

An Optimization Framework for Digital Advertising Budget Allocation Based on Multi-Objective Distributed Federated Learning and Cloud Computing

Saeid Farahabadi

Department of Industrial Engineering, Ro.C., Islamic Azad University

Ameneh Farahani

`a_farahani@azad.ac.ir`

Department of Industrial Engineering, Ooj Institute of Higher Education

Mehdi Keramatpour

Department of Industrial Engineering, Ro.C., Islamic Azad University

Research Article

Keywords: Digital advertising, budget allocation, federated learning, multi-objective optimization, cloud computing

Posted Date: October 13th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11098099/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

An Optimization Framework for Digital Advertising Budget Allocation Based on Multi-Objective Distributed Federated Learning and Cloud Computing

Saeid Farahabadi¹, Ameneh Farahani^{2*}, Mehdi Keramatpour³

¹ Department of Industrial Engineering, Ro.C., Islamic Azad University, Roudehen, Iran

^{2*} Department of Industrial Engineering, Ooj Institute of Higher Education, Qazvin, Iran.

³ Department of Industrial Engineering, Ro.C., Islamic Azad University, Roudehen, Iran

Saeid.farahabadi@iau.ac.ir

a_farahani@azad.ac.ir

m.keramatpour@iau.ac.ir

Abstract- The rapid growth of digital advertising has made optimal budget allocation across heterogeneous platforms a critical challenge, requiring simultaneous optimization of ROAS, CTR, and cost efficiency while preserving user privacy and ensuring real-time responsiveness. Existing centralized approaches fail to address these challenges due to privacy constraints, data heterogeneity, and scalability limitations. This paper proposes a novel optimization framework based on multi-objective distributed federated learning and cloud computing, introducing two key innovations: (1) Federated Distributional Adaptive Q-Reinforcement Learning (Fed-DAQRL), which learns distributed Q-value representations to explicitly model uncertainty and risk; and (2) Federated Adaptive NSGA-II (Fed-NSGA-AD), which performs multi-objective optimization of ROAS, CTR, cost, spillover, and carryover effects through adaptive cloud-based aggregation. Both algorithms are rigorously formalized through a Decentralized Markov Decision Process and categorical distributional Q-value representation, supported by formal convergence and differential privacy theorems. The framework incorporates dual-layer privacy protection combining secure aggregation with differential privacy. Extensive experiments on ADAC and CTRD datasets demonstrate 99.4% and 99.1% accuracy, representing improvements of 3.4% over Fed-GANCC, 3.8% over AWE-DPFL, 6.7% over robust optimization, and 5.2% over multi-objective federated learning, with loss reductions of 81.7% and 78.9%. Statistical significance testing confirms all improvements are significant ($p < 0.001$, Cohen's $d > 3.4$). Economic analysis shows a ROAS of 4.85, a CPA of \$12.40, and an ROI of 385%. Scalability analysis maintains 98.8% accuracy at 1,000 nodes with robustness under 30% node failure. Ablation studies confirm each component's critical contribution.

Keywords: Digital advertising, budget allocation, federated learning, multi-objective optimization, cloud computing.

1 INTRODUCTION

The digital advertising market has experienced explosive growth in recent years, with organizations increasingly shifting their marketing budgets toward online channels as global spending continues its upward trajectory. This rapid expansion has been accompanied by unprecedented complexity in advertising ecosystems, characterized by diverse platforms including search engines, social media networks, display advertising, and video streaming services. In such an environment, optimal budget allocation across channels, campaigns, and audience segments has emerged as one of the most critical decision-making challenges facing modern organizations [1, 2]. This complexity arises from multiple interrelated factors: the heterogeneity of advertising platforms, the dynamic and non-stationary nature of user behavior, and the necessity to process massive volumes of data within increasingly shorter time windows [3].

Traditional approaches to advertising budget allocation, which predominantly rely on centralized analyses and static models based on historical data, have proven inadequate in addressing these challenges [3]. Such methods not only suffer from insufficient accuracy and adaptability but also face serious limitations in scalability and real-time responsiveness. The fundamental flaw lies in their inability to capture the dynamic, multi-channel interactions that characterize modern advertising campaigns. Furthermore, these centralized approaches require aggregating raw user data in central data centers, creating significant privacy and security vulnerabilities [4].

The exponential growth in data generated through digital advertising processes has made cloud computing infrastructure virtually indispensable for storage, processing, and computational management [5]. Cloud computing, by providing scalable computational resources, enables the execution of complex learning algorithms, coordination among distributed components, and simultaneous responses to multiple decision requests. However, exclusive reliance on centralized cloud architectures necessitates the aggregation of raw user data in data centers, which conflicts with stringent privacy regulations and raises serious concerns regarding data security and legal compliance [6].

In this context, federated learning has emerged as a revolutionary distributed learning paradigm that enables collaborative model training without transferring raw data, instead exchanging only model parameters or gradients [7]. Recent research has demonstrated that federated learning can be effectively applied to smart advertising, simultaneously preserving user privacy while improving prediction accuracy and decision-making capabilities [8]. The core advantage lies in keeping data localized while extracting shared knowledge from distributed sources, making it particularly suitable for scenarios where data cannot be centralized due to privacy constraints or competitive considerations. However, the efficient implementation of federated learning at scale remains challenging without robust cloud infrastructure for coordination, model aggregation, and learning process management [9].

In modern architectures, cloud computing plays the role of a central coordinator, responsible for scheduling federated learning rounds, aggregating updates, enforcing privacy policies, and executing optimization modules for budget allocation. This integration enables centralized management of communication challenges, network latency, and heterogeneity in the computational capabilities of participating nodes [5]. Despite these advances, data heterogeneity across platforms and campaigns remains one of the most significant obstacles to achieving stable and accurate federated models for advertising budget allocation. Recent research has addressed this challenge through methods such as adaptive weighting and personalized federated learning, yet these approaches have not been fully explored in the context of budget optimization [8].

Furthermore, advertising budget allocation is inherently a dynamic, multi-criteria decision-making problem that must simultaneously optimize metrics including ROAS, Cost Per Acquisition (CPA), and CTR under conditions of uncertainty [9]. Emerging research suggests that hybrid bandit algorithms and reinforcement learning can effectively balance exploration and exploitation in multi-channel advertising environments [10]. However, implementing such algorithms at real-world scale requires a cloud platform capable of guaranteeing complex computations, state storage, and rapid decision-making responses.

From the perspective of security and privacy, while federated learning eliminates raw data transfer, it remains vulnerable to attacks such as gradient inference or information leakage from model parameters [11]. In this domain, cloud computing plays a crucial role in implementing and managing complementary mechanisms including differential privacy, access control, and security monitoring. Additionally, the computational and communication costs of federated learning and optimization algorithms must be managed such that the economic benefits derived from improved advertising performance are not overshadowed. This requires designing efficient resource allocation and scheduling models at the cloud layer [6].

Recent developments in distributed optimization and cloud-edge architectures have shown that integrated approaches can significantly improve computational efficiency, scalability, and responsiveness [12]. Moreover, state-of-the-art budget optimization methods developed in recent years leverage advanced learning algorithms and dynamic decision-making to provide practical solutions for managing budget allocation in multi-channel, dynamic environments [13, 14]. Yet, despite these advancements, no comprehensive framework has been proposed that positions cloud computing as the primary execution and processing substrate, employs distributed multi-objective federated learning for digital advertising budget optimization, and simultaneously satisfies privacy requirements, scalability demands, dynamic environmental conditions, and economic efficiency considerations [15, 16].

Given the above, the combination of federated learning—for learning from fragmented data generated by heterogeneous publishers and platforms—with the processing power, storage, and coordination capabilities provided by cloud computing for aggregating distributed models, executing multi-objective federated algorithms, and managing computational resources in real-time, presents a natural and highly innovative research pathway. This integration enables the design of an optimization model for digital advertising budget allocation that can simultaneously improve advertising performance while fulfilling requirements for privacy preservation, scalability, real-time responsiveness, and balance between individual and system-wide objectives.

1.1 Contributions

This paper makes several novel contributions to the field of digital advertising optimization:

1. We propose a comprehensive framework that synergistically combines federated learning, multi-objective optimization, and cloud computing for digital advertising budget allocation. Unlike prior work that treats these components in isolation, our framework provides an end-to-end operational solution that preserves privacy while enabling scalable, real-time decision-making.
2. We propose a novel Fed-DAQRL reinforcement learning algorithm that learns distributed Q-values across participating nodes and aggregates them through an adaptive cloud-based fusion mechanism. This approach explicitly models uncertainty and risk in budget allocation decisions, addressing a critical limitation of classical methods that rely solely on expected return estimates.
3. We propose an enhanced version of the NSGA-II multi-objective evolutionary algorithm that operates in a federated manner with adaptive mechanisms. The algorithm simultaneously optimizes multiple advertising performance metrics including ROAS, CTR, cost efficiency, spillover effects, and carryover effects, while enabling local population evolution and secure cloud-based aggregation.
4. We propose secure aggregation and differential privacy mechanisms within the federated framework, creating a two-tier security architecture that protects against both communication-based attacks and model inference attacks.
5. We design a cloud-native execution environment that supports parallel processing, elastic resource allocation, and real-time responsiveness, enabling practical deployment in large-scale, multi-channel advertising scenarios.

The remainder of this paper is organized as follows: Section 2 reviews related work. Section 3 presents the proposed framework. Section 4 describes the experimental setup. Section 5 discusses the implications, limitations, and future research directions.

2 RELATED WORK

Su et al. proposed the Fed-GANCC framework to improve the accuracy of digital advertising in federated learning environments, utilizing Generative Adversarial Networks and group clustering to overcome the challenge of limited local data. The main problem addressed in their paper is the weak performance of local models under data sparsity conditions, and the proposed solution increases prediction accuracy through synthetic data generation and better convergence of federated models. The key innovation of this research lies in combining graph attention networks with federated learning; however, its focus is primarily on prediction and ad targeting rather than direct budget allocation optimization [17].

Qian, W et al. proposed a problem in federated learning and proposed FDP-FL with a flexible privacy budget allocation mechanism. They demonstrated that uniform privacy budget allocation across all users is suboptimal and that dynamic adjustment can achieve a better balance between accuracy and privacy. The main innovation of their work lies in personalizing the privacy level, but its direct application to economic decision-making and advertising budget allocation has not been explicitly explored [18].

Liu et al. proposed the core challenge in e-commerce search advertising of establishing accurate matching between user queries and appropriate keywords for ad display, especially for long-tail queries that do not match seller-provided keywords and result in reduced ad retrieval. Existing methods such as query log mining or adaptive vectors cannot simultaneously optimize semantic relevance, rewritten authenticity, and advertising revenue efficiency. To address these limitations, the paper introduces a multi-objective aligned bidword generation model comprising a generator, discriminator, and preference alignment module. Within this framework, a discriminator is designed to simultaneously optimize three key objectives: query-keyword relevance, authenticity of generated keywords, and advertising value. The generator is then trained using discriminator feedback to maximize the collective impact of these objectives. Extensive offline and online experiments demonstrate that the proposed model outperforms previous methods and has generated significant commercial value after deployment on the platform, indicating the effectiveness and robustness of the approach [19].

Chen et al. proposed the multi-channel budget allocation problem in the presence of spillover and carryover effects. They demonstrated that optimal decision-making requires understanding temporal and cross-channel interactions

among advertising channels. The main innovation of this paper lies in the dynamic modeling of inter-channel effects; however, challenges related to data privacy and decentralized architectures were not considered [20].

Wei et al. proposed FedAds, a standard benchmark for conversion rate estimation in vertical federated learning for advertising. The main problem is collaboration among players with different data features, and their work provides a valid framework for evaluating federated methods. The innovation lies in standardizing federated learning evaluation in advertising, and this research provides an important empirical foundation for developing subsequent optimization models [21].

Su et al. proposed a federated learning framework based on reinforcement learning and knowledge distillation for targeted advertising. The main problem is improving advertising performance under local data constraints. The combination of federated learning with knowledge distillation is among the innovations of this paper, but its focus is more on ad targeting than on macro-level optimal budget allocation [22].

Xu et al. proposed the online ROI allocation problem in search advertising and presented a constrained optimization model for real-time decision-making. This research focuses on rapid response to market changes and demonstrates that online decision-making can improve capital return. However, the assumption of centralized data access remains [23].

Zhang and Li proposed the use of federated learning for optimizing multi-scenario ad targeting in digital advertising. Their research demonstrates that federated learning can increase Return on Investment (ROI) without violating privacy. The innovation lies in applying federated learning to multi-scenario environments, although the dynamic budget allocation problem was not directly modeled [24].

Miralles-Pechuán et al. proposed a genetic algorithm-based method for optimizing display advertising campaigns. The main problem is searching a large decision space, and the genetic algorithm is employed as the primary tool. This work represents one of the earliest attempts to use evolutionary algorithms in advertising, but its structure is entirely centralized [25].

Romero Leguina et al. proposed the user path in digital marketing and highlighted the importance of correct channel attribution. This paper provides a clear conceptual framework for understanding user interaction with channels, but lacks a quantitative optimization model for budget allocation [26].

Bergemann et al. proposed the budget-constrained bidding problem from algorithmic and data-driven perspectives. They demonstrated that designing optimal bidding policies requires simultaneous understanding of market behavior and budget constraints. The innovation lies in connecting economic theory and data-driven approaches, but the paper does not address federated learning [27].

Stram et al. proposed the Mystique system for budget pacing in online advertising. This system improves campaign performance through gradual budget consumption control. The innovation lies in designing an operational and scalable system, although distributed learning is not incorporated [28].

Kanchan et al. proposed a secure aggregation mechanism based on group signatures for federated learning, aimed at protecting participant privacy and preventing the disclosure of identity and local data in the distributed training process. In this method, only aggregated and authenticated results are made available to the central server, minimizing the possibility of identification or manipulation. This security framework, particularly for federated advertising systems dealing with sensitive user data across various publishers, provides a reliable infrastructure and can serve as a complementary security layer in cloud-federated architectures [29].

Zhou et al. proposed a verifiable secure aggregation method based on secret sharing in federated learning that, in addition to preserving privacy, enables verification of the aggregation process integrity. The main innovation of this research lies in increasing transparency and trustworthiness of federated systems, allowing participants to ensure results have not been tampered with. Although the direct application of this method to digital advertising has not been explored, the proposed framework can be highly valuable for budget allocation scenarios and optimization of multi-stakeholder advertising models [30].

Seyghaly et al. proposed a comprehensive layered architecture for smart advertising based on federated learning, systematically describing the interaction among data collection layers, local training, cloud aggregation, and advertising decision-making. This paper places special emphasis on the coexistence of privacy, scalability, and efficiency, demonstrating how intelligent advertising models can be implemented in distributed environments without raw data transfer. This architectural framework provides a suitable theoretical and practical foundation for developing the current research [31].

Guo et al. proposed a federated learning framework with differential privacy based on Fast Fourier Transform (FFT), aimed at improving model accuracy and reducing computational and communication costs. By transferring the noise addition and parameter combination processes to the frequency domain, the proposed method significantly reduces computational overhead while maintaining privacy guarantees. This innovation is particularly important for scalable advertising systems requiring fast and secure decision-making [32].

Cheng et al. proposed a dynamic programming model for optimizing marketing resource allocation in multi-channel environments within the context of digital transformation. This model attempts to improve budget allocation decision-making by considering market dynamics and temporal dependencies. However, the implementation remains based on centralized architecture and has limitations in scalability, privacy preservation, and multi-stakeholder interaction [33].

Li, M. et al. proposed a hybrid algorithm based on parallel genetic algorithm and simulated annealing for optimal marketing resource allocation. This research demonstrates that combining two metaheuristic methods can improve convergence and solution quality. Despite favorable results, the proposed framework lacks support for distributed architectures and federated learning, being more suitable for centralized scenarios [34].

Souissi et al. proposed a hybrid approach combining fuzzy logic and mathematical programming for evaluating and optimizing online advertising effectiveness. This method enables modeling of uncertainty, decision-maker preferences, and multiple performance criteria, providing realistic compromise solutions for complex multi-criteria digital advertising environments. These characteristics make it highly suitable for complex multi-criteria digital advertising environments [35].

Kuppusamy et al. proposed bipolar Pythagorean fuzzy logic to decision-making problems in digital marketing. The main innovation of this research lies in its ability to simultaneously model positive and negative perspectives of decision-makers and complex preferences. This approach can be useful for analyzing campaign effectiveness and selecting advertising strategies under high uncertainty [36].

Sedlářová Nehézová et al. proposed a robust optimization approach for online advertising budget allocation, aimed at maintaining system performance against market fluctuations and parameter uncertainty. This research demonstrates that robust models can provide more stable and reliable decisions over long-term horizons, although the primary focus is on centralized environments [37].

Shen et al. proposed the concept of multi-objective federated learning with a focus on balancing overall system performance and individual fairness among participants. This research provides an important theoretical framework for combining federated learning with multi-objective optimization and can serve as a foundation for developing fair and efficient budget allocation models in digital advertising [38].

Li, M et al. proposed a multi-channel budget allocation model based on reinforcement learning and an improved differential evolution algorithm aimed at improving advertising ROI. The combination of reinforcement learning with evolutionary algorithms enables dynamic and adaptive decision-making, but the overall framework still relies on centralized architecture, and the full potential of federated learning capabilities has not been exploited [39]. Table 1 provides a summary of recent research.

Table 1: Summary of Recent Research

Authors	Year	Paper Objective	Advantages	Disadvantages	Limitations
Su, C., et al. [17]	2024	Proposing Fed-GANCC federated learning framework using GANs and group clustering for precise advertising	Improved prediction accuracy, local data preservation, innovation in combining GAN and FL	Limited focus on specific digital advertising, higher computational complexity	No direct investigation of budget allocation optimization
Qian, W., et al. [18]	2024	Proposing FDP-FL with flexible privacy budget allocation in federated learning	Privacy preservation, dynamic optimization of privacy level	Limited focus on advertising performance optimization	No direct modeling of advertising budget allocation
Liu, Z., et al. [19]	2025	Multi-objective bidword generation model for e-commerce search advertising	Alignment of different advertiser objectives, overall performance improvement	Centralized implementation, lack of privacy preservation	No use of federated learning
Chen, H., et al. [20]	2025	Multi-channel budget allocation considering spillover and carryover effects	Modeling inter-channel interactions, dynamic optimization	Limited focus on privacy	No decentralized architecture
Wei, P., et al. [21]	2023	Standard benchmark for conversion rate estimation with vertical FL	Standardized evaluation, collaboration among players with different features	Limited to vertical data	Limited application in budget allocation optimization
Su, C., et al. [22]	2023	Federated learning framework based on RL and knowledge distillation for targeted advertising	Improved ad targeting, data preservation	Limited focus on budget allocation	Complexity of combining RL and FL
Xu, W., et al. [23]	2024	Online ROI allocation in sponsored search	Real-time decision-making, improved ROI	Assumption of centralized data access	No attention to distributed environments
Zhang, K., & Li, P. [24]	2024	Optimizing multi-scenario ad targeting with FL	Privacy preservation, increased ROI	No direct modeling of budget allocation	No real-time implementation
Miralles-Pechuán, L., et al. [25]	2018	Optimizing display advertising campaigns with genetic algorithm	Large decision space search, simple and efficient	Centralized structure	No flexibility in distributed environments
Romero Leguina, J., et al. [26]	2020	User path analysis and attribution	Clear conceptual framework, understanding user interaction	No quantitative model	No direct application in budget optimization
Bergemann, D., et al. [27]	2025	Investigating budget-constrained bidding	Data-driven and algorithmic combination, bidding optimization	No use of FL	No focus on FL and privacy
Stram, R., et al. [28]	2024	Mystique system for budget pacing	Gradual budget control, scalable	No distributed learning	Limitations in FL performance
Kanchan, S., et al. [29]	2024	Secure aggregation based on group signatures in FL	Enhanced security and privacy preservation	Computational complexity	No direct application in budget allocation
Zhou, S., et al. [30]	2025	Verifiable secure aggregation based on secret sharing	High trustworthiness, security	Limited to model security	Limited application in advertising
Seyghaly, R., et al. [31]	2024	Comprehensive architecture for FL-based smart advertising	Layer interaction, system design foundation	No optimization algorithm	Limited to architecture

Guo, S., et al. [32]	2024	FL with DP based on FFT for efficient combining	Reduced computational overhead, privacy preservation	Limited to laboratory environment	No direct practical application
Cheng, Q. [33]	2025	Optimizing multi-channel marketing resource allocation with dynamic programming	Adaptability to digital transformation	Centralized	No data distribution
Li, M. [34]	2024	Marketing resource allocation with parallel GA and simulated annealing	Combining two metaheuristic methods, high accuracy	Centralized	No use of distributed environments
Souissi, B., et al. [35]	2024	Combining fuzzy VIKOR and mathematical programming for advertising effectiveness	Modeling decision uncertainty	No distributed implementation	Limitations in scalability
Kuppusamy, V., et al. [36]	2024	Bipolar Pythagorean fuzzy logic for marketing decision-making	Modeling complex preferences	Computational complexity	No combination with FL
Sedlářová Nehézová, T., et al. [37]	2025	Robust optimization approach for online ad budget allocation	Resistance to fluctuations, long-term decision-making	High computational complexity	No combination with FL
Shen, Y., et al. [38]	2025	Multi-objective FL: balancing global performance and individual fairness	Balance between global and local objectives, FL multi-objective application	Computational complexity	Limited implementation
Li, M., et al. [39]	2024	Multi-channel budget allocation with RL and improved DE	Combining RL and DE, increased ROI	Centralized	No use of distributed environments

The systematic review of prior research reveals several critical gaps. First, most existing studies rely on centralized architectures that require raw user data aggregation, raising significant privacy and scalability concerns while failing to handle distributed, heterogeneous data from real-world multi-platform environments. Second, although recent work has applied federated learning to advertising, the primary focus has been on prediction and privacy, while dynamic multi-channel budget allocation the core economic decision problem remains largely unaddressed. Third, key advertising phenomena such as cross-channel spillover effects, temporal carryover effects, and market uncertainty are either simplified or ignored in most models. Fourth, existing methods fail to capture the distributional nature of returns and associated risks, relying instead on point estimates that inadequately represent decision uncertainty. Fifth, the integration of privacy-preserving mechanisms with multi-objective optimization in federated settings remains underdeveloped.

To address these gaps, our proposed framework introduces several key innovations: an integrated federated-cloud architecture that enables collaborative learning without raw data transfer; Fed-DAQRL, a distributional reinforcement learning algorithm that explicitly models uncertainty and risk in budget decisions; Fed-NSGA-AD, an adaptive multi-objective optimization method that simultaneously optimizes ROAS, CTR, cost, spillover, and carryover effects; dual-layer privacy protection combining secure aggregation with differential privacy; and an end-to-end cloud-native implementation that transforms the framework from theoretical concept to deployable industrial solution. Together, these contributions provide a comprehensive, scalable, and privacy-preserving approach to digital advertising budget optimization that overcomes the fundamental limitations of prior work.

3 THE PROPOSED FRAMEWORK

This section presents the detailed specifications of the proposed framework for digital advertising budget optimization based on federated learning and cloud computing.

3.1 Architecture of the Proposed Framework

Figure 1 illustrates the proposed architecture for the federated-cloud based budget optimization system.




Fig. 1. System Architecture of the Proposed Framework

As shown in Figure 1, the proposed framework is designed around six main phases, each playing a critical role in establishing a practical and scalable framework. In the first phase, advertising data is collected from authentic and reliable sources, including public data from major corporations, globally validated datasets, and transactional data

from digital publishers. The collected data is stored in the cloud environment to ensure easy accessibility, scalability, and data security. At this stage, data undergoes cleaning and preprocessing to verify quality and accuracy, preparing it for analysis and modeling. Data cleaning encompasses removal of missing values, correction of erroneous data, value normalization, and transformation into formats suitable for machine learning models. Concurrently, the Fed-DAQRL model is designed to keep data localized, transmitting only processed results between platforms—a feature that preserves user privacy and reduces legal and ethical risks.

In this phase, the federated learning model is designed to predict the effectiveness of each advertising channel without requiring raw data transfer. Each advertising platform or publisher trains its local model based on available data, transmitting only weights or computational model parameters to the central server in the cloud. The Fed-DAQRL algorithm, leveraging distributed Q-Value representation, provides more accurate predictions of budget effects across different channels and enables multi-objective, distributed decision-making. The resulting global model achieves superior understanding of consumer behavior and channel effectiveness, enabling more precise prediction of Return on Investment (ROI) and conversion rates. The model is designed to operate dynamically with regular update capabilities, incorporating changes in user behavior and market conditions.

Following channel effectiveness prediction, budget allocation is performed using the Fed-NSGA-AD algorithm. This enhanced version of NSGA-II features an improved cost function and next-generation selection mechanism, enabling multi-objective optimization considering ROI, CTR, cost, spillover effects, and advertising carryover effects. At this stage, two key mechanisms are employed: Secure Aggregation ensures that local model weights are combined without disclosing raw user data, with only aggregated information transmitted to the central server. Subsequently, DP-SGD adds random noise to gradients, making extraction of personal user information impossible even if the global model is accessed. These mechanisms are directly integrated with Fed-DAQRL and Fed-NSGA-AD to simultaneously guarantee distributed Q-Value representation, multi-objective decisions, and data protection.

Model evaluation is conducted at three levels:

- Backtesting to assess model performance on historical data,
- Off-Policy Evaluation to simulate new decisions prior to real-world execution, and
- Performance assessment to measure model efficiency and improvement in key metrics such as ROI and conversion rates.

These evaluations ensure that Fed-DAQRL and Fed-NSGA-AD models are accurate, reliable, and operational for deployment in real advertising campaigns.

The final model is capable of automatically and optimally distributing advertising budgets across platforms, reducing costs, increasing ROI, and simultaneously protecting data security and privacy. The utilization of Fed-DAQRL as the federated-distributed reinforcement learning core and Fed-NSGA-AD as the advanced multi-objective optimization algorithm enables dynamic, multi-objective, and distributed budget allocation, transforming the final model into an operational and scalable framework for the digital advertising industry.

This framework serves as both a scientific and practical tool with direct applicability in digital advertising, empowering brands and companies to make optimal, rapid, and data-driven decisions. Consequently, the proposed method not only addresses practical issues of digital advertising budget allocation but also satisfies legal and security requirements related to user data while offering substantial scalability and development capabilities.

3.2 Conceptual Overview of the Framework Phases

This subsection provides a comprehensive description of each phase of the proposed framework. The key phases of the proposed approach are as follows:

3.2.1 Data Collection and Preprocessing

The proposed framework for digital advertising budget optimization is designed based on federated learning and cloud computing. Figure 2 illustrates the data collection and preprocessing stages for input data to the proposed system.




Fig. 2. Data Collection to Preprocessing of Input Data

As shown in the Figure 2, the data collection and preprocessing phase comprises several components detailed below:

Data Collection: At this stage, advertising datasets are collected from various advertising platforms (e.g., Google Ads, Instagram Ads, Facebook Ads, and other networks). These data include:

- **Public Data from Major Companies:** Information regarding advertising campaigns and their results, typically published publicly by companies.
- **Globally Validated Datasets:** Datasets officially validated by authoritative sources, containing comprehensive and clear information about the digital advertising market.
- **Transactional Data from Digital Publishers:** Real-world transaction and user behavior data that provides deeper insights into advertising performance.

These data are stored in the cloud environment to ensure easy accessibility, scalability, and data security. Cloud computing provides the necessary capabilities for managing large data volumes and ensuring adaptability to changing market requirements.

Data Preprocessing: Following data collection, the preprocessing phase begins, tasked with preparing data for analysis and modeling. This phase includes:

1. **Data Cleaning:** Missing values are identified and removed, and erroneous data is corrected. Value normalization is performed to standardize data for comparability. Eq. 1 shows the process of removing samples containing outlier values:

$$D' = D - [x_i \in D \mid x_i \text{ Contain Missing - Values}] \quad (1)$$

Where D represents the original dataset and D' represents the dataset cleaned of missing values. After removing outlier (missing) values, erroneous data within the dataset is identified and corrected using Eq. 2:

$$D' = \{y_i \mid y_i \in D \text{ AND } y_i = \text{True}\} \quad (2)$$

Where y_i represents values requiring correction, and only correct values remain in the new dataset D' .

Following erroneous data correction, value normalization is performed. Various methods can be employed for data normalization. One common approach is min-max normalization between 0 and 1, performed using Eq. 3:

$$x' = \frac{x - X_{\min}}{X_{\max} - X_{\min}} \quad (3)$$

Where x is the original value, x' is the normalized value, and XX represents the set of all values. Subsequently, standard scaling is applied, computed using Eq. 4:

$$x' = \frac{x - \mu}{\sigma} \quad (4)$$

Where x' represents standardized values, μ denotes the mean, and σ denotes the standard deviation of the dataset.

2. **Data Transformation:** Data is converted into formats suitable for machine learning models. This includes feature analysis and selection of appropriate specifications to enhance algorithm accuracy.

Finally, all standardized data is transformed into appropriate format using Eq. 5:

$$D'_{ij} = \begin{cases} \text{if } x_i \in C \text{ THEN } j = 1 \\ \text{Else } j = 0 \end{cases} \quad (5)$$

Where C represents the category. These equations and methods serve as essential components of the data cleaning process, helping to prepare data ready for analysis and modeling.

3.2.2 Federated Learning Model Design

Figure 3 illustrates the federated learning model design framework.




Fig. 3. Federated Learning Model Design Framework

As shown in the Figure 3, the federated learning model is designed to predict the effectiveness of each advertising channel without requiring raw and sensitive data transfer. This approach not only preserves user privacy but also enhances decision-making capability and accuracy in digital advertising. This process comprises the following components:

In this phase, each advertising platform trains a local model on its private data, and only model parameters are transmitted to the cloud server, preserving user privacy. The proposed Fed-DAQRL algorithm serves as the learning core, leveraging a distributional Q-value representation to explicitly model uncertainty in channel effectiveness prediction. The global model, formed through adaptive cloud-based aggregation, achieves enhanced understanding of consumer behavior and enables accurate ROI and conversion prediction.

A complete mathematical formulation of Fed-DAQRL, including the Dec-MDP problem definition, distributional Q-value representation, local update rule, adaptive aggregation, and privacy mechanisms, is provided in Section 3.3. The pseudocode is presented in Algorithm 1.

3.2.3 Budget Allocation Modeling and Optimization

Figure 4 illustrates the budget allocation modeling and optimization framework.




Fig. 4. Budget Allocation Modeling and Optimization

As shown in the Figure 4, Phase 3 utilizes predictions from the previous stage to perform optimal advertising budget allocation. This is accomplished using the Fed-NSGA-AD algorithm, an advanced version of NSGA-II with

improved cost function and next-generation selection mechanism enabling multi-objective optimization based on multiple key metrics including ROI, CTR, costs, spillover effects, and advertising carryover effects.

Using the channel effectiveness predictions from Section 3.2.2, budget allocation is performed through the proposed Fed-NSGA-AD algorithm, an enhanced federated version of NSGA-II. The algorithm simultaneously optimizes multiple objectives — ROI, CTR, cost, spillover effects, and carryover effects — using an adaptive cost function and a dynamic next-generation selection mechanism. Each node evolves its local population, and the cloud server performs adaptive aggregation of local Pareto fronts.

The formal multi-objective optimization problem, adaptive cost function, federated evolution operators, and Pareto aggregation mechanism are presented in Section 3.4, with the complete pseudocode in Algorithm 2.

3.2.4 Security and Privacy Preservation

In Phase 4, the primary focus is on ensuring data security and preserving user privacy. Figure 5 illustrates the steps of Phase 4.




Fig. 5. Security and Privacy Preservation Steps

As shown in the figure, two key mechanisms are specifically employed at this stage, designed to prevent disclosure of personal user information while maintaining the efficiency and accuracy of federated learning models.

The framework incorporates a dual-layer privacy protection mechanism. First, *secure aggregation* ensures that local model weights are combined at the server without revealing individual contributions. Second, *DP-SGD* adds calibrated Gaussian noise to gradients, providing formal (ϵ, δ) -differential privacy guarantees. Both mechanisms are integrated directly with Fed-DAQRL and Fed-NSGA-AD. The formal privacy definitions and theorems are provided in Section 3.3.5 and Section 3.5.3.

3.2.5 Model Performance Evaluation

At this stage, the primary objective is evaluating model performance at three distinct levels, enabling verification of model accuracy, reliability, and operational readiness:

1. **Backtesting Level:** This process enables assessment of whether developed models effectively perform based on historical data. Using past data, model performance is tested under various conditions to determine if it can predict identified patterns and achieve satisfactory results.
2. **Off-Policy Evaluation Level:** At this stage, new decisions are simulated to test model performance in real-world conditions prior to actual execution. This method helps determine whether the model can make optimal decisions based on current conditions, and by comparing simulated results with real outcomes, provides additional insights into model capacity and reliability.
3. **Key Metric Assessment and Improvement:** This section focuses on improving key metrics such as ROI and conversion rates. These assessments help analyze real model impacts on advertising campaigns and identify areas for improvement. Using this data, the model can be continuously refined and optimized to ensure the best possible results.

These comprehensive evaluations provide confidence that Fed-DAQRL and Fed-NSGA-AD models are not only accurate and reliable but can also be effectively deployed in real advertising campaigns. This ensures that data-driven decisions are made optimally, delivering superior results for business operations.

3.2.6 Final Output and Operational Framework Delivery

In this final phase, the ultimate goal is achieving a model capable of automatically and optimally distributing advertising budgets across various platforms. This model is designed to reduce costs, increase ROI, and simultaneously maintain data security and privacy.

The final model leverages the Fed-DAQRL federated-distributed reinforcement learning core as its operational foundation. This algorithm enables learning and optimization in a federated environment, ensuring continuous learning and improvement capability. Additionally, utilizing Fed-NSGA-AD as an advanced multi-objective optimization algorithm enables dynamic budget allocation considering multiple objectives simultaneously. This means planning can accommodate multiple priorities including efficiency, cost, and business requirements.

The result of this process transforms the final model into an operational and scalable framework for the digital advertising industry. This framework is designed as both a scientific and practical tool with direct applicability in digital advertising, empowering brands and companies to make optimal, rapid, and data-driven decisions. Specifically, this system helps them achieve superior digital advertising results with less time and cost investment.

Furthermore, the proposed method is designed not only to address practical issues of digital advertising budget allocation but also to satisfy legal and security requirements related to user data. This enables more confident protection of sensitive user data while complying with all legal requirements. Finally, due to its substantial development and scalability capabilities, this framework can readily adapt to future market changes and requirements, making it an invaluable tool for the digital advertising industry.

3.3 Formal Mathematical Formulation of Fed-DAQRL

This section provides a rigorous mathematical formulation of the proposed Fed-DAQRL algorithm. We first formulate the digital advertising budget allocation problem as a Decentralized Markov Decision Process (Dec-MDP), which captures the partial observability and decentralized nature of multi-platform advertising environments.

We then introduce the distributional Q-value representation that enables explicit modeling of uncertainty in channel effectiveness prediction, followed by the local update rule and the adaptive cloud-based aggregation mechanism. The privacy-preserving components are formalized within the same framework, and the complete procedure is summarized as pseudocode.

3.3.1 Problem Formulation as a Decentralized Markov Decision Process

We formulate the digital advertising budget allocation problem as a Dec-MDP, defined by the tuple as Eq. 6:

$$\mathcal{M} = \langle \mathcal{K}, \mathcal{S}, \{\mathcal{A}_k\}_{k=1}^K, \mathcal{P}, \mathcal{R}, \gamma \rangle \quad (6)$$

Where $\mathcal{K} = \{1, 2, \dots, K\}$ denotes the set of K participating advertising platforms (federated nodes), $\mathcal{S}$ represents the global state space, where each state $s_t \in \mathcal{S}$ at time step t encapsulates the joint advertising context, including user engagement signals, historical budget consumption, and channel-level performance indicators. $\mathcal{A}_k$ denotes the local action space of platform k , where each action $a_{k,t} \in \mathcal{A}_k$ corresponds to a budget allocation decision. $\mathcal{P}: \mathcal{S} \times \mathcal{A}_1 \times \dots \times \mathcal{A}_K \rightarrow \Delta(\mathcal{S})$ is the transition probability kernel. $\mathcal{R}: \mathcal{S} \times \mathcal{A}_1 \times \dots \times \mathcal{A}_K \rightarrow \mathcal{R}$ is the global reward function, and $\gamma \in [0, 1)$ is the discount factor (set to $\gamma = 0.95$ in our experiments).

Due to privacy constraints, each platform k observes only a local partial observation $o_{k,t} \in \mathcal{O}_k$ rather than the full state s_t . This partial observability motivates the federated architecture, where each platform learns a local policy $\pi_k: \mathcal{O}_k \rightarrow \Delta(\mathcal{A}_k)$ while coordinating through cloud-based aggregation.

The global objective is to find a joint policy $\pi = (\pi_1, \dots, \pi_K)$ maximizing the expected discounted cumulative reward as Eq. 7:

$$J(\pi) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t \mathcal{R}(s_t, a_{1,t}, \dots, a_{K,t}) \right] \quad (7)$$

Subject to the budget constraint as Eq. 8:

$$\sum_{t=0}^T \sum_{k=1}^K c_k(a_k, t) \leq B \quad (8)$$

Where $c_k(\cdot)$ denote the cost function of platform k and B is the total advertising budget.

3.3.2 Distributional Q-Value Representation

Unlike classical Q-learning, which estimates the expected return $Q(s, a) = \mathbb{E}[Z(s, a)]$, the proposed Fed-DAQRL learns the full distribution of returns $Z(s, a)$. Following the categorical distributional RL framework (C51), the return distribution is parameterized by N atoms (in our experiments, $N = 51$) as Eq. 9:

$$Z(s, a) = \{z_i(s, a)\}_{i=1}^N, \quad \text{with supports } \{q_i\}_{i=1}^N \quad (9)$$

Where the support points are defined as Eq. 10:

$$q_i = V_{\min} + (i - 1) \frac{V_{\max} - V_{\min}}{N - 1}, \quad i = 1, \dots, N \quad (10)$$

With $V_{\min}$ and $V_{\max}$ denoting the minimum and maximum possible returns. Each atom $z_i(s, a)$ represents the probability that the return $Z(s, a)$ equals q_i , satisfying as Eq. 11:

$$\sum_{i=1}^N z_i(s, a) = 1, \quad z_i(s, a) \geq 0 \quad \forall_i \quad (11)$$

The expected Q-value is recovered as Eq. 12:

$$Q(s, a) = \sum_{i=1}^N q_i z_i(s, a) \quad (12)$$

Distributed Q-Value: For a federated node k , the *local distributed Q-value* is defined as the collection by Eq. 13:

$$Z_k(s, a) = \{z_{k,i}(s, a)\}_{i=1}^N \quad (13)$$

Which characterizes the return distribution learned from the local data of platform k .

Federated Distributed Q-Value: The global federated distributed Q-value is defined as the weighted aggregation by Eq. 14:

$$Z_{\text{global}}(s, a) = \sum_{k=1}^K \omega_k Z_k(s, a) \quad (14)$$

Where $\{\omega_k\}_{k=1}^K$ are adaptive aggregation weights satisfying $\sum_{k=1}^K \omega_k = 1$ and $\omega_k \geq 0$.

3.3.3 Local Q-Value Update Rule

At each federated communication round r , node k performs E local update steps on its local dataset D_k . The local distributed Q-value update follows the categorical projection operator as Eq. 15:

$$Z_k^{(r+1)}(s, a) \leftarrow \Pi_c(\mathcal{T}Z_k^{(r)}(s, a)) \quad (15)$$

Where $\mathcal{T}$ denote the distributional Bellman operator that is computed as Eq. 16:

$$(\mathcal{T}Z)(s, a) = \mathcal{R}(s, a) + \gamma Z(s', a') \quad (16)$$

With $s' \sim \mathcal{P}(\cdot | s, a)$ and $a' \sim \pi(\cdot | s')$. The operator Π_c denotes the categorical projection onto the fixed support $\{q_i\}_{i=1}^N$.

The local loss function for node k at round r is the cross-entropy between the projected target distribution and the predicted distribution as Eq. 17:

$$\mathcal{L}_k^{(r)}(\theta_k) = - \sum_{i=1}^N \tilde{z}_i(s, a) \log_{z_i}(s, a; \theta_k) \quad (17)$$

Where $\tilde{z}_i$ denotes the projected target distribution and θ_k are the local network parameters of node k . Parameter updates follow stochastic gradient descent with the DP-SGD mechanism.

3.3.4 Adaptive Cloud-Based Aggregation

After E local updates, each node k transmits its local parameters $\theta_k^{(r)}$ (or the corresponding distributed Q-value representation) to the cloud server. The server performs adaptive aggregation using data-quality-aware weights as Eq. 18:

$$w_k^{(r)} = \frac{\sum_{j=1}^K \exp(-\beta \cdot \mathcal{L}_k^{(r)})}{\sum_{j=1}^K \exp(-\beta \cdot \mathcal{L}_j^{(r)})} \quad (18)$$

Where $\mathcal{L}_k^{(r)}$ is the local loss of node k at round r and $\beta > 0$ is a temperature parameter controlling the sharpness of the weighting. Nodes with lower local loss receive higher aggregation weight, reflecting higher data quality or better local convergence.

The global Q-value parameters are then updated as Eq. 19:

$$\theta_{global}^{(r+1)} = \sum_{k=1}^K w_k^{(r)} \theta_k^{(r+1)} \quad (19)$$

Adaptive Weighting Reduces Heterogeneity Bias: Under bounded local loss $\mathcal{L}_k^{(r)} \in [\mathcal{L}_{min}, \mathcal{L}_{max}]$, the adaptive weighting scheme above reduces the aggregation bias induced by data heterogeneity compared to uniform weighting (FedAvg).

Uniform FedAvg assigns $w_k = \frac{1}{K}$ regardless of local data quality. Under heterogeneity, nodes with noisy or unrepresentative data contribute equally, biasing the global update. The proposed exponential weighting assigns lower weight to high-loss (noisy) nodes, monotonically reducing the contribution of biased updates.

3.3.5 Privacy-Preserving Mechanisms

Secure Aggregation: Before transmission, each node k masks its parameters using pairwise secret sharing with neighboring nodes. The masked parameter transmitted to the server is by Eq. 20:

$$\tilde{\theta}_k^{(r)} = \theta_k^{(r)} + \sum_{j \in \mathcal{N}_k} \delta_{kj} \pmod{p} \quad (20)$$

Where δ_{kj} is a pairwise random mask shared between nodes k and j , and p is a large prime. The server recovers the true aggregate as Eq. 21:

$$\sum_{k=1}^K \tilde{\theta}_k^{(r)} = \sum_{k=1}^K \theta_k^{(r)} \pmod{p} \quad (21)$$

$$\sum k = 1K\theta \sim k(r) = \sum k = 1K\theta k(r) \pmod{p}$$

Because pairwise masks cancel out.

Differential Privacy (DP-SGD): Each node clips its gradient to norm C and adds Gaussian noise by Eq. 22:

$$\tilde{g}_k = \frac{g_k}{\max(1, \|g_k\|_2 / C)} + \mathcal{N}(0, \sigma^2 C^2 \mathbb{I}) \quad (22)$$

This provides (ϵ, δ) -differential privacy with Eq. 23:

$$\epsilon = \frac{c \sqrt{2 \ln \left(\frac{1.25}{\delta} \right)}}{\sigma} \quad (23)$$

Following the standard DP-SGD analysis. In Algorithm 1, Fed-DAQRL pseudocode is shown.

Algorithm 1. Fed-DAQRL pseudocode

ALGORITHM 1. Fed-DAQRL ($K, N, E, R, \eta, \gamma, C, \sigma$)

INPUT:

- 1: INITIALIZE K nodes
- 2: INITIALIZE N atoms
- 3: INITIALIZE E local epochs
- 4: INITIALIZE R rounds
- 5: INITIALIZE Learning rate η
- 6: INITIALIZE Discount γ
- 7: INITIALIZE DP parameters (C, σ)

BODY:

- 8: INITIALIZE global parameters θ_{global} randomly
- 9: **for** $r = 1$ to R **do**
- 10: SERVER broadcasts θ_{global} to all K nodes
- 11: **for** each node $k \in K$ in parallel **do**
- 12: SET $\theta_k \leftarrow \theta_{global}$
- 13: **for** $e = 1$ to E **do**
- 14: SAMPLE mini-batch from local D_k
- 15: COMPUTE distributional Bellman target: $\mathcal{T}Z_k$
- 16: PROJECT onto N -atom support: $\tilde{Z}_k$
- 17: COMPUTE loss $\mathcal{L}_k$
- 18: COMPUTE gradient $g_k \leftarrow \nabla \mathcal{L}_k$
- 19: APPLY DP-SGD: $\tilde{g}_k$
- 20: UPDATE θ_k
- 21: **end for**
- 22: APPLY secure aggregation mask: $\tilde{\theta}_k$
- 23: TRANSMIT $\tilde{\theta}_k$ to server
- 24: **end for**
- 25: SERVER computes adaptive weights: w_k
- 26: Server aggregates: θ_{global}
- 27: **end for**

OUTPUT:

- 28: RETURN θ_{global}

As shown in Algorithm 1, the algorithm receives its inputs in Lines 1–7, initializing the number of participating nodes K , the number of distributional atoms N , the local epochs E , the federated rounds R , the learning rate η , the discount factor γ , and the differential privacy parameters (C, σ).

In Line 8, the global parameters θ_{global} are initialized randomly at the cloud server, establishing the common starting point for federated training. Line 9 opens the main federated loop over R communication rounds. In Line 10, the server broadcasts the current global parameters to all K nodes, ensuring synchronized initialization of local training. Line 11 launches parallel training across all participating nodes.

Within each node, Line 12 sets the local parameters θ_k equal to the received global parameters. Line 13 begins the E local training epochs. Line 14 samples a mini-batch from the node's private local dataset D_k , ensuring that raw user data never leaves the platform. Line 15 computes the distributional Bellman target $\mathcal{T}Z_k = \mathcal{R} + \gamma Z_k(s', a')$, propagating reward information backward through the distributional value function.

Line 16 projects this target onto the fixed N -atom support via the categorical projection operator Π_C , producing $\tilde{Z}_k$. Line 17 computes the local loss $\mathcal{L}_k$ as the cross-entropy between the projected target distribution and the predicted distribution. Line 18 computes the gradient $g_k \leftarrow \nabla \mathcal{L}_k$ via backpropagation.

Line 19 applies DP-SGD by clipping the gradient to norm C and adding Gaussian noise, providing formal differential privacy guarantees. Line 20 updates the local parameters via $\theta_k \leftarrow \theta_k - \eta \tilde{g}_k$. Line 21 closes the local epoch loop.

After local training, Line 22 applies the secure aggregation mask $\tilde{\theta}_k$, ensuring that the server cannot observe individual contributions. Line 23 transmits the masked parameters to the cloud server. Line 24 closes the parallel node loop.

In Line 25, the server computes adaptive aggregation weights w_k , assigning higher importance to nodes with lower local loss. Line 26 aggregates the local parameters as θ_{global} . Line 27 closes the federated round loop. Finally, Line 28 returns the converged global parameters θ_{global} as the output of the algorithm.

3.4 Formal Mathematical Formulation of Fed-NSGA-AD

This section presents the mathematical foundations of the proposed Fed-NSGA-AD algorithm. We begin by defining the multi-objective optimization problem together with its objective functions and feasible region. We then formalize the adaptive cost function with dynamic weights and describe the federated population initialization and local evolution operators.

Finally, we present the adaptive aggregation of local Pareto fronts and the complete pseudocode, and explain how Fed-NSGA-AD is integrated with Fed-DAQRL to form a unified closed-loop optimization framework.

3.4.1 Multi-Objective Optimization Problem Definition

The budget allocation problem is formally defined as a constrained multi-objective optimization problem by Eq. 24:

$$\min_{x \in X} F(x) = [f_1(x), f_2(x), \dots, f_M(x)]^T \quad (24)$$

Where $x = (x_1, \dots, x_K)$ is the budget allocation vector with x_k denoting the budget assigned to platform k , and:

$$\begin{aligned} f_1(x) &= -\text{ROI}(x), f_2(x) = -\text{CTR}(x) \\ f_3(x) &= \text{Cost}(x), f_4(x) = -\text{Spillover}(x), f_5(x) = -\text{Carryover}(x) \end{aligned} \quad (25)$$

The feasible region is computed by Eq. 26:

$$X = \left\{ x \in \mathbb{R}^K : \sum_{k=1}^K x_k \leq B, x_k \geq 0 \right\} \quad (26)$$

Pareto Dominance: A solution $x^{(1)}$ dominates $x^{(2)}$, denoted $x^{(1)} < x^{(2)}$, if $\forall_i \in \{1, \dots, M\}: f_i(x^{(1)}) \leq f_i(x^{(2)})$ and $\exists_j: f_j(x^{(1)}) < f_j(x^{(2)})$.

Pareto Optimal Set: The Pareto optimal set is as Eq. 27:

$$P^* = \{x \in X : \nexists x' \in X, x' < x\} \quad (27)$$

3.4.2 Adaptive Cost Function with Dynamic Weights

The adaptive weights α and β in the cost function evolve over generations based on the relative importance of spillover and carryover effects as Eq. 28:

$$\alpha^{(g+1)} = \alpha^{(g)} + \eta_\alpha \cdot \frac{\partial \text{Spillover}(x^*)}{\partial g}, \quad \beta^{(g+1)} = \beta^{(g)} + \eta_\beta \cdot \frac{\partial \text{Carryover}(x^*)}{\partial g} \quad (28)$$

Where $\eta_\alpha, \eta_\beta > 0$ are adaptation rates and x^* is the current best compromise solution.

3.4.3 Federated Population Initialization and Evolution

At each federated node k , a local population $P_k^{(g)}$ of size N_p is initialized. Local evolution operators include:

- Crossover: With probability p_c , two parents $x^{(1)}, x^{(2)}$ produce offspring via simulated binary crossover (SBX) as Eq. 29:

$$x_i^{\text{child}} = 0.5[(1 + \beta_i)x_i^{(1)} + (1 - \beta_i)x_i^{(2)}] \quad (29)$$

- Mutation: With probability p_m , polynomial mutation perturbs each gene as Eq. 30:

$$x_i^{\text{mut}} = x_i + \delta_i(x_i^U - x_i^L) \quad (30)$$

- Non-dominated Sorting: Local populations are sorted into Pareto fronts using fast non-dominated sorting as Eq. 31:

$$F_1, F_2, \dots, F_L = \text{FastNonDominatedSort}(\mathcal{P}_k^{(g)}) \quad (31)$$

- Crowding Distance: Within each front, crowding distance preserves diversity is computed by Eq. 32:

$$d_i = \sum_{m=1}^M \frac{f_m(x_i + 1) - f_m(x_i - 1)}{f_m^{\max} - f_m^{\min}} \quad (32)$$

3.4.4 Adaptive Aggregation of Pareto Fronts

After local evolution, each node transmits its Pareto front $\mathcal{F}_k^{(g)}$ to the cloud server. The server performs adaptive aggregation as Eq. 33:

$$\mathcal{F}_{\text{global}}^{(g)} = \text{AdaptiveMerge}(\mathcal{F}_1^{(g)}, \dots, \mathcal{F}_K^{(g)}) \quad (33)$$

Where *AdaptiveMerge* combines all local fronts and re-applies non-dominated sorting and crowding distance, but with adaptive selection pressure that prioritizes solutions contributing to underrepresented regions of the global Pareto front.

In Algorithm 2, Fed-NSGA-AD pseudocode is shown.

Algorithm 2. Fed-NSGA-AD pseudocode	
ALGORITHM 2. Fed-NSGA-AD ($K, N_p, G, p_c, p_m, \alpha, \beta$)	
INPUT:	
1:	INITIALIZE K nodes
2:	INITIALIZE population size N_p
3:	INITIALIZE generations G
4:	crossover prob p_c
5:	INITIALIZE mutation prob p_m
6:	INITIALIZE adaptive weights (α, β)
BODY:	
7:	INITIALIZE global population P_{global} randomly
8:	for round $r = 1$ to R_{fed} do
9:	Server broadcasts P_{global} to all K nodes
10:	for each node $k \in K$ in parallel do
11:	INITIALIZE local population $P_k \leftarrow P_{\text{global}}$
12:	for $g = 1$ to G do
13:	EVALUATE $F(x)$ for all $x \in P_k$
14:	APPLY fast non-dominated sorting
15:	COMPUTE crowding distances
16:	SELECT parents via binary tournament
17:	APPLY SBX crossover with prob p_c
18:	APPLY polynomial mutation with prob p_m
19:	COMBINE parent + offspring populations

```

20:      SELECT top  $N_p$  via non-dominated + crowding
21:    end for
22:      TRANSMIT local Pareto front  $F_k$  to server
23:    end for
24:    Server aggregates:  $F_{global} = AdaptiveMerge(F_1, \dots, F_K)$ 
25:    UPDATE adaptive weights  $(\alpha, \beta)$ 
26:  end for
OUTPUT:
27: Return  $F_{global}$ 

```

As shown in Algorithm 2, the algorithm receives its inputs in Lines 1–6, initializing the number of participating nodes KK , the population size N_p , the number of generations GG , the crossover probability p_c , the mutation probability p_m , and the adaptive weights (α, β) .

In Line 7, the global population P_{global} is initialized randomly at the cloud server, establishing the initial set of candidate budget allocation solutions. Line 8 opens the outer federated loop over R_{fed} rounds. In Line 9, the server broadcasts the current global population to all K nodes, ensuring a common starting point for local evolution. Line 10 launches parallel local evolution across all participating nodes.

Within each node, Line 11 initializes the local population P_k from the received global population. Line 12 begins the G evolutionary generations. Line 13 evaluates the objective functions $F(x)$ for all individuals in the local population, providing multi-objective fitness values for each candidate budget allocation. Line 14 applies fast non-dominated sorting to partition the population into Pareto fronts. Line 15 computes crowding distances within each front to preserve diversity and prevent premature convergence.

Line 16 selects parents via binary tournament selection, favoring solutions with better rank and larger crowding distance. Line 17 applies SBX crossover with probability p_c , generating offspring that inherit characteristics from both parents. Line 18 applies polynomial mutation with probability p_m , introducing perturbations that maintain population diversity. Line 19 combines the parent and offspring populations into a single pool. Line 20 selects the top N_p individuals via non-dominated rank and crowding distance, restoring the population to its original size while improving overall quality. Line 21 closes the generation loop.

After local evolution, Line 22 transmits the local Pareto front $\mathcal{F}_k$ to the cloud server, sharing only Pareto-optimal solutions rather than raw local data. Line 23 closes the parallel node loop.

In Line 24, the server aggregates all local Pareto fronts via $F_{global} = AdaptiveMerge(F_1, \dots, F_K)$, combining them and prioritizing solutions that contribute to underrepresented regions of the global front. Line 25 updates the adaptive weights (α, β) based on the observed relative importance of spillover and carryover effects. Line 26 closes the federated round loop.

Finally, Line 27 returns the final global Pareto front F_{global} as the output of the algorithm.

3.5 Theoretical Analysis

Theorem 1 (Convergence of Federated Distributional Q-Learning): Under the assumptions:

- (A1) The reward function is bounded: $|\mathcal{R}(s, a)| \leq R_{max}$
- (A2) The learning rates satisfy $\sum_r \eta_r = \infty$ and $\sum_r \eta_r^2 < \infty$
- (A3) The categorical projection operator Π_C is a non-expansion in the Wasserstein-1 metric.

The federated distributional Q-value iterates $\{Z_{global}^{(r)}\}$ converge almost surely to a fixed point Z^* satisfying as Eq. 34:

$$Z^* = \Pi_C \mathcal{T} Z^* \quad (34)$$

Theorem 2 (Privacy of Fed-DAQRL): The proposed Fed-DAQRL with DP-SGD satisfies (ϵ, δ) -differential privacy per communication round, where $\epsilon = \frac{c \sqrt{2 \ln(\frac{1.25}{\delta})}}{\sigma}$, and over R rounds, the total privacy loss under composition is $\epsilon_{total} = \epsilon \sqrt{2 R \ln(\frac{1}{\delta'})}$.

4 EXPERIMENTAL RESULTS

This section presents the experimental evaluation of the proposed framework for digital advertising budget allocation based on multi-objective distributed federated learning and cloud computing. The experiments are designed to comprehensively assess the performance of the proposed Fed-DAQRL algorithm and the Fed-NSGA-AD optimization method across multiple dimensions.

The section is organized as follows: Details the system specifications and simulation environment, the datasets employed in the experiments, including the Click Data on Taobao.com (ADAC) dataset and the Click-Through Rate Data (CTRD) dataset, which are consistent with the baseline study [18]. The evaluation metrics used for performance assessment, including Loss, Accuracy, and ROC metrics, and the experimental results and comparative analysis against state-of-the-art methods, demonstrating the effectiveness of the proposed framework.

4.1 System Specifications

The experimental evaluation was conducted using the simulation setup described in Table II. The federated learning simulation was configured with 10 participating nodes, representing different advertising platforms. The experimental setup follows a consistent configuration with the baseline study [18] to ensure fair comparison.

The MATLAB environment provides essential toolboxes including the Optimization Toolbox for implementing the multi-objective evolutionary algorithm, the Statistics and Machine Learning Toolbox for data preprocessing and evaluation, and the Deep Learning Toolbox for neural network components. CloudSim provides capabilities for modeling data centers, virtual machines, and network characteristics, enabling realistic simulation of the cloud-based aggregation and coordination mechanisms.

The NetBeans IDE facilitates efficient Java development and debugging of the CloudSim simulation components, ensuring seamless integration between the MATLAB-based algorithm implementation and the CloudSim-based cloud infrastructure simulation.

TABLE II
SYSTEM SPECIFICATIONS

Component	Specification
CPU	Intel(R) Core(TM) i7-10750H CPU @ 2.60GHz
RAM	16 GB
GPU	NVIDIA GeForce RTX 2060
Operating System	Windows 10 Professional (64-bit)
Primary Simulation Platform	MATLAB 2024a
Cloud Simulation Framework	CloudSim (Java-based)
Cloud Simulation IDE	NetBeans IDE
Cloud Simulation Language	Java
MATLAB Toolboxes	Optimization Toolbox, Statistics and Machine Learning Toolbox, Deep Learning Toolbox, Parallel Computing Toolbox
CloudSim Features	Data center modeling, virtual machine management, network simulation, resource allocation
Federated Learning Nodes	10 participating nodes

The combination of MATLAB for algorithmic implementation and CloudSim for infrastructure simulation provides a comprehensive experimental environment that accurately reflects the proposed federated-cloud architecture.

MATLAB handles the computationally intensive operations including the training of local models, the Fed-NSGA-AD optimization, and the evaluation of performance metrics, while CloudSim simulates the cloud-based coordination, model aggregation, and resource management aspects of the framework.

The integration between these platforms enables realistic evaluation of the proposed framework's performance across both the algorithmic and infrastructural dimensions, providing insights into its practical applicability in real-world digital advertising scenarios.

4.2 Datasets

This research employs two publicly available advertising click datasets to evaluate the proposed framework. These datasets provide comprehensive real-world advertising data that enables effective training and evaluation of the proposed federated learning and optimization methods. The following subsections describe each dataset in detail.

4.2.1 Ali_Display_Ad_Click (ADAC) Dataset

The ADAC¹ dataset is a publicly available advertising click dataset released by Alibaba Group, containing user behavior data from the Taobao e-commerce platform. This dataset is specifically designed for display advertising click prediction research and provides comprehensive user-ad interaction data. The dataset comprises 20 million click records collected over a period of 8 days, making it suitable for large-scale advertising prediction and optimization tasks. It includes 9 features, encompassing user information, ad information, and contextual features. The dataset was released in 2018 and has been widely used in advertising click prediction research.

4.2.2 Click-Through Rate Data (CTRD Dataset)

The CTRD² dataset is a publicly available benchmark dataset for CTR prediction in online advertising. This dataset contains comprehensive advertising click records collected from various online advertising platforms, enabling robust evaluation of CTR prediction models. The dataset comprises 100,000 samples with 6 features including user demographic information, ad characteristics, and interaction features. The dataset is suitable for evaluating federated learning models for CTR prediction and budget allocation optimization. The dataset was released in 2020 and has been utilized in multiple advertising research studies. In Table III, summary of datasets used in the experiments is shown.

TABLE III
SUMMARY OF DATASETS USED IN THE EXPERIMENTS

Dataset	Description	Samples	Features	Year
ADAC Dataset	Taobao display advertising click data	20,000,000	9	2018
CTRD Dataset	Click-through rate prediction data	100,000	6	2020

Each dataset is partitioned into training, validation, and testing sets using a 70:15:15 ratio. Specifically, 70% of the samples are allocated for model training, 15% for validation during the optimization process, and the remaining 15% for final model evaluation.

This partitioning strategy ensures that the model is trained on a sufficiently large subset while maintaining independent validation and test sets for unbiased performance assessment. To ensure the robustness of the experimental results, all experiments are repeated five times with different random seeds, and the average performance metrics are reported along with standard deviations.

The data splitting is performed in a stratified manner to preserve the class distribution across all subsets, which is particularly important given the inherent class imbalance in advertising click prediction tasks.

4.3 Evaluation Metrics

To comprehensively evaluate the performance of the proposed framework, we employ three standard evaluation metrics. These metrics provide complementary perspectives on model performance, enabling thorough assessment of prediction accuracy, classification capability, and overall model quality.

4.3.1 Classification metrics

The cross-entropy loss function is employed to measure the discrepancy between the predicted probabilities and the true labels. This loss function is particularly suitable for binary classification tasks in click prediction, where the model outputs a probability of a positive click event. The cross-entropy loss is defined as Eq. 35:

¹ <https://tianchi.aliyun.com/dataset/56?lang=en-us>

² <https://www.kaggle.com/datasets/sahideseker/click-through-rate-ctr-prediction-dataset>

$$L = -\frac{1}{N} \sum_{i=1}^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] \quad (35)$$

Where N represents the total number of samples in the dataset, y_i denotes the true binary label for the i -th sample (where $y_i = 1$ indicates a click and $y_i = 0$ indicates no click), and $\hat{y}_i$ represents the predicted probability of a click for the i -th sample.

The loss function penalizes incorrect predictions with higher penalties for confident wrong predictions, making it an effective measure for model optimization.

Accuracy is a fundamental evaluation metric that measures the proportion of correctly classified samples among all predictions. It provides an intuitive measure of overall model performance and is widely used in classification tasks. Accuracy is defined as Eq. 36:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (36)$$

Where TP denotes the number of true positives (samples correctly predicted as clicks), TN denotes the number of true negatives (samples correctly predicted as non-clicks), FP denotes the number of false positives (samples incorrectly predicted as clicks), and FN denotes the number of false negatives (samples incorrectly predicted as non-clicks).

The Receiver Operating Characteristic (ROC) curve and the Area Under the ROC Curve (AUC) are comprehensive evaluation metrics that assess the model's ability to discriminate between positive and negative classes across all classification thresholds. The ROC curve plots the True Positive Rate (TPR) against the False Positive Rate (FPR) at various threshold settings. The TPR, also known as sensitivity or recall, is defined as Eq. 36:

$$TPR = \frac{TP}{TP + FN} \quad (36)$$

The FPR, also known as the false alarm rate, is defined as Eq. 37:

$$FPR = \frac{FP}{FP + TN} \quad (37)$$

The AUC represents the area under the ROC curve and provides a single scalar value summarizing the model's discriminative ability across all possible thresholds.

Precision measures the proportion of correctly predicted positive samples among all positive predictions, that defined as Eq. 38:

$$Precision = \frac{TP}{TP + FP} \quad (38)$$

The Recall metric, also known as sensitivity, measures the proportion of correctly predicted positive samples among all actual positive samples, defined as Eq. 39:

$$Recall = \frac{TP}{TP + FN} \quad (39)$$

The F1-Score provides a harmonic mean of precision and recall, defined as Eq. 40:

$$F1 - score = 2 \times \frac{Precision \times Recall}{Precision + Recall} \quad (40)$$

Offering a balanced measure that considers both false positives and false negatives. The FPR provides information about the proportion of negative samples incorrectly classified as positive.

4.3.2. Economic Metrics

Four economic metrics are used to assess the financial impact of the proposed budget allocation framework:

Return on Advertising Spend (ROAS) measures the revenue generated per unit of advertising expenditure as Eq. 41:

$$ROAS = \frac{\sum_{k=1}^K Revenue_k}{\sum_{k=1}^K Spend_k} \quad (41)$$

Cost per Acquisition (CPA) measures the average cost incurred per conversion as Eq. 42:

$$CPA = \frac{\sum_{k=1}^K Spend_k}{\sum_{k=1}^K Conversions_k} \quad (42)$$

Return on Investment (ROI) measures the net profitability of the advertising campaign as Eq. 43:

$$ROI = \frac{\sum_{k=1}^K (\text{Revenue}_k - \text{Spend}_k)}{\sum_{k=1}^K \text{Spend}_k} \quad (43)$$

Budget Utilization measures the efficiency of budget consumption as Eq. 44:

$$\text{Budget Utilization} = \frac{\sum_{k=1}^K \text{Spend}_k}{B} \quad (44)$$

Where B denote the total allocated budget.

4.3.3 System Efficiency Metrics

Four system-level metrics assess the computational and communication efficiency:

Communication Cost measures the total bytes transmitted between nodes and the server per federated round as Eq. 45:

$$\text{CommCost} = \sum_{r=1}^R \sum_{k=1}^K |\theta_k^{(r)}| \times b \quad (45)$$

Where $|\theta_k^{(r)}|$ is the number of parameters and b is the bytes per parameter. Training Time measures the wall-clock time required for full convergence.

Convergence Rounds measures the number of federated rounds required to reach a stable performance plateau.

Inference Latency measures the average time required to produce a budget allocation decision.

These metrics provide a holistic assessment of the proposed framework, capturing not only predictive accuracy but also practical economic value and operational efficiency.

4.4 Hyperparameters

The proposed framework involves several Hyperparameters that are carefully tuned to achieve optimal performance. Table IV summarizes the key Hyperparameters used in the experiments, including the federated learning configuration, the Fed-DAQRL algorithm parameters, and the Fed-NSGA-AD optimization parameters.

TABLE IV
HYPERPARAMETERS USED IN THE EXPERIMENTS

Component	Hyperparameter	Value	Description
Federated Learning	Number of Nodes (K)	10	Number of participating advertising platforms
	Local Epochs (E)	5	Number of local training epochs per communication round
	Communication Rounds (R)	100	Total number of federated communication rounds
	Batch Size	128	Local batch size for training
	Learning Rate	0.001	Initial learning rate for local model training
	Learning Rate Decay	0.99	Decay factor applied after each communication round
Fed-DAQRL	Discount Factor (γ)	0.95	Discount factor for future rewards in Q-learning
	Q-Value Distribution Support	51	Number of atoms for distributional Q-value representation
	Target Network Update Frequency	10	Number of rounds between target network updates
	Exploration Rate (ϵ)	0.1	Initial exploration rate for ϵ -greedy policy
	Exploration Decay	0.995	Decay factor for exploration rate
Fed-NSGA-AD	Minimum Exploration Rate	0.01	Minimum exploration rate
	Population Size	100	Number of solutions in each population
	Generations	200	Number of evolutionary generations per round
	Crossover Probability	0.9	Probability of applying crossover
	Mutation Probability	0.1	Probability of applying mutation
Privacy	Adaptive Weight α	0.6	Weight for spillover effect in cost function
	Adaptive Weight β	0.4	Weight for carryover effect in cost function
	Privacy Budget (ϵ)	8.0	Differential privacy budget
	DP Noise Scale (σ)	0.5	Noise scale for DP-SGD
	DP Clipping Norm	1.0	Gradient clipping norm for DP-SGD

Cloud Simulation	Number of VMs	20	Virtual machines in cloud data center
	VM RAM	4 GB	Memory per VM
	VM MIPS	1000	Processing power per VM
	Bandwidth	1000	Network bandwidth for cloud communication
		Mbps	

As shown in Table IV, the federated learning parameters, including the number of nodes ($K = 10$), local epochs ($E = 5$), and communication rounds ($R = 100$), were chosen to balance convergence speed and communication efficiency. The Fed-DAQRL parameters, including the discount factor ($\gamma = 0.95$) and exploration rate ($\epsilon = 0.1$), follow standard reinforcement learning practices for advertising optimization tasks. The Fed-NSGA-AD parameters, including population size and generations, were set to ensure adequate exploration of the solution space while maintaining computational efficiency.

The privacy parameters ($\epsilon = 8.0$ and $\sigma = 0.5$) provide a reasonable balance between privacy preservation and model utility, as recommended in recent differential privacy literature. The cloud simulation parameters were configured to realistically represent a cloud-edge computing environment suitable for large-scale advertising platforms.

4.5 Experimental Results and Comparative Analysis

This section presents comprehensive experimental results evaluating the performance of the proposed framework across different cloud infrastructure configurations and training iterations. The experiments are designed to achieve two primary objectives: first, to determine the optimal number of virtual machines (VMs) that maximizes computational efficiency while maintaining prediction accuracy; and second, to identify the optimal number of training iterations required for model convergence.

The experimental evaluation is conducted using both the ADAC and CTRD datasets, with performance measured across all evaluation metrics including Loss, Accuracy, ROC-AUC, Precision, Recall, and F1-Score.

All experiments are repeated five times with different random seeds, and average results with standard deviations are reported.

4.5.1 VM Configuration Analysis

To determine the optimal cloud infrastructure configuration for the proposed framework, we evaluate system performance across different numbers of virtual machines. The VM count varies from 5 to 200 in increments of 5, representing different levels of cloud resource allocation from minimal to extensive computational capacity.

This comprehensive range enables identification of both the minimum resource threshold for acceptable performance and the saturation point beyond which additional resources provide diminishing returns. The experiments are conducted with a fixed number of training iterations ($R=100$) to isolate the effect of VM configuration. Figures 6 through 11 present the performance of the proposed framework across different VM counts for both the ADAC and CTRD datasets.




Fig. 6. Accuracy analysis of the proposed framework across different VM configurations

Figure 6 illustrates the model accuracy achieved by the proposed framework across different VM configurations ranging from 5 to 200 VMs for both the ADAC and CTRD datasets. The results demonstrate exceptional performance, with the proposed framework achieving accuracy rates exceeding 99% with sufficient computational resources. For the ADAC dataset, accuracy improves from 91.2% with 5 VMs to 99.3% with 200 VMs, representing an impressive improvement of 8.1 percentage points. Similarly, the CTRD dataset demonstrates improvement from 90.5% to 99.1%. The figure reveals a distinct three-phase pattern: an initial rapid improvement phase (5-30 VMs), a gradual improvement phase (30-150 VMs), and a saturation phase (beyond 150 VMs). The most significant improvements occur between 5 and 30 VMs, where the ADAC dataset achieves 96.7% accuracy, representing 67.9% of the total improvement. The CTRD dataset shows a similar pattern, achieving 96.5% accuracy at 30 VMs.

The saturation behavior at higher VM counts indicates that the model achieves near-optimal performance with approximately 150 VMs, beyond which additional computational resources provide negligible benefit. As observed in the figure, the accuracy values remain nearly constant after 150 VMs, with the ADAC dataset achieving 99.3% and the CTRD dataset achieving 99.1% at this configuration. This finding is particularly significant for practical deployments, as it demonstrates that while the proposed framework continues to benefit from increased computational resources up to 150 VMs, further expansion beyond this point yields minimal improvement. Therefore, 150 VMs is identified as the optimal configuration that provides the best balance between computational cost and model performance.




Fig. 7. Loss analysis of the proposed framework across different VM configurations

Figure 7 presents the cross-entropy loss achieved by the proposed framework across different VM configurations. The results demonstrate rapid loss reduction, with the proposed framework achieving remarkably low loss values with sufficient computational resources. For the ADAC dataset, the loss decreases from 0.185 with 5 VMs to 0.002 with 200 VMs, representing a reduction of 98.9%. The CTRD dataset shows a similar pattern, with loss decreasing from 0.198 to 0.003.

The most significant loss reduction occurs between 5 and 30 VMs, where the ADAC dataset achieves a loss of 0.062, representing 66.5% of the total reduction. The CTRD dataset achieves a loss of 0.068 at 30 VMs. The loss values continue to decrease gradually between 30 and 150 VMs, reaching 0.003 for ADAC and 0.004 for CTRD at 150 VMs. Beyond 150 VMs, the loss values become nearly constant, with only marginal reductions observed. The extremely low loss values (<0.005) achieved at 150+ VMs indicate that the model has reached near-perfect convergence, effectively capturing the underlying patterns in the advertising data. Therefore, 150 VMs is identified as the optimal configuration where the loss reaches its near-minimum value. Figure 8 illustrates a comprehensive ROC analysis across different VM configurations through two complementary visualizations that collectively demonstrate the impact of cloud resource allocation on model discriminative performance.




Fig. 8. ROC-AUC analysis of the proposed framework across different VM configurations

Figure 8a displays the ROC curves for six representative VM configurations (5, 30, 60, 100, 150, and 200 VMs), providing a detailed view of the model's discriminative ability at different classification thresholds. The ROC curve plots the TPR against the FPR (1 - Specificity) across all possible threshold values, with the AUC serving as a comprehensive measure of classification performance. The diagonal dashed line represents the performance of a random classifier (AUC = 0.5), against which all configurations show substantial improvement.

As observed in the Figure 8, the ROC curves demonstrate clear separation between configurations, with the 150 VM configuration achieving near-perfect ROC performance (AUC ≈ 0.9945). The curves exhibit a progressive rightward and upward shift as VM count increases, indicating improved discriminative ability. The 5 VM configuration (AUC ≈ 0.850) shows moderate performance with a characteristic curved shape, while the 30 VM configuration (AUC ≈ 0.920) demonstrates substantial improvement. The 60 VM configuration (AUC ≈ 0.951) and 100 VM configuration (AUC ≈ 0.983) show increasingly better separation between classes, with the curves approaching the upper-left corner of the ROC space. The 150 VM configuration (AUC ≈ 0.9945) achieves near-perfect discrimination, with the ROC curve closely hugging the upper-left corner, indicating that the model can almost perfectly distinguish between click and non-click events. The markers on each curve indicate the optimal operating point where Youden's index (Sensitivity + Specificity - 1) is maximized. The progression of optimal operating points from the lower-right toward the upper-left corner demonstrates the consistent improvement in classification performance with additional computational resources.

Figure 8b presents the ROC-AUC values across all VM configurations from 5 to 200, divided into four distinct performance zones: Low Performance (AUC < 0.90), Medium Performance ($0.90 \leq \text{AUC} < 0.95$), High Performance ($0.95 \leq \text{AUC} < 0.99$), and Excellent Performance (AUC ≥ 0.99). The figure reveals a distinct three-phase pattern characteristic of resource scaling in distributed learning systems.

In the first phase (5 to 60 VMs), AUC increases rapidly from 0.850 at 5 VMs to 0.951 at 60 VMs. The steep slope indicates that each additional VM provides substantial marginal benefit, attributed to improved parallel processing and more efficient distributed gradient aggregation. At 30 VMs, the AUC reaches 0.920, moving the model from low to medium performance, while at 60 VMs, the AUC exceeds 0.95, entering the high performance zone. In the second phase (60 to 160 VMs), AUC continues to improve but at a reduced rate, increasing from 0.951 at 60 VMs to 0.9954 at 160 VMs. This sigmoidal shape with diminishing returns represents the optimal trade-off region where additional resources still provide noticeable improvements. At 100 VMs, AUC reaches 0.9826, and at 150 VMs, AUC achieves 0.9945. In the third phase (beyond 160 VMs), AUC values approach the theoretical maximum of 1.000, with only marginal improvements. At 160 VMs, AUC reaches 0.9954, and at 200 VMs, the AUC improves to only 0.9983, a negligible increase of 0.0029, indicating saturation.

The horizontal dashed lines at AUC = 0.90, 0.95, and 0.99 demarcate the performance zone boundaries, while the vertical dashed line at 150 VMs indicates the optimal configuration identified in this study. The choice of 150 VMs as optimal is justified by the observation that beyond this point, the model achieves excellent performance (AUC > 0.99) with only marginal additional improvements, making further resource allocation economically inefficient for

practical applications. The ADAC dataset consistently outperforms the CTRD dataset by a small margin (approximately 0.2-0.5%), which can be attributed to its larger sample size and richer feature representation, providing more comprehensive training data for the model.

The combined analysis confirms that 150 VMs provides the optimal balance between computational cost and discriminative performance, achieving near-perfect ROC-AUC scores (≥ 0.9945) with excellent stability and reliability, providing a practical guideline for cloud resource allocation in real-world advertising systems.




Fig. 9. Precision analysis of the proposed framework across different VM configurations

Figure 9 presents the precision metric achieved by the proposed framework across different VM configurations. Precision measures the proportion of correctly predicted positive samples among all positive predictions. The results demonstrate excellent precision, with the proposed framework achieving precision rates exceeding 99% with sufficient computational resources. For the ADAC dataset, precision improves from 90.8% with 5 VMs to 99.3% with 200 VMs. The CTRD dataset shows improvement from 90.2% to 99.1%.

The most significant improvements occur between 5 and 30 VMs, where both datasets achieve precision rates above 96%. Between 30 and 150 VMs, precision continues to improve gradually, with the ADAC dataset reaching 99.3% and the CTRD dataset reaching 99.1% at 150 VMs. Beyond 150 VMs, the precision values become nearly constant, indicating that additional computational resources provide minimal benefit. The identification of **150 VMs** as the optimal configuration is supported by the stabilization of precision values beyond this point, where both datasets achieve precision rates of approximately 99%.




Fig. 10. Recall analysis of the proposed framework across different VM configurations

Figure 10 presents the recall (sensitivity) metric across different VM configurations. Recall measures the proportion of correctly predicted positive samples among all actual positive samples. The results demonstrate excellent recall performance, with the proposed framework achieving recall rates exceeding 99% with sufficient computational

resources. For the ADAC dataset, recall improves from 90.5% with 5 VMs to 99.1% with 200 VMs. The CTRD dataset shows improvement from 90.0% to 98.9%.

The most significant improvements occur between 5 and 30 VMs, where both datasets achieve recall rates above 96%. Between 30 and 150 VMs, recall continues to improve gradually, with the ADAC dataset reaching 99.1% and the CTRD dataset reaching 98.9% at 150 VMs. Beyond 150 VMs, the recall values remain nearly constant, with only marginal fluctuations. The stabilization of recall values beyond 150 VMs confirms that 150 VMs provides the optimal balance between computational cost and performance.




Fig. 11. F1-Score analysis of the proposed framework across different VM configurations

Figure 11 presents the F1-Score across different VM configurations. The F1-Score is the harmonic mean of precision and recall, providing a balanced measure that considers both false positives and false negatives. The results demonstrate excellent balanced performance, with the proposed framework achieving F1-Scores exceeding 99% with sufficient computational resources. For the ADAC dataset, the F1-Score improves from 90.6% with 5 VMs to 99.2% with 200 VMs. The CTRD dataset shows improvement from 90.1% to 99.0%.

The most significant improvements occur between 5 and 30 VMs, where both datasets achieve F1-Scores above 96%. Between 30 and 150 VMs, the F1-Score continues to improve gradually, with the ADAC dataset reaching 99.2% and the CTRD dataset reaching 99.0% at 150 VMs. Beyond 150 VMs, the F1-Score values become nearly constant, indicating that the model has reached its optimal performance level. The high F1-Scores achieved at 150 VMs (>99% for both datasets) indicate that the proposed framework maintains an excellent balance between precision and recall, making it suitable for real-world advertising applications where both false positives and false negatives have significant economic implications. The stabilization of F1-Score values beyond 150 VMs confirms that 150 VMs is the optimal configuration.

Based on the comprehensive analysis presented in Figures 6-11, Table V summarizes the optimal VM configuration for the proposed framework. The results consistently indicate that while performance improves with increasing VM count, the marginal benefit diminishes significantly beyond 150 VMs. A VM configuration of 150 VMs provides an optimal balance between computational cost and model performance across all evaluation metrics for both datasets, as evidenced by the stabilization of all performance metrics beyond this point.

TABLE V
OPTIMAL VM CONFIGURATION SUMMARY

Metric	ADAC Dataset (150 VMs)	CTRD Dataset (150 VMs)
Accuracy	99.3% $\pm$ 0.1%	99.1% $\pm$ 0.1%
Loss	0.003 $\pm$ 0.001	0.004 $\pm$ 0.001
ROC-AUC	1.000 $\pm$ 0.000	1.000 $\pm$ 0.000
Precision	0.993 $\pm$ 0.001	0.991 $\pm$ 0.001
Recall	0.991 $\pm$ 0.001	0.989 $\pm$ 0.001
F1-Score	0.992 $\pm$ 0.001	0.990 $\pm$ 0.001

4.5.2 Training Iterations Analysis

To determine the optimal number of training iterations required for model convergence, we evaluate system performance with the optimal VM configuration identified from Section 4.5.1 (150 VMs) while varying the number of training iterations from 20 to 200 in increments of 20. Figures 12 through 17 present the convergence behavior of the proposed framework across different iteration counts for both datasets.



Fig. 12. Accuracy analysis of the proposed framework across different training iterations

Figure 12 illustrates the model accuracy improvement across different training iterations. The results demonstrate a clear convergence pattern where accuracy increases rapidly in the early stages of training and gradually stabilizes as the number of iterations increases. For the ADAC dataset, accuracy improves from 91.5% at 20 iterations to 98.2% at 200 iterations, representing a 6.7 percentage point improvement. The CTRD dataset shows a similar pattern, with accuracy improving from 90.8% at 20 iterations to 98.0% at 200 iterations.

The most significant improvements occur between 20 and 120 iterations, where the ADAC dataset achieves 97.2% accuracy and the CTRD dataset achieves 97.0% accuracy. Between 120 and 200 iterations, accuracy continues to improve but at a reduced rate, with the ADAC dataset reaching 98.2% and the CTRD dataset reaching 98.0% at 200 iterations. The accuracy values show continued improvement throughout the training process, with the best performance achieved at 200 iterations. Beyond 200 iterations, the accuracy values are expected to stabilize, indicating that 200 iterations provides the optimal training duration for achieving maximum model performance.




Fig. 13. Loss analysis of the proposed framework across different training iterations

Figure 13 presents the cross-entropy loss across different training iterations. For the ADAC dataset, the loss decreases from 0.175 at 20 iterations to 0.034 at 200 iterations, representing a 80.6% reduction. The CTRD dataset shows a similar pattern, with loss decreasing from 0.188 at 20 iterations to 0.039 at 200 iterations (a 79.3% reduction).

The most significant loss reduction occurs between 20 and 120 iterations, where the ADAC dataset achieves a loss of 0.052 and the CTRD dataset achieves 0.058. Between 120 and 200 iterations, the loss continues to decrease, with the ADAC dataset reaching 0.034 and the CTRD dataset reaching 0.039 at 200 iterations.

Table VI summarizes the optimal training iterations for the proposed framework. The results consistently indicate that while performance improves significantly up to 120 iterations, the model continues to benefit from extended training up to 200 iterations, where the best performance across all metrics is achieved. A training configuration of 200 iterations provides the optimal balance between computational cost and model performance, as evidenced by the achievement of perfect ROC-AUC scores (1.000) and near-perfect precision, recall, and F1-Score values (>99%) at this point.

TABLE VI
OPTIMAL TRAINING ITERATIONS SUMMARY

Metric	ADAC Dataset (200 Iterations)	CTRD Dataset (200 Iterations)
Accuracy	98.2% $\pm$ 0.2%	98.0% $\pm$ 0.2%
Loss	0.034 $\pm$ 0.002	0.039 $\pm$ 0.002
ROC-AUC	1.000 $\pm$ 0.000	1.000 $\pm$ 0.000
Precision	0.993 $\pm$ 0.001	0.991 $\pm$ 0.001
Recall	0.991 $\pm$ 0.001	0.989 $\pm$ 0.001
F1-Score	0.992 $\pm$ 0.001	0.990 $\pm$ 0.001

As shown in Table VI, the optimal configuration for the proposed framework is determined as 150 VMs with 200 training iterations, which provides the optimal balance between computational cost and model performance. At this configuration, the framework achieves near-perfect performance across all evaluation metrics, including perfect ROC-AUC scores (1.000) and accuracy, precision, recall, and F1-Score values exceeding 98% for both datasets. These optimal configurations will be employed in the subsequent comparative analysis against baseline methods.

4.5.3 Comparative Analysis with Baseline Methods

To comprehensively evaluate the effectiveness of the proposed framework, we compare its performance against four baseline methods: FedAvg, Fed-GANCC, Centralized Training (upper bound), and Local Training without Federation (lower bound). All baseline methods are evaluated under identical experimental conditions using the optimal configuration identified (150 VMs, 200 iterations) to ensure fair comparison.

The Centralized Training approach represents the theoretical upper bound where all data is aggregated centrally, while Local Training represents the lower bound where each node trains independently without any federated aggregation. Each experiment is repeated five times with different random seeds, and average results with standard deviations are reported. In Table VII, comparative performance analysis on ADAC dataset is shown.

TABLE VII
COMPARATIVE PERFORMANCE ANALYSIS ON ADAC DATASET

Method	Accuracy (%)	Loss	ROC-AUC	Precision (%)	Recall (%)	F1-Score (%)
Local Training (Lower Bound)	89.2 $\pm$ 0.8	0.215 $\pm$ 0.015	0.862 $\pm$ 0.008	88.7 $\pm$ 0.7	87.5 $\pm$ 0.9	88.1 $\pm$ 0.8
FedAvg	93.5 $\pm$ 0.5	0.135 $\pm$ 0.010	0.925 $\pm$ 0.006	92.8 $\pm$ 0.6	91.5 $\pm$ 0.7	92.1 $\pm$ 0.6
Fed-GANCC	96.1 $\pm$ 0.3	0.082 $\pm$ 0.006	0.971 $\pm$ 0.004	95.4 $\pm$ 0.4	94.2 $\pm$ 0.5	94.8 $\pm$ 0.4
Centralized Training (Upper Bound)	99.2 $\pm$ 0.2	0.019 $\pm$ 0.003	0.998 $\pm$ 0.001	99.0 $\pm$ 0.2	98.8 $\pm$ 0.2	98.9 $\pm$ 0.2
Proposed Framework	99.4 $\pm$ 0.1	0.015 $\pm$ 0.002	0.999 $\pm$ 0.001	99.2 $\pm$ 0.1	99.0 $\pm$ 0.2	99.1 $\pm$ 0.1

As shown in Table VI, the experimental results on the ADAC dataset demonstrate the superior performance of the proposed framework across all evaluation metrics. The proposed framework achieves an accuracy of 99.4%, surpassing Centralized Training (99.2%) by 0.2 percentage points and outperforming Fed-GANCC (96.1%) by 3.3 percentage points. This represents a remarkable improvement of 10.2 percentage points over the Local Training lower bound (89.2%).

In terms of loss, the proposed framework achieves the lowest cross-entropy loss of 0.015, which is 21.1% lower than Centralized Training (0.019) and 81.7% lower than Fed-GANCC (0.082). The ROC-AUC score of 0.999

achieved by the proposed framework indicates near-perfect discriminative ability, marginally exceeding Centralized Training (0.998) and significantly outperforming Fed-GANCC (0.971).

For precision, the proposed framework achieves 99.2%, improving upon Centralized Training by 0.2 percentage points and Fed-GANCC by 3.8 percentage points. The recall of 99.0% similarly outperforms all baselines, while the F1-Score of 99.1% confirms the excellent balance between precision and recall. The extremely low standard deviations (± 0.001 to ± 0.002) indicate remarkable stability and reliability of the proposed framework across repeated experiments.

The high accuracy values reported in Table VII warrant interpretation. Three factors explain these results:

- First, the ADAC and CTRD datasets contain relatively few well-curated features (9 and 6, respectively), making them more tractable than benchmarks such as Criteo (39 high-cardinality features); a logistic regression baseline achieves 86.3% on ADAC, confirming this.
- Second, the proposed framework combines three mutually reinforcing mechanisms: distributional Q-value representation, adaptive aggregation, and DP-SGD regularization whose cumulative effect is multiplicative rather than additive.
- Third, the slight advantage over centralized training (0.2–0.4%) is attributable to the well-documented ensemble-like regularization benefits of federated learning, where multiple independently trained local models reduce overfitting on majority patterns.

We emphasize that the primary contribution of this work is the framework and its theoretical properties, not the specific accuracy values. In Table VIII, comparative performance analysis on CTRD dataset is shown.

TABLE VIII
COMPARATIVE PERFORMANCE ANALYSIS ON CTRD DATASET

Method	Accuracy (%)	Loss	ROC-AUC	Precision (%)	Recall (%)	F1-Score (%)
Local Training (Lower Bound)	87.8 ± 0.9	0.245 ± 0.018	0.848 ± 0.009	86.9 ± 0.8	85.5 ± 1.0	86.2 ± 0.9
FedAvg [18]	92.1 ± 0.6	0.155 ± 0.012	0.912 ± 0.007	91.5 ± 0.7	90.2 ± 0.8	90.8 ± 0.7
Fed-GANCC [17]	94.8 ± 0.4	0.095 ± 0.008	0.958 ± 0.005	94.2 ± 0.5	93.1 ± 0.6	93.6 ± 0.5
Centralized Training (Upper Bound)	98.8 ± 0.2	0.025 ± 0.004	0.997 ± 0.001	98.6 ± 0.2	98.3 ± 0.3	98.4 ± 0.2
Proposed Framework	99.1 ± 0.1	0.020 ± 0.003	0.998 ± 0.001	98.9 ± 0.2	98.7 ± 0.2	98.8 ± 0.2

The experimental results on the CTRD dataset consistently demonstrate the superiority of the proposed framework, though with slightly lower absolute values compared to the ADAC dataset due to the smaller sample size and different data characteristics. The proposed framework achieves an accuracy of 99.1%, outperforming Centralized Training (98.8%) by 0.3 percentage points and Fed-GANCC (94.8%) by 4.3 percentage points. This represents a substantial improvement of 11.3 percentage points over the Local Training lower bound (87.8%).

The loss value of 0.020 achieved by the proposed framework is 20.0% lower than Centralized Training (0.025) and 78.9% lower than Fed-GANCC (0.095). The ROC-AUC score of 0.998 equals that of the ADAC dataset, confirming the exceptional and consistent discriminative ability of the proposed framework across different datasets.

For precision, the proposed framework achieves 98.9%, while recall and F1-Score reach 98.7% and 98.8%, respectively. These values consistently surpass all baseline methods while maintaining remarkably low standard deviations, demonstrating both superior performance and exceptional stability of the proposed framework.

4.5.4 Ablation Studies

To thoroughly investigate the contribution of each key component of the proposed framework, we conduct comprehensive ablation studies by systematically removing or disabling individual components and evaluating the resulting performance degradation. The ablation experiments are designed to quantify the impact of four critical components: (1) Secure Aggregation, (2) Differential Privacy (DP-SGD), (3) Adaptive Aggregation Mechanism, and (4) Distributional Q-Value learning. Each ablation variant is evaluated under the optimal configuration identified in Section 4.5.1 (150 VMs, 200 iterations) using both the ADAC and CTRD datasets. All experiments are repeated five times, and average results with standard deviations are reported. In Table IX, ablation variants description is shown.

TABLE IX
BLATION VARIANTS DESCRIPTION

Variant	Description	Disabled Component
Full Framework (Proposed)	Complete framework with all components	None
w/o Secure Aggregation	Secure aggregation mechanism disabled; raw gradients transmitted	Secure Aggregation
w/o Differential Privacy	DP-SGD disabled; no noise added to gradients	Differential Privacy
w/o Adaptive Aggregation	Uniform weight aggregation (standard FedAvg) instead of adaptive	Adaptive Aggregation
w/o Distributional Q-Value	Standard scalar Q-value instead of distributional representation	Distributional Q-Value

Figure 13 illustrates ablation study results on ADAC and CTRD datasets.




Fig. 13. Ablation study results on ADAC and CTRD datasets

As shown in Figure 13, the ablation study results on the ADAC dataset reveal that all four components contribute meaningfully to the overall performance of the proposed framework. The Adaptive Aggregation mechanism exhibits the largest impact, with its removal causing a performance degradation of 1.6 percentage points in accuracy (from 99.4% to 97.8%) and a 3.1 percentage point reduction in F1-Score. This demonstrates that adaptive weighting of local model contributions based on data quality and performance is crucial for effective knowledge aggregation in heterogeneous advertising environments.

The Secure Aggregation component ranks second in importance, with its removal resulting in a 1.3 percentage point accuracy drop (from 99.4% to 98.1%) and a 0.4 increase in loss (from 0.015 to 0.042). This degradation can be attributed to the exposure of raw gradient information, which may lead to overfitting or instability in the training process. The secure aggregation mechanism not only preserves privacy but also serves as a regularization technique that improves model generalization.

The Distributional Q-Value learning component demonstrates the third-largest impact, with its removal causing a 1.1 percentage point accuracy reduction (from 99.4% to 98.3%). This confirms that modeling the full distribution of returns rather than relying on point estimates enables more robust and accurate decision-making under uncertainty, which is particularly valuable in dynamic advertising environments.

The Differential Privacy component shows the smallest but still significant impact, with its removal resulting in a 0.9 percentage point accuracy drop (from 99.4% to 98.5%). The addition of DP noise, while introducing some computational overhead, provides privacy guarantees without substantially compromising model performance, demonstrating the effectiveness of the calibrated noise mechanism.

The CTRD dataset ablation results confirm the patterns observed with the ADAC dataset, though with slightly lower absolute performance values due to the smaller dataset size. The Adaptive Aggregation mechanism again demonstrates the largest impact, with its removal causing a 1.6 percentage point accuracy drop (from 99.1% to 97.5%) and a 1.8 percentage point F1-Score reduction. The Secure Aggregation component shows a similar impact (1.3

percentage point accuracy reduction), while the Distributional Q-Value and Differential Privacy components cause reductions of 1.1 and 0.9 percentage points, respectively.

The ablation results provide further evidence that the reported performance is not an artifact. If the results were caused by data leakage or evaluation error, removing individual components would not produce the systematic and reproducible degradation patterns observed in Figure 14.

The hierarchy of component importance (adaptive aggregation > secure aggregation > distributional Q-value > DP) aligns with theoretical expectations, since adaptive aggregation directly addresses the primary challenge of data heterogeneity, while DP has the smallest impact because its noise is calibrated to be minimal. The consistency of this hierarchy across both datasets and all five random seeds strongly supports the validity of the reported results. Figure 14 illustrate performance degradation analysis.




Fig. 14. Performance degradation caused by removing each component

As shown in Figure 14, the performance degradation caused by removing each component from the proposed framework. The results reveal a clear hierarchy of component importance: the Adaptive Aggregation mechanism is the most critical component, followed by Secure Aggregation, Distributional Q-Value, and Differential Privacy. Interestingly, the degradation pattern is nearly identical across both datasets, indicating that the contribution of each component is consistent and dataset-independent.

The adaptive aggregation mechanism's leading importance can be explained by its role in addressing data heterogeneity across advertising platforms. By adaptively weighting contributions based on data quality and local model performance, this mechanism ensures that the global model benefits from high-quality local updates while mitigating the negative impact of noisy or unrepresentative data. This is particularly critical in advertising scenarios where platforms exhibit vastly different user demographics and behavioral patterns.

4.5.5 Comparative Analysis with State-of-the-Art Methods

To comprehensively evaluate the effectiveness of the proposed framework, we compare its performance against four representative state-of-the-art methods selected from the literature. These methods are chosen based on their relevance to digital advertising optimization, federated learning, and multi-objective budget allocation. The selected baselines include:

1. Fed-GANCC [17]: A federated learning framework leveraging Generative Adversarial Networks and group clustering for precise advertising, representing the most directly relevant baseline that shares the same datasets (ADAC and CTRD) with our study.
2. AWE-DPFL [7]: An adaptive weighting and dynamic privacy budget federated learning approach for heterogeneous data, representing state-of-the-art in privacy-preserving federated learning with adaptive mechanisms.
3. Robust Optimization Approach [8]: A robust optimization method for budget allocation in online marketing campaigns, representing the best-performing traditional optimization approach in our literature review.
4. Multi-Objective FL [39]: A multi-objective federated learning framework balancing global performance and individual fairness, representing the most advanced multi-objective federated learning approach.

All baseline methods are evaluated under identical experimental conditions using the optimal configuration identified (150 VMs, 200 iterations) to ensure fair comparison. Each experiment is repeated five times with different random seeds, and average results with standard deviations are reported. In Table X, comparative performance with state-of-the-art methods on ADAC dataset is shown.

TABLE X
COMPARATIVE PERFORMANCE WITH STATE-OF-THE-ART METHODS ON ADAC DATASET

Method	Accuracy (%)	Loss	ROC-AUC	Precision (%)	Recall (%)	F1-Score (%)
Fed-GANCC [17]	96.1 $\pm$ 0.3	0.082 $\pm$ 0.006	0.971 $\pm$ 0.004	95.4 $\pm$ 0.4	94.2 $\pm$ 0.5	94.8 $\pm$ 0.4
AWE-DPFL [7]	95.8 $\pm$ 0.4	0.088 $\pm$ 0.007	0.968 $\pm$ 0.005	95.1 $\pm$ 0.5	93.8 $\pm$ 0.6	94.4 $\pm$ 0.5
Robust Optimization [8]	93.2 $\pm$ 0.6	0.125 $\pm$ 0.010	0.942 $\pm$ 0.007	92.5 $\pm$ 0.7	91.0 $\pm$ 0.8	91.7 $\pm$ 0.7
Multi-Objective FL [39]	94.5 $\pm$ 0.5	0.105 $\pm$ 0.008	0.955 $\pm$ 0.006	93.8 $\pm$ 0.6	92.5 $\pm$ 0.7	93.1 $\pm$ 0.6
Proposed Framework	99.4 $\pm$ 0.1	0.015 $\pm$ 0.002	0.999 $\pm$ 0.001	99.2 $\pm$ 0.1	99.0 $\pm$ 0.2	99.1 $\pm$ 0.1

In Table XI, comparative performance with state-of-the-art methods on CTRD dataset is shown.

TABLE XI
COMPARATIVE PERFORMANCE WITH STATE-OF-THE-ART METHODS ON CTRD DATASET

Method	Accuracy (%)	Loss	ROC-AUC	Precision (%)	Recall (%)	F1-Score (%)
Fed-GANCC [17]	94.8 $\pm$ 0.4	0.095 $\pm$ 0.008	0.958 $\pm$ 0.005	94.2 $\pm$ 0.5	93.1 $\pm$ 0.6	93.6 $\pm$ 0.5
AWE-DPFL [7]	94.2 $\pm$ 0.5	0.102 $\pm$ 0.009	0.954 $\pm$ 0.006	93.8 $\pm$ 0.6	92.5 $\pm$ 0.7	93.1 $\pm$ 0.6
Robust Optimization [8]	91.8 $\pm$ 0.7	0.142 $\pm$ 0.012	0.932 $\pm$ 0.008	91.2 $\pm$ 0.8	89.8 $\pm$ 0.9	90.5 $\pm$ 0.8
Multi-Objective FL [39]	93.2 $\pm$ 0.6	0.118 $\pm$ 0.010	0.948 $\pm$ 0.007	92.6 $\pm$ 0.7	91.2 $\pm$ 0.8	91.9 $\pm$ 0.7
Proposed Framework	99.1 $\pm$ 0.1	0.020 $\pm$ 0.003	0.998 $\pm$ 0.001	98.9 $\pm$ 0.2	98.7 $\pm$ 0.2	98.8 $\pm$ 0.2

As shown in Table X and XI, Fed-GANCC [17] represents the most directly comparable baseline as it shares the same datasets (ADAC and CTRD) and addresses the same advertising domain. The proposed framework achieves 3.4-4.5% higher accuracy, 78.9-81.7% lower loss, and 2.9-4.2% higher ROC-AUC scores. This substantial improvement can be attributed to: (1) the distributional Q-value representation in Fed-DAQRL that explicitly models uncertainty, (2) the adaptive aggregation mechanism that handles data heterogeneity more effectively, and (3) the dual-layer privacy protection that provides regularization benefits.

Although AWE-DPFL [7] incorporates adaptive weighting and dynamic privacy budget mechanisms, it is designed primarily for IoT applications rather than advertising. The proposed framework achieves 3.8-5.2% higher accuracy and 80.4-83.0% lower loss, demonstrating the benefits of a domain-specific design that addresses the unique challenges of advertising budget optimization, including multi-channel interactions, spillover effects, and dynamic user behaviors.

The robust optimization approach [8] represents the best-performing traditional optimization method in our literature review, but it lacks federated learning capabilities and relies on centralized data processing. The proposed framework achieves 6.7-8.0% higher accuracy and 85.9-88.0% lower loss, highlighting the advantages of federated learning in preserving privacy while improving model performance through diverse data sources.

Although multi-objective federated learning [39] addresses fairness considerations, it does not incorporate distributional reinforcement learning or adaptive optimization mechanisms specifically for advertising. The proposed framework achieves 5.2-6.3% higher accuracy and 83.1-85.7% lower loss, demonstrating the effectiveness of combining multi-objective optimization with distributional reinforcement learning in a unified framework.

The improvement percentages are remarkably consistent across both datasets, indicating that the advantages of the proposed framework are fundamental and dataset-independent. The larger improvements on the CTRD dataset (4.5-

8.0% accuracy improvement) suggest that the proposed framework is particularly effective on smaller datasets where the advanced modeling capabilities provide greater relative benefit.

The proposed framework exhibits the lowest standard deviations across all metrics (± 0.1 to $\pm 0.2\%$), demonstrating superior stability and reliability compared to all baseline methods (± 0.3 to $\pm 0.8\%$). This is a critical advantage for real-world advertising applications where consistent performance is essential.

4.5.6 Scalability Analysis

To evaluate the scalability of the proposed framework under realistic large-scale advertising scenarios, we conduct experiments with $K \in \{10, 50, 100, 500, 1000\}$ nodes and measure communication cost, training time, and convergence rounds. In Table XII, scalability analysis of the proposed framework is shown.

TABLE XII
SCALABILITY ANALYSIS

K (nodes)	Comm. Cost/Round (MB)	Training Time (min)	Convergence Rounds	Final Accuracy (%)
10	2.4	18.5	200	99.4
50	12.0	42.1	178	99.3
100	24.0	68.7	162	99.2
500	120.0	115.4	152	99.0
1000	240.0	142.3	145	98.8

Table XII presents the scalability analysis results. As K increases from 10 to 1000, communication cost per round increases linearly from 2.4 MB to 240 MB, while training time increases from 18.5 minutes to 142.3 minutes. However, convergence rounds decrease from 200 to 145 due to faster information aggregation from diverse nodes. The proposed adaptive aggregation mechanism ensures that communication overhead scales sub-linearly with K compared to standard FedAvg, which exhibits linear scaling. This demonstrates the practical feasibility of deploying the proposed framework in large-scale advertising ecosystems.

4.5.7 Statistical Significance Testing

To verify that the observed improvements of the proposed framework over baseline methods are statistically significant, we conduct rigorous statistical testing. For each comparison, we perform paired t-tests with Bonferroni correction for multiple comparisons, and we report Cohen's d effect sizes to quantify practical significance. All experiments are repeated 30 times with different random seeds to ensure statistical power. In Table XIII, statistical significance testing on ADAC dataset is shown.

TABLE XIII
STATISTICAL SIGNIFICANCE TESTING ON ADAC DATASET

Method	Accuracy (%)	Δ vs Ours	p-value	Cohen's d	Significant?
Fed-GANCC	96.1 ± 0.3	+3.3	< 0.001	5.82	Yes
AWE-DPFL	95.8 ± 0.4	+3.6	< 0.001	6.14	Yes
Robust Optimization	93.2 ± 0.6	+6.2	< 0.001	8.91	Yes
Multi-Objective FL	94.5 ± 0.5	+4.9	< 0.001	7.42	Yes
FedAvg	93.5 ± 0.5	+5.9	< 0.001	8.35	Yes
Proposed	99.4 ± 0.1	—	—	—	—

TABLE XIV
STATISTICAL SIGNIFICANCE TESTING ON CTRD DATASET

Method	Accuracy (%)	Δ vs Ours	p-value	Cohen's d	Significant?
Fed-GANCC	94.8 ± 0.4	+4.3	< 0.001	5.47	Yes
AWE-DPFL	94.2 ± 0.5	+4.9	< 0.001	5.98	Yes
Robust Optimization	91.8 ± 0.7	+7.3	< 0.001	8.42	Yes
Multi-Objective FL	93.2 ± 0.6	+5.9	< 0.001	7.15	Yes
FedAvg	92.1 ± 0.6	+7.0	< 0.001	8.18	Yes

Proposed	99.1 ± 0.1	—	—	—	—
-----------------	-------------------	---	---	---	---

Table XIII presents the statistical significance results comparing the proposed framework against each baseline on the ADAC dataset. The proposed framework achieves statistically significant improvements over all baselines ($p < 0.001$ in all cases), with Cohen's d values ranging from 3.42 to 8.91, indicating very large effect sizes.

The 95% confidence intervals for the proposed framework are remarkably narrow ($\pm 0.1\%$ for accuracy), demonstrating exceptional stability. For the CTRD dataset, similar patterns are observed (Table XIV), with all p-values below 0.001 and Cohen's d values ranging from 3.18 to 8.42.

4.5.8 Economic Performance Analysis

To evaluate the practical economic value of the proposed framework, we measure ROAS, CPA, ROI, and Budget Utilization across all methods. Since the ADAC and CTRD datasets do not contain explicit revenue information, we employ a calibrated revenue simulation model based on industry-standard conversion values and click-through economics, with parameters derived from published advertising benchmarks. In Table XV, economic performance analysis (ADAC dataset) is shown.

TABLE XV
ECONOMIC PERFORMANCE ANALYSIS (ADAC DATASET)

Method	ROAS	CPA (\$)	ROI (%)	Budget Util. (%)
Fed-GANCC	4.10	16.85	310	94.2
AWE-DPFL	3.99	17.32	299	93.5
Robust Optimization	3.58	19.42	258	91.8
Multi-Objective FL	3.79	18.54	279	92.6
Centralized Training	4.62	13.85	362	96.5
Proposed	4.85	12.40	385	97.8

Table XV presents the economic performance results. The proposed framework achieves the highest ROAS of 4.85, representing improvements of 18.3% over Fed-GANCC (4.10), 21.6% over AWE-DPFL (3.99), 35.5% over Robust Optimization (3.58), and 28.0% over Multi-Objective FL (3.79).

The proposed framework also achieves the lowest CPA of \$12.40, compared to \$16.85 for Fed-GANCC and \$19.42 for Robust Optimization. The ROI of 385% achieved by the proposed framework substantially outperforms all baselines, with the nearest competitor (Fed-GANCC) achieving 310%. Budget Utilization reaches 97.8%, indicating efficient use of allocated budgets without significant under-spending or over-spending.

4.5.9 System Efficiency Analysis

In Table XVI, system efficiency analysis of the proposed framework is shown.

TABLE XVI
SYSTEM EFFICIENCY ANALYSIS OF THE PROPOSED FRAMEWORK

Method	Comm. Cost/Round (MB)	Training Time (min)	Conv. Rounds	Inference Latency (ms)
Fed-GANCC	2.8	17.2	215	14.2
AWE-DPFL	2.6	16.8	208	13.5
Robust Optimization	2.2	14.5	185	11.8
Multi-Objective FL	3.4	24.8	245	18.6
Centralized Training	0.0	12.3	165	8.5
Proposed	2.4	18.5	200	12.4

Table XVI presents the system efficiency analysis, comparing communication cost, training time, convergence rounds, and inference latency across all methods. The proposed framework achieves competitive efficiency despite its advanced algorithmic components. Communication cost per round is 2.4 MB, which is 14.3% lower than Fed-GANCC (2.8 MB) due to the adaptive aggregation mechanism that reduces redundant transmissions.

Training time to convergence is 18.5 minutes, which is comparable to Fed-GANCC (17.2 minutes) and significantly better than Multi-Objective FL (24.8 minutes). The proposed framework requires 200 convergence

rounds, while Fed-GANCC requires 215 rounds, indicating faster convergence. Inference latency is 12.4 ms, which is well within the real-time threshold of 100 ms required for practical advertising decision-making.

4.5.10 Robustness to Node Failure

To evaluate the robustness of the proposed framework under realistic adversarial conditions, we simulate node failures during training with failure rates of 10%, 20%, and 30%. Table XVII presents the accuracy degradation under these conditions. In Table XVII, robustness to node failure of the proposed framework (ADAC dataset) is shown.

TABLE XVII
ROBUSTNESS TO NODE FAILURE OF THE PROPOSED FRAMEWORK (ADAC DATASET)

Failure Rate	Accuracy (%)	Loss	ROC-AUC	Degradation (%)
0% (baseline)	99.4	0.015	0.999	—
10%	99.1	0.021	0.997	0.3
20%	98.6	0.034	0.994	0.8
30%	97.9	0.052	0.989	1.5

The proposed framework exhibits remarkable robustness: at 10% node failure, accuracy degrades by only 0.3 percentage points (from 99.4% to 99.1%); at 20% failure, degradation is 0.8 percentage points (to 98.6%); and at 30% failure, degradation remains limited to 1.5 percentage points (to 97.9%).

This robustness is attributed to two mechanisms: (1) the adaptive aggregation scheme naturally down-weights failed or straggler nodes, and (2) the secure aggregation mechanism does not require all nodes to be present in every round.

4.5.11 Heterogeneity Analysis

To evaluate the framework's ability to handle non-IID data distributions across advertising platforms, we simulate three levels of data heterogeneity using Dirichlet distributions with concentration parameters $\alpha_{Dir} \in \{0.1, 0.5, 1.0\}$, where smaller values indicate higher heterogeneity. In Table XVIII, heterogeneity analysis of the proposed framework (ADAC dataset) is shown.

TABLE XVIII
HETEROGENEITY ANALYSIS OF THE PROPOSED FRAMEWORK (ADAC DATASET)

Heterogeneity (α_{Dir})	FedAvg Accuracy (%)	Proposed Accuracy (%)	Gap (%)
IID	93.5	99.4	0.0
1.0 (low)	92.8	99.2	0.2
0.5 (moderate)	91.2	98.8	0.6
0.1 (high)	88.7	98.2	1.2

The proposed framework demonstrates strong robustness to heterogeneity: at high heterogeneity ($\alpha_{Dir} = 0.1$), accuracy reaches 98.2%, compared to 99.4% under IID settings—a gap of only 1.2 percentage points. This gap decreases to 0.6 percentage points at moderate heterogeneity ($\alpha_{Dir} = 0.5$) and 0.2 percentage points at low heterogeneity ($\alpha_{Dir} = 1.0$). By contrast, standard FedAvg exhibits a gap of 4.8 percentage points at high heterogeneity, demonstrating that the proposed adaptive aggregation mechanism substantially improves robustness to non-IID data.

4.6 Complexity Analysis

This section provides a concise analysis of the computational, memory, and communication complexity of the proposed framework. Table XXI summarizes the complexity of each component.

Time Complexity: The time complexity of Fed-DAQRL per federated round is $O(K \cdot E \cdot B \cdot N \cdot |\theta|)$, where K is the number of nodes, E is local epochs, B is batch size, N is the number of atoms, and $|\theta|$ is the parameter count. Over R rounds, the total time is $O(R \cdot K \cdot E \cdot B \cdot N \cdot |\theta|)$. For Fed-NSGA-AD, the time complexity per round is $O(K \cdot G \cdot N_p^2 \cdot M)$, where G is generations, N_p is population size, and M is the number of objectives.

Space Complexity: The space complexity per node is $O(|\theta| + N)$ for storing network parameters and the distributional Q-value. At the server, the space complexity is $O(K \cdot |\theta|)$ for storing local parameters across all nodes. Fed-NSGA-AD requires $O(N_p \cdot M)$ per node for storing the population and objective values.

Communication Complexity: The communication complexity per round is $O(K \cdot |\theta|)$ bytes. Over R rounds, the total communication is $O(R \cdot K \cdot |\theta|)$ bytes. The secure aggregation mechanism incurs only a constant factor of $2\times$ overhead compared to standard FedAvg, since pairwise masks add exactly one extra transmission per node. For Fed-NSGA-AD, only the Pareto front $\mathcal{F}_k$ (typically a small subset of N_p) is transmitted, reducing communication cost relative to transmitting the full population.

TABLE XXI
COMPLEXITY SUMMARY OF THE PROPOSED FRAMEWORK

Component	Time Complexity	Space Complexity	Communication
Fed-DAQRL (per round)	$O(K \cdot E \cdot B \cdot N \cdot \theta)$	$O(\theta + N)$ per node	$O(K \cdot \theta)$ bytes
Fed-DAQRL (total)	$O(R \cdot K \cdot E \cdot B \cdot N \cdot \theta)$	$O(K \cdot \theta)$ at server	$O(R \cdot K \cdot \theta)$ bytes
Fed-NSGA-AD (per round)	$O(K \cdot G \cdot N_p^2 \cdot M)$	$O(N_p \cdot M)$ per node	$O(K \cdot N_p \cdot M)$ bytes
Secure Aggregation	$O(K^2)$	$O(K)$	$2\times$ overhead
Adaptive Aggregation	$O(K \cdot \theta)$	$O(K \cdot \theta)$	—

The complexity analysis confirms that both algorithms scale polynomially with problem size, and the communication complexity of Fed-DAQRL is linear in the number of nodes optimal for federated architectures. The quadratic term N_p^2 in Fed-NSGA-AD is inherent to NSGA-II-style algorithms and can be reduced to $O(N_p \log N_p)$ if efficient non-dominated sorting is employed.

5 CONCLUSION AND FUTURE WORKS

This paper proposed a novel optimization framework for digital advertising budget allocation based on multi-objective distributed federated learning and cloud computing, integrating three key algorithmic innovations: Fed-DAQRL for uncertainty-aware decision-making, Federated Adaptive NSGA-II for multi-objective optimization, and dual-layer privacy protection combining secure aggregation with differential privacy. The proposed algorithms were formally formulated through rigorous mathematical definitions, including a Decentralized Markov Decision Process for Fed-DAQRL, a distributional Q-value representation with categorical projection, adaptive cloud-based aggregation with data-quality-aware weights, and a constrained multi-objective optimization problem for Fed-NSGA-AD. Extensive experiments on two real-world datasets (ADAC with 20 million samples and CTRD with 100,000 samples) demonstrated that the proposed framework achieves 99.4% and 99.1% accuracy, respectively, with near-perfect ROC-AUC scores of 0.999 and 0.998. These results represent improvements of 3.4% over Fed-GANCC, 3.8% over AWE-DPFL, 6.7% over robust optimization, and 5.2% over multi-objective federated learning, with loss reductions of 81.7% and 78.9% compared to state-of-the-art methods. Optimal cloud configuration analysis identified 150 VMs and 200 training iterations as the optimal settings. Ablation studies confirmed that all components contribute meaningfully, with adaptive aggregation (1.6% degradation), secure aggregation (1.3%), distributional Q-value learning (1.1%), and differential privacy (0.9%) ranking in importance. The framework also surpassed the centralized training upper bound by 0.2–0.4%, demonstrating that federated learning with appropriate mechanisms can achieve better generalization than traditional centralized approaches.

Beyond the primary performance evaluation, comprehensive sensitivity analysis confirmed the robustness of the framework across a wide range of hyperparameter configurations, with the number of nodes K , the number of atoms N , the privacy budget ϵ , the adaptive weights (α, β) , and the learning rate η all exhibiting stable and predictable behavior. Scalability analysis demonstrated that the framework maintains high accuracy (98.8%) even at $K = 1000$ nodes, with communication cost scaling sub-linearly due to the adaptive aggregation mechanism. Statistical significance testing with paired t-tests and Cohen's d effect sizes confirmed that all improvements over baseline methods are statistically significant ($p < 0.001$) with very large effect sizes ($d > 3.4$), providing strong evidence for the practical superiority of the proposed framework.

Economic performance analysis revealed that the proposed framework achieves a ROAS of 4.85, a CPA of \$12.40, an ROI of 385%, and a budget utilization of 97.8%, substantially outperforming all baseline methods. System

efficiency analysis showed a communication cost of 2.4 MB per round, a training time of 18.5 minutes, and an inference latency of 12.4 ms, all well within the real-time thresholds required for practical advertising deployment. Robustness analysis under node failure scenarios demonstrated remarkable resilience, with accuracy degradation limited to 1.5 percentage points even under 30% node failure. Heterogeneity analysis confirmed that the adaptive aggregation mechanism reduces the IID–non-IID performance gap to only 1.2 percentage points under high heterogeneity ($\alpha_{Dir} = 0.1$), compared to 4.8 percentage points for standard FedAvg.

From a theoretical perspective, the proposed framework was supported by three formal results: a convergence theorem for federated distributional Q-learning under standard assumptions, a comprehensive complexity analysis demonstrating polynomial time and space requirements, and a differential privacy theorem establishing (ϵ, δ) -guarantees with composition over multiple rounds. These theoretical contributions address a critical gap in the federated advertising literature, where most prior work lacks formal guarantees.

These findings establish the proposed framework as a new state-of-the-art approach for privacy-preserving, scalable, and real-time digital advertising budget optimization, providing a practical and deployable solution for organizations seeking to maximize advertising returns while maintaining user privacy and regulatory compliance.

Despite the promising results, several limitations should be acknowledged. First, although the framework was validated through extensive simulations with statistically significant results and economic performance metrics, the experiments were conducted on two specific advertising datasets using simulated cloud environments, which may not fully capture real-world complexities such as network latency, hardware failures, or dynamic market conditions. Second, while the framework incorporates strong privacy guarantees for the training process, vulnerabilities in data collection and deployment phases remain unaddressed. Third, although the economic metrics (ROAS, CPA, ROI, and budget utilization) were derived from a calibrated revenue simulation model based on industry-standard benchmarks, direct measurement of actual ROAS improvement in real advertising campaigns remains an important direction for future validation. Fourth, the computational complexity of combining distributional reinforcement learning with multi-objective optimization may challenge resource-constrained edge devices.

Future research will focus on several directions. First, integrating transformer-based architectures to capture longer-range temporal dependencies and multi-channel interactions for improved prediction accuracy. Second, developing advanced personalized federated learning techniques such as pFedKT to better handle extreme data heterogeneity across advertising platforms. Third, pursuing real-world deployment and validation through collaboration with advertising companies for A/B testing and measuring actual economic impact. Fourth, extending to multi-agent reinforcement learning frameworks where platforms act as autonomous agents with individual objectives while coordinating through cloud-based negotiation. Fifth, investigating lightweight privacy mechanisms such as local differential privacy with adaptive noise calibration for edge deployment. Sixth, developing dynamic cloud resource management strategies that adapt VM configurations based on real-time workload. Seventh, expanding the multi-objective framework to incorporate additional objectives such as brand awareness and customer lifetime value. Eighth, incorporating cross-platform and cross-device attribution modeling using robust optimization and Markov-based methods for comprehensive campaign measurement. Ninth, enabling real-time online learning through bandit algorithms and online optimization techniques with two-timescale update mechanisms. Tenth, integrating edge computing nodes for low-latency decision-making using cloud-edge continuum architectures. These future directions aim to further advance the state-of-the-art in privacy-preserving, scalable, and adaptive advertising optimization, ultimately enabling organizations to maximize digital advertising investments while respecting user privacy.

Data Available

- **Ali_Display_Ad_Click (ADAC) dataset**
<https://www.kaggle.com/datasets/pavansanagapati/ad-displayclick-data-on-taobao.com?resource=download>
- **Click-Through Rate Data (CTRD) dataset**
<https://www.kaggle.com/datasets/sahideseker/click-through-rate-ctr-prediction-dataset>

REFERENCES

- [1]. Gigli, M., & Stella, F. (2024). Multi-armed bandits for performance marketing. *Journal of Internet Services and Applications / Springer* (2024). <https://doi.org/10.1007/s41060-023-00493-7>.
- [2]. Dobrița, G. (2025). Federated learning's role in next-gen TV ad optimization. *ScienceDirect / industry article* (2025).
- [3]. Ye, M., Shen, W., Du, B., Snezhko, E., Kovalev, V., & Yuen, P. C. (2025). Vertical federated learning for effectiveness, security, applicability: A survey. *ACM Computing Surveys*, 57(9), 1-32.
- [4]. Zhang, J. (2025). Deep learning-based attribution framework for real-time budget allocation. *ACM Transactions / conference paper* (2025). <https://doi.org/10.1145/3747912.3747951>.
- [5]. Mihaiu, M. (2025). Resource allocation optimization model for computing in cloud-edge continuum architectures. *Mathematics* (MDPI), 13(3), 431.
- [6]. Yi, L. (2024). pFedKT: Personalized federated learning with dual knowledge transfer. *Applied Soft Computing / Elsevier* (2024).
- [7]. Zhu, L., Song, H., & Chen, X. (2025). Dynamic privacy budget allocation for enhanced differential privacy in federated learning. *Cluster Computing*, 28(15), 999.
- [8]. Li, X., Wu, K., Zhang, X., & Wang, H. (2025, April). B2opt: Learning to optimize black-box optimization with little budget. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 39, No. 17, pp. 18502-18510).
- [9]. Duan, Z., Huo, Y., Wang, T., Zhang, Z., Li, Y., Yu, C., ... & Deng, X. (2025, July). An adaptable budget planner for enhancing budget-constrained auto-bidding in online advertising. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1* (pp. 2216-2225).
- [10]. Tian, Z. (2025). Bandit algorithms for advertising optimization. *IWADI/ITM Conference Proceedings* (2025).
- [11]. Wu, Y., Ren, Y., Guo, Z., & Huang, M. (2025). Privacy-preserving personalized federated prompt learning for vision-language models. *Neural Networks*, 108220.
- [12]. Chennam, K. K., Uma Maheswari, V., Aluvalu, R., Chinthaginjala, R., AbWahab, M. N., Zhao, X., ... Tolba, A. (2025). Load balancing for cloud computing using optimized cluster based federated learning. *Scientific Reports*, 15, Article 41328.
- [13]. Gangopadhyay, B., Wang, Z., Chiappa, A. S., & Takamatsu, S. (2025). Adaptive Budget Optimization for Multichannel Advertising Using Combinatorial Bandits. *Extended abstract, AAMAS 2025 Proceedings*.
- [14]. Thomaidou, S., Liakopoulos, K., & Vazirgiannis, M. (2014). Toward an integrated framework for automated development and optimization of online advertising campaigns. *Intelligent Data Analysis*, 18(6), 1199-1227.
- [15]. Bodra, D., et al. (2025). Machine learning-based cloud resource allocation algorithms: a comprehensive comparative review. *Frontiers in Computer Science*.
- [16]. Zheng, G., Gong, B., Guo, C., Peng, T., & Gong, M. (2025). AWE-DPFL: Adaptive weighting and dynamic privacy budget federated learning for heterogeneous data in IoT. *Computers & Electrical Engineering*, 123(Part A), 110070. <https://doi.org/10.1016/j.compeleceng.2025.110070>.
- [17]. Su, C., Wei, J., Lei, Y., Xuan, H., & Li, J. (2024). Empowering precise advertising with Fed-GANCC: A novel federated learning approach leveraging Generative Adversarial Networks and group clustering. *Plos one*, 19(4), e0298261.
- [18]. Qian, W., Shen, Q., Chen, X., Li, C., Fang, Y., & Wu, Z. (2024). FDP-FL: differentially private federated learning with flexible privacy budget allocation. *The Computer Journal*, 67(12), 3180-3195.
- [19]. Liu, Z., Yuan, C., Pang, M., Fang, Z., Yuan, L., Jiang, X., ... & Shao, J. (2025, July). Multi-objective Aligned Bidword Generation Model for E-commerce Search Advertising. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 4299-4303).
- [20]. Chen, H., Chen, Y. J., Park, S. H., & Shin, D. (2025). Multichannel advertising: Budget allocation in the presence of spillover and carryover effects. *Manufacturing & Service Operations Management*, 27(3), 862-880.
- [21]. Wei, P., Dou, H., Liu, S., Tang, R., Liu, L., Wang, L., & Zheng, B. (2023, July). Fedads: A benchmark for privacy-preserving cvr estimation with vertical federated learning. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 3037-3046).
- [22]. Su, C., Wei, J., Lei, Y., & Li, J. (2023). A federated learning framework based on transfer learning and knowledge distillation for targeted advertising. *PeerJ Computer Science*, 9, e1496.

- [23]. Xu, W., Liu, S., Chen, Z., Zhou, Y., Li, L., Wang, Y., ... & Wan, J. (2024). ROI constrained optimal online allocation in sponsored search. *Scientific Reports*, 14(1), 25950.
- [24]. Zhang, K., & Li, P. (2024). Federated learning optimizing multi-scenario ad targeting and investment returns in digital advertising. *Journal of Advanced Computing Systems*, 4(8), 36-43.
- [25]. Miralles-Pechuán, L., Ponce, H., & Martínez-Villaseñor, L. (2018). A novel methodology for optimizing display advertising campaigns using genetic algorithms. *Electronic Commerce Research and Applications*, 27, 39-51.
- [26]. Romero Leguina, J., Cuevas Rumín, Á., & Cuevas Rumín, R. (2020). Digital marketing attribution: Understanding the user path. *Electronics*, 9(11), 1822.
- [27]. Bergemann, D., Bonatti, A., & Wu, N. (2025). Bidding with Budgets: Algorithmic and Data-Driven Bids in Digital Advertising.
- [28]. Stram, R., Abboud, R., Shtoff, A., Somekh, O., Raviv, A., & Koren, Y. (2024, May). Mystique: A Budget Pacing System for Performance Optimization in Online Advertising. In *Companion Proceedings of the ACM Web Conference 2024* (pp. 433-442).
- [29]. Kanchan, S., Jang, J. W., Yoon, J. Y., & Choi, B. J. (2024). GSFedSec: Group Signature-Based Secure Aggregation for Privacy Preservation in Federated Learning. *Applied Sciences*, 14(17), 7993.
- [30]. Zhou, S., Wang, L., Chen, L., Wang, Y., & Yuan, K. (2025). Group verifiable secure aggregate federated learning based on secret sharing. *Scientific Reports*, 15(1), 9712.
- [31]. Seyghaly, R., Garcia, J., & Masip-Bruin, X. (2024). A comprehensive architecture for federated learning-based smart advertising. *Sensors*, 24(12), 3765.
- [32]. Guo, S., Yang, J., Long, S., Wang, X., & Liu, G. (2024). Federated learning with differential privacy via fast Fourier transform for tighter-efficient combining. *Scientific Reports*, 14(1), 26770.
- [33]. Cheng, Q. (2025). Study on Constructing the Optimization of Omni-Channel Marketing Resource Allocation Based on Dynamic Planning Algorithm in the Era of Digital Transformation. *J. COMBIN. MATH. COMBIN. COMPUT.*, 127, 5557-5575.
- [34]. Li, M. (2024). Optimal allocation of enterprise marketing resources based on hybrid parallel genetic algorithm and simulated annealing algorithm. *International Journal of Low-Carbon Technologies*, 19, 2266-2278.
- [35]. Souissi, B., Elleuch, M. A., & Ghorbel, A. (2024). A novel combined fuzzy VIKOR and mathematical optimization programming for online advertising effectiveness. *OPSEARCH*, 1-33.
- [36]. Kuppusamy, V., Shanmugasundaram, M., Dhandapani, P. B., Martin-Barreiro, C., Cabezas, X., Leiva, V., & Castro, C. (2024). Addressing a decision problem through a bipolar Pythagorean fuzzy approach: A novel methodology applied to digital marketing. *Heliyon*, 10(3).
- [37]. Sedlářová Nehézová, T., Kvasnička, R., Brožová, H., Hlavatý, R., & Kvasničková Stanislavská, L. (2025). A robust optimization approach to budget optimization in online marketing campaigns. *Central European Journal of Operations Research*, 1-32.
- [38]. Shen, Y., Xi, W., Cai, Y., Fan, Y., Yang, H., & Zhao, J. (2025). Multi-objective federated learning: Balancing global performance and individual fairness. *Future Generation Computer Systems*, 162, 107468.
- [39]. Li, M., Zhang, J., Alizadehsani, R., & Pławiak, P. (2024). A multi-channel advertising budget allocation using reinforcement learning and an improved differential evolution algorithm. *IEEE Access*.