CHEMICAL ENGINEERING TRANSACTIONS

VOL. 115, 2025

A publication of AIDIC
The Italian Association of Chemical Engineering
Online at www.cetjournal.it

Guest Editors: Carlo Pirola, Antonio Espuña, Sabrina Copelli

ISBN 979-12-81206-16-8; **ISSN** 2283-9216

DOI: 10.3303/CET25115013

Effectiveness of Generative Adversarial Networks in Low Data Availability Environments for the Detection of Chemical Foam

Toon Stuyck[a,b*], Eric Demeester[b]

[a]BASF Antwerpen, BASF Group, Scheldelaan 600, 2040 Antwerpen, Belgium
[b]Department of Mechanical Engineering, KU Leuven, Wetenschapspark 27, 3590 Diepenbeek, Belgium
toon.stuyck@basf.com

The chemical sector faces significant challenges in applying deep learning techniques due to the scarcity of extensive and well-labeled datasets. This research investigates the potential of Generative Adversarial Networks (GANs) to address these limitations, exploring scenarios where limited or unlabeled data is available. One approach leverages GAN-generated synthetic data to enhance classification accuracy when dealing with small weakly labeled datasets. This technique has the potential to improve model performance without the extensive data collection and labeling traditionally required. Additionally, a second approach leverages the application of GAN-based anomaly detection algorithms, which offer the ability to identify anomalies in data without the need for manual labeling. This study demonstrates the effectiveness of augmenting small, weakly labeled datasets with synthetic data generated by GANs to significantly improve the classification of chemical foam. It is shown that trustworthy models can be developed that are able to increase the accuracy from 60% up to 91%. Conversely, the research also finds that GAN-based anomaly detection shows less impact in the context of chemical foam detection and segmentation. The method can identify anomalies but fails to completely segment them. These findings offer valuable insights into the potential and limitations of applying GANs for specific tasks within the chemical sector.
Keywords: Chemical Engineering; Machine learning; Generative Adversarial Networks

1. Introduction

Chemical industry camera systems are often monitored manually because possible AI models, while potentially very impactful, still lack sufficient trust due to limited training or validation data. For the detection of certain objects, augmentation of the data to increase the amount of available data and increase classification results, is already widely used for many deep learning techniques. The reason augmentation techniques are used is because the accuracy of deep learning classification techniques such as convolutional neural networks (CNNs) tend to dramatically drop when confronted with limited data (Stuyck and Demeester, 2024). However, it was shown by (Maharana et al., 2022) that these augmentation methods could lead to overfitting and/or loss of information. For this reason it might be more interesting to look at more sophisticated techniques such as generative adversarial networks (GANs) as introduced by (Goodfellow et al., 2014). The use of this technique allows many applications such as, e.g.: Anomaly detection or synthetic data generation.

Anomaly detection is the ability to detect anomalies, which are situations that deviate from normal. Often this is done via either the usage of variational autoencoders (VAEs) as shown by (Kiran et al., 2018) or GANs. (Akçay et al., 2019) introduced a GAN-architecture for anomaly detection they called: GANomaly. The downside of these techniques is that they are only able to label or score images as anomaly or not. (Stuyck et al., 2022) used this architecture and adapted it, so that can detect anomalies as well as locate and segment them.

Synthetic data is data that is created artificially to train deep learning models. Augmentation techniques, even though with the shortcomings described by (Maharana et al., 2022), are still frequently used due to their cost-effectiveness to increase the amount of available data with factors of thousands. In recent years, augmentation techniques were used to achieve better fire detection models (Kang et al., 2018) or for the detection of tomato

Paper Received: 29 September 2024; Revised: 09 December 2024; Accepted: 05 March 2025
Please cite this article as: Stuyck T., Demeester E., 2025, Effectiveness of Generative Adversarial Networks in Low Data Availability Environments for the Detection of Chemical Foam, Chemical Engineering Transactions, 115, 73-78 DOI:10.3303/CET25115013

leaf diseases (Agarwal et al., 2020). However, if the augmentation transformations are applied too aggressively or inappropriately, it could be that detection algorithms, instead of learning features of the original dataset, starts learning artifacts introduced due to the augmentations, leading to overfitting or a loss of information (Shorten and Khoshgoftaar., 2019). For this reason, recently more and more research has gone into the field of artificially generated synthetic data. Often GANs are used for this generation, however it is also possible to use VAEs or transformers. A large risk of using artificially generated data when dealing with low data availability is that not all features will be incorporated in the generated data, and trained models might be biased (Karras et al., 2020). One way to identify this possible bias is by introducing explainable AI (XAI). (Hassan et al., 2022) made a comparison of multiple pre-trained deep neural networks for the classification task of prostate cancer and used XAI to gain insights and understanding of the key features that led the algorithm to make a certain decision and classification. Pioneered by (Lundberg and Lee, 2017), SHAP (SHapley Additive exPlanations) has become a widely adopted explainable AI (XAI) method due to its user-friendliness and remarkable explanatory power.

2. Methods

In this paper the usability of different GAN applications for the detection of chemical foam will be validated. This will be done by comparing the classification accuracy of (i) a convolutional neural network trained on a limited dataset, with (ii) the accuracy of a classification by a GAN based anomaly detection method on the same dataset and finally with (iii) the accuracy of a convolutional neural network trained on the original dataset enlarged with GAN based synthetic data. The aim of this research is to contribute by providing a methodology when dealing with detection questions on very limited static datasets.

The dataset used in this work contains images of an active chemical production environment with (anomaly) and without (normal case) the presence of chemical foam. For the normal images only limited variation is possible due to the static nature of the setup and camera. The only possible variation in the normal case is due to residual foam that sticks on the buffer tank, or due to changes in weather phenomena. In the case of the presence of an anomaly, or foam, more variation is possible. The importance of fast detection is extremely high for this use case since the foam can start forming at any moment and is very unpredictable. If the foam starts rising and overflows from the buffer tank, the risk exists of a potential environmental impact as well as the risk of early corrosion of nearby installations, on top, a dangerous environment for operators is created.

A dataset with up to 326 images with foam and up to 424 images without foam of the same active chemical production environment have been collected over the span of multiple weeks, allowing for many variations in background, such as: day/night cycle, rain, snow, sunny, etc. All collected images have been weakly labelled, meaning that all images have been reviewed and annotated as either containing foam of no foam.

For the generation of synthetic data many different approaches exist, however, since the datasets used in this work are limited in the amount of data, the light-weight GAN structure proposed by (Liu et al., 2020) is used. The reason this light-weight approach is used, is because these models can converge within hours and more importantly, they have similar performances consistently even when only limited amounts of data are available. In this work, only the impact of GAN-based synthetic generated data on classifier accuracy will be evaluated, since (Stuyck en Demeester, 2024) showed that GAN-based synthetic generated data consistently yield superior results compared to standard augmentation methods for classification. In this work, only the light-weight approach will be used since (Lucic et al., 2018) suggest that on average, different GAN architectures have limited impact on the final generated result. The generated data in combination with the original dataset will be used to train a CNN for the classification of chemical foam. These classification results will be subjected to the SHAP explanation model to get insights in the performance and decision-making capabilities of the final proposed classification models, as well as map possible biases of the newly generated data that might negatively impact the classification algorithm results.

For the anomaly detection approach, an adaptation of the Skip-GANomaly (Akçay et al., 2019) is used. This method is trained only on images without chemical foam and tries to learn the features so that it can recreate foamless images. It does this by using a typical GAN based architecture consisting out of a generator network (G) and a discriminator network (D). The goal of the generator network is to capture the distribution of the input images and try to generate new realistic images based on the features taken from this distribution. The discriminator network has the goal to classify the original images (x) and the generated images $(\hat{x})$ to the correct class. These networks are adversarially trained where the goal of the generator network is to fool the discriminator, and the discriminator tries to continuously improve its classification accuracy. The benefit of utilizing the Skip-GANomaly proposed architecture is that it uses a combination of three different loss functions: The adversarial loss, the contextual loss and the latent loss which opens the possibility for an adaptation that allows the introduction of structural similarity for the automated segmentation of anomalies. The adversarial loss has the objective to maximize the reconstruction capability of the generator of foamless images based on a given input image to recreate realistic looking results. The discriminator has the objective to correctly classify

the original and generated images. The adversarial loss needs to be maximized for the discriminator network D, and minimized for the generator network G.
The contextual loss has the goal to ensure that not only realistic images are generated but more importantly the reconstructed images are also contextual similar to the original input image. This is done by the introduction of the L1-norm, which ensures contextual similarity between the original image (x) and the generated image ($\hat{x}$).
Finally, the latent loss makes sure that the latent representation of both the reconstructed and original image are similar. Based on these loss functions, it can be assumed that reconstructed images will be both realistic and contextually similar to an input image, expect for the part where the anomaly is located, there the reconstruction will locally fail. Due to this, anomalies can automatically be segmented using structural similarity. The effectiveness of this approach has already been confirmed in previous work (Stuyck et al., 2022).

3. Results and discussion

All the experiments in the section have been performed on a PC with an Intel I7-1080h at 2.7 GhZ and an NVIDIA Quadro RTX 4000 GPU. For the synthetic generated data approach, the training data has been collected continuously over multiple weeks. Seeing that the data collection was an ongoing process, multiple smaller subdatasets have been collected. These small datasets have also been used to simulate environments for which it is not possible to collect additional datasets. In the end, five different datasets have been collected, where the dataset size respectively represents 10%, 25%, 50%, 75% and 100% of the final complete dataset. Using these different sized datasets and the light-weight GAN architecture, different models can be trained and for each, an unlimited set of new synthetic images can be generated. Figure 1 shows some example images of the generated results. From these images it can be seen that even for small datasets, it is possible to create realistic looking new images with a small simulation-to-realism gap. However, when newly generated images for each of the models are lined up and compared, it can be noticed that some visual differences exist between them. If the amount of available data decreases, the amount of variation in the generated images decreases. This is to be expected, since the amount of variation in the training data also decreases. On top of that, it can also be observed that if the amount of training data decreases, the quality of the generated images decreases slightly. It can be observed that for the images with less training data, more noise is visible on the generated images. Based on the created datasets, two different sets of CNNs can be trained. The first set of CNNs are trained for each of the different sized dataset that only contain original data. The second set of CNNs are trained for each of the different sized datasets combined with the newly generated synthetic data, that was generated based on each of the reduced datasets. The accuracy results are shown in Figure 2. When only 10% to 50% of the total dataset was gathered, the classifier trained on the limited dataset is unusable due to very poor performance. Once more data became available, it can be observed that the accuracy quickly rises and a maximum accuracy of 87% can be achieved on the complete dataset. However if synthetic data is generated based on each of these smaller datasets, the accuracy can be increased to achieve much better performing models. When 50% of the original dataset is available, and is used for the generation of additional data, the accuracy is pushed to 91%. The highest accuracy achieved by extending the dataset is 94%. This approach shows that by introducing GAN based synthetic data, the accuracy of classifier models can be increased dramatically even when dealing with substantially less data. It can be observed that for the 75% dataset, the accuracy drops minimally. This might be explained by the random sampling of the datasets, that for the 75% dataset less representative samples might have been selected. For the total dataset, the accuracy increases again. However, to further increase trust in the developed models, and especially those with low amounts of input data, SHAP values can be calculated for all the different models where synthetic data was used to get additional insights and understanding in the respective model performance. Some example outputs are shown in Figure 3 for the model where the complete dataset was available and additional synthetic data was generated and for the model where 25% of the dataset was available and additional synthetic data was generated. Pixels indicated in red show pixels that contribute to the decision to label this image as containing foam. Blue pixels indicate pixels that contribute negatively to the decision to label this image as containing foam. It can be seen, that in the case the total dataset is available and is extended with synthetically generated data, the distribution of SHAP values is as a human observer would expect. However, when the amount of available data decreases, the distribution of the SHAP values is no longer what a human observer would expect. This is shown in Figure 3 (c and d). These figures show that the distribution seems to be completely random. This visually indicates that when the available amount of training and validation data decreases, the model might be overfitting, even though the accuracy of the model stays high. The effectiveness of the models using the SHAP values has also been validated using a quantitative approach. To do this, a performance indicator is defined using a region of interest (ROI). The ROI is defined as the area inside the chemical installation. The performance indicator is defined as follows, for each calculated SHAP value, it can be checked if it is located within the defined ROI or not, and to which category it contributed. These results, given in Table 1 quantitatively confirm previous visual findings. When little data is available and the simulation-to-realism gap is visible, the ratio of SHAP values located inside the ROI ranges between 40% and 50%, indicating that these models might not be reliable enough to use. If the amount of available data increases and the simulation-to-realism gap becomes smaller, it can be observed that the ratio increases. When using the complete dataset, a maximum ratio of 86% can be obtained. This indicates that these models are much better aligned with how a subject matter expert would make their evaluation and are thus more reasonable to trust. The benefit of using the proposed performance indicator is that it can help automate the procedure of determining how much data is needed to train useful and reliable models. By using this performance indicator it is possible to quantify the trustworthiness and no longer make it based on subjective visual observations making it much easier to interpret for end-users. The downside of this approach is that currently it is only applicable on static images with a clear predefined ROI. These results indicate that using GAN generated synthetic data to extend existing dataset is only useful when the simulation-to-realism gap can be made sufficiently small. Integrating SHAP values and the performance indicator into the workflow empowers users with additional information and insights into the real-world performance of the models. This is achieved by comparing the location of features used by the model and those a human observer would use.

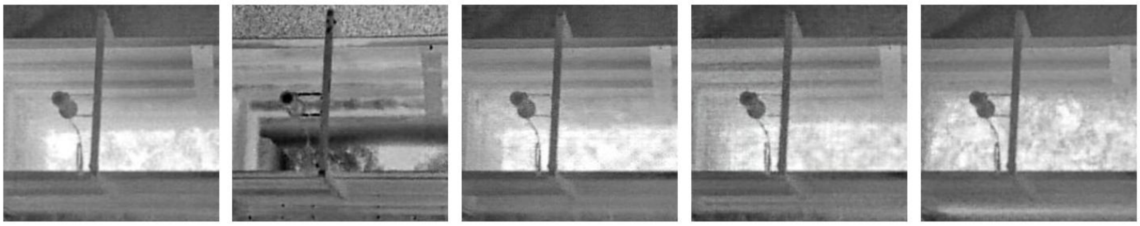

(a) 100% (b) 75% (c) 50% (d) 25% (e) 10%

Figure 1: GAN generated synthetic data for the five different datasets with decreasing amounts of training data.

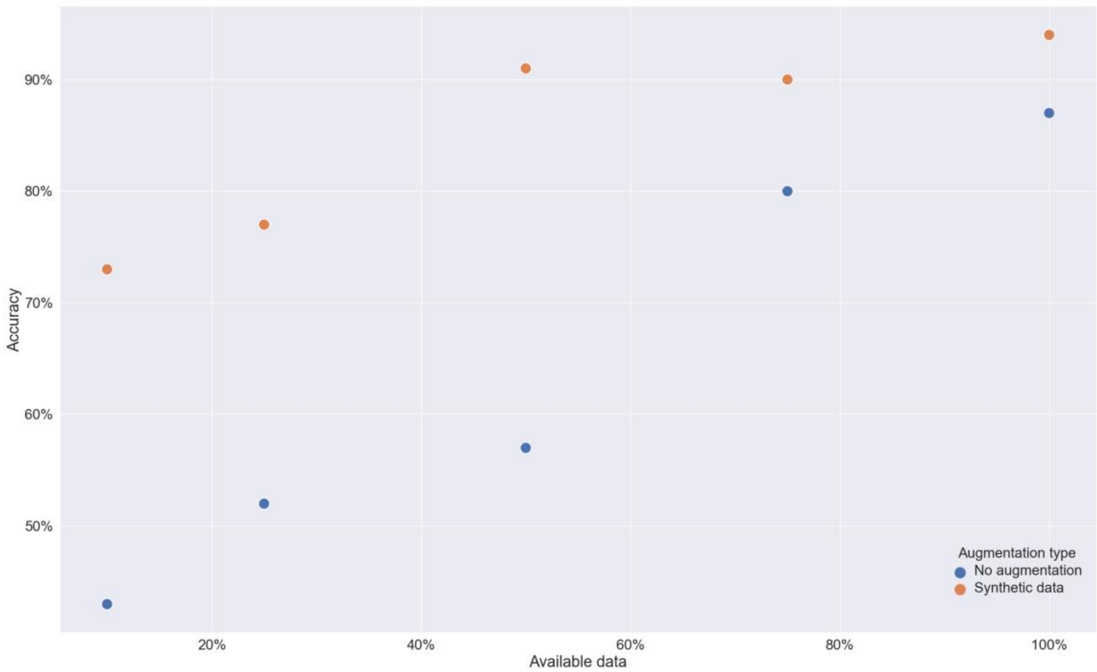

Figure 2: Overview of the accuracy of the trained convolutional neural network for the different variations in available data. The blue points indicate the accuracy results when only the original dataset is used. The red points indicate the accuracy results when the original dataset was combined with synthetic generated data.

For the anomaly detection approach the SKIP-GANomaly architecture is used. Training the model on only anomaly free images allows the model to learn the normal representation of the dataset and makes it possible to detect deviations from anomaly free samples. The loss functions that have been used, ensure that reconstructed images are visually similar to the original input image. Only when an input image contains an anomaly the expectation is that the reconstruction will locally fail where the anomaly is located. Combining this knowledge with the SSIM approach allows for automated segmentation of anomalies.

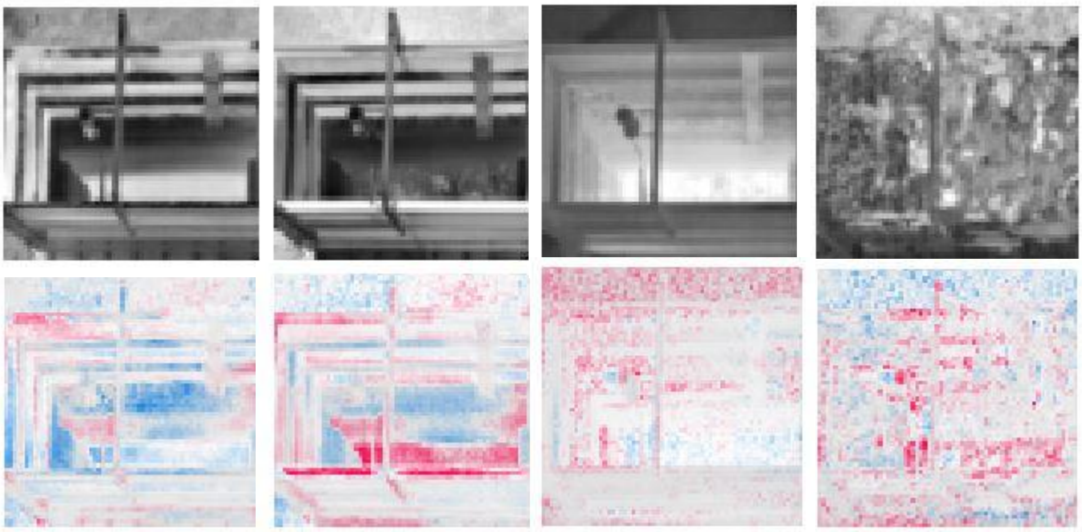

(a) 100% data, No foam (b) 100% data, Foam (c) 25% data, No foam (d) 25% data, Foam

Figure 3: Example images of SHAP values for models with high and low data availability and where the cases foam and no foam are present. Red pixels indicate a contribution to the decision to label this image as containing foam. Blue pixels indicate pixels that contribute negatively to the decision to label this image as containing foam.

Figure 4 shows some results where the reconstruction of the images with foam fails partially. The failed reconstruction is not consistent due to parts of the foam being reconstructed correctly. This has an impact on the approach, since the algorithm will not be able to detect differences between the correctly reconstructed foam, and the original foam on the image. This will lead to underestimates of the amount of present foam.

Table 1: Table showing an overview of the ratio of SHAP values that explain classification results that are also located inside the defined ROI versus outside the ROI, for the different amount trained models and different amounts of available and synthetic data.

Amount of available data from complete dataset									
10%		25%		50%		75%		100%	
Amount of synthetically extended data									
50%	100%	50%	100%	50%	100%	50%	100%	50%	100%
Ratio of SHAP values that explain classification results that are located inside the ROI									
43%	40%	48%	53%	75%	78%	81%	79%	84%	86%

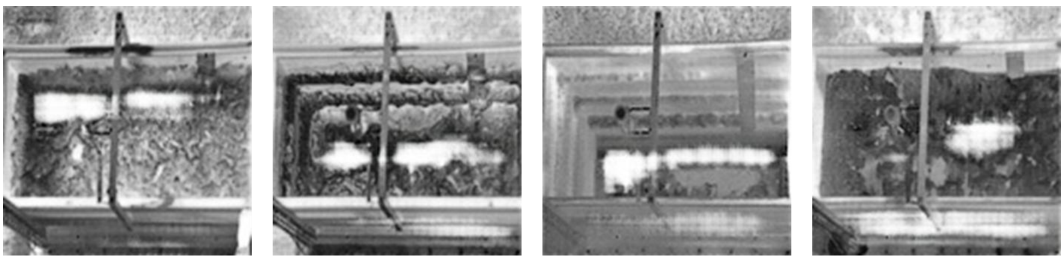

Figure 4: Example images of reconstructed images where the reconstruction of the foam fails partially. However, the reconstruction should fail for all the foam visible and not only a part.

Additional issues seem to surface when looking at the results of anomaly free images. The small dataset size seems to limit the model's ability to learn the "normal" image representation. This is more impactful than the issues that arise with only limited failed reconstruction of actual anomalies, since failure to use structural similarity only limits in the ability to apply segmentation, but correctly identifying whether foam is present in the image should not be hampered. However, since the algorithm also fails to consistently classify images correctly, the approach cannot effectively be used for classification of anomalies. Overall this approach was able to achieve an accuracy of 73% when validated on the validation dataset. In comparison, the CNN trained on the complete dataset both using labeled images with and without anomalies, was able to achieve an accuracy of 87% on the classification task.

4. Conclusions

The synthetic data generation approach successfully enhanced classification accuracy for the chemical foam dataset. Utilizing GAN-generated synthetic data, it was possible to significantly improve performance compared to solely using a small, weakly labelled dataset. Notably, even with only half of the original data, accuracy significantly rose from 57% to 91%, highlighting the potential to drastically reduce the requirement for expensive and time-consuming data collection. By introducing explainable AI methods, to further validate and explain the developed models, it was shown that when the simulation-to-realism gap is too large, the decision mainly was based on features a human observer would not use, indicating that these models are overfitting. The GAN-based anomaly detection approach proved less impactful for chemical foam detection and segmentation in low data scenarios. The results revealed the method's inability to effectively learn the complete normal data distribution, leading to inaccurate reconstructions in both normal and anomaly cases. This suggests that GAN-based anomaly detection might not be suitable for this specific application when having to deal with limited data. It was even shown that for this specific dataset, a simple CNN achieved better results. Overall, this research demonstrates the power of GANs in overcoming data scarcity within the chemical sector. By utilizing synthetic data generation, the study showcases the potential to achieve high classification accuracy and trustworthiness even with significantly reduced data collection requirements, offering a more cost-effective and time-efficient solution for deep learning applications in this domain.

Acknowledgments

We would like to thank VLAIO and BASF Antwerpen for funding the project (HBC.2020.2876).

References

Agarwal, M. Singh, A., and Gupta, S. (2020). ToLeD: Tomato leaf disease detection using convolutional neural network. Procedia Computer Science, 2020, pages 293-301.

Akçay, S., Atapour, A., and Breckon, T. P. (2019). Skip-GANanomaly: skip connected and adversarially trained encoder-decoder anomaly detection. In International Joint Conference on Neural Networks, 1-8.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. In Advances in Neural Information Processing Systems, volume 27.

Hassan, M. R., Islam, M. F., Uddin, M. Z., Ghoshal, G., Hassan, M. M., Huda, S., and Fortino, G. (2022). Prostate cancer classification from ultrasound and MRI images using deep learning based explainable artificial intelligence. Future Generation Computer Systems, 127:462-472.

Kiran, B. Ravi, Dilip Mathew Thomas, and Ranjith Parakkal. (2018). An overview of deep learning based methods for unsupervised and semi-supervised anomaly detection in videos, Journal of Imaging, 4(2):36.

Karras, T., Aittala, M., Hellsten, J., Laine, S.and Aila, T. (2020). Training generative adversarial networks with limited data. In advances in Neural Information Processing Systems, volume 33, 12104-12114.

Liu, B., Zhu, Y., Song, K., and Elgammal, A. (2020) Towards faster and stabilized GAN training for high fidelity few-shot image synthesis. In International conference on Learning Representations.

Lucic, M., Kurach, K., Michalski, M., Gelly, S., and Bousquet, O. (2018). Are GANs created equal? A large scale study. Advances in neural information processing systems, 30.

Lundberg, S. M. and Lee, S.-I. (2017). A unified approach to interpreting model predictions, Advances in neural information processing systems, 30.

Maharana, K., Mondal, S., and Nemade, B., (2022), A review: Data pre-processing and data augmentation techniques, Global Transitions Proceedings, 3(1), 91-99.

Shorten, C., and Khoshgoftaar, T. M. (2019). A survey on image data augmentation for deep learning. Journal of big data, 6(1), 1-48.

Stuyck, T. and Demeester, E., (2024), Impact of using GAN generated synthetic data for the classification of chemical foam in low data availability environments, in Proceedings of the 13th international conference on pattern recognition applications and methods, 620-627.

Stuyck, T, Rousseau, A.-J., Vallerio, M., and Demeester, E. (2022). Semi-supervised cloud detection with weakly labeled RGB aerial images using generative adversarial networks. In Proceedings of the 11th international conference on pattern recognition applications and methods, 630-635.