Journal of Artificial Intelligence and Data Mining (JAIDM), Vol. 14, No. 3, 2026, 395-406.

Shahrood University of Technology

Journal of Artificial Intelligence and Data Mining (JAIDM)
Journal homepage: http://jad.shahroodut.ac.ir

Research paper

HURA-Net: A New Model for Agricultural Field Boundary Detection

Mehdi Alizadeh, Parvin Ahmadi* and Masoumeh Azimzadeh

ICT Research Institute, Tehran, Iran.

Article Info

Article History:
Received 04 October 2025
Revised 02 May 2026
Accepted 15 June 2026

DOI:10.22044/jadm.2026.16909.2823

Keywords:
Agricultural Field, Boundary Detection, Segmentation, U-Net, Tversky Loss Function.

**Corresponding author: p.ahmadi@itrc.ac.ir (P. Ahmadi).*

Abstract

Field boundary detection is a critical task in modern agriculture, enabling precision farming, optimized resource management, and efficient crop monitoring. Despite its importance, existing deep learning models often fail to achieve high accuracy in delineating field boundaries due to challenges such as complex landscapes, varying resolutions, and noise in remote sensing images. To address these limitations, we propose HURA-Net, an advanced deep learning framework that intelligently integrates UNet++, ResUNet, and an attention mechanism into a unified architecture. By hybridizing these models, HURA-Net effectively combines their strengths—such as multi-scale feature extraction (UNet++), residual learning (ResUNet), and focus on salient regions (attention mechanism)—while minimizing their individual weaknesses. To further enhance performance, we introduce a refined loss function that not only improves segmentation precision but also addresses the class imbalance problem, which is common in boundary detection tasks. Extensive experiments on a diverse dataset of high-resolution satellite images from different regions of Iran demonstrate that HURA-Net significantly outperforms existing state-of-the-art models. Specifically, it achieves a recall of 45.85% (a 15.59% improvement over ResUNet) and an F1-score of 42.62% (7.27% higher than ResUNet), setting a new benchmark for accuracy. Moreover, our study highlights the critical role of strategic data augmentation in boosting model generalization, particularly in handling variations in lighting, crop types, and field shapes. The success of HURA-Net underscores the importance of innovative architecture design, optimized loss functions, and robust training strategies in advancing remote sensing image segmentation.

1. Introduction

Precision agriculture has changed the face of agriculture by developing and utilizing state-of-the-art technologies in managing the fields, monitoring crops, and yield prediction. Precise field boundary detection has gained a lot of attention and prominence in precision agriculture, from conventional approaches that use manual mapping to highly sophisticated methods based on high-resolution remote-sensing imagery. Still, small size, irregular shape, and complex topographic features are the challenges with which datasets come. This gets aggravated in the presence of class imbalance, where field boundaries are very under-represented as compared to the non-boundary regions, affecting the poor performance of the models. This occurs in a case where the number of pixels in one class-a representative field boundary, for instance-is much lower in number as compared to the other class-non-boundary areas-so therefore giving a very poor detection performance of that minority class.

It is recently that deep learning appeared to be powerful for image segmentation [1-3], where Convolutional Neural Networks demonstrated stunning performance in field boundary detection. Papers [4] and [5] provided basic instances of the application of deep learning approaches to this particular domain. These works brought out the potentials of CNNs and further sophisticated architectures like UNet and its variants on boundary detection. However, issues of class imbalance, fragmented boundaries, and irregular shape of the fields are still there, which limits the accuracy of the existing models. Recently, works involving inclusion of attention mechanisms [6] and specialized loss functions, including Tversky loss, have given promising results over these issues [7]. Despite these advances, even more robust models are needed to tackle the challenges posed by high-resolution remote sensing imagery.

This paper proposes the HURA-Net model, a hybrid model combining the strengths of UNet++, ResUNet, and attention mechanisms with the Tversky-Fbeta loss function to advance the state of the art in agricultural field boundary detection. Our method tries to overcome the shortcomings of the existing models and exploits the advantages of such complementarities, like class imbalance and complex shapes of the boundaries. HURA-Net embeds the nested skip connections with residual learning and attention mechanism to keep the spatial information, grasp better features, and emphasize salient regions which are very crucial for accurate boundary detection. Furthermore, class imbalance is tackled to optimize performance on rare or fragmented boundary pixels with the Tversky-Fbeta loss function.

The contributions of this work are threefold: (1) the development of a novel hybrid architecture that integrates UNet++, ResUNet, and attention mechanisms to enhance field boundary detection; (2) the introduction of the Tversky-Fbeta loss function to address class imbalance and improve segmentation accuracy; and (3) extensive experimentation on a diverse dataset of high-resolution remote sensing images from various regions in Iran, demonstrating the superior performance of HURA-Net compared to state-of-the-art models. These results show significant improvements in recall and F1-score, confirming that the model can effectively detect faint and fragmented boundaries.

2. Related Work

Various research has been conducted on the image segmentation model, ranging from medical imaging to the detection of field boundaries in agriculture. Models like U-Net and ResUNet, first proposed for medical image segmentation, were later adapted for agricultural use due to their high capabilities in solving complex segmentation tasks.

Ref. [4] proposed a boundary delineation CNN multi-task approach; while it outperforms traditional approaches, it also still suffers from class imbalance and detection of smaller or less distinctive boundaries. Ref. [5] introduce the method of agricultural field boundary detection using multispectral and Synthetic Aperture Radar (SAR) data fusion. However, this model evidenced poor performance in fragmented or even just fainter boundaries, which in turn usually happens with high-resolution imagery.

Recent works have then also given stress to eliminating these issues. For instance, [8] studied the Segment Anything Model for field boundary delineation without extra training data. Results demonstrated that SAM can capture around 58% of the field boundaries accurately, which is reasonably comparable to other methods that require huge training data. Ref. [9] proposed a multi-region transfer learning approach for the segmentation of crop field boundaries from satellite images with limited labels. Their approach outperformed state-of-the-art techniques by a margin, showcasing the potential of transfer learning in this domain.

Significant improvements in segmentation accuracy have also been achieved with architectural developments. UNet, proposed by [10], has a symmetric encoder-decoder architecture with skip connections that preserve spatial information; hence, it is very effective in segmentation tasks. Later, [11] improved this architecture by incorporating nested and dense skip connections into UNet++, thereby reducing the semantic gap between encoder and decoder feature maps. ResUNet by [12] further enhanced this performance through the integration of residual learning [13] to allow for much deeper networks with limited problems of the vanishing gradients. Ref. [14] applied U-Net, SegNet, and DenseNet deep learning models in segmenting the smallholder farming fields using a high-resolution WorldView-3 satellite that was taken over Bangladesh. Among them, DenseNet had the best performance and reached an accuracy as high as 0.8, showing its potential in small and irregular shapes of fields. Ref. [15] further expanded this work on rice paddy field segmentation using a high-resolution satellite image to show the potential of UNet and its variants for accurate crop mapping. They also investigated transfer learning and weak supervision methods to scale field delineation across diverse agricultural landscapes, emphasizing model scalability and adaptability.

ResUNet has emerged as a powerful architecture for semantic segmentation, outperforming many state-of-the-art approaches in remote sensing data analysis

[16]. Ref. [17] further enhanced ResUNet for field boundary delineation by integrating adversarial deep learning techniques, including adversarial loss functions. Fully convolutional networks have also illustrated immense potential in boundary delineation. For instance, [18] applied FCNs for the automated delineation of smallholder agricultural fields and achieved better precision and recall. As pointed out by [19], semantic segmentation in heterogeneous systems presents a trade-off between precision and scalability.

The trend of these state-of-the-art advancements puts a higher burden of continuous research for high accuracy and broader applicability in models for image segmentation for field boundary detection. Class imbalance, complex shapes of boundaries, and diversity in datasets are challenges to be faced for the generalization of models across various real-world conditions. This work presents HURA-Net, a hybrid model combining the strengths of UNet++, ResUNet, and attention mechanisms with the Tversky-Fbeta loss function, which overcomes these limitations to further advance the state of the art in agricultural field boundary detection.

3. Proposed Dataset

Current work relies on the created custom high-resolution image tiles dataset related to agricultural fields from various provinces in Iran. They were prepared carefully by applying accurate geo-referencing of the images and segmentation processes of agricultural landscapes using ArcMap software. Our dataset consists of a total of 119 images with 903×903 pixels (Figure 1), which expands the diversity and richness compared to previous works. We took multiple steps to obtain uniform size and quality images during the preprocessing stage. In particular, each large image was cropped into smaller sections of 128 × 128 pixels, which is the input size required by the model. The augmentations made during training include the following: (1) 90-degree clockwise rotation, (2) 90-degree counterclockwise rotation, (3) horizontal flipping, (4) vertical flipping, (5) horizontal and vertical flipping together and (7) inclusion of original unaltered data. This helps to increase the variability of the training data and make the model robust and adaptable under various field conditions. Hence, this should be expected to improve the model performance in terms of accurate boundary segmentation pertaining to agriculture fields.

The imagery used in this study was obtained from Google Earth and covers agricultural regions from different provinces of Iran. The dataset consists of RGB imagery with an approximate spatial resolution of 1 meter per pixel. Google Earth imagery was selected because it provides very-high-resolution observations suitable for detailed agricultural field boundary annotation. The field boundary masks were manually generated using ArcMap based on visual interpretation of the imagery. This dataset covers most of the agricultural features, including small, large, and complex-boundary fields, which better reflect reality. The information provides a strong basis for model development and testing regarding the detection of field boundaries in agriculture, where the variety of scenarios tests the functionality of these models to perform effectively. Its high quality enhances its potential to promote research in precision agriculture, among other applications in remote sensing.

In contrast, the AgriculturalField-Seg dataset [20] consists of 1,200 small 500 × 500 pixel images from orthophotos of agricultural areas in Catalonia, Spain, representing more than 3,300 agricultural fields. Although it contains several classes of agricultural fields, it is relatively geographically limited and less diverse compared to our data. The AgriculturalFieldSeg dataset exhibits class imbalance problems when there are few boundary pixels compared to the background. It also includes complex, fragmented, or faint boundaries that make segmentation tasks difficult. Moreover, this dataset is limited to only one region in Catalonia; therefore, its generalization to other agricultural regions is minimal.

The AI4Boundaries dataset [21] concerns large-scale agricultural monitoring across different parts of Europe. This dataset includes 7,831 annotated agricultural field boundaries derived from high-resolution aerial imagery and Sentinel-2 satellite data, covering seven European countries, including France, Germany, and the Netherlands, at a spatial resolution of 10 meters per pixel. Despite its extensive geographical coverage, it lacks the high-resolution granularity necessary for precise boundary detection tasks. The AI4Boundaries dataset is intended for broader agricultural analysis and does not focus on detailed field boundary delineation. Furthermore, it does not address class imbalance issues or provide preprocessing techniques relevant to deep learning models. Therefore, it is not directly suitable for tasks requiring the detection of small or irregularly shaped field boundaries with high accuracy. In contrast, our dataset covers various provinces of Iran, offering a wider geographical distribution and variation in agricultural landscapes that will enhance the generalization of our model across diverse environments.

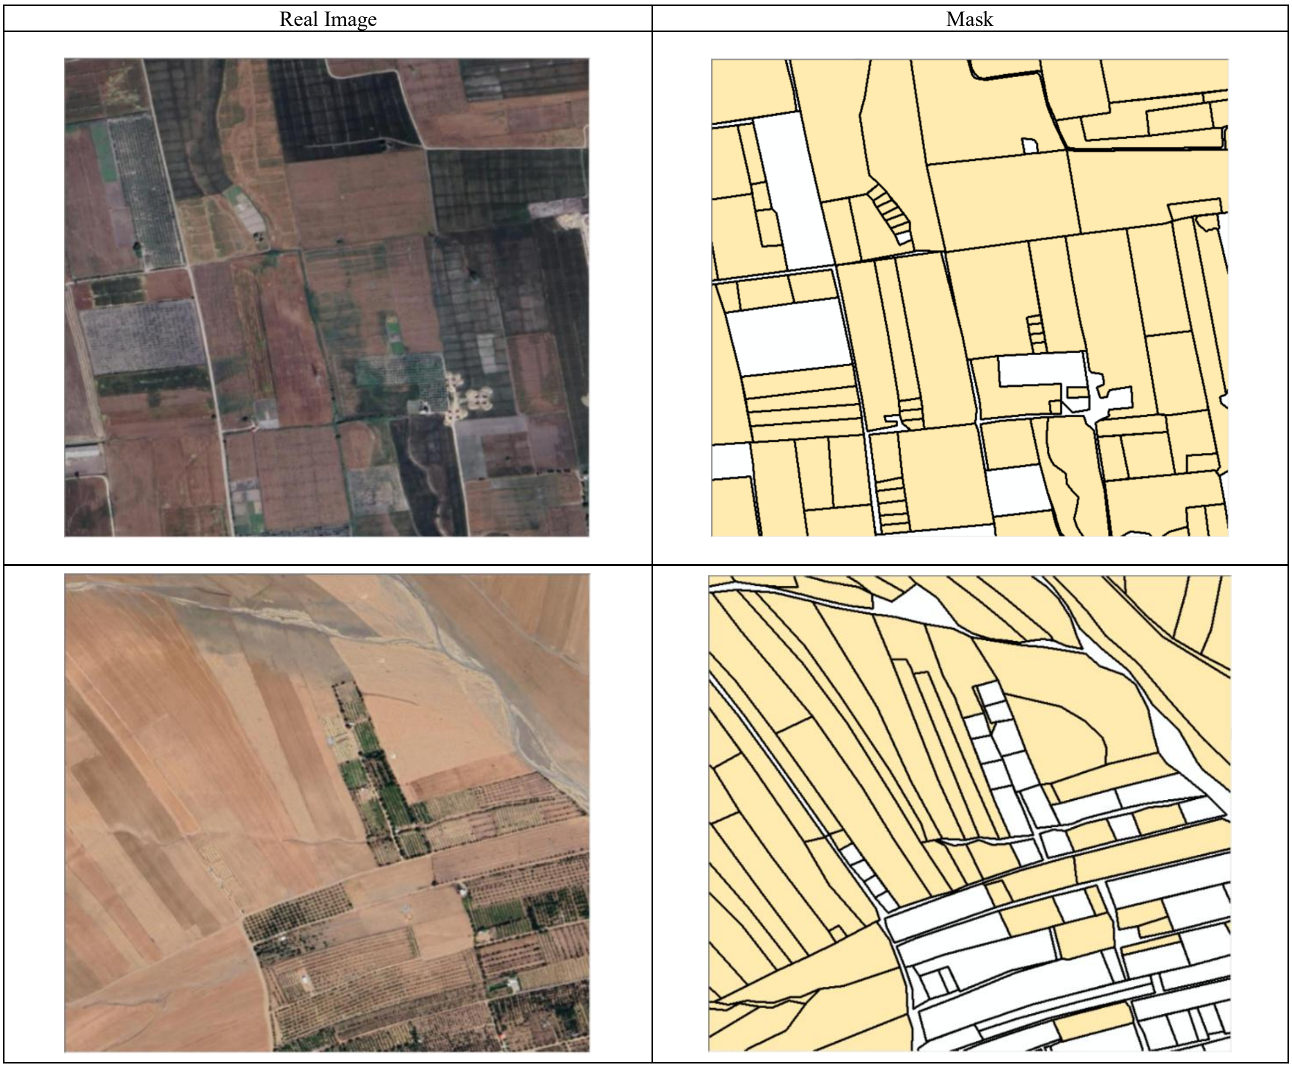

Figure 1. Samples from our dataset and their relevent mask.

Our dataset has been meticulously preprocessed and expanded to further enhance model performance. The large images were cropped into uniform sections of 128×128 pixels. The patch size of 128×128 pixels was chosen to balance spatial context and computational efficiency. Given the approximate spatial resolution of 1 m/pixel, each patch represents an area of approximately 128 m × 128 m. This size provides sufficient contextual information for identifying agricultural field boundaries while avoiding the excessive memory requirements associated with larger input patches.

In data augmentation, the following were considered: 90-degree rotations, flipping horizontally and vertically, and changes in brightness to increase variability in training data. These pre-processing strategies allow the model to handle complex boundary shapes and conditions that may be underrepresented either in the AgriculturalField-Seg or AI4Boundaries. Table 1 shows these datasets based on their features. In other words, while the AgriculturalField-Seg and AI4Boundaries datasets have their merit in contributing toward agricultural segmentation, our dataset offers superior geographical diversity, larger image sizes, and advanced preprocessing techniques for a more robust and adaptable model. These enhancements improve the performance of detecting complex field boundaries and enhance the generalizability of the model across diverse agricultural contexts.

4. Proposed Method

The approach of the present study focuses on the development of a deep learning model, namely HURA-Net, integrating UNet++, ResUNet, and attention mechanisms to achieve highly accurate segmentation of field boundaries in agriculture. This work aims to develop a network that balances maintaining details and capturing contextual information, addressing the unique challenges presented by remote sensing data. The following sections outline the architectural details of the network, the attention mechanism, the choice of loss function, and the training process to improve performance.

Table 1. Datasets and their features.

Dataset Name / features	Our Dataset	AgriculturalField-Seg	AI4Boundaries
Countries Covered	Iran (various provinces)	Catalonia, Spain	Seven European countries (France, Germany, Netherlands, etc.)
Number of Images	119 (903×903 pixels, cropped to 128×128)	1200 (500×500 pixels)	7,831 annotated field boundaries
Image Size	903×903 pixels (original), 128×128 cropped	500×500 pixels	Derived from Sentinel-2 and aerial imagery
Spatial Resolution	Approximately 1 meter per pixel (obtained from Google Earth)	Not explicitly mentioned	10 meters per pixel
Preprocessing	Cropping, augmentation (rotations, flipping, brightness changes)	Not explicitly mentioned	Not explicitly mentioned
Augmentation Techniques	90° rotations, horizontal/vertical flipping, brightness changes	None mentioned	None mentioned
Geographical Diversity	High (various provinces in Iran)	Low (limited to Catalonia, Spain)	Moderate (covers seven European countries)
Focus	Detailed field boundary segmentation	Agricultural field segmentation	Large-scale agricultural monitoring
Class Imbalance	Addressed through preprocessing and augmentation	Present (few boundary pixels compared to background)	Not addressed
Complex Boundaries	Handled through preprocessing and augmentation	Present (complex, fragmented, or faint boundaries)	Not explicitly addressed
Generalizability	High (diverse agricultural landscapes across Iran)	Low (limited to one region)	Moderate (covers multiple countries but lacks high-resolution detail)
Applications	Precision agriculture, field boundary detection	Agricultural field segmentation	Broad agricultural analysis

4.1. Network Architecture

In developing HURA-Net, the focus has been on incorporating a number of architectural building blocks with the objective of saliency feature capture while progressively refining the segmentation output. The overall architecture of the model can be elaborated as follows:

1. Input Layer: This is the starting point or entry of the model, where the raw image data in their native format are fed into the model for processing.

2. Encoder Path: The encoder consists of a series of interleaved ResUNet and UNet++ blocks to capture features at higher levels of abstraction. The UNet++ block comprises convolutional feature extraction layers connected through nested dense skip pathways between encoder and decoder stages. These dense skip connections reduce the semantic gap between feature maps at different resolutions and improve the reconstruction of fine agricultural field boundaries. The ResUNet residual connections further facilitate gradient flow, enhancing convergence stability in deeper architectures.

3. Bridge Layer: This is a transitional bridge layer that is used for extracting high-level features using the UNet++ block. These are critical for the capture of accurate boundaries. This provides the continuity of the most necessary contextual information between encoder and decoder paths.

4. Decoder Path: In the decoder path, the model reconstructs the spatial information lost due to down-sampling. This is achieved using a combination of an attention block followed by switches of ResUNet and UNet++ blocks. Attention blocks enable focusing on crucial areas that might define the field boundaries. The rest of the components ensure that both local details and global patterns are well-represented in the output.

5. Output Layer: The last output layer consists of a convolutional layer activated by the sigmoid, which outputs the probability map where each pixel gives the probability of its membership to the field boundary. This would yield a clean, continuous boundary map that reflects the understanding of the model.

4.2. Attention Block

By incorporating attention mechanisms into HURA-Net, it is ensured that the model does not treat all features equally but learns to give more importance to regions that might be more relevant for boundary detection. The attention block basically works by generating weight maps highlighting the main features present within the input image. Guided by this to focus on the more important regions and features, refinement of the segmentations can be more accurate in places where these boundaries are camouflaged or very faint.

4.3. Loss Function

For training segmentation models, choosing an appropriate loss function matters a lot. This paper thus uses the Tversky-Fbeta loss (with $\alpha = 0.7$ and $\beta = 0.3$); this is more powerful to handle the class imbalance problem that often happens in agricultural field detection; such cases often result in less than a fraction of the edge coverage of an image. It helps in fine-tuning the balance between precision and recall by regulating how much the model penalizes for false positives and false negatives, hence more accurate and dependable segmentation results. This has successfully been employed by authors in segmenting medical images with high accuracy.

4.4. Training Strategy

The training process incorporated several strategies to improve performance and reduce overfitting. The Adam optimizer was employed with a learning rate of 0.001, providing stable and smooth parameter updates during optimization. A batch size of 1 was selected due to the computational and memory requirements of the proposed hybrid architecture. Despite the small batch size, stable convergence was achieved during training, and model optimization remained consistent throughout the learning process. Training was performed for a maximum of 50 epochs, while an early stopping mechanism monitored validation performance and terminated training when no further improvement was observed for a predefined number of epochs. This strategy prevented unnecessary training and reduced the risk of overfitting. The dataset was randomly divided into training and testing subsets using an 85%/15% split. The training subset was used for model development, while the testing subset was reserved exclusively for performance evaluation.

The augmentation of data proved to be mainly important to create a robust model across various conditions. Such transformations included rotations, flipping, etc., in an amount that ensured diversity for training data, training the model to deal with several views or shapes of the field.

In particular, exciting architectures can be combined with attention mechanisms, a fitting-so-far loss function, and a well-conceived training regime to set HURA-Net (Figure 2) ready to meet some challenges confronted in agricultural field boundary detection.

5. Results

HURA-Net had a comparison performed against different high-level methods for the segmenting of boundaries of the agricultural field. These methods included UNet, UNet++, AttentionUNet, and ResUNet. A summary of these methods is provided in Table 2. The loss function for the proposed model is Tversky-Fbeta loss, whereas the binary cross-entropy (BCE) was utilized in other models.

All evaluation metrics reported in this study were computed at the pixel level by comparing the predicted segmentation masks with the corresponding ground-truth masks. Consequently, the reported Precision, Recall, Accuracy, and F1-score represent pixel-wise segmentation performance rather than boundary-specific evaluation measures.

Precision: HURA-Net minimal excludes precision at 55.94 from models UNet++ at 73.81 and AttentionUNet at 72.78; however, in general, precision is not of importance for the agricultural field boundary detection task, as the identification of boundaries is supposed to be done as comprehensively as possible. This includes attributing all the relevant regions even at the cost of some false positives. Thus, the difference in precision does not substantially undermine the efficacy of HURA-Net in the task.

Recall: HURA-Net far surpasses all the other models in terms of recall, with a rate of 45.85 which is 28.84 higher than ResUNet at 30.26 and 38.44 higher than UNet++ at 7.41. The recall metric relies heavily on boundary detection as far as the model is capable of adequately marking the complete relevant regions that are depicted in the image, however faint or fragmented they may be. On these boundary areas, HURA-Net excels to capture.

F1-score: The F1 score is found out to be 42.62 for HURA-Net against the benchmark models ResUNet at 35.35, UNet++ at 11.73, and AttentionUNet at 9.22. This derived metric provides equal emphasis on both precision and recall; it is this higher value of F1 that shows that HURA-Net has managed to find an optimal trade-off between detection of boundaries (recall) and minimization of false positives (precision) while detecting agriculture field boundary.

Algorithm 1: HURA-NET with Residual U-Net and U-Net++ and Attention Blocks.

```
Inputs:
- Learning rate α = 10e-3
- Batch size = 1
- Training epochs T = 50
- Input shape (128, 128, 3)
- Initial filters L = 16
Algorithm
1. Initialize parameters of the HURA-Net model θ.
2. for epoch = 1 to T do
3.    repeat
4.       Sample a batch of size 1 from {X, Y}.
5.       Pass the batch through the encoder path:
6.          for each down-sampling step i do
7.             if i is even, apply Residual U-Net block.
8.             if i is odd, apply U-Net++ block.
9.             Apply MaxPooling2D.
10.      Pass the output of the encoder to the bridge (U-Net++ block).
11.      Pass the output through the decoder path:
12.         for each up-sampling step i do
13.            Apply Conv2DTranspose.
14.            Concatenate with corresponding encoder features.
15.            Apply attention block.
16.      Apply the final Conv2D layer with sigmoid activation to get the output.
17.      Calculate the Tversky-Fbeta loss.
18.      Update θ using the Adam optimizer with learning rate α.
19.   until all batches are processed.
20. end for
Outputs
- Trained model parameters θ
```

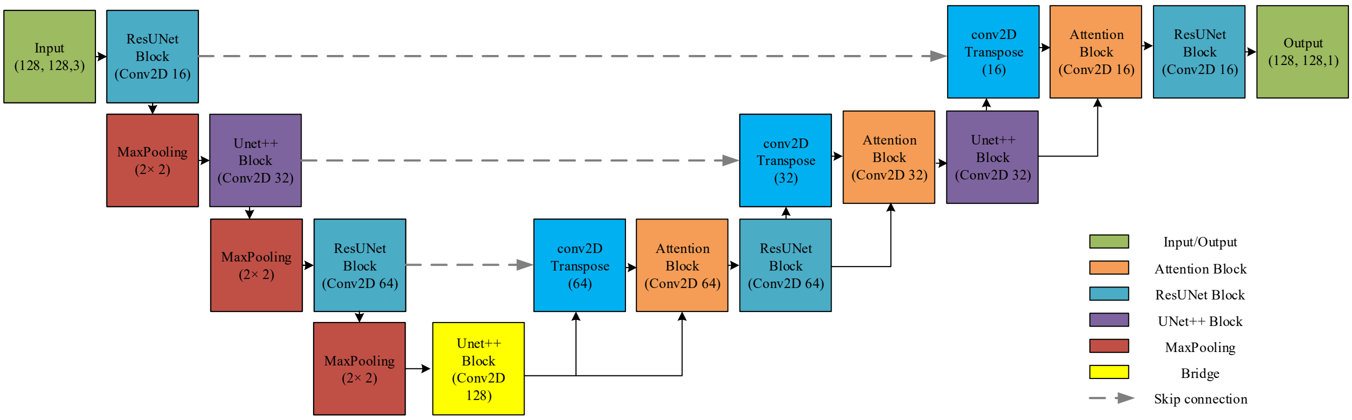

Figure 2. Proposed model diagram.

Table 2. Precision, Recall and F1score of models.

Models	Precision	Recall	F1-score
UNet	68.73	17.01	24.72
UNet++	**73.81**	7.41	11.73
AttentionUNet	72.78	5.71	9.22
ResUNet	57.29	30.26	35.35
Proposed model (HURA-Net)	55.94	**45.85**	**42.62**

We also evaluated the performance of the HURA-Net model with binary cross-entropy (BCE) loss to further understand how the loss function affects model performance. The results from BCE indicated that the model has a precision of 69, recall of 8, and F1-score of 13. These results illustrate that the Tversky-Fbeta loss function outperforms binary cross-entropy loss in terms of recall and F1-score, thereby winning kudos for class imbalance that is characteristic of agricultural field boundary detection.

Certainly, the choice of the loss for training, the Tversky-Fbeta loss function and not the standard binary-cross-entropy loss which was used by all other models, stands out as one of the key contributory factors toward what accounts for the increase in performance on the HURA-Net. Thus, this function is very important because Tversky loss is basically directed toward solving class imbalance issues, usually a recurring phenomenon among tasks like agricultural boundary detection where only a minor area of the respective image holds the boundaries. Through its employment of Tversky-Fbeta loss, HURA-Net is favored for higher recall and a better precision-recall trade-off, and thus, more effective for such complex and imbalanced

segmentation tasks like field boundary detection in agriculture.
Figure 3 intuitively shows a less noisy, more complete and precise boundary detection by HURA-Net compared with other models. Its output conformed better to true field boundaries than any other model, which is one of the major application needs of agricultural monitoring, yield prediction and land management.

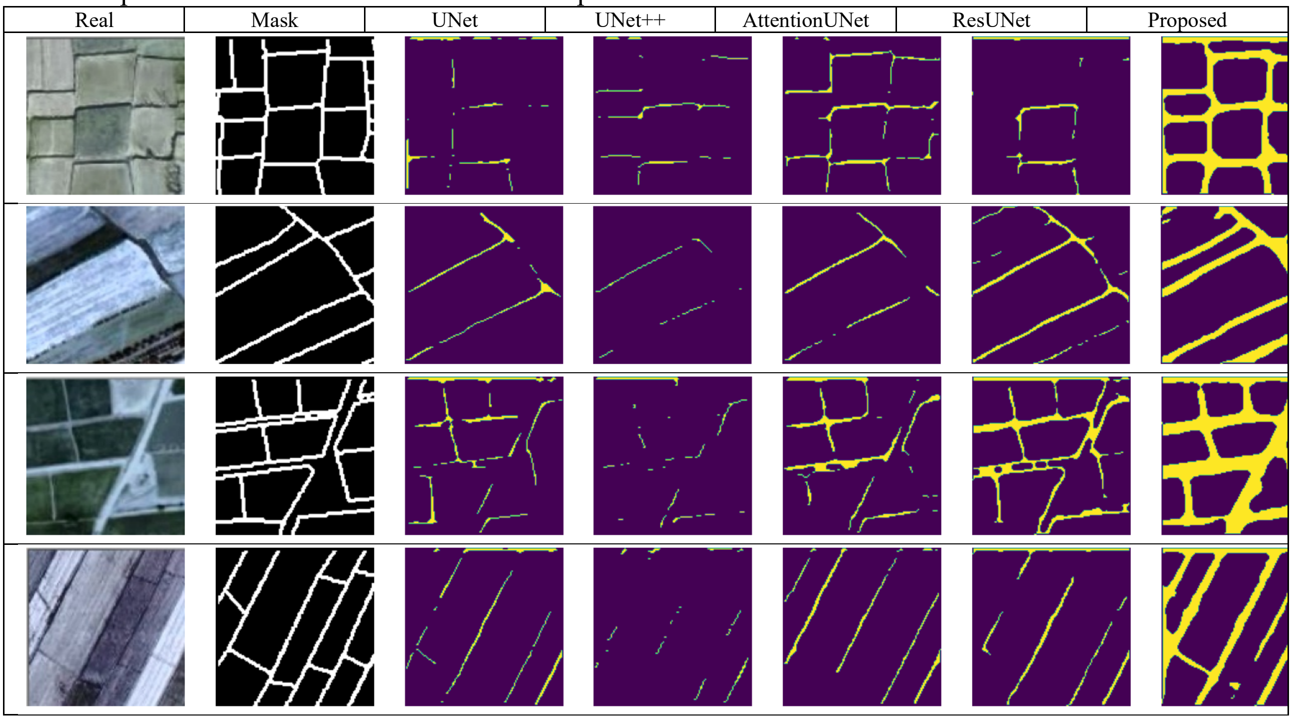

Figure 3. Proposed model prediction results beside other models.

Though UNet++ and Attention UNet, achieves superior precision, agricultural field boundary detection requires much more coherent and accurate boundary segmentation. That is, HURA-Net strives to detect as many relevant pixels as possible while avoiding significant information loss, realizing this aspiration with utmost success-due to its powerful recall and F1-score. HURA-Net has higher recall in comparison with the remaining models; thus, it is more capable of recognizing the boundaries in challenging and complicated situations. This could give applications in precision agriculture a definite edge in detecting field boundaries thoroughly and correctly, even in fragmented or unclear boundary data.

5.1. Post-processing of predicted binary mask

The purpose of post-processing is not to compensate for deficiencies in the segmentation model but rather to transform the predicted boundary masks into continuous field polygons that can be directly used in practical agricultural applications.
Several post-processing steps were applied to the predicted binary mask to refine the segmentation results and ensure boundary continuity. To achieve these objectives, edges were first detected with the Canny edge detector, which were then dilated to connect nearby segments. Subsequently, the dilated edges were processed with the Probabilistic Hough Line Transform to extract straight lines created on a blank image with the same dimensions as the mask. The line image was then overlaid with the original binary mask to enhance boundary clarity. Finally, morphological closing was applied to seal breaks and produce a smooth segmentation result, referred to as the closed mask.
The post-processing step plays a crucial role in enhancing the quality of the predicted field boundaries. It addresses common issues such as gaps, discontinuities, and noise in the raw predictions by applying morphological operations and contour refinement techniques. This step ensures that fragmented or incomplete boundaries are seamlessly connected, resulting in smooth and continuous field outlines. It also removes small artifacts and irregularities, improving the cleanliness and precision of segmentation.
Improved continuity of lines, aside from visual quality, is necessary for the correct identification of closed regions, which can be further used to extract individual fields. This refinement highlights the robustness of the proposed approach, making it well-suited for practical applications in agricultural field mapping and management.

As shown in Figure 4, the post-processing significantly enhanced the accuracy and usability of the predicted boundaries for agricultural applications. This figure illustrates the complete workflow of field boundary detection using the proposed model. The images are already in the form of tiles, and the training dataset was prepared by cropping these tiles into smaller sections. The first column presents the original input image; the second column shows the corresponding ground truth mask used for model evaluation. The third column displays the raw predicted mask from the model, highlighting its initial performance in predicting field boundaries. The fourth column presents a post-processed version of the predicted mask with more refined boundaries for better segmentation. Finally, in the fifth column, the closed regions derived from the post-processed mask are visualized, with each enclosed boundary representing a different field. This comprehensive visualization clearly demonstrates the effectiveness of the proposed approach in accurately identifying and isolating agricultural field boundaries.

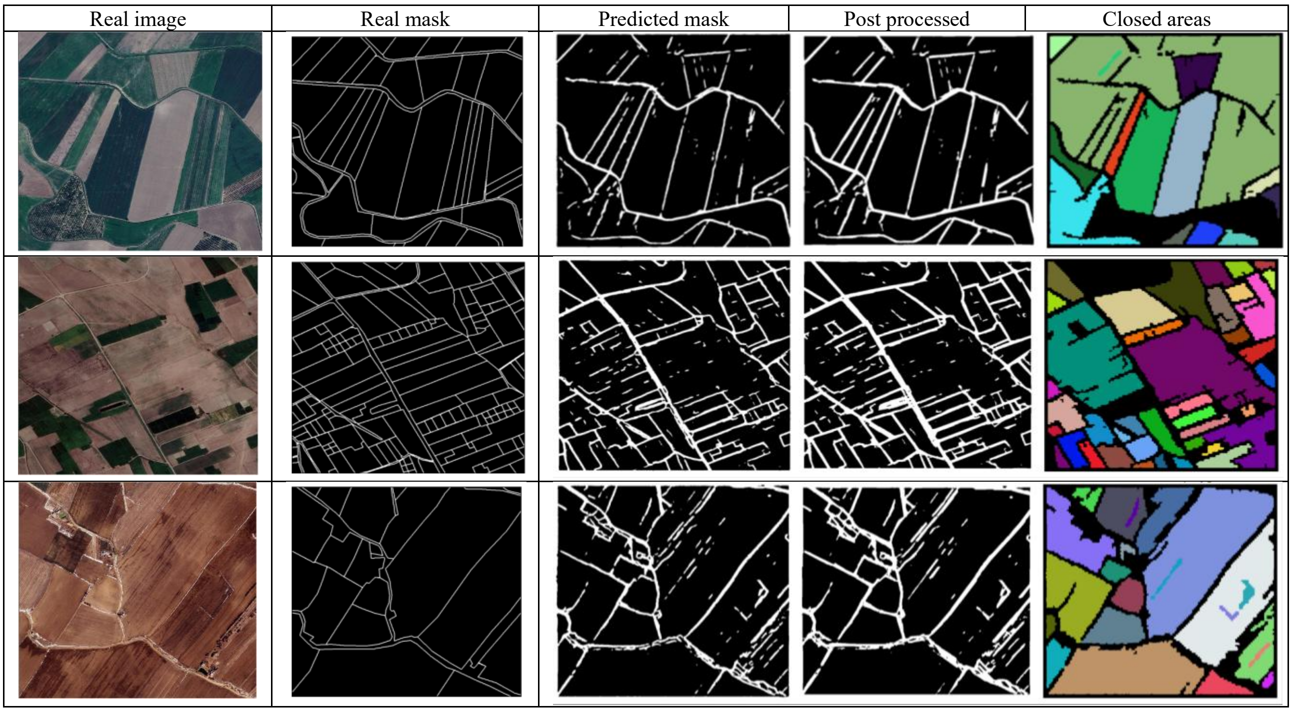

Figure 4. The first column presents the real image next to its corresponding ground truth mask. The middle column indicates the model-predicted mask and post-processing of that predicted mask. The right column represents closed areas around every boundary as a field.

It should be noted that the post-processing operations, particularly edge dilation and morphological closing, may produce boundaries that are visually thicker than the corresponding ground-truth annotations. However, the primary objective of these operations is to ensure boundary continuity and facilitate the extraction of closed field polygons. In practical agricultural mapping applications, complete and connected boundaries are often more important than achieving single-pixel boundary thickness.

5.2. Dataset Expansion

To improve the robustness and generalization capability of the proposed model, both dataset expansion and data augmentation strategies were employed. Data augmentation included image transformations such as rotations and flipping. In this experiment, brightness adjustment was added, which involves randomly changing the brightness of the image within a range of -20 to +20 pixel values. Additionally, the dataset was expanded by incorporating high-quality image tiles acquired during different seasons, enabling the model to learn from a broader range of agricultural and environmental conditions.

Concretely, there were 119 high-quality image tiles. Furthermore, 224 tiles were added, which, although derivatives of the same high-quality tiles, were captured at different times of the year. This temporal variation was introduced to enhance the robustness of the model in field boundary detection under diverse seasonal and environmental conditions. Including images from different seasons exposes the model to more factors that can alter the appearance of field boundaries, such as crop growth stages, soil moisture, and lighting conditions. This ensures that the model generalizes better to real-world scenarios, where field boundaries may look quite different depending on the time of year or environmental changes. In total, we used 327 tiles as the dataset. We enhanced the boundary detection capability of the model by introducing inputs of different dimensions. While the previous approach used fixed input dimensions of 128×128, this experiment employed expanded input sizes of 64×64, 128×128, 256×256, and 512×512. This variation ensures that the model is not dependent on a single fixed input size and can better adapt to identifying field boundaries at different scales. This multi-scale approach improves the robustness and effectiveness of the segmentation model.

Table 3 compares performance metrics for the agricultural boundary detection model before and after data increment. In the case of fixed 128 × 128 input dimensions, the model yielded a moderate performance with an F1 score of 36.47, recall of 38, and precision of 51.19. After the increase in multi-spectral input sizes (from 64 × 64 to 512 × 512), there is an overall gain in all metrics: F1 score improves by 0.7%, recall increases by 3.86%, and precision by 4.12%. It follows from here that multi-scale input dimensions enriched the model for field boundary detection from weak to strong scenarios. The results show the key roles of both diversity and scalability of data to be considered toward developing better accuracy in the boundary detection model's robustness within precision agriculture.

Table 3. Metrics of proposed model before and after data increment.

Data	Precision	Recall	F1 score
Before increment	51.19	38	36.47
After increment	**55.31**	**41.86**	**37.17**

Figure 5 shows the results of dataset augmentation and enhancement on the chosen model's predictions. The original image is shown on the right side, with the model's prediction on the left. The top one is the model prediction just using small high-quality data (Table 2), and the bottom one uses more high-quality data, along with their crop sizes. This comparison indicates that these improvements in boundary prediction for the model have been considerably enhanced by increasing the dataset diversity and including multi-scale cropped images.

6. Discussion

The results show that HURA-Net effectively overcomes key challenges in agricultural field boundary detection, including handling class imbalance and segmenting complex boundary shapes. It uses the Tversky-Fbeta loss function to better handle rare boundary pixels by balancing false positives and false negatives. This capability is critical for thorough boundary detection, where recall and F1-score metrics are of greater importance than precision. Although the precision of HURA-Net is slightly lower compared to some benchmark models, its ability to comprehensively identify all relevant boundary areas makes it highly suitable for precision agriculture applications.

The architecture of HURA-Net, which combines nested skip connections and residual blocks, bridges the semantic gap between encoder and decoder features while facilitating efficient gradient flow for deeper learning. The incorporation of attention mechanisms further enhances segmentation by emphasizing critical areas, thus enabling the model to perform well even when faced with fragmented or faint boundary conditions. Therefore, it outperforms state-of-the-art models in recall and F1-score, metrics that accurately reflect its capability to detect intricate and irregular boundary shapes.

Data augmentation and dataset expansion played pivotal roles in improving the model's generalizability and robustness. Techniques such as rotation, flipping, brightness adjustment, and multi-scale input dimensions introduced variability into the training set, ensuring the model's adaptability to diverse field shapes, sizes, and boundary complexities. The enriched dataset, which included higher-quality and more diverse samples, further enhanced the model's ability to detect boundaries effectively across varying scales and conditions. These post-processing steps involved edge detection, morphological operations, and contour refinement, which significantly improved boundary continuity and noise reduction within the predicted masks, resulting in seamless and accurate segmentation. When comparing HURA-Net with recent agricultural field boundary delineation methods, differences in datasets, spatial resolutions, annotation protocols, and evaluation procedures should be considered. The dataset used in this study includes substantial variability in field geometry, seasonal conditions, and boundary visibility, making the delineation task more challenging than in many benchmark datasets. Despite these challenges, HURA-Net demonstrated competitive performance and robust boundary extraction capabilities.

Unlike conventional semantic segmentation tasks, agricultural field boundary delineation often requires topologically valid and closed boundaries. Consequently, a lightweight post-processing stage was incorporated to ensure geometric consistency and facilitate field extraction.

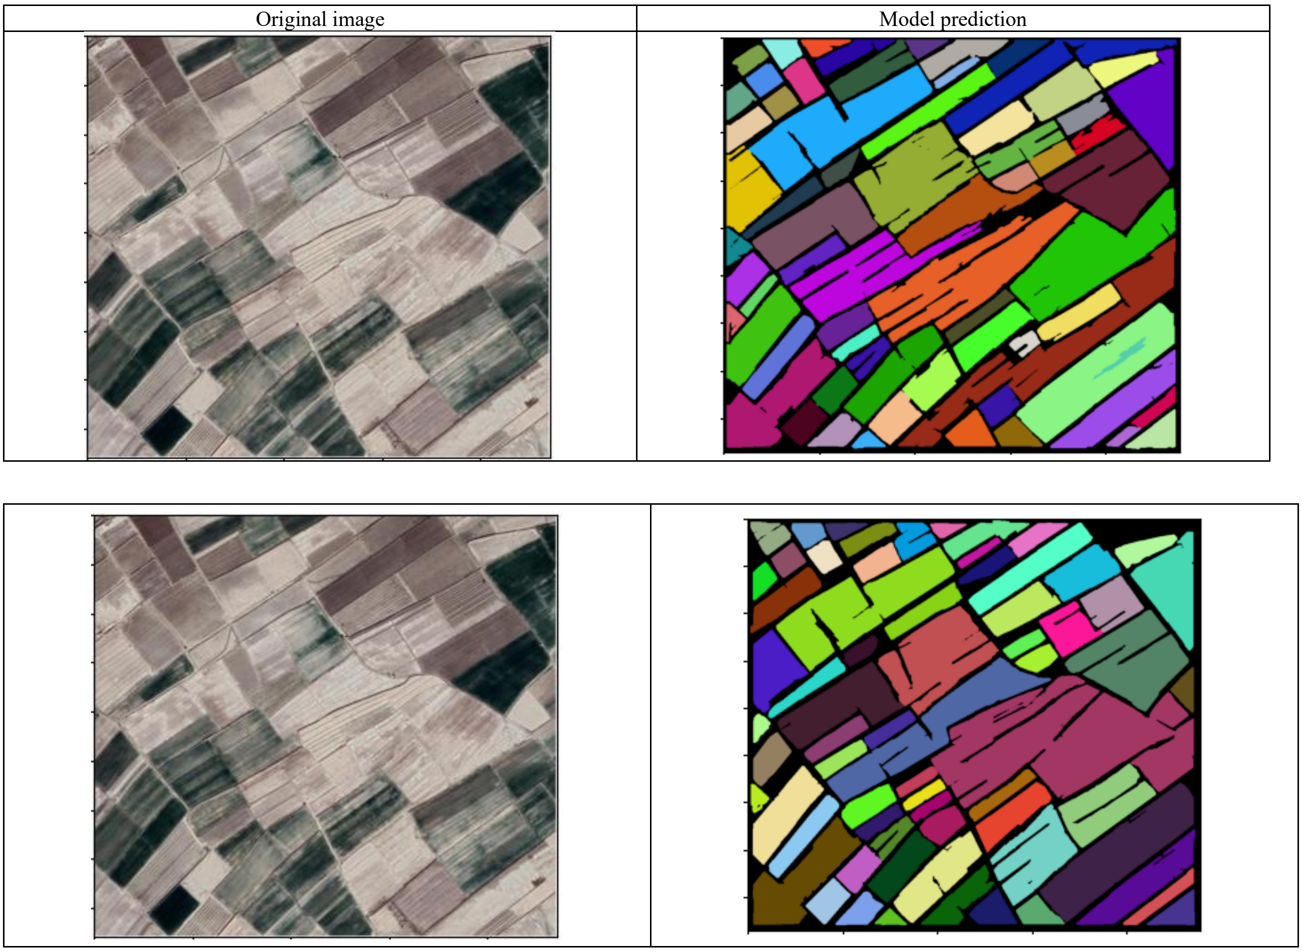

Figure 5. Prediction of the selected model on low data (top) and high data (bottom).

The proposed framework should therefore be viewed as a complete boundary delineation pipeline consisting of both segmentation and geometric refinement components. High-resolution imagery may pose a challenge when deploying HURA-Net in resource-constrained environments. Additionally, the relatively low precision necessitates post-processing to eliminate false positives in certain high-accuracy applications. A limitation of the proposed approach is that the post-processing stage may produce relatively thick boundary representations. Although this improves boundary continuity and facilitates field extraction, it may reduce geometric precision in some cases. Future work will focus on enhancing model efficiency, incorporating multi-modal data sources, expanding the dataset, and investigating boundary refinement techniques such as skeletonization, contour regularization, and boundary-aware loss functions to generate thinner and more accurate field boundaries.

7. Conclusion

This study proposes a new deep-learning architecture, the so-called HURA-Net model for detecting agricultural boundaries. The model is designed in such a way that it leverages the powers of UNet++, ResUNet, and attention mechanisms in solving major challenges like class imbalance and irregular boundary shapes found in most agricultural field data. Besides, the Tversky-Fbeta loss function integrated in it helps in further improving segmentation performance by paying more attention to correctly identifying fragmented or faint boundaries. By proposing a hybrid method that combines some of the key features of previously developed models, HURA-Net mitigates the issues arising from complex and irregular field shapes more effectively. This is very important in the context of precision agriculture, where high accuracy and the capability to handle such diverse and challenging datasets are important. It appears that HURA-Net outperforms other models in state-of-the-art, clearly improving the scores along two very critical axes for the performance metrics for a segmentation task. These HURA-Net results improve Recall by 15.59%, changing it from 30.26% to 45.85%, while increasing F1-score by 7.27%, or from 35.35% to 42.62%. The improvements such as those underlines above show improved field detection capabilities of a model even at poor conditions. The robustness and generalization capability of the model are increased with extensive data augmentation techniques and a diverse training dataset, allowing it to adapt to on-field conditions. These effects can be seen from the improved metrics: the F1-score has improved by 0.7%, recall increased by 3.86%, and precision improved by 4.12% after the increase in dataset size and multi-scale input dimensions. This adaptability, therefore, serves to underline the potential of HURA-Net in practical applications of precision agriculture. Further research is thus focused on optimization of model performance with a view to scale up for practical deployment. Additionally, efforts will be dedicated toward enabling the model to manage highly heterogeneous agricultural landscapes that will pave the way for wide-scale adoption of this technology into global precision farming practices.

References

[1] H. Pourghasem, "Comparing the semantic segmentation of high-resolution images using deep convolutional networks: Segnet, hrnet, cse-hrnet and rca-fcn," *J. Inf. Syst. Telecommun.*, vol. 4, no. 44, pp. 359–368, Dec. 2023.

[2] M. J. CM, "Convolutional neural networks for medical image segmentation and classification: A review," *J. Inf. Syst. Telecommun.*, vol. 4, no. 44, pp. 347–358, Dec. 2023.

[3] E. Yadolahi and S. Abolmaali, "Improving the Efficiency of Semantic Segmentation Implemented in Spiking Neural Networks," *Journal of AI and Data Mining*, vol. 13, no. 1, pp.25–39, 2025.

[4] F. Waldner and F. I. Diakogiannis, "Deep learning on edge: Extracting field boundaries from satellite images with a convolutional neural network," *Remote Sens. Environ.*, vol. 245, p. 111741, 2020.

[5] A. Taravat, M. P. Wagner, R. Bonifacio, and D. Petit, "Advanced fully convolutional networks for agricultural field boundary detection," *Remote Sens.*, vol. 13, no. 4, p. 722, 2021.

[6] O. Oktay et al., "Attention u-net: Learning where to look for the pancreas," *arXiv preprint arXiv:1804.03999*, 2018.

[7] S. S. M. Salehi, D. Erdogmus, and A. Gholipour, "Tversky loss function for image segmentation using 3D fully convolutional deep networks," in *Int. Workshop Mach. Learn. Med. Imag.*, 2017, pp. 379–387.

[8] P. Tripathy, K. Baylis, K. Wu, J. Watson, and R. Jiang, "Investigating the Segment Anything Foundation Model for mapping smallholder agriculture field boundaries without training labels," *arXiv preprint arXiv:2407.01846*, 2024.

[9] H. Kerner, S. Sundar, and M. Satish, "Multi-region transfer learning for segmentation of crop field boundaries in satellite images with limited labels," *arXiv preprint arXiv:2404.00179*, 2024.

[10] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Med. Image Comput. Comput.-Assist. Interv.*, 2015, pp. 234–241.

[11] Z. Zhou, M. M. Rahman Siddiquee, N. Tajbakhsh, and J. Liang, "Unet++: A nested u-net architecture for medical image segmentation," in *Deep Learn. Med. Image Anal. Multimodal Learn. Clin. Decis. Support*, 2018, pp. 3–11.

[12] Z. Zhang, Q. Liu, and Y. Wang, "Road extraction by deep residual u-net," *IEEE Geosci. Remote Sens. Lett.*, vol. 15, no. 5, pp. 749–753, 2018.

[13] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.

[14] R. Yang, Z. U. Ahmed, U. C. Schulthess, M. Kamal, and R. Rai, "Detecting functional field units from satellite images in smallholder farming systems using a deep learning based computer vision approach: A case study from Bangladesh," *Remote Sens. Appl.: Soc. Environ.*, vol. 20, p. 100413, 2020.

[15] M. Wang, J. Wang, Y. Cui, J. Liu, and L. Chen, "Agricultural field boundary delineation with satellite image segmentation for high-resolution crop mapping: A case study of rice paddy," *Agronomy*, vol. 12, no. 10, p. 2342, 2022.

[16] F. I. Diakogiannis, F. Waldner, P. Caccetta, and C. Wu, "ResUNet-a: A deep learning framework for semantic segmentation of remotely sensed data," *ISPRS J. Photogramm. Remote Sens.*, vol. 162, pp. 94–114, 2020.

[17] M. Jong, K. Guan, S. Wang, Y. Huang, and B. Peng, "Improving field boundary delineation in ResUNets via adversarial deep learning," *Int. J. Appl. Earth Observ. Geoinf.*, vol. 112, p. 102877, 2022.

[18] G. M. Musyoka, "Automatic delineation of small holder agricultural field boundaries using fully convolutional networks," M.S. thesis, Univ. Twente, 2018.

[19] Z. Du, J. Yang, C. Ou, and T. Zhang, "Smallholder crop area mapped with a semantic segmentation deep learning method," *Remote Sens.*, vol. 11, no. 7, p. 888, 2019.

[20] M. Torre, B. Remeseiro, P. Radeva, and F. Martinez, "DeepNEM: Deep network energy-minimization for agricultural field segmentation," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 13, pp. 726–737, 2020.

[21] R. d'Andrimont et al., "AI4Boundaries: An open AI-ready dataset to map field boundaries with Sentinel-2 and aerial photography," *Earth Syst. Sci. Data*, vol. 15, no. 1, pp. 317–329, 2023.

HURA-Net: مدل نوین تشخیص مرز مزارع کشاورزی

مهدی علیزاده، پروین احمدی* و معصومه عظیم‌زاده

پژوهشگاه ارتباطات و فناوری اطلاعات، تهران، ایران.

ارسال ۲۰۲۵/۱۰/۰۴؛ بازنگری ۲۰۲۶/۰۵/۰۲؛ پذیرش ۲۰۲۶/۰۶/۱۵

چکیده:

تشخیص مرز مزارع یکی از وظایف کلیدی در کشاورزی نوین است که نقش مهمی در کشاورزی دقیق، مدیریت بهینه منابع و پایش کارآمد محصولات دارد. با وجود اهمیت این موضوع، بسیاری از مدل‌های یادگیری عمیق موجود به دلیل پیچیدگی مناظر، تفاوت وضوح تصاویر و وجود نویز در تصاویر سنجش‌ازدور، در تعیین دقیق مرز مزارع با محدودیت مواجه هستند. برای رفع این چالش‌ها، مدل HURA-Net را پیشنهاد می‌کنیم؛ یک چارچوب یادگیری عمیق که معماری‌های ++UNet، ResUNet و سازوکار توجه (Attention) را به‌صورت یکپارچه ترکیب می‌کند. این مدل با بهره‌گیری از استخراج ویژگی چندمقیاسی، یادگیری باقیمانده و تمرکز بر نواحی مهم تصویر، نقاط قوت هر معماری را تلفیق کرده و ضعف‌های آن‌ها را کاهش می‌دهد. همچنین، یک تابع هزینه بهبودیافته ارائه شده است که علاوه بر افزایش دقت تفکیک، مشکل عدم توازن کلاس‌ها را نیز برطرف می‌کند. نتایج آزمایش‌ها روی مجموعه‌ای متنوع از تصاویر ماهواره‌ای با وضوح بالا از مناطق مختلف ایران نشان می‌دهد که HURA-Net عملکردی برتر نسبت به مدل‌های پیشرفته موجود دارد. این مدل به Recall برابر با ۴۵.۸۵٪ (۱۵.۵۹٪ بهتر از ResUNet) و F1-score برابر با ۴۲.۶۲٪ (۷.۲۷٪ بالاتر از ResUNet) دست یافته و معیار جدیدی در دقت تشخیص مرز مزارع ارائه کرده است. همچنین، نتایج نشان می‌دهد که راهبردهای مؤثر افزایش داده، نقش مهمی در بهبود تعمیم‌پذیری مدل و مقابله با تغییرات نور، نوع محصول و شکل مزارع دارند. این پژوهش اهمیت طراحی معماری نوآورانه، تابع هزینه بهینه و آموزش کارآمد را در پیشرفت بخش‌بندی تصاویر سنجش‌ازدور نشان می‌دهد.

کلمات کلیدی: تشخیص مرز مزارع، بخش‌بندی تصویر، U-Net، تابع زیان Tversky.