



A comprehensive survey on imbalanced data learning

Xinyi GAO¹, Dongting XIE², Yihang ZHANG³, Zhengren WANG⁴, Chong CHEN⁵, Conghui HE⁶, Hongzhi YIN¹✉, Wentao ZHANG⁷✉

1. School of Electrical Engineering and Computer Science, The University of Queensland, Brisbane QLD 4072, Australia
2. School of Information Science and Engineering, Yunnan University, Kunming 650091, China
3. School of Mathematical Sciences, Peking University, Beijing 100871, China
4. Academy for Advanced Interdisciplinary Studies, Peking University, Beijing 100871, China
5. Huawei Cloud BU, Beijing 100004, China
6. Shanghai AI Laboratory, Shanghai 200232, China
7. Center of Machine Learning Research, Peking University, Beijing 100871, China

Received March 14, 2025; accepted September 8, 2025

E-mail: h.yin1@uq.edu.au; wentao.zhang@pku.edu.cn

© The Author(s) 2025. This article is published with open access at link.springer.com and journal.hep.com.cn

Abstract

With the expansion of data availability, machine learning (ML) has achieved remarkable breakthroughs in both academia and industry. However, imbalanced data distributions are prevalent in various types of raw data and severely hinder the performance of ML by biasing the decision-making processes. To deepen the understanding of imbalanced data and facilitate the related research and applications, this survey systematically analyzes various real-world data formats and concludes existing researches for different data formats into four distinct categories: data re-balancing, feature representation, training strategy, and ensemble learning. This structured analysis helps researchers comprehensively understand the pervasive nature of imbalance across diverse data formats, thereby paving a clearer path toward achieving specific research goals. We provide an overview of relevant open-source libraries, spotlight current challenges, and offer novel insights aimed at fostering future advancements in this critical area of study.

Keywords

machine learning; imbalanced data learning; deep learning

1 Introduction

The advent of expansive data availability has propelled machine learning (ML) to the forefront of technological advancements in both academia and industry. These ML models are intricately designed to fit specific data distributions and deployed across a wide range of downstream tasks, ranging from predictive analytics to automated decision-making systems. Consequently, their performance is heavily influenced by the quality and distribution of the data used during training. When the data is representative, diverse, and carefully preprocessed, models achieve not only high accuracy but also robustness and strong generalization across different settings and challenges.

However, natural data distributions are inherently complex and frequently flawed. Among the myriad challenges they present, the issue of imbalanced data distribution is particularly prominent, reflecting the widespread and natural disparities encountered across various fields. In financial fraud detection, for example, fraudulent transactions occur far less frequently than legitimate ones, due to

long-standing financial regulations and continuous improvements in system integrity. The Credit Card Fraud Detection system exemplifies this, with fraudulent cases comprising only 0.17% of the total. In healthcare, imbalance arises from the natural prevalence of diseases in the population. Common conditions dominate medical records, while rare but critical diseases are significantly underrepresented. For instance, the NIH Chest X-ray14 dataset contains over 100,000 images, yet conditions such as Pneumothorax appear in less than 1% of cases. In industrial manufacturing, quality control systems are designed to minimize defects through efficient processes and stringent standards. As a result, defects are rare by design, leading to imbalanced datasets. The NEU Surface Defect Dataset, for instance, includes only 300 images per defect type, while real-world systems routinely process thousands of non-defective samples. Beyond these general examples, imbalance is also pervasive in broader scientific and engineering applications. In energy management, faults in power grids or renewable systems are rare compared to normal operations. In emergency management, rare but

high-impact events such as earthquakes, floods, or industrial accidents are naturally underrepresented. In environmental management, events like oil spills, algal blooms, or pollutant spikes occur sparsely compared to routine monitoring records. In agriculture and horticulture, healthy crops dominate datasets, while rare diseases, pest outbreaks, or nutrient deficiencies are significantly under-sampled. These imbalances are not mere data collection artifacts but stem from structural characteristics of the respective domains. Consequently, these scenarios not only complicate ML model training but also place greater demands on the robustness and reliability of the systems.

Generally, imbalanced data distributions significantly impact the performance and utility of ML models. In imbalanced settings, majority classes refer to categories with relatively larger numbers of samples, while minority classes denote categories with comparatively fewer samples. Models often achieve strong performance on majority classes but struggle on minority classes, leading to unclear distribution boundaries. Consequently, although ML models may exhibit satisfactory overall performance, their effectiveness is notably reduced within these minority classes. However, these minority classes often hold the greater significance in real-world scenarios. For example, in medical diagnostics, the failure to detect a rare disease due to insufficient data can lead to missed diagnoses and inadequate patient care. Similarly, in financial systems, the inability to identify infrequent instances of fraud can result in significant financial losses and compromised security. This tendency of ML models to overlook rare but critical instances diminishes the utility and safety of automated decision-making systems in practical applications.

In response to these challenges, the ML community has devised a range of methods and we organize them into four fundamental categories—data re-balancing, feature representation, training strategy, and ensemble learning—each corresponding to crucial aspects of the ML process. Data re-balancing techniques are crucial for adjusting the data distribution to better representation, utilizing approaches such as over-sampling the minority class and under-sampling the majority class. This adjustment is critical to prevent the model from favoring the majority samples, aligning with the data preparation phase of ML. Feature representation strategies enhance the ability to accurately capture and represent the information relevant to minority samples. This improvement is pivotal in the feature engineering phase, enabling models to effectively learn from and make predictions for all samples. Advanced training strategies adjust the learning algorithms to minimize their inherent bias towards the majority samples. This key adjustment during the training phase ensures that the learning process is inclusive, giving equitable consideration to all samples. Lastly, ensemble methods, which involve combining several models, correspond to the model integration part of the ML process. These methods leverage the strengths of multiple algorithms to potentially reduce biases that may arise from imbalanced data, thereby enhancing the robustness and accuracy of the final model outputs. By categorizing methods according to the foundational processes of ML, this taxonomy not only facilitates comprehensive field surveys but also clarifies the

motivations behind these strategies, facilitating the pursuit of specific objectives for specific method designation. Additionally, this survey explores the manifestation of imbalances across various data formats, including images, text, and graphs, highlighting the differences, unique challenges and required adaptations for each format. This exploration is critically important as it deepens the understanding of each data format and facilitates the development of targeted ML strategies for complex data format scenarios.

While some existing surveys on imbalanced learning have offered valuable insights, they remain limited in scope. Traditional reviews [1–3] primarily focus on classical approaches such as sampling, cost-sensitive learning, and kernel-based methods, while overlooking recent advances in deep learning. Furthermore, the survey [3] highlights open challenges and specific application scenarios across domains ranging from management science to engineering, but lacks a systematic ML methodological overview. Other surveys, such as those on deep long-tailed learning [4], are restricted to computer vision and therefore provide limited coverage of other data modalities. In contrast, our survey presents a unified taxonomy that integrates both classical and deep learning methods, with particular emphasis on feature representation and training strategies. We also discuss the data-specific challenges across multiple formats, including image, text, and graph data, and provide comprehensive coverage of evaluation metrics and open-source libraries, making this work a timely and practical resource for researchers and practitioners.

The contributions of this survey are summarized as follows:

- We provide a comprehensive survey of the literature addressing imbalanced data learning, offering a structured overview of methods that are rooted in the foundational processes of ML.
- We conduct a thorough analysis of imbalances across various data formats, including images, text, and graphs, providing an in-depth exploration of the challenges and methodologies specific to each format.
- We highlight resources available for tackling imbalanced data and explore current challenges and future directions. This discussion is designed to guide researchers struggling with imbalance issues, aiding them in developing strategies effectively and efficiently.

The structure of this survey is outlined as follows: Section 2 provides a thorough investigation of methodologies for handling imbalances, organized according to our taxonomy. Section 3 extensively discusses the manifestation of imbalances across various data formats. Section 4 offers a detailed investigation of evaluation metrics specific to imbalanced data methods. Section 5 is dedicated to summarizing the resources available for learning from imbalanced data. Finally, Section 6 presents the challenges and future directions in this field.

■ 2 Methodology

In this section, we delve into the specifics of methods for handling imbalanced data, conducting a systematic analysis across four

distinct categories: data re-balancing, feature representation, training strategy, and ensemble learning. We explore the various techniques employed within each category to achieve their respective objectives.

Additionally, we provide a detailed overview in Fig. 1, which serves to encapsulate the scope and relationships among the methods discussed.

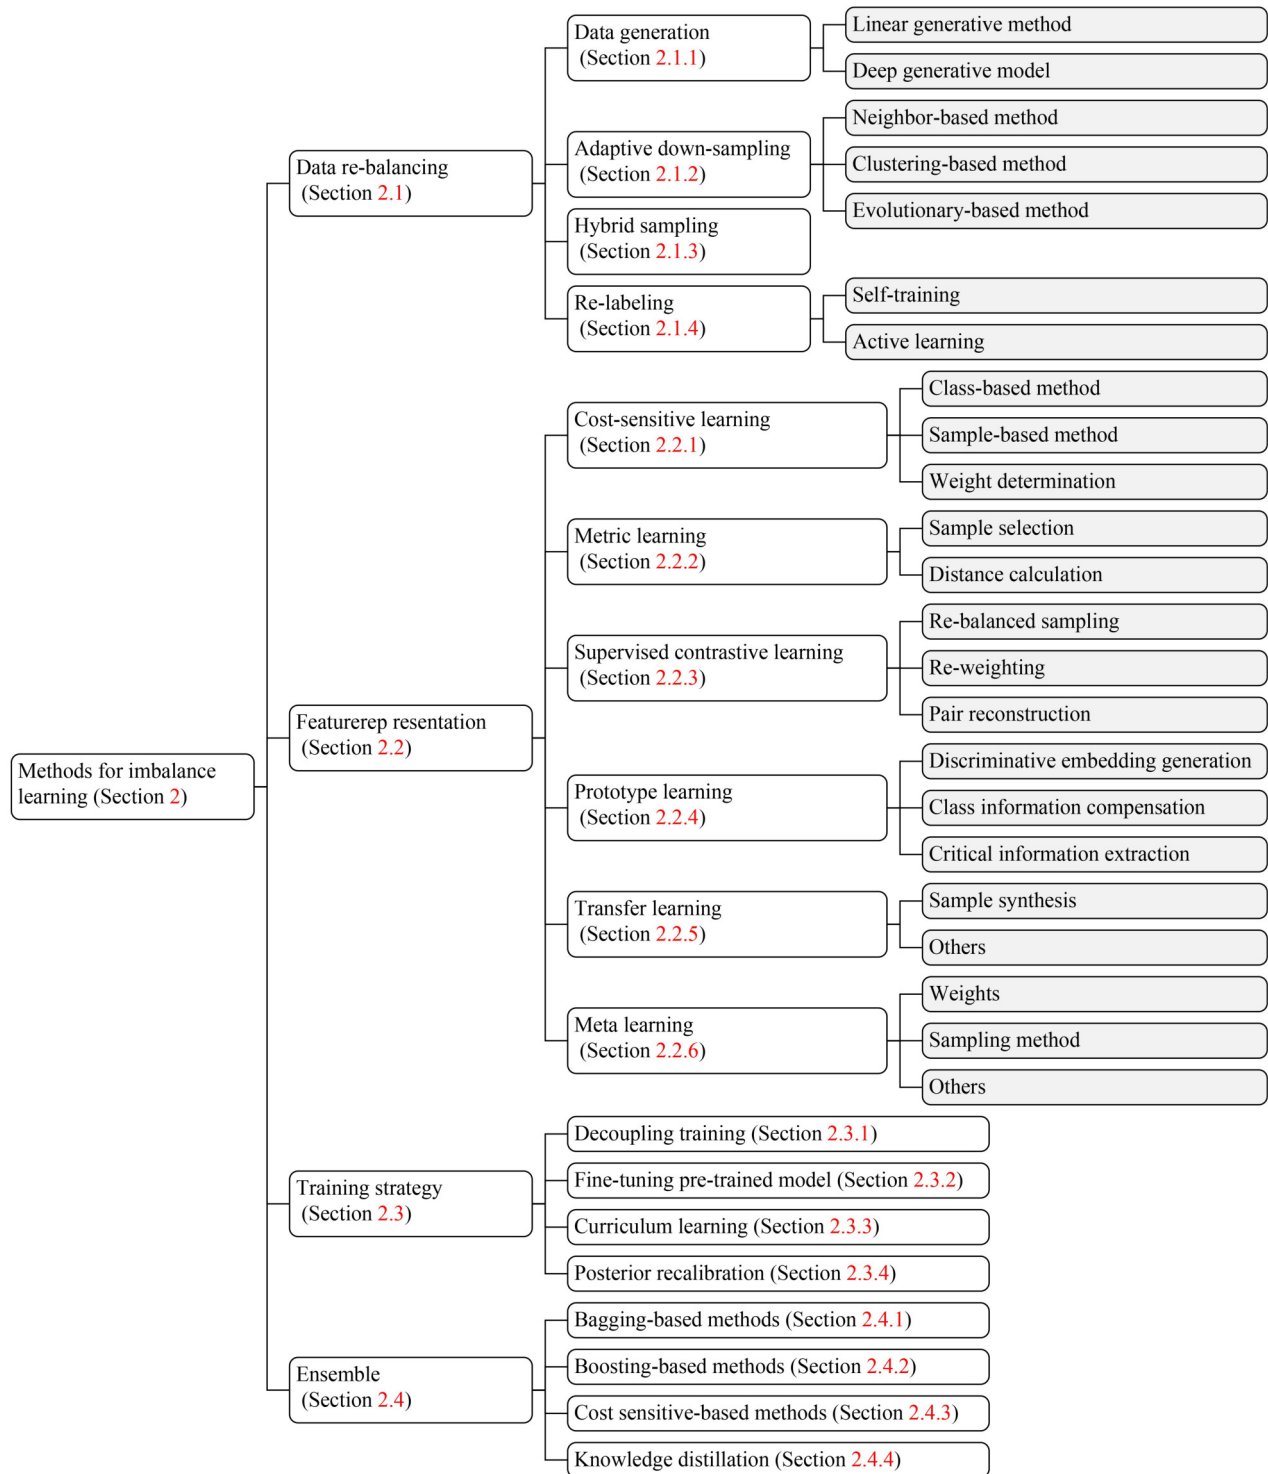


Fig. 1 A taxonomy of imbalanced learning methods. (1) Data re-balancing methods focus on modifying the input data distribution to mitigate class imbalance during the data preparation phase. (2) Feature representation methods aim to learn discriminative embeddings that better capture minority class characteristics, improving class separability in the latent space. (3) Training strategies adjust the learning process to reduce bias toward the majority class during model training. (4) Ensemble methods integrate multiple models to enhance robustness and generalization, leveraging the diversity of learners to mitigate data imbalance effects during inference

2.1 Data re-balancing

Data re-balancing is a key strategy to address class imbalance by modifying the training data distribution. It operates at the data level, aiming to equalize class representation to prevent models from being biased toward majority classes as shown in Fig. 2.

2.1.1 Data generation

One straightforward approach to address class imbalance and augment minority samples is random oversampling, which involves duplicating randomly selected minority samples and adding them to the dataset. However, this fundamental method, while effective in balancing class distributions, can inadvertently lead to classifier overfitting. This occurs because random oversampling generates multiple identical instances, resulting in overly specific rules. To overcome the limitations of random oversampling, more sophisticated techniques have been introduced. Linear generative methods and deep generative models offer a refined solution by creating diverse synthetic samples for minority classes. These advanced techniques enhance the learning process for machine learning models, enabling them to make more accurate predictions for minority samples.

Linear generative methods. Linear generative methods generate synthetic minority samples by interpolating between samples or feature vectors, which includes the Synthetic Minority Oversampling Technique (SMOTE) [5] and Mixup [6].

SMOTE selects k pairs of nearest minority neighbors and performs linear interpolation to create new synthetic minority samples. The pseudo-code is shown as Algorithm 1. This method effectively mitigates overfitting and enhances classifier learning. Nonetheless,

Algorithm 1 SMOTE oversampling

Input: minority samples \mathcal{D}_m , number of neighbors k , uniform distribution $\mathcal{U}(0, 1)$

Output: Synthetic samples \mathcal{D}_{syn}

```

1: for all  $x_i \in \mathcal{D}_m$  do
2:   Find  $k$  nearest neighbors of  $x_i$  in  $\mathcal{D}_m$ 
3:   Randomly select a neighbor  $x_j$ 
4:    $x_{new} \leftarrow x_i + \lambda \cdot (x_j - x_i)$ ,  $\lambda \sim \mathcal{U}(0, 1)$ 
5:   Add  $x_{new}$  to  $\mathcal{D}_{syn}$ 
6: end for
7: return  $\mathcal{D}_{syn}$ 

```

SMOTE generates an equal number of synthetic samples for each original minority sample, disregarding neighboring sample distribution characteristics and potentially causing over-representation of certain classes. To address this, recent advancements include D-SMOTE [7], which generates synthetic samples using mean points of nearest neighbors, and N-SMOTE [8], which employs structural information in neighbor calculations. Adaptive techniques like Borderline-SMOTE only generate synthetic samples for boundary minority samples [9], and ADASYN [10] adjusts sample generation based on density distribution. Moreover, MSMOTE categorizes minority samples into safe, border, and latent noise groups and selectively generates samples for border instances [11]. With over 100 variants developed to date, SMOTE stands out as one of the most influential and effective methods in the realm of data generation for addressing class imbalance. Given the focus of our paper on a comprehensive survey for imbalanced data, we recommend readers who are particularly interested in SMOTE and its extensive family of variants to refer to a SMOTE survey [12].

In contrast to SMOTE, Mixup introduces a distinct approach, allowing for interpolations between samples from different classes while simultaneously considering linear interpolations of associated labels using a consistent mixing factor. However, it's important to note that the vanilla Mixup method may significantly disrupt the minority class distribution, potentially blurring the decision boundary. Recognizing this challenge, Remix [13] offers a flexible solution by relaxing the constraint of employing the same mixing factor for both features and labels when generating virtual mixed samples for imbalanced data. In Remix, the mixing factors for labels are thoughtfully determined, accounting for the class sizes of interpolated samples and bestowing greater importance upon minority samples. This strategic choice serves the purpose of shifting the decision boundary toward the majority class and enriching the diversity of information pertaining to the minority class through the introduction of majority samples. Nonetheless, it is imperative to acknowledge that the random sampling of interpolated data may exacerbate data imbalance, potentially resulting in suboptimal performance. To address this concern, several sampling methods have been introduced to further bolster Remix's efficacy. MixBoost [14] incorporates an information-theoretic measure in conjunction




Fig. 2 Data re-balancing methods. Adaptive down-sampling methods reduce the number of majority samples to balance the class distribution. Linear generative methods synthesize new minority samples through linear combinations of existing instances. Deep generative methods learn the underlying data distribution and generate synthetic samples accordingly. Relabeling methods introduce new class labels. Hybrid sampling methods simultaneously reduce the majority samples and generate minority samples to achieve a balanced dataset

with the trained classifier, facilitating the selection of instances in pairs. This process ensures that one instance in each pair is in close proximity to the classifier's decision boundary, while the other is positioned further away, contributing to the iterative enhancement of classifier performance. Balanced-MixUp [15] employs a dual-strategy approach by concurrently sampling imbalanced instances and balanced classes from the training data. This facilitates the creation of a balanced training distribution, effectively mitigating the risk of under-fitting in minority classes. Label Occurrence-Balanced Mixup [16] incorporates two independent class-balanced samplers, thereby upholding statistically balanced label occurrences for each class. This meticulous approach ensures a more equitable representation of classes in mixed data pairs.

Besides linear generation methods, recent research has explored kernel-based techniques to enhance representation capacity under class imbalance. These methods introduce non-linearity, making them particularly effective when data distributions or class boundaries are complex and non-linear. For example, Kernel-based SMOTE (K-SMOTE) [17] generates synthetic minority samples in the kernel-induced feature space of an SVM, leveraging the augmented Gram matrix to better align with the decision boundary and outperforming SMOTE across 51 benchmark datasets. Similarly, KernelADASYN [18] extends ADASYN by applying kernel density estimation to model the minority class distribution and weighting samples by difficulty, producing more informative synthetic data. These findings highlight the importance of nonlinear transformations in both representation learning and data augmentation, motivating the shift toward deep generative models for more powerful imbalance handling.

Deep generative models. Conventional data generation methods typically produce synthetic samples positioned along the line segment connecting minority samples. Consequently, these techniques rely on local information and may not capture the overall distribution of the minority class effectively. Deep generative models, celebrated for their capacity to model the probability distribution of high-dimensional data, have gained widespread recognition for their ability to generate high-quality synthetic samples for minority classes.

Wan et al. [19] is a pioneering in the utilization of the generative models. Their work is dedicated to the precise modeling of high-dimensional data distributions via Variational Autoencoder (VAE), with a specific focus on generating synthetic samples for minority classes. VAE maps observed data points into a latent space while simultaneously learning the underlying data structure and probabilistic characteristics. The latent variable is represented by a probability distribution over possible values rather than a single fixed value, which is able to enforce the uncertainty of the model and lead to stable prediction. By employing VAE to generate synthetic samples through the utilization of random noise conforming to a standard normal distribution, this methodology not only ensures exceptional sample diversity but also consistently outperforms conventional data generation methods.

In addition to VAE, another noteworthy contribution to data generation methods is found in the realm of Generative Adversarial

Algorithm 2 GAN-based minority generation

Input: Real dataset $\mathcal{D} = \mathcal{D}_M \cup \mathcal{D}_m$, noise distribution P_z

Output: Generated synthetic minority dataset \mathcal{D}_{syn}

```

1: Initialize generator  $G$  and discriminator  $D$ 
2: while not converged do
3:   Sample real data  $x \sim \mathcal{D}$ 
4:   Sample noise  $z \sim P_z$ 
5:   Generate a synthetic minority sample:  $\hat{x} \leftarrow G(z)$ 
6:   Train discriminator  $D$  to distinguish real samples  $x$  from generated  $\hat{x}$ 
7:   Train generator  $G$  to improve  $\hat{x}$  so that  $D(\hat{x})$  is predicted as real
8: end while
9: Generate synthetic minority samples:  $\mathcal{D}_{\text{syn}} = \{\hat{x}_i = G(z_i) \mid z_i \sim P_z\}$ 
10: return  $\mathcal{D}_{\text{syn}}$ 

```

Networks (GAN) [20]. Douzas and Bacao [21] introduced an innovative approach that harnesses the power of conditional Generative Adversarial Networks (cGAN) [22] to approximate the genuine data distribution and produce synthetic data specifically tailored for the minority class. The pseudo-code for GAN is shown in Algorithm 2. A typical cGAN comprises two primary components: a generator and a discriminator. The generator's role is to generate synthetic data based on provided labels, while the discriminator is tasked with distinguishing real data from artificially generated data. The application of GANs to data rebalancing involves an iterative training process. Initially, both the generator and discriminator undergo training to ensure that the discriminator becomes proficient at differentiating between various data samples, and the generator becomes capable of producing synthetic data corresponding to specific class labels. Subsequently, the trained generator is employed to conduct data augmentation for the minority class. This adversarial training procedure enables the network to capture a more accurate data distribution, resulting in the production of more reliable and representative synthetic samples.

However, training a GAN directly for minority data generation poses significant challenges, primarily due to the scarcity of minority-class samples, which may be insufficient to effectively train a GAN model. In response to this challenge, BAGAN [23] is proposed to address this issue by proposing an alternative training strategy that incorporates data from both minority and majority classes simultaneously during the adversarial training process. The key innovation lies in class conditioning, which is applied within the latent space. Moreover, BAGAN leverages the discriminator not only for its traditional role of distinguishing fake samples from real ones but also as a classifier. This dual role enables the discriminator to not only discern the authenticity of generated samples but also identify the category of real samples. Consequently, BAGAN learns the underlying features of the specific classification problem by initially considering all available samples across both minority and majority classes. These shared features can then be applied universally across all classes, thereby facilitating the generation of high-quality

minorities. Similarly, SCDL-GAN [24] and RVGAN-TL [25] employ the Wasserstein auto-encoder [26] and VAE, respectively, to capture the distributive characteristics of all classes with guidance from label information before generating samples. EWGAN [27] constructs entropy-weighted label vectors for each class to characterize data imbalance and employs the Wasserstein Generative Adversarial Network [28] for stable training. EID-GANs [29] tackle extreme data imbalance by designing a novel penalty function that guides the generator to learn features of extreme minorities. Meanwhile, SMOTified-GAN [30] combines SMOTE with GAN, deploying GAN on pre-sampled minority data produced by SMOTE to enhance sample quality in minority classes. In contrast to the aforementioned GAN framework, GAMO [31] offers an alternative solution by integrating a convex generator, classifier, and discriminator into a three-player adversarial game. GAMO's convex generator creates new samples for minority classes through convex combinations of existing instances. Its primary objective is to perplex both the discriminator and the classifier, generating samples that challenge accurate classification. These synthetic samples are strategically positioned near class boundaries, thus aiding in adjusting the classifier's decision boundaries to reduce misclassification of minority class data.

Alongside GANs, diffusion models [32,33] have gained attention for their robust representation capabilities. These models present a compelling alternative to the intricate adversarial training procedures. They operate on a fundamentally distinct principle—sequentially denoising data—a methodology that has proven to be remarkably stable and effective in generating synthetic samples. Within the context of addressing data imbalance, several pioneering works have harnessed diffusion models to excellent effect. DiffMix [34] extends the utility of diffusion models to nuclei segmentation and classification tasks, which demand precise semantic image and label map pairs, alongside realistic pathology image generation. It diversifies synthetic samples by creating virtual patches and strategically shifting instance mask locations to balance class variances. Additionally, SDM [33] is employed within DiffMix to synthesize diverse, semantically realistic, and balanced pathology nuclei data, conditioned on generated label maps. DiffuASR [35] addresses data sparsity and the long-tail user problem in sequential recommendation by leveraging diffusion-based pseudo sequence generation. This approach introduces classifier-guided and classifier-free strategies to maintain the original intent of users throughout the diffusion process. MosaicFusion [36] introduces a diffusion-based data augmentation pipeline that simultaneously generates images and masks, particularly beneficial for large vocabulary instance segmentation. This method effectively mitigates distribution imbalances common in imbalanced and open-set scenarios. By running the diffusion process simultaneously on different image canvas regions and conditioning it on various text prompts, MosaicFusion controls the diffusion model to generate multiple objects at specific locations within a single image. PoGDiff [37] introduces a novel fine-tuning method for diffusion models on imbalanced datasets by replacing the ground-truth distribution with a Product of Gaussians that incorporates neighboring text embeddings,

significantly enhancing generation accuracy and quality. DiffROP [38] proposes a probabilistic contrastive learning approach to mitigate class bias in diffusion models by reducing distributional overlap among synthetic images, significantly improving generation fidelity for minority classes. LDMLR [39] utilizes a backbone model encodes the original imbalanced data into embeddings, and a Denoising Diffusion Implicit Model (DDIM) is trained to generate pseudo-features that enrich the minority representation. These generated features are combined with real embeddings to train the classifier, effectively reducing bias while keeping computation efficient. By leveraging denoising-based iterative refinement in the latent space, LDMLR produces diverse yet realistic features that improve class balance without sacrificing representation quality. Recently, BPA [40] has been proposed to address the majority class accumulation effect in diffusion models trained on imbalanced data; by adaptively assigning class-specific noise priors during training, BPA mitigates sampling bias introduced by uniform noise distributions, leading to improved generation diversity and visual fidelity across both majority and minority classes.

Discussion. To better inform practical deployment, we analyze the scalability and effectiveness trade-offs of data generation techniques. Linear generative methods, such as SMOTE and its variants (e.g., Borderline-SMOTE, ADASYN), create synthetic samples by interpolating between minority samples and their k nearest neighbors. These methods typically operate with a time complexity of $O(nkd)$, where n is the number of minority samples, and d is the feature dimensionality. Owing to their simplicity and low computational overhead, they scale well to relatively simple data structures, such as tabular datasets. However, the inherent linearity of these methods restricts their capacity to model complex data distributions.

In contrast, deep generative methods, including GANs, VAEs, and Diffusion Models, learn to synthesize samples by modeling the underlying data distribution. These approaches incur significantly higher computational costs. GANs and VAEs generally require $O(Tnd)$ operations, where T is the number of training epochs. Diffusion models are even more expensive, with time complexity of $O(TSnd)$, where S denotes the number of denoising steps per sample. While deep generative models offer improved diversity and fidelity, especially in high dimensional or unstructured data domains (e.g., images or molecular graphs), their scalability is often constrained by training time and hardware demands. To empirically evaluate this trade-off, we conduct experiments comparing both effectiveness and efficiency of Linear and Deep Generative Methods. Please refer to Section 5 for details.

2.1.2 Adaptive under-sampling

In contrast to data generation techniques that augment the original dataset, under-sampling methods involve the reduction of data within the original dataset. Random under-sampling, as the most fundamental under-sampling approach, entails the random reduction of the majority samples to achieve dataset balance. However, this method may result in the discarding of potentially valuable data, which could bear significance for the inductive process. To enhance

the efficacy of data under-sampling, researchers have explored strategies aimed at retaining the most representative majority samples. These strategies can be broadly classified into neighbor-based, clustering-based, and evolutionary-based methods.

Neighbor-based methods. Neighbor-based under-sampling methods revolve around the selection of samples based on their proximity to other samples. One notable technique is Tomek Link [41], which identifies pairs of examples as Tomek links if they belong to different classes and are each other's nearest neighbors. Under-sampling using Tomek Link can be executed by either removing all Tomek links from the dataset or specifically targeting the majority samples involved in these links. The removal of majority class nearest neighbors serves to reposition the learned decision boundary towards the majority class, thereby fostering the generation of minority samples. The pseudo-code is shown in Algorithm 3. One-sided selection (OSS) [41] extends the concept of Tomek Link to categorize majority samples into three groups: noise samples, boundary samples, and safety samples. Subsequently, it eliminates boundary and noise samples from the majority class. Condensed Nearest Neighbor (CNN) [42], on the other hand, leverages the nearest neighbor rule to iteratively decide whether a sample should be retained or removed. Central samples that can be accurately classified by the nearest neighbor rule are discarded. However, this method can be relatively slow due to its need for multiple passes over the training data, and it is sensitive to noise, potentially incorporating noisy samples into the retained subset. Edited Nearest Neighbor (ENN) [43] operates by removing samples whose class labels differ from the majority of their nearest neighbors. However, its efficacy is limited due to the prevalence of majority samples in close proximity to each other. Building upon ENN, the Neighborhood Cleaning Rule (NCR) [44] combines ENN with k -NN [45] to enhance the removal of noisy samples from the datasets. The NearMiss family of methods [46] introduces under-sampling techniques based on the distances between points in the majority class. NearMiss-1 retains majority samples whose mean distance to the k nearest minority class points is the lowest. In contrast, NearMiss-2 retains majority samples whose mean distance to the k farthest minority class points is the lowest. Finally, NearMiss-3 selects k nearest neighbors in the majority class for each point in the minority class. This approach directly controls the under-sampling ratio through the parameter k , eliminating the need for separate tuning. NB-Rec [47] investigates the presence of overlapping instances from various classes and introduces the utilization of the k -NN rule to investigate the local neighborhoods of

individual samples. It selectively eliminates majority samples from the overlapping region without impacting any minority samples, leading to enhanced detection of overlapped minority instances.

Clustering-based methods. Clustering-based under-sampling methods employ a clustering approach to select samples for under-sampling. The underlying concept is that datasets often comprise various clusters, each exhibiting distinct characteristics. This clustering strategy aims to capture both global information and sub-concepts within the dataset, thereby enhancing sampling quality. SBC [48] initiates the process by clustering the dataset into several clusters, each containing both majority and minority samples. When a cluster exhibits a higher concentration of minority samples and a reduced presence of majority samples, it signifies that this cluster aligns more with the characteristics of the minority class. Consequently, SBC identifies an appropriate number of majority samples to retain from each cluster, taking into account the ratio of majority samples to minority samples within the cluster. In contrast to clustering multiclass samples, CBMP [49] takes a direct approach by clustering the majority of samples and subsequently selecting samples that are either nearest to the centroid or chosen randomly. CentersNN [50] clusters the majority class with the number of clusters equal to the number of data points in the minority class, using cluster centers to represent the majority class. DSUS [51] introduces diversified sensitivity-based under-sampling, which selects valuable samples with high stochastic sensitivity concerning the current trained classifier. To ensure that samples are not exclusively chosen near the decision boundary and do not distort the data distribution, DSUS clusters the majority class and selects only one representative sample from each cluster for inclusion in the selection process.

However, clustering the majority class used in these methods is a computationally intensive process, which poses a significant challenge. To address this issue, Fast-CBUS [52] presents an efficient clustering-based under-sampling method that narrows its focus to clustering only the minority samples. Within each of these minority class clusters, Fast-CBUS selects majority samples that closely surround the minority instances and proceeds to create a dedicated model for each cluster. When classifying a new sample, the method prioritizes predictions made by models associated with clusters that exhibit the closest proximity to the sample. This approach is grounded in the idea that it's beneficial to train models using instances that are challenging to classify, particularly those majority samples that are situated near minority instances.

Evolutionary-based methods. EUS [53] presents a collection of techniques known as evolutionary under-sampling. These methods incorporate evolutionary algorithms into the process of under-sampling, with the primary aim of enhancing classifier accuracy by reducing instances predominantly from the majority class. To achieve this, EUS employs meticulously designed fitness functions, carefully striking a balance between data reduction, class balance, and classification accuracy. Following EUS, MapReduce [54] enhances it to make it feasible for large-scale datasets, even in the presence of extreme data imbalance. The approach involves partitioning the entire training dataset into manageable chunks. For each chunk

Algorithm 3 Tomek links under-sampling

Input: Dataset $\mathcal{D} = \{(x_i, y_i)\}$

Output: Cleaned dataset \mathcal{D}'

- 1: **for all** pairs (x_i, x_j) with $y_i \neq y_j$ **do**
 - 2: **if** x_j is the nearest neighbor of x_i and vice versa **then**
 - 3: Mark (x_i, x_j) as Tomek Link
 - 4: **end if**
 - 5: **end for**
 - 6: Remove majority samples in Tomek Links
 - 7: **return** \mathcal{D}'
-

containing minority data, the goal is to create a balanced subset using EUS, thus utilizing a representative sample from the majority set. EUSC [55] introduces a two-stage clustering-based surrogate model within the framework of evolutionary under-sampling, facilitating the expedited computation of fitness values. A key innovation is the creation of a surrogate model for binary optimization which is based on the meaning rather than their binary representation. EUSBoost [56,57] further enhances the performance of base classifiers by integrating the evolutionary under-sampling approach with Boosting algorithms. This evolutionary strategy prioritizes the selection of the most consequential samples for the classifier learning phase, based on accuracy and an additional diversity term incorporated into the fitness function. Consequently, it effectively addresses the challenges posed by imbalanced scenarios.

In addition to the three categories mentioned above, several alternative approaches have been developed from other perspectives. TU [58] introduces a meta-learning technique to parameterize the data sampler and trains it to optimize classification performance based on the evaluation metric. By incorporating evaluation metric optimization into the data sampling process, this method effectively determines which instances should be retained or discarded for a given classifier and evaluation metric. SDUS [59] utilizes supervised constructive learning, focusing on sphere neighborhoods (SPN) within the majority class. It offers two strategies: a top-down approach, treating SPN regions equally and favoring regions with smaller sample sizes, and a bottom-up approach, selecting samples proportionally based on SPN inner distribution for global pattern preservation. RIUS [60] leverages the information preservation principle and calculates cross-information information potential. It then selects the most relevant examples from the majority class, allowing for the extraction of the underlying structure of the majority class with a reduced number of samples. ECUBOOST [61] aims to prevent the loss of informative samples during data pre-processing for boosting-based ensembles. This method employs both confidence and entropy as benchmarks to ensure the validity and structural distribution of the majority samples during under-sampling.

2.1.3 Hybrid sampling

Data generation and under-sampling share the common goal of achieving a balanced class distribution, but each method comes with its own set of limitations. Under-sampling can lead to information loss due to the removal of examples, and data generation techniques can increase the dataset size with potentially synthetic and less informative samples. To achieve a balanced class distribution and optimize classifier performance, researchers have introduced hybrid sampling methods. These approaches often leverage linear generative techniques to over-sample the minority class. However, linear generative methods may fall short in capturing the complex structure of minority samples, particularly in extremely imbalanced scenarios characterized by significant class overlap. To address this challenge, some methods aim to enhance the effectiveness of linear generative techniques by integrating them with adaptive under-sampling strategies. For instance, SMOTE+Tomek Links and SMOTE+ENN [62] initially generate minority samples using SMOTE and

subsequently identify and eliminate instances involved in Tomek links or those that have been misclassified. CUSS [63] combines SMOTE with the clustering-based under-sampling method to minimize the risk of discarding valuable majority samples. SMOTE-RSB [64] and SMOTE-IPF [65] incorporate rough set theory and iterative ensemble-based noise filter known as the iterative-partitioning filter, respectively to prevent the generation of noisy samples that might disrupt the overall data distribution. BMW-SMOTE [66] controls the number of synthetic instances generated based on a weight calculated from the ratio of neighboring majority samples, which effectively reinforces the recognition of the minority class within local regions.

In addition to the approach that combines data generation and sample selection to achieve a balanced data distribution, advanced methods delve into the intricacies of sampling strategies to create diverse training data distributions. For example, BBN [67], SimCAL [68], and TLML [69] explore both the original imbalanced distribution and balanced distribution independently. Research endeavors [70–73] categorize classes into groups, each characterized by a balanced distribution. SADE [74] takes an innovative step by introducing the inversely imbalanced distribution to enhance the representation of minority classes in the dataset. Potential Anchoring [75] employs radial basis functions to measure class distribution density and maintains the original shape of class distributions by merging over- and under-sampling techniques. DLSA [76] introduces a streamlined approach by progressively adapting the label space, dynamically transitioning from imbalance to balance to facilitate classification. EHSO [77] incorporates an evolutionary algorithm to identify the optimal balance between classification performance and the replication ratio of random oversampling. Dynamic Sampling [78] leverages sample upgrading, making real-time adjustments to the dataset based on the training results. It randomly removes samples from well performing classes and duplicates samples from poorly performing classes, ensuring that classification model learns relevant information at each iteration.

Discussion. While effective, the overlapping problem remains a significant challenge in sampling methods. These methods address class overlap in imbalanced data primarily through two strategies: (1) focusing on safe regions and (2) filtering noisy or borderline samples. The first strategy generates synthetic minority samples only within areas where minority instances are densely clustered and far from decision boundaries. This reduces the risk of creating synthetic points in overlapping regions, which could mislead the classifier. The second strategy involves applying noise-filtering techniques such as Edited Nearest Neighbors or Tomek Links, which identify and remove majority samples that lie close to or within minority regions. This helps clean the decision boundary and reduces ambiguity caused by overlapping components. As for limitations, one key issue is boundary distortion. If down-sampling removes too many majority samples near the decision boundary, it can oversimplify the class separation, leading to underfitting. Additionally, there is the risk of losing valuable majority data during the cleaning process. In scenarios where class overlap is genuine, aggressively removing overlapping majority instances may discard useful information

necessary for accurate classification. These trade-offs highlight the need to balance data cleaning with decision boundary preservation in hybrid strategies.

2.1.4 Re-labeling

Apart from data generation and under-sampling, researchers have recognized the untapped potential of unlabeled data as a valuable resource. By incorporating self-training and active learning strategies, these unlabeled samples can be effectively labeled to compensate for the lack of imbalanced category supervision information and improve classification ability.

Self-training. Self-training is a widely used technique where a model is initially trained on a labeled dataset and then iteratively improves itself by using its predictions (i.e., pseudo-labels) on unlabeled data to generate more training samples. In line with this approach, Yang and Xu [79] initially experimented with using pseudo-labels from minority classes to address class imbalance and subsequently retrained the model. Expanding upon this concept, CReST [80] made an important observation that models trained on imbalanced datasets often achieve low recall but high precision on minority classes. This suggests that these models are proficient at generating high-precision pseudo-labels for minority classes. To leverage this insight, CReST introduced a stochastic update strategy for pseudo-label selection, assigning higher selection probabilities to samples predicted as minority classes, which are more likely to be accurate. To enhance the quality of pseudo-labels, SaR [81] introduced a calibration step that adjusts soft labels based on prediction bias before generating one-hot pseudo-labels. This calibration process aims to simultaneously enhance the recall of minority classes and the precision of majority classes. In addition to utilizing linear classifiers for generating pseudo-labels, DASO [82] explored the use of semantic pseudo-labels. These labels are derived by measuring the similarity between a given class center and unlabeled samples in feature space. DASO's research demonstrated that semantic pseudo-labels excel in predicting minority classes with high recall. As a result, the combination of linear and semantic pseudo-labels offers a promising approach to mitigating overall dataset bias.

Active learning. Active learning is a technique that seeks to minimize the human labeling effort by automating the selection of the most informative unlabeled samples for human annotation. In the context of imbalanced data, active learning becomes a vital tool for addressing class imbalance and mitigating the potential error labels introduced by self-training. Various active learning strategies have emerged, each focusing on different aspects of informativeness, such as margin, uncertainty, and diversity.

Margin-based selection methods, as exemplified by SVMAL [83] and AI-WSELM [84], choose the unlabeled instance closest to the classifier's decision boundary as the most informative sample. This approach is rooted in the concept that uncertain or ambiguous samples hold essential information critical for shaping the decision boundaries of a classifier. Hence, they should receive the highest priority for inclusion in the training set.

Uncertainty-based strategies for active learning rely on quantifying

sample uncertainty using prediction entropy or confidence metrics. For instance, BootOS [85] employs the entropy of prediction as an indicator of uncertainty for sample selection. It further introduces a novel concept of using max-confidence as an upper bound and min-error as a lower bound for halting the active learning process. Building upon this foundation, several other works have extended the uncertainty-based sampling approach, as evidenced in [86–89].

In contrast to these uncertainty-based approaches, CEAL [90] acknowledges a unique challenge when applying active learning in imbalanced datasets. In such scenarios, the minority class often becomes overwhelmed by the majority class during the labeling process. To mitigate this imbalance, CEAL introduces the concept of “certainty-enhanced neighborhoods.” For each unlabeled sample, it recommends considering only the local context where the majority class density is low. This strategic selection process reduces the likelihood of querying majority samples, offering a solution to the issue of minority class overshadowing during active learning.

However, these methods relying solely on a single model's output may not always perform optimally. This is particularly evident in cases where the initial model was trained on limited data and may produce biased sample selections throughout the learning process. To address these limitations, other sampling strategies focus on the diversity [86,88,91] of selected samples. For example, SAL [91] employs active learning by training a similarity model that recommends unlabeled minority class samples for manual labeling.

2.2 Feature representation

In addition to re-balancing classifier training, obtaining a high-quality and distinctive feature representation for data is essential for mitigating the impact of class imbalance on model performance as depicted in Fig. 3. Several techniques can be employed to enhance feature representation in imbalanced learning, such as cost-sensitive learning, metric learning, supervised contrastive learning, prototype learning, transfer learning, and meta-learning.

2.2.1 Cost-sensitive learning

Cost-sensitive learning has emerged as a highly effective approach for tackling the imbalance challenges. The foundation of cost-sensitive learning is rooted in the concept of cost matrices [92]. In contrast to conventional machine learning algorithms that treat all classification errors equally, cost matrices introduce the pivotal notion of assigning predefined costs to misclassified samples, i.e., false positives and false negatives. These assigned costs effectively function as weight factors, elevating certain misclassification errors




Fig. 3 Feature representation methods aim to obtain high-quality and distinctive feature representations for data

above others in terms of importance. The overarching objective is to minimize the overall expected cost, thereby steering models to place heightened emphasis on minority classes, transcending the mere fixation on average classification errors. In essence, the concept of “costs” embodies the role of “weights” in the context of reweighting errors, enabling models adjust boundaries and focus on important classes¹⁾.

In the early stages, methods focus on class-specific costs and entail assigning distinct costs or weights to different classes, with these predefined costs typically established based on experiential knowledge or domain expertise [92,93]. The representative method is MetaCost [94], a general approach aimed at enabling non-cost-sensitive algorithms to incorporate cost-sensitive considerations. The key innovation lies in the re-labeling procedure, which begins by training multiple models on bootstrap replicates of the training dataset. Subsequently, the samples are re-labeled based on the minimized cost function. Finally, the classifier is trained using these refreshed labels. Another approach, threshold-moving [95], involves adjusting the output threshold towards the less expensive classes. This adjustment makes it more challenging to misclassify examples with higher associated costs. This method employs the training set to train a neural network, and cost sensitivity is introduced during the test phase by directly modifying the output logits based on the associated costs. In addition to binary classification, Zhou and Liu [96] explored the application of cost-sensitive learning in multi-class classification tasks. They observed that the consistency of cost matrices has a significant impact on classification performance. In cases where the cost matrix lacks consistency, rescaling is recommended, particularly after decomposing the multi-class problem into a series of two-class problems. CSMLP [97] presents a novel approach involving the design of a joint objective function for Multilayer Perceptrons. This method utilizes a single cost parameter to differentiate the importance of class errors. The class-balanced loss [98] emphasizes that as the number of samples increases, the additional benefit of adding a new data point diminishes. To address this, it introduces the concept of the effective number to approximate the expected class size and leverages this effective number to calculate weights, rather than relying solely on the sample quantity.

Besides the class-specific costs, cost-sensitive learning extends its scope to encompass the weighting of individual samples. Focal loss [99] offers a perspective on addressing imbalance problems at the instance level, introducing varying weights for each training sample as shown in Algorithm 4. Given that class imbalance often amplifies the difficulty of predicting minority classes, focal loss assigns higher costs to more challenging samples while applying lower costs to those that are easier to classify during training. Influence-balanced loss [100] takes a different approach by employing influence functions to gauge the impact of individual samples on intricate and biased decision boundaries. The loss associated with each sample is proportionally weighted based on the inverse of its influence. Balanced softmax [101] extends the conventional softmax function to

Algorithm 4 Focal loss computation

Input: Logits z , true label y , hyperparameters γ and α
Output: Loss \mathcal{L}

- 1: Apply softmax to obtain the predicted probability for the true class:
- 2: $p \leftarrow \text{softmax}(z)_y$,
- 3: Down-weight well-classified examples and focus on hard ones using the modulating factor $(1 - p)^\gamma$:
- 4: $\mathcal{L} \leftarrow -\alpha \cdot (1 - p)^\gamma \cdot \log(p)$
- 5: **return** \mathcal{L}

address imbalanced distributions. It accomplishes this by incorporating the imbalanced conditional prediction probability. This generalized softmax function serves as a form of logit calibration, effectively adjusting each class’s logit value in proportion to its label frequencies, thus providing a more balanced representation in the feature space.

Another critical aspect of cost-sensitive learning revolves around the challenge of determining appropriate weights. Predefined weights often necessitate careful tuning of additional weight hyperparameters, introducing complexity to the process. To address this issue, researchers have explored diverse avenues for automatically learning the weights, considering various perspectives, including evolutionary search, gradient-based techniques, distribution transport, and meta-learning. EDS [102] employs an evolutionary search method to effectively seek optimal misclassification costs in multiclass imbalance scenarios based on an objective function defined with G-mean. This approach optimizes cost parameters through an evolutionary strategy. ECS-DBN [103] utilizes adaptive differential evolution to optimize misclassification costs based on the training data. This enhances the effectiveness of deep belief networks by incorporating evaluation measures into the objective function. For gradient-based methods, GDW [104] addresses the imbalance problem through class-level gradient manipulation. It efficiently obtains class-level weights by developing a two-stage weight generation scheme that involves storing intermediate gradients. Equalization Loss [105] and Equalization Loss v2 [106] determine weights from gradients as well. While Equalization Loss protects the learning of minority classes by blocking some negative gradients, it acknowledges that positive gradients may still be overwhelmed by the negative ones. Therefore, Equalization Loss v2 proposes balancing the positive-to-negative gradient ratio, using the accumulated gradient ratio as an indicator to up-weight positive gradients and down-weight negative gradients. OT [107] approaches the issue by viewing the training set as an imbalanced distribution over its samples. It transports this distribution to a balanced one using optimal transport. The weights of the training samples are determined as the probability mass of the imbalanced distribution, learned by minimizing the optimal transport distance between the two distributions. Besides, meta-learning methods are utilized to learn optimal weights through bi-level optimization as well. For

¹⁾ In this survey, we utilize the terms “cost” and “weight” concurrently, following the method used in the original paper.

instance, LRE [108] presents a novel meta-learning algorithm that learns to assign weights to training examples based on their gradient directions. Meta-Weight-Net [109] automatically learns an explicit loss-weight function, parameterized by a multilayer perceptron, from the data.

2.2.2 Metric learning

Metric learning aims to learn a discriminative feature space by designing task-specific distance metrics between samples, which makes samples of the same class closer together, while samples of different classes are further apart [110], as shown in Algorithm 5. It achieves great success in feature representation due to the powerful capacity of understanding the similarity relationship among samples. Thus, metric learning methods has been widely applied to mitigate the imbalanced classification problems to increase the separability between minority and majority classes and achieve a better classification performance on minorities.

Sample selection. In metric learning methods, sample selection is the most important role because it directly influences the distances calculation and the quality of the extracted similarity relationship. Therefore, the sample selection of metric learning methods for imbalanced data always emphasizes the minority class [111–115] and carefully selects samples to design the distance metrics.

To balance the training samples used in the training procedure, LMLE [111] clusters samples in each class and construct the sample sets by selecting four representative samples from neighboring clusters and classes for each anchor sample. Benefiting to clustering, LMLE captures the significant data variability within each class and greatly enriches the diversity of minority sample sets for distance evaluation. According to four distances defined on the selected sample sets, the triple-header hinge loss constrains the feature representation and enforces a local margin while handling class imbalance problems. However, the clustering is sensitive to the entire training data and class size and usually gets limited results for extremely sparse minority classes. To address this problem, class rectification loss [112] construct the triplets by identifying hard-positive and hard-negative samples for minorities in each training batch and designs a batch-wise hard mining metric based on relative, absolute, and distribution information among selected triplets. Constraining hard-positive samples facilitates the discovery of

sparsely represented minority class boundaries, whereas mining hard-negative samples enlarges the margins against confounding impostor classes. Together, these strategies alleviate intra-class compactness and inter-class sparsity in the feature space. LSTM-QDM [114] utilizes the minority sample to expend the data triplet in the triplet network [116] and constructs a quadruplet loss to balance the relationship among anchor, positive, negative and minority samples. Benefiting from the extra minority samples in quadruplet loss and optimizing the softmax loss jointly, LSTM-QDM achieves a stronger representation ability for imbalanced imbalanced time-series data. Inspired by Borderline-SMOTE, DMFBML [115] separates the samples into three categories: danger, safe and noise based on the ratio of neighbors from the different classes and generates the triplets by the guidance of three categories. Afterward, DMFBML designs a borderline margin loss function to customize the margins among different categories and assign more metric constraints for minority samples, thereby the minority and majority samples in overlapping regions could be separated.

Distance calculation. Besides introducing minority samples in the sample selection strategy, other works handle the imbalance data and constrain the minority class from distance calculation. Range loss [117] explicitly defines intra-class loss as the harmonic mean of the k -largest distances between samples within each class, and constructs inter-class loss by measuring the minimum distance among class centers. By optimizing these two components, range loss can simultaneously reduce the intra-class variations of different classes and enlarge the inter-class distance. To treat different classes equally and separate them balanced, DMBK [118] first assumes that different classes are sampled from Gaussian densities with different expected values but one identical covariance, and then defines normalized divergence using KL-divergence to represent inter-class divergence. Afterword, inspired by arithmetic-geometric average inequality, DMBK makes use of the geometric mean to constrain the normalized divergences and make all inter-class divergences balanced. This procedure separates all classes in a balanced way and avoids neglected important minorities and inaccurate similarities incurred by imbalanced class distributions. The Affinity loss function [119] proposes to measure class similarities for an input feature in terms of the Bergman divergence. Additionally, it utilizes a diversity regularizer to constrain all class centers to be uniformly distributed throughout the feature space. This approach helps the model avoid the majority classes from occupying a disproportionately larger angular space compared to the minority classes, leading to a more balanced embedding space. Imbalanced metric learning [120] designs a new weighting method for Mahalanobis distance-base metric learning methods by decomposing the pair of samples based on their labels and introduces the hyper-parameters to control the negative effect of the imbalance. Moreover, imbalanced metric learning derives a generalization bound theoretically involving the proportion of positive samples using the uniform stability framework.

2.2.3 Supervised contrastive learning

Supervised contrastive learning has emerged as a promising research

Algorithm 5 Triplet loss

Input: Anchor x_a , Positive x_p , Negative x_n , Margin m , encoder f

Output: Loss \mathcal{L}

- 1: Encode each sample into the latent space using f
 - 2: Compute the distance between the anchor and the positive (same class):
 - 3: $d_{ap} \leftarrow \|f(x_a) - f(x_p)\|^2$
 - 4: Compute the distance between the anchor and the negative (different class):
 - 5: $d_{an} \leftarrow \|f(x_a) - f(x_n)\|^2$
 - 6: Enforce that the anchor is closer to the positive than to the negative by at least margin m :
 - 7: $\mathcal{L} \leftarrow \max(0, d_{ap} - d_{an} + m)$
 - 8: **return** \mathcal{L}
-

direction in recent years, combining metric learning and self-supervised contrastive learning to leverage the advantages of both methods while overcoming their respective limitations [121]. On the one hand, conventional metric learning methods (e.g., the triplet loss [122] and N -pair loss [123]) and self-supervised contrastive learning [124] limit the number of positive samples for each anchor to one, which places more emphasis on optimizing individual pairs and hinders the extraction of richer inter-sample information [121]. On the other hand, while contrastive learning methods have been found to generate a balanced feature space and maintain stable performance, even when the degree of imbalance becomes increasingly severe [79,125], the features learned through contrastive learning lack semantic separability since they are learned without class label guidance, leading to poor classification performance. To address these limitations and handle the imbalance problem, supervised contrastive learning methods incorporate label information into the contrastive learning process and extend positive examples by including other samples from the same class along with data augmentations. This encourages the model to produce closely aligned representations of all samples from the same class and results in a more robust clustering of the representation space.

SupCon [121] first extends contrastive learning to the fully-supervised setting and surpasses the performance of the traditional supervised cross-entropy loss by leveraging class labels with a contrastive loss. The pseudo-code is shown in Algorithm 6. Although achieving performance improvement, researchers discover that the supervised contrastive loss still suffers from the imbalance of positive/negative pairs. To mitigate this issue and further enhance supervised contrastive learning for imbalanced data, the improved methods can be categorised into three classes, including re-balanced sampling, re-weighting, and pair reconstruction.

Re-balanced sampling methods aim to address the issue of imbalanced positive/negative pairs by sampling techniques. KCL [125] improves upon the positive set construction rule in SupCon by randomly selecting k instances for each anchor from the same class to form the positive set, instead of using all the instances from the same class. This modification ensures that there is an equal number of positive samples for each anchor sample, thereby avoiding the

dominance of instance-rich classes in the representation learning process. CA-SupCon [126] argues that due to the scarcity of minority samples, SupCon cannot guarantee the separability of minority classes in one mini-batch. To address this issue, CA-SupCon introduces a class-aware sampler that re-balances the data distribution within each mini-batch during training, thus allowing each minority class to have a fair chance to contribute to the feature distance optimization. GPCL [127] introduces self-training to supervised contrastive learning to provide pseudo-labels for unlabeled data. By incorporating these pseudo-labels with confidence scores, GPCL can make full use of unlabeled data to enrich the contrastive data while filtering out unreliable negative and positive pairs with the same pseudo-labels and low confidence scores in the contrastive loss, respectively.

Re-weighting methods aim to increase the importance of minority pairs in the contrastive loss. For example, PaCo [128] introduces a set of parametric class-wise learnable centers into the contrastive loss. This approach adaptively weighs the samples from different classes and enhances the intensity of pushing samples of the same class closer to their corresponding centers, effectively giving more weight to the minority class pairs. BCL and TSC [129,130] argue that supervised contrastive learning should ideally result in an embedding where different classes are uniformly distributed on a hypersphere. However, applying the contrastive loss to imbalanced data can lead to poor uniformity, with minority classes having poor separability in the feature space. BCL provides evidence for this claim through a theoretical analysis of the lower bound of the contrastive loss. The analysis shows that imbalanced data distribution heavily affects the model and the resulting representations of classes no longer attain a regular simplex configuration, which ultimately hinders the final classification performance. Therefore, to address this issue, BCL proposes three class-averaging solutions to re-weight the negative pairs in the supervised contrastive loss. This helps to reduce the gradients from negative majority classes and contributes to an asymmetrical geometry configuration. Additionally, to ensure that all classes appear in every mini-batch, BCL introduces class-center representations in the positive and negative pairing, compensating for the batch sampling bias of high-frequency classes. Although motivated by the same insight, it is worth noting that TSC proposes a distinct solution from BCL. TSC aims to maintain a uniform distribution in the feature space for all classes by urging the features of classes closer to the uniformly distributed target features on the vertices of a regular simplex. To maintain semantic relations among classes, TSC proposes an online matching method to ensure that classes that are semantically close to each other are assigned to targets that are also close to each other.

Pair reconstruction. In addition to re-balanced sampling and re-weighting techniques, researchers also attempt to leverage the pair reconstruction in supervised contrastive learning to mitigate the data imbalance. SNP-ECC [131] and MLCC-IBC [132] extend the point-to-point contrastive learning approach to a point-to-group approach, and mitigate the imbalance problem by enriching the training samples. They construct contrastive sample groups by randomly sampling from the k -nearest majority/minority neighbors of the

Algorithm 6 Supervised contrastive loss

Input: Normalized embeddings $\{z_i\}_{i=1}^N$, labels

$\{y_i\}_{i=1}^N$, temperature τ

Output: Loss \mathcal{L}

```

1:  $\mathcal{L} \leftarrow 0$ 
2: for  $i = 1$  to  $N$  do
3:   Define the positive sample:  $\mathcal{P}(i) \leftarrow \{p \neq i \mid$ 
      $y_p = y_i\}$ 
4:   Define the negative sample:  $\mathcal{N}(i) \leftarrow \{n \neq i \mid$ 
      $y_n \neq y_i\}$ 
5:    $Z \leftarrow \sum_{j \in \mathcal{P}(i) \cup \mathcal{N}(i)} \exp(\text{sim}(z_i, z_j) / \tau)$ 
6:   for all  $p \in \mathcal{P}(i)$  do
7:      $\mathcal{L} += -\frac{1}{|\mathcal{P}(i)|} \log \frac{\exp(\text{sim}(z_i, z_p) / \tau)}{Z}$ 
8:   end for
9: end for
10:  $\mathcal{L} \leftarrow \frac{1}{N} \mathcal{L}$ 
11: return  $\mathcal{L}$ 

```

anchor sample, respectively. Afterwards, the new-format training data are generated by pairing the anchor sample with various contrastive sample groups. Due to the increased diversity of the contrastive sample groups, the new training data pairs are balanced and are learned using a novel label matching task. ProCo [133] proposes to utilize the category prototype to augment representative samples and generate the adversarial proto-instance via linear interpolation in the feature space. The adversarial proto-instance is designed as a special outlier that can balance the data and encourage ProCo to rectify the decision boundaries of the minority categories during contrastive learning.

2.2.4 Prototype learning

In addition to exploring correlations among extensive and complex samples via metric learning and supervised contrastive learning, researchers have incorporated the idea of prototype learning [134] from few-shot learning into the imbalanced data scenario to enhance feature representations of minority classes. A prototype explicitly represents a summary of all instances belonging to a specific class or intrinsic sample characteristics. In the context of imbalanced data, prototypes can improve minority representation from three aspects: facilitating the generation of a more discriminative feature space, compensating for class information compensation, and aiding in the extraction of intrinsic sample information.

Discriminative embedding generation. Class-specific prototypes provide a central target for each class in the feature space. This enables sample embeddings explicitly converge toward the class center, contributing to a more distinctive feature space for classification. Hybrid-PSC [135] improves the supervised contrastive learning method by treating the prototypes as the pairing object for positive/negative pairs, forcing each sample to be pulled towards the prototypes of its class and pushed away from prototypes of all the other classes. The pseudo-code is shown in Algorithm 7. This approach allows for more robust and efficient data sampling when dealing with large-scale datasets under a limited memory budget. The Affinity loss function [119] provides a theoretical analysis of the Softmax loss with a linear prediction layer and treats the linear

weight vectors as the class prototypes. Based on this, the affinity loss function can automatically learn the prototypes for each class and conveniently measures feature similarities with class prototypes in Euclidean space using Bergman divergence. Instead of calculating prototypes from labeled data, TSC [130] manually generates a set of targets uniformly distributed on a hyper-sphere as the class prototypes and forces all classes to converge to these distinct distributed targets during training for a uniform distributed feature space. BCL [129] uses supervised contrastive loss to train the model, and specifically ensures that each sample is paired with all other class prototypes in the negative pairing. By doing so, BCL compensates for the batch sampling bias of high-frequency classes, which in turn enables comprehensive contractive information from each class to be captured during training. It is worth noting that the number of class prototypes is not limited to one, and studies [119,135] have shown that adopting a multi-prototype paradigm can capture richer within-class data distribution and promote diversity and discriminability in the feature space. For example, Prototype-based Augmentation (ProAug) [136] builds a multi-center prototype dictionary optimized with a category-aware margin loss to capture the multi-modal distribution of each class. Based on this dictionary, ProAug dynamically composes context-similar prototypes via an attention mechanism to generate informative minority features, thereby enriching the representation space and mitigating biased decision boundaries.

Class information compensation. As a summary of classes, class prototypes contain basic class information that can help compensate for the information shortage of minority classes by transferring knowledge across classes or generating minority samples directly through data augmentation. OLTR [137] utilizes an attention mechanism to adaptively merge all prototypes with each sample embedding. This helps transfer basic class concepts among all classes to the feature embeddings, thereby enhancing the minority embeddings. PAN [138] addresses the modality-imbalance problem, where the available data for text, image, and audio are all considerably imbalanced, by generating synthetic samples for the minority modalities. It learns semantically consistent prototypes that capture the semantic consistency of the same class across different modalities and uses a gated recurrent unit to fuse the nearest neighbors of the majority modality to generate new samples. MPCL [139] argues that empirical prototypes are severely biased due to the scarce data in minority classes and proposes a meta-learning method to calibrate the prototypes by contrastive loss. With the assumption that feature embeddings for each class lie in a Gaussian distribution, MPCL estimates the mean and covariance of class distribution by the calibrated prototypes. Then, MPCL samples new samples from the distribution as synthetic minorities to balance the classes. Similarly, ProCo [133] proposes to utilize the class prototype to augment representative samples and generate the adversarial proto-instance to balance the data. This encourages ProCo to rectify the decision boundaries of the minority classes.

Critical information extraction. Prototypes not only represent entire classes but can also be leveraged to extract intrinsic information from individual samples, which is crucial in handling

Algorithm 7 Prototype-based contrastive loss

Input: Embeddings $\{z_i\}_{i=1}^N$, labels $\{y_i\}_{i=1}^N$, temperature τ

Output: Loss \mathcal{L}

- 1: Compute prototype (mean embedding) for each class c :
 - 2: $p_c \leftarrow \frac{1}{|\mathcal{I}_c|} \sum_{i \in \mathcal{I}_c} z_i$, where $\mathcal{I}_c = \{i \mid y_i = c\}$
 - 3: $\mathcal{L} \leftarrow 0$
 - 4: **for** $i = 1$ to N **do**
 - 5: Let $p^+ \leftarrow p_{y_i}$ be the prototype of sample i 's true class
 - 6: Let $\mathcal{P} = \{p^+\}$ and $\mathcal{N} = \{p_c \mid c \neq y_i\}$ be positive and negative prototypes
 - 7: Compute denominator: $Z \leftarrow \sum_{p \in \mathcal{P} \cup \mathcal{N}} \exp(\text{sim}(z_i, p)/\tau)$
 - 8: $\mathcal{L} += -\log \frac{\exp(\text{sim}(z_i, p^+)/\tau)}{Z}$
 - 9: **end for**
 - 10: $\mathcal{L} \leftarrow \frac{1}{N} \mathcal{L}$
 - 11: **return** \mathcal{L}
-

typical characteristics of different classes. For instance, IEM [140] proposes to use the attention mechanism to learn multiple discriminative feature prototypes for each class, considering both local and global features, as a single prototype may not be sufficient to represent a class. By gradually updating these prototypes during the training procedure, IEM captures the representative local and global semantics of each class and incorporates more discriminative features to improve generalization on minority classes. MPNet [141] generates both instance and semantic prototypes shared by all samples, representing local representative features and semantic features, respectively. The learned prototypes are recorded in a memory module and used to represent each sample, addressing the issue of forgetting minority classes during training. Additionally, semantic memory regularization is applied for a separately distributed centroid space.

2.2.5 Transfer learning

Transfer learning is a technique that leverages salient features and patterns identified in a source domain to enhance performance in a target domain. In the context of imbalanced datasets, minority classes often suffer from a scarcity of labeled data, while majority classes benefit from an abundance of such data. Consequently, researchers concentrate on identifying shared characteristics or attributes among multiple classes and employ transfer learning to transfer common knowledge from well-represented majority classes to underrepresented minority classes.

Sample synthesis. A substantial portion of transfer learning methods focus on transferring the majority class diversity to minority classes by synthesizing new minority samples. For instance, FTL [142] presents a prototype-based transfer framework, predicated on a Gaussian prior encompassing all classes. Under this assumption, each class is partitioned into mean and variance, representing class-specific and class-generic factors, respectively. The shared variance of major classes, indicating commonalities between categories such as varying poses or lighting conditions, is conveyed to minor classes through the generation of synthetic samples. The pseudo-code is shown in Algorithm 8. LEAP [143] models the distribution of intra-class features based on the distribution of angles between features and their corresponding class centers. By assuming that these angles adhere to a Gaussian distribution, the angular variance learned from the majority class is transferred to each minority class, and synthetic samples for minority classes are drawn from the distribution with enhanced intra-class angular diversity. Similarly, FSA [144] employs class activation maps to identify class-specific features from underrepresented classes and fuses them with class-generic features to generate minority samples. RSG [145] devises a minority class sample generator module, compatible with any backbone model, to create new samples based on the variation information of majority classes. In contrast to the aforementioned methods that generate samples from minority instances, M2m [146] directly transfers the diversity of majority classes by translating majority samples into minority samples. By referring to the output of the pre-trained classifier, a small optimized noise is added to the majority sample to facilitate the translation. In this manner, the richer information of majority samples is effectively leveraged to augment the minority

classes. Sample-Adaptive Feature Augmentation (SAFA) [147] extracts transferable, class-irrelevant semantic directions from majority class features using an auto-encoder framework, then translates minority samples along these directions to generate diverse, sample-specific augmented features. A recycling training scheme enforces semantic consistency, contrastive loss ensures transferability, and mode-seeking loss promotes diversity. As a plug-in module, SAFA can be integrated with various networks and loss functions, improving decision boundary calibration without extra inference cost.

In addition to sample synthesis, several methods endeavor to transfer knowledge in alternative ways.

MetaModelNet [148], rather than transferring majority class knowledge directly, conveys the changing trend of model parameters and introduces a meta-learning method to learn the model dynamics, predicting how model parameters will alter as more training data is progressively incorporated. Specifically, MetaModelNet is trained to forecast many-shot model parameters from few-shot model parameters, which are trained on abundant and limited samples, respectively. By employing the learned model dynamics, MetaModelNet can directly enhance the model parameters trained on minority instances by predicting the parameters when additional minority samples are included.

OLTR [137] explicitly constructs class centroids and exploits the attention mechanism to amalgamate sample features with class information from various classes, yielding more comprehensive feature representations. This information aggregation not only encompasses minority class information but also transfers basic class concepts among majority classes. Furthermore, the attention mechanism enables the model to adaptively tailor the importance of class information for different samples during end-to-end training, significantly enriching the minority features that lack adequate supervision signals.

GistNet [149] posits that all classes share a common geometric structure in the feature space. With numerous samples in the majority class, GistNet can encode the class geometry into a set of structure parameters shared by all classes, and the geometric structure of the minority class is restricted via learned structure parameters.

RVGAN-TL [25] initially generates minority samples using GAN and subsequently applies the TrAdaboost algorithm [150] to assign different weights to the training samples. Synthetic minority samples that deviate from the real sample distribution are assigned extremely low weights, effectively transferring the distribution of real samples to the minority class and reducing the impact of noisy data.

Algorithm 8 Transfer learning (FTL)

Input: Labeled majority class data \mathcal{D}_M and minority class data \mathcal{D}_m

Output: Synthetic minority samples \mathcal{D}_{syn}

- 1: Estimate class-wise means μ_c and shared variance Σ for \mathcal{D}_M
 - 2: Estimate class-wise means μ_m for \mathcal{D}_m
 - 3: Transfer Σ from majority to minority classes
 - 4: Generate synthetic samples: $x \sim \mathcal{N}(\mu_m, \Sigma)$
 - 5: **return** \mathcal{D}_{syn}
-

2.2.6 Meta-learning

Meta-learning, often referred to as the process of learning to learn, aims to create a model capable of effectively generalizing across a variety of tasks rather than concentrating on instances from a single task. This approach leverages acquired meta-knowledge to address the shortage of training data for individual tasks and has been investigated for adapting models to imbalanced data distributions. Meta-learning methods achieve this goal through bi-level optimization, which consists of both inner and outer loops. The inner loop focuses on task-specific learning by training a model to excel at individual tasks. Meanwhile, the outer loop emphasizes learning across tasks and updates the model based on its performance on the inner loop tasks. To enhance the model's generalization and adaptability to imbalanced data distributions, researchers propose to optimize the model from multiple perspectives, which includes weights, sampling method, model dynamics, data augmentation and class prototypes.

Weights. The majority of meta-learning-based methods [109,151–154] concentrate on re-weighting methods for imbalanced data. The pseudo-code is depicted in Algorithm 9. These studies highlight the limitations of infeasible manually designed weighting functions and the challenges associated with applying hyper-parameters in loss functions. Consequently, these methods treat the weights for training samples as learnable parameters and optimize them accordingly.

LRE [108] first uses a clean and unbiased dataset as the meta set (i.e., development set) to guide the training procedure. By performing validation at every training iteration, LRE dynamically assigns importance weights for each batch sample. Meta-Weight-Net [109] automatically learns an explicit loss-weight function, parameterized by a multi-layer perceptron (MLP) from data, resulting in more stable learning performance. Bayesian TAML [151] learns a set of class-specific scalars to weight the learning rate of class-specific gradients for each inner loop optimization step, thus giving more attention to minority classes. MLCW [152] analyzes the imbalance problem from a domain adaptation perspective. They assume that the class-conditioned distribution of training data and test data is not the same, especially for minority classes with insufficient training data.

Algorithm 9 Meta-weight learning

Input: Training set \mathcal{D}_{train} , validation set \mathcal{D}_{val} , model f , meta-learner g

Output: Trained model f , updated weight function g

- 1: **for** each mini-batch $\{(x_i, y_i)\} \subset \mathcal{D}_{train}$ **do**
 - 2: Compute weights: $w_i \leftarrow g(x_i, y_i)$
 - 3: Compute weighted loss: $\mathcal{L}_{train} = \sum w_i \cdot \ell(f(x_i), y_i)$
 - 4: Simulate gradient step: $f' \leftarrow f - \alpha \cdot \nabla_f \mathcal{L}_{train}$
 - 5: Sample batch $\{(x_j^{val}, y_j^{val})\} \subset \mathcal{D}_{val}$
 - 6: Compute validation loss: $\mathcal{L}_{val} = \sum \ell(f'(x_j^{val}), y_j^{val})$
 - 7: Update g : $g \leftarrow g - \beta \cdot \nabla_g \mathcal{L}_{val}$
 - 8: **end for**
 - 9: **return** f, g
-

Therefore, MLCW proposes to explicitly account for the differences between conditioned distributions and utilizes meta-learning to estimate the conditional weights for each training example. FSR [153] argued against the expensive second-order computation of nested optimizations in meta-learning and proposed a sample re-weighting method. By learning from history to build proxy reward data and employing feature sharing, FSR effectively reduces optimization costs.

Sampling method. Besides learning adaptive weights for training sets, meta-learning method is deployed to learn a sampling method to re-balance the training data. Balanced Meta-Softmax [101] introduces a balanced softmax method and argues that balance sampling for the mini-batch sampling process might lead to an over-balancing problem. In this scenario, minority classes dominate the optimization process, resulting in local minimums that favor these minority classes. To address this issue, Balanced Meta-Softmax introduces a meta sampler to explicitly learn the optimal sample rate for each class under the guidance of class-balanced meta set. Instead of co-optimized with the classification model, MESA [155] decouples the model-training and meta-training process and proposes a meta sampler that serves as an adaptive under-sampling solution embedded in the iterative ensemble training process. By incorporating the classification error distribution of ensemble training as its input, MESA achieves high compatibility and is applicable to various statistical and non-statistical classification models, such as decision trees, Naïve Bayes, and k-nearest neighbor classifiers.

Model dynamics. MetaModelNet [148] utilizes meta-learning to learn model dynamics, predicting how model parameters will change as more training data is progressively incorporated. As a result, MetaModelNet can directly enhance model parameters trained on minority instances by predicting the parameters when additional minority samples are included.

Data augmentation. MetaSAug [156] employs meta-learning for data augmentation. Inspired by ISDA [157], it calculates class-wise covariance matrices to extract semantic directions and generates synthesized minority instances along these semantically meaningful directions. However, the scarcity of data for minority classes results in inaccurate covariance matrices. To address this issue, MetaSAug incorporates meta-learning and learns the most meaningful class-wise covariance by minimizing the loss on a balanced meta set. Consequently, MetaSAug trains models on an augmentation set enriched with semantically augmented samples, effectively enhancing their performance.

Class prototypes. MPDT [139] focuses on calibrating empirical prototypes that are significantly biased due to scarce data in minority classes. MPDT first introduces meta-prototype contrastive learning to automatically learn the reliable representativeness of prototypes and predict the target distribution of balanced training data based on prototypes. Subsequently, additional features are sampled from the predicted distribution to further balance the overwhelming dominance of majority classes.

2.3 Training strategy

Training strategies, as shown in Fig. 4, exploit different stages of the

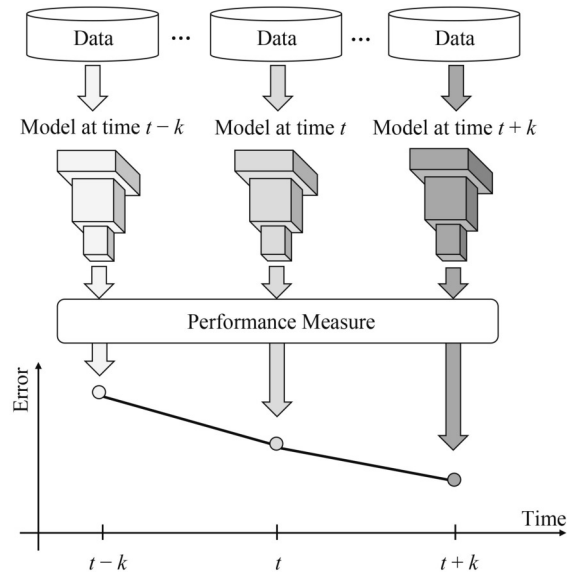


Fig. 4 Training strategy exploits different stages of the training pipeline to guide the model toward minimizing its inherent bias toward majority classes. Notably, these strategies often involve the separate use of different data subsets to enhance learning effectiveness

training pipeline to guide the model toward minimizing its inherent bias toward majority classes, which includes decoupled training, pretraining, curriculum learning, and post-hoc calibration.

2.3.1 Decoupling training

The concept of decoupling training was first formally evaluated by cRT [158], which separates the model learning procedure into two distinct processes: representation learning and classification, as shown in Algorithm 10. They systematically explored how imbalanced data affects each process individually and found that data imbalance might not be a significant issue in learning high-quality representations. Moreover, they demonstrated that the model could achieve robust classification performance by only retraining the classifier using class-balanced sampling. Similarly, LDAM [159] observes empirically that re-weighting and re-sampling are both inferior to the vanilla empirical risk minimization (ERM) algorithm (where all training examples have the same weight) before annealing the learning rate. Thus, LDAM defers the re-balancing or re-sampling training procedure to the classification stage. It initially trains the model using vanilla ERM and anneals the learning rate to obtain a better representation. Subsequently, re-weighting or re-sampling strategies are implemented with a smaller learning rate to adjust the decision boundary and fine-tune the features locally, leading to improved performance compared to one-stage training.

Following the work of cRT and LDAM, numerous methods have adopted the decoupling scheme and focused on enhancing feature representation through supervised and self-supervised learning approaches. KCL [125] leverages the superiority of supervised contrastive learning and further improve it by balancing the positive pairs, ensuring that there is an equal number of positive samples for each anchor sample. In contrast, BaCon [160] studies introduce a balance term in the supervised contrastive loss to re-weight the

negative pairs. During the classifier fine-tuning phase, they propose a generalized normalization classifier, which is more effective in preserving the length information of feature vectors and achieves more powerful classification ability. Instead of using one-hot labels, Logits Retargeting (LORT) [161] redistributes label probability mass by assigning a small portion of positive probability and spreading larger negative probabilities across other classes. This “label smooth value” encourages balanced logits magnitudes, reduces perturbations during training, and enhances class separability. Unlike conventional Label Smoothing, LORT sets a relatively large negative probability, which is crucial for improving convergence and discrimination.

Instead of utilizing the label information in representation learning, SSP [79] finds that self-supervised method could yields a great initialization that is more label-agnostic from the imbalanced dataset. rwSAM [162] systematically compares features pre-trained by self-supervised learning and supervised learning, both theoretically and empirically and finds that self-supervised learning can learn label-irrelevant but transferable features that aid in classifying minority classes better than supervised pre-training. To further reduce the influence of biased data distribution on self-supervised pre-training, rwSAM adapts sharpness-aware minimization to improve model generalization and calculates loss weight using kernel density estimation, which does not require training data labels. SCDL-GAN [24] employs the Wasserstein auto-encoder architecture to represent sample distributions previously. Guided by label information, SCDL-GAN can represent the distributive characteristics of all classes, avoiding biased minority representation and contributing to a well-performing GAN-based minority synthesis. SSD [163] proposes training an initial model under label supervision and self-supervision simultaneously using instance-balanced sampling. The pre-trained model is then used to provide informative soft labels as supervision signals during the classifier training procedure. By using self-

Algorithm 10 Decoupled training with balanced classifier fine-tuning

Input: Imbalanced training data \mathcal{D}_{imb} , balanced data \mathcal{D}_{bal} , backbone f , classifier g

Output: Final model $g \circ f$

```

1: for each epoch do
2:   for each batch  $(x, y) \sim \mathcal{D}_{\text{imb}}$  do
3:     Extract features  $z = f(x)$ 
4:     Compute loss  $\mathcal{L}_{\text{rep}} = \ell(g(z), y)$ 
5:     Update both  $f$  and  $g$  using gradient descent
6:   end for
7: end for
8: Freeze backbone  $f$ 
9: Re-initialize classifier  $g$ 
10: for each epoch do
11:   for each batch  $(x, y) \sim \mathcal{D}_{\text{bal}}$  do
12:     Extract features  $z = f(x)$  using the frozen backbone
13:     Compute loss  $\mathcal{L}_{\text{cls}} = \ell(g(z), y)$ 
14:     Update classifier  $g$  only
15:   end for
16: end for
17: return Final model  $g \circ f$ 

```

distillation, the inter-class information is integrated into the soft labels and transferred to the minority classes.

2.3.2 Fine-tuning pre-trained model

In contrast to the decoupling approach, which trains feature representations and classifiers on the same dataset, fine-tuning from a well pre-trained deep model, which is also called transfer learning, has been found to yield impressive performance for many downstream tasks. This method has been proven to improve model robustness and uncertainty estimates [164], while also benefiting imbalanced data representation. Specifically, the model is first pre-trained on a large, related dataset (source domain) to learn rich, unbiased feature representations. Subsequently, fine-tuning is conducted on the target dataset (target domain) to better adapt to the specific task.

FDM [165] initially examines the impact of imbalance in this paradigm, finding that majority classes have a more significant influence on feature learning, and achieving a more uniform distribution of sample numbers across classes is preferable. To counteract the effects of imbalance, FDM proposes a hierarchical feature learning approach that clusters objects into visually similar class groups and learns deep representations for these groups separately. This facilitates learning specific feature representations for individual classes. Similarly, taking into account the heavy computation and irrelevant features learned from large-scale datasets, DSTL [166] proposes to first measure the similarity between the target domain and source domain, selecting the most relevant samples from the source domain for pre-training the model. Subsequently, fine-tuning is performed on a more evenly-distributed subset of the target dataset, balancing the network's efforts among all categories and facilitating the transfer of learned features. It is worth noting that the source domain and target domain do not require a strict or explicit correlation. By utilizing a large-scale dataset (i.e., ImageNet) as the source domain, IBCC [167] can apply the learned basic texture knowledge to a target domain consisting of histopathological images, achieving impressive performance in handling imbalanced medical data.

2.3.3 Curriculum learning

Curriculum learning [168] is the strategy of learning from easy to hard, which can significantly improve the generalization of the deep model. To achieve this, the curriculum schedules focus on sampling, optimization, and re-weighting strategies. CASED [169] introduces a curriculum sampling method for the task of extreme imbalanced lung nodule detection in chest CT scans. Initially, CASED focuses on learning how to distinguish nodules from their immediate surroundings and gradually incorporates a greater proportion of difficult-to-classify global context until it reaches uniform sampling from the empirical data distribution. The curriculum sampling strategy effectively addresses the imbalance problem but also enhances the model's robustness against nodule annotation quality issues. CurriculumNet [170] designs the sampling curriculum by measuring the complexity of data using its distribution density in a feature space, and then ranks the complexity in an unsupervised

manner. Then the model capability is improved gradually by continuously adding the data with increasing complexity into the training process. BBN [67] proposes a bilateral-branch network structure and utilizes each branch to capture the imbalanced distribution and reverse distribution respectively. In the training procedure, a cumulative learning strategy is equipped to first learn the universal patterns of whole dataset and then pay attention to the minority class gradually.

Beside the sampling curriculum from imbalance to balance, DCL [171] further incorporates metric learning, which emphasizes learning a soft feature embedding to separate different samples in the feature space. With the expectation that the model will first learn suitable feature representations before classifying samples into correct labels, a loss scheduler is proposed to shift the learning objective from metric learning loss to classification loss.

SEDC [172] introduces a dual-curriculum learning strategy that performs feature re-weighting using adaptively updated weighting factors for each sample. By assigning different feature weight and sample weight, SEDC learns the optimal decision boundary by balancing the training contributions between majority and minority classes.

2.3.4 Posterior re-calibration

Posterior re-calibration refers to techniques that are designed from statistical perspectives to apply post-hoc adjustments to model outputs. The idea of calibrating model outputs is first leveraged in probability estimation where models are expected to generate reliable probability estimates of events (e.g., the probability of a sample falling into a certain class) [173]. Recently, significant strides have emerged in adapting posterior calibration techniques for discriminative tasks, specifically addressing challenges posed by imbalanced data. We categorize the existing model calibration methods into two main categories: Weight normalization and Logit re-balancing.

Classifier weight normalization. The weight normalization technique is initially introduced along with decoupling training in [158]. In their two-stage training approach, they observed that following joint training with instance-balanced sampling, the classifier's weight norms correlate with the cardinality of the classes. Inspired by this observation, they proposed the τ -normalization technique to re-scale the weight matrix in the second training phase or the inference phase. Following this work, DisAlign [174] was proposed to address the performance bottleneck in the biased decision boundary of the existing imbalanced methods. It extended the second training phase by introducing additional learnable parameters to the classifier and utilizing a KL divergence-based re-weighted loss. Gaussian Clouded Logit Adjustment [175] discovered that the learned representations within imbalanced datasets remain insufficient, particularly noting a severe compression of embedding spaces associated with the minority classes. To this end, it provided another solution to calibrate model output which combines normalization and random perturbation.

Logit re-balancing. This type of methods are based on one common probabilistic depiction of data imbalance: label prior shift. It

signifies a difference between the prior label distributions of the training set, denoted as $p_s(y)$, and those of the test set, $p_t(y)$ (see Table 1 for details). This difference exists while maintaining identical conditional probabilities of samples given the class label, expressed as $p_s(x|y) = p_t(x|y)$. Based on this depiction, we may consider applying an additive adjustment to the logit output during inference: $f_t(x) \leftarrow f_s(x) + \log p_t(y) - \log p_s(y)$, which forms the main idea of balanced softmax [101,176]. For the setting where labels in test set are uniformly distributed, [177] derived a similar logit adjustment formulation from another perspective. They draw inspiration from an alternative optimization objective, balanced error, which is the average of per-class error rates. Given that the target label distribution may also be imbalanced in practice, LADE [176] considered that Balanced Softmax does not involve the disentanglement of $p_s(y)$ during the training phase. It extended the original balanced softmax by utilizing regularized DV representation and importance sampling and proposed a novel loss function to further improve the performance. Ye et al. [178] and Tian et al. [179] extended the balanced softmax by adding a hyper-parameter as the scale of adjustment. UNO-IC [179] further combines this parameterized calibration method with temperature scaling and the Noisy-Or model, forming an effective multi-modal fusion algorithm.

Recently, Wang et al. [180] conducted a systematical analysis which succeeds in explaining the role of re-weighting and logit-adjustment in a unified manner, as well as some empirical results that are out of the scope of existing theories. As no additional parameters are involved, posterior recalibration manages to provide decent performance gains at a low computational cost, especially in the setting of imbalanced classification.

2.4 Ensemble

Ensemble methods combine multiple base models and leverage the strengths of diverse learners as depicted in Fig. 5. By aggregating predictions, they help mitigate the bias introduced by imbalanced data and improve the overall robustness and accuracy of the final model outputs.

2.4.1 Bagging-based ensemble

In the context of imbalanced data learning, Bagging, short for Bootstrap Aggregating, involves generating multiple balanced subsets of the imbalanced dataset by resampling techniques like

Table 1 Glossary of symbols

Symbol	Description
x	Input feature vector (sample)
y	Class label corresponding to input x
$p_s(y)$	Prior label distribution in the source (training) set
$p_t(y)$	Prior label distribution in the target (test) set
$p_s(x y)$	Conditional probability of x given label y in the source domain
$p_t(x y)$	Conditional probability of x given label y in the target domain
$f_s(x)$	Logit output for class y given input x




Fig. 5 Ensemble methods combine multiple base models and leverage the strengths of diverse learners

random undersampling and SMOTE, and training base classifiers on these subsets independently [181]. The final prediction is obtained through the aggregation of predictions made by each base classifier using majority voting or weighted voting as shown in Algorithm 11. Apart from rebalancing the class distribution, it aims to counterbalance the bias towards the majority class by introducing diversity among the base classifiers.

Over-Bagging [182] algorithm replaces the random sampling technique with the random oversampling technique in Bagging algorithm to deal with the imbalance of data. Data balance is achieved by over-sampling the minority samples. However, it may generate overfitting and redundant classifiers due to the high similarity between and within bags. This is because random oversampling can create duplicate instances. Hence, SMOTE-Bagging was proposed to alleviate the replication issue.

SMOTE-Bagging [182] overcomes this disadvantage by using SMOTE to generate synthetic minority class via interpolating between existing ones. It ensures that the synthetic instances are different from the original ones and promote the diversity of bags. By using a varying resampling rate in each iteration, SMOTE-Bagging also ensures that each bag is different from the others, which further increases the diversity of the ensemble. Specifically, in each bag, the minority samples consist of synthetic instances and resampled instances, where the resampling rate starts at 10% in the first iteration and increases to 100% in the last iteration. Therefore, SMOTE-Bagging can generate a more diverse ensemble than Over-Bagging, which can lead to better classification performance on the minority class.

Similarly, some papers suggest that the Bagging algorithm can balance data by undersampling. Thus, the Under-Bagging [182]

Algorithm 11 Bagging Ensemble

Input: Imbalanced dataset \mathcal{D} , number of models

K , base learner f_k

Output: Ensemble classifier $F(x)$

- 1: **for** $k = 1$ to K **do**
 - 2: Sample balanced subset \mathcal{D}_k from \mathcal{D} via class-wise resampling
 - 3: Train base learner f_k on \mathcal{D}_k
 - 4: **end for**
 - 5: Aggregate predictions: $F(x) = \text{Aggregate}(f_1(x), \dots, f_K(x))$
 - 6: **return** $F(x)$
-

algorithm is proposed. On contrary to Over-Bagging, it undersamples the majority samples in the process of iteratively generating training sets for base classifiers. Repeated sampling with the majority samples can also obtain a large difference in the base classifier, and compared with Over-Bagging, the number of samples per iteration will be less and the efficiency will be higher. However, in the process of under-sampling, it is easy to neglect most of the useful samples, resulting in inaccurate classification results.

Following this research line, numerous variants have been proposed from different viewpoints. For example, Roughly Balanced Bagging (RBBag) was proposed to generate more diverse bags [183]. That is, in each iteration of RBBag, the number of positive samples is kept fixed, while the number of negative samples satisfies negative binomial distribution. The diversity of the bags is achieved by varying the number of negative samples, while the rough balance of the bags ensures both the majority and minority classes are equally emphasized. From another perspective, Neighbourhood Balanced Bagging (NBBag) [184] was designed to modify the sampling probabilities of different instances in a way that focuses on the minority samples that are hard to learn while decreasing the probabilities of selecting samples from the majority class. This is achieved by analyzing the neighbourhood of samples, and taking into account the local minority class distribution. It is still an open question that how to resample instances for more accurate and diverse base classifiers and consequently better overall performance.

2.4.2 Boosting-based ensemble

Boosting-based ensemble methods aim to enhance the performance of classifiers on imbalanced data through iterative re-weighting of the training samples, to focus on the minority class, as shown in Algorithm 11. The most widely used boosting-based ensemble methods are AdaBoost [185], Gradient Boosting [186], and XGBoost [187]. Especially, in AdaBoost, each base classifier is trained on weighted samples. The weights of misclassified samples are increased with each iteration, while the weights of correctly classified samples are decreased. The final classifier is a weighted combination of base classifiers, where the weights are determined by their classification accuracy.

Based on the Adaboost and the SMOTE algorithm, the SMOTE-Boost [188] method is proposed. The SMOTE oversampling technique is introduced in each iteration. By adding synthetic minority class examples, each base classifier can pay more attention to the minority class. However, SMOTE-Boost fails to consider the

distribution of minority classes and latent noises when it generates synthetic examples, which can be blind and harmful to classification accuracy. To alleviate this problem, the MSMOTE-Boost [11] algorithm is proposed as improvement. Instead, it uses the MSMOTE algorithm to process the unbalanced data. In MSMOTE, by the distance between samples, the minority samples are divided into three groups: safety samples, boundary samples, and potential noise samples. For safety samples, the synthesis algorithm is the same as that of SMOTE; for boundary samples, the nearest sample is chosen; and for potential noise samples, no operation is performed. From another perspective, RAMOBoost [189] emphasizes to consider the learnability of minority samples. RAMOBoost employs an adaptive sampling strategy at each boosting iteration, to focus on difficult-to-learn minority samples. The sampling distribution is updated based on the pseudo-loss of the current hypothesis, which is a measure of the error rate of the hypothesis on the training data. The updated sampling distribution gives more weight to examples that are misclassified by the current hypothesis, which can help the algorithm to focus on difficult-to-learn examples.

Corresponding to the over-sampling boosting algorithms mentioned above is the boosting algorithms combined with under-sampling. The key idea of the RUS-Boost algorithm [190] is to select samples randomly from the majority samples by using the random undersampling technique (RUS) in the iteration process of the AdaBoost algorithm. Compared with the SMOTE-Boost, this algorithm has the advantages of simple implementation and short training time, but it is possible to remove many potentially useful samples in under-sampling. To avoid this problem, the EusBoost [56] algorithm is proposed. This algorithm uses the evolutionary under-sampling technique, which chooses the most representative samples in the majority samples to get more informative dataset, and it introduces the fitness function to ensure the difference of the base classifier. Another improvement to the RUSBoost algorithm is RHSBoost. The key idea of RHSBoost [191] is to combine the advantages of RUS and ROSE (Random Over-Sampling Examples) [192] sampling to generate more balanced and diverse subsets. Similar to SMOTE, at each boosting iteration, a weighted ROSE sampling algorithm is first used to oversample the minority class. Then, RUS is used to further balance the class distribution by randomly removing some instances from the majority class. Note that ROSE and SMOTE differ in the way they generate synthetic examples. SMOTE generates synthetic examples by interpolating, while ROSE samples from a multivariate kernel density estimate of the minority class, which can maintain the characteristics of the minority class and help avoid overfitting. Apart from those algorithms mentioned above, there are also many boosting-based methods proposed in recent years, such as LIUBoost [193] and HUSBoost [194], showing boosting-based methods are attractive and promising for imbalance data learning especially with undersampling technique.

2.4.3 Cost-sensitive ensemble

In scenarios such as fraud detection and medical diagnosis, the misclassification cost of different classes could vary significantly,

Algorithm 12 Boosting ensemble

Input: Dataset $\mathcal{D} = \{(x_i, y_i)\}$, rounds T , base learner f

Output: Weighted ensemble $\{(f_i, \alpha_i)\}$

- 1: Initialize weights $w_i \propto 1/N_{y_i}$
 - 2: **for** $t = 1$ to T **do**
 - 3: Normalize weights w_i
 - 4: Train f_t on \mathcal{D} with w_i
 - 5: Compute error ε_t ; set $\alpha_t = \log((1 - \varepsilon_t)/\varepsilon_t)$
 - 6: Update: $w_i \leftarrow w_i \cdot \exp(\alpha_t)$ if $f_t(x_i) \neq y_i$
 - 7: **end for**
 - 8: **return** $\{(f_t, \alpha_t)\}$
-

leading to potential risks or financial losses when traditional methods applied, which simply minimizes the average error rate or maximizes overall accuracy [195]. The key idea of cost-sensitive ensemble methods is to incorporate cost considerations into the construction and combination of base classifiers. To be more specific, these proposals usually differ in the way that they modify the weight update rule.

Among this family, AdaCost [196], CSB0, CSB1, CSB2 [197], AdaC1, AdaC2, and AdaC3 [198] are the most representative approaches. AdaCost introduces a cost adjustment function into the weight updating rule of AdaBoost. This adjustment function increases the weights of costly misclassifications more aggressively, and decreases the weights of costly correct classifications more conservatively. Theoretical analysis also provides a tighter upper bound on the cumulative training misclassification cost than AdaBoost.

CSB0 modifies the weight update rule by multiplying the weight update term by a dedicated cost function parameterized by both classifier prediction and groundtruth. CSB1 improves this vanilla version by taking the confidence of the base learner's prediction into account. CSB2, on the other hand, incorporates an additional factor into the weight update rule. This factor is a function of the misclassification cost and the prediction made by the base learner, i.e., the weight of current base learner in cost-sensitive scenario. By utilizing more informative factors, CSB2 is able to further minimize high cost errors.

Likewise the CSB family, AdaC1, AdaC2, and AdaC3 mainly differ in their weight update rules. In AdaC1, the cost function is introduced within the exponent part, multiplying by the weight of base classifier. But in AdaC2, the cost function is introduced outside the exponent part as another possible attempt. In AdaC3, this modification considers the idea of AdaC1 and AdaC2 at the same time. The weight update formula is modified by introducing the cost function both inside and outside the exponent part. Note that with the change of the cost function, the weight assignment rule for base classifiers is also changing. These details are omitted for simplicity.

2.4.4 Ensemble knowledge distillation

In the era of deep learning, numerous conventional machine learning models such as SVM, KNN, and DT have been replaced by deep learning models, which offer higher accuracy. However, these deep learning models come with a significant drawback: they require substantial computational resources and storage consumption. To address this issue, knowledge distillation was introduced as a compression technique [199]. Ensemble knowledge distillation extends the concept of traditional knowledge distillation [200]. Instead of transferring knowledge from a single teacher model, ensemble knowledge distillation transfers knowledge from an ensemble of teacher models. By leveraging the collective knowledge of the ensemble, the student model learns to mimic the collective behavior, resulting in improved robustness and accuracy.

Supervision signals from teacher models can include both high-level and intermediate-level information, namely, soft labels and intermediate layer representations, where the latter is generally

considered more informative. The diverse capabilities of teachers necessitate selecting the most suitable supervision signals or assigning importance weights based on their expertise.

For example, in [201], student model is simply distilled with the weighted average of teacher models, where weights are fixed as hyperparameters. And in [202], two strategies are investigated for utilizing labels from multiple teachers. The first one keeps the same with [201], while the second randomly selects a teacher model at the mini-batch level. For more flexible and finer grained weights, AMTML-KD [203] associates each teacher with a latent representation to adaptively learn instance-level weights, and gather intermediate-level hints from teachers using a multi-group based strategy. Concurrently, RL-KD also proposes to assign instance-level weights through a reinforcement learning process [204]. In RL-KD, the agent receives training instances and applies the teacher selection policy for distillation, where the reward is defined as the performance of the trained student model. Apart from weight assignment, CTKD [205] designs two dedicated teachers in a collaborative manner. Specifically, one teacher trained from scratch assists the student step by step using its temporary outputs, and the other pre-trained expert teacher guides the student to focus on a critical region. Inspired by the pre-training and fine-tuning paradigm, TMKD [206] first pre-trains the student, and further fine-tune the pre-trained model on downstream tasks. Both stages reduce the overfitting bias and transfer general knowledge to the student model. Collaborative Global-Local Structure Network (CGL-Net) [207] addresses the limitations of multi-expert networks for imbalanced classification, where inconsistent expert performance and underutilized data structure hinder generalization. CGL-Net decouples representation learning into two complementary stages: global structure learning, where a student network distills knowledge from diverse experts to achieve balanced category performance; and local structure learning, where augmented data guides the model to capture discriminative regional patterns of individual objects. By hierarchically integrating global and local knowledge, combined with knowledge distillation and multi-expert collaboration, CGL-Net produces robust and controllable representations. Note that ensemble knowledge distillation can be also combined with other technique such as parameter pruning and network quantization for more compact model. For instance, [208] proposes a novel framework that leverages both multi-teacher knowledge distillation and network quantization for learning low bit-width DNNs.

■ 3 Imbalance across diverse data formats

3.1 Image data

The rapid advancement of computer vision benefits greatly from the expansion of datasets. However, keeping data balanced becomes hard as data size grows in real-world image collection efforts. Consequently, data imbalance, especially in the form of imbalanced distribution, emerges as a widespread challenge across various computer vision tasks such as image classification, object detection and segmentation, leading to a decline in model performance. In the following sections, we will outline commonly used methods to mitigate the data imbalance problem within these tasks.

3.1.1 Image classification

Image classification stands as one of the foundational tasks in the realm of computer vision. Many image classification datasets as well as various domain-specific collections, exhibit imbalanced distribution patterns. In this scenario, models tend to show bias towards the majority classes, resulting in lower predicting accuracy in minority classes. Numerous studies have delved into resolving the issue of class imbalance in image classification, spanning across the categories we previously discussed. In terms of improving data qualities, resampling [101,158,171] and data augmentation [106,144,146] are frequently employed, mainly to balance the amount of data in different classes. [158] evaluates the performance of various sampling strategies for imbalanced recognition under both joint and decoupled learning schemes. Dynamic Curriculum Learning (DCL) [171] integrates the idea of curriculum learning into a sampling scheduler to train the model dynamically from imbalanced to balanced and from easy to hard. Major-to-Minor translation (M2m) [146] expands the oversampling approach by translating majority samples to minority ones using the perturbation-based optimization. Regarding model enhancements, a variety of techniques spanning model design [209–210], loss function [111–112,117], and training strategy [125,144,158] have been proposed to improve model performance when dealing with imbalanced data. GLMC [211] effectively improve the generalization of the backbone for imbalanced visual recognition by combining curated self-supervised and supervised loss functions. RoBal [210] succeed in tackle adversarial robustness under imbalanced distribution by leveraging a scale-invariant classifier and data re-balancing via both margin engineering at training stage and boundary adjustment during inference. Recently, studies like BALLAD [212] and PEL [213] employed pretrained contrastive vision-language models such as CLIP to achieve state-of-the-art performance on imbalanced datasets.

3.1.2 Object detection

Object detection involves simultaneously estimating the categories and locations of object instances within a given image. The imbalance problems in object detection can be grouped into four main categories: class imbalance, scale imbalance, spatial imbalance and objective imbalance [214]. Class imbalance refers to the unequal distribution of examples among classes, which can occur both between foreground classes and between foreground and background (negative) classes. To handle this problem, sampling methods are often utilized which re-weights bounding boxes according to their types (i.e., positive or negative) [215–216], the prediction hardness [99,217] or the IoU metric [218]. In addition to direct sampling methods, some studies introduce additional model architectures to improve the detection criterion [219–220], and some others address imbalance by producing artificial samples within the training dataset [221–222]. Scale imbalance arises from the unequal distribution of object scales and input bounding box sizes. Solutions to these issues often involve utilizing models and features with a pyramidal structure. For example, Single Shot Detector (SSD) [216] is able to exploits features at different levels through a feed-forward pyramidal

convolutional network. Its key approach to addressing scale imbalance is the use of multi-scale convolutional bounding box outputs attached to multiple feature maps at the top of the network. Feature Pyramid Networks (FPNs) [223] offer a generic solution for constructing feature pyramids within deep convolutional networks. The construction includes a bottom-up pathway, a top-down pathway, and lateral connections. By integrating these three components, FPNs empower the network to generate feature representations that are both semantically rich and spatially detailed across multiple scales.

3.1.3 Image segmentation

The goal of image segmentation is to divide an image into multiple regions and map them into classes or objects. Segmentation models often suffer from data imbalance, as real-world image segmentation datasets typically exhibit imbalanced label distributions, which can lead to a bias towards the majority classes. To mitigate this issue, some studies investigate into class-balancing loss functions. Sudre et al. [224] studied the use of different loss functions in deep learning frameworks for segmenting medical images, particularly when dealing with highly unbalanced data. The authors propose using the Generalized Dice overlap [225] as a robust loss function for such tasks. Bulò et al. [226] introduces a novel concept called loss max-pooling for handling imbalanced training data distributions in semantic image segmentation. This approach adaptively re-weights the contributions of each pixel based on their observed losses, targeting under-performing classification results often encountered for under-represented object classes. Taghanaki et al. [227] propose a new curriculum learning based loss function named Combo loss to tackle the imbalance problem in segmenting multiple organs from medical images. The Combo loss leverage Dice similarity coefficient to deter model parameters from being held at bad local minima and gradually learn better model parameters by penalizing for false positives/negatives using a cross entropy term. Zhong et al. [228] explored imbalanced semantic segmentation from the perspective of neural collapse. They introduced a regularization loss on feature centers called Center Collapse Regularizer to encourage the network to learn features closer to the appealing structure.

3.2 Text data

Data imbalance also causes problem for many NLP tasks [229–231], such as sentiment analysis, news classification, and text generation. It can cause the models trained on imbalanced data to be biased towards the majority classes and ignore the minority classes.

3.2.1 Sentiment analysis

Sentiment analysis is a task that aims to identify the emotional attitude of a text, such as positive, negative, and neutral. However, some classes may be more frequent than others in the data, such as positive reviews being more common than negative ones. This creates a problem of class imbalance, which can affect the performance of sentiment classifiers. For example, the sentiment of tweets collected from Twitter on majority of keywords related to social events, people, organization is highly imbalanced. To alleviate bias, re-sampling techniques such as SMOTE are widely applied in

[232–234] to over-sample the minority class. In addition to social psychology, there are several other factors that can greatly skew sentiment analysis results. One example comes from the work of Al-Azani and El-Alfy [235], where they examined the sentiment analysis of the Syria Tweets dataset. This dataset consisted of 1350 negative and 448 positive tweets, as reported by Mohammad et al. [236]. It is important to note that this dataset was highly imbalanced due to the ongoing regional instability in Syria. Compared with static sentiment analysis, real-time sentiment analysis is more challenging due to limited labeled data and sudden sentiment changes caused by real-world events (concept shift). To overcome this, Guerra et al. [237] acquire labels and handle concept drift using insights from self-report imbalances. That is, a preference for reporting positive feelings over negative feelings and a preference for reporting extreme feelings over average feelings, which helps generate labeled data on polarizing topics and informs a spike-based feature representation strategy. Apart from sentiment analysis of social network, another example can be product reviews. In product reviews, satisfied customers may be less likely to leave a review, while dissatisfied customers are more likely to voice their opinions, leading to a higher number of negative reviews. For instance, Obiedat et al. [238] proposed PSO-SVM with particle swarm optimization (PSO) to find the best feature weights and the best k value for each oversampling technique, and then applied SVM to classify the oversampled and weighted reviews.

3.2.2 News classification

News classification involves assigning a topic or category to a news article, such as politics, sports, and entertainment. The class imbalance phenomenon is the situation where the number of news articles that belong to different categories is not equal. Some topics may be more popular or frequent than others, such as sports news being more common than science news. For example, Dogra et al. [239] aim to extract the banking news and its most correlated news articles from the pool of financial news articles, and divide financial news articles into four classes: banking, non-banking, governmental, and global. The distribution of news articles amongst four classes are 9.16%, 70.54%, 4.21%, and 16.09% respectively, due to the different frequency and level of attention of news events. Sometimes, class imbalance phenomenon happens because the class of interest is quite rare than other classes, especially in binary classification problem. For instance, Rivera et al. [240] need to classify the news from a local newspaper based on whether they are about a traffic incident or not. However, the class of interest, which is “traffic accident”, only represents 10.30% of the instances, while the rest are other types of news. Another binary classification example can be fake news classification, in this scenario, the class imbalance refers to the number of fake news data is much smaller than the number of true news data. This happens because manual labeling of fake news data is time-consuming and difficult, while true news data is more abundant and accessible. For example, the BanFakeNews dataset [241] contains about 50,000 Bengali news articles, of which only 1,299 are labeled as fake and the rest are labeled as true. Hence, Keya et al. [242] proposed AugFake-BERT, which uses the pre-trained

BERT model to insert and substitute words for data augmentation, and train a smaller BERT model from scratch with the augmented dataset for embedding and classification. For multi-class classification problem, the imbalanced phenomenon is also evident, for example, Hasib et al. [243] study Bangla news classification with Bangla Newspaper Dataset [244]. The sample of 437,948 news articles are organized into the high imbalanced 32 categories, where the “bangladesh” category has 232,504 amounts of article and the percentage is 53.09% which is the highest and the lowest one are “askeditor” and “bs-events” categories which had only one article respectively.

3.2.3 Text generation

In addition to the aforementioned categories, the growing prevalence of large language models (LLMs) has brought the issue of data imbalance in text generation tasks to the forefront. To address this challenge, researchers have proposed several strategies. One prominent approach involves adjusting the composition of pre-training datasets to balance the proportions of data across tasks, languages, or domains. This method enhances the model’s learning efficiency, ensures a comprehensive representation of language features, and aligns the model’s capabilities with the specific requirements of the target tasks, thereby achieving balanced performance across diverse scenarios. Raffel et al. [245] demonstrated the effectiveness of two such strategies: mixing proportionally to dataset size and temperature-scaled mixing. In the former, datasets are weighted based on their sample sizes, while in the latter, a temperature parameter T is introduced to adjust task-level data distribution. This parameter increases the relative weight of smaller tasks, offering a more balanced mixture. Another promising strategy is leveraging fine-tuned models with limited samples as proxies for data selection. For example, Li et al. [246] proposed the “cherry model,” which operates through three fine-tuning stages: learning from limited experience, evaluation, and retraining based on self-guided insights. This model autonomously identifies and selects high-quality “cherry samples” from open-source datasets. Remarkably, on the Alpaca dataset, this method achieved superior performance using only approximately 5% of the data compared to models trained on the entire dataset. Custom dataset construction and task-specific fine-tuning also serve as effective methods to mitigate data imbalance. For instance, Li et al. [247] segmented raw data into question-answer pairs, translated and realigned them, and incorporated chain-of-thought (CoT) reasoning and tool-integrated reasoning (TIR) annotations. The resulting mathematical problem dataset underwent two stages of fine-tuning: one on diverse CoT-annotated datasets and another on synthetic TIR datasets, leading to significant improvements in mathematical reasoning tasks. Furthermore, iterative trial-and-error processes have proven effective in addressing data imbalance. An et al. [248] utilized multiple LLMs to collect inaccurate reasoning paths, which were subsequently corrected by GPT-4. This Correction-Centric Evolution strategy facilitated fine-tuning, enabling consistent performance improvements across various LLMs and tasks.

3.3 Graph data

Recently, graph data [249–252] has gained significant attention for

its ability to represent intricate structural relationships among diverse entities, facilitating applications across fields such as chemical analysis, social network modeling, and recommendation systems [253–256]. The extensive use of graph structures has driven considerable advancements in graph analytics, typically categorized into node-level, edge-level, and graph-level tasks based on the analytical objectives. Given these inherent objectives within graph structures, imbalance issues are prevalent across all analytical levels, manifesting in various forms. In the following, we will formulate these distinct imbalance challenges in alignment with each graph objective.

3.3.1 Node-level imbalance

Node class imbalance. Node class imbalance refers to the uneven distribution of class labels among nodes within a graph, where certain classes contain fewer labels compared to others. More importantly, the non-IID nature of graph data causes nodes to influence their neighbors through the message-passing paradigm in graph neural networks, further exacerbating the under-representation of minority classes. This interconnectedness intensifies the challenges posed by imbalanced data, as classes with fewer labels struggle to propagate meaningful signals across the graph, potentially diminishing model performance on these underrepresented classes.

To supplement the limited labels of minority classes, data generation has emerged as a widely adopted approach. In linear generative methods, for example, GraphSMOTE applies the SMOTE technique within the graph domain to synthesize new nodes for minority classes and establish edges among them. Other methods, such as GraphMixup [257], GraphENS [258], SNS [254], and GraphSHA [259], employ mixup techniques to create more challenging samples by interpolating nodes from different classes, thereby enhancing the generalization capacity of models. In contrast, deep generative methods such as ImGAGN [260] leverage GAN architectures, where a generator and discriminator work simultaneously to perform data generation and classification. Additionally, methods like GraphSR [261] and BIM [262] enhance minority class representation through self-training, expanding the labeled dataset by annotating unlabeled nodes within the original graph. This approach helps mitigate noise introduced by generative models, leading to improved and more robust performance.

Node degree imbalance. Node degrees in graphs often follow an imbalanced distribution, i.e., a significant fraction of nodes are minority nodes with a small degree. For instance, in social networks, a small number of users may have thousands of connections, while the majority maintain only a few. Similarly, in e-commerce recommendation systems, popular items frequently appear in user histories, whereas niche items receive minimal attention. This disparity, referred to as node degree imbalance, creates challenges in effectively exploring the information from low-degree minority nodes. Specifically, majority nodes with high degrees benefit from rich structural information and neighbors, allowing them to achieve better representations and stronger performance in downstream tasks. In contrast, “minority” nodes with low degrees have limited topological information, which constrains their effectiveness.

To mitigate the information scarcity in minority nodes, researchers have developed methods to transfer structural patterns from majority nodes to minority nodes. LTE4G [263] trains multiple expert models to handle different types of nodes, using knowledge distillation to transfer knowledge from majority-node experts to those focused on minority nodes. Similarly, Cold Brew leverages a strong teacher model with node-wise structural embeddings to guide a student model, enabling it to uncover missing latent neighborhoods for minority nodes. Tail-GNN [264] models variable connections between a target node and its neighbors and learns the neighborhood translation from majority nodes to enhance the structurally limited minority nodes. SAILOR proposed a structural augmentation method to enhance node homophily by adding pseudo-homophilic edges to minority nodes, allowing them to acquire more task-relevant information from neighboring nodes.

Node topology imbalance. Distinct from other imbalance issues that focus on disparities in node or degree quantity, node topology imbalance highlights the positional distribution of labeled nodes within the graph. This problem is rooted in the homophily assumption that influence boundaries created through label propagation should ideally align with true class boundaries. However, the actual placement of labeled nodes often shifts these influence boundaries away from true boundaries, leading to skewed model decision boundaries and degraded performance.

To address topology imbalance, ReNode [265] introduces a metric that quantifies this imbalance by measuring influence conflict. It then reweights labeled nodes based on their proximity to class boundaries, thereby improving the performance for nodes located near boundaries. Alternatively, PASTEL [266] proposes an anchor-based position encoding mechanism that directly optimizes information propagation paths, enhancing the connectivity among nodes within the same class to strengthen the flow of supervisory information.

3.3.2 Edge-level imbalance

Besides the node-level imbalances, real-world graphs also exhibit imbalanced edge distributions, which significantly impact various downstream tasks, including inductive recommendation and anomalous transaction detection on financial networks. Specifically, new users or items are continuously added to the system as novel inductive nodes by linking with the existing graph. However, these novel nodes, situated within a few-shot context, typically possess limited edges due to scarce information, posing challenges for effective node representation and accurate link prediction. To address this issue, a recent study [267] employs episodic learning to explore transferable knowledge and augment the representation of nodes with sparse connections. Additionally, imbalanced edge distribution prominently affects edge anomaly detection tasks, where anomalies such as abnormal transactions are modeled as minorities within the graph. This challenge is generally addressed by characterizing the normal edge distributions [267] with auxiliary information [268] and identifying deviations from these norms as indicators of anomalies [269].

3.3.3 Graph-level imbalance

Graph-level datasets, frequently employed in modeling molecular or

drug graphs, inherently exhibit an imbalanced distribution. This imbalance is primarily observed in datasets featuring naturally occurring distributions, such as those associated with rare diseases or specific chemical compounds. This scenario poses significant challenges for machine learning models, which may struggle to accurately predict properties or classifications for these infrequent but critical instances. To address these challenges, recent research [270] constructs a global graph structure that aggregates information among similar samples to enhance representation. This technique effectively boosts the predictive capabilities of models dealing with rare or infrequent graph types by pooling collective information from analogous instances. Additionally, another study [271] focuses on regression tasks and employs additional real, yet unlabeled data to compensate for minorities. By adopting an iterative training process, this model gradually annotates the unlabeled data and includes these data for model training, thereby enhancing model performance systematically over time.

■ 4 Evaluation metrics

Evaluating models for imbalanced datasets requires careful consideration of the metrics used to assess the performance. Traditional metrics like accuracy can not provide a reliable measure in this imbalanced contexts because they are biased toward the majority. In datasets where one class significantly outnumbers another, a model can achieve high accuracy simply by favoring the majority, often at the expense of poorly predicting the minority. For example, in a medical diagnosis setting where 95% of the samples are negative for a disease, a model that always predicts “negative” will achieve 95% accuracy, despite failing to identify any positive cases, which are often the most critical to detect. To counter this bias and accurately assess model effectiveness, it’s crucial to use evaluation metrics that accurately reflect performance across all classes, especially when misclassification costs are high. Below, we first describe the fundamental metrics for classification and then introduce the metrics for imbalance scenarios:

Confusion matrix is a fundamental tool that provides a detailed breakdown of classification results across the actual and predicted classes, and the prediction results are categories into 4 types:

- True Positives (TP): Correctly predicted positive observations.
- True Negatives (TN): Correctly predicted negative observations.
- False Positives (FP): Incorrectly predicted as positive.
- False Negatives (FN): Incorrectly predicted as negative.

Precision provides insights into the accuracy of positive predictions.

$$\text{Precision} = \frac{TP}{TP + FP}. \quad (1)$$

Recall, also known as **Sensitivity**, measures the proportion of actual positives that are correctly identified by the model.

$$\text{Recall} = \frac{TP}{TP + FN}. \quad (2)$$

It indicates how effective the model is at detecting positive instances from the data set.

Specificity measures the proportion of actual negatives that are correctly identified as such. In other words, it is the ability of the test to identify negative results.

$$\text{Specificity} = \frac{TN}{TN + FP}. \quad (3)$$

With these fundamental metrics, we can derive more advanced measures to enhance the evaluation of imbalanced datasets:

F1-score combines precision and recall metrics to balance their influence:

$$F1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}. \quad (4)$$

Precision and recall often have a trade-off relationship: improving precision typically reduces recall and vice versa. This occurs because increasing the threshold for classifying positives can decrease the number of false positives (increasing precision) but also miss more actual positives (decreasing recall). By combining these two metrics, the F1-score offers a single metric that balances this trade-off, providing a holistic view of the model performance.

In multi-class classification scenarios, the F1-score can be calculated in distinct ways, specifically through **micro** and **macro** averaging. Micro combines the outcomes of all classes to compute a single average metric, which is particularly useful in imbalanced datasets, as it emphasizes performance on minority classes. Conversely, macro calculates the metric separately for each class and then averages these results, ensuring that all classes are given equal importance regardless of their size in the dataset.

Balanced accuracy (BA) adjusts accuracy to reflect equitable performance across classes:

$$BA = \frac{\text{Sensitivity} + \text{Specificity}}{2}. \quad (5)$$

G-Mean is defined as the geometric mean of sensitivity and specificity, and is sensitive to the performance on both classes:

$$G\text{-Mean} = \sqrt{\text{Sensitivity} \times \text{Specificity}}. \quad (6)$$

Matthews correlation coefficient (MCC) considers all four quadrants of the confusion matrix and is highly effective for binary classification problems:

$$MCC = \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}. \quad (7)$$

Area Under the Receiver Operating Characteristic Curve (AUC-ROC) is a performance measurement for classification models at various threshold settings. It represents the probability that a classifier will rank a randomly chosen positive instance higher than a randomly chosen negative instance. The ROC curve plots the True Positive Rate (TPR, also known as recall or sensitivity) against the False Positive Rate (FPR) at different threshold levels.

$$AUC\text{-ROC} = \int_0^1 TPR(\tau) dFPR(\tau). \quad (8)$$

A model with perfect predictive accuracy would have an AUC-

ROC of 1, indicating it perfectly ranks all positive instances higher than negative ones.

These metrics provide detailed and varied perspectives for assessing the performance of models on imbalanced datasets, enabling the selection of the appropriate measures based on specific evaluation needs.

5 Experiments

To assess the effectiveness of these approaches, we evaluate 20 representative methods across seven benchmark datasets, which vary in both size and imbalance ratios. The dataset statistics are summarized in Table 2.

We begin by evaluating 12 conventional machine learning methods on five widely used imbalanced learning benchmarks. Specifically, we compare

- 3 over-sampling methods (SMOTE [5], BorderlineSMOTE [9], and ADASYN [10]);

- 4 under-sampling methods (TomekLinks [41], OneSided-Selection [41], Condensed Nearest Neighbors [42], and Edited Nearest Neighbors [43]);
- 2 hybrid sampling methods (SMOTEENN [62] and SMOTETomek [62]);
- 3 ensemble methods (Balanced Bagging [183], RUSBoost [190], and Balanced Random Forest [272]).

For a fair comparison, we use Random Forest as the base classifier, and evaluate performance using three metrics: F1 score, AUC, and Accuracy (ACC). Each experiment is repeated 10 times, and the average performance, standard deviation, and rank of each method are reported in Table 3.

Compared to the base model, all methods generally achieve improved F1-scores, indicating enhanced minority class performance. However, some methods reduce overall accuracy, suggesting diminished performance on the majority class. Among all methods, SMOTE and SMOTEENN deliver the most robust and consistent results, highlighting the effectiveness of synthetic data generation in compensating for minority class under-representation and improving classifier robustness. In contrast, while under-sampling methods can also enhance performance, they often discard a substantial portion of the majority class data, inevitably leading to information loss. It is important to note that no method consistently outperforms others across all datasets, highlighting the need to select appropriate strategies based on dataset-specific characteristics.

Besides the conventional machine learning methods, we evaluate deep learning-based methods on the widely used benchmarks MNIST and Fashion-MNIST. Following the setup in [1], we impose a 10:1 imbalance ratio to assess the models. The encoder is implemented as a convolutional neural network, followed by a MLP classifier.

Table 2 Statistic of datasets

Datasets	Imbalance Ratio	#Samples	#Features	#Classes
ecoli	8.6:1	336	7	2
optical_digits	9.1:1	5,620	64	2
pen_digits	9.4:1	10,992	16	2
car_eval_34	12:1	1,728	21	2
letter_img	26:1	20,000	16	2
MINIST	10:01	33,000	784	10
FashionMNIST	10:01	33,000	784	10

Table 3 Comparison of conventional machine learning methods, with the best-performing results highlighted in bold

Category	Dataset	ecoli			optical_digits			pen_digits			car_eval_34			letter_img			Rank
		F1/%	AUC/%	ACC/%	F1/%	AUC/%	ACC/%	F1/%	AUC/%	ACC/%	F1/%	AUC/%	ACC/%	F1/%	AUC/%	ACC/%	
Over-sampling	Base	59.31 ± 10.60	73.60 ± 5.31	92.77 ± 1.60	88.85 ± 1.97	90.00 ± 1.58	98.03 ± 0.31	96.92 ± 0.48	97.99 ± 0.43	99.61 ± 0.09	86.77 ± 4.50	91.27 ± 3.20	98.05 ± 0.65	94.90 ± 1.01	95.59 ± 1.06	99.64 ± 0.07	N.A.
	SMOTE	60.91 ± 8.47	79.10 ± 5.85	91.19 ± 1.90	94.23 ± 1.12	94.63 ± 1.01	98.91 ± 0.20	98.56 ± 0.53	98.71 ± 0.47	99.73 ± 0.10	89.88 ± 2.60	94.69 ± 2.51	98.44 ± 0.39	95.56 ± 1.06	96.14 ± 1.06	99.69 ± 0.07	4.47
	BorderlineSMOTE	60.73 ± 11.97	74.40 ± 8.33	91.21 ± 2.14	90.64 ± 1.84	91.66 ± 1.50	98.31 ± 0.31	97.84 ± 0.64	98.01 ± 0.57	99.59 ± 0.12	89.79 ± 2.07	94.79 ± 2.21	98.42 ± 0.30	95.63 ± 0.95	96.30 ± 0.88	99.69 ± 0.07	6.80
Down-sampling	ADASYN	56.56 ± 9.07	76.83 ± 6.16	90.00 ± 2.28	92.76 ± 1.35	93.49 ± 1.10	98.67 ± 0.24	98.28 ± 0.70	98.41 ± 0.68	99.68 ± 0.13	90.27 ± 3.00	94.95 ± 2.57	98.50 ± 0.45	95.60 ± 1.20	96.51 ± 1.10	99.69 ± 0.08	5.47
	TomekLinks	67.97 ± 10.09	81.12 ± 6.46	93.37 ± 2.08	89.43 ± 1.26	90.45 ± 1.03	98.12 ± 0.20	97.97 ± 0.48	98.04 ± 0.42	99.62 ± 0.09	87.56 ± 3.56	91.79 ± 2.91	98.17 ± 0.50	95.00 ± 1.06	95.62 ± 1.03	99.65 ± 0.07	7.67
	OneSidedSelection	67.16 ± 9.69	81.06 ± 7.37	93.27 ± 1.58	88.69 ± 1.18	89.85 ± 0.95	98.00 ± 0.19	97.91 ± 0.45	97.99 ± 0.43	99.61 ± 0.08	87.02 ± 3.16	91.41 ± 2.76	98.09 ± 0.43	95.09 ± 1.31	95.73 ± 1.21	99.66 ± 0.09	8.80
Hybrid sampling	CondensedNN	61.51 ± 10.17	78.58 ± 7.00	91.68 ± 2.27	91.56 ± 1.52	97.73 ± 0.43	98.23 ± 0.35	97.46 ± 1.66	98.81 ± 0.41	98.44 ± 0.36	88.40 ± 2.40	95.55 ± 2.77	98.13 ± 0.40	95.20 ± 0.99	97.98 ± 0.87	99.64 ± 0.07	6.53
	EditedNN	66.11 ± 7.64	85.14 ± 5.94	91.29 ± 2.61	89.76 ± 1.57	90.75 ± 1.32	98.17 ± 0.25	97.94 ± 0.52	98.01 ± 0.46	99.61 ± 0.10	87.56 ± 4.87	94.47 ± 2.93	97.82 ± 0.82	94.92 ± 1.08	95.75 ± 1.17	99.64 ± 0.07	8.07
	SMOTEENN	61.65 ± 6.33	86.09 ± 5.30	88.71 ± 2.66	94.28 ± 1.10	94.71 ± 0.97	98.93 ± 0.20	98.56 ± 0.66	98.72 ± 0.60	99.73 ± 0.12	88.14 ± 3.54	97.31 ± 2.34	98.00 ± 0.60	95.40 ± 1.19	96.48 ± 1.06	99.67 ± 0.08	4.47
Ensemble	SMOTETomek	60.59 ± 9.66	79.78 ± 7.73	90.99 ± 1.74	93.53 ± 1.13	94.02 ± 0.97	98.80 ± 0.20	98.46 ± 0.50	98.60 ± 0.45	99.71 ± 0.09	89.72 ± 2.83	94.45 ± 2.56	98.42 ± 0.42	95.64 ± 1.00	96.23 ± 1.00	99.70 ± 0.07	5.33
	BalancedBagging	61.46 ± 7.59	87.01 ± 5.99	88.22 ± 3.18	91.52 ± 1.30	96.42 ± 0.68	98.28 ± 0.28	97.59 ± 0.83	98.77 ± 0.50	99.54 ± 0.16	87.20 ± 2.90	98.54 ± 0.33	97.30 ± 0.61	95.37 ± 2.55	98.25 ± 0.55	98.23 ± 0.27	7.00
	RUSBoost	62.18 ± 4.28	84.82 ± 4.92	89.31 ± 3.22	91.62 ± 1.87	92.66 ± 1.13	98.89 ± 0.53	97.05 ± 1.94	98.95 ± 0.91	96.79 ± 0.46	87.11 ± 4.32	97.41 ± 2.20	97.34 ± 0.84	95.27 ± 4.57	96.08 ± 1.41	96.99 ± 0.73	7.73
	BalancedRandomForest	62.84 ± 4.28	89.00 ± 2.76	88.22 ± 2.75	92.86 ± 1.51	96.53 ± 0.86	98.58 ± 0.30	98.70 ± 0.59	99.04 ± 0.49	99.75 ± 0.11	88.09 ± 2.64	97.67 ± 0.50	95.92 ± 0.63	95.37 ± 1.81	98.53 ± 0.72	98.86 ± 0.17	4.60

According to our taxonomy, we compare 10 kinds of methods including:

- Linear generation (SMOTE [5], ADASYN [10]);
- Deep generation (GAN [31]);
- Cost-sensitive learning (Focal loss [99]);
- Metric learning (Triplet loss [122]);
- Supervised contrastive learning (SupCon [121]);
- Prototype learning (PSC [135]);
- Transfer learning (FTL [142]);
- Decoupling training (DT [158]);
- Posterior recalibration (τ -norm [158]).

As shown in Table 4, Transfer learning and GAN-based data generation achieve the best overall performance, as they can generate additional informative samples by leveraging the majority class distribution. However, despite their effectiveness, GANs entail a significantly higher computational cost due to the complexity of training both the generator and discriminator networks. As shown in Table 5, training a GAN takes approximately 15 min, while the conventional method SMOTE completes in just 2 s. Therefore, the choice of data generation strategy should balance performance and

computational efficiency, depending on the specific application and deployment constraints. Notably, decoupled training demonstrates competitive performance with minimal computational overhead. By simply fine-tuning the MLP classifier on a balanced dataset, it emphasizes the impact of training strategies on model effectiveness, offering a highly efficient alternative for imbalanced learning.

6 Resources

With the growing attention on imbalanced data, as shown in Table 6, researchers have developed an expanding array of tools to address the challenges it presents. Imbalanced-learn [273] provides a Python package incorporating a broad set of classical techniques and ensuring full compatibility with scikit-learn, thereby offering a convenient framework for assessing the effectiveness of various methods. Miriam [274] spearheaded a project summarizing many open-source contributions in this domain, including benchmark datasets, data generation and visualization tools, class overlap-based methods, data complexity measures, and software designed for learning from imbalanced data. However, some tools remain either insufficiently detailed or absent from the summary. For benchmarks, Qin et al. [275] introduced IGL-Bench, the first comprehensive benchmark dedicated to imbalanced graph learning. IGL-Bench

Table 4 The comparison of deep learning-based methods.

Dataset	MINIST			FashionMNIST			Rank
	F1/%	AUC/%	ACC/%	F1/%	AUC/%	ACC/%	
Base	90.53 ± 0.14	94.35 ± 0.07	90.63 ± 0.12	74.63 ± 0.10	87.85 ± 0.05	78.13 ± 0.09	N/A
SMOTE	91.18 ± 0.08	95.17 ± 0.05	91.30 ± 0.08	79.61 ± 0.41	88.87 ± 0.27	79.97 ± 0.49	8.33
ADASYN	91.95 ± 0.55	95.57 ± 0.30	92.03 ± 0.54	80.35 ± 0.28	89.04 ± 0.09	80.07 ± 0.17	5.67
GAN	93.46 ± 0.21	96.11 ± 0.23	92.92 ± 0.33	80.41 ± 0.52	89.48 ± 0.34	81.32 ± 0.41	2.50
Focal loss	91.35 ± 0.13	95.24 ± 0.07	91.43 ± 0.12	79.35 ± 0.17	88.81 ± 0.11	79.87 ± 0.21	8.33
Triplet loss	91.47 ± 0.29	95.56 ± 0.11	91.50 ± 0.12	79.06 ± 0.59	88.22 ± 0.65	79.80 ± 0.65	8.50
SupCon	92.96 ± 0.60	95.58 ± 0.11	92.70 ± 0.10	76.37 ± 0.22	88.89 ± 0.13	80.00 ± 0.22	6.17
PSC	92.43 ± 0.15	95.85 ± 0.10	92.53 ± 0.19	79.86 ± 0.66	90.22 ± 0.39	82.40 ± 0.71	4.33
FTL	94.08 ± 0.27	96.74 ± 0.15	94.13 ± 0.26	80.16 ± 0.74	90.37 ± 0.21	82.67 ± 0.38	1.50
DT	92.82 ± 0.70	96.04 ± 0.38	92.87 ± 0.68	80.40 ± 0.26	89.22 ± 0.18	80.60 ± 0.33	3.83
τ -norm	92.97 ± 0.11	96.13 ± 0.07	93.03 ± 0.12	78.09 ± 0.62	88.17 ± 0.49	78.70 ± 0.88	5.83

Table 5 The performance and data augmentation time comparison between representative methods.

Dataset	MINIST		FashionMNIST	
	F1/%	Aug_Time/s	F1/%	Aug_Time/s
Base	90.53 ± 0.14	N/A	74.63 ± 0.10	N/A
Focal loss	91.35 ± 0.13	N/A	79.35 ± 0.17	N/A
ADASYN	91.95 ± 0.55	2.12	80.35 ± 0.28	2.25
SMOTE	91.18 ± 0.08	1.27	79.61 ± 0.41	1.31
GAN	93.46 ± 0.21	920.41	80.41 ± 0.52	933.86

Table 6 Overview of open source tools

Project name	Content	Usage	Important information notes
imbalanced-learn [273]	Benchmark datasets, metrics, data generation, ensemble.	Python package offering classic methods.	Compatible with scikit-learn.
Open Source Contributions in Learning from Imbalanced and Overlapped Data [274]	Benchmark datasets, data generation, visualization, etc.	Study class imbalance and class overlap problems	Focus on unbalanced data distribution and feature space overlap
IGL-Bench [275]	First comprehensive benchmark for imbalanced graph learning.	Evaluate graph neural networks under class and topology imbalance	Integrates 16 datasets and 24 algorithms, used in real-world applications like social networks, e-commerce, and citation networks
Imbalanced-ensemble [276]	Python toolbox offering ensemble learning solutions for class imbalance	Provides 14 ensemble learning algorithms	Includes undersampling and SMOTE-based techniques, as well as cost-sensitive learning and visualization tools
Balance [277]	Python toolkit for adjusting sample distributions using weighting	Analyzes and adjusts sample data bias and evaluates its impact on model performance	Offers interactive visualization, supports ASMD metrics, low computational complexity, and flexible API design for quick integration

integrates 16 diverse datasets, such as Cora and PubMed, along with 24 state-of-the-art algorithms, addressing both class and topological imbalances in node-level and graph-level tasks. It ensures fair comparisons by employing unified data segmentation rules, standardized imbalance ratio definitions, and fixed splits for training, validation, and testing. Furthermore, it incorporates multi-dimensional evaluation metrics, such as accuracy, balanced accuracy, macro F1, and AUC, providing a rigorous framework for assessing algorithm performance. In terms of software, several notable tools were not included in Santos' project. One example is *imbalanced-ensemble*, developed by Liu et al. [276]. This open-source Python toolbox employs undersampling techniques in algorithms like *BalancedRandomForest* and *UnderBagging*, which randomly sample subsets from the majority class to train individual base classifiers. It also leverages SMOTE in methods such as *SmoteBoost* and *SmoteBagging* to synthesize minority samples. Additionally, the toolbox uses cost-sensitive approaches to adjust learning objectives, prioritizing minority class performance. The package includes visualization tools, such as Confusion Matrix Heatmaps, to facilitate intuitive performance analysis. Another noteworthy contribution is *Balance*, developed by Sarig et al. [277]. This Python toolkit adjusts and evaluates biased sample data using weighted approaches to align distributions with the target population. Its workflow includes diagnosing initial biases, generating weights via methods like propensity scores, and evaluating the impact of these weights on bias reduction and variance inflation. *Balance* also supports interactive visualization and diagnostic metrics, such as ASMD, offering researchers a streamlined and flexible API for bias adjustment and analysis.

■ 7 Challenges and future directions

7.1 Imbalance in open world setting

Existing researches presume that all classes present in the test set are also included in the training set. However, real-world scenarios frequently adhere to an open-world setting, where new classes may emerge exclusively in the test set. In such scenarios, the imbalance problem becomes exacerbated, as minority classes are not only

underrepresented but potentially unavailable during training, leaving models ill-equipped to detect them. Recent work [137] emphasizes the need for adaptive mechanisms that can identify and respond to unknown classes. However, two main technical bottlenecks remain: **(1) Identifying Unseen Classes:** Models must adapt to the dataset with no labels, which introduces zero-shot learning challenges under imbalance constraints. **(2) Overwhelming Confidence in Seen Classes:** Imbalanced training induces a strong bias toward seen (and especially majority) classes. This causes models to misclassify unfamiliar or low-density samples with high confidence into dominant categories, effectively suppressing the emergence of new minority classes. Future research should focus on developing uncertainty-aware learning frameworks that can detect and incorporate emerging minority classes. Promising directions include open-set recognition combined with continual learning, and the use of generative models to simulate novel class distributions during training.

7.2 Imbalanced multi-label problem

Imbalanced multi-label classification presents unique challenges, differing fundamentally from traditional multi-class scenarios due to the presence of multiple labels per instance.

This complexity is especially evident in real-world applications [278] such as image tagging or medical diagnosis, and presents unique challenges: (1) **Label Overlap Across Imbalance Spectrum:** A single instance may be associated with both majority and minority labels, making it difficult for the model to focus on underrepresented labels without being dominated by the gradients of frequent ones. This undermines conventional re-balancing strategies which treat labels independently. (2) **Semantic Entanglement Across Labels:** When labels correspond to semantically diverse or even conflicting concepts, multi-label samples introduce ambiguous supervision, making it harder for models to learn clear decision boundaries for each label.

Potential solutions could include the implementation of refined loss functions tailored to balance label importance, label-specific re-sampling strategies to adjust the representation of underrepresented

labels, or advanced regularization techniques to prevent model overfitting on prevalent labels. Moreover, leveraging insights from the existing researches on multi-label learning could provide innovative pathways for algorithmic enhancement, facilitating more accurate predictions across the diverse label distributions in real-world datasets.

7.3 Imbalanced regression problem

In contrast to classification, regression involves continuous labels and is essential in various real-world settings, such as financial forecasting, climate modeling, and healthcare analytics. However, the imbalance in regression tasks is not as extensively studied as in classification tasks [279,280]. Due to the continuous labels, two challenges are identified: (1) Undefined Minority Regions: Without discrete classes, it is difficult to explicitly define which target ranges are underrepresented. This lack of clear minority/majority boundaries complicates the design of data-balancing or cost-sensitive techniques. (2) Ordered Target Dependency: Regression labels are not independent categories but exist on an ordered scale. This structure demands methods that preserve relative distances while addressing imbalance. Future research should focus on structure-aware solutions that respect label continuity while correcting distribution skew. Promising directions include ranking-based objectives, distribution-aware sampling, and ordinal-regularized loss functions that emphasize rare but critical value regions.

7.4 Multi-modality imbalance problem

Multi-modality data is fundamental in applications such as healthcare diagnostics, autonomous driving, and multimedia analysis, where fusing diverse sources (e.g., text, image, audio, and sensor signals) enhances model robustness. However, imbalances inherent in multi-modal data [281] introduce unique challenges not present in unimodal settings. Key challenges include: (1) Cross-Modality Imbalance: Modalities often differ in data volume, quality, or annotation availability. For example, medical datasets may contain thousands of imaging records but only limited textual notes, causing models to over-rely on the dominant modality and neglect informative signals from the sparse ones. (2) Fusion-Level Imbalance Misalignment: When modalities are imbalanced, fusion layers may fail to integrate them effectively, especially if one modality dominates the representation space. This leads to suboptimal learning where underrepresented modalities are either ignored or incorrectly weighted during joint decision-making. Future research should develop imbalance-aware fusion strategies and cross-modal augmentation techniques that explicitly address representational disparities. Transfer learning and cross-modal distillation from rich to sparse modalities also hold promise for improving robustness under modality imbalance.

7.5 Large language models for imbalance problem

Large language models (LLMs) provide a powerful foundation for advancing imbalanced learning because of their generative capacity, cross-domain adaptability, and ability to leverage knowledge from large-scale pretraining corpora. Unlike conventional resampling or

augmentation methods, LLMs can generate class-conditional, semantically coherent, and label-faithful data, which makes them particularly suitable for addressing the scarcity of minority classes. First, LLM-based augmentation through prompt design and in-context learning can synthesize minority samples across different modalities such as tabular and textual data [282–284]. This enrichment of minority classes improves decision boundaries and strengthens feature correlations. Recent studies show that carefully designed prompts enable effective tabular data synthesis under imbalance [282–283] and that LLM-driven text augmentation significantly improves the performance of minority classes [284]. Besides, meta-learning with LLMs [285] opens a promising direction for imbalance-aware adaptation. By optimizing prompts or model parameters within a meta-learning framework, LLMs can learn to adapt quickly to minority classes from only a few examples. For instance, Vassilev et al. [285] propose a meta-learning approach that integrates diverse text representations and then combine it with threshold-moving, achieving robust improvements on highly imbalanced text classification tasks. This line of research highlights the potential of LLM-based meta-learning to complement traditional strategies with scalable and adaptive solutions for learning from minority classes.

8 Conclusion

This survey provides a thorough and up-to-date overview of imbalanced data learning, organizing the body of existing research into four distinct categories: data re-balancing, feature representation, training strategy, and ensemble learning. It delves into three critical real-world data formats, highlighting their characteristics and the unique challenges posed by imbalanced distributions.

Beyond synthesizing existing approaches, we emphasize imbalance learning as both a foundational component and a beneficiary of modern machine learning advancements. Crucially, the rise of generative AI and large-scale pretraining marks a transformative shift in the field. Generative models—such as GANs, VAEs, and diffusion models—enable the synthesis of high-quality, class-conditional samples, offering a principled solution to data scarcity in minority classes. At the same time, the era of large language models and multimodal foundation models introduces new challenges in distributional bias and representation imbalance across modalities. These trends call for renewed focus on scalable, modality-aware imbalance learning methods to ensure robustness, fairness, and adaptability in this new landscape.

By bridging classical techniques with modern paradigms, this survey strives to act as a guiding light for both researchers and practitioners, navigating them through the complex terrain of imbalanced data and encouraging advancements in machine learning. We hope it will act as a roadmap to inspire future research and practical innovations in building inclusive, data-efficient, and balanced intelligent systems.

Acknowledgements

This work was supported by the National Key R&D Program of China (Grant No. 2024YFA1014003), the National Natural Science

Foundation of China (Grant Nos. 92470121 and 62402016), and High-performance Computing Platform of Peking University, China.

■ Competing interests

The authors declare that they have no competing interests or financial conflicts to disclose.

■ Funding note

Open Access funding enabled and organized by CAUL and its Member Institutions.

■ Open Access

This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made.

The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

■ References

- [1] He H, Garcia E A. Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 2009, 21(9): 1263–1284
- [2] Krawczyk B. Learning from imbalanced data: open challenges and future directions. *Progress in Artificial Intelligence*, 2016, 5(4): 221–232
- [3] Guo H, Li Y, Shang J, Gu M, Huang Y, Gong B. Learning from class-imbalanced data: review of methods and applications. *Expert Systems with Applications*, 2017, 73: 220–239
- [4] Zhang Y, Kang B, Hooi B, Yan S, Feng J. Deep long-tailed learning: a survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023, 45(9): 10795–10816
- [5] Chawla N V, Bowyer K W, Hall L O, Kegelmeyer W P. SMOTE: synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 2002, 16: 321–357
- [6] Zhang H, Cisse M, Dauphin Y N, Lopez-Paz D. mixup: beyond empirical risk minimization. In: *Proceedings of the 6th International Conference on Learning Representations*. 2018
- [7] de la Calleja J, Fuentes O. A distance-based over-sampling method for learning from imbalanced data sets. In: *Proceedings of the 20th International Florida Artificial Intelligence Research Society Conference*. 2007, 634–635
- [8] García V, Sánchez J S, Mollineda R A. On the use of surrounding neighbors for synthetic over-sampling of the minority class. In: *Proceedings of the 8th Conference on Simulation, Modelling and Optimization*. 2008, 389–394
- [9] Han H, Wang W Y, Mao B H. Borderline-SMOTE: a new over-sampling method in imbalanced data sets learning. In: *Proceedings of the International Conference on Intelligent Computing*. 2005, 878–887
- [10] He H, Bai Y, Garcia E A, Li S. ADASYN: adaptive synthetic sampling approach for imbalanced learning. In: *Proceedings of 2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)*. 2008, 1322–1328
- [11] Hu S, Liang Y, Ma L, He Y. MSMOTE: improving classification performance when training data is imbalanced. In: *Proceedings of the 2nd International Workshop on Computer Science and Engineering*. 2009, 13–17
- [12] Kovács G. An empirical comparison and evaluation of minority oversampling techniques on a large number of imbalanced datasets. *Applied Soft Computing*, 2019, 83: 105662
- [13] Chou H P, Chang S C, Pan J Y, Wei W, Juan D C. Remix: rebalanced mixup. In: *Proceedings of the European Conference on Computer Vision*. 2020, 95–110
- [14] Kabra A, Chopra A, Puri N, Badjatiya P, Verma S, Gupta P, Krishnamurthy B. MixBoost: synthetic oversampling using boosted mixup for handling extreme imbalance. In: *Proceedings of 2020 IEEE International Conference on Data Mining (ICDM)*. 2020, 1082–1087
- [15] Galdran A, Carneiro G, González Ballester M A. Balanced-MixUp for highly imbalanced medical image classification. In: *Proceedings of the 24th International Conference on Medical Image Computing and Computer Assisted Intervention*. 2021, 323–333
- [16] Zhang S, Chen C, Zhang X, Peng S. Label-occurrence-balanced mixup for long-tailed recognition. In: *Proceedings of 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2022, 3224–3228
- [17] Mathew J, Luo M, Pang C K, Chan H L. Kernel-based SMOTE for SVM classification of imbalanced datasets. In: *Proceedings of the 41st Annual Conference of the IEEE Industrial Electronics Society*. 2015, 1127–1132
- [18] Tang B, He H. KernelADASYN: kernel based adaptive synthetic data generation for imbalanced learning. In: *Proceedings of 2015 IEEE Congress on Evolutionary Computation (CEC)*. 2015, 664–671
- [19] Wan Z, Zhang Y, He H. Variational autoencoder based synthetic data generation for imbalanced learning. In: *Proceedings of 2017 IEEE Symposium Series on Computational Intelligence (SSCI)*. 2017, 1–7
- [20] Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y. Generative adversarial networks. *Communications of the ACM*, 2020, 63(11): 139–144
- [21] Douzas G, Bacao F. Effective data generation for imbalanced learning using conditional generative adversarial networks. *Expert Systems with Applications*, 2018, 91: 464–471
- [22] Mirza M, Osindero S. Conditional generative adversarial nets. 2014, arXiv preprint arXiv: 1411.1784
- [23] Mariani G, Scheidegger F, Istrate R, Bekas C, Malossi C. BAGAN: data augmentation with balancing GAN. 2018, arXiv preprint arXiv: 1803.09655
- [24] Cai Z, Wang X, Zhou M, Xu J, Jing L. Supervised class distribution learning for GANs-based imbalanced classification. In: *Proceedings of 2019 IEEE International Conference on Data Mining (ICDM)*. 2019, 41–50
- [25] Ding H, Sun Y, Huang N, Shen Z, Wang Z, Iftekhar A, Cui X.

- RVGAN-TL: a generative adversarial networks and transfer learning-based hybrid approach for imbalanced data classification. *Information Sciences*, 2023, 629: 184–203
- [26] Tolstikhin I O, Bousquet O, Gelly S, Schölkopf B. Wasserstein auto-encoders. In: *Proceedings of the 6th International Conference on Learning Representations*. 2018
- [27] Ren J, Liu Y, Liu J. EWGAN: entropy-based Wasserstein GAN for imbalanced learning. In: *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*. 2019, 10011–10012
- [28] Arjovsky M, Chintala S, Bottou L. Wasserstein generative adversarial networks. In: *Proceedings of the 34th International Conference on Machine Learning*. 2017, 214–223
- [29] Li W, Chen J, Cao J, Ma C, Wang J, Cui X, Chen P. EID-GAN: generative adversarial nets for extremely imbalanced data augmentation. *IEEE Transactions on Industrial Informatics*, 2023, 19(3): 3208–3218
- [30] Sharma A, Singh P K, Chandra R. SMOTified-GAN for class imbalanced pattern classification problems. *IEEE Access*, 2022, 10: 30655–30665
- [31] Mullick S S, Datta S, Das S. Generative adversarial minority oversampling. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 2019, 1695–1704
- [32] Ho J, Jain A, Abbeel P. Denoising diffusion probabilistic models. In: *Proceedings of the 34th International Conference on Neural Information Processing Systems*. 2020, 574
- [33] Zhou W, Wang W, Bao J, Chen D, Chen D, Yuan L, Li H. Semantic image synthesis via diffusion models. 2022, arXiv preprint arXiv: 2207.00050
- [34] Oh H J, Jeong W K. DiffMix: diffusion model-based data synthesis for nuclei segmentation and classification in imbalanced pathology image datasets. In: *Proceedings of the 26th International Conference on Medical Image Computing and Computer Assisted Intervention*. 2023, 337–345
- [35] Liu Q, Yan F, Zhao X, Du Z, Guo H, Tang R, Tian F. Diffusion augmentation for sequential recommendation. In: *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 2023, 1576–1586
- [36] Xie J, Li W, Li X, Liu Z, Ong Y S, Loy C C. MosaicFusion: diffusion models as data augmenters for large vocabulary instance segmentation. *International Journal of Computer Vision*, 2025, 133(4): 1456–1475
- [37] Wang Z, Wei S, Huo X, Wang H. PoGDiff: product-of-Gaussians diffusion models for imbalanced text-to-image generation. 2025, arXiv preprint arXiv: 2502.08106
- [38] Yan D, Qi L, Hu V T, Yang M H, Tang M. Training class-imbalanced diffusion model via overlap optimization. 2024, arXiv preprint arXiv: 2402.10821
- [39] Han P, Ye C, Zhou J, Zhang J, Hong J, Li X. Latent-based diffusion model for long-tailed recognition. In: *Proceedings of 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. 2024, 2639–2648
- [40] Xu C, Yan J, Yang M, Deng C. Rethinking noise sampling in class-imbalanced diffusion models. *IEEE Transactions on Image Processing*, 2024, 33: 6298–6308
- [41] Kubat M, Matwin S. Addressing the curse of imbalanced training sets: one-sided selection. In: *Proceedings of the 14th International Conference on Machine Learning (ICML)*. 1997, 179–186
- [42] Hart P. The condensed nearest neighbor rule (Corresp.). *IEEE Transactions on Information Theory*, 1968, 14(3): 515–516
- [43] Tomek I. An experiment with the edited nearest-neighbor rule. *IEEE Transactions on Systems, Man, and Cybernetics*, 1976, SMC-6(6): 448–452
- [44] Laurikkala J. Improving identification of difficult small classes by balancing class distribution. In: *Proceedings of the 8th Conference on Artificial Intelligence in Medicine in Europe*. 2001, 63–66
- [45] Guo G, Wang H, Bell D, Bi Y, Greer K. KNN model-based approach in classification. In: *Proceedings of the OTM Confederated International Conferences on the Move to Meaningful Internet Systems 2003*. 2003, 986–996
- [46] Zhang J, Mani I. kNN approach to unbalanced data distributions: a case study involving information extraction. In: *Proceedings of Workshop on Learning from Imbalanced Datasets*. 2003, 1–7
- [47] Vuttipittayamongkol P, Elyan E. Neighbourhood-based undersampling approach for handling imbalanced and overlapped data. *Information Sciences*, 2020, 509: 47–70
- [48] Yen S J, Lee Y S. Cluster-based under-sampling approaches for imbalanced data distributions. *Expert Systems with Applications*, 2009, 36(3): 5718–5727
- [49] Zhang Y P, Zhang L N, Wang Y C. Cluster-based majority under-sampling approaches for class imbalance learning. In: *Proceedings of the 2nd IEEE International Conference on Information and Financial Engineering*. 2010, 400–404
- [50] Lin W C, Tsai C F, Hu Y H, Jhang J S. Clustering-based undersampling in class-imbalanced data. *Information Sciences*, 2017, 409–410: 17–26
- [51] Ng W W Y, Hu J, Yeung D S, Yin S, Roli F. Diversified sensitivity-based undersampling for imbalance classification problems. *IEEE Transactions on Cybernetics*, 2015, 45(11): 2402–2412
- [52] Ofek N, Rokach L, Stern R, Shabtai A. Fast-CBUS: a fast clustering-based undersampling method for addressing the class imbalance problem. *Neurocomputing*, 2017, 243: 88–102
- [53] García S, Herrera F. Evolutionary undersampling for classification with imbalanced datasets: proposals and taxonomy. *Evolutionary Computation*, 2009, 17(3): 275–306
- [54] Triguero I, Galar M, Vluymans S, Cornelis C, Bustince H, Herrera F, Saeys Y. Evolutionary undersampling for imbalanced big data classification. In: *Proceedings of 2015 IEEE Congress on Evolutionary Computation (CEC)*. 2015, 715–722
- [55] Le H L, Landa-Silva D, Galar M, Garcia S, Triguero I. EUSC: a clustering-based surrogate model to accelerate evolutionary undersampling in imbalanced classification. *Applied Soft Computing*, 2021, 101: 107033
- [56] Galar M, Fernández A, Barrenechea E, Herrera F. EUSBoost: enhancing ensembles for highly imbalanced data-sets by evolutionary undersampling. *Pattern Recognition*, 2013, 46(12): 3460–3471
- [57] Krawczyk B, Galar M, Jeleń Ł, Herrera F. Evolutionary undersampling boosting for imbalanced classification of breast cancer malignancy. *Applied Soft Computing*, 2016, 38: 714–726
- [58] Peng M, Zhang Q, Xing X, Gui T, Huang X, Jiang Y G, Ding K,

- Chen Z. Trainable undersampling for class-imbalance learning. In: Proceedings of the 33rd AAAI Conference on Artificial Intelligence. 2019, 4707–4714
- [59] Yan Y, Zhu Y, Liu R, Zhang Y, Zhang Y, Zhang L. Spatial distribution-based imbalanced undersampling. *IEEE Transactions on Knowledge and Data Engineering*, 2023, 35(6): 6376–6391
- [60] Hoyos-Osorio J, Alvarez-Meza A, Daza-Santacoloma G, Orozco-Gutierrez A, Castellanos-Dominguez G. Relevant information undersampling to support imbalanced data classification. *Neurocomputing*, 2021, 436: 136–146
- [61] Wang Z, Cao C, Zhu Y. Entropy and confidence-based undersampling boosting random forests for imbalanced problems. *IEEE Transactions on Neural Networks and Learning Systems*, 2020, 31(12): 5178–5191
- [62] Batista G E A P A, Prati R C, Monard M C. A study of the behavior of several methods for balancing machine learning training data. *ACM SIGKDD Explorations Newsletter*, 2004, 6(1): 20–29
- [63] Feng S, Zhao C, Fu P. A cluster-based hybrid sampling approach for imbalanced data classification. *Review of Scientific Instruments*, 2020, 91(5): 055101
- [64] Ramentol E, Caballero Y, Bello R, Herrera F. SMOTE-RSB*: a hybrid preprocessing approach based on oversampling and undersampling for high imbalanced data-sets using SMOTE and rough sets theory. *Knowledge and Information Systems*, 2012, 33(2): 245–265
- [65] Sáez J A, Luengo J, Stefanowski J, Herrera F. SMOTE-IPF: addressing the noisy and borderline examples problem in imbalanced classification by a re-sampling method with filtering. *Information Sciences*, 2015, 291: 184–203
- [66] Gao X, Ren B, Zhang H, Sun B, Li J, Xu J, He Y, Li K. An ensemble imbalanced classification method based on model dynamic selection driven by data partition hybrid sampling. *Expert Systems with Applications*, 2020, 160: 113660
- [67] Zhou B, Cui Q, Wei X S, Chen Z M. BBN: bilateral-branch network with cumulative learning for long-tailed visual recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2020, 9716–9725
- [68] Wang T, Li Y, Kang B, Li J, Liew J, Tang S, Hoi S, Feng J. The devil is in classification: a simple framework for long-tail instance segmentation. In: Proceedings of the 16th European Conference on Computer Vision. 2020, 728–744
- [69] Guo H, Wang S. Long-tailed multi-label visual recognition by collaborative training on uniform and re-balanced samplings. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2021, 15084–15093
- [70] Li Y, Wang T, Kang B, Tang S, Wang C, Li J, Feng J. Overcoming classifier imbalance for long-tail object detection with balanced group softmax. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2020, 10988–10997
- [71] Xiang L, Ding G, Han J. Learning from multiple experts: self-paced knowledge distillation for long-tailed classification. In: Proceedings of the 16th European Conference on Computer Vision. 2020, 247–263
- [72] Cai J, Wang Y, Hwang J N. ACE: ally complementary experts for solving long-tailed recognition in one-shot. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. 2021, 112–121
- [73] Cui J, Liu S, Tian Z, Zhong Z, Jia J. ResLT: residual learning for long-tailed recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023, 45(3): 3695–3706
- [74] Zhang Y, Hooi B, Hong L, Feng J. Self-supervised aggregation of diverse experts for test-agnostic long-tailed recognition. In: Proceedings of the 36th International Conference on Neural Information Processing Systems. 2022, 2470
- [75] Koziarski M. Potential Anchoring for imbalanced data classification. *Pattern Recognition*, 2021, 120: 108114
- [76] Xu Y, Li Y L, Li J, Lu C. Constructing balance from imbalance for long-tailed image recognition. In: Proceedings of the 17th European Conference on Computer Vision. 2022, 38–56
- [77] Zhu Y, Yan Y, Zhang Y, Zhang Y. EHSO: evolutionary hybrid sampling in overlapping scenarios for imbalanced learning. *Neurocomputing*, 2020, 417: 333–346
- [78] Pouyanfar S, Tao Y, Mohan A, Tian H, Kaseb A S, Gauen K, Dailey R, Aghajanzadeh S, Lu Y H, Chen S C, Shyu M L. Dynamic sampling in convolutional neural networks for imbalanced data classification. In: Proceedings of 2018 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR). 2018, 112–117
- [79] Yang Y, Xu Z. Rethinking the value of labels for improving class-imbalanced learning. In: Proceedings of the 34th International Conference on Neural Information Processing Systems. 2020, 1618
- [80] Wei C, Sohn K, Mellina C, Yuille A, Yang F. CREST: a class-rebalancing self-training framework for imbalanced semi-supervised learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2021, 10852–10861
- [81] Lai Z, Wang C, Cheung S C, Chuah C N. SaR: self-adaptive refinement on pseudo labels for multiclass-imbalanced semi-supervised learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). 2022, 4090–4099
- [82] Oh Y, Kim D J, Kweon I S. DASO: distribution-aware semantics-oriented pseudo-label for imbalanced semi-supervised learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2022, 9776–9786
- [83] Ertekin S, Huang J, Giles C L. Active learning for class imbalance problem. In: Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. 2007, 823–824
- [84] Qin J, Wang C, Zou Q, Sun Y, Chen B. Active learning with extreme learning machine for online imbalanced multiclass classification. *Knowledge-Based Systems*, 2021, 231: 107385
- [85] Zhu J, Hovy E. Active learning for word sense disambiguation with methods for addressing the class imbalance problem. In: Proceedings of 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL). 2007, 783–790
- [86] Aggarwal U, Popescu A, Hudelot C. Minority class oriented active

- learning for imbalanced datasets. In: Proceedings of the 25th International Conference on Pattern Recognition (ICPR). 2021, 9920–9927
- [87] Liu W, Zhang H, Ding Z, Liu Q, Zhu C. A comprehensive active learning method for multiclass imbalanced data streams with concept drift. *Knowledge-Based Systems*, 2021, 215: 106778
- [88] Dong H, Zhu B, Zhang J. A cost-sensitive active learning for imbalance data with uncertainty and diversity combination. In: Proceedings of the 12th International Conference on Machine Learning and Computing. 2020, 218–224
- [89] Jin Q, Yuan M, Wang H, Wang M, Song Z. Deep active learning models for imbalanced image classification. *Knowledge-Based Systems*, 2022, 257: 109817
- [90] Fu J H, Lee S L. Certainty-enhanced active learning for improving imbalanced data classification. In: Proceedings of the 11th International Conference on Data Mining Workshops. 2011, 405–412
- [91] Zhang C, Tavanapong W, Kijkul G, Wong J, de Groen P C, Oh J. Similarity-based active learning for image classification under class imbalance. In: Proceedings of 2018 IEEE International Conference on Data Mining (ICDM). 2018, 1422–1427
- [92] Turney P D. Cost-sensitive classification: empirical evaluation of a hybrid genetic decision tree induction algorithm. *Journal of Artificial Intelligence Research*, 1994, 2(1): 369–409
- [93] Sheng V S, Ling C X. Thresholding for making classifiers cost-sensitive. In: Proceedings of the 21st National Conference on Artificial Intelligence - Volume 1. 2006, 476–481
- [94] Domingos P. MetaCost: a general method for making classifiers cost-sensitive. In: Proceedings of the 5th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. 1999, 155–164
- [95] Zhou Z H, Liu X Y. Training cost-sensitive neural networks with methods addressing the class imbalance problem. *IEEE Transactions on Knowledge and Data Engineering*, 2006, 18(1): 63–77
- [96] Zhou Z H, Liu X Y. On multi-class cost-sensitive learning. *Computational Intelligence*, 2010, 26(3): 232–257
- [97] Castro C L, Braga A P. Novel cost-sensitive approach to improve the multilayer perceptron performance on imbalanced data. *IEEE Transactions on Neural Networks and Learning Systems*, 2013, 24(6): 888–899
- [98] Cui Y, Jia M, Lin T Y, Song Y, Belongie S. Class-balanced loss based on effective number of samples. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2019, 9260–9269
- [99] Lin T Y, Goyal P, Girshick R, He K, Dollár P. Focal loss for dense object detection. In: Proceedings of the IEEE International Conference on Computer Vision. 2017, 2999–3007
- [100] Park S, Lim J, Jeon Y, Choi J Y. Influence-balanced loss for imbalanced visual classification. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. 2021, 715–724
- [101] Ren J, Yu C, Sheng S, Ma X, Zhao H, Yi S, Li H. Balanced meta-softmax for long-tailed visual recognition. In: Proceedings of the 34th International Conference on Neural Information Processing Systems. 2020, 351
- [102] Cao P, Li B, Zhao D, Zaiane O. A novel cost sensitive neural network ensemble for multiclass imbalance data learning. In: Proceedings of 2013 International Joint Conference on Neural Networks (IJCNN). 2013, 1–8
- [103] Zhang C, Tan K C, Li H, Hong G S. A cost-sensitive deep belief network for imbalanced classification. *IEEE Transactions on Neural Networks and Learning Systems*, 2019, 30(1): 109–122
- [104] Chen C, Zheng S, Chen X, Dong E, Liu X, Liu H, Dou D. Generalized data weighting via class-level gradient manipulation. In: Proceedings of the 35th International Conference on Neural Information Processing Systems. 2021, 1080
- [105] Tan J, Wang C, Li B, Li Q, Ouyang W, Yin C, Yan J. Equalization loss for long-tailed object recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2020, 11659–11668
- [106] Tan J, Lu X, Zhang G, Yin C, Li Q. Equalization loss v2: a new gradient balance approach for long-tailed object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2021, 1685–1694
- [107] Guo D, Li Z, Zheng M, Zhao H, Zhou M, Zha H. Learning to reweight examples with optimal transport for imbalanced classification. In: Proceedings of the 36th International Conference on Neural Information Processing Systems. 2022, 1850
- [108] Ren M, Zeng W, Yang B, Urtasun R. Learning to reweight examples for robust deep learning. In: Proceedings of the 35th International Conference on Machine Learning. 2018, 4331–4340
- [109] Shu J, Xie Q, Yi L, Zhao Q, Zhou S, Xu Z, Meng D. Meta-weight-net: learning an explicit mapping for sample weighting. In: Proceedings of the 33rd International Conference on Neural Information Processing Systems. 2019, 172
- [110] Duan Y, Lu J, Feng J, Zhou J. Deep localized metric learning. *IEEE Transactions on Circuits and Systems for Video Technology*, 2018, 28(10): 2644–2656
- [111] Huang C, Li Y, Loy C C, Tang X. Learning deep representation for imbalanced classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2016, 5375–5384
- [112] Dong Q, Gong S, Zhu X. Class rectification hard mining for imbalanced deep learning. In: Proceedings of the IEEE International Conference on Computer Vision. 2017, 1869–1878
- [113] Wang N, Zhao X, Jiang Y, Gao Y, BNRist K. Iterative metric learning for imbalance data classification. In: Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI). 2018, 2805–2811
- [114] Gui X, Zhang J, Tang J, Xu H, Zou J, Fan S. A quadruplet deep metric learning model for imbalanced time-series fault diagnosis. *Knowledge-Based Systems*, 2022, 238: 107932
- [115] Yan M, Li N. Borderline-margin loss based deep metric learning framework for imbalanced data. *Applied Intelligence*, 2023, 53(2): 1487–1504
- [116] Cheng D, Gong Y, Zhou S, Wang J, Zheng N. Person re-identification by multi-channel parts-based CNN with improved triplet loss function. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2016, 1335–1344
- [117] Zhang X, Fang Z, Wen Y, Li Z, Qiao Y. Range loss for deep face

- recognition with long-tailed training data. In: Proceedings of the IEEE International Conference on Computer Vision. 2017, 5419–5428
- [118] Feng L, Wang H, Jin B, Li H, Xue M, Wang L. Learning a distance metric by balancing KL-divergence for imbalanced datasets. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 2019, 49(12): 2384–2395
- [119] Hayat M, Khan S, Zamir S W, Shen J, Shao L. Gaussian affinity for max-margin class imbalanced learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. 2019, 6468–6478
- [120] Gautheron L, Habrard A, Morvant E, Sebban M. Metric learning from imbalanced data with generalization guarantees. *Pattern Recognition Letters*, 2020, 133: 298–304
- [121] Khosla P, Teterwak P, Wang C, Sarna A, Tian Y, Isola P, Maschinot A, Liu C, Krishnan D. Supervised contrastive learning. In: Proceedings of the 34th International Conference on Neural Information Processing Systems. 2020, 1567
- [122] Weinberger K Q, Saul L K. Distance metric learning for large margin nearest neighbor classification. *The Journal of Machine Learning Research*, 2009, 10: 207–244
- [123] Sohn K. Improved deep metric learning with multi-class N-pair loss objective. In: Proceedings of the 30th International Conference on Neural Information Processing Systems. 2016, 1857–1865
- [124] Tian Y, Krishnan D, Isola P. Contrastive multiview coding. In: Proceedings of the 16th European Conference on Computer Vision. 2020, 776–794
- [125] Kang B, Li Y, Xie S, Yuan Z, Feng J. Exploring balanced feature spaces for representation learning. In: Proceedings of the 9th International Conference on Learning Representations. 2021
- [126] Zhang J, Zou J, Su Z, Tang J, Kang Y, Xu H, Liu Z, Fan S. A class-aware supervised contrastive learning framework for imbalanced fault diagnosis. *Knowledge-Based Systems*, 2022, 252: 109437
- [127] Jiang L, Shi S, Tian Z, Lai X, Liu S, Fu C W, Jia J. Guided point contrastive learning for semi-supervised point cloud semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. 2021, 6403–6412
- [128] Cui J, Zhong Z, Liu S, Yu B, Jia J. Parametric contrastive learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. 2021, 695–704
- [129] Zhu J, Wang Z, Chen J, Chen Y P P, Jiang Y G. Balanced contrastive learning for long-tailed visual recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2022, 6898–6907
- [130] Li T, Cao P, Yuan Y, Fan L, Yang Y, Feris R, Indyk P, Katabi D. Targeted supervised contrastive learning for long-tailed recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2022, 6908–6918
- [131] Gao X, Jia X, Liu J, Xue B, Huang Z, Fu S, Zhang G, Li K. An ensemble contrastive classification framework for imbalanced learning with sample-neighbors pair construction. *Knowledge-Based Systems*, 2022, 249: 109007
- [132] Gao X, Meng Z, Jia X, Liu J, Diao X, Xue B, Huang Z, Li K. An imbalanced binary classification method based on contrastive learning using multi-label confidence comparisons within sample-neighbors pair. *Neurocomputing*, 2023, 517: 148–164
- [133] Yang Z, Pan J, Yang Y, Shi X, Zhou H Y, Zhang Z, Bian C. ProCo: prototype-aware contrastive learning for long-tailed medical image classification. In: Proceedings of the 25th International Conference on Medical Image Computing and Computer Assisted Intervention. 2022, 173–182
- [134] Snell J, Swersky K, Zemel R. Prototypical networks for few-shot learning. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. 2017, 4080–4090
- [135] Wang P, Han K, Wei X S, Zhang L, Wang L. Contrastive learning based hybrid networks for long-tailed image classification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2021, 943–952
- [136] Hong Y, Zhang J, Sun Z, Yan K. ProAug: prototype-based augmentation for long-tailed image classification. In: Proceedings of 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). 2024, 3035–3039
- [137] Liu Z, Miao Z, Zhan X, Wang J, Gong B, Yu S X. Large-scale long-tailed recognition in an open world. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2019, 2532–2541
- [138] Zeng Z, Wang S, Xu N, Mao W. PAN: prototype-based adaptive network for robust cross-modal retrieval. In: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2021, 1125–1134
- [139] Fu S, Chu H, He X, Wang H, Yang Z, Hu H. Meta-prototype decoupled training for long-tailed learning. In: Proceedings of the 16th Asian Conference on Computer Vision. 2023, 252–268
- [140] Zhu L, Yang Y. Inflated episodic memory with region self-attention for long-tailed visual recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2020, 4343–4352
- [141] He T, Gong D, Tian Z, Shen C. Learning and memorizing representative prototypes for 3D point cloud semantic and instance segmentation. In: Proceedings of the 16th European Conference on Computer Vision. 2020, 564–580
- [142] Yin X, Yu X, Sohn K, Liu X, Chandraker M. Feature transfer learning for face recognition with under-represented data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2019, 5697–5706
- [143] Liu J, Sun Y, Han C, Dou Z, Li W. Deep representation learning on long-tailed data: a learnable embedding augmentation perspective. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2020, 2967–2976
- [144] Chu P, Bian X, Liu S, Ling H. Feature space augmentation for long-tailed data. In: Proceedings of the 16th European Conference on Computer Vision. 2020, 694–710
- [145] Wang J, Lukasiewicz T, Hu X, Cai J, Xu Z. RSG: a simple but effective module for learning imbalanced datasets. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2021, 3783–3792
- [146] Kim J, Jeong J, Shin J. *M2m*: imbalanced classification via major-to-minor translation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2020, 13893–13902

- [147] Hong Y, Zhang J, Sun Z, Yan K. SAFA: sample-adaptive feature augmentation for long-tailed image classification. In: Proceedings of the 17th European Conference on Computer Vision. 2022, 587–603
- [148] Wang Y X, Ramanan D, Hebert M. Learning to model the tail. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. 2017, 7032–7042
- [149] Liu B, Li H, Kang H, Hua G, Vasconcelos N. GistNet: a geometric structure transfer network for long-tailed recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. 2021, 8189–8198
- [150] Dai W, Yang Q, Xue G R, Yu Y. Boosting for transfer learning. In: Proceedings of the 24th International Conference on Machine Learning. 2007, 193–200
- [151] Lee H B, Lee H, Na D, Kim S, Park M, Yang E, Hwang S J. Learning to balance: Bayesian meta-learning for imbalanced and out-of-distribution tasks. In: Proceedings of the 8th International Conference on Learning Representations (ICLR). 2020
- [152] Jamal M A, Brown M, Yang M H, Wang L, Gong B. Rethinking class-balanced methods for long-tailed visual recognition from a domain adaptation perspective. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2020, 7607–7616
- [153] Zhang Z, Pfister T. Learning fast sample re-weighting without reward data. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. 2021, 705–714
- [154] Liu Y, Cao B, Fan J. Improving the accuracy of learning example weights for imbalance classification. In: Proceedings of the 10th International Conference on Learning Representations. 2022
- [155] Liu Z, Wei P, Jiang J, Cao W, Bian J, Chang Y. MESA: boost ensemble imbalanced learning with meta-sampler. In: Proceedings of the 34th International Conference on Neural Information Processing Systems. 2020, 1212
- [156] Li S, Gong K, Liu C H, Wang Y, Qiao F, Cheng X. MetaSAug: meta semantic augmentation for long-tailed visual recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2021, 5208–5217
- [157] Wang Y, Pan X, Song S, Zhang H, Wu C, Huang G. Implicit semantic data augmentation for deep networks. In: Proceedings of the 33rd International Conference on Neural Information Processing Systems. 2019, 1132
- [158] Kang B, Xie S, Rohrbach M, Yan Z, Gordo A, Feng J, Kalantidis Y. Decoupling representation and classifier for long-tailed recognition. In: Proceedings of the 8th International Conference on Learning Representations. 2020
- [159] Cao K, Wei C, Gaidon A, Arechiga N, Ma T. Learning imbalanced datasets with label-distribution-aware margin loss. In: Proceedings of the 33rd International Conference on Neural Information Processing Systems. 2019, 140
- [160] Kang N, Chang H, Ma B, Shan S. A comprehensive framework for long-tailed learning via pretraining and normalization. *IEEE Transactions on Neural Networks and Learning Systems*, 2024, 35(3): 3437–3449
- [161] Sun S, Lu H, Li J, Xie Y, Li T, Yang X, Zhang L, Yan J. Rethinking classifier re-training in long-tailed recognition: label over-smooth can balance. In: Proceedings of the 13th International Conference on Learning Representations. 2025
- [162] Liu H, HaoChen J Z, Gaidon A, Ma T. Self-supervised learning is more robust to dataset imbalance. In: Proceedings of the 10th International Conference on Learning Representations. 2022
- [163] Li T, Wang L, Wu G. Self supervision to distillation for long-tailed visual recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. 2021, 610–619
- [164] Hendrycks D, Lee K, Mazeika M. Using pre-training can improve model robustness and uncertainty. In: Proceedings of the 36th International Conference on Machine Learning. 2019, 2712–2721
- [165] Ouyang W, Wang X, Zhang C, Yang X. Factors in finetuning deep model for object detection with long-tail distribution. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2016, 864–873
- [166] Cui Y, Song Y, Sun C, Howard A, Belongie S. Large scale fine-grained categorization and domain-specific transfer learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2018, 4109–4118
- [167] Singh R, Ahmed T, Kumar A, Singh A K, Pandey A K, Singh S K. Imbalanced breast cancer classification using transfer learning. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 2021, 18(1): 83–93
- [168] Bengio Y, Louradour J, Collobert R, Weston J. Curriculum learning. In: Proceedings of the 26th Annual International Conference on Machine Learning. 2009, 41–48
- [169] Jesson A, Guizard N, Ghalehjegh S H, Goblot D, Soudan F, Chapados N. CASED: curriculum adaptive sampling for extreme data imbalance. In: Proceedings of the 20th International Conference on Medical Image Computing and Computer Assisted Intervention. 2017, 639–646
- [170] Guo S, Huang W, Zhang H, Zhuang C, Dong D, Scott M R, Huang D. CurriculumNet: weakly supervised learning from large-scale web images. In: Proceedings of the 15th European Conference on Computer Vision (ECCV). 2018, 139–154
- [171] Wang Y, Gan W, Yang J, Wu W, Yan J. Dynamic curriculum learning for imbalanced data classification. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. 2019, 5016–5025
- [172] Zhao R, Chen X, Chen Z, Li S. Diagnosing glaucoma on imbalanced data with self-ensemble dual-curriculum learning. *Medical Image Analysis*, 2022, 75: 102295
- [173] Wallace B C, Dahabreh I J. Class probability estimates are unreliable for imbalanced data (and how to fix them). In: Proceedings of the 12th International Conference on Data Mining. 2012, 695–704
- [174] Zhang S, Li Z, Yan S, He X, Sun J. Distribution alignment: a unified framework for long-tail visual recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2021, 2361–2370
- [175] Li M, Cheung Y M, Lu Y. Long-tailed visual recognition via Gaussian clouded logit adjustment. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2022, 6919–6928
- [176] Hong Y, Han S, Choi K, Seo S, Kim B, Chang B. Disentangling label distribution for long-tailed visual recognition. In: Proceedings of the

- IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2021, 6622–6632
- [177] Menon A K, Jayasumana S, Rawat A S, Jain H, Veit A, Kumar S. Long-tail learning via logit adjustment. In: Proceedings of the 9th International Conference on Learning Representations (ICLR). 2021
- [178] Ye H J, Chen H Y, Zhan D C, Chao W L. Identifying and compensating for feature deviation in imbalanced deep learning. 2020, arXiv preprint arXiv: 2001.01385
- [179] Tian J, Liu Y C, Glaser N, Hsu Y C, Kira Z. Posterior recalibration for imbalanced datasets. In: Proceedings of the 34th International Conference on Neural Information Processing Systems. 2020, 679
- [180] Wang Z, Xu Q, Yang Z, He Y, Cao X, Huang Q. A unified generalization analysis of re-weighting and logit-adjustment for imbalanced learning. In: Proceedings of the 37th International Conference on Neural Information Processing Systems. 2023, 2100
- [181] Breiman L. Bagging predictors. *Machine Learning*, 1996, 24(2): 123–140
- [182] Wang S, Yao X. Diversity analysis on imbalanced data sets by using ensemble models. In: Proceedings of 2009 IEEE Symposium on Computational Intelligence and Data Mining. 2009, 324–331
- [183] Hido S, Kashima H, Takahashi Y. Roughly balanced bagging for imbalanced data. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 2009, 2(5-6): 412–426
- [184] Błaszczyński J, Stefanowski J. Neighbourhood sampling in bagging for imbalanced data. *Neurocomputing*, 2015, 150: 529–542
- [185] Hastie T, Rosset S, Zhu J, Zou H. Multi-class AdaBoost. *Statistics and its Interface*, 2009, 2(3): 349–360
- [186] Friedman J H. Greedy function approximation: a gradient boosting machine. *The Annals of Statistics*, 2001, 29(5): 1189–1232
- [187] Chen T, Guestrin C. XGBoost: a scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. 2016, 785–794
- [188] Chawla N V, Lazarevic A, Hall L O, Bowyer K W. SMOTEBoost: improving prediction of the minority class in boosting. In: Proceedings of the 7th European Conference on Principles of Data Mining and Knowledge Discovery. 2003, 107–119
- [189] Chen S, He H, Garcia E A. RAMOBoost: ranked minority oversampling in boosting. *IEEE Transactions on Neural Networks*, 2010, 21(10): 1624–1642
- [190] Seiffert C, Khoshgoftaar T M, Van Hulse J, Napolitano A. RUSBoost: a hybrid approach to alleviating class imbalance. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 2010, 40(1): 185–197
- [191] Gong J, Kim H. RHSBoost: improving classification performance in imbalance data. *Computational Statistics & Data Analysis*, 2017, 111: 1–13
- [192] Lunardon N, Menardi G, Torelli N. ROSE: a package for binary imbalanced learning. *The R Journal*, 2014, 6(1): 79–89
- [193] Ahmed S, Rayhan F, Mahub A, Rafsan Jani M, Shatabda S, Farid D M. LIUBoost: locality informed under-boosting for imbalanced data classification. In: Abraham A, Dutta P, Mandal J K, Bhattacharya A, Dutta S, eds. *Emerging Technologies in Data Mining and Information Security: Proceedings of IEMIS 2018, Volume 2*. Singapore: Springer, 2019, 133–144
- [194] Popel M H, Hasib K M, Habib S A, Shah F M. A hybrid under-sampling method (HUSBoost) to classify imbalanced data. In: Proceedings of the 21st International Conference of Computer and Information Technology (ICCIT). 2018, 1–7
- [195] Ling C X, Sheng V S. Cost-sensitive learning and the class imbalance problem. In: Sammut C, ed. *Encyclopedia of Machine Learning*. New York: Springer, 2011, 231–235
- [196] Fan W, Stolfo S J, Zhang J, Chan P K. AdaCost: misclassification cost-sensitive boosting. In: Proceedings of the 16th International Conference on Machine Learning (ICML). 1999, 97–105
- [197] Ting K M. A comparative study of cost-sensitive boosting algorithms. In: Proceedings of the 17th International Conference on Machine Learning (ICML). 2000, 983–990
- [198] Sun Y, Kamel M S, Wong A K C, Wang Y. Cost-sensitive boosting for classification of imbalanced data. *Pattern Recognition*, 2007, 40(12): 3358–3378
- [199] Gou J, Yu B, Maybank S J, Tao D. Knowledge distillation: a survey. *International Journal of Computer Vision*, 2021, 129(6): 1789–1819
- [200] Asif U, Tang J, Harrer S. Ensemble knowledge distillation for learning improved and efficient networks. In: Proceedings of the 24th European Conference on Artificial Intelligence. 2020, 953–960
- [201] Chebotar Y, Waters A. Distilling knowledge from ensembles of neural networks for speech recognition. In: Proceedings of the 17th Annual Conference of the International Speech Communication Association. 2016, 3439–3443
- [202] Fukuda T, Suzuki M, Kurata G, Thomas S, Cui J, Ramabhadran B. Efficient knowledge distillation from an ensemble of teachers. In: Proceedings of the 18th Annual Conference of the International Speech Communication Association. 2017, 3697–3701
- [203] Liu Y, Zhang W, Wang J. Adaptive multi-teacher multi-level knowledge distillation. *Neurocomputing*, 2020, 415: 106–113
- [204] Yuan F, Shou L, Pei J, Lin W, Gong M, Fu Y, Jiang D. Reinforced multi-teacher selection for knowledge distillation. In: Proceedings of the 35th AAAI Conference on Artificial Intelligence. 2021, 14284–14291
- [205] Zhao H, Sun X, Dong J, Chen C, Dong Z. Highlight every step: knowledge distillation via collaborative teaching. *IEEE Transactions on Cybernetics*, 2022, 52(4): 2070–2081
- [206] Yang Z, Shou L, Gong M, Lin W, Jiang D. Model compression with two-stage multi-teacher knowledge distillation for Web question answering system. In: Proceedings of the 13th International Conference on Web Search and Data Mining. 2020, 690–698
- [207] Wu F, Liu Z, Zhang Z, Liu J, Wang L. Collaborative global-local structure network with knowledge distillation for imbalanced data classification. *IEEE Transactions on Circuits and Systems for Video Technology*, 2025, 35(3): 2450–2460
- [208] Pham C, Hoang T, Do T T. Collaborative multi-teacher knowledge distillation for learning low bit-width deep neural networks. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. 2023, 6424–6432
- [209] Wu T Y, Morgado P, Wang P, Ho C H, Vasconcelos N. Solving long-tailed recognition with deep realistic taxonomic classifier. In:

- Proceedings of the 16th European Conference on Computer Vision. 2020, 171–189
- [210] Wu T, Liu Z, Huang Q, Wang Y, Lin D. Adversarial robustness under long-tailed distribution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2021, 8655–8664
- [211] Du F, Yang P, Jia Q, Nan F, Chen X, Yang Y. Global and local mixture consistency cumulative learning for long-tailed visual recognitions. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2023, 15814–15823
- [212] Ma T, Geng S, Wang M, Shao J, Lu J, Li H, Gao P, Qiao Y. A simple long-tailed recognition baseline via vision-language model. 2021, arXiv preprint arXiv: 2111.14745
- [213] Shi J X, Wei T, Zhou Z, Han X Y, Shao J J, Li Y F. Parameter-efficient long-tailed recognition. 2023, arXiv preprint arXiv: 2309.10019
- [214] Oksuz K, Cam B C, Kalkan S, Akbas E. Imbalance problems in object detection: a review. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021, 43(10): 3388–3415
- [215] Girshick R, Donahue J, Darrell T, Malik J. Rich feature hierarchies for accurate object detection and semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2014, 580–587
- [216] Liu W, Anguelov D, Erhan D, Szegedy C, Reed S, Fu C Y, Berg A C. SSD: single shot MultiBox detector. In: Proceedings of the 14th European Conference on Computer Vision. 2016, 21–37
- [217] Li B, Liu Y, Wang X. Gradient harmonized single-stage detector. In: Proceedings of the 33rd AAAI Conference on Artificial Intelligence. 2019, 8577–8584
- [218] Cao Y, Chen K, Loy C C, Lin D. Prime sample attention in object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2020, 11580–11588
- [219] Chen J, Liu D, Luo B, Peng X, Xu T, Chen E. Residual objectness for imbalance reduction. *Pattern Recognition*, 2022, 130: 108781
- [220] Chen K, Li J, Lin W, See J, Wang J, Duan L, Chen Z, He C, Zou J. Towards accurate one-stage object detection with AP-loss. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2019, 5114–5122
- [221] Wang X, Shrivastava A, Gupta A. A-fast-RCNN: hard positive generation via adversary for object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2017, 3039–3048
- [222] Tripathi S, Chandra S, Agrawal A, Tyagi A, Rehg J M, Chari V. Learning to generate synthetic data via compositing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2019, 461–470
- [223] Lin T Y, Dollár P, Girshick R, He K, Hariharan B, Belongie S. Feature pyramid networks for object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2017, 936–944
- [224] Sudre C H, Li W, Vercauteren T, Ourselin S, Jorge Cardoso M. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In: Proceedings of the 3rd International Workshop on Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support. 2017, 240–248
- [225] Crum W R, Camara O, Hill D L G. Generalized overlap measures for evaluation and validation in medical image analysis. *IEEE Transactions on Medical Imaging*, 2006, 25(11): 1451–1461
- [226] Bulò S R, Neuhold G, Kontschieder P. Loss max-pooling for semantic image segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2017, 7082–7091
- [227] Taghanaki S A, Zheng Y, Zhou S K, Georgescu B, Sharma P, Xu D, Comaniciu D, Hamarneh G. Combo loss: handling input and output imbalance in multi-organ segmentation. *Computerized Medical Imaging and Graphics*, 2019, 75: 24–33
- [228] Zhong Z, Cui J, Yang Y, Wu X, Qi X, Zhang X, Jia J. Understanding imbalanced semantic segmentation through neural collapse. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2023, 19550–19559
- [229] Henning S, Beluch W, Fraser A, Friedrich A. A survey of methods for addressing class imbalance in deep-learning based natural language processing. In: Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics. 2023, 523–540
- [230] Zhao P, Zhang H, Yu Q, Wang Z, Geng Y, Fu F, Yang L, Zhang W, Jiang J, Cui B. Retrieval-augmented generation for AI-generated content: a survey. 2024, arXiv preprint arXiv: 2402.19473
- [231] Wang Z, Yu Q, Wei S, Li Z, Xiong F, Wang X, Niu S, Liang H, Zhang W. QAEncoder: towards aligned representation learning in question answering systems. In: Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). 2025, 4306–4332
- [232] Ghosh K, Banerjee A, Chatterjee S, Sen S. Imbalanced twitter sentiment analysis using minority oversampling. In: Proceedings of the 10th International Conference on Awareness Science and Technology (iCAST). 2019, 1–5
- [233] Flores A C, Icoy R I, Peña C F, Gorro K D. An evaluation of SVM and naive Bayes with SMOTE on sentiment analysis data set. In: Proceedings of 2018 International Conference on Engineering, Applied Sciences, and Technology (ICEAST). 2018, 1–4
- [234] Ardianto R, Rivanie T, Alkhalifi Y, Nugraha F S, Gata W. Sentiment analysis on e-sports for education curriculum using naive Bayes and support vector machine. *Jurnal Ilmu Komputer Dan Informatika*, 2020, 13(2): 109–122
- [235] Al-Azani S, El-Alfy E S M. Using word embedding and ensemble learning for highly imbalanced data sentiment analysis in short Arabic text. *Procedia Computer Science*, 2017, 109: 359–366
- [236] Mohammad S M, Salameh M, Kiritchenko S. How translation alters sentiment. *Journal of Artificial Intelligence Research*, 2016, 55: 95–130
- [237] Guerra P C, Meira W Jr, Cardie C. Sentiment analysis on evolving social streams: how self-report imbalances can help. In: Proceedings of the 7th ACM International Conference on Web Search and Data Mining. 2014, 443–452
- [238] Obiedat R, Qaddoura R, Al-Zoubi A M, Al-Qaisi L, Harfoushi O, Alrefai M, Faris H. Sentiment analysis of customers' reviews using a hybrid evolutionary SVM-based approach in an imbalanced data distribution. *IEEE Access*, 2022, 10: 22260–22273
- [239] Dogra V, Verma S, Verma K, Jhanjhi N Z, Ghosh U, Le D N. A

- comparative analysis of machine learning models for banking news extraction by multiclass classification with imbalanced datasets of financial news: challenges and solutions. *International Journal of Interactive Multimedia and Artificial Intelligence*, 2022, 7(3): 35–52
- [240] Rivera G, Florencia R, García V, Ruiz A, Sánchez-Solís J P. News classification for identifying traffic incident points in a Spanish-speaking country: a real-world case study of class imbalance learning. *Applied Sciences*, 2020, 10(18): 6253
- [241] Hossain M Z, Rahman M A, Islam M S, Kar S. BanFakenews: a dataset for detecting fake news in Bangla. In: *Proceedings of the 12th Language Resources and Evaluation Conference*. 2020, 2862–2871
- [242] Keya A J, Wadud M A H, Mridha M F, Alatiyyah M, Hamid M A. AugFake-BERT: handling imbalance through augmentation of fake news using BERT to enhance the performance of fake news classification. *Applied Sciences*, 2022, 12(17): 8398
- [243] Hasib K M, Towhid N A, Faruk K O, Al Mahmud J, Mridha M F. Strategies for enhancing the performance of news article classification in Bangla: handling imbalance and interpretation. *Engineering Applications of Artificial Intelligence*, 2023, 125: 106688
- [244] Nazi Z A. Bangla newspaper dataset. Kaggle, 2020
- [245] Raffel C, Shazeer N, Roberts A, Lee K, Narang S, Matena M, Zhou Y, Li W, Liu P J. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 2020, 21(1): 140
- [246] Li M, Zhang Y, Li Z, Chen J, Chen L, Cheng N, Wang J, Zhou T, Xiao J. From quantity to quality: boosting LLM performance with self-guided data selection for instruction tuning. In: *Proceedings of 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*. 2024, 7602–7635
- [247] Li J, Beeching E, Tunstall L, Lipkin B, Soletskyi R, Huang S, Rasul K, Yu L, Jiang A Q, Shen Z, Qin Z, Dong B, Zhou L, Fleureau Y, Lample G, Polu S. NuminaMath: the largest public dataset in AI4Maths with 860k pairs of competition math problems and solutions. 2024. Hugging Face Repository
- [248] An S, Ma Z, Lin Z, Zheng N, Lou J G, Chen W. Learning from mistakes makes LLM better reasoner. 2023, arXiv preprint arXiv: 2310.20689
- [249] Shao Y, Li H, Gu X, Yin H, Li Y, Miao X, Zhang W, Cui B, Chen L. Distributed graph neural network training: a survey. *ACM Computing Surveys*, 2024, 56(8): 191
- [250] Gao X, Zhang W, Yu J, Shao Y, Nguyen Q V H, Cui B, Yin H. Accelerating scalable graph neural network inference with node-adaptive propagation. In: *Proceedings of the 40th International Conference on Data Engineering*. 2024, 3042–3055
- [251] Gao X, Li Y, Chen T, Ye G, Zhang W, Yin H. Contrastive graph condensation: advancing data versatility through self-supervised learning. 2024, arXiv preprint arXiv: 2411.17063
- [252] Gao X, Ye G, Chen T, Zhang W, Yu J, Yin H. Rethinking and accelerating graph condensation: a training-free approach with class partition. In: *Proceedings of the ACM on Web Conference*. 2025, 4359–4373
- [253] Gao X, Yin H, Chen T, Ye G, Zhang W, Cui B. RobGC: towards robust graph condensation. *IEEE Transactions on Knowledge and Data Engineering*, 2025, 37(8): 4791–4804
- [254] Gao X, Zhang W, Chen T, Yu J, Nguyen H Q V, Yin H. Semantic-aware node synthesis for imbalanced heterogeneous information networks. In: *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 2023, 545–555
- [255] Gao X, Chen T, Zhang W, Li Y, Sun X, Yin H. Graph condensation for open-world graph learning. In: *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2024, 851–862
- [256] Gao X, Yu J, Chen T, Ye G, Zhang W, Yin H. Graph condensation: a survey. *IEEE Transactions on Knowledge and Data Engineering*, 2025, 37(4): 1819–1837
- [257] Wu L, Xia J, Gao Z, Lin H, Tan C, Li S Z. GraphMixup: improving class-imbalanced node classification by reinforcement mixup and self-supervised context prediction. In: *Proceedings of the Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. 2023, 519–535
- [258] Park J, Song J, Yang E. GraphENS: neighbor-aware ego network synthesis for class-imbalanced node classification. In: *Proceedings of the 10th International Conference on Learning Representations*. 2022
- [259] Li W Z, Wang C D, Xiong H, Lai J H. GraphSHA: synthesizing harder samples for class-imbalanced node classification. In: *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2023, 1328–1340
- [260] Qu L, Zhu H, Zheng R, Shi Y, Yin H. ImGAGN: imbalanced network embedding via generative adversarial graph networks. In: *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 2021, 1390–1398
- [261] Zhou M, Gong Z. GraphSR: a data augmentation algorithm for imbalanced node classification. In: *Proceedings of the 37th AAAI Conference on Artificial Intelligence*. 2023, 4954–4962
- [262] Zhang W, Gao X, Yang L, Cao M, Huang P, Shan J, Yin H, Cui B. BIM: improving graph neural networks with balanced influence maximization. In: *Proceedings of the 40th International Conference on Data Engineering (ICDE)*. 2024, 2931–2944
- [263] Yun S, Kim K, Yoon K, Park C. LTE4G: long-tail experts for graph neural networks. In: *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 2022, 2434–2443
- [264] Liu Z, Nguyen T K, Fang Y. Tail-GNN: tail-node graph neural networks. In: *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 2021, 1109–1119
- [265] Chen D, Lin Y, Zhao G, Ren X, Li P, Zhou J, Sun X. Topology-imbalance learning for semi-supervised node classification. In: *Proceedings of the 35th International Conference on Neural Information Processing Systems*. 2021, 2287
- [266] Sun Q, Li J, Yuan H, Fu X, Peng H, Ji C, Li Q, Yu P S. Position-aware structure learning for graph topology-imbalance by relieving under-reaching and over-squashing. In: *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 2022, 1848–1857
- [267] Zhu X, Luo P, Zhao Z, Xu T, Lizhiyu A, Yu Y, Li X, Chen E. Few-shot link prediction for event-based social networks via meta-learning. In: *Proceedings of the 28th International Conference on*

Database Systems for Advanced Applications. 2023, 31–41

[268] Liu J, Xia F, Feng X, Ren J, Liu H. Deep graph learning for anomalous citation detection. *IEEE Transactions on Neural Networks and Learning Systems*, 2022, 33(6): 2543–2557

[269] Li Z, Wang H, Zhang P, Hui P, Huang J, Liao J, Zhang J, Bu J. Live-streaming fraud detection: a heterogeneous graph neural network approach. In: *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 2021, 3670–3678

[270] Wang Y, Zhao Y, Shah N, Derr T. Imbalanced graph classification via graph-of-graph neural networks. In: *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 2022, 2067–2076

[271] Liu G, Zhao T, Inae E, Luo T, Jiang M. Semi-supervised graph imbalanced regression. In: *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2023, 1453–1465

[272] Khoshgoftaar T M, Golawala M, Van Hulse J. An empirical study of learning from imbalanced data using random forest. In: *Proceedings of the 19th IEEE International Conference on Tools with Artificial Intelligence (ICTAI 2007)*. 2007, 310–317

[273] Lemaître G, Nogueira F, Aridas C K. Imbalanced-learn: a Python toolbox to tackle the curse of imbalanced datasets in machine learning. *The Journal of Machine Learning Research*, 2017, 18(1): 559–563

[274] Miriam S. Open-source-imbalance-overlap. See github.com/miriamspantos/open-source-imbalance-overlap website, 2024

[275] Qin J, Yuan H, Sun Q, Xu L, Yuan J, Huang P, Wang Z, Fu X, Peng H, Li J, Yu P S. IGL-bench: establishing the comprehensive benchmark for imbalanced graph learning. In: *Proceedings of the 13th International Conference on Learning Representations*. 2025

[276] Liu Z, Kang J, Tong H, Chang Y. IMBENS: ensemble class-imbalanced learning in Python. 2021, arXiv preprint arXiv: 2111.12776

[277] Sarig T, Galili T, Eilat R. balance—a Python package for balancing biased data samples. 2023, arXiv preprint arXiv: 2307.06024

[278] Du G, Zhang J, Zhang N, Wu H, Wu P, Li S. Semi-supervised imbalanced multi-label classification with label propagation. *Pattern Recognition*, 2024, 150: 110358

[279] Gong Y, Mori G, Tung F. RankSim: ranking similarity regularization for deep imbalanced regression. In: *Proceedings of the 39th International Conference on Machine Learning*. 2022, 7634–7649

[280] Yang Y, Zha K, Chen Y C, Wang H, Katabi D. Delving into deep imbalanced regression. In: *Proceedings of the 38th International Conference on Machine Learning*. 2021, 11842–11851

[281] Qian Y, Zhang C, Zhang Y, Wen Q, Ye Y, Zhang C. Co-modality graph contrastive learning for imbalanced node classification. In: *Proceedings of the 36th International Conference on Neural Information Processing Systems*. 2022, 1154

[282] Kim J, Kim T, Choo J. EPIC: effective prompting for imbalanced-class data synthesis in tabular data classification via large language models. In: *Proceedings of the 38th International Conference on Neural Information Processing Systems*. 2024, 990

[283] Yang J Y, Park G, Kim J, Jang H, Yang E. Language-interfaced tabular oversampling via progressive imputation and self-authentication. In: *Proceedings of the 12th International Conference on Learning Representations*. 2024

[284] Zhao H, Chen H, Ruggles T A, Feng Y, Singh D, Yoon H J.

Improving text classification with large language model-based data augmentation. *Electronics*, 2024, 13(13): 2535

[285] Vassilev A, Jin H, Hasan M. Meta learning with language models: challenges and opportunities in the classification of imbalanced text. 2023, arXiv preprint arXiv: 2310.15019



Xinyi GAO is a Postdoctoral Researcher at the School of Electrical Engineering and Computer Science, The University of Queensland, Australia. He received his PhD in Computer Science from The University of Queensland, Australia following his BE and ME degrees in Information and Communications Engineering from Xi'an Jiaotong University, China. His research interests include data mining, graph representation learning, and multimedia signal processing.



Dongting XIE is an undergraduate student at the School of Information Science and Technology, Yunnan University, China. His research interests include data mining, large language models, and graph representation learning.



Yihang ZHANG received the BS degree in Mathematics and Computer Science from Peking University, China in 2024. He is currently pursuing the Master's degree in Data Science at Peking University, China. His research interests include machine learning and its applications in scientific discoveries.



Zhengren WANG received the BS degree in computer science from University of Electronic Science and Technology of China, China in 2023. He is currently a master student in Academy for Advanced Interdisciplinary Studies, Peking University, China advised by Prof. Wentao Zhang. His Research interests include data-centric AI, generative AI and etc. He published papers as the first author in ACL, WWW, IJCAI, etc.



Chong CHEN is a Senior Researcher and Team Leader at Huawei Cloud. He obtained his PhD and BE degrees in Computer Science and Technology from Tsinghua University, China. His research interests focus on large language model applications, including Search, RAG, NL2SQL, Recommendation, and related fields, with an emphasis on building scalable and intelligent information systems.



Conghui HE is a Young Leading Scientist at Shanghai AI Lab and an Adjunct Doctoral Supervisor at the School of AI, Shanghai Jiao Tong University, China. Recognized as a National-level Young Talent, he received his PhD from Tsinghua University, China and conducted research at Stanford University, USA and Imperial College London, UK. His work in data-centric AI and high-performance computing has led to 150+ papers in top-tier venues, 9,000+ Google Scholar citations, and open-source projects with 50,000+ GitHub stars.



Hongzhi YIN received a PhD degree in computer science from Peking University, China in 2014. He works as an ARC Future Fellow and Full Professor at The University of Queensland, Australia. He has made notable contributions to recommendation systems, graph learning, and decentralized and edge intelligence. He has published 300+ papers with an H-index of 75 and received ARC Future Fellowship 2021 and DECRA Award 2016.



Wentao ZHANG is an assistant professor in the Center for Machine Learning Research at Peking University, China. Wentao's research focuses on Graph ML, Data-centric ML, ML systems, and AI4Science. Wentao has published 60+ papers in the top DB, DM, and ML venues. Wentao is the contributor of several system projects, including Angel, SGL, and OpenBox.