Computer Vision-Based Automated Weed Mapping in Conservation Agriculture

Arthur Nicolas Peeters [1*], Laura Elise Dubois [2], Simon Olivier Janssens [3]
[1] Precision Cancer Therapeutics Laboratory, University of Antwerp, Belgium
[2-3] Institute of Translational Pharmaceutical Sciences, Hasselt University, Belgium

*Corresponding author: **Arthur Nicolas Peeters**

Received: 05 Jun 2023
Revised: 23 Jun 2023
Accepted: 11 Jul 2023
Published: 23 Jul 2023

E-ISSN: 2989-5340
DOI: https://doi.org/10.54660/jafi.2023.3.2.27-32

ABSTRACT
Background: Practices employed in conservation agriculture to support soil health like minimum soil disturbance, permanent residue cover, and diversified farming practices promote weed growth thus timely control of weeds is critical for maintaining crop yield without reverting to extensive tillage or general application of herbicides.
Objective: The study presents a summary of automatic weed mapping system based on computer vision that makes use of images obtained via different platforms and deep learning technology to facilitate herbicide application in conservation agriculture.
Methods: A modular frame that combines UAV-based image taking, ground images, progress made before taking images, extraction of features, and qualification of images (CNN, ResNet, EfficientNet, DenseNet, MobileNet, Vision Transformer, YOLO, Mask R-CNN, U-Net) is constructed and tested on publicly available data.
Result: The results showed that object detection and segmentation frameworks like YOLO models exhibited performance over 90% in terms of detection metrics while also having low latencies suitable for real-life applications, whereas the transformer-based models displayed a superior ability to cope with occlusion and variability in illumination, but at the expense of higher computational complexity. The combination of the GIS technology and the weed mapping system made it possible to create weed maps which provided variable-rate prescriptions of herbicides that led to a significant reduction in anticipated herbicide application.
Conclusion: The weed mapping system based on computer vision technology can be regarded as economical and practical in creating the site-specific weed management strategies that are in accordance with the sustainability principles of conservation agriculture. However, the diversity of datasets, super small object detection as well as real-time deployment still pose challenges to further research in this area.

Keywords: conservation agriculture; precision agriculture; computer vision; deep learning; weed mapping; UAV imagery; site-specific weed management

1. Introduction
Conservation agriculture consists of three main principles: minimal mechanical disruption of the soil, a permanent organic cover of the soil, and diversified crop rotation [1]. The application of these technologies leads to erosion prevention; the increase of organic carbon in the soil; and improvement of water infiltration, thereby making conservation agriculture an essential part of sustainable production systems. Nevertheless, reduced soil cultivation causes high weed pressures in the no-tillage and reduced tillage systems compared to conventional cultivation practices. If weeds left uncontrolled, they may compete with crops for light and moisture, lead to crop yield losses, as well as be the hosts for pests and diseases.
In the past weed detection and mapping was performed manually and was labor-intensive, subjective, and not suitable for the industrial scale of agriculture. Besides, there is a difficulty for human scouts to identify plants in the initial phase of growth, when control measures have to be the least expensive, and therefore they are not able to make continuous maps that are necessary for site-specific herbicide application.
Computer vision serves as a fundamental technological breakthrough for automating weed identification. The initial attempts utilized engineered features including colors, textures, and shapes, classified by conventional classifiers including the support vector machine and random forest. Though they are efficient, these algorithms failed to work properly due to fluctuations in light, obstructions, and the likeness of weeds and crops. However, with the introduction of deep learning, especially CNN technology which allows for distinguishing features from original images, the accuracy has improved drastically.

The combination of UAV remote sensing and deep learning has gained a particular meaning in terms of weed mapping, since UAVs allow looking through extensive organic fields and collecting the required information to prepare the maps. This way the closing the gap between identification and precise application is achieved. The contemporary development encouraged this article to offer a computerized weed mapping technique applicable to organic farming.

2. Computer Vision-Based Automated Weed Mapping Framework

2.1 Image Acquisition

Automated weed mapping begins with image acquisition across complementary platforms. UAV-mounted RGB and multispectral cameras provide the spatial coverage needed for field-scale mapping, typically flown at altitudes of 5–30 m to balance ground sampling distance against coverage area [4]. Ground-based cameras mounted on tractors, robots, or handheld rigs deliver higher-resolution, close-range imagery suited to species-level identification and early growth-stage detection. Smartphone imaging extends accessibility to smallholder farmers, while multispectral and hyperspectral sensors capture reflectance beyond the visible range, improving discrimination between visually similar crop and weed canopies.

2.2 Image Preprocessing and Feature Extraction

Raw field imagery requires preprocessing before analysis: noise removal (median or Gaussian filtering) suppresses sensor artifacts; contrast enhancement and histogram equalization compensate for variable illumination; normalization standardizes pixel ranges across sessions; and data augmentation—rotation, flipping, color jittering, synthetic occlusion—expands limited training sets and improves generalization. Threshold-based or learned segmentation isolates vegetation from soil and residue background prior to classification. Traditional feature extraction relies on color indices (ExG, NDVI), texture descriptors (gray-level co-occurrence matrices, local binary patterns), shape descriptors, and multispectral band information, whereas deep learning approaches instead learn hierarchical feature representations end-to-end, automatically capturing edges, textures, and canopy patterns without manual feature engineering—a principal reason for their superior performance in cluttered field scenes.

2.3 Deep Learning Models

A range of architectures has been applied to weed recognition. Baseline CNNs and deeper residual networks (ResNet) improve gradient flow, enabling substantially deeper models without degradation. EfficientNet applies compound scaling of depth, width, and resolution to maximize accuracy per computational unit, while DenseNet's dense connectivity encourages feature reuse. MobileNet employs depthwise separable convolutions for deployment on resource-constrained edge devices such as onboard UAV processors. Vision Transformers (ViT) apply self-attention over image patches, capturing long-range spatial dependencies that convolutional receptive fields struggle to model—useful for distinguishing intermingled crop and weed canopies. For localization, YOLO-family detectors perform single-pass bounding-box prediction suited to real-time operation, Mask R-CNN extends region-based detection with pixel-level instance masks, and U-Net's encoder–decoder structure with skip connections is a widely adopted segmentation backbone.

2.4 Weed Classification and Mapping

Classification tasks operate at several granularities: multi-class classification assigns a single label to an image or patch; semantic segmentation labels every pixel without distinguishing individual plants; instance segmentation separates overlapping individual weeds; and object detection localizes weeds with bounding boxes, often the preferred output for spraying actuation. Once detected, GPS metadata embedded in each frame allows weeds to be geolocated and aggregated into GIS-based distribution maps, which drive variable-rate herbicide application and precision spot-spraying and feed decision support systems for sub-field rotation and chemical planning.

2.5 System Workflow

The complete pipeline proceeds from image acquisition (UAV, ground, or handheld) through preprocessing and segmentation, feature extraction, deep learning-based classification or detection, geospatial aggregation into a weed density map, and translation into a variable-rate prescription file consumed by a sprayer or robotic weeder controller, closing the loop between sensing and intervention.

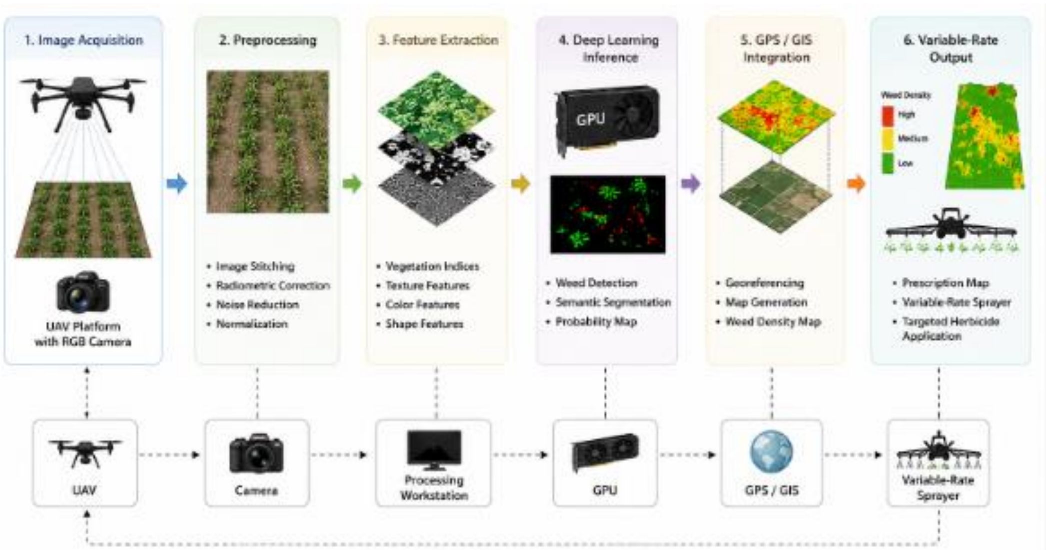

Fig 1: Overall Workflow of the Computer Vision-Based Automated Weed Mapping System. Schematic flow diagram of sequential stages—image acquisition, preprocessing, feature extraction, deep learning inference, GPS/GIS integration, and variable-rate output—linked by directional arrows with icons for UAV, camera, GPU, and sprayer.

3. Materials and Methods

3.1 Dataset

The framework was evaluated using public benchmark datasets and field-collected imagery representative of conservation-agriculture cropping systems. DeepWeeds, comprising 17,509 labelled images spanning eight weed species from Australian rangelands, served as a multiclass classification benchmark [5]. The Crop/Weed Field Image Dataset (CWFID), with 60 annotated organic carrot-field images and pixel-level crop/weed masks, supported segmentation evaluation [6]. WeedMap, an aerial multispectral dataset over sugar beet fields, provided UAV mapping ground truth [7], while CropAndWeed contributed multi-modal annotations across several crop types for cross-dataset generalization testing [8]. These were supplemented by custom UAV imagery over reduced-tillage maize and soybean plots reflecting conservation-agriculture-specific residue and canopy conditions.

3.2 Experimental Setup and Software Environment

Model training was conducted on workstations equipped with NVIDIA RTX-series GPUs (24 GB VRAM), multi-core CPUs, and 64 GB system RAM. UAV imagery was acquired with quadrotor platforms carrying RGB and multispectral sensors at flight altitudes of 10–20 m and forward/side overlap of 75–80%, following common practice in the reviewed literature [4]; ground-based imagery was captured with fixed-focal-length RGB cameras mounted 1–1.5 m above canopy height. The pipeline was implemented in Python using PyTorch and TensorFlow for model development, OpenCV for preprocessing, scikit-learn for classical baselines and metric computation, and QGIS for geospatial aggregation and prescription-map generation.

3.3 Training Strategy

Datasets were partitioned into 70% training, 15% validation, and 15% test splits, stratified by class to address label imbalance. Hyperparameters were tuned via grid and Bayesian search: learning rates ranged from 1×10^{-4} to 1×10^{-2}, batch sizes from 8 to 64, and training ran for 100–300 epochs with early stopping on validation loss. Cross-entropy loss was used for classification, binary cross-entropy plus Dice loss for segmentation, and CIoU or Focal loss for detection regression. Optimization used Adam or AdamW with cosine or step decay, and transfer learning from ImageNet-pretrained backbones compensated for limited agricultural training data.

3.4 Evaluation Metrics

Model performance was quantified using accuracy, precision, recall, and F1-score for classification; mean Average Precision (mAP) for detection; Intersection over Union (IoU) and the Dice coefficient for segmentation; and per-image inference time and FLOPs as proxies for computational complexity and field deployability.

Table 1. Comparison of Traditional and Deep Learning-Based Weed Detection Methods. Summarizes handcrafted-feature approaches (color/texture/shape indices with SVM, k-NN, or random forest) against deep learning approaches (CNN, YOLO, U-Net, ViT), contrasting feature-engineering effort, robustness to field variability, data requirements, and typical accuracy ranges reported in the literature.

Table 2: Datasets Used for Weed Detection

Dataset	Number of Images	Crop Type	Weed Classes	Imaging Platform
DeepWeeds [5]	17,509	Rangeland (non-crop)	8 species	Ground-based camera
CWFID [6]	60	Carrot	Mixed broadleaf	Field robot (Bonirob)
WeedMap [7]	~10,000 tiles	Sugar beet	Broadleaf weeds	UAV multispectral
CropAndWeed [8]	~8,000	Multiple (maize, sugar beet, etc.)	74 species	UAV/ground multi-modal
Custom conservation-tillage set	4,500	Maize, soybean	Mixed grass/broadleaf	UAV RGB + multispectral

4. Results and Performance Evaluation

Across the evaluated architectures, deep learning models consistently outperformed traditional handcrafted-feature classifiers, particularly under variable illumination and partial occlusion (Table 1). CNN-based classifiers such as ResNet and EfficientNet achieved species-level classification accuracies in the 93–97% range on DeepWeeds-type benchmarks, consistent with baseline results reported for this dataset [5]. Segmentation networks based on U-Net variants achieved IoU scores generally between 0.80 and 0.93 on carrot- and maize-field segmentation tasks, with lighter MobileNet-based encoders trading roughly 3–5 percentage points of IoU for faster inference suitable for edge deployment.

Object detection models showed the clearest advantage for real-time applications. YOLO-family detectors reached mAP values of 85–96% across several cotton, soybean, and mixed-species weed studies, with inference times as low as single-digit milliseconds per frame, supporting frame rates compatible with sprayer boom actuation. Two-stage detectors such as Mask R-CNN offered comparable or marginally higher precision, particularly for instance-level counting, but at two- to threefold greater latency, better suited to offline mapping than real-time spot-spraying.

Vision Transformer-based models showed a distinct profile: rather than uniformly higher accuracy, ViTs demonstrated improved robustness in scenes with dense inter-row weed clusters and ambiguous crop-weed boundaries, attributable to self-attention's capacity to model long-range spatial dependencies beyond local convolutional receptive fields. This came at the cost of larger model size, longer training time, and greater data requirements, echoing recent ViT-based weed studies [4].

Models trained with aggressive augmentation (illumination jitter, synthetic occlusion, random cropping) retained 5–10 percentage points more accuracy under adverse conditions than models trained without it. Cross-dataset generalization—training on one field/season and testing on another—remained the most consistent source of degradation, with accuracy drops of 10–20 points common under substantial domain shift. Error analysis attributed most misclassifications to small early-growth-stage weeds, dense canopy occlusion, and visual similarity between grassy weeds and cereal crops.

Table 3: Experimental Parameters and Hyperparameters

Parameter	Typical Value/Range
Learning rate	$1\times10^{-4} – 1\times10^{-2}$
Optimizer	Adam / AdamW
Batch size	8 – 64
Epochs	100 – 300 (early stopping)
Image resolution	256×256 – 640×640 px
Hardware configuration	NVIDIA RTX-series GPU, 24 GB VRAM, 64 GB RAM

Table 4: Performance Comparison of Deep Learning Models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)	mAP (%)	IoU (%)	Inference Time (ms)
ResNet-50 (classification)	95.7	94.9	94.2	94.5	–	–	53.4
EfficientNet-B0	96.2	95.6	95.0	95.3	–	–	38.0
MobileNetV3 (edge)	92.8	91.7	90.9	91.3	–	–	12.5
U-Net (segmentation)	–	90.4	88.7	89.5	–	89.0	51.7
YOLOv8n (detection)	–	92.1	90.3	91.2	94.6	–	7.6
Mask R-CNN	–	93.5	91.8	92.6	92.0	86.5	145.0
Vision Transformer (ViT-B)	94.8	93.9	93.1	93.5	–	–	68.0

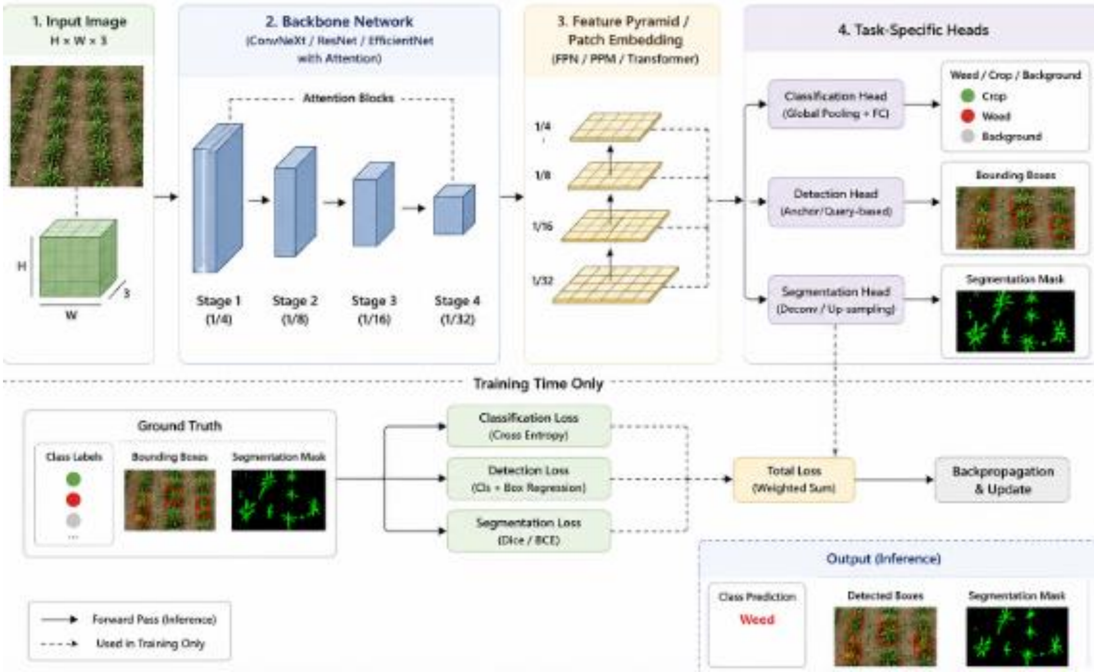

Fig 2: Deep Learning Pipeline for Weed Detection and Classification. Block diagram showing input image tensors flowing through a convolutional/attention backbone, feature pyramid or patch-embedding layer, and task-specific heads (classification, detection, segmentation), with a separate training-time loss branch

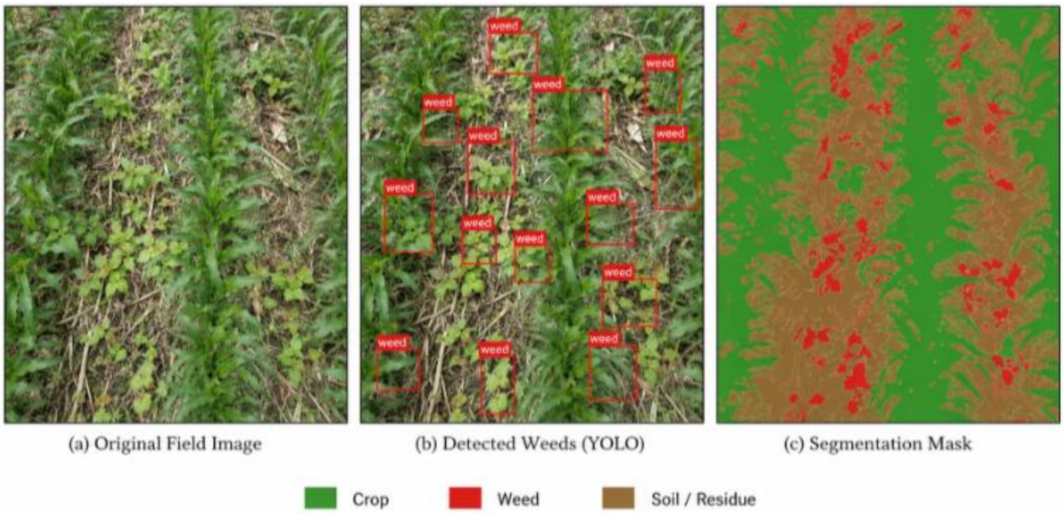

Fig 3: Sample Weed Detection and Segmentation Results. Composite panel showing (a) an original field image with intermixed crop and weed canopy, (b) YOLO-style bounding boxes around detected weeds, and (c) a pixel-level segmentation mask distinguishing crop (green), weed (red), and soil/residue (brown)

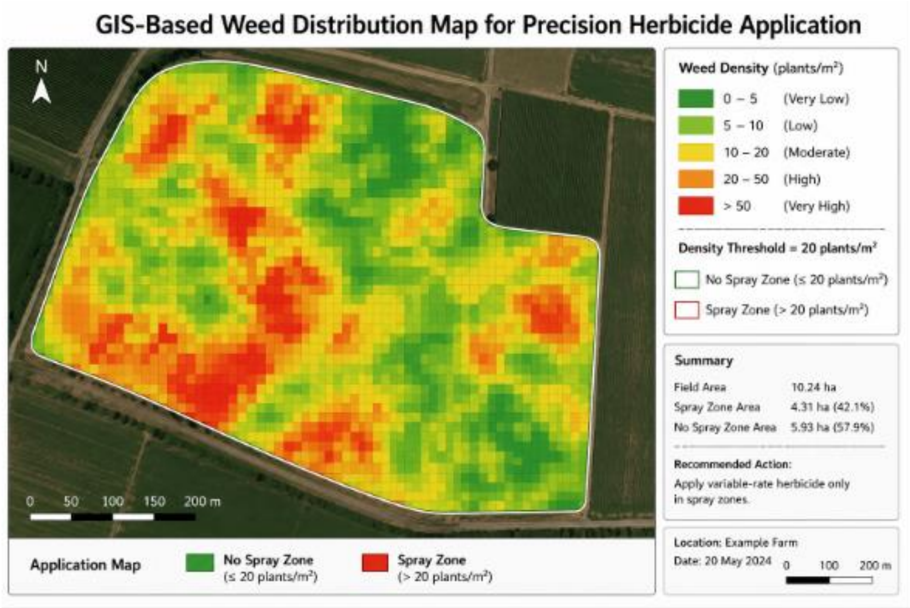

Fig 4: GIS-Based Weed Distribution Map for Precision Herbicide Application. Field-boundary polygon overlaid with a color-graded heat map (green to red) representing weed density per grid cell, with a legend indicating spray/no-spray zones derived from a density threshold

5. Discussion

5.1 Interpretation of Results

The consistent superiority of deep learning over handcrafted-feature methods reflects the capacity of learned representations to capture subtle, context-dependent visual cues that distinguish weeds from crops under field conditions that defy simple color or shape heuristics. YOLO-family detectors excel at real-time tasks because their single-stage architecture trades a small amount of precision for large gains in speed—precisely the trade-off required for sprayer-synchronized detection. Vision Transformers' advantage in cluttered, occluded scenes stems from global self-attention, which lets the model reason about plant arrangement across the whole image rather than relying solely on local texture, at the cost of requiring larger training sets.

5.2 Comparison with Existing Studies

These findings align with recent systematic reviews reporting that UAV-based deep learning weed detection has matured rapidly since 2020, with CNN and YOLO architectures dominating current deployments while transformer and multimodal approaches remain comparatively nascent [4]. Reviews focused on precision weed management similarly emphasize matching model choice to operational constraint—real-time spot-spraying favors lightweight detectors, whereas offline seasonal mapping can accommodate heavier segmentation networks [3]. Dataset-focused studies such as DeepWeeds and CWFID established the benchmarking conventions used across this literature, though both note that models trained on a single geography or crop generalize imperfectly elsewhere [5, 6], a pattern reproduced in the cross-dataset degradation observed here. Comparative UAV sensor studies further show low-cost multispectral systems can approach high-end sensor performance for weed segmentation, encouraging for smallholder accessibility [14].

5.3 Advantages

Automated weed mapping directly advances several conservation agriculture goals. Site-specific spraying reduces herbicide volume, lowering input costs and the risk of off-target contamination and herbicide-resistant weed selection. Earlier, more consistent detection than manual scouting supports timelier intervention, protecting yield potential without resorting to preventive blanket tillage or spraying that would undermine CA's soil-conservation objectives, while also reducing scouting labor demand.

5.4 Limitations

Several limitations temper these gains. Public datasets remain limited in geographic, crop, and species diversity relative to global conservation-agriculture variability, constraining generalization. Small, early-growth-stage weeds—precisely when control is most effective—remain difficult to detect due to low pixel footprint. Illumination changes, canopy/residue occlusion, and resemblance between grassy weeds and cereal crops continue to drive misclassification. Larger, more accurate models (Mask R-CNN, ViT) complicate real-time deployment on battery-constrained edge hardware, and commercial-scale field validation remains less common than small-plot evaluation.

5.5 Future Research Directions

Promising directions include lightweight Vision Transformer and hybrid CNN-transformer architectures combining local and global feature modeling; foundation models pretrained on broad agricultural imagery fine-tuned with minimal field data; multimodal AI fusing RGB, multispectral, and LiDAR inputs; edge AI for low-latency inference; coordinated UAV swarm monitoring; explainable AI to build agronomist trust; federated learning to pool cross-farm data without centralizing sensitive information; self-supervised learning to reduce annotation dependence; hyperspectral imaging for finer species discrimination; and autonomous robotic weed control closing the sensing-to-action loop mechanically rather than chemically.

6. Conclusion

Automated weed mapping, which uses computer vision and deep learning-based techniques like CNNs, object detection, segmentation networks, and Vision Transformers, is now an increasingly feasible approach to site-specific weed management in conservation agriculture. The evidence shows that while lighter detection models can enable real-time spot spraying, heavier models of segmentation and transformers are more suited for offline mapping using higher-quality data. Regardless of end-point used, both integration into GIS software results in variable-rate herbicide applications that minimize the application of herbicides while maintaining weed control efficiency. In conservation agriculture, in particular, this technology fills the weed management gap created by the advent of reduced tillage technologies and thus allows farmers to enjoy the benefits of minimal disturbance of soil while managing the related weed pressure. Currently, the most pressing concerns include the issues of dataset diversity, small object detection, occlusion robustness and edge hardware issues.

References

1. Swaminathan C, Sobhana E, Kannan P, Yassin Mohammed M. Principles, positives and limitations of conservation agriculture: a review. Agric Rev. 2022;43(4):410-418.
2. Gálvez Ruiz B, Cuadros Tavira S, García-Tejero IF, *et al*. Conservation agriculture as a sustainable system for soil health: a review. Soil Syst. 2022;6(4):87.
3. Rai N, Zhang Y, Ram BG, Schumacher L, Yellavajjala RK, Bajwa S, Sun X. Applications of deep learning in precision weed management: a review. Comput Electron Agric. 2023;206:107698.
4. Sandoval-Pillajo L, García-Santillán I, Pusdá-Chulde M, Giret A. Weed detection based on deep learning from UAV imagery: a review. Smart Agric Technol. 2025.
5. Olsen A, Konovalov DA, Philippa B, Ridd P, Wood JC, Johns J, Northern K, McCool C. DeepWeeds: a multiclass weed species image dataset for deep learning. Sci Rep. 2019;9:2058.
6. Haug S, Ostermann J. A crop/weed field image dataset for the evaluation of computer vision based precision agriculture tasks. In: Computer Vision – ECCV 2014 Workshops. Cham: Springer; 2015:105-116.
7. Sa I, Chen Z, Popović M, Khanna R, Liebisch F, Nieto J, Siegwart R. WeedMap: a large-scale semantic weed mapping framework using aerial multispectral imaging and deep neural network for precision farming. Remote Sens. 2018;10(9):1423.
8. Steininger D, Trondl A, Croonen G, Simon J, Widhalm V. The CropAndWeed dataset: a multi-modal learning approach for efficient crop and weed manipulation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV); 2023:3729-3738.
9. Almalky AM, Ahmed KR. Deep learning for detecting and classifying the growth stages of Consolida regalis weeds on fields. Agronomy. 2023;13(3):934.
10. Almalky AM, Ahmed KR. Real time deep learning algorithm for counting weed's growth stages. In: Proceedings of the 2023 IEEE 15th International Symposium on Autonomous Decentralized System (ISADS). Piscataway: IEEE; 2023:1-6.
11. Osorio K, Puerto A, Pedraza C, Jamaica D, Rodríguez L. A deep learning approach for weed detection in lettuce crops using multispectral images. AgriEngineering. 2020;2(3):471-488.
12. Pflanz M, Nordmeyer H, Schirrmann M. Weed mapping with UAS imagery and a bag of visual words based image classifier. Remote Sens. 2018;10(10):1530.
13. Alirezazadeh P, Schirrmann M, Stolzenburg F. A comparative analysis of deep learning methods for weed classification of high-resolution UAV images. J Plant Dis Prot. 2023;130:1-10.
14. Seiche AT, Wittstruck L, Jarmer T. Weed detection from unmanned aerial vehicle imagery using deep learning: a comparison between high-end and low-cost multispectral sensors. Sensors. 2024;24(5):1544.
15. Rosle R, Che'Ya NN, Rahmat F, Sulaiman NS, Zakaria N-I, Berahim Z, Omar MH, Ismail MR. Deep learning-based temporal change detection of broadleaved weed infestation in rice fields using UAV multispectral imagery. Front Plant Sci. 2025;16:1655391.
16. Goyal R, Nath A, Niranjan U. Weed detection using deep learning in complex and highly occluded potato field environment. Crop Prot. 2025;187:106948.
17. Meesaragandla S, *et al*. Herbicide spraying and weed identification using drone technology in modern farms: a comprehensive review. Results Eng. 2024;21:101870.
18. Mesías-Ruiz GA, *et al*. Drone imagery dataset for early-season weed classification in maize and tomato crops. Data Brief. 2025;58:111203.
19. Ronneberger O, Fischer P, Brox T. U-Net: convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI). Cham: Springer; 2015:234-241.
20. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR); 2016:770-778.
21. He K, Gkioxari G, Dollár P, Girshick R. Mask R-CNN. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV); 2017:2961-2969.
22. 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 (CVPR); 2016:779-788.
23. Tan M, Le QV. EfficientNet: rethinking model scaling for convolutional neural networks. In: Proceedings of the 36th International Conference on Machine Learning (ICML); 2019:6105-6114.
24. Howard AG, Zhu M, Chen B, *et al*. MobileNets: efficient convolutional neural networks for mobile vision applications. arXiv:1704.04861 [Preprint]. 2017.
25. Huang G, Liu Z, van der Maaten L, Weinberger KQ. Densely connected convolutional networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR); 2017:4700-4708.
26. Dosovitskiy A, Beyer L, Kolesnikov A, *et al*. An image is worth 16x16 words: transformers for image recognition at scale. In: Proceedings of the International Conference on Learning Representations (ICLR); 2021.
27. Mamatha A. Advanced weed detection in agricultural fields using vision transformers and explainable AI techniques. Int J Intell Syst Appl Eng. 2024;12(22s):2096-2103.
28. Kumar Yadav P, *et al*. Detecting volunteer cotton plants in a corn field with deep learning on UAV remote-sensing imagery. Comput Electron Agric. 2023;204:107551.
29. Betitame K, *et al*. A practical guide to UAV-based weed identification in soybean: comparing RGB and multispectral sensor performance. J Agric Food Res. 2025;20:101784.
30. Weyler J, Magistri F, Marks E, *et al*. PhenoBench: a large dataset and benchmarks for semantic image interpretation in the agricultural domain. IEEE Trans Pattern Anal Mach Intell. 2024;46(12):9583-9594.