XpressWeed: Meta-Inspired Few-Shot Adaption for Plant Weed Segmentation Using Texture Priors

Kiran K. Kethineni*, Rishi Raj Kanukuntla*, Saraju P. Mohanty* and Elias Kougianos†
*Department of Computer Science and Engineering, University of North Texas, USA
†Department of Electrical Engineering, University of North Texas, USA
Email: kirankumar.kethineni@unt.edu, rishirajkanukuntla@my.unt.edu, saraju.mohanty@unt.edu, elias.kougianos@unt.edu

***Abstract*—Weeds remain a persistent challenge in agriculture, competing with primary crops for resources. Recent advances in artificial intelligence have applied convolutional neural networks (CNNs) for weed detection and segmentation. However, building a single large-scale model capable of handling the vast variability in weed species is computationally infeasible and impractical. Furthermore, conventional object-centric detection and segmentation models struggle in real-world scenarios where leaves overlap, twist and exhibit considerable variations in shape and hue under changing illumination, causing traditional few-shot learning strategies to fail. To address these challenges, we introduce XpressWeed, a novel texture-prior driven segmentation framework that models plant leaves as textures, leveraging the fact that each species exhibits distinctive texture patterns which remain relatively invariant to shape deformations and occlusions. By pre-training a CNN segmentation model on a diverse collection of leaf textures, the framework learns discriminative representations that remain robust to lighting variations and morphological distortions. Subsequently, a Model-Agnostic Meta-Learning (MAML) strategy is employed to adapt the pre-trained model to novel weed species with only a few labeled samples, thereby retaining its discriminative capacity while enabling rapid transfer to unseen classes. The Fully Convolutional Network (FCN) model, pre-trained on 12 diverse plant textures, was successfully adapted to 4 previously unseen species, achieving 85% accuracy under a 12-shot learning setup using only 80 images while traditional MAML achieved 28% accuracy.**

***Keywords*—Artificial Intelligence; Texture-Based Learning; Prior; Model-Agnostic Meta-Learning (MAML); Few-Shot Learning; Image Segmentation; Model Adaptation.**

I. Introduction

Weeds are a major problem in agriculture, they grow along with primary crops and compete for essential resources such as nutrients, water, and sunlight, resulting in yield loss. Effective weed management is therefore critical for sustainable farming and food security. As a result, automated weed detection and segmentation have attracted significant interest in recent years [1]. Machine learning (ML), particularly convolutional neural networks (CNNs), has been widely adopted for plant and weed classification, detection, and segmentation. In addition, few-shot learning strategies have been explored to enable adaptation to new weed species with limited data availability.

Despite these advances, several challenges persist. Weed species exhibit considerable variability across geographical regions, environmental conditions, and evolutionary adaptations [2]. Constructing a single large-scale model that can account for all species and conditions is both computationally infeasible and impractical due to the need for an exhaustive dataset. Conventional object-centric approaches to detection and segmentation often treat plants as rigid objects and rely on learning shape or outline-based representations for each crop. In natural farm settings, however, leaves overlap, twist, and vary in orientation, creating a wide range of shapes that make such object-centric modeling unreliable. In addition, lighting conditions introduce further complexity by altering visual features and color hues. Collectively, these factors cause traditional few-shot learning strategies to fail, as they cannot rely on a limited number of images to sufficiently capture such extensive variations [3].

A key observation is that plant leaves are more naturally represented as textures rather than rigid objects [4]. Texture patterns are relatively invariant to shape deformation, occlusion, and illumination changes, making them more robust for modeling under diverse field conditions. Furthermore, in cases where texture similarity exists across species, fine-grained features such as leaf edges provide additional discriminative power. This shift in perspective from objects to textures offers a promising foundation for more reliable weed identification in complex agricultural environments. Building upon this principle, the present work "Xpress Weed" introduces a framework that leverages texture priors to improve adaptability to new weed species, which will be described in detail in the subsequent sections.

The organization of rest of the article is as follows: Section II introduces the proposed solution and its novel contributions, while Section III offers a brief review of related research. The proposed method is detailed in Section IV, with experimental results presented in Section V. The article concludes with Section VI, which also outlines potential directions for future research.

II. Novel Contributions of the Current Paper

A. Problem Statement

The main difficulty in weed identification is handling the large variations that occur in real farm conditions. Building very large object-based models is not practical, and few-shot methods fail when the small number of training images cannot cover changes in leaf shape, orientation, and lighting. Therefore, there is a need for a representation that can naturally handle these variations while still learning from only a few examples of new weed species.

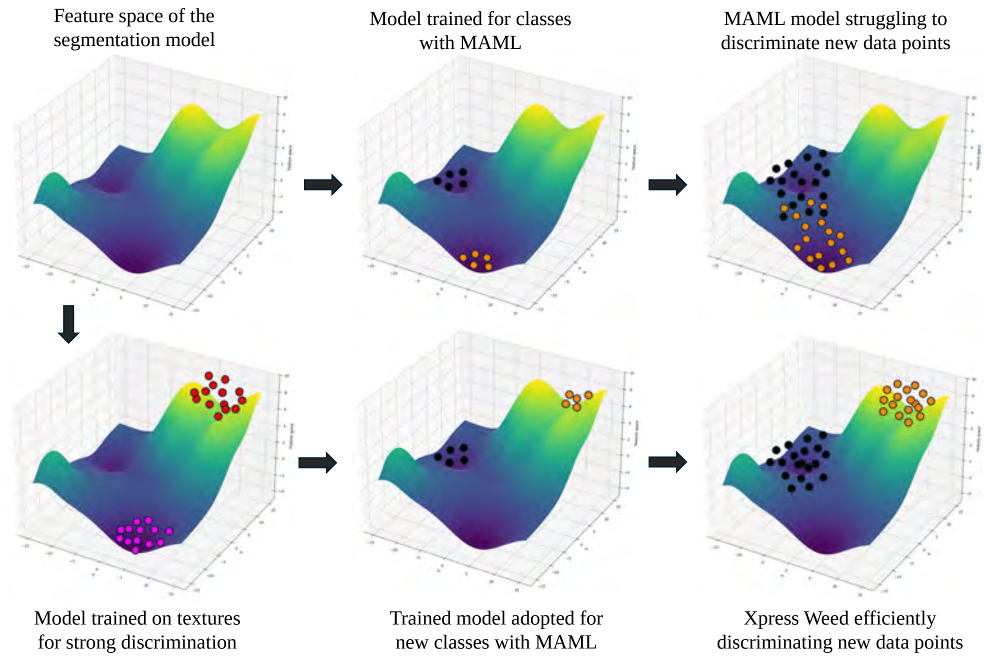

Figure 1: Illustration of feature landscapes showing robustness from texture priors and adaptation with MAML.

B. Proposed Solution of the Current Paper

We propose XpressWeed, a texture-prior driven segmentation framework that combines conventional supervised learning with meta-learning to achieve robust and resource-efficient weed segmentation. The approach consists of two main stages: (1) pre-training on diverse leaf textures using standard supervised training, and (2) adapting to new weed species using Model-Agnostic Meta-Learning (MAML) [5].

1) Texture Representation: By representing plants as textures, the model captures features that are invariant to deformation, occlusion, and illumination, making it more robust and reliable for segmentation in the field [6]. Importantly, this initial stage employs conventional supervised training, since applying meta-learning at scale would necessitate repeated second-order differentiation over a large number of training images, making the process computationally impractical.

2) Role of Meta-Learning: While texture priors offer stability, new weed species inevitably emerge across seasons and regions. Retraining or naive fine-tuning for each case is impractical and often fails with very limited samples. To address this, we apply MAML only at the adaptation stage, avoiding the heavy cost of meta-training on large datasets. This selective use allows the model to update quickly from a few labeled examples, while episodic tasks encourage implicit separation of unseen classes. In doing so, MAML complements the texture backbone with flexibility, ensuring the framework remains both efficient and adaptable.

3) Texture Priors with Meta-Learning: The strength of XpressWeed lies in the combination of texture priors and meta-learning. The supervised pre-training stage equips the model with robust, invariant texture features that generalize across deformations and illumination conditions. When adapting to a new species, MAML leverages this prior knowledge as a strong initialization, ensuring that even a few labeled samples can be integrated effectively into the learned feature space.

This dual mechanism prevents overfitting to the limited examples, preserves the robustness gained during pre-training, and accelerates adaptation. As illustrated in Fig. 1, the texture prior enables the model to settle in regions of the feature landscape where class discrimination is already strong. Starting MAML from these points ensures that even with very few images, the model maintains high separability and generalizes effectively to new data points. In summary, the texture prior provides stability, while MAML provides flexibility, together ensuring that the model remains robust to variations and retains high discriminative power when adapting to new weed classes.

C. Novelty and Significance of the Proposed Solution

The proposed XpressWeed framework introduces the following novelties and significance:

- Texture-Based Representation: Reframes weed segmentation by modeling leaves as textures rather than objects, enabling robustness to deformation, occlusion, and illumination.
- Separation of Training and Adaptation: Employs supervised learning for texture pre-training and reserves meta-learning only for adaptation, ensuring computational feasibility and efficiency.
- Discriminative Few-Shot Adaptation: Utilizes MAML in texture space to enforce contrastive separation, allowing robust adaptation to new classes even with limited and noisy samples.

Table I: Relevant literature in Weed Segmentation and Few-Shot Learning.

Related Work	Year	Method Adopted	Remark
Syed and Suganthi [7]	2023	Active contour segmentation by fuzzy-transform	Minimal data requirement but lacks deep features and scalability.
Naik [8]	2024	Fully supervised CNN, U-Net	Effective with large data but static and cannot adapt to new classes.
Hu et al. [9]	2022	Fully supervised CNN, U-Net++	Effective with large data but static and cannot adapt to new classes.
Li et al. [10]	2023	Fully supervised CNN, PSP Net	Effective with large data but static and cannot adapt to new classes.
Shorewala et al. [11]	2020	Semi-supervised binary classification	Reduces annotation cost but lacks pixel-level precision.
Kethineni et al. [12]	2023	Semi-supervised binary classification	Reduces annotation cost but fails to differentiate if crops overlap .
Amac et al. [13]	2021	Self-supervised meta-learning	Trains with pseudo-masks but depends heavily on saliency quality.
Cao et al. [14]	2019	FCN guided by Meta learner	Learns general meta-learner but cannot handle variations in plant images.
J.Logeshwaran et al. [15]	2024	Meta-stacking (transfer + meta-learning)	Applies meta-learning throughout full training with multiple datasets, effective but computationally heavy
XpressWeed	2025	Supervised texture pre-training + MAML adaptation	Resource-efficient framework, supervised pre-training captures diverse variations, while MAML enables few-shot adaptation to new classes.

- Resource-Efficient Framework: Provides a scalable solution tailored for real-world farming by achieving high accuracy with minimal data, addressing the challenge of evolving weed species.

III. Related Prior Works

Early research in weed segmentation relied on classical image processing approaches, such as fuzzy-transform-based active contour models [7], which made use morphological properties to detect weeds. With the emergence of deep learning, fully supervised CNNs like U-Net [8], U-Net++ [9] and PSP Net [10] achieved strong performance in agricultural datasets. However, these models were data-hungry and static, limiting their ability to adapt to new weed species. To reduce annotation effort, semi-supervised and hybrid approaches, such as ResNet–SVM [11] hybrids, saliency-driven methods like MaskSplit [13] and Dynamic Time Warping on profile plots of crops [12] were introduced, they often compromised on precision or depended heavily on pseudo-label quality.

More recently, meta-learning frameworks such as Meta-Seg [14], enabled few-shot segmentation by training models to adapt to unseen classes with limited samples. However, when applied to plant imagery, these methods struggle to cope with the high variability in leaf orientation, overlap, and illumination, which limits their effectiveness in real agricultural conditions. Transfer-learning-based meta-learning approaches, such as the meta-stacking [15] framework proposed for medical image segmentation, integrate pretrained models into the meta-learning process to improve adaptation. While effective, these methods apply meta-learning throughout the full training pipeline, requiring multiple large datasets and high computational cost.

The proposed Xpress Weed addresses these gaps by introducing a texture-based prior with supervised pre-training and applying meta-learning only during adaptation, enabling both stability and flexibility in weed segmentation with minimal data. A brief overview of the related works is summarized in Table I.

IV. Proposed Method

The proposed XpressWeed framework combines supervised pre-training with meta-learning adaptation in order to achieve robust weed segmentation under limited data availability. The method represented in Fig. 2 is designed to learn stable texture representations through supervised training and then adapt these features efficiently to new classes using MAML. The framework separates large-scale representation learning from lightweight adaptation, thereby achieving both robustness and computational efficiency.

A. Network Architecture

The backbone network presented in Fig. 3 is an encoder–decoder design inspired by U-Net [16] but modified to emphasize texture representation and multi-scale discrimination. The encoder progressively downsamples the input $(512 \times 768 \times 3)$ through stacked convolutional blocks, compressing it to a compact bottleneck of $(8 \times 12 \times 512)$. This embedding captures high-level texture features that remain stable under leaf deformation, orientation, and lighting variations. Intermediate feature maps at multiple resolutions are preserved during downsampling to support later fusion.

The decoder reconstructs spatial resolution by progressively up-sampling the bottleneck features. Unlike architectures that rely on transposed convolutions or bilinear refinement, the upsampled feature maps here are directly concatenated with the saved encoder feature maps at corresponding scales. This multi-scale fusion provides the classifier access to both global texture context and fine-grained structural details, such as leaf edges and venation. Thus, boundary refinement is achieved implicitly through feature fusion rather than computationally intensive deconvolution operations. Final classification is performed per pixel using a pointwise (1×1) convolution applied to the fused feature maps, producing dense segmentation masks informed by features across all scales.

B. Supervised Pre-training

The objective of the pre-training stage is to make the network a strong discriminator while keeping the architecture lightweight enough for efficient meta-learning adaptation. Since

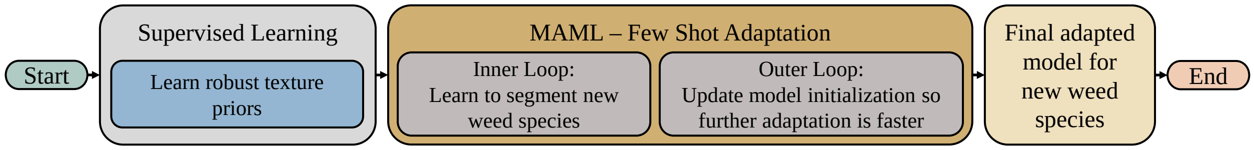

Figure 2: Workflow of the proposed XpressWeed framework.

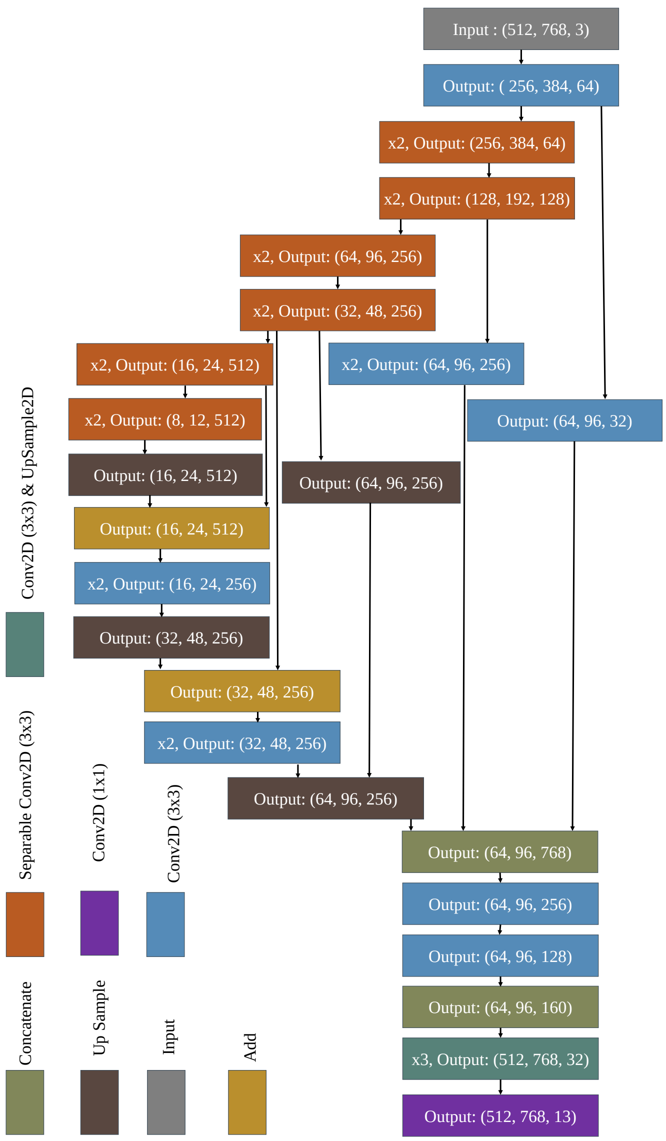

Figure 3: Architecture of the proposed CNN model.

the adaptation stage relies on gradient-based updates, a very deep or parameter-heavy model would make MAML computationally expensive and slow to converge. To avoid this, the architecture was deliberately kept simple, focusing on essential convolutional blocks, skip connections, and multi-scale fusion rather than complex modules.

The loss function was also chosen with the same principle in mind. Instead of adopting compound or adversarial objectives, we employ a weighted cross-entropy (WCE) loss, which is both effective and computationally efficient. The weighting compensates for severe class imbalance in agricultural datasets, where crop pixels dominate and weed classes may be underrepresented. By pre-training with WCE, the encoder learns to separate textures robustly under varying orientations, occlusions, and lighting, while the overall network remains lightweight enough for rapid few-shot adaptation with MAML.

C. *MAML-Based Adaptation*

After supervised pre-training, the model must adapt to unseen weed species with very limited labeled samples. Traditional fine-tuning in such low-data settings often leads to overfitting, while retraining a large model from scratch needs lot of images and computation. To address this, we use Model-Agnostic Meta-Learning (MAML) during the adaptation stage.

The key advantage of MAML is that it learns an initialization that can be quickly updated with just a few gradient steps. During adaptation, the model is exposed to small support and query sets. The support set provides the initial update, and the query set ensures that the updated parameters generalize effectively. This process forces the model to learn parameters that are both discriminative and flexible, making it suitable for few-shot adaptation.

Another reason for selecting MAML is its implicit contrastive behavior [17] arising from its *n*-way episodic training. In each episode, the model is exposed to a support set containing samples from multiple classes and must adapt its parameters to discriminate among them. The subsequent evaluation on the query set forces embeddings of the same class to cluster together while pushing different classes apart. This *n*-way setup effectively mimics contrastive learning, but without requiring explicit contrastive loss functions or additional modules, thereby keeping the adaptation pipeline simple.

By combining lightweight supervised pre-training with MAML, the framework maintains efficiency while achieving high discriminative power. The pretrained texture priors provide stability, while MAML adds the flexibility to incorporate new species with minimal labeled data. Since meta-learning is applied only during the adaptation stage rather than across the full pipeline, XpressWeed avoids heavy second-order optimization on large datasets while still retaining the benefits of few-shot adaptation, making it both practical and resource-efficient.

V. Experimental Verification

The proposed model was first trained in a supervised setting using a dataset derived from PlantVillage [18]. Healthy leaves from twelve species were arranged into collages, including overlapping leaves and background-removed variants, to introduce mixed textures and sharp boundaries as shown in Fig. 4. This design was intended to simulate real conditions (e.g., grasses) where boundaries are difficult to separate.

Figure 4: Dataset used for supervised training of texture priors.

The model achieved 96% training accuracy, 92% validation accuracy, a mean IoU of 92%, and mAP@[0.50:0.95] of 0.90. Due to pixel-level class imbalance, we present per-class PR curves in Fig. 5, which provide a more reliable assessment than accuracy or IoU alone. The curves demonstrate that the model maintains strong discrimination across all classes.

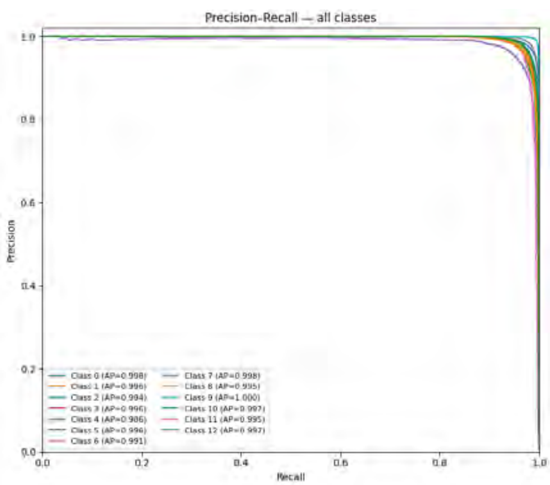

Figure 5: Precision–Recall curves of the supervised model.

To evaluate performance in real-world conditions, we used a farmland dataset [19] containing six classes: background, primary crop (soybean, previously encountered as a texture prior), and four weed species (unseen during training). Each image contained multiple crop types, and 50 images were annotated with corresponding masks, as shown in Fig. 6. From this dataset, 20 high-resolution images (2112×1600) containing more species of crops in each were used for MAML adaptation. From these, we cropped 80 patches to match the network's input size. of Since crop distributions in real farmland cannot be controlled, the prepared dataset is also imbalanced, the class distribution is illustrated in Fig. 7. Tasks were constructed with 3-way classification, 12 support samples per class, and 10 query samples per class, yielding a total of 36 support and 30 query images per episode, with training performed for a maximum of 25 epochs.

Figure 6: Dataset used for MAML adaptation.

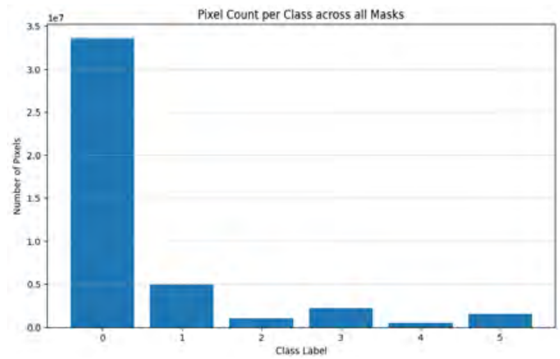

Figure 7: Pixel-wise class distribution in the MAML dataset.

The adapted model achieved 92% accuracy on the adaptation images and 85% accuracy on held-out validation images, as depicted in Fig. 8 and Fig. 9, demonstrating generalization beyond the support set while traditional MAML achieved 28% accuracy. Given the significant class imbalance, we report per-class PR curves to better assess robustness. As shown in Fig. 10, background (AP = 0.9973) and primary crop (AP = 0.9807) attained the highest areas under the PR curve. Among weed classes, Lavhala (AP = 0.7747) and Obscure Morning Glory (AP = 0.7614) performed comparatively well, while Kena (AP = 0.6966) was moderate. In contrast, Lamber (AP = 0.4869), which had the fewest annotated pixels, exhibited the weakest performance.

The overall macro mAP was 0.783, with a micro AUPRC of 0.962. These results highlight that MAML enables effective adaptation from limited support data, though extremely under-represented classes remain challenging.

VI. Conclusion

This article demonstrated the capability of a meta-learning framework to adapt effectively to real-world farmland images when supported by texture priors. The model achieved strong performance on support images and generalized well to un-

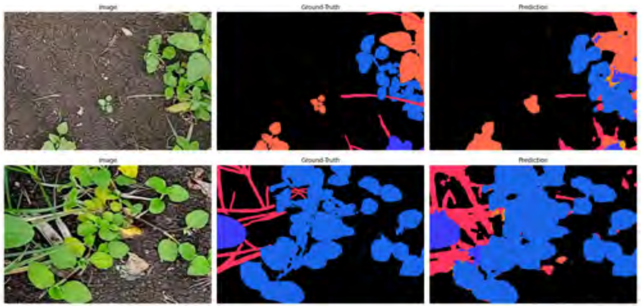

Figure 8: Training results of the MAML-adapted model.

Figure 9: Validation results of the MAML-adapted model.

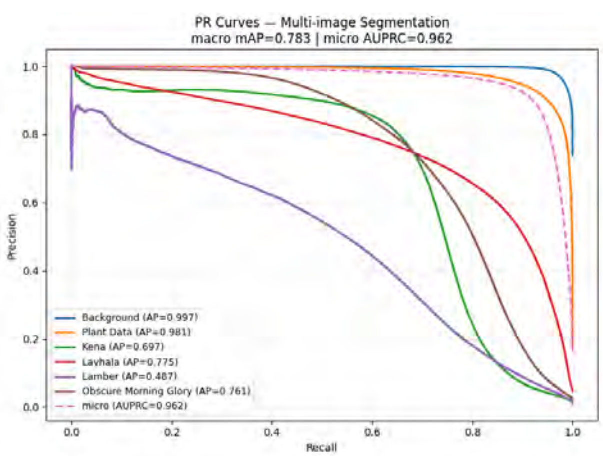

Figure 10: Precision–Recall curves of the adapted model.

seen validation images, with per-class PR curves highlighting robustness across crop and weed classes. Although one underrepresented weed class exhibited weaker performance, the overall results confirm that MAML-based adaptation provides a viable path for segmentation in precision agriculture, where labeled data is limited and field conditions are variable.

Furthermore, integrating such adaptive segmentation models with route-optimization frameworks like SprayCraft [20] would enable not only accurate detection of emerging weed species but also their precise, variable-rate spraying. By linking robust few-shot adaptation with optimized actuation paths, this combined system can minimize chemical use, reduce operational costs, and deliver scalable, resource-efficient weed management in real farmland.

REFERENCES

[1] G. R. Coleman, A. Bender, K. Hu, S. M. Sharpe, A. W. Schumann, Z. Wang, M. V. Bagavathiannan, N. S. Boyd, and M. J. Walsh, "Weed detection to weed recognition: reviewing 50 years of research to identify constraints and opportunities for large-scale cropping systems," *Weed Technology*, vol. 36, no. 6, pp. 741–757, 2022.

[2] D. R. Clements and V. L. Jones, "Ten ways that weed evolution defies human management efforts amidst a changing climate," *Agronomy*, vol. 11, no. 2, 2021.

[3] T. Qian, Y. Liu, S. Lu, L. Li, X. Zheng, Q. Ju, Y. Li, C. Xie, and G. Li, "Cucumber leaf segmentation based on bilayer convolutional network," *Agronomy*, vol. 14, no. 11, 2024.

[4] H. Kebapci, B. Yanikoglu, and G. Unal, "Plant image retrieval using color, shape and texture features," *The Computer Journal*, vol. 54, no. 9, pp. 1475–1490, 2011.

[5] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proc. International conference on machine learning*, 2017, pp. 1126–1135.

[6] Y. Zhang and M. A. Mazurowski, "Convolutional neural networks rarely learn shape for semantic segmentation," *Pattern Recognition*, vol. 146, p. 110018, 2024.

[7] Z. Syed and K. Suganthi, "Crop and weed semantic segmentation using a fuzzy transform based active contour model," in *Proc. International Conference on Innovative Data Communication Technologies and Application*, 2023, pp. 587–591.

[8] K. J. Naik, "Deep learning based segmentation of weed images in crop fields using u-net," in *Proc. 3rd International Conference for Advancement in Technology*, 2024, pp. 1–6.

[9] X.-Z. Hu, W.-S. Jeon, and S.-Y. Rhee, "Sugar beets and weed detection using semantic segmentation," in *Proc. International Conference on Fuzzy Theory and Its Applications*, 2022, pp. 1–4.

[10] X. Li, F. Duan, M. Hu, J. Hua, and X. Du, "Weed density detection method based on a high weed pressure dataset and improved psp net," *IEEE Access*, vol. 11, pp. 98 244–98 255, 2023.

[11] S. Shorewala, A. Ashfaque, R. Sidharth, and U. Verma, "Weed density and distribution estimation for precision agriculture using semi-supervised learning," *IEEE Access*, vol. 9, pp. 27 971–27 986, 2021.

[12] K. K. Kethineni, A. Mitra, S. P. Mohanty, and E. Kougianos, "Weedout: An autonomous weed sprayer in smart agriculture framework using semi-supervised non-cnn annotation," in *Proc. Internet of Things. Advances in Information and Communication Technology*, 2023, pp. 415–423.

[13] M. S. Amac, A. Sencan, O. B. Baran, N. Ikizler-Cinbis, and R. G. Cinbis, "MaskSplit: Self-supervised Meta-learning for Few-Shot Semantic Segmentation," *arXiv preprint arXiv:2110.12207*, 2021. [Online]. Available: https://arxiv.org/abs/2110.12207

[14] Z. Cao, T. Zhang, W. Diao, Y. Zhang, X. Lyu, K. Fu, and X. Sun, "Meta-seg: A generalized meta-learning framework for multi-class few-shot semantic segmentation," *IEEE Access*, vol. 7, pp. 166 109–166 121, 2019.

[15] J.Logeshwaran, N. Yuvaraj, V. Sharma, R. P. Shukla, and D. Kumar, "A meta learning approach for improving medical image segmentation with transfer learning," in *Proc. International Conference on Recent Innovation in Smart and Sustainable Technology*, 2024, pp. 1–6.

[16] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *Proc. Medical Image Computing and Computer-Assisted Intervention*, 2015, pp. 234–241.

[17] C.-H. Kao, W.-C. Chiu, and P.-Y. Chen, "Maml is a noisy contrastive learner in classification," 2022. [Online]. Available: https://arxiv.org/abs/2106.15367

[18] A. Ali, "Plantvillage dataset," 2019, last accessed April 25 2024. [Online]. Available: https://www.kaggle.com/datasets/abdallahalidev/plantvillage-dataset

[19] S. Shinde and V. Attar, "An indian annotated weed dataset for computer vision tasks in precision farming," *Data in Brief*, vol. 61, p. 111691, 2025.

[20] K. K. Kethineni, S. P. Mohanty, E. Kougianos, S. Bhowmick, and L. Rachakonda, "SprayCraft: Graph-Based Route Optimization for Variable Rate Precision Spraying," 2024. [Online]. Available: https://arxiv.org/abs/2412.12176