

Danelle Cline Sorted By Project

901103 - Biodiversity and Biooptics of Zooplankton

- Finished Argo ONNX based workflow <https://github.com/mbari-org/argo-workflows> with optional vector database integration (Week 2)
- Training new RF DETR megadetector model with targeted depth range 200-2000 meters (Week 2)
- Processing all of 2018 DocRickets dives with Argo workflow on kubernetes cluster. Roughly 2800 video files. When processing is complete, a report with associated metadata will be generated for closer inspection. (Week 4)
- Mined 2018 and 2019 DocRicketts dives at 3 second stride for species. Preparing report for Haddock. (Week 5)
- ~8K video files processed from 2018-2020 for data mining through k3. (Week 6)
- Added VARS media tap loader to unify downloads with aidata (Week 6)
- Retrained ViTS model for classes with > 100 examples; .89 balanced accuracy (Week 7)
- Tried retraining ViTS model with DinoV3 but ran into troubles with transformer library support and huggingface model gating is a pain. (Week 7)
- Repopulated vector database using ViTS model (Week 7)
- Finished argo workflows for rerunning ViTs and populating with ancillary data from the expedition database for analysis. This will take some time to rerun (classification only), but needed as the last model did not perform well for target species. (Week 7)
- Setup AWS bucket for intelligent tiering storage for Tator mantis instance storage (Week 14)
- Transcoded 2019/2018 Doc Rickett dives with 1 second keyframes; started video load which will take a while (Week 14)
- Summer intern project planning (Week 17)
- Assisting UCSC student with Velella project (Week 17)
- Megadetector exported to ONNX format and running on the modelhub. (Week 20)
- Help ticket issued to migrate data to on-site object storage. Waiting for Neil to return to test. (Week 20)
- Working on initializing VSS (Vector Search System) with four augmented velella conaria examples. Next steps are vector search in bulk, similar to how velella were mined in the planktivore archives. (Week 27)
- De-duplication code to identify individuals in tracked velella (Week 27)
- Running VSS (Vector Search System) on planktivore low-mag to search for next stage Velella conaria with augmented examples. So far, none have been found but we have a lot of data to look through. (Week 28)

- Improved memory management of REDIS to reduce storage overhead (Week 28)
- Data mining for *velella conaria* using vector database VSS (Week 29)
- Migrated data out of AWS S3 to qumulo object storage. (Week 29)

901221 - Midwater Time Series

- The RF DETR model trained with pseudo labels did not perform well. Removing generic classes like “marine organism” improved classification performance ~10%. (Week 2)
- Loaded training data into Voxel51 Community. Check out a demo on exploring i2MAP training data in Voxel51 Community. It's very intuitive and fast to interact with embeddings through lasso queries around clusters - one can see where the lines blur between annotations, where outliers are, and mistakes are made. This is a better way to interact with our data than a grid-only view in my opinion. (Week 2)
- Processed 2018 dive with Argo RF-DETR Boostrack workflow and loaded to i2map.shore for analysis (Week 4)
- Outlined three general workflows; discussed combined use of Tator+Voxel51 to quickly preview results (Week 6)
- Added embedding output endpoint to vector database service VSS; creating Voxel51 sync service for integration with Tator. This will allow one to quickly search, filter, and visualize visual data in an integrated way with the metadata, e.g. depth, temperature, etc. with a central database. See demo of Voxel51 on i2MAP training data embeddings. (Week 8)
- Deployment of new VM cortex with proxy gateways for various services, including a new Voxel51 sync service (Week 9)
- Setup Tator project config in preparation for running workflows outlined in week 6 (Week 14)
- Finalizing RF-DETR k3s (kubernetes) track pipeline. Loaded latest dive to mantis.shore. (Week 17)
- Retrained classification model to better remove many mucous strings in the latest i2MAP dive. Model exported to ONNX format and running on the modelhub. (Week 20)
- Reprocessed latest dive using kubernetes through the 3-stage RF-DETR, DINO + Boostrack pipeline. Exported to Voxel51 for review. (Week 20)
- Proposal discussion. (Week 20)
- Ran 3-stage (detect-classify-track) k3s pipeline on latest dive. Data loaded to Mantis and exported to Voxel51 for review. Used a retrained classification model with mucous strings which were very abundant in this dataset but not classified correctly. (Week 21)
- Proposal discussion and updated docs (Week 23)
- Deployment of stratagrid to visualize training data (Week 25)
- Loading i2MAP images from AWS Glacier to Qumulo S3 by depth in preparation for mining deeper depths using Voxel51 to quickly edit (Week 28)
- Export TWDA boxes to Voxel51 for review. Optimizations to fiftyone-sync service for faster processing, improved error handling, and to reduce disk and memory pressure on host doris.shore. (Week 29)

- Pulling down all frame captures from AWS glacier to load into qumulo object storage (Week 29)

901502 - Ocean Soundscape

- Realtime workflow modifications to handle OOM errors for hourly processing. (Week 2)
- Deployed Triton inference server on new VM sonus.shore.mbari.org (Week 2)
- Realtime-time workflow review with John. Triton inference server deployed and working well on new VM; 10 minute audio processes in less than 40 seconds. (Week 4)
- Revised realtime MARS pipeline with calibrated spectrogram and improved plots. Code walkthrough with J. Ryan to verify correct output per expected 5 second time resolution. Next steps will be feedback from collaborators. Triton server is working very well for this application. (Week 5)
- Bi-weekly meeting (Week 5)
- STM (Sound Topic Model) bug fix for misaligned documents (Week 6)
- CSUMB student review (Week 6)
- Starting evaluation of Whisper model embeddings (Week 6)
- Finished soundpipe which will run a series of processors on incoming MARS audio files to monitor for blue/fin/humpback presence. (Week 6)
- Reprocessed select dates with the realtime processing code soundpipe for JRyan to review. Blue, fin, and humpback model outputs agree with daily and Matlab processed results. Next step is discussion with stakeholders on outputs. (Week 7)
- CSUMB student Zoom session (Week 7)
- Thalassa outage caused some data loss (Week 7)
- CSUMB student, Gabriel Meyers, planning for service learning project on AWS (Week 8)
- Fixed issue in the syncing of backlogged data and revised the wav file size check per the latest recorder. Many thanks to Carlos Rueda for flagging problems with downstream 16K audio processing. (Week 8)
- AWS setup for CSUMB student project and project planning for 10 week program (Week 9)
- Updated recording JSON metadata in bitbucket (Week 9)
- Helped with collaborator, Franny with ROST docker build (Week 9)
- CSUMB student weekly meeting. Good progress so far. Gabriel demonstrated Athena queries on PBP generated metadata and presented a general system workflow. (Week 10)
- Upgrades to Shoutcast stream and fixes to email notifications of recording failures (Week 10)
- Mentoring CSUMB student Gabriel. He is making good progress. (Week 11)
- Setup Tator project config in preparation for loading training spectrograms (Week 14)
- Meetings to discuss realtime data use for whale monitoring(pacific coast) and interactive sound exhibits(international) (Week 14)
- Moving AWS upload to near-realtime cron instead of nightly cron (Week 14)

- Running parameter stm sweep optimization for Perch2 model. Embeddings look promising with this model. Evaluating both perplexity and ARI metrics. (Week 17)
- CSUMB student project complete. AWS Athena database with 99% coverage of soundscape metadata to provide basis for sql queries. (Week 20)
- Working on testing Perch2 model with Humpback groundtruth data. (Week 20)
- Team meeting. Discussed proposal feedback which was very positive. (Week 21)
- CSUMB student Gabriel completed his work on the Athena database. (Week 21)
- Continued work on topic model sweeps and summer intern meetings (Week 23)
- Labeling single segmented sound file for optimized STM hyperparameter sweeps. Perch2 model + Voxel51 + Raven Annotation proving useful for this task. ARI with a simple tonal benchmark reached 0.8 confirms the sweeps are good. (Week 24)
- Discussion with our new summer intern, Arie RaizmanI, Master CS student of UCSC on STM approach, hop-lite, and our overall goals. (Week 24)
- Organized team discussion for proposal input (Week 25)
- Labeled segmented units with Voxel51/Perch2 model. Refining model sweeps to optimize ARI. (Week 25)
- Oceans abstract on PBP accepted. Thanks to the team for their help with the submission. (Week 27)
- STM sweeps with Perch2 model optimizing ARI are in place. Thanks to collaborator Franny for cleaning-up the sound annotations. The single 6 and a half minute humpback recording used for testing had 32 unique sound units ! (Week 27)
- STM model sweeps with panacetacea dataset (Week 28)
- All three soundscape submissions for Oceans were accepted. Many thanks to Carlos,Duane,and John for their good work on this. (Week 28)
- Team meeting. STM optimization with panacetacea dataset. Minor maintenance on s3 cron. Drafting SMM workshop abstract. (Week 29)

901902 - Unoccupied Aerial Vehicles for Science, Operations and Outreach

- Migrated image loader to new Trinity3 platform. Some minor continued exploration of a web dashboard. (Week 2)
- Team meeting (Week 2)
- Team meeting (Week 4)
- Working on issue 23 to replace current YOLOv11 detector with better performing RF-DETR. Processed latest Trinity Pro mission in Seacliff on 20260114 with new RF-DETR model. (Week 4)
- Review with collaborator, Tim of mission mosaic processing with Kelp coverage; some issues with edge artifacts confounding SAM algorithm. Setup Kelp segmentation test project on drone.mbari.org anticipating test segmentation map loads. (Week 4)
- Working on <https://github.com/mbari-org/UAV-issues/issues/23>. Had some issues with the vGPU reporting no available memory that others have reported. A reboot fixed this issue. (Week 5)

- Shared image mosaic of a mission with collaborator T. Sananikone who is working on a SAM-based kelp coverage tool. He is evaluating the performance on a mission mosaic and looking towards the SAM3 model. (Week 5)
- Loaded new RF-DETR model results to drone.mbari.org. Appears to be a much more sensitive model than YOLOv11; some adjustments needed for threshold. (Week 6)
- Team meeting and OSM poster discussion on kelp coverage ML (Week 6)
- OSM poster review (Week 7)
- Team meeting (Week 9)
- VO assistance. Some issues with the camera so we had to cancel the flight. (Week 10)
- Summer intern review and interviews (Week 14)
- Tested revised threshold on ML pipeline which reduced false detections ~30%. Issues with running RF-DETR now need attention. perhaps misalignment with repo version. (Week 14)
- Team meeting (Week 14)
- Some work on a React web application to visualize mission results (Week 17)
- Loaded last three missions and running processing pipeline (Week 17)
- Deployed Voxel51 sync service to drone.mbari annotation server to support faster review of results. Many thanks to Adriana and Peter in IS for their assistance with this. (Week 17)
- Added two new classes to the classification model and retrained. Improved performance 5%. Models exported to ONNX format. (Week 20)
- Deployed latest classification model to MBARI modelhub (an easy to use model card for internal use only) (Week 21)
- Added FastAPI health to cortex (Week 21)
- Team meeting update (Week 21)
- Discussed SAM3 few-shot with collaborator (Week 21)
- VO on mission, continued work on mission viewer and updated docs (Week 23)
- Onboarding Sarah with AI/UAV documentation and tools. Setup account on drone for downloading annotation data, setup Slack channel, weekly meetings, and added to github repos for documentation. She is off to a great start! (Week 24)
- Visual observer for NewBrighton deployment (Week 24)
- Proposal input (Week 25)
- Deployment of stratagrid to visualize training data (Week 25)
- A few mbari-aidata pip modules updates to support train/test (Week 25)
- Oceans abstract on Kelp SAM Coverage accepted (Week 27)
- Summer intern help and VO (Week 27)
- Helping intern Sara with aidata download and workflow design (Week 28)
- Updates to mbari-aidata module to support summer intern, Sarah Williams, who is making good progress. (Week 29)

902111 - Carbon Flux Ecology

- Working with L. Chrobrak on Voxel51 enterprise evaluation. (Week 2)
- Loaded planktivore high mag into Voxel51 Enterprise and entire ISIIS video deployment into Voxel51 Community for Colleen to validate. (Week 2)
- Ongoing work with Voxel51. Working on associating depth missing from ROIs which needs to be loaded. (Week 4)
- Updated planktivore documentation on first iteration of labeling. (Week 4)
- VSS (Vector Search System) updates to add augmentations supporting Fernanda's work exploring ISIIS classification model workflows. (Week 5)
- OOD (Out of Distribution) exploration with Fernanda. (Week 6)
- Voxel51 data preparations for Planktivore use (Week 6)
- OSM presentation support (Week 7)
- Clustered OOD (Week 7)
- Planktivore support for Voxel51 (Week 7)
- Database cleaning to remove video captured during flooding events and some performance improvements to vector database service VSS to handle batching with less memory (Week 7)
- CenCOOS/Planktivore meeting. Issues with authentication support to be bumped to Voxel51. Will revisit after OSM when team is back (Week 8)
- Fernanda will be presenting our work at OSM, Thursday, Feb 26th 3 PM GMT/ 7AM PST ME43A-07: Simultaneous, automated classification of abundant marine snow particles and rare zooplankton from imagery (Week 8)
- Deployment of new VM cortex with proxy gateways for various services, including a new Voxel51 sync service (Week 9)
- SINKER deployment support; adding new physical GPU-enabled machine snowfall to workflow, switched MARS IP address, some optimization and cleaning of the ansible deployment scripts. (Week 10)
- Debugging issues with embeddings failures in vector database service VSS. New VM cortex up and running; more planned for this but working well as proxy gateway for now. (Week 10)
- SINKER deployment support, adding new physical machine snowfall to rsync and logged into instrument to confirm ADCP not working. Paul was able to fix remotely. (Week 11)
- Planktivore and ISIIS data is now syncing from Tator to Voxel51 both community and enterprise versions. Over 1 million ROIs successfully captured from video or from raw source, embeddings and similarity index computed. Laura Chrobak working with Voxel to refine plugins for our use-cases. (Week 11)
- Populated VSS (Vector Search System) with noise-only embeddings for ISIIS, computed scores, loaded into mantis, and exported through Voxel51 community sync service. Fernanda is comparing against classifier scores. (Week 14)
- Helping team with SINKER processing pipeline. (Week 17)
- Registered a few Triton videos on mantis. (Week 20)

- Developing processing pipeline to separate particles from midwater taxa. This will leverage models created for MTS and Biodiversity. Models exported to ONNX format. (Week 20)
- Working on particle detection and classification pipeline for Triton (Week 21)
- Team working on evaluating model performance of classification model on ISIS mission (Week 21)
- Fixed issue 16 (Week 21)
- Discussed Planktivore High Magnification below-surface data mining with Fred Bahr. All missions, except the CFE lab planktivore mission (awaiting PRobert's time-correction on this) has been preprocessed to parquet. (Week 21)
- Finished data mining 200K ROIs Planktivore low-mag 2024-2025 archive for model training. Stratified across depth/month/hour to represent diversity and clustered using SDCAT using high-mag model. Data loaded into mantis and Voxel51 for advanced annotation by lab members. Meeting with lab members next week to discuss data management. (Week 23)
- Fixed some issues with cropping failures (Week 23)
- SINKER storage discussion with IS (Week 23)
- Migrating syncing to a new service account (Week 24)
- 200K plankton images stratified samples across depth/time were clustered and loaded for the team to annotate (Week 24)
- Proposal input for lab and SINKER (Week 25)
- Deployment of stratagrid to visualize training data (Week 25)
- Working on active learning tutorials to teach lab how to build classifiers (Week 25)
- Successfully migrated SINKER log files using rsync with new service account, closing issue #3 (<https://github.com/mbari-org/deploy-sinker/issues/3>). (Week 27)
- Advanced Jupyter notebook tutorials for the group, including adding download functionality across media and localizations and mount to CFE lab share. (Week 27)
- Redeployed mantis with new automated SSL certificate renewal. (Week 27)
- SINKER rsync migration completed; issues with file locking and backlog seem to be resolved (Week 28)
- Advanced Jupyter notebook tutorials with download functionality; Fiftyone sync updates to add better dataset creation per users (Week 28)
- Running VSS on planktivore low-mag to search for rare taxa (Week 28)
- Issues with zmq-aggregator failures need further work. Thanks to Karen for her help with this. (Week 28)
- Data mining for rare taxa in Planktivore low-mag using vector database VSS populated with Fernanda's labels. (Week 29)
- Maintenance on deployment scripts to support both ansible and manual deployments (Week 29)

Out of Office

- Vacation: 3-12,3-13
- April 1-2, Vacation

- 21-24 September: Oceans Conference
- September 21-24 Oceans Conference, Monterey
- May 27-28, AI DevSummit, San Francisco
- May 27-28, AI DevSummit, San Francisco
- Out of office May 25-29th.
- Vacation July 8th-13th
- October 18th, SMM 2026 pre-conference Workshop: Evolving AI Applications in Modern Marine Mammal Science: From Computer Vision to Generic and Agentic AI, Puerto Rico
- Vacation July 27th-31th

Overhead

- Performance review input (Week 4)
- Annual review materials and met with supervisor to review (Week 5)
- Performance review (Week 6)
- Group meeting to discuss authentication template and AI (Week 7)
- Working on a modelhub that captures current models, metrics, and supports interactive demos. (Week 14)
- Proposals and OCEANS abstract (Week 17)
- Attended some abstract presentations (Week 20)
- Summer intern planning (Week 21)
- Travel request (Week 21)
- IS request for S3 bucket to migrate mantis media (Week 21)
- May 27-28, AI DevSummit, San Francisco (Week 23)
- Reviewing Ecosphere paper (Week 23)
- Planktivore proposal discussion (Week 23)
- AI DevSummit conference, San Francisco. Some follow-up with vendors. (Week 24)
- Finished reviewing Ecosphere paper (Week 25)
- Expansion of object storage needed for on-prem annotation server (10 TB), IS ticket pending. Tested Qumulo object as Tator object store; mostly S3 compatible, except ACL and some CORS issues. Minor code revisions needed. (Week 27)