Phenotyping System for Efficient User-centered Detection On AMIGA

Team Name: PSEUDO AMIGA
Student Team Members: Nathaniel Christie, Autonomy & Robotics
Faculty Advisor: Patrick Grenda
Project Advisor: John Reid
Reid Lab Mentors: Amey Gavale, Samuel Folorunsho, Sunny Deshpande

1 Project Description

The objective of PSEUDO-AMIGA was to begin preparations for a phenotyping system for the AMIGA robot. The goals of this preparation were thus:

1. Integrate a MicaSense RedEdge-P multispectral camera onto the AMIGA robotic platform
2. Develop an autonomous data collection pipeline.
3. Create a dataset of field images from the MicaSense camera
4. Generate vegetation indices, including Normalized Difference Vegetation Index (NDVI)[3], from collected field imagery as well as RGB and Green Filtered-NIR NDVI images[8].
5. Train and evaluate machine learning models capable of identifying vegetation and extracting plant phenotypic traits from multispectral data. [8][9][10][11][12][15]

   Other objectives were originally put in place as well, but were not finished due to delays in hardware procurement, last minute hardware issues, and restrictions on field access. These objectives include:

6. Develop a visualization and reporting system for communicating phenotypic measurements to researchers and farmers.
7. Develop an autonomous data collection pipeline that associates multispectral imagery with robot position and timestamp information.
8. Quantify plant health metrics such as canopy coverage, vegetation vigor, and stress indicators. [10] (Stretch Goal)

1.1 Integration

The MicaSense Rededge-P is not directly compatible with the Planet IGS-604HPT-M12 POE connection the AMIGA robot uses, MicaSense does have a POE conversion kit but it is only sold with the camera and does not come with the DJI SkyPort kit which was purchased by the lab. Because of this I had to make this conversion myself. The following is a wiring diagram denoting how the PWR and COMM can be converted to standard power and ethernet connections.

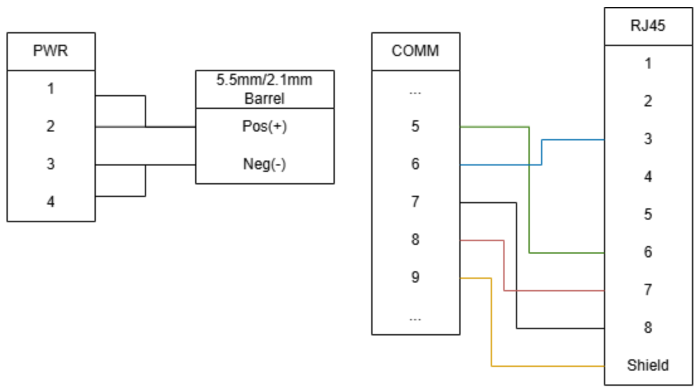

The RJ45 and 5.5mm/2.1mm connections were connected into a POE Texas GAT-12V25W Splitter which was wired into a A Coded M12 8 Position RJ45 Ethernet Cable which could then be connected to the robot directly.

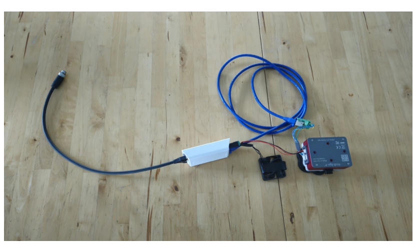

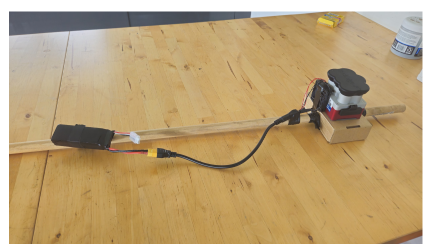

A test was performed to determine how bad the overheating issue would be by letting the camera stream continuously in direct sunlight with a thirty minute timer. After fifteen minutes the camera was extremely hot to the touch and was disconnected and brought inside. A 4486HQ/MF50101V1-1Q01U-S99 Sunon Fan was purchased and connected in parallel with the camera power. The fan was positioned in front of the exhaust and taped to the makeshift mount used for data collection.

A better mount was designed in FreeCAD and prepared to be 3D printed using the Highbay printer, but access was not granted in time for the project deadline so it was not made. In addition, there were issues connecting with the POE on the robot. Even already installed cameras could not be pinged and the issue could not be identified in time as the previously mentioned delays made integration into the robot a lower priority compared to the data collection and segmentation. The STL files for the camera attachment and robot mount are available in the PSEUDO-AMIGA Box folder and instructions on installing them are available on the PSEUDO-AMIGA Github repo.

1.2 Data Collection

The data was collected over several sessions with a test day on July 28th at 2 PM and two sessions on July 29th, around 1 PM and 7 PM for multiple lighting conditions. Over these sessions 2,983 sets of images were collected consisting of red, blue, green, NIR, rededge, and panchromatic TIFF images. They were collected using the makeshift setup displayed above, using capture.py to remotely connect to the MicaSense WiFi and accessing the built in web API to capture an image every second over a specific period of time. The reflective panel included with the camera was imaged at the start of each of these sessions, the camera had to be lowered to about a meter above the panel. For use on the robot I would recommend a small table or stool be used to raise the panel while keeping the same lighting in the field.

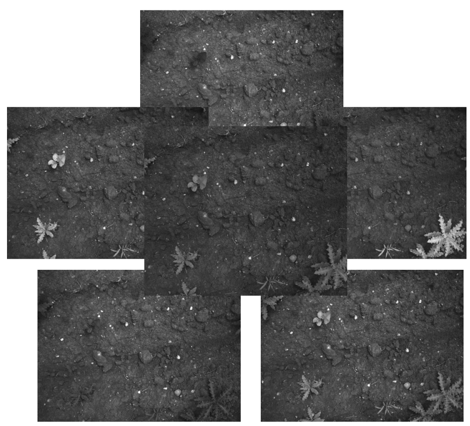

1.3 Generated Images

Two sets of images were generated from the collected sets, being RGB and Green Filtered-NIR NDVI images. The image composites were chosed based on “Segmentation of Weeds and Crops Using Multispectral Imaging and CRF-Enhanced U-Net”[8] which used the GF-NIRNDVI images to much success. I did not use the CRF Enhancement recommended by the paper as while it did increase the detection, at higher resolutions it was only marginal (around 1%). RGB was also used instead of RGB-NIR to have a more direct comparison between standard and multispectral images. Using tiff_processing.py the TIFF images were overlaid and converted to JPEG files which could be segmented by SAM V2. The wavelengths were not aligned with one another due to the location of the lenses on the camera, so they were first each fitted to the green band as reference before laying them on each other, using the initial calculations done on a reference image and applying that same fitting to every image in the batch.

After generating each image, segmentation_tool.py was used to create the masks relating to the images. Since the bands were all aligned to the same green band, both the RGB and GF-NIRNDVI composite images fit to the same masks so only one set was made for both. The masks labeled the background (soil) as [0,0,0], crops as [1,1,1], and weeds as [2,2,2] and saved it as a PNG file. These values were chosen as U-Net models prefer these kinds of channel labels and it keeps other values open for future labeling if a more complicated model was made based on this.

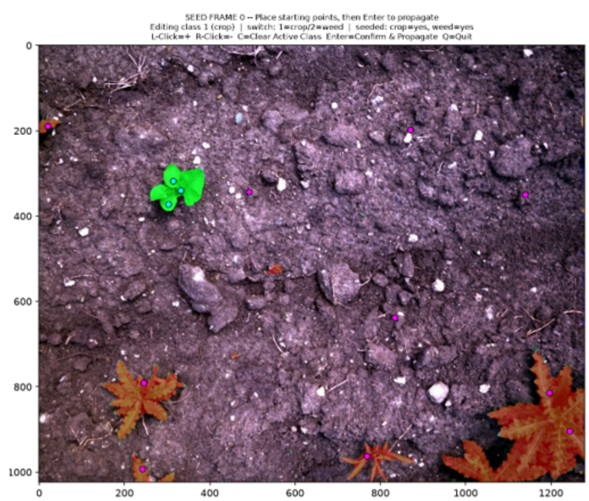

1.4 The Models

Two models were trained, one for RGB images and one for GF-NIRNDVI composite images. The same masks were found to fit both the RGB and GF-NIRNDVI images and were used for both. 2767 pairs were chosen, 2254 for training and 513 for evaluation (around 18.5%). The remaining images were set aside for prediction testing. Images were put through augmentation, like random horizontal and vertical flips as well as rotations for varied data. The model consists of the following files:

File	Purpose
`model.py`	U-Net (4 down/up stages, skip connections, configurable width/upsampling mode)
`dataset.py`	File pairing, sequence-aware train/val split, RGB-mask-to-class-index conversion
`augment.py`	Joint image+mask transforms (flip/rotate/crop, optional photometric jitter)
`losses.py`	CE+Dice loss, inverse-frequency class weight computation
`metrics.py`	IoU/Dice/pixel-accuracy via an accumulated confusion matrix
`utils.py`	Seeding, checkpoint save/load, early stopping, CSV logging, curve plotting
`train.py`	Main training loop
`predict.py`	Inference on new images, RGB mask + overlay output

The U-Net consists of four encoders that MaxPool and then DoubleConv (convolute with a 3x3 kernel, use batch normalization, activate with ReLU, all twice) into four decoders that Upsample into a DoubleConv that gets padded with the opposing encoded "skip" connections to give context to the masking. Essentially it takes the image and based on the trained weights will abstract the image and then recreate it with just three channels that correspond to the classes set previously.

The loss chosen for this model was a combination of Cross Entropy and Dice, with automatic inverse-frequency class weights. Weed pixels are often a small minority of any frame so Dice pushes on region overlap directly rather than per-pixel likelihood alone, and class weighting keeps rare-class gradients from being drowned out by background. The following table displays the results from training the two models compared with the minimum metrics set out by the original project proposal and metrics collected from the study this model was based on. [8]

	Minimum	RGB	GF-NIRNDVI (Paper)	GF-NIRNDVI (Ours)
Soil IoU	90.0%	99.66%	98.3%	99.71%
Crop IoU	80.0%	94.67%	90.5%	95.03%
Weed IoU	65.0%	73.13%	74.4%	77.08%
Mean IoU	78.3%	89.03%	87.8%	90.32%

Overall our GF-NIRNDVI won on every metric, especially on weed IoU. Both the RGB and GF-NIRNDVI composite based models made through this dataset were generally more accurate than those made by the original paper. There are a number of reasons this could be the case. The first is the size of the datasets, the PSEUDO-AMIGA dataset is nearly 3000 images while the one used in the paper was 500 images. The second is that the masks in the dataset were made using Sam V2, and while every single image was hand checked to make sure the masks fit the images they are not perfect and I certainly missed some smaller weeds trying to get through all of them. The third is the size of the images. The paper uses a 704x704 resolution while the images in this project are 1280x1024, so each image has 815,104 more pixels of information in the PSEUDO-AMIGA dataset when compared to the paper's set.

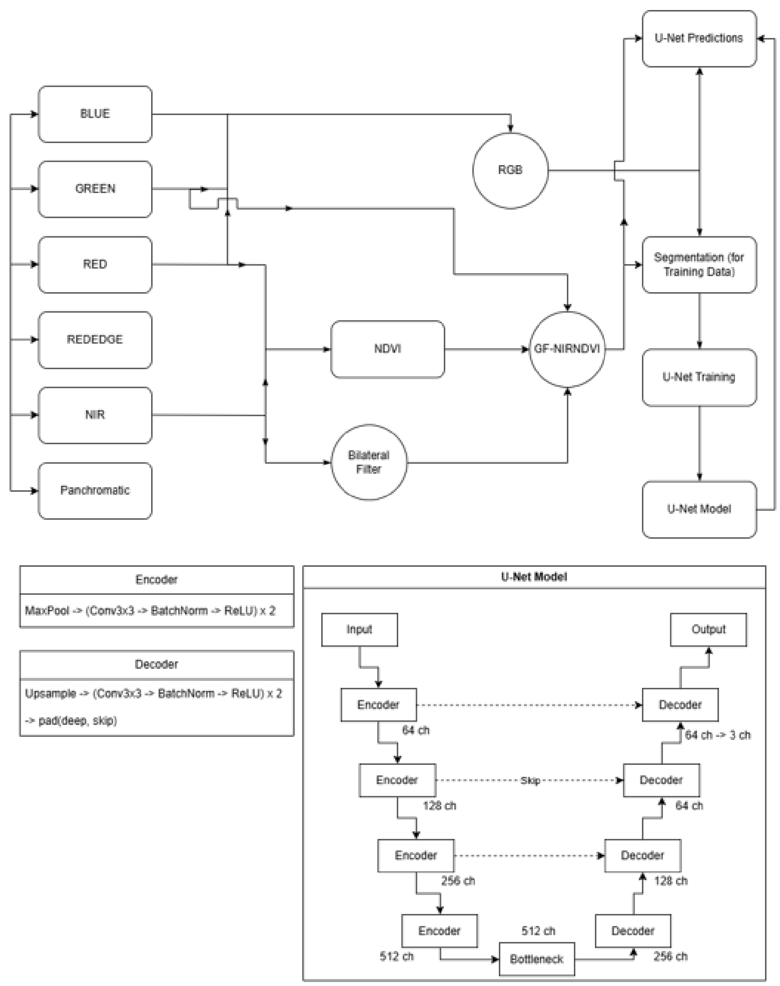

Visualized Pipeline from Data Collection to U-Net


Real Predictions from RGB and GF-NIRNDVI Models

1.5 Future Work

1. Integrate the Micasense camera into the AMIGA software and mount it on the robot
2. Convert the pipeline from an external, laptop-centric approach to one functioning on the AMIGA system
3. Create a survey function on the robot that paths through the field and sends images down the pipeline (Bands to GF-NIRNDVI to Prediction Model to Mask) and marks problem areas (high weed/crop ratio) on a GPS based UI.
4. Use collected data set to train for other phenotyping subjects (disease, canopy coverage, nitrogen levels)

2 Value of Capstone Project

2.1 Reid Group Mission for AMIGA

The goal of this group is to experiment with the Farm-ng AMIGA robotic platform with various projects that build upon existing AMIGA-based autonomy infrastructure. My work contributed to that mission by laying the groundwork for phenotyping plants using the platform, by preparing the MicaSense Rededge-P for use, creating a tool to collect and segment data, making a 3000 image dataset with this tool, and a U-Net model to identify plants as a crop or a weed. My work will allow future projects involving multi-spectral imaging to go much smoother, having already collected sample data to work with and implementing more ways to obtain it as well as including research and implementations of some tools used to phenotype plants in other ways such as canopy coverage, nitrogen levels, and disease detection, all of which is possible through a multi-spectral approach.

2.2 Connecting with My Studies

This capstone project related to my studies in numerous ways. I was able to directly apply skills learned in my Computer Vision course last semester to numerous areas of this project, from the creation of a U-Net to handing image file types to using homography to align images together, the knowledge from that class was vital to completing many of the tasks laid out in this project. My embedded systems course was also very helpful, having taught me more about how electrical signals from hardware such as sensors and other I/O are processed by a computer made it easier to understand some of the more hardware oriented problems I faced during this project, such as setting up the PoE connection.

2.3 Connecting to Current Issues in Autonomy and Robotics

My project directly examines and tests results of the paper "Segmentation of weeds and crops using multispectral imaging and CRF-enhanced U-Net"[8] published in 2023, lending evidence to its conclusion on multi-spectral imaging being more accurate when it comes to training U-Nets for segmentation on weeds and crops. In addition, one of the biggest issues in Autonomy currently is gathering data. A paper in 2026 found that using tools such as SAM V2 accelerated annotation time for masking by up to 6.1 times compared to polygon drawing while often being better than the hand segmentations [1]. This project reflects the power this tool has on relieving the bottleneck of data, allowing me to segment almost 3000 images to a relatively accurate level in just 4 days of near continuous work.

2.4 Opportunities to Network and Learn

The most significant opportunity I had to network during this project was taking a trip to AUTOMATE 2026 in Chicago during the early weeks of the semester to investigate the current status of agricultural robotics and to see if experts there had any advice or solutions to my hardware integration, or even a recommendation for a better multi-spectral camera for my project specifications. While I came up short on finding a better solution I did meet with dozens of companies and made several connections via the trip.

Learning more about hardware and how to wire and convert signals compatible but obfuscated signals in low level ways was the most difficult portion for me. I have a background in Computer Science and much of the hardware I've had to work with was simply provided for me in working order, so to be faced with an issue like that and having to solve it myself as a portion of my project was fairly refreshing and a reason I made the switch to Robotics: to be more hands on like that.

3 Reflection on Experience

3.1 How Did This Project Affect Me?

This project taught me a lot about working by myself with oversight. Most of the projects I have worked on in my educational career have been group projects, which made this one ending up as a solo endeavor a little intimidating. I learned how to keep up with my goals by scheduling and keeping track of my progress and how to communicate effectively to get what I need. I was originally not as persistent as I could have been in making sure hardware procurements would be made in time and because of that and some missed communications my project became significantly delayed due to issues getting hardware ordered. This experience taught me that auditing progress is key when communicating needs, and that when I am the sole member of a project I am the only one that can be relied on to keep it on track.

3.2 What Were My Goals?

My personal learning objectives for this project were to...

1. Develop an understanding of multispectral imaging principles and evaluate how different spectral bands can be used to characterize crop health and vegetation characteristics.

2. Analyze the effectiveness of NDVI, NDRE, GNDVI, and OSAVI as indicators of plant vigor, chlorophyll content, and crop stress.

3. Construct and annotate a multispectral agricultural image dataset suitable for training and evaluating deep learning models.

4. Investigate the impact of multispectral imagery and vegetation indices on semantic segmentation performance by comparing U-Net-based models trained with RGB imagery and derived vegetation indices.

Overall these goals were absolutely met through the course of the project with the exception of 2, which became more of a stretch goal as the project was delayed. I did, however, analyze the effectiveness of NDVI in distinguishing between weeds and crops which is a marker of crop stress, so this was partially met as well. The general theme of these goals were to just learn more about the use of computer vision in the agricultural robotics industry and I was able to not just learn but participate in experimentation with these technologies. I do wish I was able to test my experimental approach of feeding raw band data as channels into the U-Net rather than just following and confirming results from a paper, but I am satisfied with what I was able to learn from said results and I feel that is more valuable.

3.3 How Was My Perception on Autonomy and Robotics Impacted?

In researching for this project I was surprised about how little variation there was in approaches to robotic issues on the hardware side. Most solutions come in the form of what is essentially either an oddly shaped autonomous car or an arm, with small exceptions in things like soft skin crawling bots or the Boston Dynamics robotic dog (which even then, is an oddly shaped autonomous car with arms instead of wheels). It was especially clear at AUTOMATE that the industry has a clear direction and idea of what robotics will mean in the future. It has greatly changed my perception of autonomy and robotics as a concept that, rather than being mostly research and development, is one with clear standards and practices. Robotics is no longer just a technology of the future but also a technology of today, and, more surprisingly, of the past.

3.4 How Has This Prepared Me For a Career in Autonomy and Robotics?

This capstone project has prepared me for a career in Autonomy and Robotics by giving me first hand experience into planning, communicating, and executing a project and how complicated it is to translate between those skills. It has given me the opportunity to take a general idea and conceptualize it, find a way to pitch that idea in a way that sounds enticing enough to accept, and to then realistically put that idea together in the real world. It has also prepared me by teaching me to check up on requests that will require time, money, and approval to complete, making sure they are done in a timely fashion rather than being lost in an inbox somewhere, especially when I am the sole advocate for my work.

3.5 My Greatest Take Away?

My greatest takeaway from this project is that I learned what it actually means to take ownership of a robotics project. Before this, I mostly thought of robotics in terms of the technical problems. This project showed me that those things are only part of the job. I had to figure out problems that I had never dealt with before, keep myself motivated when things were delayed and make decisions about what was realistically possible with the time and resources I had.

I am also proud that, even though I did not complete everything I originally planned, I was still able to leave something useful behind. I started with the goal of preparing a multi-spectral phenotyping system and ended up with a working data collection process, a dataset of nearly 3,000 image sets, segmentation masks, and models that demonstrated the potential of multi-spectral imagery for distinguishing crops from weeds. More importantly, I created a foundation that future students can build on rather than having to start from the same problems I encountered.

The project also changed how I view my own place in robotics. Coming from a Computer Science background, I initially expected the software and machine learning portions to be the parts I was most comfortable with. Instead, some of the most valuable experiences came from the parts I knew the least about, particularly the hardware integration and the process of figuring out why something simply would not work. That experience reinforced why I wanted to move toward robotics in the first place. I want to work in a field where I am not limited to working entirely within software, but where I can understand how the different pieces of a system interact and help make the entire system work.

4 References

1. Oehme, L. H., Boysen, J., Wu, Z., Stein, A., & Müller, J. (2026). Orchestrating segment anything models to accelerate segmentation annotation on agricultural image datasets. *Frontiers in Artificial Intelligence, 8*, 1748468. https://doi.org/10.3389/frai.2025.1748468

2. farm-ng. (n.d.). *AMIGA system overview.* https://amiga.farm-ng.com/docs/concepts/system_overview/

3. NASA Earth Observatory. (n.d.). *Measuring vegetation (NDVI & EVI).* https://science.nasa.gov/earth/earth-observatory/measuring-vegetation-ndvi-evi

4. Rondeaux, D., Steven, M., & Baret, F. (1996). Optimization of soil-adjusted vegetation indices. *Remote Sensing of Environment, 55*(2), 95–107. https://www.sciencedirect.com/science/article/abs/pii/0034425795001867

5. Viña, A., & Gitelson, A. A. (2011). Sensitivity to foliar anthocyanin content of vegetation indices using green reflectance. *IEEE Geoscience and Remote Sensing Letters, 8*(3), 464–468. https://doi.org/10.1109/LGRS.2010.2086430

6. GISCarta. (n.d.). *NDRE (Normalized Difference Red Edge Index).* https://giscarta.com/gis-glossary/ndre

7. MicaSense. (n.d.). *User guide for MicaSense sensors.* https://support.micasense.com/hc/en-us/articles/360039671254-User-Guide-for-MicaSense-Sensors

8. Sahin, M., Angelov, P., & Soares, A. (2023). Segmentation of weeds and crops using multispectral imaging and CRF-enhanced U-Net. *Computers and Electronics in Agriculture, 214*, 108279. https://www.sciencedirect.com/science/article/pii/S0168169

9. Su, Y., et al. (2021). Ir-UNet: Irregular segmentation U-shape network for wheat yellow rust detection by UAV multispectral imagery. *Remote Sensing, 13*(19), 3892. https://www.mdpi.com/2072-4292/13/19/3892

10. Singh, A., Ganapathysubramanian, B., Sarkar, S., & Singh, A. (2019). A spatio-temporal spectral framework for plant stress phenotyping. *Plant Methods, 15*, 13. https://link.springer.com/article/10.1186/s13007-019-0398-8

11. Wu, J., Zhang, X., & Chen, Y. (2026). *Comparison of U-Net++, DeepLabV3+, and SegFormer for UAV multispectral disease mapping* [Preprint]. arXiv. https://arxiv.org/abs/2606.06359

12. Darbyshire, M. D., Salazar-Gomez, A. S., Gao, J. G., Sklar, E. I., & Parsons, S. P. (2023). Towards practical object detection for weed spraying in precision agriculture. *Frontiers in Plant Science, 14*. https://doi.org/10.3389/fpls.2023.1183277

13. Ultralytics. (n.d.). *SAM 2: Segment Anything Model 2*. Ultralytics Documentation. https://docs.ultralytics.com/models/sam-2

14. Haber, A. (2024, July 31). *Correctly install in Python and use Meta Segment Anything Model 2—SAM2—Computer vision model* [Video]. YouTube. https://youtu.be/MIUxiLjoA1g

15. Milesial, M. (n.d.). *Pytorch-UNet: PyTorch implementation of the U-Net for image semantic segmentation with high quality images* [Computer software]. GitHub. https://github.com/milesial/Pytorch-UNet