

Detection of Anomalous Particles from Deepwater Horizon Oil Spill Using SIPPER3 Underwater Imaging Platform

Sergiy Fefilatye¹, Kurt Kramer¹, Lawrence Hall¹, Dmitry Goldgof¹, Rangachar Kasturi¹, Andrew Remsen², Kendra Daly²

¹Department of Computer Science and Engineering, University of South Florida,
4202 E. Fowler Ave., ENB 118, Tampa, FL 33617, USA Tel.: +1-813-974-3652

²College of Marine Science, University of South Florida,
140 7th Ave. S, St. Petersburg, FL 33701, USA Tel.: +1-727-553-1130

sfefilatye@gmail.com, kurtkramer@gmail.com, hall@cse.usf.edu, goldgof@cse.usf.edu,
r1k@cse.usf.edu, aremsen@marine.usf.edu, kdaly@marine.usf.edu

ABSTRACT

The aim of this study is to investigate a data mining approach to help assess consequences of oil spills in the maritime environment. The approach under investigation is based on the visual detection of suspected oil droplets in the water column adjacent to the Deepwater Horizon oil spill. Our method detects particles in the water, classifies them and provides an interface for the visual display and detailed examination. The particles can be plankton, marine snow, oil droplets and more. The focus of this approach is to generalize the methodology utilized for plankton classification using SIPPER (Shadow Imaging Particle Profiler and Evaluation Recorder). The SIPPER, which has been in use by marine scientists for the last decade, allows the timely extraction and identification of millions of images per deployment as scanned by its underwater sensor. It can be deployed at various depths. In this paper, we report on the application of image processing and machine learning techniques to discern suspected oil droplets from plankton and other particles present in the water. We train the classifier on the data obtained during one of the first research cruises to the site of the Deepwater Horizon oil spill. Suspected oil droplets were visually identified in SIPPER images by an expert. The classification accuracy of the suspected oil droplets is reported and analyzed. Our approach reliably finds oil when it is present. It also classifies some particles (air bubbles and some marine snow), up to 2.8%, as oil in clear water. You can reliably find oil by visually looking at the examples put in the oil class ordered by probability, in which case oil will be found in the first 10% of images examined.

General Terms

Algorithms, Measurement, Documentation, Performance, Design, Verification.

Keywords

Oil-droplet detection, images, classification, oil-spill, visual inflow, analysis, machine learning, support vector machine, plankton.

1. INTRODUCTION

The Deepwater Horizon Oil Spill is the biggest environmental disaster in the United States history and is the largest marine oil spill in the history of the oil industry [1]. The

impact of the spill is still being evaluated with various estimates of the immediate damage, area affected, and longevity of its effect being generated [2]. Most of the studies are focused on the oil that has covered the surface of the ocean, addressing the fact that the oil-mixture is lighter than water and, thus, tends to float on its surface. However, there are many indications that after being affected by the chemicals aimed to disperse the oil compounds much of the oil has turned into neutrally buoyant oil droplets (see Figure 1) and has permeated the depths of the Gulf of Mexico [3]. The properties of such oil droplets have the potential to allow the particles to remain in the water for long periods of time, negatively affecting the marine habitat, fishing, and tourism industry.

In this study, we evaluated a special platform, intended for plankton research, for the use of oil droplet detection in columns of sea water. Based on a proven record for plankton population classification, we undertook a study to assess how suitable this platform is for the current efforts to detect oil droplets suspected to be in the water. Researchers from the University of South Florida's (USF) College of Marine Science collected image data during one of the first research cruises to the area affected by the Deepwater Horizon oil spill. The data represents images of particles of plankton along with suspected oil droplets. It is stressed here, that the oil droplets are only "suspected" as our data included only imagery, not physical samples of the water. However, based on the extensive experience of marine scientists involved in the manual examinations of the data it is believed that it is highly likely that the images represented oil droplets. The aim of this research is to evaluate the effectiveness of the use of image processing and machine learning techniques to process a large quantity of data obtained from our underwater research instrument to classify particles assuming the image data collected during the initial deployment indeed includes oil droplets. We are not aiming to draw any conclusions on the ecological meaning of the SIPPER data and presence of actual oil. However, this research may result in a vision-based method to assess the presence of such oil droplets in the water columns using the SIPPER tool. This paper briefly describes the hardware of the instrument used, the algorithmic process intended to discern suspected oil droplets from other plankton particles, and the results obtained on a dataset collected in the immediate vicinity of the oil spill, as well as on



Figure 1. Suspected oil droplets found in the area of Deepwater Horizon oil spill

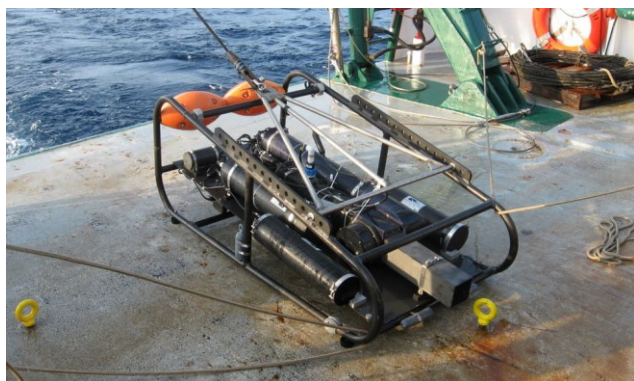


Figure 2. SIPPER is being deployed for plankton research

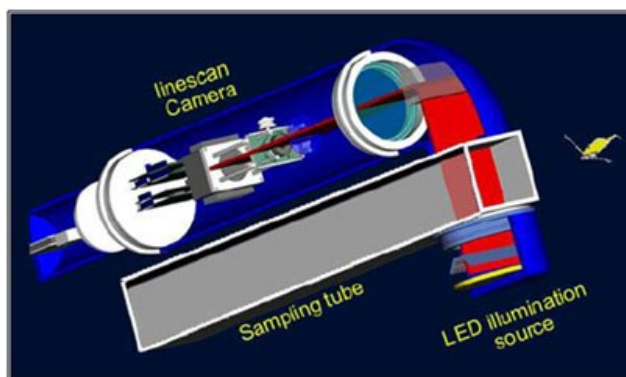


Figure 3. Interior view and optical layout of SIPPER

data from unaffected areas. We discuss our observations, limitations of the approach, and suggestions for further research.

2. DESCRIPTION OF PLATFORM

The instrument for plankton research, SIPPER [4], was developed by the Center of Ocean Technology at the University of South Florida for the purpose of monitoring the composition, distribution and size structure of plankton and other suspended particles in aquatic environments (see Figure 2). The SIPPER uses collimated LED illumination and a high speed line scan camera to continuously image particles and plankton as they pass through a $10\text{cm} \times 10\text{cm}$ sampling aperture [5, 6, 7]. A continuously scanning line scan camera captures images that are 10 cm in width and continuous in length (see Figure 3). All resolvable particles that enter the sampling tube are imaged and saved as a single large SIPPER file with concurrently collected environmental data such as temperature and depth embedded within it. A single 6 hour deployment can result in hundreds of thousands to millions of individual extracted particle images larger than 0.4 mm equivalent spherical diameter (ESD).

Custom designed software, the Plankton Imaging Classification Extraction System (PICES), was developed to quickly extract, classify, manage and analyze these discrete plankton images. A database management subsystem within PICES allows management of the large amount of data generated by SIPPER. PICES provides quick retrieval and organization of data by multiple parameters, such as, cruise, deployment, depth, salinity, temperature, classes, date-time, etc. Use of PICES results in efficient and timely processing of collected data.

The main modules of PICES incorporate image extraction, classification, active learning, feature selection, and parameter tuning. The PICES image extraction function is uniquely designed to process the continuously scanned imagery data generated by SIPPER, extracting individual plankton images and associated

embedded environmental data. Feature vectors are computed for each image, which then gets automatically classified into user defined classes by a support vector machine (SVM) [8] built using training libraries maintained by the user. The classified images with their feature vectors and environmental data are then inserted into a database. The SVM learns from the expert-labeled images to recognize the class to which the unlabeled images will be assigned. Having trained such an SVM-based classifier it is possible to classify millions of images and determine the composition of the population of the plankton in the area where the data was collected.

3. DESCRIPTION OF ALGORITHMS

The algorithms used during processing of the data include those for image extraction, feature calculation, and image classification.

PICES uses a simple algorithm to extract images of separate particles based on foreground-background segmentation and a connected components algorithm. After segmenting the image of a particle, a number of features are calculated/extracted and a feature vector is created. The features are used by a classification algorithm in order to assign a class label to the image.

The PICES SVM classifier provides a confidence or probability value [9] for its selection and, thus, gives more flexibility in the process of final decision. As a supervised classifier, the SVM requires training data in order to learn to correctly label a particular particle-class from its appearance. The training dataset is created from data labeled by an expert. One of more marine scientists views the images from one or more cruises and/or deployments and labels (some of) them. One of the inherent disadvantages of an SVM is that it can only handle two-class problems. However, it can be extended to a multiple-class problem by using several strategies, for example one-versus-all or one-versus-one. In our algorithm we used a one-versus-one

Table 1. Image features used to classify particles present in the water

Category	Sub-Category	Feature Count
Moment Features [10]	Binary	8
	Intensity weighted	8
	Edge pixels only	8
Morphological		9
Head/Tail	Pixel counts of first quarter and last quarter	2
	Length vs. width	1
	Length	1
	Width	1
Filled Area		1
Convex Area		1
Transparency	Binary/Weighted	2
Texture Using Fourier Transform [11]	According to each frequency range	5
Contour Fourier	Average of Five Frequency Domains	5
	Hybrid combinations	15
Intensity Histogram	Without white space	7
	With white space	8
Circularity	Circularity, Equivalent Diameter, Eccentricity, ratios, etc	5
Texture	Intensity statistics, Smoothness, Uniformity, Entropy	6

strategy for every possible two-class combination by selecting the winning class using voting. Features for each pair of combinations were selected separately using a Binary Feature Selection (BFS) process described in [17].

4. DESCRIPTION OF FEATURES

During its operation SIPPER records environmental data, such as water temperature, depth at which SIPPER was operating, oxygen concentration, salinity, and florescence. However, it was decided to use only image-related features as environmental data did not provide enough variety in our dataset. Some of the reasons for that are the following:

- The depth sensor is always used. However, the oil dataset was collected only for certain depth columns. Hence, this may spuriously be used to indicate the presence of oil.
- The salinity, oxygen, and florescence sensors were covered during the collection of the oil dataset because of the fear of being contaminated by oil. Thus, in the dataset collected in the vicinity of the oil spill the values for these features do not exist.
- The temperature feature is not reliable because of the fact that the oil was only detected during one of the trips spanning a very limited time in relation to all possible range of temperatures in that region during the year. Thus, temperature might be chosen by the classifier to represent oil spuriously. Oil only occurred at a couple of temperatures, but could occur at any. Generally, this feature is useful because some plankton creatures may prefer a certain temperature during a certain season.

Table 1 shows the image-related features that were used during the feature selection process. In total 93 features were used for describing the data. Those included 82 features which were previously designed for identifying the general plankton population. These features were mostly concerned with direct

measurements (pixel count, intensity), geometric and boundary properties. Another 11 features were specifically designed to aid in detection of oil droplets. They were mostly concerned with the circularity of the shape of oil droplets and their texture properties.

5. SVM PARAMETER TUNING AND FEATURE SELECTION

In this work we used a one-versus-one strategy in order to implement a multi-class classifier. The primary reason for such a decision was the faster speed of training which is shown by some studies [12]. In this strategy, all SVM classifiers for all possible binary combinations of all classes are created. A class label is selected by a majority vote. In the case of a tie among classes, the probability parameter of SVM was used to select the class label.

The feature selection process consisted of two steps: initial SVM parameter tuning and binary feature selection. The parameters (γ , C , A) of the SVM are optimized by performing a grid-search with a certain interval across the training dataset [13]. Using the SVM parameters determined in the first stage of the selection process, a binary class feature selection is performed using a wrappers approach [15, 16, 17]. Each specific combination of features and SVM parameters was evaluated using 5-fold cross validation [13] and the classification accuracy on the training set was used to guide the selection process further. In cases, where the classification accuracy is equal for several evaluated sets, the correctness of probability (CPP) [9], the inherent SVM parameter, is used to rank the sets.

6. DESCRIPTION OF DATA

The data is image data collected by the SIPPER during one of the first research cruises to the area of the Deepwater Horizon (DWH) oil spill on May 5-16 2010 on the USF research vessel Weatherbird II. Data from three deployments was collected within

Table 2. Datasets used in experiments

Dataset Identifier	Particles total	Oil droplets
Oil Original Set	8536	1072
Oil Original Replaced Set	8536	1072
Oil Large Test Set	43816	13858
Non-Oil Set	6745	0
Oil Random Set	13678	79

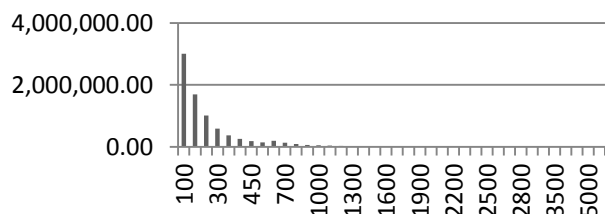


Figure 4. Size distribution of particles in the flow of water according to their size in pixels during the research cruise to the area of DWH oil spill.

5km of the original site of the DWH platform on May 14 and May 15. However, the ship was not allowed to get too close to the source of the spill per the Coast Guard’s interpretation of safety in the region. Anomalous semi-opaque spherical particles were manually detected in the SIPPER imagery in the upper 10 meters from these three deployments. These particles were imaged in areas where oil was visibly observed at the sea surface during a time of relatively strong winds and building seas. These conditions provide a possible mechanism by which surface oil could be mixed down into the water column. Based on that and because these particles were not observed in imagery collected in nearby waters where surface oil was not present, we labeled these particles as suspected oil droplets. They did not resemble other spherical particles that have been imaged by SIPPER such as fish eggs, sarcodine protists or air bubbles. Other cruises around this time resulted in the collection of data from the average plankton population in the Gulf not affected by the oil spill. The data from unaffected areas is used to assess the sensitivity of the approach to the presence of oil droplets and compare the distributions of particles between the areas. Results of such comparisons may be used for future studies of the ecological impact of the spill.

Evaluation of the observed image data suggests that the water column contained mostly small particles. Many smaller particles are found for each large particle encountered. Figure 4 shows the size distribution of particles found in the SIPPER images from the research cruise to the area affected by the oil spill. Size is the area in pixels of each particle. For this study, the dimensions of each pixel are approximately 27 μm on each side.

Particles that exceeded 100 pixels in total area were extracted. As seen in Figure 4, there were abundant small particles present, while larger particles were far less numerous. However, due to the lack of resolvable features for the smallest extracted particle images, only particles greater than 250 pixels in total area were classified by an expert. It was decided to disregard all images smaller than 250 pixels to increase the accuracy of particle classification. Images of sizes > 250 pixels, according to our observations, contained enough texture and contour information to effectively differentiate among other classes.

Table 3. Categories of classes of particles used in experiments

Category of classes of particles	# of classes in category
Crustacean Copepod	5
Crustacean Eumalacostracan	3
Detritus (including oil droplet)	5
Echinoderm	1
Elongate	2
Fish	1
Gelatinous	8
Mollusc	2
Noise	2
Phytoplankton	1
Protist	5
Radiolarian	1
TOTAL	36

We created five datasets (see Table 2) to study the data from the area affected by the oil spill, as well as unaffected areas. For all datasets the number of classes was set to 36, which represented only major classes of particles with at least 20 instances each. Table 3 shows the categories of classes used in our study.

The set of images, called *Oil-Set Original*, was obtained by selecting instances of particles that were of most interest to a marine scientist from the data obtained in the area affected by the spill. There were a number of selection criteria. First, the particles had to be identifiable in the sense that they had a high probability of being a particular plankton class or oil droplet. Second, since we were primarily interested in oil detection, oil droplets were a focus of the initial search and more likely to be labeled. This labeling was done first after the cruise. Thus, this dataset does not represent a completely random choice of particles. Overall, the set is composed of 8537 particles which represented less than 0.5% of all data during the cruise to the affected area. The oil droplet class had 1072 instances, comprising 12.49% of all the particles in the dataset. The decision to label each particle was made based on a visual analysis of the particle with the knowledge available to the expert.

The next two datasets were obtained from the data from the same cruise in the following manner. First, *Oil Original Set* was used to train an SVM classifier within PICES. The resulting classifier was used to classify all the data from the cruise. About 50,000 images that had a high confidence for the predicted class were viewed by an expert and given final class labels (which could be the same as the predicted class label). Some of the instances in this classified and validated data were part of the *Oil Original Set*, because it came from the same pool of raw data. Since our interest was mainly in oil detection, priority was given to the validation of oil droplet predictions. Out of 50,000 instances of classified and validated data about 20,000 of instances were oil droplets.

Oil Original Replaced Set is a dataset which was obtained from *Oil Original Set* by replacing the 1072 oil droplets with oil droplets randomly selected from the set of classified and validated data as described above. The *Oil Original Replaced Set* has the

Table 4. Performance of single-stage classifier. 10-fold cross validation on Oil Original Set

Oil droplet detection accuracy: 90.95%			
Absolute Performance:			
	Oil droplet	Other	Count:
Oil droplet:	975	97	1072
Other	104	7360	7464
Total:	1079	7457	8536
Relative Performance:			
	Oil droplet	Other	
Oil droplet:	90.95%	9.05%	
Other	1.39%	98.61%	

Table 6. Performance of classifier. Tested on Oil Large Test Set, Trained on Oil Original Set.

Oil droplet detection accuracy: 92.67%			
Absolute performance:			
	Oil droplet	Other	Count
Oil droplet	12842	1016	13858
other	1324	28634	29958
Total	14166	29650	43816
Relative Performance:			
	Oil droplet	Other	
Oil droplet	92.67%	7.33%	
Other	4.42%	95.58%	

same number of instances as *Oil Original Set*, but with more a diverse oil population because of the random choice of oil droplet samples.

Oil Large Test Set was obtained by including all predicted and validated data that passed through the following filters. Instances of the data that are a part of *Oil Original Set* were removed. 5000 images of oil droplets, selected randomly, were removed for future use for validation. Another 1072 oil droplets used for building *Oil Original Replaced Set* were removed as well. In summary, *Oil Large Test Set* had 36 classes, 43816 images total, of which 13858 were oil droplets.

Oil Random Set was created in a different manner. Instances of all particles in the set were randomly selected from the data from the cruise to the areas affected by the spill, not just by selecting particles of interest. The particles in the set were assigned the appropriate class based on visual analysis. Thus, this dataset has approximately the same distribution as the real distribution of particles in the water during that deployment.

Datasets *Oil Set Original*, *Oil Large Test Set* and *Oil Random Set* did not intersect. Datasets *Oil Set Original* and *Oil Set Original Replaced* intersected for instances of all classes except oil droplets.

The last dataset, *Non-Oil Set*, contains data that was collected during other cruises to areas unaffected by the BP oil spill. This data was collected from several locations in the Gulf of Mexico as well as the Caribbean Sea in 2010. The experiments conducted on this data were designed to test the classifiers sensitivity to oil droplets, i.e. detection of oil when no oil is present. The dataset

Table 5. Performance of single-stage classifier. 10-fold cross validation on Oil Original Replaced Set

Oil droplet detection accuracy: 95.80%			
Absolute performance:			
	Oil droplet	Other	Count:
Oil droplet	1027	45	1072
Other	86	7378	7464
Total:	1113	7423	8536
Relative Performance			
	Oil droplet	Other	
Oil droplet	95.80%	4.20%	
Other	1.15%	98.85%	

Table 7. Performance of classifier. Tested on Oil Large Test Set, trained on Oil Original Replaced Set.

Oil droplet detection accuracy: 94.20%			
Absolute performance:			
	Oil droplet	Other	Count
Oil droplet	13054	804	13858
other	820	29138	29958
Total	13874	29942	43816
Relative Performance:			
	Oil droplet	Other	
Oil droplet	94.20%	5.80%	
Other	2.74%	97.26%	

had 6745 particles belonging to 36 classes (the class oil droplet had 0 instances, i.e. did not occur).

7. EXPERIMENTS

In our experiments we report the accuracy of classification in the form of a 2x2 confusion matrix, as if we were doing binary classification, although the setup of experiment itself was not binary. One class was the oil droplet class – particles of particular interest for this research. The category ‘other’ represents the classification of all other particles compared against oil droplet. Thus, every prediction in favor of one of the other 36 classes of the datasets is summarized into the ‘other’ category. We do not report the accuracy among the 36 non-oil classes.

A Binary Feature Selection process was performed to select features for each of the 630 binary SVM classifiers that comprised our one-stage classifier for 36 classes. Table 4 shows the performance of the classifier using 10-fold cross validation on the *Oil Original Set*. 90% the oil was identified with a less than 2% false positive rate. Table 5 shows the results of a 10-fold cross-validation on the *Oil Original Replaced Set*. The accuracy and false positive rate in that experiment was improved, correctly identifying 95% of the oil with a 1.15% false positive rate. In all other experiments we report the performance of two classifiers, first trained on the *Oil Original Set* (called *Classifier I*) and second trained on *Oil Original Replaced Set* (called *Classifier II*) to compare their sensitivity and specificity.

The two classifiers were created and then applied to make classifications on the *Oil Large Test Set*, *Non Oil Set*, and *Oil*

Random Set. The results for these experiments are shown in Tables 6-11.

Table 8. Performance of classifier. Testing on Non Oil Set when trained on Oil Original Set.

False positive rate: 7.62%			
Absolute performance:			
	Oil droplet	Other	Count:
Oil droplet	0	0	0
Other	514	6231	6745
Total:	514	6231	6745
Relative Performance :			
	Oil droplet	Other	
Oil droplet	0.00%	0.00%	
Other	7.62%	92.38%	

Table 10. Performance of classifier. Tested on Oil Random Set, trained on Oil Original Set.

Oil droplet detection accuracy: 75.95%			
Absolute performance:			
	Oil droplet	Other	Count
Oil droplet	60	19	79
Other	691	12908	13599
Total	751	12927	13678
Relative Performance :			
	Oil droplet	Other	
Oil droplet	75.95%	24.05%	
Other	5.08%	94.92%	

For both cases of cross-validation (shown in Tables 4 and 5) the detection rate of oil droplets is above 90%. Cross Validation on the *Oil Original Replaced Set* showed a better detection rate with the true negatives comprising only half as many instances as in the cross validation of *Oil Original Set*. Such performance can be explained by the fact that the oil droplets in the *Oil Original Replaced Set* are likely more varied in their appearance. The original set of oil droplets included those that were most clear to the expert. Hence, when you look at them we can expect that they appear quite homogeneous with minimal variation, so they could be sure of the label.

For the largest of our test sets, *Oil Large Test Set*, Classifier I achieved a detection rate of 92.67%, and the Classifier II achieved 94.20%. The false positive rates were 4.42% and 2.74% correspondingly (see Tables 6 and 7). Thus, the *Classifier II* shows better performance for both accuracy of detection and false positive rate. A similar performance in relation to the false positive rate was observed in experiments with the *Non Oil Set* (see Tables 8 and 9). *Classifier II* had a false positive rate of 2.79% as opposed to 7.62% using *Classifier I*. A greater performance difference between *Classifier I* and *Classifier II* was observed while testing on *Oil Random Set*. The detection rate for oil droplets with *Classifier I* was 75.95% and the false positive rate was 5.08% (see Table 10). The detection rate with *Classifier II* was 10% lower in this case, 65.82%, with about one third as many false positives, 1.48% (see Table 11). Overall, the

detection rate was lower for the *Oil Random Set*, than in all previous test cases.

Table 9. Performance of classifier. Testing on Non Oil Set when trained on Oil Original Replaced Set.

False positive rate: 2.79%			
Absolute performance:			
	Oil droplet	Other	Count:
Oil droplet	0	0	0
Other	188	6557	6745
Total:	188	6557	6745
Relative Performance :			
	Oil droplet	Other	
Oil droplet	0.00%	0.00%	
Other	2.79%	97.21%	

Table 11. Performance of classifier. Tested on Oil Random Set, trained on Oil Original Replaced Set.

Oil droplet detection accuracy: 65.82%			
Absolute performance:			
	Oil droplet	Other	Count
Oil droplet	52	27	79
Other	201	13398	13599
Total	253	13425	13678
Relative Performance :			
	Oil droplet	Other	
Oil droplet	65.82%	34.18%	
Other	1.48%	98.52%	

Because of the way the *Oil Random Set* was built, it had a distribution of particles similar to the one expected to be in the water near the oil spill. In examining the current SIPPER data, one finds the percentage of oil droplets in the dataset was about 0.5%. The false positive rate for both classifiers was always higher (1.15-7%). Thus, for regular SIPPER data, it is not yet possible to automatically verify presence of oil droplets in water with the currently built classifiers.

So, we took the class predictions from *Classifier II* on the *Oil Random Set* and extracted probabilities for them from modified version of libsvm [17]. We then ranked the examples classified as oil by probability from highest to lowest. Figure 5 shows a plot of this. We can see that the number of oil droplets is always between 11 and 25% of the predicted oil. The good news is that if an expert looks at the images classified as oil, they will find some oil in the top 10% and top 20% of the classifications (see Figure 6). If they were to randomly search through images with 0.5% oil when there is oil they would need to look at 200 examples to find one oil sample. They will find 4 in the first 25 examined with our tool. So, the user can quit looking if no oil is found in the first 50 or so images that are highly ranked by probability of being oil.

Now, it is clear from looking at Figures 5 and 6 that many of the top probability “oil droplets” are, in fact, not oil. Air bubbles and marine snow can look very similar. In Figure 6, we see that the non-oil images are a little more elliptical in shape. However, oil

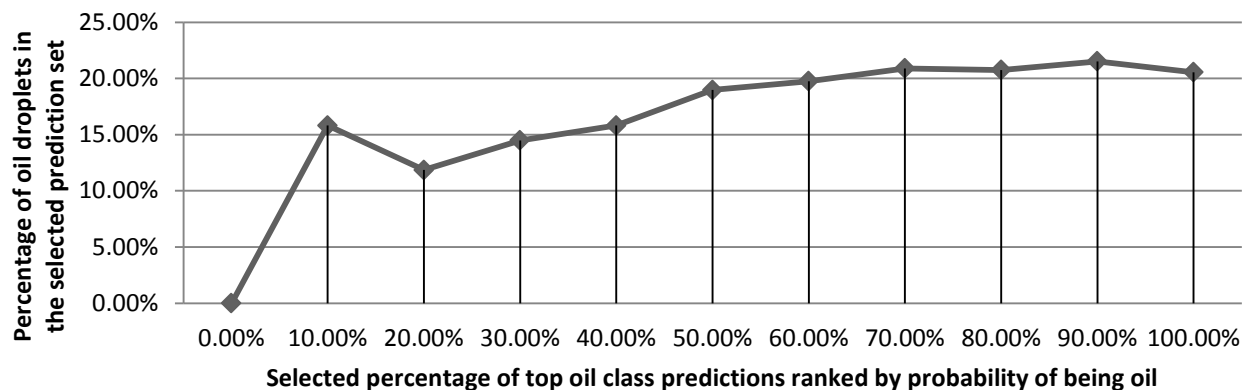


Figure 5. Percentage of oil droplets in the predictions when sorted by probability of being oil. So, 10% means the 25 highest probability predictions for oil of which 4 are actually oil. This is with Classifier II applied to the oil random dataset with results shown in Table 11.

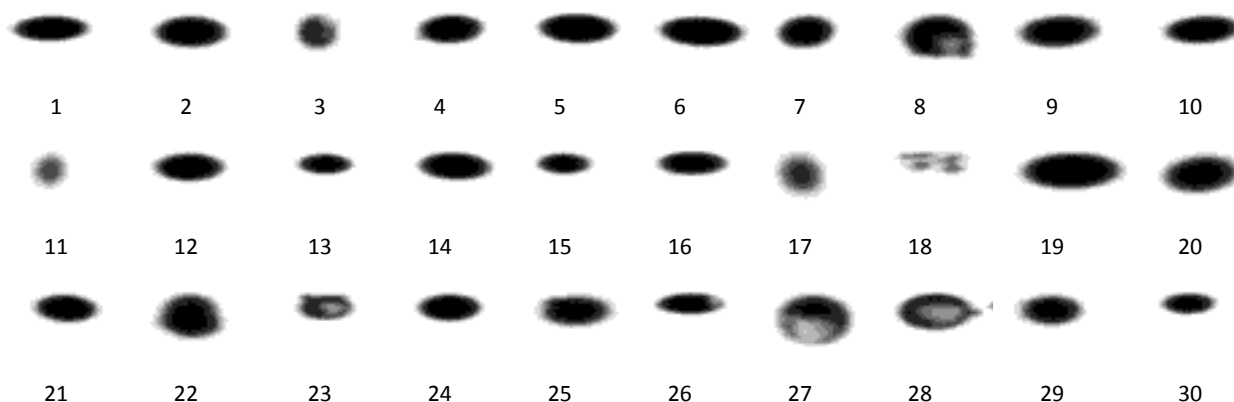


Figure 6. Top 30 particles classified as oil droplets when sorted by the probability. Particles number 3, 8, 23, 27, 28 are suspected oil droplets. Particle number 18 is detritus snow. The others are noise bubbles

does not have to be perfectly spherical as we can see from Figure 1. Reviewing the features selected for each individual binary SVM classifier that comprised this single-stage classifier it was confirmed that the most important features used to discriminate the oil droplet class from others were related to the circularity of the shape, and the texture of the particle. However, it turned out it is not quite enough for completely automatic oil detection.

8. DISCUSSION

The analysis of the particles which are confused with oil the most suggests that there are only three major classes that have an appearance similar to oil droplets: detritus snow, noise bubbles, and protist lopsided. It is possible that a two-stage classifier will allow fully automatic detection of oil droplets in water near the spill. The first stage of the classifier would be aimed at sensitivity to oil droplets, while producing many false positives. The second stage of the classifier aims at specificity to oil droplets. This setup is reported to be useful to detect very rare events and in the case many features which are costly to compute [18]. Certain improvements can be also made in relation to features used to discriminate between the most confused classes. Circularity features were found very useful to discriminate oil droplets, which are often circular in shape, from many plankton organisms. However, those features are not particularly useful with other

classes showing circularity – noise bubbles and marine snow. Further, we have found that our experts use depth to help them classify oil. We have not used this feature because many of the non-image features such as temperature and salinity were no good due to precautions made to protect the SIPPER instrument against damage from any encountered oil.

9. CONCLUSIONS

Overall, a trained SVM achieved a high detection rate for oil droplets. When tested on *Oil Large Test Set*, consisting of 43816 particles of which 13858 were oil, the accuracy of detection was about 95% which is comparable to the cross validation test on the training set. The false positive rate was less than 3% in all experiments with *Classifier II*, which was trained on a random selection of oil examples. We did an experiment with a randomly chosen test set whose distribution mimicked what would be expected during the cruise (about 0.5% oil). For that dataset oil droplet detection was just 65%. It is also the case that in water where there is no oil, our classifiers will predict that a small amount of oil exists.

We showed that by using probabilities for the class predictions and ordering them from highest to lowest, oil will regularly appear in the top 10-30% of data. So, if an expert uses our tool, PICES, to view the images that are predicted to be oil they will be

able to reliably find it much quicker than randomly searching through particles.

There is room for improved oil detection to enable the best analysis of how much oil is in the water. This will occur through classifier tuning/replacement and new features. The current results are promising in terms of finding oil under the water and getting a general count of the number of oil droplets.

10. ACKNOWLEDGEMENT

This work was supported by Baseline for Impact Assessment of Zooplankton and Imaging Oil Droplet Detection on West Florida Shelf, BP/FIO - Gulf Oil Spill Prevention, Response and Recovery Grants Program.

11. REFERENCES

- [1] Deepwater Horizon Unified Command, US Scientific Teams Refine Estimates of Oil Flow from BP's Well Prior to Capping, *Gulf of Mexico Oil Spill Response 2010*; available from: www.deepwaterhorizonresponse.com/go/doc/2931/840475
- [2] Camilli, R. Reddy, C.M. Yoerger, D.R. Van Mooy, B.A.S. Jakuba, M.V., Kinsey, J.C. McIntyre, C.P., Sylva, S.P. and Maloney, J.V., "Tracking hydrocarbon plume transport and biodegradation at Deepwater Horizon", *Science*, v.330, n.6001, pp.201, 2010
- [3] Dykes, B.M., "Researchers find thick patches of crude still on Gulf floor", *Yahoo News*. Retrieved on 2011-04-26
- [4] Samson S., Hopkins, T., Remsen, A., Langebrake, L., Sutton, T., and Patten, J., "A System for High-Resolution Zooplankton Imaging", *IEEE Journal of Oceanic Engineering*, vol. 26, no. 4, October 2001
- [5] Remsen, A., Hopkins, T., and Samson, S., "What You See is Not What You Catch: A Comparison of Concurrently Collected Net, Optical Plankton Counter, and Shadowed Image Particle Profiling Evaluation Recorder Data from Northeast Gulf of Mexico", *Deep Sea Research Part I: Oceanographic Research Papers*, vol. 51, no. 1, pp. 129-151, 2004
- [6] Remsen, A., Samson, S., Hopkins, T., and Kramer, K., "Observations of Plankton and Detrital Particle Distribution on the West Florida Shelf using SIPPER-2 and an Automated Classification System", *Journal of Plankton Research*, submitted 2010
- [7] Remsen, A., "Evolution and field application of a plankton imaging system", Ph.D. Dissertation, College of Marine Science, University of South Florida, 2008
- [8] Burges, C., "A Tutorial on Support Vector Machines for Pattern Recognition", *Data Mining and Knowledge Discovery*, vol. 2, no. 2, pp. 121-167, June 1998
- [9] Platt J., "Probabilistic Outputs for Support Vector Machines and Comparison to Regularized Likelihood Methods" in *Advances in Large Margin Classifiers*, pp. 61-74, Cambridge, MA, USA, 1999
- [10] Hu, M.K., "Visual Pattern Recognition by Moment Invariants", *IRE Transactions on Information Theory*, vol. 8, no. 2, pp. 179-187, 1962
- [11] Zhang, D. and Lu, G., "A Comparative Study on Shape Retrieval Using Fourier Descriptors with Different Shape Signatures", *Journal of Visual Communications and Image Representation*, vol. 14, no. 1, pp. 41-60, 2003
- [12] Hsu, C. and Lin C., "A Comparison of Methods for Multi-Class Support Vector Machines", *IEEE Transactions on Neural Networks*, vol. 13, no.2, pp. 415-425, March 2002
- [13] Witten, I. and Frank, E., "Data Mining: Practical Machine Learning Tools and Techniques", *Morgan Kaufmann Publishers*, 2005
- [14] Staelin, C., "Parameter Selection for Support Vector Machines", HP Laboratories Israel, Technion City, Haifa, 2002
- [15] Kohavi, J.R. and George, H., "Wrappers for Feature Subset Selection", *Artificial Intelligence*, vol. 97, no. 1., pp. 273-324, December 1997
- [16] Silva, H. and Fred, A., "Pair wise vs. Global Multi-Class Wrapper Feature Selection" in *Proceedings of the 6th Conference on 6th WSEAS Int. Conf. on Artificial Intelligence, Knowledge Engineering and Data Bases (AIKED'07)*, vol. 6, Corfu Island, Greece, 2007, pp. 1-6
- [17] Kurt, K., "System for Identifying Plankton from the SIPPER Instrument Platform", *Doctoral Dissertation*, University of South Florida, 2010.
- [18] Senator, T.E., "Multi-stage classification", *Proceedings of the Fifth IEEE International Conference on Data Mining*, 2005