Cichlid Computer Vision Project – Weekly Progress

Week ending Friday, February 7th, 2025

Time Log

Charlie Clark

What progress did you make?

- Attended weekly BioBoost meeting on Monday evening.
- Attended weekly publication seminar Tuesday evening.
- Attended weekly HAAG admin meeting Thursday.
- Continued literature review.
- Re-ran SORT on the Lindenthal dataset using Kailey's Jupyter notebook. o Found that there seemed to be a considerable number of suboptimal IoU calculations, leading to detections being dropped.
- Helped Bree work on PEARC paper. o Paper was successfully submitted before the deadline Saturday night.

What are you planning on working on next?

- Attend required weekly meetings.
- Attend optional weekly meetings.
- Explore/understand the YOLO aspects of the BioBoost project. o Move YOLO scripts from CichlidBowerTracking repo to BioBoost repo.
- Re-read BioBoost paper.
- Attend publication seminar on Tuesday.

Is anything blocking you from getting work done?

- **None**

Researcher

What did you do this week?

- (1) I attended the following 2 meetings:
    - (a) Cichlid Team Meeting on February 3rd. Inconsistent frame rates and frame skips were discussed.
    - (b) Publication Seminar on February 4th. Bree wasn't there. Random logistics discussed.
- (2) I worked on the meeting manager role requirements for the weekly meeting. I created the meeting task in Microsoft planner, created the required materials

(slides, transcripts, notes, recordings, attendance), updated the master attendance sheet, and updated the project files section on Teams. I uploaded the information into the planner section on Teams and the Slack Meeting-Manager and BioBoost channels.

- (3) I was able to take the SORT code and edit the IoU threshold to 0.15 and 0.05, resulting in 65k and 131k more detections. I spent several hours trying to fix the visualization code to represent the new changes.
- (4) I investigated some problems with the BioBoost pipeline. Specifically, I looked into the variable frame rate issue. The variable frame rate is present in the videos from Charlie prior to my SORT and OpenCV code, so it is likely a variable-frame-rate camera issue or a bag file extraction issue. It is only noticeably present in the daytime video.

What are you going to do next week?

- (1) I need to work with Charlie and Eric to get the frame inconsistencies and frame skips figured out.
- (2) I need to continue to modify the SORT code. Specifically, I need to fix the visualization code to be able to see the changes that the new IoU threshold resulted in.
- (3) I should help with the paper rewrite as needed.
- (4) I need to fulfill my meeting manager responsibilities and attend required meetings.

Is anything blocking you from getting work done?

- Yes. I still need an email from Dr. Lytle for CS 8903. My advisor is still waiting. (!!!)

Eric Iamarino

What did you do this week?

- Added Dockerfile to .bag conversion scripts so they can run in specialized OS environments on any device
- Added run_conversion.sh script which allows the Dockerfile to be reused for the python and shell conversion scripts
- Backfilling website with Weeks 1-3 Weekly Reports/Meetings
- Added Week 4 Weekly Reports to Website
- Added Week 5 Weekly Meeting Update to Website
- Attended BioBoost Weekly Meeting

What are you going to do next week?

- Attend BioBoost meeting
- Attend publication meeting
- Write & test changes to extract_videos.sh/py scripts
- Help with BioBoost rewrite

Blockers, things you want to flag, problems, etc.

- Dropbox access

Abstracts

Charlie Clark

"Forecasting insect abundance using time series embedding and machine learning", Palma et al. (2025; Ecological Informatics0

- Abstract: "Implementing insect monitoring systems provides an excellent opportunity to create accurate interventions for insect control. However, selecting the appropriate time for an intervention is still an open question due to the inherent difficulty of implementing on-site monitoring in real-time. A possible solution to enhance decisionmaking is to apply forecasting methods to predict insect abundance. However, another layer of complexity is added when other covariates are considered in the forecasting, such as climate time series collected along the monitoring system. Multiple combinations of climate time series and their lags can be used to build a forecasting method. Therefore, we propose a new approach to address this problem by combining statistics, machine learning, and time series embedding. We used two datasets containing a time series of aphids and climate data collected weekly in two municipalities in Southern Brazil for eight years. We conduct a simulation study based on a probabilistic autoregressive model with exogenous time series based on Poisson and negative binomial distributions to evaluate the performance of our approach. We pre-processed the data using our newly proposed approach and more straightforward approaches commonly used to train learning algorithms. We evaluate the performance of the selected algorithms by looking at the Pearson correlation and Root Mean Squared Error obtained using one-step-ahead forecasting. Based on Random Forests, Lasso-regularised linear regression, and LightGBM regression algorithms, we showed the feasibility of our novel approach, which yields competitive forecasts while automatically selecting insect abundances, climate time series and their lags to aid forecasting."

- AI Summary by ChatGPT (4o): “This paper investigates the impact of asynchronicity on student community-building in large-scale online education programs. Through three independent studies within an online graduate program, the authors identify a paradox: students desire synchronous interactions to foster a sense of social presence, yet the flexibility of asynchronous learning is one of the main reasons they enroll in such programs. This Synchronicity Paradox suggests that online education providers must balance flexibility with opportunities for real-time student interaction.”
    - Key Contributions
        - “Identification of the Synchronicity Paradox: While students appreciate asynchronous learning for its flexibility, they also report feeling socially isolated due to the lack of real-time engagement. This paradox challenges assumptions that online education inherently fosters weaker student communities.”
        - “Three Studies on Social Presence in Online Education: Study 1, surveys tracking social presence across a semester revealed that students felt neither deeply connected nor entirely disconnected from peers. Study 2, needfinding surveys and interviews highlighted a growing sense of isolation over time, with students expressing a desire for more real-time interactions. Study 3, a controlled experiment with immersive virtual environments (lecture halls, social lounges) found that students valued synchronous engagement, but technical barriers reduced participation.
        - “Emergent Synchronicity as a Solution: The authors propose leveraging the scale of online education to enable emergent synchronous interactions, where students naturally find real-time study partners without strict scheduling. Real-time chat tools and collaborative study groups could create spontaneous, non-mandatory synchronous interactions.”
    - Contributions to Knowledge:
        - “Challenges the Long-Held View of Online Isolation: Unlike past studies that attributed online student isolation to a lack of instructor interaction, this paper suggests asynchronicity itself is a major cause.”
        - “Bridges Online and Traditional Learning Paradigms; By proposing emergent synchronicity, the authors offer a way to retain the flexibility of online education while introducing community-building elements akin to in-person learning.”

        - "Provides Design Guidelines for Online Programs: The paper suggests using real-time messaging, virtual co-watching, and informal discussion spaces to enhance student engagement without enforcing rigid schedules."
    - Future Research Directions:
        - "Scalability of Emergent Synchronicity: Investigate how large-scale programs can implement real-time interactions without disrupting the benefits of self-paced learning."
        - "Automated Student Matchmaking for Collaboration: Develop AI-driven systems that pair students based on their learning schedules and course progress to foster organic study communities."
        - "Longitudinal Studies on Social Presence in Online Education: Analyze how student engagement evolves over multiple semesters and its impact on retention rates and academic performance."
        - "Synchronous vs. Asynchronous Engagement in Different Learning Contexts: Study whether certain disciplines (e.g., STEM vs. humanities) require different balances of synchronous and asynchronous interaction."

• Link: https://www.sciencedirect.com/science/article/pii/S157495412400476X?ref=pdf_download&fr=RR-2&rr=91008cc39dee433f

Researcher

Jalal, A., et al. "DeepFins: Capturing dynamics in underwater videos for fish detection." Ecological Informatics, 2025. https://www.sciencedirect.com/science/article/pii/S1574954125000226.

I found this paper during the publication seminar, but we didn't really have time to go through it, so I will be doing my abstract on this paper so that I get a better understanding of what is being done. This research developed a new model called DeepFins, which combines a well-known object detection system (YOLOv11) with a motion-based segmentation technique, achieving high accuracy in detecting fish in underwater videos. The OzFish and LifeCLEF 2015 datasets were used. The OzFish dataset contains over 3,000 videos with annotated frames for fish detection, while the LifeCLEF 2015 dataset includes videos of 15 different fish species. Both datasets present challenges like moving backgrounds and blurriness, which the proposed 1 hybrid feature extraction

method aims to address by combining static and motion-based detection techniques. Features from a temporal branch are merged with high-level feature maps from YOLO, and the fish blobs are selected based on confidence score. In the future work section, the authors discuss the possibility of including more spatio-temporal features to improve classification accuracy in visually challenging scenes. This paper and others like it seem to present a challenge to our BioBoost project, since they are similar and more robust than our research.

Eric Iamarino

Lingli Chen, Gang Li, Shunkai Zhang, Wenjie Mao, Mei Zhang, YOLO-SAG: An improved wildlife object detection algorithm based on YOLOv8n, Ecological Informatics, Volume 83, 2024, 102791, ISSN 1574-9541, https://doi.org/10.1016/j.ecoinf.2024.102791.

Wildlife conservation is crucial for maintaining biodiversity, ensuring ecosystem balance and stability, and fostering sustainable development. Currently, the use of infrared camera traps to monitor and capture photos of wildlife is a vital methodology in protecting and researching wildlife, and automatic detection and identification of animals within captured photographs are paramount. However, factors such as the complexity of the field environment and the varying sizes of animal targets lead to low detection accuracy, while high-precision detection models are hindered by high computational complexity and sluggish training speeds. This paper proposes a wildlife target detection algorithm based on improved YOLOv8n - YOLO-SAG, which aims to balance accuracy and speed. Training stability is enhanced by introducing the Softplus activation function, which increases detection accuracy; incorporating the AIFI enhances intra-scale feature interaction, reducing missed and false detections. Integrating the GSConv and VoV-GSCSP module lightens neck convolutions, reducing computational redundancy and balancing the computational and parametric quantities brought by the AIFI. Experimental results on a self-made wildlife dataset indicate that the YOLO-SAG achieves 94.9%, 90.9%, 96.8%, and 79.9% in Precision, Recall, mAP@0.5, and mAP@0.5–0.95, respectively, which are 3.4%, 3.3%, 3.2%, and 4.9% higher than the original YOLOv8n. Inference and post-processing times reach 1.2 ms and 0.5 ms, a speedup of 25% and 54.5%, respectively, and the computation volume is only 7.2 GFLOPs, an 11.1% decrease.

Documentation of Work

Charlie Clark

- Attended weekly BioBoost meeting on Monday evening.
- Attended weekly publication seminar Tuesday evening.
- Attended weekly HAAG admin meeting Thursday.
- Continued literature review.
- Re-ran SORT on the Lindenthal dataset using Kailey's Jupyter notebook.
    - Found that there seemed to be a considerable number of suboptimal IoU calculations, leading to detections being dropped.
- Helped Bree work on PEARC paper.
    - Paper was successfully submitted before the deadline Saturday night.

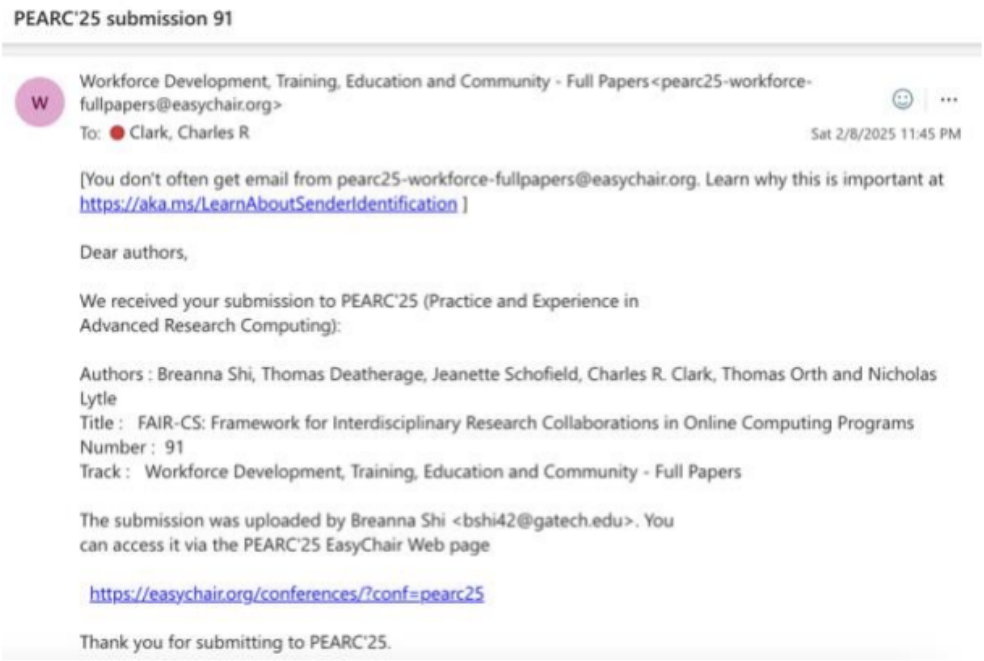

PEARC'25 submission 91

W

Workforce Development, Training, Education and Community - Full Papers<pearc25-workforce-fullpapers@easychair.org>

To: Clark, Charles R

Sat 2/8/2025 11:45 PM

[You don't often get email from pearc25-workforce-fullpapers@easychair.org. Learn why this is important at https://aka.ms/LearnAboutSenderIdentification]

Dear authors,

We received your submission to PEARC'25 (Practice and Experience in Advanced Research Computing):

Authors : Breanna Shi, Thomas Deatherage, Jeanette Schofield, Charles R. Clark, Thomas Orth and Nicholas Lytle
Title : FAIR-CS: Framework for Interdisciplinary Research Collaborations in Online Computing Programs
Number : 91
Track : Workforce Development, Training, Education and Community - Full Papers

The submission was uploaded by Breanna Shi <bshi42@gatech.edu>. You can access it via the PEARC'25 EasyChair Web page

https://easychair.org/conferences/?conf=pearc25

Thank you for submitting to PEARC'25.

4. Results Visualization: Nothing to visualize this week on my end.

Researcher

For a full list of what I did, see the time log above. The most important things I worked on this week were investigating issues with the Lindenthal dataset and modifying SORT parameters.

To investigate issues in the Lindenthal dataset, I first started by looking at the original Lindenthal videos that Charlie added to the Dropbox. I confirmed that the speedup issue existed in those videos. This means that the inconsistent frame rate issues were not caused by my SORT code or my OpenCV visualization code. In other words, this may be an extraction or a variable-frame-rate camera issue, so Charlie and Eric should continue to look into this. The dataset paper doesn't appear to mention this, but we can further inspect the paper to see if it gives the readers any hints. Next, I checked and saw that the infrared videos don't appear to have the variable frame rate problem. Or, at least it isn't as noticeable. This further points to it being a camera or extraction issue. We will need to determine if this is something that we need to fix, either by re-extraction or cutting out the first few frames of the daytime videos. If we do decide that we need to fix this, we will need to re-run YOLO and the rest of the pipeline.

The important snippets of my code this week are below. First of all, I needed to add confidence score to the bounding box function, since Charlie said that this would help with identifying some issues in our SORT code.

```
...
confidence = row['Confidence']
label = f"{class_name}: {confidence:.2f}"
font = cv2.FONT_HERSHEY_SIMPLEX
font_scale = 0.4
font_thickness = 1
text_size = cv2.getTextSize(label, font, font_scale, font_thickness)[0]
text_width, text_height = text_size
...
```

Listing 1: Add Confidence Score to SORT OpenCV Visualization

The next thing for me to do was to determine what I needed to change in SORT to produce better results. The first thing that I looked into changing was the IoU threshold in the associate detections to trackers function. The IoU threshold is the minimum overlap required for the the detection to be assigned to an existing track. Therefore, increasing the threshold will reduce incorrect matches and cause more ID switches. It can also cause objects to be lost if their position changes too much. In contrast, lowering the threshold will allow detections with small overlaps to be associated with pre-existing tracks. In our case, because we are dropping frames between YOLO and SORT, we want to probably lower the default IoU threshold of 0.3 to something like 0.2 or even 0.1. That being said, some of the frame drops when it comes to matching may be caused by the

aforementioned frame rate issue. So, if the frame rate issue is fixed, we may need to re-adjust this. To change the IoU threshold, we change the SortFish code's original SORT code in the two locations indicated below. Changing IoU threshold to 0.15 results in 65k more detections over 100 videos, and changing it to 0.05 results in 131k more detections over 100 videos. This was expected. However, what was not expected was that the visualizations looked exactly the same. I spend a few hours trying to figure out what kind of code or cache issue was causing this. I also wrote some code to compare videos side-by-side. However, it didn't help me fix the issue. I will need to spend more time figuring out why the changes aren't showing up in the visualizations.

```
...
def associate_detections_to_trackers(detections, trackers, iou_threshold=0.05):
...
```

```
class Sort(object):
    def __init__(self, max_age=5, min_hits=3, iou_threshold=0.05):
...
```

Listing 2: Changing the IoU Threshold

```
...
cap1 = cv2.VideoCapture(video1_path)
cap2 = cv2.VideoCapture(video2_path)
...
while cap1.isOpened() and cap2.isOpened():
        ret1, frame1 = cap1.read()
        ret2, frame2 = cap2.read()
        combined_frame = np.hstack((frame1, frame2))
        out.write(combined_frame)
...
```

Listing 3: Essence of Side by Side Video Comparison

And that's it! The finished scripts will be uploaded to BioBoost upon completion: https://github.com /Human-Augment-Analytics/Bio-Boost. To see the videos with confidence scores generated with OpenCV, please see the 100 videos that I added to the BioBoost channel in a zip file on February 6, 2025.

Eric Iamarino

Dockerfile

```
FROM ros:melodic

# Install dependencies
RUN apt-get update && apt-get install -y \
ros-melodic-cv-bridge \
ros-melodic-image-view \
ffmpeg \
python-opencv \
&& rm -rf /var/lib/apt/lists/*

# Set up workspace directory
WORKDIR /root/catkin_ws

# Copy conversion scripts into the image.
COPY extract_videos.py /root/catkin_ws/
COPY extract_videos.sh /root/catkin_ws/

# Make the shell script executable.
RUN chmod +x /root/catkin_ws/extract_videos.sh

# Copy custom wrapper script to select between the two scripts
COPY run_conversion.sh /root/catkin_ws/
RUN chmod +x /root/catkin_ws/run_conversion.sh

# Source ROS environment by default
ENTRYPOINT ["/bin/bash", "-c", "source /opt/ros/melodic/setup.bash && exec \"$@\"", "--"]

# Default to running the wrapper script
CMD ["/root/catkin_ws/run_conversion.sh"]
```

Dockerfile helper script:

```
#!/bin/bash

if [ "$1" == "python" ]; then
shift
echo "Running python conversion script"
python /root/catkin_ws/extract_videos.py "$@"
elif [ "$1" == "shell" ]; then
shift
echo "Running shell conversion script."
/root/catkin_ws/extract_videos.sh "$@"
else
```

```
echo "Usage: docker run ros-bag-to-mp4 {python|shell} [args]"
exit 1
fi
```

Changes live on website: https://sites.gatech.edu/cichlid-computer-vision-project/

BioBoost Meeting Recording:
https://gtvaultmy.sharepoint.com/personal/kcozart6_gatech_edu/_layouts/15/stream.aspx?id=%2Fpersonal%2Fkcozart6%5Fgatech%5Fedu%2FDocuments%2FRecordings%2FCichlid%20CV%20Weekly%20Meeting%2D20250203%5F143627%2DMeeting%20Recording%2Emp4&referrer=StreamWebApp%2EWeb&referrerScenario=AddressBarCopied%2Eview%2E6ecf7323%2D6bd6%2D40af%2D94b2%2D0216e0421a8e