

Rare Event Detection in Imbalanced Multi-Class Datasets Using an Optimal MIP-Based Ensemble Weighting Approach

Georgios Tertytchny^{1,2,*†}, Georgios L. Stavrinides^{1†}, Maria K. Michael^{1,2}

¹KIOS Research and Innovation Center of Excellence

²Department of Electrical and Computer Engineering, University of Cyprus
{tertytchny.georgios, stavrinides.georgios, mmichael}@ucy.ac.cy

Abstract

To address the challenges of imbalanced multi-class datasets typically used for rare event detection in critical cyber-physical systems, we propose an optimal, efficient, and adaptable mixed integer programming (MIP) ensemble weighting scheme. Our approach leverages the diverse capabilities of the classifier ensemble on a granular per class basis, while optimizing the weights of classifier-class pairs using elastic net regularization for improved robustness and generalization. Additionally, it seamlessly and optimally selects a predefined number of classifiers from a given set. We evaluate and compare our MIP-based method against six well-established weighting schemes, using representative datasets and suitable metrics, under various ensemble sizes. The experimental results reveal that MIP outperforms all existing approaches, achieving an improvement in balanced accuracy ranging from 0.99% to 7.31%, with an overall average of 4.53% across all datasets and ensemble sizes. Furthermore, it attains an overall average increase of 4.63%, 4.60%, and 4.61% in macro-averaged precision, recall, and F1-score, respectively, while maintaining computational efficiency.

Code — <https://github.com/gterty019/MIPENS>

Extended version — <https://arxiv.org/abs/2412.13439>

Introduction

Rare event detection in cyber-physical systems (CPS) is vital for maintaining the safety and reliability of critical infrastructures, such as water distribution networks and power grids (Taheri et al. 2024; Tertytchny, Nicolaou, and Michael 2020). Critical CPS often generate imbalanced multi-class datasets where some classes have significantly fewer instances than others, due to the infrequent occurrence of abnormal events (e.g., faults or attacks) compared to normal activities (Yin et al. 2020). Class imbalance poses several challenges, such as biased model performance towards the majority classes, poor generalization to minority classes, and thus increased difficulty in detecting rare events (Shyalika, Wickramarachchi, and Sheth 2023). These issues are further amplified by the critical nature of CPS, where misclassifying low-probability but high-impact events may lead

to safety risks or system failures. Consequently, addressing these challenges is crucial for effective rare event detection.

Ensemble learning approaches, especially weighted voting techniques, are commonly employed to handle imbalanced multi-class datasets (Khan, Chaudhari, and Chandra 2024). The aim of such methods is to improve overall performance and robustness, by leveraging the strengths of diverse classification algorithms. Specifically, a weighted voting ensemble model combines the predictions of multiple classifiers, using different weights based on their predictive performance. Higher weights are assigned to better-performing classifiers, either at an overall classifier level or on a per class basis (Dogan and Birant 2019). However, existing weighting schemes often fail to optimally balance the contribution of each classifier, especially in the presence of a highly skewed class distribution. Moreover, since smaller ensembles are typically less computationally demanding, and given that rare event detection in CPS is often performed at the network edge on resource-constrained devices, adaptable approaches are required to seamlessly and optimally select a predefined number of classifiers from a larger set while assigning their weights (Danso et al. 2022).

Aiming to fill this gap, we propose an optimal, adaptable, and efficient mixed integer programming (MIP) weighting scheme for weighted voting ensembles, suitable for rare event detection in imbalanced multi-class datasets. Our approach optimally assigns weights to classifiers per class based on their performance, leveraging their individual capabilities while selecting the predefined number of classifiers from a given set. To prevent overfitting and enhance generalization in weight assignment, we incorporate into our method the elastic net regularization technique (Zou and Hastie 2005), applying it beyond its traditional use in the training and validation phase. The proposed weighting scheme aims to enhance performance across all classes while ensuring computational efficiency, effectively addressing the limitations of existing approaches. Our method optimizes weights based on empirical accuracy results, avoiding simplifying assumptions typically required by theoretical analysis, such as independent and identically distributed instances, which may not hold in realistic scenarios involving class imbalance and diverse classifier behavior (Opitz and Maclin 1999).

Our main contributions are summarized as follows:

*Corresponding author.

†These authors contributed equally.

- We propose an optimal MIP-based ensemble weighting approach for imbalanced multi-class datasets to improve rare event detection in critical CPS. Our weighting scheme optimizes classifier weights on a per class basis, leveraging the diverse strengths of each classifier in a granular manner.
- The proposed method seamlessly and optimally selects a predefined number of classifiers from a given set while calculating their weights. Moreover, it utilizes elastic net regularization to improve generalization and robustness in the assignment of classifier weights, while maintaining computational efficiency.
- We demonstrate the improved performance, adaptability, and efficiency achieved by our approach through comparative evaluation against six well-established weighting schemes, using representative datasets and appropriate metrics under different ensemble sizes.

To our knowledge, in contrast to our approach, none of the existing ensemble weighting schemes are intrinsically adaptable to varied ensemble sizes and capable of optimally assigning classifier weights for imbalanced multi-class datasets, while attaining generalization and efficiency.

Related Work

Weighted voting ensemble methods typically apply weights per classifier or per classifier and class, based on classifier performance in the training and validation phase (Maheshwari et al. 2022). Among the fundamental techniques, UW-PC (uniform weights per classifier) and UW-PCC (uniform weights per classifier and class) assign equal weights to all classifiers and classifier-class pairs, respectively, and commonly serve as baselines in comparative studies (Dong et al. 2020). On the other hand, WA-PC (weighted average based on normalized accuracy per classifier) and WA-PCC (weighted average based on normalized accuracy per classifier and class) apply weights based on the normalized accuracy of each classifier or classifier-class pair, respectively (Sagi and Rokach 2018). Both techniques assign greater weights to classifiers or classifier-class pairs that exhibit better performance. DE (differential evolution) is an evolutionary algorithm that determines classifier weights in an iterative refining manner (Storn and Price 1997; Zhang et al. 2014, 2024). In contrast, BMA (Bayesian model averaging) assigns weights to classifier-class pairs based on their Bayesian inference posterior probabilities, which reflect the likelihood that each classifier is accurate given the observed data (Raftery, Madigan, and Hoeting 1997; Wang et al. 2022).

However, these weighting schemes do not optimally assign classifier weights, are not easily adaptable to a varied number of classifiers, and fail to address highly uneven class distributions while achieving generalization. In contrast, our MIP approach attains generalization in weight calculation through elastic net regularization, which is a regularized regression technique. It linearly combines the widely used L1 and L2 penalties of the lasso and ridge methods, respectively, to leverage the advantages of both strategies (Zou and

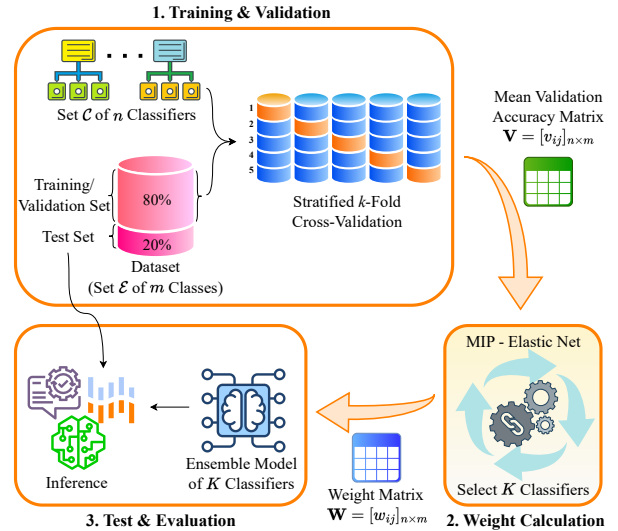


Figure 1: Overview of methodology for incorporating the proposed weighting scheme into the ensemble model.

Hastie 2005). This dual technique can lead to a weight assignment that is both sparse (less likely to overfit) and robust (less sensitive to training data noise) (Giesen et al. 2019).

Previous studies have explored different perspectives and strategies to address the challenges of imbalanced datasets. A comprehensive review of the problem, potential solutions, and applicable evaluation metrics are examined in (He and Garcia 2009). A general evaluation framework is proposed in (Khan, Chaudhari, and Chandra 2024), employing data augmentation and ensemble techniques by analyzing binary datasets with varied class imbalance. In (Lango and Stefanowski 2022), the impact of multi-class imbalanced data difficulty factors (such as class overlapping and skewed distribution) on classifier performance is examined. On the other hand, two approaches integrating sampling, data space improvement, ensemble, and self-paced learning are introduced in (Liu et al. 2020) and (Yin et al. 2020) to handle class imbalance. However, both methods are tailored to binary classification. Overall, previous studies lack an optimal and adaptable technique for classifier weight assignment to enhance ensemble performance in imbalanced multi-class datasets, while achieving generalization and computational efficiency across varied ensemble sizes.

Proposed Weighting Scheme

To incorporate the proposed weighting approach into the ensemble model, we employ the methodology outlined in Figure 1. Specifically, given a set of n classifiers $\mathcal{C} = \{C_1, C_2, \dots, C_n\}$, a set of m classes $\mathcal{E} = \{E_1, E_2, \dots, E_m\}$, and a predefined ensemble size K , our methodology consists of three phases:

1. **Training & Validation:** The input dataset corresponding to \mathcal{E} is used to train each classifier in \mathcal{C} , using stratified k -fold cross-validation to ensure a similar class distribution in each fold. This yields the mean validation accuracy matrix $\mathbf{V} = [v_{ij}]_{n \times m}$, where v_{ij} denotes the mean

validation accuracy of classifier C_i for class E_j across all folds. It is noted that $\mathcal{E}_N \cup \mathcal{E}_A = \mathcal{E}$ and $\mathcal{E}_N \cap \mathcal{E}_A = \emptyset$, where $\mathcal{E}_N \subset \mathcal{E}$ and $\mathcal{E}_A \subset \mathcal{E}$ represent the normal and abnormal (e.g., faults/attacks) classes, respectively, in \mathcal{E} .

2. **Weight Calculation:** We use \mathbf{V} to determine the weights for each classifier-class pair. The proposed approach utilizes MIP and elastic net regularization to select the predefined number of classifiers K from the initial set \mathcal{C} , while optimally calculating their weights per class. This results in a weight matrix $\mathbf{W} = [w_{ij}]_{n \times m}$, where w_{ij} denotes the weight of classifier C_i for class E_j . If a classifier C_i is not selected, then $w_{ij} = 0 \forall E_j \in \mathcal{E}$.
3. **Test & Evaluation:** The calculated weights in \mathbf{W} are assigned to the selected K classifiers to evaluate the performance of the resulting ensemble model.

The main problem addressed—determining the optimal class-based weight assignment for classifiers in a weighted voting ensemble model—corresponds to phase 2.

Problem Formulation

To solve the examined problem in phase 2, we formulate it as a MIP model by defining its decision variables, objective function, and constraints.

Decision Variables. The following variables are used to select K classifiers and to determine their optimal weights:

- x_i : binary variable corresponding to classifier C_i , such that $x_i = 1$ if C_i is selected, and $x_i = 0$ otherwise ($\mathbf{X} = [x_i]_{n \times 1}$).
- w_{ij} : non-negative continuous variable representing the (i, j) -th element of weight matrix \mathbf{W} , denoting the weight of classifier-class pair (C_i, E_j) .

Objective. We aim to maximize the overall classification accuracy of the ensemble while ensuring generalization, i.e.:

$$\max_{\mathbf{X}, \mathbf{W}} \left[\frac{1}{m} \sum_{C_i \in \mathcal{C}} \sum_{E_j \in \mathcal{E}} w_{ij} v_{ij} - \lambda \left(\alpha \sum_{C_i \in \mathcal{C}} \sum_{E_j \in \mathcal{E}} w_{ij} + \frac{1-\alpha}{2} \sum_{C_i \in \mathcal{C}} \sum_{E_j \in \mathcal{E}} w_{ij}^2 \right) \right] \quad (1)$$

subject to the constraints defined in the next subsection.

In (1), the overall accuracy is denoted by the first term, where v_{ij} represents the (i, j) -th element of mean validation accuracy matrix \mathbf{V} , denoting the mean accuracy of classifier C_i for class E_j in the training and validation phase (phase 1). To prevent the ensemble model from overly relying on specific classifier-class combinations based on the accuracy attained in phase 1, and thus achieve a more balanced and generalizable model, we incorporate into our objective the elastic net regularization technique, represented by the second term in (1). The L1 and L2 penalties are denoted by $\sum_{C_i \in \mathcal{C}} \sum_{E_j \in \mathcal{E}} w_{ij}$ and $\sum_{C_i \in \mathcal{C}} \sum_{E_j \in \mathcal{E}} w_{ij}^2$, respectively. The L1 penalty promotes sparsity by driving some weights to zero, while the L2 penalty results in smaller and more

evenly distributed weights. Consequently, the combination of both penalties leads to a sparse and robust weight assignment. In (1), λ is the regularization parameter that controls the strength of the overall penalty applied to the objective function ($\lambda \geq 0$). On the other hand, α is the mixing parameter that balances the contribution of the L1 and L2 penalties in the elastic net regularization ($0 \leq \alpha \leq 1$).

Constraints. Our objective function (1) is solved subject to the following constraints:

$$x_i \in \{0, 1\}, \forall C_i \in \mathcal{C}, \quad (2)$$

$$w_{ij} \geq 0, \forall C_i \in \mathcal{C}, \forall E_j \in \mathcal{E}, \quad (3)$$

$$\sum_{C_i \in \mathcal{C}} x_i = K, \quad (4)$$

$$\sum_{C_i \in \mathcal{C}} w_{ij} = 1, \forall E_j \in \mathcal{E}, \quad (5)$$

$$\sum_{E_j \in \mathcal{E}} w_{ij} \leq m x_i, \forall C_i \in \mathcal{C}, \quad (6)$$

$$\sum_{E_j \in \mathcal{E}} w_{ij} + M(1 - x_i) \geq \epsilon, \forall C_i \in \mathcal{C}, \quad (7)$$

$$\sum_{C_i \in \mathcal{C}} w_{ij} v_{ij} \geq \frac{1}{n} \sum_{C_i \in \mathcal{C}} v_{ij} + \epsilon, \forall E_j \in \mathcal{E}, \quad (8)$$

$$\frac{1}{m} \sum_{C_i \in \mathcal{C}} \sum_{E_j \in \mathcal{E}} w_{ij} v_{ij} \geq \frac{1}{n m} \sum_{C_i \in \mathcal{C}} \sum_{E_j \in \mathcal{E}} v_{ij} + \epsilon. \quad (9)$$

Constraints (2) and (3) ensure the binary nature and non-negativity of variables x_i and w_{ij} , respectively. Constraint (4) enforces the selection of the predefined number of classifiers K . On the other hand, constraint (5) guarantees that for each class, the sum of weights of all classifiers is equal to one. If a classifier is not selected, (6) enforces all of its weights across all classes to be equal to zero (otherwise, the sum of its weights is bounded by the total number of classes m). In contrast, constraint (7) ensures that if a classifier is selected, at least one of its weights across all classes is greater than zero. In (7), ϵ is a positive constant, sufficiently smaller than the variables and parameters in the MIP model. It is used to convert (7) to a non-strict inequality, as strict inequalities are not supported in MIP.

On the other hand, M is a positive constant, sufficiently larger than the variables and parameters in the MIP model. It is used to express the conditional aspect of (7) in linear form. Specifically, the constraint becomes relevant ($\sum_{E_j \in \mathcal{E}} w_{ij} \geq \epsilon$) only if $x_i = 1$. Otherwise, if $x_i = 0$ the constraint becomes irrelevant ($\sum_{E_j \in \mathcal{E}} w_{ij} + M \geq \epsilon$), as it is always true. Constraint (8) ensures that the selected weights per classifier and class are such that the resulting weighted average accuracy for each class is greater than the average accuracy for the particular class when uniform weights are utilized across all classifiers. Similarly, (9) guarantees that the overall weighted average accuracy of the ensemble is greater than the overall average accuracy in the case where all weights of classifier-class pairs are uniform. In (8) and (9), ϵ is utilized in the same manner as in (7).

The time complexity of MIP models depends on both the problem size and the solver used. MIP solvers typically employ proprietary algorithms with undisclosed implementation details, and thus their time complexity cannot be easily derived. However, the computational cost of the proposed MIP approach can be considered proportional to the problem size nm , where n (number of classifiers) is small in practice (Dogan and Birant 2019). Consequently, the computational cost grows approximately linearly with respect to m (number of classes). In CPS rare event detection, m is typically small (Vrachimis et al. 2018), as the aim is to distinguish between normal and rare abnormal states rather than to conduct a fine-grained classification. It is noted that our MIP model is independent of the dataset size. These factors ensure that the runtime of our MIP method remains reasonable, as demonstrated in our experimental results.

Experimental Framework

We evaluated and compared the performance and efficiency of the proposed weighting scheme against six well-established approaches. The comparative evaluation was conducted using appropriate metrics and four relevant imbalanced multi-class datasets for rare event detection in critical CPS, under varied ensemble sizes.

Experimental Setup

The three phases of our methodology were implemented using a server running CentOS 7.9, equipped with an Intel Xeon Gold 6240 processor @ 2.6GHz and 400 GB of RAM. The software environment included Python 3.12 with scipy 1.14.0, scikit-learn 1.5.1, NumPy 2.0, pandas 2.2.2, and Gurobi Optimizer 11.0.3 with its Python interface gurobipy. We utilized the Gurobi solver in phase 2 to solve our MIP-formulated model, as it is widely used to optimally solve similar optimization problems (Gurobi, LLC 2024). We split each examined dataset into 80% for training and validation, used in phase 1, and 20% for the test set, used in phase 3. In phase 1, we employed stratified k -fold cross-validation with $k = 5$ (McTavish et al. 2022). To fine-tune parameters λ and α in (1), we incremented/decremented each parameter, terminating adjustments when performance began to deteriorate. This fine-tuning resulted in the following values for (λ, α) for each of the four examined datasets: $(0.95, 0.85)$, $(0.96, 0.80)$, $(1.00, 0.82)$, and $(0.95, 0.86)$, respectively.

Existing Weighting Schemes

We compared our MIP approach against the six widely used weighting methods discussed in related work: UW-PC, UW-PCC, WA-PC, WA-PCC, DE, and BMA. In our experiments, UW-PC and UW-PCC served as baselines. Without loss of generality, we used the mean validation accuracy matrix \mathbf{V} as the likelihood matrix in BMA (Zhou 2021).

Classifiers

We considered a set of $n = 8$ commonly used classifiers $\mathcal{C} = \{\text{MLR, J48, JRIP, REPTree, MLP, SVM, GNB, IBk}\}$. These classifiers were selected for their complementary

| E_j | Class Description | # Instances | % |
|---|--------------------------------------|-------------|--------|
| Dataset D1: LeakDB ($\rho = 6059.97$) | | | |
| N_1 | Normal - No Leak/No Pressure | 17 520 | 3.14 |
| N_2 | Normal - No Leak/High Pressure | 448 438 | 80.38 |
| N_3 | Normal - No Leak/Low Pressure | 88 154 | 15.80 |
| F_1 | Abrupt - Large Leak/High Pressure | 1 721 | 0.31 |
| F_2 | Incipient - Large Leak/High Pressure | 1 181 | 0.21 |
| F_3 | Incipient - Large Leak/Low Pressure | 812 | 0.15 |
| F_4 | Incipient - Small Leak/Low Pressure | 74 | 0.01 |
| Total | | 557 900 | 100.00 |
| Dataset D2: NSL-KDD ($\rho = 1222.64$) | | | |
| N_1 | Normal | 13 449 | 53.39 |
| A_1 | Probe | 2 289 | 9.09 |
| A_2 | Denial of Service (DoS) | 9 234 | 36.65 |
| A_3 | User to Root (U2R) | 11 | 0.04 |
| A_4 | Remote to Local (R2L) | 209 | 0.83 |
| Total | | 25 192 | 100.00 |
| Dataset D3: SG-MITM ($\rho = 5.91$) | | | |
| N_1 | Normal | 455 | 65.28 |
| A_1 | MITM on Router-Controller | 84 | 12.05 |
| A_2 | MITM on Controller | 77 | 11.05 |
| A_3 | MITM on Router | 81 | 11.62 |
| Total | | 697 | 100.00 |
| Dataset D4: CIC-IDS2017 ($\rho = 2.13$) | | | |
| N_1 | Normal | 273 097 | 34.52 |
| A_1 | Distributed Denial of Service (DDoS) | 128 027 | 16.18 |
| A_2 | Denial of Service Hulk (DoS Hulk) | 231 073 | 29.20 |
| A_3 | PortScan | 158 930 | 20.10 |
| Total | | 791 127 | 100.00 |

Table 1: Class distribution in datasets D1–D4 ($N_j \in \mathcal{E}_N$ and $A_j, F_j \in \mathcal{E}_A$).

strengths in multi-class classification, as they encompass a diverse range of approaches: probabilistic (GNB), rule-based (JRIP), tree-based (J48, REPTree), function-based (MLR, MLP, SVM), and instance-based (IBk) (Hastie, Tibshirani, and Friedman 2009). All classifiers are available in the scikit-learn 1.5.1 Python library. In our experiments, we used ensemble sizes ranging from $K = 2$ to $K = 8$ to investigate a comprehensive set of scenarios.

Datasets

We experimented with four representative datasets for rare event detection in critical CPS, featuring multiple classes with varied imbalance: [D1] LeakDB (Vrachimis et al. 2018), [D2] NSL-KDD (Tavallaee et al. 2009), [D3] SG-MITM (Elrawy et al. 2023), and [D4] CIC-IDS2017 (Sharafaldin, Lashkari, and Ghorbani 2018). Benchmarks D2 and D4 pertain to intrusion detection in simulated network environments, resembling real-world network traffic captures. They are widely used to investigate anomaly detection in CPS (Zhou et al. 2020; Danso et al. 2022). On the other hand, D1 involves leakages (faults) in a simulated real-world water distribution network under varying conditions, while D3 focuses on man-in-the-middle (MITM) attacks targeting client-server protocols in the private area network (PAN) of a smart grid CPS, represented by an experimental testbed. Table 1 shows the description, number of instances, and the percentage distribution of each normal/ab-




Figure 2: Improvement achieved in balanced accuracy by proposed MIP approach over other weighting schemes (D1).

| | | K | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
|--------------|-----------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| Increase (%) | Precision | UW-PC | 4.06 | 4.59 | 4.93 | 5.34 | 5.47 | 5.59 | 5.88 |
| | | UW-PCC | 3.83 | 4.23 | 4.45 | 4.84 | 4.97 | 5.21 | 5.62 |
| | | WA-PC | 3.71 | 4.11 | 4.57 | 4.96 | 4.97 | 5.09 | 5.49 |
| | | WA-PCC | 3.48 | 3.76 | 4.33 | 4.72 | 4.85 | 4.96 | 5.10 |
| | | DE | 2.90 | 3.17 | 3.51 | 3.74 | 3.74 | 3.86 | 4.09 |
| | | BMA | 3.48 | 3.88 | 4.45 | 4.96 | 4.85 | 5.09 | 5.23 |
| | | Avg | 3.58 | 3.96 | 4.37 | 4.76 | 4.81 | 4.97 | 5.24 |
| | Recall | UW-PC | 4.02 | 4.48 | 5.04 | 5.47 | 5.57 | 5.67 | 5.78 |
| | | UW-PCC | 3.80 | 4.25 | 4.45 | 4.89 | 5.11 | 5.32 | 5.54 |
| | | WA-PC | 3.68 | 4.13 | 4.57 | 5.01 | 5.11 | 5.20 | 5.43 |
| | | WA-PCC | 3.46 | 3.79 | 4.34 | 4.77 | 4.87 | 4.97 | 5.08 |
| | | DE | 2.90 | 3.21 | 3.53 | 3.74 | 3.84 | 3.83 | 4.04 |
| | | BMA | 3.46 | 3.90 | 4.45 | 5.01 | 4.99 | 5.09 | 5.19 |
| | | Avg | 3.55 | 3.96 | 4.40 | 4.82 | 4.92 | 5.01 | 5.18 |
| | F1-Score | UW-PC | 3.99 | 4.54 | 4.99 | 5.41 | 5.53 | 5.64 | 5.72 |
| | | UW-PCC | 3.76 | 4.19 | 4.40 | 4.81 | 5.05 | 5.28 | 5.47 |
| | | WA-PC | 3.64 | 4.07 | 4.51 | 4.93 | 5.05 | 5.15 | 5.47 |
| | | WA-PCC | 3.41 | 3.72 | 4.28 | 4.69 | 4.81 | 4.91 | 5.10 |
| DE | | 2.96 | 3.14 | 3.46 | 3.74 | 3.85 | 3.85 | 4.02 | |
| BMA | | 3.41 | 3.83 | 4.40 | 4.93 | 4.93 | 5.03 | 5.23 | |
| Avg | | 3.53 | 3.92 | 4.34 | 4.75 | 4.87 | 4.98 | 5.17 | |

Table 2: Improvement attained in macro-averaged precision, recall, and F1-score by proposed MIP approach over other weighting schemes (D1).

normal class in datasets D1–D4. Furthermore, it showcases the imbalance ratio ρ for each dataset, which is defined as the ratio of the number of instances μ in the majority class to the number of instances ν in the minority class (Wang et al. 2021). Considering that a dataset is highly imbalanced when $\mu \gg \nu$, D1 and D2 are highly imbalanced, whereas D3 and D4 are moderately imbalanced (Liu et al. 2020).

Evaluation Metrics

To provide a comprehensive assessment of our approach against the examined weighting schemes, we employed in test and evaluation phase (phase 3) performance metrics well-suited for imbalanced multi-class datasets. Specifically, balanced accuracy served as our primary metric due to its ability to account for class imbalance by averaging the recall obtained for each class. Additionally, we utilized macro-averaged precision, recall, and F1-score as secondary metrics to ensure an unbiased evaluation across all classes, thus



Figure 3: Improvement achieved in balanced accuracy by proposed MIP approach over other weighting schemes (D2).

| | | K | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
|--------------|-----------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| Increase (%) | Precision | UW-PC | 3.25 | 3.84 | 4.37 | 4.87 | 4.97 | 5.06 | 5.04 |
| | | UW-PCC | 3.03 | 3.62 | 4.04 | 4.76 | 4.97 | 4.95 | 4.93 |
| | | WA-PC | 2.92 | 3.51 | 3.82 | 4.43 | 4.42 | 4.51 | 4.49 |
| | | WA-PCC | 2.69 | 3.29 | 3.60 | 4.21 | 4.09 | 4.18 | 4.17 |
| | | DE | 1.06 | 1.46 | 1.77 | 2.17 | 2.16 | 2.15 | 2.15 |
| | | BMA | 1.60 | 2.10 | 2.62 | 3.34 | 3.44 | 3.64 | 3.63 |
| | | Avg | 2.43 | 2.97 | 3.37 | 3.96 | 4.01 | 4.08 | 4.07 |
| | Recall | UW-PC | 3.23 | 3.79 | 4.36 | 4.81 | 4.80 | 5.02 | 5.13 |
| | | UW-PCC | 2.99 | 3.56 | 4.12 | 4.69 | 4.80 | 4.90 | 5.01 |
| | | WA-PC | 2.87 | 3.44 | 3.89 | 4.45 | 4.32 | 4.54 | 4.65 |
| | | WA-PCC | 2.63 | 3.20 | 3.65 | 4.21 | 4.09 | 4.19 | 4.30 |
| | | DE | 1.01 | 1.46 | 1.79 | 2.23 | 2.12 | 2.22 | 2.33 |
| | | BMA | 1.59 | 2.15 | 2.71 | 3.27 | 3.38 | 3.72 | 3.83 |
| | | Avg | 2.39 | 2.93 | 3.42 | 3.94 | 3.92 | 4.10 | 4.21 |
| | F1-Score | UW-PC | 3.24 | 3.88 | 4.43 | 4.85 | 4.84 | 5.16 | 5.15 |
| | | UW-PCC | 3.01 | 3.65 | 4.20 | 4.74 | 4.84 | 5.04 | 5.03 |
| | | WA-PC | 2.90 | 3.54 | 3.97 | 4.39 | 4.38 | 4.59 | 4.58 |
| | | WA-PCC | 2.78 | 3.31 | 3.74 | 4.16 | 4.04 | 4.24 | 4.23 |
| DE | | 1.09 | 1.52 | 1.84 | 2.15 | 2.03 | 2.24 | 2.24 | |
| BMA | | 1.65 | 2.18 | 2.72 | 3.26 | 3.36 | 3.79 | 3.78 | |
| Avg | | 2.45 | 3.01 | 3.48 | 3.93 | 3.92 | 4.18 | 4.17 | |

Table 3: Improvement attained in macro-averaged precision, recall, and F1-score by proposed MIP approach over other weighting schemes (D2).

providing a holistic comparative analysis (Grandini, Bagli, and Visani 2020; Lango and Stefanowski 2022).

Results & Discussion

Figures 2–5 demonstrate the percentage improvement achieved in balanced accuracy (primary metric) by the proposed MIP weighting scheme over all existing approaches for datasets D1–D4, respectively, for $K = 2$ to $K = 8$ classifiers selected from set \mathcal{C} . Tables 2–5 showcase the percentage improvement attained with respect to the secondary metrics (macro-averaged precision, recall, and F1-score) in each case. The average improvement per dataset is also shown in Table 6. Notably, MIP consistently outperformed the other techniques across all evaluation metrics, ensemble sizes, and datasets, demonstrating superior generalization and robustness. Specifically, MIP yielded an increase in balanced accuracy ranging from 0.99% to 7.31%, with an overall average of 4.53% across all datasets and ensemble sizes. Moreover,




Figure 4: Improvement achieved in balanced accuracy by proposed MIP approach over other weighting schemes (D3).

| | | K | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
|--------------|-----------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| Increase (%) | Precision | UW-PC | 5.02 | 5.83 | 6.48 | 6.46 | 6.93 | 6.88 | 6.99 |
| | | UW-PCC | 4.89 | 5.70 | 6.35 | 6.33 | 6.80 | 6.63 | 6.72 |
| | | WA-PC | 4.64 | 5.44 | 5.83 | 5.82 | 6.28 | 6.37 | 6.46 |
| | | WA-PCC | 4.27 | 5.06 | 5.45 | 5.31 | 5.77 | 5.86 | 5.80 |
| | | DE | 3.29 | 4.18 | 4.31 | 3.93 | 4.51 | 4.61 | 4.51 |
| | | BMA | 3.90 | 4.81 | 5.19 | 5.05 | 5.52 | 5.73 | 5.80 |
| | | Avg | 4.34 | 5.17 | 5.60 | 5.48 | 5.97 | 6.01 | 6.05 |
| | Recall | UW-PC | 5.02 | 5.84 | 6.42 | 6.41 | 6.97 | 6.82 | 6.92 |
| | | UW-PCC | 4.89 | 5.71 | 6.29 | 6.28 | 6.85 | 6.57 | 6.67 |
| | | WA-PC | 4.64 | 5.34 | 5.92 | 5.78 | 6.35 | 6.19 | 6.42 |
| | | WA-PCC | 4.27 | 5.09 | 5.42 | 5.28 | 5.85 | 5.82 | 5.80 |
| | | DE | 3.29 | 4.10 | 4.31 | 3.94 | 4.50 | 4.60 | 4.47 |
| | | BMA | 3.90 | 4.84 | 5.17 | 5.04 | 5.60 | 5.70 | 5.68 |
| | | Avg | 4.34 | 5.15 | 5.59 | 5.46 | 6.02 | 5.95 | 5.99 |
| | F1-Score | UW-PC | 5.02 | 5.90 | 6.51 | 6.50 | 6.83 | 6.91 | 6.90 |
| | | UW-PCC | 4.89 | 5.77 | 6.39 | 6.37 | 6.70 | 6.66 | 6.64 |
| | | WA-PC | 4.64 | 5.39 | 5.88 | 5.73 | 6.32 | 6.28 | 6.39 |
| | | WA-PCC | 4.27 | 5.01 | 5.50 | 5.36 | 5.81 | 5.90 | 5.75 |
| DE | | 3.29 | 4.14 | 4.25 | 4.00 | 4.45 | 4.55 | 4.37 | |
| BMA | | 3.90 | 4.89 | 5.24 | 5.11 | 5.44 | 5.78 | 5.62 | |
| Avg | | 4.34 | 5.18 | 5.63 | 5.51 | 5.93 | 6.01 | 5.95 | |

Table 4: Improvement attained in macro-averaged precision, recall, and F1-score by proposed MIP approach over other weighting schemes (D3).

it attained an overall average increase of 4.63%, 4.60%, and 4.61% in macro-averaged precision, recall, and F1-score, respectively. In addition to these metrics, MIP also outperformed all other schemes with respect to the macro-averaged area under the precision-recall curve (AUPRC), as detailed in the provided supplementary material.

Furthermore, it can be observed that the improvement provided in all performance metrics by our approach over the other schemes generally increased as the ensemble size K increased. This reveals that MIP benefits more from a larger number of classifiers compared to the other methods. By optimally assigning weights to classifier-class pairs, our approach enhances the ability of the ensemble to leverage the diverse capabilities of each classifier, leading to a better performance as the number of classifiers grows. On the other hand, the performance improvement was lower for the highly imbalanced datasets D1 and D2 (which had a higher imbalance ratio), compared to the moderately im-




Figure 5: Improvement achieved in balanced accuracy by proposed MIP approach over other weighting schemes (D4).

| | | K | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
|--------------|-----------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| Increase (%) | Precision | UW-PC | 1.68 | 4.74 | 5.13 | 5.84 | 6.54 | 7.32 | 6.89 |
| | | UW-PCC | 5.81 | 4.62 | 5.13 | 5.71 | 6.41 | 7.05 | 6.64 |
| | | WA-PC | 5.45 | 4.13 | 4.49 | 4.96 | 5.29 | 5.75 | 6.38 |
| | | WA-PCC | 5.20 | 3.76 | 3.61 | 3.74 | 3.94 | 4.22 | 5.88 |
| | | DE | 4.12 | 2.91 | 3.36 | 3.74 | 4.19 | 4.59 | 4.65 |
| | | BMA | 4.72 | 3.40 | 3.86 | 4.47 | 5.04 | 5.62 | 5.76 |
| | | Avg | 4.50 | 3.93 | 4.26 | 4.74 | 5.24 | 5.76 | 6.03 |
| | Recall | UW-PC | 4.38 | 4.82 | 5.22 | 5.82 | 6.67 | 7.31 | 6.87 |
| | | UW-PCC | 4.14 | 4.70 | 5.22 | 5.69 | 6.54 | 7.08 | 6.63 |
| | | WA-PC | 3.78 | 4.10 | 4.52 | 4.86 | 5.33 | 5.70 | 6.39 |
| | | WA-PCC | 3.54 | 3.75 | 3.72 | 3.80 | 4.04 | 4.24 | 5.90 |
| | | DE | 2.37 | 2.82 | 3.38 | 3.80 | 4.27 | 4.57 | 4.60 |
| | | BMA | 2.95 | 3.40 | 3.95 | 4.50 | 5.10 | 5.59 | 5.78 |
| | | Avg | 3.53 | 3.93 | 4.33 | 4.75 | 5.32 | 5.75 | 6.03 |
| | F1-Score | UW-PC | 3.06 | 4.67 | 5.13 | 5.71 | 6.67 | 7.23 | 6.88 |
| | | UW-PCC | 4.97 | 4.55 | 5.13 | 5.59 | 6.54 | 7.10 | 6.64 |
| | | WA-PC | 4.61 | 4.06 | 4.40 | 4.86 | 5.31 | 5.63 | 6.26 |
| | | WA-PCC | 4.37 | 3.70 | 3.68 | 3.78 | 3.99 | 4.19 | 5.90 |
| DE | | 3.30 | 2.75 | 3.32 | 3.78 | 4.23 | 4.55 | 4.69 | |
| BMA | | 3.89 | 3.34 | 3.92 | 4.37 | 5.07 | 5.50 | 5.77 | |
| Avg | | 4.03 | 3.84 | 4.26 | 4.68 | 5.30 | 5.70 | 6.02 | |

Table 5: Improvement attained in macro-averaged precision, recall, and F1-score by proposed MIP approach over other weighting schemes (D4).

balanced datasets D3 and D4 (which had a lower imbalance ratio). This highlights the negative impact of class imbalance on the ensemble performance under all weighting schemes. However, as demonstrated in Table 6, the performance gains for the highly imbalanced datasets were still significant and comparable to those observed for the moderately imbalanced datasets. The effectiveness and robustness of our method under different cases of class imbalance are further demonstrated through the ablation studies included in the supplementary material. As reported by the Gurobi solver, the number of variables and constraints required for each dataset by our MIP approach are shown in Table 6.

Weight Assignment Example

To demonstrate the effective utilization of the strengths of each classifier in \mathcal{C} by our method when assigning weights to classifier-class pairs, we consider the mean validation accuracy matrix \mathbf{V} (given in Table 7) as obtained for dataset

| \mathcal{E} | ρ | # Var/ Constr. | Avg. Increase (%) | | | |
|----------------|---------|-------------------|-------------------|-------------|-------------|-------------|
| | | | Bal. Acc. | Precision | Recall | F1-score |
| D1 | 6059.97 | 64/32 | 4.45 | 4.53 | 4.55 | 4.51 |
| D2 | 1222.64 | 48/28 | 3.54 | 3.55 | 3.56 | 3.59 |
| D3 | 5.91 | 40/26 | 5.42 | 5.52 | 5.50 | 5.51 |
| D4 | 2.13 | 40/26 | 4.72 | 4.92 | 4.81 | 4.83 |
| Overall | | | 4.53 | 4.63 | 4.60 | 4.61 |

Table 6: Average improvement achieved in all performance metrics by proposed MIP approach over other weighting schemes across all ensemble sizes for datasets D1–D4.

D2 and ensemble size $K = 8$. Although in the specific example classifier SVM (highlighted in gray) was not among the best-performing classifiers on average, it performed well for classes A_2 and A_4 (despite A_4 being among the minority classes in D2). Table 8 shows the weights assigned to each class for SVM by each weighting approach. It can be observed that MIP (highlighted in gray) assigned relatively high weights for classes A_2 and A_4 where SVM performed well, in contrast to all other methods, which assigned significantly lower weights. On the other hand, MIP assigned low weights to the classes where SVM did not perform well. Evidently, MIP leveraged the capabilities of SVM more effectively than all other approaches, even though SVM was among the lower-performing classifiers on average.

Computational Efficiency

In addition to the significant performance gains provided by our method, MIP demonstrated computational efficiency in calculating the weights for K classifiers from a set of n , compared to the other weighting schemes, despite the NP-hard nature of the examined problem. MIP seamlessly selects K classifiers while optimally calculating their weights in a single run. In contrast, the other examined approaches calculate the weights for each possible classifier combination individually (i.e., for a total of $\binom{n}{K}$ combinations) before eventually selecting the K best-performing classifiers. Figure 6 illustrates the speedup achieved in weight calculation by MIP over the other methods, for dataset D2, $K = \{3, 5, 7\}$, and $n = \{8, 16, 24\}$. It can be observed that the speedup provided by MIP increased significantly as n grew, especially for larger ensemble sizes, highlighting its practicality compared to the other techniques. Even for smaller K and n , MIP was considerably more efficient than the other weighting schemes. For example, for $K = 3$ and $n = 8$, MIP required 0.35 s to calculate the classifier weights, whereas the other approaches required 1 to 3.12 s.

Conclusion

We proposed an optimal, adaptable, and efficient MIP-based ensemble weighting scheme for imbalanced multi-class datasets to improve rare event detection in critical CPS. Our method optimizes classifier weights on a granular per class basis, while seamlessly and optimally selecting a predefined number of classifiers from a given set. Furthermore, it employs elastic net regularization to enhance robustness and generalization in weight assignment, while leveraging

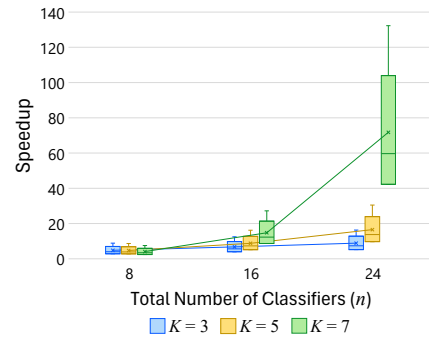


Figure 6: Speedup attained by proposed MIP approach over other schemes in calculating the weights of $K = \{3, 5, 7\}$ classifiers from a set of $n = \{8, 16, 24\}$, for dataset D2.

| C_i | E_j | | | | | Avg |
|------------|-------|-------|-------------|-------|-------------|------|
| | N_1 | A_1 | A_2 | A_3 | A_4 | |
| MLR | 0.96 | 0.92 | 0.86 | 0.99 | 0.95 | 0.94 |
| J48 | 0.89 | 0.78 | 0.85 | 0.90 | 0.90 | 0.86 |
| JRIP | 0.90 | 0.74 | 0.78 | 0.89 | 0.96 | 0.85 |
| REPTree | 0.76 | 0.86 | 0.80 | 0.98 | 0.73 | 0.83 |
| MLP | 0.90 | 0.92 | 0.81 | 0.71 | 0.79 | 0.83 |
| SVM | 0.76 | 0.73 | 0.89 | 0.76 | 0.94 | 0.82 |
| GNB | 0.90 | 0.85 | 0.81 | 0.71 | 0.73 | 0.80 |
| IBk | 0.90 | 0.72 | 0.75 | 0.74 | 0.71 | 0.76 |

Table 7: Mean validation accuracy matrix \mathbf{V} for dataset D2 and ensemble size $K = 8$.

| Weighting Scheme | E_j | | | | |
|------------------|-------|-------|-------------|-------|-------------|
| | N_1 | A_1 | A_2 | A_3 | A_4 |
| UW-PC | 0.13 | 0.13 | 0.13 | 0.13 | 0.13 |
| UW-PCC | 0.03 | 0.03 | 0.03 | 0.03 | 0.03 |
| WA-PC | 0.11 | 0.11 | 0.11 | 0.11 | 0.11 |
| WA-PCC | 0.02 | 0.02 | 0.03 | 0.02 | 0.03 |
| DE | 0.12 | 0.12 | 0.12 | 0.12 | 0.12 |
| BMA | 0.02 | 0.02 | 0.03 | 0.02 | 0.02 |
| MIP | 0.00 | 0.03 | 0.20 | 0.04 | 0.23 |

Table 8: Weights assigned to the SVM classifier (shown per class) by each weighting approach, for dataset D2 and ensemble size $K = 8$.

the strengths of each classifier. We compared our approach against six widely used weighting schemes, considering relevant datasets with multiple classes and varied imbalance ratios, under a wide range of ensemble sizes. The experimental results demonstrate the superior performance and efficiency of our technique over all existing methods, attaining an increase in balanced accuracy ranging from 0.99% to 7.31%, with an overall average of 4.53% across all datasets and ensemble sizes. Moreover, it achieved an overall average improvement of 4.63%, 4.60%, and 4.61% in macro-averaged precision, recall, and F1-score, respectively, while also exhibiting computational efficiency.

Acknowledgments

This work has been supported by the European Union's Horizon 2020 research and innovation programme under grant agreement No. 739551 (KIOS CoE) and from the Government of the Republic of Cyprus through the Cyprus Deputy Ministry of Research, Innovation and Digital Policy.

References

- Danso, P. K.; Neto, E. C. P.; Dadkhah, S.; Zohourian, A.; Molyneaux, H.; and Ghorbani, A. A. 2022. Ensemble-based intrusion detection for Internet of Things devices. In *Proceedings of the IEEE 19th International Conference on Smart Communities: Improving Quality of Life Using ICT, IoT and AI (HONET'22)*, 34–39. Marietta, GA, USA: IEEE.
- Dogan, A.; and Birant, D. 2019. A weighted majority voting ensemble approach for classification. In *Proceedings of the 4th International Conference on Computer Science and Engineering (UBMK'19)*, 1–6. Samsun, Turkey: IEEE.
- Dong, X.; Yu, Z.; Cao, W.; Shi, Y.; and Ma, Q. 2020. A survey on ensemble learning. *Frontiers of Computer Science*, 14(2): 241–258.
- Elrawy, M. F.; Hadjidemetriou, L.; Laoudias, C.; and Michael, M. K. 2023. Detecting and classifying man-in-the-middle attacks in the private area network of smart grids. *Sustainable Energy, Grids and Networks*, 36: 101167.
- Giesen, J.; Sören, L.; Andreas, L.; and Schneider, C. 2019. Using Benson's algorithm for regularization parameter tracking. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI'19)*, 3689–3696. AAAI.
- Grandini, M.; Bagli, E.; and Visani, G. 2020. Metrics for multi-class classification: An overview. arXiv:2008.05756.
- Gurobi, LLC. 2024. <https://www.gurobi.com>. Accessed: 2024-06-03.
- Hastie, T.; Tibshirani, R.; and Friedman, J. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. New York, NY: Springer, 2nd edition. ISBN 978-0-387-84858-7.
- He, H.; and Garcia, E. A. 2009. Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9): 1263–1284.
- Khan, A. A.; Chaudhari, O.; and Chandra, R. 2024. A review of ensemble learning and data augmentation models for class imbalanced problems: Combination, implementation and evaluation. *Expert Systems with Applications*, 244: 122778.
- Lango, M.; and Stefanowski, J. 2022. What makes multi-class imbalanced problems difficult? An experimental study. *Expert Systems with Applications*, 199: 116962.
- Liu, Z.; Cao, W.; Gao, Z.; Bian, J.; Chen, H.; Chang, Y.; and Liu, T.-Y. 2020. Self-paced ensemble for highly imbalanced massive data classification. In *Proceedings of the IEEE 36th International Conference on Data Engineering (ICDE'20)*, 841–852. Dallas, TX, USA: IEEE.
- Maheshwari, A.; Mehraj, B.; Khan, M. S.; and Idrisi, M. S. 2022. An optimized weighted voting based ensemble model for DDoS attack detection and mitigation in SDN environment. *Microprocessors and Microsystems*, 89: 104412.
- McTavish, H.; Zhong, C.; Achermann, R.; Karimalis, I.; Chen, J.; Rudin, C.; and Seltzer, M. 2022. Fast sparse decision tree optimization via reference ensembles. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI'22)*, 9604–9613. AAAI.
- Opitz, D.; and Maclin, R. 1999. Popular ensemble methods: An empirical study. *Journal of Artificial Intelligence Research*, 11: 169–198.
- Raftery, A. E.; Madigan, D.; and Hoeting, J. A. 1997. Bayesian model averaging for linear regression models. *Journal of the American Statistical Association*, 92(437): 179–191.
- Sagi, O.; and Rokach, L. 2018. Ensemble learning: A survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8(4): e1253.
- Sharafaldin, I.; Lashkari, A. H.; and Ghorbani, A. A. 2018. Toward generating a new intrusion detection dataset and intrusion traffic characterization. In *Proceedings of the 4th International Conference on Information Systems Security and Privacy (ICISSP'18)*, 108–116. Madeira, Portugal: SciTePress.
- Shyalika, C.; Wickramarachchi, R.; and Sheth, A. 2023. A comprehensive survey on rare event prediction. arXiv:2309.11356.
- Storn, R.; and Price, K. 1997. Differential evolution – A simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, 11(4): 341–359.
- Taheri, M.; Khorasani, K.; Shames, I.; and Meskin, N. 2024. Cyberattack and machine-induced fault detection and isolation methodologies for cyber-physical systems. *IEEE Transactions on Control Systems Technology*, 32(2): 502–517.
- Tavallaee, M.; Bagheri, E.; Lu, W.; and Ghorbani, A. A. 2009. A detailed analysis of the KDD CUP 99 data set. In *Proceedings of the 2009 IEEE Symposium on Computational Intelligence for Security and Defense Applications (CISDA'09)*, 1–6. Ottawa, ON, Canada: IEEE.
- Tertytchny, G.; Nicolaou, N.; and Michael, M. K. 2020. Classifying network abnormalities into faults and attacks in IoT-based cyber physical systems using machine learning. *Microprocessors and Microsystems*, 77: 103121.
- Vrachimis, S. G.; Kyriakou, M. S.; Eliades, D. G.; and Polycarpou, M. M. 2018. LeakDB: A benchmark dataset for leakage diagnosis in water distribution networks. In *Proceedings of the 1st International WDSA/CCWI 2018 Joint Conference*. Kingston, ON, Canada: Queen's University.
- Wang, C.; Li, X.; Sun, J.; Zhang, X.; Li, Y.; Peng, X.; Liu, J.; and Jiao, Z. 2022. Ensemble learning model of power system transient stability assessment based on Bayesian Model Averaging method. In *Proceedings of the IEEE 6th Conference on Energy Internet and Energy System Integration (EI2'22)*, 1467–1471. Chengdu, China: IEEE.
- Wang, L.; Xu, S.; Wang, X.; and Zhu, Q. 2021. Addressing class imbalance in federated learning. In *Proceedings of the*

35th AAAI Conference on Artificial Intelligence (AAAI'21), 10165–10173. AAAI.

Yin, J.; Gan, C.; Zhao, K.; Lin, X.; Quan, Z.; and Wang, Z.-J. 2020. A novel model for imbalanced data classification. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI'20)*, 6680–6687. AAAI.

Zhang, S. X.; Liu, Y. H.; Zheng, L. M.; and Zheng, S. Y. 2024. Differential evolution with collective ensemble learning. *Swarm and Evolutionary Computation*, 87: 101521.

Zhang, Y.; Zhang, H.; Cai, J.; and Yang, B. 2014. A weighted voting classifier based on differential evolution. *Abstract and Applied Analysis*, 2014(1): 376950.

Zhou, Y.; Cheng, G.; Jiang, S.; and Dai, M. 2020. Building an efficient intrusion detection system based on feature selection and ensemble classifier. *Computer Networks*, 174: 107247.

Zhou, Z.-H. 2021. Ensemble Learning. In *Machine Learning*, chapter 8, 181–210. Singapore: Springer, 1st edition.

Zou, H.; and Hastie, T. 2005. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 67(2): 301–320.