A Decade of Wheat Mapping for Lebanon

Mohamad Hasan Zahweh†
Faculty of Engineering
Lebanese University
Beirut, Lebanon

Hasan Wehbi†
Earth Observation Department
RASID SARL
Beirut, Lebanon

Hasan Nasrallah
Earth Observation Department
RASID SARL
Beirut, Lebanon

Zeinab Takach
National Center for Remote Sensing
National Council for Scientific Research
Beirut, Lebanon

Veera Ganesh Yalla
IHub Data
IIIT Hyderabad
Hyderabad, India

Ali J. Ghandour*
National Center for Remote Sensing
National Council for Scientific Research
Beirut, Lebanon

***Abstract*—Wheat accounts for approximatly 20% of the world's caloric intake making it a vital component of global food secuirty. Given this significance, mapping wheat fields plays a crucial role in enabling various stakeholders including policymakers, researchers, and agricultural organizations to make informed decisions regarding food security, supply chain management, and resource allocation. In this paper, we tackle the problem of accurately mapping wheat fields out of satellite images by enhancing our previous work on winter wheat segmentation by creating an improved pipeline for processing data as well as presenting a decade-long analysis of wheat mapping in Lebanon. We integrate a Temporal Spatial Vision Transformer (TSViT) with Parameter-Efficient Fine Tuning (PEFT) and a novel post-processing pipeline based on the FOW framework. Our enhanced pipeline addresses key challenges encountered in the previous approach such as clustering of small agricultural parcels in a single large field and sparse training labels. By merging wheat segmentation with precise field boundary extraction, our method produces geometrically coherent and semantically rich maps enabling us to perfom in-depth analysis such as calculating the total number of fields and tracking fields areas year over year. Extensive evaluations demonstrate improved boundary delineation and field-level precision, establishing the framework's potential in operational agricultural monitoring and historical trend analysis. This work lays the foundation for a range of critical studies and future advancements. Building on the accurate mapping of wheat fields, our approach provides a crucial step toward more sophisticated agricultural analyses. Future work can extend this methodology to improve yield estimation, crop monitoring and broader analysis of agricultural trends.**

***Index Terms*—crop monitoring, field deliniation, winter wheat segmentation, PEFT, TSViT, Fields of the World (FTW)**

I. Introduction

Accurate and long-term crop mapping is critical for agricultural monitoring, food security assessments, and policy formulation. In regions like Lebanon, where agricultural parcels vary widely in size, traditional pixel-based segmentation methods often lack the precise field boundaries required for operational applications. Our previous work [1] employed a Temporal Spatial Vision Transformer (TSViT) combined with Parameter-Efficient Fine Tuning (PEFT) for winter wheat segmentation, demonstrating promising results under a weakly supervised learning. However, limitations in spatial resolution and sparse label availability led the model to:

- **Inaccurate Field Boundaries and Out-of-Bounds Predictions** : The model struggled to delineate true wheat field boundaries, often resulting in under-segmentation (missing parts of fields) or over-segmentation (including non-wheat areas). This inaccuracy was particularly pronounced in areas with heterogeneous land cover or where wheat fields were adjacent to other vegetation types.
- **Merging of Adjacent Wheat Fields** : The model frequently failed to distinguish between closely spaced wheat fields, merging them into a single, larger field. This issue arose due to the limited spatial resolution.
- **High Noise in Segmentation Results** : The segmentation results exhibited significant noise, with scattered pixels or small regions incorrectly classified as wheat.

This limited our abilities to generate accurate crop field statistics. This study presents an improved processing pipeline that incorporates a dedicated field delineation model based on the open-source `ftw-baselines` [2]. This work not only builds on our earlier methodology but also provides a refined approach to reconcile pixel-level predictions with the geometric realities of agricultural fields. Our contributions include:

- **A Decade-Long Analysis of Wheat Mapping in Lebanon**: A comprehensive evaluation of wheat field trends over ten years.
- **Enhanced Segmentation with Field Boundary Extraction:** Improved integration of TSViT-PEFT outputs with precise field delineation for better field-level accuracy.
- **Refined Post-Processing for Improved Map Quality:** A novel post-processing technique applied to field delineation results. [2]
- **Addressing Previous Model Limitations:** A demonstration of how the improved pipeline overcomes past challenges, enabling more reliable agricultural monitoring.
- **Data-Driven Insights for Agricultural Decision-Making:** A methodology to generate actionable statistics

*Corresponding Author: aghandour@cnrs.edu.lb.
†These authors contributed equally to this work

for policymakers and stakeholders.

II. Method

A. Wheat Model

Our wheat segmentation approach is based on a transformer-based architecture that takes advantage of both temporal and spatial cues. At its core, the model utilizes the Temporal Spatial Vision Transformer (TSViT), which has been pre-trained on large-scale datasets like the PASTIS dataset [5] and subsequently fine-tuned using Parameter-Efficient Fine Tuning (PEFT) to adapt to our specific task. Multi-temporal Sentinel-2 imagery, processed into surface reflectance products, serves as input, capturing the full growing season of winter wheat. This temporal depth enables the model to learn distinct phenological patterns associated with wheat, while the spatial processing capabilities of TSViT ensure that broad field patterns are identified. The model architecture is particularly designed to handle the variability inherent in remote sensing data over a decade, which makes it suitable for historical trend analysis in Lebanon.

B. Limitation of the previous approach

Despite its strong performance under controlled conditions, the wheat model faces several challenges when applied in real-world scenarios.

- **Spatial Resolution Constraints:** The primary imagery, with a spatial resolution of 10 m from Sentinel-2, often fails to capture fine-scale details. This limitation is especially problematic in regions with small agricultural parcels (often less than 2 hectares), resulting in blurred or incomplete field boundaries.
- **Limited Labeled Data:** With only around 30% of the study area annotated with high-quality ground truth, the model is restricted by sparse training data. This limited supervision affects its ability to generalize across diverse conditions and leads to under-segmentation or misclassification at field borders.
- **Geometric Inconsistencies:** Transformer-based semantic segmentation, while effective at capturing broad crop patterns, tends to generate over-smooth predictions. The resulting masks may lack the crisp boundaries needed to accurately delineate individual fields, which is critical for operational agricultural monitoring.

C. Old Postprocessing Pipeline

Due to the aforementioned limitations, the old pipeline exhibited several issues. The resulting shapefile contained anomalously large polygons—one measuring up to 3907 hectares—which is unreasonably large for a wheat field in Lebanon. Moreover, field boundaries were not clearly defined. As shown in Figure 2, the basemap displays a wheat-classified polygon (outlined in red) that merges multiple wheat fields into a single polygon, thereby compromising the accuracy of the intended statistical analyses. Additionally, Figure 1 illustrates that the shapefile generated by the wheat mapping pipeline contained hundreds of small polygons, representing noise in the model's output.

D. Field Delineation Model

To improve the spatial accuracy of our winter wheat segmentation, we integrated a dedicated field delineation model based on the open-source `ftw-baselines` framework [2]. This model addresses limitations in spatial resolution and geometric inconsistencies by precisely extracting field boundaries. Key aspects include:

- **Architecture:** A convolutional neural network (CNN) with a U-net architecture optimized for instance segmentation tasks.
- **Training Data:** Training is performed on the Fields of The World (FTW) dataset, which contains approximately 1.6 million parcel boundaries across 24 countries. Each sample includes instance and semantic segmentation masks paired with multi-date, multispectral Sentinel-2 images (typically using image pairs from the tilling and harvest stages).
- **Baseline Challenges:** Although the baseline model provides a solid foundation, it suffers from ambiguous boundaries and noise. Our post-processing techniques are designed to resolve these issues.

E. Enhanced Post-Processing

In this work, we enhance our segmentation framework by introducing a novel thresholding method and integrating a dedicated field delineation model. Below, we describe the gradual thresholding approach, the field delineation model, an evaluation of threshold experiments (with figures), and our post-processing pipeline.

1) Gradual Thresholding: Traditional argmax thresholding produces binary masks that often result in undersegmentation. To address this, we introduce **gradual thresholding**.

Gradual thresholding is a technique that sets pixel values below a defined threshold to zero while preserving the original scores of pixels above the threshold, creating a smoother and more informative segmentation mask. We applied the following rules:

- **Strict Threshold for Boundary Mapping:** A threshold of 0.8 is applied to boundary predictions to minimize uncertainty and accurately extract field edges.
- **Relaxed Threshold for Field Mapping:** A threshold of 0.2 is used for field mapping to avoid undersegmentation and retain more complete field information.

2) Post-Processing Pipeline: After obtaining the gradual thresholded masks, we apply the following post-processing steps to refine the segmentation output:

- **Creation of the Basins Mask:** We combine the gradual thresholded field mask with the inverted gradual thresholded boundary mask as follows:

$$\text{basins mask} = \text{gradual\_fields} \times (1.0 - \text{gradual\_boundaries})$$

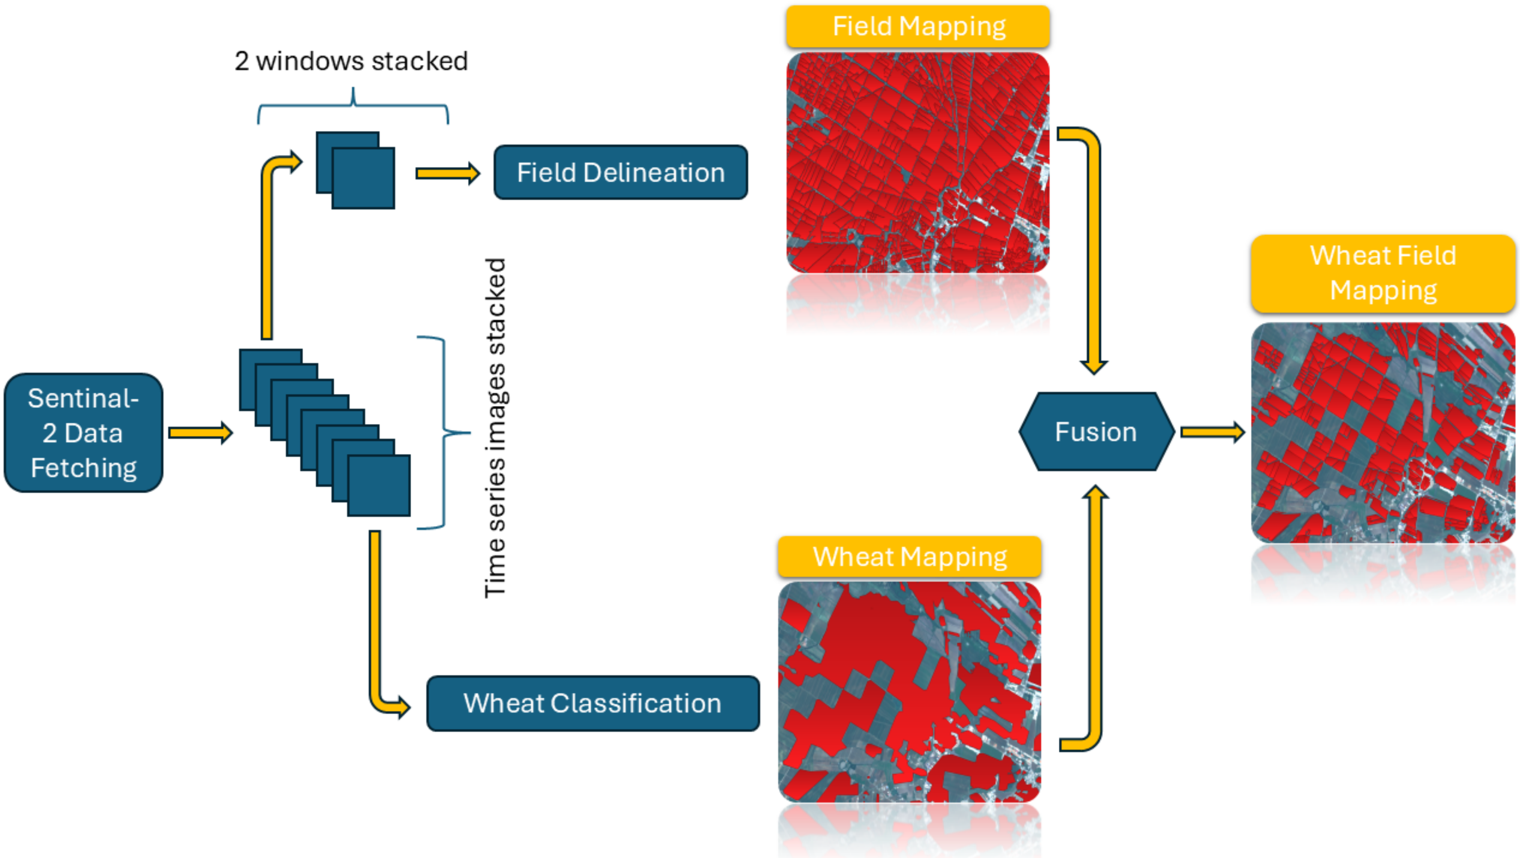

Fig. 1. Overview of the proposed wheat field mapping pipeline. Sentinel-2 time-series images are processed through two parallel branches: (1) a **field delineation model** for **field mapping** and (2) a **wheat classification** model for **wheat mapping**. The outputs from both branches are then fused to generate the final **wheat field mapping** results, ensuring improved spatial accuracy and refined segmentation.

Fig. 2. Large Wheat Polygon

- **Watershed Segmentation:** The basins mask, along with the gradual thresholded field mask, is fed into the watershed algorithm. The negative of the field probability scores is used to determine the basin depths, facilitating accurate instance segmentation.
- **Noise Filtering:** Additional filtering is applied to eliminate small, spurious parcels, thereby reducing noise and improving overall segmentation quality.

F. Evaluation of Thresholding Experiments

We conducted experiments to assess the effects of different threshold values on the segmentation quality. The following figures illustrate the results:

- **Boundary Confidence:** Figure 3 shows the boundaries confidence score map *(a)* from the baseline field delineation model.
  - A threshold of 0.5 (Figure 3 *(b)*) results in high uncertainty along the boundaries, leading to many fields being omitted.
  - In contrast, thresholding at 0.8 (Figure 3 *(c)*) retains most fields and yields better segmentation quality.
- **Field Confidence:** Figure 3 *(e)* displays the field confidence score map.
  - When thresholded at 0.5 (Figure 3 *(f)*), high uncertainty causes several fields to be removed.
  - Thresholding at 0.2 (Figure 3 *(g)*) preserves most fields and results in improved segmentation.
- **Comparison with Old Thresholding:** Figure 3 *(e)* illustrates the old approach using argmax, which led to undersegmented fields. In comparison, our gradual

thresholding method (see Figure 3 *(h)* produces more accurate segmentation with many distinct fields.

G. Pipeline

1) Automated Sentinel-2 Fetching Pipeline: We implemented an automated data fetching pipeline that utilizes the Sentinel-Hub API to request data from the Copernicus Data Space Ecosystem (CDSE). This Sentinel-2 fetching pipeline served a dual purpose: it was integrated into the full processing pipeline and was also used to generate training data. Specifically, we fetched five years of data for an Area of Interest (AOI) in Baalbek, Lebanon, and labeled them in-house. This dataset became the foundational labeled dataset for Lebanon. We then split this dataset to fine-tune the TSViT model, as described in our work. Furthermore, this pipeline will be employed for future tasks, such as automating live crop health monitoring by retrieving current and future imagery.

2) Pre-processing: Both the wheat model and the field delineation model require preprocessing steps, as the raw satellite images cannot be fed directly into the models. For the wheat model, we first read the satellite images for the target year, starting with months 11 and 12 of the previous year, followed by months 1, 2, 3, 4, 5, 6, and 7 of the target year. We then extract the necessary spectral bands: 1, 2, 3, 4, 5, 6, 7, 8, 11, and 12 (0-indexing). Each image is normalized using a predefined mean and standard deviation, ensuring that the input scale remains consistent, which helps improve model performance. After normalization, we concatenate the images into a tensor of shape (1,1,11,H,W), where:

- 1 is the batch dimension
- 1 is the time dimension (multiple months are stacked)
- 11 corresponds to 10 spectral bands plus 1 additional channel for time encoding
- H and W represent the image height and width

Next, we iterate over the stacked tensor, rearrange its dimensions, and split it into 24×24 pixel patches before feeding them into the model. The model's predictions are then post-processed into a class-wise segmentation map, where 1 represents a wheat pixel and 0 represents a non-wheat pixel. Finally, we save the segmented image as a new TIFF file.

The field model follows a different preprocessing approach. We first read the satellite images for the target year and the following year, selecting month 11 from the target year and month 6 from the following year. We then extract the necessary spectral bands: 2, 3, 4, and 8. The two selected time windows are stacked together. Each image is normalized by dividing pixel values by 3000. The images are then cropped into 256×256 pixel chunks, with any leftover areas padded with 0. The model's predictions are processed as described earlier, using thresholding, and the results are saved as a new TIFF file

3) Post-processing: We now iterate over the saved TIFF files generated by the wheat model and apply a post-processing routine. This routine utilizes the watershed algorithm to create an instance mask from the raster image. The instance mask is then passed to a noise filtering script, which removes polygons below a defined size threshold. Following this, we fill any gaps or holes within the polygons, and subsequently simplify and regularize the remaining polygons. The field model undergoes a similar post-processing routine, but without the noise filtering and regularization steps

4) Result: The two shapefiles are merged using the following criteria: any field polygon that intersects with a wheat polygon is appended to a final shapefile, designated as type "wheatfields." To account for instances where the wheat model identifies wheat but the field model does not detect a corresponding field, we also append the geometric difference between the wheat and field polygons to the final shapefile esignated as type "maybewheatfields.". The resulting shapefile accurately delineates wheat fields.

III. Statistical Results

To understand the spatial and temporal dynamics of wheat cultivation over the past decade, we analyzed changes in wheat and non-wheat areas, focusing on expansion, abandonment, and land-use transitions. The table I shows the change of wheat are in km^2 over the last 10 years. We can see how wheat increase for example from 2022 till 2023 while the most significant growth occurring between 2018 and 2019. Also the huge decrease between 2019 and 2020 showes the potential effect of the severe financial crisis that Lebanon suffered from the year of 2019. Table II shows the area of intersection between two consecative years and between two non-consecative years skipping one year in between. The consecative intersection tracks the continous land use changes. So farmers that planed wheat then a non-wheat fratile then a wheat one within adjacent years. Skipping intersections captures non-consecutive changes like wheat abandoned for a few year then reclaimed. Upon inspecting the results, we

Year	Total Wheat Area	Non Wheat Area
2016	177.0	10032.1
2017	176.8	10032.2
2018	177.6	10031.7
2019	254.4	9955.4
2020	144.4	10064.6
2021	192.7	10016.6
2022	187.9	10021.2
2023	224.7	9984.6
2024	139.1	10069.5

TABLE I
CHANGE OF AREA IN km^2 OVER THE YEARS.

Year A	Year B	Consecative Intersection	Year B	Skipping Intersection
2016	2017	51.7	2018	87.2
2017	2018	48.6	2019	93.5
2018	2019	60.0	2020	70.1
2019	2020	47.0	2021	92.6
2020	2021	43.6	2022	73.7
2021	2022	59.9	2023	97.3
2022	2023	74.8	2024	72.3
2023	2024	54.2	–	–

TABLE II
CONSECUTIVE VS. SKIPPING INTERSECTIONS WHEAT AREA IN km^2 OVER CONSECATIVE YEARS

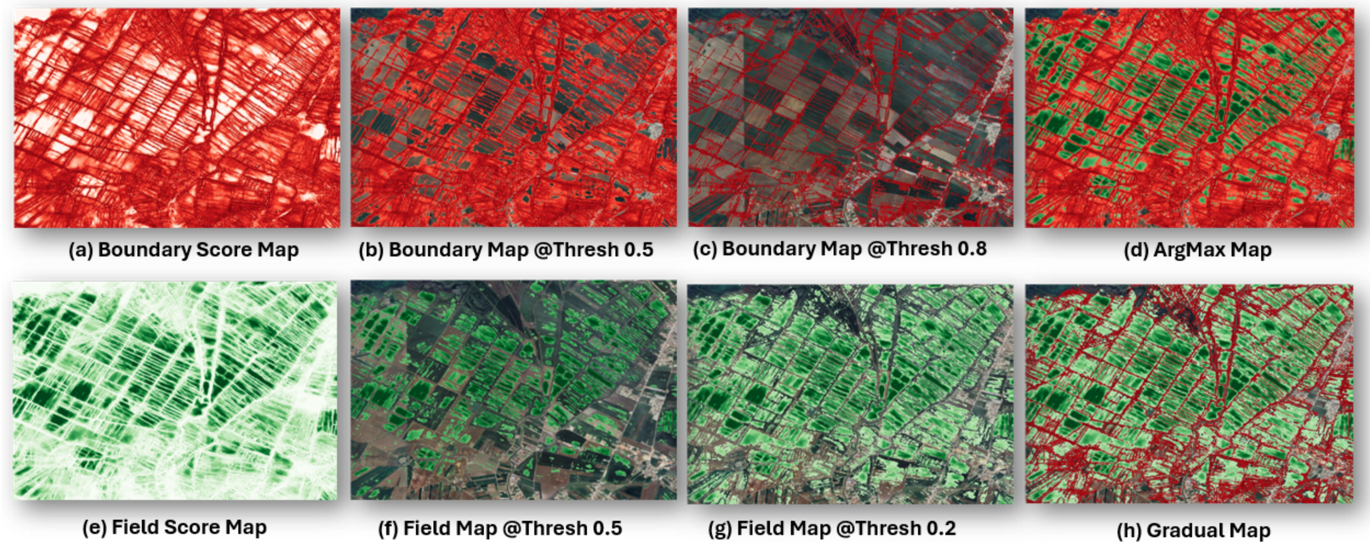

Fig. 3. Comparison of boundary and field segmentation results under different thresholding strategies. Figures (a)–(c) show boundary score maps and their thresholded outputs at 0.5 and 0.8, respectively, while (d) illustrates the traditional argmax approach. Figures (e)–(g) present the field score map and its thresholded versions at 0.5 and 0.2, and (h) demonstrates the final segmentation obtained using our gradual thresholding method.

observed that the intersection between skipped years is consistently greater than the intersection between consecutive years. This suggests a pattern of alternating wheat land abandonment in Lebanese agriculture, a practice known as crop rotation. We will examine this phenomenon in detail while presenting the Sankey diagram. Crop rotation is a strategic agricultural technique where farmers systematically alternate the types of crops planted in a field over successive growing seasons. This practice offers several key benefits like soil fertility, decreasing pest amounts and disease managemnt.

While tables presented above provides a numerical breakdown of annual wheat area changes, the Sankey diagram in Figure 4 complements this by illustrating the movement of land between different states (wheat to non-wheat and vice versa). This combined approach offers both quantitative accuracy and visual clarity

To better visualize transitions, a Sankey diagram was generated, illustrating the flow of land between different states over time. Let us examine two specific flows to better understand the results

- The white upper pillers represent the non-wheat area of a year
- The white lower pillers represent the wheat area of a year
- The red flow from a lower piller to an upper piller represents the abonded wheat areas from a year to a year
- The blue flow from an upper piller to a down piller represents new wheat areas from a year to a year
- The green flow from a lower piller to a lower piller is the wheat areas that stayed wheat from a year to a year
- The black flow from an upper piller to an upper piller is the non wheat area from a year to a year

The Sankey diagram reveals distinct patterns in wheat cultivation trends over the past decade, particularly in the abandonment of wheat fields and the introduction of new wheat areas. A notable observation is the periodic increase in abandoned wheat areas, represented by the red flow, occurring approximately every three years ($2016 \rightarrow 2019 \rightarrow 2023 - 2024$). This suggests that farmers may be gradually reducing wheat cultivation due to factors such as soil degradation, water scarcity, or shifting economic condition. The sharp increase in wheat abandonment during 2023-2024 coincides with the war on Lebanon, especially the war threats on the Bekaa Valley, the top wheat-producing region in the country Additionally, the blue flow, which represents new wheat areas, follows a cyclical pattern of expansion and contraction approximately every two years. This trend suggests that farmers may be implementing a crop rotation system, where wheat cultivation alternates with other crops or fallow periods to preserve soil health and optimize yields. Such a practice aligns with common agricultural strategies aimed at replenishing soil nutrients and mitigating pest infestations

IV. Future Work

As you may have noticed, we have not yet validated our statistics due to resource limitations. According to the USDA [3], the yield production of wheat in Lebanon for 2023 is 3.5 tons per hectare. Conversely, another study by Our World in Data [4] reports a yield of 2.5 tons per hectare for the same year. Other than that we can't simply state that the yield is 2.5 tons per hectare and multiply it by the area in hectares we generated, since our model retrieves all types of wheat breeds and land types without distinguishing between them. It is essential to note the presence of two distinct field types: rainfed fields and watershed agriculture fields

- Rainfed Fields: These rely solely on natural rainfall for irrigation. Farmers do not manually water these lands,

Fig. 4. Annual Sankey Diagram

making their yield highly dependent on seasonal precipitation patterns

- Watershed Agriculture Fields: These employ automated irrigation systems, providing farmers with greater control over water supply and resulting in more consistent yields

Furthermore, we must distinguish between differenet wheat breeds. In lebanon we have maily two breeds durum and soft wheat varieties, each with its own yield characteristics. Our model currently detects all wheat fields without differentiating between these types and irrigation methods. Consequently, we cannot accurately determine the tons per area from our statistics, as each field's yield varies depending on its type and management practices. This highlights a new research challenge: how to detect and label rainfed and watershed fields using remote sensing alone, and how to identify and distinguish between Durum and soft wheat varieties. Incorporating these labels would enable us to report more precise yield estimations. For example, we could state that we detected x1 km^2 of soft wheat, which yields y1 tons, and x2 km^2 of Durum wheat, which yields y2 tons. This would allow us to provide accurate local annual wheat production reports for Lebanon.

In our next research, we will try to fill the gap between our current model's output and the accurate, detailed agricultural statistics needed for effective policy and resource management in Lebanon. We aim to develop methodologies that can automatically differentiate between rainfed and watershed fields, as well as identify and classify Durum and soft wheat varieties, utilizing remote sensing data and advanced machine learning techniques. This will allow for more precise yield estimations and contribute to a deeper understanding of Lebanon's agricultural landscape.

References

[1] M. H. Zahweh, H. Nasrallah, M. Shukor, G. Faour, and A. J. Ghandour, "Empirical Study of PEFT techniques for Winter Wheat Segmentation," arXiv preprint arXiv:2310.01825, 2023.

[2] H. Kerner, S. Chaudhari, A. Ghosh, C. Robinson, A. Ahmad, E. Choi, N. Jacobs, C. Holmes, M. Mohr, R. Dodhia, *et al.*, "Fields of The World: A Machine Learning Benchmark Dataset For Global Agricultural Field Boundary Segmentation," arXiv preprint arXiv:2409.16252, 2024.

[3] United States Department of Agriculture (USDA) Foreign Agricultural Service, "Lebanon: Wheat Crop Summary," Available: https://ipad.fas.usda.gov/countrysummary/default.aspx?id=LE&crop=Wheat, [Accessed: Mar. 2025].

[4] Our World in Data, "Wheat Yields," Available: https://ourworldindata.org/grapher/wheat-yields?tab=chart&country=~LBN, [Accessed: Mar. 2025].

[5] Vivien Sainte Fare Garnot and Loïc Landrieu, "Panoptic Segmentation of Satellite Image Time Series with Convolutional Temporal Attention Networks," *CoRR*, vol. abs/2107.07933, 2021. [Online]. Available: https://arxiv.org/abs/2107.07933.