PROJECT ABSTRACT

Weed Detection from Field Images using YOLO

A built-to-order final-year project that trains a YOLOv8s detector on the published DeepWeeds dataset to localise and classify weeds in crop-field photographs for precision spraying, with training notebook, metrics, interactive demo and complete documentation.

01. Title

Weed Detection from Field Images using YOLO

02. Introduction / Background

Blanket herbicide spraying treats the whole field even though weeds occupy a fraction of it, wasting chemical, harming soil biology and costing farmers money. Precision spraying applies herbicide only where weeds are, but it needs a vision system that can find weeds among crops in real field photographs. YOLO (You Only Look Once) detectors predict bounding boxes and classes in a single pass, making them fast enough for real-time field use. DeepWeeds (Olsen et al., 2019) is a published benchmark of 17,509 rangeland images across 8 weed species, giving this project a solid, citable data foundation.

03. Problem Statement

Weed-vs-crop discrimination in uncontrolled field lighting, with occlusion and growth-stage variation, is beyond simple colour thresholding. Student projects often train on tiny hand-collected sets with no standard evaluation; this project instead uses a published benchmark and the standard detector metrics (mAP) so results are meaningful and viva-defensible.

04. Objectives

- Prepare DeepWeeds annotations in YOLO format with bounding boxes per weed instance.
- Fine-tune YOLOv8s (COCO pre-trained) for 8 weed species plus background.
- Evaluate with mAP@0.5, mAP@0.5:0.95 and per-species precision-recall.
- Export weights to ONNX with notes on edge deployment.
- Deliver an interactive demo drawing predicted boxes on sample field photos.

05. Proposed Methodology

Images are resized to 640×640 with mosaic and HSV augmentation (Ultralytics defaults) to simulate field variation. YOLOv8s — CSPDarknet backbone, PAN-FPN neck, anchor-free decoupled head — is fine-tuned with SGD for up to 120 epochs with early stopping. Non-maximum suppression filters overlapping predictions at inference. Evaluation follows the COCO-style mAP protocol on a held-out split.

06. System Architecture / Working

- Input: crop-field photograph.
- Preprocessing: resize 640×640, normalise.
- Model: YOLOv8s backbone + neck + anchor-free detection head.

- Post-processing: NMS, confidence thresholding.
- Output: bounding boxes with weed species label and confidence; overlay doubles as a spray map.
- Demo: web app drawing representative detections on sample photos.

07. Hardware / Software Requirements

- Hardware: free cloud GPU (T4) for training (~4–6 h expected); GPU or edge device for real-time inference.
- Software: Python 3, PyTorch, Ultralytics YOLOv8, OpenCV, NumPy, Matplotlib.
- Demo: any modern web browser.

08. Expected Outcomes

- Trained YOLOv8s weed detector with design-target mAP@0.5 ~0.88 (reported honestly after training).
- Per-species PR curves and sample detection sheets.
- ONNX export, interactive demo, report, PPT and viva Q&A.;

09. Applications

- Perception module for precision-sprayer robots and tractor implements.
- Weed-mapping surveys from drone or phone photography.
- Teaching reference for object-detection coursework.

10. Future Scope

- Sprayer actuation interface (solenoid control from detections).
- Tracking weeds across video frames for a moving platform.
- Extending to Indian crop-weed species with additional data.