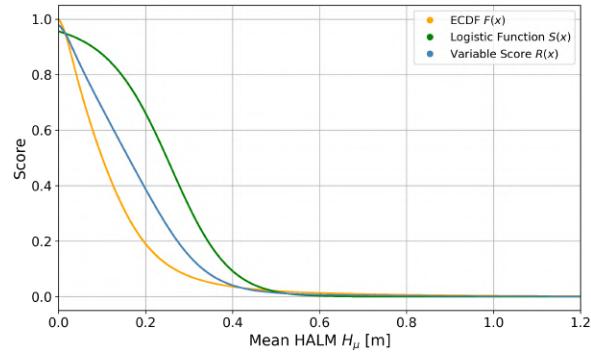


Figure 5.16: The ECDF F , Logistic Function S , and scoring function R for the mean HALM H_{μ} .

Additionally, high values of both total point density D_{all} , and ground point density D_{ground} , should correspond to higher DTM reliability. For these variables, the ECDF (Equation 4.6) is used directly, but the Logistic Function is calculated slightly differently because these variables are discrete. It was chosen to multiply the steepness parameter k by two so that it becomes steeper. The functions for the ground point density D_{ground} can be seen in Figure 5.17. Since the mean path length l_{μ} is the only variable that is also used as a constraint for the a priori high DTM reliability group (mean path length $l_{\mu} = 0$), the group has no distribution and therefore also no percentiles to use as parameters for the Logistic Function. Therefore, only the ECDF (Equation 4.6) of the full dataset is taken as the scoring function R for the mean path length, which can be seen in Figure 5.18.

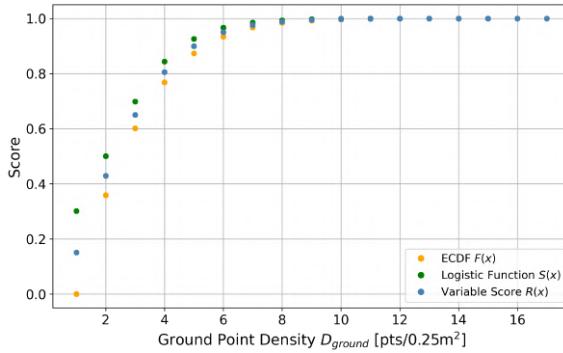


Figure 5.17: The ECDF F , Logistic Function S , and scoring function R for the ground point density D_{ground} .

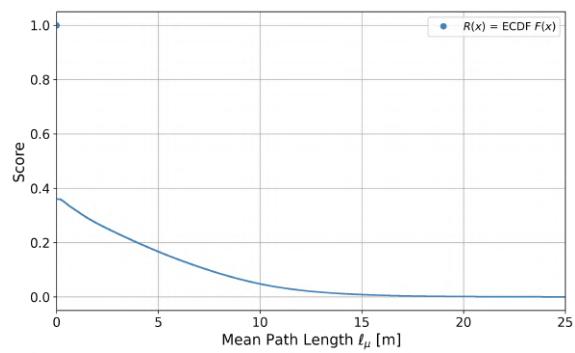


Figure 5.18: The ECDF F for the mean path length l_{μ} . The ECDF is also used as the scoring function R , since the mean path length was used as a constraint in the a priori high DTM reliability group (Subsection 4.2.2), therefore this group only has a single value for l_{μ} .

An interesting effect of the Logistic Function can be seen in the mean ground intensity $I_{\mu,ground}$. This variable has a bimodal distribution (see Figure 5.11). One mode corresponds to lower intensity values, which are likely pixels covered with vegetation or shadowed areas, and the other mode corresponds to higher intensity values, which likely correspond to bare ground. The distribution of $I_{\mu,ground}$ within

the a priori high DTM reliability group only shows the high intensity mode. Since this group was constrained to bare ground, it can indeed be inferred that the high-intensity mode corresponds to bare areas. In Figure 5.15, it can be seen that the Logistic Function correctly assigns higher S scores to this higher intensity range. The ability of the Logistic Function to correctly emphasize bare ground pixels demonstrates that its use is a justified and effective design choice within the DTM reliability algorithm.

A total overview of all functions for each variable, including the remaining variables not shown here (the std of the ground intensity $I_{\sigma,ground}$ and total point density D_{all}), can be found in Appendix I.

After the scoring function $R(x)$ (Equation 4.11) was applied to each variable, Equation 4.11 was used to obtain the DTM reliability score for each pixel. Figure 5.19 shows the PDF of the DTM reliability scores R_{DTM} for all pixels that include at least one ground point. Logically, pixels containing no ground points do not have a DTM and therefore receive a DTM reliability score of $R_{DTM} = 0$. Following

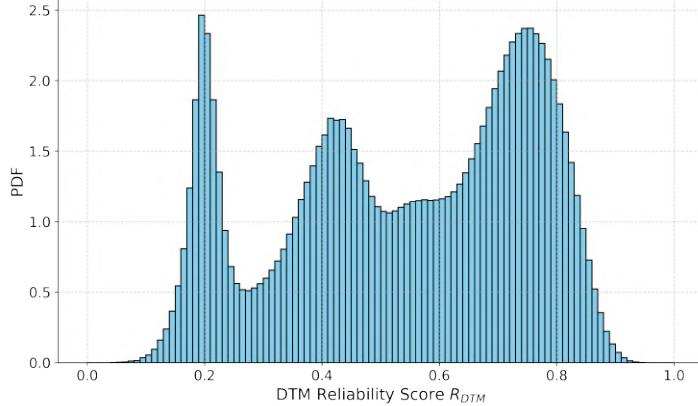


Figure 5.19: PDF of the DTM reliability scores R_{DTM} for all pixels containing at least one ground point.

the distribution, four classes were determined: (i) the group of pixels that have no ground points and therefore $R_{DTM} = 0$, (ii) low ($0 < R_{DTM} \leq 0.25$), (iii) moderate ($0.25 < R_{DTM} \leq 0.5$), and (iv) high ($0.5 \leq R_{DTM}$) reliability. A spatial map of these four classes can be seen in Figure 5.7.

5.2.5. Vegetation Cover Influences on DTM Reliability

Figure 5.20 shows the NDVI (Equation 4.12) ranges and occurrences (as a measure of vegetation cover) for each DTM reliability class shown in Figure 5.7. The first histogram in the top-left panel shows NDVI values for pixels that do not contain a ground point, and therefore have a DTM reliability of zero. This histogram serves as a representation for NDVI values in areas where, due to LiDAR limitations, ground sampling was not possible. A pronounced peak near an NDVI of 1 is visible, corresponding to extremely dense vegetation. The presence and size of this peak align with the known limitation that LiDAR signals cannot penetrate dense canopies, preventing ground detection. Aside from this peak, most pixels fall within the higher NDVI range of 0.5 to 0.9, also indicative of (dense) vegetation, as expected. However, low NDVI values near zero also occur, though less frequently. One might question how it is possible that bare areas (corresponding to an NDVI near 0) resulted in no ground sampling, and therefore a $R_{DTM} = 0$. However, these are not simply bare areas, but rather areas with little to no vegetation in steep or shadowed areas where ground sampling is hindered.

The histogram of the low DTM reliability class in the top-right panel predominantly shows high NDVI values, with the largest pixel count near an NDVI of 1. This confirms that the DTM reliability algorithm correctly identifies areas of very dense vegetation as having a low DTM reliability. In contrast, the histogram of the moderate DTM reliability class in the bottom left panel displays a broad range of NDVI values without a single dominant peak. This broad range of NDVI values suggests a mix of surface conditions, including bare ground, sparse vegetation, and dense canopy. When looking at the spatial distribution of this class (in Figure 5.7), it can be seen that areas of moderate DTM reliability is often present in places where there is also a low DTM reliability. Here, NDVI values remain relatively high, but the moderate score likely reflects partial ground visibility, which is common at forest edges or transitional zones between the forest habitats and shrubland, and grassland. Notably, the wind-sheltered outer

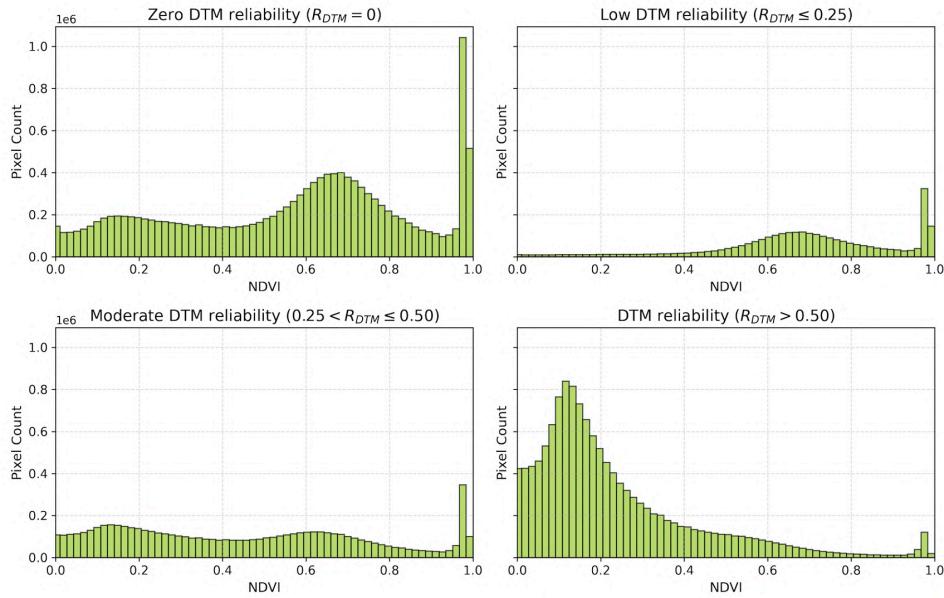


Figure 5.20: Histograms of the NDVI for each DTM reliability class in Figure 5.7. The upper left panel shows pixels that contain no ground points, and therefore have a DTM reliability of 0. The y-axis shows the pixel count, and the x-axis the NDVI. High NDVI values are associated with a dense canopy.

regions of the island (in the western part) show moderate DTM reliability as well, coinciding with more complex vegetation structures. The moderate DTM reliability class also appears along the island's cliff edges. Here, the NDVI is near zero as there is little to no vegetation growing on the cliffs. In these areas, the DTM reliability algorithm assigns a lower score because of the steep topography and unfavorable acquisition conditions.

The histogram of the high DTM reliability class in the bottom right panel shows a dense peak for low NDVI values, corresponding to areas of bare ground and sparse vegetation. These areas are easily sampled by LiDAR, and are therefore correctly assigned a high DTM reliability. However, a few pixels in this class still exhibit high NDVI values, corresponding to more densely vegetated areas. Future work should investigate the spatial extent and characteristics of these outliers, as either the characteristics of the vegetation allow adequate ground sampling despite having a high NDVI, or the DTM reliability algorithm resulted in a false positive. Overall, the NDVI evaluation supports one of the underlying assumptions of the DTM reliability: regions where the ground surface is obscured by dense canopy yield less reliable DTM values.

5.3. Assessing the Usability of LiDAR-Derived Structural Information for Land Cover Mapping

5.3.1. Vertical Structure Extraction

Normalized Height Distribution Derivation

After dividing the point cloud over the 5 m x 5 m resolution pixels, the distributions of elevations, or z value, in each pixel were normalized by their corresponding 1 m resolution DTM, as described in Sub-section 4.3.1. As a result, it was found that many pixels still contained a significant amount of negative normalized height values. This is odd because a smaller resolution DTM was purposely chosen to avoid this, as a coarser DTM results in more points being under the DTM in steep areas, leading to a negative normalized height when the DTM is subtracted. To investigate the cause of the remaining amount of negative normalized height values, it was considered whether slope effects or inaccuracies in the DTM might be responsible. Pixels from three point cloud tiles (6000_4000, 5000_3000, and 4000_2000) were analyzed at a coarser 10 m, to provide a clearer picture of the issue. The point clouds of three pixels from these three tiles can be seen in Figure 5.21. These tiles are representable for the whole

island, as the first contains cliffs and flat land, the second forested areas, and the third residential areas. The slope derived from the 10 m DTM and the mean DTM reliability are also indicated.

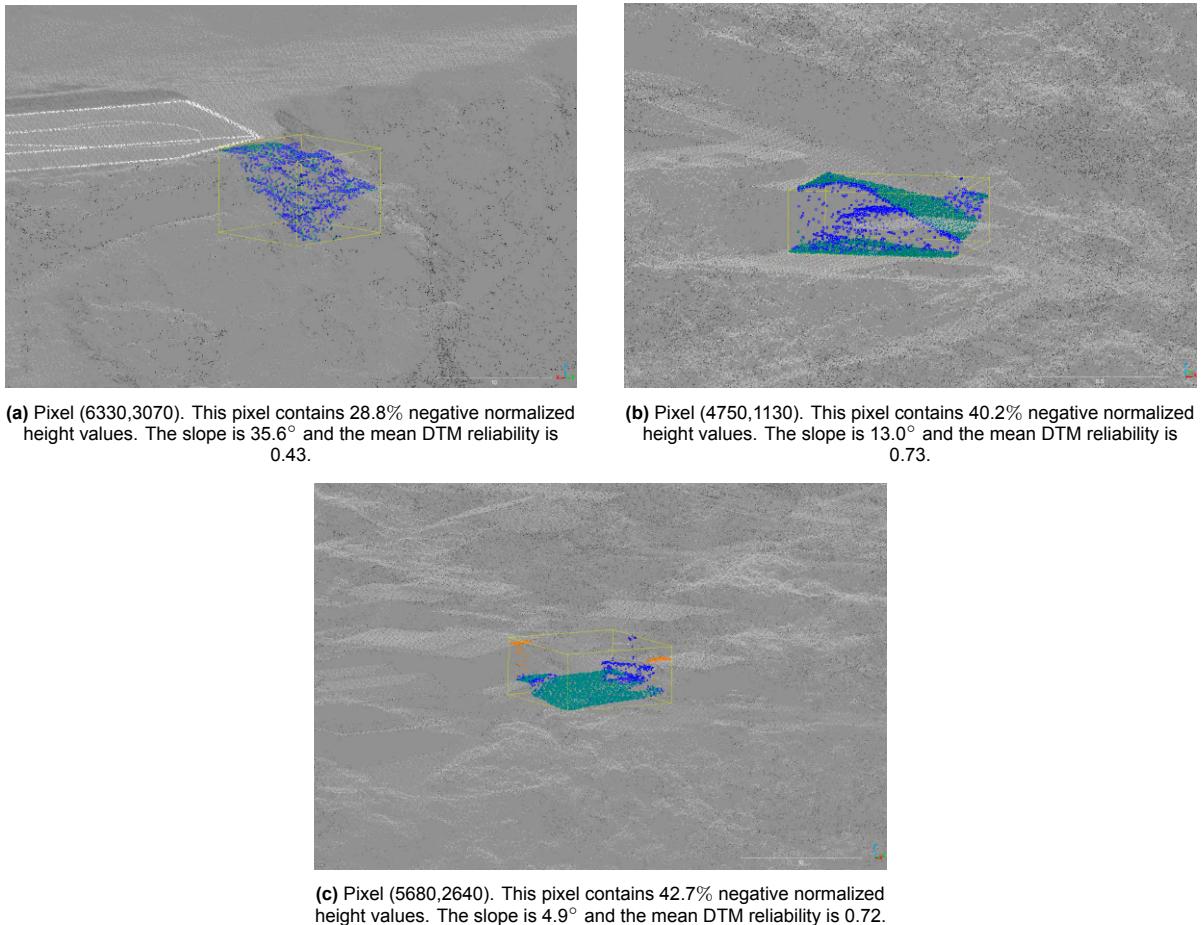


Figure 5.21: Three point clouds in an area of 10 m x 10 m. The number of negative normalized height values after normalizing with the 1 m DTM are indicated, as well as the slope at 10 m resolution and the mean DTM reliability. The point cloud is colored by classification (attribute 8 in Table 3.2: blue is unclassified, green is ground, and orange is buildings).

Figure 5.21a shows a very large slope and a relatively low percentage of negative normalized height values compared to the other two pixels. This suggests that the negative values originate from the steep slope, where some points fall below the DTM. This is also reflected in the moderate DTM reliability score.

The pixel in Figure 5.21b is located in an urbanized part of the tile, showing a high negative normalized height percentage despite its high DTM reliability score. This pixel includes a sloped road adjacent to flat vegetated terrain, which is the cause of the high percentage of negative normalized height values; if the 1 m DTM grid contains ground points from both the road and the adjacent flat ground, the vegetation points may fall below the DTM, resulting in negative normalized heights. This issue is not due to poor DTM reliability but rather the classification and elevation differences between road and vegetation.

The pixel in Figure 5.21c also contains a large percentage of negative normalized heights and a high DTM reliability, while being almost flat. It can be seen that the pixel consists almost entirely out of ground-classified points, and since the DTM represents the average z-value of these points, some will inevitably fall below the surface, producing negative normalized heights.

Besides these special cases, it was found that pixels with a low DTM reliability did not exhibit a large percentage of negative normalized heights. This is because pixels with a low DTM reliability often contain dense vegetation, which sits well above the ground and typically does not result in negative normalized heights when the DTM is subtracted. However, these pixels remain problematic, as it cannot be confirmed whether their normalized heights accurately reflect true height above ground, only the low DTM reliability offers an indication. Given all the above-mentioned observations, all negative

normalized height values were clipped to zero.

Parameterization of Normalized Height Distribution

Table 5.1 shows the pixels that were selected to determine the number of components in the GMM. Their Wageningen class code and accompanying color are also indicated.

Table 5.1: Pixel locations and their corresponding Wageningen land cover classes (Figure 2.10) with color indicators.

Pixel Coordinates	Class Code	Wageningen Class Name
(3595, 1718)	421	Forest deciduous seasonal
(3560, 1660)	461	Forest broadleaved evergreen
(3645, -260)	721	Bare rocks
(4685, 290)	323	Thorn scrub
(5385, 1680)	461	Forest broadleaved evergreen
(5905, 2745)	412	Forest dry broadleaved evergreen
(6295, 1175)	412	Forest dry broadleaved evergreen
(6435, 1135)	311	Herbaceous rangeland
(6625, 1750)	231	Pastures

The distributions of the normalized height values of the pixels can be seen in Figure 5.22. To pa-

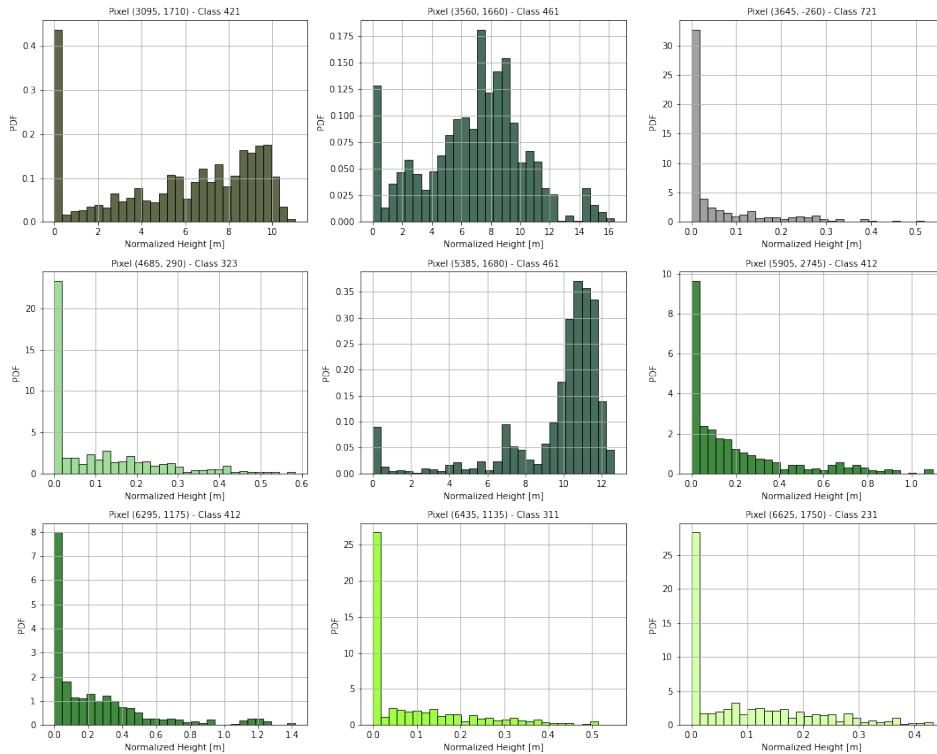


Figure 5.22: Distributions of the normalized heights for the 9 selected pixels. The distributions are colored by their Wageningen class (Figure 2.10)

parameterize these distributions, GMMs of 1 to 6 components were fitted to these distributions, and the accompanying BIC score was calculated. An overview of the BIC scores for each pixel and model can be found in Appendix J. From the BIC scores and visual inspection, a number of three components was chosen to fit the vertical profile distributions. The fit in PDF and ECDF space can be seen in Appendix K.

Figure 5.23 illustrates the parameterization of the normalized height distribution within a single 5,m × 5,m pixel. The histogram shows the distribution of LiDAR returns after DTM normalization, i.e., the vertical profile of the surface structure within this pixel. A GMM with three components was fitted to

this distribution (red curve). Each component can be ecologically interpreted as a vertical layer of the vegetation structure: Component 1 corresponds to ground and low vegetation returns, Component 2 to the understory layer, and Component 3 to the canopy. The parameters of the second component are highlighted in purple. The mean μ_2 indicates the average height of returns in this layer (the mean height of the understory). The standard deviation σ_2 captures the vertical spread of these returns, reflecting variability in the height of understory vegetation. The weight w_2 represents the relative contribution of this layer to the total distribution (the density of the understory, as it reflects the area under the curve). Since each of the three components is described by a mean, a standard deviation, and a weight, the vertical profile of this pixel can be fully summarized by $3 \times 3 = 9$ parameters (see LiDAR vertical structure parameters for $n = 3$ in Table 4.2).

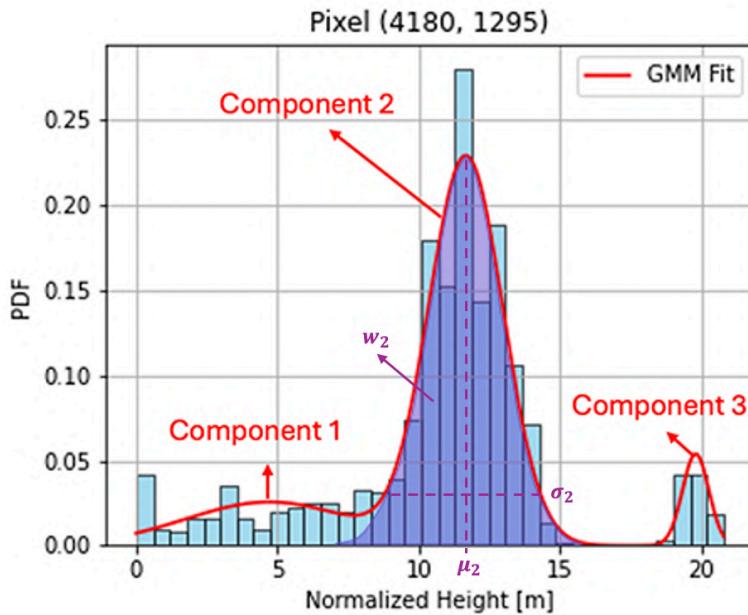


Figure 5.23: A 3-component GMM fitted on the normalized height distribution of a 5 m resolution pixel. The PDF describes the occurrences of the normalized z values of each LiDAR point that falls within the pixel. The red line indicates the GMM fit, and the 3 parameters μ_2, σ_2, w_2 for the second component are indicated.

5.3.2. Case Study Analysis

Case Study A: level-2 ecological taxonomy

For the level-2 ecological taxonomy case study, an area of 57,000 m² dominated by the broadleaved evergreen forest class from the Wageningen land cover map (Figure 2.10) was selected for analysis. This region spans from 4140 to 4320 m eastings and from 1430 to 111 m northings, comprising a total of 2,304 pixels at 5 m resolution. Within this area, 63 pixels were classified as a land cover class other than broadleaved evergreen forest. Therefore, the 5 m resolution pixels that contained these classes were excluded to ensure only broadleaved evergreen forest is assessed.

Since a GMM of three components was used to fit the normalized height distributions, each of the 2,241 broadleaved evergreen forest pixels have 9 parameters that describe their normalized height distribution: the means (μ_1, μ_2, μ_3), standard deviations ($\sigma_1, \sigma_2, \sigma_3$) and weights (w_1, w_2, w_3) of GMM component 1, 2, and 3 respectively. Figure 5.24 shows the values of the parameters for each component across all pixels. Each individual plot corresponds to one of the three GMM components with the mean height μ of the component on the x-axis, the standard deviation σ of the component on the y-axis, and the point color representing the weight w of the component. Figure 5.25 shows the normalized height distribution of 12 randomly sampled pixels, and their fitted GMM.

Component 1: ground and low vegetation layer

The first component (left panel in Figure 5.24) represents the lowest layer in the normalized height distribution of the pixels. Most of its mean height values are clustered between 0 and 5 m, and there

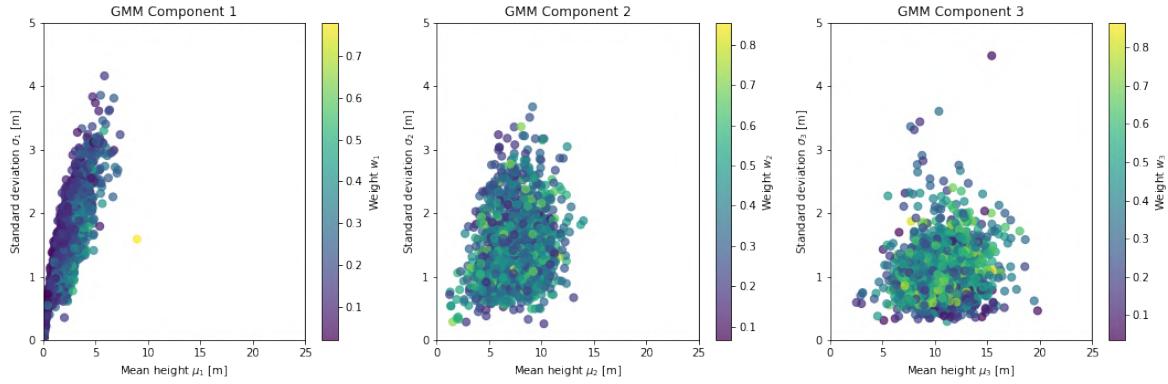


Figure 5.24: Scatter plots of the parameters from the 3-component GMM across 2,241 pixels belonging to the broadleaved evergreen forest class from Wageningen. Each subplot shows the parameters for one component: mean height μ (x-axis), standard deviation σ (y-axis), and weight w (color scale).

appears to be a linear relationship between mean height and standard deviation. This suggests that ground sampling varies across pixels. At the beginning of this relationship, there is an almost vertical line where the mean height remains low, but the standard deviation increases. These likely correspond to pixels where the ground is sampled, but the presence of low vegetation causes the first component to spread out. Moving up in the relationship, the higher the mean height, the larger the standard deviation. In this case, the ground is less sampled, and the first component mainly describes the lowest vegetation layer. This linear relationship can be seen in the top row of distributions, labeled pixel A, B, and C, in Figure 5.25. Pixel A shows a peak in its histogram at zero, corresponding to the ground. It

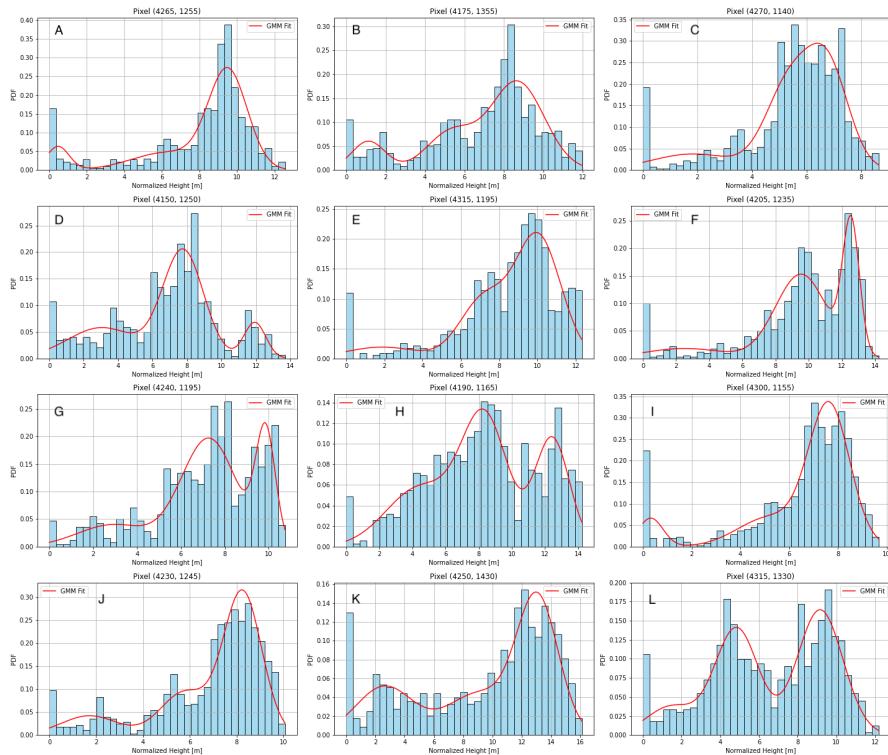


Figure 5.25: Distributions of the normalized heights for 12 pixels randomly sampled from the area and their GMM fit.

shows little low vegetation (as the first large peak occurs at a normalized height of 6 m). The mean of the first component is therefore located close to zero, with a small standard deviation. Pixel B shows a ground peak, but also contains larger peaks for low normalized height, therefore, this pixel contains more lower vegetation than pixel A. Therefore, the first component has a larger mean height and stan-

dard deviation because the first component captures the low vegetation between 0-3 m. In comparison, the histogram of pixel C also shows a peak at zero. However, the histogram is denser for larger normalized heights (indicating that this pixel contains mostly points that are 4-8 m above ground), therefore two components are assigned to describe these layers. As a result, the first component is positioned to describe a broader segment of the vegetation between 0 - 3 m, leading to both a higher mean height and a larger standard deviation for the first component.

The weights of the first component of all pixels are relatively low compared to the other components, suggesting that the upper part of the normalized height distribution is more densely populated. This is also evident from Figure 5.25, where the area under the curve of the first component is much smaller compared to those of the second or third component. This means that most of the LiDAR points in a pixel are from the canopy layer (which was also illustrated in Figure 5.2).

Component 2: mid-vegetation layer

The second component (middle panel in Figure 5.24) has mean height values mostly ranging between 5-10 m, representing the mid-vegetation layer. Its weights are generally higher than those of the first component, indicating greater density, while the lower standard deviation suggests that this layer is more clearly defined. The large range of parameter values indicates that the second component varies a lot from pixel to pixel. This variation can also be seen in Figure 5.25. For example, pixels D and H have a very clearly defined second component with the highest weight compared to their first and third components, while pixel K shows a lower weight for the second component, and is much less pronounced than the first or third component. Interestingly, all three pixels display a third component with a similar mean height, but its weight varies considerably. This suggests that these three pixels have the same canopy layer height, but pixels D and H contain more understory vegetation (higher weight for the first and second component), whereas pixel K is more dominated by taller vegetation (highest weight for the third component).

Component 3: Canopy layer

The third component (right panel in Figure 5.24) consists of mean heights ranging from 5 to 15 m, indicating that the canopy layer varies a lot in height across the pixels. It also has the lowest overall standard deviation (as also illustrated by the narrowness of the third component in pixels F and G in Figure 5.25). The weights are also generally larger for the third component, as the canopy layer is the most dense and received the most laser returns.

Overall, the scatter plots hint that there is structural heterogeneity within the class of this area, mostly in the second and third component. This is also apparent from the differing normalized height distributions of the 12 randomly drawn pixels (Figure 5.25) from the study area for this case study, although some are similar in shape like pixels A and B, and pixels I and J. The implications are further discussed in Subsection 6.1.2.

Case Study B: level-3 ecological taxonomy

The predicted coverage map of the invasive species Coralita generated in Carto can be seen in Figure 5.26. Coralita pixels with a coverage of $\geq 50\%$ are indicated in red, which sum up to a total of 392 pixels, which translates into an area of 23,378 m². The results from assessing which Wageningen land cover classes are covered with pixels of at least 50% Coralita coverage, and to which extent, can be seen in Table 5.2. Almost half of the Coralita-covered area coincides with pixels classified as "nodata", which correspond to areas obscured by cloud cover in the original dataset. After this, pastures show the highest Coralita coverage, which is also where it is known to spread most rapidly. A significant percentage of Coralita occurs in bare rock areas, which are often located within or adjacent to pastures. A large fraction of Coralita is also found in broadleaved evergreen forest, and although it can be seen from the Coralita coverage map (Figure 5.26) that it does not occur in the dense forest habitat at the center of Saba around mt. Scenery, it is instead found in small patches where these trees grow, indicating that Coralita also extends into natural habitats. Significant coverage of Coralita is also found in informal housing and roads, which is where Coralita got its naturalized status. The other forest habitats (forest dry broadleaved evergreen and forest deciduous seasonal) are also present, but in much lower occurrence. These forest classes are less abundant overall, as seen in the Wageningen land cover classification map (Figure 2.10). Overall, these results confirm that Coralita is strongly associated with disturbed land cover classes linked to human activity, while its presence is lower in natural habitats.

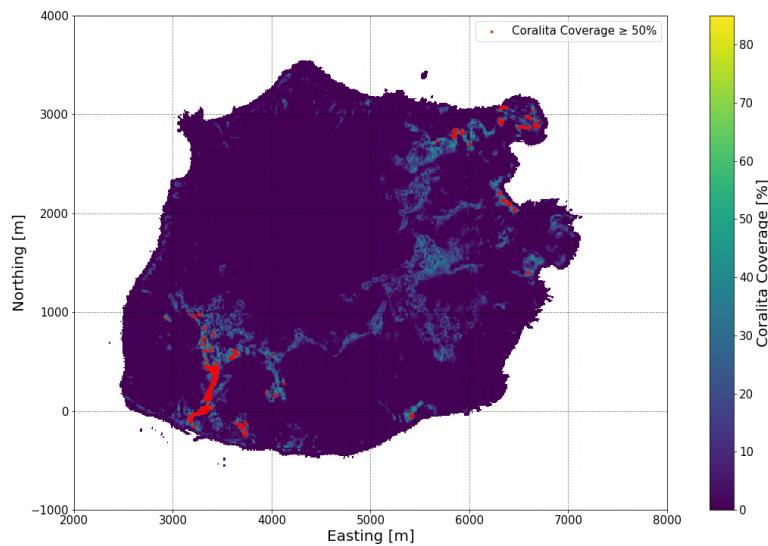


Figure 5.26: Predicted coverage of Coralita in 2024. Red pixels indicate a coverage of $\geq 50\%$.

Table 5.2: Distribution of Coralita ($\geq 50\%$ cover) across Wageningen land cover classes on Saba.

Class Code	Wageningen Class Name	Area [m ²]	Percentage of Area [%]
17	▀ Nodata (clouds)	10,620.8	45.4
231	▀ Pastures	3,441.6	14.7
721	▀ Bare rocks	2,322.6	9.9
461	▀ Forest broadleaved evergreen	2,145.3	9.2
113	▀ Informal housing	1,261.2	5.4
122	▀ Roads	1,210.3	5.2
412	▀ Forest dry broadleaved evergreen	999.4	4.3
124	▀ Airport	522.6	2.2
323	▀ Thorn scrub	511.0	2.2
712	▀ Rubble	232.4	1.0
311	▀ Herbaceous rangeland	56.1	0.2
421	▀ Forest deciduous seasonal	54.5	0.2

This pattern is consistent with previous studies, which report that Coralita spreads most effectively in degraded areas, from which it can subsequently invade natural vegetation.

Figure 5.27 and Figure 5.28 show a scatter plot of the slope vs the DTM and the distribution of aspect angles of the Coralita pixels. The scatter plot shows that most pixels are located below 250 m elevation

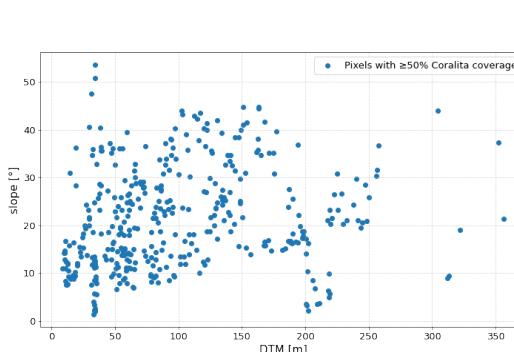


Figure 5.27: Scatter plot of the slope and DTM of Coralita pixels with $\geq 50\%$ coverage.

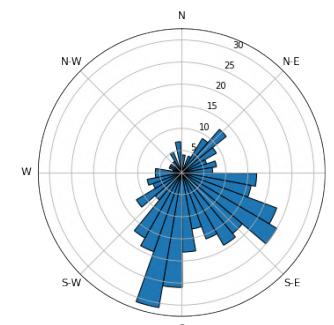


Figure 5.28: Distribution of aspect angles of Coralita pixels with $\geq 50\%$ coverage.

on slopes ranging from 10° to 40°. These findings align with the general habitat zonation of Saba mentioned in the section Habitats: Coralita is largely absent from the high-elevation Elfin and Montane forests near Mt. Scenery, and thriving in drier lowland habitats where urban areas are located. The large range of slopes indicates that it has no strict preference for slope steepness, as Coralita occurs on both relatively flat terrain, such as pastures and near houses and roads, and on steeper slopes where vegetation and rocks are present. The aspect angle distribution shows a tendency for Coralita to grow on slopes oriented from southwest to east. While this narrower aspect range could partly relate to exposure conditions such as the (dominant eastern) wind direction or sun, likely many other factors influence its spread. As mentioned in Subsection 2.3.3, Coralita can tolerate a wide range of soils, and rapidly colonizes erosion-affected sites or areas disturbed by livestock and hurricanes. Its drought tolerance and ability to exploit both natural supports (trees, shrubs) and man-made structures further expand the range of suitable habitats. This means that the observed aspect preference is unlikely to be driven by a single factor such as wind direction or sun exposure, but rather reflects a combination of environmental conditions and disturbance regimes.

6

Discussion

In this chapter, it is discussed how LiDAR-derived information can support land cover mapping on Saba. In Section 6.1, the potential of LiDAR data is assessed through both the reliability of the DTM and the case studies that demonstrate its ecological applications. This is followed by Section 6.2, which reflects on the methodological choices made in the workflow design and their implications for the robustness of the results. Finally, in Section 6.3, the societal impacts of this study, as well as the data analyzed, are discussed.

6.1. Potential of LiDAR information for Land Cover Mapping

6.1.1. The Reliability of the DTM

From the DTM reliability assessment of the AHN airborne LiDAR-derived DTM in Section 5.2, the results showed that 33% of the 50 cm resolution DTM was classified as having a high reliability. The areas of high DTM reliability are primarily concentrated in the lower-lying, outer parts of the island, mainly encompassing the dry shrubland and grassland, as well as the barren land. Because the DTM serves as the basis for many geospatial applications, it can now be used in studies that focus on the areas where the DTM reliability is high. For example, areas with high DTM reliability are also the most prone to erosion because they contain little to no vegetation. Studies on erosion risk can confidently use the DTM, as slope and elevation data are inputs for such analysis (Jagodnik et al., 2019). Similarly, hydrological modeling in the lower-lying areas of Saba can benefit from the reliable DTM, which is used as a key input for simulating water movement (Vosselman and Maas, 2010). This is particularly relevant for Saba, as it is expected that more intense rainfall events will occur under climate change scenarios (KNMI, 2023). A reliable DTM makes it possible to identify flood-prone zones, predict runoff pathways, and locate areas where water is likely to accumulate. In turn, this helps anticipate the secondary impacts of flash floods, such as erosion and slope destabilization in downstream areas.

The importance of a high-resolution and reliable terrain model for Saba has been emphasized in earlier land cover studies. In Smith et al., 2013, the available DTM contained too many artifacts to perform topographic correction, which meant that shadowed areas in the imagery could not be corrected and were therefore left unclassified. It was mentioned that a DTM with a spatial accuracy of approximately 1 m would significantly improve the land cover classification through topographic correction of shadowed areas. With the DTM from the AHN airborne LiDAR data, this limitation can be solved in areas where the DTM has been shown to be highly reliable.

Importantly, the AHN LiDAR-derived DTM also represents the most up-to-date terrain model of Saba (2024), replacing the previous DTM generated from 1991 aerial imagery photogrammetry (Mücher et al., 2014), which had never been validated and contained significant uncertainties. The LiDAR-derived DTM therefore marks an improvement in terrain data quality for Saba, providing a much more reliable foundation overall for future geospatial and ecological studies.

6.1.2. Case Study A: Level-2 Ecological Taxonomy

Although the results from the first case study in Subsection 5.3.2 only give an initial indication of structural variation within a single class of the Wageningen land cover map, they suggest a noteworthy implication: areas within the same land cover class of the Wageningen land cover map differ in their vertical vegetation structure. Future work could look at whether subgroups with similar vertical structures exist within each Wageningen class. If these subgroups exist, this could enable the identification of patches with comparable vegetation structure. In practical terms, this means two things: (i) that the Wageningen land cover map can be more detailed by integrating the vertical structure parameters examined in this study, and reveal how each class is structured; for instance, where the broadleaved evergreen forest is dominated by tall, closed-canopy vegetation or lower and more patchy growth. (ii), Areas dominated by species with similar canopy heights could be distinguished, for example, the dominant tree *Freziera undulata* in the Elfin forest that has a known height of 9 m.

A natural next step is to apply an unsupervised classification method to the nine parameters derived from the GMM fits. The parameters are used as input features for the unsupervised classification method, which will make groups of pixels that show similar vertical structure features. This could, for example, separate denser patches of forest from more open understory-dominated areas within the same Wageningen class. While this case study tested only a small area, the same approach could be applied island-wide, enabling a spatially detailed map of structurally similar patches. Combined with topographic parameters in Table 4.2, these patterns could provide new insights into how vegetation structure is distributed across Saba's complex landscape. A limitation arises from the fact that this analysis relies on the Wageningen land cover classification map to select pixels of the same class, which is likely to have changed over time, as this map is generated from 2010 data. Therefore, a classification method can be applied solely to the LiDAR data.

Of course, unsupervised classification has its limitations, since the ecological meaning of the groups with similar vertical structure features depends on the chosen algorithm and parameters. Therefore, the importance of training data is highlighted. With the use of training data, the full set of parameters identified in Table 4.2 can be used as input features in supervised classification, allowing structural patterns to be directly linked to vegetation types that have been validated in situ. Although using multispectral data on Saba alone limited the performance of classification in the Wageningen study, the combination of both LiDAR and multispectral data on Saba has not been explored yet. In many other studies, combining structural information from LiDAR with spectral information from multispectral sensors has significantly improved classification accuracy (Fassnacht et al., 2016b).

6.1.3. Case Study B: Level-3 Ecological Taxonomy

The results from this level-3 ecological taxonomy case study show that the topographic parameters extracted from the LiDAR data (Table 4.2) could be used to uncover the topographic conditions of the invasive species Coralita, which were found to be consistent with literature. Several opportunities for future work can be explored following this case study.

First, the coverage map is based on observations from 2017. As noted by the developers of Carto, adding observations from multiple years improves prediction stability and helps capture growth patterns. Since seasonal conditions and extreme events strongly influence Coralita's growth patterns, its growth pattern is not likely to be the same from year to year. Therefore, the training dataset should ideally be supplied by more observations spread over multiple years. Future work could therefore incorporate more observations spread over multiple years. These observations can for example be taken from biodiversity platforms such as *Observation.org*, where Coralita has been recorded in recent years (Observation.org, 2025b). By including this type of data, prediction maps would not only show where Coralita was present in 2017, but also where it is establishing and expanding today.

Second, the topographic parameters used in this case study (Table 4.2) are derived from a single LiDAR dataset from 2024. This means that temporal variations in Coralita's distribution cannot be captured. For example, suppose removal or management efforts were carried out in certain areas in earlier years. In that case, these areas may now appear as locations where Coralita is re-establishing, rather than as long-standing infestations. In the future, repeated LiDAR acquisitions would enable temporal analysis, which would be especially valuable for monitoring the dynamics of invasive species.

Third, future work can also include the analysis of multiple coverage thresholds in the predicted Coralita coverage map (Figure 5.26). In this case study, the threshold was chosen as 50% to ensure that Coralita is dominant within the pixel. Assessing a lower threshold range, for example, up to 10%, could provide early warning signals about areas of onset invasion where Coralita has just begun to establish, for example after an extreme event such as a hurricane, after which Coralita re-establishes quickly. The detection of areas with onset invasion enables removal efforts before the species has a chance to spread extensively and harm native vegetation. Acting at this stage not only prevents further expansion but also reduces the overall cost and effort required for removal efforts.

Finally, case studies like this are not limited to Coralita. Provided that species observations are available and its topographic preference is restricted to areas with a reliable DTM, the same workflow can be applied to other ecologically important or invasive species; for example, the other problematic invasive grass species *Bothriochloa pertusa* (Subsection 2.3.3), which is located in the dry shrubland and grassland where DTM reliability is high.

6.2. Workflow Design Impacts

6.2.1. Reflection on the DTM Quality Assessment

Excluded LiDAR Reliability Variables

During the selection of variables that are indicative of the DTM reliability, several variables were considered but eventually excluded after evaluation and are therefore not presented in Table 4.1. These variables include the mean of the local incidence angle θ_0 and its standard deviation. The reasoning for considering this variable was that near-nadir incidence angles should, in theory, maximize the probability of recording a strong ground return, as the laser beam penetrates the canopy most directly when it aligns with the surface normal. However, upon closer inspection, several limitations emerged regarding these variables.

Firstly, the local incidence angle was derived from slope and aspect maps, which themselves were calculated from the DTM of this data. Including this variable would therefore introduce a degree of circular reasoning, as the DTM reliability score would then partly depend on the surface it is supposed to evaluate. Secondly, the distributions of the mean and local incidence angle did not align with expectations, as very few points had a local incidence angle close to zero, suggesting that the assumption that a zero-degree local incidence angle is favorable for DTM reliability is too simplistic. Instead, the distribution showed a large range of local incidence angles. The likely cause for this is that vegetation characteristics such as leaf orientation strongly influence the angles under which optimal penetration of the signal can occur, and therefore deviate from 0. Due to these underlying factors and the complexity of the relationship, the local incidence angle was excluded from the selected set of variables. More information about the definition of the local incidence angle and how it's calculated in practice can be found in Appendix L.

HALM Derivation Analysis

Following the special cases for which the HALM could not be calculated in a pixel (described in Subsection 4.2.1), case 2 (the pixel contains one ground-classified point and the lowest point is not a ground-classified point) was further looked into. It was calculated for one tile (tile 3000_1000) how often this case occurs, as this case is suspicious: if the lowest point in a pixel is not classified as ground, it suggests that the ground classification algorithm has either missed the lowest ground return or incorrectly labeled vegetation or other objects as non-ground, leading to a less reliable DTM. The tile has 2,505,378 pixels with at least one ground point, and out of those, 9.8% show the considered case (the lowest point in the pixel is not classified as ground). In figure Figure 6.1, these pixels are indicated in red. It is noticeable that they seem to occur more frequently in vegetated areas, which probably contain a large slope. In these areas, the lowest point likely corresponds to vegetation, while the ground-classified point is located higher up the slope. The fact that this results in a higher HALM and subsequently a lower DTM reliability score (as a larger HALM is associated with lower DTM reliability) is justly, as in these cases the LiDAR-derived DTM would show a higher elevation than the actual ground elevation (since the only ground-classified points are higher up the slope of the pixel, while the accurate DTM represents the average of this slope).



Figure 6.1: A visual overview of pixels for which the lowest elevation of the LiDAR point within that pixel is not classified as ground (indicated in red), zoomed in on an area overlaying the RGB image.

Path Length Derivation Analysis

To address the issue in the derivation of the path length (discussed in Subsection 4.2.1), where the path length could not be calculated because the first return of a pulse lies in one LiDAR tile while the last return lies in an adjacent tile, the code to extract the pulse ID for each LiDAR point was run for a single tile (tile 6000_2000). This problem affected only 0.006% of the ground points in that tile. To check how splitting the tile into chunks affects this issue, as chunking was necessary to process tiles with a significant number of points ($>20,000,000$), the code was rerun for the same tile, now including chunking, and the number of ground points without a first return in the data chunk went from 0.006% to 0.01%. Still, this amount is negligible, considering the large amount of points in each tile, and the fact that variables are calculated on a pixel basis, where the absence of one point can still result in a valid pixel.

A Priori High DTM Reliability Group

From the a priori high DTM reliability map shown in Figure 5.13, it can be seen that the group of pixels that are assumed to have the highest DTM reliability (indicated in black) is relatively small. The group accounts for only 6.5% of all pixels containing at least one ground point. The small size of the group is a result of the strict constraints imposed on this group to ensure that the group represents truly unobstructed ground, for which the DTM quality solely relies on its LiDAR accuracy. Given the small size of the group, one could consider relaxing some of these constraints to increase the group size and thereby include more diverse terrain characteristics that could also possibly yield high reliability. For example, the ratios could have been chosen to deviate slightly from 1, or one could permit minimal path lengths, both corresponding to pixels that are *nearly* unobstructed. This way, pixels that are almost unobstructed or have very little vegetation are also included in this group. This would have shaped the Logistic Functions to also score pixels under these conditions slightly higher, as the parameters of the Logistic Functions are based on the distribution of the a priori high DTM reliability group. However, relaxing the constraints would mean deciding how much “imperfection” can still be allowed while calling a pixel highly reliable. Right now, it’s straightforward: for example, if the path length is 0, the pixel is completely unobstructed. If we also want to include pixels that are almost unobstructed (with a small but nonzero path length), then a cutoff needs to be set. In other words, one has to decide: up to what path length can a pixel still count as having a “high DTM reliability”? This is not an easy threshold to define. Importantly, since the final classification assigns pixels to low, moderate, or high DTM reliability, relaxation of the constraints is unlikely to significantly affect the overall scoring outcome. Nevertheless,

future work could investigate the constraints of the a priori high DTM reliability group and their influence on the sensitivity of the final score.

To get an idea of how the variables listed in Table 4.1 capture different aspects of DTM reliability, the Spearman's squared rank correlation coefficient (Equation 4.5) was calculated for the a priori high DTM reliability group. These include all variables listed in Table 4.1, except the ratio of first returns classified as ground out of all first returns R_{first} , the ratio of the ground-classified points out of all returns R_{ground} , and the mean path length l_μ . As these variables were used in defining this group, they show no variance since they all have the same value. Among the remaining variables, the correlation was extremely low (ranging from $0.01 \leq \rho^2 \leq 0.07$). The low correlations within this group indicate that each variable adds unique information about DTM reliability. As a result, the reliability algorithm can take into account many different aspects of what makes an area highly reliable.

Dimensionality Reduction Method

A commonly used method in dimensionality reduction is Principal Component Analysis (PCA), where a set of correlated variables is ordered according to how much variance they explain, where the first "principal component" explains the most variance (Salem and Hussein, 2019). This method was considered in the dimensionality reduction step of the DTM reliability algorithm (Subsection 4.2.3), but was rejected because the variables in Table 4.1 violate some assumptions of PCA. PCA assumes the following: the variables are continuous and normally distributed, and the variables should all have a linear relationship. PCA can still be performed if some variables lightly violate these assumptions, but it introduces uncertainty in the principal components and requires careful interpretation. It is already known that at least one of these assumptions is violated, as the ratio and point density variables in Table 4.1 are discrete. Therefore, a method of evaluating pairwise associations between variables using the squared Spearman's rank correlation coefficient is preferred. Many statistical methods, including Spearman's rank correlation, assume that the data is independent and identically distributed for statistical inference. Independent means that no data point influences the value of another data point, and identically distributed means that each data point comes from the same underlying distribution. In spatial data, such as the data used in this study, these assumptions are violated, as neighboring pixels are spatially autocorrelated and different environments on the island follow distinct distributions. However, the proposed method only uses Spearman's correlation coefficient descHAL respectively, therefore, the i.i.d. assumption is not required.

The importance of the Logistic Function $S(x)$

As mentioned in Subsection 4.2.4, the Logistic Function $S(x)$ plays an important role in the scoring process of the DTM reliability by reducing undesired effects from the ECDF, where similar values can vary greatly in score if they represent a dense part of the distribution. This is clearly visible in Figure 5.16. Since the PDF of the mean HALM H_μ is highly skewed and narrow (Figure H.1j), the ECDF shows a sharp decrease over the small region where most values are concentrated (between 0 and 0.2 m). Consequently, small differences in mean HALM H_μ result in large differences in ECDF scores, for example, a mean HALM of 0.05 m yields an ECDF score of approximately 0.7, while a mean HALM of 0.15 m yields an ECDF score of about 0.4. In the tail region ($H_\mu > 0.6$), only a few values are spread out over a wide range of H_μ , resulting in an ECDF that changes much more slowly (for example, a mean HALM of 0.4 m yields approximately the same ECDF score as for a mean HALM of 0.6 m). By contrast, the Logistic Function S , smooths this undesired effect, because it decreases less abruptly.

A similar observation can be seen in the mean path length l_μ (Figure 5.18). Here, the ECDF is used directly as the scoring function, since no distribution of the a priori high DTM reliability group exists for this variable, as it was used as a constraint. A large group of pixels has a mean path length of zero, and by definition of the ECDF, these receive a score of 1. Because the group of zero path length is quite large, the next smallest value is assigned a much lower score of ≈ 0.4 , as can be seen in Figure 5.18. The fact that pixels with $l_\mu = 0$ receive a variable score of $R = 1$ is not problematic in itself, as this group is defined to have the highest reliability. The steep drop between l_μ and the next smallest path length is more concerning. However, pixels with path lengths close to zero are still likely to receive high scores on the other variables. Furthermore, the mean path length contributes only 1/6 to the final score R_{DTM} , so the overall effect of only using the ECDF for the scoring of this variable is expected to be small. For future work, it may be useful to test how sensitive the final score is to the choice of

solely relying on the ECDF for the mean path length by using a weighted average in calculating R_{DTM} (Equation 4.11) instead of weighing all variables equally.

Another important role of the Logistic Function lies in the fact that the ECDF is bounded between $0 \leq F(x) \leq 1$, whereas the Logistic Function ranges between $0 < S(x) < 1$. Consider the ground point density D_{ground} as an example: according to the ECDF (see Figure 5.17, the pixel with the highest ground point density (which is 17 pts/0.25m²) will receive $F(D_{ground}) = 1$, and the pixel with the lowest nonzero ground point density (which is 1 pts/0.25m², since pixels with a ground point density of 0 are excluded) will receive $F(D_{ground}) = 0$. This is undesirable, as we do not want to classify a pixel as 100% unreliable in the absence of validation data. The Logistic Function $S(x)$, however, increases the score for a ground point density of 1, making its final variable score $R(x)$ nonzero.

Discreteness of the Point Densities D_{all} and D_{ground} and Their Influence on the Steepness Parameter k

As mentioned in Subsection 4.2.4, the steepness parameter k used in the Logistic Function (Equation 4.8) is multiplied by 2 for the discrete cases of D_{ground} and D_{all} . For example, the ground point density of the whole dataset ranges from 1-17 points per pixel, and from 2-17 in the a priori high DTM reliability sample. Because of this limited range and the discrete nature of the variable, applying the default calibration for k (Equation 4.9) leads to a relatively flat curve of the Logistic Function, as the increase from 0.5 at x_{10} to 0.95 at x_{90} is spread too gradually across the small set of values. As a result, adjacent densities (for example, 2 vs. 3 points per pixel) are mapped to nearly identical scores, which means that the high point densities in the upper tail region in the ECDF are weighted down, and lower point densities are going to be weighted up by the Logistic Function quite a bit. To counter this effect, k is multiplied by a factor of 2, making the Logistic Function steeper. This way, pixels with a point density of 1 get a nonzero R , and higher densities in the upper range are not significantly downweighted compared to the ECDF.

Sensitivity of the Final Score R_{DTM}

The sensitivity of the final DTM reliability score depends on the choices made during the scoring process, especially in choosing the parameters for the Logistic Function. A deliberate choice was made to set the midpoint x_0 of the Logistic Function (see Equation 4.8) to the 5th or 95th percentile of the a priori high DTM reliability group, and to set k such that $S(x)$ increases from 0.5 to 0.95 over the distance between the 95th and 5th percentiles (as described in Subsection 4.2.4). The rationale for using the 5th and 95th percentiles, rather than a lower or higher percentile, was to mitigate the influence of extremes, since many of the variable distributions are skewed. If the midpoint were, for example, set at the 1st or 99th percentile, you risk anchoring the midpoint in an extreme value. Take the mean ground intensity, for example. In the current method, the 5th percentile of the a priori high-reliability group was taken as the midpoint x_0 . If instead the first percentile were used, the midpoint corresponds to an extreme value that can be unusually low. This would cause the Logistic Function to give higher scores to pixels with very low intensity values, where high reliability may not be justified. The opposite may happen when the midpoint x_0 is set at a more central percentile of the a priori high DTM reliability group, say the 10th or the 90th. As a result, the Logistic Function shows its sharpest increase over a smaller range (associated with higher reliability) of values from the a priori high DTM reliability group, which could result in the unjustly down-weighting of values in the whole dataset.

The same trade-off applies to the choice of the steepness parameter k . With the chosen setup, the majority of values in the high reliability distributions are spread between $S = 0.5$ and $S = 0.95$. If the distance were extended, the curve of the Logistic Function would flatten. Because of the symmetry of the Logistic Function, this would also flatten the lower half of the curve, spreading many pixels over the range from $S = 0.05$ to $S = 0.5$. By contrast, shortening the distance steepens the curve, which could overemphasize small differences in variable values, much like the ECDF, which is undesirable.

To test how sensitive the final classification of the DTM reliability is to these parameter choices, an alternative setup was tested with $x_0 = x_{10}$ or $x_0 = x_{90}$ (depending on if higher or lower values for a variable are associated with a higher DTM reliability) and k chosen such that S increases from 0.50 to 0.90 over the distance between the 10th and 90th percentiles of the a priori high DTM reliability group. With these parameters, the resulting distribution of DTM reliability classes was 9% low reliability, 17%

moderate reliability, and 32% high reliability. Compared to the original findings (see Section 5.2), it was found that a total of 2% of the island changed in DTM reliability class compared to the original setup. A change of 2% in classification is relatively small, especially given the complexity of the environment and the multiple variables feeding into the reliability score R_{DTM} . This change also gives a sense of uncertainty of the DTM reliability score. It is apparent that pixels at the boundary of the DTM reliability classification thresholds are the most sensitive to parameter choices; for example a pixel with $R_{DTM} = 0.54$ is now classified as having a high DTM reliability, since the threshold for high reliability is $R_{DTM} = 0.5$, but due to classification uncertainty, it could instead belong to the moderate reliability class. These cases should therefore be interpreted accordingly. Future work could quantify this uncertainty more explicitly by testing a wider range of parameter configurations.

DTM Reliability Classification

The PDF of the DTM reliability scores R_{DTM} (Figure 5.19) of the island shows three dominant peaks and an additional smaller peak between the two right peaks, resulting in four distinguishable groups. These groups can therefore be interpreted as relatively low to high reliability with respect to each other by using thresholds. To define a threshold for the high reliability class, it can either include or exclude the intermediate peak around $R_{DTM} \approx 0.58$. This intermediate peak likely corresponds to pixels that have very favorable conditions for most variables that were used to create the DTM reliability score, while one variable has less favorable conditions. The choice for including or excluding the intermediate peak in the threshold for the high reliability class can be supported by assessing the R_{DTM} distribution of the a priori high DTM reliability group, which can be seen in Figure 6.2. Since this group was

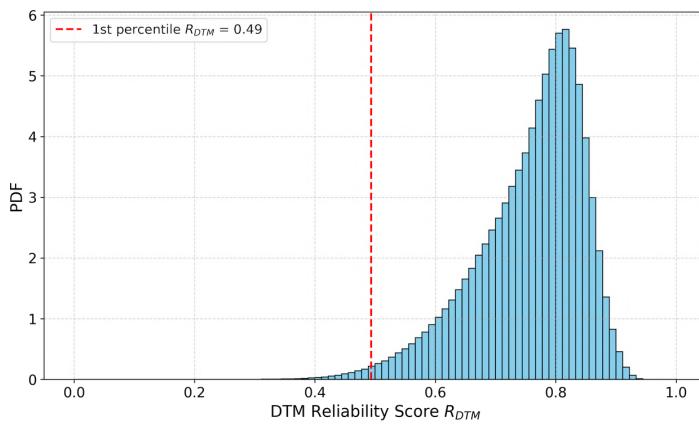


Figure 6.2: PDF of the DTM reliability scores R_{DTM} for the a priori high DTM reliability group.

defined to represent the most reliable conditions, their distribution of DTM reliability values provides a benchmark for the high reliability class. At the first percentile of this distribution, $R_{DTM} \approx 0.5$. By setting the high-DTM reliability threshold at $R_{DTM} = 0.5$, 99% of the a priori high DTM reliability pixels are included in the high-DTM reliability class. This ensures consistency with the definition of the a priori high DTM reliability group. At this threshold, the intermediate peak in the distribution of the full dataset falls mostly in the high reliability class. There is no equivalent benchmark for distinguishing between low and moderate reliability. The threshold was therefore set at $R_{DTM} = 0.25$, corresponding to the dip between the first two peaks in the PDF. This introduces uncertainty in DTM reliability classification in the region where there is overlap in the peaks. For example, a pixel with R_{DTM} just below 0.25 might still share similar characteristics with pixels just above the threshold. This reflects the probabilistic nature of the scoring, as R_{DTM} does not impose a strict categorical boundary but rather expresses the relative strength of evidence for reliable ground representation.

6.2.2. Reflection on the Vertical Structure Extraction

Parameterization of Normalized Height Distribution

To parameterize the distribution of the normalized height per pixel (as described in Subsection 4.3.1, two common approaches were considered: binning the normalized height distributions into fixed height

intervals and using the percentage of points per bin as parameters, or fitting a curve to the distribution, using the curve's parameters to describe its shape. Binning has the advantage of being non-parametric, therefore requiring no assumption about the normalized height distributions. However, a limitation is that the choice of bin size is important, as the boundary between bins is not smooth, which can lead to inconsistencies, especially in pixels where the DTM is less reliable. For example, if the DTM is reliable, the densest part of a canopy of one tree species falls in the same bin for two pixels containing the same species. However, if the DTM is less reliable for one pixel, the densest part of the canopy might fall in an adjacent bin because the DTM under- or overestimates the ground elevation, which shifts the normalized heights of the tree species upward or downward.

Another limitation of binning is that the number of bins needed depends on the height of the tallest vegetation, which can eventually lead to a large number of parameters, as each parameter corresponds to a height interval.

Considering the diverse mix of vegetation species on the island that all show distinct normalized height distributions, and the smallest vertical resolution needed to differentiate vegetation types (mentioned in the User Requirements), the binning method would likely result in a high number of parameters, which defeats the purpose of reducing computational complexity. Therefore, curve-fitting as proposed in Subsection 4.3.1 by using a GMM is the preferred approach.

Choice of Number of GMM Components

A GMM with three components was selected to fit the normalized height distributions discussed in Subsection 4.3.1. Although the BIC scores of the GMMs presented in Appendix J reach their lowest values for models with 4, 5, or 6 components, these are likely overfitting the data. From an ecological perspective, the rangeland and bare rock distributions reveal one dominant component centered around 0 (as seen in Table 5.1), corresponding to the ground, and a second, smaller component representing low vegetation or rocks, indicating that these profiles can be modeled with two components. The fact that the BIC scores indicate that more components are needed comes from fluctuations in the distribution that can be modeled with additional components, but ecologically, these fluctuations come from the same component. Overall, the BIC scores show their biggest drop between 1 and 3-4 components, indicating that 3 or 4 components should be a good fit. To minimize model complexity, a GMM with 3 components was considered first. As shown in Figure K.2, the CDF of the 3-component GMM generally follows the empirical CDFs. However, in the forest distribution of pixel (5385,1680), the peak near zero is not captured (see Figure K.1), as the first component is not centered around 0, unlike in the other classes. This did not change when fitting a GMM with 4 components. Considering this, and the fact that the rangeland and bare rocks distributions appear well-modeled with only two components, a GMM with three components provides a balanced compromise between underfitting and overfitting across the different land cover types, while being ecologically interpretable.

6.3. Societal Impacts of the Study

As discussed at the beginning of this study, progress in conservation and island management efforts on the Caribbean Netherlands depends heavily on funding. Although Saba, St. Eustatius, and Bonaire are among the most vulnerable to the impacts of global climate change, they are excluded from major climate funding mechanisms. For example, under the Paris Agreement, the Netherlands has pledged to support climate financing for developing countries. However, this commitment primarily applies to its European territory. As a result, the Caribbean Netherlands cannot claim access to these climate funding mechanisms, since they are not part of the European territory. More importantly, they are not classified as developing countries, given their status as special municipalities of the Netherlands and the fact that the European part of the Netherlands is a highly developed nation. To get funding, the Netherlands should either directly provide the funding from its own national budget, or negotiate the access of these islands to climate funding bodies (DCNA, 2024). This dependency is particularly complicated given the large differences in climate, geography, and biodiversity between the European Netherlands and the Caribbean Netherlands. As a result, the success of conservation efforts on Saba and the other islands is directly tied to the political will and financial commitment of a government operating in a completely different environmental context. This study shows how investing in modern mapping technologies, such as airborne LiDAR and multispectral ORS, can pay off in ways that directly benefit the islands and their communities.

One of the clearest societal benefits arises from the newly developed LiDAR-derived 50 cm resolution by AHN. This DTM is the most accurate and detailed terrain model available to date for Saba, replacing an outdated and unvalidated photogrammetric product from 1991. Beyond its scientific relevance, such a dataset provides a solid basis for practical applications in risk assessment. The DTM allows improved modelling of slope, runoff, and flood accumulation, thereby helping authorities to identify flood-prone or erosion-sensitive areas in advance.

This study also showed that the airborne LiDAR data likely has strong potential for improving existing land cover maps of Saba like the Wageningen land cover map from Smith et al., 2013. While a full land cover classification was not carried out here, the analysis revealed that LiDAR can capture differences in vegetation structure within areas that were previously grouped into a single class in the existing land cover map Figure 2.10. This means that LiDAR data could help create a more detailed land cover map in the future. In addition, the topographic information from the LiDAR data can be linked to the habitat preferences of certain important species that require monitoring (as seen in Table 2.2). This enables the direct connection of land cover information to species conservation, guiding practical biodiversity management on the islands.

Another impact following the island-wide airborne LiDAR acquisition is that its DTM can serve as a reference dataset for validating global, satellite-derived DTMs. Missions such as Global Ecosystem Dynamics Investigation (GEDI) (Earthdata, 2025) provide global coverage but at a much coarser resolution of about 25 m and can therefore not capture small islands in enough detail. By comparing these global DTM maps with the new airborne LiDAR-derived DTM, the Caribbean Netherlands can contribute valuable validation data that improves these global products.

While this study focused on Saba, the same LiDAR and passive multispectral ORS campaign also covered St. Eustatius and Bonaire, meaning the benefits extend across the Caribbean Netherlands and help meet national environmental reporting obligations under the NEPP (Ministry of Agriculture, Nature and Food Quality, 2020). Given the clear advantages, it would make sense for the Netherlands to expand similar acquisition efforts to include Aruba, Curaçao, and Sint Maarten. Additionally, these islands currently lack high-resolution multispectral data on freely accessible platforms such as the NSO, unlike the Caribbean Netherlands. By making detailed LiDAR and multispectral data available across the entire Dutch Caribbean, authorities, researchers, and local communities will gain the benefits of high-resolution datasets, similar to those already available in the European Netherlands. With these data, they would have far more tools at their disposal than relying solely on traditional field observations or outdated remote sensing data, which have previously proven insufficient for monitoring and managing complex island environments.

7

Conclusion and Recommendations

This chapter addresses the main research question: *What is the potential for mapping land cover of the Caribbean Netherlands using airborne LiDAR and passive optical data?* by providing a summative answer to the sub-questions presented in Section 2.8. Recommendations for future research are provided in Section 7.2.

7.1. Conclusion

1. *What land cover taxonomies are relevant for biodiversity monitoring on Saba?*

The monitoring priorities for Saba focus on tracking how species (particularly endangered, key and otherwise important species) and habitats are changing in their size, location, and overall condition. These changes are largely driven by three main biodiversity threats: climate change, roaming livestock, and invasive species. Monitoring is also needed to evaluate whether conservation measures, such as reforestation and livestock control, are having the intended effect. To ensure adequate monitoring on Saba, a land cover taxonomy needs to be defined that aligns the monitoring priorities and the data required to address them. The taxonomy system was designed by considering both the spatial scale (from island-wide patterns to specific species) and the time scale (from long-term habitat changes to rapid events such as hurricane damage) at which ecological processes occur. This led to a 3-level hierarchical classification system, presented in Table 2.2. This taxonomy system only considers terrestrial environments and focuses on flora, as the conditions of fauna depend directly on the conditions of flora. The three levels in the taxonomy system describe the following:

- Level 1: This level distinguishes broad land cover types: artificial surfaces, rangeland, forest land, and barren land. This enables the monitoring of trends such as deforestation, reforestation, or land degradation.
- Level 2: This level splits level 1 into habitats identified in the state of nature reports (A.O. Debrot, Henkens, and Verweij, 2018; A. Debrot et al., 2025) so that the size, condition, and distribution of the habitats can be assessed. One of the habitats, for example, is the Elfin forest, for which its presence is at risk of extinction under climate change projections.
- Level 3: This level focuses on species-level indicators, including both invasive species (such as *Antigonon leptopus*, or “Coralita”) and endangered species (such as *Nectandra krugii*). These species provide early warning signals about habitat health.

2. *What are the properties of the AHN LiDAR dataset?*

A First qualitative analysis (presented in Section 5.1) of the AHN LiDAR data was carried out to get an initial understanding of how well different land cover types on Saba are represented. To achieve this, three representative LiDAR point cloud tiles were selected from the data that reflect the full range of land cover conditions one might encounter. The tiles include the airport strip located on flat terrain, the peak of mt. Scenery covered by dense vegetation, and a residential area (“the Bottom”). In each

tile, their point cloud attributes were assessed. The point cloud attributes (presented in Table 3.2) correspond to (.insert general definition of what point cloud attributes are..). The analysis showed that man-made structures are very well captured in the data: buildings and roads reflect the laser strongly, making them easy to recognize. Vegetated areas, however, were mostly represented by their canopy, while the forest floor was underrepresented. In fact, in two of the densely vegetated plots in the LiDAR point cloud that were assessed, only 4.9% and 3.3% of emitted LiDAR pulses over the plot reached the ground, highlighting the difficulty of sampling under dense canopy.

The next part of the analysis looked at how the flight pattern of the aircraft the LiDAR system was mounted on, and the island's landcover affected how much detail the LiDAR system was able to capture. This detail is measured through the point density (the number of LiDAR points collected per square meter). Across the island, the average point density was 26 points/m², but values ranged widely from 1 to 228 points/m² depending on location. The value of adding a cross line to the 7 parallel flight lines was evident: the point density under this line is overall higher. Higher densities were generally found on exposed ridges, while low densities occurred along steep slopes and coastal areas where the terrain blocked the sensor's line of sight. This variation reflects the challenges of flying LiDAR surveys over rugged volcanic terrain.

The acquisition challenges on Saba are also visible in the published DTM, which is derived from LiDAR points that were classified as ground. The 50 cm resolution DTM (Figure 3.3) captures the ground well in open areas, but contains gaps in its center where the forest habitats are located. This shows the challenge of LiDAR in vegetated areas: when the vegetation is too dense, the signal dissipates before it can reach the ground. In the European Netherlands, AHN specifications require a minimum percentage of ground returns in vegetated areas to ensure reliable DTMs. On Saba, no such requirement was imposed, and the dataset falls short of this standard. As a result, while the LiDAR data provides excellent structural detail for open terrain and built-up areas, its representation of ground elevation under dense forests is incomplete. This means that the dataset is highly useful for mapping surface features and vegetation canopies, but caution is needed when relying on the terrain model in forested habitats.

3. What is the quality of the AHN LiDAR-derived DTM?

Following the qualitative analysis of the LiDAR data in sub-question 2, it was found that the DTM is of uneven quality. To better understand and quantify how much of the AHN LiDAR-derived DTM can be trusted, a DTM reliability algorithm was created (see Figure 4.2 for a schematic overview). Because no field measurements are available to directly validate the DTM, this algorithm provides a practical way to estimate reliability using the LiDAR data itself. The algorithm combines several LiDAR-derived reliability variables (presented in Table 4.1) that can provide an indication of how well the ground was captured in each 50 cm × 50 cm pixel of the DTM. Since not all variables are equally informative (as some variables might describe the same phenomena), a dimensionality reduction step was used to ensure the remaining variables provide unique information about the DTM reliability, after which the remaining variables were combined into a single reliability score R_{DTM} . The algorithm follows four main steps:

1. **Identification of LiDAR reliability variables:** such as ground point density and traveled path length of the LiDAR signal.
2. **Define an a priori high DTM reliability group:** in the absence of validation data, a group is identified that should have the highest DTM reliability: bare, open ground.
3. **Dimensionality reduction:** LiDAR reliability variables that overlapped in information were filtered out.
4. **Reliability scoring:** each pixel in the DTM was assigned a DTM reliability score based on how its characteristics compared to other pixels, and how closely it matched the characteristics of the priori high DTM reliability group, resulting in a map of zero, low, moderate, and high reliability zones.

Figure 5.7 shows the DTM reliability map, where four classes of DTM reliability are presented: zero, low, moderate, and high DTM reliability. Of the island area, 42% has zero DTM reliability, while 9% is classified as having low reliability, 16% as moderate reliability, and 33% as highly reliable. In practice,

this means that high-reliability areas can be used directly and with confidence, moderate reliability pixels should be used with caution, but can still provide accurate information about the DTM, low reliability areas are less reliable and best left out of critical analyses, while zero reliability areas represent places where the LiDAR data could not capture the ground at all. To give an indication of how vegetation cover influences the DTM reliability algorithm, a proxy for vegetation health and density (the NDVI) was derived from passive multispectral ORS from the Pleiades-Neo satellite. The NDVI was compared across the different DTM reliability classes, and it was found that the DTM reliability algorithm correctly identifies regions where the ground is covered by dense vegetation and yields less reliable DTM values.

It is important to note that this DTM reliability map does not provide an absolute truth, since no validation data was available, but rather a relative measure of confidence. Its main value lies in showing where the DTM can be used with confidence and where it should be treated more carefully. This DTM reliability map supports a more informed use of the dataset in studies where the DTM forms the basis, such as soil erosion risk mapping, landslide detection, and hydrological studies. Encouragingly, about half of the island's DTM (49%) has moderate to high reliability, meaning there is a substantial portion of Saba where the DTM provides dependable ground information.

4. What structural parameters can be derived from the AHN LiDAR data?

To make the LiDAR dataset useful for land cover mapping, the raw 3D point data first needs to be translated into structural parameters, which are measurable characteristics that describe both the vegetation and the underlying terrain. These parameters (Table 4.2) serve as input features for classification models that generate land cover maps (like the one in Smith et al., 2013). Two main groups of structural parameters were identified:

- **Vertical structure parameters**, which describe features on the terrain such as trees, shrubs, or bare ground.
- **Topographic parameters**, which describe the shape of the landscape, including its elevation (DTM), slope, and aspect.

The vertical structure parameters were derived by analyzing the vertical profile of the LiDAR data within each pixel. This vertical profile shows how LiDAR points are distributed from the ground up through the vegetation, after adjusting for the ground level provided by the DTM. Using a three-component Gaussian Mixture Model, this profile was broken down into separate layers representing the canopy, understory, and ground. In this way, the parameters capture whether these layers are present in a pixel, and if so, how much they contribute to the overall structure.

By expressing the LiDAR information in terms of these structural parameters, the dataset provides a standardized description of Saba's land cover, making it suitable for ecological analysis and mapping.

5. To what extent can the AHN LiDAR data be used to assess the land cover taxonomies?

The AHN LiDAR dataset already contains pre-classified groups: ground, building, water, and unclassified (Table 3.2). These groups align closely with the Level 1 land cover taxonomy, as buildings correspond to artificial surfaces, ground to barren land, and unclassified to forest and rangeland. This means that even in its published form, the dataset can be directly used to track large-scale landcover changes (Level 1), for example if future LiDAR surveys are carried out and compared to this dataset. To explore how the LiDAR data can be used in the assessment of the more detailed level-2 and level-3 land cover taxonomies, two small case studies were carried out.

In the case study for the level-2 taxonomy, the structural variability of areas classified as broadleaved evergreen forest in the Wageningen land cover map (see Figure 2.10 for the map) was analyzed. This was done by comparing the vertical structure parameters (see sub-question 4 above) for a subset of pixels (5 m x 5 m) that were classified as "broadleaved evergreen forest" in the Wageningen land cover map. It was found that these parameters varied significantly from pixel to pixel, therefore indicating the presence of structural variability within a single class. This implies that LiDAR data is able to provide additional information to the Wageningen land cover map that could separate the classes further into subclasses. Additionally, it might also be possible to distinguish areas within a single class with similar

vertical structures, enabling the identification of for example dominant tree species with a specific known height.

The case study for the level-3 taxonomy focused on the invasive vine *Coralita* (*Antigonon leptopus*), which thrives in lower-lying areas of Saba where the DTM reliability is high. Using an AI classification model on the Carto platform, a map was created showing *Coralita* coverage across the island. The Saba's susceptibility to *Coralita* was investigated using the Wageningen land cover map to see in which classes *Coralita* dominates; the results confirmed that *Coralita* is most common in pastures, bare rocks, and urban areas, in line with existing knowledge that it spreads in degraded or disturbed areas. The LiDAR-derived topographic parameters (see sub-question 4 above) provided further insights: *Coralita* was present on elevations between 0 and 250 m, and on a large range of slopes (0° to 50°). Interestingly, the aspect angle revealed that *Coralita* is predominantly found on slopes facing between southwest and east, likely due to a combination of wind exposure (primarily from the east), sun exposure, and soil type.

Overall, the results show that the AHN LiDAR dataset can be confidently used for Level 1 land cover taxonomies across Saba, is useful for adding structural detail to Level 2 habitats, and, in areas where the terrain model is reliable, can even provide species-level insights (Level 3). The potential of the dataset goes beyond the case studies presented here. For example, the extent can be investigated further by integrating LiDAR with other datasets, such as multispectral ORS data, where it would be possible to build land cover classification models that likely give an even more detailed picture of Saba's landscapes.

6. What limitations remain, and how can they inform future research efforts?

Although the AHN LiDAR dataset represents a significant step forward for mapping and monitoring on Saba, several limitations remain that shape how it can be applied. The most important is the uneven reliability of the DTM: in densely vegetated and steep areas, the ground is insufficiently sampled, leaving large gaps in the DTM. Another limitation is that the dataset captures only a single moment in time. Without repeated acquisitions, it is not possible to monitor ecological change, such as the spread of invasive species, recovery after storm events, or the effectiveness of management interventions. Looking ahead, much greater value could be gained from carrying out regular LiDAR surveys similar to the AHN program already established in the European Netherlands.

Another important limitation is the lack of validation data. Because no reference terrain model of the "true" ground exists for Saba, the DTM reliability algorithm had to be built using only the LiDAR data itself. This means that while the algorithm gives a strong indication of where the terrain model is trustworthy, its accuracy could not yet be directly confirmed. Similarly, the absence of detailed field data on the structure of certain habitats (see Table 2.2) meant that LiDAR information could not yet be fully used in supervised classification methods. Instead, external sources such as the Wageningen land cover map (Figure 2.10) and predicted *Coralita* coverage maps from Carto (Figure 5.26) were used to support the analysis. The positive side of this limitation is that it highlights a clear way forward: the full extent of the usability of the LiDAR data remains to be discovered with the future collection of independent ground-truth or field validation data. Such data would not only strengthen confidence in the terrain model but also allow LiDAR structural information to be directly linked to specific habitats and species. Recommendations for these steps are provided in Section 7.2.

7.2. Recommendations

Even though some areas remain difficult to sample due to Saba's topographic nature, some improvements can be made in the LiDAR acquisition plan. This is especially true for the areas currently classified as having moderate DTM reliability, which make up about 16% of the island Figure 5.7. Here, one improvement lies in the flight line design. The current survey was conducted with seven parallel lines with 55% overlap, and only one cross-line. Adding more cross-lines, flown at different orientations, would allow steep slopes and dense vegetation to be scanned from multiple angles, increasing the chance of capturing ground points. The benefits of such an approach can already be seen under line 101, where the additional cross-line produced higher point densities. This is also reflected in the

final DTM reliability score R_{DTM} , where the lower right side of the island shows a higher DTM reliability compared to the left. Although a second cross-line (line 102) was planned, it was not flown. Considering its importance, it is advised for future acquisitions to include line 102, as well as introduce more lines with a different orientation. Adjustments to the sensor settings could also improve results. The current survey operated at a PRF of 700 kHz (350 kHz for each beam), while the system can go up to 2 MHz per beam. Using a higher PRF would increase the number of laser points collected, and even without changing altitude, could roughly double the point density per square meter. Changes to altitude or flight speed could also help, but they come with risks such as longer acquisition times and a greater chance of cloud formation during acquisition.

For the low and zero reliability areas, particularly under dense forest canopies, better flight planning alone is unlikely to solve the problem. The problem of the dense canopy remains, and because most tree species are evergreen, the island is not going to have a leafless period where ground coverage significantly improves (Vosselman and Maas, 2010). Instead, it would be valuable to install (permanent) Ground Control Points (GPCs) beneath the canopy. These fixed reference markers would provide accurate elevation benchmarks that can be used to correct and calibrate the DTM in places where laser pulses cannot penetrate to the ground. While this would not improve every pixel in the dataset, it would increase confidence in critical areas.

Finally, to confirm the accuracy of the introduced DTM reliability algorithm, it could be tested in the European Netherlands, where highly accurate reference DTMs are available. Alternatively, new GPS measurements collected on Saba, particularly in low and moderate-reliability areas, could serve as ground-truth data. This would help verify that the algorithm correctly identifies areas with low to moderate DTM reliability.

References

Achsah Mitchell (2023). *Coralita smothering local vegetation*. <https://www.naturetoday.com/intl/en/nature-reports/message/?msg=30513>. Accessed: August 18, 2025.

Actueel Hoogtebestand Nederland (Sept. 2023). *Besteksvoorwaarden inwinning LiDAR Caribisch Nederland 2023*. Technical Specification Bijlage 1. Version 1.0. Definitieve versie, gepubliceerd op 28-09-2023. Netherlands: Rijkswaterstaat. url: <https://www.ahn.nl>.

Airbus Defence and Space (Aug. 2024). *Pléiades Neo User Guide*. Version 4. Airbus DS GEO SA. url: <https://intelligence.airbus.com>.

Baetens, Louis, Camille Desjardins, and Olivier Hagolle (2019). "Validation of Copernicus Sentinel-2 Cloud Masks Obtained from MAJA, Sen2Cor, and FMask Processors Using Reference Cloud Masks Generated with a Supervised Active Learning Procedure". In: *Remote Sensing* 11.4. issn: 2072-4292. doi: 10.3390/rs11040433. url: <https://www.mdpi.com/2072-4292/11/4/433>.

Beeldmateriaal Nederland (Sept. 2023). *Besteksvoorwaarden inwinning HRL Caribisch Nederland 2023*. Technical Specification. Version 1.0. Netherlands: Rijkswaterstaat. url: <https://www.beeldmateriaal.nl>.

Binding, Caren E. et al. (Mar. 2012). "The MERIS MCI and its potential for satellite detection of winter diatom blooms on partially ice-covered Lake Erie". In: *Journal of Plankton Research* 34.6, pp. 569–573. issn: 0142-7873. doi: 10.1093/plankt/fbs021. eprint: <https://academic.oup.com/plankt/article-pdf/34/6/569/4477657/fbs021.pdf>. url: <https://doi.org/10.1093/plankt/fbs021>.

Boeken, M. (2014). "Saba en zijn orchideeën". In: *Orchideeën (Tijdschrift Ned. Orchid. Verenig.)* 76.5, pp. 100–117.

Burke, Janelle M. and Antonio DiTommaso (2011). "Corallita (Antigonon leptopus): Intentional Introduction of a Plant with Documented Invasive Capability". In: *Invasive Plant Science and Management* 4.3, pp. 265–273. doi: 10.1614/IPSM-D-10-00088.1.

Cavender-Bares, Jeannine et al. (Mar. 2022). "Integrating remote sensing with ecology and evolution to advance biodiversity conservation". In: *Nature Ecology Evolution*. doi: 10.1038/s41559-022-01702-5.

Cherian, Anilla (2007). "Linkages between biodiversity conservation and global climate change in small island developing States (SIDS)". In: *Natural Resources Forum* 31.2, pp. 128–131. doi: 10.1111/j.1477-8947.2007.00138.x.

Convention on Biological Diversity (2024). *Kunming-Montreal Global Biodiversity Framework*. url: <https://www.cbd.int/gbf>.

Crooks, Jeff (Sept. 2005). "Lag times and exotic species: The ecology and management of biological invasions in slow-motion". In: *Ecoscience* 12, pp. 316–329. doi: 10.2980/i1195-6860-12-3-316.1.

DCNA (2024). *Climate Financing in the Caribbean Netherlands*. <https://dcnanature.org/climatefinancing/>. Accessed: August 13, 2025.

Debrot, A et al. (2025). *Global Register of Introduced and Invasive Species - Saba, Bonaire, Sint Eustatius and Saba*. Checklist dataset. Version 1.18. url: https://cloud.gbif.org/griis/resource?r=saba_griis&v=1.18.

Debrot, A. et al. (June 2025). *State of Nature Report for the Caribbean Netherlands 2024: A second 6-year assessment of the Conservation State, threats and management implications for habitat and species in the Caribbean Netherlands*. doi: 10.18174/684783.

Debrot, A.O., R.J.H.G. Henkens, and P.J.F.M. Verweij (Nov. 2018). *Staat van de natuur van Caribisch Nederland 2017: een eerste beoordeling van de staat (van instandhouding), bedreigingen en managementimplicaties van habitats en soorten in Caribisch Nederland*. Nederlands. Wageningen Marine Research rapport C086/17. Netherlands: Wageningen Marine Research. doi: 10.18174/426340.

Dutch Caribbean Nature Alliance (DCNA) (n.d.). *Research Monitoring Wishlist*. <https://dcnanature.org/research-monitoring-wishlist/>. Accessed on 14 August 2025.

Dutch Caribbean Species Register (July 31, 2018). *Black Sweetwood Nectandra krugii*. https://www.dutchcaribbeanspecies.org/linnaeus_ng/app/views/species/nsr_taxon.php?id=177908&cat=166. Accessed: August 10, 2025.

— (n.d.[a]). *Elephant Ear Philodendron giganteum*. https://www.dutchcaribbeanspecies.org/linnaeus_ng/app/views/species/nsr_taxon.php?id=177344&cat=CTAB_NAMES. Accessed: August 10, 2025.

— (n.d.[b]). *Prestoea acuminata var. montana*. https://www.dutchcaribbeanspecies.org/linnaeus_ng/app/views/species/nsr_taxon.php?id=178307&cat=CTAB_NAMES. Accessed: August 10, 2025.

— (n.d.[c]). *Raspberry Rubus rosifolius*. https://www.dutchcaribbeanspecies.org/linnaeus_ng/app/views/species/nsr_taxon.php?id=190004. Accessed: August 12, 2025.

Earthdata, NASA (2025). *Global Ecosystem Dynamics Investigation Lidar*. Accessed: 2025-09-21. url: <https://www.earthdata.nasa.gov/data/instruments/gedi-lidar>.

Elachi, Charles and Jakob van Zyl (2021). *Introduction to the Physics and Techniques of Remote Sensing (3rd Edition)*. John Wiley Sons. isbn: 978-1-119-52301-7. url: <https://app.knovel.com/hotlink/toc/id:kpiPTRSE17/introduction-physics/introduction-physics>.

Ernst, J.J. and P. Ketner (Mar. 2007). *Study on the ecology and possible control methods of the invasive plant species Antigonon leptopus (Corallita or Mexican Creeper)*. Tech. rep. Corallita Pilot Project St. Eustatius. url: <https://www.dcbd.nl/reposerver/api/file/3289>.

Fassnacht, Fabian Ewald et al. (2016a). "Review of studies on tree species classification from remotely sensed data". In: *Remote Sensing of Environment* 186, pp. 64–87. issn: 0034-4257. doi: <https://doi.org/10.1016/j.rse.2016.08.013>. url: <https://www.sciencedirect.com/science/article/pii/S0034425716303169>.

— (2016b). "Review of studies on tree species classification from remotely sensed data". In: *Remote Sensing of Environment* 186, pp. 64–87. issn: 0034-4257. doi: <https://doi.org/10.1016/j.rse.2016.08.013>. url: <https://www.sciencedirect.com/science/article/pii/S0034425716303169>.

Freitas, J.A. et al. (Jan. 2015). *A landscape ecological vegetation map of Saba (Lesser Antilles)*.

Gardenia.net (n.d.). *Antigonon leptopus (Coral Vine)*. url: <https://www.gardenia.net/plant/antigonon-leptopus>.

GeeksforGeeks Contributors (2025). *Bayesian Information Criterion (BIC)*. Accessed: 2025-09-06. url: <https://www.geeksforgeeks.org/machine-learning/bayesian-information-criterion-bic/>.

GeodetischeInfrastructuur (2024). *transformations*. Accessed: 2025-02-15. url: <https://github.com/GeodetischeInfrastructuur/transformations>.

Gleason, David B. (2011). *Mount Scenery*. Flickr. CC BY-SA 2.0. License: <https://creativecommons.org/licenses/by-sa/2.0/>. Accessed: 2025-10-22. url: <https://www.flickr.com/photos/71861129@N00/6585198365>.

Goepfert, Jens, Uwe Soergel, and Alexander Brzank (July 2008). "Integration of Intensity Information and Echo Distribution in the Filtering Process of LIDAR Data in Vegetated Areas". In: *Proceedings of the XXI ISPRS Congress*. International Society for Photogrammetry and Remote Sensing (ISPRS). Edinburgh, UK, pp. 17–19. url: https://www.ipi.uni-hannover.de/fileadmin/ipi/publications/goepfert_edinburgh_2008_final.pdf.

Google Earth (Sept. 2025). *Map created using Google Earth*. <https://earth.google.com/>. Imagery © Google, Maxar Technologies.

Government of the Netherlands (2025). *Controlling Invasive Alien Species*. Accessed: 2025-09-13. url: <https://www.government.nl/topics/nature-and-biodiversity/controlling-invasive-alien-species>.

Guo, Qinghua et al. (June 2010). "Effects of Topographic Variability and Lidar Sampling Density on Several DEM Interpolation Methods". In: *Photogrammetric Engineering and Remote Sensing* 76. doi: 10.14358/PERS.76.6.701.

Het Waterschapshuis (2024a). *Kwaliteitsrapportage Aerotriangulatie*. Internal Report. Unpublished internal quality report for passive multi-spectral acquisition over Saba. Accessed: 2024-12-17. Netherlands: Beeldmateriaal Nederland.

— (2024b). *Kwaliteitsrapportage LIDAR - Saba*. Internal Report. Unpublished internal quality report for LiDAR acquisition over Saba. Accessed: 2024-12-17. Netherlands: AHN.

Hu, Ronghai et al. (Mar. 2018). "Using Airborne Laser Scanner and Path Length Distribution Model to Quantify Clumping Effect and Estimate Leaf Area Index". In: *IEEE Transactions on Geoscience and Remote Sensing* 56.

Hub, Sentinel (2025). *Sentinel-2 L2A Data Documentation*. <https://docs.sentinel-hub.com/api/latest/data/sentinel-2-12a/>. Accessed: 2025-09-11.

Imaging, Vexcel (2022). *UltraCam Eagle 4.1: Push the Limits of Your Aerial Missions*. Product brochure. Graz, Austria. url: <https://www.vexcel-imaging.com/ultracam-eagle-4-1/#downloads>.

iNaturalist (2023a). *Bothriochloa pertusa*. <https://www.inaturalist.org/observations/192062977>. Accessed: August 10, 2025.

— (2023b). *Eugenia axillaris*. <https://www.inaturalist.org/photos/258521698>. Accessed: August 10, 2025.

IPCC (2019). *Summary for Policymakers. Climate Change and Land: An IPCC Special Report on Climate Change, Desertification, Land Degradation, Sustainable Land Management, Food Security, and Greenhouse Gas Fluxes in Terrestrial Ecosystems*. <https://www.ipcc.ch/site/assets/uploads/2019/08/Fullreport.pdf>. Accessed: 2025-09-23.

IUCN (n.d.). *The IUCN Red List of Threatened Species. Summary Statistics*. url: <https://www.iucnredlist.org/resources/summary-statistics#Summary%20Tables>.

Jacobs, Simone (Nov. 2024). *Highest Point in the Netherlands Found to Be Even Lower Than Assumed*. Accessed: 2025-09-14. IamExpat Media B.V. url: <https://www.iamexpat.nl/lifestyle/lifestyle-news/highest-point-netherlands-found-be-even-lower-assumed>.

Jafarbiglu, Hamid and Alireza Pourreza (June 2022). "A comprehensive review of remote sensing platforms, sensors, and applications in nut crops". In: *Computers and Electronics in Agriculture* 197, p. 106844. doi: 10.1016/j.compag.2022.106844.

Jagodnik, Petra et al. (Oct. 2019). "Identification and Mapping of Soil Erosion Processes Using the Visual Interpretation of LiDAR Imagery". In: *ISPRS International Journal of Geo-Information* 8, p. 438. doi: 10.3390/ijgi8100438.

Kattenborn, Teja (2018). "Linking Canopy Reflectance and Plant Functioning through Radiative Transfer Models". PhD thesis. Karlsruhe Institute of Technology (KIT). doi: 10.5445/IR/1000089168. url: <https://doi.org/10.5445/IR/1000089168>.

KNMI (2023). *KNMI'23 climate scenarios for the Netherlands*.

— (25 February 2025). "The state of Our Climate in 2024". In: *KNMI Publication 24-02.2*.

Koninklijk Nederlands Meteorologisch Instituut (n.d.). *Uurgegevens van het Caribisch gebied*. https://www.knmi.nl/nederland-nu/klimatologie/uurgegevens_Caribisch_gebied. Accessed: 2025-08-09.

Laughlin, Nicholas (2009a). *mt scenery and hell's gate*. Flickr. CC BY-NC-SA 2.0. License: <https://creativecommons.org/licenses/by-nc-sa/2.0/>. Accessed: 2025-10-22. url: <https://www.flickr.com/photos/nicholaslaughlin/3620118163/in/album-72157619466290866>.

— (2009b). *mt scenery bromeliad*. Flickr. CC BY-NC-SA 2.0. License: <https://creativecommons.org/licenses/by-nc-sa/2.0/>. Accessed: 2025-10-22. url: <https://www.flickr.com/photos/nicholaslaughlin/3618372730/in/album-72157619466290866>.

— (2009c). *mt scenery path*. Flickr. CC BY-NC-SA 2.0. License: <https://creativecommons.org/licenses/by-nc-sa/2.0/>. Accessed: 2025-10-22. url: <https://www.flickr.com/photos/nicholaslaughlin/3618266540/in/album-72157619466290866>.

Lesparre, Jochem, Huib de Ligt, and Lennard Huisman (Mar. 2025). *Coordinate transformation between International Terrestrial Reference System and coordinates used for the islands Bonaire, St. Eustatius and Saba (BES)*. Tech. rep. Version 24 March 2025. Creative Commons Attribution-NoDerivatives 4.0 International License. Netherlands Partnership Geodetic Infrastructure (NSGI). url: https://www.nsgi.nl/documents/1888506/69577825/BESTRANS2020_v250324.pdf/81eb5542-d7a0-ed7a-ee4c-26b96918f94f?t=1744723224098.

Mandlburger, Gottfried (Feb. 2025). *Airborne Lidar: A Tutorial for 2025 – Part II: Integrated Systems*. Accessed: 2025-09-10. LIDAR Magazine. url: <https://lidarmag.com/2025/02/09/airborne-lidar-a-tutorial-for-2025-2/>.

Mark R. Drinkwater, Helge Rebhan (2007). *Sentinel-3: Mission Requirements Document*. Tech. rep. European Space Agency.

Marselis, Suzanne et al. (Oct. 2018). "Distinguishing vegetation types with airborne waveform lidar data in a tropical forest-savanna mosaic: A case study in Lopé National Park, Gabon". In: *Remote Sensing of Environment* 216. doi: 10.1016/j.rse.2018.07.023.

Masini, Nicola et al. (2018). "Medieval Archaeology Under the Canopy with LiDAR. The (Re)Discovery of a Medieval Fortified Settlement in Southern Italy". In: *Remote Sensing* 10.10. issn: 2072-4292. doi: 10.3390/rs10101598. url: <https://www.mdpi.com/2072-4292/10/10/1598>.

Ministry of Agriculture, Nature and Food Quality (2020). *Nature and Environment Policy Plan Caribbean Netherlands 2020-2030*. <https://english.rijksdienstcn.com/agriculture-nature-and-food-quality/documents/leaflets/agriculture-horticulture-and-livestock/nature-policy-plan/nature-policy-plan/nature-and-environment-policy-plan-caribbean-netherlands-2020-2030>.

Mucher, Sander et al. (Nov. 2024). *Land Cover Classification 2020 and analysis historical land cover changes since 1900 on Aruba*. Wageningen University Research. doi: 10.18174/679738.

Mucher, Sander et al. (Oct. 2014). *Production of Digital Terrain Models for the Dutch Caribbean*. en. wageningen.

Mulatu, Kalkidan et al. (Oct. 2017). "Biodiversity Monitoring in Changing Tropical Forests: A Review of Approaches and New Opportunities". In: *Remote Sensing* 2017. doi: 10.3390/rs9101059.

Netherlands Space Office (n.d.). *Beschikbare data*. Accessed: 2025-03-02. NSO. url: <https://www.spaceoffice.nl/nl/satellietdataportaal/beschikbare-data/>.

Observation.org (2015). *Prestoea acuminata var. montana*. <https://observation.org/observation/231464587/>. Accessed: August 10, 2025.

— (2023a). *Swietenia mahagoni*. <https://observation.org/observation/264251759/>. Accessed: August 11, 2025.

— (2023b). *Swietenia mahagoni*. <https://observation.org/observation/264251689/>. Accessed: August 11, 2025.

— (2024). *Aristida adscensionis*. <https://observation.org/observation/336289797/>. Accessed: August 10, 2025.

— (2025a). *Antigonon leptopus*. <https://observation.org/observation/366373943/>. Accessed: August 18, 2025.

— (2025b). *Coralita Observations in Saba*. https://observation.org/locations/34961/observations/?date_after=2017-01-01&date_before=2025-07-30&species=198549&page=1. Accessed: 2025-08-07.

— (2025c). *Giant philodendron*. <https://observation.org/observation/340916806/>. Accessed: August 10, 2025.

Paul Illsley (2023). *Google Maps: Mount Scenery Trailhead*. <https://maps.app.goo.gl/5wjj6ESvnUUhHQm9>. Accessed: August 11, 2025.

Press, BRIT (July 2021). "A Systematic Vademecum to the Vascular Plants of Saba". In: *Journal of the Botanical Research Institute of Texas* 15.1, p. 4. doi: 10.17348/jbrit.v15.i1.1116. url: <https://journals.brit.org/jbrit/article/view/1116>.

Public Entity Saba (Feb. 2025). *Invasive Species Control Measures Ongoing*. <https://www.sabagov.nl/news/invasive-species-control-measures-ongoing>. Accessed: August 12, 2025.

RIEGL Laser Measurement Systems GmbH (2024). *RIEGL VQ-1560 II-S Dual Channel Waveform Processing Airborne LIDAR Scanning System for High Point Density Mapping and Ultra-Wide Area Mapping*. Datasheet. RIEGL Laser Measurement Systems GmbH. url: <https://www.riegl.com/>.

Rijksdienst Caribisch Nederland (2020). *Nature and environment policy plan Caribbean Netherlands 2020-2030*. <https://english.rijksdienstcn.com/agriculture-nature-and-food-quality/documents/leaflets/agriculture-horticulture-and-livestock/nature-policy-plan/nature-policy-plan/nature-and-environment-policy-plan-caribbean-netherlands-2020-2030>. Accessed: August 12, 2025.

Salem, Nema and Sahar Hussein (2019). "Data dimensional reduction and principal components analysis". In: *Procedia Computer Science* 163. 16th Learning and Technology Conference 2019Artificial Intelligence and Machine Learning: Embedding the Intelligence, pp. 292–299. issn: 1877-0509. doi: <https://doi.org/10.1016/j.procs.2019.12.111>. url: <https://www.sciencedirect.com/science/article/pii/S1877050919321507>.

Saritha, Vara et al. (2025). *Remote sensing for environmental monitoring*. Singapore: Springer. isbn: 978-981-96-5546-5 981-96-5546-3. doi: 10.1007/978-981-96-5546-5. url: <https://doi.org/10.1007/978-981-96-5546-5>.

Sentinel Hub (2024). *Sentinel-2 L2A Documentation*. Accessed: 2025-07-18. url: <https://docs.sentinel-hub.com/api/latest/data/sentinel-2-12a/>.

Smith, S. R. et al. (Nov. 2013). *Use of Satellite Data for the Monitoring of Species on Saba and St. Eustatius*. Research Report C124/13 (BO-11-011.05-019). Wageningen Research, p. 86. url: <https://library.wur.nl/WebQuery/wurpubs/fulltext/275484>.

Song, Xiao-Peng et al. (Aug. 1, 2018). "Global land change from 1982 to 2016". In: *Nature* 560.7720, pp. 639–643. issn: 1476-4687. doi: 10.1038/s41586-018-0411-9. url: <https://doi.org/10.1038/s41586-018-0411-9>.

Spheer (2025). *Spheer AI powered monitoring*. Accessed: 2025-09-09. url: <https://spheer.ai/en>.

Stoffers, A.L. (Jan. 1956). "The Vegetation of the Netherlands Antilles". In: *Studies on the Flora of Curaçao and other Caribbean Islands* 1.1, pp. 1–142.

Tarrio, Katelyn et al. (2020). "Comparison of cloud detection algorithms for Sentinel-2 imagery". In: *Science of Remote Sensing* 2, p. 100010. issn: 2666-0172. doi: <https://doi.org/10.1016/j.srs.2020.100010>. url: <https://www.sciencedirect.com/science/article/pii/S2666017220300092>.

Thayts (2011). *BES islands location map*. Accessed: 2025-01-31. url: https://commons.wikimedia.org/wiki/File:BES_islands_location_map.svg.

The American Society for Photogrammetry & Remote Sensing (Mar. 26, 2019). *LAS Specification 1.4 - R14*. url: https://www.asprs.org/wp-content/uploads/2019/03/LAS_1_4_r14.pdf.

Time and Date AS (2025). *Sun & Moon Times Today, Saba, Caribbean Netherlands*. Accessed: 2025-08-25. url: <https://www.timeanddate.com/astronomy/07610358>.

Timmermans, Joris and W. Daniel Kissling (2023). "Advancing terrestrial biodiversity monitoring with satellite remote sensing in the context of the Kunming-Montreal global biodiversity framework". In: *Ecological Indicators* 154, p. 110773. issn: 1470-160X. doi: <https://doi.org/10.1016/j.ecolind.2023.110773>. url: <https://www.sciencedirect.com/science/article/pii/S1470160X23009159>.

United Nations Conference on Environment and Development (1992). *Report of the United Nations Conference on Environment and Development (Rio de Janeiro, 3–14 June 1992)*. url: <https://digitallibrary.un.org/record/168679>.

United Nations Environment Program (Oct. 2018). *Proposed areas for inclusion in the SPAW list: Annotated format for presentation report for Mt. Scenery National Park*. Report proposing Mt. Scenery National Park for inclusion in the SPAW list. Kingdom of the Netherlands. url: <https://english.rriksdienstcn.com/agriculture-nature-and-food-quality/nature>.

van der Burg, W.J. et al. (2012). *Naturalised and invasive alien plant species in the Caribbean Netherlands: status, distribution, threats, priorities and recommendations : report of a joint IMARES/Carmabi/PRI project financed by the Dutch Ministry of Economic Affairs, Agriculture Innovation*. English. Report / IMARES Wageningen UR C185/11. Netherlands: Plant Research International.

Van Ewijk, K.Y. et al. (2014). "Predicting fine-scale tree species abundance patterns using biotic variables derived from LiDAR and high spatial resolution imagery". In: *Remote Sensing of Environment* 150. Publisher: Elsevier Inc., pp. 120–131. issn: 00344257 (ISSN). doi: 10.1016/j.rse.2014.04.026. url: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-84900840804&doi=10.1016%2fj.rse.2014.04.026&partnerID=40&md5=ba91c20ac46f8e41f316bd8abde1a86a>.

Volcanoes in the Dutch Caribbean (n.d.). <https://www.knmidc.org/volcanoes/?Mount-Scenery>. Accessed: 2025-01-08.

Vosselman, George and Hans-Gerd Maas (Jan. 2010). *Airborne and Terrestrial Laser Scanning*. Whittles Publishing. isbn: 978-1-904445-87-6.

Wang, Cheng et al. (2024). *Introduction to LiDAR remote sensing*. First edition. 1 online resource vols. Boca Raton FL: CRC Press. doi: 10.1201/9781032671512.

world-wide-mike (2021). *A blog post about a trip to Saba*. <https://world-wide-mike.travellerspot.int.com/362/>. Accessed: August 10, 2025.

Zar, Jerrold H. (2005). "Spearman Rank Correlation". In: *Encyclopedia of Biostatistics*. John Wiley Sons, Ltd. isbn: 9780470011812. doi: <https://doi.org/10.1002/0470011815.b2a15150>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/0470011815.b2a15150>. url: <https://onlinelibrary.wiley.com/doi/abs/10.1002/0470011815.b2a15150>.

Zhu, Xiao et al. (2024). "Mitigating terrain shadows in very high-resolution satellite imagery for accurate evergreen conifer detection using bi-temporal image fusion". In: *International Journal of Applied Earth Observation and Geoinformation* 134, p. 104244. issn: 1569-8432. doi: <https://doi.org/10.1016/j.jag.2024.104244>. url: <https://www.sciencedirect.com/science/article/pii/S1569843224006009>.

A

Flora of Saba

Table A.1: Flora species of Saba by habitat and occurrence taken from A. Debrot et al., 2025; Freitas et al., 2015. These are species for which enough monitoring data and/or knowledge are present.

79

Habitats	Occurrence	Species name	Notes
Elfin forest	dominant and restricted	<i>Freziera undulata</i>	Evergreen tree
		<i>Charianthus purpureus</i>	Evergreen shrub
	restricted	Epiphytes: <i>Voyria aphylla</i> , <i>Utricularia alpina</i> , <i>Notopleura guadalupensis</i> , <i>Ornithidium reflexum</i> , <i>Peperomia hernandiifolia</i> , <i>Peperomia emarginella</i> , <i>Werauhia urbaniana</i>	These plants grow on other plants. They require an environment with high humidity and mist
		<i>Besleria lutea</i>	Evergreen shrub
		<i>Hymenophyllum hirtellum</i>	Evergreen fern
		<i>Pilea obtusata</i>	Evergreen herb
		<i>Rubus rosifolius</i>	Evergreen fern, naturalized
Elfin + Montane forest		<i>Nephrolepis rivularis</i>	Evergreen fern
		<i>Cyatheaceae</i> spp.	Evergreen tree fern, characteristic of Elfin and Montane forest, as it only grows in moist and misty environments
		<i>Prestoea acuminata</i>	Evergreen palm tree
		<i>Cecropia peltata</i>	Evergreen trumpet tree, a pioneer species

Habitats	Occurrence	Species name	Notes
		<i>Begonia retusa</i>	Evergreen herb
		<i>Vriesea ringens</i>	Evergreen epiphyte
Montane forest	restricted	<i>Nectandra krugii</i>	Evergreen tree, endangered
Montane + dry tropical forest		<i>Myrcia splendens</i>	Semi deciduous tree, high cover throughout the whole dry tropical forest
		<i>Coccoloba diversifolia</i>	Evergreen tree
		<i>Cordia sulcata</i>	Tree, important in the northwestern lower slopes, a pioneer species.
		<i>Clusia major</i>	Evergreen shrub/tree, among the highest cover around the northwestern lower slopes, but also occurs from the northeastern to southwestern lower slopes, a pioneer species
		<i>Citharexylum spinosum</i>	Semi-deciduous tree, high abundance in Montane forest
		<i>Miconia laevigata</i>	Evergreen shrub, important in northwestern lower slopes
Dry tropical forest	dominant and restricted	<i>Pisonia subcordata</i>	Deciduous tree, occurs on the northern lower slopes and dominates in parts from the northeastern to southwestern lower slopes
		<i>Coccoloba uvifera</i>	Evergreen tree
		<i>Swietenia mahagoni</i>	Endangered deciduous tree, dominates in western lower slopes
		<i>Guettarda scabra</i>	Evergreen tree, dominates in the northern lower slopes and high cover in the mid section of the western and southern slopes
		<i>Coccoloba swartzii</i>	Evergreen tree, dominates in northern lower slopes and high cover from the northeastern to southwestern lower slopes
	restricted	<i>Bursera simaruba</i>	Deciduous tree
		<i>Guaiacum officinale</i>	Evergreen tree, endangered
		<i>Casearia decandra</i>	Evergreen tree/shrub, common from northern to southwestern lower slopes
		<i>Eugenia axillaris</i>	Evergreen tree, common and high cover in the western lower slopes, and frequent in the northern to southwestern lower slopes
		<i>Myrcianthes fragrans</i>	Evergreen tree/shrub

Habitats	Occurrence	Species name	Notes
		<i>Randia aculeata</i>	Deciduous tree, common in the northern lower slopes and in the mid section from northeast to southwestern slopes
		<i>Maytenus laevigata</i>	Evergreen tree, common in the northern lower slopes and in parts ranging from northeast to southwestern lower slopes
		<i>Guapira fragrans</i>	Evergreen tree/shrub
		<i>Inga laurina</i>	Evergreen tree, important in the northwestern lower slopes, and high cover in southwestern lower slopes
		<i>Byrsonima spicata</i>	Evergreen tree, important in northwestern lower slopes, a pioneer species
Elfin forest + Montane forest + dry tropical forest		<i>Philodendron giganteum</i>	Evergreen hemiepiphyte, dominant in Montane shrub/herb layer and in the northwestern lower slopes
		<i>Blechnum occidentale</i>	Evergreen fern
		<i>Heliconia bihai</i>	Evergreen shrub, most dominant in Montane forest shrub/herb layer after philodendron giganteum
		<i>Piper dilatatum</i>	Evergreen shrub
		<i>Prestoea montana</i>	Evergreen palm tree, most prominent in Montane forest and high presence in the northwestern lower slopes
Dry tropical forest + dry shrubland and grassland		<i>Lantana camara</i>	Evergreen shrub, frequent in the shrub layer of the dry tropical forest and dominant in the shrubland
		<i>Croton astroites</i>	Evergreen shrub, frequent in the shrub layer of the dry tropical forest and common in a small area on the northern lowest slopes
		<i>Plumbago scandens</i>	Evergreen shrub, common throughout the lower-elevation dry tropical forest and in the shrubland
		<i>Coccoloba uvifera</i>	Evergreen shrub, found throughout both habitats and on cliffs in coastal areas
Dry shrubland and grassland	dominant + restricted	<i>Mitracarpus polycladus</i>	Evergreen shrub, dominant from the northeastern to southwestern lower slopes
		<i>Aristida adscensionis</i>	A grass, dominant in the grassland
		<i>Dactyloctenium aegyptium</i>	A grass, dominant in both shrubland and grassland, and can also be found on cliffs

Habitats	Occurrence	Species name	Notes
restricted		<i>Botriochloa pertusa</i>	An invasive grass, dominates the dry shrubland and grassland
		<i>Jatropha gossypiifolia</i>	Semi-deciduous shrub, found on the lowest slopes and cliffs
		<i>Chloris barbata</i>	A grass, dominant in the grassland

B

Invasive Species on Saba

Table B.1: List of 46 invasive terrestrial plant species from the Global Register of Introduced and Invasive Species (GRIIS) recorded in Saba according to Debrot et al., 2025. These species have been identified as having a negative impact on biodiversity.

GRIIS Species ID	Scientific Name	GRIIS Species ID	Scientific Name
12172	<i>Albizia lebbeck</i>	12219	<i>Jasminum fluminense</i>
12173	<i>Aloe vera</i>	12220	<i>Kalanchoe daigremontiana</i>
12176	<i>Antigonon leptopus</i>	12221	<i>Kalanchoe pinnata</i>
12179	<i>Asystasia gangetica</i>	12223	<i>Lawsonia inermis</i>
12180	<i>Bambusa vulgaris</i>	12225	<i>Leucaena leucocephala</i>
12182	<i>Bothriochloa pertusa</i>	12226	<i>Mangifera indica</i>
12183	<i>Urochloa mutica</i>	12229	<i>Melinis repens</i>
12184	<i>Caesalpinia bonduc</i>	12233	<i>Moringa oleifera</i>
12186	<i>Calotropis procera</i>	12236	<i>Nephrolepis hirsutula</i>
12190	<i>Catharanthus roseus</i>	12239	<i>Oeceoclades maculata</i>
12194	<i>Cleome gynandra</i>	12244	<i>Megathyrsus maximus</i>
12195	<i>Clitoria ternatea</i>	12245	<i>Parthenium hysterophorus</i>
12199	<i>Cordia sebestena</i>	12247	<i>Philodendron giganteum</i>
12200	<i>Arivela viscosa</i>	12251	<i>Pteris tripartita</i>
12202	<i>Cryptostegia grandiflora</i>	12252	<i>Pteris vittata</i>
12204	<i>Dactyloctenium aegyptium</i>	12258	<i>Ricinus communis</i>
12205	<i>Delonix regia</i>	12259	<i>Sansevieria hyacinthoides</i>
12208	<i>Eleusine indica</i>	12260	<i>Sansevieria trifasciata</i>
12210	<i>Epipremnum aureum</i>	12262	<i>Senna bicapsularis</i>
12212	<i>Eragrostis ciliaris</i>	12263	<i>Senna italica</i>
12214	<i>Euphorbia tithymaloides</i>	12271	<i>Syngonium podophyllum</i>
12217	<i>Gossypium barbadense</i>	12275	<i>Tecoma stans</i>
12218	<i>Indigofera tinctoria</i>	12280	<i>Tithonia diversifolia</i>

C

Remote Sensing Techniques

This appendix provides a more in-depth overview of the remote sensing techniques relevant to this study: LiDAR remote sensing (Section C.1), and passive multispectral remote sensing (Section C.2). For each technique, the system components, additional principles, and factors influencing the data quality are discussed.

C.1. LiDAR Remote Sensing

C.1.1. LiDAR System Components

A LiDAR system consists of a laser scanning system and a positioning and orientation system. The laser scanning system carries out the actions of generating, directing, and receiving laser pulses. The positioning and orientation system consists of a GNSS, which provides accurate positioning data of the instrument during movement, and an Inertial Measurement Unit (IMU), which records orientation data of the instrument, measured by roll (ϕ), pitch (θ), and yaw (κ).

C.1.2. LiDAR Resolution

As mentioned in Subsection 3.1.1, one of the advantages of LiDAR compared to other remote sensing techniques is the high spatial resolution (measured in point density D) that can be obtained, particularly with airborne and terrestrial platforms. Both inherent characteristics of the instrument and flight parameters contribute to the point density. On the instrument side, the angular resolution and ranging resolution contribute to the spatial resolution, which corresponds to the smallest distinguishable horizontal distance, and the smallest distinguishable distance along the line of sight, respectively. The angular resolution is influenced by the divergence angle of the beam (which is typically only a few milliradians in LiDAR (Wang et al., 2024)), which focuses the energy into a small footprint on the surface. This footprint can be very small for terrestrial and airborne LiDAR, and is larger for spaceborne LiDAR, usually >10 m (Wang et al., 2024). The ranging resolution is influenced by the pulse duration (the length of time a pulse is emitted) and the precision of the TOF measurements.

C.1.3. LiDAR Data Quality

The quality of LiDAR measurements is determined by several factors. The geometric accuracy is determined by the TOF measurement precision, and the GNSS and IMU accuracy. Precision refers to the variability of measurements under noise, while accuracy refers to how close the measurements are to the true or reference value. The quality is also determined by the ranging precision, which is influenced by the signal-to-noise ratio (SNR). The ranging precision is usually very high due to the coherent and high-energy laser pulses (Wang et al., 2024). The reflectivity of the surface also influences the quality, as highly reflective surfaces yield stronger returns and thus more reliable points, whereas dark or water-covered surfaces often result in weak returns. In addition to surface properties, atmospheric conditions are another important factor. Bad weather, such as heavy rain, attenuates the laser signal, thereby reducing SNR and therefore ranging precision. Atmospheric turbulence can also distort the beam path, introducing stochastic noise in the measurements.

LiDAR quality is usually assessed in terms of absolute and relative accuracy, which can be separated into planimetric and vertical components. Absolute accuracy describes the agreement of the LiDAR-derived coordinates with their actual coordinates, usually determined through ground control points. The relative accuracy refers to the internal consistency of the dataset. This is often determined by comparing the point cloud coordinates of overlapping flight lines. For the vertical relative accuracy, this is done by comparing elevation differences in overlapping areas, while planimetric relative accuracy uses the alignment of features, for example, building edges, across strips. Both types of accuracy are influenced by systematic errors, which are errors inherent to the system, such as IMU misalignment or laser calibration errors, and stochastic errors, which are random errors such as noise in TOF measurements. Systematic errors can often be corrected through calibration and adjustment against control data, while stochastic errors define the inherent accuracy limits of the system.

C.1.4. Extended LiDAR Applications

A wide range of other LiDAR applications exists outside the field of this study. LiDAR is widely used in urban environments for 3D building modeling, city planning, and infrastructure monitoring (Wang et al., 2024). At smaller scales, terrestrial and mobile LiDAR enable 3D modeling and navigation, which are important for robotics and autonomous driving. In the cultural domain, LiDAR has proven valuable for archaeological discovery, where high-resolution elevation models can reveal buried or hidden structures that are invisible in conventional imagery (Masini et al., 2018).

C.2. Passive Optical Remote Sensing

C.2.1. DN to Reflectance Conversion

To convert DN to radiance L [W·sr $^{-1}$ ·m $^{-2}$ ·μm $^{-1}$], or Top of Atmosphere (TOA) radiance for spaceborne sensors, the sensor gain and bias are needed, which are sensor-specific calibration parameters. For each band b and pixel p , the conversion is a linear function:

$$L_b(p) = \text{GAIN}(b) \cdot DN_b(p) + \text{BIAS}(b) \quad (\text{C.1})$$

For spaceborne sensors, TOA radiance can be converted to TOA reflectance ρ_b , through removing illumination geometry and Earth–Sun distance effects (Airbus Defence and Space, 2024):

$$\rho_b(p) = \frac{\pi \cdot d^2 \cdot L_b(p)}{E_0(b) \cdot \cos(\theta_s)} \quad (\text{C.2})$$

Where d is the Earth–Sun distance in astronomical units (AU), $E_0(b)$ the solar spectral irradiance for the band in W·m $^{-2}$ ·μm $^{-1}$, and θ_s the solar zenith angle. The TOA reflectance is then the ratio of reflected light to the incident radiation ($0 \leq \rho_b \leq 1$). A value of 0 represents full absorption, and a value of 1 represents full reflection. For airborne systems, the sensor resides within the atmosphere, so the measured radiance only includes the atmospheric path between the surface and the airborne platform. For many applications, TOA reflectance is not sufficient, as the signal still contains atmospheric scattering and absorption. Therefore, atmospheric corrections need to be applied to obtain the surface reflectance. This way, the data can be used in time series analysis and across other sensors.

C.2.2. Passive ORS Sensor Components

A passive multispectral sensor consists of an optical system, a detector, and a scanning mechanism. The optical system focuses the incoming radiation through filters or beam splitters onto the detector. The most common detectors are Charge-Coupled Devices (CCD) and Complementary Metal-Oxide Semiconductor (CMOS). CCD detectors record the incoming photons by accumulating charge in each pixel, which is then read out and converted into a DN. CMOS detectors perform a similar conversion, but each pixel has its own readout circuit, making the data acquisition more efficient, but the data quality is lower as a result. For spaceborne sensors, two scanning mechanisms are often used: whiskbroom scanning, where a mirror sweeps across-track, and pushbroom scanning, where a linear array of detectors is used as the platform moves forward. In contrast, airborne systems can be simplified by using frame-based cameras that capture images at once. Finally, as with LiDAR, multispectral sensors often

include GNSS and IMU systems to record the position and orientation of the platform. This way, the pixels can be georeferenced.

C.2.3. Passive ORS Resolution

Spaceborne platforms can reach spatial resolutions of 10 m - 1 km while airborne platforms can reach sub-meter to 5 m resolutions (Saritha et al., 2025). The spatial resolution is defined by the Ground Sampling Distance (GSD), which quantifies the physical size of the ground area captured by a single pixel. GSD is a function of the sensor's focal length f , the altitude of the platform h , and the pixel size p of the detector:

$$\text{GSD} = \frac{h \cdot p}{f} \quad (\text{C.3})$$

A longer focal length or smaller pixel size results in a finer GSD, assuming a constant altitude. A high spectral resolution can be reached with more bands, but since fewer photons get collected, the SNR is lower, which results in a lower spatial resolution. A higher radiometric resolution allows finer distinctions between surface reflectances, but also increases data volume and can amplify noise if the signal is weak.

C.2.4. Passive ORS Data Quality

Just like LiDAR, the geometric accuracy is influenced by the GNSS and IMU accuracy. In spaceborne systems, geometric distortions may arise from orbital motion and Earth curvature, while airborne platforms are more affected by turbulence and altitude fluctuations. To correct such distortions and to ensure that each pixel corresponds to its true location on the ground, images are usually orthorectified using a Digital Elevation Model (DEM) and precise sensor position and orientation data. This step is necessary when the data is combined with other geospatial datasets or for performing quantitative analysis. The radiometric accuracy also influences the quality, which refers to how well the sensor captures true reflectance values across bands. The radiometric accuracy is influenced by the calibration accuracy and precision of three calibration processes (Saritha et al., 2025): dark current, radiometric, and spectral calibration. Dark current calibration is needed to filter out intrinsic noise in the imaging system itself. This is defined as the BIAS in Equation C.1. Radiometric calibration establishes the relationship (GAIN in Equation C.1) between the observed radiance measurements with the known radiance values of reference targets. Finally, spectral calibration ensures that the sensor accurately measures radiation across different spectral bands. Similar to LiDAR, the SNR also influences the quality. The SNR is in turn influenced by sensor design (choice between CCD and CMOS, for example), exposure time, and platform altitude. In addition, atmospheric and illumination conditions play an important role in the quality. As the atmosphere contains a mixture of gases, aerosols, and water vapor, some wavelengths will be absorbed more than others. This effect is more apparent in spaceborne sensors than in airborne sensors, as light has to travel a longer path. The quality therefore, depends on the atmospheric corrections applied to correct for this effect. Optimal illumination conditions consist of a solar angle that results in the shortest path length through the atmosphere and minimizes the appearance of shadows, which typically occurs around local solar noon when the sun is at its highest point in the sky. Spaceborne missions are typically designed to operate under consistent and optimal illumination conditions (Mark R. Drinkwater, 2007), thereby ensuring radiometric consistency. In contrast, airborne missions offer greater flexibility in acquisition timing and flight planning, but this often comes at the cost of suboptimal illumination conditions.

C.2.5. Extended Passive ORS Applications

Beyond environmental monitoring, passive ORS is also important in other fields. In hydrology, indices such as the Normalized Difference Water Index (NDWI) can be used to monitor water bodies in, for example, flood assessment or river morphodynamics. Subsequently, other spectral signatures can be used to assess, for example, water quality (Binding et al., 2012). In urban planning, multispectral imagery supports the mapping of land use changes, green cover, and infrastructure in urban areas. Conventionally, it can also be used for cadastral and land administration.

D

AHN and Beeldmateriaal Nederland Workflow and Accuracy Assessment for Saba

This appendix provides a description of the workflow and accuracy assessment of the airborne LiDAR (Section D.1) and passive multispectral ORS data (Section D.2) acquisition, which are part of the AHN and Beeldmateriaal Nederland survey for Saba. Additional system specifications on the used sensors can be found in Subsection D.1.1 and Subsection D.2.1 for the LiDAR and multispectral data, respectively.

The LiDAR data collected during the flight campaign were processed in accordance with the requirements defined in the AHN specifications for the Caribbean Netherlands (Actueel Hoogtebestand Nederland, 2023) while the delivered dataset and its quality were assessed in the quality report (Het Waterschapshuis, 2024b). Subsection D.1.4 provides a summary of the LiDAR data quality assessment. The passive multispectral ORS data collected during the flight campaign were processed in accordance with the requirements defined in the specifications for the Caribbean Netherlands defined by (Beeldmateriaal Nederland, 2023). Subsection D.2.3 provides a summary of the passive multispectral ORS data quality assessment.

D.1. AHN LiDAR Data

D.1.1. RIEGL VQ-1560 II System Specifications

The maximum operational altitude depends on the chosen PRF and laser power settings. At its maximum PRF (4 MHz), point density is maximized, but the maximum flight altitude is limited to approximately between 1700 and 2700 m. At its lowest PRF (540 kHz), the altitude extends up to 3900 m, but the point density is minimized. According to the RIEGL VQ-1560 II datasheet, average point densities can vary from 2 points/m² to 60 points/m² (RIEGL Laser Measurement Systems GmbH, 2024). In addition to discrete waveform LiDAR, the system is also capable of full waveform recording. Conventionally, the system has an integrated IMU and GNSS unit. Additional optional components include an RGB camera and even thermal and NIR cameras.

The scanning mechanism of the scanner is based on a rotating polygon mirror. The scan angle range of each channel is 60°, which results in an effective FOV of 58°. The scanning configuration can be seen in Figure D.1.

D.1.2. LiDAR Flight Execution

The exact acquisition time of the survey is not available from the data, though the total flight time was 45 minutes. The survey was required to be executed after storm season, as December marks the start of the dry season, and there should be no flooding or wet surfaces from rainfall at the time of acquisition (Actueel Hoogtebestand Nederland, 2023). The chosen flight strip overlap of 55% was sufficient to enable prescribed control tasks to be carried out according to specifications. With a FOV of 55° and

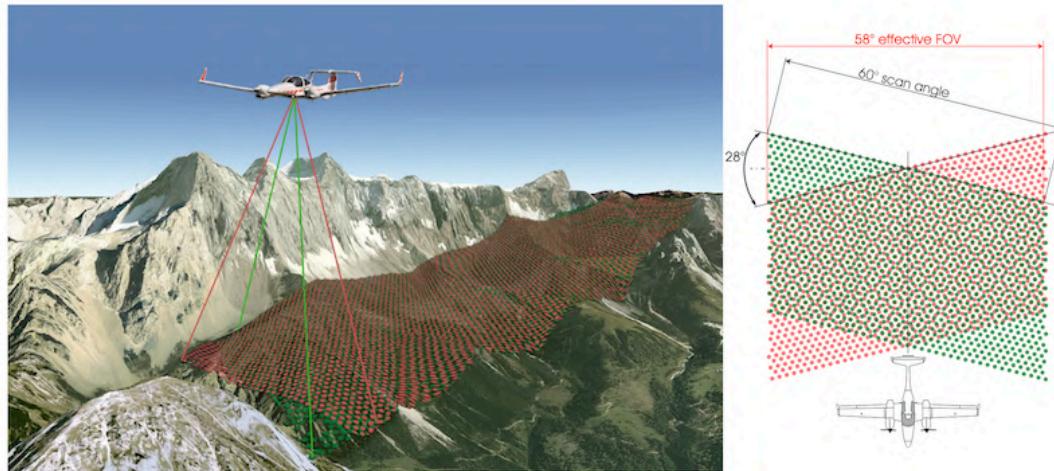


Figure D.1: The RIEGL VQ-1560 II scanning pattern. The two scanner channels are tilted 28° with respect to one another. With the rotating scanning mechanism, the effective FOV is 58° (RIEGL Laser Measurement Systems GmbH, 2024).

altitude of $h = 2000$ m, the swath width corresponds to $W \approx 2084$ m (Equation 3.1), and the effective swath width is $W_{\text{eff}} = 938$ m. Considering the longest distance across the island is roughly 6 km, this would require an up-rounded 6000 m/ 938 m ≈ 7 flight lines to cover the island. Therefore, a total of seven parallel NW–SE oriented flight lines and two cross-lines (oriented N–S and S–N, called 101 and 102, respectively) were planned. Figure D.2 shows the orientation of the flight lines.

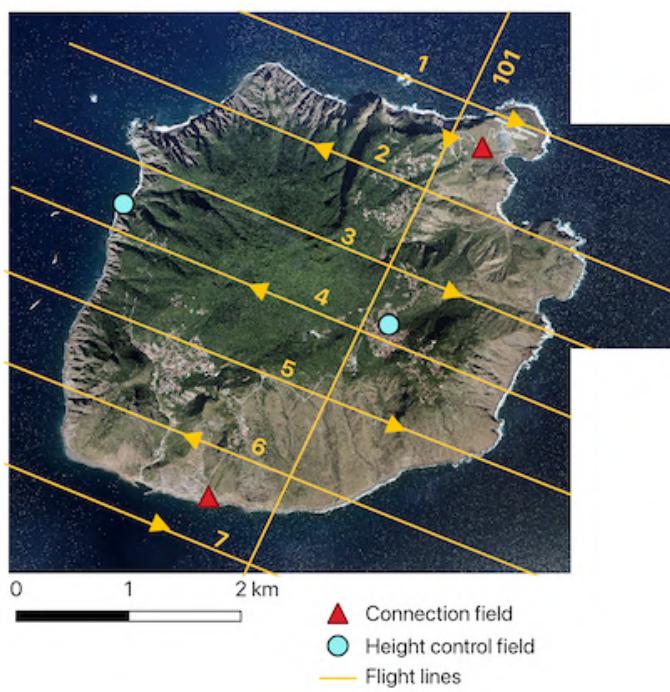


Figure D.2: Flight lines of the LiDAR acquisition estimated from the flight strips with the line ID and direction indicated. In addition, the height control fields and connection fields are also indicated. The connection fields served as reference areas for block adjustment, while the control fields were used to independently assess vertical accuracy.

As shown in the flight report in Table D.1, some flight lines were repeated due to cloud formation during the survey. The report mentions no mist around the volcano peak and only very light turbulence. Flight line 102 was flown only once, but a large cloud in the center of the line rendered the strip unusable. The reason for not re-flying line 102 remains unclear. Consequently, flight line 101 was the only line oriented perpendicular to the seven parallel flight lines.

Table D.1: Flight report summary of the LiDAR acquisition for each line.

Flight Line ID	Yaw κ	Run 1	Run 2	Run 3	Run 4
101 (N–S)	205°	clouds	ok		
1 (SE)	112°	ok			
2 (NW)	292°	thin cloud	ok		
3 (SE)	112°	cloud	cloud	ok	
4 (NW)	292°	cloud	ok		
5 (SE)	112°	ok			
6 (NW)	292°	cloud	cloud	cloud	cloud on south of line, probably outside FOV
7 (SE)	112°	small cloud on south of line, probably outside FOV			
102 (S–N)	25°	clouds			

D.1.3. LiDAR Data Processing and Product Generation

Firstly, each flight strip was processed using the IMU and GNSS data, thereby georeferencing the laser returns. From here, the point clouds are generated where each return is assigned a spatial coordinate (x, y, z) and the attributes in Table 3.2 except for attribute 8 (classification was done later). Because every line is affected by small trajectory uncertainties, the strips were then brought into alignment through a strip adjustment in *TerraMatch*. *TerraMatch* is a software package from Terrasolid, often used for the processing of LiDAR data. Subsequently, a block adjustment was carried out using the connection fields to correct for systematic vertical errors. These fields were measured with high-accuracy GPS and serve as reference areas within the dataset. By comparing the LiDAR-derived elevations to the known elevations in the connection fields, the point cloud can be adjusted to reduce systematic offsets and improve internal consistency. In contrast, the height control fields are excluded from this adjustment process and are used as independent validation sites later on to assess the absolute vertical accuracy. Once the point cloud was geometrically corrected, classification routines were executed to classify the points with a standard point class defined by the ASPRS (The American Society for Photogrammetry & Remote Sensing, 2019). Classification was done in two phases, where first an automated classification was done with *Terrasolid*, after which a manual classification followed. This manual classification was needed in places of steep terrain, along the coastline, around buildings, and in dense vegetation. The classified data was then used to generate raster products: a DTM and DSM, both at 50 cm, and resampled 5 m resolution.

D.1.4. LiDAR Data Quality

Absolute Accuracy

Absolute Vertical Accuracy

As mentioned previously, the connection fields were incorporated into the block adjustment process. Despite this adjustment, the average vertical deviation between the connection field measured heights and the LiDAR-derived heights was -11 cm. However, the stochastic errors remained low, suggesting that the dataset was internally precise. When comparing the data against the height control fields, the average deviation was consistently low (-0.005 m and +0.031 m), which is below the maximum allowed systematic vertical error of 5 cm. The standard deviations were measured as 0.021 m and 0.024 m, which meet the project's stochastic accuracy requirements. The systematic error of -11 cm from the connection fields was therefore not adjusted.

Absolute Planimetric Accuracy

The original planimetric accuracy assessment aimed to compare the horizontal positions of distinct features in the LiDAR intensity imagery, using recognizable features in their counterparts in orthophotos. However, due to insufficient signal contrast in the intensity data, reliable feature matching was not possible, even when increasing the intensity contrast. As a result, direct measurement could not be

performed. Instead, visual comparisons were conducted by overlaying the datasets and inspecting for systematic horizontal misalignments. This was done on high-contrast linear features such as the crosswalk at the airport strip. No consistent offsets were observed, indicating that the planimetric accuracy is acceptable within the constraints of the available data, though not precisely quantifiable. The maximum allowed systematic planimetric error is 8 cm, and the stochastic accuracy requirements for the horizontal component follow the same thresholds as those defined for vertical accuracy.

Relative Accuracy

Relative Vertical Accuracy

The relative vertical accuracy was assessed by analyzing elevation differences in overlapping flight strips. The project specifications required that 70% of the height differences between overlapping strips must be less than 3.5 cm, 95% less than 7 cm, and 99.5% less than 10 cm. The data showed that 95.35% of the height differences were within 3.5 cm, 98.89% within 7 cm, and 99.54% within 10 cm, thereby meeting the requirements. This was also done in flat areas, as in sloped terrain, even a small horizontal misalignment between strips can cause large vertical discrepancies. In those cases, the accuracy improved, with 96.32% of the data within 3.5 cm and 99.78% within 10 cm.

Planimetric Relative Accuracy

The planimetric relative accuracy was evaluated using roof ridge line matching, again processed through *terrismatch*. This was evaluated for both the planimetric and vertical components. The vertical component is also measured separately with an additional method, therefore there exist two relative accuracy assessments for the vertical component. For vertical relative accuracy, individual height deviations must adhere to a distribution in which 70% of the deviations are less than 3.5 cm, 95% are less than 7 cm, and 99.5% are less than 10 cm. Additionally, the standard deviation of height differences per strip overlap must not exceed $\sqrt{2} \sigma_z = 7.1$ cm. Similarly, the planimetric relative accuracy must follow a distribution where 70% of the deviations are less than 5 cm, 95% are less than 10 cm, and 99.5% are less than 16 cm. The standard deviation criterion for planimetric differences across strip overlaps is defined in the same way as for vertical accuracy. The maximum standard deviations are below the maximum allowed standard deviation, although not clearly reported, but are apparent from the standard deviation histograms. However, a distribution table of the individual deviations that confirms compliance with the required thresholds is not provided. A referenced .txt file is said to contain the relative planimetric accuracy report, but it was not available at the time of review.

Point Density

The specification for the point density required at least 10 points per square meter in 99% of 1 m x 1 m cells, where only the last and/or only returns are considered. The acquired point density met the requirements: 0.01% of cells have a point density of $D < 6$ pts/m², 0.65% fell within $6 \leq D \leq 10$ pts/m², and 99.34% of cells exceeded the threshold of 10 pts/m². According to the report, cloud cover in the southern region during acquisition led to a slight reduction in point density. This confirms that the clouds, initially believed to be outside the FOV as noted in Table D.1, were in fact within the FOV. However, it remained within acceptable limits.

D.2. Beeldmateriaal Nederland Passive Multispectral ORS

D.2.1. UltraCam Eagle 4.1 Digital Aerial Mapping System Overview

The UltraCam Eagle 4.1 digital aerial mapping system integrates both a panchromatic sensor (420–690 nm) and a multispectral CMOS sensor. Both sensors have a pixel size of 3.76 μ m. The system offers three lens options with focal lengths of 90, 120, and 150 mm, allowing a consistent GSD of, for example, 5 cm to be achieved at different flight altitudes: specifically 1205, 1596, 1995 m for the respective focal lengths (see Equation C.3). The system also contains an adaptive motion compensation, which corrects for multi-directional (ϕ, θ, κ) motion blur caused by the aircraft movement. The flight plan includes a forward overlap (along flight line) of 80% and a side overlap of 60% (between flight lines). This overlap is necessary for aerotriangulation, a method to georeference the imagery based on shared features and GPCs. A total of 10 Ground Control Points (GCPs) were selected for the campaign.

D.2.2. Passive Multispectral ORS Data Processing and Product Generation

Prior to acquisition, the camera underwent full laboratory calibration in August 2023, encompassing geometric, radiometric, shutter, sensor, and electronics calibration. Following the acquisition, the raw imagery was processed in *UltraMap*, the processing software developed by the manufacturer of the imaging system. Firstly, aerotriangulation was performed. Aerotriangulation makes use of shared features known as tie points across overlapping photographs. The tie points are automatically detected and matched between images that cover the same terrain from slightly different angles. The shift of these features between images can subsequently be used to compute the 3D coordinates of the tie points. Aerotriangulation was done in two stages: a free network adjustment and a block adjustment. A total of 651 aerial images were used for this process. In the free network adjustment, the images are geometrically linked solely through their tie points, resulting in an image block that is internally consistent. Next, the internal consistency was validated, and a full block adjustment was carried out using the GCPs. A total of 651 aerial images were used for aerotriangulation. After this, a radiometric correction was applied to ensure uniform brightness across the dataset. From the original set, 554 images were selected that did not have redundant or excessive overlap for the orthorectification process. This was done using the DTM derived from the LiDAR dataset. The locations of the image acquisitions are shown in Figure D.3.

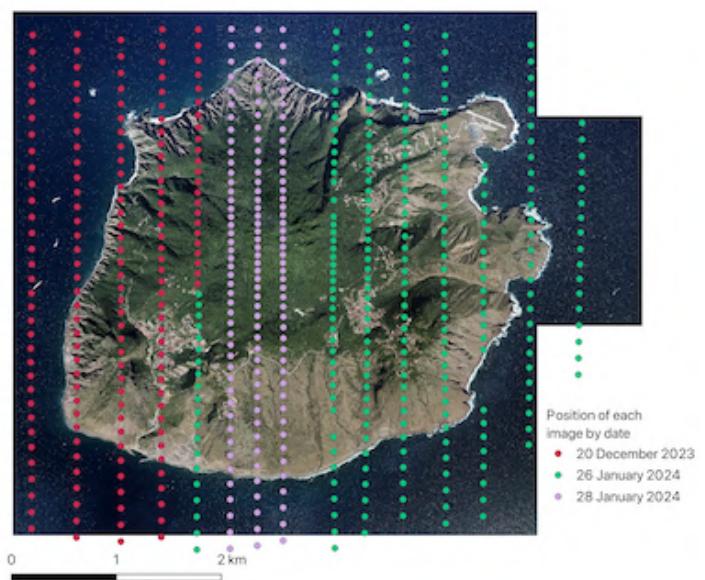


Figure D.3: Locations of the 554 image acquisitions used to produce the orthorectified images colored by acquisition date.

D.2.3. Passive Multispectral ORS Data Quality

According to Beeldmateriaal Nederland, 2023, a quality report on the aerotriangulation and the orthorectification should be provided with the data, but at the time of this study, only the quality report on the aerotriangulation (Het Waterschapshuis, 2024a) was provided. The aerotriangulation adjustment yielded root mean square (RMS) residuals on GCPs of 1.7 cm, 2.4 cm, and 5.4 cm in the x, y, and z direction, respectively. Relative geometric consistency was evaluated through tie point residuals. No systematic distortions were observed, and stereo coverage was sufficient across the island. The block adjustment demonstrated uniform accuracy across all strips, including the reflighted areas. The report also includes a checklist for quality control tied to the aerial imagery, although only indicated with a checkmark, so no additional information is provided. The delivered imagery met the following criteria: sufficient sharpness, cloud-free and cloudshadow-free ($\leq 2\%$ allowable shadow cover, not exceeded), no compression artifacts or sensor noise, natural color balance across RGB composites, sufficient contrast in both bright and dark regions, and band registration.

E

Saba's Local Coordinate Reference System (CRS)

All spatial data for this study are referenced in the local DPnet coordinate system of Saba, maintained by the Netherlands Partnership Geodetic Infrastructure (NSGI) (Lesparre, Ligt, and Huisman, 2025). DPnet Saba is a local map projection based on the International 1924 (Hayford) ellipsoid and implemented with a Transverse Mercator projection. The coordinates are expressed as eastings (x) and northings (y), measured in meters (m). In this projection, the x-coordinate (easting) increases eastward, and the y-coordinate (northing) increases northward. For the vertical reference on Saba, no geoid model is used. Instead, the ellipsoidal height in DPnet is taken directly as the physical height in the Saba height system.

When the data were published, an official EPSG code was not yet available for DPnet in QGIS or ArcGIS. An EPSG code is a unique ID used to identify a specific coordinate system, which makes it easy for mapping software to recognize and apply the correct projection automatically. Without an EPSG code, the user must define the system manually by entering the projection parameters (like projection type, ellipsoid, and units). Although documentation is provided to manually add the EPSG code for the Caribbean Netherlands to the projection database of QGIS (GeodetischelInfrastructuur, 2024), this method was not applied in practice due to compatibility issues with QGIS on macOS. Specifically, the utility `crssync`, which is required to link the projection database to QGIS, is not included in the macOS version of the application. As a workaround, the updated projection database and its associated files were directly transferred from a Windows machine to the MacBook, thereby bypassing the need for `crssync`.

F

Cloud and Shadow Identification in the Pleiades-Neo Data

The Pleiades-Neo dataset (see Subsection 3.2.3) is used in this study to calculate the NDVI (Equation 4.12). However, NDVI values from areas covered by clouds or shadows are not reliable and should be excluded. A common way to detect these cloud- or shadow-affected pixels is to compare the reflectance values from the dataset of interest with those from another dataset that is free of clouds and shadows (Tarrio et al., 2020). Under clear skies, both datasets should produce very similar reflectance values for the same surface, leading to an almost linear relationship when plotted against one another. Clouds and shadows break this relationship: clouds reflect strongly in almost all optical wavelengths, while shadows reflect very little, especially in the infrared. These deviating values can then be detected and removed by applying a threshold.

F.1. Calculation in Practice

The reference dataset used to compare the reflectance values of the Pleiades-Neo data to is from Sentinel-2 data (described in Subsection 3.2.4). In the RGB image, Sentinel-2 shows no clouds in the region where Pleiades-Neo does, making it suitable for comparison.

Before comparison, both datasets were projected into Saba's local CRS (Appendix E). To ensure consistency in spatial resolution, the Pleiades-Neo data were resampled from 10 to 20 m using bilinear interpolation, matching the resolution of the Sentinel-2 bands used. To obtain reflectance values, the DN from the Pleiades-Neo data are divided by 250, as they are stored in 8-bit format, while Sentinel-2 values were scaled by dividing by 10,000 (Hub, 2025).

Cloud detection was carried out using the visible bands (blue, green, and red), while shadow detection was based on the NIR band. For each band, Pleiades-Neo reflectance values were plotted against Sentinel-2 reflectance values. Under clear-sky conditions, points align closely along a linear relationship. In the case of the blue band (illustrated in Figure F.1), this appears as a dense diagonal cluster. The scatter plot is colored by density, with yellow indicating regions where a lot of points are located. A distinct branch with higher reflectance values for the Pleiades-Neo data is visible, corresponding to the clouds. To isolate the cloud pixels, a regression line was fitted using RANSAC algorithm via the `RANSACRegressor` function in python, which iteratively samples subsets of the data to estimate a line supported by the majority of points. RANSAC was chosen because it identifies a line supported by the majority of data points, which correspond to clear-sky conditions, while ignoring outliers caused by clouds. The fitted line, indicated in red, can be seen in the right figure of Figure F.1. Residuals (the distance of each point from the fitted line) were then calculated. For clear-sky pixels, these residuals are generally small and distributed around zero. Cloud pixels, however, yield larger positive residuals. Therefore, to separate clouds from clear-sky pixels, a percentile-based threshold was applied to the residual distribution. The 90th percentile of the residual distribution was chosen for this, and pixels with residuals above the 90th percentile were classified as cloud-contaminated, which can be seen in gray in the right figure of Figure F.1. This process was repeated for the green and red bands, and cloud pixels identified across all three bands were combined.

Shadow detection followed the same approach but applied to the NIR band. Since shadows strongly reduce reflectance in this wavelength, pixels with residuals below the 10th percentile were classified as shadow-contaminated. Initial results showed that some cloud classified pixels corresponded to roads or

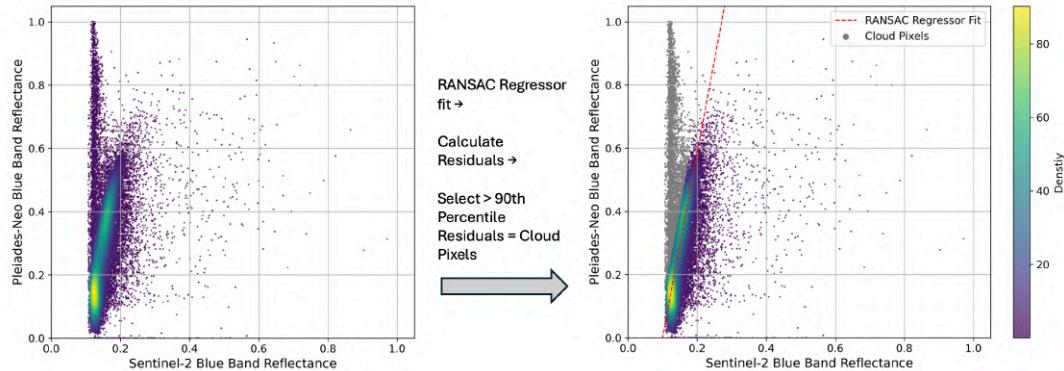


Figure F.1: Visualization of cloud detection in the blue band of the Pleiades-Neo data, using Sentinel-2 as the reference. The dense linear cluster shows clear-sky pixels, with a regression line fitted using the RANSAC method. Points that deviate strongly from this line form a separate branch, which corresponds to cloud-contaminated pixels.

urban surfaces. These false positives stem from differences in sun angle and sensor viewing geometry, which can produce variations in observed surface reflectance. In particular, roads often reflect more strongly in certain directions depending on the conditions at the time of image acquisition, which can mimic the high reflectance signature of clouds. To address this, these pixels were removed from the cloud mask.

A buffer of a commonly used distance of 200 m (Baetens, Desjardins, and Hagolle, 2019) was applied around the detected cloud pixels to ensure that cloud edges and adjacency effects of clouds in their vicinity were accounted for. The final set of identified cloud and shadow pixels can be seen in Figure F.2.

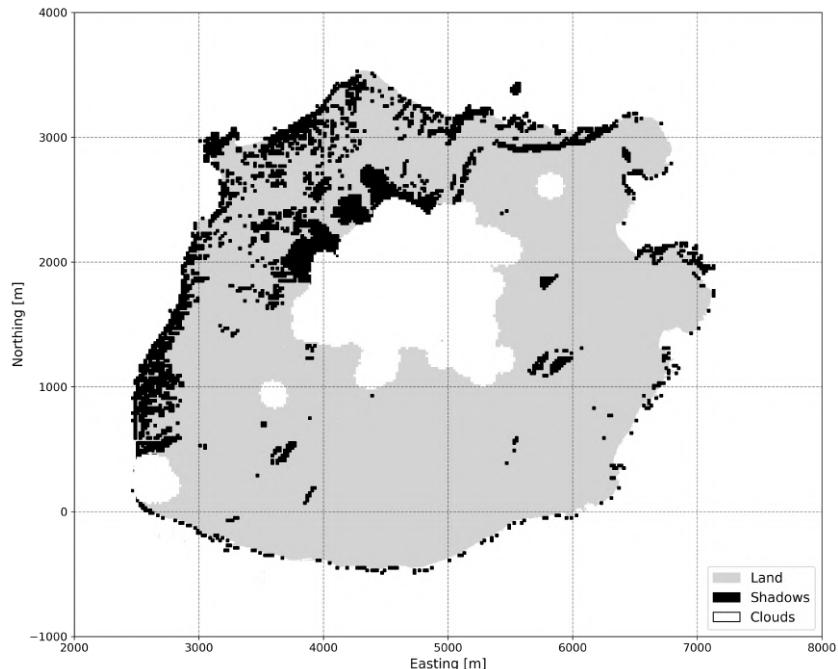


Figure F.2: The final set of identified cloud and shadow pixels in the Neo-Pleiades data of Saba with acquisition date 12-03-2024.

G

Beeldmateriaal Nederland Passive Multispectral ORS Data Properties

In Subsection C.2.4, it was mentioned that optimal illumination conditions occur around local solar noon. For the December observation date in the passive multispectral ORS data from Beeldmateriaal Nederland (described in Subsection 3.2.2), solar noon occurred at approximately 12:10 local time, and for the two January dates, around 12:25 (Time and Date AS, 2025). The acquisition times suggest that the imagery has been captured under suboptimal illumination conditions. While the choice behind the acquisition timing is not explicitly stated, optimal illumination was likely compromised in favor of cloud-free coverage, as this was one of the acquisition requirements.

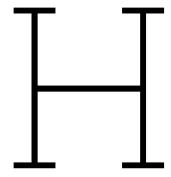
The most noticeable consequence is the presence of shaded areas in the dataset, particularly in the northern part of the island, as well as near the cliffs and ridges. The quality report states that the overall contrast between light and dark areas is sufficient for its intended use. Since the imagery was originally collected for cadastral purposes, the requirements focus on urban details such as curbstones, street drains, and building extensions, which must be clearly visible. For this type of application, the reported contrast is acceptable. However, if the imagery is to be used for land cover classification across the island, the shadows become problematic. In multispectral imagery, shadows can mask surface features and alter the apparent color of different land cover types. This can lead to confusion between classes and reduce the accuracy of classification results (Zhu et al., 2024). Additionally, visual inspection



Figure G.1: Thin clouds and haze in the RGB aerial imagery of the Beeldmateriaal Nederland passive multispectral ORS data.

reveals the presence of thin clouds and haze (Figure G.1). Although the requirements state that all imagery should be free of clouds, this likely applies primarily to urban areas, given the cadastral focus. Nonetheless, clouds are visible above the Elfin forest and along the lower northern slopes of the island

as seen in Figure G.1. For land cover classification purposes, these clouds should be considered, as they interfere with surface reflectance in much the same way as shadows.



Variable Distributions in the DTM Reliability Assessment

This appendix provides the PDFs of the 13 LiDAR-derived reliability variables shown in Table 4.1 that were used in the assessment of the DTM reliability.

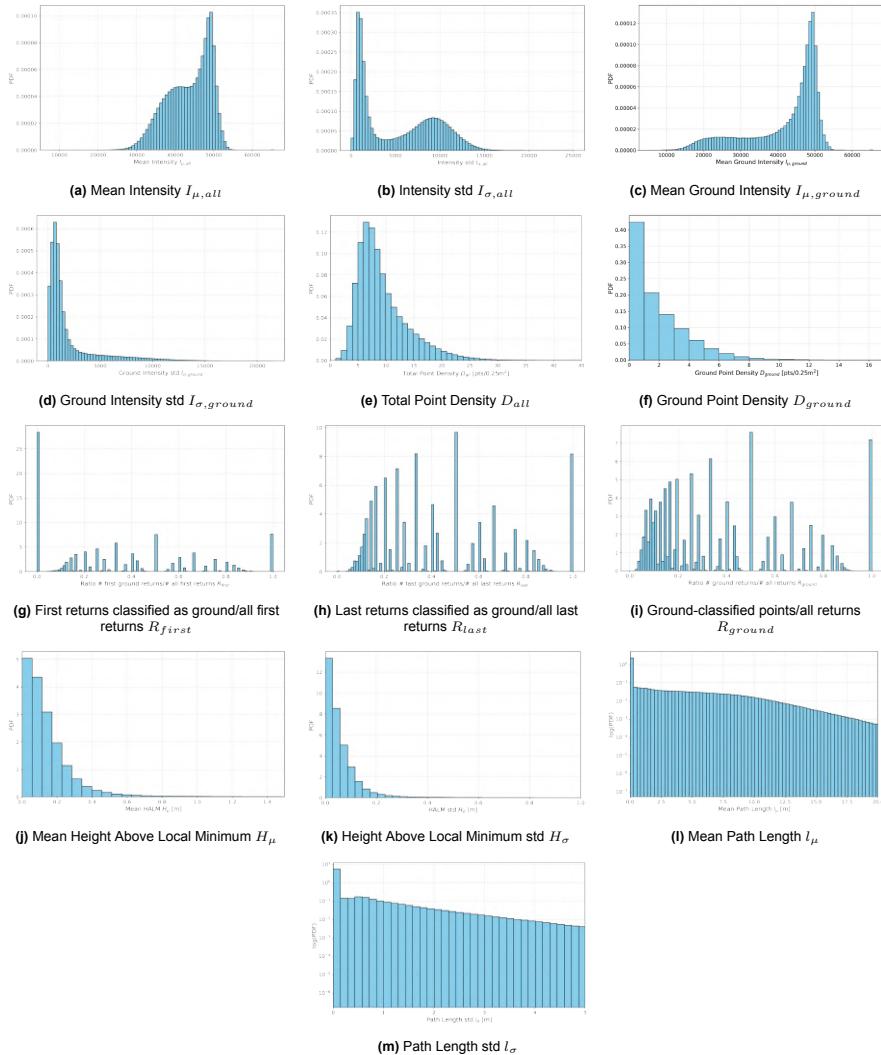


Figure H.1: PDFs of the LiDAR-derived reliability variables listed in Table 4.1.

Scoring Functions $R(x)$

This appendix provides the ECDF $F(x)$ (Equation 4.6 or Equation 4.7), the Logistic Function $S(x)$ (Equation 4.8), and the unweighted average of the two: the scoring function $R(x)$ for the LiDAR-derived reliability variables that remained after dimensionality reduction.

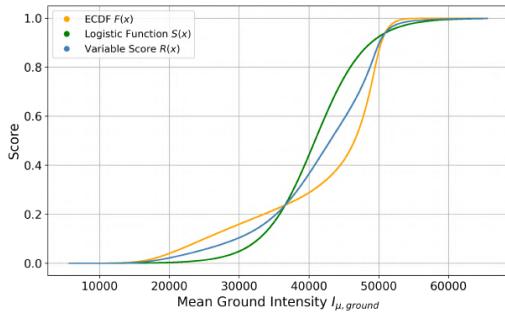


Figure I.1: The ECDF F , logistic function S , and scoring function R for the mean ground intensity $I_{\mu, \text{ground}}$.

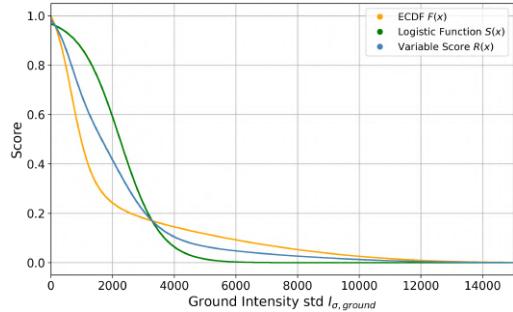


Figure I.2: The ECDF F , logistic function S , and scoring function R for the ground intensity std $I_{\sigma, \text{ground}}$.

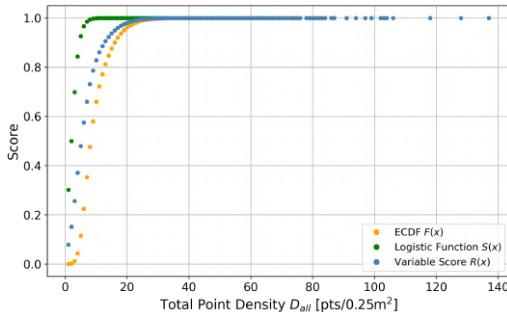


Figure I.3: The ECDF F , logistic function S , and scoring function R for the total point density D_{all} .

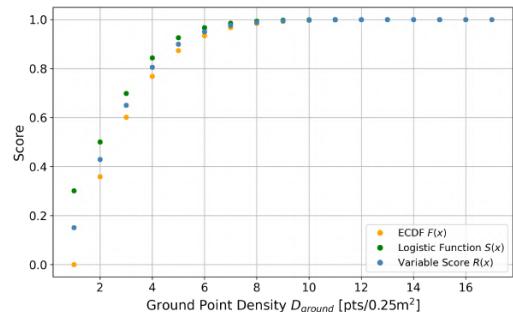


Figure I.4: The ECDF F , logistic function S , and scoring function R for the ground point density D_{ground} .

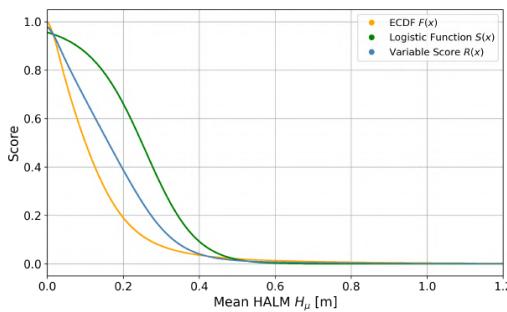


Figure I.5: The ECDF F , logistic function S , and scoring function R for the mean HALM H_{μ} .

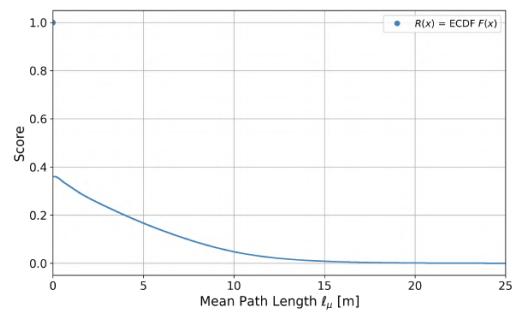


Figure I.6: The ECDF F for the mean path length ℓ_{μ} which was also used as R .

J

GMM BIC scores

Components	(3095,1710)	(3560,1660)	(3645,-260)	(4685,290)	(5385,1680)	(5905,2745)	(6295,1175)	(6435,1135)	(6625,1750)
1	7794.55	8369.06	-997.06	-821.52	5524.69	38.33	265.27	-856.52	-702.92
2	7365.03	8339.69	-2887.09	-3127.07	4321.08	-445.57	-1212.57	-3313.19	-1913.61
3	6385.10	8247.42	-3084.56	-3171.44	4145.19	-1653.96	-1387.58	-3371.30	-1932.61
4	6335.75	8251.86	-3080.19	-3173.08	4152.45	-1687.64	-1400.66	-3358.83	-1921.21
5	6345.82	8272.45	-3106.20	-3162.38	4111.14	-1668.50	-1422.29	-3349.35	-1909.36
6	6356.00	8215.18	-3091.67	-3150.51	4130.53	-1659.70	-1418.84	-3336.54	-1932.61

Table J.1: BIC scores for GMMs with 1–6 components across the selected pixels. The lowest score per pixel is highlighted.

K

GMM Fit of Normalized Height Distributions

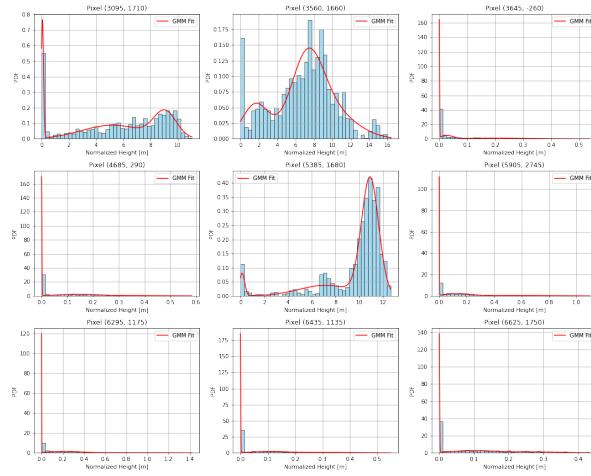


Figure K.1: GMM fit for 3 components shown on the normalized height distributions of 9 representable pixels of the island.

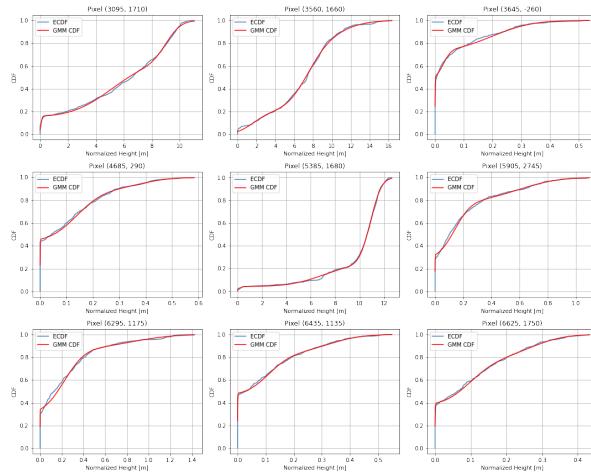
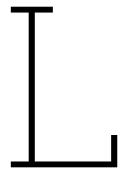


Figure K.2: GMM fit for 3 components shown on the ECDF of the normalized height distributions of 9 representable pixels of the island.



Local Incidence Angle

L.1. Definition

The local incidence angle is defined as the angle at which the laser pulse hits the surface relative to the surface normal. For ground-classified points, this angle indicates how obliquely the laser pulse struck the terrain surface, so if the incidence angle is 0° , the laser pulse hit the ground perpendicularly. When scanning vegetated areas at a larger angle, the signal has to travel a larger distance, increasing the chance that the signal dissipates before reaching the ground.

The following pixel-based variables can be derived from the local incidence angle:

- **Mean local incidence angle $\theta_{0,\mu}$:** Represents the average incidence angle at which ground points in a pixel were observed.
- **Standard deviation of local incidence angle $\theta_{0,\sigma}$:** Measures the variability of incidence angles across ground points within a pixel. A low standard deviation implies consistent observation geometry, while a high standard deviation suggests mixed viewing angles.

The local incidence angle θ_0 is calculated as:

$$\cos(\theta_0) = \cos(\theta)\cos(\alpha) - \sin(\alpha)\sin(\theta)\cos(\gamma - \beta) \quad (\text{L.1})$$

Where θ is the scan angle: the angle between the point and the nadir of the plane. α is the angle of the slope, and β is the aspect angle: the direction of the slope relative to the north. γ is the viewing angle relative to the north. The viewing angle is not given in the data, but can be calculated using the yaw angle ψ relative to the north from the plane rotation angles:

$$\gamma = \begin{cases} \psi + 90^\circ & \text{if } \theta > 0 \\ \psi - 90^\circ & \text{if } \theta < 0 \end{cases} \quad (\text{L.2})$$

L.2. Calculation in Practice

scan angle: the original data is in .laz format and when this is loaded in python it decompresses it differently, therefore the scan angles need to be multiplied by 0.006 when reading the data using laspy (The American Society for Photogrammetry & Remote Sensing, 2019) To be able to calculate the local incidence angle (Equation L.1), the slope α and aspect β are needed. Slope and aspect maps are typically derived from a DTM, but in this analysis, the reliability is being assessed of the very DTM these maps would normally be derived from, which introduces a somewhat circular problem. Using the 50 cm resolution DTM directly is not feasible anyway, as it contains substantial data gaps and, of course, even in areas with ground returns, we are specifically questioning whether those points truly represent the terrain surface or if they might result from misclassified vegetation or other artefacts. To overcome this, slope and aspect maps are derived from coarser resolution versions of the DTM. A coarser DTM will represent the broader structure of the terrain while smoothing out the local noise. However, reducing the resolution too much risks oversmoothing important terrain features.

To select the most appropriate resolution for the slope and aspect map, a comparative analysis was carried out using DTMs resampled to 10, 15, and 20 m. The slope and aspect maps can be found in ?? To choose between these three resolutions, profiles of the DTMs were compared against the 50cm DTM (Figure L.1). It can be seen that in terms of slope, the 20 m DTM indeed smooths out some terrain features, for example, the ridge in the first profile and the peak and dip in the first 20 m of the second profile. The 15 m DTM seems to over and underestimate the slope at times, for example, between 20-60 m in the second profile. The 10 m DTM seems to capture the general terrain structure of the 50 cm DTM the best.

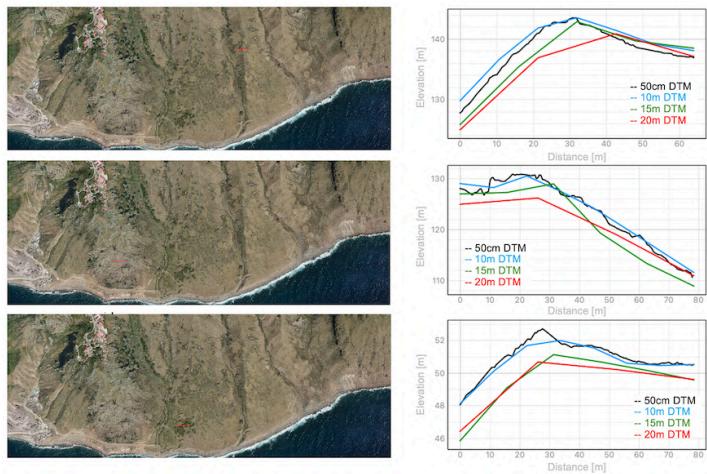


Figure L.1: Profiles of the DTM at 0.5, 10, 15, and 20 m resolution in the south of the island at three different areas: a ridge, bare rock, and vegetation (top to bottom). The red line in the imagery on the left indicates where the profile was drawn.

To be able to calculate the local incidence angle at 50 cm resolution, the 10 m resolution slope and aspect maps need to be resampled to match this resolution. For this resampling, the nearest-neighbor resampling method was used. This method assigns each (50 cm) pixel the value of the nearest coarser resolution (10 m) pixel, thereby preserving the original values without introducing interpolated ones.

Additionally, the yaw angle ψ is also not directly given in the flight report. The flight report does include the flight strips, and from these the flight lines were estimated by eye (Figure D.2) and knowing the flight direction from the flight report, the yaw angles could be estimated for each line ID. Since flight lines 1-7 are parallel, there are only three different yaw angles. The derived yaw angles can be found along the flight summary in Table D.1.

M

Dimensionality Reduction of the LiDAR Reliability Variables

The final DTM reliability score is influenced by the choice of which variables are included and which are excluded in the dimensionality reduction process described in Subsection 4.2.3, and will slightly differ depending on these choices. As mentioned in Subsection 4.2.3, the variable pairs with the highest and moderate correlation were assessed, and one from each pair was eliminated to ensure the remaining variables are independent of each other and each shows a unique aspect of DTM reliability. For the high correlation pairs (indicated in red in Figure 5.14), the following considerations were made:

- The highest correlation pair of $\rho_s^2 = 0.98$ is the mean and std of the path length l_μ and l_σ . This means that 98% of the variation in the rank order of one variable is explained by the other, which is strongly redundant. Since their correlation with other variables do not differ much between these variables, l_μ is kept because the mean path length is easier to interpret, as a low mean path length gives a higher chance of sampling the ground than a long mean path length.
- The second highly correlated pair is the mean intensity of all points $I_{\mu,all}$ and the mean intensity of ground points $I_{\mu,ground}$ with $\rho_s^2 = 0.84$. Firstly, $I_{\mu,ground}$ has a lower correlation with other variables compared to $I_{\mu,all}$, and secondly, $I_{\mu,ground}$ is a more direct measure of the actual terrain itself. Therefore, $I_{\mu,all}$ is discarded.
- The ratios (R_{first} , R_{last} , and R_{ground}) also show a high correlation with each other. Additionally, these ratios also show a moderate correlation with the ground point density D_{ground} . Considering the ratios are all calculated with a version of D_{ground} (see calculation in practice for the ratios in Subsection 4.2.1), high correlation was to be expected. Because D_{ground} directly represents the number of ground-classified points in a pixel, it is the most informative variable for assessing how well the DTM at that pixel can be estimated (as the elevation of the ground-classified points determine the value of the DTM). Therefore, all ratios were discarded.

From the moderate correlation group (indicated in orange in Figure 5.14), the standard deviation of the intensity $I_{\sigma,all}$ has a moderate correlation with two variables (standard deviation of the ground intensity $I_{\sigma,ground}$, and the mean path length l_μ). Following the same reasoning as for the intensity mean described above, $I_{\sigma,all}$ is discarded. One pair of moderate correlation remains, the HALM variables H_μ and H_σ . Following the same reasoning as for the path length mean and standard deviation, H_σ is discarded. The remaining variables reflect three aspects of the data: the acquisition conditions (D_{all} , l_μ), ground representation ($I_{\mu,ground}$, $I_{\sigma,ground}$, D_{ground}), and surface structure (H_μ).