Balancing User Experience and Monetization: An iQoE-Aware Recommendation Framework for Audio Platforms

Jie Zhao, Ziyang Cai*, Daiyang Wu
School of Management, Guangdong University of Technology
Email: zhaojie@gdut.edu.cn, 2112308019@mail2.gdut.edu.cn, 1821203795@qq.com

***Abstract*—Balancing user experience and monetization is a fundamental challenge in recommendation systems, particularly for experience-oriented content platforms with highly dynamic user interactions. This paper addresses this challenge by explicitly modeling Instantaneous Quality of Experience (iQoE) and incorporating it into a multi-objective recommendation framework. We propose an iQoE-aware multi-expert neural architecture that jointly learns payment-related behavior and iQoE through task-specific gated experts, together with a dynamic loss adjustment strategy to mitigate objective conflicts during training. In addition, a multi-objective re-ranking strategy is designed to balance experience and monetization at the list level for Top-K recommendation. Experiments on real-world datasets collected from a large-scale interactive audio platform demonstrate consistent improvements over baselines in both prediction performance and recommendation quality, validating the effectiveness of jointly optimizing user experience and monetization.**

***Index Terms*—Recommender Systems, Multi-task Learning, Quality of Experience, Re-ranking**

I. Introduction

Modern recommendation systems are increasingly expected to optimize multiple objectives simultaneously, such as user engagement, long-term satisfaction, and monetization [1]–[3]. In many real-world applications, however, these objectives are inherently conflicting [4]–[6]. Improving monetization-related outcomes, such as purchase or subscription conversion, may negatively affect user experience, while experience-oriented strategies may fail to translate into sustainable revenue [7]. Balancing user experience and monetization therefore remains a fundamental challenge in multi-objective recommendation systems.

This challenge is particularly pronounced in experience-oriented content services, where user preferences are highly subjective and dynamically evolve during content consumption [8]. Unlike static consumption scenarios, users continuously interact with content through real-time feedback mechanisms, generating rich but noisy interaction signals. Traditional recommendation models [9]–[11], typically rely on coarse-grained behavioral indicators or single-task optimization, which are insufficient to capture users' instantaneous experience states and often ignore the impact of experience quality on downstream monetization behavior. As a result, existing methods struggle to effectively coordinate experience-related objectives with revenue-driven goals.

To address this limitation, we focus on iQoE, a fine-grained representation of user experience that reflects users' real-time responses during interaction [12]–[14]. iQoE captures the dynamic and temporal nature of user experience and provides a complementary perspective to conventional behavioral features. However, explicitly modeling iQoE and integrating it into recommendation systems introduces new challenges, including task interference in multi-objective learning and the need for effective coordination between competing optimization goals.

In this paper, we propose an iQoE-aware multi-objective recommendation framework based on a multi-expert neural architecture. The proposed model jointly learns user monetization behavior (denoted as the Pay task) and iQoE prediction task by disentangling task-shared and task-specific representations through gated expert networks. To further mitigate conflicts between objectives, we introduce a dynamic loss balancing strategy based on task uncertainty, enabling adaptive optimization during training. In addition, we design a multi-objective re-ranking strategy that incorporates both Pay preference and iQoE preference for Top-K recommendation.

The main contributions of this work can be summarized as follows:

1) We introduce iQoE as an explicit learning objective for multi-objective recommendation and formulate it as a dynamic, fine-grained representation of user experience.
2) We propose a multi-expert neural framework with dynamic loss balancing to jointly optimize monetization and experience-related objectives.
3) We design a multi-objective re-ranking strategy to effectively integrate predicted preferences into recommendation lists.
4) Extensive experiments on real-world datasets demonstrate that the proposed method achieves improved recommendation quality while better balancing user experience and monetization.

*Corresponding author.
This work was supported by the National Natural Science Foundation of China under Grant Nos. 72271063 and 71871069, and by the Guangdong Province Philosophy and Social Science Planning Project under Grant No. GD24CGL45.

The source code is available at the following repository: https://github.com/lemonade019/iQoERec.

II. Related Work

A. QoE and Interactive Experience Modeling

Quality of Experience (QoE) has been widely studied as a subjective measure of users' perceived quality in multimedia services. Early work [15]–[17] primarily modeled the relationship between objective system-level indicators (e.g., latency, bitrate, and packet loss) and overall user satisfaction. With the development of data-driven methods, learning-based QoE models have incorporated user behavior, content characteristics, and contextual factors [18].

More recent studies highlight the dynamic nature of user experience during interaction and explore fine-grained signals such as real-time feedback and interaction logs to capture instantaneous perception [12]–[14]. However, most existing researches treat QoE as an auxiliary analysis metric or aggregate experience signals over long time horizons, limiting their effectiveness in recommendation optimization. In contrast, our work focuses on Interactive QoE and integrates interaction-level experience modeling directly into the learning and optimization process of recommendation systems.

B. Multi-Task and Multi-Objective Recommendation

Multi-task learning is widely adopted in recommendation systems to jointly optimize correlated objectives such as click-through rate and conversion rate [3]. Early shared-bottom architectures exploit task correlations but often suffer from negative transfer under conflicting objectives [19]–[21]. To alleviate this issue, advanced architectures such as mixture-of-experts and gated multi-expert models have been proposed to balance shared and task-specific representations [22], [23]. In parallel, multi-objective recommendation studies explore optimization strategies including weighted loss aggregation and constraint-based methods to manage trade-offs among competing goals [24]–[27].

However, most existing researches [23], [27]–[29] primarily focus on behavioral or business-oriented objectives, while experience-related factors are rarely modeled explicitly. Moreover, fixed weighting strategies lack adaptability to task difficulty and uncertainty, making it challenging to effectively coordinate experience-centric objectives with monetization.

III. Method

In this section, we present the proposed iQoE-Aware Multi-Objective Recommendation framework. We first define the problem setting and learning objectives, and then introduce the model architecture, dynamic loss balancing strategy, and multi-objective re-ranking strategies.

A. Problem Formulation

We consider a recommendation scenario with a user set $\mathcal{U} = \{u_1, \dots, u_N\}$ and a candidate item (sound) set $\mathcal{S} = \{s_1, \dots, s_M\}$. For each user–item pair (u, s), we construct a heterogeneous input feature vector

$$\mathbf{x}_{u,s} = \left[\, \mathbf{x}_{u,s}^{(d)},\ \mathbf{x}_{u,s}^{(c)} \,\right] \tag{1}$$

where $\mathbf{x}_{u,s}^{(d)}$ and $\mathbf{x}_{u,s}^{(c)}$ denote the discrete and continuous parts, respectively.

1) **Task 1**: predict whether user u performs a payment-related action on s, with label $y_{u,s}^{\text{Pay}} \in \{0, 1\}$.
2) **Task 2**: predict the iQoE during consumption, with target $y_{u,s}^{\text{iQoE}} \in \mathbb{R}^{+}$.

Given an interaction dataset $\mathcal{D} = \{(\mathbf{x}_{u,s}, y_{u,s}^{\text{Pay}}, y_{u,s}^{\text{iQoE}})\}$, our goal is to learn a multi-task model parameterized by Θ that outputs $\hat{y}_{u,s}^{\text{Pay}}$ and $\hat{y}_{u,s}^{\text{iQoE}}$ via deep feature extraction on $\mathbf{x}_{u,s}$. The overall optimization objective is defined as

$$\min_{\Theta}\ \mathcal{L}(\Theta) = \lambda_{\text{Pay}} \mathcal{L}_{\text{Pay}}(\Theta) + \lambda_{\text{iQoE}} \mathcal{L}_{\text{iQoE}}(\Theta) + \alpha \mathcal{R}(\Theta), \tag{2}$$

where λ_{Pay} and λ_{iQoE} control the importance of the two tasks, and $\mathcal{R}(\Theta)$ is a regularization term.

B. Model Architecture

The proposed model is built upon a multi-expert architecture that disentangles shared and task-specific representations. The overall structure consists of four main modules: input embedding, shared representation learning, multi-expert modeling with task-specific gating, and task-specific prediction layers.

1) Input Embedding: Considering the diversity of input data, we adopt different embedding strategies for discrete and continuous features.

For discrete features, we introduce an embedding matrix $\mathbf{E}_j \in \mathbb{R}^{n_j \times d_{\text{emb}}}$, where n_j denotes the cardinality of the j-th discrete feature and d_{emb} is the embedding dimension. The corresponding embedding vector is obtained by lookup:

$$\mathbf{e}_j = \mathbf{E}_j[x_j^{(d)}]. \tag{3}$$

All discrete feature embeddings are then concatenated as

$$\mathbf{e}^{(d)} = \text{concat}(\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_{M_d}), \tag{4}$$

where M_d denotes the number of discrete features.

For continuous features $\mathbf{x}^{(c)} \in \mathbb{R}^{d_c}$, we first apply normalization and then project them into a d_{proj}-dimensional latent space using a fully connected layer:

$$\mathbf{e}^{(c)} = \sigma(\mathbf{W}^{(c)} \mathbf{x}^{(c)} + \mathbf{b}^{(c)}), \tag{5}$$

where $\sigma(\cdot)$ denotes a nonlinear activation function, implemented as the Rectified Linear Unit (ReLU). $\mathbf{W}^{(c)} \in \mathbb{R}^{d_{\text{proj}} \times d_c}$ and $\mathbf{b}^{(c)} \in \mathbb{R}^{d_{\text{proj}}}$ are the weight matrix and bias term, respectively.

Finally, the overall embedded representation is obtained by concatenation:

$$\mathbf{e} = \text{concat}(\mathbf{e}^{(d)}, \mathbf{e}^{(c)}). \tag{6}$$

After feature embedding, we employ a L layer fully-connected encoder to extract generic representations. The generic representation output is:

$$\mathbf{h}_0 = f_{\text{enc}}(\mathbf{e}) = \sigma\left(\mathbf{W}_L \sigma\left(\cdots \sigma(\mathbf{W}_1 \mathbf{e} + \mathbf{b}_1) \cdots\right) + \mathbf{b}_L\right). \tag{7}$$

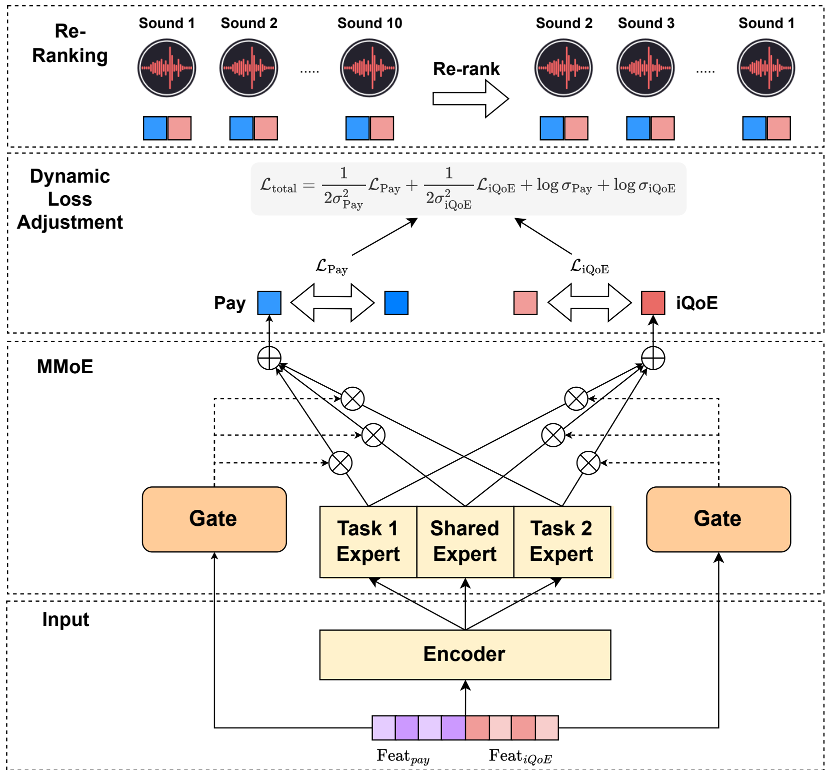

Fig. 1. Overview of the proposed iQoE-aware multi-objective recommendation framework.

This encoder captures common latent information and provides a unified input for subsequent expert networks.

2) Multi-gate Mixture-of-Experts Module: Following the MMOE design, we construct K expert networks (in our setting, $K = 2$). Each expert takes $\mathbf{h}_0$ as input and extracts complementary latent information:

$$\mathbf{E}_k = f_k(\mathbf{h}_0) = \sigma(\mathbf{W}_k\mathbf{h}_0 + \mathbf{b}_k), \quad k = 1, 2, \ldots, K. \tag{8}$$

To enable task-specific expert selection, we introduce a gating network for each task $i \in \{\text{Pay}, \text{iQoE}\}$. The gate value on expert k is:

$$g_{ik} = \frac{\exp\left((\mathbf{w}_{gk}^{(i)})^\top \mathbf{h}_0 + b_{gk}^{(i)}\right)}{\sum_{j=1}^{K}\exp\left((\mathbf{w}_{gj}^{(i)})^\top \mathbf{h}_0 + b_{gj}^{(i)}\right)}, \quad i \in \{\text{Pay}, \text{iQoE}\}. \tag{9}$$

Then the task-specific mixture representation is computed by weighted aggregation:

$$\mathbf{H}_i = \sum_{k=1}^{K} g_{ik} \cdot \mathbf{E}_k. \tag{10}$$

Finally, each task uses an independent task-specific tower $f_i(\cdot)$ to produce prediction:

$$\hat{y}_i = f_i(\mathbf{H}_i) = \text{Sigmoid}(\text{MLP}_i(\mathbf{H}_i)). \tag{11}$$

C. Dynamic Loss Adjustment Strategy

To handle differences in loss scale and learning difficulty between objectives, we adopt a task-uncertainty-based dynamic loss adjustment strategy. Let the Pay task loss be $\mathcal{L}_{\text{Pay}}$ (Binary Cross-entropy) and the iQoE task loss be $\mathcal{L}_{\text{iQoE}}$ (Mean Squared Error). The total loss is:

$$\mathcal{L}_{\text{total}} = \frac{1}{2\tau^2_{\text{Pay}}}\mathcal{L}_{\text{Pay}} + \frac{1}{2\tau^2_{\text{iQoE}}}\mathcal{L}_{\text{iQoE}} + \log\tau_{\text{Pay}} + \log\tau_{\text{iQoE}}, \tag{12}$$

where τ_{Pay} and τ_{iQoE} denote the learnable task uncertainty parameters.

D. Multi-Objective Recommendation Re-ranking

In the re-ranking stage, let $\hat{y}_{\text{Pay}}(u, s)$ and $\hat{y}_{\text{iQoE}}(u, s)$ denote the predicted Pay score and iQoE score of item s for user u, respectively, and let $R_K(u)$ denote the corresponding Top-K recommendation list.

We design three re-ranking strategies:

1. Preference Ranking Strategy (PRS)

$$R_K^{\text{PRS}}(u) = \text{TopK}_{s\in\mathcal{S}}^{\downarrow(\hat{y}_{\text{Pay}}(u,s),\, \hat{y}_{\text{iQoE}}(u,s))}, \tag{13}$$

where $R_K^{\text{PRS}}(u)$ denotes the Top-K recommendation list using PRS. Items are ranked in descending lexicographic order, first by $\hat{y}_{\text{Pay}}(u, s)$ and then by $\hat{y}_{\text{iQoE}}(u, s)$.

2. Weighted Ranking Strategy (WRS)

$$R_K^{\text{WRS}}(u) = \text{TopK}_{s\in S}^{\downarrow f(u,s)}, \tag{14}$$

where $f(u,s)$ denotes a linear combination of $\hat{y}_{\mathrm{Pay}}(u,s)$ and $\hat{y}_{\mathrm{iQoE}}(u,s)$. In our implementation, we use their average.

3. Segment Ranking Strategy (SRS)

$$R_K^{\mathrm{SRS}}(u) = A_{K/2}^{\mathrm{Pay}}(u) \oplus \mathrm{TopK}_{s \in \mathcal{S} \setminus A_{K/2}^{\mathrm{Pay}}(u)}^{\downarrow(\hat{y}_{\mathrm{Pay}}(u,s),\, \hat{y}_{\mathrm{iQoE}}(u,s))}, \quad (15)$$

where $A_{K/2}^{\mathrm{Pay}}(u) = \mathrm{TopK}_{s \in S}^{\downarrow \hat{y}_{\mathrm{Pay}}(u,s)}$ denotes the top-$K/2$ items ranked by Pay score, and $\oplus$ denotes list concatenation.

IV. Performance Evaluation

In this section, we evaluate the effectiveness of the proposed iQoE-aware multi-objective recommendation framework. The evaluation focuses on the following research questions:

- **RQ1**: Can the proposed model outperform existing baselines in the Pay and iQoE prediction tasks?
- **RQ2**: Does explicitly modeling iQoE contribute to improved recommendation performance?
- **RQ3**: How do different multi-objective re-ranking strategies affect Top-K recommendation quality?

A. Experimental Setup

1) Datasets: Experiments are conducted on seven real-world datasets collected by the authors from a large-scale interactive audio content platform and constructed using sliding time windows. Each dataset contains user–item interaction records with monetization-related signals and iQoE annotations derived from interaction behaviors. For each dataset, samples are randomly split into training, validation, and test sets with a ratio of 6:2:2. All experiments are repeated ten times with different random seeds, and average results are reported.

2) Evaluation Metrics: For prediction tasks, Pay and iQoE are evaluated using AUC and Accuracy (ACC), where AUC reflects ranking discrimination under class imbalance and ACC measures overall correctness. For Top-K recommendation, we use Hit Ratio (HR@K) and Normalized Discounted Cumulative Gain (NDCG@K), where HR indicates whether relevant items appear in the recommendation list and NDCG further accounts for ranking positions.

3) Baselines: We compare the proposed method with representative baselines, including:

- **PLE** [26]: Progressive layered extraction with shared and task-specific experts combined via gating networks to mitigate negative transfer.
- **ShareB** [19]: A hard parameter sharing multi-task model with a shared encoder and task-specific layers, which may suffer from negative transfer.
- **ESMM** [27]: A joint modeling framework for CTR and CVR that addresses sample selection bias and data sparsity.
- **DeepFM** [9]: Combines factorization machines and deep neural networks to capture low- and high-order feature interactions.
- **NFM** [10]: Extends FM with bi-interaction pooling and a DNN to model higher-order feature interactions.
- **ONN** [11]: Enhances FFM with a multilayer perceptron to capture complex feature interactions.

All methods are implemented under the same experimental settings. Hyperparameters are tuned using validation sets.

B. Multi-Objective Prediction Performance

To answer RQ1 from a multi-objective perspective, we first evaluate the proposed model under full-feature settings and compare it with representative multi-objective baselines.

As shown in Table I, the proposed model achieves consistently competitive or best performance on both the Pay and iQoE tasks across all datasets in terms of AUC and ACC.

For the Pay task, our model obtains the highest average AUC among all compared multi-objective methods, outperforming ShareBottom and ESMM by a clear margin and slightly surpassing PLE. This result indicates that the multi-expert architecture effectively captures monetization-related patterns while alleviating negative transfer between objectives. For the iQoE task, the proposed model also achieves the highest average AUC, although ShareBottom exhibits marginal advantages on a few individual datasets.

Overall, the results in Table I suggest that the proposed method can jointly optimize monetization and iQoE in a balanced manner.

C. Single-Objective Performance Comparison

To further answer RQ1 and examine whether the proposed architecture remains effective when optimized independently, we evaluate the model under single-objective settings. The results are reported in Table II.

For the Pay task, the proposed model achieves the best average performance and remains competitive across datasets, demonstrating strong capability in capturing payment-related features. For the iQoE task, performance differences among models are relatively smaller. Our model achieves comparable performance, while DeepFM and NFM obtain slightly higher ACC on average and ONN performs worse overall.

These results suggest that the shared-expert architecture retains competitive task-specific modeling ability, particularly on the Pay task.

D. Ablation Study

To answer RQ2 and investigate the contribution of explicitly modeling iQoE, we conduct feature-level ablation experiments based on different feature groups. Specifically, $\mathrm{Feat}_{\mathrm{All}}$ denotes the complete feature set used in our model. $\mathrm{Feat}_{\mathrm{iQoE}}$ and $\mathrm{Feat}_{\mathrm{Pay}}$ represent task-relevant feature subsets selected for the iQoE task and the Pay task, respectively, using a feature selection algorithm [30]. We further define $\mathrm{Feat}_{\mathrm{Cross}}$ as the intersection of $\mathrm{Feat}_{\mathrm{iQoE}}$ and $\mathrm{Feat}_{\mathrm{Pay}}$. The remaining non-overlapping features are denoted as $\mathrm{Feat}_{\text{E-C}} = \mathrm{Feat}_{\mathrm{iQoE}} \setminus \mathrm{Feat}_{\mathrm{Cross}}$ and $\mathrm{Feat}_{\text{P-C}} = \mathrm{Feat}_{\mathrm{Pay}} \setminus \mathrm{Feat}_{\mathrm{Cross}}$, respectively.

Based on these feature groups, we evaluate the impact of different feature combinations on model performance. As shown in Fig. 2, using only iQoE-related features yields

TABLE I
PERFORMANCE COMPARISON OF MULTI-OBJECTIVE MODELS ON AUC AND ACC.

Dataset	AUC								ACC							
	Task 1				Task 2				Task 1				Task 2			
	Ours	PLE	ESMM	ShareB	Ours	PLE	ESMM	ShareB	Ours	PLE	ESMM	ShareB	Ours	PLE	ESMM	ShareB
1	**91.79%**	91.15%	87.71%	90.95%	87.80%	**88.21%**	84.29%	81.89%	**83.69%**	83.19%	78.21%	83.37%	79.95%	**80.54%**	74.90%	74.54%
2	88.77%	**89.67%**	79.93%	68.99%	87.00%	**87.74%**	81.30%	84.25%	79.92%	**80.64%**	68.54%	61.91%	**80.18%**	80.12%	68.97%	78.10%
3	**82.02%**	81.38%	70.79%	71.71%	**87.49%**	86.29%	84.66%	79.78%	90.50%	89.65%	63.45%	**91.43%**	**80.64%**	79.75%	68.65%	73.89%
4	68.75%	**68.85%**	49.78%	58.26%	86.64%	**87.59%**	83.09%	86.40%	**83.06%**	82.88%	69.76%	82.98%	79.90%	**80.72%**	73.93%	79.92%
5	65.58%	**67.33%**	50.84%	63.55%	**88.82%**	88.43%	85.94%	88.45%	**80.92%**	80.92%	69.70%	80.92%	82.24%	82.61%	75.17%	**82.63%**
6	**69.94%**	68.40%	66.70%	66.18%	**89.73%**	88.79%	88.69%	89.04%	**83.22%**	83.22%	58.41%	83.22%	**83.37%**	82.77%	76.12%	83.11%
7	**83.11%**	81.41%	67.11%	84.08%	**89.27%**	87.32%	84.72%	87.74%	86.99%	85.74%	72.57%	**87.37%**	**81.39%**	81.19%	75.17%	80.98%
Avg.	**78.57%**	78.31%	67.55%	71.96%	**88.11%**	87.77%	84.67%	85.36%	**84.04%**	83.75%	68.66%	81.60%	81.10%	**81.10%**	73.27%	79.02%

TABLE II
PERFORMANCE COMPARISON OF SINGLE-OBJECTIVE MODELS ON AUC AND ACC.

Dataset	AUC								ACC							
	Task 1				Task 2				Task 1				Task 2			
	Ours	DeepFM	NFM	ONN	Ours	DeepFM	NFM	ONN	Ours	DeepFM	NFM	ONN	Ours	DeepFM	NFM	ONN
1	**91.79%**	91.10%	90.54%	90.33%	87.80%	**88.96%**	88.73%	84.62%	**83.69%**	81.11%	80.52%	77.50%	79.95%	**80.65%**	80.29%	77.96%
2	**88.77%**	88.55%	88.06%	88.18%	87.00%	88.64%	**88.81%**	85.91%	79.92%	81.01%	**81.56%**	77.64%	**80.18%**	79.23%	79.75%	75.97%
3	**82.02%**	76.78%	74.38%	80.89%	**87.49%**	86.23%	86.48%	84.09%	**90.50%**	79.75%	80.06%	76.95%	80.64%	**80.73%**	80.11%	72.36%
4	**68.75%**	67.53%	68.54%	63.92%	86.64%	**88.76%**	88.73%	83.09%	**83.06%**	81.93%	82.18%	75.88%	**79.90%**	73.41%	73.72%	66.04%
5	65.58%	66.34%	**66.58%**	61.64%	88.82%	89.62%	**89.72%**	83.98%	80.92%	82.65%	**83.46%**	77.38%	**82.24%**	70.68%	69.66%	66.29%
6	**69.94%**	67.64%	68.15%	64.01%	**89.73%**	89.44%	89.55%	84.93%	**83.22%**	82.59%	82.66%	78.80%	**83.37%**	82.62%	82.66%	72.08%
7	**83.11%**	80.67%	80.61%	80.33%	89.27%	89.19%	**89.29%**	85.68%	**86.99%**	81.98%	81.64%	78.06%	**81.39%**	80.21%	79.92%	75.55%
Avg.	**78.57%**	76.94%	76.69%	75.61%	88.11%	88.69%	**88.76%**	84.61%	**84.04%**	81.57%	81.73%	77.46%	**81.10%**	78.22%	78.02%	72.32%

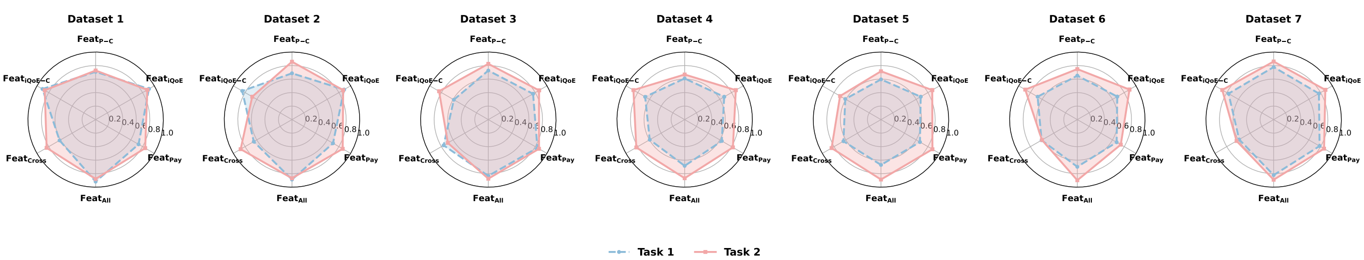

Fig. 2. Feature ablation results.

better Pay prediction performance than using only payment-related features, highlighting the strong correlation between interactive experience and monetization behavior. Moreover, removing the intersection features between Pay and iQoE leads to noticeable performance degradation on both tasks.

These findings confirm that iQoE is not merely an auxiliary signal but a critical modeling component that facilitates effective knowledge sharing and joint optimization.

E. Top-K Recommendation Performance and Re-Ranking Strategies

To answer RQ3, we further evaluate three multi-objective re-ranking strategies, namely PRS, WRS, and SRS, under the Top-K recommendation setting.

TABLE III
PERFORMANCE COMPARISON OF MODELS ON TOP-K METRICS.

Method	HR@5	NDCG@5	HR@10	NDCG@10	HR@20	NDCG@20
SharedB	56.47%	49.85%	63.03%	51.79%	71.67%	53.85%
PLE	63.23%	54.75%	68.63%	56.34%	75.28%	57.87%
ESMM	45.65%	40.66%	49.74%	41.86%	54.04%	42.75%
Ours (SRS)	**70.98%**	**65.04%**	**74.56%**	**66.00%**	**78.49%**	**66.81%**

The experimental results in Fig. 3 and Table III show that the SRS consistently outperforms PRS and WRS in terms of both HR and NDCG across different values of K on most datasets, where Table III reports the average results over seven datasets.

In addition, the results indicate that the proposed model combined with SRS achieves the best overall recommendation quality among the compared re-ranking strategies in most cases.

These results suggest that monetization preference dominates user decisions at the top of the recommendation list, while iQoE plays a more important role in later positions, which directly supports the design motivation of SRS.

V. CONCLUSION

This paper addresses the challenge of balancing user experience and monetization in recommendation systems by explicitly modeling iQoE. We proposed an iQoE-aware multi-objective recommendation framework that jointly optimizes payment-related and experience-related objectives through a

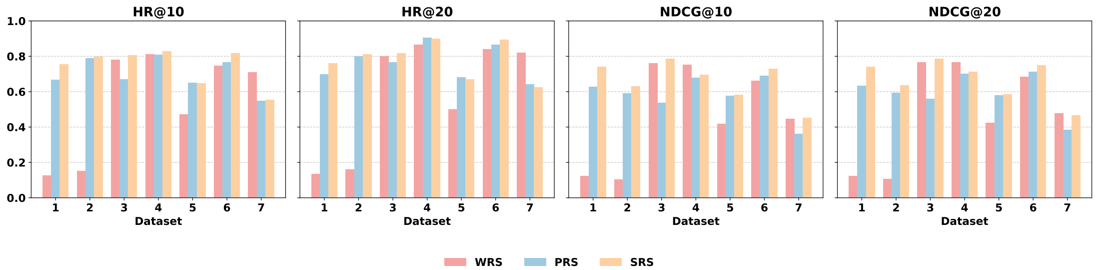

Fig. 3. Recommendation performance under different re-ranking strategies.

multi-expert neural architecture, dynamic loss adjustment, and multi-objective re-ranking. Experimental results on real-world datasets demonstrate that the proposed method achieves competitive or superior performance on both Pay and iQoE prediction tasks, while effectively improving Top-K recommendation quality. Ablation studies further verify that explicitly modeling iQoE provides complementary information that benefits both objectives. These findings highlight the importance of integrating interaction-level experience modeling into multi-objective recommendation systems and suggest a promising direction for experience-aware optimization in practical recommender applications.

References

[1] D. Jannach and H. Abdollahpouri, "A survey on multi-objective recommender systems," *Frontiers in Big Data*, 2023.

[2] F. E. Zaizi, S. Qassimi, and S. Rakrak, "Multi-objective optimization with recommender systems: A systematic review," *Information systems*, 2023.

[3] A. De Biasio, A. Montagna, F. Aiolli, and N. Navarin, "A systematic review of value-aware recommender systems," *Expert Systems with Applications*, 2023.

[4] P. Dokoupil, L. Peska, and L. Boratto, "SM-RS: single- and multi-objective recommendations with contextual impressions and beyond-accuracy propensity scores," in *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2024.

[5] Y. Wang, L. Tao, and X. X. Zhang, "Recommending for a multi-sided marketplace: A multi-objective hierarchical approach," *Marketing Science*, 2025.

[6] S. Agah, S. Schaeffer, M. Peifer, N. Sharma, A. Maheshwari, and S. Hamidian, "Pareto-optimal solution: Optimizing engagement and revenue," in *Proceedings of the Nineteenth ACM Conference on Recommender Systems*, 2025.

[7] Y. Wang, M. Sharma, C. Xu, S. Badam, Q. Sun, L. Richardson, L. Chung, E. H. Chi, and M. Chen, "Surrogate for long-term user experience in recommender systems," in *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2022.

[8] X. Yi, L. Hong, E. Zhong, N. N. Liu, and S. Rajan, "Beyond clicks: dwell time for personalization," in *Eighth ACM Conference on Recommender Systems*, 2014.

[9] H. Guo, R. Tang, Y. Ye, Z. Li, and X. He, "Deepfm: A factorization-machine based neural network for CTR prediction," in *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*, 2017.

[10] X. He and T. Chua, "Neural factorization machines for sparse predictive analytics," in *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2017.

[11] Y. Yang, B. Xu, S. Shen, F. Shen, and J. Zhao, "Operation-aware neural networks for user response prediction," *Neural Networks*, 2020.

[12] N. Eswara, S. Ashique, A. Panchbhai, S. Chakraborty, and S. S. Channappayya, "Streaming video qoe modeling and prediction: A long short-term memory approach," *IEEE Transactions on Circuits and Systems for Video Technology*, 2019.

[13] U. Bulkan and T. Dagiuklas, "Predicting quality of experience for online video service provisioning," *Multimedia Tools and Applications*, 2019.

[14] H. Z. Jahromi, D. T. Delaney, and A. Hines, "Beyond first impressions: Estimating quality of experience for interactive web applications," *IEEE Access*, 2020.

[15] H. J. Kim, D. H. Lee, J. M. Lee, K. H. Lee, and S. G. Choi, "The qoe evaluation method through the qos-qoe correlation model," *IEEE Computer Society*, 2008.

[16] P. Brooks and B. Hestnes, "User measures of quality of experience: why being objective and quantitative is important," *IEEE Netw.*, 2010.

[17] M. Alreshoodi and J. Woods, "Survey on qoe qos correlation models for multimedia services," *International Journal of Distributed and Parallel systems*, 2013.

[18] T. Zhao, Q. Liu, and C. W. Chen, "Qoe in video transmission: A user experience-driven strategy," *IEEE Communications Surveys & Tutorials*, 2017.

[19] R. Caruana, "Multitask learning," *Machine Learning*, 1997.

[20] Y. Zhang and Q. Yang, "A survey on multi-task learning," *IEEE Transactions on Automatic Control*, 2022.

[21] T. Standley, A. Zamir, D. Chen, L. Guibas, J. Malik, and S. Savarese, "Which tasks should be learned together in multi-task learning?" in *37th International Conference on Machine Learning*, 2021.

[22] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. V. Le, G. E. Hinton, and J. Dean, "Outrageously large neural networks: The sparsely-gated mixture-of-experts layer," in *5th International Conference on Learning Representations*, 2017.

[23] J. Ma, Z. Zhao, X. Yi, J. Chen, L. Hong, and E. H. Chi, "Modeling task relationships in multi-task learning with multi-gate mixture-of-experts," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018.

[24] Z. Zhang, S. Liu, J. Yu, Q. Cai, X. Zhao, C. Zhang, Z. Liu, Q. Liu, H. Zhao, and L. Hu, "M3oe: Multi-domain multi-task mixture-of experts recommendation framework," *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2024.

[25] Q. Zhang, X. Liao, Q. Liu, J. Xu, and B. Zheng, "Leaving no one behind: A multi-scenario multi-task meta learning approach for advertiser modeling," in *The Fifteenth ACM International Conference on Web Search and Data Mining*, 2022, pp. 1368–1376.

[26] H. Tang, J. Liu, M. Zhao, and X. Gong, "Progressive layered extraction (PLE): A novel multi-task learning (MTL) model for personalized recommendations," in *Fourteenth ACM Conference on Recommender Systems*, 2020.

[27] X. Ma, L. Zhao, G. Huang, Z. Wang, Z. Hu, X. Zhu, and K. Gai, "Entire space multi-task model: An effective approach for estimating post-click conversion rate," in *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, 2018.

[28] Y. Ni, D. Ou, S. Liu, X. Li, W. Ou, A. Zeng, and L. Si, "Perceive your users in depth: Learning universal user representations from multiple e-commerce tasks," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018.

[29] W. Ouyang, X. Zhang, S. Ren, C. Qi, Z. Liu, and Y. Du, "Representation learning-assisted click-through rate prediction," in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, 2019.

[30] J. Zhao, D. Wu, J. Wu, W. Ye, F. Huang, J. Wang, and E. W. K. See-To, "Consistency approximation: Incremental feature selection based on fuzzy rough set theory," *Pattern Recognition*, 2024.