

Zero-1-to-3: Domain-Level Zero-Shot Cognitive Diagnosis via One Batch of Early-Bird Students towards Three Diagnostic Objectives

Weibo Gao^{1,2}, Qi Liu^{1,2*}, Hao Wang^{1,2}, Linan Yue^{1,2}, Haoyang Bi^{1,2}, Yin Gu^{1,2}, Fangzhou Yao^{1,2}, Zheng Zhang^{1,2}, Xin Li^{1,3}, Yuanjing He⁴

¹ University of Science and Technology of China

² Anhui Province Key Laboratory of Big Data Analysis and Application & State Key Laboratory of Cognitive Intelligence

³ Artificial Intelligence Research Institute, iFLYTEK Co., Ltd

⁴ The Open University of China, China

{weibogao,lnyue,bhy0521,gy128,fangzhouyao,zhangzheng}@mail.ustc.edu.cn, {qiliuql,wanghao3,leexin}@ustc.edu.cn, heyuanjing@ouchn.edu.cn

Abstract

Cognitive diagnosis seeks to estimate the cognitive states of students by exploring their logged practice quiz data. It plays a pivotal role in personalized learning guidance within intelligent education systems. In this paper, we focus on an important, practical, yet often underexplored task: domain-level zero-shot cognitive diagnosis (DZCD), which arises due to the absence of student practice logs in newly launched domains. Recent cross-domain diagnostic models have been demonstrated to be a promising strategy for DZCD. These methods primarily focus on how to transfer student states across domains. However, they might inadvertently incorporate non-transferable information into student representations, thereby limiting the efficacy of knowledge transfer. To tackle this, we propose *Zero-1-to-3*, a domain-level zero-shot cognitive diagnosis framework via one batch of early-bird students towards three diagnostic objectives. Our approach initiates with pre-training a diagnosis model with dual regularizers, which decouples student states into domain-shared and domain-specific parts. The shared cognitive signals can be transferred to the target domain, enriching the cognitive priors for the new domain, which ensures the *cognitive state propagation* objective. Subsequently, we devise a strategy to generate simulated practice logs for cold-start students through analyzing the behavioral patterns from early-bird students, fulfilling the *domain-adaptation* goal. Consequently, we refine the cognitive states of cold-start students as diagnostic outcomes via virtual data, aligning with the *diagnosis-oriented* goal. Finally, extensive experiments on six real-world datasets highlight the efficacy of our model for DZCD and its practical application in question recommendation. The code is publicly available at <https://github.com/bigdata-ustc/Zero-1-to-3>.

1 Introduction

Intelligent education systems offer access to learning resources and tailor-made services, contributing significantly to the burgeoning popularity of online learning. These platforms cover a broad range of learning topics. As shown in




Figure 1: Illustrative procedure of personalized learning in intelligent education systems.

Figure 1, each topic includes an extensive question bank, empowering students to independently select specific questions for targeted practice. By analyzing student practice logs (e.g., correct or incorrect responses), their cognitive states (i.e., the proficiency on specific knowledge concepts) are estimated, which is referred to as the procedure of *cognitive diagnosis* (Wang et al. 2020). The diagnostic results can support further customized applications, such as question recommendation (Liu et al. 2023a) and adaptive testing (Zhuang et al. 2023). As a result, cognitive diagnosis has garnered significant attention from both the “AI for education” community and the general populace (Nguyen 2015).

Previously, a number of cognitive diagnosis models (CDMs) (Embretson and Reise 2013; Reckase 2009; Tsutsumi, Kinoshita, and Ueno 2021; Wang et al. 2020; Gao et al. 2021; Yao et al. 2023) have been developed to enhance diagnostic precision. However, many of these models encounter challenges with the “cold-start” problem. This challenge arises when an online platform launches a novel learning topic with a fresh question bank (e.g., *Physics* in Figure 1). At the initial launch, there exists only a limited collection of practice records from the early-bird learners, who form the first batch of students in this domain. However, practice logs of unseen students remain unavailable for model training. Consequently, the diagnostic performance of traditional CDMs is often impaired as they only work in mature domains where student practice logs are available. We call this task domain-level zero-shot cognitive diagnosis (DZCD). DZCD is an important and practical task aiming to diagnose cognitive states of unseen students, for whom

*Corresponding Author.

practice logs are blank in the new domain.

Recently, promising strategies have emerged to address the challenge of DZCD through cross-domain CDMs (Gao et al. 2023). They primarily focus on the problem of *how to transfer* student cognitive signals from well-established source domains to cold-start target domains. The key point is to profile cognitive states of students based on their historical behaviors from some mature source domains and represent questions in the target domain with available features. However, they tend to overlook a crucial but rarely touched concern: *what to transfer*. This issue involves identifying valuable transferable student states from different topical domains, as not all cognitive signals can be transferred. We argue that there exist two kinds of student cognitive preferences in each domain, namely *domain-shared* and *domain-specific*. Let us consider *Geometry* and *Function* domains as an illustration. Cognitive signals shared across domains (e.g., basic skills) usually provide valuable clues for cross-domain transfer. For example, students may perform similarly on simple questions in both domains because the knowledge examined in these straightforward questions is usually basic and general, e.g., *Addition* and *Subtraction*, which is suitable for cross-domain propagation. In contrast, domain-specific states (e.g., high-level concept *Cube* in *Geometry* and concept *Polynomials* in *Function*) offer deeper insights within their respective domain, while they could be irrelevant or detrimental for other domains, which is commonly non-transferable. For instance, understanding concept *Cube* might not contribute significantly to mastering *Polynomials*. Unfortunately, previous models inevitably encode domain-specific information into student states intended for cross-domain transfer, which hampers transfer performance and can even the negative transfer problem (Hu, Zhang, and Yang 2018).

In this paper, we aim to delve deeper into the problem of “what to transfer” to unlock the full potential of CDMs under DZCD, which is significant but challenging. Ideally, a desirable solution should fulfill the following three objectives (Gao et al. 2023): (1) *diagnosis-oriented*: the model should proficiently diagnose student cognitive states in DZCD scenarios. (2) *cognitive signal propagation*: the model must effectively extract domain-shared states from source domains for the cross-domain propagation of student cognitive signals. (3) *domain-adaption*: for any new cold-start domain, the model is expected to be domain-adaptive, which needs to fully leverage available domain-specific cues in cold-start scenarios, e.g., the first batch of early-bird students in the domain.

Motivated by the above considerations, we propose *Zero-1-to-3*, a domain-level *zero-shot* cognitive diagnosis framework via one batch of early-bird students towards three diagnostic objectives. Specifically, our approach begins with the pre-training of a CDM across multiple source domains to establish an initial profile of students’ states. During this phase, we separate student profiles into domain-shared and domain-specific parts as input of the CDM. Meanwhile, two well-designed regularizers are placed to guide their optimizations. After pre-training, the shared cognitive signals can be refined and can be transferred to the new do-

main to provide useful and broad experiences, which ensures the *cognitive signal propagation* objective. Next, our focus shifts to effectively adapting cold-start students whose practice logs are unavailable. To achieve this, we first use domain-shared states as initial student embeddings in the new domain. Then, we devise a strategy to generate simulated practice logs for unseen students, to fine-tune the states for unseen students. Here, we leverage the cognitive similarity between early-bird and unseen students to transfer practice patterns from the former to the latter, resulting in synthesized data. Notably, as the practice behaviors of early-bird students originate directly from the new domain, they offer distinct insights crucial for achieving the *domain-adaption* goal. After attaining warm-up states for cold-start students through fine-tuning using simulation logs, we can proceed with diagnostic predictions in the new domain characterizing our approach as *diagnosis-oriented*. Notably, it is important to highlight that Zero-1-to-3 framework, as a general framework, is applicable across a wide range of CDMs.

Finally, extensive experimental results on six real-world datasets not only prove that Zero-1-to-3 can effectively perform DZCD and outperform typical baselines, but also highlight a practical application in question recommendation.

2 Related Work

Traditional Cognitive Diagnosis Models. Cognitive diagnosis has been well-researched for decades in educational psychology (Bi et al. 2023; Chen et al. 2023). It aims to profile the cognitive states of students by exploiting their response results (e.g., correct or wrong) (Embretson and Reise 2013; Tong et al. 2022). For instance, Item Response Theory (IRT) (Embretson and Reise 2013) and Multidimensional IRT (MIRT) (Reckase 2009) use unidimensional/multidimensional latent parameters indicating student ability and question difficulty, respectively, to predict student response on this question in a logistic way. Deterministic Inputs, Noisy-And gate (DINA) (De La Torre 2009), NeuralCD (Wang et al. 2020) and RCD (Gao et al. 2021) directly model student proficiency of specific knowledge concepts.

Cross-domain Cognitive Diagnosis. Cross-domain cognitive diagnosis is proposed to address the DZCD issue which arises when an online education platform introduces new domains, resulting in unavailable practice logs for most students. DZCD is a practical task, but research in this area is almost blank (Gao et al. 2023). Existing studies (Gao et al. 2023) on DZCD primarily concentrate on effectively transferring student cognitive signals from source domains to cold-start target domains through cross-domain modeling. The primary challenge is to construct student cognitive representations based on existing domains and utilize question attributes (e.g., textual contents (Liu et al. 2019; Schmucker and Mitchell 2022) and question relational attributes (Gao et al. 2023)) as intermediaries for the cross-domain transfer. However, these methods might inadvertently incorporate non-transferable information into student representations limiting transfer performance.

3 Preliminaries

3.1 Cognitive Diagnosis Model

We first introduce the general form of cognitive diagnosis models (CDMs). We select the general form of CDMs proposed in TechCD (Gao et al. 2023) as our framework, which consists of three types of basic elements: students, questions and knowledge concepts. The diagnosis process can be abstracted as modeling student-question-concept interactions through predicting student practice performances as follows:

$$\hat{y}_{u,v} = \mathcal{M}_{CD}(\mathbf{u}, \mathbf{v}, \mathbf{C}),$$

where \mathbf{u} and \mathbf{v} are the traits of student u (e.g., cognitive states) and question v (e.g., difficulty). \mathbf{C} is the embedding matrix of all knowledge concepts. $\mathcal{M}_{CD}(\cdot)$ is the diagnostic function to predict the response result $\hat{y}_{u,v}$ of student u on question v (correct or wrong). For instance, $\mathcal{M}_{CD}(\cdot)$ is a logistic-like function for IRT/MIRT with uni-/multi-dimensional latent parameters of student ability and question difficulty, respectively, i.e., $\hat{y}_{u,v} = \text{sigmoid}(\mathbf{u} - \mathbf{v})$, where \mathbf{u} and \mathbf{v} are enhanced by fusing concept features \mathbf{C} . It is a multi-layer neural network $f(\cdot)$ in NeuralCD, i.e., $\hat{y}_{u,v} = f(\mathbf{q}_v \circ (\mathbf{p}_u - \mathbf{d}_v))$, where $\mathbf{p}_u \leftarrow f_u(\mathbf{u}, \mathbf{C})$ and $\mathbf{d}_v \leftarrow f_v(\mathbf{v}, \mathbf{C})$. Two full connection layers $f_u(\cdot)$ and $f_v(\cdot)$ are used to fuse knowledge concepts into the student/question embeddings, respectively. Each element $p_{u,c}$ of \mathbf{p}_u denotes the mastery level on concept c of student u . \mathbf{q}_v masks unrelated concepts for question v with element-wise product \circ , where $q_{v,c} = 1$ if question v associates concept c and $q_{v,c} = 0$ otherwise.

To ensure psychometric interpretability of prediction, CDMs should strictly follow the *Monotonicity* assumption: the probability of correctly answering the question monotonically increases with student cognitive proficiency, i.e., $\partial \mathcal{M}_{CD} / \partial \mathbf{u} > 0$ (Tong et al. 2022).

3.2 Problem Setup

In the domain-level zero-shot cognitive diagnosis (DZCD) scenario, we consider M source domains $\mathcal{S}^1, \mathcal{S}^2, \dots, \mathcal{S}^M$ and one target domain \mathcal{T} . In the k -th source domain \mathcal{S}^k , suppose $\mathcal{U}_{\mathcal{S}^k}$, $\mathcal{V}_{\mathcal{S}^k}$ and $\mathcal{C}_{\mathcal{S}^k}$ are the sets of students, questions and concepts. The practice logs in this domain is depicted as $L_{\mathcal{S}^k} = \{(u^k, v^k, y_{u,v}^k) | y_{u,v}^k \in \{0, 1\}, u^k \in \mathcal{U}_{\mathcal{S}^k}, v^k \in \mathcal{V}_{\mathcal{S}^k}\}$, where $y_{u,v}^k = 1$ represents student u answers question v correctly in source domain \mathcal{S}^k , and $y_{u,v}^k = 0$ otherwise. In the target domain \mathcal{T} , student set $\mathcal{U}_{\mathcal{T}}$ is split into two subsets: $\mathcal{U}_{\mathcal{T}^{(0)}}$ and $\mathcal{U}_{\mathcal{T}^{(?)}}$ denote the first batch of early-bird students and remaining unseen students, respectively. $\mathcal{V}_{\mathcal{T}}$ and $\mathcal{C}_{\mathcal{T}}$ are the sets of questions and concepts, respectively. The practice records of $\mathcal{U}_{\mathcal{T}^{(0)}}$ is available, denoted as $L_{\mathcal{T}^{(0)}} = \{(u, v, y_{u,v}) | y_{u,v} \in \{0, 1\}, u \in \mathcal{U}_{\mathcal{T}^{(0)}}, v \in \mathcal{V}_{\mathcal{T}}\}$. The student set of the target domain is the subset of all the source domains, but question sets on each domain (including source and target domains) are totally different. Besides, let $|\mathcal{U}|$ denote the number of set \mathcal{U} , and $\|\cdot\|$ denote ℓ_2 norm.

Based on the above setup, we aim to diagnose cognitive states for unseen cold-start students $\mathcal{U}_{\mathcal{T}^{(?)}}$ through fully exploiting available practice records (i.e., $L_{\mathcal{S}^k} \cup L_{\mathcal{T}^{(0)}}$) with student performance predictions.

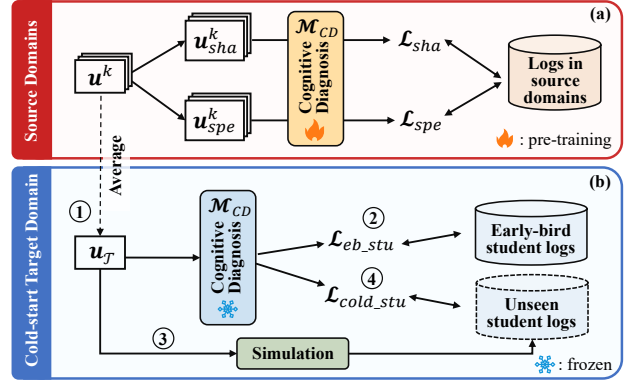


Figure 2: The main framework of Zero-1-to-3. (a) shows the pre-training stage in source domains with cognitive state decoupling. (b) is the adaptive diagnosis stage in the new domain, where ① ~ ④ denote execution sequence. ① is the initialization step, ② refines early-bird student states, ③ simulates virtual logs for unseen students. The simulated logs are used to fine-tune cold-start students with ④.

4 Methodology

Our framework contains a pre-training process with cognitive state decoupling (Figure 2 (a)) and an adaptive diagnosis with fine-tuning via simulation logs (Figure 2 (b)). Next, we introduce them starting from an embedding layer.

4.1 Embedding Layer

This layer offers initialized representations for students, questions and knowledge concepts in each domain. Specifically, it contains several parameter matrices as student embeddings, i.e., $\mathbf{U}^{\mathcal{U}_{\mathcal{S}^k}} \in \mathbb{R}^{|\mathcal{U}_{\mathcal{S}^k}| \times F}$ and $\mathbf{U}^{\mathcal{U}_{\mathcal{T}}} \in \mathbb{R}^{|\mathcal{U}_{\mathcal{T}}| \times F}$ for students in source domain \mathcal{S}^k and target domain \mathcal{T} , respectively, where F is the dimensional size. For the question, we adopt a pre-trained Bert (Devlin et al. 2018) to encode its textual content as initial representations by averaging its word-level embeddings. The question content-based embeddings in each domain includes $\mathbf{V}^{\mathcal{V}_{\mathcal{S}^k}} \in \mathbb{R}^{|\mathcal{V}_{\mathcal{S}^k}| \times F}$ and $\mathbf{V}^{\mathcal{V}_{\mathcal{T}}} \in \mathbb{R}^{|\mathcal{V}_{\mathcal{T}}| \times F}$. The concept embedding is obtained by averaging all relative question embeddings. We denote concept embeddings in source and target domains as $\mathbf{C}^{\mathcal{C}_{\mathcal{S}^k}} \in \mathbb{R}^{|\mathcal{C}_{\mathcal{S}^k}| \times F}$ and $\mathbf{C}^{\mathcal{C}_{\mathcal{T}}} \in \mathbb{R}^{|\mathcal{C}_{\mathcal{T}}| \times F}$, respectively. Note that the content-based encoder can cope with cold-start questions/concepts well by encoding their semantics (Liu et al. 2019). Our focus falls into student-side cross-domain transfer.

4.2 Cognitive State Decoupling

During this stage, a CDM is pre-trained across multiple source domains to establish an initial profile of students' states. To ensure effective cross-domain transfer, the model should extract domain-shared cognitive representations from the input student embeddings within source domains. However, prior cross-domain CDMs struggle to differentiate between general and specific signals, potentially hindering the effectiveness of cross-domain transfer. Inspired by the success of decoupling learning in various

fields (Wang et al. 2023; Liu et al. 2023b), we propose to decouple student states into domain-shared and domain-specific parts. The initial decoupled embeddings for students in each source domain can be obtained as follows:

$$\mathbf{u}_{sha}^k = f_{sha}(\mathbf{u}^k), \mathbf{u}_{spe}^k = f_{spe}(\mathbf{u}^k),$$

where \mathbf{u}^k is u -th row of $\mathbf{U}^{\mathcal{S}^k}$ denoting input trait for student u in source domain \mathcal{S}^k from embedding layer. The input can be divided into: \mathbf{u}_{sha}^k containing domain-shared states while \mathbf{u}_{spe}^k emphasizing domain-specific states. $f_{sha}(\cdot)$ and $f_{spe}(\cdot)$ are cross-domain shared full connection layers with the same input and output dimensions.

Then, we devise two mathematical regularizers to optimize these two types of signals separately, inspired by (Zhao et al. 2017; Chen et al. 2018, 2022). Unlike prior decoupling methods, our insight is rooted in educational field-specific analysis, thoughtfully tailored for the DZCD task.

Domain-shared states reveal general cognitive preference across domains, providing a panoramic perspective for inferring student preferences. Leveraging these shared states can notably enhance the cross-domain transfer capability of CDMs. However, these broad clues cannot offer a fine peek into student preference in each individual domain as they are blocked by domain-specific information. Hence, domain-shared states should fulfill the following requirement:

Requirement 1 Suppose a CDM has obtained the domain-shared student representations from a domain \mathcal{S}^k . It may exhibit strong overall predictions across source domains. Besides, its performance within \mathcal{S}^k is hindered by its absence of domain-specific information.

The above requirement can be approximatively expressed through the following regularization function:

$$\begin{aligned} \mathcal{L}_{sha} = & \sum_{k=1}^M \left(\mathbb{E}_{L_{\mathcal{S}^i}, i \in \{1, \dots, M\}} \|y_{u,v}^i - \mathcal{M}_{CD}(\mathbf{u}_{sha}^k, \mathbf{v}^i, \mathbf{C}^i)\| \right. \\ & \left. - \sum_{(u^k, v^k, y_{u,v}^k) \in L_{\mathcal{S}^k}} \|y_{u,v}^k - \mathcal{M}_{CD}(\mathbf{u}_{sha}^k, \mathbf{v}^k, \mathbf{C}^k)\| \right), \end{aligned}$$

where the first term minimizes the expectation of global prediction errors while the second term deliberately undermines prediction performance within a specific domain.

Domain-specific states contain student preferences for unique and specialized knowledge within the relevant learning topic, often supplying richer in-domain insights. Logically, by harnessing these specific cues within a particular domain, CDMs can achieve enhanced diagnostic performance as these clues hold significant in-domain value. However, this unique information might compromise the diagnostic accuracy across other domains, as it is typically irrelevant to unrelated domains. Hereby, the domain-specific cognitive states should satisfy the following requirement:

Requirement 2 Assume a CDM has obtained the domain-specific student states from \mathcal{S}^k . It probably leads to impressive predictive performance within domain \mathcal{S}^k . Conversely, prediction accuracy is expected to be suboptimal when the model is employed in other domains.

The above requirement can be further expressed through the minimization of the following regularization objective:

$$\begin{aligned} \mathcal{L}_{spe} = & \sum_{k=1}^M \sum_{(u^k, v^k, y_{u,v}^k) \in L_{\mathcal{S}^k}} \left(\|y_{u,v}^k - \mathcal{M}_{CD}(\mathbf{u}_{spe}^k, \mathbf{v}^k, \mathbf{C}^k)\| \right. \\ & \left. - \sum_{i \in \{1, \dots, M\} \setminus \{k\}} \|y_{u,v}^k - \mathcal{M}_{CD}(\mathbf{u}_{spe}^i, \mathbf{v}^k, \mathbf{C}^k)\| \right), \end{aligned}$$

where the first term narrows the gap between the actual performance $y_{u,v}^k$ and the prediction using the specific state of student u^k on question v^k from domain \mathcal{S}^k . The second term encourages misleading predictions by using domain-specific student states from other domains.

To summarize, the pre-training stage is directed by two regularizers for cognitive state decoupling as: $\mathcal{L}_{dec} = \mathcal{L}_{spe} + \mathcal{L}_{sha}$. The shared cognitive signals from different domains can offer broad experiences that are useful when students encounter new areas, making it easier to transfer knowledge between domains. Consequently, this process can preserve the *cognitive signal propagation* objective.

4.3 Domain-adaptive Cognitive Diagnosis

After completing pre-training, our focus shifts towards performing domain-adaptive cognitive diagnosis for unseen students. A simple yet effective solution is to directly use domain-shared states as initial student representations for DZCD, which has been used in previous studies (Gao et al. 2023). However, we contend that it is not domain-adaptive as it ignores in-domain considerations. For this goal, we devise a strategy to generate simulated practice logs for unfamiliar students by fully utilizing available early-bird students. Based on simulated practice logs, we can warm up student states by fine-tuning within the new domain.

Specifically, given a pre-trained CDM \mathcal{M}_{CD} and decoupled embeddings of each student u , i.e., $\{\mathbf{u}_{sha}^k, \mathbf{u}_{spe}^k\}_{k=1}^M$, we freeze model parameters and initialize the state of each student $u \in \mathcal{U}_{\mathcal{T}}$ using his/her domain-shared representations from source domains inspired by the related techniques (Yang et al. 2017; Yue et al. 2023) as follows:

$$\mathbf{u}_{\mathcal{T}} = \frac{1}{M} \sum_{k=1}^M \mathbf{u}_{sha}^k, \quad (1)$$

where an average pooling is adopted to merge shared signals to obtain initial states $\mathbf{u}_{\mathcal{T}}$ of each student u across M source domains, since this operation can augment representations compared to ones in a single domain and smooth the biases across domains (Wang et al. 2021; Zhu et al. 2023). The domain-shared states are transferable, which can provide valuable priors.

Based on the setup, we utilize practice logs of the early-bird student $u \in \mathcal{U}_{\mathcal{T}(0)}$ to refine their cognitive states, aiming to establish an initial understanding for the new domain. This process is optimized by minimizing the difference between prediction $\mathcal{M}_{CD}(\cdot)$ and actual response y_{uv} as:

$$\mathcal{L}_{eb-stu} = \sum_{(u, v, y_{u,v}) \in L_{\mathcal{T}(0)}} \|y_{u,v} - \mathcal{M}_{CD}(\mathbf{u}_{\mathcal{T}}, \mathbf{v}_{\mathcal{T}}, \mathbf{C}^{\mathcal{T}})\|.$$

Next, we aim to generate simulated practice logs for remaining unseen students by exploiting in-domain clues from early-bird students. Our basic assumption is that students with similar cognitive preferences are likely to achieve similar practice performance in a certain domain (Long et al. 2022; Yin et al. 2023). Thus, leverage the cognitive similarity between early-bird and unseen students to transfer practice patterns from the former to the latter, resulting in synthesized data. With this in mind, we introduce a strategy to extract students’ similarities from source domains as available clues in the target domain are very few.

In detail, for each early-bird student u , we first find a reference source domain S^k through the alignment of their refined embeddings in the target domain and domain-specific states originating from each source domain as follows:

$$S^k = \arg \max_{k \in \{1, \dots, M\}} \text{sim} \left(\mathbf{u}_{\mathcal{T}}, \mathbf{u}_{spe}^k \right), \quad (2)$$

where $\text{sim}(\cdot, \cdot)$ denotes the similarity function, for which we use cosine similarity. S^k is selected as a reference domain as student u has the most similar states between both S^k and the target domain. Within S^k , we compute similarity score between student u and each unseen student i by using their domain-specific representations as follows:

$$s_{u,i} \in \text{sim} \left(\mathbf{u}_{spe}^k, \mathbf{i}_{spe}^k \right), \quad (3)$$

where \mathbf{u}_{spe}^k and \mathbf{i}_{spe}^k are domain-specific states pre-trained in S^k of early-bird student u and each unseen student i . Then, we rank these unseen students via similarity scores and select top p individuals to form a similar peer set for student u , denoted as \mathcal{P}_u . Subsequently, for each selected unseen student in \mathcal{P}_u , denoted as i , their virtual data can be approximatively generated by duplicating practice records of u in the target domain as their performances are similar:

$$L_{\mathcal{P}_i} \leftarrow L_{\mathcal{T}_u} \subset L_{\mathcal{T}^{(0)}}. \quad (4)$$

By repeating the above processes (Eq. (2) - (4)) for each early-bird student, we can generate practice logs for parts of unseen students in the target domain. Note that the above simulation process might not cover every cold-start student, a concern linked to the parameter p . While a larger p could involve more students, such an approach is not recommended due to the potential introduction of noise. Based on simulation data, denoted as $L_{\mathcal{P}}$, we fine-tune the cognitive states of unseen students in the new domain as:

$$\mathcal{L}_{cold_stu} = \sum_{(u,v,y_{u,v}) \in L_{\mathcal{P}}} \left\| y_{u,v} - \mathcal{M}_{CD}(\mathbf{u}_{\mathcal{T}}, \mathbf{v}_{\mathcal{T}}, \mathbf{C}^{\mathcal{T}}) \right\|. \quad (5)$$

Through optimization with the above loss, the cognitive embeddings of cold-start students can be augmented, ensuring our framework is *diagnosis-oriented*. Concurrently, this stage maximizes the utilization of domain-specific cues from the perspective of early-bird students, aligning with the *domain-adaptation* objective.

Datasets	#students	#questions	#concepts	#logs
Geometry	15,283	2,299	251	127,570
Function	15,404	2,172	201	121,512
Probability	4,076	246	32	10,237
Physics	13,369	2,699	552	146,326
Arithmetic	14,073	1,828	200	86,699
English	5,906	409	135	24,739
Total	21,068	9,653	1,371	517,083

Table 1: Some basic statistics of the datasets.

5 Experiments

We conduct comprehensive experiments to address the following research questions:

- **RQ1** How powerful is Zero-1-to-3 for the DZCD task?
- **RQ2** How effective are the key components of the Zero-1-to-3 model?
- **RQ3** Can the Zero-1-to-3 perform cognitive diagnosis well beyond the cold-start stage?
- **RQ4** How to apply our framework to provide personalized learning guidance?

5.1 Datasets

We conduct experiments on six real-world datasets i.e., Geometry, Function, Probability, Physics, Arithmetic and English, which are collected from the iFLYTEK Learning Machine¹. All the datasets provide student practice records, question textual contents and question-concept correlations, where each question associates one knowledge concept. For each dataset, we reserve only the first attempt of each question to ensure that student states are static following the (Wang et al. 2020). Each dataset is treated as a domain. We switch their roles, each acting as the cold-start target domain and leaving the other five as the source domains. We split each source domain by randomly selecting two historical interactions from each student’s logs for validation, with the remaining data serving as the training set, similar to the widely used *leave-one-out* evaluation (Rendle et al. 2009). We randomly select some students (reported in section 5.2) as early birds (the order in which student commence new domains does not influence one another) to introduce data diversity for each domain when acting as the target domain. Besides, to train the Oracle models (section 5.2), we also split the target domain’s dataset into training (70%), validation (10%), and test sets (20%). The basic statistics of datasets are listed in Table 1.

5.2 Experimental Setup

Baselines To verify the effectiveness of our model, we apply our framework to three well-adopted CDMs, i.e., IRT (Embretson and Reise 2013), MIRT (Reckase 2009) and NeuralCD (Wang et al. 2020) (introduced in section 3.1). We call them Zero-CDM, e.g., Zero-IRT. We select several baselines for comparison. Among them, the Random and

¹<https://xxj.xunfei.cn/>

Methods		Geometry as Target			Function as Target			Probability as Target		
		ACC \uparrow	AUC \uparrow	RMSE \downarrow	ACC \uparrow	AUC \uparrow	RMSE \downarrow	ACC \uparrow	AUC \uparrow	RMSE \downarrow
Random		49.91	49.83	57.80	49.97	50.04	57.76	50.17	50.96	57.60
IRT	Oracle	73.57*	79.86*	44.54*	74.16*	80.44*	42.71*	70.00*	72.41*	45.81*
	NLP	62.27	<u>64.93</u>	48.62	<u>55.38</u>	61.92	49.15	58.51	65.28	48.73
	GCN	<u>60.49</u>	<u>62.92</u>	51.62	<u>54.23</u>	56.67	52.12	58.16	57.02	51.36
	Tech	<u>59.96</u>	62.50	49.86	52.45	59.58	55.13	<u>65.22</u>	54.24	54.58
	Ours	57.01	64.96	<u>49.51</u>	55.73	<u>60.96</u>	<u>51.95</u>	65.37	<u>58.30</u>	<u>50.47</u>
MIRT	Oracle	73.53*	80.57*	42.35*	74.03*	80.92*	41.82*	70.32*	72.76*	44.78*
	NLP	56.74	68.27	52.15	55.32	<u>68.00</u>	55.22	<u>65.67</u>	<u>64.62</u>	<u>47.77</u>
	GCN	56.50	59.20	51.95	<u>59.39</u>	63.80	<u>50.95</u>	65.22	58.19	50.48
	Tech	<u>56.66</u>	63.02	49.33	58.82	61.32	53.98	62.22	60.26	51.50
	Ours	51.48	<u>65.98</u>	<u>50.56</u>	60.91	68.01	49.36	68.27	70.46	45.60
NeuralCD	Oracle	73.62*	81.10*	41.95*	74.27*	80.97*	41.99*	72.26*	75.25*	41.81*
	NLP	58.66	58.16	49.05	<u>58.90</u>	57.85	49.31	64.36	50.28	48.73
	GCN	56.84	50.10	<u>49.74</u>	55.00	51.12	54.30	64.92	52.89	49.40
	Tech	<u>60.20</u>	<u>60.23</u>	53.44	58.33	52.83	52.63	<u>65.22</u>	<u>53.53</u>	57.23
	Ours	61.48	60.30	52.58	59.78	<u>55.11</u>	<u>52.04</u>	66.67	67.43	48.68
Methods		Physics as Target			Arithmetic as Target			English as Target		
		ACC \uparrow	AUC \uparrow	RMSE \downarrow	ACC \uparrow	AUC \uparrow	RMSE \downarrow	ACC \uparrow	AUC \uparrow	RMSE \downarrow
Random		49.94	49.99	57.75	49.99	50.03	57.69	50.30	50.23	57.58
IRT	Oracle	70.46*	77.49*	43.91*	69.57*	69.14*	48.48*	74.37*	81.78*	40.96*
	NLP	<u>56.35</u>	59.63	49.51	60.66	60.78	<u>50.83</u>	64.03	66.75	48.37
	GCN	<u>54.32</u>	52.08	<u>52.03</u>	54.04	51.12	<u>51.03</u>	54.34	50.90	55.13
	Tech	56.17	<u>57.53</u>	<u>56.93</u>	55.44	52.01	51.20	<u>56.73</u>	55.76	51.03
	Ours	56.43	52.28	56.52	<u>56.23</u>	<u>52.08</u>	50.72	54.46	<u>58.28</u>	<u>50.07</u>
MIRT	Oracle	70.43*	77.45*	44.50*	74.60*	81.10*	42.68*	72.80*	81.01*	42.54*
	NLP	54.32	<u>61.18</u>	54.96	60.01	66.12	50.77	50.69	<u>58.34</u>	56.75
	GCN	<u>54.35</u>	<u>58.17</u>	57.52	<u>60.53</u>	60.02	<u>50.60</u>	50.47	51.15	54.41
	Tech	53.33	60.31	<u>50.24</u>	59.93	54.35	55.59	<u>54.86</u>	55.08	<u>53.40</u>
	Ours	56.71	72.06	48.71	61.88	<u>63.71</u>	48.52	61.52	66.17	48.24
NeuralCD	Oracle	70.22*	77.72*	43.94*	74.45*	81.80*	41.22*	72.64*	81.53*	41.98*
	NLP	<u>56.32</u>	<u>56.45</u>	51.42	<u>60.63</u>	57.95	49.43	50.47	50.48	51.12
	GCN	<u>54.21</u>	54.54	<u>50.26</u>	<u>60.44</u>	<u>58.40</u>	52.68	50.47	52.15	50.51
	Tech	54.23	52.18	<u>50.81</u>	56.99	52.40	<u>49.57</u>	<u>57.93</u>	<u>56.12</u>	<u>50.10</u>
	Ours	61.59	68.08	50.22	60.89	58.92	51.99	60.87	60.23	48.96

Table 2: Performance comparison (%). The best zero-shot performance is highlighted in bold, and the runner-up is underlined. \uparrow (\downarrow) means the higher (lower) score the better (worse) performance, the same as below. * indicates the oracle result.

Oracle methods indicate the lower and upper bounds of performance. For each baseline (excluding Random), we also select IRT, MIRT and NeuralCD as diagnostic functions.

- Random: The Random method predicts the students' scores randomly from $Uniform(0, 1)$.
- Oracle: The Oracle baseline is trained with logs from both source and target domains using the traditional CDM.
- NLP-based (Liu et al. 2019): The NLP-based method uses learnable embeddings as student states in source domains and represents questions by encoding their texts. To implement it, we adopt Bert (Devlin et al. 2018) as the textual encoder following the setup in (Gao et al. 2023).

- GCN-based and Tech-based (Gao et al. 2023): Both these methods use a knowledge graph (Yang et al. 2023) linking each domain for transfer. The graph is constructed using a statistical method proposed in RCD (Gao et al. 2021).

Evaluation As cognitive states cannot be directly observed in practice, it is common to indirectly assess CDMs through predicting student performance on validation datasets. To evaluate prediction performance, we adopt ACC and AUC, and RMSE as metrics from the perspectives of classification and regression, respectively, following previous works (Gao et al. 2021; Zhang et al. 2023).

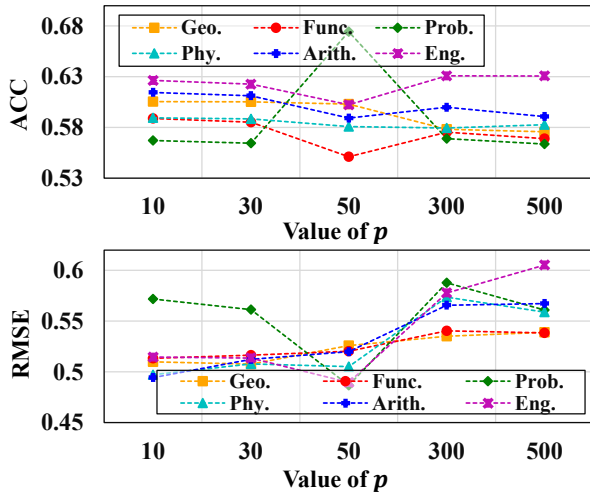


Figure 3: Performance under different peer student numbers.

Implementation Details We set the dimensions of student and question vectors, i.e., parameter F , as the total number of knowledge concepts of each domain. The dimensions of neural network layers are 512 and 256 for all NeuralCD-based models. The number of early-bird students ($|\mathcal{U}_{\mathcal{T}^{(0)}}|$) for each domain is set to be 0.01 of the number of students in the domain, and the number of unseen peer students (p) for the overall results of student performance prediction (RQ1) is 50. We set the mini-batch size as 256 and select learning rate from $\{0.001, 0.002, 0.02, 0.05\}$ for each model. Each experiment is repeated five times under consistent conditions and the best score is reported. Each model is implemented by PyTorch and optimized by Adam optimizer (Kingma and Ba 2014). All experiments are run on a Linux server with two 3.00GHz Intel Xeon Gold 5317 CPUs and one Tesla A100 GPU. The code is publicly available at <https://github.com/bigdata-ustc/Zero-1-to-3>.

5.3 Experimental Results

Student Performance Prediction (RQ1) We compare our model with several baselines on student performance predictions under DZCD. We switch the role of each dataset acting as the target domain. The overall performance is reported in Table 2. We obtain that: (1) For different diagnostic implementations (i.e., IRT, MIRT and NeuralCD as diagnostic models), our Zero-1-to-3 framework almost outperforms all baselines on each target domain, which indicates the diagnostic effectiveness of our solution under DZCD scenarios. (2) The most significant distinction between Zero-1-to-3 and NLP methods lies in our consideration of decoupling student states and only transferring domain-shared states to address the “what to transfer” issue. Our method outperforms NLP methods in most cases, underscoring the significance of adeptly capturing the domain-shared cognitive signals among students. (3) Both GCN-based and Tech-based models employ a knowledge graph linking both source and target domains for domain-adaption by joint training. However, They cannot fully utilize domain-specific student logs.




Figure 4: T-SNE visualization of student states.




Figure 5: Performance in normal diagnosis scenarios.

In contrast, Zero-1-to-3 outperforms these methods, which confirms its effectiveness.

In the following parts, we primarily present the experimental results of Zero-NeuralCD as the representative ones, since other diagnosis models can be abstracted as the special cases of NeuralCD (Wang et al. 2020).

Detailed Analysis (RQ2) This section provides some in-depth analysis of how the key component in Zero-1-to-3 contributes to solving the challenges of DZCD.

Exploration of Peer Student Number. The sampling number of early-bird students’ peer students who are unseen (p) plays a crucial role in the transfer. To study the impact of different numbers, we train Zero-NeuralCD under several settings and perform zero-shot student performance predictions on each dataset. We randomly sample the number of peer students from $\{10, 30, 50, 300, 500\}$. Figure 3 presents

the results under different settings. From the figure, we can see the diagnostic effect does not consistently improve with the increasing number of peer students during the simulation process, indicating that the violent simulation data will bring noise. We are motivated to explore more suitable methods for matching peer students in future research.

Visualization. This part visualizes student cognitive states to observe decoupling effects during pre-training. Under the setting of “English as Target”, for the pre-trained Zero-NeuralCD (without fine-tuning), we randomly select 200 students and visualize their domain-specific/-shared states in each source domain using T-SNE (Van der Maaten and Hinton 2008) in Figure 4. We can observe that domain-specific student states are more distant from each other compared to domain-shared student states, as they encompass more unique insights. Furthermore, domain-shared student states are not entirely blended, and we infer that this is because they still retain some personalized cognitive features that contribute to diagnostic performance. These findings substantiate the efficacy of decoupling regularizers.

5.4 Normal Diagnosis (RQ3)

The above examples show Zero-1-to-3 can be successfully applied in the DZCD scenarios. A subsequent question is whether it can effectively perform cognitive diagnosis after the cold-start stage. Thus, we conduct experiments by fine-tuning a Zero-NeuralCD (named Zero-NeuralCD-Normal) under the oracle setting where 70% data in the target domain can be used for training. The results under “Geometry as Target” in Figure 5 demonstrate a significant improvement in the performance of the model fine-tuned in a normal scenario compared to the model in a cold-start environment. This observation reveals that our model not only performs well in DZCD scenarios but also achieves satisfactory results in the subsequent stage, compared to Oracle models.

5.5 Question Recommendation (RQ4)

The above experiments have proved that Zero-1-to-3 can complete the DZCD task effectively. This part demonstrates one of the most popular diagnostic applications, i.e., question recommendation, that are in need of industrial practice.

We implement a simple yet effective recommendation strategy as an example to recommend x questions for the student under DZCD. Generally, a proper recommendation should not be too hard or easy to maintain students’ enthusiasm when practicing (Huang et al. 2019). Thus, with a refined CDM, we first predict each student’s performance on questions in the new domain via the $M_{CD}(\cdot)$ in Eq. (5). All questions can be divided into two sets that answer correctly or not (i.e., positive or negative samples) according to prediction results. Then, we sample $\frac{x}{2}$ questions from each type of sample, respectively, to yield the recommendation list.

Table 3 lists 6 recommended questions on target domain Geometry for a randomly selected student using a refined Zero-NeuralCD, the diagnosed student mastery levels and question difficulties of the associated concepts, and the student’s true performance on the questions as recorded in the Geometry dataset. We can see the recommended questions are tailored to the student’s proficiency, neither too easy nor

Geometry as Target	①	②	③	④	⑤	⑥
Question id	3,213	200	1,032	2,122	3,013	32
Mastery (%)	40.07	32.00	48.63	26.30	54.33	41.02
Difficulty (%)	41.23	40.04	44.34	30.10	52.01	39.99
True performance	×	×	✓	×	✓	✓

Table 3: Question recommendation via Zero-NeuralCD.

too difficult. Some of them will challenge the student, while others will serve as “gifts” that can help increase his/her engagement. It confirms the application potential of Zero-1-to-3 in cold-start scenarios.

6 Conclusion

In this paper, we have proposed a general Zero-1-to-3 framework to tackle the real-world challenge of domain-level zero-shot cognitive diagnosis (DZCD). The central obstacle in DZCD involves extracting shared information to facilitate cross-domain transfer, while simultaneously adapting to new domains. To address this, we first pre-train a diagnosis model with dual regularizers that disentangle student states into domain-shared and domain-specific parts. These shared cognitive signals can be transferred to the target domain, enriching cognitive priors for the new domain. Subsequently, we develop a strategy to generate simulated practice logs for cold-start students in the target domain by using the behavioral patterns of early-bird students. Consequently, the cognitive states of cold-start students can be adaptively refined using virtual data from the target domain, enabling the execution of DZCD. Finally, extensive experiments highlight the efficacy and potential applicability of our framework.

Acknowledgments

This research was partially supported by grants from the National Natural Science Foundation of China (Grant No. 62337001 and No. 62202443), the Anhui Provincial Natural Science Foundation (No. 2308085MG226), and the Fundamental Research Funds for the Central Universities.

References

- Bi, H.; Chen, E.; He, W.; Wu, H.; Zhao, W.; Wang, S.; and Wu, J. 2023. BETA-CD: A Bayesian Meta-Learned Cognitive Diagnosis Framework for Personalized Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 5018–5026.
- Chen, P.; Lu, Y.; Zheng, V. W.; and Pian, Y. 2018. Prerequisite-driven deep knowledge tracing. In *2018 IEEE International Conference on Data Mining (ICDM)*, 39–48. IEEE.
- Chen, X.; Liu, T.; Zhao, H.; Zhou, G.; and Zhang, Y.-Q. 2022. Cerberus transformer: Joint semantic, affordance and attribute parsing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19649–19658.
- Chen, X.; Wu, L.; Liu, F.; Chen, L.; Zhang, K.; Hong, R.; and Wang, M. 2023. Disentangling Cognitive Diagnosis

- with Limited Exercise Labels. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- De La Torre, J. 2009. DINA model and parameter estimation: A didactic. *Journal of educational and behavioral statistics*, 34(1): 115–130.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Embretson, S. E.; and Reise, S. P. 2013. *Item response theory*. Psychology Press.
- Gao, W.; Liu, Q.; Huang, Z.; Yin, Y.; Bi, H.; Wang, M.-C.; Ma, J.; Wang, S.; and Su, Y. 2021. Rcd: Relation map driven cognitive diagnosis for intelligent education systems. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 501–510.
- Gao, W.; Wang, H.; Liu, Q.; Wang, F.; Lin, X.; Yue, L.; Zhang, Z.; Lv, R.; and Wang, S. 2023. Leveraging Transferable Knowledge Concept Graph Embedding for Cold-Start Cognitive Diagnosis. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 983–992.
- Hu, G.; Zhang, Y.; and Yang, Q. 2018. Conet: Collaborative cross networks for cross-domain recommendation. In *Proceedings of the 27th ACM international conference on information and knowledge management*, 667–676.
- Huang, Z.; Liu, Q.; Zhai, C.; Yin, Y.; Chen, E.; Gao, W.; and Hu, G. 2019. Exploring multi-objective exercise recommendations in online education systems. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 1261–1270.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Liu, F.; Hu, X.; Liu, S.; Bu, C.; and Wu, L. 2023a. Meta Multi-Agent Exercise Recommendation: A Game Application Perspective.
- Liu, Q.; Huang, Z.; Yin, Y.; Chen, E.; Xiong, H.; Su, Y.; and Hu, G. 2019. Ekt: Exercise-aware knowledge tracing for student performance prediction. *IEEE Transactions on Knowledge and Data Engineering*, 33(1): 100–115.
- Liu, S.; Yu, X.; Ma, H.; Wang, Z.; Qin, C.; and Zhang, X. 2023b. Homogeneous Cohort-Aware Group Cognitive Diagnosis: A Multi-grained Modeling Perspective. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 4094–4098.
- Long, T.; Qin, J.; Shen, J.; Zhang, W.; Xia, W.; Tang, R.; He, X.; and Yu, Y. 2022. Improving knowledge tracing with collaborative information. In *Proceedings of the fifteenth ACM international conference on web search and data mining*, 599–607.
- Nguyen, T. 2015. The effectiveness of online learning: Beyond no significant difference and future horizons. *MERLOT Journal of online learning and teaching*, 11(2): 309–319.
- Reckase, M. D. 2009. Multidimensional item response theory models. In *Multidimensional item response theory*, 79–112. Springer.
- Rendle, S.; Freudenthaler, C.; Gantner, Z.; and Schmidt-Thieme, L. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, 452–461.
- Schmucker, R.; and Mitchell, T. M. 2022. Transferable Student Performance Modeling for Intelligent Tutoring Systems. *arXiv preprint arXiv:2202.03980*.
- Tong, S.; Liu, J.; Hong, Y.; Huang, Z.; Wu, L.; Liu, Q.; Huang, W.; Chen, E.; and Zhang, D. 2022. Incremental Cognitive Diagnosis for Intelligent Education. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1760–1770.
- Tsutsuni, E.; Kinoshita, R.; and Ueno, M. 2021. Deep-IRT with Independent Student and Item Networks. *International Educational Data Mining Society*.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- Wang, C.; Zhu, Y.; Sun, A.; Wang, Z.; and Wang, K. 2023. A Preference Learning Decoupling Framework for User Cold-Start Recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1168–1177.
- Wang, F.; Liu, Q.; Chen, E.; Huang, Z.; Chen, Y.; Yin, Y.; Huang, Z.; and Wang, S. 2020. Neural cognitive diagnosis for intelligent education systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 6153–6161.
- Wang, H.; Lian, D.; Tong, H.; Liu, Q.; Huang, Z.; and Chen, E. 2021. Hypersorec: Exploiting hyperbolic user and item representations with multiple aspects for social-aware recommendation. *ACM Transactions on Information Systems (TOIS)*, 40(2): 1–28.
- Yang, Y.; Yang, J.; Bao, R.; Zhan, D.; Zhu, H.; Gao, X.; Xiong, H.; and Yang, J. 2023. Corporate Relative Valuation Using Heterogeneous Multi-Modal Graph Neural Network. *IEEE Trans. Knowl. Data Eng.*, 35(1): 211–224.
- Yang, Y.; Zhan, D.; Fan, Y.; Jiang, Y.; and Zhou, Z. 2017. Deep Learning for Fixed Model Reuse. In *Proceedings of the 31 Conference on Artificial Intelligence*, 2831–2837. San Francisco, California.
- Yao, F.; Liu, Q.; Hou, M.; Tong, S.; Huang, Z.; Chen, E.; Sha, J.; and Wang, S. 2023. Exploiting non-interactive exercises in cognitive diagnosis. *Interaction*, 100(200): 300.
- Yin, M.; Wang, H.; Xu, X.; Wu, L.; Zhao, S.; Guo, W.; Liu, Y.; Tang, R.; Lian, D.; and Chen, E. 2023. APGL4SR: A Generic Framework with Adaptive and Personalized Global Collaborative Information in Sequential Recommendation. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, CIKM '23*, 3009–3019. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701245.
- Yue, L.; Liu, Q.; Du, Y.; Gao, W.; Liu, Y.; and Yao, F. 2023. FedJudge: Federated Legal Large Language Model. *arXiv preprint arXiv:2309.08173*.
- Zhang, Z.; Liu, Q.; Jiang, H.; Wang, F.; Zhuang, Y.; Wu, L.; Gao, W.; and Chen, E. 2023. FairLISA: Fair User Modeling

with Limited Sensitive Attributes Information. In *Thirty-seventh Conference on Neural Information Processing Systems*.

Zhao, H.; Lu, M.; Yao, A.; Guo, Y.; Chen, Y.; and Zhang, L. 2017. Physics inspired optimization on semantic transfer features: An alternative method for room layout estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 10–18.

Zhu, J.; Wang, Y.; Zhu, F.; and Sun, Z. 2023. Domain Disentanglement with Interpolative Data Augmentation for Dual-Target Cross-Domain Recommendation. *arXiv preprint arXiv:2307.13910*.

Zhuang, Y.; Liu, Q.; Zhao, G.; Huang, Z.; Huang, W.; Pardos, Z.; Chen, E.; Wu, J.; and Li, X. 2023. A Bounded Ability Estimation for Computerized Adaptive Testing. In *Thirty-seventh Conference on Neural Information Processing Systems*.