

On the Transferability of Agricultural Weed Detection Under Cross-Field Distribution Shift

Nikhilesh Prabhakar¹, Pranuthi Tenali¹, Wilfredo Abudeye Fernandez²,
Shekhar Borah², Athresh Karanam¹, Erik Blasch³, Prabha Sundaravadivel²,
and Sriraam Natarajan¹

¹ The University of Texas at Dallas

{nikhilesh.prabhakar, pranuthi.tenali, athresh.karanam,
sriraam.natarajan}@utdallas.edu

² The University of Texas at Tyler

wabudeyefernandez@patriots.uttyler.edu
{psundaravadivel, sborah}@uttyler.edu

³ Air Force Research Lab

erik.blasch.1@us.af.mil

Abstract. Accurate agricultural weed detection in real-world field conditions is essential for precision agriculture, enabling targeted intervention and reducing yield loss. Recent work has reported strong detection performance from UAV-based imagery across a range of crops, yet existing approaches evaluate within a single crop and field, leaving practitioners with little evidence that a model trained on one crop will generalize to a new field or crop type. In this work, we characterize where cross-dataset weed-localization performance degrades and which modeling choices recover it, reducing the need to relabel every new deployment field. We introduce a newly collected and annotated UAV image dataset for weed detection in cotton fields and use it alongside an existing soybean dataset collected under a similar protocol. Using these datasets, we evaluate the performance of several strategies for transferring a detector trained on one crop to another, comparing unsupervised domain adaptive object detection (DAOD) against pretraining on a domain-adjacent source dataset followed by few-shot fine-tuning on the target dataset. Our analysis spans target-domain label budgets from zero to the full target dataset, characterizing the trade-off between adaptation strategy and annotation effort. We find that few-shot fine-tuning with as few as 25 labeled target examples outperforms unsupervised DAOD in our cross-crop comparison, suggesting that source domain selection combined with modest target supervision is more productive than algorithmic sophistication in adaptation.

Keywords: Weed detection · Domain adaptation · Cross-Dataset transfer · Distribution shift · Object detection.

1 Introduction

Weeds are a major contributor to yield loss in row crop agriculture, competing with crops for nutrients, water, and light [19]. Timely and accurate weed detection is therefore central to effective field management and to reducing the indiscriminate use of herbicides, which carries substantial economic and environmental costs [20]. Unmanned aerial vehicles (UAVs) have emerged as a practical platform for scaling weed monitoring to entire fields [24, 18]. They offer high-resolution imagery at low operational cost and can be deployed with minimal infrastructure relative to ground-based systems.

Broadly, UAV based weed detection is an instance of the Dynamic Data-Driven Applications Systems (DDDAS) paradigm [8], in which the vehicle flies over the entire field while the detector determines the position and the class of weeds. The detector outputs could then be used to determine the parts of the field to concentrate on, plan optimal flight paths, and design effective monitoring strategies.

Modern weed detection pipelines on UAV imagery rely predominantly on deep learning approaches [1, 25]. The dominant paradigm involves training an object detector on a labeled dataset collected from the field of interest. While this approach can yield strong detection performance, it imposes high annotation costs. Producing high-quality bounding box labels requires expert knowledge of crop and weed species, and the resulting datasets are often narrow in scope, covering a single location, growth stage, or weed species [13]. Models trained under these conditions tend to generalize poorly when deployed in new fields, seasons, or to different crops, which limits the practical utility of any single labeled dataset [2, 13].

Domain Adaptive Object Detection (DAOD) offers a principled approach to this challenge by leveraging labeled data from a related source domain to improve detection performance on an unlabeled or sparsely labeled target domain. Recent frameworks such as ALDI [15] have standardized benchmarking across DAOD methods and demonstrated strong results on established benchmarks. However, the application of DAOD to UAV-based agricultural imagery in general and weed detection in particular remains underexplored. Existing work in this direction has focused primarily on weather-induced domain shift [27], leaving the more common practical scenario of cross-field, cross-crop transfer largely unexamined.

In this work, we address this gap by presenting **an empirical study of cross-crop transfer for UAV-based weed detection**. We introduce a newly collected and annotated UAV image dataset for weed detection in cotton fields and use it alongside a previously published soybean dataset collected under a similar protocol [2]. Using these datasets we evaluate a range of strategies for transferring a detector trained on one crop to another, comparing unsupervised DAOD against a simpler alternative: pretraining on a domain-adjacent source dataset and fine-tuning with a small number of labeled examples from the target domain. We thus study the data-assimilation component of the DDDAS paradigm. Our central finding is that this latter strategy is remarkably effective, with fine-tuning a detector trained on a domain-adjacent source using as few as

$k=25$ labeled target examples outperforming DAOD on the same target. A key aspect of our setup is that we examine performance across a spectrum of target domain label budgets, from $k=0$ (no examples from target domain) up to the full target dataset, which allows us to characterize the trade-off between adaptation strategy and annotation effort.

We make the following key contributions: (1) We present a newly collected and annotated UAV image dataset for weed detection in cotton fields, providing a real-world benchmark for evaluating cross-crop transfer in agricultural object detection settings. (2) We design and conduct a systematic empirical comparison of unsupervised DAOD against domain-adjacent pretraining with few-shot fine-tuning across a range of target domain label budgets.

Our analysis suggests that in settings where small amounts of target-domain supervision are practical to obtain, such as cross-crop transfer in UAV-based weed detection, **source domain selection combined with few-shot fine-tuning may be more productive in practice than algorithmic sophistication in adaptation.**

2 Background and Related Work

2.1 Weed Detection in UAV Imagery

UAVs have become an increasingly common platform for weed detection in precision agriculture, offering high-resolution imagery at field scale at a low operational cost. Over the last decade, deep learning approaches have become predominant for processing this imagery [1, 25]. A variety of architectures have been explored across crops including soybean [2] and cotton [9]. YOLO-based detectors [21] have emerged as a particularly common choice due to their balance of accuracy and real-time inference [2, 23].

A persistent challenge across this line of work is the visual similarity between weed and crop species from overhead viewpoints. This makes both inter-class separation and model generalization difficult [1]. A direct consequence of this lack of generalization is the aspect of poor transferability of models across fields, seasons, and crop types. Weed datasets are typically collected with limited variability in conditions, often from a single location, growth stage, and weed species, which hinders generalization to new field conditions and different crops [13]. Domain adaptation has been identified as a promising yet underexplored direction for weed detection [13]. Prior work in this direction has primarily examined ground-level crop-weed imagery [14], and its application to UAV-based detection pipelines remains limited. In this work, we aim to address this gap by presenting an empirical analysis of domain adaptive object detection methods on a combination of UAV-acquired weed detection datasets, including a newly collected cotton field dataset.

2.2 Domain Adaptive Object Detection

Domain Adaptive Object Detection addresses the performance degradation that arises when a detector trained on a labeled source domain is evaluated on a

(potentially unlabeled) target domain with a different data distribution. Existing approaches can be broadly classified into two categories - alignment-based methods and self-distillation or self-training methods.

Alignment-based methods try to reduce the discrepancy between source and target domains by learning domain-invariant feature representations [11]. Refinements in this direction include Strong-Weak Distribution Alignment [22], which argues for weak global alignment while enforcing stronger local alignment on detection-relevant regions, and a suite of methods that further incorporate multi-level [5], class-aware [26], scale-aware [6], foreground/background-aware [12] among others, aiming to align the aspects of the representation most relevant for object recognition rather than the entire feature distribution.

Self-distillation/self-training approaches based on the mean-teacher paradigm generate pseudo-labels on the target domain to supervise adaptation, with the teacher maintained as an exponential moving average of the student [10, 17]. Since pseudo-label quality bounds the effectiveness of this paradigm, refinements have largely focused on producing or exploiting cleaner supervision through strong-weak augmentation [17], uncertainty-aware filtering of unreliable predictions [4], and contrastive objectives that derive useful signal even from noisy bounding boxes [3] among others.

3 Datasets

Our evaluation uses four datasets of aerial and ground weed imagery that vary in crop, weed species, altitude, and sensor. The **soybean** dataset [2] was collected from a weed-infested soybean field in Stoneville, Mississippi: 637 high-resolution RGB images captured by a UAV at 25 ft, annotated with bounding boxes for pigweed and tiled into 1628 images of size 640×640 . We collected a new **cotton** dataset from weed-infested fields in the same region following the protocol of [2]: an IF1200A UAV flown at 13–15 mph and 25 ft (ground resolution ≈ 2.13 cm/pixel) over a ~ 2.3 km survey, yielding 590 images manually annotated with LabelImg and tiled into 5706 images of size 640×640 . The remaining two are public: **CoFly-WeedDB** [16] (low-altitude drone, 207 images) and **Cotton-WeedDet12** [7] (ground-level cotton-weed, 5648 images, 12 classes remapped to a single weed class). The supplementary material gives the collection and reports the splits and box scale (Table A2), which span nearly three orders of magnitude and are not equalized by preprocessing, and shows representative tiles (Figure A1).

4 Methodology

Weed detection requires carefully collecting and manually annotating large image datasets which is labor-intensive. Consequently, deploying a pretrained detector on the target field is highly desirable. However, the source and target fields are often very different. A key question in cross-field weed detection is how much target-field supervision is required to achieve satisfactory performance. Formally,

given pretrained detector on the source domain, and n examples from the target training set, of which k are labeled, the task is to adapt the source detector on the target domain. To systematically study this, we analyze the following techniques on the transferability among the 4 different domains.

4.1 Domain Adaptive Object Detection

We compare the two dominant families of domain-adaptive object detection.

Adversarial feature alignment (Align) adds image-level and instance-level domain classifiers trained through a gradient-reversal layer, encouraging domain-invariant features, without requiring a teacher.

Self-distillation (Distill) instead maintains an exponential moving average (EMA) copy of the detector as a teacher that produces soft targets on the unlabeled target domain, combined with masked-image consistency (MIC). It doesn't use adversarial alignment. In our evaluation, gradient clipping (max norm 1.0) is applied to the distillation runs for numerical stability.

We use ALDI (Align-and-Distill) [15], a unified framework for domain-adaptive object detection, to run our adaptation experiments. ALDI implements both adaptation families above on a common Faster R-CNN backbone and under a single training recipe, standardizing the components that otherwise confound cross-method comparison: strong augmentation, an EMA teacher, and training from a source-only checkpoint. Adversarial alignment and self-distillation therefore differ only in their adaptation objective, which lets us attribute differences in target performance to the adaptation mechanism rather than to incidental training choices. We use this shared recipe for every adaptation run reported below.

4.2 Few-Shot Fine Tuning

As a supervised alternative for adaptation, we fine-tune the source-only detector on a small number of labeled *target* tiles. We vary the budget $K \in \{5, 10, 25\}$ labeled images and, for each K , draw $M=5$ independent random subsets to capture the variance from which images are labeled. Fine-tuning is standard supervised training, with out any adaptation losses or unlabeled target data, for 1.5k iterations at learning rate 0.005, starting from the source checkpoint as pretrained weights.

5 Experiments

In our empirical evaluation, we aim to answer the following questions:

- Q1.** Does a detection model trained on one field effectively transfer to another?
- Q2.** Does unsupervised domain adaptation help improve the transfer?
- Q3.** When domain adaptation falls short, does few-shot learning work better?

5.1 Experimental Setup

Training: All detectors are Faster R-CNN with a ResNet-50-FPN backbone, initialized from COCO pretraining, inside the ALDI framework [15]. We use ALDI’s fair-baseline recipe throughout: strong augmentation, random erasing, and an exponential moving average (EMA) of the weights for evaluation. Source-only models train for 8k iterations (SGD, batch size 8, base learning rate 0.01, step decay at 7k). Every domain adaptation run resumes from the source-only checkpoint and trains for a further 10k iterations on labeled source plus unlabeled target. All datasets reach the detector at a common 640px input: the UAV surveys are tiled, and the others are resized to a 640px shorter edge.

Input Normalization: All datasets are presented to the detector at a common 640px input. The high-resolution UAV surveys (cotton, soybean; $\approx 5184 \times 3888$ px), in which weeds are small relative to the frame, are tiled to non-overlapping 640×640 px tiles; the remaining datasets are resized to a 640px shorter edge.

Evaluation: All numbers are mAP@50 on the held-out target *test* split. For each run, we select the checkpoint with the best target-validation mAP@50 and report its target-test score. The source-only baselines are instead selected on the source validation split. Domain-adaptation and few-shot results are reported as mean $\pm$ standard deviation over multiple training seeds (and, for few-shot, over random label draws). The training seed varies data ordering, augmentation sampling, and RPN/ROI mini-batch composition. The train/val/test partition is held fixed across all runs.

5.2 Results

We organize the results around the three questions posed above. We first report the quantitative results and then identify the conditions under which it holds and the factors that predict when it breaks.

Q1: Cross-field transfer To answer **Q1**, by default, a weed detector trained on one source does not always effectively transfer across fields, and its in-domain accuracy is not an indication of how far it will fall. (Table A1) reports the mAP@50 results for every direct model transfer for every source-to-target pair. For a couple of datasets, Soybean and CWDet12, the in-domain accuracy is strong. The numbers off the diagonal show the transfer mAP@50 scores with no supervised fine-tuning. Most zero-shot transfer collapses to single-digits. The exceptions to this are the transfer between Cotton $\rightarrow$ Soybean (57.8) and Cotton $\rightarrow$ CWDet12 (31.8), the targets that share the species of weed or scale with the source. Pairs that differ in weed-type and scale of detection drop to near zero.

Q2: Unsupervised domain adaptation. Table 1 reports *mAP*@50 scores for source-only, Distill, and Align across our source-target pairs. The results are highly variable. DAOD substantially improves over source-only in some cases (cotton $\rightarrow$ CWDet12; CoFly $\rightarrow$ soybean) but reduces performance in others, including the cotton $\rightarrow$ soybean comparison most relevant to our setting, where Distill and Align fall to 50.3 and 42.8 from a source-only baseline of 57.8. Even

Table 1. Domain-adaptation comparison for cross-field weed detection (target-test mAP₅₀, checkpoints selected on the target validation split). *In-Domain* is the target’s in-domain supervised mAP₅₀; *Source-only* is the source-trained detector applied zero-shot to the target. Distillation and adversarial alignment are reported as mean $\pm$ std over n seeds (single-seed cells show no interval); Δ is the change relative to source-only. Adversarial alignment on the 108-image CoFly source diverges (Inf/NaN) at the default learning rate and is omitted ($\dagger$). Adversarial alignment was not run for the CottonWeedDet12 target (--); we report distillation only.

Source $\rightarrow$ Target	In-Domain	Source-only	Distill (n)	Align (n)
<i>Target: soybean</i>				
cotton $\rightarrow$ soybean	76.8	57.8	50.3 ± 2.5 (5)	42.8 ± 4.5 (5)
cofly $\rightarrow$ soybean	76.8	4.1	16.7 ± 2.1 (3)	div. †
cwdet12 $\rightarrow$ soybean	76.8	6.0	9.6 ± 1.0 (3)	6.5 (1)
<i>Target: cotton</i>				
soybean $\rightarrow$ cotton	37.4	7.5	8.2 ± 0.4 (4)	7.4 ± 0.6 (5)
cofly $\rightarrow$ cotton	37.4	1.3	1.6 (1)	0.7 (1)
cwdet12 $\rightarrow$ cotton	37.4	1.5	3.9 (1)	2.5 (1)
<i>Target: CottonWeedDet12</i>				
cotton $\rightarrow$ cwdet12	94.5	31.8	42.3 ± 0.2 (3)	—
soybean $\rightarrow$ cwdet12	94.5	18.8	36.9 ± 0.5 (3)	—
cofly $\rightarrow$ cwdet12	94.5	14.3	6.3 ± 0.8 (3)	—

where DAOD helps, results often remain far below in-domain training. We note that ALDI was applied with its default configuration and a more extensive hyperparameter or augmentation search may improve results. Our observation is narrower: *applying a modern DAOD framework out of the box does not yield consistent improvements across the pairs we evaluate*. This variability matters in practice. UAV-based weed detection systems are often deployed by agronomists rather than ML specialists, and a method whose success depends on deployment-specific tuning is harder to recommend as a default. Few-shot fine-tuning on a detector pretrained on a domain-adjacent source is a better-understood procedure with fewer tunable modules and less opaque failure modes. We examine whether it offers a more reliable alternative next.

Q3: Few-shot fine-tuning Table 2 presents results for fine-tuning a detector trained on the source domain using a small number of labeled samples from the target domain. A handful of labeled data outperforms DAOD when the target is learnable. On soybean, $k=5$ already reaches 66.7 ± 4.3 , above both source-only (57.8) and the best DAOD result (50.3), and $k=25$ reaches 80.0 ± 3.2 , near the 76.8 in-domain score; full fine-tuning reaches 84.9 ± 0.5 , indicating that source pretraining adds value. The effect is larger for weak sources: CWDet12 is a weak source for Soybean (6.0 for source-only and 9.6 for best DAOD), yet five labels lift it to a mAP of 40.9. Harder datasets like cotton behave differently. A handful of labels do not effectively transfer, similar to DAOD.

Table 2. Few-shot fine-tuning vs. unsupervised domain adaptation (target-test mAP₅₀). Each K column is the mean $\pm$ std over M random label draws ($M=5$ for $K \in \{5, 10\}$, $M=4$ for $K=25$). The “All (full FT)” column fine-tunes the source detector on the *entire* target training set (full supervision, $n=3$ seeds, longer schedule) and is compared against the in-domain score to isolate the value of source pretraining. *Source* and *DAOD* (from Table 1) and the in-domain score are shown for reference. Fine-tuning on a handful of target labels exceeds both source-only and the best adapted model on the learnable target (soybean (S)), reaching the in-domain score by $K=25$, but does not move the target (cotton (C)). CoFly and CWDet12 are referred by CF and CW for brevity.

Transfer	Source	DAOD	$K=5$	$K=10$	$K=25$	Full FT	In-Domain
C $\rightarrow$ S	57.8	50.3	66.7 ± 4.3	70.8 ± 3.6	80.0 ± 3.2	84.9 ± 0.5 (3)	76.8
S $\rightarrow$ C	7.5	8.3	8.8 ± 1.1	8.2 ± 0.7	9.4 ± 1.5	35.1 ± 0.7 (3)	37.4
CW $\rightarrow$ S	6.0	9.6	40.9 ± 11.4	56.8 ± 8.2	65.4 ± 6.2	84.9 ± 1.6 (5)	76.8
CW $\rightarrow$ C	1.5	3.9	2.4 ± 0.4	3.2 ± 1.1	3.3 ± 0.4	43.2 ± 2.4 (5)	37.4
CF $\rightarrow$ S	4.1	16.7	37.2 ± 15.3	49.5 ± 13.9	62.9 ± 2.6	75.4 ± 1.0 (3)	76.8
CF $\rightarrow$ C	1.3	1.6	3.0 ± 1.1	4.1 ± 1.1	8.1 ± 2.4	41.3 ± 2.5 (3)	37.4

6 Conclusion

We presented an empirical study of cross-crop transfer for UAV-based weed detection, introducing a newly collected cotton dataset and comparing unsupervised DAOD against few-shot fine-tuning on a domain-adjacent source across a range of target label budgets. Applying a modern DAOD framework (ALDI) out of the box yielded inconsistent improvements over source-only training, while few-shot fine-tuning with as few as $k=25$ labeled target examples proved more reliable in our setting. Extending this analysis to additional crops, a deeper analysis of domain to domain discrepancy in transferability, and leveraging these observations to develop new DAOD methods are promising directions for future work.

7 Acknowledgements

The authors gratefully acknowledge the generous support from AFOSR award FA9550-23-1-0239.

References

1. Bah, M.D., Hafiane, A., Canals, R.: Deep learning with unsupervised data labeling for weed detection in line crops in uav images. *Remote sensing* **10**(11), 1690 (2018)
2. Borah, S., Anand, A., Sundaravadeivel, P., Fletcher, R., Reddy, K.: Unmanned aerial vehicle solutions for weed detection in precision agriculture. *IEEE Access* (2025)

3. Cao, S., Joshi, D., Gui, L.Y., Wang, Y.X.: Contrastive mean teacher for domain adaptive object detectors. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 23839–23848 (2023)
4. Chen, M., Chen, W., Yang, S., Song, J., Wang, X., Zhang, L., Yan, Y., Qi, D., Zhuang, Y., Xie, D., et al.: Learning domain adaptive object detection with probabilistic teacher. arXiv preprint arXiv:2206.06293 (2022)
5. Chen, Y., Li, W., Sakaridis, C., Dai, D., Van Gool, L.: Domain adaptive faster r-cnn for object detection in the wild. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3339–3348 (2018)
6. Chen, Y., Wang, H., Li, W., Sakaridis, C., Dai, D., Van Gool, L.: Scale-aware domain adaptive faster r-cnn. *IJCV* **129**(7), 2223–2243 (2021)
7. Dang, F., Chen, D., Lu, Y., Li, Z.: YOLOWeeds: A novel benchmark of YOLO object detectors for multi-class weed detection in cotton production systems. *Computers and Electronics in Agriculture* **205**, 107655 (2023)
8. Darema, F.: Dynamic data driven applications systems (dddas) — a transformative paradigm. In: Proceedings of the 8th International Conference on Computational Science, Part III. p. 5. ICCS '08, Springer-Verlag, Berlin, Heidelberg (2008). https://doi.org/10.1007/978-3-540-69389-5_2
9. Das, A., Subburaj, V.H., Yang, Y., Bednarz, C.W.: A uav image dataset for object detection with annotations generated using labeling and roboflow. *Data in Brief* p. 112483 (2026)
10. Deng, J., Li, W., Chen, Y., Duan, L.: Unbiased mean teacher for cross-domain object detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4091–4101 (2021)
11. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., March, M., Lempitsky, V.: Domain-adversarial training of neural networks. *Journal of machine learning research* **17**(59), 1–35 (2016)
12. He, X., Li, X., Guo, X.: Differential alignment for domain adaptive object detection. In: AAAI. vol. 39, pp. 17150–17158 (2025)
13. Hu, K., Wang, Z., Coleman, G., Bender, A., Yao, T., Zeng, S., Song, D., Schumann, A., Walsh, M.: Deep learning techniques for in-crop weed recognition in large-scale grain production systems: A review. *Precision Agriculture* **25**(1), 1–29 (2024)
14. Ilyas, T., Lee, J., Won, O., Jeong, Y., Kim, H.: Overcoming field variability: unsupervised domain adaptation for enhanced crop-weed recognition in diverse farmlands. *Frontiers in Plant Science* **14**, 1234616 (2023)
15. Kay, J., Haucke, T., Stathatos, S., Deng, S., Young, E., Perona, P., Beery, S., Van Horn, G.: Align and distill: Unifying and improving domain adaptive object detection. *Transactions on Machine Learning Research (TMLR)* (2025)
16. Krestenitis, M., Raptis, E.K., Kapoutsis, A.C., Ioannidis, K., Kosmatopoulos, E.B., Vrochidis, S., Kompatsiaris, I.: CoFly-WeedDB: A UAV image dataset for weed detection and species identification. *Data in Brief* **45**, 108575 (2022)
17. Li, Y.J., Dai, X., Ma, C.Y., Liu, Y.C., Chen, K., Wu, B., He, Z., Kitani, K., Vajda, P.: Cross-domain adaptive teacher for object detection. In: CVPR. pp. 7581–7590 (2022)
18. Lottes, P., Khanna, R., Pfeifer, J., Siegwart, R., Stachniss, C.: Uav-based crop and weed classification for smart farming. In: 2017 IEEE international conference on robotics and automation (ICRA). pp. 3024–3031. IEEE (2017)
19. Oerke, E.C.: Crop losses to pests. *J. Agric. Sci.* **144**(1), 31–43 (2006)
20. Owen, M.D., Zelaya, I.A.: Herbicide-resistant crops and weed resistance to herbicides. *Pest Management Science: formerly Pesticide Science* **61**(3), 301–311 (2005)

21. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 779–788 (2016)
22. Saito, K., Ushiku, Y., Harada, T., Saenko, K.: Strong-weak distribution alignment for adaptive object detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6956–6965 (2019)
23. Shahi, T.B., Dahal, S., Sitaula, C., Neupane, A., Guo, W.: Deep learning-based weed detection using uav images: A comparative study. *Drones* **7**(10), 624 (2023)
24. Torres-Sánchez, J., López-Granados, F., De Castro, A.I., Peña-Barragán, J.M.: Configuration and specifications of an unmanned aerial vehicle (uav) for early site specific weed management. *PloS one* **8**(3), e58210 (2013)
25. Veeranampalayam Sivakumar, A.N., Li, J., Scott, S., Psota, E., J. Jhala, A., Luck, J.D., Shi, Y.: Comparison of object detection and patch-based classification deep learning models on mid-to late-season weed detection in uav imagery. *Remote Sensing* **12**(13), 2136 (2020)
26. Vs, V., Gupta, V., Oza, P., Sindagi, V.A., Patel, V.M.: Mega-cda: Memory guided attention for category-aware unsupervised domain adaptive object detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4516–4526 (2021)
27. Zhang, M., Rong, Q., Jing, H.: Ttsda-yolo: A two training stage domain adaptation framework for object detection in adverse weather. *IEEE Transactions on Instrumentation and Measurement* **74**, 1–13 (2025). <https://doi.org/10.1109/TIM.2024.3497132>