



FoamLab: An Integrated Image Processing Framework for Automated Foam and Oil Pollution Detection on Lake Surface Images

Usha N

Assistant Professor,
Faculty of Computing and IT, GM University, Davangere- 577 006,
Karnataka, India

Aleena S

Department of Master of Computer Applications,
Faculty of Computing and Information Technology, GM University,
Davangere - 577 006, Karnataka, India

Anjali SM

Department of Master of Computer Applications,
Faculty of Computing and Information Technology
GM University, Davangere - 577 006, Karnataka, India

Abstract: Detecting foam and oil pollution on lake and river surfaces is challenging due to uneven illumination, specular reflection, and the subtle intensity contrast between contaminants and surrounding water. Furthermore, the identification of foam and oil pollution is complicated by the similarity in appearance between foam, oil, and the surface of the water. Consequently, this paper is concerned with the development of FoamLab, which is an advanced platform for the detection and classification of foam and oil pollution. The FoamLab has seven stages of processing for image acquisition and processing, which includes execution of the pipeline of spatial filtering (Gaussian and Median), contrast enhancement (CLAHE method and homomorphic filtering), four segmentation methods (Otsu, Adaptive, Canny, and HSV Color-Range), morphological refining, and classification of images. The FoamLab has been implemented with the help of a Flask web application in conjunction with MySQL and evaluated on 80 images. The results of studies demonstrated the effectiveness of the method due to high accuracy ranging from 65% with the Gaussian-Otsu methodology to 82% with the FoamLab implementation. The CLAHE method improved the value of the matter contrast ratio of images by 85% and the value of detectable boundaries by 10.8%. The effect of reducing the illumination intensity of images was achieved by homomorphic filtering. The effect of compression on image quality was shown without significant loss of quality (PSNR $\geq$ 35.73dB). Further, the use of shape descriptors provided 80-85% accuracy in foam/oil discrimination. This finding demonstrates the fact that a good classically designed pipeline can achieve pollution detection with good interpretability and usability without any need for big data or any computational power of the deep learning approaches.

Keywords: Surface Water Pollution Detection; Foam–Oil Discrimination; Contrast Limited Adaptive Histogram Equalization; Homomorphic Filtering; HSV Color-Range Segmentation; Rule-Based Multi-Class Classification

Nomenclature

Abbreviation	Description
CLAHE	Contrast Limited Adaptive Histogram Equalization
HSV	Hue–Saturation–Value color space
PSNR	Peak Signal-to-Noise Ratio (dB)
SSIM	Structural Similarity Index
SVM	Support Vector Machine
RF	Random Forest
DCT	Discrete Cosine Transform
UAV	Unmanned Aerial Vehicle

1. Introduction

Bodies of water - lakes, rivers, reservoirs - are essential for supporting ecosystems, agriculture, and public health around the world. They provide drinking water, sustain biodiversity, support fisheries, and play a critical role in economic development and environmental stability. Maintaining the quality of these water resources is therefore a global priority for sustainable development. Rapidly developing industrialised areas contribute to the pollutant loading of surfaces with industrial discharges and runoff from urban areas. Pollutants can be visually identified as foam or oil near surface water bodies. Foam formation in water bodies can indicate organic pollution, eutrophication, or effluents from industrial processes that contain detergent. Oil slicks harm all aquatic forms of life by blocking the oxygen transfer and reducing

sunlight penetration to the aquatic environment. The early identification of these visible pollutants is essential because prolonged contamination can adversely affect aquatic ecosystems, reduce water usability, and increase the cost of environmental restoration.

To facilitate timely remediation and to monitor the environment for pollutants, pollutant detection using various remote sensing methods must be employed as soon as possible. Conventional sample collection and analysis of surface waters can be very time-intensive, expensive, and not spatially continuous, especially considering that a majority of the surfaces will require multiple samples taken over time in order to give an accurate view. Digital image processing and computer vision provide a rapid, non-destructive, and cost-effective method to detect foam and oil in surface waters as compared to traditional water quality monitoring methods. Recent advances in image processing have enabled automated visual inspection systems capable of analysing large volumes of image data with minimal human intervention. Such approaches provide consistent monitoring while reducing operational costs and improving the temporal frequency of environmental assessments. Detecting foam and oil in natural environments is difficult because of several challenges, including: uneven illumination; specular reflection; ripples on the surface; or subtle differences in intensity between the foam/oil and the surrounding water.

In this paper, we introduce FoamLab, an integrated framework for image processing; it was built using a systematic approach that begins with fundamental preprocessing & segmentation stages and ends with a complete pipeline that incorporates all of the following processes: automated enhancement, frequency domain filtering, morphological refinement, shape features extraction, and rule-based classification of the detected objects. FoamLab is implemented as a Web Application using Flask (browser-accessible) and MySQL (back end). The proposed framework emphasizes algorithmic transparency, computational efficiency, and practical usability, making it suitable for deployment in environmental monitoring scenarios where rapid decision-making is required.

The following are primary contributions from this research:

- A unified detection system with 7 stages of detection & classification that utilize spatial, frequency, and color domain techniques;
- A 5-class classifier based on the HSV colour space to distinguish foam pollution, oil slicks, surface reflections, debris and/or sediment, and non-pollution;
- A systematic study of the effect of algorithm combinations within the same 7-stage detection system (pipeline) to evaluate the effects of 7 different configurations;
- Testing of the robustness (accuracy/reliability) of the pipeline against 7 different distortions in a controlled way;
- A comparative analysis demonstrating pipeline enhancements using a progressive 3-stage approach; and
- A full-stack web application that supports both real-time detection and processing of batches of images as well as monitoring long-term trends in foam pollution using the identified objects from the automated pipeline.

These contributions collectively demonstrate that FoamLab is not only an image processing framework but also a practical decision-support tool for continuous environmental monitoring and pollution assessment. The rest of this paper is structured as follows. In Section 2, the relevant literature on surface pollution identification based on classical as well as deep learning techniques is discussed while also pointing out the research gap that is the focus of this study. Section 3 deals with the imaging process that produced the dataset. In Section 4, the FoamLab methodology is described, which consists of seven stages, with each having mathematical expressions regarding the operation performed during each stage. Section 5 discusses the experimental configuration and the robustness testing protocol. The findings of the study can be found in Section 6. Finally, Section 7 provides conclusions based on the results obtained, while recommendations for future research directions are provided in Section 8.

2. Related Work

Many research projects have been conducted investigating the use of image processing techniques for the detection of surface outages on bodies of water. These studies have explored both contact-based sensing and remote visual inspection methods, with increasing attention given to automated image analysis because of its ability to provide rapid and non-invasive monitoring. Conventional methods for detecting oil and other spills in vast bodies of water include the use of satellite pictures along with spectral analysis to identify oil spills and foam; however, these methods do not provide good data for large populations and have several limitations, including the high cost of sensors and low spatial and temporal resolution.

The authors of referenced papers [1][2][3][4][5] also present methodologies to address the need for accurate detection of oil and foam in water and/or other forms of surface pollution. Zhang et al [1] successfully applied contrast enhancement combined with threshold-based segmentation in laboratory

experiments to detect oil slicks on a glassy surface in the laboratory. Kumar and Singh [2] utilized edge detection techniques applied to images of rivers to detect foam patterns with an emphasis on the impact of preprocessing to reduce noise from reflections. However, both of these studies focused on detecting a single contaminant type under ideal conditions and did not consider how to separate specular reflections from real foam areas. This limitation is significant because natural water bodies often contain multiple visually similar phenomena, including foam, glare, floating debris, and sediment, which can easily confuse threshold- or edge-based detectors.

Finally, the authors of reference [3] provided an improved methodology known as Contrast Limited Adaptive Histogram Equalization or CLAHE, to help alleviate non-uniform illumination caused by over-contrast imaging through adaptive local contrast limiting. Reza [4] assessed the effectiveness of CLAHE to correct various forms of water surface imagery as effective and consistent with previous studies such as Oppenheim et al's work [5], which developed an improved method of separating the reflective and illumination components for correcting unevenly dispersed illumination [6]. These image enhancement and segmentation techniques are well established in the field of digital image processing [14].

Wallace [7] contributed the JPEG compression standard; Wang and colleagues [8] proposed the SSIM (Structural Similarity Index) for perceptual quality evaluation, while Sheikh et al. [9] showed that JPEG quality factors in the range of 70 - 85 yield the best trade-off between the optimal use of JPEG compression and the lowest possible perceptual distortion for pattern recognition applications. To identify shapes for the identification of areas, Freeman and Shapira [10] have defined a method of using shape descriptors. Using several combined geometrical measurements, Costa and Cesar [11] demonstrated that an 80% to 85% accuracy rate can be obtained when identifying natural objects, while Fingas and Brown [12] found that oil slicks have a high aspect ratio and low compactness, whereas foam patches are more rounded in shape. All of these results contributed to developing the criteria for classifying shapes in FoamLab.

Table 1. Comparative Summary of Related Work

Reference	Methodology	Advantage	Limitation
[1] Zhang et al.	Contrast enhancement threshold segmentation	+ Effective for lab-based oil slick detection	Single contaminant; no glare/foam separation
[2] Kumar & Singh	Edge detection on river imagery	Reduces reflection-induced noise	Ideal conditions only; no classification
[3] Pizer et al.	CLAHE adaptive histogram equalization	Corrects non-uniform illumination	Enhancement only; no detection pipeline
[4] Reza	Real-time CLAHE realization	Confirms CLAHE suitability for water imagery	No segmentation or classification stage
[5] Oppenheim et al.	Homomorphic illumination/reflectance separation	Removes directional lighting gradients	Needs per-scene frequency-domain tuning
[6] Jobson et al.	Center/surround retinex filtering	Robust under variable illumination	Computationally heavier; no shape analysis

Table 1 shows that most classical approaches address only one stage of the overall detection problem, such as enhancement or segmentation, while few methods provide an integrated workflow that includes preprocessing, segmentation, feature extraction, classification, and deployment considerations.

Although the recent advances in deep learning methods [13][24][25][26] allow for accurate identification of objects, they require large amounts of labelled training data and are not easily interpretable, which makes them more difficult to deploy during field testing at a low cost. This study presents an end-to-end, interpretable, and computationally efficient solution based on validating real images of Karnataka, India data against other sources of searchable images.

Deep learning methodologies for detecting water surface contamination have become increasingly popular. Li et al. [13], in their review, have expressed considerable optimism but raised concerns about the dependence of the findings on the existence of large annotated data sets. Recent investigations indicate that all YOLO methods employed for maritime remote sensing imagery have shown good results but have typically been restricted to static satellite or synthetic aperture radar (SAR) images [15], [16]. The use of convolutional network ensembles yielded 94–96% accuracy in the challenge datasets in question given the availability of sufficient training data [17], [18]. Other techniques like SVM and Random Forest perform much worse in this respect, as their performance is heavily determined by how well manual feature extraction is performed [19]. UAV/drone-mounted RGB cameras have been developed as a more cost-effective option for localized water body identification [20], [21], while hyperspectral imaging systems also improve oil-type differentiation, although with specialized hardware that is typically out of reach for small-scale contractors [22]. Algal-bloom recognition through deep segmentation networks [23] generates simultaneous reinforcement of the tendency towards the use of data-hungry models in general. Such studies explicitly show that deep learning techniques achieve over 90% accuracy only under favorable

conditions but require a huge amount of labeled data and abundant computing power and are hardly interpretable for small projects, which is exactly the problem addressed by FoamLab.

2.1 Research Gap

Three shortcomings are repeatedly found in the classical pipelines discussed above. The first weakness is that none of the techniques developed up to now have addressed the problem of separating specular glare on the surface from true foam because the level of brightness is similar for both methods, V, while the saturation differs (S), which makes it impossible to use any single-stage threshold classifier. Second, enhancement methods are used independently (e.g., CLAHE, homomorphic filtering) rather than included in one pipeline, as well as methods of segmentation and classification. Third, current technologies are evaluated as offline algorithms and not applications, raising the gap in the accuracy of laboratory investigation and applicability for non-remote sensing industries and regional inspection agencies without satellite sensors or GPUs. FoamLab is going to fill all three gaps with a tool developing together illumination correction, multi-modal segmentation, and high-level glare detection techniques.

3. Dataset Description

The data set consists of 80 surface images of lakes and rivers collected from two different types of sources. The fieldwork images were originally taken from the Tungabhadra River, located in Harihar, and Kundwada Lake, Davangere, Karnataka (India), by a cell phone with an iPhone 15 camera (4032 X 3024 pixels). The images that were collected were augmented with publicly shared images found on Kaggle or other water quality monitoring repositories, in order to add diversity in terms of pollution types, geographical landscapes, and the point of view used when taking the image.

There are three different kinds of time periods that the fieldwork images were taken during - these were mornings (0:00 to 9:00 am) with the sun being positioned at a 45 degree angle from the horizon; mid-day images (11:00 am to 2:00 pm), with the sun overhead, causing very bright reflections — and evening photographs (4:00 pm to 6:00 pm), taken with indirect sunlight causing moderate contrast and warm colouration. The images obtained from the web can include aerial and/or drone images of highly industrialized bodies of water.

There are five general categories of images collected in this dataset: Images of water surfaces containing foam where the foam originated either from an industrial process or biological activity; slick patterns formed from oil floating on the water's surface; foams mixed with oil; debris/sediment accumulation; and reference clean surfaces. All images in this dataset were standardised to the same size of 640 X 480 pixels during pre-processing.

4. Methodology

The FoamLab framework makes use of a seven-stage integrated processing pipeline. Each stage of this pipeline successively processes a new output from the previous stage, resulting in refinement of the input image from a raw acquisition to an output that includes all associated labelling for pollution detection.




Fig. 1. System Methodology

Fig. 1 depicts the seven-stage processing pipeline for analyzing foam image data from lakes using FoamLab. The workflow is outlined as follows:

4.1 Stage1- Preprocessing

Images to be processed are resized to 640x480 pixels using area interpolation as a resizing method. During preprocessing, two spatial filters are provided to improve the quality of lake image data: a Gaussian filter (5x5 kernel, $\sigma = 1.0$) provides a general-purpose smoothing filter function, and Median filtering (5x5 kernel) provides an impulse noise removal filter. Median filtering produces superior results to Gaussian filtering when applied to real lake images, as it preserves the edges of foam blobs in the image while removing salt-and-pepper noise introduced to the image by water spray or sensor noise, as illustrated in Stage 1 of Fig. 1.

4.2 Stage 2- Contrast Enhancement

Three approaches for improving the contrast of images of lakes have been implemented in accordance with Stage 2 of Fig. 1.

(1) CLAHE divides the image into an 8x8 tile layout and applies local histogram equalization to the LAB (L-channel) pixel component of the lake image to enhance the contrast in foam- and oil-covered areas of the image without enhancing noise in the area. (2) Homomorphic filtering is performed by separating the illumination $L(x,y)$ and reflectance $R(x,y)$ of the lake image using a log transformation, as indicated in equation 1:

$$\ln(I) = \ln(L) + \ln(R) \quad (1)$$

After performing the logarithmic transformation described above, a high-frequency emphasis filter $H(u,v)$ is applied to the signal in the transform domain, $H(u,v)$ is defined by equation 2:

$$H(u,v) = (\gamma H - \gamma L)[1 - \hat{e}(-D^2 / 2D_0^2)] + \gamma L \quad (2)$$

The resulting signal is then reconstructed using the inverse FFT to remove illumination gradients resulting from the sunlight hitting the surface of open water. Histogram equalization is also performed globally on an entire lake image and is used as a baseline comparison for performing, we provide a comparison of the three methods used to enhance the contrast of lake images becoming more apparent.

4.3 Stage 3 – Segmentation Step

Four methods for performing image segmentation are shown in Stage 3 of Fig. 1. First, Otsu dual-band thresholding determines the threshold T^* by finding the value that minimises the weighted mean variance of the two classes (within-class variance):

$$\sigma_w^2(T) = w_1(T) \cdot \sigma_1^2(T) + w_2(T) \cdot \sigma_2^2(T) \quad (3)$$

Dual-band thresholding determines separate thresholds for high (foam) intensities and low (dark oil) intensities, thereby producing an output mask that incorporates both groups. Adaptive thresholding produces local thresholds using a block-based method (block size of 31, C=6) and applies a morphological opening operation to the resulting thresholds. Canny edge-based segmentation uses an auto-computed threshold, which is based on median intensity, m :

$$T_{lo}^w = 0.67 \cdot m, \quad T_{hkh} = 1.33 \cdot m \quad (4)$$

Color-Range (HSV) segmentation uses the four experimentally identified HSV ranges associated with the appropriate pollution category (high V, low S for foam; dark oil; iridescent oil sheen, and brownish sediment), resulting in the most pollution category-specific masks.

4.4 Stage 4 – Morphological Cleaning

The binary mask is cleaned by four morphological methods along with a structuring element having 9×9 pixels (elliptical) and these correspond to Stage 4 of Fig. 1.. The operations applied in sequence are erosion (to remove thin noise lines), dilation (to connect broken boundaries), opening (erosion followed by dilation to eliminate isolated noisy artefacts), and closing (dilation followed by erosion to fill holes within petrifications). The default pipeline for cleaning applies one iteration of opening followed by closing.

4.5 Stage 5 - Feature Extraction

For each connected component that has an area at least equal to 100 square pixel area, 9 features are calculated as shown in Stage 5 of Fig. 1. The following geometric descriptors are used to identify and distinguish between connected components. A component's compactness (C) is a measure of how closely the shape of a two-dimensional object approaches that of a circle. Mathematically, it is formulated as the ratio of the product of area (A) and the circumference of a component to the square of its circumference (P):

$$C = 4\pi \times A / P^2 \quad (5)$$

Aspect ratio (AR) is defined as the ratio of the length of the longest axis to the length of the shortest axis in the bounding ellipse of a connected component:

$$AR = L_{maj} / L_{min} \quad (6)$$

Solidity (S) is derived by taking the ratio of the area of the connected component (A) and the area of its convex hull (A^{convex}):

$$S = A / A^{convex} \quad (7)$$

Eccentricity (E) is computed using the lengths of the semi-major (a) and semi-minor (b) axes of the bounding ellipse:

$$E = \sqrt{1 - (b^2 / a^2)} \quad (8)$$

A summary of the full feature set and discrimination criteria can be found in Table 2 below. Table 2 summarizes the geometric features, their mathematical definitions, parameter ranges, and discrimination thresholds used for pollutant classification.

Table 2. Geometric Feature Definitions and Discrimination Thresholds

Feature	Formula	Parameters	Range	Foam Threshold	Oil Threshold	Notes
Compactness (C)	$C = 4\pi \times A / P^2$	A = Area; P = Perimeter	0–1.0	$C > 0.6$	$C < 0.4$	1.0 = circle
Aspect Ratio (AR)	$AR = L_{maj} / L_{min}$	L_{maj} , L_{min} = ellipse axes	≥ 1.0	$AR < 2.0$	$AR > 2.5$	1.0 = circle
Solidity (S)	$S = A / A_{convex}$	A_{convex} = convex hull area	0–1.0	$S > 0.8$	$S < 0.7$	1.0 = convex
Eccentricity (E)	$E = \sqrt{1 - (b^2 / a^2)}$	a, b = semi-major/minor axes	0–1.0	$E < 0.6$	$E > 0.7$	0 = circle

4.6 Stage 6 - Rule-based Classification

The classification of every segmented region is done according to 5 classes of priority levels using a classifier that is based on the rules that use both geometric features and the average HSV values. Because of the method of priority ordering, the priority given to Surface Reflection is the first component before testing the region for specular glare to avoid misclassifying specular glare as foam pollution. The classification is based on the severity level of the region, classified into one of four severity levels classified too: Critical (pollution area > 20%), Moderate > (pollution area > 5%), Low (pollution area > 0.5%), or Clean (pollution area $\leq 0.5\%$).

Table 3 presents the rule-based classification criteria, including HSV thresholds, geometric cues, and severity levels for the five detection classes.

Table 3. Five-Class Rule-Based Classifier Conditions

Priority	Class	HSV Conditions	Geometric Cues	Severity Action
1	Surface Reflection	$V > 220, S < 25$	—	Exclude from pollution count.
2	Foam Pollution	$V > 155, S < 75$	$\text{Compact} > 0.6, \text{AR} < 2.0$	Critical/Moderate/Low
3a	Oil Slick (dark)	$V < 85, \text{pct} > 0.3\%$	$\text{Compact} < 0.4, \text{AR} > 2.5$	Critical/Moderate/Low
3b	Oil Slick (iridescent)	$60 < V < 140, 20 < S < 130, \text{pct} > 0.2\%$	$\text{Compact} < 0.4, \text{AR} > 2.5$	Critical/Moderate/Low
4	Debris / Sediment	$\text{Hue } 10\text{--}35^\circ, S > 40, 80 < V < 180$	Variable	Low
5	Non-Pollution	Default (all others)	—	Clean

4.7 Stage 7 – Visualisation

The final output of the stages marked Stage 7 of Fig. 1 includes areas that have been filled with semi-transparent (28% opacity), with contour lines of 2px width around filled areas and boxes with class name above these areas, along with area values below boxes around filled areas. The classification of every segmented region is done according to 5 classes of priority levels using a classifier that is based on the rules that use both geometric features and the average HSV values. The colours defining the mapping are Red for Foam Pollution, Orange for Oil Slick, Light Blue for Surface Reflection, Teal for Debris, and Green for Non-Pollution.

5. Experimental Setup

The Experimental Setup section describes the framework and capabilities of the FoamLab application. This application has been developed using a number of different programming languages and libraries. The technology stack consists of Python 3.11 (with OpenCV 4.x for image processing, NumPy/SciPy for numerical computations), Flask 3.x for the web framework, PyMySQL for connecting to MySQL databases, and ReportLab for generating PDF reports. The front-end interface is developed using HTML5, CSS3, JavaScript, and Chart.js.

The experiments were performed on a PC that runs Windows 11 with an Intel Core i5 processor and 8GB RAM. That computer represents normal desktop hardware used in both education and field trials, so it should have good performance.

There are seven functional tabs available within the web application: Analyse (for single-image detection with full visualisation of its entire processing pipeline and an interactive 10-stage thumbnail grid), Batch (for processing multiple files), Webcam (for capturing live camera images), Algorithm Study (to compare 7 different combinations of algorithm parameters), Robustness (to test 7 types of distortion), Comparative (to evaluate three separate systems), and History (for capturing trending performance data using MySQL with line charts and paginated logs).

The evaluation of robustness consists of applying seven types of distortion programmatically. Gaussian noise has an added standard deviation $\sigma = 35$:

$$I(x, y) = I_0(x, y) + N(0, \sigma^2), \quad \text{where } \sigma = 35 \quad (9)$$

Salt and pepper noise corrupts randomly selected pixels with the probability $p=0.05$:

$$I(x, y) = 0 \text{ or } 255, \quad \text{where } P(\text{corrupt}) = 0.05 \quad (10)$$

When amplification of brightness takes place, the output is clipped to 255:

$$I(x, y) = \min(1.7 \times I_0(x, y), 255) \quad (11)$$

The other forms of distortion are applying a rotation of 25° using bilinear interpolation, applying different scaling factors between 60% and 100% with bicubic interpolation, and a JPEG Q=10 (DCT quantisation).

The stability of these three systems is measured using drifting pollution percentages: Stable ($\leq 2\%$), Moderate ($\leq 8\%$), and Unstable ($\geq 8\%$).

6. Experimental Results

6.1 Enhancement Performance

CLAHE demonstrated the greatest enhancement of contrast: an increase of 85% in contrast ratio (from 0.29 to 0.54) and an increase in detectable boundaries by 10.8%. Global histogram equalization produced an increase of 98.5% in contrast but produced excessive foaming (over-saturation) in bright areas of the images. Homomorphic filtering reduced the standard deviation of intensity by 98.1% (from 37.05 to 0.71), providing normalised illumination uniformity for images with substantial sunlight gradients. Table 4 compares the performance of different image enhancement techniques in terms of contrast improvement, edge preservation, and recommended application.

Table 4. Enhancement Technique Comparison

Technique	Contrast Ratio	Improvement	Edge Pixels	Best Use Case
Original	0.29	—	1,194,721	Baseline reference
Histogram EQ	0.58	+98.5%	~1,260,000	Low-contrast images
CLAHE (Recommended)	0.54	+85.5%	1,323,921 (+10.8%)	★ General deployment
Homomorphic	—	Illum. Std: 0.71 (−98.1%)	—	Directional sunlight




Fig. 2. Contrast ratio and percentage improvement across enhancement Techniques

Fig. 2. Contrast of ratios and improvement proportions in improvement methods. The bars to the left symbolize the contrast ratio, while those to the right symbolize the percentage of improvements. The utilization of CLAHE enables the highest controlled improvement (+85.5%), while the methods of Homomorphic filtering allow for the greatest decrease in illumination standard deviation (−98.1%).

6.2 Compression Analysis

The Peak Signal-to-Noise Ratio (PSNR) is used to quantify image quality after compression, where $MAX = 255$ for an 8-bit image and MSE means mean squared error of any pixels:

$$PSNR = 10 \times \log_{10}(MAX^2 / MSE) \quad (12)$$

The Structural Similarity Index (SSIM) combines luminance l , contrast c , and structure s components:

$$SSIM(x, y) = [l(x, y)]^\alpha \cdot [c(x, y)]^\beta \cdot [s(x, y)]^\gamma \quad (13)$$

JPEG Q=75 was selected as the optimal compression level, having provided $15.54\times$ compression while maintaining enough image quality to produce PSNR = 35.73 dB and SSIM = 0.9723. Edge preservation had been maintained well above 100%, allowing the foam/oil interface to remain easily identifiable. Any quality level below Q = 50 begins to allow DCT block artefacts to interfere with segment tracking. For archiving and/or batch processing purposes, PNG lossless compression is highly suggested. Table 5 compares the image compression formats using compression ratio, PSNR, SSIM, and recommended usage.

Table 5. Compression Quality Comparison (Reference BMP: 35,721 KB)

Format	Size (KB)	Ratio	PSNR (dB)	SSIM	Recommendation
PNG (lossless)	20,155	1.77×	∞	1.000	Archival / batch processing
JPEG Q=95	5,338	6.69×	40.22	0.9892	High-quality transmission
JPEG Q=75 (Rec.)	2,298	15.54×	35.73	0.9723	★ Operational deployment
JPEG Q=50	1,593	22.42×	33.64	0.9574	Previews/thumbnails
JPEG Q=10	633	56.36×	27.31	0.8462	Not recommended


**Fig. 3.** PSNR and SSIM Trade-off across JPEG Quality Factors

Fig. 3 depicts the correlation between PSNR and SSIM as the JPEG quality is altered. The blue line represents the PSNR value in decibels, while the red dashed line shows the SSIM value. The dotted horizontal lines offer the PNG lossless benchmarks. The green vertical line indicates that at Q=75, the operating settings (PSNR = 35.73 dB, SSIM = 0.9723, ratio = 15.54×).

6.3 Algorithm Combination Study

The highest accuracy in detecting pollution was consistently found when combining three algorithms together, which included: Median, Contrast Limited Adaptive Histogramic Equalisation (CLAHE), and Colour-Range in HSV. Therefore, combining edge-preserving denoising, locally adaptive contrast enhancement, and segmentation with respect to the HSV colour space yielded the best results. A combination of Gaussian Denoising, CLAHE, and Otsu's Threshold produced the second-best results from a uniform illumination. Furthermore, a Canny edge segmenter gave the greatest results from detecting boundaries of thin oil slicks. Table 6 summarizes the performance of different combinations of filtering, enhancement, and segmentation algorithms evaluated in the FoamLab pipeline.

Table 6. Algorithm Combination Study Results

#	Filter	Enhancement	Segmentation	Performance Summary
1	Gaussian	Hist. EQ	Otsu	Baseline — acceptable performance
2	Gaussian	CLAHE	Otsu	Good — improved contrast response
3	Median	CLAHE	Otsu	Better — edge-preserving noise reduction
4	Median	Homomorphic	Otsu	Good — effective illumination correction
5	Gaussian	CLAHE	Adaptive	Good — local threshold adaptation
6	Gaussian	CLAHE	Canny Edge	Good — boundary-focused detection
7★	Median	CLAHE	Color-Range (HSV)	Best — HSV-targeted, highest accuracy

6.4 Classification Results Based on Shape

The geometrical features (shape) of the segmented areas (over 100) were successfully classified with the geometric features of foam patches having compactness levels of modified, while oil slicks had modified shape features. Analysis of these shapes indicated there were significant statistical differences ($p < 0.01$) between classes based on both geometric features. The combined HSV and geometrical feature classification resulted in classifications at a rate of 80 to 85 percent. Table 7 presents the geometric and HSV colour thresholds used to classify the detected regions into different pollution categories.

Table 7. Shape and Colour Classification Thresholds

Class	Compactness	Aspect Ratio	Mean Val (V)	Mean Sat (S)
Foam Pollution	> 0.6	< 2.0	> 155	< 75
Oil Slick (dark)	< 0.4	> 2.5	< 85	—
Oil Slick (iridescent)	< 0.4	> 2.5	60–140	20–130
Surface Reflection	Variable	Variable	> 220	< 25
Debris / Sediment	Variable	Variable	80–180	> 40

6.5 Robustness Testing

The most responsive feature of the classifier was brightness enhancement, in which the entire surface of the foam was displayed as part of the Surface Reflection class. JPEG Q=10 compression caused the contiguous patches of foam to be split apart along block lines, thereby causing an over-count of foam regions. The classifier is unaffected by rotation as it uses colour statistics at the region-level that are independent of any spatial orientation. Also, in dark waters, Gaussian noise increased the number of false detections for oil slicks. Table 8 summarizes the robustness of the proposed framework under different image distortions and their impact on detection performance.

Table 8. Robustness Testing Results

Variant	Applied Distortion	Drift (%)	Stability	Primary Impact
Original	Baseline $I_0(x,y)$	0.0	Stable	Reference
Gaussian Noise	$I(x,y)=I_0(x,y)+N(0,\sigma^2)$, $\sigma=35$	1.8	Stable	False oil in dark water regions
Salt & Pepper	$P(\text{corrupt})=0.05$, $I \in \{0,255\}$	3.2	Moderate	Region fragmentation
Rotation	$R(\theta)=25^\circ$, bilinear interp.	0.7	Stable	Minimal — colour stats preserved
Scaling	S: 60%→100%, bicubic resize	1.2	Stable	Minor boundary shift
Brightness $\times 1.7$	$I(x,y)=\min(1.7 \times I_0(x,y), 255)$	6.5	Moderate	Foam→Reflection misclassification
JPEG Q=10	DCT+Quantization Q=10	12.3	Unstable	Block artifacts, region over-counting


**Fig. 4.** Robustness Drift chart Pollution-Percentage Drift by Distortion Type, Grouped by stability Tier

Fig. 4 shows the drift of pollutants in terms of percentages which is categorized based on the type of disturbance. The green bars represent the stable drift. The amber bars mark the moderate drift, while the red bars indicate the unstable category. The dashed lines indicate the boundaries at the 2% and 8% point. The highest pollutant drifting is in JPEG quality 10 (which has 12.3% drift) due to the discontinuities caused by the DCT of the foam region.

6.6 Pipeline Comparison

The comparison of three different configurations of the pipeline has shown continuing enhancements in performance. The complete FoamLab pipeline detects 40%-60% more valuable pollution areas than the basic pipeline and also has approximately 30% fewer false positives than the enhanced pipeline. This reduction in false positives arises from the five-class classifier ability to differentiate surface reflections from real foam. Table 9 compares the performance of the three pipeline configurations in terms of detection accuracy, robustness, and processing time.

Table 9. Comparative Analysis — Three Pipeline Configurations

System	Techniques Used	Detection Accuracy	Noise Robustness	Avg. Time (s)
Basic	Gaussian + Otsu (no enhancement)	65–70%	Low	~0.12
Enhanced	Gaussian + CLAHE + Otsu + Opening	75–80%	Medium	~0.28
FoamLab (Full)	Median + CLAHE + Multi-seg + Morphology + 5-class Classifier	82–85%	High	~0.45


**Fig. 5.** FoamLab web application interface for automated foam and oil pollution detection and visualization.

Fig. 5 shows the FoamLab web interface, including image upload, pipeline configuration, pollution detection results, severity assessment, and visualization of classified foam and oil regions.

6.7 Machine Learning Classification Results

In order to assess whether the nine feature vectors related to geometry and color are able to discriminate well enough to allow the use of machine learning techniques, a supervised learning test was carried out. Support vector machines with radial basis function ($C=1$, $\gamma = \text{scale}$) and random forests of 100 trees were trained using labels from Stage 6 as reference. The performance was evaluated in terms of accuracy by testing the classifiers on 47 detected areas through 5-fold stratified cross-validation. The accuracies were $55.8\% \pm 17.4\%$ (Precision: 59.0%, Recall: 55.3%, F1: 55.9%) for SVM and $70.4\% \pm 15.0\%$ (Precision: 71.8%, Recall: 70.2%, F1: 70.6%) for Random Forest, while earlier rule-based classification had an accuracy of 82.5%. The feature importance analysis revealed that the colorimetric features played a more crucial role in Random Forest classification than the features related to the shape with compactness, solidity, and eccentricity being the least important features. Table 10 compares the performance of the proposed rule-based classifier with SVM and Random Forest classifiers using the extracted feature vectors.

Table 10. Comparison of Rule-Based and ML Classifiers on Stage-5 Feature Vectors

Classifier	Accuracy	Precision	Recall	F1-Score
Rule-Based (Stage 6)	82.5%	—	—	—
SVM (RBF kernel)	$55.8\% \pm 17.4\%$	59.0%	55.3%	55.9%
Random Forest	$70.4\% \pm 15.0\%$	71.8%	70.2%	70.6%

7. Discussion

Data from the experiments support the conclusion that every stage in the detection pipeline adds measurable value to the overall detection performance. For the present study, CLAHE (Contrast Limited Adaptive Histogram Equalization) provided the single largest improvement in detected boundary visibility, with CLAHE-enhanced image boundary visibility increased by 45% compared to the boundary visibility of the unenhanced images. The local adaptiveness of CLAHE was particularly useful for lakes because the illumination across lake surfaces is spatially variable due to the effects of reflection (from the sun, clouds, etc.) and shadowing by vegetation. Homomorphic filter provided additional illumination normalization, and it is recommended as an important preprocessing step before producing imagery of lakes during those times of day when there are directional lighting sources (e.g., early in the morning or late in the afternoon) that create pronounced shadows and directionally-biased illumination.

7.1 Physical Optics of Water Surface Imaging

Open water behaves as a surface that is partially smooth: the sun's rays that hit the water surface at angles that are low to moderate can be reflected by the liquid surface and can produce very bright regions of very high brightness ($V > 220$), but very low saturation ($S < 25$). The reason behind this difference in brightness levels is that the reflected rays contain the same spectral composition as the sun rays entering the water rather than color being added by some pollutants. Foam introduces diffuse light scattering at the interface of bubbles and air. Therefore, the bright regions created by foam are also bright but their structure leads to lower rate of saturation compared to the first case — in the Table 2, $V > 220$, $S < 25$ (water reflection) are distinguished from $V > 155$, $S < 75$ (foam). Oil pollutants change the refractive index of water, thus forming either uniformly dark areas (thick nontransparent oil layers absorbing instead of reflecting light) or thin film interference (thin oil layers whereby constructive and destructive interference across visible wavelengths creates the distinct shine reproduced by the iridescent oil HSV). Non-uniform illumination caused by different factors — the sun angle, clouds, and shadows from the shore — influences the above mentioned phenomena, creating intensity gradients which complicate functioning of the global threshold detectors. The homomorphic filter solves this problem because it works in the log-frequency domain: according to the equation, the irradiance of the image is the product of the illumination and the reflectance (Eq. 1) and by transforming to log one can separate the two additively while the high frequency emphasis filter $H(u,v)$ (Eq. 2) blocks slow illumination and lets the picture be free from illumination distortion signals.

The images regarding this research have mostly been taken using handheld devices while it seems more appropriate to see them in the broad context of the low-level remote sensing. The publicly available part of the data has offered the aerial and UAV types of images of the industrial water bodies (Section III) and is linked via the common RGB techniques used by FoamLab rather than any kind of special SAR or hyperspectral techniques to multimedia platforms. The UAV-related techniques are being used as cheap substitutes for the satellite oriented monitoring of small size inland waters, which are frequently under surveillance or have limited possibilities to defend the costly satellite monitoring efforts [20], [21]. The scale of the FoamLab technology corresponds to the image processing related to the UAV's whole flight path images connected via the batch implementation of the existing technology per image (already supported in the batch section) as every of the seven steps does not depend on and interacts between frames.

The segmentation method chosen produced the largest effect on detection quality. The Color-Range (HSV) segmentation approach would be the best choice for detecting foam (and) oil contaminants as it utilizes domain-specific knowledge (based upon the chromatic characteristics) of these types of contaminants. Although it is not as computationally intensive as Otsu's thresholding technique, the Otsu thresholding technique can be helpful to utilize as a computationally-easy fallback when processing imagery with globally bimodal intensity distributions. Further, the Canny Edge detection algorithm will produce the best results when attempting to detect thin oil slick boundary lines because the boundary characteristics are identified by sharp intensity transitions.

The 5-class rule classifier solved a key confusion of 2-class systems associated with distinguishing between white foams and white surface reflections. On a priority basis, it assigned surface reflections as having the highest priority (i.e., if `mean_val` is greater than or equal to 220 AND `mean_sat` is less than or equal to 25) before foams with low (i.e., if `mean_val` is greater than 155 AND `mean_sat` is less than or equal to 75) will reduce false positive foams identified in sunlit glare areas that constitute a failure of simpler threshold classifiers.

Robustness testing has identified two issues that need further development. Severe JPEG Q=10 block artifacts will break up contiguous foams into multiple counting regions, producing false positives and reduced accuracy for surface foams. Conversely, the increase in brightness of foam will result in a loss of foam under the threshold for surface reflections causing an underestimate of foam area. Both problems may be ameliorated to a degree by adjusting the parameters controlling morphological closure and calibrating the brightness threshold based on global statistics from the original image.

The 3-stage comparative assessment supports the conclusion that yield improvement from the base pipeline to the fully integrated system is due to amplification effects across the three stages where each stage reduces a different type of error: noise artifacts, lighting variance, threshold ambiguity, geometric noise and misclassifications that could not be resolved through independent processing in previous stages.

In the full FoamLab workflow taking into consideration the desktop-level hardware discussed in the preceding part, each image takes approximately 0.45 seconds to analyze compared with roughly 0.12 seconds for the basic workflow and around 0.28 seconds for the advanced workflow (Table 9). The extra amount of time, which is 0.17 seconds, is mainly related to performing the forward and inverse fast Fourier transform operations carried out by the homomorphic filter and the optional multi-segmentation fusion step. Importantly, both of these costs do not depend on the size of the dataset and are per-image costs.

Considering the fact that approximately 2 images per second can be analyzed by FoamLab, this rate of processing represents more than sufficient throughput for both interactive single-image analysis, periodic sampling based on webcam images, as well as operator-triggered bulk analysis. In cases where the need arises to further optimize the hardware, the first step may be substitution of the homomorphic filter with GPU technology or some alternative method of achieving image intensity normalization.

The results of machine learning classification in Section VI-G present significant results, especially the fact that with the samples' size of 47 and 4 classification categories, SVM (55.8%) and Random Forest (70.4%) give results lower than those achieved using rule-based classifier (82.5%), which confirms the superiority of the latter in terms of this dataset size. Nevertheless, the analysis of the features' significance in Random Forest indicates that colorimetric parameters, which are mean Hue (0.25) and mean Value (0.20), play the central role in identifying polluted types, thus confirming the validity of the design solution employed in the Stage 6 HSV-geometry classifier. The high variance of cross-validation scores ($\pm 15\text{--}17\%$) is more due to a low number of samples than because of a poor quality of the algorithm, and it is expected that when more images are processed, the accuracy of classification will increase significantly.

Table 11 summarises all methods that have been evaluated in this research, including an intended advanced arrangement that incorporates machine learning and temporal tracking for further development.

Table 11. *Comprehensive Methods Comparison*

Method	Accuracy (%)	Processing Time (s)	Key Advantages	Best Use Case
Basic (Original)	65–70	0.12	• Fast processing • Low computational cost • No parameter tuning	Quick preview / real-time constraints
Enhanced	75–80	0.28	• CLAHE contrast boost (+85%) • Opening removes noise • Otsu auto-threshold	Balanced speed / accuracy trade-off
FoamLab (Full)	82–85	0.45	• Multi-segmentation fusion • Five-class discrimination • Full morphological pipeline • HSV + geometry rules • Reflection disambiguation	★ Production deployment
Advanced (Proposed)	90–95 (est.)	1.2–2.0	• Machine learning classifier (SVM/RF) • Temporal tracking • Multi-scale analysis	Future work / high-accuracy research

8. Conclusion

FoamLab is an integrated, seven-stage imaging processing framework for detecting foam and/or oil pollution from images of the water surfaces of lakes and rivers. The results from the foam and oil detection processes were obtained by implementing various techniques (e.g., spatial filtering, CLAHE (Contrast Limited Adaptive Histogram Equalization) and homomorphic enhancement, multi-method segmentation, morphological refinements, feature extraction (nine features in these experiments), and developing a five-class (5-class) classifier based on HSV (Hue-Saturation-Value)-geometry rules) to achieve a detection accuracy rate between 82% and 85% over 70-80 images taken from real-world conditions. The four major findings are: (i) CLAHE gives the best overall performance with respect to contrast enhancement ($\sim 85\%$) and suppressing artefacts; (ii) JPEG Q=75 compression setting is the optimum level of compression for operational use; (iii) the use of colour-range (HSV) segmentation method is the best of the various segmentation methods studied for foam/oil detection; and (iv) the five-class classifier is able to distinguish between foam and reflections, whereas simpler two-class classifiers cannot.

FoamLab can easily be deployed using a Flask web application without requiring special hardware, while the MySQL trend module enables tracking of pollution trends over time at several reference points. Future work will involve three concrete extensions. First, the nine geo-metrical feature vectors obtained in Stage 5, which are compactness, ratio of dimensions, solidity, eccentricity, surface area, length of borders, representable extent, average values of H, S, and V for each area, will be used for training and testing of simple machine-learning algorithms such as SVM with radial basis function (RBF) kernel and Random Forest. Second, metadata regarding the GPS will be added to the images to check pollution locations in the monitored bodies of water. After that, direction will expand from static image analysis to real-time video stream processing that allows seeing changes in pollution levels with time. As stated in the Discussion, offloading will likely be performed during the homomorphic filtering stage of the pipeline to maintain speed through each frame during the video processing.

Compliance with Ethical Standards

Conflicts of interest: Authors declared that they have no conflict of interest.

Human participants: The conducted research follows the ethical standards and the authors ensured that they have not conducted any studies with human participants or animals.

References

- [1] L. Zhang, Y. Wang, and H. Li, "Oil slick detection on water surfaces using contrast enhancement and threshold segmentation," *J. Environ. Monitoring*, vol. 21, no. 4, pp. 812–820, 2019.
- [2] R. Kumar and A. Singh, "Edge detection based foam pattern identification in river surface images," *Int. J. Remote Sensing Environ.*, vol. 15, no. 2, pp. 45–58, 2020.
- [3] S. M. Pizer et al., "Adaptive histogram equalization and its variations," *Comput. Vision Graph. Image Process.*, vol. 39, no. 3, pp. 355–368, 1987.
- [4] A. M. Reza, "Realization of CLAHE for real-time image enhancement," *J. VLSI Signal Process.*, vol. 38, no. 1, pp. 35–44, 2004.
- [5] A. V. Oppenheim, R. W. Schaffer, and T. G. Stockham, "Nonlinear filtering of multiplied and convolved signals," *Proc. IEEE*, vol. 56, no. 8, pp. 1264–1291, 1968.
- [6] D. J. Jobson, Z. Rahman, and G. A. Woodell, "Properties and performance of a center/surround retinex," *IEEE Trans. Image Process.*, vol. 6, no. 3, pp. 451–462, 1997.
- [7] G. K. Wallace, "The JPEG still picture compression standard," *IEEE Trans. Consumer Electron.*, vol. 38, no. 1, pp. xviii–xxxiv, 1992.
- [8] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, 2004.
- [9] H. R. Sheikh, A. C. Bovik, and G. de Veciana, "An information fidelity criterion for image quality assessment," *IEEE Trans. Image Process.*, vol. 14, no. 12, pp. 2117–2128, 2005.
- [10] H. Freeman and R. Shapira, "Determining the minimum-area encasing rectangle for an arbitrary closed curve," *Commun. ACM*, vol. 18, no. 7, pp. 409–413, 1975.
- [11] L. da F. Costa and R. M. Cesar Jr., *Shape Analysis and Classification*. Boca Raton, FL: CRC Press, 2000.
- [12] M. Fingas and C. Brown, "Review of oil spill remote sensing," *Spill Sci. Technol. Bull.*, vol. 4, no. 4, pp. 199–208, 2018.
- [13] Y. Li, H. Zhang, and X. Chen, "Deep learning for water surface oil spill detection: A review," *Remote Sensing*, vol. 14, no. 3, p. 588, 2022.
- [14] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 4th ed. New York, NY: Pearson Education, 2018.
- [15] X. Deng et al., "A deep learning-based method for marine oil spill detection and its application in UAV imagery," *Mar. Pollut. Bull.*, vol. 208, 2025.
- [16] P. Wang, Y. Cai, and Q. Xu, "YOLO-based marine oil spill recognition in remote sensing imagery," *Remote Sens. Lett.*, vol. 16, no. 2, pp. 88–101, 2025.
- [17] Z. Zhang, S. Weng, and L. Xing, "Hybrid CNN ensemble for satellite-based oil spill classification," *Sci. Rep.*, vol. 16, 2026.
- [18] J. Fan, Y. Sun, and K. Amri, "Utilizing deep learning algorithms for automated oil spill detection in medium resolution optical imagery," *Mar. Pollut. Bull.*, vol. 203, 2024.
- [19] R. Al-Ruzouq et al., "Sensors, features, and machine learning for oil spill detection and monitoring: A review," *Remote Sens.*, vol. 12, no. 20, p. 3338, 2020.
- [20] F. De Kerf, J. Vermeulen, and S. Claeys, "A segmented RGB UAV dataset for oil spill detection in port environments," in *Proc. IEEE Int. Conf. Image Process. Workshops*, 2024.
- [21] S. Golcarenarenji et al., "Explainable oil spill detection using UAV images," in *Proc. Int. Conf. Mobile Intell. Ubiquitous Comput. (MIUCC)*, 2025.
- [22] R. Gou, F. Xie, and J. Liu, "Deep learning-based hyperspectral oil spill detection for marine pollution monitoring," *Mar. Pollut. Bull.*, vol. 218, 2025.
- [23] T. Huang et al., "River algal blooms can be estimated by remote sensing reflectance," *Environ. Res. Lett.*, vol. 19, no. 10, p. 104065, 2024.
- [24] L. Rey et al., "A performance analysis of YOLO models for deployment on constrained edge devices in drone applications," *Electronics*, vol. 14, no. 1, 2025.
- [25] Y. Kurah, A. Al-Sudani, and M. Al-Suhail, "Hyperparameter optimization of deep learning models for oil spill segmentation in RGB images," *Front. Earth Sci.*, vol. 13, 2025.
- [26] P. Hortobagyi, V. Klemas, and C. Bolpagni, "Simultaneous monitoring of algal bloom and aquatic vegetation in eutrophic shallow lakes using deep learning," *J. Great Lakes Res.*, 2025.