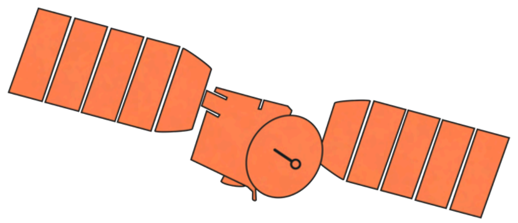

Check out our website!

TÉCNICO
UNIVERSIDADE DE LISBOA

CROP *Auditing*

Growing the Future From Above

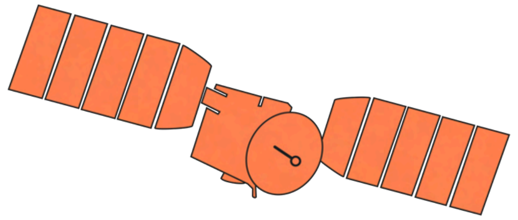

CONTENT

TÉCNICO
UNIVERSIDADE
DE LISBOA

PROJECT'S SCIENTIFIC ADVISOR AND MENTOR

- **Scientific Advisor:**
  - Luis Caldas de Oliveira (lco@tecnico.ulisboa.pt)
  - Afonso Cruz (afonsobpcruz@tecnico.ulisboa.pt)
- **Coordinator:**
  - Luis Caldas de Oliveira (lco@tecnico.ulisboa.pt)
- **Mentor:**
  - Tiago Morais (tiago.morais@virtuacrop.com)

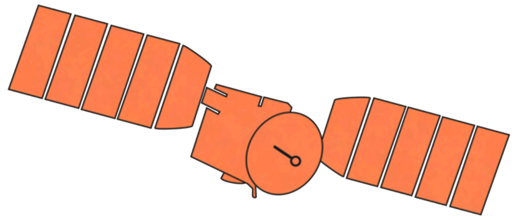

PROBLEM

TÉCNICO UNIVERSIDADE DE LISBOA

Problem

Keeping track of large-scale plantations is currently hard and inefficient, mostly done manually, drone or even helicopter which is not only time consuming but very resource wasteful, whether it being for insurance companies during audits or for any farmer trying to keep track of his plots progress

Challenge

Scaling agricultural oversight through advanced image processing solutions that guarantee accuracy across vast crop areas, ensuring maximum operational efficiency while minimizing the need for resource-intensive field inspections.

Solution

We harness the power of satellite-driven machine learning models and data science to deliver comprehensive insights into crop distribution, eliminating manual overhead while ensuring scalable monitoring throughout the growing season.

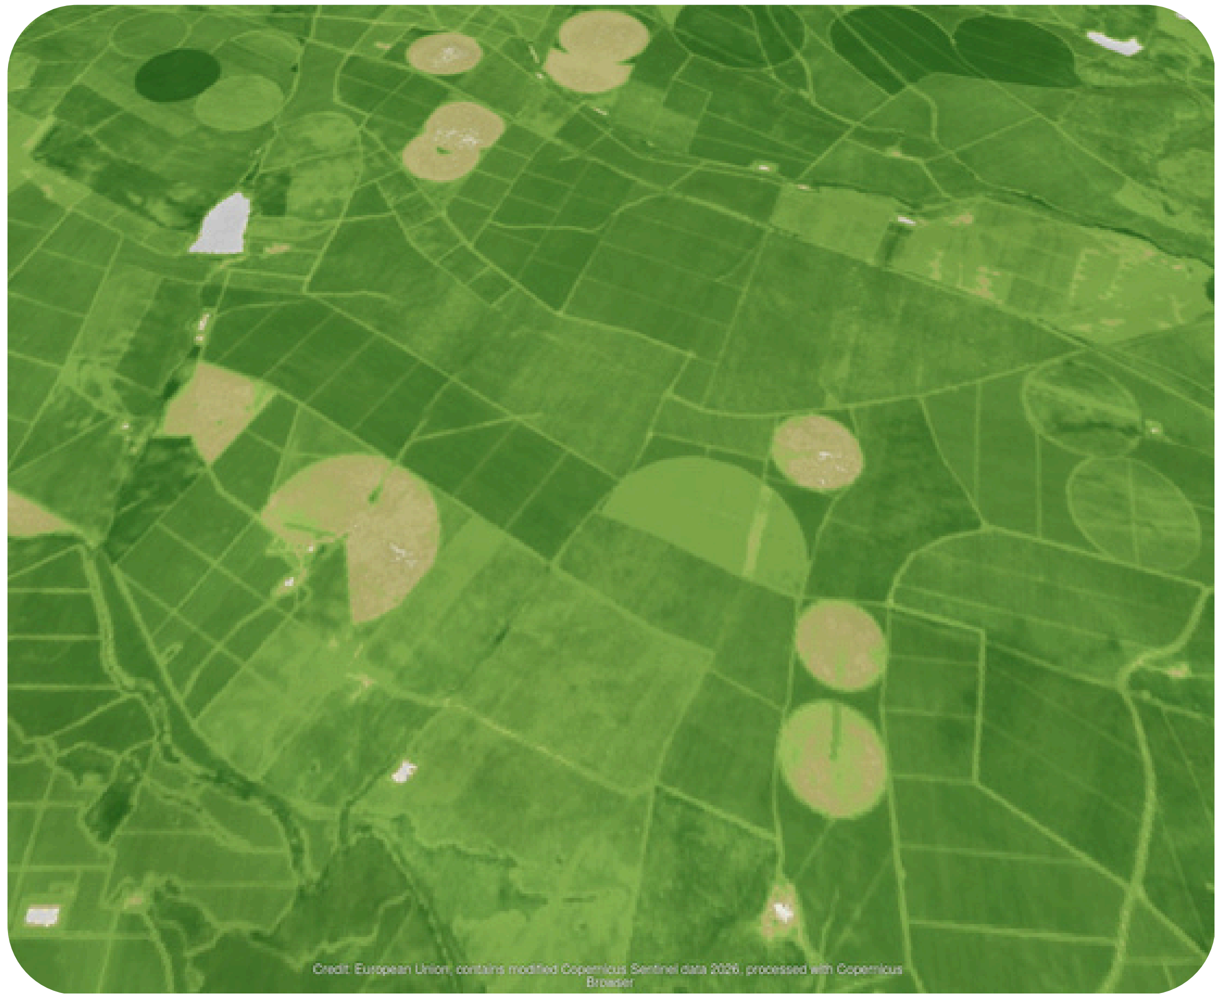

Credit: European Union, contains modified Copernicus Sentinel data 2026, processed with Copernicus Browser

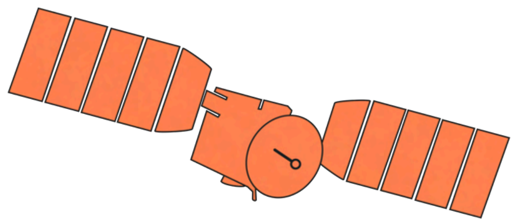

SOLUTION REQUIREMENTS

Accuracy

- Reliable crop identification and area estimation are essential for insurers, farmers and financial institutions.
- The model must be validated against parcels outside training data to ensure trustworthy results.

Performance

- The platform must be fast, simple and usable by non-technical users.
- Results should be delivered within a practical timeframe, with a clear interface for crop auditing and monitoring.

DATABASE + SATELLITE

Countries in database
Countries not in database

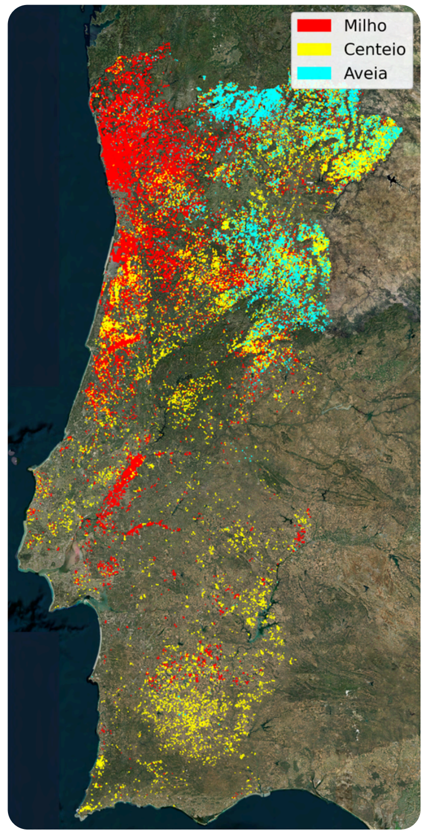

The EuroCropsV2 database, which includes information on the planted crop in a large number of parcels across Europe from 2017 to 2023, is used to train our model on a subset of corn, rye and oat plantations in Portugal.

The Sentinel-2 satellite gives us weekly time series on the NDVI, NDRE, SAVI, EVI and LSWI indexes, derived from frequency bands collected by its sensors, for each year and parcel in our database.

Database → Sentinel-2 Satellite →

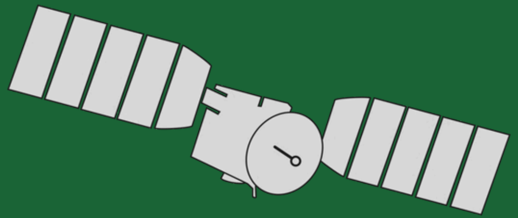

SATELLITE INDEXES

Color

NDVI - Normalized Vegetation Index

Index for qualifying the health and density of vegetation. Uses Red and Near Infrared frequency bands.

$$NDVI = \frac{NIR - RED}{NIR + RED}$$

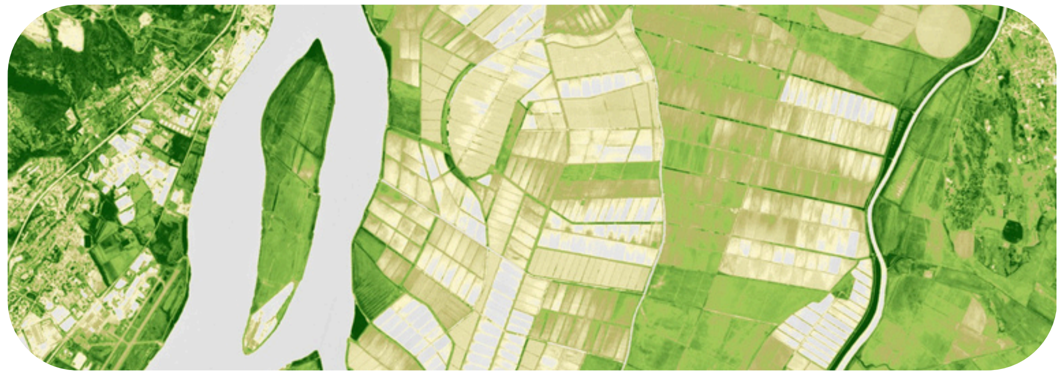

SAVI - Soil Adjusted Vegetation Index

Similar to NDVI, but with corrections to minimize bare-soil reflections.

$$SAVI = \frac{NIR - RED}{NIR + RED + L} \cdot (1 + L),$$

$$L \text{ (soil brightness correction factor)} \in [0, 1]$$

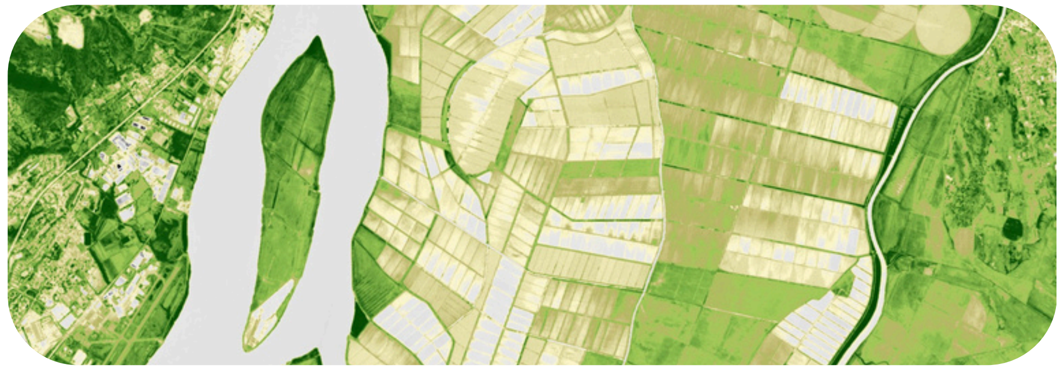

EVI - Enhanced Vegetation Index

Vegetation index for areas with high biomass. Uses Red, Near Infrared and Blue frequency bands.

$$EVI = 2.5 \cdot \left(\frac{NIR - RED}{NIR + 6 \cdot RED - 7.5 \cdot BLUE + 1} \right)$$

NDRE - Normalized Index Red Edge Index

Vegetation index that measures chlorophyll in plants. Uses Red Edge and Near Infrared frequency bands.

$$NDRE = \frac{NIR - RED\ EDGE}{NIR + RED\ EDGE}$$

LSWI - Land Surface Water Index

Measures vegetation and soil water content. Uses Near Infrared and Short-wavelength infrared frequency bands.

$$LSWI = \frac{NIR - SWIR}{NIR + SWIR}$$

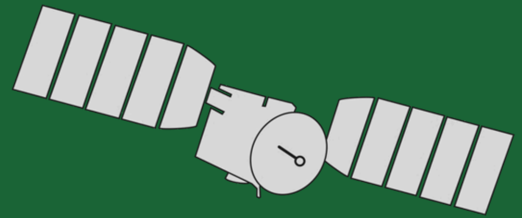

MACHINE LEARNING

LightGBM

With data on the chosen indexes for each parcel, after applying preprocessing such as Tukey's Fences for outlier rejection, filtering invalid and out-of-season values, we train a gradient boosting Machine Learning model (LightGBM) to identify the crops in our dataset.

Sentinel 2 Satellite

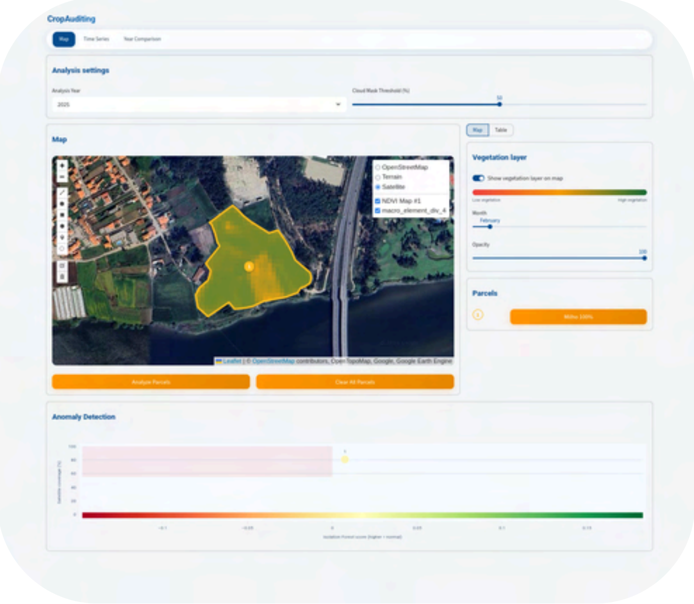

Machine Learning → Sentinel 2 Satellite → App

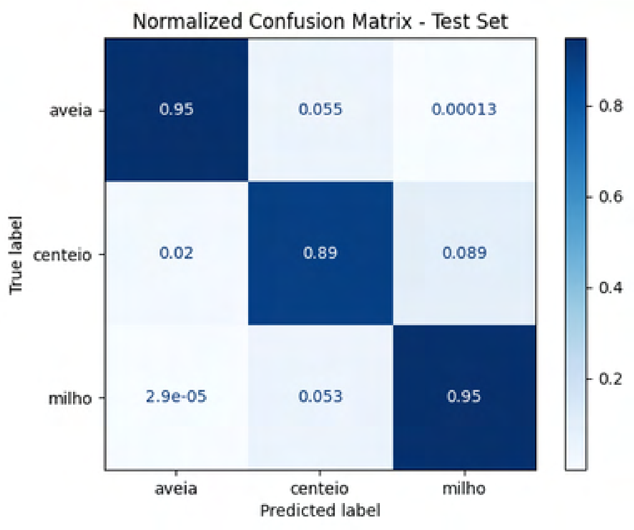

	Precision	Recall	F1-Score
Aveia	0.99	0.95	0.96
Centeio	0.86	0.89	0.87
Milho	0.94	0.95	0.94

With information on the parcel to be analyzed on user request, we collect data from the satellite, apply our Machine Learning model to identify the planted crop and display information on the plantation area, vegetation density and weekly evolution of satellite indexes.

APPLICATION

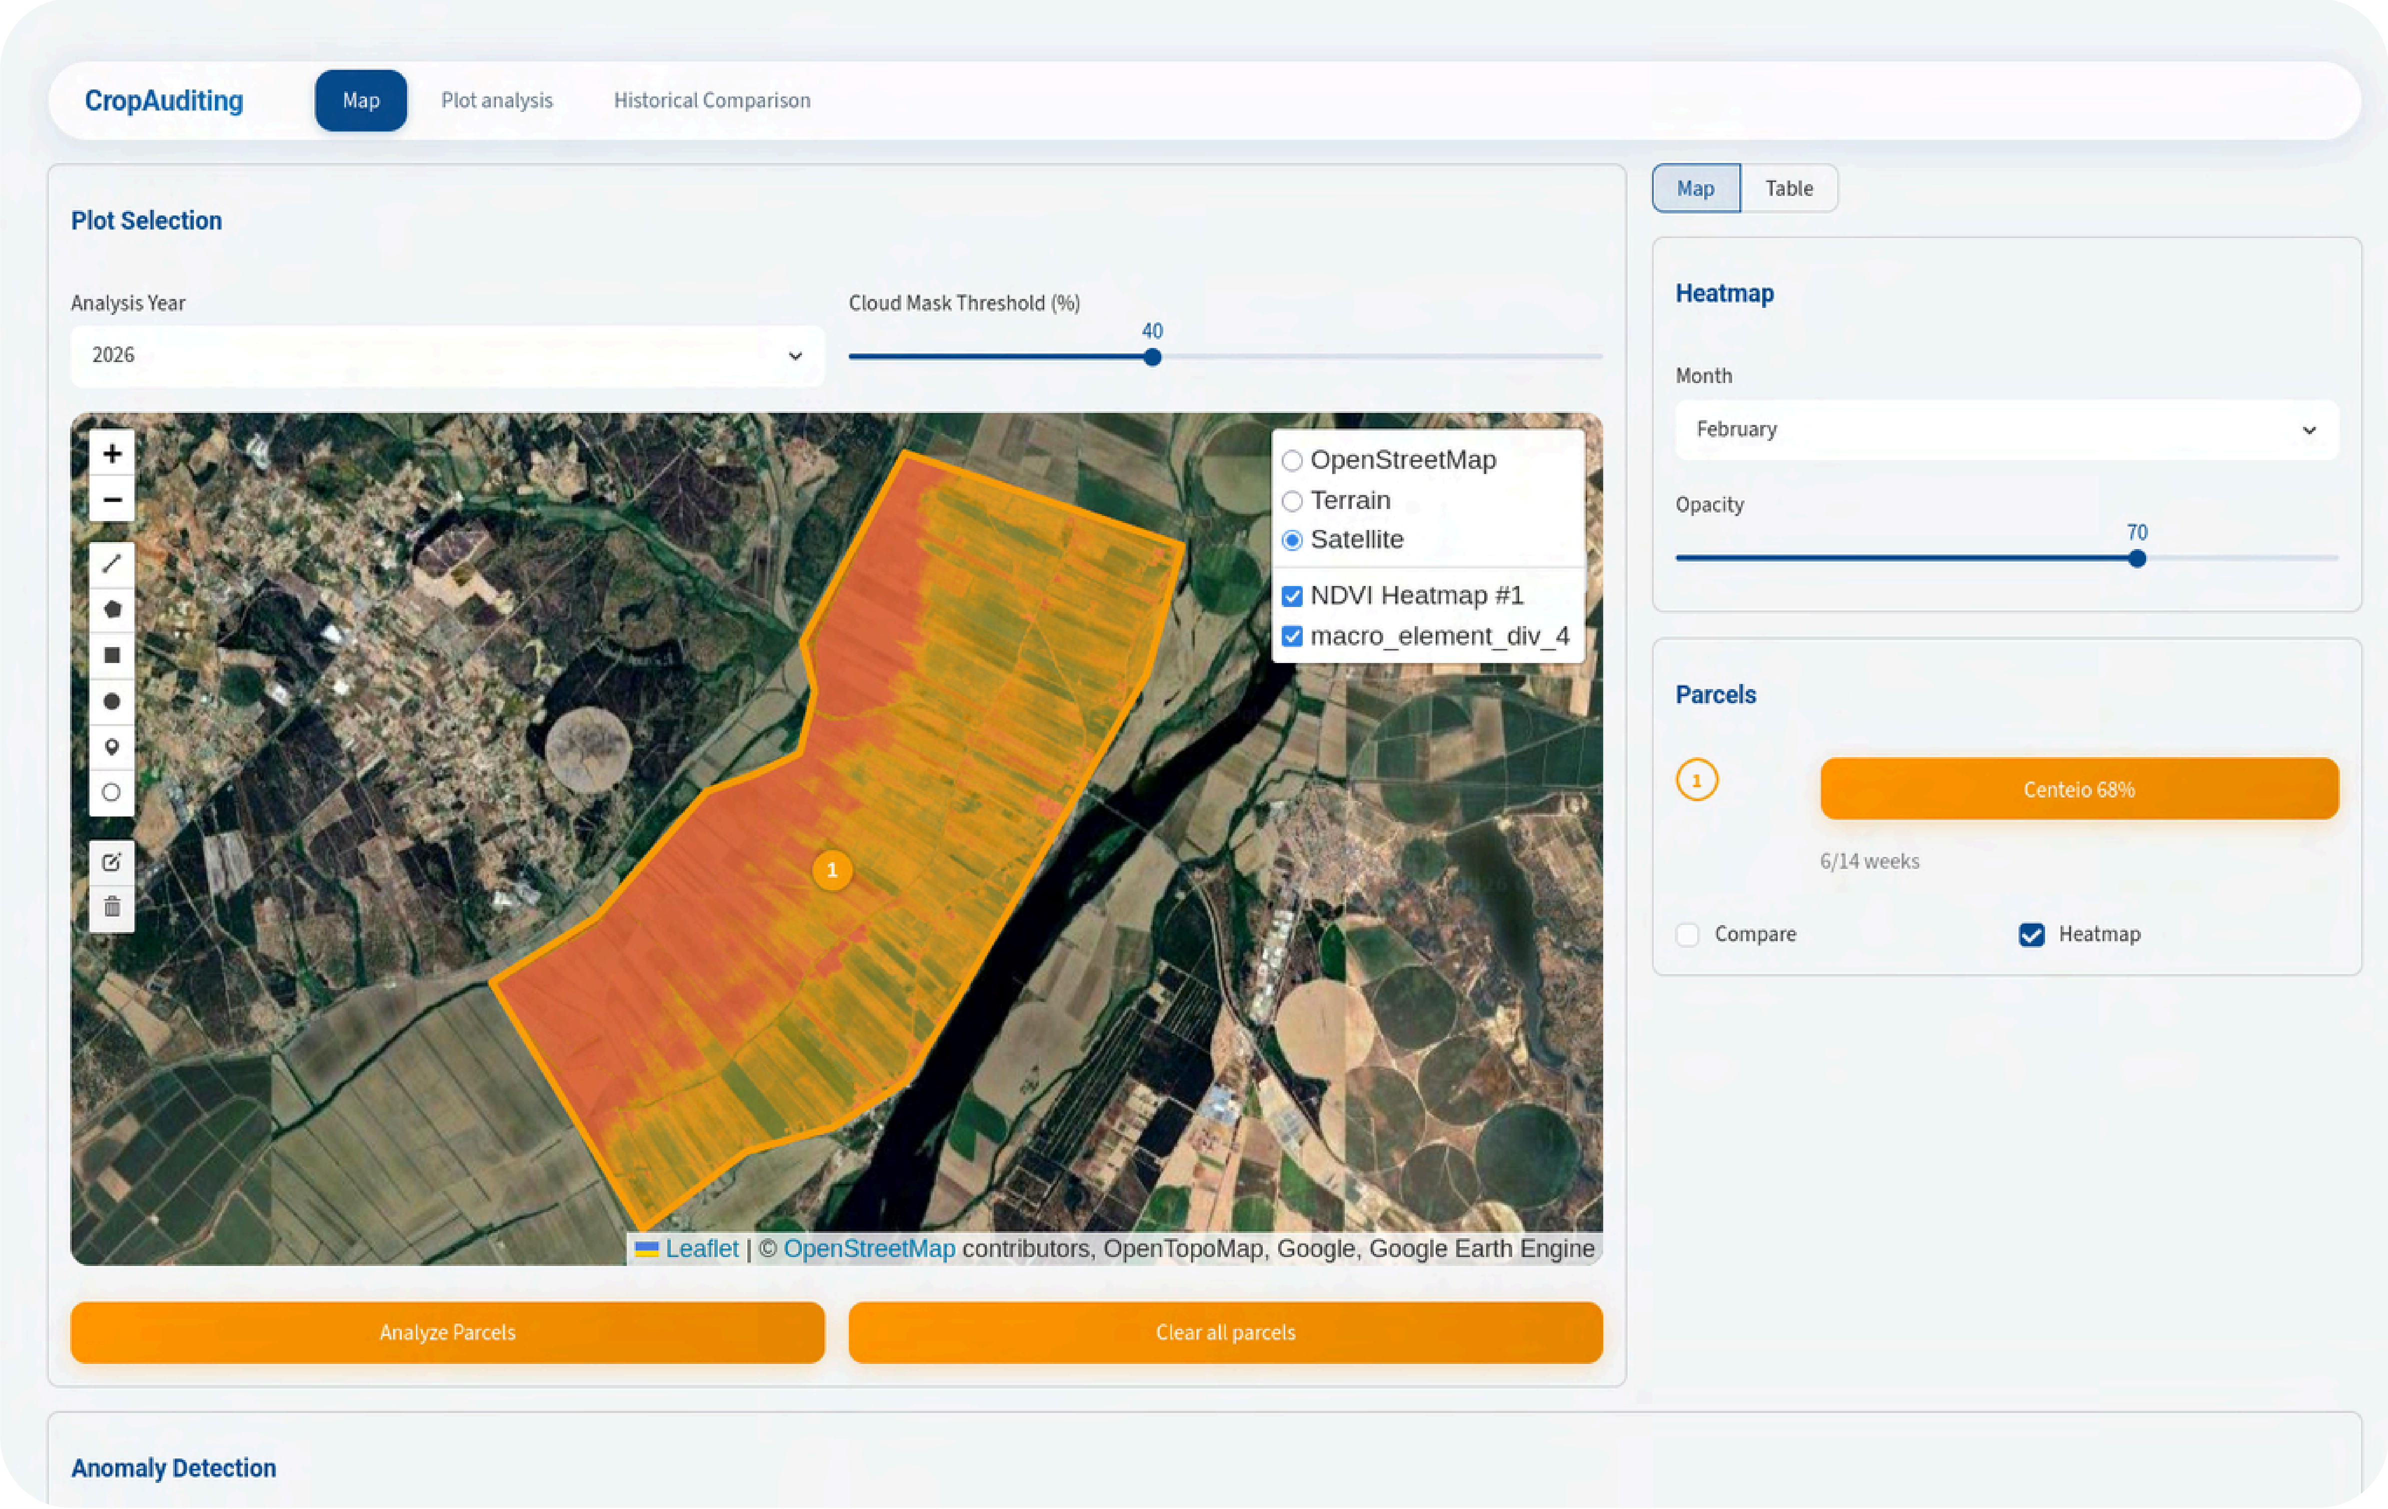

- Allows the user to select large plots of land
- Easier to verify the crop's expected growth
- HeatMap representation for easy visualization of planted area
- Can easily detect underperforming parts of the plantation

APPLICATION

- Allows the user to select specific plots identifying the correct type of crop with an accuracy of near 90%.
- Returns relevant results, with the heatmap representation.
- The user can choose how selective the data is, moving the cloud threshold slider.
- The more cloud coverage the user allows, the more data the model receives, but it might lose some accuracy.

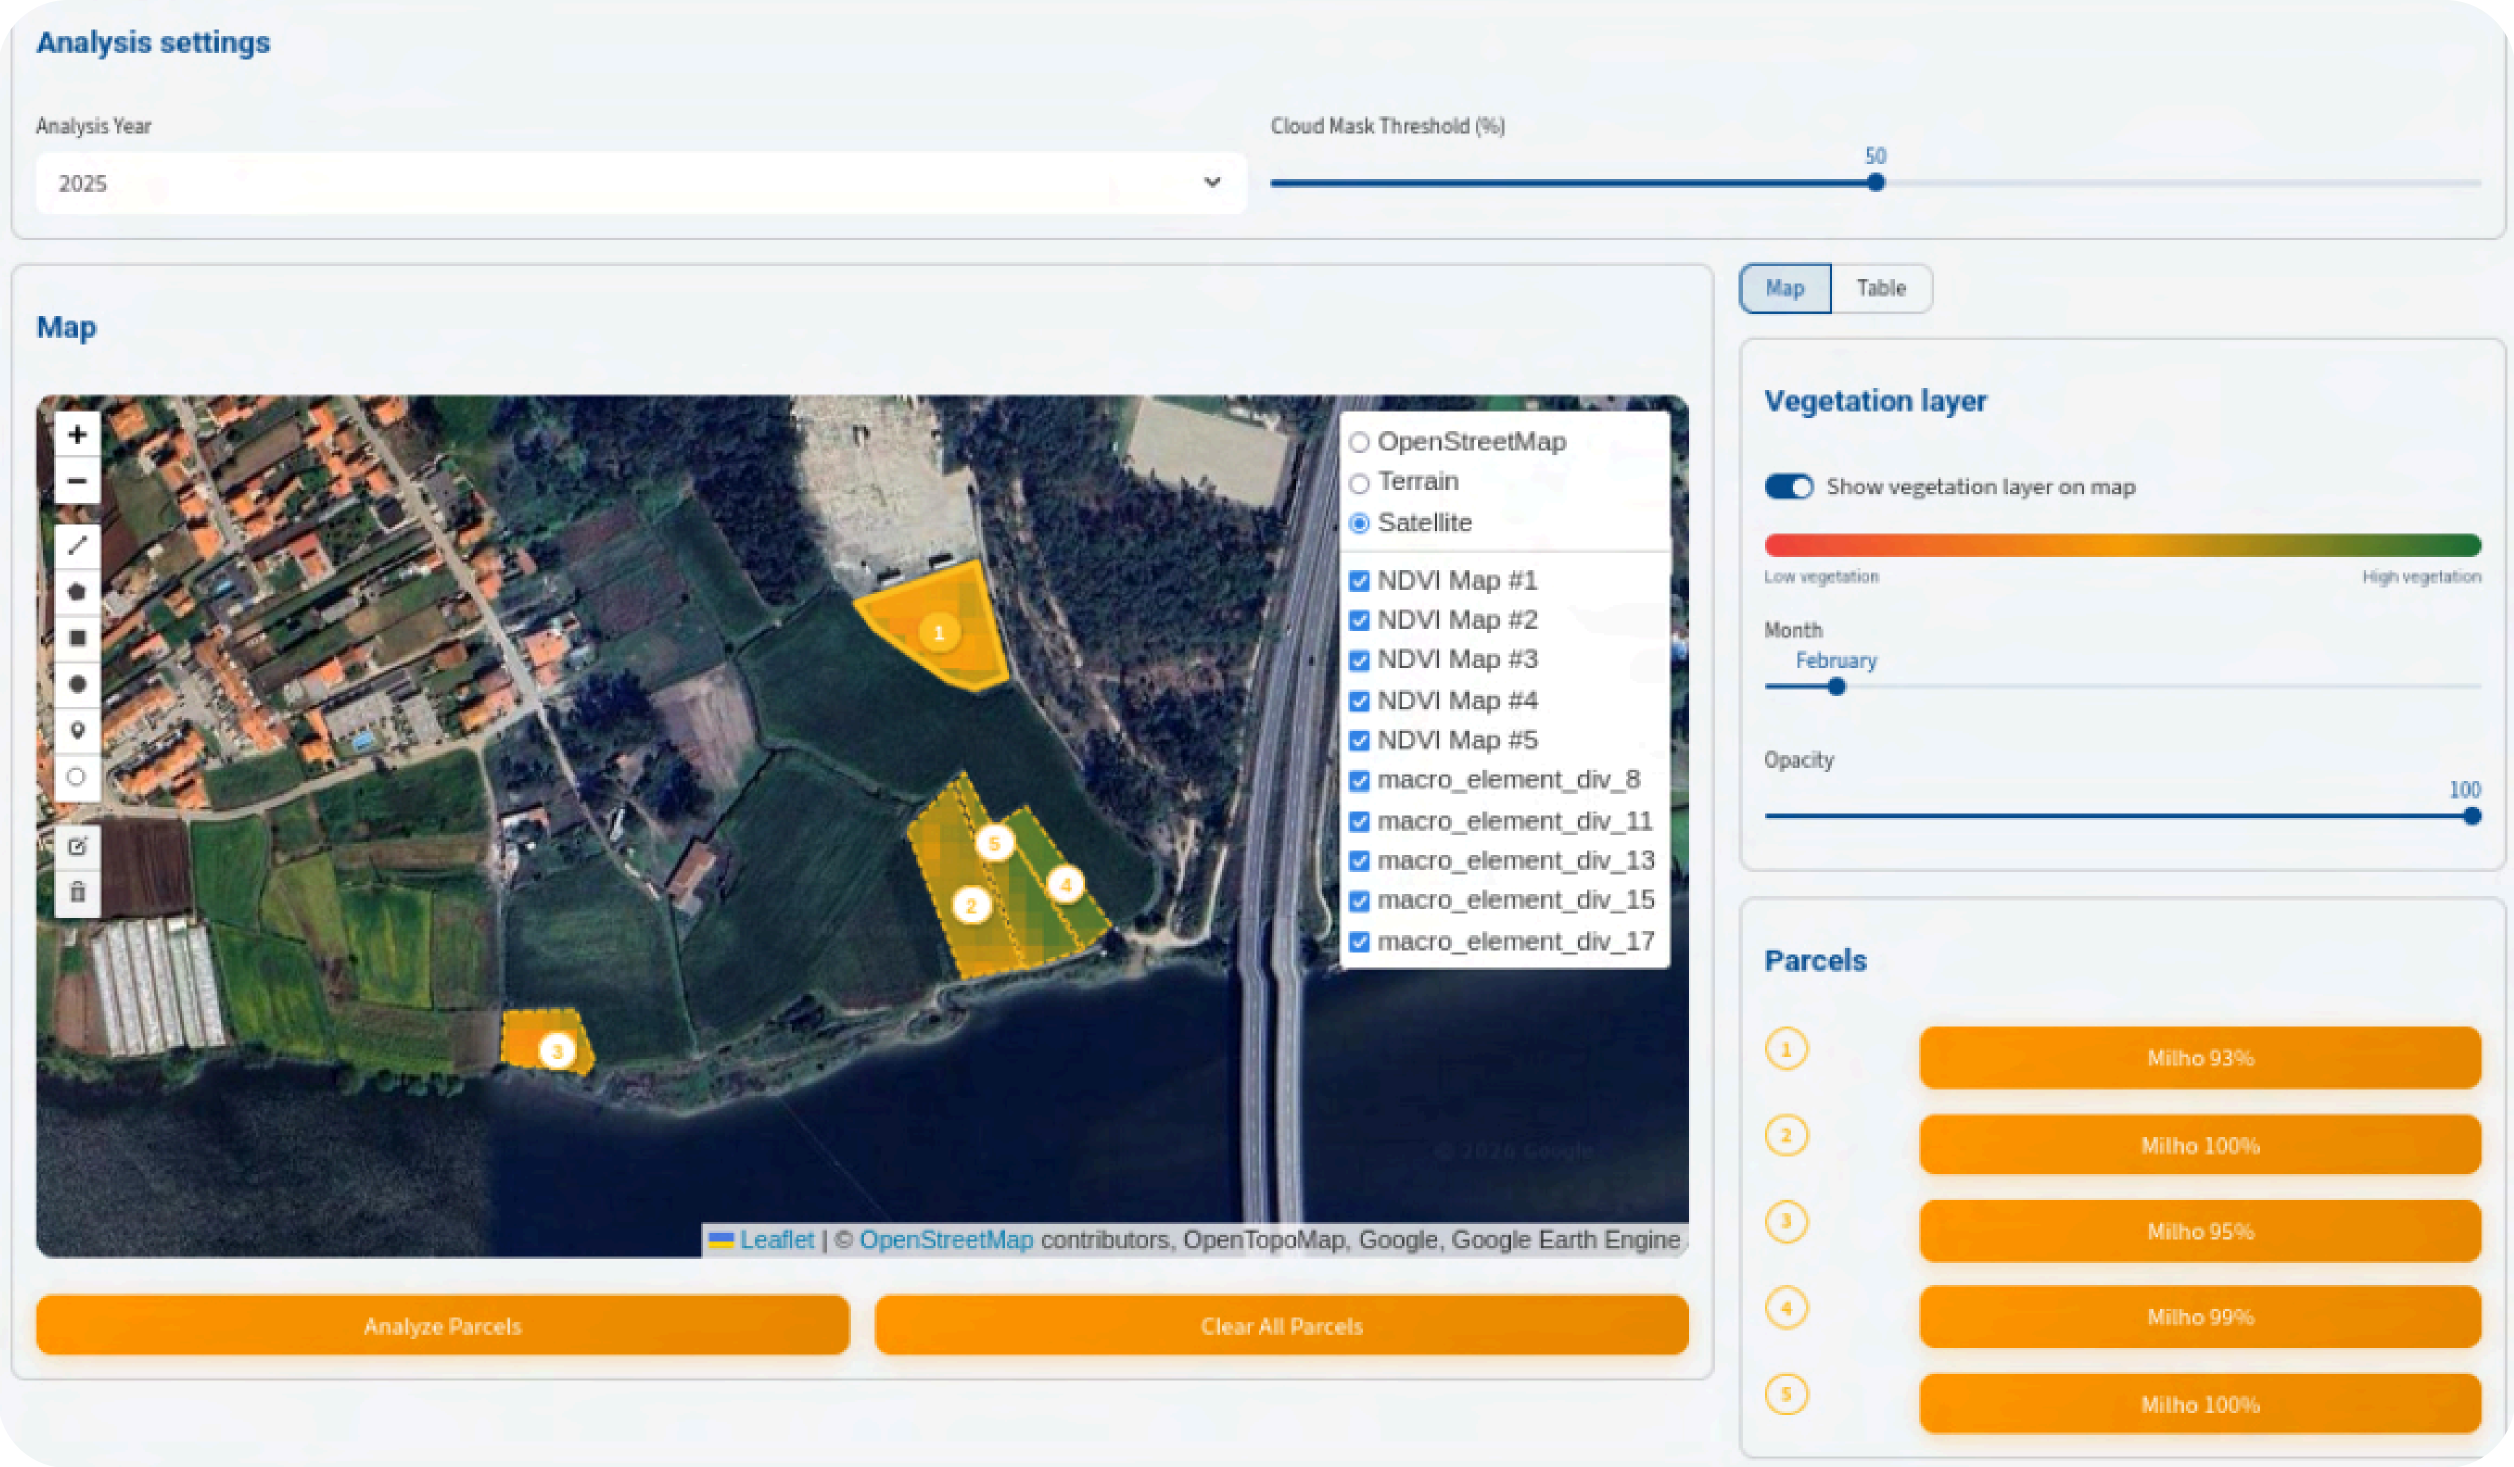

BENEFICIARIES

Farmers

- Easy access to the planted crops without needing to be on-site.
- Monitoring crop growth and detecting anomalies.
- Easier demonstration of compliance with agriculture programs

Insurance Companies & Financial Institutions

- Simpler, faster and cheaper audits
- Easy access to crop growth data
- Easier fraud detection

CONTEMPORARY PEERS

ESA

The European Space Agency provides open-access satellite imagery through the Copernicus Programme, which enables scientific research in a wide range of areas, such as this one.

IFAP

IFAP's Surface Monitoring System (SVS) uses Copernicus Programme's imagery to validate compliance for government aids, which demonstrates the practical viability of satellite based agricultural verification at an institutional level.

EOS

EOS provides a crop monitoring system which manages agricultural risks, provides insights on crop health, and determines the area of each crop, with the aim of optimizing resources and farming operations.

TEAM MEMBERS

TÉCNICO
UNIVERSIDADE
DE LISBOA


109489
Filipe Ferrão


109515
Rodrigo Barreiros


110060
Tiago Rei


110097
Maria Henriques


110164
David Freire


110234
Gonçalo Martins

OUR PARTNERS

Besquare

VirtuaCrop

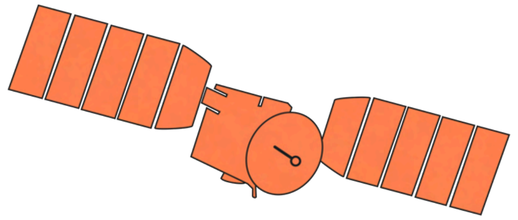

TESTING AND VALIDATION METRICS

Precision

How many predicted crops were correct.

Recall

How many real crop cases were detected.

F1-Score

Balanced score between precision and recall.

Support

Number of test samples available for each crop class.

Brier-Score

Measures the quality of predicted probabilities for each crop.

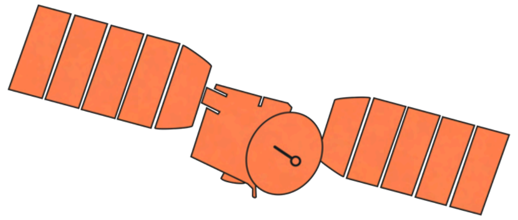

MODEL VALIDATION

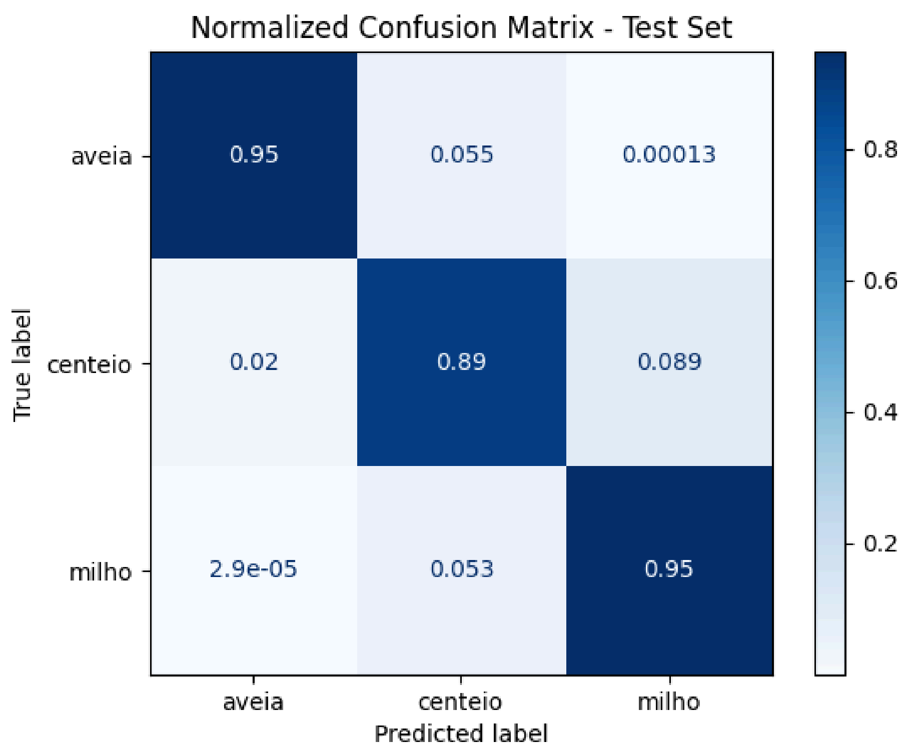

	precision	recall	fl-score	support
aveia	0.99	0.95	0.96	100992
centeio	0.86	0.89	0.87	72780
milho	0.94	0.95	0.94	103200
accuracy			0.93	276972
macro avg	0.93	0.93	0.93	276972
weighted avg	0.93	0.93	0.93	276972

Metrics

- The model achieved strong overall performance, with 0.93 accuracy out of a desired 0.96, as recommended to us by Besquare, a company with experience in the area. and macro F1-score.
- Model training was performed on a reduced dataset due to computational constraints, so performance is expected to improve with full-scale training.
- Oat and corn obtained the strongest classification results, while rye remains the most challenging crop.

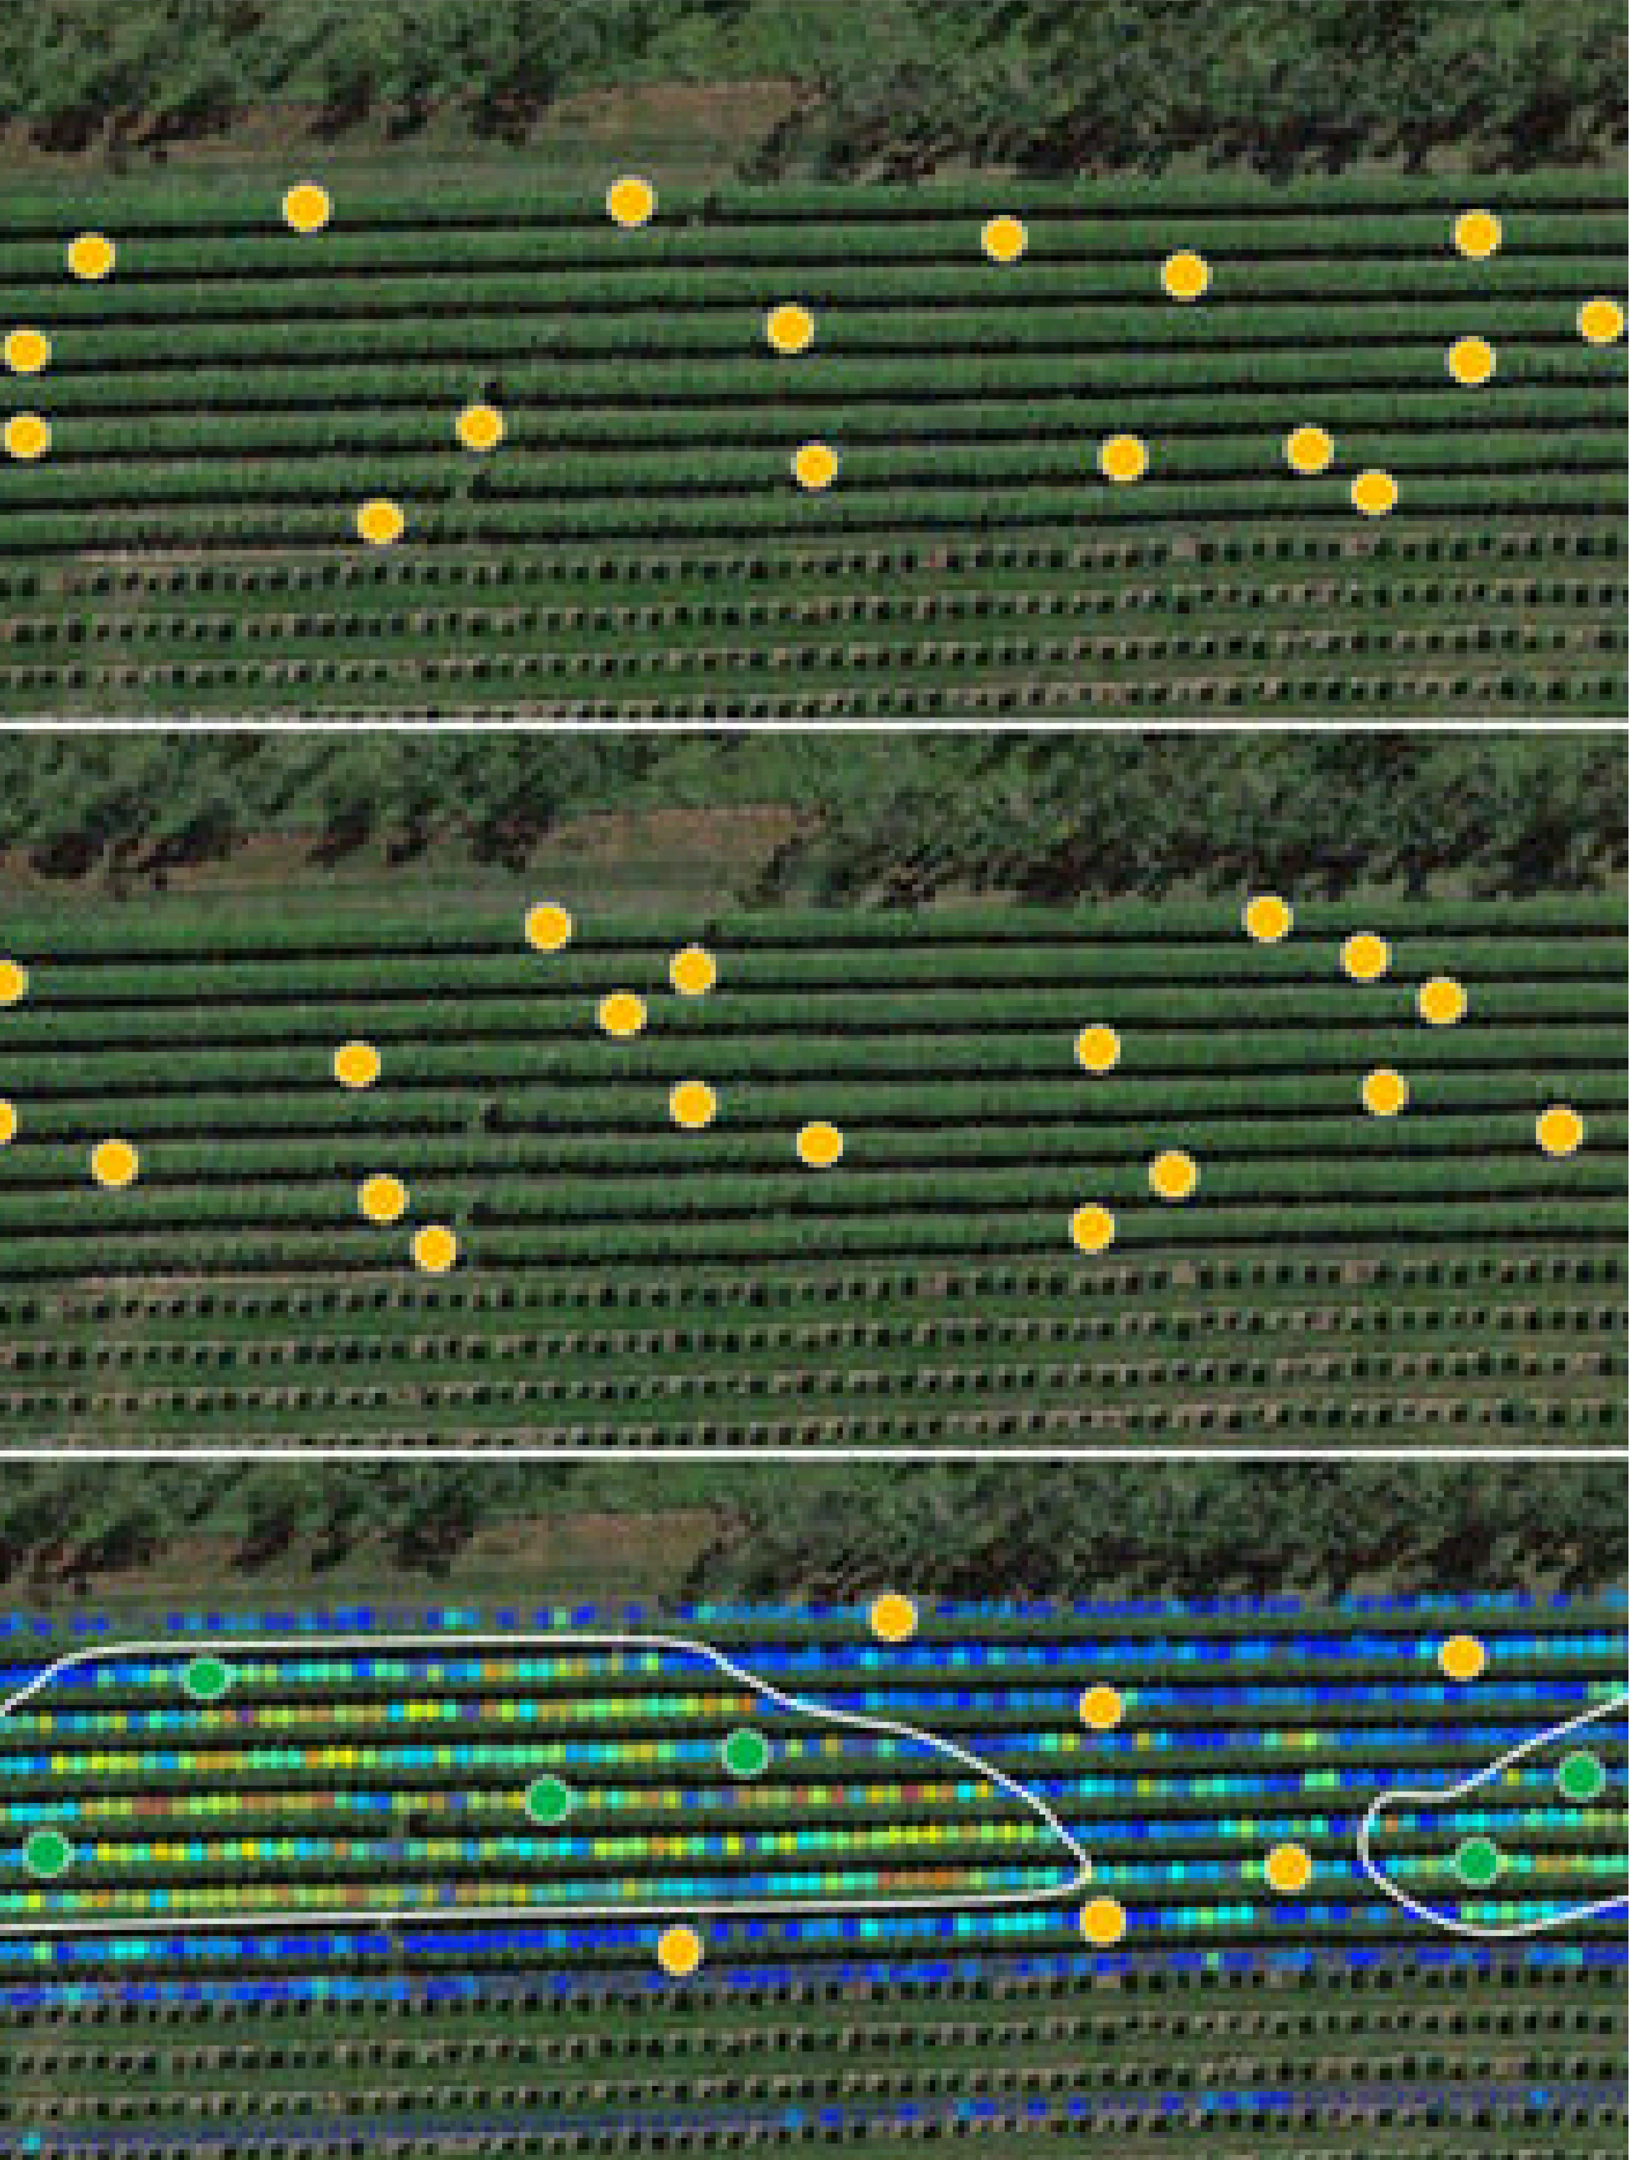

INTERVIEW ANALYSIS

VirtuaCrop

Although we were not successful in gathering answers from all the entities we would have desired, the key takeaways of the kind companies, instituitions and people behind them who did, proved to be instrumental to turn the sight of our project into another direction.

The first major contribution came for our interview with Dr. Tiago Morais, Co. Founder VirtuaCrop, switching our attention to satellites as the main vessel for gathering data, which can be far more efficiently employed and automated than a drone would suggest. The suggestion of what AI models to utilize and the concept NDVI actually were key takeaways from this interview.

INTERVIEW ANALYSIS

SmartFarmColab

Through our interview with Smart Farm Colab, we gained further insight into the harships of the satellite approach. including weather conditions and local differences in plant growth and diversity.

Luis Martins, a member of Food4Sustainability once more revealed to us the usefulness of NDVI, and how best wield it to gain correct information, analyzing leaf density with such technology.

Issac, a owner of large swaths of land in Brasil was able to give feedback in the value proposition of our project, as well as show further points we could take in the current auditing system, such as the lack of manpower and prediction of threats to the plantation.

INTERVIEW ANALYSIS

IFAP

Sr. Manuel Marques, representing IFAP, was a later interview, but one that reassured us that we were on the right track. Already following similar principals with satellite usage, they gave us access to their database and insider knowledge to be able to accelerate our work.

Our interview with the Sintra City Council was an important step in understanding how our project could be applied in real-world contexts. They provided valuable insights into local environmental monitoring practices and highlighted potential use cases for our solution. This interaction helped us validate the relevance of our approach.

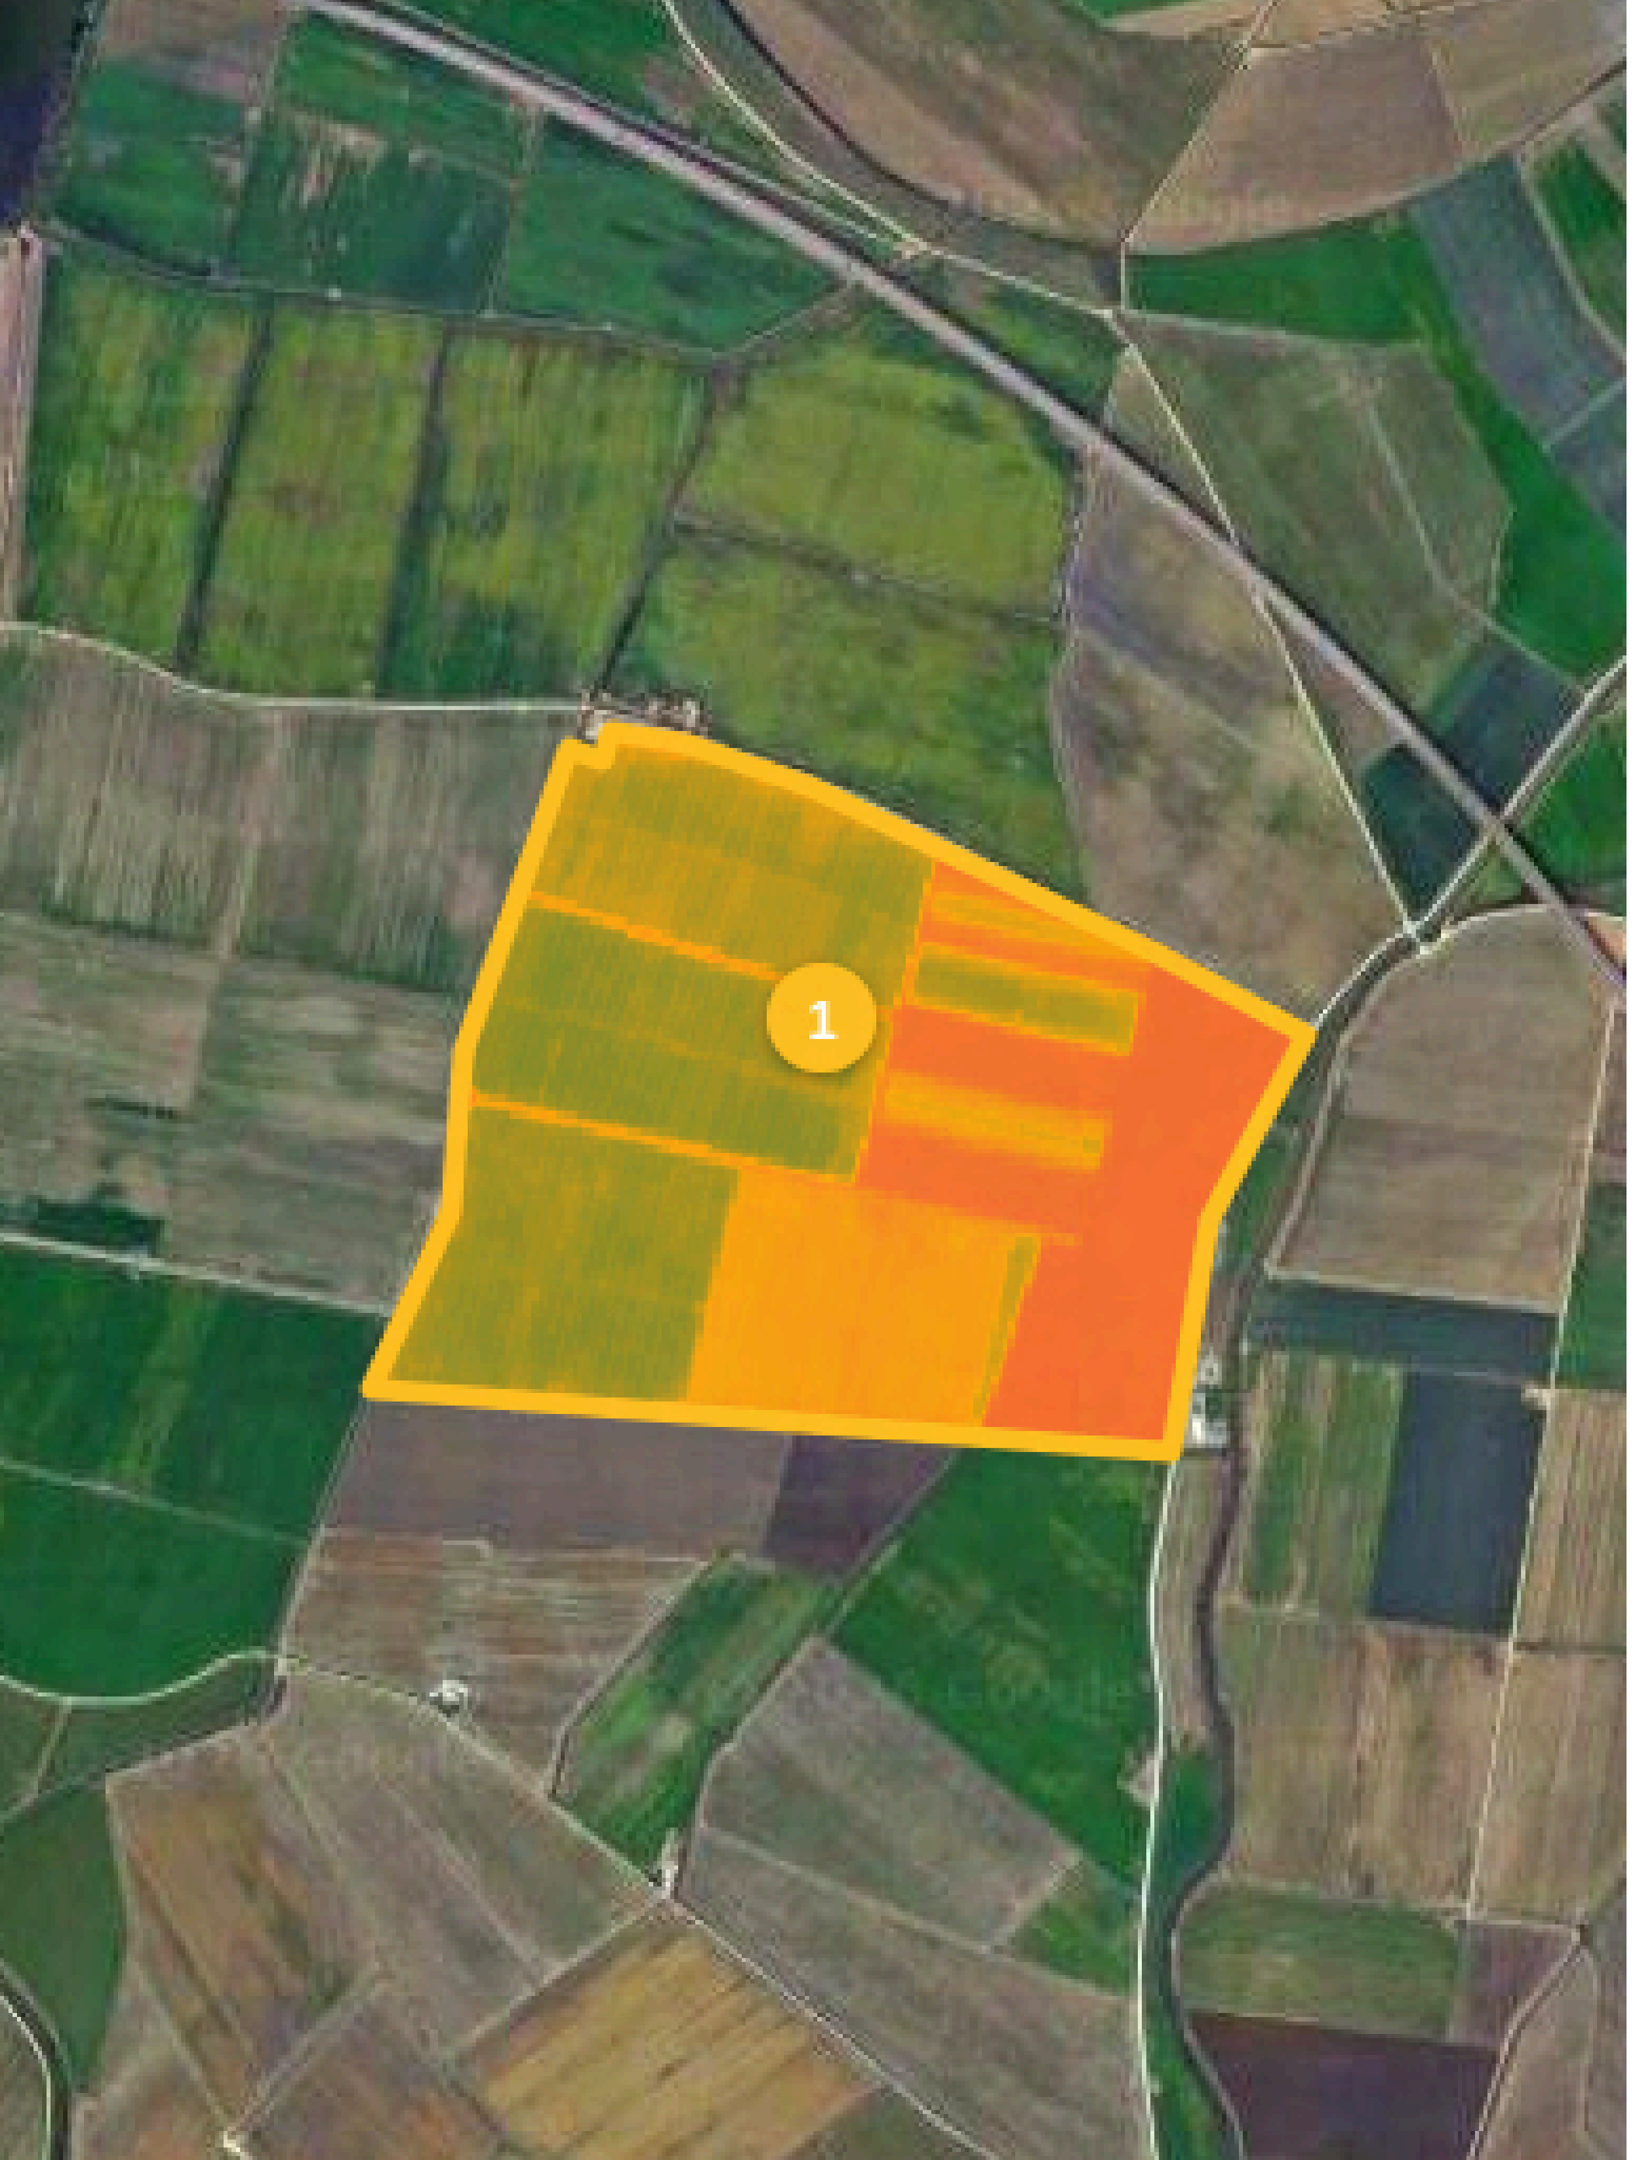

INTERVIEW ANALYSIS

Food4Sustainability

In a second round of meetings, we presented our progress to be able to tune the finer details of our algorithm.

F4S once more offered their key insight in this field, nudging us to work with colors instead of numbers to improve acessibility for all users and suggested the analysis of temperature, Irrigation, as well as some elements such as: “nitrogen, phosporous and potassium” as another possible focus of analysis in the future.

IFAP showed their interest in maintaining contact with our group in the forseeable future, given the recent developments of our project, including but not only reserved to the crop identification algorithm, but also the integration into the webapp.

DATABASE AND ML MODEL

Databse: EuroCropsV2

Source: European Commission

Technology: DuckDB

This database was chosen because it provides the locations of declared plots in various European countries over several years, including Portugal, with information on the type of crop. For our study, we chose corn, rye, and oats because of their abundance and because they are not fruit trees, in which it is more difficult to detect variations in fruit NDVI, NDRE, SAVI, EVI and LSWI.

Machine Learning Model: LightGBM

We chose this Machine Learning model based on a recommendation from our mentor, Tiago Morais.

Year	Corn	Rye	Oats
2017	179 362	44 957	27 295
2018	107 397	47 020	14 950
2019	167 708	48 642	25 352
2020	162 106	52 268	23 976
2021	157 133	52 050	23 047
2022	153 165	57 149	23 539
2023	140 835	46 982	21 380
Total	1 067 706	349 068	159 539

Declared parcels in Portugal by year

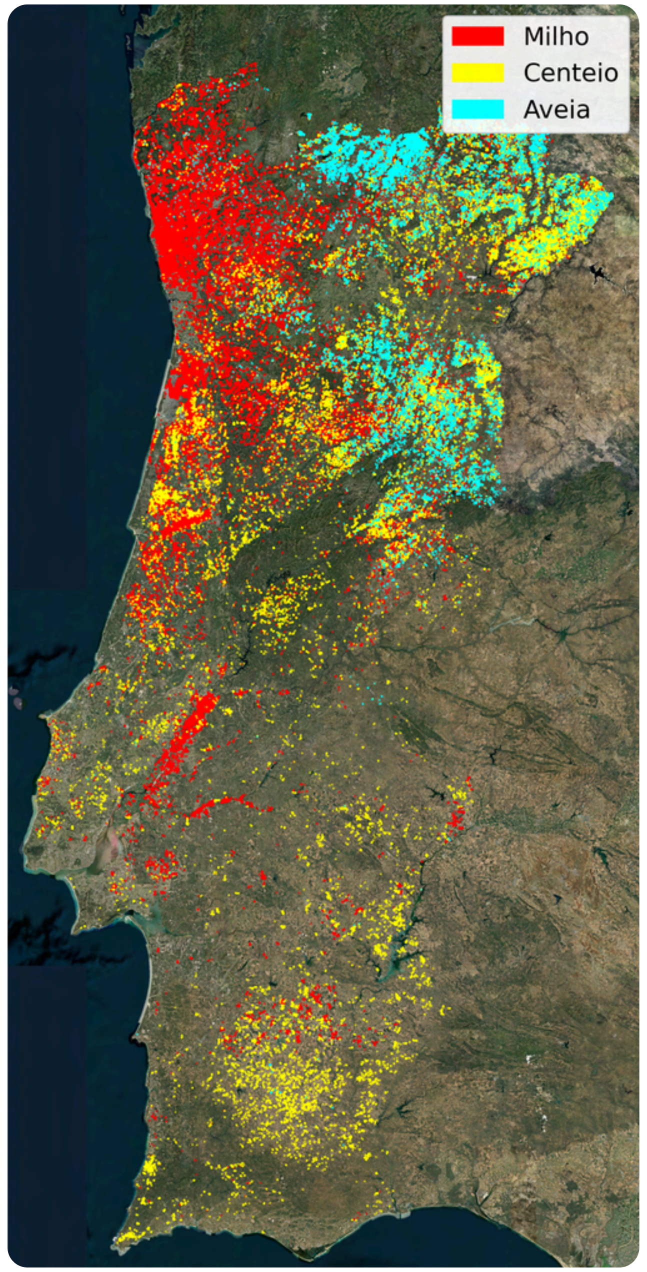

Declared parcels in Portugal in 2023

SOLUTION DIAGRAM

TÉCNICO UNIVERSIDADE DE LISBOA

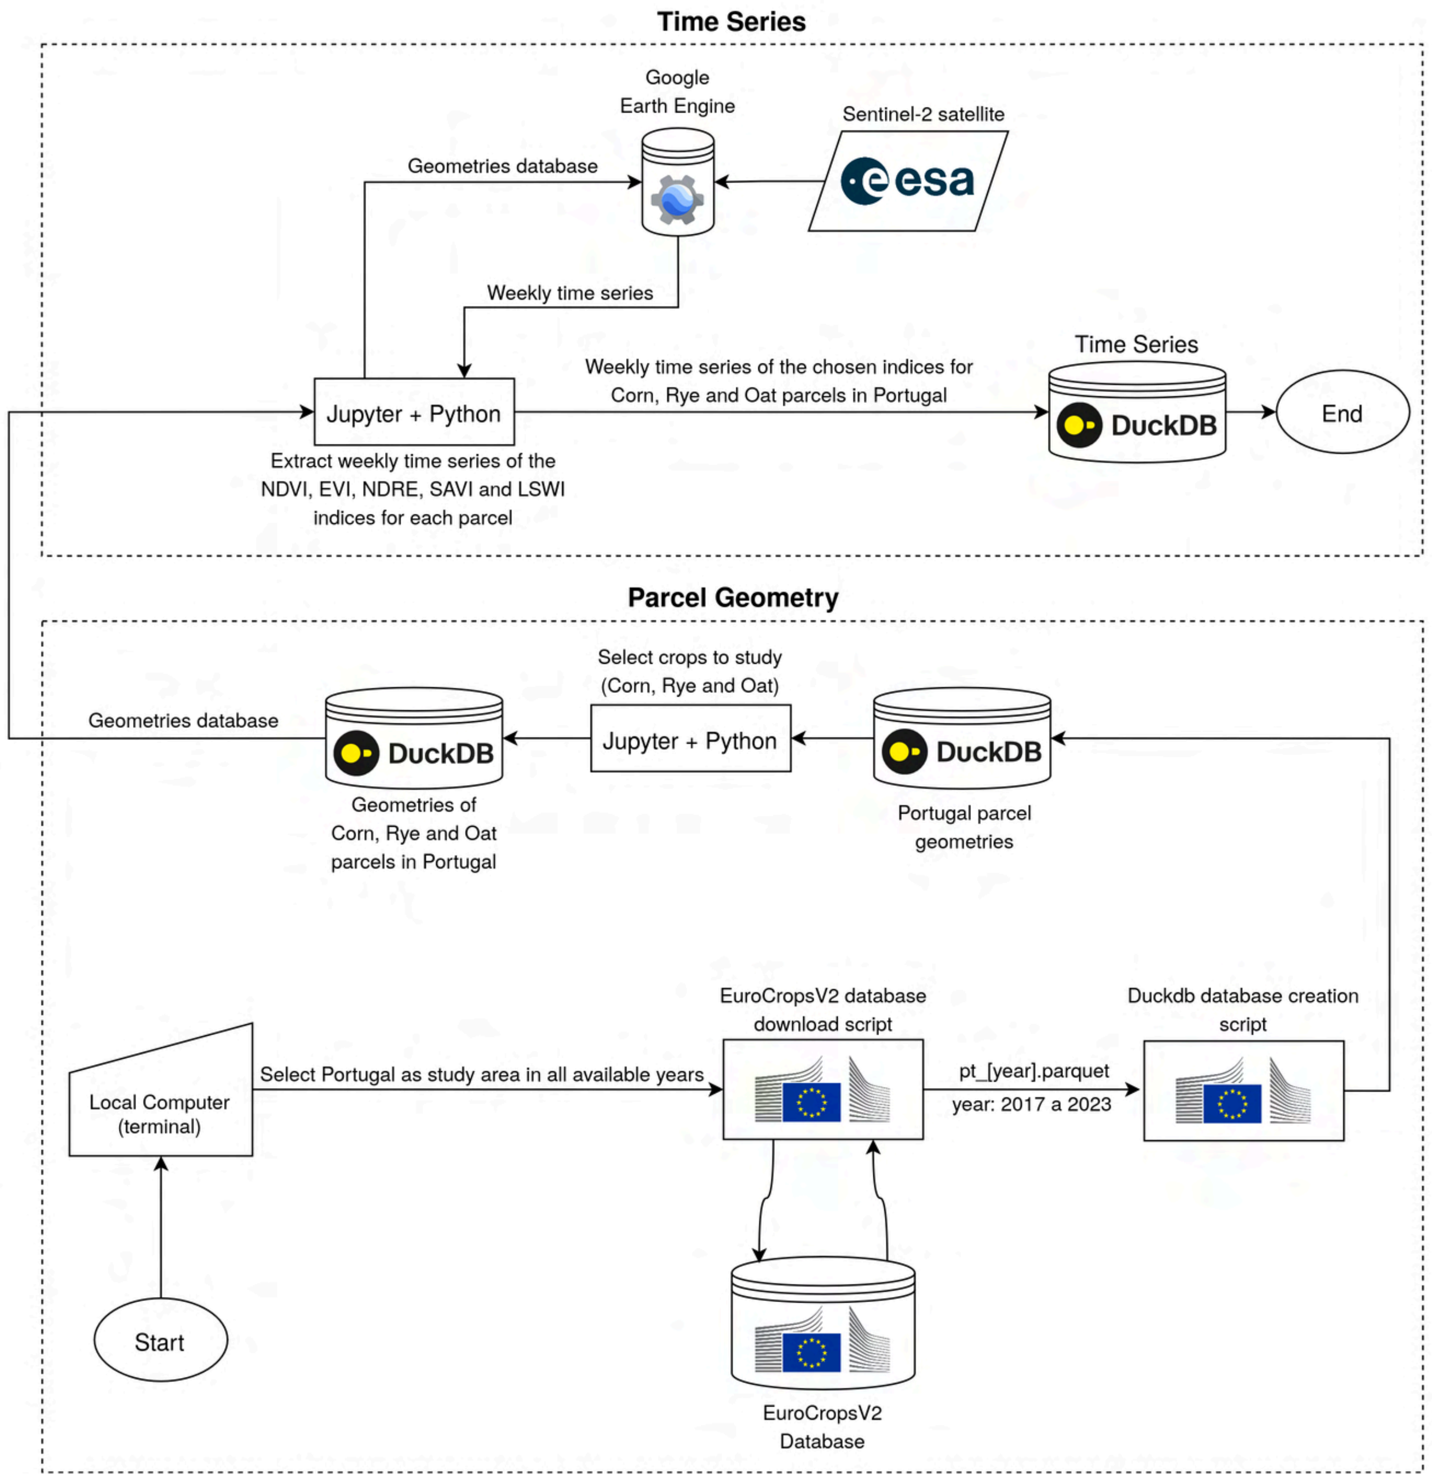

Data Collection

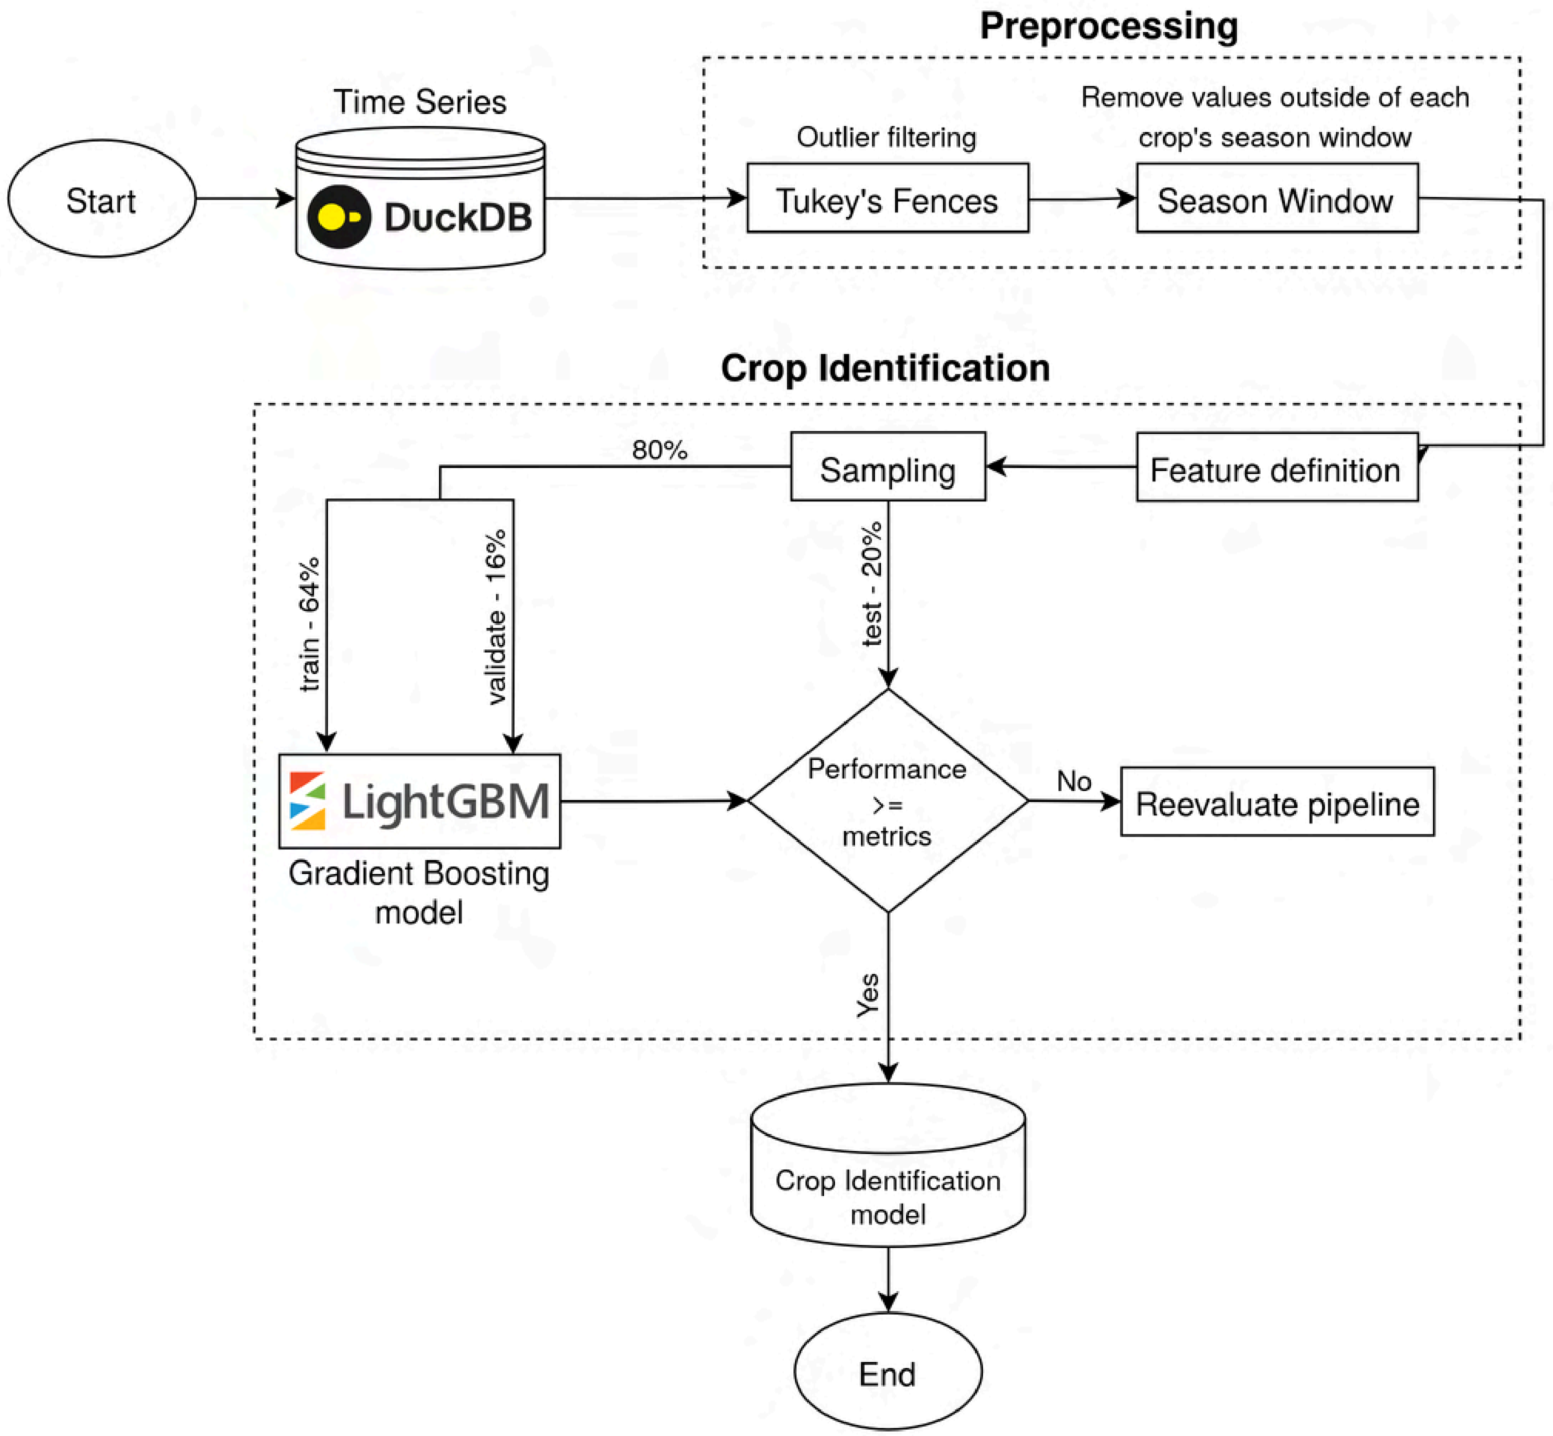

Crop Identification model

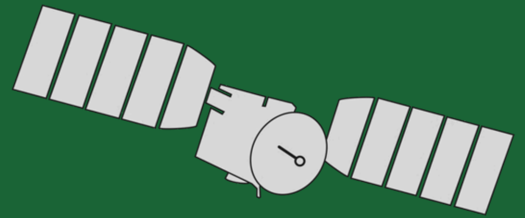

SOLUTION DIAGRAM

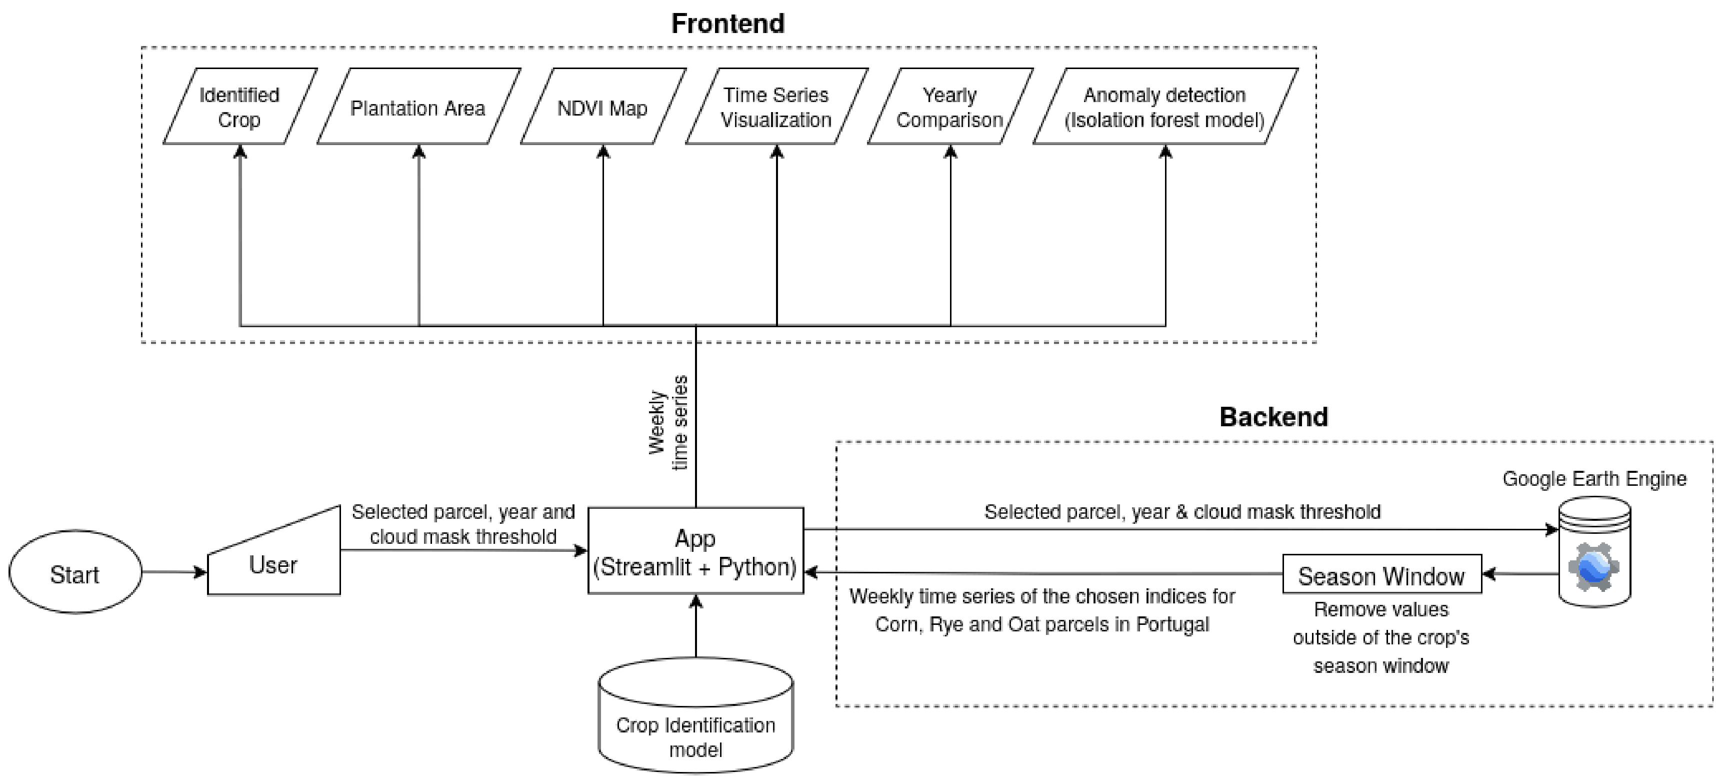

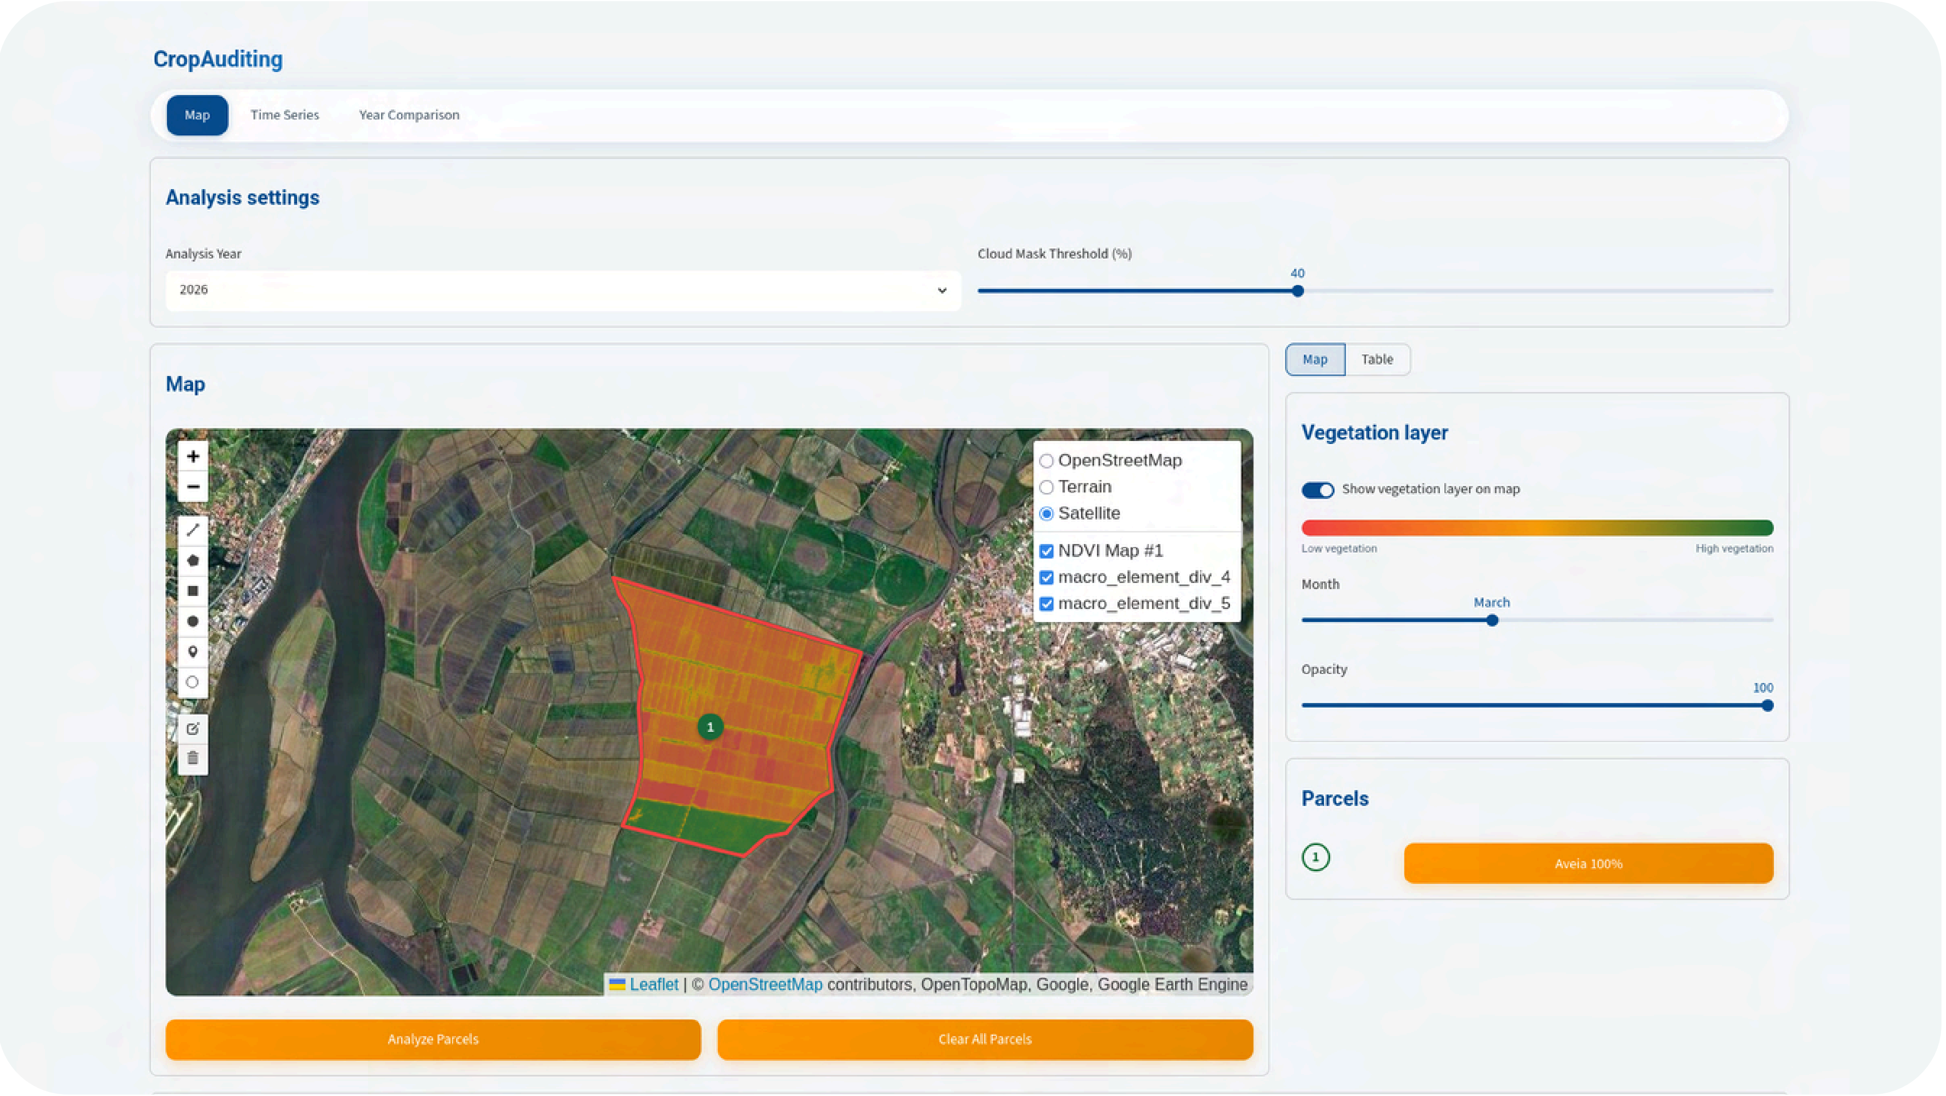

Web Application

Using an Application, we display a user-interface which analyses a selected agricultural field.

TASKS

Data Collection

Crop vs Terrain Identification
Crop Species Identification
Growth Stage Identification

Web App
Website
Blog

Interviews
Pitch Deck
Poster
Video

TEAM MEMBERS' CONTRIBUTIONS

TÉCNICO
UNIVERSIDADE DE LISBOA

Filipe Ferrão
- Growth Stage Identification*
- Crop vs Terrain Identification
- Web App
- Website*
- Poster

Rodrigo Barreiros
- Crop vs Terrain Identification*
- Crop Species Identification
- Web App*
- Pitch Deck
- Communication

Tiago Rei
- Crop Species Identification*
- Data Collection
- Growth Stage Identification
- Communication*
- Blog

Maria Henriques
- Crop Species Identification
- Data Collection
- Web App
- Blog* / Poster*
- Video

Dyanne Freire
- Growth Stage Identification
- Data Collection
- Video*
- Pitch Deck*
- Interviews

Gonçalo Martins
- Data Collection*
- Growth Stage Identification
- Interviews*
- Pitch Deck
- Website

*Responsible

COSTS & BENEFITS

- Utility: Our solution is useful and scalable to other problems such as identifying invasive plants or nutrients present in the soil.
- Information Clarity: The information in the web app is clear and easy to interpret and can be used by anyone.
- Cost and Time Reduction: It allows for a reduction in the costs associated with agricultural audits and lowers their overall cost.
- Our beneficiaries would consider using our solution in their daily work.

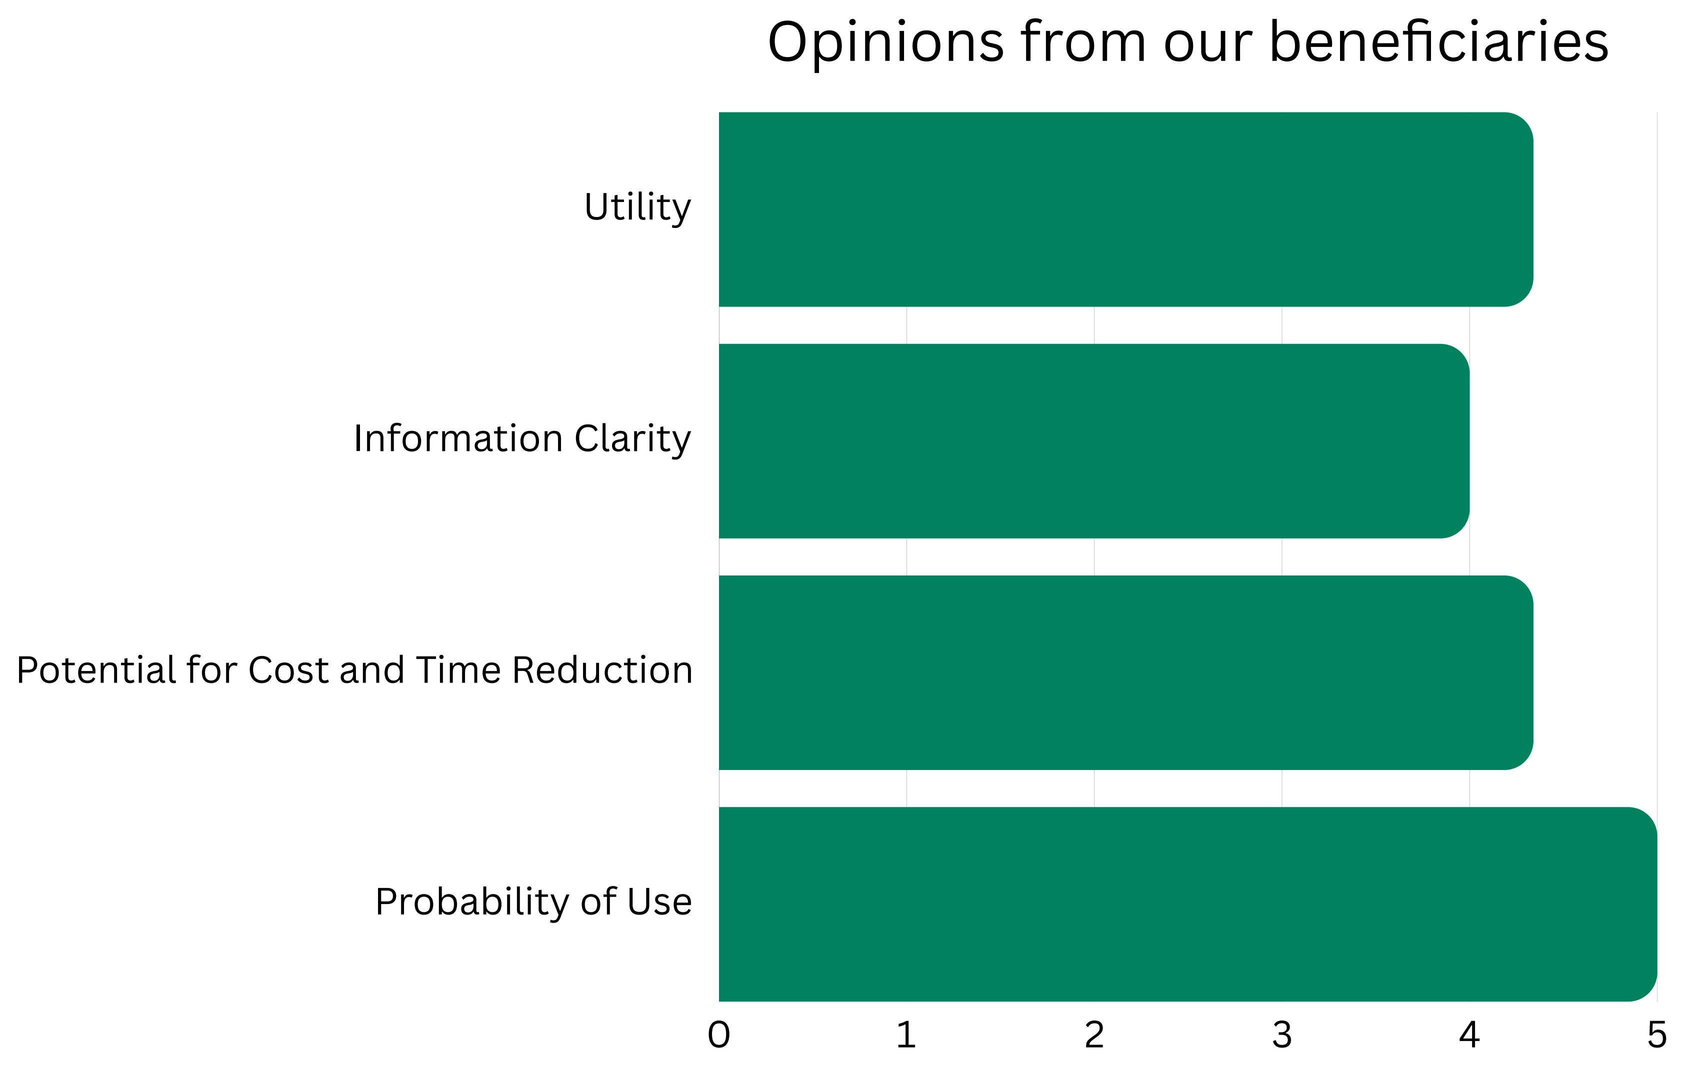

Thank You!

TÉCNICO
UNIVERSIDADE
DE LISBOA