

# Integrating machine learning and image processing for void fraction estimation in two-phase flow through corrugated channels

Stefano Passoni<sup>a,b</sup>, Riccardo Mereu<sup>a,\*</sup>, Matteo Bucci<sup>b</sup>

<sup>a</sup> Department of Energy, Politecnico di Milano, via Lambruschini 4A, 20156, Milan, Italy

<sup>b</sup> Department of Nuclear Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

## ARTICLE INFO

### Keywords:

Compact plate heat exchanger  
Convolutional Neural Networks  
U-Net  
Image segmentation  
Void fraction  
Drift flux model

## ABSTRACT

There is a substantial amount of information embedded in images of two-phase flow captured through high-speed video (HSV) or high-resolution photography. However, accurate image segmentation is necessary to unlock a meaningful analysis of the data. In this study, we discuss how to estimate the flow void fraction in chevron-type corrugated channels typical of compact plate heat exchangers (CPHE) from back-lit front-view HSV images, using machine learning (ML) algorithms and data processing techniques. A U-Net neural network was employed for image segmentation, demonstrating robust performance with evaluation metrics consistently exceeding 0.9. The binary masks (indicating gas or liquid phases) derived from segmentation were processed in MATLAB<sup>®</sup> to estimate void fraction through a 3D reconstruction algorithm of the gas cluster's volume. In contrast to conventional void fraction estimates based on the area ratio of binary masks, this algorithm models the curvature of the liquid-vapor interface through the corrugated channel. When compared to other methods, our algorithm predicts very similar void fraction contour maps. However, the average discrepancy between our algorithm and the area-ratio approach can be as high as 80%, underscoring the importance of the processing method in analyzing the data and developing correlations. Finally, a drift flux model was introduced to predict the void fraction distribution using a two-part equation accommodating the dependency of the distribution coefficient  $C_0$  on the liquid flow rate for a corrugation Froude number  $Fr_c$  larger than 1. The proposed model can predict the void fraction dataset with a mean average percentage error of 8.17%. In summary, U-Net's pixel-level accuracy facilitates deep and precise post-processing of HSV images, enabling meaningful void fraction measurements. Thanks to its generality and minimal training effort requirements, the discussed methodology can be applied to estimate void fractions in various two-phase flow experiments and operating conditions.

## 1. Introduction

Compact plate heat exchangers (CPHEs) are very well known in the industrial world. For decades they have been used effectively for operations with single-phase fluids. When the brazing process started to replace gaskets to seal the channels formed by the different plates, this technology started to be used in the heating, ventilation and air conditioning industry and operated with two-phase flows as evaporators or condensers. Nowadays, given the performance of such heat exchangers compared to their small volume footprint, they are actively being investigated for different purposes like compact steam generators for small modular reactors (Kang et al., 2022). This application is rather new and so understanding the behavior of two-phase flow in compact plate heat exchangers is crucial for optimizing their design. Researchers and engineers could in fact develop more efficient designs

by studying pressure drops, two-phase flow behavior and heat transfer characteristics.

Flow regimes play an important role in the multiphase fluid dynamics of such components. The phase distribution can highly influence the thermal-hydraulic performance of single channels and ad-hoc corrugations patterns and distribution sections are key to optimize performance (Ayub, 2003). Various research studies involved the visualization and visual categorization of two-phase flow patterns within chevron-type corrugated channels (Tribbe and Müller-Steinhagen, 2001; Vlasogiannis et al., 2002; Nilpueng and Wongwises, 2010; Grabenstein et al., 2017; Buscher, 2019; Shiomi et al., 2004). The standard way to categorize flow pattern is by naked-eye. This method is inherently inaccurate especially near flow pattern transition due to the subjectivity of the observation and scarce reproducibility of the

\* Corresponding author.

E-mail addresses: [stefano.passoni@polimi.it](mailto:stefano.passoni@polimi.it) (S. Passoni), [riccardo.mereu@polimi.it](mailto:riccardo.mereu@polimi.it) (R. Mereu), [mbucci@mit.edu](mailto:mbucci@mit.edu) (M. Bucci).

<https://doi.org/10.1016/j.ijmultiphaseflow.2024.104871>

Received 22 February 2024; Received in revised form 6 May 2024; Accepted 14 May 2024

Available online 18 May 2024

0301-9322/© 2024 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

results. Flow visualization is often performed with the aid of digital imaging, using high-speed video or high-resolution photography. There is a huge amount of information embedded in those data and using it only to visually classify the flow pattern would not be fruitful. Image analysis is therefore a very important step in the multiphase flow data processing workflow. Although this has been the common approach for different kind of geometries, image analysis has rarely been performed for two-phase flow in plate heat exchangers. To the best of our knowledge, the only two available works on the topic are very recent. Wang et al. (2021) used a CNN-based YOLO (you-only-look-once) algorithm for detection and tracking of bubbles in a dimpled PHE channel whereas (Buscher, 2022) applied bubble detection and texture analysis algorithms to images of air–water flow in a transparent cross-corrugated channel. Unfortunately, segmentation algorithms based on standard image-processing techniques like edge-detection and thresholding, do not provide an accurate output when applied to such complex flows as the one established by the cross-corrugated pattern. In addition, these methods may lack of generality as researchers are often required to manually tweak some parameter to adapt the codes to the different flow regimes or optical setups.

More advanced techniques are required to obtain a very precise image segmentation therefore enabling an accurate post-processing. To this cause, Convolutional Neural Networks (CNNs) have gained significant attention and success in various computer vision tasks, including image segmentation for two-phase flow analysis like IR images of boiling surfaces (Ravichandran et al., 2021, 2023). While CNNs have been widely used for image analysis tasks, their application to the specific problem of multiphase flow in compact plate heat exchangers is completely novel. Few of the advantages of machine learning for image segmentation are a general improvement of accuracy and precision of the outcome, possibility of automation and standardization thus minimizing human error and potential for generalization to similar flow configuration.

In this study we will discuss a methodology that integrates machine learning (ML) algorithms and data processing techniques to perform post-processing on back-lit front-view images of two-phase flow in corrugated channels. We adopted a U-Net architecture described in Falk et al. (2019) for image segmentation and leveraged its transfer-learning capabilities to minimize the training dataset resulting in the manual annotation of just few images per flow regime. This was the main driver for selecting the U-Net architecture as alternative methods would have required a complete training from scratch, thereby increasing computational costs and time therefore reducing the effectiveness of the proposed approach. The objective is to discuss a post-processing methodology to estimate the void fraction from the binarized images, highlighting how sensible the estimation is according to the different approaches adopted. In Section 2, the data acquisition and pre-processing, details of the network used for segmentation as well as the post-processing algorithms will be thoroughly presented. In Section 3, the results of the work are critically discussed. The validation of the segmentation platform is addressed and the results of void fraction estimation methodology are presented. Conclusions are drawn in Section 4.

## 2. Methodology

### 2.1. Data acquisition and pre-processing

A schematic of the experimental loop used for image acquisition is displayed in Fig. 1. The setup consists of a closed-loop water circuit connected to an open-loop compressed air circuit. Adiabatic mixing between air and water is obtained before entering the corrugated channel through a T-junction before feed the mixture to an inlet plenum. A complete description of the experimental loop is reported in Passoni et al. (2024). The facility was operated in upward configuration, and flow visualization was possible thanks to the two transparent corrugated

**Table 1**

Geometrical parameters of present study's plates.

Corrugation angle	$\phi$	63°
Corrugation depth	$b_p$	2.5 mm
Corrugation wavelength	$\lambda$	9 mm
Plate width	$B_p$	182 mm
Plate/Corrugation length	$L_p$	320 mm
Enlargement factor	$\Phi$	1.1712
Frontal area	$A_f$	455 mm <sup>2</sup>
Hydraulic diameter	$D_h$	4.3 mm

plates made from resin casted directly in a silicon mold obtained from the original steel plates. The chevron pattern on the plates featured corrugation angles of 63° and grooves 2.5 mm deep. Table 1 reports all the details about the corrugation geometry and Fig. 2 displays the main ones. The grooved channel was back-illuminated by a flicker-free panel of white LEDs powered by direct current (DC). A light diffuser sheet was used to guarantee uniformity of illumination. High-speed videos were filmed using a Phantom Miro C110 high-fps camera. The camera sampling rate was configured at 1900 frames per second, allowing for the detailed visualization of rapid shifts and intricate patterns with exceptional time accuracy. The videos were captured in a resolution of 1280 × 480 pixels. To achieve proper luminosity across the entire frame, the camera exposure time was set to 40  $\mu$ s. The recording employed a NIKKOR AF-S Micro 60 mm lens, which was manually calibrated for focus and aperture settings to maintain precise control over image sharpness and depth of field. The dataset of the observed flow conditions encompasses a liquid superficial velocity spanning from 0.006 m/s to about 0.5 m/s along with a gas superficial velocity ranging between 0.007 m/s and 3.25 m/s. Additionally, the dataset maintains a homogeneous volume fraction  $x_V$  that varies from 1.8 to 99.8%. The test matrix of the visualization data collected is displayed in Fig. 3.

Once the videos were acquired, before starting the post-processing phase, the total number of frames to be considered for each video was selected. For every video, we reduced the frame count by retaining only one out of every hundred frames. This reduced the dataset size and the computational load for processing each experimental point, resulting in only 88 frames per condition. The determination of this specific frame count was informed by an assessment of how statistical metrics related to void fraction and total gas volume per frame responded to varying frame numbers. Our chosen value ensured that the percentage difference in mean and standard deviation of these parameters, compared to the maximum frame count, remained below 0.5%.

Prior to segmentation, a preliminary processing procedure is applied to all experimental videos, enhancing and preparing the input images. The input frames are read into a MATLAB® script where they are first converted in 8-bit greyscale images and then intensity based image registration is performed to correct eventual misalignments between each frame and the background before subtraction. This latter operation is performed with the `imregister` function present in MATLAB®. The process involves comparing pixel intensities to identify corresponding regions within the images being compared. Then the algorithm can determine spatial transformations needed to overlay the different images accurately. The resulting image is then cropped to the desired size of 460-by-1180 pixels and contrast is enhanced by saturating the top 2.5% of darkest and brightest pixels by linearly remapping the contrast levels of the input image to the required output ones. All the datasets are then exported as a lossless multi-page tiff file each containing the selected number of frame for each experimental point. The frames used for training and validation of the U-Net were manually annotated with the aid of the software Fiji (Schindelin et al., 2012) by marking the areas where gas bubbles could be visually identified.

### 2.2. CNNs for image segmentation

The U-Net is an encoder–decoder-style neural network adept at solving semantic segmentation tasks. It was firstly presented in Ronneberger et al. (2015) and implemented as an ImageJ (Schneider et al.,

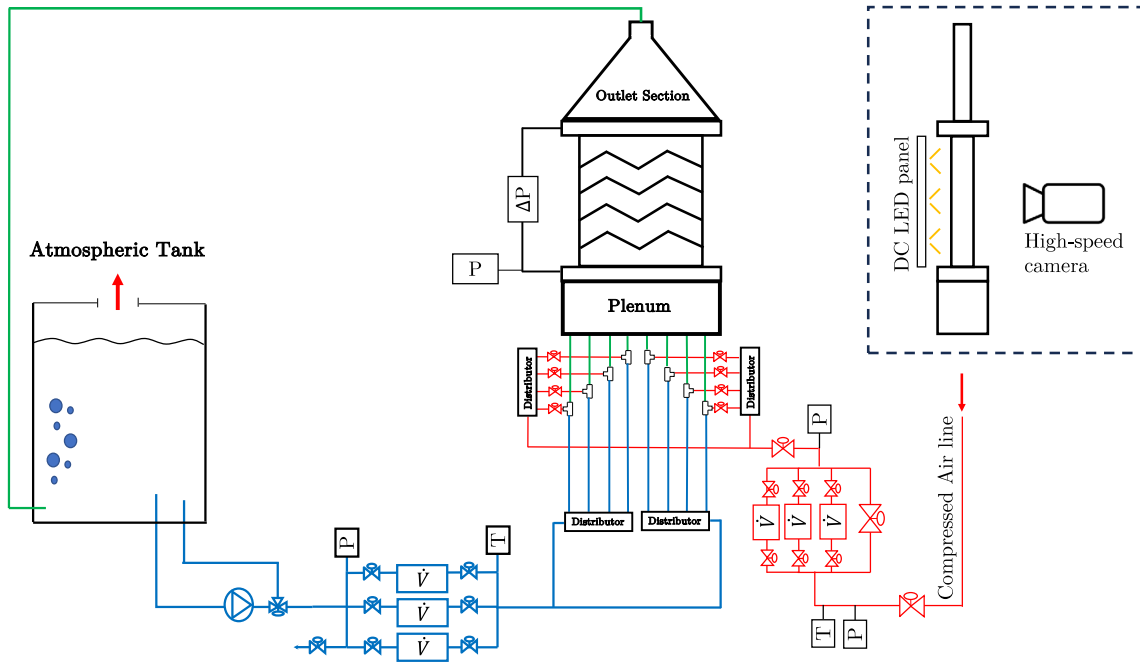


Fig. 1. Schematic of the loop adopted for the experiments with a detail on the flow visualization setup.




Fig. 2. Representation of the corrugation's main geometrical parameters.




Fig. 3. Test matrix of the flow visualization campaign. Dotted lines represent iso-lines of constant volumetric fraction  $x_v$ .

2012) plugin by Falk et al. (2019). U-Net was developed for single-cell detection and segmentation in biomedical images. Its peculiar network architecture depicted in Fig. 4 utilizes a multi-level arrangement of up-sampling and down-sampling layers, enabling it to encompass the

complete image context while remaining trainable even with limited datasets. As described in Ronneberger et al. (2015), the architecture consists of a contracting (or encoding) path on the left side and an expansive (or decoding) path on the right. The contracting path adheres to the conventional convolutional network design. It involves iteratively applying two  $3 \times 3$  convolutions, each succeeded by a rectified linear unit (ReLU), along with a  $2 \times 2$  max pooling operation. With each step of downsampling, the number of feature channels is doubled. On the contrary, each stage in the expansive path features an upsampling of the feature map, followed by a  $2 \times 2$  convolution (referred to as “up-convolution”) that reduces the feature channel count by half. This is followed by a concatenation with the corresponding cropped feature map from the contracting path, and then two  $3 \times 3$  convolutions, each followed by a ReLU activation. In the final layer, a  $1 \times 1$  convolution is employed to map the feature vector to the desired number of classes. A more in-depth description of the U-Net architecture and its features is available in Falk et al. (2019) and Ronneberger et al. (2015).

The training process is the most important part of a workflow. It is the process of teaching an algorithm to recognize patterns, make predictions or perform tasks by exposing it to a dataset and allowing it to learn iteratively. During training, the model adjusts its internal weights based on the input data and the desired output, with the goal of minimizing a loss function. To reduce the amount of training data (i.e. manually annotated images) we exploited transfer learning (Torrey and Shavlik, 2010; Weiss et al., 2016). This process leverages the prior knowledge the U-Net has gained through training on a different dataset by only performing a fine-tuning of the model's weights and not a complete retraining. By doing this, the authors in Falk et al. (2019) claim that only a few annotated images are sufficient for a good fine tuning, with the most complex cases requiring no more than ten. In our particular case, we performed transfer learning from weights trained for segmentation of front-lit HSV images of sub-cooled flow boiling (Seong et al., 2023).

In order to train our model, we progressively fine-tuned the network by exposing it to more and more complex flow regimes. The fine-tuning process is generally halted after a certain number of iterations when cross-validation ceases to show improvement. The learning rate was set  $1e-5$  and the number of iterations was generally kept within 10 to 20 thousands. Training was performed on a NVIDIA Tesla V100




Fig. 4. Example of the U-Net architecture. The computed feature hierarchy is visualized by the blocks. The numbers above each network block represent the count of feature channels, while the numbers to the right of each block indicate the spatial dimension in pixels.

GPU with 32 GB of VRAM. The fine-tuning process started by feeding the network with coarse bubbly flow images at low liquid flow rate since they display mostly round and well-separated bubbles (Fig. 5(a)). Increasing air flowrate, Taylor-like bubbly flow appears. This flow regime exhibits larger, zigzag-shaped gas structures (Fig. 5(b)) so that the network can start learning the characteristics of such gas patches. Increasing liquid flowrate instead, the higher turbulence enhances bubble breakup resulting in finer structures coexistent with bigger patches (Figs. 5(c) and 5(d)). In the end, the last fine-tuning was performed with heterogeneous flow (Fig. 5(e)) which is the most complex and chaotic flow regime. Using this method, we observed an enhanced accuracy in the final model, leading to a more precise segmentation of various flow regimes, compared to the outcome achieved by conducting a single fine-tuning using all the designated training images. While this outcome might appear counter-intuitive, it may be due to the diverse array of flow features present across the various images. Generally, the loss function (weighted soft-max cross-entropy loss, described in Falk et al., 2019) following the fine-tuning process was lower when we executed multiple successive training phases, indicating that the network benefits from being exposed to similar flow features during training. Final segmentation was achieved by applying the trained network to the input frames exploiting the “rotation averaging” feature present in the Fiji plugin to increase the accuracy of the output binary mask. With this option the network segments different rotated variants of the input data and the segmentation results are derived by calculating an average softmax score across all predefined orientations, which has been observed to enhance the overall quality of segmentation for complex data sets.

The annotation procedure was refined to minimize the manual effort and time requirement. When possible, a thresholding algorithm was employed to convert the gray-scale image into binary form followed by manual annotation adjustments to correct any inaccuracies in the mask like improper segmentation of touching objects and small bubbles. The same strategy was followed also for complex flow patterns even though the initial binarization of images was achieved through the U-net segmentation output adopting a model not yet trained for such flow regime. For the various stages of fine-tuning a total of 13 images were annotated, 8 of which were used for training and 5 for validation purposes. In order to subsequently assess the network’s segmentation precision on entirely new images, another set of 10 images was annotated, across the different flow regimes.

### 2.3. Data post-processing and void fraction estimation

After segmenting all the images, post-processing is needed to extract all the available information. In particular, the segmented image is processed with MATLAB® before proceeding with the analysis. First, morphological operations were performed to fill small gaps in the binary mask and remove objects smaller than 4 pixels which are assumed to be noise. Then, to measure all the most important geometric bubble characteristics, the MATLAB® regionprops function was used to perform blob analysis. With this function it is possible to directly measure properties of connected components such as equivalent diameter, minor and major axis length, orientation and coordinates of the blob centroid, which will be useful for further analysis. Particularly, the equivalent diameter, used to distinguish between small and large gas clusters further in the analysis, is computed as the diameter of a circle with the same area as the blob, mathematically computed as  $\sqrt{4 * Area/\pi}$ .

A typical approach to estimate the void fraction is to approximate it with the ratio of the total blobs area over the total area of the image. This “area-ratio” approach benefits from the high accuracy of the segmentation reached with the U-Net but it will be always an overestimation of the actual void fraction as small bubbles that do not occupy the entire channel depth and large gas patches that do, are given the same weight in a 2D space. To overcome this limitation we developed an algorithm to reconstruct bubble volumes in 3D exploiting the pseudo-2D shape of the corrugated channel. To do so, we first computed a channel depth map of the analyzed frame by recreating the sinusoidal shape of the Chevron-type corrugation (according to the geometrical parameters of the plate) and discretizing it on a mesh grid with the same number of pixels as the processed frames. The pattern position was then adjusted by carefully calibrating the location of the contact points according to the experimental image as shown in Fig. 6. The outcome of this process are three matrixes representing positive and negative displacement, relative to the middle plane of the channel, as well as the total channel depth.

Subsequently, the blobs within the labeled image are categorized as either small or large based on their dimensions and their specific placement within the channel. Precisely, a cluster is designated as small if its equivalent diameter is less than the average channel depth across its footprint. This classification is needed to choose the proper volume reconstruction approach. A small bubble is confined entirely within an individual corrugation pitch and it could be approximated using an




Fig. 5. Details of some of the images used for training. Flow regime complexity increases from left to right.




Fig. 6. Example of contact points position calibration for the computation of the depth map. The red regions represents the contact points between the two corrugated plates identified by a total depth lower than 0.2 mm for the sake of visualization. Due to the geometry of the corrugation, a double contact point is created at the pivot point when the corrugation changes slope. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

ellipsoid obtained by revolving its area along its minor axis. In this case, the volume is then computed as

$$V_{small} = \frac{4}{3}\pi ab^2 \quad (1)$$

where  $a$  and  $b$  indicate the length of the major and minor semiaxes.

On the contrary, for larger structures, the volume estimation could be performed by coherently summing the values of the positive and negative displacement maps (for each pixel within their 2D area) as if each pixel were extruded according to the channel depth. This technique assumes of course that the bubble is occupying the entire channel therefore neglecting the thickness of the liquid film between the bubble and the wall. In addition, to be able to account for the curvature of the bubble at its boundary and render a more physical representation of bubble shape, we introduced a smoothing function that computes the pixel-wise percentage of volume to be subtracted from the extrusion maps of large gas blobs. To serve this purpose, an exponentially decaying function based on the solution of the Laplace equation given in Gennes et al. (2004) for surface perturbations, was chosen as a proxy to describe the curvature. This was implemented into the post processing algorithm as shown in Eq. (2).

$$s = \exp(-k * distMap) \quad (2)$$

This equation describes the smoothing  $s$  as function of the euclidean distance from the border of the bubble ( $distMap$ ) and a length scale ( $k$ ). The first parameter (i.e.,  $s$ ) is computed with the function `bwdist` available in MATLAB® and, before applying to it the conversion factor from pixel to millimetres, a value of 1 is subtracted from all elements to set a value of exactly 0 for all the pixels of the perimeter of the gas cluster. The parameter  $k$  is instead a length scale equal to either the capillary length  $\lambda_c = \sqrt{\frac{\sigma}{\rho g}}$  (2.7 mm for water) or the mean channel depth over the analyzed bubble footprint, depending on which of the two is the smallest.

The volume of each blob is then stored in memory and the void fraction of each frame is computed as the sum of the volume of all the gas clusters over the total volume of the analyzed portion of channel computed from the depth map (Eq. (3)). The value of void fraction for the specific experimental point is calculated as the average of the

frame-specific void fractions over the number of frames analyzed (88) for each video.

$$\alpha = \frac{\sum V_{gas}}{V_{tot}} \quad (3)$$

Fig. 7 shows the difference between the 3D shape of a single bubble reconstructed with and without the smoothing function. Picture 7(b) shows not-so-physical vertical boundaries that lead surely to an overestimation of the volume. Fig. 7(b) exhibits instead the effect of the applied smoothing function on the same bubble. Fig. 8 displays then the volume reconstruction using the smoothing function on an entire frame. The reader can appreciate the different processing method used for small or large gas bubbles. Finally, the entire processing algorithm is summarized in the block diagram showed in Fig. 9.

### 3. Results and discussion

#### 3.1. Validation and U-net performance

The U-Net's segmentation performance was assessed by comparing the segmentation results against 10 manually annotated ground-truth images never used for model training. In this scenario, we performed a 1:1 comparison between the binary masks provided by the U-Net and the ones extracted from annotation. We then classified pixels as true positives (TP) when they were correctly segmented as gas phase, false positives (FP) if falsely segmented as gas, true negatives (TN) when pixels were correctly detected as background (i.e. liquid phase) and false negatives (FN) when incorrectly segmented as such. According to this classification, it is common practice to calculate pixel-based metrics to assess the accuracy of the model. The definition and description of such proxies is given in Table 2 and the results summarized in Table 3. In the latter, the different images used for validation are ordered for progressively increasing liquid superficial velocity.

It is easily noticeable how these metrics are generally high indicating an overall good segmentation quality. On average, all the metrics have a value larger than 0.9 with recall being the one that achieves the lowest value, stabilizing at an average score of 91.7%. This metric is in fact impacted the most by the presence of false negatives which

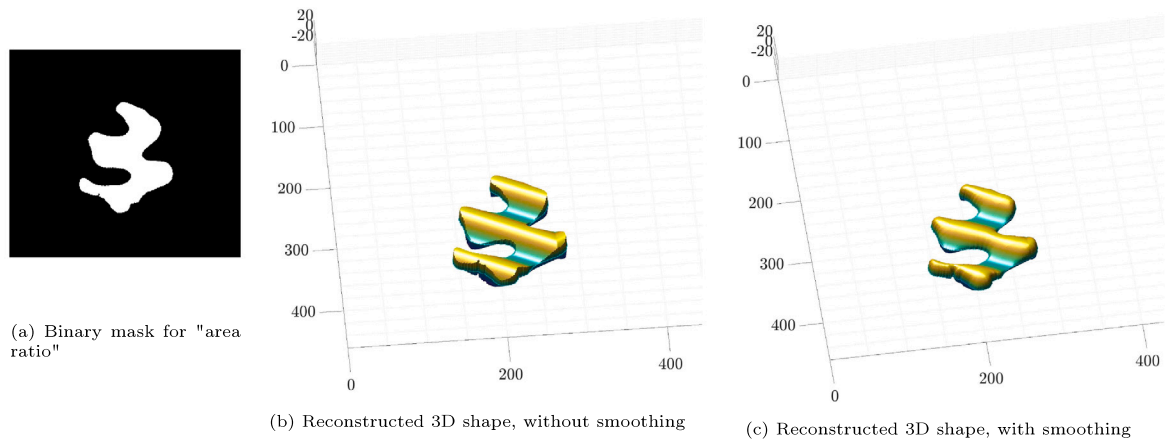


Fig. 7. Example of volume reconstruction for a large gas bubble. Dimensions in pixels.




Fig. 8. Example of gas volume reconstruction for an entire frame using the smoothing function. Dimensions in pixels.

Table 2

Description of the main metrics used to evaluate the segmentation results. ( $TP$ =True Positive,  $TN$ =True Negative,  $FP$ =False Positive,  $FN$ =False Negative).

Metric	Expression	Description
Accuracy	$\frac{TP+TN}{TP+TN+FP+FN}$	Rate of correctly classified pixels to the total number of pixels
Precision	$\frac{TP}{TP+FP}$	Amount of true positive predictions among all the positive predictions
Recall	$\frac{TP}{TP+FN}$	Amount of true positive predictions among all actual positive pixels
F1-Score	$\frac{2 \times TP}{(2 \times TP) + (FP + FN)}$	Harmonic mean of precision and recall, index of the overall model accuracy

Table 3

Evaluation of U-Net's segmentation performance for different liquid ( $J_L$ ) and gas ( $J_G$ ) superficial velocities.

Image	1	2	3	4	5	6	7	8	9	10	
$J_L$ [m/s]	0.006	0.012	0.031	0.031	0.031	0.031	0.061	0.061	0.092	0.092	
$J_G$ [m/s]	0.020	0.107	0.007	0.010	0.020	0.041	0.010	0.020	0.007	0.014	
True Positives	174 878	259 744	124 714	81 903	118 477	180 194	63 627	82 097	28 565	65 828	
True Negatives	365 090	277 350	415 339	462 038	423 131	358 371	481 419	459 503	520 364	476 511	
False Positives	750	2692	1397	260	648	1173	385	423	147	351	
False Negatives	11 282	12 214	10 550	7799	9744	12 262	6569	9977	2924	9310	<b>AVG</b>
Accuracy	0.978	0.973	0.978	0.985	0.981	0.976	0.987	0.981	0.994	0.983	0.982
Precision	0.996	0.990	0.989	0.997	0.995	0.994	0.994	0.995	0.995	0.995	0.994
Recall	0.939	0.955	0.922	0.913	0.924	0.936	0.906	0.892	0.907	0.876	0.917
F1-Score	0.967	0.972	0.954	0.953	0.958	0.964	0.948	0.940	0.949	0.932	0.954

are largely more than the false positive predictions. Fig. 10 displays few examples of the comparison between segmentation mask and ground-truth images. In these frames, the false positive and false negative pixels are highlighted in red and blue colors to visually highlight the regions of falsely classified pixels. Instead, the correct foreground segmentation (gas-phase) is shown in gray while the black background represents the liquid phase. From this figure it clearly stands out that

the false negative predictions are mainly located at the borders of the gas clusters, indicating that the U-Net tends to slightly shrink their dimension by a few pixels.

Unfortunately, when the flow is too complex due to the high superficial velocity of the mixture, breakage and fluid-dynamic instabilities are predominant and it becomes nearly impossible to manually annotate the experimental images to evaluate numerically the segmentation

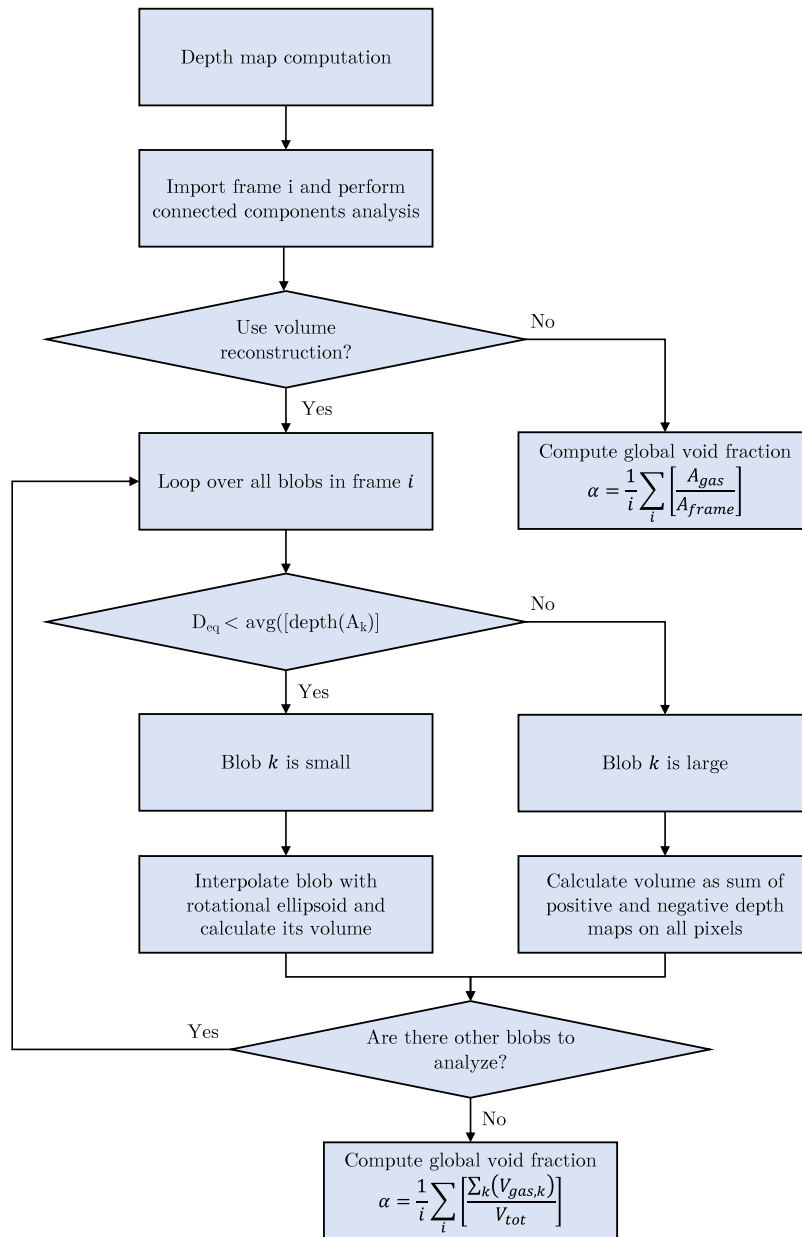


Fig. 9. Flowchart of the algorithm employed to compute the void fraction for both area-ratio and volume reconstruction approaches. In the boxes,  $i$  represents the number of frames and  $k$  the number of blobs in each frame.

quality. The only possible way to assess the segmentation accuracy is to visually compare the original frame and the binary mask obtained. As expected, at high water flow rates segmentation accuracy reduces. The gas clusters start to present no clear borders and the bubbles become so fine that superimpose one another leading to error in segmentation. The binary mask though still captures the major flow characteristics. The main concern is that the inaccuracy in the segmentation severely impacts the prediction of the void fraction. Due to the impossibility of quantifying such uncertainty, we decided to not report the results of void fraction estimation for a liquid superficial velocity higher than 0.25 m/s. Figs. 11–13 display few examples of the post-processing pipeline up to volume estimation for an increasing liquid superficial velocity of 0.006, 0.061 and 0.183 m/s, respectively, to span across the entire database. In these images, the four different post-processing steps already described are showcased and the reader can appreciate the quality of the segmentation task performed by the U-Net. At high liquid and gas flow rates (Fig. 13,  $J_G = 3.048$  [m/s]), the segmentation performance degrades according to the reasons explained above.

### 3.2. Void fraction

As introduced in Section 2.3, the most straightforward approach to estimate the void fraction when in possession of a binary mask is to make a simple area ratio between all the pixels labeled as gas and the total number of pixels. This approach has strong limitations as it implies that vertical footprint on the channel area has the same weight for small and large bubbles which is very unphysical. In this manuscript, this drawback is alleviated by reconstructing the volume of gas clusters and calculating the void fraction as a volume ratio. Also, to obtain more physical bubble curvature, we introduced a smoothing function inspired by a particular solution of Laplace equation. In the next paragraphs we present the results of void fraction estimation and modeling considering the approach based on volume smoothing as the baseline. To inquire about how sensitive this measurement is to the processing approach, we also compare these results with those obtained by estimating void fraction by area ratio and non-smoothed volume ratio.




Fig. 10. Examples of U-Net segmentation compared to ground-truth images for two different flow regimes. Entire frame on the left and detail on the right. Gray regions represents the agreement between segmentation and manual annotations while red pixels represent false positives (*FP*) and blue pixels the false negatives (*FN*). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

The results of the algorithm described in Section 2.3 are reported in the form of void fraction maps plotted against the superficial velocity of both phases in Fig. 14. These contour maps were obtained by fitting the data resulting from image analysis with a locally weighted smoothing linear regression (LOWESS) approach. In particular, Fig. 14 showcases three different maps obtained with the above-discussed methods for void fraction estimation: ratio of smoothed volumes (Fig. 14(a)), ratio of non-smoothed volumes (Fig. 14(b)) and the area ratio (Fig. 14(c)). As visible, the shape of the contour lines is very similar between all three maps but minor differences are present. The ratio of non-smoothed volumes is the approach that gives the single maximum void fraction value (0.94) while the other, smoothing or area ratio, give 0.80 and 0.83 respectively. The minimum value is instead very similar for the two approaches that use the reconstructed volumes, close to 0.008, while it is more than four times higher (0.036) for the area ratio approach. This behavior underlines the fact that the simple area ratio faces challenges in predicting very low void fractions because of its intrinsic reliance on two-dimensional bubble footprints, neglecting the crucial consideration of the actual three-dimensional volume occupied by the bubbles. These differences are also visible from the parity plots displayed in Fig. 15 where the baseline approach (ratio of smoothed volumes) is benchmarked against the other two. The usage of non-smoothed volumes has an almost constant deviation from baseline, with an average percentage difference of about 32%. Whereas, the usage of area ratio has a severe impact in the determination of the void fraction, leading to important differences especially in the region of low void fraction when pure bubbly flow is present for the reasons already explained. Fig. 16 gives more insights on the percentage differences when using these two approaches. Considering the area-based method, the percentage difference spikes to about 700% for low values of  $\alpha$  and it decreases moving to higher void fraction while remaining still quite scattered. Instead, the approach based on non-smoothed volumes, exhibits a maximum discrepancy ( $\sim 40\%$ ) for void fractions of about 5% and it is typically between 20% to 40% although the discrepancy

can be smaller for small void fractions, (i.e., small bubbles) as the smoothing has a little to no impact. This suggest the impact of smoothing is almost constant across the different void fractions. Furthermore, these data may be interpreted as representative of the upper bound of uncertainty inherent in the application of the smoothing function to the reconstruction of bubble curvature.

The void fraction data were then elaborated according to the drift flux model theory proposed in Zuber and Findlay (1965). In Fig. 17, the actual gas velocity ( $U_G = J_G/\alpha$  [m/s]) is plotted against the superficial velocity of the mixture ( $J_{mix} = J_G + J_L$  [m/s]) and colored according to the liquid phase superficial velocity (i.e. liquid flow rate). The Drift flux theory suggests these data could be fitted with a line in the form of  $U_G = C_0 * J_{mix} + U_{GJ}$  where  $C_0$  is called the distribution parameter and  $U_{GJ}$  the drift velocity (i.e., the relative velocity of the gas phase with respect to the superficial one of the mixture).

It is clearly visible from Fig. 17(a) that the higher the liquid superficial velocity, the higher the slope of the interpolating line meaning that the distribution parameter, which accounts for the non-uniform distribution of the two phases, clearly depends on such quantity. In addition, by looking more in detail at low mixture velocities, it can be seen that all the data below a certain  $J_{mix}$  collapse on a unique line. This particular threshold has been observed to be coincident with a corrugation Froude number  $Fr_{\sim}$  equal to one.

$$Fr_{\sim} = \frac{J_{mix}^2}{g \cdot b_p} \quad (4)$$

This dimensionless parameter was introduced by Buscher (2019) to quantify the relative effects of centrifugal and buoyancy forces on two-phase flow patterns in cross-corrugated channels. Buscher stated that a clear impact of buoyancy on flow patterns was observed only for  $Fr_{\sim} < 1$ . In this case, the slip between phases is dominated by buoyancy and so the distribution parameter is not dependent on liquid flow rate as it is observed to be for higher mixture superficial velocities. To model



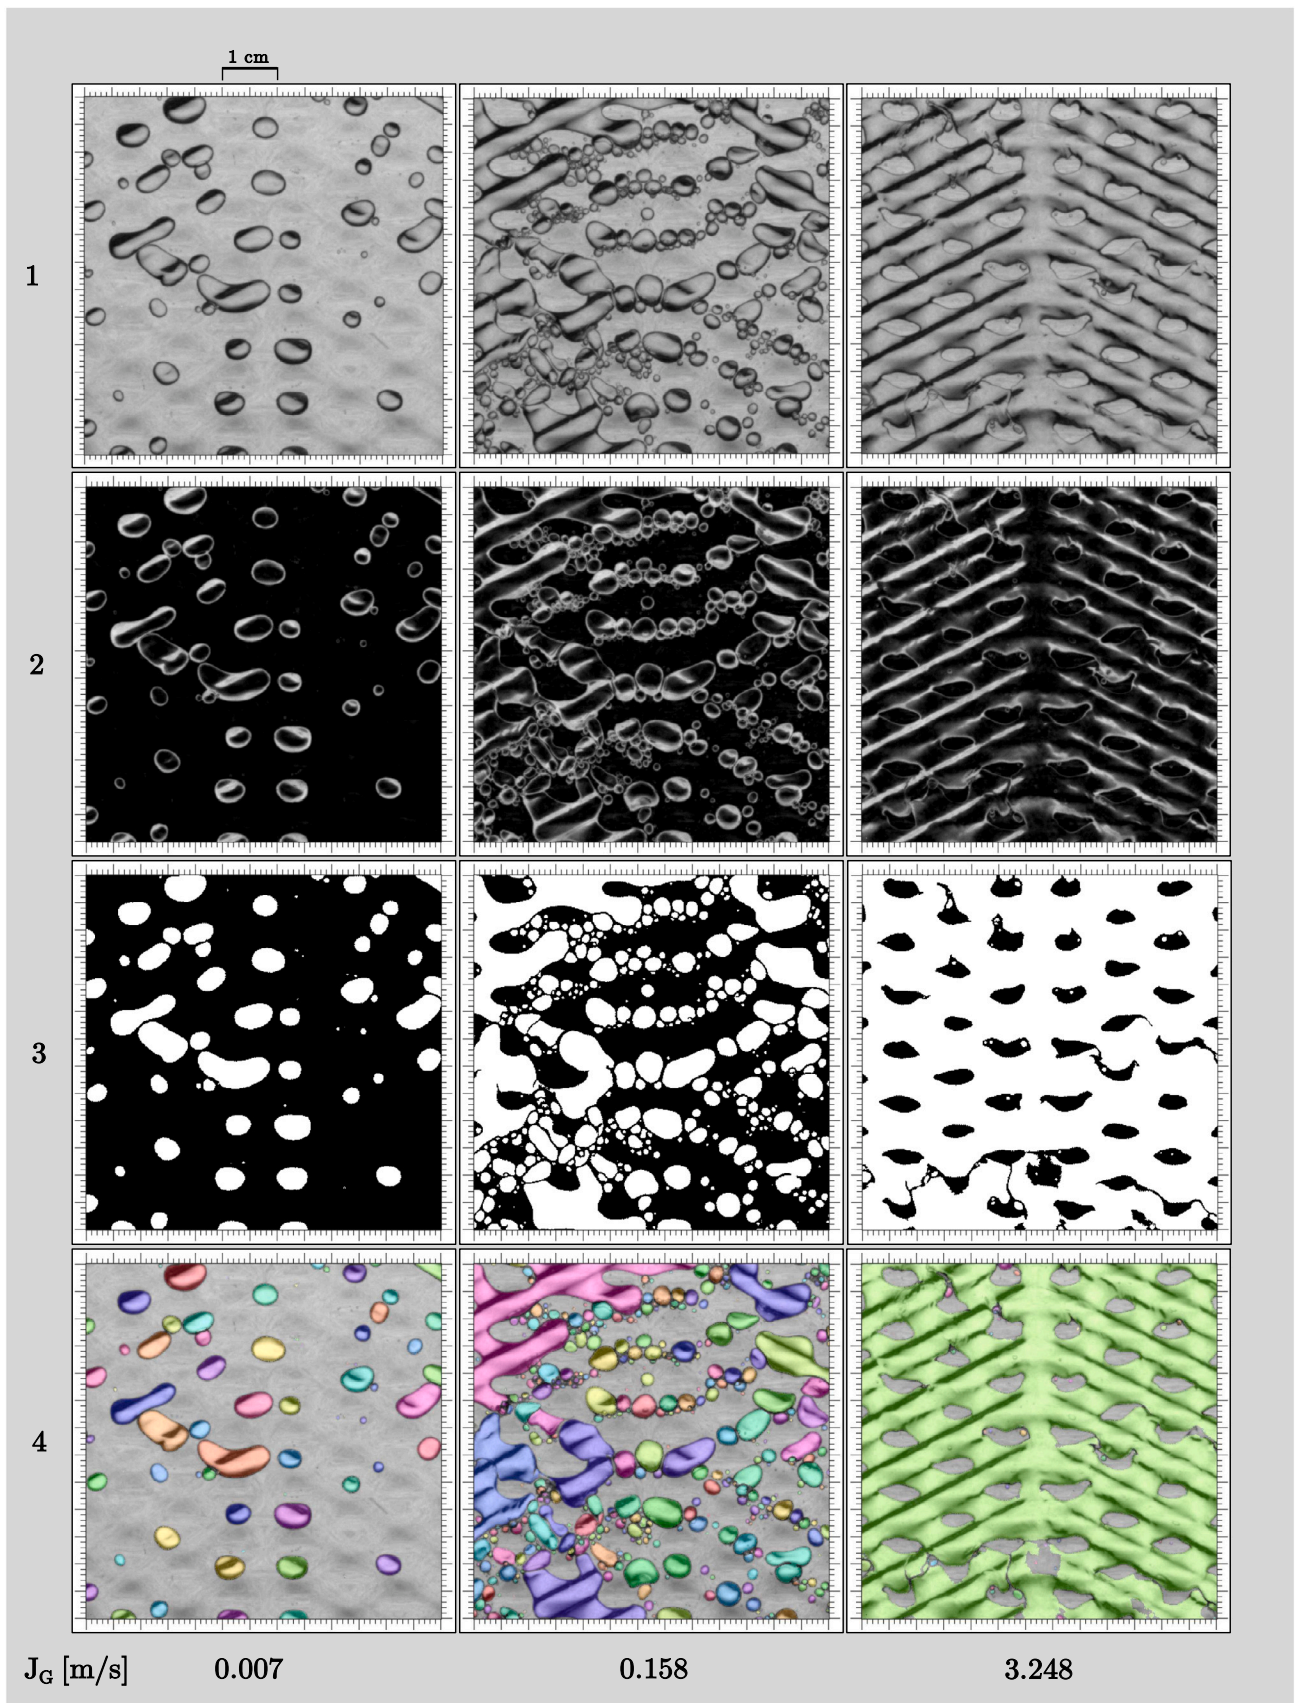


Fig. 11. Example of post-processing pipeline. Data for  $J_L = 0.006$  [m/s] at different  $J_G$  as indicated on the image. Step 1: original image, Step 2: background subtraction and contrast normalization, Step 3: final binary mask, Step 4: connected component labeling. Images cropped to fit squared aspect ratio.

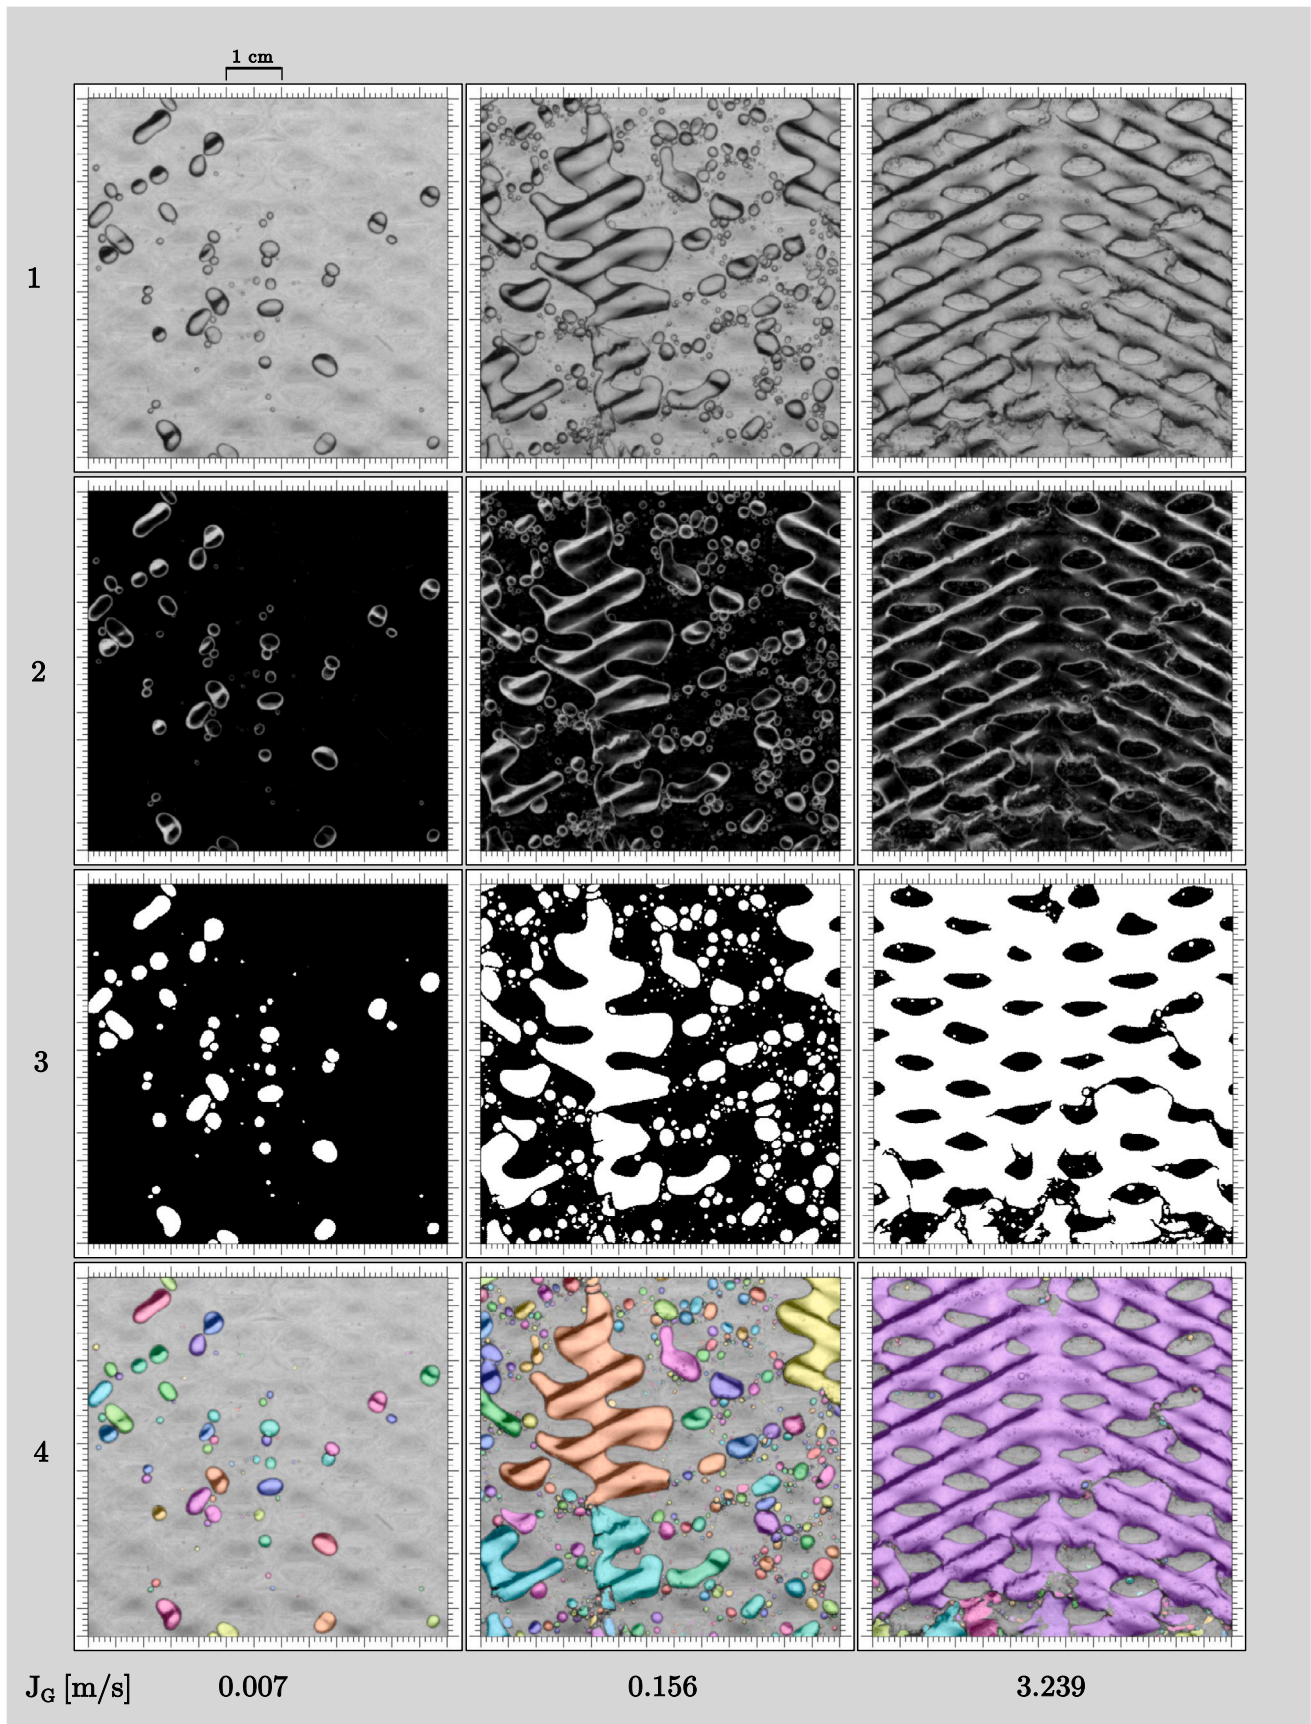


Fig. 12. Example of post-processing pipeline. Data for  $J_L = 0.061$  [m/s] at different  $J_G$  as indicated on the image. Step 1: original image, Step 2: background subtraction and contrast normalization, Step 3: final binary mask, Step 4: connected component labeling. Images cropped to fit squared aspect ratio.

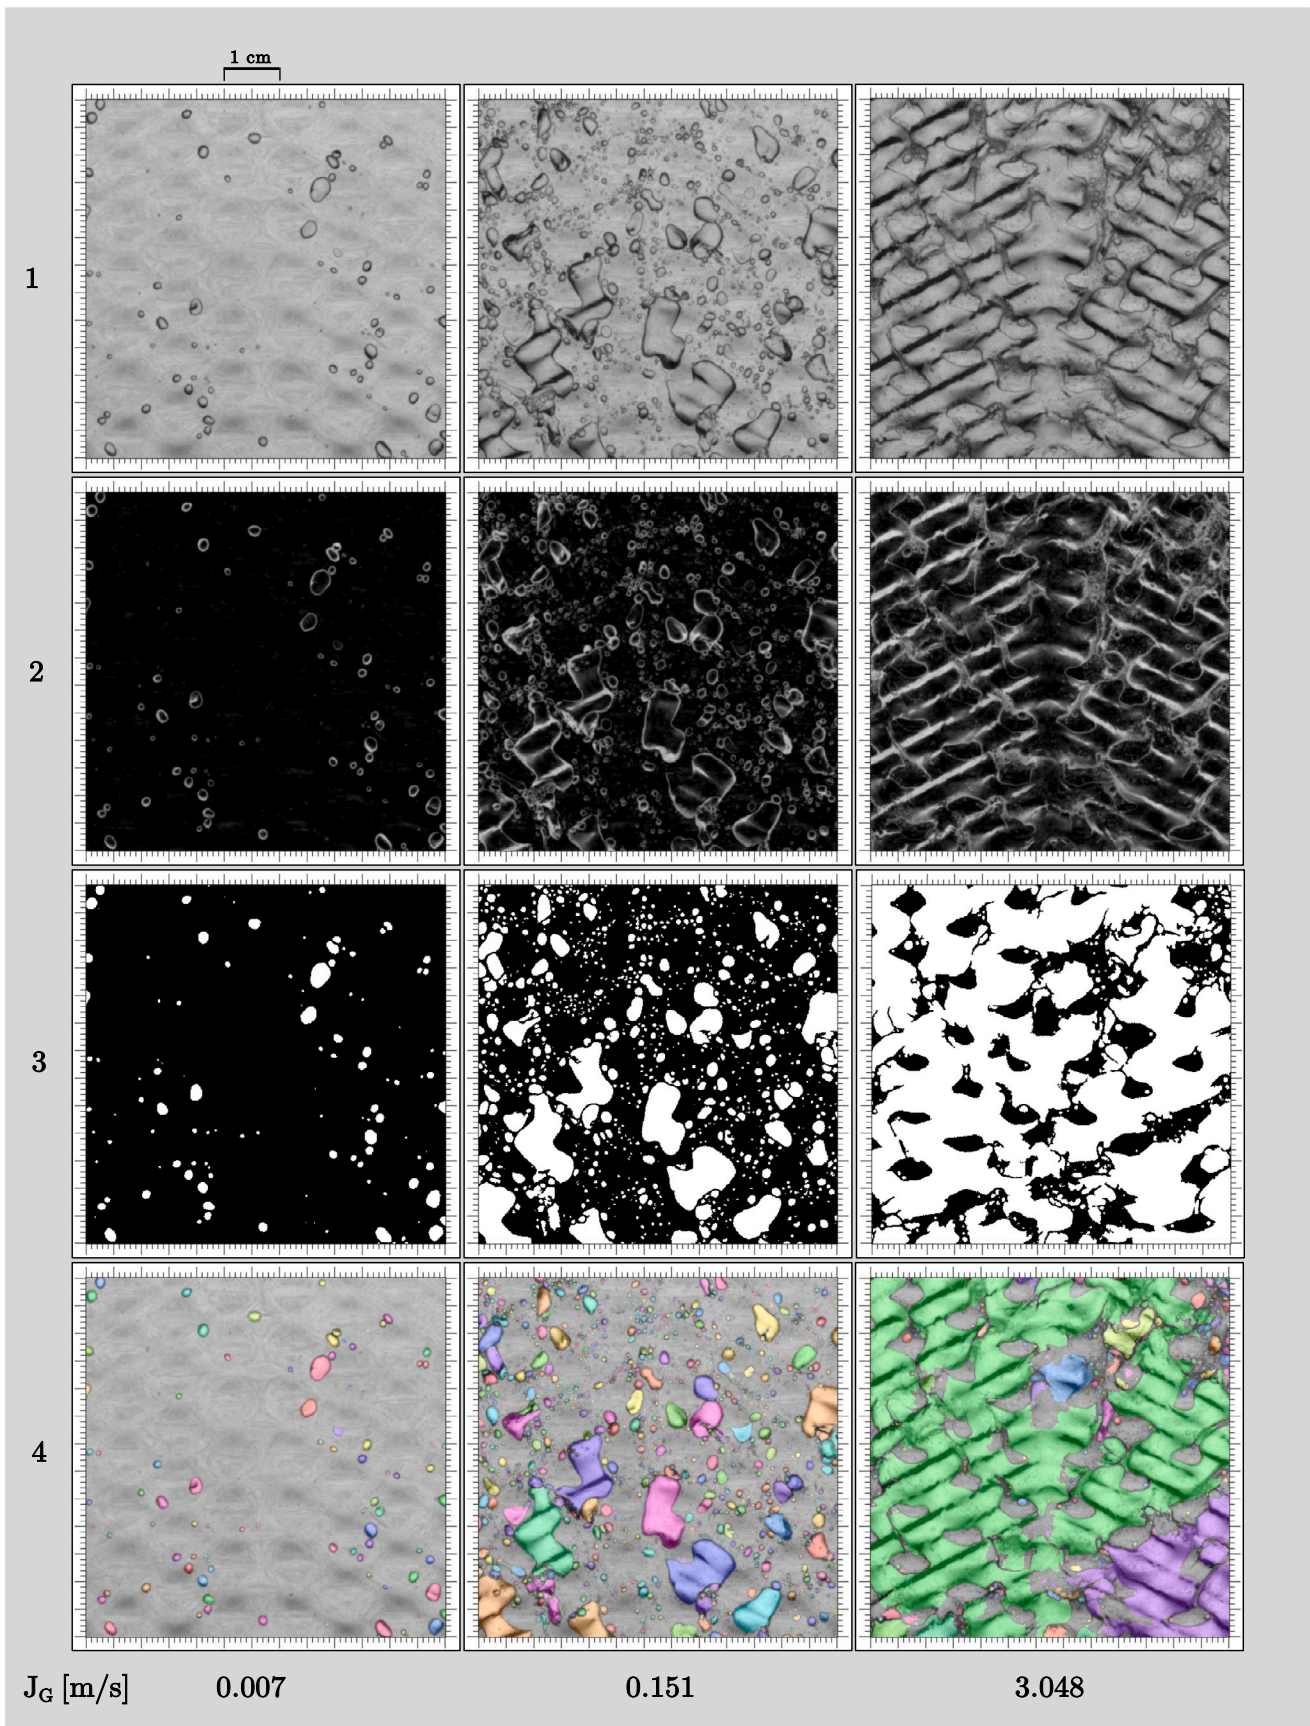


Fig. 13. Example of post-processing pipeline. Data for  $J_L = 0.183$  [m/s] at different  $J_G$  as indicated on the image. Step 1: original image, Step 2: background subtraction and contrast normalization, Step 3: final binary mask, Step 4: connected component labeling. Images cropped to fit squared aspect ratio.

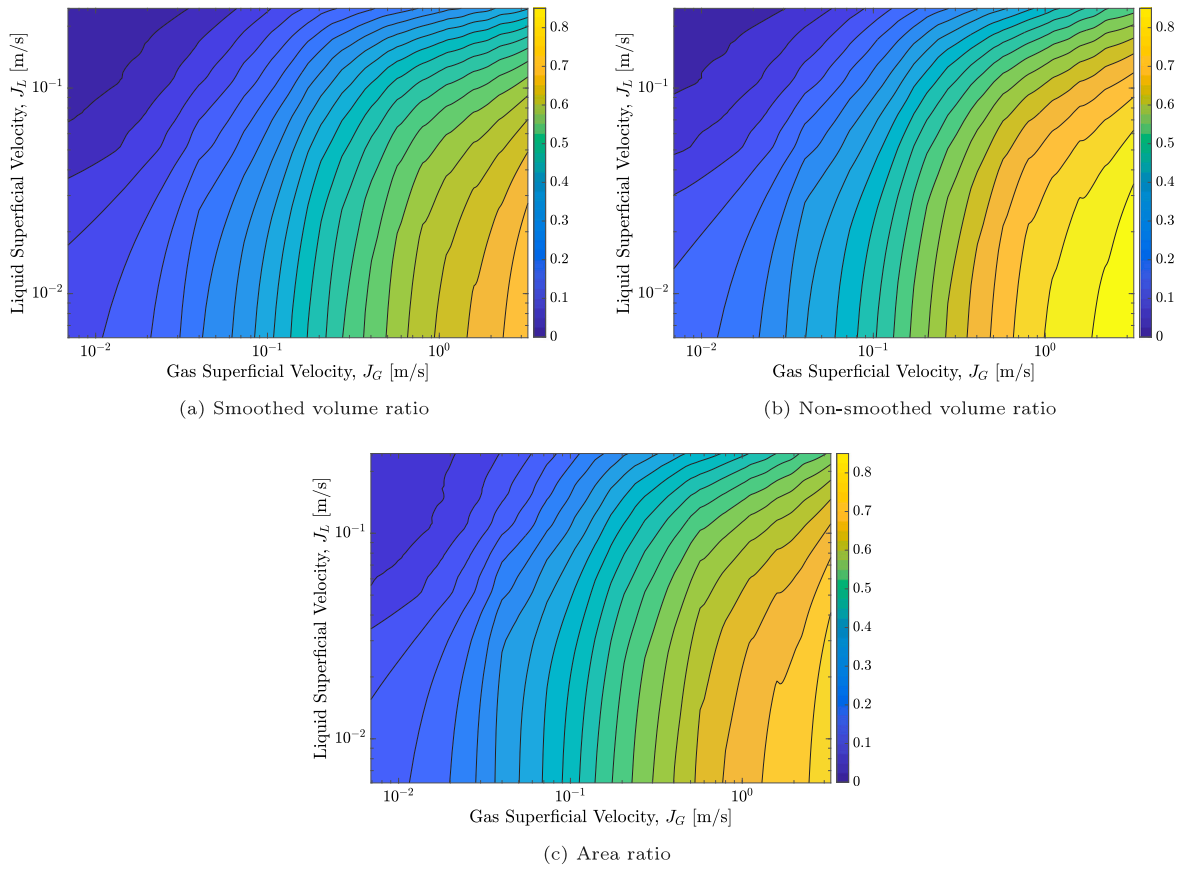


Fig. 14. Comparison between the void fraction maps obtained from image analysis by estimating the void fraction  $\alpha$  with the three different approaches discussed. All plots have the same colorbar  $\alpha \in [0; 0.85]$ .

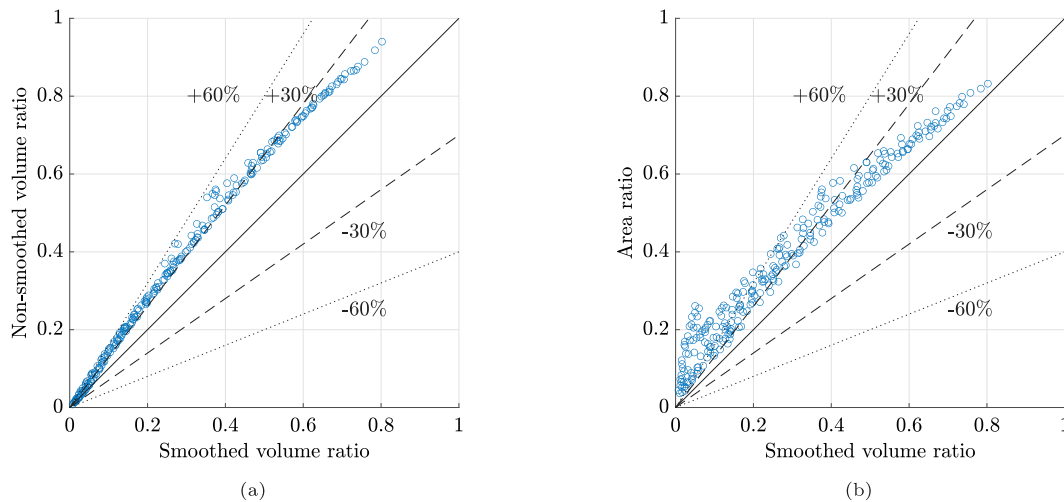


Fig. 15. Parity plot of void fraction prediction. Comparison between smoothed and non-smoothed volume ratio (a) and between smoothed volume and area ratio (b).

this behavior, two different equations are proposed to fit the database obtained using the smoothed volume ratio (i.e., Fig. 17):

$$U_G = \frac{J_G}{\alpha} = \begin{cases} 2.16J_{mix} + 0.0295 & \text{if } Fr_{\sim} \leq 1 \\ (1.25 + 4.54J_L) \cdot (J_{mix} - \sqrt{g \cdot b_p}) + 0.368 & \text{if } Fr_{\sim} > 1 \end{cases} \quad (5)$$

The first one is casted simply as a drift flux model while the second keeps into account the variability of the distribution coefficient

on liquid flow rate at corrugation Froude number greater than one. A graphical representation of the model is given in Fig. 18 where the lines corresponding to the above-mentioned equations are traced on top of the data points. The goodness of the interpolation is very high, with  $R^2$  values of about 0.94 and 0.96 for  $Fr_{\sim}$  smaller and higher than 1, respectively. Eq. (5) could be inverted to provide a predictor for the void fraction given as  $\alpha = J_G/U_G$ . When this model is benchmarked against the void fraction data obtained using smoothing of the reconstructed volumes, it yields a mean absolute percentage error

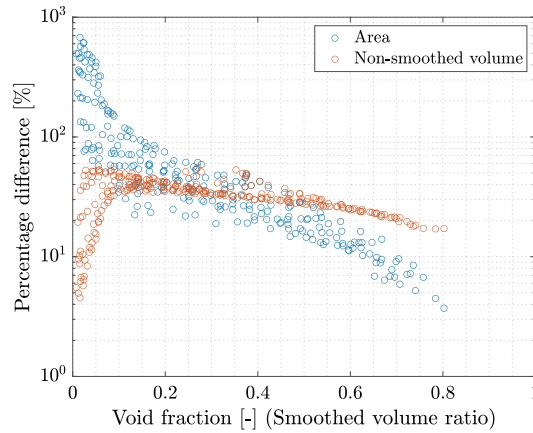


Fig. 16. Percentage difference of void fraction prediction from baseline method when using the two other approaches.

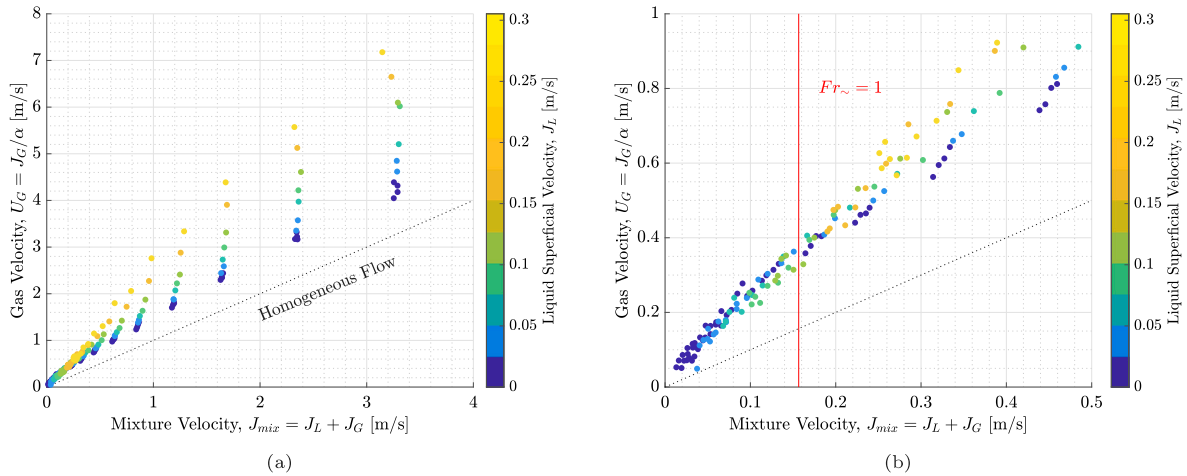


Fig. 17. Actual gas velocity against mixture velocity. Entire database (a) and detail at  $J_{mix} < 0.5$  m/s (b). Data colored with liquid superficial velocity  $J_L$  [m/s]. Void fraction data obtained using smoothed volumes ratio.

of 8.17% in predicting the parameter of interest. Fig. 19(a) displays the parity plot comparing the set of results from the model against the benchmark data.

When considering the two other methodologies for the calculation of void fraction, the behavior of the database is observed to be still the same even though some minor differences arise. It can be noticed from Fig. 20 that when reporting on the  $U_G$ - $J_{mix}$  plane the data obtained using the area ratio as a proxy for the void fraction, the results are more scattered around  $Fr_{\sim} = 1$  and some points at high liquid superficial velocities even fall below the dotted line that indicates the homogeneous flow model (representing absence of slip between phases). This evidence suggests an unphysical prediction of the void fraction, as  $\alpha$  is even higher than its homogeneous counterpart. Eqs. (6) and (7) report the fitting of the proposed drift flux equations with the two different datasets obtained with area non-smoothed volumes and area ratio respectively. By looking at the coefficients, it is of easy notice how the dependance of the distribution parameter on  $J_L$  for  $Fr_{\sim} > 1$  tend to diminish with the different approaches adopted. Figs. 19(b)

and 19(c) show the parity plot that compares the performance of the model in predicting the void fraction data obtained with these two different approaches. The mean average percentage error increase from the 8.17% of the baseline approach to 11.8% and 18.8% without the use of smoothing and when adopting the area ratio respectively.

$$U_G = \frac{J_G}{\alpha} = \begin{cases} 1.57J_{mix} + 0.0185 & (R^2 = 0.92) & \text{if } Fr_{\sim} \leq 1 \\ (1.07 + 2.58J_L) \cdot (J_{mix} - \sqrt{g \cdot b_p}) + 0.264 & (R^2 = 0.94) & \text{if } Fr_{\sim} > 1 \end{cases} \quad (6)$$

$$U_G = \frac{J_G}{\alpha} = \begin{cases} 1.38J_{mix} + 0.0297 & (R^2 = 0.78) & \text{if } Fr_{\sim} \leq 1 \\ (1.30 + 0.998J_L) \cdot (J_{mix} - \sqrt{g \cdot b_p}) + 0.245 & (R^2 = 0.97) & \text{if } Fr_{\sim} > 1 \end{cases} \quad (7)$$

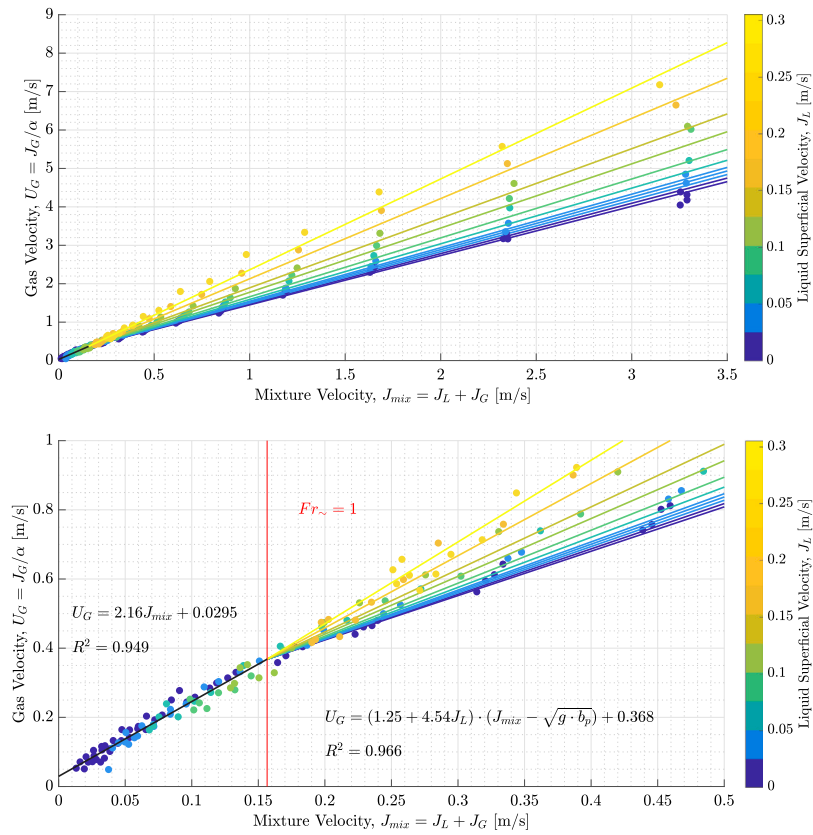


Fig. 18. Actual gas velocity against mixture velocity. Entire database on the top and detail at  $J_{mix} < 0.5$  m/s on the bottom. Lines represent the proposed model used to fit the data. Points colored with liquid superficial velocity  $J_L$  [m/s].

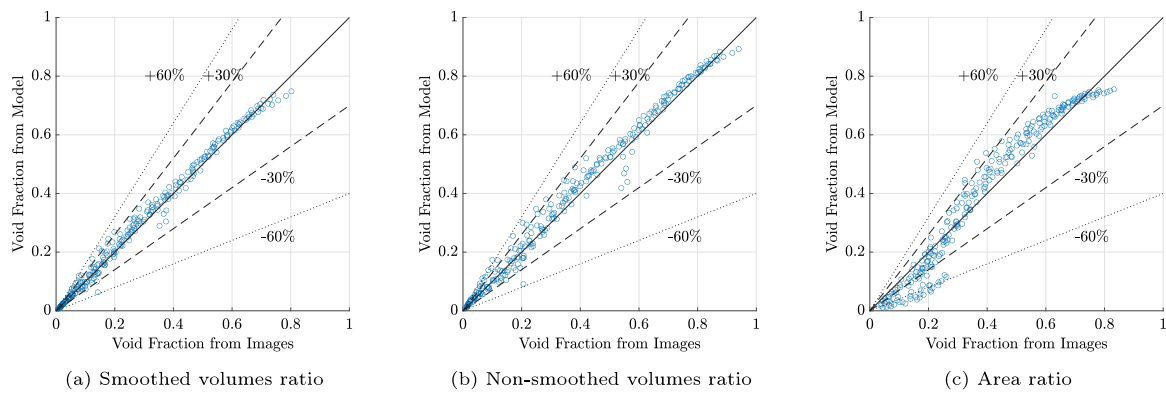


Fig. 19. Parity plot of void fraction prediction when the proposed model is fitted void fraction data obtained from the different approaches discussed in the text.

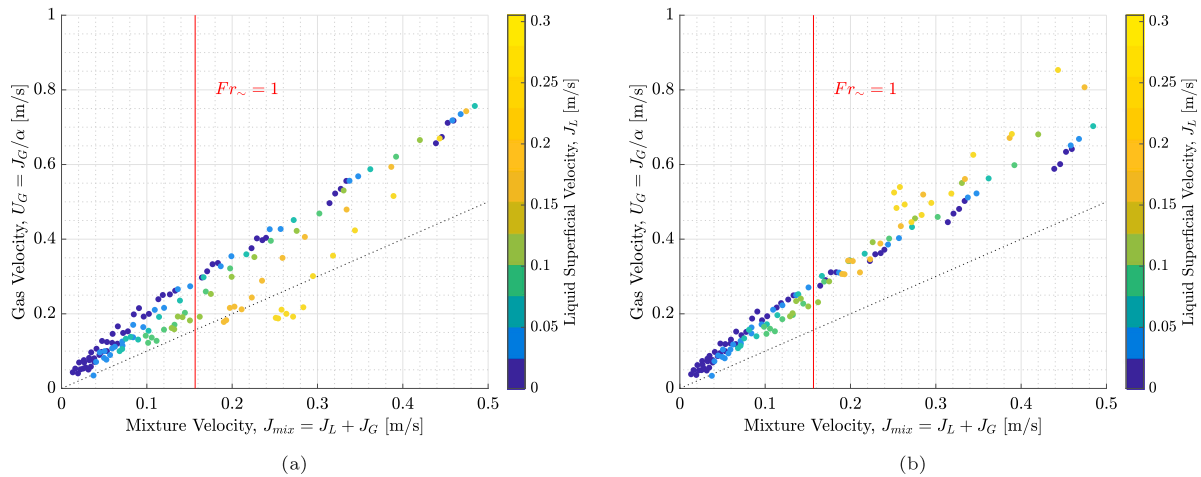


Fig. 20. Detail of actual gas velocity data against mixture velocity for  $J_{mix} < 0.5$  when using area ratio (a) of non-smoothed volumes (b) for the calculation of void fraction. Data colored with liquid superficial velocity  $J_L$  [m/s].

#### 4. Conclusions

In this work, we discussed a methodology that integrates machine learning (ML) algorithms and data processing techniques to estimate void fraction from back-lit front-view images of two-phase flow in corrugated channels. We adopted an encoder–decoder style neural network (U-Net) to perform image segmentation leveraging its transfer-learning capabilities to minimize the training set. The model was trained by exposing it to progressively more complex flow patterns with a total of 13 annotated images. The quantitative benchmark of U-Net performance in the segmentation task was numerically evaluated by computed standard metrics (accuracy, precision, recall and F1-score) that rely on the comparison between binary masks obtained from the network and “ground truth” manually annotated images. The performance showed to be very good, with a value higher than 0.9 for all the evaluation indexes. The binary masks obtained with segmentation were then processed to estimate the void fraction. We presented an algorithm to reconstruct the 3D volume of the gas clusters starting from the 2D binary masks to be able to estimate the void fraction as a volume ratio. To render a more physical bubble shape, a smoothing function was introduced to account for bubble curvature. To inquire about how sensitive this approach was, this method (i.e. usage of smoothed volume ratio) was benchmarked against the use of non-smoothed volumes and the simpler area ratio between the number of pixels labeled as gas and the total number of pixels. The results showed an agreement between all three methods in terms of the shape of void fraction contour maps. The area ratio showed to be the one with the highest discrepancy, with a mean average percentage difference close to 80%. The highest discrepancy is presented in the low void fraction region, where the use of area ratio is expected to give a strong overestimation of the void fraction. A drift flux model was proposed to fit the dataset. We introduced a two-part equation to account for the dependence of the distribution coefficient  $C_0$  on the liquid flow rate for a corrugation Froude number  $Fr_{\sim}$  larger than 1. The model proved to be a good predictor of the estimated void fraction with a mean average percentage error of 8.17% when using the proposed approach based on smoothed volume ratio.

The U-net-based approach we proposed can be applied to various two-phase flow conditions and different pseudo-2D geometries given its generality and minimal model training effort requirement. The U-net is able to provide pixel-level accuracy in segmenting the images enabling a more deep and accurate post-processing of the results. This techniques enables the user to extract a lot more information from HSV images than what is typically done with simpler segmentation algorithms that fail at providing accurate segmentation masks. Thanks to binary masks

provided by the U-net, we were able to quantify the void fraction using different approaches. The results revealed that the choice of processing technique influenced the void fraction calculation. Future work can be focused in exploiting the accurate binary masks to extract quantitative information about the different flow patterns to, for instance, provide a quantitative way to classify them.

#### CRedit authorship contribution statement

**Stefano Passoni:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Riccardo Mereu:** Writing – review & editing, Project administration, Funding acquisition, Conceptualization. **Matteo Bucci:** Writing – review & editing, Supervision, Methodology, Formal analysis, Conceptualization.

#### Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Stefano Passoni reports financial support was provided by EU Framework Programme for Research and Innovation Euratom.

#### Data availability

Data will be made available on request.

#### Acknowledgments

This work has received funding from the Euratom research and training programme 2014–2018 under Grant Agreement No. 847553. S.Passoni acknowledges the funding support of the “Progetto Roberto Rocca” doctoral fellowship. The authors acknowledge the CFDhub of Politecnico di Milano (<https://www.cfdhub.polimi.it>) for providing the computational resources needed for this work.

#### References

- Ayub, Z.H., 2003. Plate heat exchanger literature survey and new heat transfer and pressure drop correlations for refrigerant evaporators characterization of advanced heat exchanger performance of single and multi-phase flows view project.
- Buscher, S., 2019. Visualization and modelling of flow pattern transitions in a cross-corrugated plate heat exchanger channel with uniform two-phase distribution. *Int. J. Heat Mass Transfer* 144, 118643.
- Buscher, S., 2022. Digital image analysis of gas-liquid flow in a cross-corrugated plate heat exchanger channel: A feature-based approach on various two-phase flow patterns. *Int. J. Multiph. Flow* 154, 104149.

- Falk, T., Mai, D., Bensch, R., Çiçek, Ö., Abdulkadir, A., Marrakchi, Y., Böhm, A., Deubner, J., Jäckel, Z., Seiwald, K., et al., 2019. U-Net: Deep learning for cell counting, detection, and morphometry. *Nature Methods* 16, 67–70.
- Gennes, P.-G., Brochard-Wyart, F., Quéré, D., et al., 2004. *Capillarity and Wetting Phenomena: Drops, Bubbles, Pearls, Waves*. Springer.
- Grabstein, V., Polzin, A.-E., Kabelac, S., 2017. Experimental investigation of the flow pattern, pressure drop and void fraction of two-phase flow in the corrugated gap of a plate heat exchanger. *Int. J. Multiph. Flow* 91, 155–169.
- Kang, J., Bak, J.Y., Lee, B.J., Chung, C.K., Yun, B., 2022. Numerical investigation of a plate-type steam generator for a small modular nuclear reactor. *Nucl. Eng. Technol.* 54, 3140–3153.
- Nilpueng, K., Wongwises, S., 2010. Two-phase gas-liquid flow characteristics inside a plate heat exchanger. *Exp. Therm Fluid Sci.* 34, 1217–1229.
- Passoni, S., Ferrario, A., Ricotti, M.E., Mereu, R., 2024. Experimental investigation of two-phase flow in chevron-type compact plate heat exchangers: A study on pressure drops and flow regimes visualization. *Appl. Therm. Eng.* 242, 122542.
- Ravichandran, M., Kossolapov, A., Aguiar, G.M., Phillips, B., Bucci, M., 2023. Autonomous and online detection of dry areas on a boiling surface using deep learning and infrared thermometry. *Exp. Therm Fluid Sci.* 110879.
- Ravichandran, M., Su, G., Wang, C., Seong, J.H., Kossolapov, A., Phillips, B., Rahman, M.M., Bucci, M., 2021. Decrypting the boiling crisis through data-driven exploration of high-resolution infrared thermometry measurements. *Appl. Phys. Lett.* 118, 253903.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-Net: Convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (Eds.), *Medical Image Computing and Computer-Assisted Intervention. MICCAI 2015*, Springer International Publishing, Cham, pp. 234–241.
- Schindelin, J., Arganda-Carreras, I., Frise, E., Kaynig, V., Longair, M., Pietzsch, T., Preibisch, S., Rueden, C., Saalfeld, S., Schmid, B., et al., 2012. Fiji: An open-source platform for biological-image analysis. *Nature Methods* 9, 676–682.
- Schneider, C.A., Rasband, W.S., Eliceiri, K.W., 2012. NIH image to ImageJ: 25 years of image analysis. *Nature Methods* 9, 671–675.
- Seong, J.H., Ravichandran, M., Su, G., Phillips, B., Bucci, M., 2023. Automated bubble analysis of high-speed subcooled flow boiling images using U-net transfer learning and global optical flow. *Int. J. Multiph. Flow* 159, 104336.
- Shiomi, Y., Nakanishi, S., Uehara, T., 2004. Characteristics of two-phase flow in a channel formed by chevron type plates. *Exp. Therm Fluid Sci.* 28, 231–235.
- Torrey, L., Shavlik, J., 2010. Transfer learning. In: *Handbook of Research on Machine Learning Applications and Trends: Algorithms, Methods, and Techniques*. IGI global, pp. 242–264.
- Tribbe, C., Müller-Steinhagen, H.M., 2001. Gas/liquid flow in plate-and-frame heat exchangers - part II: Two-phase multiplier and flow pattern analysis. *Heat Transf. Eng.* 22, 12–21.
- Vlasogiannis, P., Karagiannis, G., Argyropoulos, P., Bontozoglou, V., 2002. Air–water two-phase flow and heat transfer in a plate heat exchanger. *Int. J. Multiph. Flow* 28, 757–772.
- Wang, Q., Li, X., Xu, C., Yan, T., Li, Y., 2021. Bubble recognizing and tracking in a plate heat exchanger by using image processing and convolutional neural network. *Int. J. Multiph. Flow* 138, 103593.
- Weiss, K., Khoshgoftaar, T.M., Wang, D., 2016. A survey of transfer learning. *J. Big Data* 3, 1–40.
- Zuber, N., Findlay, J.A., 1965. Average volumetric concentration in two-phase flow systems.