Integrating UAV Imagery into Agricultural Statistics in the Cook Islands*

Anthony Burgard[1], Takaaki Masaki[†1], Anna Christine Durante[1], and Pamela Lapitan[1]

[1]Asian Development Bank, Manila, Philippines

March 2, 2026

Abstract

Accurate agricultural area measurement underpins reliable production and yield statistics, yet conventional ground surveys in Pacific Island contexts are costly and prone to measurement error. Building on evidence from the 2021 Cook Islands Census of Agriculture and a 2022 post-enumeration survey that revealed systematic underestimation of plot sizes, this paper demonstrates how unmanned aerial vehicle (UAV) imagery can strengthen official agricultural statistics. Using a 2023 UAV survey of a census enumeration area (17.6 ha) in Rarotonga, we develop a fully reproducible workflow in R that integrates centimeter-level RGB and multispectral orthophotos with vector plot boundaries to compute plot areas and extract vegetation indicators, including mean NDVI and within-plot variability. The analysis shows that very high-resolution imagery (1.6 cm ground sample distance) substantially reduces mixed-pixel errors common in 10–30 m satellite data and enables precise boundary delineation for small, fragmented plots. Thousands of pixels per plot support robust zonal statistics and fine-scale assessment of crop vigor that are infeasible with coarser sensors. By combining OpenDroneMap and open-source R packages (terra, sf, tmap), the paper provides a practical, end-to-end framework for integrating UAV-derived data into agricultural statistics systems in resource-constrained island settings.

Keywords: UAV imagery; agricultural statistics; remote sensing; plot area measurement; Cook Islands; R
JEL codes: C81; C83; Q10

Conference submission: ICAS-X Conference 2026. An earlier version of this paper was published as a book chapter in The UN Handbook on Remote Sensing for Agricultural Statistics (https://fao-eostat.github.io/UN-Handbook/).

*The authors certify that the text, figures, tables, and all materials contained in this manuscript are original and free from any copyright violations.

[†]Corresponding author: tmasaki@adb.org

1. Introduction

Accurate land area estimates form the foundation of reliable agricultural statistics, providing critical inputs for estimating crop production, calculating yield, and informing key agricultural policy decisions. Traditional ground surveys used to collect these data, however, face significant resource constraints—especially in the Pacific Islands, where environmental conditions, physical distances between islands, and natural barriers often make data collection cost-prohibitive.

In 2022, the Asian Development Bank (ADB), in collaboration with the Cook Islands Ministry of Agriculture (MOA), conducted a post-enumeration survey on the island of Rarotonga to validate agricultural area estimates reported in the 2021 Cook Islands Census of Agriculture using Global Positioning System (GPS) measurements. The study found that farmers generally underestimated their agricultural land areas by approximately 13% relative to objective GPS measurements, with larger plots (greater than 2,000 m^2) more likely to be underestimated and smaller plots slightly overestimated [1].[1] This discrepancy highlights the salience of non-sampling errors—in particular, measurement error—that can systematically bias national agricultural production estimates.

In a follow-up activity conducted in 2023, the ADB team, together with the MOA, employed unmanned aerial vehicles (UAVs) to capture high-resolution imagery and measure agricultural land areas, aiming to reduce the resource intensity of field surveys in environmentally challenging terrain. While initial setup costs can be higher, UAV imagery enables precise plot-boundary delineation and accurate area measurement. UAVs are particularly suitable for the Pacific context, where frequent cloud cover, fragmented land plots, and difficult terrain require flexible, localized, and detailed data collection approaches.

Drawing on the 2023 Cook Islands UAV survey of a census enumeration area on Rarotonga, this paper has two aims: (i) to provide a reproducible, end-to-end workflow in R—from loading orthophotos and vector boundaries, harmonizing coordinate reference systems (CRS), and visualizing imagery, to computing plot-level zonal statistics—and (ii) to demonstrate, through a worked case study, how very high-resolution data (approximately 1.6 cm ground sampling distance) enable agricultural statistics that are difficult or impossible with 10–30 m satellite sensors. Using red, green, and blue (RGB) and multispectral imagery, we quantify plot areas and derive vegetation indicators such as the Normalized Difference Vegetation Index (NDVI)—including mean, minimum and maximum, and within-plot coefficients of variation—revealing fine-scale patterns in crop vigor and heterogeneity that are obscured at coarser resolutions. We then contrast these plot-level metrics with Sentinel-2 and Landsat-8 views to illustrate mixed-pixel effects and the implications for area estimation, monitoring, and operational decision-making.

By the end of this paper, you will be able to:

- understand the UAV imagery workflow from flight planning to map outputs;
- load and visualize RGB (true-color) and NDVI UAV imagery in R;
- load and manage vector data in R;
- extract plot-level zonal statistics from UAV imagery; and
- compare the spatial resolutions of UAV imagery versus publicly available satellite imagery.

[1] https://dx.doi.org/10.22617/TCS240326-2

2. Methods

2.1. Study area

The study area comprises a single census enumeration area located on the northeastern coast of Rarotonga Island in the Cook Islands, encompassing 17.6 ha. This area was selected to enable cross-validation with area estimates from the 2021 Agriculture Census and the 2022 post-census enumeration survey. The terrain is characterized by largely flat topography extending from inland residential areas to the beach and ocean along the eastern boundary. Land use is mixed, including residential areas, sports facilities (e.g., a rugby field), and fragmented agricultural plots scattered throughout the area.

2.2. UAV image acquisition and processing

UAV imagery of the study area was captured on 28 August 2023 at midday. Two flight missions were conducted. The first mission used an RGB camera and captured 347 aerial images (Figure 1); the second used a multispectral sensor and captured 318 images. Flight missions were planned using DJI Flight Planner software, where operators specified resolution requirements by setting flight height (approximately 100 m) and image overlap (approximately 80%) to ensure adequate coverage.

Figure 1: Positions of 347 RGB aerial images. Source: Asian Development Bank visualization using OpenDroneMap 3.6.0 (2025).

The processing of raw aerial images was conducted in OpenDroneMap, a free and open-source photogrammetry software that stitches overlapping UAV images into a geometrically corrected

single image known as an orthophoto. The software first applies a Structure-from-Motion (SfM) approach [2], which identifies common feature points across images to reconstruct camera positions and build an initial three-dimensional model of the terrain and its features (Figure 2).

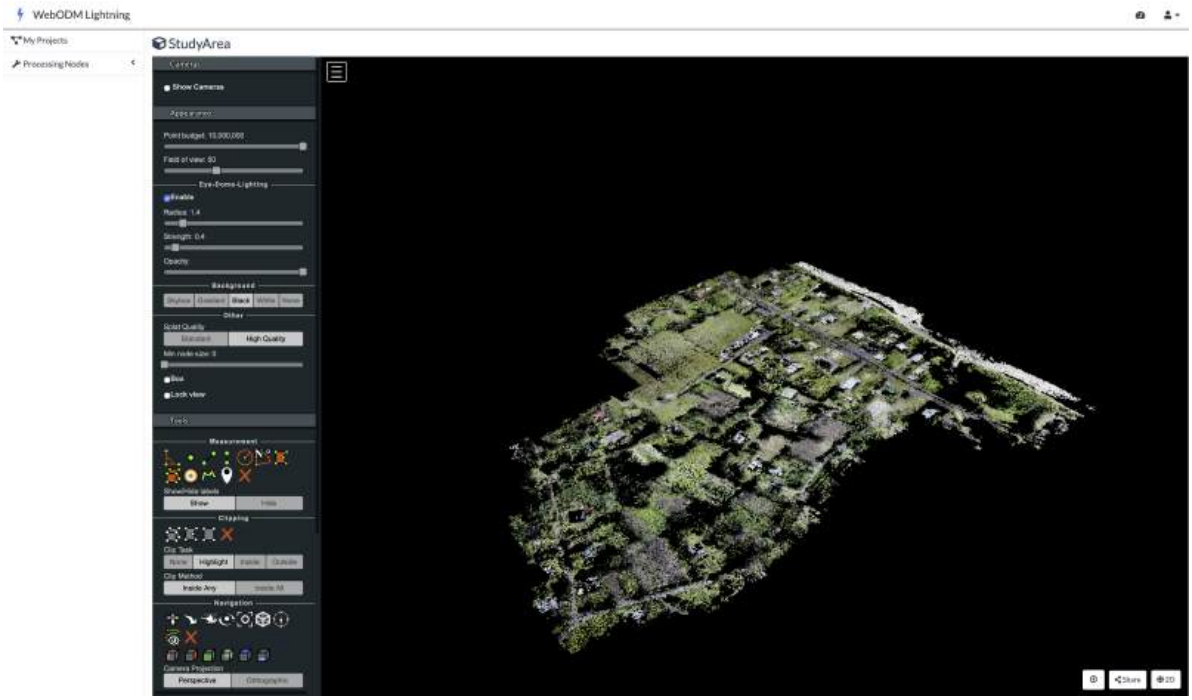

Figure 2: 3D model of the study area constructed from RGB aerial images. Source: Asian Development Bank visualization using OpenDroneMap 3.6.0 (2025).

Prior to the flights, operators established 23 ground control points (GCPs), which are markers with centimeter-level geographic coordinates distributed throughout the study area. Marker coordinates were measured using high-precision Global Navigation Satellite System receivers. During processing in OpenDroneMap, these GCPs are identified in the imagery and used to warp the reconstructed model to fit known coordinates, correcting geometric distortions and ensuring that the final orthophoto is accurately georeferenced. This step is critical for reliable area calculations because it mitigates distortions that would otherwise bias distances and areas.

This workflow was applied to both RGB and multispectral imagery, producing two orthophotos of the same study area captured by two sensors.

In addition to orthophotos, OpenDroneMap was used to compute crop health indicators. The platform implements 23 spectral indices derived from multispectral imagery (Figure 3), as documented in the OpenDroneMap repository.[2] These indices include the NDVI, which is commonly used to assess plant vigor and photosynthetic activity. Other indices can support detection of water stress, approximation of biomass, and identification of soil reflectance in fallow areas.

[2] https://www.researchgate.net/publication/370057094_Open_Drone_Map_Structure-from-Motion_Workflow

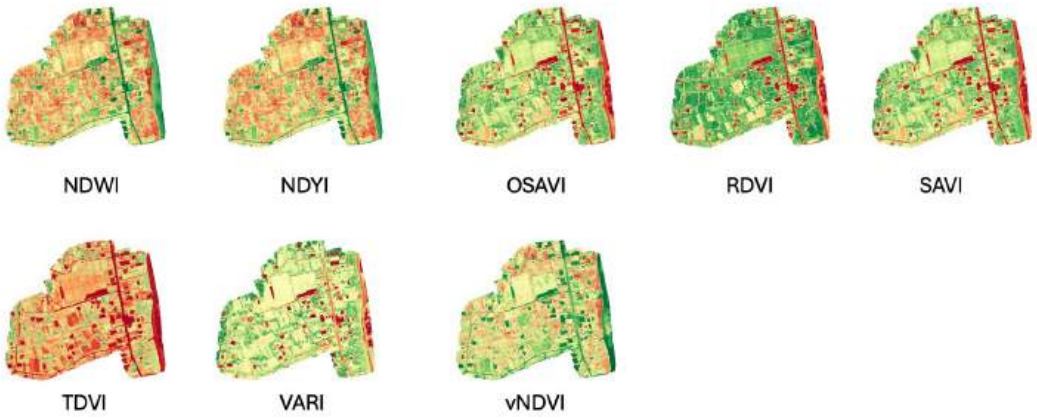

Figure 3: Crop health algorithms in OpenDroneMap. Source: Asian Development Bank visualization using OpenDroneMap 3.6.0 (2025).

2.3. Agricultural plot delineation

Using the RGB orthophoto as a basemap, a subject matter expert from the MOA manually digitized plot boundaries as polygon features in QGIS based on visual interpretation of land use and cultivated plot boundaries. This digitization produced 36 individual agricultural plot polygons.

2.4. Study objectives

This paper provides an analytical framework for working with UAV imagery in R to produce agricultural statistics. Specific objectives include importing and harmonizing raster layers (RGB and NDVI orthophotos) and vector boundaries (study area and plot polygons) to a common CRS; extracting plot-level statistics through zonal analysis of NDVI values; calculating plot areas; and comparing spatial resolution across UAV, Sentinel-2, and Landsat-8 imagery to quantify the mixed-pixel effect and its implications for area estimation accuracy.

2.5. Data and software requirements

The analysis uses four primary datasets:

1. **RGB orthophoto:** a true-color raster stitched in OpenDroneMap from 347 images captured on 28 August 2023 with a DJI Matrice 210.

2. **NDVI orthophoto:** an NDVI raster (values from -1 to 1) stitched and computed in OpenDroneMap from 318 multispectral images captured on 28 August 2023.

3. **Study area boundary:** a vector shapefile representing the boundary of a census enumeration area in Rarotonga, Cook Islands (approximately 17.6 ha).

4. **Plot boundaries:** 36 agricultural plot polygons digitized by an MOA subject matter expert using the RGB orthophoto as a basemap.

To implement the examples in this paper, we use:

- R (version 4.0+) and RStudio;
- R packages: `terra` (1.8-70), `sf` (1.0-21), `ggplot2` (4.0.0), `tmap` (4.2), and `dplyr` (1.1.4); and
- data files: RGB orthophoto, NDVI raster, and boundary shapefiles (provided as an R project).

Important: Open the provided files as an R project in RStudio to ensure correct file paths and organization. All exercises include complete code examples. Basic R familiarity is helpful but not required.

3. Results

3.1. Environment setup and data loading

To work with both raster and vector data in this paper, we use the `terra` package for raster analysis (UAV imagery) and `sf` for vector data (boundaries and plot polygons). For data manipulation and visualization, we use `dplyr`, `ggplot2`, and `tmap`.

```
# Install packages (run once)
# install.packages(c("terra", "sf", "ggplot2", "tmap", "dplyr", "knitr"))

# Load packages
library(terra)    # Raster data analysis
library(sf)       # Vector data (shapefiles, polygons)
library(ggplot2)  # Data visualization
library(tmap)     # Thematic mapping and cartography
library(dplyr)    # Data manipulation
library(knitr)    # Table formatting
```

Next, we load the 2023 Cook Islands UAV survey data. The files for this exercise are provided in an R project, which automatically sets the working directory to the project folder upon opening.

```
# Load the RGB orthophoto
rgb <- rast("data/2 - UAV Image/Map1_orthomosaic_export_WedAug30060354143404.tif")

# Load the NDVI orthophoto
ndvi <- rast("data/2 - UAV Image/ndvi.tif")

# Load the study area boundary
study_area <- st_read("data/1 - Shapefiles/AOI_EA1102.shp")

# Load plot boundaries
plots <- st_read("data/1 - Shapefiles/PLOT_DRONE.shp")
```

The RGB orthophoto contains four layers (red, green, blue, and alpha transparency). Its spatial resolution corresponds to approximately 1.6 cm ground sampling distance. The data use the World Geodetic System 1984 (WGS 84) coordinate reference system (EPSG:4326), with coordinates in decimal degrees.

3.2. Coordinate reference system harmonization

Before conducting further analysis, all layers must use a compatible CRS. Mismatched coordinate systems may result in misaligned layers and incorrect area calculations. We first verify the CRS for each layer and then reproject layers to a common projected CRS appropriate for distance and area calculations.

```
# Check coordinate reference systems
crs(rgb, describe=TRUE)$code      # rgb (raster)
crs(ndvi, describe=TRUE)$code     # ndvi (raster)
st_crs(study_area)$epsg           # study area (vector)
st_crs(plots)$epsg                # plots (vector)
```

In this case, the NDVI raster uses WGS 84 / UTM Zone 4S (EPSG:32704) with units in meters, while other layers use WGS 84 geographic coordinates (EPSG:4326) in degrees. We therefore reproject all layers to UTM Zone 4S, which is appropriate for the Cook Islands and supports meter-based area and distance computations.

```
# Define target projection
target_crs <- "EPSG:32704"

# Transform vector data to UTM Zone 4S
study_area <- st_transform(study_area, target_crs)
plots <- st_transform(plots, target_crs)

# Transform RGB and NDVI raster to UTM Zone 4S
rgb <- project(rgb, target_crs, method="bilinear")
ndvi <- project(ndvi, target_crs, method="bilinear")
```

After reprojection, the resolution of the RGB orthophoto is expressed in meters (approximately 0.016 m), and spatial extents are represented in UTM coordinates.

3.3. Initial visualization

To confirm spatial alignment of the RGB raster and vector layers, we create a map displaying all layers together. This verifies that the orthophoto, plot boundaries, and study area boundary are properly aligned after reprojection.

```
plotRGB(rgb, main="RGB Orthophoto Showing Plot Boundaries and Study Area")
plot(st_geometry(plots), add=TRUE, border="white", lwd=1)
plot(st_geometry(study_area), add=TRUE, border="yellow", lwd=2)
sbar(200, type="line", xy="bottom", divs=8, below="meters", col="black")
north(type=1)
```

Figure 4: RGB Orthophoto showing the study area boundary (yellow) and individual agricultural plot boundaries (white).

3.4. Area of the study area and plot boundaries

To quantify the extent of the study area and plot boundaries, we compute their areas. The study area is converted to hectares for reporting.

```
# Calculate study area size
study_area_m2 <- st_area(study_area)
study_area_ha <- study_area_m2 / 10000  # Convert to hectares
```

Next, we calculate areas for each plot, examine summary statistics, and visualize the distribution of plot sizes using a histogram.

```
# Calculate area for each plot in square meters
plots$area_m2 <- as.numeric(st_area(plots))

# Summary statistics in square meters
summary(plots$area_m2)

# Visualize distribution of plot sizes
```

```
hist(plots$area_m2,
     breaks=15,
     main="Distribution of Plot Sizes",
     xlab="Area (square meters)",
     col="lightblue",
     border="white")
```

3.5. Visualization of NDVI values

We now turn to the NDVI raster. NDVI ranges from -1 to 1: values below 0 typically indicate water or bare soil; values between 0.2 and 0.4 often reflect sparse or stressed vegetation; values between 0.4 and 0.6 indicate moderate vegetation vigor; and values above 0.6 indicate dense, healthy vegetation. Monitoring NDVI can help identify crop stress from water deficits, nutrient deficiencies, pests, or diseases, distinguish vegetated areas from bare soil or built structures, and track vegetation change over time.

Higher NDVI values, however, do not necessarily indicate better crop health, as dense weeds, overgrown vegetation, or mature crops can also yield high NDVI. Interpretation therefore benefits from local context and, where possible, ground validation.

Figure 5: Distribution of plot sizes (square meters) in study area

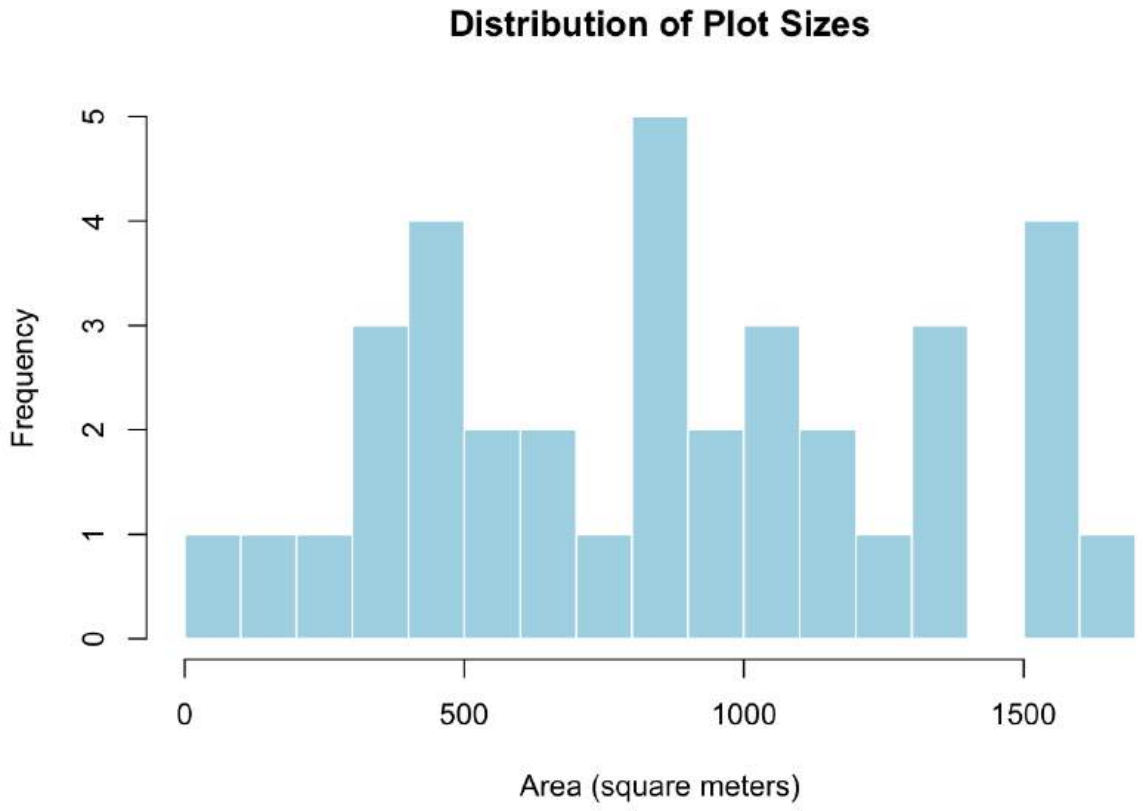

Next, we create an NDVI map with plot boundaries and the study area boundary overlaid.

```
# Create NDVI color scheme
ndvi_colors <- colorRampPalette(c(
  "#8B0000",  # Dark red (very low NDVI)
  "#FF0000",  # Bright red (low NDVI)
  "#FF6B00",  # Orange (low-medium NDVI)
  "#FFFF00",  # Yellow (medium NDVI)
  "#9ACD32",  # Yellow-green (medium-high NDVI)
  "#228B22",  # Forest green (high NDVI)
  "#006400"   # Dark green (very high NDVI)
))(100)

# NDVI Orthophoto Showing Plot Boundaries and Study Area
plot(ndvi,
     col=ndvi_colors,
     main="NDVI Orthophoto Showing Plot Boundaries and Study Area")
plot(st_geometry(plots), add=TRUE, border="white", lwd=1)
plot(st_geometry(study_area), add=TRUE, border="black", lwd=2)
sbar(200, type="line", xy="bottom", divs=8, below="meters", col="black")
north(type=1)
```

Figure 6: NDVI Orthophoto Showing Plot Boundaries and Study Area

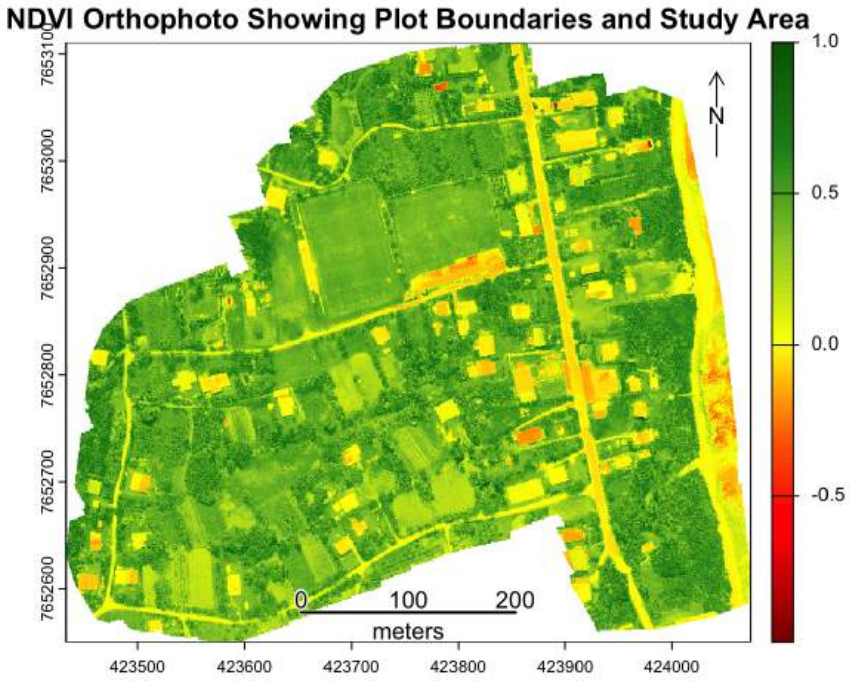

3.6. Calculating plot-level NDVI statistics

With the NDVI raster, we compute zonal statistics for NDVI within each plot. A common challenge in raster-based analysis is edge contamination: pixels near plot boundaries may contain signals from adjacent land covers (e.g., roads, paths, or neighboring plots). To reduce this issue, we apply a 0.5 m negative buffer to each plot polygon before extracting NDVI values. This shrinks each polygon slightly and helps ensure extracted values better represent plot interiors.

```
# Apply negative buffer to reduce edge contamination
plots_buffered <- st_buffer(plots, dist = -0.5)
```

We then compute mean, standard deviation, minimum, and maximum NDVI within each buffered plot, and calculate the coefficient of variation (CV).

```
# Extract mean NDVI for each buffered plot
plot_ndvi <- terra::extract(ndvi, plots_buffered, fun=mean, na.rm=TRUE)
plots$ndvi_mean <- plot_ndvi[,2]

# Additional summary statistics
plot_ndvi_sd  <- terra::extract(ndvi, plots_buffered, fun=sd,  na.rm=TRUE)
plot_ndvi_min <- terra::extract(ndvi, plots_buffered, fun=min, na.rm=TRUE)
plot_ndvi_max <- terra::extract(ndvi, plots_buffered, fun=max, na.rm=TRUE)

plots$ndvi_sd  <- plot_ndvi_sd[,2]
plots$ndvi_min <- plot_ndvi_min[,2]
plots$ndvi_max <- plot_ndvi_max[,2]

# Coefficient of variation (CV)
plots$ndvi_cv <- (plots$ndvi_sd / plots$ndvi_mean) * 100
```

We visualize mean NDVI across plots, where low NDVI may indicate harvested or fallow plots or crop stress, and higher NDVI indicates denser vegetation.

```
# Mean NDVI by Plot
tm_shape(plots) +
  tm_fill(fill="ndvi_mean",
          fill.scale=tm_scale(values="brewer.rd_yl_gn"),
          fill.legend=tm_legend(title="Mean NDVI")) +
  tm_text(text="id", size=0.5) +
  tm_scalebar(position=c("center","bottom")) +
  tm_compass(position=c("right","top")) +
  tm_title("Mean NDVI by Plot")
```

Next, we visualize the CV, which highlights within-plot heterogeneity.

Figure 7: Mean NDVI by plot

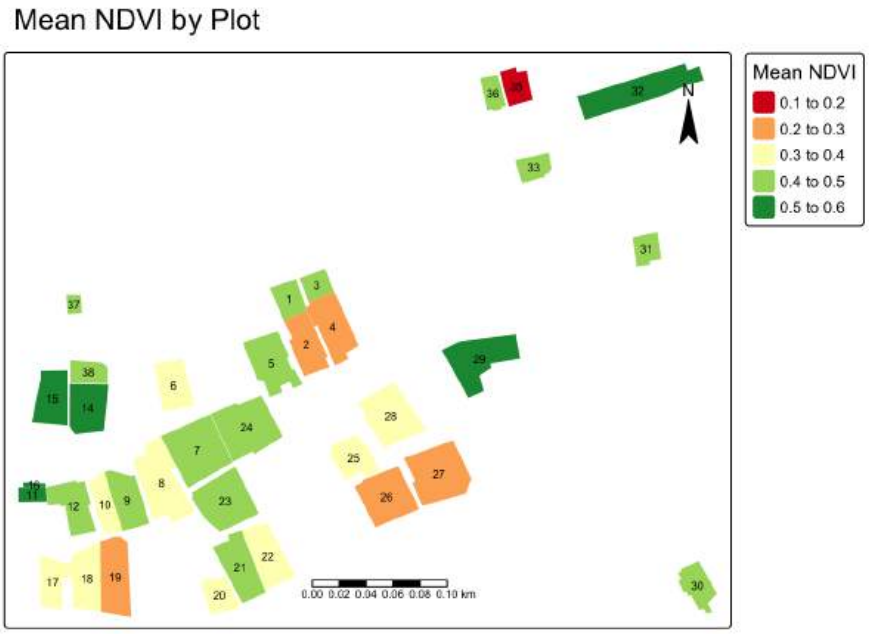


```
# Coefficient of Variation by Plot
tm_shape(plots) +
  tm_fill(fill="ndvi_cv",
          fill.scale=tm_scale(values="brewer.rd_yl_gn"),
          fill.legend=tm_legend(title="CV (%)")) +
  tm_text(text="id", size=0.5) +
  tm_scalebar(position=c("center","bottom")) +
  tm_compass(position=c("right","top")) +
  tm_title("Coefficient of Variation by Plot")
```

3.7. Resolution comparison: UAV vs. satellite imagery

UAV imagery offers much higher spatial resolution than publicly available satellite imagery. Here, we compare UAV imagery at approximately 1.6 cm resolution, Sentinel-2 imagery at 10 m resolution, and Landsat-8 imagery at 30 m resolution.

Higher resolution, however, involves tradeoffs: UAV and high-resolution commercial satellite imagery require greater financial investment, larger storage capacity, and more computational resources. The most appropriate imagery source depends on the analytical objective and scale. Satellite imagery is well suited for regional or national monitoring where broad coverage and frequent revisits are priorities, while UAVs are more appropriate for targeted, plot-level analysis, validation of satellite-based estimates, and monitoring fragmented agricultural landscapes.

We load Sentinel-2 and Landsat-8 imagery and reproject them to the common CRS.

```
# Load satellite imagery
```

Figure 8: Coefficient of Variation by Plot

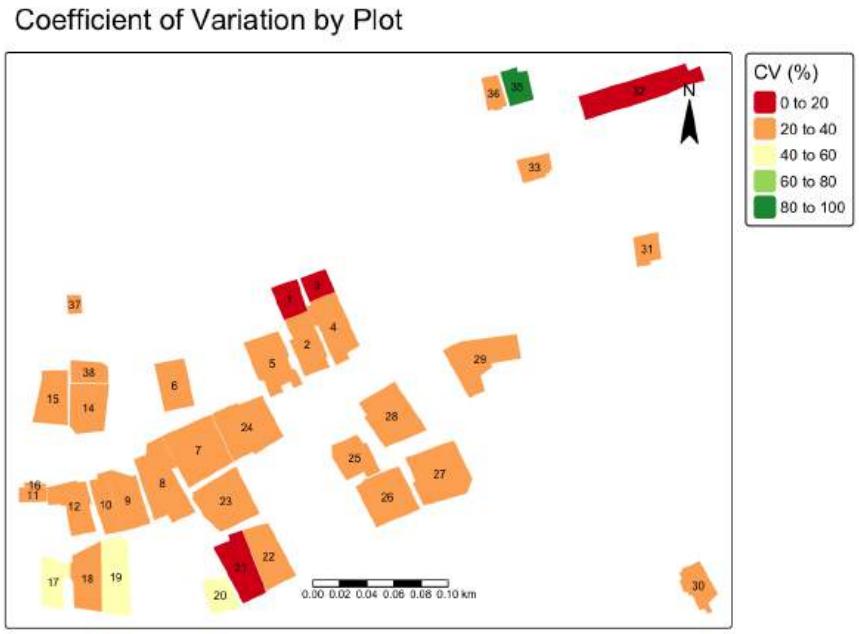


```
sen2 <- rast("data/3 - Satellite Image/sentinel2_10m_Oct.tif")
landsat8 <- rast("data/3 - Satellite Image/landsat8_30m_Oct.tif")

# Reproject to common CRS
sen2 <- project(sen2, target_crs, method="bilinear")
landsat8 <- project(landsat8, target_crs, method="bilinear")
```

Sentinel-2 (10 m resolution). Sentinel-2 provides 10 m pixels (100 m^2 each), which can be suitable for larger agricultural fields but may struggle with small, fragmented plots.

```
# Sentinel-2 True Color Image (10 m)
plotRGB(sen2, r=1, g=2, b=3, stretch="lin", main="Sentinel-2 True Color Image")
plot(st_geometry(plots), add=TRUE, border="white", lwd=1)
plot(st_geometry(study_area), add=TRUE, border="yellow", lwd=2)
sbar(200, type="line", xy="bottom", divs=8, below="meters", col="yellow")
north(type=1)
```

Landsat-8 (30 m resolution). At 30 m resolution (900 m^2 per pixel), plot-level delineation is typically impractical, though the imagery can be useful for broad agricultural zones and land-use patterns.

```
# Landsat-8 True Color Image (30 m)
plotRGB(landsat8, r=1, g=2, b=3, stretch="lin", main="Landsat-8 True Color Image")
plot(st_geometry(plots), add=TRUE, border="white", lwd=1)
```

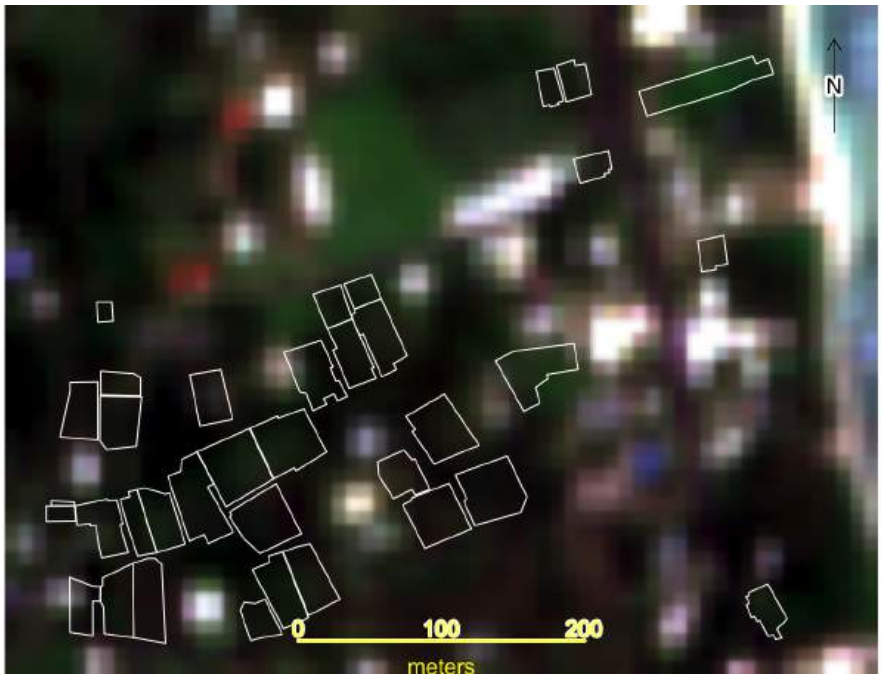

Figure 9: Sentinel-2 True Color Image

```
plot(st_geometry(study_area), add=TRUE, border="yellow", lwd=2)
sbar(200, type="line", xy="bottom", divs=8, below="meters", col="yellow")
north(type=1)
```

Figure 10: Landsat-8 True Color Image

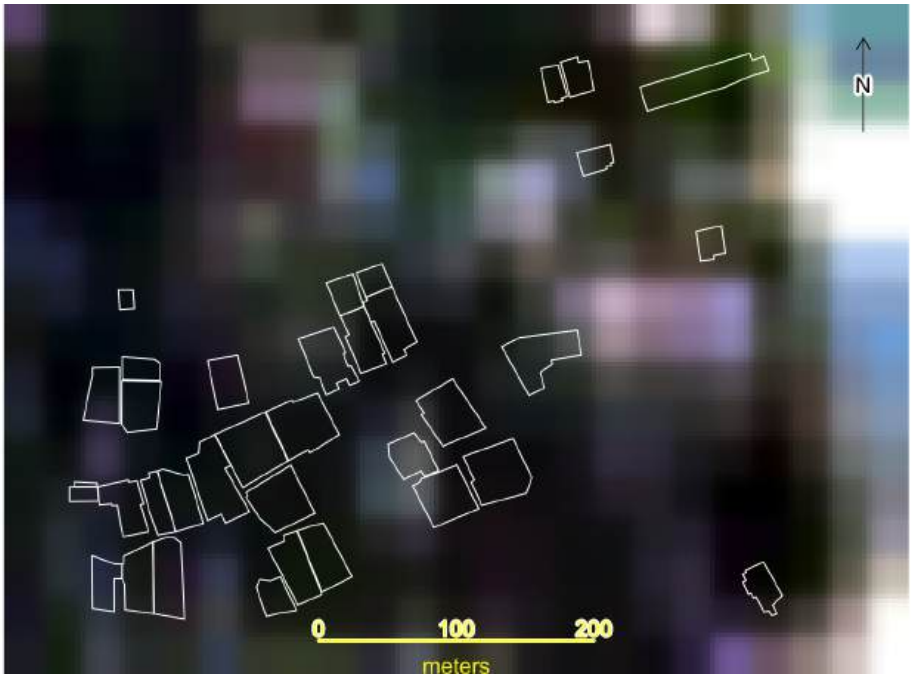

3.8. Quantifying resolution differences

To illustrate resolution differences, we count the number of pixels falling within each plot for each sensor.

```
# Count pixels per plot for each sensor
pixel_counts_rgb <- terra::extract(rgb[[1]], plots, na.rm=TRUE)
plots$pixel_count_rgb <- as.numeric(table(pixel_counts_rgb$ID))

pixel_counts_sen2 <- terra::extract(sen2[[1]], plots, na.rm=TRUE)
plots$pixel_count_sen2 <- as.numeric(table(pixel_counts_sen2$ID))

pixel_counts_landsat <- terra::extract(landsat8[[1]], plots, na.rm=TRUE)
plots$pixel_count_landsat <- as.numeric(table(pixel_counts_landsat$ID))

# Calculate UAV resolution in cm
UAV_res_m <- res(rgb)[1]
UAV_res_cm <- UAV_res_m * 100

# Create comparison data frame
comparison <- data.frame(
  Sensor = c("UAV RGB", "Sentinel-2", "Landsat-8"),
  Resolution = c(paste(round(UAV_res_cm, 2), "cm"), "10 m", "30 m"),
  Min_Pixels = c(min(plots$pixel_count_rgb, na.rm=TRUE),
                 min(plots$pixel_count_sen2, na.rm=TRUE),
                 min(plots$pixel_count_landsat, na.rm=TRUE)),
  Mean_Pixels_per_Plot = c(mean(plots$pixel_count_rgb, na.rm=TRUE),
                           mean(plots$pixel_count_sen2, na.rm=TRUE),
                           mean(plots$pixel_count_landsat, na.rm=TRUE)),
  Max_Pixels = c(max(plots$pixel_count_rgb, na.rm=TRUE),
                 max(plots$pixel_count_sen2, na.rm=TRUE),
                 max(plots$pixel_count_landsat, na.rm=TRUE))
)

knitr::kable(comparison,
             format="pipe",
             digits=1,
             col.names=c("Sensor",
             "Resolution",
             "Min Pixels",
             "Mean Pixels",
             "Max Pixels"),
             caption="Comparison of Pixel Counts per Plot Across Different Sensors")
```

Figure 11: Comparison of Pixel Counts per Plot Across Different Sensors

Sensor	Resolution	Min Pixels	Mean Pixels	Max Pixels
UAV RGB	1.6 cm	185736	3401651.4	6292073
Sentinel-2	10 m	1	9.8	18
Landsat-8	30 m	1	1.9	4

3.9. Mixed pixel problem

A major source of error in satellite-based area measurement is the mixed pixel problem, which occurs when a single pixel covers multiple land cover types. For example, a 30 m Landsat pixel at a field edge may include crop, bare soil, and a pathway within the same pixel. This creates challenges for area estimation, particularly in Pacific Island contexts where agricultural landscapes are dominated by small, irregularly shaped plots. Classifying mixed pixels as crop can overestimate crop area, while classifying them as non-crop can underestimate it, introducing systematic uncertainty and bias in agricultural area statistics.

Selecting an imagery source depends on scale, budget, and objectives. Satellite imagery excels in large-scale monitoring with regular temporal coverage, while UAV imagery is well suited for targeted plot-level analysis where boundary precision is critical.

3.10. Sensor characteristics and applications

Table 1 summarizes key characteristics of satellite and UAV imagery for agricultural statistics applications.

Table 1: Sensor characteristics and applications.

Characteristic	Satellite imagery	UAV imagery
Spatial resolution	Commercial: ~30 cm; Public: 10 m (Sentinel-2) to 1 km (MODIS)	1–5 cm (depends on flight altitude and sensor)
Spectral resolution	Multispectral and SAR; Sentinel-2: 13 bands (visible–SWIR); Landsat-8: 11 bands	RGB (standard); multispectral (5–10 bands), thermal, and LiDAR possible with specialized sensors
Temporal resolution	Sentinel-2: 5 days; Landsat: 16 days; MODIS: daily	User-controlled (on-demand); practical frequency often 2–4 flights/season due to cost and logistics
Coverage area	Large swaths (e.g., Sentinel-2: 290 km)	Limited (often 5–50 ha per flight); battery constrained (typically 15–25 min per flight)
Geometric accuracy	Sentinel-2: ~10–20 m absolute; Landsat-8: ~12–30 m; commercial: ~3–5 m	With GCPs: ±5–10 cm; with RTK: ±2–5 cm; without GCPs: ±2–5 m
Radiometric quality	High (well-calibrated; consistent across dates); surface reflectance products available	Variable; may require calibration targets; sensitive to illumination conditions
Cloud cover	Major limitation in tropical regions	Can fly below clouds; sensitive to wind and rain
Cost	Public: free; commercial: \$10–30/km^2	Initial investment: \$2,000–40,000; operational: ~\$500–2,000 per survey day
Data volume	Moderate per unit area (Sentinel-2 scene: ~1 GB)	Large (raw RGB: ~3 GB; multispectral: ~1.7 GB per survey; orthophotos: ~2–3 GB total)
Processing complexity	Low to moderate (many pre-processed products)	Moderate to high; requires photogrammetry expertise; processing can take 3–8 hours on consumer hardware
Use cases	National/regional crop mapping; temporal monitoring; land use classification; change detection	Plot delineation; ground-truthing; detailed crop assessment; infrastructure mapping

4. Discussion

While UAV imagery offers substantial advantages for agricultural monitoring in fragmented landscapes, successful integration into operational statistical systems requires careful consideration of practical constraints and strategic implementation approaches.

4.1. Operational challenges in using UAVs for agricultural statistics

Key operational challenges include:

- **Limited scalability:** consumer-grade rotary UAVs are constrained by battery capacity and line-of-sight requirements. Larger areas require multiple batteries and extended field time, increasing logistical burden and cost.
- **High costs:** beyond UAV hardware, agencies must budget for batteries and sensors, ground control equipment, software, operator training, insurance, and maintenance.
- **Regulatory barriers:** operations are subject to aviation and privacy regulations that vary across countries, including licensing, altitude limits, and restrictions near sensitive areas.

- **Privacy concerns:** community engagement and transparent communication about flight paths and coverage are essential, even where operations are legally permitted.
- **Technical capacity gaps:** effective deployment requires skilled operators for flight missions and trained analysts for data processing and interpretation.

4.2. Strategies for effective integration of UAV imagery into agricultural statistics

Strategies to support integration include:

- **Build local analytical capacity:** start with small-scale applications and open-source tools, and use public satellite imagery to understand capabilities and limitations before investing in UAV systems.
- **Match UAV use to operational scale:** prioritize applications requiring high detail over small areas, validation exercises, or monitoring priority zones; full coverage mapping across dispersed areas may be impractical.
- **Use sample-based validation:** focus on statistically selected units rather than full coverage to validate farmer-reported data, verify crop types, and identify stress patterns.
- **Follow local regulations:** establish clear operating protocols, coordinate with civil aviation authorities, and implement field safety and community engagement procedures.
- **Document methodologies:** robust documentation of workflows and use cases supports adoption, encourages innovation, and improves data quality over time.

5. Conclusion

This paper highlights the potential of UAV imagery for modernizing agricultural statistics. A primary barrier to adoption by national statistics offices has been the limited availability of documented, end-to-end use cases. By providing a reproducible workflow in R using a case study from Rarotonga, Cook Islands, this paper addresses that gap. Practical demonstrations are essential for establishing the standards and methodologies needed to integrate UAV-derived data into official agricultural statistics production.

The results underscore several key takeaways. First, centimeter-level resolution (approximately 1.6 cm) substantially reduces the mixed-pixel problem relative to coarser satellite sensors such as Sentinel-2 (10 m) and Landsat-8 (30 m). By capturing thousands of observations within even small plots, UAVs support precise boundary delineation and can reduce measurement error that otherwise biases production estimates. Second, the workflow enables extraction of detailed plot-level vegetation metrics, including mean NDVI and its coefficient of variation, which reveal fine-scale patterns of crop health and heterogeneity that are typically invisible at satellite resolutions. These metrics provide actionable information for monitoring and management that is difficult to obtain using coarser imagery alone.

By presenting an approach based on free and open-source tools such as R and OpenDroneMap, we demonstrate a reproducible and cost-effective framework for integrating UAV imagery into agricultural statistics—supporting more reliable agricultural data for evidence-based policy, improved food security, and sustainable resource management in the Pacific and beyond.

References

[1] Asian Development Bank. Enhancing Agricultural Mapping in Asia and the Pacific. Manila, Philippines: Asian Development Bank; 2024.

[2] Marie V, Zaiats A, Wickersham R, Caughlin TT. Open Drone Map: Structure-from-Motion Workflow. University of Idaho; 2023.