

PEOAT: Personalization-Guided Evolutionary Question Assembly for One-Shot Adaptive Testing

Xiaoshan Yu¹, Ziwei Huang², Shangshang Yang^{3*}, Ziwen Wang³,
Haiping Ma⁴, Xingyi Zhang³

¹School of Artificial Intelligence, Anhui University, China

²School of Computing and Information Systems, Singapore Management University, Singapore

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

⁴Institutes of Physical Science and Information Technology, Anhui University, China

{yxsleo, yangshang0308, wzwl2sir, xyzhanghust}@gmail.com, ziwei Huang@smu.edu.sg, hpma@ahu.edu.cn

Abstract

With the rapid advancement of intelligent education, Computerized Adaptive Testing (CAT) has attracted increasing attention by integrating educational psychology with deep learning technologies. Unlike traditional paper-and-pencil testing, CAT aims to efficiently and accurately assess examinee abilities by adaptively selecting the most suitable items during the assessment process. However, its real-time and sequential nature presents limitations in practical scenarios, particularly in large-scale assessments where interaction costs are high, or in sensitive domains such as psychological evaluations where minimizing noise and interference is essential. These challenges constrain the applicability of conventional CAT methods in time-sensitive or resource-constrained environments. To this end, we first introduce a novel task called one-shot adaptive testing (OAT), which aims to select a fixed set of optimal items for each test-taker in a one-time selection. Meanwhile, we propose **PEOAT**, a **P**ersonalization-guided **E**volutionary question assembly framework for **O**ne-shot **A**daptive **T**esting from the perspective of combinatorial optimization. Specifically, we began by designing a personalization-aware initialization strategy that integrates differences between examinee ability and exercise difficulty, using multi-strategy sampling to construct a diverse and informative initial population. Building on this, we proposed a cognitive-enhanced evolutionary framework incorporating schema-preserving crossover and cognitively guided mutation to enable efficient exploration through informative signals. To maintain diversity without compromising fitness, we further introduced a diversity-aware environmental selection mechanism. The effectiveness of PEOAT is validated through extensive experiments on two datasets, complemented by case studies that uncovered valuable insights.

Introduction

Computerized adaptive testing (CAT) (Wainer et al. 2000), as a significant and promising approach to personalized assessment in intelligent education (Roll and Wylie 2016; Holmes, Bialik, and Fadel 2023), has garnered increasing attention and development in recent years. Its goal is to deliver adaptive ability evaluation for students through progressive

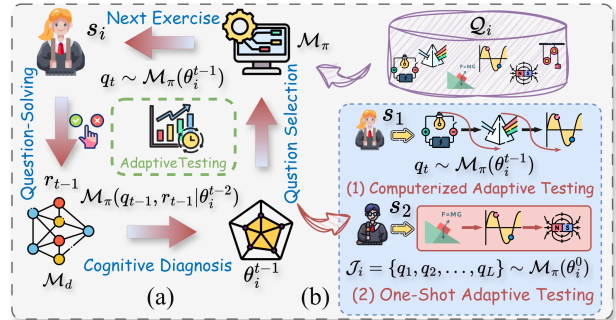


Figure 1: (a) The process of computerized adaptive testing; (b) Comparison of two adaptive testing tasks (CAT & OAT).

interaction and feedback, integrating principles from educational psychology with advances in deep learning. In general, CAT consists of two primary modules (Liu et al. 2024; Chang 2015): a question selection module \mathcal{M}_π and a cognitive diagnosis module \mathcal{M}_d . The former adaptively selects the most suitable questions based on the test taker's current ability status, while the latter diagnoses the test taker's knowledge proficiency based on response feedback, as illustrated in Figure 1(a). These two modules operate alternately until a predefined termination condition is satisfied.

Existing research on CAT (Liu et al. 2024) primarily focuses on enhancing the question selection algorithm, which are widely regarded as key determinants of assessment adaptability and effectiveness. These approaches can be broadly categorized into heuristic methods and data-driven learning methods. Heuristic approaches (Bi et al. 2020; Zhuang et al. 2023) rely on explicitly defined, interpretable rules to select items that align question characteristics with the test taker's estimated ability. For instance, BECAT (Zhuang et al. 2023) approximates full-response gradients to guide item selection, enabling accurate ability estimation with fewer questions and offering theoretical guarantees on estimation error. In contrast, data-driven methods (Ghosh and Lan 2021; Zhuang et al. 2022) seek to improve performance by learning personalized item selection policies directly from data. A representative example is NCAT (Zhuang et al. 2022), which views CAT as a bilevel reinforcement learning problem, where an attentive policy is

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

trained to select items by modeling learning behavior.

Despite existing CAT methods have demonstrated notable success, their inherent interactivity, requiring iterative item selection and ability estimation, poses significant limitations in scenarios with high interaction costs or constrained response conditions. In many real-world scenarios, such as psychological assessments (Meyer et al. 2001), post-instruction diagnostic evaluations (Holman 2000), or remote/offline testing (Haq et al. 2021), the feasibility of interactive testing is often hindered by factors such as response latency, user anxiety, or device limitations. To address this gap, this paper proposes a novel task called One-Shot Adaptive Testing (OAT), in which a fixed set of candidate items is adaptively selected beforehand and presented to the test-taker all at once, as illustrated in Figure 1(b). Considering the characteristics of this problem, we attempt to model it from a combinatorial optimization perspective.

However, this task is challenging mainly due to three issues: (1) *ensuring student adaptability during optimization*; (2) *searching effectively in a vast solution space*; (3) *mitigating encoding sparsity given a candidate pool much larger than the test length*. To this end, in this paper, we propose **PEOAT**, a **P**ersonalization-guided **E**volutionary question assembly framework for **O**ne-hot **A**daptive **T**esting. Specifically, we first propose a personalization aware-based population initialization strategy that accounts for individual student ability differences and exercise difficulty, employing multi-strategy sampling to generate a diverse initial question population and effectively construct the initial search space. Next, we develop a cognitive-enhanced evolutionary search framework, featuring the schema-preserving uniform crossover and the cognitive information-guided mutation operators that leverage informative cues throughout population evolution for efficient exploration. Finally, we design a diversity-preserving environmental selection strategy that balances diversity maintenance with fitness during offspring selection. Extensive experiments on two real-world educational datasets validate the effectiveness of the proposed PEOAT model. Additionally, we conduct insightful case studies that reveal valuable findings.

Related Work

Computerized Adaptive Testing

As a core assessment paradigm in personalized education, computerized adaptive testing (CAT) (Wainer et al. 2000) originated from educational psychology and has evolved through the incorporation of deep learning techniques (Ma et al. 2024b, 2025b; Li et al. 2025). It aims to achieve accurate ability diagnosis by interactively selecting suitable exercises in response to test-taker performance. Recent advances in CAT have predominantly focused on improving item selection strategies, generally falling into two categories (Chang 2015; Liu et al. 2024; Yu et al. 2024c): heuristic methods and data-driven approaches. The former (Chang and Ying 1996; Chang 2015; Zhuang et al. 2023; Bi et al. 2020; Ma et al. 2025a; Yang, Qin, and Yu 2024) selects questions based on explicitly defined and interpretable rules, aiming to match question characteristics with the test-taker's

estimated ability. For example, Maximum Fisher Information (MFI) (Lord 2012) minimizes ability estimation variance via local item information, whereas KLI (Chang and Ying 1996) improves robustness by incorporating global Kullback-Leibler divergence. Moreover, MAAT (Bi et al. 2020) defines the informativeness of exercises based on the expected maximum change criterion from active learning. In contrast, data-driven methods (Ghosh and Lan 2021; Zhuang et al. 2022; Wang et al. 2023; Yu et al. 2024a) aim to enhance performance by learning personalized selection policies directly from learner-exercise interaction data. Representatively, NCAT (Zhuang et al. 2022) casts CAT as a bilevel reinforcement learning problem, where an attentive neural policy is trained to select items by directly modeling student behaviors (Gao et al. 2025, 2024a; Yu et al. 2024d). Although these methods have achieved notable success, they are often impractical in resource-constrained ability assessment scenarios, highlighting the need for one-shot adaptive testing, which serves as the primary motivation for this study.

Evolutionary Optimization Application

Combinatorial optimization (Papadimitriou and Steiglitz 1998; Blum and Roli 2003) refers to the process of searching for an optimal object from a finite but often exponentially large solution space, and it plays a central role in various complex decision-making tasks (Yu et al. 2025b; Yang et al. 2025b; Ma et al. 2024a). When the solution space lacks closed-form structure or involves complex constraints, gradient-based methods (Lezcano Casado 2019) often fail, making heuristic strategies, particularly evolutionary algorithms (EAs), a compelling alternative (Yang et al. 2023b; Yu et al. 2024b). Over the past decades, a wide variety of evolutionary algorithms (Črepinšek, Liu, and Mernik 2013) have been proposed and refined. Classical examples include the Genetic Algorithm (GA) (Lambora, Gupta, and Chopra 2019), which mimics natural selection through genetic operators, and Differential Evolution (DE) (Das and Suganthan 2010), which leverages vector-based mutations for continuous and combinatorial tasks. These methods have proven effective in various domains and are gaining increasing traction in education (Yang et al. 2023a; Bu et al. 2022; Sun et al. 2022; Bu et al. 2023), where they are used to tackle complex decision-making problems. For example, PEGA (Yang et al. 2023a) employs a constrained multi-objective framework with dual co-evolution to assemble personalized exercise groups (Liu et al. 2023; Yu et al. 2024b). In the cognitive diagnosis (Yang et al. 2025a; Dong, Chen, and Wu 2025), HGA-CDM (Bu et al. 2022) applies a memetic algorithm combining genetic and adaptive local search to the DINA model, mitigating its exponential computational complexity. However, how to effectively model the OAT task from an evolutionary optimization perspective remains unexplored and presents a valuable research direction.

Preliminary

Problem Statement

In this section, we formally define the One-Shot Adaptive Testing (OAT) task. In an intelligent education system, let

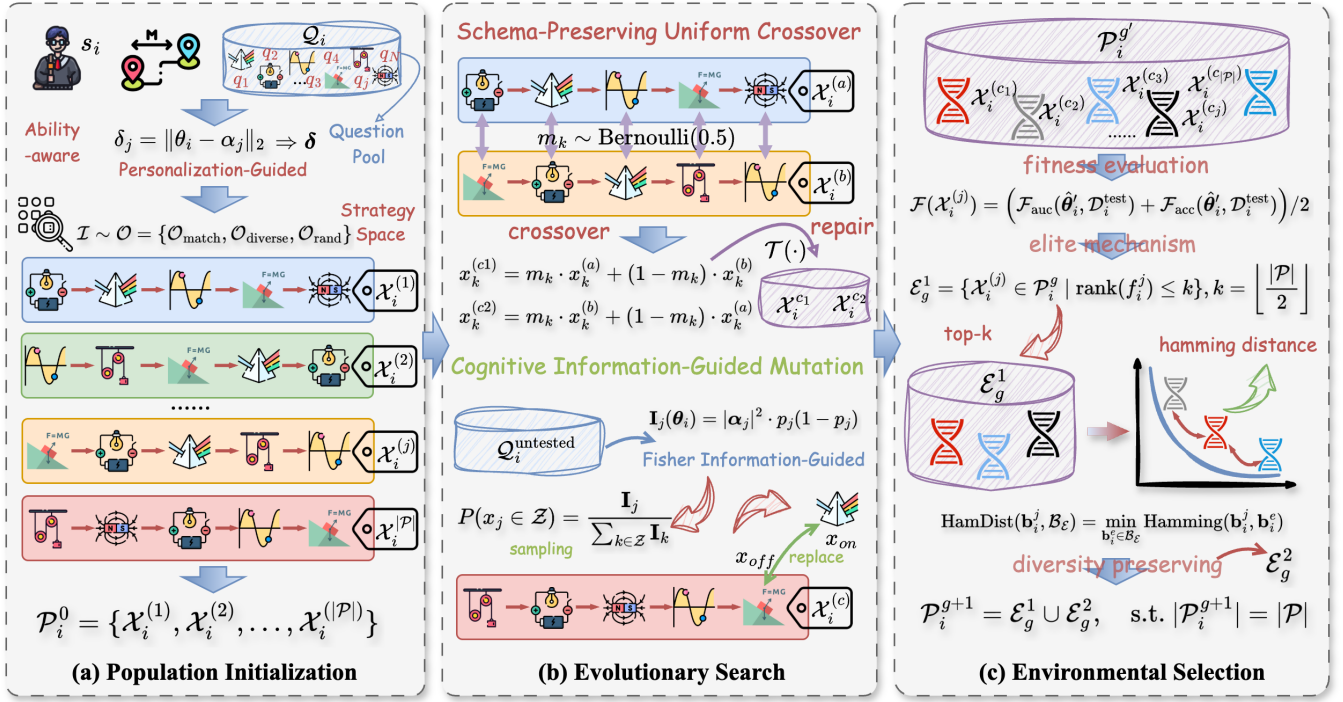


Figure 2: The overview architecture of our proposed PEOAT model. (a) The personalization-aware population initialization. (b) The the cognitive-enhanced evolutionary search. (c) The diversity-preserving environmental selection. Best viewed in color.

$\mathcal{S} = \{s_1, s_2, \dots, s_N\}$ be the set of N students, $\mathcal{Q} = \{q_1, q_2, \dots, q_M\}$ be the candidate pool of M questions, and $\mathcal{C} = \{c_1, c_2, \dots, c_K\}$ be the set of K knowledge concepts. The mapping between questions and knowledge concepts is commonly represented by a Q -matrix, denoted as $\mathbf{Q} = \{m_{ij}\}^{M \times K}$. In this matrix, an entry $m_{ij} = 1$ signifies that question e_i is linked to concept c_j , while $m_{ij} = 0$ indicates no such association. For each student $s_i \in \mathcal{S}$ with historical assessment records, their interactions can be represented as $\mathcal{R}_i = \{(s_i, q_j, r_{ij}) \mid q_j \in \mathcal{Q}, r_{ij} \in \{0, 1\}\}$, where $r_{ij} = 1$ denotes a correct response to question q_j , and $r_{ij} = 0$ otherwise. The complete **One-Shot Adaptive Testing (OAT)** system is composed of two fundamental components: (1) the *cognitive diagnosis module* \mathcal{M}_d that models the examinee’s knowledge proficiency by predicting the probability of correctly answering each question q (Gao et al. 2024b; Yu et al. 2025a), denoted as $\mathcal{M}_d(q \mid \theta) \in [0, 1]$; and (2) the *question selection module* \mathcal{M}_π that selects a subset of L questions $\mathcal{J} \subset \mathcal{Q}$ in a one-shot manner, based on an initial ability θ^0 . More specifically, given the initial ability estimate θ_i^0 of examinee s_i , the OAT selects a fixed-length question set $\mathcal{J}_i = \{q_1, q_2, \dots, q_L\} \sim \mathcal{M}_\pi(\theta_i^0)$ without any intermediate feedback during the test process. After the examinee finishes all L questions and their responses $\mathbf{r} = \{r_{i1}, r_{i2}, \dots, r_{iL}\}$ are collected, the diagnostic model \mathcal{M}_d conducts a single-step ability update to produce the final proficiency estimate θ_i^{final} . In contrast to conventional CAT, where questions are selected sequentially as $q_t \sim \mathcal{M}_\pi(\theta_i^{t-1})$ and ability estimates θ_i^t are updated iteratively after each response, OAT aims to estimate the true knowledge proficiency $\hat{\theta}_i$ as accurately and efficiently as possible using only a single batch of

adaptively selected questions, i.e., $\theta_i^{\text{final}} \rightarrow \hat{\theta}_i$.

Combinatorial Optimization Perspective

Unlike CAT, which selects questions in a sequential and feedback-driven manner (Yu et al. 2024a), OAT poses a distinct challenge: selecting an optimal fixed-length question set in a single round without any intermediate feedback. This constraint requires the selection policy to holistically consider the test-taker’s prior ability and question characteristics to maximize the diagnostic utility of the selected items. From a modeling perspective, this task can be naturally viewed as a bi-level combinatorial optimization problem, where the outer layer selects a subset of questions, and the inner layer estimates student ability based on simulated response data. The optimization objective is to ensure that the final ability estimation is as close as possible to the student’s true proficiency. Formally, from the perspective of discrete combinatorial optimization, the OAT task for each student $s_i \in \mathcal{S}$ can be characterized as follows:

$$\left\{ \begin{array}{l} \mathcal{J}_i^* = \arg \max_{\mathcal{J}_i \subseteq \mathcal{Q}_i^{\text{untested}}} \mathcal{F}(\theta_i^{\text{final}}(\mathcal{J}_i), \hat{\theta}_i), \\ \text{s.t. } \theta_i^{\text{final}}(\mathcal{J}_i) = \arg \min_{\theta_i} \sum_{(q_j, r_{ij}) \in \mathcal{R}_i(\mathcal{J}_i)} \mathcal{L}(r_{ij}, \mathcal{M}_d(q_j \mid \theta_i)), \\ \text{where } \mathcal{J}_i \sim \mathcal{M}_\pi(\theta_i^0), \text{ and} \\ \mathcal{R}_i(\mathcal{J}_i) = \{(q_1, r_{i1}), \dots, (q_L, r_{iL}) \mid q_i \in \mathcal{J}_i\}. \end{array} \right. \quad (1)$$

Methodology

In this section, we present the PEOAT framework in detail. As depicted in Figure 2, the PEOAT is composed of three key components: the personalization-aware population initialization, the cognitive-enhanced evolutionary search, and the diversity-preserving environmental selection.

Personalization-Aware Population Initialization

To effectively guide the evolutionary search in OAT, we design a personalization-aware population initialization mechanism that adaptively generates an informative and diverse initial population based on students' personal abilities and the characteristics of the candidate exercises. As mentioned earlier, the one-shot item selection process for each student can be modeled as a population-based evolutionary optimization procedure. Accordingly, for each student $s_i \in \mathcal{S}$, every individual in the population represents a candidate test form consisting of L questions, encoded as follows:

$$\mathcal{X}_i^{(j)} = [x_1, x_2, \dots, x_L] \in \mathcal{Q}_i^{\text{untested}^L}, \quad (2)$$

where $\mathcal{X}_i^{(j)}$ denotes the j -th individual in the candidate population of student s_i , and x_k represents the k -th gene in the chromosome, which indexes a question from the student's remaining question pool, i.e., $x_k \rightarrow q_k \in \mathcal{Q}_i^{\text{untested}}$, and $\mathcal{Q}_i^{\text{untested}}$ denotes the set of untested questions for student s_i . We assume that all selected indices are unique—i.e., $q_{k_1} \neq q_{k_2}$ for $k_1 \neq k_2$ —thereby satisfying the fixed-length constraint $|\mathcal{X}| = L$. This subset-based encoding not only defines the structure of each individual but also serves as a retrieval mechanism for latent features (e.g., question embeddings or difficulty parameters). Compared to sparse one-hot encodings, it offers a more compact and efficient representation, particularly suited for large-scale optimization.

To embed personalized prior knowledge into the search space while effectively balancing exploitation and exploration, we propose a multi-strategy population initialization mechanism. Specifically, we define a strategy space $\mathcal{O} = \{\mathcal{O}_{\text{match}}, \mathcal{O}_{\text{diverse}}, \mathcal{O}_{\text{rand}}\}$, representing three initialization strategies that select candidate exercises based on students' initial abilities: matching, diverse, and random, respectively—each encouraging a distinct form of exploration. For each individual, one strategy is randomly sampled from \mathcal{O} , and the process of constructing question index gene-encoded candidates can be formalized as follows:

$$\mathcal{I}_i \sim \begin{cases} \{\text{Uniform}(\text{Top}_{2L}(\delta_i^\uparrow))\}^L, & \text{if } \mathcal{O}' = \mathcal{O}_{\text{match}}, \\ \{\text{Uniform}(\text{Top}_{2L}(\delta_i^\downarrow))\}^L, & \text{if } \mathcal{O}' = \mathcal{O}_{\text{diverse}}, \\ \{\text{Uniform}(\delta_i^\rightarrow [2L : -2L])\}^L, & \text{if } \mathcal{O}' = \mathcal{O}_{\text{rand}}, \end{cases} \quad (3)$$

where $|\mathcal{I}| = L$, and δ_i^\uparrow and δ_i^\downarrow denote the ascending and descending sorted indices of δ_i , respectively. The vector $\delta_i = [\delta_1, \delta_2, \dots, \delta_{|\mathcal{Q}_i|}]$ represents a personalized distance vector that quantifies the matching quality between student s_i and the questions in \mathcal{Q}_i , where each δ_j is computed as:

$$\delta_j = \|\theta_i - \alpha_j\|_2, \quad \forall j \in \{1, 2, \dots, |\mathcal{Q}_i|\}, \quad (4)$$

where θ_i and α_j represents the ability vector of student s_i and the difficulty vector of question q_j , respectively. Each

resulting index set \mathcal{I}_i is subsequently transformed into the corresponding individual encoding, i.e., $\mathcal{I}_i \rightarrow \mathcal{X}_i$. The final initialized population of predefined size $|\mathcal{P}|$ is given by:

$$\mathcal{P}_i^0 = \{\mathcal{X}_i^{(1)}, \mathcal{X}_i^{(2)}, \dots, \mathcal{X}_i^{(|\mathcal{P}|)}\}. \quad (5)$$

Cognitive-Enhanced Evolutionary Search

To evolve high-quality question subsets tailored to individual examinees, we propose a cognitive-enhanced evolutionary search framework comprising two key operators: the schema-preserving uniform crossover operator and the cognitive information-guided mutation operator. Both operators maintain the fixed-length structure of individuals while being guided by the cognitive relevance signals.

Schema-Preserving Uniform Crossover Let two parent individuals be denoted as $\mathcal{X}_i^{(a)} = [x_1^{(a)}, x_2^{(a)}, \dots, x_L^{(a)}]$ and $\mathcal{X}_i^{(b)} = [x_1^{(b)}, x_2^{(b)}, \dots, x_L^{(b)}]$, each representing a candidate question list. To generate two offspring $\mathcal{X}_i^{(c1)}$ and $\mathcal{X}_i^{(c2)}$, we sample a binary mask vector $\mathbf{m} \in \{0, 1\}^L$ with $m_k \sim \text{Bernoulli}(0.5)$, and perform crossover as follows:

$$\begin{cases} x_k^{(c1)} = m_k \cdot x_k^{(a)} + (1 - m_k) \cdot x_k^{(b)}, \\ x_k^{(c2)} = m_k \cdot x_k^{(b)} + (1 - m_k) \cdot x_k^{(a)}, \end{cases} \quad (6)$$

where $1 \leq k \leq L$ denotes the crossover index, and the operator preserves individual structure while enabling fine-grained recombination, outperforming one-point or multi-point crossover in maintaining feasibility and diversity. To ensure that both offspring preserve uniqueness and validity (i.e., no duplicate questions and $\mathcal{X}_i^{(c)} \subset \mathcal{Q}_i^{\text{untested}}$), we apply a repair operator $\mathcal{T}(\cdot)$ that resolves duplicates by replacing them with randomly sampled non-overlapping items from the untested pool. The final offspring are given by:

$$\mathcal{X}_i^{(c1)} \leftarrow \mathcal{T}(\mathcal{X}_i^{(c1)}), \quad \mathcal{X}_i^{(c2)} \leftarrow \mathcal{T}(\mathcal{X}_i^{(c2)}). \quad (7)$$

Cognitive Information-Guided Mutation To introduce adaptive perturbation, we propose a mutation strategy that leverages personalized item information gain. For a given individual $\mathcal{X}_i = [x_1, \dots, x_L]$, we randomly select a gene x_{off} to remove, and then sample a replacement x_{on} from the unselected pool based on an information-based distribution. Specifically, let $\theta_i \in \mathbb{R}^d$ denote the ability vector of examinee s_i , and let $\alpha_j \in \mathbb{R}^d$ be the difficulty vector of item q_j . According to the item response theory (IRT) (Reckase 2009), the probability that s_i correctly answers q_j is computed as: $p_j = \sigma(\theta_i^\top \alpha_j)$, where $\sigma(\cdot) = \frac{1}{1+e^{-\cdot}}$ denotes the sigmoid function. To quantify how informative item q_j is for estimating θ_i , we refer to the Fisher information matrix (Rissanen 1996), which characterizes the expected curvature of the log-likelihood with respect to θ_i , and is defined as:

$$\begin{aligned} \mathbf{I}_j(\theta_i) &= \mathbb{E}\left[\left(\frac{\partial}{\partial \theta} \log p_j(\theta)^{r_{ij}} (1 - p_j(\theta))^{1-r_{ij}}\right) (\cdot)^\top\right], \\ &= p_j(1 - p_j) \cdot \alpha_j \alpha_j^\top \in \mathbb{R}^{d \times d}, \quad \forall j \in \mathcal{Q}_i^{\text{untested}}. \end{aligned} \quad (8)$$

However, directly manipulating this matrix in the mutation operator is computationally inefficient, especially when

comparing information across many candidate items. To address this, we approximate the information matrix using its Frobenius norm (Peng et al. 2018) as a scalar proxy, yielding the scalar information gain for item q_j as follows:

$$\begin{cases} \|\mathbf{I}_j(\boldsymbol{\theta}_i)\|_F = p_j(1 - p_j) \cdot \|\boldsymbol{\alpha}_j \boldsymbol{\alpha}_j^\top\|_F = p_j(1 - p_j) \cdot \|\boldsymbol{\alpha}_j\|^2, \\ \Rightarrow \mathbf{I}_j(\boldsymbol{\theta}_i) = |\boldsymbol{\alpha}_j|^2 \cdot p_j(1 - p_j), \quad \forall j \in \mathcal{Q}_i^{\text{untested}}. \end{cases} \quad (9)$$

Let $\mathcal{Z} = \mathcal{Q}_i^{\text{untested}} \setminus \mathcal{X}_i$ denote the pool of unselected candidate questions. We define a categorical sampling distribution over \mathcal{Z} based on normalized information gain:

$$P(x_j \in \mathcal{Z}) = \frac{\mathbf{I}_j}{\sum_{k \in \mathcal{Z}} \mathbf{I}_k}, \quad (10)$$

where the new gene x_{on} is then sampled from this distribution to replace the removed gene x_{off} , introducing a personalized, cognitively-informed mutation step that promotes high-information test composition. This mutation operator ensures that inserted genes are both personalized and cognitively informative, leading to more effective evolution.

Diversity-Preserving Environmental Selection

To ensure robust convergence and mitigate premature stagnation, we adopt a diversity-preserving environmental selection strategy. This mechanism balances fitness-oriented exploitation with diversity-aware exploration, ultimately forming the next-generation population with both high-quality and semantically diverse candidate question lists.

For each individual $\mathcal{X}_i^{(j)} = [x_1, x_2, \dots, x_L]$, its fitness is assessed by simulating the one-shot assessment process. Specifically, student s_i first completes the selected set of questions, after which the cognitive diagnosis model \mathcal{M}_d performs a virtual parameter update to estimate the personalized knowledge ability vector, following the trajectory $\hat{\boldsymbol{\theta}}_i^0 \xrightarrow{\text{update}} \hat{\boldsymbol{\theta}}_i'$. The updated ability $\hat{\boldsymbol{\theta}}_i'$ is then evaluated on the reserved test set $\mathcal{D}_i^{\text{test}}$, and the prediction quality is measured using a hybrid metric that combines AUC and accuracy:

$$\mathcal{F}(\mathcal{X}_i^{(j)}) = (\mathcal{F}_{\text{auc}}(\hat{\boldsymbol{\theta}}_i', \mathcal{D}_i^{\text{test}}) + \mathcal{F}_{\text{acc}}(\hat{\boldsymbol{\theta}}_i', \mathcal{D}_i^{\text{test}})) / 2, \quad (11)$$

where $\mathcal{F}_{\text{auc}}(\cdot)$ and $\mathcal{F}_{\text{acc}}(\cdot)$ are computed between the predicted responses (based on $\hat{\boldsymbol{\theta}}_i'$ and the true labels in $\mathcal{D}_i^{\text{test}}$). During this process, the model parameters are restored after evaluation to preserve consistency across candidates.

Let $\mathcal{P}_i^g = \{\mathcal{X}_i^{(1)}, \dots, \mathcal{X}_i^{(|\mathcal{P}|)}\}$ denote the population of student s_i at generation g , with corresponding fitness values $\mathcal{F}_i^g = [f_i^1, \dots, f_i^{|\mathcal{P}|}]$. We sort the individuals in descending order of fitness and retain the top- k elites as:

$$\mathcal{E}_g^1 = \{\mathcal{X}_i^{(j)} \in \mathcal{P}_i^g \mid \text{rank}(f_i^j) \leq k\}, \quad k = \left\lfloor \frac{|\mathcal{P}|}{2} \right\rfloor. \quad (12)$$

To preserve diversity, the remaining individuals are selected by filtering the rest of the population based on Hamming distance. Specifically, each candidate is encoded as a binary bit-string $\mathbf{b}_i^j = \text{Pack}(\mathcal{X}_i^{(j)})$ and compared against the elite pool $\mathcal{B}_\mathcal{E}$ via batch Hamming distance:

$$\text{HamDist}(\mathbf{b}_i^j, \mathcal{B}_\mathcal{E}) = \min_{\mathbf{b}_i^e \in \mathcal{B}_\mathcal{E}} \text{Hamming}(\mathbf{b}_i^j, \mathbf{b}_i^e), \quad (13)$$

Dataset	JUNYI	PTADisc
#Learners	54,564	18,768
#Exercises	565	3,262
#Knowledge concepts	30	50
#Interactions	1,711,210	5,720,582
Avg. interactions per learner	31.36	304.80
Avg. exercises per concept	18.83	70.06

Table 1: The statistics of all datasets.

where only those candidates satisfying $\text{HamDist} > \tau$ are admitted to the survivor set, and τ is a threshold (e.g., $\tau = 0.15L$). This filtering is repeated until the survivor set reaches the desired size, or a maximum number of attempts is reached. The final population is formed as:

$$\mathcal{P}_i^{g+1} = \mathcal{E}_g^1 \cup \mathcal{E}_g^2, \quad \text{s.t. } |\mathcal{P}_i^{g+1}| = |\mathcal{P}|, \quad (14)$$

where \mathcal{E}_i^g contains the diversity-preserved candidates sampled under the Hamming constraints.

Experiments

Experimental Setting

Datasets. We conducted experiments on two real-world educational datasets of different scales and characteristics, JUNYI (Chang et al. 2015) and PTADisc (Hu et al. 2023), to evaluate the effectiveness of the proposed PEOAT on the one-shot adaptive testing (OAT) task. The statistical overview of both datasets is presented in Table 1.

Baseline Approaches. To demonstrate the effectiveness of the proposed model, we compare it with a comprehensive set of computerized adaptive testing approaches, including both heuristic and data-driven methods. In total, eight CAT algorithms are considered: RAND, MKLI (Chang 2015), MAAT (Bi et al. 2020), BECAT (Zhuang et al. 2023), BOBCAT (Ghosh and Lan 2021), NCAT (Zhuang et al. 2022), GMOCAT (Wang et al. 2023), and UATS (Yu et al. 2024a).

Evaluation Metrics. The goal of the OAT task is to maximize the quality of ability assessment. Following the evaluation protocol commonly used in traditional CAT settings, we adopt two standard metrics to assess model performance: the area under the ROC curve (AUC) and accuracy (ACC).

Experimental Settings. In our experiment, we adopt MIRT (Reckase 2009) and NCD (Wang et al. 2020) as the backbone diagnosis models of the ability estimation module. During the pre-training of \mathcal{M}_d , the student and item embeddings are initialized with dimensions equal to the number of knowledge concepts. In the OAT evaluation phase, the question selection model \mathcal{M}_π adopts consistent settings, where the learning rates for MIRT and NCD updates are set to 0.02 and 0.005, respectively, with $5 * \sqrt{L}$ epochs. The one-shot selection lengths L are set $\{5, 10, 15, 20\}$. We used a population size of 20, 15 evolutionary generations, a crossover rate of 0.8, a mutation rate of 0.2, and search the distance threshold τ in $\{0.5, 0.75, 1, 1.25, 1.5\}$. All models are Xavier-initialized and optimized using Adam in PyTorch, with experiments conducted on two NVIDIA RTX 4090 GPUs.

CDM		MIRT				NCD			
Dataset/Metric		JUNYI / ACC/AUC(%) ↑							
Type	Methods	length=5	length=10	length=15	length=20	length=5	length=10	length=15	length=20
Heuristic	RNAD	67.98/68.24	74.48/73.64	79.60/77.73	82.47/80.48	67.19/79.02	69.38/80.12	71.41/81.08	73.21/81.94
	MKLI	70.14/70.27	78.03/76.64	83.26/81.39	86.07/84.27	67.12/80.45	68.39/81.29	70.63/82.34	72.59/83.26
	MAAT	68.45/69.50	74.55/73.04	77.94/75.23	79.99/77.01	68.66/80.36	70.31/80.84	72.37/81.40	74.41/82.03
	BECAT	67.85/69.24	74.69/73.72	79.92/77.59	83.80/80.90	67.71/80.33	68.91/81.06	69.73/81.60	70.73/82.02
Data-Driven	BOBCAT	69.15/71.86	77.05/ <u>77.60</u>	81.66/81.12	84.29/83.43	<u>70.98/81.62</u>	72.84/82.68	74.51/83.53	76.01/84.35
	NCAT	<u>71.19/73.48</u>	80.23/77.37	82.69/ <u>81.43</u>	<u>84.93/84.04</u>	70.71/80.95	73.18/82.82	74.20/83.46	76.77/ <u>84.59</u>
	GMOCAT	71.47/73.19	79.15/77.53	82.38/80.49	84.74/83.58	70.55/80.86	72.47/81.73	<u>74.69/83.91</u>	76.58/84.04
	UATS	70.83/ <u>74.45</u>	<u>80.33/77.19</u>	<u>83.13/81.27</u>	<u>84.38/84.65</u>	70.44/80.61	<u>73.35/83.16</u>	74.18/82.87	<u>77.52/84.15</u>
Ours	PEOAT	79.64/83.05	85.38/85.85	86.39/86.68	86.85/87.83	74.56/83.06	81.90/86.47	85.85/88.86	87.34/89.78
Dataset/Metric		PTADisc / ACC/AUC(%) ↑							
Type	Methods	length=5	length=10	length=15	length=20	length=5	length=10	length=15	length=20
Heuristic	RAND	61.29/62.93	63.89/63.89	65.93/64.78	66.74/65.52	63.83/66.71	64.57/67.11	65.24/67.47	65.92/67.84
	MKLI	62.14/65.44	64.09/66.73	67.23/ <u>68.24</u>	65.84/67.65	64.95/68.41	64.67/68.51	64.73/68.69	65.16/68.87
	MAAT	61.78/63.29	62.90/63.84	65.51/64.25	65.77/65.63	64.77/68.42	64.71/68.65	64.94/69.01	65.11/69.32
	BECAT	62.34/64.80	63.43/64.85	64.74/65.12	66.10/65.43	65.25/68.42	65.79/68.94	66.64/69.38	67.18/69.52
Data-Driven	BOBCAT	62.87/65.04	<u>65.32/66.17</u>	67.32/67.10	69.01/67.81	65.95/ <u>68.71</u>	66.70/69.13	67.39/69.52	68.08/69.94
	NCAT	63.19/65.23	64.82/ <u>67.45</u>	67.96/67.87	69.55/68.39	66.38/68.09	<u>67.64/69.22</u>	68.48/ <u>69.61</u>	<u>70.17/70.40</u>
	GMOCAT	<u>63.73/65.80</u>	65.11/66.86	67.52/67.27	69.61/68.05	<u>66.47/68.24</u>	67.48/68.97	<u>69.14/69.49</u>	69.73/70.36
	UATS	63.51/64.94	64.90/67.29	<u>68.23/67.51</u>	<u>70.25/68.67</u>	66.09/68.33	67.47/68.86	68.63/69.35	69.81/70.03
Ours	PEOAT	68.10/69.80	71.96/71.59	73.05/72.36	74.17/72.65	69.37/70.84	73.65/73.58	75.44/75.07	75.91/74.93

Table 2: Performance comparison of PEOAT and baselines on the JUNYI and PTADisc datasets in terms of ACC and AUC. **Bold** highlights the best performance (statistically significant at $p < 0.05$), and underline marks the second-best performance.

Performance Comparison

Table 2 presents the experimental results of the proposed PEOAT model for one-shot adaptive testing, compared with all baseline methods on the two datasets. The best performance for each metric is highlighted in bold, while the second-best is underlined. According to the results, there are several observations: (1) PEOAT consistently outperforms all state-of-the-art baselines across both datasets and question lengths. Specifically, compared to the second-best model, it achieves average improvements of 7.74% and 5.82% in ACC and AUC on the JUNYI dataset, and 6.97% and 5.62% in ACC and AUC on the PTADisc dataset, respectively. This consistent advantage suggests that PEOAT’s personalization-guided selection effectively aligns exercise assembly with individual diagnostic objectives; (2) The superiority of PEOAT is particularly pronounced at shorter testing lengths. For instance, on the JUNYI dataset under the MIRT diagnosis model, PEOAT outperforms the second-best baseline by 10.61% in ACC and 10.35% in AUC at $length = 5$, and by 3.77% and 6.05% in ACC and AUC at $length = 15$, respectively. These results further highlight the strong potential and practical applicability of PEOAT in fast, one-shot question assembly scenarios. In addition, we compared the performance of the basic version of PEOAT without targeted design, as presented in Table 3. The results

demonstrate that formulating the OAT task as a combinatorial optimization problem and incorporating the evolutionary algorithm significantly enhance performance, further validating the superiority of the proposed PEOAT model.

Ablation Study

We conducted a comprehensive ablation study to investigate the contribution of each module in the PEOAT framework by defining the following variants: 1) **w/o PI**: removing the personalization-aware population initialization and replacing it with random initialization only; 2) **w/o CE**: removing the cognitive-enhanced evolutionary search strategy and replacing it with basic crossover and mutation operations; 3) **w/o ES**: removing the diversity-preserving environmental selection. To conserve space, we provide the accuracy results of MIRT as a basic diagnosis model on the JUNYI dataset. As illustrated in Figure 3, the results reveal insightful observations: (1) Compared to PEOAT, all variants exhibit relative performance degradation, highlighting the contribution of the designed sub-modules to our proposed model. (2) The most significant performance drop occurs when the population initialization strategy is removed, indicating that the incorporation of personalized information substantially enhances the quality of the initial population.

Dataset/Metric		JUNYI / ACC/AUC(%) \uparrow			
CDM	Methods	length=5	length=10	length=15	length=20
MIRT	PEOAT-B	78.35/81.97	84.12/84.48	84.96/85.21	85.73/86.55
	PEOAT	79.64/83.05	85.38/85.85	86.39/86.68	86.85/87.83
NCD	PEOAT-B	73.27/81.81	80.64/85.19	84.73/87.59	86.11/88.80
	PEOAT	74.56/83.06	81.90/86.47	85.85/88.86	87.34/89.78

Dataset/Metric		PTADisc / ACC/AUC(%) \uparrow			
CDM	Methods	length=5	length=10	length=15	length=20
MIRT	PEOAT-B	66.57/68.24	70.44/70.25	71.71/71.08	72.86/71.30
	PEOAT	68.10/69.80	71.96/71.59	73.05/72.36	74.17/72.65
NCD	PEOAT-B	67.93/69.41	72.26/72.09	74.67/73.88	74.79/73.53
	PEOAT	69.37/70.84	73.65/73.58	75.44/75.07	75.91/74.93

Table 3: Performance comparison of PEOAT and its base version PEOAT-B on the JUNYI and PTADisc datasets.




Figure 3: Performance of ablation studies conducted on the JUNYI, where “w/o” means removing the target module.

Parameter Sensitivity Analysis

In this section, we conducted a parameter sensitivity analysis to examine the impact of key hyper-parameters, with a primary focus on the distance threshold τ used in the diversity-preserving environmental selection. Specifically, we set τ to $\{0.5, 0.75, 1.0, 1.25, 1.5\}$, and primarily report the experimental results on the JUNYI dataset. As shown in Figure 4, the model achieves its best performance when τ is set to 1.0, under testing lengths of 10 and 20. Notably, as the threshold varies, the model’s performance does not exhibit a strictly consistent pattern or a clear linear trend. Nevertheless, the overall tendency roughly follows an initial increase followed by a decrease, which may be impacted by the testing length.

Case Study

To further investigate the evolution of question populations and the convergence of search strategies in PEOAT’s question selection, we conduct two case studies in this section. Specifically, 20 students with similar ability levels from the JUNYI dataset are selected, and their ability estimation performance (accuracy and fitness) is tracked during population evolution under varying test lengths, using MIRT as the base model. Figure 5 presents the performance evolution with error bands under two metrics. It can be observed that the assessment performance of individual students improves significantly as the population evolves across different test lengths, particularly in terms of fitness, highlighting the effectiveness of PEOA in evolutionary search. Meanwhile, we also sampled two student groups and visualized the evolution of their overall assessment performance using cloud-rain plots. As shown in Figure 6, both groups ex-




Figure 4: Sensitivity analysis of the distance threshold τ of the environmental selection on the JUNYI dataset.




Figure 5: Case study of the performance evolution of the assembled question populations on the JUNYI dataset.




Figure 6: Case study of the generational fitness progression across varying test lengths on the JUNYI dataset.

hibit an upward performance trend under test lengths of 5 and 10, gradually converging as the number of generations increases. This indicates that student performance not only improved but also became more consistent over time.

Conclusion

In this paper, we first proposed a novel task called One-Shot Adaptive Testing (OAT). This task posed three major challenges: ensuring student adaptability during optimization, effectively searching an enormous solution space, and alleviating encoding sparsity due to a candidate pool far exceeding test length. To address these, we introduced **PEOAT**, a **P**ersonalization-guided **E**volutionary question assembly framework for **O**ne-shot **A**daptive **T**esting. We first designed a personalization-aware population initialization method that incorporated individual ability and exercise difficulty, using multi-strategy sampling to build a diverse and effective initial search space. Then, we developed a cognitive-enhanced evolutionary search incorporating schema-preserving crossover and cognitive information-guided mutation operators to enable efficient exploration. Finally, a diversity-preserving environmental selection strategy was implemented to maintain population diversity while considering fitness. Extensive experiments on two real educational datasets demonstrated the model’s effectiveness, and additional case studies provided valuable insights.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (No. U21A20512, No.62107001, No.62302010), in part by the Anhui Province Key Laboratory of Intelligent Computing and Applications (No. AFZNJS2024KF01), and in part by the Anhui Provincial Natural Science Foundation(No. 2508085MF160).

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.
- Blum, C.; and Roli, A. 2003. Metaheuristics in combinatorial optimization: Overview and conceptual comparison. *ACM computing surveys (CSUR)*, 35(3): 268–308.
- Bu, C.; Cao, Z.; He, C.; and Zhang, Y. 2023. Probabilistic model with evolutionary optimization for cognitive diagnosis. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 891–899.
- Bu, C.; Liu, F.; Cao, Z.; Li, L.; Zhang, Y.; Hu, X.; and Luo, W. 2022. Cognitive diagnostic model made more practical by genetic algorithm. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 7(2): 447–461.
- Chang, H.-H. 2015. Psychometrics behind computerized adaptive testing. *Psychometrika*, 80(1): 1–20.
- Chang, H.-H.; and Ying, Z. 1996. A global information approach to computerized adaptive testing. *Applied Psychological Measurement*, 20(3): 213–229.
- Chang, H.-S.; Hsu, H.-J.; Chen, K.-T.; et al. 2015. Modeling Exercise Relationships in E-Learning: A Unified Approach. In *EDM*, 532–535.
- Črepinšek, M.; Liu, S.-H.; and Mernik, M. 2013. Exploration and exploitation in evolutionary algorithms: A survey. *ACM computing surveys (CSUR)*, 45(3): 1–33.
- Das, S.; and Suganthan, P. N. 2010. Differential evolution: A survey of the state-of-the-art. *IEEE transactions on evolutionary computation*, 15(1): 4–31.
- Dong, Z.; Chen, J.; and Wu, F. 2025. Knowledge is power: Harnessing large language models for enhanced cognitive diagnosis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 164–172.
- Gao, W.; Liu, Q.; Wang, H.; Yue, L.; Bi, H.; Gu, Y.; Yao, F.; Zhang, Z.; Li, X.; and He, Y. 2024a. Zero-1-to-3: Domain-level zero-shot cognitive diagnosis via one batch of early-bird students towards three diagnostic objectives. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 8417–8426.
- Gao, W.; Liu, Q.; Yue, L.; Yao, F.; Lv, R.; Zhang, Z.; Wang, H.; and Huang, Z. 2025. Agent4edu: Generating learner response data by generative agents for intelligent education systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 23923–23932.
- Gao, W.; Liu, Q.; Yue, L.; Yao, F.; Wang, H.; Gu, Y.; and Zhang, Z. 2024b. Collaborative Cognitive Diagnosis with Disentangled Representation Learning for Learner Modeling. *arXiv preprint arXiv:2411.02066*.
- Ghosh, A.; and Lan, A. 2021. Bobcat: Bilevel optimization-based computerized adaptive testing. *arXiv preprint arXiv:2108.07386*.
- Haq, F. U.; Shin, D.; Nejati, S.; and Briand, L. 2021. Can offline testing of deep neural networks replace their online testing? a case study of automated driving systems. *Empirical Software Engineering*, 26(5): 90.
- Holman, L. 2000. A comparison of computer-assisted instruction and classroom bibliographic instruction. *Reference & User Services Quarterly*, 53–60.
- Holmes, W.; Bialik, M.; and Fadel, C. 2023. Artificial intelligence in education. Globethics Publications.
- Hu, L.; Dong, Z.; Chen, J.; Wang, G.; Wang, Z.; Zhao, Z.; and Wu, F. 2023. PTADisc: a cross-course dataset supporting personalized learning in cold-start scenarios. *Advances in Neural Information Processing Systems*, 36: 44976–44996.
- Lambora, A.; Gupta, K.; and Chopra, K. 2019. Genetic algorithm-A literature review. In *2019 international conference on machine learning, big data, cloud and parallel computing (COMITCon)*, 380–384. IEEE.
- Lezcano Casado, M. 2019. Trivializations for gradient-based optimization on manifolds. *Advances in Neural Information Processing Systems*, 32.
- Li, M.; Tong, J.; Huang, Y.; Ding, Y.; Qian, H.; and Zhou, A. 2025. Paper-Level Computerized Adaptive Testing for High-Stakes Examination via Multi-Objective Optimization. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, 1435–1446.
- 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. 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.; Yang, Y.; Qin, C.; Yu, X.; Yang, S.; Zhang, X.; and Zhu, H. 2024b. HD-KT: Advancing Robust Knowledge Tracing via Anomalous Learning Interaction Detection. In *Proceedings of the ACM on Web Conference 2024*, 4479–4488.
- Ma, H.; Yin, Y.; Wang, Z.; Wang, C.; Yu, X.; Yang, S.; and Zhang, X. 2025a. Learning Patterns-Guided Data Generation for Knowledge Tracing. In *Proceedings of the 31st*

- ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2, 2031–2041.
- Ma, H.; Zhou, W.; Yu, X.; Wang, C.; Yang, S.; Zhang, L.; and Zhang, X. 2025b. Reconciling Efficiency and Effectiveness of Exercise Retrieval: An Uncertainty Reduction Hashing Approach for Computerized Adaptive Testing. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- Meyer, G. J.; Finn, S. E.; Eyde, L. D.; Kay, G. G.; Moreland, K. L.; Dies, R. R.; Eisman, E. J.; Kubiszyn, T. W.; and Reed, G. M. 2001. Psychological testing and psychological assessment: A review of evidence and issues. *American psychologist*, 56(2): 128.
- Papadimitriou, C. H.; and Steiglitz, K. 1998. *Combinatorial optimization: algorithms and complexity*. Courier Corporation.
- Peng, X.; Lu, C.; Yi, Z.; and Tang, H. 2018. Connections Between Nuclear-Norm and Frobenius-Norm-Based Representations. *IEEE Transactions on Neural Networks and Learning Systems*, 29(1): 218–224.
- Reckase, M. D. 2009. Multidimensional item response theory models. In *Multidimensional item response theory*, 79–112. Springer.
- Rissanen, J. 1996. Fisher information and stochastic complexity. *IEEE Transactions on Information Theory*, 42(1): 40–47.
- Roll, I.; and Wylie, R. 2016. Evolution and revolution in artificial intelligence in education. *International Journal of Artificial Intelligence in Education*, 26: 582–599.
- Sun, S.; Hu, X.; Bu, C.; Liu, F.; Zhang, Y.; and Luo, W. 2022. Genetic algorithm for bayesian knowledge tracing: A practical application. In *International Conference on Swarm Intelligence*, 282–293. Springer.
- 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. 2020. 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.
- Yang, S.; Chen, M.; Wang, Z.; Yu, X.; Zhang, P.; Ma, H.; and Zhang, X. 2025a. DisenGCD: A Meta Multigraph-assisted Disentangled Graph Learning Framework for Cognitive Diagnosis. *Advances in Neural Information Processing Systems*, 37: 91532–91559.
- Yang, S.; Han, J.; Yu, X.; Wang, Z.; Jiang, H.; Ma, H.; Zhang, X.; and Min, G. 2025b. FedCD: A Fairness-aware Federated Cognitive Diagnosis Framework. *arXiv preprint arXiv:2508.01296*.
- 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.; Wei, H.; Ma, H.; Tian, Y.; Zhang, X.; Cao, Y.; and Jin, Y. 2023a. Cognitive diagnosis-based personalized exercise group assembly via a multi-objective evolutionary algorithm. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 7(3): 829–844.
- Yang, S.; Yu, X.; Tian, Y.; Yan, X.; Ma, H.; and Zhang, X. 2023b. Evolutionary neural architecture search for transformer in knowledge tracing. *Advances in Neural Information Processing Systems*, 36: 19520–19539.
- Yu, J.; Zhuang, Y.; Huang, Z.; Liu, Q.; Li, X.; Li, R.; and Chen, E. 2024a. A unified adaptive testing system enabled by hierarchical structure search. In *Forty-first International Conference on Machine Learning*.
- Yu, X.; Qin, C.; Shen, D.; Ma, H.; Zhang, L.; Zhang, X.; Zhu, H.; and Xiong, H. 2024b. Rdtg: enhancing group cognitive diagnosis with relation-guided dual-side graph transformer. *IEEE Transactions on Knowledge and Data Engineering*, 36(7): 3429–3442.
- Yu, X.; Qin, C.; Shen, D.; Yang, S.; Ma, H.; Zhu, H.; and Zhang, X. 2024c. 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. 2024d. DISCO: A Hierarchical Disentangled Cognitive Diagnosis Framework for Interpretable Job Recommendation. In *IEEE International Conference on Data Mining (ICDM) 2024*.
- Yu, X.; Yang, S.; Li, J.; Wang, Z.; Qin, C.; Ma, H.; and Zhang, X. 2025a. Rethinking Learner Modeling: A Feedback-Centric Cognitive Disentanglement Perspective. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, 3657–3668.
- Yu, X.; Yang, S.; Wang, Z.; Song, S.; Ma, H.; Cao, Z.; and Zhang, X. 2025b. LIGHT: Enhancing Learning Path Recommendation via Knowledge Topology-Aware Sequence Optimization. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- 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. *Advances in Neural Information Processing Systems*, 36: 2381–2402.