

Deep Computerized Adaptive Testing

Jiguang Li*, Robert Gibbons† and Veronika Ročková‡

March 17, 2026

Abstract

Computerized adaptive tests (CATs) play a crucial role in educational assessment and diagnostic screening in behavioral health. Unlike traditional linear tests that administer a fixed set of pre-assembled items, CATs adaptively tailor the test to an examinee’s latent trait level by selecting a smaller subset of items based on their previous responses. Existing CAT applications predominantly rely on item response theory (IRT) models with a single latent variable, a choice driven by both conceptual simplicity and computational feasibility. However, many real-world item response datasets exhibit complex, multi-factor structures, limiting the applicability of CATs in broader settings. In this work, we develop a novel CAT system that incorporates multivariate latent traits, building on recent advances in Bayesian sparse multivariate IRT [1]. Our approach leverages direct sampling from the latent factor posterior distributions, significantly accelerating existing information-theoretic item selection criteria by eliminating the need for computationally intensive Markov Chain Monte Carlo (MCMC) simulations. Recognizing the potential suboptimality of existing item selection rules which are often based on myopic one-step-lookahead optimization of some information-theoretic criterion, we propose a double deep Q-learning algorithm to learn an optimal item selection policy. Through simulation and real-data studies, we demonstrate that our approach not only accelerates existing item selection methods but also highlights the potential of reinforcement learning in CATs. Notably, our Q-learning-based strategy consistently achieves the fastest posterior variance reduction, leading to earlier test termination. In finite-horizon settings, it also yields final posterior distributions that more closely approximate the oracle posterior distribution—the posterior that would be obtained if test takers had responded to all items in the item bank.

Keywords: Reinforcement Learning; Multidimensional Item Response Theory; Adaptive Data Collection; Deep Q-Learning;

*Jiguang Li is a 4th-year doctoral student in Econometrics and Statistics at the Booth School of Business of the University of Chicago

†Robert Gibbons is the Blum-Reise Professor of Statistics at the University of Chicago

‡Veronika Ročková is the Bruce Lindsay Professor of Econometrics and Statistics in the Wallman Society of Fellows

1 Introduction

Multidimensional Computerized Adaptive Testing (MCAT) has revolutionized the field of educational and psychological assessments by dynamically selecting tailored items from a large test pool, thereby enhancing the efficiency and precision of latent ability estimates [2]. Powered by Multidimensional Item Response Theory (MIRT) [3], MCAT leverages multidimensional statistical inference to evaluate respondents' multidimensional latent traits, and allows for more comprehensive and efficient assessments compared to unidimensional approaches. MCAT's adaptability and accuracy are also particularly crucial in high-stakes diagnostic assessments such as in clinical psychology and psychiatry, where it can be substituted for in-person assessments by clinical professionals especially in areas with limited medical resources [4].

A substantial body of research in MCAT has focused on item selection strategies derived from experimental design principles. One prominent strategy is to select items that maximize the determinant of the Fisher information matrix evaluated at the current estimates of the latent traits [5, 6], known as the D-optimality criterion. An alternative, the A-optimality criterion, aims to minimize the trace of the asymptotic covariance matrix, thereby reducing overall estimation variance [7]. In the absence of nuisance latent abilities, both A-optimality and D-optimality have demonstrated superior accuracy relative to other common experimental design criteria [8].

Another prominent line of MCAT item selection algorithms leverages Kullback-Leibler (KL) information, often within a Bayesian framework. A common approach is to select items that produce response distributions at the true latent trait value, θ_0 , that differs maximally from the response distributions generated at the other value of θ [9, 10]. Moving beyond response distributions alone, some researchers propose maximizing the KL divergence between the current posterior distribution and the posterior distribution at the next

selection step, thereby enhancing adaptation through updated trait estimates [11]. Another promising strategy is the mutual information criterion, which aims to maximize entropy reduction of the current posterior distribution, encouraging more and more accurate posterior estimates [12]. In particular, [13] demonstrates both theoretical and empirical advantages of the Bayesian mutual information item selection rule over common experimental criteria such as D-optimality. More detailed theoretical comparison of KL information and Fisher information criteria is presented by [14].

Although numerous effective item selection rules have been proposed in the MCAT literature, they all rely on one-step lookahead optimization of an information-theoretic criterion. Despite their ease of implementation and attractive theoretical properties, these selection rules are inherently myopic: they select items based solely on immediate information gain, ignoring how current choices influence future decisions, which can lead to suboptimal policies. For example, existing methods tend to favor items with high loading parameters [15]. However, CAT researchers also recommend reserving such items for the later stages of testing to improve efficiency [16]. Integrating heuristic guidance into existing selection rules remains challenging.

Addressing these limitations, we propose a novel deep CAT system that integrates a flexible Bayesian MIRT model with a non-myopic online item selection policy, guided by reinforcement learning (RL) principles [17]. Leveraging recent advancements in Bayesian sparse MIRT [1], our framework seamlessly accommodates multiple latent factors with complex loading structures, while maintaining scalability in both the number of items and factors. To learn the optimal item selection policy that prioritizes the assessment of target factors, we draw on contemporary RL methodologies and introduce a general double deep Q-learning algorithm [18, 19]. This algorithm efficiently trains a Q-network offline using only item parameter estimates; the learned network can then be deployed online to select optimal items based on the current multivariate latent factor posterior distribution. When

the test terminates, our framework robustly characterizes the full latent factor posterior distributions rather than providing only a point estimate.

A primary contribution of our work is to show how identification of the latent-factor posterior distribution leads to substantial computational gains in online item selection. Given that such posterior distribution is deemed to be non-Gaussian and unknown, traditional Bayesian methods typically rely on expensive MCMC simulations [20] and combined with additional data augmentation to handle categorical likelihood [21, 22]. Our approach achieves substantial acceleration by directly sampling latent factor posterior distributions [1, 23, 24]. Notably, this improvement not only increases the efficiency of existing Bayesian item-selection procedures but also provides a computational foundation for training our proposed Q-learning algorithm through rapid, large-scale simulations of testing sessions.

Another critical advancement in our work is the integration of CAT within a reinforcement learning framework. This approach addresses the practical need to prioritize accurate estimation and to overcome known limitations of greedy item selection methods in sequential decision-making [25]. Building on the Bayesian MIRT foundation, our neural-network architecture incorporates the identified posterior parameters as state variables, allowing the model to learn optimal item-selection policies through a large amount of testing simulations. This formulation bridges the two paradigms: the Bayesian component provides statistically grounded representations of examinee uncertainty, while the reinforcement-learning component leverages these representations to optimize sequential decisions. The trained neural network is deployable on standard laptops without GPU acceleration, making it suitable for online adaptive testing applications. The sequential nature of CAT aligns naturally with deep Q-learning methods, which have demonstrated remarkable success across diverse application domains [26, 27]. Notably, RL has been successfully employed in educational measurement settings to design personalized learning plans [28, 29].

Finally, our work aligns with the emerging research trend of framing traditional sta-

tistical sequential decision-making problems as optimal policy learning tasks [30]. This perspective has appeared in Bayesian adaptive design [31, 32, 33] and Bayesian optimization [34, 35], where recent methods aim to move beyond one-step criteria toward policy learning that account for long-term consequences. Our contribution fits within this broader trend by providing a principled approach to cognitive and behavioral assessment.

The remainder of the paper is organized as follows. Section 2 motivates our deep CAT framework using a cognitively complex item bank from a recent clinical study [36]. Section 3 reviews existing information-theoretic item selection methods and outlines the necessity to reformulate CAT as a reinforcement learning task. Section 4 introduces a general Bayesian framework that accelerates existing CAT item-selection rules and serves as the foundation for our Q-learning algorithm. Section 5 details our reinforcement learning approach, including the neural network architecture and double Q-learning algorithm. Finally, Sections 6 and 7 evaluate our method through both simulations and real data experiments. The implementations of the baseline Bayesian CAT algorithms, our Q-learning algorithm, and the driver codes to reproduce all the experiments in the paper can be found in https://github.com/JiguangLi/deep_CAT.

2 Adaptive High-Dimensional Cognitive Assessment

Our proposed deep CAT system is motivated by the growing need for adaptive cognitive assessment in high-dimensional latent spaces. Cognitive impairment, particularly Alzheimer’s dementia (AD), is a major public health challenge, affecting 6.9 million individuals in the U.S. in 2024 [37]. Early detection of mild cognitive impairment (MCI), a precursor to AD, is crucial for slowing disease progression and improving patient outcomes [38]. However, traditional neuropsychological assessments are costly, time-consuming, and impractical for frequent use, underscoring the need for more efficient and scalable assessment methods.




Figure 1: Estimated Bifactor Factor Loading Matrix for pCAT-COG

Recently, pCAT-COG, the first computerized adaptive test (CAT) item bank based on MIRT for cognitive assessment, demonstrated its potential as an alternative to clinician-administered evaluations [36]. The data were collected from 730 participants. After careful item calibration and model selection, the final item bank consisted of $J = 57$ items covering five cognitive subdomains: episodic memory, working memory, executive function, semantic memory, and processing speed. Since each item comprised three related tasks, we used a binary score, where 1 indicates correct answers on all tasks for our analysis.

Following [36], we fit a six-factor bifactor model [39] to the 730 by 57 binary response matrix, with one primary factor representing the global cognitive ability and five secondary factors. This yields a 57 by 6 factor loading matrix, visualized in Figure 1, where rows correspond to pCAT-COG items and columns represent distinct factors.

Since items in pCAT-COG are primarily designed to measure the general cognitive factor and only partially capture subdomain information, our goal is to develop an item selection strategy that efficiently estimates the primary factor (first column) using as few items as possible while maintaining robustness to subdomain influences. Additionally, the

selection algorithm must be computationally efficient to navigate a six-dimensional latent space in real time for seamless interactive testing. Online item selection is critical, as learning an optimal sequence from 57 items offline presents an intractable combinatorial problem, even in a binary response setting.

Adaptive cognitive assessment is essential for cognitive assessment, as designing items is costly and administering all items is time-consuming. Our deep CAT system requires on average 11.2 items to reduce the posterior variance of the primary factor from 1 to 0.16 (posterior s.d. = 0.4), whereas the next-best mutual information method requires 12.6 items to achieve the same precision. The item bank is currently expanding to 500 items with a more nationally representative sample. Success on this prototype dataset paves the way for broader deployment in clinical research.

3 From One-Step Optimization to Policy Learning

We formally define the problem of designing a CAT system from a reinforcement learning (RL) perspective. Typical CAT systems consist of two components:

- **Offline Calibration:** An MIRT model is fitted to a calibrated dataset $\mathbf{Y} \in \mathbb{R}^{N \times J}$ to estimate the item characteristic parameters, where N represents the number of examinees, and J represents the number of items in the item bank.
- **Online Deployment:** Given the estimated item parameters, an item selection algorithm is deployed online to adaptively select items for future examinees.

The performance of CAT can be measured by the number of items required to estimate an examinee’s latent traits with sufficient precision.

3.1 Notation and Problem Formulation

For the entirety of the paper, we assume the calibration dataset \mathbf{Y} is binary, where each element y_{ij} represents whether subject i answered item j correctly. We consider a general two-parameter MIRT framework with K latent factors [3]. Let $\mathbf{B} \in \mathbb{R}^{J \times K}$ denote the factor loading matrix, and $\mathbf{D} \in \mathbb{R}^J$ denote the intercept vector. Throughout the paper, boldface notation is used to denote vectors and matrices (e.g., $\boldsymbol{\theta}$, \mathbf{B} , \mathbf{D}), while scalar quantity such as D_j is written in standard font. For each examinee i with multivariate latent trait $\boldsymbol{\theta}^{(i)} \in \mathbb{R}^K$, the data-generating process for y_{ij} is given by:

$$y_{ij} \mid \boldsymbol{\theta}^{(i)}, \mathbf{B}_j, D_j \sim \text{Bernoulli}(\Phi(\mathbf{B}'_j \boldsymbol{\theta}^{(i)} + D_j)), \quad (1)$$

where \mathbf{B}_j is the j -th row of \mathbf{B} , D_j is the j -th entry of \mathbf{D} , and $\Phi(\cdot)$ denotes the standard normal cumulative distribution function. The item parameters can be compactly expressed as $\{\boldsymbol{\xi}_j\}_{j=1}^J := \{(\mathbf{B}_j, D_j)\}_{j=1}^J$. This two-parameter MIRT framework is highly general, imposing no structural constraints on the loading matrix and requiring no specific estimation algorithms for item parameter calibration.

Given the estimated item parameters, we need to design an item selection algorithm that efficiently tests a future examinee with unobserved latent ability $\boldsymbol{\theta} \in \mathbb{R}^K$. The sequential nature of the CAT problem makes Bayesian approaches particularly appealing. Without loss of generality, we assume a standard multivariate Gaussian prior on $\boldsymbol{\theta} \sim \mathcal{N}(0, \mathbb{I}_K)$, and introduce the following notation:

- For any positive integer n , let $[n]$ denote all the positive integers no greater than n .
Let $j \in [J]$ be the index for the items in the item bank.
- Let j_t denote the index of the item selected at step t based on an arbitrary item selection rule. Define \mathcal{I}_t as the set of the first t administered items, and let $R_t :=$

$[J] \setminus \mathcal{I}_{t-1}$, the set of available items before the t -th item is picked.

- We use the shorthand notation $f(\boldsymbol{\theta}|\mathbf{Y}_{1:T})$ to represent the latent factor posterior distributions $f(\boldsymbol{\theta}|\mathbf{Y}_{1:T}, \boldsymbol{\xi}_{1:T})$ after T items have been selected. Here, the response history is denoted by $\mathbf{Y}_{1:T} := [y_{j_1}, \dots, y_{j_T}]'$, and the item parameters are given by $\boldsymbol{\xi}_{1:T} := (\mathbf{B}_{1:T}, \mathbf{D}_{1:T})$, where $\mathbf{B}_{1:T} := [\mathbf{B}_{j_1}, \dots, \mathbf{B}_{j_T}]'$ and $\mathbf{D}_{1:T} := [D_{j_1}, \dots, D_{j_T}]'$.

At time $(t+1)$, the algorithm takes the current posterior $f(\boldsymbol{\theta}|\mathbf{Y}_{1:t})$ as input and outputs the next item selection j_{t+1} . This process continues until at time T' , either when the posterior variance of $f(\boldsymbol{\theta}|\mathbf{Y}_{1:T'})$ falls below a predefined threshold τ^2 , or when $T' = H$, where $H \leq J$ is the maximum number of items that can be administered. Hence the goal of CAT is to minimize T' given the estimated item parameters of the item bank.

3.2 Reviews of KL Information Item Selection Rules

Given that our proposed deep CAT system is built on a general Bayesian MIRT framework [1], we briefly revisit common Bayesian item selection rules and show in Section 4.1 how our framework can be used to accelerate these baseline methods. The popular KL Expected A Priori (EAP) rule selects the t -th item based on the average KL information between response distributions on the candidate item at the EAP estimate $\hat{\boldsymbol{\theta}}_{t-1} = \int_{\boldsymbol{\theta}} \boldsymbol{\theta} f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)}) d\boldsymbol{\theta}$, and random factor $\boldsymbol{\theta}$ sampled from the posterior distribution $f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)})$ [10]:

$$\arg \max_{j_t \in R_t} \int_{\boldsymbol{\theta}} \left\{ \sum_{l=0}^1 P(y_{j_t} = l | \hat{\boldsymbol{\theta}}_{t-1}) \log \frac{P(y_{j_t} = l | \hat{\boldsymbol{\theta}}_{t-1})}{P(y_{j_t} = l | \boldsymbol{\theta})} \right\} f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)}) d\boldsymbol{\theta}. \quad (2)$$

Rather than focusing on the KL information based on the response distributions, [11] propose the MAX Pos approach by maximizing the KL information between two subsequent latent factor posteriors $f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)})$ and $f(\boldsymbol{\theta}|\mathbf{Y}_{1:t})$. Intuitively, this approach prioritizes

items that induce the largest shift in the posterior, formalized as:

$$\arg \max_{j_t \in R_t} \sum_{y_{j_t}=0}^1 f(y_{j_t} | \mathbf{Y}_{1:(t-1)}) \int_{\boldsymbol{\theta}} f(\boldsymbol{\theta} | \mathbf{Y}_{1:(t-1)}) \log \frac{f(\boldsymbol{\theta} | \mathbf{Y}_{1:(t-1)})}{f(\boldsymbol{\theta} | \mathbf{Y}_{1:t})} d\boldsymbol{\theta}, \quad (3)$$

where $f(y_{j_t} | \mathbf{Y}_{1:(t-1)})$ represents the posterior predictive probability:

$$f(y_{j_t} | \mathbf{Y}_{1:(t-1)}) = \int_{\boldsymbol{\theta}} f(y_{j_t} | \boldsymbol{\theta}) f(\boldsymbol{\theta} | \mathbf{Y}_{1:(t-1)}) d\boldsymbol{\theta}. \quad (4)$$

A third Bayesian strategy, the mutual information (MI) approach, maximizes the mutual information between the current posterior distribution and the response distribution of new item y_{j_t} [12]. Rooted in experimental design theory [40], we can also interpret mutual information as the entropy reduction of the current posterior distribution of $\boldsymbol{\theta}$ after observing new response y_{j_t} . More formally:

$$\arg \max_{j_t \in R_t} I_M(\boldsymbol{\theta}, y_{j_t}) = \arg \max_{j_t \in R_t} \sum_{y_{j_t}=0}^1 \int_{\boldsymbol{\theta}} f(\boldsymbol{\theta}, y_{j_t} | \mathbf{Y}_{1:(t-1)}) \log \frac{f(\boldsymbol{\theta}, y_{j_t} | \mathbf{Y}_{1:(t-1)})}{f(\boldsymbol{\theta} | \mathbf{Y}_{1:(t-1)}) f(y_{j_t} | \mathbf{Y}_{1:(t-1)})} d\boldsymbol{\theta}. \quad (5)$$

Given the impressive empirical success of the MI method [13], we also introduce a competitive heuristic item-selection rule that chooses items with the highest predictive variance under the current posterior estimates, serving as an additional benchmark method. Formally, write $c_{j_t} := f(y_{j_t} | \mathbf{Y}_{1:(t-1)})$ as defined in (4), and consider the following criterion:

$$\arg \max_{j_t \in R_t} \int_{\boldsymbol{\theta}} (\Phi(\mathbf{B}'_{j_t} \boldsymbol{\theta} + D_{j_t}) - c_{j_t})^2 f(\boldsymbol{\theta} | \mathbf{Y}_{1:(t-1)}) d\boldsymbol{\theta}. \quad (6)$$

This rule selects the item j_t that maximizes the variance of the predictive means, weighted by the current posterior $f(\boldsymbol{\theta} | \mathbf{Y}_{1:(t-1)})$. In Appendix A, we establish a connection between this selection rule and the MI method, showing that both favor items with high prediction uncertainty.

3.3 Reinforcement Learning Formulation

Existing item selection rules face several limitations: first, these methods rely on one-step-lookahead selection, choosing items that provide the most immediate information without considering their impact on future selections, which can result in suboptimal policies [17]. Second, they are heuristically designed to balance information across all latent factors, rather than emphasizing the most essential factors of interest. Finally, these rules do not directly minimize test length, potentially increasing test duration without proportional gains in accuracy.

A more principled approach is to formulate item selection as a reinforcement learning (RL) problem, where the optimal policy is learned using Bellman optimality principles rather than relying on heuristics that ignore long-term planning. Beyond addressing myopia, an RL-based formulation provides a direct mechanism to minimize the number of items required to reach a predefined posterior variance reduction threshold, explicitly guiding item selection toward accurately measuring the primary factors of interests.

More generally, consider a general finite horizon setting where each examinee can answer at most $H \leq J$ items, and the CAT algorithm terminates whenever the posterior variance of all factors of interest is smaller than a certain threshold τ^2 , or when H is reached. Since H can be set sufficiently large, it serves as a practical secondary stopping criterion to prevent excessively long tests. Formally:

- **State Space \mathcal{S} :** the space of all possible latent factor posteriors $\mathcal{S} := \{f(\boldsymbol{\theta}|\mathbf{Y}_{1:t}) : t \in \{1, \dots, H\}\}$. The state variable in CAT represents the Bayesian estimate of the examinee’s latent traits as a multivariate distribution at time t . We may write $s_t := f(\boldsymbol{\theta}|\mathbf{Y}_{1:t})$, with $s_0 := \mathcal{N}(0, \mathbb{I}_K)$. The state space is discrete and can be exponentially large. Since there are J items in the item bank and the responses are binary, the total number of possible states is $\sum_{t=1}^H \binom{J}{t} \cdot 2^t$.

- **Action space \mathcal{A} :** the item bank $[J] := \{1, \dots, J\}$. However, at the t -th selection time, the action space is the remaining items in the test bank that have not yet been selected $R_t := [J] \setminus \mathcal{I}_{t-1}$, as we do not select the same item twice.
- **Transition Kernel $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$,** where $\Delta(\mathcal{S})$ denotes the space of probability distributions over \mathcal{S} . The transition kernel \mathcal{P} specifies the stochastic rule that governs how the state evolves after an action is taken. In the CAT context, the kernel maps the current latent-trait estimate s_t and the selected item a_t to the next posterior state s_{t+1} . Conditional on (s_t, a_t) , the next estimate s_{t+1} depends on the examinee’s binary response $y_t \in \{0, 1\}$ to item a_t , which yields only two possible posterior updates. Because the true ability of the examinee and thus the response probability $p(y_t = 1 \mid s_t, a_t)$ are unknown, the transition kernel in CAT is also unknown and must be approximated.
- **Reward Function:** To directly minimize the test length, we assign a simple 0 – 1 reward structure, where we assign negative rewards whenever more items are needed to reduce posterior variance to a given threshold. In a K -factor MIRT model, we often prioritize a subset of factors $\mathcal{K} \subset [K]$ (e.g. $\mathcal{K} = \{1\}$ for pCAT-COG), with the test terminating once the posterior variances of all factors in \mathcal{K} fall below the predefined threshold τ^2 :

$$R^{(t)}(s_t, a, s_{t+1}) = \begin{cases} -1 & \text{if } V_{t+1} > \tau^2, \\ 0 & \text{otherwise,} \end{cases} \quad (7)$$

where $V_{t+1} = \max_{k \in \mathcal{K}} \text{Var}(\boldsymbol{\theta}_k \mid \mathbf{Y}_{1:(t+1)})$ is the maximum marginal posterior variance among the prioritized factors \mathcal{K} . This also simplifies learning the value function, as the rewards are always bounded integers within $[-H, -1]$. We further illustrate the advantages of adopting this reward structure in Appendix E.

Rather than following a heuristic rule, we learn a policy π that maps the current state (latent factor estimates) to a distribution over potential items:

$$\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A}), \quad a_t \sim \pi(\cdot | s_t).$$

Given an initial state $S_0 = s_0$, and the discount factor $\gamma \in (0, 1]$, the value function is

$$v_\pi(s_0) := \mathbb{E}_\pi \left[\sum_{t=0}^{H-1} \gamma^t R^{(t)}(s_t, \pi(s_t), s_{t+1}) | S_0 = s_0 \right], \quad (8)$$

The expectation $\mathbb{E}_\pi[\cdot]$ is taken over trajectories generated by drawing a_t from $\pi(\cdot | s_t)$ and then s_{t+1} from the transition kernel $\mathcal{P}(\cdot | s_t, a_t)$. Because direct evaluation of (8) is challenging, it is more convenient to consider the action–value function

$$Q_\pi(s_0, a) := \mathbb{E}_\pi \left[\sum_{t=0}^{H-1} \gamma^t R^{(t)}(s_t, \pi(s_t), s_{t+1}) | S_0 = s_0, A_0 = a \right], \quad (9)$$

with $v_\pi(s) = \mathbb{E}_{a \sim \pi(\cdot | s)} Q_\pi(s, a)$. The optimal policy π^* satisfies the Bellman equation [25]:

$$Q_{\pi^*}(s, a) = \mathbb{E}_{s' \sim \mathcal{P}(\cdot | s, a)} \left[R(s, a, s') + \gamma \max_{a'} Q_{\pi^*}(s', a') \right], \quad (10)$$

where s' is the next posterior estimate after applying action a to the current posterior s , and $v_{\pi^*}(s) = \max_a Q_{\pi^*}(s, a)$. Since the state space grows exponentially and the transition kernel is unknown, solving for π^* using traditional dynamic programming approaches becomes intractable [17].

In practice, our deep Q-learning approach [41] does not evaluate these expectations analytically. Let the transition (s, a, r, s') represent one testing step, where item a is selected under state s , the reward r is observed, and the posterior is updated to s' through the Bayesian MIRT model. We approximate the Bellman fixed point in (10) by fitting a

parametric function Q_w that minimizes the squared temporal-difference loss:

$$\mathcal{L}(w) = \mathbb{E}_{(s,a,r,s') \sim \mathcal{D}} [(Q_w(s, a) - y(s, a, r, s'))^2] \quad (11)$$

where

$$y(s, a, r, s') = r + \gamma \max_{a'} Q_{\bar{w}}(s', a'). \quad (12)$$

Here, \mathcal{D} (replay buffer) denotes a collection of previously observed testing steps (s, a, r, s') generated during simulation, which is used to approximate the expectation above by empirical averaging. The parameter \bar{w} corresponds to a delayed copy of the model parameters w , updated less frequently to stabilize the numerical optimization. We describe the full Q-learning algorithm in Section 5.

4 Accelerating Item Selection via Posterior Identification

The central insight of our deep CAT framework is that, by iteratively applying the E-step of the PXL-EM algorithm [1], the latent factor posteriors admit tractable posterior updates. This result is critical for reinforcement learning, as the posterior distribution is deemed to be non-Gaussian and is analytically intractable under the traditional MIRT literature. By obtaining a tractable representation of this posterior, we can parametrize the examinee’s evolving ability and uncertainty as a well-defined state variable, which is an essential prerequisite for applying Q-learning to adaptive testing. The Bayesian MIRT formulation thus not only replaces costly MCMC procedures with efficient posterior updates under a probit link for existing Bayesian item selection rules discussed in Section 3.2, but also supplies the statistical foundation that makes the subsequent reinforcement-learning framework feasible.

Specifically, we show that the latent factor posterior updates during CAT belong to an instance of the unified-skew-normal distribution [42], defined as follows:

Definition 4.1. Let $\Phi_T\{\mathbf{V}; \boldsymbol{\Sigma}\}$ represent the cumulative distribution function of a T -dimensional multivariate Gaussian distribution $N_T(0_T, \boldsymbol{\Sigma})$ evaluated at vector \mathbf{V} . A K -dimensional random vector $\boldsymbol{\theta} \sim \text{SUN}_{K,T}(\boldsymbol{\mu}, \boldsymbol{\Omega}, \boldsymbol{\Delta}, \boldsymbol{\gamma}, \boldsymbol{\Gamma})$ has the **unified skew-normal** distribution if it has the probability density function:

$$\phi_K(\boldsymbol{\theta}; \boldsymbol{\mu}, \boldsymbol{\Omega}) \frac{\Phi_T\{\boldsymbol{\gamma} + \boldsymbol{\Delta}'\bar{\boldsymbol{\Omega}}^{-1}\boldsymbol{\omega}^{-1}(\boldsymbol{\theta} - \boldsymbol{\mu}); \boldsymbol{\Gamma} - \boldsymbol{\Delta}'\bar{\boldsymbol{\Omega}}^{-1}\boldsymbol{\Delta}\}}{\Phi_T(\boldsymbol{\gamma}; \boldsymbol{\Gamma})}.$$

Here, $\phi_K(\boldsymbol{\theta}; \boldsymbol{\mu}, \boldsymbol{\Omega})$ is the density of a K -dimensional multivariate Gaussian with expectation $\boldsymbol{\mu} = (\mu_1, \dots, \mu_K)'$, and a K by K covariance matrix $\boldsymbol{\Omega} = \boldsymbol{\omega}\bar{\boldsymbol{\Omega}}\boldsymbol{\omega}$, where $\bar{\boldsymbol{\Omega}}$ is the correlation matrix and $\boldsymbol{\omega}$ is a diagonal matrix with the square roots of the diagonal elements of $\boldsymbol{\Omega}$ in its diagonal. $\boldsymbol{\Delta}$ is a K by T matrix that determines the skewness of the distribution, and $\boldsymbol{\gamma} \in \mathbb{R}^T$ controls the flexibility in departures from normality.

In addition, the $(K + T) \times (K + T)$ matrix $\boldsymbol{\Omega}^*$, having blocks $\boldsymbol{\Omega}_{[11]}^* = \boldsymbol{\Gamma}$, $\boldsymbol{\Omega}_{[22]}^* = \bar{\boldsymbol{\Omega}}$ and $\boldsymbol{\Omega}_{[21]}^* = \boldsymbol{\Omega}_{[12]}^{*\prime} = \boldsymbol{\Delta}$, needs to be a full-rank correlation matrix.

Suppose an arbitrary CAT item selection algorithm has already selected T items with item parameters $\boldsymbol{\xi}_{1:T} = (\mathbf{B}_{1:T}, \mathbf{D}_{1:T})$, where $\mathbf{B}_{1:T} := [\mathbf{B}_{j_1}, \dots, \mathbf{B}_{j_T}]'$ and $\mathbf{D}_{1:T} := [D_{j_1}, \dots, D_{j_T}]'$. Then it is possible to show the following result:

Theorem 4.2. Consider a K -factor CAT item selection procedure after selecting T items, with $\mathcal{N}(\mathbf{0}_K, \mathbb{I}_K)$ prior placed on the test taker's latent trait $\boldsymbol{\theta}$. If $\mathbf{Y}_{1:T} = (y_{j_1}, \dots, y_{j_T})'$ is conditionally independent binary response data from the two-parameter probit MIRT model defined in (1), then

$$(\boldsymbol{\theta} \mid \mathbf{Y}_{1:T}, \mathbf{B}_{1:T}, \mathbf{D}_{1:T}) \sim \text{SUN}_{K,T}(\boldsymbol{\mu}_{\text{post}}, \boldsymbol{\Omega}_{\text{post}}, \boldsymbol{\Delta}_{\text{post}}, \boldsymbol{\gamma}_{\text{post}}, \boldsymbol{\Gamma}_{\text{post}}),$$

with posterior parameters

$$\begin{aligned}\boldsymbol{\mu}_{\text{post}} &= \mathbf{0}_K, & \boldsymbol{\Omega}_{\text{post}} &= \mathbb{I}_K, & \boldsymbol{\Delta}_{\text{post}} &= \mathbf{C}'_1 \mathbf{C}_3^{-1}, \\ \boldsymbol{\gamma}_{\text{post}} &= \mathbf{C}_3^{-1} \mathbf{C}_2, & \boldsymbol{\Gamma}_{\text{post}} &= \mathbf{C}_3^{-1} (\mathbf{C}_1 \mathbf{C}'_1 + \mathbb{I}_T) \mathbf{C}_3^{-1},\end{aligned}$$

where $\mathbf{C}_1 = \text{diag}(2y_{j_1} - 1, \dots, 2y_{j_T} - 1) \mathbf{B}_{1:T}$ and $\mathbf{C}_2 = \text{diag}(2y_{j_1} - 1, \dots, 2y_{j_T} - 1) \mathbf{D}_{1:T}$. The matrix \mathbf{C}_3 is a T by T diagonal matrix, where the (t, t) -th entry is $(\|\mathbf{B}_{t,T}\|_2^2 + 1)^{\frac{1}{2}}$, with $\mathbf{B}_{t,T}$ representing the t -th row of $\mathbf{B}_{1:T}$.

Theorem 4.2 provides an exact finite-sample Bayesian characterization of the latent factor posterior. With a multivariate normal prior on $\boldsymbol{\theta}$ and a probit MIRT likelihood, the posterior $f(\boldsymbol{\theta} \mid \mathbf{Y}_{1:T})$ belongs to the unified skew-normal (SUN) family. This does not conflict with the empirical Bayes result of [43], which establishes asymptotic posterior normality as the number of items J grows; related discussion of Bayesian latent-trait estimation and asymptotic covariance properties in MIRT is given in [44]. The finite-sample form is particularly relevant for CAT, where only a small number of items has been administered and large-sample normal approximations may be unreliable. This representation enables exact posterior calculations, improving uncertainty quantification and item selection in short tests. For ordinal responses under a probit link, closely related SUN-based approximations can also be developed as discussed in [1]. The latent factor posterior can often be well approximated within the SUN family, although the representation is generally not exact.

According to [42], an arbitrary unified skew normal distribution $\boldsymbol{\theta} \sim \text{SUN}_{K,T}(\boldsymbol{\mu}, \boldsymbol{\Omega}, \boldsymbol{\Delta}, \boldsymbol{\gamma}, \boldsymbol{\Gamma})$ has a stochastic representation as a linear combination of a K -dimensional multivariate normal random variable \mathbf{V}_0 , and a T -dimensional truncated multivariate normal random variable $\mathbf{V}_{1,-\gamma}$ as follows:

$$\boldsymbol{\theta} \stackrel{d}{=} \boldsymbol{\mu} + \boldsymbol{\omega}(\mathbf{V}_0 + \boldsymbol{\Delta} \boldsymbol{\Gamma}^{-1} \mathbf{V}_{1,-\gamma}), \quad (13)$$

where $\mathbf{V}_0 \sim \mathcal{N}(0, \bar{\boldsymbol{\Omega}} - \boldsymbol{\Delta}\boldsymbol{\Gamma}^{-1}\boldsymbol{\Delta}') \in \mathbb{R}^K$, and $\mathbf{V}_{1,-\gamma}$ is a multivariate normal distribution $\mathcal{N}(0, \boldsymbol{\Gamma})$ truncated to $\{\mathbf{V}_1 \in \mathbb{R}^T : V_{1i} \geq -\gamma_i, \forall i\}$.

Based on Theorem 4.2, we have a closed-form expression for the posterior parameters $(\boldsymbol{\mu}, \boldsymbol{\Omega}, \boldsymbol{\Delta}, \boldsymbol{\gamma}, \boldsymbol{\Gamma})$ for $\boldsymbol{\theta}$. Recall that $\boldsymbol{\Omega} = \boldsymbol{\omega}\bar{\boldsymbol{\Omega}}\boldsymbol{\omega}$ is the standard covariance correlation matrix decomposition from definition 4.1, and hence the only unrealized stochastic terms in equation (13) are \mathbf{V}_0 and $\mathbf{V}_{1,-\gamma}$. This suggests that sampling from the latent factor posterior distribution $f(\boldsymbol{\theta}|\mathbf{Y}_{1:T})$ requires two independent steps: sampling from a K -dimensional multivariate normal distribution \mathbf{V}_0 , and a T -dimensional multivariate truncated normal distribution \mathbf{V}_1 . As a result, the direct sampling approach scales efficiently with the number of factors K , since generating samples from K -dimensional multivariate normal distributions is trivial. Moreover, in CAT settings, the test is typically terminated early, meaning T remains relatively small. When T is moderate (e.g., $T < 1,000$), sampling from the truncated multivariate normal distribution remains computationally efficient using the minimax tilting method [24]. The exact proof of Theorem 4.2 and the sampling details can be found in Appendix B.

The direct sampling approach provides substantial gains in both computational efficiency and numerical precision compared with traditional MCMC methods commonly used in the MIRT literature [20, 45]. In standard MIRT settings, the posterior distribution of the latent factors is typically regarded as intractable and non-Gaussian, requiring a Markov chain to be constructed via data-augmentation techniques [21, 22] so that its stationary distribution approximates the posterior. This procedure entails repeated simulation of augmented data and is inherently sequential, which limits opportunities for parallelization. Moreover, it demands additional tuning, burn-in, and convergence diagnostics to ensure that the chain adequately converges to the posterior distribution. In contrast, Theorem 4.2 establishes that the posterior distribution can be expressed exactly as a unified skew-normal distribution, allowing direct and parallel sampling without the need for iterative

convergence procedures.

Theorem 4.2 also plays a central role in our proposed deep Q-learning algorithm, as it fully characterizes the latent factor posterior distribution $f(\boldsymbol{\theta}|\mathbf{Y}_{1:T})$, enabling the parametrization of the state variable s_T , which serves as an input to the Q-network illustrated in Section 5.1. We also illustrate the application of Theorem 4.2 in accelerating the existing item selection rules below.

4.1 Accelerating Existing Rules

While Bayesian CAT criteria require multidimensional integration, these quantities are typically evaluated in practice via Monte Carlo approximation rather than analytic quadrature in moderate to high dimensions. In such Bayesian CAT implementations, the dominant computational cost lies in obtaining valid samples from the latent factor posterior. By providing exact posterior samples without Markov chain construction, our approach removes this primary bottleneck and renders the remaining Monte Carlo integration computationally efficient and easily parallelizable. For example, in computing the KL-EAP item selection rule (2), we directly sample from $f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)})$ instead of resorting to MCMC, and evaluate the integral via Monte Carlo integration. Since the EAP estimate $\hat{\boldsymbol{\theta}}_{t-1}$ remains fixed at time step t , evaluating the KL information term is straightforward. For the Max Pos item selection rule in equation (3), we again obtain i.i.d. samples from $f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)})$ and compute the posterior predictive probabilities in (4). Although the density ratio $\frac{f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)})}{f(\boldsymbol{\theta}|\mathbf{Y}_{1:t})}$ is difficult to evaluate, we can leverage the conditional independence assumption of the MIRT model. In particular, we can express the joint distribution of $(\boldsymbol{\theta}, y_{jt})$ as

$$f(\boldsymbol{\theta}, y_{jt}|\mathbf{Y}_{1:(t-1)}) = f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)}, y_{jt})f(y_{jt}|\mathbf{Y}_{1:(t-1)}) = f(y_{jt}|\boldsymbol{\theta})f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)}). \quad (14)$$

Using Equation (14), we rewrite the KL information term in Equation (3) as:

$$\begin{aligned} \int f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)}) \log \frac{f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)})}{f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)}, y_{j_t})} d\boldsymbol{\theta} &= \int f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)}) \log \frac{f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)})f(y_{j_t}|\mathbf{Y}_{1:(t-1)})}{f(\boldsymbol{\theta}, y_{j_t}|\mathbf{Y}_{1:(t-1)})} d\boldsymbol{\theta} \\ &= \int f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)}) \log \frac{f(y_{j_t}|\mathbf{Y}_{1:(t-1)})}{f(y_{j_t}|\boldsymbol{\theta})} d\boldsymbol{\theta}. \end{aligned}$$

Since $f(y_{j_t}|\boldsymbol{\theta})$ can be easily computed for each $\boldsymbol{\theta} \sim f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)})$, the online computation of Max Pos remains efficient.

Although the mutual information selection rule in Equation (5) has demonstrated strong empirical performance [13], its computational complexity remains a significant challenge. By applying Equation (14), we can rewrite mutual information as

$$\arg \max_{j_t \in R_t} \sum_{y_{j_t}=0}^1 f(y_{j_t}|\mathbf{Y}_{1:(t-1)}) \int_{\boldsymbol{\theta}} f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)}, y_{j_t}) \log \frac{f(y_{j_t}|\boldsymbol{\theta})}{f(y_{j_t}|\mathbf{Y}_{1:(t-1)})} d\boldsymbol{\theta}. \quad (15)$$

This formulation reveals that maximizing mutual information is structurally similar to Max Pos but requires much more computational effort. Unlike Max Pos, where sampling is performed from the current posterior $f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)})$, mutual information requires sampling from future posteriors $f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)}, y_{j_t})$. Since each candidate item $j_t \in R_t$ has two possible outcomes ($y_{j_t} = 0$ or $y_{j_t} = 1$), evaluating equation (15) requires obtaining samples from $|R_t| \times 2$ distinct posterior distributions. Even if sampling each individual posterior is computationally efficient, this approach becomes impractical for large item banks.

We hence propose a new approach to dramatically accelerate the computation of the mutual information quantity using the idea of importance sampling and resampling and bootstrap filter [46, 47]. Rather than explicitly sampling the future posterior $f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)}, y_{j_t})$ for each $j_t \in R_t$ and $y_{j_t} \in \{0, 1\}$, we can simply sample from the current posterior $f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)})$ once, and then perform proper posterior reweighting. Under equation (1), let $p(\boldsymbol{\theta})$ be the prior on the latent factors $\boldsymbol{\theta}$, $l_1(\boldsymbol{\theta})$ be the current data likelihood, and $l_2(\boldsymbol{\theta})$ denote the

future data likelihood after observing y_{j_t} . We have the future posterior density as follows:

$$\begin{aligned} f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)}, y_{j_t}) &\propto l_2(\boldsymbol{\theta})p(\boldsymbol{\theta}) \propto \frac{l_2(\boldsymbol{\theta})p(\boldsymbol{\theta})}{l_1(\boldsymbol{\theta})p(\boldsymbol{\theta})} f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)}) \\ &= \Phi(\mathbf{B}'_{j_t}\boldsymbol{\theta} + D_{j_t})^{y_{j_t}}(1 - \Phi(\mathbf{B}'_{j_t}\boldsymbol{\theta} + D_{j_t}))^{1-y_{j_t}} f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)}). \end{aligned} \quad (16)$$

Equation (16) suggests we can generate samples from $f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)}, y_{j_t})$ via reweighting and resampling from the current posterior samples $f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)})$. Specifically, given a sufficiently large set of posterior samples $(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M) \sim f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)})$, we assign distinct weights for each sample $m \in \{1, \dots, M\}$ and future item $j_t \in R_t$:

$$q_m = \frac{w_m}{\sum_{i=1}^M w_i}, \quad \text{where } w_i = \Phi(\mathbf{B}'_{j_t}\boldsymbol{\theta}_i + D_{j_t})^{y_{j_t}}(1 - \Phi(\mathbf{B}'_{j_t}\boldsymbol{\theta}_i + D_{j_t}))^{1-y_{j_t}}.$$

To sample from $f(\boldsymbol{\theta}|\mathbf{Y}_{1:(t-1)}, y_{j_t})$, we then draw from the discrete distribution over $(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M)$, placing weight q_m on $\boldsymbol{\theta}_m$. This approach eliminates the need to sample from $|R_t| \times 2$ distinct posteriors directly, further accelerating mutual information computation and making it scalable for large item banks.

5 Learning Optimal Item Selection Policy

Building on the Bayesian MIRT foundation established in the previous section, we now formulate the problem of CAT as an RL task. The Bayesian framework provides two key advantages that make this integration both principled and computationally feasible. First, the identified latent factor posterior distributions offer a well-defined representation of examinee knowledge, allowing their corresponding posterior parameters to be parametrized directly as state variables. Without such identification, encoding an unknown and analytically intractable posterior would be ambiguous. Second, because reinforcement learning typically requires extensive simulations to learn an optimal policy, the acceleration achieved

in online item selection enables rapid simulation of testing sessions, thereby substantially improving the efficiency of policy training.

As illustrated in Section 3.3, an RL approach addresses the myopic nature of traditional CAT selection rules, enables a more flexible reward structure, and directly minimizes the number of items required for performing online adaptive testing. Specifically, we propose a novel double Q-learning algorithm [18, 19] for online item selection in CAT. The algorithm trains a deep neural network offline using only the item parameters estimated from an arbitrary two-parameter MIRT model. The offline training phase is completed before any live CAT administration. After training, the network weights are fixed, and online CAT only requires sequentially updating the posterior state and applying the learned network to select the next item. During online item selection, the neural network takes the current posterior distribution $s_t := f(\boldsymbol{\theta} \mid \mathbf{Y}_{1:t})$ as input and outputs the next item selection for step $(t + 1)$. The neural network architecture is described in Section 5.1, while the double deep Q-learning algorithm is detailed in Section 5.2.

5.1 Deep Q-learning Neural Network Design

A potential reason that a deep reinforcement learning approach has not been proposed in the CAT literature is the ambiguity arising from unidentified latent factor distributions. However, by Theorem 4.2, it is straightforward to compactly parametrize the posterior distribution at each time step t using the parameters $\mathbf{C}_1 \in \mathbb{R}^{t \times K}$, $\mathbf{C}_2 \in \mathbb{R}^t$, and $\text{diag}(\mathbf{C}_3) \in \mathbb{R}^t$, where $\text{diag}(\mathbf{C}_3)$ represents the diagonal vector of \mathbf{C}_3 . This structured representation of the posterior enables item selection policy learning via deep neural networks, which takes posterior parameters as input and outputs item selection.

Neural networks can be regarded as flexible function approximators that learn nonlinear mappings between inputs and outputs through a composition of simple transformations [48]. In our framework, the deep neural network consists of two key components: an encoder and

a classifier. At each time step t , the encoder maps the collection of posterior parameters to a latent representation in \mathbb{R}^L , where L is a hyperparameter and represents the dimension of the latent feature space. The classifier then outputs a J -dimensional vector of estimated Q-values and selects the next item from the item bank that is expected to yield the highest Q-value.

Define the collection of the posterior parameters $\tilde{\xi}_t := \{(\mathbf{C}_{1h}, C_{2h}, C_{3h})\}_{h=1}^t$, where $\mathbf{C}_{1h} \in \mathbb{R}^K$ is the h -th row of \mathbf{C}_1 , $C_{2h} \in \mathbb{R}$ is the h -th element of \mathbf{C}_2 , and $C_{3h} \in \mathbb{R}$ is the h -th elements in the diagonal of \mathbf{C}_3 . Since the size of the posterior parameters $\tilde{\xi}_t$ grows over time, and permuting the tuples within $\tilde{\xi}_t$ still describes the same posterior, we have to design a neural network that can take inputs of growing size, and can provide output that is permutation invariant of the inputs. One solution is to consider the weight sharing idea from the Bayesian experimental design literature [33]. Let $\phi_1(\cdot) : \mathbb{R}^{K+2} \rightarrow \mathbb{R}^{L_1}$ denote an encoder component that maps each tuple $(\mathbf{C}_{1h}, C_{2h}, C_{3h}) \in \mathbb{R}^{K+2}$ to an L_1 -dimensional latent space, and consider the operation $g_1(\cdot)$ as follows:

$$g_1(\tilde{\xi}_t) := \sum_{h=1}^t \phi_1\{(\mathbf{C}_{1h}, C_{2h}, C_{3h})\}. \quad (17)$$

Observe that $g_1(\cdot)$ is capable of handling a growing number of inputs through summations of $\phi_1(\cdot)$ functions over t . More importantly, permuting the order of the tuples in $\tilde{\xi}_t$ does not change the value of $g_1(\cdot)$, since summation is permutation invariant. Although the form of g_1 may look restrictive, any function $f(\cdot)$ operating on a countable set can be decomposed into the form $\rho \circ g_1(\cdot)$, where $\rho(\cdot)$ is a suitable transformation that can be learned from another neural network (see Theorem 2 of [49]).

In principle, the network design $\rho \circ g_1(\cdot)$, coupled with the Q-learning algorithm, is sufficient to learn the optimal policy. However, to enhance learning efficiency, we further enrich the state representation $\tilde{\xi}_t$ with a matrix of prediction quartiles $\Psi_t \in \mathbb{R}^{J \times Q}$, where

each row contains a vector of quantiles of the predictive distribution for item j . We form these quantiles by first drawing samples $\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M \sim f(\boldsymbol{\theta}|\mathbf{Y}_{1:t})$ as in Section 4, and then computing the quantiles of the prediction samples $\{\Phi(\mathbf{B}'_j \boldsymbol{\theta}_i + D_j)\}_{i=1}^M$ for each item j . Since sampling from $f(\boldsymbol{\theta}|\mathbf{Y}_{1:t})$ only needs to be done once, computing the matrix $\boldsymbol{\Psi}_t$ online is computationally efficient.

This practice of augmenting the raw state variable with additional contextual features of the states ($\boldsymbol{\Psi}_t$) echoes the common strategy in the RL literature [18, 50]: learning tends to be more stable and efficient when the state representation captures not only the current ability estimate but also characteristics of the item bank. In our setting, $\boldsymbol{\Psi}_t$ offers a richer description of the predictive distributions across items given the current estimates, effectively serving as additional covariates for item selection. Empirically, we find that incorporating $\boldsymbol{\Psi}_t$ substantially accelerates convergence toward the optimal item selection policy. Including $\boldsymbol{\Psi}_t$ in the state representation is harmless, as it is a deterministic function of the posterior distribution; identical posteriors will always produce identical $\boldsymbol{\Psi}_t$, ensuring that the policy remains consistent across equivalent states.

In summary, Figure 2 describes the architecture of the policy network used in our approach. The same network is trained offline and then deployed during live CAT with fixed weights, so only the posterior state is updated sequentially during online administration. To select the $(t + 1)$ -th item, our proposed neural net takes two inputs: the posterior parameters $\tilde{\boldsymbol{\xi}}_t$ and the prediction matrix $\boldsymbol{\Psi}_t$, and outputs the selected item. Let $\phi_2 : \mathbb{R}^{J \times Q} \rightarrow \mathbb{R}^{L_2}$ represent the encoder component that maps the matrix $\boldsymbol{\Psi}_t \in \mathbb{R}^{J \times Q}$ to L_2 -dimensional space. Write $L = L_1 + L_2$ and the concatenated outputs of $g_1(\cdot)$ and $\phi_2(\cdot)$ as $[g_1(\tilde{\boldsymbol{\xi}}_t)', \phi_2(\boldsymbol{\Psi}_t)']' \in \mathbb{R}^L$, we then define the classification component of neural network as $\rho(\cdot) : \mathbb{R}^L \rightarrow \mathbb{R}^J$, which maps the concatenated outputs of $g_1(\cdot)$ and $\phi_2(\cdot)$ to a J -dimensional logit vector. Denote the final policy network as π_ϕ and recall that $\rho(\cdot)$ represents the classification layer, our proposed network selects the item at time t corresponding to the

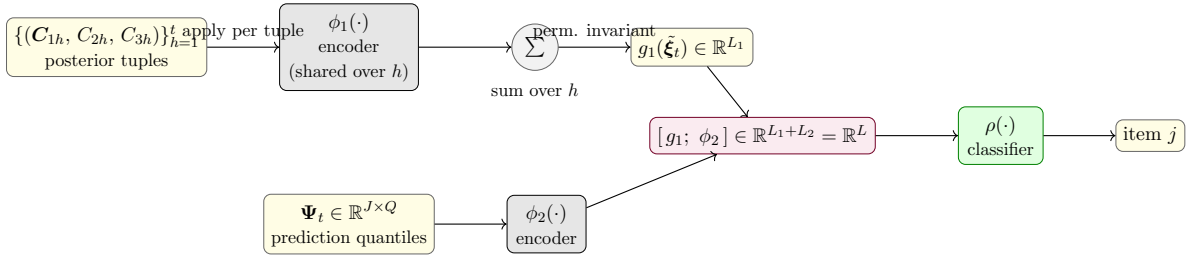


Figure 2: High-level architecture of the Q-network. The shared encoder ϕ_1 maps each tuple of posterior parameters to \mathbb{R}^{L_1} and the sum yields the permutation invariant representation $g_1(\tilde{\xi}_t)$. The matrix Ψ_t is encoded by ϕ_2 . The concatenated vector in \mathbb{R}^L is passed to the classifier ρ to select the j -th item (largest value in the J logits). This network is trained offline using Algorithm 1; during live CAT, its weights are fixed and only the posterior state is updated sequentially.

maximum value of the following function:

$$\pi_\phi(\tilde{\xi}_t, \Psi_t) := \rho \circ \left\{ [g_1(\tilde{\xi}_t)', \phi_2(\Psi_t)']' \right\}.$$

We implement both primary and target Q-networks as simple feed-forward multilayer perceptron and find that their performance is largely insensitive to the exact choice of L_1 and L_2 , and does not require very large depth. All architectural details such as layer sizes, activations, and optimizer settings are provided in Appendix D. This robustness suggests that our empirical improvements stem from the CAT-specific state design and augmentation rather than network complexity.

5.2 Double Deep Q-learning for CAT

In classical tabular Q-learning, the action-value function is stored as a Q-matrix whose rows correspond to states and whose columns correspond to actions. Such a representation is infeasible in CAT because the state space is combinatorially large and the state variable s_t is represented by continuous posterior summaries rather than a small finite index. We therefore replace the tabular Q-matrix by a neural-network approximator $Q_w(s, a)$. For a given state s_t , the network in Figure 2 outputs the full vector of estimated Q-values before

selecting the optimal item, which plays the same role as one row of a tabular Q-matrix, but is produced through function approximation rather than table lookup.

The standard Q-learning algorithm often overestimates Q-values because the same action-value approximation is used both to select the maximizing action and to evaluate it when constructing the temporal-difference target. To address this issue, we adopt the Double Q-learning approach [19], which replaces the single action-value approximation by two Q-networks with distinct roles. The primary network Q_w is updated by gradient descent and is used to select the action with the largest estimated Q-value, while the target network $Q_{\bar{w}}$ is a delayed copy used only to evaluate that selected action when forming the target. This is the neural-network analogue of maintaining two Q-matrices in Double Q-learning. The separation reduces overestimation bias and leads to more stable training. Specifically, compared with standard Q-learning, Double Q-learning minimizes the same TD loss in (11) but replaces the target in (12) with

$$y(s, a, r, s') = r + \gamma Q_{\bar{w}}(s', \arg \max_{a'} Q_w(s', a')).$$

Our proposed algorithm is presented in Algorithm 1. The algorithm trains, entirely offline, the policy network illustrated in Figure 2. Once training is complete, the learned network is deployed during live CAT with fixed weights, and only the posterior state is updated sequentially as new responses are observed. Importantly, the offline training phase in Algorithm 1 relies only on the item parameters (\mathbf{B}, \mathbf{D}) and simulated examinees drawn from a specified latent factor distribution (e.g., $\boldsymbol{\theta}_i \sim \mathcal{N}(0, \mathbb{I}_K)$), rather than on observed item response data from the item bank. If prior information about the distribution of future online examinees is available, the simulation distribution can be adapted accordingly. In this manuscript, we consistently use a standard multivariate normal distribution to ensure comparability across methods and experiments.

Note that for the Q-learning to converge to an optimal policy, it is essential to adopt an ϵ -greedy policy, where the algorithm makes random item selections with probability ϵ and gradually decreases ϵ over the course of training. For simplicity, the state variable s_t in Algorithm 1 is a compact representation of $(\tilde{\xi}_t, \Psi_t)$ defined in Section 5.1. In practice, we can terminate the training when both the rewards and the validation loss stabilize, indicating that the neural network has well approximated the optimal policy.

6 Simulation

6.1 Simulation Design

To evaluate our approach in a challenging and realistic simulation setting, we randomly generated a 5-factor, 150-item factor loading matrix $\mathbf{B} \in \mathbb{R}^{150 \times 5}$. Each column of \mathbf{B} was initialized by randomly permuting 150 equally spaced values, with the magnitude constrained to lie within $[0.3, 3]$. We then imposed a lower triangular structure to ensure identifiability. To better reflect practical datasets, where items rarely load on all five factors, each item was required to load on the first factor and on at most two additional factors beyond the first. A visualization of the loading matrix can be found in Appendix F. Item intercepts were independently drawn from $\text{Unif}(-1.5, 1.5)$.

The goal of the simulation is to accurately estimate the first three factors ($\mathcal{K} = \{1, 2, 3\}$), while accounting for the presence of factors 4 and 5. Leveraging our proposed direct sampling approach outlined in Section 4, we implemented our double Q-learning algorithm alongside all existing methods discussed in Section 3.2. These include the EAP approach (Equation 2), the Max Pos approach (Equation 3), the MI approach (Equation 5), and the Max Var approach (Equation 6). We further modified all baseline information-based criteria so that they also target only the prioritized subset of factors $\mathcal{K} = \{1, 2, 3\}$. Specifically, the EAP, Max Pos, and MI rules are now written as integrals over $\theta_{\mathcal{K}}$ rather than all 5

Algorithm 1: Double Q-learning Algorithm for Bayesian CAT (offline)

Input: Item Parameters (\mathbf{B}, \mathbf{D}) , total Episodes E , $\epsilon \in [\underline{\epsilon}, \bar{\epsilon}]$ with $\epsilon^{(0)} = \bar{\epsilon}$, buffer \mathcal{H} , target network update frequency T , batch size M .

Output: Learned Q-network Q_w for item selection.

Initialize primary Q-network Q_w and target Q-network $Q_{\bar{w}}$, with weights $w_0 = \bar{w}_0$;

for $i = 1$ **to** E **do**

- Draw $\boldsymbol{\theta}_i$ from $\mathcal{N}(0, \mathbb{I}_K)$;
- Set initial state $s_0 := [\mathcal{N}(0, \mathbb{I}_K), \boldsymbol{\Psi}_0]$ and available item set $R_0 = \{1, \dots, J\}$;
- for** $t = 0$ **to** H **do**
 - Select action a_t using ϵ -greedy policy:
$$a_t = \begin{cases} \text{Random Selection} & \text{with probability } \epsilon^{(t)} \\ \arg \max_{a \in R_t} Q_w(s_t, a) & \text{otherwise.} \end{cases}$$
 - Simulate response $y_t \sim \text{Bernoulli}(\Phi(\mathbf{B}'_{a_t} \boldsymbol{\theta}_i + D_{a_t}))$;
 - Update available item set $R_{t+1} = R_t \setminus a_t$;
 - Obtain new state s_{t+1} from $(\boldsymbol{\xi}_{t+1}, \boldsymbol{\Psi}_{t+1})$;
 - Compute reward r_t ;
 - if** $r_t = 0$ **then**
 - | Break and proceed to the next episode;
 - end**
 - Store transition $H_{i,t} := (s_t, a_t, r_t, R_t, s_{t+1})$ into buffer \mathcal{H} ;
 - if** $|\mathcal{H}| > \text{batch size } M$ **then**
 - | Randomly sample M transitions $\{H_{i,t}\}_{j=1}^M$ from \mathcal{H} ;
 - | Compute $a_j = \arg \max_{a \in R_j} Q_w(s'_j, a)$;
 - | Compute target:
$$y^{(j)} = \begin{cases} -1 & r_j = 0 \\ -1 + \gamma Q_{\bar{w}}(s'_j, a_j) & \text{otherwise.} \end{cases}$$
 - | Perform batch gradient descent on loss: $\sum_{j=1}^M (y^{(j)} - Q_w(s_j, a_j))^2$;
 - end**
 - Decrease $\epsilon^{(t)}$ if $\epsilon^{(t)} \geq \underline{\epsilon}$;
- end**
- Update the target network by setting $\bar{w} = w$ if i is a multiple of T ;

end

factors. These adjustments ensure that all competing CAT algorithms are tuned to the same estimation target, making the simulation comparisons fair and directly comparable.

For performance evaluation, we generated $N = 500$ online examinees with latent traits $\boldsymbol{\theta}_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \mathbb{I}_5)$. For each examinee, we administered a 50-item test using each of the

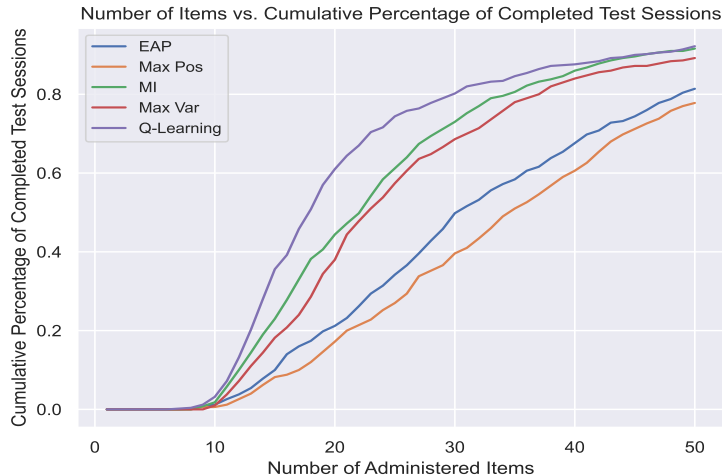


Figure 3: Number of Items Versus Cumulative Percentage of Completed Tests

CAT algorithms. We then compared performance in terms of posterior variance reduction, mean-squared error (MSE), termination efficiency, and item exposure rates based on these 500 adaptive testing sessions.

For the Q-learning algorithm, we trained the Q-network for 80,000 episodes, with exploration parameter ϵ decreasing linearly from 0.99 to 0.01 over 700,000 steps. We set a sufficiently large $H = 60$ items, with discount factor $\gamma = 0.95$. To verify that the double Q-learning algorithm indeed converges to the optimal policy π^* , we provide further details on the training dynamics of our deep Q-learning algorithm, illustrating the increase in rewards until convergence in Appendix F.2.

6.2 Simulation Results

To evaluate termination efficiency, we start with a standard multivariate Gaussian prior with unit marginal variances and terminate the test once the maximum posterior variance across *all three* target factors in \mathcal{K} falls from 1 to below $\tau^2 < 0.16$. In practice, practitioners may select the threshold τ^2 based on the specific requirements of their application. Based on 500 simulated adaptive testing sessions, Figure 3 illustrates how the percentage of completed test sessions increases as more items are administered. A faster growth rate of the

Table 1: **Comparison of Win shares (W.S), Termination, and Computation**

Algorithm	Avg Termination (items)	W.S dim1	W.S dim2	W.S dim3	Avg Time (s/item)
EAP	31.7	17.0%	16.4%	13.6%	0.037
Max Pos	34.0	16.8%	16.2%	18.6%	0.055
MI	23.2	19.0%	21.4%	20.2%	0.082
Max Var	25.7	23.6%	22.6%	23.6%	0.038
Q-learning	21.5	23.6%	23.4%	24.0%	0.064

completion percentages indicates faster posterior variance reduction. Notably, the purple Q-learning curve consistently rises more quickly than the others, demonstrating significant testing efficiency gains. As summarized in the second column of Table 1, Q-learning requires an average of only 21.5 items to reach the termination criterion for all 3 targeted factors, outperforming all other methods.

To assess early-stage performance, we compute the MSEs after 20 items for each test session and summarize results using win shares, defined as the percentage of test takers for whom each method achieves the lowest MSE after 20 items. These win shares are reported across all three latent dimensions in Table 1. Notably, Q-learning achieves the best win shares across all three factors. This advantage at early stages is consistent with the training objective of Q-learning, which explicitly rewards rapid posterior variance reduction and early test termination, with the average stopping time occurring at around 20 items.

Table 2 reports the evolution of MSEs as a function of test length up to 50 administered items. Consistent with the win share analysis, Q-learning exhibits competitive performance during the early stages of testing. In particular, at $T = 20$ and $T = 30$, Q-learning attains the smallest MSEs for Factors 2 and 3. As the test length increases, performance differences across methods diminish, and Q-learning achieves MSEs comparable to those of MI and Max Var. This pattern is expected, as the Q-learning policy is trained primarily on short-horizon testing trajectories: since most simulated tests terminate before 30 items, the algorithm is exposed less frequently to long-horizon simulations during training. Additional visualizations of the MSE trajectories are provided in Appendix F.3.

Table 2: MSEs between Posterior Mean and Ground Truth for the first three latent factors as a function of test length.

Factor	Algorithm	Number of items				
		10	20	30	40	50
1	EAP	0.175	0.123	0.090	0.073	0.062
	Max Pos	0.194	0.128	0.104	0.083	0.062
	MI	0.180	0.100	0.066	0.055	0.050
	Max Var	0.186	0.080	0.073	0.059	0.056
	Q-learning	0.221	0.096	0.071	0.060	0.054
2	EAP	0.362	0.214	0.164	0.125	0.105
	Max Pos	0.354	0.221	0.166	0.129	0.103
	MI	0.223	0.150	0.120	0.102	0.086
	Max Var	0.214	0.141	0.118	0.092	0.087
	Q-learning	0.282	0.138	0.111	0.096	0.087
3	EAP	0.317	0.226	0.152	0.112	0.085
	Max Pos	0.324	0.207	0.139	0.115	0.092
	MI	0.213	0.127	0.103	0.096	0.089
	Max Var	0.234	0.136	0.110	0.099	0.096
	Q-learning	0.233	0.114	0.098	0.091	0.085

Figure 4 summarizes the distribution of item exposure rates across the 150 item bank for each CAT algorithm, where the exposure rate of an item is defined as the proportion of examinees whose test includes that item at least once. With test length $H = 50$ and bank size $J = 150$, the expected average exposure is $H/J \approx 0.33$, and all five algorithms yield mean and median exposure rates close to this benchmark. Differences across methods primarily arise in the tails of the exposure distribution. In particular, the EAP and Q-learning policies exhibit slightly higher upper-quartile and maximum exposure rates, indicating a more selective use of highly informative items. Importantly, no method concentrates exposure excessively on a small subset of items, suggesting that all algorithms maintain reasonable item utilization in this simulation setting. For the Q-learning approach, incorporating explicit exposure control mechanisms into the learning objective, such as penalizing repeated use of highly exposed items, is a natural extension but is beyond the scope of the present




Figure 4: Distributions of Item Exposure Rates

study.

The last column of Table 1 reports the average online item selection time, highlighting the significant computational advantages of our direct sampling approach. Even the most computationally intensive MI approach requires only 0.082 seconds per item selection, via our proposed posterior reweighting strategy. Q-learning adds only a single feed-forward pass per selection, keeping online latency under a few milliseconds on standard hardware without GPU acceleration. While training the Q-network offline using Algorithm 1 took approximately 30 hours for this exercise, this is a one-time investment that can be accelerated if GPUs are available; thereafter, the virtually instantaneous online policy makes the approach highly practical for real-time CAT deployment.

7 Cognitive Function Measurements

7.1 pCAT-COG Data and Experiment Design

We revisit the problem of designing a deep CAT system for the pCAT-COG study, as outlined in Section 2. Since item response data for all $N = 730$ examinees across $J = 57$ items are available, we can directly use real item responses during evaluation rather than

simulating testing sessions. In this experiment, all $N = 730$ examinees are treated as online examinees. All CAT algorithms, including Q-learning, have access only to the estimated item parameters from the pCAT-COG study and do not observe the examinees’ binary item responses beyond those revealed sequentially during adaptive administration.

Given that pCAT-COG is designed to measure global cognitive ability (first column in Figure 1) while accounting for five cognitive subdomains, we specified the Q-learning reward function to reduce posterior variance for the primary factor ($\mathcal{K} = \{1\}$). This demonstrates the flexibility of our CAT system, as it can be tailored to the specific cognitive assessment needs. As in Section 6, we also modified all baseline CAT algorithms so that they also target only the primary factor by integrating only over the primary dimension for fair comparison.

Finally, this real-data experiment also serves as a robustness check for the proposed Q-learning approach. As described in Algorithm 1, the policy is trained using simulated examinees with latent factors drawn from a standard multivariate normal distribution. In contrast, the empirical factor correlations in pCAT-COG need not follow this distributional assumption. Evaluating Q-learning on real response data therefore provides evidence that the learned policy remains effective when the latent factor structure deviates from the training distribution used in simulation.

7.2 Results

The left subplot of Figure 5 shows that Q-learning again achieves the fastest test termination compared to other methods. As before, the test is dynamically terminated when the posterior variance drops from 1 to below $\tau^2 = 0.16$, and the cumulative percentage of completed test sessions is computed over all 730 examinees. The second column of Table 3 shows that Q-learning reaches the desired posterior variance reduction threshold after only an average of 11.2 items.

For further comparison, we consider adaptively selecting 20 items for each test taker without dynamic termination. We chose the number 20 because the Max Pos approach required the largest average of 19.8 items for termination. The third column of Table 3 shows that the Q-learning approach achieves the highest 0.959 correlation between the estimated primary factor posterior means and the ground truth after only 20 items. Additionally, the Q-learning approach attains the highest win shares in estimating the primary dimension across all $N = 730$ examinees after 20 items. As defined in Section 6, win shares denote the proportion of examinees for whom a CAT method attains the smallest mean-squared error. Since all 57 items are ultimately administered in the real-data study, exposure rates are degenerate under full-length testing. We therefore report item exposure patterns based on the first 30 administered items in Appendix G.1.

Additionally, the right subplot of Figure 5 highlights the rapid decay of mean squared error (MSE) in estimating the primary dimension with Q-learning. As summarized in Table 4, Q-learning achieves the smallest MSE at the early stages when $T = 10$ and $T = 20$. At longer horizons when $T = 30$ and $T = 40$, the gains in MSE for the Q-learning approach diminish. This behavior is consistent with the simulation findings in Section 6 and reflects the reward specification used for Q-learning, which emphasizes rapid early-stage variance reduction. Because most simulated training trajectories generated in Algorithm 1 terminate well before 20 items, such performance behavior for Q-learning is expected.

This experiment highlights the effectiveness of our Q-learning approach in high-dimensional cognitive function measurement. Unlike approaches that learn a fixed test design offline [51], our framework enables fully dynamic item selection, maximizing the utilization of the item bank by exploring diverse testing trajectories. This termination efficiency is particularly valuable in cognitive assessment, where item development is costly, and can help mitigate practice effects and preserve items for future use.

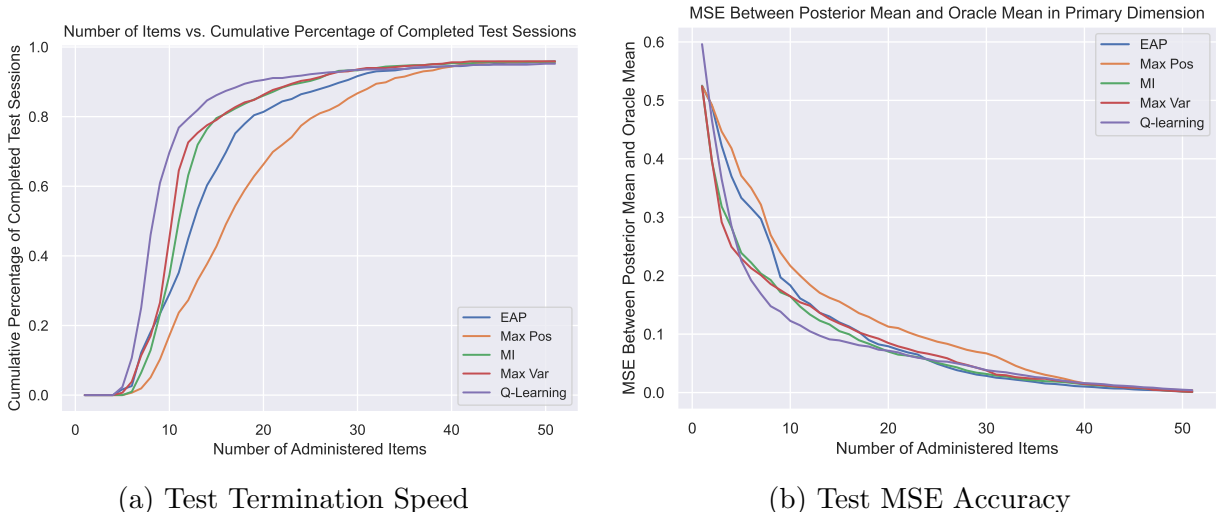


Figure 5: pCAT-COG: Primary Factor Posterior Variance Reduction (Left) and Estimation Accuracy (Right)

Table 3: Comparison of termination efficiency, primary-factor accuracy, and computation for pCAT-COG.

Algorithm	Avg Termination (items)	Correlation	Win shares	Avg Time (s/item)
EAP	15.1	0.956	17.7%	0.020
Max Pos	18.5	0.933	16.7%	0.030
MI	12.6	0.959	21.8%	0.038
Max Var	13.0	0.951	20.9%	0.021
Q-learning	11.2	0.959	22.9%	0.025

Note. Primary factor correlation and win shares (W.S) are computed after 20 administered items. Win shares denote the proportion of examinees for whom a method attains the smallest mean-squared error for the primary factor at this stage.

8 Discussion

This work advocates for a deep reinforcement learning perspective in the design of multidimensional computerized adaptive testing systems. We make two key contributions to the existing CAT literature: (1) a computational advancement that accelerates existing online item selection rules within a flexible Bayesian MIRT framework, and (2) a novel reinforcement learning-based approach that mitigates myopic decision-making and prioritizes the assessment of primary factors of interest.

Theorem 4.2 not only provides an efficient parameterization of the latent factor poste-

Table 4: Mean-squared errors (MSEs) as a function of test length

Algorithm	10 items	20 items	30 items	40 items
EAP	0.183	0.079	0.029	0.010
Max Pos	0.217	0.113	0.067	0.016
MI	0.165	0.070	0.032	0.014
Max Var	0.165	0.085	0.038	0.014
Q-learning	0.123	0.072	0.038	0.016

rior distributions for our proposed deep Q-learning approach, but also improves existing CAT item selection algorithms, as detailed in Section 4. Leveraging direct sampling from unified skew-normal distributions, our methodology scales efficiently with a large number of factors and items, achieving near-instantaneous online selection by circumventing MCMC sampling and data augmentation. Additionally, our approach naturally extends to fully Bayesian item selection by accounting for uncertainties in item parameters, which is an essential consideration when the item bank is not well-calibrated as detailed in Appendix C. Throughout the testing trajectory, our approach precisely characterizes the evolution of posterior distributions at each time step, providing a more robust measurement process beyond point estimates with difficult-to-compute standard errors.

Another key contribution of this work is the development of a robust deep double Q-learning algorithm with a customized reward structure that directly minimizes test length. As demonstrated in both simulations and real-data studies, our Q-learning algorithm consistently achieves the fastest posterior variance reduction while rapidly decreasing estimation bias. Moreover, its flexible reward function allows adaptation to different testing objectives, providing a principled framework for designing customized tests and overcoming the myopic nature of traditional CAT methods.

Our work also offers practical guidance for selecting the appropriate item selection algorithms. Unlike the existing heuristic rules, one limitation of our deep CAT system is the requirement of offline training (Algorithm 1) before online deployment. For the

pCAT-COG study, offline training took approximately 12 hours on a single GPU as a one-time investment, but no GPU is needed for subsequent online item selection. Even when offline training is undesirable, our framework significantly accelerates existing methods. Experiments show that MI (Equation 5) and our Max Var (Equation 6) approaches often outperform other online item selection rules, echoing the findings presented in [13].

A promising direction for future work is to address several limitations of the current Q-learning framework for CAT. First, the learned policy is tied to a specific item bank and must be retrained when the items are substantially modified or replenished, which may limit its immediate applicability in settings with frequent item updates. Second, as a model-free approach, Q-learning can be computationally expensive to train, particularly for long-horizon testing scenarios in which optimal policies depend on extended future trajectories. In addition, although the proposed approach does not rely on tabular exploration of the full state space and instead uses function approximation based on a compact posterior representation, scalability remains an important practical consideration as the state space grows exponentially. Our simulation results suggest that the method works well for moderate 150-item banks, but additional methodological development may be useful for extending the approach to substantially larger banks. Finally, while the proposed method yields reasonable item exposure patterns in our experiments, additional work is needed to incorporate explicit exposure control mechanisms directly into the learning objective.

Another exciting avenue is to explore alternative reward structures for our Q-learning algorithm. While the 0-1 reward structure is interpretable and stabilizes Q-network training, it provides sparse feedback, which may limit empirical performance. Alternative designs could incorporate intermediate rewards or increased penalization as more items are administered. An interesting theoretical direction is to characterize when reinforcement learning approaches are particularly advantageous compared to myopic methods. This requires careful assumptions about the reward function and the item bank properties, providing deeper

insights into the trade-offs between RL and traditional item selection strategies.

References

- [1] Jiguang Li, Robert Gibbons, and Veronika Rockova. Sparse bayesian multidimensional item response theory. *Journal of the American Statistical Association*, 2025.
- [2] Wim J. van der Linden and Cees A. W. Glas. *Elements of Adaptive Testing*. Springer, New York, 2010.
- [3] R.D. Bock and R.D. Gibbons. *Item Response Theory*. Wiley, 2021.
- [4] Robert D. Gibbons, David J. Weiss, Ellen Frank, and David Kupfer. Computerized adaptive diagnosis and testing of mental health disorders. *Annual Review of Clinical Psychology*, 12:83–104, 2016. Epub 2015 Nov 20.
- [5] Daniel O. Segall. Multidimensional adaptive testing. *Psychometrika*, 61(2):331–354, 1996.
- [6] Daniel O. Segall. Principles of multidimensional adaptive testing. In Wim J. van der Linden and Cees A. W. Glas, editors, *Computerized Adaptive Testing: Theory and Practice*, pages 53–73. Kluwer Academic, Boston, 2000.
- [7] Wim J. van der Linden. Multidimensional adaptive testing with a minimum error-variance criterion. *Journal of Educational and Behavioral Statistics*, 24(4):398–412, 1999.
- [8] Joram Mulder and Wim J. van der Linden. Multidimensional adaptive testing with optimal design criteria for item selection. *Psychometrika*, 74(2):273–296, Jun 2009. Epub 2008 Dec 23.
- [9] Hua-Hua Chang and Zhiliang Ying. A global information approach to computerized adaptive testing. *Applied Psychological Measurement*, 20:213 – 229, 1996.
- [10] Bernard Veldkamp and Wim van der Linden. Multidimensional adaptive testing with constraints on test content. *Psychometrika*, 67(4):575–588, December 2002.
- [11] Joris Mulder and Wim van der Linden. Multidimensional adaptive testing with kullback–leibler information item selection. In *Elements of Adaptive Testing*, pages 77–101. 2010.
- [12] Alexander Weissman. Mutual information item selection in adaptive classification testing. *Educational and Psychological Measurement*, 67(1):41–58, 2007.
- [13] Chun Wang and Hua-hua Chang. Item selection in multidimensional computerized adaptive testing—gaining information from different angles. *Psychometrika*, 76(3):363–384, 07 2011. Copyright - The Psychometric Society 2011; Last updated - 2023-12-03.

- [14] Chun Wang, Hua-Hua Chang, and Keith A. Boughton. Kullback–leibler information and its applications in multi-dimensional adaptive testing. *Psychometrika*, 76(1):13–39, 2011.
- [15] Hua-Hua Chang. Psychometrics behind computerized adaptive testing. *Psychometrika*, 80(1):1–20, March 2015.
- [16] Hua-Hua Chang and Zhiliang Ying. a-stratified multistage computerized adaptive testing. *Applied Psychological Measurement*, 23(3):211–222, 1999.
- [17] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, second edition, 2018.
- [18] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dhharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518:529–533, 2015.
- [19] Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 30(1), Mar. 2016.
- [20] A. A. Béguin and C. A. W. Glas. Mcmc estimation and some model-fit analysis of multidimensional irt models. *Psychometrika*, 66(4):541–561, 2001.
- [21] James H. Albert and Siddhartha Chib. Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*, 88(422):669–679, 1993.
- [22] Nicholas G. Polson, James G. Scott, and Jesse Windle. Bayesian inference for logistic models using pólya–gamma latent variables. *Journal of the American Statistical Association*, 108(504):1339–1349, 2013.
- [23] Daniele Durante. Conjugate bayes for probit regression via unified skew-normal distributions. *Biometrika*, 106(4):765–779, aug 2019.
- [24] Z. I. Botev. The normal law under linear restrictions: Simulation and estimation via minimax tilting. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 79(1):125–148, feb 2016.
- [25] Dimitri P. Bertsekas and John N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, 1st edition, 1996.
- [26] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.

- [27] Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, and Sergey Levine. Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation. *ArXiv*, abs/1806.10293, 2018.
- [28] Xiao Li, Hanchen Xu, Jinming Zhang, and Hua hua Chang. Deep reinforcement learning for adaptive learning systems. *Journal of Educational and Behavioral Statistics*, 48(2):220–243, 2023.
- [29] Chunxi Tan, Ruijian Han, Rougang Ye, and Kani Chen. Adaptive learning recommendation strategy based on deep q-learning. *Applied Psychological Measurement*, 44(4):251–266, 2020.
- [30] Tom Rainforth, Adam Foster, Desi R. Ivanova, and Freddie Bickford Smith. Modern Bayesian Experimental Design. *Statistical Science*, 39(1):100 – 114, 2024.
- [31] Kathryn Chaloner and Isabella Verdinelli. Bayesian Experimental Design: A Review. *Statistical Science*, 10(3):273 – 304, 1995.
- [32] Paola Sebastiani and Henry P. Wynn. Maximum entropy sampling and optimal bayesian experimental design. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 62(1):145–157, 2000.
- [33] Adam Foster, Desi R. Ivanova, Ilyas Malik, and Tom Rainforth. Deep adaptive design: Amortizing sequential bayesian experimental design. In *International Conference on Machine Learning*, 2021.
- [34] Niranjjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias W Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the 27th International Conference on Machine Learning (ICML)*, pages 1015–1022. Omnipress, 2010.
- [35] Remi Lam, Karen Willcox, and David H. Wolpert. Bayesian optimization with a finite budget: An approximate dynamic programming approach. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- [36] Robert Gibbons, Diane Lauderdale, Robert Wilson, David Bennett, Tesnim Arar, and David Gallo. Adaptive measurement of cognitive function based on multidimensional item response theory. *Alzheimer’s & Dementia: Translational Research & Clinical Interventions*, 10, 11 2024.
- [37] Alzheimer’s Association. 2024 alzheimer’s disease facts and figures. *Alzheimer’s & Dementia*, 20(5), 2024.
- [38] Alison R. Huang, Kiersten L. Strombotne, Elizabeth Mokyr Horner, and Susan J. Lapham. Adolescent cognitive aptitudes and later-in-life alzheimer disease and related disorders. *JAMA Network Open*, 1(5):e181726–e181726, 09 2018.

- [39] Robert D. Gibbons and Donald Hedeker. Full-information item bi-factor analysis. *Psychometrika*, 57(3):423–436, September 1992.
- [40] Alfréd Rényi. On measures of entropy and information. 1961.
- [41] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin A. Riedmiller. Playing atari with deep reinforcement learning. *ArXiv*, abs/1312.5602, 2013.
- [42] Reinaldo B. Arellano-Valle and Adelchi Azzalini. On the unification of families of skew-normal distributions. *Scandinavian Journal of Statistics*, 33(3):561–574, 2006.
- [43] Hua-Hua Chang and William Stout. The asymptotic posterior normality of the latent trait in an IRT model. *Psychometrika*, 58(1):37–52, 1993.
- [44] Chun Wang. On latent trait estimation in multidimensional compensatory item response models. *Psychometrika*, 80(2):428–449, 2015.
- [45] Z. Jiang and J. Templin. Gibbs samplers for logistic item response models via the pólya–gamma distribution: A computationally efficient data-augmentation strategy. *Psychometrika*, 84(2):358–374, 2019.
- [46] Neil J. Gordon, David Salmond, and Adrian F. M. Smith. Novel approach to nonlinear/non-gaussian bayesian state estimation. 1993.
- [47] A. F. M. Smith and A. E. Gelfand. Bayesian statistics without tears: A sampling-resampling perspective. *The American Statistician*, 46(2):84–88, 1992.
- [48] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [49] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabás Póczos, Ruslan Salakhutdinov, and Alexander J Smola. Deep sets. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 3394–3404, Red Hook, NY, USA, 2017. Curran Associates Inc.
- [50] Tom Schaul, Daniel Horgan, Karol Gregor, and David Silver. Universal value function approximators. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 1312–1320, Lille, France, 07–09 Jul 2015. PMLR.
- [51] Chelsea Krantsevich, P. Richard Hahn, Yi Zheng, and Charles Katz. Bayesian decision theory for tree-based adaptive screening tests with an application to youth delinquency. *Annals of Applied Statistics*, 17(2):1038–1063, June 2023.
- [52] Wim J. van der Linden and Hao Ren. A fast and simple algorithm for bayesian adaptive testing. *Journal of Educational and Behavioral Statistics*, 45(1):58–85, 2020.
- [53] Carri W. Chan and Vivek F. Farias. Stochastic depletion problems: Effective myopic policies for a class of dynamic optimization problems. *Mathematics of Operations Research*, 34(2):333–350, 2009.

Supplementary Material: Deep Computerized Adaptive Testing

A Mutual Information as Prediction Uncertainties

We observe that all the heuristic Bayesian item selection methods discussed in this manuscript account for prediction uncertainty. In particular, the mutual information item selection rule can be rewritten as follows:

$$\begin{aligned}
 \arg \max_{j_t \in R_t} I_M(\boldsymbol{\theta}, y_{j_t}) &= \arg \max_{j_t \in R_t} \sum_{y_{j_t}=0}^1 \int_{\boldsymbol{\theta}} f(\boldsymbol{\theta}, y_{j_t} | \mathbf{Y}_{1:(t-1)}) \log \frac{f(\boldsymbol{\theta}, y_{j_t} | \mathbf{Y}_{1:(t-1)})}{f(\boldsymbol{\theta} | \mathbf{Y}_{1:(t-1)}) f(y_{j_t} | \mathbf{Y}_{1:(t-1)})} d\boldsymbol{\theta} \\
 &= \arg \max_{j_t \in R_t} \sum_{y_{j_t}=0}^1 \int_{\boldsymbol{\theta}} f(y_{j_t} | \boldsymbol{\theta}) f(\boldsymbol{\theta} | \mathbf{Y}_{1:(t-1)}) \log \frac{f(\boldsymbol{\theta}, y_{j_t} | \mathbf{Y}_{1:(t-1)})}{f(\boldsymbol{\theta} | \mathbf{Y}_{1:(t-1)}) f(y_{j_t} | \mathbf{Y}_{1:(t-1)})} d\boldsymbol{\theta} \\
 &= \arg \max_{j_t \in R_t} \sum_{y_{j_t}=0}^1 \int_{\boldsymbol{\theta}} f(y_{j_t} | \boldsymbol{\theta}) f(\boldsymbol{\theta} | \mathbf{Y}_{1:(t-1)}) \log \frac{f(y_{j_t} | \boldsymbol{\theta})}{f(y_{j_t} | \mathbf{Y}_{1:(t-1)})} d\boldsymbol{\theta} \\
 &= \arg \max_{j_t \in R_t} \int_{\boldsymbol{\theta}} [\Phi(\mathbf{B}'_{j_t} \boldsymbol{\theta} + D_{j_t}) \log \frac{\Phi(\mathbf{B}'_{j_t} \boldsymbol{\theta} + D_{j_t})}{c_{j_t}} + \\
 &\quad (1 - \Phi(\mathbf{B}'_{j_t} \boldsymbol{\theta} + D_{j_t})) \log \frac{(1 - \Phi(\mathbf{B}'_{j_t} \boldsymbol{\theta} + D_{j_t}))}{(1 - c_{j_t})}] f(\boldsymbol{\theta} | \mathbf{Y}_{1:(t-1)}) d\boldsymbol{\theta}. \tag{18}
 \end{aligned}$$

Equation (18) suggests the mutual information criterion selects the item j_t that has the largest expected KL divergence between the Bernoulli distributions parametrized by $\Phi(\mathbf{B}'_{j_t} \boldsymbol{\theta} + D_{j_t})$ and c_{j_t} , weighted by the current posterior $f(\boldsymbol{\theta} | \mathbf{Y}_{1:(t-1)})$.

Observe that mutual information essentially favors the item j_t that has the largest prediction uncertainties around its prediction mean c_{j_t} , quantified by the KL divergence. To compare equation (18) with equation (6), note that we have replace the integral of KL divergence from c_{j_t} with variances $(\Phi(\mathbf{B}'_{j_t} \boldsymbol{\theta} + D_{j_t}) - c_{j_t})^2$ in (6). This suggests that our proposed approach shares similar theoretical properties with the mutual information criterion, but offers a much simpler formula for online computation, avoiding posterior reweighting and log transformations. Empirically, our proposed item selection rule achieves comparable performance to mutual information but with significantly reduced computational time.

B Sequential Sampling in CAT

B.1 Proof of Theorem 4.2

For any $T \geq 1$, we may write the likelihood of item response after administering T items under the two-parameter probit MIRT model as follows:

$$\begin{aligned} \prod_{t=1}^T \Phi(\mathbf{B}'_{jt} \boldsymbol{\theta} + D_{jt})^{y_{jt}} (1 - \Phi(\mathbf{B}'_{jt} \boldsymbol{\theta} + D_{jt}))^{1-y_{jt}} &= \prod_{t=1}^T \Phi\{(2y_{jt} - 1)(\mathbf{B}'_{jt} \boldsymbol{\theta} + D_{jt})\} \\ &= \Phi_T(\mathbf{C}_1 \boldsymbol{\theta} + \mathbf{C}_2; \mathbb{I}_T) \end{aligned}$$

Given the standard multivariate normal prior $\boldsymbol{\theta} \sim N(0, \mathbb{I}_K)$, the posterior can be expressed as follows:

$$\begin{aligned} \pi(\boldsymbol{\theta} | \mathbf{Y}_{1:T}, \mathbf{B}_{1:T}, \mathbf{D}_{1:T}) &\propto \phi_K(\boldsymbol{\theta}; \mathbf{0}, \mathbb{I}_K) \Phi_T\{(\mathbf{C}_1 \boldsymbol{\theta} + \mathbf{C}_2); \mathbb{I}_T\} \\ &= \phi_K(\boldsymbol{\theta}; \mathbf{0}, \mathbb{I}_K) \Phi_T\{\mathbf{C}_3^{-1}(\mathbf{C}_1 \boldsymbol{\theta} + \mathbf{C}_2); \mathbf{C}_3^{-1} \mathbf{C}_3^{-1}\} \\ &= \phi_K(\boldsymbol{\theta}; \mathbf{0}, \mathbb{I}_K) \Phi_T\{\mathbf{C}_3^{-1} \mathbf{C}_2 + \mathbf{C}_3^{-1} \mathbf{C}_1 \boldsymbol{\theta}; \mathbf{C}_3^{-1} \mathbf{C}_3^{-1}\}. \end{aligned}$$

To identify the exact posterior parameters, we draw back to our attention to the probability kernel of the unified skew-normal distribution as defined in Definition 4.1. We observe that the posterior density above aligns with the SUN density. Specifically, if $\boldsymbol{\theta} \sim \text{SUN}_{K,T}(\boldsymbol{\mu}, \boldsymbol{\Omega}, \boldsymbol{\Delta}, \boldsymbol{\gamma}, \boldsymbol{\Gamma})$, its density is

$$\phi_K(\boldsymbol{\theta}; \boldsymbol{\mu}, \boldsymbol{\Omega}) \frac{\Phi_T(\boldsymbol{\gamma} + \boldsymbol{\Delta}^\top \bar{\boldsymbol{\Omega}}^{-1} \boldsymbol{\Omega}^{-1}(\boldsymbol{\theta} - \boldsymbol{\mu}); \boldsymbol{\Gamma} - \boldsymbol{\Delta}^\top \bar{\boldsymbol{\Omega}}^{-1} \boldsymbol{\Delta})}{\Phi_T(\boldsymbol{\gamma}; \boldsymbol{\Gamma})}.$$

We can interpret this as a product of a K -variate Gaussian pdf in $\boldsymbol{\theta}$ and a T -variate Gaussian cdf in $\boldsymbol{\theta}$, with normalizing constant $\Phi_T(\boldsymbol{\gamma}; \boldsymbol{\Gamma})$, since $\Phi_T(\boldsymbol{\gamma}; \boldsymbol{\Gamma})$ is not a function of $\boldsymbol{\theta}$. But this exactly matches our derived latent-factor posterior $\pi(\boldsymbol{\theta} | \mathbf{Y}_{1:T}, \mathbf{B}_{1:T}, \mathbf{D}_{1:T})$ under the MIRT model. It remains to identify $(\boldsymbol{\mu}, \boldsymbol{\Omega}, \boldsymbol{\Delta}, \boldsymbol{\gamma}, \boldsymbol{\Gamma})$ by matching $\phi_K(\boldsymbol{\theta}; \boldsymbol{\mu}, \boldsymbol{\Omega}) \Phi_T(\boldsymbol{\gamma} + \boldsymbol{\Delta}^\top \bar{\boldsymbol{\Omega}}^{-1} \boldsymbol{\Omega}^{-1}(\boldsymbol{\theta} - \boldsymbol{\mu}); \boldsymbol{\Gamma} - \boldsymbol{\Delta}^\top \bar{\boldsymbol{\Omega}}^{-1} \boldsymbol{\Delta})$ to

$$\phi_K(\boldsymbol{\theta}; \mathbf{0}, \mathbf{I}_K) \Phi_T(\mathbf{C}_3^{-1} \mathbf{C}_2 + \mathbf{C}_3^{-1} \mathbf{C}_1 \boldsymbol{\theta}; \mathbf{C}_3^{-1} (\mathbf{C}_3^{-1})^\top).$$

Matching the pdf part immediately yields $\boldsymbol{\mu}_{\text{post}} = \mathbf{0}_K$ and $\boldsymbol{\Omega}_{\text{post}} = \mathbb{I}_K$. To match the cdf part, we need to match $\boldsymbol{\gamma} + \boldsymbol{\Delta}^\top \bar{\boldsymbol{\Omega}}^{-1} \boldsymbol{\Omega}^{-1}(\boldsymbol{\theta} - \boldsymbol{\mu})$ to $\mathbf{C}_3^{-1} \mathbf{C}_2 + \mathbf{C}_3^{-1} \mathbf{C}_1 \boldsymbol{\theta}$, which implies $\boldsymbol{\gamma}_{\text{post}} = \mathbf{C}_3^{-1} \mathbf{C}_2$.

It remains to solve the following two linear equations:

$$\begin{aligned}\Delta' \bar{\Omega}^{-1} \Omega^{-1} &= \mathbf{C}_3^{-1} \mathbf{C}_1 \\ \Gamma - \Delta' \bar{\Omega}^{-1} \Delta &= \mathbf{C}_3^{-1} \mathbf{C}_3^{-1}.\end{aligned}$$

Recall that in definition 4.1, we let $\Omega = \omega \bar{\Omega} \omega$ be the decomposition of covariance matrix into correlation matrix. Since $\Omega_{\text{post}} = \mathbb{I}_K$, we have $\bar{\Omega} = \Omega = \mathbb{I}_K$. It follows that solving the first equation yields $\Delta_{\text{post}} = \mathbf{C}'_1 \mathbf{C}_3^{-1}$. Plugging our solution of Δ_{post} in the second solution yields

$$\Gamma_{\text{post}} = \mathbf{C}_3^{-1} \mathbf{C}_3^{-1} + \mathbf{C}_3^{-1} \mathbf{C}_1 \mathbf{C}'_1 \mathbf{C}_3^{-1} = \mathbf{C}_3^{-1} (\mathbf{C}_1 \mathbf{C}'_1 + \mathbb{I}_T) \mathbf{C}_3^{-1}.$$

Finally, to show this is indeed a unified skew-normal distribution, we need to show the matrix Ω^* as defined in Definition 4.1 is indeed a full-rank correlation matrix. To see this, we may decompose Ω^* as follows:

$$\begin{bmatrix} \mathbf{C}_3^{-1} (\mathbf{C}_1 \mathbf{C}'_1 + \mathbb{I}_T) \mathbf{C}_3^{-1} & \mathbf{C}_3^{-1} \mathbf{C}_1 \\ \mathbf{C}'_1 \mathbf{C}_3^{-1} & \mathbb{I}_K \end{bmatrix} = \begin{bmatrix} \mathbf{C}_3^{-1} & 0 \\ 0 & \mathbb{I}_K \end{bmatrix} \times \begin{bmatrix} \mathbf{C}_1 \mathbf{C}'_1 + \mathbb{I}_T & \mathbf{C}_1 \\ \mathbf{C}'_1 & \mathbb{I}_K \end{bmatrix} \times \begin{bmatrix} \mathbf{C}_3^{-1} & 0 \\ 0 & \mathbb{I}_K \end{bmatrix}.$$

Observe that this is a decomposition of a correlation matrix, where the middle matrix above is the covariance matrix of a $(T + K)$ dimensional random vector $[\mathbf{z}'_1, \mathbf{z}'_2]'$, where $\mathbf{z}_1 = \mathbf{C}_1 \mathbf{z}_2 + \boldsymbol{\epsilon}$, \mathbf{z}_2 is a K -dimensional standard multivariate Gaussian vector with identity covariance matrix, and $\boldsymbol{\epsilon}$ is a T -dimensional standard multivariate Gaussian vector, independent of \mathbf{z}_2 .

B.2 Sequential Sampling Algorithms

For any Bayesian item selection algorithms, including our proposed deep CAT framework, it is essential to sample from the latent factor posterior distributions $f(\boldsymbol{\theta} | \mathbf{Y}_{1:T})$ for any $T \geq 1$. Theorem 4.2 provides a direct sampling approach to perform such sequential sampling, hence circumventing the need for MCMC algorithms, which cannot be parallelized and require additional tuning and mixing time.

The key observation is that any unified skew-normal distribution $\boldsymbol{\theta} \sim \text{SUN}_{K,T}(\boldsymbol{\mu}, \Omega, \Delta, \gamma, \Gamma)$ has the stochastic representation as $\boldsymbol{\theta} \stackrel{d}{=} \boldsymbol{\mu} + \omega(\mathbf{V}_0 + \Delta \Gamma^{-1} \mathbf{V}_{1,-\gamma})$, where $\mathbf{V}_0 \sim N(0, \bar{\Omega} - \Delta \Gamma^{-1} \Delta')$ $\in \mathbb{R}^K$ and $\mathbf{V}_{1,-\gamma}$ is obtained by component-wise truncation below $-\gamma$ of a variate

$N(0, \mathbf{\Gamma}) \in \mathbb{R}^T$. By plugging in the posterior parameters derived in Theorem 4.2, we have the representation

$$\boldsymbol{\theta} \sim \mathbf{V}_0 + \mathbf{C}'_1(\mathbf{C}_1\mathbf{C}'_1 + \mathbb{I}_T)^{-1}\mathbf{C}_3\mathbf{V}_{1,-\gamma}.$$

Hence we can conduct efficient sequential posterior sampling using the following three steps when $t > 1$:

- **Step One:** Sample from multivariate normal distribution

$$\mathbf{V}_0 \sim N(0, \mathbb{I}_K - \mathbf{C}'_1(\mathbf{C}_1\mathbf{C}'_1 + \mathbb{I}_T)^{-1}\mathbf{C}_1).$$

- **Step Two:** Leveraging the minimax tilting method [24], sample from the zero-mean T -variate truncated multivariate normal distribution $\mathbf{V}_{1,-\gamma}$ with covariance matrix $\mathbf{C}_3^{-1}(\mathbf{C}_1\mathbf{C}'_1 + \mathbb{I}_T)\mathbf{C}_3^{-1}$ and truncation below $-\mathbf{C}_3^{-1}\mathbf{C}_2$.
- **Step Three:** perform linear computation $\mathbf{V}_0 + \mathbf{C}'_1(\mathbf{C}_1\mathbf{C}'_1 + \mathbb{I}_T)^{-1}\mathbf{C}_3\mathbf{V}_{1,-\gamma}$.

Note for $t = 1$, the sampling steps above can be modified into their univariate analogue. Since steps one and step two are independent, one can simply perform i.i.d sampling from both \mathbf{V}_0 and $\mathbf{V}_{1,-\gamma}$ in parallel to draw large samples of $\boldsymbol{\theta} \sim f(\boldsymbol{\theta}|\mathbf{Y}_{1:T})$ in the probit MIRT model. The $t \times t$ matrix inverse $(\mathbf{C}_1\mathbf{C}'_1 + \mathbb{I}_T)$ is easy to compute and only needs to be computed once.

C Fully Bayesian Item Selection

C.1 Generalizing to Fully Bayesian Item Selection

Another key advantage of our direct sampling approach is its ability to enable rapid fully Bayesian item selection. Since Section 3, we have assumed that the item bank is well-calibrated, allowing the factor loading matrix and the intercept to be treated as fixed. According to [52], this assumption can lead to overly optimistic estimation of the latent traits in practice, and is only reasonable when large calibration datasets are available. A fully Bayesian framework offers a principled way to incorporate this uncertainty by integrating over the joint posterior of item parameters and latent traits. However, performing this integration in online item selection has long been considered intractable, limiting its practical adoption. Notably, [52] introduced an efficient MCMC method, but only for unidimensional item response theory model.

Our direct sampling approach extends fully Bayesian item selection to multidimensional IRT while avoiding the computational bottlenecks of MCMC. To illustrate, we consider the mutual information criterion in (15). Fully Bayesian inference involves obtaining posterior samples of item parameters $\{\boldsymbol{\xi}^{(m)}\}_{m=1}^M$, where $\boldsymbol{\xi}^{(m)} := (\mathbf{B}^{(m)}, \mathbf{D}^{(m)})$. The mutual information criterion in Equation (15) requires computing the posterior predictive probability:

$$f(y_{jt} | \mathbf{Y}_{1:(t-1)}) = \int_{\boldsymbol{\xi}} \int_{\boldsymbol{\theta}} f(y_{jt} | \boldsymbol{\theta}) f(\boldsymbol{\theta} | \boldsymbol{\xi}, \mathbf{Y}_{1:(t-1)}) d\boldsymbol{\theta} d\boldsymbol{\xi}. \quad (19)$$

Similarly, the KL divergence term must integrate out the nuisance item parameters $\boldsymbol{\xi}$:

$$\int_{\boldsymbol{\xi}} \int_{\boldsymbol{\theta}} f(\boldsymbol{\theta} | \mathbf{Y}_{1:(t-1)}, y_{jt}, \boldsymbol{\xi}) \log \frac{f(y_{jt} | \boldsymbol{\theta})}{f(y_{jt} | \mathbf{Y}_{1:(t-1)})} d\boldsymbol{\theta} d\boldsymbol{\xi}. \quad (20)$$

By leveraging Theorem 4.2, we can exactly identify the posterior distributions for any arbitrary configuration of $f(\boldsymbol{\theta} | \mathbf{Y}_{1:(t-1)}, y_{jt}, \boldsymbol{\xi}^{(m)})$, hence enabling direct sampling in parallel. It follows Equations (19) and (20) are computationally feasible via Monte Carlo integration.

In contrast, a standard MCMC approach requires constructing M independent Markov chains with additional data augmentation techniques [21, 22], each targeting $f(\boldsymbol{\theta} | \mathbf{Y}_{1:(t-1)}, y_{jt}, \boldsymbol{\xi}^{(m)})$, with sampling only possible after all chains have reached convergence. This introduces a significant computational burden, making MCMC impractical for real-time item selection.

We evaluate the effectiveness of our direct sampling approach in enabling fully Bayesian online item selection. Unlike traditional methods that treat item parameters as fixed, the fully Bayesian approach explicitly accounts for their uncertainties, which can be substantial when the item response dataset is small or poorly calibrated. To illustrate, we generate a binary item response dataset \mathbf{Y} with $N = 500$ examinees and $J = 150$ items under a 3-factor probit MIRT model. The true factor loading matrix \mathbf{B} has dimensions $J \times 3$, and the intercept vector \mathbf{D} has 150 elements.

To integrate item parameter uncertainty into online item selection, we fit a Gibbs sampler to the $N \times J$ item response data. We generated 5,000 MCMC draws, retaining the last 500 posterior samples of the item parameters $\boldsymbol{\Xi} = \{\boldsymbol{\xi}^{(m)}\}_{m=1}^{500}$ after burn-in. As described in Section C.1, the fully Bayesian approach marginalizes over the joint distributions of the nuisance item parameters $\boldsymbol{\Xi}$ and latent traits. We then implemented all item selection rules depicted in Sections 3.2 in a fully Bayesian manner.

Figure 6 illustrates the cumulative percentage of completed test sessions across all 500 simulated examinees as more items are administered. The test was dynamically terminated when the posterior standard deviations across all three factors fell below 0.4. Con-

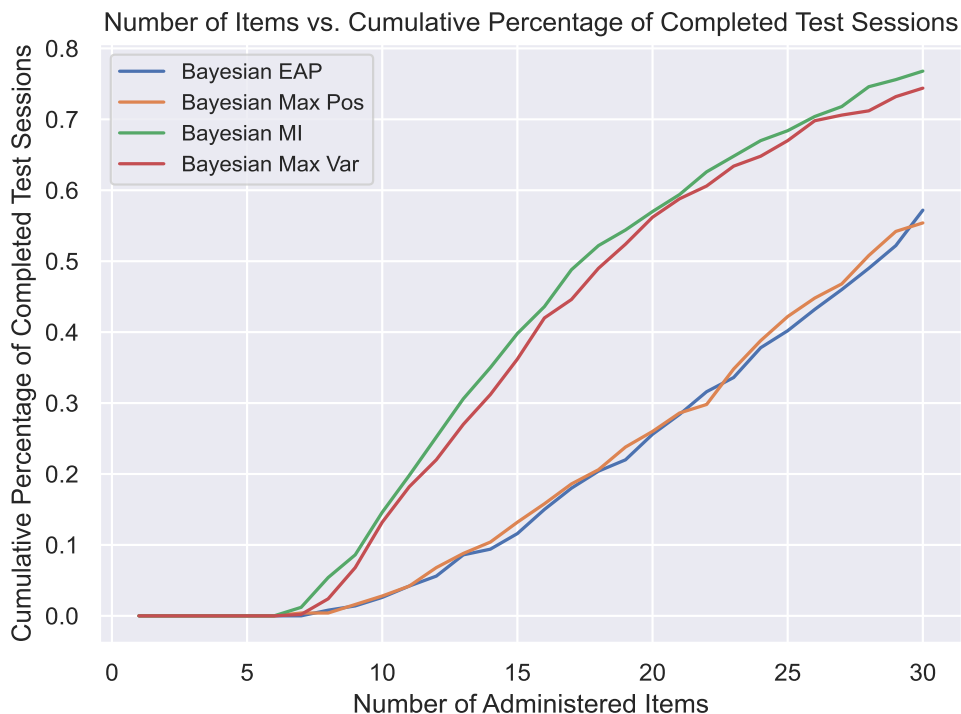


Figure 6: Fully Bayesian: Number of Items Versus Posterior Variance Reduction

Table 5: **Comparison of Winshares (W.S), Termination, and Computation**

Algorithm	Avg Termination (items)	W.S dim0	W.S dim1	W.S dim2	Avg Time (s/item)
EAP	24.4	22.2%	19.2%	21.8%	4.10
Max Pos	24.2	17.4%	18.0%	21.0%	4.10
MI	18.7	32.8%	31.0%	27.8%	4.11
Max Var	19.2	27.6%	31.8%	29.4%	4.10

sistent with the experimental results in Section 6 of the manuscript, the mutual information method and our proposed Max Var method demonstrated superior posterior variance reduction. The second column of Table 5 shows that the mutual information method terminated after an average of 18.7 items, whereas the EAP approach required 24.4 items.

Beyond variance reduction, the mutual information and maximizing prediction variance methods also exhibited superior estimation accuracy, as shown in Figure 7, which tracks the decline in mean squared error (MSE) between the posterior mean and the oracle posterior mean. Since estimation errors stabilized around $H = 30$ items, we compared MSEs between the posterior mean at $H = 20$ and the oracle posterior mean, and then summarized the win shares for each item selection rule across all 500 examinees and all three dimensions in Table 5.

Finally, we emphasize that even in a fully Bayesian framework, where nuisance pa-

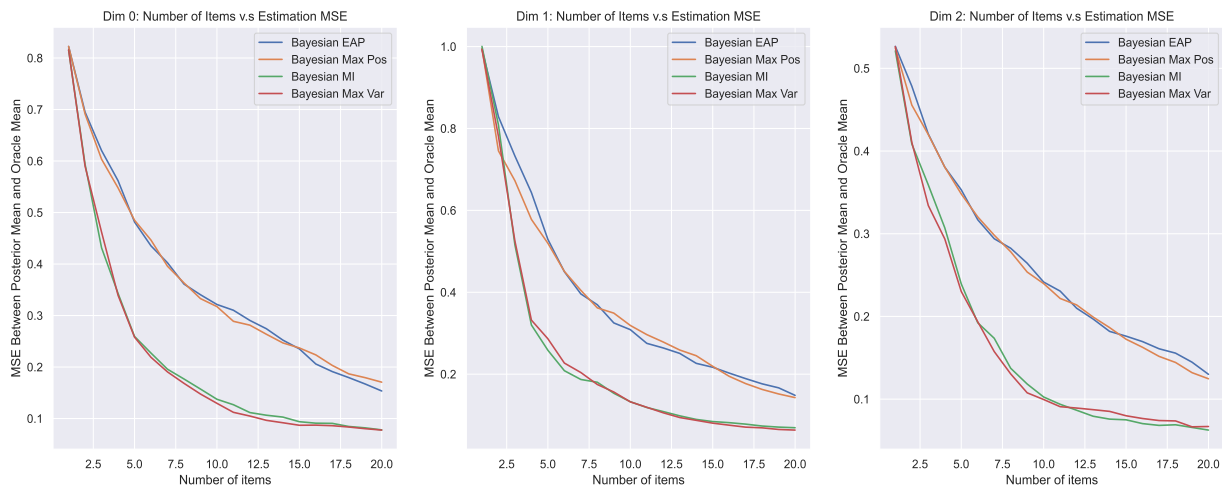


Figure 7: Fully Bayesian: MSE Between Posterior Mean and Oracle Posterior Mean

rameters are integrated rather than fixed, our approach remains computationally efficient. Across all item selection rules, online selection required an average of just 4.1 seconds per test session on a personal laptop, highlighting the scalability of our direct sampling method for Bayesian MCAT applications. Readers may notice that 4.1 seconds is significantly longer than the item selection time reported in the experimentation section, where item parameters are treated as fixed. This discrepancy arises because fully Bayesian MCAT requires sampling from $M = 500$ distinct unified-skew normal distributions for each item selection, whereas the standard case involves sampling from a single fixed distribution. This also explains the minimal differences in item selection time across algorithms, as the primary computational bottleneck lies in sampling from these distributions. When more CPUs are available, we can certainly sample from all M distributions in parallel, and hence further improving the 4.1 seconds benchmark.

D Neural Network Design Details

Both the primary and target Q-networks share the same modular feed-forward design. At each time t , the posterior parameters $\tilde{\xi}_t$ are split into two streams: a permutation-invariant path ϕ_1 (three 256-unit ReLU layers) that processes the main state input. In parallel, a RowwiseNetwork $\phi_2(\cdot)$ network applies independently to each of the J prediction-quartile feature vectors yields item-wise predictive values, which are then mapped to 256 dimensional features. The two 256-dimensional vectors are then merged by a small combiner MLP, which is then mapped to the final layers with J nodes through the classification

network ρ . All linear layers use Xavier initialization; activations are ReLU; the final output is masked to $-\infty$ for unavailable items.

Table 6: Layer-by-layer specification of `OnlineQNetworkV` (and its submodules).

Module	Input Dim	Output Dim	Activation
ϕ_1 (Phi1Network)	S	256	ReLU
	256	256	ReLU
	256	256	ReLU
ϕ_2 (RowwiseNetwork) (for Ψ_t)	11	256	ReLU
	256	256	ReLU
	256	1	Linear $\times J$ rows
ϕ_2 (SmallPhi2Network)	J	256	ReLU
	256	256	ReLU
	256	256	Linear
ρ (Classification)	256	J	ReLU
	J	J	Linear

Notes:

- S = total scalar-state dimension; Each new item selection introduced $(K + 2)$ dimensional tuple
- The RowwiseNetwork applies its three layers *independently* to each of the J rows of the $(J \times 11)$ Ψ_t matrix. There are 11 columns, because we include mean, variance, and the 10% quantile increments from 10% to 90%. Because we process each row independently, the output for the RowwiseNetwork is J dimensional.
- All weight matrices are initialized with Xavier uniform; no dropout is used.
- Optimizer: Adam with learning rate 10^{-4} , $\beta_1 = 0.9$, $\beta_2 = 0.999$.

E Discussion on the 0-1 Reward

Recall that our choice of reward function is defined in Equation (7). The simple 0-1 reward structure offers several advantages. It is interpretable as it directly minimizes the number of items required to terminate assessment tasks and prioritizes the main factors of interest. Additionally, it has demonstrated success in related applications, such as in designing deep adaptive learning system [28]. Since the cumulative reward is constrained to integer values between $-H$ and -1 , the action-value function can be more reliably approximated by the

neural network.

Our double Q-learning algorithm enjoys theoretical convergence guarantees to the optimal policy [19] and can accommodate arbitrary reward specifications. It is known that myopic policy cannot outperform the principled optimal policy obtained through RL. However, under certain reward structures, a greedy one-step lookahead (myopic) policy can perform surprisingly well.

In [53], it is shown that for stochastic depletion problems, a myopic policy achieves at least 50% of the expected reward of the optimal policy. A similar result can be proven for CAT under specific reward structures. Although this can be interpreted as a somewhat negative result for the reinforcement learning approach, we emphasize our choice of the 0-1 reward does not conform to the stochastic depletion problem, and hence is not applicable to the theorem below.

To illustrate this, we extend the state space \mathcal{S} with a time variable, defining $\bar{\mathcal{S}} = \{(f(\boldsymbol{\theta} | \mathbf{Y}_{1:T}), t) : t \in [H]\}$. Let $\pi^*(\bar{s})$ represent the optimal policy, and $\pi^m(\bar{s})$ the myopic policy that maximizes the next-step expected reward. Define the mapping $\hat{G}_{y_a} : \bar{s} \rightarrow \bar{s}'$ as the posterior update after observing the binary response random variable y_a , with an incremented time step, where $\bar{s} = (f(\boldsymbol{\theta} | \mathbf{Y}_{1:T}), t)$ and $\bar{s}' = (f(\boldsymbol{\theta} | \mathbf{Y}_{1:T}, y_a), t + 1)$. Similarly, define $\tilde{G}_{y_a} : \bar{s} \rightarrow \bar{s}'$ as the posterior update after observing y_a while keeping the time step fixed, where $\bar{s} = (f(\boldsymbol{\theta} | \mathbf{Y}_{1:T}), t)$ and $\bar{s}' = (f(\boldsymbol{\theta} | \mathbf{Y}_{1:T}, y_a), t)$.

To demonstrate the near-optimality of the myopic policy, [53] introduces two key assumptions: “value function monotonicity” and “immediate rewards”. These assumptions can be adapted to the CAT setting as follows:

- **Immediate Rewards:** for any time stamp $t < H$, suppose the myopic policy π^m chooses item a , then we must have

$$V_{\pi^*}(\bar{s}) \leq E[R(\bar{s}, a, \bar{s}')] + V_{\pi^*}(\tilde{G}_{y_a}(\bar{s})). \quad (21)$$

Essentially, the property says given a free item (without increasing the time step) chosen by the myopic policy would not deteriorate the optimal rewards.

- **Value Function Monotonicity:** for any time stamp $t < H$, suppose the optimal policy would choose item a_1 yet the myopic policy would chooses item a_2 , then

$$V_{\pi^*}(\hat{G}_{y_{a_2}}(\tilde{G}_{y_{a_1}}(\bar{s}))) \leq V_{\pi^*}(\hat{G}_{y_{a_2}}(\bar{s})). \quad (22)$$

Note both sides of inequality refer to the optimal value function evaluated at the posterior distribution at time step $t + 1$.

In Theorem E.1, we establish that a reinforcement learning approach to CAT cannot significantly outperform the one-step-lookahead greedy policy when the reward function satisfies the two assumptions outlined above.

Theorem E.1. *Consider the finite horizon CAT problem with H steps. Suppose the assumptions of immediate rewards in (21) and value function monotonicity in (22) hold, we have $V_{\pi^*}(\bar{s}) \leq 2V_{\pi^m}(\bar{s})$ for all $\bar{s} \in \bar{\mathcal{S}}$.*

Proof. We proceed by induction. Since the myopic and the optimal policy would agree on time $t = H - 1$, the inequality holds trivially for the base case. For the induction step, we assume the claimed inequality holds for all t' such that $H > t' > t$. Consider the time horizon $t < t'$ and let a_1 represent the item chosen by the optimal policy and a_2 the item chosen by the myopic policy. The inequality holds trivially for $a_1 = a_2$. Hence, we consider the case when $a_1 \neq a_2$ and let y_{a_1} and y_{a_2} represent the binary random variables for the item response. For given extended state \bar{s} , we use \bar{s}' to denote the next step after selecting an item. Then we have

$$\begin{aligned} V_{\pi^*}(\bar{s}|y_{a_1}, y_{a_2}) &= E[R(\bar{s}, a_1, \bar{s}')|y_{a_1}] + V_{\pi^*}(\hat{G}_{y_{a_1}}(\bar{s})) \\ &\leq E[R(\bar{s}, a_1, \bar{s}')|y_{a_1}] + E[R(\bar{s}, a_2, \bar{s}')|y_{a_2}] + V_{\pi^*}(\tilde{G}_{y_{a_2}}(\hat{G}_{y_{a_1}}(\bar{s}))) \end{aligned} \quad (23)$$

$$\begin{aligned} &= E[R(\bar{s}, a_1, \bar{s}')|y_{a_1}] + E[R(\bar{s}, a_2, \bar{s}')|y_{a_2}] + V_{\pi^*}(\hat{G}_{y_{a_1}}(\tilde{G}_{y_{a_2}}(\bar{s}))) \\ &\leq E[R(\bar{s}, a_1, \bar{s}')|y_{a_1}] + E[R(\bar{s}, a_2, \bar{s}')|y_{a_2}] + V_{\pi^*}(\hat{G}_{y_{a_2}}(\bar{s}))) \end{aligned} \quad (24)$$

$$\leq E[R(\bar{s}, a_1, \bar{s}')|y_{a_1}] + E[R(\bar{s}, a_2, \bar{s}')|y_{a_2}] + 2V_{\pi^m}(\hat{G}_{y_{a_2}}(\bar{s}))). \quad (25)$$

Note we applied the immediate rewards assumption in equation (23), the value function monotonicity assumption in equation (24), and the induction hypothesis in equation (25). Now, we may take expectation over y_{a_1} and y_{a_2} on both side of inequality and obtain:

$$\begin{aligned} V_{\pi^*}(\bar{s}) &\leq E[R(\bar{s}, a_1, \bar{s}')] + E[R(\bar{s}, a_2, \bar{s}')] + 2V_{\pi^m}(\hat{G}_{y_{a_2}}(\bar{s}))) \\ &\leq 2(E[R(\bar{s}, a_2, \bar{s}')] + V_{\pi^m}(\hat{G}_{y_{a_2}}(\bar{s}))) = 2V_{\pi^m}(\bar{s}). \end{aligned}$$

The last inequality is due to the fact the π^m is the myopic policy that maximizes the next step reward. \square

Theorem E.1 can be seen as a limitation of the reinforcement learning approach when the reward structure aligns with the stochastic depletion problem, as it implies that a simple myopic policy provides a 2-approximation to the optimal policy. However, it also highlights the advantage of the simple 0-1 reward function as the value function monotonicity is unlikely to hold, rendering the theorem inapplicable. Specifically, the expected posterior variance of $\hat{G}_{y_{a_2}}(\tilde{G}_{y_{a_1}}(\bar{s}))$ is expected to be smaller than that of $\hat{G}_{y_{a_2}}(\bar{s})$. Consequently, fewer items are expected for the posterior $\hat{G}_{y_{a_2}}(\tilde{G}_{y_{a_1}}(\bar{s}))$ to reach the minimum variance threshold, which implies $V_{\pi^*}(\hat{G}_{y_{a_2}}(\tilde{G}_{y_{a_1}}(\bar{s}))) \geq V_{\pi^*}(\hat{G}_{y_{a_2}}(\bar{s}))$.

F More Analysis on Simulation

F.1 Factor Loading Matrix Visualization

We visualize the factor loading matrix in our simulation exercise in Figure 8. Note all the factor loadings are bounded between $(-3, 3)$ and the matrix is relatively sparse: all the items are loaded on the first factor, but are only allowed to load for at most two more factors. Although identification of the factor loading matrix is not a primary concern in CAT, as it is typically treated as fixed, we purposefully imposed a lower-triangular structure to ensure identifiability and create a more realistic simulation setting.

F.2 Double Q-Learning Training Dynamics

In Figure 9, we present a metric illustrating the training dynamics of our Q-learning algorithm. The left subplot displays the average episode rewards, computed over every 500 training episodes, showing that rewards began converging around episode 25,000 at approximately -23 . The reward is bounded within $[-40, -20]$. Given our definition of the $0 - 1$ reward, the implication is that a random policy would take 40 items for the test to converge. After training, it would only take around 20 items.

Since we save the primary neural network every 1000 episodes, we select the checkpoint corresponding to the highest average reward for offline deployment in future item selection tasks. In this case, the network at episode 72,000 is chosen, as it achieves the peak average reward of approximately -21 . This model selection strategy is consistently applied across all experiments in Section 5.




Figure 8: Simulated Factor Loading Matrix

F.3 Matching the Oracle Posterior

Recall the "oracle posterior distribution" of a test taker is defined as the posterior obtained if a given test taker had answered all items in the item bank. To obtain the oracle distributions for each of the 500 test takers in Section 6, we simulate their item response to all 150 items in the test bank, and then calculate the posterior distribution using Theorem 4.2.

Figures 10, 12, and 13 illustrate the decrease of the estimation MSE for the first quartile, median, and the third quartiles compared to those of the oracle distribution, respectively. The quartiles of each posterior distribution are computed via 1,000 independent draws. In all these figures, we observe similar superior performances for the MI, Max Var, and the Q-learning methods over the EAP and the Max Pos approaches in estimating the entire oracle distributions. Consistent with the findings in Section 6, the Q-learning approach demonstrates the fastest error reduction rate, especially in the early stages of the test.

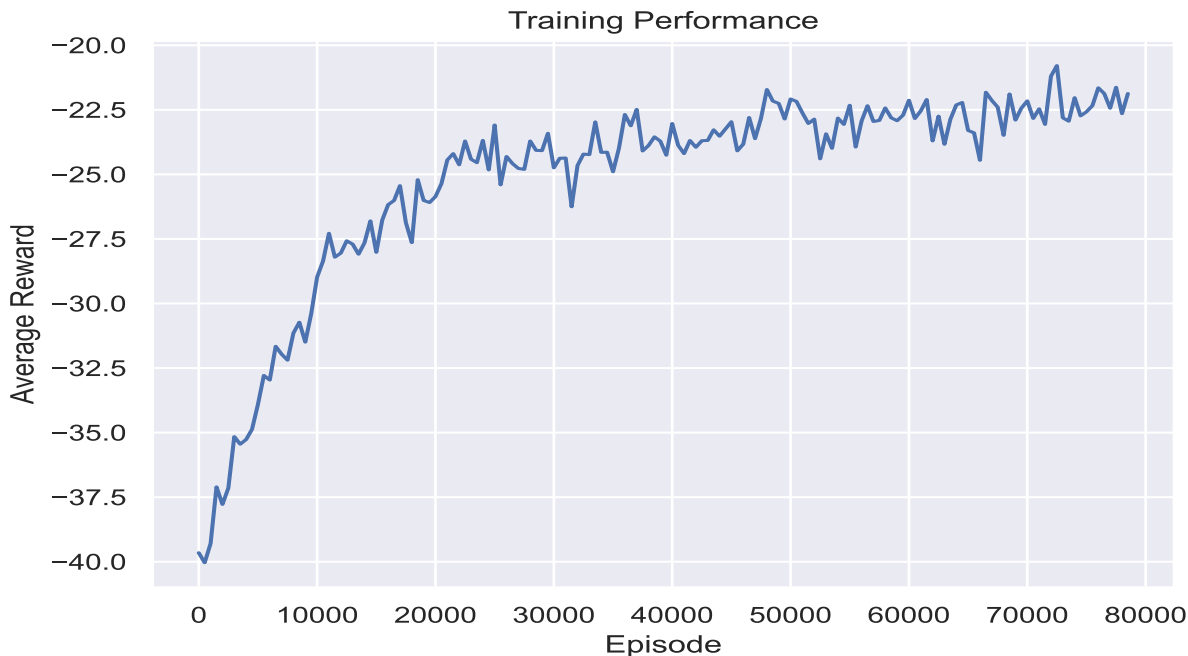


Figure 9: Training Dynamics for the Simulation Experiment in Section 5.1

G pCAT-COG Study Additional Visualization

G.1 Exposure rate

In the real-data experiment, we administer the full 57-item bank in every simulated session to compare the rate of MSE reduction as items accrue. Consequently, item exposure is trivially 100% for all items under all methods, and exposure-rate comparisons are not directly informative in this setting. Therefore, we visualize the distributions of item exposure rate when $T = 30$ in Figure 14. Since the full item bank has 57 items, we should expect the exposure rate to be centered around $30/57 \approx 52.6\%$.

G.2 Q-learning Training Dynamics

H DESE Educational Assessment

We consider an important application in educational assessment, utilizing 2022 Grade 8 student item response data from the Massachusetts Department of Elementary and Secondary Education (DESE). Excluding non-multiple-choice questions, the dataset includes responses to 24 English items and 34 math items. For efficiency, we randomly sampled 1,000 student responses to evaluate each CAT item selection algorithm.




Figure 10: MSE Between Posterior 1st Quantile and Oracle 1st Quantile

Given the unknown latent structure, we estimate the factor loading matrix via the PXL-EM algorithm [1], resulting in a 58 by 4 factor loading matrix, with visualization and the details of model fitting provided in Appendix H. When estimating the item parameters, we adopted the idea of dynamic posterior exploration, letting the regularization λ_0 parameter in the PXL-EM algorithm to gradually increase over the path $\{1, 5, 10, 20, 40, 60, 80, 100\}$. The resulting factor loading matrix \mathbf{B} is visualized in Figure 16, where dark black regions indicate exact zeros. Notably, all English items load exclusively onto the first factor, representing general ability, while math items can also load onto factors 2–4, reflecting their greater structural complexity. For instance, the final factor (“Math 3”) has nonzero loadings only on items 34, 36, and 38, all of which correspond to geometry problems.¹ Interestingly, the authors in [1] obtained very similar factor loading matrix estimation for the Grade 10 data as well.

Since all items are heavily loaded on the first primary factor, representing the global latent ability, we focus on the problem of adaptively measuring the first factor while accounting for the underlying multidimensional structure in dimensions 2-4. The left subplot of Figure 15 illustrates the cumulative percentage of completed test sessions for the 1,000 sampled examinees as more items are administered. As shown in the plot, nearly 70%

¹Exam questions are available at: <https://www.doe.mass.edu/mcas/2022/release/>.

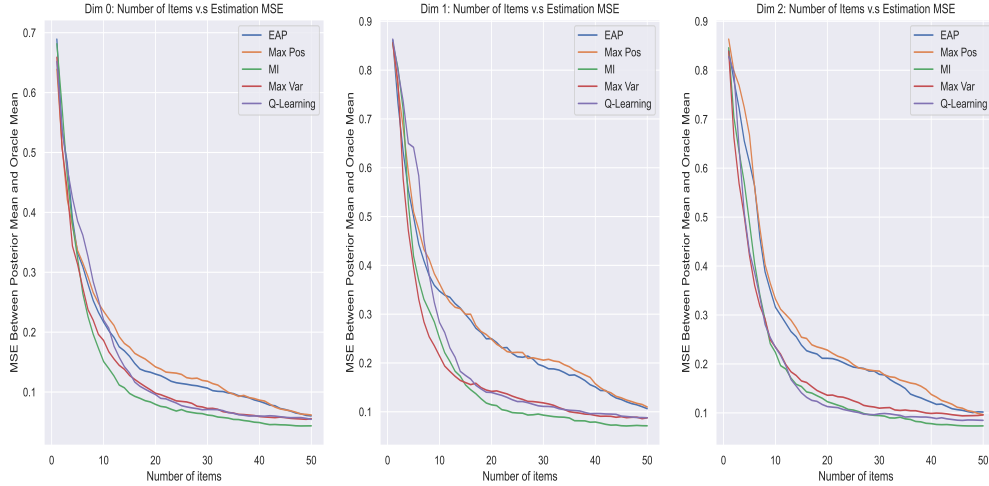


Figure 11: MSE Between Posterior Means and Oracle Posterior Means

of CAT sessions terminate after only 20 items under the Q-learning policy. On average over these 1,000 examinees, Q-learning also required only 20 items to reach the termination criterion, whereas the second-fastest method (MI) required nearly 26 items. Beyond its efficiency in variance reduction, Q-learning also achieves a faster decay in MSE when estimating the first factor, reinforcing findings from earlier experiments.

These results highlight the potential of Q-learning for educational assessments, demonstrating its ability to adapt to various testing environments and to accelerate testing while maintaining accurate estimation of students' multivariate latent traits.




Figure 12: MSE Between Posterior Median and Oracle Median




Figure 13: MSE Between Posterior 3rd Quantile and Oracle 3rd Quantile

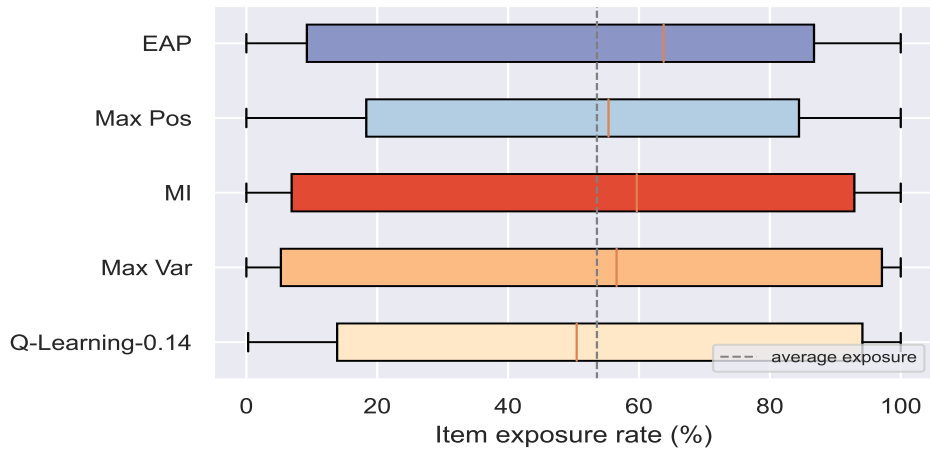
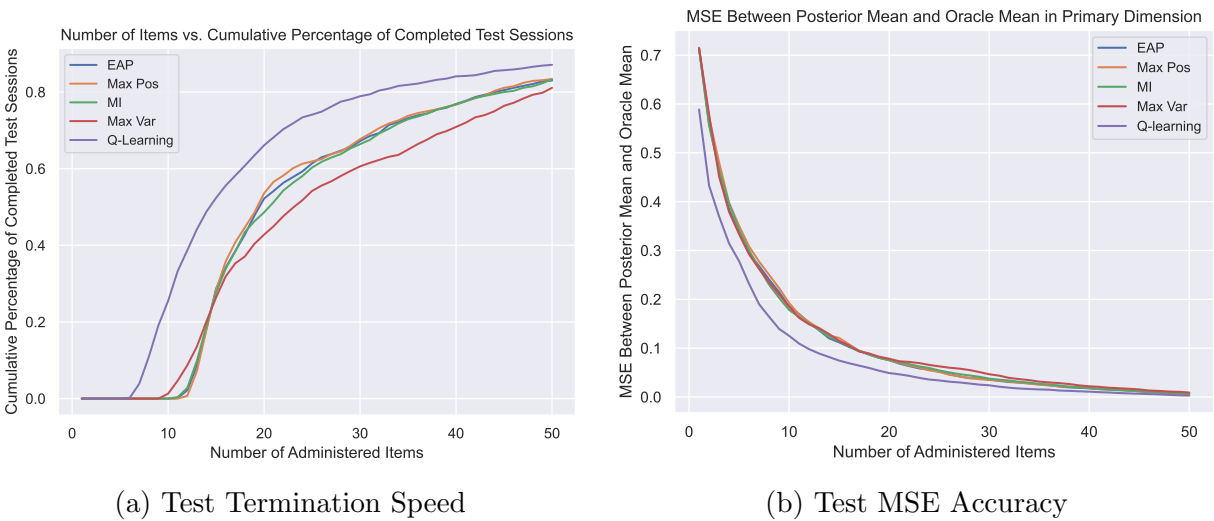


Figure 14: Distributions of Item Exposure Rates at $T = 30$: pCAT-COG



(a) Test Termination Speed

(b) Test MSE Accuracy

Figure 15: Educational Assessment: Main Factor Posterior Variance Reduction (Left) and Estimation Accuracy (Right)




Figure 16: Estimated Factor Loadings for DESE Data