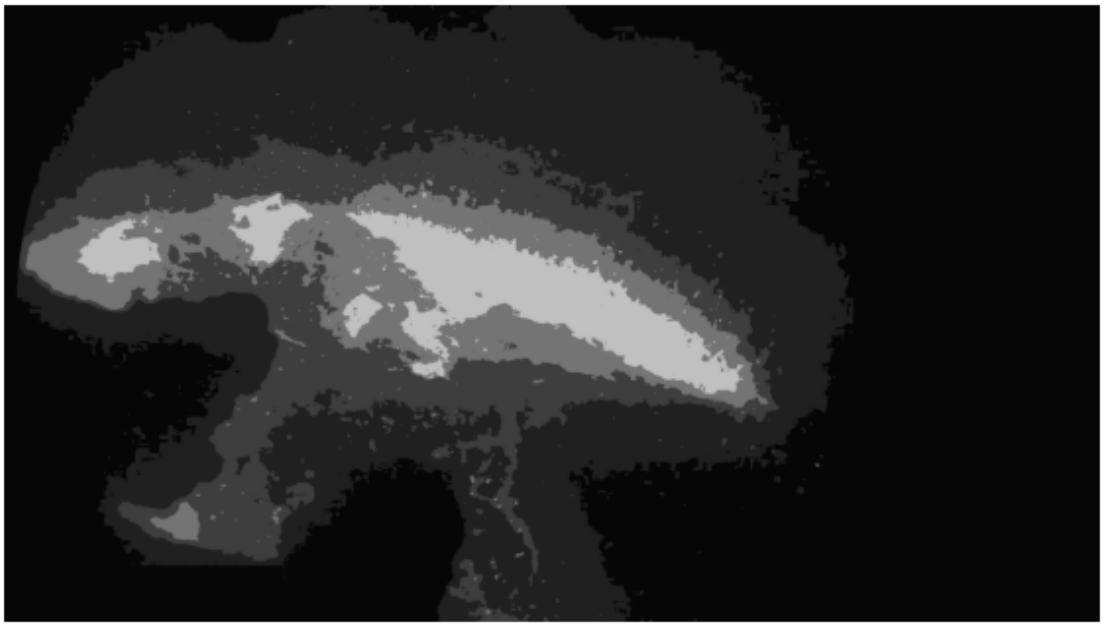

Video analysis of the degassing process

A proof of concept for information extraction from camera feed using image analysis techniques

Master's thesis in Engineering Mathematics and Computational Science

NOAH LANAI

DEPARTMENT OF MATHEMATICAL SCIENCES

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2023
www.chalmers.se

Master's thesis 2023

Video analysis of the degassing process

A proof of concept for information extraction from camera feed using image analysis techniques

NOAH LANAI

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Mathematical Sciences
Chalmers University of Technology
Gothenburg, Sweden 2023

Video analysis of the degassing process
A proof of concept for information extraction from camera feed using image analysis techniques
NOAH LANAI

Supervisors: Krister Ekström, Swerim AB. Marcus Svadling, Ovako Sweden AB.
Professor Aila Särkkä, Department of Mathematical Sciences
Examiner: Professor Aila Särkkä, Department of Mathematical Sciences

Master's Thesis 2023
Department of Mathematical Sciences
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: Segmented image from the degassing process camera feed in Ovako's steel production

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2023

Video analysis of the degassing process
A proof of concept for information extraction from camera feed using image analysis techniques
NOAH LANAI
Department of Mathematical Sciences
Chalmers University of Technology

Abstract

This thesis presented a collaborative effort between Ovako and Swerim to create a system to assist operators decision-making during the degassing stage of Ovako's steel production, using video data obtained from the degassing camera feed. The purpose of the project was to explore all potential information that could be extracted from the camera feed. This would act as proof-of-concept for such a system.

The research goal were accomplished by employing image analysis techniques and statistical methods to collect as much relevant information as possible. Traditional image analysis techniques were utilized for segmentation and feature extraction, which were combined with statistical methods to estimate key degassing properties such as fluid height/volume, fluid movement, and proportions of steel/slag in the degassing ladle.

The results demonstrated a functional system which processed saved video data, providing estimated statistics at a chosen frame rate and producing relevant plots showing the segmentation and statistics. Findings demonstrated ample opportunities for extracting valuable information from the degassing images that exceeded initial expectations set forth by engineers at Ovako and Swerim. Additionally, a correlation was seen between extracted statistics from the system and actual video footage, suggesting the reliability of the system.

This thesis demonstrated the potential of creating a comprehensive system to aid the operators during degassing. It served as a starting point for further discussions among the companies regarding refinements to align more closely with operator needs. More work must be done towards realising the goal, yet this thesis provided proof of concept that the development of such a system is feasible.

Keywords: degassing, image processing, steel production, spatial statistics, video analysis.

Acknowledgements

This thesis was in collaboration with Ovako and Swerim, where the majority of the work was conducted remotely. Nevertheless, I got the chance to visit both sites several times which was an enjoyable experience.

I would like to thank my supervisors at both of the companies, Krister Ekström and Marcus Svadling, for their continous support on the project and their genuine interest in my findings.

Furthermore, I would like to thank my supervisor and examiner at Chalmers University of Technology, Professor Aila Särkkä, for her guidance during the course of the project.

Lastly, I would like to extend my gratitude to all the other individuals who were involved in this project, including those at both companies, the university, and within my personal circle.

Noah Lanai, Gothenburg, June 2023

Contents

List of Figures

List of Tables

1
Introduction

1.1 Background

To take advantage of modern image analysis tools, Ovako and Swerim have undertaken a master's thesis project to apply such tools to video footage of an important step in steel production: degassing. This step is essential in steel-making as it ensures the correct chemical properties of the steel. However, determining when the degassing process has been completed can be a quite complex task for operators.

By applying image analysis and computer vision techniques, the goal is to develop a tool that can assist operators in making more precise and efficient decisions. If successful, such an implementation could translate into significant cost-savings for Ovako.

1.2 Limitations

This study seeks to create a system that can assist operators in decision-making based on camera feed of Ovako's steel production's degassing process. However, it should be recognized that developing such an system takes time and this thesis serves as just the start. The focus will be on learning what information could be extracted from the images of the camera feed. This will serve as a proof-of-concept that the development of such a system is feasible. Because of this, the system may not be able to realize its full potential just yet.

Another limitation of this study is due to the resources and constraints of its research environment. For instance, there may be restrictions on what data can be collected from cameras in operation, making it difficult to develop robust models. Furthermore, Ovako's steel production environment may not be generalizable to other steel productions or industries.

1.3 Research Question

This thesis seeks to explore the potential of modern image analysis tools in extracting valuable information from Ovako's degassing process camera feed. The objective

is to empower operators with some knowledge that can help them make better decisions. The research question at hand is:

- How can modern image analysis tools be employed to extract useful data from the camera feed of a degassing process?

This study will identify and define key features in the process, such as fluid level in the tank, steel pump performance, and steel/slag movement. These features will then be analyzed using computer vision techniques to extract relevant parameters which will provide quantified outcomes. Furthermore, the accuracy and feasibility of the extracted features will be assessed to determine whether modern image analysis tools could be beneficial when used on camera feed.

2

Data Acquisition and Processing

2.1 Degassing

Degassing in steel production is an essential step to produce high-quality steel. The primary goal of degassing is to eliminate dissolved gases from molten steel, which can cause various problems during solidification and subsequent processing such as cracking, porosity and reduced mechanical properties. The overall goal is for the steel to have the correct chemical properties.

Degassing can be done through several methods and Ovako's steel production uses vacuum degassing. With vacuum degassing, a vacuum is created over the surface of molten steel to reduce pressure and allow dissolved gases to boil away and escape from within the steel.

Slag is the byproduct of the steelmaking process that forms when impurities are removed from molten steel. It consists of metal oxides formed when impurities react with the added materials during removal. Slag has a lower density than molten steel, so it floats on the top.

Maintaining the correct pressure inside a degassing tank is one of the most critical elements for the success of the process. To guarantee efficient removal of dissolved gases from molten steel, the pressure inside should be carefully controlled; otherwise, lower quality steel with reduced mechanical properties could result. Conversely, low pressure may cause excessive boiling and lead to problems like casting itself with the lid of the tank.

If the steel boils too much and casts itself with the lid of the tank, the consequences can be dire. Molten steel can damage degassing equipment and even injure the operators.

2.2 Camera Feed

The camera is mounted on the wall of the tank and features a cooling lens to withstand its high temperature inside. It provides a view of the molten steel within the tank, producing 25 frames per second with 1080 x 1920 pixel images. The images have 255 x 3 RGB (red, green and blue) color channels. During degassing, its angle remains fixed and the date/timestamp information can be seen in the lower left corner.

However, the vision captured by the camera lens is limited and partially obscured. Only those areas within its scope are relevant for the analysis, while slag particles accumulate on the lens during the degassing process which causes dead spots or inaccessible areas that cannot be assessed. It must be taken into account that the degree of obstruction caused by slag deposits varies over time, so each frame must account for them accordingly.

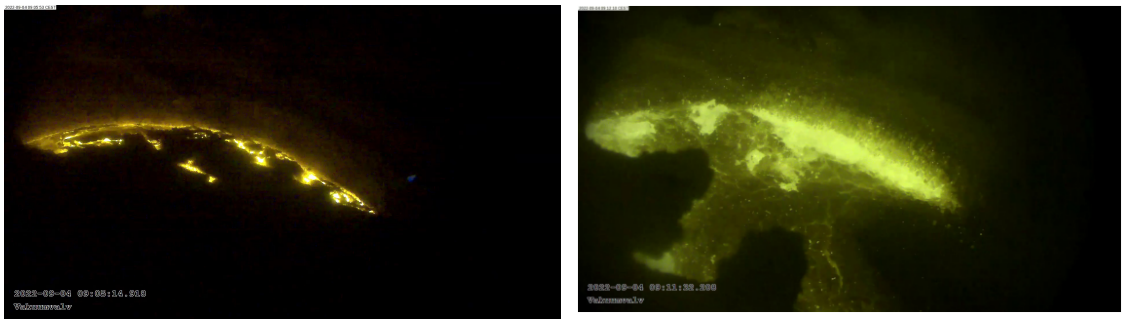

Figure 2.1: Example of a degassing process. The left image shows the early step of the process, where the steel appears as lid up areas and the black mass on top is slag; roundness of the light shows the tank geometry. The right image shows the process where boiling have commenced. The brightest areas indicate emerging steel from the two pumps in the left and right while the slightly darker areas indicate melted slag while the floating part remains unmelted.

At Ovako's production site, the degassing process typically takes around 20 minutes to complete successfully. To accurately analyze the images produced during this step, it's essential to comprehend its various stages. As visual appearance varies significantly throughout this step, different feature extraction and segmentation methods and parameters must be employed in order to capture relevant data. Figure 2.1 shows two examples of a degassing process at Ovako's steel production.

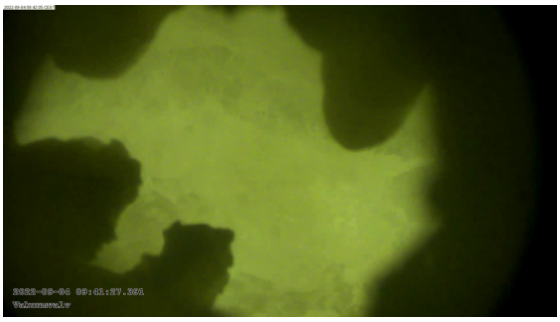

Figure 2.2: Example of a full boiling degassing process. The dark appearing areas are slag that has been tossed at the lens during the process. The green area in the middle shows the moving mass of steel and melted slag.

Once all the slag has melted, the fluid becomes turbulent and homogeneous - making it difficult to distinguish steel from slag as shown in Figure 2.2. Therefore, operators

must closely monitor the fluid behavior and keep track of the boiling level throughout this step.

2.3 Experimental System

The analysis was performed using Python version 3 along with NumPy, SciPy, Matplotlib, Skimage, Scikit-learn and OpenCV libraries. Python proved ideal for the image analysis due to its simplicity of use and access to numerous libraries supporting image processing. NumPy library provided numerical computing with efficient array objects for handling multidimensional arrays needed in the image processing while SciPy provided signal processing functions which helped with noise reduction, filtering and enhancement tasks.

Matplotlib was utilized for data visualization and plotting results. This library offers an easy, intuitive way to create plots, charts, and other visualizations that effectively convey the findings from an analysis. Furthermore, ipynb format was utilized in certain examples which enabled easy sharing of interactive notebooks that demonstrate both the results and methodology.

Skimage and OpenCV libraries were used for image processing and computer vision tasks such as segmentation, edge detection, and feature extraction. Furthermore, Scikit-learn was utilized for data preprocessing.

Python was utilized to turn each frame of a video into an object with functions and properties, making it simpler to apply different methods to each frame object and compare results between various image processing techniques.

Overall, the combination of Python and relevant libraries used in this study provided a robust and flexible platform for image analysis and computer vision tasks. Using Python in combination with these libraries allowed for efficient, reproducible research, while ipynb format made the visualization and interpretation of the results much simpler.

2.3.1 Visualization

To visualize the results of the image processing and analysis, a large plot is generated for each frame using Matplotlib library in Python. This plot includes both the original image as well as any manipulated images used to produce statistics. The goal of these statistics is to give quantified results of the interesting features of the degassing process. As explained in Section 1.3, these features are fluid level in the tank, movement of the fluid and steel/slag ratio. By bringing all this information together in one visualization for each frame, one can identify patterns or trends in the degassing process movement.

Once individual large plots have been generated for each frame, a GIF (Graphics Interchange Format) is produced that compiles all these extracted visualizations.

This facilitates an easy viewing of statistics at each frame in the video and allows some conclusions to be drawn based on the analysis.

2.3.2 Management of Results

Visualizing the results of the image processing is essential, but it's also necessary to export the data into an Excel file in order to save each statistic computed from each frame. By exporting this data, additional insights can be drawn from the analysis. The data exported can be plotted and graphed to enable further analysis and identification of trends or patterns in the process. If there is a correlation between what can be seen in the video and the extracted statistics, then this confirms that the analysis was somewhat correct; furthermore, this step provides additional assurance in its accuracy.

Exporting the data to an Excel file facilitates sharing and reusing for future studies or analyses. With these saved records, other researchers can easily access and manipulate the data, building upon what has been presented here and further developing the system.

3

Image Analysis Techniques

3.1 Preprocessing

3.1.1 Image Conversion

Image conversion is an essential preprocessing step in many image analysis tasks, as it allows for standardizing the input data and emphasizing specific features of interest. This section will discuss various image conversion techniques such as color space transformation, image resizing, and normalization that can be employed to enhance the performance of segmentation algorithms when studying degassing processes.

Images are typically represented in RGB (Red, Green, Blue) color space, where each pixel's hue is composed of these three primary colors. However, certain image analysis tasks may benefit from converting the input image to another color space such as grayscale or HSV (Hues, Saturation, Value). For instance, converting an RGB image to grayscale helps to focus on the intensity information which could be particularly helpful for segmenting the degassing process components. To convert an RGB image to a grayscale image is done by applying the following formula for each pixel:

$$I = 0.299R + 0.587G + 0.114B \tag{3.1}$$

where I represents the intensity value and R, G and B denote the red, green and blue channel values respectively.

Image resizing can be used to standardize the input image dimensions, making it compatible with various image analysis algorithms. Furthermore, shrinking images reduces computational complexity and memory requirements while still retaining the essential information necessary for segmentation.

Normalizing pixel intensity values can enhance the performance of image segmentation algorithms by scaling the input data to a common range. Normalization also reduces the impact of illumination variations, increasing robustness of the segmentation process. There are various approaches for normalizing image intensity values, such as min-max scaling and zero-mean normalization. Min-max scaling maps pixel intensities into an exact range [0, 1], usually by applying the following formula:

$$I_{norm} = \frac{I - I_{min}}{I_{max} - I_{min}} \tag{3.2}$$

where I_{norm} is the normalized intensity value, while I_{min} and I_{max} represent the minimum and maximum intensity values in the image respectively. Zero-mean normalization on the other hand centers the pixel intensity values around zero, giving them a mean value of zero and standard deviation of one:

$$I_{norm} = \frac{I - \mu}{\sigma} \tag{3.3}$$

where μ and σ represent the mean and standard deviation of the pixel intensities respectively. By incorporating these image conversion techniques into the preprocessing pipeline, the quality and compatibility of the input images can be improved, leading to a more precise segmentation for the analysis of the degassing processes.

3.1.2 Noise Reduction

Filtering is a critical pre-processing step in image analysis, as it helps remove noise, create smooth images and highlight some specific features - ultimately improving the performance of the segmentation algorithms. In this section, we review three filtering techniques that can be employed before applying image segmentation methods: Gaussian filter, bilateral filter and non-local means filter. Each filter offers its own advantages: the Gaussian filter provides general smoothing and noise reduction, the bilateral filter focuses on edge preservation, while non-local means filters excel at preserving textures and fine details.

The Gaussian filter is a linear filter commonly used to smooth images and reduce noise. It works by convolving an image with a Gaussian kernel, which has standard deviation (or scale) σ. As pixels move away from the center of the kernel, their values decrease, giving more weight to the nearby pixels in terms of output than to the distant ones. A 2D Gaussian kernel can be represented as:

$$G(x, y) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right) \tag{3.4}$$

Applying the Gaussian filter to an image helps to preserve object boundaries while smoothing it and reducing noise, making it a suitable preprocessing step for image segmentation tasks (Cabello, Leon, Iano, Arthur, 2015).

The bilateral filter is a nonlinear, edge-preserving filter that smooths images while maintaining edges by taking into account both spatial and intensity information. This works by combining a Gaussian filter in the spatial domain (for smoothing) with another Gaussian filter in the intensity domain (to preserve edges). The bilateral filter can be expressed as:

$$BF(I)(x, y) = \frac{1}{W(x, y)} \sum_{x'=-k}^{k} \sum_{y'=-k}^{k} G_s(x', y') G_r(I(x+x', y+y') - I(x, y)) I(x+x', y+y') \tag{3.5}$$

where $BF(I)(x, y)$ is the filtered output at pixel location (x, y), $W(x, y)$ a the normalization factor, G_s is a spatial Gaussian kernel and G_r an intensity Gaussian kernel; $I(x, y)$ is the intensity values and k its kernel size. Through bilateral filtering techniques such as these it is possible to effectively reduce noise while maintaining the edges - making it particularly helpful for image segmentation tasks requiring accurate boundary detection tasks (Zhang, & Allebach 2008).

The non-local means filter is an advanced non-linear filter designed to reduce noise while maintaining image structures by taking advantage of natural redundancy in images. It works by comparing similarities between patches of pixels instead of individual pixels, leading to better noise reduction and texture preservation. The non-local means filter can be described as:

$$NLM(I)(x, y) = \frac{1}{Z(x, y)} \sum_{x'=-k}^{k} \sum_{y'=-k}^{k} w(x, y, x', y') I(x + x', y + y') \tag{3.6}$$

where $NLM(I)(x, y)$ is the filtered output at the pixel location (x, y), $Z(x, y)$ is a normalization factor, $w(x, y, x', y')$ is a weight function that measures the similarity between the patches centered at (x, y) and $(x + x', y + y')$, $I(x, y)$ is the intensity values, and k is the kernel size. The non-local means filter provides excellent noise reduction while maintaining the image details, making it a suitable pre-processing step for image segmentation tasks that require high-quality feature preservation. (Buades, Coll, & Morel, 2011)

3.2 Feature extraction with Hough Transforms

Hough Transforms are a feature extraction technique commonly employed in image processing and computer vision to detect shapes, lines, and circles. Hough Transforms work by translating image coordinates into a parameter space. At the start of the degassing process, the base fluid level for melted steel typically takes on a circular shape which can be observed and used as a reference point. Using the Hough transform method to detect circles in the image may help to locate this level. To detect circles in the image use the following formula to convert the image space points to circles within the parameter space:

$$(x - a)^2 + (y - b)^2 = r^2$$

where (a, b) are the center coordinates of the circle and r is the radius of the circle. (Hassanein, Mohammad, Sameer, & Ragab, 2015)

3.3 Segmentation

3.3.1 Otsu's Threshold

Otsu's thresholding is a widely used image segmentation technique for distinguishing the foreground from the background pixels in an image. The objective of Otsu's

thresholding is to find a threshold value that minimizes the intra-class variance (difference between the foreground and background pixels) between them. This variance can be described as follows:

$$\sigma_w^2(t) = \omega_0(t)\sigma_0^2(t) + \omega_1(t)\sigma_1^2(t) \tag{3.7}$$

where $\omega_0(t)$ and $\omega_1(t)$ are the probabilities of the background and foreground pixels, respectively; and $\sigma_0^2(t)$ and $\sigma_1^2(t)$ are their variances in intensity values respectively. Otsu's thresholding algorithm computes an intra-class variance for all possible threshold values before selecting one that minimizes $\sigma_w^2(t)$. The optimal value t^*, can be computed as:

$$t^* = \arg\min_t \sigma_w^2(t) \tag{3.8}$$

Otsu's thresholding can be used as a pre-processing step to isolate interesting areas without blocking slag. (Makkar, & Pundir, 2014)

3.3.2 K-means Algorithm

The k-means algorithm is an unsupervised clustering technique that divides data points into k distinct clusters based on their similarity. When applied to image analysis, this algorithm takes into account the pixel intensity values from preprocessed images. According to Hastie, Tibshirani & Friedman (2017), the steps for the k-means algorithm are as follows:

Start by selecting $k \geq 2$ initial centroids randomly from the set of pixel intensities. Let C_i represent the centroid for cluster i, where $i \in 2, 3, ..., k$. Assign each pixel I_j to the cluster corresponding to the centroid with the smallest Euclidean distance to the pixel intensity value. This can be mathematically represented as:

$$z_j = \arg\min_i \|I_j - C_i\|^2 \tag{3.9}$$

where z_j denotes the cluster index assigned to pixel I_j. The vertical bars "$\|$" represents the Euclidean norm. This indicates that we are interested in the magnitude of the vector inside the vertical bars. The next step is to calculate the new centroids as the mean intensity value of all pixels belonging to the respective clusters. This can be expressed as:

$$C_i = \frac{\sum_{I_j \in S_i} I_j}{S_i} \tag{3.10}$$

where S_i denotes the set of all pixels assigned to cluster i, and S_i is the number of pixels in cluster i. Repeat until the centroids' positions become stable, or a predetermined number of iterations have been reached. Convergence can be assessed by monitoring the change in centroids' positions, as follows:

$$\sum_{i=1}^{k} \|C_i^{(t+1)} - C_i^{(t)}\|^2 < \epsilon \tag{3.11}$$

where $C_i^{(t)}$ and $C_i^{(t+1)}$ are the centroid positions at iteration t and $t+1$, respectively, and ϵ is a small positive value representing the convergence threshold. After convergence, create the segmented image by assigning each pixel the intensity value of its corresponding centroid.

By grouping pixels with similar intensity values, distinct regions in the image that correspond to the different components of the degassing process, such as steel, slag and wall of ladle can be effectively highlighted. Applying the k-means algorithm for image analysis enables effective segmentation of the essential parts of an image for better understanding of the degassing process.

3.3.3 Gaussian Mixture Model

The Gaussian Mixture Model (GMM) is a generative probabilistic model that represents the mixture of multiple Gaussian distributions. When applied to image segmentation, GMM helps to model the distribution of the pixel intensities within an image. According to Hastie, Tibshirani & Friedman (2017), the steps for applying GMM to image segmentation are outlined below:

Start by determining the number of Gaussian components K for the mixture model, which corresponds to the number of distinct intensity groups in the image. This value should be selected based on domain knowledge and experimental evaluation. Initialize the parameters of the Gaussian components. For each component k, where $k \in 1, 2, ..., K$, estimate the initial mean μ_k, covariance matrix Σ_k, and mixture weight π_k by using the expectation–maximization (EM) algorithm. The EM algorithm is used to estimate the parameters of the GMM by iteratively updating the means, covariances, and weights of the Gaussian distributions. In the expectation step, the probability of each pixel belonging to each Gaussian component is computed. In the maximization step, the parameters of the Gaussian distributions are updated based on the pixel probabilities. These are repeated until the estimated parameters converges.

Once the GMM parameters have converged, the image pixels are assigned labels based on the Gaussian component with the highest probability. This results in a segmentation of the image into different regions, each corresponding to a different Gaussian distribution.

This assignment effectively segments the image, grouping pixels with similar intensity values together. This segmented image clearly distinguishes distinct regions within it that correspond to different aspects of degassing such as steel, slag and the wall of the ladle.

3.4 Video Analysis

This section introduces a time-based dimension to the analysis, allowing the comparison of frames over time to gain valuable insights into features such as fluid

movement. By comparing each frame with another, we can compute the difference and derive instantaneous movement between frames by analyzing the pixel intensities. Furthermore, computing an average movement over larger intervals - say 25 frames equivalent to one second of the camera feed - could also provide useful data.

3.4.1 Difference Between Frames

The method involves analyzing the difference between consecutive frames in the video feed to quantify the degree of movement or change that occurred over time. The difference between two frames, I_i and I_{i-1}, was calculated using the formula:

$$D_i = |I_i - I_{i-1}|$$

where D_i represents the difference image between frames i and $i-1$. (Singla, 2014)

The difference images were then analyzed using various statistical measures, such as the mean and standard deviation to quantify the degree of movement or change that occurred between the frames. The mean difference between two frames was computed as:

$$\mu_D = \frac{1}{N} \sum_{i=1}^{N} D_i$$

where μ_D represents the mean difference and N is the total number of pixels in the difference image. This measure provides an estimate of the average movement or change that occurred between the frames.

3.4.2 Moving Average

To gain insights into a longer time interval, it is beneficial to analyze the average movement over an extended duration. This can be achieved using the sliding window method which computes an average over predefined frames, effectively capturing any underlying trends in the data. Hirzel, Schneider, & Tangwongsan (2017) provides an informative tutorial on different sliding window algorithms; for looking at simple moving averages we will use the "first in, first out" algorithm with one second intervals representing 25 frames each. The sliding window method involves following these steps:

Start by defining a window W of size 25 corresponding to the number of frames in the interval (in this case, $W = 25$). Calculate the differences between consecutive frames according to the steps in section 3.4.1 and store them in an array of length W.

$$A_1 = \frac{1}{25} \sum_{i=1}^{25} W_{1,i} \tag{3.12}$$

where A_1 denotes the average movement over the first 25 frames. Iterate through the remaining frames in the video, updating the sliding window and computing the

moving average at each time step $t > 25$. For each subsequent frame, perform the following operations:

a. Remove the first element of the array W, shifting all elements one position to the left.

b. Calculate the difference between the current frame and the previous one, and append the result to the end of the array W.

c. Update the moving average for the current time step t:

$$A_t = \frac{1}{25} \sum_{i=1}^{25} W_{t,i} \tag{3.13}$$

By applying this technique to the video data, we are able to effectively capture underlying trends and patterns which help to deepen the comprehension of the degassing process dynamics.

3.5 Statistical Analysis

3.5.1 Quantifying Class Proportions

After segmenting an image, it is possible to quantify the proportions of different classes such as slag, steel and the wall of a ladle during degassing. Calculating class ratios allows analysis of the relative distributions within each image and gain valuable insights into the process dynamics and fluid levels within the ladle. This section will cover the methodology for calculating class ratios as well as why having accurate ladle geometry as a ground truth for comparison is so important. To compute class ratios after image segmentation, follow these steps:

Assign unique pixel values to each class, allowing for quick identification and quantification of the segmented regions. Calculate the total number of pixels belonging to each class by counting those with their corresponding unique values in the segmented image. Divide this total number of pixels per overall number of pixels in the image, yielding the proportion of each class within it:

$$R_i = \frac{N_i}{\sum_{j=1}^{C} N_j} \tag{3.14}$$

where R_i is the ratio of class i, N_i is the number of pixels in class i, and C is the total number of classes. Class ratios can be used to estimate fluid levels within the tank by comparing pixels classified as fluid to the overall number of pixels of the vessel. To get precise readings, however, it's necessary to define the geometry of the tank first - such as its dimensions and orientation within an image - in order to compare against it accurately. Any features or boundaries that must be taken into account during the analysis should also be defined beforehand.

In the degassing process, slag can sometimes obstruct the field of view, leading to inaccurate calculations and segmentations. To account for this uncertainty and improve the class ratio computations, it is necessary to incorporate a tolerance measure that takes into account the fluid region proximity to any obstructions. One approach to estimate this tolerance is applying edge detection on segmented images in order to recognize boundaries between the classes and calculate the distance between them and any obstructions. When computing the tolerance using edge detection and distance calculations, follow these steps:

Start by applying an edge detection algorithm to the segmented image in order to recognize boundaries between different classes, such as fluid regions and obstructing slag. Then Calculate the minimum distance to the nearest slag boundary for each fluid region. This can be done by using distance transform methods or by computing Euclidean distances between pixels on the edge of the fluid region and its slag boundary. After that define a threshold distance to categorize fluid regions as either near or far from the obstructing slag. Fluid regions closer to the boundary are considered more susceptible to being affected by obstruction and should thus contribute more towards the tolerance levels. Finally add the tolerance measure to class ratio calculations, and adjust each class proportion according to how close the fluid regions are to the obstructing slag.

By including tolerance measures into class ratio computations, the accuracy and reliability of results can be increased. This information can then be used for further analysis, process optimization, and decision-making related to the degassing process.

3.5.2 Entropy

Image entropy is a statistical measure of the randomness, complexity or disorder present in an image. A study by Pal, N. R. and Pal, S. K. (1991) provides an introduction to grayscale image entropy definition and then proposes four different algorithms for various applications within this realm of analysis. Entropy measures the degree of information or uncertainty present in a pixel intensity distribution, making it a useful metric for describing structural properties of an image. To calculate entropy from an image, take these steps:

Calculate the histogram of an image to show the distribution of pixel intensities. Normalize this histogram to create a probability distribution for pixel intensities:

$$p(i) = \frac{h(i)}{\sum_{j=0}^{255} h(j)} \tag{3.15}$$

where $p(i)$ is the probability of intensity i and $h(i)$ represents the count of intensity i in the histogram. After that calculate the entropy of the image using the normalized histogram:

$$E = -\sum_{i=0}^{255} p(i) \log_2(p(i) + \epsilon) \tag{3.16}$$

where E is the image entropy and ϵ is a small positive constant (e.g., 10^{-7}) added to prevent undefined logarithm values for zero probabilities.

Image entropy can be an effective measure of complexity of the degassing process. High values indicate greater randomness or disorder in the pixel intensity distribution, suggesting a more dynamic or active degassing process; on the other hand, low values signify uniformity or less complexity.

3.5.3 Energy

Image energy is a measure of intensity variation within an image, which can provide insights into the activity or structural complexity. In the degassing process, this quantifies overall intensity changes and fluctuations within a frame, allowing us to identify specific time intervals or regions with significant variations. In this section, we'll describe how to compute the energy of an image and discuss its implications for understanding the degassing process. To compute an image's energy, follow these steps:

Calculate the gradient of an image, which represents the spatial rate of change in pixel intensity. You can do this using various methods such as the Sobel operator or Scharr operator. These are Kernel windows that are applied to each pixel which outputs two gradient values, one for vertical and one for horizontal direction (Levkine, n.d.). Compute the gradient magnitude for each pixel:

$$M(x, y) = \sqrt{S_x(x, y)^2 + S_y(x, y)^2} \tag{3.17}$$

where $M(x, y)$ is the gradient magnitude at pixel location (x, y), and $S_x(x, y)$ and $S_y(x, y)$ are the horizontal and vertical gradient components computed with a Sobel operator or Scharr operator. Calculate the energy of the image as the sum of the squared gradient magnitudes:

$$E = \sum_{x=0}^{W-1} \sum_{y=0}^{H-1} M(x, y)^2 \tag{3.18}$$

E is the image energy, while W and H represent the image width and height in pixels. The image energy can serve as a useful metric to quantify intensity variations during degassing in steel production; higher values indicate greater activity or intensity fluctuations. Conversely, lower energies signify more stability or homogeneous processes with fewer intensity variations.

3.5.4 Skewness and Kurtosis

Skewness and kurtosis are statistical measures that describe the shape of the distribution. These metrics can provide valuable insights into the overall structure, asymmetry, and tail behavior of an intensity distribution - characteristics useful for the degassing processes. This section describes methods for estimating kurtosis and skewness for images as well as discuss their implications when analyzing this degassing process.

Skewness is a statistic that measures the degree of asymmetry or deviation from symmetry of a distribution. A positive value indicates a longer right tail, while negative values indicate an extended left tail. To calculate skewness for an image, follow these steps:

Start by calculating the mean μ and standard deviation σ of the pixel intensities in the image. Compute the skewness as an estimate of the third standardized moment:

$$\beta_1 = \frac{1}{N} \sum_{i=1}^{N} \left(\frac{I_i - \mu}{\sigma} \right)^3 \tag{3.19}$$

where β_1 is the skewness, N is the total number of pixels in the image and I_i represents the intensity of the i-th pixel.

Kurtosis measures the "tailedness" or "peakiness" of a distribution. To calculate kurtosis for an image, take these steps:

Start by calculating the mean μ and standard deviation σ of the pixel intensities in the image. Compute the kurtosis as the fourth standardized moment:

$$\beta_2 = \frac{1}{N} \sum_{i=1}^{N} \left(\frac{I_i - \mu}{\sigma} \right)^4 \tag{3.20}$$

where β_2 is the kurtosis, N is the total number of pixels in the image and I_i represents the intensity of the i-th pixel. Higher values ($\beta_2 > 3$) indicate more peaked distribution with heavy tails, while lower values ($\beta_2 < 3$) suggest flatter distribution with light tails. (Sagitov, 2022)

Kurtosis and skewness, when applied to the degassing process, provide valuable insight into the distribution of the pixel intensities and the image structure. High kurtosis values may indicate more extreme intensity values or localized regions of interest while low kurtosis values could indicate homogeneous or uniform distribution patterns. Skewness also helps to detect any asymmetry in the intensity distribution.

4
Results

4.1 Visualization of images

Analyzing an image from the middle of a process is particularly instructive for analysts, since this stage involves the operators making crucial decisions. Figure 4.1 is an example image taken in the full boiling step of a degassing process. Despite the obstructions caused by a considerable amount of slag, it is still difficult to make sense of the image. This image depicts a boiling mixture of slag and steel, making it difficult to differentiate between. Lighter regions can be identified at the center, while darker patches appear above it - signifying the tank's wall.

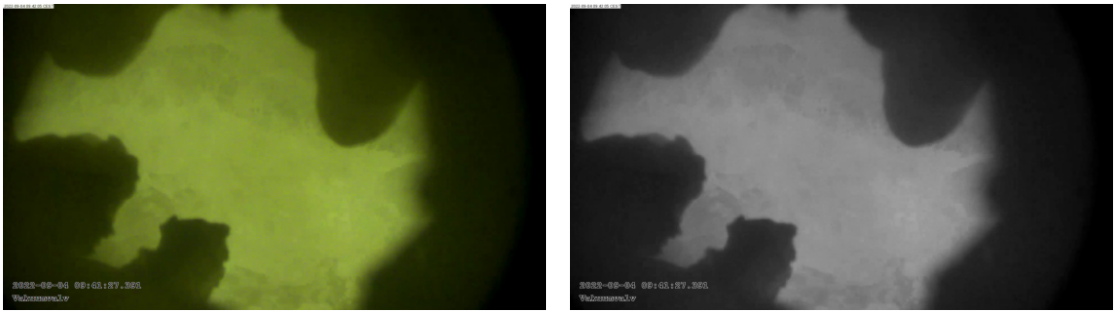

Figure 4.1: Visualization of an image. The right image shows the true color and the left image is a grayscale conversion. The dark parts in both of the images are the obstructing slag on the camera lens. The brighter parts shows the moving mass inside the degassing ladel.

Figure 4.2 presents the image visualized in various color spaces. It is evident that these images contain a wealth of useful information, particularly the value plot which highlights multiple distinct regions within the image that can be segmented.

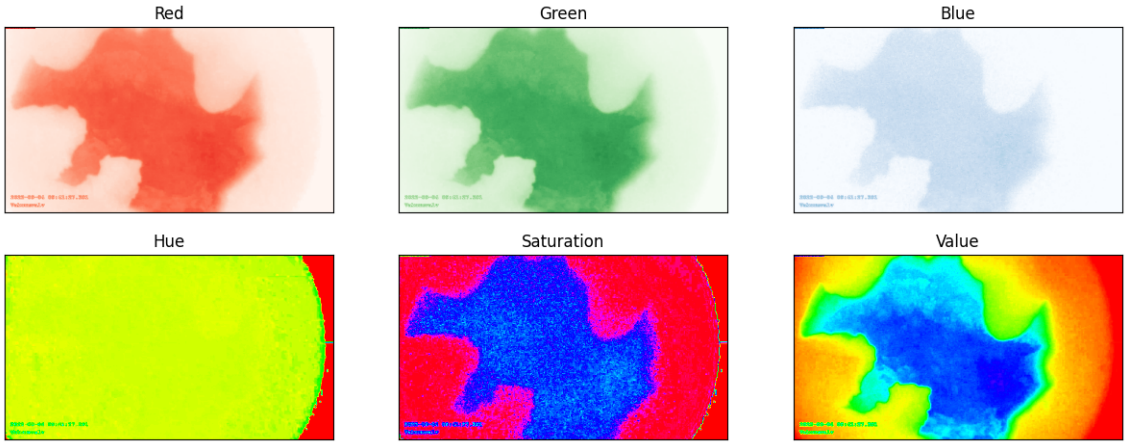

Figure 4.2: Different color scales of an image of a degassing process.

Histograms of the pixel values in the image provide valuable insights into their distribution. As shown in Figure 4.3, both the RGB color scale and grayscale images have distinct pixel values displayed. It is evident that all plots contain an outlier at 0, which could be due to the camera lens appearing completely dark. In order to gain a more precise insight into the distributions, these outliers need to be removed from the graphs. The plots demonstrate that blue color does not exhibit a distinct distribution, while red and green values show two distinct patterns which can also be observed in the grayscale image. Therefore, we can proceed the analysis with only using grayscale images.

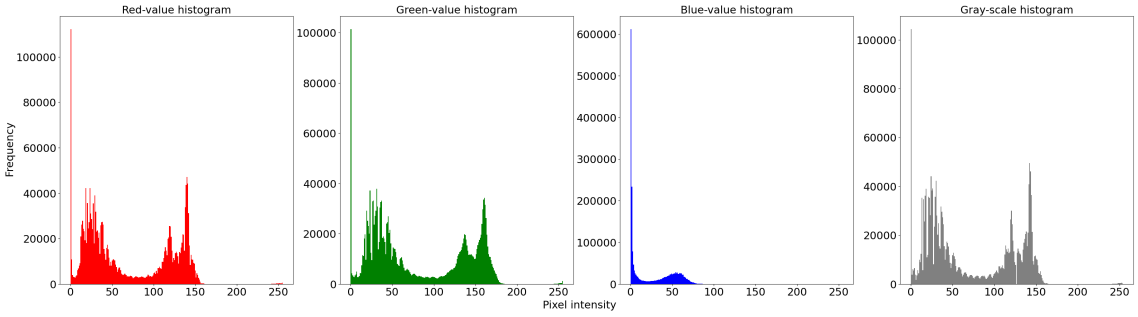

Figure 4.3: Histogram of pixel intensities for the different color scales of an image in the degassing process.

One essential part of the image that must be eliminated is the text present on the top left and bottom left because it contributes pixel values to the distributions in these regions, potentially creating problems during segmentation.

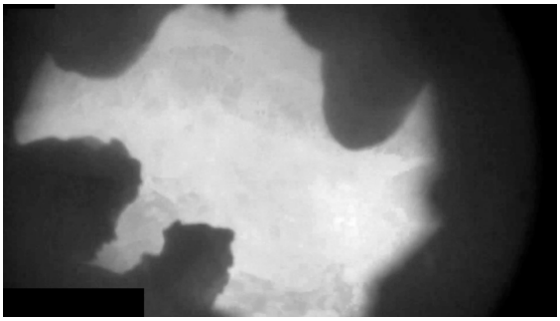

Figure 4.4: An image of the degassing process in grayscale with text information removed.

4.2 Definition of tank geometry

Manually segmenting the camera lens image involves drawing a circle around its darkest parts, as shown in red color in Figure 4.5. To improve this process, one should have knowledge on both the camera lens physical location and size as well as how it appears represented in pixels.

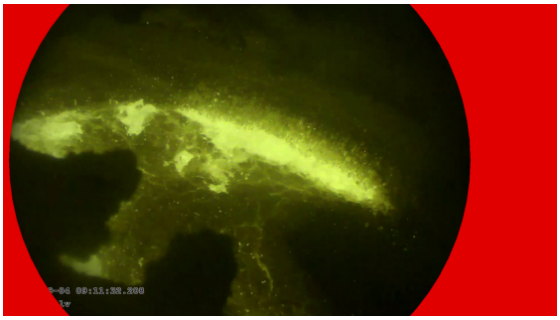

Figure 4.5: An image of the degassing process with the camera lens segmented in red color.

To determine the geometry of the tank, two levels must be defined: a lower bound which represents the base level where the process starts and an upper bound which corresponds to where the lid of the container rests. As illustrated in Figure 4.6, two circles have been drawn manually to represent these boundaries by visually fitting them to their respective positions.

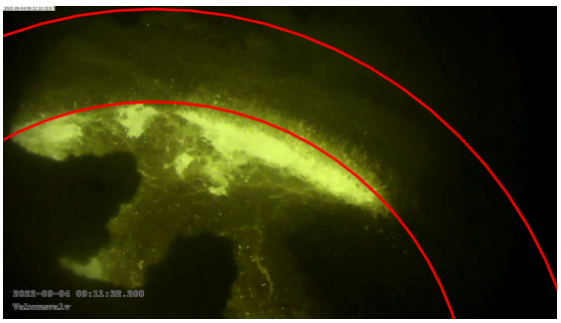

Figure 4.6: An image of a degassing process with two red lines. The upper line representing the lid of the ladel and the lower line representing the base level of mass.

To locate the lower boundary, we apply Hough transforms as described in Section 3.2. A function has been developed to identify the largest circle in the image by examining an image at the start of degassing. The algorithm works as follows: a radius is manually specified, and the Hough transform searches for a circle of that size in the image. It then automatically finds a fit for the base level of the mass in the tank. An ellipse can be drawn and rotated with the camera's angle using the coordinates of the circle center and fixed radius, producing the left image shown in Figure 4.7. To obtain more precise results, knowledge of the true camera angle would be beneficial.

The next step is to determine the upper boundary, or the lid of the ladel, which is easier to identify. Since the camera angle remains constant throughout each process, this information can be hardcoded. By looking at a frame where the lid can clearly be seen, an upper ellipse can be defined as shown in the right image of Figure 4.7.

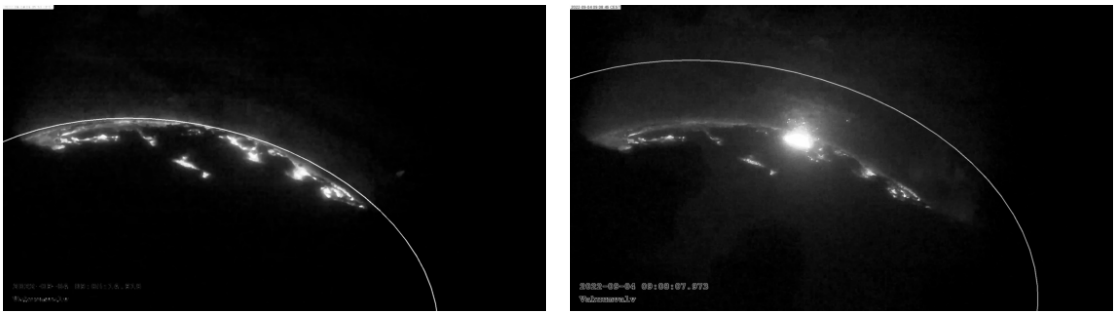

Figure 4.7: Two images of a degassing process in grayscale. The image to the left has an ellipse that is automatically drawn with Hough Transforms, that represent the lower bound. The image to the right has an ellipse that is manually drawn that represents the lid of the ladel.

4.3 Experimenting with segmentation

In this section, we will discuss the results of our segmentation methods. We have experimented with three different techniques: the Otsu threshold, the k-means al-

gorithm, and the Gaussian mixture model. Each of these methods is explained in detail in section 3.3. The ultimate objective of this study is to identify the most effective segmentation method for different types of images, which can then be utilized for further analysis.

Given the presence of two clear peaks in the histogram of a grayscale image, segmenting can be done efficiently by applying Otsu's threshold, as previously specified in Section 3.3.1. Figure 4.8 illustrates this binary image clearly by highlighting both the dark camera lens and any slag that was thrown at it during processing as they exhibit similar low intensities. This provides an effective means to remove irrelevant parts of the picture.

Figure 4.8: Binary filter extracted by using Otsu's threshold on an image of the degassing process. The black parts are obstructions to the camera vision. The white parts are the areas where the camera captures useful information.

Figure 4.9 displays the image with the binary filter applied. As a result, all the irrelevant values in the image now have the value 0.

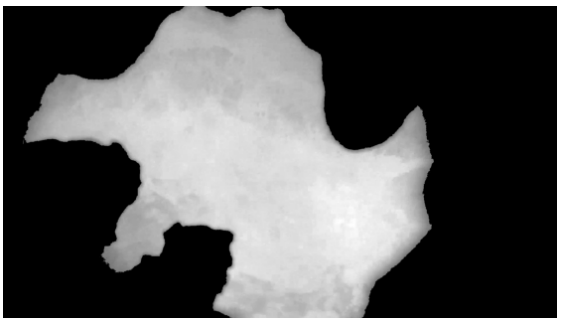

Figure 4.9: An image of a degassing process with a binary filter applied to segment the interesting parts from obstructions to the camera lens. The lid up area shows the moving mass in the degassing ladel.

Multi-Otsu's thresholding is used to segment the image, with the aim of isolating the interesting part from the rest. Unfortunately, when applied directly onto the original image, it fails to segment out this interesting area and instead places its threshold somewhere near the middle of the dark pixel value distribution - leading to an ineffective segmentation that looks strange. Figure 4.10 illustrates this histogram,

thresholds and result of the segmentation. To improve this segmentation, first use simple Otsu's thresholding to extract the interesting parts as shown in Figure 4.9 and then use that result as input for the multi-Otsu's thresholding.

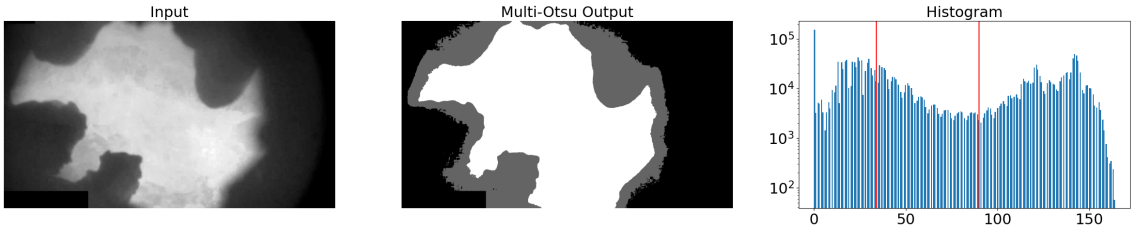

Figure 4.10: Behaviour of the multi-Otsu's thresholding algorithm applied to an image of the degassing process. The input is an original grayscale image and the red lines in the histogram show the thresholds for the segmentation.

However, when a binary filter is first applied to remove the non-interesting parts of an image and this resulting image is used as input for multi-Otsu's thresholding, better thresholds are obtained. The first threshold is placed at 0, separating the non-interesting values into one class while the second threshold divides both distributions apart creating two classes of interesting values. As shown in Figure 4.11, segmentation shows moving masses of steel and slag as white pixels - clearly isolating them from other parts of the scene.

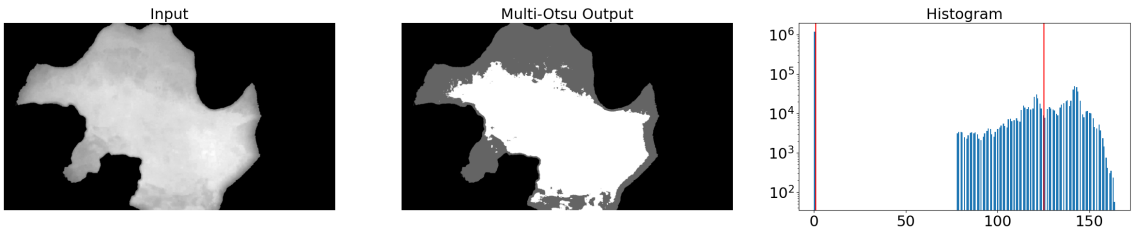

Figure 4.11: Behaviour of the multi-Otsu's thresholding algorithm applied to an image of the degassing process. The input is a binary filtered image and the red lines in the histogram show the thresholds for the segmentation.

Figure 4.12 displays the image segmented using the k-means algorithm with k = 3. The outcome is very similar to that obtained through multi-Otsu thresholding.

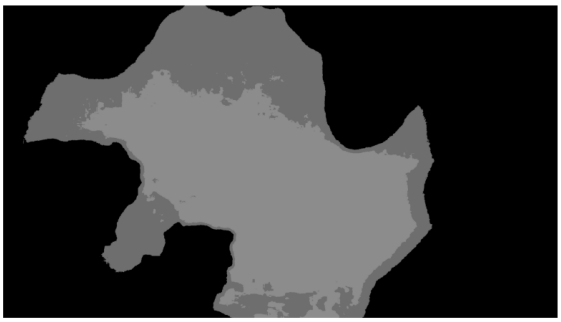

Figure 4.12: Segmenting an image of a degassing process into three classes using the k-means algorithm. The black parts are the obstructions to the camera. The dark grey parts are the wall of the tank and some miss-classified fluid. The lightest part in the middle is the moving fluid.

In the earlier stages of the degassing process it is easier to see the difference of the properties in the fluid. It can clearly be seen what is unmelted slag, melted slag and steel. Because of this, we can try to further segment an image of the earlier stage of the degassing process into five classes. The k-means algorithm performs admirably and accurately in this task. By applying k-means to segment an image into more classes with (k = 5), we obtain the classification shown in Figure 4.13.

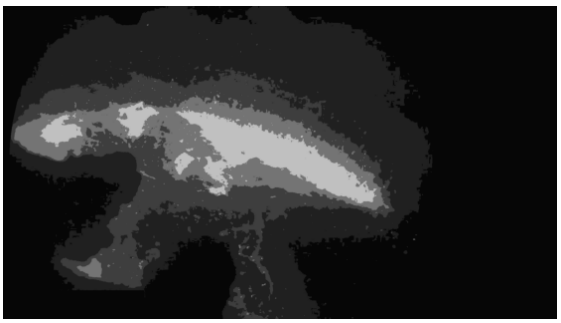

Figure 4.13: Segmenting an image of an early step of a degassing process into five classes using the k-means algorithm. The black parts are the obstructions to the camera. The dark grey parts are the wall of the tank and some miss-classified fluid. The three most lid up classes are segmenting the fluid into unmelted slag, melted slag and steel.

Figure 4.14 shows the results of using Gaussian Mixture Model (GMM) to segment an image of a full boiling process into three classes. While some details such as fluid and holes appear differently than when using the previous segmentation methods, the overall appearance remains similar.

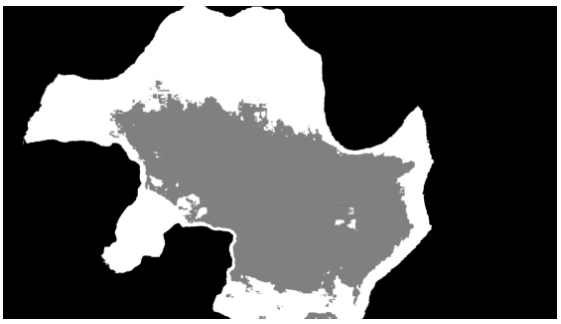

Figure 4.14: Segmenting an image of a degassing process into three classes using Gaussian Mixture Model. The black areas are the obstructions to the camera lens. The white areas are the wall of the ladle and some miss-classified fluid. The dark gray area is the moving fluid.

When using the GMM to segment fluid into five classes, the result is very similar to when running the k-means algorithm, as can be seen in Figure 4.15. It accurately segments the image into five classes.

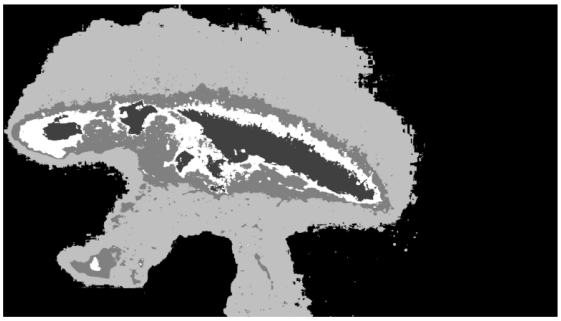

Figure 4.15: Segmenting an image of a degassing process into five classes using Gaussian Mixture Model. The black parts are the obstructions to the camera lens. The light gray areas are the wall of the ladle and some miss-classified fluid. The three areas closest to the middle are the fluid segmented into unmelted slag, melted slag and steel.

4.4 Estimating the height/volume of fluid

To determine the height or volume of fluid in the ladle, an image that has been segmented using the k-means algorithm with k=3 can help us to estimate it. This segmentation gives an estimate of the size of the fluid in the ladle but, more importantly, shows us how much its level has changed from the beginning. To account for this variation in levels since starting, we apply the assumed base level of fluid to each image to provide us with a reliable starting point when computing our estimates.

At an earlier step of this process, Hough transforms were utilized to determine the base level of the fluid. Once determined, this level was then applied to an image file, producing the dark grey ellipse at the bottom of Figure 4.16. A separate dark grey

area at the top represents everything above the lid of the ladle which falls outside its scope for computation.

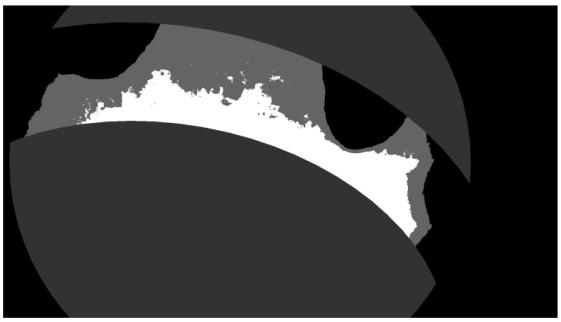

Figure 4.16: Segmented image of a degassing process. The white parts in the middle represent the moving fluid in the ladle. The dark gray ellipse on the bottom and the top represent the base fluid and the lid of the ladle respectively. The black area inside the boundaries is the obstructing slag.

This calculation seeks to measure how far fluid has traveled from its starting point to the lid of the ladle. We define height increase as the area occupied by white portions in relation to the tank walls and obstructions such as slag. Note, however, that this method is only an estimate as projecting three-dimensional fluid into a two-dimensional image results in some loss of accuracy.

We are not only interested in measuring how much fluid has risen from its initial level to the ladle lid, but also in any excess that might accumulate above it. To calculate this, we create an ellipse from the lid to bottom of the image that allows us to focus on only the parts above the lid of the ladle.

Figure 4.17 provides an example of this approach. Although no excess fluid is visible in this image, its purpose is still clear - to demonstrate the principles behind the method.

Figure 4.17: Segmented image of a degassing process. The dark gray ellipse on the bottom represents the lid of the ladle. The black area inside the boundaries is the obstructing slag.

4.5 Estimating movement of the degassing process

Movement in an image can be calculated based on the differences of the pixel values between the frames. As illustrated on the right of Figure 4.18, various statistics from the image are extracted. Brighter pixels indicate greater change from previous frames, thus suggesting some movement within it.

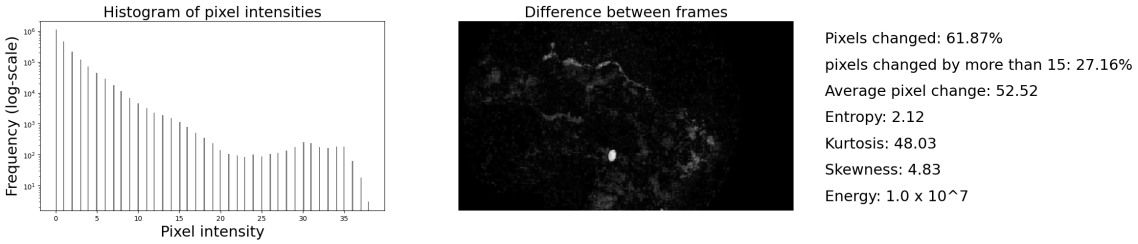

Figure 4.18: Estimated movement between two frames of a degassing process. The image to the right shows the movement between two frames with brighter pixels indicating more movement. The histogram to the left shows the frequencies of the pixel values of the right image. The statistics furthest to the right show some properties of the movement.

One of the statistics computed is the percentage of the pixel values that changed between two consecutive frames. This is defined by looking at the entire image and dividing by all the pixels with even the slightest difference in value from the previous frame. This gives information about the overall area of movement. The second statistic shows which percentage of the pixel values changed more than a certain threshold in an image that can be set by operators according to their desired level of interest. This gives some insight into the intensity of the motion such as if the fluid has moved significantly at certain frames. The entropy shows the disorder or randomness of the process. The kurtosis and skewness both gives clear values of the behaviour of the distribution of pixel intensities in the image. The Energy shows how much intensity variations there is within the image.

After running the movement algorithm for 25 frames an average is computed. Figure 4.19 shows the average movement of 25 frames. It shows promising results since it indicates in which parts of the image the fluid moves the most. The white spots in the image also shows some bubbles appearing that could be analyzed further.

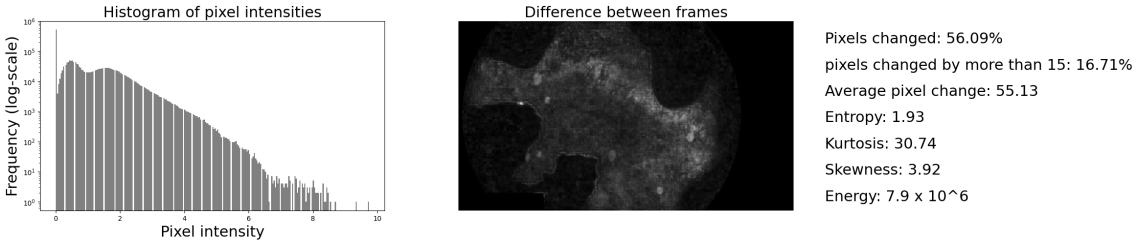

Figure 4.19: Estimated average movement between 25 frames of a degassing process. The right image shows the average of 25 frame-to-frame movement differences. Brighter pixels in the image indicate more movement. The histogram to the left shows the frequencies of the pixel values of the right image. The statistics furthest to the right shows some properties of the movement.

4.6 Generalization and testing with more data

To ensure that a system can effectively function with new degassing process data, its design must include some degree of generality. Here, we will explore any adjustments or modifications required when running the system developed in this work on other degassing processes. Furthermore, we will evaluate the performance of the system by running it on a set of entire degassing videos; allowing us to identify any areas requiring improvement while assuring it functions optimally in various scenarios.

The initial step of the system is calculating an accurate base level for the fluid. Unfortunately, the Hough transform which is used to find this level contains an error. Specifically that its radius search criteria does not work correctly in certain instances. These are instances where the circular line shown in the left image of Figure 4.7 does not appear as clearly. This is because for some videos the slag layer on top is covering the fluid in a different way or the lighting of the image is different. When this happens the Hough Transform sometimes places the circle at an incorrect spot and with an incorrect size. In order to increase the reliability and accuracy of the system we currently must manually specify a radius value per degassing video in order to accurately identify base levels. Even when manually specifying a radius the Hough Transform sometimes fail to find an accurate base level because of the different instances explained above. So currently, the best way is to manually draw the base level for each video.

Once the system was run over a set of entire videos, the data could be plotted. Since the data was exported to Excel it was easy to plot the data of the extracted statistics. Figure 4.20 shows histograms of four graphs of statistics extracted from a video of a degassing process. Inspection of the these graphs in comparison with the true video footage is proof that the computed statistics such as the Fluid above/below lid and the movement are correlated to the degassing parameters in the video footage. For example, when the fluid below lid rises in the graph of figure 4.20, the fluid also rises at a similar level in the video. This indicates that these four estimates are reasonably accurate. By studying the four graphs, it becomes evident that the process undergoes changes over time. At first, movement was minimal, with the

height/volume increasing gradually over time. Once boiling begins, the movement increased significantly and the height/volume began increasing more quickly than previously. At the end of the process, boiling becomes excessive, necessitating an operator to act decisively to control it. The height/volume statistics also reflect this escalation with an increase.

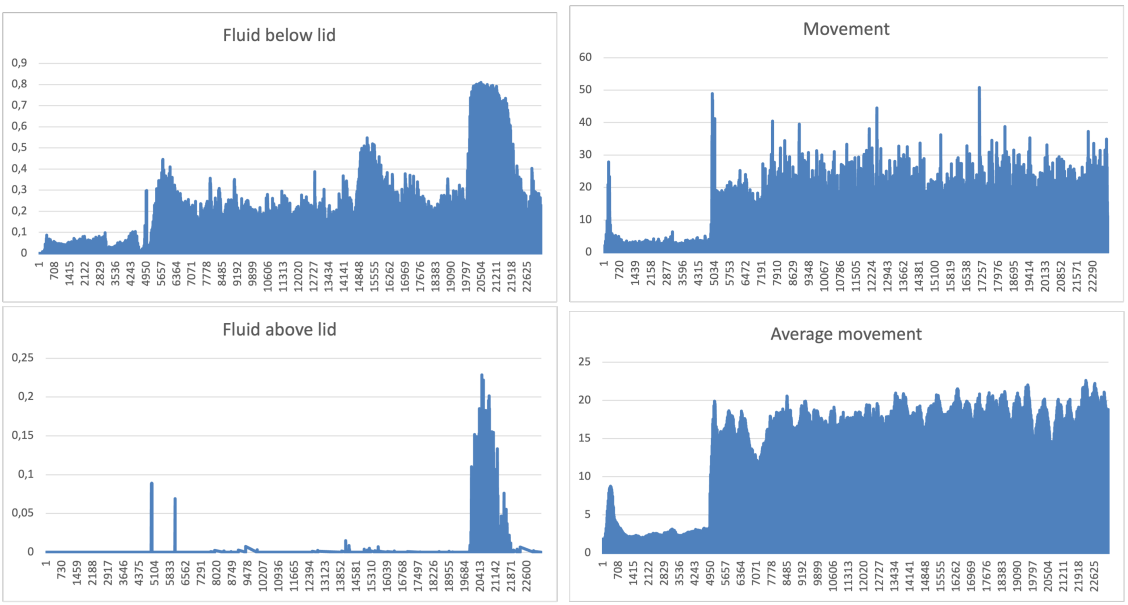

Figure 4.20: Graphed data points at every 10th frame of an entire degassing video. These data points are computed by running the developed system on the video data.

It was noted that the statistics did not always give accurate results due to some unknown errors arising in some frames. With this meaning the system did not provide fully reliable statistics at every frame. But by plotting graphs of the saved data in Excel actually produced promising results if these errors were to be fixed.

4.7 Performance of system

To maximize the efficiency of the script, it is important to monitor the performance of its individual parts. By identifying the components which take up most of the time, we can optimize them and work towards an ideal real-time system. A way of measuring performance is adding benchmarking checkpoints before each computation in the script; this approach allows us to accurately assess processing times and identify the ones in need for optimization.

Table 4.1 gives the execution time for every process executed on every frame. Computing all the statistics takes approximately 6 seconds in total. The k-means algorithms using both k=5 and k=3 consume considerable amounts of time. The slowest process is the plotting of the results. This illustrates some areas that may need further optimization to increase the processing efficiency.

Table 4.1: Time in seconds taken for each process to extract data from a frame (average of 100 frames)

Process	**Time Taken (s)**
Load image	0.001994
Preprocessing	0.125380
Create binary filter	0.027204
Extract interesting parts	0.012435
Create blocking slag image	0.009933
Segmenting with k=5	1.098930
Adding lower bound	0.012152
Adding upper bound	0.003995
Segmenting with k=3	0.767357
Adding base	0.009983
Adding base and lid	0.001990
Adding above lid boundary	0.001988
First tolerance	0.026331
Second tolerance	0.012917
Compute frame to frame movement	0.010510
Change moving average list	0.000000
Compute average movement	0.010443
Create image dictionary	0.000000
Compute fluid parameters	0.080501
Compute height parameters	0.097823
Compute movement parameters	0.206309
Plot results	3.803307
Total Time	**6.357513**

5
Conclusion

5.1 Verdict

Research question - How can modern image analysis tools be utilized to extract useful information from the camera feed of the degassing process?

Prior to this experiment, OVAKO had never conducted image analysis of its degassing process - making this experiment the first use of these techniques. In this project, more useful data than initially anticipated were produced and even more statistics from the video footage could have been provided if more time had been given. The system revealed an abundance of useful data, particularly regarding the fluid volume/height and the movement. These discoveries proved the immense potential of image analysis as a tool to aid the operators, with numerous valuable pieces of information extracted from the video. The next steps involve determining which information would most benefit the operators and how the procedure introduced here should be developed/optimized further. Undertaking image analysis could greatly help the operators by extracting essential data from the camera feed.

5.2 Contributions

In this thesis we have developed a system that takes a video of a degassing process and a user-specified frame rate as inputs, and produces the following outputs:

- A folder of plots that visualize each of the extracted statistics.
- An Excel file containing all the extracted statistics for the requested frame rate.
- A GIF video that showcases the different plots and statistics.

Behind the development of this system are many image analysis techniques and statistical methods. This system is designed to be used by researchers and engineers at OVAKO and Swerim for the analysis of new degassing videos.

5.3 Future Work

To reach the desired end goal of creating a fully live system to aid operators, there are various areas requiring further development. Below, some suggestions for the continuation of this project are listed.

First, the Hough transform used to find the base levels presents some issues; since, the radius differ between videos. To address this, an adaptive radius finder could be developed which would identify the ideal radius values. This would enable the Hough transform to more effectively operate across a wider variety of degassing videos.

Second, when the ladle walls and the lid share similar color tones as the obstructing slag in the image, inaccurate segmentation in these areas is obtained. To remedy this, image processing techniques like erosion and dilation could be implemented before segmentation of these images in such cases.

Third, OVAKO Engineers believe that the height of the steel pumps is an important variable; however, the current system only computes the relative amount of steel produced by each steel pump in the image. Implementing a method for measuring the height would prove helpful to their production.

Fourth, all statistics are calculated for every frame in the present approach. To increase the efficiency, an automatic function could be added that detects which part of the degassing process is being examined in each frame and accordingly updates plots with relevant statistics for that stage of process. An example of such a case could be in the early stage of the degassing process. Then it is easier to separate out unmelted/melted slag/steel in the image. In the earlier stages of the degassing process the plot could display this information and when boiling commences these statistics could be removed from plot.

Fifth, prior to designing a live system, it is necessary to determine its complexity level. If frequent calculations are desired, high performance programming languages might be better suited than the current approach written in Python. Otherwise, this current system might be enough.

Sixth, there are various ways the computational speed of the system could be increased. One solution would be using OpenCV library to plot the images. Also, consider whether using K-means algorithm or another faster segmentation method could provide faster segmentation times. Finally, adjusting the parameters of the algorithms for the specific cases could increase the speed.

Lastly, it would be useful to decide how frequently certain statistics need to be computed. For instance, it may not be necessary to compute the blocking slag levels at every frame as it tends to remain stationary on the camera lens for several seconds after changing.

Bibliography

[1] Hastie, T., Tibshirani, R., Friedman, J. (2017). The elements of statistical learning. Springer-Verlag.

[2] Hirzel, M., Schneider, S., & Tangwongsan, K. (2017). Sliding-Window Aggregation Algorithms: Tutorial. In Proceedings of the 11th ACM International Conference on Distributed and Event-Based Systems (DEBS '17). Association for Computing Machinery. https://doi.org/10.1145/3093742.3095107

[3] Singla, N. (2014). Motion detection based on frame difference method. International Journal of Information & Computation Technology, 4(15), 1559-1565.

[4] Pal, N. R., & Pal, S. K. (1991). Entropy: a new definition and its applications. IEEE Transactions on Systems, Man, and Cybernetics, 21(5), 1260-1270. https://doi.org/10.1109/21.120079

[5] Makkar, H., & Pundir, A. (2014). Image analysis using improved Otsu's thresholding method. International Journal on Recent and Innovation Trends in Computing and Communication, 2(8).

[6] Hassanein, A. S., Mohammad, S., Sameer, M., & Ragab, M. E. (2015). A Survey on Hough Transform, Theory, Techniques and Applications. Informatics Department, Electronics Research Institute, El-Dokki, Giza, 12622, Egypt. [Submitted on 7 Feb 2015].

[7] Cabello, F., Leon, J., Iano, Y., Arthur, R. (2015). Implementation of a fixed-point 2D Gaussian Filter for Image Processing based on FPGA. Signal Processing - Algorithms, Architectures, Arrangements, and Applications Conference Proceedings, SPA, 2015-December, 28-33. https://doi.org/10.1109/SPA.2015.7365108

[8] Zhang, B., & Allebach, J. P. (2008). Adaptive bilateral filter for sharpness enhancement and noise removal. IEEE Transactions on Image Processing, 17(5), 664-678. https://doi.org/10.1109/TIP.2008.919949

[9] Buades, A., Coll, B., & Morel, J.-M. (2011). Non-Local Means Denoising. Image Processing On Line, 1, 208-212. https://doi.org/10.5201/IPOL.2011.BCM_NLM

[10] Levkine, G. (n.d.). Prewitt, Sobel and Scharr gradient 5x5 convolution matrices. Retrieved May 3, 2023, from https://hlevkin.com/hlevkin/47articles/SobelScharrGradients5x5.pdf

[11] Sagitov, S. (2022). Introduction to Statistical Inference. Course litterature, Chalmers University of Technology.

DEPARTMENT OF MATHEMATICAL SCIENCES
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY