

Explicit and Implicit Examinee-Question Relation Exploiting for Efficient Computerized Adaptive Testing

Changqian Wang¹, Shangshang Yang², Siyu Song¹, Ziwen Wang³,
Haiping Ma^{1,4*}, Xingyi Zhang³, Bo Jin^{5,6}

¹Institutes of Physical Science and Information Technology, Anhui University, China

²School of Artificial Intelligence, Anhui University, China

³School of Computer Science and Technology, Anhui University, China

⁴Department of Information Materials and Intelligent Sensing Laboratory of Anhui Province, China

⁵School of Innovation and Entrepreneurship, Dalian University of Technology, Dalian, China

⁶Guangdong Medical University, Zhanjiang, China

{changqian.wang.dl, yangshang0308, siyusong00, wzw12sir, xyzhanghust}@gmail.com, hpma@ahu.edu.cn, jinbo@dlut.edu.cn

Abstract

Computerized adaptive testing (CAT) is a crucial task in computer-aided education, which aims to adaptively select suitable questions to diagnose examinees' ability status. Existing CAT approaches enhance selection performance by exploring examinee-question (E-Q) relation. These approaches either exclusively utilize explicit E-Q relation. For instance, policy-based approaches determine question selection based on predefined criteria. While effective in adapting to changes in question banks, these methods often entail significant computational costs in searching for suitable questions. Conversely, some studies focus solely on implicit E-Q relation. For example, learning-based approaches train agents to efficiently select questions by learning from large-scale datasets. However, they may struggle with newly introduced questions. Additionally, most of these existing question selectors are based on greedy strategies, which potentially overlooks promising questions. To bridge the above two types of approaches, we propose a novel framework named **Relation Exploiting-based CAT (RECAT)** by exploring and exploiting the implicit and explicit examinee-question relation. Specifically, we first define an examinee true ability-oriented selection objective to select more suitable questions. Then, to learn the implicit E-Q relation, we design a question selector, which explores the examinee ability and generates best-fitting questions for specific examinee ability from two aspects, including generation consistency and knowledge matching. The former aims to maximize the likelihood estimation of the implicit E-Q relation learning process, while the latter is employed to fit the distribution of real questions. To fully exploit explicit E-Q relation, we generate a high-quality candidate set for the given examinee's ability using implicit E-Q relation, which streamlines the search process, minimizing selection latency. We demonstrate the effectiveness and efficiency of our framework through comprehensive experiments on real-world datasets.

Introduction

Amid the rapid advancement of computer-aided education, computerized adaptive testing (CAT) has been widely applied in real-world scenarios, such as educational measurement, psychological assessment, and talent recruitment (Wainer et al. 2000; Qin et al. 2023; Yang et al. 2024b; Yu et al. 2024a). CAT typically consists of two components: the Cognitive Diagnosis Model (CDM) (Ma et al. 2024c,a) and the question selection algorithm, which alternates until a termination condition is met. As illustrated in Figure 1, at time step t , the CDM estimates the examinee's ability based on his/her responses on the previous $t - 1$ steps. Subsequently, the question selection algorithm, guided by the diagnostic results, selects a suitable question for the examinee at step t (Ma et al. 2024b; Liu et al. 2024).

Question selection approaches, a core component of the CAT system, have been extensively studied due to their adaptive and personalized nature. Much research delved into understanding the examinees-questions (E-Q relation) (Lord 2012; Chang and Ying 1996). By analyzing the response label of a selected question, the diagnostic value contributed by such a question to the examinee's ability can be assessed. Some studies leverage this explicit E-Q relation to establish prior selection criteria, which are referred to as policy-based. For instance, MAAT (Bi et al. 2020) and BECAT (Zhuang et al. 2023) translate the impact of a question on examinee into model expectation and seek a subset that approximates the gradient of the complete question bank, respectively. Such policy-based approaches possess a theoretical foundation and interpretability, which also demonstrate good adaptability to the continuously modified question bank, wherein new questions are added.

Simultaneously, some studies (Li et al. 2023a; Gilavert and Freire 2022; Li et al. 2023b) have explored implicit E-Q relation embedded in previous examinee-question interactions and utilized these information to train selector for adaptive question selection and are referred to as learning-based. The representatives of this category involve transforming the CAT problem into a bilevel optimization problem (Ghosh and Lan 2021; Zhuang et al. 2022). In such a

*Corresponding Authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Illustration of CAT. The CDM diagnoses Bob’s ability based on his response records from the previous t steps. The selection algorithm then selects a suitable question for Bob based on the diagnostic results at step $t+1$.

framework, the outer-level optimization problem involves learning a data-driven selection approach, which is used to optimize the examinee’s ability in the inner-level optimization problem. This coupled learning process thoroughly explores implicit E-Q relation, especially in NCAT (Zhuang et al. 2022), which considers the presence of guessing and slipping in real scenarios. In addition, these learning-based approaches have lower selection latency. However, the prior selection criteria forces policy-based approaches to traverse almost the entire question bank at each testing step to fully utilize explicit E-Q relation. This results in high computational costs and prolonged selection latency (Hong et al. 2023). Learning-based selectors that leverage implicit E-Q relation can address these drawbacks. Consequently, they exhibit poor adaptability to the dynamically modified question bank (Ma et al. 2023b). When new questions are added to the bank, the entire model retraining is often required. This retraining process incurs additional costs, which policy-based approaches excel at. Therefore, an optimal strategy is to integrate the advantages of both types of approaches, fully exploiting both explicit and implicit E-Q relationships. However, achieving this integration is challenging as existing learning-based approaches rely on Q-values and are incompatible with policy-based approaches. This gap results in selectors being unable to fully utilize both explicit and implicit E-Q relation simultaneously, which hinders the mining of examinee ability.

To bridge aforementioned approaches, we innovatively propose a novel framework, namely **Relation Exploiting-based CAT (RECAT)**, which explores and exploits implicit and explicit examinee-question relation. We first define an examinee true ability-oriented selection objective as a global directive to guide the exploration of both implicit and explicit E-Q relation. To explore implicit E-Q relation, we design a generative question selector that uncovers the examinee’s ability and generates the best-fitting question, facilitating participation in the post-optimization process. This selector is optimized from two perspectives, including generation consistency and knowledge matching. The former guides the quality of question generation, while the latter is used to fit the parameter distribution of real questions. To exploit explicit E-Q relation, we generate a high-quality candidate set based on implicit E-Q relation, which reduces selection latency and potential biases. Extensive real-world dataset experiments have demonstrated that our

proposed RECAT exhibits robust performance in question selection while significantly enhancing selection efficiency. Our code is available at <https://github.com/ChangqianWang/Intelligent-Education/tree/main/RECAT>.

Related Work

In this paper, we focus on the question selector component of CAT. Existing question selection approaches can be mainly divided into two categories: policy-based selection approaches that utilize explicit E-Q relation and learning-based selection approaches that learn implicit E-Q relation.

Policy-based Selection Approaches

Policy-based approaches explore selection results by designing a priori selection criteria based on explicit E-Q relation that do not require learning. For a long time, question selection approaches have been model-specific, which means that they are specifically designed for IRT (Lord 1980) or MIRT (Reckase 2009), such as maximum Fisher information (Lord 2012), Kullback-Leibler information index (Chang and Ying 1996), and their extended selection algorithms. However, due to the simple structure of these models, they cannot effectively capture the potential E-Q relation, which cannot meet the development of CDMs (Yang et al. 2024a; Liu et al. 2023; Yang, Qin, and Yu 2024; Yu et al. 2024b), and the performance of IRT-specific selection algorithms is also limited. To address this limitation, MAAT (Bi et al. 2020) first proposed a general model-agnostic selection algorithm, which calculates the expected model change (EMC) for each question as the explicit E-Q relation to measure the information content of the question, and selects the question with the largest EMC. MAAT can be adapted to different CDMs (Yu et al. 2024c), which greatly releases the performance of selection. BECAT (Zhuang et al. 2023) first provided a theoretical guarantee for question selection algorithms to approximate the true ability of examinees. The true ability of examinees is approximated by the diagnostic results of the examinee on the entire question bank. To fully exploit this explicit E-Q relation, BECAT defines CAT as a subset selection problem, which selects the subset so that the gradient of the responses on the subset approximates the gradient of all responses. These model-agnostic policy-based approaches typically have a theoretical foundation and exhibit good adaptability to the dynamically modified question bank. However, computing over the entire question bank can lead to high selection latency.

Learning-based Selection Approaches

Learning-based approaches make full use of the implicit E-Q relation embedded in prior examinee response records, which are utilized to train the selector, enabling an adaptive selection process. For this aim, a series of data-driven learnable approaches (Gilavert and Freire 2022; Li et al. 2023b; Ma et al. 2023c) have been proposed. BOBCAT (Ghosh and Lan 2021) first defined the CAT problem as a bilevel optimization problem. The outer optimization problem aims to learn the selector using the examinee’s estimated ability,

while the inner optimization problem optimizes the examinee’s ability using the questions selected by the outer selector. Along this line, NCAT (Zhuang et al. 2022) took into account the fact that examinees may guess or slip in real-world scenarios, and designed a deep neural network-based framework to capture the implicit E-Q relation. GMOCAT (Wang et al. 2023) conducts a fine-grained analysis of implicit E-Q relation, which considered that the diversity of knowledge concepts can make E-Q relation more coupled. Therefore, it captures these complex E-Q relations through a graph structure. Learning-based approaches reduce selection latency, but they exhibit poor adaptability to new questions added to the bank, which leads to high training costs due to the entire model retraining typically required.

Preliminaries and Problem Formulation

Preliminaries of CAT

Task Introduction For an online intelligent education platform, the CAT system is responsible for accurately revealing examinees’ true ability values by continuously selecting a question for examinees. The CAT system comprises two components: (1) a question selection algorithm Π , used to select the “best-fitting” question from the question bank as the next question for the examinee based on his/her previously estimated ability; (2) a cognitive diagnosis model (CDM) \mathcal{M} , used to diagnose the current examinee’s ability after the examinee practices the question selected by the selection algorithm Π and gets the response.

Specifically, given a examinee s_i and his true ability scalar or vector θ_i , the CAT process at each step can be denoted as

$$\begin{aligned} \min \|\theta_i^t - \theta_i\|, \quad t = 1, \dots, T, \\ \theta_i^t \leftarrow \mathcal{M}(q^t, r_{iq^t} \mid \theta_i^{t-1}), \quad q^t \leftarrow \Pi(Q_i \mid \theta_i^{t-1}), \end{aligned} \quad (1)$$

q^t denotes the question selected by Π from the question bank Q_i based on his estimated ability θ_i^{t-1} at step $t - 1$, where the question bank Q_i only contains the questions that examinee s_i has not practiced. r_{iq^t} is the response of examinee s_i on question q^t . θ_i^t is the ability of examinee s_i estimated by \mathcal{M} after practicing question q^t , where θ_i^0 is usually randomized. Note that T is the maximal number of steps for stopping the testing, and another common stopping criterion is to predefine a parameter ϵ to ensure $\|\theta_i^T - \theta_i\| \leq \epsilon$.

The Best-Fitting Question To unify the definition of the “best fit” problem in different CATs, we propose a novel definition. Specifically, the devised definition adopts a direct measure manner to select a question at each step, i.e., selecting the question that maximizes the gap improvement between the examinee’s true ability and estimated ability, which can be roughly formulated as

$$q^t = \arg \min_{q \in Q_i} \|\theta_i^t - \theta_i\|. \quad (2)$$

Such a definition directly adopts the ability gap improvement as the question selection measure to find the “best-fitting” question, and more details will be presented later.

Training and Testing Phases In a given platform, suppose there are N examinees and L questions, which can be denoted by $S = \{S_{train}, S_{val}, S_{test}\} = \{s_1, s_2, \dots, s_N\}$ and $Q = \{q_1, q_2, \dots, q_L\}$, respectively. For each examinee, their response records are denoted as $\mathcal{R} = \{R_1, \dots, R_N\}$, where $R_i = \{(s_i, q_j, r_{ij}) \mid q_j \in Q - Q_i\}$ is the record set of examinee s_i . Here $r_{ij} = 1$ means examinee s_i correctly answered question q_j and $r_{ij} = 0$ otherwise.

There are commonly two phases in the CAT process, i.e., the training phase to train the question selection algorithm based on the response records of examinees in S_{train} if needed, and the testing phase to test new examinees in S_{test} .

Examinees’ True Ability We follow prior work BECAT that the examinees’ true ability θ_i can be approximated by the diagnostic result θ_i^* over the full question bank, which can be formalized as

$$\theta_i \approx \theta_i^* = \arg \min_{\theta \in \Theta} \sum_{q_j \in E - Q_i} \mathcal{L}(r_{ij}, \mathcal{M}(q_j \mid \theta)). \quad (3)$$

$\mathcal{L}(r_{ij}, \mathcal{M}(q_j \mid \theta))$ is the loss between the true response r_{ij} and the predicted response on q_j under the ability θ , and Θ denotes the search space of θ . Before training the CAT, we generate simulated response logs for examinees based on the question bank (Gao et al. 2023; Yao et al. 2023). Once a sufficient number of responses are reached, the diagnosed abilities of the examinees approximate their true abilities (Zhuang et al. 2023), which are denoted by $\theta = \{\theta_1, \dots, \theta_N\} \approx \{\theta_1^*, \dots, \theta_N^*\}$. Besides, another set of important parameters, i.e., question-related parameters $Q^\omega = \{q_1^\omega, \dots, q_M^\omega\}$, can be also obtained, where each item $q_j^\omega \in Q^\omega$ is a vector with the length of embedding size d .

For easy reading, in subsequent sections, q_j represents not only the j -th question but also the $1 \times d$ vector representation of the question, determined by the context of using q_t .

Method

Best-Fitting Question Definition

To start with, the definition of the “best-fitting” question in RECAT needs to be clarified. The process of selecting a question for examinee s_i at step t can be formulated as the conditional probability $p(q^t \mid \theta_i^{t-1})$. Since the conditional probability cannot be directly obtained, existing CAT approaches adopt indirect manners to optimize the gap between the student’s true ability and estimated ability.

Different from them, this paper directly takes the gap between the student’s true ability and estimated ability as the criterion to select the question at each step. The definition of the “best-fitting” question for student s_i at step t in RECAT is formulated as

$$q^t = \arg \min_{q_j \in Q_i} \underbrace{\left\| \theta_i^* - \theta_i^{t-1} + \alpha \nabla_{\theta_i^{t-1}} \mathcal{L}(r_{ij}, \mathcal{M}(q_j \mid \theta_i^{t-1})) \right\|}_{\text{denoted as } \omega_{ij}}, \quad (4)$$

where θ_i^* is the approximation of true ability θ_i , α is the learning rate, $\mathcal{L}(\cdot)$ denotes the cross-entropy loss, and

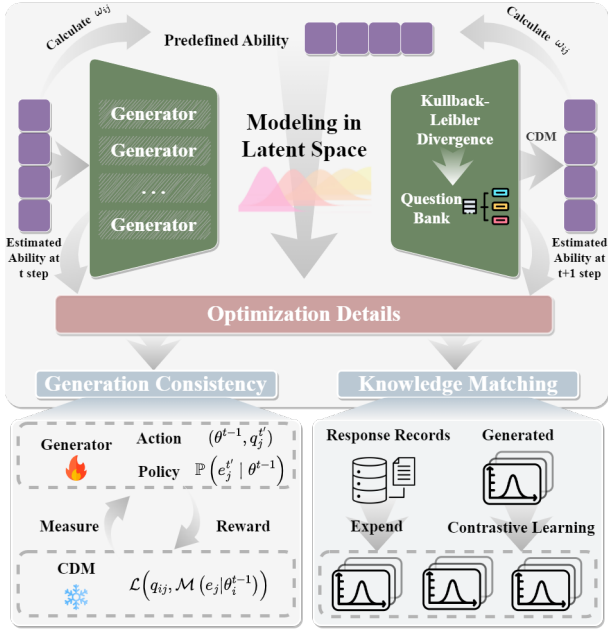


Figure 2: The overview of RECAT.

$\mathcal{M}(q_j | \theta_i^{t-1})$ denotes the predicted response r_{ij}^t of student s_i on exercise q_j under the ability θ_i^{t-1} . In such a definition, the former item aims to select a question that holds the smallest gap between the student’s true ability and estimated ability, the latter item is used to select the question that the student is easier to answer correctly.

Pseudo Feedback Note that there are no true responses/labels for student s_i on the question bank Q_i as shown in Eq.(4). Therefore, it is necessary to establish corresponding pseudo-real response labels for questions in Q_i .

To achieve this, we follow the assumption in (Ma et al. 2023a; Cao et al. 2007): *the probability of an examinee correctly answering a generated question is less than or equal to that of the examinee correctly answering the most similar question in the student’s finished questions (termed the real question) according to the question distribution.* Such the real question q_j^* can be obtained by $q_j^{t*} = \arg \min_{q_{in} \in \mathcal{R}_i} KL(q_{in} || gq_j^t)$, where gq_j^t denotes the j -th generated question at step t .

Question Generation Module

Based on the given “best-fitting” question definition, the RECAT aims to generate the question with the minimal ω_{ij} at each step. To this end, we regard the aforementioned problem of solving the conditional probability as a process of generating posterior distribution: given the examinee’s prior ability, generate the minimum ω_{ij} question for this ability value, where the conditional probability can be rewritten as $p(q^t | \theta_i^{t-1}, \theta_i^*)$.

To generate q^t , we encode it into a latent variable, serving as an implicit representation for updating the examinee’s

ability state. Given d knowledge concepts, the representations of question q_j and examinee s_i can be represented by two $1 \times d$ vector, whose k -bits can be denoted as $q_j(k)$ and $s_i(k)$, indicating the question difficulty and the examinee mastery level in k -th knowledge concept, respectively. Based on this, a CD-specific question generator is designed to generate questions.

Specifically, the examinee’s currently estimated ability θ_i^{t-1} is fed to the generator $G(\cdot)$ to analyze each dimension for generating suitable question parameters $gq^t \in \mathcal{R}^{1 \times d}$ used to denote the generated question gq^t , whose k -bit $gq^t(k)$ denotes the difficulty of the generated question in the k -th knowledge concept. The process of generating the distribution of the “best-fitting” question can be formalized as follows:

$$\begin{aligned} [\mu^t, \log(\sigma^t)^2] &= G(\theta_i^{t-1}), \\ gq^t &= \mu^t + \epsilon \odot \sigma^t, \epsilon \sim \mathcal{N}(0, 1), \end{aligned} \quad (5)$$

where the generator $G(\cdot)$ consists of multiple fully connected layers, $\mu^t \in \mathbb{R}^{1 \times d}$ and $\sigma^t \in \mathbb{R}^{1 \times d}$ represent the difficulty distribution generated for each dimension corresponding to knowledge concepts, and gq^t is the generated question for θ_i^{t-1} at step t .

Considering the diversity of examinees’ abilities and enhancing the robustness of the question selection process, K generators holding different weights are employed to generate multiple distributions to further generate multiple (K) questions, which allows for the diagnosis of examinees’ ability from different dimensions:

$$QK_i^t = \{gq_j^t | gq_j^t \sim \mathcal{N}(\mu_j^t, \sigma_j^t) = G_j(\theta_i^{t-1}), j = 1, \dots, K\}, \quad (6)$$

where $G_j(\cdot)$ is j -th generator, gq_j^t is the j -th generated question at step t , and QK_i^t is the set of K generated questions.

Question Matching Module

To better exploit the explicit E-Q relation, after obtaining K generated questions in QK_i^t for student s_i , we have to match each generated question with top- H similar real questions in the question bank, where the matching process can be denoted as

$$MQ_j^* = \arg \text{top-}H \min_{q_{in} \in \mathcal{Q}_i} KL(q_{in} || gq_j^t), \quad (7)$$

where MQ_j^* contains H questions for the generated question gq_j^t .

As a result, $K \times H$ real questions $MQ^* = \{MQ_1^*, \dots, MQ_K^*\}$ will be obtained, and the policy-based approaches can be incorporated into the RECAT to select the question by $q^t = \arg \max_{q \in MQ^*} \mathcal{I}(q | \theta_i^{t-1})$. The selection process

does not traverse the entire question bank but the candidate set, greatly reducing the question selection delay.

These policy-based question selection strategies involve the training of CDM while traversing the entire question bank (assuming a time complexity of A), with a time complexity of $O(A|E|)$. The proposed RECAT reduces the question selection scope for these combined selection strategies, and the network structure of RECAT is relatively simple (assuming a computational time complexity of B), whose time complexity is $O(A \times B \times K \times H) \ll O(A \times |E|)$.

Optimization Details

Given the examinee’s estimated ability θ_i^{t-1} at step $t - 1$ and the question generation criterion, our goal is to enhance the question generator’s capability to generate diverse best-fitting questions based on varying examinee abilities. While we can assess the quality of generated questions, this measurement occurs after the training of the CDM, making it challenging to effectively optimize the question selector. To address it, we optimize the generation objective $p(q^t | \theta_i^{t-1}, \theta_i^*)$. We view q^t from a global optimization perspective, which is a driving force generated in the $p(\theta_i^* | \theta_i^{t-1})$ process that steers the examinee state θ_i^{t-1} towards the true ability θ_i^* . Therefore, we maximize the likelihood $p(\theta_i^* | \theta_i^{t-1})$, which represents the possible marginalization result of all possible q^t . Maximizing the likelihood probability encourages the generation of "best-fitting" questions, which can be used as the optimization objective:

$$\begin{aligned} \log p_\psi(\theta_i^* | \theta_i^{t-1}) &= \int_{q^t} p_\psi(\theta_i^*, q^t | \theta_i^{t-1}) dq^t \\ &\geq \underbrace{\mathbb{E}_{c_\Psi(q^t | \theta_i^*, \theta_i^{t-1})} [\log p_\psi(\theta_i^* | q^t, \theta_i^{t-1})]}_{\text{Generation Consistency}} \\ &\quad - \underbrace{\text{KL}[c_\Psi(q^t | \theta_i^*, \theta_i^{t-1}) \| p_\psi(q^t | \theta_i^{t-1})]}_{\text{Knowledge Matching}}, \end{aligned} \quad (8)$$

where $c_\Psi(q^t | \theta_i^*, \theta_i^{t-1})$ denotes an approximate distribution to address the intractability of $p(q^t | \theta_i^*, \theta_i^{t-1})$. It is used to maximize the evidence lower bound on the right side of the inequality. Thus, we will optimize two objectives based on the CAT task, including Generation Consistency and Knowledge Matching.

Generation Consistency Optimizing the question generator through stochastic gradient descent (SGD) (Rendle et al. 2012) is a posterior event, which implies that evaluating the generator’s ability can only occur after updating examinee ability with CDM. This is not a straightforward task because ω_{ij} cannot be used to update both CDM and the question generator. To address this challenge, we introduce policy gradient-based reinforcement learning to maximize log-likelihood and enhance the generative ability:

$$\begin{aligned} \nabla_\Psi \mathcal{L}_{rl} &= \nabla_\Psi \sum_{s_i \in S} \mathbb{E}_{gq_j^t \sim QK_i^t} [\nabla_{\Phi} \mathcal{M}(gq_j^t, \theta_i^{t-1})] \\ &= \sum_{s_i \in S} \mathbb{E}_{gq_j^t \sim QK_i^t} [\nabla_{\Phi} \mathcal{M}(gq_j^t, \theta_i^{t-1}) \nabla_\Psi \mathbb{P}(gq_j^t | \theta_i^{t-1})], \end{aligned} \quad (9)$$

where ω_{ij} , generated by CDM after diagnosing each question, measures the quality of the generated questions. During the optimization process, the agent, i.e., K generators, follows a policy $\mathbb{P}(gq_j^t | \theta_i^{t-1})$. They generate questions $gq_j^t \in QK_i^t$ for each examinee to provide rich diagnostic value. For each action (θ_i^{t-1}, gq_j^t) , the environment, i.e., CDM, calculates the magnitude of the update of the examinee’s ability state towards the true ability θ_i^* for each question and computes ω_{ij} , returning the predicted probability as the reward to guide the optimization direction. To maximize

the log-likelihood, we define the policy $\mathbb{P}(gq_j^t | \theta_i^{t-1})$ as

$$\mathbb{P}(gq_j^t | \theta_i^{t-1}) = 1 - \frac{\exp(\omega_{ij})}{\sum_{gq_j^t \in QK_i^t} \exp(\omega_{ij'})}. \quad (10)$$

The smaller the ω_{ij} , the richer the diagnostic value generated, which encourages the generator to produce high-quality questions.

Knowledge Matching To align the parameter distribution of the generated question with that of a real question, we propose the use of KL divergence. Given the extensive and diverse nature of the question bank, relying solely on the parameters of the closest real questions for fitting might introduce biases. To address this issue, we suggest selecting questions from examinee historical response records \mathcal{R}_i to expand the candidate question set used for fitting, and then optimizing the distribution of generated questions using contrastive learning to reduce biases. However, due to the limited response records of an examinee, an insufficient number of questions may not form an effective candidate set. Inspired by (Gao et al. 2023), we propose expanding examinee s_i ’s response set based on similar true ability. We merge the response logs of the top- X examinees with the most similar representation to θ_i^* , forming a shared response log \mathcal{R}'_i . Then we select the top- M questions from \mathcal{R}'_i , which exhibit the minimum KL divergence with the generated question q^t , and add the corresponding response labels to the candidate set \mathcal{C}_i^t :

$$\mathcal{C}_i^t = \underset{\text{top-}M}{\operatorname{argmin}} \left\{ \left(\underset{(q_l, r_{il}) \in \mathcal{R}'_i}{\text{KL}}(q_l \| q^t), r_{il} \right) \right\}. \quad (11)$$

Consequently, we utilize the questions from the candidate set \mathcal{C}_i^t and the corresponding real response labels to update the examinee’s ability θ_i^{t-1} , obtaining the question evaluation:

$$\mathcal{W}_i^t = \left\{ \underbrace{\|\theta_i^* - \theta_i^{t-1} + \alpha \nabla \mathcal{L}(r_{ij}, \mathcal{M}(q_j, \theta_i^{t-1}))\|}_{=\omega_{ij}} \right\}_{q_j \in \mathcal{C}_i^t}, \quad (12)$$

where the smaller ω_{ij} denotes the real question q_j can offer more diagnostic value for θ_i^{t-1} . Then, we introduce contrastive learning to better optimize the parameters of generated questions. We consider the parameters of the question corresponding to the smallest ω_{ij} as the positive sample q^+ , and the parameters of the remaining $M - 1$ questions as a negative sample set $\{q^-\}_{m=1}^{M-1}$. We aim for the generated question q^t to be closer to the positive sample:

$$\mathcal{L}_{cl} = -\log \frac{\exp(\text{KL}(q^+ \| q^t) / \tau)}{\exp(\text{KL}(q^+ \| q^t) / \tau) + \sum_{m=1}^{M-1} \exp(\text{KL}(q_m^- \| q^t) / \tau)}. \quad (13)$$

Since examinees’ abilities are diagnosed through the CDM, we use the following loss function to update the entire framework:

$$\mathcal{L}_{RECAT} = \mathcal{L}_{rl}(\Psi) + \mathcal{L}_{cl}(\Psi) + \mathcal{L}_{CD}(\Phi_{\nabla_\theta}), \quad (14)$$

where Ψ represents the training parameters of the generation process, Φ_{∇_θ} denotes the examinee’s parameters of the cognitive diagnosis model while all other parameters are frozen.

Experiments

In this section, we conduct experiments on four real-world datasets to validate the effectiveness and efficiency of RECAT. We focus on the following research questions: **(RQ1)** Can RECAT outperform the performance achieved by solely utilizing either explicit or implicit relation in prediction task? **(RQ2)** How does RECAT perform in the simulation of ability estimation scenario? **(RQ3)** Can RECAT-based approaches reduce the latency of policy-based approaches? **(RQ4)** How adaptive is RECAT to new questions?

Dataset Description. We conducted experiments on four real-world datasets, namely ASSIST09, ASSIST12, NIPS-EDU, and Math. ASSIST09 (Feng, Heffernan, and Koedinger 2009) and ASSIST12 (Pardos et al. 2013) are publicly available datasets derived from the online tutoring system ASSISTments. NIPS-EDU (Wang et al. 2020b) is a dataset provided by the NeurIPS 2020 Education Challenge, which is collected from the educational platform Eedi. Math is a proprietary dataset compiled by a renowned e-learning platform, which comprises mathematics practice and examination records of elementary and secondary school examinees. We followed the preprocessing procedures from previous works (Zhuang et al. 2022, 2023; Bi et al. 2020), and the statistics of the preprocessed four datasets are summarized in Table 1.

Evaluation Metrics. We evaluated the effectiveness and efficiency of our RECAT on two CAT tasks. **(1) Examinee Score Prediction**, which is evaluated by AUC (Bradley 1997) and ACC. **(2) Simulation of Ability Estimation**, which is measured by MSE.

Baseline Methods. We compared RECAT with two categories of approaches in CAT, including policy-based (Random, MAAT (Bi et al. 2020), BECAT (Zhuang et al. 2023)) and learning-based (BOBCAT (Ghosh and Lan 2021), NCAT (Zhuang et al. 2022), GMOCAT (Wang et al. 2023)). We use RECAT-X to denote the use of X approach to exploit explicit E-Q relation. Our RECAT is compatible with existing policy-based approaches, which provides a candidate set by exploring implicit relations for these approaches that rely on explicit relation, and improves the effectiveness and efficiency of the original policy-based approaches. To ensure a fair comparison, we adopted the neural network-based cognitive diagnosis model NCD (Wang et al. 2020a) as the proficiency evaluation model.

Parameter Settings. During the training process, we initialized all the parameters in the CDM with Xavier (Glorot and Bengio 2010) initialization and used the Adam (Kingma and Ba 2014) optimizer with a fixed batch size of 256. We set the dimensions of latent features for both examinees and exercises to be equal to the number of knowledge concepts. During the testing process, we set the maximum test length $T = 10$. The number of question generators is set to $N = 5$. We set $k = 1$ as the number of real question with the minimum KL divergence for each generated question. All compared approaches follows the settings in their original papers, RECAT was implemented in Pytorch, and all experiments are conducted on an NVIDIA RTX4090 GPU.

Statistics	ASSIST09	ASSIST12	NIPS-EDU	MATH
# Examinees	1,473	3,823	4,918	6,618
# Questions	903	2,337	900	326
# Knowledge concepts	22	89	85	56
# Response logs	58,427	444,723	1,160,925	245,155
# Avg logs per student	39	116	236	37

Table 1: The statistics of the datasets.

Examinee Performance Prediction (RQ1)

To verify the effectiveness of the proposed framework, we combined the existing policy-based approaches with the RECAT, and set the question selection step to 5 and 10. Table 2 shows the comparison results on RECAT-based with policy-based and learning-based approaches. RECAT-based approaches, which exploit both explicit and implicit E-Q relations, outperform all baselines in terms of ACC and AUC over four datasets. Additionally, we can observe the following: (1) Compared to policy-based approaches solely relying on explicit E-Q relation, RECAT-based approaches enhance the selection performance of such methods by exploring implicit E-Q relation. Specifically, within NIPS-EDU, the original MAAT exhibits negative optimization of examinee ability from step 5 to step 10, possibly due to the greedy strategy leading to local optima. In contrast, BECAT-MAAT significantly improves the quality of question selection, as the additional implicit E-Q relation enhances the robustness of the selection process. (2) Compared to learning-based approaches that solely leverage implicit E-Q relation, the RECAT-Random, which does not utilize any explicit relation, demonstrates the effectiveness of our implicit E-Q relation mining approach and proves that it outperforms existing learning-based approaches.

Ability Estimation Simulation (RQ2)

To verify the performance of our proposed RECAT in examinee ability estimation, we conduct experiments on the ASSIST12 and NIPS-EDU datasets for simulated ability estimation. We use MSE to measure the impact of the question selector on examinee ability estimation. As shown in Figure 3, the original MAAT exhibits unstable decreases in estimated error as the testing steps increase. This instability may arise from the greedy strategy leading them into local optima, thus adversely affecting future question selections. In contrast, the RECAT-based approaches consistently achieve the lowest estimation error. This is attributed to RECAT of mining implicit relation with examinee’s true ability-oriented, which aids the selector in better learning this implicit relation, i.e., using low estimated error as a learning target enhances the capability of the selector to understand and predict examinee ability accurately.

Selection Latency Comparison (RQ3)

To validate the efficiency of RECAT, as shown in Table 3, we compare the selection latency of RECAT-based approaches with the original policy-based approaches on the ASSIST12 and NIPS-EDU datasets. Firstly, BECAT computes the ques-

Dataset	Assist09		NIPS-EDU		Assist12		Math		
Metric	AUC/ACC								
Step	5	10	5	10	5	10	5	10	
Policy	Random	0.6394/0.6024	0.6480/0.6057	0.6851/0.6265	0.7035/0.6305	0.6452/0.7050	0.6488/0.7085	0.7240/0.6577	0.7302/0.6733
	MAAT	0.6431/0.6129	0.6517/0.6170	0.6988/0.6349	0.7226/0.6329	0.6549/0.7120	0.6563/0.7137	0.7262/0.6653	0.7425/0.6800
	BECAT	0.6425/0.6141	0.6522/0.6191	0.7011/0.6404	0.7288/0.6491	0.6589/0.7133	0.6611/0.7172	0.7189/0.6601	0.7432/0.6822
Learnable	BOBCAT	0.6419/0.6136	0.6513/0.6213	0.7016/0.6395	0.7268/0.6513	0.6553/0.7125	0.6680/0.7197	0.7211/0.6649	0.7453/0.6864
	NCAT	0.6423/0.6184	0.6526/0.6263	0.7053/0.6421	0.7298/0.6533	0.6565/0.7154	0.6653/0.7203	0.7245/0.6673	0.7467/0.6856
	GMOCAT	<u>0.6443/0.6225</u>	0.6525/0.6283	0.7082/0.6478	0.7314/0.6546	0.6633/0.7144	0.6741/0.7188	0.7278/0.6704	0.7524/0.6933
RECAT	Random	0.6428/0.6215	0.6525/0.6322	0.7112/0.6498	0.7356/0.6615	0.6706/0.7158	0.6801/0.7211	0.7339/0.6728	0.7514/0.6923
	MAAT	0.6485/0.6249	0.6530/0.6348	0.7131/0.6525	0.7400/0.6663	0.6711/0.7170	0.6805/0.7222	0.7345/0.6737	0.7562/0.6945
	BECAT	0.6436/0.6222	0.6564/0.6307	0.7197/0.6542	0.7423/0.6632	0.6709/0.7161	0.6815/0.7213	0.7356/0.6758	0.7574/0.6956

Table 2: The AUC and ACC performance on Student Performance Prediction task on four datasets. The best results are highlighted in bold, while the second best results are underlined.




Figure 3: The MSE of Ability Estimation Simulation.

tion bank with the selected questions during the subset selection, which leads to a monotonically increasing selection latency as the testing steps increase. In contrast, RECAT-BECAT capitalizes on the implicit E-Q relation to narrow the search space, reducing computational costs and selection latency without the need for exhaustive computations. Secondly, as the question bank size increases, original MAAT and BECAT show significantly higher selection latency on ASSIST12 compared to NIPS-EDU, whereas RECAT-based approaches maintain similar computational costs. This is because exploiting implicit relations in question selection significantly reduces computational complexity compared to gradient updates, leading to substantial reductions in selection latency for RECAT-based approaches.

Adaptability to New Questions (RQ4)

To validate the adaptability of RECAT to a continuously modified question bank where we randomly removed {25, 50, 75, 100} questions along with their corresponding response records from NIPS-EDU question bank during the training process. During the testing phase, despite retaining pre-trained parameters, they were entirely unfamiliar to the selector. As depicted in Figure 4, experiments with RECAT-MAAT revealed a decline in selection performance as more questions were added. This decline can be attributed to the potentially richer diagnostic value of some removed questions during training, particularly advantageous for the Knowledge Matching process that guides the impact of im-

Datasets	Step	Approach	Latency(s)
NIPS-EDU	5	MAAT/RECAT-MAAT	14.63/ 0.12
		BECAT/RECAT-BECAT	47.56/ 1.82
	10	MAAT/RECAT-MAAT	28.04/ 0.21
		BECAT/RECAT-BECAT	239.63/ 8.41
ASSIST12	5	MAAT/RECAT-MAAT	49.04/ 0.49
		BECAT/RECAT-BECAT	137.32/ 5.72
	10	MAAT/RECAT-MAAT	94.82/ 0.81
		BECAT/RECAT-BECAT	1432.15/ 26.15

Table 3: The Latency in selection Comparison




Figure 4: Adaptability to the New Questions.

PLICIT E-Q relation on question generation. Consequently, bias in generating question parameters resulted in lower-quality candidate sets. However, unlike other learning-based approaches that solely rely on implicit E-Q relation, our RECAT does not require retraining the entire model, which saves on training costs when dealing with new questions.

Conclusion

In this work, we attempted to exploit explicit and implicit examinee-question (E-Q) relation for CAT. We proposed a novel and generic framework, **RECAT (Relation Exploiting-based CAT)**, which delves deeper into the implicit E-Q relation guided by the examinee’s true ability and generates high-quality candidate sets for the utilization of explicit E-Q relation. Experimental results on real datasets demonstrate the effectiveness and efficiency of simultaneously exploiting both types of relation for CAT selection. We intend to further investigate the exploration of both relations tailored to the characteristics of CAT task.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (NO.62107001), in part by the Enterprise Research Transformation and Industrialization Special Project (NO.2023-GX-C13), in part by the National Natural Science Foundation of China (No.62302010), in part by China Postdoctoral Science Foundation (No.2023M740015), in part by the Postdoctoral Fellowship Program (Grade B) of China Postdoctoral Science Foundation (No.GZB20240002), and in part by the Anhui Province Key Laboratory of Intelligent Computing and Applications (No.AFZJNS2024KF01).

References

- Bi, H.; Ma, H.; Huang, Z.; Yin, Y.; Liu, Q.; Chen, E.; Su, Y.; and Wang, S. 2020. Quality meets diversity: A model-agnostic framework for computerized adaptive testing. In *2020 IEEE International Conference on Data Mining (ICDM)*, 42–51. IEEE.
- Bradley, A. P. 1997. The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern recognition*, 30(7): 1145–1159.
- Cao, Z.; Qin, T.; Liu, T.-Y.; Tsai, M.-F.; and Li, H. 2007. Learning to rank: from pairwise approach to listwise approach. In *Proceedings of the 24th international conference on Machine learning*, 129–136.
- Chang, H.-H.; and Ying, Z. 1996. A global information approach to computerized adaptive testing. *Applied Psychological Measurement*, 20(3): 213–229.
- Feng, M.; Heffernan, N.; and Koedinger, K. 2009. Addressing the assessment challenge with an online system that tutors as it assesses. *User Modeling and User-Adapted Interaction*, 19(3): 243–266.
- Gao, W.; Liu, Q.; Wang, H.; Yue, L.; Bi, H.; Gu, Y.; Yao, F.; Li, Z. Z. X.; and He, Y. 2023. Zero-1-to-3: Domain-level Zero-shot Cognitive Diagnosis via One Batch of Early-bird Students towards Three Diagnostic Objectives. *arXiv preprint arXiv:2312.13434*.
- Ghosh, A.; and Lan, A. 2021. Bobcat: Bilevel optimization-based computerized adaptive testing. *arXiv preprint arXiv:2108.07386*.
- Gilavert, P.; and Freire, V. 2022. Computerized Adaptive Testing: A Unified Approach Under Markov Decision Process. In *International Conference on Computational Science and Its Applications*, 591–602. Springer.
- Glorot, X.; and Bengio, Y. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, 249–256. JMLR Workshop and Conference Proceedings.
- Hong, Y.; Tong, S.; Huang, W.; Zhuang, Y.; Liu, Q.; Chen, E.; Li, X.; and He, Y. 2023. Search-Efficient Computerized Adaptive Testing. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 773–782.
- Kingma, D.; and Ba, J. 2014. Adam: A Method for Stochastic Optimization. *Computer Science*.
- Li, X.; Xu, H.; Zhang, J.; and Chang, H.-h. 2023a. Deep reinforcement learning for adaptive learning systems. *Journal of Educational and Behavioral Statistics*, 48(2): 220–243.
- Li, X.; Xu, H.; Zhang, J.; and Chang, H.-h. 2023b. Deep reinforcement learning for adaptive learning systems. *Journal of Educational and Behavioral Statistics*, 48(2): 220–243.
- Liu, Q.; Zhuang, Y.; Bi, H.; Huang, Z.; Huang, W.; Li, J.; Yu, J.; Liu, Z.; Hu, Z.; Hong, Y.; et al. 2024. Survey of Computerized Adaptive Testing: A Machine Learning Perspective. *arXiv preprint arXiv:2404.00712*.
- Liu, S.; Yu, X.; Ma, H.; Wang, Z.; Qin, C.; and Zhang, X. 2023. 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.
- Lord, F. M. 1980. *Applications of Item Response Theory to Practical Testing Problems*. LAWRENCE ERLBAUM ASSOCIAATES.
- Lord, F. M. 2012. *Applications of item response theory to practical testing problems*. Routledge.
- Ma, H.; Song, S.; Qin, C.; Yu, X.; Zhang, L.; Zhang, X.; and Zhu, H. 2024a. DGCD: An Adaptive Denoising GNN for Group-level Cognitive Diagnosis. In *The 33rd International Joint Conference on Artificial Intelligence (IJCAI-24)*.
- Ma, H.; Wang, C.; Zhu, H.; Yang, S.; Zhang, X.; and Zhang, X. 2023a. Enhancing Cognitive Diagnosis using Un-interacted Exercises: A Collaboration-aware Mixed Sampling Approach. *arXiv preprint arXiv:2312.10110*.
- Ma, H.; Xia, A.; Wang, C.; Wang, H.; and Zhang, X. 2024b. Diffusion-Inspired Cold Start with Sufficient Prior in Computerized Adaptive Testing. *arXiv preprint arXiv:2411.12182*.
- Ma, H.; Yang, Y.; Qin, C.; Yu, X.; Yang, S.; Zhang, X.; and Zhu, H. 2024c. HD-KT: Advancing Robust Knowledge Tracing via Anomalous Learning Interaction Detection. In *Proceedings of the ACM on Web Conference 2024*, 4479–4488.
- Ma, H.; Zeng, Y.; Yang, S.; Qin, C.; Zhang, X.; and Zhang, L. 2023b. A novel computerized adaptive testing framework with decoupled learning selector. *Complex & Intelligent Systems*, 1–12.
- Ma, H.; Zeng, Y.; Yang, S.; Qin, C.; Zhang, X.; and Zhang, L. 2023c. A novel computerized adaptive testing framework with decoupled learning selector. *Complex & Intelligent Systems*, 1–12.
- Pardos, Z. A.; Baker, R. S.; San Pedro, M. O.; Gowda, S. M.; and Gowda, S. M. 2013. Affective states and state tests: Investigating how affect throughout the school year predicts end of year learning outcomes. In *Proceedings of the third international conference on learning analytics and knowledge*, 117–124.
- Qin, C.; Zhang, L.; Zha, R.; Shen, D.; Zhang, Q.; Sun, Y.; Zhu, C.; Zhu, H.; and Xiong, H. 2023. A Comprehensive Survey of Artificial Intelligence Techniques for Talent Analytics. *arXiv preprint arXiv:2307.03195*.

- Reckase, M. D. 2009. *Multidimensional Item Response Theory*. Springer New York.
- Rendle, S.; Freudenthaler, C.; Gantner, Z.; and Schmidt-Thieme, L. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618*.
- Wainer, H.; Dorans, N. J.; Flaugher, R.; Green, B. F.; and Mislevy, R. J. 2000. *Computerized adaptive testing: A primer*. Routledge.
- Wang, F.; Liu, Q.; Chen, E.; Huang, Z.; Chen, Y.; Yin, Y.; Huang, Z.; and Wang, S. 2020a. Neural cognitive diagnosis for intelligent education systems. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 6153–6161.
- Wang, H.; Long, T.; Yin, L.; Zhang, W.; Xia, W.; Hong, Q.; Xia, D.; Tang, R.; and Yu, Y. 2023. GMOCAT: A Graph-Enhanced Multi-Objective Method for Computerized Adaptive Testing. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2279–2289.
- Wang, Z.; Lamb, A.; Saveliev, E.; Cameron, P.; Zaykov, Y.; Hernández-Lobato, J. M.; Turner, R. E.; Baraniuk, R. G.; Barton, C.; Jones, S. P.; et al. 2020b. Instructions and guide for diagnostic questions: The neurips 2020 education challenge. *arXiv preprint arXiv:2007.12061*.
- Yang, S.; Chen, M.; Wang, Z.; Yu, X.; Zhang, P.; Ma, H.; and Zhang, X. 2024a. DisenGCD: A Meta Multigraph-assisted Disentangled Graph Learning Framework for Cognitive Diagnosis. *arXiv preprint arXiv:2410.17564*.
- Yang, S.; Qin, L.; and Yu, X. 2024. Endowing Interpretability for Neural Cognitive Diagnosis by Efficient Kolmogorov-Arnold Networks. *arXiv preprint arXiv:2405.14399*.
- Yang, S.; Yu, X.; Tian, Y.; Yan, X.; Ma, H.; and Zhang, X. 2024b. Evolutionary neural architecture search for transformer in knowledge tracing. *Advances in Neural Information Processing Systems*, 36.
- 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.
- Yu, X.; Qin, C.; Shen, D.; Ma, H.; Zhang, L.; Zhang, X.; Zhu, H.; and Xiong, H. 2024a. Rdgt: enhancing group cognitive diagnosis with relation-guided dual-side graph transformer. *IEEE Transactions on Knowledge and Data Engineering*.
- Yu, X.; Qin, C.; Shen, D.; Yang, S.; Ma, H.; Zhu, H.; and Zhang, X. 2024b. Rigl: A unified reciprocal approach for tracing the independent and group learning processes. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 4047–4058.
- Yu, X.; Qin, C.; Zhang, Q.; Zhu, C.; Ma, H.; Zhang, X.; and Zhu, H. 2024c. DISCO: A Hierarchical Disentangled Cognitive Diagnosis Framework for Interpretable Job Recommendation. *arXiv preprint arXiv:2410.07671*.
- Zhuang, Y.; Liu, Q.; Huang, Z.; Li, Z.; Shen, S.; and Ma, H. 2022. Fully adaptive framework: Neural computerized adaptive testing for online education. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 4734–4742.
- 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*.