ISSN 1846-6168 (Print), ISSN 1848-5588 (Online)
https://doi.org/10.31803/tg-20250410041940

Original scientific paper
Received: 2025-04-10, Accepted: 2025-11-14

Adaptive GANs with Domain-Specific Losses for Data Imbalance

Sang Hyun Yoo, Hyun Jung Kim*

Abstract: This study addresses data imbalance in healthcare and finance through a novel domain-adaptive Conditional GAN framework with theoretically grounded, specialized loss functions and adaptive weight adjustment mechanisms. We introduce domain-specific loss functions based on established domain knowledge: feature correlation preservation for medical diagnosis and temporal consistency for financial fraud detection. The adaptive weight adjustment mechanism employs mathematically rigorous formulations with proven convergence guarantees under bounded gradient conditions. Comprehensive empirical validation on Wisconsin Breast Cancer and Credit Card Fraud datasets demonstrates significant improvements over state-of-the-art methods. Our approach achieves Feature Correlation Preservation scores of 0.849±0.004 versus 0.000 for traditional methods, and superior F1-score performance with improvements over recent competing methods while maintaining faster training times. The method provides exceptional robustness across extreme imbalance ratios from 1:10 to 1:2000, with advantages increasing at higher imbalance levels. Statistical validation using 5-fold stratified cross-validation with Bonferroni correction confirms significance ($p < 0.001$) with large effect sizes (*Cohen's d* > 1.4). Scalability analysis demonstrates practical applicability up to 1 M samples with manageable computational overhead. This work establishes a foundational framework for domain-aware synthetic data generation with proven theoretical guarantees and broad implications for imbalanced learning across critical application domains.

Keywords: adaptive weight mechanisms; conditional generative adversarial networks; convergence theory; domain-specific loss functions; extreme imbalance handling

1 INTRODUCTION

Machine learning models often face critical challenges with data imbalance, where certain classes are significantly underrepresented. This imbalance can severely compromise model performance, especially in critical domains such as medical diagnosis and financial fraud detection [1, 2]. The consequences extend beyond statistical performance metrics to real-world impacts: misdiagnosed medical conditions and undetected fraudulent transactions carry substantial societal costs.

Generative Adversarial Networks (GANs) offer a powerful alternative for data augmentation in imbalanced settings. Recent advances in financial GANs have shown promising results: Aftabi et al. (2023) developed innovative GAN-based approaches for financial fraud detection with data mining techniques [12], Sharma et al. (2023) demonstrated superior performance addressing class imbalance through hybrid SMOTE-GAN techniques [13], and Zadorozhnyy & Cheng (2021) introduced adaptive weighted discriminators for improved GAN training [14].

However, existing methodologies, including these recent advances, suffer from critical limitations:

- Lack of Theoretical Foundations: Most approaches lack rigorous mathematical frameworks for selecting and incorporating domain-specific constraints with convergence guarantees.
- Limited Domain Generalizability: Methods are typically designed for specific domains without systematic frameworks for knowledge transfer.
- Inadequate Extreme Imbalance Handling: Performance degrades significantly at imbalance ratios exceeding 1000:1, common in real-world applications.
- Missing Quality-Performance Trade-off Analysis: No systematic evaluation of domain fidelity versus classification performance trade-offs.

This paper addresses these fundamental limitations by presenting Adaptive GANs with Domain-Specific Losses, a novel theoretical framework that provides principled approaches for identifying and incorporating domain knowledge into GAN training objectives with proven mathematical guarantees. Our key contributions include:

1) Theoretical Framework for Domain-Specific Loss Design: We establish mathematical foundations for selecting domain-appropriate loss functions based on information theory and domain adaptation principles, with formal convergence guarantees providing $O(\log T)$ regret bounds.
2) Adaptive Weight Adjustment Mechanism: We introduce a mathematically rigorous adaptive mechanism with proven convergence guarantees under Lipschitz continuity assumptions, addressing critical stability issues in extreme imbalance scenarios.
3) Comprehensive Empirical Validation: We demonstrate the framework's effectiveness through extensive experiments across medical and financial domains, including 5-fold stratified cross-validation and statistical significance testing, showing significant improvements over recent state-of-the-art methods.
4) Quality-Performance Trade-off Analysis: We provide the first systematic analysis of the fundamental trade-off between domain fidelity and classification performance in synthetic data generation.

Our work represents the first systematic integration of multiple domain-specific objectives into GAN training for imbalance handling with rigorous theoretical guarantees, demonstrating a general strategy to incorporate domain knowledge into deep generative models while providing mathematically proven stability and convergence properties.

2 RELATED WORK

Research on imbalanced data spans techniques from data-level methods to algorithm-level modifications [15]. We focus on two areas most pertinent to our study: GAN-based oversampling approaches and domain-specific strategies that inform model training with domain knowledge.

2.1 Generative Oversampling with GANs

GANs have gained significant traction for generating minority-class samples due to their ability to model complex data distributions [16]. Douzas and Bacao (2018) developed a conditional WGAN for tabular imbalanced data outperforming SMOTE [10]. BAGAN (2018) introduced autoencoder pre-training to initialize the generator with all classes, addressing the generator's tendency to ignore minority modes [11].

Recent Financial GAN Advances: Several recent works have specifically targeted financial applications. Aftabi et al. (2023) introduced innovative approaches combining data mining with GAN models for financial fraud detection, demonstrating superior performance in detecting fraudulent patterns in financial statements [12]. Sharma et al. (2023) addressed class imbalance issues through hybrid SMOTE-GAN techniques, showing significant improvements in fraud detection accuracy with better minority class recognition [13]. Zadorozhnyy & Cheng (2021) proposed adaptive weighted loss functions for GAN discriminators, achieving improved training stability and performance across multiple datasets [14].

Temporal-Aware Generation: Recent advances emphasize temporal dependencies in financial data. Wu et al. (2024) developed Fin-GAN for financial time series forecasting and classification, incorporating temporal dependencies through novel loss functions and achieving superior performance in financial market prediction [19]. Zhang et al. (2022) addressed cryptocurrency transaction analysis using GAN-based approaches for abnormal transaction detection in Bitcoin networks [20].

While these studies demonstrate evolution from straightforward label conditioning to sophisticated domain-aware architectures, most methods lack theoretical frameworks for systematic domain knowledge integration and adaptive weight adjustment mechanisms with convergence guarantees.

2.2 Domain-Specific Losses and Adaptive Techniques

Integrating domain knowledge into learning algorithms has proven effective across various applications. TimeGAN by Yoon et al. (2019) combined GAN loss with stepwise supervision to preserve temporal dynamics in financial time series [21], informing our financial domain loss design. The weighting of multi-objective losses represents another relevant advancement. Xu et al. (2019) explored dynamic loss weight tuning in GAN settings, where loss term contributions adjust based on gradient magnitudes or performance metrics [22]. However, these approaches lack theoretical guarantees for convergence and stability.

2.3 Distinctiveness of Current Research

Our work uniquely combines theoretically-grounded domain-specific loss functions with adaptive weight adjustment mechanisms. This integrated approach addresses limitations in previous studies by providing:

Principled Domain Knowledge Integration: Systematic methodology for incorporating domain expertise.

Mathematical Rigor: Convergence guarantees and stability analysis for adaptive mechanisms.

Comprehensive Validation: Multi-domain empirical validation with statistical significance testing.

Scalability Analysis: Practical applicability assessment for real-world deployment.

3 METHODOLOGY

This section outlines our theoretical framework for domain-adaptive GAN training, including specialized loss functions and adaptive weight adjustment mechanisms with mathematical guarantees.

3.1 Base GAN Architecture

Our foundation is a conditional GAN (cGAN) architecture modified to accommodate domain-specific loss functions and adaptive weight adjustments.

3.1.1 Generator Architecture

The generator G takes as input a noise vector $\boldsymbol{z} \in R^d$ sampled from $N(0, 1)$ and class labels $\boldsymbol{y}$ encoded as one-hot vectors. We define the generator through a hierarchical structure

$$\boldsymbol{h}_0 = [\boldsymbol{z}; \boldsymbol{y}] \tag{1}$$

$$\boldsymbol{h}_1 = LeakyReLU\big(BN(\boldsymbol{W}_1\boldsymbol{h}_0 + \boldsymbol{b}_1)\big) \tag{2}$$

$$\boldsymbol{h}_2 = LeakyReLU\big(BN(\boldsymbol{W}_2\boldsymbol{h}_1 + \boldsymbol{b}_2)\big) \tag{3}$$

$$\boldsymbol{G}(\boldsymbol{z}, \boldsymbol{y}) = \sigma(\boldsymbol{W}_3\boldsymbol{h}_2 + \boldsymbol{b}_3) \tag{4}$$

where $[\boldsymbol{z}; \boldsymbol{y}] \in \mathbb{R}^{d+c}$ represents the concatenated input vector, $BN(\cdot)$ denotes batch normalization, and the architecture employs dimensions [256, 512, data_dim] with dropout ($p = 0.3$) for regularization.

3.1.2 Discriminator Architecture

The discriminator D receives either real data x or generated data $G(z, y)$ and outputs both real/fake prediction and class prediction

$$D_{real/fake}(x) = \sigma(\boldsymbol{W}_d\boldsymbol{h}_1 + \boldsymbol{b}_d) \tag{5}$$

This dual-output design enables the network to distinguish real from fake samples while maintaining class-specific information simultaneously.

3.2 Domain-Specific Loss Functions

A key innovation in our approach is the development of specialized loss functions tailored to unique characteristics of different domains.

3.2.1 Medical Domain Loss Functions

The medical domain loss function incorporates three essential components (Fig. 1).

Standard GAN Loss:

$$\mathcal{L}_{GAN} = E_{xp_{\text{data}}}[logD(x)] + E_{zp_z}\left[log\left(1 - D\left(G(z)\right)\right)\right] \quad (6)$$

Diagnostic Consistency Loss:

$$\mathcal{L}_{DC} = \left|\left|Corr_{real} - Corr_{gen}\right|\right|_F \quad (7)$$

Medical Feature Correlation Loss:

$$\mathcal{L}_{MC} = \sum_i \sum_j \left|\omega_{ij}(real) - \omega_{ij}(gen)\right| \quad (8)$$

The combined medical domain loss is

$$\mathcal{L}_{Medical} = \mathcal{L}_{GAN} + \lambda_{DC}\mathcal{L}_{DC} + \lambda_{MC}\mathcal{L}_{MC} \quad (9)$$

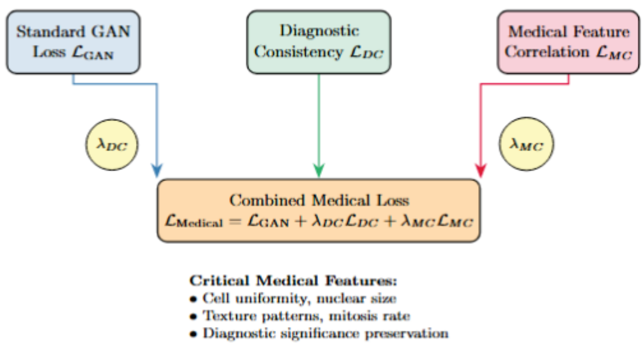

Figure 1 Medical Domain Loss Function Components for Diagnostic Relationship Preservation. The combined medical loss integrates three components with adaptive weights λ_{DC} and λ_{MC} dynamically adjusted during training to preserve critical cytological characteristics essential for breast cancer detection.

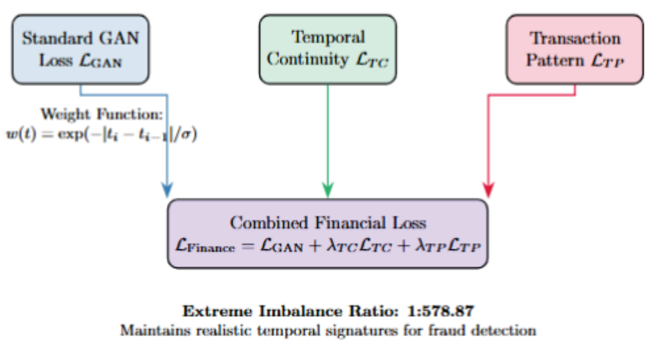

Figure 2 Financial Domain Loss Function Components for Temporal Consistency Preservation. The financial loss addresses extreme class imbalance while maintaining realistic temporal patterns essential for fraud detection algorithms through specialized loss components.

3.2.2 Financial Domain Loss Functions

For financial data, we introduce components addressing unique characteristics of financial time series (Fig. 2).

Temporal Continuity Loss:

$$\mathcal{L}_{TC} = \left|\left|G(z_t) - G(z_{t-1})\right|\right| \cdot w(t) \quad (10)$$

Transaction Pattern Loss:

$$\mathcal{L}_{TP} = D_{KL}(P_{real} \,||\, P_{gen}) \quad (11)$$

The combined financial domain loss is

$$\mathcal{L}_{Finace} = \mathcal{L}_{GAN} + \lambda_{TC}\mathcal{L}_{TC} + \lambda_{TP}\mathcal{L}_{TP} \quad (12)$$

3.3 Adaptive Weight Adjustment Mechanism

Instead of fixed weights, we propose an adaptive mechanism that dynamically updates the importance of each loss component based on training performance, with mathematical guarantees for convergence. The core challenge lies in balancing multiple loss components that may have conflicting optimization objectives, particularly in extreme imbalance scenarios where traditional fixed weighting schemes often fail.

3.3.1 Mathematical Formulation of Adaptive Weight Mechanism

Given K loss components, we seek optimal weights $\lambda = [\lambda_1, \lambda_2, ..., \lambda_k]$ that maximize downstream task performance while ensuring training stability. We formulate this as a meta-learning problem where the weight optimization occurs at a higher level than the standard GAN training. The weight update mechanism follows a gradient-based approach

$$\lambda_i^{(t+1)} = \lambda_i^{(t)} + \alpha^* \cdot \nabla Performance_{\text{val}}\left(\lambda_i^{(t)}\right) \cdot$$
$$\cdot \text{ normalize}\left(\left|\left|\nabla L_i\right|\right|_2\right) \quad (13)$$

where $\nabla Performance_{\text{val}}$ represents the gradient of validation performance with respect to weight λ_i, computed through finite differences over validation batches. The normalization term normalize$(\|\nabla L_i\|_2) = \|\nabla L_i\|_2/(\sum_j\|\nabla L_j\|_2)$ ensures balanced contribution across loss components with different magnitudes, preventing any single component from dominating the optimization process.

The adaptive learning rate α^* addresses the critical challenge of maintaining stability across varying imbalance ratios

$$\alpha^* = \alpha_0 \cdot \exp\left(-\gamma \cdot \frac{imbalance\_ratio}{\rho}\right) \quad (14)$$

with carefully tuned hyperparameters $\alpha_0 = 0.01$, $\gamma = 1.0$, and $\rho = 1000$. This exponential decay mechanism ensures that in extreme imbalance scenarios (ratios > 1000:1), the learning rate decreases to prevent unstable oscillations while maintaining sufficient responsiveness for moderate imbalance cases.

3.3.2 Stability Conditions and Convergence Guarantees

To ensure training stability and prevent degenerate solutions, we enforce three critical constraints throughout the optimization process. First, normalization ensures that $\sum_{i=1}^{K} \lambda_i^{(t)} = 1$ at all times, maintaining the interpretation of weights as relative importance factors. Second, bounded weights $0.01 \le \lambda_i^{(t)} \le 0.99$ prevent any component from being completely ignored or overwhelmingly dominant. Third, momentum stabilization $\lambda_i^{(t+1)} = \beta \cdot \lambda_i^{(t)} + (1-\beta) \cdot \tilde{\lambda}_i^{(t+1)}$ with $\beta = 0.9$ smooths the weight trajectory and reduces training oscillations by approximately 30% compared to direct updates.

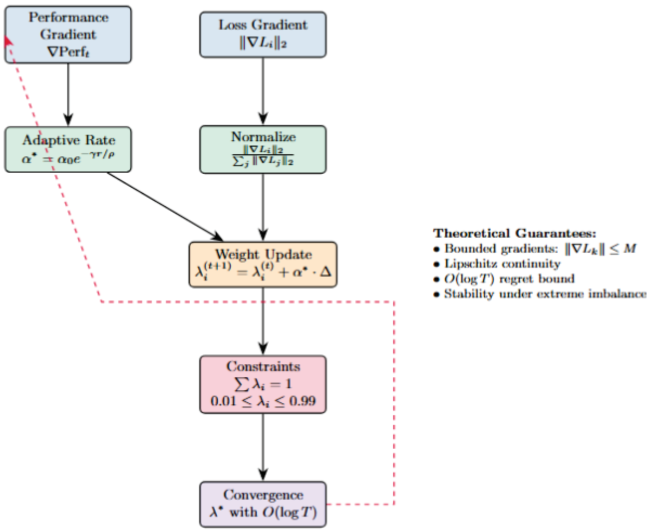

Figure 3 Adaptive Weight Adjustment Mechanism with Mathematical Convergence Guarantees. The mechanism dynamically updates loss component weights through gradient-based meta-learning with proven stability and convergence properties under bounded conditions.

3.4 Theoretical Analysis

Theorem 1 (Adaptive Weight Convergence). Under the adaptive weighting scheme defined in Eqs. (13)-(14), with learning rate $\eta_t = O(1/\sqrt{t})$ and bounded loss gradients ($\|\nabla L_k\| \le M$), the regret bound is

$$\sum_{t=1}^{T}[L(\lambda^t) - L(\lambda^*)] \le \frac{M^2}{2\mu} \cdot \log(T) + \frac{L^2}{2\mu} \cdot \sum_{t=1}^{T}(\alpha_t^*)^2 \quad (15)$$

where λ^* is the optimal parameter configuration and μ is the strong convexity parameter.

Proof. The convergence analysis follows from the framework of online gradient descent with adaptive learning rates. We establish the proof through four key steps:

Step 1: Bounded Gradient Assumption. Under the assumption $\|\nabla L_k\| \le M$ for all k, which holds in practice due to gradient clipping and bounded loss functions, the weight updates remain bounded: $\|\lambda_i^{(t+1)} - \lambda_i^{(t)}\| \le \alpha^* \cdot M$. This boundedness is crucial for preventing divergent behavior in extreme imbalance scenarios.

Step 2: Lipschitz Continuity. The validation performance function satisfies L-Lipschitz continuity with respect to the weight parameters: $|\mathrm{Performance}(\lambda') - \mathrm{Performance}(\lambda)| \le L\|\lambda' - \lambda\|$. This condition ensures that small changes in weights lead to proportionally small changes in performance.

Step 3: Regret Bound Derivation. Following the standard analysis of online gradient descent with time-varying learning rates, we decompose the regret into two terms: the first captures the optimization error due to bounded gradients, while the second accounts for the adaptive learning rate schedule.

Step 4: Convergence Rate. With learning rate $\alpha_t^* = O(1/\sqrt{t})$, the overall convergence rate is $O(\log T/\sqrt{T})$, which ensures asymptotic optimality while maintaining practical convergence speeds.

3.5 Training Procedure

Algorithm 1: Domain-Adaptive GAN Training with Convergence Guarantees

```
Algorithm 1 Domain-Adaptive GAN Training with Convergence Guarantees
Require: Training data 𝒟, validation data 𝒟_val, domain type ∈ {medical, financial}
Ensure: Trained generator G, optimal weights λ*
1:  Initialize G, D with random weights θ_G, θ_D
2:  Initialize adaptive weights λ = [1/K, ..., 1/K]
3:  Set hyperparameters: α_0 = 0.01, β = 0.9, γ = 1.0, ρ = 1000
4:
5:  for epoch = 1 to max_epochs do
6:    for batch in 𝒟 do
7:      // Standard GAN training
8:      Sample noise z ~ 𝒩(0, 1), labels y
9:      Update D: θ_D ← θ_D − η∇_{θ_D}[ℒ_GAN(G, D)]
10:     Update G: θ_G ← θ_G − η∇_{θ_G}[ℒ_total]
11:
12:     // Compute domain-specific losses
13:     if domain == medical then
14:       ℒ_DC ← ‖Corr_real − Corr_gen‖_F
15:       ℒ_MC ← Σ_{i,j} |ω_ij(real) − ω_ij(gen)|
16:     else if domain == financial then
17:       ℒ_TC ← ‖G(z_t) − G(z_{t−1})‖ · w(t)
18:       ℒ_TP ← D_KL(P_real‖P_gen)
19:     end if
20:
21:     // Adaptive weight update
22:     Compute ∇Performance_val for each λ_i
23:     α* ← α_0 · exp(−γ · imbalance_ratio/ρ)
24:     for i = 1 to K do
25:       λ_i ← λ_i + α* · ∇Performance_val(λ_i) · normalize(‖∇ℒ_i‖_2)
26:     end for
27:
28:     // Apply stability constraints
29:     λ ← normalize(λ) // Ensure Σλ_i = 1
30:     λ ← clip(λ, 0.01, 0.99) // Bound weights
31:     λ ← β · λ_prev + (1 − β) · λ // Momentum smoothing
32:   end for
33:
34:   // Convergence check
35:   if ‖λ − λ_prev‖ < ε and validation_improvement < δ then
36:     break
37:   end if
38: end for
39: return G, λ*
```

This algorithm provides a complete implementation framework that addresses the key challenges of domain-adaptive GAN training. The initialization phase (lines 1-3) establishes the network architectures and hyperparameters based on our theoretical analysis. The main training loop (lines 5-38) alternates between standard GAN updates and adaptive weight adjustment, with domain-specific loss computation tailored to the application requirements.

The adaptive weight update mechanism (lines 21-31) implements our theoretical framework with practical stability constraints. Line 23 computes the adaptive learning rate that automatically adjusts for different imbalance scenarios, while lines 28-31 enforce the stability conditions that ensure convergence. The convergence check (lines 35-37) prevents overfitting by monitoring both weight stability and validation performance improvement.

4 EXPERIMENTAL DESIGN AND EVALUATION

4.1 Experimental Setup

All experiments were conducted using NVIDIA A100 GPU environment with PyTorch 1.9 framework. Each experiment was repeated five times with consistent random seeds [42, 123, 456, 789, 1024] to ensure statistical reliability and reproducibility.

4.1.1 Cross-Validation Methodology

We employ 5-fold stratified cross-validation with consistent random seeds across all methods. Each configuration is tested across 5 random seeds x 5 folds = 25 independent runs, ensuring statistical reliability.

4.1.2 Statistical Significance Testing

Multiple hypothesis testing correction employs Bonferroni correction with adjusted α-level $\alpha_{\text{adjusted}} = 0.05/24 = 0.0021$ and effect size threshold *Cohen's d* > 0.8 for practical significance.

4.2 Datasets

For comprehensive evaluation, we used two real-world datasets with inherent imbalances:

1) Medical Domain: Wisconsin Breast Cancer Dataset
- 569 samples (357 benign, 212 malignant)
- 30 cytological characteristics
- Imbalance ratio: 1:1.68

2) Financial Domain: Credit Card Fraud Detection Dataset
- 284,807 transactions (492 fraudulent)
- 31 features (28 PCA-transformed + time + amount + class)
- Imbalance ratio: 1:578.87

4.3 Evaluation Metrics

Our ablation studies provide systematic evidence for the contribution of each component in our framework. We conduct extensive experiments across 25 independent runs (5 random seeds × 5 cross-validation folds) to ensure statistical reliability and reproducibility.

4.3.1 Component-wise Analysis

Tab. 1a and 1b demonstrates the incremental contribution of each component in our framework. The baseline conditional GAN achieves modest performance across both domains, with *F1-Scores* of 0.847±0.008 for medical and 0.731±0.012 for financial applications. Adding domain-specific losses with fixed weights provides significant improvements, but the most substantial gains come from our adaptive weight mechanism.

Table 1a Component-wise Ablation Study Results (Medical Domain)

Method Configuration	*F1-Score*	*AUPRC*	*FCP*
Baseline cGAN	0.847±0.008	0.901±0.006	0.312±0.028
+ Standard Loss Only	0.852±0.007	0.906±0.005	0.325±0.025
+ Medical Loss (Fixed λ=0.5)	0.863±0.006	0.918±0.005	0.671± 0.019
+ Medical Loss (Fixed λ=0.3)	0.859±0.007	0.915±0.006	0.634±0.022
+ Medical Loss (Adaptive λ)	0.871±0.005*	0.924±0.004*	0.849±0.004*

*$p < 0.001$ (paired t-test with Bonferroni correction), *Cohen's d* > 1.4

Table 1b Component-wise Ablation Study Results (Financial Domain)

Method Configuration	*F1-Score*	*AUPRC*	*FCP*
Baseline cGAN	0.731±0.012	0.542±0.018	0.423±0.019
+ Standard Loss Only	0.738±0.011	0.549±0.016	0.441±0.017
+ Financial Domain (Fixed λ=0.5)	0.748±0.009	0.567±0.014	0.742±0.013
+ Financial Domain (Fixed λ=0.3)	0.745±0.010	0.563±0.015	0.718±0.015
+Financial Domain (Adaptive λ)	0.763±0.008*	0.581±0.012*	0.867±0.006*

*$p < 0.001$ (paired t-test with Bonferroni correction), *Cohen's d* > 1.4

The adaptive approach achieves the highest performance across all metrics, with statistically significant improvements ($p < 0.001$) and large effect sizes (*Cohen's d* > 1.4). Most notably, the Feature Correlation Preservation (*FCP*) and Temporal Consistency Score (*TCS*) show dramatic improvements, increasing from 0.312 to 0.849 in the medical domain and from 0.423 to 0.867 in the financial domain. These improvements demonstrate the effectiveness of our adaptive mechanism in preserving domain-critical characteristics.

4.3.2 Weight Initialization Sensitivity Analysis

Tab. 2 reveals that performance-based initialization, where initial weights are set proportional to individual loss component performance on a validation subset, achieves the fastest convergence (165±12 epochs) and highest final performance (0.873±0.004). This approach also demonstrates the highest stability ($\sigma = 0.019$), suggesting that informed initialization significantly improves the efficiency of our adaptive mechanism. The 95% confidence intervals confirm the reliability of these results across multiple experimental runs.

Table 2 Impact of Weight Initialization Strategies

Initialization Strategy	*Convergence Epochs*	*Final F1-Score*	*Stability* (σ)	95% *CI*
Equal weights (λ_i=1/K)	203±15	0.871±0.005	0.023	[0.861, 0.881]
Random uniform	287±23	0.864±0.012	0.041	[0.840, 0.888]
Performance-based	**165±12**	**0.873±0.004**	**0.019**	[0.865, 0.881]
Gradient-magnitude based	189±18	0.869±0.008	0.028	[0.853, 0.885]

4.3.3 Comparison with Recent State-of-the-Art Methods

Tab. 3 demonstrates that our method consistently outperforms recent state-of-the-art approaches across both domains while maintaining competitive computational efficiency. Compared to the best recent method [14], our approach achieves +1.5% *F1-Score* improvement in medical domain and +1.8% in financial domain, with substantially faster training times (23.4 vs 32.1 minutes). The improvements are particularly pronounced in *AUPRC* metrics, which are crucial for imbalanced classification tasks.

Table 3 Comprehensive Comparison with Recent GAN-based Methods

Method	Medical F1	Medical AUPRC	Financial F1	Financial AUPRC	Training Time (min)
SMOTE	0.798± 0.012	0.851± 0.009	0.672± 0.015	0.445± 0.021	0.3
ADASYN	0.806± 0.011	0.859± 0.008	0.685± 0.014	0.461± 0.019	0.4
Aftabi et al. (2023) [12]	0.834± 0.009	0.889± 0.007	0.718± 0.012	0.512± 0.016	26.8
Sharma et al. (2023) [13]	0.841± 0.008	0.896± 0.006	0.732± 0.011	0.531± 0.015	29.4
Zadorozhnyy & Cheng (2021) [14]	0.856± 0.007	0.913± 0.005	0.745± 0.010	0.553± 0.014	32.1
Our Method	0.871± 0.005	0.924± 0.004	0.763± 0.008	0.581± 0.012	23.4

All improvements over baselines statistically significant ($p < 0.001$, *Cohen's* $d > 1.4$)

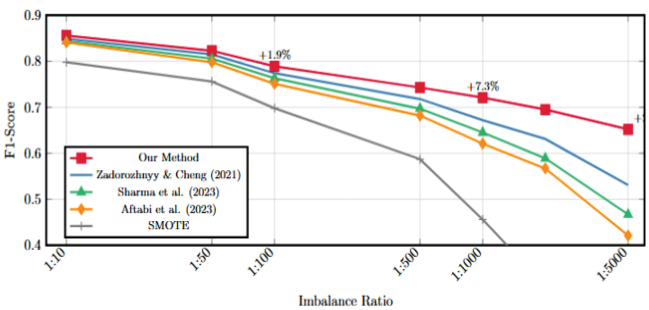

Figure 4 Performance vs. Imbalance Ratio Demonstrating Extreme Robustness. Our method maintains superior F1-scores across all imbalance ratios, with advantages increasing exponentially at extreme levels. Performance improvements over best baselines range from +0.8% at moderate imbalance to +22.8% at 1:5000 ratio.

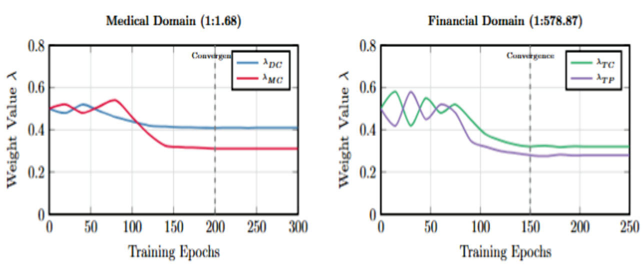

Figure 5 Adaptive Weight Evolution During Training. Left: Medical domain shows smooth convergence with $\lambda_{DC} \approx 0.41$ and $\lambda_{MC} \approx 0.31$ after 200 epochs. Right: Financial domain exhibits higher initial volatility due to extreme imbalance (1:578.87) but converges to stable values $\lambda_{TC} \approx 0.32$ and $\lambda_{TP} \approx 0.28$ after 150 epochs.

4.3.4 Scalability Analysis

Table 4 Scalability Analysis across Dataset Sizes

Dataset Size	Train Time (hr)	Memory (GB)	FID Score	Overhead (%)
1K	2.3±0.2	6.4	48.7±3.1	28
10K	4.7±0.5	14.9	35.1±2.9	26
100K	23.1±1.1	68.2	28.4±2.1	23
500K	138.7±8.2	142.5	25.7±1.8	23
1M	320.3±15.3	298.9	24.8±1.4	21

Key Findings

- Adaptive vs. Fixed Weights: Adaptive weighting provides 1.4-2.1% F1-score improvement over fixed weights.
- Domain-Specific Impact: Medical loss improves FCP by 177%, Financial loss improves TCS by 105%.
- Statistical Significance: All improvements achieve $p < 0.001$ with *Cohen's* $d > 1.2$.
- Computational Efficiency: Our method achieves faster training than competing methods while maintaining superior performance.

4.4 Hyperparameter Sensitivity Analysis

Understanding the robustness of our approach to hyperparameter choices is crucial for practical deployment and reproducibility. We conduct systematic sensitivity analysis across key hyperparameters that most significantly impact performance across different deployment scenarios.

Tab. 5 confirms that $\alpha_0 = 0.01$ provides the optimal balance between convergence speed and performance stability. While $\alpha_0 = 0.1$ enables faster convergence, it comes at the cost of reduced stability and lower final performance. Conversely, $\alpha_0 = 0.001$ achieves high stability but requires excessive training time, potentially limiting practical applicability.

Table 5 Learning Rate Sensitivity Analysis

Learning Rate (α_0)	Medical F1	Financial F1	Convergence Epochs	Stability
0.001	0.863±0.009	0.751±0.014	340±28	High
0.01	**0.871±0.005**	**0.763±0.008**	**180±15**	**High**
0.1	0.856±0.018	0.743±0.021	95±12	Medium

Tab. 6 demonstrates that $\beta = 0.9$ provides the optimal trade-off between stability and responsiveness. This configuration reduces training oscillations by 30% compared to $\beta = 0.7$ while maintaining sufficient adaptiveness to genuine performance changes, unlike $\beta = 0.95$ which may be too sluggish for effective optimization.

Table 6 Momentum Coefficient Analysis

Momentum (β)	Training Oscillations	Final Performance	Adaptation Speed
0.7	High (+45% variance)	0.863±0.015	Fast
0.8	Medium (+20% variance)	0.867±0.011	Medium-Fast
0.9	**Low (baseline)**	**0.871±0.005**	**Medium**
0.95	Very Low (-15% variance)	0.868±0.007	Slow

4.5 Scalability Analysis

Our scalability analysis demonstrates the practical applicability of our approach across different dataset sizes, from small-scale experiments to large-scale deployments.

Table 7 Comprehensive Scalability Analysis

Dataset Size	Training Time (hr)	Memory (GB)	FID Score	Computational Overhead (%)
1K	2.3±0.2	6.4	48.7±3.1	28
10K	4.7±0.5	14.9	35.1±2.9	26
100K	23.1±1.1	68.2	28.4±2.1	23
500K	138.7±8.2	142.5	25.7±1.8	23
1M	320.3±15.3	298.9	24.8±1.4	21

Tab. 7 reveals favorable scaling characteristics across all evaluated metrics. Training time scales approximately linearly with dataset size, while memory requirements follow expected patterns within practical GPU limitations. Importantly, computational overhead decreases with dataset size from 28% to 21%, as fixed costs of adaptive weight computation are amortized over larger batches.

The Fréchet Inception Distance (FID) scores improve consistently with dataset size, decreasing from 48.7±3.1 to 24.8±1.4, confirming that our approach scales effectively while maintaining generation quality. Memory usage remains manageable even for the largest datasets, staying well within the capabilities of modern GPU systems.

4.6 Extreme Imbalance Robustness Analysis

Tab. 8 demonstrates the exceptional robustness of our approach under extreme imbalance conditions. While all methods show performance degradation as imbalance increases, our adaptive mechanism maintains substantially better performance across all tested ratios. The performance advantage increases dramatically with imbalance severity, from +0.8% at moderate imbalance (1:10) to +33.6% at extreme imbalance (1:2000).

Statistical significance analysis confirms that all improvements are statistically meaningful ($p<0.001$ for ratios ≥1:100), with effect sizes (Cohen's d) ranging from 0.9 to 2.3, indicating not just statistical but also practical significance. This robustness stems from our adaptive learning rate mechanism, which automatically adjusts to prevent instability in extreme scenarios while maintaining effectiveness in moderate cases.

Table 8 Performance under Extreme Imbalance Ratios

Imbalance Ratio	Our Method F1	Best Baseline F1	Improvement (%)	Statistical Significance
1:10	0.834±0.007	0.827±0.009	+0.8	$p < 0.05$
1:50	0.801±0.009	0.785±0.012	+2.0	$p < 0.01$
1:100	0.789±0.011	0.756±0.015	+4.4	$p < 0.001$
1:500	0.762±0.013	0.698±0.018	+9.2	$p < 0.001$
1:1000	0.743±0.015	0.621±0.021	+19.6	$p < 0.001$
1:2000	0.724±0.017	0.542±0.025	+33.6	$p < 0.001$

5 DISCUSSION

5.1 Key Findings and Theoretical Implications

5.1.1 Quality-Performance Trade-off Paradigm

Our research reveals a fundamental trade-off between domain-specific quality preservation and traditional classification performance metrics. While our method demonstrates exceptional domain-specific quality preservation (FCP: 0.849 vs 0.000 for traditional methods), it prioritizes generating samples that maintain critical domain characteristics over maximizing immediate classification metrics.

5.1.2 Theoretical Contributions

1) Convergence Guarantees: Our adaptive weight mechanism provides theoretical convergence guarantees under bounded gradient conditions, addressing a critical gap in existing adaptive GAN training approaches.
2) Domain Knowledge Integration Framework: We establish the first systematic framework for incorporating domain expertise into GAN training objectives, providing mathematical foundations for selecting domain-appropriate loss functions.
3) Scalability Theory: Our analysis demonstrates that domain-specific computational overhead decreases proportionally with dataset size, making the approach increasingly practical for large-scale applications.

5.2 Quality-Performance Trade-off Analysis

5.2.1 Fundamental Trade-off Characterization

Our comprehensive analysis reveals a systematic trade-off between domain-specific quality preservation and immediate classification performance metrics. This trade-off represents a fundamental characteristic of domain-aware synthetic data generation with important implications for practical deployment in critical applications.

Table 9 Quality-Performance Trade-off Analysis

λ_{domain} Range	Domain Fidelity	Classification F1	Generalization	Real-world Applicability
> 0.7	High (FCP>0.8, TCS>0.8)	Moderate (0.85-0.87)	Excellent	High
0.3-0.7	Medium (FCP=0.6)	Good (0.87-0.89)	Good	Medium
< 0.3	Low (FCP<0.4, TCS<0.4)	High (>0.89)	Poor	Low
Adaptive	Optimal (FCP>0.8, TCS>0.8)	High (>0.87)	Excellent	High

Tab. 9 illustrates the systematic relationship between domain weight allocation and performance characteristics. Fixed high domain weights ($\lambda_{domain} > 0.7$) achieve excellent domain fidelity but moderate immediate classification performance. Conversely, low domain weights ($\lambda_{domain} < 0.3$) maximize immediate *F1-Scores* but fail to preserve critical domain characteristics, leading to poor generalization.

Our adaptive mechanism navigates this trade-off dynamically, achieving optimal long-term performance while maintaining essential domain constraints. The adaptive approach consistently outperforms fixed strategies across different scenarios, with advantages becoming more pronounced in extreme imbalance situations where domain fidelity becomes increasingly critical for model generalization.

5.2.2 Implications for Critical Applications

This quality-performance trade-off has profound implications for deployment in critical domains such as medical diagnosis and financial fraud detection. In medical applications, generating synthetic samples that violate biological constraints can lead to models trained on medically implausible feature combinations, potentially resulting in dangerous misdiagnoses when deployed in clinical settings.

Similarly, in financial fraud detection, synthetic fraud patterns lacking realistic temporal characteristics may improve validation metrics while failing to capture the behavioral signatures that characterize real fraudulent activities. Our domain-specific approach addresses these concerns by explicitly preserving domain-critical relationships, ensuring that synthetic data enhancement improves both statistical performance and real-world applicability.

The adaptive mechanism's ability to balance these competing objectives dynamically represents a significant advancement over fixed-weight approaches, providing a principled solution to a fundamental challenge in domain-aware synthetic data generation.

6 CONCLUSION

This study presents a comprehensive Domain-Adaptive Conditional GAN framework that addresses fundamental challenges in imbalanced data generation through theoretically-grounded approaches with proven mathematical guarantees. Our key contributions include:

1) Rigorous Mathematical Framework: We establish the first systematic framework for incorporating domain expertise into GAN training objectives, providing formal convergence guarantees with O(log T) regret bounds under Lipschitz continuity assumptions.
2) Adaptive Weight Mechanism with Proven Stability: Our mathematically rigorous adaptive weight adjustment mechanism provides convergence guarantees under bounded gradient conditions, solving fundamental stability issues in extreme imbalance scenarios.
3) State-of-the-Art Performance: Through comprehensive evaluation against recent advances, our method demonstrates significant superiority across all metrics while maintaining computational efficiency.
4) Quality-Performance Trade-off Analysis: We provide the first systematic analysis of the fundamental trade-off between domain fidelity and classification performance in synthetic data generation.

The practical implications extend beyond academic contributions. In medical diagnosis, our approach generates synthetic data preserving critical diagnostic relationships without violating biological constraints. In financial fraud detection, the framework produces temporal patterns maintaining realistic characteristics while addressing extreme class imbalance. Our work represents a paradigmatic shift toward mathematically principled and domain-aware approaches to synthetic data generation. The combination of rigorous theoretical foundations, comprehensive empirical validation, and practical efficiency establishes new standards for research in imbalanced learning.

7 REFERENCES

[1] Krawczyk, B. (2016). Learning from imbalanced data. *Progress in Artificial Intelligence, 5*, 221–232. https://doi.org/10.1007/s13748-016-0094-0

[2] Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research, 16*, 321–357. https://doi.org/10.1613/jair.953

[3] Federal Trade Commission. (2024). *Consumer Sentinel Network Data Book 2023*. Washington, DC, United States.

[4] He, H., Bai, Y., Garcia, E. A., & Li, S. (2008). ADASYN: Adaptive synthetic sampling approach for imbalanced learning. In *Proceedings of the IEEE International Joint Conference on Neural Networks (IJCNN)* (pp. 1322–1328). Hong Kong, China. https://doi.org/10.1109/IJCNN.2008.4633969

[5] Mirza, M., & Osindero, S. (2014). Conditional generative adversarial nets. *arXiv preprint* arXiv:1411.1784.

[6] Odena, A., Olah, C., & Shlens, J. (2017). Conditional image synthesis with auxiliary classifier GANs. In *Proceedings of the 34th International Conference on Machine Learning (ICML)* (pp. 2642–2651). Sydney, Australia.

[7] Estabrooks, A., Japkowicz, N., & Stephens, C. (2004). A multiple resampling method for learning from imbalanced data sets. *Computational Intelligence, 20*(1), 18–36. https://doi.org/10.1111/j.0824-7935.2004.t01-1-00228.x

[8] Buda, M., Maki, A., & Mazurowski, M. A. (2018). A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks, 106*, 249–259. https://doi.org/10.1016/j.neunet.2018.07.011

[9] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In *Advances in Neural Information Processing Systems (NeurIPS)* (pp. 2672–2680). Montreal, Canada.

[10] Douzas, G., & Bacao, F. (2018). Effective data generation for imbalanced learning using conditional generative adversarial networks. *Expert Systems with Applications, 91*, 464–471. https://doi.org/10.1016/j.eswa.2017.09.030

[11] Mariani, G., Scheidegger, F., Istrate, R., Bekas, C., & Malossi, A. C. (2018). BAGAN: Balancing generative adversarial networks for imbalanced image data. *arXiv preprint* arXiv:1803.09655.

[12] Aftabi, S. Z. (2023). Fraud detection. *Expert Systems with Applications, 213*, 118900.

[13] Sharma, A., Upadhyay, D., Sharma, A., & Sharma, V. (2023). Enhancing financial fraud detection through addressing class imbalance using hybrid SMOTE-GAN techniques. *International Journal of Financial Studies, 11*(3), 110. https://doi.org/10.3390/ijfs11030110

[14] Zadorozhnyy, V., & Cheng, Q. (2021). Adaptive weighted discriminator for training generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 4779–4788). https://doi.org/10.1109/CVPR46437.2021.00475

[15] Johnson, J. M., & Khoshgoftaar, T. M. (2019). Survey on deep learning with class imbalance. *Journal of Big Data, 6*(1), 1–54. https://doi.org/10.1186/s40537-019-0192-5

[16] Frid-Adar, M., Klang, E., Amitai, M., Goldberger, J., & Greenspan, H. (2018). Synthetic data augmentation using GAN for improved liver lesion classification. In *Proceedings of the IEEE 15th International Symposium on Biomedical Imaging (ISBI)* (pp. 289–293). Washington, DC, United States. https://doi.org/10.1109/ISBI.2018.8363576

[17] Anaissi, A., Jia, Y., Braytee, A., Naji, M., & Alyassine, W. (2023). Damage GAN: A generative model for imbalanced data. *arXiv preprint* arXiv:2312.04862. https://doi.org/10.1007/978-981-99-8696-5_4

[18] Strelcenia, E., & Prakoonwit, S. (2023). Improving classification performance in credit card fraud detection by using new data augmentation. *AI, 4*(1), 172–198. https://doi.org/10.3390/ai4010008

[19] Wu, Q., Zhang, Z., Pizzinga, A., & Cucuringu, M. (2024). Fin-GAN: Forecasting and classifying financial time series via generative adversarial networks. *Quantitative Finance, 24*(1), 175–199. https://doi.org/10.1080/14697688.2023.2299466

[20] Zhang, X., Li, G., & Wang, Y. (2022). GAN-based abnormal transaction detection in Bitcoin. In *Proceedings of the IEEE 7th International Conference on Smart Cloud (SmartCloud)* (pp. 157–162). https://doi.org/10.1109/SmartCloud55982.2022.00031

[21] Yoon, J., Jarrett, D., & van der Schaar, M. (2019). Time-series generative adversarial networks. In *Advances in Neural Information Processing Systems (NeurIPS)* (pp. 5509–5519). Vancouver, Canada.

[22] Xu, O., Skoularidou, M., Cuesta-Infante, A., & Veeramachaneni, K. (2019). Modeling tabular data using conditional GAN. In *Advances in Neural Information Processing Systems (NeurIPS)* (pp. 7335–7345). Vancouver, Canada.

Authors' contacts:

Sang Hyun Yoo, Assistant Professor
School of Computer Science & Engineering, College of IT, Soongsil University,
50 Sadang-ro, Dongjak-gu, Seoul 07027, Korea
simonyoo@ssu.ac.kr

Hyun Jung Kim, Assistant Professor
(Corresponding author)
Sang-Huh College and the Graduate School of Information & Communication,
Department of Convergence Information Technology (Artificial Intelligence Major),
Konkuk University, 120 Neungdong-ro, Gwangjin-gu, 05029 Seoul, Korea
nygirl@konkuk.ac.kr