Use of ensemble modeling to predict bubble velocities using an optical fibre probe

BEP Report

Author:

Luuk Wouterse (5422752)

First Supervisor: Cees Haringa
Second Examiner: Hanieh Bazyar
Daily Supervisor: Rik Volger
Date: June 25, 2025

TU Delft Delft University of Technology

Contents

1 Introduction

Bubble columns are an essential component in a wide range of industrial processes, including chemical reactors, bioreactors, and wastewater treatment systems [1]. The design of a bubble column is simple and energy-efficient. Gas is fed at the bottom of a column containing a liquid. As the gas rises in the form of bubbles, it mixes the gas and liquid well, providing excellent heat and mass-transfer[2]. Due to their lack of moving parts and ease of scale-up, bubble columns are attractive for large-scale operations [2]. Even with the high mass-transfer a bubble column provides, the gas-to-liquid mass transfer is often still the limiting factor in the reaction [3]. To be able to optimize the mass-transfer, a good understanding of the internal hydrodynamics is needed, which can be difficult to achieve [1, 2].

Bubble size and velocity are two of the most important parameters for understanding those hydrodynamics [4]. They directly influence gas holdup, mixing intensity and therefore mass-transfer rates, which in turn directly influences product yield and efficiency. Measuring bubble size and velocity reliably is therefore of high interest, especially in systems that are visually inaccessible due to e.g. broth opacity or high holdup [4]. Optical fibre probes provide a promising technique to measure bubble size and velocity in these systems. They can detect bubbles based on changes in refractive index, showing when a bubble enters and exits the probe tip, without having visual access to the system [1]. Despite their advantages, interpreting data from optical fibre probes can be very complex. In the past, multi-tip probes were used because it is relatively easy to derive velocities, but they influence the bubbles more, changing their shape, velocity or direction [5, 4]. By using a single-tip probe, smaller bubbles can be detected and the effect on the bubbles can be minimized and therefore, more accurate data can be derived. However, deriving bubble velocity is more complicated because a single-tip cannot derive velocities from a time difference between multiple tips. The velocity has to be derived from a single signal, unfortunately this method suffers from a low detection rate, this project aims to improve the detection rate.

The signals produced by a single-tip fibre probe are often noisy and the interaction with the probe happens very quickly, making it very sensitive [6]. Previously, a Fourier transform based algorithm was used to analyze the data and provide a velocity and size, but this suffers from a low detection rate; depending on the flow rate, between 40 and 20% in air-water systems. Recent research by R. Volger et al. has shown that an ensemble model of neural networks could assign velocity and size to 51% more bubbles in a bubble column with a flow rate of 60 L/min [7]. This ensemble uses multiple neural networks to reduce variance and improve robustness. However, this model does not translate very well to other flow rates. In practice, flow rate influences the mass-transfer, and is therefore varied for different processes [8]. A model that only works well at one flow rate limits its use to only that flow rate, while other flow rates might be just as, if not more important. Therefore, generalization across more flow rates is essential for making a truly useful model.

This leads to the main research question: Can we improve the detection rate of bubble velocities with a fibre probe using an ensemble of neural networks and make it able to assign velocities across a wide range of flow rates in bubble column reactors? To address this, the models will be trained on a more diverse dataset that covers multiple flow rates, and the ensemble will be expanded with additional machine learning (ML) models to enhance both prediction accuracy and detection rates.

2 Experimental

2.1 Fibre Probe Measurement

The fibre probe is a tool for measuring bubble velocity in multiphase flows, such as those found in bubble column reactors. It operates based on the difference in refractive index between gas and liquid phases by emitting light and measuring the light that gets reflected with a light sensor, which results in a voltage signal [1]. The probe has a tip with a diameter of around 8 μm that pierces bubbles [1]. When a bubble gets pierced, it results in a difference in voltage. A typical signal is shown in Figure 1.

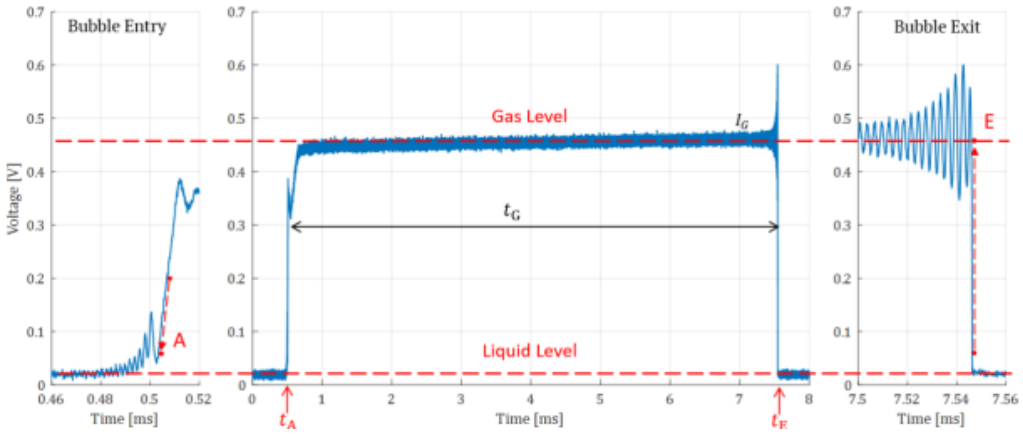

Figure 1: A typical bubble signal by Lefebvrea et al.[1]

These voltage signals contain information about the dynamics of the bubble, including the velocity and residence time. At the start of the signal the voltage is low, which indicates the probe is in water. When the probe enters the bubble there is a quick rise in voltage, which can be seen at around 0.5 ms. Then there is a period when the probe is in the bubble, where the voltage remains high. After that there is a decrease in voltage back to the liquid level, which is the bubble exit [1, 6]. This can be seen here at around 7.5 ms. The difference between the entry and exit peak is the residence time [1, 6]. The information about the velocity of the bubble can be extracted from the frequency of the oscillations in the signal because of the Doppler shift [1]. Notably, oscillatory patterns are most clear at the bubble's entry and exit, corresponding to abrupt changes in the measured signal.

Bubbles are slowed down by the probe, the assumption is that the most of the loss in velocity is accounted for by the initial impact. Although the exit velocity, captured at the exit peak, is not the the most representative of the actual bubble velocity, it is considered more representative of the bubble's velocity while interacting with the probe and thus more useful to determine the bubble size, which we ultimately want to know. For this project, experimental runs are used where the fibre probe was installed in the center of a bubble column with an inner diameter of 19.2 cm. The fibre probe from A2 Photonic Sensors [9] was used to gather data at 20.833 MHz. The runs were analyzed using a Fourier transform based algorithm supplied by A2 Photonic Sensors to determine bubble velocities. Additionally, the raw bubble signal data was stored and used for the ML approach developed in this project.

2.2 Data Processing

A pandas(2.2.3) DataFrame was constructed using python(3.13.2) for which seven, out of a total of 20, experimental runs at various flow rates were selected, ranging from 30 to 240 L/min, which correspond to a superficial velocity range of about 0.017 to 0.14 m/s. First, we detected the bubble exits and the region of interest got extracted from the rest of the voltage data. The 600 data points before the bubble exit were selected as the region of interest, which corresponds to a signal length of around 30 μs. The current ensemble focuses only on determining exit velocities, therefore bubble entry signals and residence time is not used currently. This was used for all models and is also applied when using the ensemble. An example voltage signal is shown below in Figure 2. This bubble was assigned an exit velocity of 1.46 m/s by the original method.

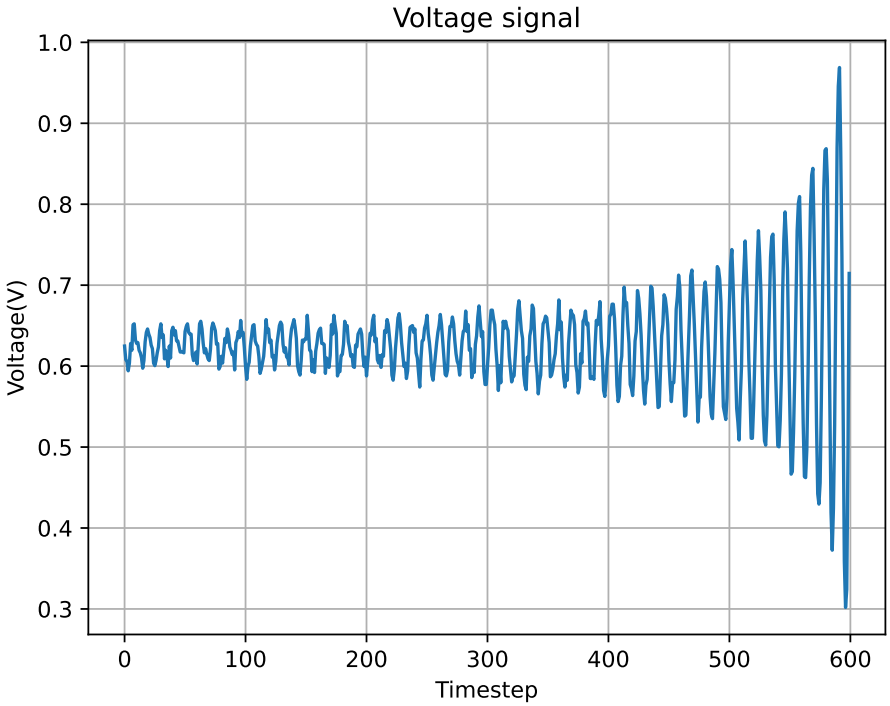

Figure 2: An example Voltage signal

The unlabeled data got separated from the data with labels. The labeled data was saved for model training. The labeled dataset was split into training, validation, and test sets using a 67.5%, 22.5%, 10% ratio, respectively. This split was chosen to ensure that a large portion of the data was available for training, which helps the models learn underlying patterns more effectively. The validation set was used to monitor performance during training and to fine-tune hyperparameters. While the test set was not very important, since runs that were not selected could serve as test sets, a small portion was kept apart as a precaution, in case an internal test set was needed. The train, validation and test set consist of a voltage signal and the label, bubble velocity. Then the train set was scaled using StandardScaler from scikit learn(1.6.1), so the mean was 0 and the standard deviation was 1. The validation and test set were scaled using the same scaler as the train set.

The train set contained a total of 3064 labeled bubbles. Random noise was added to decrease overfitting and in the hope of the model being able to predict more noisy bubbles that the original method could not predict. A randomly selected 30% of the data was duplicated,

random noise got added to the data points with a mean of 0 and a standard deviation of 0.05 for all models. After adding random noise, the train set contained 3983 bubbles.The distribution of the bubble velocities in the train set is shown in Figure 3.

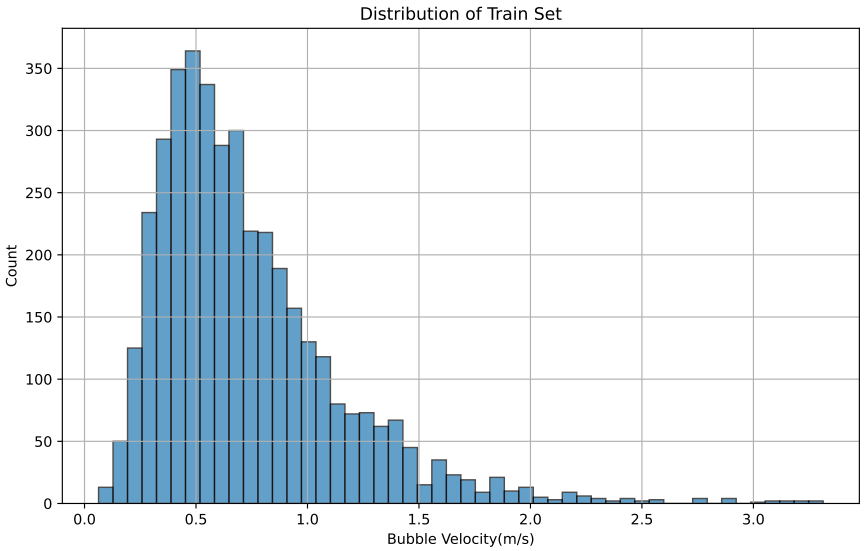

Figure 3: Distribution of the train set

2.3 Model Training

For the ensemble, two Gated Recurrent Unit (GRU) models, one Long Short-Term Memory (LSTM) model and two Convolutional Neural Network (CNN) models were selected. These architectures were chosen for their complementary strengths in modeling time-series data. The architectures that were used for these models are shown in Figure 4. The LSTM model is a Recurrent Neural Network (RNN) that uses an input, output and forget gate to memorize important features and forget less important features. With its long memory capabilities, it is well-suited for capturing long-term temporal dependencies and lower-frequency trends [10, 11]. A GRU model is also a RNN and is similar to a LSTM, but it uses only two gates, an update and reset gate. The simpler gating mechanism, is effective in modeling shorter-term dependencies and more transient patterns [11]. A CNN model employs stacked 1D convolutional layers with local receptive fields, activation functions, and pooling operations to extract features from input sequences. They excel at detecting localized patterns and high-frequency features due to their convolutional receptive fields [12]. By combining these diverse model types, the ensemble aims to improve robustness and generalization by leveraging both global and local characteristics of the input signal [13].

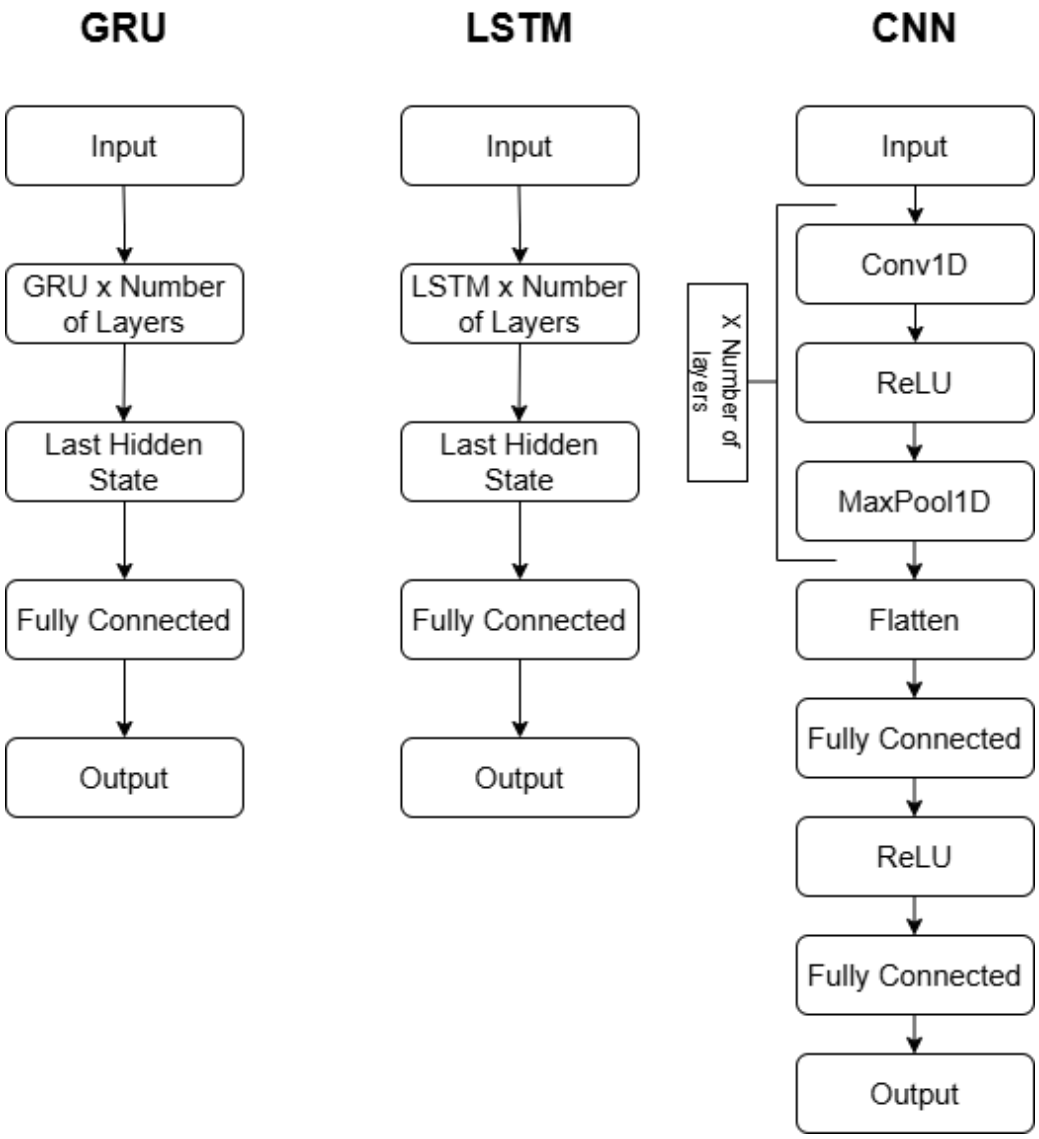

Figure 4: Architecture of the GRU, LSTM and CNN

The models were trained using PyTorch(2.7.0) using a NVIDIA GeForce RTX 4090 GPU and the method described by Bai et al. [14]. The parameters that were varied for the GRU and LSTM models are: Hidden size, number of layers, learning rate and the amount of epochs. For the CNN models, the parameters that were varied are: Hidden units, kernel size, number of layers, learning rate and the number of epochs. The train and validation loss were tracked during training to ensure the model was not overfitting, also the epoch with the least validation loss was selected as the final model. The parameters were determined by trial and error, the results are shown below.

Table 1: Hyperparameters of all models used in the ensemble.

Model	**Type**	**Hidden Size/Units**	**Kernel Size**	**Layers**	**Learning Rate**	**Epochs**
Model 1	GRU	20	–	2	0.004	7500
Model 2	GRU	10	–	3	0.005	15000
Model 3	LSTM	18	–	2	0.005	50000
Model 4	CNN	32	9	7	0.0016	10000
Model 5	CNN	64	15	5	0.001	10000

To make the ensemble, the predictions of the trained models shown above are averaged for each bubble, resulting in the final prediction. The prediction is accepted if the standard deviation is less than 10% of the prediction and thus rejected if it is bigger than 10%. The idea of this method is that each model makes different mistakes, if one model would be taken separately, it would determine a velocity for every bubble with 100% certainty. If the ensemble determines velocities for every bubble with high certainty, we cannot trust the predictions, since it should not be able to assign a velocity to every bubble, knowing that the training set contains mainly typical signals and a complete run does not. So if the ensemble rejects a portion of the bubbles because the predictions of the separate models do not align, we can trust the ones it does accept more.

3 Results

3.1 Validation

To evaluate the ensemble model, an independent test set was used, consisting of four runs with different flow rates of 40, 80, 130 and 200 L/min. This range was selected to test if the model performs well at various flow rates. On this test set, the ensemble predicted 98.0% of the labeled bubbles within 10% of their true velocity and with an uncertainty of less than 10%. Figure 5 compares the distribution of the ensemble's predictions on the labeled data compared to the true velocities determined by the original method.

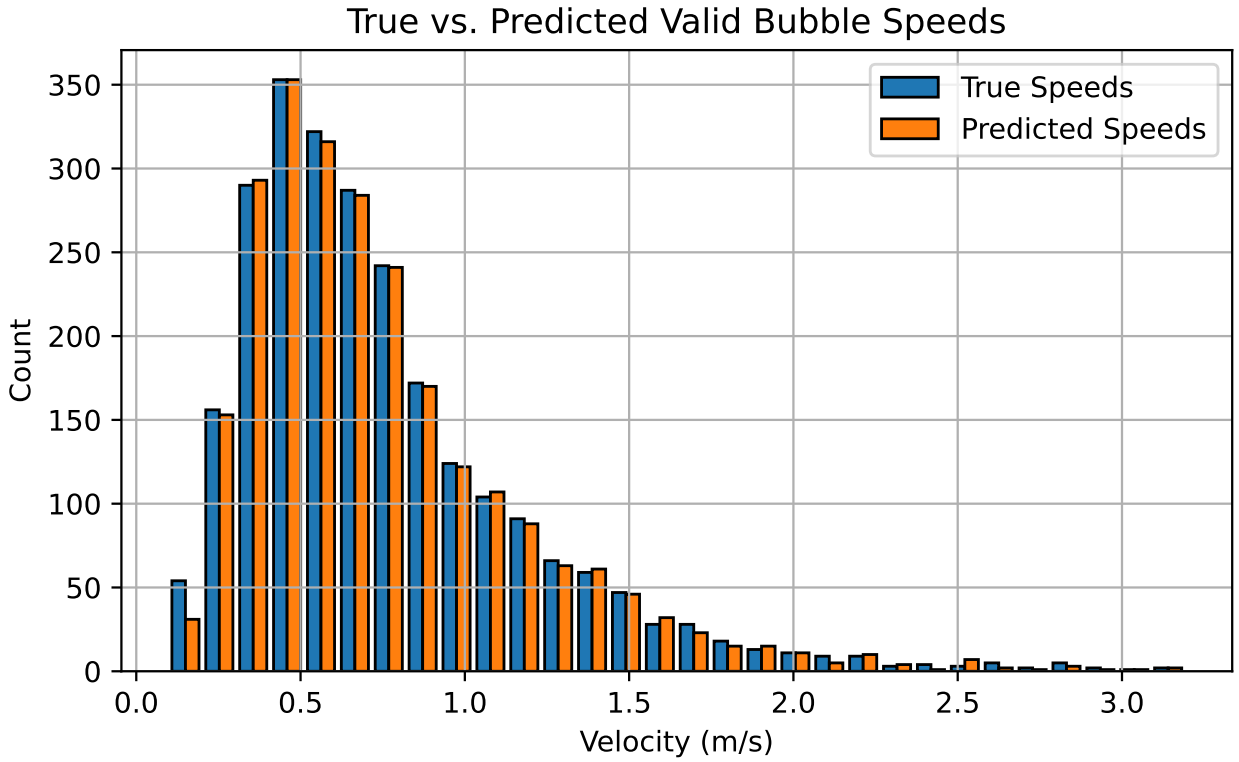

Figure 5: Comparison of the distributions of true and predicted velocities

Figure 5 shows that the distribution of the ensemble's predictions closely matches that of the true velocities, indicating a good performance on the test set. However, the plot also shows that the ensemble struggles in predicting the lowest and highest velocities. This is expected, since these groups are the least common within the training set. To gain more insight into the ensemble's performance, the predictions are shown in a parity plot in Figure 6.

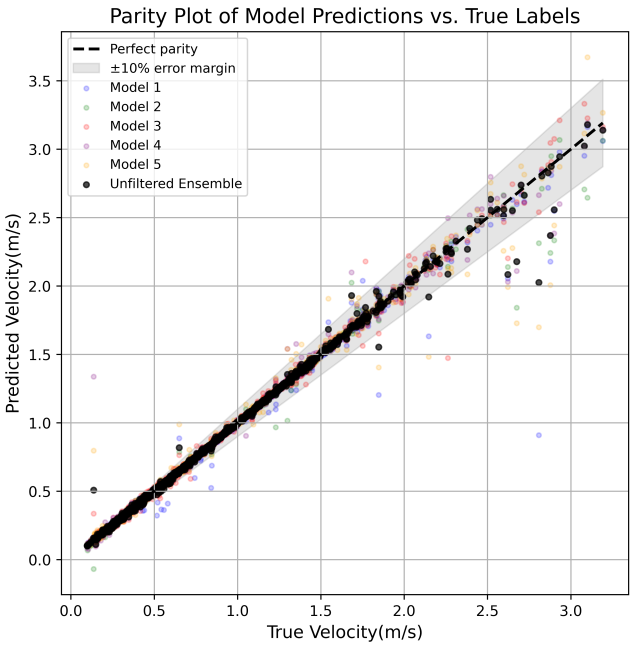

Figure 6: Parity plot of all models and the ensemble

Figure 6 shows that most predictions are within the 10% margin and tells us that the ensemble performs very well overall. However, the plot also shows that the largest errors occur at low and high bubble velocities. At low velocities the ensemble mostly overestimates and at high velocities, the ensemble underestimates. To gain a better insight into how each model performs on it's own, the parity plots of each model and the ensemble are shown separately in Figure 7.

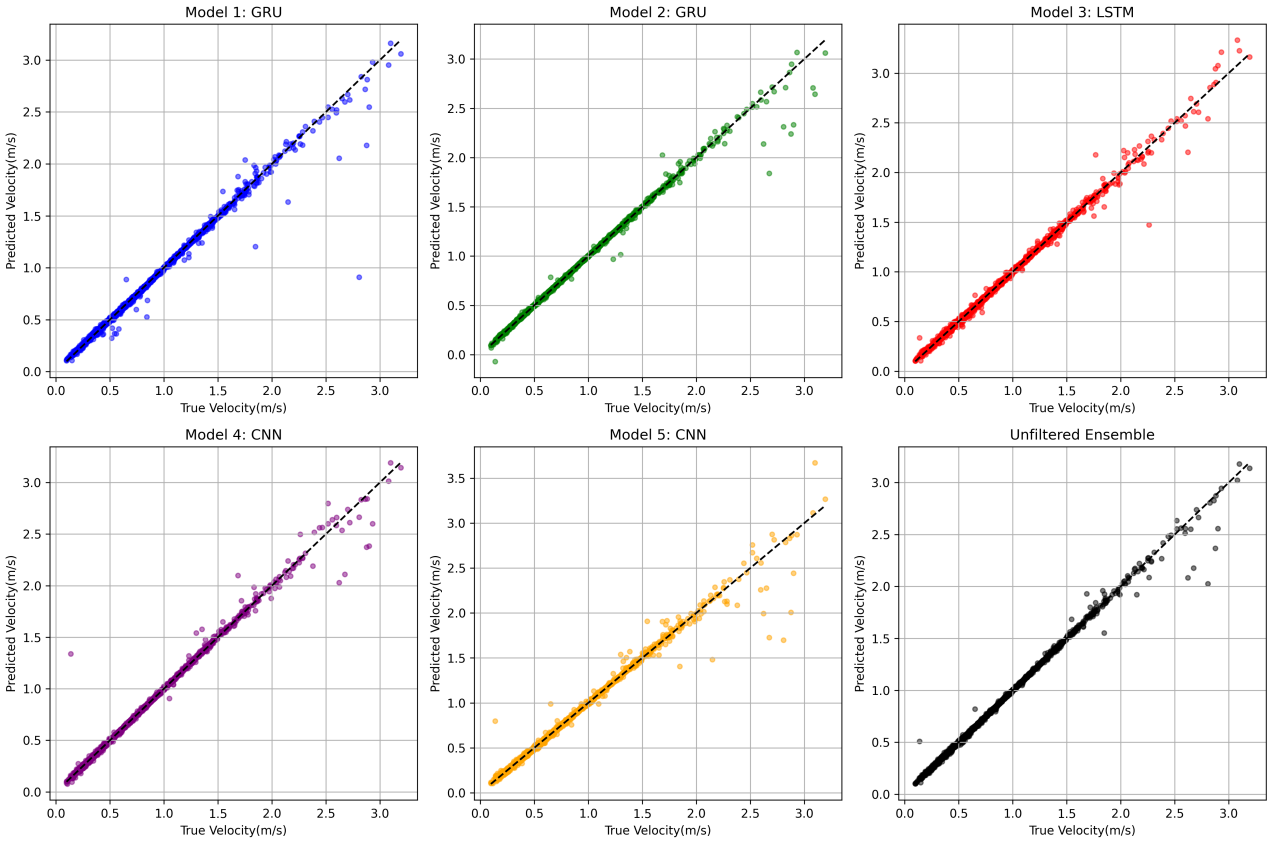

Figure 7: Parity plot of all models and the unfiltered ensemble showed separate

These plots show that each model accurately predicts most velocities, but generally tends to underestimate at higher velocities. Only the LSTM model does not have this struggle, however, it is a bit less precise. This could be due to the long-term memory of the LSTM, making it more capable of looking at the bigger picture, but lacking the ability to be precise. Even though it does not always underestimate these bubbles, it still does not predict with very high accuracy in this region.

The ensemble shown Figure 6 and Figure 7 are both the unfiltered version. Figure 8 shows the predictions of the ensemble that are accepted (uncertainty $< 10\%$).

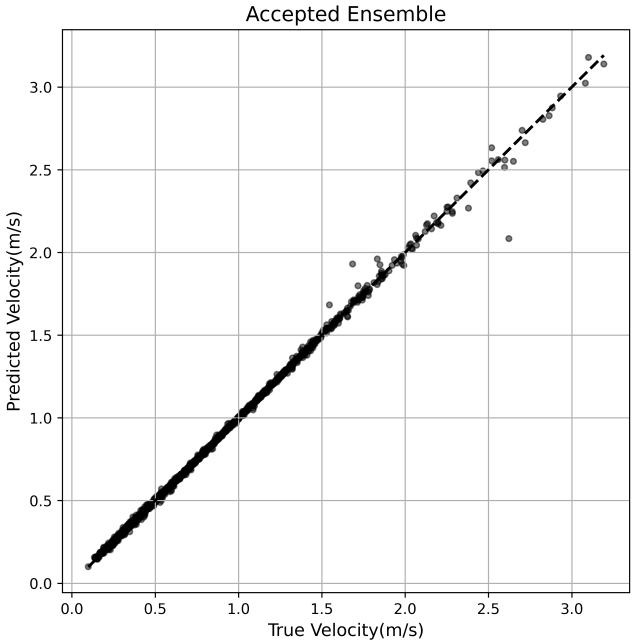

Figure 8: Parity plot of the accepted bubbles

Figure 8 shows that model performs very well on the labeled data, apart from some exceptions. When comparing Figure 8 to the unfiltered ensemble shown in Figure 7, we can see that the ensemble rejects most of the predictions that deviate a lot from the true velocity.

3.2 Detection Rate

Having evaluated the model's performance on the labeled data, the next step is to evaluate it's performance on the full dataset, which contains unlabeled data as well. The ensemble has a detection rate of 50.1% on the test set. This is an improvement of 93.3% compared to the detection rate of 25.9% of the original method. The distribution of the accepted bubbles, compared to the bubbles the original method could predict is shown in Figure 9.

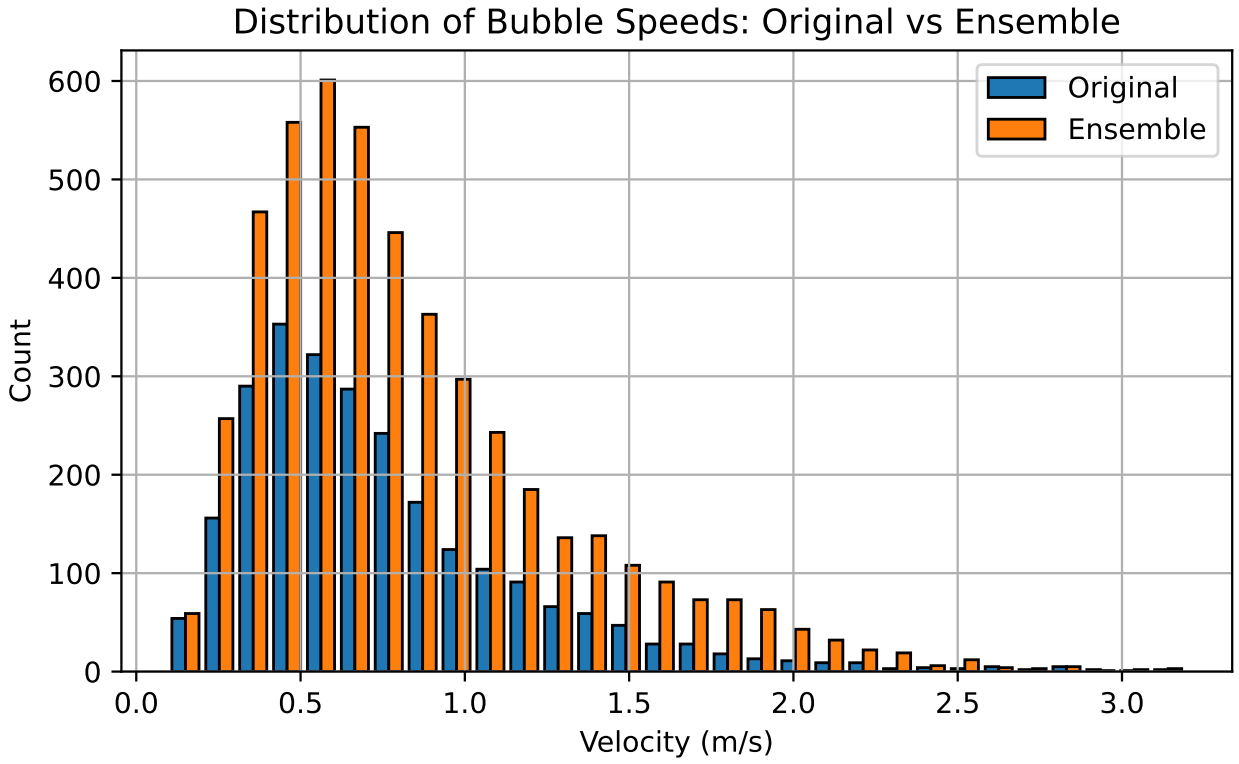

Figure 9: Comparison of the distribution of the labeled data and the ensemble predictions

The distribution of the predicted bubbles closely resembles that of the labeled data, which is desirable as it increases confidence in the model's predictions. However, the plot also indicates that the ensemble barely predicts velocities above 2.5 m/s. It was tried to flatten the data by calculating bins that contained the least amount of bubbles, duplicating these signals and adding random noise to the duplicates. This was done to try to be able to predict more higher velocities, but that led to some unexpected results, as shown in Figure 10.

Even though the distribution of the ensemble closely resembles that of the labeled data, Figure 9 shows that there is a slight shift in the distribution towards higher velocities. This could have multiple reasons; it could be that the ensemble tends to predict velocities more towards the mean of the labeled data, which is 0.59 m/s for the test set. this does align with the peak of the distribution shown in Figure 9. Another reason could be that it has to do with the accept and reject method of the ensemble. By describing the standard deviation as a percentage of the final prediction and using that to accept or reject the prediction, predictions at lower velocities are rejected with lower standard deviation values.

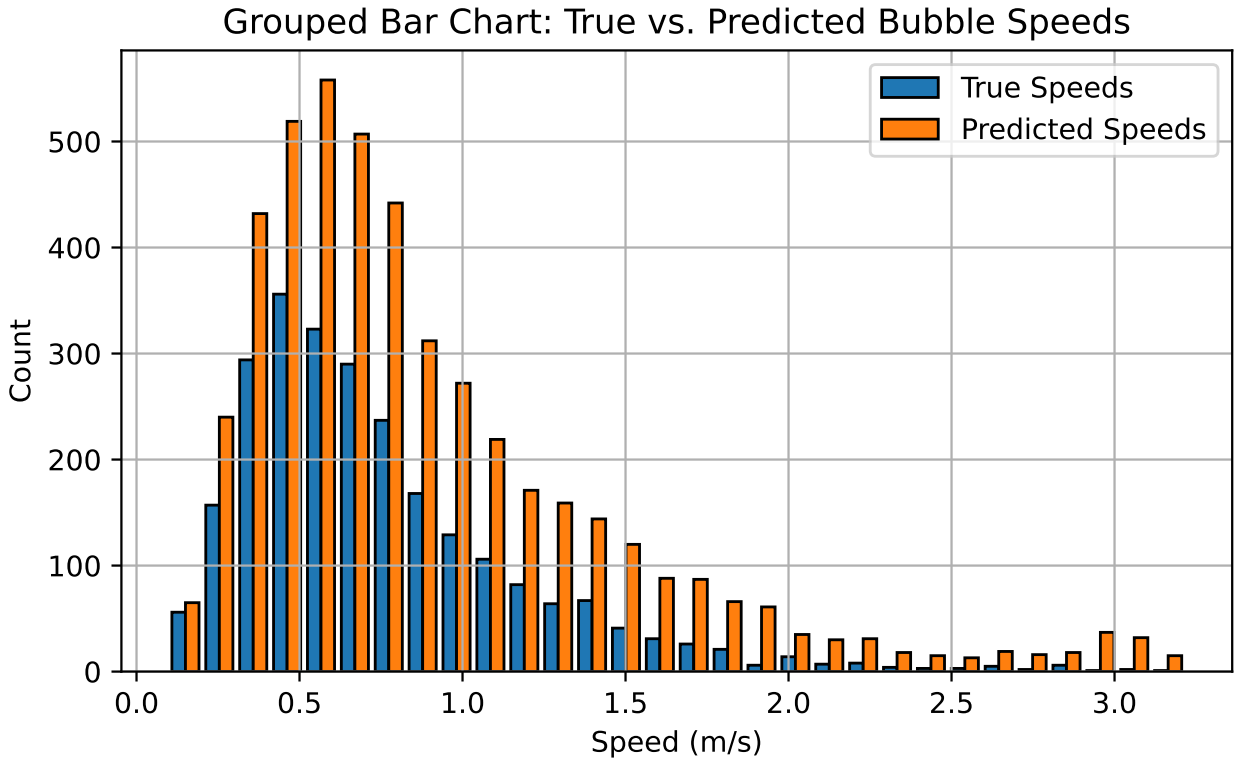

Figure 10: Distribution when flattening was used

While the ensemble trained with flattened data was able to predict more high-velocity bubbles, the resulting distribution revealed an unusual peak between 2.5 and 3.0 m/s. This does not align with the distribution of the labeled data. A possible explanation for this is that the model is overfitting on the generated data of higher velocities. Since this is not the expected distribution, it was decided that it was more reliable to use the model without data flattening.

3.3 Signal Analysis

To better understand the ensemble's performance on bubbles that the original method could not predict, the voltage signals of these previously unpredicted bubbles were examined. Figure 11 presents eight example signals for which the original method failed to assign a velocity, but the ensemble successfully assigned a velocity and Figure 12 shows eight example signals that both the original method and the ensemble failed to assign a velocity.

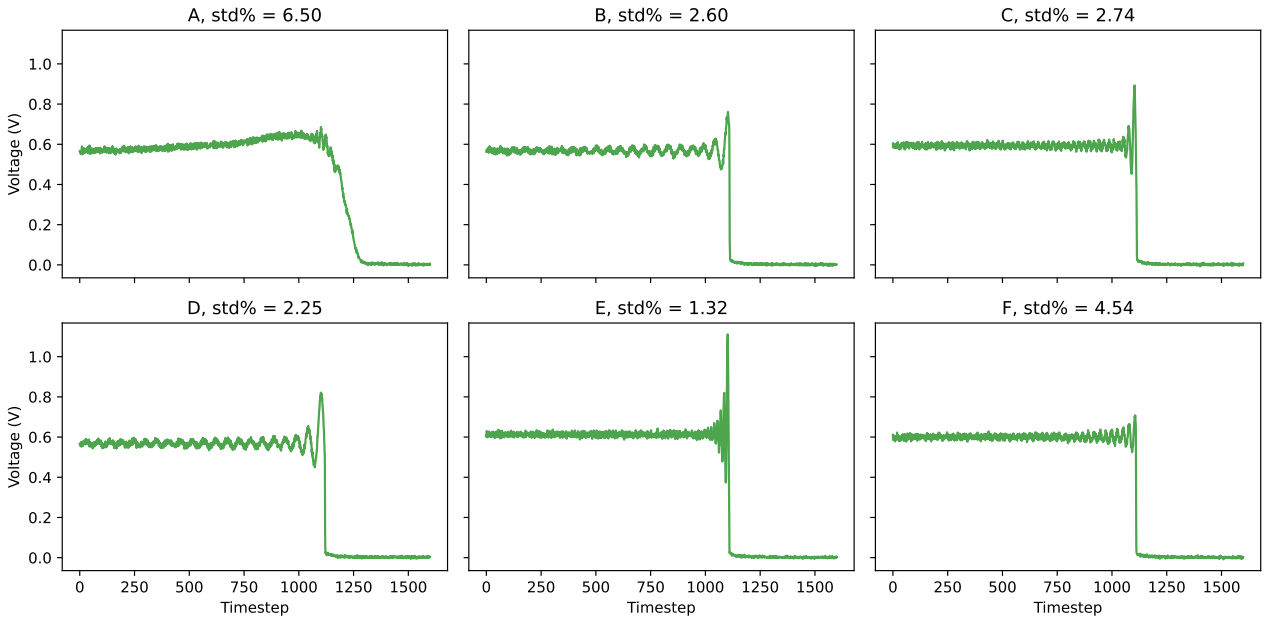

Figure 11: Bubbles that the ensemble could predict but the original method could not

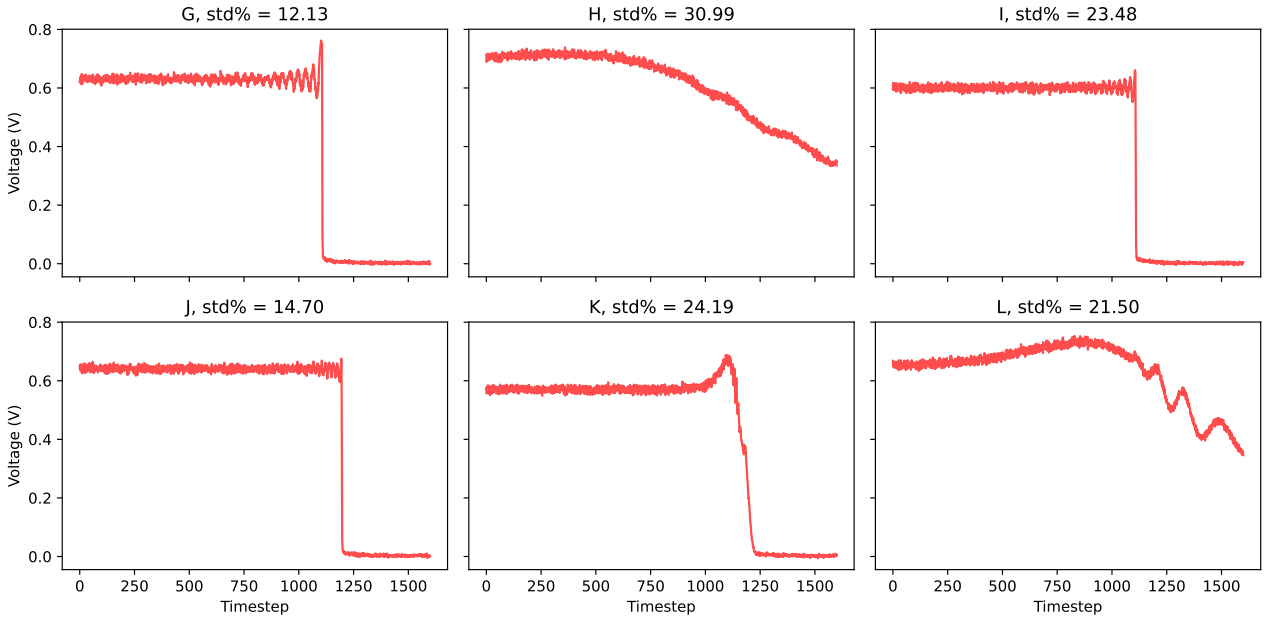

Figure 12: Bubbles that the both the ensemble and original method could not predict

Most of the accepted bubble signals showed clear oscillations like the five bubbles in Figure 11. These signals appear to contain meaningful patterns from which bubble velocity can be determined. However, there is also a bubble(A), that does not look very much like a typical signal as described in Figure 1. In some cases the models in the ensemble agree on non-typical signals too.

In contrast, Figure 12 shows examples of bubble signals that neither the original method nor the ensemble was able to predict. Most of the signals are noisy and do not look like the typical bubble signal. These are the types of signals that are expected to be rejected by both the original method and the ensemble. However, some signals, like signal G, do show oscillations,

yet the ensemble failed to assign a velocity. This suggests that even when certain signal features are present, they may not be clear enough for the ensemble to make a confident prediction.

When combining Figure 11 and Figure 12, it becomes clear that the ensemble generally accepts bubbles with signals that appear predictable and rejects those that are more complex or noisy. While there are occasional exceptions, most decisions seem reasonable. This indicates that the model is likely relying on meaningful signal features to make its predictions. This conclusion is also supported by the ensemble's 98.0% accuracy on the labeled dataset. It could be tried to change the acceptance margin, for example the margin could be increased to 12.5%. This would mean that bubble G would be accepted, which would seem reasonable, but more non-typical bubbles like bubble A would be accepted too. The margin could also be changed into the other direction, to 5% for example. This would mean that bubble A would be rejected, however, more typical signals would be rejected as well.

3.4 Detection Rate at Various Flow Rates

Besides improving the detection rate, another goal of the project was to make the ensemble predict bubble velocities in runs at various flow rates. To gain more insight into this, the ensemble was tested on runs with various flow rates separately, the runs that were not used for model training. The results of each separate set are shown below in Table 2. The accuracy, detection rates and improvement are also plotted in Figure 13.

Table 2: Performance of the model on various flow rates

Flow rate(L/min)	**Accuracy(%)**	**Detection Rate AI(%)**	**Detection Rate original(%)**	**Improvement(%)**
20	97.6	58.8	38.7	51.7
40	98.3	55.1	33.8	63.1
50	97.1	56.1	34.2	64.1
70	95.5	49.8	29.7	67.9
80	97.8	52.9	30.3	74.6
100	98.2	50.5	28.4	77.9
110	97.8	52.3	26.4	98.4
130	99.0	48.9	22.0	121.9
140	98.0	50.4	21.1	139.1
160	97.1	48.8	21.8	123.9
170	97.6	47.3	22.4	111.0
190	98.3	47.3	22.0	115.1
200	97.0	46.3	21.6	113.7

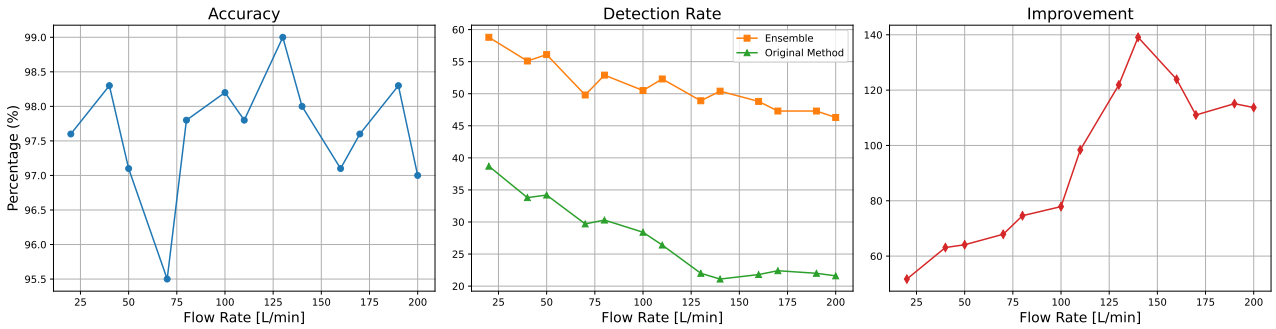

Figure 13: Accuracy, Detection Rates and Improvement plotted vs. Flow Rate

Table 2 and Figure 13 shows that the ensemble's accuracy on the labeled data is high at every flow rate. Regarding the detection rate, there are decreases in both the ensemble's and the original method's detection rates when the flow rate is increased. However, The decrease in detection rate for the original method is much steeper compared to the detection rate of the ensemble. This results in a big increase in detection rate improvement.

4 Discussion & Conclusion

The ensemble has shown an improvement in the detection rate of 93.3% across a wide range of flow rates with a high accuracy on the labeled data. These results hint that the use of ensemble modeling has the ability to boost the detection rate of an optical fibre probe. Although the ensemble does show great improvement in the detection rate, the performance on the unlabeled data is not yet validated. Figure 11 and Figure 12 support that the ensemble focuses on important features of the signals, but it also shows that the ensemble does sometimes, assign velocities to bubbles that don't look like they should be assigned a velocity. Based on that and the high accuracy on the labeled data, the assumption is made that it also has high accuracy on the unlabeled data, but if this assumption is true, remains unknown. To resolve this, it could be tried to look at other validation techniques, like camera measurements. This could be used to determine bubble sizes at low flow rates and therefore validate the currently unlabeled data.

The results show that bubbles with low velocity, often get overestimated and that bubbles with low velocity often get underestimated. This is probably because of the lower frequency of extreme bubble velocities in the training data. Data flattening was tried, but that showed unexpected results. Even though the extreme velocities are only a small percentage of our data, when they get predicted wrong, this can still give a wrong idea of the bubble size distribution in the bubble column. This could be resolved by only adding data flattening to the training data of models that do not struggle with overfitting. Since data flattening simply duplicated and added some noise to the signals of the undervalued regions, instead of generating data, it could be tried to select it from other runs, resulting in more complete data and more variation of data in these regions.

subsection 3.4 shows that the ensemble's accuracy does not depend on the flow rate. It also shows that the improvement increases when the flow rate increases. This is mainly because the detection rate of the original method decreases fast when the flow rate in creases, while that of the ensemble does not decrease as fast. A reason for this could be that when the flow rate increases, the flow regime in the bubble column becomes more chaotic. This could make bubble signals measured by the fibre probe more noisy, because less bubbles approach the probe vertically for example, therefore making the original method struggle more with predicting. The ensemble was trained to be able to assign velocities to more noisy data, and therefore it could be that it is influenced less by this extra noise, decreasing the detection rate less. It could also be tried to execute runs with a higher acquisition frequency to see if the detection rate improves for the original method.

References

[1] Anthony Lefebvrea et al. *A new, optimized Doppler optical probe for phase detection, bubble velocity and size measurements: investigation of a bubble column operated in the heterogeneous regime*. Sept. 2021. DOI: 10.48550/arXiv.2109.08100.

[2] Nigar Kantarci, Fahir Borak, and Kutlu O. Ulgen. "Bubble column reactors". In: *Process Biochemistry* 40.7 (June 2005), pp. 2263–2283. ISSN: 1359-5113. DOI: 10.1016/J.PROCBIO.2004.10.004.

[3] Lars Puiman et al. "Gas mass transfer in syngas fermentation broths is enhanced by ethanol". In: *Biochemical Engineering Journal* 185 (July 2022). ISSN: 1873295X. DOI: 10.1016/J.BEJ.2022.108505.

[4] Wei Bai. "Experimental and numerical investigation of bubble column reactors". PhD thesis. Eindhoven: Technische Universiteit Eindhoven, 2010. DOI: 10.6100/IR693280.

[5] Simone Pagliara et al. "Intrusive effects of dual-tip conductivity probes on bubble measurements in a wide velocity range". In: *International Journal of Multiphase Flow* 170 (Jan. 2024), p. 104660. ISSN: 0301-9322. DOI: 10.1016/J.IJMULTIPHASEFLOW.2023.104660.

[6] Yuki Mizushima, Akihiro Sakamoto, and Takayuki Saito. "Measurement technique of bubble velocity and diameter in a bubble column via single-tip optical-fiber probing with judgment of the pierced position and angle". In: *Chemical Engineering Science* 100 (Aug. 2013), pp. 98–104. ISSN: 0009-2509. DOI: 10.1016/J.CES.2013.01.046.

[7] Volger R et al. "Use of ensemble modelling for bubble size measurements with an optical fibre probe". In: *AI4b.io symposium*. Delft, The Netherlands: Delft University of Technology, 2025.

[8] H. Chaumat et al. "Mass transfer in bubble column for industrial conditions—effects of organic medium, gas and liquid flow rates and column design". In: *Chemical Engineering Science* 60.22 (Nov. 2005), pp. 5930–5936. ISSN: 0009-2509. DOI: 10.1016/J.CES.2005.04.026.

[9] *M2 Bubbly Flow Analyzer*. URL: https://a2photonicsensors.com/m2-bubbly-flow-analyzer/.

[10] Youru Li et al. "EA-LSTM: Evolutionary attention-based LSTM for time series prediction". In: *Knowledge-Based Systems* 181 (Oct. 2019), p. 104785. ISSN: 0950-7051. DOI: 10.1016/J.KNOSYS.2019.05.028.

[11] Peter T Yamak, Li Yujian, and Pius K Gadosey. "A Comparison between ARIMA, LSTM, and GRU for Time Series Forecasting". In: *2019 2nd International Conference on Algorithms, Computing and Artificial Intelligence*. Vol. 7. Sanya, China: ACM, 2019. ISBN: 9781450372619. DOI: 10.1145/3377713.3377722.

[12] Antonio A. Abello, Roberto Hirata, and Zhangyang Wang. "Dissecting the high-frequency bias in convolutional neural networks". In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. Nashville, TN, USA: IEEE, June 2021. ISBN: 9781665448994. DOI: 10.1109/CVPRW53098.2021.00096.

[13] M. A. Ganaie et al. "Ensemble deep learning: A review". In: *Engineering Applications of Artificial Intelligence* 115 (Oct. 2022), p. 105151. ISSN: 0952-1976. DOI: 10.1016/J.ENGAPPAI.2022.105151.

[14] Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. “An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling”. In: (Apr. 2018). DOI: 10.48550/arXiv.1803.01271.