CULC-Net: A Recipe for Tailored Creative Selection in Online Advertising

Baosheng Zhang[1,2], Liufang Sang[3], Haoran Wang[3], Wei Wang[3], Wenlong Chen[3], Changping Peng[3], Zhangang Lin[3], Jingping Shao[3], Jie He[3], Haoqian Wang[2(✉)], and Yuchen Guo[1(✉)]

[1] Beijing National Research Center for Information Science and Technology (BNRist), Tsinghua University, Beijing, China
yuchen.w.guo@gmail.com
[2] Tsinghua Shenzhen International Graduate School, Tsinghua University, Shenzhen, China
wangyizhai@sz.tsinghua.edu.cn
[3] JD.com, Beijing, China

Abstract. Online advertising is a major application of recommendation systems. The primary process first involves recommending appropriate items to users, followed by selecting suitable creatives, such as ad posters. While much research has focused on optimizing item recommendations to increase user clicks, creative selection has often been overlooked. Properly chosen creatives can significantly enhance purchasing intent by aligning with the diverse preferences, ages, and genders of users. Current state-of-the-art methods typically rely on historical Click-Through Rates (CTR), which may exhibit biases during initial exposures due to limited data. In this paper, we introduce CULC-Net, which builds detailed profiles to uncover hidden connections between users and creatives, utilizing a creative relevance score for soft-decision making. This approach improves recommendation effectiveness and reduces reliance on sparse CTR data. Furthermore, we advance beyond the traditional CTR-based “only top for training" strategy by introducing FlexiRank. Creatives are sorted based on the relative strength of their CTRs, effectively managing noise and outliers. We test CULC-Net in a real-world search ad system, demonstrating a 3.43% increase in online and a 4.01% increase in offline. Further validation on a public benchmark confirms the effectiveness of our approach.

Keywords: Recommendation System · Ad Recommendation· Creative Selection.

1 Introduction

With the growth of the Internet and mobile technologies, online advertising has become vital for the income of digital platforms. It is important to grab user attention with eye-catching visuals and clear, brief messages [1]. Ads that look good are more likely to be clicked on, which increases the Click-Through Rate

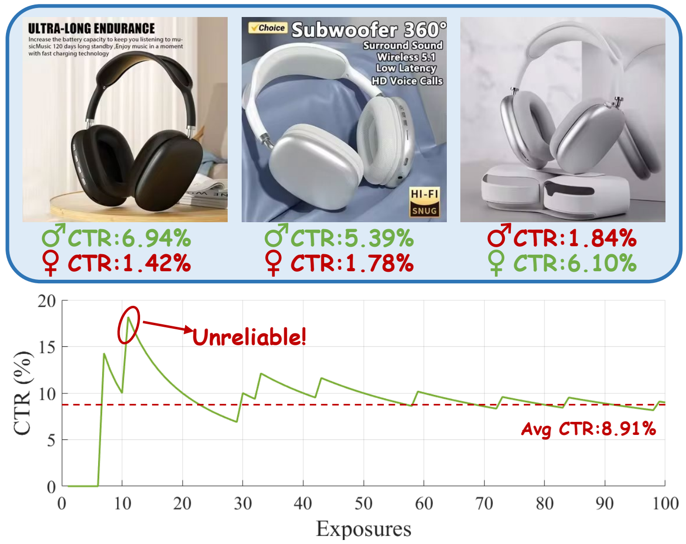

Fig. 1. CTR comparison for different ad creative examples, highlighting the impact of ad creatives. The first row shows how CTR varies based on user characteristics; for example, males may prefer descriptions of features, while females may be more attracted to visual displays for the same headphone. The second row shows how the CTR of a single creative changes with more exposure, emphasizing the unreliability of CTR with limited exposures.

(CTR) of products [26]. CTR is not only a sign of user interest; it is also a key indicator of how financially successful an online advertising campaign is. This affects the total revenue and the return on the investment made in advertising. Even a small improvement in CTR can have a big financial impact, especially for large e-commerce platforms. Therefore, improving the effectiveness of ad recommendations is crucial [4].

Current recommendation systems generally follow a multi-stage cascading structure that includes matching, ad ranking, and creative selection [13]. Initially, the matching stage reduces the extensive pool of ad candidates from billions down to thousands. The next stage, ad ranking, organizes these ads to identify the final top selections. Finally, the creative selection stage chooses the most appropriate visual and text creatives for each ad [30]. At the creative selection stage, industry practices often ignore specific creatives, either choosing them randomly or picking popular ones from past data [28]. They do not fully consider the actual impact of different visuals on user engagement. However, evidence shows that creative elements can significantly impact Click-Through Rates (CTR) [14, 16, 34]. As demonstrated in Figure 1, there are clear differences in CTR for different creatives between male and female users. For example, for the same headphone, men may prefer descriptions emphasizing utility, while women might be more drawn to simplistic visual designs. Capturing these preferences with specific rules is challenging, and effectively leveraging diverse multimodal data with neural networks remains a challenge [20, 23, 34].

Ad recommendations often involve frequently changing in creatives [7], which results in each creative being displayed an average of only 6 times. With such limited exposure, even random clicks can significantly affect the CTR [34]. For example, if a creative receives only one exposure and one click, this could lead to a misleading CTR of 100%. Such a scenario could easily be due to an accidental click rather than genuine user interest. Consequently, relying on CTR to gauge the popularity of a creative with few exposures can be misleading [11]. This underscores the need for ad recommendation systems that are better tailored to handle user interactions effectively. Current creative selection methods primarily follow two aspect [22]. The first class uses hard labels, recommending the top creative based solely on the highest CTR value [34]. In practice, when the creative with the highest CTR has only a few exposures, recommending it on actual advertising platforms is not reliable. The second class considers the ranking information of creatives [24], where the model ranks creatives by their CTR, comprehensively assess the performance of different creatives. However, relying on unstable or inaccurate ranking data can undermine the accuracy and reliability of the recommendations [6].

To address the issue of unreliable CTR data due to limited exposure, we explore the relationships between creatives and customize recommendations for users with similar features, reducing our reliance on CTR data for individual users or creatives. We designed CULC-Net (Contrastive User Learning for Creative Selection), which incorporates a soft-decision contrastive learning approach to address this challenge. By partially masking user features, CULC-Net builds profiles of users with similar attributes, allowing it to uncover hidden connections between users and creatives, and thereby mitigating the bias from sparse data. Unlike existing methods that might wrongly push all other creatives apart by focusing only on classification data, our approach introduces a creative relevance score. This score evaluates how related different creatives are to each other. During training, creatives that are less related are separated more, while those with higher similarity are kept closer together in the representation space. This method not only enhances the quality of data but also reduces the impact of unreliable CTR data from low exposure, resulting in a more robust model.

Furthermore, we proposed FlexiRank, which utilizes a soft ranking loss function in place of classification and ranking losses. It sorts creatives based on the relative ranking relationships of their CTRs, effectively managing noise and outliers in the data. FlexiRank also incorporates an instructional gradient that decays with training, progressively reducing reliance on explicit ranking information. In this way, it enhances the model's effectiveness and improves its predictive accuracy for unseen data.

We highlight our contributions in this paper as follows:

(1) We introduce CULC-Net, which changes the "focus on self" concept in CL by leveraging differences within users and creatives to obtain a creative relevance score for soft-decision making.

(2) We introduce FlexiRank, improving upon the “only top for training" method by using all CTR data for more precise ranking and better handling of noise and outliers.

(3) Our results show that CULC-Net outperforms existing models, achieving performance improvements of 4.01% offline and 3.43% online, proving its effectiveness in real-world advertising scenarios.

2 Related works

2.1 Creative Selection

In online advertising, there are two main methods used for selecting ad creatives. The first class of creative selection strategies relies on hard labels to recommend creatives based on the highest CTR [34]. These approaches straightforward suffers from reliability issues especially when creatives have limited exposures — a problem acknowledged and addressed in part by methodologies that attempt to pre-evaluate creatives [32, 33]. For instance, PEAC (Pre Evaluation of Ad Creative Model) utilizes deep learning to predict potential online performance without relying on user clicks, emphasizing the importance of offline creative quality evaluation based on comprehensive image and text content analysis [36] .Similarly, the Adaptive and Efficient ad creative Selection (AES) framework introduces an innovative ingredient tree combined with Thompson sampling for efficient selection based on predicted CTR, addressing the high variance due to limited feedback and the sparsity of user interactions, which is a common challenge in creative selection [5].

The second class involves ranking creatives based on CTR, assigning higher ranks to those with higher CTRs and lower ranks to those with lower CTRs, using these rankings to infer user preferences. However, this method struggles with accurately identifying true user preferences due to the presence of false negatives in low CTR data. Advanced hybrid models and category-specific approaches that incorporate visual and categorical data to refine the ranking process have been demonstrated to enhance creative optimization and integrate creative selection more effectively within ad ranking stages [6,24]. Furthermore, systems like HBM-VAM introduce visual priors and a flexible updated bandit method that can raise platform revenue by focusing on online assessments [31]. CACS presents a method that places the creative ranking module before the ad ranking stage, then jointly optimizes them with distillation and shared embedding, resembling our method closely and showcasing the potential for significant advancements in advertising systems [23].

2.2 Contrastive Learning

Contrastive learning is a powerful branch of self-supervised learning that focuses on encoding data by contrasting positive and negative samples [18]. This technique is particularly effective in recommendation systems, where it is used to learn representations that capture unique features of content, thereby facilitating more personalized and effective recommendations [20].

SimCLR [8] optimizes agreement between augmentations of the same image while minimizing agreement between different images. It uses data augmentation techniques and projection head designs to learn visual representations. MoCo [15] employs a momentum-based encoder and a memory bank for dynamic embedding updates, enhancing robustness.

Other methods like RotNet, BERT, and CPC expand the scope of contrastive learning. RotNet [12] improves object recognition by training on rotated images. BERT [10] predicts masked tokens in sentences, capturing deep dependencies in natural language processing. CPC [29] generates compact representations by predicting future observations in sequences, aiding sequential analysis.

3 Methods

In this section, we introduce CULC-Net (Contrastive User Learning for Creative Selection), designed to address unreliable CTR data due to limited exposure. We start with the problem formulation and the base model of creative selection. We then compare CULC-Net with the base model and provide detailed explanations of CUL and FlexiRank, emphasizing their impact on enhancing model accuracy.

3.1 Preliminary

Problem Formulation: We address a creative selection scenario within an ad recommendation system, characterized by a user request $u \in U$ and a set of ads $A = \{a_j\}_{j=1}^{n}$. Each ad a_j includes m creatives, represented as $a_j = \{c_i\}_{i=1}^{m}$, where each i-th creative c_i is described by a tuple (v_i, t_i, id_i). Here, v_i, t_i, and id_i correspond to the image feature, text feature, and unique hash ID. Our objective is to select the optimal creative for each ad by integrating visual and textual information of the creatives with the unique user preferences.

STM (Single-Tower Model): The Single-Tower Model utilizes embeddings [27], DNNs [19], and cross-entropy [9] to extract and learn features from the creative data. The image and text features are integrated at the embedding layer, subsequently processed through a DNN to derive a more abstract representation. This representation is used to predict the hard label for each creative, with cross-entropy serving as the loss function to optimize the model parameters. Base Creative selection model is the STM model.

TTM (Two-Tower Model): The Two-Tower Model employs two distinct DNNs to enhance the separation of feature processing: one for user features and another for item features. This architecture allows each tower to specialize in extracting detailed representations from its respective domain. User features and item features are embedded and processed independently through their respective DNNs. The resulting features are then combined in a fusion layer, facilitating effective interaction between user preferences and item attributes. Ad ranking stage is based on this model [2, 21].

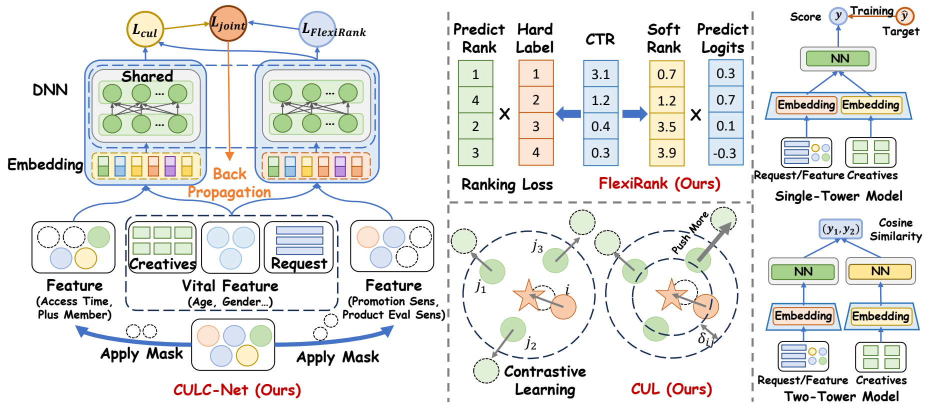

Fig. 2. Visual comparison of the base creative selection method and the detailed workings of CULC-Net. STM employs a single tower for recommendation, whereas CULC-Net integrates shared DNN and different user knowledge to improve CTR prediction. We also show the differences between Contrastive Learning and our CUL approach, as well as ranking loss and FlexiRank, highlighting the advantages of soft-decision making and the comprehensive utilization of ranking data.

3.2 Overview of CULC-Net

Inspired by [35], we adopt similar contrastive learning algorithms to learn representations of categorical features. We apply different data augmentations to the training examples to learn these representations and then use an adaptive contrastive loss function to ensure that the representations learned for the same training example are similar. Our proposed CULC-Net method enhances the STM by introducing a shared DNN and incorporating varied user knowledge for different augmentations. The shared DNN layer ensures that the model learns a unified representation for both image and text features, while the integration of diverse user knowledge allows the model to adapt to varying user preferences more effectively.

Figure 2 offers a visual comparison between STM and our proposed CULC-Net approach. The illustration underscores the differences in architecture, particularly emphasizing the shared DNN and the varied user knowledge embedded within our model. These enhancements contribute to superior creative selection and a better overall user experience.

Unlike some Two-Tower Model in advertising recommendation systems [23], where one tower handles user requests and features, and the other manages ads and creatives, our model leverages data augmentation techniques to improve ranking for similar users. This approach helps address issues like exposure bias and the long-tail distribution of data. By modifying changeable features, our method brings similar users closer together, thereby enhancing the creative selection process. The integration of a shared DNN and diverse user knowledge within our two-tower model, termed CULC-Net, significantly boosts the effectiveness of creative selection. This results in more tailored and relevant recommendations, enhancing the user experience in consuming multimedia content.

3.3 CUL Augmentation: Which and Where

The key point of CUL (Contrastive User Learning) is feature augmentation, implemented through a strategic two-stage process: which features to mask and where to apply the masking. In the first stage, we determine 'which' features to focus on by categorizing them into two types: important and changeable. Important features, such as user gender, age, and customer purchase history, are consistently shared across both towers to maintain a solid base for user profiling. The second stage addresses 'where' to apply the masking. Changeable features, including user access time, membership status (e.g., plus member), sensitivity to promotions, and sensitivity to product evaluations, are randomly masked at a rate of 35% to create two augmented data instances. This approach allows the model to effectively handle user behavior variance and improve overall performance.

3.4 From Contrastive Learning to CUL

In our CULC-Net, we design a CUL framework that enhances the outputs of the two towers to be more similar for identical inputs while minimizing the impact of inaccurate CTR data. We first present the formula for Contrastive Learning and then extend it to Contrastive User Learning.

Given a batch of N examples $a_1, \ldots, a_N$, where $a_i \in A$ denotes a set of features for example i, we define x_i and x_i' as the respective inputs for the two towers. Our goal is to learn distinct representations x_i and x_i' while ensuring the model recognizes both as originating from the same input i.

We aim to minimize the difference between x_i and x_i' while maximizing the difference between the representations learned for distinct examples i and j. Let y_i and y_i' represent the outputs of x_i and x_i'. We consider (y_i, y_i') as positive pairs and (y_i, y_j') as negative pairs for $i \neq j$. Let $s(y_i, y_j') = \frac{1}{\tau} \cdot \frac{\langle y_i, y_j' \rangle}{|y_i| \cdot |y_j'|}$, where τ is a temperature parameter that controls the concentration of the probability distribution. To promote the desired properties, we define the InfoNCE loss for a batch of N examples as:

$$L_{\mathrm{CL}} = -\frac{1}{N} \sum_{i=1}^{N} \log \left(\frac{\exp(s(y_i, y_i'))}{\sum_{j=1}^{N} \exp(s(y_i, y_j'))} \right), \tag{1}$$

This is a vanilla InfoNCE loss that facilitates the comparison of positive and negative creatives. However, due to the presence of low exposure creatives, which may be false negatives, not all $i \neq j$ in $\sum_{j=1}^{N} \exp(s(y_i, y_j')$ should be considered as negative. We define a relevance score δ_{ij} representing the similarity distance from each negative item y_j' to the positive anchor y_i. Its formula is set as:

$$s(y_i, y_j') - s(y_i, y_i') + \delta_{ij} < 0, \tag{2}$$

integrating this into the Contrastive Learning formula:

$$
\begin{aligned}
L_{\mathrm{CUL}} = -\frac{1}{N}\sum_{i=1}^{N}\log\frac{\exp(s(y_i,y_i'))}{\sum_{j=1}^{N}\exp(s(y_i,y_j')+\delta_{ij})} = \\
-\frac{1}{N}\sum_{i=1}^{N}\log\frac{\exp(s(y_i,y_i'))}{\exp(s(y_i,y_i'))+\sum\limits_{j\neq i}\exp(s(y_i,y_j')+\delta_{ij})},
\end{aligned}
\tag{3}
$$

When $\delta_{ij} = 0$, it simplifies to the vanilla InfoNCE loss function. This relevance score δ_{ij} helps display the ranking among negative creatives; a larger δ_{ij} suggests a stronger negative sample, while a smaller δ_{ij} indicates a potential false negative. Models may learn implicit relationships between feature and different negative samples under the vanilla InfoNCE loss, but we argue that modeling this relationship explicitly by δ_{ij} has a positive influence on learning better representations.This approach allows the model to give more weight to certain negative samples, considering a broader spectrum of negatives.

The specific computation of the loss function is as follows, using adversarial training [25] to learn the difficulty of specific negative samples:

$$
\begin{aligned}
&\min_{\theta} L_{\mathrm{CUL}} = \min_{\theta}\max_{\Delta\in C} \\
&-\frac{1}{N}\sum_{i=1}^{N}\log\frac{\exp(s(y_i,y_i'))}{\exp(s(y_i,y_i'))+\sum\limits_{j\neq i}\exp(\delta_{ij})\exp(s(y_i,y_j'))},
\end{aligned}
\tag{4}
$$

where C encapsulates the collective set of all δ_{ij}.for each δ_{ij} we have:

$$
\exp(\delta_{ij}) \in C(ij) = (1-\epsilon, 1+\epsilon)\,,
$$

and ϵ is a hyperparameter that regulates the upper-bound deviation of hardness. In practice, ϵ is regulated by the number of adversarial training epochs under a fixed learning rate.

In this way, we reduce the reliance on individual CTR data by exploring the relationships between creatives and customizing recommendations for users with similar features.

3.5 FlexiRank for Utilizing All CTR Data

To better learn the ranking information of creatives, not just identifying the best but understanding which is better, we introduce the FlexiRank method, which utilizes soft ranking to effectively manage the predicted outputs $\{y_1, y_2, \dots, y_m\}$ and their corresponding soft labels $\{l_1, l_2, \dots, l_m\}$.

The motivation for FlexiRank is rooted in the necessity to accurately rank creatives according to their potential to engage users. In softmax, labels($\hat{y}$) are one-hot encoded, which only predicts the accuracy for top one creative, thus losing information of the relative ranking of creatives. Inspired by seminal works [4, 23, 31], we designed FlexiRank, which normalizes CTR to obtain soft labels. FlexiRank considers the relative relationships and rankings among different creatives, exploring the potential of all creatives to attract users and providing more accurate recommendations. The probability that a creative is ranked as the top choice is defined by:

$$P_i = \frac{\exp(y_i)}{\sum_{k=1}^{m} \exp(y_k)}, \tag{5}$$

where $\exp(\cdot)$ denotes the exponential function. The corresponding probability of the soft labels is defined as:

$$\hat{P}_i = \frac{\exp(l_i/T)}{\sum_{k=1}^{m} \exp(l_k/T)}, \tag{6}$$

where T is a temperature coefficient, adjusting the scale when y_m is small.

To enable the model to learn ranking information better in the initial epochs and improve contrastive learning in the later epochs, FlexiRank includes an instructional gradient inspired by Curriculum Learning [3]:

$$L_{\text{FlexiRank}} = -\sum_{i=1}^{m} \hat{P}_i \log(P_i) + \sum_{i=1}^{m} L(x_i, \hat{P}_i, t), \tag{7}$$

where $L(x_i, \hat{P}_i, t) = -\hat{P}_i \cdot \log(p_t(x_i)) \cdot \exp(-\alpha \cdot t)$ represents the instructional gradient, $p_t(x_i)$ is P_i at step t, and α is a hyperparameter.

This objective function is designed to capture the relative ordering of creatives and to facilitate smooth knowledge acquisition within our CULC-Net. By effectively balancing the influence of predicted outputs and soft labels, the FlexiRank method leads to a more robust and generalizable model.

3.6 CULC-Net Optimization

To leverage the advantages of Contrastive User Learning and FlexiRank, CULC-Net employs a unified learning framework that jointly optimizes both the contrastive user learning loss (L_{CUL}) and the flexible ranking loss ($L_{\text{FlexiRank}}$). This joint optimization approach enables the model to learn more informative and effective parameters, enhancing creative selection performance. The training process involves minimizing a joint loss function formulated as:

$$L_{\text{joint}} = L_{\text{FlexiRank}} + \lambda L_{\text{CUL}}, \tag{8}$$

where λ is a hyperparameter that balances the adaptive regularization and contrastive user learning losses.

Table 1. Offline evaluation results of 5 creative selection models across 7 creative settings. CULC-Net demonstrates significant improvements in All setting, with a 1.06% increase in sCTR, a 1.35% increase in AUC, and a 0.85% increase in GAUC, showcasing our advantage. Non-CL also shows a clear improvement over STM, highlighting the effectiveness of the FlexiRank.

Metrics	Image	Title	STM	VAM-HBM	CACS	Non-CL (Ours)	CULC-Net (Ours)
sCTR	single	—	28.54%	28.73%	28.82%	29.12%	29.94% (+1.40%)
	single	single	27.22%	27.43%	27.51%	27.77%	28.63% (+1.41%)
	multi	—	23.12%	23.31%	23.42%	23.67%	24.98% (+1.86%)
	multi	multi	26.07%	26.21%	26.33%	26.61%	27.29% (+1.22%)
	—	single	31.82%	31.94%	32.02%	32.23%	33.82% (+2.00%)
	—	multi	24.87%	25.03%	25.14%	25.72%	26.88% (+2.01%)
	ALL	ALL	26.41%	26.63%	26.73%	26.95%	27.47% (+1.06%)
AUC	single	—	64.42%	64.74%	64.89%	65.37%	66.43% (+2.01%)
	single	single	67.15%	67.38%	67.52%	68.09%	69.12% (+1.97%)
	multi	—	70.46%	70.68%	70.84%	71.25%	72.44% (+1.98%)
	multi	multi	67.04%	67.27%	67.49%	68.01%	69.07% (+2.03%)
	—	single	62.54%	62.41%	62.23%	61.87%	63.54% (+1.00%)
	—	multi	67.91%	68.05%	68.16%	68.33%	69.90% (+1.99%)
	ALL	ALL	67.11%	67.34%	67.47%	67.84%	68.46% (+1.35%)
GAUC	ALL	ALL	58.06%	58.17%	58.24%	58.33%	58.91% (+0.85%)

4 Experiments

4.1 Datasets and Experimental Setup

Our experiments utilize a comprehensive dataset from a real-world search ad system, which includes user click history and various ad attributes collected over a one-month period. Specifically, the dataset comprises approximately 9 billion training samples and 200 million test samples, encompassing interactions from 46 million users with 12 million ads and 43 million ad creatives. The creatives include images and titles. In our offline experiments, we categorized them into Single Image, Multi-Image, Single Title, Multi-Title, Combination (combining images and titles, single or multiple), and All (including all types of creatives).

We employ the Adam optimizer with a learning rate of 0.001, beta1 of 0.9, beta2 of 0.999, and epsilon of 1e-9. For online and offline experiments, λ is 0.6, α is 0.2, and ϵ is 0.1. The model is trained with a batch size of 256. Additionally, we did not use the dropout strategy, and we utilize ReLU activation with a sigmoid output layer to ensure predictions are bound within the range of (0, 1).

4.2 Evaluation Metrics

In order to assess the performance of our creative selection model, we employ specific evaluation metrics for both offline and online experiments. For the online experiments, CTR is used as the primary evaluation metric.

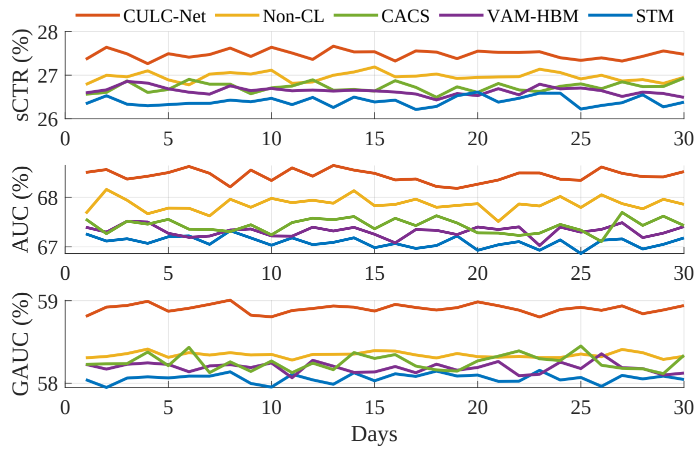

Fig. 3. Offline daily comparison of sCTR, AUC, and GAUC between STM, VAM-HBM, CACS, Non-CL, and CULC-Net over 30 days. The comparison is based on all available data, with CULC-Net consistently achieving the best performance across all metrics.

For the offline experiments, we use Simulated Click-Through Rate (sCTR) [31], Area Under the Curve (AUC), and Group AUC (GAUC) to comprehensively evaluate the model's performance in a controlled environment that mimics real-world conditions. By employing these metrics, we demonstrate the effectiveness and robustness of our proposed CULC-Net in both offline and online settings.

4.3 Baselines

To effectively evaluate the performance of CULC-Net, we introduce a range of advanced baseline models. Below is a concise overview of these models:

STM: employing a single-tower structure, feeds features directly into a multi-layer network, culminating in a classification problem for optimization. This method provides a fundamental benchmark for comparison, excluding the use of Contrastive Learning.

VAM-HBM [31]: combining the Visual-aware Ranking Model (VAM) for learning visual features related to performance and the Hybrid Bandit Model (HBM) for updating its understanding based on prior data, offering a responsive model to creative selection.

CACS [23]: employing a two-tower structure reminiscent of the widely recognized Deep Structured Semantic Model (DSSM) [17]. It maps creatives and user queries into a common semantic space, allowing it to gauge their relevance based on proximity.

Non-CL(Ours): CULC-Net excludes the CUL loss, essentially replacing softmax in STM with FlexiRank to validate the effectiveness of FlexiRank.

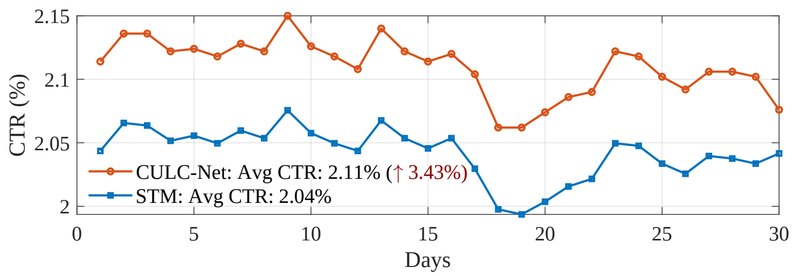

Fig. 4. Daily CTR comparison between CULC-Net and STM over 30 days in online A/B test. The average CTR for STM is 2.04%, while CULC-Net achieves 2.11%, with a 3.43% increase. The result demonstrates the effectiveness of CULC-Net in the real-world online advertising scenario.

4.4 Offline Results

In this section, we discuss the performance of our proposed method, CULC-Net, and compare it with the baseline method, the Non-CL method (without contrastive learning), as well as two recent prominent approaches, VAM-HBM and CACS. We first present a comprehensive comparison test of the creatives, including single image, multi-image, single title, multi-title, single combination, multi-combination, and all data in Table 1. We observe that CULC-Net consistently outperforms the other methods in all test scenarios, affirming its superiority in creative selection.

In terms of total performance across all data, CULC-Net achieves a 4.01% increase in sCTR, a 2.01% increase in AUC, and a 1.44% increase in GAUC compared to the baseline method. Notably, when contrastive learning is incorporated (compared to the Non-CL method), the performance improvement is significant. The sCTR increases by 1.93%, AUC by 0.91%, and GAUC by 0.99%, illustrating the effectiveness of contrastive learning in capturing the nuances between creatives, ultimately leading to a more refined understanding of the underlying patterns.

Moreover, our analysis extends to the comparison with VAM-HBM and CACS. The results suggest that while these methods exhibit competitive performance in certain scenarios, CULC-Net consistently provides a more comprehensive and accurate creative selection mechanism. Specifically, CULC-Net demonstrates a more pronounced ability to leverage the complex interplay between different creative elements, as evidenced by its superior performance in multi-image and multi-title scenarios.

Figure 3 showcases the daily comparison of sCTR, AUC, and GAUC between STM, Non-CL, VAM-HBM, CACS, and CULC-Net over a span of 30 days, where each method processed one day of data from the dataset. The results underscore that CULC-Net is not only effective but also exhibits remarkable stability, consistently outperforming the other methods across all evaluation metrics.

Fig. 5. Visualization of selected creatives by CULC-Net and STM. Each column presents one pair of creatives with the same item. CULC-Net attracts more clicks and recommends better creatives for both males and females.

4.5 Online Results

To investigate the effectiveness of CULC-Net in a real-world scenario, we conducted a 30-day A/B test comparing its performance with STM. The daily and average CTRs for both methods are presented in Figure 4. From the results, we observe that CULC-Net consistently outperforms STM in terms of daily CTR, with a relative increase in average CTR of 3.43%, demonstrating the effectiveness of our approach.

Figure 5 visualizes the creatives selected by CULC-Net and STM. Each column shows a pair of creatives associated with the same item, where CULC-Net consistently selects creatives that are better designed and more informative, leading to higher click rates. This analysis also compares male and female responses, highlighting enhanced engagement: males show a notable increase in clicks for items like headphones and refrigerators, while females demonstrate a significant uplift for footwear. These trends suggest that the improvements in creative selection by CULC-Net are influenced by gender-specific preferences. This further supports the effectiveness of CULC-Net in optimizing creative selection and enhancing advertising performance.

Our approach demonstrates that incorporating user knowledge and contrastive learning into the creative selection process can result in more effective ad creatives, which in turn, enhances platform revenue and contributes to a more engaging user experience.

4.6 Ablation Study

Offline Results on Public Dataset. Due to the lack of open-source datasets and the specific requirements for online performance, most creative selection methods [33, 34] have only been tested on proprietary datasets. The Creative

Table 2. Offline results on CreativeRanking dataset. This public dataset lacks user information, so we only test the effectiveness of FlexiRank (Non-CL). The results proving the robustness of our approach across multiple datasets.

Method	sCTR	AUC	GAUC
STM	2.95%	60.13%	54.27%
VAM-HBM	3.32%	62.75%	57.09%
CACS	3.27%	61.98%	56.35%
Non-CL	3.48% (↑17.9%)	63.42% (↑5.5%)	57.88% (↑6.7%)

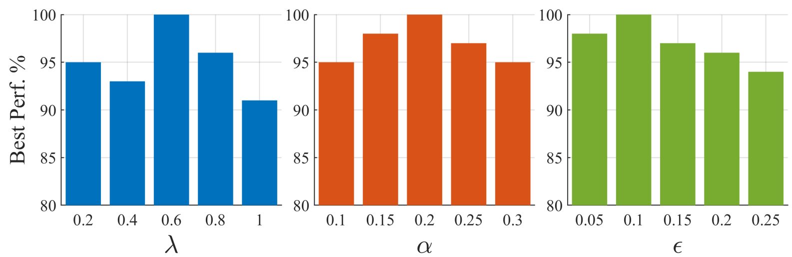

Fig. 6. Impact of key hyperparameters. We test the effects of λ, α, and ϵ in offline experiments, with the y-axis representing the percentage of the best performance.

Ranking Dataset [31] is the only public dataset we identified for the creative selection task. This dataset includes advertising creative data from Taobao, collected from July to August 2020, with 500,827 product samples, 1,204,988 different creatives, and over 200 million impressions. However, since this dataset lacks user information, we could not evaluate CULC-Net directly and instead focused on testing the FlexiRank component (Non-CL). We conducted additional offline experiments on this dataset to verify the effectiveness of FlexiRank. As shown in Table 2, Non-CL achieved the best results, with a 17.9% improvement in sCTR, a 5.5% increase in AUC, and a 6.7% rise in GAUC. These results further demonstrate the advantages of utilizing soft ranking.

Parameter Sensitivity. We analyze the impact of key hyperparameters of CULC-Net, as shown in Figure 6, where the y-axis represents the percentage of the best performance.

1.Impact of λ: The balance between FlexiRank and contrastive user learning losses is best achieved with $\lambda = 0.6$, with performance dropping off as λ increases.

2.Impact of α: The optimal decay rate is found with $\alpha = 0.2$, ensuring stable training and effective learning.

3.Impact of ϵ: The hardness of negative samples is best regulated by $\epsilon = 0.1$, striking the right balance between challenge and learnability.

5 Conclusion

In this paper, we introduced CULC-Net, a novel approach for optimizing ad creative selection by addressing the challenges of unreliable CTR data and sparse data exposures. Unlike existing methods that rely heavily on surface-level analysis, CULC-Net delves deeper into user behavior, uncovering latent connections between users and creatives. This approach enables a more personalized and effective advertising experience. Our extensive evaluation shows a 3.43% improvement in online and 4.01% in offline.

The key contribution of CULC-Net lies in its ability to enhance data robustness and relevance, even with unbalanced datasets, leading to more effective ad recommendations. This improvement is economically significant; for instance, on a large e-commerce platform with daily sales of $100 million, a 3% increase in CTR could yield an additional $3 million in revenue. Moving forward, CULC-Net sets the stage for future innovations in personalized, data-driven advertising, emphasizing the critical role of user-centric approaches in optimizing online advertising.

Acknowledgments. This research was funded by the National Science and Technology Major Project under contract No. 2021ZD0109903 and the National Natural Science Foundation of China (No. 82441013).

References

1. Auer, P., Cesa-Bianchi, N., Fischer, P.: Finite-time analysis of the multiarmed bandit problem. Machine Learning **47**, 235–256 (2002)
2. Baltescu, P., Chen, H., Pancha, N., Zhai, A., Leskovec, J., Rosenberg, C.: Itemsage: Learning product embeddings for shopping recommendations at pinterest. In: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. pp. 2703–2711 (2022)
3. Bengio, Y., Louradour, J., Collobert, R., Weston, J.: Curriculum learning. In: Proceedings of the International Conference on Machine Learning. pp. 41–48 (2009)
4. Cao, Z., Qin, T., Liu, T.Y., Tsai, M.F., Li, H.: Learning to rank: from pairwise approach to listwise approach. In: Proceedings of the International Conference on Machine Learning. pp. 129–136 (2007)
5. Chen, J., Ge, T., Jiang, G., Zhang, Z., Lian, D., Zheng, K.: Efficient optimal selection for composited advertising creatives with tree structure. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 3967–3975 (2021)
6. Chen, J., Xu, J., Jiang, G., Ge, T., Zhang, Z., Lian, D., Zheng, K.: Automated creative optimization for e-commerce advertising. In: Proceedings of the Web Conference 2021. pp. 2304–2313 (2021)
7. Chen, J., Sun, B., Li, H., Lu, H., Hua, X.S.: Deep ctr prediction in display advertising. In: Proceedings of the 24th ACM International Conference on Multimedia. pp. 811–820 (2016)
8. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: Proceedings of the International Conference on Machine Learning. pp. 1597–1607. PMLR (2020)
9. De Boer, P.T., Kroese, D.P., Mannor, S., Rubinstein, R.Y.: A tutorial on the cross-entropy method. Annals of Operations Research **134**, 19–67 (2005)
10. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. Conference of the North American Chapter of the Association for Computational Linguistics (2019)

11. Ge, T., Zhao, L., Zhou, G., Chen, K., Liu, S., Yi, H., Hu, Z., Liu, B., Sun, P., Liu, H., et al.: Image matters: Visually modeling user behaviors using advanced model server. In: Proceedings of the 27th ACM International Conference on Information & Knowledge Management. pp. 2087–2095 (2018)
12. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. Proceedings of the International Conference on Learning Representations (2018)
13. Glowacka, D.: Bandit algorithms in recommender systems. In: RecSys '19. pp. 574–575 (2019)
14. Guo, H., Tang, R., Ye, Y., Li, Z., He, X.: Deepfm: a factorization-machine based neural network for ctr prediction. Proceedings of The International Joint Conferences on Artifical Intelligence (2017)
15. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9729–9738 (2020)
16. He, X., Pan, J., Jin, O., Xu, T., Liu, B., Xu, T., Shi, Y., Atallah, A., Herbrich, R., Bowers, S., et al.: Practical lessons from predicting clicks on ads at facebook. In: Proceedings of the eighth international workshop on data mining for online advertising. pp. 1–9 (2014)
17. Huang, P.S., He, X., Gao, J., Deng, L., Acero, A., Heck, L.: Learning deep structured semantic models for web search using clickthrough data. In: Proceedings of the 22th ACM International Conference on Information & Knowledge Management. pp. 2333–2338 (2013)
18. Jaiswal, A., Babu, A.R., Zadeh, M.Z., Banerjee, D., Makedon, F.: A survey on contrastive self-supervised learning. Technologies **9**(1), 2 (2020)
19. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. Nature **521**(7553), 436–444 (2015)
20. Li, H., Zhou, X., Tuan, L.A., Miao, C.: Rethinking negative pairs in code search. Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (2023)
21. Li, S., Guo, H., Hou, L., Zhang, W., Tang, X., Tang, R., Zhang, R., Li, R.: Adaptive low-precision training for embeddings in click-through rate prediction. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 4435–4443 (2023)
22. Li, X., Wang, C., Tan, J., Zeng, X., Ou, D., Ou, D., Zheng, B.: Adversarial multimodal representation learning for click-through rate prediction. In: Proceedings of The Web Conference 2020. pp. 827–836 (2020)
23. Lin, K., Zhang, X., Li, F., Wang, P., Long, Q., Deng, H., Xu, J., Zheng, B.: Joint optimization of ad ranking and creative selection. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 2341–2346 (2022)
24. Liu, H., Lu, J., Yang, H., Zhao, X., Xu, S., Peng, H., Zhang, Z., Niu, W., Zhu, X., Bao, Y., et al.: Category-specific cnn for visual-aware ctr prediction at jd. com. In: Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery. pp. 2686–2696 (2020)
25. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. Proceedings of the International Conference on Learning Representations (2018)
26. McMahan, H.B., Holt, G., Sculley, D., Young, M., Ebner, D., Grady, J., Nie, L., Phillips, T., Davydov, E., Golovin, D., et al.: Ad click prediction: a view from the

trenches. In: Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining. pp. 1222–1230 (2013)
27. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. arXiv:1301.3781 (2013)
28. Mo, K., Liu, B., Xiao, L., Li, Y., Jiang, J.: Image feature learning for cold start problem in display advertising. In: Proceedings of The International Joint Conferences on Artifical Intelligence (2015)
29. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
30. Schwartz, E.M., Bradlow, E.T., Fader, P.S.: Customer acquisition via display advertising using multi-armed bandit experiments. Marketing Science **36**(4), 500–522 (2017)
31. Wang, S., Liu, Q., Ge, T., Lian, D., Zhang, Z.: A hybrid bandit model with visual priors for creative ranking in display advertising. In: Proceedings of the Web Conference 2021. pp. 2324–2334 (2021)
32. Wei, P., Liu, S., Yang, X., Wang, L., Zheng, B.: Towards personalized bundle creative generation with contrastive non-autoregressive decoding. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 2634–2638 (2022)
33. Wei, P., Yang, X., Liu, S., Wang, L., Zheng, B.: Creater: Ctr-driven advertising text generation with controlled pre-training and contrastive fine-tuning. arXiv preprint arXiv:2205.08943 (2022)
34. Yang, Z., Sang, L., Wang, H., Chen, W., Wang, L., He, J., Peng, C., Lin, Z., Gan, C., Shao, J.: Parallel ranking of ads and creatives in real-time advertising systems. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 9278–9286 (2024)
35. Yao, T., Yi, X., Cheng, D.Z., Yu, F., Chen, T., Menon, A., Hong, L., Chi, E.H., Tjoa, S., Kang, J., et al.: Self-supervised learning for large-scale item recommendations. In: Proceedings of the 30th ACM International Conference on Information & Knowledge Management. pp. 4321–4330 (2021)
36. Zhao, Z., Li, L., Zhang, B., Wang, M., Jiang, Y., Xu, L., Wang, F., Ma, W.: What you look matters? offline evaluation of advertising creatives for cold-start problem. In: Proceedings of the 28th ACM International Conference on Information & Knowledge Management. pp. 2605–2613 (2019)