Learning to Expand Audience via Meta Hybrid Experts and Critics for Recommendation and Advertising

Yongchun Zhu^{1,3,4}, Yudan Liu³, Ruobing Xie³, Fuzhen Zhuang^{2,5,*}, Xiaobo Hao³, Kaikai Ge³,
Xu Zhang³, Leyu Lin³ and Juan Cao^{1,4}

¹Key Lab of Intelligent Information Processing of Chinese Academy of Sciences (CAS), Institute of Computing Technology, CAS, Beijing 100190, China ²Institute of Artificial Intelligence, Beihang University, Beijing 100191, China

³WeChat Search Application Department, Tencent, China

⁴University of Chinese Academy of Sciences, Beijing 100049, China

⁵Xiamen Data Intelligence Academy of ICT, CAS, China

{zhuyongchun18s, caojuan}@ict.ac.cn, {danydliu, ruobingxie, rolyhao, kavinge, xuonezhang, goshawklin}@tencent.com, zhuangfuzhen@buaa.edu.cn

ABSTRACT

In recommender systems and advertising platforms, marketers always want to deliver products, contents, or advertisements to potential audiences over media channels such as display, video, or social. Given a set of audiences or customers (seed users), the audience expansion technique (look-alike modeling) is a promising solution to identify more potential audiences, who are similar to the seed users and likely to finish the business goal of the target campaign. However, look-alike modeling faces two challenges: (1) In practice, a company could run hundreds of marketing campaigns to promote various contents within completely different categories every day, e.g., sports, politics, society. Thus, it is difficult to utilize a common method to expand audiences for all campaigns. (2) The seed set of a certain campaign could only cover limited users. Therefore, a customized approach based on such a seed set is likely to be overfitting.

In this paper, to address these challenges, we propose a novel two-stage framework named Meta Hybrid Experts and Critics (MetaHeac) which has been deployed in WeChat Look-alike System. In the offline stage, a general model which can capture the relationships among various tasks is trained from a meta-learning perspective on all existing campaign tasks. In the online stage, for a new campaign, a customized model is learned with the given seed set based on the general model. According to both offline and online experiments, the proposed MetaHeac shows superior effectiveness for both content marketing campaigns in recommender systems and advertising campaigns in advertising platforms. Besides, MetaHeac has been successfully deployed in WeChat for the promotion of both contents and advertisements, leading to great improvement in the quality of marketing. The code has been available at <https://github.com/easezy/MetaHeac>.

*Fuzhen Zhuang is the corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '21, August 14–18, 2021, Virtual Event, Singapore

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8332-5/21/08...\$15.00

<https://doi.org/10.1145/3447548.3467093>

CCS CONCEPTS

• **Information systems** → **Recommender systems**; **Online advertising**.

KEYWORDS

Look-alike, Audience Expansion, Meta Learning, Campaign

ACM Reference Format:

Yongchun Zhu^{1,3,4}, Yudan Liu³, Ruobing Xie³, Fuzhen Zhuang^{2,5,*}, Xiaobo Hao³, Kaikai Ge³, Xu Zhang³, Leyu Lin³ and Juan Cao^{1,4}. 2021. Learning to Expand Audience via Meta Hybrid Experts and Critics for Recommendation and Advertising. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '21)*, August 14–18, 2021, Virtual Event, Singapore. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3447548.3467093>

1 INTRODUCTION

With the booming development of the mobile internet, there are more and more mobile applications in the world, and a large number of users are active on these applications every day. In such a new mobile market with billions of users, it becomes crucial for marketers to effectively deliver contents [12], advertisements [2] or products [13, 35] to potential audiences by recommender systems or advertising platforms. To achieve the goals of various promotion activities, Internet companies conduct hundreds of marketing campaigns every day. One of the keys of running marketing campaigns is the audience expansion technique (look-alike modeling) which has been deployed in many online systems, e.g., Google [8], LinkedIn [11], Pinterest [2], Ant Financial [35], WeChat [12].

For a certain campaign, given a set of audiences (seed users), the audience expansion technique aims to identify more potential audiences who are similar to the seed users and likely to finish the business goal of the target campaign. High-quality expanded audiences will benefit the companies by increasing the conversion population while reducing the operating costs for their marketing campaigns [35]. A good look-alike technique can result in a great economic benefit, but it suffers from two significant challenges. (1) The tasks of various campaigns can cover diverse contents, e.g., sports, politics, society, and the promoted contents can vary from task to task. Thus, it is hard to exploit a common method for all campaign tasks. (2) A certain campaign gives a seed set that can only cover limited users. Based on this seed set, a customized

approach for the campaign is likely to be overfitting. Especially, some marketing campaigns only have a few hundred seeds.

With designed common rules for all campaigns, previous rule-based approaches [18, 24] match similar users with specific demographic tags (age, gender, geography) or interests that are targeted by marketers, which has unsatisfying performance. Using common pre-defined similarity function objectively such as Cosine or Jaccard for all campaigns, traditional similarity-based look-alike methods [3, 11, 16, 17] usually find look-alike users by directly comparing all possible pairs between seed users and available users. However, the performance of these methods largely depends on selecting the right features and similarity functions, and it is hard to achieve satisfying performance. Recently, training customized prediction models for each campaign, model-based look-alike approaches [7, 12, 23] achieve remarkable results. However, these one-stage model-based look-alike methods [7, 12, 23] which directly train customized models from scratch for each campaign seriously suffer from the overfitting problem.

To tackle these challenges, some model-based methods [2, 13, 35] decompose the audience expansion tasks into two stages: offline stage and online stage. In the offline stage, they train a common embedding layer. In the online stage, based on the common embedding, customized models are trained for each campaign. However, these methods cannot solve the two challenges for three reasons. (1) The goal of look-alike modeling is to perform well for new campaign tasks. Nevertheless, the common embedding is just fitted on existing campaigns, and the generalization ability [28] is not taken into consideration. (2) They directly train the common embedding layer over a large user-campaign graph ignoring the relationships among various tasks. Actually, the inherent conflicts from task differences can actually harm the predictions of some tasks, e.g., the seed users of a certain task can be negative samples of another task. (3) They only pre-train the embedding layer without the deep network. Thus, for a certain campaign, the customized network could be still overfitting to the seed users.

To address these issues, we propose a novel two-stage audience expansion framework named Meta Hybrid Experts and Critics (MetaHeac). The framework of MetaHeac follows a hybrid online-offline manner. The core of the offline stage is to train the general model with various existing marketing campaign tasks. In the online stage, when a new marketing campaign comes, a customized model which serves for audience expansion is learned by fine-tuning the pre-trained general model with the given seed users. MetaHeac mainly focuses on training a better general model and adopts two high-level key ideas: 1) The general model is expected to learn the ability to expand audiences, 2) The general model should learn transferable knowledge from various marketing campaigns.

Recall that the look-alike modeling consists of two phases, (1) understand the characteristics of seed users, (2) find the potential audiences who are similar to the seed users. Thus, in order to learn how to expand audiences, we design a two-phase simulation over hundreds of marketing campaign tasks. The simulation consists of an ‘understanding’ phase and a ‘finding’ phase. At the ‘understanding’ phase, the model needs to understand the characteristics of the seed users. At the ‘finding’ phase when we have found some potential audiences, we update the model with the expanded audiences. With the advantage of meta-learning [6], the model trained

Figure 1: The architecture of WeChat Look-alike System.

on learning tasks of such two-phase simulation can learn how to expand audiences.

To learn transferable knowledge, it is important to model the relationships among tasks. Thus, MetaHeac takes a hybrid structure that consists of multiple experts and critics. MetaHeac utilizes multiple experts to capture multiple user representations. Since different experts specialize in different tasks, we propose a task-driven gate to combine user’s multiple representations into a single representation. Then, with the user representation, MetaHeac employs multiple critics to express their opinions about whether the user can be interested in the promoted contents. Having the scores given by each critic, a similar task-driven gate is exploited to count the final score of the user.

MetaHeac has been successfully deployed in WeChat Look-alike System, and it can work for WeChat to promote both contents and advertisements, leading to great improvement in the quality of marketing. According to both offline and online experiments, the proposed model shows superior effectiveness for both marketing campaigns in recommender systems and advertising campaigns in advertising platforms.

The main contributions of our work are summarized into four folds:

- We propose a novel two-stage framework named MetaHeac for the audience expansion problem, which significantly improves the performance.
- We formally define the audience expansion problem from a meta-learning perspective and use meta-learning to learn how to expand audiences.
- To learn transferable knowledge and model the task relationships, we propose a hybrid structure that consists of multiple experts and critics.
- We conduct both offline and online experiments to demonstrate the effectiveness and robustness of MetaHeac.

2 SYSTEM OVERVIEW

In this section, we will introduce the WeChat Look-alike System which contains the proposed MetaHeac framework. As shown in Figure 1, the system employs a hybrid online-offline architecture. Campaign System is a campaign management system that contains both seeds and online feedback of all campaigns.

The core of the offline stage is to maintain a pre-trained general model that can adapt fast to new campaigns. Thus, WeChat Look-alike System utilizes data of all existing campaigns from the campaign system to train the general model. Note that the data

consists of both the initial seeds and the online user feedback of all campaigns. To train the general model, the proposed MetaHeac framework is exploited. To learn how to expand audience, MetaHeac trains the general model from a meta-learning perspective. To capture task relationships and learn transferable knowledge, MetaHeac employs hybrid experts and critics. The details of MetaHeac will be introduced in the next section.

The online stage aims to find potential audiences for a certain campaign. For a new campaign, the marketers will give a certain seed set, and the campaign system will sample a negative set for training. With the data, a customized campaign model can be learned by fine-tuning the pre-trained general model. Then, with the customized model, the audience retrieval module will score all users in the candidate pool w.r.t. the target campaign, and then select the top K users with the highest scores as expanded audiences. Then, the contents of the target campaign will be promoted to the expanded audiences. The Log Parser will record the user feedback and pass it to the campaign system. The user feedback will be utilized to incrementally train the general model and the customized model.

Remark. The WeChat Look-alike System runs 400-500 marketing campaigns every day. The pre-trained general model is trained once a day with the new data of the day. The training of the customized model takes about 5 minutes. The audience retrieval of one billion users takes about 10 minutes. The customized model will be incrementally trained with real-time feedback once an hour.

3 THE PROPOSED MODEL

3.1 Definitions

For a certain marketing campaign c , the marketers want to promote a set of contents or products denoted as $\mathcal{V}$. A marketing campaign has its specific goal, e.g., promoting a category of contents. Thus, the campaign has some features which can represent the commonalities of the promoted contents, e.g., categories, tags, creative size. The campaign c can be denoted as $c = \{x_1^c, \dots, x_N^c\}$, where x represents feature and N is the number of features. Moreover, the marketers will give a set of seed users denoted as $\mathcal{S}_{[c]}$ which may be generated in different ways, e.g., selected by experts, retrieved according to the past relevant campaigns. To train a binary prediction model, both positive and negative samples are needed. Thus, the seed users $\mathcal{S}_{[c]}$ are labeled as positive, and the negative samples can come from random sampling on a candidate pool $\mathcal{U}$. And we denote the dataset consisting of both positive seeds and negative samples as $\mathcal{D}_{[c]}$. With $\mathcal{D}_{[c]}$ of a certain campaign, audience expansion aims to find a larger amount of users $\mathcal{U}_{[c]}$ (called expanded audiences) who are similar to users in $\mathcal{S}_{[c]}$ and likely to finish the business goal of the target campaign c . Both $\mathcal{D}_{[c]}$ and $\mathcal{U}_{[c]}$ come from the same candidate pool $\mathcal{U}$ denoted as $\mathcal{D}_{[c]}, \mathcal{U}_{[c]} \subset \mathcal{U}$. Note that the size of $\mathcal{U}_{[c]}$ is usually given by the marketer of the campaign c and is much larger than the size of $\mathcal{S}_{[c]}$. Each user $u \in \mathcal{U}$ may have multiple features, e.g., age, gender, city, interest tags. Therefore, a user can be denoted as $u = \{x_1^u, \dots, x_M^u\}$, where M denotes the number of user's feature.

Then we need to encode the raw features x into vectors. Look-up embedding has been widely adopted to learn dense representations from raw data for online prediction [20, 34]. Following [20, 33],

Algorithm 1 Training MetaHeac from a meta-learning perspective.

Input: Given hundreds of marketing campaign dataset $\mathcal{D}_{[c]}$.

Input: The general model f_θ .

Input: The learning rate α, β .

1. randomly initialize θ .
 2. **while** not converge **do**:
 3. sample batch of training tasks $\{\mathcal{T}_1, \dots, \mathcal{T}_n\}$.
 4. **for** $\mathcal{T}_i \in \{\mathcal{T}_1, \dots, \mathcal{T}_n\}$ **do**:
 5. $\mathcal{T}_i$ contains two disjoint sets $\mathcal{D}_{[c]}^a, \mathcal{D}_{[c]}^b$
 6. evaluate loss $\mathcal{L}_a(\theta)$ with $\mathcal{D}_{[c]}^a$
 7. compute updated parameter $\theta_{[c]} = \theta - \alpha \frac{\partial \mathcal{L}_a(\theta)}{\partial \theta}$
 8. evaluate loss $\mathcal{L}_b(\theta_{[c]})$ with $\mathcal{D}_{[c]}^b$
 9. **end**
 10. update $\theta = \theta - \beta \sum_{\mathcal{T}_i \in \{\mathcal{T}_1, \dots, \mathcal{T}_n\}} \frac{\partial \mathcal{L}_b(\theta_{[c]})}{\partial \theta}$
 11. **end**
-

we transform all features x into embeddings denoted as $\mathbf{e}$. Then, we denote the dense representations of a campaign and a user as $\mathbf{c} = \{\mathbf{e}_1^c, \dots, \mathbf{e}_N^c\}$ and $\mathbf{u} = \{\mathbf{e}_1^u, \dots, \mathbf{e}_M^u\}$, where $\mathbf{e} \in \mathbb{R}^k$ and k denotes the dimension of embeddings.

3.2 Learn to Expand Audience

A good look-alike model should be able to adapt to new marketing campaign tasks fast and alleviate the overfitting problem. Recently, meta-learning is a promising solution to achieve these goals: 1) meta-learning can learn general knowledge from lots of similar tasks [27], 2) meta-learning can achieve fast adaptation even only limited samples available [6]. Thus, we propose a novel look-alike method from a meta-learning perspective. The key idea is learning how to expand audiences.

For a certain marketing campaign c , the final goal of audience expansion is to learn a customized model $f_{[c]}(\cdot; \theta_{[c]})$ which can find the potential audiences from the candidate pool $\mathcal{U}$. The input of the model $f_{[c]}(\cdot; \theta_{[c]})$ consists of both user and campaign features. Then, the binary prediction model can be denoted as:

$$\hat{p} = f_{[c]}(\mathbf{c}, \mathbf{u}; \theta_{[c]}), \quad (1)$$

where $\theta_{[c]}$ denotes the parameters of customized model $f_{[c]}(\cdot; \theta_{[c]})$.

Note that the procedure of audience expansion can be divided into two phases: 1) 'Understanding': This phase aims to capture the characteristics of seed users $\mathcal{S}_{[c]}$ by training a customized model $f_{[c]}(\cdot)$ with the dataset $\mathcal{D}_{[c]}$. 2) 'Finding': The 'finding' phase exploits the customized model $f_{[c]}(\cdot)$ to seek an expanded audience set $\mathcal{U}_{[c]}$ from the candidate pool $\mathcal{U}$. To learn a general pre-trained model that knows how to expand audiences, we propose a training procedure to simulate the two phases.

To begin with, we denote $f(\cdot; \theta)$ as the general model, where θ is the general parameters. Note that the general model $f(\cdot; \theta)$ is different from the customized model $f_{[c]}(\cdot; \theta_{[c]})$. With the dataset $\mathcal{D}_{[c]}$ of a certain campaign, we construct a training task $\mathcal{T}$ by sampling two disjoint mini-batches from $\mathcal{D}_{[c]}$, and we denoted the two mini-batches as $\mathcal{D}_{[c]}^a$ and $\mathcal{D}_{[c]}^b$ to simulate the 'understanding' and 'finding' phases, respectively. $\mathcal{D}_{[c]}^a$ can be seen as the seed users and the randomly sampled negative examples, while $\mathcal{D}_{[c]}^b$

can be seen as the expanded audiences. The $\mathcal{D}_{[c]}^a$ and $\mathcal{D}_{[c]}^b$ are also called as the support set and the query set in the meta-learning literature.

Understanding phase. The meta-learner adapts the general θ to customized task-specific parameters $\theta_{[c]}$ w.r.t. the loss on the support set $\mathcal{D}_{[c]}^a$. Firstly, we make predictions using the general model $f(\cdot; \theta)$ on the support set $\mathcal{D}_{[c]}^a$ and get the predicted $\hat{p} = f(c, u; \theta)$. Then, the loss of this phase can be defined as:

$$\mathcal{L}_a(\theta) = \sum_{\mathcal{D}_{[c]}^a} [-y \log \hat{p} - (1 - y) \log(1 - \hat{p})], \quad (2)$$

where y is the label of the sample and $\mathcal{L}_a(\theta)$ denotes the loss computed based on θ . By minimizing the loss $\mathcal{L}_a(\theta)$, a customized model $f_{[c]}(\cdot; \theta_{[c]})$ can be learned, where $\theta_{[c]} = \theta - \alpha \frac{\partial \mathcal{L}_a}{\partial \theta}$. The α denotes the learning rate of this phase.

Finding phase. After the understanding phase, a customized model $f_{[c]}(\cdot; \theta_{[c]})$ for the campaign c is available. Similarly, based on $\theta_{[c]}$, we can compute the loss on the query set $\mathcal{D}_{[c]}^b$ as $\mathcal{L}_b(\theta_{[c]})$. In order to learn how to expand audience, we utilize the loss computed with the ‘expanded’ audience set $\mathcal{D}_{[c]}^b$ on parameters $\theta_{[c]}$ to update the general parameters θ , as:

$$\theta = \theta - \beta \frac{\partial \mathcal{L}_b(\theta_{[c]})}{\partial \theta} = \theta - \beta \frac{\partial \mathcal{L}_b(\theta_{[c]})}{\partial \theta_{[c]}} \frac{\partial \theta_{[c]}}{\partial \theta}. \quad (3)$$

The β denotes the learning rate of the finding phase. Note that the meta-optimization is performed over the model parameters θ , whereas the objective $\mathcal{L}_b(\theta_{[c]})$ is computed using the updated model parameters $\theta_{[c]}$.

Training procedure. For each iteration, we sample a batch of training tasks from different marketing campaigns. Then, we update the general model $f(\cdot; \theta)$ with all training tasks together with Equation (3). It is obvious that the setting of the training tasks is similar to the audience expansion task, and it only updates the general model with the expanded audiences $\mathcal{D}_{[c]}^b$ after understanding the seed users $\mathcal{D}_{[c]}^a$. Thus, the general model is able to learn how to expand audiences. We come to the overall training algorithm from a meta-learning perspective, see Algorithm 1.

3.3 Hybrid Experts

To select potential audiences, it is necessary to generate high-quality user representations that can represent users’ intentions toward different campaigns. Intuitively, we can employ an expert (a feed-forward network) to extract the users’ representations for all campaigns. However, it faces a serious problem with a single expert. Usually, an expert only specializes in one area. The user’s representation extracted by the single expert could only contain partial information about the user, which cannot completely cover the characteristics of the user. As a result, while other characteristics of the user have been not observed by the expert, such a user representation is only good for a part of campaigns. To address the challenge, it is intuitive to employ multiple experts. Formally, we average the user’s representation extracted by all experts, as:

$$r = \frac{1}{n} \sum_{i=1}^n h_i(u), \quad (4)$$

Figure 2: The model structure of the Meta Hybrid Experts and Critics (MetaHeac).

where $h_i(\cdot)$ denotes the i -th expert, and n denotes the number of experts. However, the operation of average can remove the task-specific information, so the user representation r may not be good for the specific marketing campaign tasks. Note that different experts specialize in different areas, and they are good at handling different tasks. Therefore, for a certain marketing campaign, we want to select experts who specialize in the specific task.

Along this line, we propose a task-driven gate with campaign features as input to guide the selection process. In addition, note that an expert only specializes in a part of tasks, and his ability to handle different users is also different. Thus, the task-driven gate should take not only campaign features as input but also the user features. Thus, the task-driven gate is defined as:

$$w^{expert} = \text{softmax}(g(c, \mathcal{G}(u))), \quad (5)$$

where $g(\cdot)$ is a feed-forward network, and c and u are the feature embeddings of the campaign task c and the user u . Besides, softmax denotes the softmax function to normalize the output of $g(\cdot)$, and $w^{expert} \in \mathbb{R}^n$ is the weight vector which can denote the importance of different experts. $\mathcal{G}(\cdot)$ denotes an aggregating function. In this paper, we utilize mean pooling as $\mathcal{G}(\cdot)$. The input of $g(\cdot)$ is the embeddings of campaign features and the mean embedding of user features. Note that the mean pooling operation on user features is extremely important. The main reason is that we mainly focus on capturing the transferable knowledge across various tasks, not users. Thus, such a mean pooling operation can not only make the task information dominate the gate, but also exploit the user information. With the task-driven gate, we can reformulate the user’s representation as:

$$r = \frac{1}{n} \sum_{i=1}^n w_i^{expert} h_i(u), \quad (6)$$

where w_i^{expert} denotes the i -th element in the vector w^{expert} .

Each expert can specialize in different tasks, and the task-driven gate can select a subset of experts who are good at handling the specific task. This is desirable for a flexible parameter sharing in the

audience expansion problem which faces hundreds of marketing campaign tasks. It is easy to understand the hybrid experts are able to model the task relationships. If the tasks are highly related, the task-driven gate will select similar experts to extract user representations. If the tasks are less related, the gate networks will learn to utilize different experts instead.

3.4 Hybrid Critics

Given the user representation $\mathbf{r}$ extracted by hybrid experts, we need to decide whether this user is interested in the contents promoted by the certain campaign. There are two direct ways: 1) a single critic; 2) for each task, employ a critic. However, both of the two methods cannot achieve satisfying performance. For the first method, since a user could have different labels in different tasks, a single critic is hard to satisfy all tasks. For the second method, since some tasks only have limited seed users, it is difficult to train a specific critic for such tasks. Based on the intuition that different critics are good at different tasks, we employ multiple critics. And a single critic is defined as:

$$\hat{p} = \sigma(t(\mathbf{r})), \quad (7)$$

where σ denotes the sigmoid function, and $t(\cdot)$ is a feed-forward network. Besides, $\hat{p}$ is the probability predicted by the critic. Based on the similar motivation as hybrid experts, we utilize another task-driven gate to guide the combination of multiple critics:

$$\hat{p} = \frac{1}{m} \sum_{i=1}^m w_i^{critic} \sigma(t_i(\mathbf{r})), \quad (8)$$

where $t_i(\cdot)$ denotes the i -th critic, and m is the number of critics. w_i^{critic} denotes the weight of the i -th critic, which is similar to w_i^{expert} in Equation (6). Note that the task-driven gates for hybrid experts and hybrid critics are two separate modules.

It is obvious that the hybrid critics can also learn the relationships among various tasks as hybrid experts do.

3.5 Overall Framework

The overall MetaHeac framework consists of the meta-learning framework, hybrid experts, and hybrid critics, as shown in Figure 2. Training the general look-alike model $f(\cdot)$ with the meta-learning framework, the model can learn how to expand audiences. In this paper, the model $f(\cdot)$ contains both hybrid experts and hybrid critics, formulated as:

$$f(\mathbf{c}, \mathbf{u}) = \frac{1}{m} \sum_{i=1}^m w_i^{critic} \sigma(t_i(\frac{1}{n} \sum_{j=1}^n w_j^{expert} h_j(\mathbf{u}))). \quad (9)$$

With such hybrid experts and critics, the look-alike model can capture the relationships among various tasks and learn transferable knowledge. The deployment of the MetaHeac framework consists of two stages: offline and online.

Offline stage. In this stage, we learn the general look-alike model with the meta-learning framework of the two-phases simulation on all existing marketing campaigns. All existing campaign tasks contain a set of seed users and the online feedback, and we combine all of them into a single training set to train the general model as Section 3.2.

Online stage. For a new marketing campaign c , given the dataset $\mathcal{D}_{[c]}$ of the new campaign, we fine-tune the general look-alike

Table 1: Statistics of the experimental datasets. The ‘tasks’ denotes the number of marketing campaigns.

	tasks	seeds	expanded audiences
Tencent Dataset	172	421,961	2,265,989
WeChat Dataset	90	286,036	3,146,396

model $f(\cdot; \theta)$ with the cross-entropy loss similar to Equation (2). And we can obtain a customized look-alike model $f_{[c]}(\cdot; \theta_{[c]})$. Then, the customized model can be directly exploited to find the potential audiences $\mathcal{U}_{[c]}$ from the whole candidate pool $\mathcal{U}$ who could be interested in the contents promoted by the specific campaign c . Finally, the look-alike systems can promote the contents to the expanded users.

4 EXPERIMENTS

In this section, we conduct extensive offline and online experiments with the aim of answering the following research questions: **RQ1:** Does our proposed MetaHeac outperform other look-alike approaches in different tasks? **RQ2:** Does this MetaHeac framework get improvement on the performance of WeChat Look-alike system? **RQ3:** What are the effects of meta-learning, hybrid experts, and hybrid critics in our proposed MetaHeac?

4.1 Experimental Settings

Datasets. We evaluate MetaHeac with baselines on both large-scale public and private datasets of audience expansion tasks.

*Tencent Look-alike Dataset*¹: The public dataset for Tencent Ads competitions in 2018 contains hundreds of seed sets and aims to find similar users. Each advertising task contains lots of seed users and non-seed users. Each advertisement has six categorical features: advertiser ID, ad category, campaign ID, product ID, product type, creative size. Each user has 14 features, including carrier, consumption ability, location, age, education, gender, house, three-level interests, two-level keywords, two-level topics.

*WeChat Look-alike Dataset*²: The private dataset is from WeChat. This is a dataset of content promotion in recommender systems. Each task is a content marketing campaign that consists of several contents (items), lots of seed users, and non-seed users. We take the first few days to train the general model and test the performance on the tasks of the rest.

The details of the two datasets are shown in Table 1. For the Tencent and WeChat datasets, seed:non-seed = positive samples: negative samples = 1:20 and 1:10, respectively. Each campaign c contains a seed set $\mathcal{S}_{[c]}$ and a non-seed set for training, and a set of expanded audiences for testing. The set of expanded audiences consists of actual audiences (positive samples) and other candidate users (negative samples).

For WeChat Look-alike Dataset, we use all data of the first two days (58 tasks) to learn the pre-trained model, and the last day (32 tasks) for the test stage. Moreover, to validate the performance on campaigns with different size, according to the size of seed lists, we divide the campaigns of the test set into two groups: ‘ $|\mathcal{S}_{[c]}| \leq T$ ’

¹<https://algo.qq.com/archive.html?>

²All data are preprocessed via data masking to protect the user’s privacy.

Table 2: Offline results (AUC, P@5% and R@5%) on two look-alike datasets. We report the mean results over five runs. Best results are in boldface, and the second bests results are with underline. ** and * indicates 0.005 and 0.05, paired t-test of MetaHeac vs. the best baselines. Improve denotes relative improvement over the best baseline.

Dataset	Method	Pre-trained		$S_{[c]} \leq T$			$S_{[c]} > T$		
		Emb	Network	AUC	P@5%	R@5%	AUC	P@5%	R@5%
Tencent Look-alike Dataset	LR	-	-	0.5942	0.1015	0.1044	0.6824	0.1910	0.2006
	MLP_one-stage	-	-	0.5928	0.1048	0.1081	0.6910	0.1797	0.1888
	MLP+emb	✓	-	0.6624	0.1881	0.1930	0.7060	0.2118	0.2224
	Pinterest	✓	-	0.6245	0.1635	0.1665	0.6802	0.1687	0.1770
	Hubble	✓	-	0.6797	0.2056	0.2110	0.7085	0.2171	0.2279
	MLP+pre-training	✓	✓	<u>0.7117</u>	<u>0.2325</u>	<u>0.2384</u>	0.7082	0.2136	0.2242
	Shared-Bottom+pre-training	✓	✓	0.6936	0.2198	0.2258	<u>0.7089</u>	0.2144	0.2250
	MMoE+pre-training	✓	✓	0.6977	0.2224	0.2280	<u>0.7088</u>	<u>0.2150</u>	<u>0.2257</u>
	MetaHeac	✓	✓	0.7239**	0.2489**	0.2554**	0.7142**	0.2244**	0.2356**
	Improve			1.7%	7.0%	7.1%	0.8%	4.7%	4.7%
WeChat Look-alike Dataset	LR	-	-	0.5654	0.1351	0.0742	0.6711	0.2166	0.1182
	MLP_one-stage	-	-	0.6663	0.2477	0.1363	0.6970	0.2605	0.1419
	MLP+emb	✓	-	0.7143	0.3058	0.1684	0.7217	0.2988	0.1628
	Pinterest	✓	-	0.6289	0.1947	0.1066	0.7044	0.2639	0.1439
	Hubble	✓	-	0.7391	<u>0.3524</u>	<u>0.1936</u>	0.7243	<u>0.3062</u>	0.1668
	MLP+pre-training	✓	✓	<u>0.7440</u>	0.3473	0.1908	0.7272	0.3030	0.1673
	Shared-Bottom+pre-training	✓	✓	0.7271	0.3093	0.1700	0.7275	0.3052	0.1663
	MMoE+pre-training	✓	✓	0.7368	0.3265	0.1797	<u>0.7292</u>	0.3051	<u>0.1675</u>
	MetaHeac	✓	✓	0.7607**	0.3839**	0.2110**	0.7323*	0.3133*	0.1707*
	Improve			2.3%	8.9%	9.0%	0.4%	2.3%	1.9%

and ' $S_{[c]} > T$ '. We set $T = 10,000$ for WeChat Dataset. For the Tencent Look-alike Dataset, we randomly select 60% campaigns as the training tasks, and others as the test tasks. We set $T = 4,000$ for Tencent Dataset. About the selection of T , we keep the campaign tasks of ' $S_{[c]} \leq T$ ': ' $S_{[c]} > T$ ' = 8:2.

Baselines. We categorize our baselines into three groups, and the first group is the traditional one-stage look-alike methods which directly train a model for the specific campaign without pre-training. (1) LR [23]: a one-stage model-based look-alike method taking LR as the classifier. (2) MLP_one-stage: an one-stage model-based method using MLP as the classifier.

The second group is the two-stage approaches which only pre-train the embeddings. (1) MLP+emb: all tasks share the same embedding layer, and a customized MLP is utilized for classification. (2) Pinterest [2]: a two-stage approach presented by Pinterest. In the offline stage, a global embedding model is trained across all tasks. In the online stage, a similarity-based classification is performed on the pre-trained embeddings. (3) Hubble [35]: a GNN-based two-stage method proposed by Ant Financial. In the offline stage, a GNN is trained on the large campaign-user graph to learn the user representations. In the online stage, based on the user representations, a customized MLP is learned for a certain campaign.

The third group is the two-stage methods that pre-train not only the embeddings but also the networks. (1) MLP+pre-training: learn a general MLP (including embeddings) across various tasks in the offline stage, and fine-tune it for a specific campaign in the online stage. (2) Shared-Bottom+pre-training [1]: Shared-Bottom is a traditional multi-task model. We modify it to fit the audience expansion problem. The bottom shared by all tasks is pre-trained in

the offline stage, and fine-tune the classification layer in the online stage. (3) MMoE+pre-training [15]: MMoE is a recent popular multi-task model that shares a mixture-of-experts (MoE) across various tasks, and each task has its specific head. We pre-train the overall structure in the offline stage, and fine-tune it in the online stage for the new campaign.

Experimental details. For all methods, the initial learning rate for the Adam [9] optimizer are tuned by grid searches within {0.001, 0.005, 0.01, 0.02, 0.1}, including the global learning rate β in MetaHeac. Following [14, 32], we set the local learning rate α as 0.001. In addition, we set the dimension of embeddings as 16. For all methods, we set mini-batch size of 512. We tune the number of experts and critics in [3, 10]. For all MLP in these methods, the ReLU activation function is employed, and the dimension of each layer is set to 64, 64, and 1. For Shared-Bottom, MMoE, and MetaHeac, for a fair comparison, the experts and bottom are two-layer MLP, and the dimension of each layer is 64, 64. Besides, the heads or classifiers of the three methods are one fully-connected layer that outputs one unit. The gate network of MMoE and MetaHeac are two-layer MLP with dimension of each layer 64 and 1. For the two-stage methods, we train the model with the training set for one epoch in the offline stage, and we fine-tune the pre-trained model on $\mathcal{D}_{[c]}$ until loss converges in the online stage. For the one-stage methods, directly train the model for each campaign until loss converges.

Evaluation metrics. In binary prediction tasks, AUC (Area Under ROC) is a widely used metric [5] (to predict whether a user is actual audience or whether a conversion interaction exists between a user and a campaign, respectively). In addition, following [13, 35],

Table 3: Online A/B testing results.

Scenarios	Exposure	Conversion	CVR
video	+3.07%	+10.18%	+7.90%
advertisement	+0.65%	+15.50%	+15.40%
article	+3.18%	+9.23%	+4.64%

we also choose another two metrics, Precision/Recall at top K percent candidates (P@K% and R@K% for short), defined as:

$$P@K\% = \frac{|\mathcal{U}_{at} \cap \mathcal{U}_{cdd,K}|}{|\mathcal{U}_{cdd,K}|}, \quad R@K\% = \frac{|\mathcal{U}_{at} \cap \mathcal{U}_{cdd,K}|}{|\mathcal{U}_{at}|}. \quad (10)$$

$\mathcal{U}_{at}$ denotes the set of actual audiences of the certain campaign, and $\mathcal{U}_{cdd,K}$ denotes the set of K% candidate users with a higher score predicted by the audience expansion model w.r.t. the certain campaign. In this paper, following [13], we set $K = 5$.

4.2 Offline Results (RQ1)

In this section, we conduct offline experiments on the two offline look-alike datasets. By treating the actual audiences as positive examples while other candidate users as negative examples, we can evaluate AUC, P@K%, and R@K%. From the offline results shown in Table 2, we have several findings.

The two-stage methods which pre-train embeddings outperform the one-stage methods (LR, MLP_one-stage) which train customized models from scratch. It demonstrates that the embedding pre-training is effective to improve the performance of look-alike modeling. Then, we can find that the third group also outperforms the second group, which validates that pre-training networks are also important.

Another observation is that the gaps between different methods in the $\mathcal{S}_{[c]} \leq T$ scenarios are larger than in the $\mathcal{S}_{[c]} > T$ scenarios. For the $\mathcal{S}_{[c]} > T$ scenarios, each test campaign has sufficient seed users, and the model can learn relatively sufficient information from the seed set. Thus, the transferable knowledge from other campaigns can only slightly improve the performance. In contrast, for the $\mathcal{S}_{[c]} \leq T$ scenarios, each campaign only has limited seeds, and it is hard to learn a satisfying customized model with the limited seeds. So the transferable knowledge can largely improve the performance of these campaigns.

With the results of the t-test, we can find that MetaHeac could outperform the best baseline significantly in all tasks, which demonstrates that MetaHeac is an effective solution to expand audiences. The main reason is that MetaHeac not only learns how to expand audiences but also takes relationships among various tasks into consideration. MetaHeac outperforms the multi-task methods (Shared-Bottom and MME), which demonstrates that the way to learn the relationships among various tasks for the audience expansion problem should be different from traditional multi-task learning. Comparing with the SOTA look-alike methods (Pinterest and Hubble), MetaHeac achieves better results, which validates that pre-training a general model is important for learning customized models for each campaign.

4.3 Online A/B Testing (RQ2)

The MetaHeac framework has already been deployed in WeChat Look-alike System to serve a billion users and conduct hundreds of

Table 4: Ablation Study on Tencent Look-alike Dataset. HC denotes Hybrid Critics, and HE denotes Hybrid Experts.

Method	$\mathcal{S}_{[c]} \leq T$		$\mathcal{S}_{[c]} > T$	
	AUC	P@5%	AUC	P@5%
MetaHeac w/o HC	0.7199	0.2472	0.7115	0.2220
MetaHeac w/o HE	0.7181	0.2419	0.7112	0.2193
MetaHeac w/o Meta	0.7173	0.2431	0.7107	0.2180
MetaHeac	0.7239	0.2489	0.7142	0.2244

campaigns. Three applications have utilized the MetaHeac framework to expand audiences, including promotion of video, advertisement, and article in WeChat. 'video' and 'article' are content recommender systems that can conduct content marketing campaigns to promote the specific contents, while 'advertisement' conducts advertising campaigns to promote advertisements. To verify the real benefits MetaHeac brings to our system, we conduct online A/B testing experiments for the three applications within a week. For each application, we randomly split all candidates into two buckets with the same size. The control group here is an MLP_one-stage classifier as mentioned in subsection 4.1, while the treatment group is the proposed MetaHeac framework. We allow them to select the same amount of top-scored users from their buckets, respectively. To evaluate the performance online, the following metrics are involved: 1) Exposure: the number of users who have been exposed by the campaign, 2) Conversion: the amount of users who have converted in the campaign, 3) CVR: the conversion rate of the campaign. For the content marketing campaigns, CVR denotes the rate of sharing. The audience expansion task aims to find users who are most likely to finish the business goal of the target campaign. Thus, a higher CVR indicates a higher quality of selected audiences. From the online results shown in Table 3, we can see the proposed MetaHeac Framework achieves a significant improvement on Exposure, Conversion, and CVR, against the MLP+pre-training classifier. Especially, the proposed MetaHeac can increase 7.90%, 15.40% and 4.64% CVR for the three scenarios, respectively, which is a very significant improvement and will result in great economic benefit to the company. Thus, our MetaHeac is effective to improve not only content marketing campaigns but also advertising campaigns.

4.4 Analysis (RQ3)

In this section, we analyze the effects of meta-learning, hybrid experts, and hybrid critics in our proposed MetaHeac. First, we conduct an ablation study to verify the effectiveness of each module. Second, we present convergence analysis to testify the meta-learning framework can converge better and achieve better performance. Third, we visualize the average representations of seeds and the gate of the hybrid critics to show that the hybrid experts and critics can capture the relationships among different tasks.

Ablation Study. To test the effectiveness of each module in MetaHeac, we present an ablation study. Namely the Hybrid Critics introduced in Section 3.4 (HC), Hybrid Experts introduced in Section 3.3 (HE) and the Meta-learning Framework introduced in Section 3.2 (Meta). First, we replace the hybrid critics (HC) with a single critic (classifier), which is denoted as MetaHeac w/o HC. Second, we replace the hybrid experts (HE) with a single expert denoted

Figure 3: (a) Training Loss and (b) AUC w.r.t. the number of iterations on WeChat Look-alike Dataset. The Visualization of (c) average representations extracted by hybrid experts, and (d) Gate of Hybrid Critics.

as MetaHeac w/o HE. Third, we replace the training framework of meta-learning with standard BP, which is denoted as MetaHeac w/o Meta. The results are shown in Table 4, and we have several findings: 1) MetaHeac outperforms MetaHeac w/o HC, which demonstrates that the hybrid critics are effective to improve the performance. 2) MetaHeac outperforms MetaHeac w/o HE, demonstrating that the hybrid experts are useful to capture multiple representations of users. 3) MetaHeac achieves better results than MetaHeac w/o Meta, which demonstrates training the general look-alike model from the meta-learning perspective could learn how to expand audiences.

Convergence. We testify the convergence of MLP_one-stage (denoted as MLP), Hubble, MMoE, and MetaHeac, with the training loss and AUC on a random test task in WeChat Look-alike Dataset. We randomly select a task from the test set, and fine-tune the pre-trained model with its training data of seeds and non-seeds (Hubble, MMoE, MetaHeac) or train a new model from scratch (MLP). The training loss and AUC are shown in Figure 3(a) and (b), respectively.

From Figure 3(a), we can find that, among these Look-alike methods, MetaHeac achieves the lowest training loss, which testifies MetaHeac has better convergence. Besides, note that MetaHeac has a smaller initial loss, which demonstrates that the meta-learning framework can learn a better initialization. From Figure 3(b), compared with Hubble and MMoE which learn pre-trained models, MetaHeac achieves a better performance, which demonstrates that MetaHeac can learn more general knowledge from existing tasks. We can find MLP has an extremely unsatisfying performance, and the reason could be overfitting on the training data. The above findings provide empirical evidence that our MetaHeac can learn more general knowledge (how to expand audience), and achieve better convergence and performance.

Visualization. To understand how hybrid experts and critics capture the relationships among various tasks, we visualize average representations of seeds and the distribution of the softmax gate of each campaign task. First, with hybrid experts, we extract seed users' representations and compute the average representation for each campaign task, and then we employ the t-SNE [4] to visualize the average users' representations of all tasks, as shown in Figure 3(c). We can find average representations of the campaigns within the same category are roughly clustering. Note that the campaigns within different categories have no classification boundary.

For example, the contents about NBA belong to the category of sports, but it is also related to entertainment. Thus, the campaigns within different categories are also overlapping. We show the average distribution of the softmax hybrid critics' gate of some tasks in Figure 3(d). Each row denotes the gating weight of a critic, and each column denotes a campaign task. It is easy to observe that the campaign tasks within the same category have a similar distribution of gate. In summary, the above findings demonstrate that the hybrid experts and critics can capture the relationships among various tasks.

5 RELATED WORK

In this section, we will introduce the related work from three aspects: Audience Expansion, Meta Learning, Multi-task Learning.

Audience Expansion: Audience expansion (Look-alike modeling) which aims to find similar audiences has been used by many companies. Look-alike approaches can be categorized into three clusters: rule-based, similarity-based, and model-based.

The rule-based methods [18, 24] match similar users with specific demographic tags (age, gender, geography) or interests that are targeted by marketers. The core technical support in the background is user profile mining, which means that the interest tags are inferred from the user behaviour [21]. However, these methods only capture the high-level feature, which is hard to achieve satisfying performance.

The similarity-based methods [3, 11, 16, 17] expand a given seed-set via calculating the similarity of all pairs between seed users and candidate users. The similarity function is usually pre-defined objectively, e.g., Cosine, Jaccard, vector-based dot product. The disadvantage of these methods is that the performance largely depends on artificially defined similarity functions.

The model-based methods [7, 23] train customized prediction models for each campaign. Our method also falls into this stream. Qu et al. [23] utilized LR to expand audiences, which is effective. In terms of audience expansion performance, [7] studied the impact of sampling ratios and sampling techniques. These one-stage methods train customized models from scratch for each campaign which is time-consuming and suffers from the overfitting problem. Recently, some two-stage approaches [2, 13, 35] have been proposed to pre-train the embedding with data of all campaigns. However, these

methods ignore the generalization ability, task relationships, and deep networks. In contrast, MetaHeac can solve these issues better.

Meta Learning: Also known as learning to learn, meta-learning intends to learn the general knowledge across similar learning tasks, so as to rapidly adapt to new tasks [6]. The works in this stream can be grouped into three clusters, i.e., metric-based [25], optimization-based [6], and parameter-generating approaches [19]. Meta-learning leads to an interest in many areas, such as recommendation [14, 20, 32, 34], natural language processing [30], and computer vision [26]. In meta-learning literature, the audience expansion problem relates to general model [10] and fast adaptation [6]. Since audience expansion is a very practical problem that involves how to update an online system, none of these methods can be directly applied to the online look-alike system.

Multi-task Learning: Multi-task models can jointly learn from several tasks so that it can result in improved efficiency for each task [15, 22, 29]. Moreover, multi-task learning is a promising method to learn relationships among different tasks [15, 31]. However, existing multi-task approaches usually serve for scenarios with fewer than five tasks. In practice, the audience expansion problem faces hundreds of marketing campaign tasks every day, so these multi-task methods cannot work well for this problem.

6 CONSLUSION

In this paper, to solve the audience expansion problem in recommender systems and advertising platforms, we proposed a novel two-stage framework named Meta Hybrid Experts and Critics (MetaHeac). In the offline stage, a general model which can capture the relationships among various tasks is trained from a meta-learning perspective with existing campaigns. In the online stage, by fine-tuning the general model with given seeds, it is easy for a certain campaign to obtain a customized model to find potential audiences who are likely to finish the business goal. We demonstrated the superior effectiveness of the proposed MetaHeac in both offline experiments and online A/B testing, compared to other state-of-the-art approaches. Moreover, the MetaHeac framework has been deployed in WeChat for the promotion of both contents and advertisements, which has a significant improvement.

ACKNOWLEDGMENTS

The research work is supported by the National Key Research and Development Program of China under Grant No. 2017YFC0820604, the National Natural Science Foundation of China under Grant No. 61773361, U1836206, U1811461.

REFERENCES

- [1] Rich Caruana. 1997. Multitask learning. *Machine learning* 28, 1 (1997), 41–75.
- [2] Stephanie deWet and Jiafan Ou. 2019. Finding Users Who Act Alike: Transfer Learning for Expanding Advertiser Audiences. In *KDD*. 2251–2259.
- [3] Khoa D Doan, Pranjul Yadav, and Chandan K Reddy. 2019. Adversarial factorization autoencoder for look-alike modeling. In *CIKM*. 2803–2812.
- [4] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. 2014. Decaf: A deep convolutional activation feature for generic visual recognition. In *ICML*. 647–655.
- [5] Tom Fawcett. 2006. An introduction to ROC analysis. *Pattern recognition letters* 27, 8 (2006), 861–874.
- [6] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*. JMLR. org, 1126–1135.
- [7] Jinling Jiang, Xiaoming Lin, Junjie Yao, and Hua Lu. 2019. Comprehensive Audience Expansion based on End-to-End Neural Prediction. In *eCOM@ SIGIR*.
- [8] Bhargav Kanagal, Amr Ahmed, Sandeep Pandey, Vanja Josifovski, Lluís Garcia-Pueyo, and Jeff Yuan. 2013. Focused matrix factorization for audience selection in display advertising. In *ICDE*. IEEE, 386–397.
- [9] Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR*, Vol. 5.
- [10] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. 2018. Learning to generalize: Meta-learning for domain generalization. In *AAAI*, Vol. 32.
- [11] Haishan Liu, David Pardoe, Kun Liu, Manoj Thakur, Frank Cao, and Chongzhe Li. 2016. Audience expansion for online social network advertising. In *KDD*. 165–174.
- [12] Yudan Liu, Kaikai Ge, Xu Zhang, and Leyu Lin. 2019. Real-time Attention Based Look-alike Model for Recommender System. In *KDD*. 2765–2773.
- [13] Zhining Liu, Xiao-Fan Niu, Chenyi Zhuang, Yize Tan, Yixiang Mu, Jinjie Gu, and Guannan Zhang. 2020. Two-Stage Audience Expansion for Financial Targeting in Marketing. In *CIKM*. 2629–2636.
- [14] Yuanfu Lu, Yuan Fang, and Chuan Shi. 2020. Meta-learning on heterogeneous information networks for cold-start recommendation. In *KDD*. ACM.
- [15] Jiaqi Ma, Zhe Zhao, Xinyang Yi, Jilin Chen, Lichan Hong, and Ed H Chi. 2018. Modeling task relationships in multi-task learning with multi-gate mixture-of-experts. In *KDD*. 1930–1939.
- [16] Qiang Ma, Eeshan Wagh, Jiayi Wen, Zhen Xia, Robert Ormandi, and Datong Chen. 2016. Score Look-Alike Audiences. In *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW)*. IEEE, 647–654.
- [17] Qiang Ma, Musen Wen, Zhen Xia, and Datong Chen. 2016. A sub-linear massive-scale look-alike audience extension system. In *Proceedings of the 5th International Workshop on Big Data, Streams and Heterogeneous Source Mining: Algorithms, Systems, Programming Models and Applications*.
- [18] Ashish Mangalampalli, Adwait Ratnaparkhi, Andrew O Hatch, Abraham Bagherjeiran, Rajesh Parekh, and Vikram Pudi. 2011. A feature-pair-based associative classification approach to look-alike modeling for conversion-oriented user-targeting in tail campaigns. In *WWW*. 85–86.
- [19] Tsendsuren Munkhdalai and Hong Yu. 2017. Meta networks. In *ICML*. 2554–2563.
- [20] Feiyang Pan, Shuokai Li, Xiang Ao, Pingzhong Tang, and Qing He. 2019. Warm up cold-start advertisements: Improving ctr predictions via learning to learn id embeddings. In *SIGIR*. 695–704.
- [21] Sandeep Pandey, Mohamed Aly, Abraham Bagherjeiran, Andrew Hatch, Peter Ciccolo, Adwait Ratnaparkhi, and Martin Zinkevich. 2011. Learning to target: what works for behavioral targeting. In *CIKM*. 1805–1814.
- [22] Zhen Qin, Yicheng Cheng, Zhe Zhao, Zhe Chen, Donald Metzler, and Jingzheng Qin. 2020. Multitask Mixture of Sequential Experts for User Activity Streams. In *KDD*. 3083–3091.
- [23] Yan Qu, Jing Wang, Yang Sun, and Hans Marius Holtan. 2014. Systems and methods for generating expanded user segments. US Patent 8,655,695.
- [24] Jianqiang Shen, Sahin Cem Geyik, and Ali Dasdan. 2015. Effective audience extension in online advertising. In *KDD*. 2099–2108.
- [25] Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. In *NeurIPS*. 4077–4087.
- [26] Jae Woong Soh, Sunwoo Cho, and Nam Ik Cho. 2020. Meta-Transfer Learning for Zero-Shot Super-Resolution. In *CVPR*. 3516–3525.
- [27] Qianru Sun, Yaoyao Liu, Tat-Seng Chua, and Bernt Schiele. 2019. Meta-transfer learning for few-shot learning. In *CVPR*. 403–412.
- [28] Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, Wenjun Zeng, and Tao Qin. 2021. Generalizing to Unseen Domains: A Survey on Domain Generalization. In *IJCAI*.
- [29] Dongbo Xi, Zhen Chen, Peng Yan, Yinger Zhang, Yongchun Zhu, Fuzhen Zhuang, and Yu Chen. 2021. Modeling the Sequential Dependence among Audience Multi-step Conversions with Multi-task Learning in Targeted Display Advertising. *arXiv preprint arXiv:2105.08489* (2021).
- [30] Ming Yan, Hao Zhang, Di Jin, and Joey Tianyi Zhou. 2020. Multi-source Meta Transfer for Low Resource Multiple-Choice Question Answering. In *ACL*. 7331–7341.
- [31] Jiejie Zhao, Bowen Du, Leilei Sun, Fuzhen Zhuang, Weifeng Lv, and Hui Xiong. 2019. Multiple Relational Attention Network for Multi-task Learning. In *KDD*. 1123–1131.
- [32] Yongchun Zhu, Kaikai Ge, Fuzhen Zhuang, Ruobing Xie, Dongbo Xi, Xu Zhang, Leyu Lin, and Qing He. 2021. Transfer-Meta Framework for Cross-domain Recommendation to Cold-Start Users. In *SIGIR*.
- [33] Yongchun Zhu, Dongbo Xi, Bowen Song, Fuzhen Zhuang, Shuai Chen, Xi Gu, and Qing He. 2020. Modeling Users’ Behavior Sequences with Hierarchical Explainable Network for Cross-domain Fraud Detection. In *WWW*. 928–938.
- [34] Yongchun Zhu, Ruobing Xie, Fuzhen Zhuang, Kaikai Ge, Ying Sun, Xu Zhang, Leyu Lin, and Juan Cao. 2021. Learning to Warm Up Cold Item Embeddings for Cold-start Recommendation with Meta Scaling and Shifting Networks. In *SIGIR*.
- [35] Chenyi Zhuang, Ziqi Liu, Zhiqiang Zhang, Yize Tan, Zhengwei Wu, Zhining Liu, Jianping Wei, Jinjie Gu, Guannan Zhang, Jun Zhou, et al. 2020. Hubble: an Industrial System for Audience Expansion in Mobile Marketing. In *KDD*. 2455–2463.