# Adaptive Contrastive Learning Framework for Real-World Applications with Class Imbalance

Hyeonseok Seo*, Junghyun Lee*, Jun Kyun Choi*, Chanhyung Lee*, Hyunseo Park†, and Ismail Uzun‡
*School of Electrical Engineering, Korea Advanced Institute of Science and Technology, Daejeon, Rep. of Korea
†Information & Electronic Research Institute, Korea Advanced Institute of Science and Technology, Daejeon, Rep. of Korea
‡inosens, Gebze, Kocaeli, Turkiye
Email: gustjr0413@kaist.ac.kr, junghyun960@kaist.ac.kr, jkchoi59@kaist.edu,
leech0529@kaist.ac.kr, tkf92001@kaist.ac.kr, ismail.uzun@inosens.com.tr.

***Abstract*—Class imbalance is a critical challenge in supervised learning, where models often become biased toward majority classes and fail to effectively learn from minority samples. This issue is particularly pronounced in real-world domains such as industrial monitoring, where abnormal or fault-related events are rare by nature. In this paper, we propose an adaptive contrastive learning framework that addresses class imbalance by dynamically combining classification loss and supervised contrastive loss, based on the degree of imbalance in the dataset. To validate the effectiveness of our method, we conduct experiments using public anomaly detection datasets and open-source classification models. The results demonstrate that the proposed framework improves minority class prediction while maintaining overall accuracy.**

***Index Terms*—Class imbalance, Contrastive learning, Deep learning.**

## I. Introduction

Supervised learning models have achieved remarkable success in various classification tasks across domains such as computer vision, natural language processing, and medical diagnosis. Their widespread success relies heavily on the availability of large-scale labeled datasets that are well balanced across classes. However, this assumption does not hold in many real-world applications, where data often exhibit substantial class imbalance. In such scenarios, models tend to overfit to the dominant class distribution, leading to poor performance on underrepresented classes [1], [2].

Class imbalance is particularly problematic in safety-critical domains, including industrial monitoring, autonomous systems, healthcare diagnostics, and disaster response [3]–[5]. In these settings, most of the collected data correspond to routine or normal events, while abnormal, anomalous, or hazardous events occur infrequently. For instance, in aerial surveillance for emergency response, images depicting disasters such as fires or floods are much rarer than those depicting normal scenes. Despite their low frequency, these minority class instances are often of highest practical importance, and failure to detect them accurately may result in significant safety, operational, or financial consequences [6].

This research was financially supported by the Ministry of Trade, Industry and Energy(MOTIE) and Korea Institute for Advancement of Technology(KIAT) through the International Cooperative R&D program. (Project No. P0027980)

Moreover, class imbalance frequently results from inherent characteristics of real-world data collection processes. Data acquisition in critical domains is often costly, time-consuming, or even dangerous, leading to limited samples for rare or hazardous events. This further exacerbates the challenge by constraining the diversity and quantity of available data for minority classes, ultimately amplifying the biases learned by supervised models. Additionally, minority class instances may exhibit greater variability or complexity, requiring models to learn discriminative representations from a comparatively small set of challenging examples. These conditions highlight the necessity of developing specialized training strategies to ensure balanced performance across all classes.

This imbalance not only affects the model's ability to generalize but also introduces bias during training [7]–[9]. Conventional classification losses such as cross-entropy treat each instance equally, thereby favoring classes with more data. As a result, minority class samples contribute less to the loss and are under-emphasized during optimization. This leads to skewed decision boundaries and performance disparities across classes.

To address this challenge, we propose a method to mitigate performance bias under class-imbalanced conditions by integrating an adaptive contrastive learning framework into existing classification model architectures, without modifying the dataset or altering the core model architecture. proposed framework dynamically combines a class-weighted cross-entropy loss and a supervised contrastive loss, where the influence of each loss component is adjusted based on the degree of class imbalance present in the data. Furthermore, we introduce a temperature-scaling mechanism in the contrastive loss to adaptively control the sharpness of similarity discrimination in proportion to imbalance severity.

We validate proposed approach through experiments on public datasets for anomaly detection using open-source classification model. By constructing synthetic scenarios with varying class imbalance ratios, we demonstrate that our method improves model robustness and prediction performance for minority classes, while maintaining high overall accuracy. The results highlight the framework's ability to address imbalance effectively and its potential for integration into real-world safety-critical applications.

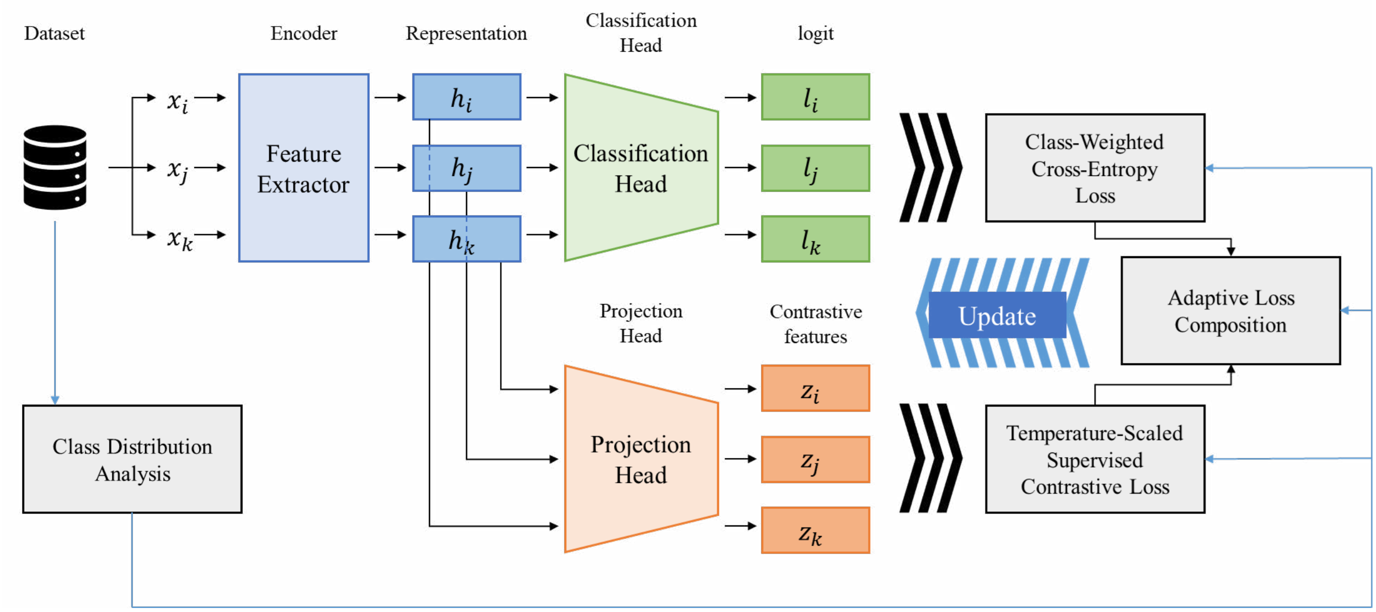


Fig. 1. Overview of the proposed adaptive contrastive learning framework. The architecture includes a general classification structure consisting of a feature extractor and a classification head, and integrates a contrastive learning structure consisting of an encoder and a prediction head. To mitigate performance bias caused by class imbalance, the framework computes a class-weighted cross-entropy loss and a temperature-scaled supervised contrastive loss based on the class distribution and the degree of imbalance, and dynamically combines them to train the model.

## II. Methodology

This study designs an adaptive contrastive learning framework aiming to enhance classification robustness under class-imbalanced conditions by jointly utilizing classification and representation learning objectives. The overall architecture of the proposed framework is illustrated in Fig. 1.

The upper part of proposed framework follows a standard supervised learning pipeline, where a labeled dataset is used to train a classification model composed of a feature extractor and a classification head. Outputs of feature extractor are passed through the original classification head to produce class logits, which are used to calculate a classification loss.

In the lower part of proposed framework, the feature extractor is reused as an encoder to generate high-level representations from the input data. These representations are then fed into the projection head that maps the encoder outputs into a lower-dimensional embedding space. This projection is used to compute a supervised contrastive loss that encourages the model to bring representations of the same class closer together while pushing those of different classes apart.

The final training objective combines both the classification and contrastive losses, which is weighted adaptively based on the degree of class imbalance. The proposed framework consists of three key components to address imbalance: a class-weighted cross-entropy loss, a temperature-scaled supervised contrastive loss, and an adaptive loss composition strategy.

### A. Class-Weighted Cross-Entropy Loss

Conventional classification models are typically trained using the cross-entropy loss, which aims to maximize the likelihood of the correct class label given the model's output. Formally, the standard cross-entropy loss for a sample $(x_i, y_i)$ is defined as:

$$\mathcal{L}_{\text{CE}} = -\frac{1}{N}\sum_{i=1}^{N}\log p_{y_i}, \quad \text{where } p_{y_i} = \frac{\exp(l_{i,y_i})}{\sum_{j=1}^{C}\exp(l_{i,j})} \tag{1}$$

Here, $l_{i,j}$ represents the logit for class $j$, $C$ is the total number of classes, and $N$ is batch size. While cross-entropy loss is effective for balanced datasets, it performs poorly in class-imbalanced settings, as it tends to bias predictions toward majority classes. To address this problem, we introduce a class-weighted cross-entropy loss $\mathcal{L}_{\text{WCE}}$ that assigns greater importance to minority classes.

$$\mathcal{L}_{\text{WCE}} = -\sum_{i=1}^{N} w_{y_i}\log p_{y_i} \tag{2}$$

where $w_{y_i}$ is the weight corresponding to the class $y_i$, defined as $w_{y_i} = \frac{1}{f_{y_i}}$, with $f_{y_i}$ representing the relative frequency of class $y_i$ in the dataset. This class-weighted cross-entropy loss can mitigate the bias toward majority class predictions.

### B. Temperature-Scaled Supervised Contrastive Loss

Supervised contrastive learning (SupCon) addresses the limitation of CE by considering relationships between samples within the same class [10]. Instead of treating each sample independently, it encourages the model to learn embeddings where samples from the same class are closer together in the feature space. The supervised contrastive loss is defined as:

$$\mathcal{L}_{\text{SupCon}} = \sum_{i=1}^{N} \frac{1}{|P(i)|} \sum_{p \in P(i)} -\log \frac{\exp(\tilde{z}_i \cdot \tilde{z}_p / \tau)}{\sum_{a=1}^{N} \mathbf{1}_{[a \neq i]} \exp(\tilde{z}_i \cdot \tilde{z}_a / \tau)} \tag{3}$$

where $\tilde{z}_i$ and $\tilde{z}_p$ are normalized contrastive features, $P(i)$ denotes the set of positive samples (same class) for sample $i$, and $\tau$ is the temperature hyperparameter. Although SupCon improves representation learning by considering inter-sample similarities, it does not explicitly address class imbalance or dynamically adapt to different data distributions. In particular, using a fixed temperature $\tau$ can limit its flexibility across datasets with varying imbalance levels. To address this, we introduce temperature scaling that adapts the temperature dynamically based on the degree of class imbalance.

Specifically, we define the imbalance level $\sigma$ as the standard deviation of the class distribution:

$$\sigma = \text{StdDev}\left(\{p_1, p_2, \ldots, p_C\}\right), \tag{4}$$

where $p_c$ denotes the relative frequency of class $c$, and $C$ is the total number of classes in the dataset.

We then compute the adaptive temperature $\tau$ as:

$$\tau = \text{clip}\left(\tau_0 \cdot (1 - \sigma),\ \tau_{min},\ \tau_{max}\right), \tag{5}$$

where $\tau_0$ is a base temperature value, and $\text{clip}\,(x, a, b)$ denotes element-wise clipping to the range $[a, b]$. The bounding values $\tau_{min}$ and $\tau_{max}$ are used to ensure numerical stability.

This adaptive scaling ensures that the contrastive objective becomes sharper when the class distribution is highly imbalanced (i.e., large $\sigma$) and more relaxed when the distribution is relatively balanced. It allows the model to more effectively differentiate between classes under varying data conditions.

### C. Adaptive Loss Composition

To integrate both the class-weighted classification loss and the temperature-scaled supervised contrastive loss, we need to formulate a unified training objective. In order to flexibly control the influence of each component, we introduce an adaptive weighting mechanism based on the degree of class imbalance.

Let $\mathcal{L}_{\text{WCE}}$ denote the class-weighted cross-entropy loss (as defined in Section II-A), and $\mathcal{L}_{\text{SupCon}}$ denote the supervised contrastive loss with temperature scaling (as described in Section II-B). We define the total loss as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{WCE}} + \lambda \cdot \mathcal{L}_{\text{SupCon}}, \tag{6}$$

where $\lambda$ is a contrastive loss weight that is dynamically scaled based on the imbalance level $\sigma$ introduced in the previous section.

The contrastive loss weight is then defined as:

$$\lambda = \text{clip}\left(\lambda_0 \cdot (1 + \sigma),\ \lambda_{min},\ \lambda_{max}\right), \tag{7}$$

where $\lambda_0$ is a base scaling factor. This formulation increases the relative importance of the contrastive loss in highly imbalanced settings, helping the model to more effectively learn from underrepresented classes. When the class distribution is more balanced, the influence of contrastive learning is reduced accordingly, allowing the classification loss to dominate. This adaptive composition strategy enables the model to respond to varying degrees of imbalance.

## III. Experimental Results

In this section, we discuss the experimental results for various class imbalance scenarios using public datasets and open source classification models to evaluate the proposed framework.

### A. Experimental Setup

We conduct experiments on the AIDER dataset (Aerial Image Database for Emergency Response applications) [11], a publicly available dataset designed for scene classification in disaster response scenarios. The dataset contains aerial imagery labeled with five semantic categories: *collapsed buildings*, *fire*, *flooded areas*, *traffic accidents*, and *normal scenes*. These categories represent a mix of emergency and non-emergency conditions typically encountered in real-world aerial monitoring.

To simulate varying degrees of class imbalance, we construct five synthetic training scenarios by sampling different proportions of each class from the original dataset. In each scenario, the training set contains 1,500 images, and the test set contains 100 images, keeping the total sample size consistent across all experiments. The class distributions for each scenario correspond to the five classes listed in the order: *collapsed buildings*, *fire*, *flooded areas*, *traffic accidents*, and *normal*, and are defined as follows:

- **Scenario 1**: [0.05, 0.05, 0.05, 0.05, 0.80]
- **Scenario 2**: [0.05, 0.05, 0.10, 0.10, 0.70]
- **Scenario 3**: [0.05, 0.05, 0.20, 0.20, 0.50]
- **Scenario 4**: [0.10, 0.10, 0.20, 0.20, 0.40]
- **Scenario 5**: [0.20, 0.20, 0.20, 0.20, 0.20]

Scenario 1 represents the most severe imbalance, where the majority class (*normal*) dominates the dataset, while Scenario 5 serves as the fully balanced baseline.

We adopt MobileNetV3-Large as the classification model and encoder in our experiments [12]. MobileNetV3 is a lightweight convolutional neural network architecture developed through neural architecture search (NAS), incorporating squeeze-and-excitation (SE) modules and the hard-swish activation function to achieve a balance between performance and efficiency. The final classification layer is removed, and the output feature vectors are passed to both a classification head and a projection head used for contrastive learning.

For all experiments, we set the lower and upper bounds of the adaptive parameters as follows: $\tau_{min} = 0.01$, $\tau_{max} = 0.1$, $\lambda_{min} = 0.1$, and $\lambda_{max} = 1.0$. These bounds are applied via clipping in the definitions of Eq. (5) and Eq. (7), respectively. This

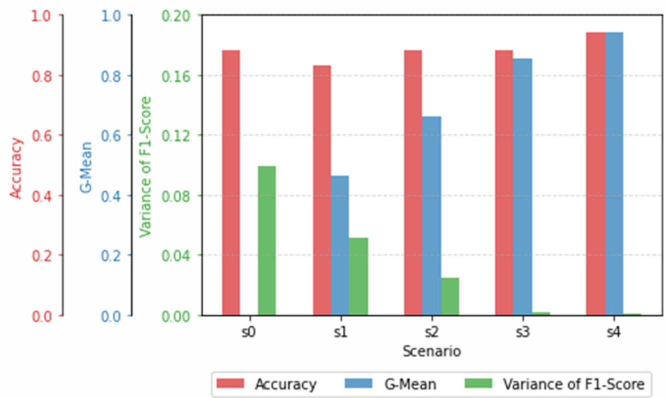


Fig. 2. Performance of the baseline model across different class-imbalance scenarios, evaluated using accuracy, G-Mean, and variance of F1-score.

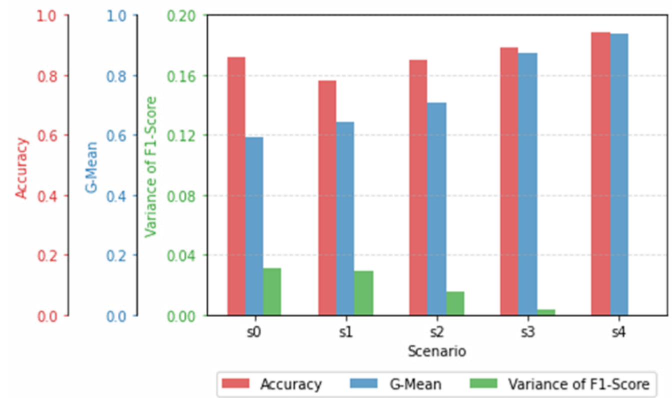


Fig. 4. Performance of the proposed framework across different class-imbalance scenarios, evaluated using various metric.

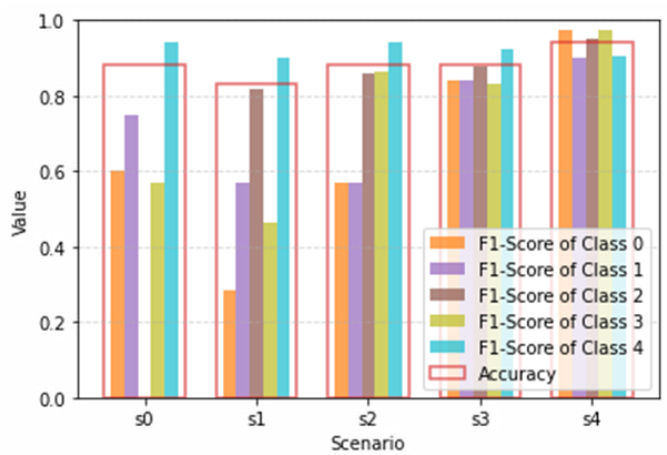


Fig. 3. Accuracy and per-class F1-scores of the baseline model across five class-imbalance scenarios.

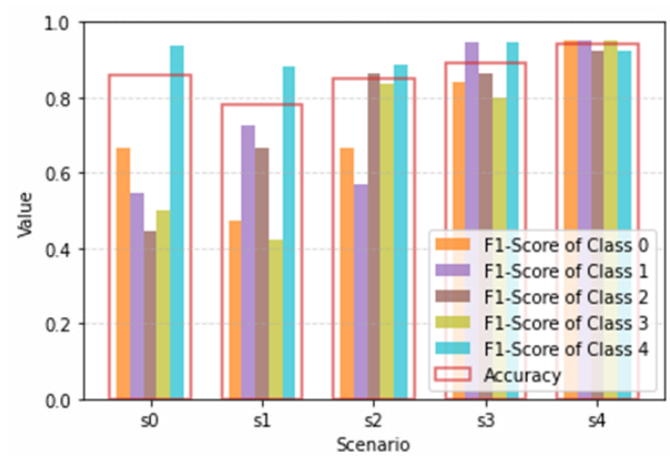


Fig. 5. Accuracy and per-class F1-scores of the proposed framework across five class-imbalance scenarios.

constraint ensures stable training behavior while maintaining sensitivity to the degree of class imbalance.

### *B. Results and Discussions*

To evaluate the effectiveness of the proposed framework, we conduct a comparative analysis between the baseline MobileNetV3 model and our enhanced model integrated with the adaptive contrastive learning framework. The evaluation is performed across all five class-imbalance scenarios using multiple performance metrics, including overall accuracy, geometric mean (G-Mean), F1-scores. These metrics collectively capture not only the general classification performance but also the model's robustness to class imbalance and its ability to perform fairly across all classes. A higher G-Mean and a lower variance of f1-score indicate less variation in detection performance between classes.

We first evaluate the baseline model without applying the proposed framework across all five class-imbalance scenarios. The results in Fig. 2 shows the overall performance of the baseline model across five different class-imbalance scenarios. While the accuracy remains consistently high, this metric alone is insufficient to reflect the model's performance under imbalance. As the class distribution becomes more skewed (e.g., Scenario 1 and 2), the G-Mean drops significantly, indicating that the model fails to maintain balanced recall across all classes. In parallel, the variance of the F1-scores increases, suggesting that the model performs inconsistently across classes. These results reveal that the baseline model is strongly biased toward the majority class, which dominates the accuracy score but leads to poor recognition of minority classes.

A more detailed view is provided in Fig. 3, which shows the per-class F1-scores along with the overall accuracy. In scenarios with high imbalance, the F1-scores of minority classes are significantly lower than that of the majority class (Class 4), which maintains high scores. This confirms that the baseline model largely fails to learn effective decision boundaries for underrepresented classes. Despite this, the overall accuracy appears unaffected due to the model's tendency to favor the majority class. This performance gap highlights the limitations of using conventional training objectives under imbalanced conditions and underscores the need for more adaptive methods.

To address the limitations observed in the baseline model, we apply the proposed adaptive contrastive learning framework to the same set of imbalance scenarios. This experiment aims to evaluate whether the framework can improve the recognition of minority classes and reduce performance disparities across classes, while maintaining overall classification accuracy.

Fig. 4 illustrates the performance of the proposed framework across the five imbalance scenarios. While the accuracy remains consistently high, as observed in the baseline, the G-Mean increases noticeably, especially in scenarios with more severe imbalance. This indicates that the model achieves more balanced recall across all classes. Furthermore, the variance of the F1-scores is significantly reduced compared to the baseline. This reduction suggests that the model performs more consistently across classes and is less biased toward the majority class. These improvements confirm that our framework effectively mitigates the adverse effects of class imbalance and enhances overall robustness.

A more detailed breakdown is provided in Fig. 5, which shows the F1-scores for each individual class under all scenarios. Compared to the baseline, the F1-scores for minority classes, which were previously low, show substantial improvements. This trend is especially clear in highly imbalanced scenarios such as Scenario 1 and 2, where the performance gap between majority and minority classes is significantly narrowed. Notably, the performance for the majority class remains stable, indicating that the improvement for underrepresented classes does not come at the cost of degrading other class performance. These results highlight the effectiveness of our adaptive strategy in improving class-wise balance while preserving overall accuracy.

In summary, the experimental results clearly demonstrate the effectiveness of the proposed adaptive contrastive learning framework in addressing class imbalance. By dynamically adjusting loss weights and contrastive temperature based on the class distribution, the model achieves more balanced performance across all classes. The improvements in G-Mean and the reduced variance of F1-scores confirm that the framework mitigates the bias toward majority classes while preserving overall accuracy. These findings validate the proposed method as a simple yet effective extension to existing classification models in imbalanced settings.

## IV. Conclusion

This paper presents an adaptive contrastive learning framework designed to mitigate performance bias in class-imbalanced classification tasks. Unlike conventional methods that rely on resampling or architectural modifications, our approach integrates seamlessly with existing models and datasets by dynamically combining a class-weighted cross-entropy loss with a temperature-scaled supervised contrastive loss. The degree of imbalance in the data is quantified and used to control key hyperparameters, including the temperature and loss weighting, in an adaptive manner.

Through extensive experiments on the AIDER dataset under various synthetic imbalance scenarios, we demonstrate that the proposed framework improves minority class performance, reduces class-wise variance, and maintains high overall accuracy. These results highlight the practical applicability of our method in real-world settings where class imbalance is prevalent.

## References


[1] Z. Li, B. Chen, S. Wu, M. Su, J. M. Chen, and B. Xu, "Deep learning for urban land use category classification: A review and experimental assessment," *Remote Sensing of Environment*, vol. 311, p. 114290, 2024.

[2] J. M. Johnson and T. M. Khoshgoftaar, "Survey on deep learning with class imbalance," *Journal of big data*, vol. 6, no. 1, pp. 1–54, 2019.

[3] K. Dimitropoulos and et al., "Machine learning in disaster management: Recent developments and challenges," *Technologies*, vol. 8, no. 4, p. 77, 2020.

[4] S. Goyal, "Handling class-imbalance with knn (neighbourhood) undersampling for software defect prediction," *Artificial Intelligence Review*, vol. 55, no. 3, pp. 2023–2064, 2022.

[5] J. Kuang, G. Xu, T. Tao, and Q. Wu, "Class-imbalance adversarial transfer learning network for cross-domain fault diagnosis with imbalanced data," *IEEE Transactions on Instrumentation and Measurement*, vol. 71, pp. 1–11, 2021.

[6] U. Perez-Cerrolaza *et al.*, "Artificial intelligence for safety-critical systems in industrial and transportation domains: A survey," *ACM Computing Surveys*, 2023.

[7] Z. Farou, Y. Wang, and T. Horváth, "Cluster-based oversampling with area extraction from representative points for class imbalance learning," *Intelligent Systems with Applications*, vol. 22, p. 200357, 2024.

[8] K. Ghosh, C. Bellinger, R. Corizzo, P. Branco, B. Krawczyk, and N. Japkowicz, "The class imbalance problem in deep learning," *Machine Learning*, vol. 113, no. 7, pp. 4845–4901, 2024.

[9] S. Rezvani and X. Wang, "A broad review on class imbalance learning techniques," *Applied Soft Computing*, vol. 143, p. 110415, 2023.

[10] P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan, "Supervised contrastive learning," in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 18 661–18 673.

[11] C. Kyrkou and T. Theocharides, "Deep-learning-based aerial image classification for emergency response applications using unmanned aerial vehicles." in *CVPR workshops*, 2019, pp. 517–525.

[12] A. Howard, M. Sandler, G. Chu, L.-C. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan *et al.*, "Searching for mobilenetv3," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 1314–1324.