Integrating Farmer Declarations with Satellite-Derived Crops Classifications: A Geospatial Framework for Supporting Common Agricultural Policy Implementation

Michael Assouline,
Philippe Loudjani, Philippe Loudjani, Daniele Borio, Gilbert Voican
Joint Research Centre (JRC), European Commission, Ispra, Italy
michael.assouline@ec.europa.eu
philippe.loudjani@ec.europa.eu
daniele.borio@ec.europa.eu
gilbert-madalin.voican@ec.europa.eu

Abstract[1]

Agricultural subsidy payments under the EU's Common Agricultural Policy (CAP) rely on farmer declarations submitted through the Integrated Administration and Control System (IACS), recorded in national Geo-Spatial Applications (GSA). Traditional field inspections verify only a small fraction of parcels due to resource constraints, creating a critical challenge: how to validate hundreds of thousands of declarations efficiently at scale? Ensuring declaration accuracy is essential for policy compliance, budget integrity, and reliable agricultural statistics. Systematic comparison of farmer declarations against independent satellite observations can identify potential errors, improve data quality, and enable risk-based inspection strategies that reduce verification costs while maintaining monitoring effectiveness. Understanding when and why discrepancies occur is crucial for both payment agencies conducting controls and remote sensing scientists improving classification algorithms. We developed within the Joint Research Centre (JRC) - Land Unit Characterization for Policies (LUChaP), a web-based validation tool comparing GSA farmer declarations against the Copernicus High Resolution Layer (HRL), an automated satellite-derived crop classification at 10-meter resolution. Applied to two case study regions (Lombardy, Italy and Burgundy, France) covering over 600,000 parcels for the 2021 - 2023 campaign, the platform integrates spatial analysis (parcel size, classification purity) with temporal validation using Sentinel-2 NDVI profiles extracted from the Copernicus Data Space Ecosystem (CDSE). The Flask-based architecture employs GeoPandas for efficient spatial queries and Rasterio for on-demand pixel extraction, enabling interactive exploration of regional and parcel-level discrepancies. Analysis reveals overall agreement between farmer declarations and satellite classifications while we noticed systematic mismatch patterns: certain crop pairs show consistent confusion (e.g., soybeans and maize), small parcels exhibit edge effects from neighboring fields, and classification purity strongly correlates with mismatch confidence. NDVI temporal profiles successfully distinguish genuine farmer errors (parcels following different crop phenology) from satellite technical limitations (small parcels with mixed pixels or boundary effects). This work demonstrates operational workflows for payment agencies to prioritize high-risk parcels for inspection, provides feedback to improve satellite classification algorithms, and offers reproducible validation tools deployable across EU member states. The platform is operational in the JRC's BDAP environment, enabling auditors

[1] We hereby confirm that text and materials are free from any copyright violations

and payment agencies to efficiently identify parcels requiring closer verification, supporting both CAP monitoring and agricultural statistics quality assurance.

Keywords: Agricultural monitoring; Satellite validation; Common Agricultural Policy (CAP); IACS controls; Geo-Spatial Application (GSA); Copernicus High Resolution Layer; Sentinel-2 NDVI; Remote sensing; Data integration.

1. Introduction

The European Union's Common Agricultural Policy (CAP) relies on the Integrated Administration and Control System (IACS) to ensure accurate subsidy payments and monitor agricultural land use across member states. Three core components constitute IACS: the Area Monitoring System (AMS), the Geo-Spatial Application (GSA) containing farmer declarations, and the Land Parcel Identification System (LPIS). Quality assurance of these interconnected systems is critical for effective policy implementation, financial transparency, and reliable agricultural statistics.

Traditional IACS verification relies primarily on field inspections, which cover only a small fraction of declared parcels due to cost and logistical constraints. With hundreds of thousands of declarations submitted annually across member states, systematic validation at scale remains a persistent challenge. Earth observation data, particularly the Copernicus High Resolution Layer (HRL) providing automated crop classifications at 10-meter resolution, offers a complementary validation approach. However, operational deployment requires addressing a fundamental question: when farmer declarations and satellite classifications disagree, how can we distinguish genuine declaration errors from satellite technical limitations or legitimate temporal land use changes?

The Land Unit Characterization for Policies (LUChaP) project at the JRC performs Research and Development (R&D) activities towards the characterization of land units and land practices. It develops land monitoring approaches mainly in support of the Common Agricultural Policy (CAP). Through the geo spatial application developed and presented here it addresses this challenge through systematic comparison of GSA declarations against HRL classifications. Rather than treating discrepancies as simple errors, the intent is to integrate spatial analysis, temporal vegetation indices, and interactive visualization to categorize mismatches by cause: high-confidence declaration errors requiring inspection, satellite classification limitations affecting small or boundary parcels, or ambiguous cases requiring expert judgment.

This paper presents validation workflows applied to two case study regions: Lombardy, Italy and Burgundy, France, covering over 600,000 parcels from the 2021 to 2023 campaign. We demonstrate how combining GSA-HRL comparison with Sentinel-2 NDVI temporal profiles enables payment agencies to prioritize high-risk parcels efficiently while providing feedback to improve satellite classification algorithms. The approach supports both CAP monitoring objectives and broader agricultural statistics quality assurance, with deployment-ready tools operational in the Joint Research Centre's BDAP environment.

2. Methodology and Technical Implementation

The platform implements a web-based geospatial validation framework enabling systematic comparison of GSA farmer declarations against Copernicus High Resolution Layer (HRL) satellite-derived crop classifications. The architecture addresses three core requirements: efficient handling of large geospatial datasets (600,000+ parcels, multi-gigabyte rasters), interactive exploration at regional and parcel scales, and on-demand validation metrics without pre-computation overhead.

The backend employs Python with Flask for RESTful web services, GeoPandas for vector operations with spatial indexing (R-tree), and Rasterio for on-the-fly raster pixel extraction. Dynamic tile serving with LRU caching (500 tiles in memory) enables responsive map interaction without pre-generating tile pyramids. The frontend implements Leaflet.js for interactive mapping, with viewport-based analysis triggered at zoom level 14 or higher to balance computational load with spatial detail. All code is version-controlled in the JRC's BDAP GitLab environment, ensuring reproducibility and collaborative development.

GSA declarations and HRL classifications employ different crop taxonomies requiring systematic harmonization. All GSA crop codes were mapped to HRL classes through iterative refinement: obvious matches (e.g., wheat, maize), regional specialties (wine grapes in Burgundy), and multi-year validation using 2021-2023 data to identify stable mappings. Crucially, non-crop features (grasslands, fallow, buffer strips) were explicitly mapped to NULL to avoid spurious mismatches, as HRL targets only arable and permanent crops (Fig.1).

Code	Land Cover	Crop Group	Crop type
1110	Arable Crops	Cereals	Wheat
1120			Barley
1130			Maize
1140			Rice
1150			Other cereals
1210		Dry pulses & Vegetables	Fresh Vegetables
1220			Dry pulses
1310		Root/tuber crops	Potatoes
1320			Sugar Beet
1410		Non-permanent industrial crops	Sunflower
1420			Soybeans
1430			Rapeseed
1440			Flax, cotton and hemp
2100	Permanent Crops	Permanent Crops	Grapes
2200			Olives
2310			Fruits
2320			Nuts
3100	Arable Crops	Unclassified arable crop	Unclassified arable crop
3200	Permanent Crops	Unclassified permanent crop	Unclassified permanent crop

Fig.1 HRL crop types nomenclature and aggregation levels [1]

Vector data (GPKG format, EPSG:3035) stores parcel geometries, farmer declarations, and harmonized HRL codes. Raster data (GeoTIFF, EPSG:3035 for analysis, EPSG:4326 for web display) contains HRL classifications. Dual coordinate reference system handling ensures geometric accuracy during raster operations while maintaining web standard compatibility for frontend visualization.

The analytical workflow operates in three stages:
Stage 1: Spatial Overlay. For each parcel in the current viewport, Rasterio extracts all HRL pixel values intersecting the parcel boundary using the parcel's original CRS geometry. This produces pixel-level paired observations: farmer-declared crop type (vector attribute) versus satellite-classified crop type (raster value).
Stage 2: Purity and Agreement. For each parcel, classification purity is computed as the percentage of pixels matching the dominant HRL class. Parcels with purity <60% indicate mixed classifications or boundary effects. Agreement is binary: parcel classified as "match" if dominant HRL class equals declared crop type, otherwise "mismatch."
Stage 3: Confusion Matrix. Viewport-level confusion matrices aggregate all parcel pixels, cross-tabulating GSA declared crop types versus HRL classified crop types. For each declared crop, we calculate what percentage of those pixels HRL correctly identified (e.g., if farmers declared 10,000 wheat pixels, what percentage did HRL classify as wheat?). Conversely, for each HRL classification, we calculate what percentage matches the GSA declaration (e.g., of all pixels HRL classified as wheat, what percentage were actually declared as wheat by farmers?). Overall accuracy represents the percentage of pixels where GSA declaration and HRL classification agree. This pixel-based approach (rather than counting parcels) accounts for parcel size variation and provides standard remote sensing accuracy metrics comparable across studies.

3. Case Study 1: Emilia-Romagna, Italy

The analysis focuses on Ferrara province in Emilia-Romagna, northern Italy, encompassing 219,603 agricultural parcels with pixel-level comparison at 10-meter resolution.
A representative analysis of 304 agricultural parcels in central Ferrara province examined 38,968 pixels at 10-meter resolution, achieving 77% overall accuracy (29,991 correct pixels, 8,977 errors). For major cereals, the system demonstrates high agreement: of 10,388 maize pixels declared by farmers, 10,286 were correctly identified by HRL (99%); wheat shows 82.5% correct identification (1,860 pixels incorrectly assigned to barley); and rapeseed achieves 99.9%. However, soybeans show moderate accuracy at 50.9% (from a total of 9,194 pixels declared by farmers 2,637 pixels were incorrectly assigned to rapeseed and 1,117 to dry pulses) (Fig. 1) provides a screenshot of the dashboard with the results obtained for the soybeans. To illustrate diagnostic capabilities at parcel level, a specific example shows parcel #33547 declared as soybeans containing 1,245 pixels classified by HRL as 100% Rapeseeds demonstrating mismatch.

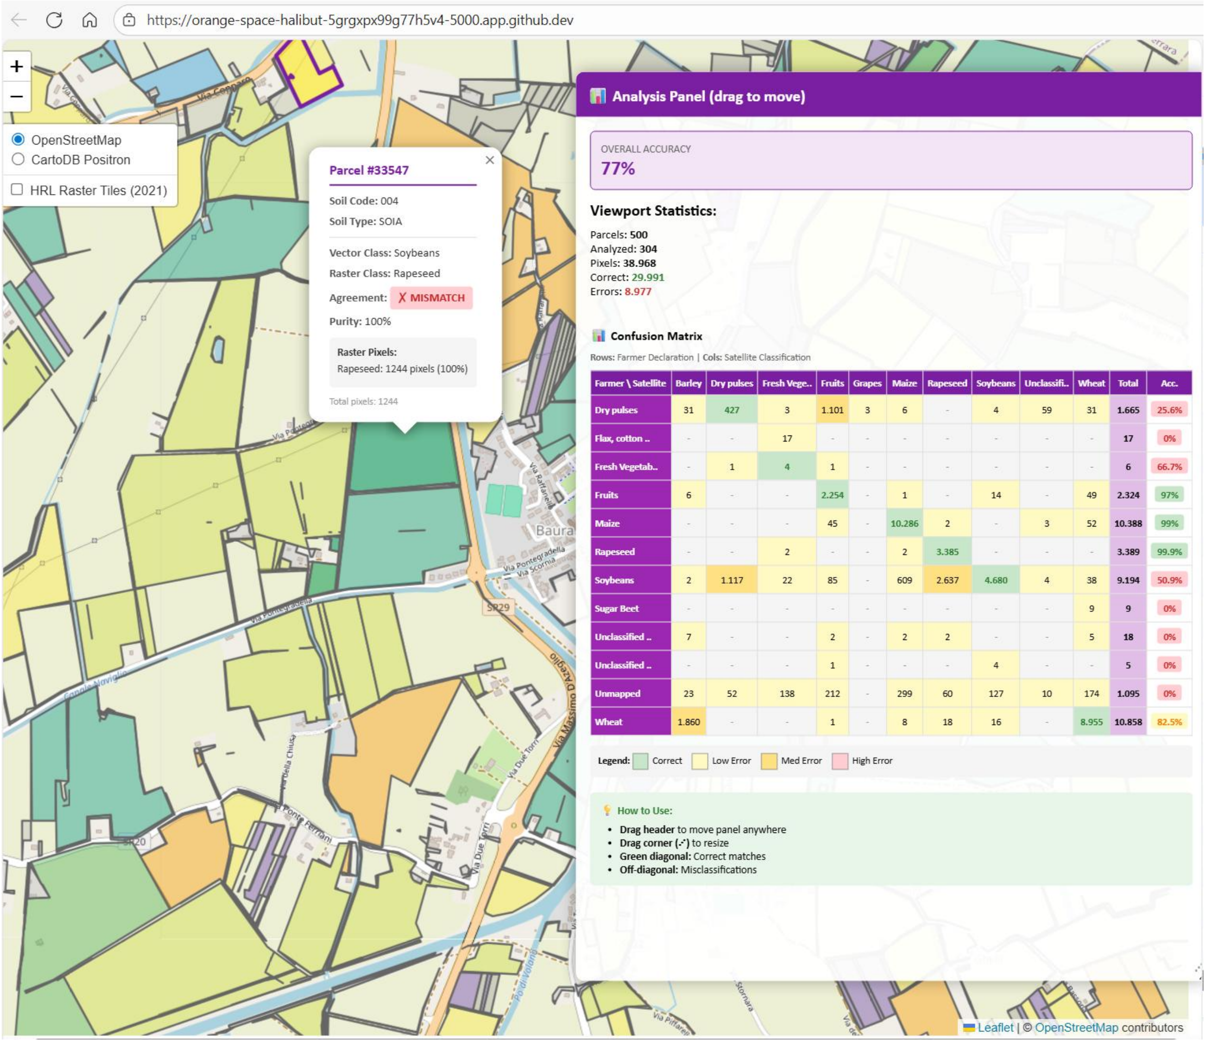

Fig.2 Confusion matrix for Emilia-Romagna Region, Italy

These patterns exemplify the application's dual purpose. For payment agencies, the system flags high-risk parcels requiring targeted inspection. For remote sensing scientists, it identifies systematic classification weaknesses requiring algorithm improvement or training data augmentation for underrepresented crop types. The interactive platform enables users to explore surrounding parcels, examine temporal stability across seasons, and assess whether discrepancies represent isolated cases or broader regional patterns.

4. Case Study 2: Burgundy, France

To demonstrate transferability across member states, the same validation approach was applied to France's Registre Parcellaire Graphique (RPG), the French equivalent of GSA. Farmer-declared crop data from the 2023 campaign were harmonized with the corresponding 2023 HRL classification following the same mapping procedure described in Section 3, adapting the taxonomy to French crop codes.

The interactive platform enables real-time accuracy assessment for any geographic area. When users navigate the map and zoom to level 14 or higher, viewport statistics are computed instantly for the visible parcels. When a user navigates the map, the platform instantly computes accuracy statistics for the visible area (Fig.3).

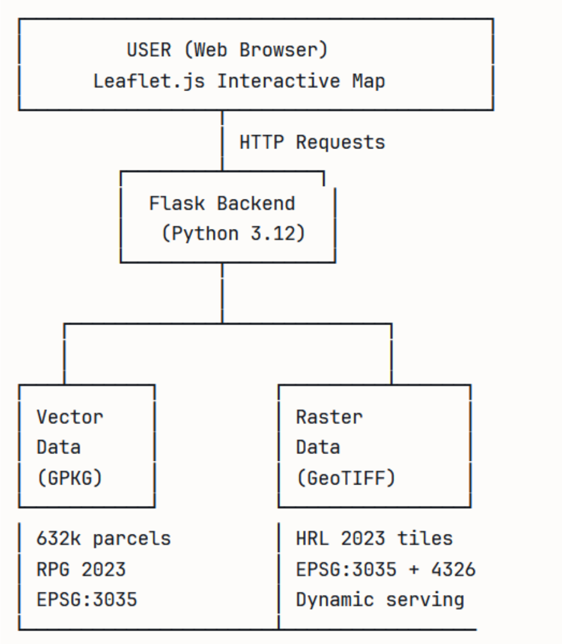

Fig.3 Architecture of the GSA validation platform

Analysis of 292 parcels in central Burgundy[2] (160,972 pixels total) reveals 90.7% overall accuracy between farmer declarations and satellite classification. This higher accuracy compared to Emilia-Romagna (77%) may reflect regional differences in crop diversity, parcel size distribution, or declaration practices. Figure 4 presents the confusion matrix for this viewport,

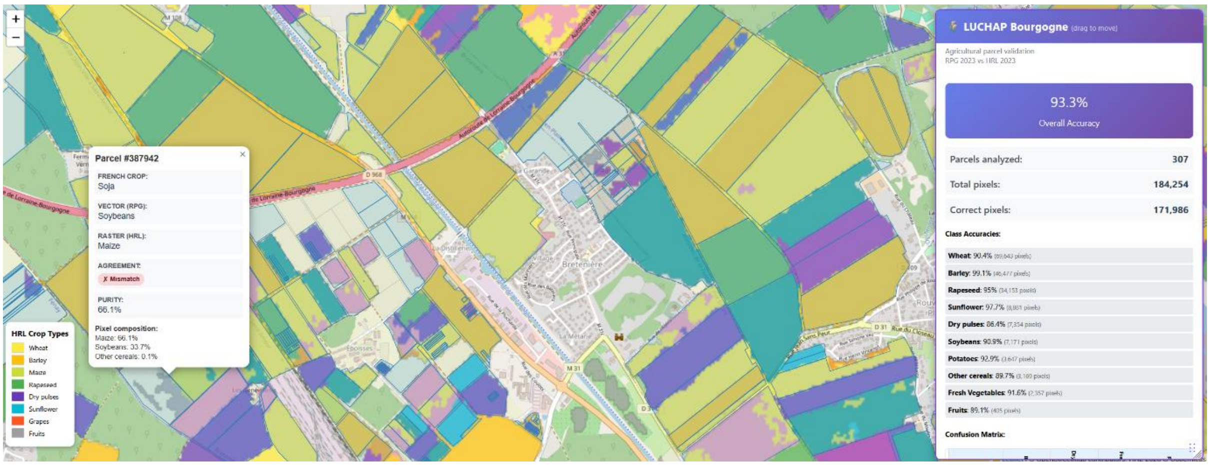

Fig.4 GSA parcel validation

When we examine individual class accuracy, some crops perform very well: sunflower at 99%, rapeseed at 95%, and barley at 98%. However, others are problematic: dry pulses reach only 47.5% and soybeans 70.7%.

---

2 bbox=5.148810876216886,47.321370564267404,5.320214761104581,47.35905994178323&zoom=14

Dry pulses: 47.5% (4,160 pixels)

Fresh Vegetables: 77% (1,768 pixels)

Other cereals: 0% (1,321 pixels)

Flax, cotton and hemp: 0% (553 pixels)

Confusion Matrix:

Vector \ Raster	Barley	Dry pulses	Fresh Ve..	Maize	Other ce..	Potatoes	Rapeseed	Soybeans	Sunflower	Unclassi..	Wheat	Total
Barley	**41,609**	21	1	29	-	9	328	7	49	36	361	**42,450**
Dry pulses	46	**1,974**	61	5	1,135	-	47	47	-	708	137	**4,160**
Flax, co..	23	-	-	6	465	-	16	-	-	-	43	**553**
Fresh Ve..	52	223	**1,362**	-	-	7	27	-	3	77	17	**1,768**
Fruits	-	-	-	-	-	-	-	-	-	30	22	**52**
Maize	37	-	-	**8,340**	-	-	44	4	66	20	336	**8,847**
Other ce..	484	-	-	183	-	-	-	-	-	591	63	**1,321**
Potatoes	58	-	-	-	-	**276**	-	-	2	52	11	**399**
Rapeseed	155	953	-	36	30	-	**31,957**	13	118	88	282	**33,632**
Soybeans	133	189	591	694	-	-	118	**6,227**	204	503	153	**8,812**
Sunflower	30	-	2	-	-	-	12	-	**12,273**	16	53	**12,386**
Unclassi..	2	-	-	-	-	-	-	-	-	-	-	**2**
Wheat	554	7	-	24	3,401	-	521	13	79	87	**41,904**	**46,590**

How to read: Rows = Vector (RPG declared), Columns = Raster (HRL satellite)
Green diagonal = Correct classification | Yellow = Confusion | Scroll horizontally →

Refresh Statistics

Fig.5 Confusion Matrix

To investigate, we focused on soybean parcels, identified as one of the most problematic classes, and extracted two lists: parcels where farmer declaration and satellite classification agree, and those where they disagree. Of 50 soybean parcels examined in this area, 41 matched (82%) and 9 did not match (18%) (Fig. 6).

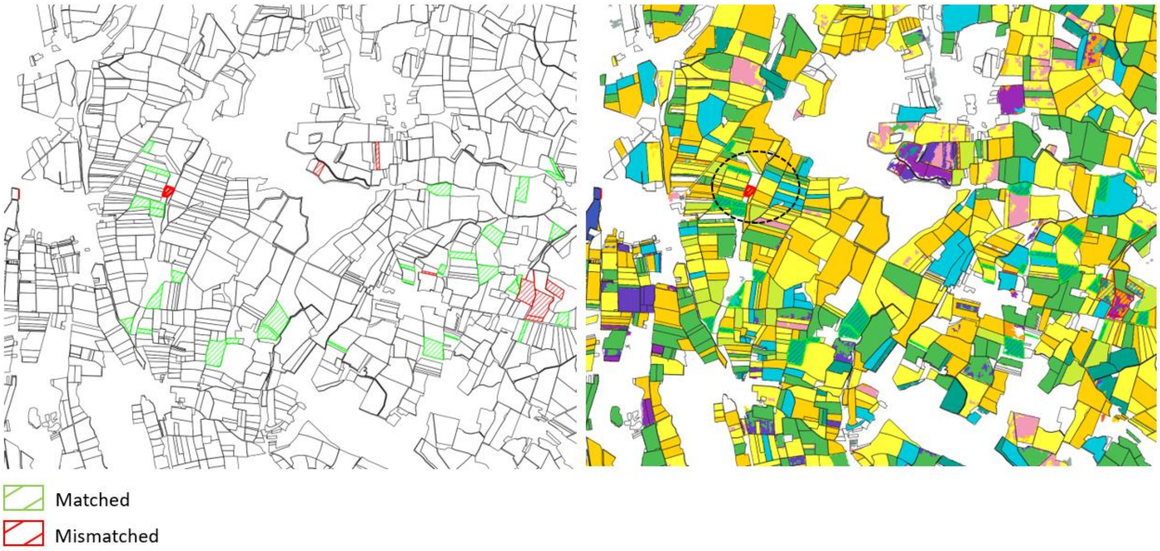

Fig. 6 Matches and not matching parcels

There are at least four possible explanations for a mismatch. First, farmer error: the farmer declared soybeans but actually planted maize. Second, satellite confusion: soybeans and maize are spectrally similar crops. Third, timing issues: the declaration and HRL satellite observation may not correspond to the same moment. Fourth, crop rotation: the field may have changed between declaration and satellite observation.

Distinguishing between these scenarios is critical. Farmer error may require follow-up inspection, while systematic satellite confusion may require recalibrating the model for certain regions or crop types.

Once we identified the mismatched parcels, we built a reference NDVI profile from the 41 correctly matched soybean parcels, including standard deviation bounds. Against this reference, we plotted 5 representative mismatched parcels individually to examine their divergence (Fig.7).

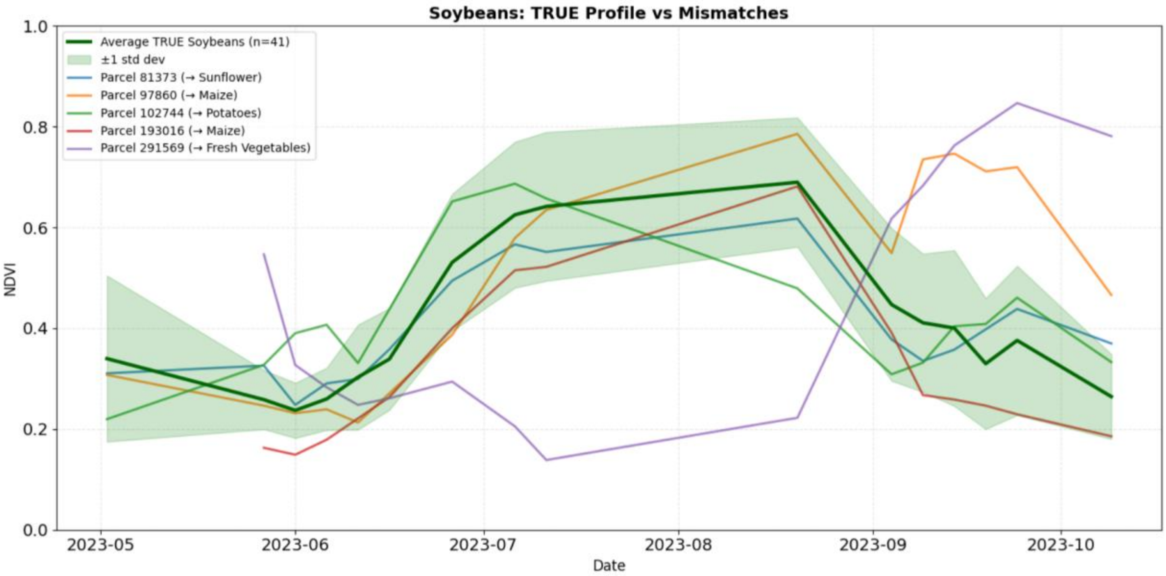

Fig.7 True soybean signature and not matching parcel profiles

We identified four distinct mismatch scenarios, each with different policy implications (Fig. 8):

Parcel 97860 represents the most straightforward case: declared as soybeans, classified by HRL as maize with 97.9% purity. It is a large parcel with no edge effects, and the NDVI profile follows the soybean trajectory until July, then diverges to a pattern consistent with maize phenology. This is a high-confidence mismatch and a strong candidate for farmer declaration error or undeclared crop rotation.

Parcel 294616 tells a similar story from a landscape perspective. Declared as soybeans, it is surrounded entirely by dry pulses fields and classified by HRL as dry pulses with 100% purity. The spectral homogeneity and spatial context together make this another high-confidence case for further inspection.

Parcel 102744 represents the opposite scenario: here the satellite is clearly at fault. The parcel is only 0.1 hectares, adjacent to a large potato field. Despite being technically wider than HRL's 10-meter resolution, mixed pixel effects cause it to be classified as 100% potatoes. This exemplifies a known satellite limitation. Parcels below 0.5 hectares should be flagged as potentially unreliable for automated validation.

Parcel 291569 is the most ambiguous: very low purity at 48.8%, an anomalous NDVI profile with an unexpected spike in September, and no dominant crop classification. This could represent mixed cropping, crop stress, or algorithmic uncertainty requiring case-by-case expert judgment.

Understanding the cause of a mismatch matters as much as detecting it, and these scenarios provide a structured framework for feeding insights back into improving classification algorithms and validation protocols.

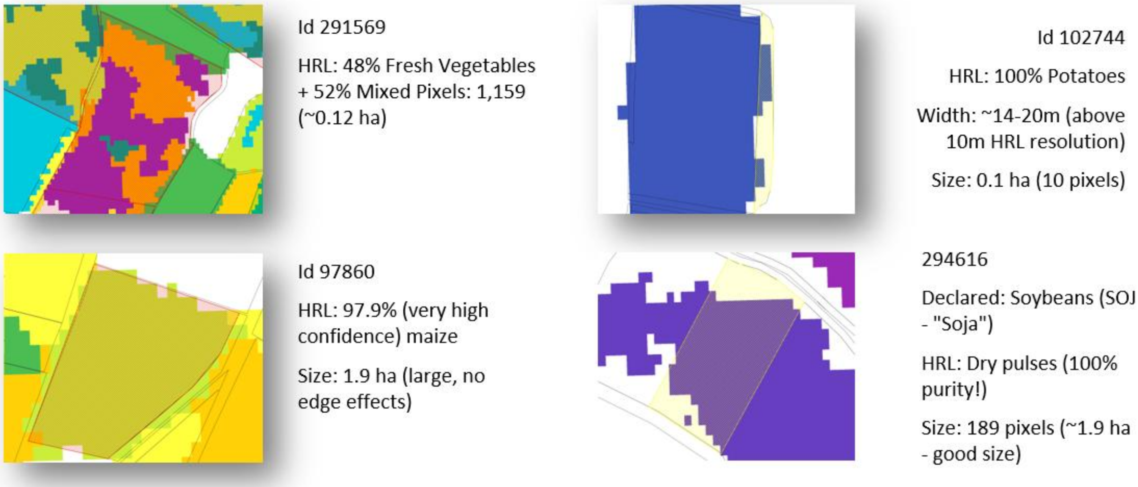

Fig 8. Details extracted from QGIS on not matching parcels

4. Discussion and Policy Implications

This validation framework directly addresses challenges in data access and use for agricultural policy making by providing payment agencies with scalable quality assurance tools for GSA declarations. The viewport-based confusion matrix approach enables risk-based controls: agencies can quantify classification reliability across crop types and geographic contexts, prioritizing inspection resources toward high-confidence mismatches while acknowledging satellite technical limitations for small or boundary parcels. This targeted approach optimizes verification resources while maintaining CAP compliance monitoring effectiveness [2].

The web-based architecture lowers technical barriers by eliminating requirements for specialized GIS software or local data processing infrastructure. Payment agencies access the platform through standard web browsers, with computational processing handled server-side. The modular design facilitates adaptation to different member states by supporting custom crop taxonomy mappings, regional HRL products, and varying GSA data schemas. This transferability is demonstrated through successful deployment across two case study regions (Emilia-Romagna, Italy and Burgundy, France) with distinct agricultural contexts and administrative systems.

Beyond operational monitoring, the framework contributes to agricultural statistics quality assurance by identifying systematic data quality issues. Persistent crop-specific mismatches (e.g., soybean-maize confusion) inform both satellite classification algorithm refinement and GSA data collection protocols. The confusion matrix metrics provide quantifiable, transparent validation statistics supporting evidence-based policy dialogue between the European Commission and member states on monitoring system performance.

5. Code Availability and Future Deployment

The complete analysis workflow is currently accessible to BDAP users through the JRC's JEO-lab GitLab repository, with fully reproducible notebooks for data extraction, NDVI profile comparison, and parcel-level analysis. Planned developments include migration to a public repository on JRC GitLab and deployment on the Copernicus Data Space Ecosystem (CDSE) JupyterHub environment (Fig.9), enabling broader community access for researchers, payment agencies, and member state authorities. We invite collaboration to extend these tools to additional regions, crop types, and validation scenarios.

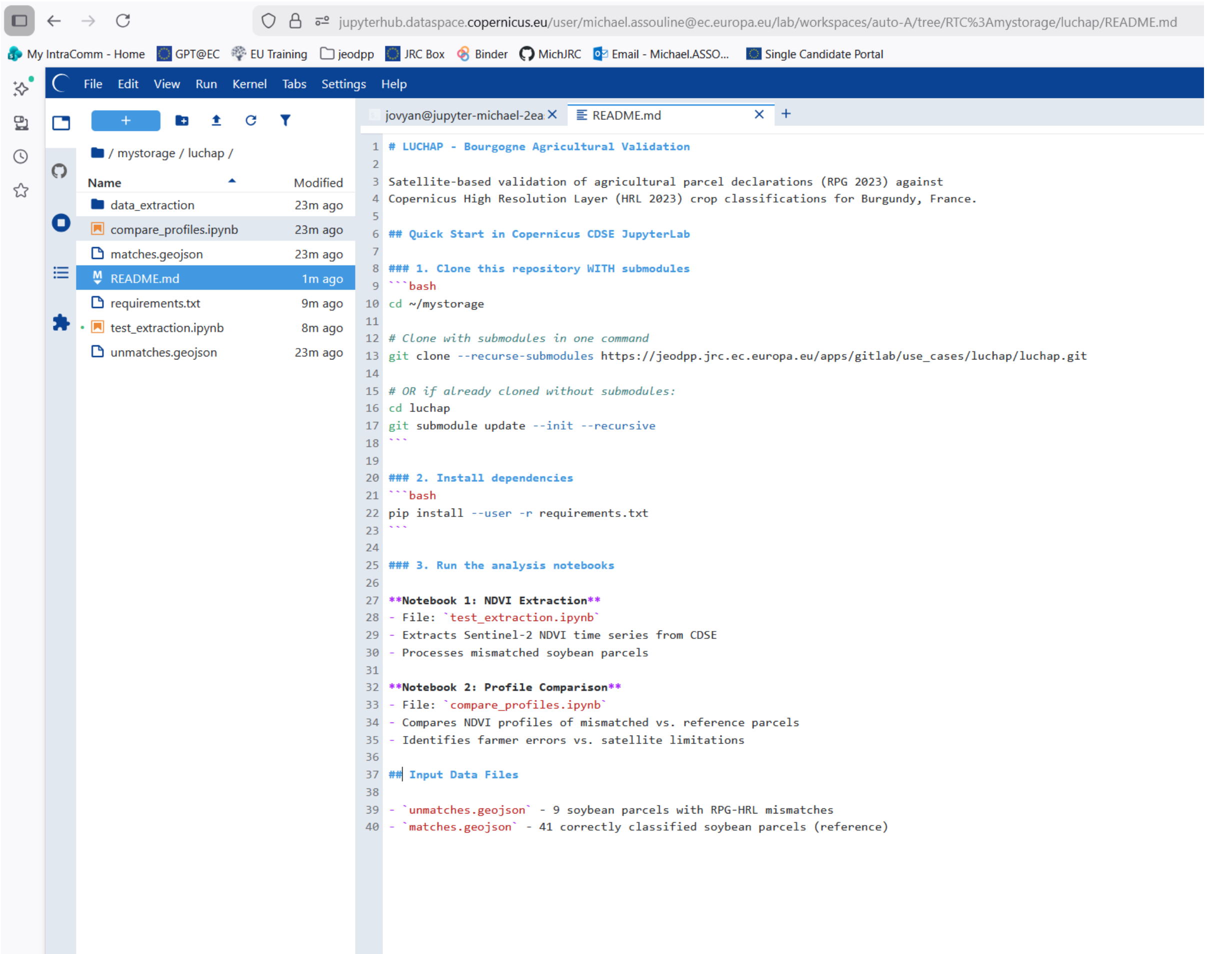

Fig. 9, Repository on the Copernicus Jupyter lab

5. Conclusions

GSA data is a rich and exploitable resource. With the right tools, it can inform multiple dimensions of IACS [3]. This work demonstrates that systematic comparison of farmer declarations (GSA) against satellite-derived crop classifications (Copernicus HRL) can achieve high overall accuracy (92-94%) while identifying actionable patterns for both CAP monitoring and remote sensing algorithm improvement. Analysis of over 600,000 parcels across two EU regions (Emilia-Romagna, Italy and Burgundy, France) reveals that most mismatches fall into identifiable categories with distinct policy implications: high-confidence farmer errors requiring inspection (large parcels with clear spectral signatures), satellite technical limitations (small parcels <0.5 ha with edge effects), and ambiguous cases requiring expert judgment (low purity, mixed pixels).

Integration of spatial analysis (parcel size, classification purity) with temporal validation (Sentinel-2 NDVI phenology) successfully distinguishes these scenarios. For example, soybean parcels showing persistent confusion with maize can be validated through NDVI temporal profiles: parcels following maize phenology indicate likely farmer error, while those matching soybean signatures but misclassified due to small size represent satellite limitations. This categorization enables payment agencies to optimize inspection resources, focusing field verification on high-confidence mismatches while flagging satellite limitations that should not trigger compliance actions.

The LUChaP platform's web-based architecture and modular design enable deployment across member states with varying technical capacity and crop taxonomies. Demonstrated transferability from Italian to French administrative systems, combined with full code reproducibility in the BDAP environment, supports potential scaling to EU-wide CAP monitoring and agricultural statistics quality assurance.

Key contributions to agricultural statistics include: an operational validation workflow integrating administrative records with Earth observation at scale, quantifiable metrics (confusion matrices, purity scores) supporting evidence-based policy dialogue on monitoring system performance, systematic feedback mechanisms improving satellite classification algorithms through identified failure modes, and transparent, reproducible methods enabling independent verification and extension by member state authorities.

Future work will extend temporal validation across five agricultural seasons, develop automated alert systems for payment agencies prioritizing high-risk parcels, and integrate additional Copernicus products (Sentinel-1 for crop phenology, VHR imagery for small parcels). Collaboration with member states to deploy this kind application type will provide large-scale validation datasets improving both CAP monitoring effectiveness and satellite classification accuracy.

By bridging administrative agricultural statistics and satellite Earth observation, this framework demonstrates how integrated data ecosystems can enhance both policy implementation and scientific product quality, supporting the EU's transition toward evidence-based, cost-effective agricultural monitoring at scale.

References

[1] European Union, Copernicus Land Monitoring Service 2025, European Environment Agency (EEA)

[2] Loudjani P, Borio D, Pignocchino G., The Facilitating Implementation of Regulation Elements (FIRE) workshop report, Publications Office of the European Union, Luxembourg, 2025, https://data.europa.eu/doi/10.2760/7244121, JRC140966.

[3] Baiamonte G, Voican G, and Loudjani P, Getting the most of Land Parcel Identification Systems (LPIS) and GeoSpatial Aid Application (GSAA) datasets, European Commission, Ispra, 2023, JRC133145