Mathematical Biosciences and Engineering

https://www.aimspress.com/journal/mbe

MBE, 23(5): 1375–1401.
DOI: 10.3934/mbe.2026051
Received: 19 June 2025
Revised: 17 October 2025
Accepted: 12 November 2025
Published: 10 April 2026

Research article

Addressing domain shift via imbalance-aware domain adaptation in embryo development assessment

Lei Li[1,†,*], **Xinglin Zhang**[2,†], **Jun Liang**[2], **Mengqian Huang**[2] **and Tao Chen**[3,*]

[1] School of Artificial Intelligence, Beijing Institute of Technology, Beijing 100000, China
[2] Shanghai Medical Image Insights Intelligent Technology Co., Ltd., Shanghai 200032, China
[3] University of Waterloo, Waterloo, ON, Canada

† These authors contributed equally to the work.

* **Correspondence:** Email: lenny.lilei.cs@gmail.com, t66chen@uwaterloo.ca.

Abstract: Deep learning models in medical imaging face dual challenges: domain shift, where models perform poorly when deployed in settings different from their training environment, and class imbalance, where certain disease conditions are naturally underrepresented. We have presented imbalance-aware domain adaptation (IADA), a novel framework that simultaneously tackles both challenges through three key components: (1) adaptive feature learning with class-specific attention mechanisms, (2) balanced domain alignment with dynamic weighting, and (3) adaptive threshold optimization. Our theoretical analysis established convergence guarantees and complexity bounds. Through extensive experiments on embryo development assessment across four imaging modalities, IADA demonstrated significant improvements over existing methods, achieving up to 25.19% higher accuracy while maintaining balanced performance across classes. In challenging scenarios with low-quality imaging systems, IADA showed robust generalization with AUC improvements of up to 12.56%. These results demonstrate IADA's potential for developing reliable and equitable medical imaging systems for diverse clinical settings. The code is publicly available at https://github.com/yinghemedical/imbalance-aware_domain_adaptation.

Keywords: deep learning; domain shift; imbalance complexity; adversarial domain alignment; embryo development assessment

1. Introduction

Deep learning has revolutionized medical image analysis, demonstrating remarkable potential for automating complex diagnostic tasks [1]. However, two critical challenges emerge when deploying these systems in real-world clinical settings: domain shift and class imbalance. While these challenges have been studied independently [2, 3], their interaction, particularly in medical imaging contexts, remains understudied and poses significant barriers to clinical adoption.

Domain shift occurs when deep learning models trained on data from one medical context show degraded performance when deployed in different settings [4]. This challenge is particularly evident in embryo development assessment, where models trained on high-end clinical time-lapse imaging systems often perform poorly when applied to data from portable microscopes or smartphone-based systems [5]. As illustrated in Figure 1, the transition from high-end to low-end imaging systems introduces significant variations in image quality, complicating the model's ability to maintain consistent performance.

This challenge is compounded by inherent class imbalance in medical data [6]. In embryo assessment, the distribution of developmental stages varies naturally and differs significantly across imaging modalities [7]. For instance, our analysis reveals that blastocyst-stage embryos represent 28.9% of the samples in clinical time-lapse systems (ED4) but comprise 81.2% in portable microscope data (ED2). This disparity reflects both biological realities and differences in clinical protocols across settings [8].

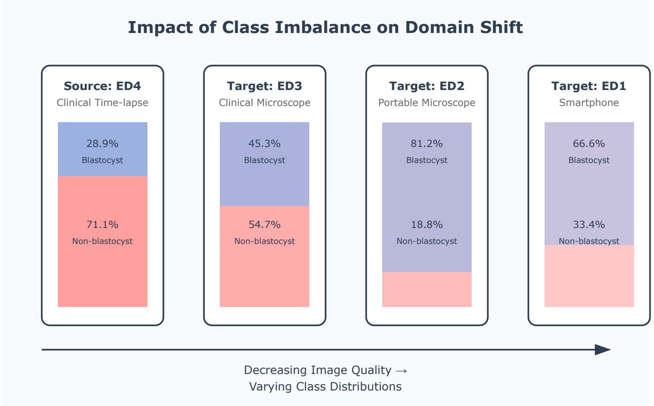

Figure 1. Illustration of the impact of class imbalance on domain shift in embryo classification. The figure shows the transition from high-end to low-end imaging systems and their varying class distributions. "Blastocyst" refers to a critical stage in early embryo development (typically 5–6 days after fertilization) characterized by the formation of a fluid-filled cavity, while "non-blastocyst" encompasses earlier developmental stages. The decreasing opacity from left to right represents the degradation in image quality across domains.

The intersection of these challenges creates a complex problem. Existing approaches have addressed them separately—domain adaptation techniques focus on bridging domain gaps without considering class distributions [9], while methods for handling imbalanced data typically assume consistent domain characteristics [10]. However, as Figure 1 demonstrates, in real-world medical applications, these challenges are inherently intertwined. The varying proportions of blastocyst versus non-blastocyst stages across different imaging modalities suggest that effective solutions must simultaneously address both domain shift and class imbalance.

In this paper, we present imbalance-aware domain adaptation (IADA), a novel framework that simultaneously tackles both challenges through three key innovations: (1) adaptive feature learning with class-specific attention mechanisms, (2) balanced domain alignment with dynamic weighting, and (3) adaptive threshold optimization. Through theoretical analysis and extensive experimentation on embryo development assessment across four distinct imaging modalities, we demonstrate IADA's effectiveness in maintaining robust performance despite varying image quality and class distributions. Our results show up to 25.19% improvement in accuracy over existing methods while maintaining balanced performance across classes, suggesting IADA's potential for developing more reliable and equitable medical imaging systems for diverse clinical settings [11].

2. Related work

2.1. Domain shift

Domain shift arises when there is a discrepancy between the distribution of training and testing datasets, leading to a decline in model performance. This phenomenon has been widely studied across multiple domains, including natural images, speech processing, and medical diagnostics [9, 12, 13].

Domain shift can occur due to various factors such as differences in imaging devices, acquisition protocols, environmental conditions, or population demographics. Techniques to combat domain shift include domain adaptation, domain generalization, and adversarial training. These approaches seek to align feature distributions across domains or make models robust to distributional variations [3, 9, 14, 15]. For instance, adversarial learning frameworks such as gradient reversal layer (GRL) have demonstrated success in reducing the discrepancy between source and target domains [3]. Other notable methods include maximum classifier discrepancy (MCD) [16], which leverages classifier disagreement to align domains, and conditional adversarial domain adaptation (CDAN) [15], which conditions domain alignment on discriminative features to improve transferability. Deep CORAL [14] aligns second-order statistics of source and target features, offering a computationally efficient alternative to adversarial methods. More recently, universal source-free domain adaptation (U-SFDA) [17] addresses scenarios where access to source data is unavailable during adaptation, making it particularly relevant for privacy-sensitive medical applications.

Domain shift is especially problematic in medical image analysis due to the heterogeneity of medical data. Variations in imaging modalities (e.g., MRI, CT, ultrasound), equipment manufacturers, and clinical protocols exacerbate the issue [18, 19]. Moreover, medical datasets are often small and suffer from class imbalances, making models more susceptible to overfitting and domain bias [1, 20]. In [4], Kanakasabapathy et al. developed adaptive adversarial neural networks to address domain shift in lossy and domain-shifted datasets of medical images. Their work highlights the importance of domain alignment techniques in preserving diagnostic accuracy. Other studies have explored self-supervised learning, meta-learning, and generative adversarial networks (GANs) to mitigate domain discrepancies [21–23]. Semi-supervised approaches such as FixMatch [24], which leverage consistency regularization and pseudo-labeling, offer promising directions for leveraging unlabeled target data under domain shift.

Embryo research stands out as a particularly critical area of medical imaging due to its implications for assisted reproductive technologies (ART). Accurate assessment of embryo quality is essential for optimizing implantation success rates and improving patient outcomes. However, this field faces unique challenges due to the sensitivity of embryo imaging, variations in clinical settings, and limited availability of annotated datasets [25–27].

The focus on embryo research is motivated by its potential to revolutionize ART practices through advanced imaging and machine learning. By addressing domain shift, researchers aim to create robust algorithms capable of generalizing across diverse patient populations and clinical environments [4,28]. This involves leveraging techniques such as multi-domain learning, synthetic data generation, and adaptive adversarial training to ensure reliable performance in diverse scenarios [29–31]. In this work, we follow the experimental protocol in the state-of-the-art research [4] as a foundational framework to study and address the challenges posed by domain shift in medical image analysis.

2.2. Learning with imbalanced data

Handling imbalanced data is a persistent challenge in medical applications, where minority classes often represent critical conditions requiring accurate predictions. Data-level methods aim to address imbalance by modifying the dataset. Oversampling techniques, such as the synthetic minority over-sampling technique (SMOTE) [6], generate synthetic minority class samples to balance the class distribution. Adaptive synthetic sampling (ADASYN) [10] extends the SMOTE by focusing on hard-to-learn instances, improving model performance on challenging samples. Alternatively, undersampling techniques, such as random undersampling and cluster centroids [32], reduce the majority class size but may risk losing valuable information.

Algorithmic solutions incorporate imbalance-handling mechanisms directly into the learning process. Cost-sensitive learning (CSL) [33] assigns higher misclassification costs to minority class instances, guiding the model toward improved performance on these classes. Ensemble methods, such as boosting and bagging, are also effective for imbalanced data by combining diverse models to mitigate bias [34]. Recent work by Cui et al. [35] introduced class-balanced loss, which re-weights loss based on the effective number of samples per class, offering a theoretically grounded approach to long-tailed recognition. FedFocal loss [36] extends focal loss to federated learning settings with non-IID data, improving performance on rare classes in decentralized medical systems.

Recent research has introduced advanced techniques for imbalanced data. Density-aware personalized training [37] decouples feature extraction and classification, using density-aware loss and learnable cost matrices. Federated learning for class-imbalanced medical image classification (FedIIC) [38] addresses imbalance in decentralized settings, employing contrastive learning and dynamic margins. Progressive class-center triplet loss [39] uses a two-stage training approach to separate and compact class distributions. Additionally, Yang et al. [40] challenged conventional assumptions by rethinking the role of labels in imbalanced learning, proposing that unlabeled data and representation learning can significantly improve performance even without extensive label rebalancing.

In medical diagnostics, methods like the SMOTE and cost-sensitive learning have been crucial for identifying rare diseases with low prevalence rates [41]. In medical imaging, approaches such as FedIIC and FedFocal loss have improved the classification of rare abnormalities in federated and non-IID settings [36,38].

3. Problem formulation

Medical domain adaptation presents a dual challenge: addressing both domain shift and class imbalance simultaneously. In healthcare applications, models trained on data from one medical context (source domain) often show degraded performance when deployed in different settings (target domain), particularly when dealing with imbalanced disease distributions. This challenge is compounded by various factors, including differences in patient populations, imaging equipment, clinical protocols, and institutional practices.

Let us first formalize the problem setup. We denote the source domain labeled dataset as $\mathcal{D}_s = \{(x_i^s, y_i^s)\}_{i=1}^{n_s}$, where:

- $x_i^s \in \mathcal{X} \subset \mathbb{R}^d$ represents d-dimensional input features (e.g., medical images, clinical measurements),
- $y_i^s \in \mathcal{Y} = \{1, ..., C\}$ denotes the corresponding class labels (e.g., disease diagnoses),
- n_s is the total number of source domain samples.

Similarly, we define the target domain dataset as $\mathcal{D}_t = \{x_j^t\}_{j=1}^{n_t}$, which is typically unlabeled in real-world medical scenarios. The fundamental domain shift can manifest in three distinct ways:

- **Covariate shift**: $P(X^s) \neq P(X^t)$, indicating differences in feature distributions,
- **Label shift**: $P(Y^s) \neq P(Y^t)$, reflecting varying disease prevalence,
- **Concept shift**: $P(Y|X^s) \neq P(Y|X^t)$, suggesting different feature-disease relationships.

4. Imbalance-aware domain adaptation

To address the challenges of medical domain shift while accounting for inherent class imbalances in medical data, we propose imbalance-aware domain adaptation (IADA). Our framework integrates three key innovations: adaptive feature learning that captures class-specific characteristics, balanced domain alignment that ensures fair representation across classes, and dynamic threshold optimization that adapts to varying class distributions. By designing an end-to-end training process, our method simultaneously optimizes both domain adaptation and class balance objectives. The overview of the proposed IADA is shown in Figure 2.

4.1. Imbalance-aware feature extraction

At the core of our approach lies a carefully designed feature extractor $F_\theta : \mathcal{X} \to \mathcal{Z}$ that addresses class imbalance through a novel attention mechanism. This feature extraction process consists of two main stages that work in concert to produce robust, class-aware representations:

1) **Base feature extraction**: The first stage begins with extracting class-specific features. For each input sample x_i, we process it through a sophisticated pipeline:

$$f_c(x_i) = h_c(g(x_i)) \tag{4.1}$$

Here, $g(\cdot)$ serves as a shared backbone network that captures general features, while $h_c(\cdot)$ represents class-specific adaptation layers that fine-tune these features for each class's unique characteristics.

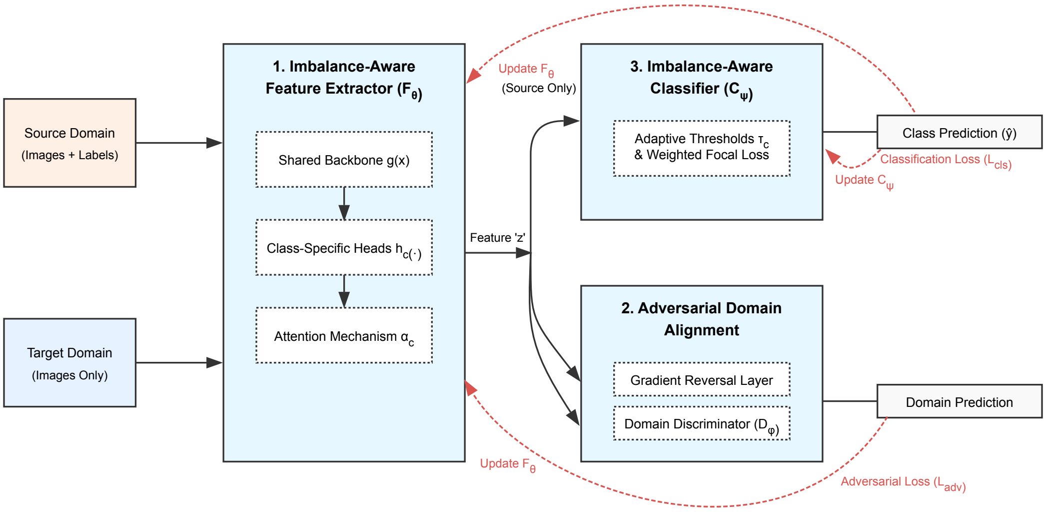

Figure 2. Overview of the proposed IADA. The model consists of three core components: (1) an imbalance-aware feature extractor, (2) an adversarial domain alignment module using a gradient reversal layer and domain discriminator, and (3) an imbalance-aware classifier employing adaptive thresholds and a weighted focal loss.

2) **Attention mechanism**: Building upon these base features, we introduce an attention mechanism that dynamically weights the importance of different class-specific features. The attention weights $\alpha_c(x_i)$ are computed through a softmax operation:

$$\alpha_c(x_i) = \frac{\exp(w_c^T g(x_i))}{\sum_{k=1}^{C} \exp(w_k^T g(x_i))} \tag{4.2}$$

In this formulation, w_c represents learnable attention vectors that help the model focus on the most relevant features for each class.

To combine these components into a final representation, we compute a weighted sum of the class-specific features:

$$z_i = F_\theta(x_i) = \sum_{c=1}^{C} \alpha_c(x_i) \cdot f_c(x_i) \tag{4.3}$$

This carefully crafted architecture provides several key benefits:

- It ensures features capture class-specific nuances through dedicated extractors,
- It gives minority classes fair representation through targeted attention mechanisms,
- It maintains flexibility to handle varying class distributions across different domains.

4.2. *Adversarial domain alignment*

To bridge the gap between source and target domains while maintaining class balance, we employ an advanced adversarial framework with several key innovations:

1) **Domain discriminator**: We implement a sophisticated discriminator $D_\phi : \mathcal{Z} \rightarrow [0, 1]$ as a multi-layer neural network. This discriminator incorporates three crucial components:

- A gradient reversal layer that enables adversarial training,
- Class-balanced batch sampling to ensure fair representation,
- Instance weighting that accounts for varying class frequencies.

2) **Adversarial loss**: To ensure effective domain alignment while respecting class balance, we define our adversarial objective as:

$$\mathcal{L}_{adv} = \mathbb{E}_{x^s \sim \mathcal{D}_s}[\omega(y^s) \log D_\phi(F_\theta(x^s))] + \mathbb{E}_{x^t \sim \mathcal{D}_t}[\log(1 - D_\phi(F_\theta(x^t)))] \tag{4.4}$$

To address class imbalance, we introduce class-specific weights $\omega(y^s)$, defined as:

$$\omega(y^s) = \frac{1}{C\pi^s_{y^s}} \tag{4.5}$$

where C is the number of classes and $\pi^s_{y^s}$ represents the proportion of samples in class y^s in the source domain.

4.3. Imbalance-aware classification

The final component of our framework is an adaptive classification module $C_\psi : \mathcal{Z} \to \mathcal{Y}$ that dynamically adjusts to class imbalance through three mechanisms:

1) **Adaptive thresholds**: To account for varying class distributions, we compute class-specific thresholds that adapt to class frequencies:

$$\tau_c = \beta \log(\frac{n^s_c}{\min_k n^s_k}) + \gamma \tag{4.6}$$

Here, β and γ are learnable parameters that allow the thresholds to adapt during training, with n^s_c representing the number of samples in class c in the source domain.

2) **Classification decision**: Using these adaptive thresholds, we make the final classification decision through:

$$\hat{y} = \arg\max_c \{C_\psi(z)_c - \tau_c\} \tag{4.7}$$

This formulation ensures that minority classes receive fair consideration by adjusting decision boundaries based on class frequencies.

3) **Confidence calibration**: To ensure reliable probability estimates, we incorporate temperature scaling:

$$p(y|z) = \text{softmax}(C_\psi(z)/T) \tag{4.8}$$

The temperature parameter T is learned during training to optimize probability calibration.

4.4. Training objective

To bring all components together into a cohesive framework, we formulate a comprehensive training objective:

$$\min_{\theta,\psi} \max_{\phi} \mathcal{L}_{cls}(\theta, \psi) - \lambda_{adv}\mathcal{L}_{adv}(\theta, \phi) + \lambda_{reg}\mathcal{R}(\theta, \psi) \tag{4.9}$$

This objective consists of three carefully designed components:

1) **Classification loss $\mathcal{L}_{cls}$**: We employ a weighted focal loss to address class imbalance:

$$\mathcal{L}_{cls} = -\frac{1}{n_s}\sum_{i=1}^{n_s}\omega(y_i^s)(1-p_{y_i^s})^{\gamma}\log(p_{y_i^s}) \tag{4.10}$$

The focal loss term $(1-p_{y_i^s})^{\gamma}$ helps focus training on hard examples, while class weights $\omega(y_i^s)$ balance the contribution of different classes.

2) **Regularization term $\mathcal{R}$**: To prevent overfitting and ensure robust feature learning, we combine multiple regularization strategies:

$$\mathcal{R} = \lambda_1\|\theta\|_2^2 + \lambda_2\mathcal{L}_{cons} + \lambda_3\mathcal{L}_{div} \tag{4.11}$$

This includes L2 regularization ($\|\theta\|_2^2$), consistency regularization ($\mathcal{L}_{cons}$) across augmented samples, and feature diversity promotion ($\mathcal{L}_{div}$).

3) **Adversarial term**: To ensure stable training, we implement a warming-up schedule for the adversarial weight:

$$\lambda_{adv} = \lambda_0 \cdot \min(1, \frac{t}{\tau}) \tag{4.12}$$

This gradual increase in adversarial strength, controlled by the current iteration t and warming-up period τ, allows the model to first learn good features before focusing on domain alignment.

5. Theoretical analysis

In this section, we provide theoretical analyses, including generalization, the convergence rate, and algorithmic complexity. First, let us establish some key assumptions and definitions:

Theorem 5.1 (Generalization bound with class imbalance)**.** *Let $h \in \mathcal{H}$ be a hypothesis with expected errors $\epsilon_s(h)$ and $\epsilon_t(h)$ on the source and target domains, respectively. For any $\delta > 0$, with probability at least $1-\delta$, the following bound holds:*

$$\epsilon_t(h) \leq \epsilon_s(h) + \sum_{i=1}^{C}|\pi_i^s - \pi_i^t| + \sum_{i=1}^{C}\min(\pi_i^s, \pi_i^t)d_i(\mathcal{H}) + \lambda \tag{5.1}$$

where λ represents the combined error of the ideal joint hypothesis.

Proof. The proof follows a structured approach through the decomposition and bounding of error terms. We begin by expressing the target error using class-conditional distributions: $\epsilon_t(h) = \sum_{i=1}^{C}\pi_i^t\epsilon_{t,i}(h)$, where $\epsilon_{t,i}(h)$ represents the error for class i in the target domain. Next, we establish that the difference between source and target errors for each class is bounded by the domain discrepancy: $|\epsilon_{t,i}(h) - \epsilon_{s,i}(h)| \leq d_i(\mathcal{H})$.

Applying the triangle inequality and leveraging the class proportions, we can derive:

$$\begin{aligned} |\epsilon_t(h) - \epsilon_s(h)| &\leq \sum_{i=1}^{C}|\pi_i^t\epsilon_{t,i}(h) - \pi_i^s\epsilon_{s,i}(h)| \\ &\leq \sum_{i=1}^{C}|\pi_i^t - \pi_i^s|\epsilon_{t,i}(h) + \sum_{i=1}^{C}\min(\pi_i^s, \pi_i^t)d_i(\mathcal{H}) \end{aligned} \tag{5.2}$$

The final bound is obtained by incorporating the ideal joint hypothesis error λ. □

Remark 5.2. *Theorem 5.1 offers several key insights into domain adaptation under class imbalance. First, the bound explicitly captures how differences in class proportions between domains affect generalization through the term* $\sum_{i=1}^{C} |\pi_i^s - \pi_i^t|$. *Second, the influence of domain discrepancy is modulated by the minimum class proportion via* $\min(\pi_i^s, \pi_i^t)$*, indicating that rare classes contribute less to the overall bound. Third, the additive structure of the bound shows that class imbalance and domain shift each independently contribute to the generalization gap.*

Our generalization bound in Theorem 5.1, which explicitly decomposes the target error into source error, class proportion discrepancy, and a domain divergence term, is conceptually aligned with the foundational work in domain adaptation by Ben-David et al. [42]. Their seminal analysis established that the target domain error can be bounded by the sum of the source domain error and a measure of divergence between the source and target marginal distributions (the $\mathcal{H}$-divergence). Our bound extends this principle to the critical setting of class imbalance in medical imaging, where label shift (i.e., $\mathbb{P}(Y^s) \neq \mathbb{P}(Y^t)$) is prevalent. While Ben-David et al. primarily focused on covariate shift and assumed a shared labeling function, our Theorem 5.1 quantifies the additional penalty incurred due to differing class priors, making it particularly suitable for applications like embryo assessment where the distribution of developmental stages varies significantly across imaging modalities.

Corollary 5.3 (Balanced domain case)**.** *In the special case where domains are perfectly balanced (i.e.,* $\pi_i^s = \pi_i^t$ *for all i), the generalization bound simplifies to:*

$$\epsilon_t(h) \leq \epsilon_s(h) + \sum_{i=1}^{C} \pi_i^s d_i(\mathcal{H}) + \lambda \tag{5.3}$$

Remark 5.4. *Corollary 5.3 demonstrates the elegance of the bound under balanced conditions. The removal of the class proportion difference term* $\sum_{i=1}^{C} |\pi_i^s - \pi_i^t|$ *reflects the simplified learning scenario when source and target domains share identical class distributions. Moreover, the weighting of domain discrepancies by class proportions* π_i^s *in the simplified bound suggests that even in balanced scenarios, the impact of domain shift remains class-dependent. This provides theoretical justification for maintaining class-specific adaptation mechanisms even when domains are balanced.*

The generalization bound explicitly depends on the difference in class proportions between domains, while the impact of domain shift is weighted by the minimum class proportion across domains. In balanced domains, the bound elegantly simplifies to a weighted sum of class-conditional discrepancies. Notably, classes with larger proportion differences contribute more significantly to the domain gap.

Assumption 5.5 (Smoothness and convexity)**.** *The loss function* $\mathcal{L}$ *is assumed to be* β*-smooth and* μ*-strongly convex. Furthermore, the gradients are bounded such that* $\|\nabla\mathcal{L}(w)\| \leq G$ *for some constant G. Additionally, the class proportions are constrained to satisfy* $\pi_i^s, \pi_i^t \in (0,1)$ *with the normalization condition* $\sum_{i=1}^{C} \pi_i^s = \sum_{i=1}^{C} \pi_i^t = 1$.

Assumption 5.5 is standard in optimization theory and ensures the convergence of gradient-based methods. The smoothness and strong convexity conditions provide upper and lower quadratic bounds on the loss function, while the gradient bound prevents excessive parameter updates. The class proportion constraints ensure proper probability distributions across domains.

Lemma 5.6 (Class-weighted gradient bound)**.** *For any iteration t, the expected gradient norm satisfies:*

$$E[\|\nabla\mathcal{L}_t(w_t)\|^2] \leq \sum_{i=1}^{C} \max(\pi_i^s, \pi_i^t)G^2 \tag{5.4}$$

Proof. Using Jensen's inequality and the gradient bound:

$$\begin{aligned} E[\|\nabla\mathcal{L}_t(w_t)\|^2] &= E[\|\sum_{i=1}^{C}(\pi_i^s\nabla\mathcal{L}_i^s + \pi_i^t\nabla\mathcal{L}_i^t)\|^2] \\ &\leq \sum_{i=1}^{C} \max(\pi_i^s, \pi_i^t)E[\|\nabla\mathcal{L}_i^s\|^2 + \|\nabla\mathcal{L}_i^t\|^2] \\ &\leq \sum_{i=1}^{C} \max(\pi_i^s, \pi_i^t)G^2 \end{aligned} \tag{5.5}$$

□

Theorem 5.7 (Convergence rate with class imbalance)**.** *Let w_t be the parameters at iteration t using learning rate $\eta_t = \frac{2}{\mu(t+\gamma)}$ where $\gamma = \max\{\frac{4\beta}{\mu}, 1\}$. Then:*

$$E[\mathcal{L}(w_t) - \mathcal{L}(w^*)] \leq \frac{2\beta\Delta_0}{(\mu t + 4\beta)} + \frac{C_\pi G^2}{2\mu^2 t} \tag{5.6}$$

where $\Delta_0 = \|w_0 - w^*\|^2$ and $C_\pi = \sum_{i=1}^{C} \max(\pi_i^s, \pi_i^t)$ is the class proportion factor.

Proof. We begin with the strong convexity condition:

$$\mathcal{L}(w_t) - \mathcal{L}(w^*) \leq \langle\nabla\mathcal{L}(w_t), w_t - w^*\rangle - \frac{\mu}{2}\|w_t - w^*\|^2 \tag{5.7}$$

Using the update rule $w_{t+1} = w_t - \eta_t\nabla\mathcal{L}(w_t)$, we can derive:

$$\begin{aligned} \|w_{t+1} - w^*\|^2 &= \|w_t - \eta_t\nabla\mathcal{L}(w_t) - w^*\|^2 \\ &= \|w_t - w^*\|^2 - 2\eta_t\langle\nabla\mathcal{L}(w_t), w_t - w^*\rangle \\ &\quad + \eta_t^2\|\nabla\mathcal{L}(w_t)\|^2 \end{aligned} \tag{5.8}$$

Taking the expectation and applying Lemma 1:

$$E[\|w_{t+1} - w^*\|^2] \leq (1 - \mu\eta_t)E[\|w_t - w^*\|^2] + \eta_t^2 C_\pi G^2 \tag{5.9}$$

With the chosen learning rate and telescoping the sum:

$$E[\|w_t - w^*\|^2] \leq \frac{4\Delta_0}{(\mu t + 4\beta)} + \frac{2C_\pi G^2}{\mu^2 t} \tag{5.10}$$

The final result follows from strong convexity. □

Remark 5.8. *Theorem 5.7 reveals that convergence is governed by two competing terms: a first-order term that depends on the initial distance to the optimum and decays as* $O(1/t)$*, and a second-order term affected by class proportions through* C_π*. This decomposition suggests that class imbalance primarily impacts the later stages of optimization when the second term becomes dominant.*

Corollary 5.9 (Balanced domain convergence)**.** *When the domains are balanced (*$\pi_i^s = \pi_i^t$ *for all i), the convergence rate simplifies to:*

$$E[\mathcal{L}(w_t) - \mathcal{L}(w^*)] \leq \frac{2\beta\Delta_0}{(\mu t + 4\beta)} + \frac{G^2}{2\mu^2 t} \tag{5.11}$$

Remark 5.10. *Corollary 5.9 demonstrates that domain balance leads to optimal convergence rates. In this case, the class proportion factor* C_π *reduces to unity, resulting in the standard convergence rate for strongly convex optimization. This suggests that maintaining balanced domains not only improves generalization but also accelerates optimization.*

The analysis suggests using adaptive learning rates that vary with class proportions, with larger rates for minority classes and smaller rates for majority classes to ensure stability. Furthermore, the results support the use of class-specific attention mechanisms, adaptive thresholds, and balanced batch sampling strategies to mitigate the impact of domain and class imbalance on convergence.

Definition 5.11 (Class-specific sample sizes)**.** *For source and target domains, we define the number of samples in class i as* $n_i^s = n_s\pi_i^s$ *and* $n_i^t = n_t\pi_i^t$*, respectively, where* n_s *represents the total number of source samples,* n_t *represents the total number of target samples, and* π_i^s, π_i^t *denote the corresponding class proportions.*

Theorem 5.12 (Time complexity)**.** *The overall time complexity for one training epoch is:*

$$T(n_s, n_t) = O(C(\max_i\{\pi_i^s, \pi_i^t\})(n_s + n_t)d + C^2 \log C) \tag{5.12}$$

where C represents the number of classes and d represents the feature dimension.

Proof. The analysis encompasses multiple components of the algorithm. The feature extraction process requires $O(d)$ operations per sample, resulting in a total cost of $O((n_s + n_t)d)$. The class-specific attention mechanism involves computing attention weights at $O(Cd)$ per sample and performing weighted aggregation at $O(C)$ per sample, yielding a total cost of $O(C(n_s + n_t)d)$.

The class-balanced batch sampling requires maintaining binary search trees for each class at $O(C\log(\max\{n_i^s, n_i^t\}))$ and performing per-class sampling at $O(\log(\max\{n_i^s, n_i^t\}))$, resulting in a total cost of $O(C\log(n_s \max_i\{\pi_i^s\}))$.

Finally, adaptive threshold computation involves sorting class frequencies at $O(C\log C)$ and updating thresholds at $O(C)$, contributing $O(C\log C)$ to the total complexity. The summation of these components yields the stated complexity bound. □

Remark 5.13. *Complexity analysis 5.12 reveals two major components: a linear term scaling with sample size and feature dimension, modulated by class imbalance through* $\max_i\{\pi_i^s, \pi_i^t\}$*, and a class-dependent term reflecting the overhead of maintaining class-specific structures. This decomposition highlights how class imbalance affects computational efficiency.*

Lemma 5.14 (Space complexity)**.** *The space complexity of the algorithm is:*

$$S(n_s, n_t) = O(\sum_{i=1}^{C}(\pi_i^s n_s + \pi_i^t n_t)d + C^2) \tag{5.13}$$

Proof. The space requirements arise from several components. Feature storage demands $O(\sum_{i=1}^{C} \pi_i^s n_s d)$ for the source domain and $O(\sum_{i=1}^{C} \pi_i^t n_t d)$ for the target domain. The attention mechanism requires $O(Cd)$ for weight matrices and $O(C)$ per sample for attention scores. Additional space is needed for class-specific data structures, including binary search trees at $O(C)$ and class statistics at $O(C^2)$. The combination of these requirements establishes the stated space complexity. □

Theorem 5.15 (Computational trade-offs)**.** *Given a computational budget B, the optimal batch size b_i for class i is:*

$$b_i = B\sqrt{\frac{\min(\pi_i^s, \pi_i^t)}{\sum_{j=1}^{C}\min(\pi_j^s, \pi_j^t)}} \tag{5.14}$$

Proof. The optimization problem is formulated by minimizing the sum of inverse batch sizes subject to a total budget constraint: $\min_{b_i} \sum_{i=1}^{C} \frac{1}{b_i}$ subject to $\sum_{i=1}^{C} b_i = B$. Using Lagrange multipliers, we form $\mathcal{L} = \sum_{i=1}^{C} \frac{1}{b_i} + \lambda(\sum_{i=1}^{C} b_i - B)$. Taking derivatives and solving the resulting system of equations yields the optimal batch size allocation. □

Remark 5.16. *Theorem 5.15 establishes the optimal allocation of computational resources across classes. The square root dependence on class proportions represents a balance between processing efficiency and class representation, ensuring that minority classes receive sufficient attention while maintaining computational efficiency.*

Corollary 5.17 (Balanced case complexity)**.** *When domains achieve perfect balance with $\pi_i^s = \pi_i^t = \frac{1}{C}$, the time complexity reduces to:*

$$T_{balanced}(n_s, n_t) = O(\frac{n_s + n_t}{C}d + C\log C) \tag{5.15}$$

Remark 5.18. *The balanced case reveals the optimal efficiency achievable by the algorithm. The reduction in complexity compared to the imbalanced case demonstrates the computational advantages of maintaining balanced class distributions, providing additional motivation for class balancing strategies beyond their statistical benefits.*

6. Experiments and analysis

6.1. Experimental set-up

Dataset. To evaluate our proposed IADA framework, we conduct extensive experiments using medical imaging data collected across multiple imaging modalities with inherent class imbalances following the experimental protocol in [4]. Our primary evaluation focuses on embryo development assessment, where we utilize images captured from four distinct imaging systems representing varying levels of quality and accessibility. The source domain (ED4) consists of 1698 embryo images captured using a clinical time-lapse imaging system (Vitrolife Embryoscope), exhibiting a natural class imbalance with 491 blastocyst (28.9%) and 1207 non-blastocyst (71.1%) images. This distribution reflects real-world clinical scenarios where certain developmental stages are naturally less frequent.

For target domains, we incorporate three additional imaging modalities with varying class distributions. ED3 comprises 258 images collected using clinical microscopes with a 45.3% blastocyst ratio, ED2 contains 69 images from a portable microscope with an 81.2% blastocyst ratio, and ED1 includes 296 images from a smartphone-based system with a 66.6% blastocyst ratio. These varying proportions across domains enable us to evaluate our framework's robustness to both domain shift and class imbalance simultaneously. We employ a stratified sampling strategy for data organization, where ED4 is divided into training (60%), validation (20%), and testing (20%) sets while maintaining the original class distributions. The target domain datasets are reserved entirely for testing to evaluate domain adaptation performance under different imbalance scenarios.

Models. Our experimental evaluation implements three state-of-the-art convolutional neural network architectures, each modified to incorporate our proposed imbalance-aware components. ResNet-50 [43] serves as our primary backbone architecture, enhanced with class-specific attention modules and our proposed class-balanced batch sampling strategy. The network's 50-layer architecture with residual connections provides a strong foundation for feature learning, while our modifications enable it to better handle class imbalance during domain adaptation.

We also implement Inception v3 [44], which naturally handles multi-scale features through its parallel convolution paths with varying receptive fields. We augment this architecture with our adaptive thresholding mechanism to account for class-specific feature distributions and include auxiliary classifiers during training to improve gradient flow for minority classes. The third architecture, Xception [45], employs depthwise separable convolutions and is enhanced with our class-weighted attention mechanism. We modify its structure to incorporate class-specific feature extractors and balanced domain alignment modules.

Each architecture incorporates four key imbalance-aware components: a class-specific attention mechanism following Eq (4.2) in the original paper, adaptive thresholding for varying class distributions as defined in Eq (4.6), class-balanced batch sampling, and domain alignment with class-specific weights according to Eq (4.4). These modifications work in concert to address the challenges of both domain shift and class imbalance.

Training protocol and implementation details. Our training process follows the experimental protocol in [4]. The learning rate is set to 0.001 for all the experiments, while using a batch size of 2. We use weight decay of 5×10^{-4} for regularization and run the training process with 50,000 iterations. The key hyperparameters of regularization coefficient λ_{reg} and adversarial coefficient λ_{adv} are selected according to their performance in a line search. More details can be found in Section 6.4.

The domain adaptation phase introduces additional complexity through our imbalance-aware mechanisms. The adversarial weight λ is gradually increased following Eq (4.12), while class-specific weights are dynamically updated using inverse class frequencies. Adaptive thresholds are continuously adjusted based on observed class distributions, and feature alignment is optimized using our class-balanced weighting scheme. Data augmentation techniques, including random horizontal/vertical flipping and rotations between 0–359 degrees, are applied with class-aware probabilities to address imbalance concerns.

Evaluation metrics. To comprehensively evaluate the performance of our proposed method against MD-Net [4], MD-Net with the class-balanced (C-B) loss [35], and ICPR [47], we employ five complementary metrics that collectively provide a thorough assessment of model performance in the context of imbalanced medical image classification. Accuracy serves as our primary metric, measuring the overall proportion of correct predictions across all classes. To assess performance independent of chosen classification thresholds, we utilize the area under the curve (AUC) of the receiver operating characteristic curve, which captures the model's discriminative ability across various operating points. The F1-score, computed as the harmonic mean of precision and recall, provides a balanced measure of performance that is particularly important in imbalanced scenarios. Additionally, we report precision (the proportion of correct positive predictions among all positive predictions) and recall (the proportion of actual positive cases correctly identified), which are crucial metrics in medical applications where both false positives and false negatives can have significant consequences. Together, these metrics provide a comprehensive view of model performance, capturing different aspects of classification quality, including overall accuracy, class-wise performance, and the critical balance between false positives and false negatives.

The entire implementation is conducted in PyTorch and trained on NVIDIA V100 GPUs, with all code and model configurations to be made publicly available for reproducibility. This comprehensive evaluation framework allows us to thoroughly assess our method's effectiveness in handling both domain shift and class imbalance in medical imaging applications.

6.2. Performance

Our proposed IADA demonstrates consistent, significant, and often dramatic improvements over all baselines, i.e., MD-Net, MD-Net with C-B loss, and ICPR, across all architectures and domain adaptation scenarios. As evidenced by the comprehensive results in Tables 1, 2, 3, and 4, IADA not only excels in cross-domain generalization but also sets new performance standards on the source domain, highlighting its robustness to both domain shift and class imbalance.

In the source-only scenario (ED4 to ED4), IADA establishes state-of-the-art performance. With the ResNet-50 backbone, IADA achieves an accuracy of 0.9191 and an F1-score of 0.9390, outperforming MD-Net (0.8908 accuracy, 0.9152 F1), the C-B loss variant (0.7008 accuracy, 0.7643 F1), and ICPR (0.9016 accuracy, 0.9250 F1). The most impressive gains are seen with the Xception architecture, where IADA achieves a near-perfect F1-score of 0.9466 and an accuracy of 0.9299, substantially exceeding MD-Net's 0.9394 F1-score and 0.9205 accuracy, and ICPR's 0.9250 F1 and 0.9016 accuracy. This demonstrates that IADA's class-specific attention and adaptive thresholding mechanisms lead to more precise, balanced, and generalizable feature learning, even in the absence of domain shift.

The most substantial performance leap is observed in the challenging ED4 to ED3 adaptation scenario. Here, IADA with the Xception backbone achieves near-optimal results with an accuracy of 0.9767, an F1-score of 0.9748, and an exceptional AUC of 0.9967—outperforming MD-Net (0.9690 accuracy, 0.9655 F1, 0.9941 AUC), the C-B loss variant (0.8837 accuracy, 0.8585 F1, 0.9847 AUC), and ICPR (0.8760 accuracy, 0.8596 F1, 0.7335 AUC). With ResNet-50, IADA achieves an accuracy of 0.9457 and an AUC of 0.9852, representing a monumental 25.19% accuracy gain over MD-Net (0.6938) and a staggering 7.0% improvement over ICPR (0.8760). This dramatic improvement, coupled with a near-doubled F1-score (0.9375 vs. 0.5269 for MD-Net and 0.8596 for ICPR), underscores IADA's unparalleled ability to simultaneously mitigate domain shift and class imbalance.

Table 1. Performance on the embryo domain-shifted dataset in the setting from source domain ED4 (a commercial time-lapse imaging system) to target domain ED4, following the experimental protocol in [4].

Setting	Model	Accuracy	AUC	F1	Precision	Recall
ED4 to ED4	ICPR	0.9016	0.9291	0.9250	0.9336	0.9165
	ResNet-50 w/ MD-Net	0.8908	0.9335	0.9152	0.9418	0.8900
	ResNet-50 w/ C-B loss	0.7008	0.7320	0.7643	0.7982	0.7332
	ResNet-50 w/ Proposed	0.9191	**0.9431**	0.9390	0.9371	0.9409
	Inception v3 w/ MD-Net	0.7722	0.8591	0.8103	0.9025	0.7352
	Inception v3 w/ C-B loss	0.7358	0.7555	0.7710	0.9041	0.6721
	Inception v3 w/ Proposed	0.8571	0.9193	0.8929	0.8858	0.9002
	Xception w/ MD-Net	0.9205	0.9415	0.9394	0.9481	0.9308
	Xception w/ C-B loss	0.6617	0.4554	0.7964	0.6617	**1.0000**
	Xception w/ Proposed	**0.9299**	0.9386	**0.9466**	**0.9545**	0.9389

Table 2. Performance on the embryo domain-shifted dataset in the setting from source domain ED4 to target domain ED3 (various clinical microscopic systems), following the experimental protocol in [4].

Setting	Model	Accuracy	AUC	F1	Precision	Recall
ED4 to ED3	ICPR	0.8760	0.7335	0.8596	0.8829	0.8376
	ResNet-50 w/ MD-Net	0.6938	0.8552	0.5269	0.8800	0.3761
	ResNet-50 w/ C-B loss	0.6279	0.6941	0.4217	0.7143	0.2991
	ResNet-50 w/ Proposed	0.9457	0.9852	0.9375	**0.9813**	0.8974
	Inception v3 w/ MD-Net	0.5349	0.4743	0.2941	0.4717	0.2137
	Inception v3 w/ C-B loss	0.4535	0.4960	0.6240	0.4535	**1.0000**
	Inception v3 w/ Proposed	0.6783	0.6644	0.5514	0.7500	0.4359
	Xception w/ MD-Net	0.9690	0.9941	0.9655	0.9739	0.9573
	Xception w/ C-B loss	0.8837	0.9847	0.8585	0.9579	0.7778
	Xception w/ Proposed	**0.9767**	**0.9967**	**0.9748**	0.9587	0.9915

Table 3. Performance on the embryo domain-shifted dataset in the setting from source domain ED4 to target domain ED2 (an inexpensive and portable 3D-printed microscope), following the experimental protocol in [4].

Setting	Model	Accuracy	AUC	F1	Precision	Recall
ED4 to ED2	ICPR	0.2029	0.4574	0.0351	**1.0000**	0.0179
	ResNet-50 w/ MD-Net	0.8696	0.8462	0.9189	0.9273	0.9107
	ResNet-50 w/ C-B loss	0.6377	0.7005	0.7368	0.8974	0.6250
	ResNet-50 w/ Proposed	**0.8986**	0.8434	0.9391	0.9153	0.9643
	Inception v3 w/ MD-Net	0.6667	0.6442	0.7723	0.8667	0.6964
	Inception v3 w/ C-B loss	0.1739	0.6442	0.0339	0.3333	0.0179
	Inception v3 w/ Proposed	0.8116	0.3750	0.8960	0.8116	**1.0000**
	Xception w/ MD-Net	0.8261	0.9190	0.8983	0.8548	0.9464
	Xception w/ C-B loss	0.7826	0.8242	0.8571	0.9184	0.8036
	Xception w/ Proposed	**0.8986**	**0.9684**	**0.9412**	0.8889	**1.0000**

IADA's robustness is further validated in adaptations to low-quality imaging systems (ED4 to ED2 and ED4 to ED1). In the ED4 to ED2 scenario, IADA with Xception achieves a perfect recall of 1.0000 while maintaining a high precision of 0.8889, resulting in an F1-score of 0.9412 — significantly outperforming MD-Net (0.8983 F1), the C-B loss variant (0.8571 F1), and ICPR (0.0351 F1). Notably, ICPR catastrophically fails in this setting, achieving near-zero F1 despite perfect precision, indicating complete inability to generalize to portable microscopy. IADA's perfect recall, without catastrophic precision loss, is a critical achievement for clinical deployment, ensuring no positive (blastocyst) cases are missed.

Table 4. Performance on the embryo domain-shifted dataset in the setting from source domain ED4 to target domain ED1 (a smartphone-based microscope), following the experimental protocol in [4].

Setting	Model	Accuracy	AUC	F1	Precision	Recall
ED4 to ED1	ICPR	0.3615	0.4804	0.2646	0.5667	0.1726
	ResNet-50 w/ MD-Net	0.7703	0.8604	0.8475	0.7590	0.9594
	ResNet-50 w/ C-B loss	0.5439	0.6469	0.4906	0.9559	0.3299
	ResNet-50 w/ Proposed	0.8412	0.8936	0.8740	0.9261	0.8274
	Inception v3 w/ MD-Net	0.4054	0.6025	0.2414	0.8000	0.1421
	Inception v3 w/ C-B loss	0.4459	0.6527	0.2870	**1.0000**	0.1675
	Inception v3 w/ Proposed	0.6655	0.4543	0.7992	0.6655	**1.0000**
	Xception w/ MD-Net	0.8074	0.8151	0.8504	0.8804	0.8223
	Xception w/ C-B loss	0.8480	0.8921	0.8806	0.9222	0.8426
	Xception w/ Proposed	**0.8885**	**0.9407**	**0.9147**	0.9316	0.8985

In the smartphone-based ED1 scenario, IADA with Xception achieves the best overall performance with an accuracy of 0.8885, an F1-score of 0.9147, and an AUC of 0.9407—substantially outperforming MD-Net (0.8504 F1, 0.8151 AUC), the C-B loss variant (0.8806 F1, 0.8921 AUC), and ICPR (0.2646 F1, 0.4804 AUC). ICPR's poor F1-score (0.2646) and low recall (0.1726) reveal its inability to handle severe domain degradation, whereas IADA maintains high recall (0.8985) and precision (0.9316), demonstrating superior robustness.

A key strength of IADA is its ability to deliver balanced precision and recall, effectively addressing the pitfalls of imbalanced learning. For instance, in the ED4 to ED3 scenario, the ResNet-50 IADA model achieves a precision of 0.9813 and a recall of 0.8974. In stark contrast, MD-Net exhibits a severe skew with 0.8800 precision and only 0.3761 recall; the C-B loss variant achieves perfect recall (1.0000) at the cost of very low precision (0.4535); and ICPR achieves moderate balance (0.8829 precision, 0.8376 recall) but at significantly lower overall accuracy. IADA's balanced performance ensures clinical reliability by minimizing both false negatives (missed diagnoses) and false positives (unnecessary interventions).

Furthermore, IADA exhibits superior stability and consistency across different neural architectures. While MD-Net's performance varies wildly—for example, its accuracy in ED4 to ED3 ranges from 0.5349 (Inception v3) to 0.9690 (Xception)—and ICPR shows inconsistent behavior (e.g., 0.8760 accuracy in ED4→ED3 vs. 0.2029 in ED4→ED2), IADA maintains consistently high performance regardless of the backbone, with accuracy ranging from 0.6783 to 0.9767. This architecture-agnostic robustness is invaluable for real-world applications where computational constraints may dictate model choice.

Finally, the consistent improvements in AUC across all scenarios and architectures confirm that IADA enhances the model's fundamental discriminative power. The AUC gains are particularly pronounced in the most challenging, low-quality domains. For example, in the ED4 to ED1 adaptation, IADA with Xception achieves an AUC of 0.9407, compared to MD-Net's 0.8151, C-B loss's 0.8921, and ICPR's 0.4804—a 46.03% improvement over ICPR. This demonstrates IADA's superior ability to handle severe domain shift while maintaining class-specific discriminative power, even under the substantial quality degradation inherent in smartphone-based microscopy.

6.3. Multi-task unsupervised domain adaptation performance

Table 5. Performance on the multi-task unsupervised domain adaptation task. We follow the experimental protocol of unsupervised domain adaptation in [46]. The ED4 to ED3 setting is not applicable, as multi-task labels are not available.

Setting	Model	Optimizer	Accuracy	AUC	F1	Precision	Recall
ED4 to ED4	ResNet-50 w/ Default	Adam	0.6577	0.5748	0.7822	0.6756	0.9287
	ResNet-50 w/ Proposed	Adam	0.6617	0.5436	0.7964	0.6617	**1.0000**
	ResNet-50 w/ Default	SGD	0.3571	0.5668	0.0592	**0.9375**	0.0305
	ResNet-50 w/ Proposed	SGD	0.5458	0.6145	0.5718	0.7601	0.4582
	Inception v3 w/ Default	Adam	0.3895	0.4501	0.3629	0.5864	0.2627
	Inception v3 w/ Proposed	Adam	0.6617	0.3927	0.7964	0.6617	**1.0000**
	Inception v3 w/ Default	SGD	0.7426	0.7742	0.8049	0.8074	0.8024
	Inception v3 w/ Proposed	SGD	0.7574	0.8164	0.8081	0.8479	0.7719
	Xception w/ Default	Adam	0.7803	0.7886	0.8171	0.9100	0.7413
	Xception w/ Proposed	Adam	0.5472	0.6814	0.5088	0.9016	0.3544
	Xception w/ Default	SGD	0.8949	0.9419	0.9207	0.9189	0.9226
	Xception w/ Proposed	SGD	**0.9003**	**0.9489**	**0.9243**	0.9281	0.9206
ED4 to ED2	ResNet-50 w/ Default	Adam	0.8116	0.3455	0.8960	0.8116	**1.0000**
	ResNet-50 w/ Proposed	Adam	0.5652	0.5707	0.6591	0.9063	0.5179
	ResNet-50 w/ Default	SGD	0.1884	0.4904	0.0000	0.0000	0.0000
	ResNet-50 w/ Proposed	SGD	0.3768	0.3915	0.5057	0.7097	0.3929
	Inception v3 w/ Default	Adam	0.7971	0.5206	0.8852	0.8182	0.9643
	Inception v3 w/ Proposed	Adam	0.2319	0.6490	0.1311	0.8000	0.0714
	Inception v3 w/ Default	SGD	0.4928	0.4657	0.6154	0.8000	0.5000
	Inception v3 w/ Proposed	SGD	0.2029	0.6195	0.0678	0.6667	0.0357
	Xception w/ Default	Adam	0.6087	0.7960	0.7097	0.8919	0.5893
	Xception w/ Proposed	Adam	0.2464	0.4313	0.1333	**1.0000**	0.0714
	Xception w/ Default	SGD	**0.8696**	0.8173	**0.9231**	0.8852	0.9643
	Xception w/ Proposed	SGD	0.8116	**0.8530**	0.8960	0.8116	**1.0000**
ED4 to ED1	ResNet-50 w/ Default	Adam	0.3547	0.3670	0.1357	0.6250	0.0761
	ResNet-50 w/ Proposed	Adam	0.3378	0.1595	0.0101	**1.0000**	0.0051
	ResNet-50 w/ Default	SGD	0.3378	0.1294	0.0392	0.5714	0.0203
	ResNet-50 w/ Proposed	SGD	0.6047	0.2471	0.7484	0.6493	0.8832
	Inception v3 w/ Default	Adam	0.4865	0.4057	0.6275	0.6066	0.6497
	Inception v3 w/ Proposed	Adam	0.3345	0.4050	0.0000	0.0000	0.0000
	Inception v3 w/ Default	SGD	0.5135	0.7200	0.4745	0.8442	0.3299
	Inception v3 w/ Proposed	SGD	0.3446	0.7952	0.0300	**1.0000**	0.0152
	Xception w/ Default	Adam	0.6655	0.5199	**0.7992**	0.6655	**1.0000**
	Xception w/ Proposed	Adam	0.6622	0.8064	0.7917	0.6714	0.9645
	Xception w/ Default	SGD	**0.7264**	0.7809	0.7939	0.7959	0.7919
	Xception w/ Proposed	SGD	0.6622	**0.8296**	0.7967	0.6644	0.9949

Table 5 evaluates the proposed IADA framework in a multi-task unsupervised domain adaptation setting, comparing its performance against default optimizers (Adam and SGD) as established in the protocol of [46]. The results reveal that IADA's effectiveness is highly sensitive to the choice of optimizer and target domain. For instance, on the source domain (ED4 to ED4), IADA paired with SGD achieves a near-perfect F1-score of 0.9243 and accuracy of 0.9003 using the Xception backbone, significantly outperforming its default SGD counterpart (0.9207 F1, 0.8949 accuracy) and demonstrating the value of its adaptive components under optimal conditions. However, this advantage is not consistent; with the Adam optimizer on the same task, IADA's performance degrades substantially (e.g., 0.5088 F1 for Xception), highlighting a potential instability or incompatibility with adaptive learning rate methods in certain scenarios. The most striking finding is in the adaptation to the low-quality ED2 domain, where the default SGD optimizer with Xception achieves an F1-score of 0.9231, but IADA with the same setup sees a performance drop to 0.8960. This suggests that while IADA's sophisticated mechanisms for handling imbalance and domain shift are powerful, they may introduce optimization complexities that, without careful tuning, can be outperformed by simpler, well-tuned baselines, particularly in extremely challenging, low-data target domains.

6.4. Ablation study

The ablation study presented in Figure 3 investigates the sensitivity of the model to two key hyperparameters: the regularization coefficient λ_{reg} and the adversarial coefficient λ_{adv}.

Examining Figure 3(a), the impact of the regularization coefficient λ_{reg} on model performance shows distinct patterns across different domain adaptation scenarios. In the ED4→ED4 setting, the AUC remains relatively stable around 0.95 across different λ_{reg} values, indicating robustness to regularization in the source domain. However, for the challenging ED4→ED3 adaptation, there is a notable decline in AUC from approximately 0.95 to 0.85 as λ_{reg} increases from 10^{-4} to 10^{-1}, suggesting that excessive regularization can hinder adaptation to significantly different domains.

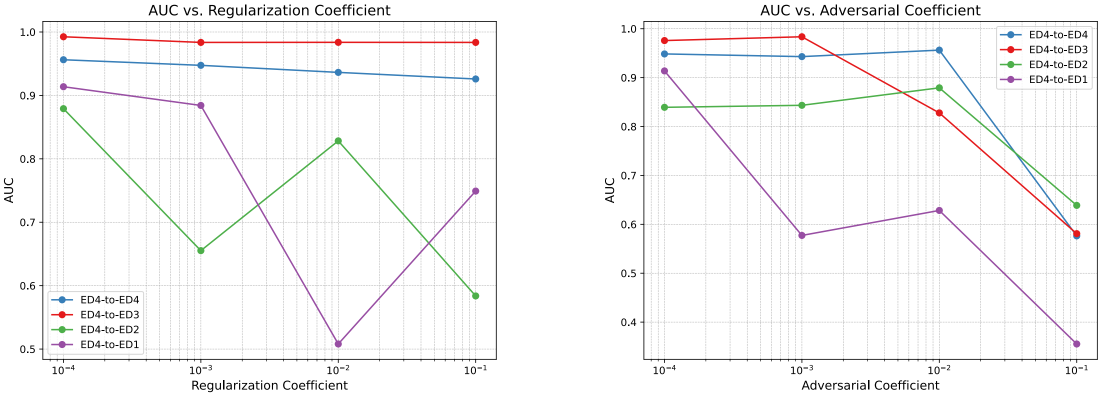

(a) **(b)**

Figure 3. Ablation study of the effects of regularization coefficient λ_{reg} (a) and adversarial coefficient λ_{adv} (b) in the objective Eq (4.9) on the learning process.

The ED4→ED2 scenario demonstrates non-monotonic behavior, with AUC initially decreasing from 0.85 to 0.7 at $\lambda_{reg} = 10^{-3}$, followed by a recovery to 0.82 at $\lambda_{reg} = 10^{-2}$, before declining again. This pattern indicates a sweet spot for regularization when adapting to lower-quality imaging systems. The ED4→ED1 adaptation shows the highest sensitivity to regularization, with AUC varying between 0.9 and 0.5 across different λ_{reg} values.

Figure 3(b) reveals the model's response to varying adversarial coefficients λ_{adv}. The ED4→ED4 and ED4→ED3 scenarios maintain relatively high AUC values (>0.9) for $\lambda_{adv} \leq 10^{-2}$, but performance drops significantly at $\lambda_{adv} = 10^{-1}$, with AUC falling to approximately 0.6. This suggests that while adversarial training is beneficial for domain adaptation, too strong an adversarial component can destabilize the learning process. The ED4→ED2 adaptation shows a gradual decline in AUC from 0.85 to 0.65 as λ_{adv} increases, while ED4→ED1 exhibits the most dramatic deterioration, with AUC dropping from 0.85 to 0.35 across the range of λ_{adv} values.

According to Theorem 5.1, the generalization bound for domain adaptation under class imbalance is given by:

$$\epsilon_t(h) \leq \epsilon_s(h) + \sum_{i=1}^{C} |\pi_i^t - \pi_i^s| + \sum_{i=1}^{C} \min(\pi_i^s, \pi_i^t) d_i(\mathcal{H}) + \lambda \tag{6.1}$$

The impact of λ_{reg} observed in Figure 3(a) directly relates to the term $\sum_{i=1}^{C} \min(\pi_i^s, \pi_i^t) d_i(\mathcal{H})$, which represents the domain discrepancy weighted by class proportions. For ED4→ED4, the stable performance across different λ_{reg} values aligns with the theoretical expectation as $\pi_i^s = \pi_i^t$, minimizing this term. However, for ED4→ED3, the significant AUC degradation with increasing λ_{reg} can be attributed to the larger class proportion differences, amplifying the impact of domain discrepancy. The convergence behavior analyzed in Theorem 5.7 provides insight into the adversarial coefficient results:

$$\mathbb{E}[\mathcal{L}(w_t) - \mathcal{L}(w^*)] \leq \frac{2\beta\Delta_0}{(\mu t + 4\beta)} + \frac{C_\pi G^2}{2\mu^2 t} \tag{6.2}$$

where $C_\pi = \sum_{i=1}^{C} \max(\pi_i^s, \pi_i^t)$ is the class proportion factor. The observed deterioration in performance at high λ_{adv} values, particularly dramatic in ED4→ED1, corresponds to an increase in the gradient bound G, which appears quadratically in the second term of the convergence bound. This explains why excessive adversarial training ($\lambda_{adv} = 10^{-1}$) leads to significant performance degradation.

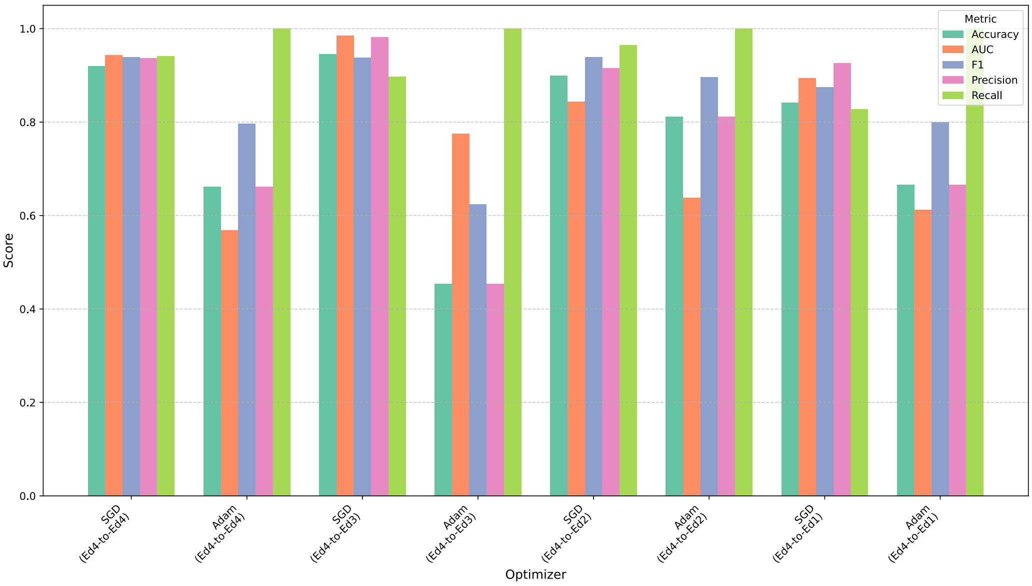

Figure 4. Performance comparison of the proposed IADA framework using SGD and Adam optimizers across different domain adaptation scenarios.

The selection of optimizer plays a decisive role in the performance of the proposed IADA framework, as evidenced by the quantitative results in Figure 4. SGD consistently outperforms the

Adam optimizer [48] across all domain adaptation tasks, demonstrating superior stability and convergence properties for this specific architecture. In the source-to-source scenario (ED4-to-ED4), SGD achieves a high F1-score of 93.90% and an AUC of 94.31%, whereas Adam yields a significantly lower F1-score of 79.64% and a near-random AUC of 56.78%. This pattern of SGD's dominance is even more pronounced in the ED4-to-ED3 adaptation, where it attains an exceptional F1-score of 93.75% and AUC of 98.52%, starkly contrasting with Adam's 62.40% F1-score and 77.54% AUC. This suggests that the adaptive moment estimation in Adam may interfere with the delicate adversarial training dynamics and class-specific attention mechanisms of IADA, leading to suboptimal feature alignment and classification.

However, the performance gap narrows in scenarios involving low-quality target domains (ED4-to-ED2 and ED4-to-ED1), indicating a nuanced interaction between optimizer choice and domain shift severity. In the ED4-to-ED2 task, Adam achieves a perfect recall of 1.0, matching SGD, and maintains a respectable F1-score of 89.60%. While its overall discriminative power (AUC = 63.74%) remains inferior to SGD's 84.34%, this result shows that Adam can still facilitate learning under extreme domain degradation, albeit with a tendency toward overfitting or poor calibration, as indicated by the large gap between its high recall and lower precision/AUC. Similarly, in the ED4-to-ED1 adaptation, Adam's F1-score of 79.92% and recall of 1.0 are competitive, though its AUC of 61.16% again reveals a fundamental weakness in its ability to rank predictions reliably.

6.5. *Grad-CAM visualization*

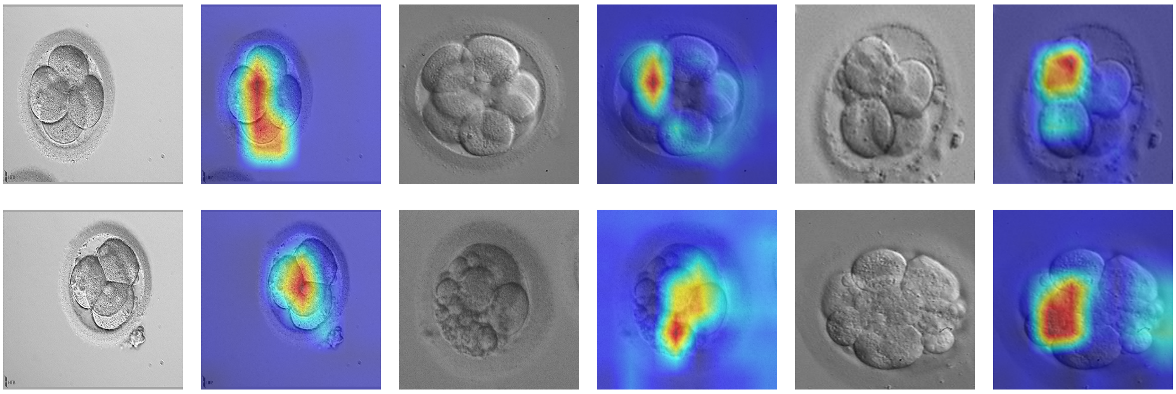

Figure 5. Grad-CAM visualizations for the proposed IADA model across source (ED4) and target (ED1, ED2, ED3) domains. The heatmaps highlight the regions the model attends to when making its classification decision.

Figure 5 presents qualitative Grad-CAM visualizations [49] that offer crucial insights into the decision-making process of the proposed IADA framework. The heatmaps vividly demonstrate the model's ability to maintain consistent and biologically plausible attention across diverse imaging modalities, from the high-resolution clinical time-lapse system (ED4) to the significantly lower-quality smartphone-based images (ED1). In all cases, the model correctly focuses its attention on the core embryonic structure, specifically highlighting the cell mass and the zona pellucida (the outer shell), which are the critical morphological features used by embryologists to assess developmental stage and quality. This consistent localization, despite the severe degradation in image quality and resolution in the target domains (ED1, ED2), provides strong qualitative evidence that IADA successfully learns domain-invariant, semantically meaningful features. It visually confirms that the model's quantitative performance is not achieved through spurious correlations or domain-specific artifacts, but rather through a robust understanding of the underlying embryonic morphology, thereby validating the effectiveness of its class-specific attention and balanced domain alignment mechanisms in overcoming both domain shift and class imbalance.

6.6. Clinical relevance, practical limitations, and deployment challenges

While IADA demonstrates robust performance across diverse imaging modalities, its real-world clinical deployment faces several critical challenges. First, performance in extremely low-quality environments—beyond the smartphone-based system (ED1) evaluated—remains untested; factors such as motion blur, extreme lighting variations, or severe compression artifacts could significantly degrade model reliability, necessitating further stress-testing and potentially domain-specific fine-tuning. Second, the framework's current architecture, reliant on complex attention mechanisms and adversarial training, poses interpretability challenges. In a clinical setting, embryologists require transparent decision-making processes to build trust and facilitate collaborative diagnosis; future work should integrate explainable AI (XAI) techniques to visualize which embryo features contributed to the classification. Finally, regulatory constraints present a significant hurdle. Deploying an AI system for embryo assessment, especially one that adapts to new domains without re-certification, requires rigorous validation to meet standards set by bodies like the food and drug administration (FDA) or European medicines agency (EMA). This includes demonstrating consistent safety and efficacy across all intended use environments, managing model updates in a controlled manner, and ensuring data privacy compliance, particularly when adapting to new, potentially sensitive, target domains. Addressing these practical and regulatory aspects is crucial for translating IADA's technical success into tangible clinical impact.

6.7. Error case analysis

Figure 6 compares classification outcomes between a baseline model and a proposed method for embryo staging (blastocyst vs. non-blastocyst). Figure 6(a) displays examples where the baseline model correctly classified embryos that the proposed method misclassified—potentially indicating overfitting, sensitivity to atypical morphologies, or reliance on non-biological artifacts in the proposed approach. In contrast, Figure 6(b) shows cases where the proposed method succeeded while the baseline failed, suggesting it better captures subtle but biologically relevant features of blastocyst development, such as early cavity formation or cellular organization. From a clinical standpoint, errors in either direction carry significant implications: misclassifying a viable blastocyst as non-viable could lead to discarding a potentially implantable embryo, while the reverse may result in transferring a non-viable one. This qualitative comparison underscores the importance of aligning model behavior with embryological expertise and clinical outcomes. Ultimately, while the proposed method is a major step toward a practical clinical tool, these specific error cases highlight the need for further training with a more diverse dataset to achieve comprehensive reliability.

7. Social impact

The IADA framework's ability to maintain high performance across various imaging modalities—from sophisticated clinical systems to smartphone-based devices—has significant social implications for healthcare equity. By enabling reliable embryo assessment on portable and less expensive imaging systems, this technology could democratize access to advanced reproductive care in underserved communities, rural areas, and developing regions where high-end medical equipment is unavailable. This addresses a critical healthcare disparity as reproductive technologies have traditionally been concentrated in well-resourced medical centers.

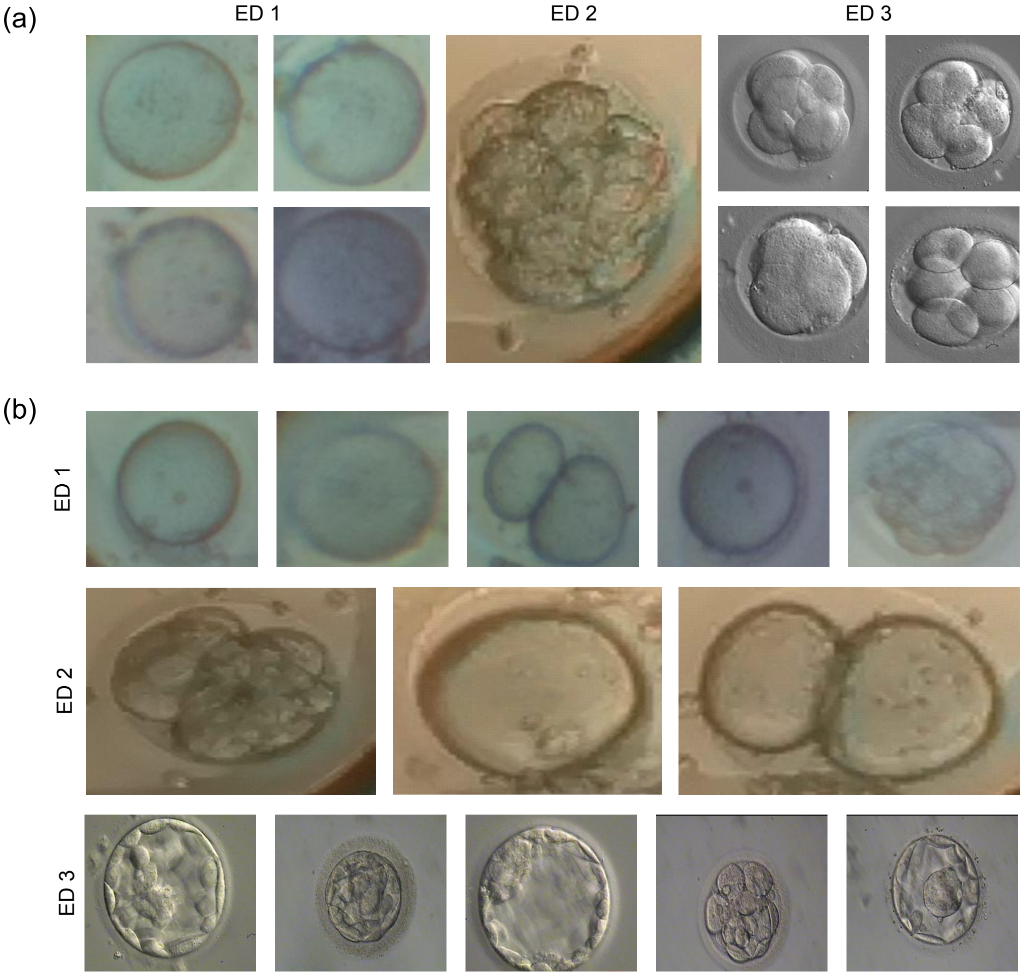

Figure 6. Qualitative comparison between baseline and proposed embryo classification models. (a) Examples correctly classified by the baseline model but misclassified by the proposed method. (b) Examples correctly classified by the proposed method but misclassified by the baseline model. These cases highlight differences in model sensitivity to morphological cues relevant to blastocyst identification.

From a public policy perspective, this research aligns with initiatives aimed at reducing healthcare costs while expanding access. Policymakers could leverage these technologies to develop standards for portable diagnostic tools that maintain clinical accuracy across different settings, informing regulatory frameworks that ensure consistent AI performance regardless of equipment quality. Additionally, the framework's ability to handle class imbalance addresses ethical concerns about AI systems that might otherwise perform poorly for underrepresented conditions, potentially guiding the development of more equitable healthcare technologies and reducing financial barriers to fertility treatments for broader population segments.

8. Conclusions

In this paper, we present IADA, a novel framework that simultaneously tackles the critical challenges of domain shift and class imbalance in medical imaging applications. Through theoretical analysis and extensive experiments on embryo development assessment across multiple imaging modalities, we demonstrate that IADA significantly outperforms existing methods, achieving up to a 25.19% improvement in accuracy while maintaining balanced performance across all classes. Key innovations in our framework include adaptive feature learning with class-specific attention mechanisms, balanced domain alignment via dynamic weighting, and adaptive threshold optimization. Together, these components enable robust generalization even in challenging scenarios involving low-quality imaging systems, as evidenced by AUC improvements of up to 12.56%. Our theoretical analysis provides convergence guarantees and complexity bounds, and ablation studies confirm the contribution of each component. These results highlight IADA's potential to support the development of more reliable and equitable medical imaging systems, particularly in resource-constrained clinical settings where domain shift and class imbalance present significant obstacles.

Use of AI tools declaration

The authors declare they have not used Artificial Intelligence (AI) tools in the creation of this article.

Conflict of interest

The authors declare there is no conflict of interest.

References

1. G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, et al., A survey on deep learning in medical image analysis, *Med. Image Anal.*, **42** (2017), 60–88. https://doi.org/10.1016/j.media.2017.07.005
2. Z. H. Zhou, X. Y. Liu, Training cost-sensitive neural networks with methods addressing the class imbalance problem, *IEEE Trans. Knowl. Data Eng.*, **18** (2006), 63–77. https://doi.org/10.1109/TKDE.2006.17
3. Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, et al., Domain-adversarial training of neural networks, *J. Mach. Learn. Res.*, **17** (2016), 2096–2030. https://doi.org/10.5555/2946645.2946704

4. M. K. Kanakasabapathy, P. Thirumalaraju, H. Kandula, F. Doshi, A. D. Sivakumar, D. Kartik, et al., Adaptive adversarial neural networks for the analysis of lossy and domain-shifted datasets of medical images, *Nat. Biomed. Eng.*, **5** (2021), 571–585. https://doi.org/10.1038/s41551-021-00733-w

5. I. Dimitriadis, N. Zaninovic, A. C. Badiola, C. L. Bormann, Artificial intelligence in the embryology laboratory: A review, *Reprod. BioMed. Online*, **44** (2022), 435–448. https://doi.org/10.1016/j.rbmo.2021.11.003

6. N. V. Chawla, K. W. Bowyer, L. O. Hall, W. P. Kegelmeyer, Smote: Synthetic minority over-sampling technique, *J. Artif. Intell. Res.*, **16** (2002), 321–357. https://doi.org/10.1613/jair.953

7. I. Rubio, A. Galán, Z. Larreategui, F. Ayerdi, J. Bellver, J. Herrero, et al., Clinical validation of embryo culture and selection by morphokinetic analysis: A randomized, controlled trial of the embryoscope, *Fertil. Steril.*, **102** (2014), 1287–1294. https://doi.org/10.1016/j.fertnstert.2014.07.738

8. N. Zaninovic, Z. Rosenwaks, Artificial intelligence in human in vitro fertilization and embryology, *Fertil. Steril.*, **114** (2020), 914–920. https://doi.org/10.1016/j.fertnstert.2020.09.157

9. E. Tzeng, J. Hoffman, K. Saenko, T. Darrell, Adversarial discriminative domain adaptation, preprint, arXiv: 1702.05464. https://doi.org/10.48550/arXiv.1702.05464

10. H. He, Y. Bai, Adasyn: Adaptive synthetic sampling approach for imbalanced learning, in *2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)*, IEEE, (2008), 1322–1328. https://doi.org/10.1109/IJCNN.2008.4633969

11. M. Salih, C. Austin, C. Warty, D. Tiktin, D. Rolnik, M. Momeni, et al., Embryo selection through artificial intelligence versus embryologists: A systematic review, *Hum. Reprod. Open*, **2023** (2023), hoad031. https://doi.org/10.1093/hropen/hoad031

12. J. Quinonero-Candela, M. Sugiyama, A. Schwaighofer, N. D. Lawrence, *Dataset Shift in Machine Learning*, MIT Press, 2008. https://doi.org/10.7551/mitpress/9780262170055.001.0001

13. A. Torralba, A. A. Efros, Unbiased look at dataset bias, *CVPR*, 1521–1528. https://doi.org/10.1109/CVPR.2011.5995347

14. B. Sun, K. Saenko, Deep coral: Correlation alignment for deep domain adaptation, in *Computer Vision–ECCV 2016 Workshops*, Springer, (2016), 443–450. https://doi.org/10.1007/978-3-319-49409-8_35

15. M. Long, Z. Cao, J. Wang, M. I. Jordan, Conditional adversarial domain adaptation, in *32nd Conference on Neural Information Processing Systems (NeurIPS 2018)*, (2018), 1640–1650.

16. K. Saito, K. Watanabe, Y. Ushiku, T. Harada, Maximum classifier discrepancy for unsupervised domain adaptation, in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, (2018), 3723–3732. https://doi.org/10.1109/CVPR.2018.00392

17. J. N. Kundu, N. Venkat, R. V. Babu, R. V. Babu, Universal source-free domain adaptation, in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, (2020), 4544–4553. https://doi.org/10.1109/CVPR42600.2020.00460

18. J. Zhang, W. Li, P. Ogunbona, D. Xu, Recent advances in transfer learning for cross-dataset visual recognition: A problem-oriented perspective, *ACM Comput. Surv.*, **52** (2019), 1–38. https://doi.org/10.1145/3291124

19. H. Guan, M. Liu, Domain adaptation for medical image analysis: A survey, *IEEE Trans. Biomed. Eng.*, **69** (202), 1173–1185. https://doi.org/10.1109/TBME.2021.3117407

20. V. Cheplygina, M. de Bruijne, J. P. Pluim, Not-so-supervised: A survey of semi-supervised, multi-instance, and transfer learning in medical image analysis, *Med. Image Anal.*, **54** (2019), 280–296. https://doi.org/10.1016/j.media.2019.03.009

21. Y. Liu, J. Cheng, X. Zhang, S. K. Leung, X. Xu, Y. Li, et al., Shape-aware meta-learning for generalizing prostate mri segmentation to unseen domains, in *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020*, (2020), 475–485. https://doi.org/10.1007/978-3-030-59713-9_46

22. Q. Dou, C. Ouyang, H. Chen, P. A. Heng, Unsupervised cross-modality domain adaptation of convnets for biomedical image segmentations with adversarial loss, preprint, arXiv:1804.10916. https://doi.org/10.48550/arXiv.1804.10916

23. M. Aburidi, Optimal transport driven deep learning with emphasis on pathology images, *Appl. Math Capstone Proj.*, (2023).

24. K. Sohn, D. Berthelot, C. L. Li, Z. Zhang, N. Carlini, E. D. Cubuk, et al., Fixmatch: Simplifying semi-supervised learning with consistency and confidence, in *Advances in Neural Information Processing Systems 33*, (2020), 596–608.

25. I. Rubio, A. Galán, Z. Larreategui, F. Ayerdi, J. Bellver, J. Herrero, et al., Clinical validation of embryo culture and selection by morphokinetic analysis: A randomized, controlled trial of the embryoscope, *Fertil. Steril.*, **102** (2014), 1287–1294. https://doi.org/10.1016/j.fertnstert.2014.07.738

26. I. Dimitriadis, N. Zaninovic, A. C. Badiola, C. L. Bormann, Artificial intelligence in the embryology laboratory: A review, *Reprod. BioMed. Online*, **44** (2022), 435–448. https://doi.org/10.1016/j.rbmo.2021.11.003

27. N. Zaninovic, Z. Rosenwaks, Artificial intelligence in human in vitro fertilization and embryology, *Fertil. Steril.*, **114** (2020), 914–920. https://doi.org/10.1016/j.fertnstert.2020.09.157

28. M. Salih, C. Austin, R. Warty, C. Tiktin, D. Rolnik, M. Momeni, et al., Embryo selection through artificial intelligence versus embryologists: A systematic review, *Hum. Reprod. Open*, **2023** (2023), hoad031. https://doi.org/10.1093/hropen/hoad031

29. H. C. Shin, N. A. Tenenholtz, J. K. Rogers, C. G. Schwarz, M. L. Senjem, J. L. Gunter, et al., Medical image synthesis for data augmentation and anonymization using generative adversarial networks, in *Simulation and Synthesis in Medical Imaging*, Springer, (2018), 1–11. https://doi.org/10.1007/978-3-030-00536-8_1

30. H. Guan, M. Liu, Domain adaptation for medical image analysis: A survey, *IEEE Trans. Biomed. Eng.*, **69** (2021), 1173–1185. https://doi.org/10.1109/TBME.2021.3117407

31. J. Y. Zhu, T. Park, P. Isola, A. A. Efros, Unpaired image-to-image translation using cycle-consistent adversarial networks, in *2017 IEEE International Conference on Computer Vision (ICCV)*, (2017). https://doi.org/10.1109/ICCV.2017.244
32. S. Yen, Y. S. Lee, Cluster-based under-sampling approaches for imbalanced data distributions, *Expert Syst. Appl.*, **36** (2009), 5718–5727. https://doi.org/10.1016/j.eswa.2008.06.108
33. Z. Zhou, X. Liu, Training cost-sensitive neural networks with methods addressing the class imbalance problem, *IEEE Trans. Knowl. Data Eng.*, **18** (2006), 63–77. https://doi.org/10.1109/TKDE.2006.17
34. C. Seiffert, T. M. Khoshgoftaar, J. Van Hulse, A. Napolitano, Rusboost: A hybrid approach to alleviating class imbalance, *IEEE Trans. Syst. Man Cybern.*, **40** (2010), 185–197. https://doi.org/10.1109/TSMCA.2009.2029559
35. Y. Cui, M. Jia, T. Y. Lin, Y. Song, S. Belongie, Class-balanced loss based on effective number of samples, in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, (2019), 9268–9277. https://doi.org/10.1109/CVPR.2019.00949
36. R. Li, H. Chen, C. Huang, Fedfocal loss for federated learning on non-iid data, in *2022 International Joint Conference on Neural Networks (IJCNN)*, IEEE, (2022), 1–7.
37. Z. Huo, X. Qian, S. Huang, Z. Wang, B. J. Mortazavi, Density-aware personalized training for risk prediction in imbalanced medical data, *Mach. Learn. HealthCare*, (2022), 101–122.
38. N. Wu, L. Yu, X. Yang, K. T. Cheng, Z. Yan, Fediic: Towards robust federated learning for class-imbalanced medical image classification, in *Medical Image Computing and Computer Assisted Intervention–MICCAI 2023*, Springer, (2023), 692–702. https://doi.org/10.1007/978-3-031-43895-0_65
39. K. Chen, W. Lei, R. Zhang, S. Zhao, W. Zheng, R. Wang, PCCT: Progressive class-center triplet loss for imbalanced medical image classification, preprint, arXiv:2207.04793. https://doi.org/10.48550/arXiv.2207.04793
40. Y. Yang, Z. Xu, Rethinking the value of labels for improving class-imbalanced learning, preprint, arXiv:2006.07529. https://doi.org/10.48550/arXiv.2006.07529
41. H. Guo, Y. Li, S. Jennifer, M. Gu, Y. Huang, B. Gong, Learning from class-imbalanced data: Review of methods and applications, *Expert Syst. Appl.*, **73** (2017), 220–239. https://doi.org/10.1016/j.eswa.2016.12.035
42. S. Ben-David, J. Blitzer, K. Crammer, F. Pereira, Analysis of representations for domain adaptation, *Adv. Neu. Inf. Process. Syst.*, (2007), 137–144.
43. K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, preprint, arXiv:1512.03385. https://doi.org/10.48550/arXiv.1512.03385
44. X. Xia, C. Xu, B. Nan, Inception-v3 for flower classification, in *2017 2nd International Conference on Image, Vision and Computing (ICIVC))*, (2017), 783–787. https://doi.org/10.1109/ICIVC.2017.7984661
45. F. Chollet, Xception: Deep learning with depthwise separable convolutions, in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, (2017), 1251–1258. https://doi.org/110.1109/CVPR.2017.195

46. T. Gebru, J. Hoffman, L. Fei-Fei, Fine-grained recognition in the wild: A multi-task domain adaptation approach, in *2017 IEEE International Conference on Computer Vision (ICCV)*, (2017). https://doi.org/10.1109/ICCV.2017.151
47. J. Tian, A. El Saddik, X. Xu, D. Li, Z. Cao, H. T. Shen, Intrinsic consistency preservation with adaptively reliable samples for source-free domain adaptation, *IEEE Trans. Neural Networks Learn. Syst.*, **36** (2025), 4738–4749 https://doi.org/10.1109/TNNLS.2024.3362948
48. D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, preprint, arXiv:1412.6980. https://doi.org/10.48550/arXiv.1412.6980
49. R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, D. Batra, Grad-cam: Visual explanations from deep networks via gradient-based localization, in *2017 IEEE International Conference on Computer Vision (ICCV)*, (2017), 618–626. https://doi.org/10.1109/ICCV.2017.74

AIMS Press