


AI for Smart Farming: Detecting and Identifying Cattle from Drone Images



By Eyad Ahmed Abdelaziz Abdelnabi Mohamed
eyadmohamedswe@gmail.com

Supervised by Mingqi Gao
m.gao@sheffield.ac.uk

Overview of the project with a background



- 4K drone images from farm in Nottingham
- How accurately can we re-identify cows?
- How scalable can it be?
- How efficient can the ML pipeline become in terms of storage, memory, and computation?



Data and Methods used

- Overall framework
 - Cropping: YOLO26n
 - Segmentation: SAM3
 - Alignment: Image processing on the segmentation mask
 - Cropping Label: Custom fine-tuned YOLO26n model
 - ID Recognition: Gemini (over cloud)
- Interface between stages
 - Used to be raw images
 - Now it's JSON Objects


Preprocessing



Cropping Cattle




Segmentation



Alignment



Cropping Label



ID Recognition


Detected Number: 3
Confidence: 95%

Data and Methods used



```
{  
  orig_img_name: 'DJI_20250731121343_0030_W.JPG',  
  bounding_box: tensor([0.51, 0.38, 0.59, 0.53]),  
  tensor: tensor(..., dtype=torch.uint8),  
  mask: tensor(..., dtype=torch.bool),  
  ccw_rotation_deg: -150.56,  
  id_bounding_box: tensor([279, 301, 335, 397]),  
  detected_label: 3,  
  detection_conf: 0.95 # i.e. Confidence  
}
```


Preprocessing



Cropping Cattle




Segmentation



Alignment



Cropping Label



ID Recognition

Detected Number: 3
Confidence: 95%

Data and Methods used

- Evolution in number of cows detected
 - The first prototype detected about 40 cows
 - The latest release detected over 110 cows
- Development Stack
 - Python
 - Jupyter Notebook
 - Ultralytics (YOLO and SAM models)
 - Pytorch (tensors)
 - OpenCV
 - JSONLines



Results


- Overall performance
- Ablations (accuracy & storage efficiency)
 - Accuracy of full pipeline: 82%
 - Accuracy of ablation from stage 1 straight to stage 5 (i.e. without preprocessing): 55%

Example cows that were misidentified only during the ablation (not in full pipeline):



Discussion

- In summary:
 - The Full ML Pipeline delivered the intended results
 - Room for future improvement
 - Adding Graphical User Interface
 - Enabling seamless swapping for different implementations/models for every stage in the pipeline
 - Collecting more images for better training and an even higher degree accuracy
- The following are examples of difficult scenarios that the ML pipeline could not identify



Lastly, Many Thanks to my supervisor, Mingqi Gao!