



ColdNAS: Search to Modulate for User Cold-Start Recommendation

Shiguang Wu
Department of Electronic Engineering
Tsinghua University
Beijing, China
wsg19@mails.tsinghua.edu.cn

Yaqing Wang*
Baidu Inc.
Beijing, China
wangyaqing01@baidu.com

Qinghe Jing
Baidu Inc.
Beijing, China
jingqinghe@baidu.com

Daxiang Dong
Baidu Inc.
Beijing, China
dongdaxiang@baidu.com

Dejing Dou
Baidu Inc.
Beijing, China
doudejing@baidu.com

Quanming Yao
Department of Electronic Engineering
Tsinghua University
Beijing, China
qyaoaa@tsinghua.edu.cn

ABSTRACT

Making personalized recommendation for cold-start users, who only have a few interaction histories, is a challenging problem in recommendation systems. Recent works leverage hypernetworks to directly map user interaction histories to user-specific parameters, which are then used to modulate predictor by feature-wise linear modulation function. These works obtain the state-of-the-art performance. However, the physical meaning of scaling and shifting in recommendation data is unclear. Instead of using a fixed modulation function and deciding modulation position by expertise, we propose a modulation framework called ColdNAS for user cold-start problem, where we look for proper modulation structure, including function and position, via neural architecture search. We design a search space which covers broad models and theoretically prove that this search space can be transformed to a much smaller space, enabling an efficient and robust one-shot search algorithm. Extensive experimental results on benchmark datasets show that ColdNAS consistently performs the best. We observe that different modulation functions lead to the best performance on different datasets, which validates the necessity of designing a searching-based method. Codes are available at <https://github.com/LARS-research/ColdNAS>.

CCS CONCEPTS

• Information systems → Recommender systems; • Computing methodologies → Supervised learning; Neural networks.

KEYWORDS

User-Cold Start Recommendation, Neural Architecture Search, Few-Shot Learning, Meta-Learning, Hypernetworks

*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 the author(s) 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](https://permissions.acm.org).

WWW '23, April 30–May 04, 2023, Austin, TX, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9416-1/23/04...\$15.00

<https://doi.org/10.1145/3543507.3583344>

ACM Reference Format:

Shiguang Wu, Yaqing Wang, Qinghe Jing, Daxiang Dong, Dejing Dou, and Quanming Yao. 2023. ColdNAS: Search to Modulate for User Cold-Start Recommendation. In *Proceedings of the ACM Web Conference 2023 (WWW '23)*, April 30–May 04, 2023, Austin, TX, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3543507.3583344>

1 INTRODUCTION

Recommendation systems (RSs) [32] target at providing suggestions of items that are most pertinent to a particular user, such as movie recommendation [12] and book recommendation [47]. Nowadays, RSs are abundant online, offering enormous users convenient ways to shopping regardless of location and time, and also providing intimate suggestions according to their preferences. However, user cold-start recommendation [33] remains a severe problem in RSs. On the one hand, the users in RSs follow long tail effect [27], some users just have a few interaction histories. On the other hand, new users are continuously emerging, who naturally have rated a few items in RSs. Such a problem is even more challenging as modern RSs are mostly built with over-parameterized deep networks, which needs a huge amount of training samples to get good performance and can easily overfit for cold-start users [36].

User cold-start recommendation problem can naturally be modeled as a few-shot learning problem [38], which targets at quickly generalize to new tasks (i.e. personalized recommendation for cold-start users) with a few training samples (i.e. a few interaction histories). A number of works [6, 19, 24, 37, 44] adopt the classic gradient-based meta-learning strategy called model-agnostic meta-learning (MAML) [8], which learns a good initialized parameter from a set of tasks and adapts it to a new task by taking a few steps of gradient descent updates on a limited number of labeled samples. This line of models has demonstrated high potential of alleviating user cold-start problem. However, gradient-based meta-learning strategy require expertise to tune the optimization procedure to avoid over-fitting. Besides, the inference time can be long.

Instead of adapting to each user by fine-tuning via gradient descent, another line of works uses hypernetworks [11] to directly map user interaction history to user-specific parameters [6, 7, 21, 26]. These modulation-based methods consist of embedding layer,




Figure 1: Architecture used in user cold-start models, which consists of embedding layer, adaptation network and predictor. Modulation structure in red dotted lines are the searching target of our paper.

adaptation network and predictor. The adaptation network generates user-specific parameters, which are used to modulate the predictor in the form of a modulation function. Particularly, they all adopt feature-wise linear modulation function (FiLM) [28], which modulates the representation via scaling and shifting based on the conditioning information, to modulate the user cold-start models to obtain user-specific representation. Although FiLM has been proved to be highly effective in on images [31] and graphs such molecules and protein-protein interaction graphs [4], applying scaling and shifting on user interaction history has rather obscure physical meaning. Moreover, choosing where to modulate is hard to decide. Existing works modulate different parts of the model, such as last layers of the decoder [21] and most layers in the model [26]. How to modulate well for different users, and how to choose the right functions at the right positions to modulate, are still open questions.

In this paper, we propose ColdNAS to find appropriate modulation structure for user cold-start problem by neural architecture search (NAS). Although NAS methods have been applied in RSs [9, 39], the design of NAS methods are problem-specific. For user cold-start problem, it is still unknown how to (i) design a search space that can cover effective cold-start models with good performance for various datasets, and (ii) design an efficient and robust search algorithm. To solve the above challenges, we design a search space of modulation structure, which can cover not only existing modulation-based user cold-start models, but also contain more flexible structures. We theoretically prove that the proposed search space can be transformed to an equivalent space, where we search efficiently and robustly by differentiable architecture search. Our main contributions are summarized as follows:

- We propose ColdNAS, a modulation framework for user cold-start problem. We use a hypernetwork to map each user’s history interactions to user-specific parameters which are then used to modulate the predictor, and formulate how to modulate and where to modulate as a NAS problem.
- We design a search space of modulation structure, which can cover not only existing modulation-based user cold-start models, but also contain more expressive structures. As this search space can be large to search, we conduct search space transformation to transform the original space to an equivalent but much smaller space. Theoretical analysis is provided to validate its correctness. Upon the transformed space, we then can search efficiently and robustly by differentiable architecture search algorithm.
- We perform extensive experiments on benchmark datasets for user cold-start problem, and observe that ColdNAS consistently obtains the state-of-the-art performance. We also validate the design consideration of search spaces and algorithms, demonstrating the strength and reasonableness of ColdNAS.

2 RELATED WORKS

2.1 User Cold-Start Recommendation

Making personalized recommendation for cold-start users is particularly challenging, as these users only have a few interaction histories [33]. In the past, collaborative filtering (CF)-based [13, 18, 34] methods, which make predictions by capturing interactions among users and items to represent them in low-dimensional space, obtain leading performance in RSs. However, these CF-based methods make inferences only based on the user’s history. They cannot handle user cold-start problem. To alleviate the user cold-start problem, content-based methods leverage user/item features [33] or even user social relations [20] to help to predict for cold-start users. The recent deep model DropoutNet [36] trains a neural network with dropout mechanism applied on input samples and infers the missing data. However, it is hard to generalize these content-based methods to new users, which usually requires model retraining.

A recent trend is to model user cold-start problem as a few-shot learning problem [38]. The resultant models learn the ability to quickly generalize to recommend for new users with a few interaction histories. Most works mainly follow the classic gradient-based meta-learning strategy MAML [8], which first learns a good initialized parameter from training tasks, then locally update the parameter on the provided interaction history by gradient descent. In particular, existing works consider different directions to improve the performance: MeLU [19] selectively adapts model parameters to the new task in the local update stage, MAMO [6] introduces external memory to guide the model to adapt, MetaHIN [24] uses heterogeneous information networks to leverage the rich semantics between users and items, REG-PAML [44] proposes to use user-specific learning rate during local update, and PAML [37] leverages social relations to share information among similar users. While these approaches can adapt models to training data, they are computationally inefficient at test-time, and usually require expertise to tune the optimization procedure to avoid over-fitting.

Another line of few-shot learning methods leverages hypernetworks [11] to directly map training data to the desired parameters. Existing works [7, 21, 26] in this line all choose to adopt the

feature-wise linear modulation function (FiLM) [28]. By design, FiLM modulates the representation via scaling and shifting based on the conditioning information. It has been shown to be highly effective in other domains including images [31] and graphs such molecules and protein-protein interaction graphs [4]. In particular, TaNP [21] learns to map user interaction history to generate user-specific parameters, CMML [7] uses the same set of user-specific parameters across different layers, and PNMTA [26] further leverages pretrained encoder to capture generalized representation. By using FiLM, these methods can obtain user-specific representation, which eliminates the necessity of fine-tuning by gradient descent. However, unlike images and graphs, scaling and shifting operations have unknown physical meaning in user interaction history. In other words, FiLM may not be the ideal modulation function in RSs.

2.2 Neural Architecture Search

Neural architecture search (NAS) targets at finding an architecture with good performance without human tuning [16]. Recently, NAS methods have been applied in RSs. SIF [41] searches for interaction function in collaborative filtering, AutoCF [10] further searches for basic components including input encoding, embedding function, interaction function, and prediction function in collaborative filtering. AutoFIS [22], AutoCTR [35] and FIVES [39] search for effective feature interaction in click-through rate prediction. AutoLoss [46] searches for loss function in RSs. Due to different problem settings, search spaces needs to problem-specific and cannot be shared or transferred. Therefore, none of these works can be applied for user cold-start recommendation problem. To search efficiently on the search space, one can choose reinforcement learning methods [1], evolutionary algorithms [30], and one-shot differentiable architecture search algorithms [23, 29, 42]. Among them, one-shot differentiable architecture search algorithms have demonstrated higher efficiency. Instead of training and evaluating different models like classical methods, they optimize only one supernet where the model parameters are shared across the search space and co-adapted.

3 PROPOSED METHOD

In this section, we present the details of ColdNAS, whose overall architecture is shown in Figure 1. In the sequel, we first provide the formal problem formulation of user cold-start problem (Section 3.1). Then, we present our search space and theoretically show how to transform it (Section 3.2). Finally, we introduce the search algorithm to search for a good user cold-start model (Section 3.3).

3.1 Problem Formulation

Let user set be denoted $\mathcal{U} = \{u_i\}$, where each user u_i is associated with user features. The feature space is shared across all users. Let item set be denoted $\mathcal{V} = \{v_j\}$ where each item v_j is also associated with item features. When a user u_i rates an item v_j , the rating is denoted $y_{i,j}$. In user cold-start recommendation problem, the focus is to make personalized recommendation for user u_i who only has rated a few items.

Following recent works [3, 21, 24], we model the user cold-start recommendation problem as a few-shot learning problem. The target is to learn a model from a set of training user cold-start tasks $\mathcal{T}^{\text{train}}$ and generalize to provide personalized recommendation for

new tasks. Each task T_i corresponds to a user u_i , with a support set $\mathcal{S}_i = \{(v_j, y_{i,j})\}_{j=1}^N$ containing existing interaction histories and a query set $\mathcal{Q}_i = \{(v_j, y_{i,j})\}_{j=1}^M$ containing interactions to predict. N and M are the number of interactions in \mathcal{S}_i . \mathcal{Q}_i and N are small.

3.2 Search Space

Existing modulation-based user cold-start works can be summarized into the following modulation framework consisting of three parts, i.e., embedding layer, adaptation network and predictor, as plotted in Figure 1.

- The embedding layer E with parameter θ_E embeds the categorical features from users and items into dense vectors, i.e., $(\mathbf{u}_i, \mathbf{v}_j) = E(\mathbf{u}_i, \mathbf{v}_j; \theta_E)$.
- The adaptation network A with parameter θ_A takes the support set \mathcal{S}_i for a specific user u_i as input and generates user-specific adaptive parameters, i.e.,

$$\Phi_i = \{\phi_i^k\}_{k=1}^C = A(\mathcal{S}_i; \theta_A), \quad (1)$$

where C is the number of adaptive parameter groups for certain modulation structure.

- The predictor P with parameter θ_P takes user-specific parameters ϕ_i and $v_q \in \mathcal{Q}_i$ from the query set as input, and generate predictions by

$$\hat{y}_{i,j} = P((\mathbf{u}_i, \mathbf{v}_q), \Phi_i; \theta_P). \quad (2)$$

Comparing with classical RS models, the extra adaptation network is introduced to handle cold-start users. To make personalized recommendation, for each u_i , the support set \mathcal{S}_i is first mapped to user-specific parameter ϕ_i by (1). Then taking the features of target item $v_q \in \mathcal{Q}_i$, user features u_i , and the ϕ_i , prediction is made as (2). Subsequently, how to use the user-specific parameter ϕ_i to change the prediction process in (2) can be crucial to the performance. Usually, a multi-layer perception (MLP) is used as P [5, 13, 21]. Assume a L -layer MLP is used and \mathbf{h}^l denotes its output from the l th layer, and let $\mathbf{h}^0 = (\mathbf{u}_i, \mathbf{v}_q)$ for notation simplicity. For the i th user, \mathbf{h}^{l+1} is modulated as

$$\mathbf{h}_i^l = M^l(\mathbf{h}^l, \Phi_i), \quad (3)$$

$$\mathbf{h}^{l+1} = \text{ReLU}(\mathbf{W}_P^l \mathbf{h}_i^l + \mathbf{b}_P^l), \quad (4)$$

where M^l is the modulation function, \mathbf{W}^l and \mathbf{b}^l are learnable weights at the l th layer. This M^l controls how \mathbf{h}^l is personalized w.r.t. the i th user. The recent TaNP [21] directly lets M^l in (3) adopt the form of FiLM for all L MLP layers:

$$\mathbf{h}_i^l = \mathbf{h}^l \odot \phi_i^1 + \phi_i^2. \quad (5)$$

FiLM applies a feature-wise affine transformation on the intermediate features, and has been proved to be highly effective in other domains [4, 31]. However, users and items can have diverse interaction patterns. For example, both the inner product and summation have been used in RS to measure the preference of user over items [15, 18].

In order to find the appropriate modulation function, we are motivated to search M^l for different recommendation tasks. We design the following space for M^l :

$$\mathbf{h}_i^l = \mathbf{h}^l \circ_{\text{op}^1} \phi_i^1 \circ_{\text{op}^2} \phi_i^2 \cdots \circ_{\text{op}^C} \phi_i^C, \quad (6)$$

where op^i are defined as

$$\circ_{\text{op}_i} \in \mathcal{O} \equiv \{\max, \min, \odot, /, +, -\}.$$

They are all commonly used simple dimension-preserving binary operations. We choose them to avoid the resultant M^l being too complex, which can easily overfit for cold-start users.

The search space in (6) can be viewed as a binary tree, as shown in Figure 2(a). Since M^l can be different for each layer, the size of this space is $6^{C \times L}$. A larger C leads to a larger search space, which has higher potential of containing appropriate modulation function but is also more challenging to search effectively.

3.3 Search Strategy

We aim to find an efficient and practicable search algorithm on the proposed search space. However, the search space can be very large, and differentiable NAS methods are known to be fragile on large spaces [43]. In the sequel, we propose to first transform the original search space to an equivalent but much smaller space, as shown in Figure 2(a), where the equivalence is inspired by some similarities between the operations and the expressiveness of deep neural networks. On the transformed space, we design a supernet structure to conduct efficient and robust differentiable search.

3.3.1 Search Space Transformation. Though the aforementioned search space can be very large, we can transform it to an equivalent space of size $2^{4 \times L}$ which is invariant with C , as proved in Proposition 3.1¹ below.

PROPOSITION 3.1 (SEARCH SPACE TRANSFORMATION). *Assume the adaptation network A is expressive enough. Any M^l with a form of (6) where $\circ_{\text{op}^k} \in \mathcal{O}$, C is any non-negative integer, and $\phi_i^k \in \Phi_i = A(\mathcal{S}_i, \theta_A)$, can be represented as*

$$\mathbf{h}_i^l = \min(\max(\mathbf{h}^l, \hat{\phi}_i^1), \hat{\phi}_i^2) \odot \hat{\phi}_i^3 + \hat{\phi}_i^4, \quad (7)$$

and the above four operations are permutation-invariant.

The intuitions are as follows. First, operations in \mathcal{O} can be divided into four groups: $G_1 = \{\max\}$, $G_2 = \{\min\}$, $G_3 = \{+, -\}$, $G_4 = \{\odot, /\}$. Then, with mild assumption on adaptation network, we can prove two important properties: (i) inner-group consistence: operations that in the same group can be associated; and (ii) inter-group permutation-invariance: operations that are not in the same group are permutation-invariant which means any two operations can switch with another. Thus, we can recurrently commute operations until operations in the same group are neighbors, and associate operations in the four groups respectively.

REMARK 1. *To better understand Proposition 3.1, let us check two examples.*

- (1) $\min(\max(\mathbf{h}^l, \phi_i^1) + \phi_i^2 - \phi_i^3, \phi_i^4) \odot \phi_i^5$ equals to (7) where $\hat{\phi}_i^1 = \phi_i^1$, $\hat{\phi}_i^2 = \phi_i^2 - \phi_i^3 + \phi_i^4$, $\hat{\phi}_i^3 = \phi_i^5$, $\hat{\phi}_i^4 = (\phi_i^2 - \phi_i^3) \odot \phi_i^5$; and
- (2) $\max(\min(\mathbf{h}^l + \phi_i^1, \phi_i^2), \phi_i^3) \odot \phi_i^4$ also equals to (7) where $\hat{\phi}_i^1 = \phi_i^3 - \phi_i^1$, $\hat{\phi}_i^2 = \phi_i^2 - \phi_i^1$, $\hat{\phi}_i^3 = \phi_i^4$, $\hat{\phi}_i^4 = \phi_i^1 \odot \phi_i^4$.

Note that due to the universal approximation ability of deep network [14], the assumption in this proposition can be easily

¹The proof is in Appendix B.

satisfied. For example, we implement A based on two-layer MLP in experiments (see Appendix A.2), which can already ensure a good performance (see Section 4.3.3). After the transformation, the space in (7) also spans a permutation-invariant binary tree, as plotted in Figure 2(b). Such a space can be significantly smaller than the original one, as explained in Remark 2 below.

REMARK 2. *The space transformation plays an essential role in ColdNAS, Table 1 helps to better understand to what extent the proposition can help reduce the search space, we take layer number $L = 4$ and the ratio is calculated as $\frac{\text{original space}}{\text{transformed space}} = \frac{6^{C \times 4}}{2^{4 \times 4}}$.*

Table 1: The reduction ratio w.r.t different C .

C	1	2	3	4	5	6
Ratio	2.0×10^{-2}	2.6×10^1	3.3×10^4	4.3×10^7	5.6×10^{10}	7.2×10^{13}

Note that when $C = 1$, the transformation will not lead to a reduction on the space. However, such a case is not meaningful as it suffers from poor performance due to lack of flexibility for the modulation function (see Section 4.3.2). The transformed spaces enable efficient and robust differentiable search, via reducing the number of architecture parameter from $6 \times C \times L$ to $4 \times L$. Meanwhile, the space size is transformed from $6^{C \times L}$ to $2^{4 \times L}$, which is a reduction for any $C > 1$.

3.3.2 Construction of the Supernet. For each M^l , since there are 4 operations at most and they are permutation-invariant, we only need to decide whether to take the operation or not by any order. We search by introducing differentiable parameters to weigh the operations and optimize the weights. For the l th layer of the predictor, we have

$$\hat{\mathbf{h}}^{l,k+1} = \alpha^{l,k+1} (\hat{\mathbf{h}}^{l,k} \circ_{\text{op}^{k+1}} \phi_i^{k+1}) + (1 - \alpha^{l,k+1}) \hat{\mathbf{h}}^{l,k}, \quad (8)$$

where $\alpha^{l,k+1}$ is a weight to measure operation $\circ_{\text{op}^{k+1}}$ in M^l , $k \in \{0, 1, 2, 3\}$ and $\{\circ_{\text{op}^{k+1}}\}_{k=0}^3 = \{\max, \min, \odot, +\}$. For notation simplicity, we let $\hat{\mathbf{h}}^{l,0} = \mathbf{h}^l$, $\mathbf{h}_i^l = \hat{\mathbf{h}}^{l,4}$. We construct the supernet by replacing (3) with (8), i.e., replacing every M^l in red dotted lines in Figure 1 with the structure shown in Figure 2 (c).

3.3.3 Complete Algorithm. The complete algorithm is summarized in Algorithm 1. We first optimize the supernet and make selection to determine the structure, then reconstruct the model with determined structure and retrain it to inference.

Benefited from the great reduction brought by the space transformation, while conventional differentiable architecture search [23] optimizes the supernet w.r.t. the bilevel objective with the upper-level variable α and lower-level variable θ :

$$\min_{\alpha} \mathcal{L}^{\text{val}}(\theta^*(\alpha), \alpha), \text{ s.t. } \theta^*(\alpha) = \operatorname{argmin}_{\theta} \mathcal{L}^{\text{train}}(\theta, \alpha), \quad (9)$$

where \mathcal{L}^{val} and $\mathcal{L}^{\text{train}}$ represent the loss obtained on validation set and training set respectively, we only need to optimize the supernet w.r.t. objective $\mathcal{L}^{\text{train}}$ only, in an end-to-end manner by episodic training. For every task $T_i \in \mathcal{T}^{\text{train}}$, we first input \mathcal{S}_i to the adaptation network A to generate Φ_i by (1), and then for every item $v_j \in \mathcal{Q}_i$, we take (u_i, v_j) and Φ_i as input of the predictor P and




Figure 2: Illustration of our proposition 3.1, and the structure of the supernet to search on the reduced space.

Algorithm 1 Training procedure of ColdNAS.

- Input:** Learning rate β , number of operations to keep K .
- 1: Construct the supernet by (8) and randomly initialize all parameters $\Theta = \{\{\alpha^{l,k}\}_{k=1,l=0}^{4,L-1}, \theta_E, \theta_A, \theta_P\}$.
 - 2: **while** Not converge **do**
 - 3: **for** Every $T_i \in \mathcal{T}^{\text{train}}$ **do**
 - 4: Calculate Φ_i by (1).
 - 5: Calculate $\hat{y}_{i,j}$ for every v_j in Q_i by (2).
 - 6: Calculate loss \mathcal{L}_i by (10).
 - 7: **end for**
 - 8: $\mathcal{L}^{\text{train}} = \frac{1}{|\mathcal{T}^{\text{train}}|} \sum_{i=1}^{|\mathcal{T}^{\text{train}}|} \mathcal{L}_i$
 - 9: Update all parameters $\Theta \leftarrow \Theta - \beta \nabla_{\Theta} \mathcal{L}^{\text{train}}$.
 - 10: **end while**
 - 11: Determine the modulation structure by keeping operations corresponding to Top- K $\alpha^{l,k}$ and remove the others.
 - 12: Construct the model with determined modulation structure and randomly initialize all parameters $\Theta = \{\theta_E, \theta_A, \theta_P\}$.
 - 13: Train the model in the same way as Step 2 ~ 10.
 - 14: **Return:** The trained model.
-

make prediction by (2). Then, we use mean squared error (MSE) between the prediction $\hat{y}_{i,j}$ and true label $y_{i,j}$ as loss function:

$$\mathcal{L}_i = \frac{1}{M} \sum_{j=1}^M (y_{i,j} - \hat{y}_{i,j})^2, \quad (10)$$

and $\mathcal{L}^{\text{train}} = \sum_{T_i \in \mathcal{T}^{\text{train}}} \mathcal{L}_i$. We update all parameters by gradient descent. Once the supernet converges, we determine all M^l 's jointly by keeping the operation corresponding to the Top- K largest values among the $4 \times L$ values in $\{\alpha^{l,k+1}\}_{k=1,l=0}^{4,L-1}$. We then retrain the model to obtain the final user cold-start model with searched modulation structure. During inference, a new set of tasks $\mathcal{T}^{\text{test}}$ is given, which is disjoint from $\mathcal{T}^{\text{train}}$. For $T_i \in \mathcal{T}^{\text{test}}$, we take the whole S_i and (u_i, v_j) as input to the trained model to obtain prediction $\hat{y}_{i,j}$ for each item $v_j \in Q_i$.

3.4 Discussion

Existing works in space transformation can be divided into three types. One is using greedy search strategy [9]. Methods of this kind explore different groups of architectures in the search space greedily. Thus, they fail to explore the full search space and can easily fall into bad local optimal. Another is mapping the search space into a low-dimensional one. For examples, using auto-encoder [25] or sparse coding [40]. However, these types of methods do not consider special properties of the search problem, e.g., the graph

structure of the supernet. Thus, the performance ranking of architectures may suffer from distortion in the low-dimensional space. ColdNAS belongs to the third type, which is to explore architecture equivalence in the search space. The basic idea is that if we can find a group of architectures that are equivalent with each other, then evaluation any one of them is enough for all architectures in the same group. This type of methods is problem-specific. For example, perturbation equivalence for matrices is explored in [45], morphism of networks is considered in [17]. The search space of ColdNAS is designed for modulation structures in cold-start problem, which has not been explored before. We theoretically prove the equivalence between the original space and the transformed space, which then significantly reduces the space size (Remark 2). Other methods for general space reduction cannot achieve that.

4 EXPERIMENTS

We perform experiments on three benchmark datasets with the aim to answer the following research questions:

- **RQ1:** What is the modulation structure selected by ColdNAS and how does ColdNAS perform in comparison with the state-of-the-art cold-start models?
- **RQ2:** How can we understand the search space and algorithm of ColdNAS?
- **RQ3:** How do hyperparameters affect ColdNAS?

Results are averaged over five runs.

4.1 Datasets

We use three benchmark datasets (Table 4): (i) **MovieLens** [12]: a dataset containing 1 million movie ratings of users collected from MovieLens, whose features include gender, age, occupation, Zip code, publication year, rate, genre, director and actor; (ii) **BookCrossing** [47]: a collection of users' ratings on books in BookCrossing community, whose features include age, location, publish year, author, and publisher; and (iii) **Last.fm**: a collection of user's listening count of artists from Last.fm online system, whose features only consist of user and item IDs. Following Lin et al. [21], we generate negative samples for the query sets in Last.fm.

Data Split. Following Lin et al. [21], the ratio of $\mathcal{T}^{\text{train}} : \mathcal{T}^{\text{val}} : \mathcal{T}^{\text{test}}$ is set as 7 : 1 : 2. \mathcal{T}^{val} is used to judge the convergence of supernet. $\mathcal{T}^{\text{train}}, \mathcal{T}^{\text{val}}, \mathcal{T}^{\text{test}}$ contain no overlapping users. For MovieLens and Last.fm, we keep any user whose interaction history length lies in [40, 200]. Each support set contains $N = 20$ randomly selected interactions of a user, and query set contains the rest interactions of the same user. As for BookCrossing with severe long-tail distribution of user-item interactions, we particularly put any

Table 2: Test performance (%) obtained on benchmark datasets. The best results are highlighted in bold and the second-best in italic. For MSE and MAE, smaller value is better. For nDCG₃ and nDCG₅, larger value is better.

Dataset	Metric	DropoutNet	MeLU	MetaCS	MetaHIN	MAMO	TaNP	ColdNAS-Fixed	ColdNAS
MovieLens	MSE	100.90 _(0.70)	95.02 _(0.03)	95.05 _(0.04)	91.89 _(0.06)	90.20 _(0.22)	<u>89.11</u> _(0.18)	91.05 _(0.13)	87.96 _(0.12)
	MAE	85.71 _(0.48)	77.38 _(0.25)	77.42 _(0.26)	75.79 _(0.27)	75.34 _(0.26)	<u>74.78</u> _(0.14)	75.65 _(0.30)	74.29 _(0.20)
	nDCG ₃	69.21 _(0.76)	74.43 _(0.59)	74.46 _(0.78)	74.69 _(0.32)	74.95 _(0.13)	<u>75.60</u> _(0.07)	75.11 _(0.09)	76.16 _(0.03)
	nDCG ₅	68.43 _(0.48)	73.52 _(0.41)	73.45 _(0.56)	73.63 _(0.22)	73.84 _(0.16)	<u>74.29</u> _(0.12)	73.89 _(0.12)	74.74 _(0.09)
BookCrossing	MSE	15.38 _(0.23)	15.15 _(0.02)	15.20 _(0.08)	14.76 _(0.07)	14.82 _(0.05)	14.75 _(0.05)	<u>14.44</u> _(0.16)	14.15 _(0.08)
	MAE	3.75 _(0.01)	3.68 _(0.01)	3.66 _(0.01)	3.50 _(0.01)	3.51 _(0.02)	<u>3.48</u> _(0.01)	3.49 _(0.02)	3.40 _(0.01)
	nDCG ₃	77.66 _(0.18)	<u>77.69</u> _(0.15)	77.68 _(0.12)	77.66 _(0.19)	77.68 _(0.09)	77.48 _(0.06)	77.65 _(0.09)	77.83 _(0.01)
	nDCG ₅	80.87 _(0.15)	81.10 _(0.15)	80.97 _(0.09)	80.95 _(0.04)	81.01 _(0.05)	<u>81.16</u> _(0.21)	81.12 _(0.06)	81.32 _(0.10)
Last.fm	MSE	21.91 _(0.38)	21.69 _(0.34)	21.68 _(0.12)	<u>21.43</u> _(0.23)	21.64 _(0.10)	21.58 _(0.20)	21.62 _(0.16)	20.91 _(0.05)
	MAE	43.02 _(0.52)	42.28 _(1.21)	42.28 _(0.76)	<u>42.07</u> _(0.49)	42.30 _(0.28)	42.15 _(0.56)	42.32 _(0.34)	41.78 _(0.24)
	nDCG ₃	75.13 _(0.48)	80.15 _(2.09)	80.81 _(0.97)	<u>82.01</u> _(0.56)	80.73 _(0.80)	81.03 _(0.36)	80.77 _(0.32)	82.80 _(0.69)
	nDCG ₅	69.03 _(0.31)	75.03 _(0.68)	75.01 _(0.64)	<u>75.98</u> _(0.33)	75.45 _(0.29)	<u>75.98</u> _(0.41)	75.48 _(0.21)	76.77 _(0.10)

Table 3: Modulation structure with Top-4 operations searched on the three benchmark datasets respectively. We also show the modulation structure of ColdNAS-Fixed, which is the same regardless of the dataset used.

	M^0	M^1	M^2	M^3
MovieLens	$\min(\max(\mathbf{h}^0, \phi_i^{0,1}), \phi_i^{0,2}) + \phi_i^{0,3}$	$\mathbf{h}^1 + \phi_i^{1,1}$	\mathbf{h}^2	\mathbf{h}^3
BookCrossing	$\min(\mathbf{h}^0, \phi_i^{0,1})$	$\mathbf{h}^1 + \phi_i^{1,1}$	$\mathbf{h}^2 \odot \phi_i^{2,1} + \phi_i^{2,2}$	\mathbf{h}^3
Last.fm	$\mathbf{h}^0 + \phi_i^{0,1}$	$\mathbf{h}^1 + \phi_i^{1,1}$	$\max(\mathbf{h}^2, \phi_i^{2,1}) + \phi_i^{2,2}$	\mathbf{h}^3
ColdNAS-Fixed	$\mathbf{h}^0 \odot \phi_i^{0,1} + \phi_i^{0,2}$	$\mathbf{h}^1 \odot \phi_i^{1,1} + \phi_i^{1,2}$	$\mathbf{h}^2 \odot \phi_i^{2,1} + \phi_i^{2,2}$	$\mathbf{h}^3 \odot \phi_i^{3,1} + \phi_i^{3,2}$

Table 4: Summary of datasets used in this paper.

Dataset	# User (Cold)	# Item	# Rating	# User Feat.	# Item Feat.
MovieLens	6040 (52.3%)	3706	1000209	4	5
BookCrossing	278858 (18.6%)	271379	1149780	2	3
Last.fm	1872 (15.3%)	3846	42346	1	1

user whose interaction history length lies in [50, 1000] into $\mathcal{T}^{\text{train}}$. Then, we divide users with interaction history length in [2, 50] into 70%, 10% and 20% to be put in $\mathcal{T}^{\text{train}}$, \mathcal{T}^{val} , $\mathcal{T}^{\text{test}}$ respectively. The proportion of cold users in each dataset is also shown in Table 4. Then, we randomly select half of each user’s interaction history as support set and take the rest as query set.

Evaluation Metric. Following [19, 26], we evaluate the performance by mean average error (MAE), mean squared Error (MSE), normalized discounted cumulative gain nDCG₃ and nDCG₅. MAE and MSE evaluate the numerical gap between the prediction and the ground-truth rating, lower value is better. For nDCG₃ and nDCG₅, the higher value is better, representing the proportion between the discounted cumulative gain of the predicted item list and the ground-truth list.

4.2 Performance Comparison (RQ1)

We compare ColdNAS with the following representative user cold-start methods: (i) traditional deep cold-start model **DropoutNet** [36] and (ii) FSL based methods include **MeLU** [19], **MetaCS** [3], **MetaHIN** [24], **MAMO** [6], and **TaNP** [21]. We run the public codes provided by the respective authors. PNMTA [26], CMML [7], PAML [37] and REG-PAML [44] are not compared due to the lack of public codes. We choose a 4-layer predictor, more details of our model and parameter setting are provided in Appendix C.1. We also compare with a variant of ColdNAS called **ColdNAS-Fixed**, which uses the fixed FiLM function in (5) at every layer rather than our searched modulation function.

Table 2 shows the overall user-cold start recommendation performance for all methods. We can see that ColdNAS significantly outperforms the others on all the datasets and metrics. Among all compared baselines, DropoutNet performs the worst as it is not a few-shot learning method that the model has no ability to adapt to different users. Among meta-learning based methods, MeLU, MetaCS, MetaHIN and MAMO adopt gradient-based meta-learning strategy, which may suffer from overfitting during local-updates. In contrast, TaNP and ColdNAS learn to generate user-specific parameters to guide the adaptation. TaNP uses a fixed modulation structure which may not be optimal for different datasets, while ColdNAS automatically finds the optimal structure. Further, the




Figure 3: Testing MSE vs clock time of different search strategies. ColdNAS, ColdNAS-bilevel and Random-t operate on the transformed space, while Random operates on the original space with $C = 4$ in every M^l .

consistent performance gain of ColdNAS over ColdNAS-Fixed validates the necessity of searching modulation structure to fit datasets instead of using a fixed one.

Table 3 shows the searched modulation structures. We find that the searched modulation structures are different across the datasets. No modulation should be taken at the last layer, as directly use user-specific preference possibly leads to overfitting. Note that all the operations in the transformed search space have been selected at least once, which means all of them are helpful.

Comparing time consumption to other cold-start models, ColdNAS only additionally optimizes a supernet with the same objective and comparable size. Table 5 shows the clock time taken by ColdNAS and TaNP which obtains the second-best in Table 2. As can be observed, the searching in ColdNAS is very efficient, as it is only slightly more expensive than retraining.

Table 5: Clock time taken by ColdNAS and TaNP.

Clock time (min)		MovieLens	BookCrossing	Last.fm
TaNP		15.5	44.2	4.1
ColdNAS	Search	16.2	45.5	3.9
	Retrain	12.7	35.5	3.5

4.3 Searching in ColdNAS (RQ2)

4.3.1 Choice of Search Strategy. In ColdNAS, we optimize (10) by gradient descent. In this section, we compare ColdNAS with other search strategies to search the modulation structure, including **Random ($C = 4$)** which conducts random search [2] on the original space with operation number $C = 4$ in each M^l , **Random (T)** which conducts random search [2] on the transformed space, **ColdNAS-Bilevel** which optimizes the supernet with the bilevel objective 9. For random search, we record its training and evaluation time. For ColdNAS and ColdNAS-bilevel, we sample the architecture parameter of the supernet and retrain from scratch, and exclude the time for retraining and evaluation. Figure 3 shows the performance of the best searched architecture. One can observe that searching on the transformed space allows higher efficiency and accuracy, comparing Random (T) to Random ($C = 4$). In addition, differentiable search on the supernet structure has brought more efficiency, as both ColdNAS and ColdNAS-Bilevel outperform random search. ColdNAS and ColdNAS-Bilevel converge to similar testing MSE, but ColdNAS is faster. This validates the efficacy of directly using gradient descent on all parameters in ColdNAS.

4.3.2 Necessity of Search Space Transformation. Now that the search algorithm has been chosen, we now particularly examine the necessity of search space transformation in terms of time and performance. As shown in Table 1, a larger C which constrains the number of operations will lead to a larger original space. And the space reduction ratio grows exponentially along with the original space size. Figure 4 plots the testing MSE vs clock time of random search on original spaces of different size (Random ($C = c$)), random search on transformed space (Random (T)), and our ColdNAS. When C is small, modulation functions would not be flexible enough, though it is easy to find the optimal candidate in this small search space. When C is large, the search space is large, where random search would be too time-consuming to find a good modulation structure. In contrast, the transformed search space has a consistent small size, and is theoretically proved to be equal to the space with any C . As can be seen, both Random (T) and ColdNAS can find good structure more effective than Random ($C = 4$), and ColdNAS searches the fastest via differentiable search.

4.3.3 Understanding Proposition 3.1. Here, we first show that the assumption of adaptation network A being expressive enough can be easily satisfied. Generally, deeper networks are more expressive. Figure 6 plots the effect of changing the depth of the neural network in A on Last.fm. As shown in Figure 6 (a-d), although different number of layers in A are used, the searched modulation operations are the same. Consequently, the performance difference is small, as shown in Figure 6 (d). Thus, we use 2 layers which is expressive enough and has smaller parameter size.

Further, we validate the inner-group consistence and permutation-invariance properties proved in Proposition 3.1. Comparing Figure 6 (a) with Figure 7 (a-b), one can see that changing \odot to $/$ and $+$ to $-$ obtains the same results, which validates inner-group consistency. Finally, comparing Figure 6 (a) with Figure 7 (c), one can observe that inter-group permutation invariant also holds.

4.4 Sensitivity Analysis (RQ3)

Finally, we conduct sensitivity analysis of ColdNAS on Last.fm.

Effects of K . Recall that we keep modulation operations corresponding to the Top- K largest $\alpha^{l,k}$ s in modulation structure. Figure 5(a) plots the effect of changing K . As shown, K cannot be too small that the model is unable to capture enough user-specific preference, nor too large that the model can overfit to the limited interaction history of cold-start users.




Figure 4: Testing MSE vs clock time on various search spaces: transformed space, and original spaces with different C .




Figure 5: Model sensitivity analysis of ColdNAS on Last.fm.

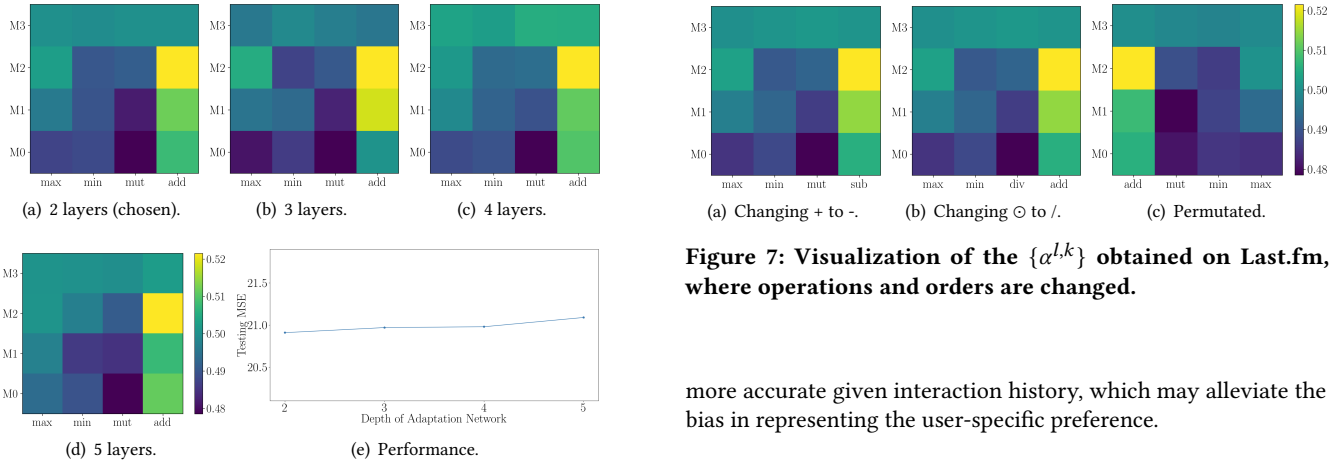


Figure 6: Effect of changing the depth of adaptation network on Last.fm. Figure 6 (a-d) shows the visualization of the $\{\alpha^{l,k}\}$ obtained with varying depth, and Figure 6 (e) shows the testing MSE obtained correspondingly.

Effect of the Depth of Predictor. Figure 5(b) plots the effect of using predictors with different L number of layers. One can observe that choosing different L in a certain range has low impact on the performance, while choosing $L = 4$ is already good enough to obtain the state-of-the-art results as shown in Table 2.

Effects of Support Set Size. Figure 5(c) shows the performance of users with different length of history. During inference, we randomly sample only a portion of interactions from the original support set S_i of each T_i in T^{test} . As can be seen, prediction would be

Figure 7: Visualization of the $\{\alpha^{l,k}\}$ obtained on Last.fm, where operations and orders are changed.

more accurate given interaction history, which may alleviate the bias in representing the user-specific preference.

5 CONCLUSION

We propose a modulation framework called ColdNAS for user cold-start recommendation. In particular, we use a hypernetwork to map each user’s history interactions to user-specific parameters, which are then used to modulate the predictor. We design a search space for modulation functions and positions, which not only covers existing modulation-based models but also has the ability to find more effective structures. We theoretically prove the space could be transformed to a smaller space, where we can search for modulation structure efficiently and robustly. Extensive experiments show that ColdNAS performs the best on benchmark datasets. Besides, ColdNAS can efficiently find proper modulation structure for different data, which make it easy to be deployed in recommendation systems.

ACKNOWLEDGMENTS

Q. Yao was in part supported by NSFC (No. 92270106) and CCF-Baidu Open Fund.

REFERENCES

- [1] Bowen Baker, Otkrist Gupta, Nikhil Naik, and Ramesh Raskar. 2017. Designing neural network architectures using reinforcement learning. In *International Conference on Learning Representations*.
- [2] James Bergstra and Yoshua Bengio. 2012. Random search for hyper-parameter optimization. *Journal of Machine Learning Research* 13, 2 (2012).
- [3] Homanga Bharadhwaj. 2019. Meta-learning for user cold-start recommendation. In *International Joint Conference on Neural Networks*. 1–8.
- [4] Marc Brockschmidt. 2020. GNN-FiLM: Graph neural networks with feature-wise linear modulation. In *International Conference on Machine Learning*. 1144–1152.
- [5] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishikesh Aradhye, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ispir, et al. 2016. Wide & deep learning for recommender systems. In *Workshop on Deep Learning for Recommender Systems*. 7–10.
- [6] Manqing Dong, Feng Yuan, Lina Yao, Xiwei Xu, and Liming Zhu. 2020. MAMO: Memory-augmented meta-optimization for cold-start recommendation. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 688–697.
- [7] Xidong Feng, Chen Chen, Dong Li, Mengchen Zhao, Jianye Hao, and Jun Wang. 2021. CMML: Contextual modulation meta learning for cold-start recommendation. In *ACM International Conference on Information and Knowledge Management*. 484–493.
- [8] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*. 1126–1135.
- [9] Chen Gao, Yinfeng Li, Quanming Yao, Depeng Jin, and Yong Li. 2021. Progressive feature interaction search for deep sparse network. In *Advances in Neural Information Processing Systems*, Vol. 34. 392–403.
- [10] Chen Gao, Quanming Yao, Depeng Jin, and Yong Li. 2021. Efficient Data-specific Model Search for Collaborative Filtering. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 415–425.
- [11] David Ha, Andrew Dai, and Quoc V Le. 2017. Hypernetworks. In *International Conference on Learning Representations*.
- [12] F Maxwell Harper and Joseph A Konstan. 2015. The MovieLens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems* 5, 4 (2015), 1–19.
- [13] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *International Conference on World Wide Web*. 173–182.
- [14] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. 1989. Multilayer feed-forward networks are universal approximators. *Neural Networks* 2, 5 (1989), 359–366.
- [15] Cheng-Kang Hsieh, Longqi Yang, Yin Cui, Tsung-Yi Lin, Serge Belongie, and Deborah Estrin. 2017. Collaborative metric learning. In *International Conference on World Wide Web*. 193–201.
- [16] Frank Hutter, Lars Kotthoff, and Joaquin Vanschoren. 2019. *Automated machine learning: methods, systems, challenges*. Springer Nature.
- [17] Haifeng Jin, Qingquan Song, and Xia Hu. 2019. Auto-Keras: An efficient neural architecture search system. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1946–1956.
- [18] Yehuda Koren. 2008. Factorization meets the neighborhood: A multifaceted collaborative filtering model. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 426–434.
- [19] Hoyeop Lee, Jinbae Im, Seongwon Jang, Hyunsouk Cho, and Sehee Chung. 2019. MeLU: Meta-learned user preference estimator for cold-start recommendation. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1073–1082.
- [20] Jovian Lin, Kazunari Sugiyama, Min-Yen Kan, and Tat-Seng Chua. 2013. Addressing cold-start in app recommendation: latent user models constructed from twitter followers. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*. 283–292.
- [21] Xixun Lin, Jia Wu, Chuan Zhou, Shirui Pan, Yanan Cao, and Bin Wang. 2021. Task-adaptive neural process for user cold-start recommendation. In *The Web Conference*. 1306–1316.
- [22] Bin Liu, Chenxu Zhu, Guilin Li, Weinan Zhang, Jincai Lai, Ruiming Tang, Xiquang He, Zhenguo Li, and Yong Yu. 2020. AutoFIS: Automatic feature interaction selection in factorization models for click-through rate prediction. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2636–2645.
- [23] Hanxiao Liu, Karen Simonyan, and Yiming Yang. 2019. DARTS: Differentiable architecture search. In *International Conference on Learning Representations*.
- [24] Yuanfu Lu, Yuan Fang, and Chuan Shi. 2020. Meta-learning on heterogeneous information networks for cold-start recommendation. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1563–1573.
- [25] Renqian Luo, Fei Tian, Tao Qin, Enhong Chen, and Tie-Yan Liu. 2018. Neural architecture optimization. In *Advances in Neural Information Processing Systems*, Vol. 31. 7827–7838.
- [26] Haoyu Pang, Fausto Giunchiglia, Ximing Li, Renchu Guan, and Xiaoyue Feng. 2022. PNMETA: A pretrained network modulation and task adaptation approach for user cold-start recommendation. In *The Web Conference*. 348–359.
- [27] Yoon-Joo Park and Alexander Tuzhilin. 2008. The long tail of recommender systems and how to leverage it. In *ACM Conference on Recommender Systems*. 11–18.
- [28] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. 2018. FiLM: Visual reasoning with a general conditioning layer. In *AAAI Conference on Artificial Intelligence*, Vol. 32. 3942–3951.
- [29] Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. 2018. Efficient neural architecture search via parameters sharing. In *International Conference on Machine Learning*. 4095–4104.
- [30] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. 2019. Regularized evolution for image classifier architecture search. In *AAAI Conference on Artificial Intelligence*, Vol. 33. 4780–4789.
- [31] James Requeima, Jonathan Gordon, John Bronskill, Sebastian Nowozin, and Richard E Turner. 2019. Fast and flexible multi-task classification using conditional neural adaptive processes. In *Advances in Neural Information Processing Systems*, Vol. 32. 7957–7968.
- [32] Paul Resnick and Hal R Varian. 1997. Recommender systems. *Commun. ACM* 40, 3 (1997), 56–58.
- [33] Andrew I Schein, Alexandrin Popescul, Lyle H Ungar, and David M Pennock. 2002. Methods and metrics for cold-start recommendations. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*. 253–260.
- [34] Suvash Sedhain, Aditya Krishna Menon, Scott Sanner, and Lexing Xie. 2015. AutoRec: Autoencoders meet collaborative filtering. In *International Conference on World Wide Web*. 111–112.
- [35] Qingquan Song, Dehua Cheng, Hanning Zhou, Jiyan Yang, Yuandong Tian, and Xia Hu. 2020. Towards automated neural interaction discovery for click-through rate prediction. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 945–955.
- [36] Maksims Volkovs, Guangwei Yu, and Tomi Poutanen. 2017. DropoutNet: Addressing cold start in recommender systems. In *Advances in Neural Information Processing Systems*, Vol. 30. 4957–4966.
- [37] Li Wang, Binbin Jin, Zhenya Huang, Hongke Zhao, Defu Lian, Qi Liu, and Enhong Chen. 2021. Preference-adaptive meta-learning for cold-start recommendation. In *International Joint Conference on Artificial Intelligence*. 1607–1614.
- [38] Yaqing Wang, Quanming Yao, James T Kwok, and Lionel M Ni. 2020. Generalizing from a few examples: A survey on few-shot learning. *Comput. Surveys* 53, 3 (2020), 1–34.
- [39] Yuexiang Xie, Zhen Wang, Yaliang Li, Bolin Ding, Nezihe Merve Gürel, Ce Zhang, Minlie Huang, Wei Lin, and Jingren Zhou. 2021. FIVES: Feature interaction via edge search for large-scale tabular data. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 3795–3805.
- [40] Yibo Yang, Hongyang Li, Shan You, Fei Wang, Chen Qian, and Zhouchen Lin. 2020. ISTA-NAS: Efficient and consistent neural architecture search by sparse coding. In *Cross attention network for few-shot classification*, Vol. 33. 10503–10513.
- [41] Quanming Yao, Xiangning Chen, James T Kwok, Yong Li, and Cho-Jui Hsieh. 2020. Efficient neural interaction function search for collaborative filtering. In *The Web Conference*. 1660–1670.
- [42] Quanming Yao, Ju Xu, Wei-Wei Tu, and Zhanxing Zhu. 2020. Efficient neural architecture search via proximal iterations. In *AAAI Conference on Artificial Intelligence*, Vol. 34. 6664–6671.
- [43] Kaicheng Yu, Christian Sciufo, Martin Jaggi, Claudiu Musat, and Mathieu Salzmann. 2019. Evaluating the search phase of neural architecture search. In *International Conference on Learning Representations*.
- [44] Runsheng Yu, Yu Gong, Xu He, Bo An, Yu Zhu, Qingwen Liu, and Wenwu Ou. 2021. Personalized adaptive meta learning for cold-start user preference prediction. In *AAAI Conference on Artificial Intelligence*, Vol. 35. 10772–10780.
- [45] Yongqi Zhang, Quanming Yao, and James T Kwok. 2023. Bilinear scoring function search for knowledge graph learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45, 2 (2023), 1458–1473.
- [46] Xiangyu Zhao, Haochen Liu, Wenqi Fan, Hui Liu, Jiliang Tang, and Chong Wang. 2021. AutoLoss: Automated loss function search in recommendations. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 3959–3967.
- [47] Cai-Nicolas Ziegler, Sean M McNee, Joseph A Konstan, and Georg Lausen. 2005. Improving recommendation lists through topic diversification. In *International Conference on World Wide Web*. 22–32.

A DETAILS OF MODEL STRUCTURE

As mentioned in Section 3.2, existing models [6, 21, 26] share an architecture consisting of three components: embedding layer E , adaptation network A , and predictor P . Here we provide the details of these components used in ColdNAS.

A.1 Embedding Layer

We first embed user and item’s one-hot categorical features into dense vectors through the embedding layer. Taking u_i for example, we generate a content embedding for each categorical content feature and concatenate them together to obtain the initial user embedding. Given B user contents, the embedding is defined as:

$$u_i = [W_E^1 c_i^1 \mid W_E^2 c_i^2 \mid \dots \mid W_E^B c_i^B] \quad (11)$$

where $[\cdot \mid \cdot]$ is the concatenation operation, c_i^b is the one-hot vector of the b th categorical content of u_i , and W_E^b represents the embedding matrix of the corresponding feature in the shared user feature space. The parameters of embedding layer are collectively denoted as $\theta_E = \{W_E^b\}_{b=1}^B$.

A.2 Adaptation Network

We first use a fully connected (FC) layer to get hidden representation of interaction history $r_{i,j}$, which is calculated as

$$r_{i,j} = \text{ReLU}(W_A^1 [u_i \mid v_j \mid y_{i,j}] + b_A^1). \quad (12)$$

We use mean pooling to aggregate the interactions in S_i as the task context information of u_i , i.e., preference, $c_i = \frac{1}{N} \sum_1^N r_{i,j}$. Then, we use another FC layer to generate the user-specific parameters as $\phi_i = W_A^2 c_i + b_A^2$. The parameters of adaptation network is denoted as $\theta_A = \{W_A^1, b_A^1, W_A^2, b_A^2\}$.

A.3 Predictor

We use a L -layer MLP as the predictor. Denote output of the l th layer hidden units as h^l , the modulation function at the l th layer as M^l , we obtain prediction by

$$\begin{aligned} h_i^l &= M_l(h^l, \Phi_i), \\ h^{l+1} &= \text{ReLU}(W_P^l h_i^l + b_P^l), \end{aligned} \quad (13)$$

where $l \in [0, 1, \dots, L-1]$, $h^0 = [u_i \mid v_j]$ and $\hat{y}_{i,j} = h^L$. The parameters of predictor are denoted as $\theta_P = \{W_P^l, b_P^l\}_{l=0}^{L-1}$.

B PROOF OF THE PROPOSITION 3.1

PROOF. To prove Proposition 3.1, we need to use the two conditions below:

Condition 1: If ϕ_i^k is the input of operation \odot or $/$, then every element in ϕ_i^k is non-negative. We implement this as ReLU function.

Condition 2: The adaptation network A is expressive enough, that it can approximate $\hat{\phi}_i = \phi_i^p \circ_{\text{op}^1} \phi_i^q$, where $\circ_{\text{op}^1} \in \mathcal{O}$ and $\phi_i^p, \phi_i^q \in \Phi_i$ learned from the data. This naturally holds due to the assumption.

Divide \circ_{op} into 4 groups: $G_1 = \{\max\}$, $G_2 = \{\min\}$, $G_3 = \{+, -\}$, $G_4 = \{\odot, /\}$. We can prove two important properties below.

Property 1: Inner-group consistence. We can choose an operation in each group as the group operation \circ_{g_i} , e.g. $\circ_{g_1} = \max$, $\circ_{g_2} = \min$, $\circ_{g_3} = +$, $\circ_{g_4} = \odot$. Then any number of successive operations that belong to the same group G_i can be associated into one operation \circ_{g_i} , i.e.,

$$\begin{aligned} x \circ_{\text{op}^k} \phi_i^k \circ_{\text{op}^{k+1}} \phi_i^{k+1} \dots \circ_{\text{op}^{k+m}} \phi_i^{k+m} &= x \circ_{g_i} \hat{\phi}_i^k, \\ \text{s.t. } \circ_{\text{op}^k}, \circ_{\text{op}^{k+1}} \dots \circ_{\text{op}^{k+m}} &\in G_i \end{aligned}$$

It’s trivial to prove with condition 2, e.g., $x + \phi_i^1 - \phi_i^2 + \phi_i^3 = x + \hat{\phi}_i^1$, where $\hat{\phi}_i^1 = \phi_i^1 - \phi_i^2 + \phi_i^3$.

Property 2: Inter-group permutation-invariance. The operations in different groups are permutation-invariant, i.e.,

$$x \circ_{g_i} \phi_i^k \circ_{g_j} \phi_i^{k+1} = x \circ_{g_j} \hat{\phi}_i^k \circ_{g_i} \hat{\phi}_i^{k+1}$$

It is also trivial to prove with condition 1 and condition 2, e.g.,

$$\max(x, \phi_i^1) + \phi_i^2 = \max(x + \hat{\phi}_i^1, \hat{\phi}_i^2),$$

where $\hat{\phi}_i^1 = \phi_i^2$ and $\hat{\phi}_i^2 = \phi_i^1 + \phi_i^2$.

With the above two properties, we can recently commute operations until operations in the same group are gathered, and merge operations in the four groups respectively. So

$$h_i^l = h^l \circ_{\text{op}^1} \phi_i^1 \circ_{\text{op}^2} \phi_i^2 \dots \circ_{\text{op}^C} \phi_i^C$$

equals to

$$h_i^l = h^l \circ_{g_1} \hat{\phi}_i^1 \circ_{g_2} \hat{\phi}_i^2 \circ_{g_3} \hat{\phi}_i^3 \circ_{g_4} \hat{\phi}_i^4,$$

and the four group operations are permutation-invariant.

Note that since the identity element ($\mathbf{0}$ or $\mathbf{1}$) of each group operation can be learned from condition 2, a modulation function that doesn’t cover all groups can also be represented in the same form. \square

C EXPERIMENTS

C.1 Experiment Setting

Experiments were conducted on a 24GB NVIDIA GeForce RTX 3090 GPU, with Python 3.7.0, CUDA version 11.6.

C.1.1 Hyperparameter Setting. We find hyperparameters using the \mathcal{T}^{val} via grid search for existing methods. In ColdNAS, the batch size is 32, W_E^b in (11) has size $32 \times |c_i^b|$ where $|c_i^b|$ is the length of c_i^b . The dimension of hidden units in (13) is set as $h^1 = 128$, $h^2 = 64$, $h^3 = 32$ for all three datasets. The learning rate β is chosen from $\{5 \times 10^{-6}, 1 \times 10^{-5}, 5 \times 10^{-5}, 1 \times 10^{-4}\}$ and the dimension of $r_{i,j}$ in (12) is chosen from $\{128, 256, 512, 1024\}$. The final hyperparameters chosen on the three benchmark datasets are shown in Table 6.

C.1.2 URLs of Datasets and Baselines. We use three benchmark datasets (Table 4): (i) **MovieLens**² [12]: a dataset containing 1 million movie ratings of users collected from MovieLens, whose features include gender, age, occupation, Zip code, publication year, rate, genre, director and actor; (ii) **BookCrossing**³ [47]: a collection of users’ ratings on books in BookCrossing community, whose features include age, location, publish year, author, and publisher; and (iii) **Last.fm**⁴: a collection of user’s listening count of artists

²<https://grouplens.org/datasets/movielens/1m/>

³<http://www2.informatik.uni-freiburg.de/~cziegler/BX/>

⁴<https://grouplens.org/datasets/hetrec-2011/>




Figure 8: Model sensitivity analysis on MovieLens.




Figure 9: Model sensitivity analysis on BookCrossing.

Table 6: The final hyperparameters chosen in ColdNAS.

Dataset	MovieLens	BookCrossing	Last.fm
β	5×10^{-5}	5×10^{-6}	1×10^{-4}
$r_{i,j}$	1024	512	256
W_A^1	1024×289	512×161	256×65
W_A^2	2052×1024	1540×512	1156×256
b_A^1	1024	512	256
b_A^2	2052	1540	1156
W_P^0	128×289	128×161	128×65
W_P^1	64×128	64×128	64×128
W_P^2	32×64	32×64	32×64
W_P^3	1×32	1×32	1×32
b_P^0	128	128	128
b_P^1	64	64	64
b_P^2	32	32	32
b_P^3	1	1	1
maximum iteration number	50	50	100

from Last.fm online system, whose features only consist of user and item IDs. In experiments, we compare **ColdNAS** with the following representative user cold-start methods: (i) traditional

deep cold-start model **DropoutNet**⁵ [36] and (ii) FSL based methods include **MeLU**⁶ [19], **MetaCS** [3], **MetaHIN**⁷ [24], **MAMO**⁸ [6], and **TaNP**⁹ [21]. MetaCS is very similar to MeLU, except that it updates all parameters during meta-learning. Hence, we implement MetaCS based on the codes of MeLU.

C.2 More Experimental Results

Here, we show empirical results of Section 4.4 of MovieLens and BookCrossing in Figure 8 and 9. The results of all three datasets show similar patterns and the same analysis and conclusions in Section 4.4 are applicable.

C.3 Complexity Analysis

Denote the layer number of adaptation network as L_A , the layer number of predictor as L_P , the support set size as N , and the query set size as M . For notation simplicity, we denote the hidden size of each layer is the same as D . In the search phase (Algorithm 1 step 2~10), the calculation of Φ_i of a task is $O(NL_A D^3)$, and predicting each item costs $O(L_P D^3)$, so the average complexity is $O((L_P + \frac{NL_A}{M})D^3)$. Similarly, in the retrain (Algorithm 1 step 12~13), the time complexity is also $O((L_P + \frac{NL_A}{M})D^3)$. In total, the time complexity is $O((L_P + \frac{NL_A}{M})D^3)$.

⁵<https://github.com/layer6ai-labs/DropoutNet>

⁶<https://github.com/hoyeoplee/MeLU>

⁷<https://github.com/rootlu/MetaHIN>

⁸<https://github.com/dongmanqing/Code-for-MAMO>

⁹<https://github.com/IIEdm/TaNP>