

Exponentially Weighted Instance-Aware Repeat Factor Sampling for Long-Tailed Object Detection Model Training in Unmanned Aerial Vehicles Surveillance Scenarios

Taufiq Ahmed*, Abhishek Kumar[†], Constantino Álvarez Casado*, Anlan Zhang[§],
Tuomo Hänninen*, Lauri Loven*, Miguel Bordallo López*, Sasu Tarkoma^{¶*}

*University of Oulu, Oulu, Finland

[†]University of Jyväskylä, Jyväskylä, Finland

[¶]University of Helsinki, Helsinki, Finland

[§]University of Southern California, Los Angeles, United States

Abstract—Object detection models often struggle with class imbalance, where rare categories appear significantly less frequently than common ones. Existing sampling-based rebalancing strategies, such as Repeat Factor Sampling (RFS) and Instance-Aware Repeat Factor Sampling (IRFS), mitigate this issue by adjusting sample frequencies based on image and instance counts. However, these methods are based on linear adjustments, which limit their effectiveness in long-tailed distributions. This work introduces Exponentially Weighted Instance-Aware Repeat Factor Sampling (E-IRFS), an extension of IRFS that applies exponential scaling to better differentiate between rare and frequent classes. E-IRFS adjusts sampling probabilities using an exponential function applied to the geometric mean of image and instance frequencies, ensuring a more adaptive rebalancing strategy. We evaluate E-IRFS on a dataset derived from the Fireman-UAV-RGBT Dataset and four additional public datasets, using YOLOv11 object detection models to identify fire, smoke, people and lakes in emergency scenarios. The results show that E-IRFS improves detection performance by 22% over the baseline and outperforms RFS and IRFS, particularly for rare categories. The analysis also highlights that E-IRFS has a stronger effect on lightweight models with limited capacity, as these models rely more on data sampling strategies to address class imbalance. The findings demonstrate that E-IRFS improves rare object detection in resource-constrained environments, making it a suitable solution for real-time applications such as UAV-based emergency monitoring. The code is available at: <https://github.com/futurians/E-IRFS>.

Index Terms—Object detection, class imbalance, rebalancing strategies, instance-aware sampling, long-tailed distributions, UAV-based monitoring, emergency detection, YOLOv11

I. INTRODUCTION

Effective object detection is essential in applications such as wildfire monitoring and emergency response [1]. UAV-based monitoring extends the capability of fixed-camera networks by providing flexible aerial perspectives, enhancing situational awareness in dynamic and remote environments. In wildfire scenarios, timely detection of fire, smoke, and




Fig. 1. Class distribution of the training and validation sets in the custom benchmark dataset. Bars represent the number of images and instances per class, with vertically aligned counts. The imbalance highlights the datasets challenge for object detection models.

affected individuals is crucial for an effective response. However, deploying object detection models on UAVs introduces computational constraints, requiring lightweight models that can operate efficiently on resource-limited edge devices. This highlights the need for training strategies that enable compact models to achieve competitive detection performance.

Training object detection models for such applications is challenging due to the scarcity of representative datasets and the presence of class imbalance, where critical categories (e.g., fire, people in distress) appear far less frequently than common ones (e.g., buildings, vehicles) [2], [3]. This imbalance affects model learning, leading to biased predictions toward frequent classes and reduced accuracy for rare ones. Addressing this issue is essential in applications where the detection of rare objects is critical to decision making. Sampling-based methods, such as Repeat Factor Sampling (RFS) [4], mitigate class imbalance by increasing the probability of selecting images that contain rare objects. Instance-Aware Repeat Factor Sampling (IRFS) improves upon RFS by incorporating both the image count and the instance count when determining the sampling weights [5].

In this context, this work introduces Exponentially Weighted Instance-Aware Repeat Factor Sampling (E-IRFS), a rebalancing strategy for improving rare-class detection in long-tailed object detection datasets. Unlike linear adjustments in existing methods, E-IRFS applies an exponential scaling mechanism to the geometric mean of image and instance frequencies, amplifying rare-class sampling while maintaining stable training. The evaluation is conducted on a custom dataset based on the Fireman-UAV-RGBT Dataset [1], [6] and four publicly available small datasets, covering fire, smoke, people, and lakes in emergency scenarios. Figure 1 illustrates the class distribution, emphasizing the disparity in image and instance counts. The method is compared against RFS, IRFS, and a baseline without rebalancing using YOLOv11 object detection models. The main contributions of this work are:

- The introduction of E-IRFS, an extension of IRFS that applies an exponential weighting mechanism to improve sampling-based rebalancing for rare classes.
- A detailed evaluation of E-IRFS on an emergency response dataset, showing its effectiveness in improving detection accuracy for underrepresented categories.
- An analysis of how E-IRFS benefits lightweight object detection models, where rebalancing strategies have a stronger effect due to the models limited capacity.

These contributions refine class imbalance handling in object detection. E-IRFS enhances rare-class representation without excessive oversampling, benefiting lightweight models that struggle with imbalance. This makes it well-suited for UAV-based monitoring and edge applications. Table I summarizes the symbols and notations used throughout the paper for consistency. The code for E-IRFS and the implementation details of the custom benchmark dataset are available at: <https://github.com/futurians/E-IRFS>.

TABLE I
SUMMARY OF SYMBOLS AND NOTATIONS

Symbol	Definition
c	Class index
i	Image index
f_c	Fraction of images containing class c
$f_{i,c}$	Fraction of images with at least one instance of c
$f_{b,c}$	Fraction of bounding boxes belonging to c
r_c	Repeat factor for class c
r_i	Repeat factor for image i
t	Threshold for oversampling activation
α	Scaling parameter for E-IRFS
p_i	Probability of selecting image i for training

II. RELATED WORK

Class imbalance in object detection datasets affects model performance by introducing biases toward frequent classes while reducing the ability to detect underrepresented objects. Many real-world datasets, particularly in autonomous systems, surveillance, and emergency response, follow long-tailed distributions where rare but critical categories appear infrequently [2], [3]. This imbalance is especially problematic

in safety-critical applications, where missing rare objects can impact decision-making. Studies have shown that standard training pipelines often neglect underrepresented classes due to conventional loss formulations and the dominance of frequent classes in training data [7]. The issue is more pronounced in single-stage detectors, where class distribution affects anchor assignments and training dynamics, whereas two-stage detectors demonstrate greater robustness but still require adjustments to handle skewed datasets. Approaches to address class imbalance in object detection can be categorized into data-level methods, algorithm-level modifications, and hybrid techniques. Data-level methods adjust the dataset distribution by modifying how samples are selected or generated during training. Algorithm-level modifications introduce adjustments to the model architecture or loss function to enhance learning for underrepresented classes. Hybrid methods combine both strategies, integrating data augmentation with reweighting techniques to improve representation.

A. Data Augmentation for Imbalance Mitigation

Among data-level strategies, data augmentation is commonly used to artificially increase the number of rare-class samples. Techniques such as copy-paste augmentation [8] insert rare objects into new images to enhance their representation. Similarly, Mosaic augmentation [9] merges multiple images into a single training sample to increase the variability of object occurrences. While data augmentation methods improve the diversity of training samples, they do not directly correct class distribution imbalances. Augmented samples remain subject to the same class frequencies as the original dataset, which means that rare classes may still be underrepresented during model training.

B. Sampling-Based Methods for Rebalancing

Sampling-based rebalancing adjusts the selection frequency of training samples without modifying the dataset itself. By increasing the presence of rare classes in training batches, these methods help mitigate class imbalance. Unlike loss reweighting, which alters optimization, they directly influence sample distribution. Early methods, such as class-balanced sampling, assigned equal probabilities to all classes but often led to overfitting due to excessive rare-class upsampling. More refined strategies constrain class occurrence to maintain stability. Among them, RFS [4] and IRFS [5] have shown effectiveness.

1) *Repeat Factor Sampling (RFS)*: Repeat Factor Sampling (RFS) [4] was introduced to rebalance long-tailed object detection datasets by assigning higher sampling probabilities to images containing rare classes. The method defines a repeat factor based on the number of images in which a class appears, ensuring that classes with lower frequency are oversampled during training. For a given category c , the category-level repeat factor is defined as:

$$r_c = \max \left(1, \sqrt{\frac{t}{f_c}} \right) \quad (1)$$

For each image i , the image-level repeat factor is computed as $r_i = \max_{c \in i} r_c$, where $c \in i$ represents the categories in image i . Images are repeated during training based on r_i , increasing the presence of rare-class instances. However, RFS only considers image-level frequency, ignoring instance counts, which can lead to inconsistencies when classes with similar image frequencies have different instance distributions.

2) *Instance-Aware Repeat Factor Sampling (IRFS)*: To address the previous issue, IRFS method [5] extends RFS by incorporating both image count and instance count into the sampling process. The method defines the repeat factor based on the geometric mean of these two values, allowing the sampling strategy to reflect both the presence and density of objects in the dataset. The repeat factor for a given class c is defined as:

$$r_c = \max \left(1, \sqrt{\frac{t}{\sqrt{f_{i,c} \cdot f_{b,c}}}} \right), \quad (2)$$

determining how frequently images containing class c should be sampled during training, t is a predefined threshold that controls the degree of rebalancing by determining how aggressively rare classes should be oversampled. $\sqrt{f_{i,c} \cdot f_{b,c}}$ is the geometric mean of the image frequency and instance frequency, ensuring a balanced measure of class presence and density.

III. PROPOSED METHODOLOGY

Existing sampling-based rebalancing techniques, such as RFS and IRFS, have improved rare-class representation in long-tailed object detection. However, these methods rely on linear weighting functions, which do not sufficiently amplify the presence of extremely rare classes. IRFS improves upon RFS by incorporating both image count and instance count, but its geometric mean formulation still applies only a limited adjustment for highly underrepresented categories. To address this limitation, we introduce Exponentially Weighted Instance-Aware Repeat Factor Sampling (E-IRFS), which applies an exponential weighting function to the geometric mean of image and instance frequencies. This adjustment ensures that rare classes receive a stronger sampling adjustment while preventing excessive overrepresentation of frequent classes. The exponential function enhances differentiation between rare and frequent categories, making the sampling process more adaptive to highly skewed distributions.

A. Exponentially Weighted Instance-Aware Repeat Factor Sampling (E-IRFS)

To further refine instance-aware sampling, we propose Exponentially Weighted Instance-Aware Repeat Factor Sampling (E-IRFS). This method amplifies the oversampling strength for rare classes by applying exponential scaling to the repeat factor. The new class-level repeat factor is formulated as:

$$r_c = \exp \left(\alpha \cdot \sqrt{\frac{t}{\sqrt{f_{i,c} \cdot f_{b,c}}}} \right), \quad (3)$$

where α is a scaling parameter that controls the aggressiveness of sampling adjustments, $\exp(\cdot)$ ensures that small values of $f_{i,c} \cdot f_{b,c}$ result in exponentially larger repeat factors, and the square root mean $\sqrt{f_{i,c} \cdot f_{b,c}}$ balances both image and instance counts.

B. Image-Level Sampling Using E-IRFS

Once the category-level repeat factors r_c are computed, the image-level repeat factor r_i is given by $r_i = \max_{c \in i} r_c$, where r_i determines how often an image is repeated during training. Higher r_i values indicate more frequent sampling of images containing rare objects. The final sampling probability of selecting each image for training is given by equation 4, ensuring that rare classes receive higher sampling priority to counteract class imbalance:

$$p_i = \frac{r_i}{\sum_j r_j}, \quad (4)$$

where $\sum_j r_j$ computes the total sum of all image-level repeat factors across the entire dataset and p_i represents the normalized probability of selecting image i for training. This ensures that the sampling probabilities sum to 1, maintaining training stability while emphasizing rare categories effectively.

C. Theoretical Analysis

This section examines the properties of E-IRFS, including its scaling behavior, impact on class distribution, and effect on training convergence.

1) *Scaling Behavior and Growth Rate*: The effectiveness of sampling-based rebalancing depends on how it adjusts the sampling probability for rare classes. RFS and IRFS follow sublinear growth patterns, providing limited reinforcement for underrepresented categories. E-IRFS introduces an exponential scaling mechanism that increases the sampling probability more aggressively for rare classes while maintaining stability for frequent ones. The repeat factor r_c determines how strongly rare classes are oversampled. In RFS, it follows an inverse square root relationship with image frequency $f_{i,c}$, while IRFS incorporates instance frequency $f_{b,c}$ using a geometric mean:

- a) RFS Growth Rate: $r_c \propto f_{i,c}^{-1/2}$
- b) IRFS Growth Rate: $r_c \propto (f_{i,c} \cdot f_{b,c})^{-1/4}$
- c) E-IRFS Growth Rate: $r_c \propto \exp(\alpha \cdot (f_{i,c} \cdot f_{b,c})^{-1/4})$

E-IRFS extends IRFS by applying an exponential function to the geometric mean, increasing the contrast between rare and frequent categories. Since the exponential function grows faster than polynomial functions, this approach increases model exposure to underrepresented categories while maintaining stable sampling for moderately frequent ones.

2) *Sampling Distribution and Training Balance*: E-IRFS amplifies the sampling probability of images containing rare objects, redistributing the probability mass to improve model exposure to underrepresented categories. The probability of selecting an image follows equation 4, where the exponential weighting enhances the inclusion of rare-class instances while preserving representation of frequent categories. The resulting class distribution in training is given by $P_{\text{train}}(c) \propto r_c \cdot P_{\text{data}}(c)$, where $P_{\text{data}}(c)$ is the original dataset distribution. For long-tailed datasets, $P_{\text{data}}(c)$ follows a power-law distribution [10]: $P_{\text{data}}(c) \approx c_0 f_{i,c}^{-\gamma}$ (for some constant c_0 and exponent γ). For different sampling strategies, the training distribution is as follows:

- a) RFS: $P_{\text{train}}(c) \propto f_{i,c}^{-(\gamma+1/2)}$
- b) IRFS: $P_{\text{train}}(c) \propto (f_{i,c} \cdot f_{b,c})^{-(\gamma+1/4)}$
- c) E-IRFS: $P_{\text{train}}(c) \propto \exp(\alpha \cdot (f_{i,c} \cdot f_{b,c})^{-1/4}) f_{i,c}^{-\gamma}$

Since the exponential function in E-IRFS grows faster than polynomial adjustments, it significantly enhances the presence of rare classes in training. This ensures that underrepresented categories receive stronger sampling reinforcement while maintaining balance across all classes.

3) *Computational Complexity*: E-IRFS introduces an additional exponential computation and a geometric mean calculation per class. However, these operations remain in $\mathcal{O}(N)$ per epoch, where N is the number of training images, similar to RFS and IRFS. Since the sampling process itself remains unchanged, the added computational overhead is minimal.

4) *Effect on Convergence and Sampling Stability*: In long-tailed object detection, rare classes contribute weaker gradient updates, making them harder to learn. E-IRFS mitigates this by increasing their representation, balancing gradient updates, and stabilizing training. This is especially beneficial for lightweight models with limited capacity to capture long-tailed distributions. The stability of this adjustment is supported by the mathematical properties of r_c , with its first derivative given by:

$$\frac{d}{df_{i,c}} r_c = \alpha \cdot \exp\left(\alpha \cdot \sqrt{\frac{t}{\sqrt{f_{i,c} \cdot f_{b,c}}}}\right) \cdot \left(-\frac{t}{2\sqrt{f_{b,c} f_{i,c}^3}}\right) \quad (5)$$

Since the exponential function is always positive and the second term is negative, r_c is monotonically decreasing. This ensures that as the image frequency $f_{i,c}$ decreases, the sampling rate increases, reinforcing rare-class representation in training. Furthermore, the second derivative is as follows:

$$\frac{d^2}{df_{i,c}^2} r_c = \alpha \cdot \exp\left(\alpha \cdot \sqrt{\frac{t}{\sqrt{f_{i,c} \cdot f_{b,c}}}}\right) \cdot \left(\frac{3t}{4\sqrt{f_{b,c} f_{i,c}^5}}\right) \quad (6)$$

Since this expression is always positive, r_c is convex. This ensures a smooth transition between frequent and rare categories, preventing abrupt sampling shifts and maintaining stable training dynamics while improving detection of underrepresented classes.

D. Benchmark Database

We created a custom benchmark for object detection by merging five different datasets ([1], [11], [12], [13], [14]). To construct the training and validation sets of our custom dataset, we used the original training and validation splits from four of these datasets [11]–[14] without applying any selection or filtering. From the dataset in [1], we selectively used only the Unimodal-RGB subset from the Multimodal collection, excluding the thermal modality. Furthermore, we included only a curated set of video frames that contained at least one of the following objects: fire, smoke, humans, or lakes. The resulting custom dataset comprises 40,384 training images and 11,953 validation images. The validation set was formed by concatenating the official validation splits from the five merged datasets rather than selecting a separate subset. This setup ensures a diverse, yet intentionally unbalanced distribution for evaluation purposes. The class imbalance provides a challenging scenario for evaluating fire, smoke, human presence, and water body detection, fostering robust model training under diverse conditions. We also evaluate E-IRFS on two well-known balanced datasets, CIFAR-10 [15] and Caltech-101 [16], to assess its behavior when class distributions are uniform. This allows us to verify whether the exponential reweighting introduces performance degradation outside long-tailed scenarios.

TABLE II
CLASS DISTRIBUTION IN TRAINING AND VALIDATION SETS IN OUR CUSTOM DATASET, ILLUSTRATING CLASS IMBALANCE.

Class	Training Set			Validation Set		
	Images	Instances	Percentage	Images	Instances	Percentage
Fire	16,915	33,773	23.0%	1,436	2,336	8.5%
Smoke	28,769	32,538	22.1%	6,735	7,090	25.8%
Human	18,525	67,992	46.3%	4,804	16,612	60.5%
Lake	12,646	12,646	8.6%	1,087	1,426	5.2%
Total	40,384	146,949	100%	11,953	27,464	100%

E. Evaluation Metrics and Protocols

The effectiveness of the YOLOv11 object detection models was evaluated using standard metrics provided by the Ultralytics validation system. These metrics assess the accuracy and reliability of the model in identifying objects across different conditions. To quantify detection performance, we used mean Average Precision at an Intersection over Union (IoU) threshold of 0.5 (mAP_{50}) and mean Average Precision across multiple IoU thresholds ranging from 0.5 to 0.95 in steps of 0.05 (mAP_{50-95}). The mAP_{50} metric evaluates detection performance by computing the precision-recall curve at a fixed IoU threshold of 0.5, allowing for some localization tolerance. In contrast, mAP_{50-95} provides a more strict assessment by averaging precision across multiple IoU thresholds, thereby offering a more comprehensive measure of detection accuracy and localization robustness. These metrics facilitate a standardized comparison of model performance across different object scales and levels of complexity within the dataset, ensuring a rigorous evaluation of the proposed

approach. For classification tasks on CIFAR-10 and Caltech-101, we report Top-1 and Top-5 accuracy, which measure whether the correct class is the model first or among its five highest-ranked predictions. These metrics are standard for evaluating classification without localization.

IV. EXPERIMENTAL ANALYSIS AND RESULTS

A. Experimental Setup

All experiments used the Ultralytics YOLOv11 framework [17]. The YOLOv11-Nano model (2.6M parameters) and YOLOv11-Large (25.3M parameters) were trained on the custom dataset using an NVIDIA A30 GPU (24GB VRAM) to compare the impact of model size on performance. Training was performed using an image resolution of 640 640 pixels, batch size 16, and 100 epochs. The optimizer was auto-selected with momentum 0.937, weight decay 0.0005, and an initial learning rate of 0.01, gradually reduced to 0.001. Data augmentation was applied in all sampling-based experiments unless explicitly omitted, including Mosaic (0.9), Mixup (0.2), Copy-Paste (0.3), AutoAugment, and geometric transformations. Model performance was assessed using mAP_{50} and mAP_{50-95} metrics at different threshold values and values of the scaling parameter. The evaluation included comparisons between RFS, IRFS, and the proposed E-IRFS method, as well as baseline experiments with default hyperparameters and augmentation configurations.

B. Performance Comparison of Re-Balancing Methods

This experiment compares the RFS, IRFS, and E-IRFS with a baseline in the YOLOv11-Nano setup to evaluate their impact on rare-class detection and overall model performance with and without data augmentation. During training, we set $\alpha = 2.0$ for E-IRFS, and threshold $t = 0.0001$ for all methods. The results presented in Table III, firstly, highlight the performance of re-balancing methods without the influence of data augmentation. Both RFS and IRFS improve mAP_{50} by 8% compared to the baseline but do not affect the mAP_{50-95} metric. In contrast, E-IRFS achieves a 17% increase in mAP_{50} and a 13% improvement in mAP_{50-95} , demonstrating its stronger ability to improve detection accuracy across different object sizes. The significant gain in mAP_{50-95} suggests that E-IRFS not only enhances class balance but also contributes to better localization precision for rare objects.

TABLE III

COMPARISON OF mAP_{50} AND mAP_{50-95} IN THE YOLOV11-NANO SETUP, TRAINED WITH AND WITHOUT DATA AUGMENTATION.

Method	No Data Augmentation		Data Augmentation	
	mAP_{50}	mAP_{50-95}	mAP_{50}	mAP_{50-95}
Baseline	0.45	0.22	0.49 (+8%)	0.23 (+4%)
RFS	0.49 (+8%)	0.22 (+0%)	0.50 (+11%)	0.24 (+9%)
IRFS	0.49 (+8%)	0.22 (+0%)	0.50 (+11%)	0.24 (+9%)
E-IRFS	0.53 (+17%)	0.25 (+13%)	0.55 (+22%)	0.25 (+13%)

To further assess the effect of re-balancing with data augmentation, Table III presents results with augmentation

applied across all methods, including the baseline. While RFS and IRFS improve mAP_{50} by 11%, their gains remain marginal compared to training without augmentation. In contrast, E-IRFS achieves a 22% increase in mAP_{50} and a 13% improvement in mAP_{50-95} , further demonstrating its advantage in addressing class imbalance, as shown in Figure 2. These results suggest that while augmentation benefits all methods, E-IRFS benefits the most due to its exponential weighting, which amplifies the effect of balanced sample exposure. The performance differences highlight a key limitation of RFS and IRFS, where their linear weighting does not sufficiently enhance rare-class representation.

C. Threshold and Scaling Factor Analysis in E-IRFS

This experiment analyzes the impact of the threshold parameter and the scaling factor α on the E-IRFS method, evaluating their effect on model performance. The formulation in Equation 4 demonstrates that α controls the aggressiveness of exponential scaling, while the threshold t determines the starting point for rebalancing adjustments. By systematically varying these parameters, we assess their influence on model accuracy. The results in Table IV and Table V show that increasing α improves the emphasis on rare categories, but only when combined with a sufficiently low threshold.

TABLE IV
 mAP_{50} OF E-IRFS ON YOLOV11-NANO WITH VARYING α AND THRESHOLD VALUES

$\alpha \setminus$ Threshold	0.1	0.01	0.001	0.0001
0.5	0.53 (+17%)	0.49 (+8%)	0.53 (+17%)	0.48 (+6%)
1.0	0.52 (+15%)	0.49 (+8%)	0.52 (+15%)	0.49 (+8%)
2.0	0.48 (+7%)	0.50 (+11%)	0.47 (+4%)	0.55 (+22%)

For example, for the small model YOLOv11-Nano (Table IV), setting $\alpha = 2.0$ with $t = 0.0001$ produces a 22% improvement in mAP_{50} , outperforming lower α values where the effect of exponential scaling is weaker. A similar trend is observed in YOLOv11-Large, where the same configuration leads to an increase 18%. Lower values of α , such as 0.5 or 1.0, result in more conservative adjustments, which reduce the ability to compensate for extreme class imbalances. The improvement at $\alpha = 2.0$ and a low threshold is due to the exponential function in Equation 4. When the geometric mean is small, indicating a rare class, exponentiation amplifies the repeat factor more significantly, ensuring stronger sampling weights and better representation in training.

TABLE V
 mAP_{50} OF E-IRFS ON YOLOV11-LARGE WITH VARYING α AND THRESHOLD VALUES

$\alpha \setminus$ Threshold	0.1	0.01	0.001	0.0001
0.5	0.52 (+4%)	0.56 (+12%)	0.57 (+14%)	0.54 (+8%)
1.0	0.55 (+10%)	0.58 (+16%)	0.53 (+6%)	0.57 (+14%)
2.0	0.55 (+10%)	0.54 (+8%)	0.57 (+14%)	0.59 (+18%)

To evaluate the effect of threshold selection in E-IRFS against RFS and IRFS, we fixed $\alpha = 2.0$ and varied the

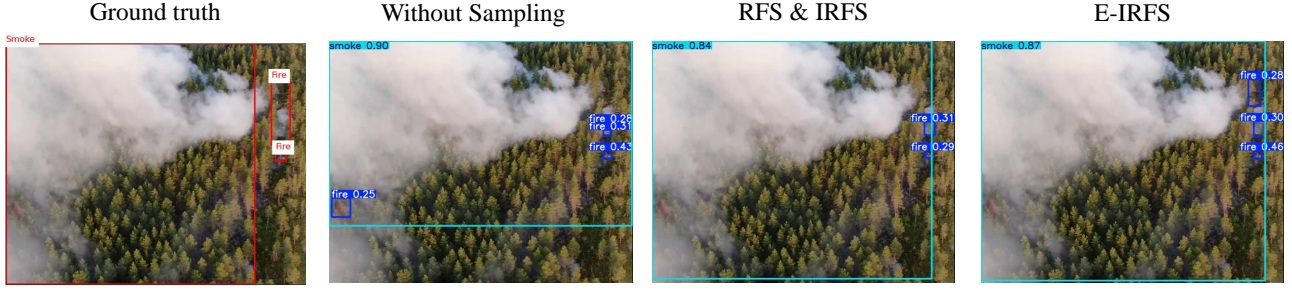


Fig. 2. Comparison of rare category detection performance on our custom testset for YOLOv11-Nano model.

threshold across different values, comparing the results with traditional sampling strategies. The results in Table VI show that RFS and IRFS consistently achieve an 11% increase in mAP_{50} regardless of the threshold, indicating that their sampling adjustments remain unaffected by changes in this parameter.

In contrast, E-IRFS is highly sensitive to threshold selection, with performance improving as the threshold decreases. This sensitivity arises from its exponential weighting mechanism, which amplifies rare-class sampling more dynamically than the linear adjustments in RFS and IRFS. Lower thresholds enhance differentiation between class frequencies, improving rare-class representation without excessive oversampling. As shown in Table IV, E-IRFS achieves its highest mAP_{50} at $t = 0.0001$, where the exponential function optimally balances sampling rates. At higher thresholds, sampling becomes more uniform, reducing rebalancing benefits and causing performance fluctuations. This highlights the importance of carefully selecting α and t for long-tailed object detection.

TABLE VI
COMPARISON OF mAP_{50} FOR RFS, IRFS, AND E-IRFS WITH $\alpha = 2.0$ AT DIFFERENT THRESHOLD VALUES ON YOLOv11-NANO.

Threshold	RFS	IRFS	E-IRFS ($\alpha = 2.0$)
0.1	0.50 (+11%)	0.50 (+11%)	0.48 (+7%)
0.01	0.50 (+11%)	0.50 (+11%)	0.50 (+11%)
0.001	0.50 (+11%)	0.50 (+11%)	0.47 (+4%)
0.0001	0.50 (+11%)	0.50 (+11%)	0.55 (+22%)

D. Impact of Re-Balancing Methods on Class-Wise Performance

The evaluation of re-balancing methods at the class level provides a more precise insights into their effectiveness in improving detection across underrepresented categories. The results in Tables VII and VIII indicate that E-IRFS achieves the most significant improvements, particularly for minority classes such as Lake and Fire.

For the Lake category, mAP_{50} increased from 0.02 to 0.09, corresponding to a relative improvement of 350%, while $mAP_{50} - 95$ increased from 0.009 to 0.032, reflecting a gain of 255%. Similarly, the Fire class showed an increase

TABLE VII
CLASS-WISE mAP_{50} PERFORMANCE IN YOLOv11-NANO SETUP

Class	Baseline	Aug	Aug+RFS	Aug+IRFS	Aug+E-IRFS
Fire	0.32	0.43(+34%)	0.44(+37%)	0.44(+37%)	0.53(65%)
Smoke	0.80	0.81(+1%)	0.82(+2%)	0.82(+2%)	0.82(+2%)
Human	0.67	0.65(-3%)	0.67(+0%)	0.67(+0%)	0.78(+16%)
Lake	0.02	0.05(+150%)	0.06(+200%)	0.06(+200%)	0.09(+350%)

of 65% in mAP_{50} and 31% in $mAP_{50} - 95$. The Human class benefited from an increase 16% in mAP_{50} and an 18% increase in $mAP_{50} - 95$, indicating balanced improvements. In contrast, the Smoke class exhibited only marginal improvements, suggesting that it was already well-represented in the dataset.

TABLE VIII
CLASS-WISE $mAP(50-95)$ PERFORMANCE IN YOLOv11-NANO SETUP

Class	Baseline	Aug	Aug+RFS	Aug+IRFS	Aug+E-IRFS
Fire	0.165	0.191(+15%)	0.213(+29%)	0.213(+29%)	0.217(31%)
Smoke	0.387	0.465(+20%)	0.456(+17%)	0.456(+17%)	0.404(+4%)
Human	0.299	0.282(-5%)	0.310(+3%)	0.310(+3%)	0.353(+18%)
Lake	0.009	0.012(+33%)	0.016(+77%)	0.016(+77%)	0.032(+255%)

These results confirm that E-IRFS effectively enhances model performance on underrepresented classes, making it more adaptable to long-tailed distributions. Unlike RFS and IRFS, which provide uniform improvements across categories, E-IRFS dynamically adjusts sampling weights, leading to a more pronounced effect on classes with fewer instances. This supports its suitability for applications where rare object detection is critical.

E. Generalization to Balanced Datasets

To evaluate the applicability of E-IRFS beyond long-tailed scenarios, we tested it on two balanced datasets: CIFAR-10 [15] and Caltech-101 [16]. These experiments aim to determine whether the exponential weighting mechanism introduces performance degradation or bias when class distributions are approximately uniform. We trained the YOLOv11-Nano classification model on both datasets using baseline and E-IRFS configurations with $\alpha = 0.5$, $\alpha = 1.0$, and $\alpha = 2.0$, fixing the threshold at $t = 0.1$.

TABLE IX
ACCURACY ON BALANCED DATASETS (TOP-1 / TOP-5)

Dataset	Method	Top-1 Acc	Top-5 Acc
CIFAR-10	Baseline	0.802	0.989
	E-IRFS ($\alpha = 0.5, t = 0.1$)	0.802	0.989
	E-IRFS ($\alpha = 1.0, t = 0.1$)	0.802	0.989
	E-IRFS ($\alpha = 2.0, t = 0.1$)	0.802	0.989
Caltech-101	Baseline	0.920	0.986
	E-IRFS ($\alpha = 0.5, t = 0.1$)	0.927	0.986
	E-IRFS ($\alpha = 1.0, t = 0.1$)	0.924	0.985
	E-IRFS ($\alpha = 2.0, t = 0.1$)	0.920	0.984

As shown in Table IX, E-IRFS maintains performance parity with the baseline across all settings on CIFAR-10 and Caltech-101. Minor fluctuations in Caltech-101 are within acceptable bounds and do not suggest systematic degradation. These results confirm that E-IRFS does not introduce bias or degrade performance on balanced datasets, supporting its robustness and generalization beyond imbalanced scenarios.

V. CONCLUSION

This paper introduces E-IRFS, a novel rebalancing strategy designed to improve rare-class detection in long-tailed object detection. The experimental results demonstrate that E-IRFS outperforms existing methods in detecting rare objects such as fires and water bodies in UAV-based monitoring. E-IRFS is especially beneficial for lightweight models with fewer parameters, such as YOLOv11-Nano, since they lack the capacity to learn long-tailed distributions effectively, making them more dependent on data sampling strategies. E-IRFS benefits from a lower threshold and a higher scaling factor, allowing more adaptive sampling adjustments. Beyond its empirical improvements, it provides a structured approach to addressing class imbalance, with differentiation between frequent and rare categories, while highlighting the importance of adaptive sampling strategies. E-IRFS shows sensitivity to selection of α and t , but our results indicate that effective values can generalize across datasets with minimal tuning. Future work may explore the extension of E-IRFS to other recognition tasks and the integration of it with additional augmentation techniques to improve its performance in highly imbalanced datasets. Furthermore, a critical direction involves investigating methodologies to determine the optimal α and t parameters based on statistical priors from new datasets. This would reduce the need for extensive manual tuning, thus enhancing the practicality and applicability of the method across diverse domains.

ACKNOWLEDGMENT

This research was supported by the Research Council of Finland 6G Flagship program (Grant Number 369116), FIREMAN project (Grant Number 348008), and Business Finland Neural Publish Subscribe for 6G Project (diary number 8754/31/2022).

REFERENCES

- [1] S. Kularatne, C. Álvarez Casado, J. Rajala, T. Hnninen, M. Bordallo Lpez, and L. Nguyen, "Fireman-uav-rgbt: A novel uav-based rgb-thermal video dataset for the detection of wildfires in the finnish forests," in *2024 IEEE 29th International Conference on Emerging Technologies and Factory Automation (ETFA)*, 2024, pp. 1–8.
- [2] N. Elasal, D. M. Swart, and N. Miller, "Frame augmentation for imbalanced object detection datasets," *Journal of Computational Vision and Imaging Systems*, vol. 4, no. 1, pp. 3–3, 2018.
- [3] J. Farooq, S. Fatima, and N. Aafaq, "Synthetic Randomized Image Augmentation (SRIA) to Address Class Imbalance Problem," in *2023 3rd International Conference on Computing and Information Technology (ICCIIT)*. IEEE, sep 13 2023, pp. 308–313.
- [4] A. Gupta, P. Dollar, and R. Girshick, "Lvis: A dataset for large vocabulary instance segmentation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, p. 5356.
- [5] B. Yaman, T. Mahmud, and C.-H. Liu, "Instance-aware repeat factor sampling for long-tailed object detection," *arXiv preprint arXiv:2305.08069*, 2023.
- [6] W. Kularatne, K. Sedaghat Shayegan, C. Álvarez Casado, J. Rajala, T. Hnninen, M. Bordallo Lpez, and N. Le Nguyen, "Fireman-uav-rgbt," September 2024. [Online]. Available: <https://doi.org/10.5281/zenodo.13732947>
- [7] H. Gu, H. Dong, N. Konz, and M. A. Mazurowski, "A systematic study of the foreground-background imbalance problem in deep learning for object detection," *arXiv.org*, 2023.
- [8] G. Ghiasi, Y. Cui, A. Srinivas, R. Qian, T.-Y. Lin, E. D. Cubuk, Q. V. Le, and B. Zoph, "Simple copy-paste is a strong data augmentation method for instance segmentation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 2918–2928.
- [9] C. Zhang, T.-Y. Pan, Y. Li, H. Hu, D. Xuan, S. Changpinyo, B. Gong, and W.-L. Chao, "Mosaicos: a simple and effective use of object-centric images for long-tailed object detection," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 417–427.
- [10] A. Clauset, C. R. Shalizi, and M. E. Newman, "Power-law distributions in empirical data," *SIAM review*, vol. 51, no. 4, pp. 661–703, 2009.
- [11] "forest fire and smoke dataset," <https://universe.roboflow.com/master-candidate/forest-fire-and-smoke-rf4pd>, jan 2024, visited on 2025-02-15.
- [12] Saus, "Saus dataset," <https://universe.roboflow.com/saus/saus>, jun 2024, visited on 2025-02-15.
- [13] U. of Engineering and Technology, "Sar custom drone dataset," https://universe.roboflow.com/university-of-engineering-and-technology-huotg/sar_custom_drone, dec 2023, visited on 2025-02-15.
- [14] Diplomatic, "Merged thesis dataset," <https://universe.roboflow.com/diplomatic/merged-thesis>, jan 2023, visited on 2025-02-15.
- [15] A. Krizhevsky, G. Hinton, *et al.*, "Learning multiple layers of features from tiny images," 2009.
- [16] L. Fei-Fei, R. Fergus, and P. Perona, "One-shot learning of object categories," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 4, pp. 594–611, 2006.
- [17] G. Jocher and J. Qiu, "Ultralytics yolo11," 2024. [Online]. Available: <https://github.com/ultralytics/ultralytics>