

DOSI New Technologies WG Marine Imaging Workshop Report

December 2015

The Deep Ocean Stewardship Initiative (DOSI) working group on Technology recently held a Workshop on Marine Imaging, with the main discussions concerning Image Acquisition, Real-Time and Post-Cruise Image Annotation, Data Sharing and Archiving. Experts from the Australian Centre for Field Robotics, the Monterey Bay Aquarium Research Institute (MBARI), the University of Tokyo, the Japan Agency for Marine-Earth Science and Technology (JAMSTEC), the Schmidt Ocean Institute (SOI) and commercial enterprise (e.g. axle Video, Greybits Engineering) contributed to two days of in-depth discussion of the pros and cons of various systems, the identification of best practices, and the specifications for a cross-platform software solution for managing imagery and video data. Over the coming weeks these discussions will be translated into a list of functional requirements that will be assessed against real-world User scenarios. Please contact working group head Dhugal Lindsay (dhugal@jamstec.go.jp) if you have a user scenario that you feel we may not already have covered. A list of the user scenarios that were considered appear below. We plan to present outcomes of this workshop and a preliminary list of requirements at the Marine Video Workshop scheduled for March in Rhode Island.

User stories that outline potential use-cases for the proposed tool.

Live video annotation

A researcher is at sea and wants to be able to collaboratively annotate video data on the fly being collected by his ROV. The data is being logged locally on the ship, being displayed on multiple screens/tablets, with annotation input by 4-5 members of the science team on the ship. Remote scientists are also logging in and providing annotations.

Data Registration / Upload

A scientist has just returned from a trip to Scott Reef. BRUVs, Towed Video and AUV surveys were conducted during the campaign. She wants to add the new data to the system and make it available to other users by registering the video data with the system. Another researcher, on the other hand, wants to upload his video data to YouTube.

Accessing Data

A scientist wants to access and download all available towed video and AUV images, with annotations, from Scott Reef for research purposes, and anticipates that correctly identifying the current owners / holders of that data will be time consuming. In addition, from his experience, working out user agreements will also be challenging.

Query the data

One scientist wants to identify available datasets by date, platform, instrument, geographic location (with some notion of accuracy), depth, altitude(?), institution, PI, etc. Another wants to find all images or video sequences above a certain resolution taken at Scott Reef of anemones in temperatures above 25°C in depths between 15 and 45m and with extended tentacles.

Subsetting the Imagery for Download or Annotation

A scientist wants to work with a number of AUV dense grids off of the NSW coast. She would like to download 100 images from each of these dives drawn randomly from a uniform distribution over the spatial extent of these grids or subsetting by depth strata. Another scientist wants to do this for one minute long video sequences from a seafloor observatory.

Viewing Individual Images

A scientist has selected a Collection of imagery and is interested in looking at individual images from within the set he has selected. He would like to use a map interface to view the location of the dives and to select a particular image which he can look at in full resolution. He would also like to have the ability to tweak the image to enhance contrast or colour.

Viewing video sequences

A researcher has selected a video sequence collected by an ROV and is interested in looking at frames from within the sequence. She would like to use a map interface to view the location within the dive and to select a particular set of frames which she can view. She would like to view the particular frames, to be able to jog the video around a particular species of interest, and to export short subclips at the original resolution to send electronically to taxonomic specialists for identification. She would also like to see associated temperature and depth data and to have the ability to tweak the video to enhance contrast or colour.

Video plankton data

A scientist would like to annotate photographs and scans of bulk plankton samples for samples that were collected at an offshore location. He would like to be able to upload subsampled elements of the large images to be annotated. He has an approximate location where the water samples were collected.

GIS map-based data

A researcher has a broad scale photo-realistic mosaic and would like to view it and annotate it through the system.

Point Annotation of Subsetted Images

Rather than downloading the data, a scientist wants to be able to annotate a random set of points in these images using tools, like CPCe (<https://eur03.safelinks.protection.outlook.com/?url=http%3A%2F%2Fwww.nova.edu%2Foceanocean%2Fcpce%2F&data=01%7C01%7Chls1e18%40southamptonalumni.ac.uk%7C78a137dfc7084164eef908d8182b1785%7C4a5378f929f44d3ebe89669d03ada9d8%7C0&sd=ata=WmYNQKddI5Hmxd1ipn%2BXa%2FxFQ1HBDz2GMXsrhjlfPI%3D&reserved=0>), online. She would also like to be able to identify specific organisms of interest through a point and click interface and to distinguish randomly assigned points from those which were selected by a user. She would like to share this Workset and the resulting annotations with some of her students so they can actually do the labelling.

Region Annotation of Images

A scientist wants to be able to annotate regions in a set of images or video frames. He would like to be able to identify a set of pixels within the image that correspond to a particular organism.

Video Annotation

A researcher wants to be able to annotate frames from within a video sequence. She wants to be able to assign points on the screen or to select regions from within the frame to identify as a particular organism. She would also like these annotations to be translated to adjacent frames and to identify the period over which the organism is in view. She would also want to be able to add comments related to the audio stream which was recorded while the data was being collected.

Stereo Annotation

A scientist has borrowed a stereo camera for his ROV. He would like users to be able to annotate the stereo images together and to take measurements of organisms using an available stereo calibration to resolve distances within the images. The registration of points from one frame to another should be semi automated but tweaked in such a way that the user can refine the registration of the points.

Annotation Schemes

A scientist wants to be able to define an annotation scheme specific to his study. He would like to be able to leverage existing labels but define a specific set of labels that he will use for his study. He would like the scheme to capture some of the hierarchical nature of the taxonomy. He may also need to be able to change the labels in light of new developments, changing understanding of the organisms he is seeing or in response to feedback from expert taxonomists. The system should also capture synonyms describing the same organism using scientific or more common names.

Multiple labels

In certain cases a specialist may be able to identify an organism to be one of a number of alternative species, which may not be closely related in terms of their taxonomy. He would like to capture the uncertainty in the labelling by assigning multiple labels to a point with some notion of the uncertainty of the label assignments. He suggests alpha OR beta-siblings NOT delta-siblings as one way to label them.

Modifiers

Annotations should be assigned based on a selected annotation scheme. The labelling scheme should also feature the ability to include modifiers (eating, bleached, etc.) as well as free-form comments. A student wants to tag an anemone eating a Discidicus. He also wants to add a modifier that says that there was an audio-comment and include free form text about the comment.

'Depth'-Dependent Species Counts

A student wants to count the number of anemones in an AUV campaign in the Great Barrier Reef. A number of AUV dives are involved in this study from the year 2007. He knows that these anemones typically live in water shallower than 80m, and they only live on reef and

not on sand. He would like a way of only being presented with AUV images taken over reef that include anemone annotations, between 10m and 80m depth.

Habitat Identification from Towed Video Data

A scientist has 20 towed video transects she wants to share with others. She would like to be able to assign labels to the entire frame concerning video quality, substrate type or the presence of organisms of interest, either from live video sequences or from captured data. She is also interested in classifying still images from these transects into coarse habitat and substrate types. Unfortunately many of the images are poorly illuminated, are quite turbid in appearance or were collected when the towed video was moving too fast to have good quality data. She would like to be able to filter these poor quality video sequences out before annotating the rest. She might also want to be able to identify a region of interest within the video stream.

Cluster Deployment Imagery

A researcher is particularly interested in focusing on images that feature kelp from within the Collection of the data he has selected. He wishes to have AUV deployment imagery from his selected dives automatically clustered into visually similar groups and to use these groups to further refine the subset of images he will analyse in more detail.

Anomaly Detection

A scientist has a long sequence of video collected by a deepsea mooring. The video is mostly of a sandy, unchanging environment. However there are short sequences of video that feature animals previously not seen at the depths he is surveying. He would like a way to automatically identify these unusual events within the video stream.

View Label Output

A researcher wishes to view data associated with a deployment and with its associated manual labels, auto-clustering or classification output.

User interface requirements

Users may need to add information in a variety of languages and using local descriptions of the organisms they are working on.

User Registration

A Principal Investigator wants to know who annotated a set of images so that she can evaluate the quality of annotations. She also wants to register her datasets as her own so that she can administer the access rights to her data and be acknowledged in publications by others using the data. She would also like to register her areas of expertise so others can evaluate how confident she is of her labels and perhaps seek out his opinion on particular organisms they are unsure about.

Quality Control

A researcher does not trust the quality of labels entered by his unpaid interns. He would like to be able to view the annotations assigned by his team and to flag errors in the annotations. He would then like to be able to view statistics on the quality of the

annotators. He would also like to be able to register quality control information associated with the annotations.

Export Imagery and Annotation Output

A Chief Scientist would like to export data collected by JAMSTEC during their recent cruise. He would like to be able to access imagery and video sequences as well as associated meta-data and labels that have been entered by the team. He would like to be able to see draft figures summarising the labels, including depth and location distributions of organisms and to export the raw data to csv with associated image and video sequences.

Import existing labels and video/images

MBARI has warehouses full of existing video sequences and associated labelled data. They would like to be able to share this data with the world and to register it in the system.

Synchronising field data with global repository

The JAMSTEC team have returned from their latest campaign. They have produced a large number of annotations in the field and wish to synchronise these with their existing repository. This could also happen over a satellite link while still in the field.

Transcribing labelling data between media streams

Given registered data between streams (synchronised video streams for example), a scientist would like labels in one stream to be suggested in related video streams collected together based on the time within the video stream. Locations within the frame could be provided by the user or suggested using registration techniques.

Classification of Imagery using Annotations (Supervised Classification)

A researcher has been annotating a number of images from her selected Collection. She only has time and resources available to label a small fraction of the data in her Collection as a Workset. In order to increase the efficiency of the annotation process and to extend the number of images that are labelled from the Collection she wishes to use automated classification tools to pre-suggest labels and ultimately to label full images.

Active Learning to Select Images to Annotate

A scientist is preparing to annotate a number of images from her selected Collection. She is interested in using Active Learning to allow the system to suggest images which are problematic to classify automatically. As further images are annotated, the automated classification system will incorporate these labels and improve automated classification performance for the Collection.

Validation of Automated Classification Output

A scientist has used the annotations from her Workset to train a classifier over her collection of images. She would like to validate the performance of the classifier by examining automatically derived measures of classification performance as well as examining images that the classifier has labelled based on the learned classifier.