Danelle Cline Sorted By Project

300057 - Mooring Maintenance

- Proposal input (Week 27)

500055 - VARS, M3/RS, DSG

- Spun-up M3 stack per the quick start instructions https://github.com/mbari-media-management/m3-quickstart (Week 3)
- Deployed the new mbari315K model server in AWS (Week 44)

901009 - CANON

901502 - Soundscape Research and Education

- Testing machine learning workflow on a subset from 2017 on 53K images for final approval from team, then final 6 year/2 million blue whale A detection proposals. This workflow uses AWSBatch to generate spectrograms and SageMaker BatchTransform Jobs to quickly run inference. (Week 3)
- Project pitch to Cal Poly ML class (Week 3)
- Our livestream audience is growing; the service has almost exceeded its 2500 THL limit in Shoutcast and may bump to the next tier automagically. (Week 3)
- AWS OpenData 2022 updates (Week 5)
- Discussion with grad student at University of Denver about potential collaboration on miniML sound detection (Week 5)
- Debugging an issue with blue A inference workflow (Week 5)
- S3 bucket sync management for logs access (Week 17)
- Setup Amazon Athena database and working on queries to assess OpenData program data usage (Week 17)
- Group meeting discussion on click train detection (Week 17)
- Completed setup and deployment of AWS Athena queries for our Pacific Sound opendata use. Queries revealed in April, 2022, our sound data is being used in 5 countries, with our raw 256kHz recordings from 2020 being the most accessed with 1 Petabyte of downloads in April. (Week 19)
- Discussions with the team about the upcoming blog for World Oceans Day on June 8th. (Week 19)
- Working on blog materials for upcoming World Ocean Day with John, AWS partners and ITD group. (Week 20)
- Switching audio streaming to a new production Windows server shoutcast.shore.mbari.org in response to changes in our streaming provider. (Week 20)
- Finished ansible deployment of Aurora data use workflow (Week 20)
- Shoutcast stream DNAS software installed and tested; more work to open network ports and bring some old python code up to date (Week 21)
- World oceans day blog writing in coordination with ITD/AWS (Week 21)
- Finished deployment of shoutcast stream. Peak listeners at 88 simultaneous; average listening time of 5 hours, 15 minutes and 14 seconds thus far. Zen listening to help calm the mind. (Week 24)
- Maintenance on shoutcast port 8000 (Week 27)
- Project proposal input (Week 27)
- Setup shoutcast Windows streaming service for restart on reboot on shoutcast.mbari.org (Week 31)
- Team meeting (Week 31)
- AthenaDB monthly query running as expected, but daily bump in S3 cost remains a mystery (Week 31)

- Finished demo Colab jupyter notebook on running blue D ResNet classifier to kick-start efforts with with Stanford student, Kanoe Makalii Aiu, who will be working with Will Oestrich and Danelle on improvements to the blue D call ML workflow. (Week 39)
- Meeting with Kanoe to do walk through (Week 39)
- Potential upcoming opportunty to present work https://www.frontiersin.org/research-topics/45485/deep-learning-for-marine-science (Week 39)
- Updates to the open data documentation (Week 42)
- Team meeting and meeting with Stanford student Kanoe on continued ML work with blue whale calls. (Week 42)
- Investigated live stream interruption and submitted a help ticket with details. IS has resolved this issue by updating a web certificate on the streaming server shoutcast.mbari.org. (Week 43)
- Assisting collaborator, Karl Watkins on potential export of Google Humpback model to his embedded platform (Week 43)
- Assisting collaborator, Kanoe Makalii on blue whale D call exploration; installed and test drove the MorphoCluster tool on some sound spectrograms. We are looking for ways to help organize calls faster for her training experiments. (Week 43)
- Participated in SoundCOOP - Python Sniffers telecon, focused on PyPAM module which is a potential community module for creating standard sound measurements similar to the Manta Matlab system. Signed up for new slack channel OceanObservationData for further discussions. (Week 44)
- Weekly meeting. Thanks Carlos and Duane for improvements to procedures to monitor the live stream with an external service. Our stream fan club will be so pleased. (Week 44)
- Working on debugging an issue with blue whale spectrogram color map inversion RGB to BGR which is causing poor model performance. (Week 44)
- Assisted student Karl Watkins with some feedback on how to evaluate his humpback model designed to work on tiny devices (e.g. Jetson Nano). (Week 44)

901602 - Data TAG

901603 - Video TAG

- Finished draft of abstract for Marine Imaging Workshop and initial cut at a graphic. (Week 19)
- Read project abstracts (Week 20)
- Abstract submission for 2022 Marine Imaging Workshop (Week 20)
- Reviewing 2023 proposals (Week 31)

901915 - Improving the Impact of MBARI's Data

- Meeting to discuss M1 mooring data and page migration (Week 39)
- Working on FathomNet web page for new web design (Week 43)
- Worked on FathomNet page upgrade; more work to finish user profiles then pass on for team feedback. (Week 44)
- Need to start DeepSea Guide page upgrades. (Week 44)

902005 - VARS Annotation Assistance

- Pushing hard to get a webapp up and running on deepsea-ai. This will be our ML dashboard with models, processing jobs, etc. all in one place. Vue3 front-end and postgres db backend design. Hope to get some help from Carlos to take it further. (Week 3)
- Finished bulk loader to quickly load machine results to summarize and manage the job data and experimental tests. (Week 3)
- A preview of the machine generating tracks in BoxJelly look promising! (Week 3)
- Planned discussion with AWS ProServe team next week (Week 3)

- Working on infrastructure roll-out and back-end deployment (Week 5)
- Finished dockerized deepsort tracker which the team is testing (Week 5)
- So far so good. Hoping to have bulk batch testing finished next week. (Week 5)
- Tracking pipeline benchmarking and optimization with Duane. Current MOTA is ~53%. Need more fully annotated videos to measure accurately. BoxJelly is a potential tool for this. (Week 17)
- Syslog setup on production server deepsea-ai to start logging dashboard and database API use. (Week 17)
- Fix issues with postgres after sql server migration to M3_TRACKS database broke a few things. Improved handling of AWS security keys and S3 track data sync/loads. (Week 17)
- Next steps are to return the batch processing, inserting the optimized tracking docker container into the elastic EC2 pool that I paused on to focus on database migration. (Week 17)
- Automated deployment of backend/frontend code using Ansible now in place. (Week 19)
- Working on adding supporting code for batch processing to the Dockerized deployment of the video processing pipeline (Week 19)
- Added support for mirrored Titan prefixes with unique ISO tags in processing workflow S3 buckets. Testing looks good on that. (Week 20)
- Moved YOLOv5 detection-tracking docker image to docker hub instead of ECR; added linux/amd64 and linux/arm64 builds. (Week 20)
- Testing batch CDK stack. Working through issues with IAM, but expect this work to be finished soon. (Week 20)
- Finished up ansible deployment of deepsea-ai stack. (Week 20)
- Drafting jupyter notebooks to help the team with integration/editing against tracks database (Week 21)
- Ongoing work to test micro EC2 auto-scaling CDK stack is going well with short 1-minute video clips processing. (Week 21)
- Zoom call with summer intern Manual Gozzi (Week 21)
- Batch processing going well. AWS elastic compute cluster is scaling and fine-tuning almost complete. (Week 24)
- Added support for Job deletion and autoloading by user of processed data. (Week 24)
- Added more detail to jupyter notebooks on track database use (Week 24)
- (virtual) team meeting due to COVID (Week 24)
- Intern support (Week 27)
- Added cluster tab to dashboard http://deepsea-ai.shore.mbari.org/#/clusters to monitor high-level information about AWS ECR clusters like tasks and queue messages; some other updates to dash to improve track auto playback. (Week 27)
- Improvements to batch processing graceful termination (Week 27)
- Finished PR5 https://bitbucket.org/mbari/vaa-track-pipeline/pull-requests/5 a minor performance improvement to batch process faster. (Week 31)
- Merged PR13 https://bitbucket.org/mbari/deepsea-ai-backend/pull-requests/13 a series of performance improvements to the CDK stack to speed-up tear up/down. Cloud processing time now at 15-40 minutes per 15 minute video with minimal overhead for spin up/down of EC2 GPU instances. (Week 31)
- No team meeting this week :((Week 31)
- MIW22 Presentation finished and dry-run with team (thanks, Kevin, Brian, and Lonny for the good feedback and shared materials!) (Week 39)
- Moved all docker repositories from private to public ECR for use; finally figured out how to run automated github action docker releases. One small step for womankind. (Week 39)
- Fixed issue 3 in strongsort pipeline (Week 39)
- Exploring demo notebook using FathomNet data with deepsea-ai pip module (Week 39)
- Working on tutorial on FathomNet object detection training in the deepsea-ai module (Week 42)
- Presentation on cloud-enabled benthic annotation at the Marine Imaging Workshop; PDF of that is here. This was well received with questions on cost, time, data augmentation efforts, etc. (Week 42)
- Mostly testing in new AWS subaccount (using MBARI login) which gives us better cost reporting and control of our production environment: (Week 43)

    - Tested setup and training command in DeepSeaAI pip module in our new AWS subaccount; data prepared and uploaded, but waiting on quota increase to try full training run on all of the training data (315K localizations) on an 8 GPU instance. (Week 43)
    - CDK object detection and StrongSort tracking pipeline stack deploys ok in new AWS subaccount; added support for more flexible stack config through environment variables. (Week 43)
- Set up the deepsea-ai python module in our new VAA production subaccount (thanks Kevin G. and Adriana, especially for setting this up!). After patiently waiting for quota to be approved, used this to train the new mbari315K model on an 8 GPU p3.xlarge instance. (Week 44)
- Tested the new mbari315K model on short 1 second videos per the instructions. Fixed a few bugs. Release notes can be found in the CHANGELOG. Duane is doing further testing with some select dives. (Week 44)

902202 - MBARI Cloud Integration Exploration

- Discussion with Duane about a brainstorming session with Kevin. Possible idea to use STOQS parquet data in data lake as proof of concept. (Week 20)

Out of Office

- May 27th, 28th,
- Marine Imaging workshop 3 - 7 Oct 2022 CEST Brest, France (GMT+2)
- Vacation July 7-28th, back in August. 1st.
- Marine Imaging workshop 2022 (MIW22) 3 - 7 Oct 2022 CEST Brest, France (GMT+2)
- Vacation 9/28 - 9/29, 10/10-11

Overhead

- Performance review summary (Week 3)
- Updated htcmastervm-c upgrade (need to ask Pat to change machine name) (Week 5)
- Attended AWS Summit Conference in San Francisco, Moscone Center. Met with non-profit representative, attended sessions on AWS RoboFlow to support Giancarlo, database and SageMaker sessions. (Week 17)
- Discussion with Giancarlo about cloud best practices and emerging technologies (Week 19)
- Moving my office is going slow; I hope to finish by early next week. (Week 20)
- Project update presentation (Week 21)
- 2023 project proposals drafts (Week 27)
- Coordination on logistics with ITD for upcoming code.org interview and location shoot. (Week 31)
- Coordination with RoboFlow on community demonstration. (Week 31)
- Code.org video is complete & some social media activity to follow (Week 39)
- Code.org video released this was part of a series of career highlights targeting K-5, promoting diversity in computer science. (Week 42)