Leveraging AI and Remote Sensing for Conservation

A Replicable Workflow for Habitat Analysis in Nyakweri Forest, Kenya

Ruari Bradburn, Chief Technology Officer, Langland Conservation

1. Executive Summary

This report details Langland Conservation's support to The Pangolin Project in Nyakweri Forest, Kenya, focusing on the implementation of an AI-driven remote sensing project to aid in the conservation of Giant Pangolins and their habitat. The project, supported by the Connected Conservation Foundation, and Airbus, demonstrates a viable workflow for applying accessible, high-tech solutions to pressing conservation challenges.

Langland Conservation, in collaboration with The Pangolin Project, has tested an innovative, accessible AI-driven process to map indicators of habitat degradation in Kenya's Nyakweri Forest. This project, made possible by the Satellites for Biodiversity Award, aims to help protect the recently discovered Giant Pangolin population facing severe habitat loss and signpost a replicable set of steps for other conservation projects.

Key achievements and impacts:

- Developed a replicable, low-code deep learning workflow using ArcGIS Pro and high-resolution Pleiades Neo imagery, operable on modest hardware ($4,000 system).
- Quantified habitat degradation: detected x of linear land partitions, x buildings, and x of remaining forest cover in a 1000km² area.
- Created an AI model to detect key landscape features, overcoming challenges like seasonal vegetation changes and lightweight fence structure identification.
- Demonstrated that sophisticated AI analysis can be performed without relying on extensive cloud computing or tech giant partnerships, making advanced conservation technologies more accessible to local projects.
- Established a template for other conservation efforts worldwide, showcasing how existing machine learning techniques and high-resolution imagery can address urgent environmental challenges with limited resources

2. Introduction

2.1. Background

Nyakweri Forest, located in Kenya's Greater Mara Ecosystem, is a critical Afromontane Forest habitat that has recently gained significant conservation attention. In 2022, The Pangolin Project made a groundbreaking discovery that dramatically altered our understanding of Giant Pangolin distribution in East Africa.

Giant Pangolins (Smutsia gigantea), the largest of all pangolin species, were previously thought to have their eastern range limit in Uganda. However, the research published by The Pangolin Project revealed the presence of these elusive creatures in Nyakweri Forest, extending their known range by 500 kilometres15 eastward. This discovery not only underscored the ecological importance of Nyakweri Forest but also emphasized its potential role in pangolin conservation.

Unfortunately, the Nyakweri Forest landscape faces severe environmental challenges. The area has experienced dramatic habitat loss over the past decade, with more than half of its forest cover disappearing. This rapid deforestation is primarily driven by the conversion of viable habitat into partitioned farmland, a process that inflicts irreversible damage on the ecosystem.

The key threats to Nyakweri Forest and its Giant Pangolin population are the following:

- Deforestation: Large-scale clearing of forest for agriculture, charcoal production, and human settlement. Giant Pangolin have a preference for dense forest habitats with abundant populations of ants and termites.
- Habitat Fragmentation: The remaining forest is increasingly divided into smaller, isolated patches. The construction of field boundaries, roads and buildings further encroaches on natural habitats and destroys the connectivity between “islands” of remaining viable habitat.
- Electric Fences: While intended to protect crops from elephants and other raiding species, these pose a significant physical threat to pangolins.
- Poaching: All pangolin species are highly valued for their meat and scales. Pangolin are the most trafficked wild mammal globally.

2.2. The Pangolin Project

The Pangolin Project is a non-profit organization based in Kenya, led by Dr. Claire Okell. Dedicated to pangolin conservation research and protection, the organization focuses on securing a future for African Pangolins in their native landscapes.

The Pangolin Project's work encompasses four main areas:

Conservation Research: Conducting practical conservation science to provide evidence-based strategies for conserving pangolins in the wild.

Sustainable Protection: Supporting rangers and anti-poaching teams by developing specific skills and knowledge for protecting pangolins and their habitats.

Community Partnerships: Empowering local communities to act as pangolin custodians through education, conservation training, and support activities.

Awareness and Advocacy: Increasing awareness about pangolins among partners, the conservation community, and society at large.

2.3. Langland Conservation

Langland Conservation is a UK-registered charity that specializes in using data analytics and technology to support a range of conservation partners. Their work focuses on three core aspects:

- Using data-driven insights to help decision-makers achieve greater results in conservation.
- Empowering others to leverage technology in conservation efforts.
- Supporting investigations to tackle organized wildlife crime worldwide.

In the Nyakweri Forest project, Langland Conservation's role was to conduct a thorough analysis of the area and develop and implement an AI-driven remote sensing solution to quantify and visualize habitat degradation.

Their support to The Pangolin Project was led by Chief Technology Officer, Ruari Bradburn, and Head of Analytics, Dr. Alice Ball.

2.4 The Satellites for Biodiversity Award

The Satellites for Biodiversity Award, a collaboration between Connected Conservation Foundation and Airbus, played a crucial role in this project. This initiative, which bridges conservation organizations and satellite imagery providers, provided Langland Conservation and The Pangolin Project with access to high-resolution Airbus Pleiades Neo 0.3m RGB imagery. This imagery was essential for developing the AI model and analysing the Nyakweri Forest landscape.

The Connected Conservation Foundation, particularly through director Sophie Maxwell, offered invaluable support throughout the project. This collaboration exemplifies how partnerships between technology providers and conservation organizations can significantly enhance the effectiveness of biodiversity protection efforts.

3. Project Objectives

3.1. Quantify Habitat Degradation

The primary aim was to understand and quantify the extent and rate of habitat degradation in the Nyakweri Forest Landscape, part of the Greater Mara Ecosystem in Kenya. This objective was crucial given the recent discovery of Giant Pangolins in the area and the rapid conversion of viable habitat to farmland.

3.2. Develop AI-Driven Conservation Tools

A core objective was to develop a Deep Learning model capable of detecting key landscape features:

1. Artificial Land Partitions & Fence Lines
2. Tree Cover
3. Buildings

This AI model would serve as a powerful tool for analysing large areas quickly and accurately, providing vital data for conservation efforts.

3.3. Inform Conservation Strategies

By providing detailed, up-to-date information on the landscape, the project aimed to raise awareness, help inform protection strategies such as fence de-electrification, and help focus efforts on conserving remaining areas of woodland.

3.4. Demonstrate Accessible Technology Use

A key objective of this project was to showcase how high-resolution satellite imagery, and AI technologies could be leveraged effectively by conservation projects with limited resources. This approach aimed to serve as a model for conservation initiatives worldwide, demonstrating that sophisticated analysis can be performed without relying on extensive cloud computing resources.

The project aimed to test and establish a low-code, end-to-end deep learning workflow that could be applied to other conservation challenges. This workflow was designed to be accessible to projects with modest hardware and technical skills, democratizing the use of advanced technology in conservation.

4. Methodology

4.1. Imagery

4.1.2 Data Acquisition

The Project utilised Airbus Pleaides Neo imagery. All imagery was provided by the Connected Conservation Foundation through partnership with Airbus.

4.1.3. Imagery Coverage:

The imagery was received in multiple tranches:

a. Tranche 1: Covered approximately 450km² of the western side of the study area. Comprised of 3 images. Dated 24/08/2023.

b. Tranche 2: Covered approximately 500km² of the eastern side of the study area. Comprised of 2 images. Dated 02/10/2023.

c. Tranche 3: Covered approximately 130km², including all of the designated Priority Conservation Area on the eastern edge of the study area. Comprised of 1 image. Dated 01/03/2024.

It's worth noting that the imagery tranches were taken under different seasonal conditions, with Tranches 2 and 3 appearing much greener than Tranch 1.

Raw Imagery Products		
	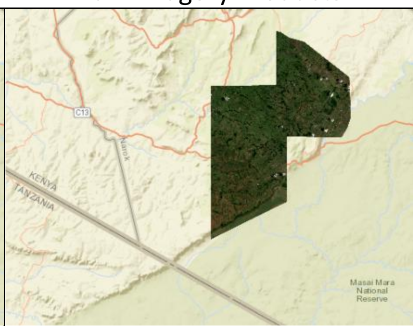	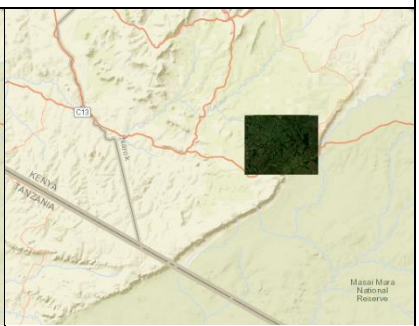
Tranche 1	Tranche 2	Tranche 3

Two merged products were created:

a. Merge 2x1: Created by combining Tranche 2 on top of Tranche 1.

b. Merge 3x2x1: Created by combining Trance 3 on top of Merge 2x1

Merged Imagery Products

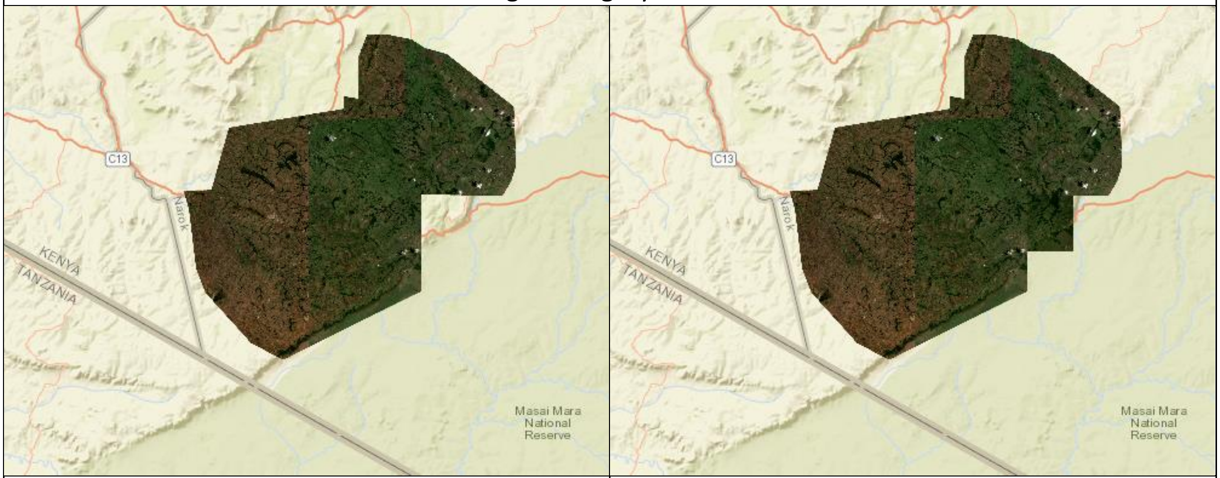

Merge 2x1 | Merge 3x2x1

The resultant file rasters were very large in size. The .tif file for Merge 3x2x1 that was passed to the deep learning model for inference totalled 291GB.

4.1.4. Imagery Characteristics

The imagery was a 4-band pansharpened product, consisting of:

a. RGB (Red, Green, Blue) bands

b. NIR (Near-Infrared) band

The resolution of the pansharpened imagery was 0.3m, resulting from the fusion of:

a. 1.2m multispectral (RGB + NIR) data

b. 0.3m panchromatic data

Imagery Examples

Tranche 1 | Tranche 2 | Tranche 3

4.1.5. Multispectral Capabilities

A single band Normalised Difference Vegetation Index (NDVI) version of Merge 3x2x1 was created, totalling 145GB.

Comparison between RGB and NDVI Enhanced views	
	
RGB Imagery from 3x2x1	3x2x1 “NDVI Enhanced” View

An “NDVI” enhanced view of the rasters were created by overlaying a semi-transparent NDVI view of 3x2x1 over the RGB version of 3x2x1. This emphasised the areas of negative values in red, areas of positive values in green. Areas between -0.2 and 0.2 were left transparent. This was mainly used as a working view to assist with the manual production of the training dataset.

Most deep learning models do not support 4-band imagery, meaning the 3-band RGB images were used for the creation of training datasets and deep learning models.

Whilst it fell out of the scope of this study, it would be highly desirable to export the “NDVI Enhanced” View to its own 3-band RGB raster and create a dedicated training dataset and deep-learning model. The performance of this model would then be compared against the performance against the raw RGB version.

4.2. Hardware

The project aimed to demonstrate that sophisticated AI analysis for conservation can be performed using a consumer desktop without relying on external cloud computing resources. This approach significantly simplified the workflow, but it's important to consider the trade-offs between desktop and cloud systems.

Considerations for Desktop Hardware

- Essential for general data handling, pre-processing, post-processing, and GIS analysis. Such a system, (less the GPU) would be likely required regardless.
- One-time purchase cost; NVIDIA GPU pays for itself in approximately 135 days of continual use.
- Provides a single tool for end-to-end execution of deep learning workflows and all supporting steps

Considerations for of Cloud Systems:

- Scalable on demand
- Supports multiple GPU clusters for faster training and inference
- Access to high grade commercial processors and rapid training and inference
- Portable, can be run remotely
- Significantly higher skill requirement to configure and execute
- Likely cheaper inference for a single project
- Many providers offer limited disk quota (e.g., RunPod's 40GB), which can be prohibitive for large imagery datasets. In most cases the raster will be loaded directly into memory using libraries like rasterio or GDAL. Such memory availability is usually only available on which would high-end, multiple GPU configurations which would have significant cost implications.
- High costs for active GPU-enabled pods during basic configuration and file upload and download. Idle uptime must be minimised.

Given the experimental nature of this project, in addition to the desire to identify a transferrable and replicable workflow within conservation, a local configuration was selected for this project.

The total cost of the system was $4,025.

The specifications of the system used are provided here, with prices to illustrate accessibility implications:

Processor: 13th Gen Intel(R) Core(TM) i5-13600K, 3.50 GHz - $320

- This mid-range modern processor provided sufficient computational power for general data processing and model deployment.

Memory: 64GB Corsair Dominator Platinum DDR4 6600 MT/s – $400

- Ample RAM allowed for efficient handling of large datasets and image processing tasks.
- However, being able to load the whole raster into memory would be highly desirable. This would require a dedicated build that would go beyond a consumer-grade desktop.

Storage: Samsung 990 Pro 4TB NVME SSD - $350

- High-speed storage facilitated quick data access and write operations, crucial for working with large imagery files. Use of an NVME SSD drastically improved performance, particularly with handling of the training datasets, which are comprised of a very large number of small files.

Graphics Processing Unit: NVIDIA GeForce RTX 4090 24GB GDDR6 (Gigabyte Aero version) - $2,250

- The GPU was the cornerstone of the system, providing the necessary computational power for training and running deep learning models.
- When training on larger training datasets and running inference on large imagery the use of a single NVIDIA 4090 did result in longer than desirable compute times.

- A system using multiple NVIDIA 4090s would offer a significant improvement and could present a cheaper alternative (4 x $2,250 = $9,000) to using a commercial grade alternative like the NVIDIA H100 ($42,000).

Other - $705

- The system was built in a consumer desktop case ($150), complete with a reliable 1200W power supply ($285), CPU watercooler ($160) and an ample number of high-performance case fans (5 x $22 = $110).

4.3. Software

The project primarily utilized ArcGIS Pro as its core software environment, a choice that significantly influenced the workflow and accessibility of the project:

4.3.1. ArcGIS Pro Selection:

- ArcGIS Pro was chosen as the primary software environment for its comprehensive toolset and accessibility to non-profit organizations.

- While private enterprises and governments pay a premium for its features, charitable enterprises can access it at a significantly reduced cost, making it a highly attractive option for conservation use.

- ArcGIS Pro provided an integrated environment for the entire workflow. It contains all the necessary tools to prepare training data, train deep-learning models, and deploy them.

- This choice of software contributes to the replicability of the workflow, making it accessible to other conservation projects with limited resources.

4.3.2. Deep Learning Capabilities:

- ArcGIS Pro has a dedicated Deep Learning Library, which includes PyTorch-based deep learning algorithms. This includes a range of RCNN (Region-Based Convolutional Neural Network) type architecture used in this project.

- Various versions of the ResNet base architecture were used during research, with the final model produced using ResNext-101

4.3.3. Post-processing Capabilities:

- ArcGIS Pro provided tools for post-processing the model outputs, including dissolving overlapping output features, removing features below certain size and confidence thresholds and smoothing, simplifying, and regularizing output polygons.

4.3.4. Visualization and Analysis:

- The software's GIS capabilities were utilized for visualizing results and conducting spatial analyses on the detected features.

4.4. Deep Learning Model

4.4.1 Deep Learning Approaches

In the field of deep learning for computer vision, several approaches are available for analysing satellite and aerial imagery, each with its own strengths and use cases. The main approaches include:

- Image Classification: Assigns a single label or multiple labels to an entire image.
- Semantic Segmentation: Classifies each pixel in an image into a predefined category.
- Object Detection: Identifies and locates multiple objects within an image.
- Instance Segmentation: Combines object detection and semantic segmentation, identifying individual instances of objects.
- Regression: Predicts continuous values (e.g., wind speed, tree height, soil moisture) by analysing series of images.
- Change Detection: Identifies differences between images of the same scene taken at different times.

For our project in the Nyakweri Forest, we selected Object Detection as the most appropriate approach. While Semantic Segmentation could have been another suitable method, offering pixel-level precision for land cover types, we chose object detection due to its ability to provide both classification and localization of specific features. This choice was driven by several factors:

- Feature-specific analysis: Object detection allows for easier identification and analysis of specific landscape features (e.g., land partitions, fences, buildings, tree cover).
- Quantification of discrete elements: Enables straightforward counting of individual instances (e.g., buildings) and measurement of dimensions (e.g., fence lengths).
- Handling of linear features: Better captures orientation and extent of linear elements like fences and land partitions, crucial for assessing habitat fragmentation.
- Spatial relationships: Facilitates analysis of spatial relationships, including overlap, between different landscape elements, aiding in habitat connectivity assessment.
- Ease of analysis: Object detection outputs (bounding boxes with labels) are easier to post-process and integrate into GIS workflows compared to pixel-level segmentation data.

4.4.2 Model Architecture Selection

Within the object detection paradigm, we chose to use the ResNet (Residual Network) series of architecture, specifically ResNet-152 initially and later ResNext-101, as the backbone forour model. This choice was based on several factors:

- Deep Architecture: ResNets can be very deep without suffering from the vanishing gradient problem, allowing for more complex feature learning.
- Residual Learning: The skip connections in ResNets help in training deeper networks more effectively.
- Performance: ResNets have consistently shown strong performance in various computer vision tasks, including object detection.

4.4.3. Acknowledging Transformer Approaches

It's important to note the distinction between two major approaches in computer vision: Convolutional Neural Networks (CNNs) and Transformers.

CNNs:

- Use convolutional layers to extract features from images.
- Are generally faster and more memory-efficient for object detection and segmentation tasks.
- Work well with smaller datasets and are more interpretable.

Transformers:

- Originally designed for natural language processing, now adapted for vision tasks.
- Use self-attention mechanisms to process the entire image at once.
- Can capture long-range dependencies in images more effectively.
- Often require larger datasets and more computational resources.

The trade-offs between these approaches are significant:

- Performance: Transformers can achieve higher accuracy in some tasks but may struggle with small objects or fine-grained details where CNNs excel.
- Hardware Requirements: Transformers typically require more memory and computational power, eliminating the possibility of their use in a consumer desktop environment for this kind of workload.
- Training Data: Transformers often need larger datasets to perform well, while CNNs can achieve good results with smaller, more focused datasets.
- Inference Speed: CNNs are generally faster for inference, which is crucial when processing large satellite images.
- Flexibility: Transformers are more adaptable to different types of input data and can handle variable-sized inputs more naturally.

Given our project's constraints and objectives - particularly the aim to develop a workflow operable on modest hardware - we opted for the CNN-based approaches using R-CNN architectures.

This choice allowed us to achieve a balance between model performance and hardware efficiency, making the technology more accessible to conservation projects with limited resources. These approaches also aligned well with our available training data size and the specific nature of our segmentation and detection tasks in satellite imagery.

4.5. Training Dataset Preparation

The preparation of a high-quality training dataset was crucial for the success of the deep learning model. This process involved several key steps and considerations. Image chips measuring 224x224 pixels (the recommended size for our base architecture).

Once features were fully labelled, a training dataset with a structure tailored to our desired architecture was created using the “Export Training Data for Deep Learning” tool within ArcGIS Pro.

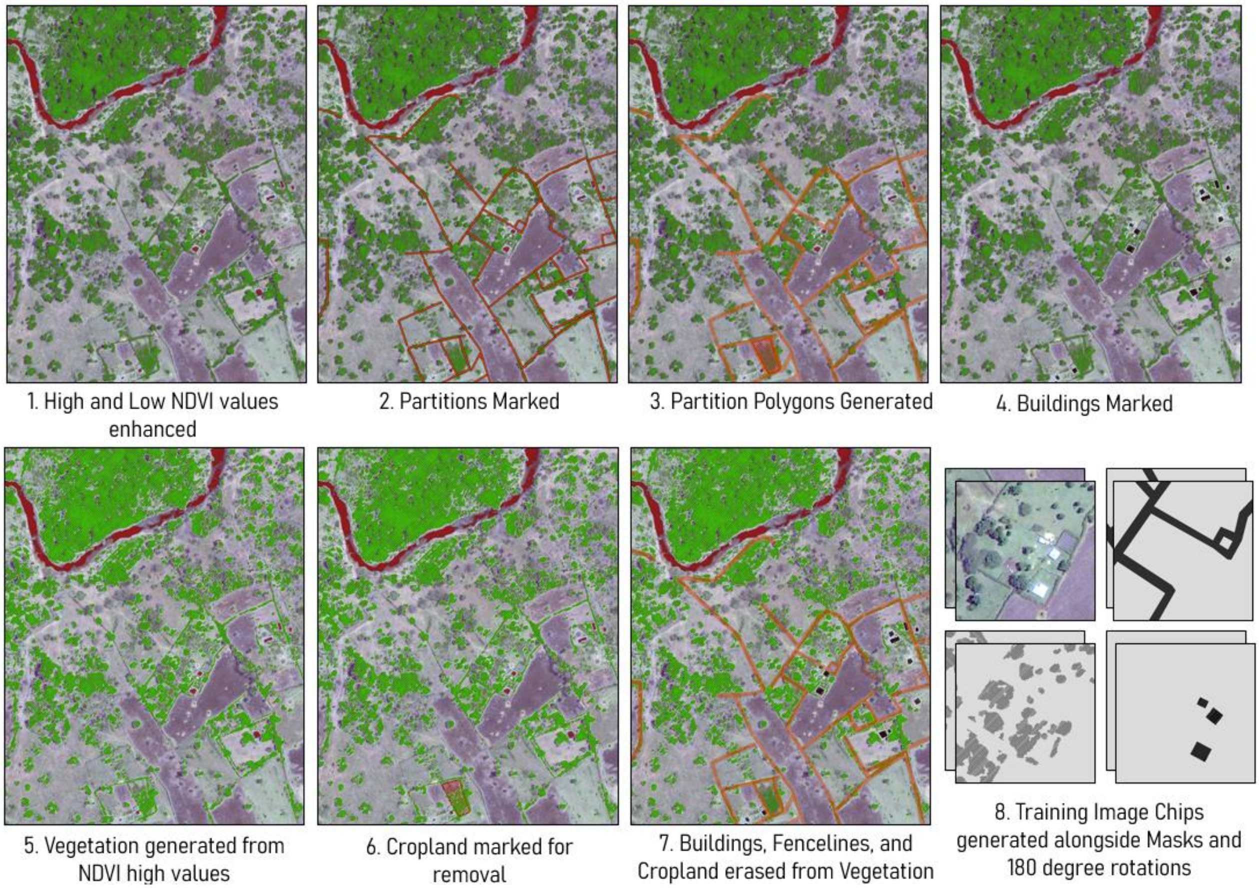

4.5.1. Dataset Scope, Size and Composition

- Approximately 7.5% of the total imagery provided was labelled for training.
- This was distributed across 7 sample areas, totalling about 78km² across the three imagery tranches. These areas were manually selected for the a high density and diversity in the desired classes, as well as having a high level of representation of the various locale types found within the target geography.
- The labelling process took over 80 hours to complete, with approximately half conducted by two volunteers.
- Partially labelled image chips (where some features were labelled but others were not) were found to have a strong negative impact model performance in early tests.

4.5.2. Feature Classes

The training dataset included the following distinct classes:

1. Tree cover
2. Buildings
3. Linear land partition features
4. Fencelines (separated from other partitions in the final version, with visible posts and lack of vegetation)
5. Roads (added in later iterations)
6. Surface Water (added in the final version)

7. Clouds (added in the final version)

In the context of this study, we distinguished between two types of linear land division features: partitions and fences. This distinction was important for understanding the landscape structure and potential impacts on wildlife movement.

Partitions are primarily solid, vegetated structures that serve as boundaries between different land areas. Key characteristics include:

- Typically appear as hedges or dense vegetation lines
- Often wider and more visually prominent in satellite imagery
- May include a mix of trees, shrubs, and other plants
- Can provide some habitat and cover for wildlife
- Usually detectable through their continuous, linear vegetation signature

Fences, on the other hand, are artificial structures with different visual and ecological characteristics:

- Primarily detected through posts repeating at regular intervals
- Generally thinner and less visually prominent than partitions
- Often lack significant vegetation along their length
- More likely to be electrified, posing a direct threat to wildlife
- Detectable through the pattern of posts and, in some cases, visible wire or netting

The distinction between these two types of features is crucial for several reasons:

- Wildlife impact: Fences, especially when electrified, can pose a more significant barrier to wildlife movement compared to vegetated partitions.
- Habitat value: Partitions may offer some habitat value and connectivity for certain species, while fences generally do not.
- Detection challenges: The visual differences between partitions and fences necessitated different approaches in the AI model for accurate detection and classification.
- Conservation strategies: Understanding the prevalence and distribution of fences versus partitions can inform different conservation strategies, such as removal or de-electrification of the bottom strands of wire.
- It's worth noting that despite these distinctions, the detection and classification of these features posed significant challenges for the AI model, particularly for fences due to their subtle visual signature in satellite imagery. This led to the development of a specialized fence detection model to improve accuracy in identifying these critical landscape features.
- In many cases, there may be an underlying fence beneath another partition, but the structure has been overgrown with vegetation. This creates a hybrid feature that combines characteristics of both fences and vegetated partitions, making clear classification difficult.

4.5.3. Labelling Methodology

- NDVI values were used to assist with the preparation of the training dataset, particularly for vegetation-related features. These were then manually altered for quality control.
- Polygons were manually drawn around buildings, clouds, and water features.
- Linear features like land partitions were initially drawn as lines and later converted to polygons for more efficient dataset production. All fences were given a buffer of 2m either side.
- For partitions comprised primarily of thick vegetation, a dynamic approach was used to test and adjust buffer sizes based on their NDVI values, saving time manually draining polygons around each linear feature, but also being sensitive to the anticipated thickness of the feature.

1. Partitions Drawn | 2. Mean NDVI calculated within buffer | 3. New buffer size interpolated according to mean NDVI | 4. Polygons simplified

- In the final model a 7-class polygon dataset was used alongside the large Merge 3x2x1 .tif file.

4.5.4. Data Augmentation

The model trained on Tranche 1 saw a noticeable drop in performance (particularly when applied to vegetation) when applied to Tranche 2. This is due to the different seasonal conditions.

To enhance the model's performance and generalization capabilities, several data augmentation techniques were employed:

180-degree rotations were produced within the “Train Deep Learning Model” to make the model agnostic to lighting / shadow conditions.

When training the model additional augmentations were applied. With each chip (including the 180-degree rotations) being passed to a further five augmentations (zoom, crop, brightness, rotation, saturation applied at random levels between set thresholds). This meant for each original imagery chip, an additional 11 variations were trained within the model.

These techniques, combined with additional training data taken from Tranche 2, and Tranche 3, made the later models developed far more robust to new and unseen data.

4.6 Performance

Generally, performance of the output model is strong for the wider area. Several classes had performance that makes them highly suitable for accurate mapping. Vegetation and Buildings, and to a lesser extent, Water and Roads, all performed well enough that the resultant layers can be considered as reasonably reliable layers of the area for mapping and analysis.

Partitions and Fencelines did not perform to the same standard, and while the model is useful for highlighting areas where Fences and Other partitions are likely, it did not perform strongly enough to produce authoritative mapping of these features. Nonetheless it plays a useful role in their detection and general mapping of their prevalence across the area in a range of more general visualisations.

The final model, trained over 20 epochs, achieved an average precision of 44% across all classes. Allowing the model to run for a further 10 epochs would likely have marginally improved the performance by a further 1-3% precision but this fell outside the scope of this study.

However, it's crucial to understand what precision means in the context of object detection and how it relates to the model's practical performance.

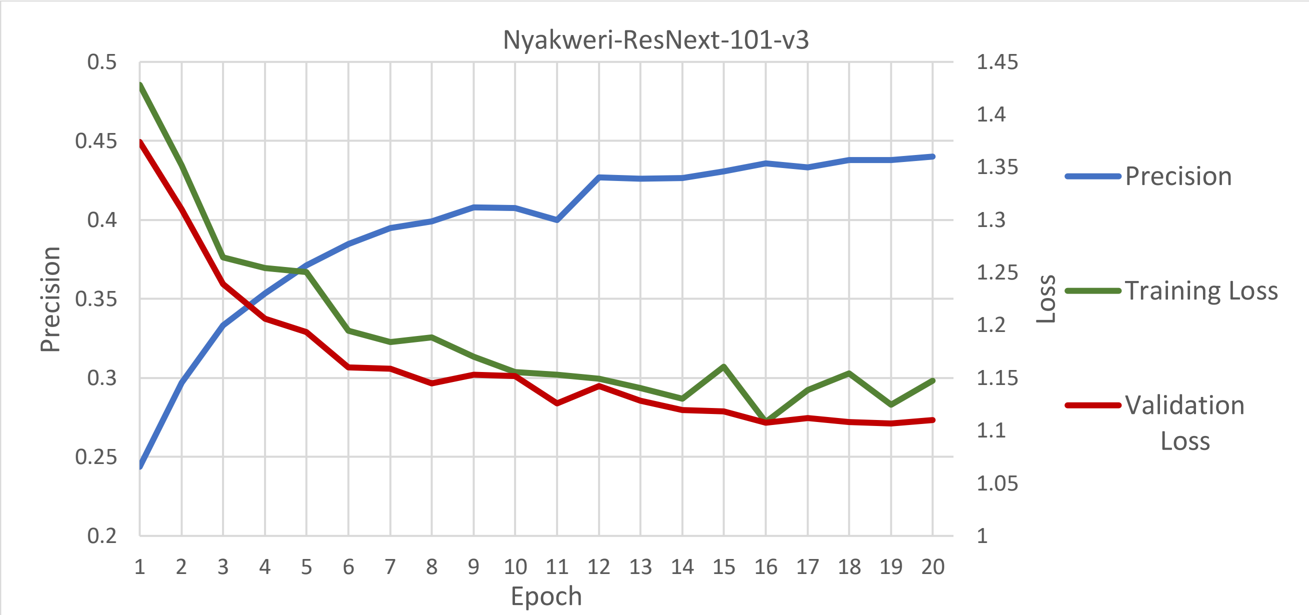

4.6.1 Understanding Precision and IoU

In object detection tasks, precision is intimately tied to the concept of Intersection over Union (IoU). IoU measures the overlap between the predicted bounding box and the ground truth bounding box. Precision, in this context, isn't simply about the number of correct detections, but rather how accurately the output polygons align with the actual feature geometries.

IoU is calculated as:

IoU = Area of Overlap / Area of Union

A higher IoU indicates a better match between the predicted and actual object boundaries. Typically, a detection is considered "correct" if its IoU exceeds a certain threshold (often 0.5, but this can vary).

4.6.2 Class-Specific Performance

The performance of each class, within the 7-class general model is as follows:

1. Vegetation: 50.7%
2. Buildings: 71.6%
3. Partitions: 16%
4. Fencelines: 4%
5. Roads: 27%
6. Water Bodies: 23.3%
7. Clouds: 55.8%

The dedicated fence detection model performed at 16%, which will be further improved upon further in a subsequent iteration, utilising the newly detected fences as additional training data.

4.6.3. Interpreting Performance

It's important to note that the seemingly low precision for some features doesn't always indicate poor detection. In reality, fencelines are very thin, linear features. For practical detection purposes, larger buffer zones were used around these features allow nearby contextual information to support their detection. While this approach successfully brings fencelines to our attention, it significantly impacts the precision score due to the mismatch between the oversized output polygons and the actual thin features.

This scenario underscores a key point: the precision score, as calculated based on IoU, may not always reflect the practical utility of the model for certain feature types. In our use case, detecting the presence and location of fencelines is sufficient, even if the precise boundaries of the resultant polygons don't perfectly match the physical footprint of the fenceline across its thickness.

Nonetheless a score of 4% for fencelines fell significantly below what was hoped and these features were retrained for 30 cycles in a dedicated model with fine-tuned augmentation parameters. An expanded training dataset with 65,000 fence image chips was used. This improved the precision score from 4% to 16%. Whilst this performance makes the model much more useful for detection, the performance still falls short of a definitive mapping solution, particularly due to the presence of small tracks frequently triggering false positives.

This could potentially be offset in the future by using a higher number of negative samples during training, exposing the model to a greater range of images that do not contain fencing but do contain visually similar objects.

Comparison between Raw Imagery and Successive Fence-Detection Models

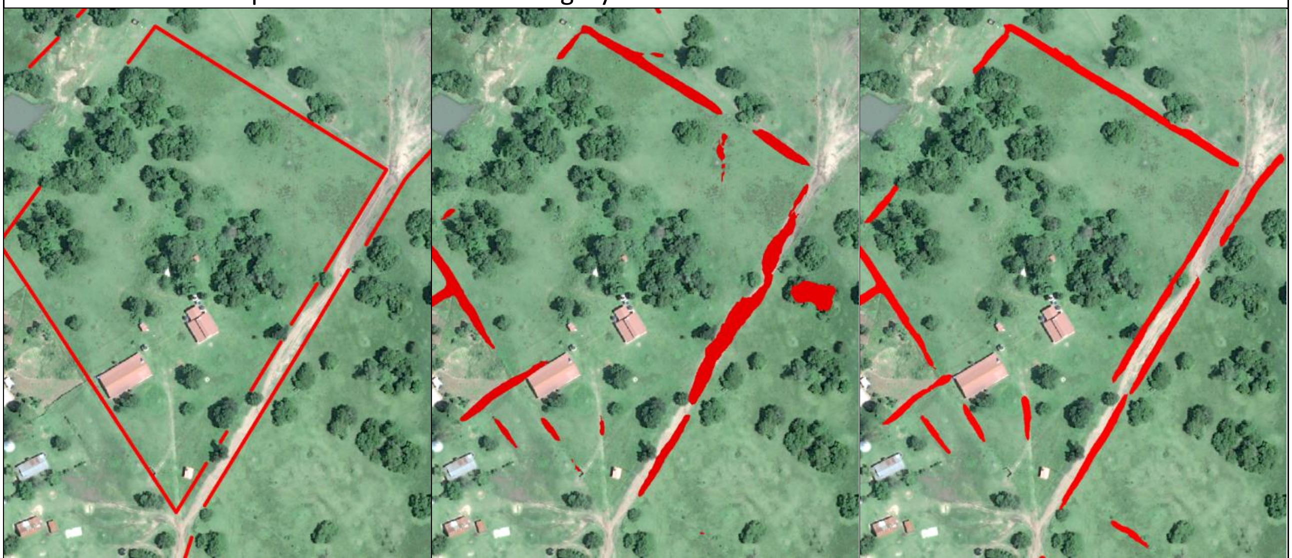

Manually Drawn Fences	General Model (4% precision)	Dedicated Model (16% precision)

4.6.3. Practical Implications

When considering the overall precision of 44%, it's essential to consider:

- The complexity of satellite imagery analysis, particularly with the feature types chosen
- Visual similarity between linear feature types
- The specific requirements of our use case
- Top performing RCNN models placed within the 30 – 55% precision range against the COCO benchmarking dataset. DETR Models, using transformers, performed as high as 66, but were considered impractical for this study due to their computational cost[1].

The model's strong performance in detecting buildings (71.6% precision) and vegetation (50.7% precision) suggests it's particularly effective for providing definitive footprints of these larger, more distinct features.

For linear features like roads (27% precision) and fencelines (4% precision), the lower scores likely reflect the challenges in precisely delineating these narrow objects rather than a failure to detect them.

In practice, the model's performance at a low confidence threshold has been observed to be very good, indicating that it's effectively identifying features of interest, even if the exact boundaries aren't perfectly aligned.

This discrepancy between numerical precision and practical utility highlights the importance of considering domain-specific evaluation criteria and the end-user requirements when assessing model performance in specialized applications like satellite imagery analysis.

4.6. Limitations

4.6.1. Feature Detection Challenges

- Reliable detection of lightweight fence line structures proved difficult, necessitating a broader approach to detecting linear partitions. This may have led to some inaccuracies in identifying specific types of land partitions. Overall the model detections of fences cannot be relied on to provide definitive mapping of features, but are useful for providing a more generalised risk map for detections.
- The high degree of variation in fencing structures and the incorporation of vegetation in partitions complicated the classification process, potentially leading to misclassifications or missed detections.
- In many cases many features were detected as a mixture of both hedge partitions and fencelines (visible posts, no vegetation).
- In the case of linear shapes, the output polygons are usually significantly larger than the shape itself. Where two partitions run in parallel (e.g. either side of a small road) they frequently

1 https://paperswithcode.com/task/object-detection

overlap, making calculations such as total fence length impossible to definitively calculate. In future testing with various buffer sizes would be highly desirable to determine an optimum buffer size.

- In some cases, very small tracks and boundaries between neighbouring fields were erroneously classified as linear partitions. Harsh earth banks to the sides of larger roads also were frequently misclassified. Generally, performance with a high confidence threshold (e.g. 50%) resulted in lots of missing linear detections and inference with a low confidence threshold (e.g. 25%) resulted in a large number of false positives, mostly from tracks and road edges.

Example of Erroneous Fence Detections

- In some cases, the output correctly detected a feature but the mapping of the polygon to the features boundaries was coarse, requiring additional post-processing and manual review.

Examples of Fences alongside satellite viewpoint

 	
Example 1	Example 2

4.6.2. Data Limitations

- The imagery tranches were taken under different seasonal conditions, which initially posed challenges for model generalization across different vegetation states. This could affect the accuracy of vegetation-related classifications across seasons.

- The training dataset, while comprehensive, only covered about 7.5% of the total imagery provided and covers a very specific geography. This limited sample size may restrict the model's performance in areas with significantly different characteristics from the training data. Future tests will apply the same model to imagery of the greater Mara landscape to tests its applicability in neighbouring areas.

- The project provides a snapshot of the landscape at specific points in time. Without regular updates to the imagery and model, the analysis may quickly become outdated in rapidly changing environments like Nyakweri Forest.

- Only a few areas overlapped between Tranche 1, 2, and 3, preventing a comprehensive change analysis being conducted over the whole area without additional imagery at a later date.

4.6.3. Practicality Limitations

- The raw output cannot be relied upon as a definitive map without further processing and interpretation. Whilst vegetations and buildings are provided with a high degree of confidence, linear features, particularly fence lines proved challenging to reliably detect.

- While the model can detect physical features, it may not directly translate to habitat suitability for species like the Giant Pangolin without additional contextualisation and analysis.

- Confidence levels vary across detected objects. A more nuanced visualization could be created calculating probability mass (as a function of polygon area x confidence and then aggregating these values within binned geographic areas.

- Without ground-truthing and validation, there is potential for false positives or negatives in the feature detection, which could impact conservation decision-making if not properly accounted for.

- The output should be considered a tool to guide further investigation and field work, rather than a standalone product for making definitive conservation decisions.

4.7. Optimal Confidence Thresholds

Whilst detecting objects at a low threshold was desirable to establish a baseline feature set, subsequent analysis demonstrates that further refinement above these thresholds is desirable for a high-quality feature output. Whilst these thresholds were not applied within our analysis for consistently, they have been applied to output feature datasets that will be shared for mapping purposes. The optimal thresholds above which features are retained for mapping for each class is as follows:

1. Vegetation: 65%
2. Buildings: 45%
3. Partitions: 55%
4. Fencelines: 22%
5. Roads: 60%
6. Water Bodies: 50%
7. Clouds: 85%

4.8. Post-Processing.

As standard all features under $5m^2$ were deleted to eliminate very small standalone trees and small fence and road misdetections

Buildings were regularised to form more coherent shapes around the target features. The result is not perfect, particularly with complex shapes, but is a noticeable improvement upon the original geometry.

Building Regularization

Before Regularisation | After Regularisation

A process also was derived to convert the linear feature polygons to line features.

- Polygons were generalised to capture their general direction and remove irregularities

- The Collapse Hydro polygons tool was run to capture centrelines
- Very short segments were deleted to improve output

Example Output of Linear Features

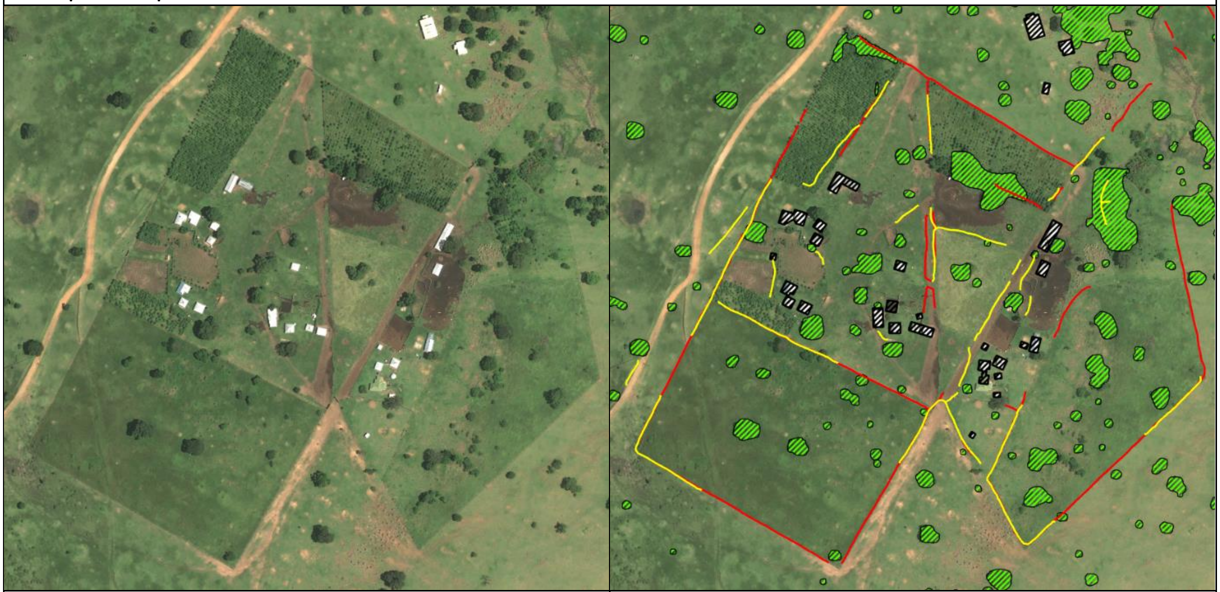

Raw Imagery	Detected Partitions (Yellow) and (Red)

4.9. Dissolved Features with Retained Confidence

Many output features are larger than the 224x224 pixel processing size however the model produced output polygons no larger than 224x224 pixels in size. This left many larger objects comprised of a number of constituent output polygons. In order to smooth output, a degree of overlap was maintained.

Comparison Between Raw Output and Dissolved Features

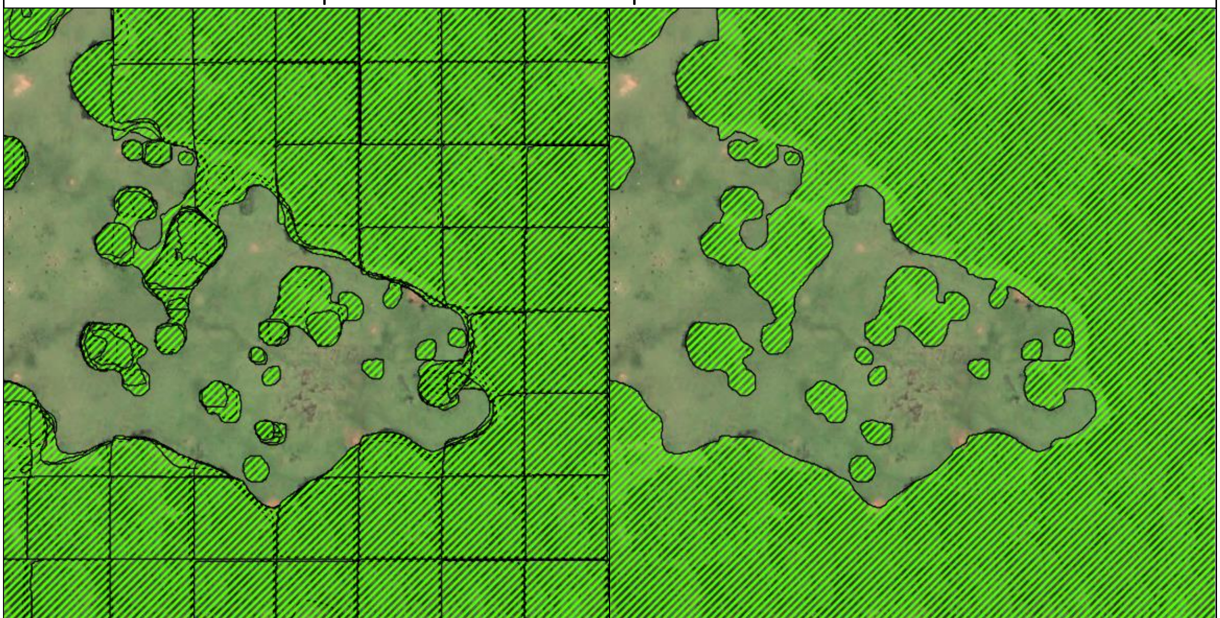

Raw Output Features	Dissolved Features

However, in order to achieve clean output polygons of the feature classes, objects were dissolved by class so overlapping features formed one feature class. A Summarize Within process was then applied across the dissolved features and the constituent features to establish a mean confidence of the constituent polygons. This was weighted by the proportion of the summarized constituent layer within the dissolved polygons and grouped by class.

4.10. Test Time Augmentation (TTA)

Test Time Augmentation is a technique used to improve the accuracy and robustness of model predictions during the inference phase.

- Image Transformations: The input image was subjected to multiple transformations flip and rotate transformations.
- Multiple Inferences: The model was run on each transformed version of the image.
- Aggregation: The results from all transformations were combined, taking the mean shape and confidence of the outputs.
- Thresholding: Outputs falling below the process confidence threshold (set at 50% for the PCA) were discarded.

In our Nyakweri Forest analysis, TTA was employed to refine the output polygons for several key classes, particularly in the Priority Conservation Area (PCA).

TTA resulted in smoother, more accurate polygon boundaries for detected features and helped eliminate many spurious detections. This had a dramatic effect on reducing object “sprawl” at lower confidence thresholds. TTA was particularly beneficial for classes with complex shapes or high variability, such as vegetation boundaries and linear features. While TTA improved overall quality, it sometimes resulted in the loss of smaller or less confidently detected features.

It is important to note TTA significantly increased processing time, making it impractical to apply to the entire study area within the project's scope. The performance impact of TTA increases exponentially at lower confidence threshold.

Due to these computational constraints, TTA was primarily applied to the 130 km2 Priority Conservation Area and run with a 50% confidence threshold, as opposed to the 25% confidence threshold employed across the entire 1018km2 study area without TTA.

The most effective use of TTA however, would have been to run inference at a low confidence threshold, such as 15% and then delete output polygons below class specific thresholds prior to running the dissolve and summary process.

Comparison of Output using TTA

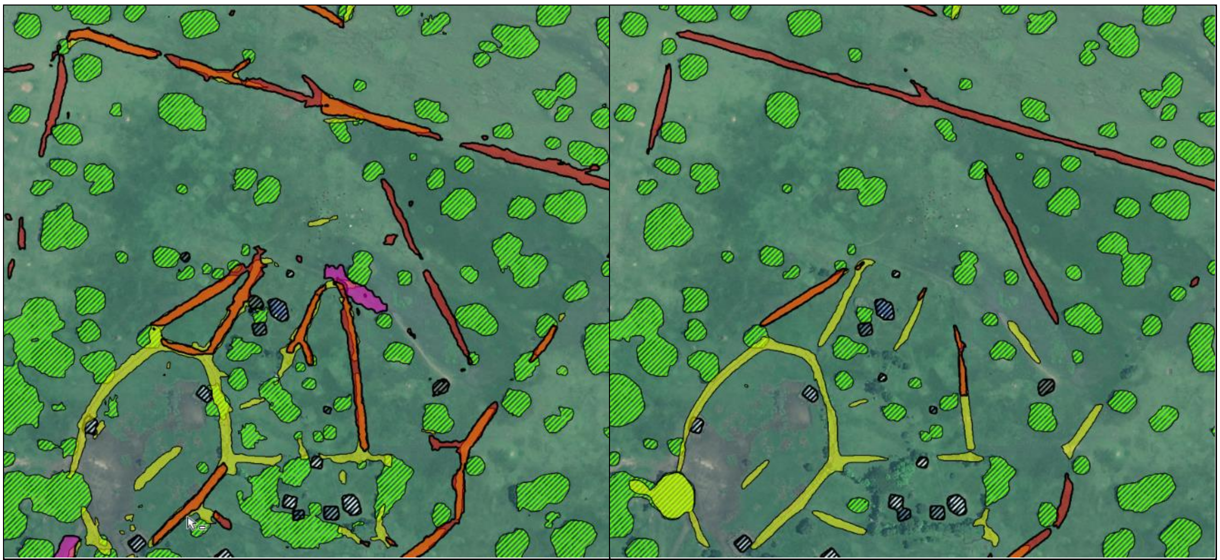

Without TTA | With TTA

Example Output Using TTA

Raw Imagery | Detection of Wire Fence using TTA

Leveraging AI and Remote Sensing for Conservation

Part 2: Analysis of Data

Ruari Bradburn, Chief Technology Officer, Langland Conservation

1. Introduction

The main purpose of the analysis was to identify indicators of habitat degradation and their implications on Giant Pangolin conservation.

The two key areas of focus were forest loss and fence electrification. Recent imagery shows that both within and on the peripheries of the priority conservation area an alarming amount of forest is still being lost in a relatively short time, posing serious questions about the long-term ecological outlook of this area.

Understanding what areas of forest remain, and how fences can be targeted to improve connectivity between remaining areas of woodland, forms the central basis of this analysis.

In addition to protection and de-electrification, this study suggests priority areas to replant trees to counter habitat fragmentation. This is usually targeted towards reinforcing thin veins of connecting forest or bridging gaps between major disconnected forest blocks across the shortest or most undeveloped path.

2. Two-Stage Analysis

Analysis was conducted in two stages. First a general analysis of the wider area was conducted using the data resulting from running the model at 25% confidence without using Test Time Augmentation (TTA). Then a detailed analysis of the Priority Conservation Area was conducting using data resulting from the model at 50% confidence with TTA enabled. Whilst some less confidently predicted features may be absent from this iteration, the data is provided with a greater degree of confidence and a significantly lower rate of false positives

In the PCA area a calculation was done to weight scores by confidence but this caused a negligible variation in the output visualisations so was omitted here.

3. Remaining Forest Cover

The PCA represents the largest concentration of remaining intact forest cover within the wider area. There remains substantial forest to the immediate north of the PCA, however signs of landscape change are widespread, such as felled trees or scars on the soil from charcoal production. In the Priority Conservation Area 21.3km2 of tree cover remained in March 2024 out of the area total of 49.2km, representing 43.2% of the total area. Of this:

- 18.4km2 was comprised of areas of tree cover 1000m2 or more in size, or 37.3% of the total area.

- 12.1km2 was comprised of areas of tree cover of 0.5km2 or more in size, or 24.6% of the total area.
- 8.3km2 was comprised of areas of tree cover 1km2 or more in size, or 16.9% of the total area.

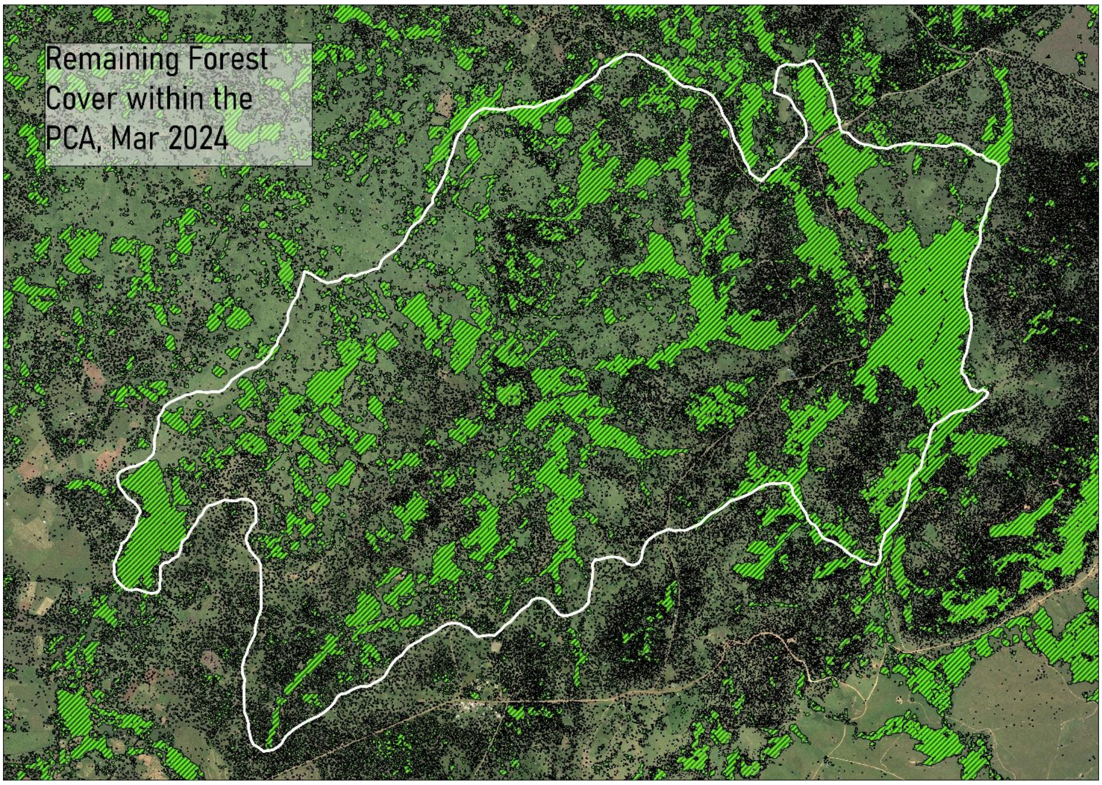

4. Forest Loss

Forest loss stands as the most pressing conservation challenge in the Nyakweri Forest landscape, directly impacting the Giant Pangolin population and the overall ecosystem health. The Nyakweri Forest landscape has been eviscerated by deforestation over the last two decades, losing more than 60% of its tree cover since 2001. Most of this change has occurred since 2012 and follows a steady pattern of loss moving from the north of the area to the south.

Once a vibrant network of large contiguous forest, the area has been reduced to a mosaic of thin veins of tree cover and disconnected forest blocks. In addition to widespread conversion of large blocks to agricultural land, there is also a consistent pattern of trees being felled in small numbers across the area, "nibbling" away at remaining forest blocks. The sum of this activity across landscape is the consistent and continued erosion of its viability as a pangolin habitat.

Our analysis of forest cover change using high-resolution satellite imagery and AI-driven forest detection provides the most accurate and up-to-date available data on remaining forest cover in this area.

It also confirms that deforestation and habitat fragmentation is continuing to occur at an unsustainable rate, including within the PCA.

Google Earth Imagery Over 10 Years		
		
2012	2013	2014
		
2015	2016	2017
		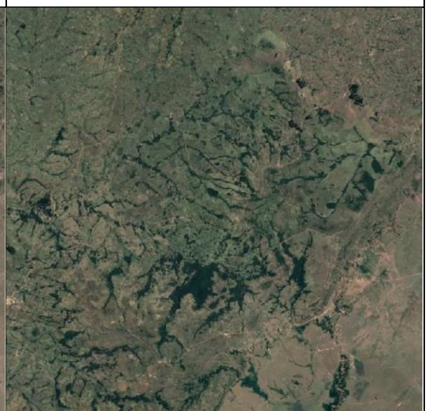
2018	2019	2020
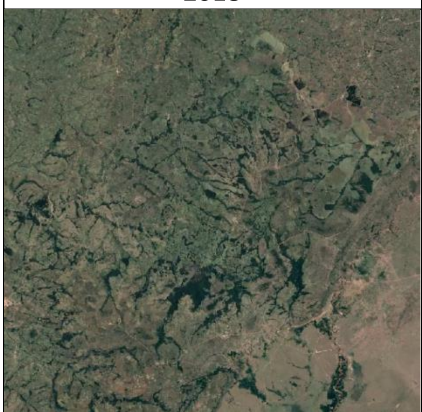	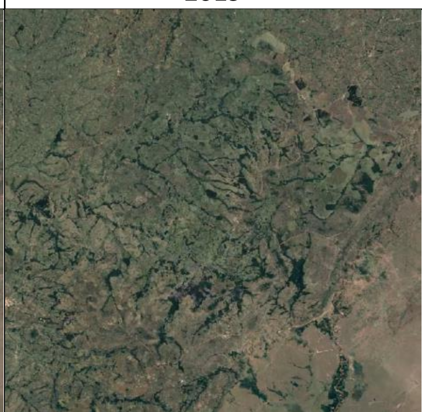	
2021	2022	

Examples of Recent Forest Loss within the Priority Conservation Area (North-West Corner)

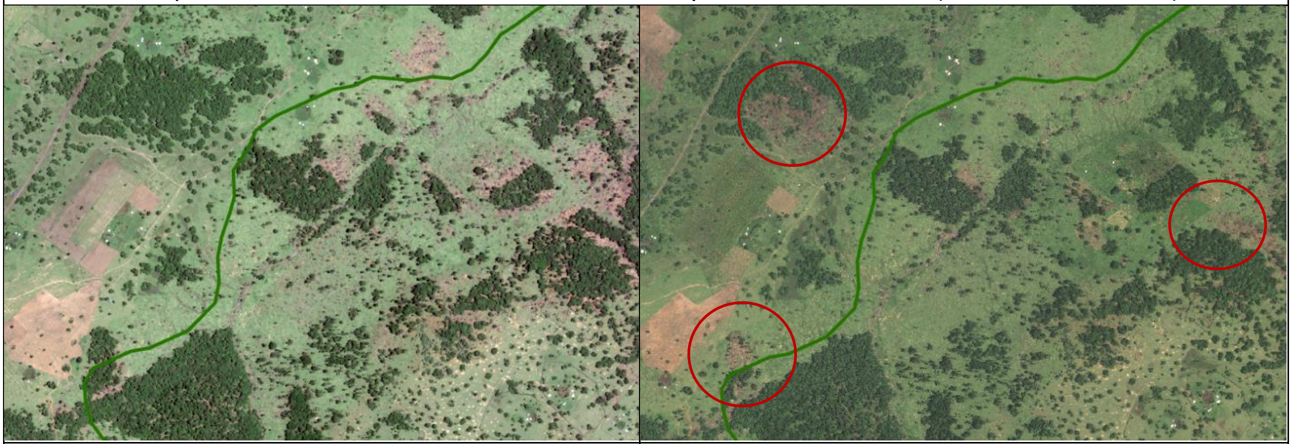

Oct 2023 | Mar 2024

Close-Up Example Recent Forest Loss within the Priority Conservation Area

Oct 2023 | Mar 2024

Our imagery covers different areas of the Nyakweri Forest at different times, making a comprehensive analysis at high resolution impossible. However, between the imagery days there are key areas of overlap, including 29.8km2 of the Priority Conservation Area, allowing an effective change analysis in 72% of the PCA.

Seasonal variability between the two dates complicates the accurate measurement of loss due to the rainy season meaning vegetation is much lush and full during the spring. The total area of detected tree cover was virtually identical across both dates.

However, when small trees were removed from the data and blocks at least 100mx100m in size were counted, it proved an extremely useful step in highlighting specific areas of loss for further manual confirmation.

A difference map was created with areas of loss between the October imagery marked in red.

Manually Confirmed Sites of Forest Block Loss between October 2023 and Mar 2024

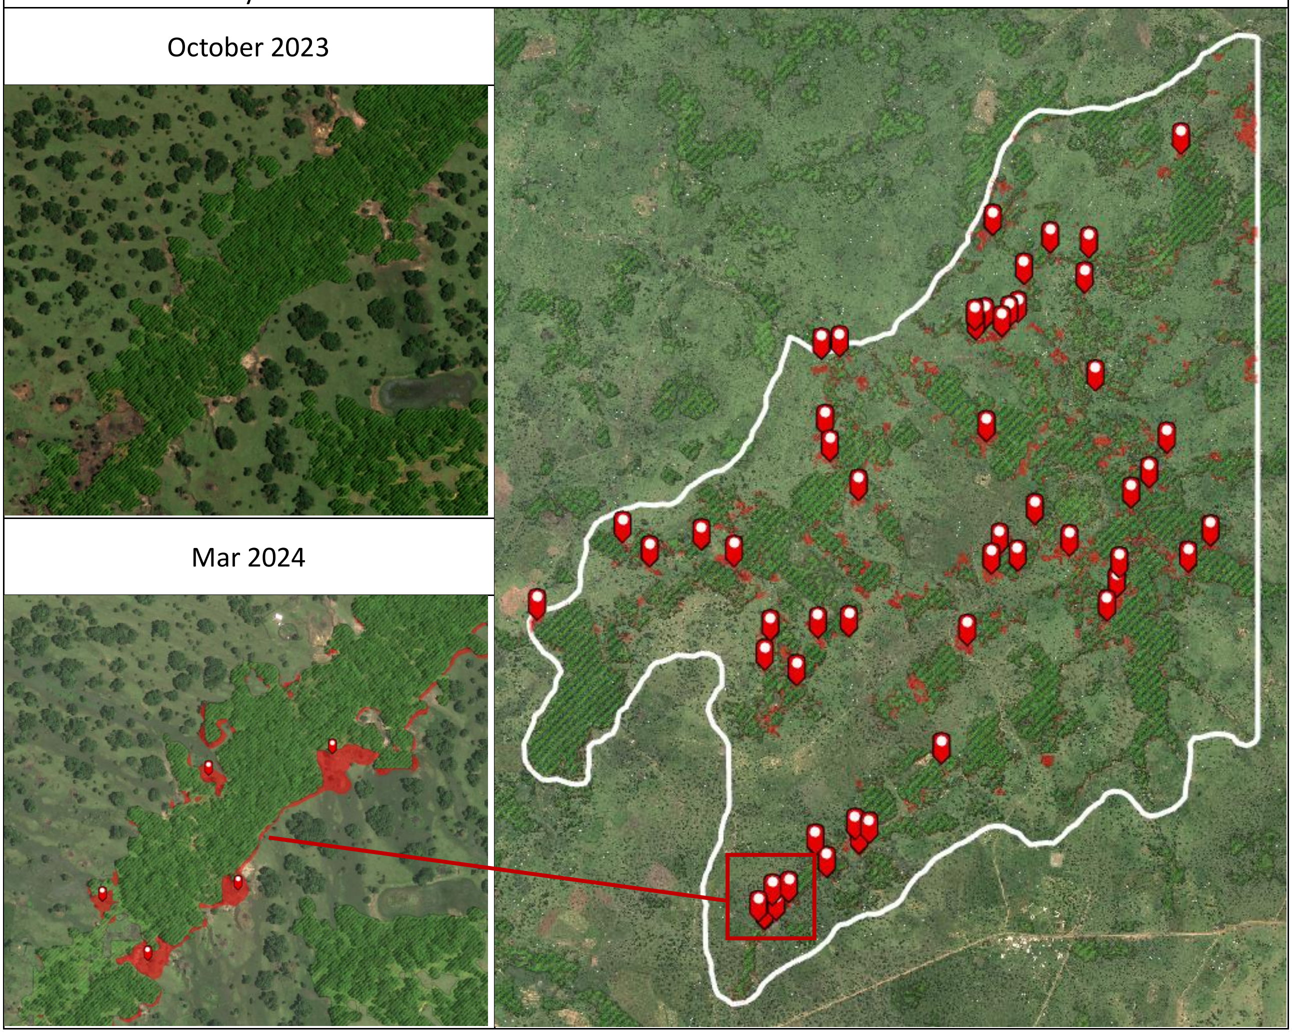

Most large area loss detections usually correlated to ongoing felling of the remaining forest blocks and rapidly improved the rate and reliability with which they could be detected. After manually verifying felling activity

The prevalence of scars left behind by other previous felling activity underline a persistent and historic pattern of loss.

Using the AI-driven object detection model, we identified buildings within the PCA. We then analyzed their proximity to areas of confirmed recent deforestation. The process involved:

- Identifying clusters of buildings (potential settlements or homesteads)
- Creating buffer zones around recent deforestation sites (100m radius)
- Intersecting these buffer zones with the identified building clusters

This analysis revealed several key findings 70% of the recent deforestation sites were within 200m of buildings

- There remains very little space within the PCA where there remain large gaps between buildings – 87% of the PCA falls within 300m of a building. The 13% of the PCA that falls out of 300m proximity of a building accounts for 20% of the remaining tree cover and 22% of the remaining tree cover in blocks larger than 100x100m.
- In the areas where buildings were further away, there were visible signs of historic deforestation between the buildings and the deforestation sites, indicating a continued pattern of erosion.

Due to the fragmented nature of the habitat and the fact communities are interspersed throughout the priority area, a community-centric solution, as has been adopted by The Pangolin Project, must play a central role within wider conservation efforts. Working alongside communities, fostering positive attitudes towards conservation, and incentivising the safekeeping of the remaining forest is the principal way that further damage can be mitigated.

This analysis has identified which communities sit in close proximity to active sites of deforestation where trees were felled between October 2023 and March 2024.

Clusters of housholds in proximity to recent deforestation

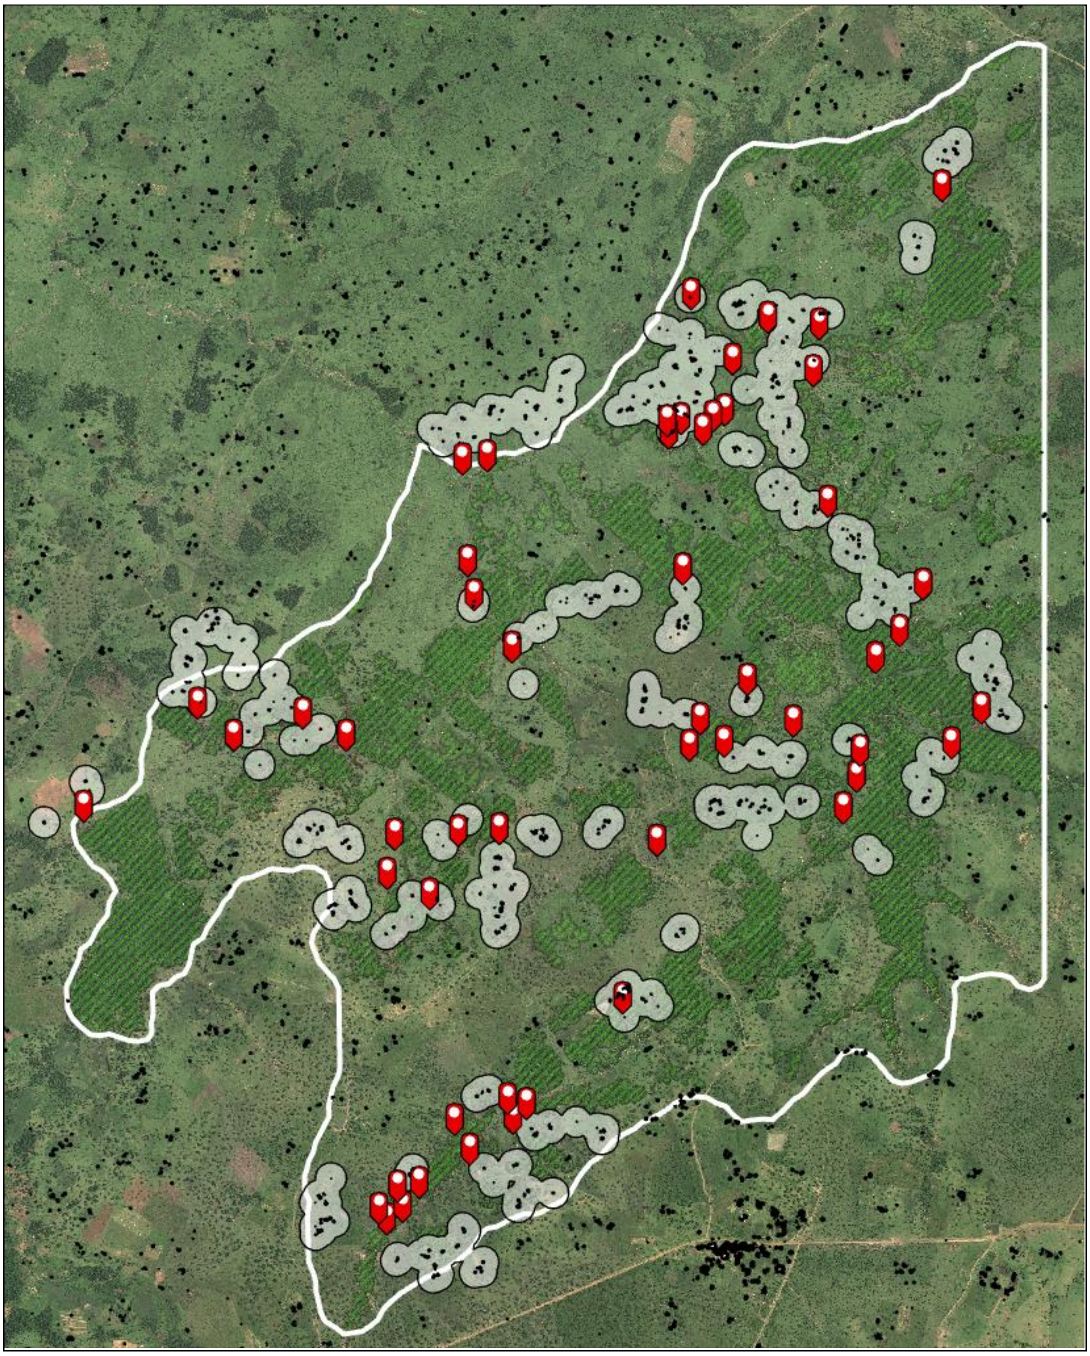

5. Connectivity Analysis

The connectivity analysis aims to identify and prioritize areas where limited potential reforestation efforts to improve habitat connectivity for wildlife, particularly the Giant Pangolin.

The process was conducted as follows:

Identify significant forest patches:

- Tree cover areas larger than 100x100m² (10,000m²) were identified as significant forest patches.
- A 50m buffer was drawn around these significant forest patches.

Identify potential connection areas:

- The intersecting areas of these buffers were kept, as they represent potential corridors between forest patches.
- Areas where these intersections crossed roads were highlighted in red, indicating potential wildlife crossing points, as well as were reforestation efforts are likely unfeasible due the importance of the road to local communities.

Account for human settlements:

- A 50m buffer was created around households (buildings).
- This household buffer was then erased from the tree cover intersect areas.

Identify priority areas for community engagement:

- The resulting areas represent potential reforestation zones that could connect existing forest patches while minimizing conflict with human settlements.
- These areas are highlighted as priority locations for engaging local communities in reforestation efforts.

This analysis could help conservation efforts by:

- Identifying gaps in forest cover that, if filled, could significantly improve habitat connectivity. The biggest opportunities for reforestation likely fall within the centre and west of the PCA. Whilst the most intact forest block falls to the east, a major road likely prohibits reconnecting it to the other areas.
- Highlighting areas where wildlife might be at risk when crossing between forest patches (e.g., road crossings)..
- Providing a data-driven approach to prioritize areas for reforestation and community engagement efforts.

Connectivity Analysis	
Tree cover areas above 100x100m2 identified	50m Buffer Drawn

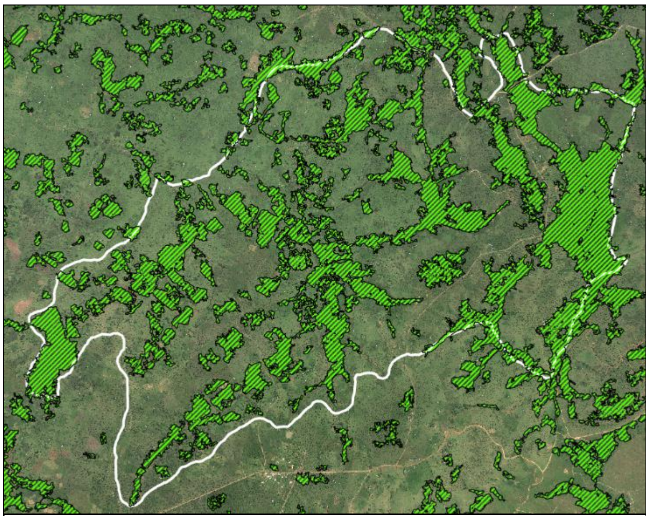

Intersecting Areas Kept (red where over roads)

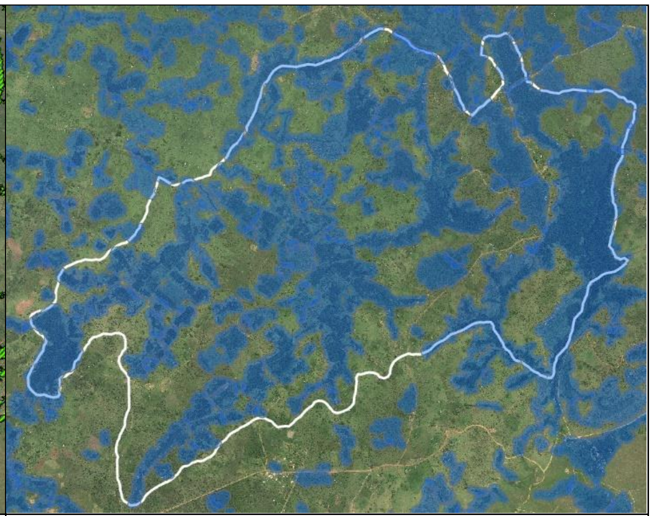

50m Buffer over Households

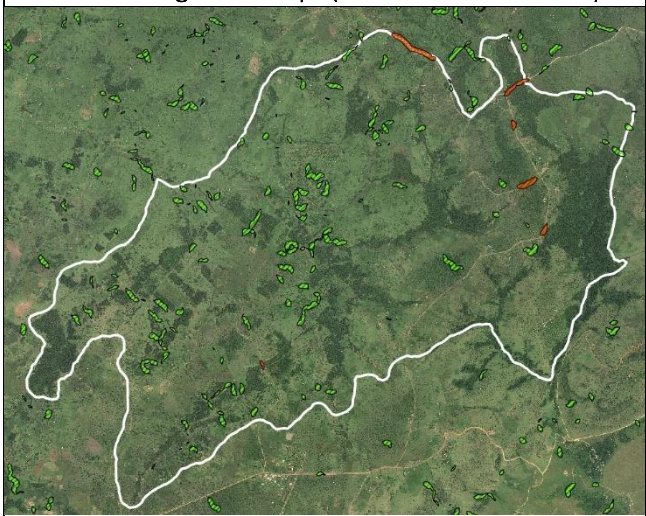

Household Buffer Erased from Tree Cover Intersect

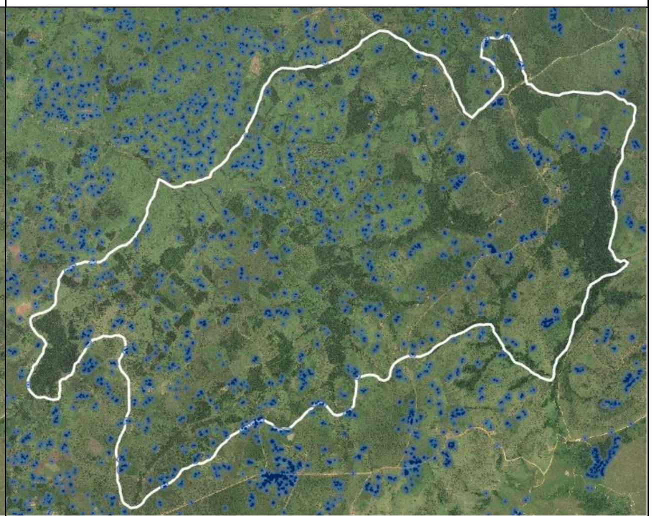

Communities to engage regarding reforestation

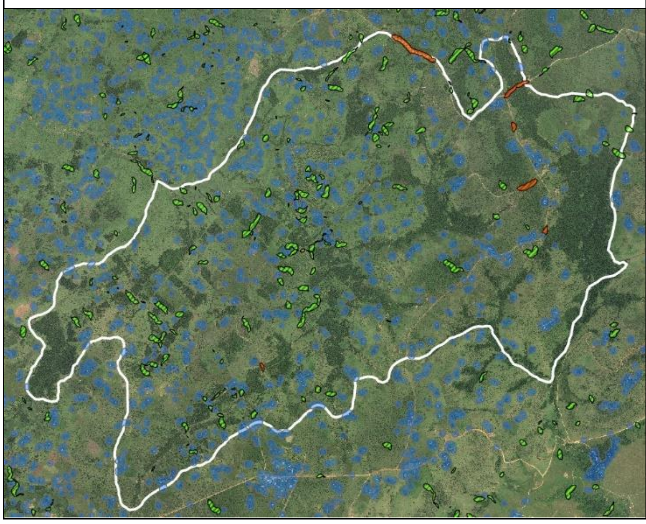

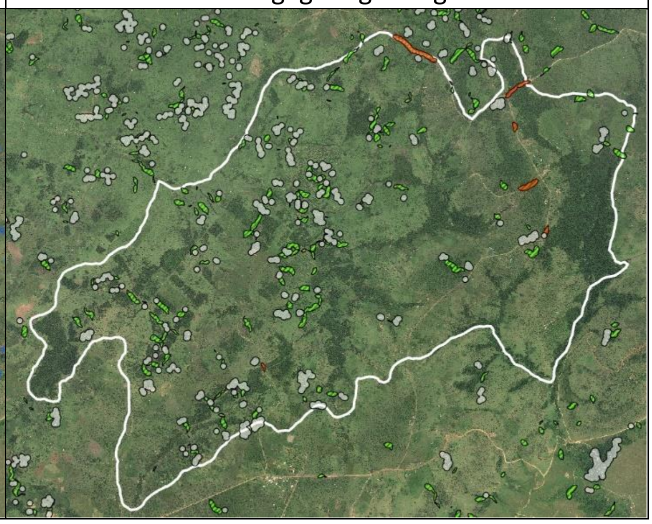

6. Analysis of Fencing

The first iterations of the model, as well as the final 7-class model had poor performance for fences, necessitating a bespoke fence-specific model with a greatly expanded training dataset for this class.

Due to the difficulty accurately mapping fences using the model, the model was used primarily for general fence detection. Once the fences were detected, they would be hand-drawn within the PCA and surrounding area to create a comprehensive map of fence lines. This had the added advantage of generating high-quality training data that will be used to train successive versions of the model.

Comprehensive Fenceline Map (Manually Drawn, Based off Detections)

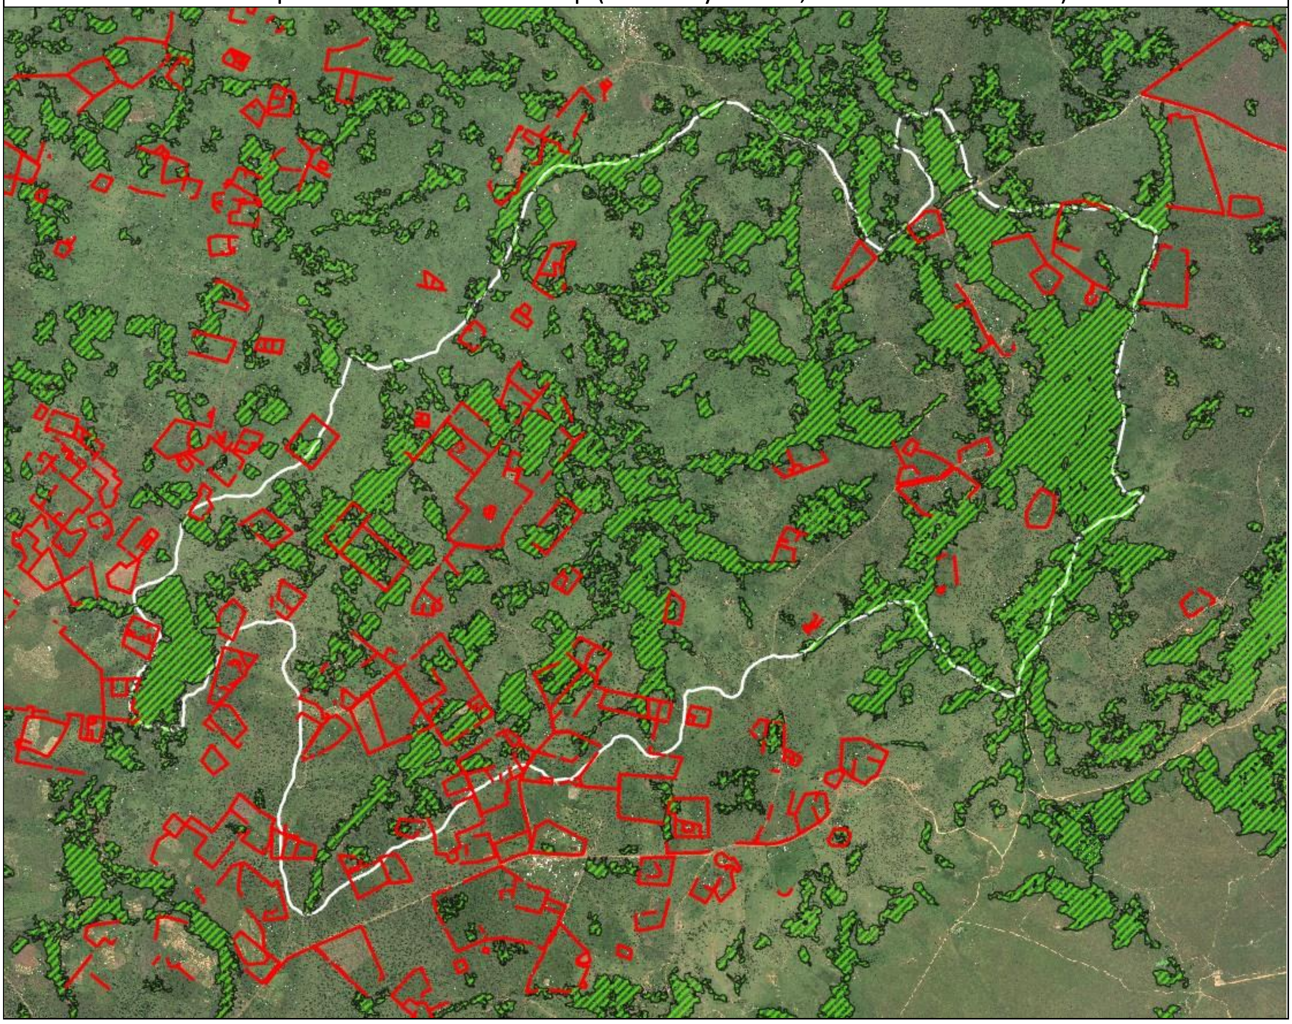

Whilst this does not indicate whether the fences are electrified, it provides a useful starting point to identify priority areas for ground verification and community engagement.

In order to achieve this fence lines within 100m of remaining forest were filtered. From here nearby homesteads near those fence networks were grouped into communities and highlighted for engagement.

Most such communities lie within the western area of the PCA, which is has the highest concentration of fences.

Fenceline Analysis	
Fences within 100m of forest filtered	Fencing Areas buffered to 200m
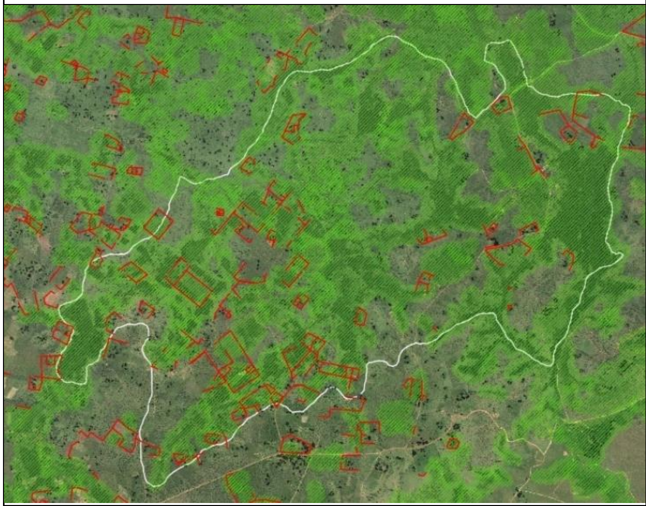	
Communities within 200m of a fence identified	Example of identified communities / homesteads
	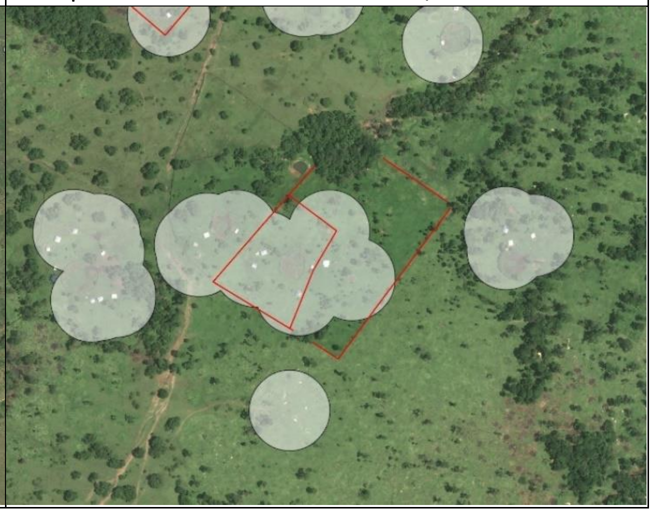

7. Annex

This annex provides a series of visualisations for the output data. It covers a series of visualisations for both the wider analysis performed at 25% confidence without TTA, and the detailed analysis of the PCA conducted at 50% confidence with TTA enabled.

Note that whilst the area coverage for fences and partitions is a very good indicator of relative density, due to the oversized bounding boxes, the numbers of area coverage to not correlate to an accurate real-world metric.

Full resolution images and raw data are available on request.

7.1.1. Wide Area Analysis - Tree Coverage and Feature Count

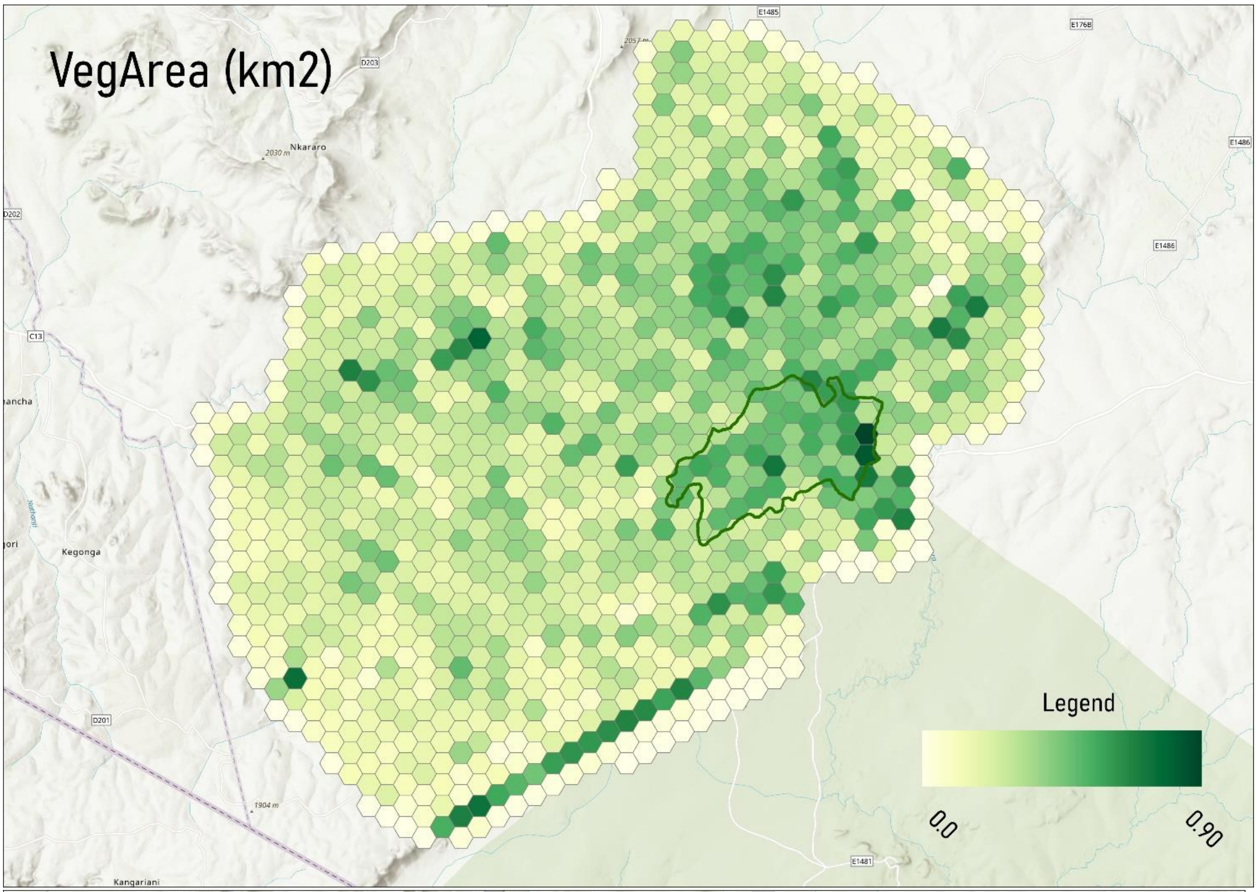

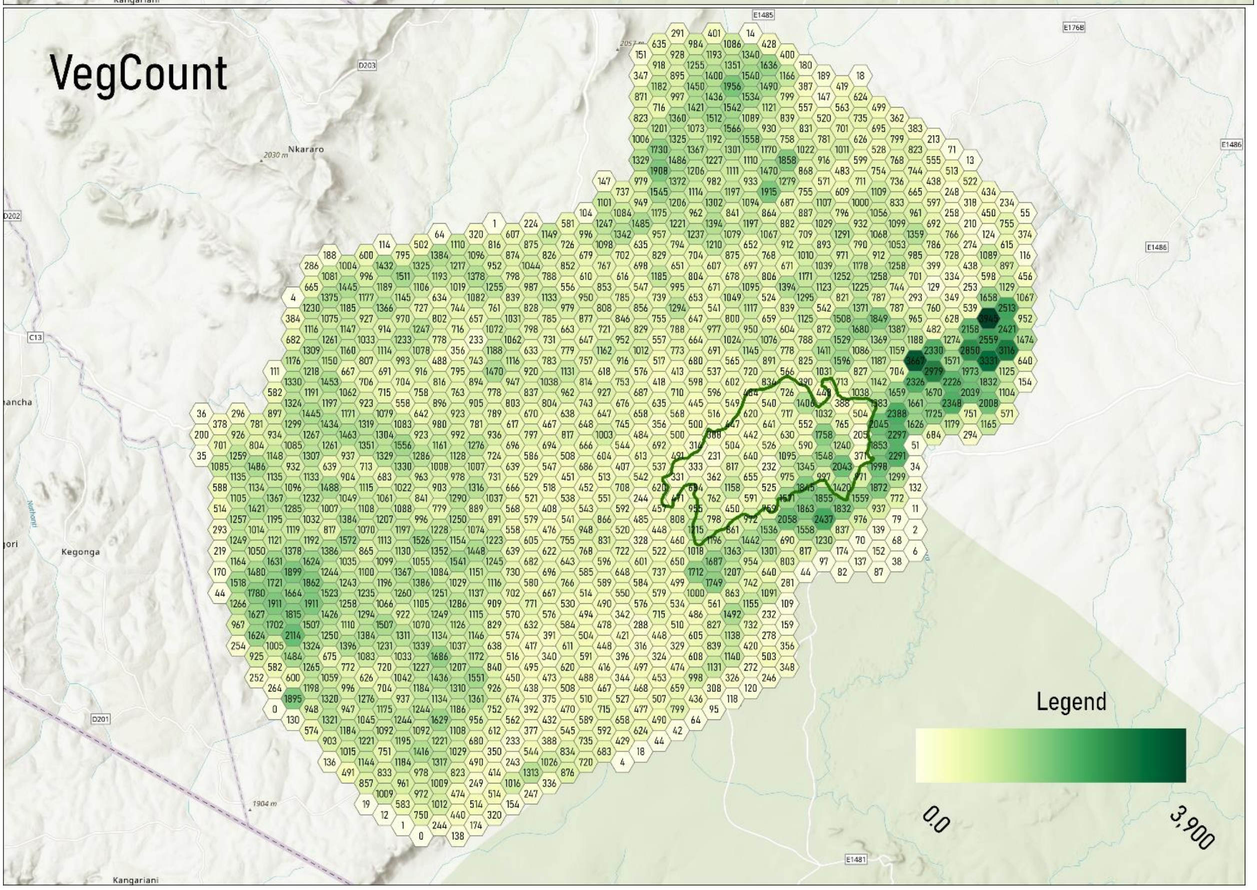

7.1.2. Wide Area Analysis - Building Coverage and Feature Count

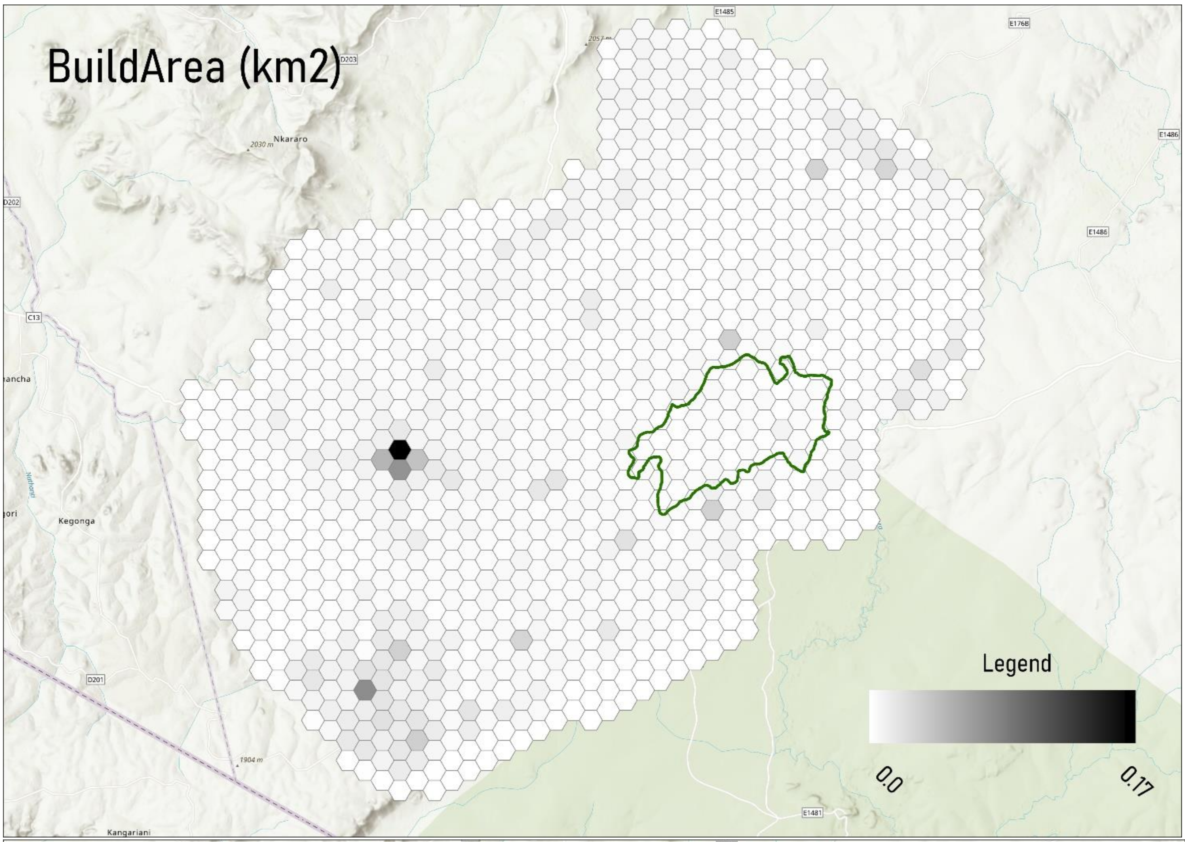

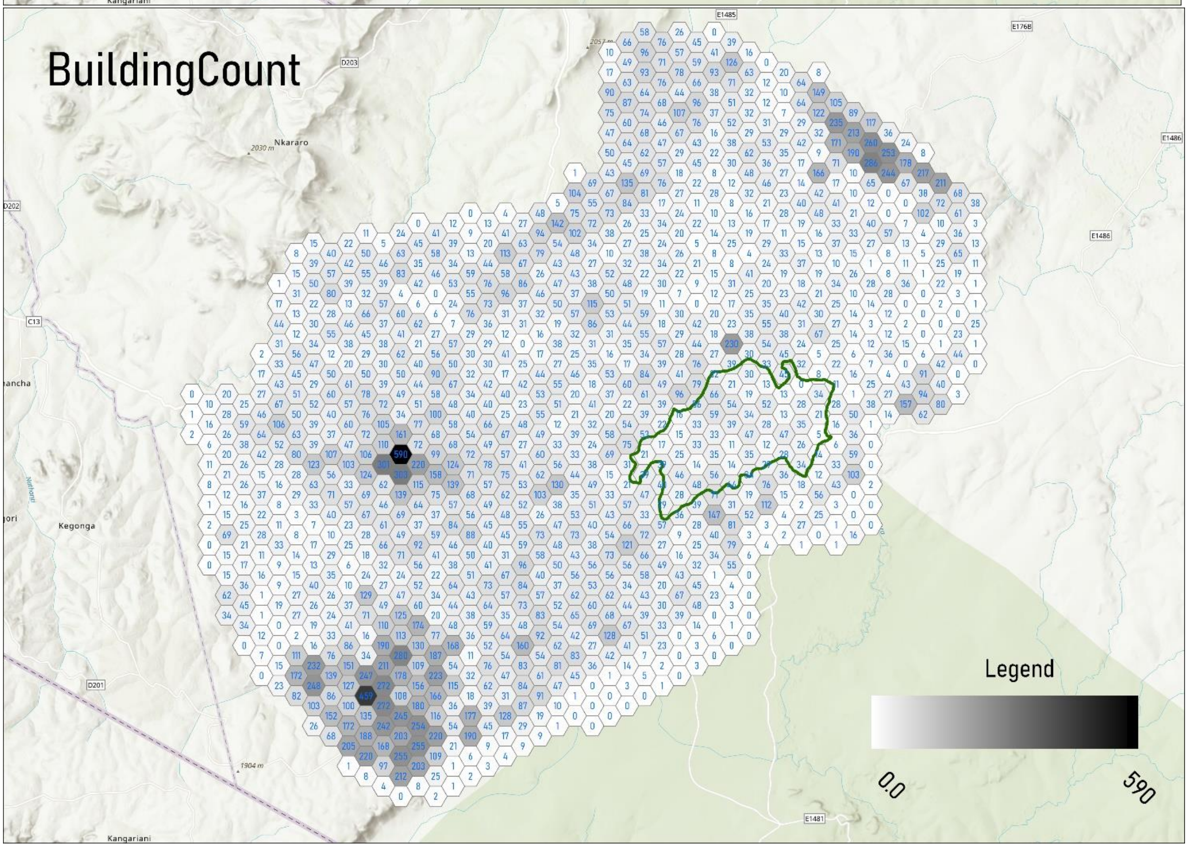

7.1.3. Wide Area Analysis – Partition Coverage and Feature Count

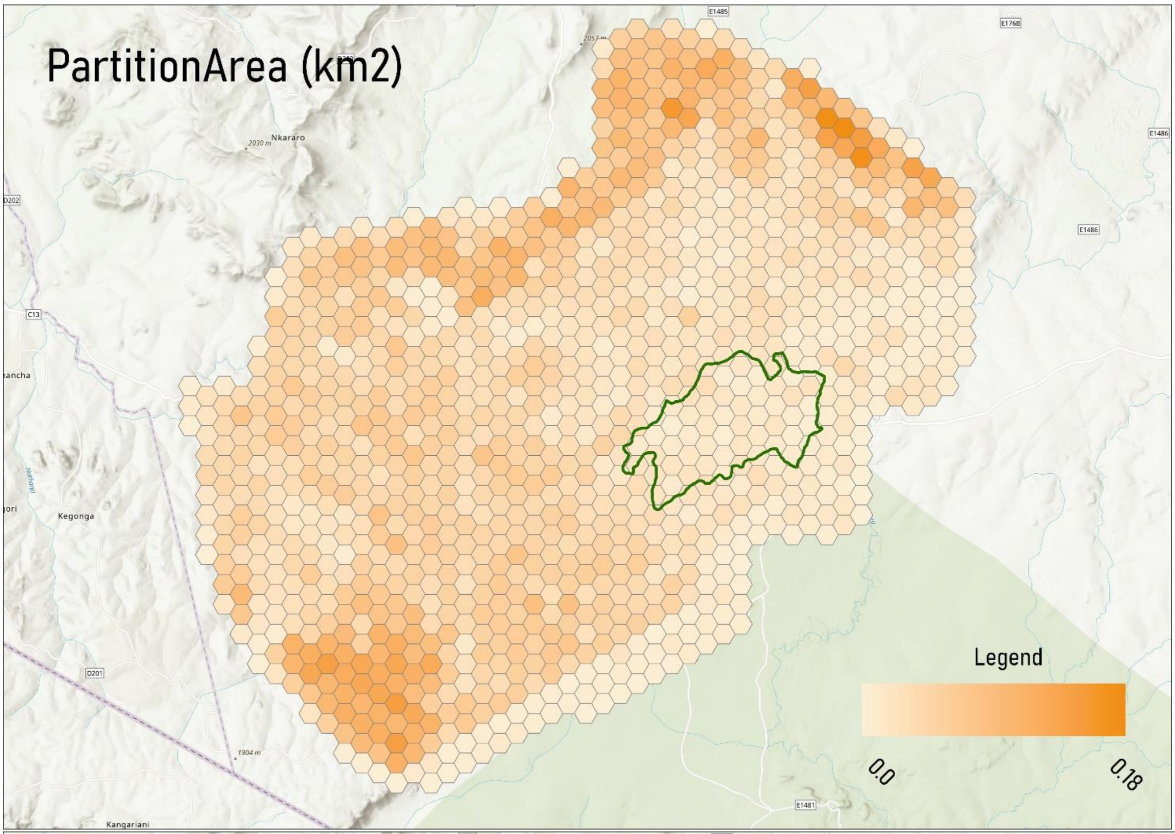

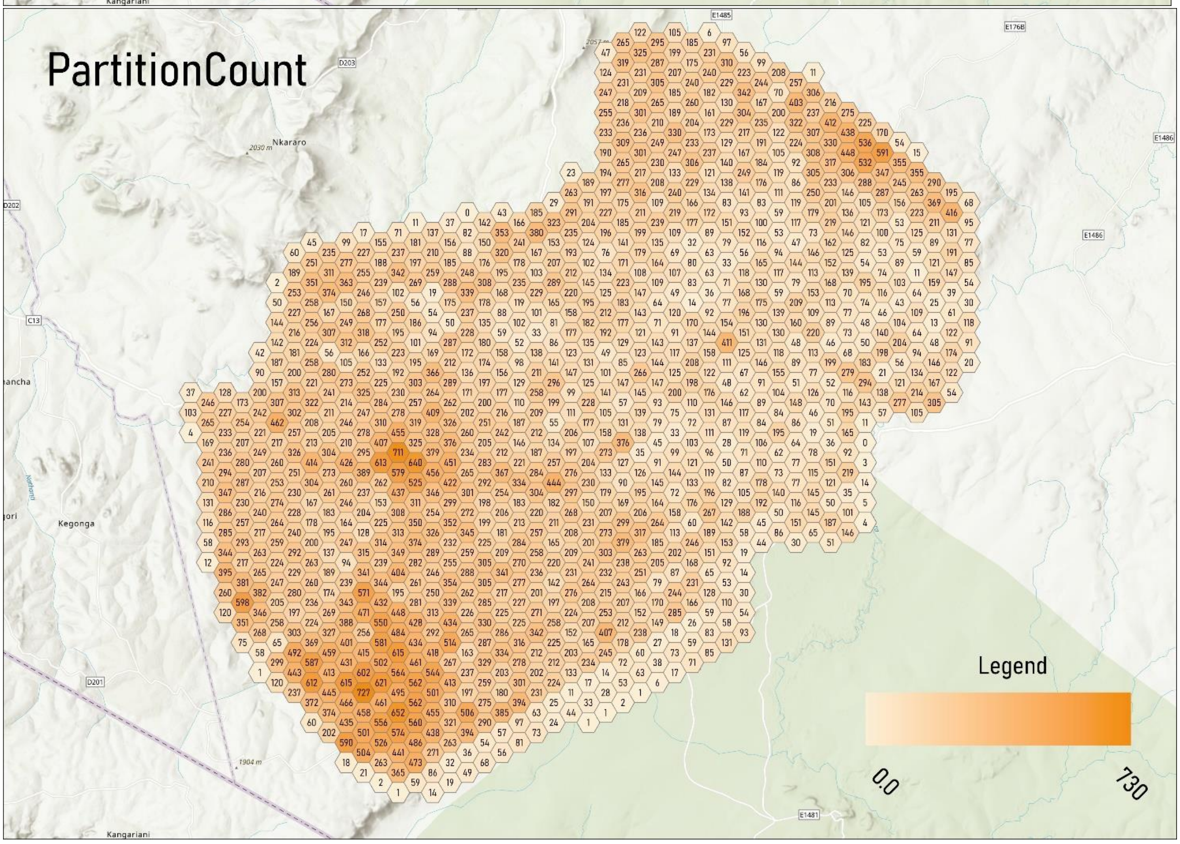

7.1.4. Wide Area Analysis – Fence Coverage and Feature Count

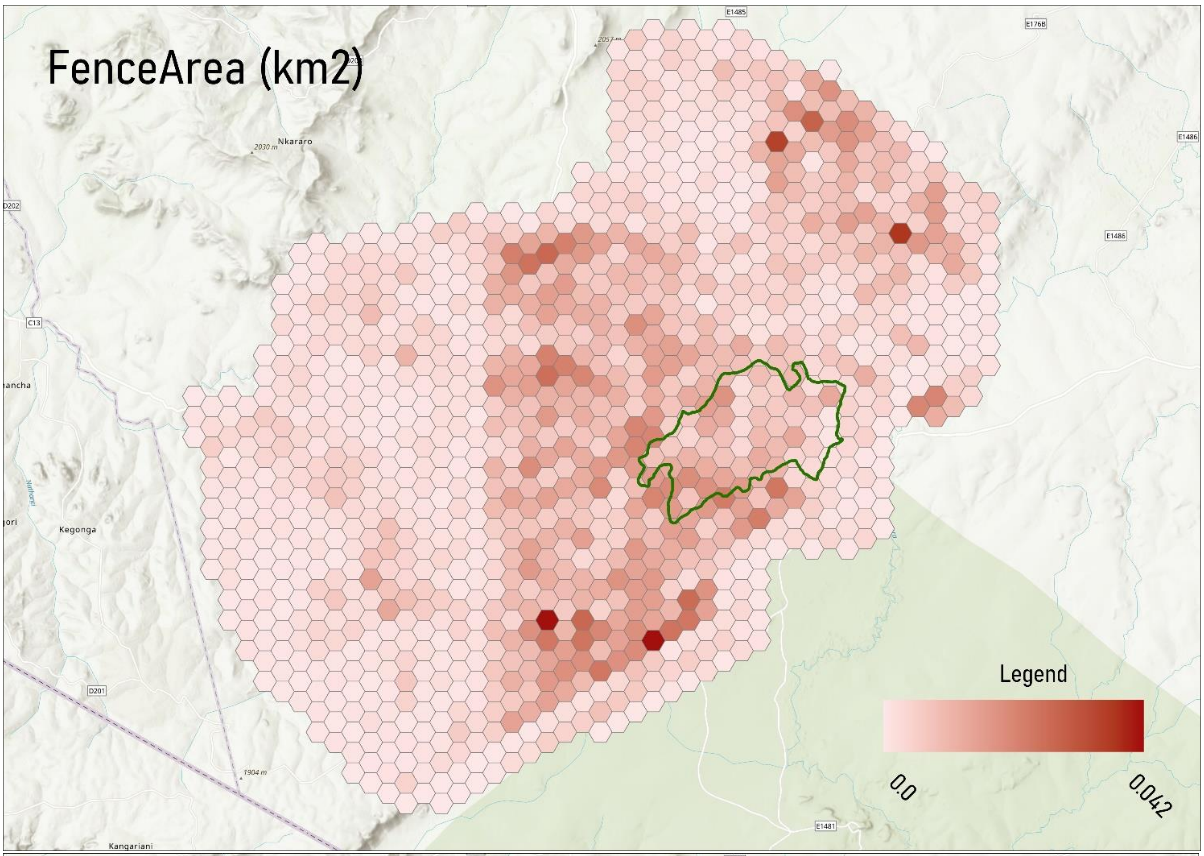

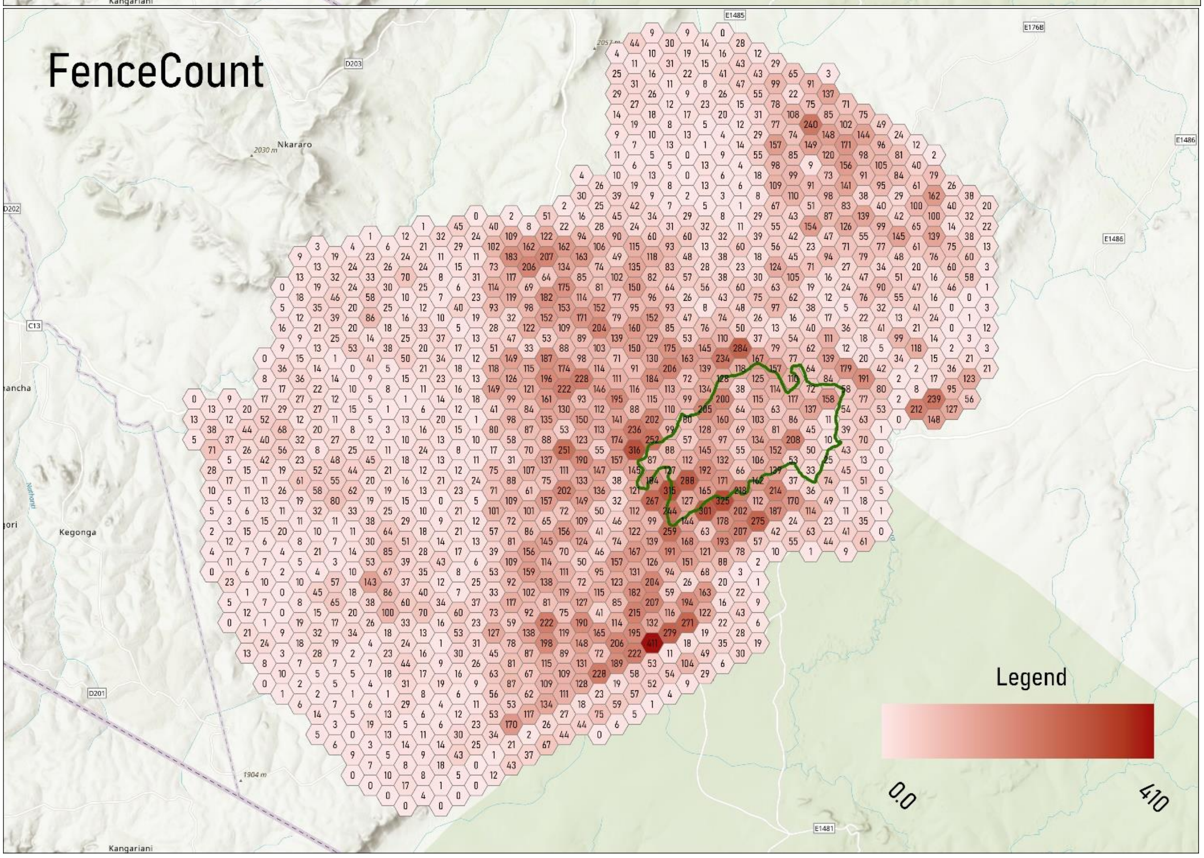

7.1.5. Wide Area Analysis – Road Coverage and Feature Count

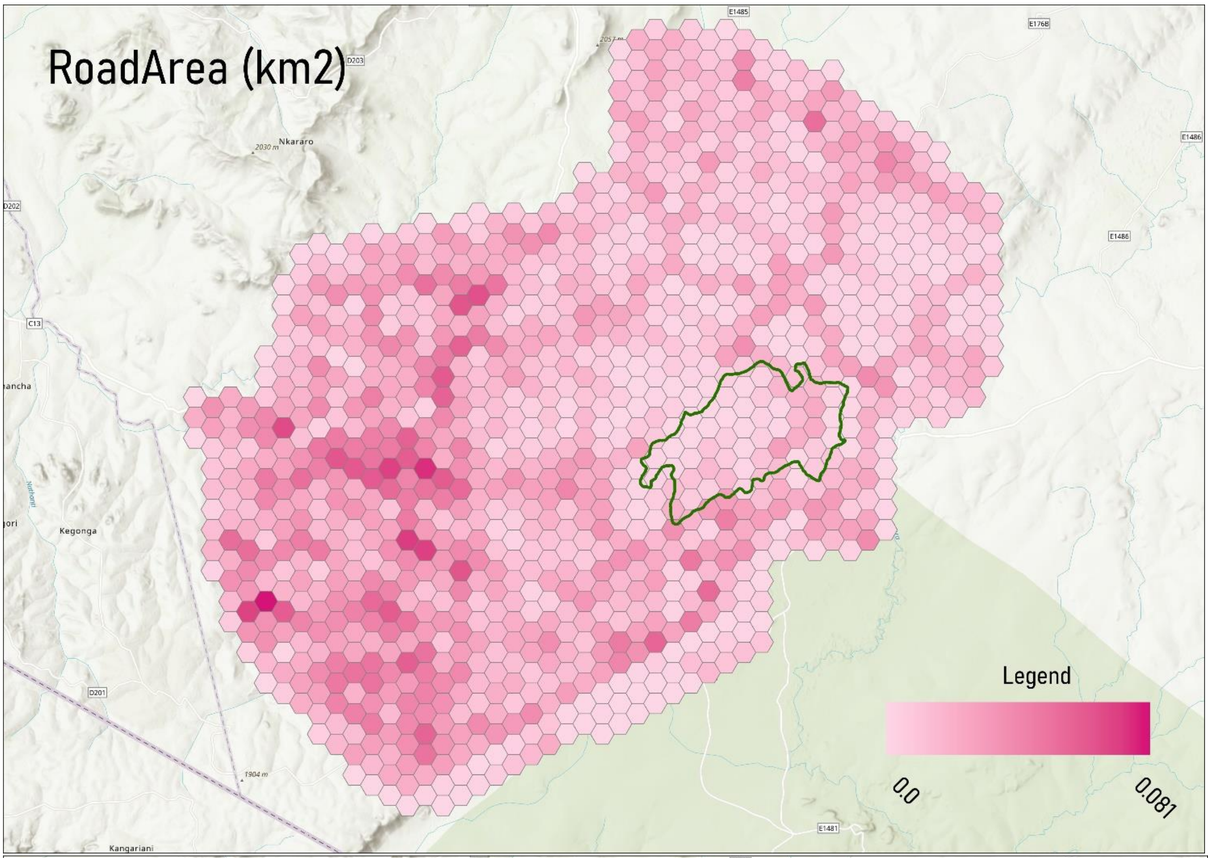

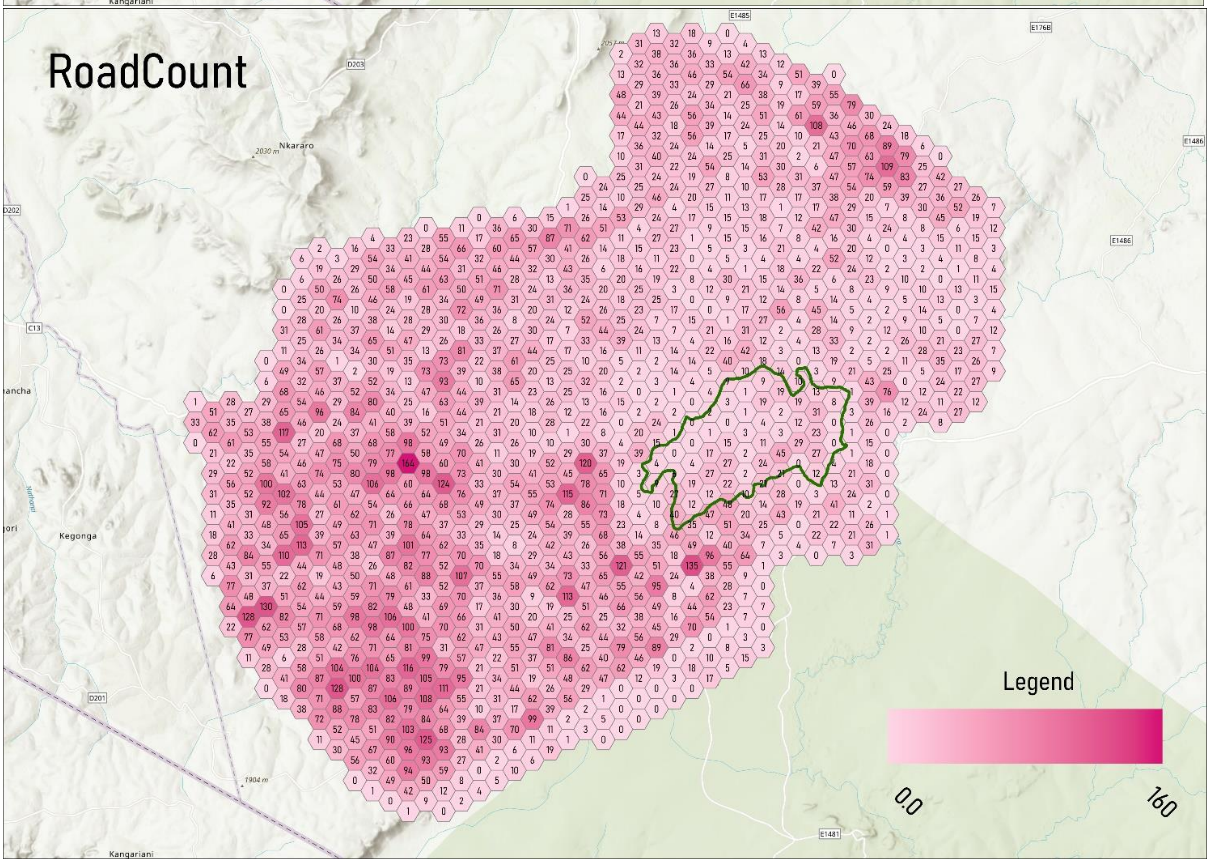

7.1.6. Wide Area Analysis – Road Coverage and Feature Count

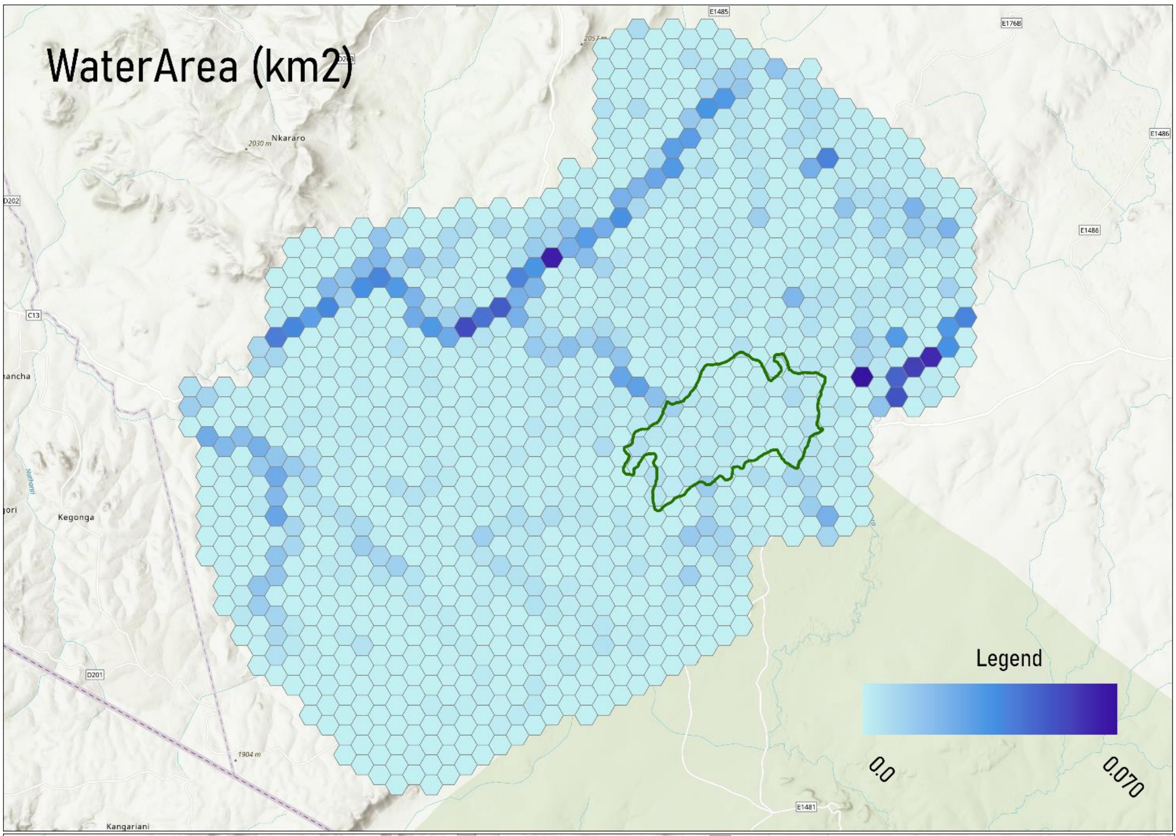

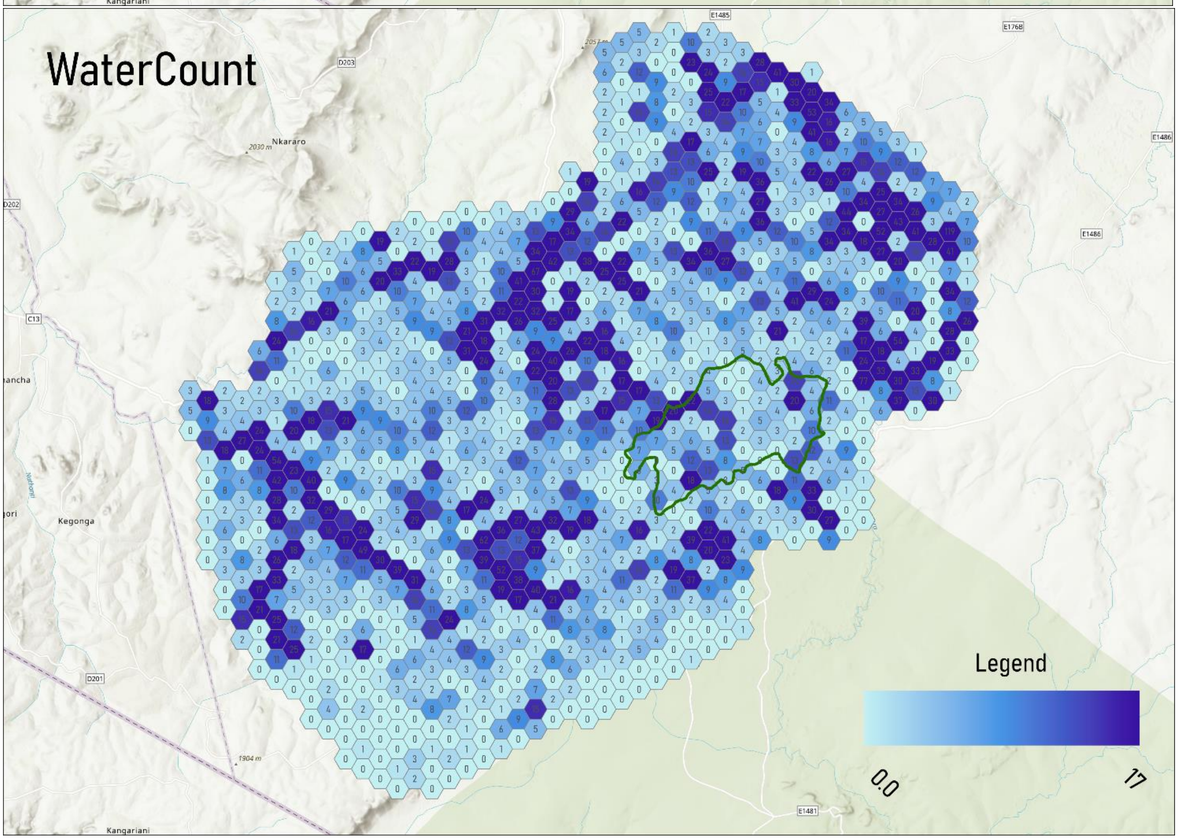

7.2.1. Priority Conservation Area Analysis – Tree Coverage and Feature Count

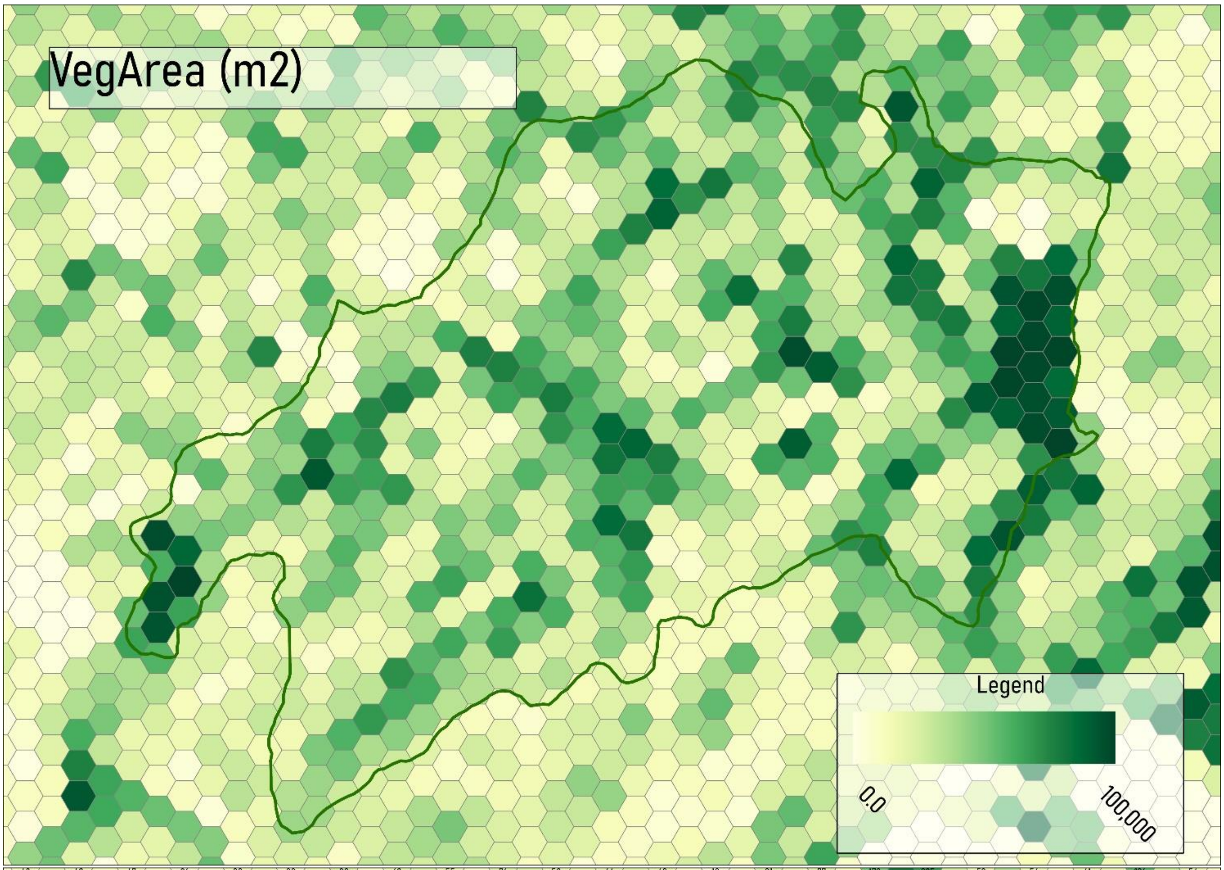

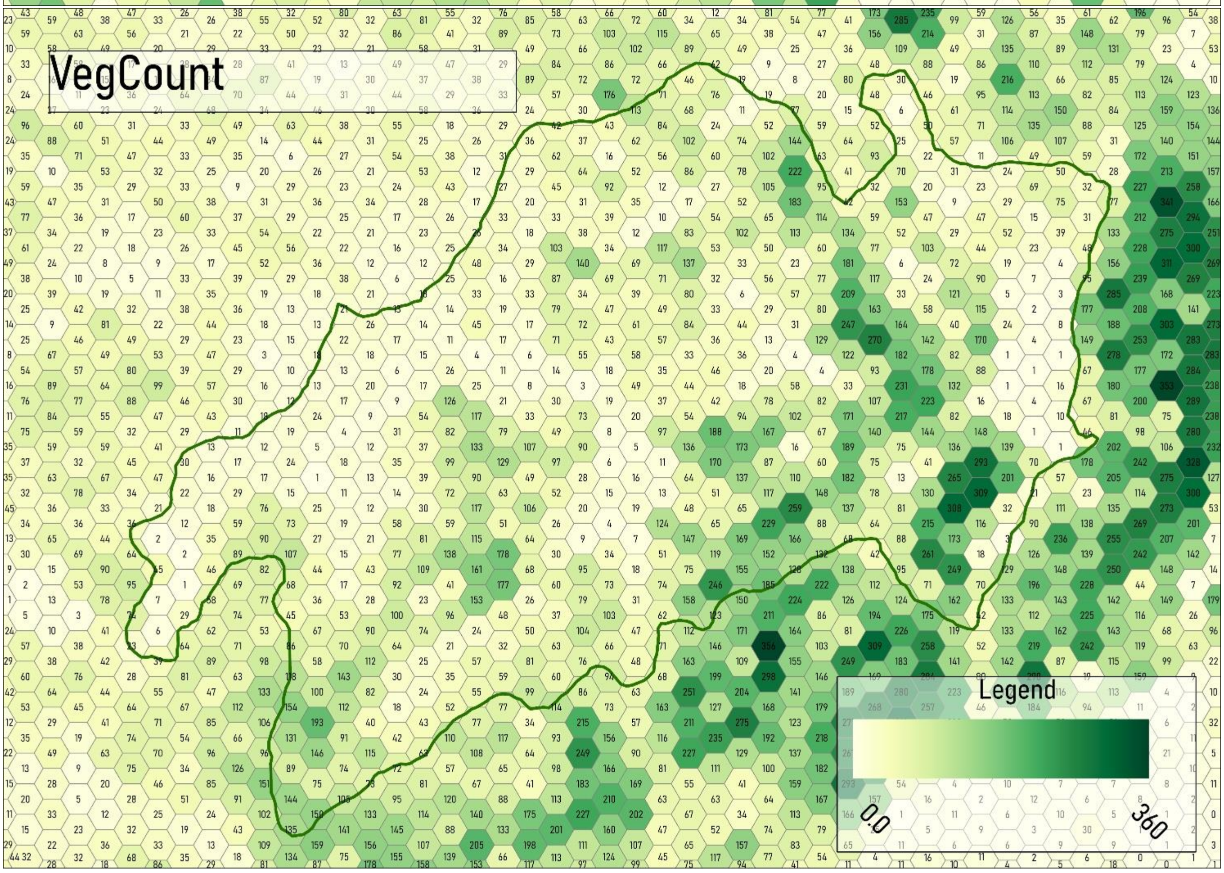

7.2.2. Priority Conservation Area Analysis – Building Coverage and Feature Count

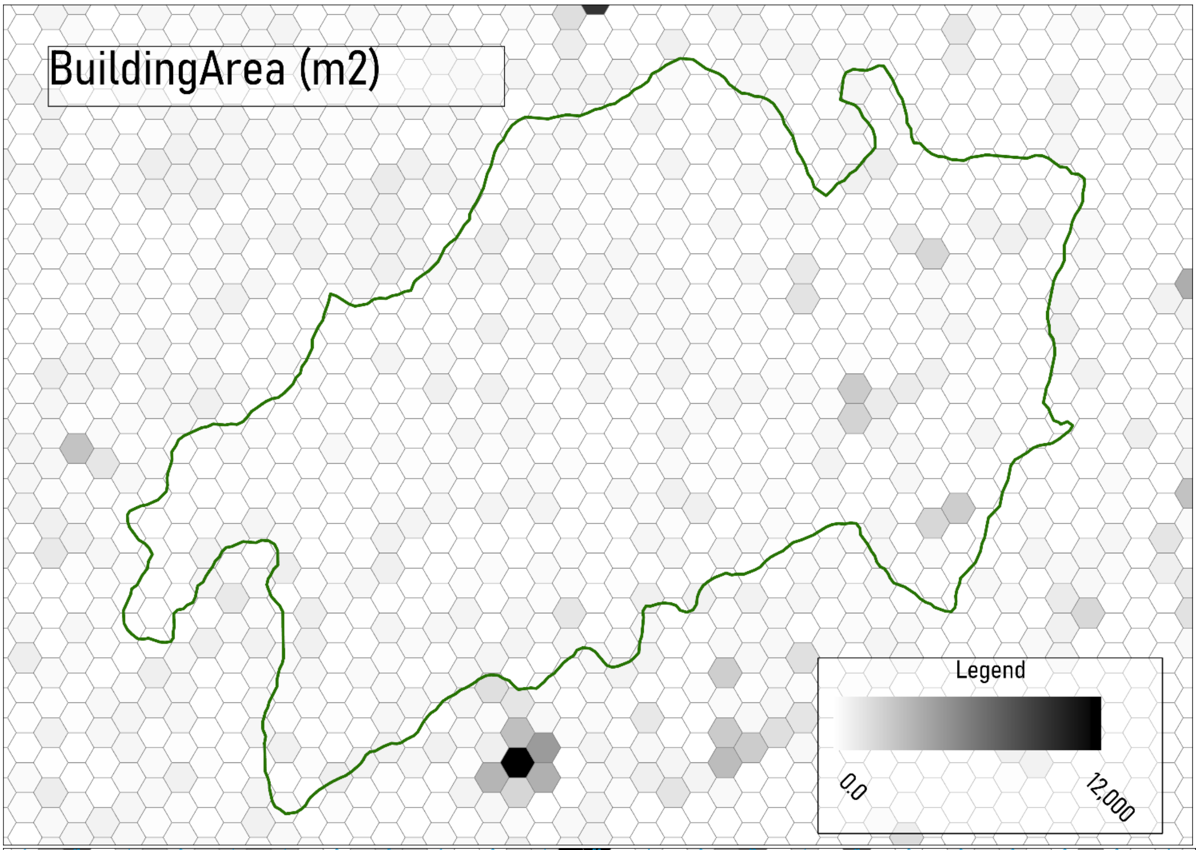

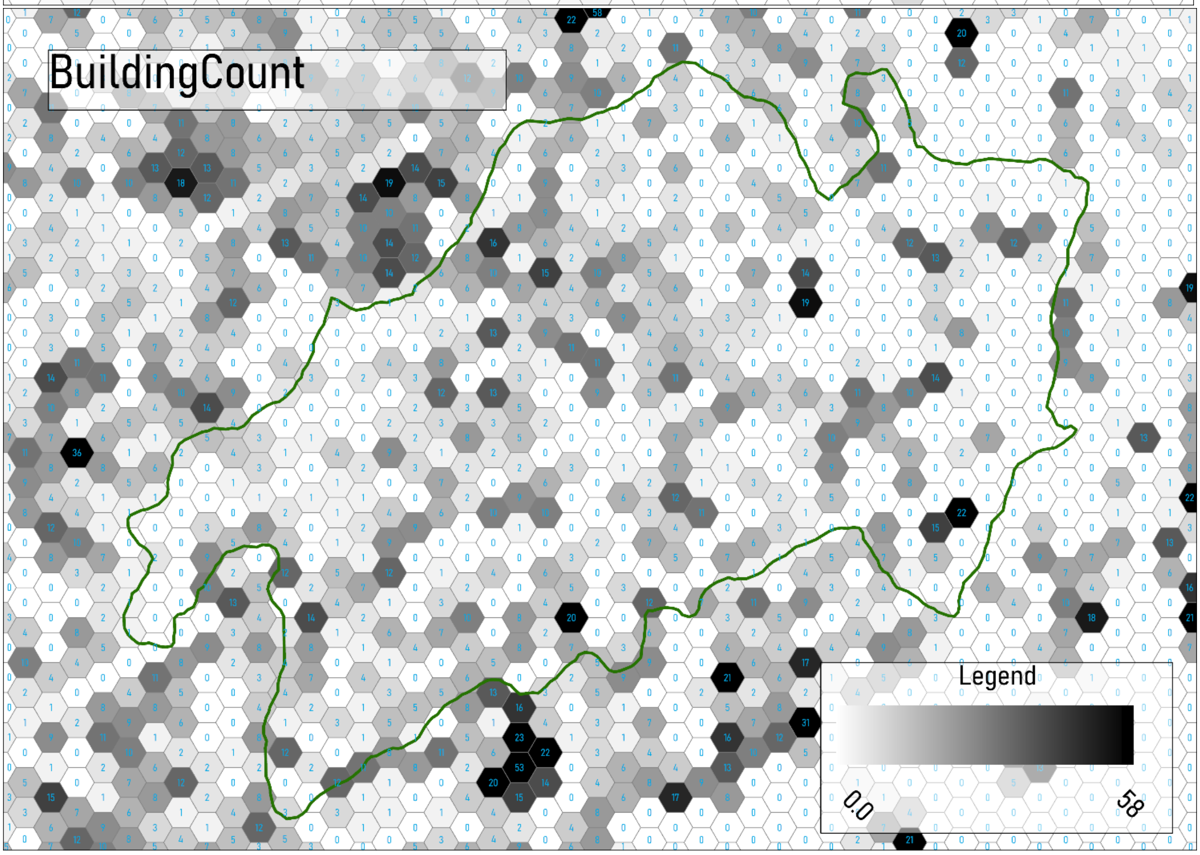

7.2.3. Priority Conservation Area Analysis – Partition Coverage and Feature Count

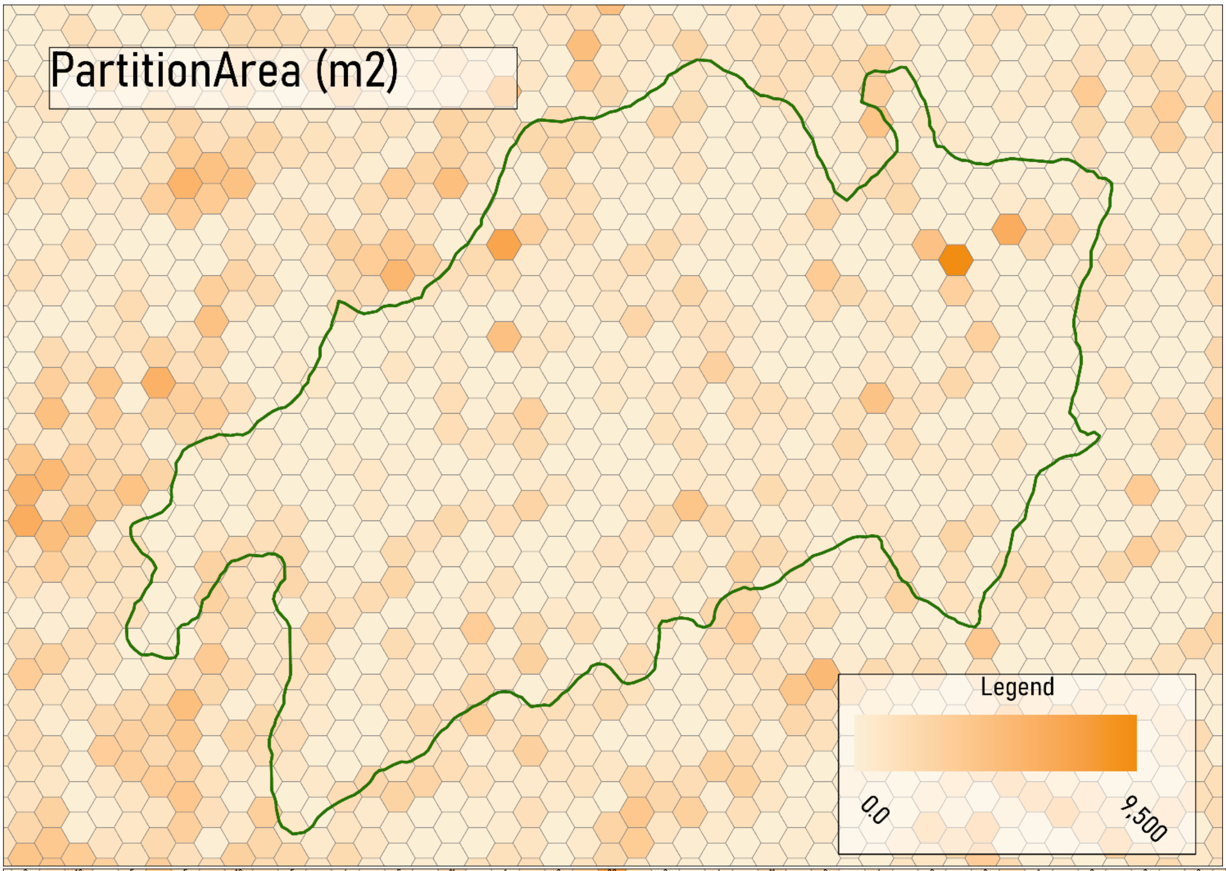

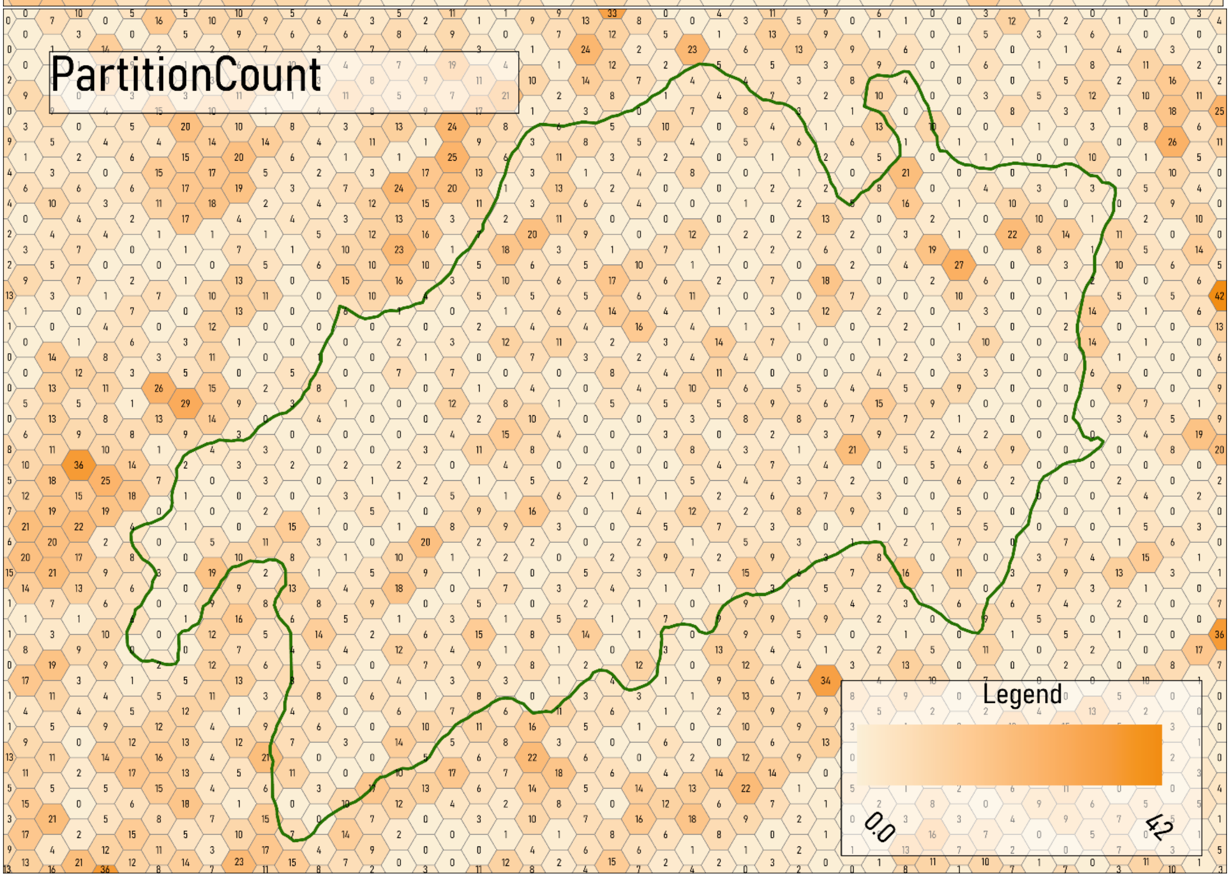

7.2.4. Priority Conservation Area Analysis – Fence Coverage and Feature Count

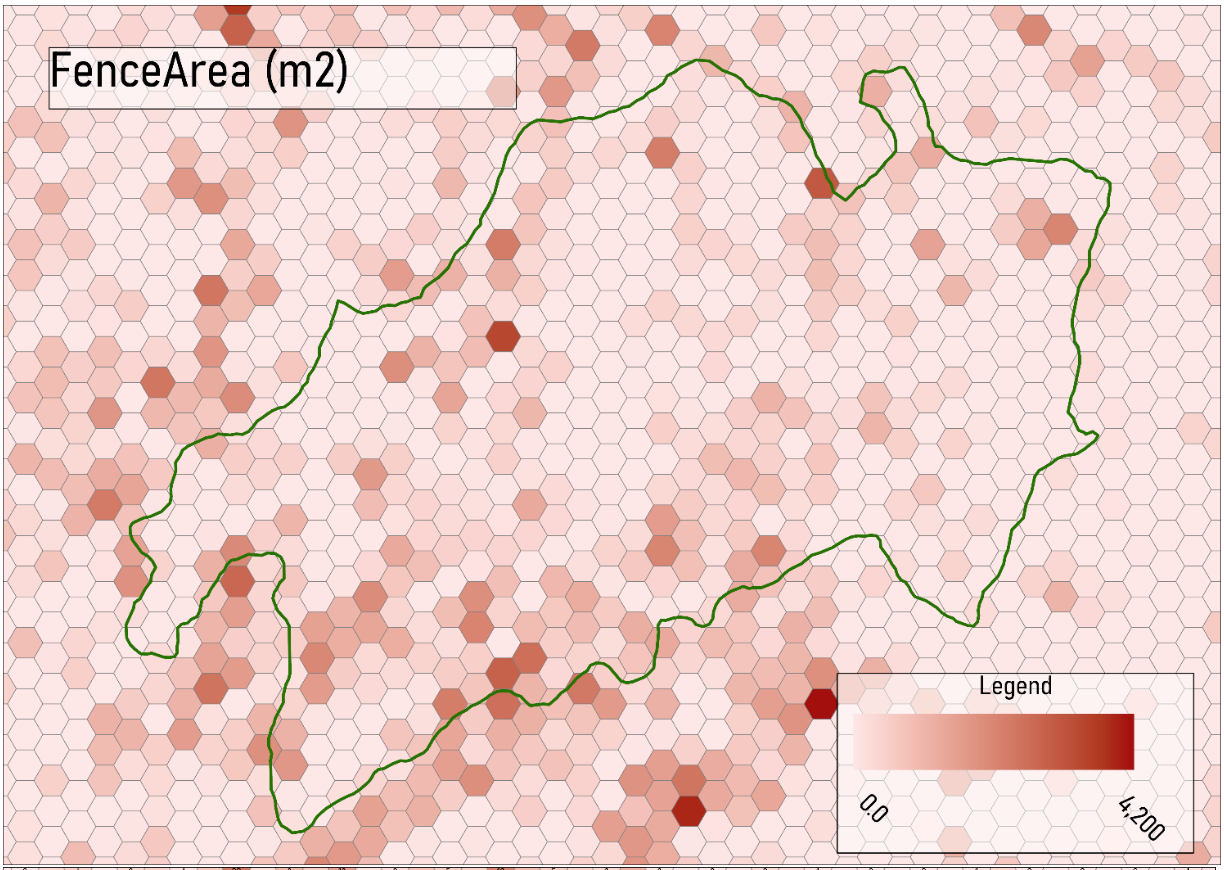

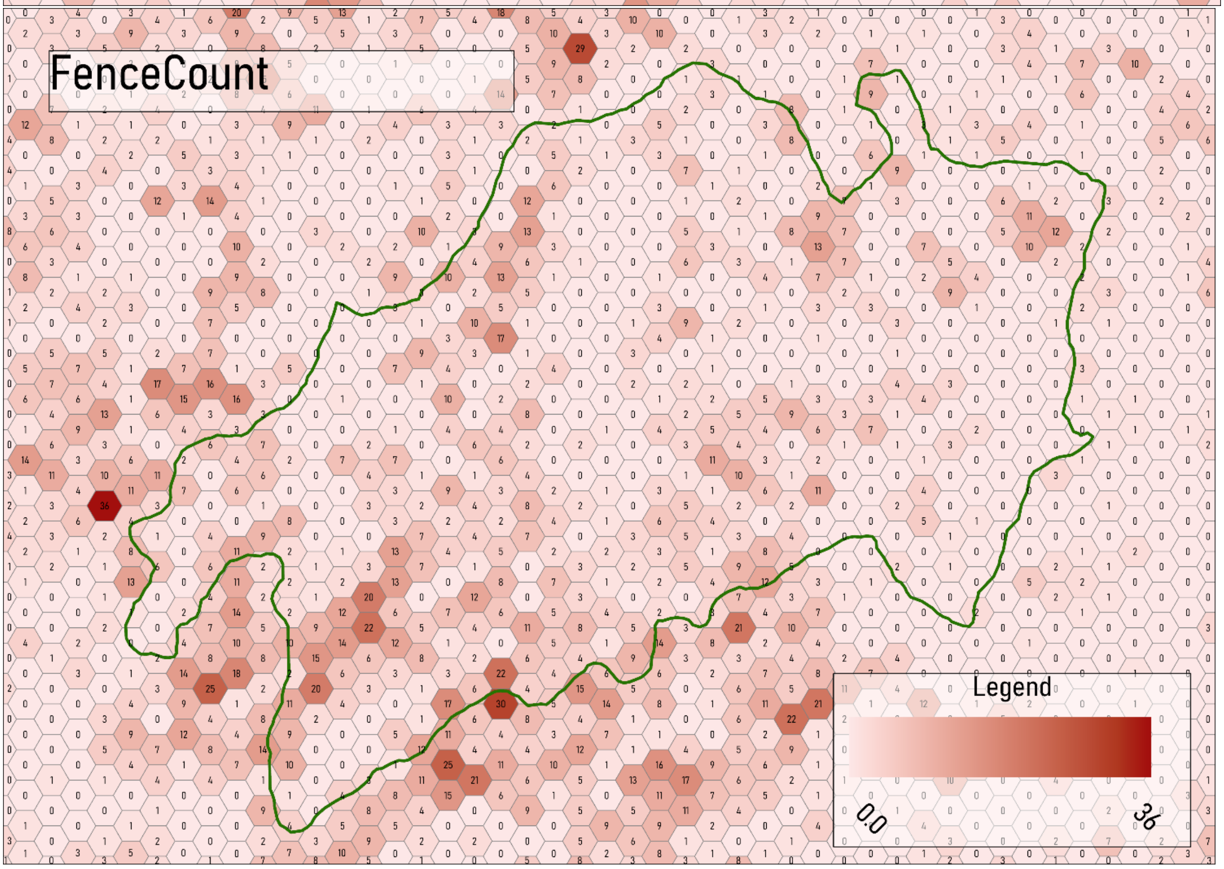

7.2.5. Priority Conservation Area Analysis – Road Coverage and Feature Count

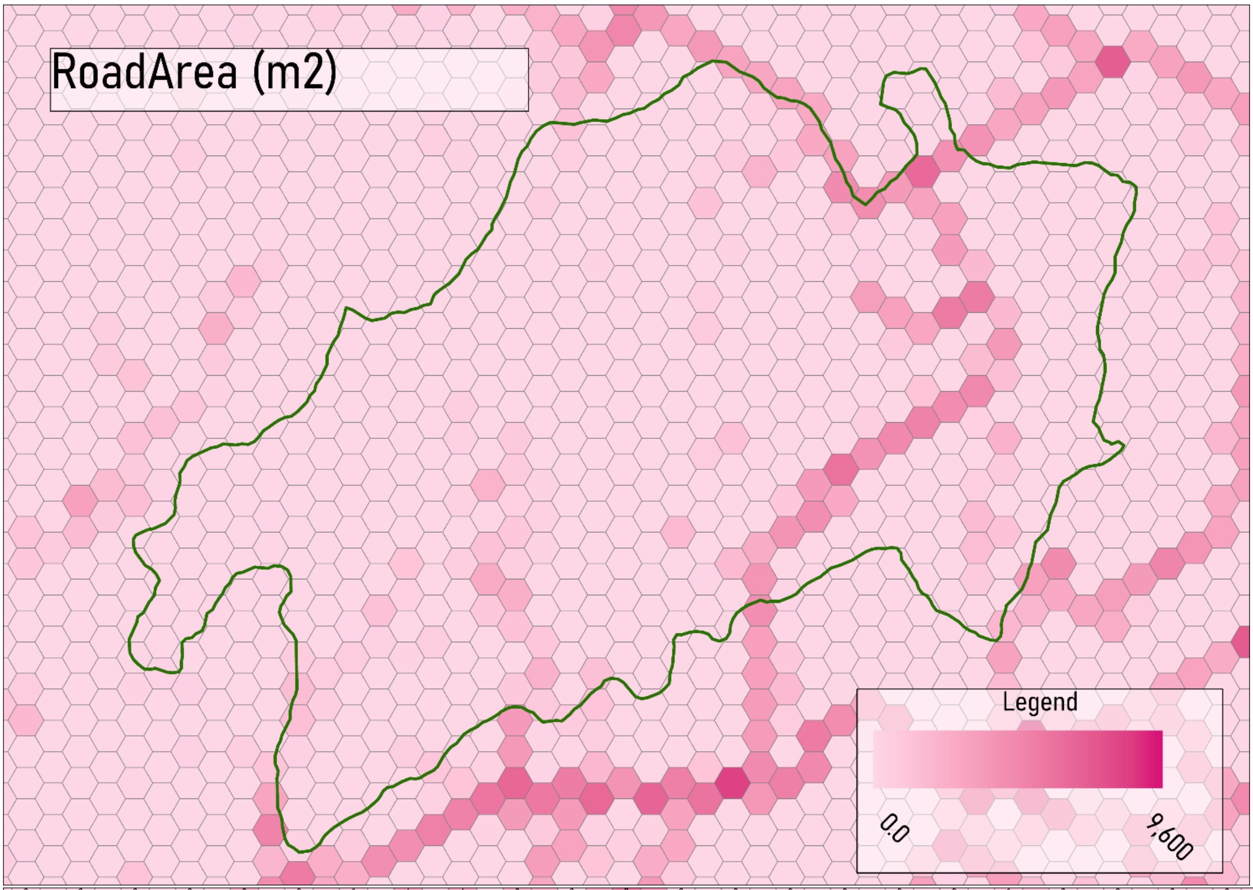

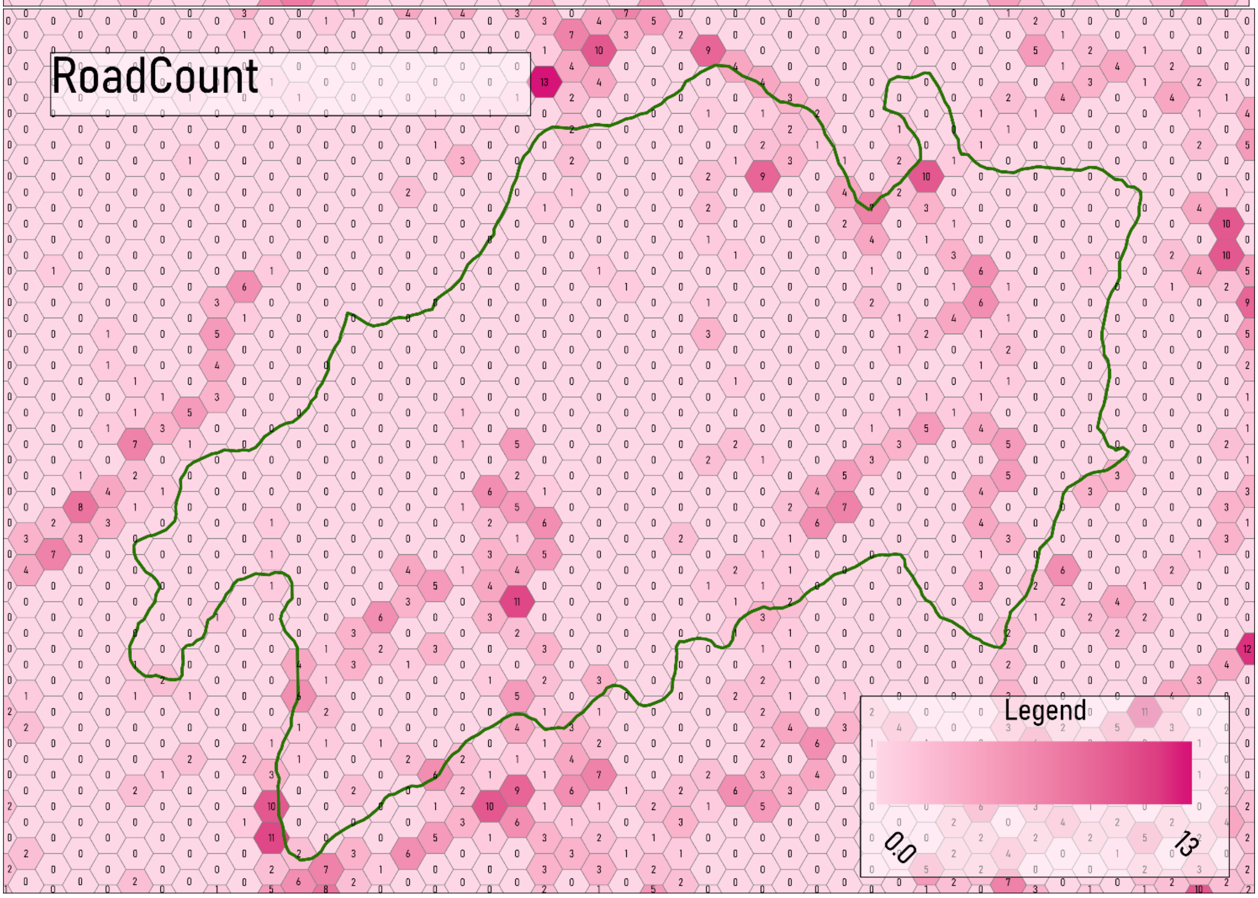

7.2.6. Priority Conservation Area Analysis – Water Coverage and Feature Count

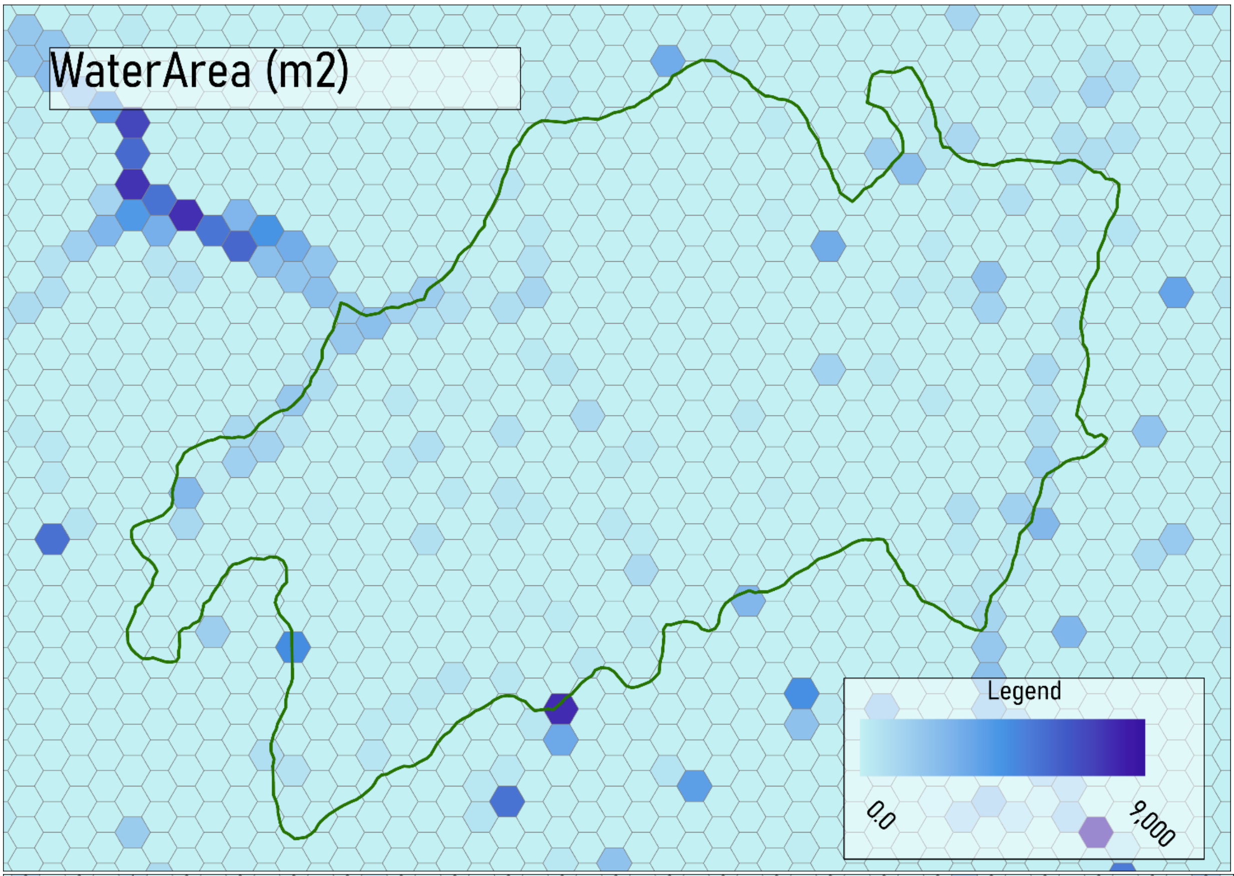
WaterArea (m2)
Legend
0.0
9,000

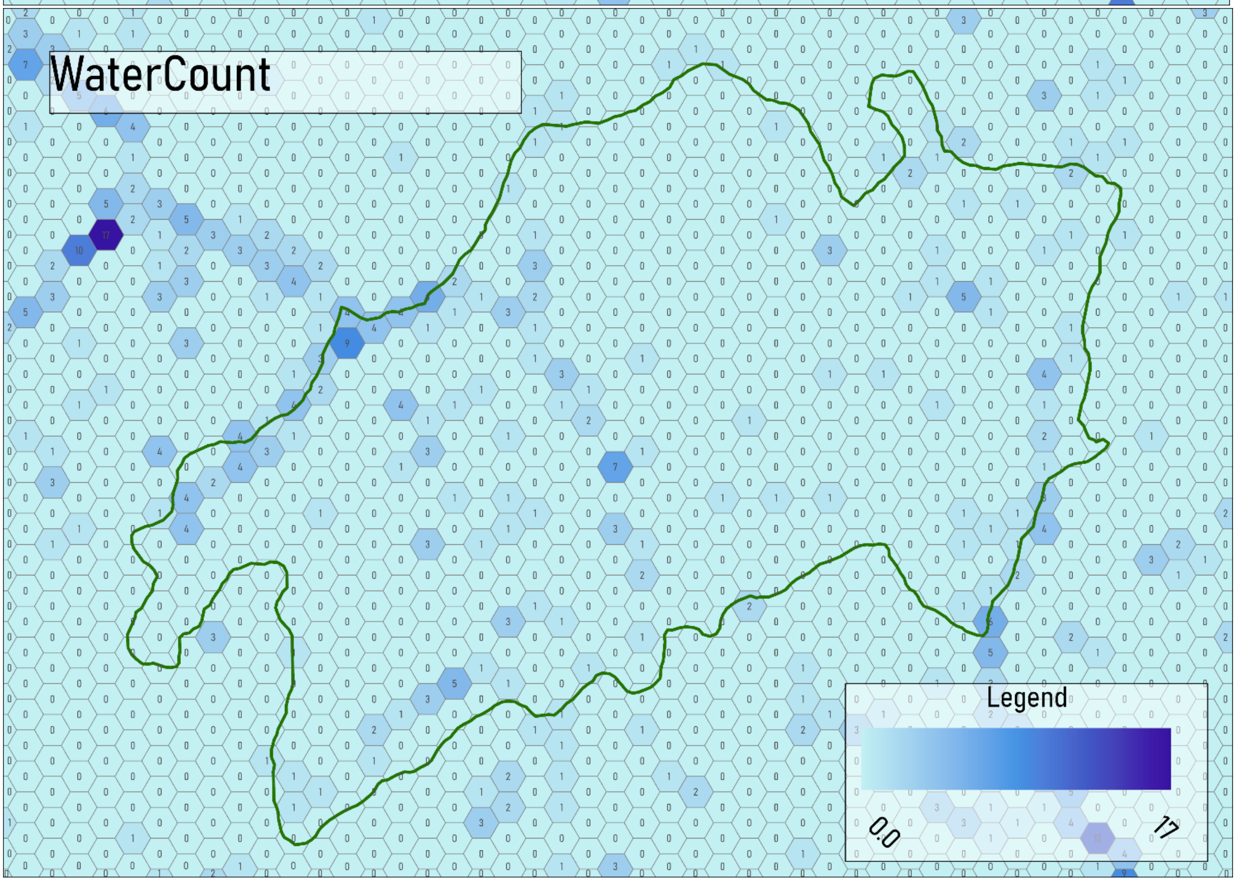
WaterCount
Legend
0.0
17